

SPIN POLARIZATION, MODELING, AND BEAM
CONTROL IN HADRON ACCELERATORS: FOR
RHIC AND THE EIC

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Eiad Hamwi

December 2025

© 2025 Eiad Hamwi
ALL RIGHTS RESERVED

SPIN POLARIZATION, MODELING, AND BEAM CONTROL IN HADRON
ACCELERATORS: FOR RHIC AND THE EIC

Eiad Hamwi, Ph.D.

Cornell University 2025

This dissertation develops theory, models, and operational methods for preserving and optimizing spin polarization in hadron accelerators spanning the RHIC injector chain and the Electron–Ion Collider (EIC) Hadron Storage Ring (HSR). Unlike electrons, hadrons do not self-polarize at high energy; source polarization must survive long acceleration ramps that traverse hundreds of depolarizing resonances and multiple machines. The work integrates high-fidelity spin–orbit modeling, snake lattice design, and data-driven beam tuning into a reproducible framework.

Key contributions are: (1) Application of a Maxwellian field representation and symplectic tracking algorithm for AGS combined-function dipoles and partial Siberian snakes, removing non-symplectic tracking artifacts and enabling stable million-turn spin-orbit simulations, yielding realistic polarization transmission over most of the AGS cycle. (2) Conception, and optimization of the Doubly Lee–Courant (DLC) six-snake scheme for the HSR, achieving near-perfect polarization preservation for (protons and) helions and establishing a resilient baseline configuration. (3) First experimental deployment of Bayesian optimization for injection tuning in the BtA transfer line (Booster → AGS), improving transmission from $\approx 65\% \rightarrow 90+\%$ with no operator intervention. (4) Experimental validation of resonant slow extraction modeling at NSRL, confirming normal-form predictions and achieving one to two orders of

magnitude dispersion reduction.

Together these results unify snake design, orbit control, and machine-learning operations, providing a practical path to EIC polarization goals and a computational foundation for future polarized hadron facilities and their automated optimization.

BIOGRAPHICAL SKETCH

Eiad Hamwi was born in Syria, but grew up in Al Ain, UAE. When he was 15, he moved to the US with his family to start college. After working on exoplanet atmospheric characterization during his undergraduate studies and completing an NSF REU at UWashington, he received his B.S. in Physics and Math from Missouri State University in 2017. He then joined Cornell University as a graduate student, where he has explored various topics ranging from experimental evolutionary biophysics to moduli spaces of Calabi-Yau compactifications, before discovering an exciting community in polarized beam dynamics and control for hadron accelerators.

To my parents, for instilling in me a deep curiosity and love for learning.

This is for you.

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to my advisor, Professor Georg Hoffstaetter, for his unwavering mentorship, guidance, and encouragement throughout my Ph.D. journey. His contagious enthusiasm for accelerator physics has deeply shaped my own passion for the field... although I still wonder whether he would pick accelerator physics over sailing! Thank you Georg for being a great role model to look up to.

I would like to thank Drs. David Sagan and Desmond Barber for their invaluable scientific support and mentorship of young researchers, both have been a cherished source of information to me. I am also thankful to my committee members, Drs. Matthias Liepe and Tomas Arias, for their openness and helpfulness. Thank you to Matthias for introducing me to Georg. Thank you to Drs. Andrea Giometto and Liam McAllister for their support during my early years at Cornell. I would like to sincerely thank my undergraduate advisor, Dr. David Cornelison for fostering my enthusiasm and encouraging me to pursue a graduate education in physics. Finally, I would like to thank my father, Farid Hamwi, for instilling in me a boundless curiosity about the natural world and a deep drive to pursue knowledge.

I am very grateful to my collaborators and mentors at BNL: Vincent Schoefer, Kiel Hock, Kevin Brown, Haixin Huang, Vadim Ptitsyn, Chuyu Liu, and Nick Tsoupas, for sharing their expertise and making this work possible. Thank you to Vincent for being a constant fountain of knowledge on all things AGS and RHIC.

To my group members and colleagues: Joe, Lucy, Ningdong, Jonathan, Matt, Trevor, Leve, Bhawin — thank you for the many hours of stimulating discussions, collaborations, and friendships.

To my friends outside of work: Jooheon, Brandon, Zepoor, Namitha, Shattique, Daniel, Ali, Blue, Chloe, Dany, Dave, Grant, Hannah, Jared, Kaylee, Kevin, Kieran, Maciej, Marianna, Maya, Michael, Rana, Shadeed, Shaza, Steven, Vada, Zina. I could not have gone through this journey without any of you. Thank you for all the fun times together.

To Chris and Erica, thank you both for being extremely supportive throughout this journey. You've both been like second parents to me.

To my family, I am eternally grateful for your unconditional love and support.

Mom and Dad, thank you for patiently listening to me go on random physics tangents and endlessly supporting me along the way.

Badri and Carla, you know how much I love you. We're all in this together, so lets kill it together!

Finally, to my partner: Isabel, thank you for your endless patience and encouragement. Your presence in my life has been a well of strength and joy beyond words.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vii
List of Tables	xi
List of Figures	xii
1 Introduction and Motivation	1
1.1 The Scientific Motivation for Polarized Beams	1
1.2 The Challenge of Spin Dynamics in Hadron Accelerators	2
1.3 The Role of Siberian Snakes and Resonance Control	3
1.4 Limitations of Existing Models and Instrumentation	5
1.5 Computational Demands	6
1.6 Broader Structure of the Dissertation	8
1.7 Tying the Big Picture	10
2 Mathematical Foundations	14
2.1 Equations of Motion	15
2.1.1 Physical Interpretation of the T-BMT Equation	16
2.1.2 Hamiltonian Formulation and Symplectic Structure	17
2.1.3 Spin-Orbit Coupling and Extended Phase Space	19
2.2 Normal Forms and the Structure of Invariant Motion	20
2.2.1 From the s -Dependent Hamiltonian to the One-Turn Map	20
2.2.2 Linearization of the One-Turn Map	22
2.2.3 Normal Form and the Interpolating Hamiltonian	23
2.2.4 Spin Normal Form	25
2.2.5 Resonant Terms	27
2.3 Spin Resonance Overview	28
2.3.1 Spin Tune and Resonance Conditions	28
2.3.2 Intrinsic Resonances	30
2.3.3 Imperfection Resonances	32
2.3.4 Higher-Order Resonances	33
2.4 Field Modeling and Map Construction	34
2.4.1 Maxwell's Equations and Field Constraints	35
2.4.2 Generalized Gradient Field Equations	36
2.4.3 Generalized Gradient Fitting Algorithm	37
2.4.4 Vector Potentials and Symplectic Maps	38
2.4.5 From Local Models to Global Tracking	40
2.5 The Big Picture	41

