Revision: August 4, 2026

The Bmad Reference Manual

David Sagan

Overview

Bmad (Otherwise known as “Baby MAD" or “Better MAD" or just plain “Be MAD!") is a subroutine library for charged–particle and X-Ray simulations in accelerators and storage rings. *Bmad* has been developed at the Cornell Laboratory for Accelerator-based ScienceS and Education (CLASSE) and has been in use since 1996.

Prior to the development of *Bmad*, simulation programs at Cornell were written almost from scratch to perform calculations that were beyond the capability of existing, generally available software. This practice was inefficient, leading to much duplication of effort. Since the development of simulation programs was time consuming, needed calculations where not being done. As a response, the *Bmad* subroutine library, using an object oriented approach and written in Fortran 2008, were developed. The aim of the *Bmad* project was to:

- Cut down on the time needed to develop programs.
- Cut down on programming errors.
- Provide a simple mechanism for lattice function calculations from within control system programs.
- Provide a flexible and powerful lattice input format.
- Standardize sharing of lattice information between programs.

Bmad can be used to study both single and multi–particle beam dynamics as well as X-rays. Over the years, *Bmad* modules have been developed for simulating a wide variety of phenomena including intra beam scattering (IBS), coherent synchrotron radiation (CSR), Wakefields, Touschek scattering, higher order mode (HOM) resonances, etc., etc. *Bmad* has various tracking algorithms including Runge–Kutta and symplectic (Lie algebraic) integration. Wakefields, and radiation excitation and damping can be simulated. *Bmad* has routines for calculating transfer matrices, emittances, Twiss parameters, dispersion, coupling, etc. The elements that *Bmad* knows about include quadrupoles, RF cavities (both storage ring and LINAC accelerating types), solenoids, dipole bends, Bragg crystals etc. In addition, elements can be defined to control the attributes of other elements. This can be used to simulate the “girder” which physically support components in the accelerator or to easily simulate the action of control room “knobs” that gang together, say, the current going through a set of quadrupoles.

To be able to extend *Bmad* easily, *Bmad* has been developed in a modular, object oriented, fashion to maximize flexibility. As just one example, each individual element can be assigned a particular tracking method in order to maximize speed or accuracy and the tracking methods can be assigned via the lattice file or at run time in a program.

Introduction

As a consequence of *Bmad* being a software library, this manual serves two masters: The programmer who wants to develop applications and needs to know about the inner workings of *Bmad*, and the user who simply needs to know about the *Bmad* standard input format and about the physics behind the various calculations that *Bmad* performs.

To this end, this manual is divided into three parts. The first two parts are for both the user and programmer while the third part is meant just for programmers.

Part I
Part I discusses the *Bmad* lattice input standard. The *Bmad* lattice input standard was developed using the *MAD* [Grote96, Iselin94]. lattice input standard as a starting point but, as *Bmad* evolved, *Bmad*'s syntax has evolved with it.

Part II
part II gives the conventions used by *Bmad*— coordinate systems, magnetic field expansions, etc. — along with some of the physics behind the calculations. By necessity, the physics documentation is brief and the reader is assumed to be familiar with high energy accelerator physics formalism.

Part III
Part III gives the nitty–gritty details of the *Bmad* subroutines and the structures upon which they are based.

More information, including the most up–to–date version of this manual, can be found at the *Bmad* web site[Bmad]. Errors and omissions are a fact of life for any reference work and comments from you, dear reader, are therefore most welcome. Please send any missives (or chocolates, or any other kind of sustenance) to:

```
David Sagan <dcs16@cornell.edu>
```

The *Bmad* manual is organized as reference guide and so does not do a good job of instructing the beginner as to how to use *Bmad*. For that there is an introduction and tutorial on *Bmad* and *Tao* (§1.2) concepts that can be downloaded from the *Bmad* web page. Go to either the *Bmad* or *Tao* manual pages and there will be a link for the tutorial.

It is my pleasure to express appreciation to people who have contributed to this effort, and without whom, *Bmad* would only be a shadow of what it is today: To David Rubin and Georg Hoffstaetter for their support all these years, to Étienne Forest (aka Patrice Nishikawa) for use of his remarkable PTC/FPP library (not to mention his patience in explaining everything to me), to Desmond Barber for very useful discussions on how to simulate spin, to Jonathan Laster, Mark Palmer, Matt Rendina, and Attilio De Falco for all their work maintaining the build system and for porting *Bmad* to different platforms, to Frank Schmidt and CERN for permission to use the *MAD* tracking code. To Hans Grote and CERN for granting permission to adapt figures from the *MAD* manual for use in this one, to Martin Berz for his DA package, and to Dan Abell, Jacob Asimow, Ivan Bazarov, Moritz Beckmann, Scott Berg, Oleksii Beznosov, Kevin Brown, Joel Brock, Sarah Buchan, Avishek Chatterjee, Jing Yee Chee, Christie Chiu, Joseph Choi, Robert Cope, Jim Crittenden, Laurent Deniau, Bhawin Dhital, Gerry Dugan, Michael Ehrlichman, Jim Ellison, Ken Finkelstein, Mike Forster, Thomas Gläßle, Juan Pablo Gonzalez-Aguilera, Sam Grant, Colwyn Gulliford, Eiad Hamwi, Klaus Heinemann, Richard Helms, Lucy Lin, Henry Lovelace III, Chris Mayes, Vasiliy Morozov, Karthik Narayan, Katsunobu Oide, Tia Plautz, Matt Randazzo, Robert Ryne, Michael Saelim, Jim Shanks, Matthew Signorelli, Hugo Slepicka, Jeff Smith, Jonathan Unger, Jeremy Urban, Ningdong Wang, Suntao Wang, Mark Woodley, and Demin Zhou for their help.

Contents

I Language Reference **23**

1 Orientation **25**
1.1 What is Bmad? . . . 25
1.2 Tao and Bmad Distributions . . . 25
1.3 Resources: More Documentation, Obtaining Bmad, etc. . . . 26
1.4 PTC: Polymorphic Tracking Code . . . 27

2 Bmad Concepts and Organization **29**
2.1 Lattice Elements . . . 29
2.2 Lattice Branches . . . 29
2.3 Lattice . . . 30
2.4 Lord and Slave Elements . . . 30

3 Lattice File Statements **35**
3.1 File Example and Syntax . . . 35
3.2 Digested Files . . . 36
3.3 Element Sequence Definition . . . 37
3.4 Lattice Elements . . . 37
3.5 Lattice Element Names . . . 38
3.6 Matching to Lattice Element Names . . . 39
3.7 Lattice Element Parameters . . . 42
3.8 Nonstandard Parameter Syntax . . . 43
3.9 Custom Element Attributes . . . 44
3.10 Parameter Types . . . 46
3.11 Particle Species Names . . . 46
3.12 Units and Constants . . . 48
3.13 Arithmetic Expressions . . . 49
3.14 Intrinsic functions . . . 51
3.15 Statement Order . . . 52
3.16 Setenv Statement . . . 53
3.17 Print Statement . . . 53
3.18 Title Statement . . . 54
3.19 Call Statement . . . 54
3.20 Inline Call . . . 54
3.21 Undefined_vars_evaluate_to_zero Statement . . . 55
3.22 Use_local_lat_file Statement . . . 55
3.23 No_Superimpose Statement . . . 55
3.24 Return and End_File Statements . . . 56
3.25 Expand_Lattice Statement . . . 56

3.26 Lattice Expansion . . . 56
3.27 Calc_Reference_Orbit Statement . . . 57
3.28 Merge_Elements Statement . . . 57
3.29 Combine_Consecutive_Elements Statement . . . 58
3.30 Remove_Elements Statement . . . 58
3.31 Slice_Lattice Statement . . . 58
3.32 Start_Branch_At Statement . . . 59
3.33 Debugging Statements . . . 60

4 Lattice Elements **61**
4.1 AB_Multipole . . . 63
4.2 AC_Kicker . . . 64
4.3 BeamBeam . . . 66
4.4 Beginning_Ele . . . 69
4.5 Bends: Rbend and Sbend . . . 70
4.6 Capillary . . . 76
4.7 Collimators: Ecollimator and Rcollimator . . . 77
4.8 Converter . . . 78
4.9 Crab_Cavity . . . 81
4.10 Crystal . . . 83
4.11 Custom . . . 86
4.12 Detector . . . 87
4.13 Diffraction_Plate . . . 88
4.14 Drift . . . 90
4.15 E_Gun . . . 91
4.16 ELseparator . . . 93
4.17 EM_Field . . . 94
4.18 Feedback . . . 95
4.19 Fiducial . . . 96
4.20 Fixer . . . 98
4.21 Floor_Shift . . . 100
4.22 Foil . . . 102
4.23 Fork and Photon_Fork . . . 105
4.24 Girder . . . 108
4.25 GKicker . . . 111
4.26 Group . . . 112
4.27 Hybrid . . . 115
4.28 Instrument, Monitor, and Pipe . . . 116
4.29 Kickers: Hkicker and Vkicker . . . 117
4.30 Kicker . . . 118
4.31 Lcavity . . . 119
4.32 Marker . . . 122
4.33 Mask . . . 123
4.34 Match . . . 125
4.35 Mirror . . . 129
4.36 Multipole . . . 130
4.37 Multilayer_mirror . . . 131
4.38 Null_Ele . . . 132
4.39 Octupole . . . 133
4.40 Overlay . . . 134
4.41 Patch . . . 136

4.42 Photon_Init . . . 140
4.43 Quadrupole . . . 144
4.44 Ramper . . . 145
4.45 RF_bend . . . 147
4.46 RFcavity . . . 149
4.47 Sad_Mult . . . 151
4.48 Sample . . . 153
4.49 Sextupole . . . 154
4.50 Sol_Quad . . . 155
4.51 Solenoid . . . 156
4.52 Taylor . . . 157
4.53 Thick_Multipole . . . 161
4.54 Wiggler and Undulator . . . 162
4.54.1 Periodic Type Wigglers . . . 163
4.54.2 Map Type Wigglers . . . 164
4.54.3 Old Wiggler Cartesian Map Syntax . . . 165

5 Element Attributes **167**
5.1 Dependent and Independent Attributes . . . 167
5.2 Field_Master and Normalized Vs. Unnormalized Field Strengths . . . 167
5.3 Type, Alias and Descrip Attributes . . . 168
5.4 Group, Overlay, and Ramper Element Syntax . . . 168
5.5 Reference Energy, Time, and Wavelength . . . 172
5.6 Orientation: Offset, Pitch, Tilt, and Roll Attributes . . . 174
5.6.1 Global Random Misalignment of Elements . . . 174
5.6.2 Straight Line Element Orientation . . . 175
5.6.3 Bend Element Orientation . . . 176
5.6.4 Photon Reflecting Element Orientation . . . 177
5.6.5 Reference Orbit Manipulator Element Orientation . . . 177
5.6.6 Fiducial Element Orientation . . . 178
5.6.7 Girder Orientation . . . 178
5.7 Hkick, Vkick, and Kick Attributes . . . 179
5.8 Aperture and Limit Attributes . . . 179
5.8.1 Apertures and Element Offsets . . . 181
5.8.2 Aperture Placement . . . 181
5.8.3 Apertures and X-Ray Generation . . . 182
5.9 X-Rays Crystal & Compound Materials . . . 182
5.10 X-Ray Reflectivity Tables . . . 185
5.11 Surface Properties for X-Ray elements . . . 185
5.11.1 Displacement, H_Misalign, and Segmented Surface Grids . . . 187
5.12 Walls: Vacuum Chamber, Capillary and Mask . . . 189
5.12.1 Wall Syntax . . . 189
5.12.2 Wall Sections . . . 190
5.12.3 Interpolation Between Sections . . . 192
5.12.4 Capillary Wall . . . 193
5.12.5 Vacuum Chamber Wall . . . 193
5.12.6 Mask Wall For Diffraction Plate and Mask Elements . . . 195
5.13 Length Attributes . . . 197
5.14 Is_on Attribute . . . 197
5.15 Multipole Attributes: Magnetic and Electric . . . 198
5.16 Field Maps . . . 199

5.16.1 Field Map Common attributes . . . 200
5.16.2 Cartesian_Map Field Map . . . 202
5.16.3 Cylindrical_Map Field Map . . . 203
5.16.4 Grid_Field Field Map . . . 204
5.16.5 Gen_Gradients . . . 207
5.17 RF Couplers . . . 209
5.18 Field Extending Beyond Element Boundary . . . 209
5.19 Automatic Phase and Amplitude Scaling of RF Fields . . . 210
5.20 Wakefields . . . 211
5.20.1 Short-Range Wakes . . . 211
5.20.2 Short-Range Longitudinal Wake With a Wake Table . . . 213
5.20.3 Short-Range Wakes — Old Format . . . 214
5.20.4 Long-Range Wakes . . . 214
5.20.5 Long-Range Wakes – Old Format . . . 216
5.21 Fringe Fields . . . 216
5.21.1 Turning On/Off Fringe Effects . . . 217
5.21.2 Fringe Types . . . 217
5.22 Instrumental Measurement Attributes . . . 219

6 Tracking, Spin, and Transfer Matrix Calculation Methods 221
6.1 Particle Tracking Methods . . . 221
6.2 Linear Transfer Map (Mat6) Calculation Methods . . . 225
6.3 Spin Tracking Methods . . . 228
6.4 Integration Methods . . . 230
6.5 CSR and Space Charge Methods . . . 230
6.5.1 ds_step and num_steps Parameters . . . 231
6.5.2 Field_calc Parameter . . . 231
6.5.3 PTC Integration . . . 232
6.6 Symplectic Versus Non-Symplectic Tracking . . . 233
6.7 Symplectify Attribute . . . 234
6.8 taylor_map_include_offsets Attribute . . . 234

7 Beam Lines and Replacement Lists 235
7.1 Branch Construction Overview . . . 235
7.2 Beam Lines and Lattice Expansion . . . 235
7.3 Line Slices . . . 237
7.4 Element Orientation Reversal . . . 237
7.5 Beam Lines with Replaceable Arguments . . . 238
7.6 Lists . . . 238
7.7 Use Statement . . . 239
7.8 Tagging Lines and Lists . . . 239

8 Superposition 241
8.1 Superposition Fundamentals . . . 241
8.2 Superposition Wrapping . . . 244
8.3 Superposition of Zero Length Elements . . . 245
8.4 Superposition and Drifts . . . 245
8.5 Superposition and Element Reversal . . . 246
8.6 Superposition and Sub-Lines . . . 246
8.7 Superposition Order . . . 247
8.8 Jumbo Super_Slaves . . . 247

8.9 Changing Element Lengths when there is Superposition . . . 248

9 Multipass 249
9.1 Multipass Fundamentals . . . 249
9.2 Reference Energy of Multipass Lords and Slaves . . . 252
9.3 Phi0_multipass . . . 252

10 Lattice File Global Parameters 255
10.1 Parameter Statements . . . 255
10.2 Particle_Start Statements . . . 258
10.3 Beam Statement . . . 258
10.4 Beginning and Line Parameter Statements . . . 259

11 Parameter Structures 261
11.1 What is a Structure? . . . 261
11.2 Bmad_Common_Struct . . . 262
11.3 PTC_Common_Struct . . . 267
11.4 Bmad_Com . . . 267
11.5 Space_Charge_Common_Struct . . . 268
11.6 Opti_DE_Param_Struct . . . 270
11.7 Dynamic Aperture Simulations: Aperture_Param_Struct . . . 271

12 Beam Initialization 273
12.1 Beam_Init_Struct Structure . . . 273
12.2 File Based Beam Initialization . . . 278
12.2.1 Old Beam ASCII Format . . . 279

13 Lattice Examples 283
13.1 Example: Injection Line . . . 283
13.2 Example: Chicane . . . 284
13.3 Example: Energy Recovery Linac . . . 285
13.4 Example: Colliding Beam Storage Rings . . . 286
13.5 Example: Rowland Circle X-Ray Spectrometer . . . 288
13.6 Example: Backward Tracking Through a Lattice . . . 290

14 Lattice File Conversion 291
14.1 Bmad to PALS Conversion . . . 291
14.2 MAD Conversion . . . 291
14.2.1 Convert MAD to Bmad . . . 291
14.2.2 Convert Bmad to MAD . . . 292
14.3 Convert to PTC . . . 292
14.4 SAD Conversion . . . 292
14.5 Elegant Conversion . . . 292
14.6 Astra, Blender, CSRTrack, GPT, and Merlin Conversion . . . 292

15 List of Element Attributes 293
15.1 !PTC_Com Element Attributes . . . 293
15.2 !Space_Charge_Com Element Attributes . . . 293
15.3 AB_multipole Element Attributes . . . 293
15.4 AC_Kicker Element Attributes . . . 294
15.5 BeamBeam Element Attributes . . . 294
15.6 Beginning Statement Attributes . . . 295

15.7 Bends: Rbend and Sbend Element Attributes . . . 296
15.8 Bmad_Com Statement Attributes . . . 296
15.9 Capillary Element Attributes . . . 297
15.10 Collimators: Ecollimator and Rcollimator Element Attributes . . . 297
15.11 Converter Element Attributes . . . 298
15.12 Crab_Cavity Element Attributes . . . 298
15.13 Crystal Element Attributes . . . 299
15.14 Custom Element Attributes . . . 299
15.15 Detector Element Attributes . . . 300
15.16 Diffraction_Plate Element Attributes . . . 300
15.17 Drift Element Attributes . . . 301
15.18 ELSeparator Element Attributes . . . 301
15.19 EM_Field Element Attributes . . . 302
15.20 E_Gun Element Attributes . . . 302
15.21 Feedback Element Attributes . . . 303
15.22 Fiducial Element Attributes . . . 303
15.23 Fixer Element Attributes . . . 303
15.24 Floor_Shift Element Attributes . . . 304
15.25 Foil Element Attributes . . . 304
15.26 Fork and Photon_Fork Element Attributes . . . 304
15.27 GKicker Element Attributes . . . 305
15.28 Girder Element Attributes . . . 305
15.29 Group Element Attributes . . . 305
15.30 Hybrid Element Attributes . . . 305
15.31 Instrument, Monitor, and Pipe Element Attributes . . . 306
15.32 Kicker Element Attributes . . . 306
15.33 Kickers: Hkicker and Vkicker Element Attributes . . . 307
15.34 Lcavity Element Attributes . . . 308
15.35 Lens Element Attributes . . . 308
15.36 Line Statement Attributes . . . 309
15.37 Marker Element Attributes . . . 309
15.38 Mask Element Attributes . . . 310
15.39 Match Element Attributes . . . 310
15.40 Mirror Element Attributes . . . 311
15.41 Multilayer_Mirror Element Attributes . . . 311
15.42 Multipole Element Attributes . . . 312
15.43 Octupole Element Attributes . . . 312
15.44 Overlay Element Attributes . . . 312
15.45 Parameter Statement Attributes . . . 313
15.46 Particle_Start Statement Attributes . . . 313
15.47 Patch Element Attributes . . . 313
15.48 Photon_Init Element Attributes . . . 314
15.49 Pickup Element Attributes . . . 314
15.50 Quadrupole Element Attributes . . . 315
15.51 RFCavity Element Attributes . . . 316
15.52 RF_Bend Element Attributes . . . 317
15.53 Ramper Element Attributes . . . 317
15.54 Sad_Mult Element Attributes . . . 318
15.55 Sample Element Attributes . . . 318
15.56 Sextupole Element Attributes . . . 319
15.57 Sol_Quad Element Attributes . . . 319

15.58 Solenoid Element Attributes . . . 320
15.59 Taylor Element Attributes . . . 320
15.60 Thick_Multipole Element Attributes . . . 321
15.61 Wiggler and Undulator Element Attributes . . . 321

II Conventions and Physics 323

16 Coordinates and Physical Layout 325
16.1 Laboratory Coordinates and Reference Orbit . . . 326
16.1.1 The Reference Orbit . . . 326
16.1.2 Element Entrance and Exit Coordinates . . . 327
16.1.3 Reference Orbit and Laboratory Coordinates Construction . . . 327
16.1.4 Patch Element Local Coordinates . . . 329
16.2 Global Coordinates . . . 330
16.2.1 Lattice Element Positioning . . . 332
16.2.2 Position Transformation When Transforming Coordinates . . . 333
16.2.3 Crystal and Mirror Element Coordinate Transformation . . . 334
16.2.4 Patch and Floor_Shift Elements Entrance to Exit Transformation . . . 334
16.2.5 Fiducial and Girder Elements Origin Shift Transformation . . . 335
16.2.6 Reflection Patch . . . 335
16.3 Transformation Between Laboratory and Element Body Coordinates . . . 335
16.3.1 Straight Element Misalignment Transformation . . . 336
16.3.2 Bend Element Misalignment Transformation . . . 336

17 Phase Space Coordinates 339
17.1 Reference Particle, Reference Energy, and Reference Time . . . 339
17.2 Charged Particle Phase Space Coordinates . . . 340
17.3 Time-based Phase Space Coordinates . . . 341
17.4 Photon Phase Space Coordinates . . . 342
17.5 MAD and Bmad Energy-Time Phase Space Coordinates . . . 342

18 Electromagnetic Fields 345
18.1 Magnetostatic Multipole Fields . . . 345
18.2 Electrostatic Multipole Fields . . . 347
18.3 Exact Multipole Fields in a Bend . . . 348
18.4 Map Decomposition of Magnetic and Electric Fields . . . 350
18.5 Cartesian Map Field Decomposition . . . 351
18.6 Cylindrical Map Decomposition . . . 353
18.6.1 DC Cylindrical Map Decomposition . . . 353
18.6.2 AC Cylindrical Map Decomposition . . . 355
18.7 Generalized Gradient Field Modeling . . . 357
18.8 RF fields . . . 358

19 Fringe Fields 359
19.1 Bend Second Order Fringe Map . . . 359
19.2 SAD Dipole Soft Edge Fringe Map . . . 361
19.3 Sad_Mult Dipole Hard Edge Fringe Map . . . 362
19.4 Linear Dipole Hard Edge Fringe Map . . . 362
19.5 Exact Dipole Hard Edge Fringe Map . . . 362
19.6 Quadrupole Soft Edge Fringe Map . . . 363
19.7 Magnetic Multipole Hard Edge Fringe Map . . . 363

19.8 Electrostatic Multipole Hard Edge Fringe Map . . . 364
19.9 RF Fringe Fields . . . 364

20 Wakefields **367**
20.1 Short–Range Wakes . . . 367
20.2 Long–Range Wakes . . . 368

21 Multiparticle Simulation **371**
21.1 Bunch Initialization . . . 371
21.1.1 Elliptical Phase Space Distribution . . . 371
21.1.2 Kapchinsky-Vladimirsky Phase Space Distribution . . . 373
21.2 Touschek Scattering . . . 374
21.3 Macroparticles . . . 375
21.4 Space Charge and Coherent Synchrotron Radiation . . . 376
21.4.1 1_Dim CSR Calculation . . . 377
21.4.2 Slice Space Charge Calculation . . . 377
21.4.3 FFT_3D Space Charge Calculation . . . 379
21.5 High Energy Space Charge . . . 379

22 Synchrotron Radiation **381**
22.1 Radiation Damping and Excitation . . . 381
22.2 Transport Map with Radiation Included . . . 382
22.3 Synchrotron Radiation Integrals . . . 384

23 Linear and Non-Linear Optics **389**
23.1 Coupling and Normal Modes . . . 389
23.2 Tunes From One-Turn Matrix Eigen Analysis . . . 391
23.3 Linear Action-Angle Coordinates . . . 393
23.4 Dispersion Calculation . . . 394
23.5 Twiss Calculation . . . 395
23.5.1 Twiss Calculation of Off-Energy Particles . . . 395
23.5.2 Twiss Calculation With Time Varying Fields . . . 396
23.6 Montague W-function . . . 396

24 Spin Dynamics **397**
24.1 Equations of Motion . . . 397
24.2 Quaternions and Biquaternions . . . 398
24.3 Invariant Spin Field . . . 400
24.4 Polarization Limits and Polarization/Depolarization Rates . . . 401
24.5 Bmad Tune and n_0 convention . . . 403
24.6 Linear $\partial \mathbf{n}/\partial \delta$ Calculation . . . 403
24.7 Single Resonance Model . . . 406
24.8 Linear Resonance Analysis . . . 407
24.9 SLIM Formalism . . . 409
24.10 Spinor Notation . . . 413

25 Taylor Maps **415**
25.1 Taylor Maps . . . 415
25.2 Spin Taylor Map . . . 416
25.3 Symplectification . . . 417
25.4 Map Concatenation and Feed-Down . . . 418
25.5 Symplectic Integration . . . 418

26 Tracking of Charged Particles **421**
26.1 Relative Versus Absolute Time Tracking . . . 421
26.2 Element Coordinate System . . . 423
26.3 Hamiltonian . . . 423
26.4 Symplectic Integration . . . 425
26.5 BeamBeam Tracking . . . 425
26.6 Bend: Exact Body Tracking with k1 = 0 . . . 427
26.7 Bend: Body Tracking with finite k1 . . . 429
26.8 Bend: Fiducial Point Calculations . . . 430
26.9 Converter Tracking . . . 432
26.10 Drift Tracking . . . 436
26.11 ElSeparator Tracking . . . 436
26.12 Foil Tracking . . . 437
26.12.1 Scattering in a Foil . . . 437
26.12.2 Energy Loss in a Foil . . . 439
26.13 Kicker, Hkicker, and Vkicker Tracking . . . 439
26.14 LCavity Tracking . . . 439
26.14.1 Integrating through RF fields . . . 439
26.14.2 kick-drift-kick model . . . 439
26.14.3 Edge kick . . . 440
26.14.4 Standing wave focusing: the pondermotive force . . . 441
26.15 Magnus Spin Tracking . . . 441
26.16 Octupole Tracking . . . 442
26.17 Patch Tracking . . . 442
26.18 Quadrupole Tracking . . . 443
26.19 RFcavity Tracking . . . 444
26.20 Sad_Mult Tracking . . . 445
26.21 Sextupole Tracking . . . 445
26.22 Sol_Quad Tracking . . . 445
26.23 Solenoid Tracking . . . 447
26.24 Sprint Spin Tracking . . . 448
26.24.1 SBend Body, $k_1 = 0$ . . . 448
26.24.2 Sbend Body, $k_1 \neq 0$ . . . 449
26.24.3 Sbend Entrance Fringe . . . 449
26.24.4 Quadrupole . . . 449
26.24.5 Solenoid Element Body . . . 450
26.24.6 Solenoid Entrance Fringe . . . 450
26.25 Symplectic Tracking with Cartesian Modes . . . 450

27 Tracking of X-Rays **453**
27.1 Coherent and Incoherent Photon Simulations . . . 453
27.1.1 Incoherent Photon Tracking . . . 453
27.1.2 Coherent Photon Tracking . . . 454
27.1.3 Partially Coherent Photon Simulations . . . 456
27.2 Element Coordinate System . . . 456
27.2.1 Transform from Laboratory Entrance to Element Coordinates . . . 456
27.2.2 Transform from Element Exit to Laboratory Coordinate . . . 457
27.3 Mirror and Crystal Element Transformation . . . 457
27.3.1 Transformation from Laboratory to Element Coordinates . . . 457
27.3.2 Transformation from Element to Laboratory Coordinates . . . 458
27.4 Crystal Element Tracking . . . 460

27.4.1 Calculation of Entrance and Exit Bragg Angles . . . 460
27.4.2 Crystal Coordinate Transformations . . . 462
27.4.3 Laue Reference Orbit . . . 464
27.4.4 Crystal Surface Reflection and Refraction . . . 465
27.4.5 Bragg Crystal Tracking . . . 466
27.4.6 Coherent Laue Crystal Tracking . . . 468
27.5 X-ray Targeting . . . 468

28 Simulation Modules **471**
28.1 Instrumental Measurements . . . 471
28.1.1 Orbit Measurement . . . 471
28.1.2 Dispersion Measurement . . . 472
28.1.3 Coupling Measurement . . . 472
28.1.4 Phase Measurement . . . 473

29 Using PTC/FPP **475**
29.1 PTC Tracking Versus Bmad Tracking . . . 475
29.2 PTC / Bmad Interfacing . . . 475

III Programmer's Guide **477**

30 Bmad Programming Overview **479**
30.1 Manual Notation . . . 479
30.2 The Bmad Libraries . . . 479
30.3 Using getf and listf for Viewing Routine and Structure Documentation . . . 481
30.4 Precision of Real Variables . . . 483
30.5 Programming Conventions . . . 483
30.6 Using Modules . . . 484

31 An Example Bmad Based Program **485**
31.1 Programming Setup . . . 485
31.2 A First Program . . . 485
31.3 Explanation of the Simple_Bmad_Program . . . 487

32 The ele_struct **491**
32.1 Initialization and Pointers . . . 494
32.2 Miscellaneous . . . 494
32.3 Element Attribute Bookkeeping . . . 494
32.4 String Components . . . 495
32.5 Element Key . . . 495
32.6 The %value(:) array . . . 495
32.7 Connection with the Lat_Struct . . . 496
32.8 Limits . . . 497
32.9 Twiss Parameters, etc. . . . 497
32.10 Element Lords and Element Slaves . . . 497
32.11 Group and Overlay Controller Elements . . . 497
32.12 Coordinates, Offsets, etc. . . . 498
32.13 Transfer Maps: Linear and Non-linear (Taylor) . . . 499
32.14 Reference Energy and Time . . . 500
32.15 EM Field Maps . . . 500
32.16 Wakes . . . 500

32.17 Wiggler and Undulator Fields . . . 501
32.18 Spin Transport . . . 502
32.19 Radiation Maps . . . 502
32.20 High-Energy Space Charge . . . 502
32.21 Converter, Foil, and RF Components . . . 503
32.22 Multipoles . . . 503
32.23 Tracking Methods . . . 503
32.24 Custom and General Use Attributes . . . 503
32.25 Bmad Reserved Variables . . . 504

33 The lat_struct 505
33.1 Initializing . . . 506
33.2 Pointers: Allocation and Deallocation . . . 506
33.3 Branches in the lat_struct . . . 506
33.4 Param_struct Component . . . 508
33.5 Elements Controlling Other Elements . . . 509
33.6 Lattice Bookkeeping . . . 514
33.7 Intelligent Bookkeeping . . . 515
33.8 particle_start Component . . . 516
33.9 Custom Parameters . . . 516

34 Lattice Element Manipulation 519
34.1 Creating Element Slices . . . 519
34.2 Adding and Deleting Elements From a Lattice . . . 519
34.3 Finding Elements . . . 520
34.4 Accessing Named Element Attributes . . . 521

35 Reading and Writing Lattices 523
35.1 Reading in Lattices . . . 523
35.2 Digested Files . . . 523
35.3 Writing Lattice files . . . 524

36 Normal Modes: Twiss Parameters, Coupling, Emittances, Etc. 525
36.1 Components in the Ele_struct . . . 525
36.2 Tune and Twiss Parameter Calculations . . . 526
36.3 Tune Setting . . . 527
36.4 Emittances & Radiation Integrals . . . 527
36.5 Chromaticity Calculation . . . 528

37 Tracking and Transfer Maps 529
37.1 The coord_struct . . . 529
37.2 Tracking Through a Single Element . . . 532
37.3 Tracking Through a Lattice Branch . . . 533
37.4 Forking from Branch to Branch . . . 535
37.5 Multi-turn Tracking . . . 536
37.6 Closed Orbit Calculation . . . 536
37.7 Partial Tracking through elements . . . 536
37.8 Apertures . . . 537
37.9 Custom Tracking . . . 537
37.10 Tracking Methods . . . 538
37.11 Using Time as the Independent Variable . . . 538
37.12 Absolute/Relative Time Tracking . . . 538

37.13 Taylor Maps . . . 538
37.14 Tracking Backwards . . . 539
37.15 Reversed Elements and Tracking . . . 540
37.16 Beam (Particle Distribution) Tracking . . . 540
37.17 Spin Tracking . . . 541
37.18 X-ray Targeting . . . 541

38 Miscellaneous Programming **543**
38.1 Custom and Hook Routines . . . 543
38.2 Custom Calculations . . . 544
38.3 Hook Routines . . . 546
38.4 Nonstandard Init_Custom Use . . . 547
38.5 Physical and Mathematical Constants . . . 547
38.6 Global Coordinates and S-positions . . . 548
38.7 Reference Energy and Time . . . 548
38.8 Global Common Structures . . . 549
38.9 Parallel Processing . . . 549
38.10 Expression Trees and Expression Stacks . . . 550

39 PTC/FPP Programming **551**
39.1 Phase Space . . . 551
39.2 PTC Initialization . . . 552
39.3 PTC Structures Compared to Bmad's . . . 552
39.4 Variable Initialization and Finalization . . . 553
39.5 Correspondence Between Bmad Elements and PTC Fibres . . . 553
39.6 Taylor Maps . . . 554
39.7 Patches . . . 554
39.8 Number of Integration Steps & Integration Order . . . 554
39.9 Creating a PTC layout from a Bmad lattice . . . 555
39.10 Internal_State . . . 555

40 OPAL **557**
40.1 Phase Space . . . 557

41 C++ Interface **559**
41.1 C++ Classes and Enums . . . 559
41.2 Conversion Between Fortran and C++ . . . 560

42 Quick_Plot Plotting **563**
42.1 An Example . . . 565
42.2 Plotting Coordinates . . . 566
42.3 Length and Position Units . . . 567
42.4 Y2 and X2 axes . . . 568
42.5 Text . . . 568
42.6 Styles . . . 568
42.7 Structures . . . 573

43 HDF5 **575**
43.1 HDF5 Particle Beam Data Storage . . . 575
43.2 HDF5 Grid_Field Data Storage . . . 577

44 Bmad Library Routine List **579**
44.1 Beam: Low Level Routines . . . 581
44.2 Beam: Tracking and Manipulation . . . 581
44.3 Branch Handling Routines . . . 582
44.4 Coherent Synchrotron Radiation (CSR) . . . 582
44.5 Collective Effects . . . 583
44.6 Custom Routines . . . 583
44.7 Electro-Magnetic Fields . . . 584
44.8 HDF Read/Write . . . 584
44.9 Helper Routines: File, System, and IO . . . 586
44.10 Helper Routines: Math (Except Matrix) . . . 587
44.11 Helper Routines: Matrix . . . 588
44.12 Helper Routines: Miscellaneous . . . 588
44.13 Helper Routines: String Manipulation . . . 589
44.14 Helper Routines: Switch to Name . . . 591
44.15 Inter-Beam Scattering (IBS) . . . 591
44.16 Lattice: Element Manipulation . . . 592
44.17 Lattice: Geometry . . . 594
44.18 Lattice: Informational . . . 595
44.19 Lattice: Low Level Stuff . . . 597
44.20 Lattice: Manipulation . . . 597
44.21 Lattice: Miscellaneous . . . 599
44.22 Lattice: Nametable . . . 599
44.23 Lattice: Reading and Writing Files . . . 600
44.24 Matrices . . . 600
44.25 Matrix: Low Level Routines . . . 602
44.26 Measurement Simulation Routines . . . 602
44.27 Multipass . . . 602
44.28 Multipoles . . . 603
44.29 Nonlinear Optimizers . . . 603
44.30 Overloading the equal sign . . . 604
44.31 Particle Coordinate Stuff . . . 604
44.32 Photon Routines . . . 605
44.33 PTC Interface Routines . . . 605
44.34 Quick Plot Routines . . . 606
44.34.1 Quick Plot Page Routines . . . 606
44.34.2 Quick Plot Calculational Routines . . . 606
44.34.3 Quick Plot Drawing Routines . . . 607
44.34.4 Quick Plot Set Routines . . . 608
44.34.5 Informational Routines . . . 610
44.34.6 Conversion Routines . . . 610
44.34.7 Miscellaneous Routines . . . 611
44.34.8 Low Level Routines . . . 611
44.35 Spin . . . 612
44.36 Transfer Maps: Routines Called by make_mat6 . . . 613
44.37 Transfer Maps: Complex Taylor Maps . . . 613
44.38 Transfer Maps: Taylor Maps . . . 614
44.39 Transfer Maps: Driving Terms . . . 616
44.40 Tracking and Closed Orbit . . . 616
44.41 Tracking: Low Level Routines . . . 618
44.42 Tracking: Mad Routines . . . 618

44.43 Tracking: Routines called by track1 . . . 619
44.44 Twiss and Other Calculations . . . 620
44.45 Twiss: 6 Dimensional . . . 621
44.46 Wakefields . . . 621
44.47 C/C++ Interface . . . 622

IV Appendices **623**

A Obsolete Straight-Coordinate Generalized Gradients **625**

V Bibliography and Index **629**

Bibliography **631**

Routine Index **637**

Index **645**

List of Figures

2.1 Superposition example. . . . 31

4.1 Coordinate systems for rbend and sbend elements. . . . 71
4.2 Crystal element geometry. . . . 83
4.3 Foil geometry. . . . 102
4.4 Example with photon_fork elements. . . . 105
4.5 Girder example. . . . 109
4.6 Patch Element. . . . 136

5.1 Geometry of Pitch and Offset attributes . . . 175
5.2 Geometry of a Tilt . . . 176
5.3 Geometry of a Bend . . . 176
5.4 Geometry of a photon reflecting element orientation . . . 177
5.5 Apertures for ecollimator and rcollimator elements. . . . 179
5.6 Surface curvature geometry. . . . 184
5.7 Capillary or vacuum chamber wall. . . . 190
5.8 Convex cross-sections do not guarantee a convex volume. . . . 192
5.9 vacuum chamber crotch geometry. . . . 194
5.10 Example mask wall . . . 195
5.11 `Field map`coordinates when used with a bend element. . . . 200

6.1 Dark current tracking. . . . 222

8.1 Superposition example. . . . 241
8.2 Superposition Offset. . . . 242

13.1 Injection line into a dipole magnet. . . . 283
13.2 Four bend chicane. . . . 284
13.3 Example Energy Recovery Linac. . . . 286
13.4 Dual ring colliding beam machine . . . 286
13.5 Rowland circle spectrometer . . . 288

16.1 The three coordinate system used by *Bmad*. . . . 325
16.2 The local Reference System. . . . 326
16.3 Lattice elements as LEGO blocks. . . . 327
16.4 Laboratory coordinates construction. . . . 328
16.5 The local reference coordinates in a `patch`element. . . . 329
16.6 The Global Coordinate System . . . 330
16.7 Orientation of a Bend. . . . 332
16.8 Mirror and crystal geometry . . . 333

17.1 Interpreting phase space z at constant velocity. . . . 340

21.1 CSR Calculation . . . 377

23.1 Illustration of a positive tune . . . 392

26.1 Element Coordinate System. . . . 423
26.2 Geometry for the exact bend calculation. . . . 427
26.3 Geometry with `fiducial_pt`set to (a) `entrance_end`and (b) `center`. In both cases, $\mathbf{r}_1$ and $\mathbf{r}_2$ are the entrance and exit reference points before and $\mathbf{r}_1'$ and $\mathbf{r}_2$ are the entrance and exit points after variation of one of `rho`, `g`, `b_field`, or `angle`. Similarly, ρ and α are the bending radius and bending angle before variation while ρ' and α' are the bending radius afterwards. Finally, e_1 e_2 are the face angles and rectangular length before variation, and L_r' and $\mathbf{r}_0'$ are the rectangular length and center of curvature after variation. . . . 431
26.4 Converter geometry. . . . 433
26.5 ElSeparator electric field. . . . 436
26.6 Standard patch transformation. . . . 443
26.7 Solenoid with a hard edge. . . . 447

27.1 Crystal, Mirror, and Multilayer_Mirror Element Coordinates. . . . 457
27.2 Reference trajectory reciprocal space diagram for crystal diffraction. . . . 460
27.3 Reference energy flow for Laue diffraction . . . 464

31.1 Example Bmad program . . . 486
31.2 Output from the example program . . . 489

32.1 The `ele_struct`(part 1). . . . 492
32.2 The `ele_struct`(part 2). . . . 493

33.1 Definition of the `lat_struct`. . . . 505
33.2 Definition of the `param_struct`. . . . 508
33.3 Example of multipass combined with superposition . . . 511

37.1 Condensed track_all code. . . . 534

39.1 PTC structure relationships . . . 553

41.1 Example Fortran routine calling a *C++* routine. . . . 560
41.2 Example *C++* routine callable from a Fortran routine. . . . 560

42.1 *Quick Plot* example program. . . . 564
42.2 Output of plot_example.f90. . . . 565
42.3 A Graph within a Box within a Page. . . . 566
42.4 Continuous colormap used by `qp_continuous_color`. Typical usage: `call qp_routine(..., color = qp_continuous_color(0.25_rp), ...)` . . . 569

List of Tables

3.1 Physical units used by *Bmad*. . . . 48
3.2 Physical and mathematical constants recognized by *Bmad*. . . . 49

4.1 Table of element types suitable for use with charged particles. Also see Table 4.3 61
4.2 Table of element types suitable for use with photons. Also see Table 4.3 62
4.3 Table of controller elements. . . . 62

5.1 Table of dependent variables. . . . 167
5.2 Example normalized and unnormalized field strength attributes. . . . 168

6.1 Table of available tracking_method switches. . . . 224
6.2 Actual `mat6_calc_method`used with `auto`setting. . . . 225
6.3 Table of available mat6_calc_method switches. . . . 227
6.4 Table of available spin_tracking_method switches. . . . 229

18.1 F and $n_{\rm ref}$ for various elements. . . . 347

33.1 Bounds of the root branch array. . . . 507
33.2 Possible element %lord_status/%slave_status combinations. . . . 511

42.1 Plotting Symbols at Height = 40.0 571
42.2 PGPLOT Escape Sequences. . . . 572
42.3 Roman to Greek Character Conversion 572

Part I

Language Reference

Chapter 1

Orientation

1.1 What is Bmad?

Bmad is an open-source software library (aka toolkit) for simulating charged particles and X-rays. *Bmad* is not a program itself but is used by programs for doing calculations. The advantage of *Bmad* over a stand-alone simulation program is that when new types of simulations need to be developed, *Bmad* can be used to cut down on the time needed to develop such programs with the added benefit that the number of programming errors will be reduced.

Over the years, *Bmad* has been used for a wide range of charged-particle and X-ray simulations. This includes:

```
Lattice design                                  X-ray simulations
Spin tracking                                   Wakefields and HOMs
Beam breakup (BBU) simulations in ERLs          Touschek Simulations
Intra-beam scattering (IBS) simulations         Dark current tracking
Coherent Synchrotron Radiation (CSR)            Frequency map analysis
```

1.2 Tao and Bmad Distributions

The strength of *Bmad* is that, as a subroutine library, it provides a flexible framework from which sophisticated simulation programs may easily be developed. The weakness of *Bmad* comes from its strength: *Bmad* cannot be used straight out of the box. Someone must put the pieces together into a program. To remedy this problem, the *Tao* program[Tao] has been developed. *Tao*, which uses *Bmad* as its simulation engine, is a general purpose program for simulating particle beams in accelerators and storage rings. Thus *Bmad* combined with *Tao* represents the best of both worlds: The flexibility of a software library with the ease of use of a program.

Besides the *Tao* program, an ecosystem of *Bmad* based programs has been developed. These programs, along with *Bmad*, are bundled together in what is called a *Bmad* `Distribution` which can be downloaded from the web. The following is a list of some of the more commonly used programs.

bmad_to_mad_sad_elegant
: The `bmad_to_mad_sad_elegant` program converts *Bmad* lattice format files to `MAD8`, `MADX`, `Elegant` and `SAD` format.

bbu
The `bbu` program simulates the beam breakup instability in Energy Recovery Linacs (ERLs).

dynamic_aperture
The `dynamic_aperture` program finds the dynamic aperture through tracking.

ibs_linac
The `ibs_linac` program simulates the effect of intra-beam scattering (IBS) for beams in a Linac.

ibs_ring
The `ibs_ring` program simulates the effect of intra-beam scattering (IBS) for beams in a ring.

long_term_tracking
The `long_term_tracking_program` is for long term tracking of a particle or beam possibly including tracking of the spin.

lux
The `lux` program simulates X-ray beams from generation through to experimental end stations.

mad8_to_bmad.py, madx_to_bmad.py
These python programs will convert `MAD8` and `MADX` lattice files to to *Bmad* format.

moga
The `moga` (multiobjective genetic algorithms) program does multiobjective optimization.

synrad
The `synrad` program computes the power deposited on the inside of a vacuum chamber wall due to synchrotron radiation from a particle beam. The calculation is essentially two dimensional but the vertical emittance is used for calculating power densities along the centerline. Crotch geometries can be handled as well as off axis beam orbits.

synrad3d
The `synrad3d` program tracks, in three dimensions, photons generated from a beam within the vacuum chamber. Reflections at the chamber wall is included.

tao
Tao is a general purpose simulation program.

1.3 Resources: More Documentation, Obtaining Bmad, etc.

More information and download instructions are readily available at the *Bmad* web site:

`www.classe.cornell.edu/bmad/`

Links to the most up-to-date *Bmad* and *Tao* manuals can be found there as well as manuals for other programs and instructions for downloading and setup.

The *Bmad* manual is organized as reference guide and so does not do a good job of instructing the beginner as to how to use *Bmad*. For that there is an introduction and tutorial on *Bmad* and *Tao* (§1.2) concepts that can be downloaded from the *Bmad* web page. Go to either the *Bmad* or *Tao* manual pages and there will be a link for the tutorial.

1.4 PTC: Polymorphic Tracking Code

The PTC/FPP library of Étienne Forest handles Taylor maps to any arbitrary order. This is also known as Truncated Power Series Algebra (TPSA). The core Differential Algebra (DA) package used by PTC/FPP was developed by Martin Berz[Berz89]. The PTC/FPP libraries are interfaced to *Bmad* so that calculations that involve both *Bmad* and PTC/FPP can be done in a fairly seamless manner.

Basically, the FPP ("Fully Polymorphic Package") part of the PTC/FPP code handles Taylor map manipulation. This is purely mathematical. FPP has no knowledge of accelerators, magnetic fields, particle tracking etc. PTC ("Polymorphic Tracking Code") implements the physics and uses FPP to handle the Taylor map manipulation. Since the distinction between `FPP` and `PTC` is irrelevant to the non-programmer, "PTC" will be used to refer to the entire PTC/FPP package.

PTC is used by *Bmad* when constructing Taylor maps and when the `tracking_method` §6.1) is set to `symp_lie_ptc`. All Taylor maps above first order are calculated via PTC. No exceptions.

For more discussion of PTC see Chapter §29. For the programmer, also see Chapter §39.

For the purposes of this manual, PTC and FPP are generally considered one package and the combined PTC/FPP library will be referred to as simply "PTC".

Chapter 2

Bmad Concepts and Organization

This chapter is an overview of some of the nomenclature used by *Bmad*. Presented are the basic concepts, such as `element`, `branch`, and `lattice`, that *Bmad* uses to describe such things as LINACs, storage rings, X-ray beam lines, etc.

2.1 Lattice Elements

The basic building block *Bmad* uses to describe a machine is the `lattice element`. An element can be a physical thing that particles travel "through" like a bending magnet, a quadrupole or a Bragg crystal, or something like a `marker` element (§4.32) that is used to mark a particular point in the machine. Besides physical elements, there are `controller` elements (Table 4.3) that can be used for parameter control of other elements.

Chapter §4 lists the complete set of different element types that *Bmad* knows about.

In a lattice `branch` (§2.2), The ordered array of elements are assigned a number (the element index) starting from zero. The zeroth `beginning_ele` (§4.4) element, which is always named `BEGINNING`, is automatically included in every branch and is used as a marker for the beginning of the branch. Additionally, every branch will, by default, have a final marker element (§4.32) named `END`.

2.2 Lattice Branches

The next level up from a `lattice element` is the `lattice branch`. A `lattice branch` contains an ordered sequence of lattice elements that a particle will travel through. A branch can represent a LINAC, X-Ray beam line, storage ring or anything else that can be represented as a simple ordered list of elements.

Chapter §7 shows how a `branch` is defined in a lattice file with `line`, `list`, and `use` statements.

A `lattice` (§2.3), has an array of `branches`. Each `branch` in this array is assigned an index starting from 0. Additionally, each `branch` is assigned a name which is the `line` that defines the branch (§7.7).

Branches can be interconnected using `fork` and `photon_fork` elements (§4.23). This is used to simulate forking beam lines such as a connections to a transfer line, dump line, or an X-ray beam line. A `branch` from which other `branches` fork but is not forked to by any other `branch` is called a `root` branch. A

branch that is forked to by some other branch is called a `downstream` branch.

2.3 Lattice

an array of `branches` that can be interconnected together to describe an entire machine complex. A `lattice` can include such things as transfer lines, dump lines, x-ray beam lines, colliding beam storage rings, etc. All of which can be connected together to form a coherent whole. In addition, a lattice may contain `controller elements` (Table 4.3) which can simulate such things as magnet power supplies and lattice element mechanical support structures.

Branches can be interconnected using `fork` and `photon_fork` elements (§4.23). This is used to simulate forking beam lines such as a connections to a transfer line, dump line, or an X-ray beam line. The `branch` from which other `branches` fork but is not forked to by any other `branch` is called a `root` branch.

A lattice may contain multiple `root branches`. For example, a pair of intersecting storage rings will generally have two `root` branches, one for each ring. The `use` statement (§7.7) in a lattice file will list the `root branches` of a lattice. To connect together lattice elements that are physically shared between branches, for example, the interaction region in colliding beam machines, `multipass` lines (§9) can be used.

The root branches of a lattice are defined by the `use` (§7.7) statement. To further define such things as dump lines, x-ray beam lines, transfer lines, etc., that branch off from a root branch, a forking element is used. `Fork` elements can define where the particle beam can branch off, say to a beam dump. `photon_fork` elements can define the source point for X-ray beams. Example:

```
erl: line = (..., dump, ...)              ! Define the root branch
use, erl
dump: fork, to_line = d_line               ! Define the fork point
d_line: line = (..., q3d, ...)             ! Define the branch line
```

Like the root branch *Bmad* always automatically creates an element with `element index` 0 at the beginning of each branch called `beginning`. The longitudinal `s` position of an element in a branch is determined by the distance from the beginning of the branch.

Branches are named after the line that defines the `branch`. In the above example, the branch line would be named `d_line`. The root branch, by default, is called after the name in the `use` statement (§7.7).

The "branch qualified" name of an element is of the form

```
branch_name>>element_name
```

where `branch_name` is the name of the branch and `element_name` is the "regular" name of the element. Example:

```
root>>q10w
xline>>cryst3
```

When parsing a lattice file, branches are not formed until the lattice is expanded (§3.26). Therefore, an `expand_lattice` statement is required before branch qualified names can be used in statements. See §3.6 for more details.

2.4 Lord and Slave Elements

A real machine is more than a collection of independent lattice elements. For example, the field strength in a string of elements may be tied together via a common power supply, or the fields of different elements may overlap.

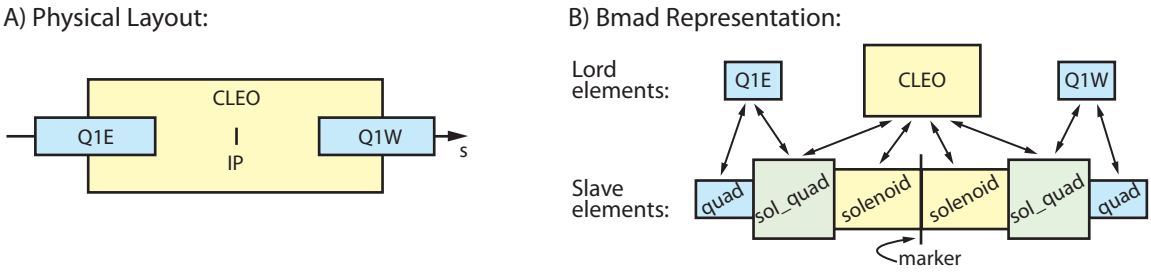

Figure 2.1: Superposition Example. A) Interaction region layout with quadrupoles overlapping a solenoid. B) The Bmad lattice representation has a list of split elements to track through and the undivided "lord" elements. Pointers (double headed arrows), keep track of the correspondence between the lords and their slaves.

Bmad tries to capture these interdependencies using what are referred to as `lord` and `slave` elements. The `lord` elements may be divided into two classes. In one class are the `controller` elements. These are `overlay` (§4.40), `group` (§4.26), `ramper` (§4.44), and `girder` (§4.24) elements that control the attributes of other elements which are their slaves.

The other class of `lord` elements embody the separation of the physical element from the track that a particle takes when it passes through the element. There are two types

An example will make this clear. `Superposition` (§8) is the ability to overlap lattice elements spatially. Fig. 2.1 shows an example which is a greatly simplified version of the IR region of Cornell's CESR storage ring when CESR was an e+/e– collider. As shown in Fig. 2.1A, two quadrupoles named `q1w` and `q1e` are partially inside and partially outside the interaction region solenoid named `cleo`. In the lattice file, the IR region layout is defined to be

```
cesr: line = (... q1e, dft1, ip, dft1, q1w ...)
cleo: solenoid, l = 3.51, superimpose, ref = ip
```

The line named `cesr` ignores the solenoid and just contains the interaction point marker element named `ip` which is surrounded by two drifts named `dft1` which are, in turn, surrounded by the `q1w` and `q1e` quadrupoles. The solenoid is added to the layout on the second line by using superposition. The "ref = ip" indicates that the solenoid is placed relative to `ip`. The default, which is used here, is to place the center of the superimposed `cleo` element at the center of the `ip` reference element. The representation of the lattice in *Bmad* will contain two branch `sections` ("sections" is explained more fully later): One section, called the `tracking section`, contains the elements that are needed for tracking particles. In the current example, as shown in Fig. 2.1B, the first IR element in the tracking section is a quadrupole that represents the part of `q1e` outside of the solenoid. The next element is a combination solenoid/quadrupole, called a `sol_quad`, that represents the part of `q1e` inside `cleo`, etc. The other branch section that Bmad creates is called the `lord section` This section contain the undivided "physical" `super_lord` elements (§8) which, in this case are `q1e`, `q1w`, and `cleo`. Pointers are created between the lords and their `super_slave` elements in the tracking section so that changes in parameters of the lord elements can be transferred to their corresponding slaves.

`super_lord`s are used when there are overlapping fields between elements, the other case where there is a separation between the physical (lord) element and the (slave) element(s) used to track particles through comes when a particle passes through the same physical element multiple times such as in an Energy Recovery Linac or where different beams pass through the same element such as in an interaction region. In this case, `multipass_lords` representing the physical elements and `multipass_slaves` elements which are used for tracking can be defined (§9). Superposition and multipass can be combined in situations where there are overlapping fields in elements where the particle passes through

Each lattice element is assigned a `slave_status` indicating what kind of slave it is and a `lord_status` indicating what kind of lord it is. Normally a user does not have to worry about this since these status attributes are handled automatically by *Bmad*. The possible `lord_status` settings are:

girder_lord
A `girder_lord` element is a `girder` element (§4.24).

multipass_lord
`multipass_lord` elements are created when multipass lines are present (§9).

overlay_lord
An `overlay_lord` is an `overlay` element (§4.40).

group_lord
A `group_lord` is a `group` element (§4.26).

super_lord
A `super_lord` element is created when elements are superimposed on top of other elements (§8).

not_a_lord
This element does not control anything.

Any element whose `lord_status` is something other than `not_a_lord` is called a `lord` element. In the `tracking part` of the branch, `lord_status` will always be `not_a_lord`. In the `lord section` of the branch, under normal circumstances, there will never be any `not_a_lord` elements.

Lord elements are divided into two classes. A `major` lord represents a physical element which the slave elements are a part of. `super_lord`s and `multipass_lord`s are `major` lords. As a consequence, a `major` lord is a lord that controls nearly all of the attributes of its slaves. The other lords — `girder_lord`s, `group_lord`s and `overlay_lord`s — are called `minor` lords. These lords only control some subset of a slaves attributes.

The possible `slave_status` settings are

multipass_slave
A `multipass_slave` element is the slave of a `multipass_lord` (§9).

slice_slave
A `slice_slave` element represents a longitudinal slice of another element. Slice elements are not part of the lattice but rather are created on-the-fly when, for example, a program needs to track part way through an element.

super_slave
A `super_slave` element is an element in the tracking part of the branch that has one or more `super_lord` lords (§8).

minor_slave
`minor_slave` elements are elements that are not `slice_slave`s and are only controlled by `minor` lords (`overlay_lord`s, `group_lord`s, or `girder_lord`s).

free
A `free` element is an element with no lords.

For historical reasons, each `branch` in a lattice has a `tracking section` and a `lord section` and the `tracking section` is always the first (lower) part of the element array and the `lord section` inhabits the second (upper) part of the array. All the `lord` elements are put in the `lord section` of branch 0 and all the other `lord sections` of all the other branches are empty.

As a side note, Étienne Forest's PTC code (§1.4) uses separate structures to separate the physical element, which PTC calls an `element` from the particle track which PTC call a `fibre`. [Actually, PTC has two structures for the physical element, `element` and `elementp`. The latter being the "polymorph" version.] This `element` and `fibre` combination corresponds to *Bmad* `multipass_lord` and `multipass_slave` elements. PTC does not handle overlapping fields as *Bmad* does with `superposition` (§8).

Chapter 3

Lattice File Statements

A lattice (§2) defines the sequence of elements that a particle will travel through along with the attributes (length, strength, orientation, etc.) of the elements. A lattice file (or files) is a file that is used to describe an accelerator or storage ring.

When *Bmad* was first developed, The *Bmad* lattice file syntax was modeled upon the format defined for the *MAD* program[Grote96]. Since then, the *Bmad* format has been developed to meet ever increasing simulation needs so currently there are many differences between the two formats. One difference, which has been present from the very start, is that there are no "action" commands (action commands tell the program to calculate the Twiss parameters, do tracking, etc.) in a *Bmad* lattice file. The reason for this is due to the fact that *Bmad* is a software library and not a program. That is, interacting with the user to determine what actions a *Bmad* based program should take is left to the program itself and is not part of the *Bmad* standard.

3.1 File Example and Syntax

The following (rather silly) example shows some of the features of a *Bmad* lattice file:

```
! This is a comment
parameter[E_TOT] = 5e9                    ! Parameter definition
pa1 = sin(3.47 * pi / c_light)                  ! Constant definition
bend1: sbend, type = "arc bend", l = 2.3,       ! An element definition
    g = 2*pa1, tracking_method = bmad_standard
bend2: bend1, l = 3.4                           ! Another element def
bend2[g] = 105 - exp(2.3) / 37.5                ! Redefining an attribute
ln1: line = (ele1, ele2, ele3)                  ! A line definition
ln2: line = (ln1, ele4, ele5)                   ! Lines can contain lines
arg_ln(a, b): line = (ele1, a, ele2, b)         ! A line with arguments.
use, ln2                                        ! Which line to use for the lattice
```

A *Bmad* lattice file consists of a sequence of statements. An exclamation mark (!) denotes a comment and the exclamation mark and everything after the exclamation mark on a line are ignored.

Bmad is generally case insensitive. Most names are converted to uppercase. Exceptions are file names and atomic formulas for materials used in crystal diffraction. Also `type`, `alias`, and `descrip` string labels, which can be set for any element, are not converted (§5.3). Notice that commands for the `Tao`

program are case sensitive. The reason for this difference is due to one of history. *Bmad* was developed back in the days where case sensitivity was not important.

Normally a statement occupies a single line in the file. Several statements may be placed on the same line by inserting a semicolon (";") between them. A long statement can occupy multiple lines by putting an ampersand ("&") at the end of each line of the statement except for the last line. Additionally, lines that end with an "implicit end-line continuation character" are automatically continued to the next line and lines that begin with an "implicit begin-line continuation character are automatically appended to the previous line. The implicit end-line continuation characters are:

```
,   (   {   [   =
```

The implicit begin-line continuation characters are:

```
,   )   }   ]   =
```

The following example shows a command extending over four lines

```
wall = {                ! End-line continuation character.
  section = {s = 0.45
  }                     ! Begin-line continuation character.
}                       ! Begin-line continuation character.
```

Note: The symbols "+", "-", "*", and "/" are *not* valid implicit continuation characters. The reason for this is to avoid confusion when a species name (for example "He++") comes at the end of a line.

A single line in a lattice file is limited to 500 characters. Commands have no restrictions as to the number of lines they may be continued over and there is no limit to the length of a command.

Names of constants, elements, lines, etc. are limited to 40 characters. The first character must be a letter (`A` — `Z`). The other characters may be a letter, a digit (`0` — `9`) or an underscore (`_`). Other characters may appear but should be avoided since they are used by Bmad for various purposes. For example, the backslash (\) character is used to by Bmad when forming the names of superposition slaves (§8) and dots (`.`) are used by Bmad when creating names of `tagged` elements (§7.8). Also use of special characters may make the lattice files less portable to non-Bmad programs.

The following example constructs a linear lattice with two elements:

```
parameter[geometry] = open
parameter[e_tot] =2.7389062E9
parameter[particle] = POSITRON
beginning[beta_a] = 14.5011548
beginning[alpha_a] = -0.53828197
beginning[beta_b] = 31.3178048
beginning[alpha_b] = 0.25761815
q: quadrupole, l = 0.6, b1_gradient = 9.011
d: drift, l = 2.5
t: line = (q, d)
use, t
```

here `parameter[geometry]` (§10.1) is set to `open` which specifies that the lattice is not circular. In this case, the beginning Twiss parameters need to be specified and this is done by the `beginning` statements (§10.4). A quadrupole named `q` and a drift element named `d` are specified and the entire lattice consists of element `q` followed by element `d`.

3.2 Digested Files

Normally the *Bmad* parser routine will create what is called a "digested file" after it has parsed a lattice file so that when a program is run and the same lattice file is to be read in again, to save time, the

digested file can be used to load in the lattice information. This digested file is in binary format and is not human readable. The digested file will contain the transfer maps for all the elements. Using a digested file can save considerable time if some of the elements in the lattice need to have Taylor maps computed. (this occurs typically with map–type wigglers).

Bmad creates the digested file in the same area as the lattice file. If *Bmad* is not able to create a digested file (typically because it does not have write permission in the directory), an error message will be generated but otherwise program operation will be normal.

Digested files contain the names of the lattice files used to create them. If a lattice file has been modified since the digested file has been created then the lattice files will be reread and a new digested file will be generated.

Note: If any of the random number functions (§3.14) are used in the process of creating the lattice, the digested file will be ignored. In this case, each time the lattice is read into a program, different random numbers will be generated for expressions that use such random numbers.

Digested files can also be used for easy transport of lattices between programs or between sessions of a program. For example, using one program you might read in a lattice, make some adjustments (say to model shifts in magnet positions) and then write out a digested version of the lattice. This adjusted lattice can now be read in by another program.

3.3 Element Sequence Definition

A `line` defines a sequence of elements. `lines` may contain other `lines` and so a hierarchy may be established. One line is selected, via a `use` statement, that defines the lattice. For example:

```
l3: line = (l1, l2)   ! Concatenate two lines
l1: line = (a, b, c)  ! Line with 3 elements
l2: line = (a, z)     ! Another line
use, l3               ! Use l3 as the lattice definition.
```

In this case the lattice would be

```
(a, b, c, a, z)
```

`Lines` can be defined in any order. See Chapter 7 for more details.

The `superimpose` construct allows elements to be placed in a lattice at a definite longitudinal position. What happens is that after a lattice is expanded, there is a reshuffling of the elements to accommodate any new superimpose elements. See §8 for more details.

3.4 Lattice Elements

The syntax for defining a lattice element roughly follows the *MAD* [Grote96] program:

```
ele_name: keyword [, attributes]
```

where `ele_name` is the element name, `keyword` is the type of element, and `attributes` is a list of the elements attributes. Chapter 4 gives a list of elements types with their attributes. `Overlay` and `group` type elements have a slightly different syntax:

```
ele_name: keyword = { list }, master-attribute [= value] [, attributes]
```

and `Girder` elements have the syntax

```
ele_name: keyword = { list } [, attributes]
```

For example:

```
q01w: quadrupole, type = "A String", l = 0.6, tilt = pi/4
h10e: overlay = { b08e, b10e }, var = {hkick}
```

The `keyword` specifying the element type can be the name of an element that has already been declared. In this case, the element being defined will inherit the attributes that have been set for the element associated with `keyword`. Example:

```
qa, quad, l = 0.6, tilt = pi/4  ! Define QA.
qb: qa                          ! QB Inherits from QA.
qa[k1] = 0.12                   ! QB unaffected by modifications of QA.
```

In this example, element `QB` inherits the attributes of `QA` which, in this case, are the length and tilt parameters of `QA`. Once `QB` is defined, the elements are separate so modifications of the parameters of `QA` after `QB` is defined will not affect `QB`.

Bmad allows element names to be abbreviations of element types. For example, "`Q`", `QU`", and "`QUADRUPOL`" are all valid element names but "`QUADRUPOLE`", being an exact match to the corresponding element type, is not. Care must be used when using elements that are abbreviations of element types since *Bmad* allows element type names to be abbreviated (but element names may not be abbreviated when using inheritance). For example:

```
q1: quad          ! Q1 is a quadrupole since it comes before the next line.
quad: sextupole   ! QUAD element is defined to be a sextupole.
q2: quad          ! Q2 is a sextupole as it inherits from QUAD.
q3: qua           ! Q3 is a quadrupole. Inherit from names cannot be abbreviated.
```

3.5 Lattice Element Names

A valid element name may be up to 40 characters in length. The first character of the name must be a letter [A-Z]. After that, the rest of the name can contain only letters, digits [0-9], underscore "_", period ".", backslash "\", or a hash mark "#". A double hash mark "##" is not permitted since this interferes with the notation for finding the N^{th} element with a given name (§3.6). It is best to avoid these last three symbols since *Bmad* uses them to denote "relationships". Periods are used for tagging (§7.8), and backslash and hash marks are used for to compose names for superposition (§8) and multipass (§9) slave elements.

There is a short list of names that cannot be used as an element, line or list name. These reserved names are:

```
beam                            no_digested                 root
beginning                       no_superimpose              setenv
call                            parameter                   slice_lattice
calc_reference_orbit            parser_debug                start_branch_at
combine_consecutive_elements    particle_start              superimpose
debug_marker                    print                       title
end_file                        redef                       undefined_vars_evaluate_to_zero
expand_lattice                  remove_elements             use
merge_elements                  return                      use_local_lat_file
                                                            write_digested
```

Note: The one exception is if `end` is used to define a `marker`. This exception is allowed since *Bmad* uses the `end` name to define a marker in any case.

It is perfectly acceptable for multiple lattice elements to have the same name. Example:

```
q: quadrupole, ...
aline: line = (5*q)
use, aline
```

This will produce a lattice with five elements named "q". The exception is `group` (§4.26) and `overlay` (§4.40) controller elements will always have unique names. It is important to keep in mind that elements with the same name do not necessarily have the same parameter values. See Section §3.6 for an example.

3.6 Matching to Lattice Element Names

Where appropriate, for example when setting element attributes (§3.7), the wild card characters **"*"** and **"%"** can be used to select multiple elements. The **"*"** character will match any number of characters (including zero) while **"%"** maches to any single character. Additionally, matching can be restricted to a certain element class using the syntax:

```
class::element_name
```

where `class` is a class (EG: sextupole). For example:

```
m*              ! Match to all elements whose name begins with "m".
a%c             ! Match to "abc" but not to "ac" or "azzc".
quadrupole::*w  ! Match to all quadrupoles whose name ends in "w"
```

Note: The character **"%"** can be used in expressions used for setting lattice element parameter values. In this case, the **"%"** character represents the name of the lattice element whose parameter is being set. This is discussed in section §3.7.

Instead of matching to an element name, matches may be made to `type`, `alias`, or `descrip` attributes (§5.3) using the syntax:

```
attribute_name::string
```

where `attribute_name` is one of:

```
type
alias
descrip
```

and `string` is a string to match to which can include wild card characters **"*"** and **"%"**. If `string` contains blank characters, the string must be enclosed with single or double quotation marks. Note: An element attribute that is blank will never match. Also note that while `type`, `alias`, or `descrip` strings may have lower case characters (unlike element names, these strings are not converted to upper case), matching is always case insensitive. For Example:

```
type::"det bpm*"   ! Match to all elements whose type string starts with "det bpm".
alias::*           ! Match to all elements whose alias string is not blank.
```

After lattice expansion (§3.26), the general syntax to specify a set of elements is:

```
{branch_id>>}{class::}element_id
{branch_id>>}element_name##N                 ! N^th instance of element_name in
                                             !        branch with branch_id
{branch_id>>}{class::}element_id{##N}+M  ! M^th element after element
{branch_id>>}{class::}element_id{##N}-M  ! M^th element before element
```

where `{...}` marks an optional component, `class` is a class name, `branch_id` is a branch name or index (§2.2), `element_id` is and element name or element index (§7.2), and `##N` indicates that the N^{th} matching element in a branch is to be used. Examples:

```
x_br>>quad::q*        ! All quadrupoles of branch "x_br" whose name begins with "q".
2>>45                 ! element #45 of branch #2.
q01##3                ! The 3rd element named q01 in each branch where there are at
                      ! least three elements named q01.
q01-1                 ! Elements just before all elements named q01.
q01##3+0              ! Same as q01##3.
```

Note: `Group` and `overlay` elements have unique names so using `##` is unnecessary.

Note: When using the `##` construct, `super_lord` and `girder_lord` elements are considered to be situated where their slave elements are situated in the lattice. This is independent of where they actually exist which is in the lord part of branch 0 (§2.4).

Note: When using `+M` and `-M` offsets, no space should be put around the plus or minus signs. Offsets cannot be used with `overlay`, `group`, `ramper_lord` or `multipass_lord` elements. If used with a `super_lord` or `girder` element, for a `-M` offset the element selected is the M^{th} element before the first `slave` element of the lord and for a `+M` offset the element selected is the M^{th} element after the last `slave` element of the lord. Also, independent of the `geometry` of the branch, the element selected will "wrap" around the ends of the branch. For example, if a lattice branch looks like:

```
Index  Name
  0    Beginnning
  1    A
  2    B\1       ! First super_slave of B super_lord
  3    M
  4    B\2       ! Second super_slave of B super_lord
  5    C
```

then

```
Name  Translates to     Name  Translates to
----  -------------     ----  -------------
B-1   A                 B+1   C
B-2   Beginning         B+2   Beginning
B-3   C                 B+3   A
B-4   B\2               B+4   B\1
```

It is advised to avoid setting the parameters of differing elements that share the same name to differing values since this can lead problems later on. For example, consider this in a lattice file named, say, `lat.bmad`:

```
q1: quadrupole, ...
a: line = (..., q1, ...)
b: line = (..., q1, ...)
c: line = (a, b)  ! There are two q1 elements.
                  ! One from A-line and one from B-line.
use, c
expand_lattice    ! Expand the lattice
q1##1[k1] = 0.1   ! Set first element whose name is q1
q1##2[k1] = 0.2   ! Set second element whose name is q1
```

Now if later on someone wants to study just the `B` line that person could try to do this by creating a second file with just two lines:

```
call, file = lat.bmad
use, b
```

Normally this would work but in this case the lattice is invalid since there is only one `q1` element in line `B`. A more flexible solution would be to use unique names for the two `q1` elements.

Multiple elements in a lattice may share the same name. When multiple branches are present, to differentiate elements that appear in different branches, the "branch qualified" element name may be used. The branch qualified element name is of the form

```
branch_name>>element_name
```

where `branch_name` is the name of the branch and `element_name` is the "regular" name of the element. Example:

```
root>>q10w
x_branch>>crystal3
```

For `branch` lines (§2.2), the full "branch qualified" name of an element is of the form

```
branch_name>>element_name
```

where `branch_name` is the name of the branch and `element_name` is the "regular" name of the element. Example:

```
root>>q10w
xline>>cryst3
```

Using the full name is only needed to distinguish elements that have the same regular name in separate branches. When parsing a lattice file, branches are not formed until the lattice is expanded (§3.26). Therefore an `expand_lattice` statement is required before full names can be used in statements.

After lattice expansion (§3.26), when setting element attributes (§3.7, a comma delimited list of names can be used (the commas are actually optional). Each item in a list is either the name of an element or an element range. An element range has the syntax:

```
{branch>>}{class::}ele1:ele2
```

where

```
branch   ! Optional branch name or index.
class    ! Optional element class name ("quadrupole", "sbend", etc.)
ele1     ! Starting element of the range which includes ele1.
ele2     ! Ending element of the range which includes ele2.
```

For example:

```
3,15:17          ! Elements with index 3, 15, 16, and 17 in branch 0.
3 15:17          ! Same as above (commas are optional).
2>>45:51         ! Elements 45 through 51 of branch 2.
q1:q5            ! Elements between q1 and q5.
sbend::q1:q5     ! All sbend elements between q1 and q5 including q1 and q5 if they.
                 ! are sbends. Notice that q1 and q5 do not have to be sbends
```

if the element index of `ele1` is greater than `ele2` then the range wraps around the end of the lattice. For example, if branch0 has 360 tracking elements (§2.4), the range `321:72` is equivalent to `372:360, 0:72`.

With a comma delimited list of names, a tilde prefix can be used to remove elements from the list. Adding and subtracting elements is done left to right. For example:

```
*::*, ~quadrupole::*  ! All elements except quadrupoles
b*, ~b1*, b13         ! All elements with names beginning with B except
                      !   elements with names beginning with B1. However,
                      !   elements named B13 are retained.
```

With a list of names, an ampersand "`&`" can be used to form the intersection of two sets.

```
100:200 & sbend::*    ! All sbend elements whose index is between 100 and 200.
q* & quad::*          ! Equivalent to "quadrupole::q*".
q* & quad::* b1       ! Equivalent to "quadrupole::q*, b1".
```

When lattice expansion occurs during parsing of a lattice (§3.25), all `rbend` elements are converted to `sbend` elements (§4.5) but there is a `sub_key` element parameter that is used so *Bmad* knows which bend elements were defined as `rbend`s in the lattice file. After lattice expansion, the string `sbend::*` will match to all bend elements irregardless of whether a bend was defined to be a `sbend` or a `rbend`. On the other hand, after lattice expansion the string `rbend::*` will match to all bend elements that were defined as `rbend`s before expansion.

3.7 Lattice Element Parameters

Any lattice element has various attributes like its name, its length, its strength, etc. The values of element attributes can be specified when the element is defined. For example:

```
b01w: sbend, l = 6.0, rho = 89.0 ! Define an element with attributes.
```

After an element's definition, most attributes may be referred to using the syntax

```
class::element_name[attribute_name]
```

Examples:

```
q01[k1]                      ! K1 attribute of element Q01.
sbend::b0%[dg]               ! DG attribute of sbend elements whose name
                             !   has three characters starting with "B0"
```

Some element parameters have a more complicated syntax and are listed in section §3.8.

Element attributes can be set or used in an algebraic expression:

```
b01w[roll] = 6.5                 ! Set an attribute value.
b01w[L] = 6.5                    ! Change an attribute value.
b01w[L] = b01w[rho] / 12         ! OK to reset an attribute value.
my_const = b01w[rho] / b01w[L]   ! Use of attribute values in an expression.
```

Notice that there can be no space between the element name and the `[` opening bracket.

Chapter Chapter 4 lists the attributes appropriate for each element class.

When setting an attribute value, if more than one element has the `element_name` then *all* such elements will be set. When setting an attribute value, if `element_name` is the name of a type of element, all elements of that type will be set. For example

```
q_arc[k1] = 0.234                    ! Set all elements named Q_ARC.
rfcavity::*[voltage] = 3.7           ! Set all RFcavity elements.
```

To set an attribute for multiple element at one time, The wild cards "`*`", and "`%`" can be used in element names (§3.6). Examples:

```
*[tracking_method] = bmad_standard  ! Matches all elements.
quadrupole::Q*[k1] = 0.234    ! Matches all quadrupoles with names beginning with Q.
Q%1[k1] = 0.234               ! Matches to "Q01" but not "Q001".
```

Unlike when there are no wild cards used in a name, it is not an error if a name with wild cards does not match to any element. Note: A name with wild cards will never match to the `BEGINNING` element (§7.7).

The "`%`" symbol can be used in expression to represent the lattice element whose parameter is being set. For example:

```
s%z[k2] = %[k2] + 0.03 * ran_gauss()
```

The "`s%z`" on the left hand side matches to all three letter elements whose name begins with "`s`" and ends with "`z`". For each element that is matched, the "`%[k2]`" in the expression on the right hand side will be the `k2` value of that element. Thus, if there are three elements, named `s0z`, `saz`, and `s.z`, in the lattice that match `s%z`, the above command is equivalent to the following three commands

```
s0z[k2] = s0z[k2] + 0.03 * ran_gauss()
saz[k2] = saz[k2] + 0.03 * ran_gauss()
s.z[k2] = s.z[k2] + 0.03 * ran_gauss()
```

After lattice expansion (§3.26), the attributes of specific elements may be set using the syntax as discussed in Section §3.6. Example:

```
expand_lattice              ! Expand the lattice.
97[x_offset] = 0.0023       ! Set x_offset attribute of 97th element
b2>>si_cryst##2[tilt] = 0.1 ! Tilt the 2nd instance of "si_cryst" in branch "b2"
5:32[x_limit] = 0.3         ! Sets elements with indexes 5 through 32 in branch 0.
```

3.8 Nonstandard Parameter Syntax

This section lists parameters that have a "nonstandard" syntax (not in the form "`ename[parameter]`"). In the following, "`ename`" is an element name, and "`N`", "`N1`", "`N2`", "`N3`" and "`M`" along with `<out>`, `<n1>`, and `<n2>` are integers.

AC_kicker (§4.2):

```
ename[amp_vs_time(N)%time]
ename[amp_vs_time(N)%amp]
ename[frequencies(N)%freq]
ename[frequencies(N)%amp]
ename[frequencies(N)%phi]
```

Cartesian map (§5.16.2):

```
ename[cartesian_map(N)%field_scale]
ename[cartesian_map(N)%r0(1)]
ename[cartesian_map(N)%r0(2)]
ename[cartesian_map(N)%r0(3)]
ename[cartesian_map(N)%master_parameter]
ename[cartesian_map(N)%t(M)%A]                  -- M^th term in N^th map.
ename[cartesian_map(N)%t(M)%kx]
ename[cartesian_map(N)%t(M)%ky]
ename[cartesian_map(N)%t(M)%kz]
ename[cartesian_map(N)%t(M)%x0]
ename[cartesian_map(N)%t(M)%y0]
ename[cartesian_map(N)%t(M)%phi_z]
```

Controller knot points (§5.4):

```
ename[x_knot(N)]                        -- N^th x_knot point.
ename[slave(M)%y_knot(N)]               -- N^th y_knot point for M^th slave.
```

Custom attributes (§3.9):

```
ename[r_custom(N1, N2, N3)]
ename[r_custom(N1, N2)]                 -- Equivalent to ename[r_custom(N1, N2,  0)]
ename[r_custom(N1)]                     -- Equivalent to ename[r_custom(N1,  0,  0)]
```

Cylindrical map (§5.16.3):

```
ename[cylindrical_map(N)%phi0_fieldmap]
ename[cylindrical_map(N)%theta0_azimuth]
ename[cylindrical_map(N)%field_scale]
ename[cylindrical_map(N)%dz]
ename[cylindrical_map(N)%r0(1)]
ename[cylindrical_map(N)%r0(2)]
ename[cylindrical_map(N)%r0(3)]
ename[cylindrical_map(N)%master_parameter]
```

Gen_gradients map (§5.16.5):

```
ename[gen_gradients(N)%field_scale]
ename[gen_gradients(N)%dz]
ename[gen_gradients(N)%g_ref]
ename[gen_gradients(N)%r0(1)]
ename[gen_gradients(N)%r0(2)]
ename[gen_gradients(N)%r0(3)]
ename[gen_gradients(N)%master_parameter]
```

Grid_field (§5.16.4):

```
ename[grid_field(N)%phi0_fieldmap]
ename[grid_field(N)%interpolation_order]
ename[grid_field(N)%harmonic]
ename[grid_field(N)%geometry]
ename[grid_field(N)%ename_anchor_pt]
ename[grid_field(N)%phi0_fieldmap]
ename[grid_field(N)%field_scale]
ename[grid_field(N)%dr(1)]
ename[grid_field(N)%dr(2)]
ename[grid_field(N)%dr(3)]
ename[grid_field(N)%r0(1)]
ename[grid_field(N)%r0(2)]
ename[grid_field(N)%r0(3)]
ename[grid_field(N)%master_parameter]
```

Long range wake (§5.20.4):

```
ename[lr_wake%amp_scale]
ename[lr_wake%time_scale]
ename[lr_wake%freq_spread
ename[lr_wake%mode(N)%freq_in]
ename[lr_wake%mode(N)%freq]
ename[lr_wake%mode(N)%r_over_q]
ename[lr_wake%mode(N)%damp]
ename[lr_wake%mode(N)%phi]
ename[lr_wake%mode(N)%polar_angle]
ename[lr_wake%mode(N)%polarized]
```

Surface curvature (§5.11):

```
ename[curvature%spherical]
ename[curvature%elliptical_x]
ename[curvature%elliptical_y]
ename[curvature%elliptical_z]
ename[curvature%x(N1)y(N2)]
```

Taylor terms (§4.52):

```
ename[tt<out><n1><n2>...]      ! Orbital terms. EG: rot[tt13] -> M13 matrix term
ename[ttS0<n1><n2><n3>...]     ! S0 spin quaternion terms.
ename[ttSx<n1><n2><n3>...]     ! Sx spin quaternion terms.
ename[ttSy<n1><n2><n3>...]     ! Sy spin quaternion terms.
ename[ttSz<n1><n2><n3>...]     ! Sz spin quaternion terms.
```

Wall for vacuum chambers and masks (§5.12):

```
ename[wall%section(N)%s]
ename[wall%section(N)%wall%dr_ds]
ename[wall%section(N)%v(M)%x]
ename[wall%section(N)%v(M)%y]
ename[wall%section(N)%v(M)%radius_x]
ename[wall%section(N)%v(M)%radius_y]
ename[wall%section(N)%v(M)%tilt]
```

3.9 Custom Element Attributes

Real scalar and vector custom element attributes may be defined for any class of element and real scalar parameters can be defined for the lattice as a whole. Custom element attributes are useful with programs

that need to associate "extra" information with particular lattice elements or the lattice itself and it is desired that this extra information be settable from within a lattice file. For example, a program might need an error tolerance for the strength of quadrupoles.

Adding custom attributes will not disrupt programs that are not designed to use the custom attributes. Currently, up to 40 named custom attributes may be defined for any given element class. The syntax for defining custom attributes is:

```
parameter[custom_attributeN] = {class_name::}attribute_name
```

Where "N" is an integer between 1 and 40 and "attribute_name" is the name of the attribute. To restrict the custom attribute to a particular element class, the element class can be prefixed to the attribute name. To define a global parameter for the lattice, use parameter" as the class name. Examples:

```
parameter[custom_attribute1] = quadrupole::error_k1
parameter[custom_attribute1] = mag_id
parameter[custom_attribute1] = sextupole::error_k2
parameter[custom_attribute2] = color
parameter[custom_attribute2] = parameter::quad_mag_moment
```

The first line in the example assigns, for the first custom attribute group (custom_attribute1), a name of error_k1 to all quadrupoles. The second line in the example assigns to the first custom attribute group the name mag_id to all element classes except quadrupoles since that class of element already has an assigned name. The third line assigns, for the first custom attribute group, a name of error_k2 to all sextupoles overriding the mad_id name. The fourth line in the above example assigns, for the second custom attribute group, a name of color to all element classes. Finally, the last line defines a global parameter called quad_mag_moment.

Once a custom attribute has been defined it may be set for an element of the correct class. Example:

```
parameter[custom_attribute2] = lcavity::rms_phase_err
parameter[custom_attribute3] = parameter::cost
...
parameter[cost] = 140000000
l2a: lcavity, rms_phase_err = 0.0034, ...
```

Notice that defining the name for a custom attribute must come before its use.

Custom attributes that are assigned to an individual element class, like error_k1 above, are called "class-specific" attributes. Custom attributes, like mag_id above, that are assigned to all element classes, are called "common" attributes. For a given custom attribute group, The setting of a class-specific attribute will take precedence over the setting of a common attribute. Thus, in the above example, the fact that quadrupole::error_k1 comes before mag_id and sextupole::error_k2 appears after does not affect anything. Once a common attribute is defined for a given custom attribute group, it cannot be changed. Similarly, once a class-specific attribute is defined for a given class for a given custom attribute group it cannot be changed. Trying to redefine a given custom attribute using a new name that is the same as the old name is not considered an error. For example, the following is OK:

```
parameter[custom_attribute2] = color
parameter[custom_attribute2] = color   ! OK since the same name is used.
```

Custom attributes are global in a program and not lattice-specific. That is, if a program reads in two different lattices the custom attribute settings of both lattices will be combined.

For someone creating a program, section §32.24 describes how to make the appropriate associations.

Note: If custom string information needs to be associated with an element, the type, alias and descrip element components (§5.3) are available.

Besides the named custom attributes described above, there is a three dimensional vector, called r_custom, associated with each element that can be used to store numbers. For example:

```
qq: quadrupole, r_custom(-2,1,5) = 34.5, r_custom(-3) = 77.9
```

Negative indices are accepted and if only one or two indices are present, the others are assumed to be zero. Thus `r_custom(-3)` is equivalent to `r_custom(-3,0,0)`.

Note: When there is a superposition (§8), the `super_slave` elements that are formed do *not* have any custom attributes assigned to them even when their `super_lord` elements have custom attributes. This is done since the *Bmad* bookkeeping routines are not able to handle the situation where a `super_slave` element has multiple `super_lord` elements and thus the custom attributes from the different `super_lord` elements have to be combined. Proper handling of this situation is left to any custom code that a program implements to handle custom attributes.

3.10 Parameter Types

There are five types of parameters in *Bmad*: reals, integers, switches, logicals (booleans), and strings. Acceptable logical values are

```
true    false
t       f
```

For example

```
rf1[is_on] = False
```

String literals can be quoted using double quotes (") or single quotes (’). If there are no blanks or commas within a string, the quotes can be omitted. For example:

```
Q00W: Quad, type = "My Type", alias = Who_knows, &
                                  descrip = "Only the shadow knows"
```

Unlike most everything else, strings are not converted to uppercase.

Switches are parameters that take discrete values. For example:

```
parameter[particle] = positron
q01w: quad, tracking_method = bmad_standard
```

The name “switch” can refer to the parameter (for example, `tracking_method`) or to a value that it can take (for example, `bmad_standard`). The name “method” is used interchangeably with switch.

3.11 Particle Species Names

For the purpose of assigning names to simulated particles, particles are divided into four groups. One group are are “fundamental particles”. These are:

```
electron,  positron
muon,      antimuon
proton,    antiproton
neutron    anti_neutron
deuteron   anti_deuteron
pion+,     pion0,      pion-
helion     anti_helion        ! #3He
photon
```

For historical reasons, names for the fundamental particles are *not* case sensitive.

Another group are atoms. The general syntax for atoms is:

```
{#nnn}AA{ccc}
```

The curly brackets {...} denote optional prefixes and suffixes. `AA` here is the atomic symbol, `#nnn` is the number of nucleons, and `ccc` is the charge. Examples:

```
parameter[particle] = #12C+3       ! Triply charged carbon-12.
parameter[p0c] = 12 * 500e6         ! Reference momentum is total momentum for particle.
parameter[particle] = He--          ! Doubly charged He.
```

If the number of nucleons is given, the appropriate weight for that isotope is used. If the number of nucleons is not present, the mass is an average weighted by the isotopic abundances of the element. The charge may be given by using the appropriate number of plus (+) or minus (-) signs or by using a plus or minus sign followed by a number. Thus "`---`" is equivalent to "-3". Names here are case sensitive. "@M" must be used and not "@m" for specifying the mass. The mass of an atom is adjusted by the number of electrons relative to neutral. That is

$$m_{\text{atom}} = m_{\text{neutral atom}} - C \cdot m_{\text{electron}} \tag{3.1}$$

where C is the charge in units number of electrons relative to neutral. No adjustment is made for mass shifts due to finite electron binding energies. This shift is small typically being well less than 1% of the mass of the electron.

Anti-atoms made with antimatter have names using the prefix "`anti`". For example, a bare gold anti-atom nucleus would be designated `antiAu-79`.

Note: When setting the reference momentum `parameter[p0c]`, or reference total energy `parameter[E_tot]`, the total for the whole particle is used. *Not* the value per nucleon.

Another group of particles are the "known" molecules. The syntax for these are:

```
BBB{@Mxxxx}{ccc}
```

`@Mxxxx` is the mass in AMU, `ccc` is the charge, and `BBB` is the molecular formula. The mass may to specified to hundredths of an AMU. The known molecules are:

```
CO      CO2
D2      D2O
OH      O2
H2      H2O     HF
N2      NH2     NH3
CH2     CH3     CH4
C2H3    C2H4    C2H5
```

Like with atoms, if the mass is not specified, the average isotopic mass is used. Examples:

```
C2H3@M28.4+     ! Singly charged C2H3 with mass of 28.4
CH2             ! Neutral CH2
```

Like the atomic formulas, molecular formulas are case sensitive. Like atoms, the mass of a known molecule is adjusted by the number of electrons relative to neutral.

The last group of particle are particles where only the mass and charge are specified. The syntax for these are:

```
@Mxxxx{ccc}
```

Example:

```
@M37.54++    ! Doubly charged molecule of mass 37.54 AMU.
```

Note: When setting the value of a variable to be a particle species ID, use the `species` function as discussed in §3.14.

3.12 Units and Constants

Bmad uses SI (Système International) units as shown in Table 3.1. Note that *MAD* uses different units. For example, *MAD*'s unit of Particle Energy is GeV not eV.

Note: For compatibility with *MAD*, the `beam, energy = xxx` command (§10.3) uses GeV and the `emass` and `pmass` constants (see below this section) also use GeV. It is recommended that the use of these constructs be avoided.

Quantity	*Units*
Angles	radians
Betatron Phase	radians
Current	Amps
Frequency	Hz
Kick	radians
Length	meters
Magnetic Field	Tesla
Particle Energy	eV
RF Phase Angles	radians/2π
Voltage	Volts

Table 3.1: Physical units used by *Bmad*.

Bmad defines commonly used physical and mathematical constants shown in Table 3.2. All symbols use straight SI units except for `emass` and `pmass` which are provided for compatibility with *MAD* and *should be avoided*.

As an alternative, the `mass_of`, and `anomalous_moment_of` functions (§3.14) may be used in place of the defined constants for mass and anomalous magnetic moment.

Note: The standard definition of the magnetic moment g-factor for spin 1/2 fundamental particles is

$$\mu = g \frac{q}{2\,m} \mathbf{S} \tag{3.2}$$

where μ is the magnetic moment, q is the particle charge, and m is the mass. The anomalous moment a is then defined to be

$$a = \frac{g-2}{2} \tag{3.3}$$

For nuclei and other composite baryonic particles, it is conventional to define the g-factor using

$$\mu = g \frac{e}{2\,m_p} \mathbf{S} \tag{3.4}$$

where m_p is the mass of the proton. This is inconvenient for calculations since an equation like Eq. (24.2) would not work for all particles. To get around this, the g-factors used by Bmad are always derived from Eq. (3.2) (think of this as an "effective" g-factor).

Symbol	*Value*	*Units*	*Name*
pi	3.141592653589793		
twopi	2 * pi		
fourpi	4 * pi		
e	2.718281828459045		
e_log	2.718281828459045		
sqrt_2	1.414213562373095		
degrad	180 / pi		From rad to deg
degrees	pi / 180		From deg to rad
raddeg	pi / 180		From deg to rad
anom_moment_deuteron	-0.1425617662		Deuteron anomalous magnetic moment*
anom_moment_electron	0.00115965218128		Electron anomalous magnetic moment
anom_moment_muon	0.00116592089		muon anomalous magnetic moment
anom_moment_proton	1.792854734463		proton anomalous magnetic moment
anom_moment_he3	-4.184153686		He^3 anomalous magnetic moment*
fine_struct_const	0.0072973525693		Fine structure constant
m_deuteron	$1.87561294257 \cdot 10^9$	eV	Deuteron mass
m_electron	$0.51099895000 \cdot 10^6$	eV	Electron mass
m_neutron	$0.93956542052 \cdot 10^9$	eV	Neutron mass
m_muon	$105.6583755 \cdot 10^6$	eV	Muon mass
m_pion_0	$134.9766 \cdot 10^6$	eV	π^0 mass
m_pion_charged	$139.57018 \cdot 10^6$	eV	π^+, π^- mass
m_proton	$0.93827208816d \cdot 10^9$	eV	Proton mass
c_light	$2.99792458 \cdot 10^8$	m/sec	Speed of light
r_e	$2.8179403262 \cdot 10^{-15}$	m	Electron radius
r_p	$1.5346982647 \cdot 10^{-18}$	m	Proton radius
e_charge	$1.602176634 \cdot 10^{-19}$	Coul	Electron charge
h_planck	$4.135667696 \cdot 10^{-15}$	eV*sec	Planck's constant
h_bar_planck	$6.582118990 \cdot 10^{-16}$	eV*sec	Planck / 2π
emass	$0.51099895000 \cdot 10^{-3}$	GeV	Electron mass (please avoid using)
pmass	0.93827208816	GeV	Proton mass (please avoid using)

* Effective anomalous moments. See the discussion after Eq. (3.2).

Table 3.2: Physical and mathematical constants recognized by *Bmad*.

3.13 Arithmetic Expressions

Arithmetic expressions can be used in a place where a real value is required. The standard operators are defined:

$a + b$	Addition
$a - b$	Subtraction
$a * b$	Multiplication
a / b	Division
$a \wedge b$	Exponentiation

Bmad also has a set of intrinsic functions. A list of these is given in §3.14.

Literal constants can be entered with or without a decimal point. An exponent is marked with the letter E. For example

```
1, 10.35, 5E3, 314.159E-2
```

Symbolic constants can be defined using the syntax

```
constant_name = expression
```

Alternatively, to be compatible with *MAD*, using ":=" instead of "=" is accepted

```
constant_name := expression
```

Examples:

```
my_const = sqrt(10.3) * pi^3
abc      := my_const * 23
```

Unlike *MAD*, *Bmad* uses immediate substitution so that all constants in an expression must have been previously defined. For example, the following is *not* valid:

```
abc      = my_const * 23      ! No: my_const needs to be defined first.
my_const = sqrt(10.3) * pi^3
```

here the value of `my_const` is not known when the line "`abc` = ..." is parsed. Note: To get the effect of delayed evaluation, use `overlay` (§4.40) or `group` (§4.26) controller elements.

Once defined, symbolic constants cannot be redefined. For example:

```
my_const = 1
my_const = 2  ! No! my_const cannot be redefined.
```

The restriction against redefining constants was implemented to avoid hard to find problems. That is, the value of an expression that uses a redefined constant will depend upon if the expression is evaluated before or after the redefinition. On very rare occasions, it is convenient to be able to redefine constants so if the redefinition has a `redef:` prefix, a constant can be redefined

```
my_const = 1
redef: my_const = 2  ! OK!
```

It is advised not to use `redef` unless there a very good reason for its use.

`group` (§4.26) and `overlay` (§4.40) controller elements are an exception to the immediate evaluation rule. Since controller elements may control elements that do not exist until `lattice expansion` (§3.26), the arithmetic expressions associated with controller elements are not evaluated until lattice expansion. Example:

```
s_20W: sextupole, l = 0.27
sk: overlay = {s_20W[a1]:-2*s_20W[L]}, var = {k1}, k1 = 0.2
s_20W[L] = 0.34
s_30E: s_20W
...
expand_lattice
```

Here the expression of overlay `sk` is evaluated, when the lattice is expanded, to be `-0.68 = -2*0.34`. This uses the length of element `s_20W` at the point when the lattice is expanded and not at the point when `sk` was defined. Additionally, the element `s_30E`, which inherits the attributes of `s_20W`, inherits a value of zero for `a1` (skew multipole moment) since inheritance uses immediate evaluation just like the setting of constants.

Element attributes can be used after they have been defined but not before. Example:

```
sa: sextupole, l = 0.3, k2 = 0.01 * sa[L]  ! Good
sb: sextupole, k2 = 0.01 * sb[L], l = 0.3  ! BAD SET OF K2. L IS DEFINED AFTER.
```

In this example, the `k2` attribute of element `sa` is correctly set since `k2` is defined after `l`. On the other hand, `k2` of element `sb` will have a value of zero since `l` of `sb` defaults to zero before it is set.

One potential pitfall with immediate substitution is that when an element attribute changes, it does not affect prior evaluations. Example:

```
s1: sextupole, k2 = 2.3
aa = s1[k2]              ! aa = 2.3
s1[k2] = 1.7             ! value of aa does not change
```

Here the value of constant `aa` will remain fixed at 2.3 no matter how the value of `s1[k2]` is altered after `aa` is defined.

Another potential pitfall is when using dependent element attributes (§5.1). For example:

```
b01w: sbend, l = 0.5, angle = 0.02
a_const = b01w[g]    ! No: bend g has not yet been computed!
```

Here the bend strength `g` (§4.5) will eventually be computed to be 0.04 (= angle / l) but that computation does not happen until lattice expansion (§3.26). In this case, the value of `a_const` will be the default value of `g` which is zero. As a rule of thumb, never rely on dependent attributes having their correct value.

3.14 Intrinsic functions

The following intrinsic functions are recognized by *Bmad*:

`sqrt`(x)	Square Root
`log`(x)	Logarithm
`exp`(x)	Exponential
`sin`(x), `cos`(x)	Sine and cosine
`tan`(x), `cot`(x)	Tangent and cotangent
`sinc`(x)	Sin(x)/x Function
`asin`(x), `acos`(x)	Arc sine and Arc cosine
`atan`(x)	Arc tangent
`atan2`(y, x)	Arc tangent of y/x
`sinh`(x), `cosh`(x)	Hyperbolic sine and cosine
`tanh`(x), `coth`(x)	Hyperbolic tangent and cotangent
`asinh`(x), `acosh`(x)	Hyperbolic arc sine and Arc cosine
`atanh`(x), `acoth`(x)	Hyperbolic arc tangent and cotangent
`abs`(x)	Absolute Value
`factorial`(n)	Factorial
`ran`()	Random number between 0 and 1
`ran_gauss`()	Gaussian distributed random number
`ran_gauss`(sig_cut)	Gaussian distributed random number
`int`(x)	Nearest integer with magnitude less then x
`nint`(x)	Nearest integer to x
`sign`(x)	1 if x positive, -1 if negative, 0 if zero
`floor`(x)	Nearest integer less than x
`ceiling`(x)	Nearest integer greater than x
`modulo`(a, p)	a - floor(a/p) * p. Will be in range [0, p].
`mass_of`(A)	Mass of particle A
`charge_of`(A)	Charge, in units of the elementary charge, of particle A
`anomalous_moment_of`(A)	Anomalous magnetic moment of particle A
`species`(A)	Species ID of A

`ran_gauss` is a Gaussian distributed random number with unit RMS. Both `ran` and `ran_gauss` use a seeded random number generator. To choose the seed set

```
parameter[ran_seed] = <Integer>
```

A `value` of zero will set the seed using the system clock so that different sequences of random numbers will be generated each time a program is run. The default behavior if `parameter[ran_seed]` is not present is to use the system clock for the seed.

The `ran_gauss(cut)` function with an argument truncates the distribution so that no values are returned with an absolute value greater than `cut`. If `cut` is non-positive, it is ignored so that, for example, `ran_gauss(-1)` is equivalent to `ran_gauss()`.

If an element is used multiple times in a lattice, and if `ran` or `ran_gauss` is used to set an attribute value of this element, then to have all instances of the element have different attribute values the setting of the attribute must be after the lattice has been expanded (§3.26). For example:

```
a: quad, ...
a[x_offset] = 0.001*ran_gauss()
my_line: line = (a, a)
use, my_line
```

Here, because *Bmad* does immediate evaluation, the `x_offset` values for `a` gets set in line 2 and so both copies of `a` in the lattice get the same value. This is probably not what is wanted. On the other hand if the attribute is set after lattice expansion:

```
a: quad, ...
my_line: line = (a, a)
use, my_line
expand_lattice
a[x_offset] = 0.001*ran_gauss()
```

Here the two `a` elements in the lattice get different values for `x_offset`.

The following functions take a species ID as the argument:

```
mass_of(A),               charge_of(A)
anomalous_moment_of(A),   species(A)
```

See §3.11 for the syntax of naming particles.

The `mass_of`, `charge_of`, and `anomalous_moment_of` functions give the mass of, charge of (in units of the elementary charge), and anomalous moment of, a particle. Example:

```
parameter[particle] = deuteron
am = anomalous_moment_of(parameter[particle])^2
my_particle = species(He++)      ! my_particle now represents He++
chg1 = charge_of(my_particle)    ! chg = charge of He++
chg2 = charge_of(He++)           ! Same as previous line
chg3 = charge_of(species(He++))  ! Same as previous line
```

The `species` function is needed in the definition of `my_particle` in the above example so that *Bmad* knows that the string "He++" represents a type of particle. Inside functions like `mass_of`, the use of `species` is optional since, in this case, *Bmad* can correctly parse the argument.

The value returned by the `mass_of` function accounts for the ionization state of a particle in that there is a correction for the change in the number of electrons a particle has. Thus the values of `mass_of#3He` will be heavier than `mass_of#3He++` by two electron masses. This correction only involves multiples of the electron mass and variations in the particle mass due to electron binding energies is not accounted for. These binding energy corrections are generally very small. If the isotopic state is not specified for an atom, the average weighted by the natural abundance is used.

3.15 Statement Order

With some exceptions, statements in a lattice file can be in any order. For example, the lines (§7.2) specified in a `use` statement (§7.7) can come after the `use` statement. And `group` (§4.26) and `overlay` (§4.40) controller elements may be defined before the slave elements whose parameters they control are defined.

The exceptions to this rule are:

- If there is an `expand_lattice` statement (§3.25), everything necessary for lattice expansion must come before. In particular, all `lines` (§7.2), `lists` (§7.6), and `use` (§7.7) statements necessary for lattice expansion must come before.
- Immediate evaluation of arithmetic expressions (§3.13) mandates that values be defined before use.
- A lattice element must be defined before any of its parameters are set. Example:
  ```
  pp[z_offset] = 0.1     ! WRONG! PP HAS NOT BEEN DEFINED YET!
  pp: patch              ! Here PP is defined
  ```
  In this example, the `z_offset` of the element `pp` is set before `pp` has been defined. This is an error. As a corollary to this rule, element parameters that are set using wild card characters will only affect those parameters that have been already defined. For example:
  ```
  crystal::*[b_param] = 0.2
  c5: crystal
  ```
  In this example, the `b_param` of all `crystal` elements is set to `0.2` *except* for `c5` and all other crystal elements that are defined after the set.

3.16 Setenv Statement

The `setenv` statement is used to set an environment variable from within a lattice file. This is useful to easily switch what directory is searched for `called` (§3.19) files. The syntax is:

```
setenv <env_name> = <env_value>
```

where `<env_name>` is the name of the environment variable and `<env_value>` is the value. Example:

```
setenv SUB_DIR = sub/version2
call, file = $SUB_DIR/lat2.bmad
```

In this example, the name of the called file expands to `sub/version2/lat2.bmad`.

An environment variable `LATTICE_ROOT_DIR` is automatically set to the directory that contains the top-level lattice file. If multiple lattices are parsed by a program (for example, with multiple universes in *Tao*), the value of `LATTICE_ROOT_DIR` is reset each time a top-level lattice file is parsed. However, environment variables that have defined will be retained even when a new top-level lattice file is parsed.

Environment variables that are set are local to the execution of the program that calls the lattice parser and do not persist after the program terminates.

3.17 Print Statement

The `print` statement prints a message at the terminal when the lattice file is parsed by a program. Syntax:

```
print <string>
```

Where `<string>` is the string to be printed. Variable values can be printed by using **back-tick** characters. For example:

```
print Remember! Q01 quad strength of `q01[k1]` not yet optimized!
print Optimization is as easy as 2 + 2 = `2+2`.
```

will result in the following being printed:

```
Message in Lattice File: Remember! Q01 quad strength of 0.4526 not yet optimized!
Message in Lattice File: Optimization is as easy as 2 + 2 = 4.
```

The `print` statement is useful to remind someone using the lattice of important details.

3.18 Title Statement

The `title` statement sets a title string which can be used by a program. For consistency with *MAD* there are two possible syntaxes

```
title, <String>
```

or the statement can be split into two lines

```
title
<String>
```

For example

```
title
"This is a title"
```

3.19 Call Statement

It is frequently convenient to separate the lattice definition into several files. Typically there might be a file (or files) that define the layout of the lattice (something that doesn't change often) and a file (or files) that define magnet strengths (something that changes more often). The `call` is used to read in separated lattice files. The syntax is

```
call, filename = <file-name>   ! or
call, no_abort_on_open_error, file = <file-name>
```

If the `no_abort_on_open_error` option is present, a missing file will not abort the parsing. Examples:

```
call, filename = "../layout/my_layout.bmad"      ! Relative pathname
call, filename = "/nfs/cesr/lat/my_layout.bmad"  ! Absolute pathname
call, filename = "$LATDIR/my_layout.bmad"        ! Absolute pathname
```

Environment variables in the file name will be expanded. *Bmad* will read the called file until a `return` or `end_file` statement is encountered or the end of the file is reached.

For filenames that have a relative pathname, the called file will be searched for relative to the directory of the calling file. Thus, in the above example, if the file containing the call statements is in the directory `/path/to/lat_dir`, the first call will open the file:

```
/path/to/lat_dir/../layout/my_layout.bmad
```

To call a file relative to the current working directory, use the environment variable `PWD`. Example:

```
call, filename = $PWD/here.bmad
```

An environment variable `LATTICE_ROOT_DIR` is automatically set to the directory that contains the head lattice file.

Where a called file is searched for may be modified by using a `use_local_lat_file` statement. See Section §3.22 for more details. Additionally, the `setenv` command (§3.16) may be used to define environment variable on the fly.

3.20 Inline Call

Any lattice elements will have a set of attributes that need to be defined. As a convenience, it is possible to segregate an element attribute or attributes into a separate file and then "call" this file using an "inline call". The inline call has three forms:

```
<ele_name>: <ele_type>, ..., call::<file_name>, ...                        ! or
<ele_name>: <ele_type>, ..., <attribute_name> = call::<file_name>, ...     ! or
<ele_name>[<attribute_name>] = call::<file_name>
```

where <attribute_name> is the name of the attribute and <file_name> is the name of the where the attribute structure is given. The Environment variables in the file name will be expanded. Example:

```
c: crystal, call::$AB/my_curvature.bmad, h_misalign = call::my_surface.bmad, ...
```

For grid_fields, which can take some time to parse, a HDF5 binary file can to created and then the HDF5 file, which must have a .h5 or .hdf5 suffix, can be read in with an inline call Example:

```
qq: quadrupole, grid_field = call::my_grid.h5, ...
```

To create hdf5 files, first create a lattice with the grid_field defined with plain text. Next read the lattice into any program that can create *Bmad* lattice files (for example, the *Tao* program (§1.2) can do this) and have the program then generate a lattice file.

3.21 Undefined_vars_evaluate_to_zero Statement

Normally, using an undefined variable in an expression is treated as an error. This helps prevent spelling bugs. This is in contrast to *MAD* where undefined variables are treated as zero. Occasionally, it is convenient to follow the *MAD* convention for lattice files translated from *MAD*. To do this, put at the top of the lattice file the command:

```
undefined_vars_evaluate_to_zero
```

Note that use of this command is discouraged. The superior fix is just to define the variables to be zero before any use.

3.22 Use_local_lat_file Statement

It is sometimes convenient to override where *Bmad* looks for called files (see §3.19). For example, suppose it is desired to temporarily override the settings in a called file without modifying the called file itself. In this case, the use_local_lat_file statement can be used. When this statement is encountered in a lattice file, the local directory (that is, the directory from which the program is run) is searched first for the called file and if a file of the correct name is found, that file is used.

An example will make this clear. Suppose lattice file /A/lat.bmad contains the call:

```
call, filename = "/B/sub.bmad"
```

Now suppose that you want to use lat.bmad with a modified sub.bmad but you do not want to modify /A/lat.bmad or /B/sub.bmad. The solution is to create two new files. One file, call it new.bmad, which can be situated in any directory, has two lines in it:

```
use_local_lat_file
call, filename = "/A/lat.bmad"
```

The second new file is the modified sub.bmad and it must be in the directory from which the program is run.

3.23 No_Superimpose Statement

In certain cases it is useful to turn off superposition (§8). The no_superposition statement will do this. To turn off all superpositioning, this statement can appear anywhere as long as it is before any expand_lattice (§3.25) statement. If the lattice has an expand_lattice statement, and the no_superposition statement appears after, the no_superposition statement will only block superpositions that are defined after the no_superposition statement.

3.24 Return and End_File Statements

`Return` and `end_file` have identical effect and tell *Bmad* to ignore anything beyond the `return` or `end_file` statement in the file.

3.25 Expand_Lattice Statement

Normally, lattice expansion happens automatically at the end of the parsing of the lattice file but an explicit `expand_lattice` statement in a lattice file will cause immediate expansion. See §3.26 for details. Subsequent `expand_lattice` statements after the first one are ignored and have no effect on the lattice.

3.26 Lattice Expansion

At some point in parsing a lattice file, the ordered sequence (or sequences if there are multiple branches) of elements that form a lattice must be constructed. This process is called `lattice expansion` since the element sequence can be built up from sub–sequences (§7). Normally, lattice expansion happens automatically at the end of the parsing of the lattice file (or files) but an explicit `expand_lattice` statement in a lattice file will cause immediate expansion. The reason why lattice expansion may be necessary before the end of the file is due to the fact that some operations need to be done after lattice expansion. This includes:

- The `ran` and `ran_gauss` functions, when used with elements that show up multiple times in a lattice, generally need to be used after lattice expansion. See §3.14.
- Some dependent parameters may be set as if they are independent parameters but only if done before lattice expansion. See §5.1.
- Setting the `phi0_multipass` attribute for an `Lcavity` or `RFcavity` multipass slave may only be done after lattice expansion (§9).
- Setting individual element attributes for tagged elements can only be done after lattice expansion (§7.8).

Notice that all `lines` (§7.2), `lists` (§7.6), and `use` (§7.7) statements necessary for lattice expansion must come before an `expand_lattice` statement.

Lattice expansion is only done once so it is an error if multiple `expand_lattice` statements are present.

The steps used for lattice expansion are:

1. Instantiate all of the lines listed in the last `use` statement (§7.7). If an instantiated line has `fork` or `photon_fork` (§4.23) elements, instantiate the lines connected to the fork elements if the `fork` or `photon_fork` is connected to a new branch. Instantiation of a given line involves:
   (a) Line expansion (§7) where the element sequence is constructed from the line and sub-lines.
   (b) Adding any superpositions (§8).
2. Form multipass lords and mark the appropriate multipass slaves (§9).
3. Add girder control elements (§4.24).
4. Add group (§4.26) and overlay (§4.40) control elements.

A lattice file where all the statements are post lattice expansion valid is called a "`secondary lattice file`". To promote flexibility, *Bmad* has methods for parsing lattices in a two step process: First, a "primary" lattice file that defines the basic lattice is read. After the primary lattice has been parsed and lattice expansion has been done, the second step is to read in one or more secondary lattice files. Such secondary lattice files can be used, for example, to set such things as element misalignments. The point here is that there are no calls (§3.19) of the secondary files in the primary file so the primary lattice file does not have to get modified when different secondary files are to be used.

3.27 Calc_Reference_Orbit Statement

The `calc_reference_orbit` statement triggers the computation of the "standard" reference orbit which is defined to be the closed orbit if the geometry (set by `parameter[geometry]` (§10.1)) is closed and which is defined to be the orbit as calculated from the starting position (set by any `particle_start` statements (§10.2)) if the geometry is open.

The `calc_reference_orbit` statement is used before a `merge_elements` statement (§3.28) to signal that the maps of the Taylor elements produced by the `merge_elements` statement are computed from the standard reference orbit and not the zero orbit.

The `calc_reference_orbit` statement must come after an `expand_lattice` command (§3.26).

The `calc_reference_orbit` statement can only be used if the standard reference orbit can be computed. For example, for a lattice with closed geometry, the closed orbit must exist.

The `calc_reference_orbit` statement has no arguments. Example:

```
expand_lattice          ! Expand the lattice.
calc_reference_orbit    ! Compute the reference orbit.
```

3.28 Merge_Elements Statement

The `merge_elements` statement is used to merge groups of consecutive elements into single taylor elements (§4.52) for faster tracking. The syntax of this statement is

```
merge_elements <list>
```

where `<list>` is a list of elements that are *not* to be combined.

Example:

```
call, file = my_ring.bmad     ! Read in a lattice
parameter[taylor_order] = 3   ! Set the taylor order for the maps
expand_lattice                ! Expand the lattice
calc_reference_orbit          ! ref orbit calc.
merge_elements sbend::*, bb ! Merge everything but sbend elements and
                              !   any elements named BB.
```

In this example, groups of elements that are between bends (with the exception of any elements named "BB"), are replaced by `taylor` elements. The order of the Taylor maps is set by `parameter[taylor_order]` (§10.1).

The `merge_elements` statement must come after an `expand_lattice` command (§3.26).

If there is a `calc_reference_orbit` statement before the `merge_elements` statement, the "standard" reference orbit (§3.27) is used for the computation of the Taylor maps. Otherwise the zero orbit is used as the reference orbit.

3.29 Combine_Consecutive_Elements Statement

The `combine_consecutive_elements` statement is used to combine consecutive elements with the same name into a single element. If a marker element has been placed in between two element with the same name, the marker element will be discarded. This can be a useful statement to add when given a lattice where elements have been split into two.[1] The `combine_consecutive_elements` statement must come after lattice expansion (§3.26). Example:

```
m: marker
myline: line = (q1, m, q1)
use, myline
expand_lattice
combine_consecutive_elements
```

In this case the finished line will will have a single `q1` element whose length will be twice the length of q1

3.30 Remove_Elements Statement

The `remove_elements` statement is used to remove elements from the lattice. The `remove_elements` statement must come after lattice expansion (§3.26). The syntax of the `remove_elements` statement is

```
remove_elements <element-list>
```

where `<element-list>` is a list of elements (§3.6). For example

```
expand_lattice              ! Lattice expansion must happen first
remove_elements overlay::*
```

In this example, all overlay elements are to be removed. This is useful, for example, when direct control of overlay slave parameters is desired.

3.31 Slice_Lattice Statement

The `slice_lattice` statement is used to remove elements from the lattice. The `slice_lattice` is useful when analysis of only part of the lattice is desired and the analysis of the entire lattice can take a significant amount of time.

The `slice_lattice` statement must come after lattice expansion (§3.26). The syntax of the `slice_lattice` statement is

```
slice_lattice <element-list>
```

where `<element-list>` is a list of elements (§3.6). For example

```
expand_lattice           ! Lattice expansion must happen first
slice_lattice q1##2:357,end
```

In this example, all elements outside of the range from element `q1##2` (the second instance of `q1` in the lattice) to element with index number 357 are discarded except for the element named `end` (which is typically the last element in any lattice branch). Additionally, the lord elements (§2.4) of any elements that remain are retained and the `beginning` element at the start of any branch is also retained.

For any lattice branch where elements are removed, the Twiss parameters and reference energy is computed, and the Twiss parameters and reference energy at the entrance end of the first element that is not removed is transferred to the beginning branch element. The branch `geometry` is also set to `open`.

[1] This is common in lattices translated from MAD.

For a lattice branch with a closed geometry, the Twiss parameters are computed with the RF on. That is, the reference momentum at the beginning of the sliced lattice branch, which is computed from the closed orbit phase space p_z of the unsliced lattice, may be non-zero. This will affect the Twiss calculation. If this is not what is wanted, the RF can be turned off before the `slice_lattice` command which will ensure the reference momentum is zero at the beginning of the lattice branch. Example:

```
expand_lattice
rfcavity::*[is_on] = False    ! Turn RF off
slice_lattice q1##2:357       ! Slice the lattice
! This shows how to reset the RF and geometry if needed.
rfcavity::*[is_on] = True     ! Turn RF back on
parameter[geometry] = closed  ! Change the geometry.
... etc ...
```

To create a lattice slice that wraps around the lattice ends, that is, joins a section at the end of the lattice followed by a section at the beginning of the lattice, use a `start_branch_at` statement before a slice_lattice statement. Example:

```
expand_lattice
start_branch_at Q9
slice_lattice Q9:Q1   ! With Q1 being before Q9 in the original lattice.
```

3.32 Start_Branch_At Statement

The `start_branch_at` statement is used to shift the starting point of a lattice branch while keeping the relative order of the elements the same. The syntax of the `start_branch_at` statement is

```
start_branch_at <lattice-element>                    ! or
start_branch_at, move_end_marker <lattice-element>
```

where `<lattice-element>` is the name or index of a lattice element to be moved to the start of the branch the element is in. The `start_branch_at` statement must come after lattice expansion (§3.26).

The shifting only applies to elements in the tracking part of the lattice (§2.4). The `BEGINNING` element of the branch (§7.7) and any lord elements are unaffected. The `end` marker element (§7.1), if it is present (that is, if no `parameter[no_end_marker]` is used), will also remain at the end of the branch except if the `move_end_marker` option is used with `start_branch_at`. For example:

```
expand_lattice           ! Lattice expansion must happen first
start_branch_at Q3
```

In this example, the elements in the lattice branch containing `Q3` would be shifted so that the `Q3` element is the first element in the branch. Thus, if the lattice branch originally consisted of the elements

```
Beginning, Q1,  Q2,  Q3,  Q4,  Q5,  Q6, End
```

then the shifted lattice would be

```
Beginning, Q3,  Q4,  Q5,  Q6, Q1,  Q2, End
```

Elements that originally come before the new starting point are always wrapped around to the end of the branch. If the move_end_marker is present:

```
start_branch_at, move_end_marker Q3
```

then the shifted lattice would be

```
Beginning, Q3,  Q4,  Q5,  Q6, End, Q1,  Q2
```

Also see the `slice_lattice` statement (§3.31).

3.33 Debugging Statements

There are a few statements which can help in debugging the *Bmad* lattice parser itself. That is, these statements are generally only used by programmers. These statements are:

```
debug_marker
no_digested
parser_debug
write_digested
```

The `debug_marker` statement is used for marking a place in the lattice file where program execution is to be halted. This only works when running a program in conjunction with a program debugging tool.

The `no_digested` statement if present, will prevent *Bmad* from creating a digested file (§3.2). That is, the lattice file will always be parsed when a program is run. The `write_digested` statement will cancel a `no_digested` statement.

The `parser_debug` statement will cause information about the lattice to be printed out at the terminal. The syntax is

```
parser_debug <switches>
```

Valid `<switches>` are

```
particle_start      ! Print the particle_start information.
const               ! Print table of constants defined in the lattice file.
ele <n1> <n2> ...   ! Print full info on selected elements.
lattice             ! Print a list of lattice element information.
lord                ! Print full information on all lord elements.
seq                 ! Print sequence information.
slave               ! Print full information on all slave elements.
time                ! Print timing information.
```

Here $< n1 >$, $< n2 >$, etc. are the index of the selected elements in the lattice. Example

```
parser_debug var lat ele 34 78
```

Chapter 4

Lattice Elements

A lattice is made up of a collection of elements — quadrupoles, bends, etc. This chapter discusses the various types of elements available in *Bmad*.

Element	*Section*	*Element*	*Section*
AB_Multipole	4.1	Mask	4.33
AC_Kicker	4.2	Match	4.34
BeamBeam	4.3	Monitor	4.28
Beginning_Ele	4.4	Multipole	4.36
Converter	4.8	Null_Ele	4.38
Crab_Cavity	4.9	Octupole	4.39
Custom	4.11	Patch	4.41
Drift	4.14	Photon_Fork	4.23
E_Gun	4.15	Pipe	4.28
Ecollimator	4.7	Quadrupole	4.43
ElSeparator	4.16	Rbend	4.5
EM_Field	4.17	Rcollimator	4.7
Feedback	4.18	RF_bend	4.45
Fiducial	4.19	RFcavity	4.46
Fixer	4.20	Sad_Mult	4.47
Floor_Shift	4.21	Sbend	4.5
Foil	4.22	Sextupole	4.49
Fork	4.23	Sol_Quad	4.50
GKicker	4.25	Solenoid	4.51
HKicker	4.29	Taylor	4.52
Hybrid	4.27	Thick_Multipole	4.53
Instrument	4.28	Undulator	4.54
Kicker	4.30	VKicker	4.29
Lcavity	4.31	Wiggler	4.54
Marker	4.32		

Table 4.1: Table of element types suitable for use with charged particles. Also see Table 4.3

Most element types available in *MAD* are provided in *Bmad*. Additionally, *Bmad* provides a number of element types that are not available in *MAD*. A word of caution: In some cases where both *MAD* and

Bmad provide the same element type, there will be an overlap of the attributes available but the two sets of attributes will not be the same. The list of element types known to *Bmad* is shown in Table 4.1, 4.2, and 4.3. Table 4.1 lists the elements suitable for use with charged particles, Table 4.2 which lists the elements suitable for use with photons, and finally Table 4.3 lists the `controller` element types that can be used for parameter control of other elements. Note that some element types are suitable for both particle and photon use.

Element	*Section*	*Element*	*Section*
`Beginning_Ele`	4.4	`Marker`	4.32
`Capillary`	4.6	`Mask`	4.33
`Crystal`	4.10	`Match`	4.34
`Custom`	4.11	`Monitor`	4.28
`Detector`	4.12	`Mirror`	4.35
`Diffraction_Plate`	4.13	`Multilayer_Mirror`	4.37
`Drift`	4.14	`Patch`	4.41
`Ecollimator`	4.7	`Photon_Fork`	4.23
`Fiducial`	4.19	`Photon_Init`	4.42
`Floor_Shift`	4.21	`Pipe`	4.28
`Fork`	4.23	`Rcollimator`	4.7
`GKicker`	4.25	`Sample`	4.48
`Instrument`	4.28		

Table 4.2: Table of element types suitable for use with photons. Also see Table 4.3

Element	*Section*	*Element*	*Section*
`Group`	4.26	`Overlay`	4.40
`Girder`	4.24	`Ramper`	4.44

Table 4.3: Table of controller elements.

For a listing of element attributes for each type of element, see Chapter §15.

4.1 AB_Multipole

An `ab_multipole` is a thin magnetic multipole lens up to 21^{st} order. The basic difference between this and a `multipole` (§4.36) is the input format. See section §18.1 for how the multipole coefficients are defined.

General `ab_multipole` attributes are:

Attribute Class	§	*Attribute Class*	§
a*n*, b*n* multipoles	5.15	Length	5.13
Aperture limits	5.8	Offsets & tilt	5.6
Chamber wall	5.12	Reference energy	5.5
Custom Attributes	3.9	Superposition	8
Description strings	5.3	Tracking & transfer map	6
Is_on	5.14		

See §15.3 for a full list of element attributes along with a their units.

The length `l` is a fictitious length that is used for synchrotron radiation computations and affects the longitudinal position of the next element but does not affect any tracking or transfer map calculations. The `x_pitch` and `y_pitch` attributes are not used in tracking.

When an `ab_multipole` is superimposed (§8) on a lattice, it is treated as a zero length element and in this case it is an error for the length of the `ab_multipole` to be set to a nonzero value.

Unlike a `multipole`, an `ab_multipole` will *not* affect the reference orbit if there is a dipole component.

Example:

```
abc: ab_multipole, a2 = 0.034e-2, b3 = 5.7, a11 = 5.6e6/2
```

4.2 AC_Kicker

An `ac_kicker` element simulates a time dependent kicker element.

General `ac_kicker` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Is_on	5.14
Chamber wall	5.12	Length	5.13
Custom Attributes	3.9	Mag & Elec multipoles	5.15
Description strings	5.3	Offsets, pitches & tilt	5.6
Field Maps	5.16	Reference energy	5.5
Fringe Fields	5.21	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Tracking & transfer map	6

See §15.4 for a full list of element attributes along with a their units.

Attributes specific to a `ac_kicker` element are:

```
t_offset        = <Real>              ! Time offset of field waveform.
interpolation   = <switch>            ! cubic (default) or linear.
amp_vs_time = {(<time1>, <Amp1>), (<time2>, <Amp2>), ...}  ! Field amp vs Time.
frequencies = {(<freq1>, <Amp1>, <phi1>),
                            (<freq2>, <Amp2>, <phi2>), ...} ! Freq components.
```

Note: The `frequencies` attribute phases `phi1`, `phi2`, `phi3`, etc., have units of `radians/2pi`.

An `ac_kicker` element is like a `kicker` (§4.30) element except that the field varies in time. The field is calculated in two steps:

1. Calculate the field the same as for a `kicker` element (§4.30).
2. Scale the field using the function $A(\delta t)$ (discussed below)

$$B(\delta t) = A(\delta t)\, B_0, \qquad E(\delta t) = A(\delta t)\, E_0 \tag{4.1}$$

   where B and E are the applied magnetic and electric fields, and B_0 and E_0 are the fields as calculated as if the element where a `kicker` ignoring the time dependence.

$\delta t = t_{eff} - t_0$ where t_{eff} is the effective time as discussed in §26.1 and t_0 is the value of the `t_offset` attribute.

There are two ways to specify the dimensionless time variation $A(\delta t)$ of the field. One way is to specify points $(A, \delta t)$ using the `amp_vs_time` attribute. Example:

```
mk: ac_kicker, l = 0.3, scale_multipoles = F, b1 = 0.27, t_offset = 3.6e-8,
        amp_vs_time = {(-1.2e-6, 0.02), ... }
```

The element in this example is an AC quadrupole kicker. The times (in seconds) must be in ascending order and no two times may be the same. The method used to interpolate between the time points is determined by the setting of the `interpolate` parameter which may be one of

```
linear      ! Linear interpolation.
cubic       ! Cubic spline interpolation (default).
```

For times outside of the range specified by `amp_vs_time`, the amplitude will be extrapolated. For the cubic spline, extrapolation is only permitted over a distance outside the time range equal to the time difference between an end point and the next nearest point.

The second way to specify the waveform is to specify the frequencies in the spectrum using the `frequencies` attribute. In this case the amplitude is:

$$A(t) = \sum_i A_i \cos(2\pi(f_i \delta t + \phi_i)) \tag{4.2}$$

Example:

```
mk: ac_kicker, l = 0.3, field_calc = fieldmap, cartesian_map = {...},
        frequencies = {(3.4e6, 0.34, 0.12), ...}, a0_elec = 1e5
```

When using a frequency spectrum, the `interpolate` attribute is ignored. Note: The units of the phases `phi` with the `frequencies` attribute are `rad/2pi`.

To specify an `amp_vs_time` component after an `ac_kicker` element has been defined, use the syntax

```
name[AMP_VS_TIME(i)%time]      ! Time of i^th point.
name[AMP_VS_TIME(i)%amp]       ! Amplitude of i^th point.
```

where `name` is the name of the element and `i` is the index of the point. To specify a `frequencies` component after an `ac_kicker` element has been defined, use the syntax

```
name[FREQUENCIES(i)%freq]     ! Frequency of i^th spectrum point.
name[FREQUENCIES(i)%amp]      ! Amplitude of i^th spectrum point.
name[FREQUENCIES(i)%phi]      ! Phase (rad/2pi) of i^th spectrum point.
```

Example:

```
mk: ac_kicker, amp_vs_time = {(-1.2e-6, 0.02), ... }
mk[amp_vs_time(2)%amp] = 0.03  ! Change 2nd point amplitude.
```

When specifying the time dependent using a set of frequencies, it is generally advisable to use absolute time tracking (§26.1). This can be done in the lattice file by setting

```
bmad_com[absolute_time_tracking] = T
```

Note: The calculated field will only obey Maxwell's equations in the limit that the time variation of the field is "slow":

$$\omega \ll \frac{c}{r} \tag{4.3}$$

where ω is the characteristic frequency of the field variation, c is the speed of light, and r is the characteristic size of the `ac_kicker` element. That is, the fields at opposite ends of the element must be able to "communicate" (which happens at the speed of light) in a time scale short compared to the time scale of the change in the field.

4.3 BeamBeam

A `beambeam` element simulates an interaction with an opposing ("strong") beam traveling in the opposite direction. The strong beam is assumed to be Gaussian in shape. In the `bmad_standard` calculation the beam–beam kick is computed using the Bassetti–Erskine complex error function formula[Talman87]

General `beambeam` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Is_on	5.14
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Reference energy	5.5
Description strings	5.3	Superposition	8
Is_on	5.14	Tracking & transfer map	6

See §15.5 for a full list of element attributes along with a their units.

Attributes specific to a `beambeam` element are:

```
sig_x          = <Real>    ! Horizontal strong beam sigma at the center
sig_y          = <Real>    ! Vertical strong beam sigma at the center
sig_z          = <Real>    ! Strong beam length
charge         = <Real>    ! Strong beam charge. Default = -1
n_particle     = <Real>    ! Number of particles in strong beam.
n_slice        = <Int>     ! Number of strong beam slices
crab_x1        = <Real>    ! Crabbing linear coefficient.
crab_x2        = <Real>    ! Crabbing quadratic coefficient.
crab_x3        = <Real>    ! Crabbing cubic coefficient.
crab_x4        = <Real>    ! Crabbing 4th order coefficient.
crab_x5        = <Real>    ! Crabbing 5th order coefficient.
crab_tilt      = <Real>    ! Crabbing tilt.
species_strong = <Species> ! Strong beam species
E_tot_strong   = <Real>    ! Strong beam particle energy.
pc_strong      = <Real>    ! Strong beam momentum.
beta_a_strong  = <Real>    ! Strong beam $a$-mode beta Twiss parameter
alpha_a_strong = <Real>    ! Strong beam $a$-mode alpha Twiss parameter
beta_b_strong  = <Real>    ! Strong beam $b$-mode beta Twiss parameter
alpha_b_strong = <Real>    ! Strong beam $b$-mode alpha Twiss parameter
bbi_constant               ! See below. Dependent attribute (§5.1).
ks             = <Real>    ! Solenoid strength.
bs_field       = <Real>    ! Solenoid field strength.
field_master   = <T/F>     ! Is ks or bs_field value the master (§5.2)?
s_beta_ref     = <Real>    ! Reference position of strong beam Twiss.
crossing_time  = <Real>    ! Time when strong beam center reaches IP.
repetition_frequency = <Real>  ! Strong beam repetition rate.
```

The strength of the strong beam is set by:

```
charge * n_particle
```

The default The default value of `charge` is -1 which indicates that the strong beam has the opposite charge of the weak beam. The default for `n_particle` is 0.

For historical reasons, the global parameter `parameter[n_part]` (§10.1) will be used in place of `n_particle` if `n_particle` has a value of 0.

`sig_z` are the strong beam's longitudinal sigma. The strong beam is divided up into `n_slice` equal charge (not equal thickness) slices. Propagation through the strong beam involves a kick at the charge center of each slice with propagation between slice centers. A solenoid field can be set for the regions in between the slice centers. The kicks are calculated using the standard Bassetti–Erskine complex error function formula[Talman87]. Even though the strong beam can have a finite `sig_z`, the length of the `beambeam` element is zero. This is achieved by propagating a particle at the beginning and at the end of tracking so that the longitudinal starting point and ending points are at the `beambeam` element. Documentation of how a particle is tracked through a `beambeam` element is given in §26.5.

The `ks` and `bs_field` parameters are the normalized and unnormalized solenoid strengths (§5.2) related through Eq. (18.3). If the beambeam element is superimposed on top of a solenoid, the beambeam element will inherit the solenoid field strength from the solenoid element instead.

The strong beam Twiss parameters `beta_a_strong`, `beta_b_strong`, `alpha_a_strong`, and `alpha_b_strong` are the Twiss parameters of the strong beam at the s-position given by `s_twiss_ref`. Additionally, `sig_x`, `sig_y` are the transverse sigmas of the strong beam at this point. `S_beta_ref` is measured relative to the position of the`beambeam` element. If `beta_a_strong` is zero (the default), the a-mode Twiss parameters as calculated from the lattice is used. Similarly, if `beta_b_strong` is zero (the default), the b-mode Twiss parameters as calculated from the lattice is used. To calculate the sigmas of any given slice, `sig_x` and `sig_y` are extrapolated using the Twiss parameters at `s_twiss_ref`.

The `x_offset`, `y_offset`, and `z_offset` attributes (§5.6) are used to offset the strong beam. The `x_pitch` and `y_pitch` parameters orient the strong beam with respect to the laboratory coordinate system. This will be give the beam–beam interaction a crossing angle. The full crossing angle is the angle of the strong beam (set by `x_pitch` and `y_pitch`) with respect to the trajectory of the weak beam centroid. The weak beam centroid orbit will be the closed orbit if the lattice geometry is closed. If the lattice geometry is open, the weak beam centroid orbit is determined by the beginning centroid orbit of the weak beam (which can be program dependent) and details of the lattice between the beginning of the lattice and the `beambeam` element.

To curve the strong beam centroid to simulate crabbing, the following parameters can be used:

```
crab_x1,    crab_x2,    crab_x3
crab_x4,    crab_x5,    crab_tilt
```

If `crab_tilt` is zero (the default), the strong beam centroid (x_c, y_c) will have y_c zero and

$$x_c(z) = \text{crab_x1} \cdot z + \text{crab_x2} \cdot z^2 + \text{crab_x3} \cdot z^3 + \text{crab_x4} \cdot z^4 + \text{crab_x5} \cdot z^5 \tag{4.4}$$

where positive x_c and y_c are the same as positive x and positive y for the weak beam and z is the longitudinal position with respect to the strong beam center with positive z being towards the head of the strong bunch (and remember that since the strong bunch is going in the opposite direction, the head of the strong bunch is opposite that of the weak bunch). With a finite `crab_tilt`, the curvature is rotated around the z axis as shown in figure 5.2.

The `bbi_constant` is a measure of the beam–beam interaction strength. It is a dependent variable and is calculated from the equation

$$C_{bbi} = N\, m_e\, r_e / (2\,\pi\,\gamma\,(\sigma_x + \sigma_y)) \tag{4.5}$$

In the linear region, near $x = y = 0$, the beam–beam kick is approximately

$$\begin{aligned} k_x &= -4\,\pi\,x\,C_{bbi}/\sigma_x \\ k_y &= -4\,\pi\,y\,C_{bbi}/\sigma_y \end{aligned} \tag{4.6}$$

and the linear beam–beam tune shift is

$$\begin{aligned} dQ_x &= C_{bbi}\,\beta_x/\sigma_x \\ dQ_y &= C_{bbi}\,\beta_y/\sigma_y \end{aligned} \tag{4.7}$$

The `species_strong`, `pc_strong`, and `E_tot_strong` give the particle species, momentum and energy of the strong beam. This is only relevant if the velocity of the strong beam is not equal to the velocity of the weak beam. Only one of `pc_strong` or `E_tot_strong` needs to be set as one will be calculated from the other. If both are set, the setting of `pc_strong` will be used.

The `crossing_time` parameter sets when, relative to the reference time, the center of the strong beam crosses the plane of the beambeam element (the IP). For example, if tracking is done with radiation damping on, the (weak beam) closed orbit at the IP will shift in time. To have the weak beam and strong beam centers cross the plane of the `beambeam` element at the same time, the value of `crossing_time` should be set to the value $-z/c\beta$ at the IP where z is the closed orbit phase space z and β is the closed orbit velocity.

When absolute time tracking (§26.1) is in use, the `repetition_frequency` parameter (along with the `crossing_time` parameter) is used to calculate the time that the strong bunch crosses the plane of the `beambeam` element. Generally, this frequency should be set equal to the fundamental RF frequency or some harmonic thereof. If this frequency is zero (the default), *Bmad* will assume that the repetition period is a harmonic of the reference particle oscillation time so that in this case a particle's phase space z coordinate will be used.

Example:

```
bbi: beambeam, sig_x = 3e-3, sig_y = 3e-4, x_offset = 0.05, n_particle = 1.3e9
```

4.4 Beginning_Ele

A `beginning_ele` element, named "`BEGINNING`", is placed at the beginning of every branch (§2.2) of a lattice to mark the start of the branch. The `beginning_ele` always has element index 0 (§2). The creation of this `beginning_ele` element is automatic and it is not permitted to define a lattice with `beginning_ele` elements at any other position.

The attributes of the `beginning_ele` element in the root branch are are generally set using `beginning` (§10.4) statements or line parameter (§10.4) statements. [The attributes of other `beginning_ele` elements are set solely with line parameter statements.]

If the first element after the `beginning_ele` element at the start of a branch is reversed (§7.4), the `beginning_ele` element will be marked as reversed so that a reflection patch is not needed in this circumstance.

See §15.6 for a full list of element attributes.

Note: To set Twiss and/or orbit parameters at a point not at the beginning of a branch, use a `fixer` (§4.20) element.

4.5 Bends: Rbend and Sbend

Rbends and sbends are dipole bends. The difference is that rbend elements use a Cartesian ("rectangular") coordinate system to describe the shape of the magnet while sbend elements use a polar ("sector") coordinate system.

For any given sbend element it is possible to construct an equivalent rbend element that has the same shape and vice versa. Given this, and to simplify internal bookkeeping, all rbend elements are converted to sbend elements when a lattice is read in to a program. In order to preserve the information as to whether a bend element was originally specified as an sbend or an rbend in the lattice file, all bend elements have a sub_key parameter which is appropriately set when the lattice is parsed. This sub_key parameter does not affect tracking and is only used if a new lattice file is generated by the program.

General rbend and sbend attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Mag & Elec multipoles	5.15
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Overlapping Fields	5.18
Description strings	5.3	Reference energy	5.5
Fringe Fields	5.21	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Is_on	5.14	Field Maps	5.16
Integration settings	6.4	Tracking & transfer map	6
Length	5.13		

See §15.7 for a full list of element attributes along with a their units.

Attributes specific to rbend and sbend elements are:

```
angle              = <Real>   ! Design bend angle. Dependent var (§5.1).
b_field            = <Real>   ! Design field strength (= P_0 g / q) (§5.1).
db_field           = <Real>   ! Actual - Design bending field difference (§5.1).
b_field_tot                   ! Net field = b_field + db_field. Dependent param (§5.1).
b1_gradient        = <Real>   ! Quadrupole field strength (§5.1).
b2_gradient        = <Real>   ! Sextupole field strength (§5.1).
e1, e2             = <Real>   ! Face angles.
exact_multipoles   = <Switch> ! Curved coordinate correction? off is default.
fint, fintx        = <Real>   ! Face field integrals.
g                  = <Real>   ! Design bend strength (= 1/rho).
dg                 = <Real>   ! Actual - Design bend strength difference (§5.1).
g_tot                         ! Net design strength = g + dg Dependent param (§5.1).
h1, h2             = <Real>   ! Face curvature.
hgap, hgapx        = <Real>   ! Pole half gap.
k1                 = <Real>   ! Quadrupole strength.
k2                 = <Real>   ! Sextupole strength (§5.1).
l                  = <Real>   ! "Length" of bend. See below.
l_arc              = <Real>   ! Arc length. For rbends only.
l_chord            = <Real>   ! Chord length. See §5.13.
l_rectangle        = <Real>   ! "Rectangular" length.
l_sagitta                     ! Sagittal length. Dependent param (§5.1).
ptc_field_geometry  = <Switch> ! See below. Default is sector.
```

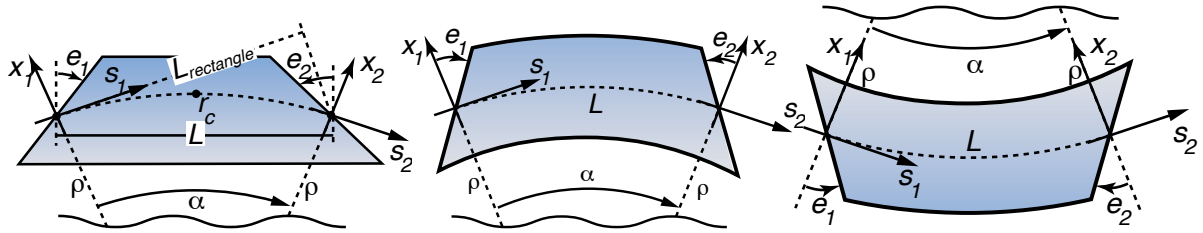

(a) Rbend (b) Sbend (c) Reversed sbend

Figure 4.1: Coordinate systems for (a) normal (non-reversed)`rbend`, (b) normal `sbend`, and (c) `reversed sbend` elements. The bends are viewed from "above" (viewed from positive y). Normal bends have `g`, `angle`, and `rho` all positive. Reversed bends have `g`, `angle`, and `rho` all negative. The face angles `e1` and `e2` are drawn for `reference_pt` set to `none` or `center`. For (a) and (b), as drawn, the `e1` and `e2` face angles are both positive. For (c), as drawn, `e1` and `e2` are both negative. In all cases, `L` is positive. Notice that for reversed bends, the x-axis points towards the center of the bend while for normal bends the x-axis points towards the outside.

```
ptc_fringe_geometry = <Switch> ! Sec. §5.21.2
rho                 = <Real>   ! Design bend radius. Dependent param (§5.1).
roll                = <Real>   ! See 5.6.
fiducial_pt         = <switch> ! See below. Default is none.
field_master        = <T/F>    ! See 5.2.
```

angle

The total design bend angle. A positive `angle` represents a bend towards negative x values (see Fig. 16.2).

B_field, dB_field

The `B_field` parameter is the design magnetic bending field which determines the reference orbit and the placement of lattice elements downstream from the bend. The `dB_field` parameter is the difference between the actual ("total") field and the design field. Thus:

```
Actual B-field = B_field + dB_field
```

See the discussion of `g` and `dg` below for more details.

e1, e2

The rotation angle of the entrance pole face is `e1` and at the exit face it is `e2`. Zero `e1` and `e2` for an `rbend` gives a rectangular magnet (Fig. 4.1a). Zero `e1` and `e2` for an `sbend` gives a wedge shaped magnet (Fig. 4.1b). An `sbend` with an `e1` = `e2` = `angle`/2 is equivalent to an `rbend` with `e1` = `e2` = 0. This formula holds for both positive and negative angles. For `rbend` elements, the above discussion is true if `fiducial_pt` is set to `none` or `center`. If `fiducial_pt` is set to `entrance_end`, then the face angles are measured with respect to the entrance coordinates $(s_1 x_1)$. If the `fiducial_pt` is set to `exit_end`, the face angles are measured with respect to the exit coordinates (s_2, x_2). Thus

```
e1(f_pt=none) = e1(f_pt=entrance_end) - angle/2 = e1(f_pt=exit_end) + angle/2
e2(f_pt=none) = e2(f_pt=entrance_end) + angle/2 = e2(f_pt=exit_end) - angle/2
```

Note: The correspondence between `e1` and `e2` and the corresponding parameters used in the SAD program [SAD] is:

```
e1(Bmad) =  e1(SAD) * angle + ae1(SAD)
e2(Bmad) =  e2(SAD) * angle + ae2(SAD)
```

exact_multipoles

The `exact_multipoles` switch can be set to one of:

```
off                  ! Default
vertically_pure
horizontally_pure
```

This switch determines if the multipole fields, both magnetic and electric, and including the `k1` and `k2` components, are corrected for the finite curvature of the reference orbit in a bend. See §18.3 for a discussion of what `vertically` pure versus `horizontally` pure means. Setting `exact_multipoles` to `vertically_pure` means that the individual a_n and b_n multipole components are used with the vertically pure solutions

$$\mathbf{B} = \sum_{n=0}^{\infty} \left[\frac{a_n}{n+1} \nabla\phi_n^r + \frac{b_n}{n+1} \nabla\phi_n^i \right], \qquad \mathbf{E} = \sum_{n=0}^{\infty} \left[\frac{a_{en}}{n+1} \nabla\phi_n^i + \frac{b_{en}}{n+1} \nabla\phi_n^r \right] \tag{4.8}$$

and if `exact_multipoles` is set to `horizontally_pure` the horizontally pure solutions ψ_n^r and ψ_n^i are used instead of the vertically pure solutions ϕ_n^r and ϕ_n^i.

To use exact multipoles with PTC based tracking (§6), the PTC exact model tracking must be turned on. That is, in the lattice file set:

```
ptc_com[exact_model] = T
```

With exact model tracking, PTC always assumes that multipole coefficient values correspond to `horizontally_pure`. In this case, *Bmad* will convert `vertically_pure` to `horizontally_pure` as needed when passing multipole coefficients to PTC. Note that in the case where PTC is doing exact model tracking (§6.4) but the `exact_multipoles` switch is set to `off`, PTC will still be treating the multipoles as `horizontally_pure` even though *Bmad* tracking will be treating them as straight line multipoles. Note: If the bend has an associated electric field, PTC will always be doing exact modeling.

fint, fintx, hgap, hgapx

The field integrals for the entrance pole face is given by the product of the `fint` and `hgap` parameters with `hgap` being the half gap between poles at the entrance face

$$F_{H1} \equiv F_{int}\, H_{gap} = \int_{pole} ds\, \frac{B_y(s)\,(B_{y0} - B_y(s))}{2\,B_{y0}^2} \tag{4.9}$$

For the exit pole face there is a similar equation using `fintx` and `hgapx` which defines F_{H2}. In the above equation B_{y0} is the field in the interior of the dipole. The values of `fint`, `fintx`, `hgap`, and `hgapx` are never used in isolation when tracking. Only the values for F_{H1} and F_{H2} matter.

If `fint` or `fintx` is given without a value then a value of 0.5 is used. If `fint` or `fintx` is not present, the default value of 0 is used. Note: *MAD* does not have the `fintx` and `hgapx` attributes. *MAD* just assumes that the values are the same for the entrance and exit faces. For compatibility with *MAD*, if `fint` is given but `fintx` is not, then `fintx` is set equal to `fint`. Similarly, `hgapx` will be set to `hgap` if `hgapx` is not given. Note that this setting of `fintx` or `hgapx` using the value of `fint` or `hgap` will only be done before lattice expansion (§3.26).

Note: To have an effect, both `fint` and `hgap` (or `fintx` and `hgapx`) must be non-zero.

Note: The SAD program uses `fb1+f1` for the entrance fringe and `fb2+f1` for the exit fringe. The correspondence between the two is

F_{H1} `= fint  * hgap  = (fb1 + f1) / 12`
F_{H2} `= fintx * hgapx = (fb2 + f1) / 12`

`fint` and `hgap` can be related to the Enge function which is sometimes used to model the fringe field. The Enge function is of the form

$$B_y(s) = \frac{B_{y0}}{1 + \exp[P(s)]} \tag{4.10}$$

where

$$P(s) = C_0 + C_1\, s + C_2\, s^2 + C_3\, s^3 + \ldots \tag{4.11}$$

The C_0 term simply shifts where the edge of the bend is. If all the C_n are zero except for C_0 and C_1 then

$$C_1 = \frac{1}{2\, H_{gap}\, F_{int}} \tag{4.12}$$

fiducial_pt

The `fiducial_pt` parameter sets a fiducial point which can be used to keep the shape of the bend constant when, in a program, the parameters `rho`, `g`, `b_field` or `angle` are varied. Varying these parameters typically happens when doing machine design. Using a fiducial point can be helpful when designing a machine using bend magnets that already exist.

The `fiducial_pt` parameter has four possible settings:

```
none          ! No fiducial point (default).
entrance_end  ! The entrance point is the fiducial point.
center        ! The center of the reference curve is the fiducial point.
exit_end      ! The exit point is the fiducial point.
```

With `fiducial_pt` set to `none` (the default). The bend shape is not held constant. With the other three settings, the bend shape will be held constant as discussed in §26.8. With `fiducial_pt` set to `entrance_end`, the reference trajectory at the entrance end is held fixed in both position and orientation with respect to the bend face and `g`, `l` and `e2`, along with the other dependent parameters, are adjusted to both give the desired change in what was varied (which is one of `rho`, `g`, `b_field` or `angle`) and to keep the shape of the bend unchanged. See Fig. 26.3a. Similarly, if `fiducial_pt` is set to `center`, the center of the reference trajectory is held fixed in both position and orientation and if `fiducial_pt` is set to `exit_end`, the exit point is held fixed in both position and orientation.

g, dg, rho

The design bending radius which determines the reference coordinate system is `rho` (see §16.1.1). `g` = `1/rho` is the curvature function and is proportional to the design dipole magnetic field. `g` is related to the design magnetic field `B_field` via

$$\mathrm{g} = \frac{q}{p_0}\, \mathrm{B_field} \tag{4.13}$$

where q is the charge of the reference particle and p_0 is the reference momentum. It is important to keep in mind that changing `g` will change the design orbit (§16) and hence will move all downstream lattice elements in space.

The parameter `dg` is the difference between the actual and the design bending strengths. The relationship between `dg` and `dB_field` is analogous to the relationship between `g` and `B_field` in Eq. (4.13). The actual ("total") field strength is given by the sum:

```
Actual g = g + dg
```

Changing `dg` leaves the design orbit and the positions of all downstream lattice elements unchanged but will vary a particle's orbit. One common mistake when designing lattices is to vary `g` and not `dg` which results in downstream elements moving around. See Sec. §13.2 for an example.

Note: A positive `g`, which will bend particles and the reference orbit in the $-x$ direction represents a field of opposite sign as the field due a positive `hkick`.

h1, h2

The attributes `h1` and `h2` are the curvature of the entrance and exit pole faces. They are present for compatibility with MAD but are not yet implemented in terms of tracking and other calculations.

k1, b1_gradient
The normalized and unnormalized (§5.2) quadrupole strength. See Eqs. (18.2) and (18.3).

k2, b2_gradient
The normalized and unnormalized (§5.2) sextupole strength. See Eqs. (18.2) and (18.3).

l, l_arc, l_chord, l_sagitta
For compatibility with MAD, for an `rbend`, `l` is the chord length and not the arc length as it is for an `sbend`. After reading in a lattice, *Bmad* will internally convert all `rbend`s into `sbend`s, and the `l_chord` attribute of the created `sbend` will be set to the input `l`. The `l` of the created `sbend` will be set to the true path length (see above). Alternatively for an `rbend`, instead of setting `l`, the `l_arc` attribute can be used to set the true arc length.

For `sbend` elements, `l_chord` will be set to the calculated chord length. For both types of bends, the `l_sagitta` parameter will be set to the sagitta length (The sagitta is the distance from the midpoint of the arc to the midpoint of the chord). `l_sagitta` can be negative and will have the same sign as the `g` parameter.

l_rectangle
The `l_rectangle` parameter is the "rectangular" length defined to be the distance between the entrance and exit points. The coordinate system used for the calculation is defined by the setting of `fiducial_pt`. Fig. 4.1a shows `l_rectangle` for `fiducial_pt` set to `entrance_end` (the coordinate system corresponds to the entrance coordinate system of the bend). In this case, and in the case where `fiducial_pt` is set to `exit_end`, the rectangular length will be $\rho \sin\alpha$. If `fiducial_pt` is set to `none` or `center`, `l_rectangle` is the same as the chord length.

ref_tilt
The `ref_tilt` attribute rotates a bend about the longitudinal axis at the entrance face of the bend. A bend with `ref_tilt` of $\pi/2$ and positive `g` bends the element in the $-y$ direction ("downward"). See Fig. 16.7. It is important to understand that `ref_tilt`, unlike the `tilt` attribute of other elements, bends both the reference orbit along with the physical element. Note that the MAD `tilt` attribute for bends is equivalent to the *Bmad* `ref_tilt`. Bends in *Bmad* do not have a `tilt` attribute.

Important! Do not use `ref_tilt` when doing misalignment studies for a machine. Trying to misalign a dipole by setting `ref_tilt` will affect the positions of all downstream elements! Rather, use the `roll` parameter.

The difference between `rbend` and `sbend` elements is the way the `l`, `e1`, and `e2` attributes are interpreted. To ease the bookkeeping burden, after reading in a lattice, *Bmad* will internally convert all `rbend`s into `sbend`s. This is done using the following transformation on `rbend`s:

```
l_chord(internal) = l(input)
l(internal) = 2 * asin(l_chord * g / 2) / g
e1(internal) = e1(input) + theta / 2
e2(internal) = e2(input) + theta / 2
```

The attributes `g`, `angle`, and `l` are mutually dependent. If any two are specified for an element *Bmad* will calculate the appropriate value for the third. After reading in a lattice, `angle` is considered the dependent variable so if `l` or `g` is varied, the value of `angle` will be set to `g * l`. if `theta` is varied, `l` will be set accordingly.

Since internally all `rbend`s are converted to `sbend`s, if one wants to vary the `g` attribute of a bend and still keep the bend rectangular, an overlay (§4.40) can be constructed to maintain the proper face angles. For example:

```
l_ch = 0.54
g_in = 1.52
l_coef = asin(l_ch * g_in / 2) / g_in
my_bend: rbend, l = l_ch, g = g_in
my_overlay: overlay = {my_bend, my_bend[e1]:l_coef, my_bend[e2]:l_coef},
              var = {g}, g = g_in
```

Notice that `l_coef` is just `arc_length/2`.

In the local coordinate system (§16.1.1), looking from "above" (bend viewed from positive y), and with `ref_tilt` = 0, a positive `angle` represents a particle rotating clockwise. In this case. `g` will also be positive. For counterclockwise rotation, both `angle` and `g` will be negative but the length `l` is always positive. Also, looking from above, a positive `e1` represents a clockwise rotation of the entrance face and a positive `e2` represents a counterclockwise rotation of the exit face. This is true irregardless of the sign of `angle` and `g`. Also it is always the case that the pole faces will be parallel when

```
e1 + e2 = angle
```

Example bend specification:

```
b03w: sbend, l = 0.6, k1 = 0.003, fint  ! gives fint = fintx = 0.5
```

`ptc_field_geometry` determines what reference coordinates PTC uses within a bend for calculating higher order fields. This only affects tracking if PTC is being used and if `ptc_com[exact_model]` is set to True (§11.4). Possible values for `ptc_field_geometry` are:

```
sector      ! Default
straight
```

For `sector` reference coordinates, the field coordinate reference frame is with respect to the arc of the reference trajectory. For `straight` coordinates the coordinate reference frame is with respect to the chord line. For a bend where there are no other fields besides the basic dipole field, tracking is essentially unaffected.[1] When there are quadrupole or higher order fields, the fields are centered about the reference frame set by `ptc_field_geometry`. Since *Bmad* based tracking does not implement `straight` geometry tracking, *Bmad* and PTC tracking will show marked differences when `ptc_field_geometry` is set to `straight`.

[1] There will be a small difference due to the fact that with a `straight` geometry tracking uses a coordinate system with the z-axis along the chord and with a `sector` geometry an integration step uses the curvilinear coordinate system with the z-axis along the arc of the bend. If the length of an integration step is made small, this difference will go to zero.

4.6 Capillary

A `capillary` element is a glass tube that is used to focus x-ray beams.

General `capillary` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Offsets, Pitches & Tilt	5.6
Capillary Wall	5.12	Reference energy	5.5
Custom Attributes	3.9	Tracking & transfer map	6
Description strings	5.3		

See §15.9 for a full list of element attributes along with a their units.

Attributes specific to a `capillary` element are:

```
critical_angle_factor = <Real>     ! Critical angle * Energy (rad * eV)
```

The critical angle above which photons striking the capillary surface are refracted into the capillary material scales as 1/Energy. The constant of critical angle * energy is given by the `critical_angle_factor`.

The inside wall of a capillary is defined using the same syntax used to define the chamber wall for other elements (§5.12).

The length of the capillary is a dependent variable and is given by the value of `s` of the last wall cross-section (§5.12.4).

4.7 Collimators: Ecollimator and Rcollimator

An ecollimator is a drift with elliptic collimation. An rcollimator is a drift with rectangular collimation.

Alternatively, for defining a collimator with an arbitrary shape, a mask element (§4.33) may be used.

General ecollimator and rcollimator attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Offsets, Pitches & Tilt	5.6
Chamber wall	5.12	Overlapping Fields	5.18
Custom Attributes	3.9	Reference energy	5.5
Description strings	5.3	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Field Maps	5.16
Is_on	5.14	Tracking & transfer map	6
Length	5.13		

Attributes specific to a capillary element are:

```
px_aperture_width2 = <real>  ! px aperture half width
px_aperture_center = <real>  ! px aperture center
py_aperture_width2 = <real>  ! py aperture half width
py_aperture_center = <real>  ! py aperture center
z_aperture_width2  = <real>  ! z aperture half width
z_aperture_center  = <real>  ! z aperture center
pz_aperture_width2 = <real>  ! pz aperture half width
pz_aperture_center = <real>  ! pz aperture center
```

Note: Collimators are the exception to the rule that the aperture is independent of any tilts. See §5.8 for more details. Additionally, the default setting of offset_moves_aperture is True for collimators (§5.8.1).

Besides the standard aperture settings 5.8 that can be used to limit x and y phase space coordinates, collimators can be used to limit the other four phase space coordinates as well. For rcollimator elements, particles are collimated if px_aperture_width2 is greater than zero and

```
px < px_aperture_center - px_aperture_width2  or
px > px_aperture_center + px_aperture_width2
```

with similar equations for py, z, and pz. For ecollimator elements, if px_aperture_width2 and py_aperture_width2 are both nonzero, particles are collimated if

```
((px - px_aperture_center) / px_aperture_width2)^2 +
      ((py - py_aperture_center) / py_aperture_width2)^2 < 1
```

If one or both of px_aperture_width2 or py_aperture_width2 are zero, the computation is the same as for an rcollimator. A similar situation occurs for z and pz.

Example:

```
d21: ecollimator, l = 4.5, x_limit = 0.09, y_limit = 0.05,
            px_aperture_width2 = 0.3, py_aperture_width = 0.1
```

4.8 Converter

A converter element represents a target (plate) onto which particles are slammed in order to generate particles of a different type. For example, a tungsten plate which is bombarded with electrons to generate positrons.

General custom attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Is_on	5.14
Chamber wall	5.12	Length	5.13
Custom Attributes	3.9	Offsets, pitches & tilt	5.6
Description strings	5.3	Reference energy	5.5
Integration settings	6.4	Superposition	8
		Tracking & transfer map	6

The attributes specific to an converter are

```
  distribution    = <Struct>    ! Outgoing particle distribution.
  pc_out_min      = <Real>      ! Minimum outgoing particle momentum (eV).
  pc_out_max      = <Real>      ! Maximum outgoing particle momentum (eV).
  angle_out_max   = <Real>      ! Maximum outgoing angle.
  species_out     = <SpeciesID> ! Output species.
  p0c             = <Real>      ! Output ref momentum.
  E_tot           = <Real>      ! Output ref energy. Dependent var (§5.1).
```

The species of the outgoing particles is specified by the species_out parameter (§3.11).

The converter must be the last element in a lattice branch (§2.2) except for possible fork, photon_fork or marker elements. A fork or photon_fork element (§4.23) after the converter is used to connect to the line containing the elements that come after the converter Example:

```
  parameter[particle] = electron
  parameter[geometry] = open
  to_after: fork, to_line = after_cvter
  cvter, species_out = positron, p0c = 3e6, distribution = ...
  pre_linac: line = (..., cvter, to_after)
  after_cvter: line = (...)                ! Everything after the converter.
  after_cvter[beta_a] = 27; after_cvter[beta_b] = 32
  use, pre_linac
```

The line up to the fork element, pre_linac, has the converter just before the fork element. The fork element, called to_after, connects to the line named after_cvter which contains all the elements after the converter. The reference particle and reference momentum for the after_cvter line is set to positron and 3e6 respectively to agree with the setting of species_out and p0c set in the converter element.

Since *Bmad* cannot calculate the appropriate Twiss and dispersion values after the converter, values must be set in the lattice file. Thus, in the above example, the starting beta function at the beginning of the after_cvter line is set to be $\beta_a = 27$ m and $\beta_b = 32$ m.

The p0c and E_tot attributes of the converter set the reference momentum or energy at the exit end of the converter. At least one of these attributes must be set. If both are set, E_tot is calculated to be consistent with p0c.

The distribution parameter of a converter element specifies the distribution of outgoing particles for a given converter thickness. Multiple distribution instances with differing thicknesses may be

present in an element. The actual thickness of the converter will be taken to be the element's, length `L` parameter. During tracking, the outgoing distribution will be computed by interpolating between the two distributions that bracket the actual thickness. The exception is when there is only one`distribution` present. In this case, the calculation will just use that distribution for the calculation independent of the element length. Example:

```
cvter: converter, ..., distribution = {
    material = tungsten,      ! Optional. Not used in tracking.
    thickness = 0.003,        ! Converter thickness for this distribution.
    sub_distribution = {...}, ! Distribution at one incoming momentum.
    sub_distribution = {...}, ! Distribution at another incoming momentum.
    ...                       ! etc.
  }
```

The `material` component is optional and is only for recording the converter material. Each `distribution` is made up of a number of `sub_distribution` components. Each one specified the outgoing distribution for a given incoming particle momentum. During tracking, interpolation is used to compute the distribution appropriate for an incoming particle with a given momentum. It is an error in the momentum of the incoming particle is outside the range of the momentums specified in the `sub_distributions`. A given `sub_distribution` will look like:

```
  sub_distribution = {
    pc_in = 3e8,            ! Incoming momentum*c (eV)>
    prob_pc_r = {...},      ! Momentum and radius probability table
    direction_out = {...},  ! Momentum orientation probability coefs
  }
```

A `sub_distribution` has three components: The `pc_in` component specifies the incoming particle momentum appropriate for the `sub_distribution`, the `prob_pc_r` component holds a two-dimensional table of the probability $P(p_{\rm out}, r)$ (Eq. (26.65)), and `direction_out` holds the coefficients for calculating the outgoing particle direction. `prob_pc_r` look like:

```
  prob_pc_r = {
    r_values = [0.0, 4.9e-5, 1.25e-4, ...],
    row = {pc_out = 1.55e6, prob = [0.0, 6.1e-6, 1.23e-5, ...]},
    row = {pc_out = 3.96e6, prob = [0.0, 1.1e-5, ...]},
    ...                   ! More rows
  }
```

A `prob_pc_r` has one `r_values` component and multiple `row` components. The `r_values` component is a vector of radius values for the columns of the probability table. The value for the first column is always zero and the radius values are strictly increasing. Each `row` component represents one row of the table. Each row has a momentum value `pc_out` in eV along with a `prob` component which is a vector of probability values. The length of a `prob` vector is always equal to the length of the `r_values` vector which is the number of columns in the table. The probability value of the first column is always zero which reflects the fact that there is vanishing area in an annulus of width dr as r tends to zero.

The `direction_out` component of `sub_distribution` look like:

```
  direction_out = {
    c_x = {...},
    alpha_x = {...},
    alpha_y = {...},
    beta = {...},
    dxds_min = {...},
    dxds_max = {...},
    dyds_max = {...}
  }
```

The `c_x`, `alpha_x`, `alpha_y`, and `beta`, components of `direction_out` give the coefficients for calculating c_x, α_x, α_y, and β respectively in Eq. (26.70). The other three components give, `dxds_min`, `dxds_max`, and `dyds_max` give the range for x' and y' over which Eq. (26.70) is valid. By symmetry, `dyds_min` will be equal to `-dydx_max`. The form of all these components is similar. For example:

```
dxds_min = {
  fit_1d_r = {pc_out = 1.5e+06, poly = [-2.48, -658.4, -2.26e5, 1.71e+8]},
  fit_1d_r = {pc_out = 3.9e+06, poly = [...]},
  ...,
  C = 2.99394,
  fit_2d_pc = {k = 1.96e-8, poly = [1.0, -4.10e-10, 3.7e-16, 2.77e-27]},
  fit_2d_r = {k = 4.2e-4, poly = [-4.50, 400.2, -108985, 9.18e+06]},
}
```

Here there are multiple `fit_1d_r` components, one for each fit Γ_i fit function (Eq. (26.72)). The `pc_out` sub-component of a `fit_1d_r` component gives the momentum p_i at which the fit function fits the data and the `poly` sub-component of `fit_1d_r` gives the polynomial coefficients needed for Eq. (26.72). The C, `fit_2d_pc` and `fit_2d_r` components are used for computing Ξ (Eq. (26.73)). The `k` sub-components of `fit_2d_pc` and `fit_2d_r` give k_p and k_r respectively in Eq. (26.73) and the `poly` sub-components of `fit_2d_pc` and `fit_2d_r` give the polynomial coefficients w_n and k_n respectively.

To calculate the distributions and output the appropriate `distribution` structures which then can be incorporated into a *Bmad* lattice, there is modeling code that is distributed with *Bmad*. Specifically, it is in the directory

```
$ACC_ROOT_DIR/util_programs/converter_element_modeling
```

[See your local *Bmad* Guru if you don't know how to find this directory.] There is documentation for running the program in this directory. The distribution modeling is based upon the `Geant` simulation toolkit for the simulation of the passage of particles through matter.

The mechanics of how *Bmad* generates outgoing particles is discussed in Sec. §26.9. In a tracking simulation, a single outgoing particle is generated for each incoming particle. All outgoing particles will be assigned a `weight` that represent how many actual outgoing particles a single actual incoming particle will generate. For example, if an actual incoming particle with a particular momentum would generate, on average, 0.42 particles, the outgoing particle in the simulation will have a weight of 0.42. To make simulations more efficient, the `pc_out_min`, `pc_out_max`, and `angle_out_max` parameters can be set to restrict the momentum and angle range of outgoing particles. If the outgoing particles are restricted in momentum or angle, the weight of the outgoing particles will be appropriately adjusted such that the weighted distribution of outgoing particles within the momentum and/or angle restricted range is independent of the whether or not there is are restrictions. A value of zero (the default) for any one of these parameters means that that parameter is ignored.

4.9 Crab_Cavity

A `crab_cavity` is an RF cavity that gives a z-dependent kick. This is useful in colliding beam machines, where there is a finite crossing angle at the interaction point, to rotate the beams near the IP.

General `crab_cavity` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Mag & Elec multipoles	5.15
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Reference energy	5.5
Description strings	5.3	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Field Maps	5.16
Is_on	5.14	Tracking & transfer map	6
Length	5.13		

See §15.12 for a full list of element attributes along with a their units.

The attributes specific to an `crab_cavity` are

```
gradient        = <Real>    ! Accelerating gradient (V/m).
phi0            = <Real>    ! Phase (rad/2π) of the reference particle with
                            !   respect to the RF. phi0 = 0 is on crest.
phi0_multipass  = <Real>    ! Phase (rad/2π) with respect to a multipass lord (§9).
rf_frequency    = <Real>    ! RF frequency (Hz).
rf_wavelength               ! RF wavelength. Dependent parameter.
harmon          = <Real>    ! Harmonic number
harmon_master   = <Logic>   ! Is harmon or rf_frequency the dependent variable
                            !   with reference energy changes?
voltage                     ! Cavity voltage. Dependent attribute (§5.1).
```

The Hamiltonian $H_{\rm crab}$ for a thin crab cavity in energy-time phase space coordinates (§17.5) is[Sun10]:

$$H_{\rm crab} = \frac{-r_q\, e\, V}{c P_0}\, x\, \sin(2\,\pi\,(f_{\rm rf}\, t + \phi_0)) \tag{4.14}$$

where x is the particle position, t is the time, r_q is the charge relative to the reference particle, V is the "effective" cavity voltage, ϕ_0 is a user settable phase, and $f_{\rm rf}$ is the RF frequency. The kick is:

$$\begin{aligned} \Delta p_x &= -\frac{\partial H_{\rm crab}}{\partial x} = \frac{r_q\, e\, V}{c\, P_0}\, \sin(2\,\pi\,(f_{\rm rf}\, t + \phi_0)) \\ \Delta E &= -c\, P_0\, \frac{\partial H_{\rm crab}}{\partial t} = 2\,\pi\, r_q\, e\, V\ f_{\rm rf}\, x\, \cos(2\,\pi\,(f_{\rm rf}\, t + \phi_0)) \end{aligned} \tag{4.15}$$

Note: The sign of $H_{\rm crab}$ used by Authors in the literature is not standardized. *Bmad* uses the convention such that a particle with the charge of the reference particle and with phase space z (which has opposite sign to t) and V positive will have a positive Δp_x.

In the above equations r_q is the relative charge between the reference particle (set by the `parameter[particle]` parameter in a lattice file) and the particle being tracked through the cavity. For example, if the reference particle and the tracked particle are the same, r_q is unity independent of the type of particle tracked.

The equations of motion can also be derived from analysis of a TM110 cavity mode for particles near the centerline[Kim11]. With this mode, the transverse kick is due to the magnetic field and the longitudinal kick is due to the electric field. Using this, the integrated electric and magnetic fields needed for spin tracking are:

$$\begin{aligned} \int B_y &= \frac{-V}{c}\ \sin(2\,\pi\,(f_{\rm rf}\,t+\phi_0)) \\ \int E_s &= 2\,\pi\,\beta\,V\ f_{\rm rf}\,x\ \cos(2\,\pi\,(f_{\rm rf}\,t+\phi_0)) \end{aligned} \tag{4.16}$$

where $\beta = v/c$ is the normalized speed of the particle.

The relationship between `harmon` and $f_{\rm rf}$ is `harmon` $= t_0 \cdot f_{\rm rf}$ where t_0 is the reference 1-turn transit time.

The `rf_wavelength` is computed to be $c/f_{\rm rf}$. Notice that this is independent of the velocity of the reference particle. In particular, It will only be true that `harmon` times `rf_wavelength` will equal the branch length if the reference particle travels at the speed of light.

Note: Multipoles 5.15 do not vary in time.

4.10 Crystal

A `crystal` element represents a crystal used for photon diffraction.

General `crystal` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Surface Properties	5.11
Custom Attributes	3.9	Symplectify	6.7
Description strings	5.3	Offsets, Pitches & Tilt	5.6
Reference energy	5.5	Tracking & transfer map	6
Reflection tables	5.10		

See §15.13 for a full list of element attributes along with a their units.

Attributes specific to a `crystal` element are:

```
b_param            = <Real>       ! b parameter for photons with the reference energy.
crystal_type       = <String>     ! Crystal material (§5.9) and reflection plane.
psi_angle          = <Real>       ! Rotation of H-vector about the surface normal.
thickness          = <Real>       ! Thickness of crystal for Laue diffraction.
ref_orbit_follows  = <which_beam> ! Reference orbit aligned with what outgoing beam?
graze_angle_in     = <Real>       ! Angle between incoming ref orbit and surface.
graze_angle_out    = <Real>       ! Angle between outgoing ref orbit and surface.
```

Dependent variables (§5.1) specific to a `crystal` element are:

```
alpha_angle                ! Angle of H-vector with respect to the surface normal.
bragg_angle                ! Nominal Bragg angle at the reference wave length.
bragg_angle_in             ! Incoming grazing angle for Bragg diffraction.
bragg_angle_out            ! Outgoing grazing angle for Bragg diffraction.
d_spacing                  ! Lattice plane spacing.
darwin_width_pi            ! Darwin width for pi polarized light (radians).
darwin_width_sigma         ! Darwin width for sigma polarized light (radians).
dbragg_angle_de            ! Variation of the Bragg angle with energy (radians/eV).
```

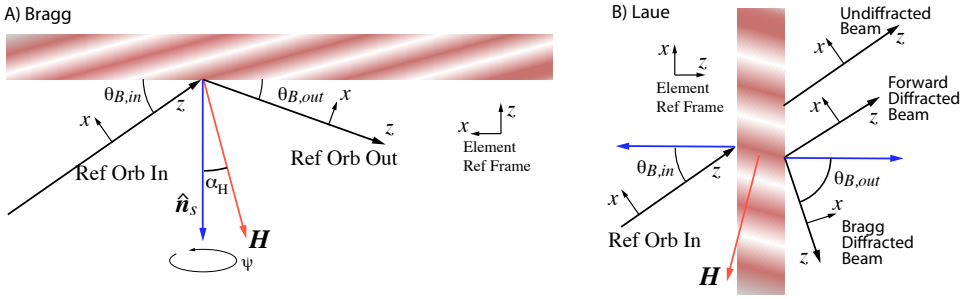

Figure 4.2: Crystal element geometry. A) Geometry for Bragg diffraction. The geometry shown is for `ref_tilt` = 0 (reference trajectory in the x-z plane). The angle α_H (`alpha_angle`) is the angle of the $\mathbf{H}$ vector with respect to the surface normal $\hat{\mathbf{n}}$. For ψ (`psi_angle`) zero, the incoming reference orbit, the outgoing reference orbit, $\hat{\mathbf{n}}$, and $\mathbf{H}$ are all coplanar. B) Geometry for Laue diffraction. In this case there are three outgoing beams: The Bragg diffracted beam, the forward diffracted beam, and the undiffracted beam.

```
l                          ! Length of reference orbit.
pendellosung_period_pi     ! Pendellosung period for pi polarized light.
pendellosung_period_sigma  ! Pendellosung period for sigma polarized light.
ref_wavelength             ! Reference wavelength (§5.5). Dependent attribute (§5.1).
ref_cap_gamma              ! Γ at the reference wavelength.
tilt_corr                  ! Tilt correction due to a finite psi_angle.
v_unitcell                 ! Unit cell volume.
```

The `crystal_type` attribute defines the crystal material and diffraction lattice plane. The syntax is `"ZZZ(ijk)"` where `ZZZ` is the material name and `ijk` are the Miller indices for the diffraction plane. Alternatively, for double digit indices, commas can be used to separate the indices. For example,

```
a_cryst: crystal, crystal_type = "Si(111)", b_param = -1, ...
b_cryst: crystal, crystal_type = "Si(2,4,10)", b_param = -1, ...
```

The atomic formula is case sensitive so, for example, `"SI(111)"` is not acceptable. The list of known crystal materials is given in §5.9. Given the `crystal_type`, the spacing between lattice planes (`d_spacing`), the unit cell volume (`v_unitcell`), and the structure factor[Bater64] values can be computed.

The `b_param` is the standard b asymmetry factor

$$b = \frac{\sin(\alpha_H + \theta_B)}{\sin(\alpha_H - \theta_B)} \tag{4.17}$$

where θ_B is the Bragg angle (`bragg_angle`)

$$\theta_B = \sin^{-1}\left(\frac{\lambda}{2\,d}\right) \tag{4.18}$$

and α_H (`alpha_angle`) is the angle of the reciprocal lattice $\mathbf{H}$ vector with respect to the surface normal as shown in Fig. 4.2A. If `b_param` is set to -1 then there is Bragg reflection and `alpha_H` is zero. If `b_param` is set to 1 then there is Laue diffraction again with `alpha_H` zero. With the orientation shown in Fig. 4.2A, `alpha_H` is positive.

The `ref_orbit_follows` parameter sets how the outgoing reference orbit is constructed. This is only relevant with Laue diffraction. The possible settings of this parameter are:

```
bragg_diffracted      ! Default
forward_diffracted
undiffracted
```

The geometry of this situation is shown in Fig. 4.2B. The reference orbit for the `undiffracted` beam is just a straight line extension of the incoming reference trajectory. This trajectory is that trajectory that photons whose energy is far from the Bragg condition (that is, far from the reference energy) will follow. The `forward_diffracted` reference orbit is parallel to the `undiffracted` trajectory and is the trajectory of the forward diffracted photons whose energy is the reference energy and whose incoming orbit is on the incoming reference trajectory. Finally, the `bragg_diffracted` reference orbit (the default) is the backward diffracted orbit.

Note: Changing the setting of `ref_orbit_follows` will change the reference orbit downstream of the crystal which, in turn, will change the placement all downstream elements.

The value of the element reference orbit length `l` is calculated by *Bmad*. `L` will be zero for Bragg diffraction. For Laue diffraction, `l` will depend upon the crystal `thickness` and the setting of `ref_orbit_follows`.

If `psi_angle` is zero, the incoming reference orbit, the outgoing reference orbit, $\hat{\mathbf{n}}$ and $\mathbf{H}$ are all coplanar. A non-zero `psi_angle` Rotates the $\mathbf{H}$ vector around the $+\hat{\mathbf{x}}$ axis of the `Element Reference Frame` (See Fig. 4.2A).

To keep the outgoing reference trajectory independent of the value of `psi_angle`, the crystal will be automatically tilted by the appropriate "tilt correction" `tilt_corr`. The calculation of `tilt_corr` is outlined in §27.4.2. `tilt_corr` will be zero if `psi_angle` is zero.

The reference trajectory for a Bragg `crystal` is that of a zero length bend (§16.2.3) and hence the length (`l`) parameter of a crystal is fixed at zero. If the `graze_angle_in` and `graze_angle_out` angles are zero (the default), the orientation of the reference trajectory with respect to the crystal surface is specified by the incoming Bragg angle `bragg_angle_in` ($\theta_{g,in}$) and outgoing Bragg angle `bragg_angle_out` ($\theta_{g,out}$) as shown in Fig. 4.2A. These angles are computed from the photon reference energy and the other crystal parameters such that a photon with the reference energy traveling along the reference trajectory will be in the center of the Darwin curve (§27.4). It is sometimes convenient to be able to specify the angles that the reference trajectory makes with respect to the crystal independent of the Bragg angles. To do this, set `graze_angle_in` and `graze_angle_out` to the desired angles.

Notice that due to refraction at the surface, the computed `bragg_angle` from Eq. (4.18) will deviate slightly from the average of `bragg_angle_in` and `bragg_angle_out`.

The reference trajectory in the global coordinate system (§16.2) is determined by the value of the `ref_tilt` parameter along with the value of `bragg_angle_in` + `bragg_angle_out`. These bragg angles take into account refraction so that the reference trajectory downstream of the crystal will be properly centered with respect to the reference photon. A positive `bragg_angle_in` + `bragg_angle_out` bends the reference trajectory in the same direction as a positive `g` for a bend element. The

A `crystal` may be offset and pitched (5.6). The incoming local reference coordinates are used for these misalignments.

When a crystal is bent (§5.11), the **H** vector is assumed follow the surface curvature. That is, it is assumed that the lattice planes are curved by the bending.

Example:

```
crystal_ele: crystal, crystal_type = "Si(111)", b_param = -1
```

The `darwin_width_sigma` and `darwin_width_pi` parameters are the computed Darwin width, in radians, for sigma and pi polarized light respectively. Here the Darwin width $d\theta_D$ is defined as the width at the $\eta = \pm 1$ points (cf. Batterman[Bater64] Eq (32))

$$d\theta_D = \frac{2\,\Gamma\,|P|\,\mathrm{Re}\big([F_H\,F_{\overline{H}}]^{1/2}\big)}{|b|^{1/2}\,\sin\theta_{tot}} \tag{4.19}$$

where

θ_{tot} `= bragg_angle_in + bragg_angle_out`

The `pendellosung_period_sigma` and `pendellosung_period_pi` are the pendellosung periods for Laue diffraction. If the crystal is set up for Bragg diffraction then the values for these parameters will be set to zero.

The `dbragg_angle_de` parameter is the variation in Bragg angle with respect to the photon energy and is given by the formula

$$\frac{d\theta_B}{dE} = -\frac{\lambda}{2\,d\,E\,\cos(\theta_B)} \tag{4.20}$$

See Section §13.5 for an example lattice that can be used to simulate a Rowland circle spectrometer.

4.11 Custom

A `custom` element is an element whose properties are defined outside of the standard *Bmad* subroutine library. That is, to use a custom element, some programmer must write the appropriate custom routines which are then linked with the *Bmad* subroutines into a program. *Bmad* will call the custom routines at the appropriate time to do tracking, transfer matrix calculations, etc. See the programmer who wrote the custom routines for more details! See §38.2 on how to write custom routines.

General `custom` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Is_on	5.14
Chamber wall	5.12	Length	5.13
Custom Attributes	3.9	Offsets, pitches & tilt	5.6
Description strings	5.3	Reference energy	5.5
Field Maps	5.16	Superposition	8
Fringe fields	5.21	Symplectify	6.7
Integration settings	6.4	Tracking & transfer map	6

See §15.14 for a full list of element attributes along with a their units.

As an alternative to defining a custom element, standard elements can be "customized" by setting one or more of the following attributes to `custom`:

```
tracking_method      §6.1
mat6_calc_method     §6.2
field_calc           §6.4
aperture_type        §5.8
```

As with a custom element, setting one of these attributes to `custom` necessitates the use of custom code to implement the corresponding calculation.

Attributes specific to a `custom` element are

```
val1, ..., val12 = <Real>  ! Custom values
delta_e_ref      = <Real>  ! Change in energy.
```

`delta_e_ref` is the energy gain of the *reference* particle between the starting edge of the element and the ending edge.

Example:

```
c1: custom, l = 3, val4 = 5.6, val12 = 0.9, descrip = "params.dat"
```

In this example the `descrip` string is being used to specify a file that contains parameters for the element.

4.12 Detector

A `detector` element is used to detect particles and X-rays. A `detector` is modeled as a grid of pixels which detect particles and x-rays impinging upon them.

General `detector` element attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Offsets, pitches & tilt	5.6
Chamber Wall	5.12	Reference energy	5.5
Custom Attributes	3.9	Superposition	8
Description strings	5.3	Tracking & transfer map	6
Detector Geometry	5.11.1		

See §15.15 for a full list of element attributes along with a their units.

Attributes specific to a `detector` element are:

```
pixel          = {...}   ! Define detector pixel grid.
```

The detector pixels are are arranged in a rectangular grid. The general syntax for defining a detector pixel grid is

```
pixel = {
    ix_bounds = (<ix_min>, <ix_max>),  ! Min/max index bounds in x-direction
    iy_bounds = (<iy_min>, <iy_max>),  ! Min/max index bounds in y-direction
    r0 = (<x0>, <y0>),                 ! (x,y) coordinates at grid origin
    dr = (<dx>, <dy>)                  ! Spacing between grid points.
        }
```

See Sec. §5.11.1 for an explanation of the various pixel parameters.

Example:

```
det: detector, pixel =
        {ix_bounds = (-4,5), iy_bounds = (-10,10), dr = (0.01, 0.01)}
```

This example defines a detector with 1 cm x 1 cm pixels.

The `aperture_type` (§5.8) parameter of a `detector` will default to `auto` which will set the aperture limits to define a rectangular aperture that just cover the area of the pixel grid.

A curved detector can be constructed by setting the appropriate surface curvature parameters (§5.11). It is assumed that any curvature is only in one dimension (x or y). This allows a straight forward mapping of the rectangular pixel grid onto the curved surface.

4.13 Diffraction_Plate

Note! Due to priorities being elsewhere, development of coherent photon tracking is moribund. As a consequence, tracking with diffraction_plate elements is currently disabled.

A `diffraction_plate` element is a flat surface oriented, more or less, transversely to a x-ray beam through which photon can travel. A `diffraction_plate` can be used, for example, to model a Fresnel zone plate or Young's double slits. A `diffraction_plate` element is used in places where diffraction effects must be taken into account. This is in contrast to setting an aperture attribute (§5.8) for other elements where diffraction effects are ignored.

A `diffraction_plate` element is similar to a `mask` (§4.33) element except that with a `mask` element coherent effects are ignored. Additionally, a `mask` element can be used with charged particles while a `diffraction_plate` cannot.

General `diffraction_plate` element attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Mask geometry	5.12
Description strings	5.3	Reference energy	5.5
Is_on	5.14	Tracking & transfer map	6

See §15.16 for a full list of element attributes along with a their units.

Attributes specific to a `diffraction_plate` element are:

```
mode                = <Type>   ! Reflection or transmission
field_scale_factor = <Real>   ! Factor to scale the photon field
ref_wavelength                ! Reference wavelength (§5.5). Dependent attribute (§5.1).
```

The `mode` switch sets whether X-rays are transmitted through the `diffraction_plate` or or reflected. Possible values for the `mode` switch are:

```
reflection
transmission        ! Default
```

The geometry of the plate, that is, where the openings (in transmission mode) or reflection regions are, is defined using the "wall" attribute. See (§5.12) for more details.

In transmission mode, a `diffraction_plate` is nominally orientated transversely to the beam. Like all other elements, the `diffraction_plate` can be reoriented using the element's offsets, pitches and tilt attributes (§5.6).

The `aperture_type` (§5.8) parameter of a `diffraction_plate` will default to `auto` which will set the aperture limits to define a rectangular aperture that just cover the clear area of the plate.

The `field_scale_factor`, if set to a non-zero value (zero is the default) will be used to scale the field of photons as they pass through the `diffraction_plate` element:

```
field -> field * field_scale_factor
```

Scaling is useful since the electric field of photons traveling through a `diffraction_plate` are renormalized (see Eqs. (27.10) and (27.11)). This can lead to large variation of the photon field and can, for example, make visual interpretation of plots of field verses longitudinal position difficult to interpret. `field_scale_factor` can be used to keep the field more or less constant.

A `diffraction_plate` that is "turned off" (`is_on` attribute set to False), does not diffract at all and transmits through all the light incident on it.

Example:

```
fresnel: diffraction_plate, wall = {...}
```

4.14 Drift

A `drift` element is a space free and clear of any fields.

General `drift` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Reference energy	5.5
Description strings	5.3	Symplectify	6.7
Length	5.13	Tracking & transfer map	6

See §15.17 for a full list of element attributes along with a their units.

Example:

```
d21: drift, l = 4.5
```

Note: If a chamber wall (§5.12) is needed for a field free space, use a `pipe` element instead of a `drift` [a wall for a drift is not allowed due to the way drifts are treated with superposition. That is, drifts "disappear" when superimposed upon. (§8)].

4.15 E_Gun

An e_gun element represents an electron gun and encompasses a region starting from the cathode were the electrons are generated. General e_gun attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Length	5.13
Chamber wall	5.12	Mag & Elec multipoles	5.15
Custom attributes	3.9	Offsets, pitches & tilt	5.6
Description strings	5.3	Overlapping Fields	5.18
Field autoscaling	5.19	Reference energy	5.5
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Field Maps	5.16
Is_on	5.14	Tracking & transfer map	6

See §15.20 for a full list of element attributes along with a their units.

The attributes specific to an e_gun are

```
gradient     = <Real>    ! Gradient.
phi0         = <Real>    ! Phase (rad/2π) of the reference particle with
                         !   respect to the RF. phi0 = 0 is on crest.
rf_frequency = <Real>    ! Frequency of the RF field.
voltage      = <Real>    ! Voltage. Dependent attribute (§5.1).
delta_E_ref  = <Real>    ! Change in reference energy.
```

The voltage is simply related to the gradient via the element length L:

```
voltage = gradient * L
```

If the voltage is set to a non-zero value, the length L must also be non-zero to keep the gradient finite. A particle with the charge as the reference particle will have a positive energy gain if the voltage and gradient are positive and vice versa.

The reference energy change from the entrance to the exit end of the element is set by delta_E_ref. The exception is if this change would make the total energy negative at the exit end. In this case, the change in reference energy is calculated via:

```
dE_ref = Voltage              ! For DC (rf_frequency = 0) guns.
       = Voltage * cos(phi0)  ! For AC guns.
```

In the first instance, the change in reference energy is set by delta_E_ref independent of the gun voltage. This is generally what is wanted since the reference energy can set the strengths of elements downstream from the gun and typically having the gun voltage affect downstream element strengths not desireable. *It is thus recommended that, unless there is a good reason not to, to set* ***delta_E_ref*** *appropriately.*

field_autoscale The voltage and gradient are scaled by field_autoscale and, if there is a finite rf_frequency, the phase of the frequency is shifted by phi0_autoscale as discussed in Section §5.19. Autoscaling can be toggled on/off by using the autoscale_phase and autoscale_amplitude toggles.

An e_gun may either be DC if the rf_frequency component is zero of AC if not. For an AC e_gun, the phase of the e_gun, The phase ϕ_{ref} is

ϕ_{ref} = phi0 + phi0_autoscale

Electrons generated at the cathode can have zero initial momentum and this presents a special problem (§5.5). As a result, the use of e_gun elements are restricted and they can only be used in a "linear"

(non-recirculating) lattice branch. Only one `e_gun` can be present in a lattice branch and, if it is present, it must be, except for possibly `marker` or `null_ele` elements, the first element in any branch.

Note: In order to be able to avoid problems with a zero reference momentum at the beginning of the `e_gun`, the reference momentum and energy associated with an `e_gun` element (§17.1) is calculated as outlined in Section §5.5. Additionally, the reference momentum at the exit end of the `e_gun`, that is `p0c`, must be non-zero. Thus, for example, if `p0c` is zero at the start of the lattice, the `e_gun` voltage must be non-zero.

Additionally, in order to be able to avoid problems with a zero reference momentum at the beginning of the `e_gun`, absolute time tracking (§26.1) is always used in an `e_gun` element independent of the setting of `bmad_com[absolute_time_tracking]` (§11.2).

Note: The default `tracking_method` (§6.1) setting for an `e_gun` is `time_runge_kutta` and the default `mat6_calc_method` is `tracking`.

In this example the field of an e_gun is given by a grid of field values (§5.16.4):

```
apex: e_gun, l = 0.23, field_calc = fieldmap, rf_frequency = 187e6,
              grid_field = call::apex_gun_grid.bmad, delta_E_ref = 1000
```

with the file `apex_gun_grid.bmad` being:

```
{
  m = 0, harmonic = 1,
  master_scale = voltage,
  geometry = rotationally_symmetric_rz,
  r0 = (0, 0),
  dr = (0.001, 0.001),
  pt(0,0) = ( (0, 0), (0, 0), (1, 0),  (0, 0), (0, 0), (0, 0)),
  pt(0,1) = ( (0, 0), (0, 0), (0.99, 0),  (0, 0), (0, 0), (0, 0)),
  ... }
```

Note: Before the parameter `delta_E_ref` was introduced as an `e_gun` parameter, the following `e_gun` parameters were defined:

```
gradient_err, gradient_tot
voltage_err, voltage_tot
phi0_err
```

If the *Bmad* parser encounters these parameters in a lattice file, an error will be thrown.

4.16 ELseparator

An elseparator is an electrostatic separator.

General elseparator attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Mag & Elec multipoles	5.15
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Overlapping Fields	5.18
Description strings	5.3	Reference energy	5.5
Fringe Fields	5.21	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Field Maps	5.16
Is_on	5.14	Tracking & transfer map	6
Length	5.13		

See §15.18 for a full list of element attributes along with a their units.

Attributes specific to an elseparator element are:

```
gap = <Real> ! Distance between electrodes
voltage      ! Voltage between electrodes. This is a settable dependent variable (§5.1).
e_field      ! Electric field. This is a settable dependent variable (§5.1).
```

For an elseparator, the kick for a positively charged particle, with the magnitude of the charge that is the same as that of the reference particle (set by parameter[particle] §10.1), is determined by hkick and vkick. The kick for a negatively charged particle is opposite this. The gap for an Elseparator is used to compute the voltage for a given kick

```
e_field (V/m) = sqrt(hkick^2 + vkick^2) * P0 * c_light / L
voltage (V) = e_field * gap
```

Specifying a e_field or voltage with no tilt results in a vertical kick.

Examples:

```
h_sep1: elsep, l = 4.5, hkick = 0.003, gap = 0.11
h_sep2: elsep, l = 4.5, e_field = 1e5, tilt = pi/2
```

4.17 EM_Field

An em_field element can contain general electro-magnetic (EM) fields. Both AC and DC fields are accommodated. General em_field attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Is_on	5.14
Chamber wall	5.12	Length	5.13
Custom Attributes	3.9	Offsets, pitches & tilt	5.6
Description strings	5.3	Reference energy	5.5
Field Maps	5.16	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Tracking & transfer map	6

See §15.19 for a full list of element attributes along with a their units.

Attributes specific to an em_field element are:

```
constant_ref_energy = <Logical> ! Is the reference energy constant? Default = True.
polarity    = <Real>  ! For scaling the field.
```

The polarity value is used to scale the magnetic field. By default, polarity has a value of 1.0. Example:

```
wig1: wiggler, l = 1.6, polarity = -1, cartesian_map = {...}
```

If the constant_ref_energy logical is set to True (the default), the reference energy (§17.1) at the exit end of the element is set equal to the entrance end reference energy. This is the same behavior for most other elements. If the constant_ref_energy logical is set to False, the reference energy at the exit end is calculated like it is in a lcavity or e_gun element.

Note: em_field elements will be created when elements are superimposed (§8) and there is no other suitable element class.

4.18 Feedback

A `feedback` element is a lord element with two types of slaves called the `input` slaves and `output` slaves. The feedback element gathers information about particle trajectories from the input slaves and uses this to either adjust beam trajectories in the output slaves and/or adjust parameters in the output slaves. A feedback element could be used, for example, to simulate RF feedback systems or beam position feedback, or cooling of a proton beam by a beam of electrons.

General `feedback` element attributes are:

Attribute Class	Section	*Attribute Class*	Section
Custom Attributes	3.9	Description strings	5.3

See §15.21 for a full list of element attributes along with a their units.

NOTE! 2024/3 The `feedback` element is currently under development so changes can be expected in the future.

Attributes specific to a `feedback` element are:

```
input_ele  = <list>   ! Lattice element(s) feedback element gets information from
output_ele = <list>   ! Lattice elements(s) where the feedback element can influence
                      !   particle trajectories or element parameters.
```

The `input_ele` parameter defines a list of lattice elements that specify, when tracking particles, where the `feedback` element will monitor particle trajectories. The `output_ele` parameter defines a list of lattice elements that specify the points at which the `feedback` element can either modify particle trajectories and/or modify lattice element parameters.

The `<list>` of lattice elements uses the standard *Bmad* name matching conventions as given in §3.6. If commas are used in the `<list>`, the list must be enclosed in curly brackets `{...}` to avoid ambiguities. Curly brackets are optional when commas are not used. Examples:

```
fff: feedback, input_ele = bpm5,
                output_ele = kicker3      ! Single input and output elements.
fe2: feedback, input_ele = {bpm5},
                output_ele = rfcavity::*  ! Output to all RF cavities.
g7: feedback, input_ele = {mon1, mon2},
                output_ele = type::IRQ    ! Match to element's type field.
```

The set of `input_ele` and `output_ele` elements will be `minor_slave` slaves of the `control_lord` `feedback` element. Like other lord elements (`groups`, `overlays`, etc.), particles are never tracked through a `feedback` element itself.

Since feedback systems vary greatly in how they work, a generic feedback element is not currently planned (but could be in the future as more experience is gained developing feedback simulation code). This being the case, a program must be specifically setup to handle feedback elements. In particular, feedback elements in a lattice will not affect any calculation when using the `Tao` program (but `Tao` can still be used to inspect feedback elements and their slaves).

Currently, the only program that handles `feedback` elements is the `e_cooling` program that is located in the `bsim` directory of a `Bmad` Release (and this program is currently under development and not ready for doing simulations). A feedback aware program will handle the task of feedback related parameter setup so the program documentation should be consulted for specifics.

4.19 Fiducial

A `fiducial` element is used to fix the position and orientation of the reference orbit within the global coordinate system at the location of the `fiducial` element. A `fiducial` element will affect the global floor coordinates (§16.2) of elements both upstream and downstream of the fiducial element.

Other elements that are used to shift the lattice in the global coordinate frame are `floor_shift` (§4.21) and `patch` (§4.41).

General `fiducial` element attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Reference energy	5.5
Custom Attributes	3.9	Superposition	8
Description strings	5.3	Tracking & transfer map	6

See §15.22 for a full list of element attributes along with a their units.

Attributes specific to a `fiducial` elements are:

```
origin_ele        = <Name>     ! Reference element.
origin_ele_ref_pt = <location> ! Reference pt on reference ele.
dx_origin         = <Real>     ! x-position offset
dy_origin         = <Real>     ! y-position offset
dz_origin         = <Real>     ! z-position offset
dtheta_origin     = <Real>     ! orientation angle offset.
dphi_origin       = <Real>     ! orientation angle offset.
dpsi_origin       = <Real>     ! orientation angle offset.
```

For tracking purposes, the `fiducial` element is considered to be a zero length marker. That is, the transfer map through a `fiducial` element is the unit map.

A `fiducial` element sets the global floor coordinates (§16.2) of itself and of the elements, both upstream and downstream, around it. This can be thought of as a two step process. The first step is to determine the global coordinates of the `fiducial` element itself, and the second step is to shift the coordinates of the elements around it. That is, shifting the position of a `fiducial` element shifts the lattice elements around it as one solid body.

The floor coordinates of the `fiducial` element are determined starting with an `origin_ele` element. If `origin_ele` is not specified, the origin of the global coordinates (§16.2 is used. If the `origin_ele` has a finite length, the reference point may be chosen using the `origin_ele_ref_pt` attribute which may be set to one of

```
entrance_end
center         ! Default
exit_end
```

Once the origin reference position is determined, the reference position of the `fiducial` element is calculated using the offset attributes

```
[dx_origin,  dy_origin,  dz_origin]
[dtheta_origin,  dphi_origin,  dpsi_origin]
```

The transformation between origin and fiducial positions is given in §16.2.4.

Once the position of the `fiducial` element is calculated, all elements of the lattice branch the `fiducial` element is contained in, *both* the upstream and downstream elements, are shifted so that everything

is consistent. That is, the `fiducial` element orients the entire lattice branch. The exception here is that if there are `flexible patch` elements (§4.41) in the lattice branch, the `fiducial` element will only determine the positions up to the `flexible patch` element.