3	Modeling Magnetic Fields and Symplectic Tracking in the AGS	45
3.1	Field Modeling Challenges in the AGS	46
3.1.1	Combined-Function Dipole Magnets	47
3.1.2	Partial Siberian Snakes and Local Coupling	50
3.1.3	The Impact of Diagnostic Limitations	53
3.2	From Generalized Gradient Expansions to Symplectic Maps	55
3.2.1	Manifestations of Non-Symplectic Tracking	55
3.2.2	Maxwell Multipole Expansions	56
3.2.3	Construction of Symplectic Maps	57
3.2.4	Validation and Performance	60
3.3	Simulated Field Models	62
3.3.1	Field Modeling Workflow	63
3.3.2	From Field Maps to Vector Potentials	64
3.3.3	Symplectic Integration with Interpolated Fields	66
3.3.4	Validation Against Measurements	67
3.3.5	Future Directions	69
3.4	Validation and Lattice Matching to Machine Settings	70
3.5	The Big Picture	72
4	Spin Dynamics in RHIC and the Hadron Storage Ring	75
4.1	Early Work: RHIC Snake Re-Optimization	76
4.1.1	The Snake Damage Event and Operational Response	76
4.1.2	Simulation Studies and Key Findings	77
4.1.3	Physical Interpretation and Implications	78
4.2	The EIC Hadron Storage Ring: Design and Polarization Challenge	79
4.2.1	HSR Design Parameters and Snake Configuration	80
4.2.2	Symmetry Breaking and Resonance Driving	81
4.2.3	Baseline Tracking and the Need for Optimization	82
4.3	Method I: Optimization of Betatron Phase Advance	83
4.3.1	Theoretical Basis for Phase Advance Optimization	84
4.3.2	Baseline Behavior and Tune Scans	85
4.3.3	Optimization Results	86
4.3.4	Practical Implementation and Limitations	89
4.4	Method II: Optimization of Snake Rotation Axes	92
4.4.1	Exploration of the General Snake Axis Space	92
4.4.2	The Statistical Advantage of Lee-Courant Schemes	93
4.4.3	The Doubly Lee-Courant (DLC) Scheme	94
4.4.4	Performance Results and the DLC Discovery	96
4.5	Comparative Analysis and Path Forward	100
4.5.1	Implications for EIC Commissioning	101
4.6	Discussion: Closed-Orbit and Sextupole Feed-Down	102
4.6.1	Sextupole Feed-Down	103
4.6.2	Operational Guidance	104
4.6.3	Chapter Summary	104

5	Machine Learning for Injection Optimization	106
5.1	Injector Chain Instrumentation	107
5.1.1	Beam Position Monitors	108
5.1.2	Emittance and Profile Measurements	109
5.1.3	Closed-Orbit Distortions	110
5.1.4	Implications for Optimization Strategies	111
5.2	Bayesian Optimization for Booster-to-AGS Injection	113
5.2.1	Bayesian Optimization Framework	114
5.2.2	Acquisition Functions and Exploration-Exploitation	116
5.2.3	Kernel Selection and Dimensionality Reduction	118
5.2.4	Experimental Results	120
5.2.5	Limitations and Insights	125
5.3	Reinforcement Learning with Multiwire Profiles	127
5.3.1	Markov Decision Process Formulation	127
5.3.2	Policy Representation and Deep Reinforcement Learning	130
5.3.3	Sim-to-Real Transfer Challenges	131
5.4	Summary	133
6	Slow Extraction to the NSRL	134
6.1	Extraction Mechanics	136
6.1.1	First-Order Resonant Normal Form à la Kobayashi	136
6.1.2	Lattice Elements and Extraction Channel	142
6.1.3	The Extraction Cycle and Spill Structure	144
6.2	Simulation and Validation	146
6.2.1	Booster Lattice Model	146
6.2.2	Model Simulation and Normal Form Validation	148
6.2.3	Dispersion Suppression in the Extracted Beam	149
6.2.4	Experimental Validation of Dispersion Suppression	151
6.3	Spill Quality Control and Future Directions	155
6.3.1	Power Supply Ripple and Mitigation Strategies	156
6.3.2	Machine Learning and Future Control Strategies	156
6.3.3	Summary and Connections to Dissertation Themes	157
7	Conclusions and Future Directions	158
7.1	Key Contributions	158
7.1.1	Mathematical and Computational Infrastructure	158
7.1.2	Field Realism and Symplectic Tracking in the AGS	159
7.1.3	Spin Dynamics from RHIC Operations to EIC Design	159
7.1.4	Machine Learning for Injector Operations	160
7.1.5	Extending the Toolkit to Controlled Extraction	160
7.2	Future Directions	161
7.2.1	Toward an End-to-End EIC Digital Twin	161
7.2.2	Implementing the Doubly Lee–Courant Baseline	161
7.2.3	Autonomous Data-Driven Control	162

7.2.4	Collective Effects	162
7.3	Closing Remarks	163
	Bibliography	164

LIST OF TABLES

3.1	AGS Snake Strength Calculation Results	62
3.2	AGS Snakes Tracking Times (seconds)	62
4.1	Species-Dependent Linear Resonance Spectrum in the HSR . . .	81
4.2	Snake Axis Schemes for a 2N-Snake Ring	97

LIST OF FIGURES

2.1	First-order resonance strength spectrum of RHIC ε_κ in the absence of snakes.	31
2.2	Jumps in the amplitude-dependent spin tune $\nu(\vec{J})$ across 3 rd -order resonances; jump size = $2\varepsilon_\kappa$	34
2.3	On-axis wiggler field agreement: grid value vs. GG fits using Elegant vs. Bmad algorithms (surface vs volume fits). In blue/cyan are the absolute errors of the fits (left axis), and in red is the actual value of the grid field points (right axis).	37
2.4	Off-axis wiggler field agreement at max grid radius: grid value vs. GG fits (surface vs volume fits). In blue/cyan are the absolute errors of the fits (left axis), and in red is the actual value of the grid field points (right axis).	39
3.1	Simulated field of an AGS Type-B main magnet at 2650 A, calculated via 2D FEM with 1–5 mm grid spacing.	47
3.2	Locations of two partial snakes in the AGS ring.	50
3.3	Simulated orbit through the warm snake at 2 GeV, from [1].	52
3.4	Stability of phase space motion with grid-table integration (red) vs symplectic tracking of cylindrical expansion (green)	61
3.5	Optics and orbit at 2.5 GeV showing closed-orbit bumps around both snakes after local compensation.	71
3.6	Polarization transmission through the AGS cycle: 97% achieved with symplectic snake models and optics compensation.	72
4.2	Schematic comparison of the intrinsic resonance spectrum of RHIC (left) versus the HSR (right). The HSR's broken symmetry from the complex interaction region creates both stronger systematic resonances and a higher non-systematic resonance floor, making polarization preservation significantly more challenging.	82
4.3	Baseline spin tracking of uncooled, round, helion beams of $3\ \mu\text{m}$ emittance in the injection HSR lattice with $\pm 15^\circ$ snakes. Each colored trace represents the vertical spin component of each particle in the bunch. The ensemble polarization is lost rapidly as the beam is accelerated through strongly resonant energies around $G\gamma \approx 705$ and $G\gamma \approx 745$	84
4.4	Tune scan of the limiting polarization P_{lim} for the baseline HSR lattice at $G\gamma \approx 704$. The dark bands indicate strong higher-order resonance doublets whose splitting reveals the magnitude of the ADST spread. The nominal working tune is marked by the crosshairs. The wide doublet splitting indicates poor polarization preservation with the baseline configuration.	86

4.5	Tune scan for the HSR lattice after optimizing the inter-snake betatron phase advances at $G\gamma \approx 704$. The resonance doublets are significantly narrowed compared to the baseline (Fig. 4.4), creating a much larger area of high polarization (bright yellow regions). This demonstrates the effectiveness of phase advance optimization in suppressing ADST spread.	87
4.6	Spin tracking with an energy-dependent interpolation of optimized phase advances. Each curve shows the vertical spin component of individual particles. Polarization is successfully preserved through the challenging energy range with $P_{\text{dyn}} > 99\%$, demonstrating that carefully orchestrated phase advance variations can navigate the dense resonance structure.	88
4.7	Spin tracking using a single, constant set of optimized phase advances, found via multi-objective optimization targeting multiple energies simultaneously. This provides a robust solution without requiring dynamic lattice changes during the ramp, achieving $P_{\text{dyn}} > 96\%$ while being operationally simpler than the dynamic solution.	89
4.8	Final tracking result with a single, optimized set of quadrupole currents in the fully matched lattice with realistic optics constraints. The solution achieves $P_{\text{dyn}} > 97\%$ while being implementable with actual machine hardware, demonstrating that Method I provides a viable path to high polarization transmission.	90
4.9	Tune scan after optimization using realistic quadrupole matching with actual IR optics constraints. The vertical resonance doublets are narrowed at the target energy, though the tune space is more complex than with idealized phase trombones. A safe working point (marked) can be identified that provides high polarization transmission.	91
4.10	Histogram of $\min(P_{\text{lim}})$ over the energy range $G\gamma \in (700, 800)$ for a random sampling of general snake axis configurations (blue) versus the subset satisfying the Lee-Courant conditions (orange). The L-C schemes are statistically far more likely to have higher equilibrium polarization, demonstrating the power of symmetry-based design principles.	94
4.11	Schematic of the Doubly Lee-Courant (DLC) snake configuration. Snakes are shown as diamond shapes around the ring. Complete segments (each enforcing a π spin phase advance) are indicated by colored dashed lines, with each segment spanning two consecutive snakes. The defining feature of DLC is that each snake simultaneously participates in two overlapping segments (note how each diamond is part of two different color schemes), providing maximal local spin cancellation.	96

4.12	Histogram of $\min(P_{\text{lim}})$ for general snake axis configurations (blue), Lee-Courant schemes (orange), and Doubly Lee-Courant schemes (green). The DLC schemes show statistically superior equilibrium polarization, with the vast majority achieving $P_{\text{lim}} > 0.95$. This demonstrates the power of maximal local spin symmetry.	97
4.13	Histogram of $\max(\Delta\nu)$ (maximum ADST spread) for general configurations (blue), Lee-Courant schemes (orange), and Doubly Lee-Courant schemes (green). Better performance is indicated by smaller ADST spread. DLC schemes achieve a factor of 4-8 reduction compared to general configurations, demonstrating superior resonance suppression.	98
4.14	Spin tracking result for the standard DLC scheme with $\pm 45^\circ$ snake axes. The configuration shows significantly improved polarization preservation compared to the baseline (Figure 4.3), though some polarization loss still occurs around $G\gamma \approx 745$	99
4.15	Spin tracking result for the optimal DLC scheme with maximal P_{dyn} found by optimizing the single free parameter. The configuration provides near-perfect polarization transmission throughout the energy range, with all particle spins remaining tightly aligned. This represents the best achievable performance through snake axis optimization alone.	100
5.1	Schematic diagram of the BtA line and its connections with the Booster and the AGS.	114
5.2	BO results of the ratio of AGS injection intensity to Booster extraction intensity after spoiling the initial settings, while varying two horizontal and two vertical steering magnets.	120
5.3	Beam brightness is a secondary target not included in optimization, but shows strong recovery following spoilage of the settings.	120
5.4	BO results of the ratio of AGS extraction intensity to Booster extraction intensity after spoiling the initial settings, while varying two horizontal and two vertical steering magnets.	121
5.5	BO results while varying three focusing quadrupoles and two horizontal steering magnets to compensate unintentional quadrupole-induced steering (horizontal optics matching). . . .	122
5.6	Real-world BO results of the ratio of AGS injection intensity to Booster extraction intensity starting from an untuned machine (about 65% transmission), while varying two horizontal and two vertical steering magnets. Convergence in approximately 15 minutes.	124

6.1	Comparison of normal form separatrix prediction for a particle with $\delta = 0.35\%$ and $J = 8.45\mu\text{m}$ with symplectic tracking through the exact s -dependent lattice Hamiltonian shows remarkable agreement. Also shown is the extracted bunch distribution. The lattice parameters are $\beta_x = 8.02\text{ m}$, $\alpha_x = -1.31$, $Q_x = 4.364$, $\xi_x = -7.6$, and $S_v = 9.27 + 8.63i$. This validates both the normal form analysis and the constant-optics extraction mechanism that enables dispersion suppression.	142
6.2	Extracted particles' horizontal position vs momentum deviation relative to bunch average at the D3 thin magnetic septum shows almost zero correlation, indicating dramatically suppressed beam dispersion. The near-vertical scatter demonstrates that particles of all momenta exit along common unstable trajectories defined by betatron amplitude, not momentum. This validates the normal form prediction that the constant-optics extraction scheme produces dispersion-free beams despite non-zero periodic dispersion at the septum.	150
6.3	D6 flag showing two overlaid beam spots corresponding to radial offsets of $dR = -13.5\text{ mm}$ (blue) and $dR = -6.5\text{ mm}$ (grey). The red "L" and "R" lines are 1.11" apart, providing calibrated position reference. The measured horizontal displacement between the two spots is $\Delta x \approx 0.7 \pm 0.5\text{ mm}$ despite a momentum difference of $\Delta\delta = 0.93\%$, demonstrating dramatic dispersion suppression compared to the periodic value of $\eta_x = 2.93\text{ m}$ at this location.	154

CHAPTER 1

INTRODUCTION AND MOTIVATION

The pursuit of understanding the fundamental structure of matter has driven the development of increasingly sophisticated particle accelerators and experimental techniques. Among the most powerful tools in nuclear and particle physics are polarized hadron beams, which provide unique access to spin-dependent phenomena that remain hidden in experiments with unpolarized particles. This dissertation addresses the complex challenge of preserving and controlling spin polarization in hadron accelerators, from injection through acceleration to extraction, with particular emphasis on current facilities like RHIC and the AGS, and future machines like the HSR for the EIC.

The central motivation driving this work is the recognition that polarized beams are not merely advantageous for nuclear physics research. They are essential for unlocking fundamental questions about the spin structure of hadrons and nuclei. Yet maintaining polarization through the complex accelerator chains required to reach collision energies presents formidable technical challenges that demand both deep theoretical understanding and sophisticated computational tools. This dissertation contributes to both domains, developing the physical models and computational infrastructure needed to achieve the polarization requirements of next-generation facilities.

1.1 The Scientific Motivation for Polarized Beams

Polarized hadron beams are essential to the EIC physics program: they enable flavor-separated, spin-dependent measurements and precision symmetry tests

that unpolarized beams cannot provide. Meeting these goals requires high-luminosity collisions of polarized protons and light nuclei (notably helium-3 as an effective neutron proxy), with stringent polarization targets at the interaction points.