Example: A lattice branch with elements 0 through 103 has a `fiducial` element at position 34 and a `flexible patch` at position 67. In this case the `fiducial` element will determine the reference orbit for elements 0 through 66.

Rules:

- If an `origin_ele` is specified, the position of this element must to calculated before the position of the `fiducial` element is calculated (§16.1.1). This means, the `origin_ele` must be in a prior lattice branch from the branch the `fiducial` element is in or the `origin_ele` in the same branch as the `fiducial` element but is positioned upstream from the `fiducial` element and there is a `flexible patch` in between the two elements.
- If a `fiducial` element affects the position of element 0 in the lattice branch (that is, there are no flexible `patch` elements in between), any positioning of element 0 via `beginning` or `line parameter` statements (§10.4) are ignored.
- `Fiducial` elements must not over constrain the lattice geometry. For example, two `fiducial` elements may not appear in the same lattice branch unless separated by a `flexible patch`.

  Another example is that if there are no flexible `patch` elements in the lattice, and if branch `A` has a `branch` element connecting to branch `B`, the geometry of branch `A` will be calculated first and the geometry of branch `B` can then be calculated from the known coordinates of the `fork` element. If branch `B` contains a `fiducial` element then this is an error since the coordinate calculation never backtracks to recalculate the coordinates of the elements of a branch once the calculation has finished with that branch.

Example:

```
f1: fiducial, origin_ele = mark1, x_offset = 0.04
```

See §13.4 for an example where a fiducial element is used to position the second ring in a dual ring colliding beam machine.

4.20 Fixer

This element is in construction. Status: Fixer elements can be used but testing is not yet complete.

A `fixer` element is used to set the particle orbit and Twiss parameters at any point in an open geometry lattice branch. The orbit and Twiss parameters can then be propagated from the fixer element both forward and backwards to the ends of the branch. `Fixed` elements are convenient, for example, when it is desired to design a beam line with Twiss and/or orbit parameters that are fixed at some point that is not the beginning of the line. Another example is when the orbit or Twiss parameters have been measured at some point in a beam line and a calculation of the orbit or Twiss parameters at other places in the line is desired consistent with the measured values.

Note: For the purposes of this discussion, the term "Twiss" parameters include coupling and dispersion parameters and "orbit" includes both phase space and spin coordinates.

General `floor_shift` element attributes are:

Attribute Class	Section	*Attribute Class*	Section
Custom Attributes	3.9	Superposition	8
Description strings	5.3	Tracking & transfer map	6
Reference energy	5.5		

See §15.24 for a full list of element attributes along with a their units.

Attributes specific to a `fixer` elements are:

```
spin_x_stored, spin_y_stored, spin_z_stored     ! Particle spin

x_stored, px_stored
y_stored, py_stored                             ! Particle phase space coordinates
z_stored, pz_stored

beta_a_stored, alpha_a_stored, phi_a_stored     ! a-mode Twiss
dbeta_dpz_a_stored, dalpha_dpz_a_stored

beta_b_stored, alpha_b_stored, phi_b_stored     ! b-mode Twiss
dbeta_dpz_b_stored, dalpha_dpz_b_stored

eta_x_stored, etap_x_stored                     ! Horizontal dispersion
deta_dpz_x_stored, detap_dpz_x_stored

eta_y_stored, etap_y_stored                     ! Vertical dispersion
deta_dpz_y_stored, detap_dpz_y_stored

mode_flip_stored                                ! Logical: Normal modes flipped?
cmat_11_stored, cmat_12_stored                  ! Coupling
cmat_21_stored, cmat_22_stored
```

A `fixer` element has zero length. For tracking particles, a `fixer` element acts like a `marker` element. That is, there is no change in phase space coordinates when a particle is tracked through a `fixer` element.

If the geometry of a branch is `closed` (§10.1), `fixer` elements have no effect. Thus, in the following discussion, it is assumed that the branch geometry is set to `open`. Additionally, for this discussion, the

term "`fixer`" element will include `beginning_ele` elements (§4.4) since, for the purposes of setting orbit and/or Twiss parameters, `beginning_ele` elements essentially act the same as a fixer element.

A fixer element is "active" if the value of its `is_on` parameter (§5.14) is True. One and only one fixer element in a branch can be active at any given time. The *Bmad* bookkeeping routines will ensure that this rule is enforced and turning on a given fixer will result in all others being turned off. Turning off the active fixer will result in the `beginning_ele` element being turned on. The default for a fixer element defined in a lattice file is to be non-active except for the `beginning_ele` which defaults to being active. When a lattice is read in, for a given branch, the last fixer that is set active will be the active fixer and if no `fixer` elements are set active, the `beginning_ele` will be active.

A fixer elements has two classes of Twiss and orbit parameters. There are the "stored" parameters listed above. These parameters may be set in a lattice file. The other class are the "actual" parameters which cannot be set by the User. For a `fixer` element that is active, the actual parameters are set equal to the corresponding stored parameters and these actual parameters are then used by *Bmad* to propagate the Twiss and orbit backwards and forwards through the rest of the branch. For fixer elements that are not active, the actual Twiss and orbit is calculated in the standard fashion by propagating the Twiss and orbit from the active fixer. The stored Twiss and orbit parameter values in a non-active fixer do not affect the Twiss and orbit calculation.

Example:

```
f2: fixer, beta_a_stored = 11.37, beta_b_stored = 34.72, is_on = T
```

4.21 Floor_Shift

A `floor_shift` element shifts the reference orbit in the global coordinate system without affecting particle tracking. That is, in terms of tracking, a `floor_shift` element is equivalent to a `marker` (§4.32) element.

Also see `patch` (§4.41) and `fiducial` (§4.19) elements.

General `floor_shift` element attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Reference energy	5.5
Custom Attributes	3.9	Superposition	8
Description strings	5.3	Tracking & transfer map	6
Length	5.13		

See §15.24 for a full list of element attributes along with a their units.

Attributes specific to a `floor_shift` elements are:

```
l                = <Real>     ! Length
x_offset         = <Real>     ! x offset from origin point.
y_offset         = <Real>     ! y offset from origin point.
z_offset         = <Real>     ! z offset from origin point.
x_pitch          = <Real>     ! rotation of the reference coords.
y_pitch          = <Real>     ! rotation of the reference coords.
tilt             = <Real>     ! rotation of the reference coords.
origin_ele       = <Name>     ! Reference element.
origin_ele_ref_pt = <location> ! Reference pt on the reference ele.
```

The `floor_shift` element sets the reference orbit at the exit end of the `floor_shift` element as follows: Start with the reference orbit at the `origin_ele` reference point (see below). This coordinate system is shifted using the offset, pitch and tilt parameters of the `floor_shift` element. The shifted coordinate system is used as the coordinate system at the exit end of the `floor_shift` element. The reference position transformation through a `floor_shift` element is given in Section §16.2.4. In this respect, the `floor_shift` element is similar to the `fiducial` element. The difference being that the `fiducial` element affects the global floor coordinates of elements both upstream and downstream of the `fiducial` element while a `floor_shift` element only affects the floor position of elements downstream from it.

Like a `fiducial` element, the transfer map through a `floor_shift` element will be the unit map. That is, the phase space coordinates of a particle will not change when tracking through a `floor_shift` element.

The `l` attribute can be used to adjust the longitudinal s position.

The `floor_shift` element can be used, for example, to restore the correct global geometry when a section of the lattice is represented by, say, a `taylor` type element.

If an `origin_ele` is not specified, the default `origin_ele` is the lattice element before the `floor_shift` element. If an `origin_ele` is specified, *Bmad* needs to be able to calculate the position of this element before the position of the `fiducial` element is calculated. See the discussion of the `origin_ele` for `fiducial` elements (§4.19). Notice that if the `origin_ele` is specified, and is different from the element upstream from the `floor_shift` element, the coordinates at the exit end of the `floor_shift` element is independent of the coordinates of the upstream element.

If the `origin_ele` has a finite length, the reference point may be chosen using the `origin_ele_ref_pt` attribute which may be set to one of

```
entrance_end
center
exit_end          ! Default
```

PTC does not have an analogous element for the `Floor_shift` element. When converting to PTC, a `floor_shift` element will be treated as a `marker` element.

Example:

```
floor: floor_shift, z_offset = 3.2
```

This offsets the element after the `floor_shift` 3.2 meters from the previous element.

4.22 Foil

A `foil` element represents a planar sheet of material which can strip electrons from a particle. In conjunction, there will be scattering of the particle trajectory as well as an associated energy loss.

General `foil` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Reference energy	5.5
Description strings	5.3	Superposition	8
Integration settings	6.4	Tracking & transfer map	6
Is_on	5.14		

See §15.58 for a full list of element attributes along with a their units.

Attributes specific to a `foil` element are:

```
material_type       = <String>  ! Foil material.
thickness           = <Real>    ! Material thickness (m).
density             = <Real>    ! Input material density (kg/m^3).
density_used                    ! Density value used in tracking (kg/m^3).
radiation_length    = <Vector>  ! Input material radiation length (m).
radiation_length_used           ! Radiation length used in tracking (m).
area_density        = <Vector>  ! Input material area density (kg/m^2).
area_density_used               ! Area density used in tracking (kg/m^2).
F_factor            = <Real>    ! lynch_dahl scattering F factor. Default: 0.98.
final_charge        = <Integer> ! Final charge state
scatter_test        = <Logic>   ! For testing scattering. Default: False.
scatter_method      = <Switch>  ! Scattering algorithm. Default: highland.
dthickness_dx       = <Real>    ! Wedge slope when the foil is wedge shaped.
x1_edge             = <Real>    ! Foil edge in the x-direction. Default: -99 m.
x2_edge             = <Real>    ! Foil edge in the x-direction. Default:  99 m.
y1_edge             = <Real>    ! Foil edge in the y-direction. Default: -99 m.
y2_edge             = <Real>    ! Foil edge in the y-direction. Default:  99 m.
```

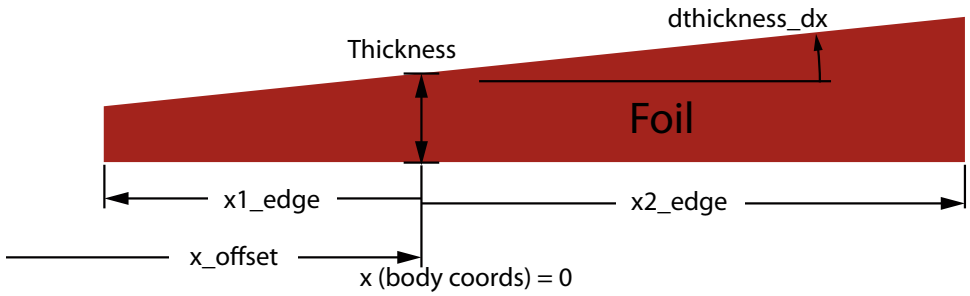

Figure 4.3: Foil geometry. Parameters `thickness`, `dthickness_dx`, `x1_edge`, `x2_edge`, `y1_edge`, and `y2_edge` determine the foil geometry in body coordinates (see Fig. 16.1). For the drawing, `x1_edge` is negative (the associated edge is to the left of $x = 0$) and the `x2_edge` is positive. Orientation parameters like `x_offset` (§5.6) orient the foil with respect to laboratory coordinates but do not change the foil shape.

Scattering is simulated to be Gaussian distributed with a sigma calculated in one of two methods. The two scattering algorithms are given in section §26.12.1. Which algorithm is used is determined by the `scatter_method` parameter which can be set to:

```
highland        ! Default
lynch_dahl
off             ! No scattering
```

Additionally, the `scatter_test` logical may be used for testing. If set to True (default is False), the random numbers used in Eq. (26.83) are set to 1.

Energy loss is calculated using the Bethe-Bloch formula as discussed in section §26.12.2.

The `material_type` is the type of material which can be elemental or a compound material.

The radiation length used in the scattering calculation is given by the `radiation_length_used` parameter. For compound materials, this parameter is a vector with each value of the vector being the radiation length of the corresponding component. `Radiation_length_used` cannot be set directly. Rather, if the `radiation_length` parameter is set non-zero, the value (or values for a compound material) will be transferred to `radiation_length_used`. If `radiation_length` is zero (the default), the value of `radiation_length_used` will be set by *Bmad* using measured values from the published literature.

Similarly, the `area_density_used` (density of the material per unit of surface area) value (or values for a compound material) needed for the calculation is not set directly but is set in one of two ways depending upon if the material `thickness` is non-zero or not. If `thickness` is non-zero, `area_density_used` is set by the product of `thickness` and `density_used` while the value of `density_used` is set by *Bmad* to be either the value `density` if the `density` is non-zero or by the measured density of the material as given in the published literature. If `thickness` is zero (the default), the value of `area_density_used` is set equal to the value of `area_density`. Note that the value of `density_used` is only used to set `area_density_used` when `thickness` is non-zero and otherwise does not affect the tracking calculation.

Example:

```
f1: foil, thickness = 1e-2, material_type = "B4C",
        density = (2e3, 1e3), radiation_length = (5.49, 4.26)
f1: foil, thickness = 1e-2, material_type = "B4C",
        density = 9e3, radiation_length = (5.49, 4.26)  ! Same as above
```

Here, since `material_type` is set to `B4C`, there are two components: Boron and Carbon. If `radiation_length` is used, it must have the same number of values as the number of material components. If `density`, or `area_density` is used they either must have the same number of values as the number of material components or they can have a single value. If there is a single value, this value represents the sum of the individual component values.

Values are set in order so, in the above example, the Carbon component has a density of 1e3 and a radiation length of 4.26.[2] In the case where there is a component material, either `density` or `area_density` needs to be set since *Bmad* is not able to calculate appropriate values in this case. When there is only one component, the parentheses may be omitted. Example: Example:

```
stripper: foil, material_type = "Cu", thickness = 0.127, &
                                radiation_length = 12.3, x1_edge = -0.3
```

In terms of element placement, The length of a `foil` element (Eq. (16.7)) is considered to be zero. This is similar to the `beambeam` element which is also considered to have zero length but the interaction occurs over a finite length.

The 6x6 transfer matrix of a `foil` element is the unit matrix. That is, scattering does not affect the transfer matrix.

[2] From a computational standpoint it does not matter which parameter is associated with which component.

Particles going through the `foil` are stripped to have a final charge given by `final_charge`. The default is to fully strip the particle except if the particle has no electrons to strip, there will be no change in charge state.

The foil has a rectangular shape and particles will only be considered to have hit the foil if:

```
x1_edge < x_particle < x2_edge  and
y1_edge < y_particle < y2_edge
```

where (x_particle, y_particle) are the coordinates of the particle in the element body (not the laboratory) coordinate system (Fig. 16.1). See Fig. 4.3. Particles that do not hit the foil pass through the element without a change in charge nor a change in trajectory. The default for `x1_edge` and `y1_edge` is -99 meters and for `x2_edge` and `y2_edge` the default is +99 meters.

The `dthickness_dx` parameter can be used to get a varying foil thickness. The thickness t at a point (in body coordinates) (x, y) on the foil will be

$$t = t_0 + x \frac{dt}{dx} \tag{4.21}$$

where t_0 is the thickness given by the `thickness` parameter and dt/dx is given by the `dthickness_dx` parameter. To orient the wedge in the transverse plane, use the `tilt` orientation parameter (§5.6). If `dthickness_dx` is non-zero, the `area_density` and `thickness` parameters are defined to be the area density and thickness at $(x, y) = (0, 0)$. If `dthickness_dx` is non-zero, and if `area_density` is set (as opposed to setting the `density`), then the `thickness` must be non-zero since otherwise the calculation of the area density at the point where a particle passes through the foil is singular.

If a `foil` element is part of a lattice branch with a closed geometry, the closed orbit calculation will temporarily set the `scatter` parameter to false since scattering is a random process and the closed orbit is not well defined in the presence of any random processes (similarly, radiation fluctuations are also turned off for the closed orbit calculation).

4.23 Fork and Photon_Fork

A `fork` or `photon_fork` element marks the start of an alternative `branch` for the beam (or X-rays or other particles generated by the beam) to follow.

Collectively `fork` and `photon_fork` elements are called forking elements. An example geometry is shown in Fig. 4.4. The `branch` containing a forking element is called the "`base branch`". The `branch` that the forking element points to is called the "`forked-to branch`".

The only difference between `fork` and `photon_fork` is that the default particle type for the forked-to branch forked from a `fork` element is the same particle type as the base branch. The default particle type for the forked-to branch from a `photon_fork` element is a photon. The actual particle associated with a branch can be set by setting the `particle` attribute of the forking element.

General `fork` and `photon_fork` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Length	5.13
Chamber wall	5.12	Reference energy	5.5
Custom Attributes	3.9	Superposition	8
Description strings	5.3	Tracking & transfer map	6
Is_on	5.14		

See §15.26 for a full list of element attributes along with a their units.

Attributes specific to `fork` and `photon_fork` elements are:

```
direction    = <+/- 1>      ! Particles are entering or leaving?
to_line      = <LineName>   ! What line to fork to.
to_element   = <ElementID>  ! What element to attach to in the line being forked to.
new_branch   = <T/F>        ! Make a new branch from the to_line? Default = True.
```

Branch lines can themselves have forking elements. A branch line always starts out tangential to the line it is branching from. A `patch` element (§4.41) can be used to reorient the reference orbit as needed. Example:

```
from_line: line = (... A, PB, B, ...)  ! Defines base branch
pb: photon_fork, to_line = x_line
x_line: line = (X_PATCH, X1, X2, ...)     ! Defines forked-to branch
x_line[p0c] = 1e3                         ! Photon reference momentum
x_patch: patch, x_offset = 0.01
use, from_line
```

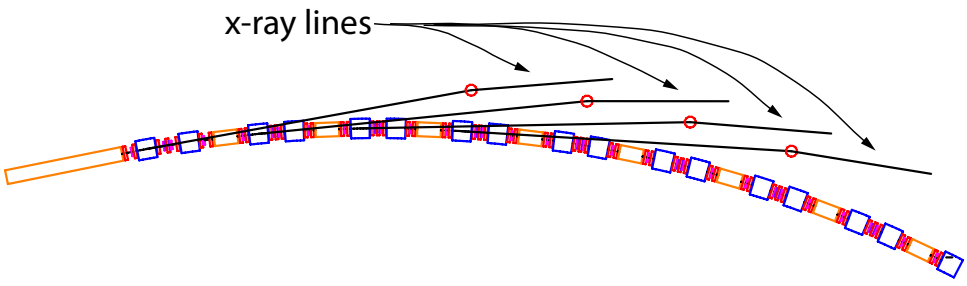

Figure 4.4: Example use of `photon_fork` elements showing four X-ray lines (branches) attached to a machine.

In this example, a photon generated at the fork element PB with $x = 0$ with respect to the from_line reference orbit through PB will, when transferred to the x_line, and propagated through X_PATCH, have an initial value for x of -0.01.

Forking elements have zero length and, like marker elements, the position of a particle tracked through a forking element does not change.

Forking elements do not have orientational attributes like x_pitch and tilt (5.6). If the orientation of the forked-to branch needs to be modified, this can be accomplished using a patch element at the beginning of the line.

The is_on attribute, while provided for use by a program, is ignored by *Bmad* proper.

If the reference orbit needs to be shifted when forking from one ring to another ring, a patch can be placed in a separate "transfer" line to isolate it from the branches defining the rings. Example:

```
ring1: line = (... A, F1, B, ...)     ! First ring
x_line: line = (X_F1, X_PATCH, X_F2)  ! "Transfer" line
ring2: line = (... C, F2, D, ...)     ! Second ring
use, ring1

f1: fork, to_line = x_line
f2: fork, to_line = x_line, direction = -1
x_patch: patch, x_offset = ...
x_f1: fork, to_line = ring1, to_element = f1, direction = -1
x_f2: fork, to_line = ring2, to_element = f2
```

Here the fork F1 in ring1 forks to x_line which in turn forks to ring2.

The above example also illustrates how to connect machines for particles going in the reverse direction. In this case, rather than using a single fork element to connect lines, pairs of fork elements are used. Ring2 has a fork element f2 that points back through x_line and then to ring1 via the x_f1 fork. Notice that both f2 and x_f2 have their direction attribute set to -1 to indicate that the fork is appropriate for particles propagating in the -s direction. Additionally, since f2 has direction set to -1, it will, by default, connect to the downstream end of the x_line. The default setting of direction is 1.

It is important to note that the setting of direction does not change the placement of elements in the forked line. That is, the global position (§16.2) of any element is unaffected by the setting of direction. To shift the global position of a forked line, patch elements must be used. In fact, the direction parameter is merely an indicator to a program on how to treat particle propagation. The direction parameter is not used in any calculation done by *Bmad*.

The to_element attribute for a forking element is used to designate the element of the forked-to branch that the forking element connects to. To keep things conceptually simple, the to_element must be a "marker-like" element which has zero length and unit transfer matrix. Possible to_element types are:

```
beginning_ele
fiducial
fork and photon_fork
marker
```

When the to_element is not specified, the default is to connect to the beginning of the forked-to branch if direction is 1 and to connect to the end of the downstream branch if direction is -1. In this case, there is never a problem connecting to the beginning of the forked-to branch since all branches have a beginning_ele element at the beginning. When connecting to the end of the forked-to branch the last element in the forked-to branch must be a marker-like element. Note that, by default, a marker element is placed at the end of all branches (§7.1)

The default reference particle type of a branch line will be a `photon` is using a `photon_fork` or will be the same type of particle as the base branch if a `fork` element is used. If the reference particle of a branch line is different from the reference particle in the base branch, the reference energy (or reference momentum) of a forked-to branch line needs to be set using line parameter statements (§10.4). If the reference particle of a branch line is the same as the reference particle in the base branch, the reference energy will default to the reference energy of the base branch if the reference energy is not set for the branch.

Example showing an injection line branching to a ring which, in turn, branches to two x-ray lines:

```
inj: line = (..., br_ele, ...)            ! Define the injection line
use, inj                                  ! Injection line is the root
br_ele: fork, to_line = ring              ! Fork element to ring
ring: line = (..., x_br, ..., x_br, ...)  ! Define the ring
ring[E_tot] = 1.7e9                       ! Ring ref energy.
x_br: photon_fork, to_line = x_line       ! Fork element to x-ray line
x_line: line = (...)                      ! Define the x-ray line
x_line[E_tot] = 1e3
```

The `new_branch` attribute is, by default, `True` which means that the lattice branch created out of the `to_line` line is distinct from other lattice branches of the same name. Thus, in the above example, the two lattice branches made from the `x_line` will be distinct. If `new_branch` is set to `False`, a new lattice branch will not be created if a lattice branch created from the same line already exists. This is useful, for example, when a chicane line branches off from the main line and then branches back to it.

When a lattice is expanded (§3.26), the branches defined by the `use` statement (§7.7) are searched for fork elements that branch to new forked-to branches. If found, the appropriate branches are instantiated and the process repeated until there are no more branches to be instantiated. This process does *not* go in reverse. That is, the lines defined in a lattice file are not searched for fork elements that have forked-to instantiated branches. For example, if, in the above example, the use statement was:

```
use, x_line
```

then only the `x_line` would be instantiated and the lines `inj` and `ring` would be ignored.

If the forked-to branch and base branch both have the same reference particle, and if the element forked into is the beginning element, the reference energy and momentum of the forked-to branch will be set to the reference energy and momentum at the fork element. In this case, neither the reference energy nor reference momentum of the forked-to branch should be set. If it is desired to have the reference energy/momentum of the forked-to branch different from what is inherited from the fork element, a patch element (§4.41) can be used at the beginning of the forked-to branch. In all other cases, where either the two branches have different reference particles or the fork connects to something other than the beginning element, there is no energy/momentum inheritance and either the reference energy or reference momentum of the forked-to branch must be set.

How to analyze a lattice with multiple branches can be somewhat complex and will vary from program to program. For example, some programs will simply ignore everything except the root branch. Hopefully any program documentation will clarify the matter.

4.24 Girder

A girder is a support structure that orients the elements that are attached to it in space. A girder can be used to simulate any rigid support structure and there are no restrictions on how the lattice elements that are supported are oriented with respect to one another. Thus, for example, optical tables can be simulated.

General girder attributes are:

Attribute Class	Section	*Attribute Class*	Section
Custom Attributes	3.9	Length	5.13
Description strings	5.3	Offsets, pitches & tilt	5.6
Is_on	5.14		

See §15.28 for a full list of element attributes along with a their units.

Attributes specific to a girder are:

```
girder = {<List>}   ! List of elements on the Girder
origin_ele        = <Name>     ! Reference element.
origin_ele_ref_pt = <location> ! Reference pt on reference ele.
dx_origin         = <Real>     ! x-position offset
dy_origin         = <Real>     ! y-position offset
dz_origin         = <Real>     ! z-position offset
dtheta_origin     = <Real>     ! orientation angle offset.
dphi_origin       = <Real>     ! orientation angle offset.
dpsi_origin       = <Real>     ! orientation angle offset.
l                 ! Girder "Length" (5.13). Dependent attribute (§5.1).
```

Attributes specific to girder elements are:

A simple example of a girder is shown in Fig. 4.5. Here a girder supports three elements labeled A, B, and C where B is a bend so the geometry is nonlinear. Such a girder may specified in the lattice file like:

```
g1: girder = {A, B, C}
```

The girder statement can take one of two forms:

```
<element_name>: GIRDER = {<ele1>, <ele2>, ..., <eleN>}, ...
```

or

```
<element_name>: GIRDER = {<ele_start>:<ele_end>}, ...
```

With the first form, a girder element will be created for each section of the lattice where there is a "consecutive" sequence of "slave" elements <ele1> through <eleN>. This section of the lattice from <ele1> through <eleN> is called the "girder support region". "Consecutive" here means there are no other elements in the girder support region except for possibly drift and/or marker elements. Drift elements cannot be controlled by a girder[3] but may appear in the girder slave list. If a drift does appear in the slave list, drift elements will not be ignored when determining if elements are consecutive. Note: If a drift-like element is desired to be supported by a girder, use a pipe element instead. Marker elements present in a girder support region, but not mentioned in the girder slave list, are simply ignored.

The second form of a girder statement specifies the first and last elements in the sequence of elements to be supported. Everything in between except drift elements will be supported by the girder.

Wild card characters (§3.7) can be used in any element name in the girder slave list. Additionally, beam line names (§7.2) can be used. In this case, any drift elements within a beam line will be ignored.

[3] This policy was created to avoid the problem where the superposition of marker elements on top of drifts would prevent girder formation.

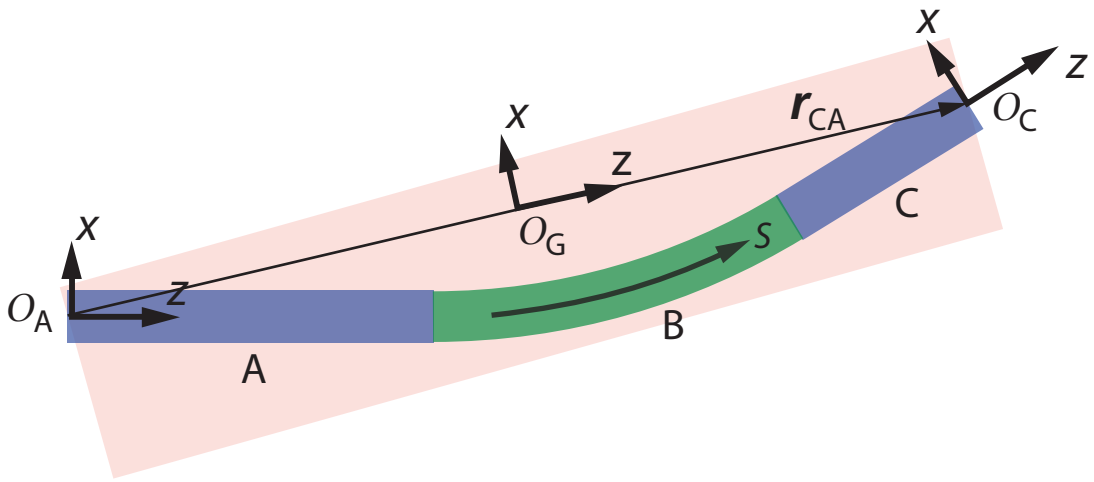

Figure 4.5: Girder supporting three elements labeled `A`, `B`, and `C`. $\mathcal{O}_A$ is the reference frame at the upstream end of element `A` (§16.1.3), $\mathcal{O}_C$ is the reference frame at the downstream end of element `C`, and $\mathcal{O}_G$ is the default `origin` reference frame of the girder if the `origin_ele` parameter is not set. $\mathbf{r}_{CA}$ is the vector from $\mathcal{O}_A$ to $\mathcal{O}_C$. The length `l` of the girder is set to be the difference in s between points $\mathcal{O}_C$ and $\mathcal{O}_A$.

A lattice element may have at most one `girder` supporting it. However, a `girder` can be supported by another `girder` which in turn can be supported by a third `girder`, etc. Girders that support other girders must be defined in the lattice file after the supported girders are defined. Example:

```
g1: girder = {A, B, C}
g2: girder = {g1}       ! g2 must come after g1!
```

A `girder` may not directly support `multipass_slave` (§9) or `super_slave` (§8) elements. Rather, a `girder` may support the corresponding lord elements.

The reference frame from which the girder's offset, pitch, and tilt attributes (§5.6) are measured is constructed as follows: A reference frame, called the "`origin`" reference frame may be defined using the attributes `origin_ele` and `origin_ele_ref_pt` which constructs the girder's `origin` frame to be coincident with the reference frame of another element. Example:

```
g2: girder = {...}, origin_ele = Q, origin_ele_ref_pt = entrance_end
```

In this example, girder `g2` has an `origin` reference frame coincident with the entrance end frame of an element named `Q`. Valid values for `origin_ele_ref_pt` are

```
entrance_end
center         ! Default
exit_end
```

For `crystal`, `mirror`, and `multilayer_mirror` elements, setting `origin_ele_ref_pt` to `center` results in the reference frame being the frame of the surface (cf. Fig. 5.6).

To specify that the global coordinates (§16.2) are to be used for a girder set `origin_ele` to

```
global_coordinates
```

Typically this is the same as using the `beginning` element (§4.4) as the `origin_ele` except when the `beginning` element is offset or reoriented (§10.4).

If `origin_ele` is not given, the default `origin` frame is used. The default `origin` frame is constructed as follows: Let $\mathcal{O}_A$ be the reference frame of the upstream end of the first element in the list of supported elements. In this example it is the upstream end of element `A` as shown in the figure. Let $\mathcal{O}_C$ be

the downstream end of the last element in the list of supported elements. In this example this is the downstream end of element `C`. The origin of the `girder`'s reference frame, marked $\mathcal{O}_G$ in the figure, will be half way along the vector r_{CA} from the origin of $\mathcal{O}_A$ to the origin of $\mathcal{O}_B$. The orientation of $\mathcal{O}_G$ is constructed by rotating the $\mathcal{O}_A$ coordinates about an axis perpendicular to the z-axis of $\mathcal{O}_A$ and $\mathbf{r}_{CA}$ such that the z-axis of $\mathcal{O}_G$ is parallel with r_{CA}.

Once the `origin` reference frame is established, the reference frame of the girder can be offset from the `origin` frame using the parameters

```
dx_origin     dtheta_origin
dy_origin     dphi_origin
dz_origin     dpsi_origin
```

The orientation of the `girder`'s reference frame from the `origin` frame is given in §16.2.4. Example:

```
g3: girder = { ... }, dx_origin = 0.03
```

This offsets girder `g3`'s reference frame 3 cm horizontally from the default `origin` frame. If no offsets are given, the `origin` frame is the same as the girder's reference frame.

The length `l` of a girder, which is not used in any calculations, is a dependent attribute computed by *Bmad* and set equal to the s path length between points $\mathcal{O}_C$ and $\mathcal{O}_A$.

The physical orientation of the girder with respect to it's reference frame is, like other elements, determined by the offset, pitch and tilt orientation attributes as outlined in §5.6 and §16.2.4. When a girder is shifted in space, the elements it supports are also shifted. In this case, the orientation attributes (`x_offset`, `y_pitch`, etc.) give the orientation of the element with respect to the `girder`. The orientation with respect to the local reference coordinates is given by `x_offset_tot`, which are computed from the orientation attributes of the element and the `girder`. An example will make this clear:

```
q1: quad, l = 2
q2: quad, l = 4, x_offset = 0.02, x_pitch = 0.01
d: drift, l = 8
g4: girder = {q1, q2}, x_pitch = 0.002, x_offset = 0.03
this_line: line = (q1, d, q2)
use, this_line
```

In this example, `g4` supports quadrupoles `q1` and `q2`. Since the supported elements are colinear, the computation is greatly simplified. The reference frame of `g4`, which is the default `origin` frame, is at $s = 7$ meters which is half way between the start of `q1` at at $s = 0$ meters and the end of `q2`) which is at $s = 14$. The reference frames of `q1` and `q2` are at their centers so the s positions of the reference frames is

```
Element        S_ref   dS_from_g4
q1              1.0     -6.0
g4              7.0      0.0
q2             12.0      5.0
```

Using a small angle approximation to simplify the calculation, the `x_pitch` of `g4` produces an offset at the center of `q2` of $0.01 = 0.002 * 5$. This, added to the offsets of `g4` and `q2`, give the total `x_offset`, denoted `x_offset_tot` of `q2` is $0.06 = 0.01 + 0.03 + 0.02$. The total `x_pitch`, denoted `x_pitch_tot`, of `q2` is $0.022 = 0.02 + 0.001$.

A `girder` that has its `is_on` attribute set to False is considered to be unsifted with respect to it's reference frame.

4.25 GKicker

A gkicker element is a "general" zero length kicker element that can displace a particle in all six phase space dimensions.

General group attributes are:

Attribute Class	Section	*Attribute Class*	Section
Custom Attributes	3.9	Description strings	5.3

See §15.29 for a full list of element attributes along with a their units.

Attributes specific to a gkicker are:

```
x_kick        = <Real>  ! X-position kick
px_kick       = <Real>  ! X-momentum kick
y_kick        = <Real>  ! Y-position kick
py_kick       = <Real>  ! Y-momentum kick
z_kick        = <Real>  ! Z-position kick
pz_kick       = <Real>  ! Momentum kick
```

Example:

```
gk: gkicker, x_kick = 0.003, pz_kick = 0.12
```

4.26 Group

Group elements are a type of control element (§2.4) used to make variations in the attributes of other elements (called "slave" attributes) during execution of a program. For example, to simulate the action of a control room knob that changes the beam tune in a storage ring, a group element can be used to vary the strength of selected quads in a specified manner. Also see overlay (§4.40) The difference between group and overlay elements is that overlay elements set the values of the attributes directly while group elements make delta changes to attribute values.

General group attributes are:

Attribute Class	Section	*Attribute Class*	Section
Custom Attributes	3.9	Is_on	5.14
Description strings	5.3		

See §15.29 for a full list of element attributes along with a their units.

There are two types of group elements: Expression based and knot based. The general syntax for a expression based group element is

```
name: GROUP = {ele1[attrib1]:exp1, ele2[attrib2]:exp2, ...},
              VAR = {var1, var2, ...}, var1 = init_val1, old_var1 = init_val_old1, ...
```

where ele1[attrib1], ele2[attrib2], etc. specify the slave attributes and exp1, exp2, etc. are the arithmetical expressions, that are functions of var1, var2, etc., and are used to determine a value for the slave attributes.

The general syntax for a knot based group element is

```
name: GROUP = {ele1[attrib1]:{y_knot_points1}, ele2[attrib2]:{y_knot_points2}, ...},
              VAR = {var1}, X_KNOT = {x_knot_points}, INTERPOLATION = {type},
              var1 = init_val1, old_var1 = init_val_old1, ...
```

When using knot points, the group may only have one variable parameter.

See Section §5.4 for a detailed description of this syntax.

Example of a expression based group element:

```
gr1: group = {q[k1]:a+b^2}, var = {a, b}, a = 1, old_a = 2
gr1[old_b] = 2
```

There are two numbers associated with each variable in a group: One number is the value of the variable (also called the "present" value) and the other number is the "old" value. To refer to these old values prepend the string "old_" to the variable name. Thus, in the above example, the old variable values have names old_a and old_b and these old values can be set in the same manner as the present values.

Example of a knot based group element:

```
gr2: group = {beginning[E_tot]]:{4e6,...}},
     var = {time}, x_knot = {...}, interpolate = cubic
```

Here the function used to translate from the group's variables to the slave attribute values is a cubic spline interpolation based upon the knot points specified (§5.4).

A group element is like an overlay element in that a group element controls the attribute values of other "slave" elements. The difference is that the value of a slave attribute that is controlled by (one or more) overlay elements is uniquely determined by the controlling overlay elements. A group element, on the other hand, is used to make changes in value. An example will make this clear:

```
gr: group = {q1[k1]:0.1*a^2}, var = {a}, a = 2, old_a = 1
q, quad, k1 = 0.5
```

When a program reads the lattice file, initially the value of q[k1] will be 0.5 as set in the definition of q. Later, during lattice expansion (§3.26), the group elements are added to the lattice. When the group element gr is added, the fact that old_a and a are different causes the value of q[k1] to be modified. The delta value is

```
delta = 0.1*a^2 - 0.1*old_a^2
      = 0.3
```

And this is added to the existing value of 0.5 so that the value of q[k1] becomes 0.8. After the value of q[k1] has been updated, the value of old_a is automatically update to be the present value of a so that the value of q[k1] will not be further modified.

In general, deltas used to modify slave attributes are computed as the difference between the arithmetic expression evaluated with the present variable values and the arithmetic expression evaluated with the old variable values.

Notice that in a lattice file the value of a slave attribute after the lattice is read in is independent of whether the group is defined before or after elements whose attributes are controlled by the group. This is true since the effect of a group element happens when the lattice is expanded, not when parser reads the group definition. On the other hand, after the lattice has been read in, if a program varies both a group variable and a slave attribute, the value of the slave attribute will be dependent upon the order of which is modified first. For example, consider a lattice containing:

```
gr: group = {q[k1]:a^2}, var = {a}
q, quad
```

Now if a program first sets gr[a] to 0.3 and then sets q[k1] to 0.5, the result is that q[k1] will have a value of 0.5. That is, the value of q[k1] will be independent of gr[a]. If the setting is reversed so that q[k1] is set first, the value of q[k1] will be 0.59. Since the result is order dependent, trying to "simultaneously" vary the attributes of both group variables and slave attributes can lead to unpredictable results. For example, consider lattice "optimization" where a program varies a set of lattice parameters to achieve certain goals (for example, minimum beta at some point in the lattice, etc.). If the list of parameters to be varied contains both group variables and slave attributes, the actual changes to slave attributes may be different from what the user expects when the program varies its list of parameters.

During running of a program, If a group element has been turned off (is_on parameter set false), a change to a variable of the group element is saved but the old variable value is not updated and slave parameters are not affected. Subsequently, turning on the group element will result in the appropriate changes to the slave parameters and the old variable value.

Different group elements may control the same slave attribute and a group element may control other group, overlay or girder element attributes. However, It does not make sense, and it is not permitted, for a group element to control the same attribute as an overlay element or for a group element to control a dependent attribute (§5.1). To setup a group element to control the same slave attribute as an overlay, define an intermediate overlay. For example:

```
ov: overlay = {qk1  q2[k1], ...}, var = {a}
q, quad
gr: group = {ov_q[k1]:a^2}, var = {a}  ! New
ov_q: overlay = {q}, var = {k1}        ! New
```

In this example, the overlay ov controls the attribute q[k1] so it is not permitted for q[k1] to be a slave of a group element. To have group control of q[k1], two elements are introduced: the group gr is setup controlling ov_q[k1] and overlay ov_q is an overlay that controls q[k1]. Notice that trying to control ov directly by a group element will not work since ov controls multiple elements.

A group can be used to control an elements position and length using one of the following attributes:

```
accordion_edge  ! Element grows or shrinks symmetrically
```

```
start_edge      ! Varies element's upstream edge s-position
end_edge        ! Varies element's downstream edge s-position
s_position      ! Varies element's overall s-position. Constant length.
```

With `accordion_edge`, `start_edge`, `end_edge`, and `symmetric_edge` the longitudinal position of an elements edges are varied. This is done by appropriate control of the element's length and the lengths of the elements to either side. In all cases the total length of the lattice is kept invariant.

As an example, consider `accordion_edge` which varies the edges of an element so that the center of the element is fixed but the length varies:

```
gr: group = {Z[accordion_edge]:1}, var = {offset}
```

A change of, say, 0.1 `gr`'s `offset` variable moves both edges of element `Z` by 0.1 meters so that the length of `Z` changes by 0.2 meters but the center of `Z` is constant. To keep the total lattice length invariant, the lengths of the elements to either side are decreased by 0.1 meters to keep the total lattice length constant.

```
q10: quad, l = ...
q11: quad, l = ...
d1: drift, l = ...
d2: drift, l = ...
this_line: line = (... d1, q10, d2, q11, ...)
gr2: group = {q10[start_edge]:1}, var = {a}, a = 0.1
```

The effect of `gr2[a]` will be to lengthen the length of `q10` and shorten the length of `d1`.

A lattice file may contain lines and lattice elements that are not part of the actual finished lattice when the lattice is constructed. `Group` elements where *none* of its slave elements are part of the finished lattice are ignored and are also not part of the finished lattice. When a `group` element has some slave elements that are part of the finished lattice and some slave elements that are not, the `group` element as implemented in the finished lattice will only control slave elements that actually exist in the finished. In any case, a slave element that a `group` element references must be defined (but not necessarily be used in the finished lattice) in the lattice file. This rule is enforced in order to catch spelling mistakes.

If the arithmetical expression used for an `group` contains an element attribute, care must be taken if that element attribute is changed. This is discussed in §3.13 and §5.4.

4.27 Hybrid

A `hybrid` element is an element that is formed by a program by concatenating other element together. `Hybrid` elements are used to reduce the number of elements in a lattice to speed up a simulation. In terms of tracking a `hybrid` element is essentially the same as a `taylor` element.

4.28 Instrument, Monitor, and Pipe

Essentially *Bmad* treats `instrument`, `monitor`, and `pipe` elements like a `drift`. There is a difference, however, when superimposing elements (§8). For example, a `quadrupole` superimposed on top of a `drift` results in a free `quadrupole` element in the tracking part of the lattice and no lord elements are created. On the other hand, a `quadrupole` superimposed on top of a `monitor` results in a `quadrupole` element in the tracking part of the lattice and this `quadrupole` element will have two lords: A `quadrupole` superposition lord and a `monitor` superposition lord. The exception is if a `instrument`, `monitor`, and `pipe` is superimposed with an element with non-constant reference energy like a `lcavity`. In this case no `instrument`, `monitor`, or `pipe` `super_lord` element is made.

General `instrument`, `monitor`, and `pipe` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Is_on	5.14
Chamber wall	5.12	Length	5.13
Custom Attributes	3.9	Offsets, pitches & tilt	5.6
Description strings	5.3	Reference energy	5.5
Hkick & Vkick	5.7	Superposition	8
Instrumental variables	5.22	Symplectify	6.7
Integration settings	6.4	Tracking & transfer map	6

See §15.31 for a full list of element attributes along with a their units.

The `offset`, `pitch`, and `tilt` attributes are not used by any *Bmad* routines. If these attributes are used by a program they are typically used to simulate such things as measurement offsets. The `is_on` attribute is also not used by *Bmad* proper. Example:

```
d21: instrum, l = 4.5
```

4.29 Kickers: Hkicker and Vkicker

An `hkicker` gives a beam a horizontal kick and a `vkicker` gives a beam a vertical kick. Also see the `kicker` (§4.30) element.

General `hkicker vkicker` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Is_on	5.14
Chamber wall	5.12	Length	5.13
Custom Attributes	3.9	Mag & Elec multipoles	5.15
Description strings	5.3	Offsets, pitches & tilt	5.6
Field Maps	5.16	Reference energy	5.5
Fringe Fields	5.21	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Tracking & transfer map	6

See §15.33 for a full list of element attributes along with a their units.

Note that `hkicker` and `vkicker` elements use the `kick` attribute while a `kicker` uses the `hkick` and `vkick` attributes. Example:

```
h_kick: hkicker, l = 4.5, kick = 0.003
```

4.30 Kicker

A `kicker` can deflect a beam in both planes. Note that a `kicker` uses the `hkick` and `vkick` attributes while `hkicker` and `vkicker` elements use the `kick` attribute. In addition, a `kicker` can apply a displacement to a particle using the `h_displace` and `v_displace` attributes.

General `kicker` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Mag & Elec multipoles	5.15	Length	5.13
Aperture limits	5.8	Offsets, pitches & tilt	5.6
Chamber wall	5.12	Overlapping Fields	5.18
Custom Attributes	3.9	Reference energy	5.5
Description strings	5.3	Superposition	8
Fringe Fields	5.21	Symplectify	6.7
Hkick & Vkick	5.7	Field Maps	5.16
Integration settings	6.4	Tracking & transfer map	6
Is_on	5.14		

See §15.32 for a full list of element attributes along with a their units.

Example:

```
a_kick: kicker, l = 4.5, hkick = 0.003
```

4.31 Lcavity

An `lcavity` is a LINAC accelerating cavity. The main difference between an `rfcavity` and an `lcavity` is that, unlike an `rfcavity`, the reference energy (§17.1) through an `lcavity` is not constant.

General `lcavity` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Mag & Elec multipoles	5.15
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Overlapping Fields	5.18
Description strings	5.3	Reference energy	5.5
Field autoscaling	5.19	RF Couplers	5.17
Fringe Fields	5.21	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Field Maps	5.16
Is_on	5.14	Tracking & transfer map	6
Length	5.13	Wakes	5.20

See §15.34 for a full list of element attributes along with a their units.

The attributes specific to an `lcavity` are

```
cavity_type     = <Switch>  ! Type of cavity.
gradient        = <Real>    ! Accelerating gradient (V/m).
gradient_err    = <Real>    ! Accelerating gradient error (V/m).
gradient_tot                ! Net gradient = gradient + gradient_err. Dependent param (§5.1).
phi0            = <Real>    ! Phase (rad/2π) of the reference particle with
                            !   respect to the RF. phi0 = 0 is on crest.
phi0_autoscale              ! Set by Bmad when autoscaling is turned on §5.19.
phi0_multipass  = <Real>    ! Phase (rad/2π) with respect to a multipass lord (§9).
phi0_err        = <Real>    ! Phase error (rad/2π)
e_loss          = <Real>    ! Loss parameter for short range wakefields (V/Coul).
rf_frequency    = <Real>    ! RF frequency (Hz).
field_autoscale             ! Set by Bmad when autoscaling is turned on §5.19.
voltage                     ! Cavity voltage. Dependent attribute (§5.1).
voltage_err                 ! Error voltage
voltage_tot                 ! Net voltage = voltage + voltage_err. Dependent param (§5.1).
l_active        = <Real>    ! Active region length. Dependent attribute (§5.1).
n_cell          = <Int>     ! Number of cavity cells. Default is -1.
n_rf_steps      = <Int>     ! Number of steps in kick-drift-kick model.
ks              = <Real>    ! Normalized solenoid strength.
bs_field        = <Real>    ! Solenoid field strength.
longitudinal_mode = <Int>   ! Longitudinal mode. Default is 0. May be 0 or 1.
```

The `voltage` and `voltage_err` attributes can be used in place of `gradient` and `gradient_err`. The relationship between gradient and voltage is

```
voltage     = L * gradient
voltage_err = L * gradient_err
```

The energy kick felt by a particle, assuming no phase slippage, is

$$\mathrm{dE} = r_q \cdot \mathrm{gradient_tot} \cdot L \cdot \cos(2\,\pi\,(\phi_t + \phi_{\mathrm{ref}})) \tag{4.22}$$

where r_q is the charge of the particle relative to the charge of the reference particle and where the total gradient is

```
gradient_tot = gradient + gradient_err
```

ϕ_t is the part of the phase due to when the particle arrives at the cavity and depends upon whether `absolute time tracking` or `relative time tracking` is being used as discussed in §26.1. The phase $\phi_{\rm ref}$ is

$\phi_{\texttt{ref}}$ `= phi0 + phi0_multipass + phi0_err`

In the above equation r_q is the relative charge between the reference particle (set by the `parameter[particle]` parameter in a lattice file) and the particle being tracked through the cavity. For example, if the reference particle and the tracked particle are the same, r_q is unity independent of the type of particle tracked.

`phi0_multipass` is only to be used with multipass elements to shift the phase of the cavity from pass to pass. See §9 for details.

The `phi0_autoscale` and `field_autoscale` parameters are calculated by *Bmad*'s auto-scale module to adjust the RF amplitude and phase so that during tracking, the energy gain of a particle on the reference orb is what the user expects based upon the set set voltage and phase. See Section §5.19 for more details. In particular, the gradient and phase are adjusted by

```
gradient (tracking) = gradient_tot * field_autoscale
```

$\phi_{\texttt{ref}}$ `(tracking) =` $\phi_{\texttt{ref}}$ `+ phi0_autoscale`

These autoscale shifts are not used when an element's `tracking_method` is `bmad_standard` since, in this case, the tracking algorithm is such that no adjustments are needed. The calculation of autoscale values can be toggled on/off by setting the `autoscale_phase` and `autoscale_amplitude` logicals.

The energy change of the reference particle is just the energy change for a particle with $z = 0$ and no phase or gradient errors. Thus

`dE(reference) = gradient * L *` $\cos(2\pi *$ $\phi_{\texttt{ref}})$

The energy kick for a *Bmad* `lcavity` is consistent with MAD. Note: The MAD8 documentation for an `lcavity` has a wrong sign. Essentially the MAD8 documentation gives

`dE = gradient * L *` $\cos(2\pi *$ `(`$\phi_{\texttt{ref}}$ `- phi(z))) ! WRONG`

This is incorrect.

The `n_rf_steps` parameter sets the number of drifts used in the kick-drift-kick model (§26.14) used when the `tracking_method` is set to `bmad_standard`. When using a fixed step Runge-Kutta `tracking_method`, the step size is determined by the setting of `num_steps` or `ds_step`. The reason why there are two different parameters to choose a step size is due to the fact that, in general, the number of drifts needed for an accurate result is much smaller than the number Runge-Kutta steps.

When short-range wakefields are being simulated, with `bmad_com%sr_wakes_on = True` (§11.4), the `e_loss` attribute can be used to modify the gradient in order to maintain a constant average energy gain. That is, `e_loss` can be used to simulate the effect of a feedback circuit that attempts to maintain the average energy of the bunch after the element constant. The energy kick is then

```
dE(with wake) = dE + e_loss * n_part * e_charge
```

`n_part` is set using the `parameter` statement (§10.1) and represents the number of particles in a bunch. `e_charge` is the magnitude of the charge on an electron (Table 3.2). Notice that the `e_loss` term is independent of the sign of the charge of the particle.

The `cavity_type` is the type of cavity being simulated. Possible settings are:

```
ptc_standard
standing_wave    ! Default
traveling_wave
```

The cavity_type switch is ignored if a field map is used. The traveling_wave cavity has an electromagnetic wave traveling (in body coordinates) in the forward direction. The standing_wave cavity has a standing wave which can be decomposed into a forward traveling wave and a backwards traveling wave.

With the standing_wave setting, the transverse trajectory through an lcavity is modeled using equations developed by Rosenzweig and Serafini[Rosen94]. See Section §26.14 for more details. Note: The transfer matrix for an lcavity with finite gradient is never symplectic if the reference energy at either end is different. See §17.2.

With field_calc set to bmad_standard, the electromagnetic fields in a lcavity are described in Section (§18.8). With cavity_type set to standing_wave, the longitudinal mode is set by the longitudinal_mode parameter. The possible values are 0 or 1 and the default setting is 0. The electromagnetic fields are used for tracking when an element's tracking_method is set to runge_kutta or any similar tracking method. If an element's cavity_type parameter is set to standing_wave, with field_calc set to bmad_standard, the "active region" over which there is a finite field is n_cell half-wave pillbox resonators where each pillbox has length $\lambda/2$ (§18.8). The default setting for n_cell is -1 which will result in the active region being as large as possible without exceeding the length of the element.

The dependent parameter l_active is set to the length of the active region. The active region is limited to not exceed the length of the element. If the length of the element is not equal to the active region, the active region is centered in the element and the regions to either side are treated as field free.

When an element's tracking_method is set to bmad_standard, settings for the parameters n_cell, and longitudinal_mode, num_steps, and ds_step are ignored.

Example:

```
lwf: lcavity, l = 2.3, rf_frequency = 500e6, voltage = 20e6,
       tracking_method = runge_kutta, n_cell = 3, longitudinal_mode = 1
```

Lcavity elements can model couplers (§5.17) and HOM wakes (§5.20) as well as multipoles (5.15). Note that multipoles in a lcavity do not vary in time.

4.32 Marker

A `marker` is a zero length element meant to mark a position.

General `marker` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Is_on	5.14
Chamber wall	5.12	Offsets & tilt	5.6
Custom Attributes	3.9	Reference energy	5.5
Description strings	5.3	Superposition	8
Instrumental variables	5.22	Tracking & transfer map	6

See §15.37 for a full list of element attributes along with a their units.

The `x_offset`, `y_offset` and `tilt` attributes are not used by any *Bmad* routines. Typically, if these attributes are used by a program, they are used to simulate things like BPM offsets. The `is_on` attribute is also not used by *Bmad* proper.

Example:

```
mm: mark, type = "BPM"
```

4.33 Mask

A mask element defines an aperture where the mask area can essentially have an arbitrary shape.

For X-ray tracking, a mask element is similar to a diffraction_plate (§4.13) element except that with a diffraction_plate element, coherent effects are taken into account while, with a mask element, coherent effects are ignored. Also a mask element can be used with charged particles while a diffraction_plate cannot.

General mask element attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Reference energy	5.5
Description strings	5.3	Superposition	8
Is_on	5.14	Tracking & transfer map	6

See §15.38 for a full list of element attributes along with a their units.

Notice that, unlike a rcollimator or a ecollimator, a mask element has zero length.

Attributes specific to a mask element are:

```
mode                = <Type>   ! Reflection or transmission (photon tracking only).
field_scale_factor = <Real>   ! Factor to scale the photon field.
ref_wavelength                ! Reference wavelength (§5.5). Dependent attrib (§5.1).
wall                = {...}   ! Defines mask geometry (5.12, 5.12.6).
```

Note: These attributes are only pertinent for photon tracking. Charged particle tracking assumes transmission mode and does not use field_scale_factor and ref_wavelength attributes.

The mode switch, which is only used for photon tracking, sets whether X-rays are transmitted through the mask or or reflected. Possible values for the mode switch are:

```
reflection
transmission        ! Default
```

The geometry of the mask, that is, where the openings (in transmission mode) or reflection regions are, is defined using the "wall" attribute. See §5.12 and 5.12.6 for more details.

In transmission mode, a mask is nominally orientated transversely to the beam. Like all other elements, the mask can be reoriented using the element's offsets, pitches and tilt attributes (§5.6).

The aperture_type (§5.8) parameter of a mask will default to auto which will set the aperture limits to define a rectangular aperture that just cover the clear area of the mask.

The field_scale_factor, if set to a non-zero value (zero is the default) will be used to scale the field of photons as they pass through the mask element:

```
field -> field * field_scale_factor
```

Scaling is useful since the electric field of photons traveling through a mask are renormalized (see Eqs. (27.10) and (27.11)). This can lead to large variation of the photon field and can, for example, make visual interpretation of plots of field verses longitudinal position difficult to interpret. field_scale_factor can be used to keep the field more or less constant.

A mask that is "turned off" (is_on attribute set to False), does not mask at all and transmits everything.

Example:

```
scrapper: mask, mode = transmission, wall = {
    section = {type = clear, v(1) = {0.9, 0.5}},
    section = {type = opaque, r0 = (0, 0.4), v(1) = {0.1, 0.1}}
  }
```

4.34 Match

A match element is used to match the Twiss parameters between two points.

General match attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Length	5.13
Custom Attributes	3.9	Reference energy	5.5
Description strings	5.3	Superposition	8
Is_on	5.14		

See §15.39 for a full list of element attributes along with a their units.

Attributes specific to a match element are:

```
beta_a0, beta_b0   = <Real>            ! Entrance betas
beta_a1, beta_b1   = <Real>            ! Exit betas
alpha_a0, alpha_b0 = <Real>            ! Entrance alphas
alpha_a1, alpha_b1 = <Real>            ! Exit alphas
eta_x0, eta_y0     = <Real>            ! Entrance etas
eta_x1, eta_y1     = <Real>            ! Exit etas
etap_x0, etap_y0   = <Real>            ! Entrance momentum dispersion
etap_x1, etap_y1   = <Real>            ! Exit eta'
c11_mat0, c12_mat0, c21_mat0, c22_mat0 = <Real>  ! Entrance coupling.
c11_mat1, c12_mat1, c21_mat1, c22_mat1 = <Real>  ! Exit coupling.
mode_flip0, mode_flip1 = <T/F>         ! Mode flip status (§23.1). Default: False.
dphi_a, dphi_b             = <Real>    ! Phase advances (radians).
x0, px0, y0, py0, z0, pz0 = <Real>     ! Entrance coordinates
x1, px1, y1, py1, z1, pz1 = <Real>     ! Exit coordinates
delta_time                 = <Real>    ! Change in time.
matrix                     = <Switch>  ! Matrix calculation. Default: standard.
kick0                      = <Switch>  ! Zeroth order calc. Default: standard.
spin_tracking_model        = <Switch>  ! How to track the spin?
recalc                     = <Logical> ! Calculate transfer map? Default is True.
```

The transfer map for a match element is a first order transformation:

$$r_1 = \mathbf{M}\, r_0 + \mathbf{V} \tag{4.23}$$

where r_1 is the output coordinates, and r_0 are the input coordinates. The matrix $\mathbf{M}$ is the linear part of the map and the vector $\mathbf{V}$ is the zeroth order ("kick") part of the map.

Nomenclature: The parameters beta_a0, alpha_a0, etc. of the match element are called the entrance (upstream) "element" Twiss parameters. The parameters beta_a1, alpha_a1, etc. of the match element are called the exit (downstream) "element" Twiss parameters. Similarly, c11_mat0, etc. are the entrance components of the $\mathbf{C}$ coupling matrix (§23.1) and c11_mat1 are the exit end element values. These parameters in general will be different from the actual computed Twiss and coupling parameters at the ends of the match element.

The matrix switch determines how the linear $\mathbf{M}$ matrix is calculated. Possible settings of this parameter are:

```
match_twiss      ! Element entrance Twiss set to actual Twiss to match exit end values.
identity         ! M is set to the unit matrix.
```

```
phase_trombone   ! Element Twiss are set so that M produces a pure phase shift.
standard         ! Use element values to calc matrix (default).
```

standard
: With `matrix` set to `standard` (the default), the matrix **M** is calculated such that if (and only if) the actual Twiss and coupling parameters at the entrance of the `match` element are equal to the element entrance Twiss and coupling parameters, then the computed Twiss and coupling parameters at the exit end of the `match` element will be the element end Twiss and coupling parameters. Additionally, the phase advances (in radians) will be `dphi_a` and `dphi_b`. Exception: If `beta_a0`, `beta_b0`, `beta_a1`, and `beta_b1` are all zero, the matrix will be set to the identity.

identity
: With `matrix` set to `identity` the transfer matrix will be set to the unit matrix independent of the element Twiss and coupling settings.

match_twiss
: The `match_twiss` setting for `matrix` instructs *Bmad*, when a program is run, to set the element entrance Twiss and coupling values to the computed Twiss and coupling values from the exit end of the previous element. This ensures that the computed Twiss and coupling at the element's exit end will correspond to the element Twiss and coupling values. This is only done if `recalc` is set to True (the default). If `recalc` is False, no element Twiss and coupling parameters are modified and the transfer matrix is calculated from the element Twiss and coupling parameters the same as the `standard` setting.

  `match_twiss` with `recalc` set to True can only be used with lattices with an `open` geometry (§10.1) since, for a `closed` lattice, it is not possible to calculate the Twiss parameters at the previous element independently of the element end Twiss parameters at the `match` element.

  When running a program, if a `match` element initially has `matrix` set to `match_twiss` and `recalc` is set to True, the *Bmad* bookkeeping routines will ensure that the `match` element's entrance element Twiss parameters are appropriately set as explained above. If `recalc` is now toggled to False (which is done automatically, for example, by the *Tao* program), the entrance Twiss attribute values, and hence the transfer matrix for the `match` element, will be frozen. Thereafter, variation of any parameter in the lattice that affects the calculated Twiss parameters at the entrance of the `match` element will not affect the `match` element's transfer matrix.

phase_trombone
: The `phase_trombone` setting for `matrix` is conceptually similar to the `match_twiss`. The difference is that with `phase_trombone`, *Bmad* will modify both the entrance and exit element parameters so that the actual entrance Twiss and coupling equals the actual exit Twiss and coupling and there will be a phase advance through the match element that is set by `dphi_a` and `dphi_b` for the a and b modes respectively.

  Like the `match_twiss` setting, the `recalc` parameter determines if *Bmad* will modify the element parameters.

Note: With `match_twiss` and `phase_trombone` settings the element's Twiss and coupling parameters are modified. With `identity` and `standard` the element parameters are not varied.

Note: There is an old notation with logical parameters `match_end` and `phase_trombone` instead of `matrix` and `recalc`. The correspondence is

match_end	phase_trombone	matrix	recalc
False	False	standard	False or True
True	False	match_twiss	True
False	True	phase_trombone	True

The setting of `match_end` and `phase_trombone` both True is not allowed.

The setting of the `kick0` parameter determines how the zeroth order transfer map vector **V** is constructed. Possible settings of this parameter are:

```
match_orbit    ! Element entrance orbit set to actual orbit to match exit end values.
zero           ! Set  V zero.
standard       ! Use element orbit values to calc  V (default).
```

standard

With `kick0` set to `standard` (the default), the vector **V** is set so that if a particle enters the `match` element with position `(x0, px0, y0, py0, z0, pz0)` the element position at the exit end will be `(x1, px1, y1, py1, z1, pz1)`. With this, **V** will be:

$$\mathbf{V} = \begin{pmatrix} \text{x1} \\ \text{px1} \\ \text{y1} \\ \text{py1} \\ \text{z1} \\ \text{pz1} \end{pmatrix} - \mathbf{M} \begin{pmatrix} \text{x0} \\ \text{px0} \\ \text{y0} \\ \text{py0} \\ \text{z0} \\ \text{pz0} \end{pmatrix} \tag{4.24}$$

match_orbit

The `match_orbit` setting for `kick0` instructs *Bmad*, when a program is run, to set the element entrance position to the computed orbit from the exit end of the previous element. This ensures that the computed orbit at the element's exit end will correspond to the position set in the element. This is only done if `recalc` is set to True (the default). If `recalc` is False, no element position values are modified and **V** is calculated from the element position the same as the `standard` setting.

Like the situation with `match_twiss` set with `matrix`, `match_orbit` with `recalc` set to True can only be used with lattices with an `open` geometry (§10.1) since, for a `closed` lattice, it is not possible to calculate the Twiss parameters at the previous element independently of the element end Twiss parameters at the `match` element.

zero

The `zero` setting for `kick0` sets **V** to zero independent of the element position values.

Note: With the `match_orbit` setting the element's entrance orbit values will be modified. With `zero` and `standard` settings the elements parameters will not be varied by *Bmad*.

The `delta_time` parameter adds a constant to the particle's time. This will also affect the z phase space coordinate through Eq. (17.3) and the transfer map though the element. If `delta_time` is zero, the transfer map through the element will be the $\mathbf{M}$ matrix as discussed above. With a finite `delta_time`, the transfer map will be different from $\mathbf{M}$. The order of operations, is the effect of `delta_time` is applied first and the linear transformation above is applied afterwards. Since using `match_twiss` or `match_orbit` with a finite `delta_time` can be confusing, such a situation is not allowed. Use two separate `match` elements if needed.

A `match` element that is "turned off" (`is_on` attribute set to False), is considered to be like a `marker` element. That is, the orbit, spin, and Twiss parameters are unchanged when tracking through a `match` element that is turned off.

The length attribute `l` is not used in the transfer matrix calculation. The length `l` is used to compute the time it takes to go through a match element.

Example:

```
mm: match, beta_a1 = 12.5, beta_b1 = 3.4, eta_x1 = 1.0, matrix = match_twiss
```

4.35 Mirror

A `mirror` reflects photons.

General `mirror` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Reflection tables	5.10
Custom Attributes	3.9	Superposition	8
Description strings	5.3	Surface Properties	5.11
Offsets, pitches & tilt	5.6	Tracking & transfer map	6
Reference energy	5.5		

See §15.40 for a full list of element attributes along with a their units.

Attributes specific to a `mirror` element are:

```
graze_angle     = <Real>   ! Angle between incoming beam and mirror surface.
critical_angle  = <Real>   ! Critical angle.
ref_wavelength             ! Reference wavelength (§5.5). Dependent attribute (§5.1).
```

The reference trajectory for a `mirror` is that of a zero length bend (§16.2.3) and hence the length (`l`) parameter of a mirror is fixed at zero. The reference trajectory is determined by the values of the `graze_angle` and `ref_tilt` parameters. A positive `graze_angle` bends the reference trajectory in the same direction as a positive `g` for a bend element.

A `mirror` may be offset and pitched (5.6). The incoming local reference coordinates are used for these misalignments.

4.36 Multipole

A `multipole` is a thin magnetic multipole lens up to 21^{st} order. The basic difference between this and an `ab_multipole` is the input format. See section §18.1 for how the multipole coefficients are defined.

General `multipole` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Reference energy	5.5
Custom Attributes	3.9	Is_on	5.14
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Description strings	5.3	Tracking & transfer map	6
K*n*L, K*n*SL, T*n* multipoles	5.15		

See §15.42 for a full list of element attributes along with a their units.

Attributes specific to a `mirror` element are:

```
KOL_status      = <Switch> ! not_allowed (default), bends_reference, or straight_reference.
```

The length `l` is a fictitious length that is used for synchrotron radiation computations and affects the longitudinal position of the next element but does not affect any tracking or transfer map calculations.

When an `multipole` is superimposed (§8) on a lattice, it is treated as a zero length element and in this case it is an error for the length of the `multipole` to be set to a nonzero value.

The `KOL` parameter presents a compatibility problem since MAD treats a finite `KOL` as a zero length bend with a corresponding change in reference orbit (§16.1.1). On the other hand, `PTC` does not do this. An additional consideration is that this can be very confusing to users. To minimize problems, a finite `KOL` is only allowed if the `KOL_status` switch is set appropriately. Possible settings of `KOL_status` are:

```
not_allowed          ! Default. KOL not allowed to be set.
bends_reference      ! Compatible with MAD.
straight_reference   ! Compatible with PTC.
```

The tracking with `KOL_status` set to `bends_reference` follows MAD. Here a paraxial approximation is used. That is, `KOL` is assumed small compared to one. With this approximation, a particle with zero p_z will exit the multipole with the same phase space coordinates as it enters with. Notice that since the reference orbit is shifted, the setting of `KOL` will affect the placement of elements downstream of the multipole in the global coordinate system (analogous to the shift in downstream element positions after `sbend` and `rbend` elements).

With `KOL_status` set to `straight_reference`, a finite `KOL` acts as a simple dipole kick analogous to the `a0` and `b0` dipole kicks when using an `ab_multipole`.

Example:

```
m1: multipole, k1l = 0.034e-2, t1, k3sl = 4.5, t3 = 0.31*pi
```

4.37 Multilayer_mirror

A multilayer_mirror is a substrate upon which multiple layers of alternating substances have been deposited. The idea is similar to crystal diffraction: light reflected at each interface constructively interferes with light reflected from other interfaces. The amplified reflection offsets losses due to absorption.

General crystal attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Symplectify	6.7
Custom Attributes	3.9	Offsets, pitches & tilt	5.6
Description strings	5.3	Superposition	8
Reference energy	5.5	Tracking & transfer map	6
Surface Properties	5.11		

The attributes specific to a multilayer_mirror are

```
material_type    = <String>  ! Materials in each layer.
d1_thickness     = <Real>    ! Thickness of layer 1
d2_thickness     = <Real>    ! Thickness of layer 2
n_cell           = <Integer> ! Number of cells (= Number of layers / 2)
ref_wavelength               ! Reference wavelength (§5.5). Dependent attribute (§5.1).
```

See §15.41 for a full list of element attributes along with a their units.

Dependent attributes (§5.1) are

```
graze_angle      ! Angle between incoming beam and mirror surface.
v1_unitcell      ! Unit cell volume for layer 1
v2_unitcell      ! Unit cell volume for layer 2
```

A multilayer_mirror is constructed of a number of "cells". The number of cells is set by n_cell. Each cell consists of two layers of dielectric material. The materials used is given by the material_type attribute. The format for this is

```
material_type = "<material_1>:<material_2>"
```

where <material_1> and <material_2> are the material names for the first and second layers of the cell respectively. The first layer is the bottom layer and the second layer is the top layer of the cell. Material names are case sensitive. So "FE" cannot be used in place of "Fe" A list of materials is given in §5.9 and can include crystal materials or elemental materials.

Example:

```
mm: multilayer_mirror, material_type = "W:BORON_CARBIDE", n_cell = 100, &
          d1_thickness = 1e-9, d2_thickness = 1.5e-9
m2: multi    ! This is a multipole element!!!
```

Note: Due to the fact that multilayer_mirrors where introduced much later than multipole elements, if there is an ambiguity in the name as shown in the above example, an element will be considered to be of type multipole.

4.38 Null_Ele

A `null_ele` is a special type of element. It is like a `marker` but it has the property that when the lattice is expanded (§7.2) all `null_ele` elements are removed. The primary use of a `null_ele` is in computer generated lattices where it can be used to serve as a reference point for element superpositions (§8). Another use is to split an element using superposition while avoiding having to add a marker element to the lattice. Example:

```
N: null_ele, superimpose, ref = quadrupole::*
```

This will split all quadrupoles in the lattice in two.

`Null_ele` elements are not generally useful otherwise.

4.39 Octupole

An octupole is a magnetic element with a cubic field dependence with transverse offset (§18.1). The bmad_standard calculation treats an octupole using a kick–drift–kick model.

General octupole attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Mag & Elec multipoles	5.15
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Overlapping Fields	5.18
Description strings	5.3	Reference energy	5.5
Fringe Fields	5.21	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Field Maps	5.16
Is_on	5.14	Tracking & transfer map	6
Length	5.13		

See §15.43 for a full list of element attributes along with a their units.

Attributes specific to an octupole element are:

```
k3           = <Real>   ! Octupole strength.
b3_gradient  = <Real>   ! Field strength. (§5.1).
field_master = <T/F>    ! See §5.2.
```

The normalized octupole k3 strength is related to the unnormalized b3_gradient field strength through Eq. (18.3).

If the tilt attribute is present without a value then a value of $\pi/8$ is used. Example:

```
oct1: octupole, l = 4.5, k3 = 0.003, tilt ! same as tilt = pi/8
```

4.40 Overlay

Overlay elements are a type of control element (§2.4) used to make variations in the attributes of other elements (called "slave" attributes) while a program is running. For example, to simulate the action of a magnet power supply that controls a string of magnets. Also see group (§4.26) The difference between group and overlay elements is that overlay elements set the values of the attributes directly while group elements make delta changes to attribute values.

General overlay attributes are:

Attribute Class	Section	*Attribute Class*	Section
Custom Attributes	3.9	Is_on	5.14
Description strings	5.3		

See §15.44 for a full list of element attributes along with a their units.

There are two types of overlay elements: Expression based and knot based. The general syntax for a expression based overlay element is

```
name: OVERLAY = {ele1[attrib1]:exp1, ele2[attrib2]:exp2, ...},
            VAR = {var1, var2, ...}, var1 = init_val1, old_var1 = init_val_old1, ...
```

where ele1[attrib1], ele2[attrib2], etc. specify the slave attributes and exp1, exp2, etc. are the arithmetical expressions, that are functions of var1, var2, etc., and are used to determine a value for the slave attributes.

The general syntax for a knot based overlay element is

```
name: OVERLAY = {ele1[attrib1]:{y_knot_points1}, ele2[attrib2]:{y_knot_points2}, ...},
            VAR = {var1}, X_KNOT = {x_knot_points}, INTERPOLATION = {type},
            var1 = init_val1, ...
```

See Section §5.4 for a detailed description of this syntax.

An overlay element is used to control the attributes of other elements. If multiple overlays control the same slave parameter, the parameter value will be the sum of the values set by the individual overlays. For example:

```
over1: overlay = {a_ele, b_ele:2.0}, var = {hkick}, hkick = 0.003
over2: overlay = {b_ele}, var = {hkick}
over2[hkick] = 0.9
a_ele: quad, hkick = 0.05   ! NO: Cannot control slave attributes of overlays
b_ele: rbend, ...
this_line: line = (... a_ele, ... b_ele, ...)
use, this_line
```

In the example the overlay over1 controls the hkick attribute of the "slave" elements a_ele and b_ele. over2 controls the hkick attribute of just b_ele. over1[hick] has a value of 0.003 and over2[hkick] has been assigned a value of 0.9. Thus:

```
a_ele[hkick] = over1[hkick]
             = 0.003
b_ele[hkick] = over2[hkick] + 2 * over1[hkick]
             = 0.906
```

Overlays completely determine the value of the attributes that are controlled by the overlay. in the above example, the hkick of 0.05 assigned directly to a_ele is overwritten by the overlay action of over1.

The default value for an overlay is 0 so for example

```
over3: overlay = {c_ele}, var = {k1}
```

will make `c_ele[k1]` = 0.

As illustrated above, different `overlay` elements may control the same element attribute. And an `overlay` element may control other `overlay`, `group` or `girder` elements. However, It does not make sense for an `overlay` element to control the same attribute as a `group` element or for an `overlay` element to control a `dependent` attribute (§5.1).

The `is_on` parameter may be set for an overlay. If set to `False`, the `overlay` will be ignored. If all the overlays controlling a given attribute are turned off, the attribute can be set directly just like if there were no controlling overlays to begin with. Example:

```
abc: overlay = { ... }, ...
abc[is_on] = F
```

A lattice file may contain lines and lattice elements that are not part of the actual finished lattice when the lattice is constructed. `Group` elements where *none* of its slave elements are part of the finished lattice are ignored and are also not part of the finished lattice. When a `group` element has some slave elements that are part of the finished lattice and some slave elements that are not, the `group` element as implemented in the finished lattice will only control slave elements that actually exist in the finished. In any case, a slave element that a `group` element references must be defined (but not necessarily be used in the finished lattice) in the lattice file. This rule is enforced in order to catch spelling mistakes.

If the arithmetical expression used for an `overlay` contains an element attribute, care must be taken if that element attribute is changed. This is discussed in §3.13 and §5.4.

4.41 Patch

A patch element shifts the reference orbit and time. Also see floor_shift (§4.21) and fiducial (§4.19) elements. A common application of a patch is to orient two lines with respect to each other. For example, to orient an injection line with the ring it is injecting into (§13.1).

General patch element attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Offsets, pitches & tilt	5.6
Chamber wall	5.12	Reference energy	5.5
Custom Attributes	3.9	Superposition	8
Description strings	5.3	Tracking & transfer map	6
Length	5.13		

See §15.47 for a full list of element attributes along with a their units.

Attributes specific to a patch elements are:

```
x_offset        = <Real>    ! Exit face offset from Entrance.
y_offset        = <Real>    ! Exit face offset from Entrance.
z_offset        = <Real>    ! Exit face offset from Entrance.
t_offset        = <Real>    ! Reference time offset.
x_pitch         = <Real>    ! Exit face orientation from Entrance.
y_pitch         = <Real>    ! Exit face orientation from Entrance.
tilt            = <Real>    ! Exit face orientation from Entrance.
E_tot_offset    = <Real>    ! Reference energy offset (eV).
E_tot_set       = <Real>    ! Reference energy at exit end (eV).
flexible        = <T/F>     ! Default: False.
p0c_set         = <Real>    ! Reference momentum at exit end (eV).
ref_coords      = <Switch>  ! Coordinate system defining the length.
user_sets_length = <T/F>    ! User sets element length? Default is F.
l               = <Real>    ! Reference length.
```

A straight line element like a drift or a quadrupole has the exit face parallel to the entrance face. With a patch element, the entrance and exit faces can be arbitrarily oriented with respect to one another as

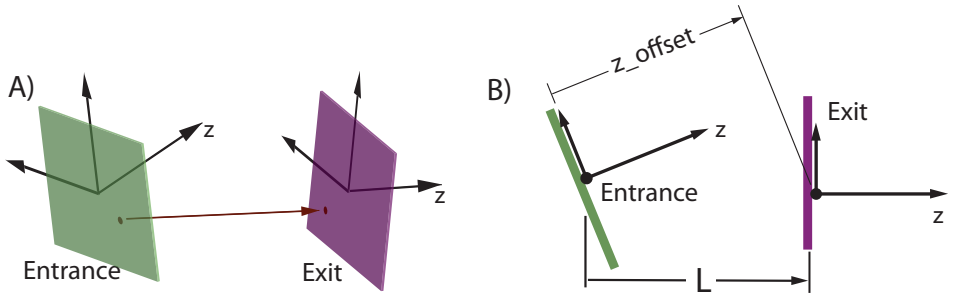

Figure 4.6: A) A patch element can align its exit face arbitrarily with respect to its entrance face. The red arrow illustrates a possible particle trajectory form entrance face to exit face. B) The reference length of a patch element, if ref_coords is set to the default value of exit_end, is the longitudinal distance from the entrance origin to the exit origin using the reference coordinates at the exit end as shown. If ref_coords is set to entrance_end, the length of the patch will be equal to the z_offset.

shown in Fig. 4.6A.

There are two different ways the orientation of the exit face is determined. Which way is used is determined by the setting of the `flexible` attribute. With the `flexible` attribute set to `False`, the default, The exit face of the `patch` will be determined from the offset, tilt and pitch attributes as described in §16.2.4. This type of `patch` is called "rigid" or "inflexible" since the geometry of the `patch` is solely determined by the `patch`'s attributes as set in the lattice file and is independent of everything else. Example:

```
pt: patch, z_offset = 3.2   ! Equivalent to a drift
```

With `flexible` set to `True`, the exit face is taken to be the reference frame of the entrance face of the next element in the lattice. In this case, it must be possible to compute the reference coordinates of the next element before the reference coordinates of the `patch` are computed. A `flexible patch` will have its offsets, pitches, and tilt as dependent parameters (§5.1) and these parameters will be computed appropriately. Here the `patch` is called "flexible" since the geometry of the patch will depend upon the geometry of the rest of the lattice and, therefore, if the geometry of the rest of the lattice is modified (is "flexed"), the geometry of the `patch` will vary as well. See Section §13.3 for an example.

The coordinates of the lattice element downstream of a `flexible patch` can be computed if there is a `fiducial` element (§4.19) somewhere downstream or if there is a `multipass_slave` (§9) element which is just downstream of the `patch` or at most separated by zero length elements from the `patch`. In this latter case, the `multipass_slave` must represent an N^{th} pass slave with N greater than 1. This works since the first pass slave will be upstream of the `patch` and so the first pass slave will have its coordinates already computed and the position of the downstream slave will be taken to be the same as the first pass slave. Notice that, without the `patch`, the position of multipass slave elements are independent of each other.

With `bmad_standard` tracking (§6.1) A particle, starting at the upstream face of the `patch`, is propagated in a straight line to the downstream face and the suitable coordinate transformation is made to translate the particle's coordinates from the upstream coordinate frame to the downstream coordinate frame (§26.17). In this case the `patch` element can be thought of as a generalized `drift` element.

If there are magnetic or electric fields within the `patch`, the tracking method through the `patch` must be set to either `runge_kutta` or `custom`. Example:

```
pa2: patch, tracking_method = runge_kutta, field_calc = custom,
            mat6_calc_method = tracking, ...
```

In order to supply a custom field when `runge_kutta` tracking is used, `field_calc` (§6.4) needs to be set to `custom`. In this case, custom code must be supplied for calculating the fields as a function of position (§38.2).

The `E_tot_offset` attribute offsets the reference energy:

```
E_tot_ref(exit) = E_tot_ref(entrance) + E_tot_offset (eV)
```

Setting the `E_tot_offset` attribute will affect a particle's p_x, p_y and p_z coordinates via Eqs. (17.2) and (17.6). Notice that `E_tot_offset` does not affect a particle's actual energy, it just affects the difference between the particle energy and the reference energy.

Alternatively, to set the reference energy, the `E_tot_set` or `p0c_set` attributes can be used to set the reference energy/momentum at the exit end. It is is an error if more than one of `E_tot_offset`, `E_tot_set` and `p0c_set` is nonzero.

`Important`: *Bmad* may apply the energy transformation either before or after the coordinate transformation. This matters when the speed of the reference particle is less than c. For this reason, and due to complications involving PTC, it is recommended to use two patches in a row when both the orbit and energy are to be patched.

A patch element can have an associated electric or magnetic field (§5.16). This can happen, for example, if a patch is used at the end of an injection line to match the reference coordinates of the injection line to the line being injected into (§13.1) and the patch element is within the field generated by an element in the line being injected into. In such a case, it can be convenient to set what the reference coordinates are since the orientation of any fields that are defined for a patch element will be oriented with respect to the patch element's reference coordinates. For this, the ref_coords parameter of a patch can be used. Possible settings are: ref_coords are:

```
entrance_end  !
exit_end      ! Default
```

The default setting of ref_coords is exit_end and with this the reference coordinates are set by the exit end coordinate system (see Fig. 4.6). If ref_coords is set to entrance_end, the reference coordinates are set by the entrance end coordinate system. Example:

```
p1: patch, x_offset = 1, x_pitch = 0.4   ! L = 0.289418 see below
p2: p1, ref_coords = entrance_end        ! L = 0
```

Here p1 has ref_coords set to exit_end (the default). p2 inherits the parameters of p1 and sets ref_coords to entrance_end.

It is important to keep in mind that if there are multiple patches in a row, while two different configurations may be the same in a geometrical sense the total length may not be the same. For example:

```
pA: patch, x_offset = 1    ! L = 0
pB: patch, x_pitch = 0.4   ! L = 0
sum: line = (pA, pB)
```

The configuration of pA followed by pB is equivalent geometrically to the p1 patch above but the total length of the (pA, pB) line is zero which is different from the length of p1.

Unfortunately, there is no intuitive way to define the "length" L of a patch. This is important since the transit time of the reference particle is the element length divided by the reference velocity. And the reference transit time will affect how the phase space z coordinate changes through the patch via Eq. (17.3). If the parameter user_sets_length is set to True, the value of l set in the lattice file will be used (default is zero). user_sets_length is set to False (the default), the length of a patch is calculated depending upon the setting of ref_coords. If ref_coords is set to exit_end, the length of the patch is calculated as the perpendicular distance between the origin of the patch's entrance coordinate system and the exit face of the patch as shown in Fig. 4.6B. If ref_coords is set to entrance_end, the length is calculated as the perpendicular distance between the entrance face and the origin of the exit coordinate system. In this case, the length will be equal to z_offset.

To provide flexibility, the t_offset attribute can be used to offset the reference time. The reference time at the exit end of the patch t_ref(exit) is related to the reference time at the beginning of the patch t_ref(entrance) via

```
t_ref(exit) = t_ref(entrance) + t_offset + dt_travel_ref
```

where dt_travel_ref is the time for the reference particle to travel through the patch. dt_travel_ref is defined to be:

```
dt_travel_ref = L / beta_ref
```

Where L is the length of the patch and beta_ref is the reference velocity/c at the exit end of the element. That is, the reference energy offset is applied *before* the reference particle is tracked through the patch. Since this point can be confusing, it is recommended that a patch element be split into two consecutive patches if the patch has finite l and E_tot_offset values.

While a finite t_offset will affect the reference time at the end of a patch, a finite t_offset will *not* affect the time that is calculated for a particle to reach the end of the patch. On the other hand, a finite t_offset will affect a particle's z coordinate via Eqs. (17.3). The change in z, δz will be

$$\delta z = \beta \cdot c \cdot \text{t_offset} \tag{4.25}$$

where β is the normalized particle speed (which is independent of any energy patch). Another way of looking at this is to note that In a drift, if the particle is on-axis and on-energy, t and t_ref change but z does not change. In a time patch (a patch with only `t_offset` finite), t_ref and z change but t does not.

When a lattice branch contains both normally oriented and reversed elements (§16.1.3), a `patch`, or series of `patches`, which reflects the z direction must be placed in between. Such a `patch`, (or patches) is called a `reflection patch`. See Section §16.2.6 for more details on how a reflection patch is defined. In order to avoid some confusing conceptual problems involving the coordinate system through a reflection patch, Runge-Kutta type tracking is prohibited with a reflection patch.[4]

Since the geometry of a `patch` element is complicated, interpolation of the chamber wall in the region of a patch follows special rules. See section §5.12.5 for more details.

[4] In general, Runge-Kutta type tracking through a patch is a waste of time unless electric or magnetic fields are present.

4.42 Photon_Init

A photon_init element is used as a starting element for x-ray tracking. A photon_init element can be used to define such things as the initial energy spectrum and angular orientation. As explained below, a photon_init element can be a "stand alone" photon source or it can have an associated "physical source" element.

Note: There is a utility program called photon_init_plot that comes with a Bmad Distribution that will plot initial photon distributions and can be used as a check.

General photon_init attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Length	5.13
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Reference energy	5.5
Description strings	5.3	Tracking & transfer map	6

See §15.48 for a full list of element attributes along with a their units.

Attributes specific to an photon_init element are:

```
  ds_slice                  = <Real>
  E_center                  = <Real>    ! Average init photon energy of 1st mode (eV).
  E2_center                 = <Real>    ! Average init photon energy of 2nd mode (eV).
  E2_probability            = <Real>    ! Probability of 2nd mode.
  E_center_relative_to_ref  = <T/F>     ! E_center relative to reference E? Default True.
  e_field_x                 = <Real>    ! Polarization. x & y = 0 -> random
  e_field_y                 = <Real>
  energy_distribution       = <Switch>  ! Gaussian, uniform, or curve.
  energy_probability_curve  = {...}   ! Used with energy_distribution = curve. See below.
  physical_source           = <String>  ! physical source of x-rays
  ref_wavelength                        ! Ref wavelength (§5.5). Dep attribute (§5.1).
  sig_x                     = <Real>
  sig_y                     = <Real>
  sig_z                     = <Real>
  sig_vx                    = <Real>
  sig_vy                    = <Real>
  sig_E                     = <Real>    ! Init photon energy width of 1st mode (eV).
  sig_E2                    = <Real>    ! Init photon energy width of 2nd mode (eV).
  spatial_distribution      = <Switch>  ! Gaussian or uniform.
  transverse_sigma_cut      = <Real>
  velocity_distribution     = <Switch>  ! Gaussian, spherical, or uniform.
```

When the energy_distribution is set to gaussian or uniform, the distribution of photons is bimodal. The first mode is characterized by the parameters E_center, and sig_E, the second mode is characterized by the parameters E2_center and sig_E2. The probability of emitting a photon in the second mode is given by E2_probability.

ds_slice
: Used when there is an associated physical source element. The physical source element is sliced into pieces of thickness ds_slice and each slice is tested to see if photons from the slice can possibly

pass through the first aperture. When photons are generated, photons will only be generated from slices where they have a hope of passing through the first aperture. This makes the simulation more efficient. The default value of `ds_slice` is 0.01 meter.

`E_center`, `E2_center`
: Average initial photon energy in eV. If `E_center_relative_to_ref` is set to True, `E_center` and `E2_center` will be relative to the reference energy.

`E_center_relative_to_ref`
: With a setting of True (the default), `E_center` and `E2_center` are taken to be with respect to the reference energy (§17.1). That is, if True, the center energy `<E>` is

```
<E-1st-mode> = E1_center + Reference_Energy
<E-2nd-mode> = E2_center + Reference_Energy
```

If `E_center_relative_to_ref` is set to False, `E_center` and `E2_center` are taken to be the center energy values independent of the reference energy.

`E2_probability`
: Probability of emitting a photon from the 2nd mode. A value of 0 (the default) will mean that all photons will be emitted from the 1st mode and a value of 1 will mean that all photons will be emitted from the 2nd mode.

`e_field_x`, `e_field_y`
: Electric field component of initial photons in the x and y planes. If both are set to 0 then a random field is chosen with unit intensity $E_x^2 + E_y^2 = 1$.

`energy_distribution`
: Sets the type of energy spectrum for emitted photons. If there is an associated physical element then this parameter is ignored and the energy distribution is calculated from the properties of the physical element. Possible settings are:

```
gaussian   ! Default
uniform
curve
```

The `gaussian` setting gives Gaussian distributions for the two modes with width set by `sig_E` and `sig_E2`. The `uniform` setting gives a flat distribution in the range:

```
[-sig_E, sig_E]    ! For the 1st mode
[-sig_E2, sig_E2]  ! For the 2nd mode
```

The `curve` setting uses the energy probability curve set by the `energy_probability_curve` component.

`energy_probability_curve`
: The `energy_probability_curve` attribute provides a way to specify the energy probability distribution when an Gaussian or uniform distribution is not suitable. The probability curve is defined by specifying the curve at a number of points. The syntax is:

```
energy_probability_curve = {E1 p1, E2 p2, ..., EN pN}
```

where the `E p` pairs are the energy and photon emission probability at that energy. The commas between `E p` pairs is optional. The probability curve does not have to be normalized, *Bmad* will take care of that. *Bmad* will use cubic spline interpolation between points.

`physical_source`
: Used to specify the "physical" source of the photons. See below for more details

`sig_E`, `sig_E2`
: Energy width of the two modes in eV. See `energy_distribution` for more details.

`sig_vx, sig_vy`
: Width of emitted photons in v_x/c and v_y/c directions. See `velocity_distribution` for more details.

`sig_x, sig_y, sig_z`
: Width of emitted photons in x, y and z directions. See `spatial_distribution` for more details.

`spatial_distribution`
: Sets spacial (x, y, z) spectrum of emitted photons. If there is an associated physical element then this parameter is ignored and the energy distribution is calculated from the properties of the physical element. Possible settings are:
```
gaussian    ! Default
uniform
```
The `gaussian` setting gives a Gaussian distribution with width σ where σ is
```
sig_x     ! for x distribution
sig_y     ! for y distribution
sig_z     ! for z distribution
```
The `uniform` setting gives a flat distribution in the range: $[-\sigma, \sigma]$.

`velocity_distribution`
: Sets the transverse $(v_x/c, v_y/c)$ velocity spectrum of emitted photons. If there is an associated physical element then this parameter is ignored and the energy distribution is calculated from the properties of the physical element. The longitudinal velocity is always computed to make $v_x^2 + v_y^2 + v_z^2 = c^2$ Possible settings are:
```
gaussian    ! Default
spherical
uniform
```
The `gaussian` setting gives a Gaussian distribution with width σ where *sigma* is
```
sig_vx     for vx/c distribution
sig_vy     for vy/c distribution
```
The `uniform` setting gives a flat distribution in the range: $[-\sigma, \sigma]$. The `spherical` setting gives flat distribution in all directions. With the `spherical` setting is used, and the next downstream element excluding drifts and markers is an element with aperture limits (§5.8), *Bmad* can optimize photon emission to only emitting photons that are very likely to be within the aperture when they hit the downstream element. This cuts down on computation time.

For the purposes of positioning the elements in the lattice around it, a `photon_init` element is considered to have zero length.

`photon_init` elements are used in one of two modes: With or without an associated physical source element specified by the `physical_source` attribute. Without an associated physical source, the `photon_init` element completely specifies the initial photon distribution. With an associated physical source element, the photon distribution is determined by the properties of the physical source but the shape of the energy spectrum can be modified by setting attributes in the `photon_init` element. Example:

```
b05w: sbend, l = 3.2, angle = 0.1
pfork: photon_fork, to_line = c_line, superimpose, ref = b05w, offset = 0.4
bend_line: line = (..., b05w, ...)
use bend_line

c_line: line = (pinit, ...)
c_line[E_tot] = 15e3
pinit: photon_init, physical_source = "b05w", sig_E = 2.1
```

In this example, the bend b05w is a bend producing photons. It is part of the line bend_line. bend_line also contains a photon_fork element named pfork which branches to the line c_line. c_line contains the photon_init element pinit which references b03w as the associated physical source element. When photons are tracked, they are generated in b05w and then propagated to the pfork fork. After this they are propagated through c_line. The pinit element acts like a zero length marker element when photons propagate through it. That is, the pinit element essentially serves to associate c_line with b03w for the purposes of photon tracking. Also, in this example, pinit modifies the photon energy spectrum so that only photons whose energy is within 2.1 eV are generated

It is important to note that in the above example, with the photon_init element having an associated physical source, the setting of things like the spatial shape sig_z, etc. in the photon_init element will be ignored.

See Section §13.5 for an example lattice that can be used to simulate a Rowland circle spectrometer using a photon_init element.

4.43 Quadrupole

A `quadrupole` is a magnetic element with a linear field dependence with transverse offset (§18.1).

General `quadrupole` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Mag & Elec multipoles	5.15
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Description strings	5.3	Overlapping Fields	5.18
Fringe Fields	5.21	Reference energy	5.5
Hkick & Vkick	5.7	Superposition	8
Integration settings	6.4	Symplectify	6.7
Is_on	5.14	Field Maps	5.16
Length	5.13	Tracking & transfer map	6

See §15.50 for a full list of element attributes along with a their units.

Attributes specific to a `quadrupole` element are:

```
b1_gradient  = <Real>    ! Field strength. (§5.1).
k1           = <Real>    ! Quadrupole strength.
fq1          = <Real>    ! Soft edge fringe parameter.
fq2          = <Real>    ! Soft edge fringe parameter.
field_master = <T/F>     ! See §5.2.
```

The normalized quadrupole `k1` strength is related to the unnormalized `b1_gradient` field strength through Eq. (18.3).

If the `tilt` attribute is present without a value then a value of $\pi/4$ is used.

For a quadrupole with zero `tilt` and a positive `k1`, the quadrupole is horizontally focusing and vertically defocusing (§18.1).

The `fq1` and `fq2` parameters are used to specify the quadrupolar "soft" edge fringe. See §19.6 for more details. The `fringe_at` and `fringe_type` settings (§5.21) determine if the fringe field is used in tracking (§5.21).

Example:

```
q03w: quad, l = 0.6, k1 = 0.003, tilt  ! same as tilt = pi/4
```

4.44 Ramper

A ramper element is a type of control element (§2.4). That is, a ramper element can be used to make variations in the attributes of other elements while a program is running. The ramper element is similar to an overlay element except that ramper elements are designed to control large sets of elements. Also ramper elements can be used to smoothly vary parameters as particle are propagated through a lattice. Ramper elements where implemented to solve the problem of simulating machine ramping where the strength of many elements in a machine are varied continuously as a function of time. The drawback of ramper elements is that they can be only be used with programs that that are designed to handle them.[5] Ramper elements will be ignored in programs that are not designed to handle them. How a program handles ramper elements will be program dependent and the program documentation should be consulted for details.

Note: ac_kicker elements can also be used for simulating a time dependent element.

General ramper attributes are:

Attribute Class	Section	*Attribute Class*	Section
Custom Attributes	3.9	Is_on	5.14
Description strings	5.3		

See §15.53 for a full list of element attributes.

The syntax for ramper elements is exactly the same as for overlay (§5.4) elements except that ramper elements do not have a gang attribute.

Like overlays, There are two types of ramper elements: Expression based and knot based. The general syntax for a expression based ramper element is

```
name: RAMPER = {ele1[attrib1]:coef1, ele2[attrib2]:coef2, ...}, VAR = {var1}
```

where ele1[attrib1], ele2[attrib2], etc. specify the slave attributes and exp1, exp2, etc. are the arithmetical expressions, that are functions of var1, var2, etc., and are used to determine a value for the slave attributes.

The general syntax for a knot based ramper element is

```
name: RAMPER = {ele1[attrib1]:{y_knot_points1}, ele2[attrib2]:{y_knot_points2}, ...},
            VAR = {var1}, X_KNOT = {x_knot_points}, INTERPOLATION = {type},
            var1 = init_val1, ...
```

See Section §5.4 for a detailed description of this syntax.

Examples:

```
ramp_e: ramper = {*[e_tot]:{4e+08, 4.00532e+08, 4.01982e+08, ...}},
              var = {time}, x_knot = {0, 0.001, 0.002, ...}, interpolation = cubic

amp = 1e9;  omega = 0.167;  t0 = 0.053
ramp_rf: ramper = {rfcavity::*[voltage]:amp*sin(omega *(t + t0)),
      rfcavity::*[phi0]:0.00158*t^2 + 3*q}, var = {t, q}
```

Ramp_e uses a cubic spline fit to interpolate between the knot points specified in the element definition. The "*[e_tot]" construct in the definition of ramp_e means that the ramper will be applied to the e_tot attribute (§5.5) of all elements (since the wild card character "*" (§3.6) will match to all element names).

[5] In particular, the long_term_tracking program that is bundled with the *Bmad* software (§1.2) can handle ramper elements.

The `ramp_rf` ramper in the above example varies the `voltage` and `phase` (`phi0`) attributes of all elements that match to `rfcavity::*`. That is, all `rfcavity` elements. Here mathematical expressions are used instead of knot points.

`Ramper` elements can control the variables of other controller element except `rampers` are not allowed to control `rampers`. When a `ramper` controls variables in other controller elements it is not permitted to use wild card characters. That is, in the above example, "`*`" will not match to any controller elements.

If a slave name contains wild card characters, for a given lattice element that the slave name matches to, it is not required that the controlled attribute be a valid attribute of the element. In the case where the controlled attribute is not valid for a given lattice element, no attributes of the given lattice element are varied when and the ramper is varied. For example:

```
rz: ramper = {*[k1]: ...
```

In this example the `k1` attribute of all those elements that have a `k1` attribute will be controlled but something like a `sextupole` element which does not have a `k1` attribute will not be controlled.

Due to the way bookkeeping is done for ramper elements, and unlike `group` or `overlay` elements, it is not permitted for different `ramper` elements to control the same parameter of a given slave element. Additionally, parameters that `ramper` elements control must not be controlled by any `overlay` (but a `ramper` can control an `overlay`).

Note: There is a program to plot controller response curves bundled with the *Bmad* software (§1.2) called `controller_function_plot`. Documentation on this can be found at:

```
util_programs/controller_function_plot
```

4.45 RF_bend

An rf_bend is an RF cavity with the geometry of an sbend (§4.5). This element is currently considered to be experimental so please contact a *Bmad* maintainer if you want to use this type of element.

General rfcavity attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Offsets, pitches & tilt	5.6
Chamber wall	5.12	Overlapping Fields	5.18
Custom Attributes	3.9	Reference energy	5.5
Description strings	5.3	Superposition	8
Symplectify	6.7	Field Maps	5.16
Integration settings	6.4	Tracking & transfer map	6
Is_on	5.14	Wakes	5.20
Length	5.13		

See §15.52 for a full list of element attributes along with a their units.

Attributes specific to an rf_bend are:

```
! Bend-like attributes:
angle              = <Real>   ! Design bend angle. Dependent var (§5.1).
b_field            = <Real>   ! Design field strength (= P_0 g / q) (§5.1).
g                  = <Real>   ! Design bend strength (= 1/rho).
l                  = <Real>   ! "Length" of bend. See below.
l_arc              = <Real>   ! Arc length. For rbends only.
l_chord            = <Real>   ! Chord length. See §5.13.
l_sagitta                     ! Sagittal length. Dependent param (§5.1).
rho                = <Real>   ! Design bend radius. Dependent param (§5.1).
roll               = <Real>   ! See 5.6.
field_master       = <T/F>    ! See 5.2.

! RF-like attributes:
rf_frequency    = <Real>    ! Frequency
rf_wavelength               ! RF wavelength. Dependent parameter.
harmon          = <Real>    ! Harmonic number
harmon_master   = <Logic>   ! Is harmon or rf_frequency the dependent var with ref energy changes?
phi0            = <Real>    ! Cavity phase (rad/2pi).
phi0_multipass  = <Real>    ! Phase variation with multipass elements (rad/2pi) (§9).
```

Tracking through an rf_bend is limited to Runge Kutta like tracking methods (§6.1). The default tracking_method is runge_kutta. Fields must be specified using a grid field map (§5.16). The default field_calc is fieldmap (§6.5.2).

The geometry of the rf_bend is the same as an sbend. An rf_bend has a sector shape which is equivalent to e1 and e2 being zero for an sbend. Since the fields are specified using a grid field map, there are no fringe attributes to set (for any elements where a grid field is used, it is always assumed that the fringe fields are included as part of the grid field).

The relationship between harmon and $f_{\rm rf}$ is harmon $= t_0 \cdot f_{\rm rf}$ where t_0 is the reference 1-turn transit time.

The rf_wavelength is computed to be $c/f_{\rm rf}$. Notice that this is independent of the velocity of the reference particle. In particular, It will only be true that harmon times rf_wavelength will equal the

branch length if the reference particle travels at the speed of light.

4.46 RFcavity

An `rfcavity` is an RF cavity without acceleration generally used in a storage ring. The main difference between an `rfcavity` and an `lcavity` is that, unlike an `lcavity`, the reference energy (§17.2) through an `rfcavity` is constant.

General `rfcavity` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Mag & Elec multipoles	5.15
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Overlapping Fields	5.18
Description strings	5.3	Reference energy	5.5
Field autoscaling	5.19	RF Couplers	5.17
Fringe Fields	5.21	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Field Maps	5.16
Is_on	5.14	Tracking & transfer map	6
Length	5.13	Wakes	5.20

See §15.51 for a full list of element attributes along with a their units.

Attributes specific to an `rfcavity` are:

```
rf_wavelength               ! RF wavelength. Dependent parameter.
rf_frequency    = <Real>    ! Frequency
harmon          = <Real>    ! Harmonic number
harmon_master   = <Logic>   ! Is harmon or rf_frequency the dependent var with ref energy changes?
voltage         = <Real>    ! Cavity voltage
phi0            = <Real>    ! Cavity phase (rad/2pi).
phi0_multipass  = <Real>    ! Phase variation with multipass (rad/2pi).
phi0_autoscale  = <Real>    ! Set by Bmad if autoscaling is turned on (rad/2pi). (§9).
gradient        = <Real>    ! Accelerating gradient (V/m). Dependent attribute (§5.1).
longitudinal_mode = <Int>   ! Longitudinal mode. Default is 0. May be 0 or 1.
```

The integrated energy kick felt by a particle, assuming no phase slippage, is

$$\texttt{dE = -e_charge * } r_q \texttt{ * voltage * sin(2}\pi\texttt{ * (}\phi_{\texttt{t}}\texttt{ - }\phi_{\texttt{ref}}\texttt{))}$$

where

$$\phi_{\texttt{ref}} \texttt{ = phi0 + phi0_multipass}$$

and ϕ_t is the part of the phase due to when the particle arrives at the cavity and depends upon whether `absolute time tracking` or `relative time tracking` is being used as discussed in §26.1.

In the above equation r_q is the relative charge between the reference particle (set by the `parameter[particle]` parameter in a lattice file) and the particle being tracked through the cavity. For example, if the reference particle and the tracked particle are the same, r_q is unity independent of the type of particle tracked.

The correspondence between the *Bmad* `phi0` attribute and the `lag` attribute of *MAD* is

```
phi0 = mad + 0.5
```

`phi0_multipass` is only to be used to shift the phase with respect to a `multipass` lord. See §9 for more details. `e_charge` is the magnitude of the charge on an electron (Table 3.2). Notice that the energy kick is independent of the sign of the charge of the particle

The phi0_autoscale and field_autoscale parameters are calculated by *Bmad*'s auto-scale module to adjust the RF amplitude and phase so that during tracking, the energy gain of a particle on the reference orb is what the user expects based upon the set set voltage and phase. See Section §5.19 for more details. In particular, the gradient and phase are adjusted by

gradient (tracking) = gradient * field_autoscale
$\phi_{\texttt{ref}}$ (tracking) = $\phi_{\texttt{ref}}$ + phi0_autoscale

These autoscale shifts are not used when an element's tracking_method is bmad_standard since, in this case, the tracking algorithm is such that no adjustments are needed. The calculation of autoscale values can be toggled on/off by setting the autoscale_phase and autoscale_amplitude logicals.

Note: Zero phase for ϕ_{ref} corresponds to the stable fixed point above transition.

Note: Phi0 is not to be confused with the synchronous phase. The synchronous phase is the phase of the particle as it passes through the cavity with respect to the RF waveform. The synchronous phase is not something that is set by the User but rather is established by the balance between the energy gain of the particle as it goes through the cavities in the ring versus the energy lost to synchrotron radiation. In fact, for a ring with a single cavity, the synchronous phase is independent of phi0. Changing phi0 in such a situation will result in the closed orbit phase space z to vary in lock step.

If harmon is non–zero the rf_frequency is calculated by

rf_frequency = harmon * c_light * beta0 / L_lattice

where L_lattice is the total lattice length and beta0 is the velocity of the reference particle at the start of the lattice. After the lattice has been read in, rf_frequency will be the independent variable (§5.1).

Couplers (§5.17) and HOM wakes (§5.20) can be modeled. In addition, if a field map is specified (§5.16), tracking using an integrator is possible.

If a field map is specified (§5.16), tracking using an integrator is possible. A field map is only used for runge_kutta, fixed_step_runge_kutta, and symp_lie_bmad tracking (§6.1). Only the fundamental mode has an analytical formula for the symplectic tracking. With cavity_type set to standing_wave, the longitudinal mode is set by the longitudinal_mode parameter. The possible values are 0 or 1 and the default setting is 0.

The cavity_type is the type of cavity being simulated. Possible settings are:

```
ptc_standard
standing_wave    ! Default
traveling_wave
```

The cavity_type switch is ignored if a field map is used.

Example:

```
rf1: rfcav, l = 4.5, harmon = 1281, voltage = 5e6
```

The relationship between harmon and f_{rf} is harmon $= t_0 \cdot f_{\text{rf}}$ where t_0 is the reference 1-turn transit time.

The rf_wavelength is computed to be c/f_{rf}. Notice that this is independent of the velocity of the reference particle. In particular, It will only be true that harmon times rf_wavelength will equal the branch length if the reference particle travels at the speed of light.

Note: Multipoles 5.15 do not vary in time.

4.47 Sad_Mult

A sad_mult element is used to emulate a SAD[SAD] MULT element with the restriction that sad_mult elements cannot be used to simulate RF cavities or bends.

General sample attributes are:

Attribute Class	Section	*Attribute Class*	Section
a*n*, b*n* multipoles	5.15	Length	5.13
Aperture limits	5.8	Offsets, pitches & tilt	5.6
Chamber wall	5.12	Reference energy	5.5
Custom Attributes	3.9	Superposition	8
Description strings	5.3	Tracking & transfer map	6
Fringe Fields	5.21		

See §15.54 for a full list of element attributes along with a their units.

Attributes specific to an sad_mult element are:

```
bs_field        = <Real>    ! Solenoid field. SAD equivalent: BZ.
ks              = <Real>    ! Solenoid strength.
e1, e2          = <Real>    ! Bend face angles.
eps_step_scale  = <Real>    ! Step size scale. Default = 1. SAD equivalent: EPS.
fq1, fq2        = <Real>    ! Quadrupole fringe parameters. SAD equivalents: F1, F2.
fb1, fb2        = <Real>    ! Bend edge fringe parameters. SAD equivalents: FB1, FB2.
x_offset_mult   = <Real>    ! Mult component offset. SAD equivalent: DX.
y_offset_mult   = <Real>    ! Mult component offset. SAD equivalent: DY.
fringe_type     = <Switch>  ! Type of fringe. SAD equivalent: DISFRIN.
fringe_at       = <Switch>  ! Where fringe is applied. SAD equivalent: FRINGE.
```

One difference between SAD and *Bmad* is that SAD defines the solenoid field by what are essentially a set of marker elements so that the solenoid field at a SAD mult element is not explicitly declared in the mult element definition. *Bmad*, on the other hand, requires a sad_mult element to explicitly declare the solenoid parameters.

Another difference between SAD and *Bmad* is that, within a solenoid, the reference trajectory is aligned with the solenoid axis (and not aligned with the axis of the elements within the solenoid region).

The SAD mult element uses normal Kn and skew KSn multipole components. The *Bmad* sad_mult element used normal an and skew bn multipole components. As can be seen from the equations in §18.1, there is a factor of $n!$ between the two representations.

The a0 or b0 multipole moments give a dipole kick (just like a kicker element). The face angles e1 and e2 are used with the dipole kick in calculating fringe effects.

The fq1 and fq2 parameters are used to specify the quadrupolar "soft" edge fringe. See §19.6 for more details.

The fringe_at and fringe_type settings determine if the fringe field is used in tracking. See Sec §5.21 for the translation between these two switches and the fringe and disfrin switches of SAD.

The x_offset_mult and y_offset_mult orients the non-solenoid components of the field while leaving the solenoid component unshifted.

Unlike other elements, the ds_step and num_steps attributes (§6.4) of a sad_mult are dependent attributes (§5.1) and are not directly settable. Rather these attributes are calculated using SAD's own

algorithm for setting the step size. To vary the calculated step size for a single `sad_mult` element, the attribute `eps_step_scale` may be set. To vary the step size for all `sad_mult` elements, the global parameter `bmad_com[sad_eps_scale]` (§11.4) may be set. The default values for these parameters are:

```
eps_step_scale          = 1
bmad_com[sad_eps_scale] = 5e-3
```

SAD conventions to be aware of when comparing SAD to Bmad:

- A SAD `rotate` or `chi3` rotation is opposite to a *Bmad* `tilt`
- SAD element offsets (`dx`, `dy`, `dz`) are with respect to the entrance end of the element as opposed to *Bmad*'s convention of referencing to the element center.
- The *Bmad* `sad_mult` element does not have any attributes corresponding to the following SAD `mult` element attributes:
  `angle, harmon, freq, phi, dphi, volt, dvolt`

Example:

```
qs1: sad_mult, l = 0.1, fringe_type = full, b2 = 0.6 / factorial(2)
```

4.48 Sample

A `sample` element is used to simulate a material sample which is illuminated by x-rays.

General `sample` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Offsets, pitches & tilt	5.6
Chamber wall	5.12	Reference energy	5.5
Custom Attributes	3.9	Surface Properties	5.11
Description strings	5.3	Superposition	8
Integration settings	6.4	Tracking & transfer map	6
Length	5.13		

See §15.55 for a full list of element attributes along with a their units.

This element is in development.

Attributes specific to an `solenoid` element are:

```
mode       = <Switch> ! Reflection or transmission.
material   = <type>   ! Type of material. §5.9
```

The `mode` parameter can be set to:

```
reflection
transmission
```

With `mode` set to `reflection`, photons will be back scattered from the sample surface isotropically. In this case the material properties will not matter. Additionally, a `patch` (§4.41) element will be needed after the `sample` element to properly reorient the reference orbit.

With `mode` set to `transmission`, photons will be transmitted through the sample. In this case `material` will be used to determine the attenuation and phase shift of the photons.

Example:

```
formula409: sample, x_limit = 10e-3, y_limit = 20e-3, mode = reflection
```

4.49 Sextupole

A `sextupole` is a magnetic element with a quadratic field dependence with transverse offset (§18.1).

General `sextupole` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Mag & Elec multipoles	5.15
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Overlapping Fields	5.18
Description strings	5.3	Reference energy	5.5
Fringe Fields	5.21	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Field Maps	5.16
Is_on	5.14	Tracking & transfer map	6
Length	5.13		

See §15.56 for a full list of element attributes along with a their units.

Attributes specific to an `sextupole` element are:

```
k2          = <Real>   ! Sextupole strength.
b2_gradient = <Real>   ! Field strength. (§5.1).
field_master = <T/F>    ! See §5.2.
```

The normalized sextupole `k2` strength is related to the unnormalized `b2_gradient` field strength through Eq. (18.3).

The `bmad_standard` calculation treats a sextupole using a kick–drift–kick model.

If the `tilt` attribute is present without a value then a value of $\pi/6$ is used. Example:

```
q03w: sext, l = 0.6, k2 = 0.3, tilt  ! same as tilt = pi/6
```

4.50 Sol_Quad

A sol_quad is a combination solenoid/quadrupole. Alternatively, the sad_mult element can also be used. The advantage of the sad_mult element is that it can simulate a quadrupole field that is canted with respect to the solenoid field.

General sol_quad attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Mag & Elec multipoles	5.15
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Overlapping Fields	5.18
Description strings	5.3	Reference energy	5.5
Fringe Fields	5.21	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Field Maps	5.16
Is_on	5.14	Tracking & transfer map	6
Length	5.13		

See §15.57 for a full list of element attributes along with a their units.

Attributes specific to a sol_quad element are:

```
k1           = <Real>    ! Quadrupole strength.
ks           = <Real>    ! Solenoid strength.
bs_field     = <Real>    ! Solenoid Field strength.
b1_gradient  = <Real>    ! Quadrupole Field strength.
field_master = <T/F>     ! See §5.2.
```

The normalized quadrupole k1 and solenoid ks field strengths are related to the unnormalized b1_gradient and bs_field field strengths through Eq. (18.3).

Example:

```
sq02: sol_quad, l = 2.6, k1 = 0.632, ks = 1.5e-9*parameter[p0c]
```

4.51 Solenoid

A `solenoid` is an element with a longitudinal magnetic field.

General `solenoid` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Mag & Elec multipoles	5.15
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Overlapping Fields	5.18
Description strings	5.3	Reference energy	5.5
Fringe Fields	5.21	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Field Maps	5.16
Is_on	5.14	Tracking & transfer map	6
Length	5.13		

See §15.58 for a full list of element attributes along with a their units.

Attributes specific to an `solenoid` element are:

```
ks           = <Real>  ! Normalized solenoid strength.
bs_field     = <Real>  ! Solenoid field strength.
field_master = <T/F>   ! See §5.2.
l_soft_edge  = <Real>  ! For modeling a ‘‘soft" fringe.
r_solenoid   = <Real>  ! Solenoid radius.
```

The `ks` and `bs_field` parameters are the normalized and unnormalized solenoid strengths related through Eq. (18.3).

The `bmad_standard` tracking model (§6.1) uses a "hard edge" model where the field goes from zero to full strength. Example:

```
cleo_sol: solenoid, l = 2.6, ks = 1.5e-9 * parameter[p0c]
```

"Soft edge" end fields may be simulated by using Runge-Kutta tracking and setting the `field_calc` parameter of the element to `soft_edge`. Equations for the soft edge model are taken from Derby and Olbert [Derby09]. The equations used are for the exact ideal azimuthally symmetric solenoid model (not the near-axis approximation model). In this case `ks` and `bs_field` are the field of an infinite pipe with the same current density (equal to $\mu_0 \, n \, I$ in the notation of Derby and Olbert). Example:

```
soft_sol: solenoid, l = 1.0, field_calc = soft_edge, l_soft_edge = 0.5,
            r_solenoid = 0.1, tracking_method = runge_kutta, mat6_calc_method = tracking
```

Here the solenoid is modeled as a perfect current carrying cylinder with a length of 0.5 meters and a cylinder radius of 0.1 meters. The length of the element, 1.0 meters, must be greater than `l_soft_edge` since particles need to be tracked through the non-zero fringe field that extends outside of the cylinder. In fact, since the field is always finite everywhere, to the extent that the field is nonzero at the edges of the element places a bound on the accuracy of the simulation. Note that `bmad_standard` tracking will always ignore the setting of `field_calc`. That is, with `bmad_standard` tracking the field always extends to the edges of the element and the value of `l_soft_edge` is ignored.

4.52 Taylor

A `taylor` is a Taylor map (§25.1) that maps the input orbital phase space and possibly spin coordinates of a particle at the entrance end of the element to the output orbital and spin coordinates at the exit end of the element. This can be used in place of the *MAD* `matrix` element.

General `taylor` attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Offsets & tilt	5.6
Custom Attributes	3.9	Reference energy	5.5
Description strings	5.3	Superposition	8
Is_on	5.14	Symplectify	6.7
Length	5.13	Tracking & transfer map	6

See §15.59 for a full list of element attributes along with a their units.

Attributes specific to a `taylor` element are:

```
ref_orbit = (<x>, <px>, <y>, <py>, <z>, <pz>)     ! Reference orbit.
x_ref  = <Real>                                   ! $x$ reference orbit component.
px_ref = <Real>                                   ! $p_x$ reference orbit component.
y_ref  = <Real>                                   ! $y$ reference orbit component.
py_ref = <Real>                                   ! $p_y$ reference orbit component.
z_ref  = <Real>                                   ! $z$ reference orbit component.
pz_ref = <Real>                                   ! $p_z$ reference orbit component.
{<out>: <coef>, <e1> <e2> <e3> <e4> <e5> <e6>}    ! Taylor term. First form.
{<out>: <coef> | <n1> <n2> ...}                   ! Taylor term. Second form.
tt<out><n1><n2>...  = <Coef>                      ! Taylor term. Third form.
delta_ref_time = <Real>                           ! Change in the reference time.
delta_e_ref = <Real>                              ! Change in the reference energy.
```

For historical reasons, there are three different forms that can be used to specify a taylor term. Notice that the first form (above) uses a comma "," to separate the `<coef>` from `<e1>`, while the second form uses a vertical bar "|" to separate `<coef>` from `<n1>`.

The orbital (x, p_x, y, p_y, z, p_z) part of the Taylor map, $\mathcal{M}$, maps input orbital coordinates $\mathbf{r}(\text{in})$ to the output orbital coordinates $\mathbf{r}(\text{out})$

$$\mathbf{r}(\text{out}) = \boldsymbol{\mathcal{M}}(\mathbf{r}(\text{in})) \tag{4.26}$$

Notice that Stern-Gerlach effects are ignored so that the output coordinates are independent of the spin. $\boldsymbol{\mathcal{M}}$ has six components $\mathcal{M}_j$ one for each output coordinate $r_j(\text{out})$

$$r_j(\text{out}) = \mathcal{M}_j(\mathbf{r}(\text{in})) \tag{4.27}$$

Each $\mathcal{M}_j$ is made up of a number of terms

$$\mathcal{M}_j = \sum_{k=1}^{N_j} M_{jk} \tag{4.28}$$

and each term M_{jk} is a polynomial in the input orbital coordinates with respect to the reference orbit.

$$M_{jk}(\mathbf{r}(\text{in})) = C_{jk} \cdot \Pi_{i=1}^{6} \delta r_i^{e_{ijk}} \tag{4.29}$$

where C_{jk} is the coefficient for the term, the e_{ijk} are integer exponents, and $\delta\mathbf{r} = \mathbf{r}(\text{in}) - \mathbf{r}_{\text{ref}}$ with $\mathbf{r}_{\text{ref}}$ with being the reference orbit.

A term in a Taylor map can be specified by one of three forms as shown above. The first form is

```
{<out>: <coef>, <e1> <e2> <e3> <e4> <e5> <e6>}
```

`<Out>` is an integer in the range 1 to 6 corresponding to the index j in Eq. (4.29) (`<out>` = 1 for x, etc.). `<coef>` corresponds to C_{jk} in Eq. (4.29), and `<e1>`, `<e2>`, `<e3>`, `<e4>`, `<e5>`, and `<e6>` correspond to $e_{ijk}, i = 1 \ldots 6$. For example, the Taylor map

$$p_y(\text{out}) = 0.9 \cdot \delta x + 2.73 \cdot \delta y^2(\text{in})\, \delta p_z(\text{in}) + \ldots \tag{4.30}$$

would be written as

```
{4: 0.9, 1 0 0 0 0 0}, {4: 2.73, 0 0 2 0 0 1}, ...
```

The second form for specifying a Taylor term uses the syntax:

```
{<out>: <coef> | <n1> <n2> ...}
```

The set of integers `<n1>`, `<n2>` each must be between 1 and 6 inclusive. The value of the i^{th} exponent e_{ijk} in Eq. (4.29) is equal to the number of integers that are equal to i. For example, the above Taylor map would be written using the second form as

```
{4: 0.9 | 1}, {4: 2.73 | 336}, ...
```

Notice that with the second form, spaces between exponent integers is optional.

The third form is like the second form and has the syntax:

```
tt<out><n1><n2>...  = <Coef>                    ! Taylor term. Third form.
```

For example, the Taylor map above would be written using the third form as:

```
tt41 = 0.9, tt4336 = 2.73, ...
```

The spin (§24.1) part of the transport map $\boldsymbol{\mathcal{Q}}$ (§25.2) gives the spin rotation quaternion $\mathbf{q}$ (§24.2) as a function of input orbital coordinates (the form of the T-BMT equation assures that $\boldsymbol{\mathcal{Q}}$ cannot depend upon the spin coordinates):

$$\mathbf{q} = \boldsymbol{\mathcal{Q}}(\mathbf{r}(\text{in})) \tag{4.31}$$

$\mathbf{q}$ has four components and in analogy to Eq. (4.27) one writes

$$q_j = \mathcal{Q}_j(\mathbf{r}(\text{in})) \tag{4.32}$$

Each $\mathcal{Q}_j$ is made up of a number of terms

$$\mathcal{Q}_j = \sum_{k=1}^{N_j} Q_{jk} \tag{4.33}$$

and each term Q_{jk} is a polynomial in the input orbital coordinates with respect to the reference orbit.

$$Q_{jk}(\mathbf{r}(\text{in})) = C_{jk} \cdot \Pi_{i=1}^{6}\, \delta r_i^{e_{ijk}} \tag{4.34}$$

Rather than using an integer index, the four components of a quaternion are labeled (`S1`, `Sx`, `Sy`, `Sz`). The syntax for the spin part uses the three forms as described above. For example

```
{Sx: 0.43 | 13 }          ! or
{Sx: 0.43, 1 0 1 0 0 0}   ! or
ttSx13 = 0.43
```

is equivalent to the term

$$S_x = 0.43 \cdot \delta x(\text{in})\, \delta y(\text{in}) \tag{4.35}$$

By default, a `taylor` element starts out with the unit phase space map. That is, a `taylor` element starts with the following 6 terms

```
{1: 1.0, 1 0 0 0 0 0}, {2: 1.0, 0 1 0 0 0 0},
{3: 1.0, 0 0 1 0 0 0}, {4: 1.0, 0 0 0 1 0 0}
{5: 1.0, 0 0 0 0 1 0}, {6: 1.0, 0 0 0 0 0 1}
```

Which is equivalent to

```
{1: 1.0 | 1}, {2: 1.0 | 2}, {3: 1.0 | 3}
{4: 1.0 | 4}, {5: 1.0 | 5}, {6: 1.0 | 6}
```

If there are no `Sx` spin terms are present, the `Sx` quaternion component will always evaluate to zero. This is equivalent to a single term {Sx: 0.0 |}. Similarly for the `Sy` and `Sz` components. If no `S1` term is present, it is considered an error if any `Sx`, `Sy`, or `Sz` term is present. If no `S1`, `Sx`, `Sy`, nor `Sz` spin terms are present, `S1` component will be given a default term of {S1: 1.0 |}. Thus, if no spin terms are present at all, the spin map will be the unit map.

The `ref_orbit` attribute specifies the phase space (x, px, y, py, z, pz) reference orbit at the start of the element used to construct the Taylor map. Alternatively, the individual components of the reference orbit may be specified by using the attributes `x_ref`, `px_ref`, `y_ref`, `py_ref`, `z_ref`, or `pz_ref`.

Note: when converting the map from *Bmad* to PTC (§29), the *Bmad*/PTC interface code will convert from Bmad phase space coordinates to PTC phase space coordinates and will convert the map to using the reference orbit as the map zero orbit. This does not affect tracking but will affect map analysis.

A term in a `taylor` element will override any previous term with the same `out` and `e1` through `e6` indexes. For example the term:

```
my_tlr: Taylor, {1: 4.5, 1 0 0 0 0 0}
```

will override the default `{1:  1.0, 1 0 0 0 0 0}` term.

The `l` length attribute of a `taylor` element does not affect phase space coordinates but will affect the longitudinal s position of succeeding elements and will affect the time it takes a particle to track through the element The calculation involves first calculating the change in reference time which is the time a particle with the reference energy would take to transverse the element. Next, Eq. (17.3) is used with the change in the phase space z coordinate to calculate the time a particle takes to traverse the element.

The time a particle takes to track through a `taylor` element can also be controlled by setting the `delta_ref_time` attribute which sets the travel time for the reference particle. `delta_ref_time` is a dependent attribute so that if both `l` and `delta_ref_time` are set, the value of `delta_ref_time` will be modified by *Bmad* to correspond to the setting of `l`.

The `delta_e_ref` attribute can be used to modify the reference energy at the exit end of the `taylor` element. The phase space transport is completely determined by the Taylor map and is independent of `delta_e_ref`. For example, with a unit Taylor map, the phase space coordinates p_x and p_y constant through the element independent of `delta_e_ref`. However, a finite `delta_e_ref` will modify the reference momentum P_0 and hence through Eq. (17.2) will affect the transport downstream of the `Taylor` element. This behavior is in contrast to how `delta_e_ref` is handled in a `patch` element. In a `patch` element, the transformation used when `delta_e_ref` is non-zero is to hold as constant the actual transverse momenta P_x and P_y and then p_x and p_y are modified using Eq. (17.2).

A `taylor` element that is "turned off" (`is_on` attribute set to False), is considered to be like a `marker` element. That is, the orbit and Twiss parameters are unchanged when tracking through a `taylor` element that is turned off.

Example `taylor` element definitions:

```
mtlr: Taylor, {4:  2.7, 0 0 2 0 0 1}, {2:  1.9 | 1 1 2},
            {S1: 0.43 | 2 }, ...,
            ref_orbit = (0.01, 0.003, 0.002, 0.001, 0.0, 0.2)
t_unit: taylor {s1: 1 | }  ! This is the identity spin/orbital map.
```

And

Note: When tracking a particle's spin through a map, the quaternion used to rotate the spin is always normalized to one so that the magnitude of the spin will be invariant.

Note: Tracking through a `taylor` elements using `symp_lie_ptc` is the same as tracking with the `taylor` tracking method. That is, the Taylor map is simply evaluated and no effort at symplectification is done. Furthermore, evaluating the Taylor map of a `taylor` element using the `taylor` method is faster than evaluation using `symp_lie_ptc`. Thus the `taylor` tracking method should always be used with `taylor` elements.

4.53 Thick_Multipole

A thick_multipole element is like a sextupole or octupole element except that the thick_multipole does not have a K2 sextupole like parameter nor a K3 octupole like parameter. Rather, thick_multipoles, like sextupole or octupole elements, have a0, a1, a2, etc. and b0, b1, b2, etc. multipoles (§18.1). In terms of tracking, given equivalent multipole values, thick_multipoles are indistinguishable from sextupoles or octupoles. thick_multipole elements are useful for differentiating elements that only have higher order multipole moments.

General thick_multipole attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Mag & Elec multipoles	5.15
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Overlapping Fields	5.18
Description strings	5.3	Reference energy	5.5
Fringe Fields	5.21	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Field Maps	5.16
Is_on	5.14	Tracking & transfer map	6
Length	5.13		

See §15.60 for a full list of element attributes along with a their units.

Example:

```
tm1: thick_multipole, l = 4.5, tilt, x_pitch = 0.34, a7 = 1.23e3, b8 = 7.54e5
```

4.54 Wiggler and Undulator

A wiggler or undulator element is basically a periodic array of alternating bends. The difference between wigglers and undulators is in the x-ray emission spectrum. Charged particle tracking will be the same.

Henceforth, the term "wiggler" will denote either a wiggler or undulator

General wiggler attributes are:

Attribute Class	Section	*Attribute Class*	Section
Aperture limits	5.8	Mag & Elec multipoles	5.15
Chamber wall	5.12	Offsets, pitches & tilt	5.6
Custom Attributes	3.9	Overlapping Fields	5.18
Description strings	5.3	Reference energy	5.5
Fringe Fields	5.21	Superposition	8
Hkick & Vkick	5.7	Symplectify	6.7
Integration settings	6.4	Field Maps	5.16
Is_on	5.14	Tracking & transfer map	6
Length	5.13		

See §15.61 for a full list of element attributes along with a their units.

There are three types of wigglers. Wigglers that are described using a magnetic field map are called "map type" and are discussed in §4.54.2. Wigglers that are described assuming a periodic field are called "periodic type" and are described in §4.54.1. The third type of wiggler has a custom field. The different wiggler types are distinguished by the setting of the element's field_calc parameter as discussed in section §6.5.2. For example:

```
wig1: wiggler, l = 1.6, field_calc = fieldmap, ...
```

In this example wig1 is a map type wiggler.

Attributes specific to wiggler and undulator elements are:

```
b_max      = <Real>  ! Maximum magnetic field (in T) on the wiggler centerline.
l_period   = <Real>  ! Length over which field vector returns to the same orientation.
n_period   = <Real>  ! The number of periods. Dependent attribute (§5.1).
l_pole     = <Real>  ! Wiggler pole length. DEPRECATED. USE L_PERIOD INSTEAD.
n_pole     = <Real>  ! The number of poles. DEPRECATED. USE N_PERIOD INSTEAD.
polarity   = <Real>  ! For scaling the field.
kx         = <Real>  ! Planar wiggler horizontal wave number.
delta_ref_time_user_set = <Logic>  ! delta_ref_time set in lattice file.
k1x                  ! Planar wiggler horizontal defocusing strength. Dep attribute (§5.1).
k1y                  ! Planar wiggler vertical focusing strength. Dep attribute (§5.1).
g_max                ! Maximum bending strength. Dependent attribute.
osc_amplitude        ! Amplitude of the particle oscillations. Dependent attribute.
```

The polarity value is used to scale the magnetic field. By default, polarity has a value of 1.0. Example:

```
wig1: wiggler, l = 1.6, polarity = -1, cartesian_map = {...}
```

In this example the wiggler field is defined by a Cartesian map (§5.16.2) and the field is reversed from what it would be with polarity set to 1.

4.54.1 Periodic Type Wigglers

`Periodic type` wigglers are modeled assuming the field is periodic longitudinally. `Periodic type` wigglers have their `field_calc` parameter set to one of

```
planar_model     ! Default
helical_model
```

For historic purposes, if there is no fieldmap defined for the element (that is, it is not a `map type` wiggler), and `field_calc` is not set, then `field_calc` will default to `planar_model`.

Example:

```
wig2: wiggler, l = 1.6, b_max = 2.1, n_period = 8
```

This defines a periodic type wiggler with `field_type` defaulting to `planar_model`.

For the `planar_model`, wigglers use a simplified model where the wiggler has field components

$$\begin{aligned} B_x &= -\text{b_max}\,\frac{k_x}{k_y}\,\sin(k_x\,x)\,\sinh(k_y\,y)\,\cos(k_z\,z+\phi_z) \\ B_y &= \quad\text{b_max}\,\cos(k_x\,x)\,\cosh(k_y\,y)\,\cos(k_z\,z+\phi_z) \\ B_z &= -\text{b_max}\,\frac{k_z}{k_y}\,\cos(k_x\,x)\,\sinh(k_y\,y)\,\sin(k_z\,z+\phi_z) \end{aligned} \tag{4.36}$$

with $k_y^2 = k_x^2 + k_z^2$. Here z is the distance from the beginning of the wiggler, the input parameter `b_max` is the maximum field on the centerline, and k_z is given in terms of the period length (`l_period`) by

$$k_z = \frac{2\pi}{l_{\text{period}}} \tag{4.37}$$

The phase ϕ_z is chosen so that B_y is symmetric about the center of the wiggler

$$\phi_z = \frac{-k_z\,L}{2} \tag{4.38}$$

Note: Originally k_z was calculated using `l_pole` — the length of a pole — with the period length being twice the pole length. When the helical model option was introduced this became problematical since the period of a helical wiggler could be either 2 or 4 times the pole length depending upon the geometry. As a result, using the pole length was deprecated and instead the period length or number should be used.

The `helical_model` for the field is

$$\begin{aligned} B_x &= -\text{b_max}\,\cosh(k_z\,x)\,\sin(k_z\,z+\phi_z) \\ B_y &= \quad\text{b_max}\,\cosh(k_z\,y)\,\cos(k_z\,z+\phi_z) \\ B_z &= -\text{b_max}\,[\sinh(k_z\,x)\,\cos(k_z\,z+\phi_z) + \sinh(k_z\,y)\,\sin(k_z\,z+\phi_z)] \end{aligned} \tag{4.39}$$

With `field_calc` set to `planar_model`, and with `bmad_standard` tracking (§6), the horizontal and vertical focusing is assumed small. The vertical motion is modeled as a combination focusing quadrupole and focusing octupole giving a kick (modified from [Corbett99])

$$\frac{dp_y}{dz} = k_{1y}\left(y + \frac{2}{3}\,k_y^2\,y^3\right) \tag{4.40}$$

where

$$g_{\max} = \frac{e\,B_{\max}}{P_0\,(1+p_z)} \tag{4.41}$$

$$k_{1y} = \frac{-k_y^2}{2\,k_z^2}\,g_{\max}^2 \tag{4.42}$$

with `k1y` (a dependent element attribute) being the linear focusing constant.

The averaged horizontal motion is

$$\frac{dp_x}{dz} = k_{1x}\, x \tag{4.43}$$

with

$$k_{1x} = \frac{k_x^2}{2\, k_z^2}\, g_{\max}^2 \tag{4.44}$$

With `field_calc` set to `helical_model`, and with `bmad_standard` tracking, the transport in the vertical and horizontal planes is the same as with the transport in the vertical plane with `planar_model` (Eq. (4.40)).

While `bmad_standard` tracking uses an averaged trajectory, the actual trajectory has oscillations that look like

$$x = A\, \cos(k_z\, z) \tag{4.45}$$

with the amplitude A given by

$$A = \frac{g_{\max}}{k_z^2} \tag{4.46}$$

The value of A, computed for an on-energy ($p_z = 0$) particle, is calculated and stored in the dependent parameter `osc_amplitude`.

With `field_calc` set to `planar_model` and `bmad_standard` tracking, the phase ϕ_z in Eqs. (4.38) is irrelevant. When the tracking involves Taylor maps and symplectic integration, the choice of phase is such that, with an integer number of periods, a particle that enters the wiggler on-axis will leave the wiggler on-axis provided there is an integer number of periods. Notice that with `field_calc` set to `helical_model` it is not possible to set the phase so that a particle that enters the wiggler on-axis will leave the wiggler on-axis.

When using a tracking through a periodic wiggler with a tracking method that integrates through the magnetic field (§6.4), The magnetic field is approximated using a single wiggler `term` as if the wiggler were a `map type` wiggler. This wiggler model has unphysical end effects and will give results that are different from the results obtained when using the `bmad_standard` tracking method.

With `delta_ref_time_user_set` set to `False` (the default), *Bmad* will calculate the change in reference time through the element, `delta_ref_time`, by tracking a particle through the wiggler starting the particle with zero phase-space coordinates (§5.5). Sometimes this calculation does not produce a value for `delta_ref_time` that is appropriate for a simulation. In this case, `delta_ref_time_user_set` can be set to `True` and `delta_ref_time` can then be set in the lattice file.

4.54.2 Map Type Wigglers

`Map type` wigglers are modeled using a field map as described in section §5.16. `Map type` wigglers have their `field_calc` parameter set to `fieldmap`. Note: For historic reasons, unlike other types of elements, `field_calc` will default to `fieldmap` if there is a field map present in a `wiggler`.

Unlike `periodic type` wigglers, the `b_max` attribute for a `map type wiggler` is a dependent attribute and is set by *Bmad* to be the maximum field on-axis computed for `polarity` = 1.

Note: There is no `bmad_standard` tracking for a `map_type wiggler`.

4.54.3 Old Wiggler Cartesian Map Syntax

When the wiggler model was first developed, the only type of map that could be used for `map type` wigglers was a Cartesian map (§5.16). The syntax for specifying this Cartesian map was different from what it is currently. The old syntax for a Cartesian map term was:

$$\text{term(i)} = \{C, k_x, k_y, k_z, \phi_z\} \tag{4.47}$$

Example:

```
wig1: wiggler, l = 1.6,
      term(1) = {0.03, 3.00, 4.00, 5.00, 0.63},
      term(2) = ...
```

The old syntax was limited to using the `cartesian_map y` family (§18.5) with $x_0 = y_0 = 0$. There was also a different normalization convention. The old style `hyper-y` form was

$$\begin{aligned} B_x &= -C \frac{k_x}{k_y} \sin(k_x x) \sinh(k_y y) \cos(k_z z + \phi_z) \\ B_y &= \quad C \quad \cos(k_x x) \cosh(k_y y) \cos(k_z z + \phi_z) \qquad \text{! Old style} \\ B_s &= -C \frac{k_z}{k_y} \cos(k_x x) \sinh(k_y y) \sin(k_z z + \phi_z) \\ &\text{with } k_y^2 = k_x^2 + k_z^2 \ . \end{aligned} \tag{4.48}$$

The old style `hyper-xy` form was

$$\begin{aligned} B_x &= \quad C \frac{k_x}{k_y} \sinh(k_x x) \sinh(k_y y) \cos(k_z z + \phi_z) \\ B_y &= \quad C \quad \cosh(k_x x) \cosh(k_y y) \cos(k_z z + \phi_z) \qquad \text{! Old style} \\ B_s &= -C \frac{k_z}{k_y} \cosh(k_x x) \sinh(k_y y) \sin(k_z z + \phi_z) \\ &\text{with } k_y^2 = k_z^2 - k_x^2 \ , \end{aligned} \tag{4.49}$$

The old style `hyper_x` form was

$$\begin{aligned} B_x &= \quad C \frac{k_x}{k_y} \sinh(k_x x) \sin(k_y y) \cos(k_z z + \phi_z) \\ B_y &= \quad C \quad \cosh(k_x x) \cos(k_y y) \cos(k_z z + \phi_z) \qquad \text{! Old style} \\ B_s &= -C \frac{k_z}{k_y} \cosh(k_x x) \sin(k_y y) \sin(k_z z + \phi_z) \\ &\text{with } k_y^2 = k_x^2 - k_z^2 \ . \end{aligned} \tag{4.50}$$

The correspondence between C in the above equations and A in the new equations is given by comparing Eqs. (4.48), (4.49), and (4.50) with Eqs. (18.34).

When the cartesian_map construct was being developed, an intermediate hybrid syntax was used defined:

$$term(i) = \{A, k_x, k_y, k_z, x_0, y_0, \phi_z, \text{family}\} \tag{4.51}$$

The parameters here directly correspond to the `cartesian_map` forms (see Eqs. (18.31) through (18.36)).

For example, the old style syntax:

```
term(1) = {0.03*4/5, 3.00, 4.00, 5.00, 0.63 }    ! Old style
```

is equivalent to the hybrid syntax:

```
term(2) = {0.03, 3.00, 4.00, 5.00, 0, 0, 0.63, y}  ! Hybrid style
```

Note: When converting from the old or hybrid styles to the new syntax, the `field_calc` parameter must be set to `fieldmap`.

Chapter 5

Element Attributes

For a listing of element attributes for each type of element, see Chapter §15.

5.1 Dependent and Independent Attributes

For convenience, *Bmad* computes the values of some attributes based upon the values of other attributes. Some of these dependent variables are listed in Table 5.1. Also shown in Table 5.1 are the independent variables they are calculated from.

Element	*Independent Variables*	*Dependent Variables*
All elements	`ds_step`	`num_steps`
`BeamBeam`	`charge`, `sig_x`, `sig_y`, `e_tot`	`bbi_constant`
`Elseparator`	`hkick`, `vkick`, `gap`, `l`, `e_tot`	`e_field`, `voltage`
`Lcavity`	`gradient`, `l`	`e_loss`, `voltage`
`Rbend`, `Sbend`	`g`, `l`	`rho`, `angle`, `l_chord`
`Wiggler` (map type)	`term(i)`	`b_max`, `k1`, `rho`
`Wiggler` (periodic type)	`b_max`, `e_tot`	`k1`, `rho`

Table 5.1: Partial listing of dependent variables and the independent variables they are calculated from.

For electric and magnetic field strength parameters, the `field_master` parameter (§5.2) can be used to determine if the be normalized or unnormalized, values are dependent or independent.

5.2 Field_Master and Normalized Vs. Unnormalized Field Strengths

The `field_master` attribute of an element sets whether the element's normalized (normalized by the reference energy) field strengths or the unnormalized strengths are the independent variables (§5.1). See §18.1 for details as to how they are related. The setting of `field_master` also sets whether an element's magnetic multipoles (§5.15) are interpreted as normalized or unnormalized (electric multipoles are always treated as unnormalized).

Table 5.2 shows some normalized and unnormalized field strength attributes. The default value of

field_master for an element is False if there are no field values set in the lattice file for that element. If normalized field values are present then the default is also False. If there are unnormalized field values present then the default is True.

For example:
```
Q1: quadrupole, b1_gradient = 0   ! Field strengths are the independent variables
Q1: quadrupole, field_master = T  ! Same as above
Q2: quadrupole        ! Define Q2.
Q2[b1_gradient] = 0   ! Field strengths now the independent variables.
Q2[field_master] = T  ! Same as above.
```
Specifying both normalized and unnormalized strengths for a given element is not permitted. For example:
```
Q3: quadrupole, k1 = 0.6, bl_hkick = 37.5  ! NO. Not VALID.
```

Element	*Normalized*	*Unnormalized*
Sbend, Rbend	g	b_field
Sbend, Rbend	dg	db_field
Solenoid, Sol_quad	ks	bs_field
Quadrupole, Sol_quad, Sbend, Rbend	k1	b1_gradient
Sextupole, Sbend, Rbend	k2	b2_gradient
Octupole	k3	b3_gradient
HKicker, VKicker	kick	bl_kick
Most	hkick	bl_hkick
Most	vkick	bl_vkick

Table 5.2: Example normalized and unnormalized field strength attributes.

5.3 Type, Alias and Descrip Attributes

There are three string labels associated with any element:
```
type    = "<String>"
alias   = "<String>"
descrip = "<String>"
```
The type and alias attributes can be up to 40 characters in length and descrip can be up to 200 characters. If the attribute string does not contain a blank, comma, or a semicolon, the quotation marks may be omitted. Unlike the element name, these strings are not converted to upper case.

These labels have a number of purposes. For example, a *Bmad* based program that is used as an online machine model could associate a label string with a machine database element to facilitate machine control. Also these labels can be used with pattern matching when selecting elements (§3.6).

Example:
```
Q00W: Quad, type = "My Type", alias = Who_knows, &
                                  descrip = "Only the shadow knows"
```

5.4 Group, Overlay, and Ramper Element Syntax

The syntax for specifying group (§4.26), overlay (§4.40), or ramper (§4.44) elements are virtually identical and is discussed below. The name "controller" will be used to denote either a group, overlay, or ramper element.

Controller elements have a set of one or more "`variables`" that are used to control the values of attributes of other elements (called "`slave`" attributes).

There are two types of controllers in terms of how slave attribute values are computed — `Expression` based controllers and `knot` based controllers. `Expression` based controllers use mathematical expressions (§3.13) to define how slave attribute values are calculated based upon the values of the controller variables. `Knot` based controllers use an array of points (called "`knots`") with either linear or cubic spline interpolation (in particular, the cubic non-smoothing Akima spline[Akima70]) to determine the relationship between variables and slave attributes. With a knot based controller, the number of controller variables is restricted to be one.

The general syntax for an expression based `group` element is

```
name: GROUP = {ele1[attrib1]:exp1, ele2[attrib2]:exp2, ...},
        VAR = {var1, var2, ...}, var1 = init_val1,
        old_var1 = old_init_val1, ..., GANG = logical, ...
```

For an expression based `overlay` element, the syntax is identical except OVERLAY is substituted for GROUP and there are no old values to set:

```
name: OVERLAY = {ele1[attrib1]:exp1, ele2[attrib2]:exp2, ...},
        VAR = {var1, var2, ...}, var1 = init_val1, GANG = logical, ...
```

For an expression based `ramper` element, the syntax is identical to the `overlay` element except `RAMPER` is used in place of `OVERLAY` and there is no `GANG` attribute.

In the above, `Name` is the name of the controller element, `ele1`, `ele2`, ... are the elements whose attributes are to be controlled, `attrib1`, `attrib2`, etc. are the controlled attributes (called "slave" attributes), `var1`, `var2`, etc. are the control variables, and `exp1`, `exp2`, etc. are the arithmetical expressions that define the relationship between the variables and the slave attributes. Example:

```
gr1: group = {q1[k1]:-tan(a)*b, q2[tilt]:b^2}, var = {a, b}
```

The function `ran()` and `ran_gauss()` may be used in expressions with `ramper` elements but may not be used for other types of controllers. These functions can be useful with long-term tracking to simulate such things as ground noise or RF jitter. Example

```
rf_noise: ramper = {rfcavity::*[phi0]: 1e-4*ran_gauss()}, var = { }
```

Notice that the expression does not depend upon any variables and the variable list is empty. An empty variable list is permissible only if all slave expressions involve `ran()` or `ran_gauss()`. With an empty variable list, *Bmad* will define a variable for the `ramper` called `null` to serve as a place holder for bookkeeping purposes.

To define an expression based `controller` element, two lists are needed: One list defines the slave attributes along with the arithmetic expressions used for computing the value of the slave attributes. The other list defines the variables within the `controller` element that can be varied. In the above example, the variables are `a` and `b`, and the slave attributes are the `k1` attribute of element `q1` and the `tilt` attribute of element `q2`. The arithmetic expressions used for the control are `-tan(a)*b` and `b^2`.

The general syntax for For a `knot` based `group` element is:

```
name: GROUP = {ele1[attrib1]:{y_knot_points1}, ele2[attrib2]:{y_knot_points2}, ...},
            VAR = {var1}, X_KNOT = {x_knot_points}, INTERPOLATION = {type},
            var1 = init_val1, old_var1 = init_val_old1, ..., GANG = logical, ...
```

and the general syntax for a `knot` based `overlay` element is similar:

```
name: OVERLAY = {ele1[attrib1]:{y_knot_points1}, ele2[attrib2]:{y_knot_points2}, ...},
      VAR = {var1}, X_KNOT = {x_knot_points}, var1 = init_val1, ...,
      GANG = logical, INTERPOLATION = {type}, ...
```

For ramper elements, the syntax is similar to the overlay syntax except that OVERLAY is replaced by RAMPER and there is no GANG attribute. The x_knot array is combined with a given y_knot array to give a series of N knot points where N is the size of the x_knot and y_knot arrays. As such, the size of all knot arrays must be the same with the additional restriction that the x_knot array must have values that are strictly ascending. Example:

```
ov: overlay = {q1[k1]:{-0.12, 0.23, 0.53}, q2[k1], q3[k1], q4[k1]:{0.37, 0.12, 0.92}},
      var = time, x_knot = {0.0, 0.7, 1.0}
```

K In this example there are three knot points for each slave parameter. For the slave parameter q4[k1] the knot points are $(0.0, 0.37)$, $(0.7, 0.12)$, and $(1.0, 0.92)$. If a y_knot_points array is not present for a particular slave attribute, the knot array of the previous slave attribute is used. Thus, in the above example, the y_knot array used for q2[k1] and q3[k1] are the same as the y_knot array for q1[k1].

The method used to interpolate between the knot points is determined by the setting of the interpolate parameter which may be one of

```
linear      ! Linear interpolation.
cubic       ! Cubic spline interpolation (default).
```

For variable values outside of the range specified by x_knot, the slave parameter value will be extrapolated using the interpolation function calculated from the three points nearest the end. That is, extrapolation will be linear if interpolate is set to linear and will be cubic if interpolate is set to cubic. For the cubic spline, extrapolation is only permitted over a distance outside the x_knot range equal to the difference between the value of the x_knot end point and its neighbor point. Thus in the above example, extrapolation will be allowed in the intervals $[-0.7, 0.0]$ and $[1.0, 1.7]$.

After a controller has been defined, the individual knot points may be referenced using the syntax

```
name[x_knot(i)]           ! I^th x_knot point
name[slave(j)%y_knot(i)]  ! I^th y_knot point of j^th slave.
```

For example:

```
ov: overlay = {q1[k1]:{-0.12, 0.23, 0.53}, q2[tilt]:{0.37, 0.12, 0.92}},
                      var = time, x_knot = {0.0, 0.7, 1.0}
ov[x_knot(3)] = 2.0           ! Changes from 1.0 to 2.0
ov[slave(2)%y_knot(2)] = 0.2  ! Changes q2[tilt] y_knot(2) from 0.12 to 0.2
```

If the gang attribute is True (which is the default) a controller will control all elements of a given name. Thus, in the above example, if there are multiple elements named q1 then gr1 will control the k1 attribute of all of them. If gang is set to False, a separate controller is created for each element in the lattice of a given name. For example:

```
gr1: group = {q1[k1]:-tan(a)*b, q2[tilt]:b^2}, var = {a, b}, gang = False
```

In this example, suppose there are five q1 and five q2 elements in the lattice. In this case, there will be five gr1 group elements created. The first gr1 will control the first q1 and q2 to appear in the lattice, etc. With gang set to False, it is an error if the number of instances in the lattice for a given slave name is different from any other slave name. In this example, it would be an error if the number of q1 elements in the lattice is different from the number of q2 elements in the lattice.

To vary coefficients in expressions, use variables in place of the coefficients and then variables can be varied. For example:

```
gr1: group = {s1[k2]: 0.03+1.07*vv}, var = {vv}                     ! Instead of this...
gr2: group = {s1[k2]: a+b*vv}, var = {vv, a, b}, a = 0.03, b = 1.07 !   Use this.
```

Wild card characters can be used in element names. Example:

```
rf_control: overlay = {rfcavity::*[voltage]:v}, var = {v}
```

The syntax for specifying an attribute attrib of element ele to be controlled is ele[attrib]. The attribute part [attrib] may be omitted and in this case the name of the attribute will be taken to be the name of the first variable. Example:

```
ov1: overlay = {sex1:-tan(k2)^b}, var = {k2, b}
ov2: overlay = {sex1[k2]: -tan(k2)^b}, var = {k2, b} ! Equivalent to ov1.
```

In this example, the controlled attribute of element `sex1` is `k2`. Except in cases where this default attribute syntax is used, the names of the variables are arbitrary and do not have to correspond to the name of any actual attribute. If the slave attribute name is set to "`*`", multiple slaves will be generated, one for each variable. Example:

```
zz1: overlay = {sb0[*]}, var = {g, e1}
zz2: overlay = {sb0[g]:g, sb0[e1]:e1}, var = {g, e1} ! Equivalent to zz1.
```

The arithmetic expressions used to evaluate controlled attribute value changes may be a constant. In this case, the actual expression used is this constant times the first variable. If the expression is omitted entirely, along with the separating ":", the constant will be taken to be unity. Example:

```
gr1: group = {b1, b3:-pi}, var = {angle}
```

This is equivalent to

```
gr1: group = {b1[angle]:angle, b3[angle]:-pi*angle}, var = {angle}
```

The exception is when `ran()` or `ran_gauss()` functions are present in the expression. In this case, the expression will `not` be modified. Example:

```
quake: ramper = {*[y_offset]: 1e-3*ran_gauss()}, var = { }
quake: ramper = {*[y_offset]: 1e-3*ran_gauss()}, var = {time}  ! Equivalent to above
```

Arithmetic expressions may themselves contain element attributes. Example:

```
sk_q20W: overlay = {sex_20W[a1]:-sex_20W[L]}, k1
```

Here the `sk_q20w` overlay controls the `a1` multipole attribute of element `sex_20w` and the length of `sex_20w` is used as a scale factor between the overlay's variable `k1` and the controlled attribute `a1`. The potential problem here is that, to keep the internal bookkeeping simple, the value of `sex_20w[L]` is evaluated once during parsing of the lattice file and never reevaluated (§3.13). If it is desired to use a variable element attribute in an expression, this may be effectively done by defining a control variable to take its place. Thus the above overlay may be recast as:

```
sk_q20W: overlay = {sex_20W[L]:ll, sex_20W[a1]:ll*k1}, var = {k1, ll}
```

Initial values can be assigned to the variables from within the definition of the controller element. Example:

```
ov1: overlay = {...}, var = {a, b}, a = 7, b = 2
```

Here the initial values 7 and 2 are assigned to `a` and `b` respectively. Alternatively, variables can be set after a controller element has been defined. Example:

```
ov1: overlay = {...}, var = {a, b}
ov1[a] = 7
gr1[b] = 2
```

Note: There is an old deprecated syntax. For `group` elements the syntax was:

```
name: GROUP = {ele1[attrib1]:coef1, ele2[attrib2]:coef2, ...},
                      attrib = init_value  ! DO NOT USE THIS SYNTAX!
```

and for `overlay` elements the old syntax was identical except that `GROUP` was replaced by `OVERLAY`:

```
name: OVERLAY = {ele1[attrib1]:coef1, ele2[attrib2]:coef2, ...},
                      attrib = init_value  ! DO NOT USE THIS SYNTAX!
```

With this old syntax, there is only one variable. Additionally, there are no arithmetic expressions. Rather, attribute changes are linear in the `command` variable with the constant of proportionality given by a specified coefficient. For example with the old syntax

```
ov1: overlay = {sq1:3.7, sq2[tilt]}, k0 = 2  ! DO NOT USE THIS SYNTAX!
```

is equivalent, in the present syntax, to:

```
ov1: overlay = {sq[k0]:3.7, sq2[tilt]}, var = {k0}, k0 = 2
```

Note: In this old syntax the colon ":" separating the controlled attribute from the linear coefficient may be replaced by a slash "/". For `group` elements, there was an added wrinkle that, with the old syntax, the variable's name is fixed to be `command`. For example, with the old syntax

```
gr1: group = {sq1:3.7, sq2[tilt]}, k0 = 2  ! DO NOT USE THIS SYNTAX!
```

is equivalent, in the present syntax, to:

```
gr1: group = {sq[k0]:3.7, sq2[tilt]}, var = {command}, command = 2
```

5.5 Reference Energy, Time, and Wavelength

The attributes that define the reference energy and momentum at an element are:

```
e_tot  = <Real>  ! Total energy in eV.
p0c    = <Real>  ! Momentum in eV.
```

The energy and momentum are defined at the exit end of the element. For ultra–relativistic particles, and for photons, these two values are the same (§17.2). Except for multipass elements (§9), `e_tot` and `p0c` are dependent attributes and, except for multipass elements, any setting of `e_tot` and `p0c` in the lattice input file is an error. The value of `e_tot` and `p0c` for an element is calculated by *Bmad* to be the same as the previous element except for `e_gun`, `lcavity`, `converter`, and `patch` elements. To set the `e_tot` or `p0c` at the start of the lattice use the `beginning` or `parameter` statements. See §10.1. Since the reference energy changes from the start to the end of an `lcavity`, `converter` or `em_field` element, these elements have the dependent attributes

```
e_tot_start   and
p0c_start
```

which are just the reference energy and momentum at the start of the element.

The `beginning_ele` element (§4.4) also has associated `e_tot_start` and `p0c_start` attributes as well as `e_tot` and `p0c`. Generally, for a `beginning_ele`, `p0c_start` and `p0c` are the same and `e_tot_start` and `e_tot` are the same and the values for these attributes are set in the lattice file with the appropriate `parameter` (§10.1) or `beginning` (§10.4) statement. The exception occurs when there is an `e_gun` element in the lattice (§4.15). In this case, the `p0c_start` and `e_tot_start` attributes of the `beginning_ele` are set to the values as set in the lattice file and `e_tot` is set to

```
e_tot = e_tot_start + voltage
```

and `p0c` is calculated from `e_tot` and the mass of the particle being tracked. For example, if the lattice file contained:

```
beginning[p0c] = 0
gun: e_gun, voltage = 0.5e6
injector: line = (gun, ...)
```

Then the following energy values will be set for the beginning `beginning_ele` element:

```
p0c_start   = 0
e_tot_start = mc2
e_tot       = mc2 + 0.5e6
p0c         = Sqrt(e_tot - mc2^2)
```

where `mc2` is the particle rest mass. The reason for using this convoluted convention is to allow the setting, in the lattice file, of a zero reference momentum at the start of the lattice, while avoiding the calculational problems that would occur if the `e_gun` element truly had a starting reference momentum of zero. Specifically, the problem with zero reference momentum is that the phase space momentum would be infinity as can be seen from Eqs. (17.2).

For lattice branches other than the root branch, the reference energy or momentum is set using the line name of the branch (§10.4). For example:

```
ff: fork, to_line = line_2nd, ...
line_2nd: line = (...)
line_2nd[particle] = He+
line_2nd[p0c] = 1e9
```

For `multipass` elements, the reference energy may be set by specifying one of `e_tot`, `p0c`, as described in §9.2.

The value of the reference time (cf Eq. (17.3)) at the downstream end of an element is given by the element parameter `ref_time`. At the upstream end of the element, the reference time is given by

`ref_time_start`. The change in reference time is given by delta_ref_time. The value of `ref_time` for a given element in a given branch is computed starting from the `ref_time` set for the `beginning` element of the branch and summing all the `delta_ref_time` values of the elements up to the element in question. Notice that when multi-turn tracking is involved, the `ref_time` of an element is appropriate for the first turn but to calculate the reference time for subsequent turns involves knowing how many turns a particle has traveled along with the reference time for each turn (which may not be the same turn-for-turn if there is ramping and energy changes).

The value of `delta_ref_time` for an element is generally computed using the time-of-flight of the `reference particle` assuming that the reference particle is following the reference orbit. Exceptions are `wiggler` elements which uses the time-of-flight of the actual trajectory starting from a particle on the zero orbit (phase space parameters all zero). Note that if `delta_ref_time_user_set` for a `wiggler` is set `True`, `delta_ref_time` may be set in the lattice file (as opposed to being computed by *Bmad*).

For photons, the reference wavelength, `ref_wavelength` is also a dependent attribute calculated from the reference energy.

5.6 Orientation: Offset, Pitch, Tilt, and Roll Attributes

By default, an element, like a quadrupole, is aligned in space coincident with the reference orbit running through it (§16.1.3). A quadrupole can be displaced in space using the quadrupole's "`orientational`" attributes. For a quadrupole, the orientational attributes only affect the physical element and not the reference orbit. However, the orientational attributes of some other elements, like the `fiducial` element, do affect the reference orbit. To sort all this out, lattice elements can be divided into seven classes:

1. `Straight line elements` (§5.6.2)
   Straight line elements are elements where the reference orbit is a straight line. Examples include `quadrupoles`, and `sextupoles` as well as zero length elements like `markers`.
2. `Dipole bends` (§5.6.3)
   Dipole bends are:
   `sbend & rbend`
3. `Photon reflecting elements` (§5.6.4)
   The reflecting elements are
   `crystal`
   `mirror`
   `multilayer_mirror`
   These elements have a kink in the reference orbit at the nominal element surface.
4. `Reference orbit manipulator elements` (§5.6.5)
   Elements that are used to manipulate the reference orbit are
   `fork & photon_fork`
   `floor_shift`
   `patch`
5. `Fiducial Element` (§5.6.6)
6. `Girder Elements` (§5.6.7)
7. `Control Elements`
   Control elements are elements that control attributes of other elements. Except for `girder` control elements, these elements do not have orientational attributes. Control elements that fall into this list are:
   `group`
   `overlay`
   `ramper`

5.6.1 Global Random Misalignment of Elements

It is often convenient to randomly misalign sets of elements. This can be done using the `ran` and `ran_gauss` functions (§3.14). For example:

```
quadrupole::bnd10:bnd20[y_offset] = 1.4e-6*ran_gauss(5)
```

The above line sets the `y_offset` of all the quadrupoles in the range from element `BND10` to element `BND20`.

When `ran` or `ran_gauss` is used in a lattice file, each time the file is read in, a new set of random numbers are generated unless `parameter[ran_seed]` is set to a non-zero value in the lattice file before `ran` or `ran_gauss` is used. To save a particular set of generated random values, write out the lattice file (the `write bmad` command can be used if running *Tao*) after it has been read in.

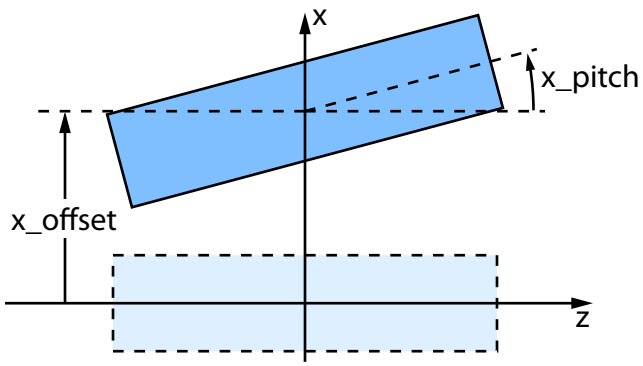

Figure 5.1: Geometry of Pitch and Offset attributes

5.6.2 Straight Line Element Orientation

The straight line elements have the following orientational attributes:

```
x_offset = <Real>
y_offset = <Real>
z_offset = <Real>
x_pitch  = <Real>
y_pitch  = <Real>
tilt     = <Real>
```

For straight line elements the orientational attributes only shift the physical element and do not affect the reference orbit.

`x_offset` translates an element in the local x–direction as shown in Fig. 5.1. Similarly, `y_offset` and `z_offset` translate an element along the local y and z–directions respectively.

The `x_pitch` attribute rotates an element about the element's center such that with a positive `x_pitch` the exit face of the element is displaced in the $+x$–direction as shown in figure 5.1. [One way to visualize the effect of an `x_pitch` is to think of the element as an airplane pointing in the $+z$ direction. A positive `x_pitch` would then move the front of the plane in the $+x$–direction.] An`x_pitch` represents a rotation around the positive y-axis.

Similarly, the `y_pitch` attribute rotates an element about the element's center using the negative x–axis as the rotation axis so that, with a positive `y_pitch` the exit face of the element is displaced in the $+y$–direction.

Note: the `x_pitch` and `y_pitch` rotations are about the center of the element which is in contrast to the `dtheta` and `dphi` misalignments of *MAD* which rotate around the entrance point. The sense of the rotation between *Bmad* and MAD is:

```
x_pitch (Bmad) =  dtheta (MAD)
y_pitch (Bmad) = -dphi (MAD)
```

The tilt attribute rotates the element in the (x, y) plane as shown in figure 5.2. The rotation axis is the positive z-axis. For example

```
q1: quad, l = 0.6, x_offset = 0.03, y_pitch = 0.001, tilt
```

Like MAD, *Bmad* allows the use of the `tilt` attribute without a value to designate a skew element. The default tilt is $\pi/(2(n+1))$ where n is the order of the element:

```
sol_quad       n = 1
quadrupole     n = 1
```

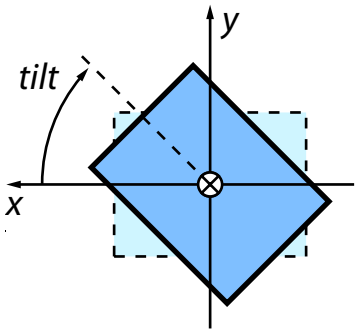

Figure 5.2: Geometry of a Tilt

```
sextupole      n = 2
octupole       n = 3
```

Note that `hkick` and `vkick` attributes are not affected by `tilt` except for `kicker` and `elseparator` elements.

5.6.3 Bend Element Orientation

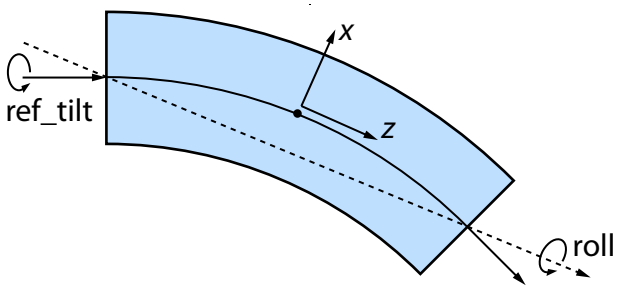

Figure 5.3: Geometry of a Bend. Like straight line elements, offsets and pitches are calculated with respect to the coordinates at the center of the bend. The exception is the `roll` attribute which is a rotation around the axis passing through the entrance and exit points. Shown here is the geometry for a bend with `ref_tilt` = 0. That is, the bend is in the $x - z$ plane.

The orientation attributes for `sbend`, `rbend` and `rf_bend` elements is

```
x_offset = <Real>
y_offset = <Real>
z_offset = <Real>
x_pitch  = <Real>
y_pitch  = <Real>
ref_tilt = <Real>    ! Shifts and reference orbit rotation axis.
roll     = <Real>
```

The geometry for orienting a bend is shown in Fig. 5.3. Like straight line elements, the offset and pitch attributes are evaluated with respect to the center of the element.

Unlike the straight line elements, bends do not have a `tilt` attribute. Rather they have a `ref_tilt` and a `roll` attribute. The `roll` attribute rotates the bend along an axis that runs through the entrance point and exit point as shown in figure 5.3. A `roll` attribute, like the offset and pitch attributes does

not affect the reference orbit. The major effect of a `roll` is to give a vertical kick to the beam. For a bend with positive bend angle, a positive `roll` will move the outside portion ($+x$ side) of the bend upward and the inside portion (-x side) downward. Much like car racetracks which are typically slanted towards the inside of a turn.

The `ref_tilt` attribute of a bend rotates the bend about the z axis at the upstream end of the bend as shown in Fig. 5.3. Unlike `rolls` and `tilts`, `ref_tilt` also shifts the rotation axis of the reference orbit along with the physical element. A `bend` with a `ref_tilt` of $\pi/2$ will bend a beam vertically downward (§16.2). Note that the `ref_tilt` attribute of *Bmad* is the same as the MAD `tilt` attribute.

Important! Do not use `ref_tilt` when doing misalignment studies for a machine. Trying to misalign a dipole by setting `ref_tilt` will affect the positions of all downstream elements! Rather, use the `roll` parameter.

5.6.4 Photon Reflecting Element Orientation

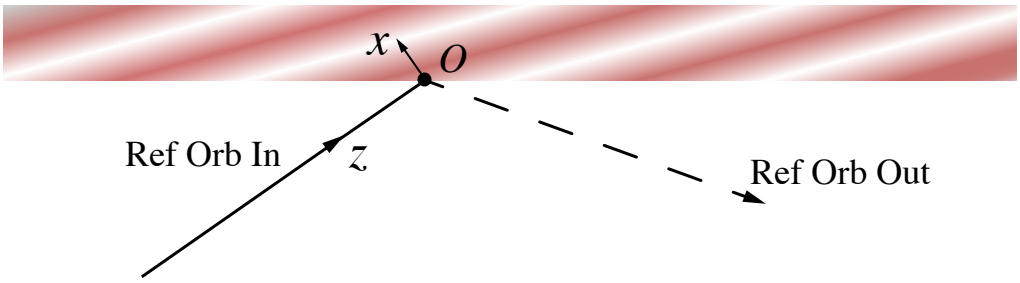

Figure 5.4: Geometry of a photon reflecting element orientation. The reference coordinates used for defining the orientational attribute is the entrance reference coordinates.

Photon reflecting elements have the following orientational attributes:

```
x_offset = <Real>
y_offset = <Real>
z_offset = <Real>
x_pitch  = <Real>
y_pitch  = <Real>
ref_tilt = <Real>    ! Shifts both element and reference orbit.
tilt     = <Real>
```

Roughly, these elements can be viewed as zero length bends except, since there is no center position, the orientational attributes are defined with respect to the entrance coordinates as shown in Fig. 5.4. Like bend elements, the `ref_tilt` attribute rotates both the physical element and the reference coordinates. The `tilt` attribute rotates just the physical element. Thus the total rotation of the physical element about the entrance z axis is the sum `tilt` + `ref_tilt`.

Frequently, it is desired to orient reflecting elements with respect to the element's surface. This can be done using a `girder` element (§4.24) which supports the reflecting element and with the `girder`'s `origin_ele_ref_pt` attribute set to `center`.

5.6.5 Reference Orbit Manipulator Element Orientation

The `fork`, `photon_fork`, `floor_shift`, and `patch` elements use the following attributes to orient their exit edge with respect to their entrance edge:

```
x_offset = <Real>
y_offset = <Real>
z_offset = <Real>
x_pitch  = <Real>
y_pitch  = <Real>
tilt     = <Real>
```

Here "exit" edge for `fork` and `photon_fork` elements is defined to be the start of the line being branched to. [Within the line containing the fork, the `fork` element is considered to have zero length so the exit face in the line containing the fork is coincident with the entrance face.] The placement of the exit edge for these elements defines the reference orbit. Thus, unlike the corresponding attributes for other elements, the orientational attributes here directly control the reference orbit.

5.6.6 Fiducial Element Orientation

The `fiducial` element (§4.24) uses the following attributes to define its position:

```
origin_ele        = <Name>     ! Reference element.
origin_ele_ref_pt = <location> ! Reference pt on reference ele.
dx_origin         = <Real>     ! x-position offset
dy_origin         = <Real>     ! y-position offset
dz_origin         = <Real>     ! z-position offset
dtheta_origin     = <Real>     ! orientation angle offset.
dphi_origin       = <Real>     ! orientation angle offset.
dpsi_origin       = <Real>     ! orientation angle offset.
```

See Section §4.19 for more details.

5.6.7 Girder Orientation

A `girder` (§4.24) element uses the same attributes as a `fiducial` element (§4.19) to orient the reference girder position. In addition, the following attributes are used to move the girder physically from the reference position:

```
x_offset = <Real>
y_offset = <Real>
z_offset = <Real>
x_pitch  = <Real>
y_pitch  = <Real>
tilt     = <Real>
```

Shifting the girder from its reference position shifts all the elements that are supported by the girder. See Section §4.24 for more details.

If an element is supported by a `girder` element (§4.24), the orientational attributes of the element are with respect to the orientation of the `girder`. The computed offsets, pitches and tilt with respect to the local reference coordinates are stored in the dependent attributes

```
x_offset_tot
y_offset_tot
z_offset_tot
x_pitch_tot
y_pitch_tot
tilt_tot
roll_tot
```

A `*_tot` attribute will only be present if the corresponding non `*_tot` attribute is present. For example, only `sbend` and `rbend` elements have a `roll_tot` attribute since only these elements have a `roll` attribute.

If an element is not supported by a `girder`, the values of the `*_tot` attributes will be the same value as the values of the corresponding non `*_tot` attributes.

5.7 Hkick, Vkick, and Kick Attributes

The kick attributes that an element may have are:

```
kick,  bl_kick  = <Real>  ! Used only with a Hkicker or Vkicker
hkick, bl_hkick = <Real>
vkick, bl_vkick = <Real>
```

`kick`, `hkick`, and `vkick` attributes are the integrated kick of an element in radians. `kick` is only used for `hkicker` and `vkicker` elements. All other elements that can kick use `hkick` and `vkick`. The `tilt` attribute will only rotate a kick for `hkicker`, `vkicker`, `elseparator` and `kicker` elements. This rule was implemented so that, for example, the `hkick` attribute for a skew quadrupole would represent a horizontal steering. The `bl_kick`, `bl_hkick`, and `bl_vkick` attributes are the integrated field kick in `meters-Tesla`. Normally these are dependent attributes except if they appear in the lattice file (§5.1).

For an `elseparator` element, the `hkick` and `vkick` are appropriate for a positively charged particle. The kick for a negatively charged particle is opposite this.

5.8 Aperture and Limit Attributes

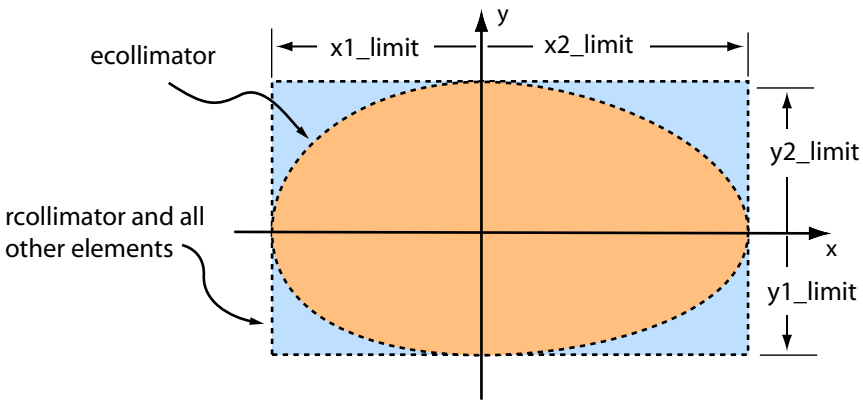

Figure 5.5: Apertures for ecollimator and rcollimator elements. [note: positive z points up, out of the page.] As drawn, all limits `x1_limit`, `x2_limit`, `y1_limit`, `y2_limit` are positive.

The aperture attributes are:

```
x1_limit      = <Real>      ! Horizontal, negative side, aperture limit
x2_limit      = <Real>      ! Horizontal, positive side, aperture limit
y1_limit      = <Real>      ! Vertical, negative side, aperture limit
y2_limit      = <Real>      ! Vertical, positive side, aperture limit
x_limit       = <Real>      ! Alternative to specifying x1_limit and x2_limit
y_limit       = <Real>      ! Alternative to specifying y1_limit and y2_limit
```

```
aperture      = <Real>      ! Alternative to specifying x_limit and y_limit
aperture_at   = <Switch>    ! What end aperture is at. (§5.8.2)
aperture_type = <Switch>    ! What type of aperture it is
offset_moves_aperture = <Logical> ! Element offsets affect aperture position (§5.8.1)
```

`x1_limit`, `x2_limit`, `y1_limit`, and `y2_limit` specify the half–width of the aperture of an element as shown in figure 5.5. Note: Normally all of these will be zero or positive. A zero `x1_limit`, `x2_limit`, `y1_limit`, or `y2_limit` is interpreted as no aperture in the appropriate plane.

For convenience, `x_limit` can be used to set `x1_limit` and `x2_limit` to a common value. Example:

```
s: sextupole, x1_limit = 0.09, x2_limit = 0.09
s: sextupole, x_limit = 0.09   ! Same as above
```

Similarly, `y_limit` can be used to set `y1_limit` and `y2_limit`. The `aperture` attribute can be use to set all four `x1_limit`, `x2_limit`, `y1_limit` and `y2_limit` to a common value. Internally, the *Bmad* code does *not* store `x_limit`, `y_limit`, or `aperture`. This means that using `x_limit`, `y_limit` or aperture in arithmetic expressions is an error:

```
q1: quad, aperture = 0.09
q2: quad, aperture = q1[aperture]   ! THIS IS AN ERROR!
q2: quad, aperture = q1[x1_limit]   ! Correct
```

By default, apertures are assumed to be rectangular except that an `ecollimator` has a elliptical aperture. This can be changed by setting the `aperture_type` attribute. The possible values of this attribute are:

```
auto         ! Default for detector, mask and diffraction_plate elements
custom
elliptical   ! Default for ecollimator elements.
rectangular  ! Default for most elements.
wall3d       ! Vacuum chamber wall (§5.12).
```

The `custom` setting is used in the case where programs have been compiled with custom, non-Bmad, code to handle the aperture calculation. The `auto` setting is used for automatic calculation of a rectangular aperture. For `diffraction_plate` and `mask` elements, the `auto` setting causes the four aperture limits to be set to just cover the clear area of element (§5.12.6). For all other elements, the `auto` setting is only to be used when there is an associated surface grid (§5.11.1) for the element and, in this case, *Bmad* to set the four limits to just cover the surface grid.

The `wall3d` setting uses the vacuum chamber wall as specified by a `wall` attribute (§5.12). Using the `wall` construct allows for complex apertures to be constructed. Note that The wall thickness and material type are not used when calculating if a particle has hit the wall. That is, the wall is considered to be infinitely thin. Also note that a wall must cover the entire length of the element longitudinally. This is done in order to be able to spot errors in specifying the wall geometry.

For elliptical apertures, all four `x1_limit`, `x2_limit`, `y1_limit`, and `y2_limit` must be non-zero.

For rectangular apertures, the limits `x1_limit`, `x2_limit`, `y1_limit`, or `y2_limit` may be negative. For example:

```
s: sextupole, x1_limit = -0.02, x2_limit = 0.09
```

In this case, particles will hit the aperture if their x-coordinate is outside the interval [0.02, 0.09]. That is, particles at the origin will be lost.

To avoid numerical overflow and other errors in tracking, a particle will be considered to have hit an aperture in an element, even if there are no apertures set for that element, if its orbit exceeds 1000 meters. Additionally, there are other situations where a particle will be considered lost. For example, if a particle's trajectory does not intersect the output face in a bend.

Examples:

```
q1, quadrupole, y1_limit = 0.03
q1[y2_limit] = 0.03
q1[y_limit] = 0.03  ! equivalent to the proceeding 2 lines.
q1[aperture_at] = both_ends
```

5.8.1 Apertures and Element Offsets

Normally, whether a particle hits an aperture or not is evaluated independent of any element offsets (§5.6). This is equivalent to the situation where a beam pipe containing an aperture is independent of the placement of the physical element the beam pipe passes through. That is, the beam pipe does not "touch" the physical element. This can be changed by setting the `offset_moves_aperture` attribute to `True`. In this case any offsets or pitches will be considered to have shifted the aperture boundary. The exceptions here is that the default for the following elements is for `offset_moves_aperture` to be `True`:

```
rcollimator,
ecollimator,
multilayer_mirror,
mirror, and
crystal
```

Even with `offset_moves_aperture` set to `True`, `tilt`s will not affect the aperture calculation. This is done, for example, so that the tilt of a skew quadrupole does not affect the aperture. The exception here is that tilting an `rcollimator` or `ecollimator` element will tilt the aperture. Additionally, when the aperture is at the `surface` (see below), any `tilt` will be used in the calculation.

Example:

```
q1: quad, l = 0.6, x1_limit = 0.045, offset_moves_aperture = T
```

5.8.2 Aperture Placement

By default, for most elements, the aperture is evaluated at the exit face of the element. This can be changed by setting the `aperture_at` attribute. Possible settings for `aperture_at` are:

```
both_ends
continuous
entrance_end
exit_end       ! Default for most elements
no_aperture
surface
wall_transition
```

The `exit_end` setting is the default for most elements except for the following elements who have a default of `surface`:

```
crystal
detector
diffraction_plate
mask
mirror
multilayer_mirror
sample
```

In fact, for the following elements:

```
mirror,
multilayer_mirror
crystal
```

The `surface` setting for `aperture_at` must be used. Additionally, due to the complicated geometry of these elements, to keep things conceptionally simple, the rule is imposed that, for an aperture at the surface, the `offset_moves_aperture` setting must be left in its default state of True. Additionally, For `entrance_end` or `exit_end` apertures, `offset_moves_aperture` must be set to False.

Note: The entrance and exit ends of an element are independent of which direction particles are tracked through an element. Thus if a particle is tracked backwards it enters an element at the "exit end" and exits at the "entrance end". The `continuous` setting indicates that the aperture is continuous along the length of the element. This only matters when particle tracking involves stepping through an element a little bit at a time. For example, as in Runge-Kutta tracking (§6.1). For tracking where a formula is used to transform the particle coordinates at the entrance of an element to the coordinates at the exit end, the aperture is only checked at the end points so, in this situation, a `continuous` aperture is equivalent to the `both_ends` setting.

The `wall_transition` setting is like the `continuous` setting in that the aperture boundary is considered to be continuous along the element's length. However, unlike the `continuous` setting, with the `wall_transition` setting a particle outside the wall is considered alive and it is only when a particle moves through the wall that it is lost. The `wall_transition` setting is used for things like septum magnets where a particle may be safely outside or inside the wall. Note to programmers: By supplying a custom `wall_hit_handler_custom` routine, scattering of particles through a wall may be simulated.

Examples:

```
q2: quad, aperture_type = elliptical, aperture_at = continuous
q1: quad, l = 0.6, x1_limit = 0.045, offset_moves_aperture = T
```

5.8.3 Apertures and X-Ray Generation

With X-ray simulation apertures can be used by *Bmad* to limit the directions in which photons are generated. This can greatly decrease simulation times. For example, a photon passing through a `diffraction_plate` element will diffract in an arbitrary direction. If a *downstream* element has an aperture set, *Bmad* can restrict the velocity directions so that the photons will fill the downstream aperture and the amount of time wasted tracking photons that ultimately would be collimated is minimal.

5.9 X-Rays Crystal & Compound Materials

For basic crystallographic and X-ray matter interaction cross-sections, *Bmad* uses the XRAYLIB[Schoon11] library. Crystal structure parameters in XRAYLIB are mainly from R. W. G. Wyckoff[Wyckoff65] with some structure parameters coming from NIST. The list of available structures is:

AlphaAlumina	GaP	KCl	Platinum
AlphaQuartz	GaSb	KTP	RbAP
Aluminum	Ge	LaB6	Sapphire
Be	Gold	LaB6_NIST	Si
Beryl	Graphite	LiF	Si_NIST
Copper	InAs	LiNbO3	Si2
CsCl	InP	Muscovite	SiC
CsF	InSb	NaCl	Titanium
Diamond	Iron	PET	TlAP
GaAs	KAP		

These names are case sensitive

Besides the above crystal list, *Bmad* can calculate structure factors for all the elements and the following list of materials. Material properties are from NIST. These names are case sensitive. That is, the NIST materials all use upper case. As noted in the table, several of the materials may be specified using the appropriate chemical formula. For example, liquid water may be referenced using the name `H2O`.

A_150_TISSUE_EQUIVALENT_PLASTIC
ACETONE
ACETYLENE
ADENINE
ADIPOSE_TISSUE_ICRP
AIR_DRY_NEAR_SEA_LEVEL
ALANINE
ALUMINUM_OXIDE, Al2O3
AMBER
AMMONIA, NH3
ANILINE
ANTHRACENE
B_100_BONE_EQUIVALENT_PLASTIC
BAKELITE
BARIUM_FLUORIDE
BARIUM_SULFATE
BENZENE, C6H6
BERYLLIUM_OXIDE
BISMUTH_GERMANIUM_OXIDE
BLOOD_ICRP
BONE_COMPACT_ICRU
BONE_CORTICAL_ICRP
BORON_CARBIDE, B4C
BORON_OXIDE, B2O3
BRAIN_ICRP
BUTANE
N_BUTYL_ALCOHOL
C_552_AIR_EQUIVALENT_PLASTIC
CADMIUM_TELLURIDE
CADMIUM_TUNGSTATE
CALCIUM_CARBONATE
CALCIUM_FLUORIDE
CALCIUM_OXIDE
CALCIUM_SULFATE
CALCIUM_TUNGSTATE
CARBON_DIOXIDE
CARBON_TETRACHLORIDE
CELLULOSE_ACETATE_CELLOPHANE
CELLULOSE_ACETATE_BUTYRATE
CELLULOSE_NITRATE
CERIC_SULFATE_DOSIMETER_SOLUTION
CESIUM_FLUORIDE
CESIUM_IODIDE
CHLOROBENZENE
CHLOROFORM
CONCRETE_PORTLAND
CYCLOHEXANE
12_DDIHLOROBENZENE
DICHLORODIETHYL_ETHER
12_DICHLOROETHANE
DIETHYL_ETHER
NN_DIMETHYL_FORMAMIDE
DIMETHYL_SULFOXIDE
ETHANE
ETHYL_ALCOHOL
LITHIUM_TETRABORATE
LUNG_ICRP
M3_WAX
MAGNESIUM_CARBONATE
MAGNESIUM_FLUORIDE
MAGNESIUM_OXIDE
MAGNESIUM_TETRABORATE
MERCURIC_IODIDE
METHANE
METHANOL
MIX_D_WAX
MS20_TISSUE_SUBSTITUTE
MUSCLE_SKELETAL
MUSCLE_STRIATED
MUSCLE_EQUIVALENT_LIQUID_WITH_SUCROSE
MUSCLE_EQUIVALENT_LIQUID_WITHOUT_SUCROSE
NAPHTHALENE
NITROBENZENE
NITROUS_OXIDE
NYLON_DU_PONT_ELVAMIDE_8062
NYLON_TYPE_6_AND_TYPE_6_6
NYLON_TYPE_6_10
NYLON_TYPE_11_RILSAN
OCTANE_LIQUID
PARAFFIN_WAX
N_PENTANE
PHOTOGRAPHIC_EMULSION
PLASTIC_SCINTILLATOR_VINYLTOLUENE_BASED
PLUTONIUM_DIOXIDE
POLYACRYLONITRILE
POLYCARBONATE_MAKROLON_LEXAN
POLYCHLOROSTYRENE
POLYETHYLENE
POLYETHYLENE_TEREPHTHALATE_MYLAR
POLYMETHYL_METHACRALATE_LUCITE_PERSPEX
POLYOXYMETHYLENE
POLYPROPYLENE
POLYSTYRENE
POLYTETRAFLUOROETHYLENE_TEFLON
POLYTRIFLUOROCHLOROETHYLENE
POLYVINYL_ACETATE
POLYVINYL_ALCOHOL
POLYVINYL_BUTYRAL
POLYVINYL_CHLORIDE
POLYVINYLIDENE_CHLORIDE_SARAN
POLYVINYLIDENE_FLUORIDE
POLYVINYL_PYRROLIDONE
POTASSIUM_IODIDE
POTASSIUM_OXIDE
PROPANE
PROPANE_LIQUID
N_PROPYL_ALCOHOL
PYRIDINE
RUBBER_BUTYL
RUBBER_NATURAL

Continued on next page

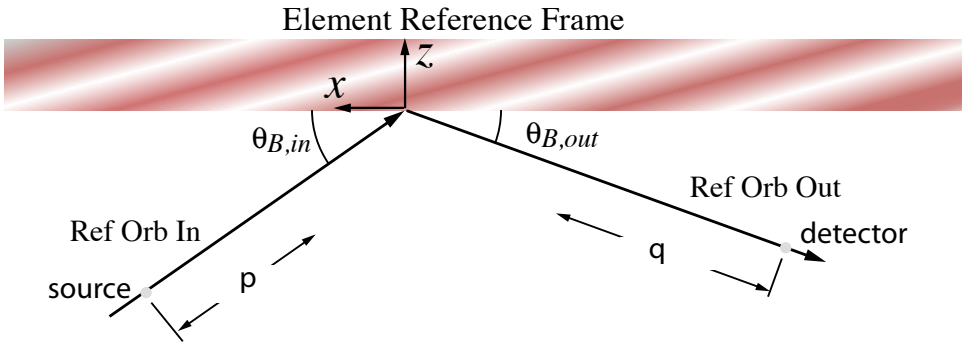

Figure 5.6: Surface curvature geometry. The element reference frame used to describe surface curvature has the z axis pointing towards the interior of the element, and the x axis in the plane defined by the entrance and exit reference orbit.

ETHYL_CELLULOSE
ETHYLENE
EYE_LENS_ICRP
FERRIC_OXIDE
FERROBORIDE
FERROUS_OXIDE
FERROUS_SULFATE_DOSIMETER_SOLUTION
FREON_12
FREON_12B2
FREON_13
FREON_13B1
FREON_13I1
GADOLINIUM_OXYSULFIDE
GALLIUM_ARSENIDE
GEL_IN_PHOTOGRAPHIC_EMULSION
GLASS_PYREX
GLASS_LEAD
GLASS_PLATE
GLUCOSE
GLUTAMINE
GLYCEROL
GUANINE
GYPSUM_PLASTER_OF_PARIS
N_HEPTANE
N_HEXANE
KAPTON_POLYIMIDE_FILM
LANTHANUM_OXYBROMIDE
LANTHANUM_OXYSULFIDE
LEAD_OXIDE
LITHIUM_AMIDE
LITHIUM_CARBONATE
LITHIUM_FLUORIDE
LITHIUM_HYDRIDE
LITHIUM_IODIDE
LITHIUM_OXIDE
RUBBER_NEOPRENE
SILICON_DIOXIDE
SILVER_BROMIDE
SILVER_CHLORIDE
SILVER_HALIDES_IN_PHOTOGRAPHIC_EMULSION
SILVER_IODIDE
SKIN_ICRP
SODIUM_CARBONATE
SODIUM_IODIDE
SODIUM_MONOXIDE
SODIUM_NITRATE
STILBENE
SUCROSE
TERPHENYL
TESTES_ICRP
TETRACHLOROETHYLENE
THALLIUM_CHLORIDE
TISSUE_SOFT_ICRP
TISSUE_SOFT_ICRU_FOUR_COMPONENT
TISSUE_EQUIVALENT_GAS_METHANE_BASED
TISSUE_EQUIVALENT_GAS_PROPANE_BASED
TITANIUM_DIOXIDE
TOLUENE
TRICHLOROETHYLENE
TRIETHYL_PHOSPHATE
TUNGSTEN_HEXAFLUORIDE
URANIUM_DICARBIDE
URANIUM_MONOCARBIDE
URANIUM_OXIDE
UREA
VALINE
VITON_FLUOROELASTOMER
WATER_LIQUID, H2O
WATER_VAPOR
XYLENE

5.10 X-Ray Reflectivity Tables

Reflectivity tables are used to define reflectivity probabilities as functions of incidence angle and photon energy for `crystal` elements.

The general syntax is:

```
m: crystal, use_reflectivity_table = <T/F>,
  reflectivity_table = {
    polarization = <polarization-type>
    angles = (<ang1> <ang2> ... <angN>),
    p_reflect = {
      <energy1> <P_11> <P_12> <P_13> ... <P_1N>,
      <energy2> <P_21> <P_22> <P_23> ... <P_2N>,
      ...
    }
  }, ...
```

The `use_reflectivity_table` element parameter can be used to turn on or off use of the reflectivity table. The `p_reflect` table inside the `reflectivity_table` construct must come last. Angles are in radians and energy is in eV. Angle and energy points do not need to be evenly spaced. Probabilities `P_ij` are between 0 and 1 and must be real numbers (as opposed to expressions). For photons that fall outside the range of angles or energies, the probability is taken to be zero. The `p_reflect` portion of the reflectivity table must come last. Notice that angles and probability values are separated by blank characters.

Possible `polarization` types are:

```
pi        ! Table is for pi mode
sigma     ! Table is for sigma mode
both      ! Table is for both polarizations. Default.
```

An element needs a single table with `polarization` marked as `both` or two tables, one for `sigma` and the other for `pi` reflections.

For `crystal` elements, angles are always with respect to the Bragg angle for the energy where the Bragg angle is calculated without any refraction corrections.

Example:

```
cryst1: crystal, use_reflectivity_table = T,
    reflectivity_table = {
      polarization = both,
      angles = (0   0.5*pi),
      p_reflect = {
          12000 1 1,
          14000 1 1
        }
    }, ...
```

This gives 100% reflection for photons whose energy is in the range 12 KeV to 14 KeV and whose incidence angle is in the range between 0 and 0.5π. Photons outside of either of these ranges have zero probability of reflection.

5.11 Surface Properties for X-Ray elements

The following X-ray elements have a surface which X-rays impinge upon:

`crystal`	§4.10
`detector`	§4.12
`diffraction_plate`	§4.13
`mask`	§4.33
`mirror, and`	§4.35
`multilayer_mirror`	§4.37
`sample`	§4.48

[There is also the `capillary` element but this element specifies its surface differently.]

The coordinate system used for characterizing the curvature of a surface is the element reference frame as shown in Fig. 5.6). This coordinate system has the z axis pointing towards the interior of the element, and the x axis in the plane defined by the entrance and exit reference orbit. In this coordinate system, the surface is an ellipsoid plus a fourth order polynomial in x and y plus a possible "figure error" contribution z_{off} defined by a surface grid:

$$-z = \frac{1}{g_z}\left[1 - \sqrt{1-(g_x\,x)^2-(g_y\,y)^2}\right] + \frac{1}{g_{sp}}\left[1 - \sqrt{1-(g_{sp}\,x)^2-(g_{sp}\,y)^2}\right] + \sum_{2\le i+j\le 6} c_{ij}\,x^i\,y^j - z_{\text{off}} \tag{5.1}$$

z_{off} is discussed in section §5.11.1 and is only present when the surface grid `type` is set to `Displacement`. In Eq. (5.1) the c_{ij} coefficients parameterize the fourth order polynomial, g_{sp} parameterize the spherical curvature, and g_x, g_y, and g_z parameterize ellipsoid curvature. [In principle, the spherical curvature is not needed since the elliptical curvature is more general. In practice, it is sometimes convenient to be able to specify spherical curvature.] If g_z is zero, the elliptical curvature is ignored. If g_z or g_{sp} is positive, the curvature is concave towards the incoming photon. If negative, the curvature is convex.

The spherical, ellipsoid and polynomial parameters are set for an element by setting the element's curvature parameter. The syntax is

```
curvature = {
    spherical    = <Real>,   ! g_sp
    elliptical_x = <Real>,   ! g_x
    elliptical_y = <Real>,   ! g_y
    elliptical_z = <Real>,   ! g_z
    xMyN         = <Real>    ! c_ij
```

The polynomial coefficients c_{ij} are set via the `xMyN` components where `M` and `N` are integers in the range 0 through 6 with the restriction

```
2 ≤ M + N ≤ 6
```

Example:

```
c2: crystal, curvature = {spherical = 1/4.7, x2y0 = 0.37, ...}
```

in this example, `x2y0` corresponds to the c_{20} term in Eq. (5.1). To get the effect of a nonzero $x^0\,y^0$, $x^1\,y^0$, or $x^0\,y^1$ terms (since corresponding `curvature xNyM` are not permitted), element offsets and pitches can be used (§5.6).

Some useful formulas: Series expansion for a sphere of radius R:

$$-z = \frac{x^2}{2\,R} + \frac{x^4}{8\,R^3} + \frac{x^6}{16\,R^5} + \frac{y^2}{2\,R} + \frac{y^4}{8\,R^3} + \frac{y^6}{16\,R^5} + \frac{x^2\,y^2}{4\,R^3} + \frac{3\,x^4\,y^2}{16\,R^5} + \frac{3\,x^2\,y^4}{16\,R^5} \tag{5.2}$$

For a torus with equation

$$\left(\sqrt{x^2+(z+r+R)^2} - R\right)^2 + y^2 = r^2 \tag{5.3}$$

The series expansion is:

$$-z = \frac{x^2}{2\,(r+R)} + \frac{x^4}{8\,(r+R)^3} + \frac{x^6}{16\,(r+R)^5} + \frac{y^2}{2\,r} + \frac{y^4}{8\,r^3} + \frac{y^6}{16\,r^5} + \frac{x^2\,y^2}{4\,r\,(r+R)^2} + \frac{3\,x^4\,y^2}{16\,r\,(r+R)^4} + \frac{(3\,r+R)\,x^2\,y^4}{16\,r^3\,(r+R)^3} \tag{5.4}$$

For a curved crystal, if p is the distance from the source to the crystal, and q is the distance from the crystal to the detector, the needed radius of curvature R_s in the sagittal (transverse) plane to give focusing is [Rio98]:

$$\frac{1}{p} + \frac{1}{q} = \frac{\sin\theta_{g,in} + \sin\theta_{g,out}}{R_s} \tag{5.5}$$

where $\theta_{g,in}$ and $\theta_{g,out}$ are the entrance and exit graze angles. In the tangential (meridional) plane, the radius R_t needed for focusing is

$$\frac{\sin^2\theta_{g,in}}{p} + \frac{\sin^2\theta_{g,out}}{q} = \frac{\sin\theta_{g,in} + \sin\theta_{g,out}}{R_t} \tag{5.6}$$

The above formulas assume that the crystal is constructed so that the orientation of the Bragg planes follows the orientation of the surface. Mirrors have similar formulas with $\theta_{g,in} = \theta_{g,out} = \theta$.

Example:

```
t_bragg = 1.3950647
rt = 1  ! Crystal transverse radius
rs = rt*(sin(t_bragg))^2
c: crystal, crystal_type =  "Si(553)", b_param = -1, curvature = {x0y2 = 1 / (2 * rs),
                    x0y4 = 1 / (8 * rs^3), x2y0 = 1 / (2 * rt), x4y0 = 1 / (8 * rt^3)}
c[curvature%x0y2] = 1.01 * c[curvature%x0y2]  ! Put in an error
```

5.11.1 Displacement, H_Misalign, and Segmented Surface Grids

A surface can be broken up into a grid of rectangles. This is useful, for example, in simulating crystal surface roughness. The case of the `pixel` grid for a `detector` element is discussed in Sec. §4.12. Here the other three types of grids are explained. These are:

```
Displacement       ! Mesh defines an offset from the nominal surface.
H_Misalign         ! Misalignment of crystal H vector
Segmented          ! Surface is a matrix of flat rectangles
```

All grids have the following common parameters:

```
active = <T/F>                      ! Turn on/off effect of grid. Default = True
ix_bounds = (<ix_min>, <ix_max>),   ! Min/max grid index bounds in x-direction
iy_bounds = (<iy_min>, <iy_max>),   ! Min/max grid index bounds in y-direction
r0 = (<x0>, <y0>),                  ! (x,y) coordinates at grid origin
dr = (<dx>, <dy>),                  ! Spacing between grid points.
```

Example:

```
ccd: crystal, h_misalign = {
          r0 = (0.0, 0.01), dr = (0.005, 0.005),
          ix_bounds = (1, 57), iy_bounds = (-30, 10),
          pt(1,-30) = (0.001, -0.002, 0, 0),
          pt(1,-29) = ...,
        }
```

The grid is a two dimensional rectangular mesh with bounds given by the `ix_bounds` and `iy_bounds` components. In the above example the grid is 57 pixels in x and 41 pixels in y.