Delivering such beams is a system-level challenge. Polarization must be preserved from source through the Booster, AGS and RHIC or ultimately the HSR, while traversing hundreds of depolarizing resonances. This dissertation frames the problem around design choices (snake configurations and optics), operational control (orbits, tunes, and injection), and computational fidelity (symplectic models and scalable tracking), setting the stage for the results synthesized in later chapters.

1.2 The Challenge of Spin Dynamics in Hadron Accelerators

Spin motion is governed by the Thomas–Bargmann–Michel–Telegdi (T-BMT) equation, described in Chapter 2. Here we summarize the resonance landscape that drives polarization loss during acceleration.

Intrinsic resonances arise from the periodic focusing structure of the lattice and occur when the spin precession couples to betatron or synchrotron motion (e.g., $G\gamma = kP \pm Q_y$, $\forall k \in \mathbb{Z}$). They are densely spaced for species with large G and during long ramps.

Imperfection resonances stem from closed-orbit distortions and field errors (condition $G\gamma = k \in \mathbb{Z}$). Even when less frequent, they can be strong if orbit control is poor.

Maintaining polarization requires reducing resonance strengths (via optics and orbit control), avoiding dangerous working points, and, critically, altering the spin tune structure with Siberian snakes so that large classes of resonances are suppressed by design. The practical methods this dissertation develops combine these levers into robust operating strategies.

1.3 The Role of Siberian Snakes and Resonance Control

The concept of Siberian snakes, first proposed by Derbenev and Kondratenko [2, 3], represents a paradigm shift in hadron spin dynamics. A Siberian snake is a specialized insertion device that rotates the spin vector by 180° while leaving the orbital motion largely unaffected. When properly implemented, snakes fundamentally alter the spin dynamics by fixing the spin tune at exactly $1/2$, independent of energy.

The impact of this seemingly simple modification is profound. With the spin tune locked at $1/2$, all intrinsic resonances are automatically avoided with appropriate choice of betatron tunes away from $1/2$. This self-protecting mechanism allows hadron beams to be accelerated through energy ranges that would otherwise be impossible to traverse with significant polarization.

RHIC, the world's first and most successful polarized hadron collider [4], employs two full Siberian snakes in each ring, with diagonally horizontal spin rotation axes at $\pm 45^\circ$ relative to the longitudinal direction. This configuration provides almost complete compensation for both intrinsic and imperfection resonances under ideal conditions, enabling proton polarization of 50-60% at top energy (275 GeV per beam). The success of RHIC has validated the snake con-

cept and demonstrated the feasibility of high-energy polarized hadron collisions.

However, the AGS, which serves as RHIC's injector, operates with only two partial Siberian snakes due to space and field constraints. The warm and cold snakes provides $\sim 11^\circ$ and $\sim 22^\circ$ spin rotation around the longitudinal axis [5,6], respectively, rather than the ideal 180° , leaving the machine vulnerable to many residual resonances. This compromise, necessitated by the retrofit nature of the AGS snake installation, requires sophisticated resonance analysis and correction techniques to maintain acceptable polarization transmission.

The planned HSR for the EIC will employ six full Siberian snakes [7], providing even greater resonance protection than RHIC. However, the HSR must accommodate not only protons but also light nuclei with different gyromagnetic ratios. Helium-3, with its large anomalous magnetic moment, has $G \approx -4.18$ and presents particular challenges due to its high resonance density and the resulting strong depolarization of spin motion.

Recent theoretical work has explored advanced snake configurations beyond the traditional alternating-axis approach. The "Lee-Courant" scheme arranges consecutive pairs of snakes to produce local 180° rotations [8], potentially offering advantages for species with high resonance density. This dissertation investigates a novel extension, the "Doubly Lee-Courant (DLC)" configuration [9], where overlapping pairs of snakes simultaneously produce localized full spin flips, providing enhanced protection for challenging species like helium-3.

The optimization of snake configurations requires sophisticated computational tools that can model millions of turns of beam evolution while account-

ing for realistic magnetic field maps, orbit distortions, and statistical variations in machine parameters. This computational challenge has driven much of the methodological development described in subsequent chapters.

Building on classical Lee–Courant ideas, Chapter 4 reveals a new arrangement of six full snakes for the HSR, the DLC scheme, that produces overlapping local 180° spin rotations across the ring. This structure suppresses both intrinsic and imperfection resonances over the full ramp, yielding near-perfect polarization transmission for realistic beam emittances prior to the introduction of errors. The DLC scheme emerges as a practical, commissioning-friendly baseline for the EIC and is a flagship result of this dissertation.

By embedding strong symmetry directly in the spin transport, DLC turns a hard online tuning problem into an offline design problem. Furthermore, hybrid operating modes are outlined that trade global optimality for localized tolerance to specific resonances which could be susceptible to errors.

1.4 Limitations of Existing Models and Instrumentation

Significant gaps remain between idealized models and operational reality in polarized hadron machines, and they directly affect the reliability of polarization predictions.

First, **field realism** is challenging in lattices with combined-function dipoles and partial snakes: thin-lens or slice-based generalized-gradient models can break symplectic structure and create artificial resonances. Second, **instrumentation** in injector complexes is sparse (low-frequency BPM coverage and limited

polarimetry), complicating validation and closing the loop from simulation to operation. Third, **operational constraints** limit fast measurement and actuation of key elements, so robust, largely passive strategies are preferred.

Chapter 3 addresses these limitations by developing Maxwell-consistent field representations and symplectic tracking methods that remain stable over million-turn ramps and match real machine settings. Applied to the AGS, this symplectic framework authentically simulates polarization transmission through the full acceleration cycle, resolving long-standing discrepancies at low energies and enabling practical optimization studies.

1.5 Computational Demands

The computational requirements for hadron spin dynamics modeling are fundamentally different from those encountered in electron accelerator physics, reflecting the absence of radiation damping and the resulting need for long-term tracking studies. While electron beam polarization can often be evaluated after a few thousand turns once equilibrium is established, hadron spin dynamics require tracking over millions of turns to accurately assess polarization survival during energy ramps.

Consider a typical AGS acceleration cycle lasting 1.4 seconds with a revolution frequency of approximately 370 kHz. The beam makes roughly 500,000 revolutions during acceleration, crossing hundreds of resonances with varying strengths and characteristics. To obtain statistically meaningful results, ensembles of hundreds or thousands of particles must be tracked, leading to computational requirements of 10^8 to 10^9 individual turn-by-turn calculations for a single

acceleration scenario.

The situation becomes even more demanding when optimization studies are considered. Investigating the performance of different snake configurations, correction strategies, or injection parameters requires repeated ensemble simulations across multi-dimensional parameter spaces. Machine learning approaches, which show great promise for complex optimization problems, may require thousands of individual simulations during training, multiplying the computational burden by additional orders of magnitude.

These requirements have driven the adoption of modern high-performance computing approaches within the accelerator physics community. This dissertation primarily leverages the *Bmad* ecosystem [10], a Fortran toolkit for accelerator physics simulation and analysis purposes. Nevertheless, much work has been done on developing a next-generation ecosystem, *SciBmad*, written using the Julia programming language which is specifically designed for scientific computing. Julia's combination of high-level expressiveness and near-C performance makes it particularly well-suited for the large-scale simulations required in hadron spin dynamics. Much of the work done in *Bmad* has been reproduced with *SciBmad* for confirmation and further analysis.