The physical placement of the grid on the element is determined by the `r0` and `dr` components. `r0` is optional and gives the (x, y) coordinates of the center of the pixel with index $(0, 0)$. The `dr` component, which must be present, gives the pixel width and height. Thus the center of the (i, j) pixel is:

```
(x,y) = (r0(1), r0(2)) + (i*dr(1), j*dr(2))
```

The different grid types are:

Displacement grid

The general syntax for a `displacement` grid is:

```
displacement = {
    active = <T/F>                      ! Turn on/off effect of grid. Default = True
    ix_bounds = (<ix_min>, <ix_max>),   ! Min/max index bounds in x-direction
    iy_bounds = (<iy_min>, <iy_max>),   ! Min/max index bounds in y-direction
    r0 = (<x0>, <y0>),                  ! (x,y) coordinates at grid origin
    dr = (<dx>, <dy>),                  ! Spacing between grid points.
    pt(<i>,<j>) = (<z>),                                    ! or
    pt(<i>,<j>) = (<z>, <dz/dx>, <dz/dy>),                  ! or
    pt(<i>,<j>) = (<z>, <dz/dx>, <dz/dy>, <d2z/dxdy>),
        }
```

With a `displacement` grid, an offset $z_{\rm off}$ is added to the surface position defined in Eq. (5.1). The z offset at a given (x, y) position is determined by a spline interpolation of the z values of the grid of points defined by `pt(<i>,<j>)`. For each point `pt(<i>,<j>)`, the z value along with the dz/dx, dz/dy, $d2z/dxdy$ slopes can be specified or only the z value needs to be specified. With the later option the slopes will be computed by taking finite differences of nearest neighbors.

H_Misalign grid

The general syntax for a `h_misalign` grid:

```
h_misalign = {
    active = <T/F>                      ! Turn on/off effect of grid. Default = True
    ix_bounds = (<ix_min>, <ix_max>),   ! Min/max index bounds in x-direction
    iy_bounds = (<iy_min>, <iy_max>),   ! Min/max index bounds in y-direction
    r0 = (<x0>, <y0>),                  ! (x,y) coordinates at grid origin
    dr = (<dx>, <dy>)                   ! Spacing between grid points.
    pt(<i>,<j>) = (<rot_y, <rot_x>, <rot_y_rms>, <rot_x_rms>)  ! For Bragg diffraction
    pt(<i>,<j>) = (<rot_y, <rot_z>, <rot_y_rms>, <rot_z_rms>)  ! For Laue diffraction
        }
```

An `h_misalign` grid is used with crystals only. With `H_Misalign`, the grid defines misalignment of the **H** vector which is the normal to the diffracting planes of the crystal (§27.4). When using `H_Misalign`, each `pt(i,j)` component gives the angular misalignment of **H** for the region around the point. for Bragg diffraction where **H** is oriented approximately along the $-z$-axis, the misalignment of **H** is characterized by rotations about the y-axis and x-axis

```
rot_y_tot = <rot_y> + r2 * <rot_y_rms>
rot_x_tot = <rot_x> + r1 * <rot_x_rms>
```

where `rot_x_tot` and `rot_y_tot` are the rotational misalignment (in radians) of **H** around the y-axis and x-axis respectively. The small angle approximation is used, That is, it is assumed that both `rot_x_tot` and `rot_y_tot` are small compared to one. the quantities in brackets `<...>` are components of `pt`, and `r1` and `r2` are Gaussian distributed random numbers with unit rms. These random numbers are regenerated for each photon.

For Laure diffraction, since the **H** vector is approximately aligned with the $-x$-axis, the misalignment is characterized by rotations about the y-axis and z-axis. Notice that for both Bragg and Laue diffraction, the rotation around the y-axis misaligns (approximately) **H** in the plane of the diffraction.

Segmented grid

The general syntax for a segmented grid:

```
segmented = {
    active = <T/F>                      ! Turn on/off effect of grid. Default = True
    ix_bounds = (<ix_min>, <ix_max>),   ! Min/max index bounds in x-direction
    iy_bounds = (<iy_min>, <iy_max>),   ! Min/max index bounds in y-direction
    r0 = (<x0>, <y0>),                  ! (x,y) coordinates at grid origin
    dr = (<dx>, <dy>)                   ! Spacing between grid points.
        }
```

With a segmented grid the crystal surface is modeled as a grid of flat "rectangles" (the actual shape is very close but not quite rectangular). Using a segmented surface only makes sense when the surface is curved (see Eq. (5.1)). There is one rectangle for each grid point. Each rectangle has an extent in the (x, y) transverse dimensions equal to the spacing between grid points. Eq. (5.1) is used to calculate the z coordinate of the vertices of a given rectangle and then these z values are adjusted so that

1) The rectangle is flat (the vertices all lie on a plane).
2) The rectangle contacts the unsegmented surface (Eq. (5.1)) at two diagonally opposite vertices.
3) The other two diagonally opposite vertices will be as close as possible in the least squares sense from the unsegmented surface.

Note: The pt component is not used here.

5.12 Walls: Vacuum Chamber, Capillary and Mask

The wall attribute for an element is used to define:

vacuum chamber wall
capillary element (§4.6) inside wall
diffraction_plate (§4.13) geometry

The topics of the following subsections are:

§5.12.1 General wall syntax.
§5.12.2 Cross-section construction.
§5.12.3 Capillary and vacuum chamber wall interpolation.
§5.12.4 Capillary wall.
§5.12.5 Vacuum chamber wall.
§5.12.6 Diffraction_plate and mask element geometries.

5.12.1 Wall Syntax

The syntax of the wall attribute is:

```
wall = {
  superimpose = <T/F>,                ! Chamber wall only
  thickness = <real>                  ! Default thickness.
  opaque_material = <material_type>   ! Default opaque material.
  clear_material = <material_type>    ! Default clear material.
  section = {
    type = <section_type>,            ! Chamber Mask, and Diffraction_plate only
    s = <longitudinal_position>,      ! Relative to beginning of element.
                                      !   Not used for mask or diffraction_plate.
```

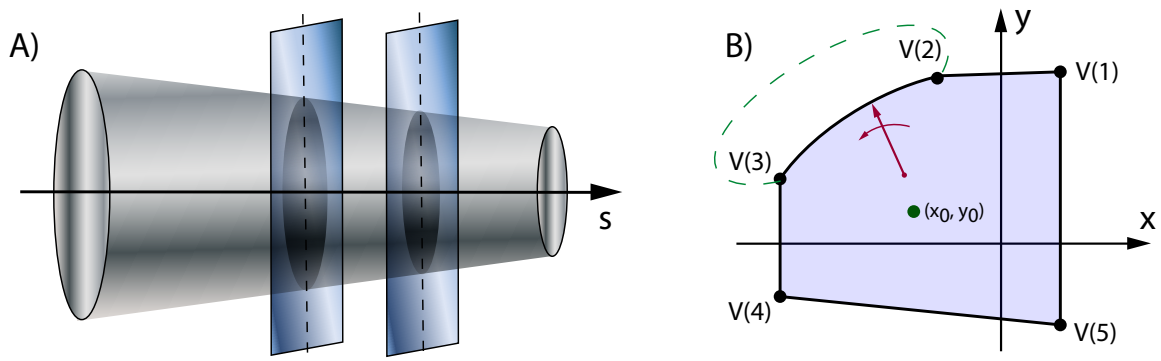

Figure 5.7: A) The inside wall of a capillary or the vacuum chamber wall of a non-capillary element is defined by a number of cross-sectional slices. B) Each cross-section is made up of a number of vertices. The segments between the vertices can be either a line segment, the arc of a circle, or a section of an ellipse.

```
    r0 = (<x0>, <y0>),                ! section (x,y) origin. Default = (0, 0).
    absolute_vertices = <T/F>,        ! Vertex relative to r0? Default = F.
    material = <material_type>,       ! Mask and Diffraction_plate only.
    thickness = <real>,               ! Mask and Diffraction_plate only.
    dr_ds = <value>,                  ! Capillary and Chamber only
    v(1) = {<x>, <y>, <radius_x>, <radius_y>, <tilt>},
    v(2) = { ... },
    ...},
  section = {
    s = <longitudinal_position>,
    v(1) = {... },
    ... },
  ... }
```

A `wall` begins with "`wall = {`" and ends with a "`}`". In between are a number of individual cross-section structures. Each individual cross-section begins with "`section = {`" and ends with a "`}`". The `s` parameter of a cross-section gives the longitudinal position of the cross-section. Example:

```
this_cap: capillary,
  wall = {
    section = { ! cross-section with top/bottom symmetry
      s = 0, v(1) =  {0.02, 0.00},
      v(2) = {0.00, 0.02, 0.02}, v(3) = {-0.01, 0.01} },
    section = {  ! Cross-section that is a tilted ellipse.
      s = 0.34,
      v(1) = {0.003, -0.001, 0.015, 0.008, 0.2*pi} } }
```

In this example an element called `this_cap` is a `capillary` whose wall is defined by two cross-sections.

5.12.2 Wall Sections

The wall is defined by a number of cross-sectional slices. For Fig. 5.7A shows the geometry for `capillary` or vacuum chamber walls. Each cross-section is defined by a longitudinal position s relative to the beginning of the element and a number of vertices. The arc between each vertex may be either a straight line, an arc of a circle, or a section of an ellipse. For a capillary it is mandatory that a cross-section be

convex. That is, given any two points within the cross-section, all points on the line segment connecting them must be within the cross-section.

The `v(<j>)` within a cross-section define the vertices for each cross-section. The vertices are defined with respect to the section origin given by `r0`. Each `v(<j>)` has five parameters. It is mandatory to specify the first two parameters `<x>` and `<y>`. Specifying the rest, `<radius_x>`, `<radius_y>`, and `<tilt>`, is optional. The default values, if not specified, is zero. The point (`<x>`, `<y>`) defines the position of the vertex. The parameters `<radius_x>`, `<radius_y>`, and `<tilt>` define the shape of the segment of the cross-section between the given vertex and the preceding one.

```
<radius_x>  = 0, <radius_y>  = 0   --> Straight line segment.
<radius_x> != 0, <radius_y>  = 0   --> Circular arc with radius = radius_x
<radius_x>  = 0, <radius_y> != 0   --> Illegal!
<radius_x> != 0, <radius_y> != 0   --> Ellipse section.
```

When an ellipse is specified, `<radius_x>`, and `<radius_y>` are the half width and half height of the semi-major axes and the `<tilt>` parameter gives the tilt of the ellipse. `<radius_x>` and `<radius_y>` must not be negative.

In the example above, for the first cross-section, `v(2)` specifies a non-zero `<radius_x>` and, by default, `<radius_y>` is zero. Thus the segment of the cross-section between `v(1)` and `v(2)` is circular in nature with a radius of 0.02. Since `v(3)` does not specify `<radius_x>` nor `<radius_y>`, the cross-section between `v(2)` and `v(3)` is a straight line segment.

The vertex points must be arranged in a "counter clockwise manner". For vertices `<v(i)>` and `<v(i+1)>` connected by a line segment this translates to

$$0 < \theta_{i+1} - \theta_i \pmod{2\pi} < \pi \tag{5.7}$$

where (r_n, θ_n) are the polar coordinates of the n^{th} vertex with respect to the point `r0`. For vertices connected by an arc, "counter clockwise manner" means that the line segment with one end at the center of the arc and the other end traversing the arc from `<v(i)>` to `<v(i+1)>` rotates in counter clockwise as shown in Fig. 5.7B.

The red line segment with one end at the center of the arc and the other end traversing the arc from, in this case, $V(2)$ to $V(3)$, rotates in counter clockwise manner. In general, there are two solutions for constructing such an arc. For positive radii, the solution chosen is the one whose center is closest to the section origin (x_0, y_0). If the radii are negative, the center point will be the point farthest from the origin (the dashed line between $V(2)$ and $V(3)$ in the figure).

A restriction on cross-sections is that the section origin (x_0, y_0) must be in the interior of any cross-section and that for any cross-section a line drawn from the origin at any given angle θ will intersect the cross-section at exactly one point as shown in Fig. 5.7B. This is an important point in the construction of the wall between cross-sections as explained below.

The last vertex specified, call it `<v(n)>`, should not have the same `<x>`, `<y>` values as the first vertex `<v(1)>`. That is, there will be a segment of the cross-section connecting `<v(n)>` to `<v(1)>`. The geometry of this segment is determined by the parameters of `<v(1)>`.

If there is mirror symmetry about the x or y axis for a cross-section, the "mirrored" vertices, on the "negative" side of the mirror plane, do not have to be specified. Thus if all the vertex points of a cross-section are in the first quadrant, that is, all `<x>` and `<y>` are zero or positive, mirror symmetry about both the x and y axes is assumed. If all the `<y>` values are zero or positive and some `<x>` values are positive and some are negative, mirror symmetry about the x axis is assumed. Finally, if all the `<x>` values are zero or positive but some `<y>` values are positive and some are negative, symmetry about the y axis is assumed. For example, for the first in the above example, since all the `<y>` values are non-negative and there are positive and negative `<x>` values, symmetry about the x axis is assumed.

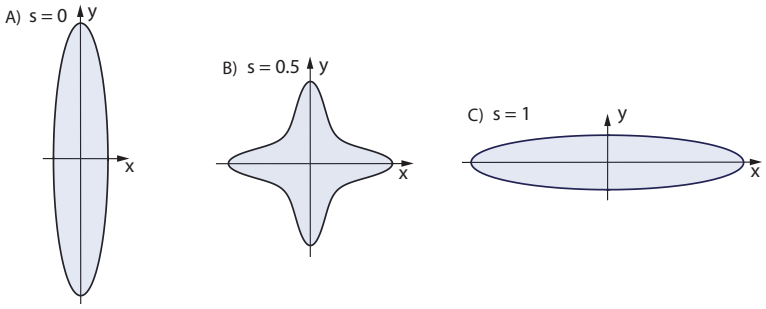

Figure 5.8: Example where convex cross-sections do not produce a convex volume. Cross-sections (A) and (C) are ellipses with a 5 to 1 aspect ratio. Half way in between, linear interpolation produces a convex cross-section as shown in (B).

The one exception to the above rule that (<x>, <y>) is the vertex center is when a single vertex v(1) is specified for a cross-section with a non-zero <radius_x>. In this case, (<x>, <y>) are taken to be the center of the circle or ellipse. For example, if a single vertex is specified for a cross-section as:

```
section = {s = 0.3, v(1) = {0.03, -0.01, 0.15, 0.08, 0.2}}
```

the cross-section will be an ellipse with center at $(0.03, -0.01)$ with a tilt of 0.2 and axes radii of 0.15 and 0.08. If a cross-section has a single vertex and <radius_x> is not specified, the cross-section is a rectangle. For example

```
section = {s = 0.3, v(1) = {0.03, 0.01}}
```

The vertices are defined with respect to the local sector origin r_0 except if absolute_vertices is set to True in which case the vertex numbers are taken as absolute. For example, the following two cross-sections are identical and describe a rectangle with edges at $x = 1$ and 5 and $y = -6$ and 6

```
section = {absolute_vertices = T, r0 = (4, 0),
                v(1) = {5, 6}, v(2) = {1, 6}, v(3) = {1, -6}, v(4) = {5, -6}}
section = {r0 = (4, 0), v(1) = {2, 6}}
```

Notice that while $r0$ is not needed in the first section for positioning of the vertices, it is needed in the first section to make Eq. (5.7) true.

5.12.3 Interpolation Between Sections

For capillary and vacuum chamber walls, the wall between cross-sections, is defined by interpolation. At a given s position, the r, θ coordinate system in the transverse x, y plane is defined with respect to an origin $\mathbf{r}_O(s)$ given by a linear interpolation of the origins of the cross-sections to either side of the given s position. Let s_1 denote the position of the cross-section just before s and s_2 denote the position of the cross-section just after s. Let $\mathbf{r}_{01}$ be the (x_0, y_0) origin defined for the cross section at s_1 and $\mathbf{r}_{02}$ be the (x_0, y_0) origin defined for the cross section at s_2. Then

$$\mathbf{r}_O(s) = (1 - \widetilde{s})\,\mathbf{r}_{01} + \widetilde{s}\,\mathbf{r}_{02} \tag{5.8}$$

where

$$\widetilde{s} \equiv \frac{s - s_1}{s_2 - s_1} \tag{5.9}$$

Let $r_{c1}(\theta)$ and $r_{c2}(\theta)$ be the radius of the wall as a function of θ for the cross-sections at $s = s_1$ and $s = s_2$ respectively. The wall $r_c(\theta, s)$ at any point s between s_1 and s_2 is then defined by the equation

$$r_c(\theta, s) = p_1(\widetilde{s})\, r_{c1}(\theta) + p_2(\widetilde{s})\, r_{c2}(\theta) \tag{5.10}$$

where p_1 and p_2 are cubic polynomials parameterized by

$$\begin{aligned} p_1 &= 1 - \widetilde{s} + a_1\,\widetilde{s} + a_2\,\widetilde{s}^2 + a_3\,\widetilde{s}^3 \\ p_2 &= \widetilde{s} + b_1\,\widetilde{s} + b_2\,\widetilde{s}^2 + b_3\,\widetilde{s}^3 \end{aligned} \tag{5.11}$$

If $a_i = b_i = 0$ for all $i = 1, 2, 3$, the interpolation is linear and this is the default if either of the parameters `dr_ds1` and `dr_ds2` are not given in the wall definition. These parameters are the slopes of the wall with respect to s at the end points

$$\text{dr_ds1} \equiv \left.\frac{d\overline{r}}{ds}\right|_{s=s_1}, \qquad \text{dr_ds2} \equiv \left.\frac{d\overline{r}}{ds}\right|_{s=s_2} \tag{5.12}$$

where $\overline{r}$ is the average r averaged over all θ. When *both* `dr_ds1` and `dr_ds2` are specified, the a_i and b_i are calculated so that the slopes of the wall match the values of `dr_ds1` and `dr_ds2` along with the constraints.

$$\begin{aligned} &p_1(0) = 1\,, \qquad p_1(1) = 0 \\ &p_2(0) = 0\,, \qquad p_2(1) = 1 \\ &M \equiv a_1^2 + a_2^2 + a_3^2 + b_1^2 + b_2^2 + b_3^2 \text{ is a minimum} \end{aligned} \tag{5.13}$$

The last constraint ensures a "smooth" transition between the two cross-sections.

To refer to a cross-section parameters after an element has been defined, the following syntax is used:

```
ele_name[wall%section(n)%v(j)%x]   ! x value of j^th vertex of n^th cross-section
```

5.12.4 Capillary Wall

For a `capillary`, `s` must be zero for the first cross-section and the length of the capillary is given by the value of `s` of the last cross-section.

For a `capillary`, in order for *Bmad* to quickly track photons, *Bmad* assumes that the volume between the cross-sections is convex. The volume will be convex if each cross-section $r_c(\theta, s)$ at any given s is convex. Note that it is *not* sufficient for $r_c(\theta, s)$ to be convex at the specified cross-sections as shown in Fig. 5.8. Also note that it is perfectly fine for the total capillary volume to not be convex.

5.12.5 Vacuum Chamber Wall

The vacuum chamber wall is independent of the element apertures (§5.8). Unless a program is specifically constructed, the presence of a vacuum chamber wall will not affect particle tracking.

The vacuum chamber wall defined for an element may be shorter or longer than the element. The vacuum chamber wall for a particular lattice branch is the sum of all the chamber walls of the individual elements. That is, the chamber wall at any given point is determined by interpolation of the nearest sections upstream and downstream to the point. Thus a given lattice element need not contain a `wall` component for the chamber wall to be well defined at the element.

The exception to the above rule is when a `section` has its `type` component set to either:

```
wall_start
wall_end
```

`wall_start` and `wall_end` sections must come in pairs. The next section after a `wall_end` section (if this section is not the last section in the lattice) must be a `wall_start` section. If a section has a `type`

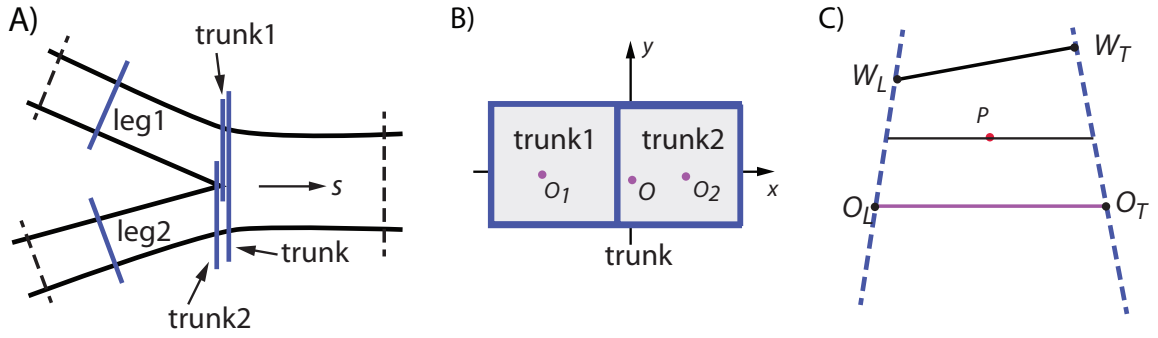

Figure 5.9: A) Crotch geometry: Two pipes labeled "leg1" and "leg2" merge into a single pipe called the "trunk" pipe. Five wall sections are used to define the crotch geometry (solid lines). Dashed lines represent sections not involved in defining the crotch. For purposes of illustration, the three trunk sections are displaced longitudinally but in reality must have the same longitudinal coordinate. B) Example layout of the trunk1, trunk2 and trunk wall sections. O_1, O_2 and O are the x_0, y_0 origins of the sections.

of `wall_start`, the region between that section and the previous section (which must be a `wall_end` section) will be considered to have no wall. If the `wall_start` section is the first section of the lattice branch, the region of no wall will start at the beginning of the branch. Similarly, if a section has a `type` of `wall_end`, the region between that section and the next section (or the end of the lattice branch if there is no next section) will not have a wall.

The chamber walls of any two elements may not overlap. The exception is when the `superimpose` attribute for a wall of an element is set to True. In this case, any other wall cross-sections from any other elements that overlap the superimposed wall are discarded. Superposition of a wall is useful, for example, in introducing mask regions into the wall.

If a branch has a closed geometry (§10.1), wall sections that extend beyond the ends of the branch are "wrapped" around.

If a particle is past the last wall cross-section or before the first wall cross-section, The following rules are used: If the branch has a `closed geometry`, the wall will be interpolated between the last and first cross-sections. If the branch has an `open` geometry, the wall is taken to have a constant cross-section in these regions.

The chamber wall is defined with respect to the local coordinate system (§16.1.1). That is, in a bend a wall that has a constant cross section is a section of a torus.

`Patch` elements (§4.41) complicate the wall geometry since the coordinate system at the end of the `patch` may be arbitrarily located relative to the beginning of the patch. To avoid confusion as to what coordinate system a wall section belongs to, `patch` elements are not allowed to define a wall. The wall through a patch is determined by the closest wall sections of neighboring elements.

Each section has a `type` attribute. This attribute is not used for `capillary` elements. For a vacuum chamber wall, the `type` attribute is used to describe a "crotch" geometry where two pipes merge into one pipe. The possible values for the `type` attribute are:

```
normal     ! default
leg1
leg2
trunk1
```

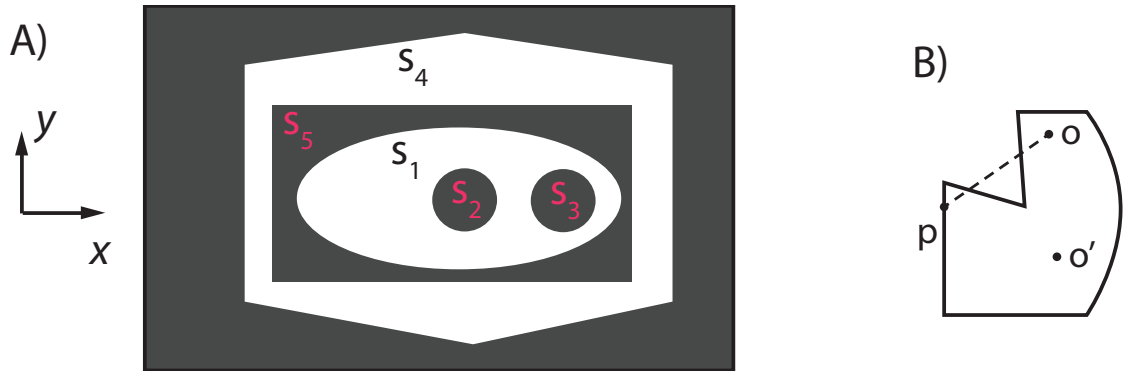

Figure 5.10: A) The surface of a `diffraction_plate` or `mask` element is divided into "clear" (white) and "opaque" (black) areas. As explained in the text, these areas are defined by five `sections` labeled s_1 through s_5. B) All wall sections must be star shaped with respect to the section's origin. In this example, The section is *not* star shaped since a line drawn from the origin point o to the point p on the boundary intersects the boundary twice in between. In this case the section can be made star shaped by moving the origin to o'.

```
trunk2
trunk
```

The geometry of a crotch is shown in Fig. 5.9A. Two pipes, called "leg1" and "leg2", merge into one pipe called the "trunk" pipe. The trunk pipe can be either upstream or downstream of the leg pipes. To describe this situation, five sections are needed: One section in each leg pipe which need to have their `type` attribute set to `leg1` and `leg2`, and three sections in the trunk with one having a `type` attribute of `trunk1`, another having a `type` attribute of `trunk2`, and the third having a `type` attribute of `trunk`. There can be no sections between the leg sections and the trunk sections.

All three trunk sections must be associated with the same element and have the same `s` value. In the list of sections of the element containing the trunk elements, the `trunk1` and `trunk2` sections must be listed first if the leg pipes are upstream of the trunk pipe (the situation shown in the figure) and must be listed last if the leg pipes are downstream. That is, the `trunk1` and `trunk2` sections are "between" the leg sections and the `trunk` section. It does not matter if `trunk1` is before or after `trunk2`.

The `trunk1` and `trunk2` sections must not overlap and the `trunk` section must be constructed so that its area is the union of the areas of `trunk1` and `trunk2`. An example is illustrated in Fig. 5.9B. Here the `trunk1` and `trunk2` sections are squares with origins labeled O_1 and O_2 in the figure. By necessity, these origins must be different since each must lie within the boundaries of their respective areas. The `trunk` section is a rectangle encompassing the two squares and has an origin labeled O.

Between `leg1` and `trunk1` sections the wall is interpolated using these two section. Similarly for the region between `leg2` and `trunk2` sections. Away from these regions interpolation is done as outlined in §5.12.3. However, these two regions need a different interpolation scheme since, `leg1` and `trunk1`, as well as `leg2` and `trunk2` sections do not have to be parallel to each other.

5.12.6 Mask Wall For Diffraction Plate and Mask Elements

The `wall` attribute (§5.12) of a `diffraction_plate` or `mask` element specifies what areas of the element will transmit or reflect particle and what areas will not. Here the "wall" defines a 2-dimensional area where particles (or X-rays) impinge upon and not a 3-dimensional surface. As such, the `s` longitudinal

position parameter of wall `sections` is not used with `mask` and `diffraction_plate` elements.

The algorithm used to decide if a particle hitting a given point will be transmitted or not is as follows. A `wall` is comprised of an ordered list of `sections` as discussed in §5.12.2. Sections will be labeled $s_1, s_2, s_3, \ldots$ with s_i being the i^{th} section defined in the `wall` structure. Each section of the wall must have its `type` attribute set to one of:

```
clear
opaque
```

A section is called "clear" or "opaque" depending upon the setting of its `type` attribute. The first section s_1 must be labeled `clear`. Each `clear` section has zero or more associated `opaque` sections that follow the `clear` section. An opaque section s_j is associated with clear section s_i if and only if i is less than j, and all the sections k between i and j (that is, $i < k < j$) are opaque sections. For example, if the `wall` defines 6 sections:

s_1	`Clear`
s_2	`Clear`
s_3	`Opaque`
s_4	`Opaque`
s_5	`Clear`
s_6	`Opaque`

with this arrangement, clear section s_1 does not have any associated opaque sections, clear section s_2 has two associated opaque section s_3 and s_4, and clear section s_5 has one associated opaque section s_6.

Each `section` covers an area specified by the vertex list associated with the `section`. To decide if a particle is transmitted, each clear section, starting from s_1, is tested to see if the particle's (x, y) coordinates is within the section. If the particle position is not within any clear section, the particle is considered to have hit an opaque region. If a particle is within one or more clear regions, let s_i be the clear region with the smallest index i that the particle is within. The particle is transmitted if the particle is outside of all associated opaque region of s_i.

When tracking photons, any clear section can be given a `material` and `thickness`. Available materials are listed in §5.9. A photon transversing a clear area with a defined material will be attenuated and have a phase shift. Note that `material` and `thickness` properties are not to be assigned to `opaque` sections.

To enable *Bmad* to quickly calculate whether a particle has landed on a clear or opaque section, All sections, both clear and opaque, must be "star shaped" with respect to the (x_0, y_0) origin used by the section. That is, a line drawn from the section origin to any point on the section boundary must not pass through any boundary points of the section in between. This is illustrated in Fig. 5.10B where the section is not star shaped since a line drawn from the origin o to the point p on the boundary passes through two boundary points in between. In this case the section can trivially be made star shaped by moving the origin to point o'. If it is not possible to make a section star shaped by moving the origin, the section must be divided into multiple sections.

An example geometry is shown in Fig. 5.10A. A wall that constructs this geometry is:

```
z_plate: diffraction_plate, wall = {
  section = {          ! s1
    type = clear,
    v(1) = {0, 0, 0.03, 0.013},
  section = {          ! s2
    type = opaque,
    v(1) = {0, 0, 0.005},
  section = {          ! s3
    type = opaque,
    r0 = (0.02, 0.00),
```

```
    v(1) = {0, 0, 0.005} }
  section = {          ! s4
    type = clear,
    v(1) = {0.04, 0}, v(2) = {0.04, 0.022},
    v(3) = {0, 0.03},
  section = {          ! s5
    type = opaque,
    v(1) = {0.032, 0.016},
```

There are two clear sections s_1 and s_4. Clear section s_1 has an oval shape and has two associated circular opaque sections s_2 and s_3. Clear section s_4 has a hexagonal shape and has one associated rectangular opaque section s_5. Notice that because opaque section s_5 is associated with clear section s_4, and since clear section s_4 comes after clear section s_1, opaque section s_5 does not affect clear section s_1 even though s_5 (completely) overlaps s_1. This shows that the ordering of the clear sections is important but the ordering of the opaque sections associated with a given clear section is not important.

5.13 Length Attributes

The length attributes are

```
l           = <Real>  !
l_chord     = <Real>  ! Chord length of a bend. Dependent attribute.
l_rectangle = <Real>  ! Rectangular length. See §4.5
```

The length `l` is the path length of the reference particle. The one exception is for an `rbend`, the length `l` set in the lattice file is the chord length (§4.5). internally, *Bmad* converts all `rbend`s to `sbend`s and stores the chord length under the `l_chord` attribute. Example:

```
b: rbend, l = 0.6   ! For rbends, l will be converted to l_chord
```

For a `girder` element the length `l` is a dependent attribute and is set by *Bmad* to be the difference in longitudinal position s of the downstream end of the last element supported relative to the upstream end of the first element.

For `wiggler`s, the length `l` is not the same as the path length for a particle with the reference energy starting on the reference orbit. See §16.1.1.

For `patch` elements the `l` length is, by definition, equal to `z_offset`. For `patch` elements, `l` is a dependent attribute and will be automatically set to `z_offset` by *Bmad*.

The length of a `capillary` element is a dependent variable and is given by the value of `s` of the last wall cross-section (§5.12.4).

The length of a crystal is zero for Bragg diffraction and is a dependent attribute dependent upon the crystal thickness for Laue diffraction. See §4.10 for more details.

5.14 Is_on Attribute

The `is_on` attribute

```
is_on = <Logical>
```

is used to turn an element off. Turning an element off essentially converts it into a drift. Example

```
q1: quad, l = 0.6, k1 = 0.95
q1[is_on] = False
```

is_on does not affect any apertures that are set. Additionally, is_on does not affect the reference orbit (§16.1.1) or reference energy (§17.1). For example, turning off an lcavity will not affect the reference energy.

The following elements cannot be "turned off:"

capillary	null_ele
crystal	hybrid
drift	mirror
fiducial	multilayer_mirror
floor_shift	photon_init
patch	sample

A related parameter is multipoles_on (§5.15).

5.15 Multipole Attributes: Magnetic and Electric

Multipole formulas for are given in §18.1 and §18.2. Note that the setting of field_master (§5.2) will determine if multipoles are interpreted as normalized or unnormalized.

A multipole (§4.36) element specifies its magnetic multipole components using normal and skew components with a tilt

```
KnL  = <Real>  ! Normal component. n = Integer.
KnSL = <Real>  ! Skew component. n = Integer.
Tn   = <Real>  ! Tilt. n = Integer. Default is $pi$/(2n + 2)
```

Where n is an integer in the range from 0 (dipole component) through 21. If Tn is given without a value, a default of $pi/(2n + 2)$ will be used producing a skew field. Example:

```
m: multipole, k1l = 0.32, t1  ! Skew quadrupole of strength 0.32
```

Following MAD, a non-zero dipole (K0L component will affect the reference orbit (just like a normal dipole will). This is not true for any other element.

An ab_multipole (§4.1) specifies magnetic multipoles using normal (Bn) and skew (An) components:

```
An = <Real>
Bn = <Real>
```

Here n ranges from 0 (dipole component) through 21. Example:

```
q1: ab_multipole, b2 = 0.12, a20 = 1e7, field_master = T
```

Note that there is a factor of n-factorial difference between An, Bn and KnL, KnSL multipoles §18.1.

Elements like quadrupoles and sextupoles can have assigned to them both magnetic and electric multipole fields. In this case, the magnetic fields are specified using the same convention as the ab_multipole. For such non-multipole elements, the magnetic multipole strength is scaled via (Eq. (18.16))

$$a_n(\text{actual}) = F\,\frac{r_0^{n_{\text{ref}}}}{r_0^n}\,a_n(\text{input}), \qquad b_n(\text{actual}) = F\,\frac{r_0^{n_{\text{ref}}}}{r_0^n}\,b_n(\text{input}) \tag{5.14}$$

where "input" denotes the input value set in the lattice file, "actual" denotes the value that is used to compute the field, F is the strength of the element (for example F is $K1 \cdot L$ for a quadrupole), and r_0 is the "measurement radius" and is set by the r0_mag attribute. The default value of r_0 is 0 in which case the factor of $r_0^{n_{\text{ref}}}/r_0^n$ is omitted. The scaling may be turned off altogether by setting the scale_multipoles attribute. Example:

```
q1: quadrupole, b2 = 0.12, a20 = 1e7, scale_multipoles = F
```

Electric multipoles are specified using normal (Bn_elec) and skew (An_elec) components.

```
An_elec = <Real>
Bn_elec = <Real>
```

Here `n` ranges from 0 (dipole component) through 21. Like the magnetic multipoles, a measurement radius `r0_elec` can be used to scale the multipoles as explained in §18.2. Example:

```
q1: quadrupole, l = 1.2, b2_elec = 1e6, r0_elec = 0.034
```

See §18.2 for how electric multipoles are defined. Notice that Electric multipoles are never scaled by the element's field strength as they are with magnetic multipoles. If the value of `r0_elec` is zero (the default) the multipoles will not be scaled.

Unlike magnetic multipoles, there are no factors of the reference momentum nor the element length in the definition for electric multipoles. That is, electric multipole values represent the field and not the normalized integrated field. Thus an electric multipole associated with a zero length element will have no effect on tracking. This being the case, *Bmad* does not allow electric multipole values to be specified for `multipole` and `ab_multipole` elements. Indeed, in the limit of zero element length at constant integrated electric field strength, the equations of motion are singular since, unlike the magnetic case, the infinite fringe fields give rise to infinite energy shifts.

The magnetic and electric multipole kick can be toggled on or off using the `multipoles_on` attribute. Example:

```
call, file = "lattice.bmad"             ! Read in a lattice file
quadrupole::*[multipoles_on] = False    ! But I want the multipoles off.
q1[k1] = 0.3                            ! k1 attribute not affected.
```

`multipoles_on` only effect multipoles specified by `An`, `Bn`, `An_elec`, or `Bn_elec`. Other multipoles, like the `k2` multipole of a `sextupole`, are not affected. The exception is `multipole` and `ab_multipole` elements do not have the `multipoles_on` attribute. Rather they can be toggles on/off using the `is_on` attribute.

5.16 Field Maps

There are two general ways to specify complicated electro-magnetic field configurations that cannot be simply modeled using multipoles. One way is to use `custom` fields. Specifying a custom field is done by using custom code and linking this code with *Bmad* into a program. That is, custom fields are defined outside of the *Bmad* software (§6.4).

The other way to specify a complicated field is to use a "`field map`". There are four types of field maps:

```
cartesian_map       ! §5.16.2
cylindrical_map     ! §5.16.3
grid_field          ! §5.16.4
gen_gradients        ! §5.16.5
```

Essentially, `cylindrical_map` and `cartesian_map` define fields using a set of functions with user defined coefficients with the functions formulated to obey Maxwell's equations. The `grid_field` type defines the field on a grid of points and interpolation is used to evaluate the field in between the points. Finally, the `gen_gradients` type defines a set of "generalized gradients" (§18.7).

The `cylindrical_map` and `grid_field` types can be used with both RF and DC fields. The other two types can only be used with DC fields. RF fields may only be used with the following element classes:

```
e_gun         ! §4.15
em_field      ! §4.17
lcavity       ! §4.31
rfcavity      ! §4.46
```

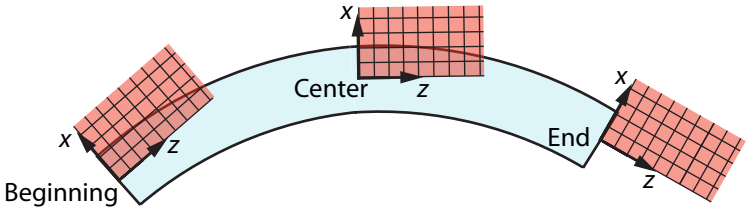

Figure 5.11: When used with a bend element, by default, field map coordinates will be Cartesian and not curved like the reference orbit. The orientation of the field map coordinates is determined by the setting of `ele_anchor_pt`. To use curvilinear coordinates instead, `curved_ref_frame` must be set to `True` [Available in `grid_field` and `gen_gradients` only].

An element may specify multiple fields of a given type and/or may define multiple fields of different types. In both these cases, the field in the element is taken to be the sum of the individual fields. For example:

```
sb: sbend, field_calc = fieldmap, cylindrical_map = {...},  cylindrical_map = {...}
```

In this example an element has two `cylindrical_map` fields and the total field is the sum of the fields of each one. Separating fields like this can be useful, for example, to decouple the specification of electric from magnetic fields, or to decouple the specification of AC and DC fields.

The field of one element can overlap onto other elements. This is explained in Sec. §5.18.

Field maps are used with integration type tracking methods (§6.4). It is important to note that field maps are *ignored* by `bmad_standard` tracking. Additionally, `grid_field` field maps cannot be used with `symp_lie_ptc`.

Field maps may extend longitudinally beyond the ends of an element. See Sec (§5.18) for more details.

In a lattice file, once a field map is defined for an element, components of the field map may be redefined using the notation

```
ele_name[fieldmap_name(index)%component_name] = value
```

where `ele_name` is the name of the element, `fieldmap_name` is the name of the type of field map, `index` is the index of the field map which is "1" for the first field map defined for an element, etc., `component_name` is the name of the component, and `value` is the value to set to. Example:

```
qq, quadrupole, grid_field = {field_scale = 0.5, ...}, ...
qq[grid_field(1)%field_scale] = 0.7  ! Change field_scale value
```

5.16.1 Field Map Common attributes

This section explains some of the attributes that are common to the `field map` types. Not all attributes are used in all field map types. See the documentation on the individual types for a list of the attributes pertinent to that type.

curved_ref_frame

For bends, the coordinates of the field are, by default, Cartesian and do not follow the curved bend coordinates. The orientation of the field map coordinates with respect to the bend is determined by the placement of the anchor point (specified by `ele_anchor_pt`) as shown in Fig. 5.11. In this case, when tracking a particle, *Bmad* will convert particle coordinates (which are expressed in the bend's curvilinear coordinate system defined by the reference orbit) to the Cartesian coordinates of the field map and will rotate the computed field from the field map coordinates back to the particle coordinates.

For grid_field and gen_gradients types only, this default behavior can be changed by setting the curved_ref_frame component of the field map to True. In this case, the field grid coordinates will follow curved bend coordinates. The curved_ref_frame parameter is only pertinent for bend elements (sbends, rbends). The setting of curved_ref_frame is ignored for non-bend elements.

ele_anchor_pt

The ele_anchor_pt, along with r0, determines the origin of the field with respect to the lattice element. Possible settings are:

```
beginning   ! Beginning of element (default).
center      ! Center of element.
end         ! Exit end of element.
```

Example:

```
rfc0: rfcavity, gen_gradients = {ele_anchor_pt = center, ...}, ...
```

field_type

The field_type attribute sets the type of field described. Possible settings for field_type are:

```
electric      ! Pure electric field. For DC fields only.
magnetic      ! Pure magnetic field. For DC fields only.
mixed         ! Mixed EM fields. Used for grid_field only.
```

Example:

```
bb: sbend, cartesian_map = {field_type = electric, ...}, ...
```

The cylindrical_map type does not have a field_type since it has explicit arrays for the electric and magnetic fields.

field_scale

The field_scale attribute is used to scale the overall field magnitude. The default value is 1. A value of -1 will reverse the field. If the master_parameter is defined, it is multiplied with the field_scale to give the overall scale. Example:

```
qq, quadrupole, grid_field = {field_scale = 0.5, ...}, ...
qq[grid_field(1)%field_scale] = 0.7     ! Change value after element def.
qq[grid_field(1)%master_parameter] = k1 ! Change value after element def.
```

harmonic

The harmonic attribute, along with rf_frequency element attribute, sets the oscillation frequency of the field map. The harmonic attribute is only used with cylindrical_map and grid_field types. The default value of harmonic is 0. The harmonic number needs to be 0 for DC fields. Example:

```
lc1: lcavity, rf_frequency = 500e6, grid_field = {harmonic = 2, ...}, ...
```

Notice that rf_frequency is set outside of any field map and is common to all field maps.

master_parameter

The master_parameter defines a "master" element attribute for scaling the field. Example:

```
qq: quadrupole, gen_gradients = {master_parameter = "K1", ...}, k1 = ...
```

This example defines the master_parameter for the gen_gradients to be the quadrupole strength k1. By using the same master parameter for a set of field map instances within a given lattice element, the sum field of the set can be controlled by a single attribute. The master_parameter must be set to a valid element attribute. If the name is blank (""), no master parameter is used. The master_parameter, if defined, is multiplied with the field_scale to give the value used to scale the fields. The default master_parameter is blank ("") except for wiggler elements where, for historical reasons, the default is polarity.

phi0_fieldmap

For AC fields, phi0_fieldmap is the phase of the field map field relative to the fundamental mode. The phase phi0_fieldmap is relative to the fundamental frequency and not the frequency of the field map mode. That is, the "zero crossing" point of the field map is shifted by a time phi0_fieldmap$/f_0$ where f_0 is the fundamental mode frequency.

r0

The r0 attribute is the $(x0, y0, z0)$ vector specifying the offset of the origin point that defines the field relative to the anchor point defined by ele_anchor_pt. The origin position of the field (r_origin) is determined by

```
r_origin = r0 + r_anchor
```

where r_anchor is determined by the setting of ele_anchor_pt. In the reference coordinates (§16.1.1) with respect to the element r_anchor is:

```
ele_anchor_pt       r_anchor
-------------       ---------
beginning           (0, 0, 0)      ! Default
center              (0, 0, L/2)
end                 (0, 0, L)
```

with L being the length of the element.

Example:

```
rfc0: rfcavity, gen_gradients = {r0 = (-0.23, ...), ...}, ...
```

5.16.2 Cartesian_Map Field Map

The cartesian_map is used to describe DC fields. Each term of a cartesian_map is a solution of Laplace's equation in cartesian coordinates. as described in Sec. §18.5.

The lattice file syntax for the cartesian_map type is:

```
cartesian_map = {
  field_type       = <String>, ! Type of field: Default = Magnetic.
  field_scale      = <Real>,   ! Scale factor for the E & B fields.
  master_parameter = <Name>,   ! Master scaling parameter for E & B fields.
  ele_anchor_pt    = <Real>,   ! Anchor position: Beginning (default), Center, or End.
  r0               = (<x0>, <y0>, <z0>), ! Anchor offset. Default is 0.
  term = {<A>, <k_x>, <k_y>, <k_z>, <x_0>, <y_0>, <phi_z>, <family>},
  term = {....},
  ... more terms ...
}
```

The possible settings of <family> are explained in Sec. §18.5. Example:

```
q01: quadrupole, l = 0.6, field_calc = fieldmap,
       cartesian_map = {
         term = {0.03, 3.00, 4.00, 5.00, 0, 0, 0.63, y},
         term = {...}, ...    }
```

See Sec. §5.16.1 for an explanation of the attributes that are common with other field map types.

See Sec. §3.8 for the syntax of setting Cartesian map components.

To use with PTC dependent tracking methods (§6.4) there are a number of restrictions:

- There can be only one cartesian_map field map and there cannot be any other field maps of any kind.
- cartesian_map may not be used with a bend.
- Only magnetic fields may be used.
- The transverse terms in r0 must be zero.
- Since PTC evaluates the vector potential (§26.25), and since k_z appears in the denominator of some terms, k_z must be non-negative for all terms.

5.16.3 Cylindrical_Map Field Map

The `cylindrical_map` is used to describe both DC and AC fields. Each term of a `cylindrical_map` is a solution of Laplace's equation in cylindrical coordinates. as described in Sec. §18.6.

The lattice file syntax for the `cylindrical_map` type is:

```
cylindrical_map = {
  field_scale      = <Real>,    ! Scale factor for the E & B fields.
  master_parameter = <Name>,    ! Master scaling parameter for E & B fields.
  ele_anchor_pt    = <Real>,    ! Anchor position: Beginning (default), Center, or End.
  m                = <Integer>, ! Azimuthal mode number
  harmonic         = <Integer>, ! RF frequency harmonic number
  phi0_fieldmap    = <Real>,    ! Phase of oscillations.
  theta0_azimuth   = <Real>,    ! Azimuthal orientation.
  r0               = (<x0>, <y0>, <z0>), ! Anchor offset. Default is 0.
  dz        = <Real>,                    ! Distance between sampled field points.
  e_coef_re = (<Real>, <Real>, ....),    ! Real part of E.
  e_coef_im = (<Real>, <Real>, ....),    ! Imaginary part of E.
  b_coef_re = (<Real>, <Real>, ....),    ! Real part of B.
  b_coef_im = (<Real>, <Real>, ....),    ! Imaginary part of B.
}
```

See Sec. §5.16.1 for an explanation of the attributes that are common with other field map types.

For DC fields, the `e` coefficients specify the electric fields and the `b` coefficients specify the magnetic fields. For AC fields, the `e` coefficients specify modes that have finite longitudinal electric fields while the modes associated with the `b` coefficients do not.

To specify the RF frequency, specify the `rf_frequency` *element* attribute along with the `harmonic` attribute. See the discussion of the `harmonic` attribute in Sec. §5.16.1.

The basic equations used for the `cylindrical_map` decomposition of the fields are given in Section §18.6. A lattice element may have multiple `cylindrical_map` components with each `cylindrical_map` being associated with a particular azimuthal mode m.

`e_re` and `e_im` give the real an imaginary part of e and `b_re` and `b_im` give the real and imaginary part of b. All of these vectors must be present and have the same length. The exception is with an $m = 0$ mode either the e or b arrays can be omitted and will default to zero. The number of terms N for the e or b vectors must be a power of 2 and all modes must have the same number of terms. The n^{th} element in the e or b arrays, with n running from 0 to $N-1$, is associated with a wavelength k_n

$$k_n = \begin{cases} \frac{2\,\pi\,n}{N\,dz} & 0 \le n < \frac{N}{2} \\ \frac{2\,\pi\,(n-N)}{N\,dz} & \frac{N}{2} \le n \le N-1 \end{cases} \tag{5.15}$$

This convention produces less high frequency components then the convention of using $k_n = 2\,\pi\,n/Ndz$.

The longitudinal length of the field is

$$L_{\rm field} = \frac{N-1}{dz} \tag{5.16}$$

this may be different from the length `l` specified for the element.

For AC fields, the time t in Eq. (18.46) is computed depending upon whether `absolute time tracking` or `relative time tracking` is being used as discussed in §26.1. For `rfcavity` elements, the phase factor ϕ_{0j} in Eq. (18.46) is computed by

ϕ_{0j} `= harmonic(j) * [0.25 - (phi0 + phi0_multipass + phi0_err +`
`phi0_autoscale + phi0_fieldmap(j))]`

where `phi0_fieldmap(j)` and `harmonic(j)` are specific to the jth grid field while the other factors are element parameters and so will be the same for all grid field maps of a given element. For non `rfcavity` elements the phase is

ϕ_{0j} = `harmonic(j) * [phi0 + phi0_multipass + phi0_err + phi0_autoscale + phi0_fieldmap(j)]`

where `phi0_fieldmap(j)` and `harmonic(j)` are specific to the jth cylindrical field map while the other factors are element parameters and so will be the same for all cylindrical field maps of a given element.

Example:

```
m1: lcavity, rf_frequency = 1e6, voltage = 2e6, cylindrical_map = {
  m = 2,                   harmonic = 3,
  r0 = (0, 0, 0.001),      dz = 0.1,
  theta0_azimuth = 0.3,    field_scale = 0.7,
  ele_anchor_pt = center,  master_parameter = voltage,
  e_coef_re = (...),       e_coef_im = (...),
  b_coef_re = (...),       b_coef_im = (...)}, field_calc = fieldmap
```

See Sec. §3.8 for the syntax of setting Cylindrical map components.

Note: When using PTC based tracking (§6), the following restrictions apply:

- The fields must be DC.
- all the `e_coef` and `b_coef` arrays must have the same length.
- `r0(1)` and `r0(2)` (the transverse offsets) must be zero.
- The element containing the map cannot be an `sbend` or `rbend`.
- May not be combined with other field map types.

5.16.4 Grid_Field Field Map

A `grid_field` is grid of field points specified using the syntax:

```
grid_field = {
  geometry         = <String>,    ! Geometry of the grid.
  field_type       = <String>,    ! Type of field: Default = Mixed.
  field_scale      = <Real>,      ! Scale factor for the E & B fields.
  phi0_fieldmap    = <Real>,      ! Phase of oscillations.
  harmonic         = <Int>,       ! RF frequency harmonic number
  interpolation_order = <Int>,    ! 1 (default) or 3 interpolation polynomial order.
  master_parameter = <Name>,      ! Master scaling parameter for E & B fields.
  curved_ref_frame = <Logical>,   ! Use a curved reference frame with bends?
  r0   = (...),                   ! Grid origin. Syntax is geometry dependent.
  dr   = (...),                   ! Grid spacing. Syntax is geometry dependent.
  ele_anchor_pt = <Position>      ! BEGINNING, CENTER, or END
  { ... },                      ! Field table. Syntax is geometry dependent.
  pt(<Integer>, ...) = ( ... ), ! Field table point. Old style.
}
```

See Sec. §5.16.1 for an explanation of the attributes that are common with other field map types.

To specify the RF frequency, specify the `rf_frequency` *element* attribute along with the `harmonic` attribute. See the discussion of the `harmonic` attribute in Sec. §5.16.1.

For `field_type` set to `electric` or `magnetic`, the field is DC. That is, For `field_type` set to `electric` or `magnetic`, the value of `harmonic` must be 0. For `field_type` set to `mixed`, the field may be DC or AC.

For AC fields, the individual field components are complex. the syntax for specifying a complex number is:

```
(<Re> <Im>)
```

Example:

```
{
  0 0 -7: (0.34 -4.3) (2.37 9.34) ...,    ! Complex field
  0 0 -7: 0.12 -0.33 ...,                 ! Imaginary components are zero
  ...
}
```

The actual fields $\mathbf{E}$ and $\mathbf{B}$ are computed from the complex fields $\mathbf{E}_c$ and $\mathbf{B}_c$ via

$$\mathbf{E} = \Re\Big[\mathbf{E}_c \exp\left(-2\,\pi\,i\,(\phi_t + \phi_{\rm ref})\right)\Big] \tag{5.17}$$

with a similar equation for $\mathbf{B}$. ϕ_t is the part of the phase due to when the particle arrives at the cavity and depends upon whether `absolute time tracking` or `relative time tracking` is being used as discussed in §26.1. The phase $\phi_{\rm ref}$ for the j^{th} grid field in an `rfcavity` element is

$\phi_{\texttt{ref,j}}$ `= harmonic(j) * [0.25 - (phi0 + phi0_multipass + phi0_err +`
`phi0_autoscale + phi0_fieldmap(j))]`

where `phi0_fieldmap(j)` and `harmonic(j)` are specific to the jth grid field while the other factors are element parameters and so will be the same for all grid field maps of a given element. For non `rfcavity` elements the phase is

$\phi_{\texttt{ref,j}}$ `= harmonic(j) * [phi0 + phi0_multipass + phi0_err +`
`phi0_autoscale + phi0_fieldmap(j)]`

The `geometry` switch sets the type of the grid and must come before the field table. The possible settings of `geometry` are:

```
rotationally_symmetric_rz
xyz
```

The `rotationally_symmetric_rz` setting for `geometry` is for fields that are rotationally symmetric around the z axis. The format for this type of `grid_field` is

```
grid_field = {
  geometry = rotationally_symmetric_rz,
  r0   = (<x0>, <y0>, <z0>),        ! Grid origin
  dr   = (<dr>, <dz>),              ! Grid spacing
  {                                ! Field table...
    <ir> <iz>: <E_r> <E_phi> <E_z>,   ! For field_type = Electric
    <ir> <iz>: <B_r> <B_phi> <B_z>,   ! For field_type = Magnetic
    <ir> <iz>: <E_r> <E_phi> <E_z> <B_r> <B_phi> <B_z>,
                                      ! For field_type = Mixed.
    ...
  }
}
```

where `<iz>` can be negative but `<ir>` must be non-negative. Notice that commas are only used in the field table to demarcate between field points.

There is an old style syntax for the field table that looks like:

```
    pt(<ir>, <iz>) = (<E_r>, <E_phi>, <E_z>) ! For field_type = Electric
    pt(<ir>, <iz>) = (<B_r>, <B_phi>, <B_z>) ! For field_type = Magnetic
    pt(<ir>, <iz>) = (<E_r>, <E_phi>, <E_z>, <B_r>, <B_phi>, <B_z>)
                                             ! For field_type = Mixed.
```

While the old style is accepted, parsing the old syntax is almost a factor of two slower. The only possible advantage to the old style is that expressions can be used for field values.

The `xyz` setting for `geometry` can be used for all rectangular field grids. The format for this type of `grid_field` is

```
  grid_field = {
    geometry = xyz,
    r0   = (<x0>, <y0>, <z0>),    ! Grid origin
    dr   = (<dx>, <dy>, <dz>),    ! Grid spacing
    }
      <ix> <iy> <iz>: <E_x> <E_y> <E_z>,  ! For field_type = Electric
      <ix> <iy> <iz>: <B_x> <B_y> <B_z>,  ! For field_type = Magnetic
      <ix> <iy> <iz>: <E_x> <E_y> <E_z> <B_x> <B_y> <B_z>,
                                                ! For field_type = Mixed.
      ...
    }
  }
```

where `<ix>`, `<iy>`, and `<iz>` can be negative.

The old style syntax for the `xyz` field table is:

```
    pt(<ix>, <iy>, <iz>) = (<E_x>, <E_y>, <E_z>),  ! For field_type = Electric
    pt(<ix>, <iy>, <iz>) = (<B_x>, <B_y>, <B_z>),  ! For field_type = Magnetic
    pt(<ix>, <iy>, <iz>) = (<E_x>, <E_y>, <E_z>, <B_x>, <B_y>, <B_z>),
                                                ! For field_type = Mixed.
```

[For clarity sake, the following discusses the `xyz` case. Extension to other cases is straight forward.] There is no restriction on the bounds of the indexes `(ix, iy, iz)` of the `pt(ix, iy, iz)` array. A point (ix, iy, iz) corresponds in space to the point (x, y, z):

```
  (x, y, z) = dr * (ix, iy, iz) + r0 + r_anchor
```

where `z` is measured from the beginning of the element and `r_anchor` is determined by the setting of `ele_anchor_pt`:

```
  ele_anchor_pt       r_anchor
  -------------       ---------
  beginning           (0, 0, 0)      ! Default
  center              (0, 0, L/2)
  end                 (0, 0, L)
```

with `L` being the length of the element.

Example:

```
  apex: e_gun, l = 0.23, field_calc = fieldmap, rf_frequency = 187e6,
                                  grid_field = call::apex_gun_grid.bmad
```

with the file `apex_gun_grid.bmad` being:

```
  {
    geometry = rotationally_symmetric_rz,
    harmonic = 1,
    master_parameter = voltage,
    r0 = (0, 0),
```

```
    dr = (0.001, 0.001),
    {
      0 0:  (0 0) (0 0) (1 0)     (0 0) (0 0) (0 0),
      0 1:  (0 0) (0 0) (0.99 0)  (0 0) (0 0) (0 0),
      ...
    }
  }
```

The order of the polynomials used to interpolate the field is determined by the setting of `interpolation_order`. The value of this integer may be either `1` or `3`. A value of `1` is the default and gives linear interpolation. A value of `3` will give cubic interpolation which will be better if the field values on the grid are smooth and well behaved. However, a cubic interpolation will be slower and will magnify gird field errors so care should be taken in choosing the interpolation order.

It is considered an error if the field of the grid is evaluated for a point that is transversely outside of the grid. That is, a grid must extend transversely to the aperture or at least beyond the trajectory of any particle. [Actually, to prevent problems when the aperture is set at the grid boundary, if the distance between the particle and the grid boundary is within $1/2$ of the spacing between grid points, no error is generated and the field will be calculated using extrapolation.] On the other hand, it is acceptable to evaluate the grid field at a point that is longitudinally outside of the grid. In this case, the field is assumed to go to zero. This is done by effectively adding to the grid two planes of zero field longitudinally to either side of the grid. So a particle traveling outside of the grid longitudinally will see the field drop to zero within one longitudinal grid spacing length.

Grid fields may stored in `HDF5` binary format which may then be called using an `inline call` (§3.20). For example:

```
qq: quadrupole, grid_field = call::my_grid.h5, ...
```

See Sec. §3.8 for the syntax of setting `grid_field` components.

5.16.5 Gen_Gradients

The `gen_gradients` attribute characterizes DC magnetic or electric fields using "generalized gradients" (GG). *Bmad* uses the curved-curvilinear-coordinate formalism (which reduces to the straight-coordinate case when the reference-frame curvature is zero) as described by Van der Schueren et al.[Schueren24]. The field is characterized by the midplane derivative functions $a_n(s)$ (skew), $b_n(s)$ (normal), and $b_s(s)$ (solenoidal); formulas are given in Sec. §18.7.

Note: Before any tracking is done with GG, the GG coefficients need to be computed and the values stored in a lattice file. Typically, GG coefficients are computed by fitting field data. A fitting program is provided by the SciBmad package `GeneralizedGradients.jl`[GenGrad]. This package also provides a translator to translate from SciBmad format to Bmad format.

> **Note:** The older straight-coordinate `gen_grad_map` format (Venturini–Dragt azimuthal harmonics $C_{m,\sin/\cos}$) has been removed. Old lattice files that use the old `gen_grad_map` syntax will generate an error and must be converted to the new `gen_gradients` format using the `gen_grad_map_to_gen_gradients` utility (§A). Also note that while the old and new formalism will agree exactly on the fit planes, interpolation of the GG coefficients between fit planes will differ. Thus there will be shifts in tracking results between old and new.

An element can store a set of GGs. For example, a GG for the magnetic field and a GG for the electric field. The syntax for describing a single GG is:

```
gen_gradients = {
  field_type       = <String>,    ! Type of field: Default = magnetic.
  field_scale      = <Real>,      ! Scale factor for the E & B fields.
  master_parameter = <Name>,      ! Master scaling parameter for E & B fields.
  g_ref            = <Real>,      ! Reference-frame curvature 1/rho. Default = g of element.
  ele_anchor_pt    = <Real>,      ! Anchor position: Beginning (default), Center, or End.
  r0               = (<x0>, <y0>, <z0>), ! Anchor offset. Default is 0.
  dz               = <Real>,             ! Distance between sampled field points.
  curve = {                       ! Generalized gradient curve
    kind = <a-or-b-or-bs>,        ! Curve family: a (skew), b (normal), or bs (solenoid).
    n = <Int>,                    ! Harmonic index (n = 0 for bs).
    derivs = {
      <z0>: <d0> <d1> <d2> ... <dN>,
      <z1>: ...,
      ...
    }
  },
  curve = {
    ...
  }
}
```

Example:

```
t1: sbend, g = 0.37, gen_gradients = {
  field_type = magnetic,   ele_anchor_pt = beginning,
  dz = 0.1,                r0 = (0, 0, 0),
  field_scale = 1.3,       g_ref = 0.37,
  curve = { kind = b, n = 1, derivs = { 0: 0.5 0.1, 0.1: 0.4 0.05 } },
  ...  }, field_calc = fieldmap, integrator_order = 6, num_steps = 70,
tracking_method = taylor, mat6_calc_method = taylor
```

See Sec. §5.16.1 for an explanation of the attributes that are common with other field map types. The `g_ref` component is the curvature $1/\rho$ of the reference frame in which the GGs are defined; it is zero for a straight frame and, for a bend, is normally equal to the element's bending strength `g`. In general, there will be multiple `curve` instances. A given `curve` instance specifies one generalized gradient with a `kind` (`a` for skew, `b` for normal, or `bs` for the solenoidal term) and a harmonic index `n` (`n` = 1 is dipole, `n` = 2 is quadrupole, etc.; `n` = 0 for `bs`). Each `curve` has a table of derivatives given by the `derivs` component. Each row in the derivative table starts with the z-position of the derivatives `<z0>`, `<z1>`, `z2`, ... etc. The z-positions must be in increasing order and be multiples of `dz` with no gaps (that is, `<z1>` = `<z0>` + `dz`, etc.). Each line in the `derivs` table specifies the derivatives `<d0>`, `<d1>`, etc. at the given z-position where `<d0>` is the zeroth derivative (the value of the generalized gradient itself), `<d1>` is the first derivative, etc. For a given `curve` instance, all derivative rows must have the same number of derivatives. Different `curves` can have differing number of derivatives. All curves must specify the same z-positions but different curves may have differing number of derivatives.

Like `grid_field`s, the edges of a `gen_gradients` map can be defined to be outside of the edges of the element. These areas will be ignored while tracking unless field overlap is defined (§5.18)).

See Sec. §3.8 for the syntax of setting `gen_gradients` components.

To use with PTC dependent tracking methods (§6.4) there are a number of restrictions:

- There can be only one `gen_gradients` map and there cannot be any other field maps of any kind.
- Only magnetic fields may be used.

The integrator_order for PTC dependent tracking can be either set to 4 or 6 with 4 being the default§6.5.3.

5.17 RF Couplers

For lcavity and rfcavity elements, the attributes that characterize the dipole transverse kick due to a coupler port are:

```
coupler_at       = <Switch> ! What end the coupler is at
coupler_strength = <Real>   ! Normalized strength
coupler_angle    = <Real>   ! Polarization angle (rad/2π)
coupler_phase    = <Real>   ! Phase angle with respect to the RF (rad/2π)
```

The possible coupler_at settings are:

```
entrance_end
exit_end  ! default
both_ends
```

The kick due to the coupler is

```
dP_x = amp * cos(phase) * cos(angle)
dP_y = amp * cos(phase) * sin(angle)
dE   = amp * (cos(angle) * x + sin(angle) * y) * sin(phase) * twopi * rf_frequency / c_light
```

where dP_x and dP_y are the transverse momentum kicks, dE is an energy kick, and

```
amp   = gradient * coupler_strength
phase = twopi * (phase_particle + phase_ref + coupler_phase)         ! For lcavity §4.31
      = pi/2 + twopi * (phase_particle - phase_ref + coupler_phase)  ! For rfcavity §4.31
angle = twopi * coupler_angle
```

The energy kick is needed to keep things symplectic.

Example:

```
rf1: lcav, l = 4.5, gradient = 1.2e6, coupler_at = both_ends,
                                              coupler_strength = 0.037
```

5.18 Field Extending Beyond Element Boundary

The field_overlaps element attribute can be used to indicate that the electric or magnetic fields of one element overlap another element. The syntax is:

```
<overlapping_ele>: ... field_overlaps = {<overlapped_ele1>, <overlapped_ele2>, ...}
```

The {} braces are optional if there is only one overlapped element.

Example:

```
b1: sbend, l = 2.3, field_overlaps = {q1, s2}, ...
inj_line: line = (..., s2, b1, mark3, q1, ...)
```

In this example, the field of element b1 extends beyond the ends of b1 and overlaps elements q1 and s2. There is no limit to the number of elements that are overlapped by any given element and overlapped elements do not have to be next to the overlapping element in the line. If there are multiple elements whose name matches the name of a overlapped element, the element closest to the overlapping element is chosen. Thus in the above example, if there are multiple elements named q1, the closest q1 to b1 is designated as the overlapped element.

There can be multipole field_overlaps = ... constructs for an overlapping element. Thus the following is equivalent to the above example:

```
b1: sbend, l = 2.3, field_overlaps = q1, field_overlaps = s2
```

Note: When the field overlaps elements that are superimposed (§8), the overlapped elements must be the `super_lord` elements and never the slaved elements.

The field, when `field_calc` (§6.4) is set to `bmad_standard`, never extends beyond the element boundary and so a `bmad_standard` field will never overlap another element.

5.19 Automatic Phase and Amplitude Scaling of RF Fields

Elements that have accelerating fields are:

```
e_gun       ! §4.15
em_field    ! §4.17
lcavity     ! §4.31
rfcavity    ! §4.46
```

[Notice that `rfcavity` elements by definition, have a constant reference energy while with all the other elements the entrance end reference energy will, in general, be different from the exit end reference energy.]

The problem that arises with accelerating fields is how to set the overall amplitude (and phase if the fields are oscillating) of the field so that a particle, starting on the reference orbit and starting with the reference energy, has the desired energy gain at the exit end of the element where the "desired" is set by the `voltage` or `gradient` attribute of the element plus a `phi0` phase attribute for AC fields.

The scaling problem is not present when `bmad_standard` tracking (§6.1) is used since `bmad_standard` tracking uses an integrated formula that is designed to give the proper acceleration. Rather it is a problem for Runge-Kutta and other integration methods. In consequence of this, the `field_autoscale` and `phi0_autoscale` parameters are ignored with `bmad_standard` tracking.

The problem becomes even more complicated at non-ultra relativistic energies where the particle velocity is not a constant. In this case, the proper amplitude and/or phase settings will depend upon what the incoming energy of the reference particle is.

To help with the scaling problem, *Bmad* has the capability to automatically scale an accelerating field's amplitude and/or phase. The two lattice element parameters that turn on/off auto scaling are (§10.1):

```
autoscale_phase      = <Logical>  ! Automatic phase scaling.
autoscale_amplitude  = <Logical>  ! Automatic amplitude scaling.
```

The default value is True for both parameters. Example:

```
rf2: rfcavity, autoscale_phase = F
```

Scaling takes place during program execution when a lattice is initially created (that is, when the lattice file is parsed) and when parameters in the lattice that would change the scaling are varied. The element parameters varied when autoscaling is done are:

```
field_autoscale       ! Amplitude scale
phi0_autoscale        ! phase scale
```

For an `rfcavity` element, the `field_autoscale` parameter is set such that when the phase is adjusted for maximum acceleration, the voltage gain of a particle on the zero orbit is equal to the value of the element's `voltage` parameter.

For `rfcavity` elements, `phi0_autoscale` is set so that with `phi0` equal to zero, a particle on the zero orbit will not see any energy gain through the cavity. There are two values for `phi0_autoscale` where, with `phi0` equal zero, the energy gain of the zero orbit particle is zero. If the `bmad_com` global parameter `rf_phase_below_transition_ref` (§11.2) is set to `False` (the default), `phi0_autoscale` will be set

such that, with `phi0` equal zero, the zero orbit for a particle above transition will be at the stable zero-crossing. For a particle below transition, setting `rf_phase_below_transition_ref` to `True`, will result in `phi0_autoscale` being set such that, for particles below transition, and with `phi0` equal zero, a particle on the zero orbit will be at the stable zero-crossing.

For `lcavity` elements, `phi0_autoscale` is set such that a particle entering the cavity with zero phase space coordinates has an energy change that is equal to change in reference energy

```
dE = voltage * cos(2 * pi * phi0)
```

With this, the ending p_z phase space coordinate is zero.

It is important to keep in mind that with a ring, the closed orbit will be the equilibrium orbit. For example, for a ring with a single cavity, changes in `phi0` will just result in variations in the closed orbit z. The voltage kick that the closed orbit particle gets going through the cavity is independent of `phi0` and will be equal to the radiation losses throughout the ring (and if radiation of off, the kick will be zero).

It is not possible to autoscale if the voltage is very small. *Bmad* sets a lower limit of 10 Volts and if the voltage is set less than this no scaling is done. If no autoscaling is done, the default setting of `field_autoscale` is 1 and the default setting of `phi0_autoscale` is 0.

Note: If a field map (§5.16) is used to define the cavity fields, in order for the field amplitude to vary with changes in the value of the `voltage`, the `master_parameter` of the field map must be set to `voltage` (or `gradient` since `gradient` and `voltage` are linked).

`field_autoscale` and `phi0_autoscale` are not needed and therefore ignored when `bmad_standard` tracking is done.

5.20 Wakefields

Wakefield modeling is discussed in Chapter §20. The syntax for specifying short-range wakefields for an element is given in Sec. §5.20.1. The syntax for specifying long-range wakefields is given in Sec. §5.20.4.

Bmad has two modes for tracking particles. One mode tracks individual particles one at a time. The other mode tracks bunches of particles. Which mode is used for a given program is decided by the program. wakefields are ignored when tracking individual particles and only used when tracking bunches.

5.20.1 Short-Range Wakes

The short-range wakes for a lattice element are specified via a set of "pseudo" modes. Equations for short-range wakefields are given in Sec. §20.1. The `sr_wake` attribute is used to set wakefield parameters. The general form of this attribute is:

```
sr_wake = {
    z_max = <real>, z_scale = <real>, amp_scale = <real>,
    scale_with_length = <logical>,
    longitudinal = {<amp>, <damp>, <k>, <phi>, <position_dependence>},
    ...
    longitudinal = {...},
    transverse = {<amp>, <damp>, <k>, <phi>, <polarization>, <particle_dependence>},
    ...
    transverse = {...}
    z_long = {...}
```

```
}
```

The sr_wake structure has optional components z_max, z_scale, amp_scale, and scale_with_length. To specify wakes, zero or more longitudinal sub-structures each one specifying a single longitudinal mode, zero or more transverse sub-structures each one specifying a single transverse mode and zero or one z_long sub-structures that can be used to specify the longitudinal wake as a function of z-position (§5.20.2). Example:

```
cav9: lcavity, ..., sr_wake = {z_max = 1.3e-3, scale_with_length = F,
      longitudinal = {3.23e14, 1.23e3, 3.62e3, 0.123, none},
      longitudinal = {6.95e13, 5.02e2, 1.90e3, -1.503, x_leading},
      transverse = {4.23e14, 2.23e3, 5.62e3, 0.789, none,  trailing},
      transverse = {8.40e13, 5.94e2, 1.92e3, 1.455, x_axis, none} }
```

Note: After an element has been defined, to refer to a given component use the notation:

```
<element-name>[sr_wake%<component-name>]
```

where <element-name> is the name of the element and <component-name> is the name of the component. For example, after the cav9 element has been defined, the z_scale component can be changed via:

```
cav9[sr_wake%z_scale] = 0.4 * cav9[sr_wake%z_max]
```

The first four components of both the longitudinal and transverse sub-structures give A_i, d_i, k_i, and $\phi_i/2\pi$ of Eq. (20.6). The units for these components are:

	Monopole	*Dipole*
A	V/Coul/m	V/Coul/m^2
d	1/m	1/m
k	1/m	1/m
$\phi/2\pi$	Radians/2π	Radians/2π

Monopole modes are modes that are independent of transverse position and dipole modes are modes that are linear in the transverse position.

For the longitudinal sub-structures, there is a 5^{th} component, called position_dependence, gives the transverse position dependence of the wake. Possible values are:

```
none          ! No position dependence
x_leading     ! Linear in the leading particle x-position
y_leading     ! Linear in the leading particle y-position
x_trailing    ! Linear in the trailing particle x-position
y_trailing    ! Linear in the trailing particle y-position
```

"x_leading", for example, means that the wake left by a "leading" particle is linear in the x-position of the particle while the kick to a "trailing" particle is independent of the the trailing particle's transverse position. "y_trailing", on the other hand, means that the wake left by a leading particle is independent of the leading particle's transverse position but that the kick felt by a trailing particle is proportional to the trailing particle's y-position.

For the transverse sub-structures, there is a 5^{th} component giving the polarization and a 6^{th} component specifying if the kick is dependent upon the leading or trailing particle transverse position. Possible values for the polarization are

```
none          ! Kick in both x and y-planes
x_axis        ! Kick in only the x-plane
y_axis        ! Kick in only the y-plane
```

and possible values for the particle dependence are:

```
none          ! No transverse dependence
leading       ! Depends linearly on the leading particle position.
trailing      ! Depends linearly on the trailing particle position.
```

The z_max component of sr_wake specifies the maximum z value at which the pseudo mode fit is valid. The z_max component is optional and if present and positive, *Bmad* will check that the distance between particles does not exceed z_max. If it does, *Bmad* will report an error. If z_max is not positive it is ignored.

If scale_with_length is False (the default is True), the length factors in Equations like Eqs. (20.1) and (20.4) are dropped. This is convenient for using zero length elements with wake simulations.

The amp_scale component is used to scale the amplitude of the modes. This corresponds to A_{amp} in Eq. (20.6). The default value is 1.0.

The z_scale is used to scale the z distance in the wake equations:

```
z(used in equations) = z_scale * z(actual)
```

The default value is 1.0.

Note: In a beam chamber with circular symmetry, the linear terms in the longitudinal wake are zero and the transverse wake has no terms independent of the transverse offsets nor terms that depend upon the trailing particle offset.

5.20.2 Short-Range Longitudinal Wake With a Wake Table

The z_long substructure of sr_wake (§5.20.1) is an alternative to the longitudinal substructure. With the z_long substructure, the wake is specified using a table of equally spaced points in z-position or as a function of time. The syntax for z_long is

```
z_long = {
  time_based = <T/F>,
  smoothing_sigma = <sigma>,
  position_dependence = <dep>,
  w = {
    <z1>  <w1>,            ! Wake <w> table as a function of z-position (or time).
    <z2>  <w2>,
    ...
  }
}
```

The w component gives the single particle wake as a function of either z-position or time depending upon the setting of the optional time_based logical. The default is False. The z (or time) values in the w table must be equally spaced and in increasing order. Unlike the longitudinal and transverse pseudo-mode description, the wake can be finite for both positive and negative z (or time). When the wake is parsed, if time_based is True, the table is converted to be z-based using the conversion $z = -c\,t$ where c is the speed of light. If the table is not symmetric around $z = 0$ (if there are more points on one size of zero than the other), the table is extended to be symmetric (this is needed for tracking). During tracking, the half-width of the wake (the magnitude of the maximum or minimum z of the symmeterized table) must be greater than the full width of the bunch (excluding 1% outlier particles). If not, an error message is generated and tracking is stopped.

The optional position_dependence component determines if the wake depends upon the transverse coordinates of the bunch particles. See §5.20.1 for details. The default is none.

The wake is applied to a bunch by convoluting the bunch distribution with the wake. The optional smoothing_sigma component, if set non-zero, applies a Gaussian smoothing filter to the convolution. The value of smoothing_sigma is in meters if time_based is False or seconds otherwise.

The z_scale and amp_scale parameters of the sr_wake structure (§5.20.1) are used with z_long.

5.20.3 Short-Range Wakes — Old Format

Note: Short range wakes use distance (z) as the dependent parameter instead of time which is what the long range wakes use.

There is an old, deprecated, format for specifying short-range wakes where the wake data is contained in a separate file whose name is given by the `lr_wake_file` attribute. Example:

```
abc: lcavity, sr_wake_file = "sr.wake", lr_freq_spread = 0.0023, lr_self_wake_on = F
```

Example file:

```
! Pseudo Wake modes:
!                        Amp       Damp          K      Phase  Polar-    Transverse_
! Longitudinal:       [V/C/m]     [1/m]      [1/m]     [rad]  ization   Dependence
! Transverse:       [V/C/m^2]     [1/m]      [1/m]     [rad]

&short_range_modes
  longitudinal(1) = 3.23e14     1.23e3     3.62e3     0.123
  longitudinal(2) = 6.95e13     5.02e2     1.90e3    -1.503
  .. etc ..
  transverse(1) =   4.23e14     2.23e3     5.62e3     0.789   none   linear_trailing
  transverse(2) =   8.40e13     5.94e2     1.92e3     1.455
   .. etc ..
  z_max = 1.3e-3
/
```

Notice that the old format uses radians and not radians/2π for the phase.

Possible settings for the polarization parameter are: parameter are:

```
none    ! Default
x_axis
y_axis
```

The `polarization` name may be abbreviated.

The `transverse_dependence` parameter sets whether the wake kick is linear in the offset of the leading or trailing particle or is independent of the transverse offset. Possible settings of this parameter are:

```
none              ! Default for longitudinal modes
linear_leading    ! Default for transverse modes
linear_trailing
```

The `transverse_dependence` parameter may be abbreviated. Note: Due to the way the wake file is parsed, if `transverse_dependence` is specified for a particular mode, `polarization` must also be specified.

For `longitudinal` modes: If the `transverse_dependence` is `none` (the default), the `polarization` must also be `none` (other combinations do not make sense). If the `transverse_dependence` is *not* `none` for a `longitudinal` mode, the `polarization` must be set to `x_axis` or `y_axis`.

5.20.4 Long-Range Wakes

Note: Long range wakes use time as the dependent parameter instead of distance (z) which is what the short range wakes use.

The `lr_wake` attribute is used to set the long-range wakefield parameters for a lattice element. Equations for long-range wakes is given in Sec. §20.2. The general form of this attribute is:

```
lr_wake = {time_scale = <real>, amp_scale = <real>, freq_spread = <real>,
    self_wake_on = <logical>, t_ref = <real>,
    mode = {<freq_in>, <R_over_Q>, <damp>, <phi>, <m_order>, <polar_angle>,
                                          <b_sin>, <b_cos>, <a_sin>, a_cos>},
    ...
    mode = {...} }
```

The `lr_wake` structure has optional components `time_scale`, `amp_scale`, `freq_spread`, `self_wake_on`, and `t_ref` along with one or more `mode` sub-structures each specifying a single long-range mode. Example:

```
f1 = 1.65e9;  q1 = 7e4
cav9: lcavity, ..., lr_wake = {time_scale = 1.7, freq_spread = 0.001,
     mode = {f1, 0.76, f1/(2*q1), 0, 1, unpolarized},
     mode = {-1, 0.57, 3e4, 0, 2, 0.15} }
```

The first six components of the `mode` sub-structure correspond to $\omega/2\pi$, R/Q, *damp*, $\phi/2\pi$, m, and $\theta_p/2\pi$ in §20.2. The units of R/Q are Ω/meter^{2m} and the units of *damp* are $1/\text{sec}$. The last four components, which are optional, correspond to $b_{\sin}$, $b_{\cos}$, $a_{\sin}$, and $a_{\cos}$. These four components can be used as a convenient way to save the state of the long-range wakes. These components will typically not be present in a lattice file except for lattice files generated by a program that does wake simulations.

A negative frequency is used to designate wakes that are part of the fundamental accelerating mode. That is, the frequency of such a mode is set to the value of `rf_frequency` for the lattice element the wake is associated with. It is an error to have a negative frequency for a mode for elements that do not have a `rf_frequency` attribute. *Bmad* needs to know if a wake is part of the fundamental mode due to timing issues as discussed in §26.1.

The `freq_spread` component of the `lr_wake` structure is used to randomly vary the long-range mode frequencies. This can be used to spread out the long-range mode frequencies among different cavities. The default value is zero which means there is no varying the mode frequencies. The fractional difference between of the mode frequencies used in a simulation and the input mode frequencies will have a Gaussian distribution with an RMS given by `freq_spread`. For example, a value of 0.01 for `freq_spread` gives a 1% variation in frequency. Note: Wake modes that are locked to the fundamental accelerating mode (§5.20.5), are not shifted.

The `self_wake_on` component can be used to turn off the long-range self-wake which is the longitudinal kick given a particle due to the wake generated by that same particle. [The transverse self wake is always zero.] The default setting of `self_wake_on` is `True`. Turning off the self-wake, for example, can be done to avoid double counting if both long-range and short-range wakes are defined. Also, The standard formulas for the long–range resistive wall wake are not valid over short time scales. In this case, the self-wake should be turned off.

The `amp_scale` component is used to scale the amplitude of the modes. This corresponds to A_{amp} in Eq. (20.6). The default value is 1.0.

The `time_scale` component is used to scale the time distance in the wake equations:

```
t (used in equations) = time_scale * t(actual)
```

The default value is 1.0.

The `t_ref` component is the reference time used in computing the wake (Eq. (20.20)). Like the mode components `b_sin`, etc., This component is typically not set when a lattice file is constructed.

After the long–range wake has been defined, components can be referenced or redefined using the notation

```
lr_wake%mode(n)%freq_in      ! Input Frequency
```

```
lr_wake%mode(n)%freq          ! Actual Frequency (set by Bmad)
lr_wake%mode(n)%r_over_q      ! R/Q
lr_wake%mode(n)%damp          ! damp
lr_wake%mode(n)%phi           ! Actually phi/2pi
lr_wake%mode(n)%polar_angle   ! Polarization Angle
lr_wake%mode(n)%polarized     ! Logical
lr_wake%freq_spread           ! Frequency spread
lr_wake%amp_scale             ! Amplitude scale
lr_wake%time_scale            ! Time scale
```

`freq_in` is the input frequency set in the `lr_wake` structure and `freq` is the actual frequency set by *Bmad*. `Freq` will only be different from `freq_in` if `freq_spread` is nonzero. Example:

```
cav9[lr_wake%mode(2)%freq_in] = 1.1 * cav9[lr_wake%mode(2)%freq_in] ! Raise frequency by 10%
```

5.20.5 Long-Range Wakes – Old Format

There is an old, deprecated, format for specifying long-range wakes where the wake data is contained in a separate file whose name is given by the `lr_wake_file` element attribute. Example:

```
abc: lcavity, lr_wake_file = "lr.wake", lr_freq_spread = 0.0023, lr_self_wake_on = F
```

The file gives the wake modes by specifying the frequency ($\omega/2\pi$), R/Q, Q, and m (order number), and, optionally, the polarization angle $\theta_p/2\pi$ for each cavity mode. The input uses Fortran90 namelist syntax: The data begins with the string `&long_range_modes` and ends with a slash `/`. Everything outside this is ignored. Each mode is labeled `lr(i)` where `i` is the mode index. An example input file is:

```
              Freq      R/Q      Q    m   Polar   b_sin  b_cos a_sin  a_cos  t_ref
                       [Ohm/              Angle
              [Hz]     m^(2m)]          [Rad/2pi]
&long_range_modes
  lr(1) = 1.650e9    0.76    7.0e4  1    "unpol"
  lr(2) = 1.699e9   11.21    5.0e4  1    "0.15"
  lr(3) =    -1      0.57    1.1e6  0    "unpol"
/
```

[Note: The quotation marks are needed with some compilers and not with others.] If the polarization angle is set to "`unpolarized`" the mode is taken to be unpolarized.

Two deprecated element attributes that affect the long-range wake are `lr_freq_spread` and `lr_self_wake_on` which correspond to `lr_wake%freq_spread` and `lr_wake%self_wake_on` (5.20.4

Notice that with the old format it is not possible to specify the phase offset ϕ.

5.21 Fringe Fields

Some lattice elements can have fringe fields at the element edges. Whether *Bmad* tries to model the fringe fields using the models described below first depends upon what kind of tracking is done. Fringe effects are *not* applied when an element's `tracking_method` is set to:

```
custom
mad
```

Additionally, no fringe effects will be used if the `tracking_method` is `runge_kutta` or `time_runge_kutta`, and the element's `field_calc` (§6.4) is *not* `bmad_standard`. This is done since, for non-`bmad_standard` field calculations, it is assumed that the field profile includes the fringe regions.

5.21.1 Turning On/Off Fringe Effects

For elements that have a fringe, whether fringe fields are ignored or not is determined by the setting of the `fringe_at` element parameter. The possible settings are

```
no_end
both_ends             ! Default
entrance_end
exit_end
```

This is particularly useful in vetoing the fringe effect in the interior of split elements. The setting of attributes like `fringe_type` (see below) are ignored for boundaries where the fringe has been turned off.

When a particle's spin is being tracked, the `spin_fringe_on` logical attribute of the element determines how the spin tracking is handled through a fringe region. Example:

```
q: quad, spin_fringe_on = T, fringe_at = exit_end
```

Here, there is no fringe effect at the entrance of the element and the fringe at the exit end of the element will affect the spin. The default setting of `spin_fringe_on` is `True`. The `spin_fringe_on` attribute is useful for examining how much fringe fields affect spin precession.

5.21.2 Fringe Types

Bmad and PTC have several fringe field models for static magnetic fields. Which fringe model is used is set by two element attributes: `fringe_type` and `ptc_fringe_geometry` with `ptc_fringe_geometry` only being used with bends and PTC dependent tracking. The `fringe_type` switch is used to select how a fringe field is simulated. The possible settings of `fringe_type` are:

```
none              ! No fringe effect.
soft_edge_only
hard_edge_only
full
linear_edge       ! Sbend, rbend only.
basic_bend        ! Sbend, rbend only.
sad_full          ! Sbend, rbend, sad_mult only.
```

Default settings for `fringe_type` are

Element Type	Default fringe_type
e_gun	full
lcavity, rfcavity	full
sbend, rbend	basic_bend
All others	none

Some fringe fields can be divided into two pieces. The first piece is called the `hard edge` fringe kick and is the kick in the limit that the longitudinal extent of the fringe is zero. The second piece is the `soft edge` fringe kick which is the fringe kick with the fringe having a finite longitudinal extent minus the hard edge fringe kick. That is

```
fringe kick = hard fringe kick + soft fringe kick
```

The advantage of separating the fringe kick in this way is that the hard fringe can be used without having to know anything about the longitudinal extent of the fringe (which happens when simulating magnets that have not yet been fully designed). In many cases, this is a good enough approximation. Note that using the soft fringe without the hard fringe is not physical but can be useful in understanding how the soft edge component affects tracking. See §19 for details.

For bend elements the following fringe maps are used:

fringe_type	Fringe model
none	No fringe effect
soft_edge_only	SAD dipole soft edge (§19.2)
hard_edge_only	Bend Second Order (§19.1) with fint and hgap ignored.
full	Exact bend
linear_edge	Linear dipole hard edge (§19.4)
basic_bend	Bend Second Order (§19.1)
sad_full	

The basic_bend setting for bend elements, which is the default, is essentially the basic vertical focusing effect that is present when there is a finite e1 or e2 face angle. With bmad_standard tracking, basic_bend also includes second order terms (§19.1). The linear_edge setting ignores these second order terms. In some cases, for instance in a chicane, basic_bend is not good enough. With fringe_type set to full, higher order effects are taken into account.

PTC does not have a linear_edge fringe model. With PTC tracking, basic_bend tracking is used if linear_edge is chosen.

Additionally, for use with PTC, the ptc_fringe_geometry switch can be used to define the symmetry of the fringe fields. Possible settings are:

```
x_invariant
multipole_symmetry
```

The difference between x_invariant and multipole_symmetry is that with multipole_symmetry the fringe field for an n^{th} order multipole is assumed to have the same rotational symmetry as the multipole. With this assumption, the fringe field has $n + 1^{st}$ order terms. With x_invariant, the fringe field is calculated assuming that there is translational invariance along the horizontal x axis. This differs from multipole_symmetry by adding terms of increasing order consistent with the translational invariance. See Étienne Forest's book[Forest98] for more details. Which setting of ptc_fringe_geometry is appropriate depends upon how the dipole under consideration is constructed. The two settings of ptc_fringe_geometry represent two points of a continuum of possible fringe field geometries.

When using PTC tracking (§1.4), the ptc_com[max_fringe_order] (§10.1) determines the maximum order of the calculated fringe fields.

Example:

```
b1: rbend, angle = pi/4, g = 0.3, fringe_type = full
```

The soft_edge_only, hard_edge_only and sad_full settings of fringe_type emulate the fringe field tracking used in the SAD program[SAD]. The soft_edge_only setting only uses the linear part of the fringe, hard_edge_only ignores the linear part of the fringe, and sad_full uses the full fringe. For an sbend or rbend element, these SAD fringe fields are in addition to the fringe fields that occurs with a finite e1 or e2 face angle.

Equivalent settings of SAD fringe and disfrin for sbend and rbend elements:

fringe_type	fringe	disfrin
soft_edge_only	1	1
hard_edge_only	0	0
basic_bend	0	1
sad_full	1	0

For quadrupole and sad_mult elements, the translation between the fringe_at and fringe_type settings and the fringe and disfrin switches of SAD is:

	fringe_at:			
fringe_type	no_end	entrance_end	exit_end	both_ends*
none	[0, $\neq 0$]	[0, $\neq 0$]	[0, $\neq 0$]	[0, $\neq 0$]
soft_edge_only	[0, $\neq 0$]	[1, $\neq 0$]	[2, $\neq 0$]	[3, $\neq 0$]
hard_edge_only*	[0, $\neq 0$]	No SAD Equiv	No SAD Equiv	[0, $= 0$]
full	[0, $\neq 0$]	[1, $= 0$]	[2, $= 0$]	[3, $= 0$]

*Default value. [fringe, disfrin]

Each entry is the table is of the form [fringe, disfrin]. The soft_edge_only fringe kick is a kick that is linear in the transverse (x, p_x, y, p_y) coordinates and comes from the finite width of the quadrupolar fringe field. The width of the quadrupolar fringe field is characterized by the f1 and f2 attributes. The sad_nonlinear_only fringe kick comes from the nonlinear part of the quadrupolar field plus the fringes of the other multipoles.

For sad_mult elements, the soft fringe quadrupole (§19.6) parameters fq1 and fq2 (§19.6) are related to the corresponding SAD parameters f1 and f2 via

```
f1 = -sign(fq1) * sqrt(24 * |fq1|)
f2 = fq2
```

Note: With SAD bend elements, the f1 parameter is associated with the dipole fringe (§19.2) and not the quadrupolar fringe! Also note that in the SAD documentation, the soft edge is called the "linear" fringe.

For programmers who deal with PTC directly: The translation between ptc_fringe_geometry on the *Bmad* side and bendfringe on the PTC side is:

ptc_fringe_geometry	bendfringe
x_invariant	True
multipole_symmetry	False

5.22 Instrumental Measurement Attributes

instrument, monitor, detector, and marker elements have special attributes to describe errors associated with orbit, betatron phase, dispersion and coupling measurements. These attributes are:

Attribute	*Symbol* (See: §28.1)	
`tilt`	θ_t	See §5.6
`x_offset`	x_{off}	See §5.6
`y_offset`	y_{off}	See §5.6
`x_gain_err`	$dg_{x,\text{err}}$	Horizontal gain error
`y_gain_err`	$dg_{y,\text{err}}$	Vertical gain error
`crunch`	ψ_{err}	Crunch angle
`tilt_calib`	θ_{err}	tilt angle calibration
`x_offset_calib`	x_{calib}	Horizontal offset calibration
`y_offset_calib`	y_{calib}	Vertical offset calibration
`x_gain_calib`	$dg_{x,\text{calib}}$	Horizontal gain calibration
`y_gain_calib`	$dg_{y,\text{calib}}$	Vertical gain calibration
`crunch_calib`	ψ_{calib}	Crunch angle calibration
`noise`	n_f	Noise factor
`de_eta_meas`	dE/E	Percent change in energy
`x_dispersion_err`	$\eta_{x,\text{err}}$	Horizontal dispersion error
`y_dispersion_err`	$\eta_{y,\text{err}}$	Vertical dispersion error
`x_dispersion_calib`	$\eta_{x,\text{calib}}$	Horizontal dispersion calibration
`y_dispersion_calib`	$\eta_{y,\text{calib}}$	Vertical dispersion calibration
`n_sample`	N_s	Number of sampling points
`osc_amplitude`	A_{osc}	Oscillation amplitude

A program can use these quantities to calculate "measured" values from the "laboratory" values. Here, "laboratory" means as calculated from some model lattice. See §28.1 for the conversion formulas.

Chapter 6

Tracking, Spin, and Transfer Matrix Calculation Methods

Bmad allows for a number of methods that can be use to "track" a particle through a lattice element. Here "track" can mean one of three things:

```
1) Calculate a particle's phase space coordinates at the exit
   end of the element given the coordinates at the entrance end.
2) Calculate the linear transfer map (Jacobian) through an element
   about a given reference orbit.
3) Calculate the a particle's spin orientation at the exit end
   of the element given the coordinates at the beginning.
```

The different tracking methods that are available have different advantages and disadvantages in terms of speed, symplecticity, etc. What tracking method is used, is selected on an element–by–element basis using the attributes:

```
tracking_method      = <Switch>   ! phase space tracking method.
mat6_calc_method     = <Switch>   ! 6x6 transfer matrix calculation.
spin_tracking_method = <Switch>   ! Spin tracking method.
```

Example:

```
q2: quadrupole, tracking_method = symp_lie_ptc
q2[tracking_method] = symp_lie_ptc
quadrupole::*[tracking_method] = symp_lie_ptc
```

The first two lines of this example have exactly the same effect in terms of setting the `tracking_method`. The third line shows how to set the `tracking_method` for an entire class of elements.

These switches are discussed in more detail in the following sections.

6.1 Particle Tracking Methods

The `tracking_method` attribute of an element sets the algorithm that is used for single particle tracking through that element. Table 6.1 gives which methods are available for each type of element. Note: Table 6.1 pertains to charged-particle tracking only. When tracking photons, only `bmad_standard` and `custom` tracking method are available.

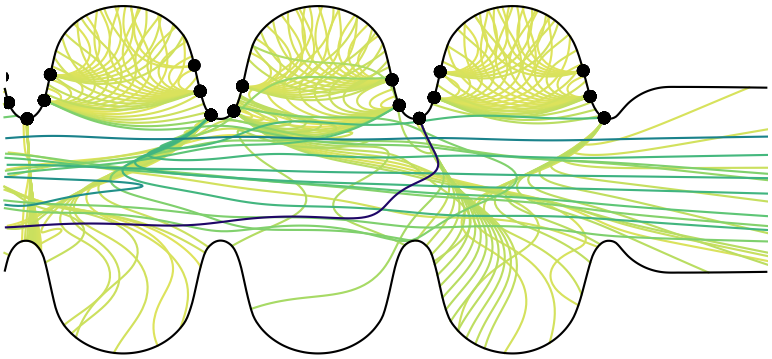

Figure 6.1: Dark current tracking. Example of where a time based tracker (`time_runge_kutta`) is useful for simulating particles that can reverse their longitudinal velocity. Here the tracks drawn are from a simulation of "dark current" electrons generated at the walls of an RF cavity due to the large electromagnetic fields.

A note on terminology: Adaptive step size control used with the `Runge_Kutta` integrator means that instead of taking fixed step sizes the integrator chooses the proper step size so that the error in the tracking is below the maximum allowable error set by `rel_tol_adaptive_tracking` and `abs_tol_adaptive_tracking` tolerances. The advantage of step size control is that the integrator uses a smaller step size when needed (the fields are rapidly varying), but makes larger steps when it can. The disadvantage is that a step is more computationally intensive since the error in a step is estimated by repeating a step using two mini steps. Except for testing purposes, it is recommended that adaptive stepping be used over fixed step tracking since experience has shown that adaptive stepping is almost always faster. It is also recommended that `runge_kutta` be used over `time_runge_kutta` since `runge_kutta` does not have the overhead of switching between time-coordinates and z-coordinates. The exceptions are cases where `time_runge_kutta` must be used like with an `e_gun` where the particles start with zero momentum and in cases where particles may reverse their longitudinal direction (EG: dark current electrons).

`Bmad_Standard` Uses formulas for tracking. The formulas generally use the paraxial approximation. The emphasis here is on speed. It is important to note that field maps (§5.16) are *ignored* by `bmad_standard` tracking. The tracking is non-symplectic but the non-symplectic errors tend to be small so that `bmad_standard` can be used in the vast majority of cases (§6.6).

`Custom` This method can be used to call custom tracking code supplied by the User. See `s:custom.ele` for more details.

`fixed_step_runge_kutta` The `fixed_step_runge_kutta` method is similar to `runge_kutta` tracking except that `fixed_step_runge_kutta` does not use adaptive step size control but instead takes steps of fixed size using the setting of `ds_step` or `num_steps` for the element being tracked through (§6.4). Generally, using adaptive step control will be much more efficient so it is recommended that `fixed_step_runge_kutta` *not* be used unless there is a compelling reason not to. This method is non-symplectic (§6.6).

`fixed_step_time_runge_kutta` The `fixed_step_time_runge_kutta` method is similar to `time_runge_kutta` tracking except that `fixed_step_time_runge_kutta` does not use adaptive step size control but instead takes steps of fixed size using the setting of `ds_step` or `num_steps` for the element being tracked through (§6.4). Generally, using adaptive step control will be much more efficient so it is recommended that `fixed_step_time_runge_kutta` *not* be used unless there is a compelling reason not to. This method is non-symplectic (§6.6).

`Linear` The `linear` method just tracks particles using the 0th order vector with the 1st order 6x6 transfer matrix of an element. Depending upon how the transfer matrix was generated this may or may not be symplectic. Since there would be a circular dependency to have the orbital tracking dependent upon the transfer matrix and the transfer matrix dependent upon the determination of the reference orbit, the calculation of the transfer matrix when the `tracking_method` is set to `linear` will always use the zero orbit as the reference orbit.

Additionally, a `linear` tracking method may not be used with `mat6_calc_method` set to `tracking` since this would also give a circular dependency. Note: setting the `tracking_method` to `linear` does not affect PTC calculations (§1.4). In particular, Taylor maps will not be affected.

`MAD` This uses the MAD 2nd order transfer map. This method is not able to handle element misalignments or kicks, and becomes inaccurate as the particle energy deviates from the reference energy. MAD tracking should only be used for testing purposes. Note: Thanks to CERN and Frank Schmidt for permission to use the MAD tracking code within *Bmad*.

`runge_kutta` This uses a 4^{th} order Runge Kutta integration algorithm with adaptive step size control. This is essentially the Cash-Karp formulation. This method will be slow compared to non-Runge-Kutta methods so only use this if it is not possible to use something like `bmad_standard`. This method is accurate but non-symplectic (§6.6). `Warning:` When using `custom` fields, if the fields do not obey Maxwell's equation, there is the possibility of the `runge_kutta` tracking halting mid-way through an element. See section §6.4 for more details.

`Symp_Lie_Bmad` Symplectic tracking using a Hamiltonian with Lie operation techniques. This is similar to `Symp_Lie_PTC` (see below) except this uses a *Bmad* routine. By bypassing some of the generality inherent in PTC (§1.4), `Symp_Lie_Bmad` achieves about a factor of 10 improvement in speed over `Symp_Lie_PTC`.

`Symp_Lie_PTC` Symplectic tracking using a Hamiltonian with Lie operator techniques. This uses Étienne Forest's PTC (§1.4) software for the calculation. This method is symplectic but can be slow. Exceptions: The tracking is not symplectic when tracking through and element with an associated electric field and when tracking through a `taylor` element.

`Taylor` The tracking uses a Taylor map. The map is either explicitly given in the lattice file, that is, the element must be of type `taylor` (§4.52), or the Taylor map is generated from the PTC (§1.4) package. Generating the map may take time but once you have it it should be very fast. One possible problem with using a Taylor map is that you have to worry about the accuracy if you do tracking at points that are far from the expansion point about which the map was made. This method is non-symplectic away from the expansion point. Whether the Taylor map is generated taking into account the offset an element has is governed by the `taylor_map_includes_offsets` attribute (§6.8).

The order of a Taylor map is set by the `parameter[taylor_order]` parameter (§10.1).

`Time_Runge_Kutta` This method uses time as the independent variable instead of the longitudinal z position. The advantage of this method is that it can handle particles which reverse direction longitudinally. One use for this method is "dark current" tracking where, as illustrated in Fig. 6.1, low energy particles generated at the vacuum chamber walls can be found traveling in all directions. Notice that `time_runge_kutta` is different from using `absolute time tracking` as explained in §26.1. This method is non-symplectic (§6.6).

Element Class	Bmad_Standard	Custom	Linear	MAD	Runge_Kutta[a]	Symp_Lie_Bmad	Symp_Lie_PTC	Taylor	Time_Runge_Kutta[a]
ab_multipole and multipole	D	X	X				X	X	
ac_kicker	D	X	X		X				X
beambeam	D	X	X						
bends: rbend and sbend	D	X	X	X	X		X	X	X
converter	D	X							
crab_cavity	D	X	X						
custom		D	X		X				X
drift	D	X	X	X	X		X	X	X
e_gun		X			X[b]				D
ecollimator and rcollimator	D	X	X		X		X	X	X
elseparator	D	X	X	X	X		X	X	X
em_field		X			D		X	X	X
fiducial	D	X					X		
fixer	D	X							
floor_shift	D	X					X		
fork	D	X	X						
gkicker	D	X	X				X	X	
hkicker	D	X	X		X		X	X	X
instrument, monitor, and pipe	D	X	X		X		X	X	X
kicker	D	X	X		X		X	X	X
lcavity and rfcavity	D	X	X		X		X	X	X
marker	D	X	X				X	X	
match	D	X						X	
octupole	D	X	X		X		X	X	X
patch	D	X			X[c]		X	X	
photonic elements	D	X							
quadrupole	D	X	X	X	X	X	X	X	X
rf_bend		X	X		X				X
sad_mult	D	X	X				X	X	
sextupole	D	X	X	X	X		X	X	X
solenoid	D	X	X	X	X	X	X	X	X
sol_quad	D	X	X		X	X	X	X	X
taylor		X	X				X	D	
vkicker	D	X	X		X		X	X	X
wiggler (map type)		X	X		X	X	X	X	X
wiggler (periodic type)	D	X	X		X[d]	X[d]	X[d]	D[d]	

[a] Includes fixed step versions.
[b] Only if the beginning energy is non-zero.
[c] Only available for non-reflection patch elements.
[d] See §4.54.1 for more details.

Table 6.1: Table of valid tracking_method switches. "D" denotes the default method. "X" denotes a valid method. Photonic elements are elements in Table 4.2 that cannot be used for charged particle tracking (Table 4.1).

6.2 Linear Transfer Map (Mat6) Calculation Methods

The `mat6_calc_method` attribute sets how the 6x6 Jacobian transfer matrix for a given element is computed. Table 6.3 gives which methods are available for each type of element. Note: Table 6.3 is for charged-particle tracking only. When tracking photons, transfer matrices (which are not very useful) are not computed.

If an element's `static_linear_map` parameter is set to `True` (the default is `False`), this prevents the linear map, which consists of the transfer matrix and the zeroth order part of the map, from being recomputed. For example, if somewhere in a lattice a steering is changed, this will shift the reference orbit and the linear transfer map in elements where the reference orbit changes will, in general, vary. However, having `static_linear_map` set to `True` will prevent this variation.

In addition to the `mat6_calc_method` switch, two element attributes that can affect the way the transfer matrix is calculated are `symplectify` and `taylor_map_includes_offsets`. These are discussed in sections §6.7 and §6.8 respectively.

For methods that do not necessarily produce a symplectic matrix the `symplectify` attribute of an element can be set to `True` to solve the problem. See §25.3.

Symplectic integration is like ordinary integration of a function f(x) but what is integrated here is a Taylor map. Truncating the map to 0^{th} order gives the particle trajectory and truncating to 1^{st} order gives the transfer matrix (Jacobian). The order at which a Taylor series is truncated at is set by `taylor_order` (see §10.1. Like ordinary integration there are various formulas that one can use to do symplectic integration.

`Auto` With `auto` the `mat6_calc_method` appropriate for the element's setting of `tracking_method` is used. The correspondence is:

Element's `tracking_method`	`Mat6_calc_method` used
bmad_standard	`bmad_standard`
`linear`	`bmad_standard`
`custom`	`custom`
`mad`	`mad`
`symp_lie_bmad`	`symp_lie_bmad`
`symp_lie_ptc`	`symp_lie_ptc`
`taylor`	`taylor`
All `Runge-Kutta` types	`tracking`

Table 6.2: Actual `mat6_calc_method` used when when the `mat6_calc_method` is set to `auto`.

`Bmad_Standard` Uses formulas for the calculation. The formulas generally use the paraxial approximation. The emphasis here is on speed.

`Custom` This method can be used to call custom matrix calculation code supplied by the User. See `s:custom.ele` for more details.

`MAD` This uses the MAD 2nd transfer map. This method is not able to handle element misalignments or kicks, and becomes inaccurate as the particle energy deviates from the reference energy. MAD tracking is generally only used for testing purposes. Thanks must be given to CERN and Frank Schmidt for permission to use the MAD tracking code within *Bmad*.

`Symp_Lie_Bmad` A symplectic calculation using a Hamiltonian with Lie operator techniques. This is similar to `Symp_Lie_PTC` (see below) except this uses a *Bmad* routine. By bypassing some of the generality inherent in PTC, `Symp_Lie_Bmad` achieves about a factor of 10 improvement in speed over `Symp_Lie_PTC`. However, `Symp_Lie_Bmad` cannot generate maps above first order.

`Symp_Lie_PTC` Symplectic integration using a Hamiltonian and Lie operators. This uses the PTC (§1.4) software for the calculation. This method is symplectic but can be slow. Exceptions: The tracking is not symplectic when tracking through and element with an associated electric field and when tracking through a `taylor` element.

`Taylor` This uses a Taylor map generated from Étienne's PTC package. Generating the map may take time but once you have it it should be very fast. One possible problem with using a Taylor map is that you have to worry about the accuracy if you do a calculation at points that are far from the expansion point about which the map was made. This method is non-symplectic away from the expansion point. Whether the Taylor map is generated taking into account the offset an element has is governed by the `taylor_map_includes_offsets` attribute (§6.8). `bmad_standard` and `taylor` tracking methods are identical. Note: Taylor maps for `match`, and `patch` elements are limited to first order.

The order of a Taylor map is set by the `parameter[taylor_order]` parameter (§10.1).

`Tracking` This uses the tracking method set by `tracking_method` to track 6 particles around the central orbit. This method is susceptible to inaccuracies caused by nonlinearities. Furthermore this method is almost surely slow. While non–symplectic, the advantage of this method is that it is directly related to any tracking results. Note: a `linear` tracking method may not be used with `mat6_calc_method` set to `tracking` since this would give a circular dependency. The two parameters that affect this calculation are `bmad_com%d_orb(6)` (§11.2) which sets the six deltas used for displacing the initial particle coordinates from the reference orbit.

	Auto	Bmad_Standard	Custom	MAD	Static	Symp_Lie_Bmad	Symp_Lie_PTC	Taylor	Tracking
ab_multipole and multipole	D	X	X		X		X	X	X
ac_kicker			X		X		X	X	D
beambeam	D	X	X		X				X
bends: rbend and sbend	D	X	X	X	X		X	X	X
converter	D	X	X						
crab_cavity			X		X				D
custom			D		X				X
drift	D	X	X	X	X		X	X	X
e_gun			X		X				D
ecollimator and rcollimator	D	X	X		X		X	X	X
elseparator	D	X	X	X	X		X	X	X
em_field			X		X		X	X	D
fiducial	D	X	X		X		X		X
floor_shift	D	X	X		X		X		X
hkicker	D	X	X		X		X	X	X
instrument, monitor, and pipe	D	X	X		X		X	X	X
kicker	D	X	X		X		X	X	X
lcavity and rfcavity	D	X	X		X		X	X	X
marker	D	X	X		X		X	X	X
match	D	X	X		X				X
octupole	D	X	X		X		X	X	X
patch	D	X	X		X		X	X	X
quadrupole	D	X	X	X	X	X	X	X	X
rf_bend	D	X	X		X				X
sad_mult	D	X	X		X		X	X	X
sextupole	D	X	X	X	X		X	X	X
solenoid	D	X	X	X	X	X	X	X	X
sol_quad	D	X	X	X	X	X	X	X	X
taylor			X		X		X	D	
vkicker	D	X	X		X		X	X	X
wiggler (map type)	D	X	X		X	X[a]	X[a]	X[a]	X
wiggler (periodic type)			X		X	X[a]	X[a]	D[a]	X

[a] See §4.54.1 for more details

Table 6.3: Table of available mat6_calc_method switches. When tracking photons, transfer matrices are not computed. "D" denotes the default method. "X" denotes an available method.

6.3 Spin Tracking Methods

The `spin_tracking_method` attribute of an elements sets the algorithm that is used for tracking a particle's spin (§24.1) through that element. Table 6.4 gives which methods are available for each type of element. Note: This table is only for charged-particle tracking since photons do not have spin.

Possible `spin_tracking_method` settings are:

`Custom`
: This method will call a routine `track1_spin_custom` which must be supplied by the programmer implementing the custom spin tracking calculation. See `s:custom.ele` for more details.

`Magnus`
: This method uses a second-order Magnus expansion of the T-BMT equation to track spins (§26.15). It is fast and in most cases closely matches numerical integration.

`Off`
: No spin tracking is done and the spin at the end of the element is the same as at the beginning.

`Sprint`
: The `sprint` algorithm (§26.24) uses first order transfer spin maps to track the spin through lattice elements. This method is very fast at the cost of accuracy for particles away from the zero orbit. The algorithm is also limited in what elements it can handle and it ignores higher order multipoles that may be present.

`Symp_Lie_PTC`
: Symplectic integration using a Hamiltonian and Lie operators. This uses Étienne's PTC software for the calculation. This method is symplectic but can be slow.

`Tracking`
: How spin is tracked here will depend also on the setting of `tracking_method`. If `tracking_method` is set to `runge_kutta` or `time_runge_kutta` the spin will be tracked along with the phase space particle coordinates using the local fields. For `tracking_method` set to `symp_lie_ptc`, the spin tracking will use `PTC`. For all other `tracking_methods`, the spin will be tracked using the "`bmad_standard`" spin tracking method which involves Romberg integration of the spin rotation matrix.

  The `runge_kutta` and `time_runge_kutta` spin tracking uses the same fourth order integrator as is used for the orbital coordinates to track the spin rotation vector.

`Transverse_Kick`
: With a setting of `transverse_kick`, it is assumed that the orbital transfer matrix is due solely to transverse magnetic fields so that the integrated spin rotation (Eqs. (24.2), (18.4), and (18.5)) can be related to the orbital transport via

$$\int \mathbf{\Omega}_{BMT} = \frac{1 + a\gamma}{1 + p_z} (-\Delta p_y, \Delta p_x, 0) \tag{6.1}$$

  where Δp_x and Δp_y are the change in p_x and p_y. Here the small angle approximation ($p_x, p_y \ll 1$) has been used.

Since speed may be an issue, *Bmad* has an global parameter called `spin_tracking_on` which is part of the `bmad_com` instance (§11.4) that determines whether spin is tracked or not.

The `spin_fringe_on` element attribute (§5.21.1) can be used to toggle whether the fringe fields of an element affect the spin.

Example:

q: quadrupole, spin_tracking_method = symp_lie_ptc

	Off	Custom	Magnus	Sprint	Symp_Lie_PTC	Tracking	Transverse_Kick
ab_multipole and multipole	X	X				D	
ac_kicker	X	X				D	
beambeam	X	X				D	
bends: rbend and sbend	X	X	X	X	X	D	
converter	X	X				D	
crab_cavity	X	X				D	
custom	X	D				X	
drift	X	X	X	X	X	D	
e_gun	X	X				D	
ecollimator and rcollimator	X	X		X	X	D	
elseparator	X	X			X	D	
em_field	X	X				D	
fiducial	X	X			X	D	
floor_shift	X	X			X	D	
hkicker	X	X	X	X	X	D	
instrument, monitor and pipe	X	X	X	X	X	D	
kicker	X	X	X	X	X	D	
lcavity and rfcavity	X	X			X	D	
marker	X	X				D	
match	D						X
octupole	X	X		X	X	D	
patch	X	X			X	D	
quadrupole	X	X	X	X	X	D	
sad_mult	X	X				D	
sextupole	X	X	X	X	X	D	
solenoid	X	X	X	X	X	D	
sol_quad	X	X			X	D	X
taylor	X					D	
vkicker	X	X	X	X	X	D	
wiggler	X	X			X	D	

Table 6.4: Table of available spin_tracking_method switches. "D" denotes the default method. "X" denotes an available method. Note: Photon tracking does not involve spin.

6.4 Integration Methods

"Integration methods" are tracking methods that involve integrating through an element's magnetic and electric fields. Integration methods are split into two classes: Those that can track Taylor maps and those that simply track a particle's position. The Taylor map methods are

```
symp_lie_bmad   ! Only to first order
symp_lie_ptc    ! Uses PTC
taylor          ! Uses PTC
```

See section §25.1 for more information on Taylor maps and symplectic integration. The latter two methods involve using the PTC library (§1.4).

The methods that do not involve Taylor maps are

```
fixed_step_runge_kutta
fixed_step_time_runge_kutta
runge_kutta
time_runge_kutta
```

there are a number of element attributes that can affect the calculation. They are

```
ds_step             = <Real>     ! Integration step length (§6.5.1)
num_steps           = <Integer>  ! Number of integration steps. (§6.5.1)
integrator_order    = <Integer>  ! Integrator order (§6.5.3)
field_calc          = <Switch>   ! How the field is calculated (§6.5.2)
```

Example:

```
q1: quadrupole, l = 0.6, tracking_method = bmad_standard,
      mat6_calc_method = symp_lie_ptc, ds_step = 0.2, field_calc = custom
```

6.5 CSR and Space Charge Methods

When doing beam tracking through an element (§21), Coherent Synchrotron Radiation (CSR) and Space Charge (SC) effects can be included by setting the appropriate method switches in that element. These switches are:

```
csr_method          = <Switch>   ! Coherent Synchrotron Radiation
space_charge_method = <Switch>   ! Space charge method
```

Note: For CSR or SC effects to be included in tracking the `bmad_com` logical `csr_and_space_charge_on` must be set to `True` (§11.2).

The possible settings for `csr_method` are

```
off             ! No CSR. Default.
1_dim           ! One dimensional calculation (§21.4.1).
```

The `1_dim` setting cannot be used when `space_charge_method` is set to `cathode_fft_3d`.

The possible settings of `space_charge_method` are

```
off             ! No SC. Default.
slice           ! SC using slices (§21.4.2).
fft_3d          ! SC using a 3D grid (§21.4.3).
cathode_fft_3d  ! Same as fft_3d with cathode image charge included (§21.4.3).
```

The `cathode_fft_3d` setting can only be used with `csr_method` set to `off`. Additionally, the `cathode_fft_3d` setting can only be used with the element `tracking_method` set to `time_runge_kutta` or `fixed_step_time_runge_kutta`.

Example:

```
q1: quadrupole, l = 0.6, csr_method = 1_dim, space_charge_method = slice, ...
```

Also see the space_charge_com structure (§11.5) which contains parameters used in space charge and CSR calculations.

Note: There is also high energy space charge calculation that can be used with single particle tracking and is discussed in §21.5.

6.5.1 ds_step and num_steps Parameters

One way to create a transfer map through an element is to divide the element up into slices and then to propagate the transfer map slice by slice. There are several ways to do this integration. The runge_kutta type methods integrate the equations of motion to give the 0^{th} order Taylor map which just represents a particle's orbit. Symplectic integration using Lie algebraic techniques, on the other hand, can generate Taylor maps to any order. The ds_step attribute determines the slice thickness. Alternatively, num_steps attribute can be used in place of ds_step to specify the number of slices. This is applicable to symp_lie_bmad and symp_lie_ptc integration. Example:

```
q: quadrupole, l = 0.6, ds_step = 0.1  ! 10 cm step size.
sbend::*[ds_step] = 0.2                ! Set the step_size for all sbend elements.
```

When tracking using maps or element-by-element with PTC there are a few points to keep in mind. First is that PTC tracks through a lattice element step by step. This is true for both map creation and symplectic integration. This means that the setting of the element parameter integrator_order (§6.5.3) or num_steps (or ds_step) for each element will affect the accuracy and speed of the computations. Bmad tries to choose reasonable default settings for the integrator order and number of steps however the calculation is not perfect. To make sure that the integrator order and number of steps is set properly, vary both and choose values (which can be different for different elements) such that the number of steps and integrator order is minimal (to minimize computation time) while at the same time is large enough so that results do not change significantly if the number of steps or is varied. Generally it is much better to use a large integrator order and a small step size rather than vice versa with the proviso that for elements with a longitudinally varying field (think wigglers or undulators), the step size must be small compared to the typical longitudinal length scale over which the field is varying (this length scale is the pole period length with with wigglers and undulators).

The default value for ds_step for a given element is calculated based upon the element's field strength. One should consider the default as more of a guesstimate.

The runge_kutta and time_runge_kutta tracking uses adaptive step control independent of the setting of the elements ds_step parameter. These methods use three bmad_com parameters §11.4) namely:

```
bmad_com[rel_tol_adaptive_tracking]
bmad_com[abs_to_adaptive_tracking]
bmad_com[max_num_runge_kutta_step]
```

The estimated error of the integration is then bounded by

```
error < abs_tol + |orbit| * rel_tol
```

lowering the error bounds makes for greater accuracy (as long as round-off doesn't hurt) but for slower tracking.

6.5.2 Field_calc Parameter

The runge_kutta type tracking methods all use as input the electric and magnetic fields of an element. How the EM fields are calculated is determined by the field_calc attribute for an element. For all lattice elements, except wigglers and undulators, possible values for field_calc are:

```
bmad_standard     ! This is the default except for custom elements
custom            ! Default for custom elements.
fieldmap
```

For `wigglers` and `undulators`, possible values for `field_calc` are:

```
planar_model
helical_model
custom
fieldmap
```

For historical reasons, the default setting for `field_calc` for `wigglers` and `undulators` is `planar_model` except if there is a field map present (§5.16) in which case the default is `fieldmap`. Note that with `bmad_standard` tracking, the setting of `field_calc` is ignored except in the case of `wigglers` and `undulators` where `field_calc` must be set to either `planar_model` or `helical_model`.

`Custom` means that the field calculations are done outside of the *Bmad* software. A program doing `custom` field calculations will need the appropriate custom routine (§38.2). Elements that set `field_calc` to `fieldmap` need to have a field map defined (§5.16).

`Warning:` When tracking a particle through a custom field using `runge_kutta`, it is important that the field obey Maxwell's equations. Fields that do not obey Maxwell's Equations may cause the `runge_kutta` adaptive step size control algorithm to take smaller and smaller steps until the step size becomes so small the tracking will stop. What happens is that the step size control algorithm takes a step and then takes two half steps over the same region and from this estimates the error in the calculation. If the error is larger than the allowed tolerance the control algorithm shortens the step and tries again. A field that does not obey Maxwell's equations can fool the control algorithm into thinking that the error is always larger than the allowed tolerance for any finite step size. A typical situation is where the field has an unphysical step across some boundary.

6.5.3 PTC Integration

The `integrator_order` element attribute is the order of the integration formula for `Symp_Lie_PTC` and is used for constructing Taylor maps. Possible values are

```
integrator_order = 2, 4, 6, or 8
```

Essentially, an integrator order of n means that the error in an integration step scales as dz^{n+1} where dz is the slice thickness. For a given number of steps a higher order will give more accurate results but a higher order integrator will take more time per step. It turns out that for wigglers, after adjusting `ds_step` for a given accuracy, the order 2 integrator is the fastest. This is not surprising given the highly nonlinear nature of a wiggler. Note that `symp_lie_bmad` always uses an order 2 integrator independent of the setting of `integrator_order`. The setting of `8` is not implemented for all elements. If `8` is set for a given element type that does not support it, a value of `6` will be used instead.

When tracking uses the `PTC` library (§1.4), there are two global parameters that can be set in the lattice file that affect the calculation. These are:

```
ptc_com[exact_model]    = <Logical>  ! "exact" tracking? Default: False
ptc_com[exact_misalign] = <Logical>  ! "exactly" misalign elements? Default: True
```

The default for `exact_model` is `True` and the default for `exact_misalign` is `True`.

The `exact_model` parameter sets whether PTC uses an "exact" model for tracking. Essentially this means that the paraxial approximation (§26.3) is made for `exact_model` set to `False` and is not made if set to `True`. This can be important, for example, for bend tracking when the bend radius is small.

In PTC, exact modeling can be set on an element-by-element basis. Currently *Bmad* does not support specifying element-by-element setting of exact modeling. However, PTC does not have a non-exact

tracking option for elements that have an electric field. In this case, PTC tracking will always be exact independent of the setting of `exact_model`. Additionally, for elements with an electric field, tracking will not be symplectic.

The `exact_misalign` parameter determines whether misalignments are handled exactly or whether approximations are made that will speed up the calculation.

In addition to the above parameters, how the Hamiltonian is split when tracking with `PTC` can be set for individual elements using the `ptc_integration_type` parameter. Possible settings of this parameter are

```
drift_kick    ! See Eq. (125) of [Forest06]
matrix_kick   ! See Eq. (132) of [Forest06]. Default
ripken_kick   ! See Eq. (130) of [Forest06]
```

Example:

```
q2: quad, l = 0.6, k1 = 0.34, ptc_integration_type = drift_kick
```

A discussion of the different types of integration schemes is given by Forest[Forest06]. The equation that shows the appropriate splitting of the Hamiltonian for each integration type is referenced in the above list. The `ripken_kick` type is for benchmarking with the `SixTrack` program and is not otherwise generally useful. The difference between `drift_kick` and `matrix_kick` is that with `drift_kick` the quadrupolar part of the magnetic multipole is is included in the applied kick between drifts while in the `matrix_kick` method the quadrupolar component is used for the "matrix" tracking between kicks. With the `matrix_kick` method the tune of a machine tends to be insensitive to how many integration steps (set by `ds_step` or `n_steps`) are used.

PTC does not implement `matrix_kick` tracking for elements with an electric field. In this case, the setting of `ptc_integration_type` is ignored and tracking will be `drift_kick`. Thus, if an electric field is introduced into an element, more integration steps may be required to get the correct tune.

6.6 Symplectic Versus Non-Symplectic Tracking

When selecting tracking methods for lattice elements, there are several factors to consider, including symplecticity. Despite its emphasis in accelerator textbooks, symplecity (or the lack therof) is typically only relevant for long-term tracking when there is minimal radiation emission over many turns. That is, the potential problem with non-symplectic tracking is the buildup of errors over many turns. Thus, computations that involve only tracking through the lattice from beginning to end – like calculating Twiss functions or tracking through a linac – generally do not benefit from symplectic tracking. More important in these cases is the speed of the calculation, which can be obtained with the cost of non-symplecticity.

The motion of particles that radiate is not symplectic. Thus, symplectic tracking for the non-radiative part of the motion may not be needed if radiation is large enough. For example, for simulations of the Cornell CESR ring with electron and positron beam energies of order 1 GeV to 10 GeV and with damping times on the order of 10,000 turns, the `bmad_standard` tracking has proved quite adequate. However in other cases with radiation, the symplectic error may cause an extra damping or anti-damping effect, giving equilibrium beam sizes that are an under/overestimate of the actual beam sizes. When opting for speed versus symplecticity in long term tracking over many turns, care should be taken to ensure that the effects of the non-symplecticity are minimal.

6.7 Symplectify Attribute

The `symplectify` attribute

```
symplectify = <Logical>
```

is used to make the transfer matrix for an element symplectic. The linear transport matrix may be non–symplectic for a number of reasons. For example, the linear matrix that comes from expanding a Taylor Map around any point that is not the origin of the map is generally not symplectic. The transfer matrix for an element can be symplectified by setting the `symplectify` attribute to True. See section §25.3 for details on how a matrix is symplectified. The default value of `symplectify`, if it is not present, is `False`. If it is present without a value then it defaults to true. Examples:

```
s1: sextupole, l = 0.34                       ! symplectify = False
s1: sextupole, symplectify = True, l = 0.34   ! symplectify = True
s1: sextupole, symplectify, l = 0.34          ! symplectify = True
```

Note that for elements like an `lcavity` where the reference momentum at the downstream end of the element is different from the upstream end, the transfer matrix is never symplectic. In this case, "symplectification" involves first transforming the transfer matrix so that the reference momentum is the same upstream and downstream, then performing symplectification, and finally back transforming the reference momentum to their original values.

6.8 taylor_map_include_offsets Attribute

The `taylor_map_includes_offsets` attribute sets whether the Taylor map generated for an element includes the affect due to the elements (mis)orientation in space. That is, the affect of any pitches, offsets or tilt (§5.6). The default is `True` which means that the Taylor map will include such effects.

How `taylor_map_includes_offsets` is set will not affect the results of tracking or the Jacobian matrix calculation. What is affected is the speed of the calculations. With `taylor_map_includes_offsets` set to `True` the Taylor map will have to be recalculated each time an element is reoriented in space. On the other hand, with `taylor_map_includes_offsets` set to `False` each tracking and Jacobian matrix calculation will include the extra computation involving the effect of the orientation. Thus if an element's orientation is fixed it is faster to set `taylor_map_includes_offsets` to `True` and if the orientation is varying it is faster to set `taylor_map_includes_offsets` to `False`.

If the global parameter `bmad_com%conserve_taylor_maps` (§11.4) is set to True (the default), then, if an element is offset within a program, and if `taylor_map_include_offsets` is set to True for that element, *Bmad* will toggle `taylor_map_include_offsets` to False to conserve the map.

Chapter 7

Beam Lines and Replacement Lists

This chapter describes how to define the ordered list of elements that make up a lattice branch (§2.2). In a lattice, branches may be connected together using `fork` or `photon fork` elements (`s:fork`), or by using `multipass` (§9).

7.1 Branch Construction Overview

A lattice branch (§2.2) is defined in a lattice file using what are called `beam lines` (§7.2) and `replacement lists` (§7.6). The `beam lines` are divided into two types - lines with (§7.5) and lines without (§7.2) `replacement arguments`. This essentially corresponds to the *MAD* definition of lines and lists. There can be multiple `beam lines` and `replacement lists` defined in a lattice file and lines and lists can be nested inside other lines and lists.

Since lines can be nested within other lines, The same element name may be repeated multiple times in a branch. To distinguish between multiple elements of the same name, lines and lists may be `tagged` (§7.8) to produce unique element names.

A marker element named `END` will, by default, be placed at the ends of all the branches unless a `parameter[no_end_marker]` statement (§10.1) is used to suppress the insertion. Additionally, if an ending marker named `END` is already present in the lattice file, no extra marker will be created.

Branches are ordered in an array (§2.3) and each branch is assigned an index number starting with index 0. When there are multiple branches in a lattice, the reference orbit (§16.1.1) of a branch must not depend upon details of branches later on in the array. *Bmad* depends upon this and calculates the reference orbits of the branches one at a time starting with the first branch.

7.2 Beam Lines and Lattice Expansion

A `beam line` without arguments has the format

```
label: line = (member1, member2, ...)
```

where `member1`, `member2`, etc. are either elements, other `beam lines` or `replacement lists`, or sublines enclosed in parentheses. Example:

```
line1: line = (a, b, c)
```

```
line2: line = (d, line1, e)
use, line2
```

The `use` statement is explained in Section §7.7. This example shows how a `beam line` member can refer to another `beam line`. This is helpful if the same sequence of elements appears repeatedly in the lattice.

The process of constructing the ordered sequences of elements that comprise the branches of the lattice is called `lattice expansion`. In the example above, when `line2` is expanded to form the lattice (in this case there is only one branch so `lattice` and `branch` can be considered synonymous), the definition of `line1` will be inserted in to produce the following lattice:

```
beginning, d, a, b, c, e, end
```

The `beginning` and `end` marker elements are automatically inserted at the beginning and end of the lattice. The `beginning` element will always exist but insertion of the `end` element can be suppressed by inserting into the lattice:

```
parameter[no_end_marker] = T    ! See: §10.1
```

Lattice expansion occurs either at the end after the lattice file has been parsed, or, during parsing, at the point where an `expand_lattice` statement (§3.26) is found.

Each element is assigned an `element index` number starting from 0 for the `beginning` element, 1 for the next element, etc.

In the expanded lattice, any `null_Ele` type elements (§4.38) will be discarded. For example, if element `b` in the above example is a `null_Ele` then the actual expanded lattice will be:

```
beginning, d, a, c, e, end
```

A member that is a line or list can be "reflected" (elements taken in reverse order) if a negative sign is put in front of it. For example:

```
line1: line = (a, b, c)
line2: line = (d, -line1, e)
```

`line2` when expanded gives

```
d, c, b, a, e
```

It is important to keep in mind that line reflection is `not` the same as going backwards through elements. For example, if an `sbend` or `rbend` element (§4.5) is reflected, the face angle of the upstream edge (§16.1.3) is still specified by the `e1` attribute and not the `e2` attribute. True element reversal can be accomplished as discussed in Sec. §7.4.

Reflecting a subline will also reflect any sublines of the subline. For example:

```
line0: line = (y, z)
line1: line = (line0, b, c)
line2: line = (d, -line1, e)
```

`line2` when expanded gives

```
d, c, b, z, y, e
```

A repetition count, which is an integer followed by an asterisk, means that the member is repeated. For example

```
line1: line = (a, b, c)
line2: line = (d, 2*line1, e)
```

`line2` when expanded gives

```
d, a, b, c, a, b, c, e
```

Repetition count can be combined with reflection. For example

```
line1: line = (a, b, c)
line2: line = (d, -2*line1, e)
```

`line2` when expanded gives

```
d, c, b, a, c, b, a, e
```

Instead of the name of a line, subline members can also be given as an explicit list using parentheses. For example, the previous example could be rewritten as

```
line2: line = (d, -2*(a, b, c), e)
```

Lines can be defined in any order in the lattice file so a subline does not have to come before a line that references it. Additionally, element definitions can come before or after any lines that reference them.

A line can have the `multipass` attribute. This is covered in §9.

7.3 Line Slices

A line "`slice`" is a section of a line from some starting element to some ending element. A line slice can be used to construct a new line similar to how an unsliced line is used to construct a new line. An example will make this clear:

```
line1: line = (a, b, c, d, e)
line2: line = (z1, line1[b:d], z2)
```

The line slice `line1[b:d]` that is used to construct `line2` consists of the elements in `line1` from element `b` to element `d` but not elements `a` or `e`. When `line2` is expanded, it will have the elements:

```
z1, b, c, d, z2
```

The general form for line slices is

```
line_name[element1:element2]
```

where `line_name` is the name of the line and `element1` and `element2` delimit the beginning and ending positions of the slice. The beginning and ending element names may be omitted and, if not present, the default is the beginning element and ending element of the line respectively. Thus, for example, "`line4[:q1]`" represents the list of elements from the start of `line4` up to, and including the element `q1`.

If there are multiple elements of the same name, the double hash `##` symbol (§3.6) can be use to denote the N^{th} element of a given name. If double hash is not used, the first instance of a given element name is assumed. That is, something like "`q1`" is equivalent to "`q1##1`".

Wild card characters and `class::element_name` syntax (§3.6) are not allowed with slice element names.

Line slicing of a given line occurs after the line has been expanded (all sublines and line slices substituted in). Thus, the following makes sense:

```
line1: line = (a, b, c, d, e)
line2: line = (z1, line1, z2)
line3: line = (line2[z1:c])
```

7.4 Element Orientation Reversal

An element's orientation is `reversed` if particles traveling through it enter at the "exit" end and leave at the "entrance" end. Being able to reverse elements is useful, for example, in describing the interaction region of a pair of rings where particles of one ring are going in the opposite direction relative to the particles in the other ring.

Element reversal is indicated by using a double negative sign "−−" prefix. The double negative sign prefix can be applied to individual elements or to a line. If it is applied to a line, the line is both reflected (same as if a single negative sign is used) and each element is reversed. For example:

```
line1: line = (a, b, --c)
line2: line = (--line1)
line3: line = (c, --b, --a)
```

In this example, `line2` and `line3` are identical. Notice that the reversal of a reversed element makes the element unreversed.

Another example involving element reversal is given in Section §13.6.

Reversed elements, unlike other elements, have their local z-axis pointing in the opposite direction to the local s-axis (§16.1.3). This means that there must be a `reflection patch` (§16.2.6) between reversed and unreversed elements. Since this complicates matters, it is generally only useful to employ element reversal in cases where there are multiple intersecting lines with particle beams going in opposite directions through some elements (for example, colliding beam interaction regions). In this case, element reversal is typically used with `multipass` (§9) and the lattice will contain a branch of unreversed elements for simulating particles going in one direction along with a branch of reversed elements to simulate particle going in the other direction.

Where reversed elements are not needed, it is simple to define elements that are effectively reversed. For example:

```
b00: bend, angle = 0.023, e1 = ...
b00_rev: b00, angle = -b00[angle], e1 = -b00[e2], e2 = -b00[e1]
```

and `b00_rev` serves as a reversed version of `b00`.

Internally, *Bmad* associates an `orientation` attribute with each element. This attribute is set to -1 for reversed elements and 1 for unreversed elements.

7.5 Beam Lines with Replaceable Arguments

`Beam lines` can have an argument list using the following syntax

```
line_name(dummy_arg1, dummy_arg2, ...): LINE = (member1, member2, ...)
```

The dummy arguments are replaced by the actual arguments when the line is used elsewhere. For example:

```
line1(DA1, DA2): line = (a, DA2, b, DA1)
line2: line = (h, line1(y, z), g)
```

When `line2` is expanded the actual arguments of `line1`, in this case (y, z), replaces the dummy arguments `(DA1, DA2)` to give for `line2`

```
h, a, z, b, y, g
```

Unlike *MAD*, `beam line` actual arguments can only be elements or `beam lines`. Thus the following is not allowed

```
line2: line = (h, line1(2*y, z), g)   ! NO: 2*y NOT allowed as an argument.
```

7.6 Lists

When a lattice is expanded, all the lattice members that correspond to a name of a `list` are replaced successively, by the members in the `list`. The general syntax is

```
label: LIST = (member1, member2, ...)
```

For example:

```
my_list1 list = (a, b, c)
line1: line = (z1, my_list, z2, my_list, z3, my_list, z4, my_list)
use, line1
```

When the lattice is expanded the first instance of `my_list` in `line1` is replaced by `a` (which is the first element of `my_list`), the second instance of `my_list` is replaced by `b`, etc. If there are more instances of `my_list` in the lattice then members of `my_list`, the replacement starts at the beginning of `my_list` after the last member of `my_list` is used. In this case the lattice would be:

```
z1, a, z2, b, z3, c, z4, a
```

members of a `replacement list` can only be simple elements and not other lines or lists. For example, the following is not allowed:

```
line1: line = (a, b)
my_list: list = (2*line1)  ! Lines cannot be list members.
```

A repetition count is permitted

```
my_list1: list = (2*a, b)
my_list2: list = (a, a, b) ! Equivalent to my_list1
```

7.7 Use Statement

The particular line or lines that defines the root branches (§2.3) to be used in the lattice is selected by the `use` statement. The general syntax is

```
use, line1, line2 ...
```

For example, `line1` may correspond to one ring and `line2` may correspond to the other ring of a dual ring colliding beam machine. In this case, `multipass` (§9) will be needed to describe the common elements of the two rings. Example

```
use, e_ring, p_ring
```

would pick the lines `e_ring` and `p_ring` for analysis. These will be the `root` branches.

`use` statements can come anywhere in the lattice, even before the definition of the lines they refer to. Additionally, there can be multiple `use` statements. The last `use` statement in the file defines which `line` to use.

The total number of branches in the lattice is equal to the number of lines that appear on the `use` statement plus the number of `fork` and `photon_fork` elements that branch to a new branch.

To set such things as the geometry of a branch, beginning Twiss parameters, etc., see Section `s:beginning`.

7.8 Tagging Lines and Lists

When a lattice has repeating lines, it can be desirable to differentiate between repeated elements. This can be done by tagging lines with a `tag`. An example will make this clear:

```
line1: line = (a, b)
line2: line = (line1, line1)
use, line2
```

When expanded the lattice would be:

```
a, b, a, b
```

The first and third elements have the same name "a" and the second and fourth elements have the same name "b". Using tags the lattice elements can be given unique names. lines or lists are tagged using the at (@) sign. The general syntax is:

```
tag_name@line_name                          ! Syntax for lines
tag_name@list_name                          ! Syntax for lists
tag_name@replacement_line(arg1, arg2, ...)  ! Syntax for replacement lines.
```

Thus to differentiate the lattice elements in the above example `line2` needs to be modified using tags:

```
line1: line = (a, b)
line2: line = (t1@line1, t2@line1)
use, line2
```

In this case the lattice elements will have names of the form:

```
tag_name.element_name
```

In this particular example, the lattice with tagging will be:

```
t1.a, t1.b, t2.a, t2.b
```

Of course with this simple example one could have just as easily not used tags:

```
t1.a: a;   t2.a: a
t1.b: b;   t2.b: b
line1: line = (t1.a, t1.b, t2.a, t2.b)
use, line2
```

But in more complicated situations tagging can make for compact lattice files.

When lines are nested, the name of an element is formed by concatenating the tags together with dots in between in the form:

```
tag_name1.tag_name2. ... tag_name_n.element_name
```

An example will make this clear:

```
list1 = (g, h)
line1(y, z) = (a, b)
line2: line = (t1@line1(a, b))
line3: line = (line2, hh@list1)
line4: line = (z1@line3, z2@line3)
use, line4
```

The lattice elements in this case are:

```
z1.t1.a, z1.t1.b, z1.hh.g, z2.t1.a, z2.t1.b, z1.hh.h
```

To modify a particular tagged element the lattice must be expanded first (§3.26). For example:

```
line1: line = (a, b)
line2: line = (t1@line1, t2@line1)
use, line2
expand_lattice
t1.b[k1] = 1.37
b[k1] = 0.63       ! This statement generates an error
```

After the lattice has been expanded there is no connection between the original `a` and `b` elements and the elements in the lattice like `t1.b`. Thus the last line in the example where the `k1` attribute of`b` is modified generates an error since there are no elements named `b` in the lattice.

Chapter 8

Superposition

This chapter covers the concept of `superposition`. `Superposition` is used when elements overlap spatially. With `superposition`, `lord` and `slave` elements (§2.4) are constructed by *Bmad* to hold the necessary information. The `lord` elements will represent the "physical" element while the `slave` elements will embody the "beam path".

8.1 Superposition Fundamentals

In practice the field at a particular point in the lattice may be due to more than one physical element. One example of this is a quadrupole magnet inside a larger solenoid magnet as shown in Fig. 8.1A. *Bmad* has a mechanism to handle this using what is called "superposition". A simple example shows how this works (also see section §2.4):

```
Q: quad, l = 4
D: drift, l = 12
S: solenoid, l = 8, superimpose, ref = Q, ele_origin = beginning
M: marker, superimpose, ref = S, offset = 1
lat: line = (Q, D)
use, lat
```

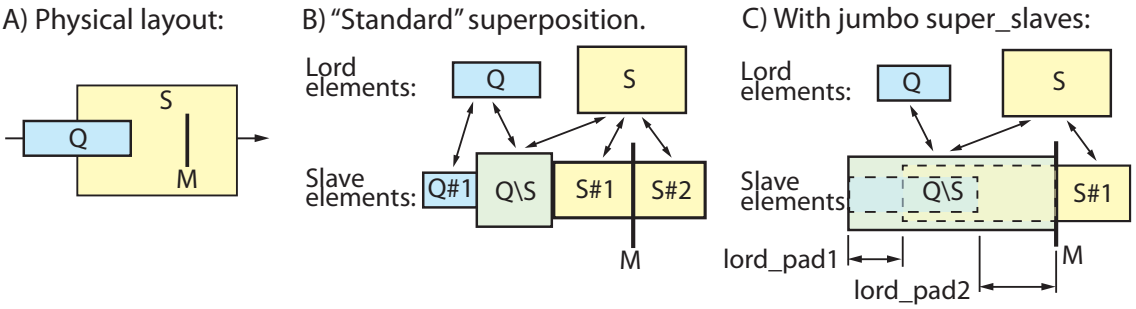

Figure 8.1: Superposition example. A) The physical layout involves a quadrupole partially inside a solenoid. B) The standard superposition procedure involves creating `super_slave` elements whose edges are at the boundaries where the physical elements overlap. C) When jumbo `super_slaves` are created, the `super_slaves` span the entire space where elements overlap.

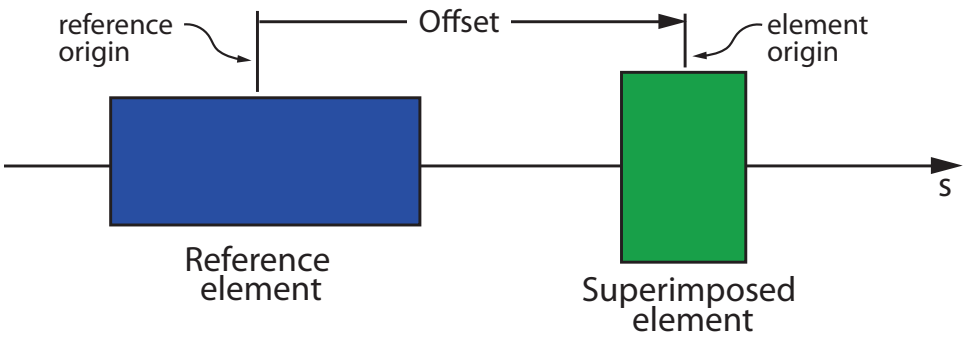

Figure 8.2: The superposition offset is the distance along the local reference orbit from the origin point of the reference element to the origin point of the element being superimposed.

The `superimpose` attribute of element `S` superimposes `S` over the lattice `(Q, D)`. The placement of `S` is such that the beginning of `S` is coincident with the center of `Q` (this is is explained in more detail below). Additionally, a marker `M` is superimposed at a distance of +1 meter from the center of `S`. The tracking part of the lattice (§2.4) looks like:

```
        Element    Key          Length   Total
  1)    Q#1        Quadrupole   2          2
  2)    Q\S        Sol_quad     2          4
  3)    S#1        Solenoid     3          7
  4)    M          Marker       0
  4)    S#2        Solenoid     3         10
  5)    D#2        Drift        4         14
```

What *Bmad* has done is to split the original elements `(Q, D)` at the edges of `S` and then `S` was split where `M` is inserted. The first element in the lattice, `Q#1`, is the part of `Q` that is outside of `S`. Since this is only part of `Q`, *Bmad* has put a `#1` in the name so that there will be no confusion. (a single `#` has no special meaning other than the fact that *Bmad* uses it for mangling names. This is opposed to a double `##` which is used to denote the N^{th} instance of an element (§3.6). The next element, `Q\S`, is the part of `Q` that is inside `S`. `Q\S` is a combination solenoid/quadrupole element as one would expect. `S#1` is the part of `S` that is outside `Q` but before `M`. This element is just a solenoid. Next comes `M`, `S#1`, and finally `D#2` is the rest of the drift outside `S`.

In the above example, `Q` and `S` will be `super_lord` elements (`s:lord.slave`) and four elements in the tracking part of the lattice will be `super_slave` elements. This is illustrated in Fig. 8.1B.

Notice that the name chosen for the `sol_quad` element `Q\S` is dependent upon what is being superimposed upon what. If `Q` had been superimposed upon `S` then the name would have been `S\Q`.

When *Bmad* sets the element class for elements created from superpositions, *Bmad* will set the class of the element to something other than an `em_field` element (§4.17) if possible. If no other possibilities exist, *Bmad* will use `em_field`. For example, a `quadrupole` superimposed with a `solenoid` will produce a `sol_quad super_slave` element but a `solenoid` superimposed with a `rfcavity` element will produce an `em_field` element since there is no other class of element that can simultaneously handle solenoid and RF fields. An `em_field super_slave` element will also be created if any of the superimposing elements have a non-zero orientation (§5.6) since it is not, in general, possible to construct a slave element that properly mimics the effect of a non-zero orientation.

With the lattice broken up like this *Bmad* has constructed something that can be easily analyzed. However, the original elements `Q` and `S` still exist within the lord section of the lattice. *Bmad* has bookkeeping routines so that if a change is made to the `Q` or `S` elements then these changes can get

propagated to the corresponding slaves. It does not matter which element is superimposed. Thus, in the above example, `S` could have been put in the Beam Line (with a drift before it) and `Q` could then have been superimposed on top and the result would have been the same (except that the split elements could have different names).

If an element has zero length (for example, a `marker` element), is superimposed, or is superimposed upon, then the element will remain in the tracking part of the lattice and there will be no corresponding lord element. See Fig. 8.1.

Superimpose syntax:

```
Q: quad, superimpose, ...       ! Superimpose element Q.
Q: quad, superimpose = T, ...   ! Same as above.
Q: quad, ...                    ! First define element Q ...
Q[superimpose] = T              !   ... and then superimpose.
Q[superimpose] = F              ! Suppress superposition.
```

Superposition happens at the end of parsing so the last set of the `superimpose` for an element will override previous settings.

It is also possible to superimpose an element using the `superimpose` command which has the syntax:

```
superimpose, element = <ele-name>, ...
```

With the same optional superposition parameters (`ref`, `offset`, etc.) given below. Example:

```
superimpose, element = Q1, ref = B12, offset = 1.3,
                              ele_origin = beginning, ref_origin = end
```

Note: Superposition using the superimpose statement allows superimposing the same element with multiple reference elements and/or multiple offsets. The drawback is that superposition using the superimpose statement may not be switched off later in the lattice file.

The placement of a superimposed element is illustrated in Fig. 8.2. The placement of a superimposed element is determined by three factors: An origin point on the superimposed element, an origin point on the reference element, and an offset between the points. The attributes that determine these three quantities are:

```
create_jumbo_slave = <Logical>     ! See §8.8
wrap_superimpose   = <Logical>     ! Wrap if element extends past lattice ends?
ref          = <lattice_element>
offset       = <length>            ! default = 0
ele_origin   = <origin_location>   ! Origin pt on element.
ref_origin   = <origin_location>   ! Origin pt on ref element.
```

`ref` sets the reference element. If `ref` is not present then the start of the lattice is used (more precisely, the start of branch 0 (§2.2)). Wild card characters (§3.6 can be used with `ref`. If `ref` matches to multiple elements (which may also happen without wild card characters if there are multiple elements with the name given by `ref`) in the lattice a superposition will be done, one for each match.

The location of the origin points are determined by the setting of `ele_origin` and `ref_origin`. The possible settings for these parameters are

```
beginning       ! Beginning (upstream) edge of element
center          ! Center of element. Default.
end             ! End (downstream) edge of element
```

`center` is the default setting. `Offset` is the longitudinal offset of the origin of the element being superimposed relative to the origin of the reference element. The default offset is zero. A positive offset moves the element being superimposed in the `downstream` direction if the reference element has a normal longitudinal `orientation` (§7.4) and vice versa for the reference element has a reversed longitudinal orientation.

Note: There is an old syntax, deprecated but still supported for now, where the origin points were specified by the appearance of:

```
ele_beginning         ! Old syntax. Do not use.
ele_center            ! Old syntax. Do not use.
ele_end               ! Old syntax. Do not use.
ref_beginning         ! Old syntax. Do not use.
ref_center            ! Old syntax. Do not use.
ref_end               ! Old syntax. Do not use.
```

For example, "ele_origin = beginning" in the old syntax would be "ele_beginning".

The element begin superimposed may be any type of element except `drift`, `group`, `overlay`, and `girder` control elements. The reference element used to position a superimposed element may be a `group` or `overlay` element as long as the `group` or `overlay` controls the attributes of exactly one element. In this case, the controlled element is used as the reference element.

When the attributes of a super_slave are computed from the attributes of its super_lords, some types of attributes may be "missing". For example, it is, in general, not possible to set appropriate aperture attributes (§5.8) of a super_slave if the lords of the slave have differing aperture settings. When doing calculations, *Bmad* will use the corresponding attributes stored in the lord elements to correctly calculate things.

The `no_superposition` statement (§3.23) can be used to turn off superpositioning

8.2 Superposition Wrapping

By default, a superimposed element that extends beyond either end of the lattice will be "wrapped around" so part of the element will be at the beginning of the lattice and part of the element will be at the end. For consistency's sake, this is done even if the `geometry` is set to `open` (for example, it is sometimes convenient to treat a circular lattice as linear). Example:

```
d: drift, l = 10
q: quad, l = 4, superimpose, offset = 1
machine: line = (d)
use, machine
```

The lattice will have five elements in the tracking section:

```
      Element     Key             Length
0)    BEGINNING   Beginning_ele   0
1)    Q#2         Quadrupole      3   ! Slave representing beginning of Q element
2)    D#1         Drift           6
3)    Q#1         Quadrupole      1   ! Slave representing end of Q element
4)    END         Marker          0
```

And the lord section of the lattice will have the element `Q`.

To not wrap an element that is being superimposed, set the `wrap_superimpose` logical to `False`. Following the above example, if the definition of`q` is extended by adding `wrap_superimpose`:

```
q: quad, l = 4, superimpose, offset = 1, wrap_superimpose = F
```

In this instance there are four elements in the tracking section:

```
      Element     Key             Length
0)    BEGINNING   Beginning_ele   0
1)    Q           Quadrupole      4
2)    D#1         Drift           7
4)    END         Marker          0
```

And the lord section of the lattice will not have any elements.

8.3 Superposition of Zero Length Elements

To superimpose a zero length element "S" next to a zero length element "Z", and to make sure that S will be on the correct side of Z, set the ref_origin appropriately. For example:

```
S1: marker, superimpose, ref = Z, ref_origin = beginning
S2: marker, superimpose, ref = Z, ref_origin = end
Z: marker
```

The order of the elements in the lattice will be

```
S1, Z, S2
```

If ref_origin is not present or set to center, the ordering of the elements will be arbitrary.

If a zero length element is being superimposed at a spot where there are other zero length elements, the general rule is that the element will be placed as close as possible to the reference element. For example:

```
S1: marker, superimpose, offset = 1
S2: marker, superimpose, offset = 1
```

In this case, after S1 is superimposed at $s = 1$ meter, the superposition of S2 will place it as close to the reference element, which in this case is the BEGINNING elements at $s = 0$, as possible. Thus the final order of the superimposed elements is:

```
S2, S1
```

To switch the order while still superimposing S2 second one possibility is to use:

```
S1: marker, superimpose, offset = 1
S2: marker, superimpose, ref = S1, ref_origin = end
```

If a superposition uses a reference element, and there are N elements in the lattice with the reference element name, there will be N superpositions. For example, the following will split in two all the quadrupoles in a lattice:

```
M: null_ele, superimpose, ref = quadrupole::*
```

A null_ele (§4.38) element is used here so that there is no intervening element between split quadrupole halves as there would be if a marker element was used.

8.4 Superposition and Drifts

When a superposition is made that overlaps a drift, the drift, not being a "real" element, vanishes. That is, it does not get put in the lord section of the lattice. Note that if aperture limits (§5.8) have been assigned to a drift, the aperture limits can "disappear" when the superposition is done. Explicitly, if the exit end of a drift has been assigned aperture limits, the limits will disappear if the superimposed element overlays the exit end of the drift. A similar situation applies to the entrance end of a drift. If this is not desired, use a pipe element instead.

To simplify bookkeeping, a drift element may not be superimposed. Additionally, since drifts can disappear during superposition, to avoid unexpected behavior the superposition reference element may not be the N^{th} instance of a drift with a given name. For example, if there are a number of drift elements in the lattice named a_drft, the following is not allowed:

```
my_oct: octupole, ..., superimpose, ref = a_drft##2  ! This is an error
```

8.5 Superposition and Element Reversal

When superposition is done in a line where there is `element reversal` (§7.4), the calculation of the placement of a superimposed element is also "reversed" to make the relative placement of elements independent of any element reversal. An example will make this clear:

```
d1: drift, l = 1
d2: d1
q1: quad, l = 0.1, superimpose, ref = d1, offset = 0.2,
          ref_origin = beginning, ele_origin = beginning
q2: q1, ref = d2
p: patch, x_pitch = pi  ! Needed to separate reversed and unreversed.
this_line: line = (d1, p, --d2)
use, this_line
```

Since the reference element of the `q2` superposition, that is `d2`, is a reversed element, `q2` will be reversed and the sense of `offset`, `ref_origin`, and `ele_origin` will be reversed so that the position of `q2` with respect to `d2` will be the mirror image of the position of `q1` with respect to `d1`. The tracking part of the lattice will be:

```
Element:           d1#1    q1  d1#2   d2#2    q2   d2#1
Length:             0.2   0.1   0.7    0.7   0.1    0.3
Reversed element?:   No    No    No    Yes   Yes    Yes
```

Superposition with `line reflection` (§7.2) works the same way as line reversal.

8.6 Superposition and Sub-Lines

Sometimes it is convenient to do simulations with only part of a lattice. The rule for how superpositions are handled in this case is illustrated in the following example. Consider a lattice file which defines a `line` called `full` which is defined by two sublines called `sub1` and `sub2`:

```
sub1: line = ..., ele1, ...
sub2: line = ...
full: line = sub1, sub2
m1: marker, superimpose, ref = ele1, offset = 3.7
use, full
```

Now suppose you want to do a simulation using only the `sub2` line. Rather than edit the original file, one way to do this would be to create a second file which overrides the used line:

```
call, file = "full.bmad"
use, sub2
```

where `full.bmad` is the name of the original file. What happens to the superposition of `m1` in this case? Since `m1` uses a reference element, `ele1`, that is not in `sub1`, *Bmad* will ignore the superposition. Even though *Bmad* will ignore the superposition of `m1` here, *Bmad* will check that `ele1` has been defined. If `ele1` has not been defined, *Bmad* will assume that there is a typographic error and issue an error message.

Notice that in this case it is important for the superposition to have an explicit reference element since without an explicit reference element the superposition is referenced to the beginning of the lattice. Thus, in the above example, if the superposition were written like:

```
m1: marker, superimpose, offset = 11.3
```

then when the `full` line is used, the superposition of `m1` is referenced to the beginning of `full` (which is the same as the beginning of `sub1`) but when the `sub2` line is used, the superposition of `m1` is referenced to the beginning of `sub2` which is not the same as the beginning of `full`.

8.7 Superposition Order

Superpositions are done in the same order as the order in which the elements to be superimposed are defined. The superposition order typically does not affect the finished lattice but an effect can be be seen when superpositions use reference elements that are themselves superimposed. An example will make this clear:

```
A: quad, ...
B: octupole, ...
C: sbend, ...
B[superimpose] = T
B[ref] = C
A[superimpose] = T
A[ref] = B
my_line: line = (C)
use, my_line
```

Here element `A` is defined first (the order of definitions is not affected by element modifications afterwards) and then `B`.

Since, when element `A` is to be superimposed, there are no instances of `B` in the expanded lattice, and element `A` uses `B` as a reference, the finished lattice will not contain any instances of `A`. On the other hand, when element `B` is to be superimposed, there is an instance of element `C` to be used as a reference, and so the finished lattice will contain an instance of element `B`.

8.8 Jumbo Super_Slaves

The problem with the way `super_slave` elements are created as discussed above is that edge effects will not be dealt with properly when elements with non-zero fields are misaligned. When this is important, especially at low energy, a possible remedy is to instruct *Bmad* to construct "`jumbo`" super_slave elements. The general idea is to create one large `super_slave` for any set of overlapping elements. Returning to the superposition example at the start of Section §8, If the superposition of solenoid `S` is modified to be

```
S: solenoid, l = 8, superimpose, ref = Q, ele_origin = beginning,
             create_jumbo_slave = T
```

The result is shown in Fig. 8.1C. The tracking part of the lattice will be

```
       Element   Key         Length  Total
1)     Q\S       Sol_quad     2        4
2)     M         Marker       0
3)     S#2       Solenoid     3       10
4)     D#2       Drift        4       14
```

`Q` and part of `S` have been combined into a jumbo `super_slave` named `Q\S`. Since the `super_lord` elements of a jumbo `super_slave` may not completely span the slave two attributes of each lord will be set to show the position of the lord within the slave. These two attributes are

```
lord_pad1    ! offset at upstream end
lord_pad2    ! offset at downstream end
```

`lord_pad1` is the distance between the upstream edge of the jumbo `super_slave` and a `super_lord`. `lord_pad2` is the distance between the downstream edge of a `super_lord` and the downstream edge of the jumbo `super_slave`. With the present example, the lords have the following padding:

```
         lord_pad1    lord_pad2
Q            0            3
S            2            0
```

The following rule holds for all super lords with and without jumbo slaves:

```
Sum of all slave lengths = lord length + lord_pad1 + lord_pad2
```

One major drawback of jumbo `super_slave` elements is that the `tracking_method` (§6.1) will, by necessity, have to be `runge_kutta`, or `time_runge_kutta` and the `mat6_calc_method` (§6.2) will be set to `tracking`.

Notice that the problem with edge effects for non-jumbo `super_slave` elements only occurs when elements with nonzero fields are superimposed on top of one another. Thus, for example, one does not need to use jumbo elements when superimposing a `marker` element.

Another possible way to handle overlapping fields is to use the `field_overlaps` element attribute as discussed in §5.18.

8.9 Changing Element Lengths when there is Superposition

When a program is running, if `group` (§4.26) or `overlay` (§4.40) elements are used to vary the length of elements that are involved in superimposition, the results are different from what would have resulted if instead the lengths of the elements where changed in the lattice file. There are two reasons for this. First, once the lattice file has been parsed, lattices can be "mangled" by adding or removing elements in a myriad of ways. This means that it is not possible to devise a general algorithm for adjusting superimposed element lengths that mirrors what the effect of changing the lengths in the lattice file.

Second, even if a lattice has not been mangled, an algorithm for varying lengths that is based on the superimpose information in the lattice file could lead to unexpected results. To see this consider the first example in Section §8. If the length of `S` is varied in the lattice file, the upstream edge of `S` will remain fixed at the center of `Q` which means that the length of the `super_slave` element `Q#1` will be invariant. On the other hand, if element `S` is defined by

```
S: solenoid, l = 8, superimpose, offset = 6
```

This new definition of `S` produces exactly the same lattice as before. However, now varying the length of `S` will result in the center of `S` remaining fixed and the length of `Q#1` will not be invariant with changes of the length of `S`. This variation in behavior could be very confusing since, while running a program, one could not tell by inspection of the element positions what should happen if a length were changed.

To avoid confusion, *Bmad* uses a simple algorithm for varying the lengths of elements involved in superposition: The rule is that the length of the most downstream `super_slave` is varied. With the first example in Section §8, the `group` `G` varying the length of `Q` defined by:

```
G: group = {Q}, var = {l}
```

would vary the length of `Q\S` which would result in an equal variation of the length of `S`. To keep the length of `S` invariant while varying `Q` the individual `super_slave` lengths can be varied. Example:

```
G2: group = {Q#1, S#1:-1}, var = {l}
```

The definition of `G2` must be placed in the lattice file after the superpositions so that the super slaves referred to by `G2` have been created.

In the above example there is another, cleaner, way of achieving the same result by varying the downstream edge of `Q`:

```
G3: group = {Q}, var = {end_edge}
```

Chapter 9

Multipass

This chapter covers the concept of `multipass`. `Multipass` is used when an element is "shared" between branches such as the interaction region shared by two storage rings, or when a beam goes through the same physical element in a branch multiple times as in an energy recovery linac. With `multipass`, `lord` and `slave` elements (§2.4) are constructed by *Bmad* to hold the necessary information. The `lord` elements will represent the "physical" element while the `slave` elements will embody the "beam path".

9.1 Multipass Fundamentals

`Multipass` lines are a way to handle the bookkeeping when different elements being tracked through represent the same physical element. For example, consider the case where dual ring colliding beam machine is to be simulated. In this case the lattice file might look like:

```
ring1: line = (..., IR_region, ...)
ring2: line = (..., --IR_region, ...)
IR_region: line = (Q1, ....)
use, ring1, ring2
```

[The "–" construct means go through the line backwards (§7.4)] In this case, the `Q1` element in `ring1` represents the same physical element in `ring2`. Thus the parameters of both the `Q1`s should be varied in tandem. This can be done automatically using `multipass`. The use of multipass simplifies lattice and program development since the bookkeeping details are left to the *Bmad* bookkeeping routines.

To illustrate how `multipass` works, consider the example of an Energy Recovery Linac (ERL) where the beam will recirculate back through the LINAC section to recover the energy in the beam before it is dumped. In *Bmad*, this situation can simulated by designating the LINAC section as `multipass`. The lattice file might look like:

```
RF1: lcavity
linac: line[multipass] = (RF1, ...)
erl: line = (linac, ..., linac)
use, erl
expand_lattice
RF1\2[phi0_multipass] = 0.5
```

The line called `linac` is designated as `multipass`. This `linac` line appears twice in the line `erl` and `erl` is the root line for lattice expansion. The lattice constructed from `erl` will have two `RF1` elements in the tracking part of the lattice:

```
RF1\1, ..., RF1\2, ...
```

Since the two elements are derived from a `multipass` line, they are given unique names by adding a `\N` suffix where `N` is an integer indicating the "pass index". These types of elements are known as `multipass_slave` elements. In addition, to the `multipass_slave` elements, there is a `multipass_lord` element (that doesn't get tracked through) called `RF1` in the lord part of the lattice (§2.4). Changes to attributes of the lord `RF1` element will be passed to the slave elements by *Bmad*'s bookkeeping routines. As explained in detail in §9.3, assuming that the phase of `RF1\1` gives acceleration, to make the `RF1\2` ele to decelerate, the `phi0_multipass` attribute of `RF1\2` is set to 0.5. This is an attribute that *Bmad*'s bookkeeping routines will not touch when transferring attribute values from the lord `RF1` to its slaves. That is, `phi0_multipass` can be set differently for different slaves. Notice that the `phi0_multipass` attribute had to be set after the `expand_lattice` command (§3.26) is used to expand the lattice. This is true since *Bmad* does immediate evaluation and `RF1\2` does not exist before the lattice is expanded.

"Intrinsic" attributes are attributes that must, to make sense physically, be the same for all slaves of a given multipass lord. The element length is one such example. The following non-intrinsic attributes can be set in a multipass slave and will not affect the corresponding attributes in the lord or the other slaves of the lord:

```
csr_ds_step           num_steps
csr_method            ptc_integration_type
ds_step               spin_tracking_method
field_calc            space_charge_method
integrator_order      tracking_method
mat6_calc_method      phi0_multipass
```

Multiple elements of the same name in a multipass line are considered physically distinct. Example:

```
m_line: line[multipass] = (A, A, B)
u_line: line = (m_line, m_line)
use, u_line
```

In this example the tracking part of the lattice is

```
A\1, A\1, B\1, A\2, A\2, B\2
```

In the control section of the lattice there will be two multipass lords called `A` and one called `B`. [That is, *Bmad* considers the lattice to have three physically distinct elements.] The first `A` lord controls the 1^{st} and 4^{th} elements in the tracking part of the lattice and the second `A` lord controls the 2^{nd} and 5^{th} elements. If `m_line` was *not* marked `multipass`, the tracking part of the lattice would have four `A` and two `B` elements and there would be no lord elements.

Sublines contained in a multipass line that are themselves not marked multipass act the same as if the elements of the subline where substituted directly in place of the subline in the containing line. For example:

```
a_line: line = (A)
m_line: line[multipass] = (a_line, a_line, B)
u_line: line = (m_line, m_line)
use, u_line
```

In this example, `a_line`, which is a subline of the multipass `m_line`, is *not* designated `multipass` and the result is the same as the previous example where `m_line` was defined to be `(A, A, B)`. That is, there will be three physical elements represented by three multipass lords.

Multipass lines do not have to be at the same "level" in terms of nesting of lines within lines. Additionally, multipass can be used with line reversal (§7.4). Example:

```
m_line: line[multipass] = (A, B)
m2_line: line = (m_line)
P: patch, ...
arc: line = (..., P)
u_line: line = (m_line, arc, --m2_line)
use, u_line
```

Here the tracking part of the lattice is

```
A\1, B\1, ..., B\2 (r), A\2 (r)
```

The "(r)" here just denotes that the element is reversed and is not part of the name. The lattice will have a multipass lord `A` that controls the two `A\n` elements and similarly with `B`. This lattice represents the case where a particle goes through the m_line in the "forward" direction, gets turned around in the `arc` line, and then passes back through `m_line` in the reverse direction. While it is possible to use reflection "−" (§7.2) instead of reversal "−−" (§7.4), reflection here does not make physical sense. Needed here is a reflection patch `P` (§4.41) between reversed and unreversed elements.

The procedure for how to group lattice elements into multipass slave groups which represent the same physical element is as follows. For any given element in the lattice, this element has some line it came from. Call this line L_0. The L_0 line in turn may have been contained in some other line L_1, etc. The chain of lines L_0, L_1, ..., L_n ends at some point and the last (top) line L_n will be one of the root lines listed in the `use` statement (§7.7) in the lattice file. For any given element in the lattice, starting with L_0 and proceeding upwards through the chain, let L_m be the *first* line in the chain that is marked as `multipass`. If no such line exists for a given element, that element will not be a multipass slave. For elements that have an associated L_m multipass line, all elements that have a common L_m line and have the same element index when L_m is expanded are put into a multipass slave group (for a given line the element index with respect to that line is 1 for the first element in the expanded line, the second element has index 2, etc.). For example, using the example above, the first element of the lattice, `A\1`, has the chain:

```
m_line, u_line
```

The last element in the lattice, (`A\2`), has the chain

```
m_line, m2_line, u_line
```

For both elements the L_m line is `m_line` and both elements are derived from the element with index 1 with respect to `m_line`. Therefore, the two elements will be slaved together.

As a final example, consider the case where a subline of a multipass line is also marked `multipass`:

```
a_line: line[multipass] = (A)
m_line: line[multipass] = (a_line, a_line, B)
u_line: line = (m_line, m_line)
use, u_line
```

In this case the tracking part of the lattice will be:

```
A\1, A\2, B\1, A\3, A\4, B\2
```

There will be two lord elements representing the two physically distinct elements `A` and `B`. The `A` lord element will will control the four `A\n` elements in the tracking part of the lattice. The `B` lord will control the two `B\n` elements in the tracking part of the lattice.

To simplify the constructed lattice, if the set of lattice elements to slave together only contains one element, a multipass lord is not constructed. For example:

```
m_line: line[multipass] = (A, A, B)
u_line: line = (m_line)
use, u_line
```

In this example no multipass lords are constructed and the lattice is simply

```
A, A, B
```

It is important to note that the global coordinates (§16.2) of the slaves of a given multipass lord are not constrained by *Bmad* to be the same. It is up to the lattice designer to make sure that the physical positions of the slaves makes sense (that is, are the same).

An example of an ERL lattice with multipass can be found in Section §13.3.

9.2 Reference Energy of Multipass Lords and Slaves

Consider the lattice where the tracking elements are

```
A\1, C, A\2
```

where `A\1` and `A\2` are multipass slaves of element `A` and `C` is an `lcavity` element with some finite voltage. In this case, the reference energy calculation (§5.5) where the reference energy of an element is inherited from the previous element, assigns differing reference energies to `A\1` and `A\2`. In such a situation, what should be the assigned reference energy for the multipass lord element `A`? *Bmad* calculates the lord reference energy in one of two ways. If, in the lattice file, `e_tot` or `p0c` is set for the multipass lord element, that setting will be used. Exception: For `em_field`, `lcavity`, and `custom` elements where the reference energy may change, set `e_tot_start` or `p0c_start` instead of `e_tot` or `p0c`. If the reference energy (or reference momentum) is not set in the lattice file, the reference energy of the lord is set equal to the reference energy of the first pass slave element.

It is important to keep this convention in mind if the normalized field strength (k1, for a quadrupole, etc.) for the lord element is set in the lattice file. To be physical, the unnormalized strength (the actual field) has to be the same for all slave elements. *Bmad* therefore calculates the unnormalized strength for the lord and sets the slave unnormalized strengths to be equal to the lord unnormalized strength. After this, the normalized strength for the slaves is calculated. Notice that the normalized strengths for the slaves will differ from each other. For `sbend` and `rbend` elements the calculation is a bit trickier. Here the `g` bending strength must be the same for all slaves since the setting of `g` determines the reference geometry. In this case, `dg` for each slave is adjusted accordingly so that the total normalized field, `g` + `dg`, gives the same unnormalized field for all slaves. Note that since the normalized field is calculated from the unnormalized field for the slaves, the setting of `field_master` (§5.2) is set to True for all the slaves independent of the setting of `field_master` for the lord.

To keep track of how the reference energy has been calculated for an element, *Bmad* sets an internal element switch called `multipass_ref_energy` which is set to "`user_set`" if the energy is explicitly set in the lattice file and is set to "`first_pass`" if the reference energy is calculated from the standard reference energy calculation of the first pass slave element.

Note: Historically, there was an element parameter `n_ref_pass` that could be set to control the reference energy. This parameter may be seen in old lattice files but will be ignored.

9.3 Phi0_multipass

As discussed above, `phi0_multipass` is an `lcavity` parameters that can be set individually for the different slaves of a given multipass lord. `phi0_multipass` is to calculate the change in reference energy through an `lcavity`. Additionally, `phi0_multipass` is used in relative time tracking (§26.1), but not absolute time tracking, to calculate the RF phase when tracking a particle. An example will make this clear:

```
LC: lcavity, phi0 = 0, L = 1.3, rf_frequency = 500e6, voltage = 10e6
linac: line[multipass] = (LC)
d: drift, l = ...
erl: line = (linac, d, linac)
use, erl
expand_lattice
LC\2[phi0_multipass] = 0.5
```

The expanded line for tracking will be

```
LC\1, d, LC\2
```

The first pass slave LC\1 is "accelerating" (since phi0 is zero) and the reference energy at the exit end of the element will be 10 MeV greater than the reference energy at the entrance end. Since LC\2 has phi0_multipass set to 0.5, LC\2 will be "decelerating" and the reference energy will drop by 10 MeV from entrance to exit end.

When tracking a particle with relative time tracking, phi0_multipass is used to calculate the phase of the RF field so a particle with phase space $z = 0$ will see an accelerating field in LC\1 and a decelerating field in LC\2 independent of the travel distance between the two slave cavities. This, of course, is "unphysical" and to "correctly" simulate the recirculating beam, absolute time tracking should be used. When absolute time tracking is used, phi0_multipass is ignored in tracking and, to get deceleration, the transit time from the entrance end of LC\1 to the entrance end of LC\2 should be adjusted to be $N \cdot t_{\rm rf}/2$ where N is an odd integer and $t_{\rm rf}$ is the RF period. Notice that this transit time includes the transit time through LC\1.

Chapter 10

Lattice File Global Parameters

This chapter deals with statements that can be used to set "global" parameter values. That is, parameter values that are associated with the lattice as a whole and not simply associated with a single element.

Discussed elsewhere are the global structures (§11.1) `bmad_com` and `ptc_com` (§11.4).

10.1 Parameter Statements

`Parameter` statements are used to set a number of global variables. If multiple branches are present (§2.2), these variables pertain to the `root` branch. The variables that can be set by `parameter` are

```
parameter[custom_attributeN]        = <string>    ! Defining custom attributes (§3.9).
parameter[default_tracking_species] = <Switch>    ! Default type of tracked particle.
                                                  !    Default is ref_particle.
parameter[e_tot]                    = <Real>      ! Reference total Energy.
                                                  !      Default: 1000 * rest_energy.
parameter[electric_dipole_moment]   = <Real>      ! Particle electric dipole moment.
parameter[geometry]                 = <Switch>    ! Open or closed
parameter[lattice]                  = <String>    ! Lattice name.
parameter[live_branch]              = <Logical>   ! Is branch fit for tracking?
parameter[machine]                  = <String>    ! Machine name.
parameter[n_part]                   = <Real>      ! Number of particles in a strong beam-beam bunch.
parameter[no_end_marker]            = <Logical>   ! Default: False.
parameter[parser_make_xfer_mats]    = <Logical>   ! Make element transfer maps during lattice parsing?
parameter[p0c]                      = <Real>      ! Reference momentum.
parameter[particle]                 = <speciesID> ! Reference species: positron, proton, etc.
parameter[photon_type]              = <Switch>    ! Incoherent or coherent photons?
parameter[ran_seed]                 = <Integer>   ! Random number generator init.
parameter[taylor_order]             = <Integer>   ! Default: 3
```

Examples

```
parameter[lattice]      = "L9A19C501.FD93S_4S_15KG"
parameter[geometry]     = closed
parameter[taylor_order] = 5
parameter[E_tot]        = 5.6e9    ! eV
```

parameter[custom_attributeN]
Here N is an integer between 1 and 40. For more information on defining custom attributes, see §3.9.

parameter[default_tracking_species]
The parameter[default_tracking_species] switch establishes the default type of particles to be tracked. See §3.11 for the syntax for naming particle species. In addition, this switch can be set to:

```
ref_particle     ! default
anti_ref_particle
```

By default, default_tracking_species is set to ref_particle so that the particle being tracked is the same as the reference particle set by param[particle]. In the case, for example, where there are particles going one way and antiparticles going the another, default_tracking_species can be used to switch between tracking the particles or antiparticles.

parameter[e_tot], parameter[p0c]
The parameter[e_tot] and parameter[p0c] are the reference total energy and momentum at the start of the lattice. Each element in a lattice has an individual reference e_tot and p0c attributes which are dependent parameters. The reference energy and momentum will only change between LCavity or Patch elements. The starting reference energy, if not set, will be set to 1000 time the particle rest energy. Note: beginning[e_tot] and beginning[p0c] (§10.4) are equivalent to parameter[e_tot] and parameter[p0c].

parameter[electric_dipole_moment]
The electric_dipole_moment sets the electric dipole moment value η for use when tracking with spin (§24.1).

parameter[geometry]
Valid geometry settings are

```
closed  ! Default w/o LCavity element present.
open    ! Default if LCavity elements present.
```

A machine with a closed geometry is something like a storage ring where the particle beam recirculates through the machine. A machine with an open geometry is something like a linac. In this case, if the reference particle is not a photon, the initial Twiss parameters need to be specified in the lattice file using the beginning statement (§10.4). If the geometry is not specified, closed is the default. The exception is that if there is an Lcavity element present or the reference particle is a photon, open will be the default.

Notice that by specifying a closed geometry it does *not* mean that the downstream end of the last element of the lattice has the same global coordinates (§16.2) as the global coordinates at the beginning. Setting the geometry to closed simply signals to a program to compute closed orbits and periodic Twiss parameters as opposed to calculating orbits and Twiss parameters based upon initial orbit and Twiss parameters at the beginning of the lattice. Indeed, it is sometimes convenient to treat lattices as closed even though there is no closure in the global coordinate sense. For example, when a machine has a number of repeating "periods", it may be convenient to only use one period in a simulation. Since *Bmad* ignores closure in the global coordinate sense, it is up to the lattice designer to ensure that a lattice is truly geometrically closed if that is desired.

parameter[lattice]
Used to set the lattice name. The lattice name is stored by *Bmad* for use by a program but it does not otherwise effect any *Bmad* routines.

parameter[live_branch]
Setting live_branch to False (default is True) indicates to a program that no tracking or other

analysis of the root branch should be done. This can be useful if the lattice has multiple branches and analysis of the root branch is not necessary. Other branches can also be marked as alive/dead using line parameter statements (§10.4). Note that the *Bmad* library itself ignores the setting of `live_branch` and it is up to the program being run to decide if this parameter is ignored or not. In particular, the *Tao* program (§1.2) *will* respect the setting of `live_branch`. Name of the machine the lattice simulates. Example: "LHC".

parameter[n_part]
The `parameter[n_part]` is the number of particle in a bunch. This parameter is used in a number of calculations, for example, with intrabeam scattering and to calculate the change in energy through an `Lcavity` (§4.31). Historically, this parameter has been used to set the number of strong beam particle with `BeamBeam` elements but it is strongly recommended to use the `beambeam` element's `n_particle` parameter instead.

parameter[no_end_marker]
Setting `parameter[no_end_marker]` to True will suppress the automatic inclusion of a marker named `END` at the end of the lattice (§7.1).

parameter[p0c]
See `parameter[e_tot]`.

parameter[parser_make_xfer_mats]
Default is True. If True, element-by-element transfer matrices will be constructed and put in the digested lattice file (§3.2) to save time next time the lattice is parsed. These matrices are constructed around the zero orbit. Occasionally, it is not possible to construct such matrices and this will generate error messages even though the analysis program being used will otherwise run normally. To suppress such messages, set this parameter to False.

parameter[particle]
The `parameter[particle]` switch sets the reference species. See §3.11 for the syntax for naming particle species.

The setting of the reference particle is used, for example, to determine the direction of the field in a magnet and given the normalized field strength (EG: `k1` for a quadrupole). Generally, the particles that by default are tracked through a lattice are the same as the reference particle. This default behavior can be altered by setting `parameter[default_tracking_species]`.

parameter[photon_type]
The `photon_type` switch is used to set the type of photons that are used in tracking. Possible settings are:

```
incoherent    ! Default
coherent
```

The general rule is use incoherent tracking except when there is a `diffraction_plate` element in the lattice.

parameter[ran_seed]
For more information on `parameter[ran_seed]` see §3.14.

parameter[taylor_order]
The Taylor order (§25.1) is set by `parameter[taylor_order]` and is the maximum order for a Taylor map.

10.2 Particle_Start Statements

particle_start statements are used, among other things to set the starting coordinates for particle tracking. If multiple branches are present (§2.2), these variables pertain to the root branch.

```
particle_start[x]                = <Real>   ! Horizontal position.
particle_start[px]               = <Real>   ! Horizontal momentum.
particle_start[y]                = <Real>   ! Vertical position.
particle_start[py]               = <Real>   ! Vertical momentum.
particle_start[z]                = <Real>   ! Longitudinal position.
particle_start[pz]               = <Real>   ! Momentum deviation. Only for non-photons.
particle_start[direction]        = +/-1     ! Longitudinal direction of travel.
particle_start[E_photon]         = <Real>   ! Energy (eV). Only used for photons.
particle_start[emittance_a]      = <Real>   ! A-mode emittance
particle_start[emittance_b]      = <Real>   ! B-mode emittance
particle_start[emittance_z]      = <Real>   ! Z-mode emittance
particle_start[sig_z]            = <Real>   ! Beam sigma in z-direction
particle_start[sig_pz]           = <Real>   ! Beam Sigma pz
particle_start[field_x]          = <Real>   ! Photon beam field along x-axis
particle_start[field_y]          = <Real>   ! Photon beam field along y-axis
particle_start[phase_x]          = <Real>   ! Photon beam phase along x-axis
particle_start[phase_y]          = <Real>   ! Photon beam phase along y-axis
particle_start[t]                = <Real>   ! Absolute time
particle_start[spin_x]           = <Real>   ! Spin polarization x-coordinate
particle_start[spin_y]           = <Real>   ! Spin polarization y-coordinate
particle_start[spin_z]           = <Real>   ! Spin polarization z-coordinate
```

Normally the absolute time, set by particle_start[t], is a dependent parameter set by solving Eq. (17.3) for t. The exception is when the initial velocity is zero. (This can happen if there is an e_gun (§4.15) element in the lattice). In this case, z must be zero and t is an independent parameter that can be set.

The longitudinal direction of travel is set by particle_start[direction]. This can be set to +1 (travel in the +s direction) or -1 for the reverse. +1 is the default. Generally particle_start[direction] should not be set to -1 since most programs will not be constructed to handle this situation. To track a particle in the reverse direction see §13.6.

For particles with spin, the spin can be specified using Cartesian coordinates with spin_x, spin_y, and spin_z.

For photons, px, py, and pz are the normalized velocity components (Cf. Eq. (17.12)). For photons pz is a dependent parameter which will be set so that Eq. (17.13) is obeyed.

Note: particle_start used to be called beam_start. Since this was confusing (beam initialization parameters are stored in a separate beam_init_struct structure (§12.1)), the name was changed. Currently the use of the beam_start name is deprecated but still supported for backwards compatibility.

Example

```
particle_start[y] = 2 * particle_start[x]
```

10.3 Beam Statement

NOTE: THE BEAM STATEMENT IS DEPRECATED. PLEASE DO NOT USE!

The beam statement is provided for compatibility with *MAD*. The syntax is

```
beam, energy = GeV, pc = GeV, particle = <Switch>, n_part = <Real>
```

For example

```
beam, energy = 5.6  ! Note: GeV to be compatible with MAD
beam, particle = electron, n_part = 1.6e10
```

Setting the reference energy using the `energy` attribute is the same as using `parameter[e_tot]`. Similarly, setting `pc` is equivalent to setting `parameter[p0c]`. Valid `particle` switches are the same as `parameter[particle]`.

10.4 Beginning and Line Parameter Statements

For non–circular lattices, the `beginning` statement can be used to set the Twiss parameters, beam energy, and other quantities at the beginning of the first lattice branch.

```
beginning[alpha_a]       = <Real>  ! "a" mode alpha
beginning[alpha_b]       = <Real>  ! "b" mode alpha
beginning[beta_a]        = <Real>  ! "a" mode beta
beginning[beta_b]        = <Real>  ! "b" mode beta
beginning[cmat_ij]       = <Real>  ! C coupling matrix. i, j = {``1'', or ``2''}
beginning[mode_flip]     = <logic> ! Set the mode flip status (§23.1). Default is False.
beginning[e_tot]         = <Real>  ! Reference total energy in eV.
beginning[eta_x]         = <Real>  ! x-axis dispersion
beginning[eta_y]         = <Real>  ! y-axis dispersion
beginning[etap_x]        = <Real>  ! x-axis momentum dispersion.
beginning[etap_y]        = <Real>  ! y-axis momentum dispersion.
beginning[deta_x_ds]     = <Real>  ! x-axis dispersion derivative.
beginning[deta_y_ds]     = <Real>  ! y-axis dispersion derivative.
beginning[p0c]           = <Real>  ! Reference momentum in eV.
beginning[phi_a]         = <Real>  ! "a" mode phase.
beginning[phi_b]         = <Real>  ! "b" mode phase.
beginning[ref_time]      = <Real>  ! Starting reference time.
beginning[s]             = <Real>  ! Longitudinal starting position.
beginning[spin_dn_dpz_x] = <Real>  ! Spin dn/dpz x-coordinate
beginning[spin_dn_dpz_y] = <Real>  ! Spin dn/dpz y-coordinate
beginning[spin_dn_dpz_z] = <Real>  ! Spin dn/dpz z-coordinate
beginning[dalpha_dpz_a]  = <Real>  ! dalpha_a/dpz derivative
beginning[dalpha_dpz_b]  = <Real>  ! dalpha_b/dpz derivative
beginning[dbeta_dpz_a]   = <Real>  ! dbeta_a/dpz derivative
beginning[dbeta_dpz_b]   = <Real>  ! dbeta_b/dpz derivative
beginning[deta_dpz_x]    = <Real>  ! deta_x/dpz derivative
beginning[deta_dpz_y]    = <Real>  ! deta_y/dpz derivative
beginning[detap_dpz_x]   = <Real>  ! detap_x/dpz derivative
beginning[detap_dpz_y]   = <Real>  ! detap_y/dpz derivative
```

The `gamma_a`, `gamma_b`, and `gamma_c` (the coupling gamma factor) will be kept consistent with the values set. If not set the default values are all zero. `beginning[e_tot]` and `parameter[e_tot]` are equivalent and one or the other may be set but not both. Similarly, `beginning[p0c]` and `parameter[p0c]` are equivalent.

Setting either momentum dispersion `etap_x` or `etap_y` also sets the corresponding dispersion derivative `deta_x_ds` or `deta_y_ds` (§23.4). If a momentum dispersion is set in the lattice file, or during program running the dispersion derivatives are "slaved" to the momentum dispersion. That is, if the reference

phase space momentum p_z changes, the momentum dispersions will be keept constant and the dispersion derivatives will be calculated from Eq. (23.38). Similarly, if a dispersion derivative is set in the lattice file or during program running, the momentum dispersions are slaved to the dispersion derivative. Which is slaved to which is determined by the last derivative set. If no derivatives are set, the dispersion derivatives are slaved to the momentum dispersions.

For any lattice the `beginning` statement can be used to set the starting floor position of the first lattice branch (see §16.2). The syntax is

```
beginning[x_position]     = <Real>  ! X position
beginning[y_position]     = <Real>  ! Y position
beginning[z_position]     = <Real>  ! Z position
beginning[theta_position] = <Real>  ! Angle on floor
beginning[phi_position]   = <Real>  ! Angle of attack
beginning[psi_position]   = <Real>  ! Roll angle
```

If the floor position is not specified, the default is to place beginning element at the origin with all angles set to zero.

The `beginning` statement is useful in situations where only parameters for the first branch need be specified. If this is not the case, the parameters for any branch can be specified using a statement of the form

```
line_name[parameter] = <Value>
```

This construct is called a `line parameter` statement Here `line_name` is the name of a line and `parameter` is the name of a parameter. The parameters that can be set here are the same parameters that can be set with the `beginning` statement with the additional parameters from the `parameter` statement:

```
default_tracking_species
geometry
live_branch
particle
inherit_from_fork
```

Example:

```
x_ray_fork: fork, to_line = x_ray
x_ray: line = (...)
x_ray[E_tot] = 100
```

The `inherit_from_fork` logical is used to determine if the reference energy and Twiss parameters as calculated from the `fork` element defining the branch is used to set the beginning element of the branch. This parameter is ignored if the fork element does not fork to the beginning element of the branch. The default is True. If any reference energy or momentum or any Twiss parameter is set, `inherit_from_fork` is implicitly set to False.

Rules:

1. The floor position of a line can only be set if the line is used for a `root branch`.
2. Line parameters statements must come after the associated line. This rule is similar to the rule that element attribute redefinitions must come after the definition of the element.

Chapter 11

Parameter Structures

11.1 What is a Structure?

A "structure" is a collection of parameters. *Bmad* has various structures which can be used for various tasks. For example, the `beam_init_struct` structure (§12.1) is used to set parameters used to initialize particle beams.

A given program may give the user access to some of these structures so, in order to allow intelligent parameter setting, this chapter gives an in-depth description of the most common ones.

Each structure has a **''structure name''** (also called a **''type name''**) which identifies the list of parameters (also called "components") in the structure. Associated with a structure there will be an **''instance''** of this structure and this instance will have an **''instance name''** which is what the user uses to set parameters. It is possible to have multiple instances of a structure. For example, in the situation where a program is simulating multiple particle beams, there could be multiple `beam_init_struct` (§12.1) instances with one for each beam.

Bmad defines uses some structures to hold global parameters. That is, parameters that shared by all code. Each of these structures has a single associated instance. These are:

Structure	Instance
bmad_common_stuct	bmad_com
space_charge_common_stuct	space_charge_com

All other structures will have instance names that are program specific. That is, see the program documentation for the instance name(s) used.

For historical reasons, There are two syntaxes used for setting structure components. The syntax when setting in a lattice file uses square brackets:

```
instance_name[parameter_name] = value
```

When setting a component in a program initialization file the syntax uses the percent "%" character:

```
instance_name%parameter_name = value
```

Examples:

```
bmad_com[max_aperture_limit] = 10   ! Lattice file syntax.
bmad_com%max_aperture_limit = 10    ! Program initialization file syntax.
```

this sets the `max_aperture_limit` parameter of `bmad_com` which is an instance name of the `bmad_common_struct`. Note: A program is free to set the instance name for any structure. This should be documented in the program manual.

Note: Thought must be given to setting `bmad_com` and other global parameters in a lattice file (§11.4) since that will affect every program that uses the lattice.

11.2 Bmad_Common_Struct

The `bmad_common_struct` structure contains a set of global parameters. There is only one global instance (§11) of this structure and this instance has the name `bmad_com`. The components of this structure along with the default values are:

```
type bmad_common_struct
  abs_tol_adaptive_tracking = 1e-10   ! Runge-Kutta tracking absolute tolerance.
  abs_tol_tracking = 1e-11            ! Closed orbit absolute tolerance.
  autoscale_amp_abs_tol = 0.1_rp      ! Autoscale absolute amplitude tolerance (eV).
  autoscale_amp_rel_tol = 1e-6        ! Autoscale relative amplitude tolerance
  autoscale_phase_tol = 1e-5          ! Autoscale phase tolerance.
  d_orb(6)           = 1e-5           ! for the make_mat6_tracking routine.
  default_ds_step    = 0.2            ! Integration step size.
  electric_dipole_moment = 0.0        ! Particle's EDM.
  fatal_ds_adaptive_tracking = 1e-8   ! Threshold for loosing particles.
  init_ds_adaptive_tracking = 1e-3    ! Initial step size.
  max_aperture_limit = 1e3            ! Max Aperture.
  min_ds_adaptive_tracking = 0        ! Minimum step size to use.
  ptc_cut_factor = 0.006              ! Cut factor for PTC tracking
  rel_tol_adaptive_tracking = 1e-8    ! Runge-Kutta tracking relative tolerance.
  rel_tol_tracking = 1e-8             ! Closed orbit relative tolerance.
  sad_amp_max = 5.0e-2                ! Used in sad_mult step length calc.
  sad_eps_scale = 5.0e-3              ! Used in sad_mult step length calc.
  significant_length = 1e-10          ! meter
  synch_rad_scale = 1.0               ! Synch radiation kick scale. 1 => normal

  default_integ_order = 2             ! PTC integration order
  max_num_runge_kutta_step = 10000    ! Max num RK steps before particle is lost.
  ptc_max_fringe_order = 2            ! PTC max fringe order (2 => Quadrupole !).
  sad_n_div_max = 1000                ! Used in sad_mult step length calc.
  taylor_order = 3                    ! 3rd order is default

  absolute_time_ref_shift = T         ! Absolute time referenced to element ref time?
  absolute_time_tracking = F          ! Set absolute or relative time tracking.
  aperture_limit_on = T               ! Use aperture limits in tracking.
  auto_bookkeeper = T                 ! No longer used.
  conserve_taylor_maps = T            ! Enable bookkeeper to set
                                      ! ele%taylor_map_includes_offsets = F?
  csr_and_space_charge_on = F         ! CSR and space charge (separate from HE SC).
  debug = F                           ! Used for code debugging.
  high_energy_space_charge_on = F     ! High energy space charge calc toggle.
  lr_wakes_on = T                     ! Long range wakefields
  normalize_twiss = F                 ! Normalize Twiss calc for off-energy particles?
```

```
    ptc_use_orientation_patches = T     ! Ele orientation translated to PTC patches?
    radiation_damping_on = F            ! Radiation damping toggle.
    radiation_fluctuations_on = F       ! Radiation fluctuations toggle.
    radiation_zero_average = F          ! Shift so that average radiation kick is zero?
    rf_phase_below_transition_ref = F   ! Autoscale around phase phi0 = 0.5
    spin_sokolov_ternov_flipping_on = F ! Spin flipping during radiation emission?
    spin_tracking_on = F                ! spin tracking?
    sr_wakes_on = T                     ! Short range wakefields?
  end type
```

Parameter description:

`abs_tol_adaptive_tracking`
: Absolute tolerance to use in adaptive tracking. This is used in `runge-kutta` and `time_runge_kutta` tracking (§6.4).

`abs_tol_tracking`
: Absolute tolerance to use in tracking. Specifically, Tolerance to use when finding the closed orbit.

`absolute_time_tracking`
: The `absolute_time_tracking` switch[1] sets whether the clock for the `lcavity` and `rfcavity` elements is tied to the reference particle or to uses the absolute time (§26.1). A value of `False` (the default) mandates relative time and a value of `True` mandates absolute time. The exception is that for an `e_gun` element (§4.15), absolute time tracking is always used in order to be able to avoid problems with a zero reference momentum at the beginning of the element.

`absolute_time_ref_shift`
: When absolute time tracking is used (§26.1), if `absolute_time_ref_shift` is True (the default), then the value of the time used to calculate RF phases and other time dependent parameters is shifted by the reference time of the lattice element under consideration. If set to False, no time shift is done. The advantage of `absolute_time_ref_shift` set to True is that (at least on the first turn of tracking) there is no phase shift between relative time and absolute time tracking. The advantage of `absolute_time_ref_shift` set to False is that when trying to compare tracking in *Bmad* with tracking in programs that use absolute time tracking but do not implement a reference shift (for example, the IMPACT and GPT programs), it is convenient not to have to worry about the reference shift.

  Note: If `absolute_time_ref_shift` is True, then the reference time shift for a multipass slave element, will be based upon the reference time of the first pass slave. This is done since each slave of a given multipass lord represents the same physical element.

`aperture_limit_on`
: Setting `aperture_limit_on` to `False` will disable the aperture limits set for individual elements (§5.8). The default is `True`.

`auto_bookkeeper`
: No longer used. Previously, toggled automatic or intelligent bookkeeping. Now intelligent bookkeeping is assumed. See section §33.6 for more details.

`autoscale_amp_abs_tol`
: Used when *Bmad* autoscales (§5.19) an elements field amplitude. This parameter sets the absolute tolerance for the autoscale amplitude parameter.

[1] An old, deprecated notation for this switch is `parameter[absolute_time_tracking]`.

`autoscale_amp_rel_tol`
: Used when *Bmad* autoscales (§5.19) an elements field amplitude. This parameter sets the relative tolerance for the autoscale amplitude parameter. Used when *Bmad* autoscales (§5.19) an elements AC phase. This parameter sets the absolute tolerance for the autoscale parameter.

`autoscale_phase_tol`

`init_ds_adaptive_tracking`
: Initial step to use for adaptive tracking. This is used in `runge-kutta` and `time_runge_kutta` tracking (§6.4).

`conserve_taylor_maps`
: Toggle to determine if the Taylor map for an element include any element "misalignments". See Section §6.8 for more details.

`csr_and_space_charge_on`
: Turn on or off the coherent synchrotron radiation and space charge calculations. (§21.4). The space charge calculation here is not to be confused with the high energy space charge calculation (§21.5)

`d_orb`
: Sets the orbit displacement used in the routine that calculates the transfer matrix through an element via tracking. That is, when the `mat6_calc_method` (§6.2) is set to `tracking`. `d_orb` needs to be large enough to avoid significant round-off errors but not so large that nonlinearities will affect the results. The default value is 10^{-5}.

`debug`
: Used for communication between program units for debugging purposes.

`default_ds_step`
: Step size for tracking code §6 that uses a fixed step size. For example, `symp_lie_ptc` tracking.

`default_integ_order`
: Order of the integrator used by Étienne Forest's PTC code (§30.2). The order of the PTC integrator is like the order of a Newton-Cotes method. Higher order means the error term involves a higher order derivative of the field.

`electric_dipole_moment`
: The electric dipole moment value used in tracking a particle's spin (§24.1).

`fatal_ds_adaptive_tracking`
: This is used in `runge-kutta` and `time_runge_kutta` tracking (§6.4). If the step size falls below the value set for `fatal_ds_adaptive_tracking`, a particle is considered lost. This prevents a program from "hanging" due to taking a large number of extremely small steps. The most common cause of small step size is an "unphysical" magnetic or electric field.

`high_energy_space_charge_on`
: Default is `False`. Toggle to turn on or off the ultra-relativistic space charge effect in particle tracking (§21.5). Unlike other space charge calcuflations, the high energy space charge can be done with single particle tracking (that is, tracking a beam of particle is not necessary). Computationally, this is separate from the lower energy space charge and CSR calculation (§21.4).

`lr_wakes_on`
: Toggle for turning on or off long-range higher order mode wakefield effects.

`max_aperture_limit`
: Sets the maximum amplitude a particle can have during tracking. If this amplitude is exceeded, the particle is lost even if there is no element aperture set. Having a maximum aperture limit helps prevent numerical overflow in the tracking calculations.

`max_num_runge_kutta_step`
: The maximum number of steps to take through an element with `runge_kutta` or `time_runge_kutta` tracking. The default value is 10,000. If the number of steps reaches this value, the particle being tracked is marked as lost and a warning message is issued. Under "normal" circumstances, a particle will take far fewer steps to track through an element. If a particle is not through an element after 10,000 steps, it generally indicates that there is a problem with how the field is defined. That is, the field does not obey Maxwell's Equations. Especially: discontinuities in the field can cause problems.

`min_ds_adaptive_tracking`
: This is used in `runge-kutta` and `time_runge_kutta` tracking (§6.4). Minimum step size to use for adaptive tracking. If To be useful, `min_ds_adaptive_tracking` must be set larger than the value of `fatal_ds_adaptive_tracking`. In this case, particles are never lost due to taking too small a step.

`normalize_twiss`
: Normalize transfer matrix when calculating Twiss parameters with off-energy particles? See §23.5.1 for details.

ptc_use_orientation_patches
: With *Bmad*, there is no distinction whether an element's orientation attributes (`offsets`, `pitches`, or `tilt` (§5.6)) is deliberate (part of the "design" of the machine) or an error (a "misalignment"). With PTC this is not true. If the `ptc_use_orientation_patches` switch is set to True (the default), when a *Bmad* element is translated to PTC, the element's orientation attributes are stored as patches. That is, "design" values. If set to False, these parameters are stored as misalignments. This will generally not make any difference to a calculation. The exception comes with PTC centric programs that vary machine parameters.[2]

ptc_max_fringe_order
: Maximum order for computing fringe field effects in PTC.

`rf_phase_below_transition_ref`
: Used when *Bmad* autoscales (§5.19) an `rfcavity` and when *Bmad* calculates the reference time through a cavity (which affects calculation of phase space z via Eq. (17.3)). If True, the reference phase will be taken to be at `phi0` = 0.5 which is appropriate for a ring below transition. Default is False in which case autoscaling will be around the phase `phi0` = 0.

`radiation_damping_on`
: Toggle to turn on or off effects due to radiation damping in particle tracking (§22.1). The default is `False`. Note: The standard *Bmad* emittance calculation, which involves calculating synchrotron radiation integrals (§22.3) can be done without a problem when `radiation_damping_on` is set to False. However, since the closed orbit will be affected by whether `radiation_damping_on` is set or not, the calculated emittances will depend upon the setting of `radiation_damping_on`.

`radiation_fluctuations_on`
: Toggle to turn on or off effects due to radiation fluctuations in particle tracking (§22.1). The default is `False`. Note: The standard *Bmad* emittance calculation, which involves calculating synchrotron

[2]None of the programs that come bundled with *Bmad* (a *Bmad* `Distribution`) are PTC centric.

radiation integrals (§22.3) can be done without a problem when `radiation_damping_on` is set to False. And since the calculation of the closed orbit ignores the fluctuating part of the radiation, the setting of `radiation_damping_on`, unlike the setting of `radiation_damping_on`, will not affect the emittance calculation.

`radiation_zero_average`
: As discussed in Section §22.1, it is sometimes convenient to shift the emitted radiation spectrum so that the average energy emitted along the closed orbit is zero. This gets rid of the "sawtooth" effect. To shift the average emitted energy to zero, set `radiation_zero_average` to `True`. The default is `False`. Currently, the shifting of the spectrum only works for non PTC dependent tracking. That is, the shifting is not applicable to tracking with Taylor maps and with `symp_lie_ptc` (§6.1) tracking.

`rel_tol_adaptive_tracking`
: Relative tolerance to use in adaptive tracking. This is used in `runge_kutta` and `time_runge_kutta` tracking (§6.4).

`rel_tol_tracking`
: Relative tolerance to use in tracking. Specifically, Tolerance to use when finding the closed orbit.

`significant_length`
: Sets the scale to decide if two length values are significantly different. For example, The superposition code will not create any super_slave elements that have a length less then this.

`sr_wakes_on`
: Toggle for turning on or off short-range higher order mode wakefield effects.

`spin_n0_direction_user_set`
: The closed orbit invariant spin $\mathbf{n}_0$ (§24.3) is degenerate such that if $\mathbf{n}_0(s)$ is a solution then so is $-\mathbf{n}_0(s)$. When calculating $\mathbf{n}_0(s)$, if `spin_n0_direction_user_set` is false (the default), *Bmad* will choose the solution which has the tune in the range $[0, \pi]$. This is sometimes inconvenient so if `spin_n0_direction_user_set` is set true, *Bmad* will choose $\mathbf{n}_0$ such that $\mathbf{n}_0$ at the beginning of the lattice branch is as nearly parallel to what is set using `particle_start` (§10.2). For branches with open geometry, `spin_n0_direction_user_set` is ignored and $\mathbf{n}_0$ at the beginning of the branch is set to the value given by `particle_start`.

`spin_sokolov_ternov_flipping_on`
: This determines if the Sokolov-Ternov effect is included in a simulation. The Sokolov-Ternov effect[Barber99] is the self-polarization of charged particle beams due to asymmetric flipping of a particle's spin when the particle is bent in a magnetic field. Also, spin flipping will *not* be done if spin tracking is off or both radiation damping and excitation are off.

`spin_tracking_on`
: Determines if spin tracking is performed or not.

`synch_rad_scale`
: This parameter is a multiplier for the kick given particles when radiation damping or excitation is turned on. This parameter is useful for artificially speeding up (or slowing down) the effect of radiation. The default value is one. Values greater than one will give larger kicks and will reduce the equilibrium settling time.

`taylor_order`
: Cutoff Taylor order of maps produced by `sym_lie_ptc`.

11.3 PTC_Common_Struct

The ptc_common_struct structure contains a set of global parameters that effect tracking when PTC is involved. There is only one global instance (§11) of this structure and this instance has the name ptc_com. The components of this structure along with the default values are:

```
type ptc_common_struct
  max_fringe_order  = 2    ! 2 => Quadrupole.
  complex_ptc_used = True ! Complex PTC code in use?
  use_totalpath   = False  ! phase space z = time instead of time - ref_time?
  old_integrator = True    ! PTC OLD_INTEGRATOR.
  exact_model     = True   ! PTC EXACT_MODEL.
  exact_misalign = True    ! PTC ALWAYS_EXACTMIS.
  translate_patch_drift_time = True
  pancake_symplectic = False ! Symplectic pancake integrator for gen_gradient elements?
  pancake_canonical  = False ! Canonical pancake tracking?
end type
```

Note: To set the Taylor map order for PTC, set the taylor_order parameter of bmad_com.

parameter[ptc_exact_model]
Deprecated. Replaced by ptc_com[exact_model] (§11.4).

The ptc_exact_model and ptc_exact_misalign switches affect tracking using the PTC library. See §6.4 for more details.

ptc_com[max_fringe_order]
When using PTC tracking (§1.4), the parameter[ptc_max_fringe_order] determines the maximum order of the calculated fringe fields. The default is 2 which means that fringe fields due to a quadrupolar field. These fields are 3^{rd} order in the transverse coordinates.

ptc_com[translate_patch_drift_time]
If translate_patch_drift_time is set True (the default) the patch in PTC that is setup to correspond to a *Bmad* patch is given a reference time offset equal to the *Bmad* reference time through the patch. This is generally what is wanted but for a PTC expert who knows what they are doing and really wants no time offset, translate_patch_drift_time can be set False.

ptc_com[pancake_symplectic]
When a gen_gradient element (§5.16.5) is tracked through PTC, the field map is translated into a set of PTC pancakes. If pancake_symplectic is set True, the symplectic pancake integrator is used (the pancakes carry the magnetic vector potential in addition to the field). The default is False. When pancake_symplectic is True the allowed integrator orders (§6.4) are 1, 2, 4, and 6; otherwise the allowed orders are 4 and 6.

ptc_com[pancake_canonical]
pancake_canonical selects canonical pancake tracking. This flag is only used when pancake_symplectic is False and is forced to True when pancake_symplectic is True. The default is False.

11.4 Bmad_Com

The parameters of the bmad_com instance of the bmad_common_struct structure (§11.2) can be set in the lattice file using the syntax

```
bmad_com[parm-name] = value
```

where parm-name is the name of a component of bmad_common_struct. For example:

```
bmad_com[rel_tol_tracking] = 1e-7
```

A similar situation holds for the ptc_com instance of the ptc_common_struct structure.

Be aware that setting either a bmad_com or ptc_com parameter value in a lattice file will affect all computations of a program even if the program reads in additional lattice files. That is, setting of bmad_com or ptc_com components is "sticky" and persists even when other lattice files are read in. There are two exceptions: A program is always free to override settings of these parameters. Additionally, a second lattice file can also override the setting made in a prior lattice file.

11.5 Space_Charge_Common_Struct

The space_charge_common_struct structure holds parameters for space charge (including CSR (§21.4)) calculations.[3] The setting of the csr_method and space_charge_method element parameters (§6.5) will also affect space charge calculations as well as the setting of the bmad_com logical csr_and_space_charge_on (§11.2).

Besides the parameters discussed below, the csr_and_space_charge_on parameter of bmad_com (§11.4) must be set True to enable the space charge/CSR calculations. Additionally, tracking with CSR will only be done through elements where the element parameter csr_method (§6.4) has been set to something other than off and tracking with space charge will only be done through elements where the element parameter space_charge_method is set to something other than off. This is done so that the computationally intensive space charge and CSR calculations can be restricted to places where the effects are significant.

The space charge / CSR parameter structure has a type name of space_charge_common_struct and an instance name of space_charge_com. This structure has components

```
type space_charge_common_struct
  ds_track_step = 0                   ! Tracking step size
  dt_track_step = 0                   ! Time based space charge step
  beam_chamber_height = 0             ! Used in shielding calculation.
  cathode_strength_cutoff = 0.01      ! Cutoff for the cathode field calc.
  rel_tol_tracking = 1e-8
  abs_tol_tracking = 1e-10
  lsc_sigma_cutoff = 0.1              ! Cutoff for the lsc calc. If a bin sigma
                                      !  is < cutoff * sigma_ave then ignore.
  particle_sigma_cutoff = -1          ! Veto particles that are far from the bench with 3D SC.
  n_bin = 0                           ! Number of bins used
  particle_bin_span = 2               ! Longitudinal particle length / dz_bin
  n_shield_images = 0                 ! Chamber wall shielding. 0 = no shielding.
  sc_min_in_bin = 10                  ! Min number of particle needed to compute sigmas.
  space_charge_mesh_size = [32,32,64] ! Mesh size with fft_3d space charge calc.
  csr_3d_mesh_size = [32,32,64]       ! Mesh size for 3D CSR calc.
  print_taylor_warning = T            ! Print Taylor element warning?
  diagnostic_output_file = ""         ! Wake file name
end type
```

The values for the various quantities shown above are their default values.

[3] This structure was formally called the csr_parameter_struct. The name was changed to reflect the fact that the structure has parameters for space charge calculations that do not involve CSR.

ds_track_step
: The `ds_track_step` parameter sets the nominal longitudinal distance traveled by the bunch between CSR kicks if the lattice element in which the bunch is being tracked has not set element's `csr_ds_track` parameter. The actual distance between kicks within a lattice element is adjusted so that there is an integer number of steps from the element entrance to the element exit. Either `ds_track_step` or the element's `csr_track_step` must be set to something positive otherwise an error will result when doing CSR or space charge tracking. Larger values will speed up the calculation at the expense of accuracy.

dt_track_step
: The `dt_track_step` parameter is used when the `tracking_method` of the lattice element the bunch is passing through is set to `time_runge_kutta` or `fixed_step_time_runge_kutta`.

beam_chamber_height
: `beam_chamber_height` is the height of the beam chamber in meters. This parameter is used when shielding is taken into account. See also the description of the parameter `n_shield_images`.

cathode_strength_cutoff
: When tracking through an element whose `space_charge_method` is set to `cathode_fft_3d` (§6.5, The value of `cathode_strength_cutoff` is used to determine at how far from the cathode the cathode image field is included. If the image field is less than `cathode_strength_cutoff` * bunch field, the image field will be ignored.

lsc_sigma_cutoff
: `lsc_sigma_cutoff` is used in the longitudinal space charge (LSC) calculation and is used to prevent bins with only a few particles in them to give a large contribution to the kick when the computed transverse sigmas are abnormally low.

n_bin
: `n_bin` is the number of bins used. The bind width is dynamically adjusted at each kick point so that the bins will span the bunch length. This parameter must be set to something positive. Larger values will slow the calculation while smaller values will lead to inaccuracies and loss of resolution. `n_bin` should also not be set so large that the average number of particles in a bin is too small. "Typical" values are in the range 100 — 1000.

particle_bin_span
: The `particle_bin_span` parameter is the width of a particle's triangular density distribution (cf. §21.4) in multiples of the bin width. A larger span will give better smoothing of the computed particle density with an attendant loss in resolution.

particle_sigma_cutoff
: The 3D space charge calculation uses a particle-in-cell algorithm. If there are halo particles far from the bunch center the grid spacing for the particle-in-cell may become too course. To help remedy this, particles far from the bunch center may be vetoed by setting `particle_sigma_cutoff` to a positive value. When set positive, particles will be ignored in the space charge calc when

$$\max\left(\frac{|dx|}{\sigma_x}, \frac{|dy|}{\sigma_y}, \frac{|dz|}{\sigma_z}\right) > \text{particle_sigma_cutoff} \tag{11.1}$$

: where dx, dy, and dz are the distances along the x, y, and z axis of the particle from the bunch centroid, and σ_x, σ_y, and σ_z are the bunch beam sizes.

`%space_charge_mesh_size`
: The `%space_charge_mesh_size` sets the size of the grid used when an element's `space_charge_method` is set to `fft_3d` (§6.5). The value of this parameter is a 3-element array (n_x, n_y, n_z) giving the mesh size in the x, y, and z directions respectively. Default values are $(32, 32, 64)$.

%csr3d_mesh_size
: The %csr3d_mesh_size sets the size of the grid used when an element's csr_method is set to steady_state_3d (§6.5). The value of this parameter is a 3-element array (n_x, n_y, n_z) giving the mesh size in the x, y, and z directions respectively. Default values are $(32, 32, 64)$.

n_shield_images
: n_shield_images is the number of shielding current layers used in the shielding calculation. A value of zero results in no shielding. See also the description of the parameter beam_chamber_height. The proper setting of this parameter depends upon how strong the shielding is. Larger values give better accuracy at the expense of computation speed. "Typical" values are in the range 0 — 5.

sc_min_in_bin
: the sc_min_in_bin parameter sets the minimum number of particle in a bin needed to compute the transverse beam sigmas for that bin. If the number of particles is less than this number, the beam sigmas are taken to be equal to the beam sigmas of a nearby bin where there are enough particle to compute the sigma. The beam sigmas are needed for the CS calculation but not need for the CSR calculation.

diagnostic_output_file
: If set non blank, an output file of this name is created that contains a table of the CSR wake at each track step (the track step size is set by ds_track_step). If tracking is done through multiple lattice elements, the wake tables for the elements are appended to the file. This file is useful for visualization of the wake.

Note: Taylor map elements (§4.52) that have a finite length cannot be subdivided for the CSR calculation. *Bmad* will ignore any taylor elements present in the lattice but will print a warning that it is doing so. To suppress the warning, print_taylor_warning should be set to False.

11.6 Opti_DE_Param_Struct

The Differential Evolution (DE) optimizer is used in nonlinear optimization problems. This optimizer is based upon the work of Storn and Price[Storn96]. There are a number of parameters that can be varied to vary how the optimizer works. These parameters are are contained in a structure named opti_de_param_struct. the instance name is opti_de_param. This structure has components

```
                              Default
  type opti_de_param_struct
    CR             = 0.8    ! Crossover Probability.
    F              = 0.8    !
    l_best         = 0.0    ! Percentage of best solution used.
    binomial_cross = False  ! IE: Default = Exponential.
    use_2nd_diff   = False  ! use F * (x_4 - x_5) term
    randomize_F    = False  !
    minimize_merit = True   ! F => maximize the Merit func.
  end type
```

The "perturbed vector" is v = x_1 + l_best * (x_best - x_1) + F * (x_2 - x_3) + F * (x_4 - x_5) The last term F * (x_4 - x_5) is only used if use_2nd_diff = T.

The crossover can be either "Exponential" or "Binary". Exponential crossover is what is described in the paper. With Exponential crossover the crossover parameters from a contiguous block and the average number of crossover parameters is approximately average crossovers ~ min(D, CR / (1 - CR)) where D

is the total number of parameters. With Binary crossover the probability of crossover of a parameter is uncorrelated with the probability of crossover of any other parameter and the average number of crossovers is average crossovers = D * CR

`randomize_F` = True means that the F that is used for a given generation is randomly chosen to be within the range [0, 2*F] with average F.

11.7 Dynamic Aperture Simulations: Aperture_Param_Struct

The `dynamic_aperture_struct` is used for dynamic aperture calculations. This structure has components:

```
type aperture_param_struct
  min_angle = 0
  max_angle = pi
  n_angle   = 9
  n_turn = 100                ! Number of turns a particle must survive
  x_init = 1e-3              ! Initial estimate for horizontal aperture
  y_init = 1e-3              ! Initial estimate for vertical aperture
  accuracy = 1e-5            ! Resolution of bracketed aperture.
end type
```

Chapter 12

Beam Initialization

Some *Bmad* based programs track beams of particles instead of tracking individual particles one-by-one. This can be useful for several reasons. For example, tracking beams is useful when inter-bunch or intra-bunch effects are to be simulated. Also tracking beams can simplify the bookkeeping a program needs to do to calculate such quantities such as the bunch size.

A *Bmad* based program has two standard ways to specify the initial distribution of a beam. One is using a `beam_init_struct structure` (§11.1) which holds parameters (for example, the beam emittances) from which a distribution of particles can be constructed. The `beam_init_struct` structure is explained in Section §12.1. The other way is to specify the initial beam distribution via a file that has the individual particle positions. This is covered in Section §12.2.

12.1 Beam_Init_Struct Structure

The `beam_init_struct structure` (§11.1) holds parameters which are used to initialize a beam. The parameters of this structure, shown with their default values, are:

```
type beam_init_struct
  character(200) :: position_file = ""       ! Initialization file name.
  character distribution_type(3)             ! "ELLIPSE", "KV", "GRID", "" (default).
  type (ellipse_beam_init_struct) ellipse(3) ! For ellipse beam distribution
  type (kv_beam_init_struct) KV              ! For KV beam distribution
  type (grid_beam_init_struct) grid(3)       ! For grid beam distribution
  logical use_particle_start = F     ! Use particle_start instead of %center and %spin?
  character random_engine            ! "pseudo" (default) or "quasi".
  character random_gauss_converter   ! "exact" (default) or "quick".
  real center(6) = 0                 ! Bench phase space center offset.
  real t_offset = 0                  ! Time offset.
  real center_jitter(6) = 0          ! Bunch center rms jitter
  real emit_jitter(2)   = 0          ! %RMS a and b-mode emittance jitter
  real sig_z_jitter     = 0          ! bunch length RMS jitter
  real sig_pz_jitter     = 0         ! pz energy spread RMS jitter
  real random_sigma_cutoff = -1      ! -1 => no cutoff used.
  integer n_particle = 0             ! Number of simulated particles per bunch.
  logical renorm_center = T          ! Renormalize centroid?
  logical renorm_sigma = T           ! Renormalize sigma?
```

```
    real spin(3) = 0, 0, 0            ! Spin (x, y, z)
    real a_norm_emit = 0              ! a-mode normalized emittance (= β γ ε_unnorm)
    real b_norm_emit = 0              ! b-mode normalized emittance (= β γ ε_unnorm)
    real a_emit = 0                   ! a-mode unnormalized emittance (= ε_norm/βγ)
    real b_emit = 0                   ! b-mode unnormalized emittance (= ε_norm/βγ)
    real dpz_dz = 0                   ! Correlation of pz with longitudinal position.
    real dt_bunch = 0                 ! Time between bunches.
    real sig_z = 0                    ! Z sigma in m.
    real sig_pz = 0                   ! pz sigma.
    real bunch_charge = 0             ! Charge in a bunch.
    integer n_bunch = 0               ! Number of bunches.
    integer ix_turn = 0               ! Turn index.
    character species = ""            ! Species. Default is reference particle.
    logical full_6D_coupling_calc = F  ! Use 6x6 1-turn matrix to match distribution?
    logical use_t_coords = F  ! If true, the distributions will be
                              !   calculated using time coordinates (§17.3).
    logical use_z_as_t   = F  ! Only used if use_t_coords = T:
                              !   If True,  the z coordinate stores the time.
                              !   If False, the z coordinate stores the s-position.
  end type
```

Note: The `z` coordinate value given to particles of a bunch is with respect to the nominal center of the bunch. Therefore, if there are multiple bunches, and there is an RF cavity whose frequency is not commensurate with the spacing between bunches, absolute time tracking (§26.1) must be used.

%position_file
: `%position_file` sets the name of the file to be read in containing the particle coordinates. Input from a file is triggered if not-blank. The format of the file is discussed in Section §12.2.

%a_emit, %b_emit, %a_norm_emit, %b_norm_emit
: Normalized and unnormalized emittances. Either `a_norm_emit` or `a_emit` may be set but not both. similarly, either `b_norm_emit` or `b_emit` may be set but not both.

  When simulating a ring, if any of these parameters is set negative, and if the *Bmad* based program being run has enabled it, the value of the parameter will be set the value as calculated from the lattice using synchrotron radiation integral formulas (§22.3).

%bunch_charge
: The `%bunch_charge` paramter sets the charge of a bunch. If reading from a file, the bunch charge will be set to the value of `%bunch_charge` except if `%bunch_charge` has a value of zero in which case the bunch charge as specified in the file is used.

%center(6)
: The `%center` parameter is used to offset the center position of a bunch. Exception: When `%use_particle_start` is set to True, the `particle_start` values are used instead for the center position offset. See the description of `%use_particle_start` below for more details.

%center_jitter, %emit_jitter, %sig_z_jitter, %sig_pz_jitter
: These components can be used to provide a bunch-to-bunch random variation in the emittance and bunch center.

%distribution_type(3)
: The `%distribution_type(:)` array determines what algorithms are used to generate the particle

distribution for a bunch. Note: If %position_file is not blank, the beam distribution will be read from the appropriate file and %distribution_type will be ignored.

%distribution_type(1) sets the distribution type for the (x, p_x) 2D phase space, etc. Possibilities for %distribution_type(:) are:

```
"", or "RAN_GAUSS"  ! Random distribution (default).
"ELLIPSE"           ! Ellipse distribution (§21.1.1)
"KV"                ! Kapchinsky-Vladimirsky distribution (§21.1.2)
"GRID"              ! Uniform distribution.
```

Since the Kapchinsky-Vladimirsky distribution is for a 4D phase space, if the Kapchinsky-Vladimirsky distribution is used, "KV" must appear exactly twice in the %distribution_type(:) array.

Unlike all other distribution types, the GRID distribution is independent of the Twiss parameters at the point of generation. For the non-GRID distributions, the distributions are adjusted if there is local x-y coupling (§23.1). For lattices with a closed geometry, if full_6D_coupling_calc is set to True, the full 6-dimensional coupling matrix is used. If False, which is the default, The 4-dimensional **V** matrix of Eq. (23.5) is used.

Note: The total number particles generated is the product of the individual distributions. For example:

```
type (beam_init_struct) bi
bi%distribution_type = ELLIPSE", "ELLIPSE", "GRID"
bi%ellipse(1)%n_ellipse = 4
bi%ellipse(1)%part_per_ellipse = 8
bi%ellipse(2)%n_ellipse = 3
bi%ellipse(2)%part_per_ellipse = 100
bi%grid(3)%n_x = 20
bi%grid(3)%n_px = 30
```

The total number of particles per bunch will be $32 \times 300 \times 600$. The exception is that when RAN_GAUSS is mixed with other distributions, the random distribution is overlaid with the other distributions instead of multiplying. For example:

```
type (beam_init_struct) bi
bi%distribution_type = RAN_GAUSS", "ELLIPSE", "GRID"
bi%ellipse(2)%n_ellipse = 3
bi%ellipse(2)%part_per_ellipse = 100
bi%grid(3)%n_x = 20
bi%grid(3)%n_px = 30
```

Here the number of particle is 300×600. Notice that when RAN_GAUSS is mixed with other distributions, the value of beam_init%n_particle is ignored.

%dPz_dz
Correlation between p_z and z phase space coordinates.

%dt_bunch
Time between bunches

%ellipse(3)
The %ellipse(:) array sets the parameters for the ellipse distribution (§21.1.1). Each component of this array looks like

```
type ellipse_beam_init_struct
  integer part_per_ellipse  ! number of particles per ellipse.
  integer n_ellipse         ! number of ellipses.
  real sigma_cutoff         ! sigma cutoff of the representation.
end type
```

%full_6D_coupling_calc

If set `True`, coupling between the transverse and longitudinal modes is taken into account when calculating the beam distribution. The default `False` decouples the transverse and longitudinal calculations.

%grid(3)

The `%grid` component of the `beam_init_struct` sets the parameters for a uniformly spaced grid of particles. The components of `%grid` are:

```
type grid_beam_init_struct
  integer n_x        ! number of columns.
  integer n_px       ! number of rows.
  real x_min         ! Lower x limit.
  real x_max         ! Upper x limit.
  real px_min        ! Lower px limit.
  real px_max        ! Upper px limit.
end type
```

%ix_turn

Turn index. This affects how particle time is calculated. Particle time is calculated from phase space z and t_0 via Eq. (17.3). t_0 is the reference time at the lattice element that the beam is initialized at (§17.1). For simulations where the beam is circulating in a ring over many turns, it may be desired to initialize the beam appropriate for some turn after the first turn. That is, set the particle time with the reference time t_{0n} associated with the n^{th} (set by `ix_turn`) turn at which the beam is being initialized at

$$t_{0n} = t_0 + n\, t_{rev} \tag{12.1}$$

where t_{rev} is the revolution time.

%KV

The `%kv` component of the `beam_init_struct` sets the parameters for the Kapchinsky-Vladimirsky distribution (§21.1.2). The components of `%KV` are:

```
type kv_beam_init_struct
  integer part_per_phi(2)    ! number of particles per angle variable.
  integer n_I2               ! number of I2
  real A                     ! A = I1/e
end type
```

%n_bunch

The number of bunches in the beam is set by `n_bunch`. If reading the distribution from a file, if `%n_bunch` is zero, the number of bunches created is the number of defined in the file and if `%n_bunch` is not zero, the number created is `%n_bunch`. It is an error if `%n_bunch` is greater than the number of bunches defined in the file. If not reading from a file, if `%n_bunch` is zero, one bunch is created.

%n_particle

Number of particles generated when the `%distribution_type` is `"RAN_GAUSS"`. Ignored for other distribution types. When reading the distribution from a file, if `%n_particle` is zero, the number of particles in a bunch will be the number of particles defined in the file. If `%n_particle` is non-zero when reading from a file, the number of particles in a bunch will be `%n_particle`. It is an error if `%n_particle` is non-zero and the number of particles defined in the file is less than `%n_particle`.

%random_engine

This component sets the algorithm to use in generating a uniform distribution of random numbers in the interval [0, 1]. `"pseudo"` is a pseudo random number generator and "quasi" is a quasi random generator. "quasi random" is a misnomer in that the distribution generated is fairly uniform.

%random_gauss_converter, %random_sigma_cutoff
To generate Gaussian random numbers, a conversion algorithm from the flat distribution generated according to `%random_engine` is needed. `%random_gauss_converter` selects the algorithm. The `"exact"` conversion uses an exact conversion. The `"quick"` method is somewhat faster than the `"exact"` method but not as accurate. With either conversion method, if `%random_sigma_cutoff` is set to a positive number, this limits the maximum sigma generated.

%renorm_center, %renorm_sigma
If set to True, these components will ensure that the actual beam center and sigmas will correspond to the input values. Otherwise, there will be fluctuations due to the finite number of particles generated.

%sig_pz, %sig_z
Longitudinal sigmas. `%sig_pz` is the fractional energy spread dE/E. This, along with `%dPz_dz` determine the longitudinal profile.

When simulating a ring, if any of these parameters is set negative, and if the *Bmad* based program being run has enabled it, the value of the parameter will be set the value as calculated from the lattice using synchrotron radiation integral formulas (§22.3).

%species
Name of the species tracked. If not set then the default tracking particle type is used.

%spin
Particle spin in Cartesian (x, y, z) coordinates. Only used when not reading in particle positions from a file. Also: When `%use_particle_start` is set to True, and not reading in particle positions form a file, the `particle_start` spin values are used instead. See the description of `%use_particle_start` below for more details.

%t_offset
Time offset. Always applied independent of the setting of `use_particle_start`.

%use_particle_start
If `%use_particle_start` is set to `False` (the default) the center of the bunch is determined by the setting of the `%center` component. If set to `True`, the center is determined by the setting of:

```
particle_start[x], particle_start[px],
particle_start[y], particle_start[py],
particle_start[z], particle_start[pz]
```

`%use_particle_start = True` will also affect the setting of the spin. If set to `True`, the spin is determined by:

```
particle_start[spin_x], particle_start[spin_y], particle_start[spin_z]
```

Additionally, if `%use_particle_start = True`, the setting of `beam_init%t_offset` will be ignored in favor of `particle_start[t]` set in the lattice file.

See §10.2 for details about the `particle_start` structure. Using the `particle_start` structure allows setting center and spin values in the lattice file rather than the setting of `%center` and `%spin` in the `beam_init_struct`. In this case, `%center` is a dependent parameter and will be set to the value of `particle_start`.

%use_t_coords, %use_z_as_t
The problem with handling particle distributions at low energies is that when a particle's velocity is zero, the phase space z-coordinate is zero. What is needed here is to generate distributions in time and or in longitudinal s-position. To do this, two switches, `use_to_coords` and `use_z_as_t`, are used.

If `use_t_coords` is True (default is False), the values in the distributions are taken as describing particles using "time coordinates" (§17.3).

The `use_z_as_t` parameter is only used if `use_t_coords` is set to True. When `use_t_coords` is True, if `use_z_as_t` is also True, the value of the time (in seconds) assigned to a particle will be set equal to the `z`-coordinate (in meters) and a new `z`-coordinate will be calculated based upon Eq. (17.3). This is useful for describing the situation, for example, where particles may originate at a cathode at the same s-position, but different times. If `use_z_as_t` is False (the default), the z coordinate is taken to describe the s-coordinate. This is useful for modeling particles that have the same time but different s positions. In this case, particles are "born" inside the lattice element at the location of the particle's s-coordinate. To properly track the particles, the bunch will need to be tracked through the element where they are born with a tracking method that can handle inside particles. Currently the only such tracking method is `time_runge_kutta`. When the particles get to the end of the element, the particle positions are converted to standard s-coordinates.

Since HDF5 beam files store complete information about a particle, `use_to_coords` and `use_z_as_t` are not needed and are ignored.

12.2 File Based Beam Initialization

A beam initialization file specifies the coordinates of all the particles in a beam. If a *Bmad* based program uses a `beam_init_struct` (§12.2) for inputting initialization parameters, the file name for file based beam initialization can be set using the `%position_file` component of the structure. Also the bunch charge, bunch number, and number of particles per bunch can be set in the `beam_init_struct`. Additionally, the bunch centroid can be offset by setting `beam_init%center` and `beam_init%center_jitter`.

There are two formats for the beam initialization file: ASCII and binary. The binary file format for beam position storage is based on the `HDF5` standard. More information is in chapter §43.

The new ASCII format describes a particle bunch with a `header` section followed by a `table` of particle parameters. Multiple bunches that comprise a beam can be specified by multiple header section / particle parameter table pairs, one pair for each bunch. An example header section:

```
# The header field lines all start with a pound "#" sign.
# Any line in the header field that does not start with a recognized parameter or
# does not have an equal sign is ignored.
# my_param = 1.23    ! Custom parameters can be defined and will be ignored by Bmad.
# species = proton   ! This parameter will be read by Bmad.
# spin = 1, 0, 0     !   and this one too.
# The last line in the header field starts with "#!" and defines the table columns.
#!   index     x         px       y     py   ... etc...
```

The header section lines all start with a pound "#" sign. The last line in the header section must start with "#!" and this line defines the particle parameter table columns. With the exception of the last line, all header lines will be ignored except ones that begin with a recognized parameter followed by an equal sign. Recognized parameters are the components of the `coord_struct`, as documented in §37.1, that describes individual particles. In addition, the following parameters are recognized:

```
  charge_tot      ! Total bunch charge (including dead particles).
  s_position      ! Longitudinal position of bunch. Can be used in place of "s".
  time            ! Time particles of bunch are at. Can be used in place of "t".
```

These parameters, if present, will be used to set the corresponding particle parameter in all the particles of the corresponding bunch. Possible `location` and `state` parameter settings are documented in §37.1. The string equivalent to any setting is obtained by removing the trailing "$" from the variable name.

For example, the variable `alive$` which is a possible `state` setting becomes the string `alive`. Note: `charge_tot` (total bunch charge) will superceed `charge` (the charge per particle).

The particle parameter table follows the header section. Each row gives the parameters for one particle. Not all particle parameters must be specified. If a particular parameter is not present, its default value will be used or, if present, the value given in the header section. Column order is irrelavent and what determines what particle parameter is associated with a given column is the last line of the header section which starts with the characters `#!`. Except for vector parameters, column names correspond to `coord_struct` parameters the same as in the header section. For parameters that are vectors, the mapping from column name to parameter is:

```
Column Name             Corresponding coord_struct components
----------------------  -------------------------------------
x, px, y, py, z, pz     vec(1), ..., vec(6)
spin_x, spin_y, spin_z  spin(1), spin(2), spin(3)
field_x, field_y        field(1), field(2)
phase_x, phase_y        phase(1), phase(2)
```

Additionally, there can be (but is not required) an `index` column giving the particle index in the array of particles of a bunch. This column is ignored so, for example, values in the first line of the table will always be used to set the particle with index 1 independent of the value given in the index column. This behavior is implemented so that a beam file can be edited to add or remove particles without worrying about reindexing.

Note: Example beam files (both ASCII and HDF5) can be generated using the program in the Bmad Distribution `code_examples/beam_track_example`.

Note: Beam files can be converted between ASCII and HDF5 binary using the program in the Bmad Distribution `util_programs/beam_file_translate_format`.

12.2.1 Old Beam ASCII Format

There is an old ASCII format that is still accepted by *Bmad* but is deprecated and should be avoided. The format is:

```
<ix_ele>         ! Lattice element index. This is ignored.
<n_bunch>        ! Number of bunches.
<n_particle>     ! Number of particles per bunch to use
[bunch loop: ib = 1 to n_bunch]
  BEGIN_BUNCH    ! Marker to mark the beginning of a bunch specification block.
  <species_name> ! Species of particle
  <charge_tot>   ! Total charge of bunch (alive + dead). 0 => Use <macro_charge>.
  <z_center>     ! z position at center of bunch.
  <t_center>     ! t position at center of bunch.
  [particle loop: Stop when END_BUNCH marker found]
    <x> <px> <y> <py> <z> <pz> <macro_charge> <state> <spin_x> <spin_y> <spin_z>
  [end particle loop]
  END_BUNCH      ! Marker to mark the end of the bunch specification block
[end bunch loop]
```

Example:

```
0       ! ix_ele
1       ! n_bunch
25000   ! n_particle
BEGIN_BUNCH
```

```
  POSITRON
  3.2E-9   ! charge_tot
  0.0      ! z_center
  0.0      ! t_center
 -6.5E-3  9.6E-3 -1.9E-2  8.8E-3  2.2E-2 -2.4E-2  1.2E-13  1 1.0 0.0 0.0
  8.5E-3  5.5E-3  4.0E-2 -1.9E-2 -4.9E-3  2.1E-2  1.2E-13  1 1.0 0.0 0.0
  1.1E-2 -1.9E-2 -2.5E-2  1.0E-2 -1.8E-2 -7.1E-3  1.2E-13  1 1.0 0.0 0.0
 -3.4E-2 -2.7E-3 -4.1E-3  1.3E-2  1.3E-2  1.0E-2  1.2E-13  1 1.0 0.0 0.0
  6.8E-3 -4.5E-3  2.5E-3  1.4E-2 -2.3E-3  7.3E-2  1.2E-13  1 1.0 0.0 0.0
  1.2E-2 -9.8E-3  1.7E-3  6.4E-3 -9.8E-3 -7.2E-2  1.2E-13  1 1.0 0.0 0.0
  1.1E-2 -3.5E-4  1.2E-2  1.8E-2  5.4E-3  1.4E-2  1.2E-13  1 1.0 0.0 0.0
     ... etc. ...
END_BUNCH
```

The first line of the file gives <ix_ele>, the index of the lattice element at which the distribution was created. This is ignored when the file is Read.

The second line gives <n_bunch>, the number of bunches. This can be overridden by a non-zero setting of beam_init%n_bunch.

The third line gives <n_particle> the number of particles in a bunch. The actual number rows specifying particle coordinates may be more then <n_particle>. In this case, particles will be discarded so that the beam has <n_particle> particles. The setting of <n_particle> can be overridden by a non-zero setting of beam_init%n_particle.

After this, there are <n_bunch> blocks of data, one for each bunch. Each one of these blocks starts with a BEGIN_BUNCH line to mark the beginning of the block and ends with a END_BUNCH marker line. In between, the first four lines give the <species_name> name, <charge_tot>, <z_center>, and <t_center> values. The <species_name> name may be one of:

```
positron  ! default
electron
proton
antiproton
muon
antimuon
photon
```

The lines following the <t_center> line specify particle coordinates. One line for each particle. Only the first six numbers, which are the phase space coordinates, need to be specified for each particle. If the <macro_charge> column is not present, or is zero, it defaults to <charge_tot>/<n_particle>.

The <state> parameter indicates whether a particle is alive or dead. Values are

```
1     ! Alive
2-7   ! Dead
8     ! Pre-born
```

The pre-born state indicates that the particle is waiting to be emitted from the cathode of an electron gun (§21.4.3).

The particle spin is specified by x, y and z components.

Each particle has an associated <macro_charge>. If <charge_tot> is set to a non-zero value, the charge of all the particles will be scaled by a factor to make the total macro charge equal to <charge_tot>. The macro charge is ignored in tracking. The charge of the particle used in tracking is the charge as calculated for the particle species. On the other hand, the macro charge is used to calculate such things as the total charge in a particular region or the field produced by a particle. That is, the macro charge

acts as a weighting factor for a particle when the particle's field or the particle's effect on other particles is calculated.

When the particle coordinates are read in the centroid will be shifted by the setting of `beam_init%center` (unless `beam_init%use_particle_start` is set True) and `beam_init%center_jitter`.

Chapter 13

Lattice Examples

This chapter gives some examples of how lattice files can be constructed to describe various machine geometries.

13.1 Example: Injection Line

An injection line is illustrated in Fig. 13.1. In this example, The path of an injected particle after it leaves the last element `X` of the injection line (dashed blue line) partially goes through the field of the dipole `BND` in the storage ring. One way to simulate this is:

```
INJ_L: line = (..., X, P, BND2, BR)
RING_L: line = (..., BND, M, ...)
P: patch, x_offset = -0.5, x_pitch = 0.15, z_offset = 0.3
BND: sbend, l = 6.2, g = 1/52
BND2: BND, l = 4.7, tracking_method = runge_kutta,
        field_calc = fieldmap, grid_field = {...}
BR: fork, to_line = RING_L, to_element = M
M: marker
use, INJ_L
```

In order to properly track particles through the fringe field of the dipole `BND`, a partial section of `BND`, called `BND2`, is placed in the injection line `INJ_L`. The `tracking_method` for `BND2` is set to `runge_kutta` since the default `bmad_standard` tracking is not able to handle these fringe fields. Additionally, the

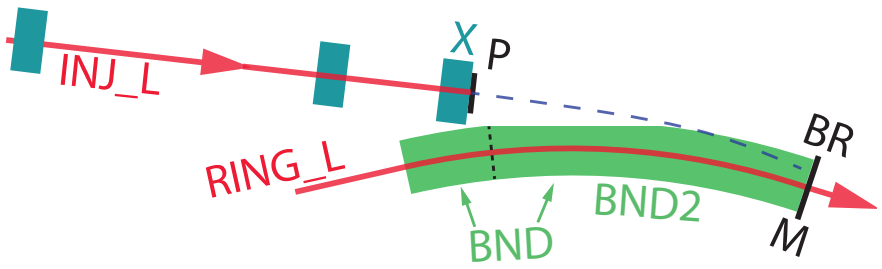

Figure 13.1: Injection line into a dipole magnet.

field_calc parameter of BND2 is set to field map so that the actual field profile of this particular magnet can be used in the tracking. The field is specified in the grid_field parameter (§5.16).

After traversing element X in the injection line, the particle goes through the patch P which offsets the reference trajectory so that the following element, BND2, is properly positioned. The beginning of BND2 is marked by a black dashed line in the figure. At the end of BND2 the fork element BR connects INJ_L with the marker M in RING_L.

13.2 Example: Chicane

A chicane is a series of bending magnets that shifts the beam trajectory in the transverse plane while keeping the starting and ending trajectories constant. Chicanes are useful for doing things like bunch compression or as a variable bunch time delay. An example four bend chicane is illustrated in Fig. 13.2A. The lattice file for this is:

```
parameter[geometry] = open
beginning[beta_a] = 20
beginning[beta_b] = 20
parameter[p0c] = 1e7

bnd1: rbend, l = 2
bnd2: rbend, l = 2
d1: drift, l = 3
d2: drift, l = 2

chicane: overlay = bnd1[dg]:-ge, bnd2[dg]:ge, var = ge, ge = 2e-3

c_line: line = (bnd1, d1, bnd2, d2, bnd2, d1, bnd1)
use, c_line
```

The chicane is controlled by an overlay (§4.40) called chicane. This overlay controls the dg attribute of the bends (alternatively the hkick or b0 attributes could have been used).

It is important to note that using g instead of dg to control the chicane strength would be a mistake. [Or, alternatively, b_field instead of db_field if controlling the unnormalized field.] This is illustrated in Fig. 13.2B where the chicane overlay was replaced by:

```
chicane: overlay = bnd1[g]:-ge, bnd2[g]:ge, var = ge, ge = 1e-1
```

The problem here is that g sets the reference orbit so varying g will vary the physical layout of the machine which is not what is wanted. Since the actual normalized field is g + dg, it is dg that should

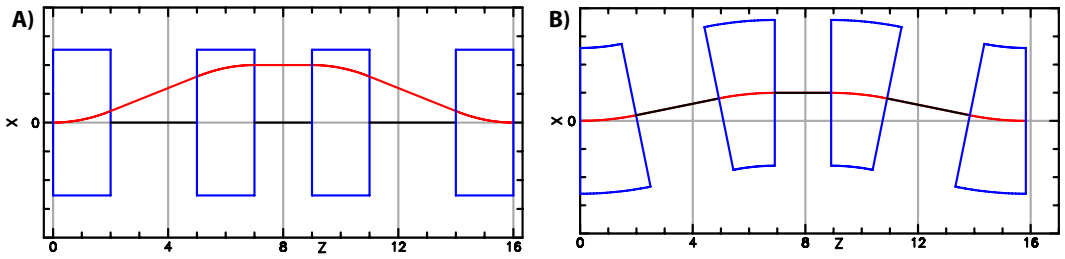

Figure 13.2: Four bend chicane. A) Correctly implemented chicane. The red line is the beam orbit magnified by a factor of 100. B) Incorrectly implemented chicane where the g parameter of the bends is used to control the chicane strength instead of dg.

be varied. Notice that in the case where g is varied, the bends are no longer rectangular. This is true since rbend elements, when read in from the lattice file, are converted to sbend elements. In this case the converted bends have e1 and e2 face angles of zero. When the chicane overlay varies g, the face angle remain zero. That is, the bends will always be pure sector bends. [In a situation where, indeed, g is to be varied and it is desired to keep bends rectangular, the appropriate variation of the e1 and e2 can be put in the controlling overlay.]

13.3 Example: Energy Recovery Linac

An Energy Recovery Linac (ERL) is illustrated in Fig. 13.3A. The ERL starts with an injection line that feeds a linac which accelerates the beam to some energy. The beam then transverses a return arc which reinjects the bunches into the linac. The length of the return arc is such that, on the second pass, the beam is decelerated giving its energy back to the RF cavities. Finally, the decelerated beam is steered through a dump line where, at the end, an absorber stops the beam.

A lattice file for modeling this ERL:

```
parameter[geometry] = open
bmad_com[absolute_time_tracking] = T

BEND_L1: sbend, angle = -25*pi/180, l = 0.2, ...
BEND_L2: BEND_L1

A_PATCH: patch, flexible = T
D_PATCH: patch, x_offset = 0.034, x_pitch = asin(0.32)
INJECT: line = (...)
LINAC: line[multipass] = (BEND_L1, ..., BEND_L2)
ARC: line = (..., BEND_A7)
DUMP: line = (...)

ERL: line = (INJECT, LINAC, ARC, A_PATCH, LINAC, D_PATCH, DUMP)
```

Fig. 13.3B shows the injector and arc merging into the beginning of the linac. The first element of the linac is a bend named BEND_L1. The bending angle for BEND_L1 has been set at the appropriate value for injection from the injector. To get the correct geometry for injection from the arc, a patch element, named A_PATCH, is placed in the ERL line between the arc and the linac. A_PATCH is a flexible patch which means that the exit edge of A_PATCH will automatically be aligned with the entrance edge of the next element which is BEND_L1.

Note that this use of a flexible patch works since the orientation of BEND_L1 has been determined before the orientation of A_PATCH is determined. The orientation of elements is determined in order starting from the first element in the line (the exception to this rule is if there is a floor_position element) and the orientation of BEND_L1 is thus determined right after the injector section on the first pass through the linac.

Fig. 13.3C shows the end of the linac splitting off into the dump and arc sections. The D_PATCH is used to orient the reference trajectory so that the dump is correctly positioned. Here it is not possible to make the D_PATCH flexible since the position of the dump is unknown when the orientation of the D_PATCH is calculated. However, the D_PATCH could be made flexible if a floor_position element is used in the dump line (*Bmad* will work both forward and backwards from a floor_position element so that a floor_position element may be placed anywhere in the dump line).

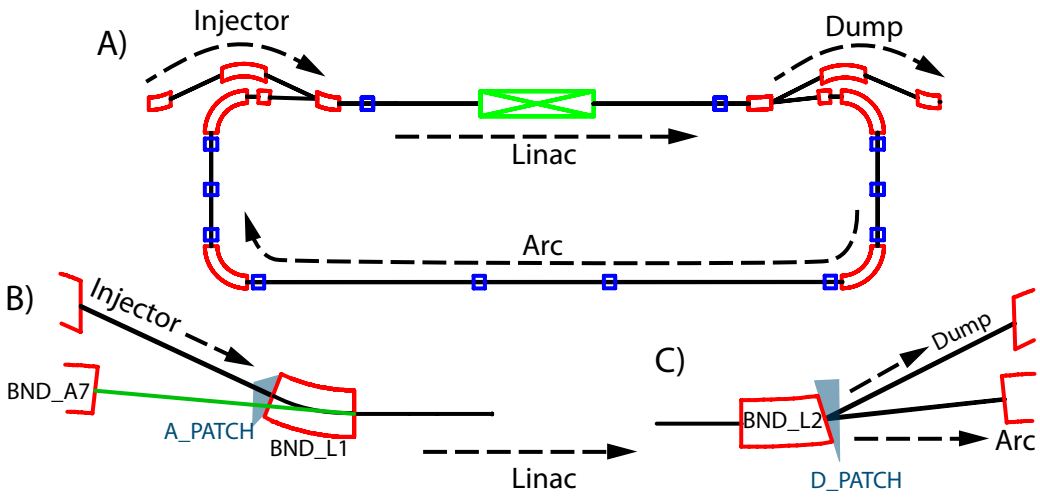

Figure 13.3: Example Energy Recovery Linac. A) The ERL consists of an injection line, accelerating linac, return arc, deceleration linac, and finally a beam dump. B) Close up of the section where the end of the injector and the end of the arc inject into the beginning of the linac. C) Close up of the end of the linac which injects into the dump and the beginning of the arc.

13.4 Example: Colliding Beam Storage Rings

The idealized layout of a pair of storage rings used for colliding counter rotating beams of electrons and gold is shown in Fig. 13.4. Rings `A` and `B` intersect at two interaction regions labeled `ir1` and `ir2` where the beams collide. The basic lattice description is:

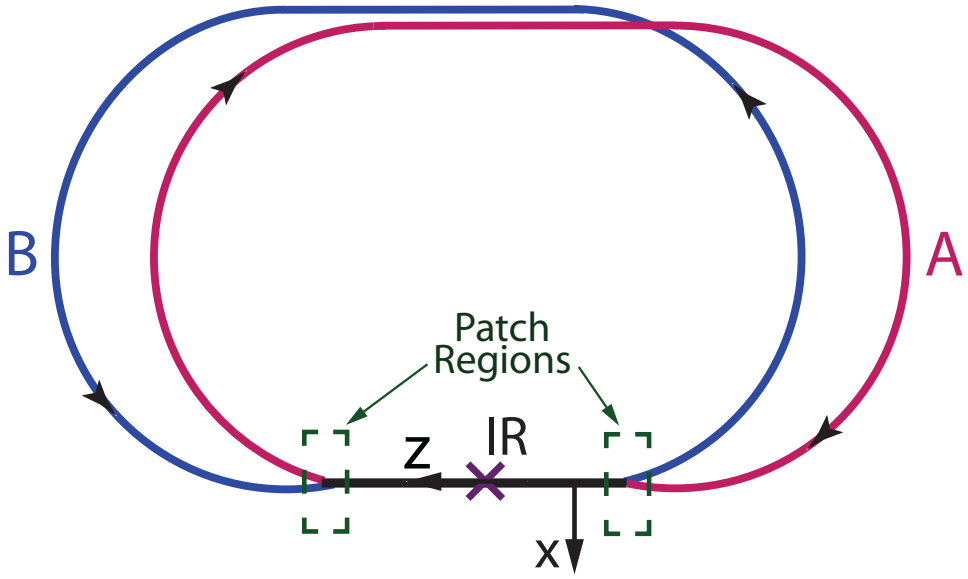

Figure 13.4: Dual ring colliding beam machine. The beam in the `A` ring rotates clockwise and in the `B` ring counterclockwise.

```
ir: line[multipass] = (...)
pa_in; patch, ... ;  pa_out; patch, ...
pb_in; patch, ... ;  pb_out; patch, ...
m: marker
fid: fiducial, origin_ele = m
...
A: line = (arc_a, pa_in, ir, m, pa_out)
A[particle] = electron
B_rev: line = (arc_b, pb_in, ir, fid, pb_out)
B: line = (--B_rev)
B[particle] = Au+79
use, A, B
```

Lines `ir` is the interaction region line which is declared `multipass` since they are shared by the two rings. Line `A` represents ring `A`. In ring `A` where the electron beam which, by definition, travels in the same direction as increasing s, rotates clockwise. Line `B_rev` is a "reversed" line of ring `B` and, like `a`, represents a beam rotating clockwise. Line `B`, which represents ring `B`, is the reverse of `B_rev` and here the gold beam rotates counterclockwise. In this construction, all elements of `B` are reversed. While this is not mandatory (only the interaction regions must be reversed in `B`), having all of `B` reversed simplifies the geometry since this means that the local coordinate systems of both lines `A` and `b` will be "aligned" with the x-axis pointing to the outside of the ring and the y-axis pointing up, out of the page. Having non-aligned coordinate systems is possible but potentially very confusing.

The two rings are physically aligned using a marker `m` in `A` and a `fiducial` element `fid` in `B` that aligns with `m`. Each ring has two rigid `patch` elements, `pa_in` and `pa_out` for the `A` ring, and `pb_in` and `bp_out` for the `B` ring, on either side of the interaction region. The dashed, green rectangles in the figure show the regions where the patches are.

The finished lattice will have two branches, The first branch (with index 0) will be derived from line `A` (and hence will be named "A") and the second branch (with index 1) will be derived from line `B` (and hence will be named "B"). The multipass lords representing the physical IR elements will be in the "lord section" of branch 0.

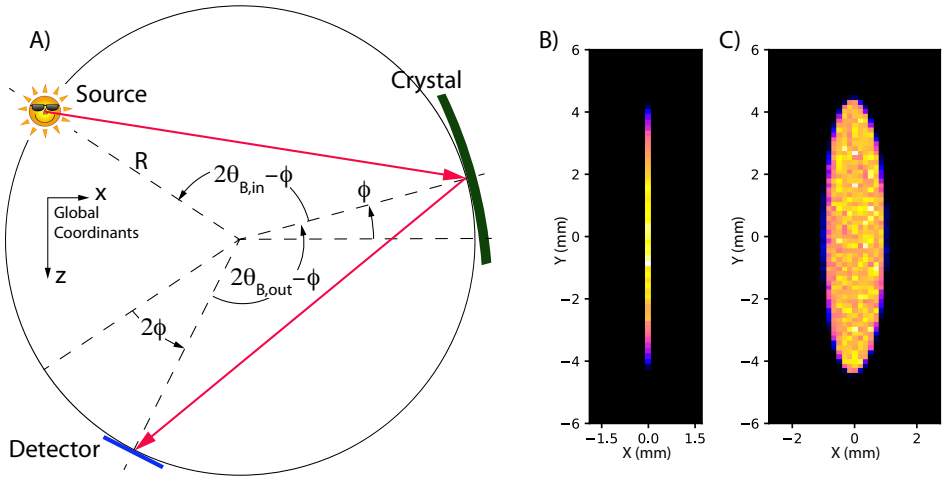

Figure 13.5: Rowland circle spectrometer: A) X-rays scattered from a sample (labeled source in the figure) illuminates a crystal with X-rays. Some of the X-rays are reflected from the crystal onto the detector. Note: For clarity's sake the center of the global coordinate system as shown is shifted from the true center at the Source element. B) The detector image when the radius of curvature of the bent crystal is "perfect". That is, twice the radius of the Rowland circle. C) The detector image when the radius of curvature of the crystal is shifted 1% from perfect.

13.5 Example: Rowland Circle X-Ray Spectrometer

This example shows how *Bmad* can be used to simulate X-rays. In this case, the present example is taken from a case study where simulations were done in order to understand how imperfections in a Rowland circle spectrometer would affect measurements.

A Rowland circle spectrometer is illustrated in Fig. 13.5A. The `source` was a sample that is illuminated with X-rays. Some of the X-rays scatter from the `sample` and are reflected from the `crystal` to the `detector`. To properly focus the X-rays onto the `detector`, the `source`, `crystal` and `detector` lie on a circle, called the Rowland circle. The `crystal` is bent and the radius of curvature of the crystal is $2R$ where R is the radius of the Rowland circle.

The angle from the `source` to the Rowland circle center to the crystal is $2\theta_{\rm B,in} - \phi$ where $\theta_{\rm B,in}$ is the entrance Bragg angle for photons whose energy matches the given reference energy and ϕ is an angle that will be varied when doing an energy scan of the scattered X-ray spectrum. Similarly, the angle from the `crystal` to the Rowland circle center to the `detector` is $2\theta_{\rm B,out} - \phi$ where $\theta_{\rm B,out}$ is the exit Bragg angle at the given reference energy.

The lattice for this simulation is:

```
beginning[e_tot] = 8.955e3     ! Reference photon energy
parameter[particle] = photon

phi = 0
err = 0
r_rowland = 0.5                ! Rowland circle radius

source: photon_init, sig_x = 5e-5, sig_y = 5e-5, spatial_distribution = uniform,
        E_center_relative_to_ref = T, sig_E = 2, energy_distribution = gaussian,
        velocity_distribution = spherical
```

```
drift1: drift

cryst: crystal, crystal_type = "Si(553)", b_param = -1, aperture = 0.050,
 curvature = {spherical = (1+err) / (2 * r_rowland)}, aperture_type = elliptical

drift2: drift

det: detector, pixel = {ix_bounds = (-97, 97),
                                  iy_bounds = (-243, 243), dr = (172e-6, 172e-6)}

daves_line: line = (source, drift1, cryst, drift2, det)
use, daves_line

!------------------

expand_lattice ! Calculates the Bragg angles needed below.

theta_in  = cryst[bragg_angle_in]  ! 78.2759 * pi / 180
theta_out = cryst[bragg_angle_out] ! 78.2759 * pi / 180

cryst[graze_angle_in]  = theta_in - phi/2
cryst[graze_angle_out] = theta_out - phi/2
drift1[L] = 2 * r_rowland * sin(theta_in-phi/2)
drift2[L] = 2 * r_rowland * sin(theta_out-phi/2)

beginning[theta_position] = theta_in + phi/2
det[x_pitch] = pi/2 - theta_out + phi/2
```

The reference photon energy is 8.995 KeV and the Rowland circle radius is 0.5 m. The simulation uses a `photon_init` element (§4.42) for the `source` having a Gaussian energy spread with a sigma of 2 eV. The initial velocity distribution of the photons, set by the `velocity_distribution` parameter, is taken to be uniform in all directions ("`spherical`" distribution). Since the element that is downstream from the source (which is the crystal element) has a defined aperture, *Bmad* is able to use this to not generate photons that will be lost at the crystal. This reduces the simulation time.

The crystal is `Silicon` 553 crystal which is symmetrically cut (`b_param` = -1) so that in this example the entrance Bragg angle is equal to the exit Bragg angle. The detector has a segmented surface with pixels spaced 172μm apart. Along the x-axis, which is the coordinate along the detector surface in the plane of Fig. 13.5A, the pixel index is in the range $[-97, 97]$. Along the y-axis, which is the out of plane coordinate, the pixel index is in the range $[-243, 243]$.

The `expand_lattice` command (§3.25) is used to command *Bmad* to construct the lattice which includes calculating the Bragg angles. After lattice expansion, the variables `theta_in` and `theta_out` are set to the Bragg angle entrance and exit Bragg angles respectively. The entrance and exit graze angles of the crystal, which are used to determine the reference trajectory (§16.2.3), can be set to `theta_in - phi` and `theta_out - phi` respectively. Note that if these graze angles had not been explicitly set, the graze angles would be automatically set to the Bragg angles which is not what is wanted when doing an energy scan with finite `phi`.

In the actual experimental setup that this example is modeled on, the source and Rowland circle were fixed in the global coordinate system (§16.2) while the crystal and detector move with changing `phi` (see Fig. 13.5A). To mimic this, the `beginning[theta_position]` (§10.4) is set to give the desired

orientation of the beginning reference trajectory within the global coordinate system. This does not affect photon tracking since changing the initial orientation of the reference trajectory just shifts all the lattice elements as one rigid body. Additionally, the detector orientation is fixed so that the detector surface normal always points towards the Rowland circle center. To get the correct orientation for the detector, the detector's `x_pitch` attribute, which rotates the detector (§5.6), is set appropriately.

The effect of varying the `crystal` curvature is shown in Fig. 13.5B and Fig. 13.5C. A *Bmad* based program called `Lux` was used for the simulation. The `Lux` program generates a set of photons and records the statistics at the detector. In Fig. 13.5B the `crystal` is correctly bent with the parameter `err` in the lattice set to zero. This produces a well focused spot on the `detector`. In Fig. 13.5C the `crystal` curvature is shifted by 1% by setting `err` equal to 0.01. This error degrades the focusing and leads to a spot that is enlarged along the x-axis.

13.6 Example: Backward Tracking Through a Lattice

By creating a reversed lattice, one can track a particles backwards through the lattice. For example, assume that you have a lattice file called `original_lattice.bmad` which defines a line called `original_line` tracking electrons. To create a reversed lattice, create a new file with the following:

```
call, file = original_lattice.bmad
reversed_line: line = (--original_line)
parameter[default_tracking_species] = positron
use, reversed_line
```

The "−−" reverses the line and reverses the elements (§7.4). Tracking through `reversed_line` is equivalent to tracking backwards through `original_line`.

The default for the type of particle tracked is set by `parameter[default_tracking_species]` (§10.1). [A *Bmad* based program can always override this default but it will be assumed here that this is not the case.] In this case, the default species to use for tracking is set to the antiparticle of the reference particle species. If the `original_line` lattice had just static magnetic fields and no electric fields, by tracking with the anti-particle in the reversed lattice, the anti-particle will follow the same path (but backward) as the particle in the original lattice. For this to work, the anti-particle must be started with the appropriate phase space coordinates. If (x, p_x, y, p_y, z, p_z) is the phase space coordinates of the particle at the end of the original lattice, the anti-particle must be initialized with phase space coordinates of $(x, -p_x, y, -p_y, \text{immaterial}, p_z)$.

It should be keept in mind that tracking backwards in the lattice is not exactly the same as tracking backwards in time. In particular, the two are different if there are electric fields or if radiation damping and/or excitation is turned on.

Chapter 14

Lattice File Conversion

A *Bmad* Distribution (§1.2) contains a number of translation programs between *Bmad* and other formats. Generally, conversion from *Bmad* is done using Fortran code to take advantage of the existing routines that parse *Bmad* lattice files. As such, conversion from *Bmad* is built into the *Tao* program. Conversion to *Bmad* is generally done using Python scripts to take advantage of Python's superior string handling abilities. As such, many programs to convert to *Bmad* are stored in the `util_programs` directory of the *Bmad* repo.

Conversions have been developed for:

`Astra`	`GPT`	`PALS`
`Blender`	`MAD-8`	`SAD`
`CSRtrack`	`MAD-x`	`SciBmad`
`Elegant`	`Merlin`	`SlickTrack`

14.1 Bmad to PALS Conversion

The *Bmad* to PALS conversion creates a YAML file with the PALS syntax. Extra information is included to aid in any back convertion.

14.2 MAD Conversion

14.2.1 Convert MAD to Bmad

Python scripts to convert from MAD8 and MADX are available at:

```
util_programs/mad_to_bmad
```

Due to differences in language definitions, conversions must be done with some care. The following differences should be noted:

- *Bmad*, unlike *MAD*, does not have any "action" commands. An action command is a command that makes a calculation. Examples include *MAD*'s `SURVEY` and `TWISS` commands.
- In *Bmad* all variables must be defined before being used (§3.13) while *MAD* will simply take a variable's value to be zero if it is not defined.
- *Bmad*, unlike *MAD*, does not allow variable values to be redefined.

14.2.2 Convert Bmad to MAD

To convert to MAD8 or MADX, the *Tao* program can be used. Additionally, there is the program
`util_programs/bmad_to_mad_sad_elegant`
Since *MAD* does not have a `wiggler` or a `sol_quad` element, this conversion routine makes "equivalent" substitution. For a `sol_quad`, the equivalent substitution will be a drift-matrix-drift series of elements. For a `wiggler`, a series of bend and drift elements will be used (the program can also use a drift-matrix-drift model here but that is not as accurate). The bends and drifts for the `wiggler` model are constructed so that the global geometry of the lattice does not change. Additionally the bends and drifts are constructed to most nearly match the wiggler's
`Transfer matrix`
`$I_2$ and $I_3$ synchrotron radiation integrals` (§22.3)
Note that the resulting model will not have the vertical cubic nonlinearity that the actual wiggler has.

14.3 Convert to PTC

A PTC "`flatfile`" can be constructed using the *Tao* program with the following commands:
```
Tao> ptc init
Tao> write ptc
```

14.4 SAD Conversion

Conversion from `SAD`[SAD] to *Bmad* is accomplished using the Python script
`util_programs/sad_to_bmad/sad_to_bmad.py`
Currently, the following restrictions on SAD lattices apply:

- SAD `mult` elements cannot have an associated RF field
- Misalignments in a `sol` element with `geo` = 1 cannot be handled.

Bmad to `SAD` to conversion can be done with the *Tao* program or the program
`util_programs/bmad_to_mad_sad_elegant`

14.5 Elegant Conversion

Conversion from `Elegant`[Elegant] to *Bmad* is accomplished using the Python script
`util_programs/elegant_to_bmad/elegant_to_bmad.py`
Bmad to `Elegant` to conversion can be done with the *Tao* program or the program
`util_programs/bmad_to_mad_sad_elegant`

14.6 Astra, Blender, CSRTrack, GPT, and Merlin Conversion

Conversion programs to Astra, Blender, CSRTrack, GPT, and Merlin exist in the `util_programs` directory. Some conversion code is still in beta development so if you encounter problems please contact a *Bmad* maintainer.

Chapter 15

List of Element Attributes

Alphabetical list of element attributes for each type of element.

Note for programmers: The program that generates a file of attributes indexed by the internal reference number is:

`util_programs/element_attributes.f90`

15.1 !PTC_Com Element Attributes

exact_misalign	max_fringe_order	vertical_kick
exact_model	old_integrator	

15.2 !Space_Charge_Com Element Attributes

15.3 AB_multipole Element Attributes

a0 - a20, b0 - b20	is_on	spin_tracking_method	x_offset [m]
alias	l [m]	superimpose	x_offset_tot [m]
aperture [m]	mat6_calc_method	tilt [rad]	y1_limit [m]
aperture_at	multipoles_on	tilt_tot [rad]	y2_limit [m]
aperture_type	offset [m]	tracking_method	y_limit [m]
create_jumbo_slave	offset_moves_aperture	type	y_offset [m]
delta_ref_time [sec]	p0c [eV]	wall	y_offset_tot [m]
descrip	ptc_integration_type	wrap_superimpose	z_offset [m]
e_tot [eV]	ref_origin	x1_limit [m]	z_offset_tot [m]
ele_origin	ref_time_start [sec]	x2_limit [m]	
field_master	reference	x_limit [m]	

15.4 AC_Kicker Element Attributes

a0 - a20, b0 - b20	frequencies	p0c [eV]	type
alias	fringe_at	phi0_multipass [rad/2pi]	vkick
amp_vs_time	fringe_type	ptc_integration_type	wall
aperture [m]	gen_gradients	r0_elec [m]	wrap_superimpose
aperture_at	grid_field	r0_mag [m]	x1_limit [m]
aperture_type	hkick	ref_origin	x2_limit [m]
bl_hkick [T*m]	integrator_order	ref_time_start [sec]	x_limit [m]
bl_vkick [T*m]	interpolation	reference	x_offset [m]
cartesian_map	is_on	scale_multipoles	x_offset_tot [m]
create_jumbo_slave	l [m]	space_charge_method	x_pitch [rad]
csr_ds_step [m]	lord_pad1 [m]	spin_fringe_on	x_pitch_tot [rad]
csr_method	lord_pad2 [m]	spin_tracking_method	y1_limit [m]
cylindrical_map	lr_freq_spread [Hz]	sr_wake	y2_limit [m]
delta_ref_time [sec]	lr_self_wake_on	sr_wake_file	y_limit [m]
descrip	lr_wake	superimpose	y_offset [m]
ds_step [m]	lr_wake_file	symplectify	y_offset_tot [m]
e_tot [eV]	mat6_calc_method	t_offset [sec]	y_pitch [rad]
ele_origin	multipoles_on	taylor_map_includes_offsets	y_pitch_tot [rad]
field_calc	num_steps	tilt [rad]	z_offset [m]
field_master	offset [m]	tilt_tot [rad]	z_offset_tot [m]
field_overlaps	offset_moves_aperture	tracking_method	

15.5 BeamBeam Element Attributes

alias	crab_x4 [1/m^3]	pc_strong [eV]	wrap_superimpose
alpha_a_strong	crab_x5 [1/m^4]	ptc_integration_type	x1_limit [m]
alpha_b_strong	create_jumbo_slave	ref_origin	x2_limit [m]
aperture [m]	crossing_time [sec]	ref_time_start [sec]	x_limit [m]
aperture_at	delta_ref_time [sec]	reference	x_offset [m]
aperture_type	descrip	repetition_frequency [Hz]	x_offset_tot [m]
bbi_constant	e_tot [eV]	s_twiss_ref [m]	x_pitch [rad]
beta_a_strong [m]	e_tot_strong [eV]	sig_x [m]	x_pitch_tot [rad]
beta_b_strong [m]	ele_origin	sig_y [m]	y1_limit [m]
bs_field [T]	field_calc	sig_z [m]	y2_limit [m]
charge	is_on	species_strong	y_limit [m]
cmat_11	ks [1/m]	spin_tracking_method	y_offset [m]
cmat_12 [m]	l [m]	superimpose	y_offset_tot [m]
cmat_21 [1/m]	mat6_calc_method	symplectify	y_pitch [rad]
cmat_22	n_particle	tilt [rad]	y_pitch_tot [rad]
crab_tilt [rad]	n_slice	tilt_tot [rad]	z_offset [m]
crab_x1	offset [m]	tracking_method	z_offset_tot [m]
crab_x2 [1/m]	offset_moves_aperture	type	
crab_x3 [1/m^2]	p0c [eV]	wall	

15.6 Beginning Statement Attributes

alpha_a	dbeta_dpz_a [m]	eta_y [m]	pz_stored
alpha_a_stored	dbeta_dpz_a_stored [m]	eta_y_stored [m]	ref_species
alpha_b	dbeta_dpz_b [m]	eta_z [m]	ref_time [sec]
alpha_b_stored	dbeta_dpz_b_stored [m]	etap_x	s [m]
beta_a [m]	dcmat_dpz_11_stored [???]	etap_x_stored	spin_dn_dpz_x
beta_a_stored [m]	dcmat_dpz_12_stored [???]	etap_y	spin_dn_dpz_y
beta_b [m]	dcmat_dpz_21_stored [???]	etap_y_stored	spin_dn_dpz_z
beta_b_stored [m]	dcmat_dpz_22_stored [???]	inherit_from_fork	spin_x_stored
cmat_11	deta_dpz_x [m]	mode_flip	spin_y_stored
cmat_11_stored	deta_dpz_x_stored [m]	mode_flip_stored	spin_z_stored
cmat_12 [m]	deta_dpz_y [m]	p0c [eV]	theta_position [rad]
cmat_12_stored [m]	deta_dpz_y_stored [m]	p0c_start [eV]	x_position [m]
cmat_21 [1/m]	detap_dpz_x	phi_a [rad]	x_stored [m]
cmat_21_stored [1/m]	detap_dpz_x_stored	phi_a_stored	y_position [m]
cmat_22	detap_dpz_y	phi_b [rad]	y_stored [m]
cmat_22_stored	detap_dpz_y_stored	phi_b_stored	z_position [m]
dalpha_dpz_a	e_tot [eV]	phi_position [rad]	z_stored [m]
dalpha_dpz_a_stored	e_tot_start [eV]	psi_position [rad]	
dalpha_dpz_b	eta_x [m]	px_stored	
dalpha_dpz_b_stored	eta_x_stored [m]	py_stored	

15.7 Bends: Rbend and Sbend Element Attributes

a0 - a20, b0 - b20	fiducial_pt	lr_freq_spread [Hz]	spin_tracking_method
alias	field_calc	lr_self_wake_on	sr_wake
angle [rad]	field_master	lr_wake	sr_wake_file
aperture [m]	field_overlaps	lr_wake_file	superimpose
aperture_at	fint	mat6_calc_method	symplectify
aperture_type	fintx	multipoles_on	taylor_map_includes_offsets
b1_gradient [T/m]	fringe_at	num_steps	tracking_method
b2_gradient [T/m^2]	fringe_type	offset [m]	type
b_field [T]	g [1/m]	offset_moves_aperture	vkick
b_field_tot [T]	g_tot [1/m]	p0c [eV]	wall
bl_hkick [T*m]	gen_gradients	ptc_canonical_coords	wrap_superimpose
bl_vkick [T*m]	grid_field	ptc_field_geometry	x1_limit [m]
cartesian_map	h1 [1/m]	ptc_fringe_geometry	x2_limit [m]
create_jumbo_slave	h2 [1/m]	ptc_integration_type	x_limit [m]
csr_ds_step [m]	hgap [m]	r0_elec [m]	x_offset [m]
csr_method	hgapx [m]	r0_mag [m]	x_offset_tot [m]
cylindrical_map	hkick	ref_origin	x_pitch [rad]
db_field [T]	integrator_order	ref_tilt [rad]	x_pitch_tot [rad]
delta_ref_time [sec]	is_on	ref_tilt_tot [rad]	y1_limit [m]
descrip	k1 [1/m^2]	ref_time_start [sec]	y2_limit [m]
dg [1/m]	k2 [1/m^3]	reference	y_limit [m]
ds_step [m]	l [m]	rho [m]	y_offset [m]
e1 [rad]	l_chord [m]	roll [rad]	y_offset_tot [m]
e2 [rad]	l_rectangle [m]	roll_tot [rad]	y_pitch [rad]
e_tot [eV]	l_sagitta [m]	scale_multipoles	y_pitch_tot [rad]
ele_origin	lord_pad1 [m]	space_charge_method	z_offset [m]
exact_multipoles	lord_pad2 [m]	spin_fringe_on	z_offset_tot [m]

15.8 Bmad_Com Statement Attributes

15.9 Capillary Element Attributes

alias	mat6_calc_method	tilt [rad]	x_pitch [rad]
aperture [m]	n_slice_spline	tilt_tot [rad]	x_pitch_tot [rad]
aperture_at	offset [m]	tracking_method	y1_limit [m]
aperture_type	offset_moves_aperture	type	y2_limit [m]
create_jumbo_slave	p0c [eV]	wall	y_limit [m]
critical_angle_factor [rad*eV]	ptc_integration_type	wrap_superimpose	y_offset [m]
delta_ref_time [sec]	ref_origin	x1_limit [m]	y_offset_tot [m]
descrip	ref_time_start [sec]	x2_limit [m]	y_pitch [rad]
e_tot [eV]	reference	x_limit [m]	y_pitch_tot [rad]
ele_origin	spin_tracking_method	x_offset [m]	z_offset [m]
l [m]	superimpose	x_offset_tot [m]	z_offset_tot [m]

15.10 Collimators: Ecollimator and Rcollimator Element Attributes

alias	integrator_order	pz_aperture_center	wrap_superimpose
aperture [m]	is_on	pz_aperture_width2	x1_limit [m]
aperture_at	l [m]	ref_origin	x2_limit [m]
aperture_type	lord_pad1 [m]	ref_time_start [sec]	x_limit [m]
bl_hkick [T*m]	lord_pad2 [m]	reference	x_offset [m]
bl_vkick [T*m]	lr_freq_spread [Hz]	space_charge_method	x_offset_tot [m]
create_jumbo_slave	lr_self_wake_on	spin_fringe_on	x_pitch [rad]
csr_ds_step [m]	lr_wake	spin_tracking_method	x_pitch_tot [rad]
csr_method	lr_wake_file	sr_wake	y1_limit [m]
delta_ref_time [sec]	mat6_calc_method	sr_wake_file	y2_limit [m]
descrip	num_steps	superimpose	y_limit [m]
ds_step [m]	offset [m]	symplectify	y_offset [m]
e_tot [eV]	offset_moves_aperture	taylor_map_includes_offsets	y_offset_tot [m]
ele_origin	p0c [eV]	tilt [rad]	y_pitch [rad]
field_calc	ptc_integration_type	tilt_tot [rad]	y_pitch_tot [rad]
field_overlaps	px_aperture_center	tracking_method	z_aperture_center [m]
fringe_at	px_aperture_width2	type	z_aperture_width2 [m]
fringe_type	py_aperture_center	vkick	z_offset [m]
hkick	py_aperture_width2	wall	z_offset_tot [m]

15.11 Converter Element Attributes

alias	l [m]	spin_tracking_method	x_pitch [rad]
angle_out_max [rad]	mat6_calc_method	superimpose	x_pitch_tot [rad]
aperture [m]	offset [m]	tilt [rad]	y1_limit [m]
aperture_at	offset_moves_aperture	tilt_tot [rad]	y2_limit [m]
aperture_type	p0c [eV]	tracking_method	y_limit [m]
create_jumbo_slave	p0c_start [eV]	type	y_offset [m]
delta_ref_time [sec]	pc_out_max [eV]	wall	y_offset_tot [m]
descrip	pc_out_min [eV]	wrap_superimpose	y_pitch [rad]
distribution	ptc_integration_type	x1_limit [m]	y_pitch_tot [rad]
e_tot [eV]	ref_origin	x2_limit [m]	z_offset [m]
e_tot_start [eV]	ref_time_start [sec]	x_limit [m]	z_offset_tot [m]
ele_origin	reference	x_offset [m]	
is_on	species_out	x_offset_tot [m]	

15.12 Crab_Cavity Element Attributes

a0 - a20, b0 - b20	gradient [eV/m]	phi0 [rad/2pi]	voltage [Volt]
alias	grid_field	phi0_multipass [rad/2pi]	wall
aperture [m]	harmon	ptc_integration_type	wrap_superimpose
aperture_at	harmon_master	ref_origin	x1_limit [m]
aperture_type	hkick	ref_time_start [sec]	x2_limit [m]
bl_hkick [T*m]	integrator_order	reference	x_limit [m]
bl_vkick [T*m]	is_on	rf_frequency [Hz]	x_offset [m]
cartesian_map	l [m]	rf_wavelength [m]	x_offset_tot [m]
create_jumbo_slave	lord_pad1 [m]	scale_multipoles	x_pitch [rad]
csr_ds_step [m]	lord_pad2 [m]	space_charge_method	x_pitch_tot [rad]
csr_method	lr_freq_spread [Hz]	spin_tracking_method	y1_limit [m]
cylindrical_map	lr_self_wake_on	sr_wake	y2_limit [m]
delta_ref_time [sec]	lr_wake	sr_wake_file	y_limit [m]
descrip	lr_wake_file	superimpose	y_offset [m]
ds_step [m]	mat6_calc_method	symplectify	y_offset_tot [m]
e_tot [eV]	multipoles_on	taylor_map_includes_offsets	y_pitch [rad]
ele_origin	n_rf_steps	tilt [rad]	y_pitch_tot [rad]
field_calc	num_steps	tilt_tot [rad]	z_offset [m]
field_master	offset [m]	tracking_method	z_offset_tot [m]
field_overlaps	offset_moves_aperture	type	
gen_gradients	p0c [eV]	vkick	

15.13 Crystal Element Attributes

alias	e_tot [eV]	ref_orbit_follows	wrap_superimpose
alpha_angle	ele_origin	ref_origin	x1_limit [m]
aperture [m]	graze_angle_in [rad]	ref_tilt [rad]	x2_limit [m]
aperture_at	graze_angle_out [rad]	ref_tilt_tot [rad]	x_limit [m]
aperture_type	h_misalign	ref_time_start [sec]	x_offset [m]
b_param	is_mosaic	ref_wavelength [m]	x_offset_tot [m]
bragg_angle [rad]	l [m]	reference	x_pitch [rad]
bragg_angle_in [rad]	mat6_calc_method	reflectivity_table	x_pitch_tot [rad]
bragg_angle_out [rad]	mosaic_angle_rms_in_plane [rad]	segmented	y1_limit [m]
create_jumbo_slave	mosaic_angle_rms_out_plane [rad]	spin_tracking_method	y2_limit [m]
crystal_type	mosaic_diffraction_num	superimpose	y_limit [m]
curvature	mosaic_thickness [m]	thickness [m]	y_offset [m]
curvature_x0_y2 [1/m]	offset [m]	tilt [rad]	y_offset_tot [m]
d_spacing [m]	offset_moves_aperture	tilt_corr [rad]	y_pitch [rad]
darwin_width_pi [rad]	p0c [eV]	tilt_tot [rad]	y_pitch_tot [rad]
darwin_width_sigma [rad]	pendellosung_period_pi [m]	tracking_method	z_offset [m]
dbragg_angle_de [rad/eV]	pendellosung_period_sigma [m]	type	z_offset_tot [m]
delta_ref_time [sec]	psi_angle [rad]	use_reflectivity_table	
descrip	ptc_integration_type	v_unitcell [m^3]	
displacement	ref_cap_gamma	wall	

15.14 Custom Element Attributes

alias	lord_pad1 [m]	superimpose	wrap_superimpose
aperture [m]	lord_pad2 [m]	symplectify	x1_limit [m]
aperture_at	lr_freq_spread [Hz]	taylor_map_includes_offsets	x2_limit [m]
aperture_type	lr_self_wake_on	tilt [rad]	x_limit [m]
create_jumbo_slave	lr_wake	tilt_tot [rad]	x_offset [m]
csr_ds_step [m]	lr_wake_file	tracking_method	x_offset_tot [m]
csr_method	mat6_calc_method	type	x_pitch [rad]
delta_e_ref [eV]	num_steps	val1	x_pitch_tot [rad]
delta_ref_time [sec]	offset [m]	val10	y1_limit [m]
descrip	offset_moves_aperture	val11	y2_limit [m]
ds_step [m]	p0c [eV]	val12	y_limit [m]
e_tot [eV]	p0c_start [eV]	val2	y_offset [m]
e_tot_start [eV]	ptc_integration_type	val3	y_offset_tot [m]
ele_origin	ref_origin	val4	y_pitch [rad]
field_calc	ref_time_start [sec]	val5	y_pitch_tot [rad]
field_master	reference	val6	z_offset [m]
field_overlaps	space_charge_method	val7	z_offset_tot [m]
integrator_order	spin_tracking_method	val8	
is_on	sr_wake	val9	
l [m]	sr_wake_file	wall	

15.15 Detector Element Attributes

alias	mat6_calc_method	tilt_tot [rad]	y1_limit [m]
aperture [m]	n_sample	tracking_method	y2_limit [m]
aperture_at	noise	type	y_dispersion_calib [m]
aperture_type	offset [m]	wall	y_dispersion_err [m]
create_jumbo_slave	offset_moves_aperture	wrap_superimpose	y_gain_calib [m]
crunch [rad]	osc_amplitude [m]	x1_limit [m]	y_gain_err [m]
crunch_calib [rad]	p0c [eV]	x2_limit [m]	y_limit [m]
curvature	pixel	x_dispersion_calib [m]	y_offset [m]
curvature_x0_y2 [1/m]	ptc_integration_type	x_dispersion_err [m]	y_offset_calib [m]
de_eta_meas	ref_origin	x_gain_calib [m]	y_offset_tot [m]
delta_ref_time [sec]	ref_time_start [sec]	x_gain_err [m]	y_pitch [rad]
descrip	reference	x_limit [m]	y_pitch_tot [rad]
displacement	segmented	x_offset [m]	z_offset [m]
e_tot [eV]	spin_tracking_method	x_offset_calib [m]	z_offset_tot [m]
ele_origin	superimpose	x_offset_tot [m]	
is_on	tilt [rad]	x_pitch [rad]	
l [m]	tilt_calib [rad]	x_pitch_tot [rad]	

15.16 Diffraction_Plate Element Attributes

alias	is_on	superimpose	x_pitch_tot [rad]
aperture [m]	mat6_calc_method	tilt [rad]	y1_limit [m]
aperture_at	mode	tilt_tot [rad]	y2_limit [m]
aperture_type	offset [m]	tracking_method	y_limit [m]
create_jumbo_slave	offset_moves_aperture	type	y_offset [m]
curvature	p0c [eV]	wall	y_offset_tot [m]
curvature_x0_y2 [1/m]	ptc_integration_type	wrap_superimpose	y_pitch [rad]
delta_ref_time [sec]	ref_origin	x1_limit [m]	y_pitch_tot [rad]
descrip	ref_time_start [sec]	x2_limit [m]	z_offset [m]
displacement	ref_wavelength [m]	x_limit [m]	z_offset_tot [m]
e_tot [eV]	reference	x_offset [m]	
ele_origin	segmented	x_offset_tot [m]	
field_scale_factor	spin_tracking_method	x_pitch [rad]	

15.17 Drift Element Attributes

alias	l [m]	superimpose	x_pitch [rad]
aperture [m]	lord_pad1 [m]	symplectify	x_pitch_tot [rad]
aperture_at	lord_pad2 [m]	taylor_map_includes_offsets	y1_limit [m]
aperture_type	mat6_calc_method	tilt [rad]	y2_limit [m]
create_jumbo_slave	num_steps	tilt_tot [rad]	y_limit [m]
csr_ds_step [m]	offset [m]	tracking_method	y_offset [m]
csr_method	offset_moves_aperture	type	y_offset_tot [m]
delta_ref_time [sec]	p0c [eV]	wall	y_pitch [rad]
descrip	ptc_integration_type	wrap_superimpose	y_pitch_tot [rad]
ds_step [m]	ref_origin	x1_limit [m]	z_offset [m]
e_tot [eV]	ref_time_start [sec]	x2_limit [m]	z_offset_tot [m]
ele_origin	reference	x_limit [m]	
field_calc	space_charge_method	x_offset [m]	
integrator_order	spin_tracking_method	x_offset_tot [m]	

15.18 ELSeparator Element Attributes

a0 - a20, b0 - b20	fringe_type	ptc_canonical_coords	vkick
alias	gap	ptc_integration_type	voltage [Volt]
aperture [m]	gen_gradients	r0_elec [m]	wall
aperture_at	grid_field	r0_mag [m]	wrap_superimpose
aperture_type	hkick	ref_origin	x1_limit [m]
cartesian_map	integrator_order	ref_time_start [sec]	x2_limit [m]
create_jumbo_slave	is_on	reference	x_limit [m]
csr_ds_step [m]	l [m]	scale_multipoles	x_offset [m]
csr_method	lord_pad1 [m]	space_charge_method	x_offset_tot [m]
cylindrical_map	lord_pad2 [m]	spin_fringe_on	x_pitch [rad]
delta_ref_time [sec]	lr_freq_spread [Hz]	spin_tracking_method	x_pitch_tot [rad]
descrip	lr_self_wake_on	sr_wake	y1_limit [m]
ds_step [m]	lr_wake	sr_wake_file	y2_limit [m]
e_field [V/m]	lr_wake_file	superimpose	y_limit [m]
e_tot [eV]	mat6_calc_method	symplectify	y_offset [m]
ele_origin	multipoles_on	taylor_map_includes_offsets	y_offset_tot [m]
field_calc	num_steps	tilt [rad]	y_pitch [rad]
field_master	offset [m]	tilt_tot [rad]	y_pitch_tot [rad]
field_overlaps	offset_moves_aperture	tracking_method	z_offset [m]
fringe_at	p0c [eV]	type	z_offset_tot [m]

15.19 EM_Field Element Attributes

alias	fringe_at	phi0_err [rad/2pi]	wall
aperture [m]	fringe_type	polarity	wrap_superimpose
aperture_at	gen_gradients	ptc_canonical_coords	x1_limit [m]
aperture_type	grid_field	ptc_integration_type	x2_limit [m]
autoscale_amplitude	integrator_order	ref_origin	x_limit [m]
autoscale_phase	is_on	ref_time_start [sec]	x_offset [m]
cartesian_map	l [m]	reference	x_offset_tot [m]
constant_ref_energy	lord_pad1 [m]	rf_frequency [Hz]	x_pitch [rad]
create_jumbo_slave	lord_pad2 [m]	rf_wavelength [m]	x_pitch_tot [rad]
csr_ds_step [m]	lr_freq_spread [Hz]	space_charge_method	y1_limit [m]
csr_method	lr_self_wake_on	spin_fringe_on	y2_limit [m]
cylindrical_map	lr_wake	spin_tracking_method	y_limit [m]
delta_ref_time [sec]	lr_wake_file	sr_wake	y_offset [m]
descrip	mat6_calc_method	sr_wake_file	y_offset_tot [m]
ds_step [m]	num_steps	superimpose	y_pitch [rad]
e_tot [eV]	offset [m]	symplectify	y_pitch_tot [rad]
e_tot_start [eV]	offset_moves_aperture	taylor_map_includes_offsets	z_offset [m]
ele_origin	p0c [eV]	tilt [rad]	z_offset_tot [m]
field_autoscale	p0c_start [eV]	tilt_tot [rad]	
field_calc	phi0 [rad/2pi]	tracking_method	
field_overlaps	phi0_autoscale [rad/2pi]	type	

15.20 E_Gun Element Attributes

alias	field_calc	p0c [eV]	type
aperture [m]	field_overlaps	phi0 [rad/2pi]	voltage [Volt]
aperture_at	fringe_at	phi0_autoscale [rad/2pi]	wall
aperture_type	fringe_type	ptc_integration_type	wrap_superimpose
autoscale_amplitude	gen_gradients	ref_origin	x1_limit [m]
autoscale_phase	gradient [eV/m]	ref_time_start [sec]	x2_limit [m]
cartesian_map	grid_field	reference	x_limit [m]
create_jumbo_slave	integrator_order	rf_frequency [Hz]	x_offset [m]
csr_ds_step [m]	is_on	rf_wavelength [m]	x_offset_tot [m]
csr_method	l [m]	space_charge_method	x_pitch [rad]
cylindrical_map	lord_pad1 [m]	spin_fringe_on	x_pitch_tot [rad]
delta_e_ref [eV]	lord_pad2 [m]	spin_tracking_method	y1_limit [m]
delta_ref_time [sec]	lr_freq_spread [Hz]	sr_wake	y2_limit [m]
descrip	lr_self_wake_on	sr_wake_file	y_limit [m]
ds_step [m]	lr_wake	superimpose	y_offset [m]
dt_max [sec]	lr_wake_file	symplectify	y_offset_tot [m]
e_tot [eV]	mat6_calc_method	taylor_map_includes_offsets	y_pitch [rad]
ele_origin	num_steps	tilt [rad]	y_pitch_tot [rad]
emit_fraction	offset [m]	tilt_tot [rad]	z_offset [m]
field_autoscale	offset_moves_aperture	tracking_method	z_offset_tot [m]

15.21 Feedback Element Attributes

alias	input_ele	type
descrip	output_ele	

15.22 Fiducial Element Attributes

alias	dy_origin [m]	origin_ele	spin_tracking_method
delta_ref_time [sec]	dz_origin [m]	origin_ele_ref_pt	superimpose
descrip	e_tot [eV]	p0c [eV]	tracking_method
dphi_origin [rad]	ele_origin	ptc_integration_type	type
dpsi_origin [rad]	l [m]	ref_origin	wrap_superimpose
dtheta_origin [rad]	mat6_calc_method	ref_time_start [sec]	
dx_origin [m]	offset [m]	reference	

15.23 Fixer Element Attributes

alias	dbeta_dpz_a_stored [m]	eta_y_stored [m]	ref_origin
alpha_a	dbeta_dpz_b [m]	eta_z [m]	ref_time_start [sec]
alpha_a_stored	dbeta_dpz_b_stored [m]	etap_x	reference
alpha_b	dcmat_dpz_11_stored [???]	etap_x_stored	s [m]
alpha_b_stored	dcmat_dpz_12_stored [???]	etap_y	spin_dn_dpz_x
beta_a [m]	dcmat_dpz_21_stored [???]	etap_y_stored	spin_dn_dpz_y
beta_a_stored [m]	dcmat_dpz_22_stored [???]	is_on	spin_dn_dpz_z
beta_b [m]	delta_ref_time [sec]	mat6_calc_method	spin_tracking_method
beta_b_stored [m]	descrip	mode_flip	spin_x_stored
cmat_11	deta_dpz_x [m]	mode_flip_stored	spin_y_stored
cmat_11_stored	deta_dpz_x_stored [m]	offset [m]	spin_z_stored
cmat_12 [m]	deta_dpz_y [m]	p0c [eV]	superimpose
cmat_12_stored [m]	deta_dpz_y_stored [m]	phi_a [rad]	theta_position [rad]
cmat_21 [1/m]	detap_dpz_x	phi_a_stored	tracking_method
cmat_21_stored [1/m]	detap_dpz_x_stored	phi_b [rad]	type
cmat_22	detap_dpz_y	phi_b_stored	wrap_superimpose
cmat_22_stored	detap_dpz_y_stored	phi_position [rad]	x_position [m]
dalpha_dpz_a	e_tot [eV]	psi_position [rad]	x_stored [m]
dalpha_dpz_a_stored	ele_origin	ptc_integration_type	y_position [m]
dalpha_dpz_b	eta_x [m]	px_stored	y_stored [m]
dalpha_dpz_b_stored	eta_x_stored [m]	py_stored	z_position [m]
dbeta_dpz_a [m]	eta_y [m]	pz_stored	z_stored [m]

15.24 Floor_Shift Element Attributes

alias	l [m]	reference	x2_limit [m]
aperture [m]	mat6_calc_method	spin_tracking_method	x_limit [m]
aperture_at	offset [m]	superimpose	x_offset [m]
aperture_type	offset_moves_aperture	tilt [rad]	x_pitch [rad]
create_jumbo_slave	origin_ele	tracking_method	y1_limit [m]
delta_ref_time [sec]	origin_ele_ref_pt	type	y2_limit [m]
descrip	p0c [eV]	upstream_ele_dir	y_limit [m]
downstream_ele_dir	ptc_integration_type	wall	y_offset [m]
e_tot [eV]	ref_origin	wrap_superimpose	y_pitch [rad]
ele_origin	ref_time_start [sec]	x1_limit [m]	z_offset [m]

15.25 Foil Element Attributes

alias	is_on	spin_tracking_method	x_pitch [rad]
aperture [m]	l [m]	superimpose	x_pitch_tot [rad]
aperture_at	mat6_calc_method	thickness [m]	y1_edge [m]
aperture_type	material_type	tilt [rad]	y1_limit [m]
area_density [kg/m^2]	num_steps	tilt_tot [rad]	y2_edge [m]
area_density_used [kg/m^2]	offset [m]	tracking_method	y2_limit [m]
create_jumbo_slave	offset_moves_aperture	type	y_limit [m]
delta_ref_time [sec]	p0c [eV]	wall	y_offset [m]
density [kg/m^3]	ptc_integration_type	wrap_superimpose	y_offset_tot [m]
density_used [kg/m^3]	radiation_length [kg/m^2]	x1_edge [m]	y_pitch [rad]
descrip	radiation_length_used [kg/m^2]	x1_limit [m]	y_pitch_tot [rad]
dthickness_dx	ref_origin	x2_edge [m]	z_offset [m]
e_tot [eV]	ref_time_start [sec]	x2_limit [m]	z_offset_tot [m]
ele_origin	reference	x_limit [m]	
f_factor	scatter_method	x_offset [m]	
final_charge	scatter_test	x_offset_tot [m]	

15.26 Fork and Photon_Fork Element Attributes

alias	is_on	ref_origin	wrap_superimpose
aperture [m]	ix_to_branch	ref_species	x1_limit [m]
aperture_at	ix_to_element	ref_time_start [sec]	x2_limit [m]
aperture_type	l [m]	reference	x_limit [m]
create_jumbo_slave	mat6_calc_method	spin_tracking_method	y1_limit [m]
delta_ref_time [sec]	new_branch	superimpose	y2_limit [m]
descrip	offset [m]	to_element	y_limit [m]
direction	offset_moves_aperture	to_line	
e_tot [eV]	p0c [eV]	tracking_method	
ele_origin	ptc_integration_type	type	

15.27 GKicker Element Attributes

alias	mat6_calc_method	ref_time_start [sec]	x_kick [m]
aperture [m]	offset [m]	reference	x_limit [m]
aperture_at	offset_moves_aperture	spin_tracking_method	y1_limit [m]
aperture_type	p0c [eV]	superimpose	y2_limit [m]
create_jumbo_slave	ptc_integration_type	tracking_method	y_kick [m]
delta_ref_time [sec]	px_kick	type	y_limit [m]
descrip	py_kick	wrap_superimpose	z_kick [m]
e_tot [eV]	pz_kick	x1_limit [m]	
ele_origin	ref_origin	x2_limit [m]	

15.28 Girder Element Attributes

alias	dz_origin [m]	tilt [rad]	y_offset [m]
descrip	is_on	tilt_tot [rad]	y_offset_tot [m]
dphi_origin [rad]	l [m]	type	y_pitch [rad]
dpsi_origin [rad]	origin_ele	x_offset [m]	y_pitch_tot [rad]
dtheta_origin [rad]	origin_ele_ref_pt	x_offset_tot [m]	z_offset [m]
dx_origin [m]	ref_tilt [rad]	x_pitch [rad]	z_offset_tot [m]
dy_origin [m]	ref_tilt_tot [rad]	x_pitch_tot [rad]	

15.29 Group Element Attributes

accordion_edge [m]	gang	slave	x_knot
alias	interpolation	start_edge	y_knot
descrip	is_on	type	
end_edge [m]	s_position [m]	var	

15.30 Hybrid Element Attributes

alias	e_tot [eV]	ref_origin	x1_limit [m]
aperture [m]	ele_origin	ref_time_start [sec]	x2_limit [m]
aperture_at	l [m]	reference	x_limit [m]
aperture_type	mat6_calc_method	spin_tracking_method	y1_limit [m]
create_jumbo_slave	offset [m]	superimpose	y2_limit [m]
delta_e_ref [eV]	offset_moves_aperture	tracking_method	y_limit [m]
delta_ref_time [sec]	p0c [eV]	type	
descrip	ptc_integration_type	wrap_superimpose	

15.31 Instrument, Monitor, and Pipe Element Attributes

alias	integrator_order	spin_fringe_on	x_offset [m]
aperture [m]	is_on	spin_tracking_method	x_offset_calib [m]
aperture_at	l [m]	sr_wake	x_offset_tot [m]
aperture_type	lord_pad1 [m]	sr_wake_file	x_pitch [rad]
bl_hkick [T*m]	lord_pad2 [m]	superimpose	x_pitch_tot [rad]
bl_vkick [T*m]	lr_freq_spread [Hz]	symplectify	y1_limit [m]
create_jumbo_slave	lr_self_wake_on	taylor_map_includes_offsets	y2_limit [m]
crunch [rad]	lr_wake	tilt [rad]	y_dispersion_calib [m]
crunch_calib [rad]	lr_wake_file	tilt_calib [rad]	y_dispersion_err [m]
csr_ds_step [m]	mat6_calc_method	tilt_tot [rad]	y_gain_calib [m]
csr_method	n_sample	tracking_method	y_gain_err [m]
de_eta_meas	noise	type	y_limit [m]
delta_ref_time [sec]	num_steps	vkick	y_offset [m]
descrip	offset [m]	wall	y_offset_calib [m]
ds_step [m]	offset_moves_aperture	wrap_superimpose	y_offset_tot [m]
e_tot [eV]	osc_amplitude [m]	x1_limit [m]	y_pitch [rad]
ele_origin	p0c [eV]	x2_limit [m]	y_pitch_tot [rad]
field_calc	ptc_integration_type	x_dispersion_calib [m]	z_offset [m]
field_overlaps	ref_origin	x_dispersion_err [m]	z_offset_tot [m]
fringe_at	ref_time_start [sec]	x_gain_calib [m]	
fringe_type	reference	x_gain_err [m]	
hkick	space_charge_method	x_limit [m]	

15.32 Kicker Element Attributes

a0 - a20, b0 - b20	fringe_type	ptc_integration_type	wall
alias	gen_gradients	r0_elec [m]	wrap_superimpose
aperture [m]	grid_field	r0_mag [m]	x1_limit [m]
aperture_at	h_displace [m]	ref_origin	x2_limit [m]
aperture_type	hkick	ref_time_start [sec]	x_limit [m]
bl_hkick [T*m]	integrator_order	reference	x_offset [m]
bl_vkick [T*m]	is_on	scale_multipoles	x_offset_tot [m]
cartesian_map	l [m]	space_charge_method	x_pitch [rad]
create_jumbo_slave	lord_pad1 [m]	spin_fringe_on	x_pitch_tot [rad]
csr_ds_step [m]	lord_pad2 [m]	spin_tracking_method	y1_limit [m]
csr_method	lr_freq_spread [Hz]	sr_wake	y2_limit [m]
cylindrical_map	lr_self_wake_on	sr_wake_file	y_limit [m]
delta_ref_time [sec]	lr_wake	superimpose	y_offset [m]
descrip	lr_wake_file	symplectify	y_offset_tot [m]
ds_step [m]	mat6_calc_method	taylor_map_includes_offsets	y_pitch [rad]
e_tot [eV]	multipoles_on	tilt [rad]	y_pitch_tot [rad]
ele_origin	num_steps	tilt_tot [rad]	z_offset [m]
field_calc	offset [m]	tracking_method	z_offset_tot [m]
field_master	offset_moves_aperture	type	
field_overlaps	p0c [eV]	v_displace [m^3]	
fringe_at	ptc_canonical_coords	vkick	

15.33 Kickers: Hkicker and Vkicker Element Attributes

a0 - a20, b0 - b20	fringe_at	p0c [eV]	wall
alias	fringe_type	ptc_canonical_coords	wrap_superimpose
aperture [m]	gen_gradients	ptc_integration_type	x1_limit [m]
aperture_at	grid_field	ref_origin	x2_limit [m]
aperture_type	integrator_order	ref_time_start [sec]	x_limit [m]
bl_kick [T*m]	is_on	reference	x_offset [m]
cartesian_map	kick	scale_multipoles	x_offset_tot [m]
create_jumbo_slave	l [m]	space_charge_method	x_pitch [rad]
csr_ds_step [m]	lord_pad1 [m]	spin_fringe_on	x_pitch_tot [rad]
csr_method	lord_pad2 [m]	spin_tracking_method	y1_limit [m]
cylindrical_map	lr_freq_spread [Hz]	sr_wake	y2_limit [m]
delta_ref_time [sec]	lr_self_wake_on	sr_wake_file	y_limit [m]
descrip	lr_wake	superimpose	y_offset [m]
ds_step [m]	lr_wake_file	symplectify	y_offset_tot [m]
e_tot [eV]	mat6_calc_method	taylor_map_includes_offsets	y_pitch [rad]
ele_origin	multipoles_on	tilt [rad]	y_pitch_tot [rad]
field_calc	num_steps	tilt_tot [rad]	z_offset [m]
field_master	offset [m]	tracking_method	z_offset_tot [m]
field_overlaps	offset_moves_aperture	type	

15.34 Lcavity Element Attributes

a0 - a20, b0 - b20	field_autoscale	n_rf_steps	tracking_method
alias	field_calc	num_steps	type
aperture [m]	field_master	offset [m]	vkick
aperture_at	field_overlaps	offset_moves_aperture	voltage [Volt]
aperture_type	fringe_at	p0c [eV]	voltage_err [Volt]
autoscale_amplitude	fringe_type	p0c_start [eV]	voltage_tot [Volt]
autoscale_phase	gen_gradients	phi0 [rad/2pi]	wall
bl_hkick [T*m]	gradient [eV/m]	phi0_autoscale [rad/2pi]	wrap_superimpose
bl_vkick [T*m]	gradient_err [eV/m]	phi0_err [rad/2pi]	x1_limit [m]
bs_field [T]	gradient_tot [eV/m]	phi0_multipass [rad/2pi]	x2_limit [m]
cartesian_map	grid_field	ptc_integration_type	x_limit [m]
cavity_type	hkick	ref_origin	x_offset [m]
coupler_angle [rad]	integrator_order	ref_time_start [sec]	x_offset_tot [m]
coupler_at	is_on	reference	x_pitch [rad]
coupler_phase [rad/2pi]	ks [1/m]	rf_frequency [Hz]	x_pitch_tot [rad]
coupler_strength	l [m]	rf_wavelength [m]	y1_limit [m]
create_jumbo_slave	l_active [m]	scale_multipoles	y2_limit [m]
csr_ds_step [m]	longitudinal_mode	space_charge_method	y_limit [m]
csr_method	lord_pad1 [m]	spin_fringe_on	y_offset [m]
cylindrical_map	lord_pad2 [m]	spin_tracking_method	y_offset_tot [m]
delta_ref_time [sec]	lr_freq_spread [Hz]	sr_wake	y_pitch [rad]
descrip	lr_self_wake_on	sr_wake_file	y_pitch_tot [rad]
ds_step [m]	lr_wake	superimpose	z_offset [m]
e_loss [eV]	lr_wake_file	symplectify	z_offset_tot [m]
e_tot [eV]	mat6_calc_method	taylor_map_includes_offsets	
e_tot_start [eV]	multipoles_on	tilt [rad]	
ele_origin	n_cell	tilt_tot [rad]	

15.35 Lens Element Attributes

alias	mat6_calc_method	tilt [rad]	x_pitch [rad]
aperture [m]	offset [m]	tilt_tot [rad]	x_pitch_tot [rad]
aperture_at	offset_moves_aperture	tracking_method	y1_limit [m]
aperture_type	p0c [eV]	type	y2_limit [m]
create_jumbo_slave	ptc_integration_type	wall	y_limit [m]
delta_ref_time [sec]	radius [m]	wrap_superimpose	y_offset [m]
descrip	ref_origin	x1_limit [m]	y_offset_tot [m]
e_tot [eV]	ref_time_start [sec]	x2_limit [m]	y_pitch [rad]
ele_origin	reference	x_limit [m]	y_pitch_tot [rad]
focal_strength [1/m]	spin_tracking_method	x_offset [m]	z_offset [m]
l [m]	superimpose	x_offset_tot [m]	z_offset_tot [m]

15.36 Line Statement Attributes

alpha_a	default_tracking_species	geometry	ref_time [sec]
alpha_b	deta_dpz_x [m]	high_energy_space_charge_on	s [m]
beta_a [m]	deta_dpz_y [m]	inherit_from_fork	spin_dn_dpz_x
beta_b [m]	detap_dpz_x	live_branch	spin_dn_dpz_y
cmat_11	detap_dpz_y	mode_flip	spin_dn_dpz_z
cmat_12 [m]	e_tot [eV]	p0c [eV]	theta_position [rad]
cmat_21 [1/m]	e_tot_start [eV]	p0c_start [eV]	x_position [m]
cmat_22	eta_x [m]	particle	y_position [m]
dalpha_dpz_a	eta_y [m]	phi_a [rad]	z_position [m]
dalpha_dpz_b	eta_z [m]	phi_b [rad]	
dbeta_dpz_a [m]	etap_x	phi_position [rad]	
dbeta_dpz_b [m]	etap_y	psi_position [rad]	

15.37 Marker Element Attributes

alias	lr_wake_file	tilt [rad]	x_pitch [rad]
aperture [m]	mat6_calc_method	tilt_calib [rad]	x_pitch_tot [rad]
aperture_at	n_sample	tilt_tot [rad]	y1_limit [m]
aperture_type	noise	tracking_method	y2_limit [m]
create_jumbo_slave	offset [m]	type	y_dispersion_calib [m]
crunch [rad]	offset_moves_aperture	wall	y_dispersion_err [m]
crunch_calib [rad]	osc_amplitude [m]	wrap_superimpose	y_gain_calib [m]
de_eta_meas	p0c [eV]	x1_limit [m]	y_gain_err [m]
delta_ref_time [sec]	ptc_integration_type	x2_limit [m]	y_limit [m]
descrip	ref_origin	x_dispersion_calib [m]	y_offset [m]
e_tot [eV]	ref_species	x_dispersion_err [m]	y_offset_calib [m]
ele_origin	ref_time_start [sec]	x_gain_calib [m]	y_offset_tot [m]
is_on	reference	x_gain_err [m]	y_pitch [rad]
l [m]	spin_tracking_method	x_limit [m]	y_pitch_tot [rad]
lr_freq_spread [Hz]	sr_wake	x_offset [m]	z_offset [m]
lr_self_wake_on	sr_wake_file	x_offset_calib [m]	z_offset_tot [m]
lr_wake	superimpose	x_offset_tot [m]	

15.38 Mask Element Attributes

alias	mode	tilt_tot [rad]	y1_limit [m]
aperture [m]	offset [m]	tracking_method	y2_limit [m]
aperture_at	offset_moves_aperture	type	y_limit [m]
aperture_type	p0c [eV]	wall	y_offset [m]
create_jumbo_slave	ptc_integration_type	wrap_superimpose	y_offset_tot [m]
delta_ref_time [sec]	ref_origin	x1_limit [m]	y_pitch [rad]
descrip	ref_time_start [sec]	x2_limit [m]	y_pitch_tot [rad]
e_tot [eV]	ref_wavelength [m]	x_limit [m]	z_offset [m]
ele_origin	reference	x_offset [m]	z_offset_tot [m]
field_scale_factor	spin_tracking_method	x_offset_tot [m]	
is_on	superimpose	x_pitch [rad]	
mat6_calc_method	tilt [rad]	x_pitch_tot [rad]	

15.39 Match Element Attributes

alias	c22_mat1	l [m]	spin_tracking_method
alpha_a0	create_jumbo_slave	mat6_calc_method	superimpose
alpha_a1	delta_ref_time [sec]	matrix	tracking_method
alpha_b0	delta_time [sec]	mode_flip0	type
alpha_b1	descrip	mode_flip1	wrap_superimpose
aperture [m]	dphi_a [rad]	offset [m]	x0 [m]
aperture_at	dphi_b [rad]	offset_moves_aperture	x1 [m]
aperture_type	e_tot [eV]	p0c [eV]	x1_limit [m]
beta_a0 [m]	ele_origin	ptc_integration_type	x2_limit [m]
beta_a1 [m]	eta_x0 [m]	px0	x_limit [m]
beta_b0 [m]	eta_x1 [m]	px1	y0 [m]
beta_b1 [m]	eta_y0 [m]	py0	y1 [m]
c11_mat0	eta_y1 [m]	py1	y1_limit [m]
c11_mat1	etap_x0	pz0	y2_limit [m]
c12_mat0 [m]	etap_x1	pz1	y_limit [m]
c12_mat1 [m]	etap_y0	recalc	z0 [m]
c21_mat0 [1/m]	etap_y1	ref_origin	z1 [m]
c21_mat1 [1/m]	is_on	ref_time_start [sec]	
c22_mat0	kick0	reference	

15.40 Mirror Element Attributes

alias	l [m]	spin_tracking_method	x_pitch [rad]
aperture [m]	mat6_calc_method	superimpose	x_pitch_tot [rad]
aperture_at	offset [m]	tilt [rad]	y1_limit [m]
aperture_type	offset_moves_aperture	tilt_tot [rad]	y2_limit [m]
create_jumbo_slave	p0c [eV]	tracking_method	y_limit [m]
critical_angle [rad]	ptc_integration_type	type	y_offset [m]
curvature	ref_origin	use_reflectivity_table	y_offset_tot [m]
curvature_x0_y2 [1/m]	ref_tilt [rad]	wall	y_pitch [rad]
delta_ref_time [sec]	ref_tilt_tot [rad]	wrap_superimpose	y_pitch_tot [rad]
descrip	ref_time_start [sec]	x1_limit [m]	z_offset [m]
displacement	ref_wavelength [m]	x2_limit [m]	z_offset_tot [m]
e_tot [eV]	reference	x_limit [m]	
ele_origin	reflectivity_table	x_offset [m]	
graze_angle [rad]	segmented	x_offset_tot [m]	

15.41 Multilayer_Mirror Element Attributes

alias	graze_angle [rad]	reference	x_limit [m]
aperture [m]	l [m]	segmented	x_offset [m]
aperture_at	mat6_calc_method	spin_tracking_method	x_offset_tot [m]
aperture_type	material_type	superimpose	x_pitch [rad]
create_jumbo_slave	n_cell	tilt [rad]	x_pitch_tot [rad]
curvature	offset [m]	tilt_tot [rad]	y1_limit [m]
curvature_x0_y2 [1/m]	offset_moves_aperture	tracking_method	y2_limit [m]
d1_thickness [m]	p0c [eV]	type	y_limit [m]
d2_thickness [m]	ptc_integration_type	v1_unitcell [m^3]	y_offset [m]
delta_ref_time [sec]	ref_origin	v2_unitcell [m^3]	y_offset_tot [m]
descrip	ref_tilt [rad]	wall	y_pitch [rad]
displacement	ref_tilt_tot [rad]	wrap_superimpose	y_pitch_tot [rad]
e_tot [eV]	ref_time_start [sec]	x1_limit [m]	z_offset [m]
ele_origin	ref_wavelength [m]	x2_limit [m]	z_offset_tot [m]

15.42 Multipole Element Attributes

alias	k0l - k20l, t0 - t20	spin_tracking_method	x_offset [m]
aperture [m]	k0l_status	superimpose	x_offset_tot [m]
aperture_at	l [m]	tilt [rad]	y1_limit [m]
aperture_type	mat6_calc_method	tilt_tot [rad]	y2_limit [m]
create_jumbo_slave	offset [m]	tracking_method	y_limit [m]
delta_ref_time [sec]	offset_moves_aperture	type	y_offset [m]
descrip	p0c [eV]	wall	y_offset_tot [m]
e_tot [eV]	ptc_integration_type	wrap_superimpose	z_offset [m]
ele_origin	ref_origin	x1_limit [m]	z_offset_tot [m]
field_master	ref_time_start [sec]	x2_limit [m]	
is_on	reference	x_limit [m]	

15.43 Octupole Element Attributes

a0 - a20, b0 - b20	fringe_at	ptc_canonical_coords	wall
alias	fringe_type	ptc_integration_type	wrap_superimpose
aperture [m]	gen_gradients	r0_elec [m]	x1_limit [m]
aperture_at	grid_field	r0_mag [m]	x2_limit [m]
aperture_type	hkick	ref_origin	x_limit [m]
b3_gradient [T/m^3]	integrator_order	ref_time_start [sec]	x_offset [m]
bl_hkick [T*m]	is_on	reference	x_offset_tot [m]
bl_vkick [T*m]	k3 [1/m^4]	scale_multipoles	x_pitch [rad]
cartesian_map	l [m]	space_charge_method	x_pitch_tot [rad]
create_jumbo_slave	lord_pad1 [m]	spin_fringe_on	y1_limit [m]
csr_ds_step [m]	lord_pad2 [m]	spin_tracking_method	y2_limit [m]
csr_method	lr_freq_spread [Hz]	sr_wake	y_limit [m]
cylindrical_map	lr_self_wake_on	sr_wake_file	y_offset [m]
delta_ref_time [sec]	lr_wake	superimpose	y_offset_tot [m]
descrip	lr_wake_file	symplectify	y_pitch [rad]
ds_step [m]	mat6_calc_method	taylor_map_includes_offsets	y_pitch_tot [rad]
e_tot [eV]	multipoles_on	tilt [rad]	z_offset [m]
ele_origin	num_steps	tilt_tot [rad]	z_offset_tot [m]
field_calc	offset [m]	tracking_method	
field_master	offset_moves_aperture	type	
field_overlaps	p0c [eV]	vkick	

15.44 Overlay Element Attributes

alias	interpolation	type	y_knot
descrip	is_on	var	
gang	slave	x_knot	

15.45 Parameter Statement Attributes

absolute_time_tracking	high_energy_space_charge_on	n_part	photon_type
default_tracking_species	lattice	no_end_marker	ptc_exact_misalign
e_tot [eV]	lattice_type	p0c [eV]	ptc_exact_model
electric_dipole_moment	live_branch	parser_make_xfer_mats	ran_seed
geometry	machine	particle	taylor_order

15.46 Particle_Start Statement Attributes

e_photon [eV]	field_y	pz	spin_z
emittance_a [m*rad]	phase_x [rad]	sig_pz	t [sec]
emittance_b [m*rad]	phase_y [rad]	sig_z [m]	x [m]
emittance_z [m*rad]	px	spin_x	y [m]
field_x	py	spin_y	z [m]

15.47 Patch Element Attributes

alias	e_tot_start [eV]	ref_origin	wrap_superimpose
aperture [m]	ele_origin	ref_time_start [sec]	x1_limit [m]
aperture_at	field_calc	reference	x2_limit [m]
aperture_type	flexible	space_charge_method	x_limit [m]
create_jumbo_slave	l [m]	spin_tracking_method	x_offset [m]
csr_ds_step [m]	mat6_calc_method	superimpose	x_pitch [rad]
csr_method	offset [m]	t_offset [sec]	y1_limit [m]
delta_ref_time [sec]	offset_moves_aperture	tilt [rad]	y2_limit [m]
descrip	p0c [eV]	tracking_method	y_limit [m]
downstream_ele_dir	p0c_set [eV]	type	y_offset [m]
e_tot [eV]	p0c_start [eV]	upstream_ele_dir	y_pitch [rad]
e_tot_offset [eV]	ptc_integration_type	user_sets_length	z_offset [m]
e_tot_set [eV]	ref_coords	wall	

15.48 Photon_Init Element Attributes

alias	energy_distribution	sig_vx [m/s]	x1_limit [m]
aperture [m]	energy_probability_curve	sig_vy [m/s]	x2_limit [m]
aperture_at	l [m]	sig_x [m]	x_limit [m]
aperture_type	mat6_calc_method	sig_y [m]	x_offset [m]
create_jumbo_slave	offset [m]	sig_z [m]	x_offset_tot [m]
delta_ref_time [sec]	offset_moves_aperture	spatial_distribution	x_pitch [rad]
descrip	p0c [eV]	spin_tracking_method	x_pitch_tot [rad]
ds_slice [m]	physical_source	superimpose	y1_limit [m]
e2_center [eV]	ptc_integration_type	tilt [rad]	y2_limit [m]
e2_probability	ref_origin	tilt_tot [rad]	y_limit [m]
e_center [eV]	ref_time_start [sec]	tracking_method	y_offset [m]
e_center_relative_to_ref	ref_wavelength [m]	transverse_sigma_cut	y_offset_tot [m]
e_field_x [V/m]	reference	type	y_pitch [rad]
e_field_y [V/m]	scale_field_to_one	velocity_distribution	y_pitch_tot [rad]
e_tot [eV]	sig_e [eV]	wall	z_offset [m]
ele_origin	sig_e2 [eV]	wrap_superimpose	z_offset_tot [m]

15.49 Pickup Element Attributes

alias	mat6_calc_method	tilt_tot [rad]	y1_limit [m]
aperture [m]	num_steps	tracking_method	y2_limit [m]
aperture_at	offset [m]	type	y_limit [m]
aperture_type	offset_moves_aperture	wall	y_offset [m]
create_jumbo_slave	p0c [eV]	wrap_superimpose	y_offset_tot [m]
delta_ref_time [sec]	ptc_integration_type	x1_limit [m]	y_pitch [rad]
descrip	ref_origin	x2_limit [m]	y_pitch_tot [rad]
ds_step [m]	ref_time_start [sec]	x_limit [m]	z_offset [m]
e_tot [eV]	reference	x_offset [m]	z_offset_tot [m]
ele_origin	spin_tracking_method	x_offset_tot [m]	
is_on	superimpose	x_pitch [rad]	
l [m]	tilt [rad]	x_pitch_tot [rad]	

15.50 Quadrupole Element Attributes

a0 - a20, b0 - b20	fq1 [m]	offset_moves_aperture	type
alias	fq2 [m]	p0c [eV]	vkick
aperture [m]	fringe_at	ptc_canonical_coords	wall
aperture_at	fringe_type	ptc_integration_type	wrap_superimpose
aperture_type	gen_gradients	r0_elec [m]	x1_limit [m]
b1_gradient [T/m]	grid_field	r0_mag [m]	x2_limit [m]
bl_hkick [T*m]	hkick	ref_origin	x_limit [m]
bl_vkick [T*m]	integrator_order	ref_time_start [sec]	x_offset [m]
cartesian_map	is_on	reference	x_offset_tot [m]
create_jumbo_slave	k1 [1/m^2]	scale_multipoles	x_pitch [rad]
csr_ds_step [m]	l [m]	space_charge_method	x_pitch_tot [rad]
csr_method	lord_pad1 [m]	spin_fringe_on	y1_limit [m]
cylindrical_map	lord_pad2 [m]	spin_tracking_method	y2_limit [m]
delta_ref_time [sec]	lr_freq_spread [Hz]	sr_wake	y_limit [m]
descrip	lr_self_wake_on	sr_wake_file	y_offset [m]
ds_step [m]	lr_wake	superimpose	y_offset_tot [m]
e_tot [eV]	lr_wake_file	symplectify	y_pitch [rad]
ele_origin	mat6_calc_method	taylor_map_includes_offsets	y_pitch_tot [rad]
field_calc	multipoles_on	tilt [rad]	z_offset [m]
field_master	num_steps	tilt_tot [rad]	z_offset_tot [m]
field_overlaps	offset [m]	tracking_method	

15.51 RFCavity Element Attributes

a0 - a20, b0 - b20	field_autoscale	n_cell	tilt [rad]
alias	field_calc	n_rf_steps	tilt_tot [rad]
aperture [m]	field_overlaps	num_steps	tracking_method
aperture_at	fringe_at	offset [m]	type
aperture_type	fringe_type	offset_moves_aperture	vkick
autoscale_amplitude	gen_gradients	p0c [eV]	voltage [Volt]
autoscale_phase	gradient [eV/m]	phi0 [rad/2pi]	wall
bl_hkick [T*m]	grid_field	phi0_autoscale [rad/2pi]	wrap_superimpose
bl_vkick [T*m]	harmon	phi0_multipass [rad/2pi]	x1_limit [m]
cartesian_map	harmon_master	ptc_integration_type	x2_limit [m]
cavity_type	hkick	ref_origin	x_limit [m]
coupler_angle [rad]	integrator_order	ref_time_start [sec]	x_offset [m]
coupler_at	is_on	reference	x_offset_tot [m]
coupler_phase [rad/2pi]	l [m]	rf_frequency [Hz]	x_pitch [rad]
coupler_strength	l_active [m]	rf_wavelength [m]	x_pitch_tot [rad]
create_jumbo_slave	longitudinal_mode	scale_multipoles	y1_limit [m]
csr_ds_step [m]	lord_pad1 [m]	space_charge_method	y2_limit [m]
csr_method	lord_pad2 [m]	spin_fringe_on	y_limit [m]
cylindrical_map	lr_freq_spread [Hz]	spin_tracking_method	y_offset [m]
delta_ref_time [sec]	lr_self_wake_on	sr_wake	y_offset_tot [m]
descrip	lr_wake	sr_wake_file	y_pitch [rad]
ds_step [m]	lr_wake_file	superimpose	y_pitch_tot [rad]
e_tot [eV]	mat6_calc_method	symplectify	z_offset [m]
ele_origin	multipoles_on	taylor_map_includes_offsets	z_offset_tot [m]

15.52 RF_Bend Element Attributes

alias	g [1/m]	phi0 [rad/2pi]	type
angle [rad]	grid_field	phi0_multipass [rad/2pi]	vkick
aperture [m]	harmon	ptc_integration_type	wall
aperture_at	harmon_master	ref_origin	wrap_superimpose
aperture_type	hkick	ref_tilt [rad]	x1_limit [m]
b_field [T]	integrator_order	ref_tilt_tot [rad]	x2_limit [m]
bl_hkick [T*m]	is_on	ref_time_start [sec]	x_limit [m]
bl_vkick [T*m]	l [m]	reference	x_offset [m]
create_jumbo_slave	l_chord [m]	rf_frequency [Hz]	x_offset_tot [m]
csr_ds_step [m]	l_rectangle [m]	rf_wavelength [m]	x_pitch [rad]
csr_method	l_sagitta [m]	rho [m]	x_pitch_tot [rad]
delta_ref_time [sec]	lord_pad1 [m]	roll [rad]	y1_limit [m]
descrip	lord_pad2 [m]	roll_tot [rad]	y2_limit [m]
ds_step [m]	lr_freq_spread [Hz]	space_charge_method	y_limit [m]
e_tot [eV]	lr_self_wake_on	spin_fringe_on	y_offset [m]
ele_origin	lr_wake	spin_tracking_method	y_offset_tot [m]
fiducial_pt	lr_wake_file	sr_wake	y_pitch [rad]
field_calc	mat6_calc_method	sr_wake_file	y_pitch_tot [rad]
field_master	num_steps	superimpose	z_offset [m]
field_overlaps	offset [m]	symplectify	z_offset_tot [m]
fringe_at	offset_moves_aperture	taylor_map_includes_offsets	
fringe_type	p0c [eV]	tracking_method	

15.53 Ramper Element Attributes

alias	is_on	var
descrip	slave	x_knot
interpolation	type	y_knot

15.54 Sad_Mult Element Attributes

a0 - a20, b0 - b20	fb2 [m]	ref_origin	x_limit [m]
alias	field_calc	ref_time_start [sec]	x_offset [m]
aperture [m]	fq1 [m]	reference	x_offset_mult [m]
aperture_at	fq2 [m]	rho [m]	x_offset_tot [m]
aperture_type	fringe_at	space_charge_method	x_pitch [rad]
bs_field [T]	fringe_type	spin_fringe_on	x_pitch_tot [rad]
create_jumbo_slave	integrator_order	spin_tracking_method	y1_limit [m]
csr_ds_step [m]	is_on	superimpose	y2_limit [m]
csr_method	ks [1/m]	symplectify	y_limit [m]
delta_ref_time [sec]	l [m]	taylor_map_includes_offsets	y_offset [m]
descrip	lord_pad1 [m]	tilt [rad]	y_offset_mult [m]
ds_step [m]	lord_pad2 [m]	tilt_tot [rad]	y_offset_tot [m]
e1 [rad]	mat6_calc_method	tracking_method	y_pitch [rad]
e2 [rad]	num_steps	type	y_pitch_tot [rad]
e_tot [eV]	offset [m]	wall	z_offset [m]
ele_origin	offset_moves_aperture	wrap_superimpose	z_offset_tot [m]
eps_step_scale [m]	p0c [eV]	x1_limit [m]	
fb1 [m]	ptc_integration_type	x2_limit [m]	

15.55 Sample Element Attributes

alias	l [m]	spin_tracking_method	x_offset_tot [m]
aperture [m]	mat6_calc_method	superimpose	x_pitch [rad]
aperture_at	material_type	tilt [rad]	x_pitch_tot [rad]
aperture_type	mode	tilt_tot [rad]	y1_limit [m]
create_jumbo_slave	offset [m]	tracking_method	y2_limit [m]
curvature	offset_moves_aperture	type	y_limit [m]
curvature_x0_y2 [1/m]	p0c [eV]	wall	y_offset [m]
delta_ref_time [sec]	ptc_integration_type	wrap_superimpose	y_offset_tot [m]
descrip	ref_origin	x1_limit [m]	y_pitch [rad]
displacement	ref_time_start [sec]	x2_limit [m]	y_pitch_tot [rad]
e_tot [eV]	reference	x_limit [m]	z_offset [m]
ele_origin	segmented	x_offset [m]	z_offset_tot [m]

15.56 Sextupole Element Attributes

a0 - a20, b0 - b20	fringe_at	ptc_canonical_coords	wall
alias	fringe_type	ptc_integration_type	wrap_superimpose
aperture [m]	gen_gradients	r0_elec [m]	x1_limit [m]
aperture_at	grid_field	r0_mag [m]	x2_limit [m]
aperture_type	hkick	ref_origin	x_limit [m]
b2_gradient [T/m^2]	integrator_order	ref_time_start [sec]	x_offset [m]
bl_hkick [T*m]	is_on	reference	x_offset_tot [m]
bl_vkick [T*m]	k2 [1/m^3]	scale_multipoles	x_pitch [rad]
cartesian_map	l [m]	space_charge_method	x_pitch_tot [rad]
create_jumbo_slave	lord_pad1 [m]	spin_fringe_on	y1_limit [m]
csr_ds_step [m]	lord_pad2 [m]	spin_tracking_method	y2_limit [m]
csr_method	lr_freq_spread [Hz]	sr_wake	y_limit [m]
cylindrical_map	lr_self_wake_on	sr_wake_file	y_offset [m]
delta_ref_time [sec]	lr_wake	superimpose	y_offset_tot [m]
descrip	lr_wake_file	symplectify	y_pitch [rad]
ds_step [m]	mat6_calc_method	taylor_map_includes_offsets	y_pitch_tot [rad]
e_tot [eV]	multipoles_on	tilt [rad]	z_offset [m]
ele_origin	num_steps	tilt_tot [rad]	z_offset_tot [m]
field_calc	offset [m]	tracking_method	
field_master	offset_moves_aperture	type	
field_overlaps	p0c [eV]	vkick	

15.57 Sol_Quad Element Attributes

a0 - a20, b0 - b20	field_overlaps	offset_moves_aperture	type
alias	fringe_at	p0c [eV]	vkick
aperture [m]	fringe_type	ptc_canonical_coords	wall
aperture_at	gen_gradients	ptc_integration_type	wrap_superimpose
aperture_type	grid_field	r0_elec [m]	x1_limit [m]
b1_gradient [T/m]	hkick	r0_mag [m]	x2_limit [m]
bl_hkick [T*m]	integrator_order	ref_origin	x_limit [m]
bl_vkick [T*m]	is_on	ref_time_start [sec]	x_offset [m]
bs_field [T]	k1 [1/m^2]	reference	x_offset_tot [m]
cartesian_map	ks [1/m]	scale_multipoles	x_pitch [rad]
create_jumbo_slave	l [m]	space_charge_method	x_pitch_tot [rad]
csr_ds_step [m]	lord_pad1 [m]	spin_fringe_on	y1_limit [m]
csr_method	lord_pad2 [m]	spin_tracking_method	y2_limit [m]
cylindrical_map	lr_freq_spread [Hz]	sr_wake	y_limit [m]
delta_ref_time [sec]	lr_self_wake_on	sr_wake_file	y_offset [m]
descrip	lr_wake	superimpose	y_offset_tot [m]
ds_step [m]	lr_wake_file	symplectify	y_pitch [rad]
e_tot [eV]	mat6_calc_method	taylor_map_includes_offsets	y_pitch_tot [rad]
ele_origin	multipoles_on	tilt [rad]	z_offset [m]
field_calc	num_steps	tilt_tot [rad]	z_offset_tot [m]
field_master	offset [m]	tracking_method	

15.58 Solenoid Element Attributes

a0 - a20, b0 - b20	fringe_at	p0c [eV]	type
alias	fringe_type	ptc_canonical_coords	vkick
aperture [m]	gen_gradients	ptc_integration_type	wall
aperture_at	grid_field	r0_elec [m]	wrap_superimpose
aperture_type	hkick	r0_mag [m]	x1_limit [m]
bl_hkick [T*m]	integrator_order	r_solenoid [m]	x2_limit [m]
bl_vkick [T*m]	is_on	ref_origin	x_limit [m]
bs_field [T]	ks [1/m]	ref_time_start [sec]	x_offset [m]
cartesian_map	l [m]	reference	x_offset_tot [m]
create_jumbo_slave	l_soft_edge [m]	scale_multipoles	x_pitch [rad]
csr_ds_step [m]	lord_pad1 [m]	space_charge_method	x_pitch_tot [rad]
csr_method	lord_pad2 [m]	spin_fringe_on	y1_limit [m]
cylindrical_map	lr_freq_spread [Hz]	spin_tracking_method	y2_limit [m]
delta_ref_time [sec]	lr_self_wake_on	sr_wake	y_limit [m]
descrip	lr_wake	sr_wake_file	y_offset [m]
ds_step [m]	lr_wake_file	superimpose	y_offset_tot [m]
e_tot [eV]	mat6_calc_method	symplectify	y_pitch [rad]
ele_origin	multipoles_on	taylor_map_includes_offsets	y_pitch_tot [rad]
field_calc	num_steps	tilt [rad]	z_offset [m]
field_master	offset [m]	tilt_tot [rad]	z_offset_tot [m]
field_overlaps	offset_moves_aperture	tracking_method	

15.59 Taylor Element Attributes

alias	mat6_calc_method	symplectify	x_pitch [rad]
aperture [m]	offset [m]	taylor_map_includes_offsets	x_pitch_tot [rad]
aperture_at	offset_moves_aperture	tilt [rad]	x_ref [m]
aperture_type	p0c [eV]	tilt_tot [rad]	y1_limit [m]
create_jumbo_slave	ptc_integration_type	tracking_method	y2_limit [m]
delta_e_ref [eV]	px_ref	tt<out><n1><n2>...	y_limit [m]
delta_ref_time [sec]	py_ref	type	y_offset [m]
descrip	pz_ref	wall	y_offset_tot [m]
e_tot [eV]	ref_orbit	wrap_superimpose	y_pitch [rad]
ele_origin	ref_origin	x1_limit [m]	y_pitch_tot [rad]
is_on	ref_time_start [sec]	x2_limit [m]	y_ref [m]
l [m]	reference	x_limit [m]	z_offset [m]
lord_pad1 [m]	spin_tracking_method	x_offset [m]	z_offset_tot [m]
lord_pad2 [m]	superimpose	x_offset_tot [m]	z_ref [m]

15.60 Thick_Multipole Element Attributes

a0 - a20, b0 - b20	fringe_at	ptc_canonical_coords	wrap_superimpose
alias	fringe_type	ptc_integration_type	x1_limit [m]
aperture [m]	gen_gradients	ref_origin	x2_limit [m]
aperture_at	grid_field	ref_time_start [sec]	x_limit [m]
aperture_type	hkick	reference	x_offset [m]
bl_hkick [T*m]	integrator_order	scale_multipoles	x_offset_tot [m]
bl_vkick [T*m]	is_on	space_charge_method	x_pitch [rad]
cartesian_map	l [m]	spin_fringe_on	x_pitch_tot [rad]
create_jumbo_slave	lord_pad1 [m]	spin_tracking_method	y1_limit [m]
csr_ds_step [m]	lord_pad2 [m]	sr_wake	y2_limit [m]
csr_method	lr_freq_spread [Hz]	sr_wake_file	y_limit [m]
cylindrical_map	lr_self_wake_on	superimpose	y_offset [m]
delta_ref_time [sec]	lr_wake	symplectify	y_offset_tot [m]
descrip	lr_wake_file	taylor_map_includes_offsets	y_pitch [rad]
ds_step [m]	mat6_calc_method	tilt [rad]	y_pitch_tot [rad]
e_tot [eV]	multipoles_on	tilt_tot [rad]	z_offset [m]
ele_origin	num_steps	tracking_method	z_offset_tot [m]
field_calc	offset [m]	type	
field_master	offset_moves_aperture	vkick	
field_overlaps	p0c [eV]	wall	

15.61 Wiggler and Undulator Element Attributes

a0 - a20, b0 - b20	fringe_type	offset_moves_aperture	tracking_method
alias	g_max [1/m]	osc_amplitude [m]	type
aperture [m]	gen_gradients	p0c [eV]	vkick
aperture_at	grid_field	polarity	wall
aperture_type	hkick	ptc_canonical_coords	wrap_superimpose
b_max [T]	integrator_order	ptc_integration_type	x1_limit [m]
bl_hkick [T*m]	is_on	r0_elec [m]	x2_limit [m]
bl_vkick [T*m]	k1x [1/m^2]	r0_mag [m]	x_limit [m]
cartesian_map	k1y [1/m^2]	ref_origin	x_offset [m]
create_jumbo_slave	kx [1/m]	ref_time_start [sec]	x_offset_tot [m]
csr_ds_step [m]	l [m]	reference	x_pitch [rad]
csr_method	l_period [m]	scale_multipoles	x_pitch_tot [rad]
cylindrical_map	lord_pad1 [m]	space_charge_method	y1_limit [m]
delta_ref_time [sec]	lord_pad2 [m]	spin_fringe_on	y2_limit [m]
delta_ref_time_user_set	lr_freq_spread [Hz]	spin_tracking_method	y_limit [m]
descrip	lr_self_wake_on	sr_wake	y_offset [m]
ds_step [m]	lr_wake	sr_wake_file	y_offset_tot [m]
e_tot [eV]	lr_wake_file	superimpose	y_pitch [rad]
ele_origin	mat6_calc_method	symplectify	y_pitch_tot [rad]
field_calc	multipoles_on	taylor_map_includes_offsets	z_offset [m]
field_master	n_period	term	z_offset_tot [m]
field_overlaps	num_steps	tilt [rad]	
fringe_at	offset [m]	tilt_tot [rad]	

Part II

Conventions and Physics

Chapter 16

Coordinates and Physical Layout

Bmad uses three coordinate systems as illustrated in Fig. 16.1. First, the `global` (also called "`floor`") coordinates are independent of the accelerator. Such things as the building the accelerator is in may be described using `global` coordinates.

It is not convenient to describe the position of the beam using the `global` coordinate system so a "`local`" coordinate system is used (§16.1). This curvilinear coordinate system defines the nominal position of the lattice elements. The relationship between the `local` and `global` coordinate systems is described in §16.2.

The "nominal" position of a lattice element is the position of the element without what are called "misalignments" (that is, position and orientation shifts). Each lattice element has "`element body`" coordinates which are attached to the physical element. That is, the electric and magnetic fields of an element are described with respect to `element` coordinates. If there are no misalignments, the `element` coordinates are aligned with the `local` coordinates. The transformation between `local` and `element` coordinates is given in §16.3.

When discussing `local` vs `element` coordinates, it can be less confusing to use the name "`laboratory`" cordinates instead of `local` coordinates. The $x = y = 0$ curved line of the laboratory coordinate system is known as the "`reference orbit`".

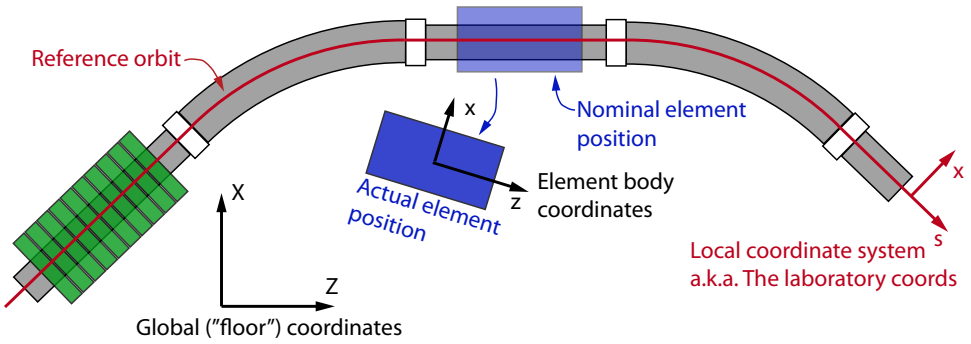

Figure 16.1: The three coordinate systems used by *Bmad*: The `global` (or "`floor`") coordinate system is independent of the accelerator. The `local` curvilinear coordinate system follows the bends of the accelerator. Each lattice element has `element body` coordinates which, if the element is not "misaligned" is the same as the `local` coordinates. The $x = y = 0$ curved line of the laboratory coordinate system is known as the "reference orbit".

16.1 Laboratory Coordinates and Reference Orbit

16.1.1 The Reference Orbit

The `local reference orbit` is the curved path used to define a coordinate system for describing a particle's position as shown in Fig. 16.2. The reference orbit is also used for orientating lattice elements in space. At a given time t, a particle's position can be described by a point $\mathcal{O}$ on the reference orbit a distance s relative to the reference orbit's zero position plus a transverse (x, y) offset. The point $\mathcal{O}$ on the reference orbit is used as the origin of the local (x, y, z) coordinate system with the z–axis tangent to the reference orbit. The z–axis will generally be pointing in the direction of increasing s (Fig. 16.2A) but, as discussed below, will point counter to s for elements that are `reversed` (Fig. 16.2B). The x and y–axes are perpendicular to the reference orbit and, by construction, the particle is always at $z = 0$. The coordinate system so constructed is called the **''local coordinate system''** or sometimes the **''laboratory coordinate system''** when there is need to distinguish it from the **''element coordinate system''** (§16) which is attached to the physical element. There is a separate reference orbit for each branch (§2.2) of a lattice.

Notice that, in a `wiggler`, the reference orbit, which is a straight line, does *not* correspond to the orbit that any actual particle could travel. Typically the physical element is centered with respect to the reference curve. However, by specifying offsets, pitches or a tilt (See §5.6), the physical element may be arbitrarily shifted with respect to its reference curve. Shifting a physical magnet with respect to its reference curve generally means that the reference curve does *not* correspond to the orbit that any actual particle could travel.

Do not confuse this reference orbit (which defines the local coordinate system) with the "map reference orbit" about which the transfer maps are calculated (§36.2). The former is fixed by the lattice while the latter can be any arbitrary orbit. For the map reference orbit, *Bmad* uses the closed orbit for lattice branches with a closed geometry and, for lattice branches with an open geometry, the map reference orbit is the orbit computed starting at the user defined initial position at the beginning of the branch (§10.4).

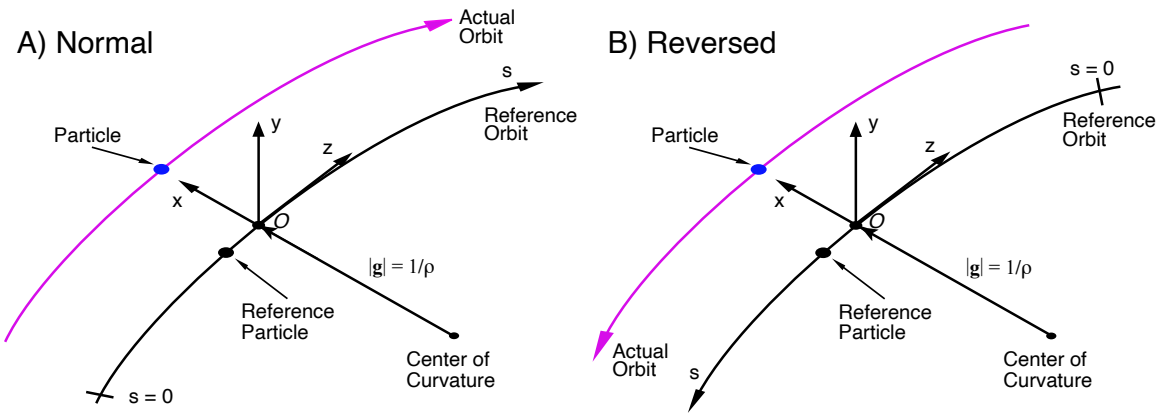

Figure 16.2: The local reference coordinate system. By construction, a particle's z coordinate is zero. This is not to be confused with the phase space z coordinate (§17.2). The curvature vector $\mathbf{g}$ lies in the x-y plane and has a magnitude of $1/\rho$ where ρ is the bending radius. A) The z-axis will normally be parallel to the s-axis. B) For `reversed` elements it will be antiparallel. In both cases, the particle and reference particle are traveling in the direction of greater s.

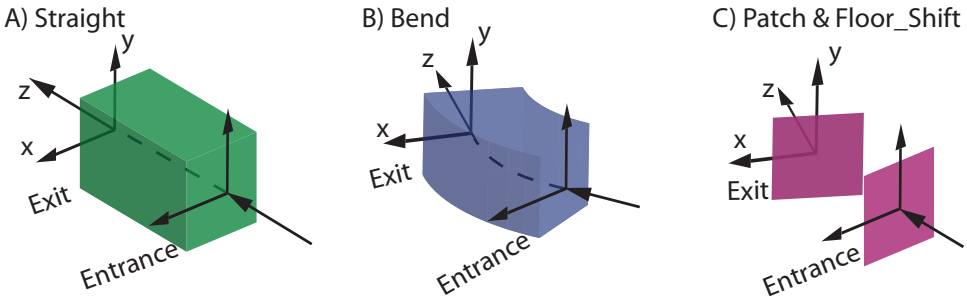

Figure 16.3: Lattice elements can be imagined as "LEGO blocks" which fit together to form the reference orbit along with the laboratory coordinate system. How elements join together is determined in part by their entrance and exit coordinate frames. A) For straight line elements the entrance and exit frames are colinear. B) For bends elements, the two frames are rotated with respect to each other. C) For `patch` and `floor_shift` elements the exit frame may be arbitrarily positioned with respect to the entrance frame.

16.1.2 Element Entrance and Exit Coordinates

One way of thinking about the reference orbit and the laboratory coordinates is to imagine that each element is like a LEGO block with an "`entrance`" and an "`exit`" coordinate frame as illustrated in Fig. 16.3[1]. These coordinate frames are attached to the element. that is, things like electric and magnetic fields, apertures, etc., are described with respect to the entrance and exit coordinates. Thus, for example, the `e1` edge of a bend (§4.5) is always at the `entrance` face and the `e2` is always at the `exit` face. Most elements have a "straight" geometry as shown in Fig. 16.3A. That is, the reference orbit through the element is a straight line segment with the x and y axes always pointing in the same direction. For a `bend` element (§4.5), the reference orbit is a segment of a circular arc as shown in Fig. 16.3B. With the `ref_tilt` parameter of a bend set to zero, the rotation axis between the entrance and exit frames is parallel to the y-axis (§16.2). For `patch` (§4.41), and `floor_shift` (§4.21) elements, the exit face can can arbitrarily oriented with respect to the entrance end. In this case, the reference orbit between the entrance and exit faces is not defined.

16.1.3 Reference Orbit and Laboratory Coordinates Construction

Assuming for the moment that there are no `fiducial` elements present, the construction of the reference orbit starts at the `beginning_ele` element (§4.4) at the start of a branch. If the branch is a `root` branch (§2.3), The orientation of the beginning element within the global coordinate system (§16) can be set via the appropriate positioning statements (§10.4). If the branch is not a `root` branch, the position of the beginning element is determined by the position of the `fork` or `photon_fork` element from which the branch forks from. Unless set otherwise in the lattice file, $s = 0$ at the `beginning_ele` element.

If there are `fiducial` elements, the laboratory coordinates are constructed beginning at these elements.

Once the beginning element in a branch is positioned, succeeding elements are concatenated together to form the laboratory coordinates. All elements have an "`upstream`" and a "`downstream`" end as shown in Fig. 16.4A. The `downstream` end of an element is always farther (at greater s) from the beginning element than the `upstream` end of the element. Particles travel in the $+s$ direction, so particles will enter an element at the upstream end and exit at the downstream end.

[1] Thanks to Dan Abell for this analogy.

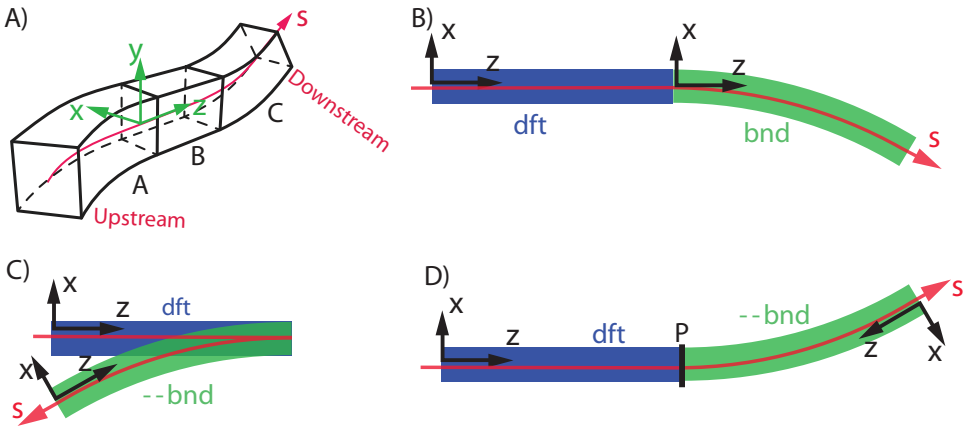

Figure 16.4: A) The laboratory coordinates are constructed by connecting the `downstream` reference frame of one element with the `upstream` reference frame of the next element in the branch. Coordinates shown is for the mating of element `A` to element B. B) Example with drift element `dft` followed by a bend `bnd`. Both elements are unreversed. C) Similar to (B) but in this case element `bnd` is reversed. D) Similar to (C) but in this case a reflection patch has been added in between `dft` and `bnd`. In (B), (C), and (D) the (x, z) coordinates are drawn at the `entrance` end of the elements. The y coordinate is always out of the page.

Normally, the `upstream` end is the element's `entrance` end (Fig. 16.3) and the `downstream` end is the element's `exit` exit. This corresponds to particles entering at the entrance end and exiting the element at the exit end. However, if an element is reversed (§7.4), the element's `exit` end will be `upstream` end and the element's `entrance` end will be the `downstream` end. That is, for a reversed element, particles will enter at the element's `exit` end and will exit at the `entrance` end.

The procedure to connect elements together to form the laboratory coordinates is to mate the downstream reference frame of the element with the upstream reference frame of the next element in the branch so that, without misalignments, the (x, y, z) axes coincide[2]. This is illustrated in Fig. 16.4. Fig. 16.4A shows the general situation with the downstream frame of element `A` mated to the upstream frame of element `B`. Figures 16.4B-C show branches constructed from the following lattice file:

```
DFT: drift, l = 2
BND: sbend, l = 2, g = pi/12
P: patch, x_pitch = pi          ! Reflection patch.
B_line: line = (DFT, BND)       ! No reversal.
C_line: line = (DFT, --BND)     ! Illegal. Do not use!
D_line: line = (DFT, P, --BND)  ! Valid.
```

The (x, z) coordinates are drawn at the entrance end of the elements and z will always point towards the element's exit end. Fig. 16.4B shows the branch constructed from `B_line` containing an unreversed drift named `dft` connected to an unreversed bend named `bnd`. Fig. 16.4C shows the branch constructed from `C_line`. This is like `B_line` except here element `bnd` is reversed. This gives an unphysical situation since a particle traveling through `dft` will "fall off" when it gets to the end. Fig. 16.4D shows the branch constructed from `D_line`. Here a "`reflection`" patch `P` (§16.2.6) has been added to get a plausible geometry. The patch rotates the coordinate system around the y-axis by 180°(leaving the y-axis invariant). It is always the case that a reflection patch is needed between reversed and unreversed elements

[2]If there are misalignments, the `entrance` and `exit` frames will move with the element. However, the `upstream` and `downstream` frames, along with the reference orbit and laboratory coordinates, will not move.

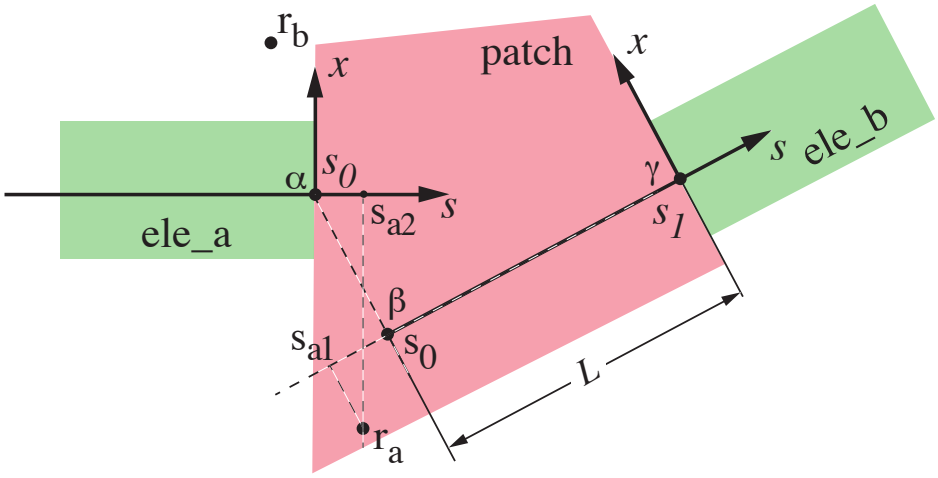

Figure 16.5: The local reference coordinates in a `patch` element. The `patch` element, shown schematically as an irregular quadrilateral, is sandwiched between elements `ele_a` and `ele_b`. `L` is the length of the `patch`. In this example, the `patch` has a finite `x_pitch`.

Notes:

- If the first element after the `beginning_ele` element at the start of a branch is reversed, the `beginning_ele` element will be marked as reversed so that a reflection patch is not needed in this circumstance.
- Irrespective of whether elements are reversed or not, the laboratory (x, y, z) coordinate system at all s-positions will always be a right-handed coordinate system.
- Care must be take when using reversed elements. For example, if the field of the `bnd` element in `B_line` is appropriate for, say, electrons, that is, electrons will be bent in a clockwise fashion going through `bnd`, then an electron going through `D_line` will be lost in the bend (the y-axis and hence the field is in the same direction for both cases so electrons will still be bent in a clockwise fashion but with `D_line` a particle needs to be bent counterclockwise to get through the bend). To get a particle through the bend, positrons must be used.
- A reflection patch that rotated the coordinates, for example, around the x-axis by 180°(by setting `y_pitch` to `pi`) would also produce a plausible geometry.

16.1.4 Patch Element Local Coordinates

Generally, if a particle is reasonably near the reference orbit, there is a one-to-one mapping between the particle's position and (x, y, s) coordinates. A `patch` (§4.41) elements with a non-zero `x_pitch` or non-zero `y_pitch` breaks the one-to-one mapping. This is illustrated in Fig. 16.5. The `patch` element, shown schematically as an, irregular quadrilateral, is sandwiched between elements `ele_a` and `ele_b`. The local coordinate system with origin at α are the coordinates at the end of `ele_a`. The coordinates at the end of the `patch` has its origin labeled γ. By convention, the length of the patch `L` is taken to be the longitudinal distance from α to γ with the `patch`'s exit coordinates defining the longitudinal direction. The "beginning" point of the `patch` on the reference orbit a distance `L` from point γ is labeled β in the figure.

In the local (x, y, s) coordinate system a particle at α will have some value $s = s_0$. A particle at point β will have the same value $s = s_0$ and a particle at γ will have $s = s_1 = s_0 + L$. A particle at point r_a in Fig. 16.5 illustrates the problem of assigning (x, y, s) coordinates to a given position. If the particle is considered to be within the region of `ele_a`, the particle's s position will be s_{a2} which is greater than the value s_0 at the exit end of the element. This contradicts the expectation that particles within `ele_a` will have $s \leq s_0$. If, on the other hand, the particle is considered to be within the `patch` region, the particle's s position will be s_{a1} which is less than the value s_0 at the entrance to the patch. This contradicts the expectation that a particles within the `patch` will have $s \geq s_0$.

To resolve this problem, *Bmad* considers a particle at position r_a to be within the `patch` region. This means that there is, in theory, no lower limit to the s-position that a particle in the `patch` region can have. This also implies that there is a discontinuity in the s-position of a particle crossing the exit face of `ele1`. Typically, when particles are translated from the exit face of one element to the exit face of the next, this `patch` problem does not appear. It only appears when the track between faces is considered.

Notice that a particle at position r_b in Fig. 16.5 can simultaneously be considered to be in either `ele_a` or the `patch`. While this creates an ambiguity it does not complicate tracking.

16.2 Global Coordinates

The Cartesian `global` coordinate system, also called the 'floor" coordinate system, is the coordinate system "attached to the earth" that is used to describe the local coordinate system. Following the *MAD* convention, the `global` coordinate axis are labeled (X, Y, Z). Conventionally, Y is the "vertical"

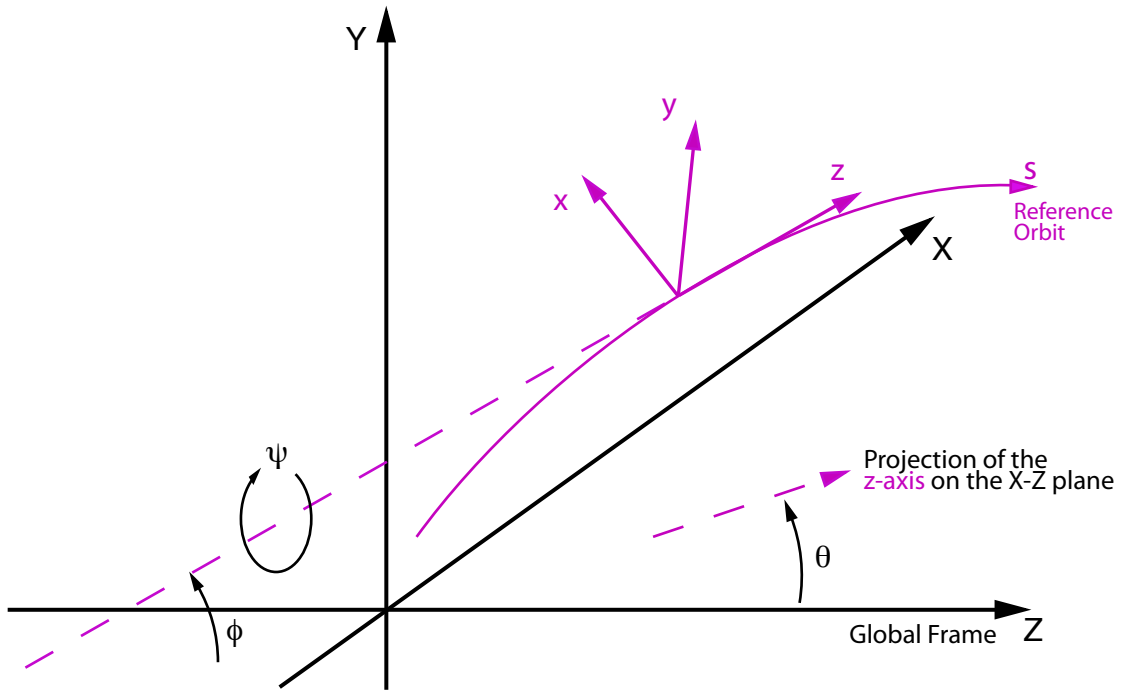

Figure 16.6: The local (reference) coordinate system (purple), which is a function of s along the reference orbit, is described in the global coordinate system (black) by a position $(X(s), Y(s), Z(s))$ and and by angles $\theta(s)$, $\phi(s)$, and $\psi(s)$.

coordinate and (X, Z) are the "horizontal" coordinates. To describe how the local coordinate system is oriented within the global coordinate system, each point on the s-axis of the local coordinate system is characterized by its (X, Y, Z) position and by three angles $\theta(s)$, $\phi(s)$, and $\psi(s)$ that describe the orientation of the local coordinate axes as shown in Fig. 16.6. These three angles are defined as follows:

$\theta(s)$ **Azimuth (yaw) angle:** Angle in the (X, Z) plane between the Z–axis and the projection of the z–axis onto the (X, Z) plane. Corresponds to the `x_pitch` element attribute (§5.6). A positive angle of $\theta = \pi/2$ corresponds to the projected z–axis pointing in the positive X direction.

$\phi(s)$ **Pitch (elevation) angle:** Angle between the z–axis and the (X, Z) plane. Corresponds to the `y_pitch` element attribute (§5.6). A positive angle of $\phi = \pi/2$ corresponds to the z–axis pointing in the positive Y direction.

$\psi(s)$ **Roll angle:** Angle of the x–axis with respect to the line formed by the intersection of the (X, Z) plane with the (x, y) plane. Corresponds to the `tilt` element attribute (§5.6). A positive ψ forms a right–handed screw with the z–axis.

By default, at $s = 0$, the reference orbit's origin coincides with the (X, Y, Z) origin and the x, y, and z axes correspond to the X, Y, and Z axes respectively. If the lattice has no vertical bends (the `ref_tilt` parameter (§4.5) of all bends are zero), the y–axis will always be in the vertical Y direction and the x–axis will lie in the horizontal (X, Z) plane. In this case, θ decreases as one follows the reference orbit when going through a horizontal bend with a positive bending angle. This corresponds to x pointing radially outward. Without any vertical bends, the Y and y axes will coincide, and ϕ and ψ will both be zero. The `beginning` statement (§10.4) in a lattice file can be use to override these defaults.

Following *MAD*, the global position of an element is characterized by a vector $\mathbf{V}$

$$\mathbf{V} = \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} \tag{16.1}$$

The orientation of an element is described by a unitary rotation matrix $\mathbf{W}$. The column vectors of $\mathbf{W}$ are the unit vectors spanning the local coordinate axes in the order (x, y, z). $\mathbf{W}$ can be expressed in terms of the orientation angles θ, ϕ, and ψ via the formula

$$\begin{aligned} \mathbf{W} &= \mathbf{R}_y(\theta)\, \mathbf{R}_{-x}(\phi)\, \mathbf{R}_z(\psi) \\ &= \begin{pmatrix} \cos\theta\cos\psi - \sin\theta\sin\phi\sin\psi & -\cos\theta\sin\psi - \sin\theta\sin\phi\cos\psi & \sin\theta\cos\phi \\ \cos\phi\sin\psi & \cos\phi\cos\psi & \sin\phi \\ -\cos\theta\sin\phi\sin\psi - \sin\theta\cos\psi & \sin\theta\sin\psi - \cos\theta\sin\phi\cos\psi & \cos\theta\cos\phi \end{pmatrix} \end{aligned} \tag{16.2}$$

where

$$\mathbf{R}_y(\theta) = \begin{pmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{pmatrix}, \quad \mathbf{R}_{-x}(\phi) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\phi & \sin\phi \\ 0 & -\sin\phi & \cos\phi \end{pmatrix}, \quad \mathbf{R}_z(\psi) = \begin{pmatrix} \cos\psi & -\sin\psi & 0 \\ \sin\psi & \cos\psi & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{16.3}$$

Notice that $\mathbf{R}_{-x}(\phi)$, for positive ϕ, represents a rotation around the negative x-axis. Also notice that these are Tait-Bryan angles and not Euler angles.

An alternative representation of the $\mathbf{W}$ matrix (or any other rotation matrix) is to specify the axis $\mathbf{u}$ (normalized to 1) and angle of rotation β

$$\mathbf{W} = \begin{pmatrix} \cos\beta + u_x^2\,(1-\cos\beta) & u_x\,u_y\,(1-\cos\beta) - u_z\sin\beta & u_x\,u_z\,(1-\cos\beta) + u_y\sin\beta \\ u_y\,u_x\,(1-\cos\beta) + u_z\sin\beta & \cos\beta + u_y^2\,(1-\cos\beta) & u_y\,u_z\,(1-\cos\beta) - u_x\sin\beta \\ u_z\,u_x\,(1-\cos\beta) - u_y\sin\beta & u_z\,u_y\,(1-\cos\beta) + u_x\sin\beta & \cos\beta + u_z^2\,(1-\cos\beta) \end{pmatrix} \tag{16.4}$$

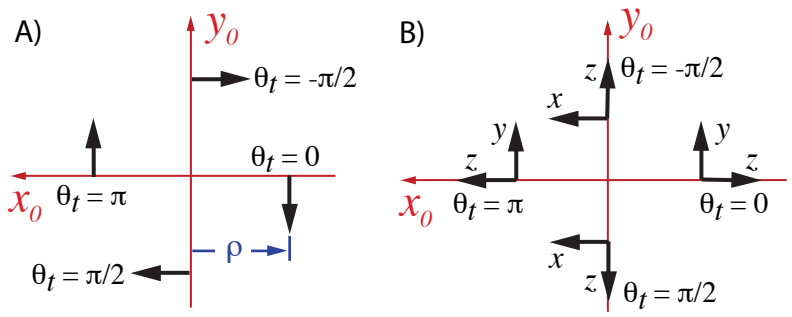

Figure 16.7: A) Rotation axes (bold arrows) for four different `ref_tilt` angles of $\theta_t = 0, \pm\pi/2$, and π. (x_0, y_0, z_0) are the local coordinates at the entrance end of the bend with the z_0 axis being directed into the page. Any rotation axis will be displaced by a distance of the bend radius `rho` from the origin. B) The (x, y, z) coordinates at the exit end of the bend for the same four `ref_tilt` angles. In this case the bend angle is taken to be $\pi/2$.

16.2.1 Lattice Element Positioning

Bmad, again following *MAD*, computes $\mathbf{V}$ and $\mathbf{W}$ by starting at the first element of the lattice and iteratively using the equations

$$\mathbf{V}_i = \mathbf{W}_{i-1}\,\mathbf{L}_i + \mathbf{V}_{i-1}, \tag{16.5}$$
$$\mathbf{W}_i = \mathbf{W}_{i-1}\,\mathbf{S}_i \tag{16.6}$$

$\mathbf{L}_i$ is the displacement vector for the i^{th} element and matrix $\mathbf{S}_i$ is the rotation of the local reference system of the exit end with respect to the entrance end. For clarity, the subscript i in the equations below will be dripped. For all elements whose reference orbit through them is a straight line, the corresponding $\mathbf{L}$ and $\mathbf{S}$ are

$$\mathbf{L} = \begin{pmatrix} 0 \\ 0 \\ L \end{pmatrix}, \quad \mathbf{S} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{16.7}$$

Where L is the length of the element.

For a `bend`, the axis of rotation is dependent upon the bend's `ref_tilt` angle (§5.6) as shown in Fig. 16.7A. The axis of rotation points in the negative y_0 direction for `ref_tilt` $= 0$ and is offset by the bend radius `rho`. Here (x_0, y_0, z_0) are the local coordinates at the entrance end of the bend with the z_0 axis being directed into the page in the figure. For a non-zero `ref_tilt`, the rotation axis is itself rotated about the z_0 axis by the value of `ref_tilt`. Fig. 16.7B shows the exit coordinates for four different values of `ref_tilt` and for a bend angle `angle` of $\pi/2$. Notice that for a bend in the horizontal $X - Z$ plane, a positive bend `angle` will result in a decreasing azimuth angle θ.

For a bend, $\mathbf{S}$ is given using Eq. (16.4) with

$$\begin{aligned} \mathbf{u} &= (\sin\theta_t, -\cos\theta_t, 0) \\ \beta &= \alpha_b \end{aligned} \tag{16.8}$$

where θ_t is the `ref_tilt` angle. The $\mathbf{L}$ vector for a `bend` is given by

$$\mathbf{L} = \mathbf{R}_z(\theta_t)\,\widetilde{\mathbf{L}}, \quad \widetilde{\mathbf{L}} = \begin{pmatrix} \rho(\cos\alpha_b - 1) \\ 0 \\ \rho\,\sin\alpha_b \end{pmatrix} \tag{16.9}$$

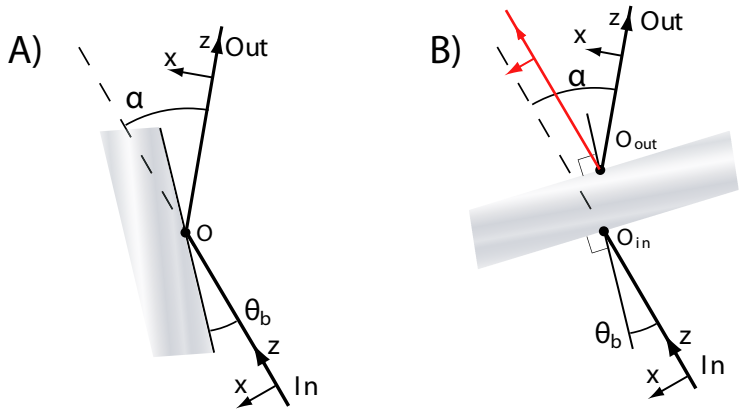

Figure 16.8: Mirror and crystal geometry. The geometry shown here is appropriate for a `ref_tilt` angle of $\theta_t = 0$. θ_g is the bend angle of the incoming (entrance) ray, and α_b is the total bend angle of the reference trajectory. A) Geometry for a mirror or a Bragg crystal. Point $\mathcal{O}$ is the origin of both the local coordinates just before and just after the reflection/diffraction. B) Geometry for a Laue crystal. Point $\mathcal{O}_{out}$ is the origin of the coordinates just after diffraction is displaced from the origin $\mathcal{O}_{in}$ just before diffraction due to the finite thickness of the crystal. here the bend angles are measured with respect to the line that is in the plane of the entrance and exit coordinates and perpendicular to the surface. For Laue diffraction, the user has the option of using the undiffracted beam (shown in red) as the reference trajectory.

where α_b is the bend `angle` (§4.5) and ρ being the bend radius (`rho`). Notice that since $\mathbf{u}$ is perpendicular to z, the curvilinear reference coordinate system has no "torsion". That is, it is a Frenet-Serret coordinate system.

Note: An alternative equation for $\mathbf{S}$ for a bend is

$$\mathbf{S} = \mathbf{R}_z(\theta_t)\, \mathbf{R}_y(-\alpha_b)\, \mathbf{R}_z(-\theta_t) \tag{16.10}$$

The bend transformation above is so constructed that the transformation is equivalent to rotating the local coordinate system around an axis that is perpendicular to the plane of the bend. This rotation axis is invariant under the bend transformation. For example, for $\theta_t = 0$ (or π) the y-axis is the rotation axis and the y-axis of the local coordinates before the bend will be parallel to the y-axis of the local coordinates after the bend as shown in Fig. 16.7. That is, a lattice with only bends with $\theta_t = 0$ or π will lie in the horizontal plane (this assuming that the y-axis starts out pointing along the Y-axis as it does by default). For $\theta_t = \pm\pi/2$, the bend axis is the x-axis. A value of $\theta_t = +\pi/2$ represents a downward pointing bend.

16.2.2 Position Transformation When Transforming Coordinates

A point $\mathbf{Q}_g = (X, Y, Z)$ defined in the global coordinate system, when expressed in the coordinate system defined by $(\mathbf{V}, \mathbf{W})$ is

$$\mathbf{Q}_{VW} = \mathbf{W}^{-1}\,(\mathbf{Q}_g - \mathbf{V}) \tag{16.11}$$

This is essentially the inverse of Eq. (16.5). That is, vectors propagate inversely to the propagation of the coordinate system.

Using Eq. (16.11) with Eqs. (16.5), and (16.6), the transformation of a particle's position $\mathbf{q} = (x, y, z)$ and momentum $\mathbf{P} = (P_x, P_y, P_z)$ when the coordinate frame is transformed from frame $(\mathbf{V}_{i-1}, \mathbf{W}_{i-1})$

to frame $(\mathbf{V}_i, \mathbf{W}_i)$ is

$$\mathbf{q}_i = \mathbf{S}_i^{-1} \left(\mathbf{q}_{i-1} - \mathbf{L}_i\right), \tag{16.12}$$

$$\mathbf{P}_i = \mathbf{S}_i^{-1} \mathbf{P}_{i-1} \tag{16.13}$$

Notice that since $\mathbf{S}$ (and $\mathbf{W}$) is the product of orthogonal rotation matrices, $\mathbf{S}$ is itself orthogonal and its inverse is just the transpose

$$\mathbf{S}^{-1} = \mathbf{S}^T \tag{16.14}$$

16.2.3 Crystal and Mirror Element Coordinate Transformation

A `crystal` element (§4.35) diffracts photons and a `mirror` element (§4.35) reflects them. For a crystal setup for Bragg diffraction, and for a mirror, the reference orbit is modeled as a zero length bend with $\widetilde{\mathbf{L}} = (0, 0, 0)$, as shown in Fig. 16.8A. Shown in the figure is the geometry appropriate for a `ref_tilt` angle of $\theta_t = 0$ (the rotation axis is here the y-axis). Since the mirror or crystal element is modeled to be of zero length, the origin points (marked $\mathcal{O}$ in the figure) of the entrance and exit local coordinates are the same. For Laue diffraction, the only difference is that $\widetilde{\mathbf{L}}$ is non-zero due to the finite thickness of the crystal as shown in Fig. 16.8B. This results in a separation between the entrance coordinate origin $\mathcal{O}_{in}$ and the exit coordinate origin $\mathcal{O}_{out}$.

In all cases, the total bending angle is

$$\begin{aligned}
\alpha_b &= \text{bragg_angle_in} + \text{bragg_angle_out} && \text{! Crystal, graze_angle_in} = 0 \\
\alpha_b &= \text{graze_angle_in} + \text{graze_angle_out} && \text{! Crystal, graze_angle_in} \neq 0 \\
\alpha_b &= 2\,\text{graze_angle} && \text{! Mirror}
\end{aligned} \tag{16.15}$$

With a mirror or Bragg diffraction, the bend angles are measured with respect to the surface plane. With Laue diffraction the bend angles are measured with respect to the line in the bend plane perpendicular to the surface.

For Laue diffraction, the user has the option of using the undiffracted beam (shown in red) as the reference trajectory.

The orientation of the exit coordinates (the local coordinates after the reflection) are only affected by the element's `ref_tilt` and bend angle parameters and is independent of all other parameters such as the radius of curvature of the surface, etc. The local z-axis of the entrance coordinates along with the z-axis of the exit coordinates define a plane which is called the element's `bend plane`. For a mirror, the graze angle is a parameter supplied by the user. For a crystal, the Bragg angles are calculated so that the reference trajectory is in the middle of the Darwin curve. Calculation of the Bragg angles for a crystal is given in Section §27.4.1.

16.2.4 Patch and Floor_Shift Elements Entrance to Exit Transformation

For `patch` (§4.41) and `floor_shift` (§4.21) elements, the shift in the exit end reference coordinates is given by Eqs. (16.5) and (16.6) with

$$\begin{aligned}
\mathbf{L} &= \begin{pmatrix} \text{x_offset} \\ \text{y_offset} \\ \text{z_offset} \end{pmatrix} \\
\mathbf{S} &= \mathbf{R}_y(\text{x_pitch})\, \mathbf{R}_{-x}(\text{y_pitch})\, \mathbf{R}_z(\text{tilt})
\end{aligned} \tag{16.16}$$

The difference here between `patch` and `floor_shift` elements is that, with a `patch` element, the shift is relative to the exit end of the previous element while, for a `floor_shift` element, the shift is relative to the reference point on the origin element specified by the `origin_ele` parameter of the `floor_shift`.

16.2.5 Fiducial and Girder Elements Origin Shift Transformation

For `fiducial` and `girder` elements, the alignment of the reference coordinates with respect to "`origin`" coordinates is analogous to Eqs. (16.16). Explicitly:

$$\mathbf{L} = \begin{pmatrix} \text{dx_origin} \\ \text{dy_origin} \\ \text{dz_origin} \end{pmatrix}$$
$$\mathbf{S} = \mathbf{R}_y(\text{dtheta_origin})\ \mathbf{R}_{-x}(\text{dphi_origin})\ \mathbf{R}_z(\text{dpsi_origin}) \tag{16.17}$$

16.2.6 Reflection Patch

A `Patch` (or a series of patches) that reflects the direction of the z-axis is called a `reflection patch`. By "reflected direction" it is meant that the dot product $\mathbf{z}_1 \cdot \mathbf{z}_2$ is negative where $\mathbf{z}_1$ is the z-axis vector at the `entrance` face and $\mathbf{z}_2$ is the z-axis vector at the `exit` face. This condition is equivalent to the condition that the associated $\mathbf{S}$ matrix (see Eq. (16.16)) satisfy:

$$S(3,3) < 0 \tag{16.18}$$

Using Eq. (16.16) gives, after some simple algebra, this condition is equivalent to

$$\cos(\text{x_pitch})\ \cos(\text{y_pitch}) < 0 \tag{16.19}$$

When there are a series of patches, The transformations of all the patches are concatenated together to form an effective $\mathbf{S}$ which can then be used with Eq. (16.18).

16.3 Transformation Between Laboratory and Element Body Coordinates

The `element body` coordinates are the coordinate system attached to an element. Without any misalignments, where **"misalignments"** are here defined to be any offset, pitch or tilt (§5.6), the `laboratory` coordinates (§16.1.1) and `element body` coordinates are the same. With misalignments, the transformation between `laboratory` and `element body` coordinates depends upon whether the local coordinate system is straight (§16.3.1) or bent (§16.3.2).

When tracking a particle through an element, the particle starts at the `nominal` (§16) upstream end of the element with the particle's position expressed in laboratory coordinates. Tracking from the nominal upstream end to the actual upstream face of the element involves first transforming to element body coordinates (with $s = 0$ in the equations below) and then propagating the particle as in a field free drift space from the particle's starting position to the actual element face. Depending upon the element's orientation, this tracking may involve tracking backwards. Similarly, after a particle has been tracked through the physical element to the actual downstream face, the tracking to the nominal downstream end involves transforming to laboratory coordinates (using $s = L$ in the equations below) and then propagating the particle as in a field free drift space to the nominal downstream edge.

16.3.1 Straight Element Misalignment Transformation

For straight line elements, given a laboratory coordinate frame Λ_s with origin a distance s from the beginning of the element, misalignments will shift the coordinates to a new reference frame denoted E_s. Since misalignments are defined with respect to the middle of the element, the transformation between Λ_s and E_s is a three step process:

$$\Lambda_s \longrightarrow \Lambda_{\rm mid} \longrightarrow E_{\rm mid} \longrightarrow E_s \tag{16.20}$$

where $\Lambda_{\rm mid}$ and $E_{\rm mid}$ are the laboratory and element reference frames at the center of the element.

The first and last transformations from Λ_s to $\Lambda_{\rm mid}$ and from $E_{\rm mid}$ to E_s use Eqs. (16.5), (16.6), and (16.7) with the replacement $L \to L/2 - s$ for the first transformation and $L \to s - L/2$ for the third transformation. The middle transformation, by definition of the offset, pitch and tilt parameters is

$$\begin{aligned} \mathbf{L} &= \begin{pmatrix} \text{x_offset} \\ \text{y_offset} \\ \text{z_offset} \end{pmatrix} \\ \mathbf{S} &= \mathbf{R}_y(\text{x_pitch})\, \mathbf{R}_{-x}(\text{y_pitch})\, \mathbf{R}_z(\text{tilt}) \end{aligned} \tag{16.21}$$

Notice that with this definition of how elements are misaligned, the position of the center of a non-bend misaligned element depends only on the offsets, and is independent of the pitches and tilt.

16.3.2 Bend Element Misalignment Transformation

For `rbend` and `sbend` elements there is no `tilt` attribute. Rather, there is the `roll` attribute and a `ref_tilt` attribute. The latter affects both the reference orbit and the bend position (§5.6.3). Furthermore, `ref_tilt` is calculated with respect to the coordinates at the beginning of the bend while, like straight elements, `roll`, offsets, and pitches are calculated with respect to the center of the bend. The different reference frame used for `ref_tilt` versus everything else means that five transformations are needed to get from the laboratory frame to the element body frame (see Eq. (16.20)). Symbolically:

$$\Lambda_s \longrightarrow \Lambda_{\rm mid} \longrightarrow \Omega_{\rm mid} \longrightarrow \Omega_0 \longrightarrow E_0 \longrightarrow E_s \tag{16.22}$$

The first transformation, Λ_s to $\Lambda_{\rm mid}$, from laboratory coordinates at a distance s from the beginning of the element to laboratory coordinates at the center the bend is a rotation around the center of curvature of the bend and is given by Eqs. (16.5) and (16.6) with Eqs. (16.8) and (16.9) with the substitution $\alpha_b \to (L/2 - s)/\rho$.

The second transformation $\Lambda_{\rm mid}$ to $\Omega_{\rm mid}$ at the center of the element adds in the misalignments (Note that the coordinate frame $\Omega_{\rm mid}$ is neither a laboratory frame or an element frame so hence the use of a different symbol Ω). Explicitly, the $\Lambda_{\rm mid} \longrightarrow \Omega_{\rm mid}$ transformation is

$$\begin{aligned} \mathbf{L} &= \mathbf{L}_{\rm off} + [\mathbf{R}_z(\text{roll}) - \mathbf{1}]\ \mathbf{R}_z(\theta_t)\, \mathbf{R}_y(\alpha_b/2)\, \mathbf{L}_c \\ \mathbf{S} &= \mathbf{R}_y(\text{x_pitch})\, \mathbf{R}_{-x}(\text{y_pitch})\, \mathbf{R}_z(\text{roll}) \end{aligned} \tag{16.23}$$

where

$$\mathbf{L}_c = \begin{pmatrix} \rho(\cos(\alpha_b/2) - 1) \\ 0 \\ \rho\, \sin(\alpha_b/2) \end{pmatrix}, \qquad \mathbf{L}_{\rm off} = \begin{pmatrix} \text{x_offset} \\ \text{y_offset} \\ \text{z_offset} \end{pmatrix} \tag{16.24}$$

The reason why $\mathbf{L}$ has a different form from straight line elements is due to the fact that the axis of rotation for a `roll` is displaced from the z-axis of the coordinate system at the center of the bend (see Fig. 5.3).

The third transformation from Ω_{mid} to Ω_0 is like the first transformation and rotates from the center of the bend to the beginning. Again Eqs. (16.8) and (16.9) are used with the substitution $\alpha_b \to -L/2\rho$.

The fourth transformation Ω_0 to E_0 tilts the reference frame by an amount `ref_tilt`:

$$\mathbf{L} = 0, \quad \mathbf{S} = \mathbf{R}_z(\theta_t) \tag{16.25}$$

The fifth and final transformation, E_0 to E_s, like the first and third, rotates around the center of the bend but in this case, since we are dealing with element coordinates, the `ref_tilt` is ignored. That is, Eqs. (16.8) and (16.9) are used with the substitutions $\theta_t \to 0$ and $\alpha_b \to L/\rho$.

Notice that with this definition of how elements are misaligned, the position of the center of a misaligned element depends only on the offsets and `roll`, and is independent of the pitches and tilt. Also the orientation of an element depends only on the pitches roll, and ref_tilt, and is independent of the offsets.

Chapter 17

Phase Space Coordinates

17.1 Reference Particle, Reference Energy, and Reference Time

The `reference energy` and `reference time` are needed in evaluating the phase space coordinates of charged particles (§17.2).

All lattice elements, except for controller elements, have an associated `reference energy` energy. The reference energy at the start of a lattice's `root branch` (§2.2) is set in the lattice file by setting the reference momentum (`p0c`) or total energy (`E_tot`) using a `parameter` (§10.1) or `beginning` (§10.4) statement. For other branches, the energy at the start of the branch is set using the appropriate line parameter (§10.4) statement.

Note that the reference momentum `p0c` is actually the reference momentum times the speed of light so that the reference momentum has the same unit (eV) as the reference energy.

For most elements, the reference energy is the same as the reference energy of the proceeding element. The following elements are exceptions:

```
custom
em_field
hybrid
lcavity
patch
```

The reference energy of these elements is determined by tracking a particle (the "`reference particle`") through the element with the particle starting on the reference orbit and whose energy is equal to the reference energy. The energy of the particle at the downstream end is the reference energy of the element. Note: Tracking through an element to determine the reference energy is always done with the element turned on independent of the setting of the element's `is_on` (§5.14) parameter. Reference energy tracking is also done ignoring any orientation attributes (§5.6) and errors like `voltage_err`.

Besides the reference energy, lattice elements have an associated `reference time` which is computed, for most elements, by the time-of-flight of the `reference particle` assuming that the reference particle is following the reference orbit. Exceptions are `wiggler` elements which uses the time-of-flight of the actual undulating trajectory. See §5.5 for more details.

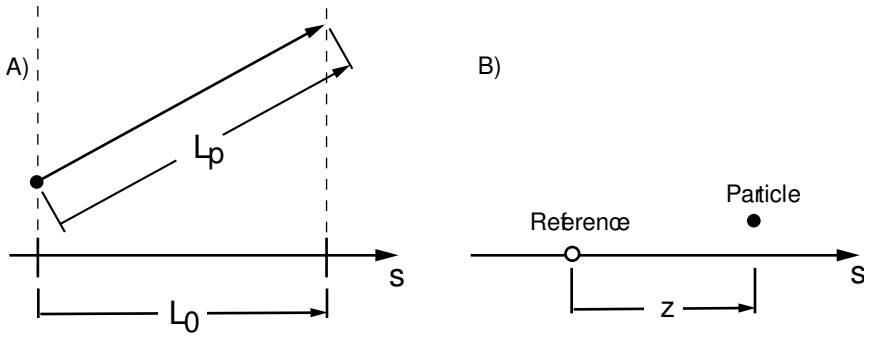

Figure 17.1: Interpreting phase space z at constant velocity: A) The change in z going through an element of length L_0 is $L_0 - L_p$. B) At constant time, z is the longitudinal distance between the reference particle and the particle.

17.2 Charged Particle Phase Space Coordinates

For charged particles (more correctly, for everything but photons (§17.4)), *Bmad* uses the canonical phase space coordinates

$$\mathbf{r}(s) = (x, p_x, y, p_y, z, p_z) \tag{17.1}$$

The longitudinal position s is the independent variable instead of the time. x and y, are the transverse coordinates of the particle as shown in Fig. 16.2A. Note that x and y are independent of the position of the reference particle.

The phase space momenta p_x and p_y are normalized by the reference (sometimes called the design) momentum P_0

$$p_x = \frac{P_x}{P_0}, \qquad p_y = \frac{P_y}{P_0} \tag{17.2}$$

where P_x and P_y are respectively the x and y momentums.

The phase space z coordinate is

$$\begin{aligned} z(s) &= -\beta(s)\, c\, (t(s) - t_0(s)) \\ &\equiv -\beta(s)\, c\, \Delta t(s) \end{aligned} \tag{17.3}$$

$t(s)$ is the time at which the particle is at position s, $t_0(s)$ is the time at which the reference particle is at position s, and β is v/c with v being the particle velocity (and not the reference velocity). The reference particle is, by definition, "synchronized" with elements whose fields are oscillating and therefore the actual fields a particle will see when traveling through such an element will depend upon the particle's phase space z. For example, the energy change of a particle traveling through an `lcavity` (§4.31) or `rfcavity` (§4.46) element is z dependent. Exception: With absolute time tracking (§26.1) fields are tied to the absolute time and not z.

If the particle's velocity is constant, and is the same as the velocity of the reference particle (for example, at high energy where $\beta = 1$ for all particles), then $\beta\, c\, t$ is just the path length. In this case, the change in z going through an element is

$$\Delta z = L_0 - L_p \tag{17.4}$$

where, as shown in Fig. 17.1A, L_0 is the path length of the reference particle (which is just the length of the element) and L_p is the path length of the particle in traversing the element. Another way of interpreting phase space z is that, at constant β, and constant time, z is the longitudinal distance

between the particle and the reference particle as shown in Fig. 17.1B. with positive z indicating that the particle is ahead of the reference particle.

Do not confuse the phase space z with the z that is the particle's longitudinal coordinate in the local reference frame as shown in Fig. 16.2. By construction, this latter z is always zero.

Notice that if a particle gets an instantaneous longitudinal kick so that β is discontinuous then, from Eq. (17.3), phase space z is discontinuous even though the particle itself does not move in space. In general, from Eq. (17.3), The value of z for a particle at s_2 is related to the value of z for the particle at s_1 by

$$z_2 = \frac{\beta_2}{\beta_1}\, z_1 - \beta_2\, c\, (\Delta t_2 - \Delta t_1) \tag{17.5}$$

$\Delta t_2 - \Delta t_1$ can be interpreted as the difference in transit time, between the particle and the reference particle, in going from s_1 to s_2.

The longitudinal phase space momentum p_z is given by

$$p_z = \frac{\Delta P}{P_0} \equiv \frac{P - P_0}{P_0} \tag{17.6}$$

where P is the momentum of the particle. For ultra–relativistic particles p_z can be approximated by

$$p_z = \frac{\Delta E}{E_0} \tag{17.7}$$

where E_0 is the reference energy (energy here always refers to the total energy) and $\Delta E = E - E_0$ is the deviation of the particle's energy from the reference energy. For an `Lcavity` element (§4.31) the reference momentum is *not* constant so the tracking for an `Lcavity` is not canonical.

The relationship, between the phase space momenta and the slopes $x' \equiv dx/ds$ and $y' \equiv dy/ds$ is

$$\begin{aligned} x' &= \frac{p_x}{\sqrt{(1+p_z)^2 - p_x^2 - p_y^2}}\,(1+gx) \\ y' &= \frac{p_y}{\sqrt{(1+p_z)^2 - p_x^2 - p_y^2}}\,(1+gx) \end{aligned} \tag{17.8}$$

$g = 1/\rho$ is the curvature function with ρ being the radius of curvature of the reference orbit and it has been assumed that the bending is in the x–z plane.

With the paraxial approximation, and in the relativistic limit, the change in z with position is

$$\frac{dz}{ds} = -g\,x - \frac{1}{2}(x'^2 + y'^2) \tag{17.9}$$

This shows that in a linac, without any bends, the z of a particle always decreases.

A particle can also have a spin. The spin is characterized by the spinor $\Psi = (\psi_1, \psi_2)^T$ where $\psi_{1,2}$ are complex numbers (§24.1).

17.3 Time-based Phase Space Coordinates

Some specialized routines (for example, time Runge Kutta tracking) use the time t as the independent variable for charged particle tracking. This is useful when particles can reverse direction since the normal z based tracking cannot handle this. Direction reversal can happen, for example, with low energy "dark current" electrons that are generated at the walls of the vacuum chamber.

When the tracking is time based the phase space coordinates are:

$$(x, c\,p_x, y, c\,p_y, z, c\,p_s) \tag{17.10}$$

The positions x, y, and z are the same as with phase space coordinates (§17.2). The momenta are defined as

$$\begin{aligned} cp_x &\equiv mc^2\gamma\beta_x \\ cp_y &\equiv mc^2\gamma\beta_y \\ cp_s &\equiv mc^2\gamma\beta_s, \end{aligned} \tag{17.11}$$

and internally are stored in units of eV.

17.4 Photon Phase Space Coordinates

The phase space coordinates discussed above implicitly assume that particles are traveling longitudinally in only one direction. That is, the sign of the s component of the momentum cannot be determined from the phase space coordinates. This is generally fine for tracking high energy beams of charged particles but for photon tracking this would oftentimes be problematical. For photons, therefore, a different phase space is used:

$$(x, \beta_x, y, \beta_y, z, \beta_z) \tag{17.12}$$

Here $(\beta_x, \beta_y, \beta_z)$ is the normalized photon velocity with

$$\beta_x^2 + \beta_y^2 + \beta_z^2 = 1 \tag{17.13}$$

and (x, y, z) are the reference orbit coordinates with z being the distance from the start of the lattice element the photon is in.

In *Bmad*, the information associated with a photon include its phase space coordinates and time along with the photon energy and four parameters E_x, ϕ_x, and E_y, ϕ_y specifying the intensity and phase of the field along the x and y axes transverse to the direction of propagation. the field in the vicinity of the photon is

$$\begin{aligned} E_x(\mathbf{r}, t) &\sim E_x\, e^{i(k\,(z-z_0)-\omega\,(t-t_{\rm ref})+\phi_x)} \\ E_y(\mathbf{r}, t) &\sim E_y\, e^{i(k\,(z-z_0)-\omega\,(t-t_{\rm ref})+\phi_y)} \end{aligned} \tag{17.14}$$

where z_0 is the photon z position and $t_{\rm ref}$ is the reference time.

The normalization between field and intensity is dependent upon the particular parameters of any given simulation and so must be determined by the program using *Bmad*.

17.5 MAD and Bmad Energy-Time Phase Space Coordinates

MAD uses a different coordinate system where (z, p_z) is replaced by (τ, p_τ) where

$$\tau \equiv -c\,\Delta t = \frac{z}{\beta} \tag{17.15}$$

$$p_\tau \equiv \frac{\Delta E}{P_0 c} = \frac{E}{P_0\,c} - \frac{1}{\beta_0} \tag{17.16}$$

Where E is the total energy, and β_0 is the reference velocity. For highly relativistic particles the two coordinate systems are identical. This difference must be keept in mind when comparing results between *Bmad* and *MAD*. For example, the calculated dispersion will not be the same except in the ultra relativistic limit.

It is sometimes convenient to use energy-time phase space cocordinates in cases where there is an electric field like in an RF cavity. In this case, *Bmad* uses a differnt energy coordinate p_E

$$p_E \equiv \frac{E}{P_0\, c} \tag{17.17}$$

Chapter 18

Electromagnetic Fields

18.1 Magnetostatic Multipole Fields

Start with the assumption that the local magnetic field has no longitudinal component (obviously this assumption does not work with, say, a solenoid). Following *MAD*, ignoring skew fields for the moment, the vertical magnetic field along the $y = 0$ axis is expanded in a Taylor series

$$B_y(x,0) = \sum_n B_n \frac{x^n}{n!} \tag{18.1}$$

Assuming that the reference orbit is locally straight (there are correction terms if the reference orbit is curved (§18.3)), the field is

$$\begin{aligned} B_x &= \quad B_1 y + \quad B_2\, xy \quad\quad + \frac{1}{6}B_3(3x^2y - y^3) \;+\; \ldots \\ B_y &= B_0 + B_1 x + \frac{1}{2}B_2(x^2 - y^2) + \frac{1}{6}B_3(x^3 - 3xy^2) \;+\; \ldots \end{aligned} \tag{18.2}$$

The relation between the field B_n and the normalized field K_n is:

$$K_n \equiv \frac{q\, B_n}{P_0} \tag{18.3}$$

where q is the charge of the reference particle (in units of the elementary charge), and P_0 is the reference momentum (in units of eV/c). Note that P_0/q is sometimes written as $B\rho$. This is just an old notation where ρ is the bending radius of a particle with the reference energy in a field of strength B. Notice that P_0 is the local reference momentum at the element which may not be the same as the reference energy at the beginning of the lattice if there are `lcavity` elements (§4.31) present.

The kicks Δp_x and Δp_y that a particle experiences going through a multipole field is

$$\Delta p_x = \frac{-q\, L\, B_y}{P_0} \tag{18.4}$$

$$= -K_0 L - K_1 L\, x \;+\; \frac{1}{2}K_2 L(y^2 - x^2) + \frac{1}{6}K_3 L(3xy^2 - x^3) \;+\; \ldots$$

$$\Delta p_y = \frac{q\, L\, B_x}{P_0} \tag{18.5}$$

$$= \quad\quad K_1 L\, y + \quad K_2 L\, xy \quad\quad + \frac{1}{6}K_3 L(3x^2y - y^3) \;+\; \ldots$$

A positive $K_1 L$ quadrupole component gives horizontal focusing and vertical defocusing. The general form is

$$\Delta p_x = \sum_{n=0}^{\infty} \frac{K_n L}{n!} \sum_{m=0}^{2m \le n} \binom{n}{2m} (-1)^{m+1} x^{n-2m} y^{2m} \tag{18.6}$$

$$\Delta p_y = \sum_{n=0}^{\infty} \frac{K_n L}{n!} \sum_{m=0}^{2m \le n-1} \binom{n}{2m+1} (-1)^{m} x^{n-2m-1} y^{2m+1} \tag{18.7}$$

where $\binom{a}{b}$ ("a choose b") denotes a binomial coefficient.

The above equations are for fields with a normal component only. If a given multipole field of order n has normal B_n and skew S_n components and is rotated in the (x, y) plane by an angle T_n, the magnetic field at a point (x, y) can be expressed in complex notation as

$$B_y(x, y) + i B_x(x, y) = \frac{1}{n!}(B_n + i S_n)\, e^{-i(n+1)T_n}\, e^{in\theta}\, r^n \tag{18.8}$$

where (r, θ) are the polar coordinates of the point (x, y).

Note that, for compatibility with MAD, the $K0L$ component of a `Multipole` element rotates the reference orbit essentially acting as a zero length bend. This is not true for multipoles of any other type of element.

Instead of using magnitude K_n and rotation angle θ_n, Another representation is using normal $\widetilde{K}_n$ and skew $\widetilde{S}_n$. The conversion between the two are

$$\begin{aligned} \widetilde{K}_n &= K_n \cos((n+1)\, T_n) \\ \widetilde{S}_n &= K_n \sin((n+1)\, T_n) \end{aligned} \tag{18.9}$$

Another representation of the magnetic field used by *Bmad* divides the fields into normal b_n and skew a_n components. In terms of these components the magnetic field for the n^{th} order multipole is

$$\frac{q\, L}{P_0}\, (B_y + i B_x) = (b_n + i a_n)\, (x + iy)^n \tag{18.10}$$

The a_n, b_n representation of multipole fields can be used in elements such as quadrupoles, sextupoles, etc. to allow "error" fields to be represented. The conversion between (a_n, b_n) and $(K_n L, S_n L, T_n)$ is

$$b_n + i a_n = \frac{1}{n!}\, (K_n L + i\, S_n L)\, e^{-i(n+1)T_n} \tag{18.11}$$

In the case where $S_n L = 0$

$$K_n L = n!\, \sqrt{a_n^2 + b_n^2} \tag{18.12}$$

$$\tan[(n+1)T_n] = \frac{-a_n}{b_n} \tag{18.13}$$

To convert a normal magnet (a magnet with no skew component) into a skew magnet (a magnet with no normal component) the magnet should be rotated about its longitudinal axis with a rotation angle of

$$(n+1)T_n = \frac{\pi}{2} \tag{18.14}$$

For example, a normal quadrupole rotated by 45° becomes a skew quadrupole.

The multipole fields can be "`reference energy`" scaled and/or "`element strength`" scaled. Scaling here means that the a_n and b_n values used in tracking are scaled from the input values given in the lattice file.

`Reference energy` scaling is applied if the `field_master` attribute (§5.2) is True for an element so that the multipole values specified in the lattice file are not reference energy normalized

$$[a_n, b_n] \longrightarrow [a_n, b_n] \cdot \frac{q}{P_0} \tag{18.15}$$

`Element strength` scaling is applied when the multipoles are associated with a non `AB_Multipole` element and if the `scale_multipoles` attribute (§5.15) is `True`. This scaling uses a measurement radius r_0 and a scale factor F:

$$[a_n, b_n] \longrightarrow [a_n, b_n] \cdot F \cdot \frac{r_0^{n_{\rm ref}}}{r_0^n} \tag{18.16}$$

r_0 is set by the `r0_mag` attribute of an element. F and $n_{\rm ref}$ are set automatically depending upon the type of element as shown in Table 18.1. The γ_p term is

Element	F	$n_{\rm ref}$
`Elseparator`	$\sqrt{\texttt{Hkick}^2 + \texttt{Vkick}^2}$	0
`Hkicker`	Kick	0
`Kicker,AC_Kicker`	$\sqrt{\texttt{Hkick}^2 + \texttt{Vkick}^2}$	0
`Rbend`	G * L	0
`Sbend`	G * L	0
`Vkicker`	Kick	0
`Wiggler`	$\frac{2\,c\,\texttt{L_pole}\,B_{max}}{\pi\,\texttt{p0c}}$	0
`Quadrupole`	K1 * L	1
`Sol_Quad`	K1 * L	1
`Solenoid`	KS * L	1
`Sextupole`	K2 * L	2
`Octupole`	K3 * L	3

Table 18.1: F and $n_{\rm ref}$ for various elements.

18.2 Electrostatic Multipole Fields

Except for the `elseparator` element, *Bmad* specifies DC electric fields using normal b_{en} and skew a_{en} components (§5.15). The potential ϕ_n for the n^{th} order multipole is

$$\phi_n = -\,\mathrm{Re}\left[\frac{b_{en} - i a_{en}}{n+1}\,\frac{(x+iy)^{n+1}}{r_0^n}\right] \tag{18.17}$$

where r_0 is a "measurement radius" set by the `r0_elec` attribute of an element (§5.15).