Key advantages of the Julia framework include:

- **Performance:** Julia's just-in-time compilation allows high-level code to achieve performance comparable to optimized C or Fortran, critical for million-turn tracking studies.
- **Parallelization:** Native support for multi-threading, distributed computing, and GPU acceleration enables efficient scaling across modern compu-

tational resources.

- **Symplectic integration:** On top of the in-house developed custom symplectic-unitary integrators for spin-orbit motion in `BeamTracking.jl`, open-source packages like `DifferentialEquations.jl` provide a preponderance similarly sophisticated integrators that preserve the Hamiltonian structure of single-particle dynamics.
- **Automatic differentiation:** Built-in support for automatic differentiation enables efficient computation of gradients needed for optimization and perturbation theory analysis through `GTPSA.jl` and `NonlinearNormalForm.jl`.
- **Ecosystem:** Growing libraries for accelerator physics (`Beamlines.jl`, `BeamTracking.jl`), nonlinear dynamics (`NonlinearNormalForm.jl`), and machine learning (`Flux.jl`), as well as a simple Python interface (`PythonCall.jl/juliacall`), provide a comprehensive toolkit.

1.6 Broader Structure of the Dissertation

This dissertation is organized to provide a comprehensive treatment of hadron spin dynamics, progressing from fundamental theory through practical applications to computational implementation. The structure reflects the multi-faceted nature of the problem, which requires integration of electromagnetic theory, nonlinear dynamics, and high-performance computing.

Chapter 2: Mathematical Foundations establishes the theoretical framework underlying all subsequent work. Beginning with the Thomas-BMT equation and its integration into Hamiltonian mechanics, the chapter develops the

symplectic integration methods essential for long-term tracking studies. The treatment of resonance theory and normal form analysis provides the mathematical tools needed to understand and predict depolarization mechanisms. The chapter concludes with field expansion techniques that enable Maxwell-consistent modeling of complex magnet geometries.

Chapter 3: Field Realism in the Alternating Gradient Synchrotron addresses the critical challenge of accurately modeling electromagnetic fields in machines with Siberian snakes and combined-function magnets. The chapter demonstrates the failure of traditional slice-by-slice modeling approaches and develops alternative methods based on cylindrical and Cartesian multipole expansions. These techniques ensure symplectic integration while maintaining fidelity to Maxwell's equations, resolving long-standing discrepancies between simulations and experimental observations.

Chapter 4: Spin Dynamics in RHIC and the Hadron Storage Ring explores the performance of different Siberian snake configurations in current and future facilities. Initial investigations of RHIC operations during broken-snake scenarios provide insights into the relative importance of snake optimization versus orbit control. The bulk of the chapter focuses on the planned HSR, investigating both conventional six-snake configurations and novel "Doubly Lee-Courant (DLC)" arrangements that provide enhanced protection for high-anomalous-moment species like helium-3.

Chapter 5: Machine Learning for Injection Optimization applies modern optimization techniques to address the diagnostic limitations of injector systems. Bayesian optimization methods are first applied to improve transmission efficiency in Booster-to-AGS injections. The approach is then conceptually ex-

tended using reinforcement learning frameworks that can incorporate partial beam loss information from wire scanners and other diagnostics. The ultimate goal is to integrate injection optimization directly with polarization preservation objectives.

Chapter 6: Slow Extraction to the NASA Space Radiation Laboratory provides a detailed case study of beam extraction physics, demonstrating the application of the computational tools developed for injection and acceleration to the equally challenging problem of controlled beam delivery. The study models third-integer resonance extraction from the Booster to the NSRL line, including the effects of electrostatic and magnetic septa and beam scattering in extraction foils. Normal form analysis predictions are validated through detailed particle tracking across extraction separatrices.

Chapter 7: Conclusions and Future Directions synthesizes the results from all previous chapters and outlines the implications for EIC commissioning and operations. The chapter identifies key areas for future research, including advanced correction algorithms, improved diagnostic techniques, and next-generation computational methods for multi-physics accelerator modeling.

1.7 Tying the Big Picture

This dissertation represents more than a collection of individual studies in accelerator physics. It addresses the fundamental challenge of maintaining spin polarization across the entire hadron accelerator chain at the upcoming EIC, from the Booster to collision energy. The work recognizes that this challenge cannot be solved by focusing on any single aspect in isolation, but requires an

integrated approach that combines theoretical understanding, practical engineering solutions, and sophisticated computational tools.

The unifying theme throughout this work is the development of both the physics understanding and computational infrastructure needed to achieve the demanding polarization requirements of next-generation facilities like the EIC. This requires advancing the state of the art in multiple domains simultaneously:

Fundamental theory provides the foundation for understanding complex spin dynamics phenomena, from the basic T-BMT equation through advanced resonance analysis and normal form theory. However, theory alone is insufficient without accurate models of real accelerator components and their electromagnetic fields.

Realistic modeling bridges the gap between idealized theory and experimental reality by developing field descriptions that respect Maxwell's equations while remaining computationally tractable. The symplectic integration methods developed ensure that long-term tracking studies preserve the underlying Hamiltonian structure of the dynamics.

Correction strategies provide practical tools for mitigating residual resonances and optimizing machine performance. The work develops both passive corrections (snake configuration, lattice design) and active corrections (injection tuning) that work together to maximize polarization survival.

Computational infrastructure enables the large-scale simulations required to confidently predict machine performance and optimize design parameters. The Bmad ecosystem, written in Fortran90, as well as the newly developed Julia-based SciBmad toolkit provide both the performance needed for million-turn

tracking and the flexibility needed for complex optimization studies.

Machine learning integration represents a new frontier in accelerator control, offering the potential to automatically optimize complex systems that would be intractable for traditional approaches. The reinforcement learning frameworks developed provide a pathway toward intelligent control systems that can adapt to changing machine conditions.

The scientific return on this integrated approach extends far beyond polarized beam applications. The symplectic integration methods, field modeling techniques, and optimization frameworks developed for spin dynamics have broader applications throughout accelerator physics. Similarly, the computational infrastructure provides a foundation for next-generation accelerator modeling that can address increasingly complex multi-physics phenomena.

Perhaps most importantly, this work demonstrates that the challenges of future accelerator facilities can be addressed through the systematic application of modern computational and theoretical tools. The EIC represents the most ambitious polarized hadron facility ever attempted, with luminosity and polarization requirements that push the boundaries of current technology. The work presented in this dissertation provides confidence that these requirements can be met, laying the groundwork for a new era of discoveries in nuclear and particle physics.

The path from fundamental equations to operating accelerators is long and complex, requiring careful attention to countless technical details while maintaining sight of the ultimate scientific goals. This dissertation attempts to traverse that path systematically, developing the tools and understanding needed

to ensure that the remarkable scientific potential of polarized hadron beams can be fully realized in the coming decades.

CHAPTER 2

MATHEMATICAL FOUNDATIONS

The preservation of spin polarization in hadron accelerators rests upon a foundation of interconnected mathematical frameworks that span classical electrodynamics, Hamiltonian mechanics, and nonlinear dynamics. This chapter establishes the theoretical backbone for all subsequent work, developing the essential equations, methods, and computational approaches that enable accurate modeling and control of spin dynamics in complex accelerator systems.