The electric field for the n^{th} order multipole is

$$E_x - iE_y = (b_{en} - i a_{en})\,\frac{(x+iy)^n}{r_0^n} \tag{18.18}$$

Notice that the magnetic multipole components a_n and b_n are normalized by the element length, reference charge, and reference momentum (Eq. (18.10)) while their electric counterparts are not.

Using the paraxial approximation, The kick given a particle due to the electric field is

$$\frac{dp_x}{ds} = \frac{q\,E_x}{\beta\,P_0\,c}, \qquad \frac{dp_y}{ds} = \frac{q\,E_y}{\beta\,P_0\,c} \tag{18.19}$$

Where β is the normalized velocity.

18.3 Exact Multipole Fields in a Bend

For static magnetic and electric multipole fields in a bend, the spacial dependence of the field is different from multipole fields in an element with a straight geometry as given by Eqs. (18.10) and (18.18). The analysis of the multipole fields in a bend here follows McMillan[McMill75].

In the rest of this section, normalized coordinates $\widetilde{r} = r/\rho$, $\widetilde{x}/ = x/\rho$, and $\widetilde{y} = y/\rho$ will be used where ρ is the bending radius of the reference coordinate system, r is the distance, in the plane of the bend, from the bend center to the observation point, x is the distance in the plane of the from the reference coordinates to the observation point and y is the distance out-of-plane. With this convention $\widetilde{r} = 1 + \widetilde{x}$.

An electric or magnetic multipole can be characterized by a scalar potential ϕ with the field given by $-\nabla\phi$. The potential is a solution to Laplace's equation

$$\frac{1}{\widetilde{r}} \frac{\partial}{\partial \widetilde{r}} \left(\widetilde{r} \frac{\partial \phi}{\partial \widetilde{r}} \right) + \frac{\partial^2 \phi}{\partial \widetilde{y}^2} = 0 \tag{18.20}$$

As McMillian shows, it is also possible to calculate the magnetic field by constructing the appropriate vector potential. However, from a practical point of view, it is simpler to use the scalar potential for both the magnetic and electric fields.

Solutions to Laplace's equation can be found in form

$$\phi_n^r = \frac{-1}{1+n} \sum_{p=0}^{2p \le n+1} \binom{n+1}{2p} (-1)^p \, F_{n+1-2p}(\widetilde{r}) \, \widetilde{y}^{2p} \tag{18.21}$$

and in the form

$$\phi_n^i = \frac{-1}{1+n} \sum_{p=0}^{2p \le n} \binom{n+1}{2p+1} (-1)^p \, F_{n-2p}(\widetilde{r}) \, \widetilde{y}^{2p+1} \tag{18.22}$$

where $\binom{a}{b}$ ("a choose b") denotes a binomial coefficient, and n is the order number which can range from 0 to infinity.[1]

In Eq. (18.22) the $F_p(\widetilde{r})$ are related by

$$F_{p+2} = (p+1)\,(p+2) \int_1^{\widetilde{r}} \frac{d\widetilde{r}}{\widetilde{r}} \left[\int_1^{\widetilde{r}} d\widetilde{r}\, \widetilde{r}\, F_p \right] \tag{18.23}$$

with the "boundary condition":

$$\begin{aligned} F_0(\widetilde{r}) &= 1 \\ F_1(\widetilde{r}) &= \ln \widetilde{r} \end{aligned} \tag{18.24}$$

This condition ensures that the number of terms in the sums in Eqs. (18.21) and (18.22) are finite. With

[1] Notice that here n is related to m in McMillian's paper by $m = n + 1$. Also note that the ϕ^r and ϕ^i here have a normalization factor that is different from McMillian.

this condition, all the F_p can be constructed:

$$\begin{aligned}
F_1 &= \ln \widetilde{r} = \widetilde{x} - \frac{1}{2}\widetilde{x}^2 + \frac{1}{3}\widetilde{x}^3 - \dots \\
F_2 &= \frac{1}{2}(\widetilde{r}^2 - 1) - \ln \widetilde{r} = \widetilde{x}^2 - \frac{1}{3}\widetilde{x}^3 + \frac{1}{4}\widetilde{x}^4 - \dots \\
F_3 &= \frac{3}{2}[-(\widetilde{r}^2 - 1) + (\widetilde{r}^2 + 1)\ln \widetilde{r}] = \widetilde{x}^3 - \frac{1}{2}\widetilde{x}^4 + \frac{7}{20}\widetilde{x}^5 - \dots \\
F_4 &= 3[\frac{1}{8}(\widetilde{r}^4 - 1) + \frac{1}{2}(\widetilde{r}^2 - 1) - (\widetilde{r}^2 + \frac{1}{2})\ln \widetilde{r}] = \widetilde{x}^4 - \frac{2}{5}\widetilde{x}^5 + \frac{3}{10}\widetilde{x}^6 - \dots \\
&\text{Etc...}
\end{aligned} \tag{18.25}$$

Evaluating these functions near $\widetilde{x} = 0$ using the exact $\widetilde{r}$-dependent functions can be problematical due to round off error. For example, Evaluating $F_4(\widetilde{r})$ at $\widetilde{x} = 10^{-4}$ results in a complete loss of accuracy (no significant digits!) when using double precision numbers. In practice, *Bmad* uses a Padé approximant for $\widetilde{x}$ small enough and then switches to the $\widetilde{r}$-dependent formulas for $\widetilde{x}$ away from zero.

For magnetic fields, the "real" ϕ_n^r solutions will correspond to skew fields and the "imaginary" ϕ_n^i solutions will correspond to normal fields

$$\mathbf{B} = -\frac{P_0}{q\,L}\sum_{n=0}^{\infty} \rho^n \left[a_n\,\widetilde{\nabla}\phi_n^r + b_n\,\widetilde{\nabla}\phi_n^i\right] \tag{18.26}$$

where the gradient derivatives of $\widetilde{\nabla}$ are with respect to the normalized coordinates. In the limit of infinite bending radius ρ, the above equations converge to the straight line solution given in Eq. (18.10).

For electric fields, the "real" solutions will correspond to normal fields and the "imaginary" solutions are used for skew fields

$$\mathbf{E} = -\sum_{n=0}^{\infty} \rho^n \left[a_{en}\,\widetilde{\nabla}\phi_n^i + b_{en}\,\widetilde{\nabla}\phi_n^r\right] \tag{18.27}$$

And this will converge to Eq. (18.18) in the straight line limit.

In the vertical plane, with $\widetilde{x} = 0$, the solutions ϕ_n^r and ϕ_n^i have the same variation in $\widetilde{y}$ as the multipole fields with a straight geometry. For example, the field strength of an $n = 1$ (quadrupole) multipole will be linear in $\widetilde{y}$ for $\widetilde{x} = 0$. However, in the horizontal direction, with $\widetilde{y} = 0$, the multipole field will vary like $dF_2/d\widetilde{x}$ which has terms of all orders in $\widetilde{x}$. In light of this, the solutions ϕ_n^r and ϕ_n^i are called "vertically pure" solutions.

It is possible to construct "horizontally pure" solutions as well. That is, it is possible to construct solutions that in the horizontal plane, with $\widetilde{y} = 0$, behave the same as the corresponding multipole fields with a straight geometry. A straight forward way to do this, for some given multipole of order n, is to construct the horizontally pure solutions, ψ_n^r and ψ_n^i, as linear superpositions of the vertically pure solutions

$$\psi_n^r = \sum_{k=n}^{\infty} C_{nk}\,\phi_k^r, \qquad \psi_n^i = \sum_{k=n}^{\infty} D_{nk}\,\phi_k^i \tag{18.28}$$

with the normalizations $C_{nn} = D_{nn} = 1$. The C_{nk} and D_{nk} are chosen, order by order, so that ψ_n^r and ψ_n^i are horizontally pure. For the real potentials, the C_{nk}, are obtained from a matrix $\mathbf{M}$ where M_{ij} is the coefficient of the $\widetilde{x}^j$ term of $(dF_i/d\widetilde{x})/i$ when F_i is expressed as an expansion in $\widetilde{x}$ (Eq. (18.25)). C_{nk}, $k = 0, \dots, \infty$ are the row vectors of the inverse matrix $\mathbf{M}^{-1}$. For the imaginary potentials, the D_{nk} are constructed similarly but in this case the rows of $\mathbf{M}$ are the coefficients in $\widetilde{x}$ for the functions F_i. To

convert between field strength coefficients, Eqs. (18.26) and (18.27) and Eqs. (18.28) are combined

$$a_n = \sum_{k=n}^{\infty} \frac{1}{\rho^{k-n}} C_{nk}\, \alpha_k, \quad a_{en} = \sum_{k=n}^{\infty} \frac{1}{\rho^{k-n}} D_{nk}\, \alpha_{ek},$$
$$b_n = \sum_{k=n}^{\infty} \frac{1}{\rho^{k-n}} D_{nk}\, \beta_k, \quad b_{en} = \sum_{k=n}^{\infty} \frac{1}{\rho^{k-n}} D_{nk}\, \beta_{ek} \tag{18.29}$$

where α_k, β_k, α_{ek}, and β_{ek} are the corresponding coefficients for the horizontally pure solutions.

When expressed as a function of $\widetilde{r}$ and $\widetilde{y}$, the vertically pure solutions ϕ_n have a finite number of terms (Eqs. (18.21) and (18.22)). On the other hand, the horizontally pure solutions ψ_n have an infinite number of terms.

The vertically pure solutions form a complete set. That is, any given field that satisfies Maxwell's equations and is independent of z can be expressed as a linear combination of ϕ_n^r and ϕ_n^i. Similarly, the horizontally pure solutions form a complete set. [It is, of course, possible to construct other complete sets in which the basis functions are neither horizontally pure nor vertically pure.]

This brings up an important point. To properly simulate a machine, one must first of all understand whether the multipole values that have been handed to you are for horizontally pure multipoles, vertically, pure multipoles, or perhaps the values do not correspond to either horizontally pure nor vertically pure solutions! Failure to understand this point can lead to differing results. For example, the chromaticity induced by a horizontally pure quadrupole field will be different from the chromaticity of a vertically pure quadrupole field of the same strength. With *Bmad*, the `exact_multipoles` (§4.5) attribute of a bend is used to set whether multipole values are for vertically or horizontally pure solutions. [Note to programmers: PTC always assumes coefficients correspond to horizontally pure solutions. The *Bmad* PTC interface will convert coefficients as needed.]

18.4 Map Decomposition of Magnetic and Electric Fields

Electric and magnetic fields can be parameterized as the sum over a number of functions with each function satisfying Maxwell's equations. These functions are also referred to as "maps", "modes", or "terms". *Bmad* has three parameterizations:

```
Cartesian Map              ! §18.5.
Cylindrical Map            ! §18.6
Generalized Gradient Map   ! §18.7
```

These parameterizations are three of the four `field map` parameterizations that *Bmad* defines §5.16.

The `Cartesian map` decomposition involves a set of terms, each term a solution the Laplace equation solved using separation of variables in Cartesian coordinates. This decomposition can be used for DC but not AC fields. See §18.5. for more details. The syntax for specifying the `Cartesian map` decomposition is discussed in §5.16.2.

The `cylindrical map` decomposition can be used for both DC and AC fields. See §18.6 for more details. The syntax for specifying the `cylindrical map` decomposition is discussed in §5.16.3.

The `generalized gradient map` start with the cylindrical map decomposition but then express the field using coefficients derived from an expansion of the scalar potential in powers of the radius (§18.7).

18.5 Cartesian Map Field Decomposition

Electric and magnetic fields can be parameterized as the sum over a number of functions with each function satisfying Maxwell's equations. These functions are also referred to as "maps", "modes", or "terms". *Bmad* has two types. The "`Cartesian`" decomposition is explained here. The other type is the `cylindrical` decomposition (§18.6).

The `Cartesian` decomposition implemented by *Bmad* involves a set of terms, each term a solution the Laplace equation solved using separation of variables in Cartesian coordinates. This decomposition is for DC electric or magnetic fields. No AC Cartesian Map decomposition is implemented by *Bmad*. In a lattice file, a `Cartesian` map is specified using the `cartesian_map` attribute as explained in Sec. §5.16.2.

The `Cartesian` decomposition is modeled using an extension of the method of Sagan, Crittenden, and Rubin[Sagan03]. In this decomposition, the magnetic(or electric field is written as a sum of terms B_i (For concreteness the symbol B_i is used but the equations below pertain equally well to both electric and magnetic fields) with:

$$\mathbf{B}(x,y,z) = \sum_i \mathbf{B}_i(x,y,z; A, k_x, k_y, k_z, x_0, y_0, \phi_z, family) \tag{18.30}$$

Each term B_i is specified using seven numbers $(A, k_x, k_y, k_z, x_0, y_0, \phi_z)$ and a switch called `family` which can be one of:

```
x,  qu
y,  sq
```

Roughly, taking the offsets x_0 and y_0 to be zero (see the equations below), the `x family` gives a field on-axis where the y component of the field is zero. that is, the `x` family is useful for simulating, say, magnetic vertical bend dipoles. The `y family` has a field that on-axis has no x component. The `qu family` has a magnetic quadrupole like (which for electric fields is skew quadrupole like) field on-axis and the `sq family` has a magnetic skew quadrupole like field on-axis. Additionally, assuming that the x_0 and y_0 offsets are zero, the `sq` family, unlike the other three families, has a nonzero on-axis z field component.

Each family has three possible forms These are designated as "`hyper-y`", "`hyper-xy`", and "`hyper-x`".

For the `x family` the `hyper-y` form is:

$$\begin{aligned}
B_x &= \quad A \frac{k_x}{k_y} \cos(k_x(x+x_0)) \cosh(k_y(y+y_0)) \cos(k_z z + \phi_z) \\
B_y &= \quad A \quad \sin(k_x(x+x_0)) \sinh(k_y(y+y_0)) \cos(k_z z + \phi_z) \\
B_s &= -A \frac{k_z}{k_y} \sin(k_x(x+x_0)) \cosh(k_y(y+y_0)) \sin(k_z z + \phi_z) \\
&\qquad \text{with } k_y^2 = k_x^2 + k_z^2
\end{aligned} \tag{18.31}$$

The `x family hyper-xy` form is:

$$\begin{aligned}
B_x &= \quad A \frac{k_x}{k_z} \cosh(k_x(x+x_0)) \cosh(k_y(y+y_0)) \cos(k_z z + \phi_z) \\
B_y &= \quad A \frac{k_y}{k_z} \sinh(k_x(x+x_0)) \sinh(k_y(y+y_0)) \cos(k_z z + \phi_z) \\
B_s &= -A \quad \sinh(k_x(x+x_0)) \cosh(k_y(y+y_0)) \sin(k_z z + \phi_z) \\
&\qquad \text{with } k_z^2 = k_x^2 + k_y^2
\end{aligned} \tag{18.32}$$

And the `x family` `hyper-x` form is:

$$\begin{aligned} B_x &= \quad A \quad \cosh(k_x(x+x_0))\,\cos(k_y(y+y_0))\,\cos(k_z z+\phi_z) \\ B_y &= -A\,\frac{k_y}{k_x}\,\sinh(k_x(x+x_0))\,\sin(k_y(y+y_0))\,\cos(k_z z+\phi_z) \\ B_s &= -A\,\frac{k_z}{k_x}\,\sinh(k_x(x+x_0))\,\cos(k_y(y+y_0))\,\sin(k_z z+\phi_z) \\ &\qquad \text{with } k_x^2 = k_y^2 + k_z^2 \end{aligned} \tag{18.33}$$

The relationship between k_x, k_y, and k_z ensures that Maxwell's equations are satisfied. Notice that which form `hyper-y`, `hyper-xy`, and `hyper-x` a particular $\mathbf{B}_i$ belongs to can be computed by *Bmad* by looking at the values of k_x, k_y, and k_z.

Using a compact notation where Ch $\equiv$ cosh, subscript x is $k_x(x+x_0)$, subscript z is $k_z z+\phi_z$, etc., the `y family` of forms is:

Form	hyper-y	hyper-xy	hyper-x
B_x	$-A\,\frac{k_x}{k_y}\,\mathrm{S}_x\,\mathrm{Sh}_y\,\mathrm{C}_z$	$A\,\frac{k_x}{k_z}\,\mathrm{Sh}_x\,\mathrm{Sh}_y\,\mathrm{C}_z$	$A \quad \mathrm{Sh}_x\,\mathrm{S}_y\,\mathrm{C}_z$
B_y	$A \quad \mathrm{C}_x\,\mathrm{Ch}_y\,\mathrm{C}_z$	$A\,\frac{k_y}{k_z}\,\mathrm{Ch}_x\,\mathrm{Ch}_y\,\mathrm{C}_z$	$A\,\frac{k_y}{k_x}\,\mathrm{Ch}_x\,\mathrm{C}_y\,\mathrm{C}_z$
B_z	$-A\,\frac{k_z}{k_y}\,\mathrm{C}_x\,\mathrm{Sh}_y\,\mathrm{S}_z$	$-A \quad \mathrm{Ch}_x\,\mathrm{Sh}_y\,\mathrm{S}_z$	$-A\,\frac{k_z}{k_x}\,\mathrm{Ch}_x\,\mathrm{S}_y\,\mathrm{S}_z$
with	$k_y^2 = k_x^2 + k_z^2$	$k_z^2 = k_x^2 + k_y^2$	$k_x^2 = k_y^2 + k_z^2$

(18.34)

the `qu family` of forms is:

Form	hyper-y	hyper-xy	hyper-x
B_x	$A\,\frac{k_x}{k_y}\,\mathrm{C}_x\,\mathrm{Sh}_y\,\mathrm{C}_z$	$A\,\frac{k_x}{k_z}\,\mathrm{Ch}_x\,\mathrm{Sh}_y\,\mathrm{C}_z$	$A \quad \mathrm{Ch}_x\,\mathrm{S}_y\,\mathrm{C}_z$
B_y	$A \quad \mathrm{S}_x\,\mathrm{Ch}_y\,\mathrm{C}_z$	$A\,\frac{k_y}{k_z}\,\mathrm{Sh}_x\,\mathrm{Ch}_y\,\mathrm{C}_z$	$A\,\frac{k_y}{k_x}\,\mathrm{Sh}_x\,\mathrm{C}_y\,\mathrm{C}_z$
B_z	$-A\,\frac{k_z}{k_y}\,\mathrm{S}_x\,\mathrm{Sh}_y\,\mathrm{S}_z$	$-A \quad \mathrm{Sh}_x\,\mathrm{Sh}_y\,\mathrm{S}_z$	$-A\,\frac{k_z}{k_x}\,\mathrm{Sh}_x\,\mathrm{S}_y\,\mathrm{S}_z$
with	$k_y^2 = k_x^2 + k_z^2$	$k_z^2 = k_x^2 + k_y^2$	$k_x^2 = k_y^2 + k_z^2$

(18.35)

the `sq family` of forms is:

Form	hyper-y	hyper-xy	hyper-x
B_x	$-A\,\frac{k_x}{k_y}\,\mathrm{S}_x\,\mathrm{Ch}_y\,\mathrm{C}_z$	$A\,\frac{k_x}{k_z}\,\mathrm{Sh}_x\,\mathrm{Ch}_y\,\mathrm{C}_z$	$-A \quad \mathrm{Sh}_x\,\mathrm{C}_y\,\mathrm{C}_z$
B_y	$A \quad \mathrm{C}_x\,\mathrm{Sh}_y\,\mathrm{C}_z$	$A\,\frac{k_y}{k_z}\,\mathrm{Ch}_x\,\mathrm{Sh}_y\,\mathrm{C}_z$	$A\,\frac{k_y}{k_x}\,\mathrm{Ch}_x\,\mathrm{S}_y\,\mathrm{C}_z$
B_z	$-A\,\frac{k_z}{k_y}\,\mathrm{C}_x\,\mathrm{Ch}_y\,\mathrm{S}_z$	$-A \quad \mathrm{Ch}_x\,\mathrm{Ch}_y\,\mathrm{S}_z$	$A\,\frac{k_z}{k_x}\,\mathrm{Ch}_x\,\mathrm{C}_y\,\mathrm{S}_z$
with	$k_y^2 = k_x^2 + k_z^2$	$k_z^2 = k_x^2 + k_y^2$	$k_x^2 = k_y^2 + k_z^2$

(18.36)

The singular case where $k_x = k_y = k_z = 0$ is not allowed. If a uniform field is needed, a term with very small k_x, k_y, and k_z can be used. Notice that since k_y must be non-zero for the `hyper-y` forms

(remember, $k_y^2 = k_x^2 + k_z^2$ for these forms and not all k's can be zero), and k_z must be non-zero for the `hyper-xy` forms, and k_x must be nonzero for the `hyper-x` forms. The magnetic field is always well defined even if one of the k's is zero.

Note: The vector potential for these fields is given in §26.25.

18.6 Cylindrical Map Decomposition

Electric and magnetic fields can be parameterized as the sum over a number of functions with each function satisfying Maxwell's equations. These functions are also referred to as "maps", "modes", or "terms". *Bmad* has two types. The "`cylindrical`" decomposition is explained here. The other type is the `Cartesian` decomposition (§18.6).

In a lattice file, a `cylindrical` map is specified using the `cylindrical_map` attribute as explained in Sec. §5.16.3.

The `cylindrical` decomposition takes one of two forms depending upon whether the fields are time varying or not. The DC decomposition is explained in Sec. §18.6.1 while the RF decomposition is explained in Sec. §18.6.2.

18.6.1 DC Cylindrical Map Decomposition

The DC `cylindrical` parametrization used by *Bmad* essentially follows Venturini et al.[Venturini98]. See Section §5.16 for details on the syntax used to cylindrical maps in *Bmad*. The electric and magnetic fields are both described by a scalar potential[2]

$$\mathbf{B} = -\nabla \psi_B, \qquad \mathbf{E} = -\nabla \psi_E \tag{18.37}$$

The scalar potentials both satisfy the Laplace equation $\nabla^2 \psi = 0$. The scalar potentials are decomposed as a sum of modes indexed by an integer m

$$\psi_B = \mathrm{Re}\left[\sum_{m=0}^{\infty} \psi_{Bm}\right] \tag{18.38}$$

[Here and below, only equations for the magnetic field will be shown. The equations for the electric fields are similar.] The ψ_{Bm} are decomposed in z using a discrete Fourier sum.[3] Expressed in cylindrical coordinates the decomposition of ψ_{Bm} is

$$\psi_{Bm} = \sum_{n=-N/2}^{N/2-1} \psi_{Bmn} = \sum_{n=-N/2}^{N/2-1} \frac{-1}{k_n}\, e^{i\,k_n\,z}\, \cos(m\,\theta - \theta_{0m})\, b_m(n)\, I_m(k_n\,\rho) \tag{18.39}$$

where I_m is a modified Bessel function of the first kind, and the $b_m(n)$ are complex coefficients. [For electric fields, $e_m(n)$ is substituted for $b_m(n)$] In Eq. (18.39) k_n is given by

$$k_n = \frac{2\pi\, n}{N\, dz} \tag{18.40}$$

[2] Notice the negative sign here and in Eq. (18.39) compared to Venturini et al.[Venturini98]. This is to keep the definition of the electric scalar potential ψ_E consistent with the normal definition.

[3] Venturini uses a continuous Fourier transformation but *Bmad* uses a discrete transformation so that only a finite number of coefficients are needed.

where N is the number of "sample points", and dz is the longitudinal "distance between points". That is, the above equations will only be accurate over a longitudinal length $(N-1)\,dz$. Note: Typically the sum in Eq. (18.39) and other equations below runs from 0 to $N-1$. Using a sum from $-N/2$ to $N/2-1$ gives exactly the same field at the sample points ($z = 0, dz, 2\,ds, \ldots$) and has the virtue that the field is smoother in between.

The field associated with ψ_{Bm} is for $m = 0$:

$$\begin{aligned} B_\rho &= \text{Re}\left[\sum_{n=-N/2}^{N/2-1} e^{i\,k_n\,z}\,b_0(n)\,I_1(k_n\,\rho)\right] \\ B_\theta &= 0 \\ B_z &= \text{Re}\left[\sum_{n=-N/2}^{N/2-1} i\,e^{i\,k_n\,z}\,b_0(n)\,I_0(k_n\,\rho)\right] \end{aligned} \tag{18.41}$$

And for $m \neq 0$:

$$\begin{aligned} B_\rho &= \text{Re}\left[\sum_{n=-N/2}^{N/2-1} \frac{1}{2}\,e^{i\,k_n\,z}\,\cos(m\,\theta-\theta_{0m})\,b_m(n)\left[I_{m-1}(k_n\,\rho)+I_{m+1}(k_n\,\rho)\right]\right] \\ B_\theta &= \text{Re}\left[\sum_{n=-N/2}^{N/2-1} \frac{-1}{2}\,e^{i\,k_n\,z}\,\sin(m\,\theta-\theta_{0m})\,b_m(n)\left[I_{m-1}(k_n\,\rho)-I_{m+1}(k_n\,\rho)\right]\right] \\ B_z &= \text{Re}\left[\sum_{n=-N/2}^{N/2-1} i\,e^{i\,k_n\,z}\,\cos(m\,\theta-\theta_{0m})\,b_m(n)\,I_m(k_n\,\rho)\right] \end{aligned} \tag{18.42}$$

While technically ψ_{Bm0} is not well defined due to the $1/k_n$ factor that is present, the field itself is well behaved. Mathematically, Eq. (18.39) can be corrected if, for $n = 0$, the term $I_m(k_n\,\rho)/k_n$ is replaced by

$$\frac{I_m(k_0\,\rho)}{k_0} \to \begin{cases} \rho & \text{if } m = 0 \\ \rho/2 & \text{if } m = 1 \\ 0 & \text{otherwise} \end{cases} \tag{18.43}$$

The magnetic vector potential for $m = 0$ is constructed such that only A_θ is non-zero

$$\begin{aligned} A_\rho &= 0 \\ A_\theta &= \text{Re}\left[\sum_{n=-N/2}^{N/2-1} \frac{i}{k_n}\,e^{i\,k_n\,z}\,b_0(n)\,I_1(k_n\,\rho)\right] \\ A_z &= 0 \end{aligned} \tag{18.44}$$

For $m \neq 0$, the vector potential is chosen so that A_θ is zero.

$$\begin{aligned} A_\rho &= \text{Re}\left[\sum_{n=-N/2}^{N/2-1} \frac{-i\,\rho}{2\,m}\, e^{i\,k_n\,z}\, \cos(m\,\theta - \theta_{0m})\, b_m(n)\, \Big[I_{m-1}(k_n\,\rho) - I_{m+1}(k_n\,\rho)\Big]\right] \\ A_\theta &= 0 \\ A_z &= \text{Re}\left[\sum_{n=-N/2}^{N/2-1} \frac{-i\,\rho}{m}\, e^{i\,k_n\,z}\, \cos(m\,\theta - \theta_{0m})\, b_m(n)\, I_m(k_n\,\rho)\right] \end{aligned} \tag{18.45}$$

Note: The description of the field using ‘‘`generalized gradients`’’[Newton99] is similar to the above equations. The difference is that, with the generalized gradient formalism, terms in θ and ρ are expanded in a Taylor series in x and y.

18.6.2 AC Cylindrical Map Decomposition

For RF fields, the `cylindrical` mode parametrization used by *Bmad* essentially follows Abell[Abell06]. The electric field is the real part of the complex field

$$\mathbf{E}(\mathbf{r}) = \sum_{j=1}^{M} \mathbf{E}_j(\mathbf{r})\, \exp[-2\pi i\,(f_j\,t + \phi_{0j})] \tag{18.46}$$

where M is the number of modes. Each mode satisfies the vector Helmholtz equation

$$\nabla^2 \mathbf{E}_j + k_{tj}^2\, \mathbf{E}_j = 0 \tag{18.47}$$

where $k_{tj} = 2\,\pi\, f_j/c$ with f_j being the mode frequency.

The individual modes vary azimuthally as $\cos(m\,\theta - \theta_0)$ where m is a non-negative integer. [in this and in subsequent equations, the mode index j has been dropped.] For the $m = 0$ modes, there is an accelerating mode whose electric field is in the form

$$\begin{aligned} E_\rho(\mathbf{r}) &= \sum_{n=-N/2}^{N/2-1} -e^{i\,k_n\,z}\, i\,k_n\, e_0(n)\, \widetilde{I}_1(\kappa_n, \rho) \\ E_\theta(\mathbf{r}) &= 0 \\ E_z(\mathbf{r}) &= \sum_{n=-N/2}^{N/2-1} e^{i\,k_n\,z}\, e_0(n)\, \widetilde{I}_0(\kappa_n, \rho) \end{aligned} \tag{18.48}$$

where $\widetilde{I}_m$ is

$$\widetilde{I}_m(\kappa_n, \rho) \equiv \frac{I_m(\kappa_n\,\rho)}{\kappa_n^m} \tag{18.49}$$

with I_m being a modified Bessel function first kind, and κ_n is given by

$$\kappa_n = \sqrt{k_n^2 - k_t^2} = \begin{cases} \sqrt{k_n^2 - k_t^2} & |k_n| > k_t \\ -i\,\sqrt{k_t^2 - k_n^2} & k_t > |k_n| \end{cases} \tag{18.50}$$

with

$$k_n = \frac{2\pi\, n}{N\, dz} \tag{18.51}$$

N is the number of points where E_{zc} is evaluated, and dz is the distance between points. The length of the field region is $(N-1)\,dz$. When κ_n is imaginary, $I_m(\kappa_n\,\rho)$ can be evaluated through the relation

$$I_m(-i\,x) = i^{-m}\,J_m(x) \tag{18.52}$$

where J_m is a Bessel function of the first kind. The e_0 coefficients can be obtained given knowledge of the field at some radius R via

$$e_0(n) = \frac{1}{\widetilde{I}_0(\kappa_n, R)}\,\frac{1}{N}\,\sum_{p=0}^{N-1} e^{-2\pi i n p/N}\,E_z(R, p\,dz) \tag{18.53}$$

The non-accelerating $m = 0$ mode has an electric field in the form

$$\begin{aligned} E_\rho(\mathbf{r}) &= E_z(\mathbf{r}) = 0 \\ E_\theta(\mathbf{r}) &= \sum_{n=-N/2}^{N/2-1} e^{ik_n z}\,b_0(n)\,\widetilde{I}_1(\kappa_n, \rho) \end{aligned} \tag{18.54}$$

where the b_0 coefficients can be obtained given knowledge of the field at some radius R via

$$b_0(n) = \frac{1}{\widetilde{I}_1(\kappa_n, R)}\,\frac{1}{N}\,\sum_{p=0}^{N-1} e^{-2\pi i\,n\,p/N}\,E_\theta(R, p\,dz) \tag{18.55}$$

For positive m, the electric field is in the form

$$\begin{aligned} E_\rho(\mathbf{r}) &= \sum_{n=-N/2}^{N/2-1} -i\,e^{i\,k_n\,z}\,\left[k_n\,e_m(n)\,\widetilde{I}_{m+1}(\kappa_n, \rho) + b_m(n)\,\frac{\widetilde{I}_m(\kappa_n, \rho)}{\rho}\right]\cos(m\,\theta - \theta_{0m}) \\ E_\theta(\mathbf{r}) &= \sum_{n=-N/2}^{N/2-1} -i\,e^{i\,k_n\,z}\,\Big[k_n\,e_m(n)\,\widetilde{I}_{m+1}(\kappa_n, \rho) + \\ &\qquad\qquad b_m(n)\,\left(\frac{\widetilde{I}_m(\kappa_n, \rho)}{\rho} - \frac{1}{m}\,\widetilde{I}_{m-1}(\kappa_n, \rho)\right)\Big]\sin(m\,\theta - \theta_{0m}) \\ E_z(\mathbf{r}) &= \sum_{n=-N/2}^{N/2-1} e^{i\,k_n\,z}\,e_m(n)\,\widetilde{I}_m(\kappa_n, \rho)\cos(m\,\theta - \theta_{0m}) \end{aligned} \tag{18.56}$$

The `e_m` and `b_m` coefficients can be obtained given knowledge of the field at some radius R via

$$\begin{aligned} e_m(n) &= \frac{1}{\widetilde{I}_m(\kappa_n, R)}\,\frac{1}{N}\,\sum_{p=0}^{N-1} e^{-2\,\pi\,i\,n\,p/N}\,E_{zc}(R, p\,dz) \\ b_m(n) &= \frac{R}{\widetilde{I}_m(\kappa_n, R)}\,\left[\frac{1}{N}\,\sum_{p=0}^{N-1} i\,e^{-2\,\pi\,i\,n\,p/N}\,E_{\rho c}(R, p\,dz) - k_n\,e_m(n)\,\widetilde{I}_{m+1}(\kappa_n, R)\right] \end{aligned} \tag{18.57}$$

where $E_{\rho c}$, $E_{\theta s}$, and E_{zc} are defined by

$$\begin{aligned} E_\rho(R, \theta, z) &= E_{\rho c}(R, z)\,\cos(m\,\theta - \theta_{0m}) \\ E_\theta(R, \theta, z) &= E_{\theta s}(R, z)\,\sin(m\,\theta - \theta_{0m}) \\ E_z(R, \theta, z) &= E_{zc}(R, z)\,\cos(m\,\theta - \theta_{0m}) \end{aligned} \tag{18.58}$$

The above mode decomposition was done in the gauge where the scalar potential ψ is zero. The electric and magnetic fields are thus related to the vector potential $\mathbf{A}$ via

$$\mathbf{E} = -\partial_t \mathbf{A}, \qquad \mathbf{B} = \nabla \times \mathbf{A} \tag{18.59}$$

Using Eq. (18.46), the vector potential can be obtained from the electric field via

$$\mathbf{A}_j = \frac{-i\,\mathbf{E}_j}{2\,\pi\,f_j} \tag{18.60}$$

Symplectic tracking through the RF field is discussed in Section §26.4. For the fundamental accelerating mode, the vector potential can be analytically integrated using the identity

$$\int dx\,\frac{x\,I_1(a\,\sqrt{x^2+y^2})}{\sqrt{x^2+y^2}} = \frac{1}{a}\,I_0(a\,\sqrt{x^2+y^2}) \tag{18.61}$$

18.7 Generalized Gradient Field Modeling

Bmad has a number of `field map` models that can be used to model electric or magnetic fields (§5.16). One model involves what are called `generalized gradients` (GG). This model is restricted to modeling DC magnetic or electric fields. In a lattice file, the generalized gradient field model is specified using the `gen_gradients` attribute as explained in Sec. §5.16.5.

Bmad uses the curved-curvilinear-coordinate GG formalism of Van der Schueren et al.[Schueren24], which handles a reference frame of constant curvature $g_{\rm ref} = 1/\rho$ (and reduces to the straight-coordinate case when $g_{\rm ref} = 0$). The field is characterized by three families of midplane derivative functions of the longitudinal coordinate s: the normal (non-skew) functions $b_n(s)$, the skew functions $a_n(s)$, and the solenoidal function $b_s(s)$, defined by the on-midplane field

$$\begin{aligned} B_x(x,0,s) &= \sum_{n\geq 1} a_n(s)\,\frac{x^{n-1}}{(n-1)!} \\ B_y(x,0,s) &= \sum_{n\geq 1} b_n(s)\,\frac{x^{n-1}}{(n-1)!} \\ B_s(0,0,s) &= b_s(s) \end{aligned} \tag{18.62}$$

where n is the multipole harmonic index ($n = 1$ dipole, $n = 2$ quadrupole, etc.). The full field $\mathbf{B} = (B_x, B_y, B_s)$ and the magnetic vector potential $\mathbf{A} = (A_x, A_y, A_s)$ off the midplane are reconstructed from these functions and their s-derivatives $a_n^{[m]}, b_n^{[m]}, b_s^{[m]}$ using a precomputed coefficient table: each output component is a polynomial in (x, y) whose coefficients are sums of the form $\sum \text{coef}\cdot g_{\rm ref}^k\,x^p\,y^q\,G^{[m]}$ (with $G \in \{a_n, b_n, b_s\}$), so the curvature enters through explicit powers of $g_{\rm ref}$. The vector potential is constructed so that $\mathbf{B} = \nabla \times \mathbf{A}$ holds exactly (in the curved metric of the reference frame). This construction, the coefficient table, and the field-fitting tools are provided by the `GeneralizedGradients.jl` package[GenGrad]; *Bmad*'s coefficient table (`bmad/modules/gg_coef_table_mod.f90`) is generated from it.

The functions are characterized by specifying each function and its s-derivatives at equally spaced planes z_i, up to some maximum derivative order chosen by the user. Interpolation between planes is done with a two-point Hermite ("non-smoothing spline") polynomial that matches the value and all stored derivatives at z_i and z_{i+1}. As with the older formalism, the "locality" property — the field at a point depends only on the two straddling planes — makes the coefficients easy to fit and the evaluation fast, at the cost

that, because the derivatives are truncated, the reconstructed field satisfies Maxwell's equations only up to the truncation order.

The obsolete straight-coordinate cylindrical-harmonic formalism (Venturini and Dragt[Venturini99]), used by the removed `gen_grad_map` lattice format, is documented for reference in Appendix §A.

18.8 RF fields

The following describes the how RF fields are calculated when the `field_calc` attribute of an RF element is set to `bmad_standard`.[4] Also see Section §19.9 for how fringe fields are calculated.

With `cavity_type` set to `traveling_wave`, the setting of `longitudinal_mode` is ignored and the field is given by

$$\begin{aligned} E_s(r,\phi,s,t) &= G\,\cos\bigl(\omega\,t - k\,s + 2\,\pi\,\phi\bigr) \\ E_r(r,\phi,s,t) &= -\frac{r}{2}\,G\,k\,\sin\bigl(\omega\,t - k\,s + 2\,\pi\,\phi\bigr) \\ B_\phi(r,\phi,s,t) &= -\frac{r}{2\,c}\,G\,k\,\sin\bigl(\omega\,t - k\,s + 2\,\pi\,\phi\bigr) \end{aligned} \tag{18.63}$$

where G is the accelerating gradient, $k = \omega/c$ is the wave number with ω being the RF frequency.

For standing wave cavities, with `cavity_type` set to `standing_wave`, the RF fields are modeled as N half-wave cells, each having a length of $\lambda/2$ where $\lambda = 2\,\pi/k$ is the wavelength. If the length of the RF element is not equal to the length of N cells, the "active region" is centered in the element and the regions to either side are treated as field free.

The field in the standing wave cell is modeled either with a $p = 0$ or $p = 1$ longitudinal mode (set by the `longitudinal_mode` element parameter). The $p = 1$ longitudinal mode models the fields as a pillbox with the transverse wall at infinity as detailed in Chapter 3, Section VI of reference [Lee99]

$$\begin{aligned} E_s(r,\phi,s,t) &= 2\,G\,\cos(k\,s)\,\cos(\omega\,t + 2\,\pi\,\phi) \\ E_r(r,\phi,s,t) &= r\,G\,k\,\sin(k\,s)\,\cos(\omega\,t + 2\,\pi\,\phi) \\ B_\phi(r,\phi,s,t) &= -\frac{r}{c}\,G\,k\,\cos(k\,s)\,\sin(\omega\,t + 2\,\pi\,\phi) \end{aligned} \tag{18.64}$$

The overall factor of 2 in the equation is present to ensure that an ultra-relativistic particle entering with $\phi = 0$ will experience an average gradient equal to G.

For the $p = 0$ longitudinal mode (which is the default), a "pseudo TM_{010}" mode is used that has the correct symmetry:

$$\begin{aligned} E_s(r,\phi,s,t) &= 2\,G\,\sin(k\,s)\,\sin(\omega\,t + 2\,\pi\,\phi) \\ E_r(r,\phi,s,t) &= -r\,G\,k\,\cos(k\,s)\,\sin(\omega\,t + 2\,\pi\,\phi) \\ B_\phi(r,\phi,s,t) &= \frac{r}{c}\,G\,k\,\sin(k\,s)\,\cos(\omega\,t + 2\,\pi\,\phi) \end{aligned} \tag{18.65}$$

[4] Notice that the equations here are only relavent with the `tracking_method` for an RF element set to a method like `runge_kutta` where tracking through the field of an element is done. For `bmad_standard` tracking, Equations for `lcavity` tracking are shown in §26.14 and `rfcavity` tracking in §26.19.

Chapter 19

Fringe Fields

It can be convient to divide the fringe field kick into two pieces. The first piece is called the `hard edge` fringe kick and is the kick in the limit that the longitudinal extent of the fringe is zero. The second piece is the `soft edge` fringe kick which is the fringe kick with the fringe having a finite longitudinal extent minus the hard edge fringe kick. That is

```
Total fringe kick = hard fringe kick + soft fringe kick
```

The advantage of separating the fringe kick in this way is that the hard fringe can be used without having to know anything about the longitudinal extent of the fringe field. In many cases, this is a good enough approximation.

Discussion of the fringe parameters of an element are detailed in Sec. §5.21.

19.1 Bend Second Order Fringe Map

The bend fringe kick is a combination of the equations developed by Hwang and Lee[Hwang15] modified to include quadrupole terms as given in Section 5.3.1 of Iselin[Iselin94]. The Lie map generator Ω_M given by Hwang and Lee in Eqs. (35) and (36) is used under the conditions that

$$K_{0h} = K_{1h} = K_{3h} = K_{4h} = K_{5h} = K_{6h} = 0 \tag{19.1}$$

Here the subscript "`h`" has been added so as to not confuse these parameters with the magnetic multipole coefficients K_1, K_2, etc. Note: Hwang and Lee do not present an equation for the change in the longitudinal phase space z coordinate in their paper.

The generator used by *Bmad* for the entrance fringe is:

$$\begin{aligned}\Omega_{M1} = &\frac{(x^2 - y^2)\, g_{\rm tot}\, \tan(e_1)}{2} + \frac{y^2\, g_{\rm tot}^2\, \sec^3(e_1)\, [1 + \sin^2(e_1)]\, f_{int}\, h_{gap}}{(1 + p_z)} \\ &+ \frac{x^3\, [4\, K_1\, \tan(e_1) - g_{\rm tot}^2\, \tan^3(e_1)]}{12\, (1 + p_z)} + \frac{x\, y^2\, [-4\, K_1\, \tan(e_1) + g_{\rm tot}^2\, \tan(e_1)\, \sec^2(e_1)]}{4\, (1 + p_z)} \\ &+ \frac{(x^2\, p_x - 2\, x\, y\, p_y)\, g_{\rm tot}\, \tan^2(e_1)}{2\, (1 + p_z)} - \frac{y^2\, p_x\, g_{\rm tot}\, \sec^2(e_1)}{2\, (1 + p_z)}\end{aligned} \tag{19.2}$$

where $g_{\rm tot}$ is the total bending strength (design + error) and K_1 is the quadrupole moment of the bend.

The generator for the exit fringe is

$$\begin{aligned}\Omega_{M2} = &\frac{(x^2-y^2)\,g_{\rm tot}\,\tan(e_2)}{2} + \frac{y^2\,g_{\rm tot}^2\,\sec^3(e_2)\,[1+\sin^2(e_2)]\,f_{int}\,h_{gap}}{(1+p_z)} \\ &+ \frac{x^3\,[4\,K_1\,\tan(e_2) - g_{\rm tot}^2\,\tan^3(e_2)]}{12\,(1+p_z)} + \frac{x\,y^2\,[-4\,K_1\,\tan(e_2) + g_{\rm tot}^2\,\tan(e_2)\,\sec^2(e_2)]}{4\,(1+p_z)} \\ &- \frac{(x^2\,p_x - 2\,x\,y\,p_y)\,g_{\rm tot}\tan^2(e_2)}{2\,(1+p_z)} + \frac{y^2\,p_x\,g_{\rm tot}\,\sec^2(e_2)}{2\,(1+p_z)}\end{aligned} \tag{19.3}$$

The map $\mathcal{M}$ is obtained from the equation $\mathcal{M} = \exp[:\Omega_M:]$. To second order in the transverse coordinates the map can be obtained by expanding the exponential to second order

$$\mathcal{M} \simeq 1 + :\Omega_M: + \frac{1}{2} :\Omega_M::\Omega_M: \tag{19.4}$$

The transport for the entrance fringe is then

$$\begin{aligned}\Delta x &= \frac{g_{\rm tot}}{2\,(1+p_z)}\,\left[-x^2\,\tan^2(e_1) + y^2\sec^2(e_1)\right] \\ \Delta p_x &= x\,g_{\rm tot}\,\tan(e_1) + \frac{y^2\,g_{\rm tot}^2\,[\tan(e_1) + 2\,\tan^3(e_1)]}{2\,(1+p_z)} \\ &\quad + \frac{(x^2-y^2)\,K_1\tan(e_1)}{1+p_z} + \frac{(x\,p_x - y\,p_y)\,g_{\rm tot}\,\tan^2(e_1)}{1+p_z} \\ \Delta y &= \frac{x\,y\,g_{\rm tot}\,\tan^2(e_1)}{1+p_z} \\ \Delta p_y &= y\left[-g_{\rm tot}\,\tan(e_1) + \frac{2\,g_{\rm tot}^2\,[1+\sin^2(e_1)]\,\sec^3(e_1)}{1+p_z}\,f_{int}\,h_{gap}\right] \\ &\quad - \frac{(x\,p_y\,g_{\rm tot}\,\tan^2(e_1)}{1+p_z} - \frac{y\,p_x\,g_{\rm tot}\,[1+\tan^2(e_1)]}{1+p_z} - \frac{2\,x\,y\,K_1\tan(e_1)}{(1+p_z)} \\ \Delta z &= \frac{\Omega_{M1} - (x^2-y^2)\,g_{\rm tot}\,\tan(e_1)/2}{1+p_z}\end{aligned} \tag{19.5}$$

The transport for the exit fringe is

$$\begin{aligned}\Delta x &= \frac{g_{\rm tot}}{2\,(1+p_z)}\,\left[x^2\,\tan^2(e_2) - y^2\sec^2(e_2)\right] \\ \Delta p_x &= x\,g_{\rm tot}\,\tan(e_2) - \frac{(x^2+y^2)\,g_{\rm tot}^2\,\tan^3(e_2)}{2\,(1+p_z)} \\ &\quad + \frac{(x^2-y^2)\,K_1\tan(e_2)}{1+p_z} + \frac{(-x\,p_x + y\,p_y)\,g_{\rm tot}\,\tan^2(e_2)}{1+p_z} \\ \Delta y &= -\frac{x\,y\,g_{\rm tot}\,\tan^2(e_2)}{1+p_z} \\ \Delta p_y &= y\left[-g_{\rm tot}\,\tan(e_2) + \frac{2\,g_{\rm tot}^2\,[1+\sin^2(e_2)]\,\sec^3(e_2)}{1+p_z}\,f_{int}\,h_{gap}\right] + \frac{x\,y\,g_{\rm tot}^2\,\sec^2(e_2)\,\tan(e_2)}{1+p_z} \\ &\quad + \frac{(x\,p_y\,g_{\rm tot}\,\tan^2(e_2)}{1+p_z} + \frac{y\,p_x\,g_{\rm tot}\,[1+\tan^2(e_2)]}{1+p_z} - \frac{2\,x\,y\,K_1\tan(e_2)}{(1+p_z)} \\ \Delta z &= \frac{\Omega_{M2} - (x^2-y^2)\,g_{\rm tot}\,\tan(e_2)/2}{1+p_z}\end{aligned} \tag{19.6}$$

19.2 SAD Dipole Soft Edge Fringe Map

The SAD dipole soft edge fringe model is adapted from the SAD program[SAD]. This model is only used with `sbend`, `rbend`, and `sad_mult` elements. For `sbend` and `rbend` elements, the fringe map is defined in terms of the `fint` and `hgap` for the entrance end and `fintx` and `hgapx` at the exit end (§4.5). The field integral F_{H1} for the entrance end given is given by (see Eq. (4.9))

$$F_{H1} \equiv F_{int}\, H_{gap} = \int_{pole} ds\, \frac{B_y(s)\,(B_{y0} - B_y(s))}{2\,B_{y0}^2} \tag{19.7}$$

With a similar equation for F_{H2} for the exit end. The conversion between `f1`, `fb1`, and `fb2` parameters for a SAD `bend` element and a *Bmad* bend element is

```
fq1 (Bmad) = f1 + fb1 (SAD)
fq2 (Bmad) = f1 + fb2 (SAD)
```

Note! To be confusing, the SAD `mult` element also has an `f1` parameter but for the `mult` element the `f1` parameter is associated with the soft quadrupole fringe (§19.6).

For a `sad_mult` element the corresponding parameters are `fb1` and `fb2`. the conversion between the bend and `sad_mult` parameters is

$$\text{fb1} = 12\,F_{H1}, \qquad \text{fb2} = 12\,F_{H2} \tag{19.8}$$

The map itself is

$$\begin{aligned} x_2 &= x_1 + c_1\,p_z \\ p_{y2} &= p_{y1} + c_2\,y_1 - c_3\,y_1^3 \\ z_2 &= z_1 + \frac{1}{1+p_{z1}}\left(c_1\,p_{x1} + \frac{1}{2}\,c_2\,y_1^2 - \frac{1}{4}\,c_3\,y_1^4\right) \end{aligned} \tag{19.9}$$

For the entrance face the map parameters are:

$$\begin{aligned} c_1 = \frac{g_{\rm tot}\,\text{fb1}^2}{24\,(1+p_z)} = \frac{6\,g_{\rm tot}\,F_{H1}^2}{1+p_z}, \qquad c_2 = \frac{g_{\rm tot}^2\,\text{fb1}}{6\,(1+p_z)} = \frac{2\,g_{\rm tot}^2\,F_{H1}}{1+p_z}, \\ c_3 = \frac{2\,g_{\rm tot}^2}{3\,\text{fb1}\,(1+p_z)} = \frac{g_{\rm tot}^2}{18\,F_{H1}\,(1+p_z)} \end{aligned} \tag{19.10}$$

And for the exit face, the appropriate equations can be derived using the substitution

$$\begin{aligned} F_{H1} &\to F_{H2} \\ g_{\rm tot} &\to -g_{\rm tot} \end{aligned} \tag{19.11}$$

In the above equations, for a bend, `g_tot` is the total bending strength

$$g_{\rm tot} = \text{g} + \text{dg} \tag{19.12}$$

`g` being the reference bend strength and `dg` being bend the difference between the actual and reference bend strengths (§4.5). For a `sad_mult` element `g_tot` is calculated from the equation

$$g_{\rm tot} = \sqrt{a_0^2 + b_0^2} \tag{19.13}$$

The SAD dipole soft edge map is "incomplete" and for a realistic fringe map the SAD dipole soft edge fring map must be combined with a "hard edge" map (§5.21).

It might seem strange that c_3 diverges to infinity as F_H goes to zero since naively one would expect the soft edge kick to vanish in the hard edge limit where the fringe has no longitudinal extent. However, in the hard edge limit, the field does not obey Maxwell's equations. The limiting map, as F_H goes to zero, has fields that diverge to infinity and this explains why the full (hard + soft) limiting map is not the same as the hard edge map at the limit of zero longitudinal extent.

19.3 Sad_Mult Dipole Hard Edge Fringe Map

For `sad_mult` elements, the hard dipole edge kick is adapted from SAD. The dipole normalized field $g = \sqrt{\mathrm{a0}^2 + \mathrm{b0}^2}$ is calculated from the `a0` and `b0` multipoles. Before the fringe kick is applied, the particle position is rotated in the (x, y) plane so that the dipole kick is in the horizontal direction. The dipole edge kick is then given by

$$\begin{aligned} \Delta x &= g\, y^2\ (1 - f_{yg})\ \frac{1 + p_z^2}{2\, p_{zy}^3} \\ \Delta p_y &= -g\, p_x\, y\, \frac{1 - 2 f_{yg}}{p_{zy}} \\ \Delta z &= -g\, y^2\, p_x\, (1 - f_{yg}) \frac{1 + p_z}{2\, p_{zy}^3} \end{aligned} \tag{19.14}$$

where

$$f_{yg} = \frac{y^2\, g^2}{12}, \quad \text{and} \qquad p_{zy} = \sqrt{(1 + p_z)^2 - p_x^2} \tag{19.15}$$

This is used in place of the dipole hard edge fringe kick given in §19.1.

19.4 Linear Dipole Hard Edge Fringe Map

The linear dipole hard edge fringe model is adapted from MAD[Grote96] and only includes linear terms. The fringe transport is

$$\begin{aligned} \Delta p_x &= g_{\mathrm{tot}}\ \tan(e) \cdot x \\ \Delta p_y &= -g_{\mathrm{tot}}\ \tan\left(e - \frac{2\, f_{int}\, h_{gap}\, g_{\mathrm{tot}}\, (1 + \sin^2(e))}{\cos(e)} \right) \cdot y \end{aligned} \tag{19.16}$$

where $g_{\mathrm{tot}} = g + dg$ is the actual field and e is `e1` if the particle is entering the dipole and `e2` if the particle is exiting the dipole.

19.5 Exact Dipole Hard Edge Fringe Map

The exact dipole hard edge fringe is the exact transport in the wedge region of a dipole when there is a finite `e1` or `e2` as shown in Fig. 4.1b. This model assumes that there are no higher order multipole fields. The transport is done in two stages. For a particle entering the dipole the propagation is

1. Drift (propagate in a straight line) the particle from the sector edge to the actual bend edge. The propagation may be forward or backwards depending upon on the geometry.

2. Propagate the particle as if it were in the dipole field from the actual bend edge to the sector edge.

The body of the dipole is treated as a sector bend. At the exit end, the propagation through the wedge is the reverse of the above.

19.6 Quadrupole Soft Edge Fringe Map

The quadrupole soft edge fringe model is adapted from SAD[SAD]. This fringe is only used with `sad_mult` and `quadrupole` elements. The fringe map is:

$$\begin{aligned} x_2 &= x_1\, e^{g_1} + g_2\, p_{x1} \\ p_{x2} &= p_{x1} e^{-g_1} \\ y_2 &= y_1\, e^{-g_1} - g_2\, p_{y1} \\ p_{y2} &= p_{y1} e^{g_1} \\ z_2 &= z_1 - \left[g_1\, x_1\, p_{x1} + g_2 \left(1 + \frac{g_1}{2}\right) e^{-g_1}\, p_{x1}^2 \right] + \left[g_1\, y_1\, p_{y1} + g_2 \left(1 - \frac{g_1}{2}\right) e^{g_1}\, p_{y1}^2 \right] \end{aligned} \tag{19.17}$$

where

$$g_1 = \frac{K_1\, \mathrm{fq1}}{1 + p_z}, \qquad g_2 = \frac{K_1\, \mathrm{fq2}}{1 + p_z} \tag{19.18}$$

K_1 is the quadrupole strength, and `fq1` and `fq2` are the fringe quadrupole parameters. These parameters are related to the field integral I_n via

$$\mathrm{fq1} = I_1 - \frac{1}{2}\, I_0^2, \qquad \mathrm{fq2} = I_2 - \frac{1}{3}\, I_0^3 \tag{19.19}$$

where I_n is defined by

$$I_n = \frac{1}{K_1} \int_{-\infty}^{\infty} \left(K_1(s) - H(s - s_0)\, K_1\right) (s - s_0)^n\, ds \tag{19.20}$$

and $H(s)$ is the step function

$$H(s) = \begin{cases} 1 & s > 0 \\ 0 & s < 0 \end{cases} \tag{19.21}$$

and it is assumed that the quadrupole edge is at s_0 and the interior is in the region $s > s_0$.

See Sec. §5.21 for the relation between `fq1` / `fq2` and the corresponding `f1` and `f2` parameters of SAD.

19.7 Magnetic Multipole Hard Edge Fringe Map

The magnetic multipole hard edge fringe field is modeled using the method shown in Forest[Forest98]. For the m^{th} order multipole the Lee transform is (Forest Eq. (13.29)):

$$\begin{aligned} f_{\pm} &= \mp \Re \left[\frac{(b_m + i\, a_m)\, (x + i\, y)^{m+1}}{4\, (m+2)\, (1+\delta)} \left\{ x\, p_x + y\, p_y + i \frac{m+3}{m+1} (x\, p_x - y\, p_y) \right\} \right] \\ &\equiv \frac{p_x\, f^x + p_y\, f^y}{1 + \delta} \end{aligned} \tag{19.22}$$

The multipole strengths a_m and b_m are given by (18.8) and the second equation defines f^x and f^y. On the right had side of the first equation, the minus sign is appropriate for particles entering the magnet and the plus sign is for particle leaving the magnet. Notice that here the multipole order m is equivalent to $n-1$ in Forest's notation.

With this, the implicit multipole map is (Forest Eq. (13.31))

$$\begin{aligned} x^f &= x - \frac{f^x}{1+\delta} \\ p_x &= p_x^f - \frac{p_x^f\,\partial_x f^x + p_y^f\,\partial_x f^y}{1+\delta} \\ y^f &= y - \frac{f_y}{1+\delta} \\ p_y &= p_y^f - \frac{p_x^f\,\partial_y f^x + p_y^f\,\partial_y f^y}{1+\delta} \\ \delta^f &= \delta \\ z^f &= \frac{p_x^f\,f^x + p_y^f\,f^y}{(1+\delta)^2} \end{aligned} \tag{19.23}$$

19.8 Electrostatic Multipole Hard Edge Fringe Map

The electric multipole hard edge fringe field, to lowest order, consists of just a longitudinal field. The integrated longitudinal field at constant (x, y) for the n^{th} order multipole is simply obtained by requiring that the curl of the field is zero. This gives:

$$\int E_s(x,y)\,ds = \phi_n(x,y) \tag{19.24}$$

where ϕ_n is given in Eq. (18.17). [For a magnetic multipole there is an analogous equation.]

The effect on the spin when tracking through the fringe field of a multipole field tends to be weak. As such, this hard edge model is sufficient. and the spin is tracked using the T-BMT equation (Eq. (24.1)).

19.9 RF Fringe Fields

Assuming cylindrical symmetry, the radial and azimuthal fields near the axis can be related to the longitudinal electric field via Maxwell's equations[Hartman93]

$$E_r = -\frac{r}{2}\frac{\partial E_z}{\partial z}, \qquad B_\phi = \frac{r}{2\,c}\frac{\partial E_z}{\partial t} \tag{19.25}$$

Assuming the particle velocity is c, these equations can be combined with the force equation

$$F_r = q\,(E_r - c\,B_\phi) \tag{19.26}$$

to give[Rosen94]

$$F_r = -\frac{q\,r}{2}\frac{dE_z}{dz} \tag{19.27}$$

where the total derivative is used here[1]. From Eq. (19.27), the fringe field kick in the horizontal plane at the entrance end, valid for both traveling wave and standing wave cavities, is (cf. Rosenzweig[Rosen94] Eq (10))

$$\Delta p_x = -\frac{q\,\widehat{E}_z}{2\,c\,P_0}\,x \tag{19.28}$$

with a similar equation in the vertical. Here $\widehat{E}_z$ is the longitudinal electric field just inside the fringe, and P_0 is the reference momentum. At the exit end, the kick is the negative of Eq. (19.28). This fringe kick is built into `xxpc`.

The integrated fringe fields, needed to calculate the spin precession, are at the entrance end

$$\begin{aligned} \int E_r\,ds &= -\frac{r}{4}\,\widehat{E}_z \\ \int B_\phi\,ds &= \frac{r}{4\,c}\,\widehat{E}_z \end{aligned} \tag{19.29}$$

The integrated fields at the exit end are obtained by negating the RHS of these equations.

[1] Hartman[Hartman93] Eq (16) is not valid for a forward propagating wave component of the field. Thus Hartman Eq (17) is only valid for a backward propagating wave component. Eq. (19.27), on the other hand, is valid for all wave components.

Chapter 20

Wakefields

Wakefield effects are divided into short–range (§20.1 within a bunch) and long–range (§20.2 between bunches). Short–range wakes are described in Sec. §20.1 and long–range wakes are described in Sec. §20.2. The syntax for describing wakes in a lattice file is given in Sec. §5.20.

Note: The short range wakes uses z (distance) as the dependent variable while the long range wakes use time as the dependent variable.

20.1 Short–Range Wakes

The syntax for assigning short–range wakes to a lattice element is described in Sec. §5.20.1. Only the monopole and dipole wakefields are modeled.

Short–range wakes are divided into three classes: Those that are dependent linearly upon transverse offset of the leading particle (but independent of the position of the trailing particle), those that are dependent linearly upon the transverse offset of the trailing particle (but independent of the position of the leading particle, and those wakes that are independent of the offset.

The longitudinal monopole energy kick dE for the i^{th} (trailing) macroparticle due to the wake from the j^{th} (leading) macroparticle, assuming the kick is independent of the transverse positions, is computed from the equation

$$\Delta p_z(i) = \frac{-e\,L}{v\,P_0}\left(\frac{1}{2}W_{\parallel}^{SR}(0)\,|q_i| + \sum_{j\neq i} W_{\parallel}^{SR}(dz_{ij})\,|q_j|\right) \tag{20.1}$$

where v is the particle velocity, e is the charge on an electron, q is the macroparticle charge, L is the cavity length, dz_{ij} is the longitudinal distance of the i^{th} particle with respect to the j^{th} macroparticle (positive z means i^{th} particle is in front of the j^{th} particle), $W_{\parallel}^{SR}$ is the short–range longitudinal wakefield function.

If the beam chamber has azimuthal symmetry the energy kick will be independent of the transverse positions. If this is not true, there can be some dependence. There are four cases that *Bmad* simulates: Linear in the x or y-position of the leading particle, or linear in the x or y-position of the trailing particle. For example, if the kick is linear in the x-position of the leading particle the kick is

$$\Delta p_z(i) = \frac{-e\,L\,x_i}{v\,P_0}\left(\frac{1}{2}W_{\parallel}^{SR}(0)\,|q_i| + \sum_{j\neq i} W_{\parallel}^{SR}(dz_{ij})\,|q_j|\right) \tag{20.2}$$

And if the kick is linear in the y-position of the trailing particle the kick is

$$\Delta p_z(i) = \frac{-e\,L}{v\,P_0}\left(\frac{1}{2}W_{\parallel}^{SR}(0)\,|q_i|\,y_i + \sum_{j\neq i} W_{\parallel}^{SR}(dz_{ij})\,|q_j|\,y_j\right) \tag{20.3}$$

The kick $\Delta p_x(i)$ due to the transverse wake for the i^{th} particle is modeled with the equation

$$\Delta p_x(i) = \frac{-e\,L\,\sum_j |q_j|\,x\,W_{\perp}^{SR}(dz_{ij})}{v\,P_0} \tag{20.4}$$

Where $W_{\perp}^{SR}$ is the transverse short–range wake function and x is the horizontal displacement of the leading or trailing particle as appropriate. There is a similar equation for $\Delta p_y(i)$. If the beam chamber has azimuthal symmetry, the only wakes present are those that are dependent upon the offset of the leading particle. If the transverse wake is modeled as being independent of position the above equation is modified:

$$\Delta p_x(i) = \frac{-e\,L\,\sum_j |q_j|\,W_{\perp}^{SR}(dz_{ij})}{v\,P_0} \tag{20.5}$$

With either the longitudinal wake $W_{\parallel}^{SR}$ or the transverse $W_{\perp}^{SR}$ wake, the wake can be approximated as a sum of what are called "pseudo" modes $W_i(z)$, $i = 1 \ldots M$:

$$W(z) = A_a \sum_{i=1}^{M} W_i(z) = A_a \sum_{i=1}^{M} A_i\, e^{d_i z}\, \sin(k_i\, z + \phi_i) \tag{20.6}$$

This is similar to approximating any function as a sum of Fourier terms. The parameters (A_i, d_i, k_i, ϕ_i) are chosen by the person constructing the lattice to fit the calculated wake potential. Since z is negative for trailing particles, d_i should be positive to get the wake to decay exponentially with distance. The dimensionless overall amplitude scale A_a is introduced as a convenient way to scale the overall wake. The reason why the pseudo mode approach is used in *Bmad* is due to the fact that, with pseudo modes, the calculation time scales as the number of particles N while a calculation based upon a table of wake vs z would scale as N^2. [The disadvantage is that initially must perform a fit to the wake potential to generate the mode parameter values.]

20.2 Long–Range Wakes

The lattice syntax for defining long-range wakes is discussed in Sec. §5.20.4.

Following Chao[Chao93] Eq. 2.88, the long–range wakefields are characterized by a set of cavity modes. The wake function W_i for the i^{th} mode is

$$W_i(t) = -c\,A_a \left(\frac{R}{Q}\right)_i \exp(-d_i\,t)\,\sin(\omega_i\,t + \phi_i) \tag{20.7}$$

The order of the mode m_i does not come into this equation but will appear in equations below. The dimensionless overall amplitude scale A_a is introduced as a convenient way to scale the amplitude of all the wakes with just one parameter. Normally, for a wake that has a well defined mode, The phase factor ϕ_i is zero. Finite ϕ_i is used for simulations of such things as the long-range resistive wall wake. In this case, the resistive wall wake needs to be modeled as the sum of a number of modes since the resistive wall wake is not well modeled by a single damped sinusoid.

The mode strength $(R/Q)_i$ in the above equation has units of Ohms/meter^{2m_i}. Notice that R/Q is defined so that it includes the cavity length. Thus the long–range wake equations, as opposed to the short–range ones, do not have any explicit dependence on L. To make life more interesting, different people define R/Q differently. A common practice is to define an R/Q "at the beam pipe radius". In this case the above equations must be modified to include factors of the beam pipe radius. Another convention uses a "linac definition" which makes R/Q twice as large and adds a factor of 2 in Eq. (20.7) to compensate.

Note: Originally, *Bmad* characterized the damping factor d_i using the quality factor Q_i via the relationship

$$d_i = \frac{\omega_i}{2\,Q_i} \tag{20.8}$$

This proved to be inconvenient when modeling such things as the resistive wall wake (where it is convenient to have modes where $\omega_i = 0$) so the lattice file syntax was modified to use d_i directly.

Assuming that the macroparticle generating the wake is offset a distance r_w along the x–axis, a trailing macroparticle at transverse position (r, θ) will see a kick

$$\begin{aligned} \Delta\mathbf{p}_\perp &= -C\,I_m\,W(t)\,m\,r^{m-1}\,\left(\hat{\mathbf{r}}\cos m\theta - \hat{\theta}\sin m\theta\right) &(20.9)\\ &= -C\,I_m\,W(t)\,m\,r^{m-1}\,\left(\hat{\mathbf{x}}\cos[(m-1)\theta] - \hat{\mathbf{y}}\sin[(m-1)\theta]\right) \\ \Delta p_z &= -C\,I_m\,W'(t)\,r^m\,\cos m\theta &(20.10)\end{aligned}$$

where in this, and other equations below, the subscript i has been dropped. C is given by

$$C = \frac{e}{c\,P_0} \tag{20.11}$$

and

$$I_m = q_w\,r_w^m \tag{20.12}$$

with q_w being the magnitude of the charge on the particle. Generalizing the above, a macroparticle at (r_w, θ_w) will generate a wake

$$\begin{aligned} -\Delta p_x + i\Delta p_y &= C\,I_m\,W(t)\,m\,r^{m-1}\,e^{-im\theta_w}\,e^{i(m-1)\theta} &(20.13)\\ \Delta p_z &= C\,I_m\,W'(t)\,r^m\,\cos[m(\theta-\theta_w)] &(20.14)\end{aligned}$$

Comparing Eq. (20.13) to (18.8), and using the relationship between kick and field as given by (18.4) and (18.5), shows that the form of the wakefield transverse kick is the same as for a multipole of order $n = m - 1$.

The wakefield felt by a particle is due to the wakefields generated by all the particles ahead of it. If the wakefield kicks are computed by summing over all particle pairs, the computation will scale as N^2 where N is the number of particles. This quickly becomes computationally exorbitant. A better solution is to keep track of the wakes in a cavity. When a particle comes through, the wake it generates is simply added to the existing wake. This computation scales as N and makes simulations with large number of particles practical.

To add wakes together, a wake must be decomposed into its components. Spatially, there are normal and skew components and temporally there are sin and cosine components. This gives 4 components which will be labeled $a_{\cos}$, $a_{\sin}$, $b_{\cos}$, and $b_{\sin}$. For a mode of order m, a particle passing through at a

time t_w with respect to the reference particle will produce wake components

$$\begin{aligned}
\delta a_{\sin} &\equiv \quad c\,A_a \left(\frac{R}{Q}\right) e^{d\,t_w}\, \cos(\omega\, t_w)\, I_m\, \sin(m\theta_w) \\
\delta a_{\cos} &\equiv -\,c\,A_a \left(\frac{R}{Q}\right) e^{d\,t_w}\, \sin(\omega\, t_w)\, I_m\, \sin(m\theta_w) \\
\delta b_{\sin} &\equiv \quad c\,A_a \left(\frac{R}{Q}\right) e^{d\,t_w}\, \cos(\omega\, t_w)\, I_m\, \cos(m\theta_w) \\
\delta b_{\cos} &\equiv -\,c\,A_a \left(\frac{R}{Q}\right) e^{d\,t_w}\, \sin(\omega\, t_w)\, I_m\, \cos(m\theta_w)
\end{aligned} \tag{20.15}$$

These are added to the existing wake components. The total is

$$a_{\sin} = \sum_{\text{particles}} \delta a_{\sin} \tag{20.16}$$

with similar equations for $a_{\cos}$ etc. Here the sum is over all particles that cross the cavity before the kicked particle. To calculate the kick due to wake, the normal and skew components are added together

$$\begin{aligned}
a_{tot} &= e^{-d\,t}\, (a_{\cos}\, \cos(\omega\, t + \phi) + a_{\sin}\, \sin(\omega\, t + \phi)) \\
b_{tot} &= e^{-d\,t}\, (b_{\cos}\, \cos(\omega\, t + \phi) + b_{\sin}\, \sin(\omega\, t + \phi))
\end{aligned} \tag{20.17}$$

Here t is the passage time of the particle with respect to the reference particle. In analogy to Eq. (20.13) and (20.14), the kick is

$$-\Delta p_x + i\Delta p_y = C\,m\,(b_{tot} + i a_{tot})\, r^{m-1}\, e^{i(m-1)\theta} \tag{20.18}$$

$$\Delta p_z = -C\, r^m\, \left((b'_{tot} + i a'_{tot}) e^{im\theta} + (b'_{tot} - i a'_{tot}) e^{-im\theta}\right) \tag{20.19}$$

where $a' \equiv da/dt$ and $b' \equiv db/dt$.

When simulating trains of bunches, the exponential factor $d\,t_w$ in Eq. (20.15) can become very large. To prevent numerical overflow, *Bmad* uses a reference time $z_{\rm ref}$ so that all times t in the above equations are replaced by

$$t \longrightarrow t - t_{\rm ref} \tag{20.20}$$

The above equations were developed assuming cylindrical symmetry. With cylindrical symmetry, the cavity modes are actually a pair of degenerate modes. When the symmetry is broken, the modes no longer have the same frequency. In this case, one has to consider a mode's polarization angle θ_p. Equations (20.17) and (20.18) are unchanged. In place of Eq. (20.15), the contribution of a particle to a mode is

$$\begin{aligned}
\delta a_{\sin} &= \quad c\,A_a \left(\frac{R}{Q}\right) e^{d\,t_w}\, \cos(\omega\, t_w)\, I_m\, \left[\sin(m\theta_w)\, \sin^2(m\theta_p) + \cos(m\theta_w)\, \sin(m\theta_p)\, \cos(m\theta_p)\right] \\
\delta a_{\cos} &= -\,c\,A_a \left(\frac{R}{Q}\right) e^{d\,t_w}\, \sin(\omega\, t_w)\, I_m\, \left[\sin(m\theta_w)\, \sin^2(m\theta_p) + \cos(m\theta_w)\, \sin(m\theta_p)\, \cos(m\theta_p)\right] \\
\delta b_{\sin} &= \quad c\,A_a \left(\frac{R}{Q}\right) e^{d\,t_w}\, \cos(\omega\, t_w)\, I_m\, \left[\cos(m\theta_w)\, \cos^2(m\theta_p) + \sin(m\theta_w)\, \sin(m\theta_p)\, \cos(m\theta_p)\right] \\
\delta b_{\cos} &= -\,c\,A_a \left(\frac{R}{Q}\right) e^{d\,t_w}\, \sin(\omega\, t_w)\, I_m\, \left[\cos(m\theta_w)\, \cos^2(m\theta_p) + \sin(m\theta_w)\, \sin(m\theta_p)\, \cos(m\theta_p)\right]
\end{aligned} \tag{20.21}$$

Note: Technically an unpolarized mode is actually two polarized modes perpendicular to each other. The axes of these two normal modes can be chosen arbitrary as long as they are at right angles.

Chapter 21

Multiparticle Simulation

Bmad has routines for tracking two types of objects called "`particles`" and "`macroparticles`". `Particles` are characterized by a six-vector representing the particle's phase space coordinates and a pair of complex numbers characterizing the particle's spin. A macroparticle is like a particle with the addition of a 6×6 "sigma" matrix characterizing the size of the macroparticle.

Macroparticle tracking was implemented in *Bmad* in order to simulate particle bunches. The idea was that far fewer macroparticles than particles would be needed to characterize a bunch. In practice, it was found that the complexity of handling the macroparticle sigma matrix more than offset the reduction in the number of particles needed. Hence, while the basic macroparticle tracking routines still exist, macroparticle tracking is not currently maintained and the use of this code is discouraged. However macroparticle tracking could be revived in the future if there is a demonstrated need for it.

Particle tracking can be divided into "single particle" tracking and "beam" tracking. Single particle tracking is simply tracking a single particle. Beam tracking is tracking an ensemble of particles divided up into a number of bunches that make up a "beam".

21.1 Bunch Initialization

[Developed by Michael Saelim]

To better visualize the evolution of a particle beam, it is sometimes convenient to initialize the beam with the particles regularly spaced. The following two algorithms are implemented in *Bmad* for such a purpose.

See Chapter `c:beam.init` for details on the standard input format used by *Bmad* based programs for reading in bunch initialization parameters.

21.1.1 Elliptical Phase Space Distribution

To observe nonlinear effects on the beam, it is sometimes convenient to initialize a bunch of particles in a way that puts more particles in the tails of the bunch than one would normally have with the standard method of seeding particles using a Gaussian distribution. In order to preserve the emittance, a distribution with more particles in the tail needs to decrease the charge per tail particle relative to the core. This feature, along with a regular distribution, are contained in the following **`"ellipse"`**

distribution algorithm.

Consider the two dimensional phase space (x, p_x). The transformation to action-angle coordinates, (J, ϕ), is

$$J = \frac{1}{2}[\gamma x^2 + 2\alpha\, x\, p + \beta p^2] \tag{21.1}$$

$$\tan\phi = \frac{-\beta\,(p + \alpha\, x)}{x} \tag{21.2}$$

The inverse is

$$\begin{pmatrix} x \\ p \end{pmatrix} = \sqrt{2J} \begin{pmatrix} \sqrt{\beta} & 0 \\ -\frac{\alpha}{\sqrt{\beta}} & -\frac{1}{\sqrt{\beta}} \end{pmatrix} \begin{pmatrix} \cos\phi \\ \sin\phi \end{pmatrix}. \tag{21.3}$$

In action-angle coordinates, the normalized Gaussian phase space distribution, $\rho(J, \phi)$, is

$$\rho(J, \phi) = \frac{1}{2\pi\varepsilon} e^{-\frac{J}{\varepsilon}}. \tag{21.4}$$

where the emittance ε is just the average of J over the distribution

$$\varepsilon = \langle J \rangle \equiv \int dJ\, d\phi\, J\rho(J, \phi). \tag{21.5}$$

The beam sizes are:

$$\sigma_x^2 \equiv \langle x^2 \rangle = \varepsilon\beta \tag{21.6}$$

$$\sigma_p^2 \equiv \langle p^2 \rangle = \varepsilon\gamma, \tag{21.7}$$

and the covariance is

$$\langle x\, p \rangle = -\varepsilon\alpha. \tag{21.8}$$

The `ellipse` algorithm starts by partitioning phase space into regions bounded by ellipses of constant $J = B_n$, $n = 0, \ldots N_J$. The boundary values B_n are chosen so that, except for the last boundary, the $\sqrt{B_n}$ are equally spaced

$$B_n = \begin{cases} \frac{\varepsilon}{2}\left(\frac{n_\sigma\, n}{N}\right)^2 & \text{for } 0 \le n < N_J \\ \infty & \text{for } n = N_J \end{cases} \tag{21.9}$$

where n_σ is called the **`'boundary sigma cutoff'`**. Within each region, an elliptical shell of constant J_n is constructed with N_ϕ particles equally spaced in ϕ. The charge q_n of each particle of the n^{th} ellipse is chosen so that the total charge of all the particles of the ellipse is equal to the total charge within the region

$$N_\phi\, q_n = \int_{B_{n-1}}^{B_n} dJ \int_0^{2\pi} d\phi\, \rho(J, \phi) = \exp\left(-\frac{B_{n-1}}{\varepsilon}\right) - \exp\left(-\frac{B_n}{\varepsilon}\right) \tag{21.10}$$

The value of J_n is chosen to coincide with the average J within the region

$$N_\phi\, q_n\, J_n = \int_{B_{n-1}}^{B_n} dJ \int_0^{2\pi} d\phi\, J\, \rho(J, \phi) = \varepsilon(\xi + 1)e^{-\xi}\Big|_{\frac{B_n}{\varepsilon}}^{\frac{B_{n-1}}{\varepsilon}} \tag{21.11}$$

The `ellipse` phase space distribution is thus

$$\rho_{model}(J, \phi) = q_{tot} \sum_{n=1}^{N_J} q_n\, \delta(J - J_n) \sum_{m=1}^{N_\phi} \delta(\phi - 2\pi\frac{m}{N_\phi}) \tag{21.12}$$

where q_{tot} is the total charge. At a given point in the lattice, where the Twiss parameters are known, the input parameters needed to construct the `ellipse` phase space distribution is n_σ, N_J, N_ϕ, and q_{tot}.

The `ellipse` distribution is two dimensional in nature but can easily be extended to six dimensions.

21.1.2 Kapchinsky-Vladimirsky Phase Space Distribution

The Kapchinsky-Vladimirsky (`KV`) distribution can be thought of as a four dimensional analog of the `ellipse` distribution with only one elliptical shell. Consider a 4D phase space (x, x', y, y'). Using this framework, a 4D Gaussian distribution is

$$\rho(J_x, \phi_x, J_y, \phi_y) = \frac{1}{(2\pi)^2 \varepsilon_x \varepsilon_y}\, exp(-\frac{J_x}{\varepsilon_x})\, exp(-\frac{J_y}{\varepsilon_y}) \tag{21.13}$$

$$= \frac{1}{(2\pi)^2 \varepsilon_x \varepsilon_y}\, exp(-\frac{I_1}{\varepsilon}), \tag{21.14}$$

where the orthogonal action coordinates are:

$$I_1 = \left(\frac{J_x}{\varepsilon_x} + \frac{J_y}{\varepsilon_y}\right)\varepsilon \tag{21.15}$$

$$I_2 = \left(-\frac{J_x}{\varepsilon_y} + \frac{J_y}{\varepsilon_x}\right)\varepsilon \tag{21.16}$$

with $\varepsilon = (\frac{1}{\varepsilon_x^2} + \frac{1}{\varepsilon_y^2})^{-1/2}$. The reverse transformation is:

$$J_x = \left(\frac{I_1}{\varepsilon_x} - \frac{I_2}{\varepsilon_y}\right)\varepsilon \tag{21.17}$$

$$J_y = \left(\frac{I_1}{\varepsilon_y} + \frac{I_2}{\varepsilon_x}\right)\varepsilon. \tag{21.18}$$

The `KV` distribution is

$$\rho(I_1, I_2, \phi_x, \phi_y) = \frac{1}{A}\delta(I_1 - \xi), \tag{21.19}$$

where $A = \frac{\varepsilon_x \varepsilon_y}{\varepsilon^2}\xi(2\pi)^2$ is a constant which normalizes the distribution to 1. By choosing a particular ξ, and iterating over the domain of the three remaining coordinates, one can populate a 3D subspace of constant density.

The range in I_2 to be iterated over is constrained by $J_x,\ J_y \geq 0$. Thus $I_2 is in the range [-\frac{\varepsilon_x}{\varepsilon_y} I_1, \frac{\varepsilon_y}{\varepsilon_x} I_1]$. This range is divided into N regions of equal size, with a ring of particles placed in the middle of each region. The angle variables are also constrained to $\phi_x, \phi_y \in [0, 2\pi]$, with each range divided into M_x and M_y regions, respectively. Each of these regions will have a particle placed in its center.

The weight of a particle is determined by the total weight of the region of phase space it represents. Because the density ρ is only dependent on I_1,

$$q = \int_0^\infty dI_1 \int_{I_2}^{I_2+\Delta I_2} dI_2 \int_{\phi_x}^{\phi_x+\Delta\phi_x} d\phi_x \int_{\phi_y}^{\phi_y+\Delta\phi_y} d\phi_y\ \frac{1}{A}\delta(I_1 - \xi) \tag{21.20}$$

$$= \frac{1}{A}\Delta I_2 \Delta\phi_x \Delta\phi_y. \tag{21.21}$$

To represent the distribution with particles of equal weight, we must partition (I_2, ϕ_x, ϕ_y)-space into regions of equal volume.

The weight of each particle is

$$q = \frac{1}{N M_x M_y} = \frac{1}{N_{tot}} \tag{21.22}$$

where N_{tot} is the total number of particles

21.2 Touschek Scattering

[Developed by Michael Ehrlichman]

Touschek scattering occurs when a single scattering event between two particles in the same beam transfers transverse momentum to longitudinal momentum, and the resulting change in longitudinal momentum results in the loss of one or both particles. In the case of storage rings, these losses impose a beam lifetime. In low-emittance storage rings, Touschek scattering can be the dominant mechanism for particle loss. In the case of linear accelerators, these losses generate radiation in the accelerator tunnel. When the scattered particles collide with the beam chamber, x-rays are produced which can damage equipment and impose a biohazard. Studies of Touschek scattering typically look at beam lifetime and locations where scattering occurs and where particles are lost.

A commonly utilized theory for studying Touschek scattering is from Piwinski [Piwin98]. A basic outline of the derivation is,

1. Scatter two particles from a bunch in their COM frame using the relativistic Moller cross-section.
2. Boost from COM frame to lab frame. Changes to longitudinal momentum end up amplified by a factor of γ.
3. Integrate over 3D Gaussian distribution of particle positions and angles.

During the derivation many approximations are made which lead to a relatively simple formula. The integration is set up such that only those collisions which will result in particle loss are counted. The formula takes the momentum aperture as a parameter. The resulting formula is reproduced here to give the reader an idea of what influences the scattering rate, and how one might go about evaluating the formula,

$$R = \frac{r_e^2 c \beta_x \beta_y \sigma_h N_p^2}{8\sqrt{\pi}\beta^2\gamma^4\sigma_{x\beta}^2\sigma_{y\beta}^2\sigma_s\sigma_p}\int_{\tau_m}^{\infty}\left(\left(2+\frac{1}{\tau}\right)^2\left(\frac{\tau/\tau_m}{1+\tau}-1\right)+1-\frac{\sqrt{1+\tau}}{\sqrt{\tau/\tau_m}}\right.$$
$$\left.-\frac{1}{2\tau}\left(4+\frac{1}{\tau}\right)\ln\frac{\tau/\tau_m}{1+\tau}\right)\frac{\sqrt{\tau}}{\sqrt{1+\tau}}e^{-B_1\tau}I_0\left[B_2\tau\right]d\tau, \quad (21.23)$$

where $\tau_m = \beta^2\delta_m^2$ and δ_m is the momentum aperture. This formula gives the rate at which particles are scattered out of the bunch. It is assumed that two particles are lost per scattering event, one with too much energy and one with too little energy. If a machine with an asymmetric momentum aperture is being studied, then the formula should be evaluated twice, once for each aperture, and the results averaged. Refer to [Piwin98] for definitions of the parameters involved. This formula is implemented in BMAD as part of the `touschek_mod` module.

Different formulas for calculating the Touschek scattering rate exist elsewhere in the literature. For example, Wiedemann [Wiede99], presents a formula with a simpler integrand. This formula, originally from a paper by LeDuff [Duff87], is derived in a fashion similar to Piwinski except that the formula does not take dispersion into account and uses a non-relativistic scattering cross-section. Since Piwinski's formula is the most robust, it is the one used in *Bmad*.

Particles are lost from Touschek scattering due to two effects. In storage rings, there is a momentum aperture defined by the RF system that is often referred to as the RF bucket. If the δp imparted by a Touschek scattering event exceeds this RF bucket, then the particle will no longer undergo synchrotron oscillations with the rest of the bunch and will coast through the accelerator. Second, if the Touschek scattering event occurs in a dispersive region, the scattered particles will take on a finite J and undergo

betatron oscillations. These oscillations can be large in amplitude and may cause the particles to collide with the beam pipe. To first order, the amplitude of J due to a scattering event that imparts a momentum deviation of Δp is,

$$J \approx \gamma_0 \mathcal{H}_0 \frac{\Delta p^2}{2}, \tag{21.24}$$

where γ_0 is relativistic γ and $\mathcal{H}_0$ is the dispersion invariant.

21.3 Macroparticles

Note: The macroparticle tracking code is not currently maintained in favor of tracking an ensemble of particles where each particle is specified by a position without a sigma matrix. The following is present for historical reference only.

A macroparticle[Brown77] is represented by a centroid position $\bar{\mathbf{r}}$ and a 6×6 $\boldsymbol{\sigma}$ matrix which defines the shape of the macroparticle in phase space. $\sigma_i = \sqrt{\boldsymbol{\sigma}(i,i)}$ is the RMS sigma for the i^{th} phase space coordinate. For example $\sigma_z = \sqrt{\boldsymbol{\sigma}(5,5)}$.

$\boldsymbol{\sigma}$ is a real, non-negative symmetric matrix. The equation that defines the ellipsoid at a distance of n–sigma from the centroid is

$$(\mathbf{r} - \bar{\mathbf{r}})^t \boldsymbol{\sigma}^{-1} (\mathbf{r} - \bar{\mathbf{r}}) = n \tag{21.25}$$

where the t superscript denotes the transpose. Given the sigma matrix at some point $s = s_1$, the sigma matrix at a different point s_2 is

$$\boldsymbol{\sigma}_2 = \mathbf{M}_{12}\, \boldsymbol{\sigma}_1\, \mathbf{M}_{12}^t \tag{21.26}$$

where $\mathbf{M}_{12}$ is the Jacobian of the transport map from point s_1 to s_2.

The Twiss parameters can be calculated from the sigma matrix. The dispersion is given by

$$\begin{aligned} \sigma(1,6) &= \eta_x\, \sigma(6,6) \\ \sigma(2,6) &= \eta_x'\, \sigma(6,6) \\ \sigma(3,6) &= \eta_y\, \sigma(6,6) \\ \sigma(4,6) &= \eta_y'\, \sigma(6,6) \end{aligned} \tag{21.27}$$

Ignoring coupling for now, the betatron part of the sigma matrix can be obtained from the linear equations of motion. For example, using

$$x = \sqrt{2\, \beta_x\, \epsilon_x} \cos \phi_x + \eta_x\, p_z \tag{21.28}$$

Solving for the first term on the RHS, squaring and averaging over all particles gives

$$\beta_x\, \epsilon_x = \sigma(1,1) - \frac{\sigma^2(1,6)}{\sigma(6,6)} \tag{21.29}$$

It is thus convenient to define the betatron part of the sigma matrix

$$\sigma_\beta(i,j) \equiv \sigma(i,j) - \frac{\sigma(i,6)\, \sigma(j,6)}{\sigma(6,6)} \tag{21.30}$$

and in terms of the betatron part the emittance is

$$\epsilon_x^2 = \sigma_\beta(1,1)\, \sigma_\beta(2,2) - \sigma_\beta^2(1,2) \tag{21.31}$$

and the Twiss parameters are

$$\epsilon_x \begin{pmatrix} \beta_x & -\alpha_x \\ -\alpha_x & \gamma_x \end{pmatrix} = \begin{pmatrix} \sigma_\beta(1,1) & \sigma_\beta(1,2) \\ \sigma_\beta(1,2) & \sigma_\beta(2,2) \end{pmatrix} \tag{21.32}$$

If there is coupling, the transformation between the 4×4 transverse normal mode sigma matrix $\boldsymbol{\sigma}_a$ and the 4×4 laboratory matrix $\boldsymbol{\sigma}_x$ is

$$\boldsymbol{\sigma}_x = \mathbf{V}\, \boldsymbol{\sigma}_a \mathbf{V}^t \tag{21.33}$$

where $\mathbf{V}$ is given by Eq. (23.5).

The sigma matrix is the same for all macroparticles and is determined by the local Twiss parameters:

$$\begin{aligned} \sigma(1,1) &= \epsilon_x\, \beta_x \\ \sigma(1,2) &= -\epsilon_x \alpha_x \\ \sigma(2,2) &= \epsilon_x\, \gamma_x = \epsilon_x\, (1+\alpha_x^2)/\beta_x \\ \sigma(3,3) &= \epsilon_y\, \beta_y \\ \sigma(3,4) &= -\epsilon_y \alpha_b \\ \sigma(3,4) &= \epsilon_y\, \gamma_y = \epsilon_y\, (1+\alpha_b^2)/\beta_y \\ \sigma(i,j) &= 0 \quad \text{otherwise} \end{aligned} \tag{21.34}$$

The centroid energy of the k^{th} macroparticle is

$$E_k = E_b + \frac{(n_{mp} - 2\,k + 1)\, \sigma_E\, N_{\sigma E}}{n_{mp}} \tag{21.35}$$

where E_b is the central energy of the bunch, n_{mp} is the number of macroparticles, σ_E is the energy sigma, and $N_{\sigma E}$ is the number of sigmas in energy that the range of macroparticle energies cover. The charge of each macroparticle is, within a constant factor, the charge contained within the energy region $E_k - dE_{mp}/2$ to $E_k + dE_{mp}/2$ assuming a Gaussian distribution where the energy width dE_{mp} is

$$dE_{mp} = \frac{2\, \sigma_E\, N_{\sigma E}}{n_{mp}} \tag{21.36}$$

21.4 Space Charge and Coherent Synchrotron Radiation

The electric field $\mathbf{E}$ felt by particle A due to particle B can be described using the Liénard-Wiechert formula [Sagan09]. The field is singular as the distance between particles goes to zero so one approach to handling this is to decompose the field into two parts: One part, called the "space charge" (SC) or "Coulomb" term, $\mathbf{E}_{SC}$ is the field that would result if the particles where moving without acceleration along a straight line. The "Coherent Synchrotron Radiation" (CSR) term $\mathbf{E}_{CSR}$ is everything else $\mathbf{E}_{CSR} \equiv \mathbf{E} - \mathbf{E}_{SC}$. Generally, the longitudinal component of the SC kick is negligible compared to the CSR kick at large enough particle energies.

The SC term is singular at small distances while the CSR term is not. This being the case, it is possible to model the CSR term using a 1-dimensional formalism where the beam is approximated as a line charge[Sagan09, Sagan17]. In this formalism, the CSR kick is strictly longitudinal.

Transport through a lattice element with SC and CSR involves a beam of particles. The lattice element is divided up into a number of slices. Transport through a slice is a two step process. The first step is to give all the particles a kick due to SC and CSR. The second step is transport of all particles from one slice to the next without any interaction between particles. User settable parameters pertinent to the CSR calculation are listed in §11.5.

21.4.1 1_Dim CSR Calculation

When an element's `csr_method` is set to `1_dim` (§6.5), The particle-particle CSR kick is calculated by dividing the bunch longitudinally into a number of bins. To smooth the computed bin densities, each particle of the bunch is considered to have a triangular density distribution as shown in Fig. 21.1. The particle density of a bin is calculated by summing the contribution from all the particles. The contribution of a given particle to a given bin is calculated from the overlap of the particle's triangular density distribution with the bin. For the CSR kick, the density is actually calculated for a second set of staggered bins that have been offset by 1/2 the bin width with respect to the first set. This gives the density at the edges of the original set of bins. The density is considered to vary linearly between the computed density points. For a description of the parameters that affect the CSR calculation see Section §11.5.

21.4.2 Slice Space Charge Calculation

When an element's `space_charge_method` is set to `slice` (§6.5), the calculation of the SC kick uses, the same particle binning as is used with the `1_dim` CSR calculation (§21.4.1). The kick is divided into longitudinal and a transverse parts. The transverse part uses the same Bassetti–Erskine complex error function formula[Talman87] as with the beam-beam interaction (§26.5) except here, since all the particles are moving in the same direction, the kicks due to the electric and magnetic fields generated

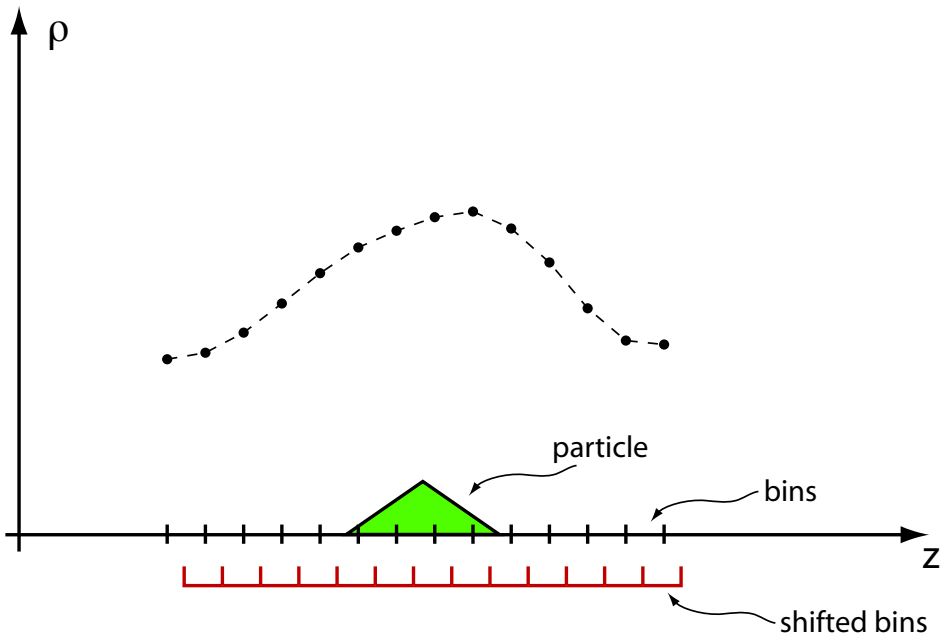

Figure 21.1: The Coherent Synchrotron Radiation kick is calculated by dividing longitudinally a bunch into a number of bins. To smooth the computed densities, each particle of the bunch is considered to have a triangular density distribution.

by a given particle tend to cancel

$$K_y(\text{CS}) + i\,K_x(\text{CS}) = \frac{r_e\,\rho(z)}{\gamma^3\,e}\cdot\sqrt{\frac{2\,\pi\,(\sigma_x+\sigma_y)}{\sigma_x-\sigma_y}} \left\{ w\left[\frac{x+i\,y}{\sqrt{2(\sigma_x^2-\sigma_y^2)}}\right] - \exp\left[-\frac{x^2}{2\,\sigma_x^2}-\frac{y^2}{2\,\sigma_y^2}\right]\cdot w\left[\frac{x\,\frac{\sigma_y}{\sigma_x}+i\,y\,\frac{\sigma_x}{\sigma_y}}{\sqrt{2(\sigma_x^2-\sigma_y^2)}}\right]\right\} \tag{21.37}$$

where $K(\text{CS})$ is the CS kick per unit length of travel of the beam, $\rho(z)$ is the density of particles per unit length evaluated at the z position of the kicked particle, e is the charge on the electron, and w is the complex error function.

The longitudinal SC kick is given by Eq. (31) of Sagan[Sagan09]

$$dK_{\text{SC}} = \frac{r_c mc^2\,\text{sign}(\zeta)\rho(z')dz'}{\sigma_x\,\sigma_y\,\exp\left[\frac{x^2}{2\,\sigma_x^2}+\frac{y^2}{2\,\sigma_y^2}\right]+\frac{\sigma_x^2+\sigma_y^2}{\sigma_x+\sigma_y}\,\gamma|\zeta|+\gamma^2\zeta^2}\ , \tag{21.38}$$

where ζ is the longitudinal distance between the kick point and the slice doing the kicking. There are two simulation modes for the longitudinal SC kick. In both these modes, the kick is evaluated at the center plane of each slice. The kick is a sum kicks from all the slices. Since the thickness of the slices is, in general, not negligible, the integral over a slice is used to calculate the kick. The total kick $K_{\text{SC}}(j)$ at slice j is

$$K_{\text{SC}}(j) = \sum_i \int_{\zeta_{ij}-dz_s/2}^{\zeta_{ij}+dz_s/2} d\zeta\,dK_{\text{SC}} \tag{21.39}$$

where the sum is over all slices i, ζ_{ij} is the distance between slices i and j, and dz_s is the slice thickness. An analytic expression of the above integral is easily calculated assuming that the charge density $\rho(z)$ is linearly varying within a given slice. For brevity's sake, the calculation is not explicitly presented here. Once the kick at the slice center planes is calculated, the kick given to a particle is calculated using linear interpolation.

One mode for calculating the transverse SC kick which is computationally fast, ignores the transverse dependence of the kick and just evaluates the kick on the beam centerline. The other simulation mode represents the kick due to a given slice using a Padé approximant of form

$$\int_{\zeta_{ij}-dz_s/2}^{\zeta_{ij}+dz_s/2} d\zeta\,dK_{\text{SC}} \simeq \frac{1}{a_{00}+a_{20}x^2+a_{40}x^4+a_{02}y^2+a_{04}y^4+a_{22}x^2\,y^2} \tag{21.40}$$

the `a_mn` are calculated from an analytic formula derived from integrating Eq. (21.38). The reason for using this form is that it is a reasonable approximation even for very large x or y in that the actual and approximate kick both go to zero in this limit. That this Padé approximant is reasonable is dependent upon the fact that all the a_{mn} for a slice are either all positive or all negative. Kicks from different slices can be combined using standard Differential Algebra techniques to give a summed kick in the same form as above. To avoid divergences, for a given j where the kick is evaluated, all the kicks from slices with negative coefficients are combined together and all the kicks from slices with positive coefficients are combined together and the total kick is then the sum of the "positive kick" part and the "negative kick" part. The kick applied to a particle is calculated by first evaluating the kick, at the particle's x and y, at the neighboring slices and then using linear interpolation.

Note: Match elements (§4.34) can have orbit shifts which are not well handled by the CSR algorithm. For this reason, match elements are ignored in the CSR calculation.

21.4.3 FFT_3D Space Charge Calculation

When an element's `space_charge_method` is set to `fft_3d` or `cathode_fft_3d` (§6.5), the space charge calculation uses code from the OpenSC package developed by Rob Ryne and Christopher Mayes [Ryne18]. The method works by calculating the field due to the particles deposited on a 3D grid using an integrated Green function method for the Poisson equation. The steps are:

1. Deposit weighted charged particles on a 3D rectangular grid.
2. Calculate the space charge fields on this grid by FFT convolution.
3. Interpolate the field to an arbitrary point within its domain.

The FFT convolution is done using FFTs from the FFTW package, is parallelized using OpenMP. Special options allow the consideration of image charges at a cathode. This method will be able to handle lower energy bunches than the `slice` method (§21.4.2) the disadvantage is that the `fft_3d` and `cathode_fft_3d` methods will be slower.

Note: The mesh size is set by the `csr_pram` parameter `space_charge_mesh_size` (§11.5).

21.5 High Energy Space Charge

Bmad has a code module for simulating the effect of space charge (SC) at high energies. This is separate from the regular space charge calculation of §21.4. Thus it should be noted that turning on of both the regular space charge and the high energy space charge in the same element will result in double counting of the space charge effect.

The advantage of the high energy space charge algorithm is that the kick on a given particle is computed assuming a Gaussian beam with the beam size calculated using emittances supplied by the user. Thus the high energy space charge calculation can be done in single particle tracking (§21) as opposed to the beam tracking that must be used for the regular space charge calculation. The other advantage is that the high energy space charge calculation is quick since it is assumed that the kick is small enough so that the kick is only applied once per lattice element. The disadvantage of the high energy space charge calculation is that there is the assumption that the beam distribution is Gaussian which is generally acceptable for storage rings at relatively high energy but will not accurate in other situations.

If a *Bmad* based program has been constructed to use the high energy space charge module (the documentation for the program should indicate if this is true), the high energy space charge force can be turned on or off by setting the `bmad_com[high_energy_space_charge_on]` parameter (§11.2, §10.4). Currently, the `long_term_tracking` program implements high energy space charge while `tao` does not (`tao` is not constructed for long term tracking where these effects would be noticeable).

The high energy space charge kick is computed assuming a gaussian bunch shape

$$K_y + i\,K_x = \frac{r_e\,N}{\gamma^3\,\sigma_z}\,\exp\left[\frac{-z^2}{2\,\sigma_z^2}\right]\cdot\sqrt{\frac{\sigma_x+\sigma_y}{\sigma_x-\sigma_y}} \left\{ w\left[\frac{x+i\,y}{\sqrt{2(\sigma_x^2-\sigma_y^2)}}\right] - \exp\left[-\frac{x^2}{2\,\sigma_x^2}-\frac{y^2}{2\,\sigma_y^2}\right]\cdot w\left[\frac{x\,\frac{\sigma_y}{\sigma_x}+i\,y\,\frac{\sigma_x}{\sigma_y}}{\sqrt{2(\sigma_x^2-\sigma_y^2)}}\right]\right\} \tag{21.41}$$

where N is the number of particles in the bunch. This equation is similar to Eq. (21.37) except that $\rho(z)$ has been replaced assuming that the longitudinal distribution is Gaussian. For particles close to the bunch core the kick is linear with displacement giving rise to a tune shift [Decking00].

The high energy space charge calculation ignores any CSR effects and ignores any longitudinal kicks and is thus not a good approximation at lower energies. See the discussion in [Sagan09] for more details.

Chapter 22

Synchrotron Radiation

22.1 Radiation Damping and Excitation

Emission of synchrotron radiation by a particle can be decomposed into two parts. The deterministic average energy emitted produces damping while the stochastic fluctuating part of the energy spectrum produces excitation[Jowett87].

The treatment of radiation damping by *Bmad* essentially follows Jowett[Jowett87]. The energy loss at a given location is modeled via

$$\frac{\Delta E}{E_0} = -k_E \equiv -\left[k_d \left\langle g^2 \right\rangle L_p + \sqrt{k_f \left\langle g^3 \right\rangle L_p}\, \xi\right] (1 + p_z)^2 \tag{22.1}$$

where L_p is the actual path length, g is the bending strength ($1/g$ is the bending radius), and $\langle \ldots \rangle$ is an average over the actual path. In the above equation k_d gives the deterministic part of the emission, ξ is a Gaussian distributed random number with unit sigma and zero mean, and k_f is the amplitude of the stochastic part of the emission. Values for k_d and k_f are calculated via the equations

$$k_d = \frac{2\, r_c}{3}\, \gamma_0^3 \tag{22.2}$$

$$k_f = \frac{55\, r_c\, \hbar}{24\, \sqrt{3}\, m\, c}\, \gamma_0^5 \tag{22.3}$$

where γ_0 is the energy factor of an on-energy particle and r_c is the particles "classical radius" given by

$$r_c = \frac{q^2}{4\, \pi\, \epsilon_0\, m\, c^2} \tag{22.4}$$

where q is the particle's charge and m is the particle's mass.

Ignoring the finite opening angle of the emission for now, the angular orientation of the particle motion is invariant for forward directed emission which leads to the following equations for the changes in momentum phase space coordinates

$$\Delta p_x = -\frac{k_E}{1 + p_z}\, p_x, \qquad \Delta p_y = -\frac{k_E}{1 + p_z}\, p_y, \qquad \Delta p_z \approx \frac{\Delta E}{E_0} = -k_E \tag{22.5}$$

Synchrotron radiation emission involves energy loss and this energy loss leads to what is known as the energy "sawtooth" effect where the curve of particle energy on the closed orbit as a function of longitudinal

position has a sawtooth shape. A sawtooth pattern can also be generally seen in the horizontal orbit. It is sometimes convenient in simulations to eliminate the sawtooth effect. This can be done by shifting the photon emission spectrum at any given element to have zero average energy loss along the closed orbit. For this calculation the closed orbit should be the closed orbit as calculated without radiation damping (in other words the closed orbit without the sawtooth). In this case, k_E is calculated by

$$k_E = \left[k_d \left\langle g^2 \right\rangle L_p + \sqrt{k_f \left\langle g^3 \right\rangle L_p}\, \xi\right] (1+p_z)^2 - k_d \left\langle g_0^2 \right\rangle L_p \tag{22.6}$$

where g_0 is g evaluated along the closed orbit. In practice, for the calculation, *Bmad* approximates the closed orbit as the zero orbit.

The deterministic and stochastic parts of the emission can be included or excluded from a tracking simulation by setting in a lattice file the *Bmad* global parameters (§11.2)

```
bmad_com[radiation_damping_on]      = True or False  ! Deterministic part on/off.
bmad_com[radiation_fluctuations_on] = True or False  ! Stochastic part on/off.
bmad_com[radiation_zero_average]    = True or False  ! Make ave radiation kick zero.
```

The global parameter `bmad_com[radiation_zero_average]` controls the shifting of the photon spectrum to have zero average. Currently, the shifting of the spectrum only works for non PTC dependent tracking. That is, the shifting is not applicable to tracking with Taylor maps and with `symp_lie_ptc` (§6.1) tracking.

The fact that an emitted photon is not exactly colinear with the particle direction (often called the "vertical opening angle") can be modeled as a separate process from the energy loss. With this approximation, the change $\Delta p_\perp$ in the momentum transverse to the bending plane is given by

$$\Delta p_\perp = \sqrt{k_v \left\langle g^3 \right\rangle L_p}\, \xi \tag{22.7}$$

where the ξ in Eq. (22.7) is independent of the ξ in Eq. (22.1) and

$$k_v = \frac{13\, r_c\, \hbar}{24\sqrt{3}\, m\, c}\, \gamma_0^3 \tag{22.8}$$

22.2 Transport Map with Radiation Included

Transport maps which include radiation effects can be constructed[Ohmi94]. The first step is to calculate the reference orbit which is the closed orbit for lattices with a closed geometry and for lattices with an open geometry the reference orbit is the orbit for some given initial position. Orbits here are calculated with radiation damping but ignoring stochastic effects. The transfer map from s_1 to s_2 will be of the form

$$\delta\mathbf{r}_2 = \mathcal{M}_{21}(\delta\mathbf{r}_1) + \mathcal{S}_{21}\boldsymbol{\Xi} \tag{22.9}$$

where $\delta\mathbf{r}_1$ and $\delta\mathbf{r}_2$ are the particle positions with respect to the reference orbit at s_1 and s_2 respectively and $\mathcal{M}_{21}$ is the transfer map with damping. The stochastic radiation part is represented by a 6×6 matrix $\mathcal{S}$ times a 6-vector

$$\boldsymbol{\Xi} = (\xi_1, \xi_2, \xi_3, \xi_4, \xi_5, \xi_6) \tag{22.10}$$

with each ξ_i being an independent Gaussian distributed random number with unit sigma and zero mean. The stochastic transport (second term in Eq. (22.9)) is treated here only in lowest order. This is a good approximation as long as the radiation emitted is small enough in the region between s_1 and s_2. This is true for nearly all practical cases. In the case where this approximation fails, the equilibrium beam distribution would not be Gaussian and the standard radiation integral treatment (§22.3), which relies on this approximation, would not be valid.

The transfer map with damping $\mathcal{M}$ is calculated by adding in the effect of the damping (Eqs. (22.5)) when integrating the equations of motion to form the map. Through a given lattice element, it is generally very safe to assume that the change in energy is small compared to the energy of a particle. Thus the matrix $\mathbf{M}$ through an element, which is the first order part of $\mathcal{M}$, can is computed via first order perturbation theory to be

$$\mathbf{M} = \mathbf{T} + \mathbf{Z} \tag{22.11}$$

where $\mathbf{T}$ is the transfer matrix without damping and $\mathbf{Z}$ is the change in $\mathbf{T}$ due to damping computed via

$$\mathbf{Z} = \int_{s_1}^{s_2} ds\, \mathbf{T}_{2,s}\, \mathbf{d}(s)\, \mathbf{T}_{s,1} \tag{22.12}$$

where s_1 and s_2 are the longitudinal positions at the ends of the element and the local damping matrix $\mathbf{d}$ is computed from Eqs. (22.5)

$$\mathbf{d} = -k_d \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{dg^2}{dx} p_x (1+p_z) & g^2(1+p_z) & \frac{dg^2}{dy} p_x (1+p_z) & 0 & 0 & g^2 p_x \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{dg^2}{dx} p_y (1+p_z) & 0 & \frac{dg^2}{dy} p_y (1+p_z) & g^2(1+p_z) & 0 & g^2 p_y \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{dg^2}{dx} (1+p_z)^2 & 0 & \frac{dg^2}{dy} (1+p_z)^2 & 0 & 0 & 2g^2(1+p_z) \end{pmatrix} \tag{22.13}$$

All quantities are evaluated on the closed orbit. Notice that since $\mathcal{M}_{21}$ is computed with respect to the beam centroid orbit, there is no constant part to the map. Since $\mathbf{T}_{21}$ is invertible, Eq. (22.11) can be written in the form

$$\mathbf{M}_{21} = \left(\mathbf{1} + \mathbf{Z}_{21}\, \mathbf{T}_{21}^{-1}\right) \mathbf{T}_{21} \equiv \mathbf{D}_{21}\, \mathbf{T}_{21} \tag{22.14}$$

$\mathbf{D}$ is defined by this equation. The 1-turn damping decrement α for each mode a, b, and c of oscillation can be calculated from $\mathbf{D}$ using Eq. (86) of Ohmi[Ohmi94].

The $\mathcal{S}$ matrix (Eq. (22.9)) is calculated by first noting that, to linear order, the distribution of $\delta\mathbf{r}_2$ due to stochastic radiation over some length ds as some point s is

$$\delta\mathbf{r}_2 = \sqrt{ds}\, \mathbf{M}_{2,s}\, \left(\mathbf{F}_f(s)\, \xi_1 + \mathbf{F}_v\, \xi_2\right) \tag{22.15}$$

where $\mathbf{M}_{2,s}$ is the first order part (matrix) of the map $\mathcal{M}_{2,s}$ from s to s_2, ξ_1 and ξ_2 are two independent Gaussian random numbers with unit sigma and zero mean, and $\mathbf{F}_f$ and $\mathbf{F}_v$ are (see §22.1)

$$\mathbf{F}_f = \sqrt{k_f g_0^3}\, (0, p_x\,(1+p_z), 0, p_y\,(1+p_z), 0, (1+p_z)^2) \tag{22.16}$$

$$\mathbf{F}_v = \sqrt{k_v g_0}\, (0, -g_y, 0, g_x, 0, 0) \tag{22.17}$$

where k_f p_x, p_y and p_z are to be evaluated on the reference orbit and (g_x, g_y) is the curvature vector which points away from the center of curvature of the particle's orbit. Notice that since $\delta\mathbf{r}$ is, by definition, the deviation from the reference orbit, $p_x = r_2$ and $p_y = r_4$ will be zero on the reference orbit. The covariance matrix $\boldsymbol{\sigma}_\gamma$ is defined by $\sigma_{\gamma ij} \equiv \langle r_i\, r_j \rangle_\gamma$ where $\langle \ldots \rangle_\gamma$ is an average over the photon emission spectrum. The contribution, $\boldsymbol{\sigma}_{\gamma 21}$, to the covariance matrix at s_2 due to the stochastic emission over the region between s_1 and s_2, is

$$\boldsymbol{\sigma}_{\gamma 21} = \int_{s_1}^{s_2} ds\, \mathbf{M}_{2,s} \left[\mathbf{F}_f(s)\, \mathbf{F}_f^t(s) + \mathbf{F}_v(s)\, \mathbf{F}_v^t(s)\right] \mathbf{M}_{2,s}^t \tag{22.18}$$

where the t superscript indicates transpose. $\boldsymbol{\sigma}_{\gamma 21}$ is related to $\mathcal{S}$ via

$$\boldsymbol{\sigma}_{\gamma 21} = \mathcal{S}_{21}\, \mathcal{S}_{21}^t \tag{22.19}$$

The calculation of $\mathcal{S}_{21}$ involves calculating $\boldsymbol{\sigma}_{\gamma 21}$ via Eq. (22.18) and then using Eq. (22.19) to solve for $\mathcal{S}_{21}$ using, say, a Cholesky decomposition. Notice that while Eq. (22.19) does not have a unique solution, what matters here is that $\mathcal{S}_{21}\,\boldsymbol{\Xi}$ (see Eq. (22.9)) gives the correct distribution. The $\mathcal{S}_{21}$ matrix may contain columns or rows that are all zero. This can happen if there are vectors $\mathbf{z}$ where $\mathbf{z}^t \boldsymbol{\sigma}_{\gamma 21} \mathbf{z}$ is zero. For example, in a planer ring where the vertical emittance is zero there will be rows that are zero.

The covariance matrix $\boldsymbol{\sigma}_\gamma(s_2)$ at s_2 relative to the covariance matrix at s_1 is

$$\boldsymbol{\sigma}_\gamma(s_2) = \boldsymbol{\sigma}_{\gamma 21} + \mathbf{M}_{21}\,\boldsymbol{\sigma}_\gamma(s_1)\,\mathbf{M}_{21}^t \tag{22.20}$$

The beam size matrix $\boldsymbol{\sigma}$ is not the same as the covariance matrix since the beam size matrix is an average over the particles of a beam and not an average over the photon emission spectrum. However, in equilibrium, the two are the same. To calculate the equilibrium beam size matrix, Eq. (22.20) is recast. For any symmetric 6×6 matrix $\mathbf{A}$, define the 21-vector $\mathbf{V}(\mathbf{A})$ by

$$\mathbf{V}(\mathbf{A}) \equiv (A_{11}, A_{12}, \ldots, A_{16}, A_{22}, A_{23}, \ldots, A_{56}, A_{66}) \tag{22.21}$$

With $s_1 = s_2 = s$, and using Eq. (22.20), the equilibrium beam size matrix can be calculated via

$$\mathbf{V}(\boldsymbol{\sigma}(s)) = \mathbf{V}(\boldsymbol{\sigma}_{\gamma ss}) + \widetilde{\mathbf{M}}\,\mathbf{V}(\boldsymbol{\sigma}(s)) \tag{22.22}$$

where the 21×21 matrix $\widetilde{\mathbf{M}}$ is defined so that for any symmetric $\mathbf{A}$, $\widetilde{\mathbf{M}}\,\mathbf{V}(\mathbf{A}) = \mathbf{V}(\mathbf{M}\mathbf{A}\mathbf{M}^t)$. That is

$$\widetilde{\mathbf{M}} = \begin{pmatrix} M_{11}^2 & 2M_{11}M_{12} & \cdots & 2M_{15}M_{16} & M_{16}M_{16} \\ M_{11}M_{21} & 2M_{11}M_{22} & \cdots & 2M_{15}M_{26} & M_{16}M_{26} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ M_{51}M_{61} & 2M_{51}M_{62} & \cdots & 2M_{55}M_{66} & M_{56}M_{66} \\ M_{61}M_{61} & 2M_{61}M_{62} & \cdots & 2M_{65}M_{66} & M_{66}M_{66} \end{pmatrix} \tag{22.23}$$

Eq. (22.22) is linear in the unknown $\mathbf{V}(\boldsymbol{\sigma})$ and is easily solved.

The emittances can be calculated from the eigenvalues of the matrix $\boldsymbol{\sigma}\,\mathbf{S}$ (Wolski[Wolski06] Eq. 30) where $\mathbf{S}$ is given by Eq. (23.25). Specifically, the eigenvalues of $\boldsymbol{\sigma}\,\mathbf{S}$ are pure imaginary and, using the eigenvector ordering given by Eq. (23.24) (which is opposite that of Wolski), the emittances are the imaginary part of the odd eigenvalues $(\epsilon_a, \epsilon_b, \epsilon_c) = (\mathrm{Im}\,\lambda_1, \mathrm{Im}\,\lambda_3, \mathrm{Im}\,\lambda_5)$.

Unlike the case where radiation is ignored and the motion is symplectic, the calculated emittances along with the beam size will vary from point to point in a manner similar to the variation of the average beam energy (sawtooth effect) around the ring (§22.1).

The emittance calculation makes a number of approximations. One approximation is embodied in Eq. (22.12) which assumes that the damping is weak enough so that second order and higher terms can be neglected. Another approximation is that, within the extent of the beam, the damping as a function of transverse position is linear. That is, the effect of the damping is well represented by the matrix $\mathbf{d}$ in Eq. (22.12). The third major assumption is that, within the extent of the beam, the stochastic kick coefficient $\mathbf{F}_f$ (Eq. (22.17)) is independent of the transverse coordinates. Other approximations involve the assumption of linearity of the guide fields and the ignoring of any resonance or wakefield effects. To the extent that these assumptions are violated, this will lead to a non-Gaussian beam shape.

22.3 Synchrotron Radiation Integrals

The synchrotron radiation integrals can be used to compute emittances, the energy spread, etc. However, using the 6D damped and stochastic transport matrices (§22.2) has a number of advantages:

- Unlike the radiation integrals, the 6D calculation does not make the approximation that the synchrotron frequency is negligible. Therefore, the 6D calculation will be more accurate.
- The 6D calculation is simpler: Not as many integrals needed (only 2) and the 6D calculation does not depend upon calculation of the Twiss parameters.
- When doing any lattice design which involves constraining the emittances: Since the integrals of the 6D calculation are local (the integrations through any given lattice element are only dependent upon the properties of that lattice element), by caching integrals element-by-element, the computation of the emittances can be speeded up. That is, in a design problem, only the parameters of some subset of all the lattice elements will be varied (for example, a design may only involve varying the strength of quadrupoles), only this subset of elements needs to have their integrals recomputed. On the other hand, the radiation integrals are dependent on the Twiss, dispersion, and coupling parameters which make the integrals nonlocal.
- The 6D formalism can be used to construct transport maps with radiation damping and excitation for efficient particle tracking.

The standard radiation formulas assume no coupling between the horizontal and vertical plains[Helm73, Jowett87]. With coupling, the equations need to be generalized and this is detailed below.

In the general case, the curvature vector $\mathbf{g} = (g_x, g_y)$, which points away from the center of curvature of the particle's orbit and has a magnitude of $|\mathbf{g}| = 1/\rho$, where ρ is the radius of curvature (see Fig. 16.2), does not lie in the horizontal plane. Similarly, the dispersion $\boldsymbol{\eta} = (\eta_x, \eta_y)$ will not lie in the horizontal plane. With this notation, the synchrotron integrals for coupled motion are:

$$I_0 = \oint ds\, \gamma_0\, g \tag{22.24}$$

$$I_1 = \oint ds\, \mathbf{g} \cdot \boldsymbol{\eta} \equiv \oint ds\, (g_x\, \eta_x + g_y\, \eta_y) \tag{22.25}$$

$$I_2 = \oint ds\, g^2 \tag{22.26}$$

$$I_3 = \oint ds\, g^3 \tag{22.27}$$

$$I_{4a} = \oint ds\, \left[g^2\, \mathbf{g} \cdot \boldsymbol{\eta}_a + \nabla g^2 \cdot \boldsymbol{\eta}_a\right] \tag{22.28}$$

$$I_{4b} = \oint ds\, \left[g^2\, \mathbf{g} \cdot \boldsymbol{\eta}_b + \nabla g^2 \cdot \boldsymbol{\eta}_b\right] \tag{22.29}$$

$$I_{4z} = \oint ds\, \left[g^2\, \mathbf{g} \cdot \boldsymbol{\eta} + \nabla g^2 \cdot \boldsymbol{\eta}\right] \tag{22.30}$$

$$I_{5a} = \oint ds\, g^3\, \mathcal{H}_a \tag{22.31}$$

$$I_{5b} = \oint ds\, g^3\, \mathcal{H}_b \tag{22.32}$$

$$I_{6b} = \oint ds\, g^3\, \beta_b \tag{22.33}$$

where γ_0 is that usual relativistic factor and $\mathcal{H}_a$ is

$$\mathcal{H}_a = \gamma_a\, \eta_a^2 + 2\, \alpha_a\, \eta_a\, \eta_a' + \beta_a \eta_a'^2 \tag{22.34}$$

with a similar equation for $\mathcal{H}_b$. Here $\boldsymbol{\eta}_a = (\eta_{ax}, \eta_{ay})$, and $\boldsymbol{\eta}_b = (\eta_{bx}, \eta_{by})$ are the dispersion vectors for the a and b modes respectively in x–y space (these 2–vectors are not to be confused with the dispersion

4–vectors used in the previous section). The position dependence of the curvature function is:

$$\begin{aligned} g_x(x,y) &= g_x + x\,k_1 + y\,s_1 \\ g_y(x,y) &= g_y + x\,s_1 - y\,k_1 \end{aligned} \tag{22.35}$$

where k_1 is the quadrupole moment and s_1 is the skew–quadrupole moment. Using this gives on–axis ($x = y = 0$)

$$\nabla g^2 = 2\,(g_x k_1 + g_y s_1,\; g_x s_1 - g_y k_1) \tag{22.36}$$

Note: The above equations must be modified in places in the lattice where there are mode flips (§23.1 since an individual integral must be evaluated using the same physical mode throughout the lattice.

I_0 is not a standard radiation integral. It is useful, though, in calculating the average number of photons emitted. For electrons:

$$\mathcal{N} = \frac{5\,r_e m\,c^2}{2\sqrt{3}\,\hbar\,c}\,I_0 \tag{22.37}$$

where $\mathcal{N}$ is the average number of photons emitted by a particle over one turn, and r_e is the particle's "classical radius" given by Eq. (22.4).

In a dipole a non–zero e_1 or e_2 gives a contribution to I_4 via the $\nabla g^2 \cdot \boldsymbol{\eta}$ term. The edge field is modeled as a thin quadrupole of length δ and strength $k = -g\,\tan(e)/\delta$. It is assumed that $\mathbf{g}$ rises linearly within the edge field from zero on the outside edge of the edge field to its full value on the inside edge of the edge field. Using this in Eq. (22.36) and integrating over the edge field gives the contribution to I_4 from a non–zero e_1 as

$$I_{4z} = -\tan(e_1)\,g^2\,(\cos(\theta)\,\eta_x + \sin(\theta)\,\eta_y) \tag{22.38}$$

With an analogous equation for a finite e_2. The extension to I_{4a} and I_{4b} involves using $\boldsymbol{\eta}_a$ and $\boldsymbol{\eta}_b$ in place of $\boldsymbol{\eta}$. In Eq. (22.38) θ is the reference `tilt` angle which is non–zero if the bend is not in the horizontal plane. Here use of the fact has been made that the $\mathbf{g}$ vector rotates as θ and the quadrupole and skew quadrupole strengths rotate as $2\,\theta$.

The above integrals are invariant under rotation of the (x, y) coordinate system and reduce to the standard equations when $g_y = 0$ as they should.

There are various parameters that can be expressed in terms of these integrals. The I_1 integral can be related to the momentum compaction α_p via

$$I_1 = L\,\frac{dL/L}{dp/p} = L\,\alpha_p \tag{22.39}$$

where p is the momentum and L is the ring circumference. The can be related to the time slip factor η_p by

$$\eta_p = \frac{dt/t}{dp/p} = \alpha_p - \frac{1}{\gamma^2} \tag{22.40}$$

The energy loss per turn is related to I_2 via

$$U_0 = \frac{2\,r_e E_0^4}{3\,(mc^2)^3}\,I_2 \tag{22.41}$$

where E_0 is the nominal energy.

The damping partition numbers are related to the radiation integrals via

$$J_a = 1 - \frac{I_{4a}}{I_2}\,, \quad J_b = 1 - \frac{I_{4b}}{I_2}\,, \text{ and } \quad J_z = 2 + \frac{I_{4z}}{I_2}\,. \tag{22.42}$$

Since

$$\boldsymbol{\eta}_a + \boldsymbol{\eta}_b = \boldsymbol{\eta}\,, \tag{22.43}$$

Robinson's theorem, $J_a + J_b + J_z = 4$, is satisfied. Alternatively, the exponential damping coefficients per turn are

$$\alpha_a = \frac{U_0\, J_a}{2E_0}\,, \quad \alpha_b = \frac{U_0\, J_b}{2E_0}\,, \text{ and } \quad \alpha_z = \frac{U_0\, J_z}{2E_0}\,. \tag{22.44}$$

The energy spread is given by

$$\sigma_{pz}^2 = \left(\frac{\sigma_E}{E_0}\right)^2 = C_q \gamma_0^2 \frac{I_3}{2I_2 + I_{4z}} \tag{22.45}$$

where γ_0 is the usual energy factor and

$$C_q = \frac{55}{32\sqrt{3}}\,\frac{\hbar}{mc} = 3.832 \times 10^{-13}\,\text{meter for electrons} \tag{22.46}$$

If the synchrotron frequency is not too large, the bunch length is given by

$$\sigma_z^2 = \frac{I_1}{M(6,5)}\,\sigma_{pz}^2 \tag{22.47}$$

where $M(6,5)$ is the $(6,5)$ element for the 1–turn transfer matrix of the storage ring. Finally, the emittances are given by

$$\begin{aligned} \epsilon_a &= \frac{C_q}{I_2 - I_{4a}}\,\gamma_0^2\, I_{5a} \\ \epsilon_b &= \frac{C_q}{I_2 - I_{4b}}\,\left(\gamma_0^2\, I_{5b} + \frac{13}{55}\, I_{6b}\right) \end{aligned} \tag{22.48}$$

The I_{6b} term come from the finite vertical opening angle of the radiation[Rauben91]. Normally this term is very small compared to the emittance due to coupling or vertical kicks due to magnet misalignment.

For a non-circular machine, radiation integrals are still of interest if there are bends or steering elements. However, in this case, the appropriate energy factors must be included to take account any changes in energy due to any `lcavity` elements. For a non-circular machine, the I_1 integral is not altered and the I_4 integrals are not relevant. The other integrals become

$$L_2 = \int ds\, g^2\, \gamma_0^4 \tag{22.49}$$

$$L_3 = \int ds\, g^3\, \gamma_0^7 \tag{22.50}$$

$$L_{5a} = \int ds\, g^3\, \mathcal{H}_a\, \gamma_0^6 \tag{22.51}$$

$$L_{5b} = \int ds\, g^3\, \mathcal{H}_b\, \gamma_0^6 \tag{22.52}$$

In terms of these integrals, the energy loss through the lattice is

$$U_0 = \frac{2\, r_c\, mc^2}{3} L_2 \tag{22.53}$$

The energy spread assuming σ_E is zero at the start and neglecting any damping is

$$\sigma_E^2 = \frac{4}{3}\, C_q\, r_c\, \left(mc^2\right)^2\, L_3 \tag{22.54}$$

The above equation is appropriate for a linac. In a storage ring, where there are energy oscillations, the growth of σ_E^2 due to quantum excitation is half that. One way to explain this is that in a storage ring, the longitudinal motion is "shared" between the z and pz coordinates and, to preserve phase space volume, this reduces σ_E^2 by a factor of 2.

Again neglecting any initial beam width, the transverse beam size at the end of the lattice is

$$\begin{aligned} \epsilon_a &= \frac{2}{3} C_q \, r_c \, \frac{L_{5a}}{\gamma_f} \\ \epsilon_b &= \frac{2}{3} C_q \, r_c \, \frac{L_{5b}}{\gamma_f} \end{aligned} \tag{22.55}$$

Where γ_f is the final gamma.

Chapter 23

Linear and Non-Linear Optics

23.1 Coupling and Normal Modes

The coupling formalism used by *Bmad* is taken from the paper of Sagan and Rubin[Sagan99]. The main equations are reproduced here.

The analysis starts with the map $\mathbf{T}(s)$ for the transverse two–dimensional phase space coordinates $\mathbf{x} = (x, x', y, y')$. In ring, with a closed geometry, this map will be a one-turn map starting and ending at some point s. For a machine with open geometry, $\mathbf{T}(0)$ can be computed from the initial Twiss and coupling parameters and $\mathbf{T}(s)$ can then be computed by propagating with the transfer map $\mathbf{M}_{0s}$ from 0 to s:

$$\mathbf{T}(s) = \mathbf{M}_{0s}\, \mathbf{T}(0)\, \mathbf{M}_{0s}^{-1} \tag{23.1}$$

$\mathbf{T}$ can be decomposed using a similarity transformation can be written as

$$\mathbf{T} = \mathbf{V}\, \mathbf{U}\, \mathbf{V}^{-1}, \tag{23.2}$$

where $\mathbf{V}$ is symplectic, and $\mathbf{U}$ is of the form

$$\mathbf{U} = \begin{pmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{B} \end{pmatrix}. \tag{23.3}$$

Since $\mathbf{U}$ is uncoupled, the standard Twiss analysis can be performed with the $\mathbf{A}$ and $\mathbf{B}$ matrices being parameterized using the standard form:

$$\mathbf{A} = \begin{pmatrix} \cos\theta_a + \alpha_a \sin\theta_a & \beta_a \sin\theta_a \\ -\gamma_a \sin\theta_a & \cos\theta_a - \alpha_a \sin\theta_a \end{pmatrix} \tag{23.4}$$

with a similar equation for $\mathbf{B}$.

The $\mathbf{V}$ "coupling" matrix is written in the form:[1]

$$\mathbf{V} = \begin{pmatrix} \gamma\mathbf{I} & \mathbf{C} \\ -\mathbf{C}^{+} & \gamma\mathbf{I} \end{pmatrix}, \tag{23.5}$$

[1] The form of $\mathbf{V}$ and $\mathbf{U}$ is not unique. The form of $\mathbf{V}$ and $\mathbf{U}$ used here essentially follows the form given by Edwards and Teng[Edwards73].

where $\mathbf{C}$ is a 2x2 matrix and + superscript denotes the symplectic conjugate:

$$\mathbf{C}^{+} = \begin{pmatrix} C_{22} & -C_{12} \\ -C_{21} & C_{11} \end{pmatrix}. \tag{23.6}$$

Since we demand that $\mathbf{V}$ be symplectic we have the condition

$$\gamma^2 + |\mathbf{C}| = 1, \tag{23.7}$$

and $\mathbf{V}^{-1}$ is given by

$$\mathbf{V}^{-1} = \begin{pmatrix} \gamma\mathbf{I} & -\mathbf{C} \\ \mathbf{C}^{+} & \gamma\mathbf{I} \end{pmatrix}. \tag{23.8}$$

$\mathbf{C}$ is a measure of the coupling. $\mathbf{T}$ is uncoupled if and only if $\mathbf{C} = \mathbf{0}$.

It is useful to normalize out the $\beta(s)$ variation in the above analysis. Normalized quantities being denoted by a bar above them. The normalized normal mode matrix $\overline{\mathbf{U}}$ is defined by

$$\overline{\mathbf{U}} = \mathbf{G}\,\mathbf{U}\,\mathbf{G}^{-1}, \tag{23.9}$$

Where $\mathbf{G}$ is given by

$$\mathbf{G} \equiv \begin{pmatrix} \mathbf{G}_a & \mathbf{0} \\ \mathbf{0} & \mathbf{G}_b \end{pmatrix}, \tag{23.10}$$

with

$$\mathbf{G}_a = \begin{pmatrix} \frac{1}{\sqrt{\beta_a}} & 0 \\ \frac{\alpha_a}{\sqrt{\beta_a}} & \sqrt{\beta_a} \end{pmatrix}, \tag{23.11}$$

with a similar equation for $\mathbf{G}_b$. With this definition, the corresponding $\overline{\mathbf{A}}$ and $\overline{\mathbf{B}}$ (cf. Eq. (23.3)) are just rotation matrices. The relationship between $\mathbf{T}$ and $\overline{\mathbf{U}}$ is

$$\mathbf{T} = \mathbf{G}^{-1}\,\overline{\mathbf{V}}\,\overline{\mathbf{U}}\,\overline{\mathbf{V}}^{-1}\,\mathbf{G}, \tag{23.12}$$

where

$$\overline{\mathbf{V}} = \mathbf{G}\,\mathbf{V}\,\mathbf{G}^{-1}. \tag{23.13}$$

Using Eq. (23.10), $\overline{\mathbf{V}}$ can be written in the form

$$\overline{\mathbf{V}} = \begin{pmatrix} \gamma\mathbf{I} & \overline{\mathbf{C}} \\ -\overline{\mathbf{C}}^{+} & \gamma\mathbf{I} \end{pmatrix}, \tag{23.14}$$

with the normalized matrix $\overline{\mathbf{C}}$ given by

$$\overline{\mathbf{C}} = \mathbf{G}_a\,\mathbf{C}\,\mathbf{G}_b^{-1}. \tag{23.15}$$

The two normal modes of oscillation are denoted a and b with the a-mode associated with the $\mathbf{A}$ matrix and the b-mode associated with the $\mathbf{B}$ matrix. The normal mode phase space coordinates are denoted $\mathbf{a} = (a, p_a, b, p_b)$. If the one–turn matrix $\mathbf{T}$ is uncoupled then the a-mode is associated with horizontal horizontal motion and b-mode is associated with vertical motion.

The normal mode coordinates $\mathbf{a}$ are related to the laboratory frame via

$$\mathbf{a} = \mathbf{V}^{-1}\,\mathbf{x}. \tag{23.16}$$

In particular the normal mode dispersion $\boldsymbol{\eta}_a = (\eta_a, \eta'_a, \eta_b, \eta'_b)$ is related to the laboratory frame dispersion $\boldsymbol{\eta}_x = (\eta_x, \eta'_x, \eta_y, \eta'_y)$ via

$$\boldsymbol{\eta}_a = \mathbf{V}^{-1}\, \boldsymbol{\eta}_x. \tag{23.17}$$

When there is no coupling ($\mathbf{C} = 0$), $\boldsymbol{\eta}_a$ and $\boldsymbol{\eta}_x$ are equal to each other.

In highly coupled lattices there is the possibility of "`mode flips`". An example will make this clear. Suppose that at one point in a lattice, which will be labeled s_1, the 1-turn matrix $\mathbf{T}_1$ is uncoupled ($\mathbf{V}_1$ is the unit matrix). The two normal modes at this point will be labeled a_1 and b_1. and $\mathbf{T}_1$ can be written in the form

$$\mathbf{T}_1 = \begin{pmatrix} \mathbf{A}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_1 \end{pmatrix} \tag{23.18}$$

Further assume that the transfer matrix $\mathbf{M}_{12}$ between point s_1 and some other point s_2 is of the form

$$\mathbf{M}_{12} = \begin{pmatrix} \mathbf{0} & \mathbf{E} \\ \mathbf{F} & \mathbf{0} \end{pmatrix} \tag{23.19}$$

The 1-turn matrix $\mathbf{T}_2$ at s_2 will be

$$\mathbf{T}_2 = \mathbf{M}_{12}^{-1}\, \mathbf{T}_1\, \mathbf{M}_{12} = \begin{pmatrix} \mathbf{F}^{-1}\, \mathbf{B}_1 \mathbf{F} & \mathbf{0} \\ \mathbf{0} & \mathbf{E}^{-1}\, \mathbf{A}_1\, \mathbf{E} \end{pmatrix} = \begin{pmatrix} \mathbf{A}_2 & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_2 \end{pmatrix} \tag{23.20}$$

This shows that the at s_2 the a_2 normal mode is associated with the b_1 mode and the b_2 mode is associated with the a_1 mode! This is a mode flip. What this means is that in this highly coupled lattice the excitation of a given "physical" mode will be described using the a-mode in some places of the lattice and the b-mode in other places. In particular, it is important to keep track of where there are mode flips when evaluating synchrotron radiation integrals like I_{4a} and I_{4b} (§22.3) since an individual integral must be evaluated using the same physical mode throughout.

At any point where *Bmad* evaluates the Twiss parameters, a `mode_flip` parameter is set. By default, *Bmad* sets the `mode_flip` at the beginning of the lattice to `False` (§10.4) and then calculates the `mode_flip` parameter appropriately for any other point. For a lattice with a closed geometry, if the lattice is stable, the `mode_flip` state at the end of the lattice will be equal to the state at the beginning of the lattice.

23.2 Tunes From One-Turn Matrix Eigen Analysis

Given the 6×6 one-turn matrix for a storage ring, one issue is how to extract the tunes. If there is no coupling the analysis is simple but with coupling things get more complicated. In the general case, calculating with eigenvectors and eigenvalues gives, assuming that the lattice is stable, three pairs of eigenvalues with the two eigenvalues of a given pair being complex conjugates and all eigenvalues having unit amplitude. That is, the eigenvalues λ_i, $i = 1, \ldots 6$ can be ordered in pairs:

$$\begin{aligned} \lambda_1,\, \lambda_2 &= \exp(i\,\theta_a),\, \exp(-i\,\theta_a) \\ \lambda_3,\, \lambda_4 &= \exp(i\,\theta_b),\, \exp(-i\,\theta_b) \\ \lambda_5,\, \lambda_6 &= \exp(i\,\theta_c),\, \exp(-i\,\theta_c) \end{aligned} \tag{23.21}$$

where θ_a, θ_b, and θ_c are the three tunes. To associate λ_1 and λ_2, along with their associated eigenvectors $\mathbf{v}_1$ and $\mathbf{v}_2$, with the "horizontal-like" mode, all the eigenvectors are compared to one another and the eigenvector pair with the largest values for the x and p_x components are used for $\mathbf{v}_1$ and $\mathbf{v}_2$. Similarly, for the "vertical-like" mode, eigenvector pair with the largest values for the y and p_y components are

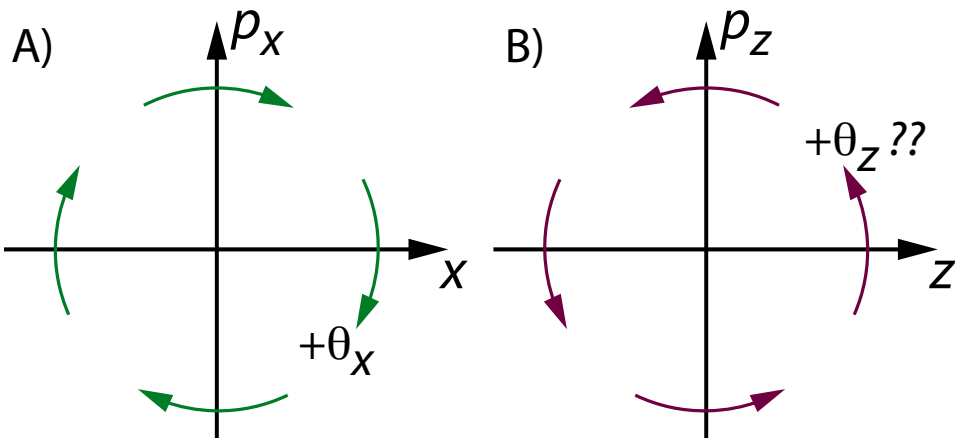

Figure 23.1: A) The standard accelerator physics convention is that a clockwise rotation in (x, p_x) or (y, p_y) space represents a positive tune. B) For longitudinal oscillations, it is sometimes conventional to take counterclockwise rotation as positive if a machine is always running above transition.

associated with $\mathbf{v}_3$ and $\mathbf{v}_4$, and finally for the "longitudinal-like" mode the eigenvector pair with the largest values for the z and p_z components are associated with $\mathbf{v}_5$ and $\mathbf{v}_6$.

It can be useful to arrange the eigenvalues such that the odd numbered eigenvalues (1, 3, and 5) are associated with the tune and the even numbered eigen values (2, 4, and 6) are associated with the negative of the tune as arranged in Eq. (23.21). The algorithm for doing this can be deduced by first considering the case where the motion is in one-dimension only. Here taken to be (x, p) as shown in Fig. 23.1. Notice that, by the standard accelerator physics convention, a positive tune represents a clockwise rotation in the transverse dimensions. For the longitudinal mode what counts as positive tune can depend upon whether the machine is above transition or not. To keep the mathematics consistent, positive tune for all modes will be taken to be clockwise.

Assuming that the motion is circular, the one-turn matrix $\mathbf{M}$ with tune θ is

$$\mathbf{M} = \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix} \tag{23.22}$$

The eignvalues and eigenvectors are

$$\begin{aligned} \lambda_1 &= \exp(i\,\theta), \qquad & \mathbf{v}_1 &= \frac{1}{\sqrt{2}}\,(1, i) \\ \lambda_2 &= \exp(-i\,\theta), \qquad & \mathbf{v}_2 &= \frac{1}{\sqrt{2}}\,(1, -i) \end{aligned} \tag{23.23}$$

Thus, for the eigenvector $(1, i)/\sqrt{2}$, were the momentum component is rotated by a factor of $\pi/2$ counterclockwise from the position coordinate, the rotation angle is the tune. The rotation angle associated with the eigenvector $(1, -i)/\sqrt{2}$ is associated with the negative of the tune.

In the general case, each $\mathbf{v}_k$, $k = 1, \ldots 6$, is a vector in (x, p_x, y, p_y, z, p_z) space with each component of the vector being a complex number. The criterion that an eigenvector is associated with the tune is that the phase of the momentum components are on average rotated clockwise from the position coordinates is

$$\begin{aligned} \widetilde{\mathbf{v}}_k^*\,\mathbf{S}\,\mathbf{v}_k &= i, \qquad & k &= 1, 3, 5 \\ \widetilde{\mathbf{v}}_k^*\,\mathbf{S}\,\mathbf{v}_k &= -i, \qquad & k &= 2, 4, 6 \end{aligned} \tag{23.24}$$

where the tilde means transpose and $\mathbf{S}$ is the matrix

$$\mathbf{S} = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & -1 & 0 \end{pmatrix} \tag{23.25}$$

Eq. (23.24), along with the condition

$$v_{k+1} = v_k^*, \qquad k = 1, 3, 5 \tag{23.26}$$

makes sure that the $\mathbf{v}_k$ are properly normalized.[2]

23.3 Linear Action-Angle Coordinates

The transformation from the one-turn 6×6 matrix $\mathbf{T}$ in laboratory coordinates to action-angle coordinates uses the similarity transformation

$$\mathbf{T} = \mathbf{N}\,\mathbf{R}\,\mathbf{N}^{-1} \tag{23.27}$$

where $\mathbf{N}$ is a symplectic matrix. $\mathbf{N}^T\,\mathbf{S}\,\mathbf{N} = \mathbf{S}$ with S given in Eq. (23.25) and $\mathbf{R}$ is a rotation matrix

$$\mathbf{R} = \begin{pmatrix} \mathbf{R}_2(\theta_a) & & \\ & \mathbf{R}_2(\theta_b) & \\ & & \mathbf{R}_2(\theta_c) \end{pmatrix} \tag{23.28}$$

with each 2×2 rotation submatrix $\mathbf{R}_2$ being of the form as $\mathbf{M}$ in Eq. (23.22). The transformation from laboratory coordinates $\mathbf{x}$ to normal mode $\mathbf{a}$ coordinates is

$$\mathbf{a} = \mathbf{N}^{-1}\,\mathbf{x} \tag{23.29}$$

In action-angle coordinates $\mathbf{a}$ looks like

$$\mathbf{a} = \Big(\sqrt{2\,J_a}\,\cos(\phi_a), -\sqrt{2\,J_a}\,\sin(\phi_a), \sqrt{2\,J_b}\,\cos(\phi_b), -\sqrt{2\,J_b}\,\sin(\phi_b), \sqrt{2\,J_c}\,\cos(\phi_c), -\sqrt{2\,J_c}\,\sin(\phi_c)\Big) \tag{23.30}$$

where J_a, J_b, and J_c are the actions and ϕ_a, ϕ_b, and ϕ_c are the angles.

When the motion is uncoupled, the action and angle of a mode is related to the laboratory coordinates, up to an overall phase factor via:

$$\begin{aligned} x &= \sqrt{2\,J\,\beta}\,\cos(\phi) \\ p &= -\sqrt{\frac{2\,J}{\beta}}\,\left(\alpha\,\cos(\phi) + \sin(\phi)\right) \end{aligned} \tag{23.31}$$

[2]Different Authors use different conventions. For example, the $\mathbf{S}$ matrix in the paper by Chao[Chao81], is the negative of the $\mathbf{S}$ matrix defined here and in the paper by Ohmi, Hirata, and Oide [Ohmi94] the phases are reversed (positive phase is counterclockwise rotation) as can be seen from Eqs. (77) and (79) in their paper.

With the eigenvectors normalized with Eq. (23.24), the particle position $\mathbf{x}$ on turn m when a single mode k is excited can be written in the form

$$\mathbf{x}(m) = \sqrt{J}\,\mathbf{v}_k\, e^{i\,(\theta_k\, m+\phi_0)} + \text{C.C.} \tag{23.32}$$

where the phase ϕ_0 is set by the initial particle position $\mathbf{x}(0)$ and `C.C.` means complex conjugate.

The $\mathbf{N}$ matrix can be constructed using the eigenvectors of $\mathbf{T}$ (§23.2)

$$\mathbf{N} = \frac{1}{\sqrt{2}}\,((\widetilde{\mathbf{v}}_1 + \widetilde{\mathbf{v}}_2), -i\,(\widetilde{\mathbf{v}}_1 - \widetilde{\mathbf{v}}_2), (\widetilde{\mathbf{v}}_3 + \widetilde{\mathbf{v}}_4), -i\,(\widetilde{\mathbf{v}}_3 - \widetilde{\mathbf{v}}_4), (\widetilde{\mathbf{v}}_5 + \widetilde{\mathbf{v}}_6), -i\,(\widetilde{\mathbf{v}}_5 - \widetilde{\mathbf{v}}_6)) \tag{23.33}$$

where the tilde means transpose. That is, the $\widetilde{\mathbf{v}}_k$ are column vectors.

The $\mathbf{v}_k$ vectors, $k = 1, 3, 5$ can each be multiplied by an arbitrary complex phase factor z with unit magnitude. The corresponding $\mathbf{v}_k$ must then be multiplied by z^* to keep Eq. (23.26) satisfied. To recover the standard Twiss parameters, without coupling $\mathbf{N}$ should have the form

$$\mathbf{N} = \begin{pmatrix} \mathbf{N}_a & & \\ & \mathbf{N}_b & \\ & & \mathbf{N}_c \end{pmatrix} \tag{23.34}$$

where the 2×2 submatrices have the standard form

$$\mathbf{N}_a = \begin{pmatrix} \sqrt{\beta_a} & 0 \\ \dfrac{-\alpha_a}{\sqrt{\beta_a}} & \dfrac{1}{\sqrt{\beta_a}} \end{pmatrix} \tag{23.35}$$

with similar equations for $\mathbf{N}_b$ and $\mathbf{N}_c$. To make the $(1,2)$ component of these submatrices zero, along with having $\sqrt{\beta}$ positive, the k^{th} component of $\mathbf{v}_k$, $k = 1, 3, 5$ must be positive real. This fixes the overall phase of the eigenvectors.

23.4 Dispersion Calculation

The dispersion η is defined in the standard way

$$\text{eta_x} = \eta_x(s) \equiv \left.\frac{dx}{dp_z}\right|_s, \qquad \text{eta_y} = \eta_y(s) \equiv \left.\frac{dy}{dp_z}\right|_s \tag{23.36}$$

The associated momentum dispersion is:

$$\text{etap_x} = \eta_{px} \equiv \left.\frac{dp_x}{dp_z}\right|_s, \qquad \text{etap_y} = \eta_{py} \equiv \left.\frac{dp_y}{dp_z}\right|_s, \tag{23.37}$$

The momentum dispersion is useful when constructing particle bunch distributions and for various calculations like for calculating radiation integrals.

To calulate the normal mode dispersions, Eq. (23.16) is used to transform from laboratory to normal mode coordinates.

The one drawback with the momentum dispersion is that it is not always simply related to the derivative of the dispersion $d\eta/ds$. This becomes a factor when designing lattices where, if some section of the lattice needs to be dispersion free, it is convienient to be able to optimize $d\eta/ds$ to zero. The dispersion

derivative is related to the momentum dispersion by

$$\begin{aligned}
\text{deta_x_ds} &\equiv \frac{d\eta_x}{ds} = \frac{d}{dp_z}\left(\frac{dx}{ds}\right) = \frac{d}{dp_z}\left(\frac{p_x}{1+p_z}\right) = \frac{1}{1+p_z}\,\eta_{px} - \frac{p_x}{(1+p_z)^2} \\
\text{deta_y_ds} &\equiv \frac{d\eta_y}{ds} = \frac{d}{dp_z}\left(\frac{dy}{ds}\right) = \frac{d}{dp_z}\left(\frac{p_y}{1+p_z}\right) = \frac{1}{1+p_z}\,\eta_{py} - \frac{p_y}{(1+p_z)^2}
\end{aligned} \tag{23.38}$$

For a lattice branch with an open (non-circular) geometry, the dispersion of the z phase space coordinate, η_z can be defined similar to the dispersion of the other coordinates. In this case, the dispersion vector $\boldsymbol{\eta}$ is defined by

$$\boldsymbol{\eta} = (\eta_x, \eta_{px}, \eta_y, \eta_{py}, \eta_z, 1) \tag{23.39}$$

and this vector is propagated via

$$\boldsymbol{\eta}(s_2) = \mathbf{M}_{21}\,\boldsymbol{\eta}(s_1) \tag{23.40}$$

where $\mathbf{M}_{21}$ is the transfer matrix between points s_1 and s_2.

For an open geometry lattice branch, there are two ways one can imagine defining the dispersion: Either with respect to changes in energy at the beginning of the machine or with respect to the local change in energy at the point of measurement. The former definition will be called "non-local dispersion" and the latter definition will be called "local dispersion" which what *Bmad* calculates. The non-local dispersion $\widetilde{\boldsymbol{\eta}}(s_1)$ at some point s_1 is related to the local dispersion $\boldsymbol{\eta}(s_1)$ via

$$\widetilde{\boldsymbol{\eta}}(s_1) = \frac{dp_{z1}}{dp_{z0}}\,\boldsymbol{\eta}(s_1) \tag{23.41}$$

where s_0 is the s-position at the beginning of the machine. The non-local dispersion has the merit of reflecting what one would measure if the starting energy of the beam is varied. The local dispersion, on the other hand, reflects the correlations between the particle energy and particle position within a beam.

For a closed geometry lattice branch, defining the dependence of z on p_z is problematical. With the RF off, z is not periodic so a closed orbit z cannot be defined. With the RF on, the dispersion of any of the phase space components is not well defined. This being the case, η_z is just treated as zero for a closed branch.

Note: For a closed geometry branch with RF on, it is possible to define dispersions. If $\mathbf{v}$ is the eigenvector of the eigenmode associated with longitudinal oscillations, the dispersion η_x can be defined by $\mathbf{v}(1)/\mathbf{v}(6)$ with similar definitions for the other dispersion components. With this definition, the dispersion become complex. In the low RF limit, the dispersions η_x, η_{px}, η_y, η_{py} converge to the standard (real) values and η_z diverges to infinity.[3]

23.5 Twiss Calculation

23.5.1 Twiss Calculation of Off-Energy Particles

The Twiss parameter calculation for off-energy particles is done in one of two ways depending upon the setting of the `bmad_com` parameter `normalize_twiss` (§11.2). Both calculations start with the transfer matrix which is the Jacobian matrix evaluated with respect to the map reference orbit which is the closed orbit for lattice branches with a closed geometry and, for lattice branches with an open geometry, the

[3] This is assuming a linear system. In practice, the motion will become unstable due to the finite size of the RF bucket.

map reference orbit is the orbit computed starting at the user defined initial position at the beginning of the branch (§16.1.1).

If `normalize_twiss` is set to False (the default), the Twiss calculation uses the transfer matrix without modification. If `normalize_twiss` is set to True, before the Twiss calculation is performed, the transfer matrix is "energy normalized" by converting the matrix from *Bmad* standard phase space (x, p_x, y, p_y, z, p_z) coordinates to "angle" coordinates (x, x', y, y', z, p_z) (§17.8) using the paraxial approximation:

$$x', y' = \frac{p_x}{1+p_z}, \frac{p_y}{1+p_z} \tag{23.42}$$

The advantage of normalizing the matrix is that for "linear" elements like a drift, the calculated Twiss parameters are independent of the reference p_z about which the matrix is calculated ([Carli22]). The disadvantage is that in an element where the map reference energy is changing, like it is in an `e_gun` or `lcavity` element, the Twiss parameters at the end of the element will be different if the whole element transfer matrix is used or the element is sliced and the Twiss parameters are propagated slice by slice.

23.5.2 Twiss Calculation With Time Varying Fields

It should always be keept in mind that the 4D Twiss calculation in *Bmad* essentially ignores coupling between transverse and longitudinal motion. This is a common practice that can be found in other simulation programs like MAD and Elegant. The full 6D treatment ([Wolski06]) is considerably more complicated than any 4D analysis. For example, Wolski's algorithm characterizes each normal mode with a 6x6 matrix.

The bottom line is that Twiss parameters calculated when there is transverse-longitudinal coupling must be taken with a grain of salt.

23.6 Montague W-function

The Montague W-funciton[Montague79] is a way to characterize nonlinearities. The W-function is

$$W = \sqrt{A^2 + B^2} \tag{23.43}$$

where

$$A = \frac{\partial \alpha}{\partial p_z} - \frac{\alpha}{\beta}\frac{\partial \beta}{\partial p_z}, \qquad B = \frac{1}{\beta}\frac{\partial \beta}{\partial p_z} \tag{23.44}$$

The W-function can be calculated even if the geometry of the lattice branch has an open geometry. In this case, dbeta/dpz and dalpha/dpz at the beginning of the branch can be set in the lattice file by setting `beginning[dbeta_dpz_a]`, `beginning[dbeta_dpz_b]`, `beginning[dalpha_dpz_a]` and `beginning[dalpha_dpz_b]` (§10.4).

While Montague does not explicitly state it, the W-function is formulated using the energy normalized Twiss calculation (§23.5.1) and using the non-normalized calculation will lead to misleading results. For example, the W-function will not be invariant within a drift with no normalization.

Chapter 24

Spin Dynamics

24.1 Equations of Motion

The propagation of the classical spin vector $\mathbf{S}$ is described in the local reference frame (§16.1.1) by a modified Thomas-Bargmann-Michel-Telegdi (T-BMT) equation[Hoff06]

$$\frac{\mathrm{d}}{\mathrm{d}s}\mathbf{S} = \left\{\frac{(1+\mathbf{r}_t\cdot\mathbf{g})}{c\,\beta_z}\,(\boldsymbol{\Omega}_{BMT}+\boldsymbol{\Omega}_{EDM}) - \mathbf{g}\times\widehat{\mathbf{z}}\right\}\times\mathbf{S} \tag{24.1}$$

where $\mathbf{g}$ is the bend curvature function which points away from the center of curvature of the particle's reference orbit (see Fig. 16.2), $\mathbf{r}_t = (x, y)$ are the transverse coordinates, $c\,\beta_z$ is the longitudinal component of the velocity, and $\widehat{\mathbf{z}}$ is the unit vector in the z-direction. $\boldsymbol{\Omega}_{BMT}$ is the usual T-BMT precession vector due to the particle's magnetic moment

$$\begin{aligned}\boldsymbol{\Omega}_{BMT}(\mathbf{r},\mathbf{P},t) &= -\frac{q}{m\,c}\left[\left(\frac{1}{\gamma}+a\right)c\,\mathbf{B} - \frac{a\,\gamma\,c}{1+\gamma}\,(\boldsymbol{\beta}\cdot\mathbf{B})\,\boldsymbol{\beta} - \left(a+\frac{1}{1+\gamma}\right)\boldsymbol{\beta}\times\mathbf{E}\right] \\ &= -\frac{q}{m\,c}\left[\left(\frac{1}{\gamma}+a\right)c\,\mathbf{B}_\perp + \frac{(1+a)\,c}{\gamma}\,\mathbf{B}_\parallel - \left(a+\frac{1}{1+\gamma}\right)\boldsymbol{\beta}\times\mathbf{E}\right]\end{aligned} \tag{24.2}$$

and $\boldsymbol{\Omega}_{EDM}$ is the precession vector due to a finite Electric Dipole Moment (EDM) [Silenko08][1]

$$\boldsymbol{\Omega}_{EDM}(\mathbf{r},\mathbf{P},t) = -\frac{q\,\eta}{2\,m\,c}\left[\mathbf{E} - \frac{\gamma}{1+\gamma}\,(\boldsymbol{\beta}\cdot\mathbf{E})\,\boldsymbol{\beta} + c\,\boldsymbol{\beta}\times\mathbf{B}\right] \tag{24.3}$$

Here $\mathbf{E}(\mathbf{r},t)$ and $\mathbf{B}(\mathbf{r},t)$ are the electric and magnetic fields, $\mathbf{B}_\perp$ and $\mathbf{B}_\parallel$ are the components perpendicular and parallel to the particle's momentum, γ is the particle's relativistic gamma factor, q, and m are the particle's charge and mass, $\boldsymbol{\beta}$ is the normalized velocity, $a = (g-2)/2$ is the particle's anomalous magnetic moment (values given in Table 3.2), and η is the normalized electric dipole moment which is related to the dipole moment $\mathbf{d}$ via

$$\mathbf{d} = \frac{\eta}{2}\,\frac{q}{m\,c}\,\mathbf{S} \tag{24.4}$$

Note: Some authors define η without the factor of c is the denominator.

It is important to keep in mind that the a and g-factors used here are defined using Eq. (3.2) which, in the case of nuclei and other composite baryonic particles, differs from the conventional definition Eq. (3.4). See the discussion after Eq. (3.2).

[1] Note: The value for the EDM is set by `bmad_com[electric_dipole_moment]` (§11.2).

24.2 Quaternions and Biquaternions

Bmad uses a quaternion representation for spin rotations[Quat] and for the single resonance analysis (§24.7) it is convenient to use biquaternions which are quaternions with complex coefficients. The following is a brief introduction to quaternions and biquaternions. For more information, the reader is referred to the literature[Sangwine].

WARNING! It is important when looking at software or documentation to keep in mind that there are multiple conventions that are used to define quaternions! The "`Hamilton`" convention used here is widely used in mathematics, physics and other fields. The "`JPL`" (Jet Propulsion Laboratory) and "`STS`" ("space shuttle") conventions are generally confined to use in areospace and robotics contexts[Yazell09, Sommer18].

A quaternion $\mathbf{q}$ is a 4-component object:

$$\mathbf{q} = q_0 + q_x\,\mathbf{i} + q_y\,\mathbf{j} + q_z\,\mathbf{k} \tag{24.5}$$

where q_0, q_x, q_y, and q_z are real numbers for quaternions and complex numbers for biquaternions. $\mathbf{i}$, $\mathbf{j}$, and $\mathbf{k}$ are the *fundamental quaternion units* with the properties under multiplication

$$\mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = \mathbf{i}\,\mathbf{j}\,\mathbf{k} = -1, \quad \mathbf{i}\,\mathbf{j} = \mathbf{k}, \quad \mathbf{j}\,\mathbf{i} = -\mathbf{k}, \quad \text{etc.} \tag{24.6}$$

$\mathbf{i}$, $\mathbf{j}$, and $\mathbf{k}$ do not commute among themselves but do commute with real or complex numbers. It is important to keep in mind the difference between the quaternion unit $\mathbf{i}$ and the imaginary number i which is not the same.

Explicitly, the product of two bi/quaternions (that is, two quaternions or biquaternions) is

$$\begin{aligned}\mathbf{b}\,\mathbf{a} =&(a_0b_0 - a_xb_x - a_yb_y - a_zb_z) + (a_0b_x + a_xb_0 + a_yb_z - a_zb_y)\,\mathbf{i} + \\ &(a_0b_y - a_xb_z + a_yb_0 + a_zb_x)\,\mathbf{j} + (a_0b_z + a_xb_y - a_yb_x + a_zb_0)\,\mathbf{k}\end{aligned} \tag{24.7}$$

The q_0 component of a quaternion $\mathbf{q}$ is called the `scalar` part and will be denoted by $S(\mathbf{q})$. the $q_x\mathbf{i} + q_y\mathbf{j} + q_z\mathbf{k}$ part is called the `vector` part and will be denoted $\mathbf{V}(\mathbf{q})$. That is:

$$\mathbf{q} = S(\mathbf{q}) + \mathbf{V}(\mathbf{q}) \tag{24.8}$$

A bi/quaternion with $S(\mathbf{q})$ equal to zero is called a "`pure`" quaternion. the quaternion product in terms of S and $\mathbf{V}$ is

$$\mathbf{p}\,\mathbf{q} = S(\mathbf{p})S(\mathbf{q}) + S(\mathbf{p})\mathbf{V}(\mathbf{q}) + S(\mathbf{q})\mathbf{V}(\mathbf{p}) + \mathbf{V}(\mathbf{p})\mathbf{V}(\mathbf{q}) \tag{24.9}$$

and, if $\mathbf{V}(\mathbf{p})$ and $\mathbf{V}(\mathbf{q})$ are considered as vectors, the last term $\mathbf{a} \equiv \mathbf{V}(\mathbf{p})\mathbf{V}(\mathbf{q})$ in the above equation can be written as

$$S(\mathbf{a}) = -\mathbf{V}(\mathbf{p}) \cdot \mathbf{V}(\mathbf{q}), \qquad \mathbf{V}(\mathbf{a}) = \mathbf{V}(\mathbf{p}) \times \mathbf{V}(\mathbf{q}) \tag{24.10}$$

In the literature, the scalar part is sometimes called the `real` part and the vector part is sometimes called the `imaginary` part. In this manual, this nomenclature is avoided to avoid confusion with the real and imaginary parts of complex numbers. Here $\mathrm{Re}(\mathbf{q})$ and $\mathrm{Im}(\mathbf{q})$ are defined to be the real and imaginary parts of the biquaternion.

$$\begin{aligned}\mathrm{Re}(\mathbf{q}) &\equiv \mathrm{Re}(q_0) + \mathrm{Re}(q_x)\,\mathbf{i} + \mathrm{Re}(q_y)\,\mathbf{j} + \mathrm{Re}(q_z)\,\mathbf{k} \\ \mathrm{Im}(\mathbf{q}) &\equiv -i\,[\mathrm{Im}(q_0) + \mathrm{Im}(q_x)\,\mathbf{i} + \mathrm{Im}(q_y)\,\mathbf{j} + \mathrm{Im}(q_z)\,\mathbf{k}]\end{aligned} \tag{24.11}$$

The dot product (inner product) of two bi/quaternions is defined to be the standard Euclidean dot product in 4D:

$$\mathbf{a} \cdot \mathbf{b} = a_0\,b_0 + a_x\,b_x + a_y\,b_y + a_z\,b_z \tag{24.12}$$

and the bi/quaternion norm $||\mathbf{q}||$ is given by

$$||\mathbf{q}|| \equiv \sqrt{\mathbf{q}\cdot\mathbf{q}} = \sqrt{q_0^2 + q_x^2 + q_y^2 + q_z^2} \tag{24.13}$$

Notice that for quaternions the norm is positive or zero and real. On the other hand, the norm of a biquaternion will, in general, be complex. For pure bi/quaternions, the `complex norm` is defined by

$$|\mathbf{q}| \equiv \sqrt{|q_x|^2 + |q_y|^2 + |q_z|^2} \tag{24.14}$$

with the convention that quaternion norm uses double bars and complex norm uses a single bar. For pure quaternions the two are equal.

The "`quaternion conjugate`", $\overline{\mathbf{q}}$, valid for both quaternions and biquaternions, is defined by[2]

$$\overline{\mathbf{q}} = q_0 - q_x\,\mathbf{i} - q_y\,\mathbf{j} - q_z\mathbf{k} \tag{24.15}$$

The quaternion conjugate has the properties

$$\overline{\mathbf{q}_1\,\mathbf{q}_2} = \overline{\mathbf{q}}_2\,\overline{\mathbf{q}}_1, \qquad \overline{S(\mathbf{q})} = S(\mathbf{q}), \qquad \overline{\mathbf{V}(\mathbf{q})} = -\mathbf{V}(\mathbf{q}) \tag{24.16}$$

Along with the properties

$$\mathbf{p}\cdot\mathbf{q} = \frac{1}{2}\left(\overline{\mathbf{p}}\mathbf{q} + \overline{\mathbf{q}}\mathbf{p}\right) \tag{24.17}$$

The bi/quaternion inverse $\mathbf{q}^{-1}$ is related to the quaternion conjugate by

$$\mathbf{q}^{-1} = \frac{\overline{\mathbf{q}}}{||\mathbf{q}||^2} \tag{24.18}$$

The biquaternion complex conjugate $\mathbf{q}^*$ is just the complex conjugate of the components

$$\mathbf{q}^* = q_0^* + q_x^*\,\mathbf{i} + q_y^*\,\mathbf{j} + q_z^*\,\mathbf{k} \tag{24.19}$$

When a quaternion represents a rotation, $\mathbf{i}$, $\mathbf{j}$, and $\mathbf{k}$ can be thought of as representing unit vectors along the three Cartesian axes $\mathbf{x}$, $\mathbf{y}$, and $\mathbf{z}$ respectively. A rotation through an angle θ around the unit axis $\mathbf{u} = (u_x, u_y, u_z)$ is represented by the quaternion

$$\mathbf{q} = \cos\frac{\theta}{2} + (u_x\,\mathbf{i} + u_y\,\mathbf{j} + u_z\,\mathbf{k})\sin\frac{\theta}{2} \tag{24.20}$$

A rotation quaternion $\mathbf{q}$ is a `unit` quaternion since its norm is unity $||\mathbf{q}|| = 1$.

Given an ordinary spatial vector $(\mathbf{r}_x, \mathbf{r}_y, \mathbf{r}_z)$, this vector is represented by a pure quaternion $\mathbf{r} = (0, \mathbf{r}_x, \mathbf{r}_y, \mathbf{r}_z)$. The rotation of $\mathbf{r}$ through a rotation represented by quaternion $\mathbf{q}$ to position $\mathbf{r}'$ is given by

$$\mathbf{r}' = \mathbf{q}\,\mathbf{r}\,\overline{\mathbf{q}} \tag{24.21}$$

The rotation matrix $\mathbf{R}$ corresponding to Eq. (24.21) so that $\mathbf{r}' = \mathbf{R}\,\mathbf{r}$ is

$$\mathbf{R} = \begin{pmatrix} q_0^2 + q_x^2 - q_y^2 - q_z^2 & 2\,q_x\;q_y - 2\,q_0\,q_z & 2\,q_x\;q_z + 2\,q_0\,q_y \\ 2\,q_x\;q_y + 2\,q_0\,q_z & q_0^2 - q_x^2 + q_y^2 - q_z^2 & 2\,q_y\;q_z - 2\,q_0\,q_x \\ 2\,q_x\;q_z - 2\,q_0\,q_y & 2\,q_y\;q_z + 2\,q_0\,q_x & q_0^2 - q_x^2 - q_y^2 + q_z^2 \end{pmatrix} \tag{24.22}$$

[2]Different symbols are used in the literature. Here the use of "*" to denote the quaternion conjugate is avoided to avoid confusion with complex conjugate. Notice that even with biquaternions, there is no complex conjugation when computing $\overline{\mathbf{q}}$.

Define functions $I_-(\mathbf{n})$ and $I_+(\mathbf{n})$ which map pure quaternions $\mathbf{n} = (0, n_x, n_y, n_z)$ to biquaternions $\mathbf{n^-}$ and $\mathbf{n^+}$ via

$$\begin{aligned}\mathbf{n^+} \equiv I_+(\mathbf{n}) &\equiv \frac{1}{2} + \frac{i}{2}\big(n_x\,\mathbf{i} + n_y\,\mathbf{j} + n_z\,\mathbf{k}\big)\\ \mathbf{n^-} \equiv I_-(\mathbf{n}) &\equiv \frac{1}{2} - \frac{i}{2}\big(n_x\,\mathbf{i} + n_y\,\mathbf{j} + n_z\,\mathbf{k}\big)\end{aligned} \tag{24.23}$$

It will be assumed throughout that $\mathbf{n}$ has unit magnitude $||\mathbf{n}|| = 1$ and has real components (is a quaternion).

For later use, the following identities are useful

$$\begin{aligned}&\mathbf{n^+n^+} = \mathbf{n^+}, && \mathbf{n^-n^-} = \mathbf{n^-},\\ &\mathbf{n} = i(\mathbf{n^-} - \mathbf{n^+}), && \mathbf{n^+n^-} = \mathbf{n^-n^+} = \mathbf{0}\\ &\overline{\mathbf{n^+}} = \mathbf{n^-}, && \overline{\mathbf{n^-}} = \mathbf{n^+}\end{aligned} \tag{24.24}$$

The relations on the first row mean that $I_+(\mathbf{n})$ and $I_-(\mathbf{n})$ produce "idempotent" biquaternions[Sangwine].

For any biquaternions $\mathbf{q}$, $\mathbf{n^+} = I_+(\mathbf{n})$, and $\mathbf{n^-} = I_-(\mathbf{n})$, the biquaternion $\mathbf{Q} = \mathbf{n^+qn^-}$ is pure. This follows from

$$2\,S(\mathbf{Q}) = \mathbf{Q} + \overline{\mathbf{Q}} = \mathbf{n^+}(S(\mathbf{q}) + \overline{S(\mathbf{q})})\mathbf{n^-} + \mathbf{n^+}(\mathbf{V}(\mathbf{q}) + \overline{\mathbf{V}(\mathbf{q})})\mathbf{n^-} = 2\,S(\mathbf{q})\,\mathbf{n^+n^-} = 0 \tag{24.25}$$

Furthermore, the real and imaginary parts of $\mathbf{Q}$ have the following properties

$$\begin{aligned}&\mathrm{Re}(\mathbf{Q}) \cdot \mathrm{Im}(\mathbf{Q}) = 0\\ &\mathrm{Re}(\mathbf{Q}) \cdot \mathbf{n} = \mathrm{Im}(\mathbf{Q}) \cdot \mathbf{n} = 0\\ &|\,\mathrm{Re}(\mathbf{Q})| = |\,\mathrm{Im}(\mathbf{Q})|\end{aligned} \tag{24.26}$$

These relations are fairly easy to derive. For example, the last one is proved via

$$4\,|\,\mathrm{Re}(\mathbf{Q})|^2 = (\mathbf{Q} + \mathbf{Q}^*) \cdot (\mathbf{Q} + \mathbf{Q}^*) = \mathbf{n^+q}^2\mathbf{n^+} + \mathbf{n^-q}^2\mathbf{n^-} = -(\mathbf{Q} - \mathbf{Q}^*) \cdot (\mathbf{Q} - \mathbf{Q}^*) = 4\,|\,\mathrm{Im}(\mathbf{Q})|^2 \tag{24.27}$$

and since the magnitude of both $\mathrm{Re}(\mathbf{Q})$ and $\mathrm{Im}(\mathbf{Q})$ must be positive (since Re and Im produce real numbers), it follows that $|\,\mathrm{Re}(\mathbf{Q})| = |\,\mathrm{Im}(\mathbf{Q})|$.

24.3 Invariant Spin Field

In a storage ring, the `invariant spin field` $\mathbf{n}(\mathbf{r}, s) = (n_x, n_y, n_z)$ [Hoff06, Duan15], which is a function of phase space position $\mathbf{r} = (x, p_x, y, p_y, z, p_z)$ and longitudinal position s, is the *continuous* function with unit amplitude that satisfies

$$\mathbf{n}(\mathcal{M}_r\mathbf{r}, s) = \mathcal{M}_s(\mathbf{r})\,\mathbf{n}(\mathbf{r}, s) \tag{24.28}$$

where $\mathcal{M}_r$ is the orbital part of the 1-turn transfer map and $\mathcal{M}_s(\mathbf{r})$, derived from T-BMT equation, is the spin part of the map which is a function of $\mathbf{r}$. That is, the invariant spin field (ISF) obeys the T-BMT equation. Thus, a particle whose spin points in the direction of $\mathbf{n}(\mathbf{r}, s)$ at some time t will, in the absence of radiation effects, always have its spin pointing in the direction of $\mathbf{n}(\mathbf{r}, s)$. When there are no resonances, $\mathbf{n}(\mathbf{r}, s)$ is unique up to a flip of sign.

In general, it is not straightforward to calculate $\mathbf{n}$. The exceptional case (besides the cases where there is a resonance) is if the particle is on the closed orbit $\mathbf{r}_0$. In this case, since $\mathcal{M}_r\mathbf{r}_0 = \mathbf{r}_0$, and since $\mathcal{M}_s(\mathbf{r})$ is a rotation matrix, Eq. (24.28) can be solved to give the invariant spin field on the closed orbit denoted

by $\mathbf{n}_0$. Over one turn, a spin on the closed orbit rotates around $\mathbf{n}_0$ by the angle $2\pi\nu_0$ where ν_0 is the closed-orbit spin tune.

Only the fractional part of the closed orbit spin tune ν_0 is generally well defined. That is, it is not in general possible to distinguish between tunes $\nu_0 + m$ where m is an integer.[3] Additionally, if $\mathbf{n}_0$ is a valid closed orbit spin field then so is $-\mathbf{n}_0$ and the spin tune associated with $-\mathbf{n}_0$ is the negative of the spin tune associated with $\mathbf{n}_0$.

Since there are two possible anti-parallel solutions for $\mathbf{n}_0$, *Bmad* allows the user to specify which solution to pick if the parameter

`bmad_com%spin_n0_direction_user_set` (§11.2)

is set `True`. If set `False` (which is the default), the direction is choosen by *Bmad* using the convention that chosen spin tune is in the range $[0, \pi]$. From the 1-turn spin transport quaternian $\mathbf{q}$ (Eqs. (24.5) and (24.20)), the equation to calculate the tune is

$$\theta = 2 \cdot \text{atan2}\left(\left|(q_x, q_y, q_z)\right|, |q_0|\right) \tag{24.29}$$

where atan2 is the standard arc tangent function. With this, $\mathbf{n}_0$ is given by

$$\mathbf{n}_0 = \frac{\text{sgn}(q_0)}{|(q_x, q_y, q_z)|}\,(q_x, q_y, q_z) \tag{24.30}$$

where sgn is the sign function. Note that with Eq. (24.30) of $\mathbf{n}_0$ is independent of whether $\mathbf{q}$ or $-\mathbf{q}$ is used in the equation.

24.4 Polarization Limits and Polarization/Depolarization Rates

Once the invariant spin field (§24.3) has been calculated, various quantities of interest can be computed. For example, given some initial distribution of spins in a beam, the maximum possible time averaged polarization $\langle \mathbf{S} \rangle_{\max}$ is

$$\langle \mathbf{S} \rangle_{\max} = \int d\mathbf{r}\, \rho(\mathbf{r})\, \mathbf{n}(\mathbf{r}) \tag{24.31}$$

where the integral is over the beam phase space space density ρ and the longitudinal s-dependence is implicit. The above equation neglects any single spin polarization or depolarization processes. Notice that what is calculated is a time averaged quantity. Instantaneously, the beam can be fully polarized but the average over many turns, at some given position s, cannot exceed $\langle \mathbf{S} \rangle_{\max}$.

Another quantity that can be computed from knowledge of $\mathbf{n}$ is the equilibrium polarization of a beam. The `Baier-Katkov-Strakhovenko` polarization P_{bks} (generalized from `Sokolov-Ternov` to include non-vertical fields) is calculated by ignoring deviations of the beam from the closed orbit[Barber99]

$$P_{bks} = \pm \frac{8}{5\sqrt{3}} \frac{\oint ds\, g^3\, \widehat{\mathbf{b}} \cdot \mathbf{n}_0}{\oint ds\, g^3 \left(1 - \frac{2}{9}(\mathbf{n}_0 \cdot \widehat{\mathbf{s}})^2\right)} \tag{24.32}$$

[3] This is in contrast to the orbital tunes where the integer part is generally well defined (but there are exceptions when there is strong coupling). The reason for this is that at all points in the ring, the orbital normal mode axes are fairly well defined (see, for example, Eq. (23.16)). This means that the phase angle of an oscillating particle with respect to the axes is fairly well defined and counting full oscillations is unambiguous. With spin, the closed orbit spin oscillations are in the plane transverse to $\mathbf{n}_0$ and here there is no non-arbitrary way to define the transverse plane coordinate axes (there is an exception here if $\mathbf{n}_0$ always is pointing in the same general direction). This makes the integer part of the spin tune ambiguous.

where $g = 1/\rho$ is the bending strength (ρ is the bending radius), $\hat{\mathbf{s}}$ is the unit vector in the direction of motion, and $\hat{\mathbf{b}}$ is defined to be

$$\hat{\mathbf{b}} \equiv \frac{\hat{\mathbf{s}} \times d\hat{\mathbf{s}}/ds}{|d\hat{\mathbf{s}}/ds|} \tag{24.33}$$

Notice that $\hat{\mathbf{b}}$ is the direction of the magnetic field when $\hat{\mathbf{s}}$ is perpendicular to the magnetic field and when there is no electric field. In the above equation, the plus sign is for positrons (polarized parallel to the field) and the minus sign is for electrons. Since the above equation is only valid when the anomalous magnetic moment is small[Jackson76], this formula is not valid for protons and anti-protons.

The corresponding BKS polarization build-up rate τ_{bks}^{-1} is

$$\tau_{bks}^{-1} = \frac{5\sqrt{3}}{8} \frac{r_e \gamma^5 \hbar}{m} \frac{1}{C} \oint ds\, g^3 \left(1 - \frac{2}{9}(\mathbf{n}_0 \cdot \hat{\mathbf{s}})^2\right) \tag{24.34}$$

If the stochastic excitation of the beam is taken into account, the generalized Sokolov-Ternov polarization is called the `Derbenev-Kondratenko` formula

$$P_{dk} = \pm \frac{8}{5\sqrt{3}} \frac{\oint ds \left\langle g^3 \hat{\mathbf{b}} \cdot \left(\mathbf{n} - \frac{\partial \mathbf{n}}{\partial \delta}\right)\right\rangle}{\oint ds \left\langle g^3 \left(1 - \frac{2}{9}(\mathbf{n} \cdot \hat{\mathbf{s}})^2 + \frac{11}{18}\left|\frac{\partial \mathbf{n}}{\partial \delta}\right|^2\right)\right\rangle} \tag{24.35}$$

where <> denotes an average over phase space, and δ is the fractional energy deviation which, for ultra-relativistic particles, is the same as phase space p_z. Since, away from any resonances, $\mathbf{n}$ is very close to $\mathbf{n}_0$, and since generally machines are tuned away from any resonances, the $\mathbf{n} - \partial\mathbf{n}/\partial\delta$ and $\mathbf{n} \cdot \hat{\mathbf{s}}$ terms can be replaced by $\mathbf{n}_0 - \partial\mathbf{n}/\partial\delta$ and $\mathbf{n}_0 \cdot \hat{\mathbf{s}}$ when evaluating P_{dk}.

The $\hat{\mathbf{b}} \cdot \partial\mathbf{n}/\partial\delta$ term in Eq. (24.35) is called the "`kinetic polarization`" term.

The time dependence of the polarization is [Barber99])

$$\mathbf{P}(t) = \mathbf{P}_{dk}\ (1 - \exp(-t/\tau_{dk})) + \mathbf{P}_0\ \exp(-t/\tau_{dk}) \tag{24.36}$$

where $\mathbf{P}_0$ is the initial polarization and the polarization rate τ_{dk}^{-1} is

$$\tau_{dk}^{-1} = \frac{5\sqrt{3}}{8} \frac{r_e \gamma^5 \hbar}{m} \frac{1}{C} \oint ds \left\langle g^3 \left(1 - \frac{2}{9}(\mathbf{n} \cdot \hat{\mathbf{s}})^2 + \frac{11}{18}\left|\frac{\partial \mathbf{n}}{\partial \delta}\right|^2\right)\right\rangle \tag{24.37}$$

τ_{dk}^{-1} can be decomposed into two parts:

$$\tau_{dk}^{-1} = \tau_{pol}^{-1} + \tau_{dep}^{-1} \tag{24.38}$$

τ_{pol}^{-1} is the polarization rate given by the first two terms on the RHS in Eq. (24.37) and the depolarization rate τ_{dep}^{-1} is given by the third term:

$$\begin{aligned} \tau_{pol}^{-1} &= \frac{5\sqrt{3}}{8} \frac{r_e \gamma^5 \hbar}{m} \frac{1}{C} \oint ds \left\langle g^3 \left(1 - \frac{2}{9}(\mathbf{n} \cdot \hat{\mathbf{s}})^2\right)\right\rangle \\ \tau_{dep}^{-1} &= \frac{5\sqrt{3}}{8} \frac{r_e \gamma^5 \hbar}{m} \frac{1}{C} \oint ds \left\langle g^3 \frac{11}{18}\left|\frac{\partial \mathbf{n}}{\partial \delta}\right|^2\right\rangle \end{aligned} \tag{24.39}$$

τ_{pol}^{-1} is generally well approximated by the `Baier-Katkov-Strakhovenko` polarization rate (Eq. (24.34)). The difference being that τ_{bks}^{-1} is evaluated along the closed orbit while τ_{pol}^{-1} involves an average over the transverse beam size.

The calculation of P_{dk} (Eq. (24.35)) and τ_{dep}^{-1} (Eq. (24.39)) involve integrating $\partial\mathbf{n}/\partial\delta$ around the ring. The calculation of $\partial\mathbf{n}/\partial\delta$ at any point in the ring involves a sum of eigenvectors $\mathbf{n}_k$, $k = 1,\ldots,6$ (Eq. (24.53)) with pairs of eigenvectors being associated with the three orbital modes of oscillation. Insight into the decoherence process can had by considering what P_{dk} and τ_{dep}^{-1} would be if only one oscillation mode was being excited. This is done by calculating $\partial\mathbf{n}/\partial\delta$ with Eq. (24.53) by using the two $\mathbf{n}_k$ for one particular mode and taking the other $\mathbf{n}_k$ to be zero. This is then used in Eq. (24.35) and Eq. (24.39).

One problem in evaluating Eq. (24.35) is that accurately evaluating the $\partial\mathbf{n}/\partial\delta$ terms over the transverse bunch distribution is complicated due to lattice nonlinearities. One way to evaluate P_{dk} with nonlinearities included is by tracking a bunch of particles over some number of turns. To help minimize the needed tracking, the spin flip process can be neglected. With this, and starting with 100% polarization, a turn-by-turn plot of the polarization will give the depolarization rate τ_{dep}^{-1}. The integrals of $g^3\,\widehat{\mathbf{b}}\cdot\mathbf{n}$ and $g^3\,(1-2(\mathbf{n}\cdot\widehat{\mathbf{s}})^2/9)$ can be well approximated by the integrals over the closed orbit ignoring the finite beam size. Finally, for most rings, the integral of $g^3\widehat{\mathbf{b}}\cdot\partial\mathbf{n}/\partial\delta$ is generally small compared to the integral of $g^3\,\widehat{\mathbf{b}}\cdot\mathbf{n}$ since, in most of the machine, $\widehat{\mathbf{b}}$ and $\mathbf{n}$ will point in the vertical direction and $\partial\mathbf{n}/\partial\delta$ will be perpendicular to the vertical (see Eq. (24.46)). Putting this all together, the equilibrium polarization can be computed from

$$P_{dk} \approx P_{bks}\frac{\tau_{bks}^{-1}}{\tau_{bks}^{-1}+\tau_{dep}^{-1}} \tag{24.40}$$

where τ_{dep}^{-1} is calculated via particle tacking and the other quantities are calculated by integrals over the closed orbit ignoring beam size effects.

24.5 Bmad Tune and n_0 convention

As mentioned above, at any point in a ring the direction of the closed orbit invariant spin $\mathbf{n}_0$ is ambiguous since, if $\mathbf{n}_0$ is a valid invariant spin direction, then so is $-\mathbf{n}_0$. The same is true of the spin tune ν_0 and if ν_0 is the spin tune associated with $\mathbf{n}_0$, $-\nu_0$ is the spin tune associated with $-\mathbf{n}_0$. This ambiguity complicates various calculations. For example, to do the integral in Eq. (24.32), it is necessary to make sure that the direction of $\mathbf{n}_0(s_1)$, the invariant spin at s_1, must be consistent with the direction of the invariant spin at s_2, $\mathbf{n}_0(s_2)$. That is, the two must be related via

$$\mathbf{n}_0(s_2) = \mathbf{q}_{21}\,\mathbf{n}_0(s_1)\,\overline{\mathbf{q}}_{21} \tag{24.41}$$

where $\mathbf{q}_{21}$ is the closed orbit spin transport quaternion from s_1 to s_2. The tune must also be computed consistent with the choice of invariant spin direction. This is important when calculating sum and difference resonance strengths (Eqs. (24.73) and (24.74)).

To ensure a consistent invariant spin direction, Eq. (24.41) is used when calculating integrals involving $\mathbf{n}_0$. To ensure a consistent tune, *Bmad* uses the following convention that the spin tune will always be in the range $[0,\pi]$, and the direction of $\mathbf{n}_0$ will be chosen to be consistent with this choice in tune (it is left as an exercise for the reader to prove that there will always be exactly one spin tune in the range $[0,\pi]$.

24.6 Linear $\partial\mathbf{n}/\partial\delta$ Calculation

When evaluating the equations in the previous section, in many situations it is sufficient to just use the value of $\partial\mathbf{n}/\partial\delta$ as calculated in the linear regime. In the linear regime, $\partial\mathbf{n}/\partial\delta$ is only dependent upon

the s-position and is independent of the phase space position. The calculation of $\partial\mathbf{n}/\partial\delta$ starts with the linearized transport equations which are characterized by a 6×6 orbital 1-turn transfer matrix $\mathbf{M}$ along with the spin transport which can be written in the form

$$\mathbf{q}_s(\mathbf{r}) = \mathbf{q}_0 + \overrightarrow{\mathbf{q}}\cdot\mathbf{r} \tag{24.42}$$

The quaternion map $\mathbf{q}_s$, evaluated at the orbital phase space point $\mathbf{r}$, has a zeroth order part $\mathbf{q}_0$ (the spin rotation for a particle on the closed orbit) and the first order part $\overrightarrow{\mathbf{q}} = (\mathbf{q}_1,\ldots,\mathbf{q}_6)$ which is a vector of six quaternions and which is evaluated in Eq. (24.42) by taking the dot product with the vector $\mathbf{r}$.

The closed orbit invariant spin $\mathbf{n}_0$ which has unit amplitude satisfies the equation

$$\mathbf{q}_0\,\mathbf{n}_0\,\overline{\mathbf{q}}_0 = \mathbf{n}_0 \tag{24.43}$$

The solution to this equation, normalized to one, is

$$\mathbf{n}_0 = \frac{(q_{0,x}, q_{0,y}, q_{0,z})}{\|(q_{0,x}, q_{0,y}, q_{0,z})\|} \tag{24.44}$$

Another way of writing this is using Eq. (24.20)

$$\mathbf{q}_0 = \cos(\pi\,\nu_0) + (n_{0,x}\,\mathbf{i} + n_{0,y}\,\mathbf{j} + n_{0,z}\,\mathbf{k})\,\sin(\pi\,\nu_0) \tag{24.45}$$

where ν_0 is the spin tune.

Since $\mathbf{q}_s$ is a rotation quaternion, the magnitude of $\mathbf{q}_s(\mathbf{r})$ must remain one. Using Eq. (24.42) in Eq. (24.13), to keep the magnitude equal to one to linear order gives the condition

$$\mathbf{q}_0\cdot\mathbf{q}_j = 0, \qquad j = 1,\ldots,6 \tag{24.46}$$

for all $\mathbf{q}_j$ components of $\overrightarrow{\mathbf{q}}$.

To calculate $\partial\mathbf{n}/\partial\delta$, the first step is to compute the eigenvectors $\mathbf{v}_k$ and eigenvalues λ_k, $k = 1,\ldots,6$ of the 1-turn orbital matrix. The corresponding spin eigenvectors $\mathbf{n}_k$ are computed from the equation

$$\mathbf{q}_s(\mathbf{v}_k)\;(\mathbf{n}_0 + \mathbf{n}_k)\;\overline{\mathbf{q}}_s(\mathbf{v}_k) = \mathbf{n}_0 + \lambda_k\,\mathbf{n}_k \tag{24.47}$$

These eigenvectors are perpendicular to $\mathbf{n}_0$. This can be easily shown by noting that $\mathbf{n}(\mathbf{r})$ in Eq. (24.52) must have unit magnitude to linear order for any arbitrary choice of $A_k(\mathbf{r})$. Using Eqs. (24.42) and (24.43), and keeping only linear terms gives

$$\lambda_k\,\mathbf{n}_k - \mathbf{q}_0\,\mathbf{n}_k\,\overline{\mathbf{q}}_0 = (\overrightarrow{\mathbf{q}}\cdot\mathbf{v}_k)\,\mathbf{n}_0\,\overline{\mathbf{q}}_0 + \mathbf{q}_0\,\mathbf{n}_0\,(\overrightarrow{\overline{\mathbf{q}}}\cdot\mathbf{v}_k) \tag{24.48}$$

This equation is linear in the unknown $\mathbf{n}_k$ and so may be solved using standard linear algebra techniques. One small problem with Eq. (24.48) is that it is degenerate along the $\mathbf{n}_0$ axis in the limit when there is no RF voltage since the eigen mode associated with the longitudinal motion will have an eigenvalue of one. In this case, round-off errors can cause large inaccuracies. To get around this, Eq. (24.48) is projected onto the plane perpendicular to $\mathbf{n}_0$ by first constructing vectors $\mathbf{c}_1$ and $\mathbf{c}_2$ which are orthogonal to $\mathbf{n}_0$ and orthogonal to each other. Eq. (24.48) is projected onto the $(\mathbf{c}_1,\mathbf{c}_2)$ plane to give

$$\bigl(\lambda_k\,\mathbf{n}_k - \mathbf{q}_0\,\mathbf{n}_k\,\overline{\mathbf{q}}_0\bigr)\cdot\mathbf{c}_m = \bigl((\overrightarrow{\mathbf{q}}\cdot\mathbf{v}_k)\,\mathbf{n}_0\,\overline{\mathbf{q}}_0 + \mathbf{q}_0\,\mathbf{n}_0\,(\overrightarrow{\overline{\mathbf{q}}}\cdot\mathbf{v}_k)\bigr)\cdot\mathbf{c}_m, \quad m = 1,2 \tag{24.49}$$

Now using the fact that $\mathbf{n}_k$ is perpendicular to $\mathbf{n}_0$ means that $\mathbf{n}_k$ can be written as a linear combination of $\mathbf{c}_1$ and $\mathbf{c}_2$

$$\mathbf{n}_k = a_1\,\mathbf{c}_1 + a_2\,\mathbf{c}_2 \tag{24.50}$$

Using this in Eq. (24.49) gives linear coupled equations in the unknowns (a_1, a_2) which is easily solved.

Since the eigenvectors $\mathbf{v}_k$ span phase space, for any given phase space position $\mathbf{r}$, there exist a set of coefficients $A_k(\mathbf{r})$, $k = 1, \ldots, 6$, such that

$$\mathbf{r} = \sum_{k=1}^{6} A_k(\mathbf{r})\, \mathbf{v}_k \tag{24.51}$$

Define the function $\mathbf{n}$ by

$$\mathbf{n}(\mathbf{r}) \equiv \mathbf{n}_0 + \sum_{k=1}^{6} A_k(\mathbf{r})\, \mathbf{n}_k \tag{24.52}$$

This function obeys the T-BMT equation and is continuous and thus is the solution (up to a flip in sign and a normalization constant) for the invariant spin field. From this, $\partial\mathbf{n}/\partial\delta$, which is computed taking the derivative at constant x, p_x, y, p_y, and z, is obtained via

$$\frac{\partial \mathbf{n}}{\partial \delta} = \sum_{k=1}^{6} A_k\, \mathbf{n}_k \tag{24.53}$$

with the A_k being computed by inverting the equation

$$(0, 0, 0, 0, 0, 1)^t = \sum_{k=1}^{6} A_k\, \mathbf{v}_k \tag{24.54}$$

where the superscript t means transpose. Notice that for $\partial\mathbf{n}/\partial\delta$, as well as any other partial derivative, the component in the direction of $\mathbf{n}_0$ will be zero since, to first order, the amplitude of $\mathbf{n}$ must be constant (since the equation for $\mathbf{n}$ is only valid to first order, the computed amplitude will have non-zero higher order terms).

A problem arises if the machine that is being simulated does not have any RF cavities or the voltage in the cavities is zero. In this case, there are no synchrotron oscillations which results in degenerate eigenvectors and the eigenvectors will not span all of phase space. The solution here is to reduce the dimensionality of phase space to five by removing the z coordinate. The above equations then can be used with the sums over k ranging from 1 to 5.

It is sometimes informative to compute the contribution of $\partial\mathbf{n}/\partial\delta$ due to just one or two modes of oscillation. That is, to compute $\partial\mathbf{n}/\partial\delta$ with the sum in Eq. (24.53) restricted to be over one or two corresponding eigen states that comprise the oscillation modes of interest. This information can help guide lattice design.

Another way for analyzing where in the lattice contributions to $\partial\mathbf{n}/\partial\delta$ are coming from is to consider the spin transport maps for individual elements. These maps will be of the form given in Eq. (24.42). The contribution to $\partial\mathbf{n}/\partial\delta$ from an element is due to the non-zero terms in $\overrightarrow{\mathbf{q}} = (\mathbf{q}_1, \ldots, \mathbf{q}_6)$. If terms are selectively zeroed, and this significantly changes the polarization, this is a clue for designing a lattice. For example, if setting $\mathbf{q}_1$ and $\mathbf{q}_2$ to zero for a set of elements in a certain region of the machine leads to significantly greater polarizations, this indicates that the polarization is sensitive to horizontal excitation in this region. Similarly, zeroing elements in $\mathbf{q}_3$ and $\mathbf{q}_4$ is associated with vertical excitation and $\mathbf{q}_5$ and $\mathbf{q}_6$ is associated with longitudinal excitation. With the slim formalism (§24.9), the 2×6 matrix $\mathbf{G}$ can be decomposed into three 2×2 sub-matrices:

$$\mathbf{G} = (\mathbf{G}_x, \mathbf{G}_y, \mathbf{G}_z) \tag{24.55}$$

Zeroing $\mathbf{q}_1$ and $\mathbf{q}_2$ is equivalent to zeroing $\mathbf{G}_x$, etc.

24.7 Single Resonance Model

Resonances occur when the spin tune ν_s is an integer ("imperfection" resonances) and when the spin tune ν_s in combination with the three orbital tunes ν_x, ν_y, and ν_z is an integer ("intrinsic" or "spin-orbit" resonances)

$$\nu_s + m_x\,\nu_x + m_y\,\nu_y + m_z\,\nu_z = m_0 \tag{24.56}$$

where m_x, m_y, m_z, and m_0 are integers. As discussed below, in the linear approximation (§24.6), resonances only occur if one and exactly one of the m_x, m_y, or m_z has a value of one and the other two are zero.

Generally, in rings where the synchrotron radiation is large (think electrons), the depolarization due to radiation will tend to dominate (§24.4). For rings where the synchrotron radiation is small (think protons), resonances will be more important. Notice that synchrotron radiation, being a stochastic process results in depolarization. Resonances, on the other hand, are not stochastic and even if a resonance tilts the polarization direction there can be the possibility of recovery.

To calculate the effect of a resonance it is helpful to know the resonance strength. The resonance strength calculation can be motivated by considering the Single Resonance Model (SRM)[Hoff06] where it is assumed that only one orbital mode is excited and that there is a single dominating resonance. In this case the spin equation of motion is

$$\frac{d\mathbf{s}}{d\theta} = \mathbf{\Omega} \times \mathbf{s}, \qquad \mathbf{\Omega} = \begin{pmatrix} \epsilon_r \cos\Phi \\ \epsilon_r \sin\Phi \\ \nu_0 \end{pmatrix} \tag{24.57}$$

where ϵ_r is the resonance strength, ν_0 is the closed orbit spin tune and θ is the longitudinal angle with $\theta = 2\,\pi$ representing one turn in the circular accelerator. In the above equation the phase Φ is related to the orbital mode tunes Q via

$$\Phi = (j_0 + j_x Q_x + j_y Q_y + j_z Q_z)\theta + \Phi_0 \equiv \kappa\theta + \Phi_0 \tag{24.58}$$

with j_0, j_x, j_y, and j_z being integers and Φ_0 being the starting phase. Eq. (24.57) defines κ. Eq. (24.57) can be solved by transforming to a coordinate system rotating about the z-axis with a rotational frequency κ. This is called the "`resonance`" coordinate system or the "`rotating`" coordinate system. In this frame, the spin vector s_R is

$$\mathbf{s}_R = \mathbf{q}_z(-\Phi/2)\ \mathbf{s}\ \overline{\mathbf{q}}_z(-\Phi/2) \tag{24.59}$$

where

$$\mathbf{q}_z(\phi) \equiv \big(\cos(\phi/2) + \sin(\phi/2)\big)\,\mathbf{k} \tag{24.60}$$

Using Eq. (24.59) in Eq. (24.57) gives

$$\frac{d\mathbf{s}_R}{d\theta} = \mathbf{\Omega}_R \times \mathbf{s}_R, \qquad \mathbf{\Omega}_R = \begin{pmatrix} \epsilon_r \\ 0 \\ \delta \end{pmatrix}, \qquad \delta = \nu_0 - \kappa \tag{24.61}$$

Spins rotate around around $\mathbf{\Omega}_R$. A spin initially aligned along the z-axis will be tilted a maximum of $\tan^{-1}(2\epsilon_r/\delta)$ away from the z-axis. This shows that the characteristic frequency width of the resonance is set by $\delta \approx \epsilon_r$. That is, the frequency width scales linearly with ϵ_r.

In the rotating coordinate system, a spin oriented parallel to $\mathbf{\Omega}_R$ remains parallel to $\mathbf{\Omega}_R$ so $\mathbf{\Omega}_R$ is in the direction of the invariant spin field. Transforming back to the laboratory frame and normalizing to one, $\mathbf{n}$ is

$$\mathbf{n}(\Phi) = \frac{\mathrm{sgn}(\delta)}{\Lambda} \begin{pmatrix} \epsilon_r \cos\Phi \\ \epsilon_r \sin\Phi \\ \delta \end{pmatrix}, \qquad \Lambda = \sqrt{\delta^2 + \epsilon_r^2} \tag{24.62}$$

where the sign factor $\text{sgn}(\delta)$ is chosen so that on the closed orbit with $\epsilon_r = 0$ the $\mathbf{n}$ axis is in the positive z-direction.

Eq. (24.62) shows that resonance occurs when $\delta = 0$ or $\nu_0 = \kappa$. The use of a sinusoidal perturbation in the single resonance model results in the suppression of all other resonances. In an actual machine, the presence of a perturbation at some specific point in the machine will produce a comb of resonances of the form $m_0 \pm Q$.

At resonance, with $\delta = 0$, Eq. (24.61) shows that if a spin in the resonance coordinate system is initially aligned along the $\mathbf{n}_0$ axis ($\mathbf{s}_R = (0, 0, 1)$), After N turns, in the linear approximation, the spin will be

$$\mathbf{s}_R(N, \mathbf{n}_0) = (0, -N \cdot 2\,\pi\,\epsilon_r, 1) \tag{24.63}$$

Eq. (24.63) gives a physical interpretation to the resonance strength ϵ_r. The resonance strength is the angle (modulo $2\,\pi$) that a spin is tipped away from $\mathbf{n}_0$ in one turn. The linear approximation will be valid for N such that $N\,\epsilon_r \ll 1$. For the linear approximation to be valid for any positive N, ϵ_r must satisfy $\epsilon_r \ll 1$.

24.8 Linear Resonance Analysis

Within the linear approximation, the resonance strength ϵ_r (Eq. (24.57)) can be related to the first order spin transport (Eq. (24.42)). The analysis starts by inversion of Eq. (24.63) and letting N go to infinity

$$\epsilon_r = \lim_{N\to\infty} \frac{1}{2\,\pi\,N} \left|\mathbf{s}_R(N, \mathbf{n}_0) - \mathbf{n}_0\right| \tag{24.64}$$

Here the restriction $N \ll 1/\epsilon_r$ can be ignored since the analysis below will use the linearized spin transport and it is the nonlinear terms in the spin transport which makes the restriction necessary. In fact, with linear spin transport, the magnitude of $\mathbf{s}_R$ is not constant and at resonance is unbounded. Eq. (24.64) will still be valid even when there are multiple resonances present since Eq. (24.64) must be evaluated on a particular resonance and, it will be seen, the contribution to Eq. (24.64) from non-resonant resonances is zero.

$\mathbf{s}_R(N, \mathbf{n}_0)$ is the spin after N turns given an initial spin of $\mathbf{n}_0$. $\mathbf{s}_R(N, \mathbf{n}_0)$ can be computed via

$$\mathbf{s}_R(N, \mathbf{n}_0) = \mathbf{Q}_N\,\mathbf{n}_0\,\overline{\mathbf{Q}}_N \tag{24.65}$$

The rotation quaternion $\mathbf{Q}_N$ over N turns is the product of one-turn rotation quaternions $\mathbf{q}_s(\mathbf{r})$ with an extra factor of $\mathbf{q}_0^{-N}$ to convert from laboratory coordinates to resonance coordinates since $\mathbf{q}_s$ is the transport in laboratory coordinates (by definition, resonance coordinates are the same as laboratory coordinates at $N = 0$).

$$\mathbf{Q}_N = \mathbf{q}_0^{-N}\,\mathbf{q}_s(\mathbf{r}_{N-1}) \ldots \mathbf{q}_s(\mathbf{r}_1)\,\mathbf{q}_s(\mathbf{r}_0) \tag{24.66}$$

In the above equation, $\mathbf{r}_j$ is the orbital position after j turns. Using Eq. (24.42) and expanding to linear order in $\overrightarrow{\mathbf{q}}$ gives

$$\mathbf{Q}_N = \mathbf{q}_0^{-1} \left(\mathbf{q}_0 + \sum_{j=0}^{N-1} \mathbf{q}_0^{-j} \left(\overrightarrow{\mathbf{q}} \cdot \mathbf{r}_j\right) \mathbf{q}_0^{j} \right) \tag{24.67}$$

Using this with Eq. (24.65) in Eq. (24.64) and keeping only linear terms, gives

$$\epsilon_r = \left| \mathbf{q}_0^{-1}\,\mathbf{Z}\,\mathbf{n}_0 + \mathbf{n}_0\,\overline{\mathbf{Z}}\,\mathbf{q}_0 \right| \tag{24.68}$$

where

$$\mathbf{Z} = \lim_{N\to\infty} \frac{1}{2\pi N} \sum_{j=0}^{N-1} \mathbf{q}_0^{-j} \left(\overrightarrow{\mathbf{q}} \cdot \mathbf{r}_j\right) \mathbf{q}_0^j \tag{24.69}$$

To evaluate $\mathbf{Z}$, the quantity $\mathbf{q}_0^j$ is decomposed into two parts using the fact that $\kappa = \nu_0 + \delta$

$$\mathbf{q}_0^j = e^{i\pi j(\nu_0+\delta)+i\Phi_0}\, \mathbf{n}_0^- + e^{-i\pi j(\nu_0+\delta)-i\phi_0}\, \mathbf{n}_0^+ \tag{24.70}$$

where $\mathbf{n}_0^- \equiv I_-(\mathbf{n}_0)$ and $\mathbf{n}_0^+ \equiv I_+(\mathbf{n}_0)$ are given by Eq. (24.23).

For a particular mode k of oscillation, $\mathbf{r}_j$ is given by Eq. (23.32). The end result will be independent of ϕ_0 in Eq. (23.32) so without loss of generality, Φ_0 will be set to zero. Using this along with the above equations gives

$$\begin{aligned}\mathbf{Z} = \lim_{N\to\infty} \frac{\sqrt{J}}{2\pi N} \sum_{j=0}^{N-1} \Big[&\mathbf{n}_0^+(\overrightarrow{\mathbf{q}} \cdot \mathbf{v}_k)\mathbf{n}_0^-\, e^{i2\pi j(Q_k+\nu_0+\delta)} + \mathbf{n}_0^-(\overrightarrow{\mathbf{q}} \cdot \mathbf{v}_k^*)\mathbf{n}_0^+\, e^{-i2\pi j(Q_k+\nu_0+\delta)} + \\ &\mathbf{n}_0^-(\overrightarrow{\mathbf{q}} \cdot \mathbf{v}_k)\mathbf{n}_0^+\, e^{i2\pi j(Q_k-\nu_0-\delta)} + \mathbf{n}_0^+(\overrightarrow{\mathbf{q}} \cdot \mathbf{v}_k^*)\mathbf{n}_0^-\, e^{-i2\pi j(Q_k-\nu_0-\delta)} + \\ &\mathbf{n}_0^-(\overrightarrow{\mathbf{q}} \cdot \mathbf{v}_k)\mathbf{n}_0^-\, e^{i2\pi jQ_k} + \mathbf{n}_0^+(\overrightarrow{\mathbf{q}} \cdot \mathbf{v}_k^*)\mathbf{n}_0^+\, e^{-i2\pi jQ_k} + \\ &\mathbf{n}_0^+(\overrightarrow{\mathbf{q}} \cdot \mathbf{v}_k)\mathbf{n}_0^+\, e^{i2\pi jQ_k} + \mathbf{n}_0^-(\overrightarrow{\mathbf{q}} \cdot \mathbf{v}_k^*)\mathbf{n}_0^-\, e^{-i2\pi jQ_k} \Big]\end{aligned} \tag{24.71}$$

At a resonance, $\delta = 0$. The first two terms on the right hand side of Eq. (24.71) will be nonzero only at the sum resonance with $Q_k + \nu_0 = p$ for p an integer. The third and fourth terms will be nonzero only at the difference resonance with $Q_k - \nu_0 = p$. And the last four terms will be nonzero only at an integer resonance where $Q_k = p$ which can be ignored since an accelerator can never be operated stably at an integer tune.

The normalized resonance strength ξ_r is defined by

$$\epsilon_r \equiv \sqrt{J}\, \xi_r \tag{24.72}$$

Combining the above equations along with Eqs. (24.24) through (24.27), the strength of the sum resonance for the k mode of oscillation is

$$\xi_{r+} = \frac{\sqrt{2}}{2\pi} \left| \mathcal{G} \cdot \mathbf{v_k} \right| \tag{24.73}$$

and for the difference resonance the strength is

$$\xi_{r-} = \frac{\sqrt{2}}{2\pi} \left| \mathcal{G} \cdot \mathbf{v_k^*} \right| \tag{24.74}$$

where $\mathcal{G}$ is given by

$$\mathcal{G} = \mathbf{2}\, \mathbf{n_0^+} \overrightarrow{\mathbf{q}}\, \mathbf{n_0^-} \tag{24.75}$$

It should be noted that sum resonances are not physically "distinct" from difference resonances. That is, sum resonances become difference resonances and vice versa if the sign of $\mathbf{n}$ and ν are flipped in the analysis (remember that if $\mathbf{n}$ is a valid spin field then so is $-\mathbf{n}$).

The resonance strength can also be calculated within the `SLIM` formalism (§24.9) via Eq. (24.92) or Eq. (24.94). Comparing Eqs. (24.73) or (24.74) with Eq. (24.94) it is seen that $\mathcal{G}$ is the quaternion equivalent of the SLIM $\mathbf{G}$ 2×6 matrix.

24.9 SLIM Formalism

The SLIM formalism[4] [Chao81, Barber99], introduced by Alex Chao, is a way to represent the linearized (that is, first order) orbital and spin transport as an 8×8 matrix which then can be analyzed using standard linear algebra techniques. The idea is to expand the transport map around the closed orbit $(\mathbf{r}_0, \mathbf{n}_0)$ where $\mathbf{r}_0$ is the orbital closed orbit and $\mathbf{n}_0$ is the "spin closed orbit". Namely the unit-vector, one-turn periodic solution of the Thomas-BMT equation on $\mathbf{r}_0$[5]. $\mathbf{n}_0$ is just the invariant spin field on the closed orbit. The formalism provides estimates of the equilibrium spin polarization and the rate of depolarization in electron storage rings, both under the restriction of the aforementioned linearization. Moreover, a procedure known as spin-matching, for minimizing depolarization driven by the noise injected into synchro-betatron motion by synchrotron radiation, and which involves optimizing the layout of the ring, can be executed in a simple and elegant way via the SLIM formalism. The formalism can also give insights into proton spin dynamics in regimes where the linearization approximation suffices.

The SLIM formalism expresses spin components using two right-hand coordinate systems: [6]

$$\begin{aligned} &\big(\mathbf{l}(s), \mathbf{n}_0(s), \mathbf{m}(s)\big) \qquad \text{and} \\ &\big(\mathbf{l}_0(s), \mathbf{n}_0(s), \mathbf{m}_0(s)\big) \end{aligned} \tag{24.76}$$

The axes $\mathbf{l}_0(s)$ and $\mathbf{m}_0(s)$ are solutions of the Thomas-BMT equation on the closed orbit and, generally, are not one-turn periodic. The axes $\mathbf{l}(s)$ and $\mathbf{m}(s)$ are chosen to be one-turn periodic but can have an arbitrary s dependence which can be chosen for convenience otherwise. The axes $\mathbf{l}_0(s)$ and $\mathbf{m}_0(s)$ are used for spin-matching and $\mathbf{l}(s)$ and $\mathbf{m}(s)$ are used for calculating polarization and depolarization. With respect to these axes, a unit-length spin $\mathbf{S}$ can be written as

$$\begin{aligned} &\mathbf{S} = \sqrt{1 - \alpha_0^2 - \beta_0^2}\ \mathbf{n}_0 + \alpha_0\, \mathbf{l}_0 + \beta_0\, \mathbf{m}_0 \qquad \text{or} \\ &\mathbf{S} = \sqrt{1 - \alpha^2 - \beta^2}\ \mathbf{n}_0 + \alpha\, \mathbf{l} + \beta\, \mathbf{m} \end{aligned} \tag{24.77}$$

To linearize the transport, it is assumed that α_0, and β_0 (and hence α and β) are small compared to one. To first order, the variation from unity of the spin component along the $\mathbf{n}_0$ axis will be second order and can be ignored:

$$\begin{aligned} &\mathbf{S} \approx \mathbf{n}_0 + \alpha_0\, \mathbf{l}_0 + \beta_0\, \mathbf{m}_0 \qquad \text{or} \\ &\mathbf{S} \approx \mathbf{n}_0 + \alpha\, \mathbf{l} + \beta\, \mathbf{m} \end{aligned} \tag{24.78}$$

The $\mathbf{n}_0$ coordinate is dropped since the spin component along $\mathbf{n}_0$ is a constant. With this, the eight-dimensional spin-orbit phase space used in the SLIM formalism is

$$\begin{aligned} &(x, p_x, y, p_y, z, p_z, \alpha_0, \beta_0) \qquad \text{or} \\ &(x, p_x, y, p_y, z, p_z, \alpha, \beta) \end{aligned} \tag{24.79}$$

where the orbital part x, p_x, etc. is taken with respect to the closed orbit.

[4] The name references an early computer program that implemented the formalism.

[5] Warning: The symbol $\hat{n}$ or $\vec{n}$ used in [Chao81, Barber85] and other early literature to denote the periodic solution of the T–BMT equation on the closed orbit should be replaced by the symbol $\hat{n}_0$ to conform to the modern convention [Barber99] and thereby avoid confusion with the symbol $\hat{n}$ which denotes the invariant spin field. In addition, the symbols $\vec{m}$ and $\vec{l}$ appearing, for example, in the formulae for the matrix $\mathbf{G}$ in [Barber85], should be replaced by the symbols $\hat{m}_0$ and $\hat{l}_0$, namely by the modern symbols for the two (normally) non-periodic solutions of the T-BMT equation, which together with $\hat{n}_0$, form an orthonormal coordinate system.

[6] Different authors will use different conventions for the ordering of the axes The ordering used here puts $\mathbf{n}_0$ second reflecting the fact that in many rings the $\mathbf{n}_0$ axis will point in the vertical y-direction in the arcs.

The first order map between two any points s_1 and s_2 is an 8×8 matrix $\widetilde{\mathbf{M}}$ which is written in the form

$$\widetilde{\mathbf{M}}(s1, s2) = \begin{pmatrix} \mathbf{M}_{6\times6} & \mathbf{0}_{6\times2} \\ \mathbf{G}_{2\times6} & \mathbf{D}_{2\times2} \end{pmatrix} \tag{24.80}$$

where $\mathbf{M}(s_1, s_2)$ is the 6×6 orbital phase space transport matrix, and $\mathbf{G}(s_1, s_2)$ contains the coupling of the spin coordinates (α_0, β_0) or (α, β) to the orbital motion. The upper right block $\mathbf{0}_{6\times2}$ in the $\widetilde{\mathbf{M}}$ matrix is zero since Stern-Gerlach effects are ignored. When $\mathbf{G}$ is calculated with respect to the $(\mathbf{l}_0, \mathbf{m}_0)$ axes, large spin precessions on the closed orbit due to dipole and solenoid fields are eliminated. That leaves small precessions due to synchro-betatron motion. The $\mathbf{G}$ matrix then represents the dominating linear dependence of the small precessions on the six synchro-betatron coordinates and it then provides a good framework for analysis [Barber85, Barber99]. In Eq. (24.80), $\mathbf{D}$ is a 2×2 rotation matrix for the spin transport of a particle on the closed orbit. In this case, since the $\mathbf{l}_0(s)$ and $\mathbf{m}_0(s)$ are solutions to the T-BMT equation, $\mathbf{D}$ is the unit matrix.

To compute $\widetilde{\mathbf{M}}$ for a section of the ring, the first step is to find the 6×6 orbital matrix for the section. To calculate $\mathbf{n}_0$, $\mathbf{l}_0$ and $\mathbf{m}_0$, first $\mathbf{n}_0$ at some starting point s is calculated (section 18.3) and propagated around the ring. This $\mathbf{n}_0$ is then available for calculations involving the whole ring. If only part of the ring is being analyzed, the orientation of $\mathbf{n}_0$ at the start of the section can be an input parameter. That is, it can be given by the User and not calculated. However, for spin-matching, it usually only makes physical sense to use the $\mathbf{n}_0$ at the start of the section that corresponds to the $\mathbf{n}_0$ calculated for the whole ring. After $\mathbf{n}_0$ is known at some s-position, $\mathbf{l}_0$ and $\mathbf{m}_0$ at that s-position can be chosen somewhat arbitrarily to form the right handed coordinate system. Sometimes it is possible to make a special choice of the initial $\mathbf{l}_0$ and $\mathbf{m}_0$ in order to simplify the $\mathbf{G}$ matrices. For example, in a section where there are only drifts and quadrupoles so that there is no spin rotation for a particle traveling on the centerline, with $\mathbf{n}_0$ pointing vertically, a choice of $\mathbf{m}_0$ pointing in the longitudinal s-direction results in the first line of the $\mathbf{G}$ matrix for the section being zero. After the initial $\mathbf{l}_0$ and $\mathbf{m}_0$ axes have been specified at some initial s, the axes can be transported along the closed orbit of the section.

If the one-turn $\mathbf{G}$ were zero everywhere, the spin motion would be completely decoupled from the orbital motion (at least to first order) and the depolarization rate τ_{dep}^{-1} given by Eq. (24.39) would be zero since $\partial\mathbf{n}/\partial\delta$ would be zero. Therefore, spin-matching analysis for a section of the ring involves adjusting the parameters (quadrupole strengths, drift lengths etc) of the section so as to minimize elements in appropriate columns of the $\mathbf{G}$ matrix. This decreases the rate of depolarization by minimizing $\partial\mathbf{n}/\partial\delta$ at the dipole magnets (where g in Eq. (24.39) is nonzero) [Barber99]. Such adjustments are made while simultaneously maintaining acceptable Courant-Snyder parameters and for this the closed orbit should be taken to be the design orbit. This optimization can be carried out using standard facilities in *Bmad*. The calculation of $\partial\mathbf{n}/\partial\delta$ in the SLIM approximation is described below.

The process for calculating electron polarization and the rate of depolarization in the `SLIM` formalism is as follows. First, the 8×8 matrix $\widetilde{\mathbf{M}}$ for one-turn is calculated as described above. After this, using the closed-orbit spin tune ν_0, a specific version of $(\mathbf{l}, \mathbf{m})$ is constructed by rotating the vectors $\mathbf{l}_0$ and $\mathbf{m}_0$ backwards around $\mathbf{n}_0$ by the angle $2\pi\nu_0$ in a drift space right at the end of the turn, thereby transforming α_0 and β_0 into α and β and transforming the $\mathbf{G}$ matrix correspondingly[7]. The original one-turn $\mathbf{G}$ matrix is not one-turn periodic but the transformed $\mathbf{G}$ matrix is one-turn periodic and the matrix $\mathbf{D}$ for one-turn becomes the 2×2 rotation matrix with rotation angle $2\pi\nu_0$ [Chao81]. The new matrix $\widetilde{\mathbf{M}}$ is then also one-turn periodic and its eigenvectors are used as described after Eq. (24.88) for calculating the derivative $\partial\mathbf{n}/\partial\delta$ used in Eq. (24.39). Note that if the elements in the appropriate columns of the non-periodic $\mathbf{G}$ matrix for the one-turn map at a dipole have been minimized by spin-matching, the corresponding elements of the periodic one-turn $\mathbf{G}$ matrix have been minimized too. As a consequence it can be seen, via Eqs. (24.89) and (24.90), that $\partial\mathbf{n}/\partial\delta$ has been minimized as required.

[7] Adding a rotation at the end is just for convenience. For some other applications it is useful to choose axes $\mathbf{l}$ and $\mathbf{m}$ with respect to which spins precess at the constant rate with a phase advance of $2\pi\nu_0$ per turn.

In contrast to the approach in [Barber85, Barber99], *Bmad* calculates the $\mathbf{G}$ and $\mathbf{D}$ matrices from the quaternion of the spin transport map (which *Bmad* calculates via PTC (§29)). After the $(\mathbf{l}_0, \mathbf{n}_0, \mathbf{m}_0)$ coordinates have been calculated (or set by the User) at some initial point, the spin axes can be transported using the $\mathbf{q}_0$ quaternion (Eq. (24.42)). When analyzing only a section of a ring, there is no identifiable spin tune so nothing further needs to be done. In this case, the $\mathbf{D}$ matrix is just a unit matrix. When analyzing one-turn maps, if the $\mathbf{l}$ and $\mathbf{m}$ axes are set to be the $\mathbf{l}_0(s)$ and $\mathbf{m}_0(s)$ axes except at the end of the lattice, the spin phase advance as a function of s will be zero except just before the starting position where there will be a discontinuous jump in phase.

Once the $(\mathbf{l}_0, \mathbf{n}_0, \mathbf{m}_0)$ axes have been calculated, the matrices $\mathbf{G}$ and $\mathbf{D}$ can be calculated from the spin transport map (which *Bmad* calculates via PTC (§29)). The first order transport map Eq. (24.42) is used. Let $\mathbf{q}_{lnm}(s)$ be the quaternion that transforms from $(\mathbf{l}_0, \mathbf{n}_0, \mathbf{m}_0)$[8] coordinates to (x, y, z) coordinates at a given point s. With this, the spin transport $\widehat{\mathbf{q}}$ from s_1 to s_2 in the $(\mathbf{l}_0, \mathbf{n}_0, \mathbf{m}_0)$ coordinate system is

$$\widehat{\mathbf{q}}_s(s_1, s_2) = \mathbf{q}_{lnm}(s_2)\, \mathbf{q}_s(s_1, s_2)\, \mathbf{q}^{-1}_{lnm}(s_1) \tag{24.81}$$

The zeroth order part of $\mathbf{q}_s$ gives:

$$\widehat{\mathbf{q}}_0(s_1, s_2) = \mathbf{q}_{lnm}(s_2)\, \mathbf{q}_0(s_1, s_2)\, \mathbf{q}^{-1}_{lnm}(s_1) \tag{24.82}$$

represents a rotation around the $\mathbf{n}_0$ axis.

To calculate the $\mathbf{D}$ matrix, $\widehat{\mathbf{q}}_0$ is converted into a 3×3 rotation matrix $\mathbf{R}_0$ via Eq. (24.22)). The second row and second column of this rotation matrix corresponds to the $\mathbf{n}_0$ axis. Since the component of the spin along this axis does not vary to first order, $\mathbf{R}_0$ has the form

$$\mathbf{R}_0 = \begin{pmatrix} R_0(1,1) & 0 & R_0(1,3) \\ 0 & 1 & 0 \\ R_0(3,1) & 0 & R_0(3,3) \end{pmatrix} \tag{24.83}$$

That is, the rotation is around the $\mathbf{n}_0$ axis. Since the $\mathbf{n}_0$ spin component is ignored in the SLIM formalism (Eq. (24.79)), the 2×2 $\mathbf{D}$ matrix is simply $\mathbf{R}_0$ with the second row and second column removed.

$$\mathbf{D}(s_1, s_2) = \begin{pmatrix} R_0(1,1) & R_0(1,3) \\ R_0(3,1) & R_0(3,3) \end{pmatrix} \tag{24.84}$$

In particular, when using the $(\mathbf{l}_0, \mathbf{n}_0, \mathbf{m}_0)$ coordinate system, $\widehat{\mathbf{q}}_0$ represents the identity ($\equiv (1, 0, 0, 0)$) and $\mathbf{D}$ is a unit matrix as expected.

The rows of the $\mathbf{G}$ matrix encode the first-order dependence of the changes of the angles α_0 and β_0 (or of the angles α and β). The $\mathbf{G}$ matrix can therefore be calculated from $\widehat{\mathbf{q}}_i$ which is the first order part of $\widehat{\mathbf{q}}_s$

$$\widehat{\mathbf{q}}_i = \mathbf{q}_{lnm}(s_2)\, \mathbf{q}_i(s_1, s_2)\, \mathbf{q}^{-1}_{lnm}(s_1) \tag{24.85}$$

Using Eq. (24.42) in Eq. (24.22) and keeping only first order terms gives

$$\begin{aligned} \mathbf{G}(1, i) &= 2(\widehat{q}_{0,y}\, \widehat{q}_{i,x} - \widehat{q}_{0,0}\, \widehat{q}_{i,z}) \\ \mathbf{G}(2, i) &= 2(\widehat{q}_{0,0}\, \widehat{q}_{i,x} + \widehat{q}_{0,y}\, \widehat{q}_{i,z}), \quad i = 1, \ldots, 6 \end{aligned} \tag{24.86}$$

where the fact that $\widehat{q}_{0,x} = \widehat{q}_{0,z} = 0$ has been used.[9] [10]

[8] Such a formalism works also with the $(\mathbf{l}, \mathbf{n}_0, \mathbf{m})$

[9] Do not be confused by the x, y and z subscripts which refer to the components of $\widehat{q}$ as defined in Eq. (24.5). $\widehat{q}$ rotates spins in the $(\mathbf{l}_0, \mathbf{n}_0, \mathbf{m}_0)$ coordinate system. Not the (x, y, z) coordinate system.

[10] Unlike the operation of going from the linearized quaternion transport (Eq. (24.42) to the $\widetilde{\mathbf{M}}$ matrix, given an $\widetilde{\mathbf{M}}$

The calculation of the derivative $\partial\mathbf{n}/\partial\delta$ within the SLIM formalism is similar to the calculation using quaternions (§24.6). The following follows Barber[Barber99]. The calculation starts with the one-turn periodic 8×8 matrix $\widetilde{\mathbf{M}}$ [Here the periodic $(\mathbf{l}(s),\mathbf{n}_0(s),\mathbf{m}(s))$ coordinate system must be used since the ending coordinates for $\mathbf{M}$ must be the same as the starting coordinates.] The eigenvectors $\mathbf{u}_k$ and eigenvalues λ_k $(k=1,\ldots 8)$ of $\widetilde{\mathbf{M}}$ are of the form

$$\mathbf{u}_k = \begin{pmatrix}\mathbf{v}_k\\ \mathbf{w}_k\end{pmatrix}, \quad k=1,\ldots,6 \tag{24.88}$$
$$\mathbf{u}_k = \begin{pmatrix}\mathbf{0}_6\\ \mathbf{w}_k\end{pmatrix}, \quad k=7,8$$

where $\mathbf{v}_k$ are eigenvectors of the orbital submatrix $\mathbf{M}$, and for the first six eigenvectors the $\mathbf{w}_k$ are computed via (compare with Eqs. (24.48))

$$\mathbf{w}_k = \left[\lambda_k\,\mathbf{I}_2 - \mathbf{D}\right]^{-1}\mathbf{G}\,\mathbf{v}_k, \quad k=1,\ldots,6 \tag{24.89}$$

where $\mathbf{I}_2$ is the 2×2 unit matrix. These eigenvectors, computed at the chosen starting point $s1$ and are then propagated to other s-positions $s2$ using $\widetilde{\mathbf{M}}(s1,s2)$.

The derivative $\partial\mathbf{n}/\partial\delta$ is computed analogously to Eq. (24.90)

$$\frac{\partial\mathbf{n}}{\partial\delta} = \left(\frac{\partial\alpha}{\partial\delta},\frac{\partial\beta}{\partial\delta}\right) = \sum_{k=1}^{6} A_k\,\mathbf{w}_k \tag{24.90}$$

with the A_k being computed from Eq. (24.54).

Alternatively, Chao [Chao79] gives a an analytical formulation where the eigenvectors are normalized in the form Eq. (23.24). With this, $\partial\mathbf{n}/\partial\delta$ is computed via

$$\frac{\partial\mathbf{n}}{\partial\delta} = i\sum_{k=1}^{6}\mathbf{v}_{k5}^{*}\,\mathbf{w}_k \tag{24.91}$$

The strength of linear resonances can be calculated from the $\mathbf{G}$ matrix. The corresponding equation to Eqs. (24.73) and (24.74) is[Hoff06]

$$\xi_r = \frac{1}{2\,\pi}\left|\vec{\mathbf{G}}\cdot\mathbf{v}_k\right| \tag{24.92}$$

where

$$\vec{\mathbf{G}} = \mathbf{G}(1,:) + i\,\mathbf{G}(2,:) \tag{24.93}$$

with $\mathbf{G}(1,:)$ being the first row of the 2×6 $\mathbf{G}$ matrix and $\mathbf{G}(2,:)$ being the second row. It can be shown that an equivalent way of writing Eq. (24.92) is

$$\xi_r = \frac{\sqrt{2}}{2\,\pi}\left|(\mathbf{G}(1,:)\cdot\mathbf{v}_k,\mathbf{G}(2,:)\cdot\mathbf{v}_k)\right| \tag{24.94}$$

matrix, it is not possible to uniquely construct the $\widehat{\mathbf{q}}_i$ quaternions. $\widehat{q}_{i,x}$ and $\widehat{q}_{i,z}$ can be determined by Eq. (24.86). However, $\widehat{q}_{i,0}$ and $\widehat{q}_{i,y}$ can only be determined via Eq. (24.43) up to an unknown factor κ:

$$\widehat{q}_{i,0} = \kappa\,\widehat{q}_{0,y}, \qquad \widehat{q}_{i,y} = -\kappa\,\widehat{q}_{0,0} \tag{24.87}$$

A finite κ represents a variation of the spin tune with a particle's orbital phase space position. This is, a finite $q_{i,0}$ and $q_{i,y}$ represent a non-linear effect which will average to zero over many turns as a particle with constant orbital amplitude samples different points on the phase space torus it is on.

24.10 Spinor Notation

The following describes the old spinor representation formally used by *Bmad* to represent spins. This documentation is kept as an aid for comparison with the spin tracking literature.

In the SU(2) representation, a spin $\mathbf{S}$ is written as a spinor $\Psi = (\psi_1, \psi_2)^T$ where $\psi_{1,2}$ are complex numbers. The conversion between SU(2) and SO(3) is

$$\mathbf{S} = \Psi^\dagger \boldsymbol{\sigma}\, \Psi \qquad \longleftrightarrow \qquad \Psi = \frac{e^{i\xi}}{\sqrt{2\,(1+s_3)}} \begin{pmatrix} 1+s_3 \\ s_1 + i s_2 \end{pmatrix} \tag{24.95}$$

Where ξ is an unmeasurable phase factor, and $\boldsymbol{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ are the three Pauli matrices

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \qquad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \tag{24.96}$$

In polar coordinates

$$\Psi = \begin{pmatrix} \psi_1 \\ \psi_2 \end{pmatrix} = e^{i\xi} \begin{pmatrix} \cos\frac{\theta}{2} \\ e^{i\phi} \sin\frac{\theta}{2} \end{pmatrix} \qquad \longleftrightarrow \qquad \mathbf{S} = \begin{pmatrix} \sin\theta\cos\phi \\ \sin\theta\sin\phi \\ \cos\theta \end{pmatrix} \tag{24.97}$$

Due to the unitarity of the spin vector, $|\psi_1|^2 + |\psi_2|^2 = 1$. The spinor eigenvectors along the x, y and z axes are

$$\begin{aligned} \Psi_{x+} &= \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}, & \Psi_{x-} &= \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix}, \\ \Psi_{y+} &= \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix}, & \Psi_{y-} &= \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -i \end{pmatrix}, \\ \Psi_{z+} &= \begin{pmatrix} 1 \\ 0 \end{pmatrix}, & \Psi_{z-} &= \begin{pmatrix} 0 \\ -1 \end{pmatrix}. \end{aligned} \tag{24.98}$$

In spinor notation, the T-BMT equation can be written as

$$\frac{\mathrm{d}}{\mathrm{d}t}\Psi = -\frac{i}{2}\,(\boldsymbol{\sigma}\cdot\boldsymbol{\Omega})\,\Psi = -\frac{i}{2} \begin{pmatrix} \Omega_z & \Omega_x - i\,\Omega_y \\ \Omega_x + i\,\Omega_y & -\Omega_z \end{pmatrix} \Psi \tag{24.99}$$

The solution over a time interval Δt, assuming constant $\boldsymbol{\Omega}$, leads to a rotation of the spin vector by an angle $\alpha = |\boldsymbol{\Omega}|\,\Delta t$ around a unit vector $\widehat{\mathbf{a}}$ pointing in the same direction as $\boldsymbol{\Omega}$

$$\begin{aligned} \Psi_f &= \exp\left[-i\frac{\alpha}{2}\widehat{\mathbf{a}}\cdot\boldsymbol{\sigma}\right]\Psi_i \\ &= \left[\cos\left(\frac{\alpha}{2}\right)\mathbf{I}_2 - i\,(\widehat{\mathbf{a}}\cdot\boldsymbol{\sigma})\,\sin\left(\frac{\alpha}{2}\right)\right]\Psi_i \\ &= \mathbf{A}\Psi_i. \end{aligned} \tag{24.100}$$

where Ψ_i is the initial spin state, Ψ_f is the final spin state, and $\mathbf{A}$, describes the spin transport. The Pauli matrices constitute a 2x2 Hermitian-matrix representation of the quaternion components $\mathbf{i}$, $\mathbf{j}$, and $\mathbf{k}$ in Eq. (24.5) and $\mathbf{A}$ is the SU(2) matrix representation of the quaternion $(a_0, \mathbf{a}) = (\cos(\alpha/2), -\sin(\alpha/2)\,\widehat{\mathbf{a}})$. $\mathbf{A}$ has the normalization condition $a_0^2 + \boldsymbol{a}^2 = 1$.

With spinors, the matrix representation of the observable $S_{\mathbf{u}}$ corresponding to the measurement of the spin along the unit vector $\mathbf{u}$ is

$$S_{\mathbf{u}} \equiv \frac{\hbar}{2}\,\boldsymbol{\sigma}\cdot\mathbf{u} \tag{24.101}$$

$$= \frac{\hbar}{2} \begin{pmatrix} u_z & u_x - i\,u_y \\ u_x + i\,u_y & u_z \end{pmatrix} \tag{24.102}$$

The expectation value of this operator, $\Psi^\dagger \mathbf{S}_u \Psi$, representing the spin of a particle, satisfies the equation of motion of a classical spin vector in the particle's instantaneous rest frame.

For a distribution of spins, the polarization P_s along the unit vector $\mathbf{u}$ is defined as the absolute value of the average expectation value of the spin over all N particles times $\frac{2}{\hbar}$,

$$P_s = \frac{2}{\hbar}\frac{1}{N}\sum_{k=1}^{N} \Psi_k^\dagger S_{\mathbf{u}} \Psi_k \tag{24.103}$$

See § §19.8 for formulas for tracking a spin through a multipole fringe field.

Chapter 25

Taylor Maps

25.1 Taylor Maps

A transport map $\mathcal{M} : \mathcal{R}^6 \to \mathcal{R}^6$ through an element or a section of a lattice is a function that maps the starting phase space coordinates $\mathbf{r}(\text{in})$ to the ending coordinates $\mathbf{r}(\text{out})$

$$\mathbf{r}(\text{out}) = \mathcal{M}\,(\delta\mathbf{r}) \tag{25.1}$$

where

$$\delta\mathbf{r} = \mathbf{r}(\text{in}) - \mathbf{r}_{\text{ref}} \tag{25.2}$$

$\mathbf{r}_{\text{ref}}$ is the reference orbit at the start of the map around which the map is made. In many cases the reference orbit is the zero orbit. For a storage ring, the closed orbit is commonly used for the reference orbit. For a lattice with an open geometry the reference orbit may be the orbit as computed from some given initial coordinates.

$\mathcal{M}$ in the above equation is made up of six functions $\mathcal{M}_i : \mathcal{R}^6 \to \mathcal{R}$. Each of these functions maps to one of the $r(\text{out})$ coordinates. Each of these functions can be expanded in a Taylor series and truncated at some order. Each Taylor series is in the form

$$r_i(\text{out}) = \sum_{j=1}^{N} C_{ij} \prod_{k=1}^{6} (\delta r_k)^{e_{ijk}} \tag{25.3}$$

Where the C_{ij} are coefficients and the e_{ijk} are integer exponents. The order of a given term associated with index i, j is the sum over the exponents

$$\text{order}_{ij} = \sum_{k=1}^{6} e_{ijk} \tag{25.4}$$

The order of the entire map is the order at which the map is truncated.

The standard *Bmad* routine for printing a Taylor map might produce something like this:

```
Taylor Terms:
Out      Coef              Exponents          Order       Reference
-------------------------------------------------
 1:     -0.600000000000   0  0  0  0  0  0       0       0.200000000
 1:      1.000000000000   1  0  0  0  0  0       1
```

```
 1:      0.145000000000   2  0  0  0  0  0       2
-------------------------------------------------------
 2:     -0.185000000000   0  0  0  0  0  0       0       0.000000000
 2:      1.300000000000   0  1  0  0  0  0       1
 2:      3.800000000000   2  0  0  0  0  1       3
-------------------------------------------------------
 3:      1.000000000000   0  0  1  0  0  0       1       0.100000000
 3:      1.600000000000   0  0  0  1  0  0       1
 3:    -11.138187077310   1  0  1  0  0  0       2
-------------------------------------------------------
 4:      1.000000000000   0  0  0  1  0  0       1       0.000000000
-------------------------------------------------------
 5:      0.000000000000   0  0  0  0  0  0       0       0.000000000
 5:      0.000001480008   0  1  0  0  0  0       1
 5:      1.000000000000   0  0  0  0  1  0       1
 5:      0.000000000003   0  0  0  0  0  1       1
 5:      0.000000000003   2  0  0  0  0  0       2
-------------------------------------------------------
 6:      1.000000000000   0  0  0  0  0  1       1       0.000000000
```

Each line in the example represents a single `Taylor term`. The Taylor terms are grouped into 6 `Taylor series`. There is one series for each of the output phase space coordinate. The first column in the example, labeled "out", (corresponding to the i index in Eq. (25.3)) indicates the Taylor series: $1 = x(out)$, $2 = p_x(out)$, etc. The 6 exponent columns give the e_{ijk} of Eq. (25.3). In this example, the second Taylor series (`out` = 2), when expressed as a formula, would read:

$$p_x(out) = -0.185 + 1.3\,\delta p_x + 3.8\,\delta x^2\,\delta p_z \tag{25.5}$$

The reference column in the above example shows the input coordinates around which the Taylor map is calculated. In this case, the reference coordinates where

$$(x, p_x, y, p_y, z, p_z)_{\rm ref} = (0.2, 0, 0.1, 0, 0, 0, 0) \tag{25.6}$$

The choice of the reference point will affect the values of the coefficients of the Taylor map. As an example, consider the 1-dimension map

$$x(out) = A\,\sin(k\,\delta x) \tag{25.7}$$

Then a Taylor map to 1^{st} order is

$$x(out) = c_0 + c_1\,\delta x \tag{25.8}$$

where

$$c_1 = A\,k\,\cos(k\,x_{\rm ref}) \tag{25.9}$$

$$c_0 = A\,\sin(k\,x_{\rm ref}) \tag{25.10}$$

Taylor maps using complex numbers is also used by *Bmad*. The output of such maps is similar to the output for real maps as shown above except that the coefficient has a real and imaginary part.

25.2 Spin Taylor Map

A Taylor map that fully describes spin (§24.1) and orbital motion, would consist of nine Taylor series (six for the orbital phase space variables and three for the spin components) and each Taylor series would be a polynomial in nine variables.

To simplify things, *Bmad* assumes that the effect on the orbital phase space due to the spin orientation is negligible. That is, Stern-Gerlach effects are ignored. With this assumption, the orbital part of the map is only dependent on the six orbital variables. This means that $\mathbf{\Omega}_{BMT}$ and $\mathbf{\Omega}_{EDM}$ in the Thomas-Bargmann-Michel-Telegdi equation (§24.1), are assumed independent of the spin. Thus the spin transport is just a rotation. *Bmad* represents this rotation using a quaternion (§24.2). Each of the four components of the quaternion is a Taylor series and the full phase space plus spin transport uses 10 (= 6 orbital + 4 spin) Taylor series with each Taylor series only being dependent on the six orbital phase space coordinates.

Spin transport involves:

1. Using the six orbital coordinates, evaluate the four spin Taylor series to produce a quaternion $\mathbf{q}$.
2. Normalize the quaternion to one: $\mathbf{q} \longrightarrow \mathbf{q}/|\mathbf{q}|$.
3. Rotate the spin vector: $\mathbf{S} \longrightarrow \mathbf{q}\,\mathbf{S}\,\mathbf{q}^{-1}$.

The normalization of the quaternion is needed since the truncation of the map to a finite order will produce errors in the magnitude of the quaternion.

The standard *Bmad* routine for printing a spin Taylor map will produce a result that is very similar as that produced for the orbital phase space. The difference is that there will only be four Taylor series labeled (`S1`, `Sx`, `Sy`, `Sz`) for the four components of the quaternion. Also the reference orbit will not be shown (it is exactly the same as the orbital phase space reference orbit).

Note: When tracking a particle's spin through a map, the quaternion used to rotate the spin is always normalized to one so that the magnitude of the spin will be invariant.

25.3 Symplectification

If the evolution of a system can be described using a Hamiltonian then it can be shown that the linear part of any transport map (the Jacobian) must obey the symplectic condition. If a matrix $\mathbf{M}$ is not symplectic, Healy[Healy86] has provided an elegant method for finding a symplectic matrix that is "close" to $\mathbf{M}$. The procedure is as follows: From $\mathbf{M}$ a matrix $\mathbf{V}$ is formed via

$$\mathbf{V} = \mathbf{S}(\mathbf{I} - \mathbf{M})(\mathbf{I} + \mathbf{M})^{-1} \tag{25.11}$$

where $\mathbf{S}$ is the matrix

$$\mathbf{S} = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 & -1 & 0 \end{pmatrix} \tag{25.12}$$

$\mathbf{V}$ is symmetric if and only if $\mathbf{M}$ is symplectic. In any case, a symmetric matrix $\mathbf{W}$ near $\mathbf{V}$ can be formed via

$$\mathbf{W} = \frac{\mathbf{V} + \mathbf{V}^t}{2} \tag{25.13}$$

A symplectic matrix $\mathbf{F}$ is now obtained by inverting (25.11)

$$\mathbf{F} = (\mathbf{I} + \mathbf{S}\mathbf{W})(\mathbf{I} - \mathbf{S}\mathbf{W})^{-1} \tag{25.14}$$

25.4 Map Concatenation and Feed-Down

Of importance in working with Taylor maps is the concept of `feed-down`. This is best explained with an example. To keep the example simple, the discussion is limited to one phase space dimension so that the Taylor maps are a single Taylor series. Take the map M_1 from point 0 to point 1 to be

$$M_1 : x_1 = x_0 + 2 \tag{25.15}$$

and the map M_2 from point 1 to point 2 to be

$$M_2 : x_2 = x_1^2 + 3\,x_1 \tag{25.16}$$

Then concatenating the maps to form the map M_3 from point 0 to point 2 gives

$$M_3 : x_2 = (x_0 + 2)^2 + 3(x_0 + 2) = x_0^2 + 7\,x_0 + 10 \tag{25.17}$$

However if we are evaluating our maps to only 1^{st} order the map M_2 becomes

$$M_2 : x_2 = 3\,x_1 \tag{25.18}$$

and concatenating the maps now gives

$$M_3 : x_2 = 3(x_0 + 2) = 3\,x_0 + 6 \tag{25.19}$$

Comparing this to Eq. (25.17) shows that by neglecting the 2^{nd} order term in Eq. (25.16) leads to 0^{th} and 1^{st} order errors in Eq. (25.19). These errors can be traced to the finite 0^{th} order term in Eq. (25.15). This is the principal of feed–down: Given M_3 which is a map produced from the concatenation of two other maps, M_1, and M_2

$$M_3 = M_2(M_1) \tag{25.20}$$

Then if M_1 and M_2 are correct to nth order, M_3 will also be correct to nth order as long as M_1 has no constant (0^{th} order) term. [Notice that a constant term in M_2 does not affect the argument.] What happens if we know there are constant terms in our maps? One possibility is to go to a coordinate system where the constant terms vanish. In the above example that would mean using the coordinate $\widetilde{x}_0$ at point 0 given by

$$\widetilde{x}_0 = x_0 + 2 \tag{25.21}$$

25.5 Symplectic Integration

Symplectic integration, as opposed to concatenation, never has problems with feed–down. The subject of symplectic integration is too large to be covered in this guide. The reader is referred to the book "Beam Dynamics: A New Attitude and Framework" by Étienne Forest[Forest98]. A brief synopsis: Symplectic integration uses as input 1) The Hamiltonian that defines the equations of motion, and 2) a Taylor map M_1 from point 0 to point 1. Symplectic integration from point 1 to point 2 produces a Taylor map M_3 from point 0 to point 2. Symplectic integration can produce maps to arbitrary order. In any practical application the order n of the final map is specified and in the integration procedure all terms of order higher than n are ignored. If one is just interested in knowing the final coordinates of a particle at point 2 given the initial coordinates at point 1 then M_1 is just the constant map

$$M_1 : x_1 = c_i \tag{25.22}$$

where c_i is the initial starting point. The order of the integration is set to 0 so that all non–constant terms are ignored. The final map is also just a constant map

$$M_3 : x_2 = c_f \tag{25.23}$$

If the map from point 1 to point 2 is desired then the map M_1 is just set to the identity map

$$M_1 : x_1 = x_0 \tag{25.24}$$

In general it is impossible to exactly integrate any non–linear system. In practice, the symplectic integration is achieved by slicing the interval between point 1 and point 2 into a number of (generally equally spaced) slices. The integration is performed, slice step by slice step. This is analogous to integrating a function by evaluating the function at a number of points. Using more slices gives better results but slows down the calculation. The speed and accuracy of the calculation is determined by the number of slices and the `order` of the integrator. The concept of integrator order can best be understood by analogy by considering the trapezoidal rule for integrating a function of one variable:

$$\int_{y_a}^{y_b} f(y)\,dy = h\left[\frac{1}{2}f(y_a) + \frac{1}{2}f(y_b)\right] + o(h^3\,f^{(2)}) \tag{25.25}$$

In the formula $h = y_b - y_a$ is the slice width. $0(h^3\,f^{(2)})$ means that the error of the trapezoidal rule scales as the second derivative of f. Since the error scales as $f^{(2)}$ this is an example of a second order integrator. To integrate a function between points y_1 and y_N we slice the interval at points $y_2 \ldots y_{N-1}$ and apply the trapezoidal rule to each interval. The concept of integrator order in symplectic integration is analogous.

The optimum number of slices is determined by the smallest number that gives an acceptable error. The slice size is given by the `ds_step` attribute of an element (§6.4). Integrators of higher order will generally need a smaller number of slices to achieve a given accuracy. However, since integrators of higher order take more time per slice step, and since it is computation time and not number of slices which is important, only a measurement of error and calculation time as a function of slice number and integrator order will unambiguously give the optimum integrator order and slice width. In doing a timing test, it must be remembered that since the magnitude of any non-nonlinearities will depend upon the starting position, the integration error will be dependent upon the starting map M_1. *Bmad* has integrators of order 2, 4, and 6 (§6.4). Timing tests performed for some wiggler elements (which have strong nonlinearities) showed that, in this case, the 2^{nd} order integrator gave the fastest computation time for a given accuracy. However, the higher order integrators may give better results for elements with weaker nonlinearities.

Chapter 26

Tracking of Charged Particles

Bmad can track both charged particles and X-rays. This chapter deals with charged particles and X-rays are handled in chapter §27.

For tracking and transfer map calculations (here generically called "tracking"), *Bmad* has various methods that can be applied to a given element (Cf. Chapter §6). This chapter discusses the `bmad_standard` calculation that is the default for almost all element types and the `symp_lie_bmad` calculation that does symplectic integration.

Generally, it will be assumed that tracking is in the forward direction.

26.1 Relative Versus Absolute Time Tracking

Unlike other elements, the kick given a particle going through an `lcavity`, `rfcavity`, or possibly an `em_field` element depends upon the time that the particle enters the element relative to some "RF clock". *Bmad* has two modes for calculating this time called "`relative time tracking`" and "`absolute time tracking`". The switch to set the type of tracking for a lattice is `bmad_com[absolute_time_tracking]` (§11.2).[1] The phase of the RF, $\phi_{\rm rf}$, is determined by

$$\phi_{\rm rf} = \phi_{\rm t} + \phi_{\rm ref} \tag{26.1}$$

where $\phi_{\rm t}$ is the part of the phase that depends upon the time t and $\phi_{\rm ref}$ is a fixed phase offset (generally set in the lattice file) and independent of the particle coordinates. See Eqs. (4.31) and (4.46)

The phase $\phi_{\rm t}$ is

$$\phi_{\rm t} = f_{\rm rf}\, t_{\rm eff} \tag{26.2}$$

where $f_{\rm rf}$ is the RF frequency, and $t_{\rm eff}$ is the effective time. With `relative time tracking`, which *Bmad* uses by default, $t_{\rm eff}$ is a function of the phase space coordinate z (§17.2) via

$$t_{\rm eff}(s) = t_0(s) - t_0(s_{\rm ent}) - \frac{z(s)}{\beta\, c} \tag{26.3}$$

where t_0 is the reference time (see Eq. (17.3)) and $s_{\rm ent}$ is the s-position at the upstream end of the element. $t_{\rm eff}$ is defined such that a particle entering an element with $z = 0$ has $t_{\rm eff} = 0$.

[1] An old, deprecated notation for this switch is `parameter[absolute_time_tracking]`.

With `absolute time tracking`, and `bmad_com[absolute_time_ref_shift]` set to True (the default), $t_{\rm eff}$ is defined by

$$t_{\rm eff}(s) = t(s) - t_0(s_{\rm ent}) \tag{26.4}$$

$t_0(s_{\rm ent})$, by definition, equal to the time of the reference particle at the entrance end of the element. With multipass §9, $t_0(s_{\rm ent})$ is set by the time of the reference particle at the entrance end of the element on the first pass. For absolute time tracking, it is important to keep in mind that $t_0(s_{\rm ent})$ is a property of the element independent of how tracking is done. Thus, if a particle goes through a particular element multiple times, the value of $t_{\rm ent}$ will be the same for each transit. If `bmad_com[absolute_time_ref_shift]` set to False, $t_{\rm eff}$ is simply

$$t_{\rm eff}(s) = t(s) \tag{26.5}$$

To understand the difference between relative and absolute time tracking, consider a particle traveling on the reference orbit (zero orbit) along side the reference particle in a circular ring with one RF cavity. This particle always has $z = 0$ and thus, with `relative time tracking`, $t_{\rm eff}$ will always be zero (assuming `bmad_com[absolute_time_ref_shift]` is set to True). On the other hand, with `absolute time tracking`, while the particle, on the first turn, will have $t_{\rm eff}$ equal to zero, on subsequent turns (or subsequent passes if using multipass), the time will increase on each turn by the zero orbit revolution time. If the RF frequency $f_{\rm rf}$ is some multiple of the revolution harmonic, the RF phase with absolute vs relative time tracking will be some multiple of $2\,\pi$ and thus RF kick given the particle will be the same in both cases. However, if the RF frequency is not some multiple of the revolution harmonic, there will be a difference in the RF kicks (except for the kick on the first turn).

There are advantages and disadvantages to using either relative or absolute time tracking. Absolute time tracking is more correct since RF cavities may have frequencies that are not commensurate with the reference revolution time. The problem with absolute time tracking is that the transfer map through the cavity is now a function of time and therefore is a function of z *and* the turn number. This complicates lattice analysis. For example, standard element transfer maps use phase space coordinates so with absolute time tracking, one has a different map for each turn.

With relative time tracking the transfer map problem is swept under the rug. The penalty for using relative time tracking is that results can be unphysical. For example, with relative time tracking, the closed orbit is essentially independent of the RF frequency. In an actual machine, the closed orbit will vary with RF frequency and, in fact, varying the RF frequency is one way of measuring the dispersion. From a different angle this independence can be viewed as a desirable feature since if one is only interested in, say, calculating the Twiss parameters, it can be an annoyance to have to worry that the ring one has constructed have a length that is exactly commensurate with the RF frequency. And it is potentially confusing to see non-zero closed orbits when one is not expecting it due to a mismatch between the ring circumference and the RF frequency or due to RF cavities not being spaced a multiple of the RF wavelength apart.

The above discussion is limited to the cavity fundamental mode. Long-range wakefields, on the other hand, cannot be synchronized to the z coordinate since, in general, their frequencies are not commensurate with the fundamental mode frequency. For simulating the long-range wakes, the kick is thus, by necessity, tied to the absolute time.

Do not confuse absolute time tracking with the `time_runge_kutta` tracking method (§6.1). The `time_runge_kutta` method uses time as the independent variable instead of z. Absolute time tracking just means that the RF phase is dependent upon the time instead of z. It is perfectly possible to use absolute time tracking with code that uses z as the independent variable.

One important point to always keep in mind is that any PTC based tracking (§1.4) will always use relative time tracking independent of the setting of `bmad_com[absolute_time_tracking]`.

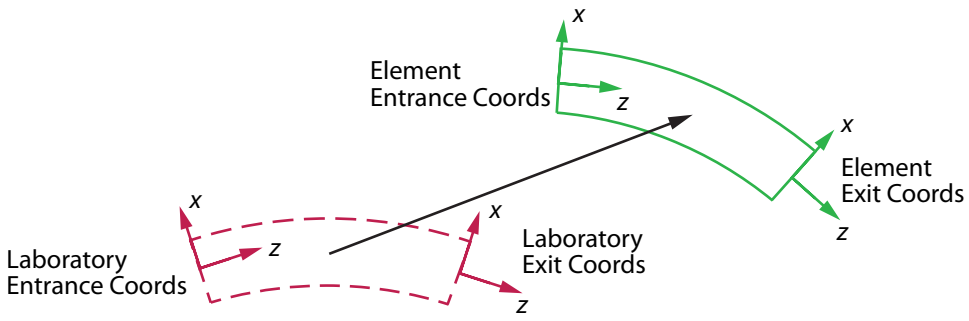

Figure 26.1: `Element` coordinates are coordinates attached to the physical element (solid green outline). The `laboratory` coordinates are fixed at the nominal position of the element (red dashed outline).

.

26.2 Element Coordinate System

The general procedure for tracking through an element makes use of `element reference` coordinates (also called just `element` coordinates). Without any offsets, pitches or tilt (§5.6), henceforth called "misalignments", the `element` coordinates are the same as the `laboratory reference` coordinates (or simply `laboratory` coordinates) (§16.1.1). The `element` coordinates stay fixed relative to the element. Therefore, if the element is misaligned, the `element coordinates` will follow as the element shifts in the laboratory frame as shown in Fig. 26.1.

Tracking a particle through an element is a three step process:

1. At the entrance end of the element, transform from the `laboratory` coordinates to the entrance `element` coordinates.
2. Track through the element ignoring any misalignments.
3. At the exit end of the element, transform from the exit `element` reference frame to the `laboratory` reference frame.

The transformation between `laboratory` and `element` reference frames is given in §16.3.1 and §16.3.2.

26.3 Hamiltonian

The time dependent Hamiltonian H_t in the curvilinear coordinate system shown in Fig. 16.2 is ([Ruth87])

$$H_t = \widetilde{\psi} + \left[\left(\frac{p_s - a_s}{1 + g\,x} \right)^2 + \widetilde{m}^2 + (p_x - a_x)^2 + (p_y - a_y)^2 \right]^{1/2} \tag{26.6}$$

where $(p_x, p_y, p_s/(1+gx))$ are the momentum normalized by P_0, ρ being the local radius of curvature of the reference particle, and $\widetilde{m}$, $\mathbf{a}$ and $\widetilde{\psi}$ are the normalized mass, vector, and scalar potentials:

$$\widetilde{m} = \frac{m\,c^2}{c\,P_0} \qquad \left(a_x, a_y, \frac{a_s}{1+g\,x} \right) = \frac{q\,\mathbf{A}}{P_0} \qquad \widetilde{\psi}(x,y,z) = \frac{q\,\psi}{P_0\,c} \tag{26.7}$$

In terms of the normalized velocities β_x, β_y, the canonical momentum are

$$p_x = \frac{m\,c^2}{P_0\,c}\,\gamma\,\beta_x + a_x, \qquad p_y = \frac{m\,c^2}{P_0\,c}\,\gamma\,\beta_y + a_y \tag{26.8}$$

The s-dependent Hamiltonian is obtained from H_t by solving for $-p_s$ and using a contact transformation to convert to *Bmad* coordinates (§17.2). For particles propagating in the positive s direction, the s-dependent Hamiltonian is, assuming $\widetilde{\psi}$ is zero

$$H \equiv H_s = -(1+g\,x)\sqrt{(1+p_z)^2-(p_x-a_x)^2-(p_y-a_y)^2} - a_s + \frac{1}{\beta_0}\sqrt{(1+p_z)^2+\widetilde{m}^2} \tag{26.9}$$

where β_0 is the reference velocity and the equality $(1+p_z)^2 = (E/c\,P_0)^2 - \widetilde{m}^2$ has been used. The last term on the RHS of Eq. (26.9) accounts for the fact that the *Bmad* canonical z (Eq. (17.3)) has an "extra" term $\beta\,c\,t_0$ so that *Bmad* canonical z is with respect to the reference particle's z.

The equations of motion are

$$\frac{dq_i}{ds} = \frac{\partial H}{\partial p_i} \qquad \frac{dp_i}{ds} = -\frac{\partial H}{\partial q_i} \tag{26.10}$$

Without an electric field, ψ is zero. Assuming a non-curved coordinate system ($g = 0$), and using the paraxial approximation (which expands the square root in the Hamiltonian assuming the transverse momenta are small) (§17.2), Eq. (26.9) becomes

$$H = \frac{(p_x-a_x)^2}{2(1+p_z)} + \frac{(p_y-a_y)^2}{2(1+p_z)} - (1+p_z) - a_s + \frac{1}{\beta_0}\sqrt{(1+p_z)^2+\widetilde{m}^2} \tag{26.11}$$

Once the transverse trajectory has been calculated, the longitudinal position z_2 at the exit end of an element is obtained from symplectic integration of Eq. (26.11)

$$z_2 = z_1 - \frac{1}{2(1+p_{z1})^2}\int ds\,\left[(p_x-a_x)^2+(p_y-a_y)^2\right] - \int ds\,g\,x \tag{26.12}$$

where z_1 is the longitudinal position at the entrance end of the element. Using the equations of motion Eqs. (26.10) this can also be rewritten as

$$z_2 = z_1 - \frac{1}{2}\int ds\,\left[\left(\frac{dx}{ds}\right)^2+\left(\frac{dy}{ds}\right)^2\right] - \int ds\,g\,x \tag{26.13}$$

For some elements, `bmad_standard` uses a truncated Taylor map for tracking. For elements without electric fields where the particle energy is a constant, the transfer map for a given coordinate r_i may be expanded in a Taylor series

$$r_{i,2} \to m_i + \sum_{j=1}^{4} m_{ij}\,r_{j,1} + \sum_{j=1}^{4}\sum_{k=j}^{4} m_{ijk}\,r_{j,1}\,r_{k,1} + \ldots \tag{26.14}$$

where the map coefficients $m_{ij\ldots}$ are functions of p_z. For linear elements, the transfer map is linear for the transverse coordinates and quadratic for $r_i = z$.

Assuming mid–plane symmetry of the magnetic field, so that a_x and a_y can be set to zero[Iselin94], The vector potential up to second order is (cf. Eq. (18.1))

$$a_s = -k_0\left(x - \frac{g\,x^2}{2(1+g\,x)}\right) - \frac{1}{2}k_1\left(x^2-y^2\right) \tag{26.15}$$

For backwards propagation, where particle are traveling in the $-\mathbf{s}$ direction and where p_s is negative, solving for p_s involves using a different part of the square root branch. There is also an overall negative sign coming from switching from using s as the independent variable to $\widetilde{s} \equiv -s$ as the independent variable. the Hamiltonian $H_{\widetilde{s}}$ is then

$$H_{\widetilde{s}} = -(1+g\,x)\sqrt{(1+p_z)^2-(p_x-a_x)^2-(p_y-a_y)^2} + a_s + \frac{1}{\beta_0}\sqrt{(1+p_z)^2+\widetilde{m}^2} \tag{26.16}$$

26.4 Symplectic Integration

Using Eq. (26.11) the Hamiltonian is written in the form

$$H = H_x + H_y + H_z \tag{26.17}$$

where

$$H_x = \frac{(p_x - a_x)^2}{2(1+\delta)}, \qquad H_y = \frac{(p_y - a_y)^2}{2(1+\delta)}, \qquad H_s = -a_s \tag{26.18}$$

For tracking, the element is broken up into a number of slices set by the element's `ds_step` attribute. For each slice, the tracking uses a quadratic symplectic integrator I:

$$I = T_{s/2}\; I_{x/2}\; I_{y/2}\; I_s\; I_{y/2}\; I_{x/2}\; T_{s/2} \tag{26.19}$$

$T_{s/2}$ is just a translation of the s variable:

$$s \rightarrow s + \frac{ds}{2} \tag{26.20}$$

And the other integrator components are

$$\begin{aligned} I_{x/2} &= \exp\left(: -\frac{ds}{2} H_x :\right) \\ I_{y/2} &= \exp\left(: -\frac{ds}{2} H_y :\right) \\ I_s &= \exp\left(: -ds\, H_s :\right) \end{aligned} \tag{26.21}$$

The evaluation of $I_{x/2}$ and $I_{y/2}$ is tricky since it involves both transverse position and momentum variables. The trick is to split the integration into three parts. For $I_{x/2}$ this is

$$\begin{aligned} I_{x/2} &= \exp\left(: -\frac{ds}{2} \frac{(p_x - A_x)^2}{2(1+\delta)} :\right) \\ &= \exp\left(: -\int A_x\, dx :\right) \exp\left(: -\frac{ds}{2} \frac{p_x^2}{2(1+\delta)} :\right) \exp\left(: \int A_x\, dx :\right) \end{aligned} \tag{26.22}$$

With an analogous expression for $I_{y/2}$.

For magnetic elements that do not have longitudinal fields (quadrupoles, sextupoles, etc.), a_x and a_y can be taken to be zero (cf. Eq. (26.15)).

For `lcavity` and `rfcavity` elements, the vector potential is computed from Eq. (18.60).

26.5 BeamBeam Tracking

A beam-beam element (§4.3) simulates the effect on a tracked particle of an opposing beam of particles moving in the opposite direction. The opposing beam, called the "strong" beam, is assumed to be Gaussian in shape.

The strong beam is divided up into `n_slice` equal charge (not equal thickness) slices. Propagation through the strong beam involves a kick at the charge center of each slice with drifts in between the kicks. The kicks are calculated using the standard Bassetti–Erskine complex error function formula[Talman87].

Even though the strong beam can have a finite `sig_z`, the length of the element is always considered to be zero. This is achieved by adding drifts at either end of any tracking so that the longitudinal starting point and ending point are identical. The longitudinal s–position of the `BeamBeam` element is at the center of the strong bunch. For example, with `n_slice` = 2 and with a solenoid field, the calculation would proceed as follows:

1. Start with the particle longitudinally at the `beambeam` element (which is considered to have zero longitudinal length) in laboratory coordinates (§16).
2. Propagate backwards through the solenoid field so that the particle is in the plane of the first beambeam slice. The fact that the plane of the slice may be, due to finite `x_pitch` or `y_pitch` values, canted with respect to the laboratory x-y plane is taken into account.
3. Transform the particle coordinates to the `beambeam` element body coordinates (§16.3).
4. Apply the beam–beam kick due to the first slice including a spin rotation.
5. Transform back to laboratory coordinates.
6. Propagate forwards so that the particle is in the plane of the second slice.
7. Transform the particle coordinates to the `beambeam` element body coordinates.
8. Apply the beam–beam kick due to the second slice.
9. Transform back to laboratory coordinates.
10. Propagate backwards through the solenoid field to end up with the particle longitudinally at the `beambeam` element.

There is an energy kick due to the motion of the strong beam. There are two parts to this $dp_z = dp_{z,s} + dp_{z,h}$. One part, $dp_{z,s}$, is similar to the gravitational slingshot in orbital mechanics. The slingshot energy kick is simply calculated using conservation of 4-momentum of the tracked particle and the strong beam where the mass of the strong beam is assumed to be large compared to the mass of the tracked particle.[2] After a little bit of algebra. The energy kick dE_w to lowest order in the angle of the weak particle with respect to the axis defined by the of motion of the strong beam is

$$dE_w = \frac{c\,P_w}{2\,(1/\beta_w + 1/\beta_s)}\,\left(\theta_{w2}^2 - \theta_{w1}^2\right) \tag{26.23}$$

where P_w is the momentum of the weak particle, β_w and β_s are the weak and strong beam velocities, and θ_{w1} and θ_{w2} are angles of the weak particle trajectory with respect to the strong beam motion before and after the interaction. Converting to phase space coordinates, the momentum kick dp_z is

$$dp_{z,s} = \frac{1}{2\,\beta_w\,(1/\beta_w + 1/\beta_s)\,(1+p_z)}\,\left(dp_x\,(dp_x + 2p_{x1}) + dp_y\,(dp_y + 2p_{y1})\right) \tag{26.24}$$

where dp_x and dp_y are the transverse kicks and p_{x1} and p_{y1} are the initial phase space momenta.

The other part of the energy kick, $dp_{z,h}$, happens when the strong beam's cross-section is changing due to the hourglass effect. The hourglass longitudinal kick relative to the transverse kicks can derived using Eq. (5) of Sagan [Sagan91]. In the relativistic limit the result is

$$dp_{z,h} = \frac{\sigma_x}{2}\,\frac{d\sigma_x}{ds}\,\frac{dp_x}{dx} + \frac{\sigma_y}{2}\,\frac{d\sigma_y}{ds}\,\frac{dp_y}{dy} \tag{26.25}$$

where σ_x and σ_y are the strong beam sizes, and the factor of two is due to the relative velocity ($2c$) between the beams.

[2] This assumption breaks down if a tracked particle is deflected due to a single scattering event with a particle of the strong beam. But particle-particle scattering is outside of the assumption of a strong beam that is unaffected by the weak beam.

26.6 Bend: Exact Body Tracking with k1 = 0

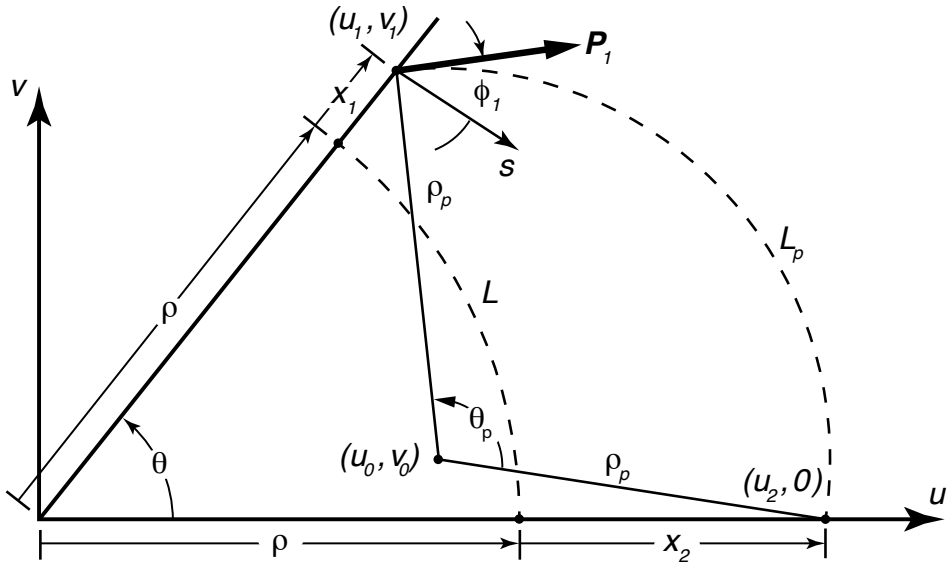

Figure 26.2: Geometry for the exact bend calculation.

Function definitions:

$$\text{sinc}(x) \equiv \frac{\sin(x)}{x} \tag{26.26}$$

$$\text{cosc}(x) \equiv \frac{1-\cos(x)}{x^2} \tag{26.27}$$

These functions cannot be directly evaluated at $x = 0$ and are defined at $x = 0$ using the $x \to 0$ limit. The point to keep in mind here is that these functions are well behaved and can be easily coded in software.

Referring to Figure 26.2, at point 1 where the particle enters a sector bend, the angle ϕ_1 of the particle trajectory in the (x, s) plane with respect to the s axis is

$$\sin(\phi_1) = \frac{p_{x1}}{\sqrt{(1+p_z)^2 - p_y^2}} \tag{26.28}$$

where the subscript "1" for p_z and p_y is dropped since these quantities are invariant.

The (u, v) coordinate system in the plane of the bend is defined with the u-axis along the exit edge of the bend and the v-axis is perpendicular to the u-axis. The origin is at the design center of the bend. In this coordinate system the point (u_1, v_1) where the particle enters the bend is given by

$$u_1 = (\rho + x_1)\,\cos(\theta) \tag{26.29}$$

$$v_1 = (\rho + x_1)\,\sin(\theta) \tag{26.30}$$

where ρ is the design radius of curvature, x_1 is the offset of the particle from the design at the entrance point, and θ is the design bend angle

$$\theta = \frac{L}{\rho} = g\,L \tag{26.31}$$

with L being the design arc length and $g \equiv 1/\rho$.

The coordinates (u_0, v_0) of the center of curvature of the particle trajectory is

$$u_0 = u_1 - \rho_p \cos(\theta + \phi_1) \tag{26.32}$$

$$v_0 = v_1 - \rho_p \sin(\theta + \phi_1) \tag{26.33}$$

where ρ_p is the radius of curvature of the particle trajectory in the (u, v) plane (see Eq. (26.39)).

The coordinates of the particle at the exit face is $(u_2, 0)$ where

$$u_2 = u_0 + \sqrt{\rho_p^2 - v_0^2} \tag{26.34}$$

After some manipulation, the offset of the particle x_2 from the design point at the exit face is

$$x_2 = u_2 - \rho = x_1 \cos(\theta) - L^2 \, g \, \mathrm{cosc}(\theta) + \xi \tag{26.35}$$

where ξ can be expressed in two different ways

$$\xi = \frac{\alpha}{\left[\cos^2(\theta + \phi_1) + g_p \, \alpha\right]^{1/2} + \cos(\theta + \phi_1)} \quad \text{or} \tag{26.36}$$

$$= \frac{\left[\cos^2(\theta + \phi_1) + g_p \, \alpha\right]^{1/2} - \cos(\theta + \phi_1)}{g_p} \tag{26.37}$$

where

$$\alpha = 2 \, (1 + g \, x_1) \, \sin(\theta + \phi_1) \, L \, \mathrm{sinc}(\theta) - g_p \, (1 + g \, x_1)^2 \, L^2 \, \mathrm{sinc}^2(\theta) \tag{26.38}$$

$$g_p = \frac{1}{\rho_p} = \frac{g_{\mathrm{tot}}}{\sqrt{(1 + p_z)^2 - p_y^2}} \tag{26.39}$$

In the above equation g_{tot} is the bending strength of the actual field. Both Eq. (26.36) and Eq. (26.37) are needed since Eq. (26.36) is singular when $\alpha = 0$ and $\theta + \phi_1 = \pi$ (which happens when the particle is bent by 180°), and Eq. (26.37) is singular when g_p is zero. A simple way to implement the calculation for x_2 is to use Eq. (26.36) when $|\theta + \phi_1| < \pi/2$ and otherwise use Eq. (26.37).

Once x_2 is computed, the arc length of the particle L_p is

$$L_p = \frac{|\mathbf{L}_c|}{\mathrm{sinc}(\theta_p/2)} \tag{26.40}$$

where $\mathbf{L}_c$ is the vector (chord) from point 1 and point 2

$$\mathbf{L}_c = (L_{cu}, L_{cv}) = (\xi, -L \, \mathrm{sinc}(\theta) - x_1 \, \sin(\theta)) \tag{26.41}$$

and θ_p is the angle made by the particle trajectory which is twice the angle between the initial particle trajectory $\mathbf{P}$ and the vector $\mathbf{L}_c$

$$\theta_p = 2 \, (\theta + \phi_1 - \mathrm{atan2}\,(L_{cu}, -L_{cv})) \tag{26.42}$$

where $\mathrm{atan2}(y, x)$ is the standard two argument arctangent function.

Once L_p is computed, p_{x2}, y_2 and z_2 are easily derived from

$$p_{x2} = \sqrt{(1 + p_z)^2 - p_y^2} \, \sin(\theta + \phi_1 - \theta_p) \tag{26.43}$$

$$y_2 = y_1 + \frac{p_y \, L_p}{\sqrt{(1 + p_z)^2 - p_y^2}} \tag{26.44}$$

$$z_2 = z_1 + \frac{\beta \, L}{\beta_{\mathrm{ref}}} - \frac{(1 + p_z) \, L_p}{\sqrt{(1 + p_z)^2 - p_y^2}} \tag{26.45}$$

where β is the normalized velocity of the particle and $\beta_{\rm ref}$ if the normalized velocity of the reference particle.

Using the above equation, round-off error will give a non-zero final position even if the initial position is zero. Even though the round-off error will be very small, a non-zero result can be confusing. To avoid this, the standard linear transfer matrix for a bend is used if all the following conditions are satisfied:

$$|x\,g|, |p_x|, |p_y|, |p_z| < 10^{-9}, \quad \text{and}, \quad g_{\rm tot} = g \tag{26.46}$$

The matrix is:

$$\begin{pmatrix} \cos(\theta) & L\,\mathrm{sinc}(\theta) & 0 & 0 & 0 & g\,L^2\,\mathrm{cosc}(\theta) \\ -g\,\sin(\theta) & \cos(\theta) & 0 & 0 & 0 & g\,L\,\mathrm{sinc}(\theta) \\ 0 & 0 & 1 & L & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ -g\,L\,\mathrm{sinc}(\theta) & -g\,L^2\,\mathrm{cosc}(\theta) & 0 & 0 & 1 & L\,\left(\frac{1}{\gamma^2} - g^2\,L^2\,\mathrm{sincc}(\theta)\right) \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \tag{26.47}$$

where

$$\mathrm{sincc}(\theta) \equiv \frac{x - \sin(x)}{x^3} \tag{26.48}$$

26.7 Bend: Body Tracking with finite k1

For a bend with a finite `k1`, the Hamiltonian for the body of an `sbend` is

$$H = (g_{\rm tot} - g)\,x - g\,x\,p_z + \frac{1}{2}\left((k_1 + g\,g_{\rm tot})x^2 - k_1\,y^2\right) + \frac{p_x^2 + p_y^2}{2(1+p_z)} \tag{26.49}$$

This is simply solved

$$\begin{aligned} x_2 &= c_x\,(x - x_c) + s_x\,\frac{p_{x1}}{1+p_{z1}} + x_c \\ p_{x2} &= \tau_x\,\omega_x^2\,(1+p_{z1})\,s_x\,(x - x_c) + c_x\,p_{x1} \\ y_2 &= c_y\,y_1 + s_y\,\frac{p_{y1}}{1+p_{z1}} \\ p_{y2} &= \tau_y\,\omega_y^2\,(1+p_{z1})\,s_y\,y_1 + c_y\,p_{y1} \\ z_2 &= z_1 + m_5 + m_{51}(x - x_c) + m_{52}p_{x1} + m_{511}\,(x - x_c)^2 + \\ &\qquad m_{512}\,(x - x_c)\,p_{x1} + m_{522}\,p_{x1}^2 + m_{533}\,y^2 + m_{534}\,y_1\,p_{y1} + m_{544}\,p_{y1}^2 \\ p_{z2} &= p_{z1} \end{aligned} \tag{26.50}$$

where

$$\begin{aligned} k_x &= k_1 + g\,g_{\rm tot} & \omega_x &\equiv \sqrt{\frac{|k_x|}{1+p_{z1}}} \\ x_c &= \frac{g\,(1+p_{z1}) - g_{\rm tot}}{k_x} & \omega_y &\equiv \sqrt{\frac{|k_1|}{1+p_{z1}}} \end{aligned} \tag{26.51}$$

and

$$\begin{array}{lll} & k_x > 0 & k_x < 0 \\ c_x = & \cos(\omega_x L) & \cosh(\omega_x L) \\ s_x = & \dfrac{\sin(\omega_x L)}{\omega_x} & \dfrac{\sinh(\omega_x L)}{\omega_x} \\ \tau_x = & -1 & +1 \end{array} \qquad \begin{array}{lll} & k_1 > 0 & k_1 < 0 \\ c_y = & \cosh(\omega_y L) & \cos(\omega_y L) \\ s_y = & \dfrac{\sinh(\omega_y L)}{\omega_y} & \dfrac{\sin(\omega_y L)}{\omega_y} \\ \tau_y = & +1 & -1 \end{array} \tag{26.52}$$

and

$$\begin{aligned} m_5 &= -g\, x_c\, L \\ m_{51} &= -g\, s_x & m_{52} &= \frac{\tau_x\, g}{1+p_{z1}} \frac{1-c_x}{\omega_x^2} \\ m_{511} &= \frac{\tau_x\, \omega_x^2}{4}\,(L - c_x\, s_x) & m_{533} &= \frac{\tau_y\, \omega_y^2}{4}\,(L - c_y\, s_y) \\ m_{512} &= \frac{-\tau_x\, \omega_x^2}{2\,(1+p_{z1})}\, s_x^2 & m_{534} &= \frac{-\tau_y\, \omega_y^2}{2\,(1+p_{z1})}\, s_y^2 \\ m_{522} &= \frac{-1}{4\,(1+p_{z1})^2}\,(L + c_x\, s_x) & m_{544} &= \frac{-1}{4\,(1+p_{z1})^2}\,(L + c_y\, s_y) \end{aligned}$$

26.8 Bend: Fiducial Point Calculations

When the `fiducial_pt` switch for a bend is set to something other than `none`, changing one of `rho`, `g`, `b_field` or `angle` in a program (that is, changing after the lattice has been read in and the bend parameters calculated) involves adjustment to the other three parameters along with adjustment to `e1`, `e2`, `l`, `l_chord`, and `l_rectangle`. This is done to keep the shape of the bend invariant. Invariance is not maintained with variation of any other parameter (EG variation of `e1`).

Fig. 26.3 shows the situation when the `fiducial_pt` is set to either `entrance_end` or `center` (the situation for the `exit_end` setting is analogous to the `entrance_end` setting and so is not discussed). For any one of the `fiducial_pt` settings discussed there are essentially two cases. One case is direct variation of the bend field via variation of `rho`, `g`, or `b_field`. This is called "`g`-variation". The other type of variation is variation of `angle`. This is called "angle-variation". The discussion below shows how, with `g`-variation, `l`, `e1`, and `e2` are calculated. With `angle`-variation, `l`, `g`, `e1`, and `e2` need to be calculated. Once `l` and `g` are know, the other parameters `l_chord`, `l_rectangle`, `l_sagitta` (and `angle` for the `g`-variation case) can be readily computed.

The `entrance_end` analysis is as follows (Fig. 26.3a). The entrance end coordinates around the point $\mathbf{r}_1$ are held fixed and as as a result $\mathbf{r}_1' = \mathbf{r}_1$ and `e1` does not vary as well. $\mathbf{r}_2$ is the exit point before variation and $\mathbf{r}_3$ is the exit point after. The position of $\mathbf{r}_3$ is calculated by first calculating the position of $\mathbf{r}_1$ in a coordinate system centered at $\mathbf{r}_2$ and with axes parallel to the $(\mathbf{s}_1, \mathbf{x}_1)$ axes of the coordinate system at $\mathbf{r}_1$

$$\bar{\mathbf{r}}_1 = (-l_{\text{rectangle}}, \rho\,(1 - \cos\alpha)) \tag{26.53}$$

Where the bar denotes that the coordinates are in the $(\mathbf{s}_1, \mathbf{x}_1)$ system. The coordinates of $\mathbf{r}_1$ in the $(\mathbf{e}_s, \mathbf{e}_x)$ coordinate system with origin at $\mathbf{r}_2$ and with $\mathbf{e}_x$ along the bend edge and $\mathbf{e}_s$ perpendicular to $\mathbf{e}_x$ is a rotation $\mathbf{R}(\theta)$

$$\mathbf{r}_1 = \mathbf{R}(\alpha - e_2)\, \bar{\mathbf{r}}_1 \tag{26.54}$$

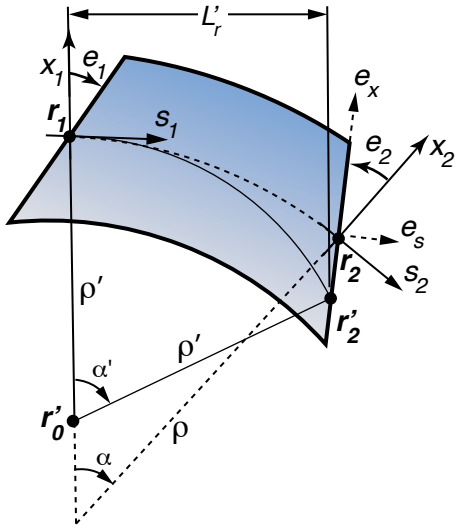

(a) With `fiducial_pt` set to `entrance_end`, $\mathbf{r}_1$ is the fiducial point at the entrance end. By construction, the entrance point $\mathbf{r}_1$ and the slope of the reference curve at $\mathbf{r}_1$ is invariant with the reference curve before (dashed line) and after (solid line) being tangent to $\mathbf{s}_1$ where $\mathbf{s}_1$ being the perpendicular to $\mathbf{x}_1$.

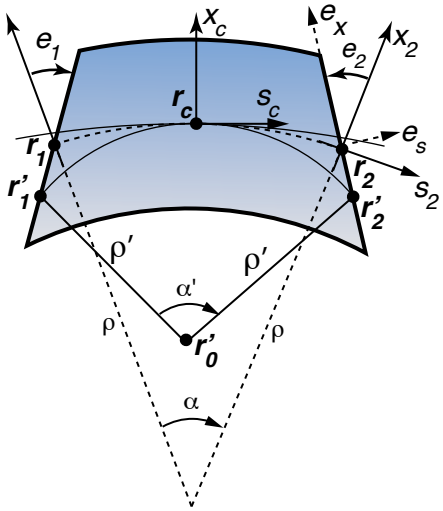

(b) With `fiducial_pt` set to `center`, $\mathbf{r}_c$ is the fiducial point at the center. By construction, the reference curve always goes through $\mathbf{r}_c$ and the tangent of the reference curve at $\mathbf{r}_c$ is invariant.

Figure 26.3: Geometry with `fiducial_pt` set to (a) `entrance_end` and (b) `center`. In both cases, $\mathbf{r}_1$ and $\mathbf{r}_2$ are the entrance and exit reference points before and $\mathbf{r}'_1$ and $\mathbf{r}_2$ are the entrance and exit points after variation of one of `rho`, `g`, `b_field`, or `angle`. Similarly, ρ and α are the bending radius and bending angle before variation while ρ' and α' are the bending radius afterwards. Finally, e_1 e_2 are the face angles and rectangular length before variation, and L'_r and $\mathbf{r}'_0$ are the rectangular length and center of curvature after variation.

The angle θ_1 of the vector $\mathbf{s}_1$, which is the invariant tangent of the reference curve at the point $\mathbf{r}_1$, in the $(\mathbf{e}_s, \mathbf{e}_x)$ coordinate system (which is used from here on) is

$$\theta_1 = \alpha - e_1 \tag{26.55}$$

The center of curvature after variation $\mathbf{r}'_0$ is

$$\mathbf{r}'_0 = \mathbf{r}_1 + \rho' \ (\sin\theta_1, -\cos\theta_1) \tag{26.56}$$

The reference trajectory after variation $\mathbf{r}'$ is a circular arc subject to the condition

$$|\mathbf{r}' - \mathbf{r}'_0| = \rho'^2 \tag{26.57}$$

With g-variation, the value of ρ' is set (perhaps indirectly) by the User. To find the point $\mathbf{r}'_2$, it is noted that in the (e_s, e_x) coordinate system, the s coordinate of $\mathbf{r}'_2$, r'_{2s} is zero. so using this inEq. (26.57) and throwing away the unphysical root gives for the x coordinate

$$r'_{2x} = r_{1x} + \frac{2\,c}{-b - \sqrt{b^2 - 4\,a\,c}} \tag{26.58}$$

where

$$\begin{aligned} a &= g' \\ b &= 2\,\cos\theta_1 \\ c &= g'\,r_{1s}^2 + 2\,r_{1s}\,\sin\theta_1 \end{aligned} \tag{26.59}$$

where $g' = 1/\rho'$. The rectangular length after variation L_r' is then

$$L_r' = L_r + r_{2x}' * \sin\theta_1 \tag{26.60}$$

where L_r is the rectangular length before variation. Finally, the length L' after variation is

$$L' = \text{asinc}\,(g'\,L_r')\;L_r' \tag{26.61}$$

where asinc is the function

$$\text{asinc}(\theta) = \frac{\sin^{-1}(\theta)}{\theta} \tag{26.62}$$

For `fiducial_pt` set to `entrance_end` and with `angle`-variation, α' is know and g' can be computed via

$$g' = \frac{\sin(\alpha' - \theta_1) + \sin(\theta_1)}{r_{1s}} \tag{26.63}$$

With this, all other parameters can be created. In both `angle`-variation and `g`-variation the new face angle e_2' is given by

$$e_2' = e_2 + \alpha' - \alpha \tag{26.64}$$

For `fiducial_pt` set to `center`, The center point $\mathbf{r}_c$ (see Fig. 26.3b) is held constant. Here the `g`-variation analysis is similar to the `g`-variation analysis with `fiducial_pt` set to `entrance_end` (or `exit_end` except in this case the reference orbit to the right and left of bfr_c are analyzed separately and the two lengths for each piece are added together. For `angle`-variation, the only situation where it is possible to keep $\mathbf{r}_c$ fixed while varying the angle is when `e1` and `e2` are equal. In this instance, the calculation is again similar to the `angle`-variation analysis with the `fiducial_pt` set to either end. If `e1` and `e2` are not equal, a calculation is done that gives the desired angle but the center point will shift.

26.9 Converter Tracking

Tracking through a `converter` element involves generating five random numbers [3] and then using these numbers with the outgoing particle distribution to generate the position and orientation of the outgoing particle. The outgoing particle distribution is pre-computed by a program `converter_element_modeling` and the distribution parameters are included in the converter element description in the *Bmad* lattice file (§4.8). The accuracy of the converter modeling will depend in part upon the granularity of the probability tables generated by the `converter_element_modeling` program and on other approximations made during tracking. Generally, inaccuracies in the 1% to 10% range are to be expected.

In a tracking simulation, a single outgoing particle is generated for each incoming particle. Since, in a real machine, the number of outgoing particles will not be equal to the number of incoming particles, each outgoing particle is assigned a weight such that the weighted distribution of outgoing particles is correct. This weight will be the same for all outgoing particles. The weight will depend upon whether the momentum or angular range of the outgoing particles is restricted using the element parameters (§4.8).

[3] Since the outgoing particle starts at the exit surface of the converter only five numbers are needed to generate the 6-dimensional particle phase space position.

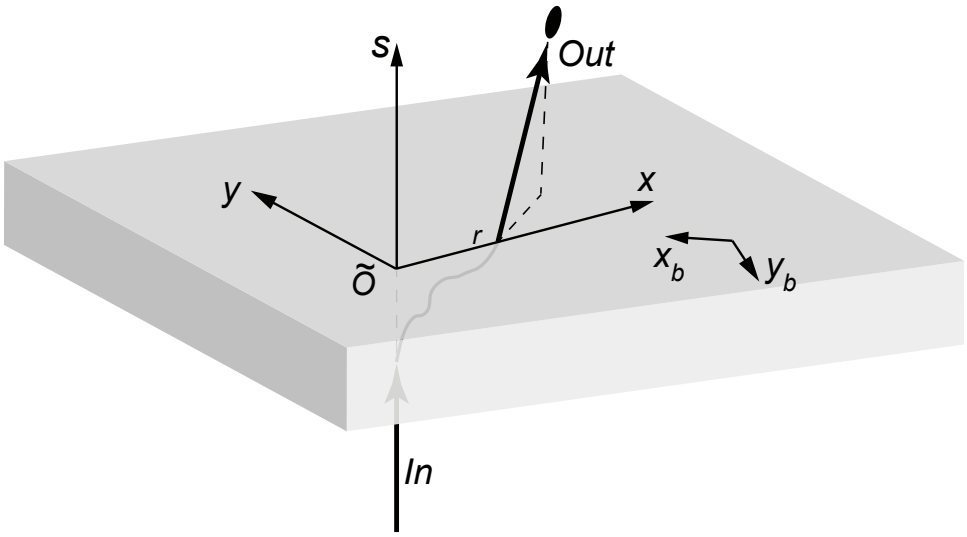

Figure 26.4: An incoming particle strikes the bottom of the converter. At some point within the interior, a new particle is generated and this new particle exits the top surface. To calculate the position and orientation of the outgoing particle, a coordinate system is established where the origin point $\widetilde{\mathcal{O}}$ is the point that the incoming particle would strike the top surface if it went straight through and the (x, y) axes are randomly rotated with respect to the (x_b, y_b) body coordinate axes. By construction, the position of the outgoing particle will be along the x-axis.

.

```
pc_out_min    ! Minimum momentum of generated outgoing particles (eV).
pc_out_max    ! Maximum momentum of generated outgoing particles (eV).
angle_out_max ! Maximum angle to the surface perpendicular (rad).
```

The geometry of the converter is shown in Fig. 26.4. On the top surface, where the outgoing particle emerges, (x_b, y_b) are the axes for the element body coordinate system (§16). To generate the position and orientation of the outgoing particle, another coordinate system is used with axes labeled by (x, y). Each outgoing particle will be assigned its own (x, y) axes. The origin $\widetilde{\mathcal{O}}$ of this coordinate system is constructed by placing $\widetilde{\mathcal{O}}$ at the point where the incoming particle under consideration would strike the top surface if the incoming particle would pass straight through the converter. The angular orientation of the (x, y) axes with respect to the (x, y) axes is chosen using a random number with a uniform probability distribution in the interval $[0, \pi]$. By construction, the outgoing particle, at the surface of the converter will be generated at a point a distance r along the x-axis.

The particle distribution is calculated at a number of converter thickness t_i, $i = 1 \ldots N_t$. It is an error if the actual converter thickness is outside the range of these thicknesses. [The exception is if only one distribution for a given thickness is present, this distribution is used to generate the outgoing particle coordinates independent of the converter thickness.] The particle distribution is also calculated within a certain incoming particle momentum range. It is also an error if an incoming particle has a momentum outside of this range.

The first step is to choose the value of the outgoing particle's momentum p_{out}. For each thickness t_i, the pre-computed particle distribution parameters includes a two-dimensional table of $P(p_{\text{out}}, r)$ — the probability density of creating an outgoing particle versus p_{out} and r. $P(p_{\text{out}}, r)$ is normalized so that the integrated probability is equal to the average number of outgoing particles created for each incoming

particle $N_{\rm out}/N_{in}$

$$\frac{N_{\rm out}}{N_{in}} = \int\int dp_{\rm out}\, dr\, P(p_{\rm out}, r) \tag{26.65}$$

The integrals are done using linear interpolation between grid points. From a $P(p_{\rm out}, r)$ probability table, a "normalized" probability $P_{\rm n}(p_{\rm out}, r)$ table is computed where $P_{\rm n}$ is the probability of generating a particle at given $p_{\rm out}$ and r with an angular range restricted by `angle_out_max`. If `angle_out_max` is not set, $P_{\rm n}$ will be equal to P. This calculation is part of a "setup" computation done before tracking which, to save time, is only done if one of the three element parameters, `pc_out_min`, `pc_out_max`, or `angle_out_max`, changes. Additionally, the setup includes creating a table of $I(p_{\rm out})$ which is the integrated probability for generating a particle with momentum less than $p_{\rm out}$

$$I_p(p_{\rm out}) = \frac{\displaystyle\int_{p_{\rm min}}^{p_{\rm out}} d\widetilde{p}_{\rm out} \int dr\, P_{\rm n}(\widetilde{p}_{\rm out}, r)}{\displaystyle\int_{p_{\rm min}}^{p_{\rm max}} d\widetilde{p}_{\rm out} \int dr\, P_n(\widetilde{p}_{\rm out}, r)} \tag{26.66}$$

where $p_{\rm min}$ is the minimum momentum in the $P(p_{\rm out}, r)$ table or the value of `pc_out_min` which ever is greatest and $p_{\rm max}$ is the maximum momentum in the $P(p_{\rm out}, r)$ table or the value of `pc_out_max` which ever is smallest. $I(p_{\rm out})$ is normalized such that $I(p_{\rm max}) = 1$. A value for $p_{\rm out}$ is generated by solving numerically for $p_{\rm out}$ the equation

$$I_p(p_{\rm out}) = R_1 \tag{26.67}$$

where R_1 is a random number with uniform distribution in the interval $[0, 1]$. This calculation is done for the two t_i thicknesses that straddle the actual thickness. The value of $p_{\rm out}$ assigned to the outgoing particle is obtained via linear interpolation between the two computed values. Note that for both thicknesses the same random number needs to be used.

The next step is to choose a value for r. This is done by solving the equation

$$I_r(r) = R_2 \tag{26.68}$$

where R_2 is another random number with uniform distribution in the interval $[0, 1]$ and I_r is

$$I_r(r) = \frac{\displaystyle\int_0^{r} d\widetilde{r}\, P_{\rm n}(p_{\rm out}, \widetilde{r})}{\displaystyle\int_0^{r_{\rm max}} d\widetilde{r}\, P_n(p_{\rm out}, \widetilde{r})} \tag{26.69}$$

with $p_{\rm out}$ being the momentum chosen for the particle and $r_{\rm max}$ being the maximum radius the probability table goes out to[4]. Like $p_{\rm out}$, this calculation is done for the two t_i thicknesses that straddle the actual thickness. The value of r assigned to the outgoing particle is obtained via linear interpolation between the two computed values. Note that for both thicknesses the same random number needs to be used.

Once $p_{\rm out}$ and r have been chosen, the next steps are to choose values for the angular orientation of the outgoing particle. The angular orientation is characterized by the distribution parameters using the derivatives $x' = dx/ds$ and $y' = dy/ds$ in the form of a skewed Lorentzian probability distribution P_d

$$P_d\,(x', y'; p_{\rm out}, r) = A_d\, \frac{1 + \beta\, x'}{1 + \alpha_x^2\ (x' - c_x)^2 + \alpha_y^2\ (y')^2} \tag{26.70}$$

[4]The range $[0, r_{\rm max}]$ encompasses nearly all of the outgoing particles. In principle, the integral could be extended by extrapolating the values in the table but this could potentially lead to inaccuracies in determining the outgoing orientation. Generally the inaccuracy in truncating the distribution at $r_{\rm max}$ should be small.

where the parameters A_d, β, c_x, α_x, and α_y all depend upon $p_{\rm out}$ and r. Notice that by construction, with the outgoing particle generated on the x-axis, the distribution is symmetric about y'-axis. The pre-computed distribution characterizes each of these parameters by a set of one or more fits which are functions of $p_{\rm out}$ and r. There are also four functions of $p_{\rm out}$ and r that give the range over which Eq. (26.70) is valid $x'_{\rm min}$, $x'_{\rm max}$, $y'_{\rm min}$, and $y'_{\rm max}$. By symmetry, $y'_{\rm min} = -y'_{\rm max}$. Also A_d can be computed from knowledge of β, c_x, α_x, and α_y using the normalization condition that at any given $p_{\rm out}$ and r

$$1 = \int_{x'_{\rm min}}^{x'_{\rm max}} dx' \int_{-y'_{\rm max}}^{y'_{\rm max}} dy'\, P_d\,(x', y') \tag{26.71}$$

Thus there are only seven independent parameters that need to be fitted. The fit functions for all seven have the same form. The fit is divided into two regions. For $p_{\rm out}$ lower than some cutoff, a parameter is fit using a set of one-dimensional functions $\Gamma_i(r)$ at discrete momentum p_i, $i = 1, \ldots, N_\beta$ with

$$\Gamma_i(r) = \sum_{n=1}^{M} c_{n,i}\, r^n \tag{26.72}$$

The polynomial cutoff M is 4 for c_x and β and is 3 for the other five. To evaluate a parameter at momenta lower than p_{N_β}, the Γ_i are used with linear interpolation in p between functions of different p_i. At higher energies, the parameter variation is smoother so a two dimensional fit Xi is used

$$\Xi(p_{\rm out}, r) = e^{-(k_p\, p_{\rm out} + k_r\, r)} \left(\sum_{n=0}^{3} k_n\, r^n \right) \left(1 + \sum_{n=1}^{3} w_n\, p_{\rm out}^n \right) + C \tag{26.73}$$

The C parameter is only nonzero for $x'_{\rm min}$.

Once A_d, β, c_x, α_x, and α_y have been calculated for a given $p_{\rm out}$ and r, The calculation of x' starts with integrating P_d in Eq. (26.70) over y'

$$\begin{aligned} I_{xd}(x') &\equiv \int_{-y'_{\rm lim}}^{y'_{\rm lim}} dy'\, P_d(x', y') \\ &= 2\, A_d\, \frac{1 + \beta\, x'}{\alpha_y \sqrt{1 + \alpha_x^2\,(x' - c_x)^2}}\, \tan^{-1}\left(\frac{\alpha_y\, y'_{\rm lim}}{\sqrt{1 + \alpha_x^2\,(x' - c_x)^2}} \right) \end{aligned} \tag{26.74}$$

where $y'_{\rm lim}$ is either the lesser of $y'_{\rm max}$ and $\tan^{-1}$(angle_out_max). A spline fit is used to integrate I_{xd} and this is used to choose a value for x'. Once x' is known, The integral of $P_d(x', y')$ over y' is used to choose a value for y'.

Except for the placement of $\widetilde{\mathcal{O}}$, the above algorithm for calculating the position and orientation of the outgoing particle will be independent of the angular orientation of the incoming particle. This is valid for incoming particles that are traveling perpendicular to the converter surface. To the extent that the incoming particles are not perpendicular to the converter, this will introduce inaccuracies. Typically, however, the incoming particles will be fairly close to being perpendicular. Considering this, and considering the approximations used to calculate the distribution parameters, the neglect of incoming particle orientation effects is usually justified.

26.10 Drift Tracking

Bmad uses the exact map for a drift This gives the map

$$\begin{aligned} x_2 &= x_1 + \frac{L\,p_{x1}}{p_l} \\ p_{x2} &= p_{x1} \\ y_2 &= y_1 + \frac{L\,p_{y1}}{p_l} \\ p_{y2} &= p_{y1} \\ z_2 &= z_1 + \left(\frac{\beta}{\beta_{\rm ref}} - \frac{1+p_{z1}}{p_l}\right) L \\ p_{z2} &= p_{z1} \end{aligned} \tag{26.75}$$

where β is the normalized particle velocity, $\beta_{\rm ref}$ is the reference particle's normalized velocity, and p_l is the normalized longitudinal momentum

$$p_l \equiv \frac{P_l}{P_0} = \sqrt{(1+p_z)^2 - p_x^2 - p_y^2} \tag{26.76}$$

with P_l being the longitudinal momentum.

26.11 ElSeparator Tracking

[Thanks to Étienne Forest for the derivation of the elseparator equation of motion.]

The Hamiltonian for an electric separator is

$$H = -p_s = -\left\{\left(\frac{1}{\beta_0} + \delta + k_E\,x\right)^2 - \widetilde{m}^2 - p_x^2 - p_y^2\right\}^{1/2} \tag{26.77}$$

Here the canonical coordinates $(-c\,t, \delta$ are being used, $\widetilde{m}$ is defined in Eq. (26.7), and $p_s = -H$ is just the longitudinal momentum. In the above equation, k_E is the normalized field

$$k_E = \frac{q\,E}{P_0\,c} \tag{26.78}$$

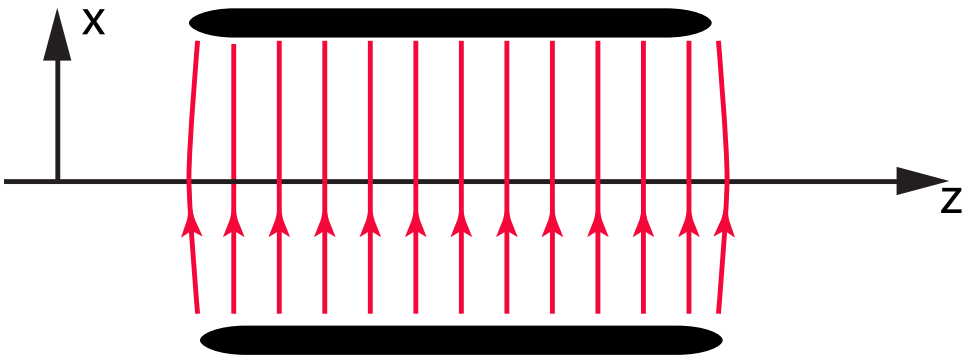

Figure 26.5: Elseparator Electric field. The fringe field lines break the translational invariance in x.

.

The field is taken to be pointing along the x-axis with positive k_E accelerating a particle in the positive x direction. To solve the equations of motion, a "hard edge" model is used where k_E is constant inside the separator and the field ends abruptly at the separator edges.

Since, as shown in Fig. 26.5, the fringe fields break the translational invariance in x, it is important here that the $x = 0$ plane be centered within the separator plates. With this, the canonical momentum δ just outside the separator assumes its free space form of $\delta = (E - E_0)/E_0$). This is analogous to the case of a `solenoid` where, to ensure that the canonical transverse momenta assume their free space form just outside the solenoid, the `z`-axis must be along the centerline of the solenoid.

The solution of the equations of motion is:

$$\begin{aligned}
x &= (x_0 - x_c)\,\cosh\left(\frac{k_E\,L}{p_s}\right) + \frac{p_{x0}}{k_E}\,\sinh\left(\frac{k_E\,L}{p_s}\right) + x_c \\
p_x &= k_E\,(x_0 - x_c)\,\sinh\left(\frac{k_E\,L}{p_s}\right) + p_{x0}\,\cosh\left(\frac{k_E\,L}{p_s}\right) \\
y &= y_0 + L\,\frac{p_{y0}}{p_s} \\
p_y &= p_{y0} \\
c\,\delta t &= \int_0^L -\frac{\partial H}{\partial \delta} = (x_0 - x_c)\,\sinh\left(\frac{k_E\,L}{p_s}\right) + \frac{p_{x0}}{k_E}\,\left[\cosh\left(\frac{k_E\,L}{p_s}\right) - 1\right]
\end{aligned} \tag{26.79}$$

where the critical position x_c is

$$x_c = -\frac{\widetilde{E}}{k_E} \tag{26.80}$$

and

$$\widetilde{E} \equiv \frac{1}{\beta_0} + \delta = \frac{E}{P_0\,c} \tag{26.81}$$

Eqs. (26.79) predict that for $x < x_c$ and $p_{x0} = 0$ a particle will, unphysically, accelerate in the negative x direction. In actuality, a particle in this instance will be reflected backwards by the longitudinal component of the edge field. Specifically, the argument of the square root in Eq. (26.77) must be non-negative and a particle will only make it through the separator if

$$x_0 > \frac{1}{k_E}\,\left(\sqrt{\widetilde{m}^2 + p_{x0}^2 + p_{y0}^2} - \widetilde{E}\right) \tag{26.82}$$

26.12 Foil Tracking

A particle going through a `foil` element is scattered both in angle and in energy, and the charge of the particle may be affected. The following two subsections give the formulas used for scattering and energy loss. Currently, the final charge is a fixed number but that may change in the future.

26.12.1 Scattering in a Foil

For the angle scattering, the user can select between one of two algorithms, both of which are given in the paper by Peralta and Louro[Peralta12] (also see Lynch and Dahl[Lynch90]) Both methods vary the phase space p_x and p_y coordinates using:

$$(dp_x, dp_y) = \frac{p\,\sigma}{P_0}\,(r_1, r_2) \tag{26.83}$$

where p is the particle momentum, P_0 is the reference momentum, r_1 and r_2 are Gaussian random numbers with unit sigma and zero mean, and σ is the sigma of the angular scattering distribution. The factor of p/P_0 is due to a translation between change in angle and change in phase space momenta (see Eq. (17.8)).

The `Highland` algorithm uses Eq. (32) of Peralta and Louro[Peralta12] to calculate the scattering sigma:

$$\sigma = \frac{(13.6 \cdot 10^6\ eV)\, z}{p\, c\, \beta} \sqrt{\frac{X}{X_0}} \left[1 + 0.038\, \ln\left(\frac{X\, z^2}{X_0\, \beta^2}\right)\right] \tag{26.84}$$

where X_0 is the material radiation "length" in kg/m^2, z is the particle charge, β is the particle relativistic beta, c is the speed of light, and X is the foil area density in kg/m^2 equal to $\rho\, t$ where ρ is the material density and t is the foil thickness.

The `Lynch_Dahl` algorithm uses Eq. (33) of Peralta and Louro:

$$\sigma^2 = \frac{\chi_c^2}{1 + F^2} \left[\frac{1 + \nu}{\nu} \ln(1 + \nu) - 1\right] \tag{26.85}$$

where

$$\begin{aligned}
\nu &= \frac{0.5\, \Omega}{(1 - F)} \\
\Omega &= \frac{\chi_c^2}{1.167\, \chi_\alpha^2} \\
\chi_c^2 &= \left(1.57 \cdot 10^{10}\, \frac{eV^2\, m^2}{kg}\right) \frac{Z(Z+1)X}{A} \left[\frac{z}{p\, \beta}\right]^2 \\
\chi_\alpha &= (2.007 \cdot 10^7\, eV^2)\, \frac{Z^{2/3}}{(p\, c)^2} \left[1 + 3.34 \left(\frac{Z\, z\, \alpha}{\beta}\right)^2\right]
\end{aligned} \tag{26.86}$$

and the A is the atomic weight, p is the particle momentum, α is the fine structure constant, and F is a fit parameter representing the percent of the central angular distribution that is used. F is a settable parameter with a default value of 0.98.

For compound materials, the value of X/X_0 in Eq. (26.84) is computed from

$$\frac{X}{X_0} = \sum_{i=1}^{N} \frac{X_i}{X_{0i}} \tag{26.87}$$

where the summation is over all constituents in the material.

Also for compound materials, χ_c^2 in Eqs. (26.85) and (26.86) is replaced by the sum of the constituent χ_{ci}^2, and χ_α is computed from Lynch and Dahl Eq. (11)

$$\ln(\chi_\alpha) = \sum_{i=1}^{N} \frac{Z_i(Z_i + 1)X_i}{A_i} \ln(\chi_{\alpha i}) \Bigg/ \sum_{i=1}^{N} \frac{Z_i(Z_i + 1)X_i}{A_i} \tag{26.88}$$

The actual scattering distribution has $1/\theta^4$ tails (θ is the scattering angle) due to single event large angle scattering (Rutherford scattering). By assuming a Gaussian distribution, these tails are not present in a simulation. It is also important to note that with both the `Highland` and `Lynch_Dahl` algorithms, simulating the passage of particles through a single foil versus two foils with half the thickness as the single foil will not give exactly the same results. This is just a reflection that both algorithms are trying to model an inherently non-Gaussian process.

26.12.2 Energy Loss in a Foil

The particle energy loss per unit length dE/dx through a foil is calculated using the `Bethe-Bloch` formula

$$-\left\langle \frac{dE}{dx} \right\rangle = \frac{4\pi}{m_e c^2} \cdot \frac{n z^2}{\beta^2} \cdot \left(\frac{e^2}{4\pi\varepsilon_0} \right)^2 \cdot \left[\ln\left(\frac{2 m_e c^2 \beta^2}{I \cdot (1-\beta^2)} \right) - \beta^2 \right] \tag{26.89}$$

where n is the material electron density, I is the mean excitation energy, z is the particle charge, c is the speed of light, ϵ_0 is the vacuum permittivity, $\beta = v/c$, is the normalized velocity, and e and m_e the electron charge and rest mass respectively.

Note that to keep the direction of travel of the particle constant when energy is lost, this implies that $p_x/(1+p_z)$ and $p_y/(1+p_z)$ are to be held constant (Eq. (17.8)).

26.13 Kicker, Hkicker, and Vkicker Tracking

The Hamiltonian for a horizontally deflecting kicker or separator is

$$H = \frac{p_x^2 + p_y^2}{2(1+p_z)} - k_0\, x \tag{26.90}$$

This gives the map

$$\begin{aligned}
x_2 &= x_1 + \frac{1}{1+p_{z1}} \left(L\, p_{x1} + \frac{1}{2} k_0\, L^2 \right), & p_{x2} &= p_{x1} + k_0\, L, \\
y_2 &= y_1 + \frac{L\, p_{y1}}{1+p_{z1}}, & p_{y2} &= p_{y1}, \\
z_2 &= z_1 - \frac{L}{2(1+p_{z1})^2} \left(p_{x1}^2 + p_{y1}^2 + p_{x1}\, k_0\, L + \frac{1}{3} k_0^2\, L^2 \right), & p_{z2} &= p_{z1}
\end{aligned} \tag{26.91}$$

The generalization when the kick is not in the horizontal plane is easily derived.

26.14 LCavity Tracking

26.14.1 Integrating through RF fields

When track by integrating through RF fields with `field_calc` set to `bmad_standard` and `tracking_method` set to something like `runge_kutta`, the fields are modeled by the equations given in Sections §18.8 and §19.9. The `bmad_standard` fields, with a phase velocity equal to the speed of light, are appropriate for high energy particles traveling near the speed of light. However, for lower energy particles, the actual fields in a cavity are likely to have been designed with a lower phase velocity and this will cause a mismatch between simulation and reality.

26.14.2 kick-drift-kick model

With `bmad_standard` tracking, a kick-drift-kick model is used where the particle trajectory is modeled such that the reference particle has a constant RF phase at the kick points even at low velocities.[5] The

[5] Developing a 6D symplectic model which continuous acceleration and constant RF phase for the reference particle proved too messy mainly due to the fact that at low energies it was found to be cumbersome to calculate the appropriate RF phase at any given s-position.

cavity is divided up into equal width sections with the number of sections N_s equal to the setting of the `n_rf_steps` parameter. Within a section particles are tracked as if in a field free region (or with a solenoid component if the element's solenoid strength is non-zero). Energy kicks are applied at the ends of the element and in between sections with the kicks at the ends being half that of the kicks applied in the interior. The voltage (integrated field) at a kick point is:

$$V_{\rm k} = \frac{r_q\,\kappa\,V_{tot}}{N_s}\,\cos(2\pi\,(\phi_t + \phi_{\rm ref})) \tag{26.92}$$

See Eq. (4.22) for a description of the parameters in the equation. Here κ is 0.5 for the end kick points and is unity for interior kick points.

The reference energy in each section is adjusted to match the energy of the reference particle (which has $\phi_t = 0$ by definition). The reference RF phase $\phi_{\rm ref}(n)$ at a n^{th} kick point is obtained from the reference RF phase $\phi_{\rm ref}(0)$ at the beginning of the element with a phase offset due to the transit time from the beginning to the kick point of the reference particle:

$$\phi_{\rm ref}(n) = \phi_{\rm ref}(0) + f_{\rm rf}\sum_{j=1}^{n} t_{0j} \tag{26.93}$$

where $t_{0j} = L/c\beta_{0j}\,N_s$ is the transit time of the reference particle across the j^{th} section with L being the element length, and β_{0j} the velocity of the reference particle in the j^{th} section. With this construction, a particle on the zero orbit will see a constant RF phase.

Other than affecting the reference time at each kick point, the adjusting of the reference energy in each section will affect the plotting of phase space momentum coordinates since these are normalized to P_0 (§17.2).

For spin tracking through the energy kick, the BMT equation is used with the integrated electric field given by Eq. (26.92) modeled as a field in the z-direction.

26.14.3 Edge kick

Cavities have a first order transverse kick in the fringe field regions at both the entrance and exit ends due to an electric radial field

Following the model of Rosenzweig and Serafini[Rosen94] (See their Eq. (12)), to make the fringe kick symplectic, the following Hamiltonian is used

$$H_f = \mp\frac{q}{2\,P_0\,c}G_t\,\cos(\phi_t + \phi_{\rm ref})\,(x^2 + y^2) \tag{26.94}$$

where the minus sign is for the entrance end and the plus sign is for the exit end, q is the particle charge, P_0 is the reference momentum. Since $\phi_t = 2\pi f_{\rm rf}\,t$ is time dependent, it is easiest to use this Hamiltonian in energy-time phase space coordinates (§17.5) to give the fringe kicks:

$$\begin{aligned}
\Delta p_x &= \frac{\partial H_f}{\partial x} = \mp\frac{q}{P_0\,c}G_t\,\cos(\phi_t + \phi_{\rm ref})\,x \\
\Delta p_y &= \frac{\partial H_f}{\partial y} = \mp\frac{q}{P_0\,c}G_t\,\cos(\phi_t + \phi_{\rm ref})\,y \\
\Delta p_E &= \frac{\partial H_f}{\partial \tau} = \mp\frac{\pi\,f_{\rm rf}\,q}{P_0\,c^2}G_t\,\sin(\phi_t + \phi_{\rm ref})\,(x^2 + y^2)
\end{aligned} \tag{26.95}$$

For spin tracking through the edge field, the BMT equation is used with the integrated edge field (voltage) $\overline{\mathbf{E}}_f$ parallel to the radius vector:

$$\overline{\mathbf{E}}_f = -G_t \cos(\phi_t + \phi_{\rm ref}) (x, y, 0) \tag{26.96}$$

26.14.4 Standing wave focusing: the pondermotive force

A `LCavity`'s `cavity_type` can be set to either `traveling_wave` or `standing_wave`. The traveling wave cavity is modeled as single forward propagating wave while with a the standing wave cavity, *Bmad* models the cavity as the superposition of a single forward propagating wave and a single backward propagating wave. In both cases there are no harmonics. In both cases, it is the forward traveling wave that gives both acceleration and the edge kick. The difference between the two cavity types is that with a standing wave cavity, there is an extra transverse focusing force called the "pondermotive force".

The time averaged transverse radial kick is modeled using R&S Eq (4)[6] where, since there are no harmonics, $\eta(\phi)$ is unity (see discussion after R&S Eq. (13)). In the *Bmad* kick-drift-kick model, the pondermotive force is included by adding a transverse kick at the kick points. To keep the pondermotive kick symplectic, A Hamiltonian H_p is used:

$$H_p = -\frac{\kappa\, L\, q^2\, G_t^2}{16\, N_s\, P_0^2\, c^2\, (1 + p_z)} (x^2 + y^2) \tag{26.97}$$

Which gives kicks (using standard *Bmad* phase space coordinates):

$$\begin{aligned}
\Delta p_x &= \frac{\partial H_p}{\partial x} = -\frac{\kappa\, L\, q^2\, G_t^2}{8\, N_s\, P_0^2\, c^2\, (1 + p_z)}\, x \\
\Delta p_y &= \frac{\partial H_p}{\partial y} = -\frac{\kappa\, L\, q^2\, G_t^2}{8\, N_s\, P_0^2\, c^2\, (1 + p_z)}\, y \\
\Delta z &= -\frac{\partial H_p}{\partial p_z} = -\frac{\kappa\, L\, q^2\, G_t^2}{16\, N_s\, P_0^2\, c^2\, (1 + p_z)^2} (x^2 + y^2)
\end{aligned} \tag{26.98}$$

The kick is always focusing and is independent of any RF phase. The reason for the phase independence is due to the fact that the phase of the backwards propagating wave with respect to the forward propagating particle is rapidly varying. The formulas above are for the kick averaged over all phases and therefore the phase dependence has been eliminated. Also note that the kick is quadratic in its dependence upon G_t. This is similar to the transverse focusing in FODO cells where the averaged kick is focusing and depends quadratically on the quadrupole field strength.

The time shift (change in z) can be thought of as due to the fact that the actual particle trajectory (and not the averaged trajectory) has a transverse sinusodial component synced to the phase of the backwards traveling wave. The sinusodial component is needed to produce the pondermotive force but this oscillation also delays the particle in time compared to the case where there is no oscillation component.

For spin tracking, to zeroth order, the averaged field is zero for the backwards wave so no spin kick is applied for the backwards wave.

26.15 Magnus Spin Tracking

The `Magnus` spin tracking method uses an analytical Magnus expansion of the T-BMT equation to track spins. Currently, the supported lattice elements are bends (including bends with $k_1 \neq 0$), quadrupoles,

[6] Notice that R&S Eq (4) is missing a minus sign.

solenoids, and sextupoles. Many other elements are treated as drifts §6.3.

One way to obtain approximate spin-transfer maps is to truncate the flow of the T-BMT equation to a given order in the entrance phase-space coordinates. The SPRINT method §26.24 uses this idea at first order. The resulting map will generally not be a rotation, e.g., if quaternions are used, the map will not have unit norm. The idea behind the Magnus expansion is to find an approximate map which is also a rotation. Since the exact solution must be an element of the Lie group SU(2), it can be written as $\exp(F)$, where F belongs to the Lie algebra of SU(2). The Magnus expansion provides a power expansion of F which belongs to the Lie algebra of SU(2) at every order [Magnus54].

In SU(2), the T-BMT equation is given by Eq. (24.99). This can be written in the standard form $\Psi'(s) = A(s)\Psi(s)$. The Magnus expansion gives the second-order solution as $\Psi(L) = \exp(F_1 + F_2)\Psi(0)$, where

$$F_1 = \int_0^L A(s)\,ds \quad \text{and} \quad F_2 = \frac{1}{2}\int_0^L ds_1 \int_0^{s_1} ds_2\,[A(s_1), A(s_2)]. \tag{26.99}$$

Once the trajectory through a magnet is known, $A(s)$ is known as well. We then expand $A(s)$ as a power series in the entrance coordinates which change through the magnet (i.e, all orders of p_z are retained in magnetic elements and all orders of p_y are retained in a pure bend). The truncation to first- and second-order will be written as $A^{(1)}(s)$ and $A^{(2)}(s)$, respectively. The truncation of the exponent F to second-order is then

$$F^{(2)} = \int_0^L A^{(2)}(s)\,ds + \frac{1}{2}\int_0^L ds_1 \int_0^{s_1} ds_2\,[A^{(1)}(s_1), A^{(1)}(s_2)]. \tag{26.100}$$

The spin-transfer map is found by exponentiating $F^{(2)}$, which yields an element of SU(2) by nature of the Magnus expansion.

26.16 Octupole Tracking

The Hamiltonian for an upright octupole is

$$H = \frac{p_x^2 + p_y^2}{2(1+p_z)} + \frac{k_3}{24}(x^4 - 6\,x^2\,y^2 + y^4) \tag{26.101}$$

An octupole is modeled using a kick-drift-kick model.

26.17 Patch Tracking

The transformation of the reference coordinates through a "standard" patch (a patch where custom fields are not used) is given by Eqs. (16.5) and (16.6). At the entrance end of the patch, a particle's position and momentum in the entrance coordinate system will be

$$\begin{aligned} \mathbf{r} &= (x, y, 0) \\ \mathbf{P} &= (P_x, P_y, P_z) = \left(p_x, p_y, \pm\sqrt{(1+p_z)^2 - p_x^2 - p_y^2}\right)\, P_{0\text{ent}} \end{aligned} \tag{26.102}$$

where p_x, p_y and p_z are the phase space momenta, and z, which is coordinate z and not phase space z, is always zero by construction as shown in Fig. 26.6 [Also see Fig. 16.2 and the discussion in §17.2.] The sign of the longitudinal momentum P_z is determined by whether the particle is traveling in the positive s or negative s direction (which will occur when an element is flipped longitudinally).

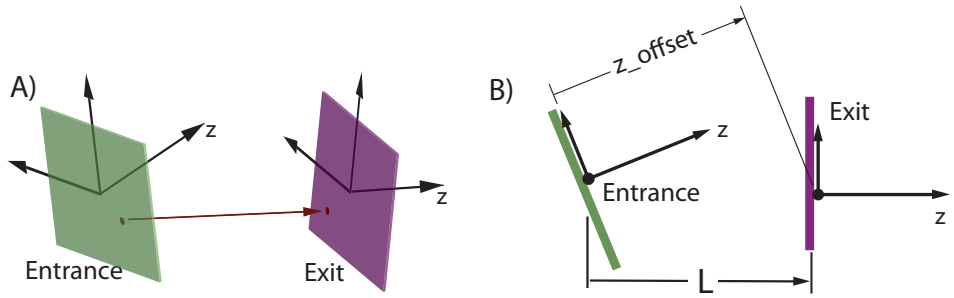

Figure 26.6: Standard tracking through a patch element. A particle's starting coordinate at the entrance end of the patch has, by construction, coordinate $z = 0$. The particle is drifted, as in a field free region, between the entrance $z = 0$ plane and the exit $z = 0$ plane.

The transformation between entrance and exit coordinate systems is given by Eqs. (16.12) and (16.13)

$$\begin{aligned} \mathbf{r} &\to \mathbf{S}^{-1}\,(\mathbf{r} - \mathbf{L}_{\rm off}) \\ \mathbf{P} &\to \mathbf{S}^{-1}\,\mathbf{P} \end{aligned} \tag{26.103}$$

where $\mathbf{L}_{\rm off}$ is given by Eq. (16.16)

After this transformation, the particle must be propagated by a longitudinal length $-r_z$ to intersect the $r_z = 0$ plane of the exit face.

$$\begin{aligned} \mathbf{r} &\to (r_x - r_z\,\frac{P_x}{P_z}, r_y - r_z\,\frac{P_y}{P_z}, 0) \\ \mathbf{P} &\to \mathbf{P} \end{aligned} \tag{26.104}$$

The final $\mathbf{r}$ and $\mathbf{P}$ can now be used compute the particles phase space coordinates, along with the time t and the reference time $t_{\rm ref}$ at the exit end.

$$\begin{aligned} x &\to r_x & p_x &\to \frac{P_x}{P_{0\rm exi}} \\ y &\to r_y & p_y &\to \frac{P_y}{P_{0\rm exi}} \\ z &\to z + r_z\,\frac{|\mathbf{P}|}{P_z} + L_0\,\frac{\beta}{\beta_0} + \beta\,{\rm t_offset} & p_z &\to \frac{(1+p_z)\,P_{0\rm ent} - P_{0\rm exi}}{P_{0\rm exi}} \\ t &\to t - r_z\,\frac{|\mathbf{P}|}{P_z\,\beta} & t_{\rm ref} &\to t_{\rm ref} + {\rm t_offset} + L_0\,\frac{1}{\beta_0} \end{aligned} \tag{26.105}$$

where the exit reference momentum $P_{0\rm exi}$ is related to the entrance reference momentum $P_{0\rm ent}$ through `e_tot_offset`. In the above equation, β is the particle velocity, β_0 is the velocity of the reference particle, and L_0 is the drift length of the reference particle

$$L_0 = \frac{1}{S^{-1}_{33}}\,\left(S^{-1}_{31}\,{\rm x_offset} + S^{-1}_{32}\,{\rm y_offset} + S^{-1}_{33}\,{\rm z_offset}\right) \tag{26.106}$$

26.18 Quadrupole Tracking

The `bmad_standard` calculates the transfer map through an upright quadrupole and then transforms that map to the laboratory frame.

The Hamiltonian for an upright quadrupole is

$$H = \frac{p_x^2 + p_y^2}{2(1+p_z)} + \frac{k_1}{2}(x^2 - y^2) \tag{26.107}$$

This is simply solved

$$\begin{aligned}
x_2 &= c_x\, x_1 + s_x\, \frac{p_{x1}}{1+p_{z1}} \\
p_{x2} &= \tau_x\, \omega^2\ (1+p_{z1})\, s_x\, x_1 + c_x\, p_{x1} \\
y_2 &= c_y\, y_1 + s_y\, \frac{p_{y1}}{1+p_{z1}} \\
p_{y2} &= \tau_y\, \omega^2\ (1+p_{z1})\, s_y\, y_1 + c_y\, p_{y1} \\
z_2 &= z_1 + m_{511}\, x_1^2 + m_{512}\, x_1\, p_{x1} + m_{522}\, p_{x1}^2 + m_{533}\, y_1^2 + m_{534}\, y_1\, p_{y1} + m_{544}\, p_{y1}^2 \\
p_{z2} &= p_{z1}
\end{aligned} \tag{26.108}$$

where

$$\omega \equiv \sqrt{\frac{|k_1|}{1+p_{z1}}} \tag{26.109}$$

and

$$\begin{array}{lll|lll}
 & k_1 > 0 & k_1 < 0 & & k_1 > 0 & k_1 < 0 \\
c_x = & \cos(\omega\, L) & \cosh(\omega\, L) & c_y = & \cosh(\omega\, L) & \cos(\omega\, L) \\
s_x = & \dfrac{\sin(\omega\, L)}{\omega} & \dfrac{\sinh(\omega\, L)}{\omega} & s_y = & \dfrac{\sinh(\omega\, L)}{\omega} & \dfrac{\sin(\omega\, L)}{\omega} \\
\tau_x = & -1 & +1 & \tau_y = & +1 & -1
\end{array} \tag{26.110}$$

with this

$$\begin{aligned}
m_{511} &= \frac{\tau_x\ \omega^2}{4}\,(L - c_x\, s_x) & m_{533} &= \frac{\tau_y\ \omega^2}{4}\,(L - c_y\, s_y) \\
m_{512} &= \frac{-\tau_x\ \omega^2}{2\,(1+p_{z1})}\, s_x^2 & m_{534} &= \frac{-\tau_y\ \omega^2}{2\,(1+p_{z1})}\, s_y^2 \\
m_{522} &= \frac{-1}{4\,(1+p_{z1})^2}\,(L + c_x\, s_x) & m_{544} &= \frac{-1}{4\,(1+p_{z1})^2}\,(L + c_y\, s_y)
\end{aligned} \tag{26.111}$$

26.19 RFcavity Tracking

For tracking using something like `runge_kutta`, with `field_calc` set to `bmad_standard`, the fields are modeled by the equations given in Sections §18.8 and §19.9.

With `bmad_standard` tracking, a kick-drift-kick model is used. The kick is a pure energy kick (see equations in §4.46) and the phase of the RF is calculated under the assumption that the waveform moves at a phase velocity equal to the velocity of the reference particle.

With `bmad_standard` tracking, the transverse forces due to the RF are ignored. This is generally a reasonable approximation when the acceleration is small as is standard in rings. `Lcavity` elements should be used in place of `rfcavity` elements when this is not so.

26.20 Sad_Mult Tracking

The "hard edge" fringe field kick is taken from Forest[Forest98] Eqs. (13.29) and onward. In the notation of *Bmad*, and taking into account both normal and skew terms, Eq. (13.29) is for the $m$þorder multipole (what Forest labels $n+1$)

$$f_{\pm} = \mp\Re\frac{(b_m + i\,a_m)\,(x+i\,y)^{(m+1)}}{4\,(m+2)\,(1+p_z)}\left[x\,p_x + y\,p_y + i\frac{m+3}{m+1}(x\,p_y - y\,p_x)\right] \tag{26.112}$$

The "soft edge" dipole fringe for `sad_mult` elements is a generalization of the soft edge dipole fringe for a SAD bend element. For the entrance kick the equations are:

$$\begin{aligned}
x_2 &= x_1 + \frac{\delta_1}{1+\delta_1}\,\Delta x_{fx}, \qquad p_{x2} = p_{x1} + \frac{1}{1+\delta_1}\,\left[\Delta x_{fy}\,v - \Delta x_{fay}\,v^3\right] \\
y_2 &= y_1 - \frac{\delta_1}{1+\delta_1}\,\Delta y_{fy}, \qquad p_{y2} = p_{y1} + \frac{1}{1+\delta_1}\,\left[\Delta y_{fx}\,w - \Delta y_{fax}\,w^3\right] \\
z_2 &= z_1 + \frac{1}{(1+\delta_1)^2}\left[\Delta x_{fx}\,p_{x1} - \Delta y_{fy}\,p_{y1} + \frac{1}{2}\,(\Delta y_{fx} + \Delta x_{fy})\,w^2 - \frac{1}{4}(\Delta y_{fax} + \Delta x_{fay})\,w^4\right]
\end{aligned} \tag{26.113}$$

where

$$\begin{aligned}
\Delta x_{fx} &= \frac{K_0\,F_B^2}{24\,L}, & \Delta y_{fx} &= \frac{K_0^2\,F_B}{6\,L^2}, & \Delta y_{fax} &= \frac{2\,K_0^2}{3\,F_B\,L^2}, \\
\Delta y_{fy} &= \frac{SK_0\,F_B^2}{24\,L}, & \Delta x_{fy} &= \frac{SK_0^2\,F_B}{6\,L^2}, & \Delta x_{fay} &= \frac{2\,SK_0^2}{3\,F_B\,L^2}, \\
v &= \cos\theta\,x_1 + \sin\theta\,y_1, & w &= -\sin\theta\,x_1 + \cos\theta\,y_1, & \tan\theta &= \frac{-SK_0}{K_0}
\end{aligned} \tag{26.114}$$

26.21 Sextupole Tracking

The Hamiltonian for an upright sextupole is

$$H = \frac{p_x^2 + p_y^2}{2(1+p_z)} + \frac{k_2}{6}(x^3 - 3\,x\,y^2) \tag{26.115}$$

Tracking through a sextupole uses a kick-drift-kick model.

26.22 Sol_Quad Tracking

The Hamiltonian is

$$H = \frac{(p_x + \frac{k_s}{2}\,y)^2}{2(1+p_z)} + \frac{(p_y - \frac{k_s}{2}\,x)^2}{2(1+p_z)} + \frac{k_1}{2}(x^2 - y^2) \tag{26.116}$$

Solving the equations of motion gives

$$\begin{aligned} x_2 &= m_{11}\, x_1 + m_{12}\, p_{x1} + m_{13}\, y_1 + m_{14}\, p_{y1} \\ p_{x2} &= m_{21}\, x_1 + m_{22}\, p_{x1} + m_{23}\, y_1 + m_{24}\, p_{y1} \\ y_2 &= m_{31}\, x_1 + m_{32}\, p_{x1} + m_{33}\, y_1 + m_{34}\, p_{y1} \\ p_{y2} &= m_{41}\, x_1 + m_{42}\, p_{x1} + m_{43}\, y_1 + m_{44}\, p_{y1} \\ z_2 &= z_1 + \sum_{j=1}^{4} \sum_{k=j}^{4} m_{5jk}\, r_j\, r_k \\ p_{z2} &= p_{z1} \end{aligned} \tag{26.117}$$

where

$$\begin{aligned} m_{11} &= \frac{1}{2\,f}\,(f_{0+}\,c + f_{0-}\,c_h) & m_{31} &= -m_{24} \\ m_{12} &= \frac{1}{2\,f\,(1+p_{z1})}\left(\frac{f_{++}}{\omega_+}\,s + \frac{f_{--}}{\omega_-}\,s_h\right) & m_{32} &= -m_{14} \\ m_{13} &= \frac{\widetilde{k}_s}{4\,f}\left(\frac{f_{+-}}{\omega_+}\,s + \frac{f_{-+}}{\omega_-}\,s_h\right) & m_{33} &= \frac{1}{2\,f}\,(f_{0-}\,c + f_{0+}\,c_h) \\ m_{14} &= \frac{\widetilde{k}_s}{f\,(1+p_{z1})}\,(-c + c_h) & m_{34} &= \frac{1}{2\,f\,(1+p_{z1})}\left(\frac{f_{+-}}{\omega_+}\,s + \frac{f_{-+}}{\omega_-}\,s_h\right) \\ m_{21} &= \frac{-(1+p_{z1})}{8\,f}\left(\frac{\xi_{1+}}{\omega_+}\,s + \frac{\xi_{2+}}{\omega_-}s_h\right) & m_{41} &= -m_{23} \\ m_{22} &= m_{11} & m_{42} &= -m_{13} \\ m_{23} &= \frac{\widetilde{k}_s^3\,(1+p_{z1})}{4\,f}\,(c - c_h) & m_{43} &= \frac{-(1+p_{z1})}{8\,f}\left(\frac{\xi_{1-}}{\omega_+}\,s + \frac{\xi_{2-}}{\omega_-}\,s_h\right) \\ m_{24} &= \frac{\widetilde{k}_s}{4\,f}\left(\frac{f_{++}}{\omega_+}\,s + \frac{f_{--}}{\omega_-}\,s_h\right) & m_{44} &= m_{33} \end{aligned} \tag{26.118}$$

and

$$\begin{aligned} \widetilde{k}_1 &= \frac{k_1}{1+p_{z1}} & \widetilde{k}_s &= \frac{k_s}{1+p_{z1}} \\ f &= \sqrt{\widetilde{k}_s^4 + 4\,\widetilde{k}_1^2} & f_{\pm 0} &= f \pm \widetilde{k}_s^2 \\ f_{0\pm} &= f \pm 2\,\widetilde{k}_1 & f_{\pm\pm} &= f \pm \widetilde{k}_s^2 \pm 2\,\widetilde{k}_1 \\ \omega_+ &= \sqrt{\frac{f_{+0}}{2}} & \omega_- &= \sqrt{\frac{f_{-0}}{2}} \\ s &= \sin(\omega_+\,L) & s_h &= \sinh(\omega_-\,L) \\ c &= \cos(\omega_+\,L) & c_h &= \cosh(\omega_-\,L) \\ \xi_{1\pm} &= \widetilde{k}_s^2\,f_{+\mp} \pm 4\,\widetilde{k}_1\,f_{+\pm} & \xi_{2\pm} &= \widetilde{k}_s^2\,f_{-\pm} \pm 4\,\widetilde{k}_1\,f_{-\mp} \end{aligned} \tag{26.119}$$

The m_{5jk} terms are obtained via Eq. (26.12)

$$\begin{aligned} m_{5jk} = -\frac{\tau_{jk}}{2(1+p_{z1})^2}\int ds\,\Big[&\left(m_{2j} + \frac{k_s}{2}\,m_{3j}\right)\left(m_{2k} + \frac{k_s}{2}\,m_{3k}\right) + \\ &\left(m_{4j} - \frac{k_s}{2}\,m_{1j}\right)\left(m_{4k} - \frac{k_s}{2}\,m_{1k}\right)\Big] \end{aligned} \tag{26.120}$$

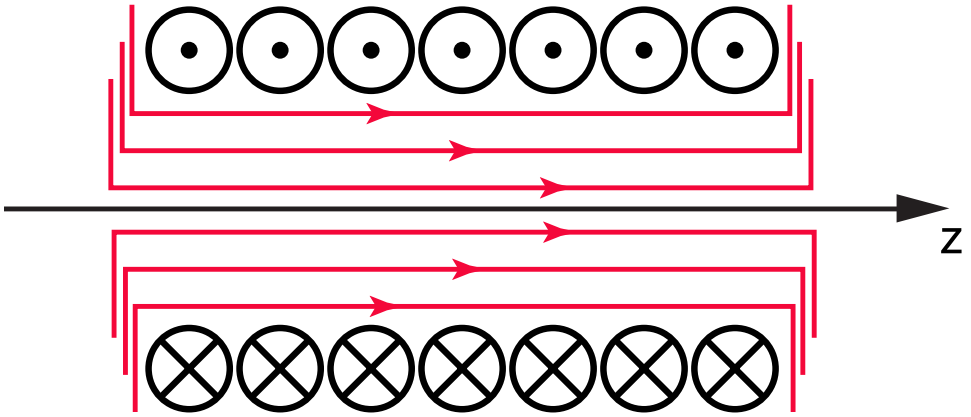

Figure 26.7: Solenoid with a hard edge. The field is assumed to end abruptly at the edges of the solenoid. Here, for purposes of illustration, the field lines at the ends are displaced from one another.

.

where

$$\tau_{jk} = \begin{cases} 1 & j = k \\ 2 & j \neq k \end{cases} \tag{26.121}$$

The needed integrals involve the product of two trigonometric or hyperbolic functions. These integrals are trivial to do but the explicit equations for m_{5jk} are quite long and in the interests of brevity are not reproduced here.

26.23 Solenoid Tracking

The `bmad_standard` solenoid tracking does not make the small angle approximation. The transfer map for the solenoid is:

$$\begin{aligned}
x_2 &= \frac{1+c}{2}\, x_1 + \frac{s}{k_s}\, p_{x1} + \frac{s}{2}\, y_1 + \frac{1-c}{k_s}\, p_{y1} \\
p_{x2} &= \frac{-k_s\, s}{4}\, x_1 + \frac{1+c}{2}\, p_{x1} - \frac{k_s\,(1-c)}{4}\, y_1 + \frac{s}{2}\, p_{y1} \\
y_2 &= \frac{-s}{2}\, x_1 - \frac{1-c}{k_s}\, p_{x1} + \frac{1+c}{2}\, y_1 + \frac{s}{k_s}\, p_{y1} \\
p_{y2} &= \frac{k_s\,(1-c)}{4}\, x_1 + \frac{-s}{2}\, p_{x1} - \frac{k_s\, s}{4}\, y_1 + \frac{1+c}{2}\, p_{y1} \\
z_2 &= z_1 + L\left(\frac{\beta}{\beta_0} - \frac{1+p_z}{p_r}\right) \\
p_{z2} &= p_{z1}
\end{aligned} \tag{26.122}$$

where $k_s = q\,B/P_0$ is the normalized field and

$$\begin{aligned}
c &= \cos\left(k_s L/p_r\right) \\
s &= \sin\left(k_s L/p_r\right)
\end{aligned} \tag{26.123}$$

with

$$p_r = \sqrt{(1+p_z)^2 - (p_x + y_1\, k_s/2)^2 - (p_y - x_1\, k_s/2)^2} \tag{26.124}$$

To be useful, the canonical momenta p_x and p_y in the above equations must be connected to the canonical momenta used for other elements (drifts, quadrupoles, etc.) that may be placed to either side of the solenoid. These side elements use zero a_x and a_y (cf. Eq. (26.8)). The vector potential used in the solenoid canonical momenta may be made zero at the edges of the solenoid if the solenoid fringe field is assumed to end abruptly at the edges of the solenoid (as shown in Fig. 26.7), and the reference axis **z**-axis (at $x = y = 0$) is placed along the centerline of the solenoid so that there is cylindrical symmetry around the **z**-axis.

26.24 Sprint Spin Tracking

The `sprint` spin tracking method is named after the `SPRINT` program developed by Matthias Vogt. The `sprint` algorithm Uses a first order spin map evaluated with respect to the zero orbit to track through elements. This method is much faster than PTC integration, and its run-time does not increase proportionally to element length. Currently, the supported lattice elements are bends (including bends with $k_1 \neq 0$), quadrupoles, and solenoids §6.3.

Elements with fringe field contributions are split into three quaternions representing the entrance, body, and exit of the element. Before propagation, the exit fringe quaternion is always equivalent to the entrance fringe quaternion, with all field strengths multiplied by -1, and e_1 replaced with $-e_2$. The exit quaternion is then propagated to the end of the element via Bmad mapping tools. Appropriate quaternions are concatenated according to the values of `spin_fringe_on` and `fringe_at`.

$$\begin{array}{llll} d = g\,l & e = a\,g\,l\gamma & s = a\,k_s l & t = (1+a)k_s l \\ c_d = \cos(d) & s_{e2} = \sin(\frac{e}{2}) & c_s = \cos(s) & c_t = \cos(t) \\ s_d = \sin(d) & c_{e2} = \cos(\frac{e}{2}) & s_s = \sin(s) & s_{t2} = \sin(\frac{t}{2}) \\ \chi = 1 + a\gamma & \zeta = \gamma - 1 & \psi = \gamma^2 - 1 & c_{t2} = \cos(\frac{t}{2}) \end{array} \tag{26.125}$$

26.24.1 SBend Body, $k_1 = 0$

	q_0	q_y	q_z
1	c_{e2}	$-s_{e2}$	
x	$-\frac{1}{2}g\chi s_d s_{e2}$	$-\frac{1}{2}g\chi\, s_d c_{e2}$	
p_x	$\frac{1}{2}\chi\,(c_d - 1)s_{e2}$	$\frac{1}{2}\chi\,(c_d - 1)c_{e2}$	
p_y			$\frac{1}{\gamma}\zeta\, s_{e2}$
p_z	$\frac{1}{2\gamma}\,(\gamma\,\chi\,s_d - a\,\psi\,d)\,s_{e2}$		

26.24.2 Sbend Body, $k_1 \neq 0$

$$
\begin{aligned}
k_x &= k_1 + g^2 & \alpha &= 2(a^2g^2\gamma^2 + k_1) \\
\omega_x &= \sqrt{|k_x|} & \beta &= agk_1(\gamma\chi - \zeta) \\
\omega_y &= \sqrt{|k_1|} & \sigma &= \omega_y(k_1 + ak_1\gamma + a^2g^2\zeta\gamma) \\
& & \xi &= \omega_y(k_1\chi + a^2g^2\zeta\gamma)
\end{aligned}
$$

$k_x > 0$	$k_x < 0$	$k_1 > 0$	$k_1 < 0$
$s_x = \sin(l\omega_x)$	$\sinh(l\omega_x)$	$s_y = \sinh(l\omega_y)$	$\sin(l\omega_y)$
$c_x = \cos(l\omega_x)$	$\cosh(l\omega_x)$	$c_y = \cosh(l\omega_y)$	$\cos(l\omega_y)$
$\tau_x = -1$	$+1$	$\tau_y = +1$	-1

	q_0	q_x	q_y	q_z
1	c_{e2}		$-s_{e2}$	
x	$\frac{-k_x\chi}{2\omega_x} s_x s_{e2}$		$\frac{-k_x\chi}{2\omega_x} s_x c_{e2}$	
p_x	$\frac{k_x\chi}{2\omega_x^2}\tau_x(1-c_x)s_{e2}$		$\frac{k_x\chi}{2\omega_x^2}\tau_x\,(1-c_x)\,c_{e2}$	
y		$\frac{-1}{\alpha}\big[\beta(1+c_y)s_{e2} + \tau_y\sigma s_y c_{e2}\big]$		$\frac{1}{\alpha}\big[\beta(1-c_y)c_{e2} + \tau_y\sigma s_y s_{e2}\big]$
p_y		$\frac{1}{\omega_y\alpha}\big[\xi(1-c_y)c_{e2} - \beta s_y s_{e2}\big]$		$\frac{1}{\omega_y\alpha}\big[\xi(1+c_y)s_{e2} - \beta s_y c_{e2}\big]$
p_z	$\frac{g}{2}\left(\frac{\chi s_x}{\omega_x} - \frac{al\psi}{\gamma}\right)s_{e2}$		$\frac{g}{2}\left(\frac{\chi s_x}{\omega_x} - \frac{al\psi}{\gamma}\right)c_{e2}$	

26.24.3 Sbend Entrance Fringe

To calculate the exit fringe, multiply all field strengths g by -1, and replace all entrance face angles e_1 with exit face angles $-e_2$. The negative exit face angle is used due to Bmad convention.

	q_0	q_x	q_y	q_z
1	1			
x			$\frac{1}{2}\chi g\tan(e_1)$	
y		$\frac{1}{2}(1+a)g\sin(e_1)$		$-\frac{1}{2}(1+a)g\cos(e_1)$

26.24.4 Quadrupole

$$\omega = \sqrt{|k_1|}$$

$k_1 > 0$	$k_1 < 0$	$k_1 > 0$	$k_1 < 0$
$s_x = \frac{\sin(l\omega)}{\omega}$	$\frac{\sinh(l\omega)}{\omega}$	$s_y = \frac{\sinh(l\omega)}{\omega}$	$\frac{\sin(l\omega)}{\omega}$
$c_x = \frac{1-\cos(l\omega)}{\omega^2}$	$\frac{-1+\cosh(l\omega)}{\omega^2}$	$c_y = \frac{-1+\cosh(l\omega)}{\omega^2}$	$\frac{1-\cos(l\omega)}{\omega^2}$

	q_0	q_x	q_y
1	1		
x			$-\frac{1}{2}k_1\chi s_x$
p_x			$-\frac{1}{2}k_1\chi c_x$
y		$-\frac{1}{2}k_1\chi s_y$	
p_y		$-\frac{1}{2}k_1\chi c_y$	

26.24.5 Solenoid Element Body

	q_0	q_x	q_y	q_z
1	c_{t2}			$-s_{t2}$
x		$\frac{1}{4}k_s\zeta((1-c_s)c_{t2}-s_s s_{t2})$	$\frac{1}{4}k_s\zeta((-1+c_s)s_{t2}-s_s c_{t2})$	
p_x		$\frac{1}{2}\zeta((1-c_s)st2+s_s c_{t2})$	$\frac{1}{2}\zeta((1-c_s)c_{t2}-s_s s_{t2})$	
y		$\frac{1}{4}k_s\zeta((1-c_s)s_{t2}+s_s c_{t2})$	$\frac{1}{4}k_s\zeta((1-c_s)c_{t2}-s_s s_{t2})$	
p_y		$\frac{1}{2}\zeta((-1+c_s)c_{t2}+s_s s_{t2})$	$\frac{1}{2}\zeta((1-c_s)s_{t2}+s_s c_{t2})$	
p_z	$\frac{1}{2}ts_{t2}$			$\frac{1}{2}tc_{t2}$

26.24.6 Solenoid Entrance Fringe

To calculate the exit fringe, multiply all field strengths k_s by -1.

	q_0	q_x	q_y
1	1		
x		$\frac{1}{4}k_s\chi$	
y			$\frac{1}{4}k_s\chi$

26.25 Symplectic Tracking with Cartesian Modes

The method for symplectic integration for elements that define the magnetic field using a Cartesian mode decomposition (§5.16.2) is outlined in §26.4. The vector potential is constructed to avoid singularities when one of the wave vectors k_x, k_y, or k_z is zero.

For the `x family` the vector potential is:

Form	hyper-y	hyper-xy	hyper-x
A_x	$A\,\frac{k_z}{k_y^2}\,\mathrm{S}_x\,\mathrm{Sh}_y\,\mathrm{S}_z$	$A\,\frac{1}{k_y}\,\mathrm{Sh}_x\,\mathrm{Sh}_y\,\mathrm{S}_z$	$A\,\frac{k_z}{k_x\,k_y}\,\mathrm{Sh}_x\,\mathrm{S}_y\,\mathrm{S}_z$
A_y	0	0	0
A_z	$A\,\frac{k_x}{k_y^2}\,\mathrm{C}_x\,\mathrm{Sh}_y\,\mathrm{C}_z$	$A\,\frac{k_x}{k_y\,k_z}\,\mathrm{Ch}_x\,\mathrm{Sh}_y\,\mathrm{C}_z$	$A\,\frac{1}{k_y}\,\mathrm{Ch}_x\,\mathrm{S}_y\,\mathrm{C}_z$

For the `y family` the vector potential is:

Form	hyper-y	hyper-xy	hyper-x
A_x	0	0	0
A_y	$-A\,\dfrac{k_z}{k_x\,k_y}\,\mathrm{S}_x\,\mathrm{Sh}_y\,\mathrm{S}_z$	$-A\,\dfrac{1}{k_x}\,\mathrm{Sh}_x\,\mathrm{Sh}_y\,\mathrm{S}_z$	$-A\,\dfrac{k_z}{k_x^2}\,\mathrm{Sh}_x\,\mathrm{S}_y\,\mathrm{S}_z$
A_z	$-A\,\dfrac{1}{k_x}\,\mathrm{S}_x\,\mathrm{Ch}_y\,\mathrm{C}_z$	$-A\,\dfrac{k_y}{k_x\,k_z}\,\mathrm{Sh}_x\,\mathrm{Ch}_y\,\mathrm{C}_z$	$-A\,\dfrac{k_y}{k_x^2}\,\mathrm{Sh}_x\,\mathrm{C}_y\,\mathrm{C}_z$

For the `qu family` the vector potential is:

Form	hyper-y	hyper-xy	hyper-x
A_x	$A\,\dfrac{1}{k_z}\,\mathrm{S}_x\,\mathrm{Ch}_y\,\mathrm{S}_z$	$A\,\dfrac{k_y}{k_z^2}\,\mathrm{Sh}_x\,\mathrm{Ch}_y\,\mathrm{S}_z$	$A\,\dfrac{k_y}{k_x\,k_z}\,\mathrm{Sh}_x\,\mathrm{C}_y\,\mathrm{S}_z$
A_y	$-A\,\dfrac{k_x}{k_y\,k_z}\,\mathrm{C}_x\,\mathrm{Sh}_y\,\mathrm{S}_z$	$-A\,\dfrac{k_x}{k_z^2}\,\mathrm{Ch}_x\,\mathrm{Sh}_y\,\mathrm{S}_z$	$-A\,\dfrac{1}{k_z}\,\mathrm{Ch}_x\,\mathrm{S}_y\,\mathrm{S}_z$
A_z	0	0	0

For the `sq family` the vector potential is:

Form	hyper-y	hyper-xy	hyper-x
A_x	$A\,\dfrac{1}{k_z}\,\mathrm{C}_x\,\mathrm{Sh}_y\,\mathrm{S}_z$	$A\,\dfrac{k_y}{k_z^2}\,\mathrm{Ch}_x\,\mathrm{Sh}_y\,\mathrm{S}_z$	$A\,\dfrac{k_y}{k_x\,k_z}\,\mathrm{Ch}_x\,\mathrm{S}_y\,\mathrm{S}_z$
A_y	$A\,\dfrac{k_x}{k_y\,k_z}\,\mathrm{S}_x\,\mathrm{Ch}_y\,\mathrm{S}_z$	$-A\,\dfrac{k_x}{k_z^2}\,\mathrm{Sh}_x\,\mathrm{Ch}_y\,\mathrm{S}_z$	$A\,\dfrac{1}{k_z}\,\mathrm{Sh}_x\,\mathrm{C}_y\,\mathrm{S}_z$
A_z	0	0	0

Chapter 27

Tracking of X-Rays

Bmad can track both charged particles and X-rays. This chapter deals with X-rays. Charged particles are handled in chapter §26.

27.1 Coherent and Incoherent Photon Simulations

Bmad can track photons either `coherently` or `incoherently`. In both cases, the photon has a transverse electric field

$$(E_x, E_y) \tag{27.1}$$

E_x and E_y are complex and therefore have both amplitude and phase information. When photons are tracked incoherently, the phase information is not used for calculating X-ray intensities.

In addition to coherent and incoherent tracking, partially coherent simulations can be done by using sets of photons with the photons in any one set treated as coherent and the photons between sets being treated as incoherent.

27.1.1 Incoherent Photon Tracking

In a simulation with incoherent photons, some number of photons, N_0, will be generated and the ith photon ($i = 1, \ldots, N_0$) will have a initial "electric field" components $E_{x0}(i), E_{y0}(i)$ assigned to it. The field amplitude E_0 will be $\sqrt{E_{x0}^2 + E_{y0}^2}$.

At some an observation point, the power S per unit area falling on some small area dA due to either x or y component of the electric field is

$$S_{x,y} = \frac{\alpha_p}{N_0\, dA} \sum_{j \in \text{hits}} E_{x,y}^2(j) \tag{27.2}$$

where α_p is a constant that can be chosen to fit the simulation against experimental results, and the sum is over photons who intersect the area. The factors of N_0 and dA in the above equation make, within statistical fluctuations, S independent of N_0 and, for dA small enough, S will be independent of dA as it should be. The total power is just $S_x + S_y$.

When traveling through vacuum, the electric field of a photon is a constant. As an example, consider a point source radiating uniformly in 4π solid angle with each photon having the same initial field E_0. An

observation area dA situated a distance R from the source will intercept $N_0\,dA/4\,\pi\,R^2$ photons which gives a power of

$$S_w = \frac{\alpha_p\,E_0^2}{4\,\pi\,R^2} \tag{27.3}$$

which falls off as $1/R^2$ as expected.

At some places the light may be split into various "channels". An example is Laue diffraction where X-rays can excite the α and β branches of the dispersion surface. Or a partially silvered mirror where some of the light is reflected and some is transmitted. In such a case, the probability P_i of a photon traveling down the i^{th} channel is

$$P_i\,\widehat{E}_i^2 = \frac{S_i}{S_0} \tag{27.4}$$

where S_i is the power flowing into channel i, S_0 is the power flowing into the junction, and $\widehat{E}_i = E_i/E_0$ is the ratio of the electric field amplitudes of any photon just before and just after being shunted into the i^{th} channel. The probabilities must be properly normalized

$$\sum P_i = 1 \tag{27.5}$$

If the ratio of the electric field of any photon just before and just after being shunted into the i^{th} channel is not a constant, than $\widehat{E}_i$ must be adjusted so that $\widehat{E}_i^2$ is equal to the average of $\widehat{E}_i^2(j)$ for all photons j channeled into channel i.

As long as Eqs. (27.4) and (27.5) are satisfied, the choice of the P_i, and $\widehat{E}_i$ are arbitrary. This freedom allows simulation to be optimized for efficiency. For example, In an actual experiment much of the light can be lost never to reach a detector and be counted. To decrease the simulation time, simulated photons may be limited to be generated with a direction to be within some solid angle Ω_1 if photons with a direction outside this solid angle will not contribute to the simulation results. In this case, there are two channels. Channel 1 consists of all photons whose direction is within Ω_1 and channel 2 is all the other photons. To limit the photons to channel 1, P_1 is taken to be 1 and P_2 is taken to be 0. Additionally, if the light, say, is being generated isotropically from a surface into a $\Omega_0 = 2\,\pi$ solid angle then

$$\widehat{E}_1 = \sqrt{\frac{\Omega_1}{\Omega_0}} \tag{27.6}$$

$\widehat{E}_2$ is infinite here but since no photons are generated in channel 2 this is not a problem.

27.1.2 Coherent Photon Tracking

In a simulation with coherent photons, some number of photons, N_0, will be generated and the ith photon $(i = 1, \ldots, N_0)$ will have an initial electric field $E_{x0}(i), E_{y0}(i)$ assigned to it. These quantities will be complex.

At some an observation point, the field E at some small area dA due to either x or y component of the electric field is

$$E = \frac{\alpha_p}{N_0\,dA}\sum_{j\in\text{hits}} E(j) \tag{27.7}$$

where α_p is a constant that can be chosen to fit the simulation against experimental results, and the sum is over photons who intersect the area. In the above equation $E(j)$ is either the x or y component of the electric field as is appropriate. The factors of N_0 and dA in the above equation make, within statistical fluctuations, E independent of N_0 and, for dA small enough, E will be independent of dA as it should be.

When traveling through a a vacuum, the photons travel ballistically in straight lines. This is justified by using the stationary phase approximation with Kirchhoff's integral. the electric field of a photon varies with the propagation length. There is nothing physical in this and is just a way to make the bookkeeping come out correctly. As an example, consider a point source radiating uniformly in 4π solid angle with each photon having the same initial field component (either x or y) E_1. An observation area dA situated a distance R from the source will intercept $N_0\, dA/4\,\pi\, R^2$ photons and each photon will have a field of $E_1\, R\, \exp(i\, k\, R)$ where k is the photon wave number (all photons must have the same k to be coherent). This gives an electric field at the observation point of

$$E = \frac{\alpha_p\, E_1\, \exp(i\, k\, R)}{4\,\pi\, R} \tag{27.8}$$

which falls off as $1/R$ as expected.

At a `diffraction_plate` element where diffraction effects are to be simulated, the following procedure is used:

1. The electric field components are multiplied by the propagation length L:

$$E \to E\, L \tag{27.9}$$

   The propagation length is reset to zero so that the at the next point where the propagation length is factored into the electric field the propagation length will be the length starting at the aperture.

2. Depending upon the program, the photon is is either given a random direction over $2\,\pi$ solid angle or the photon's direction is restricted to be within some solid angle chosen to increase the probability that the photon will make it through some downstream aperture.

   If the photon is restricted to some aperture dependent solid angle of area Ω, the photon's electric field is scaled by

$$E \to E\, \frac{\Omega}{4\,\pi} \tag{27.10}$$

3. The electric field components are scaled by

$$E \to E\, \frac{k}{4\,\pi\, i}\, (\cos\theta_1 + \cos\theta_2) \tag{27.11}$$

   where θ_1 and θ_2 are the direction cosines of the incoming and outgoing directions of the photon with respect to the longitudinal reference axis.

This algorithm is designed so that the resulting fields at points downstream from the aperture as computed from a simulation will, to within statistical errors, be the same as one would get using Kirchoff's integral. That is, the simulation is constructed to be a Monte Carlo integration of Kirchhoff's integral.

What is, and what is not considered a place where there are diffraction effects is dependent upon the problem. For example, there are diffraction effects associated with light reflecting from a mirror (or any other object) of finite size. If these effects are important to the experiment, then a procedure similar to the one above must be followed.

At places where there are no diffraction effects a simulation can treat the photons ballistically or can use the aperture procedure outlined above. While in theory it is possible to choose what to do, in practice the aperture procedure increases the number of photons that must be tracked for a given resolution. Thus, from a practical standpoint the ballistic alternative should always be used.

As explained in §27.1.1, at some places the light may be split into various "channels". With coherent photons, the analog to Eq. (27.4) is

$$P_i \widehat{E}_i = \frac{E_i}{E_0} \tag{27.12}$$

where here $\widehat{E}_i$ can be complex to take into account phase shifts. The same considerations about choosing the P_i and $\widehat{E}_i$ apply to coherent photons as incoherent photons. In particular, $\widehat{E}_1$ for the case of isotropic emission from a surface as in the example in §27.1.1 (cf. Eq. (27.6)) is

$$\widehat{E}_1 = \frac{\Omega_1}{\Omega_0} \tag{27.13}$$

27.1.3 Partially Coherent Photon Simulations

When there is partial coherence the photons must be divided into sets. All of the photons of a given set are considered coherent while the photons of different sets are treated incoherently.

The procedure is to track all the photons of one set coherently and calculate the field using equation Eq. (27.7). The fields of different sets are then combined to calculate a power using Eq. (27.2).

27.2 Element Coordinate System

The general procedure for tracking through an element makes use of `element reference` coordinates (also called just `element` coordinates). Without any offsets, pitches or tilt (§5.6), henceforth called "misalignments", the `element` coordinates are the same as the `laboratory reference` coordinates (or simply `laboratory` coordinates) (§16.1.1). The `element` coordinates stay fixed relative to the element. Therefore, if the element is misaligned, the `element coordinates` will follow as the element shifts in the laboratory frame as shown in Fig. 26.1.

For `crystal` (§4.10), `mirror` (§4.35), and `multilayer_mirror` (§4.37) elements, the "kinked" reference trajectory through the element complicates the calculation. For these elements, there are three coordinate systems attached to the element as shown in Fig. 27.1. Besides the `element entrance` and `element exit` coordinates, there are `element surface` coordinates with z perpendicular to the surface pointing inward.

Tracking a particle through an element is therefore a three step transformation:

1. At the entrance end of the element, transform from the laboratory reference coordinates to the element's `entrance` or `surface` coordinates.
2. Track through the element ignoring any misalignments.
3. At the exit end of the element, transform from the element coordinates to the `laboratory exit` coordinates.

27.2.1 Transform from Laboratory Entrance to Element Coordinates

For elements that have a reference orbit kink (§27.2), the element coordinates here are the `surface` coordinates. Otherwise the element coordinates are the entrance coordinates.

1. Apply offsets, pitches and tilt using the formulas in §16.2.2 along with Eqs. (16.6), and (16.16).

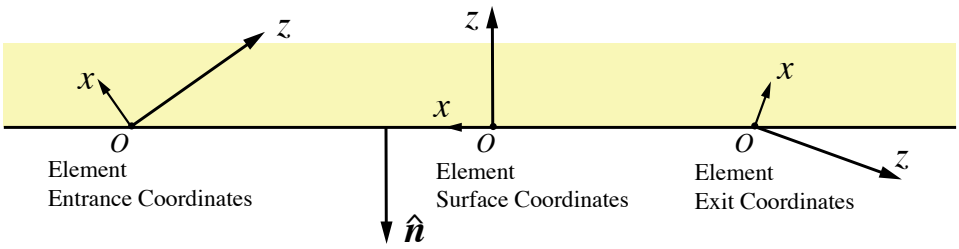

Figure 27.1: The three element coordinate systems for `crystal` (Bragg configuration), `mirror`, and `multilayer_mirror` elements. The origin **O** of all three are the same but are shown spread out for clarity. $\hat{\mathbf{n}}$ is the normal to the element surface.

.

2. Apply the `tilt` to the electric field (Eq. (17.14)).

3. For `crystal`, `mirror`, and `multilayer_mirror` elements rotate to element surface coordinates.

4. Transform the photon's position as if in a drift by a distance $-z$ where z is the photon's longitudinal coordinate. That is, z will be zero at the end of the transform to element coordinates (remember that z is the distance from the start of the element (§17.4)).

27.2.2 Transform from Element Exit to Laboratory Coordinate

The back transformation from element to laboratory coordinates is accomplished by the transformation

1. For `crystal`, `mirror`, and `multilayer_mirror` elements rotate to element from element surface coordinates to element exit coordinates

2. Apply the reverse `tilt` to the electric field (Eq. (17.14)).

3. Apply reverse offsets, pitches and tilt using the formulas in §16.2.2 along with Eqs. (16.6), and (16.16).

27.3 Transformation for Mirror and Crystal Elements Between Laboratory and Element Coordinates

27.3.1 Transformation from Laboratory to Element Coordinates

With photons, the intensities must also be transformed. The transformation from the entrance laboratory coordinates to the entrance element coordinates is:

1. Track as in a drift a distance `z_offset_tot`.

2. Apply offsets and pitches: The effective "length" of the element is zero (§16.2.3) so the origin of

the element coordinates is the same point around which the element is pitched so

$$\begin{aligned} x_1 &= x_0 - x_{\text{off}} \\ p_{x1} &= p_{x0} - (1 + p_{z0})\, x'_{pitch} \\ y_1 &= y_0 - y_{\text{off}} \\ p_{y1} &= p_{x0} - (1 + p_{z0})\, y'_{pitch} \\ z_1 &= z_0 + x'_{pitch}\, x_1 + y'_{pitch}\, y_1 \end{aligned} \tag{27.14}$$

where $x_{\text{off}} \equiv$ `x_offset`, $x'_{pitch} \equiv$ `x_pitch`, etc.

3. Apply `ref_tilt` and `tilt`:

$$\begin{aligned} \begin{pmatrix} x_2 \\ y_2 \end{pmatrix} &= \mathbf{R}(\theta_{tot}) \begin{pmatrix} x_1 \\ y_1 \end{pmatrix} \\ \begin{pmatrix} p_{x2} \\ p_{y2} \end{pmatrix} &= \mathbf{R}(\theta_{tot}) \begin{pmatrix} p_{x1} \\ p_{y1} \end{pmatrix} \\ \begin{pmatrix} \mathbf{E}_{x2} \\ \mathbf{E}_{y2} \end{pmatrix} &= \mathbf{R}(\theta_{tot}) \begin{pmatrix} \mathbf{E}_{x1} \\ \mathbf{E}_{y1} \end{pmatrix} \end{aligned} \tag{27.15}$$

where $\mathbf{E}$ is shorthand notation for

$$\mathbf{E} \equiv E\, e^{i\,\phi} \tag{27.16}$$

with E being the field intensity and ϕ being the field phase angle. In the above equations $\mathbf{R}$ is the rotation matrix

$$\mathbf{R}(\theta) = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \tag{27.17}$$

with θ_{tot} being

$$\theta_{tot} = \begin{cases} \texttt{ref_tilt} + \texttt{tilt} + \texttt{tilt_corr} & \texttt{for crystal elements} \\ \texttt{ref_tilt} + \texttt{tilt} & \texttt{for mirror elements} \end{cases} \tag{27.18}$$

The `tilt_corr` correction is explained in §27.4.2.

27.3.2 Transformation from Element to Laboratory Coordinates

The back transformation from exit element coordinates to exit laboratory coordinates is accomplished by the transformation

1. Apply `ref_tilt` and `tilt`: `ref_tilt` rotates the exit laboratory coordinates with respect to the exit element coordinates in the same way `ref_tilt` rotates the entrance laboratory coordinates with respect to the entrance element coordinates. The forward and back transformations are thus just inverses of each other. With `tilt`, this is not true. `tilt`, unlike `ref_tilt`, does not rotate the output laboratory coordinates. There is the further complication in that `tilt` is a rotation about the *entrance* laboratory coordinates. The first step is to express `tilt` with respect to the exit coordinates. This is done with the help of the $\mathbf{S}$ matrix of Eq. (16.8) with α_t given by Eq. (16.15). The effect of the `tilt` can be modeled as a rotation vector $\mathbf{e}_{in}$ in the entrance laboratory coordinates pointing along the z-axis

$$\mathbf{e}_{in} = (0, 0, \text{tilt}) \tag{27.19}$$

In the exit laboratory coordinates, the vector $\mathbf{e}_{out}$ is

$$\mathbf{e}_{out} = \mathbf{S}\,\mathbf{e}_{in} \tag{27.20}$$

The z component of $\mathbf{e}_{out}$ combines with `ref_tilt` to give the transformation

$$\begin{aligned}
\begin{pmatrix} x_2 \\ y_2 \end{pmatrix} &= \mathbf{R}(-\theta_t) \begin{pmatrix} x_1 \\ y_1 \end{pmatrix} \\
\begin{pmatrix} p_{x2} \\ p_{y2} \end{pmatrix} &= \mathbf{R}(-\theta_t) \begin{pmatrix} p_{x1} \\ p_{y1} \end{pmatrix} \\
\begin{pmatrix} \mathbf{E}_{x2} \\ \mathbf{E}_{y2} \end{pmatrix} &= \mathbf{R}(-\theta_t) \begin{pmatrix} \mathbf{E}_{x1} \\ \mathbf{E}_{y1} \end{pmatrix}
\end{aligned} \tag{27.21}$$

where θ_t is ref_tilt $+ \mathbf{e}_{out,z}$. The x and y components of $\mathbf{e}_{out}$ give rotations around the x and y axes

$$\begin{aligned}
p_{x3} &= p_{x2} - \mathbf{e}_{out,y} \\
p_{y3} &= p_{y2} + \mathbf{e}_{out,x} && (27.22) \\
z_3 &= z_2 + x_2\,\mathbf{e}_{out,y} - y_2\,\mathbf{e}_{out,x} && (27.23)
\end{aligned}$$

2. Apply pitches: Since pitches are defined with respect to the entrance laboratory coordinates, they have to be translated to the exit laboratory coordinates

$$\mathbf{P}_{out} = \mathbf{S}\,\mathbf{P}_{in} \tag{27.24}$$

where $\mathbf{P}_{in} = (x'_{pitch}, y'_{pitch}, 0)$ is the pitch vector in the entrance laboratory frame and $\mathbf{P}_{out}$ is the vector in the exit laboratory frame. The transformation is then

$$\begin{aligned}
p_{x4} &= p_{x3} - \mathbf{P}_{out,y} \\
p_{y4} &= p_{y3} + \mathbf{P}_{out,x} && (27.25) \\
z_4 &= z_3 + x_3\,\mathbf{P}_{out,y} - y_3\,\mathbf{P}_{out,x} && (27.26)
\end{aligned}$$

3. Apply offsets: Again, offsets are defined with respect to the entrance laboratory coordinates. Like pitches, the translation is

$$\mathbf{O}_{out} = \mathbf{S}\,\mathbf{O}_{in} \tag{27.27}$$

where $\mathbf{O}_{in} = (x_{\rm off}, y_{\rm off}, s_{\rm off})$ is the offset in the entrance laboratory frame. The transformation is

$$\begin{aligned}
x_5 &= x_4 + \mathbf{O}_{out,x} - p_{x4}\,\mathbf{O}_{out,z} \\
y_5 &= y_4 + \mathbf{O}_{out,y} - p_{y4}\,\mathbf{O}_{out,z} && (27.28) \\
z_5 &= z_4 + \mathbf{O}_{out,z} && (27.29)
\end{aligned}$$

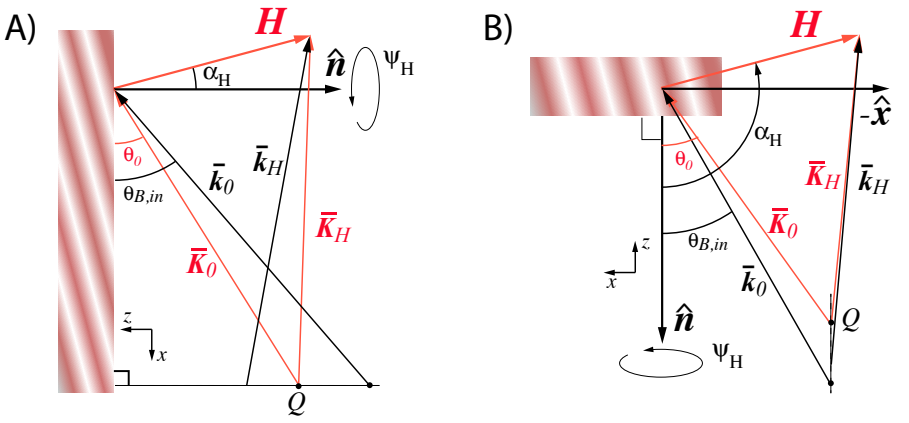

Figure 27.2: Reference trajectory reciprocal space diagram for for A) Bragg diffraction and B) Laue diffraction. The bar over the vectors indicates that they refer to the reference trajectory. The x-z coordinates shown are the element surface coordinates. All points in the diagram are in the plane of the paper except for the tip of $\mathbf{H}$. $\overline{\mathbf{K}}_0$, and $\overline{\mathbf{K}}_H$ are the wave vectors inside the crystal and $\overline{\mathbf{k}}_0$ and $\overline{\mathbf{k}}_H$ are the wave vectors outside the crystal. The reference photon traveling along the reference trajectory has $\overline{\mathbf{K}}_0$ and $\overline{\mathbf{K}}_H$ originating at the Q point. For Laue diffraction, the crystal faces are assumed parallel. For Bragg diffraction the crystal normal is in the $-\hat{\mathbf{x}}$ direction while for Laue diffraction the crystal normal is in the $-\hat{\mathbf{z}}$ direction

27.4 Crystal Element Tracking

[Crystal tracking developed by Jing Yee Chee, Ken Finkelstein, and David Sagan]

Crystal diffraction is modeled using dynamical diffraction theory. The notation here follows Batterman and Cole[Bater64]. The problem can be divided up into two parts. First the reference trajectory must be calculated. This means calculating the incoming grazing angle $\theta_{B,in}$ and outgoing grazing angle $\theta_{B,out}$ as well as calculating the transformations between the various coordinate systems. This is done in §27.4.1, §27.4.2, and §27.4.3. The second part is the actual tracking of the photon and this is covered in §27.4.5 and §27.4.6

27.4.1 Calculation of Entrance and Exit Bragg Angles

Fig. 27.2 shows the geometry of the problem. The bar over the vectors indicates that they refer to the reference trajectory. The reference trajectory is calculated such that the reference photon will be in the center of the Darwin curve. That is, the internal wave vectors $\overline{\mathbf{K}}_0$ and $\overline{\mathbf{K}}_H$ originate from the Q point (See [Bater64] Figs. 8 and 29).

The external wave vectors $\mathbf{k_0}$, and $\mathbf{k_H}$ and the internal wave vectors have magnitude

$$|\mathbf{k_0}| = |\mathbf{k_H}| = \frac{\mathbf{1}}{\lambda} \tag{27.30}$$

$$|\overline{\mathbf{K}}_0| = |\overline{\mathbf{K}}_H| = \frac{1-\delta}{\lambda} \tag{27.31}$$

where λ is the wavelength, and δ is

$$\delta = \frac{\lambda^2 r_e}{2\,\pi\,V}\,F_0' = \frac{\Gamma}{2}\,F_0' = \frac{1}{2}\,\Gamma\,F_0' \tag{27.32}$$

with r_e being the classical electron radius, V the unit cell volume, and F_0' is the real part of the F_0 structure factor.

In element surface coordinates (which will be the coordinate system used henceforth), $\overline{\mathbf{k}}_0$ lies in the x-z plane. $\overline{\mathbf{K}}_0$ is related to $\overline{\mathbf{k}}_0$ via Batterman Eq. (25)

$$\mathbf{K}_0 = \mathbf{k}_0 + q_0\,\widehat{\mathbf{n}} \tag{27.33}$$

where the value of q_0 is to be determined. Here, and in equations below, if the equation is true in general, and not just for the reference trajectory, the bar superscript is dropped.

Since $\widehat{\mathbf{n}}$ is in the $-\widehat{\mathbf{x}}$ direction, $\overline{\mathbf{K}}_0$ is also in the x-z plane. Thus $\overline{\mathbf{k}}_0$ and $\overline{\mathbf{K}}_0$ can be written in the form

$$\begin{aligned}
\overline{\mathbf{k}}_0 &= \frac{1}{\lambda}\begin{pmatrix} -\cos\theta_{B,in} \\ 0 \\ \sin\theta_{B,in}\end{pmatrix}, \qquad \overline{\mathbf{K}}_0 = \frac{1-\delta}{\lambda}\begin{pmatrix} -\cos\theta_0 \\ 0 \\ \sin\theta_0\end{pmatrix} \qquad \text{[Bragg]} \\
\overline{\mathbf{k}}_0 &= \frac{1}{\lambda}\begin{pmatrix} \sin\theta_{B,in} \\ 0 \\ \cos\theta_{B,in}\end{pmatrix}, \qquad \overline{\mathbf{K}}_0 = \frac{1-\delta}{\lambda}\begin{pmatrix} \sin\theta_0 \\ 0 \\ \cos\theta_0\end{pmatrix} \qquad \text{[Laue]}
\end{aligned} \tag{27.34}$$

Where, as shown in Fig. 27.2, $\theta_{B,in}$, and θ_0 are the angles of $\overline{\mathbf{k}}_0$ and $\overline{\mathbf{K}}_0$ with respect to the x-axis for Bragg reflections and with respect to the z-axis for Laue reflection.

α_H (`alpha_angle`) is the angle that $\mathbf{H}$ makes with respect to the $-\widehat{\mathbf{z}}$ axis and ψ_H (`psi_angle`) is the rotation of $\mathbf{H}$ around the $-\widehat{\mathbf{z}}$ axis such that for $\psi_H = 0$, $\mathbf{H}$ is in the x-z plane and oriented as shown in Fig. 27.2. Thus

$$\mathbf{H} \equiv \frac{1}{d}\,\widehat{\mathbf{H}} = \frac{1}{d}\begin{pmatrix} -\sin\alpha_H\,\cos\psi_H \\ \sin\alpha_H\,\sin\psi_H \\ -\cos\alpha_H\end{pmatrix} \tag{27.35}$$

where $\widehat{\mathbf{H}}$ is $\mathbf{H}$ normalized to 1. α_H is determined via the setting of `b_param` and via Eq. (4.17).

The vectors $\mathbf{K}_0$ and $\mathbf{H}$ must add up to the reciprocal lattice vector $\mathbf{K}_H$

$$\mathbf{K}_H = \mathbf{K}_0 + \mathbf{H} \tag{27.36}$$

Taking the length of both sides of this equation and using Eqs. (27.31), (27.34), and (27.35) gives for θ_0

$$\sin\theta_0 = \begin{cases} \dfrac{-\beta\,\widehat{H}_z - \widehat{H}_x\sqrt{\widehat{H}_x^2 + \widehat{H}_z^2 - \beta^2}}{\widehat{H}_x^2 + \widehat{H}_z^2} & \texttt{Bragg} \\[2ex] \dfrac{-\beta\,\widehat{H}_x + \widehat{H}_z\sqrt{\widehat{H}_x^2 + \widehat{H}_z^2 - \beta^2}}{\widehat{H}_x^2 + \widehat{H}_z^2} & \texttt{Laue} \end{cases} \tag{27.37}$$

where

$$\beta \equiv \frac{\lambda}{2\,d\,(1-\delta)} \tag{27.38}$$

Once θ_0 has been calculated, $\theta_{B,in}$ can be calculated from Eq. (27.33)

$$\cos\theta_{B,in} = (1-\delta)\,\cos\theta_0 \quad \text{[Bragg]} \tag{27.39}$$

$$\sin\theta_{B,in} = (1-\delta)\,\sin\theta_0 \quad \text{[Laue]} \tag{27.40}$$

The outgoing reference wave vector k_H is computed using the equation

$$\mathbf{K}_H = \mathbf{k}_H + q_H\,\widehat{\mathbf{n}} \tag{27.41}$$

Using this with Eqs. (27.35) and (27.36) gives

$$\begin{aligned}\overline{k}_{H,x} &= \overline{K}_{H,z} = \frac{1}{d}\,\widehat{H}_x + \overline{k}_{0,x} \\ \overline{k}_{H,y} &= \overline{K}_{H,y} = \frac{1}{d}\,\widehat{H}_y \\ \overline{k}_{H,z} &= \sqrt{\frac{1}{\lambda^2} - \overline{k}_{H,x}^2 - \overline{k}_{H,y}^2}\end{aligned} \tag{27.42}$$

The total bending angle of the reference trajectory is then

$$\theta_{bend} = \tan^{-1}\left(\frac{|\overline{\mathbf{k}}_0 \times \overline{\mathbf{k}}_H|}{\overline{\mathbf{k}}_0 \cdot \overline{\mathbf{k}}_H}\right) \tag{27.43}$$

The outgoing Bragg angle $\theta_{B,out}$ is then *defined* to be the difference between the total bend angle and the entrance Bragg angle.

$$\theta_{B,out} \equiv \theta_{bend} - \theta_{B,in} \tag{27.44}$$

27.4.2 Crystal Coordinate Transformations

There are four transformations needed between coordinates denoted by $\mathbf{\Sigma}_1$, $\mathbf{\Sigma}_2$, $\mathbf{\Sigma}_3$, and $\mathbf{\Sigma}_4$

$\mathbf{\Sigma}_1$ `Transform from laboratory entrance to element entrance coordinates.`
$\mathbf{\Sigma}_2$ `Transform from element entrance to surface coordinates.`
$\mathbf{\Sigma}_3$ `Transform from surface to element exit coordinates.`
$\mathbf{\Sigma}_4$ `Transform from element exit to laboratory exit coordinates.`

The total transformation is just the map represented by $\mathbf{S}$ and $\mathbf{V}$ of Eqs. (16.5) and (16.6)

$$[\mathbf{S}, \mathbf{V}] = \mathbf{\Sigma}_4\,\mathbf{\Sigma}_3\,\mathbf{\Sigma}_2\,\mathbf{\Sigma}_1 \tag{27.45}$$

The transformation $\mathbf{\Sigma}_1$ is given in §27.3.1 and the transformation $\mathbf{\Sigma}_4$ is given in §27.3.2. In general, the transformation $\mathbf{\Sigma}_1$ needs a "tilt correction" (Eq. (27.18)), as explained below, when ψ_H is nonzero. [The exception is when the `undiffracted` or `forward_diffracted` beam is tracked with Laue geometry. In these cases, no tilt correction is needed.] Since this tilt correction is independent of any misalignments, the tilt correction calculation proceeds assuming here that there are no misalignments. The finite $\mathbf{V}$ due to the finite crystal thickness in Laue diffraction will also be ignored for the moment.

Without misalignments, and with ψ_H zero, the transformation $\mathbf{\Sigma}_1$ is, as it is for every other type of element, just the unit matrix.

$$\mathbf{\Sigma}_1 = \mathbf{I} \tag{27.46}$$

That is, the two coordinate systems are identical. Furthermore, the transformation $\mathbf{\Sigma}_2$ from element entrance coordinates to surface coordinates is a rotation around the y axis

$$\begin{aligned}\mathbf{\Sigma}_2 = \mathbf{R}_y(\theta_{B,in}) &\equiv \begin{pmatrix} \cos\theta_{B,in} & 0 & \sin\theta_{B,in} \\ 0 & 1 & 0 \\ -\sin\theta_{B,in} & 0 & \cos\theta_{B,in} \end{pmatrix} && \text{[Laue]} \\ &= \mathbf{R}_y(\theta_{B,in} - \frac{\pi}{2}) && \text{[Bragg]}\end{aligned} \tag{27.47}$$

The transformation from element surface coordinates to element exit coordinates, $\mathbf{\Sigma}_3$, is another rotation around the y axis

$$\begin{aligned}\mathbf{\Sigma}_3 &= \mathbf{R}_y(\theta_{B,out}) && \text{[Laue]} \\ &= \mathbf{R}_y(\theta_{B,out} + \frac{\pi}{2}) && \text{[Bragg]}\end{aligned} \tag{27.48}$$

and the transformation from element exit coordinates to laboratory exit coordinates, $\boldsymbol{\Sigma}_{out}$ is the unity matrix

$$\boldsymbol{\Sigma}_4 = \mathbf{I} \tag{27.49}$$

Thus, the combined transformation $\mathbf{S}$ from laboratory entrance to laboratory exit coordinates is a rotation around the y axis of $\theta_{B,in} + \theta_{B,out}$ as explained in section §16.2

$$\mathbf{S} = \boldsymbol{\Sigma}_4\, \boldsymbol{\Sigma}_3\, \boldsymbol{\Sigma}_2\, \boldsymbol{\Sigma}_1 = \mathbf{R}_y(\theta_{B,in} + \theta_{B,out}) \tag{27.50}$$

When ψ_H is non-zero, the situation is complicated since, if $\mathbf{S}$ as calculated above is used, the vector $\overline{\mathbf{k}}_H$ would be bent out of the x-z plane even though it has been assumed that the `ref_tilt` θ_t is zero. But $\overline{\mathbf{k}}_H$ points in the same direction as the z axis of the outgoing reference trajectory. Furthermore, by *definition*, the reference trajectory has the form given by Eq. (16.8) with the $\mathbf{R}_z(\theta_t)$ matrix depending only upon the `ref_tilt` parameter (which is here taken to be zero). To satisfy Eq. (16.8), the crystal must be reoriented to keep the $\mathbf{k}_H$ vector in the x-z plane of the laboratory entrance coordinates. The reorientation is done by rotating the crystal about the laboratory entrance $\mathbf{z}$ axis by an amount θ_{corr} (`tilt_corr`).

With this tilt correction the transformation $\boldsymbol{\Sigma}_1$ is a rotation about the z axis

$$\boldsymbol{\Sigma}_1 = \begin{pmatrix} \cos\theta_{corr} & -\sin\theta_{corr} & 0 \\ \sin\theta_{corr} & \cos\theta_{corr} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{27.51}$$

To calculate a value for θ_{corr}, note that the transformation $\boldsymbol{\Sigma}_2$ from element entrance coordinates to element surface coordinates is not affected by a finite ψ_H and so Eq. (27.47) is unmodified. The $\mathbf{k}_H$ vector, expressed in laboratory entrance coordinates, is $\boldsymbol{\Sigma}_1^{-1}\, \boldsymbol{\Sigma}_2^{-1}\, \mathbf{k}_H$ where the components of $\mathbf{k}_H$ are given by Eq. (27.42). To satisfy Eq. (16.8), this vector must have zero y component

$$\left(\boldsymbol{\Sigma}_1^{-1}\, \boldsymbol{\Sigma}_2^{-1}\, \mathbf{k}_H\right) \cdot \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} = 0 \tag{27.52}$$

Solving gives

$$\theta_{corr} = \tan^{-1} \frac{k_{H,y}}{k_{H,z}\, \sin\theta_{B,in} - k_{H,x}\, \cos\theta_{B,in}} \tag{27.53}$$

The transformation $\boldsymbol{\Sigma}_3$ from element surface coordinates to element exit coordinates is now obtained by requiring that the total transformation from laboratory entrance to laboratory exit coordinates be the $\mathbf{R}_y(-\alpha_b)$ matrix given in Eq. (16.8)

$$\boldsymbol{\Sigma}_3\, \boldsymbol{\Sigma}_2\, \boldsymbol{\Sigma}_1 = \begin{pmatrix} \cos\theta_{bend} & 0 & -\sin\theta_{bend} \\ 0 & 1 & 0 \\ \sin\theta_{bend} & 0 & \cos\theta_{bend} \end{pmatrix} \tag{27.54}$$

In the above equation, the transformation $\boldsymbol{\Sigma}_4$ has been dropped since it is the unit matrix independent of ψ_H.

For Laue diffraction when the non-diffracted beam is tracked, the exit coordinate system corresponds to the entrance coordinate system. That is, $\mathbf{V}$ is the unit matrix. In this case, there is no tilt correction and $\boldsymbol{\Sigma}_3 = \mathbf{R}_y(-\theta_{B,in})$ is just the inverse of $\boldsymbol{\Sigma}_2$.

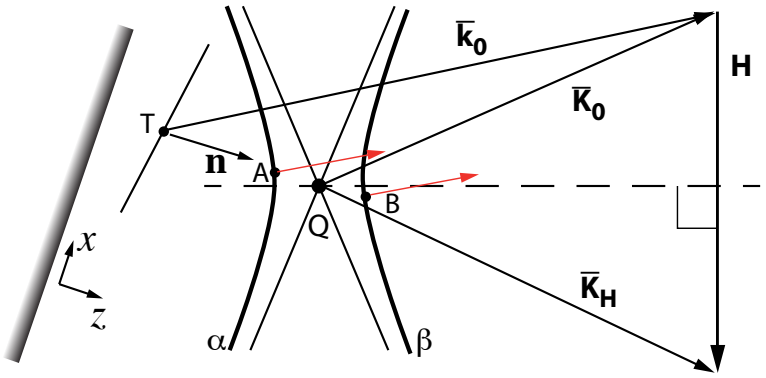

Figure 27.3: Energy flow used to determine the reference orbit for Laue diffraction.

27.4.3 Laue Reference Orbit

For Laue diffraction, with the reference orbit following the `undiffracted` beam, the reference orbit at the exit surface is just the extension of the reference orbit at the entrance surface. Since the reference orbit's direction is $\bar{\mathbf{k}}_0$. the reference orbit displacement vector $\mathbf{L}$ (cf. Eq. (16.5)) is given by

$$\mathbf{L} = \frac{t^2}{d\bar{\mathbf{k}}_0 \cdot \mathbf{t}} d\bar{\mathbf{k}}_0 \qquad \text{[undiffracted]} \tag{27.55}$$

where

$$\mathbf{t} = \begin{pmatrix} 0 \\ 0 \\ t \end{pmatrix} \tag{27.56}$$

with t being the crystal thickness and the z-axis pointing into the crystal as illustrated in Fig. 27.3. The $\mathbf{S}$ rotation matrix (Eq. (16.6)) for the undiffracted beam will be the unit matrix.

With the reference orbit following the `forward_diffracted` or `Bragg_diffracted` beam, the displacement vector $\mathbf{L}$ follows the energy flow associated with the tie points labeled A or B in Fig. 27.3. These tie points are defined by the intersection of the dispersion surfaces and the vector $\mathbf{n}$ originating from the point T as shown in the figure. The energy flow is perpendicular to the dispersion surface and it can be shown that since, by construction, $\mathbf{n}$ goes through the Q point, and since the dispersion surfaces are hyperboles, the energy flows from A and B tie points are colinear. The direction of the energy flow is given by:

$$\overline{\mathbf{K}}_f = \xi_H \overline{\mathbf{K}}_H + \xi_0 \overline{\mathbf{K}}_0 \tag{27.57}$$

where ξ_H and ξ_0 are given by [Bater64] Eq. (18) (See section §27.4.5 below). $\mathbf{L}$ is thus

$$\mathbf{L} = \frac{t^2}{\overline{\mathbf{K}}_f \cdot \mathbf{t}} \overline{\mathbf{K}}_f \tag{27.58}$$

At the exit surface, if the reference orbit is following the `forward_diffracted` beam, the orientation of the `element exit` coordinates will be the same as the orientation of the `element entrance` coordinates. That is, $\mathbf{S}$ (Eq. (16.6)) is the unit matrix. If the reference orbit is following the `Bragg diffracted` beam, $\mathbf{S}$ is the same as for Bragg diffraction

$$\mathbf{S} = \begin{pmatrix} \cos\theta_{bend} & 0 & -\sin\theta_{bend} \\ 0 & 1 & 0 \\ \sin\theta_{bend} & 0 & \cos\theta_{bend} \end{pmatrix} \tag{27.59}$$

27.4.4 Crystal Surface Reflection and Refraction

There are corrections to the field amplitude and phase when a photon reflects or refracts from the surface of a crystal. A plane wave is incident on a crystal surface with

$$E = \widehat{E}_0 \exp(i\,\mathbf{k}_0\,\mathbf{r}) \tag{27.60}$$

An outgoing plane wave has a field

$$E = \widehat{E}_1 \exp(i\,\mathbf{k}_1\,\mathbf{r}) \tag{27.61}$$

A simulation of this condition will start with a number of photons with wave vector $\mathbf{k}_0$ and electric field E_0. After reflecting from the surface, the photons will have wave vector $\mathbf{k}_1$. Now imagine a set of N photons that flow through an planar area dA_0, perpendicular to the incoming beam, before being reflected from the surface.

Since the electric field is $\widehat{E}_0$, when tracking incoherent photons

$$\widehat{E}_0^2 = \frac{\alpha_p\,E_0^2\,N}{dA_0} \tag{27.62}$$

where α_p is the simulation constant (cf. Eq. (27.2). After the photons are reflected they will have some field E_1 and thus

$$\widehat{E}_1^2 = \frac{\alpha_p\,E_1^2\,N}{dA_1} \tag{27.63}$$

Where dA_1 is the area that the photons flow through which is related to dA_0 via

$$\frac{dA_1}{dA_0} = \frac{\mathbf{k}_1 \cdot \mathbf{z}}{\mathbf{k}_0 \cdot \mathbf{z}} \equiv |b| \tag{27.64}$$

Combining the above three equations, the change in field for a photon as it reflects from the surface is

$$\frac{E_1}{E_0} = \frac{\widehat{E}_1}{\widehat{E}_0}\sqrt{|b|} \qquad \text{Incoherent} \tag{27.65}$$

For coherent photon tracking the electric field at dA_0 is

$$\widehat{E}_0 = \frac{\alpha_p\,E_0\,N}{dA_0} \tag{27.66}$$

After the photons are reflected they will have some field E_1 and thus

$$\widehat{E}_1 = \frac{\alpha_p\,E_1\,N}{dA_1} \tag{27.67}$$

Combining these equations the change in field for a photon as it reflects from the surface is

$$\frac{E_1}{E_0} = \frac{\widehat{E}_1}{\widehat{E}_0}\,|b| \qquad \text{Coherent} \tag{27.68}$$

Additionally, for coherent tracking, all photons in a plane wave must have the same phase when passing through an area transverse to the wave. Thus the two photons labeled a and b in Fig. **??** must have the same phase advance in going from dA_0 to dA_1. The difference in the phase advance for photon b relative to a from dA_0 to the surface is $\mathbf{k}_0 \cdot \mathbf{r}$ where $\mathbf{r}$ is the vector between where photon b hits the surface relative to photon a. Similarly, the difference in the phase advance for photon b relative to a from the

surface to dA_0 is $-\mathbf{k}_1 \cdot \mathbf{r}$. Since the total phase advance for both photons is the same from dA_0 to dA_1 the phase shift $d\phi_b$ of photon b as it is reflected from the surface relative to the phase shift $d\phi_a$ is

$$d\phi_b = d\phi_a - (\mathbf{k}_1 - \mathbf{k}_0) \cdot \mathbf{r} \tag{27.69}$$

This shift in the reflection phase can be related to the lattice diffraction planes. The wave vector difference can be written

$$\mathbf{k}_1 - \mathbf{k}_0 = \mathbf{H} + q\,\hat{\mathbf{n}} \tag{27.70}$$

where $\hat{\mathbf{n}}$ is perpendicular to the surface. Combining Eqs. (27.69) and (27.70) and since $\mathbf{r}$ is in the plane of the surface

$$d\phi_b = d\phi_a - \mathbf{H} \cdot \mathbf{r} \tag{27.71}$$

This shows that the reflection shift has the same periodicity as the pattern of the lattice planes at the surface of the crystal. Notice that for a mirror, where one point on the surface is the same as any other, $d\phi_b$ must be equal to $d\phi_a$. Using this in Eq. (27.69) gives

$$\mathbf{k}_1 \cdot \mathbf{r} = \mathbf{k}_0 \cdot \mathbf{r} \tag{27.72}$$

and since $|\mathbf{k}_1| = |\mathbf{k}_0|$ this proves that the angle of incidence is equal to the angle of reflection for a mirror.

In practice, the registration of the surface planes with respect to the surface is not specified in a simulation. Thus the reflection phase shift can only be calculated up to a constant offset.

27.4.5 Bragg Crystal Tracking

The starting photon coordinates are specified in the laboratory entrance coordinates. The transformation from laboratory entrance coordinates to element entrance coordinates $\widetilde{\mathbf{k}}_0$ is given in §27.3. The transformation to element surface coordinates $\mathbf{k}_0$ is

$$\mathbf{k}_0 = \boldsymbol{\Sigma}_2\,\widetilde{\mathbf{k}}_0 \tag{27.73}$$

with $\boldsymbol{\Sigma}_2$ given by Eq. (27.47). The outgoing wave vector $\mathbf{k}_H$ is related to $\mathbf{k}_0$ via

$$\mathbf{k}_H = \mathbf{k}_0 + \mathbf{H} + q_t\,\hat{\mathbf{n}} \tag{27.74}$$

where q_t is determined by using Eqs. (27.34) and (27.35) in Eq. (27.30)

$$\begin{aligned} k_{H,x} &= k_{0,x} + H_x \\ k_{H,y} &= k_{0,y} + H_y \\ k_{H,z} &= \sqrt{\lambda^2 - k_{H,x}^2 - k_{H,y}^2} \end{aligned} \tag{27.75}$$

To compute the field amplitude of the outgoing photon, the equation to be solved is ([Bater64] Eq. (21))

$$\xi_0\,\xi_H = \frac{1}{4}\,k^2\,P^2\,\Gamma^2\,F_H\,F_{\bar{H}} \tag{27.76}$$

where ξ_0 and ξ_H are given by [Bater64] Eq. (18) and P is the polarization factor

$$P = \begin{cases} 1 & \sigma \text{ polarization state} \\ \cos 2\theta_g & \pi \text{ polarization state} \end{cases} \tag{27.77}$$

$2\theta_g$ is the angle between $\mathbf{K}_0$ and $\mathbf{K}_H$ which is well approximated by $\theta_{B,in} + \theta_{B,out}$.

The solution to Eq. (27.76) is ([Bater64] Eq. (31))

$$\begin{aligned}\xi_0 &= \frac{1}{2}\,k\,|P|\,\Gamma\,[F_H\,F_{\bar{H}}]^{1/2}\,|b|^{1/2}\,[\eta \pm (\eta^2 + \mathrm{sgn}(b))^{1/2}] \\ \xi_H &= \frac{1}{2}\,k\,|P|\,\Gamma\,[F_H\,F_{\bar{H}}]^{1/2}\,\frac{1}{|b|^{1/2}\,[\eta \pm (\eta^2 + \mathrm{sgn}(b))^{1/2}]}\end{aligned} \tag{27.78}$$

where the + part of $\pm$ is for the α branch and the $-$ part of $\pm$ is for the β branch and sgn is the sign function

$$\mathrm{sgn}(b) \equiv \begin{cases} 1 & b > 0 \\ -1 & b < 0 \end{cases} \tag{27.79}$$

and η is given by [Blas94] Eq. (5)

$$\eta = \frac{-b\,a + \Gamma\,F_0\,(1-b)}{2\,\Gamma\,|P|\,\sqrt{|b|\,F_H\,F_{\bar{H}}}} \tag{27.80}$$

with the asymmetry factor b for the photon being tracked being given by [Blas94] Eq. (3)

$$b \equiv \frac{\widehat{\mathbf{n}} \cdot \widehat{\mathbf{k}}_0}{\widehat{\mathbf{n}} \cdot \widehat{(\mathbf{k}_0 + \mathbf{H})}} \tag{27.81}$$

and the angular deviation variable a is given by [Blas94] Eq. (4)

$$a \equiv \frac{H^2 + 2\,\mathbf{k}_0 \cdot \mathbf{H}}{k_0^2} = -2\,\Delta\theta\,\sin(2\theta_B) \tag{27.82}$$

Once ξ_0 and ξ_H are determined, the ratio of the incoming and outgoing fields for the α or β branches can be computed via ([Bater64] Eq. (24))

$$r_E \equiv \frac{\mathbf{E}_H}{\mathbf{E}_0} = \frac{-\,2\,\xi_0}{k\,P\,\Gamma\,F_{\bar{H}}} = \frac{-\,k\,P\,\Gamma\,F_H}{2\,\xi_H} \tag{27.83}$$

where the α or β subscript has been suppressed. The total field which is the sum of the fields on the branches is computed using the boundary conditions

$$\mathbf{E}_0 = \mathbf{E}_{0\alpha} + \mathbf{E}_{0\beta}, \qquad 0 = \mathbf{E}_{H\alpha} + \mathbf{E}_{H\beta} \tag{27.84}$$

Using the above two equations gives

$$\begin{aligned}\mathbf{E}_{0\alpha} &= \mathbf{E}_0\,\frac{r_{E\beta}}{r_{E\beta} - r_{E\alpha}} & \mathbf{E}_{H\alpha} &= \mathbf{E}_0\,\frac{r_{E\alpha}\,r_{E\beta}}{r_{E\beta} - r_{E\alpha}} \\ \mathbf{E}_{0\beta} &= -\mathbf{E}_0\,\frac{r_{E\alpha}}{r_{E\beta} - r_{E\alpha}} & \mathbf{E}_{H\beta} &= -\mathbf{E}_0\,\frac{r_{E\alpha}\,r_{E\beta}}{r_{E\beta} - r_{E\alpha}}\end{aligned} \tag{27.85}$$

As can be seen from Battermann and Cole Figs. (8) and (29), the α tie point is excited and the β tie point is not if $\xi_{0\alpha} < \xi_{0\beta}$ and vice versa. Since only one tie point is excited, The external field ratio is equal to the internal field ratio

$$\frac{E_H^e}{E_0^i} = \frac{E_{Hj}}{E_{0j}} \tag{27.86}$$

where j is α or β as appropriate.