Unlike electron storage rings where synchrotron radiation provides natural polarization through the Sokolov-Ternov effect, hadron machines must preserve polarization “by hand” through careful design and active control. This fundamental difference drives the need for sophisticated mathematical tools that can accurately predict spin evolution over millions of turns while accounting for the complex electromagnetic environments found in real accelerators. The mathematical foundations presented here provide both the physical understanding and computational infrastructure necessary to meet this challenge.

The chapter progresses through three interconnected layers of mathematical description. We begin with the fundamental equations of motion that govern both orbital and spin dynamics, emphasizing their coupled nature and the critical importance of symplectic integration for long-term tracking. We then develop resonance theory and normal form analysis, which provide the conceptual framework for understanding depolarization mechanisms and designing mitigation strategies. Finally, we address the modeling of electromagnetic fields, showing how Maxwell-consistent field descriptions can be constructed and integrated with symplectic tracking codes.

Throughout this development, we maintain focus on the practical requirements of hadron spin dynamics: the need for million-turn tracking accuracy, the importance of realistic field modeling, and the computational demands of optimization studies. These requirements have shaped both the choice of mathematical methods and their implementation, leading to approaches that balance theoretical rigor with computational efficiency.

2.1 Equations of Motion

The motion of charged particles in electromagnetic fields is governed by the Lorentz force equation, which in the presence of electric field \vec{E} and magnetic field \vec{B} takes the form:

$$\frac{d\vec{p}}{dt} = q(\vec{E} + \vec{v} \times \vec{B}) \quad (2.1)$$

where $\vec{p} = \gamma m \vec{v}$ is the relativistic momentum, q is the particle charge, m is the rest mass, and $\gamma = (1 - v^2/c^2)^{-1/2}$ is the Lorentz factor. This equation, combined with the relativistic energy-momentum relation, completely determines the orbital motion of spinless, charged particles in electromagnetic fields.

However, particles with intrinsic spin experience additional dynamics that couple to their orbital motion through the electromagnetic fields. The evolution of the spin vector \vec{S} is governed by the Thomas-Bargmann-Michel-Telegdi (T-BMT) equation [11–13]:

$$\frac{d\vec{S}}{dt} = \frac{q\vec{S}}{m\gamma} \times \left[(G\gamma + 1)\vec{B} - G(\gamma - 1)\vec{B}_{\parallel} - \left(G\gamma + \frac{\gamma}{\gamma + 1} \right) \left(\vec{\beta} \times \frac{\vec{E}}{c} \right) \right] \quad (2.2)$$

where $G = (g - 2)/2$ is the anomalous magnetic moment, $g \equiv \frac{2m}{qs}\mu$ is the gyromagnetic ratio, $\vec{\beta} = \vec{v}/c$ is the normalized velocity vector, and $\vec{B}_{\parallel} = (\vec{v} \cdot \vec{B}) \vec{v}/v^2$ is the projection of \vec{B} along \vec{v} . For protons, $G \approx 1.793$, while for helium-3, $G \approx -4.184$, reflecting the very different spin dynamics these species exhibit in the same electromagnetic environment [14].

2.1.1 Physical Interpretation of the T-BMT Equation

The T-BMT equation encodes several distinct physical effects that contribute to spin precession. The first term, proportional to $(G + 1/\gamma)\vec{B}$, derives from the precession due to the magnetic field in the particle's rest frame. In the co-moving rest frame of the beam, the anomalous magnetic moment contribution GB_0 is all that is left. Since the precession of momentum subtracted by the co-moving frame scales with B_0/γ , we find that the closed-orbit spin tune for a flat ring $\nu_0 = f_{s,0}/f_{\text{rev}} = G\gamma$ grows linearly with energy.

The second term, proportional to $(G - G/\gamma)\vec{B}_{\parallel}$, arises from the relativistic Lorentz boost of the magnetic field and counter-rotates spin around the parallel component of the magnetic field. At high energies this term combines with the first to asymptotically dampen the effects of longitudinal magnetic fields.

The third term, proportional to $(G + \frac{1}{\gamma+1})(\vec{\beta} \times \frac{\vec{E}}{c})$, represents the coupling between spin and electric fields. In most hadron accelerators, electric fields are primarily used for acceleration and are typically too weak for bending. However, due to the anomalous magnetic moment, electric fields generate an energy-independent precession of spin around some axis perpendicular to $\vec{\beta}$.

The key insight from the T-BMT equation is that spin dynamics cannot be separated from orbital dynamics. The magnetic and electric fields that appear in the spin equation are evaluated along the particle's trajectory, which is itself determined by the Lorentz force equation. This coupling means that orbit distortions, betatron oscillations, and synchrotron motion all influence spin evolution, creating the complex resonance structure that characterizes hadron spin dynamics.

On the other hand, orbital dynamics *can* be and usually are separated from spin dynamics due to the extremely weak nature of Stern-Gerlach effects in most accelerator scenarios.

2.1.2 Hamiltonian Formulation and Symplectic Structure

For long-term tracking studies spanning millions of turns, it is essential to preserve the underlying Hamiltonian structure of the particle dynamics. The Hamiltonian formulation provides a natural framework for this preservation while enabling the use of symplectic integration methods tailored to conservative systems.

In an electromagnetic field with scalar potential ϕ and vector potential \vec{A} , the relativistic Hamiltonian of a charged particle is

$$H = c \sqrt{m^2 c^2 + (\vec{P} - q\vec{A})^2} + q\Phi, \quad (2.3)$$

where $\vec{P} = \gamma m \dot{\vec{r}} + q\vec{A}$ is the canonical momentum. To describe motion relative to a curved reference trajectory $\vec{R}(s)$ of curvature $\kappa = 1/\rho$, we introduce curvilinear coordinates (x, y, s) such that $\vec{r} = \vec{R}(s) + x\hat{e}_x + y\hat{e}_y$, with $h = 1 + \kappa x$ the metric factor.

The corresponding canonical transformation

$$P_x = \hat{e}_x \cdot \vec{P}, \quad P_y = \hat{e}_y \cdot \vec{P}, \quad P_s = h \hat{e}_s \cdot \vec{P},$$

yields the Hamiltonian

$$H = c \sqrt{m^2 c^2 + (P_x - qA_x)^2 + (P_y - qA_y)^2 + \left(\frac{P_s}{h} - qA_s\right)^2} + q\Phi. \quad (2.4)$$

Choosing the path length s as the independent variable defines a new Hamiltonian $K = -P_s$, generating evolution via $d(\cdot)/ds = \{\cdot, K\}$. Introducing the normalized momentum deviation $\delta = (|\vec{p}| - p_0)/p_0$ and the longitudinal coordinate $z = -\beta c(t - t_0)$ gives the canonical pair (z, δ) , leading to the standard s -based form [15]

$$K = -h \sqrt{(1 + \delta)^2 - (p_x - a_x)^2 - (p_y - a_y)^2} - a_s + \phi + \frac{1}{\beta_0} \sqrt{(1 + \delta)^2 + \left(\frac{mc}{p_0}\right)^2}, \quad (2.5)$$

where $p_i = P_i/p_0$, $a_{x,y} = qA_{x,y}/p_0 c$, $a_s/h = qA_s/p_0 c$, and $\phi = q\Phi/p_0 c$.

This Hamiltonian preserves the symplectic structure in the six-dimensional phase space $(x, p_x, y, p_y, z, \delta)$, ensuring Liouville's theorem holds exactly under numerical integration. Numerical integration schemes that preserve this symplectic structure are called symplectic integrators.

Unlike general-purpose integration methods such as Runge-Kutta, symplectic integrators ensure that long-term tracking studies do not exhibit artificial growth or decay of phase space volumes. This property is essential for hadron accelerator simulations, where small systematic errors can accumulate over millions of turns and lead to completely incorrect predictions of beam behavior.

2.1.3 Spin-Orbit Coupling and Extended Phase Space

The coupling between spin and orbital motion through the T-BMT equation suggests extending the orbital phase space to include spin degrees of freedom. However, because the spin vector satisfies the constraint $|\vec{S}|^2 = \text{constant}$, it resides on a two-sphere rather than in a linear space, and thus does not admit a direct canonical representation. To retain an authentic symplectic structure, the spin must be expressed in terms of a canonical conjugate pair.

A convenient canonical parameterization is obtained by expressing the spin vector in spherical coordinates,

$$S_x = S \sin \theta \cos \phi, \quad (2.6)$$

$$S_y = S \sin \theta \sin \phi, \quad (2.7)$$

$$S_z = S \cos \theta, \quad (2.8)$$

and identifying the canonical pair as (ϕ, S_z) with Poisson bracket $\{\phi, S_z\} = 1$. The spin Hamiltonian,

$$H_s = \vec{\Omega} \cdot \vec{S},$$

then generates the standard precession equation $\dot{\vec{S}} = \vec{\Omega} \times \vec{S}$ through Hamilton's equations

$$\dot{\phi} = \frac{\partial H_s}{\partial S_z}, \quad \dot{S}_z = -\frac{\partial H_s}{\partial \phi}.$$

This canonical form embeds spin dynamics naturally within the overall symplectic framework of the orbital motion.

In relativistic accelerator applications, the feedback of spin dynamics on the orbital motion is extremely small (of order $e\hbar/mc$) and can be safely neglected. Consequently, the spin evolution can be treated as a passive "spectator" subsystem driven by the local electromagnetic fields along the reference trajectory.

For numerical implementation, this canonical formulation offers a compact and symplectic representation, though quaternion or rotation-matrix schemes remain advantageous for avoiding coordinate singularities near the poles.

Using a quaternion representation \mathbb{H} of spin rotations as an alternative also provides a natural framework for composing multiple rotations. A unit quaternion $q = q_0 + q_1\mathbf{i} + q_2\mathbf{j} + q_3\mathbf{k}$ with $|q|^2 = q_0^2 + q_1^2 + q_2^2 + q_3^2 = 1$ represents a rotation, and the composition of rotations corresponds to quaternion multiplication.

For practical implementations, the most robust approach often involves direct integration of the spin motion using steps with constant T-BMT-precession. When used in combination with geometric integration or Magnus integration schemes leveraging quadrature methods, such as Gauss-Legendre, high accuracy and performance can be simultaneously achieved [16].

2.2 Normal Forms and the Structure of Invariant Motion

2.2.1 From the s -Dependent Hamiltonian to the One-Turn Map

The equations of motion for a particle moving through a periodic focusing lattice are derived from the s -dependent Hamiltonian

$$K(s, \vec{z}) = -h \sqrt{(1 + \delta)^2 - (p_x - a_x)^2 - (p_y - a_y)^2} - a_s + \phi + \frac{1}{\beta_0} \sqrt{(1 + \delta)^2 + \left(\frac{mc}{p_0}\right)^2}, \quad (2.9)$$

which is periodic in the independent variable s with period L ,

$$K(s + L, \vec{z}) = K(s, \vec{z}).$$

The canonical equations

$$\frac{d\vec{z}}{ds} =: K(s) : \vec{z} = \{\vec{z}, K(s)\}$$

define a non-autonomous Hamiltonian flow. For a given starting position s_0 , the map that advances the coordinates by one full revolution is the *path-ordered exponential*

$$\mathcal{M}_{s_0} = \mathcal{P} \exp \left[\int_{s_0}^{s_0+L} ds : K(s) : \right], \quad (2.10)$$

acting on phase-space points as $\vec{z} \mapsto \mathcal{M}_{s_0}(\vec{z})$.

Because $K(s)$ is periodic, one-turn maps at different starting points are related by canonical conjugation:

$$\mathcal{M}_{s_1} = \mathcal{M}(s_1, s_0) \circ \mathcal{M}_{s_0} \circ \mathcal{M}(s_1, s_0)^{-1},$$

where $\mathcal{M}(s_1, s_0)$ is the partial map from s_0 to s_1 . Thus, while \mathcal{M}_{s_0} depends on the chosen reference location s_0 , all such maps are symplectically equivalent (in the same conjugacy class), thus share the same invariants (actions, tunes). For notational convenience, we write $\mathcal{M}_0(\vec{z})$ for the map based at $s_0 = 0$.

A fixed point \vec{z}_0 of \mathcal{M}_0 corresponds to a periodic orbit of the continuous flow. In its neighborhood, the dynamics can be locally expanded and analyzed by means of a truncated power-series algebra (TPSA). The existence of a fixed point guarantees local integrability [17]; if the full system is completely integrable, then canonical action–angle variables $(\vec{J}, \vec{\phi})$ exist [18] such that

$$\mathcal{M}_0 : (\vec{J}, \vec{\phi}) \mapsto (\vec{J}, \vec{\phi} + 2\pi\vec{Q}(\vec{J})),$$

so that the equations of motion reduce to

$$\frac{d\tilde{J}_i}{dn} = 0, \quad \frac{d\tilde{\phi}_i}{dn} = 2\pi Q_i(\vec{J}).$$

In these invariant coordinates the motion resides on nested tori, and the differential equations are diagonalized into independent rotations. Constructing such coordinates is the objective of the normal-form analysis developed below.

where $\vec{Q}(\vec{J})$ are the amplitude-dependent orbital tunes. These variables are the target of normal-form analysis, since they reduce the equations of motion to their simplest diagonal form and make the invariant structure of phase space explicit.

2.2.2 Linearization of the One-Turn Map

Given the one-turn map at a reference location $s_0 = 0$,

$$\mathcal{M}_0 = \mathcal{P} \exp \left[\int_0^L ds : K(s) : \right],$$

we assume that $\vec{z} = \vec{0}$ is a fixed point of \mathcal{M}_0 and expand the map in a truncated power-series algebra (TPSA) about this point:

$$\mathcal{M}_0(\vec{z}) = M_1 \vec{z} + \mathcal{O}(|\vec{z}|^2), \quad (2.11)$$

where M is the Jacobian matrix of \mathcal{M}_0 at the origin. Since \mathcal{M}_0 is symplectic, $M \in \text{Sp}(6, \mathbb{R})$.