27.4.6 Coherent Laue Crystal Tracking

Laue diffraction has two interior wave fields (branches), labeled α and β, corresponding to the two tie points that are excited on the two dispersion surfaces. For coherent tracking, a photon has some probability to be channeled to follow the α or β branch. The electric field ratios $\widehat{E}_\alpha$ and $\widehat{E}_\beta$ (cf. Eq. (27.12)) are taken to be equal to each other. With this choice, the probabilities P_α and P_β for being channeled to the α or β branches are such that a branch with a greater intensity will have a greater number of photons channeled down it.

When a crystal's `ref_orbit_follows` parameter is set to `bragg_diffracted`, The branching probabilities are

$$P_\alpha = \frac{|E_{H\alpha}|}{|E_{H\alpha}| + |E_{H\beta}|}, \qquad P_\beta = \frac{|E_{H\beta}|}{|E_{H\alpha}| + |E_{H\beta}|}, \qquad \widehat{E}_{H\alpha} = \widehat{E}_{H\beta} = \frac{|E_{H\alpha}| + |E_{H\beta}|}{|E_0^i|} \tag{27.87}$$

where (see Battermann and Cole[Bater64] Eqs (42)),

$$\begin{aligned} E_{H\alpha} &= -E_0^i \, \frac{|b|^{1/2}}{2\,\cosh v} \, \frac{|P|}{P} \, \frac{[F_H\, F_{\overline{H}}]^{1/2}}{F_H} \, \exp(-2\,\pi\, i\, \mathbf{K}'_{H\alpha} \cdot \mathbf{r}_\alpha) \, \exp(-2\,\pi\, \mathbf{K}''_{H\alpha} \cdot \mathbf{r}_\alpha) \\ E_{H\beta} &= E_0^i \, \frac{|b|^{1/2}}{2\,\cosh v} \, \frac{|P|}{P} \, \frac{[F_H\, F_{\overline{H}}]^{1/2}}{F_H} \, \exp(-2\,\pi\, i\, \mathbf{K}'_{H\beta} \cdot \mathbf{r}_\beta) \, \exp(-2\,\pi\, \mathbf{K}''_{H\beta} \cdot \mathbf{r}_\beta) \end{aligned} \tag{27.88}$$

where $\mathbf{r}_\alpha$ and $\mathbf{r}_\beta$ are the vectors from the entrance surface to the exit surface for the α and β wave fields

$$\mathbf{r}_\alpha = \frac{t^2}{\mathbf{S}_\alpha \cdot \mathbf{t}}\, \mathbf{S}_\alpha, \qquad \mathbf{r}_\beta = \frac{t^2}{\mathbf{S}_\beta \cdot \mathbf{t}}\, \mathbf{S}_\beta \tag{27.89}$$

with

$$\begin{aligned} \mathbf{S}_\alpha &= e^{-2\,v}\, \mathbf{s}_0 + \left| b\, \frac{F_H\, F_{\overline{H}}}{F_H^2} \right| \mathbf{s}_H \\ \mathbf{S}_\beta &= e^{2\,v}\, \mathbf{s}_0 + \left| b\, \frac{F_H\, F_{\overline{H}}}{F_H^2} \right| \mathbf{s}_H \end{aligned} \tag{27.90}$$

The phase shift of the electric field is obtained from the phase of $E_{H\alpha}$ if the photon is channeled into the α branch and $E_{H\beta}$ if the photon is channeled into the β branch.

When a crystal's `ref_orbit_follows` parameter is set to `forward_diffracted` or `undiffracted`, the algorithm is similar to the `bragg_diffracted` case except $E_{0\alpha}$ and $E_{0\beta}$ are used in place of $E_{H\alpha}$ and $E_{H\beta}$ with

$$\begin{aligned} E_{0\alpha} &= E_0^i \, \frac{e^{-v}}{2\,\cosh v} \, \exp(-2\,\pi\, i\, \mathbf{K}'_{0\alpha} \cdot \mathbf{r}_\alpha) \, \exp(-2\,\pi\, \mathbf{K}''_{0\alpha} \cdot \mathbf{r}_\alpha) \\ E_{0\beta} &= E_0^i \, \frac{e^{-v}}{2\,\cosh v} \, \exp(-2\,\pi\, i\, \mathbf{K}'_{0\beta} \cdot \mathbf{r}_\alpha) \, \exp(-2\,\pi\, \mathbf{K}''_{0\beta} \cdot \mathbf{r}_\beta) \end{aligned} \tag{27.91}$$

Since a simulation photon has two polarization components, the above equations are used for one polarization component and for the second polarization component the same branch is used as for the first with an appropriately scaled $\widehat{E}$.

27.5 X-ray Targeting

X-rays can have a wide spread of trajectories resulting in many "doomed" photons that hit apertures or miss the detector with only a small fraction of "successful" photons actually contributing to the

simulation results. The tracking of doomed photons can therefore result in an appreciable lengthening of the simulation time. To get around this, *Bmad* can be setup to use what is called "targeting" to greatly reduce the number of doomed photons generated.

Photons can be generated either at a source like a `wiggler` element or at a place where diffraction is simulated like at a `diffraction_plate` element. To be able to do targeting, an element with apertures defined must be present downstream from the generating element. The idea is to only generate photons that are going in the general direction of the "target" which is the space within the aperture.

A necessary restriction for targeting to work is that the photon must travel in a straight line through all elements between the generating element and the element with the apertures. So, for example, a `crystal` element would not be allowed between the two two elements. A `crystal` element could be the aperture element as long as the aperture was defined before photons were diffracted. That is, if the aperture was at the upstream end of the crystal or was defined with respect to the `crystal` surface.

The target is defined by the four corners of the aperture. In `element` coordinates, the (x, y, z) values of the corners are:

```
(-x1_limit, -y1_limit, z_lim)
(-x1_limit,  y2_limit, z_lim)
( x2_limit, -y1_limit, z_lim)
( x2_limit,  y2_limit, z_lim)
```

where `x1_limit`, etc. are the aperture limits (§5.8) and `z_lim` will be zero except if the element's `aperture_at` parameter is set to `entrance_end` in which case `z_lim` will be set to `-L` where `L` is the length of the element.

If the aperture is associated with a curved surface (for example with a `crystal` element), four extra corner points are also used to take into account that the aperture is not planar. These extra points have (x, y, z) values in `element` coordinates of

```
(-x1_limit, -y1_limit, z_surface(-x1_limit, -y1_limit))
(-x1_limit,  y2_limit, z_surface(-x1_limit,  y2_limit))
( x2_limit, -y1_limit, z_surface( x2_limit, -y1_limit))
( x2_limit,  y2_limit, z_surface( x2_limit,  y2_limit))
```

where `z_surface(x,y)` is the z value of the surface at the particular (x, y) point being used. Notice that in this case `z_lim` is zero.

The coordinates of the four or eight corner points are converted from `element` coordinates of the aperture element to `element` coordinates of the photon generating element. Additionally, the approximate center of the aperture, which in `element` coordinates of the aperture element is $(0, 0, z_lim)$, is converted to `element` coordinates of the photon generating element.

The above calculation only has to be done once at the beginning of a simulation.

When a photon is to be emitted from a given point $(x_{emit}, y_{emit}, z_{emit})$, the problem is how to restrict the velocity vector $(\beta_x, \beta_y, \beta_z)$ (which is normalized to 1) to minimize the number of doomed photons generated. The problem is solved by constructing a vector $\mathbf{r}$ for each corner point:

$$\mathbf{r} = (x_{lim}, y_{lim}, z_{lim}) - (x_{emit}, y_{emit}, z_{emit}) \tag{27.92}$$

The direction of each $\mathbf{r}$ is characterized in polar coordinates (ϕ, y) defined by

$$\begin{aligned} y &= \frac{r_y}{|r|} \\ \tan\phi &= \frac{r_x}{r_z} \end{aligned} \tag{27.93}$$

For now make the assumption that r_z is positive and larger than r_x and r_y for all $\mathbf{r}$. Let ϕ_{max} and ϕ_{min} be the maximum and minimum ϕ values over all the $\mathbf{r}$. Similarly, let y_{min} and y_{max} be the minimum and maximum y values over all the $\mathbf{r}$. The rectangle in (ϕ, y) space defined by these four min and max values almost covers the projection of the aperture onto the unit sphere. There is a correction that must be made due to the fact that a straight line of constant y in (x, y, z) space projects to a curved line when projected onto (ϕ, y) space. Therefore a correction must be made to y_{min} when $y_{min} < 0$:

$$y_{min} \to \frac{y_m in}{\sqrt{(1 - y_{min}^2)\,\cos^2(phi_{max} - \phi_{min})/2 + y_{min}^2}} \tag{27.94}$$

with a similar correction for y_{max} that must be made when $y_{max} > 0$.

The above prescription works as long as the projection of the aperture onto (ϕ, y) space does not touch the branch cut at $\phi = \pi$ or cover the singular points $y = \pm 1$. Generally these restrictions are fullfilled since $\mathbf{z}$ is the direction of the reference orbit. If this is not the case, a transformation can be made where rotation matrices are constructed to transform between the `element` coordinates of the emitting element and what are called `target` coordinates defined so that $\mathbf{r}$ for the `center` point has the form $(0, 0, |r|)$. The procedure for calculating the photon velocity vector is now

1. Rotate all the corner $\mathbf{r}$ from `element` to `target` coordinates.
2. Calculate min and max values for ϕ and y.
3. Calculate the velocity vector such that the (ϕ, y) of this vector falls within the min and max values in the last step.
4. Rotate the velocity vector back to `element` coordinates.

Chapter 28

Simulation Modules

In the *Bmad* "ecosystem", various modules have been developed to simulate machine hardware. This chapter provides documentation.

28.1 Instrumental Measurements

Bmad has the ability to simulate instrumental measurement errors for orbit, dispersion, betatron phase, and coupling measurements. The appropriate attributes are listed in §5.22 and the conversion formulas are outlined below.

28.1.1 Orbit Measurement

For orbits, the relationship between measured position $(x, y)_{\rm meas}$ and true position $(x, y)_{true}$ is

$$\begin{pmatrix} x \\ y \end{pmatrix}_{\rm meas} = n_f \begin{pmatrix} r_1 \\ r_2 \end{pmatrix} + \mathbf{M}_m \left[\begin{pmatrix} x \\ y \end{pmatrix}_{true} - \begin{pmatrix} x \\ y \end{pmatrix}_0 \right] \tag{28.1}$$

where the Gaussian random numbers r_1 and r_2 are centered at zero and have unit width and the factor n_f represents the inherent noise in the measurement. In the above equation, $(x, y)_0$ represents a measurement offset and $\mathbf{M}_g$ is a "gain" matrix written in the form

$$\mathbf{M}_m = \begin{pmatrix} (1 + dg_x)\cos(\theta + \psi) & (1 + dg_x)\sin(\theta + \psi) \\ -(1 + dg_y)\sin(\theta - \psi) & (1 + dg_y)\cos(\theta - \psi) \end{pmatrix} \tag{28.2}$$

Here dg_x and dg_y represent gain errors and the angles θ and ψ are tilt and "crunch" errors.

In the above equations, various quantities are written as a difference between an "error" quantity and a "calibration" quantity:

$$\begin{aligned} x_0 &= x_{\rm off} - x_{\rm calib} \\ y_0 &= y_{\rm off} - y_{\rm calib} \\ \psi &= \psi_{\rm err} - \psi_{\rm calib} \\ \theta &= \theta_{\rm err} - \theta_{\rm calib} \\ dg_x &= dg_{x,\rm err} - dg_{x,\rm calib} \\ dg_y &= dg_{y,\rm err} - dg_{y,\rm calib} \end{aligned} \tag{28.3}$$

See §5.22 for the element attribute names that correspond to these quantities.

The calibration component is useful in a simulation where initally the error quantities are set to represent the errors in the monitors. After this, analysis of orbit data with the machine in various states can be used to calculate a best guess as to what the errors are. The calculated error values can then be put in the calibration quantities. This represents a correction in software of the errors in the monitors. Further simulations of orbit measurements will show how well the actual orbit can be deduced from the measured orbit.

28.1.2 Dispersion Measurement

A dispersion measurement is considered to be the result of measuring the orbit at two different energies. The measured values are then

$$\begin{pmatrix}\eta_x\\ \eta_y\end{pmatrix}_{\text{meas}} = \frac{\sqrt{2}\,n_f}{dE/E}\begin{pmatrix}r_1\\ r_2\end{pmatrix} + \mathbf{M}_m\left[\begin{pmatrix}\eta_x\\ \eta_y\end{pmatrix}_{true} - \left(\begin{pmatrix}\eta_x\\ \eta_y\end{pmatrix}_{err} - \begin{pmatrix}\eta_x\\ \eta_y\end{pmatrix}_{calib}\right)\right] \tag{28.4}$$

The factor of $\sqrt{2}$ comes from the fact that there are two measurements. $\mathbf{M}_m$ is given in Eq. (28.2).

28.1.3 Coupling Measurement

The coupling measurement is considered to be the result of measuring the beam at a detector over N_s turns while the beam oscillates at a normal mode frequency with some amplitude A_{osc}. The measured coupling is computed as follows. First, consider excitation of the a-mode which can be written in the form:

$$\begin{pmatrix}x_i\\ y_i\end{pmatrix}_{\text{true}} = A_{\text{osc}}\begin{pmatrix}\cos\phi_i\\ K_{22a}\,\cos\phi_i + K_{12a}\sin\phi_i\end{pmatrix}_{\text{true}} \tag{28.5}$$

i is the turn number and ϕ_i is the oscillation phase on the i^{th} turn. The coefficients K_{22a} and K_{12a} are related to the coupling $\overline{\mathbf{C}}$ via Sagan and Rubin[Sagan99] Eq. 54:

$$\begin{aligned} K_{22a} &= \frac{-\sqrt{\beta_b}}{\gamma\sqrt{\beta_a}}\,\overline{\mathbf{C}}_{22}\\ K_{12a} &= \frac{-\sqrt{\beta_b}}{\gamma\sqrt{\beta_a}}\,\overline{\mathbf{C}}_{12}\end{aligned} \tag{28.6}$$

To apply the measurement errors, consider the general case where the beam's oscillations are split into two components: One component being in-phase with some reference oscillator (which is oscillating with the same frequency as the beam) and a component oscillating out-of-phase:

$$\begin{pmatrix}x_i\\ y_i\end{pmatrix}_{\text{true}} = \begin{pmatrix}q_{a1x}\\ q_{a1y}\end{pmatrix}_{\text{true}} A_{\text{osc}}\,\cos(\phi_i + d\phi) + \begin{pmatrix}q_{a2x}\\ q_{a2y}\end{pmatrix}_{\text{true}} A_{\text{osc}}\,\sin(\phi_i + d\phi) \tag{28.7}$$

where $d\phi$ is the phase of the reference oscillator with respect to the beam. Comparing Eq. (28.5) with Eq. (28.7) gives the relation

$$\begin{aligned} K_{22a} &= \frac{q_{a1x}\,q_{a1y} + q_{a2x}\,q_{a2y}}{q_{a1x}^2 + q_{a2x}^2}\\ K_{12a} &= \frac{q_{a1x}\,q_{a2y} - q_{a2x}\,q_{a1y}}{q_{a1x}^2 + q_{a2x}^2}\end{aligned} \tag{28.8}$$

This equation is general and can be applied in either the true or measurement frame of reference. Eq. (28.1) can be used to transform $(x_i, y_i)_{\rm true}$ in Eq. (28.5) to the measurement frame of reference. Only the oscillating part is of interest. Averaging over many turns gives

$$\begin{pmatrix} q_{a1x} \\ q_{a1y} \end{pmatrix}_{\rm meas} = \mathbf{M}_m \begin{pmatrix} q_{a1x} \\ q_{a1y} \end{pmatrix}_{\rm true} , \qquad \begin{pmatrix} q_{a2x} \\ q_{a2y} \end{pmatrix}_{\rm meas} = \mathbf{M}_m \begin{pmatrix} q_{a2x} \\ q_{a2y} \end{pmatrix}_{\rm true} \tag{28.9}$$

This neglects the measurement noise. A calculation shows that the noise gives a contribution to the measured K_{22a} and K_{12a} of

$$K_{22a} \to K_{22a} + r_1 \frac{n_f}{N_s\, A_{\rm osc}} , \qquad K_{12a} \to K_{12a} + r_2 \frac{n_f}{N_s\, A_{\rm osc}} \tag{28.10}$$

Using the above equations, the transformation from the true coupling to measured coupling is as follows: From a knowledge of the true $\overline{\mathbf{C}}$ and Twiss values, the true K_{22a} and K_{12a} can be calculated via Eq. (28.6). Since the value of $d\phi$ does not affect the final answer, $d\phi$ in Eq. (28.7) is chosen to be zero. Comparing this to Eq. (28.5) gives

$$\begin{pmatrix} q_{a1x} \\ q_{a1y} \end{pmatrix}_{\rm true} = \begin{pmatrix} 1 \\ K_{22a} \end{pmatrix}_{\rm true} , \qquad \begin{pmatrix} q_{a2x} \\ q_{a2y} \end{pmatrix}_{\rm true} = \begin{pmatrix} 0 \\ K_{12a} \end{pmatrix}_{\rm true} \tag{28.11}$$

Now Eq. (28.9) is used to convert to the measured q's and Eq. (28.8) then gives the measured K_{22a} and K_{12a}. Finally, Applying Eq. (28.10) and then Eq. (28.6) gives the measured $\overline{\mathbf{C}}_{22}$ and $\overline{\mathbf{C}}_{12}$.

A similar procedure can be applied to b-mode oscillations to calculate values for the measured $\overline{\mathbf{C}}_{11}$ and $\overline{\mathbf{C}}_{12}$. K_{11b} and K_{12b} are defined by

$$\begin{pmatrix} x_i \\ y_i \end{pmatrix}_{\rm true} = A_{\rm osc} \begin{pmatrix} K_{11b}\, \cos\phi_i + K_{12b} \sin\phi_i \\ \cos\phi_i \end{pmatrix}_{\rm true} \tag{28.12}$$

Comparing this to Sagan and Rubin[Sagan99] Eq. 55 gives

$$\begin{aligned} K_{11b} &= \frac{\sqrt{\beta_a}}{\gamma\sqrt{\beta_b}}\, \overline{\mathbf{C}}_{11} \\ K_{12b} &= \frac{-\sqrt{\beta_a}}{\gamma\sqrt{\beta_b}}\, \overline{\mathbf{C}}_{12} \end{aligned} \tag{28.13}$$

The q_{x1b}, q_{y1b}, q_{x2b} and q_{y2b} are defined by using Eq. (28.7) with the "a" subscript replaced by "b". The relationship between K and q is then

$$\begin{aligned} K_{11b} &= \frac{q_{b1y}\, q_{b1x} + q_{b2y}\, q_{b2x}}{q_{b1y}^2 + q_{b2y}^2} \\ K_{12b} &= \frac{q_{b1y}\, q_{b2x} - q_{b2y}\, q_{b1x}}{q_{b1y}^2 + q_{b2y}^2} \end{aligned} \tag{28.14}$$

28.1.4 Phase Measurement

Like the coupling measurement, the betatron phase measurement is considered to be the result of measuring the beam at a detector over N_s turns while the beam oscillates at a normal mode frequency with some amplitude $A_{\rm osc}$. Following the analysis of the previous subsection, the phase ϕ is

$$\begin{pmatrix} \phi_a \\ \phi_b \end{pmatrix}_{\rm meas} = \begin{pmatrix} \phi_a \\ \phi_b \end{pmatrix}_{true} + \frac{n_f}{N_s\, A_{\rm osc}} \begin{pmatrix} r_1 \\ r_2 \end{pmatrix} - \begin{pmatrix} \tan^{-1}\left(\frac{q_{a2x}}{q_{a1x}}\right) \\ \tan^{-1}\left(\frac{q_{b2y}}{q_{b1y}}\right) \end{pmatrix}_{\rm meas} \tag{28.15}$$

Chapter 29

Using PTC/FPP

The PTC/FPP library of Étienne Forest handles Taylor maps to any arbitrary order. This is also known as Truncated Power Series Algebra (TPSA). The core Differential Algebra (DA) package used by PTC/FPP was developed by Martin Berz[Berz89]. The PTC/FPP libraries are interfaced to *Bmad* so that calculations that involve both *Bmad* and PTC/FPP can be done in a fairly seamless manner.

The FPP ("Fully Polymorphic Package") part of the code handles Taylor map manipulation and Lie algebraic operations. This is purely mathematical. FPP has no knowledge of accelerators, magnetic fields, particle tracking etc. PTC ("Polymorphic Tracking Code") implements the physics and uses FPP to handle the Taylor map manipulation. Since the distinction between `FPP` and `PTC` is irrelevant to the non-programmer, "PTC" will be used to refer to the entire package.

PTC is used by *Bmad* when constructing Taylor maps and when the `tracking_method` §6.1) is set to `symp_lie_ptc`. All Taylor maps above first order are calculated via PTC. No exceptions.

For the programmer, see Chapter §39 for more information.

29.1 PTC Tracking Versus Bmad Tracking

While such things as magnet strengths will be the same, the model that PTC uses when it tracks through a lattice element is independent of *Bmad*. That is, what approximations are made can be different. Generally the agreement between PTC and *Bmad* tracking is quite good. But there are situations where there is a noticeable difference. In such cases, one thing to do is to vary parameters that affect PTC tracking. Two parameters that affect PTC accuracy are the integration step size and the order of the integrator. The step size is set by each lattice element's `ds_step` parameter (§6.5.1).

29.2 PTC / Bmad Interfacing

Bmad interfaces to PTC in two ways: One way, called "single element" mode, uses PTC on a per element basis. In this case, the method used for tracking a given element can be selected on an element-by-element basis so non-PTC tracking methods can be mixed with PTC tracking methods to optimize speed and accuracy. [PTC tends to be accurate but slow.] The advantage of single element mode is the flexibility it affords. The disadvantage is that it precludes using PTC's analysis tools which rely on the entire lattice

being tracked via PTC. Such tools include normal form analysis beam envelope tracking, etc.

The alternative to single element mode is "whole lattice" mode where a series of PTC `layout`s (equivalent to a *Bmad* branch) are created from a *Bmad* lattice. Whether single element or whole lattice mode (or both) is used is determined by the program being run.

Part III

Programmer's Guide

Chapter 30

Bmad Programming Overview

30.1 Manual Notation

Bmad defines a number of structures and these structures may contain components which are structures, etc. In order to keep the text in this manual succinct when referring to components, the enclosing structure name may be dropped. For example, the `lat_struct` structure looks like

```
type lat_struct
  character(40) name
  type (mode_info_struct) a, b, z
  type (lat_param_struct) param
  type (ele_struct), pointer ::  ele(:)
  type (branch_struct), allocatable :: branch(:)
  ... etc. ...
end type
```

In this example, "`%a`" could be used to refer to, the `a` component of the `lat_struct`. To make it explicit that this is a component of a `lat_struct`, "`lat_struct%a`" is an alternate possibility. Since the vast majority of structures have the "_struct" suffix, this may be shortened to "`lat%a`". A similar notation works for subcomponents. For example, a `branch_struct` looks like

```
type branch_struct
  character(40) name
  integer ix_from_ele                  ! Index of branching element
  integer, pointer :: n_ele_track      ! Number of tracking elements
  integer, pointer :: n_ele_max
  type (ele_struct), pointer :: ele(:) ! Element array
  ... etc. ...
end type
```

The `ele` component of the `branch` component of the `lat_struct` can be referred to using "`lat%branch%ele`", "`%branch%ele`", or "`%ele`". Potentially, the last of these could be confused with the "`lat%ele`" component so "`%ele`" would only be used if the meaning is unambiguous in the context.

30.2 The Bmad Libraries

The code that goes into a program based upon *Bmad* is divided up into a number of libraries. The *Bmad* libraries are divided into two groups. One group of libraries contains "in-house" developed code.

The other ‘‘package’’ libraries consist of “external” code that *Bmad* relies upon.

The in-house developed code libraries are:

bmad
The `bmad` library contains the routines for charged particle simulation including particle tracking, Twiss calculations, symplectic integration, etc., etc.

cpp_bmad_interface The `cpp_bmad_interface` library is for interfacing *Bmad* with C++. This library defines a set of C++ classes corresponding to the major *Bmad* structures. Along with this, the library contains conversion routines to move information between the C++ classes and the corresponding *Bmad* structures.

sim_utils
The `sim_utils` library contains a set of miscellaneous helper routines. Included are routines for string manipulation, file manipulation, matrix manipulation, spline fitting, Gaussian random number generation, etc.

The `package` libraries are:

forest
This is the PTC/FPP (Polymorphic Tracking Code / Fully Polymorphic Package) library of Étienne Forest that handles Taylor maps to any arbitrary order (this is also known as Truncated Power Series Algebra (TPSA)). See Chapter 29 for more details. FPP/PTC is a very general package. For more information see the FPP/PTC manual[Forest02]. The core Differential Algebra (DA) package used by PTC was developed by Martin Berz[Berz89].

fftw
FFTW is a C subroutine library for computing the discrete Fourier transform in one or more dimensions. FFTW has a Fortran 2003 API.

gsl / fgsl
The Gnu Scientific Library (GSL), written in C, provides a wide range of mathematical routines such as random number generators, special functions and least-squares fitting. There are over 1000 functions in total. The FGSL library provides a Fortran interface to the GSL library.

hdf5
`hdf5` is a library for for storing and managing data[HDF5]. In particular, *Bmad* uses this library for storing particle position data and field grid data.

lapack / lapack95
`lapack` is a widely used package of linear algebra routines written in Fortran77. The `lapack95` library provides a Fortran95 interface to `lapack`.

mad_tpsa
`mad_tpsa` is a subset of the MAD-NG (Next Generation MAD) code[MAD-NG]. Specifically, the `mad_tpsa` library implements TPSA (Truncated Power Series Algebra). This is similar to the `FPP` code (see above). There are several advantages to using `mad_tpsa` over `FPP`. For one, using `mad_tpsa` is independent of `PTC` so `mad_tpsa` can be used along side `PTC`. Another reason is that `mad_tpsa` is more flexible and better structured.

open_spacecharge
The `open_spacecharge` library provides low energy tracking with space charge effects.

PGPLOT

The `pgplot` Graphics Subroutine Library is a Fortran or C-callable, device-independent graphics package for making simple scientific graphs. Documentation including a user's manual may be obtained from the `pgplot` web site at

```
www.astro.caltech.edu/~tjp/pgplot.
```

One disadvantage of `pgplot` for the programmer is that it is not the most user friendly. To remedy this, there is a set of Fortran90 wrapper subroutines called `quick_plot`. The `quick_plot` suite is part of the `sim_utils` library and is documented in Chapter 42.

plplot

The `plplot` library is an updated version of `pgplot`. The `plplot` library can be used as a replacement for `pgplot`. The `quick_plot` suite, which is part of the `sim_utils` library and is documented in Chapter 42, provides wrapper routines for `plplot` to make things more programmer friendly.

xraylib

The xraylib library provides routines for obtaining parameters pertinent to the X-ray interaction with matter. xraylib is developed by Tom Schoonjans and is hosted on github[Schoon11]

30.3 Using getf and listf for Viewing Routine and Structure Documentation

As can be seen from the program example in Chapter 31 there is a lot going on behind the scenes even for this simple program. This shows that programming with *Bmad* can be both easy and hard. Easy in the sense that a lot can be done with just a few lines. The hard part comes about since there are many details that have to be kept in mind in order to make sure that the subroutines are calculating what you really want them to calculate.

To help with the details, all *Bmad* routines have in their source files a comment block that explains the arguments needed by the subroutines and explains what the subroutine does. To help quickly access these comments, there are two Python scripts that are supplied with the *Bmad* distribution that are invoked with the commands `listf` ("list function") and `getf` ("get function").

The `getf` command is used to locate routines and structures, and to type out information on them. The form of the command is

```
getf <name>
```

This searches for any routine or structure with the name `<name>`. `<name>` may contain the wild–cards "*" and "." where "*" matches to any number of characters and "." matches to any single character. For example:

```
getf bmad_parser
getf lat_struct
getf twiss_at_.
```

The third line in this example will match to the routine `twiss_at_s` but not the routine `twiss_at_start`. You may or may not have to put quotation marks if you use wild card characters. As an example, the command `getf twiss_struct` produces:

```
/home/cesrulib/cesr_libs/devel/cvssrc/bmad/modules/twiss_mod.f90
  type twiss_struct
```

```
    real(rp) beta, alpha, gamma, phi, eta, etap
    real(rp) sigma, emit
  end type
```

The first line shows the file where the structure is located (This is system and user dependent so don't be surprised if you get a different directory when you use `getf`). The rest of the output shows the definition of the `twiss_struct` structure. The result of issuing the command `getf relative_tracking_charge` is:

```
File: ../../bmad/modules/bmad_utils_mod.f90
!+
! Function relative_tracking_charge (orbit, param) result (rel_charge)
!
! Routine to determine the relative charge/mass of the particle being
! tracked relative to the charge of the reference particle.
!
! Input:
!   orbit -- coord_struct: Particle position structure.
!   param -- lat_param_struct: Structure holding the reference particle id.
!
! Output:
!   rel_charge -- real(rp): Relative charge/mass
!-
function relative_tracking_charge (orbit, param) result (rel_charge)
```

The first line again shows in what file the subroutine is located. The rest of the output explains what the routine does and how it can be called.

The `getf` command can also be used to search for global integer and real parameter constants. For example

```
getf c_light
```

will give the result:

```
File: ../sim_utils/interfaces/physical_constants.f90
    real(rp), parameter :: c_light = 2.99792458d8              ! speed of light
```

[Global constants are constants defined in a module that have global scope (defined before the `contains` statement).] For parameters whose name ends with a dollar sign "`$`" character (§30.5), the dollar sign suffix may be omitted in the search string. For example the search

```
getf quadrupole
```

will give the result

```
File: ../../bmad/modules/bmad_struct.f90
    integer, parameter :: drift$ = 1, sbend$ = 2, quadrupole$ = 3, group$ = 4, ...
```

Since the dollar sign is a special character for the Python regexp module used by `getf`, to include a dollar sign in the search string the dollar sign must be prefixed by three back slashes. Thus the search

```
getf quadrupole\\\$
```

will also locate the value of `quadrupole$`.

The `listf` command is like the `getf` command except that only the file name where a routine or structure is found is printed. The `listf` command is useful if you want to just find out where a routine or structure definition lives. For example, the `listf relative*` command would produce

```
File: ../../bmad/code/relative_mode_flip.f90
    function relative_mode_flip (ele1, ele2) result (rel_mode)

File: ../../bmad/modules/bmad_utils_mod.f90
    function relative_tracking_charge (orbit, param) result (rel_charge)
```

The way `getf` and `listf` work is that they search a list of directories to find the `bmad`, `sim_utils`, and `tao` libraries. Currently the libraries in the *Bmad* distribution that were not developed at Cornell are not searched. This is primarily due to the fact that, to save time, `getf` and `listf` make assumptions about how documentation is arranged in a file and the non–Cornell libraries do not follow this format.

30.4 Precision of Real Variables

Historically, *Bmad* come in two flavors: One version where the real numbers are single precision and a second version with double precision reals. Which version you are working with is controlled by the kind parameter `rp` (Real Precision) which is defined in the `precision_def` module. On most platforms, single precision translates to `rp` = 4 and double precision to `rp` = 8. The double precision version is used by default since round-off errors can be significant in some calculations. Long–term tracking is an example where the single precision version is not adequate. Changing the precision means recompiling all the libraries except `PTC` and `pgplot`. You cannot mix and match. Either you are using the single precision version or you are using the double precision version. Currently, *Bmad* is always compiled double precision and it is a near certainty that there would have to be some fixes if there was ever a need for compiling single precision.

To define floating point variables in Fortran with the correct precision, use the syntax ‘`real(rp)`’. For example:

```
real(rp) var1, var2, var3
```

When you want to define a literal constant, for example to pass an argument to a subroutine, add the suffix `_rp` to the end of the constant. For example

```
var1 =  2.0_rp * var2
call my_sub (var1, 1.0e6_rp)
```

Note that `2_rp` is different from `2.0_rp`. `2_rp` is an integer of kind `rp`, not a real.

Independent of the setting of `rp`, the parameters `sp` and `dp` are defined to give single and double precision numbers respectively.

30.5 Programming Conventions

Bmad subroutines follow the following conventions:

A “$” suffix denotes a parameter: A dollar sign “$” at the end of a name denotes an parameter. For example, in the above program, to check whether an element is a quadrupole one would write:

```
if (lat%ele(i)%key == quadrupole$) ...
```

Checking the source code one would find in the module `bmad_struct`

```
integer, parameter :: drift$ = 1, sbend$ = 2, quadrupole$ = 3, group$ = 4
```

One should always use the parameter name instead of the integer it represents. That is, one should never write

```
if (lat%ele(i)%key == 3) ...  ! DO NOT DO THIS!
```

For one, using the name makes the code clearer. However, more importantly, the integer value of the parameters may at times be shuffled for practical internal reasons. The use of the integer value could thus lead to disastrous results.

By convention all names ending in "$" are parameters. And most "dollar sign" parameters are integers but there are exceptions. For example, the parameter `real_garbage$` is a real number. To find the value of a dollar sign parameter, the `getf` or `listf` (§30.3) commands can be used.

Structure names have a "_struct" suffix: For example: `lat_struct`, `ele_struct`, etc. Structures without a `_struct` are usually part of Étienne's PTC/FPP package.

30.6 Using Modules

When constructing a program unit,[1] the appropriate `use` statement(s) need to appear at the top of the unit to import the appropriate modules. This should be a simple matter but, due to historical reasons, *Bmad* module dependencies are somewhat convoluted. At some point in the future this will be straightened out but, for now, the following serves as a guide on how to handle things.

In many cases, a "`use bmad`" statement is all that is needed. If there is a problem, an error message will be generated at the program linking stage. The error message might look like:

```
[100%] Linking Fortran executable /home/dcs16/linux_lib/production/bin/test
CMakeFiles/test-exe.dir/test.f90.o: In function ‘MAIN__’:
test.f90:(.text+0x22e): undefined reference to ‘taylor_coef_’
gmake[2]: *** [/home/dcs16/linux_lib/production/bin/test] Error 1
gmake[1]: *** [CMakeFiles/test-exe.dir/all] Error 2
```

The "`undefined reference`" line shows that the subroutine `taylor_coef` was called in the program ("`MAIN`") without the necessary `use` statement. To find the correct `use` statement, use the `getf` or `listf` command (§30.3) to find the file that the subroutine is in. Example:

```
> listf taylor_coef

File: ../bmad/modules/taylor_mod.f90
    subroutine taylor_coef (bmad_file, lat, err, output_form, orbit0)
```

This shows that the subroutine lives in the file `taylor_mod.f90`. The coding rule that *Bmad* follows is that the module name is the file name minus the `.f90` suffix.[2] Thus the needed `use` statement in this case is:

```
use taylor_mod
```

Note: if the linking error looks like:

```
[100%] Linking Fortran executable /home/dcs16/linux_lib/production/bin/test
CMakeFiles/test-exe.dir/test.f90.o: In function ‘MAIN__’:
test.f90:(.text+0x230): undefined reference to
                                    ‘taylor_mod_mp_taylor_coef_’
```

here the error message references

```
taylor_mod_mp_taylor_coef_
```

The "`taylor_mod_mp`" prefix shows that the linker knows to look in the module `taylor_mod` for the subroutine. Here the problem is not a missing use statement but rather the problem is that the linker cannot find the module itself. This indicates something like a library is missing in the list of libraries to link to or the order of libraries to link to is wrong. This kind of problem generally happens with libraries other than the *Bmad* library itself. Further documentation on how the list of libraries to link to is defined is contained on the *Bmad* web site[Bmad].

[1] A program unit is a module such as a subroutine, function, or program.
[2] Be aware. PTC code does not follow this logic.

Chapter 31

An Example Bmad Based Program

To get the general feel for how *Bmad* works before getting into the nitty–gritty details in subsequent chapters, this chapter analyzes an example test program.

31.1 Programming Setup

Information on how to setup your work environment for the compiling and linking of programs can be obtained from the *Bmad* web site at:

https://www.classe.cornell.edu/bmad In particular, instructions for compiling, linking, and running of the example program can be obtained by going to the above web page and clicking on the button labeled "Compiling Custom Programs".

31.2 A First Program

Consider the example program shown in Fig. 31.1. The source code for this program is provided with *Bmad* in the directory:

```
$ACC_ROOT_DIR/code_examples/simple_bmad_program
```

[Note: `$ACC_ROOT_DIR` is the root directory of the *Bmad* `Distribution` when you are running "off-site" and is root directory of the `Release` when you are running "on-site".]

```
1  program test
2
3  use bmad                 ! Define the structures we need to know about.
4  implicit none
5  type (lat_struct), target :: lat   ! This structure holds the lattice info
6  type (ele_struct), pointer :: ele, cleo
7  type (ele_pointer_struct), allocatable :: eles(:)
8  type (all_pointer_struct) a_ptr
9  integer i, ix, n_loc
10 logical err
11
12 ! Read in a lattice, and modify the ks solenoid strength of "cleo_sol".
13
14 call bmad_parser ("lat.bmad", lat)  ! Read in a lattice.
15
16 call lat_ele_locator ('CLEO_SOL', lat, eles, n_loc, err)  ! Find element
17 cleo => eles(1)%ele                        ! Point to cleo_sol element.
18 call pointer_to_attribute (cleo, "KS", .true., a_ptr, err) ! Point to KS attribute.
19 a_ptr%r = a_ptr%r + 0.001         ! Modify KS component.
20 call set_flags_for_changed_attribute (cleo, a_ptr%r)
21 call lattice_bookkeeper (lat)
22 call lat_make_mat6 (lat, cleo%ix_ele)      ! Remake transfer matrix
23
24 ! Calculate starting Twiss params if the lattice is closed,
25 ! and then propagate the Twiss parameters through the lattice.
26
27 if (lat%param%geometry == closed$) call twiss_at_start (lat)
28 call twiss_propagate_all (lat)      ! Propagate Twiss parameters
29
30 ! Print info on the first 11 elements
31
32 print *, " Ix  Name              Ele_type                    S      Beta_a"
33 do i = 0, 10
34   ele => lat%ele(i)
35   print "(i4,2x,a16,2x,a,2f12.4)", i, ele%name, key_name(ele%key), ele%s, ele%a%beta
36 enddo
37
38 ! print information on the CLEO_SOL element.
39
40 print *
41 print *, "!-----------------------------------------------------------"
42 print *, "! Information on element: CLEO_SOL"
43 print *
44 call type_ele (cleo, .false., 0, .false., 0, .true., lat)
45
46 deallocate (eles)
47
48 end program
```

Figure 31.1: Example Bmad program

31.3 Explanation of the Simple_Bmad_Program

§33.6).

A line by line explanation of the example program follows. The `use bmad` statement at line 3 defines the *Bmad* structures and defines the interfaces (argument lists) for the *Bmad* subroutines. In particular, the `lat` variable (line 5), which is of type `lat_struct` (§33.3), holds all of the lattice information: The list of elements, their attributes, etc. The call to bmad_parser (line 14) causes the lattice file `lat.bmad` to be parsed and the lattice information is stored the `lat` variable. Note: To get a listing of the `lat_struct` components or to find out more about `bmad_parser` use the `getf` command as discussed in §30.3.

The routine lat_ele_locator (§34.3) is used in line 16 to find the element in the lattice with the name `CLEO_SOL`. Line 17 defines a pointer variable named `cleo` which is used here as shortcut notation rather than having to write `eles(1)%ele` when refering to this element. The call to `pointer_to_attribute` in line 18 sets up a pointer structure `a_ptr%r` to point to the `KS` solenoid strength component of `cleo`. [`a_ptr` also has components `a_ptr%i` and `a_ptr%l` to point to integrer or logical components if needed. See §34.4.]

Line 19 changes the `ks` solenoid strength of `cleo`. Since an element attribute has been changed, the call to set_flags_for_changed_attribute in line 20 is needed for *Bmad* to inform *Bmad* that this attribute has changed and the call to lattice_bookkeeper does the necessary lattice bookkeeping (§33.6).

The call to lat_make_mat6 in line 22 recalculates the linear transfer matrix for the `CLEO_SOL` element.

In line 27, the program checks if the lattice is circular (§10.1) and, if so, uses the routine twiss_at_start to multiply the transfer matrices of the individual elements together to form the 1–turn matrix from the start of the lat back to the start. From this matrix `twiss_at_start` calculates the Twiss parameters at the start of the lattice and puts the information into `lat%ele(0)` (§36.2). The next call, to twiss_propagate_all, takes the starting Twiss parameters and, using the transfer matrices of the individual elements, calculates the Twiss parameters at all the elements. Notice that if the lattice is not circular, The starting Twiss parameters will need to have been defined in the lattice file.

The program is now ready output some information. Lines 21 through 25 of the program print information on the first 11 elements in the lattice. The do-loop is over the array `lat%ele(:)`. Each element of the array holds the information about an individual lattice element as explained in Chapter 33. The `lat%ele(0)` element is basically a marker element to denote the beginning of the array (§7). Using the pointer `ele` to point to the individual elements (line 34) makes for a cleaner syntax and reduces typing. The table that is produced is shown in lines 1 through 12 of Fig. 31.2. The first column is the element index i. The second column, `ele%name`, is the name of the element. The third column, key_name(elethe name of the element class. `ele%key` is an integer denoting what type of element (quadrupole, wiggler, etc.) it is. `key_name` is an array that translates the integer key of an element to a printable string. The fourth column, `ele%s`, is the longitudinal position at the exit end of the element. Finally, the last column, `ele%x%beta`, is the a–mode (nearly horizontal mode) beta function.

The type_ele routine on line 44 of the program is used to type out the `CLEO_SOL`'s attributes and other information as shown on lines 14 through 38 of the output (more on this later).

This brings us to the lattice file used for the input to the program. The call to `bmad_parser` shows that this file is called `simple_bmad_program/lat.bmad`. In this file there is a call to another file

```
call, file = "layout.bmad"
```

It is in this second file that the layout of the lattice is defined. In particular, the `line` used to define the element order looks like

```
cesr: line = (IP_L0, d001, DET_00W, d002, Q00W, d003, ...)
```

```
use, cesr
```

If you compare this to the listing of the elements in Fig. 31.2 you will find differences. For example, element #2 in the program listing is named `CLEO_SOL\3`. From the definition of the `cesr` line this should be `d001` which, if you look up its definition in `layout.bmad` is a drift. The difference between lattice file and output is due to the presence the `CLEO_SOL` element which appears in `lat.bmad`:

```
ks_solenoid    := -1.0e-9 * clight * solenoid_tesla / beam[energy]
cleo_sol: solenoid, l = 3.51, ks = ks_solenoid, superimpose
```

The solenoid is 3.51 meters long and it is superimposed upon the lattice with its center at $s = 0$ (this is the default if the position is not specified). When `bmad_parser` constructs the lattice list of elements the superposition of `IP_L0`, which is a zero–length marker, with the solenoid does not modify `IP_L0`. The superposition of the `d001` drift with the solenoid gives a solenoid with the same length as the drift. Since this is a "new" element, `bmad_parser` makes up a name that reflects that it is basically a section of the solenoid it came from. Next, since the `CLEO_SOL` element happens to only cover part of the `Q00W` quadrupole, `bmad_parser` breaks the quadrupole into two pieces. The piece that is inside the solenoid is a `sol_quad` and the piece outside the solenoid is a regular quadrupole. See §8 for more details. Since the center of the `CLEO_SOL` is at $s = 0$, half of it extends to negative s. In this situation, `bmad_parser` will wrap this half back and superimpose it on the elements at the end of the lattice list near $s = s_{lat}$ where s_{lat} is the length of the lattice. As explained in Chapter 33, the lattice list that is used for tracking extends from `lat%ele(0)` through `lat%ele(n)` where `n = lat%n_ele_track`. The `CLEO_SOL` element is put in the section of `lat%ele(n)` with `n > lat%n_ele_track` since it is not an element to be tracked through. The `Q00W` quadrupole also gets put in this part of the list. The bookkeeping information that the `cleo_sol\3` element is derived from the `cleo_sol` is put in the `cleo_sol` element as shown in lines 30 through 38 of the output. It is now possible in the program to vary, say, the strength of the `ks` attribute of the `CLEO_SOL` and have the `ks` attributes of the dependent ("`super_slave`") elements updated with one subroutine call. For example, the following code increases the solenoid strength by 1% lattice_bookkeeper lat_ele_locator

```
call lat_ele_locator ('CLEO_SOL', lat, eles, n_loc, err)
eles(1)%ele(ix)%value(ks$) = eles(1)%ele%value(ks$) * 1.01
call lattice_bookkeeper (lat)
```

Bmad takes care of the bookkeeping. In fact `control_bookkeeper` is automatically called when transfer matrices are remade so the direct call to `control_bookkeeper` may not be necessary.

Running the program (§31.1) gives the output as shown in Fig. 31.2.

```
    Ix  Name              Ele_type                  S      Beta_a
     0  BEGINNING         BEGINNING_ELE        0.0000      0.9381
     1  IP_L0             MARKER               0.0000      0.9381
     2  CLEO_SOL#3        SOLENOID             0.6223      1.3500
     3  DET_00W           MARKER               0.6223      1.3500
     4  CLEO_SOL#4        SOLENOID             0.6380      1.3710
     5  Q00W\CLEO_SOL     SOL_QUAD             1.7550      7.8619
     6  Q00W#1            QUADRUPOLE           2.1628     16.2350
     7  D003              DRIFT                2.4934     27.4986
     8  DET_01W           MARKER               2.4934     27.4986
     9  D004              DRIFT                2.9240     46.6018
    10  Q01W              QUADRUPOLE           3.8740     68.1771

   !----------------------------------------------------------
   ! Information on element: CLEO_SOL

   Element #              871
   Element Name: CLEO_SOL
   Key: Solenoid
   S_start, S:    766.671421,      1.755000
   Ref_time:  5.854050E-09

   Attribute values [Only non-zero/non-default values shown]:
    1   L                          =  3.5100000E+00 m
    5   KS                         = -8.5023386E-02 1/m
   10   FRINGE_TYPE                =  None (1)
   11   FRINGE_AT                  =  Both_Ends (3)
   13   SPIN_FRINGE_ON             =  T (1)
   47   PTC_CANONICAL_COORDS       =  1.0000000E+00
   49   BS_FIELD                   =  1.5000000E+00 T
   50   DELTA_REF_TIME             =  1.1708100E-08 sec
   53   P0C                        =  5.2890000E+09 eV      BETA            =  0.999999995
   54   E_TOT                      =  5.2890000E+09 eV      GAMMA           =  1.0350315E+04
   66   NUM_STEPS                  = 18
   67   DS_STEP                    =  2.0000000E-01 m

        TRACKING_METHOD            =  Bmad_Standard       APERTURE_AT               =  Exit_End
        MAT6_CALC_METHOD           =  Bmad_Standard       APERTURE_TYPE             =  Rectangular
        SPIN_TRACKING_METHOD       =  Tracking            OFFSET_MOVES_APERTURE     =  F
        PTC_INTEGRATION_TYPE       =  Matrix_Kick         SYMPLECTIFY               =  F
        CSR_METHOD                 =  Off                 FIELD_MASTER              =  F
        SPACE_CHARGE_METHOD        =  Off
        FIELD_CALC                 =  Bmad_Standard

  Slave_status: Free

  Lord_status:  Super_Lord
  Slaves:
     Index   Name           Type
       865   Q00E\CLEO_SOL  Sol_Quad
       866   CLEO_SOL#1     Solenoid
       868   CLEO_SOL#2     Solenoid
         2   CLEO_SOL#3     Solenoid
         4   CLEO_SOL#4     Solenoid
         5   Q00W\CLEO_SOL  Sol_Quad

  Twiss at end of element:
                           A              B           Cbar                        C_mat
  Beta (m)          7.73293815    88.01448113  | -0.16691726   0.00910908     -0.05360747   0.23764236
  Alpha            -6.94661336    -1.53007417  |  2.24946759  -0.02027746      0.03925732   0.14506787
  Gamma (1/m)       6.36956306     0.03796110  |  Gamma_c =   1.00851669       Mode_Flip = F
  Phi (rad)         1.00744612     1.55387058            X              Y              Z
  Eta (m)          -0.08500243    -0.00073095    -0.08600823    -0.00846284     0.00029134
```

Chapter 32

The ele_struct

This chapter describes the `ele_struct` which is the structure that holds all the information about an individual lattice element: quadrupoles, separators, wigglers, etc. The `ele_struct` structure is shown in Figs. 32.1 and 32.2. This structure is somewhat complicated, however, in practice, a lot of the complexity is generally hidden by the *Bmad* bookkeeping routines.

As a general rule, for variables like the Twiss parameters that are not constant along the length of an element, the value stored in the corresponding component in the `ele_struct` is the value at the downstream end of the element.

For printing information about an element, the type_ele routine can be used (§31.2). When called without the optional `lines` and `n_lines` arguments, `type_ele` prints to the terminal window; when those arguments are supplied, the element information is returned as an array of strings instead.

```
type ele_struct
  character(40) name                      ! name of element \sref{c:ele.string}.
  character(40) type                      ! type name \sref{c:ele.string}.
  character(40) alias                     ! Another name \sref{c:ele.string}.
  character(40) component_name            ! Used by overlays, multipass patch, etc.
  character(200), pointer :: descrip      ! Description string.
  type (twiss_struct)  a, b, z            ! Twiss parameters at end of element \sref{c:normal.modes}.
  type (xy_disp_struct) x, y              ! Projected dispersions \sref{c:normal.modes}.
  type (ac_kicker_struct), pointer :: ac_kick  ! ac_kicker element parameters.
  type (bookkeeping_state_struct) bookkeeping_state ! Element attribute bookkeeping
  type (branch_struct), pointer :: branch      ! Pointer to branch containing element.
  type (controller_struct), pointer :: control ! For group, overlay, and ramper elements.
  type (converter_struct), pointer :: converter  ! Positron converter, etc.
  type (rf_ele_struct), pointer :: rf            ! RF element stair-step parameters.
  type (foil_struct), pointer :: foil            ! Foil material parameters.
  type (ele_struct), pointer :: lord           ! Pointer to a slice lord.
  type (fibre), pointer :: ptc_fibre           ! PTC tracking.
  type (floor_position_struct) floor           ! Global floor position.
  type (high_energy_space_charge_struct), pointer :: high_energy_space_charge
                                          ! High-energy space charge parameters.
  type (mode3_struct), pointer :: mode3        ! Full 6-dimensional normal mode decomposition.
  type (multipole_cache_struct), allocatable :: multipole_cache  ! Cached multipole values.
  type (photon_element_struct), pointer :: photon
  type (rad_map_ele_struct), pointer :: rad_map ! Radiation damping and stochastic kick maps.
  type (taylor_struct) :: taylor(6)            ! Orbital Taylor map.
  real(rp) :: spin_taylor_ref_orb_in(6)        ! Reference orbit for spin Taylor map.
  type (taylor_struct) :: spin_taylor(0:3)     ! Spin Taylor map (quaternion representation).
  type (wake_struct), pointer :: wake          ! Wakes.
  type (wall3d_struct), pointer :: wall3d(:)   ! Chamber or capillary wall array.
  ! E/M field map components (see Chapter on Electromagnetic Fields):
  type (cartesian_map_struct), pointer :: cartesian_map(:)     ! Used to define E/M fields.
  type (cylindrical_map_struct), pointer :: cylindrical_map(:) ! Used to define E/M fields.
  type (gen_gradients_struct), pointer :: gen_gradients(:)       ! Used to define E/M fields.
  type (grid_field_struct), pointer :: grid_field(:)           ! Used to define DC and AC fields.
    ele_struct definition continued on next figure...
```

Figure 32.1: The `ele_struct`. structure definition. The complete structure is shown in this and the following figure.

```
    ... ele_struct definition continued from previous figure.
  type(coord_struct) map_ref_orb_in           ! Transfer map ref orbit at upstream end of element.
  type(coord_struct) map_ref_orb_out          ! Transfer map ref orbit at downstream end of element.
  type(coord_struct) time_ref_orb_in          ! Reference orbit at upstream end for ref_time calc.
  type(coord_struct) time_ref_orb_out         ! Reference orbit at downstream end for ref_time calc.
  real(rp) value(num_ele_attrib$)             ! attribute values.
  real(rp) old_value(num_ele_attrib$)         ! Used to see if %value(:) array has changed.
  real(rp) spin_q(0:3,0:6)                    ! 0th and 1st order spin transport quaternion.
  real(rp) vec0(6)                            ! 0th order transport vector.
  real(rp) mat6(6,6)                          ! 1st order transport matrix.
  real(rp) c_mat(2,2)                         ! 2x2 C coupling matrix
  real(rp) dc_mat_dpz(2,2)                    ! d(c_mat)/dpz chromatic variation.
  real(rp) gamma_c                            ! gamma associated with C matrix
  real(rp) s_start                            ! longitudinal ref position at entrance_end
  real(rp) s                                  ! longitudinal position at the downstream end.
  real(rp) ref_time                           ! Time ref particle passes downstream end.
  real(rp), pointer :: r(:,:,:)               ! For general use. Not used by Bmad.
  real(rp), pointer :: a_pole(:)              ! multipole a_n (skew) components.
  real(rp), pointer :: b_pole(:)              ! multipole b_n (normal) components.
  real(rp), pointer :: a_pole_elec(:)         ! Electrostatic multipoles.
  real(rp), pointer :: b_pole_elec(:)         ! Electrostatic multipoles.
  real(rp), pointer :: custom(:)              ! Custom attributes
  integer key                    ! key value
  integer sub_key                ! Records bend input type (rbend$, sbend$).
  integer ix_ele                 ! Index in lat%branch(n)%ele(:) array [n = 0 <==> lat%ele(:)].
  integer ix_branch              ! Index in lat%branch(:) array [0 => In lat%ele(:)].
  integer lord_status            ! overlay_lord$, etc.
  integer n_slave                ! Number of slaves (not counting field overlap slaves).
  integer n_slave_field          ! Number of field slaves.
  integer ix1_slave              ! Pointer to lat%control array.
  integer slave_status           ! super_slave$, etc.
  integer n_lord                 ! Number of lords (not counting field overlap and ramper lords).
  integer n_lord_field           ! Number of field lords.
  integer n_lord_ramper          ! Number of ramper lords.
  integer ic1_lord               ! Pointer to lat%ic array.
  integer ix_pointer             ! For general use. Not used by Bmad.
  integer ixx, iyy, izz          ! Index for Bmad internal use.
  integer mat6_calc_method       ! bmad_standard$, taylor$, etc.
  integer tracking_method        ! bmad_standard$, taylor$, etc.
  integer spin_tracking_method   ! bmad_standard$, symp_lie_ptc$, etc.
  integer csr_method             ! off$, one_dim$, steady_state_3d$, etc.
  integer space_charge_method    ! off$, slice$, fft_3d$, cathode_fft_3d$, etc.
  integer ptc_integration_type   ! drift_kick$, matrix_kick$, etc.
  integer field_calc             ! bmad_standard$, no_field$, fieldmap$, custom$, etc.
  integer ref_species            ! Reference species (not_set$ by default).
  integer aperture_at            ! Aperture location: exit_end$, entrance_end$, etc.
  integer aperture_type          ! Type of aperture: rectangular$, elliptical$, or custom$.
  integer orientation            ! -1 -> Element is longitudinally reversed. +1 -> Normal.
  logical symplectify            ! Symplectify mat6 matrices.
  logical mode_flip              ! Have the normal modes traded places?
  logical multipoles_on          ! For turning multipoles on/off.
  logical scale_multipoles       ! Multipole components scaled by strength of element?
  logical taylor_map_includes_offsets  ! Taylor map calculated with element misalignments?
  logical field_master           ! Calculate strength from the field value?
  logical is_on                  ! For turning element on/off.
  logical logic                  ! For general use. Not used by Bmad.
  logical bmad_logic             ! For Bmad internal use only.
  logical select                 ! For element selection. Used by make_hybrid_ring, etc.
  logical offset_moves_aperture  ! Element offsets affect aperture position?
end type
```

Figure 32.2: The `ele_struct`. The complete structure is shown in this and the preceding figure.

32.1 Initialization and Pointers

The ele_struct has a number of components and subcomponents that are pointers. This is done to minimize memory size. The use of pointers raises a deallocation issue. Generally, most ele_struct elements are part of a lat_struct variable (§33.2) and such elements in a lat_struct are handled by the lat_struct allocation/deallocation routines. In the case where a local ele_struct variable is used within a subroutine or function, the ele_struct variable must either be defined with the save attribute

```
type (ele_struct), save :: ele          ! Use the save attribute
logical, save :: init_needed = .false.
...
if (init_needed) then
  call init_ele (ele, quadrupole$)     ! Initialize element once
  init_needed = .false.
endif
```

or the pointers within the variable must be deallocated with a call to deallocate_ele_pointers:

```
type (ele_struct) ele
...
call init_ele (ele, sbend$)            ! Initialize element each time
...
call deallocate_ele_pointers (ele)     ! And deallocate.
```

In the "normal" course of events, the pointers of an ele_struct variable should not be pointing to the same memory locations as the pointers of any other ele_struct variable. To make sure of this, the equal sign in the assignment ele1 = ele2 is overloaded by the routine ele_equal_ele. The exception here are the "Electro-magnetic field component" pointer arrays ele%cartesian_map, ele%cylindrical_map, ele%gen_gradients, and ele%grid_field. Since these components potentially contain large arrays, and since the individual sub-components are not likely to be individually modified, the field component pointers of ele1 and ele2 after the set ele1 = ele2 will point at the same memory locations.

Note: The assignment ele1 = ele2 will not modify ele1%ix_ele or ele1%ix_branch. If ele1 is associated with a lattice then ele1%lat will also be unaffected.

32.2 Miscellaneous

The %ref_species integer component stores the reference particle species for an element This can differ from element to element in a branch due to converter or foil elements. Note: The setting of branch%param%particle gives the reference species at the beginning of a branch and this will correspond to the setting of %ref_species of the beginning element of that branch.

32.3 Element Attribute Bookkeeping

When a value of an attribute in an element changes, the values of other attributes may need to be changed (§5.1). Furthermore, in a lattice, changes to one element may necessitate changes to attribute values in other elements. For example, changing the accelerating gradient in an lcavity will change the reference energy throughout the lattice.

The attribute bookkeeping for a lattice can be complicated and, if not done intelligently, can cause programs to be slow if attributes are continually being changed. In order to keep track what bookkeeping

has been done, the ele%status component is used by the appropriate bookkeeping routines for making sure the bookkeeping overhead is keep to a minimum. See §33.6 for more details.

32.4 String Components

The %name, %type, %alias, and %descrip components of the ele_struct all have a direct correspondence with the name, type, alias, and descrip element attributes in an input lattice file (§5.3). On input (§35.1), The name attribute is converted to uppercase before being loaded into an ele_struct. The other three are not. To save memory, since %descrip is not frequently used, %descrip is a pointer that is only allocated if descrip is set for a given element.

When a lattice is constructed by bmad_parser, a "nametable" is constructed to enable the fast lookup of element names. Therefore, It is important that if element names are modified, that the nametable is updated appropriately. This is done by calling create_lat_ele_nametable after any element name modifications.

32.5 Element Key

The %key integer component gives the class of element (quadrupole, rfcavity, etc.). In general, to get the corresponding integer parameter for an element class, just add a "$" character to the class name. For example quadrupole$ is the integer parameter for quadrupole elements. The key_name array converts from integer to the appropriate string. For example:

```
type (ele_struct) ele
if (ele%key == wiggler$) then       ! Test if element is a wiggler.
print *, "This element: ", key_name(ele%key) ! Prints, for example, "WIGGLER"
```

Note: The call to init_ele is needed for any ele_struct defined outside of a lat_struct structure.

The %sub_key component is only used for bend element. When a lattice file is parsed, (§35.1), all rbend elements are converted into sbend elements (§4.5). To keep track of what the original definition of the element was, the %sub_key component will be set to sbend$ or rbend$ whatever is appropriate. The %sub_key component does not affect any calculations and is only used in the routines that recreate lattice files from a lat_struct (§35.3).

32.6 The %value(:) array

Most of the real valued attributes of an element are held in the %value(:) array. For example, the value of the k1 attribute for a quadrupole element is stored in %value(k1$) where k1$ is an integer parameter that *Bmad* defines. In general, to get the correct index in %value(:) for a given attribute, add a "$" as a suffix. To convert from an attribute name to its index in the %value array use the attribute_index routine. To go back from an index in the %value array to a name use the attribute_name routine. Example:

```
type (ele_struct) ele
call init_ele (ele, quadrupole$)    ! Initialize element
ele%value(k1$) = 0.3                                    ! Set K1 value
print *, "Index for Quad K1:  ", attribute_index(ele, "K1") ! prints: '4' (= k1$)
print *, "Name for Quad k1$: ", attribute_name (ele, k1$)   ! prints: 'K1'
```

The list of attributes for a given element type is given in the writeup for the different element in Chapter 4.

To obtain a list of attribute names and associated `%value(:)` indexes, the program `element_attributes` can be used. This program is included in the standard *Bmad* distribution.

Besides real valued attributes, the `value(:)` array also holds logical, integer, and, as explained below, "switch" attributes. To find out the type of a given attribute, use the function attribute_type. See the routine `type_ele` for an example of how `attribute_type` is used.

An example of a logical attribute is the `flexible` logical of match elements which is stored in `%value(flexible$)`. To evaluate logical attributes, the functions is_true(param) or is_false(param) should be used.

Integer attributes stored in the `value(:)` array include `n_slice` (stored in `%value(n_slice$)`). With integer attributes, the `nint(param)` Fortran instrinsic should be used for evaluation.

A switch attribte is an attribute whose value is one of a certain set of integers where each integer corresponds to some "state". For example, the `fringe_at` switch which, as explained in §5.21, may have one of four values. Generally, the integer parameters that correspond to the states of a switch can be constructed by putting a "$" after the associated name. Thus, with the `fringe_at` switch, the four integer parameters are `no_end$`, `both_ends$`, `entrance_end$`, and `exit_end$`. For example:

```
if (nint(ele%value(fringe_type$)) == soft_edge_only$) then
  ...
```

For printing purposes, to convert a switch value to the appropriate string, use the routineswitch_attrib_value_name can be used.

The `%field_master` logical within an element sets whether it is the normalized strength or field strength that is the independent variable. See §5.1 for more details.

The `%old_value(:)` component of the `ele_struct` is used by the attribute_bookkeeper routine to check for changes for changes in the `%value(:)` array since the last time the `attribute_bookkeeper` routine had been called. If nothing has been changed, the `attribute_bookkeeper` routine knows not to waste time recalculating dependent values. Essentially what this means is that the `%old_value(:)` array should not be modified outside of `attribute_bookkeeper`.

32.7 Connection with the Lat_Struct

If an element is part of a `lat_struct` (§33), the `%ix_ele` and `%ix_branch` components of the `ele_struct` identify where the element is. Additionally, the `%branch` component points to the `branch_struct` containing the element, and the enclosing `lat_struct` is accessed via `ele%branch%lat`. That is

```
type (lat_struct), pointer :: lat
type (ele_struct), pointer :: ele2
if (ele%ix_ele > -1) then
  ie = ele%ix_ele
  ib = ele%ix_branch
  lat => ele%branch%lat
  ele2 => lat%branch(ib)%ele(ie)
  print *, associated(ele2, ele)  ! Will print True.
endif
```

In this example the `ele2` pointer is constructed to point to the `ele` element. The test (`ele%ix_ele` > -1) is needed since `ele_struct` elements may exist outside of any `lat_struct` instance. Such "external" elements always have `%ix_ele` < 0. A value for `%ix_ele` of `ix_slice_slave$` (= -2) is special in that it prevents the deallocate_ele_pointers routine from deallocating the pointers of an element which has its `%ix_ele` set to `ix_slice_slave$`.

An element "slice" is an example of an element that exists external to any `lat_struct` instance. A slice is an `ele_struct` instance that represents some sub-section of a given element. Element slices are useful when tracking particles only part way through an element (§37.7).

32.8 Limits

The aperture limits (§5.8) in the `ele_struct` are:

```
%value(x1_limit$)
%value(x2_limit$)
%value(y1_limit$)
%value(y2_limit$)
```

The values of these limits along with the `%aperture_at`, `%aperture_type`, and `%offset_moves_aperture` components are used in tracking to determine if a particle has hit the vacuum chamber wall. See Section §37.8 for more details.

32.9 Twiss Parameters, etc.

The components `%a`, `%b`, `%z`, `%x`, `%y`, `%c_mat`, `%gamma_c`, `%mode_flip`, and `%mode3` hold information on the Twiss parameters, dispersion, and coupling at the downstream end of the element. See Chapter 36 for more details.

The `%dc_mat_dpz(2,2)` component stores the chromatic derivative of the 2×2 C-coupling matrix: $\partial C / \partial p_z$. This is computed along with the C matrix during the normal mode analysis.

32.10 Element Lords and Element Slaves

In *Bmad*, elements in a lattice can control other elements. The components that determine this control are:

```
%slave_status
%n_slave          ! Number of slaves (not counting field overlap slaves)
%n_slave_field    ! Number of field overlap slaves
%ix1_slave
%lord_status
%n_lord           ! Number of lords (not counting field overlap and ramper lords)
%n_lord_field     ! Number of field overlap lords
%n_lord_ramper    ! Number of ramper lords
%ic1_lord
%component_name
```

This is explained fully in the chapter on the `lat_struct` (§33).

32.11 Group and Overlay Controller Elements

`Group`, `overlay`, and `ramper` elements use the `%control` pointer for storing controller information. `%control` points to a `controller_struct` which holds the control variables and ramp expressions:

```
type controller_struct
  type (control_var1_struct), allocatable :: var(:)      ! Control variables
  type (control_ramp1_struct), allocatable :: ramp(:)    ! For ramper lord elements
  type (ramper_lord_struct), allocatable :: ramper_lord(:) ! Ramper lord info for this slave
  real(rp), allocatable :: x_knot(:)
end type
```

The `var(:)` component is an array of `control_var1_struct` elements, one per control variable:

```
type control_var1_struct
  character(40) :: name = ""
  real(rp) :: value = 0
  real(rp) :: old_value = 0
end type
```

The `%old_value` component of `control_var1_struct` is only used for `group` elements.

See Section §34.2 for an example of setting up a controller element within a program.

32.12 Coordinates, Offsets, etc.

The "upstream" and "downstream" ends of an element are, by definition, where the physical ends of the element would be if there were no offsets. In particular, if an element has a finite `z_offset`, the physical ends will be displaced from upstream and downstream ends. See §37.2 for more details.

The `%floor` component gives the "laboratory" global "floor" coordinates (§16.2) at the downstream end of the element. These coordinates are computed without misalignments. That is, the coordinates are not "body" coordinates. The components of the `%floor` structure are

```
type floor_position_struct
  real(rp) r(3)               ! Offset from origin
  real(rp) w(3,3)             ! Orientation matrix (Eq. (16.2))
  real(rp) theta, phi, psi    ! Angular orientation
end type
```

The routine ele_geometry will calculate an element's floor coordinates given the floor coordinates at the beginning of the element. In a lattice, the lat_geometry routine will calculate the floor coordinates for the entire lattice using repeated calls to `ele_geometry`.

The positional offsets (§5.6) for an element from the reference orbit are stored in

```
%value(x_offset$)
%value(y_offset$)
%value(z_offset$)
%value(x_pitch$)
%value(y_pitch$)
%value(tilt$)
```

If the element is supported by a `girder` element (§4.24) then the `girder` offsets are added to the element offsets and the total offset with respect to the reference coordinate system is stored in:

```
%value(x_offset_tot$)
%value(y_offset_tot$)
%value(z_offset_tot$)
%value(x_pitch_tot$)
%value(y_pitch_tot$)
%value(tilt_tot$)
```

If there is no `girder`, the values for `%value(x_offset_tot$)`, etc. are set to the corresponding values in `%value(x_offset$)`, etc. Thus, to vary the position of an individual element the values of `%value(x_offset$)`, etc. are changed and to read the position of an element a program should look at `%value(x_offset_tot$)`, etc.

The longitudinal position at the downstream end of an element is stored in `%s` and the reference time is stored in `%ref_time`. This reference time is calculated assuming that the reference time is zero at the start of the lattice. Also stored in the `ele_struct` is the reference time at the start of the element and the differnece in the reference time between the end and the beginning. These are given in `%value(ref_time_start$)` and `%value(delta_ref_time$)` respectively.

Notice that the reference time used to calculate the z phase space coordinate (Eq. (17.3)) may be different from `%ref_time`. For example, with multiple bunches the z phase space coordinate is generally taken to be with respect to a reference particle at the center of the bunch the particle is in. And, at a given element, the reference time of the different bunch reference particles will be different. Another example happens when a particle is tracked through multiple turns. In this case the reference time at a given element will depend upon the turn number.

32.13 Transfer Maps: Linear and Non-linear (Taylor)

The routine make_mat6 computes the linear transfer matrix (Jacobian) along with the zeroth order transfer vector. This matrix is stored in `%mat6(6,6)` and the zeroth order vector is stored in `%vec0(6)`. The reference orbit at the upstream end of the element about which the transfer matrix is computed is stored in `%map_ref_orb_in` and the reference orbit at the downstream end is stored in `%map_ref_orb_out`. In the calculation of the transfer map, the vector `%vec0` is set so that

```
map_ref_orb_out = %mat6 * map_ref_orbit_in + %vec0
```

The reason redundant information is stored in the element is to save computation time.

To compute the transfer maps for an entire lattice use the routine lat_make_mat6.

The Taylor map (§6) for an element is stored in `%taylor(1:6)`. Each `%taylor(i)` is a `taylor_struct` structure that defines a Taylor series:

```
type taylor_struct
  real (rp) ref
  type (taylor_term_struct), pointer :: term(:) => null()
end type
```

Each Taylor series has an array of `taylor_term_struct` terms defined as

```
type taylor_term_struct
  real(rp) :: coef
  integer :: expn(6)
end type
```

The coefficient for a Taylor term is stored in `%coef` and the six exponents are stored in `%exp(6)`.

To see if there is a Taylor map associated with an element the association status of `%taylor(1)%term` needs to be checked. As an example the following finds the order of a Taylor map.

```
type (ele_struct) ele
...
if (associated(ele%taylor(1)%term) then  ! Taylor map exists
  taylor_order = 0
  do i = 1, 6
    do j = 1, size(ele%taylor(i)%term)
```

```
      taylor_order = max(taylor_order, sum(ele%taylor(i)%term(j)%exp)
    enddo
  enddo
else  ! Taylor map does not exist
  taylor_order = -1  ! flag non-existence
endif
```

The Taylor map is made up around some reference phase space point corresponding to the coordinates at the upstream of the element. This reference point is saved in `%taylor(1:6)%ref`. Once a Taylor map is made, the reference point is not needed in subsequent calculations. However, the Taylor map itself will depend upon what reference point is chosen (§25.1).

32.14 Reference Energy and Time

The reference energy and reference time are computed around a reference orbit which is different from the reference orbit used for computing transfer maps (§32.13). The energy and time reference orbit for an element is stored in

```
ele%time_ref_orb_in        ! Reference orbit at upstream end
ele%time_ref_orb_out       ! Reference orbit at downstream end
```

Generally `ele%time_ref_orb_in` is the zero orbit. The exception comes when an element is a `super_slave`. In this case, the reference orbit through the super_slaves of a given `super_lord` is constructed to be continuous. This is done for consistancey sake. For example, to ensure that when a marker is superimposed on top of a wiggler the reference orbit, and hence the reference time, is not altered.

32.15 EM Field Maps

The electromagnetic field of an element may be described using one or more field map representations. These are stored in four pointer array components of the `ele_struct`:

```
ele%cartesian_map(:)     ! Cartesian harmonic field decomposition
ele%cylindrical_map(:)   ! Cylindrical harmonic field decomposition
ele%gen_gradients(:)      ! Generalized gradient field model
ele%grid_field(:)        ! Sampled-grid field (DC or AC)
```

Since each of these is a pointer array, its association status must be tested before accessing any subcomponents. Multiple field map types may be present simultaneously on a single element. Full documentation of these structures and the underlying field models is given in Chapter 18.

32.16 Wakes

The `ele%wake` component holds information on the wakes associated with an element. Since `ele%wake` is a pointer, its association status must be tested before any of its sub–components are accessed.

```
type (ele_struct) ele
...
if (associated(ele%wake)) then
  ...
```

Bmad observes the following rule: If `%wake` is associated, it is assumed that all the sub–components (`%wake%sr_table`, etc.) are associated. This simplifies programming in that you do not have to test directly the association status of the sub–components.

See §20 for the equations used in wakefield calculations. wakefields are stored in the %wake struct:

```
type wake_struct
  type (wake_sr_struct) :: sr     ! Short-range wake.
  type (wake_lr_struct) :: lr     ! Long-range wake.
  real(rp) :: amp_scale = 1       ! Wake amplitude scale factor.
  real(rp) :: time_scale = 1      ! Wake time scale factor.
end type
```

The short–range wake parameterization uses pseudo–modes (§20.1). This parameterization utilizes the %wake%sr%long(:), and %wake%sr%trans(:) arrays for the longitudinal and transverse modes respectively. The structure used for the elements of these arrays are:

```
type wake_sr_mode_struct  ! Pseudo-mode short-range wake struct
  real(rp) amp        ! Amplitude
  real(rp) damp       ! Damping factor.
  real(rp) freq       ! Frequency in Hz
  real(rp) phi        ! Phase in radians/2pi
  real(rp) b_sin      ! non-skew sin-like component of the wake
  real(rp) b_cos      ! non-skew cos-like component of the wake
  real(rp) a_sin      ! skew sin-like component of the wake
  real(rp) a_cos      ! skew cos-like component of the wake
end type
```

The wakefield kick is calculated from Eq. (20.6). %amp, %damp, %freq, and %phi are the input parameters from the lattice file. the last four components (%norm_sin, etc.) store the accumulated wake: Before the bunch passes through these are set to zero and as each particle passes through the cavity the contribution to the wake due to the particle is calculated and added the components.

%wake%z_sr_mode_max is the maximum z value beyond which the pseudo mode representation is not valid. This is set in the input lattice file.

The %wake%lr array stores the long–range wake modes. The structure definition is:

```
type wake_lr_struct   ! Long-Range Wake struct
  real(rp) freq       ! Actual Frequency in Hz
  real(rp) freq_in    ! Input frequency in Hz
  real(rp) R_over_Q   ! Strength in V/C/m^2
  real(rp) Q          ! Quality factor
  real(rp) angle      ! polarization angle (radians/2pi).
  integer m           ! Order (1 = dipole, 2 = quad, etc.)
  real(rp) b_sin      ! non-skew sin-like component of the wake
  real(rp) b_cos      ! non-skew cos-like component of the wake
  real(rp) a_sin      ! skew sin-like component of the wake
  real(rp) a_cos      ! skew cos-like component of the wake
  logical polarized   ! Polarized mode?
end type
```

This is similar to the sr_mode_wake_struct. %freq_in is the actual frequency in the input file. bmad_parser will set %freq to %freq_in except when the lr_freq_spread attribute is non-zero in which case bmad_parser will vary %freq as explained in §4.31. %polarized is a logical that indicates whether the mode has a polarization angle. If so, then %angle is the polarization angle.

32.17 Wiggler and Undulator Fields

Wiggler and undulator elements (wiggler and undulator element types) define their magnetic fields using the standard field map components of the ele_struct (see §32.15): %cartesian_map, %cylindrical_map,

%gen_gradients, or %grid_field. The field decomposition models are described in detail in Chapter 18.

32.18 Spin Transport

The linear spin transport through an element is stored in the %spin_q(0:3,0:6) array, which gives the zeroth and first order quaternion transfer map. The reference orbit about which this map is computed is the same as %map_ref_orb_in/%map_ref_orb_out (§32.13).

For higher-order spin transport, the full spin Taylor map is stored in %spin_taylor(0:3), a quaternion representation as an array of four taylor_struct elements. The reference orbit for the spin Taylor map may differ from the orbital Taylor map reference orbit; it is stored separately in %spin_taylor_ref_orb_in(6).

32.19 Radiation Maps

The %rad_map pointer stores radiation damping and stochastic kick matrices for an element. It is of type rad_map_ele_struct:

```
type rad_map_ele_struct
  type (rad_map_struct) rm0, rm1  ! Upstream and downstream half matrices.
  logical :: stale = .true.
end type
```

Each rad_map_struct contains:

```
type rad_map_struct
  real(rp) :: ref_orb(6)           ! Reference orbit for the map.
  real(rp) :: damp_dmat(6,6)       ! Damping correction matrix.
  real(rp) :: xfer_damp_vec(6)     ! 0th order transfer vector with damping.
  real(rp) :: xfer_damp_mat(6,6)   ! 1st order matrix with damping.
  real(rp) :: stoc_mat(6,6)        ! Stochastic (Cholesky-decomposed) kick matrix.
end type
```

Since ele%rad_map is a pointer, its association status must be tested before use.

32.20 High-Energy Space Charge

The %high_energy_space_charge pointer stores the beam parameters needed for high-energy space charge calculations. It is of type high_energy_space_charge_struct:

```
type high_energy_space_charge_struct
  type (coord_struct) closed_orb   ! Beam orbit
  real(rp) kick_const
  real(rp) sig_x, sig_y, sig_z
  real(rp) phi                     ! Rotation angle, lab to rotated frame.
  real(rp) sin_phi, cos_phi
end type
```

Since ele%high_energy_space_charge is a pointer, its association status must be tested before accessing sub-components.

32.21 Converter, Foil, and RF Components

Three pointer components of the `ele_struct` are used by specialized element types:

`%converter` is used by elements that convert one particle species to another (for example, a positron converter target in a linac). It points to a `converter_struct` holding output species information and a table of particle distributions as a function of incoming energy and material thickness.

`%foil` is used by `foil` elements and points to a `foil_struct` describing the material composition of the foil.

`%rf` is used by RF cavity elements to store the stair-step energy model (`rf_ele_struct`). This model divides the cavity into a sequence of energy kicks interspersed with drifts. Note that `ele%rf` is not allocated for slice and super slaves.

Since all three are pointers, their association status must be tested before use.

32.22 Multipoles

The multipole components of an element (See §18.1) are stored in the pointers `%a_pole(:)` and `%b_pole(:)`. If `%a_pole` and `%b_pole` are allocated they always have a range `%a_pole(0:n_pole_maxx)` and `%b_pole(0:n_pole_maxx)`. Currently `n_pole_maxx` = 20. For a `Multipole` element, the `%a_pole(n)` array stores the integrated multipole strength `KnL`, and the `%b_pole(n)` array stores the tilt `Tn`.

A list of *Bmad* routines for manipulating multipoles can be found in §44.28.

32.23 Tracking Methods

A number of `ele_struct` components control tracking and transfer map calculations. These are:

```
%mat6_calc_method           ! Integer switch: bmad_standard$, taylor$, etc.
%tracking_method            ! Integer switch: bmad_standard$, taylor$, linear$, etc.
%spin_tracking_method       ! Integer switch: bmad_standard$, symp_lie_ptc$, etc.
%field_calc                 ! Integer switch: bmad_standard$, no_field$, fieldmap$, etc.
%csr_method                 ! Integer switch: off$, one_dim$, steady_state_3d$, etc.
%space_charge_method        ! Integer switch: off$, slice$, fft_3d$, cathode_fft_3d$, etc.
%symplectify                ! Logical: symplectify mat6 after calculation?
%multipoles_on              ! Logical: include multipole components?
%taylor_map_includes_offsets ! Logical: Taylor map includes element misalignments?
%is_on                      ! Logical: element on or off?
%offset_moves_aperture      ! Logical: does element offset move the aperture?
```

Note that `%csr_method` and `%space_charge_method` are `integer` switch components (not logicals). See Chapter §37 for more details.

32.24 Custom and General Use Attributes

There are four components of an `ele_struct` that are guaranteed to never be used by any *Bmad* routine and so are available for use by someone writing a program. These components are:

```
%r(:,:,:)                   ! real(rp), pointer.
```

```
%custom(:)                  ! real(rp), pointer.
%ix_pointer                 ! integer.
%logic                      ! logical.
```

Values for `ele%r` and `ele%custom` can be set in the lattice file (§3.9). If values are set for `ele%r` or `ele%custom`, these arrays will be expanded in size if needed.

Accessing the `ele%custom` array should be done using the standard accessor routines (§34.4). For example, if the lattice file being used defines a custom attribute called `rise_time` (§3.9):

```
parameter[custom_attribute1] = "rise_time"
```

then a program can access the `rise_time` attribute via:

```
type (ele_struct), poiner :: ele
type (all_pointer_struct) a_ptr
...
ele => ...            ! Point ele to some element in the lattice
call pointer_to_attribute (ele, "RISE_TIME", .true., a_ptr, err_flag)
print *, "RISE_TIME Attribute has value:", a_ptr%r
```

Note: Even if there are custom attributes associated with a given type of element (say, all quadrupoles), a given element of that type may not have its `%custom(:)` array allocated. [In this case, none of the custom values have been set so are zero by definition.] In the above example, the `%custom` array will be allocated if needed in the call to `pointer_to_attribute`.

If not defined through a lattice file, custom attributes can also be defined directly from within a program using the set_custom_attribute_name routine. For example:

```
logical err_flag
...
call set_custom_attribute_name ('QUADRUPOLE::ERROR_CURRENT', err_flag)
```

Note: When there is a superposition (§8), the `super_slave` elements that are formed do *not* have any custom attributes assigned to them even when their `super_lord` elements have custom attributes. This is done since the *Bmad* bookkeeping routines are not able to handle the situation where a `super_slave` element has multiple `super_lord` elements and thus the custom attributes from the different `super_lord` elements have to be combined. Proper handling of this situation is left to any custom code that a program implements to handle custom attributes.

32.25 Bmad Reserved Variables

A number of `ele_struct` components are reserved for *Bmad* internal use only. These are:

```
%ixx
%iyy
%izz
%bmad_logic
%select
```

To avoid conflict with multiple routines trying to use these components simultaneously, these components are only used for short term bookkeeping within individual routines.

Chapter 33

The lat_struct

The `lat_struct` is the structure that holds of all the information about a lattice (§2.3). The components of a `lat_struct` are listed in Fig. 33.1.

```
type lat_struct
  character(40) use_name                    ! Name in USE statement
  character(40) lattice                     ! Lattice name
  character(80) input_file_name             ! Lattice input file name
  character(80) title                       ! From TITLE statement
  character(100), allocatable :: print_str(:) ! Saved print statements
  type (expression_atom_struct), allocatable :: constant(:)  ! Constants defined in the lattice
  type (mode_info_struct) a, b, z             ! Tunes, etc.
  type (lat_param_struct) param               ! Parameters
  type (bookkeeping_state_struct) lord_state  ! lord bookkeeping status.
  type (ele_struct) ele_init                  ! For use by any program
  type (ele_struct), pointer ::  ele(:) => null()  ! Array of elements [=> branch(0)].
  type (branch_struct), allocatable :: branch(:)   ! Branch(0:) array
  type (control_struct), allocatable :: control(:) ! Control list
  type (photon_reflect_surface_struct), pointer :: surface(:) => null()
  type (coord_struct) particle_start      ! Starting coords
  type (pre_tracker_struct) pre_tracker   ! For OPAL/IMPACT-T
  type (nametable_struct) nametable       ! For quick searching by element name.
  real(rp), allocatable :: custom(:)      ! Custom attributes.
  integer version                         ! Version number
  integer n_ele_track                     ! Number of lat elements to track through.
  integer n_ele_max                       ! Index of last valid element in %ele(:) array
  integer n_control_max                   ! Last index used in control_array
  integer n_ic_max                        ! Last index used in ic_array
  integer input_taylor_order              ! As set in the input file
  integer, allocatable :: ic(:)           ! Index to %control(:)
  integer :: photon_type = incoherent$    ! Or coherent$. For X-ray simulations.
  integer :: creation_hash                ! Number to determine if lattice is different.
end type
```

Figure 33.1: Definition of the `lat_struct`.

The %ele_init component within the lat_struct is not used by *Bmad* and is available for general program use.

33.1 Initializing

Normally initialization of a lat_struct lattice is done by bmad_parser when a lattice file is parsed and does not have to be done by the programmer. When a programmer needs to initialize a lattice, however, init_lat is used to initialize the lattice with a single branch. After this initial setup, the routines allocate_branch_array and allocate_lat_ele_array can be used to set up additional branches. Example:

```
type (lat_struct) lat
...
call init_lat (lat, 1000)           ! Branch(0) has 1001 elements.
call allocate_branch_array (lat, 2) ! Allocate Branch(1) and Branch(2).
call allocate_lat_ele_array (lat, 20, 1)  ! Branch(1) has 21 elements
call allocate_lat_ele_array (lat, 30, 2)  ! Branch(2) has 31 elements.
```

33.2 Pointers: Allocation and Deallocation

Since the lat_struct has pointers within it, there is an extra burden on the programmer to make sure that allocation and deallocation is done properly. To this end, the equal sign has been overloaded by the routine lat_equal_lat so that when one writes

```
type (lat_struct) lattice1, lattice2
! ... some calculations ...
lattice1 = lattice2
```

the pointers in the lat_struct structures will be handled properly. The result will be that lattice1 will hold the same information as lattice2 with all the lattice elements in the same place but the pointers in lattice1 will point to different locations in physical memory so that changes to one lattice will not affect the other.

Initial allocation of the pointers in a lat_struct instance is generally handled by the bmad_parser and lat_equal_lat routines. Once allocated, local lat_struct variables must have the save attribute or the pointers within must be appropriately deallocated before leaving the routine.

```
type (lat_struct), save :: lattice     ! Either do this at the start or ...
...
call deallocate_lat_pointers (lattice) ! ... Do this at the end.
```

Using the save attribute will generally be faster but will use more memory. Typically using the save attribute will be the best choice.

33.3 Branches in the lat_struct

The lattice is divided up into the "root branch" (§7.7) and, if there are fork or photon_fork elements, a number "forked" branches.

The branches of a lattice is contained in the lat%branch(0:) array. The %branch(0:) array is always indexed from 0 with the 0 branch being a root branch. The definition of the branch_struct structure is

	Element index	
section	*min*	*max*
tracking	0	%n_ele_track
control	%n_ele_track+1	%n_ele_max

Table 33.1: Bounds of the tracking and control parts of the root branch (lat%branch(0)%ele(:)) array.

```
type branch_struct
  character(40) name
  integer ix_branch                    ! Index in lat%branch(:) array.
  integer ix_from_branch               ! -1 => No forking element to this branch.
  integer ix_from_ele                  ! Index of forking element
  integer, pointer :: n_ele_track      ! Number of tracking elements
  integer, pointer :: n_ele_max
  type (mode_info_struct), pointer :: a, b, z
  type (ele_struct), pointer :: ele(:)
  type (lat_param_struct), pointer :: param
  type (wall3d_struct), pointer :: wall3d(:)
  type (ptc_branch1_info_struct) ptc
  type (normal_form_struct) normal_form_with_rf, normal_form_no_rf
end type
```

The value of the %branch(i)%ix_branch conponent is the branch index and will thus have the value i. This can be useful when passing a branch to a subroutine. The %branch(i)%ix_from_branch component gives the branch index of the branch that the i^{th} branch branched off from. %branch(i)%ix_from_ele gives the index in the %branch(j)%ele(:) array of the fork or photon_fork element that marks the beginning of the i^{th} branch. Example:

```
type (lat_struct), target :: lat
type (ele_struct), pointer :: ele
...
ib = lat%branch(3)%ix_from_branch
ie = lat%branch(3)%ix_from_ele
! ele is the fork or photon_fork element for lat%branch(3)
ele => lat%branch(ib)%ele(ie)
! This is the same as the above.
ele => pointer_to_ele(lat%branch(3)%ix_from_branch, lat%branch(3)%ix_from_ele)
```

The %branch%ele(:) array holds the array of elements in the branch. Historically, the lat_struct was developed at the start of the *Bmad* project and branches were implemented well after that. To maintain compatibility with older code, the following components point to the same memory blocks

```
lat%ele(:)              <--->  lat%branch(0)%ele(:)
lat%n_ele_track         <--->  lat%branch(0)%n_ele_track
lat%n_ele_max           <--->  lat%branch(0)%n_ele_max
lat%param               <--->  lat%branch(0)%param
```

All %branch%ele(:) arrays are allocated with zero as the lower bound. The %ele(0) element of all branches is an beginning_ele element with its %name component set to "BEGINNING". %ele(0)%mat6 is always the unit matrix. For the root branch, the %branch(0)%ele(0:) array is divided up into two parts: The "tracking" part and a "control" part (also called the "lord" part). The tracking part of this array holds the elements that are tracked through. The control part holds elements that control attributes of other elements (§33.5). The bounds of these two parts is given in Table 33.1. Only the root

branch has a lord section so %branch%n_ele_track and %branch%n_ele_max are the same for all other branches. Since the root branch can also be accessed via the lat%ele(:) array, code that deals with the lord section of the lattice may use lat%ele(:) in place of lat%branch(0)%ele(:).

for a given fork or photon_fork element, the index of the branch that is being forked to and the index of the element that is being forked to is stored in:

```
ix_branch  = nint(branch_ele%value(ix_branch_to$))  ! branch index
ix_element = nint(branch_ele%value(ix_element_to$)) ! element index
direction  = nint(branch_ele%value(direction$))
```

The direction will be +1 for forward forking and -1 for backward forking.

33.4 Param_struct Component

The %param component within each lat%branch(:) is a lat_param_struct structure whose definition is shown in Fig. 33.2 This structure would be more aptly named branch_param_struct but is named otherwise for historical reasons.

%param%total_length is the length of the branch that a beam tracks through defined by

```
%param%total_length = %ele(n_ele_track)%s - %ele(0)%s
```

Normally %ele(0)%s = 0 so %param%total_length = %ele(n_ele_track)%s but this is not always the case.

%param%n_part is the number of particles in a bunch and is used in various calculations. Historically, this parameter has been used to set the number of strong beam particle with BeamBeam elements but it is strongly recommended to use the beambeam element's n_particle parameter instead.

For closed branches, %param%t1_with_RF and %param%t1_no_RF are the 1–turn transfer matrices from the start of the branch to the end. %param%t1_with_RF is the full transfer matrix with RF on. %param%t1_no_RF is the transverse transfer matrix with RF off. %param%t1_no_RF is used to compute the Twiss parameters. When computing the Twiss parameters %param%stable is set according to whether the matrix is stable or not. If the matrix is not stable the Twiss parameters cannot be computed. If unstable, %param%unstable_factor will be set to the growth rate per turn of the unstable mode.

```
type lat_param_struct
  real(rp) n_part             ! Particles/bunch.
  real(rp) total_length       ! total_length of lattice
  real(rp) unstable_factor    ! closed branch: growth rate/turn.
                              !   all branches: |orbit/limit|
  real(rp) t1_with_RF(6,6)    ! Full 1-turn 6x6 matrix
  real(rp) t1_no_RF(6,6)      ! Transverse 1-turn 4x4 matrix (RF off).
  integer particle            ! +1 = positrons, -1 = electrons, etc.
  integer geometry            ! open$, etc...
  integer ixx                 ! Integer for general use
  logical stable              ! For closed branch. Is lat stable?
  type (bookkeeper_status_struct) bookkeeping_state
                                      ! Overall status for the branch.
end type
```

Figure 33.2: Definition of the param_struct.

Besides being set when the 1–turn transfer matrix is calculated, %param%unstable_factor will be set if a particle is lost in tracking to:

```
orbit_amplitude / limit - 1
```

The particle type for a branch is stored in the integer variable %param%particle. The value of this variable will encode for a fundamental particle, atom, or molecule. See the file particle_species_mod.f90 for more details. If the particle corresponds to a fundamental particle, %param%particle will correspond to one of the following constants:

```
electron$,     positron$,
muon$,         antimuon$,
proton$,       antiproton$,
photon$,       pion_0$,
pion_minus$,   pion_plus$
deuteron$      deuteron_0$
```

To print the name of the particle use the function species_name. A particles mass and charge can be obtained from the functions mass_of and charge_of respectively. charge_of returns the particle's charge in units of e. Example:

```
type (lat_struct) lat
...
print *, "Beam Particles are: ", species_name(lat%param%particle)
if (lat%param%particle == proton$) print *, "I do not like protons!"
print *, "Particle mass (eV):    ", mass_of(lat%param%particle)
print *, "Particle charge:       ", charge_of(lat%param%particle)
```

33.5 Elements Controlling Other Elements

In the lat_struct structure, certain elements in the %ele(:) array (equivalent to the %branch(0)%ele(:) array), called lord elements, can control the attributes (component values) of other %branch(:)%ele(:) elements. Elements so controlled are called slave elements. The situation is complicated by the fact that a given element may simultaneously be a lord and a slave. For example, an overlay element (§4.40) is a lord since it controls attributes of other elements but an overlay can itself be controlled by other overlay and group elements. In all cases, circular lord/slave chains are not permitted.

The lord and slave elements can be divided up into classes. What type of lord an element is, is set by the value of the element's ele%lord_status component. Similarly, what type of slave an element is is set by the value of the element's ele%slave_status component. Nomenclature note: An element may be referred to by it's %lord_status or %slave_status value. For example, an element with ele%lord_status set to super_lord$ can be referred to as a "super_lord" element.

The value of the ele%lord_status component can be one of:

super_lord$
: A super_lord element is created when elements are superimposed on top of other elements (§8). super_lords (along with multipass_lords), are called major lords since the attribute values of a super_slave are entirely determined by the attribute values of the super_lord(s) of the slave.

girder_lord$
: A girder_lord element is a girder element (§4.24). That is, the element will have ele%key = girder$.

multipass_lord$
: multipass_lord elements are created when multipass lines are present (§9). multipass_lords

(along with super_lords) are called major lords since most of the attribute values of a multipass_slave are entirely determined by the attribute values of the multipass_lord of the slave. The few exceptions are parameters like phi0_multipass which can be set for individual slave elements.

overlay_lord$
An overlay_lord is an overlay element (§4.40). That is, such an element will have ele%key = overlay$.

ramper_lord$
A ramper_lord is a ramper element (§4.44). That is, such an element will have ele%key = ramper$. Note that a ramper_lord will not have pointers to its slaves. That is, ele%n_slave will be zero.

group_lord$
A group_lord is a group element (§4.26). That is, such an element will have ele%key = group$.

not_a_lord$
This element does not control anything.

Any element whose %lord_status is something other than not_a_lord$ is called a lord element. In the tracking part of the branch (§33.3), %lord_status will always be not_a_lord$. In the lord section of the branch, under normal circumstances, there will never be any not_a_lord elements. However, it is permissible, and sometimes convenient, for programs to set the %lord_status of a lord element to not_a_lord$.

The possible values for the ele%slave_status component are:

multipass_slave$
A multipass_slave element is the slave of a multipass_lord (§9).

slice_slave$
A slice_slave element represents a longitudinal slice of another element. Slice elements are not part of the lattice but rather are created on-the-fly when, for example, a program needs to track part way through an element.

super_slave$
A super_slave element is an element in the tracking part of the branch that has one or more super_lord lords (§8).

minor_slave$
A minor_slave element is an element that is not a slice_slave and does not have a major lord. Major lords are super_lords and multipass_lords. A minor_slave element will some have attributes that are controlled by overlay_lords, group_lords, girder_lords, or ramper_lords.

free$
A free element is one that has no lords except for perhaps ramper_lords. Additionally, there still might be field overlap from other elements.

super_slave elements always appear in the tracking part of the branch. The other types can be in either the tracking or control parts of the branch.

Only some combinations of %lord_status values and %slave_status values are permissible for a given element. Table 33.2a lists the valid combinations. Thus, for example, it is *not* possible for an element to be simultaneously a super_lord and a super_slave.

ele%slave_status	ele%lord_status: not_a_lord$	group_lord$	girder_lord$	overlay_lord$	multipass_lord$	super_lord$
free$	X	X	X	X	X	X
minor_slave$	X	X	X	X	X	X
multipass_slave$	X					X
slice_slave$	X					
super_slave$	X					

(a) Possible ele%lord_status and ele%slave_status combinations within an individual element.

slave%slave_status	lord%lord_status: not_a_lord$	group_lord$	girder_lord$	overlay_lord$	multipass_lord$	super_lord$
free$						
minor_slave$		X	X	X		
multipass_slave$		X		X	1	
super_slave$		X		X		X

(b) Possible %lord_status and %slave_status combinations for any lord/slave pair.

Table 33.2: Possible %lord_status/%slave_status combinations. "X" marks a possible combination. "1" indicates that the slave will have exactly one lord of the type given in the column.

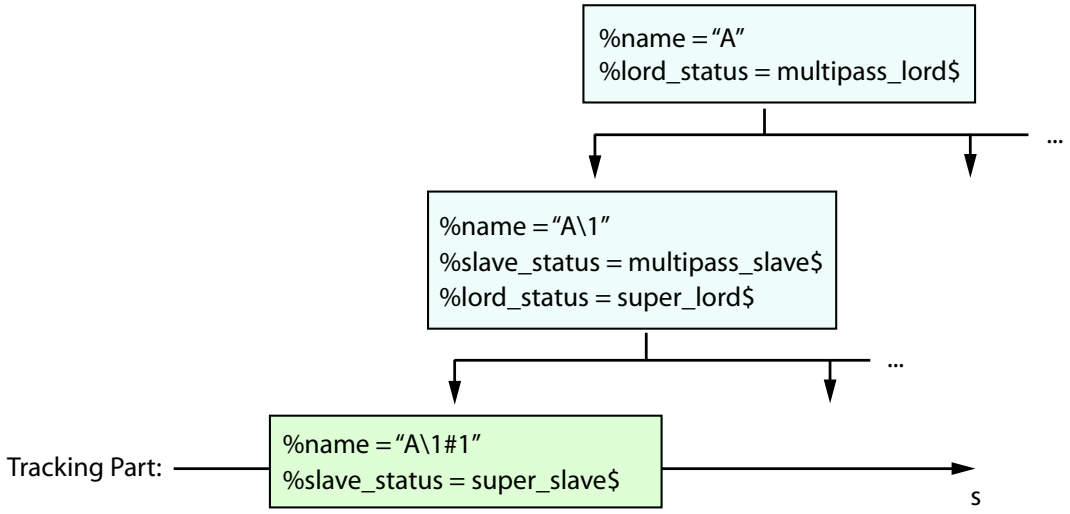

Figure 33.3: Example of multipass combined with superposition. A `multipass_lord` element named `A` controls a set of `multipass_slaves` (only one shown). The `multipass_slave` elements are also `super_lord` elements and they will control `super_slave` elements in the tracking part of the branch.

For lord/slave pairs, Table 33.2b lists the valid combinations of `%lord_status` values in the lord element and `%slave_status` values in the slave element. Thus, for example, a `super_slave` may only be controlled by a `super_lord`. In the example in Section §9, element `A` would be a `multipass_lord` and `A\1` and `A\2` would be `multipass_slaves`. When superposition is combined with multipass, the elements in the tracking part of the branch will be `super_slaves`. These elements will be controlled by `super_lords` which will also be `multipass_slaves` and these `super_lord`/`multipass_slave` elements will be controlled by `multipass_lords`. This is illustrated in Fig. 33.3.

The number of slave elements that a lord controls is given by the value of the lord's `%n_slave` component. Additionally, the number of lord elements that the slave has is given by the value of the slave's. `%n_lord` component. To find the slaves and lords of a given element, use the routines pointer_to_slave and pointer_to_lord. Example:

```
type (lat_struct), target :: lat
type (ele_struct), pointer :: this_ele, lord_ele, slave_ele
...
this_ele => lat%ele(321)    ! this_ele points to a given element in the lattice

do i = 1, this_ele%n_lord   ! Loop over all lords of this_ele
  ! lord_ele points to the i^th lord element of this_ele
  lord_ele => pointer_to_lord (this_ele, i)
  ...
enddo

do i = 1, this_ele%n_slave  ! Loop over all slaves of this_ele
  ! slave_ele points to the i^th slave element of this_ele
  slave_ele => pointer_to_slave (this_ele, i)
  ...
enddo
```

For non-ramper elements, the lord/slave bookkeeping is bidirectional. That is, for any given element, call it `this_ele`, consider the i^{th} lord:

```
lord_ele_i => pointer_to_lord (this_ele, i)
```

then there will always be some index j such that the element pointed to by

```
pointer_to_slave(lord_ele_i, j)
```

is the original element `this_ele`. The same is true for the slaves of any given element. That is, for the i^{th} slave

```
slave_ele_i => pointer_to_slave (this_ele, i)
```

there will always be some index j such that the element pointed to by

```
pointer_to_lord(slave_ele_i, j)
```

The following ordering of slaves and lords is observed:

Slaves of a super_lord:
: The associated `super_slave` elements of a given `super_lord` element are ordered from the entrance end of the `super_lord` to the exit end. That is, in the code snippet above, `pointer_to_slave (this_ele, 1)` will point to the slave at the start of the `super_lord` and `pointer_to_slave (this_ele, this_ele%n_lord)` will point to the slave at the exit end of the `super_lord`.

Slaves of a multipass_lord:
: The associated `multipass_slave` elements of a `multipass_lord` element are ordered by pass number. That is, in the code snippet above, `pointer_to_slave (this_ele, i)` will point to the slave of the i^{th} pass.

Lord of a multipass_slave:
: A `multipass_slave` will have exactly one associated `multipass_lord` and this lord will be the first one. That is, `pointer_to_lord (this_ele, 1)`.

The element control information is stored in the `lat%control(:)` array. Each element of this array is a `control_struct` structure

```
type control_struct
  type (expression_atom_struct), allocatable :: stack(:) ! Evaluation stack
  type (lat_ele_loc_struct) slave ! Slave location
  type (lat_ele_loc_struct) lord  ! Lord location
  integer ix_attrib               ! index of controlled attribute
end type
```

Each element in the `lat%control(:)` array holds the information on one lord/slave pair. The `%lord` component gives the location of the lord element which is always in the root branch — branch 0. The `%slave` component give the element location of the slave element. The `%stack` and `%ix_attrib` components are used to store the arithmetic expression and attribute index for `overlay` and `group` control. The appropriate control_struct for a given lord/slave pair can be obtained from the optional fourth argument of the pointer_to_lord and pointer_to_slave functions. Example: The following prints a list of the slaves, along with the attributes controlled and coefficients, on all group elements in a lattice.

```
type (lat_struct), target :: lat
type (ele_struct), pointer :: lord, slave
type (control_struct), pointer :: con
...
do i = lat%n_ele_track+1, lat%n_ele_max  ! loop over all lords
  lord => lat%ele(i)
  if (lord%lord_status = group_lord$) then
    print *, "Slaves for group lord: ", lord%name
    do j = 1, lord%n_slave
      slave => pointer_to_slave (lord, j, con)
      attrib_name = attribute_name (slave, con%ix_attrib)
      print *, i, slave%name, attrib_name, con%coef
    enddo
  endif
enddo
```

The elements in the `lat%control(:)` array associated with the slaves of a given lord are in the same order as the slaves and the index of the associated `lat%control(:)` element of the first slave is given by the `%ix1_slave` component of the lord Example:

```
type (lat_struct), target :: lat
type (ele_struct), pointer :: lord, slave
type (control_struct), pointer :: con1, con2
...
lord => lat%ele(i)                    ! Point to some lord
do j = 1, lord%n_slave
  slave => pointer_to_slave (lord, j, con1)
  con2 => lat%control(lord%ix1_slave+j-1) ! con1 and con2 are the same.
enddo
```

Except for a `slice_slave`, the `%ic1_lord`, `%n_lord`, and `%n_lord_field` components of a given slave element, along with the `lat%ic(:)` array, can be used to find the lords of the slave. Simplified, the code for the pointer_to_lord function is:

```
function pointer_to_lord (slave, ix_lord, con, ...) result (lord_ptr)
  implicit none
  type (lat_struct), target :: lat
  type (ele_struct) slave
  type (ele_struct), pointer :: lord_ptr
  type (control_struct), pointer, optional :: control
  integer ix_lord, icon
  !
  icon = lat%ic(slave%ic1_lord + ix_lord - 1)
  lord_ptr => lat%ele(lat%control(icon)%lord%ix_ele)
  if (present(con)) con => lat%control(icon)
end function
```

This method for finding the lords of an element is considered "private". That is, no code outside of the official *Bmad* library should rely on this.

`slice_slave` element bookkeeping has is different depending upon whether the element being sliced is a `super_slave` or not. If the element being sliced is a `super_slave`, a `slice_slave` element that is created is, for bookkeeping purposes, considered to be a slave of the `super_slave`'s lords. In this case, the bookkeeping is exactly the same as that of any `super_slave`, and `pointer_to_lord` will return a pointer to one of the `super_slave`'s lords.

On the other hand, if a non `super_slave` element is being sliced, the `%lord` pointer component of the `slice_slave` will be set to point to the element being sliced.

33.6 Lattice Bookkeeping

The term "lattice bookkeeping" refers to the updating of the appropriate parameter values when a given parameter in the lattice is changed. For example, if the accelerating gradient of an `lcavity` element is modified, the reference energy parameter of all elements downstream of the `lcavity` will need to be changed and this can also alter the transfer maps of the `lcavity` and downstream elements. *Bmad* divides the lattice bookkeeping into a "core" part and everything else. The core part itself is divided into five parts:

Attribute bookkeeping
: This refers to intra-element dependent attribute bookkeeping (§5.1).

Control bookkeeping
: This refers to Lord/Slave bookkeeping for `overlay` (§4.40) and `group` (§4.26)elements, and for `superposition` (§8) and `multipass` (§9) lords.

Floor Position bookkeeping This refers to bookkeeping to keep track of an elements global "floor" position stored in the `ele%floor` structure.

Length bookkeeping This refers to bookkeeping to keep track of the longitudinal s-position of an element stored in the `ele%s` component.

Reference Energy bookkeeping This refers to the reference energy assigned to each element (§38.7). `ele%value(E_tot$)` and `ele%value(p0c$)`

Lattice elements have a `bookkeeper_status` component which is of type `bookkeeper_status_struct` which looks like

```
type bookkeeper_status_struct
  integer attributes      ! Intra element dependent attribute status
  integer control         ! Lord/slave bookkeeping status
  integer floor_position  ! Global (floor) geometry status
  integer length          ! Longitudinal position status
  integer ref_energy      ! Reference energy status
  integer mat6            ! Linear transfer map status
  integer rad_int         ! Radiation integrals cache status
end type
```

All components of this structure give the status of some lattice bookkeeping aspect. The first five components of this structure correspond to the five core bookkeeping parts discussed above. The other two components are discussed below.

Possible values for the status components are

```
super_ok$
ok$
stale$
```

The `set_flags_for_changed_attribute` routine sets the appropriate status components of an element to `stale$` which marks that element for the appropriate bookkeeping. When the bookkeeping is done by `lattice_bookkeeper`, the `stale$` status components are set to `ok$`. The `super_ok$` value is reserved for use by any program that needs to do its own custom bookkeeping. How this works is as follows: The *Bmad* bookkeeping routines will never convert a status component with value `super_ok$` to `ok$` without first doing some needed bookkeeping. Thus if a program sets a status component to `super_ok$` and then later on finds that the status component is set to `ok$`, the program knows that bookkeeping has been done. An example will make this clear. Suppose a program needs to keep track of a collection of high order transfer maps between various points in a lattice. Suppose that the constant calculation of these maps would slow the program done so it is desired to recalculate a given map only when necessary. To implement this, the program could set the `ele%status%mat6` attribute of all the element to `super_ok$` when the maps are calculated. If the program subsequently finds a `ele%status%mat6` attribute of an element set to `ok$` it knows that it should recalculate any transfer maps that span that element.

It is guaranteed that when `lattice_bookkeeper` is run, all five core status components will not be `stale$`. The routines used by `lattice_bookkeeper` are:

```
attribute_bookkeeper      ! Intra-element attributes
control_bookkeeper        ! Lord/slave control
s_calc                    ! Longitudinal element s-position
lat_geometry              ! Global (floor) positions.
lat_compute_ref_energy_and_time  ! Reference energy
```

In general, these routines should not be called directly since the correct way to do things is not always straight forward. See the code for `lattice_bookkeeper` for more details.

After the core bookkeeping is done, a program can call lat_make_mat6 to remake the transfer matrices. `lat_make_mat6` will remake the transfer matrices if either the `ele%status%mat6` flag is `stale$` or the reference orbit around which the existing transfer matrix was computed has shifted. `lat_make_mat6` will set the `ele%status%mat6` flag to `ok$` for all elements whose transfer matrices are recomputed.

33.7 Intelligent Bookkeeping

Historically, as the code for lattice bookkeeping (§33.6) was being developed calls to bookkeeping routines were added to calculational routines such as the tracking routine track1 and the routine for calculating the linear transfer map make_mat6. This "automatic" bookkeeping system is inefficient since there is no good way to keep track of what element attributes have been modified which leads to redundant bookkeeping calculations. Eventually, as *Bmad* developed and became more complicated, it was found that the unnecessary bookkeeping load was generally causing a significant slowdown in program execution time — even in programs where no element attributes were changed. To avoid this, an "intelligent" bookkeeping system was developed.

The automatic bookkeeping system has been retired and old programs need to be upgraded if needed. Rule: A program that does not "directly modify" element attributes does not have to modified. Modification of element attributes via *Bmad* routines (for example, using the `set_on_off` routine) is "indirect". A direct modification is something like the following appearing in the program:

```
lat%ele(ie)%value(hkick$) = ...
```

To use intelligent bookkeeping, a program must set the global `bmad_com%auto_bookkeepper` to false. This only needs to be done once at the start of the program before `bmad_parser` is called. If lattice parameters are not modified in the program, this is the only thing that needs to be done.

When a set of attributes needs to be modified, the set_flags_for_changed_attribute routine must be called for each element attribute that is set. After all the attributes have been set, lattice_bookkeeper is called to do the core bookkeeping. Example

```
type (lat_struct) lat
...
bmad_com%auto_bookkeeper = .false.    ! Done once. Put this before the call to bmad_parser.
...
lat%ele(i)%value(gradient$) = 1.05e6  ! Change, say, the gradient of an RFCavity
call set_flags_for_changed_attribute (lat%ele(i), lat%ele(i)%value(gradient$))

... Set attributes of other elements ...

call lattice_bookkeeper (lat)         ! Do once after all attribute sets done.
```

The argument list for `set_flags_for_changed_attribute` is

```
set_flags_for_changed_attribute (ele, attribute)
```

The `attribute` argument may be either real, integer, or logical.

The `set_flags_for_changed_attribute` routine sets flags in the `ele%status` structure §33.6.

Historical note: Up through the end of 2025, to keep things back compatible with existing programs, the automatic bookkeeping system was set as the default. The intelligent bookkeeping could be switched on by setting the parameter:

```
bmad_com%auto_bookkeeper = .false.   ! Default was .true.
```

Currently, intelligent bookkeeping is assumed and the `auto_bookkeeper` parameter is ignored by *Bmad*.

33.8 particle_start Component

The `lat%particle_start` component is a `coord_struct` structure for holding the information obtained from `particle_start` statements (§10.2) in a *Bmad* lattice file.

This component is not used in any standard *Bmad* calculation. It is up to an individual program to use as desired. Use init_coord to initalize a particle position. Example:

```
type (lat_struct) lat
type (coord_struct) orbit_start
...
call bmad_parser("lat.bmad", lat)  ! Read in a lattice.
...
! orbit_start is initalized for tracking from the beginning of the lattice
call init_coord(orbit_start, lat%particle_start, lat%ele(0), downstream_end$)
```

33.9 Custom Parameters

Custom parameters defined for the lattice as a whole (§3.9 are stored in `lat%custom`. The following shows how to print a table of the custom parameters

```
type (lat_struct) lat
character(80) aname
...
if (allocated(lat%custom)) then
```

```
    do i = 1, size(lat%custom)
      aname = attribute_name(def_parameter$, i+custom_attribute0$)
      if (aname(1:1) == "!") cycle    ! Ignore non-existant parameters
      print "(a, es12.4)", "  parameter[" // trim(aname) // "] = ", lat%custom(i)
    enddo
  endif
```

Chapter 34

Lattice Element Manipulation

34.1 Creating Element Slices

It is sometimes convenient to split an element longitudinally into "slices" that represent a part of the element. This is complicated by the fact that elements are not necessarily uniform. For example, map type wigglers are nonuniform and bend elements have end effects. Furthermore, attributes like `hkick` need to be scaled with the element length.

To create an element slice, the routine create_element_slice can be used. Example:

```
type (ele_struct) ele, sliced_ele
...
sliced_ele = ele
sliced_ele%value(l$) = l_slice ! Set the sliced element's length
call create_element_slice (sliced_ele, ele, l_start, param, ...)
```

See the documentation on `create_element_slice` for more details (§30.3).

34.2 Adding and Deleting Elements From a Lattice

Modifying the number of elements in a lattice involves a bit of bookkeeping. To help with this there are a number of routines.

The routine remove_eles_from_lat is used to delete elements from a lattice.

For adding elements there are three basic routines: To add a lord element, the new_control routine is used. To add a new element to the tracking part of the lattice, use the insert_element routine. Finally, to split an element into two pieces, the routine split_lat is used. These basic routines are then used in such routines as create_overlay that creates overlay elements, create_group which creates group elements, add_superimpose which superimposes elements, etc. Example:

```
type (lat_struct), target :: lat
type (ele_struct), pointer :: g_lord, slave

type (control_struct) con(1)
integer ix, n
logical err_flag
```

```
...
call new_control (lat, ix)
g_lord => lat%ele(ix)
allocate (ele%control_var(1))
ele%control_var(1)%name = "A"
call reallocate_expression_stack(con(1)%stack, 10))
call expression_string_to_stack ('3.2*A^2', con(1)%stack, n, err_flag)
con(1)%ix_attrib = k1$
call lat_ele_locator ('Q1W', lat, eles)
con(1)%slave = ele_to_lat_loc(eles(1)%ele)
call create_group (g_lord, con, err_flag)
```

This example constructs a group element with one variable with name `A` controlling the `K1` attribute of element `Q1W` using the expression "$3.2 \cdot A^2$" where `A` is the name of the control variable.

For constructing `group` elements (but not `overlay` elements), the controlled attribute (set by `con(1)%ix_attrib` in the above example) can be set to, besides the set of element attributes, any one in the following list:

```
accordion_edge$  ! Element grows or shrinks symmetrically
start_edge$      ! Varies element's upstream edge s-position
end_edge$        ! Varies element's downstream edge s-position
s_position$      ! Varies element's overall s-position. Constant length.
```

See Section §4.26 for the meaning of these attributes

34.3 Finding Elements

The routine lat_ele_locator can be used to search for an element in a lattice by name or key type or a combination of both. Example:

```
type (lat_struct) lat
type (ele_pointer_struct), allocatable :: eles(:)
integer n_loc; logical err
...
call lat_ele_locator ("quad::skew*", lat, eles, n_loc, err)
print *, "Quadrupole elements whose name begins with the string "SKEW":"
print *, "Name                 Branch_index        Element_index"
do i = 1, n_loc  ! Loop over all elements found to match the search string.
  print *, eles(i)%ele%name, eles(i)%ele%ix_branch, eles(i)%ele%ix_ele
enddo
```

This example finds all elements where `ele%key` is `quadrupole$` and `ele%name` starts with "`skew`". See the documentation on `lat_ele_locator` for more details on the syntax of the search string.

The `ele_pointer_struct` array returned by `lat_ele_locator` is an array of pointers to `ele_struct` elements

```
type ele_pointer_struct
  type (ele_struct), pointer :: ele
end type
```

The `n_loc` argument is the number of elements found and the `err` argument is set True on a decode error of the search string.

Once an element (or elements) is identified in the lattice, it's attributes can be altered. However, care must be taken that an element's attribute can be modified (§5.1). The function `attribute_free` will check if an attribute is free to vary.

```
type (lat_struct) lat
integer ix_ele
...
call lat_ele_locator ('Q10W', lat, eles, n_loc, err)   ! look for an element "Q10W"
free = attribute_free (eles(i)%ele, "K1", lat, .false.)
if (.not. free) print *, "Cannot vary k1 attribute of element Q10W"
```

34.4 Accessing Named Element Attributes

A "named" parameter of the `ele_struct` structure is a parameter that has an associated name that can be used in a lattice file. For example, the quadrupole strength is named `K1` (§4.43). This parameter is stored in the `ele%value(:)` array. Specifically at `ele%value(k1$)`.

Historically, named parameters where always accessed directly but this has proved to be somewhat problematical for a number of reasons. For one, something like `ele%value(k1$)` will always have a value even if the associated lattice element does not have an associated `K1` parameter (For example, a sextupole does not have a `K1` parameter). Another issue involves allocation since components like `ele%a_pole(:)` are pointers that are not necessarily allocated.

To get around some of these issues, accessor functions have been developed for all non-character based named attributes. These accessor functions are:

```
pointer_to_attribute    ! pointer_to_attribute
pointers_to_attribute   ! pointers_to_attribute
set_ele_attribute       ! set_ele_attribute
value_of_attribute      ! value_of_attribute
```

The workhorse is `pointer_to_attribute` that returns a pointer to the appropriate attribute. The returned pointer argument is actually an instance of an `all_pointer_struct` which looks like:

```
type all_pointer_struct
  real(rp), pointer :: r => null()
  integer, pointer :: i => null()
  logical, pointer :: l => null()
end type
```

When the `all_pointer_struct` argument is returned, one (or zero if the attribute name is not recognized) of the pointer components will be associated. For example:

```
type (ele_struct) ele
type (all_pointer_struct) attrib_ptr
...
call pointer_to_attribute (ele, "A3_ELEC", .true., attrib_ptr, err)
attrib_ptr%r = 0.34
call attribute_set_bookkeeping (ele, "A3_ELEC", err_flag, attrib_ptr)
call lattice_bookkeeper (lat)  ! Bookkeeping needed due to parameter change
```

Also see the example program in §31.2.

The `set_ele_attribute` routine is useful when there is user input since this routine can evaluate expressions. For example:

```
type (lat_struct) lat
type (ele_pointer_struct), allocatable :: eles(:)
integer n_loc, n
logical err_flag, make_xfer_mat
...
```

```
call lat_ele_locator ('Q01W', lat, eles, n_loc, err_flag)
do n = 1, n_loc
  call set_ele_attribute (eles(n)%ele, "K1 = 0.1*c_light", lat, err_flag)
enddo
```

This example sets the `K1` attribute of all elements named `Q01W`. `set_ele_attribute` checks whether an element is actually free to be varied and sets the `err_flag` logical accordingly. An element's attribute may not be freely varied if, for example, the attribute is controlled via an `Overlay`.

Chapter 35

Reading and Writing Lattices

35.1 Reading in Lattices

There are two subroutines in *Bmad* to read in a *Bmad* standard lattice file: bmad_parser and bmad_parser2. `bmad_parser` is used to initialize a `lat_struct` (§33) structure from scratch using the information from a lattice file. Unless told otherwise, after reading in the lattice, `bmad_parser` will compute the 6x6 transfer matrices for each element and this information will be stored in the `digested file` (§3.2) that is created. Notice that `bmad_parser` does *not* compute any Twiss parameters.

bmad_parser2 is typically used after `bmad_parser` if there is additional information that needs to be added to the lattice. For example, consider the case where the aperture limits for the elements is stored in a file that is separate from the main lattice definition file and it is undesirable to put a `call` statement in one file to reference the other. To read in the lattice information along with the aperture limits, there are two possibilities: One possibility is to create a third file that calls the first two:

```
! This is a file to be called by bmad_parser
call, file = ""lattice_file""
call, file = ""aperture_file""
```

and then just use `bmad_parser` to parse this third file. The alternative is to use `bmad_parser2` so that the program code looks like:

```
! program code to read in everything
type (lat_struct) lat
call bmad_parser ('lattice_file', lat)       ! read in a lattice.
call bmad_parser2 ('aperture_file', lat)     ! read in the aperture limits.
```

35.2 Digested Files

Since parsing can be slow, once the `bmad_parser` routine has transferred the information from a lattice file into the `lat_struct` it will make what is called a digested file. A digested file is an image of the `lat_struct` in binary form. When `bmad_parser` is called, it first looks in the same directory as the lattice file for a digested file whose name is of the form:

```
""digested_"" // LAT_FILE
```

where `LAT_FILE` is the lattice file name. If `bmad_parser` finds the digested file, it checks that the file is not out–of–date (that is, whether the lattice file(s) have been modified after the digested file is made). `bmad_parser` can do this since the digested file contains the names and the dates of all the lattice files that were involved. Also stored in the digested file is the "*Bmad* `version number`". The *Bmad* version number is a global parameter that is increased (not too frequently) each time a code change involves modifying the structure of the `lat_struct` or `ele_struct`. If the *Bmad* version number in the digested file does not agree with the number current when `bmad_parser` was compiled, or if the digested file is out–of–date, a warning will be printed, and `bmad_parser` will reparse the lattice and create a new digested file.

Since computing Taylor Maps can be very time intensive, `bmad_parser` tries to reuse Taylor Maps it finds in the digested file even if the digested file is out–of–date. To make sure that everything is OK, `bmad_parser` will check that the attribute values of an element needing a Taylor map are the same as the attribute values of a corresponding element in the digested file before it reuses the map. Element names are not a factor in this decision.

This leads to the following trick: If you want to read in a lattice where there is no corresponding digested file, and if there is another digested file that has elements with the correct Taylor Maps, then, to save on the map computation time, simply make a copy of the digested file with the digested file name corresponding to the first lattice.

read_digested_bmad_file write_digested_bmad_file The digested file is in binary format and is not human readable but it can provide a convenient mechanism for transporting lattices between programs. For example, say you have read in a lattice, changed some parameters in the `lat_struct`, and now you want to do some analysis on this modified `lat_struct` using a different program. One possibility is to have the first program create a digested file

```
call write_digested_bmad_file ('digested_file_of_mine', lat)
```

and then read the digested file in with the second program

```
call read_digested_bmad_file ('digested_file_of_mine', lat)
```

An alternative to writing a digested file is to write a lattice file using `write_bmad_lattice_file`

35.3 Writing Lattice files

write_bmad_lattice_file To create a *Bmad* lattice file from a `lat_struct` instance, use the routine write_bmad_lattice_file. *MAD*–8, *MAD*–X, or `SAD` compatible lattice files can be created from a `lat_struct` variable using the routine write_lattice_in_foreign_format:

```
type (lat_struct) lat             ! lattice
...
call bmad_parser (bmad_lat_file, lat)              ! Read in a lattice
call write_lattice_in_foreign_format ("lat.mad", "MAD-8", lat)  ! create MAD file
```

Information can be lost when creating a *MAD* or `SAD` file. For example, neither *MAD* nor `SAD` has the concept of things such as `overlays` and `groups`.

Chapter 36

Normal Modes: Twiss Parameters, Coupling, Emittances, Etc.

36.1 Components in the Ele_struct

The `ele_struct` (§32) has a number of components that hold information on the Twiss parameters, dispersion, and coupling at the exit end of the element. The Twiss parameters of the three normal modes (§23.1) are contained in the `ele%a`, `ele%b`, and `ele%z` components which are of type `twiss_struct`:

```
type twiss_struct
  real(rp) beta         ! Twiss Beta function
  real(rp) alpha        ! Twiss Alpha function
  real(rp) gamma        ! Twiss gamma function
  real(rp) phi          ! Normal mode Phase advance
  real(rp) eta          ! Normal mode dispersion
  real(rp) etap         ! Normal mode momentum dispersion.
  real(rp) deta_ds      ! Dispersion derivative
  real(rp) sigma        ! Normal mode beam size
  real(rp) sigma_p      ! Normal mode beam size derivative
  real(rp) emit         ! Geometric emittance
  real(rp) norm_emit    ! Energy normalized emittance (= β  γ  ϵ)
end type
```

The projected horizontal and vertical dispersions in an `ele_struct` are contained in the `ele%x` and `ele%y` components. These components are of type `xy_disp_struct`:

```
type xy_disp_struct
  real(rp) eta     ! Projected dispersion
  real(rp) etap    ! Projected momentum dispersion
  real(rp) deta_ds ! Projected dispersion derivative dη_x/ds or dη_y/ds.
end type
```

Section §23.4 discussed the relationship between `etap` and `deta_ds`.

The components `ele%emit`, `ele%norm_emit`, `ele%sigma`, `ele%sigma_p` are not set by the standard *Bmad* routines and are present for use by any program.

The relationship between the projected and normal mode dispersions are given by Eq. (23.16). The 2x2 coupling matrix $\mathbf{C}$ (Eq. (23.5)) is stored in the `ele%c_mat(2,2)` component of the `ele_struct` and the

γ factor of Eq. (23.5) is stored in the `ele%gamma_c` component. There are several routines to manipulate the coupling factors. For example:

```
c_to_cbar(ele, cbar_mat)              ! Form Cbar(2,2) matrix
make_v_mats(ele, v_mat, v_inv_mat)    ! Form V coupling matrices.
```

See §44.24 for a complete listing of such routines.

Since the normal mode and projected dispersions are related, when one is changed within a program the appropriate change must be made to the other. To make sure everything is consistent, the set_flags_for_changed_attribute routine can be used. Example:

```
type (lat_struct), target :: lat
real(rp), pointer :: attrib_ptr
...
attrib_ptr => lat%ele(ix_ele)%value(k1$) ! Point to some attribute.
attrib_ptr = value                       ! Change the value.
call set_flags_for_changed_attribute (lat%ele(ix_ele), attrib_ptr)
```

The `%mode_flip` logical component of an `ele_struct` indicates whether the a and b normal modes have been flipped relative to the beginning of the lattice. See Sagan and Rubin[Sagan99] for a discussion of this. The convention adopted by *Bmad* is that the `%a` component of all the elements in a lattice will all correspond to the same physical normal mode. Similarly, the `%b` component of all the elements will all correspond to some (other) physical normal mode. That is, at an element where there is a mode flip (with `%mode_flip` set to True), the `%a` component actually corresponds to the **B** matrix element in Eq. (23.3) and vice versa. The advantage of this convention is that calculations of mode properties (for example the emittance), can ignore whether the modes are flipped or not.

The normal mode analysis of Sagan and Rubin, while it has the benefit of simplicity, is strictly only applicable to lattices where the RF cavities are turned off. The full 6-dimensional analysis is summarized by Wolski[Wolski06]. The normal_mode3_calc routine perform the full analysis. The results are put in the `%mode3` component of the `ele_struct` which is of type `mode3_struct`:

```
type mode3_struct
  real(rp) v(6,6)
  type (twiss_struct) a, b, c
  type (twiss_struct) x, y
end type
```

The 6-dimensional `mode3%v(6,6)` component is the analog of the 4-dimensional **V** matrix appearing in Eq. (23.2).

36.2 Tune and Twiss Parameter Calculations

A calculation of the Twiss parameters starts with the Twiss parameters at the beginning of the lattice. For linear machines, these Twiss parameters are generally set in the input lattice file (§10.4). For circular machines, the routine twiss_at_start may be used (§10.4)

```
type (lat_struct) lat
...
if (lat%param%geometry == closed$) call twiss_at_start(lat)
```

In either case, the initial Twiss parameters are placed in `lat%ele(0)`. The tune is placed in the variables `lat%a%tune` and `lat%b%tune`.

To propagate the Twiss, coupling and dispersion parameters from the start of the lattice to the end, the routine, twiss_propagate_all can be used. This routine works by repeated calls to twiss_propagate1 which does a single propagation from one element to another. The Twiss propagation depends upon

the transfer matrices having already computed (§37). `twiss_propagate_all` also computes the Twiss parameters for all the lattice branches.

Before any Twiss parameters can be calculated, the transfer matrices stored in the lattice elements must be computed. bmad_parser does this automatically about the zero orbit. If, to see nonlinear effects, a different orbit needs to be used for the reference, The routine twiss_and_track can be used. For example:

```
type (lat_struct) lat
type (coord_struct), allocatable :: orbit(:)
call bmad_parser ('my_lattice', lat)
call twiss_and_track (lat, orb, ok)
```

Once the starting Twiss parameters are set, twiss_propagate_all can be used to propagate the Twiss parameters to the rest of the elements

The routine twiss_and_track_at_s can be used to calculate the Twiss parameters at any given longitudinal location. Alternatively, to propagate the Twiss parameters partially through a given element use the routine twiss_and_track_intra_ele.

36.3 Tune Setting

The routine set_tune can be used to set the transverse tunes:

```
set_tune (phi_a_set, phi_b_set, dk1, lat, orb_, ok)
```

`set_tune` varies quadrupole strengths until the desired tunes are achieved. As input,`set_tune` takes an argument `dk1(:)` which is an array that specifies the relative change to be make to the quadrupoles in the lattice.

To set the longitudinal (synchrotron) tune, the routine set_z_tune can be used. set_z_tune works by varying rf cavity voltages until the desired tune is achieved.

36.4 Emittances & Radiation Integrals

See Section §22.3 for details on the radiation integral formulas.

The routine radiation_integrals is used to calculate the normal mode emittances along with the radiation integrals:

```
type (lat_struct) lat
type (normal_modes_struct) modes
type (rad_int_all_ele_struct) ele_rad_int
...
call radiation_integrals (lat, orbit, modes, rad_int_by_ele = ele_rad_int)
```

The `modes` argument, which is of type `normal_modes_struct`, holds the radiation integrals integrated over the entire lattice.

```
type normal_modes_struct
  real(rp) synch_int(0:3) ! Synchrotron integrals I0, I1, I2, and I3
  real(rp) sigE_E         ! SigmaE/E
  real(rp) sig_z          ! Sigma_Z
  real(rp) e_loss         ! Energy loss / turn (eV)
  real(rp) rf_voltage     ! Total rfcavity voltage (eV)
```

```
    real(rp) pz_aperture    ! pz aperture limit
    type (anormal_mode_struct)  a, b, z
    type (linac_normal_mode_struct) lin
  end type
```

In particular, the %a, %b, and %z components, which are of type anormal_mode_struct hold the emittance values:

```
  type anormal_mode_struct
    real(rp) emittance        ! Beam emittance
    real(rp) synch_int(4:6)   ! Synchrotron integrals
    real(rp) j_damp           ! damping partition number
    real(rp) alpha_damp       ! damping per turn
    real(rp) chrom            ! Chromaticity
    real(rp) tune             ! "Fractional" tune in radians
  end type
```

The ele_rad_int argument, which is is of type rad_int_all_ele_struct, holds the radiation integrals on an element-by-element basis.

```
  type rad_int_all_ele_struct
    type (rad_int1_struct), allocatable :: ele(:) ! Array is indexed from 0
  end type
```

36.5 Chromaticity Calculation

For a circular lattice, chrom_calc calculates the chromaticity by calculating the tune change with change in beam energy.

chrom_tune sets the chromaticity by varying the sextupoles. This is a very simple routine that simply divides the sextupoles into two families based upon the local beta functions at the sextupoles.

Chapter 37

Tracking and Transfer Maps

37.1 The coord_struct

The `coord_struct` holds the coordinates of a particle The definition of the `coord_struct` is

```
  type coord_struct
    real(rp) vec(6)     ! (x, px, y, py, z, pz)
    real(rp) s          ! Longitudinal position.
    real(rp) t          ! Absolute time (not relative to reference).
    real(rp) spin(3)    ! (x, y, z) Spin vector
    real(rp) field(2)   ! Photon (x, y) field intensity.
    real(rp) phase(2)   ! Photon (x, y) phase.
    real(rp) charge     ! charge in a particle (Coul).
    real(rp) dt_dref    ! path length (used by coherent photons).
    real(rp) r          ! For general use. Not used by Bmad.
    real(rp) p0c        ! For non-photons: Reference momentum. Negative -> going backwards.
                        !     For photons: Photon momentum (not reference).
    real(rp) beta       ! Velocity / c_light.
    integer ix_ele      ! Index of the lattice element the particle is in.
                        !   May be -1 or -2 if element is not associated with a lattice.
    integer ix_branch   ! Index of the lattice branch the particle is in.
    integer ix_user     ! Not used by Bmad
    integer state       ! alive$, lost$, lost_neg_x$, etc.
    integer direction   ! +1 or -1. Sign of longitudinal direction of motion (ds/dt).
                        !  This is independent of the element orientation.
    integer time_dir    ! +1 or -1. Time direction. -1 => Traveling backwards in time.
    integer species     ! Positron$, proton$, etc.
    integer location    ! upstream_end$, inside$, or downstream_end$
end type
```

Definitions:

Direction of Travel
: The "`direction of travel`", also called the "`direction of motion`" is the direction that the particle is moving in when traveling forward in time.

Propagation Direction

The "`propagation direction`" is the direction that a particle will be propagated in during tracking. The propagation direction will be in the same direction as the direction of travel when propagating a particle forward in time and will be opposite the direction of travel when propagating a particle backwards in time.

Reverse Tracking
: "`Reverse tracking` refers to tracking a particle with `%direction` set to -1. That is, tracking in the reverse direction longitudinally. The opposite to reverse tracking is called "`forward direction`" tracking.

Backwards Tracking
: "`Backwards Tracking` refers to tracking a particle backwards in time. That is, with `%time_dir` = -1. The opposite to backwards tracking is called "`forward time`" tracking.

The components of the `coord_struct`:

%beta
: The normalized velocity v/c is stored in `%beta`. `%beta` is always positive.

%direction
: Longitudinal forward time "direction of travel". A setting of +1 (the default) is in the forward +s (downstream) direction and a setting of -1 is in the reverse -s (upstream) direction (§16.1.3). Notice that the setting of `direction` is independent of the orientation of the lattice element the particle is traveling through. That is, for an element with reversed `orientation` (`ele%orientation` = -1), a particle with `direction` = 1 will be traveling towards the `entrance` end of the element (-z direction in body coordinates) and with `direction` = -1 the particle will be traveling towards the `exit` (+z direction in body coordinates) end (§16.1.3). See `%time_dir`.

%time_dir
: Time direction that a particle is propagated through. A value of +1 (the default) is forward time and a value of -1 is backwards time.

%field_x, %field_y
: The `%field_x` and `%field_y` components are for photon tracking and are in units of field/sqrt(cross-section-area). That is, the square of these units is an intensity. It is up to individual programs to define an overall scaling factor for the intensity if desired.

%ix_branch
: The `%ix_branch` component gives the index of the lattice branch in the `lat%branch(ib)` array that the particle is in.

%ix_ele
: The `%ix_ele` component gives the index of the element in the `lat%branch(ib)%ele(:)` array that the particle is in. If the element is not associated with a lattice, `%ix_ele` is set to -1. When initializing a `coord_struct` (see below), `%ix_ele` will be initialized to `not_set$`.

%ix_user
: The `%ix_user` component is for use by code outside of the *Bmad* library. This component will not be modified by *Bmad*.

%location
: The `%location` component indicates where a particle is longitudinally with respect to the element being tracked. `%location` will be on of:

```
entrance_end$
inside$
exit_end$
```

`entrance_end$` indicates that the particle is at the element's entrance ($-s$) end and `exit_end$` indicates that the particle is at the element's exit ($+s$) end. `inside$` indicates that the particle is in between. If the element has edge fields (for example, the `e1` and `e2` edge fields of a bend), a particle at the `entrance_end$` or `exit_end$` is considered to be just outside the element.

%p0c

For charged-particles, the reference momentum in eV is stored in the `%p0c` component. For photons, `%p0c` is the actual (not reference) momentum. For charged-particles, `%p0c` may be negative if the particle is traveling backwards longitudinally. For photons, `%vec(6)` (β_z) will be negative if the photon is going backward.

%r

The `%r` component is for use by code outside of the *Bmad* library. *Bmad* will not modify this component.

%s

The `%s` component gives the absolute s-position of the particle. When tracking through an element (say with Runge-Kutta tracking), and when the particle coordinates is expressed in element body coordinates (§16.3), the s-position at any point within the element, by definition, is independent of any misalignments the element has as long as the element is not reversed. If the element is reversed, the s-position is reversed as well.

%spin(3)

The `%spin(3)` component gives a particle's (x, y, z) spin vector (§24.1).

%state

The `%state` component will be one of:

```
not_set$
pre_born$
alive$
lost$
lost_neg_x$
lost_pos_x$
lost_neg_y$
lost_pos_y$
lost_z$
lost_pz$
```

The `not_set$` setting indicates that the `coord_struct` has not yet been used in tracking. The `alive$` setting indicates that the particle is alive. If a particle is "dead", the `%state` component will be set to one of the other settings. The `lost_neg_x$` setting indicates that the particle was lost at an aperture on the $-x$ side of the element. The `lost_z$` setting is used to indicate that the particle tried to "turn around". This can happen, for example, with strong magnetic fields or when a particle has been decelerated too much. The reason why the particle is marked lost in this case is due to the fact that s-based tracking algorithms cannot handle particles that reverse direction. The exception is that the `time_runge_kutta` (§6.1) tracking method can handle particle reversal so in this case, particles will not be declared lost if they reverse direction.

The `lost$` setting is used when neither of the other `lost_*$` settings are not appropriate. For example, `lost$` is used in Runge-Kutta tracking when the adaptive step size becomes too small (this may happen if the fields do not obey Maxwell's equations).

To convert the integer value of `%state` to a string that can be printed, use the function coord_state_name

```
type (coord_struct) orbit
print *, "State of the orbit: ", coord_state_name(orbit%state)
```

%t

`%t` gives the absolute time.

%vec(:)

The `%vec(:)` array defines the phase space coordinants (§17.2). Note that for photons, the definition of the phase space coordinates (§17.4) is different from that used for charged particles. The signs of `%vec(2)` and `%vec(4)` are such that, for the signs of the change in `%vec(1)` and `%vec(3)` during propagation will be equal to the product `%direction * %time_dir * sign_of(%vec(2)` and `%direction * %time_dir * sign_of(%vec(2)` respectively.

To initialize a `coord_struct` so it can be used as the start of tracking, the init_coord routine can be used:

```
type (coord_struct) start_orb
real(rp) phase_space_start(6)
...
phase_space_start = [...]
call init_coord (start_orb, phase_space_start, lat%ele(i), lat%param%particle)
```

Here `init_coord` will initialize `start_orb` appropriately for tracking through element `lat%ele(i)` with the particle species set to the species of the reference particle given by `lat%param%particle`.

37.2 Tracking Through a Single Element

track1 is the routine used for tracking through a single element

```
type (coord_struct), start_orb, end_orb
type (ele_struct) ele
real(rp) start_phase_space(6)
logical err
...
start_phase_space = [...]
call init_coord (start_orb, start_phase_space, ele, photon$)
call track1 (start_orb, ele, end_orb, err_flag = err)
if (.not. particle_is_moving_forward(end_orb)) then
  print *, "Particle is lost and gone forever..."
endif
```

To check if a particle is still traveling in the forward direction, the particle_is_moving_forward routine can be used as shown in the above example.

The "virtual" entrance and exit ends of a lattice element are, by definition, where the physical ends of the element would be if there were no offsets. In particular, if an element has a finite `z_offset` (§32.12), the physical ends will be displaced from the virtual ends. The position `ds` of a particle with respect to the physical entrance end of the element is

```
ds = coord%s - (ele%s + ele%value(z_offset_tot$) - ele%value(l$))
```

When tracking through an element, the starting and ending positions always correspond to the virtual ends. If there is a finite `z_offset`, the tracking of the element will involve tracking through drifts just before and just after the tracking of the body of the element so that the particle ends at the proper virtual exit end.

Note: The z phase space component of the orbit (`%vec(5)`) is independent of the value of `ele%ref_time` even though the reference time is used to define z (See Eq. (17.3)). This is true since the starting reference time that is used for a particle is arbitrary. For example, when tracking multiple bunches, the reference time is typically set so that a particle at the center of a bunch has $z = 0$. Also, in a ring, `ele%ref_time` is only the reference time for the first turn through an element. Since *Bmad* does not keep track of turn number, there is no way for *Bmad* to know what the true reference time is other than to calculate it from the value of z!

37.3 Tracking Through a Lattice Branch

When tracking through a lattice branch, one often defines an array of `coord_structs` – one for each element of the lattice branch. In this case, the i^{th} `coord_struct` corresponds to the particle coordinates at the end of the i^{th} element. Since the number of elements in the lattice is not known in advance, the array must be declared to be allocatable. The lower bound of the array must be set to zero to match a `lat%branch(i)%ele(:)` array. The upper bound should be the upper bound of the `%branch(i)%ele(:)` array. The routine reallocate_coord will allocate an array of `coord_structs`:

```
type (coord_struct), allocatable :: orbit(:)
type (lat_struct) lat
...
call reallocate_coord (orbit, lat, ix_branch)
```

Alternatively, the `save` attribute can be used so that the array stays around until the next time the routine is called

```
type (coord_struct), allocatable, save :: orb(:)
```

Saving the `coord_stuct` is faster but leaves memory tied up. Note that in the main program, the `save` attribute is not permitted If a `coord_struct` array is passed to a routine, the routine must explicitly set the lower bound to zero if the array is not declared as allocatable:

```
subroutine my_routine (orbit1, orbit2)
  use bmad
  implicit none
  type (coord_struct), allocatable :: orbit1(:)  ! OK
  type (coord_struct) orbit2(0:)                 ! Also OK
  ...
```

Declaring the array allocatable is mandatory if the array is to be resized or the array is passed to a routine that declares it allocatable.

For an entire lattice, the `coord_array_struct` can be used to define an array of `coord_array` arrays:

```
type coord_array_struct
  type (coord_struct), allocatable :: orb(:)
end type
```

The routine reallocate_coord_array will allocate an `coord_array_struct` instance

```
type (coord_array_struct), allocatable :: all_orbit(:)
type (lat_struct) lat
...
call reallocate_coord_array (all_orbit, lat)
...
```

Once an array of `coord_struct` elements is defined, the track_all routine can be used to track through a given lattice branch

```
type (coord_struct), allocatable :: orbit(:)
```

```
integer ib, track_state
...
ib = 1                        ! Branch to track through
call init_coord(orbit(0), init_phase_space, lat%branch(ib)%ele(0), proton$)
call track_all (lat, orbit, ib, track_state, err_flag)
if (track_state /= moving_forward$) then
  print *, "Particle lost at element:", track_state
  print *, "Aperture lost at: ", coord_state_name(orbit(track_state)%state)
```

After tracking, `orbit(i)` will correspond to the particles orbit at the end of `lat%branch(ib)%ele(i)`.

For routines like `track_all` where an array of `coord_structs` is used, an integer `track_state` argument is provided that is set to `moving_forward$` if the particle survives to the end, or is set to the index of the element at which the particle either hit an aperture or the particle's longitudinal velocity is reversed.

The reason why the reversal of the particle's longitudinal velocity stops tracking is due to the fact that the standard tracking routines, which are s-based (that is, use longitudinal position s as the independent coordinate), are not designed to handle particles that reverse direction. To properly handle this situation, time-based tracking needs to be used (§37.11). Notice that this is different from tracking a particle in the reversed ($-s$) direction.

Alternatively to `track_all`, the routine track_many can be used to track through a selected number of elements or to track backwards (See §37.14).

The `track_all` routine serves as a good example of how tracking works. A condensed version of the code is shown in Fig. 37.1. The call to `track1` (line 18) tracks through one element from the exit end of the $n-1^{st}$ element to the exit end of the n^{th} element.

```
1     subroutine track_all (lat, orbit, ix_branch, track_state, err_flag)
2       use bmad
3       implicit none
4       type (lat_struct), target :: lat
5       type (branch_struct), pointer :: branch
6       type (coord_struct), allocatable :: orbit(:)
7       integer, optional :: ix_branch, track_state
8       logical, optional :: err_flag
9       logical err
10
11      !
12
13      branch => lat%param(integer_option(0, ix_branch))
14      branch%param%ix_track = moving_forward
15      if (present(track_state)) track_state = moving_forward\$
16
17      do n = 1, branch%n_ele_track
18        call track1 (orbit(n-1), branch%ele(n), branch%param, orbit(n), err_flag = err)
19        if (.not. particle_is_moving_forward(orbit(n))) then
20          if (present(track_state)) track_state = n
21          orbit(n+1:)%status = not_set$
22          return
23        endif
24      enddo
25    end subroutine
```

Figure 37.1: Condensed track_all code.

37.4 Forking from Branch to Branch

Tracking from a `fork` or `photon_fork` (§4.23) element to the downstream `branch` is not "automatic". That is, since the requirements of how to handle forking can vary greatly from one situation to the next, *Bmad* does not try to track from one `branch` to the next in any one of its tracking routines.

The discussion here is restricted to the case where the particle being tracked is simply transferred from the forking element to the downstream branch. [Thus the subject of photon generation is not covered here.]

There are two cases discussed here. The first case is when a given branch (called `to_branch`) has an associated forking element in the `from_branch` that forks to the beginning of the `to_branch`. Appropriate code is:

```
type (lat_struct), target :: lat   ! Lattice
type (branch_struct) :: to_branch  ! Given forked-to branch
type (branch_struct), pointer  :: from_branch ! Base branch
type (ele_struct), pointer :: fork_ele
type (coord_struct), allocatable :: from_orbit(:), to_orbit(:)
integer ib_from, ie_from

ib_from  = to_branch%ix_from_branch

if (ib_from < 0) then
  ! Not forked to ...

else
  from_branch => lat%branch(ib_from)
  ie_from = to_branch%ix_from_ele
  fork_ele => from_branch%ele(ie_from)
  to_orbit(0) = from_orbit(ie_from)
  call transfer_twiss (fork_ele, to_branch%ele(0))
endif
```

`from_orbit(0:)` and `to_orbit(0:)` are arrays holding the orbits at the exit end of the elements for the `from_branch` and `to_branch` respectively. The call to transfer_twiss transfers the Twiss values to the `to_branch` which can then be propagated through the `to_branch` using `twiss_propagate_all`.

The second case starts with the `fork_ele` forking element. This is similar to the first case but is a bit more general since here the element, called `to_ele` in the `to_branch` that is connected to `fork_ele` need not be the starting element of `to_branch`.

```
type (lat_struct), target :: lat   ! Lattice
type (branch_struct), pointer :: to_branch  ! forked-to branch
type (ele_struct), pointer :: to_ele
type (coord_struct), allocatable :: from_orbit(:), to_orbit(:)
integer ib_to, ie_to

ib_to  = nint(fork_ele%value(ix_to_branch$))
ie_to  = nint(fork_ele%value(ix_to_element$))

to_branch => lat%branch(ib_to)
to_ele => to_branch%ele(ie_to)
to_orbit(to_ele%ix_ele) = from_orbit(fork_ele%ix_ele)
```

Notice that, by convention, the transferred orbit is located at the exit end of the `to_ele`.

37.5 Multi-turn Tracking

Multi-turn tracking over a branch is simply a matter of setting the coordinates at the beginning zeroth element equal to the last tracked element within a loop:

```
type (lat_struct) lat             ! lattice to track through
type (coord_struct), allocatable :: orbit(:)
...
call reallocate_coord (orbit, lat, ix_branch = 1)
orbit(0)%vec = [0.01, 0.2, 0.3, 0.4, 0.0, 0.0] ! init
do i = 1, n_turns
  call track_all (lat, orbit, 1)
  orbit(0) = orbit(lat%branch(1)%n_ele_track)
end do
```

Often times it is only the root branch, `branch(0)`, that is to be tracked. In this case, the above reduces to

```
type (lat_struct) lat             ! lattice to track through
type (coord_struct), allocatable :: orbit(:)
...
call reallocate_coord (orbit, lat%n_ele_max)
orbit(0)%vec = [0.01, 0.2, 0.3, 0.4, 0.0, 0.0] ! init
do i = 1, n_turns
  call track_all (lat, orbit)
  orbit(0) = orbit(lat%n_ele_track)
end do
```

37.6 Closed Orbit Calculation

For a circular lattice the closed orbit may be calculated using `closed_orbit_calc`. By default this routine will track in the forward direction which is acceptable unless the particle you are trying to simulate is traveling in the reverse direction and there is radiation damping on. In this case you must tell `closed_orbit_calc` to do backward tracking. This routine works by iteratively converging on the closed orbit using the 1–turn matrix to calculate the next guess. On rare occasions if the nonlinearities are strong enough, this can fail to converge. An alternative routine is `closed_orbit_from_tracking` which tries to do things in a more robust way but with a large speed penalty.

37.7 Partial Tracking through elements

There are several routines for tracking partially through an element:

```
twiss_and_track_at_s
twiss_and_track_intra_ele
track_from_s_to_s
twiss_and_track_from_s_to_s
mat6_from_s_to_s
```

These routines make use of element "slices" (§34.1) which are elements that represent some sub-section of an element. There are two routines for creating slices:

```
create_element_slice
create_uniform_element_slice
```

It is important to note that to slice up a given element, the s_to_s tracking routines will not always work. For example, consider the case where a given element is followed by a zero length multipole. If track_from_s_to_s is called with a value for s2 (the value at the end of the track) which corresponds to the exit end of this element, the result will also include tracking through the zero length multipole. Thus, in the case where a given element is to be sliced, one or the other of the two slice routines given above must be first used to create an element slice then this slice can be used for tracking.

37.8 Apertures

The routine check_aperture_limit checks the aperture at a given element. The ele%aperture_type component determines the type of aperture. Possible values for ele%aperture_type are

```
rectangular$
elliptical$
custom$
```

With custom$, a program needs to be linked with a custom version of check_aperture_limit_custom.

The logical bmad_com%aperture_limit_on determines if element apertures (See §5.8) are used to determine if a particle has been lost in tracking. The default bmad_com%aperture_limit_on is True. Even if this is False there is a "hard" aperture limit set by bmad_com%max_aperture_limit. This hard limit is used to prevent floating point overflows. The default hard aperture limit is 1000 meters. Additionally, even if a particle is within the hard limit, some routines will mark a particle as lost if the tracking calculation will result in an overflow.

lat%param%lost is the logical to check to see if a particle has been lost. lat%param%ix_lost is set by track_all and gives the index of the element at which a particle is lost. %param%end_lost_at gives which end the particle was lost at. The possible values for lat%param%end_lost_at are:

```
entrance_end$
exit_end$
```

When tracking forward, if a particle is lost at the exit end of an element then the place where the orbit was outside the aperture is at orbit(ix) where ix is the index of the element where the particle is lost (given by lat%param%ix_lost). If the particle is lost at the entrance end then the appropriate index is one less (remember that orbit(i) is the orbit at the exit end of an element).

To tell how a particle is lost, check the lat%param%plane_lost_at parameter. Possible values for this are:

```
x_plane$
y_plane$
z_plane$
```

x_plane$ and y_plane$ indicate that the particle was lost either horizontally, or vertically. z_plane$ indicates that the particle was turned around in an lcavity element. That is, the cavity was decelerating the particle and the particle did not not have enough energy going into the cavity to make it to the exit.

37.9 Custom Tracking

Custom code can be used for tracking. This is discussed in detail in sections §38.2 and §38.3.

37.10 Tracking Methods

For each element the method of tracking may be set either via the input lattice file (see §6.1) or directly in the program by setting the `%tracking_method` attribute of an element

```
type (ele_struct) ele
...
ele%tracking_method = symp_lie_ptc$  ! for symp_lie_ptc, tracking
print *, "Tracking_method: ", calc_method_name(ele%tracking_method)
```

To form the corresponding parameter to a given tracking method just put "$" after the name. For example, the `bmad_standard` tracking method is specified by the `bmad_standard$` parameter. To convert the integer `%tracking_method` value to a string suitable for printing, use the `tracking_method_name` array.

It should be noted that except for `linear` tracking, none of the *Bmad* tracking routines make use of the `ele%mat6` transfer matrix. The reverse, however, is not true. The transfer matrix routines (`lat_make_mat6`, etc.) will do tracking.

For determining what tracking methods are valid for a given element, use valid_tracking_method and valid_mat6_calc_method functions

```
print *, "Method is valid: ", valid_tracking_method(ele, symp_lie_ptc$)
```

Bmad simulates radiation damping and excitation by applying a kick just before and after each element.

37.11 Using Time as the Independent Variable

Time tracking uses time as the independent variable as opposed to the standard s based tracking. Time tracking is useful when a particle's trajectory can reverse itself longitudinally. For example, low energy particles generated when a relativistic particle hits the vacuum chamber wall are good candidates for time tracking.

Currently, the only `ele%tracking_method` available for time tracking is `time_runge_kutta$`. Time tracking needs extra bookkeeping due to the fact that the particle may reverse directions. See the `dark_current_tracker` program as an example.

Note: Using time as the independent variable can be used with both absolute and relative time tracking (§26.1).

37.12 Absolute/Relative Time Tracking

Absolute or relative time tracking (§26.1) can be set after the lattice file is parsed, by setting the `%absolute_time_tracking` component of the `lat_struct`. when `%absolute_time_tracking` is toggled, the autoscale_phase_and_amp must be called to reset the appropriate phase offsets and scale amplitudes.

37.13 Taylor Maps

A list of routines for manipulating Taylor maps is given in §44.38. The order of the Taylor maps is set in the lattice file using the `parameter` statement (§10.1). In a program this can be overridden using the

routine set_ptc. The routine taylor_coef can be used to get the coefficient of any given term in a Taylor map.

```
type (taylor_struct) t_map(6)
...
print *, "out(4)=coef * in(1)^2:", taylor_coef(t_map(4), 1, 1)
print *, "out(4)=coef * in(1)^2:", taylor_coef(t_map(4), [2,0,0,0,0,0])
```

Transfer Taylor maps for an element are generated as needed when the `ele%tracking_method` or `ele%mat6_calc_method` is set to `Symp_Lie_Bmad`, `Symp_Lie_PTC`, or `Taylor`. Since generating a map can take an appreciable time, *Bmad* follows the rule that once generated, these maps are never regenerated unless an element attribute is changed. To generate a Taylor map within an element irregardless of the `ele%tracking_method` or `ele%mat6_calc_method` settings use the routine ele_to_taylor. This routine will kill any old Taylor map before making any new one. To kill a Taylor map (which frees up the memory it takes up) use the routine kill_taylor.

To test whether a `taylor_struct` variable has an associated Taylor map. That is, to test whether memory has been allocated for the map, use the Fortran associated function:

```
type (bmad_taylor) taylor(6)
...
if (associated(taylor(1)%term)) then  ! If has a map ...
  ...
```

To concatenate the Taylor maps in a set of elements the routine concat_taylor can be used

```
type (lat_struct) lat          ! lattice
type (taylor_struct) taylor(6)  ! taylor map
...
call taylor_make_unit (taylor)  ! Make a unit map
do i = i1+1, i2
  call concat_taylor (taylor, lat%ele(i)%taylor, taylor)
enddo
```

The above example forms the transfer Taylor map starting at the end of element `i1` to the end of element `i2`. Note: This example assumes that all the elements have a Taylor map. The problem with concatenating maps is that if there is a constant term in the map "feed down" can make the result inaccurate (§25.1. To get around this one can "track" a taylor map through an element using symplectic integration.

```
type (lat_struct) lat          ! lattice
type (taylor_struct) taylor(6)  ! taylor map
...
call taylor_make_unit (taylor)  ! Make a unit map
do i = i1+1, i2
  call call taylor_propagate1 (taylor, lat%ele(i), lat%param)
enddo
```

Symplectic integration is typically much slower than concatenation. The width of an integration step is given by `%ele%value(ds_step$`. The attribute `%ele%value(num_steps$)`, which gives the number of integration steps, is a dependent variable (§5.1) and should not be set directly. The order of the integrator (§25.1) is given by `%ele%integrator_order`. PTC (§29) currently implements integrators of order 2, 4, or 6.

37.14 Tracking Backwards

Tracking backwards happens when a particle goes in the direction of decreasing `s`. This is indicated in the `coord_struct` by `coord%direction = -1`.

The `time_runge_kutta` `tracking_method` is able to handle the situation where a particle would reverse direction due to string electric or magnetic fields. All other tracking methods are not able to handle this since they are position (s) based, instead of time based. With non `time_runge_kutta` tracking methods, the equations of motion become singular when a particle "tries" to reverse direction. In such a situation, the particle will be marked as lost and the `coord_struct` will have `%status` $/=$ `alive$`.

The "problem" with tracking backwards is that the reference time $t_0(s)$ that is used to compute the z phase space coordinate (Eq. (17.3)) is independent of the motion of any particle. That is, a particle traveling backwards will have a large negative z. As an alternative to tracking backwards, reversing the lattice and tracking forwards is possible (§37.15).

One restriction with backwards tracking is that, for simplicity's sake, *Bmad* does not compute transfer matrices for propagation in the backwards direction. Tracking with reversed elements does not have this restriction.

37.15 Reversed Elements and Tracking

With a lattice element that is reversed (`s:ele.reverse`), the transfer map and transfer matrix that is stored in the element is, just like for a non-reversed element, appropriate for a particle traveling in the $+s$ direction.

37.16 Beam (Particle Distribution) Tracking

Tracking with multiple particles is done with a `beam_struct` instance:

```
type beam_struct
  type (bunch_struct), allocatable :: bunch(:)
end type
```

A `beam_struct` is composed of an array of bunches of type `bunch_struct`:

```
type bunch_struct
  type (coord_struct), allocatable :: particle(:)
  integer, allocatable :: ix_z(:)  ! bunch%ix_z(1) is index of head particle, etc.
  real(rp) charge_tot  ! Total charge in bunch (Coul).
  real(rp) charge_live ! Total charge of live particles in bunch (Coul).
  real(rp) z_center    ! Longitudinal center of bunch (m). Note: Generally, z_center of
                       !   bunch #1 is 0 and z_center of the other bunches is negative.
  real(rp) t_center    ! Center of bunch creation time relative to head bunch.
  integer species      ! electron$, proton$, etc.
  integer ix_ele       ! Element this bunch is at.
  integer ix_bunch     ! Bunch index. Head bunch = 1, etc.
end type
```

The `bunch_struct` has an array of particles of type `coord_struct` (§37.1).

Initializing a `beam_struct` to conform to some initial set of Twiss parameters and emittances is done using the routine init_beam_distribution:

```
type (lat_struct) lat
type (beam_init_struct) beam_init
type (beam_struct) beam
...
call init_beam_distribution (lat%ele(0), lat%param, beam_init, beam)
```

The `lat%ele(0)` argument, which is of type `ele_struct`, gives the twiss parameters to initialize the beam to. In this case, we are starting tracking from the beginning of the lattice. The `beam_init` argument which is of type `beam_init` gives additional information, like emittances, which is needed to initialize the beam. See chapter §12 for more details.

Tracking a beam is done using the track_beam routine

```
type (lat_struct) lat
type (beam_struct) beam
...
call track_beam (lat, beam)
```

or, for tracking element by element, track1_bunch can be used.

For analyzing a bunch of particles, that is, for computing such things as the sigma matrix from the particle distribution, the calc_bunch_params routine can be used.

Notice that when a particle bunch is tracked to a given longitudinal position in the lattice, all the particles of the bunch are at that longitudinal position (this is no different if particles are tracked individually independent of the bunch). Given that the bunch has a non-zero bunch length, the current time $t(s)$ associated with the particles will be different for different particles (See Eq. (17.3)). If it is desired to reconstruct the shape of the bunch at *constant time*, each particle must be tracked either forward or backwards by an appropriate amount. Since this tracking generally involves only very short distances, it is usually acceptable to ignore any fields and to propagate the particles as if they were in a field free region.

37.17 Spin Tracking

See Section §6.3 for a list of spin tracking methods available. To turn spin tracking on, use the `bmad_com%spin_tracking_on` flag. `ele%spin_tracking_method` sets the method used for spin tracking. After properly initializing the spin in the `coord_struct`, calls to `track1` will track both the particle orbit and the spin.

The Sokolov-Ternov effect[Barber99] is the self-polarization of charged particle beams due to asymmetric flipping of a particle's spin when the particle is bent in a magnetic field. Whether this effect is included in a simulation is determined by the setting of `bmad_com%spin_sokolov_ternov_flipping_on`. Also, spin flipping will *not* be done if spin tracking is off or both radiation damping and excitation are off.

37.18 X-ray Targeting

X-rays can have a wide spread of trajectories resulting in many "doomed" photons that hit apertures or miss the detector with only a small fraction of "successful" photons actually contributing to the simulation results. The tracking of doomed photons can therefore result in an appreciable lengthening of the simulation time. To get around this, *Bmad* can be setup to use what is called "targeting" to minimize the number of doomed photons generated.

This is explained in detail in §27.5. The coordinates of the four or eight corner points and the center target point are stored in:

```
gen_ele%photon%target%corner(:)%r(1:3)
gen_ele%photon%target%center%r(1:3)
```

where `gen_ele` is the generating element (not the element with the aperture).

Chapter 38

Miscellaneous Programming

38.1 Custom and Hook Routines

Bmad calculations, like particle tracking through a lattice element, can be customized using what are called "`custom`" and "`hook`" routines. The general idea is that a programmer can implement custom code which is linked into a program and this custom code will be called at the appropriate time by *Bmad*. For example, custom code can be created for Runge-Kutta tracking that calculates the electromagnetic field of some complicated electromagnet. Prototype custom and hook routines are available in the `/bmad/custom` directory and are discussed in detail below.

To enable *Bmad* to be able to call customized code, `function pointers` are defined, one for each custom or hook routine. At certain places in the *Bmad* code, the appropriate function pointer will be checked. If the function pointer is associated with a routine, that routine will be called. By default, the function pointers are not associated with any functions and the only way there will be an association by user code modification.

The function pointers are defined in the file `/bmad/modules/bmad_routine_interface.f90`. The convention followed is that for any given custom or hook routine there is a base name, for example, `track1_custom`, and in `bmad_routine_interface.f90` there will be an abstract interface which with the base name with a `_def` suffix (`track1_custom_def` for this example). Additionally the corresponding function pointer uses a `_ptr` suffix and is defined like:

```
procedure(track1_custom_def), pointer :: track1_custom_ptr => null()
```

To implement custom code for, say, `track1_custom`:

- Copy the file `/bmad/custom/track1_custom.f90` to the area where the program is to be compiled.
- Customize file as desired. The name of the routine can be changed if desired as long as that name is used consistently throughout the program. In fact, multiple custom routines can be created and switched in and out as desired in the program.
- In the program, define an interface for the custom routine like:
  ```
  procedure(track1_custom_def) :: track1_custom
  ```
  This procedure statement must go in the declaration section above the executable section of the code. Note: If the custom routine has been put in a module this will not be needed.
- Somewhere near the beginning of the program (generally before `bmad_parser` is called), set the function pointer to point to your custom routine:

```
track_custom_ptr => track1_custom
```

- If needed modify the compile script(s) (typically named something like "cmake.XXX") to compile the file the custom routine is in.
- Compile the program using the `mk` command.

While coding a custom routine, it is important to remember that it is *not* permissible to modify any routine argument that does not appear in the list of output arguments shown in the comment section at the top of the file.

Note: Custom and hook entry points are added to *Bmad* on an as-needed basis. If you have a need that is not met by the existing set of entry points, please contact a *Bmad* maintainer.

Note: The custom and hook routines in `/bmad/custom/` are not compiled with the *Bmad* library. Their only purpose is to serve as prototypes for code development.

38.2 Custom Calculations

There are essentially two ways to do `custom` (as opposed to `hook`) calculations. One way involves using a `custom` element (§4.11). The other way involves setting the appropriate `method` component of an element to `custom`. An appropriate method component is one of

`tracking_method`	§6.1
`mat6_calc_method`	§6.2
`field_calc`	§6.4
`aperture_type`	§5.8

There are eight routines that implement custom calculations:

```
check_aperture_limit_custom
em_field_custom
init_custom
make_mat6_custom
radiation_integrals_custom
track1_custom
track1_spin_custom
wall_hit_handler_custom
```

[Use `getf` for more details about the argument lists for these routines.]

The init_custom routine is called by bmad_parser at the end of parsing for any lattice element that is a `custom` element or has set any one of the element components as listed above to `custom`. The `init_custom` routine can be used to initialize the internals of the element. For example, consider a `custom` element defined in a lattice file by

```
my_element: custom, val1 = 1.37, descrip = "field.dat", mat6_calc_method = tracking
```

In this example, the `descrip` (§5.3) component is used to specify the name of a file that contains parameters for this element. When `init_custom` is called for this element (see below), the file can be read and the parameters stored in the element structure. Besides the `ele%value` array, parameters may be stored in the general use components given in §32.24.

The make_mat6_custom routine is called by the track1 routine when calculating the transfer matrix through an element.

The track1_custom routine is called by the track1 routine when the `tracking_method` for the element is set to `custom`. Further customization can be set by the routines track1_preprocess and track1_-postprocess. See Section §38.3 for more details.

A potential problem with `track1_custom` is that the calling routine, that is `track1`, does some work like checking aperture, etc. (see the `track1` code for more details). If this is not desired, the `track1_preprocess` routine (§38.3) can be used to do custom tracking and to make sure that `track1` does not do any extra calculations. This is accomplished by putting the custom tracking code in `track1_preprocess` and by setting the `finished` argument of `track1_preprocess` to True.

The check_aperture_limit_custom routine is used to check if a particle has hit an aperture in tracking. It is called by the standard *Bmad* routine check_aperture_limit when `ele%aperture_type` is set to `custom$`. A `custom` element has the standard limit attributes (§5.8) so a `custom` element does not have to implement custom aperture checking code.

The em_field_custom routine is called by the electro-magnetic field calculating routine em_field_calc when `ele%field_calc` is set to `custom$`. As an alternative to `em_field_custom`, a `custom` element can use a field map (§5.16) to characterize the element's electromagnetic fields.

Note: When tracking through a `patch` element, the first step is to transform the particle's coordinates from the entrance frame to the exit frame. This is done since it simplifies the tracking. [The criterion for stopping the propagation of a particle through a `patch` is that the particle has reached the exit face and the calculation to determine if a particle has reached the exit face is simplified if the particle's coordinates are expressed in the coordinate frame of the exit face.] Thus for `patch` elements, unlike all other elements, the particle coordinates passed to em_field_custom are the coordinates with respect to the exit coordinate frame and not the entrance coordinate frame. If field must be calculated in the entrance coordinate frame, a transformation between entrance and exit frames must be done:

```
subroutine em_field_custom (ele, param, s_rel, time, orb, &
                                  local_ref_frame, field, calc_dfield, err_flag)
use lat_geometry_mod
...
real(rp) w_mat(3,3), w_mat_inv(3,3), r_vec(3), r0_vec(3)
real(rp), pointer :: v(:)
...
! Convert particle coordinates from exit to entrance frame.
v => ele%value   ! v helps makes code compact
call floor_angles_to_w_mat (v(x_pitch$), v(y_pitch$), v(tilt$), w_mat, w_mat_inv)
r0_vec = [v(x_offset$), v(y_offset$), v(z_offset$)]
r_vec = [orb%vec(1), orb%vec(3), s_rel]  ! coords in exit frame
r_vec = matmul(w_mat, r_vec) + r0_vec      ! coords in entrance frame

! Calculate field and possibly field derivative
...

! Convert field from entrance to exit frame
field%E = matmul(w_mat_inv, field%E)
field%B = matmul(w_mat_inv, field%B)
if (logic_option(.false., calc_dfield)) then
  field%dE = matmul(w_mat_inv, matmul(field%dE, w_mat))
  field%dB = matmul(w_mat_inv, matmul(field%dB, w_mat))
endif
```

The wall_hit_handler_custom routine is called when the Runge-Kutta tracking code odeint_bmad detects that a particle has hit a wall (§5.12). [This is separate from hitting an aperture that is only defined at the beginning or end of an lattice element.] The dummy `wall_hit_handler_custom` routine does nothing. To keep tracking, the particle must be marked as alive

```
subroutine wall_hit_handler_custom (orb, ele, s, t)
```

```
...
orb%state = alive$   ! To keep on truckin'
...
```

Note: `odeint_bmad` normally does not check for wall collisions. To change the default behavior, the `runge_kutta_com` common block must modified. This structure is defined in `runge_kutta_mod.f90`:

```
type runge_kutta_common_struct
  logical :: check_wall_aperture = .false.
  integer :: hit_when = outside_wall$   ! or wall_transition$
end type

type (runge_kutta_common_struct), save :: runge_kutta_com
```

To check for wall collisions, the `%check_wall_aperture` component must be set to true. The `%hit_when` components determines what constitutes a collision. If this is set to `outside_wall$` (the default), then any particle that is outside the wall is considered to have hit the wall. If `%hit_when` is set to `wall_transition$`, a collision occurs when the particle crosses the wall boundary. The distinction between `outside_wall$` and `wall_transition$` is important if particles are to be allowed to travel outside the wall.

38.3 Hook Routines

A `hook` routine is like a `custom` routine in that a `hook` routine can be used for customizing a *Bmad* calculation by replacing the `dummy` version of a `hook` routine with customized code. The difference is that the `hook` routine is always called at the appropriate time without regard to the type of lattice element under consideration or what tracking method is being used. The `hook` routines that are available are:

```
apply_element_edge_kick_hook
ele_geometry_hook
ele_to_fibre_hook
time_runge_kutta_periodic_kick_hook
track1_bunch_hook
track1_preprocess
track1_postprocess
track1_wake_hook
```

The `apply_element_edge_kick_hook` routine can be used for custom tracking through a fringe field. See the documentation in the file `apply_element_edge_kick_hook.f90` for more details.

The `ele_geometry_hook` routine can be used for custom calculations of the global geometry of an element. This is useful, for example, for a support table on a kinematic mount since *Bmad* does not have the knowledge to calculate the table orientation from the position of the mount points. See the documentation in the file `ele_geometry_hook.f90` for more details.

The ele_to_fibre_hook routine can be used to customize how the PTC fibre corresponding to a *Bmad* lattice element is constructed.

The `time_runge_kutta_periodic_kick_hook` routine can be used to introduce a time dependent kick when doing tracking with `time_runge_kutta`. This routine could be used, for example, to add the kick due to a passing beam ! on a residual gas ion that is being tracked. See the documentation in the file `time_runge_kutta_periodic_kick_hook.f90` for more details.

The `track1_bunch_hook` routine can be used for custom bunch tracking through an element.

The `track1_preprocess` and `track1_postprocess` routines are called by the track1 routine. [Additionally, if the element being tracked through has its tracking method set to `custom`, the `track1_custom` routine is called.] The `track1_preprocess` and `track1_postprocess` routines are useful for a number of things. For example, if the effect of an electron cloud is to be modeled, these two routines can be used to put in half the electron cloud kick at the beginning of an element and half the kick at the end.

The routine `track1_preprocess` has an additional feature in that it has an argument `radiation_included` that can be set to `True` if the routine `track1_custom` will be called and `track1_custom` will be handling radiation damping and excitation effects.

The `track1_wake_hook` can be used to apply custom wakes.

38.4 Nonstandard Init_Custom Use

For back compatibility reasons, a nonstandard method for using `init_custom` is supported. This method involves not bothering to set the `init_custom_ptr` (§38.1) and simply including a custom `init_custom` routine that is included in the list of routines to link to. The problem with this method is that it is not guaranteed to work on all platforms. In particular, this method is known to fail with MacOS and the `gcc` compiler. Due to this, this method is discouraged.

38.5 Physical and Mathematical Constants

Common physical and mathematical constants that can be used in any expression are defined in the file:

```
sim_utils/interfaces/physical_constants.f90
```

The following constants are defined

```
pi = 3.14159265358979d0
twopi = 2 * pi
fourpi = 4 * pi
sqrt_2 = 1.41421356237310d0
sqrt_3 = 1.73205080757d0
complex: i_imaginary = (0.0d0, 1.0d0)

e_mass = 0.51099906d-3   ! DO NOT USE! In GeV
p_mass   = 0.938271998d0   ! DO NOT USE! In GeV

m_electron = 0.51099906d6  ! Mass in eV
m_proton   = 0.938271998d9 ! Mass in eV

c_light = 2.99792458d8             ! speed of light
r_e = 2.8179380d-15                ! classical electron radius
r_p = r_e * m_electron / m_proton  ! proton radius
e_charge = 1.6021892d-19           ! electron charge

h_planck = 4.13566733d-15          ! eV*sec Planck's constant
h_bar_planck = 6.58211899d-16      ! eV*sec h_planck/twopi

mu_0_vac = fourpi * 1e-7                   ! Permeability of free space
```

```
eps_0_vac = 1 / (c_light**2 * mu_0_vac)    ! Permittivity of free space

classical_radius_factor = r_e * m_electron ! Radiation constant

g_factor_electron = 0.001159652193    ! Anomalous gyro-magnetic moment
g_factor_proton   = 1.79285           ! Anomalous gyro-magnetic moment
```

38.6 Global Coordinates and S-positions

The routine lat_geometry will compute the global floor coordinates at the end of every element in a lattice. `lat_geometry` works by repeated calls to ele_geometry which takes the floor coordinates at the end of one element and calculates the coordinates at the end of the next. For conversion between orientation matrix **W** (§16.2) and the orientation angles θ, ϕ, ψ, the routines floor_angles_to_w_mat and floor_w_mat_to_angles can be used.

The routine s_calc calculates the longitudinal s positions for the elements in a lattice.

38.7 Reference Energy and Time

The reference energy and time for the elements in a lattice is calculated by lat_compute_ref_energy_and_time. The reference energy associated with a lattice element is stored in

```
ele%value(E_tot_start$)   ! Total energy at upstream end of element (eV)
ele%value(p0c_start$)     ! Momentum * c_light at upstream end of element (eV)
ele%value(E_tot$)         ! Total energy at downstream end (eV)
ele%value(p0c$)           ! Momentum * c_light at downstream end(eV)
```

Generally, the reference energy is constant throughout an element so that `%value(E_tot_start$` = `%value(E_tot$` and `%value(p0c_start$` = `%value(p0c$`. Exceptions are elements of type:

```
custom,
em_field,
hybrid, or
lcavity
```

In any case, the starting `%value(E_tot_start$` and `%value(p0c_start$` values of a given element will be the same as the ending `%value(E_tot$` and `%value(p0c$` energies of the previous element in the lattice.

The reference time and reference transit time is stored in

```
ele%ref_time                ! Ref time at downstream end
ele%value(delta_ref_time$)
```

The reference orbit for computing the reference energy and time is stored in

```
ele%time_ref_orb_in        ! Reference orbit at upstream end
ele%time_ref_orb_out       ! Reference orbit at downstream end
```

Generally `ele%time_ref_orb_in` is the zero orbit. The exception comes when an element is a `super_slave`. In this case, the reference orbit through the super_slaves of a given `super_lord` is constructed to be continuous. This is done for consistency sake. For example, to ensure that when a marker is superimposed on top of a wiggler the reference orbit, and hence the reference time, is not altered.

`group` (§4.26), `overlay` (§4.40), and `super_lord` elements inherit the reference from the last slave in their slave list (§33.5). For `super_lord` elements this corresponds to inheriting the reference energy of

the slave at the downstream end of the `super_lord`. For `group` and `overlay` elements a reference energy only makes sense if all the elements under control have the same reference energy.

Additionally, photonic elements like `crystal`, `capillary`, `mirror` and `multilayer_mirror` elements have an associated photon reference wavelength

```
ele%value(ref_wavelength$)      ! Meters.
```

38.8 Global Common Structures

There are two common variables used by Bmad for communication between routines. These are `bmad_com`, which is a `bmad_common_struct` structure, and `global_com` which is a `global_common_struct` structure. The `bmad_com` structure is documented in Section §11.4.

The `global_common_struct` is meant to hold common parameters that should not be modified by the user.

```
type global_common_struct
  logical mp_threading_is_safe = T    ! MP threading safe?
  logical exit_on_error  = T          ! Exit program on error?
end type
```

A global variable `global_com` is defined in the `sim_utils` library:

```
type (global_common_struct), save :: global_com
```

And various routines use the settings in `global_com`.

`%mp_threading_is_safe`
: Toggle to prevent MP threading optimizations from being done. See Sec. §38.9 for more details.

`%exit_on_error`
: The `%exit_on_error` component tell a routine if it is OK to stop a program on a severe error. Stopping is generally the right thing when a program is simply doing a calculation and getting a wrong answer is not productive. In control system programs and in interactive programs like `Tao`, it is generally better not to stop on an error.

38.9 Parallel Processing

Bmad was initially developed without regard to parallel processing. When a demand for multithreading capability arose, *Bmad* was modified to meet the need and uses both MP (Multi-Processing) and MPI (Message Passing Interface) type threading. And sometimes both will be used within the same program.

The general rule at present is that *Bmad* can be run multi-threaded as long as either lattice parameters are not varied or an array of lattices is used, one for each thread. Thus multi-threading with MPI is generally thread safe since, by default, different MPI threads do not share memory. With MP, things are more complicated. For example, tracking a particle through a lattice is generally thread safe with MP. The exception is if there are `ramper` (§4.44) elements since ramping involves the modifying lattice element parameters while tracking and is thus not MP thread safe. Another exception is that PTC code (§29) is not thread safe.

In order to signal routines that have MP code whether it is safe to using threading, there is the global switch `global_com%mp_threading_is_safe`. The default is True.

38.10 Expression Trees and Expression Stacks

Expression trees and expression stacks are how *Bmad* and *Tao* handle arithmetical expressions. There is a bit of history here. Originally, there was only expression stacks. An expression stack is an array of `expression_atom_structs`. The main routines to manipulate stacks are:

```
expression_string_to_stack  ! Convert expression string to a reverse Polish stack.
express_stack_value         ! Evaluate expression stack.
expression_value            ! Combines expression_string_to_stack and expression_stack_value.
expression_stack_to_string  ! Creates an expression string from a stack.
```

After the stack code was developed, and when *Tao* was created, since the needs of *Tao* were more extensive (needed to handle arrays, needed to handle variable, data and lattice element names that used the asterisk as a wild card while, at the same time the asterisk was used as the multiplication symbol), a different evaluation method using expression trees was implemented. Note: While trees are more flexible than stacks, they also are more complicated and for this reason, and the fact that stacks work perfectly well for *Bmad*, the stacks that are built into *Bmad* have been left untouched.

The main routines to manipulate trees are:

```
expression_string_to_tree  ! Convert expression string to tree form.
type_expression_tree       ! Print an expression tree in tree form.
expression_tree_to_string  ! Convert a tree to a string.
```

Note These routines are in *Bmad*. *Tao* implements other routines. for evaluation.

Chapter 39

PTC/FPP Programming

The PTC/FPP library of Étienne Forest handles Taylor maps to any arbitrary order. this is also known as Truncated Power Series Algebra (TPSA). The core Differential Algebra (DA) package used by PTC/FPP was developed by Martin Berz[Berz89]. The PTC/FPP code is interfaced to *Bmad* so that calculations that involve both *Bmad* and PTC/FPP can be done in a fairly seamless manner.

FPP
: The "Fully Polymorphic Package" (`FPP`) library implements Differential Algebra (DA) for the manipulation of Taylor maps. Thus in FPP you can define a Hamiltonian and then generate the Taylor map for this Hamiltonian. FPP is very general. It can work with an arbitrary number of dimensions. `FPP` is purely mathematical in nature. It has no knowledge of accelerators, magnetic fields, particle tracking, Twiss parameters, etc.

PTC
: The "Polymorphic Tracking Code" `PTC` library is for accelerator simulation. It uses `FPP` as a back end for calculating such things as one turn maps.

PTC is used by *Bmad* when constructing Taylor maps and when the `tracking_method` §6.1) is set to `symp_lie_ptc`. All Taylor maps above first order are calculated via PTC. No exceptions.

For information on using PTC within *Bmad*, see Chapter §29. For more information on PTC/FPP in general the PTC/FPP manual[Forest02].

39.1 Phase Space

PTC uses different longitudinal phase space coordinates compared to *Bmad*. *Bmad*'s phase space coordinates are (§17.2)

$$(x, p_x, y, p_y, z, p_z) \tag{39.1}$$

In PTC one can choose between several different coordinate systems. The one that Bmad uses is

$$(x, p_x, y, p_y, p_t, c\Delta t) \tag{39.2}$$

where

$$p_t = \frac{\Delta E}{c\, P_0} \tag{39.3}$$

This choice of phase space is set in set_ptc. Specifically, the PTC global variable DEFAULT, which is of type internal_states, has the %time switch set to True.

vec_bmad_to_ptc and vec_ptc_to_bmad are conversion routines that translate between the two. Actually there are a number of conversion routines that translate between *Bmad* and PTC structures. See §44.33 for more details.

39.2 PTC Initialization

One important parameter in PTC is the order of the Taylor maps. By default *Bmad* will set this to 3. The order can be set within a lattice file using the parameter[taylor_order] attribute. In a program the order can be set using set_ptc. In fact set_ptc must be called by a program before PTC can be used. bmad_parser will do this when reading in a lattice file. That is, if a program does not use bmad_parser then to use PTC it must call set_ptc. Note that resetting PTC to a different order reinitializes PTC's internal memory so one must be careful if one wants to change the order in mid program.

39.3 PTC Structures Compared to Bmad's

Bmad uses a lat_struct structure to hold the information on a machine and a lat_struct has an array of branch_structs (the %branch(:) component) with each branch_struct holding an array of ele_structs (the %ele(:) component). The ele_struct holds the information on the individual elements. An ele_struct holds information about both the physical element and the reference orbit through it.

PTC has a somewhat different philosophy as illustrated in Fig. 39.1. A PTC mad_universe structure is very roughly equivalent to a *Bmad* lat_struct. That is, both structures can contain the description for an entire accelerator complex. Note that it is standard in PTC to use two mad_universe structures called m_u and m_t. These two are defined globally. The difference between m_u and m_t is that m_u is used as a bookkeeping device for convenient accessing of all lattice elements. On the other hand, m_t contains the layouts that can be used for tracking.

equivalent to a *Bmad* branch_struct. A layout has a pointer to a linked list of fibre structures. Each fibre has a pointer to a magnet structure which holds the information about the physical element and each fibre holds information about the reference orbit through the element.

With PTC, The top level structure mad_universe has two components called %first and %last which are pointers to the ends of an array of layout_array structures. Each layout_array holds a layout structure. A layout structure has pointers to the previous and next layouts making a linked list of layouts indicated by the horizontal arrows. Each layout has pointers to a linked list of fibre structures. The fibre structures represent the reference trajectory through an element. Each fibre structure has a pointer to a element and an elementp structures which represent the physical element. With *Bmad*, the lat_struct roughly corresponds to the PTC layout_array(:), the branch_struct roughly corresponds to the PTC layout and the element_struct roughly corresponds to the PTC fibre, element and elementp structures.

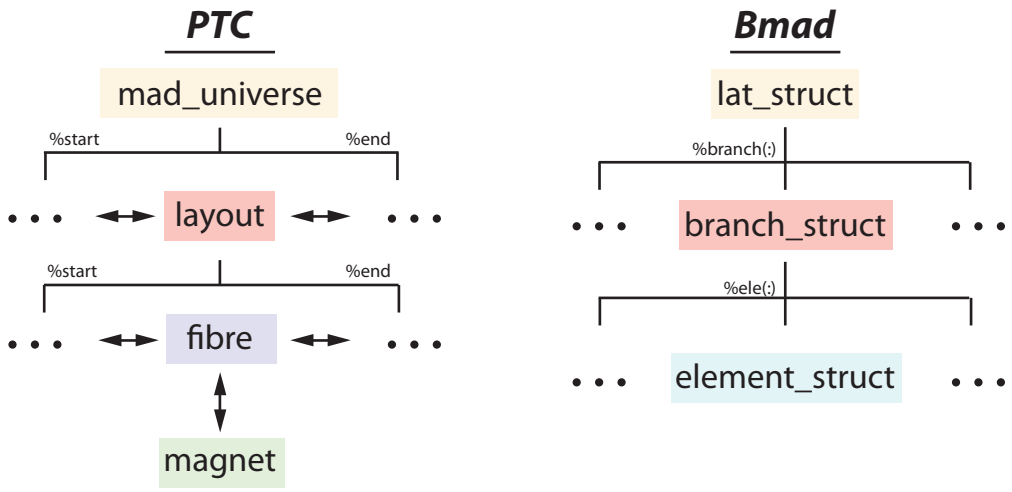

Figure 39.1: Simplified diagram showing the organization of the major PTC structures involved in defining a lattice contrasted with *Bmad*.

39.4 Variable Initialization and Finalization

PTC variables must be initialized and finalized. This is done with the`alloc()` and `kill()` routines. In addition, the `real_8_init` routine can initialize a `real_8` array:

```
type (real_8) y8(6)
...
call real_8_init (y8)
call kill (y8)
```

39.5 Correspondence Between Bmad Elements and PTC Fibres

.

When a PTC `layout` is created from a *Bmad* `lat_struct` instance using the routine lat_to_ptc_layout, the correspondence between the *Bmad* elements and the PTC fibres is maintained through the `ele%ptc_fibre` pointer. The following rules apply:

1. There will be marker `fibre`s at the beginning and end of the `layout`. The beginning `fibre` will correspond to `branch%ele(0)`. The end `fibre` will not have a corresponding *Bmad* element.

2. Generally there will be a one-to-one correspondence between `fibre`s and `branch%ele` elements. The exception is where a "hard edge" model is used for tracking. In this case, there will be three `fibre`s for the *Bmad* element: Two drift `fibre`s with a `fibre` of the appropriate type in between. In this case, `ele%ptc_fibre` will point to the last (drift) `fibre`.

Remember: The attributes like reference energy, etc. for a *Bmad* `ele_struct` instance are referenced to the exit end of the element. For PTC the reference edge for a `fibre` is the entrance end.

39.6 Taylor Maps

FPP stores its `real_8` Taylor maps in such a way that it is not easy to access them directly to look at the particular terms. To simplify life, Étienne has implemented the `universal_taylor`structure:

```
type universal_taylor
  integer, pointer  :: n       ! Number of coefficients
  integer, pointer  :: nv      ! Number of variables
  real(dp), pointer :: c(:)    ! Coefficients C(N)
  integer, pointer  :: j(:,:)  ! Exponents of each coefficients J(N,NV)
end type
```

Bmad always sets `nv` = 6. *Bmad* overloads the equal sign to call routines to convert between Étienne's `real_8` Taylor maps and `universal_taylor`:

```
type (real_8) tlr(6)           ! Taylor map
type (universal_taylor) ut(6)  ! Taylor map
...
tlr = ut                       ! Convert universal_taylor -> real_8
ut = tlr                       ! Convert real_8 -> universal_taylor
```

39.7 Patches

There is a significant difference between how patches are treated in PTC and *Bmad*. In PTC, a patch is just though of as a coordinate transformation for propagating a particle from one `fibre` to the next. As such, the `patch` is part of a `fibre`. That is, any `fibre` representing tracking through quadrupoles, bends, etc. will have patches for the entrance and exit ends of the `fibre`.

With *Bmad*, on the other hand, a `patch` is a "first class" element on par with all other elements be they quadrupoles, bends, etc. When translating a `patch` from *Bmad* to PTC, the `patch` is represented in PTC as a `marker` element with a patch at the exit end.

39.8 Number of Integration Steps & Integration Order

"Drift like" elements in PTC will use, by default, only one integration step. *Bmad* uses the default when translating from *Bmad* lattice elements to PTC fibres. The *Bmad* lattice elements that are drift like are:

```
drift
ecollimator
instrument
monitor
pipe
rcollimator
```

When tracking, there is a trade-off between step size and integrator order. Higher order means fewer steps are needed to get the same accuracy. But one higher order step is computationally more intensive then one lower order step so what is the optimum order and number of steps is dependent upon various factors like magnet strength and how fast the field is varying. Generally, when the field is varying, such as in a wiggler, lower order and more steps are favored. Also spin tracking is always 2nd order in PTC. So going to higher order for the orbital tracking with less steps will cause the spin tracking to be less accurate.

The way PTC "resplitting" routines work is that, for a given element, they start by assuming that the tracking will be done using a 2^{nd} order integrator, They then compute the number of steps needed based upon the electric and magnetic field strengths. This number is compared to a crossover limit point here named C_1. If the number of steps is less than or equal to C_1 then the resplitting routine stops and tracking will thereafter be done with a 2^{nd} order integrator with the calculated number of steps. On the other hand, if the number of steps is greater than C_1, the resplitting routine will redo the calculation assuming 4^{th} order integration. With 4^{th} order integration, the number of calculated steps will compared to a different crossover limit point here called C_2. Again, if the number of steps is less than or equal to C_2, the routine will assign 4^{th} order tracking to the element. Otherwise, the routine will assign 6^{th} order tracking to the element with an appropriate number of steps.

The default crossover limit points are

$$[C_1, C_2] = [30, 60] \qquad \text{For wiggler type elements.}$$
$$[C_1, C_2] = [4, 18] \qquad \text{For all other elements.}$$

The greater number for wigglers is a reflection of the fact that the wiggler field is not constant.

39.9 Creating a PTC layout from a Bmad lattice

For a programmer, it is sometimes useful to feed a *Bmad* lattice into PTC and then use PTC for all the calculations. As an example of how to do this, the following minimal program creates a PTC `layout` from a *Bmad* lattice:

```
use pointer_lattice, dummy => lat
use ptc_layout_mod, dum1 => dp
implicit none
type (lat_struct), target :: lat
type(layout), pointer:: als
!
call bmad_parser ('lat.bmad', lat)
call lat_to_ptc_layout (lat, .true.)
als => lat%branch(0)%ptc%m_t_layout
```

39.10 Internal_State

The `internal_state` structure looks like:

```
type internal_state
   integer totalpath      ! total time or path length is used
   logical(lp) time       ! Time is used instead of path length
   logical(lp) radiation  ! Radiation damping (but not excitation) is turned on
   logical(lp) nocavity   ! Cavity is turned into a drift
   logical(lp) fringe     ! fringe fields are turned on (mainly for quadrupoles)
   logical(lp) stochastic ! Random Stochastic kicks to x(5)
   logical(lp) envelope   ! Stochastic envelope terms tracked in probe_8
   logical(lp) para_in    ! If true, parameters in the map are included
   logical(lp) only_4d    ! REAL_8 Taylor in (x,p_x,y,p_y)
   logical(lp) delta      ! REAL_8 Taylor in (x,p_x,y,p_y,delta)
   logical(lp) spin       ! Spin is tracked
   logical(lp) modulation ! One modulated family tracked by probe
```

```
   logical(lp) only_2d    ! REAL_8 Taylor in (x,p_x)
   logical(lp) full_way   !
end type internal_state
```

Chapter 40

OPAL

OPAL (Object Oriented Parallel Accelerator Library) is a tool for charged-particle optic calculations in large accelerator structures and beam lines including 3D space charge. OPAL is built from first principles as a parallel application, OPAL admits simulations of any scale: on the laptop and up to the largest High Performance Computing (HPC) clusters available today. Simulations, in particular HPC simulations, form the third pillar of science, complementing theory and experiment.

OPAL includes various beam line element descriptions and methods for single particle optics, namely maps up to arbitrary order, symplectic integration schemes and lastly time integration. OPAL is based on IPPL (Independent Parallel Particle Layer) which adds parallel capabilities. Main functions inherited from IPPL are: structured rectangular grids, fields and parallel FFT and particles with the respective interpolation operators. Other features are, expression templates and massive parallelism (up to 8000 processors) which makes is possible to tackle the largest problems in the field.

The manual can be obtained at

`amas.web.psi.ch/docs/opal/`

40.1 Phase Space

OPAL uses different longitudinal phase space coordinates compared to *Bmad*. *Bmad*'s phase space coordinates are

$$(x, p_x/p_0, y, p_y/p0, -\beta c(t - t_0), (p - p_0)/p_0) \tag{40.1}$$

OPAL uses

$$(x, \gamma\beta_x, y, \gamma\beta_y, z, \gamma\beta_z) \tag{40.2}$$

`convert_particle_coordinates_s_to_t` and `convert_particle_coordinates_s_to_t` are conversion routines ...

Chapter 41

C++ Interface

To ease the task of using *C++* routines with *Bmad*, there is a library called `cpp_bmad_interface` which implements a set of *C++* classes in one–to–one correspondence with the major *Bmad* structures. In addition to the *C++* classes, the *Bmad* library defines a set of conversion routines to transfer data values between the *Bmad* Fortran structures and the corresponding *C++* classes.

The list of all classes is given in the file

```
cpp_bmad_interface/include/cpp_bmad_classes.h
```

The general rule is that the equivalent class to a *Bmad* structure named `xxx_struct` will be named `CPP_xxx`. Additionally, for each *Bmad* structure, there is a opaque class named `Bmad_xxx_class` for use in the translation code discussed below. The names of these opaque classes have the form `Bmad_xxx_class` and are used to define pointer instances in routine argument lists.

41.1 C++ Classes and Enums

Generally, The *C++* classes have been set up to simply mirror the corresponding *Bmad* structures. For example, the `CPP_lat` class has a string component named `.version` that mirrors the `%version` component of the `lat_struct` structure. There are some exceptions. For example, structure components that are part of `PTC` (§1.4) are not present in the classes.

While generally the same component name is used for both the *Bmad* structures and the *C++* classes, in the case where there is a *C++* reserved word conflict, the *C++* component name will be different.

A header file `bmad_enums.h` defines corresponding *Bmad* parameters for all *C++* routine. The *Bmad* parameters are in a namespace called `Bmad`. The convention is that the name of a corresponding *C ++* parameter is obtained by dropping the ending `$` (if there is one) and converting to uppercase. For example, `electron$` on the Fortran side converts to `Bmad::ELECTRON` in *C++*.

All of the *C++* class components that are arrays or matrices are zero based so that, for example, the index of the `.vec[i]` array in a `CPP_coord` runs from 0 through 5 and not 1 through 6 as on the Fortran side. Notice that for a `lat_struct` the `%ele(0:)` component has a starting index of zero so there is no off–by–one problem here. The exception to this rule is the `%value(:)` array of the `ele_struct` which has a span from 1 to `num_ele_attrib$`. In this case, To keep the conversion of the of constructs like `ele%value(k1$)` consistant, the corresponding `ele.value[]` array has goes from 0 to `Bmad::NUM_ELE_ATTRIB` with the 0th element being unused.

```
1     subroutine f_test
2       use bmad_cpp_convert_mod
3       implicit none
4
5       interface
6         subroutine cpp_routine (f_lat, c_coord) bind(c)
7           import f_lat, c_ptr
8           type (lat_struct) :: f_lat
9           type (c_ptr), value :: c_coord
10        end subroutine
11      end interface
12
13      type (lat_struct), target :: lattice   // lattice on Fortran side
14      type (coord_struct), target :: orbit
15      type (c_ptr), value :: c_lat
16      ! ...
17      call lat_to_c (c_loc(lattice), c_lat)    ! Fortran side convert
18      call cpp_routine (c_lat, c_loc(orbit))   ! Call C++ routine
19      call lat_to_f (c_lat, c_loc(lattice))    ! And convert back
20    end subroutine
```

Figure 41.1: Example Fortran routine calling a *C++* routine.

```
1     #include "cpp_bmad_classes.h"
2
3     using namespace Bmad;
4
5     extern "C" cpp_routine (CPP_lat& c_lat, Bmad_coord_class* f_coord,  f_lat) {
6       CPP_coord c_coord;
7       coord_to_c (f_coord, c_coord);        // C++ side convert
8       // ... do calculations ...
9       cout << c_lat.name << "  " << c_lat.ele[1].value[K1] << endl;
10      coord_to_f (c_coord, f_coord);        // And convert back
11    }
```

Figure 41.2: Example *C++* routine callable from a Fortran routine.

41.2 Conversion Between Fortran and C++

A simple example of a Fortran routine calling a *C++* routine is shown in Figs. 41.1 and 41.2. Conversion between structure and classes can happen on either the Fortran side or the *C++* side. In this example, the `lat_struct` / `CPP_lat` conversion is on the Fortran side and the `coord_struct` / `CPP_coord` is on the *C++* side.

On the Fortran side, the interface block defines the argument list of the *C++* routine being called.

On the *C++* side, `f_coord` is an instance of the `Bmad_coord_class` opaque class.

A *C++* routine calling a Fortran routine has a similar structure to the above example. The interface block in Fig. 41.1 can be used as a prototype. For additional examples of conversion between Fortran and *C++*, look at the test code in the directory

```
cpp_bmad_interface/interface_test
```

Chapter 42

Quick_Plot Plotting

Quick Plot is an interface layer to either the `PGPLOT`[PGPLOT] or `PLPLOT`[PLPLOT] plotting libraries. Whether `PGPLOT` or `PLPLOT` is used depends upon an environmental switch set when the *Bmad* library and other associated libraries are compiled (§30.2). [Note: *Quick Plot* lives in the `sim_utils` library which comes with the Bmad distribution.] A quick reference guide can be seen online by using the command `getf quick_plot`. For identification in a program, all *Quick Plot* subroutines start with a `qp_` prefix. Also, by convention, all PGPLOT subroutines start with a `pg` prefix.

```
1   program example_plot
2     use quick_plot
3     integer id
4     character(1) ans
5
6     ! Generate PS and X-windows plots.
7     call qp_open_page ("PS-L")  ! Tell \quickplot to generate a PS file.
8     call plot_it              ! Generate the plot
9     call qp_close_page        ! quick_plot.ps is the file name
10    call qp_open_page ("X", id, 600.0_rp, 470.0_rp, "POINTS")
11    call plot_it
12    write (*, "(a)", advance = "NO") " Hit any class to end program: "
13    accept "(a)", ans
14
15  !-----------------------------------------------------------------------
16  contains
17  subroutine plot_it                            ! This generates the plot
18    real(rp), allocatable :: x(:), y(:), z(:), t(:)
19    real(rp) x_axis_min, x_axis_max, y_axis_min, y_axis_max
20    integer x_places, x_divisions, y_places, y_divisions
21    character(80) title
22    logical err_flag
23    namelist / parameters / title
24
25    ! Read in the data
26    open (1, file = "plot.dat", status = "old")
27    read (1, nml = parameters)                  ! read in the parameters.
28    call qp_read_data (1, err_flag, x, 1, y, 3, z, 4, t, 5) ! read in the data.
29    close (1)
30
31    ! Setup the margins and page border and draw the title
32    call qp_set_page_border (0.01_rp, 0.02_rp, 0.2_rp, 0.2_rp, "%PAGE")
33    call qp_set_margin (0.07_rp, 0.05_rp, 0.05_rp, 0.05_rp, "%PAGE")
34    call qp_draw_text (title, 0.5_rp, 0.85_rp, "%PAGE", "CT")
35
36    ! draw the left graph
37    call qp_set_box (1, 1, 2, 1)
38    call qp_calc_and_set_axis ("X", minval(x), maxval(x), 4, 8, "ZERO_AT_END")
39    call qp_calc_and_set_axis ("Y", minval(z), maxval(z), 4, 8, "GENERAL")
40    call qp_draw_axes ("X\dlab\u", "\gb(\A)")
41    call qp_draw_data (x, y, symbol_every = 0)
42
43    call qp_save_state (.true.)
44    call qp_set_symbol_attrib ('times', color = "blue", height = 20.0_rp)
45    call qp_set_line_attrib ("PLOT", color = "blue", style = "dashed")
46    call qp_draw_data (x, z, symbol_every = 5)
47    call qp_restore_state
48
49    ! draw the right graph. "star5_filled" is a five pointed star.
50    call qp_save_state (.true.)
51    call qp_set_box (2, 1, 2, 1)
52    call qp_set_graph_attrib (draw_grid = .false.)
53    call qp_set_symbol_attrib ('star5_filled', height = 10.0_rp)
54    call qp_set_axis ("Y", -0.1_rp, 0.1_rp, 4, 2)
55    call qp_set_axis ('Y2', 1.0_rp, 100.0_rp, label = "Y2 axis", &
56                               draw_numbers = .true., ax_type = "LOG")
57    call qp_draw_axes ("\m1 \m2 \m3 \m4 \m5 \m6 \m7", "\fsLY\fn", title = "That Darn Graph")
58    call qp_draw_data (x, t, draw_line = .false., symbol_every = 4)
59    call qp_restore_state
60  end subroutine
61  end program
```

Figure 42.1: *Quick Plot* example program.

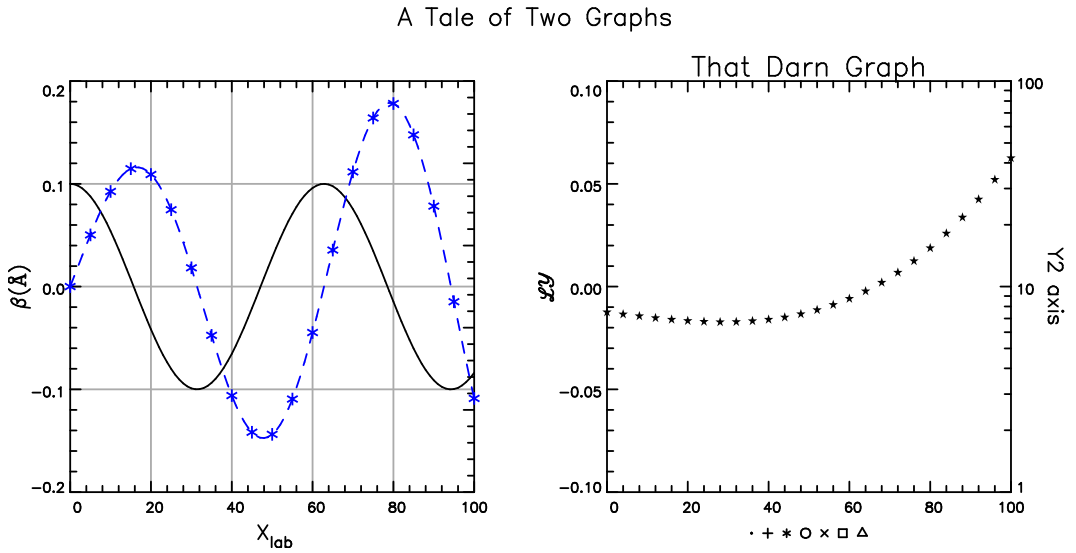

Figure 42.2: Output of plot_example.f90.

42.1 An Example

An example of how *Quick Plot* can be used in a program is shown in Fig. 42.1. In the *Bmad* distribution a copy of this program is in the file

```
sim_utils/plot_example/plot_example.f90
```

The `plot_example.f90` program generates the figure shown in Fig. 42.2 from the input file named `plot.dat`. The first few lines of the data file are

```
&parameters
  title = "A Tale of Two Graphs"
/

Any junk here...

Col1      Col2      Col3      Col4      Col5
   0    0.0000    0.1000    0.0000   -0.0125
   1    0.0001    0.0995    0.0101   -0.0127
   2    0.0004    0.0980    0.0203   -0.0130
   3    0.0009    0.0955    0.0304   -0.0132
   ...
```

The program first creates a PostScript file for printing on lines 7 through 9 and then makes an X–windows plot on lines 10 and 11. The write/accept lines 12 and 13 are to pause the program to prevent the X-window from immediately closing upon termination of the program.

The heart of the plotting is in the subroutine `plot_it` beginning on line 17. The namelist read on line 27 shows how both parameters and data can be stored in the same file so that a plotting program can be automatically told what the appropriate plot labels are. The qp_draw_text call on line 34 draws the title above the two graphs.

The qp_read_data call on line 28 will skip any "header" lines (lines that do not begin with something that looks like a number) in the data file. In this instance qp_read_data will read the first, third forth and fifth data columns and put them into the `x`, `y`, `z`, and `t` arrays.

qp_set_page_border, qp_set_box, and qp_set_margin sets where the graph is going to be placed. qp_set_box(1, 1, 2, 1) on line 37 tells *Quick Plot* to put the first graph in the left box of a 2 box grid. The qp_set_margin on line 33 sets the margins between the box and the graph axes.

qp_calc_and_set_axis on lines 38 and 39 are used to scale the axes. "ZERO_AT_END" ensures that the x–axis starts (or stops) at zero. qp_calc_and_set_axis is told to restrict the number of major divisions to be between 4 and 8. For the horizontal axis, as can be seen in Fig. 42.2, it chooses 5 divisions.

After drawing the first data curve (the solid curve) in the left graph, the routines qp_set_symbol_attrib and qp_set_line_attrib are called on lines 44 and 45 to plot the next data curve in blue with a dashed line style. By default, this curve goes where the last one did: in the left graph. To keep the setting of the line and symbol attributes from affecting other plots the routines qp_save_state and qp_restore_state on lines 43 and 47 are used. qp_save_state saves the current attributes in an attribute stack. qp_restore_state restores the saved attributes from the attribute stack. qp_draw_axes is called on line 40 to draw the x and y-axes along, and qp_draw_data is called on lines 41 and 46 to draw the two data curves.

Lines 50 through 60 draw the third curve in the right hand graph. The qp_set_axis call on lines 55/56 sets a log scale for the y2 (right hand) axis. The syntax of the string arguments of qp_draw_axes in lines 40 and 57/58 comes from PGPLOT and allows special symbols along with subscripts and superscripts.

42.2 Plotting Coordinates

Quick Plot uses the following concepts as shown in Fig. 42.3

```
PAGE  -- The entire drawing surface.
BOX   -- The area of the page that a graph is placed into.
GRAPH -- The actual plotting area within the bounds of the axes.
```

In case you need to refer to the PGPLOT routines the correspondence between this and PGPLOT is:

```
QUICK_PLOT    PGPLOT
----------    ------
PAGE          VIEW SURFACE
BOX           No corresponding entity.
```

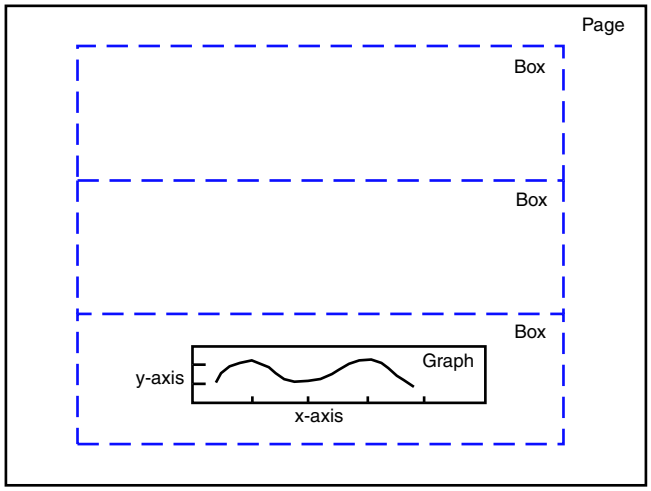

Figure 42.3: A Graph within a Box within a Page.

```
GRAPH         VIEWPORT and WINDOW
```

Essentially the VIEWPORT is the region outside of which lines and symbols will be clipped (if clipping is turned on) and the WINDOW defines the plot area. I'm not sure why PGPLOT makes a distinction, but VIEWPORT and WINDOW are always the same region.

qp_open_page determines the size of the `page` if it is settable (like for X–windows). The page is divided up into a grid of boxes. For example, in Fig. 42.3, the grid is 1 box wide by 3 boxes tall. The border between the grid of boxes and the edges of the page are set by qp_set_page_border. The box that the graph falls into is set by qp_set_box. The default is to have no margins with 1 box covering the entire page. The qp_set_margin routine sets the distance between the box edges and the axes (See the PGPLOT manual for more details).

42.3 Length and Position Units

Typically there is an optional `units` argument for *Quick Plot* routines that have length and/or position arguments. For example, using `getf` one can see that the arguments for qp_draw_rectangle are

```
Subroutine qp_draw_rectangle (x1, x2, y1, y2, units, color, width, style, clip)
```

The `units` argument is a character string which is divided into three parts. The syntax of the `units` argument is

```
unit_type/ref_object/corner
```

The first part `unit_type` gives the type of units

```
"%"       -- Percent.
"DATA"    -- Data units. (Draw default)
"MM"      -- millimeters.
"INCH"    -- Inches. (Set default)
"POINTS"  -- Printers points. NOT PIXELS. (72 points = 1 inch).
```

Note: For displays with a resolution of 72 pixels / inch, `POINTS` corresponds to pixels but many displays have a higher resolution. The second and third parts give the reference point for a position. The second part specifies the reference object

```
"PAGE"  -- Relative to the page (Set default).
"BOX"   -- Relative to the box.
"GRAPH" -- Relative to the graph (Draw default).
```

The third part gives corner of the reference object that is the reference point

```
"LB"    -- Left Bottom (Set and Draw default).
"LT"    -- Left Top.
"RB"    -- Right Bottom.
"RT"    -- Right Top.
```

Notes:

- The `DATA` unit type, by definition, always uses the lower left corner of the `GRAPH` as a reference point.
- For the `%` `unit_type` the `/` between `unit_type` and `ref_object` can be omitted.
- If the `corner` is specified then the `ref_object` must appear also.
- Everything must be in upper case.
- For some routines (qp_set_margin, etc.) only a relative distance is needed. In this case the `ref_object/corner` part, if present, is ignored.

- The `units` argument is typically an optional argument. If not present the default units will be used. There are actually two defaults: The draw default is used for drawing text, symbols, or whatever. The set default is used for setting margins, and other lengths. Initially the draw default is `DATA/GRAPH/LB` and the set default is `INCH/PAGE/LB`. Use qp_set_parameters to change this.

Examples:

```
"DATA"          -- This is the draw default.
"DATA/GRAPH/LB" -- Same as above.
"DATA/BOX/RT"   -- ILLEGAL: DATA must always go with GRAPH/LB.
"%PAGE/LT"      -- Percentage of page so (0.0, 1.0) = RT of page.
"%BOX"          -- Percentage of box so (1.0, 1.0) = RT of box.
"INCH/PAGE"     -- Inches from LB of page.
```

42.4 Y2 and X2 axes

The top and right axes of a graph are known as `X2` and `Y2` respectively as shown in Fig. 42.3. Normally the `X2` axis mirrors the `X` axis and the `Y2` axis mirrors the `Y` axis in that the tick marks and axis numbering for the `X2` and `Y2` axes are the same as the `X` and `Y` axes respectively. qp_set_axis can be used to disable mirroring. For example:

```
call qp_set_axis ("Y2", mirror = .false.)  ! y2-axis now independent of y.
```

qp_set_axis can also be used to set `Y2` axis parameters (axis minimum, maximum, etc.) and setting the `Y2` or `X2` axis minimum or maximum will, by default, turn off mirroring.

Note that the default is for the `X2` and `Y2` axis numbering not to be shown. To enable or disable axis numbering again use qp_set_axis. For example:

```
call qp_set_axis ("Y2", draw_numbers = .true.)  ! draw y2 axis numbers
```

To plot data using the `X2` or `Y2` scale use the qp_use_axis routine. For example:

```
call qp_save_state (.true.)
call qp_use_axis (y = "Y2")
! ... Do some data plotting here ...
call qp_restore_state
```

42.5 Text

PGPLOT defines certain escape sequences that can be used in text strings to draw Greek letters, etc. These escape sequences are given in Table 42.2.

PGPLOT defines a text background index:

```
        -1 - Transparent background.
         0 - Erase underlying graphics before drawing text.
  1 to 255 - Opaque with the number specifying the color index.
```

42.6 Styles

Symbolic constants have been defined for *Quick Plot* subroutine arguments that are used to choose various styles. As an example of this is in lines 44 and 45 of Fig. 42.1. The numbers in the following are the PGPLOT equivalents.

The *Quick Plot* line styles are:

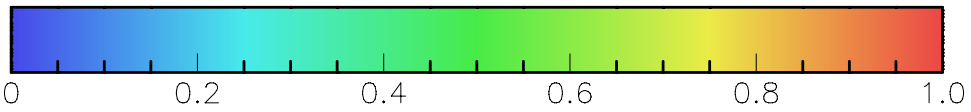

Figure 42.4: Continuous colormap used by `qp_continuous_color`. Typical usage: `call qp_routine(..., color = qp_continuous_color(0.25_rp), ...)`

```
1 -- solid$                  Solid
2 -- dashed$                 Dashed
3 -- dash_dot$               Dash--dot
4 -- dotted$                 Dotted
5 -- dash_dot3$              Dash--dot--dot--dot
```

The color styles in *Quick Plot* are:

```
 0 -- White$   (actually the background color)
 1 -- Black$   (actually the foreground color)
 2 -- Red$
 3 -- Green$
 4 -- Blue$
 5 -- Cyan$
 6 -- Magenta$
 7 -- Yellow$
 8 -- Orange$
 9 -- Yellow_Green$
10 -- Light_Green$
11 -- Navy_Blue$
12 -- Purple$
13 -- Reddish_Purple$
14 -- Dark_Grey$
15 -- Light_Grey$
16 -- Transparent$ (not drawn)
```

In addition, all 138 CSS4 (CSS Color Level 4) named colors are available. These occupy color indices 17 through 154 and can be specified by name (e.g., `"salmon"`, `"darkslategray"`, `"mediumspringgreen"`). The CSS4 specification defines 148 named colors; 10 of these share a name with an original *Quick Plot* color and are not duplicated in the CSS4 range since they are already accessible at their original indices: `black`, `blue`, `cyan`, `green`, `magenta`, `orange`, `purple`, `red`, `white`, and `yellow`. Note that CSS4 aliases with different names but the same RGB value (e.g., `aqua` = `cyan`, `fuchsia` = `magenta`) are included. The CSS4 colors are ordered by visual distinctiveness so that devices with limited color index support still get the most distinguishable colors at lower indices. To see what the CSS4 colors look like, the Wikipedia article on "Web colors" has a handy table of "Extended colors" sorted by color types.

Note: The RGB values of `Green`, `Orange`, and `Purple` now match the CSS4 specification. Users of the previous values can use these CSS4 equivalents:

```
Old Green  (0, 255, 0)   -> "lime"
Old Orange (255, 127, 0) -> (no exact match; nearest is "darkorange")
Old Purple (127, 0, 255) -> (no exact match; nearest is "blueviolet")
```

Integers from 155 onward represent continuous colors. The function `qp_continuous_color` maps a real value in [0.0, 1.0] to these integers. The colormap uses HLS interpolation with hue going from 240°(blue) to 0°(red), saturation = 0.8, and lightness = 0.6. See Fig. 42.4. The inverse function `qp_continuous_color_inverse` maps from integer back to a real in [0, 1].

The fill styles are:

```
1 -- solid_fill$
2 -- no_fill$
3 -- hatched$
4 -- cross_hatched$
```

The symbol types are:

```
 0 -- square_sym$
 1 -- dot_sym$
 2 -- plus_sym$
 3 -- times_sym$
 4 -- circle_sym$
 5 -- x_sym$
 7 -- triangle_sym$
 8 -- circle_plus_sym$
 9 -- circle_dot_sym$
10 -- square_concave_sym$
11 -- diamond_sym$
12 -- star5_sym$
13 -- triangle_filled_sym$
14 -- red_cross_sym$
15 -- star_of_david_sym$
16 -- square_filled_sym$
17 -- circle_filled_sym$
18 -- star5_filled_sym$
```

Beside this list, PGPLOT maps other numbers onto symbol types. The PGPLOT list of symbols is:

```
-3 ... -31 - a regular polygon with abs(type) edges.
         -2 - Same as -1.
         -1 - Dot with diameter = current line width.
 0 ...   31 - Standard marker symbols.
32 ... 127 - ASCII characters (in the current font).
                 E.G. to use letter F as a marker, set type = ICHAR("F").
     > 127 - A Hershey symbol number.
```

Table 42.1 shows some of the symbols and there associated numbers. Note: At constant height PGPLOT gives symbols of different size. To partially overcome this, *Quick Plot* scales some of the symbols to give a more uniform appearance. Table 42.1 was generated using a height of 40 via the call

```
call qp_draw_symbol (0.5_rp, 0.5_rp, "%BOX", k, height = 40.0_rp)
```

Table 42.3 shows how the character string "\g<r>", where "<r>" is a Roman letter, map onto the Greek character set.

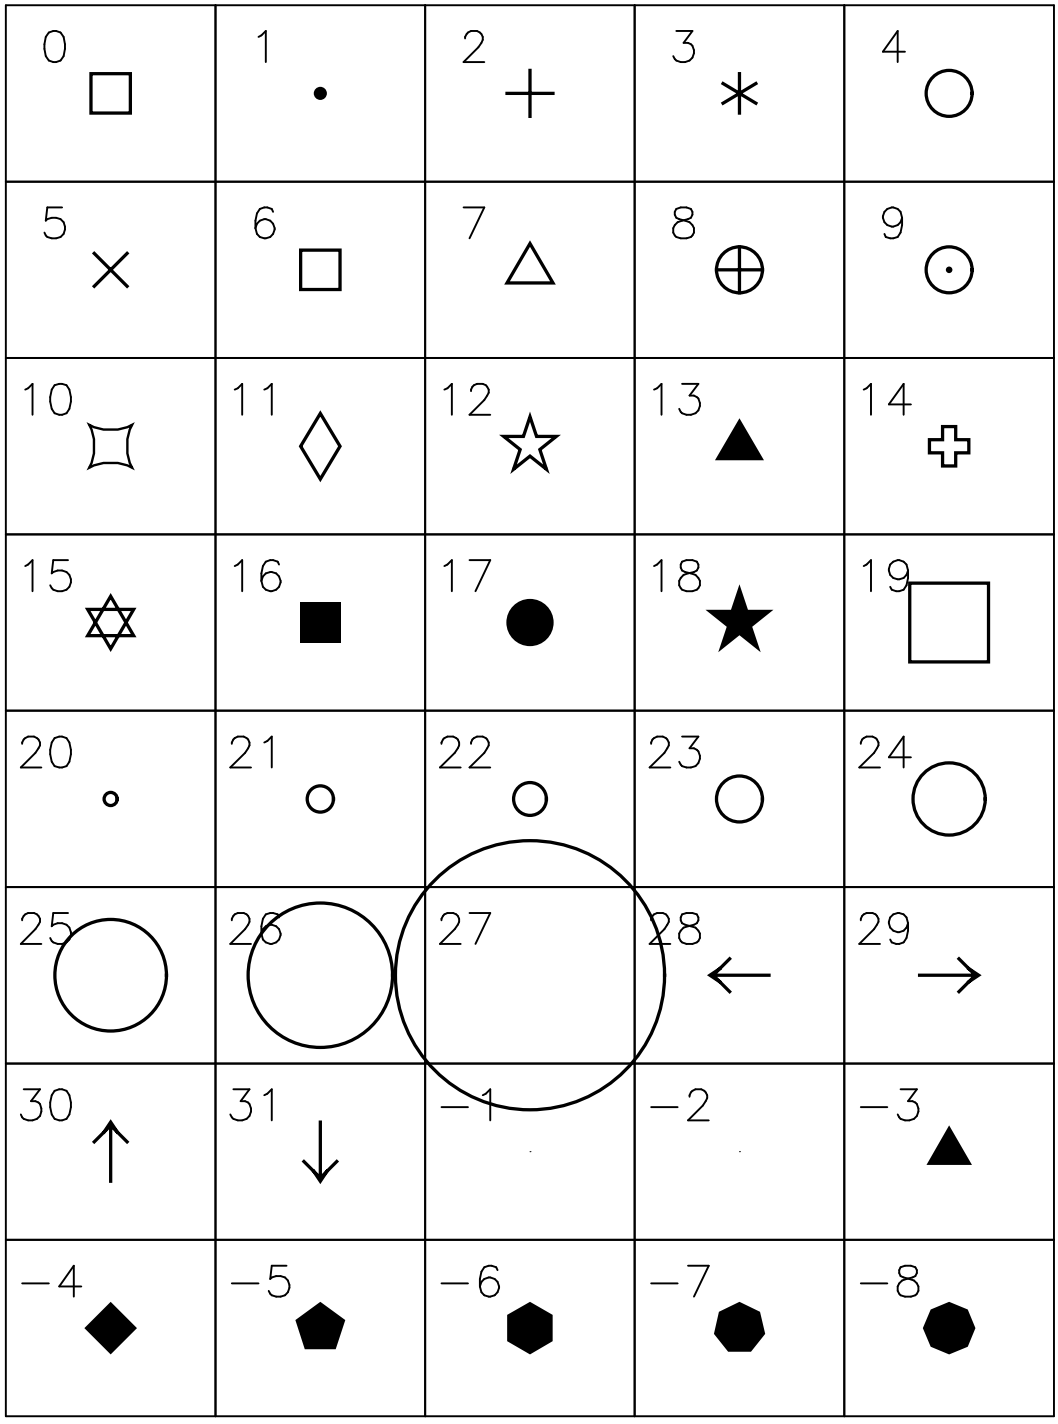

Table 42.1: Plotting Symbols at Height = 40.0

\u	Start a superscript or end a subscript
\d	Start a subscript or end a superscript. \u and \d must always be used in pairs
\b	Backspace (i.e., do not advance text pointer after plotting the previous character)
\fn	Switch to Normal font (1)
\fr	Switch to Roman font (2)
\fi	Switch to Italic font (3)
\fs	Switch to Script font (4)
\\	Backslash character (\)
\x	Multiplication sign (×)
\.	Centered dot (·)
\A	Angstrom symbol (Å)
\gx	Greek letter corresponding to roman letter x
\mn \mnn	Graph marker number n or nn (1-31)
\(nnnn)	Character number nnnn (1 to 4 decimal digits) from the Hershey character set; the closing parenthesis may be omitted if the next character is neither a digit nor ")". This makes a number of special characters (e.g., mathematical, musical, astronomical, and cartographical symbols) available.

Table 42.2: PGPLOT Escape Sequences.

Roman	a	b	g	d	e	z	y	h	i	k	l	m
Greek	α	β	γ	δ	ϵ	ζ	η	θ	ι	κ	λ	μ
Roman	n	c	o	p	r	s	t	u	f	x	q	w
Greek	ν	ξ	o	π	ρ	σ	τ	υ	ϕ	χ	ψ	ω
Roman	A	B	G	D	E	Z	Y	H	I	K	L	M
Greek	A	B	Γ	Δ	E	Z	H	Θ	I	K	Λ	M
Roman	N	C	O	P	R	S	T	U	F	X	Q	W
Greek	N	Ξ	O	Π	P	Σ	T	Υ	Φ	X	Ψ	Ω

Table 42.3: Conversion for the string "\g<r>" where "<r>" is a Roman character to the corresponding Greek character.

42.7 Structures

Quick Plot uses several structures to hold data. The structure that defines a line is a `qp_line_struct`

```
type qp_line_struct
  integer width         ! Line width.   Default = 1
  character(24) color   ! Line color.   Default = "black"
  character(16) pattern ! line pattern. Default = "solid"
end type
```

The `qp_symbol_struct` defines how symbols are drawn

```
type qp_symbol_struct
  character(16)  type        ! Default = "circle_dot"
  real(rp) height            ! Default = 6.0 (points)
  character(24)  color       ! Default = "black"
  character(16)  fill        ! Default = "solid_fill"
  integer  line_width        ! Default = 1
end type
```

The `qp_axis_struct` defines how axes are drawn

```
type qp_axis_struct
  character(80) label       ! Axis label.
  real(rp) min              ! Axis range left/bottom number.
  real(rp) max              ! Axis range right/top number.
  real(rp) number_offset    ! Offset in inches of numbering from the axis line.
                            !  Default = 0.05
  real(rp) label_offset     ! Offset in inches of the label from the numbering.
                            !  Default = 0.05
  character(24) label_color ! Default = "black"
  real(rp) major_tick_len   ! Length of the major ticks in inches. Def = 0.10
  real(rp) minor_tick_len   ! Length of the minor ticks in inches. Def = 0.06
  integer major_div         ! Number of major divisions. Default = 5
  integer major_div_nominal ! Nominal value. Def = 5.
  integer minor_div         ! Number of minor divisions. 0 = auto-choose. Default = 0
  integer minor_div_max     ! Maximum number for auto choose. Default = 5
  integer places            ! Places after the decimal point. Default = 0
  character(16) type        ! "LINEAR" (default), "LOG", or "CUSTOM".
  character(16) bounds      ! "GENERAL" (default), "ZERO_AT_END", etc.
  integer tick_side         ! +1 = draw to the inside (def), 0 = both, -1 = outside.
  integer number_side       ! +1 = draw to the inside, -1 = outside (default).
  logical draw_label        ! Draw the label? Default = True.
  logical draw_numbers      ! Draw the numbering? Default = True.
end type
```

The `%bounds` parameter sets how axis min and max values are calculated. Possible settings are:

```
"ZERO_AT_END"      ! Min or max value is set to zero.
"ZERO_SYMMETRIC"   ! Min and max chosen so that max = -min.
"GENERAL"          ! No restrictions.
"EXACT"            ! The inputted data min/max is used.
```

Finally, the `qp_plot_struct` is a container for the axis that make up a plot

```
type qp_plot_struct
  character(80) :: title = " "
  type (qp_axis_struct) x, y, x2, y2
```

```
    type (qp_axis_struct), pointer :: xx, yy  ! Pointer to axes used for plotting.
    logical :: draw_box    = .true.
    logical :: draw_title  = .true.
    logical :: draw_grid   = .true.
    logical :: x2_mirrors_x = .true.
    logical :: y2_mirrors_y = .true.
    logical :: xx_points_to_x
    logical :: yy_points_to_y
  end type
```

Chapter 43

HDF5

HDF5, which stands for "Hierarchical Data Format" version 5[HDF5], is a set of file formats designed to store and organize large amounts of data. HDF5 has been developed by scientists from a number of institutions including the National Center for Supercomputing Applications, the University of Illinois at Urbana-Champaign, and Sandia National Laboratories. Tools for viewing and editing HDF5 files are available from the HDF Group[HDF5]. Programs include `h5dump` and `HDFView` which can be used to directly view files. Interfaces so that HDF5 files can accessed via Java or Python also exist.

Bmad uses HDF5 for storing beam particle (positions, spin, etc.) and `grid_field` (§5.16.4) data. Storage details are given in sections §43.1 and §43.2 respectively. While `HDF5` defines how data is formatted, `HDF5` does not define the syntax for how data is to be stored. For that, *Bmad* uses the syntax defined by the `Beam Physics` extension to the `openPMD` standard[OpenPMD]. To understand the rest of this chapter, the reader should familiarize themselves with the `openPMD` and `Beam Physics` standards.

43.1 HDF5 Particle Beam Data Storage

The code for reading and writing beam data to/from HDF5 files is contained in the routines hdf5_read_beam and hdf5_write_beam.

As per the `openPMD`/`Beam Physics` standard, particle beam data is stored in a tree structure within a data file. The root "`group`" (tree node) for each bunch of the beam has the path within the file:

```
/data/%T/particles/
```

where `%T` is an integer.

For any bunch, parameters ("attributes") stored in the bunch root group are:

```
speciesType     ! The name of the particle species using the SpeciesType syntax.
totalCharge     ! Total bunch charge.
chargeLive      ! Charge of live particles.
numParticles    ! Number of particles.
```

The `SpeciesType` syntax defined by the `SpeciesType` extension to the `openPMD` standard is similar to the *Bmad* standard (§10.1) but there are differences. For one, the `SpeciesType` standard does not have an encoding for the charge state of atoms and molecules. Another difference is that for fundamental particles the names are case sensitive while for *Bmad* they are not (Note that atom and molecule names in *Bmad* are case sensitive).

What per-particle data is stored is determined by whether the bunch particles are photons or not. The following particle parameters are common for both types:

Beam Physics Parameter	*Bmad Equivalent*	*Notes*
time	-%vec(5) / (c %beta)	time - ref_time. See Eq. (17.3)
timeOffset	%t - time (beam physics)	reference time
totalMomentumOffset	%p0c	
sPosition	%s	See Fig. 16.2
weight	%charge	Macro bunch charge
branchIndex	%ix_branch	
elementIndex	%ix_ele	
locationInElement	%location	See below
particleStatus	%state	See the %state table in §37.1

The Bmad Equivalent column gives the conversion between the Beam Physics parameters and the coord_struct (§37.1) structure components (the coord_struct structure contains the particle position information). Parameters with a "%" suffix are coord_struct components and %vec(5) corresponds to the phase space z coordinate. The particleState is an integer which corresponds to the coord_struct %state component. A value of 1 indicates that the particle is alive (corresponding to the value of alive$) and any other value indicates that the particle is dead.

The locationInElement Beam Physics parameter is related to the coord_struct %location parameter via the following transformation:

locationInElement Value	*%location Value*
-1	upstream_end$
0	inside$
1	downstream_end$

For photons, additional per-particle data is:

Beam Physics Parameter	*Bmad Equivalent*
velocity/x, y, z	(%vx, %vy, %vz)
position/x, y, z	(%x, %y, %z)
pathLength	%path_len
photonPolarizationAmplitude/x, y	%field
photonPolarizationPhase/x, y	%phase

For clarity's sake, the %vec(1) through %vec(6) phase space coordinate components in the coord_struct have been replaced by %x, %vx, ..., %z, %vz in the above table

For non-photons, additional per-particle data is:

Beam Physics Parameter	*Bmad Equivalent*
momentum/x, y, z	%p0c×(%px, %py, sqrt((1 + %pz)^2 - %px^2 - %py^2))
totalMomentum	%p0c×%pz
position/x, y, z	(%x, %y, 0)
spin/x, y, z	%spin
chargeState	Derived from %species

For clarity's sake, the %vec(1) through %vec(6) phase space coordinate components in the coord_struct have been replaced by %x, %px, ..., %z, %pz in the above table. Notice that the Beam Physics z position (not to be confused with phase space z) is always zero by construction as shown in Fig. 16.2.

43.2 HDF5 Grid_Field Data Storage

The code for reading and writing grid_field data to/from HDF5 files is contained in the routines hdf5_read_grid_field and hdf5_write_grid_field.

As per the openPMD/Beam Physics standard, grid_field (§5.16.4 data is stored in a tree structure within a data file. The root "group" (tree node) for each grid_field has the path within the file:

/ExernalFieldmesh/%T/

where %T is an integer.

For any grid_field, parameters stored in the grid_field root group are:

Parameter in File	*Bmad Equivalent*	
gridGeometry	%geometry	
masterParameter	%master_parameter	
componentFieldScale	%field_scale	
fieldScale	{ %field_scale×master param value	If master parameter set.
	{ %field_scale	Otherwise.
harmonic	%harmonic	
RFphase	{ %harmonic×%phi0_fieldmap	For lcavity elements
	{ %harmonic×(0.25 - %phi0_fieldmap)	For all others.
eleAnchorPt	%ele_anchor_pt	
gridOriginOffset	%r0	
gridSpacing	%dr	
interpolationOrder	%interpolation_order	
gridLowerBound	%ptr%pt lower bound	
gridSize	%ptr%pt size	
fundamentalFrequency	ele%value(rf_frequency$)	
gridCurvatureRadius	ele%value(rho$)	

The Bmad Equivalent column gives the conversion between the Beam Physics parameters and the grid_field_struct structure components (that have a "%" prefix). The value for gridCurvatureRadius is set to the value of rho of the associated lattice element if %curved_ref_frame is True.

Notice that the masterParameter attribute is not part of the standard. If not present, which could happen if a file is created by non-*Bmad* code, the default is a blank string indicating no master parameter. If masterParameter is set in the data file, there is a potential problem in that it may not be possible to calculate %field_scale if the value of the master parameter is not equal to the value when the data was written. To get around this, if the non-standard masterParameter is present, the value of the non-standard componentFieldScale (which has a default value of one) will be used to set %field_scale and the fieldScale parameter will be ignored. If masterParameter is not present, componentFieldScale is ignored and %field_scale is set from the value of fieldScale.

When reading a data file, the setting of grid_field%field_type is determined by what data is stored in the file. If both electric and magnetic field data is present, %field_type is set to mixed$. Otherwise, %field_type is set to magnetic$ if magnetic field data is present or electric$ if electric field data is present.

The correspondence between the `gridGeometry` parameter and the `grid_field%geometry` component is

gridGeometry Value	*%geometry Value*
"rectangular"	xyz$
"cylindrical"	rotationally_symmetric_rz$

Chapter 44

Bmad Library Routine List

Below are a list of *Bmad* and sim_utils routines sorted by their functionality. Use the `getf` and `listf` (§30.3) scripts for more information on individual routines. This list includes low level routines that are not generally used in writing code for a program but may be useful in certain unique situations. Excluded from the list are very low level routines that are solely meant for *Bmad* internal use.

Routine Type	*Section*
Beam: Low Level Routines	44.1
Beam: Tracking and Manipulation	44.2
Branch Handling	44.3
Coherent Synchrotron Radiation (CSR)	44.4
Collective Effects	44.5
Custom and Hook Routines	44.6
Electro-Magnetic Fields	44.7
HDF Read/Write	44.8
Helper Routines: File, System, and IO	44.9
Helper Routines: Math (Except Matrix)	44.10
Helper Routines: Matrix	44.11
Helper Routines: Miscellaneous	44.12
Helper Routines: String Manipulation	44.13
Helper Routines: Switch to Name	44.14
Inter-Beam Scattering (IBS)	44.15
Lattice: Informational	44.18
Lattice: Element Manipulation	44.16
Lattice: Geometry	44.17
Lattice: Low Level Stuff	44.19
Lattice: Manipulation	44.20
Lattice: Miscellaneous	44.21
Lattice: Nametable	44.22
Lattice: Reading and Writing Files	44.23
Matrices	44.24
Matrix: Low Level Routines	44.25
Measurement Simulation Routines	44.26
Multipass	44.27
Multipoles	44.28
Optimizers (Nonlinear)	44.29
Overload Equal Sign	44.30
Particle Coordinate Stuff	44.31
Photon Routines	44.32
PTC Interface Routines	44.33
Quick Plot	44.34
Spin	44.35
Transfer Maps: Routines Called by make_mat6	44.36
Transfer Maps: Complex Taylor Maps	44.37
Transfer Maps: Taylor Maps	44.38
Transfer Maps: Driving Terms	44.39
Tracking: Tracking and Closed Orbit	44.40
Tracking: Low Level Routines	44.41
Tracking: Mad Routines	44.42
Tracking: Routines Called by track1	44.40
Twiss and Other Calculations	44.44
Twiss: 6-Dimensional	44.45
Wakefields	44.46
C/C++ Interface	44.47

44.1 Beam: Low Level Routines

The following helper routines are generally not useful for general use.

bend_edge_kick (ele, param, particle_at, orb, mat6, make_matrix, track_spin)
Subroutine to track through the edge field of an sbend. Reverse tracking starts with the particle just outside the bend and

init_spin_distribution (beam_init, bunch, ele)
Initializes a spin distribution according to init_beam%spin

order_particles_in_z (bunch)
Routine to order the particles longitudinally in terms of decreasing %vec(5). That is from large z (head of bunch) to small z.

track1_bunch (bunch, ele, err, centroid, direction, bunch_track)
Routine to track a bunch of particles through an element.

track1_bunch_hom (bunch, ele, direction, bunch_track)
Routine to track a bunch of particles through an element.

44.2 Beam: Tracking and Manipulation

See §37.16 for a discussion of using a collection of particles to simulate a bunch.

bbi_kick (x, y, sigma, nk, dnk)
Routine to compute the normalized kick due to the beam-beam interaction using the normalized position for input.

calc_bunch_params (bunch, bunch_params, error, print_err, n_mat, is_time_coords, ele)
Finds all bunch parameters defined in bunch_params_struct, both normal-mode and projected

calc_bunch_params_slice (bunch, bunch_params, plane, slice_center, slice_spread, err, print_err, is_time_coords, ele)
Finds all bunch parameters for a slice through the beam distribution.

calc_bunch_sigma_matrix_etc (particle, charge, bunch_params, is_time_coords, ele)
Routine to find the sigma matrix elements of a particle distribution.

calc_emittances_and_twiss_from_sigma_matrix(sigma_mat, bunch_params, error, print_err, n_mat)
Routine to calc emittances and Twiss function from a beam sigma matrix.

init_beam_distribution (ele, param, beam_init, beam, err_flag, modes, beam_init_set, print_p0c_shift_warning, conserve_momentum)
Routine to initialize a distribution of particles matched to the Twiss parameters, centroid position, and Energy - z correlation

init_bunch_distribution (ele, param, beam_init, ix_bunch, bunch, err_flag, modes, beam_init_used, print_p0c_shift_warning, conserve_momentum)
Routine to initialize either a random or tail-weighted distribution of particles.

read_beam_file (file_name, beam, beam_init, err_flag, ele, print_mom_shift_warning, conserve_momentum)
Subroutine to read in a beam definition file.

reallocate_beam (beam, n_bunch, n_particle, extend)
Routine to reallocate memory within a beam_struct.

reallocate_bunch (bunch, n_particle, save)
Subroutine to reallocate particles within a bunch_struct.

track_beam (lat, beam, ele1, ele2, err, centroid, direction, bunch_tracks)
Routine to track a beam of particles from the end of lat%ele(ix1) Through to the end of lat%ele(ix2).

track_bunch (lat, bunch, ele1, ele2, err, centroid, direction, bunch_track)
Subroutine to track a particle bunch from the end of ele1 Through to the end of ele2. Both must be in the same lattice branch.

track_bunch_time (bunch, branch, t_end, s_end, dt_step, extra_field)
Routine to track a particle bunch for a given time step (or if the ! particle position exceeds s_end).

write_beam_file (file_name, beam, new_file, file_format, lat, alive_only)
Routine to write a beam file.

write_beam_floor_positions (file_name, beam, ele, new_file)
Routine to write a file of beam positions in global floor coordinates.

44.3 Branch Handling Routines

allocate_branch_array (lat, upper_bound)
Routine to allocate or re-allocate an branch array. The old information is saved.

transfer_branch (branch1, branch2)
Routine to set branch2 = branch1. This is a plain transfer of information not using the overloaded equal.

transfer_branches (branch1, branch2)
Routine to set branch2 = branch1. This is a plain transfer of information not using the overloaded equal.

44.4 Coherent Synchrotron Radiation (CSR)

csr_bin_particles (ele, particle, csr, err_flag)
Routine to bin the particles longitudinally in s.

csr_bin_kicks (ele, ds_kick_pt, csr, err_flag)
Routine to cache intermediate values needed for the csr calculations.

i_csr (kick1, i_bin, csr) result (i_this)
Routine to calculate the CSR kick integral.

44.5 Collective Effects

touschek_lifetime (mode, Tl, lat)
Routine to calculate the Touschek lifetime for a lat.

44.6 Custom Routines

apply_element_edge_kick_hook (orb, fringe_info, track_ele, param, finished, mat6, make_matrix, rf_time)
Routine that can be customized to track through the edge field of an element. This routine is always called by apply_element_edge_kick.

check_aperture_limit_custom (orb, ele, particle_at, param, err_flag)
Routine to check if an orbit is outside an element's aperture. Used when `ele%aperture_type` is set to `custom$`

ele_geometry_hook (floor0, ele, floor, finished, len_scale)
Routine that can be customized to calculate the floor position of an element.

ele_to_fibre_hook (ele, ptc_fibre, param, use_offsets, err_flag)
Routine that can be customized for creating a PTC fibre from a Bmad element. This routine is always called by ele_to_fibre.

em_field_custom(ele, param, s_rel, orbit, local_ref_frame, field, calc_dfield, err_flag, calc_potential, use_overlap, grid_allow_s_out_of_bounds, rf_time, used_eles)
Custom routine for calculating fields.

init_custom (lat)
Routine for initializing custom elements or elements that do custom calculations.

make_mat6_custom (ele, param, start_orb, end_orb, err_flag)
Routine for custom calculations of the 6x6 transfer matrices.

radiation_integrals_custom (lat, ir, orb, rad_int1, err_flag)
User supplied routine to calculate the synchrotron radiation integrals for a custom element.

time_runge_kutta_periodic_kick_hook (orbit, ele, param, stop_time, init_needed)
Custom routine to add a kick to a particle at periodic times.

track1_bunch_hook (bunch, ele, err, centroid, direction, finished, bunch_track)
Routine that can be customized for tracking a bunch through a single element.

track1_custom (orbit, ele, param, err_flag, finished, track)
Dummy routine for custom tracking.

track1_postprocess (start_orb, ele, param, end_orb)
Dummy routine for post processing after the track1 routine is done.

track1_preprocess (start_orb, ele, param, err_flag, finished, radiation_included, track)
Dummy routine for pre processing at the start of the track1 routine.

track1_spin_custom (start, ele, param, end, err_flag, make_quaternion)
Dummy routine for custom spin tracking. This routine needs to be replaced for a custom calculation.

track1_wake_hook (bunch, ele, finished)
Routine that can be customized for tracking through a wake.

wall_hit_handler_custom (orb, ele, s)
This routine is called by the Runge-Kutta integrator odeint_bmad when a particle hits a wall.

44.7 Electro-Magnetic Fields

em_field_calc (ele, param, s_pos, orbit, local_ref_frame, field, calc_dfield, err_flag, calc_potential, use_overlap, grid_allow_s_out_of_bounds, rf_time, used_eles, print_err)
Routine to calculate the E and B fields for an element.

em_field_custom(orbit, ele, param, s1_body, s2_body, err_flag, track, mat6, make_matrix)
Custom routine for calculating fields.

44.8 HDF Read/Write

hdf5_write_attribute_string(root_id, attrib_name, string, error)
Routine to create an HDF5 attribute whose value is a string.

hdf5_open_file (file_name, action, file_id, error, verbose)
Routine to open an HDF5 file.

hdf5_open_object(root_id, object_name, info, error, print_error) result (obj_id)
Routine to open an existing group or dataset.

hdf5_close_object(obj_id, info)
Routine to close a group or dataset.

hdf5_exists (root_id, object_name, error, print_error) result (exists)
Routine to check if a object with object_name exists relative to root_id.

hdf5_open_group (root_id, group_name, error, print_error) result (g_id)
Routine to open an existing group.

hdf5_open_dataset(root_id, dataset_name, error, print_error) result (obj_id)
Routine to open an existing group or dataset.

hdf5_num_attributes(root_id) result (num)
Routine to return the number of attributes associated with a group or dataset.

hdf5_get_attribute_by_index(root_id, attrib_indx, attrib_id, attrib_name)
Routine to return the ID and name of an attribute given the attribute's index number. This routine is useful for looping over all the attributes in a group or dataset.

hdf5_attribute_info(root_id, attrib_name, error, print_error) result (info)
Routine to return information on an attribute given the attribute name and encompassing group.

hdf5_object_info (root_id, obj_name, error, print_error) result (info)
Routine to get information on an object (group or dataset).

hdf5_read_attribute_int(root_id, attrib_name, attrib_value, error, print_error, dflt_value)
Routine to read a integer attribute value or array.

hdf5_read_attribute_real(root_id, attrib_name, attrib_value, error, print_error, dflt_value)
Routine to read a real attribute value or array.

hdf5_read_attribute_alloc_string(root_id, attrib_name, string, error, print_error)
Routine to read a string attribute. Also see: hdf5_read_attribute_string

hdf5_read_attribute_string(root_id, attrib_name, string, error, print_error)
Routine to read a string attribute. Also see: hdf5_read_attribute_alloc_string

hdf5_write_dataset_real(root_id, dataset_name, value, error)
Routine to create a dataset of reals.

hdf5_write_dataset_int(root_id, dataset_name, value, error)
Routine to create a dataset of integers.

hdf5_read_beam (file_name, beam, error, ele, pmd_header, print_mom_shift_warning, conserve_momentum)
Routine to read a beam data file.

hdf5_write_grid_field (file_name, ele, g_field, err_flag)
Routine to create an hdf5 file encoding an array of grid_field structures. Note: Conventionally, the file name should have an ".h5" suffix.

pmd_write_int_to_dataset (root_id, dataset_name, bmad_name, unit, array, error)

pmd_write_int_to_pseudo_dataset(root_id, dataset_name, descrip_name, unit, value, v_shape, error)

pmd_write_real_to_dataset (root_id, dataset_name, bmad_name, unit, array, error)

pmd_write_real_to_pseudo_dataset (root_id, dataset_name, descrip_name, unit, value, v_shape, error)

pmd_write_complex_to_dataset (root_id, dataset_name, bmad_name, unit, array, error)

pmd_write_units_to_dataset (root_id, dataset_name, descrip_name, unit, error)

pmd_read_int_dataset (root_id, name, conversion_factor, array, error)

pmd_read_real_dataset (root_id, name, conversion_factor, array, error)

pmd_read_complex_dataset (root_id, name, conversion_factor, array, error)

hdf5_read_grid_field (file_name, ele, g_field, err_flag, pmd_header, combine)
Routine to read an hdf5 file that encodes an array of grid_field structures.

hdf5_write_beam (file_name, bunches, append, error, lat, alive_only)
Routine to write particle positions of a beam to an HDF5 binary file.

44.9 Helper Routines: File, System, and IO

append_subdirectory (dir, sub_dir, dir_out, err)
Routine to combine a directory specification with a subdirectory specification to form a complete directory

cesr_iargc ()
Platform independent function to return the number of command line arguments. Use this with cesr_getarg.

cesr_getarg (i_arg, arg)
Platform independent function to return the i'th command line argument. Use this with cesr_iargc.

dir_close ()
Routine to close a directory that was opened with dir_open. Also see dir_read.

dir_open (dir_name) result (opened)
Routine to open a directory to obtain a list of its files. Use this routine with dir_read and dir_close.

dir_read (file_name) result (valid)
Routine to get the names of the files in a directory. Use this routine with dir_open and dir_close.

file_suffixer (in_file_name, out_file_name, suffix, add_switch)
Routine to add/replace a suffix to a file name.

get_tty_char (this_char, wait, flush)
Routine for getting a single character from the terminal. Also see: get_a_char

get_a_char (this_char, wait, ignore_this)
Routine for getting a single character from the terminal. Also see: get_tty_char

get_file_time_stamp (file, time_stamp)
Routine to get the "last modified" time stamp for a file.

lunget()
Function to return a free file unit number to be used with an open statement.

milli_sleep (milli_sec)
Routine to pause the program for a given number of milli-seconds.

out_io (...)
Routine to print to the terminal for command line type programs. The idea is that for programs with a gui this routine can be easily replaced with another routine.

out_io_called (level, routine_name)
Dummy routine for linker. See out_io for more details.

out_io_end ()
Dummy routine for linker. See out_io for more details.

out_io_line (line)
Dummy routine for linker. See out_io for more details.

output_direct (file_unit, print_and_capture, min_level, max_level, set, get)
Routine to set where the output goes when out_io is called. Output may be sent to the terminal screen, written to a file, or both. Also can be used to restrict output verbosity.

read_a_line (prompt, line_out, trim_prompt, prompt_color, prompt_bold, history_file)
Routine to read a line of input from the terminal. The line is also add to the history buffer so that the up-arrow

skip_header (ix_unit, error_flag)
Routine to find the first line of data in a file.

splitfilename(filename, path, basename, is_relative) result (ix_char)
Routine to take filename and splits it into its constituent parts, the directory path and the base file name.

system_command (line, err_flag)
Routine to execute an operating system command from within the program.

type_this_file (filename)
Routine to type out a file to the screen.

44.10 Helper Routines: Math (Except Matrix)

complex_error_function (wr, wi, zr, zi)
This routine evaluates the function w(z) in the first quadrant of the complex plane.

cross_product (a, b)
Returns the cross product of a x b

linear_fit (x, y, n_data, a, b, sig_a, sig_b)
Routine to fit to y = A + B x

modulo2 (x, amp)
Function to return y = x + 2 * n * amp, n is an integer, such that y is in the interval [-amp, amp].

ran_engine (set, get, ran_state)
Routine to set what random number generator algorithm is used. If this routine is never called then pseudo_random$ is used.

ran_gauss (harvest)
Routine to return a Gaussian distributed random number with unit sigma.

ran_gauss_converter (set, set_sigma_cut, get, get_sigma_cut, ran_state)
Routine to set what conversion routine is used for converting uniformly distributed random numbers to Gaussian distributed random numbers.

ran_seed_put (seed, mpi_offset)
Routine to seed the random number generator.

ran_seed_get (seed)
Routine to return the seed used for the random number generator.

ran_uniform (harvest)
Routine to return a random number uniformly distributed in the interval [0, 1]. This routine uses the same algorithm as ran from

spline_akima (spline, ok)
Given a set of (x,y) points we want to interpolate between the points. This routine computes the semi-hermite cubic spline developed by akima

spline_evaluate (spline, x, ok, y, dy)
Routine to evaluate a spline at a set of points.

super_ludcmp (a,indx,d, err)
This routine is essentially ludcmp from Numerical Recipes with the added feature that an error flag is set instead of bombing the program when there is a problem.

44.11 Helper Routines: Matrix

mat_eigen (mat, eigen_val, eigen_vec, error, print_err)
Routine for determining the eigen vectors and eigen values of a matrix.

mat_inverse (mat, mat_inv, vec0, vec0_inv, ok, print_err)
Routine to take the inverse of a square matrix.

mat_make_unit (mat)
routine to create a unit matrix.

mat_rotation (mat, angle, bet_1, bet_2, alph_1, alph_2)
Routine to construct a 2x2 rotation matrix for translation from point 1 to point 2.

mat_symplectify (mat_in, mat_symp, p0_ratio, r_root)
Routine to form a symplectic matrix that is approximately equal to the input matrix.

mat_symp_error (mat, p0_ratio, err_mat) result (error)
Routine to check the symplecticity of a square matrix

mat_symp_conj (mat) result (mat_conj)
Routine to take the symplectic conjugate of a square matrix.

mat_symp_decouple (t0, stat, u, v,
ubar, vbar, g, twiss1, twiss2, gamma, type_out)
Routine to find the symplectic eigen–modes of the one turn 4x4 coupled transfer matrix T0.

mat_type (mat, nunit, header, num_form, lines, n_lines)
Routine to output matrices to the terminal or to a file

44.12 Helper Routines: Miscellaneous

date_and_time_stamp (string, numeric_month, include_zone)
Routine to return the current date and time in a character string.

err_exit(err_str)
Routine to first show the stack call list before exiting. This routine is typically used when a program detects an error condition.

integer_option (integer_default, opt_integer)
Function to return True or False depending upon the state of an optional integer.

is_false (param) result (this_false)
Routine to translate from a real number to a boolian True or False. Translation: 0 = False, nonzero = True.

is_true (param) result (this_true)
Routine to translate from a real number to a boolian True or False. Translation: 0 = False, nonzero = True.

logic_option (logic_default, opt_logic)
Function to return True or False depending upon the state of an optional logical.

re_allocate (ptr_to_array, n, exact)
Function to reallocate a pointer to an array of strings, integers, reals, or logicals.

re_associate (array, n)
Function to reassociate an allocatable array of strings, integers, reals, or logicals.

real_option (real_default, opt_real)
Function to return True or False depending upon the state of an optional real.

string_option (string_default, opt_string)
Routine to return True or False depending upon the state of an optional string.

44.13 Helper Routines: String Manipulation

all_pointer_to_string (a_ptr, err) result (str)
Routine to turn the value pointed to by an all_pointer_struct into a string for printing.

downcase (str_in) result (str_out)
Routine to convert a string to lower case.

downcase_string (string)
Routine to convert a string to lowercase:

index_nocase (string, match_str) result (indx)
Function to look for a sub-string of string that matches match_str. This routine is similar to the fortran INDEX function

int_str(int, width) result (str)
Routine to return a string representation of an integer.

is_alphabetic (string, valid_chars) result (is_alpha)
Function to tell if a string has all alphabetical characters. Spaces are counted as not alphabetic

is_integer (string, int, delims, ix_word)
Function to tell if the first word in a string is a valid integer.

is_logical (string, ignore) result (good)
Function to test if a string represents a logical. Accepted possibilities are (individual characters can be either case):

is_real (string, ignore, real_num) result (good)
Function to test if a string represents a real number.

location_decode (string, array, ix_min, num, names, exact_case, can_abbreviate, print_err)
Subroutine to set a list of locations in a logical array to True.

location_encode1 (string, loc, exists, ix_min, names, separator, err_flag)
Routine to encode a list of locations. This routine is overloaded by the routine: location_encode.

logic_str(logic) result (str)
Routine to return a string representation (T/F) of a logical.

match_reg (str, pat)
Function for matching with regular expressions. Note: strings are trimmed before comparison.

match_wild (string, template) result (this_match)
Function to do wild card matches. Note: trailing blanks will be discarded before any matching is done.

match_word (string, names, ix, exact_case, can_abbreviate, matched_name)
Routine to match the first word in a string against a list of names. Abbreviations are accepted.

on_off_logic (logic, true_str, false_str) result (name)
Function to return the string "ON" or "OFF".

ordinal_str(n) result (str)
Routine to return a string representing the ordinal position of n. EG n = 1 -> "1st", n = 2 -> "2nd", etc.

parse_fortran_format (format_str, n_repeat, power, descrip, width, digits)
Routine to parse a Fortran edit descriptor. This routine assumes that format_str will be a edit descriptor for a single entity like '3f10.6'.

quote(str) result (q_str)
Function to put double quote marks (") around a string. The string will not be modified if there are already quote marks of either kind.

quoten(str, delim) result (q_str)
Function to put double quote marks (") around each string in an array and return the concatenated string with a delimitor between the strings.

real_num_fortran_format (number, width, n_blanks) result (fmt_str)
Routine to find a "nice" edit descriptor format for a real number. The format will either be "es" for numbers that are very small or very large or "f".

real_str(r_num, n_signif, n_decimal) result (str)
Routine to return a string representing a real number. Trailing zeros will be supressed.

real_to_string (real_num, width, n_signif, n_decimal) result (str)
Routine to turn a real number into a string for printing. Printing the number without an explicit exponent is preferred.

reals_to_string (real_arr, width, n_blank, n_signif, n_decimal) result (str)
Routine to turn a n array of reals into a string for printing.

reals_to_table_row (real_arr, width, n_decimal, n_blank) result (str)
Routine to turn an array of real numbers into a string for printing tables. Fixed format is preferred and floating format will only be used if necessary.

str_match_wild(str, pat) result (a_match)
Function to match a character string against a regular expression pattern.

str_set(str_out, str_in)
Routine to set a variable length string. Trailing blanks will be trimmed.

string_to_int (line, default, err_flag, err_print_flag)
Routine to convert a string to an integer.

string_to_real (line, default, err_flag, err_print_flag)
Routine to convert a string to an real.

string_trim(in_string, out_string, word_len)
Routine to trim a string of leading blanks and/or tabs and also to return the length of the first word.

string_trim2 (in_str, delimitors, out_str, ix_word, delim, ix_next)
Routine to trim a string of leading delimiters and also to return the length of the first word.

str_downcase (dst, src)
Routine to convert a string to down case.

str_substitute (string, str_match, str_replace, do_trim, ignore_escaped)
Routine to substitute all instances of one sub-string for another in a string

to_str(num, max_signif) result (string)
Routine to return the string representation of a number.

unquote (str_in) result (str_out)
Routine to remove quotation marks at the ends of a string. Quotation marks will only be removed if they match at both ends.

upcase (str_in) result (str_out)
Routine to convert a string to upper case.

upcase_string (string)
Routine to convert a string to uppercase:

44.14 Helper Routines: Switch to Name

coord_state_name (coord_state, one_word) result (state_str)
Routine to return the string representation of a coord%state state.

44.15 Inter-Beam Scattering (IBS)

bane1(ele, coulomb_log, rates, n_part)
This is an implementation of equations 10-15 from "Intrabeam scattering formulas for high energy beams" Kubo,Mtingwa,Wolski.

bjmt1(ele, coulomb_log, rates, n_part)
This is an implementation of equations 1-9 from "Intrabeam scattering formulas for high energy beams" Kubo,Mtingwa,Wolski.

cimp1(ele, coulomb_log, rates, n_part)
This is an implementation of equations 34,38-40 from "Intrabeam scattering formulas for high energy beams" Kubo,Mtingwa,Wolski.

ibs_lifetime(lat,ibs_sim_params,maxratio,lifetime,granularity)
This module computes the beam lifetime due to the diffusion process.

mpxx1(ele, coulomb_log, rates, n_part)
Modified Piwinski, further modified to treat Coulomb Log in the same manner as Bjorken-Mtingwa, CIMP, Bane, Kubo & Oide, etc.

mpzt1(ele, coulomb_log, rates, n_part)
Modified Piwinski with Zotter's integral. This is Piwinski's original derivation, generalized to take the derivatives of the optics functions.

44.16 Lattice: Element Manipulation

These routine are for adding elements, moving elements, etc.

add_lattice_control_structs (ele, n_add_slave, n_add_lord, n_add_slave_field, n_add_lord_field, add_at_end)
Routine to adjust the control structure of a lat so that extra control elements can be added.

add_superimpose (lat, super_ele_in, ix_branch, err_flag, super_ele_out, save_null_drift, create_jumbo_slave, ix_insert, mangle_slave_names, wrap)
Routine to make a superimposed element.

attribute_bookkeeper (ele, force_bookkeeping)
Routine to make sure the attributes of an element are self-consistent.

autoscale_phase_and_amp(ele, param, err_flag, scale_phase, scale_amp, call_bookkeeper)
Routine to set the phase offset and amplitude scale of the accelerating field. This routine works on lcavity, rfcavity and e_gun elements.

create_element_slice (sliced_ele, ele_in, l_slice, offset, param, include_upstream_end, include_downstream_end, err_flag, old_slice, orb_in)
Routine to transfer the %value, %wig_term, and %wake%lr information from a superposition lord to a slave when the slave has only one lord.

create_field_overlap (lat, lord_name, slave_name, err_flag)
Subroutine to add the bookkeeping information to a lattice for an element's field overlapping another element.

create_group (lord, contrl, err)
Routine to create a group control element.

create_girder (lat, ix_girder, contrl, girder_info, err_flag)
Routine to add the controller information to slave elements of an girder_lord.

create_overlay (lord, contrl, err)
Routine to add the controller information to slave elements of an overlay_lord.

create_planar_wiggler_model (wiggler_in, lat, err_flag, print_err)
Routine to create series of bend and drift elements to serve as a model for a wiggler. This routine uses the mrqmin nonlinear optimizer to vary the parameters in the wiggler

insert_element (lat, insert_ele, ix_ele, ix_branch, orbit)
Routine to Insert a new element into the tracking part of the lat structure.

make_hybrid_lat (lat_in, lat_out, use_taylor, orb0_arr)
Routine to concatenate together elements to make a hybrid lat

new_control (lat, ix_ele, ele_name)
Routine to create a new control element.

pointer_to_attribute (ele, attrib_name, do_allocation, a_ptr, err_flag, err_print_flag, ix_attrib, do_unlink)
Returns a pointer to an attribute of an element with name attrib_name.

pointers_to_attribute (lat, ele_name, attrib_name, do_allocation, ptr_array, err_flag, err_print_flag, eles, ix_attrib, do_unlink)
Returns an array of pointers to an attribute with name attrib_name within elements with name ele_name.

pointer_to_branch
Routine to return a pointer to a lattice branch.

pointer_to_next_ele (this_ele, offset, skip_beginning, follow_fork) result (next_ele)
Function to return a pointer to the Nth element relative to this_ele in the array of elements in a lattice branch.

pointer_to_ele (lat, ix_ele, ix_branch) result (ele_ptr)
pointer_to_ele (lat, ele_loc_id) result (ele_ptr)
Routine to point to a given element.

pointer_to_ele_multipole (ele, a_pole, b_pole, ksl_pole, pole_type)
Routine to point to the appropriate magnetic or electric poles in an element.

pointer_to_element_at_s (branch, s, choose_max, err_flag, s_eff, position, print_err) result (ele)
Function to return a pointer to the element at position s.

remove_eles_from_lat (lat, check_sanity)
Routine to remove an elements from the lattice.

set_ele_attribute (ele, set_string, err_flag, err_print_flag, set_lords, err_id)
Routine to set an element's attribute.

set_ele_status_stale (ele, status_group, set_slaves)
Routine to set a status flags to stale in an element and the corresponding ones for any slaves the element has.

set_status_flags (bookkeeping_state, stat)
Routine to set the bookkeeping status block.

split_lat (lat, s_split, ix_branch, ix_split, split_done, add_suffix, check_sanity, save_null_drift, err_flag, choose_max, ix_insert)
Routine to split a lat at a point.

value_of_attribute (ele, attrib_name, err_flag, err_print_flag, err_value) result (value)
Returns the value of an element attribute.

44.17 Lattice: Geometry

bend_shift(position1, g, delta_s, w_mat, ref_tilt) result(position2)
Function to shift a particle's coordinates from one coordinate frame of reference to another within a bend with curvature g and reference tilt ref_tilt.

coords_curvilinear_to_floor (xys, branch, err_flag) result (global)
Routine to find the global position of a local lab (x, y, s) position. s = position from beginning of lattice branch.

coords_floor_to_curvilinear (floor_coords, ele0, ele1, status, w_mat) result (local_coords)
Given a position in global "floor" coordinates, return local curvilinear (ie element) coordinates for an appropriate element, ele1, near ele0. That is, the s-position of local_coords will be within

coords_floor_to_relative (floor0, global_position, calculate_angles, is_delta_position) result (local_position)
Returns local floor position relative to floor0 given a global floor position. This is an essentially an inverse of routine coords_relative_to_floor.

coords_relative_to_floor (floor0, dr, theta, phi, psi) result (floor1)
Starting from a given reference frame and given a shift in position, return the resulting reference frame.

coords_local_curvilinear_to_floor (local_position, ele, in_body_frame, w_mat, calculate_angles, relative_to) result (global_position)
Given a position local to ele, return global floor coordinates.

coords_floor_to_local_curvilinear(global_position, ele, status, w_mat, relative_to) result(local_position)
Given a position in global coordinates, return local curvilinear coordinates in ele relative to floor0

ele_geometry (floor_start, ele, floor_end, len_scale, ignore_patch_err)
Routine to calculate the physical (floor) placement of an element given the placement of the preceding element. This is the same as the MAD convention.

ele_geometry_with_misalignments (ele, len_scale) result (floor)
Routine to calculate the element body coordinates (that is, coordinates with misalignments) for an element at some distance s_offset from the upstream end.

ele_misalignment_l_s_calc (ele, l_mis, s_mis)
Calculates transformation vector L_mis and matrix S_mis due to misalignments for an ele Used to transform coordinates and vectors relative to the center of the element

floor_angles_to_w_mat (theta, phi, psi, w_mat, w_mat_inv)
Routine to construct the W matrix that specifies the orientation of an element in the global "floor" coordinates. See the Bmad manual for more details.

floor_w_mat_to_angles (w_mat, theta, phi, psi, floor0)
Routine to construct the angles that define the orientation of an element in the global "floor" coordinates from the W matrix. See the Bmad manual for more details.

lat_geometry (lat)
Routine to calculate the physical placement of all the elements in a lattice. That is, the physical machine layout on the floor.

orbit_to_local_curvilinear (orbit, ele, z_direction, relative_to) result (local_position)
Routine to return the local curvilinear position and orientation of a particle.

patch_flips_propagation_direction (x_pitch, y_pitch) result (is_flip)
Routine to determine if the propagation direction is flipped in a patch. This is true if the tranformation matrix element S(3,3) = cos(x_pitch) * cos(y_pitch)

s_calc (lat)
Routine to calculate the longitudinal distance S for the elements in a lat.

update_floor_angles (floor, floor0)
Routine to calculate floor angles from its W matrix.

w_mat_for_bend_angle (angle, ref_tilt, r_vec) result (w_mat)
Routine to compute the W matrix for the angle transformation in a bend. Using the notation in the Bmad manual:

w_mat_for_x_pitch (x_pitch, return_inverse)
Routine to return the transformation matrix for an x_pitch.

w_mat_for_y_pitch (y_pitch, return_inverse)
Routine to return the transformation matrix for an y_pitch.

w_mat_for_tilt (tilt, return_inverse)
Routine to return the transformation matrix for an tilt.

44.18 Lattice: Informational

at_this_ele_end (now_at, where_at) result (is_at_this_end)
Routine to determine if an aperture or fringe field is present. When tracking use this routine conjunction with physical_ele_end.

attribute_free (ix_ele, attrib_name, lat, err_print_flag, except_overlay) result (free)
attribute_free (ele, attrib_name, lat, err_print_flag, except_overlay) result (free)
attribute_free (ix_ele, ix_branch, attrib_name, lat, err_print_flag, except_overlay) result (free)
Overloaded function to check if an attribute is free to vary.

attribute_index (ele, name, full_name)
Function to return the index of an attribute for a given element type and the name of the attribute

attribute_name (ele, ix_att)
Function to return the name of an attribute for a particular type of element.

attribute_type (attrib_name, ele) result (attrib_type)
Routine to return the type (logical, integer, real, or named) of an attribute.

branch_name(branch) result (name)
Routine to return a string with the lattice branch name encoded. This routine is useful for error messages.

check_if_s_in_bounds (branch, s, err_flag, translated_s, print_err)
Routine to check if a given longitudinal position s is within the bounds of a given branch of a lattice.

element_at_s (lat, s, choose_max, ix_branch, err_flag, s_eff, position) result (ix_ele)
Routine to return the index of the element at position s.

ele_loc_to_string (ele, show_branch0, parens) result (str)
Routine to encode an element's location into a string.

equivalent_taylor_attributes (ele_taylor, ele2) result (equiv)
Routine to see if two elements are equivalent in terms of their attributes so that their Taylor Maps, if they existed, would be the same.

find_element_ends (ele, ele1, ele2, ix_multipass)
Routine to find the end points of an element.

get_slave_list (lord, slaves, n_slave)
Routine to get the list of slaves for an element.

key_name (key_index)
Translate an element key index (EG: quadrupole$, etc.) to a character string.

key_name_to_key_index (key_str, abbrev_allowed) result (key_index)
Function to convert a character string (eg: "drift") to an index (eg: drift$).

lat_ele_locator (loc_str, lat, eles, n_loc, err, above_ubound_is_err, ix_dflt_branch, order_by_index)
Routine to locate all the elements in a lattice that corresponds to loc_str.

lat_sanity_check (lat, err_flag)
Routine to check the control links in a lat structure, etc.

n_attrib_string_max_len () result (max_len)
Routine to return the maximum number of characters in any attribute name known to bmad.

num_lords (slave, lord_type) result (num)
Routine to return the number of lords of a given type for a given lattice element.

num_lords (slave, lord_type) result (num)
Routine to return the number of lords of a lattice element of a certain type.

pointer_to_indexed_attribute (ele, ix_attrib, do_allocation, a_ptr, err_flag, err_print_flag)
Returns a pointer to an attribute of an element ele with attribute index ix_attrib.

pointer_to_lord (slave, ix_lord, control, ix_slave_back, lord_type, ix_control, ix_ic) result (lord_ptr)
Function to point to a lord of a slave.

pointer_to_multipass_lord (ele, ix_pass, super_lord) result (multi_lord)
Routine to find the multipass lord of a lattice element. A multi_lord will be found for:

pointer_to_slave (lord, ix_slave, control, lord_type, ix_lord_back, ix_control, ix_ic) result (slave_ptr)
Function to point to a slave of a lord.

rf_is_on (branch, ix_ele1, ix_ele2) result (is_on)
Routine to check if any rfcavity is powered in a branch.

switch_attrib_value_name (attrib_name, attrib_value, ele, is_default, name_list) result (val_name)
Routine to return the name corresponding to the value of a given switch attribute.

type_ele (ele, type_zero_attrib, type_mat6, type_taylor, twiss_out, type_control, type_wake, type_floor_coords, type_field, type_wall, type_rad_kick, lines, n_lines)
Subroutine to print or put in a string array information on a lattice element.

type_twiss (ele, frequency_units, compact_format, lines, n_lines)
Subroutine to print or put in a string array Twiss information from an element.

valid_tracking_method (ele, species, tracking_method) result (is_valid)
Routine to return whether a given tracking method is valid for a given element.

valid_mat6_calc_method (ele, species, mat6_calc_method) result (is_valid)
Routine to return whether a given mat6_calc method is valid for a given element.

44.19 Lattice: Low Level Stuff

bracket_index (s, s_arr, i_min, dr, restrict)
Routine to find the index ix so that s(ix) $\leq$ s $<$ s(ix+1). If s $<$ s(1) then ix = 0

check_controller_controls (ele_key, contrl, name, err)
Routine to check for problems when setting up group or overlay controllers.

deallocate_ele_pointers (ele, nullify_only, nullify_branch, dealloc_poles)
Routine to deallocate the pointers in an element.

re_allocate_eles (eles, n, save_old, exact)
Routine to allocate an array of ele_pointer_structs.

twiss1_propagate (twiss1, mat2, ele_key, length, twiss2, err)
Routine to propagate the twiss parameters of a single mode.

44.20 Lattice: Manipulation

allocate_element_array (ele, upper_bound)
Routine to allocate or re-allocate an element array.

allocate_lat_ele_array (lat, upper_bound, ix_branch)
Routine to allocate or re-allocate an element array.

apply_all_rampers (lat, err_flag)
Routine to apply all the rampers in a lattice. Also see: apply_rampers_to_slave.

control_bookkeeper (lat, ele, err_flag)
Routine to calculate the combined strength of the attributes for controlled elements.

deallocate_ele_array_pointers (eles)
Routine to deallocate the pointers of all the elements in an element array and the array itself.

deallocate_lat_pointers (lat)
Routine to deallocate the pointers in a lat.

init_ele (ele, key, sub_key, ix_ele, branch)
Routine to initialize an element.

init_lat (lat, n, init_beginning_ele)
Routine to initialize a Bmad lat.

lattice_bookkeeper (lat, err_flag)
Routine to do bookkeeping for the entire lattice.

reallocate_coord (coord, n_coord)
Routine to reallocate an allocatable coord_struct array to at least: coord(0:n_coord).

reallocate_coord_array (coord_array, lat)
Routine to allocate an allocatable coord_array_struct array to the proper size for a lattice.

set_custom_attribute_name (custom_name, err_flag, custom_index)
Routine to add custom element attributes to the element attribute name table.

set_ele_defaults (ele, do_allocate)
Subroutine to set the defaults for an element of a given type.

set_on_off (key, lat, switch, orb, use_ref_orb, ix_branch, saved_values, attribute, set_val)
Routine to turn on or off a set of elements (quadrupoles, RF cavities, etc.) in a lat.

slice_lattice (lat, ele_list, error, do_bookkeeping)
Routine to discard from the lattice all elements not in ele_list. This is useful if simulations on just a section of a lattice is desired.

start_branch_at (lat, ele_start, move_end_marker, error)
Routine to, while preserving element order, rearrange the elements in a branch such that the first element is ele_start. In the finished lattice, the original first element will

taper_mag_strengths (lat, ref_lat, except, err_flag)
Routine to "correct" magnet strengths in a ring to counteract local radiation damping energy offsets.

transfer_ele (ele1, ele2, nullify_pointers)
Routine to set ele2 = ele1. This is a plain transfer of information not using the overloaded equal.

transfer_eles (ele1, ele2)
Routine to set ele2(:) = ele1(:). This is a plain transfer of information not using the overloaded equal.

transfer_ele_taylor (ele_in, ele_out, taylor_order)
Routine to transfer a Taylor map from one element to another.

transfer_lat (lat1, lat2)
Routine to set lat2 = lat1. This is a plain transfer of information not using the overloaded equal.

transfer_lat_parameters (lat_in, lat_out)
Routine to transfer the lat parameters (such as lat%name, lat%param, etc.) from one lat to another.

zero_ele_kicks (ele)
Subroutine to zero any kick attributes like hkick, $bl_v kick$, etc. See also: ele_has_kick, ele_has_-offset, zero_ele_offsets.

zero_ele_offsets (ele)
Routine to zero the offsets, pitches and tilt of an element.

44.21 Lattice: Miscellaneous

c_multi (n, m, no_n_fact, c_full)
Routine to compute multipole factors: c_multi(n, m) = +/- ("n choose m")/n!

ele_compute_ref_energy_and_time (ele0, ele, param, err_flag)
Routine to compute the reference energy and reference time at the end of an element given the reference enegy and reference time at the start of the element.

lat_compute_ref_energy_and_time (lat, err_flag)
Routine to compute the reference energy for each element in a lattice.

field_interpolate_3d (position, field_mesh, deltas, position0)
Function to interpolate a 3d field.

order_super_lord_slaves (lat, ix_lord)
Routine to make the slave elements of a super_lord in order.

release_rad_int_cache (ix_cache)
Routine to release the memory associated with caching wiggler values.

set_flags_for_changed_attribute (ele, attrib)
Routine to mark an element as modified.

44.22 Lattice: Nametable

create_lat_ele_nametable (lat, nametable)
Routine to create a sorted nametable of element names for a lattice.

ele_nametable_index(ele) result(ix_nt)
Routine to return the index in the nametable corresponding to ele. The reverse routine is: pointer_to_ele.

nametable_add (nametable, name, ix_name)
Routine to add a name to the nametable at index ix_name.

nametable_bracket_indexx (nametable, name, n_match) result (ix_max)
Routine to find the index ix_max such that: nametable%name(nametable%indexx(ix_max)) <= name < nametable%name(nametable%indexx(ix_max+1))

nametable_change1 (nametable, name, ix_name)
Routine to change one entry in a nametable.

nametable_init(nametable, n_min, n_max)
Routine to initialize a nametable_struct instance.

nametable_remove (nametable, ix_name)
Routine to remove a name from the nametable at index ix_name.

44.23 Lattice: Reading and Writing Files

bmad_parser (lat_file, lat, make_mats6, digested_read_ok, use_line, err_flag, parse_lat)
Routine to parse (read in) a Bmad input file.

bmad_parser2 (lat_file, lat, orbit, make_mats6, err_flag, parse_lat)
Routine to parse (read in) a Bmad input file to modify an existing lattice.

write_lattice_in_foreign_format (out_type, out_file_name, lat, ref_orbit, use_matrix_model, include_apertures, dr12_drift_max, ix_branch, err)
Routine to write a mad or sad lattice file.

combine_consecutive_elements (lat, error)
Routine to combine consecutive elements in the lattice that have the same name. This allows simplification, for example, of lattices where elements have been split to compute the beta function at the center.

create_sol_quad_model (sol_quad, lat)
Routine to create series of solenoid and quadrupole elements to serve as a replacement model for a sol_quad element.

create_unique_ele_names (lat, key, suffix)
Routine to give elements in a lattice unique names.

read_digested_bmad_file (digested_file, lat, inc_version, err_flag, parser_calling, lat_files)
Routine to read in a digested file.

write_bmad_lattice_file (bmad_file, lat, err, output_form, orbit0)
Routine to write a Bmad lattice file using the information in a lat_struct.

write_digested_bmad_file (digested_name, lat, n_files, file_names, extra, err_flag)
Routine to write a digested file.

44.24 Matrices

c_to_cbar (ele, cbar_mat)
Routine to compute Cbar from the C matrix and the Twiss parameters.

cbar_to_c (cbar_mat, a, b, c_mat)
Routine to compute C coupling matrix from the Cbar matrix and the Twiss parameters.

clear_lat_1turn_mats (lat)
Clear the 1-turn matrices in the lat structure.

concat_transfer_mat (mat_1, vec_1, mat_0, vec_0, mat_out, vec_out)
Routine to concatinate two linear maps

determinant (mat) result (det)
Routine to take the determinant of a square matrix This routine is adapted from Numerical Recipes.

do_mode_flip (ele, err_flag)
Routine to mode flip the Twiss parameters of an element

make_g2_mats (twiss, g2_mat, g2_inv_mat)
Routine to make the matrices needed to go from normal mode coords to coordinates with the beta function removed.

make_g_mats (ele, g_mat, g_inv_mat)
Routine to make the matrices needed to go from normal mode coords to coordinates with the beta function removed.

make_mat6 (ele, param, start_orb, end_orb, err_flag)
Routine to make the 6x6 transfer matrix for an element.

make_v_mats (ele, v_mat, v_inv_mat)
Routine to make the matrices needed to go from normal mode coords to X-Y coords and vice versa.

mat6_from_s_to_s (lat, mat6, vec0, s1, s2, ref_orb_in, ref_orb_out, ix_branch, one_turn, unit_start, err_flag, ele_save)
Subroutine to calculate the transfer map between longitudinal positions s1 to s2.

mat6_to_taylor (vec0, mat6, bmad_taylor, ref_orb)
Routine to form a first order Taylor map from the 6x6 transfer matrix and the 0th order transfer vector.

match_ele_to_mat6 (ele, start_orb, mat6, vec0, err_flag, include_delta_time, set_trombone)
Routine to make the 6 x 6 transfer matrix from the twiss parameters.

multi_turn_tracking_to_mat (track, n_var, map1, map0, track0, chi)
Routine to analyze 1-turn tracking data to find the 1-turn transfer matrix and the closed orbit offset.

transfer_matrix_calc (lat, xfer_mat, xfer_vec, ix1, ix2, ix_branch, one_turn)
Routine to calculate the transfer matrix between two elements. If ix1 and ix2 are not present the full 1–turn matrix is calculated.

one_turn_mat_at_ele (ele, phi_a, phi_b, mat4)
Routine to form the 4x4 1-turn coupled matrix with the reference point at the end of an element.

lat_make_mat6 (lat, ix_ele, ref_orb, ix_branch, err_flag)
Routine to make the 6x6 linear transfer matrix for an element

taylor_to_mat6 (a_taylor, r_in, vec0, mat6, r_out)
Routine to calculate the linear (Jacobian) matrix about some trajectory from a Taylor map.

transfer_mat2_from_twiss (twiss1, twiss2, mat)
Routine to make a 2 x 2 transfer matrix from the Twiss parameters at the end points.

transfer_mat_from_twiss (ele1, ele2, orb1, orb2, m)
Routine to make a 6 x 6 transfer matrix from the twiss parameters at the beginning and end of the element.

twiss_from_mat2 (mat_in, twiss, stat, type_out)
Routine to extract the Twiss parameters from the one-turn 2x2 matrix

twiss_from_mat6 (mat6, orb0, ele, stable, growth_rate, status, type_out)
Routine to extract the Twiss parameters from the one-turn 6x6 matrix

twiss_to_1_turn_mat (twiss, phi, mat2)
Routine to form the 2x2 1-turn transfer matrix from the Twiss parameters.

44.25 Matrix: Low Level Routines

Listed below are helper routines that are not meant for general use.

sol_quad_mat6_calc (ks_in, k1_in, tilt, length, ele, orbit, mat6, make_matrix)
Routine to calculate the transfer matrix for a combination solenoid/quadrupole element.

tilt_mat6 (mat6, tilt)
Routine to transform a 6x6 transfer matrix to a new reference frame that is tilted in (x, Px, y, Py) with respect to the old reference frame.

44.26 Measurement Simulation Routines

Routines to simulate errors in orbit, dispersion, betatron phase, and coupling measurements

to_eta_reading (eta_actual, ele, axis, add_noise, reading, err)
Compute the measured dispersion reading given the true dispersion and the monitor offsets, noise, etc.

to_orbit_reading (orb, ele, axis, add_noise, reading, err)
Calculate the measured reading on a bpm given the actual orbit and the BPM's offsets, noise, etc.

to_phase_and_coupling_reading (ele, add_noise, reading, err)
Find the measured coupling values given the actual ones

44.27 Multipass

multipass_all_info (lat, info)
Routine to put multipass to a multipass_all_info_struct structure.

multipass_chain (ele, ix_pass, n_links, chain_ele, use_super_lord)
Routine to return the chain of elements that represent the same physical element when there is multipass.

pointer_to_multipass_lord (ele, lat, ix_pass, super_lord) result (multi_lord)
Routine to find the multipass lord of a lattice element. A multi_lord will be found for:

44.28 Multipoles

ab_multipole_kick (a, b, n, ref_species, ele_orientation, coord, kx, ky, dk, pole_type, scale)
Routine to put in the kick due to an ab_multipole.

multipole_kicks (knl, tilt, ele, orbit, pole_type, ref_orb_offset)
Routine to put in the kick due to a multipole.

mexp (x, m) result (this_exp)
Returns x**m with 0**0 = 0.

multipole_ab_to_kt (an, bn, knl, tn)
Routine to convert ab type multipoles to kt (MAD standard) multipoles.

multipole_ele_to_ab (ele, use_ele_tilt, ix_pole_max, a, b, pole_type, include_kicks, b1)
Routine to put the scaled element multipole components (normal and skew) into 2 vectors.

multipole_ele_to_kt (ele, use_ele_tilt, ix_pole_max, knl, tilt, pole_type, include_kicks)
Routine to put the scaled element multipole components (strength and tilt) into 2 vectors.

multipole_init(ele, who, zero)
Routine to initialize the multipole arrays within an element.

multipole_kick (knl, tilt, n, ref_species, ele_orientation, coord, pole_type, ref_orb_offset)
Routine to put in the kick due to a multipole.

multipole_kt_to_ab (knl, knsl, tn, an, bn)
Routine to convert kt (MAD standard) multipoles to ab type multipoles.

44.29 Nonlinear Optimizers

opti_lmdif (vec, n, merit, eps) result(this_opti)
Function which tries to get the merit function(s) as close to zero as possible by changing the values in vec. Multiple merit functions can be used.

initial_lmdif()
Routine that clears out previous saved values of the optimizer.

suggest_lmdif (XV, FV, EPS, ITERMX, at_end, reset_flag)
Reverse communication routine.

super_mrqmin (y, weight, a, chisq, funcs, storage, alamda, status, maska, print_err)
Routine to do non-linear optimizations. This routine is essentially mrqmin from Numerical Recipes with some added features.

opti_de (v_best, generations, population, merit_func, v_del, status)
Differential Evolution for Optimal Control Problems. This optimizer is based upon the work of Storn and Price.

44.30 Overloading the equal sign

These routines are overloaded by the equal sign so should not be called explicitly.

branch_equal_branch (branch1, branch2)
Routine that is used to set one branch equal to another.

bunch_equal_bunch (bunch1, bunch2)
Routine that is used to set one macroparticle bunch to another. This routine takes care of the pointers in bunch1.

coord_equal_coord (coord1, coord2)
Routine that is used to set one coord_struct equal to another.

ele_equal_ele (ele_out, ele_in)
Routine that is used to set one element equal to another. This routine takes care of the pointers in ele1.

lat_equal_lat (lat_out, lat_in)
Routine that is used to set one lat equal to another. This routine takes care of the pointers in lat1.

lat_vec_equal_lat_vec (lat1, lat2)
Routine that is used to set one lat array equal to another. This routine takes care of the pointers in lat1(:).

universal_equal_universal (ut1, ut2)
Routine that is used to set one PTC universal_taylor structure equal to another.

44.31 Particle Coordinate Stuff

angle_to_canonical_coords (orbit, coord_type)
Routine to convert from angle (x, x', y, y', z, z') coordinates to canonical (x, px, y, py, z, pz) coordinates.

convert_coords (in_type_str, coord_in, ele, out_type_str, coord_out, err_flag)
Routine to convert between lab frame, normal mode, normalized normal mode, and action-angle coordinates.

convert_pc_to (pc, particle, E_tot, gamma, kinetic, beta, brho, beta1, err_flag)
Routine to calculate the energy, etc. from a particle's momentum.

convert_total_energy_to (E_tot, particle, gamma, kinetic, beta, pc, brho, beta1, err_flag, print_err)
Routine to calculate the momentum, etc. from a particle's total energy.

init_coord (orb, vec, ele, element_end, particle, direction, E_photon, t_offset, shift_vec6, spin)
Routine to initialize a coord_struct.

type_coord (coord)
Routine to type out a coordinate.

44.32 Photon Routines

bend_photon_init (g_bend_x, g_bend_y, gamma, orbit, E_min, E_max, E_integ_prob, vert_angle_min, vert_angle_max, vert_angle_symmetric, emit_probability)
Routine to initialize a photon generated by a charged particle in a bend.

bend_photon_vert_angle_init (E_rel, gamma, r_in, invert) result (r_in)
Routine to convert a "random" number in the interval [0,1] to a photon vertical emission angle for a simple bend.

44.33 PTC Interface Routines

concat_real_8 (y1, y2, y3, r2_ref, keep_y1_const_terms)
Routine to concatenate two real_8 taylor series.

ele_to_fibre (ele, ptc_fibre, use_offsets, err_flag, integ_order, steps, for_layout, ref_in)
Routine to convert a Bmad element to a PTC fibre element.

kill_ptc_layouts (lat)
Routine to kill the layouts associated with a Bmad lattice.

kind_name (this_kind)
Function to return the name of a PTC kind.

lat_to_ptc_layout (lat)
Routine to create a PTC layout from a Bmad lat.

map_coef (y, i, j, k, l)
Function to return the coefficient of the map y(:) up to 3rd order.

normal_form_rd_terms(one_turn_map, normal_form, rf_on, order)
Calculates driving terms à la [Bengt97] from the one-turn map.

ptc_transfer_map_with_spin (branch, t_map, s_map, orb0, err_flag, ix1, ix2, one_turn, unit_start)
Subroutine to calculate the transfer map between two elements. To calculate just the first order transfer matrices see the routine:

remove_constant_taylor (taylor_in, taylor_out, c0, remove_higher_order_terms)
Routine to remove the constant part of a taylor series.

set_ptc (e_tot, particle, taylor_order, integ_order, n_step, no_cavity, force_init)
Routine to initialize PTC.

sort_universal_terms (ut_in, ut_sorted)
Routine to sort the taylor terms from "lowest" to "highest".

taylor_to_genfield (bmad_taylor, ptc_genfield, c0)
Routine to construct a genfield (partially inverted map) from a taylor map.

taylor_to_real_8 (bmad_taylor, beta0, beta1, ptc_re8, ref_orb_ptc, exi_orb_ptc)
Routine to convert from a taylor map in Bmad to a real_8 taylor map in Étienne's PTC.

type_layout (lay)
Routine to print the global information in a PTC layout.

type_map (y)
Routine to type the transfer maps of a real_8 array.

type_map1 (y, type0, n_dim)
Routine to type the transfer map up to first order.

type_fibre (ptc_fibre, print_coords, lines, n_lines)
Routine to print the global information in a fibre.

type_real_8_taylors (y)
Routine to type out the taylor series from a real_8 array.

universal_to_bmad_taylor (u_taylor, bmad_taylor)
Routine to convert from a universal_taylor map in Étienne's PTC to a taylor map in Bmad.

44.34 Quick Plot Routines

44.34.1 Quick Plot Page Routines

qp_open_page (page_type, i_chan, x_len, y_len, units, plot_file, scale)
Routine to Initialize a page (window) for plotting.

qp_select_page (iw)
Routine to switch to a particular page for drawing graphics.

qp_close_page()
Routine to finish plotting on a page.

44.34.2 Quick Plot Calculational Routines

qp_axis_niceness (imin, imax, divisions) result (score)
Routine to calculate how "nicely" an axis will look. The higher the score the nicer.

qp_calc_and_set_axis (axis_str, data_min, data_max, div_min, div_max, bounds, axis_type, slop_factor)
Routine to calculate a "nice" plot scale given the minimum and maximum of the data.

qp_calc_axis_params (data_min, data_max, div_min, div_max, axis)
Routine to calculate a "nice" plot scale given the minimum and maximum of the data. This is similar to calc_axis_scale.

qp_calc_axis_places (axis)
Routine to calculate the number of decimal places needed to display the axis numbers.

qp_calc_axis_scale (data_min, data_max, axis, niceness_score)
Routine to calculate a "nice" plot scale given the minimum and maximum of the data.

qp_calc_minor_div (delta, div_max, divisions)
Routine to calculate the number of minor divisions an axis should have.

qp_convert_rectangle_rel (rect1, rect2)
Routine to convert a "rectangle" (structure of 4 points) from one set of relative units to another

44.34.3 Quick Plot Drawing Routines

qp_clear_box()
Routine to clear the current box on the page.

qp_clear_page()
Routine to clear all drawing from the page.

qp_draw_circle (x0, y0, r, angle0, del_angle,
units, width, color, line_pattern, clip)
Routine to plot a section of an ellipse.

qp_draw_ellipse (x0, y0, r_x, r_y, theta_xy,
angle0, del_angle, units, width, color, line_pattern, clip)
Routine to plot a section of an ellipse.

qp_draw_axes(x_lab, y_lab, title, draw_grid)
Routine to plot the axes, title, etc. of a plot.

qp_draw_data (x_dat, y_dat, draw_line, symbol_every, clip)
Routine to plot data, axes with labels, a grid, and a title.

qp_draw_graph (x_dat, y_dat, x_lab, y_lab, title,
draw_line, symbol_every, clip)
Routine to plot data, axes with labels, a grid, and a title.

qp_draw_graph_title (title)
Routine to draw the title for a graph.

qp_draw_grid()
Routine to draw a grid on the current graph.

qp_draw_histogram (x_dat, y_dat, fill_color, fill_pattern, line_color, clip)
Routine to plot data, axes with labels, a grid, and a title.

qp_draw_curve_legend (origin, legend, line, symbol, text)
Routine to draw a legend with each line in the legend having a line, a symbol, some text.

qp_draw_text_legend (text, x_origin, y_origin, units)
Routine to draw a legend of lines of text.

qp_draw_main_title (lines, justify)
Routine to plot the main title at the top of the page.

qp_draw_polyline (x, y, units, width, color, line_pattern, clip, style)
Routine to draw a polyline.

qp_draw_polyline_no_set (x, y, units)
Routine to draw a polyline. This is similar to qp_draw_polyline except qp_set_line_attrib is not called.

qp_draw_polyline_basic (x, y)
Routine to draw a polyline. See also qp_draw_polyline

qp_draw_line (x1, x2, y1, y2, units, width, color, line_pattern, clip, style)
Routine to draw a line.

qp_draw_rectangle (x1, x2, y1, y2, units, color, width, line_pattern, clip, style)
Routine to draw a rectangular box.

qp_draw_symbol (x, y, units, type, height, color, fill_pattern, line_width, clip)
Draws a symbol at (x, y)

qp_draw_symbols (x, y, units, type, height, color,
fill_pattern, line_width, clip, symbol_every)
Draws a symbol at the (x, y) points.

qp_draw_text (text, x, y, units, justify, height, color, angle,
background, uniform_spacing, spacing_factor)
Routine to draw text.

qp_draw_text_no_set (text, x, y, units, justify, angle)
Routine to display on a plot a character string. See also: qp_draw_text.

qp_draw_text_basic (text, len_text, x0, y0, angle, justify)
Routine to display on a plot a character string. See also: qp_draw_text.

qp_draw_x_axis (who, y_pos)
Routine to draw a horizontal axis.

qp_draw_y_axis (who, x_pos)
Routine to draw a horizontal axis.

qp_paint_rectangle (x1, x2, y1, y2, units, color, fill_pattern)
Routine to paint a rectangular region a specified color. The default color is the background color (white$).

qp_to_axis_number_text (axis, ix_n, text)
Routine to form the text string for an axis number.

44.34.4 Quick Plot Set Routines

qp_calc_and_set_axis (axis, data_min, data_max,
div_min, div_max, bounds, axis_type, slop_factor)
Routine to calculate a "nice" plot scale given the minimum and maximum of the data.

qp_eliminate_xy_distortion(axis_to_scale)
This routine will increase the x or y margins so that the conversion between data units and page units is the same for the x and y axes.

qp_set_axis (axis_str, a_min, a_max, div, places, label, draw_label,
draw_numbers, minor_div, minor_div_max, mirror, number_offset,
label_offset, major_tick_len, minor_tick_len, ax_type,
tick_min, tick_max, dtick, set_ticks, axis)
Routine to set (but not plot) the min, max and divisions for the axes of the graph.

qp_set_box (ix, iy, ix_tot, iy_tot)
Routine to set the box on the physical page. This routine divides the page into a grid of boxes.

qp_set_graph (title)
Routine to set certain graph attributes.

qp_set_graph_limits()
Routine to calculate the offsets for the graph. This routine also sets the PGPLOT window size equal to the graph size.

qp_set_graph_placement (x1_marg, x_graph_len, y1_marg, y_graph_len, units)
Routine to set the placement of the current graph inside the box. This routine can be used in place of qp_set_margin.

qp_set_layout (x_axis, y_axis, x2_axis, y2_axis,
x2_mirrors_x, y2_mirrors_y, box, margin, page_border)
Routine to set various attributes. This routine can be used in place of other qp_set_* routines.

qp_set_line (who, line)
Routine to set the default line attributes.

qp_set_margin (x1_marg, x2_marg, y1_marg, y2_marg, units)
Routine to set up the margins from the sides of the box (see QP_SET_BOX) to the edges of the actual graph.

qp_set_page_border (x1_b, x2_b, y1_b, y2_b, units)
Routine to set the border around the physical page.

qp_set_page_border_to_box ()
Routine to set the page border to correspond to the region of the current box. This allows qp_set_box to subdivide the current box.

qp_set_clip (clip)
Routine to set the default clipping state.

qp_set_parameters (text_scale, default_draw_units, default_set_units,
default_axis_slop_factor)
Routine to set various quick plot parameters.

qp_subset_box (ix, iy, ix_tot, iy_tot, x_marg, y_marg)
Routine to set the box for a graph. This is the same as qp_set_box but the boundaries of the page are taken to be the box boundaries.

qp_set_symbol (symbol)
Routine to set the type and size of the symbols used in plotting data. See the pgplot documentation for more details.

qp_set_symbol_attrib (type, height, color, fill_pattern, line_width, clip)
Routine to set the type and size of the symbols used in plotting data.

qp_set_line_attrib (style, width, color, pattern, clip)
Routine to set the default line attributes.

qp_set_graph_attrib (draw_grid, draw_title)
Routine to set attributes of the current graph.

qp_set_text_attrib (who, height, color,
background, uniform_spacing, spacing_factor)
Routine to set the default text attributes.

qp_use_axis (x, y)
Routine to set what axis to use: X or X2, Y or Y2.

44.34.5 Informational Routines

qp_get_axis_attrib (axis_str, a_min, a_max, div, places, label, draw_label, draw_numbers, minor_div, mirror, number_offset, label_offset, major_tick_len, minor_tick_len, ax_type, tick_min, tick_max, dtick)
Routine to get the min, max, divisions etc. for the X and Y axes.

qp_get_layout_attrib (who, x1, x2, y1, y2, units)
Routine to get the attributes of the layout.

qp_get_line (style, line)
Routine to get the default line attributes.

qp_get_parameters (text_scale, default_draw_units, default_set_units, default_axis_slop_factor)
Routine to get various quick_plot parameters.

qp_get_symbol (symbol)
Routine to get the symbol parameters used in plotting data. Use qp_set_symbol or qp_set_symbol_attrib to set symbol attributes.

qp_text_len (text)
Function to find the length of a text string.

44.34.6 Conversion Routines

qp_from_inch_rel (x_inch, y_inch, x, y, units)
Routine to convert from a relative position (an offset) in inches to other units.

qp_from_inch_abs (x_inch, y_inch, x, y, units)
Routine to convert to absolute position (x, y) from inches referenced to the Left Bottom corner of the page

qp_text_height_to_inches(height_pt) result (height_inch)
Function to convert from a text height in points to a text height in inches taking into account the text_scale.

qp_to_inch_rel (x, y, x_inch, y_inch, units)
Routine to convert a relative (x, y) into inches.

qp_to_inch_abs (x, y, x_inch, y_inch, units)
Routine to convert an absolute position (x, y) into inches referenced to the Left Bottom corner of the page.

qp_to_inches_rel (x, y, x_inch, y_inch, units)
Routine to convert a relative (x, y) into inches.

qp_to_inches_abs (x, y, x_inch, y_inch, units)
Routine to convert an absolute position (x, y) into inches referenced to the left bottom corner of the page.

44.34.7 Miscellaneous Routines

qp_read_data (iu, err_flag, x, ix_col, y, iy_col, z, iz_col, t, it_col)
Routine to read columns of data.

44.34.8 Low Level Routines

qp_clear_box_basic (x1, x2, y1, y2)
Routine to clear all drawing from a box. That is, white out the box region.

qp_clear_page_basic()
Routine to clear all drawing from the page.

qp_close_page_basic()
Routine to finish plotting on a page. For X this closes the window.

qp_convert_point_rel (x_in, y_in, units_in, x_out, y_out, units_out)
Routine to convert a (x, y) point from from one set of relative units to another.

qp_convert_point_abs (x_in, y_in, units_in, x_out, y_out, units_out)
Routine to convert a (x, y) point from from one set of absolute units to another.

qp_draw_symbol_basic (x, y, symbol)
Routine to draw a symbol.

qp_init_com_struct ()
Routine to initialize the common block qp_state_struct. This routine is not for general use.

qp_join_units_string (u_type, region, corner, units)
Routine to form a units from its components.

qp_justify (justify)
Function to convert a justify character string to a real value representing the horizontal justification.

qp_open_page_basic (page_type, x_len, y_len, plot_file,
x_page, y_page, i_chan, page_scale)
Routine to Initialize a page (window) for plotting.

qp_paint_rectangle_basic (x1, x2, y1, y2, color, fill_pattern)
Routine to fill a rectangle with a given color. A color of white essentially erases the rectangle.

qp_pointer_to_axis (axis_str, axis_ptr)
Routine to return a pointer to an common block axis.

qp_restore_state()
Routine to restore saved attributes. Use qp_save_state to restore the saved state.

qp_restore_state_basic (buffer_basic)
Routine to restore the print state.

qp_save_state (buffer_basic)
Routine to save the current attributes. Use qp_restore_state to restore the saved state.

qp_save_state_basic ()
Routine to save the print state.

qp_select_page_basic (iw)
Routine to switch to a particular page for drawing graphics.

qp_set_char_size_basic (height)
Routine to set the character size.

qp_set_clip_basic (clip)
Routine to set the clipping state. Note: This affects both lines and symbols.

qp_set_color_basic (ix_color)
Routine to set the color taking into account that GIF inverts the black for white.

qp_set_graph_position_basic (x1, x2, y1, y2)
Routine to set the position of a graph. Units are inches from lower left of page.

qp_set_line_width_basic (line_width)
Routine to set the line width.

qp_set_symbol_fill_basic (fill)
Routine to set the symbol fill style.

qp_set_symbol_size_basic (height, symbol_type, uniform_size)
Routine to set the symbol_size

qp_set_text_background_color_basic (color)
Routine to set the character text background color.

qp_split_units_string (u_type, region, corner, units)
Routine to split a units string into its components.

qp_text_len_basic (text)
Function to find the length of a text string.

44.35 Spin

calc_spin_params (bunch, bunch_params)
Rotine to calculate spin averages

spinor_to_polar (spinor) result (polar)
Routine to convert a spinor into polar coordinates.

polar_to_vec (polar) result (vec)
Routine to convert a spin vector from polar coordinates to Cartesian coordinates.

polar_to_spinor (polar) result (coord)
Routine to convert a spin vector in polar coordinates to a spinor.

vec_to_polar (vec, phase) result (polar)
Routine to convert a spin vector from Cartesian coordinates to polar coordinates preserving the complex phase.

spinor_to_vec (spinor) result (vec)
Routine to convert a spinor to a spin vector in Cartesian coordinates.

vec_to_spinor (vec, phase) result (coord)
Routine to convert a spin vector in Cartesian coordinates to a spinor using the specified complex phase.

angle_between_polars (polar1, polar2)
Function to return the angle between two spin vectors in polar coordinates.

spin_omega (field, orbit, sign_z_vel, phase_space_coords), result (omega)
Return the modified T-BMT spin omega vector.

track1_spin (start_orb, ele, param, end_orb, make_quaternion)
Routine to track the particle spin through one element.

44.36 Transfer Maps: Routines Called by make_mat6

`Make_mat6` is the routine for calculating the transfer matrix (Jacobin) through an element. The routines listed below are used by `make_mat6`. In general a program should call `make_mat6` rather than using these routines directly.

make_mat6_bmad (ele, param, start_orb, end_orb, err)
Routine to make the 6x6 transfer matrix for an element using closed formulas.

make_mat6_custom (ele, param, c0, c1, err_flag)
Routine for custom calculations of the 6x6 transfer matrices.

make_mat6_symp_lie_ptc (ele, start_orb, end_orb)
Routine to make the 6x6 transfer matrix for an element using the PTC symplectic integrator.

make_mat6_taylor (ele, start_orb, end_orb, err_flag)
Routine to make the 6x6 transfer matrix for an element from a Taylor map.

make_mat6_tracking (ele, param, start_orb, end_orb, err_flag, spin_only)
Routine to make the 6x6 transfer matrix for an element by tracking 7 particle with different starting conditions.

44.37 Transfer Maps: Complex Taylor Maps

complex_taylor_coef (bmad_taylor, expn)
Function complex_taylor_coef (bmad_complex_taylor, i1, i2, i3, i4, i5, i6, i7, i8, i9) Function to return the coefficient for a particular complex_taylor term from a complex_taylor Series.

complex_taylor_equal_complex_taylor (complex_taylor1, complex_taylor2)
Subroutine that is used to set one complex_taylor equal to another. This routine takes care of the pointers in complex_taylor1.

complex_taylor_make_unit (complex_taylor)
Subroutine to make the unit complex_taylor map: r(out) = Map * r(in) = r(in)

complex_taylor_exponent_index(expn) result(index)
Function to associate a unique number with a complex_taylor exponent.

complex_taylor_to_mat6 (a_complex_taylor, r_in, vec0, mat6, r_out)
Subroutine to calculate, from a complex_taylor map and about some trajectory: The 1st order (Jacobian) transfer matrix.

complex_taylors_equal_complex_taylors (complex_taylor1, complex_taylor2)
Subroutine to transfer the values from one complex_taylor map to another: complex_taylor1 <= complex_taylor2

init_complex_taylor_series (complex_taylor, n_term, save)
Subroutine to initialize a Bmad complex_taylor series (6 of these series make a complex_taylor map). Note: This routine does not zero the structure. The calling

kill_complex_taylor (complex_taylor)
Subroutine to deallocate a Bmad complex_taylor map.

mat6_to_complex_taylor (vec0, mat6, complex_taylor)
Subroutine to form a first order complex_taylor map from the 6x6 transfer matrix and the 0th order transfer vector.

sort_complex_taylor_terms (complex_taylor_in, complex_taylor_sorted)
Subroutine to sort the complex_taylor terms from "lowest" to "highest" of a complex_taylor series.

track_complex_taylor (start_orb, complex_taylor, end_orb)
Subroutine to track using a complex_taylor map.

truncate_complex_taylor_to_order (complex_taylor_in, order, complex_taylor_out)
Subroutine to throw out all terms in a complex_taylor map that are above a certain order.

type_complex_taylors (complex_taylor, max_order, lines, n_lines, file_id, out_type, clean)
Subroutine to output a Bmad complex_taylor map.

44.38 Transfer Maps: Taylor Maps

add_taylor_term (bmad_taylor, coef, expn, replace)
add_taylor_term (bmad_taylor, coef, i1, i2, i3, i4, i5, i6, i7, i8, i9, replace)
Overloaded routine to add a Taylor term to a Taylor series.

concat_ele_taylor (taylor1, ele, taylor3, err_flag)
Routine to concatenate two taylor maps.

concat_taylor (taylor1, taylor2, taylor3)
Routine to concatenate two taylor series: taylor3(x) = taylor2(taylor1(x))

ele_to_taylor (ele, orb0, taylor_map_includes_offsets, include_damping, orbital_taylor, spin_taylor)
Routine to make a Taylor map for an element. The order of the map is set by set_ptc.

equivalent_taylor_attributes (ele1, ele2) result (equiv)
Routine to see if to elements are equivalent in terms of attributes so that their Taylor Maps would be the same.

init_taylor_series (bmad_taylor, n_term, save_old)
Routine to initialize a Bmad Taylor series.

kill_taylor (bmad_taylor)
Routine to deallocate a Bmad Taylor map.

mat6_to_taylor (mat6, vec0, bmad_taylor)
Routine to form a first order Taylor map from the 6x6 transfer matrix and the 0th order transfer vector.

sort_taylor_terms (taylor_in, taylor_sorted, min_val)
Routine to sort the taylor terms from "lowest" to "highest" of a Taylor series.

taylor_coef (bmad_taylor, expn)
Function to return the coefficient for a particular taylor term from a Taylor Series.

taylor_equal_taylor (taylor1, taylor2)
Routine to transfer the values from one taylor map to another: Taylor1 $\leq$ Taylor2

transfer_map_calc (lat, t_map, err_flag, ix1, ix2, ref_orb, ix_branch, one_turn, unit_start, concat_if_possible)
Routine to calculate the transfer map between two elements.

transfer_map_from_s_to_s (lat, t_map, s1, s2, ref_orb_in, ref_orb_out, ix_branch, one_turn, unit_start, err_flag, concat_if_possible)
Subroutine to calculate the transfer map between longitudinal positions s1 to s2.

taylor_minus_taylor (taylor1, taylor2) result (taylor3)
Routine to add two taylor maps.

taylor_plus_taylor (taylor1, taylor2) result (taylor3)
Routine to add two taylor maps.

taylors_equal_taylors (taylor1, taylor2)
Routine to transfer the values from one taylor map to another.

taylor_make_unit (bmad_taylor, ref_orbit)
Routine to make the unit Taylor map

taylor_to_mat6 (a_taylor, c0, mat6, c1)
Routine to calculate the linear (Jacobian) matrix about some trajectory from a Taylor map.

taylor_inverse (taylor_in, taylor_inv, err)
Routine to invert a taylor map.

taylor_propagate1 (bmad_taylor, ele, param, err_flag, ref_in)
Routine to track a real_8 taylor map through an element. The alternative routine, if ele has a taylor series, is concat_taylor.

track_taylor (start_orb, bmad_taylor, ref_orb)
Routine to track using a Taylor map.

transfer_ele_taylor (ele_in, ele_out, taylor_order)
Routine to transfer a Taylor map from one element to another.

truncate_taylor_to_order (taylor_in, order, taylor_out)
Routine to throw out all terms in a taylor map that are above a certain order.

type_taylors (bmad_taylor, max_order, lines, n_lines, file_id, out_style, clean)
Routine to output a Bmad taylor map.

44.39 Transfer Maps: Driving Terms

srdt_calc (lat, srdt_sums, order, n_slices_gen_opt, n_slices_sxt_opt, per_ele_out)
Calculates driving terms using summations over sextupole moments like those in [Bengt97] and [Wang12]. Often called resonance driving terms (RDTs), though strictly speaking not all terms drive resonances. Terms that are first and second order in sextupole moment are included. See srdt_mod for a list of available driving terms.

make_srdt_cache(lat, n_slices_gen, n_slices_sxt, eles, cache)
Used to speed up calculation of the 2nd order driving terms. Makes an $N_{k2} \times N_{k2} \times 11$ array that contains the precomputed cross-products of the linear optics at each sextupole moment.

srdt_calc_with_cache(lat, srdt_sums, order, n_slices_gen_opt, n_slices_sxt_opt, cache, per_ele_out)
Same as 44.39, but makes use of a cache precomputed by 44.39 to speed up the calculations.

srdt_lsq_solution(lat, var_indexes, ls_soln, n_slices_gen_opt, n_slices_sxt_opt, chrom_set_x_opt, chrom_set_y_opt, weight_in)
Given a lattice and the indices of the sextupole moments to use as variables, finds the least squares solution to the sextupole moments that minimize the 1st order driving terms. If there are more variables than driving terms, then the solution sets the driving terms to zero and minimizes the sum of the squares of the variables. If there are fewer variables than driving terms, then the solution is that which minimizes the sum of the squares of the driving terms.

44.40 Tracking and Closed Orbit

The following routines perform tracking and closed orbit calculations.

check_aperture_limit (orb, ele, particle_at, param, old_orb, check_momentum)
Routine to check if an orbit is outside an element's aperture.

check_aperture_limit_custom (orb, ele, particle_at, param, err_flag)
Routine to check if an orbit is outside an element's aperture. Used when `ele%aperture_type` is set to `custom$`

closed_orbit_calc (lat, closed_orb, i_dim, direction, ix_branch, err_flag, print_err)
Routine to calculate the closed orbit at the beginning of the lat.

closed_orbit_from_tracking (lat, closed_orb, i_dim, eps_rel, eps_abs, init_guess, err_flag)
Routine to find the closed orbit via tracking.

multi_turn_tracking_analysis (track, i_dim, track0, ele, stable, growth_rate, chi, err_flag)
Routine to analyze multi-turn tracking data to get the Twiss parameters etc.

multi_turn_tracking_to_mat (track, i_dim, mat1, track0, chi)
Routine to analyze 1-turn tracking data to find the 1-turn transfer matrix and the closed orbit offset.

offset_particle (ele, set, orbit, set_tilt, set_hvkicks, drift_to_edge, s_pos, s_out, set_spin, mat6, make_matrix, spin_qrot, time)
Routine to effectively offset an element by instead offsetting the particle position to correspond to the local element coordinates.

offset_photon (ele, orbit, set, offset_position_only, rot_mat)
Routine to effectively offset an element by instead offsetting the photon position to correspond to the local crystal or mirror coordinates.

orbit_amplitude_calc (ele, orb, amp_a, amp_b, amp_na, amp_nb)
Routine to calculate the "invariant" amplitude of a particle at a particular point in its orbit.

particle_is_moving_backwards (orbit) result (is_moving_backward)
Routine to determine if a particle is moving in the backward -s direction. If not moving backward it is dead or is moving backward.

particle_is_moving_forward (orbit, dir) result (is_moving_forward)
Routine to determine if a particle is moving in the forward +s direction. If not moving forward it is dead or is moving backward.

tilt_coords (tilt_val, coord, mat6, make_matrix)
Routine to effectively tilt (rotate in the x-y plane) an element by instead rotating the particle position with negative the angle.

track1 (start_orb, ele, param, end_orb, track, err_flag, ignore_radiation, make_map1, init_to_edge)
Routine to track through a single element.

track1_bunch_csr (bunch, ele, centroid, err, s_start, s_end, bunch_track)
Routine to track a bunch of particles through the element lat%ele(ix_ele) with csr radiation effects.

track1_spin_custom (start, ele, param, end, err_flag, track, make_quaternion)
Dummy routine for custom spin tracking. This routine needs to be replaced for a custom calculation.

track_all (lat, orbit, ix_branch, track_state, err_flag, orbit0)
Routine to track through the lat.

track_from_s_to_s (lat, s_start, s_end, orbit_start, orbit_end, all_orb, ix_branch, track_state, ix_ele_end)
Routine to track a particle between two s-positions.

track_many (lat, orbit, ix_start, ix_end, direction, ix_branch, track_state)
Routine to track from one element in the lat to another.

twiss_and_track (lat, orb, ok)
twiss_and_track (lat, orb_array, ok)
Routine to calculate the twiss parameters, transport matrices and orbit.

twiss_and_track_at_s (lat, s, ele_at_s, orb, orb_at_s, ix_branch, err, use_last, compute_floor_coords)
Routine to calculate the Twiss parameters and orbit at a particular longitudinal position.

twiss_and_track_from_s_to_s (branch, orbit_start, s_end, orbit_end, ele_start, ele_end, err, compute_floor_coords, compute_twiss)
Routine to track a particle from one location to another.

twiss_and_track_intra_ele (ele, param, l_start, l_end, track_upstream_end, track_downstream_end, orbit_start, orbit_end, ele_start, ele_end, err, compute_floor_coords, compute_twiss, reuse_ele_end)
Routine to track a particle within an element.

twiss_from_tracking (lat, ref_orb0, symp_err, err_flag, d_orb)
Routine to compute from tracking the Twiss parameters and the transfer matrices for every element in the lat.

wall_hit_handler_custom (orb, ele, s)
This routine is called by the Runge-Kutta integrator odeint_bmad when a particle hits a wall.

44.41 Tracking: Low Level Routines

absolute_time_tracking (ele) result (is_abs_time)
Routine to return a logical indicating whether the tracking through an element should use absolute time or time relative to the reference particle.

odeint_bmad (orbit, ele, param, s1_body, s2_body, err_flag, track, mat6, make_matrix)
Routine to do Runge Kutta tracking.

track_a_drift (orb, length, mat6, make_matrix, ele_orientation, include_ref_motion, time)
Routine to track through a drift.

track_a_bend (orbit, ele, param, mat6, make_matrix)
Particle tracking through a bend element.

44.42 Tracking: Mad Routines

make_mat6_mad (ele, param, c0, c1)
Routine to make the 6x6 transfer matrix for an element from the 2nd order MAD transport map. The map is stored in ele%taylor.

make_mad_map (ele, param, energy, map)
Routine to make a 2nd order transport map a la MAD.

mad_add_offsets_and_multipoles (ele, map)
Routine to add in the effect of element offsets and/or multipoles on the 2nd order transport map for the element.

mad_drift (ele, energy, map)
Routine to make a transport map for a drift space. The equivalent MAD-8 routine is: TMDRF

mad_elsep (ele, energy, map)
Routine to make a transport map for an electric separator. The equivalent MAD-8 routine is: TMSEP

mad_sextupole (ele, energy, map)
Routine to make a transport map for an sextupole. The equivalent MAD-8 routine is: TMSEXT

mad_sbend (ele, energy, map)
Routine to make a transport map for a sector bend element. The equivalent MAD-8 routine is: TMBEND

mad_sbend_fringe (ele, energy, into, map)
Routine to make a transport map for the fringe field of a dipole. The equivalent MAD-8 routine is: TMFRNG

mad_sbend_body (ele, energy, map)
Routine to make a transport map for the body of a sector dipole. The equivalent MAD-8 routine is: TMSECT

mad_tmfoc (el, sk1, c, s, d, f)
Routine to compute the linear focusing functions. The equivalent MAD-8 routine is: TMFOC

mad_quadrupole (ele, energy, map)
Routine to make a transport map for an quadrupole element. The equivalent MAD-8 routine is: TMSEXT

mad_rfcavity (ele, energy, map)
Routine to make a transport map for an rfcavity element. The equivalent MAD-8 routine is: TMRF

mad_solenoid (ele, energy, map)
Routine to make a transport map for an solenoid. The equivalent MAD-8 routine is: TMSEXT

mad_tmsymm (te)
routine to symmetrize the 2nd order map t. The equivalent MAD-8 routine is: tmsymm

mad_tmtilt (map, tilt)
Routine to apply a tilt to a transport map. The equivalent MAD-8 routine is: TMTILT

mad_concat_map2 (map1, map2, map3)
Routine to concatenate two 2nd order transport maps.

mad_track1 (c0, map, c1)
Routine to track through a 2nd order transfer map. The equivalent MAD-8 routine is: TMTRAK

track1_mad (orbit, ele, param)
Routine to track through an element using a 2nd order transfer map. Note: If map does not exist then one will be created.

mad_map_to_taylor (map, energy, taylor)
Routine to convert a mad order 2 map to a taylor map.

taylor_to_mad_map (taylor, energy, map)
Routine to convert a Taylor map to a mad order 2 map. If any of the Taylor terms have order greater than 2 they are ignored.

make_unit_mad_map (map)
Routine to initialize a 2nd order transport map to unity.

44.43 Tracking: Routines called by track1

Note: Unless you know what you are doing do not call these routines directly. Rather use track1.

symp_lie_bmad (ele, param, orbit, track, mat6, make_matrix, offset_ele)
Symplectic integration through an element to 0th or 1st order.

track1_bmad (orbit, ele, param, err_flag, track, mat6, make_matrix)
Particle tracking through a single element BMAD_standard style.

track1_custom (orbit, ele, param, err_flag, finished, track)
Dummy routine for custom tracking.

track1_linear (orbit, ele, param)
Particle tracking through a single element using the transfer matrix..

track1_postprocess (start_orb, ele, param, end_orb)
Dummy routine for post processing after the track1 routine is done.

track1_preprocess (start_orb, ele, param, err_flag, finished, radiation_included, track)
Dummy routine for pre processing at the start of the track1 routine.

track1_radiation (orbit, ele, edge)
Routine to put in radiation damping and/or fluctuations.

track1_runge_kutta (orbit, ele, param, err_flag, track, mat6, make_matrix)
Routine to do tracking using Runge-Kutta integration.

track1_symp_lie_ptc (orbit, ele, param, track)
Particle tracking through a single element using a Hamiltonian and a symplectic integrator.

track1_taylor (orbit, ele, taylor, mat6, make_matrix)
Routine to track through an element using the elements taylor series.

track1_time_runge_kutta(orbit, ele, param, err_flag, track, t_end, dt_step)
Routine to track a particle through an element using Runge-Kutta time-based tracking.

44.44 Twiss and Other Calculations

calc_z_tune (branch)
Routine to calculate the synchrotron tune from the full 6X6 1 turn matrix.

chrom_calc (lat, delta_e, chrom_a, chrom_b, err_flag, pz, low_E_lat, high_E_lat, low_E_orb, high_E_orb, ix_branch, orb0)
Routine to calculate the chromaticities by computing the tune change when then energy is changed.

chrom_tune (lat, delta_e, target_x, target_y, err_tol, err_flag)
Routine to set the sextupole strengths so that the lat has the desired chromaticities.

radiation_integrals (lat, orbit, mode, ix_cache, ix_branch, rad_int_by_ele)
Routine to calculate the synchrotron radiation integrals, the emittance, and energy spread.

radiation_integrals_custom (lat, ir, orb, err_flag)
User supplied routine to calculate the synchrotron radiation integrals for a custom element.

relative_mode_flip (ele1, ele2)
Function to see if the modes of ELE1 are flipped relative to ELE2.

set_tune (phi_a_set, phi_b_set, dk1, eles, branch, orb, print_err)
Routine to Q_tune a lat. This routine will set the tunes to within 0.001 radian (0.06 deg).

set_z_tune (branch, z_tune, ok, print_err)
Routine to set the longitudinal tune by setting the RF voltages in the RF cavities.

transfer_twiss (ele_in, ele_out, reverse)
Routine to transfer the twiss parameters from one element to another.

twiss_and_track (lat, orb)
Routine to calculate the Twiss and orbit parameters. This is not necessarily the fastest routine.

twiss_at_element (ele, start, end, average)
Routine to return the Twiss parameters at the beginning, end, or the average of an element.

twiss_and_track_at_s (lat, s, ele, orb_, here)
Routine to calculate the Twiss parameters and orbit at a particular longitudinal position.

twiss_at_start (lat, status, ix_branch, type_out)
Routine to calculate the Twiss parameters at the start of the lat.

twiss_from_tracking (lat, closed_orb_, d_orb, error)
Routine to compute from tracking, for every element in the lat, the Twiss parameters and the transfer matrices.

twiss_propagate1 (ele1, ele2, err_flag)
Routine to propagate the Twiss parameters from the end of ELE1 to the end of ELE2.

twiss_propagate_all (lat, ix_branch, err_flag, ie_start, ie_end)
Routine to propagate the Twiss parameters from the start to the end.

twiss_to_1_turn_mat (twiss, phi, mat2)
Routine to form the 2x2 1-turn transfer matrix from the Twiss parameters.

44.45 Twiss: 6 Dimensional

normal_mode3_calc (t6, tune, B, HV, above_transition, abz_tunes)
Decompose a 2n x 2n symplectic matrix into normal modes. For more details see:

twiss3_propagate_all (lat, ix_branch)
Routine to propagate the twiss parameters using all three normal modes.

twiss3_propagate1 (ele1, ele2, err_flag)
Routine to propagate the twiss parameters using all three normal modes.

twiss3_at_start (lat, err_flag, ix_branch, tune3)
Routine to propagate the twiss parameters using all three normal modes.

44.46 Wakefields

init_wake (wake, n_sr_long, n_sr_trans, n_sr_z, n_lr_mode, always_allocate)
Routine to initialize a wake struct.

randomize_lr_wake_frequencies (ele, set_done)
Routine to randomize the frequencies of the lr wake HOMs.

track1_sr_wake (bunch, ele, test_particle)
Routine to apply the short range wakefields to a bunch.

track1_lr_wake (bunch, ele)
Routine to put in the long-range wakes for particle tracking.

zero_lr_wakes_in_lat (lat)
Routine to zero the long range wake amplitudes for the elements that have long range wakes in a lattice.

44.47 C/C++ Interface

fscalar2scalar (f_scalar, n) result (c_scalar)
Function to translate a scalar from Fortran form to C form.

fvec2vec (f_vec, n) result (c_vec)
Function to translate a vector from Fortran form to C form.

mat2vec (mat, n) result (vec)
Function to take a matrix and turn it into an array in C standard row-major order.

tensor2vec (tensor, n) result (vec)
Function to take a tensor and turn it into an array in C standard row-major order::

vec2mat (vec, mat)
Routine to take a an array in C standard row-major order and turn it into a matrix.

vec2tensor (vec, tensor)
Routine to take a an array in C standard row-major order and turn it into a tensor.

remove_null_in_string (str_in, str_out
Routine to convert a null character in a string to a blank.

f_logic (logic) result (f_log)
Function to convert from a C logical to a Fortran logical.

f_logic_int (logic) result (f_log)
Function to convert from a C logical to a Fortran logical. This function is overloaded by f_logic.

f_logic_bool (logic) result (f_log)
Function to convert from a C logical to a Fortran logical. This function is overloaded by f_logic.

remove_null_in_string_arr (str_in, str_out)
This routine overloaded by: remove_null_in_string

remove_null_in_string_char (str_in, str_out)
This routine overloaded by: remove_null_in_string

to_c_str (f_str, c_str)
Subroutine to append a null (0) character at the end of a string (trimmed of trailing blanks) so it will look like a C character array.

to_f_str (c_str, f_str)
Subroutine to append a null (0) character at the end of a string (trimmed of trailing blanks) so it will look like a C character array.

Part IV

Appendices

Appendix A

Obsolete Straight-Coordinate Generalized Gradients

This appendix documents the obsolete straight-coordinate generalized-gradient (GG) formalism of Venturini and Dragt[Venturini99], which was used by the removed `gen_grad_map` lattice format. It is retained here for reference and to document the field convention used by lattices that predate the curved-coordinate formalism (§18.7). Old `gen_grad_map` lattice files can be converted to the current `gen_gradients` format with the `gen_grad_map_to_gen_gradients` utility.

The electric and magnetic fields are both described by a scalar potential[1]

$$\mathbf{B} = -\nabla \psi_B, \qquad \mathbf{E} = -\nabla \psi_E \tag{A.1}$$

The scalar potential is then decomposed into azimuthal components

$$\psi = \sum_{m=1}^{\infty} \psi_{m,s} \sin(m\theta) + \sum_{m=0}^{\infty} \psi_{m,c} \cos(m\theta) \tag{A.2}$$

where the $\psi_{m,\alpha}$ ($\alpha = c, s$) are characterized by a using functions $C_{m,\alpha}(z)$ which are functions along the longitudinal z-axis.

$$\psi_{m,\alpha} = \sum_{n=0}^{\infty} \frac{(-1)^{n+1} m!}{4^n\, n!\, (n+m)!}\, \rho^{2n+m}\, C_{m,\alpha}^{[2n]}(z) \tag{A.3}$$

The notation $[2n]$ indicates the $2n^{th}$ derivative of $C_{m,\alpha}(z)$.

[1] Notice the negative sign here and in Eq. (A.3) compared to Venturini et al.[Venturini99]. This is to keep the definition of the electric scalar potential ψ_E consistent with the normal definition.

From Eq. (A.3) the field is given by

$$\begin{aligned}
B_\rho &= \sum_{m=1}^{\infty}\sum_{n=0}^{\infty}\frac{(-1)^n\, m!\,(2n+m)}{4^n\, n!\,(n+m)!}\rho^{2n+m-1}\left[C_{m,s}^{[2n]}(z)\,\sin m\theta + C_{m,c}^{[2n]}(z)\,\cos m\theta\right] + \\
&\qquad\qquad \sum_{n=1}^{\infty}\frac{(-1)^n\, 2n}{4^n n!\, n!}\rho^{2n-1}\, C_{0,c}^{[2n]}(z) \\
B_\theta &= \sum_{m=1}^{\infty}\sum_{n=0}^{\infty}\frac{(-1)^n\, m!\, m}{4^n\, n!\,(n+m)!}\rho^{2n+m-1}\left[C_{m,s}^{[2n]}(z)\,\cos m\theta - C_{m,c}^{[2n]}(z)\,\sin m\theta\right] \\
B_z &= \sum_{m=0}^{\infty}\sum_{n=0}^{\infty}\frac{(-1)^n\, m!}{4^n\, n!\,(n+m)!}\rho^{2n+m}\left[C_{m,s}^{[2n+1]}(z)\,\sin m\theta + C_{m,c}^{[2n+1]}(z)\,\cos m\theta\right]
\end{aligned} \tag{A.4}$$

Even though the scalar potential only involves even derivatives of $C_{m,\alpha}$, the field is dependent upon the odd derivatives as well. The multipole index m is such that $m = 0$ is for solenoidal fields, $m = 1$ is for dipole fields, $m = 2$ is for quadrupolar fields, etc. The `sin`–like generalized gradients represent normal (non-skew) fields and the `cos`–like one represent skew fields. The on-axis fields at $x = y = 0$ are given by:

$$(B_x, B_y, B_z) = (C_{1,c}, C_{1,s}, -C_{0,c}^{[1]}) \tag{A.5}$$

The magnetic vector potential for $m = 0$ is constructed such that only A_θ is non-zero

$$\begin{aligned}
A_\rho &= 0 \\
A_\theta &= \sum_{n=1}^{\infty}\frac{(-1)^{n+1}\, 2n}{4^n\, n!\, n!}\rho^{2n-1}\, C_{0,c}^{[2n-1]} \\
A_z &= 0
\end{aligned} \tag{A.6}$$

For $m \neq 0$, the vector potential is chosen so that A_θ is zero.

$$\begin{aligned}
A_\rho &= \sum_{m=1}^{\infty}\sum_{n=0}^{\infty}\frac{(-1)^n\,(m-1)!}{4^n\, n!\,(n+m)!}\rho^{2n+m+1}\left[C_{m,s}^{[2n+1]}\cos(m\theta) - C_{m,c}^{[2n+1]}\,\sin(m\theta)\right] \\
A_\theta &= 0 \\
A_z &= \sum_{m=1}^{\infty}\sum_{n=0}^{\infty}\frac{(-1)^n\,(m-1)!\,(2n+m)}{4^n\, n!\,(n+m)!}\rho^{2n+m}\left[-C_{m,s}^{[2n]}\cos(m\theta) + C_{m,c}^{[2n]}\,\sin(m\theta)\right]
\end{aligned} \tag{A.7}$$

The functions $C_{m,\alpha}(z)$ are characterized by specifying $C_{m,\alpha}(z_i)$ and derivatives at equally spaced points z_i, up to some maximum derivative order $N_{m,\alpha}$ chosen by the user. Interpolation is done by constructing an interpolating polynomial ("non-smoothing spline") for each GG of order $2N_{m,\alpha}+1$ for each interval $[z_i, z_{i+1}]$ which has the correct derivatives from 0 to $N_{m,\alpha}$ at points z_i and z_{i+1}. The coefficients of the interpolating polynomial are easily calculated by inverting the appropriate matrix equation.

The advantages of a generalized gradient map over a cylindrical or Cartesian map decomposition come from the fact that with generalized gradients the field at some point (x, y, z) is only dependent upon the value of $C_{m,\alpha}(z)$ and derivatives at points z_i and z_{i+1} where z is in the interval $[z_i, z_{i+1}]$. This is in contrast to the cylindrical or Cartesian map decomposition where the field at any point is dependent upon *all* of the terms that characterize the field. This "`locality`" property of generalized gradients means that calculating coefficients is easier (the calculation of $C_{m,\alpha}(z)$ at z_i can be done using only the field near z_i independent of other regions) and it is easier to ensure that the field goes to zero at the longitudinal ends of the element. Additionally, the evaluation is faster since only coefficients to

either side of the evaluation point contribute. The disadvantage of generalized gradients is that since the derivatives are truncated at some order $N_{m,\alpha}$, the resulting field does not satisfy Maxwell's equations with the error as a function of radius scaling with the power $\rho^{m+N_{m,\alpha}}$.

It is sometimes convenient to express the fields in terms of Cartesian coordinates. For sine like even derivatives $C_{m,s}^{[2n]}$ the conversion is

$$\begin{aligned}(B_x, B_y) &= (\cos\theta\, B_\rho - \sin\theta\, B_\theta,\, \sin\theta\, B_\rho + \cos\theta\, B_\theta) \\ &= \frac{(-1)^n\, m!}{4^n n!\,(n+m)!}\, C_{m,s}^{[2n]} \Big[(n+m)\,(x^2+y^2)^n\ (S_{xy}(m-1),\, C_{xy}(m-1)) + \\ &\qquad\qquad n\,(x^2+y^2)^{n-1}\ (S_{xy}(m+1),\, -C_{xy}(m+1))\Big]\end{aligned} \tag{A.8}$$

and for the sine like odd derivatives $C_{m,s}^{[2n+1]}$

$$B_z = \frac{(-1)^n\, m!}{4^n n!\,(n+m)!}(x^2+y^2)^n\, C_{m,s}^{[2n+1]}(z)\, S_{xy}(m) \tag{A.9}$$

where the last term in Eq. (A.8) is only present for $n > 0$.

$$\begin{aligned}S_{xy}(m) \equiv \rho^m\ \sin m\theta &= \sum_{r=0}^{2r\le m-1} (-1)^r \binom{m}{2r+1}\, x^{m-2r-1}\, y^{2r+1} \\ C_{xy}(m) \equiv \rho^m\ \cos m\theta &= \sum_{r=0}^{2r\le m} (-1)^r \binom{m}{2r}\, x^{m-2r}\, y^{2r}\end{aligned} \tag{A.10}$$

The conversion for the cosine like derivatives is:

$$\begin{aligned}(B_x, B_y) &= \frac{(-1)^n\, m!}{4^n n!\,(n+m)!}\, C_{m,c}^{[2n]} \Big[(n+m)\,(x^2+y^2)^n\ (C_{xy}(m-1),\, -S_{xy}(m-1)) + \\ &\qquad\qquad n\,(x^2+y^2)^{n-1}\ (C_{xy}(m+1),\, S_{xy}(m+1))\Big] \\ B_z &= \frac{(-1)^n\, m!}{4^n n!\,(n+m)!}(x^2+y^2)^n\, C_{m,c}^{[2n+1]}(z)\, C_{xy}(m)\end{aligned} \tag{A.11}$$

Part V

Bibliography and Index

Bibliography

[Abell06] Dan Abell, "Numerical computation of high-order transfer maps for rf cavities", Phys. Rev. ST Accel. Beams, **9**, pp. 052001, (2006).

[Akima70] H. Akima, "A New Method of Interpolation and Smooth Curve Fitting Based on Local Procedures," J. Assoc. Comp. Mach., **17** pp. 589-602 (1970).

[Alianelli04] L. Alianeli, M. Sanchez del Rio, R. Fe1ici, "I. A Monte Carlo algorithm for the simulation of Bragg scattering by imperfect crystals; and II. Application to mosaic copper," Proc. SPIE 5536, Advances in Computational Methods for X-Ray and Neutron Optics, (2004).

[Barber85] D. P. Barber, J. Kewisch, G. Ripken, R. Rossmanith, and R. Schmidt, "A solenoid spin rotator for large electron storage rings," Particle Accelerators, **17** 243–262 (1985).

[Barber99] D. P. Barber and G. Ripken, "Computer Algorithms and Spin Matching," *Handbook of Accelerator Physics and Engineering,* 1st Edition, 3rd printing, World Scientific Pub. Co. Inc. (1999). Note: 1st edition has more detail than the 2nd edition.

[Bater64] B. Batterman, and H. Cole, "Dynamical Diffraction of X Rays by Perfect Crystals", Rev. Mod. Phys.,**36**, 3, pp. 681–717, (1964).

[Bengt97] J. Bengtsson, "The Sextupole Scheme for the Swiss Light Source (SLS): An Analytic Approach," SLS Note 9/97, Paul Scherrer Institut, (1997).

[Berz89] M. Berz, "Differential Algebraic Description of Beam Dynamics to Very High Orders," Particle Accelerators, Vol. 24, pp. 109-124, (1989).

[Blas94] R. C. Blasdell and A. T. Macrander, "Modifications to the 1989 SHADOW ray-tracing code for general asymmetric perfect-crystal optics," Nuc. Instr. & Meth. A **347**, 320 (1994).

[Bmad] The Bmad web site:
https://www.classe.cornell.edu/bmad

[Brown77] K. L. Brown, F. Rothacker, D. C. Carey, and Ch. Iselin, "TRANSPORT Appendix," Fermilab, unpublished, (December 1977).

[Cai16] Yunhai Cai, "Analytical approach to chromatic correction in the final focus system of circular colliders," Phys. Rev. Accel. Beams **19**, 111006 (2016).

[Carli22] "On the Evolution of the chromatic Functions around an IP computed by MAD-X," CERN-ACC-NOTE-2022-0011 (2022).

[Chao79] Alexander W. Chao, "Evaluation of beam distribution parameters in an electron storage ring," Journal of Applied Physics **50**, 595 (1979).

[Chao81] Alexander W. Chao, "Evaluation of Radiative Spin Polarization in an Electron Storage Ring," Nucl. Instrum. & Meth. **180**, 29, (1981).

[Chao93] Alexander Chao, *Physics of Collective Beam Instabilities in High Energy Accelerators*, Wiley, New York (1993).

[Corbett99] J. Corbett and Y. Nosochkov, "Effect of Insertion Devices in SPEAR–3," Proc. 1999 Part. Acc. Conf., p. 238, (1999).

[Decking00] W. Decking, R. Brinkmann, "Space Charge Problems in the TESLA Damping Ring", Europ. Accel. Phys. Conf., Vienna, (2000).

[Derby09] Norman Derby and Stanislaw Olbert, "Cylindrical Magnets and Ideal Solenoids", arXiv:0909.3880v1 [physics.class-ph], https://doi.org/10.48550/arXiv.0909.3880, (2009)

[Duan15] Zhe Duan, Mei Bai, Desmond P. Barber, Qing Qin, "A Monte-Carlo simulation of the equilibrium beam polarization in ultra-high energy electron (positron) storage rings," Nuc. Instrum. and methods in Phys. Research A **793**, pp 81-91, (2015).

[Duff87] J. Le Duff, *Single and Multiple Touschek Effects*. Proc. CAS Berlin 1987, CERN 89-01, (1987).

[Edwards73] D. A. Edwards and L. C. Teng, "Parametrization of Linear Coupled Motion in Periodic Systems", IEEE Trans. Nucl. Sci. **20**, 3, (1973).

[Elegant] The Elegant web site:
https://ops.aps.anl.gov/elegant.html

[Forest02] É. Forest, F. Schmidt, E. McIntosh, *Introduction to the Polymorphic Tracking Code*, CERN-SL-2002-044 (AP), and KEK-Report 2002-3 (2002). Can be obtained at:
https://frs.web.cern.ch/frs/report/sl-2002-044.pdf

[Forest06] Étienne Forest, 'Geometric integration for particle accelerators," J. Phys. A: Math. Gen. **39** (2006) 5321–5377.

[Forest88] É. Forest, J. Milutinovic, "Leading Order Hard Edge Fringe Fields Effects Exact in $(1+\delta)$ and Consistent with Maxwell's Equations for Rectilinear Magnets," Nuc. Instrum. and Methods in Phys. Research A **269**, pp 474-482, (1988).

[Forest98] É. Forest, *Beam Dynamics: A New Attitude and Framework*, Harwood Academic Publishers, Amsterdam (1998).

[Grote96] H. Grote, F. C. Iselin, *The MAD Program User's Reference Manual*, Version 8.19, CERN/SL/90-13 (AP) (REV. 5) (1996). Can be obtained at:
https://mad.web.cern.ch/mad

[Hartman93] S. C. Hartman and J. B. Rosenzweig, "Ponderomotive focusing in axisymmetric rf linacs", Phys. Rev. E **47**, 2031 (1993).

[HDF5] Hierarchical Data Format Version 5. Web page:
https://www.hdfgroup.org/

[Healy86] L. M. Healy, *Lie Algebraic Methods for Treating Lattice Parameter Errors in Particle Accelerators*. Doctoral thesis, University of Maryland, unpublished, (1986).

[Helm73] R. H. Helm, M. J. Lee, P. L. Morton, and M. Sands, "Evaluation of Synchrotron Radiation Integrals," IEEE Trans. Nucl. Sci. NS-20, 900 (1973).

[Hoff06] G. Hoffstaetter, *High-Energy Polarized Proton Beams, A Modern View*, Springer. Springer Tracks in Modern Physics Vol 218, (2006).

[Hwang15] Hwang and S. Y. Lee, "Dipole Fringe Field Thin Map for Compact Synchrotrons", Phys. Rev. ST Accel. Beams, **12**, 122401, (2015).

[Iselin94] F. C. Iselin, *The MAD program Physical Methods Manual*, unpublished, (1994). Can be obtained at:
https://mad.home.cern.ch/mad

[Jackson76] J. D. Jackson, "On understanding spin-flip synchrotron radiation and the transverse polarization of electrons in storage rings", Rev. Modern Phys., **48**, 3, 1976.

[Jowett87] J. M. Jowett, "Introductory Statistical Mechanics for Electron Storage Rings," AIP Conf. Proc. 153, Physics of Part. Acc., M. Month and M. Dienes Eds., pp. 864, (1987).

[Kim11] Hung Jin Kim, "Symplectic map of crab cavity" Unpublished, FERMILAB-TM-2523-APC.

[Kohn95] V. G. Kohn, "On the Thcory of Reflectivitlby an X-Ray Multilaler Mirror" physica status solidi (b), **187**, 61, (1995).

[Lee99] S. Y. Lee, *Accelerator Physics*, World Scientific, Singapore, (1999).

[Lynch90] Gerald Lynch and Orin Dahl, "Approximations to Multiple Coulomb Scattering", Nucl. Instrum. Meth., B58, 6-10 (1991).

[MAD-NG] Next Generation MAD program at:
https://github.com/MethodicalAcceleratorDesign/MAD

[Magnus54] W. Magnus, "On the exponential solution of differential equations for a linear operator", Commun. Pure Appl. Math, 7, 649-673 (1954).

[McMill75] E. M. McMillan, "Multipoles in Cylindrical Coordinates," Nucl. Instrum. Meth. 127, 471 (1975).

[Montague79] B. W. Montague, "Linear Optics for Improved Chromaticity Correction," CERN LEP note 165 (1979).

[Newton99] D. Newton and A. Wolski, "Fast, Accurate Calculation of Dynamical Maps from Magnetic Field Data Using Generalised Gradients," Proc. PAC09 (2009).

[Ohmi94] K. Ohmi, K. Hirata, and K. Oide, "From the Beam-Envelope Matrix to Synchrotron-Radiation Integrals", Phys. Rev. E **49**, 751 (1994).

[OpenPMD] openPMD: Open standard for particle-mesh data files. Documentation at:
https://github.com/openPMD/openPMD-standard

[Peralta12] Luis Peralta and Alina Louro, "AlfaMC: a fast alpha particle transport Monte Carlo code," arXiv:1211.5960v3 [physics.comp-ph], 2012.

[PGPLOT] PGPLOT graphics library developed by Tim Pearson. Documentation at:
http://www.astro.caltech.edu/~tjp/pgplot/

[Piwin98] Anton Piwinski, *The Touschek Effect in Strong Focusing Storage Rings*. DESY Report 98-179, 1998.

[PLPLOT] PLPLOT graphics library developed primarily by Maurice J. LeBrun and Geoffrey Furnish. Documentation at:
http://plplot.sourceforge.net/

[Quat] Wikipedia article, "Quaternions and spatial rotation",
https://en.wikipedia.org/wiki/Quaternions_and_spatial_rotation.

[Rauben91] T. Raubenheimer, "Tolerances to Limit the Vertical Emittance in Future Storage Rings", Particle Accelerators, 1991, **36**, pp.75-119. SLAC-PUB-4937 Rev., (1991).

[Rio98] Manuel Sanchez del Rio, "Ray tracing simulations for crystal optics," Proc. SPIE 3448, Crystal and Multilayer Optics, **230** (1998).

[Rosen94] J. Rosenzweig and L. Serafini, "Transverse Particle Motion in Radio–Frequency Linear Accelerators," Phys Rev E, Vol. 49, p. 1599, (1994).

[Ruth87] R. D. Ruth, "Single-Particle Dynamics in Circular Accelerators," in AIP Conference Proceedings **153**, *Physics of Particle Accelerators*, pp. 152–235, M. Month and M. Dienes editors, American Institute of Physics, New York (1987).

[Ryne18] C. E. Mayes, R. D. Ryne, D. C. Sagan, "3D SPACE CHARGE IN BMAD", 9^{th} Int. Part. Accel. Conf., Vancouver, pp. 3428–3430, (2018). Software available at: https://github.com/RobertRyne/OpenSpaceCharge.

[SAD] D. Zhou and K. Oide, "Maps Used in SAD" (unpublished). Also see:
https://acc-physics.kek.jp/SAD/

[Sagan91] D. Sagan, "Some aspects of the longitudinal motion of ions in electron storage rings," Nuc. Instr. & Meth. A **307** 2, 171 (1991).

[Sagan03] D. Sagan, J. Crittenden, and D. Rubin, "A Symplectic Model for Wigglers," Part. Acc. Conf. (2003).

[Sagan99] D. Sagan and D. Rubin, "Linear Analysis of Coupled Lattices," Phys. Rev. ST Accel. Beams **2**, 074001 (1999).
https://link.aps.org/doi/10.1103/PhysRevSTAB.2.074001

[Sagan09] D. Sagan, G. Hoffstaetter, C. Mayes, and U. Sae-Ueng, "Extended one-dimensional method for coherent synchrotron radiation including shielding," Phys. Rev. Accel. & Beams **12**, 040703 (2009).

[Sagan17] D. Sagan and C. Mayes, "Coherent Synchrotron Radiation Simulations for Off-Axis Beams Using the Bmad Toolkit", Proc. Europ. Part. Accel. Conf. p. 2829 — 31 (2006).

[Sangwine] Stephen J. Sangwine, Todd A. Ell, Nicolas Le Bihan, "Fundamental Representations and Algebraic Properties of Biquaternions or Complexified Quaternions", Advances in Applied Clifford Algebras, 21, **21**, pp. 607, (2011).

[Schoon11] T. Schoonjans et al. "The xraylib library for X-ray-matter interactions. Recent developments," Spectrochimica Acta Part B: Atomic Spectroscopy **66**, pp. 776-784 (2011). Code for xraylib is at:
https://github.com/tschoonj/xraylib

[Silenko08] Alexander Silenko, "Equation of Spin Motion in Storage Rings in a Cylindrical Coordinate System," Phys. Rev. ST Accel. Beams, **9** pp. 034003, (2006).

[Sommer18] Hannes Sommer, Igor Gilitschenski, Michael Bloesch, Stephan Weiss, Roland Siegwart, and Juan Nieto, "Why and How to Avoid the Flipped Quaternion Multiplication", Aerospace **5**, 72, (2018), https://doi.org/10.3390/aerospace5030072

[Storn96] R. Storn, and K. V. Price, "Minimizing the real function of the ICEC'96 contest by differential evolution" IEEE conf. on Evolutionary Computation, 842-844 (1996).

[Sun10] Yi-Peng Sun, Ralph Assmann, Rogelio Tomás, and Frank Zimmermann, "Crab Dispersion and its Impact on the CERN Large Hadron Collider Collimation", Phys. Rev. ST Accel. Beams, **13**, pp. 031001, (2010).

[Talman87] R. Talman, "Multiparticle Phenomena and Landau Damping," in AIP Conf. Proc. **153**, *Physics of Particle Accelerators*, pp. 789–834, M. Month and M. Dienes editors, American Institute of Physics, New York (1987).

[Tao] D. Sagan, J. Smith, *The Tao Manual*. Can be obtained at: https://www.classe.cornell.edu/bmad/tao_entry_point.html

[Venturini98] M Venturini, D Abell, A Dragt, "Map Computation from Magnetic Field Data and Application to the LHC High-Gradient Quadrupoles," Proc. 1998 International Comp. Accel. Phys. Conf. (ICAP '98), **980914**, 184 (1998).

[Venturini99] M Venturini, A Dragt, "Accurate Computation of Transfer Maps from Magnetic Field Data", Nuc. Instr. & Meth. A **427**, 387–392 (1999).

[Schueren24] S. Van der Schueren et al., "Magnetic Field Modelling and Symplectic Integration of Magnetic Fields on Curved Reference Frames for Improved Synchrotron Design: First Steps".

[GenGrad] `GeneralizedGradients.jl`, Bmad-sim project, https://github.com/bmad-sim/GeneralizedGradients.jl.

[Wang12] C. Wang, "Explicit formulas for 2nd-order driving terms due to sextupoles and chromatic effects of quadrupoles," ANL/APS/LS-330, Argonne National Laboratory, (2012).

[Wiede99] H. Wiedemann, *Particle Accelerator Physics*, Springer, New York, 3rd Edition (2007).

[Wolski06] A. Wolski, "Alternative approach to general coupled linear optics," Phys. Rev. ST Accel. Beams 9, 024001 (2006).

[Wyckoff65] R. W. G. Wyckoff, *Crystal Structures*, Interscience Publ. (1965).

[Yazell09] Douglas J. Yazell, "Origins of the Unusual Space Shuttle Quaternion Definition", 47th AIAA Aerospace Sciences Meeting, January, Orlando, Florida (2009).

Routine Index

ab_multipole_kick, 603
absolute_time_tracking, 618
add_lattice_control_structs, 592
add_superimpose, 519, 592
add_taylor_term, 614
all_pointer_to_string, 589
allocate_branch_array, 506, 582
allocate_element_array, 597
allocate_lat_ele_array, 506, 597
angle_between_polars, 613
angle_to_canonical_coords, 604
append_subdirectory, 586
apply_all_rampers, 597
apply_element_edge_kick_hook, 546, 583
at_this_ele_end, 595
attribute_bookkeeper, 496, 515, 592
attribute_free, 595
attribute_index, 495, 595
attribute_name, 495, 595
attribute_type, 496, 595
autoscale_phase_and_amp, 538, 592

bane1, 591
bbi_kick, 581
bend_edge_kick, 581
bend_photon_init, 605
bend_photon_vert_angle_init, 605
bend_shift, 594
bjmt1, 591
bmad_parser, 486, 487, 495, 501, 506, 523, 527, 544, 600
bmad_parser2, 523, 600
bracket_index, 597
branch_equal_branch, 604
branch_name, 595
bunch_equal_bunch, 604

c_multi, 599
c_to_cbar, 526, 600
calc_bunch_params, 541, 581
calc_bunch_params_slice, 581
calc_bunch_sigma_matrix_etc, 581
calc_emittances_and_twiss_from_sigma_matrix, 581
calc_spin_params, 612
calc_z_tune, 620
cbar_to_c, 600
cesr_getarg, 586
cesr_iargc, 586
check_aperture_limit, 537, 545, 616
check_aperture_limit_custom, 537, 544, 545, 583, 616
check_controller_controls, 597
check_if_s_in_bounds, 596
chrom_calc, 528, 620
chrom_tune, 528, 620
cimp1, 592
clear_lat_1turn_mats, 600
closed_orbit_calc, 616
closed_orbit_from_tracking, 616
combine_consecutive_elements, 600
complex_error_function, 587
complex_taylor_coef, 613
complex_taylor_equal_complex_taylor, 613
complex_taylor_exponent_index, 613
complex_taylor_make_unit, 613
complex_taylor_to_mat6, 613
complex_taylors_equal_complex_taylors, 614
concat_ele_taylor, 614
concat_real_8, 605
concat_taylor, 539, 614
concat_transfer_mat, 601
control_bookkeeper, 515, 597
convert_coords, 604
convert_pc_to, 604
convert_total_energy_to, 604
coord_equal_coord, 604
coord_state_name, 531, 591
coords_curvilinear_to_floor, 594
coords_floor_to_curvilinear, 594
coords_floor_to_local_curvilinear, 594
coords_floor_to_relative, 594

coords_local_curvilinear_to_floor, 594
coords_relative_to_floor, 594
create_element_slice, 519, 536, 592
create_field_overlap, 592
create_girder, 592
create_group, 519, 592
create_lat_ele_nametable, 495, 599
create_overlay, 519, 593
create_planar_wiggler_model, 593
create_sol_quad_model, 600
create_uniform_element_slice, 536
create_unique_ele_names, 600
cross_product, 587
csr_bin_kicks, 582
csr_bin_particles, 582

date_and_time_stamp, 588
deallocate_ele_array_pointers, 598
deallocate_ele_pointers, 494, 496, 597
deallocate_lat_pointers, 598
determinant, 601
dir_close, 586
dir_open, 586
dir_read, 586
do_mode_flip, 601
downcase, 589
downcase_string, 589

ele_compute_ref_energy_and_time, 599
ele_equal_ele, 494, 604
ele_geometry, 498, 548, 594
ele_geometry_hook, 546, 583
ele_geometry_with_misalignments, 594
ele_loc_name, 596
ele_misalignment_l_s_calc, 594
ele_nametable_index, 599
ele_to_fibre, 605
ele_to_fibre_hook, 546, 583
ele_to_taylor, 539, 614
element_at_s, 596
em_field_calc, 545, 584
em_field_custom, 544, 545, 583, 584
equivalent_taylor_attributes, 596, 614
err_exit, 588

f_logic, 622
f_logic_bool, 622
f_logic_int, 622
field_interpolate_3d, 599
file_suffixer, 586
find_element_ends, 596
floor_angles_to_w_mat, 548, 594
floor_w_mat_to_angles, 548, 594
fscalar2scalar, 622
fvec2vec, 622

get_a_char, 586
get_file_time_stamp, 586
get_slave_list, 596
get_tty_char, 586

hdf5_attribute_info, 584
hdf5_close_object, 584
hdf5_exists, 584
hdf5_get_attribute_by_index, 584
hdf5_num_attributes, 584
hdf5_object_info, 584
hdf5_open_dataset, 584
hdf5_open_file, 584
hdf5_open_group, 584
hdf5_open_object, 584
hdf5_read_attribute_alloc_string, 585
hdf5_read_attribute_int, 584
hdf5_read_attribute_real, 585
hdf5_read_attribute_string, 585
hdf5_read_beam, 575, 585
hdf5_read_grid_field, 577, 585
hdf5_write_attribute_string, 584
hdf5_write_beam, 575, 585
hdf5_write_dataset_int, 585
hdf5_write_dataset_real, 585
hdf5_write_grid_field, 577, 585

i_csr, 582
ibs_lifetime, 592
index_nocase, 589
init_beam_distribution, 540, 581
init_bunch_distribution, 581
init_complex_taylor_series, 614
init_coord, 516, 532, 604
init_custom, 544, 583
init_ele, 598
init_lat, 506, 598
init_spin_distribution, 581
init_taylor_series, 614
init_wake, 621
initial_lmdif, 603
insert_element, 519, 593
int_str, 589
integer_option, 588
is_alphabetic, 589
is_false, 589
is_false(param), 496
is_integer, 589

is_logical, 589
is_real, 589
is_true, 589
is_true(param), 496

key_name, 596
key_name_to_key_index, 596
kill_complex_taylor, 614
kill_ptc_layouts, 605
kill_taylor, 539, 614
kind_name, 605

lat_compute_ref_energy_and_time, 515, 548, 599
lat_ele_locator, 486–488, 520, 596
lat_equal_lat, 506, 604
lat_geometry, 498, 515, 548, 595
lat_make_mat6, 487, 499, 515, 601
lat_sanity_check, 596
lat_to_ptc_layout, 553, 605
lat_vec_equal_lat_vec, 604
lattice_bookkeeper, 487, 488, 516, 598
linear_fit, 587
location_decode, 590
location_encode1, 590
logic_option, 589
logic_str, 590
lunget, 586

mad_add_offsets_and_multipoles, 618
mad_concat_map2, 619
mad_drift, 618
mad_elsep, 618
mad_map_to_taylor, 619
mad_quadrupole, 619
mad_rfcavity, 619
mad_sbend, 618
mad_sbend_body, 619
mad_sbend_fringe, 618
mad_sextupole, 618
mad_solenoid, 619
mad_tmfoc, 619
mad_tmsymm, 619
mad_tmtilt, 619
mad_track1, 619
make_g2_mats, 601
make_g_mats, 601
make_hybrid_lat, 593
make_mad_map, 618
make_mat6, 499, 515, 601
make_mat6_bmad, 613
make_mat6_custom, 544, 583, 613
make_mat6_mad, 618
make_mat6_symp_lie_ptc, 613
make_mat6_taylor, 613
make_mat6_tracking, 613
make_srdt_cache, 616
make_unit_mad_map, 619
make_v_mats, 526, 601
map_coef, 605
mat2vec, 622
mat6_from_s_to_s, 536, 601
mat6_to_complex_taylor, 614
mat6_to_taylor, 601, 615
mat_eigen, 588
mat_inverse, 588
mat_make_unit, 588
mat_rotation, 588
mat_symp_conj, 588
mat_symp_decouple, 588
mat_symp_error, 588
mat_symplectify, 588
mat_type, 588
match_ele_to_mat6, 601
match_reg, 590
match_wild, 590
match_word, 590
mexp, 603
milli_sleep, 586
modulo2, 587
mpxx1, 592
mpzt1, 592
multi_turn_tracking_analysis, 616
multi_turn_tracking_to_mat, 601, 616
multipass_all_info, 602
multipass_chain, 602
multipole_ab_to_kt, 603
multipole_ele_to_ab, 603
multipole_ele_to_kt, 603
multipole_init, 603
multipole_kick, 603
multipole_kicks, 603
multipole_kt_to_ab, 603

n_attrib_string_max_len, 596
nametable_add, 599
nametable_bracket_indexx, 599
nametable_change1, 599
nametable_init, 600
nametable_remove, 600
new_control, 519, 593
normal_form_rd_terms, 605
normal_mode3_calc, 526, 621

num_lords, 596

odeint_bmad, 545, 618
offset_particle, 616
offset_photon, 616
on_off_logic, 590
one_turn_mat_at_ele, 601
opti_de, 603
opti_lmdif, 603
orbit_amplitude_calc, 617
orbit_to_local_curvilinear, 595
order_particles_in_z, 581
order_super_lord_slaves, 599
ordinal_str, 590
out_io, 586
out_io_called, 586
out_io_end, 586
out_io_line, 586
output_direct, 586

parse_fortran_format, 590
particle_is_moving_backwards, 617
particle_is_moving_forward, 532, 617
patch_flips_propagation_direction, 595
pmd_read_complex_dataset, 585
pmd_read_int_dataset, 585
pmd_read_real_dataset, 585
pmd_write_complex_to_dataset, 585
pmd_write_int_to_dataset, 585
pmd_write_int_to_pseudo_dataset, 585
pmd_write_real_to_dataset, 585
pmd_write_real_to_pseudo_dataset, 585
pmd_write_units_to_dataset, 585
pointer_to_attribute, 521, 593
pointer_to_branch, 593
pointer_to_ele, 593
pointer_to_ele_multipole, 593
pointer_to_element_at_s, 593
pointer_to_indexed_attribute, 596
pointer_to_lord, 511, 513, 596
pointer_to_multipass_lord, 596, 602
pointer_to_next_ele, 593
pointer_to_slave, 511, 513, 596
pointers_to_attribute, 521, 593
polar_to_spinor, 612
polar_to_vec, 612
ptc_transfer_map_with_spin, 605

qp_axis_niceness, 606
qp_calc_and_set_axis, 564, 566, 606, 608
qp_calc_axis_params, 606
qp_calc_axis_places, 606
qp_calc_axis_scale, 606
qp_calc_minor_div, 606
qp_clear_box, 607
qp_clear_box_basic, 611
qp_clear_page, 607
qp_clear_page_basic, 611
qp_close_page, 564, 606
qp_close_page_basic, 611
qp_convert_point_abs, 611
qp_convert_point_rel, 611
qp_convert_rectangle_rel, 606
qp_draw_axes, 564, 566, 607
qp_draw_circle, 607
qp_draw_curve_legend, 607
qp_draw_data, 564, 607
qp_draw_ellipse, 607
qp_draw_graph, 607
qp_draw_graph_title, 607
qp_draw_grid, 607
qp_draw_histogram, 607
qp_draw_line, 607
qp_draw_main_title, 607
qp_draw_polyline, 607
qp_draw_polyline_basic, 607
qp_draw_polyline_no_set, 607
qp_draw_rectangle, 567, 608
qp_draw_symbol, 608
qp_draw_symbol_basic, 611
qp_draw_symbols, 608
qp_draw_text, 564, 565, 608
qp_draw_text_basic, 608
qp_draw_text_legend, 607
qp_draw_text_no_set, 608
qp_draw_x_axis, 608
qp_draw_y_axis, 608
qp_eliminate_xy_distortion, 608
qp_from_inch_abs, 610
qp_from_inch_rel, 610
qp_get_axis_attrib, 610
qp_get_layout_attrib, 610
qp_get_line_attrib, 610
qp_get_parameters, 610
qp_get_symbol_attrib, 610
qp_init_com_struct, 611
qp_join_units_string, 611
qp_justify, 611
qp_open_page, 564, 567, 606
qp_open_page_basic, 611
qp_paint_rectangle, 608
qp_paint_rectangle_basic, 611
qp_pointer_to_axis, 611

qp_read_data, 564, 565, 611
qp_restore_state, 564, 566, 611
qp_restore_state_basic, 611
qp_save_state, 564, 566, 611
qp_save_state_basic, 611
qp_select_page, 606
qp_select_page_basic, 611
qp_set_axis, 564, 568, 608
qp_set_box, 564, 566, 567, 608
qp_set_char_size_basic, 612
qp_set_clip, 609
qp_set_clip_basic, 612
qp_set_color_basic, 612
qp_set_graph, 608
qp_set_graph_attrib, 564, 609
qp_set_graph_limits, 609
qp_set_graph_placement, 609
qp_set_graph_position_basic, 612
qp_set_layout, 609
qp_set_line, 609
qp_set_line_attrib, 564, 566, 609
qp_set_line_width_basic, 612
qp_set_margin, 564, 566, 567, 609
qp_set_page_border, 564, 565, 567, 609
qp_set_page_border_to_box, 609
qp_set_parameters, 568, 609
qp_set_symbol, 609
qp_set_symbol_attrib, 564, 566, 609
qp_set_symbol_fill_basic, 612
qp_set_symbol_size_basic, 612
qp_set_text_attrib, 609
qp_set_text_background_color_basic, 612
qp_split_units_string, 612
qp_subset_box, 609
qp_text_height_to_inches, 610
qp_text_len, 610
qp_text_len_basic, 612
qp_to_axis_number_text, 608
qp_to_inch_abs, 610
qp_to_inch_rel, 610
qp_to_inches_abs, 610
qp_to_inches_rel, 610
qp_use_axis, 568, 609
quote, 590
quoten, 590

radiation_integrals, 527, 620
radiation_integrals_custom, 544, 583, 620
ran_engine, 587
ran_gauss, 587
ran_gauss_converter, 587
ran_seed_get, 587
ran_seed_put, 587
ran_uniform, 587
randomize_lr_wake_frequencies, 621
re_allocate, 589
re_allocate_eles, 597
re_associate, 589
read_a_line, 587
read_beam_file, 581
read_digested_bmad_file, 524, 600
real_num_fortran_format, 590
real_option, 589
real_str, 590
real_to_string, 590
reallocate_beam, 582
reallocate_bunch, 582
reallocate_coord, 533, 598
reallocate_coord_array, 533, 598
reals_to_string, 590
reals_to_table_row, 590
relative_mode_flip, 620
release_rad_int_cache, 599
remove_constant_taylor, 605
remove_eles_from_lat, 519, 593
remove_null_in_string, 622
remove_null_in_string_arr, 622
remove_null_in_string_char, 622
rf_is_on, 597

s_calc, 515, 548, 595
set_custom_attribute_name, 504, 598
set_ele_attribute, 521, 593
set_ele_defaults, 598
set_ele_status_stale, 593
set_flags_for_changed_attribute, 487, 516, 526, 599
set_on_off, 598
set_ptc, 539, 552, 605
set_status_flags, 593
set_tune, 527, 620
set_z_tune, 527, 620
skip_header, 587
slice_lattice, 598
sol_quad_mat6_calc, 602
sort_complex_taylor_terms, 614
sort_taylor_terms, 615
sort_universal_terms, 605
spin_omega, 613
spinor_to_polar, 612
spinor_to_vec, 612
spline_akima, 588

spline_evaluate, 588
split_lat, 519, 593
splitfilename, 587
srdt_calc, 616
srdt_calc_with_cache, 616
srdt_lsq_solution, 616
start_branch_at, 598
str_downcase, 591
str_match_wild, 591
str_set, 591
str_substitute, 591
string_option, 589
string_to_int, 591
string_to_real, 591
string_trim, 591
string_trim2, 591
suggest_lmdif, 603
super_ludcmp, 588
super_mrqmin, 603
switch_attrib_value_name, 496, 597
symp_lie_bmad, 619
system_command, 587

taper_mag_strengths, 598
taylor_coef, 539, 615
taylor_equal_taylor, 615
taylor_inverse, 615
taylor_make_unit, 615
taylor_minus_taylor, 615
taylor_plus_taylor, 615
taylor_propagate1, 615
taylor_to_genfield, 605
taylor_to_mad_map, 619
taylor_to_mat6, 601, 615
taylor_to_real_8, 605
taylors_equal_taylors, 615
tensor2vec, 622
tilt_coords, 617
tilt_mat6, 602
time_runge_kutta_periodic_kick_hook, 546, 583
to_c_str, 622
to_eta_reading, 602
to_f_str, 622
to_orbit_reading, 602
to_phase_and_coupling_reading, 602
to_str, 591
touschek_lifetime, 583
track1, 515, 532, 544, 547, 617
track1_bmad, 619
track1_bunch, 541, 581
track1_bunch_csr, 617
track1_bunch_hom, 581
track1_bunch_hook, 546, 583
track1_custom, 544, 583, 620
track1_linear, 620
track1_lr_wake, 621
track1_mad, 619
track1_postprocess, 544, 546, 583, 620
track1_preprocess, 544, 546, 583, 620
track1_radiation, 620
track1_runge_kutta, 620
track1_spin, 613
track1_spin_custom, 544, 583, 617
track1_sr_wake, 621
track1_symp_lie_ptc, 620
track1_taylor, 620
track1_time_runge_kutta, 620
track1_wake_hook, 546, 583
track_a_bend, 618
track_a_drift, 618
track_all, 533, 617
track_beam, 541, 582
track_bunch, 582
track_bunch_time, 582
track_complex_taylor, 614
track_from_s_to_s, 536, 617
track_many, 534, 617
track_taylor, 615
transfer_branch, 582
transfer_branches, 582
transfer_ele, 598
transfer_ele_taylor, 598, 615
transfer_eles, 598
transfer_lat, 598
transfer_lat_parameters, 598
transfer_map_calc, 615
transfer_map_from_s_to_s, 615
transfer_mat2_from_twiss, 601
transfer_mat_from_twiss, 601
transfer_matrix_calc, 601
transfer_twiss, 535, 620
truncate_complex_taylor_to_order, 614
truncate_taylor_to_order, 615
twiss1_propagate, 597
twiss3_at_start, 621
twiss3_propagate1, 621
twiss3_propagate_all, 621
twiss_and_track, 527, 617, 621
twiss_and_track_at_s, 527, 536, 617, 621
twiss_and_track_from_s_to_s, 536, 617
twiss_and_track_intra_ele, 527, 536, 617
twiss_at_element, 621

twiss_at_start, 486, 487, 526, 621
twiss_from_mat2, 602
twiss_from_mat6, 602
twiss_from_tracking, 617, 621
twiss_propagate1, 526, 621
twiss_propagate_all, 486, 487, 526, 527, 621
twiss_to_1_turn_mat, 602, 621
type_complex_taylors, 614
type_coord, 604
type_ele, 486, 487, 491, 597
type_map, 606
type_map1, 606
type_ptc_fibre, 606
type_ptc_layout, 605
type_real_8_taylors, 606
type_taylors, 615
type_this_file, 587
type_twiss, 597

universal_equal_universal, 604
universal_to_bmad_taylor, 606
unquote, 591
upcase, 591
upcase_string, 591
update_floor_angles, 595

valid_mat6_calc_method, 538, 597
valid_tracking_method, 538, 597
value_of_attribute, 521, 594
vec2mat, 622
vec2tensor, 622
vec_to_polar, 612
vec_to_spinor, 612

w_mat_for_bend_angle, 595
w_mat_for_tilt, 595
w_mat_for_x_pitch, 595
w_mat_for_y_pitch, 595
wall_hit_handler_custom, 544, 545, 584, 618
write_beam_file, 582
write_beam_floor_positions, 582
write_bmad_lattice_file, 524, 600
write_digested_bmad_file, 524, 600
write_lattice_in_foreign_format, 524, 600

zero_ele_kicks, 599
zero_ele_offsets, 599
zero_lr_wakes_in_lat, 622

Index

! comment symbol, 35
$
    character to denote a parameter, 483
& continuation symbol, 36

ab_kicker, **64**
ab_multipole, **63**, 198, 223, 226, 229
abs, 51
abs_tol_adaptive_tracking, 231, 262
abs_tol_tracking, 262
absolute time tracking, 421, 538
absolute_time_tracking, 262
ac_kicker, 347
accordion_edge, 113
acos, 51
action-angle, **393**
alias, **168**
alpha_a, 66, 259
alpha_a_strong, 67
alpha_angle, 83
alpha_b, 66, 259
alpha_b_strong, 67
an, *see* multipole, an
angle, 70, 71, 151, 167, 332
anomalous_moment_of, 48, 51
antimuon, 46
antimuon$, 509
antiproton, 46
antiproton$, 509
aperture, **179**, 198, 497
aperture_at, **179**, 181
aperture_limit_on, 263
aperture_type, 87, 88, 123, 179
apply_element_edge_kick_hook, 546
arithmetic expressions, 49
    constants, 49
    intrinsic functions, *see* intrinsic functions
asin, 51
Astra, 292
atan, 51
auto
    mat6_calc_method, 225
automatic field scaling, **210**

b1_gradient, 73, 144, 155, 168
b2_gradient, 74, 154, 168
b3_gradient, 133, 168
b_field, 70, 71, 168
b_field_tot, 70
b_max, 162, 167
b_param, 83
bbi_constant, 66, 167
beam, 38, 371
beam initialization parameters, **273**, **278**
beam line, *see* line
beam statement, **258**
beam tracking
    list of routines, 581
beam_init, 281
beam_init_struct, **273**
beambeam, **66**, 167, 223, 226, 229, 257, 425, 508
beginning, 38
beginning element, 29, 42
beginning statement, 36, **259**, 331
beginning_ele, **69**, 106, 172
bendfringe, 219
beta_a, 66, 259
beta_a0, 125
beta_a1, 125
beta_a_strong, 67
beta_b, 66, 259
beta_b0, 125
beta_b1, 125
beta_b_strong, 67
bl_hkick, 168, **179**
bl_kick, 168, **179**
bl_vkick, 168, **179**
Blender, 292
Bmad, **2**
    distribution, 479
    error reporting, 3
    general parameters, **262**

information, 3
lattice file format, 35
lattice format, *see* lattice file format
statement syntax, 35
bmad version number, 523
bmad_com, 267, 549
bmad_common_struct, **262**
absolute_time_tracking, 538
aperture_limit_on, 537
auto_bookkeeper, 515
max_aperture_limit, 537
bmad_parser, 523
bmad_standard
mat6_calc_method, 225
tracking method, 341
tracking_method, 222
bn, *see* multipole, bn
bookkeeper_status_struct, 516
bookkeeping
automatic, 515
intelligent, 515
both_ends, **181**
bragg_angle, 83
bragg_angle_in, 83
bragg_angle_out, 83
branch, 29, 30, 235, 260
root, 506
branch_struct, 506
bs_field, 151, 155, 168
bs_gradient, 156
bunch, 371
bunch initialization, **371**

C++ interface, 559
classes, 559
Fortran calling C++, 560
C/C++ interface
list of routines, 622
calc_reference_orbit, 38
call, 38
inline, 54
call statement, **54**
canonical coordinates, *see* phase space coordinates
capillary, **76**, 197, 223, 226, 229
wall, 193
cartesian_map, 199, 202, 351
cavity_type, 81, 119
change, 281
charge, 66, 167
charge_of, 51
chromaticity, 528
closed, 256
closed orbit, 536
cmat_ij, 259
coherent synchrotron radiation, *see* CSR
coherent tracking, 453
combine_consecutive_elements, 38
comment symbol (!), 35
complex taylor map
list of routines, 613
constant_ref_energy, 94
constants, **48**, 547
continuation symbol (&), 36
continuous, **181**
control_struct, 512
controller element, 29
conversion to other lattice formats, 291
converter, **78**, 432
coord_array_struct, 533
coord_struct, 529
coordinates, **325**
global, *see* global coordinates
list of routines, 604
phase space, *see* phase space coordinates
reference, *see* reference orbit
cos, 51
coupler_angle, 209
coupler_at, 209
coupler_phase, 209
coupler_strength, 209
coupling, *see* normal mode
crab_cavity, **81**
critical_angle_factor, **76**
crotch chamber geometry, 194
crunch, 219
crunch_calib, 219
crystal, **83**, 181, 197, 226, 229, 327, 334, 456, 457
tilt correction, 462
crystal_type, 83
CSR, **376**, **379**
csr and space charge methods, 230
csr_and_space_charge_on, 262
csr_method, 230
CSRTrack, 292
custom, **86**, 223, 226, 229, 232, 543, 544
mat6_calc_method, 225
reference energy, 339, 548
spin_tracking_method, 228
tracking_method, 222
custom_attributeN, 255

cylindrical map, 203, 353
cylindrical_map, 199

d1_thickness, 131
d2_thickness, 131
d_orb(6), 262
d_spacing, 83
db_field, 70, 71, 168
dbragg_angle_de, 83
de optimizer parameters, **270**
de_eta_meas, 219
debug_marker, 38
debug_marker statement, 60
default_ds_step, 262
default_integ_order, 262
default_tracking_species, 256, 260
delta_e, 167
delta_e_ref, 86
dependent attribute, 167
descrip, **168**, 544
detector, **87**, 219
dg, 70, 73, 168
diffraction_plate, **88**, 181
digested files, 36, 523
dispersion, 375, 385, **394**
downstream element end, 327
dphi_a, 125
dphi_b, 125
drift, **90**, 223, 226, 229, 244, 436
  superposition, 245
driving terms
  list of routines, 616
ds_slice, 140
ds_step, 167, 230–232, 539
dynamic_aperture_struct, **271**

e1, 70, 71, 151
e2, 70, 71, 151
E2_center, 141
E2_probability, 141
E_center, 141
E_center_relative_to_ref, 141
e_field, 93, 167
e_field_x, 141
e_field_y, 141
e_gun, **91**, 172, 210, 258
e_loss, 81, 119, 167
e_tot, 167, 172, 252, 255, 256, 259
e_tot_offset, 136
e_tot_set, 136
e_tot_start, 172
ecollimator, **77**, 179, 181, 223, 226, 229
ele
  %status, 516
ele_geometry_hook, 546
ele_origin, 243
ele_pointer_struct, 486
ele_struct, 486, **491**
  %a, 497, 525
  %a_pole(:), 503
  %alias, 495
  %aperture_type, 537
  %b, 497, 525
  %b_pole(:), 503
  %bmad_logic, 504
  %branch, 496
  %c, 525
  %c_mat, 497, 525
  %cartesian_map, 500
  %component_name, 497
  %converter, 503
  %cylindrical_map, 500
  %dc_mat_dpz, 497
  %descrip, 495
  %emit, 525
  %field_master, 496
  %floor, 498
  %foil, 503
  %gamma_c, 497, 525
  %gen_gradients, 500
  %grid_field, 500
  %high_energy_space_charge, 502
  %ic1_lord, 497, 513
  %ix1_slave, 497
  %ix_branch, 496
  %ix_ele, 496, 520
  %ix_pointer, 503
  %ixx, 504
  %iyy, 504
  %izz, 504
  %key, 487, 495
  %logic, 503
  %lord_status, 497, 509, 510
  %map_ref_orb_in, 499
  %map_ref_orb_out, 499
  %mat6, 499, 538
  %mode3, 497, 526
  %mode_flip, 497
  %n_lord, 497, 511, 513
  %n_lord_field, 497, 513
  %n_lord_ramper, 497
  %n_slave, 497, 511

%n_slave_field, 497
%name, 495
%norm_emit, 525
%old_value(:), 496
%r, 503
%rad_map, 502
%ref_time, 499
%rf, 503
%s, 487, 499
%sigma, 525
%sigma_p, 525
%slave_status, 497, 509, 510
%spin_q, 502
%spin_taylor, 502
%spin_taylor_ref_orb_in, 502
%sub_key, 495
%tracking_method, 538
%type, 495
%value(:), 495
%vec0, 499
%wake, 500
%x, 487
%z, 497, 525
attribute values, 495
components not used by Bmad, 503
in lat_struct, 507
initialization, 494
multipoles, 503
pointer components, 494
Taylor maps, 499
transfer maps, 499
ele_to_fibre_hook, 546
electric fields, 347
map decomposition, 350
electric_dipole_moment, 256
electron, 46
electron$, 509
Elegant, 292
element, 29, **61**
class, 221
matching to names, 39
name, 38
names, 38
table of class types, 223
element attribute, 167
dependent and independent, 167
Element attribute bookkeeping, 494
element attributes, **42**
defining custom attributes, 44
element body coordinates, **335**
element coordinates, 327, 423, 456
element reversal, 237
elseparator, **93**, 167, 179, 223, 226, 229, 347, 436, 439
em_field, **94**, 172, 210, 421
reference energy, 339
emittance_a, 258
emittance_b, 258
emittance_z, 258
end element, 29, 257
end_edge, 113
end_file, 38
end_file statement, **56**
energy, 258
energy_distribution, 141
energy_probability_curve, 141
Enge function, 72
entrance element end, 327
entrance_end, **181**, 327, 537
eps_step_scale, 151
eta_x, 259
eta_x0, 125
eta_y, 259
eta_y0, 125
etap_x, 259
etap_x0, 125
etap_y, 259
etap_y0, 125
exact_misalign, 232, 267
exact_model, 232
exact_multipoles, 70, 71, 350
exit element end, 327
exit_end, **181**, 327, 537
exp, 51
expand_lattice, 38, 52, **56**, 240

F (multipole scale factor), 347
f1, 144, 151
f2, 144, 151
factorial, 51
feedback, **95**
fftw
library, 480
fgsl
library, 480
fibre, 553
fiducial, **96**, 106, 335
field maps, 199
field_calc, 86, 230, 231
field_master, 167
field_overlaps, 209, 248
field_scale, 200

field_scale_factor, 88, 123
field_type, 200
field_x, 258
field_y, 258
fint, 70, 72
fintx, 70, 72
fixed_step_runge_kutta
    tracking_method, 222
fixed_step_time_runge_kutta
    tracking_method, 222
flexible, 136
flexible patch, 96, 137
floor coordinates, *see* global coordinates
floor_position_struct, 498
floor_shift, **100**, 327
    coordinate transformation, 334
foil, **102**, 437
follow_diffracted_beam, 83
Forest, Étienne, *see* PTC/FPP
fork, **105**, 106, 506, 535
FPP, *see* PTC/FPP
free$, 510
fringe fields, 216
fringe_at, 217
fringe_type, 151, 217
functions, *see* intrinsic functions

g, 70, 73, 151, 167, 168
g_max, 162
g_tot, 70
gang, **168**
gap, 93, 167
gen_gradients, 199, 207
generalized gradient field modeling, 357
geometry, 244, 255, 256, 260, 486
getf, 481
girder, **108**, 113, 135, 197, 244, 335, 498, 507, 513
girder_lord, 31
girder_lord$, 509
gkicker, **111**
global coordinates, **330**, 548
    in ele_struct, 498
    list of routines, 594
    reference orbit origin, 331
global_com, 549
    parallel processing, 549
global_common_struct, **549**
    parallel processing, 549
GPT, 292
grad_loss_sr_wake, 262
gradient, 81, 91, 119, 167
gradient_err, 91
graze_angle, 131
grid field, 204
grid_field, 199
group, 50, **112**, 135, 244, 497, 507, 513
    reference energy, 548
    syntax, 168
group_lord, 31, 509
gsl
    library, 480

h1, 70, 73
h2, 70, 73
h_displace, 118
harmon, 147, 149, 151, 496
hdf5, 575
    library, 480
hdf5 and grid_field data, 577
hdf5 and particle beam data, 575
hgap, 70
hgapx, 70, 72
high_energy_space_charge_on, 262
hkick, 93, 117, 118, 167, 168, **179**
hkicker, **117**, 179, 223, 226, 229, 347, 439
hybrid, **115**
    reference energy, 339, 548

incoherent tracking, 453
inflexible patch, 137
instrument, **116**, 219, 223, 226, 229
integration methods, 230
integrator_order, 230, 232, 539
intrinsic functions, 51
is_on, 106, 122, **197**

k1, 70, 73, 144, 155, 167, 168
k1x, 162
k1y, 162
k2, 74, 154, 168
k3, 133, 168
kick, 117, 168, **179**
kicker, **118**, 179, 223, 226, 229, 347, 439
kill_fringe, 151
knl, *see* multipole, knl
ks, 151, 155, 156, 168

l, 74, 83, 100, 128, 130, 167, **197**
l_arc, 70, 74
l_chord, 70, 74, 167, **197**
l_period, 162
l_pole, 162

l_rectangle, 74
l_sagitta, 70, 74
laboratory coordinates, **326**, 335
lapack
  library, 480
lapack95
  library, 480
lat_param_struct, 508
  %n_part, 508
  %stable, 508
  %t1_no_RF, 508
  %t1_with_RF, 508
  %total_length, 508
  end_lost_at, 537
  ix_lost, 537
  ix_track, 533
  lost, 537
lat_struct, 486, 497, **505**
  %branch(:), 506
  %control, 513
  %ele(:), 487, 507
  %ele_init, 506
  %ic, 513
  %ix1_slave, 513
  %n_ele_max, 507
  %n_ele_track, 507
  %n_slave, 513
  %n_slave_field, 513
  %param, *see* lat_param_struct
  %particle_start, 516
  example use of, 487
  initializing, 506
  pointers, 506
lattice, 30, 255, 256
  expansion, **56**
lattice element, 29
lattice expansion, 50
lattice files, **35**
  MAD files, 524
  name syntax, 36
  parser debugging, 60
  reading, 523
  reading and writing routines, 600
lcavity, **119**, 167, 172, 209, 210, 223, 226, 229, 252, 256, 257, 340, 341, 421, 425, 439, 500
  and geometry, 256
  and param%n_part, 508
  reference energy, 339, 548
length of elements, 197
limit, **179**
line, 37, **235**, 487
  with arguments, 238
line slice, 237
linear, 538
  tracking_method, 222
linear_leading, 214
linear_trailing, 214
list, 235, **238**
listf, 481
live_branch, 255, 256
log, 51
logicals, **46**
lord, **509**
lord_pad1, 247
lord_pad2, 247
lord_status, 31
lords
  ordering, 512
lr_freq_spread, 501
lr_wakes_on, 262

machine, 257
macroparticles, **375**
  tracking, 371
MAD, **3**, 61, 238, 259, 331, 332, 345, 524
  beam statement, 258
  conversion, 291
  delayed substitution, 50
  element rotation origin, 175
  MAD-8, 524
  mat6_calc_method, 225
  phase space convention, 342
  syntax compatibility with BMAD, 50
  tracking_method, 223
  units, 48
mad_tpsa
  library, 480
magnetic fields, **345**
  map decomposition, 350
map, *see* transfer map
  with radiation included, 382
marker, 91, 106, **122**, 219, 223, 226, 229
mask, **123**
mass_of, 48, 51
master_parameter, 200
mat6_calc_method, 86, 221, **225**
match, **125**, 223, 226, 229
material_type, 131
matrix
  list of routines, 600
max_aperture_limit, 262

max_fringe_order, 218, 267
measurement, 471
measurement simulations
  list of routines, 602
merge_elements, 38
Merlin, 292
minor_slave$, 510
mirror, **129**, 181, 327, 334, 456, 457
mode, 88, 123
mode3_struct, 526
modules, 484
monitor, **116**, 219, 223, 226, 229
mp_threading_is_safe, 262
multilayer_mirror, **131**, 181, 456
multipass, 38, 56, 149, 249
  list of routines, 602
multipass_lord, 31, 249, 512
multipass_lord$, 509
multipass_ref_energy, 252
multipass_slave, 31, 249, 512
multipass_slave$, 510
multipole, **130**, 198, 223, 226, 229, 346
  %scale_multipoles, 503
  an, bn, **198**, **346**
    in ele_struct, 503
  KnL, Tn, **345**, **346**
    in ele_struct, 503
  knl, tn, **198**
  list of routines, 603
multipoles_on, 199
muon, 46
muon$, 509

n_cell, 81, 119, 131
n_part, 255, 257, 258
  in BeamBeam element, 66
n_pole, 162
n_sample, 219
n_slice, 66
no_aperture, 181
no_digested, 38
no_digested statement, 60
no_end_marker, 255, 257
no_major_lord, 31
no_superimpose, 38
no_superimpose statement, 55
noise, 219
none, 214
normal mode
  a–mode, 389
  b–mode, 389
  Coupling, 389
not_a_lord, 31
not_a_lord$, 509
null_ele, 91, **132**
num_steps, 167, 230, 231

octupole, **133**, 168, 223, 226, 229, 347, 442
  tilt default, 175
offset, 243
offset_moves_aperture, **179**
OPAL, 557
  phase space, 557
open, 244, 256
open_spacecharge
  library, 480
opti_de_param_struct, 270
orbit
  measurement, 471
origin_ele, 96, 100, 108
origin_ele_ref_pt, 96, 100, 108
osc_amplitude, 219
overlay, 50, 110, 113, **134**, **145**, 244, 497, 507, 513, 522
  reference energy, 548
  syntax, 168
overlay_lord, 31

p0c, 172, 252, 255, 256, 259
p0c_set, 136
p0c_start, 172
pancake_canonical, 267
pancake_symplectic, 267
parameter, 38
parameter statement, 35, 36, 167, 172, **255**
parameter types, 46
paraxial approximation, 341
parser_debug, 38
parser_debug statement, 60
particle, 257, 258, 260
particle species name, 46
particle_start, 38, 516
particle_start statement, **258**
patch, 96, 105, **136**, 172, 197, 223, 226, 229, 256, 327, 329, 442, 512
  and chamber wall, 194
  coordinate transformation, 334
  example, 285
  reference energy, 339
  reflection, 335
pendellosung_period_pi, 83
pendellosung_period_sigma, 83

permfringe, 219
pgplot
  and Quick_Plot, 563
  library, 480
phase space coordinates, **339**, **340**, 529
  MAD convention, 342
phase_x, 258
phase_y, 258
phi0, 81, 119, 147, 149, 151
phi0_autoscale, 91
phi0_multipass, 81, 119, 147, 149
phi_a, 259
phi_b, 259
phi_origin, 96, 108
phi_position, 260
photon, 46
  phase space coordinates, **342**
photon_fork, **105**, 106, 506, 535
photon_init, **140**
photon_type, 255, 257
photons
  list of routines, 605
physical_source, 141
pion+, 46
pion-, 46
pion0, 46
pion_0$, 509
pion_minus$, 509
pion_plus$, 509
pipe, **116**
  superposition, 245
plplot
  library, 481
polarity, 162
positron, 46
positron$, 509
print, 38
print statement, **53**
programming
  conventions, 483
  example program, 485
  precision (rp), 483
proton, 46
proton$, 509
psi_angle, 83, 463
psi_origin, 96, 108
psi_position, 260
PTC, 27
  single element mode, 475
  whole lattice mode, 476
PTC integration, 232
PTC/FPP, 475, 551
  initialization, 552
  library, 480
  list of routines, 605
  patch, 554
  phase space, 551
  real_8, 554
  Taylor Maps, 554
  universal_taylor, 554
PTC/FPP variable
  initialization, 553
ptc_exact_model, 267
px, 258
px0, 127
px1, 127
px_kick, 111
py, 258
py0, 127
py1, 127
py_kick, 111
pz, 258
pz0, 127
pz1, 127
pz_kick, 111

qp_axis_struct, 573
qp_line_struct, 573
qp_symbol_struct, 573
quadrupole, **144**, 168, 223, 226, 229, 347, 425, 443
  tilt default, 175
quick plot
  list of routines, 606
quick_plot, **563**
  axes, 568
  color styles, 569
  fill styles, 569
  line styles, 568
  position units, 567
  structures, 573
  symbol styles, 570
  symbol table, 570

r0_mag, 198, 347
r_custom, 45
radiation
  damping and excitation, **381**
radiation damping and excitation, *see* synchrotron radiation
radiation_damping_on, 262
radiation_fluctuations_on, 262

radiation_zero_average, 262
ramper, 507
ran, 37, 51, 56
ran_gauss, 37, 51, 56
ran_seed, 51, 255, 257
rbend, **70**, 167, 168, 179, 197, 218, 223, 226, 229, 236, 327, 332, 347, 495, 519
  coordinate transformation, 336
rcollimator, **77**, 179, 181, 223, 226, 229
redef, 38
ref, 243
ref_origin, 243
ref_tilt, 74, 334, 458, 463
ref_time, 259
ref_wave_length, 83, 172
reference energy, 198, **339**, 514, 548
reference orbit, 198, **326**
  construction, 327
  origin in global coordinates, 331
reference particle, **339**
reference time, **339**
reflection of elements, 236
rel_tol_adaptive_tracking, 231, 262
rel_tol_tracking, 262
relative time tracking, 421
remove_elements, 38
replacement list, *see* list
reserved names, 38
return, 38
return statement, **56**
reversed elements, 540
RF field map, 150
RF fields, **358**
rf fields, **355**
rf_bend, **147**, 223, 226
rf_frequency, 81, 119, 147, 149, 151
rfcavity, **149**, 209, 210, 223, 226, 229, 340, 421, 425, 444, 500
rho, 73, 151, 167, 332
rigid patch, 137
roll, 70, **174**
  coordinate transformation, 336
roll_tot, **178**
root, 38
root branch, 260, 327
root branches, 30
rp, 483
runge_kutta, 231, 232
  and field maps, 150
  and Taylor maps, 230
  tracking_method, 223
s-positions, 548
SAD, 292
sad, 218
sad_mult, **151**, 445
sample, **153**, 181
sbend, **70**, 167, 168, 179, 197, 218, 223, 226, 229, 236, 327, 332, 347, 427, 429, 495, 519
  coordinate transformation, 336
scale_multipoles, 198
secondary lattice file, 57
setenv statement, **53**
sextupole, **154**, 168, 223, 226, 229, 347, 425, 445
  tilt default, 175
sig_E, 141
sig_E2, 141
sig_vx, 141
sig_vy, 141
sig_x, 66, 67, 142, 167
sig_y, 66, 67, 142, 167
sig_z, 66, 142
sim_utils library, 480
sin, 51
slave, **509**
  ordering, 512
slave_status, 31
slice_lattice, 38
slice_slave, 31
slice_slave$, 510
sol_quad, **155**, 168, 223, 226, 229, 347, 445
  conversion to MAD, 292
  tilt default, 175
solenoid, **156**, 168, 223, 226, 229, 347, 437, 447
space_charge_common_struct, **268**
space_charge_mesh_size, 268
space_charge_method, 230
spatial_distribution, 142
species, 51
spin, **397**, 529
spin taylor map, 416
spin tracking
  list of routines, 612
  methods, 228
spin_fringe_on, 217, 228
spin_sokolov_ternov_flipping_on, 262
spin_tracking_method, 221, 228
spin_tracking_on, 262
spin_x, 258
spin_y, 258
spin_z, 258
sprint spin tracking, 448
sqrt, 51

sr_wakes_on, 262
ss:coher, 454
start_branch_at, 38
start_edge, 113
statement order, **52**
strings, **46**
structure, **261**
structures, 484
super_lord, 31, 512
super_lord$, 509
super_slave, 31, 512
super_slave$, 510
superimpose, 37, 38
  example, 488
superposition, 241
  reference energy, 548
surface, **181**
surface grid, 187
switches, **46**
symmetric_edge, 113
symp_lie_Bmad, 539
  tracking_method, 223
symp_lie_bmad, 231, 232
  and field maps, 150
  and Taylor maps, 230
  mat6_calc_method, 225
symp_lie_PTC, 539
symp_lie_ptc, 231
  and Taylor maps, 230
  mat6_calc_method, 226
  spin_tracking_method, 228
  tracking_method, 223
symplectic
  conjugate, 390
  integration, 231
symplectic integration, 415, 418, 425
symplectification, 417
symplectify, 233, **234**
sympliectify, 225
synchrotron radiation
  calculating, 538
  integrals, 384

t_offset, 136
tags for Lines and Lists, 56
tags for lines and lists, **239**
tan, 51
Tao, 25
taylor, **157**, 223, 226, 229, 231, 539
  and Taylor maps, 230
  deallocating, 539
  mat6_calc_method, 226
  tracking_method, 223
taylor Map, 538
taylor map, **415**
  feed-down, 418
  list of routines, 614
  reference coordinates, 416
  structure in ele_struct, 499
  with digested files, 524
taylor_map_includes_offsets, 225, **234**
taylor_order, 255, 257, 262
theta_origin, 96, 108
theta_position, 260
thick_multipole, **161**
thickness, 83
tilt, 67, 100, 110, 116, 122, 133, 136, 144, **174**, 181, 219, 332, 458, 498
  coordinate transformation, 336
tilt_calib, 219
tilt_corr, 83, 458, 462
tilt_err_tot, **178**
tilt_tot, **178**, 498
time
  phase space coordinates, **341**
time_runge_kutta, 231
  tracking_method, 223
title, 38
title statement, **54**
tn, *see* multipole, tn
to_element, 106
Touschek Scattering, 374
track1_postprocess, 546
track1_preprocess, 546
tracking, 529
  apertures, 537
  backwards, 539
  list of routines, 616
  Macroparticles, 371
  mat6_calc_method, 226
  partial, 536
  particle distributions, 540
  spin, 541
  spin_tracking_method, 228
tracking methods, **221**
tracking_method, 86, 221
transfer map
  in ele_struct, 499
  mat6_calc_method, *see* mat6_calc_method
  Taylor map, *see* Taylor map
translate_patch_drift_time, 267
tune

calculation, 526
setting, 527
twiss
list of routines, 620, 621
twiss parameters, 526
calculation, 526
twiss_struct, 525
type, **168**

undefined_vars_evaluate_to_zero, **55**
undulator, **162**
units
with MAD, 48
upstream element end, 327
use, 38
use statement, 35, **37**, **239**
use_local_lat_file, 38
use_local_lat_file_statement, **55**

v1_unitcell, 131
v2_unitcell, 131
v_displace, 118
v_unitcell, 83
val1, ..., Val12, 86
variables, *see* lattice file format, variables
velocity_distribution, 142
vkick, 93, 117, 118, 167, 168, **179**
vkicker, **117**, 179, 223, 226, 229, 347, 439
voltage, 81, 91, 93, 119, 147, 149, 151, 167
voltage_err, 91

wake_lr_struct, 501
wake_sr_mode_struct, 501
wakefields, **367**
in ele_struct, 500
list of routines, 621
long-range, 368
short-range, 367
wakes
short-range, 211
wall, 139, 189
wall_transition, **181**
wiggler, **162**, 167, 197, 223, 226, 229, 326, 425, 495, 519
conversion to MAD, 292
field maps, 501
tracking, 450
write_digested, 38
write_digested statement, 60

x, 258
x0, 127
x1, 127
x1_limit, **179**, 497
x2_limit, **179**, 497
x_axis, 214
x_gain_calib, 219
x_gain_err, 219
x_half_length, 140
x_kick, 111
x_limit, **179**, 497
x_offset, 67, 100, 110, 116, 122, 136, **174**, **179**, 181, 219, 326, 457, 498
x_offset_calib, 219
x_offset_mult, 151
x_offset_tot, **178**, 498
x_origin, 96, 108
x_pitch, 63, 67, 100, 110, 116, 136, **174**, 181, 326, 457, 498
x_pitch_mult, 151
x_pitch_tot, **178**, 498
x_position, 260
xraylib
library, 481
xy_disp_struct, 525

y, 258
y0, 127
y1, 127
y1_limit, **179**
y2_limit, **179**
y_axis, 214
y_gain_calib, 219
y_gain_err, 219
y_half_length, 140
y_kick, 111
y_limit, **179**
y_offset, 67, 100, 110, 116, 122, 136, **174**, 181, 219, 326, 457, 498
y_offset_calib, 219
y_offset_mult, 151
y_offset_tot, **178**, 498
y_origin, 96, 108
y_pitch, 63, 67, 100, 110, 116, 136, **174**, 181, 326, 457, 498
y_pitch_mult, 151
y_pitch_tot, **178**, 498
y_position, 260

z, 258
z0, 127
z1, 127
z_kick, 111

z_offset, 67, 100, 113, 136, **174**
z_offset_tot, **178**, 457
z_origin, 96, 108
z_position, 260