The linear dynamics are governed by

$$\vec{z}_{n+1} = M \vec{z}_n,$$

which represent the discrete analogue of a periodic linear Hamiltonian system. By Floquet's theorem, there exists a linear canonical transformation

$$\vec{z} = A_0^{-1} \vec{\zeta},$$

such that in the transformed coordinates the one-turn linear map takes the block-diagonal rotational form

$$\widetilde{M} \equiv A_0 M A_0^{-1} = e^{iH_2},$$

where

$$H_2(\vec{\zeta}) = \sum_i \omega_i J_i, \quad J_i = \frac{1}{2}(\zeta_{2i-1}^2 + \zeta_{2i}^2),$$

and $\omega_i = 2\pi Q_i$ are the small-amplitude phase advances per turn. The linearized equations of motion in these Floquet coordinates are

$$\frac{dJ_i}{dn} = 0, \quad \frac{d\phi_i}{dn} = 2\pi Q_i, \quad (2.12)$$

where n is the turn index. Thus the (J_i, ϕ_i) form linear action–angle variables describing invariant ellipses in each phase plane, and e^{iH_2} is the one-turn map written as a pure rotation in these coordinates.

2.2.3 Normal Form and the Interpolating Hamiltonian

In the same Floquet-normalized coordinates, the full one-turn map \mathcal{M}_0 admits a Lie factorization, known as the Dragt-Finn factorization [19]

$$\widetilde{\mathcal{M}}_0 \equiv A_0 \mathcal{M}_0 A_0^{-1} =_N e^{iH_2} e^{if_3} e^{if_4} \dots e^{if_N}, \quad (2.13)$$

where each generator f_m is a homogeneous polynomial of degree m in the Floquet variables $A_0(x, p_x, y, p_y, z, \delta)^T$ and the equivalence class $=_N$ is understood in the truncated power-series algebra of order N .

A further canonical transformation [20]

$$\mathcal{T} = e^{iG_N} \dots e^{iG_4} e^{iG_3}$$

can be constructed such that, in the new variables $(\tilde{J}_i, \tilde{\phi}_i) = \mathcal{T}^{-1}(J_i, \phi_i)$, the one-turn map simplifies to a pure Lie exponential with an angle-independent effective Hamiltonian,

$$\mathcal{M}_{\text{NF}} \equiv \mathcal{T} \widetilde{\mathcal{M}}_0 \mathcal{T}^{-1} =_N \exp(: H_{\text{eff}}(\vec{J}) :), \quad (2.14)$$

in the nonresonant case. Explicitly,

$$H_{\text{eff}}(\vec{J}) = \sum_i \nu_i \tilde{J}_i + \sum_{|\vec{k}| \geq 2} a_{\vec{k}} \tilde{J}_1^{k_1} \tilde{J}_2^{k_2} \tilde{J}_3^{k_3},$$

defines the *interpolating Hamiltonian* of the one-turn map in normal-form coordinates. The corresponding discrete equations of motion are

$$\frac{d\tilde{J}_i}{dn} = 0, \quad \frac{d\tilde{\phi}_i}{dn} = \frac{\partial H_{\text{eff}}}{\partial \tilde{J}_i} \equiv 2\pi Q_i(\vec{J}). \quad (2.15)$$

The \tilde{J}_i are nonlinear invariants (to the truncation order) and are most naturally expressed in the complex (phasor) basis of the original action–angle coordinates,

$$c_i^\pm = \frac{x_i \pm ip_i}{\sqrt{2}} = \sqrt{J_i} e^{\pm i\phi_i}.$$

In this representation the canonical transformation \mathcal{T} acts on the phasors as a near-identity Lie transformation, yielding

$$\tilde{J}_i = \frac{1}{2} \tilde{c}_i^+ \tilde{c}_i^- = \frac{1}{2} c_i^+ c_i^- + \sum_{\vec{m}, \vec{n}} \alpha_{\vec{m}, \vec{n}}^{(i)} (c^+)^{\vec{m}} (c^-)^{\vec{n}}, \quad (2.16)$$

where $(c^+)^{\vec{m}} (c^-)^{\vec{n}} \equiv (c_1^+)^{m_1} (c_2^+)^{m_2} (c_3^+)^{m_3} (c_1^-)^{n_1} (c_2^-)^{n_2} (c_3^-)^{n_3}$, and the coefficients $\alpha_{\vec{m}, \vec{n}}^{(i)}$ are determined by the generating functions defining \mathcal{T} .

Thus the nonlinear invariant tori are smooth deformations of the linear Floquet tori. Equation (2.15) has the same diagonal form as Eq. (2.12), but now with amplitude-dependent tunes $\partial H_{\text{eff}}/\partial \tilde{J}_i$ and fully nonlinear invariants \tilde{J}_i .

2.2.4 Spin Normal Form

The spin motion along a particle trajectory is governed by the Thomas–Bargmann–Michel–Telegdi equation

$$\frac{d\vec{S}}{ds} = \vec{\Omega}(\vec{z}(s), s) \times \vec{S}, \quad (2.17)$$

where $\vec{\Omega}(\vec{z}, s)$ is the local spin–precession vector evaluated along the orbital trajectory $\vec{z}(s)$ generated by the Hamiltonian flow of $K(s, \vec{z})$. For a given initial phase–space point \vec{z}_0 at azimuth s_0 , the one–turn spin rotation is [21]

$$\mathcal{R}_{s_0}(\vec{z}_0) = \mathcal{P} \exp \left[\int_{s_0}^{s_0+L} ds \widehat{\Omega}(\vec{z}(s; \vec{z}_0), s) \right], \quad (2.18)$$

where $\widehat{\Omega}$ is the 3×3 antisymmetric matrix representing the cross product $\vec{\Omega} \times (\cdot)$. The dependence on \vec{z}_0 emphasizes that the spin transport is determined by the corresponding orbital trajectory. The combined one–turn spin–orbit map therefore acts as

$$(\vec{z}, \vec{S}) \mapsto (\mathcal{M}_0(\vec{z}), \mathcal{R}_0(\vec{z})\vec{S}),$$

where \mathcal{M}_0 is the orbital one–turn map defined in Eq. (2.10).

Invariant Spin Field and Frame. For each invariant torus of the orbital motion, labeled by the nonlinear actions \vec{J} and angles $\vec{\phi}$, the spin dynamics can be viewed as an $\text{SO}(3)$ fibre over that torus [22, 23]. The spin transport along the orbit defines a quasiperiodic map

$$(\vec{\phi}, \vec{S}) \mapsto (\vec{\phi} + 2\pi\vec{Q}(\vec{J}), \mathcal{R}_0(\vec{J}, \vec{\phi})\vec{S}),$$

which constitutes the discrete analogue of a skew product flow on $\mathbb{T}^3 \times S^2$.

A continuous unit vector field $\mathbf{n}(\vec{J}, \vec{\phi}, s)$ satisfying the stationarity condition

$$\mathbf{n}(\mathcal{M}_0(\vec{z})) = \mathcal{R}_0(\vec{z})\mathbf{n}(\vec{z}), \quad (2.19)$$

is called an *Invariant Spin Field* (ISF). If a particle's spin is initially aligned with \mathbf{n} , it remains aligned after each full turn, and the scalar invariant $J_S = \vec{S} \cdot \mathbf{n}$ is conserved away from resonance.

The ISF provides a natural choice of orthonormal basis $U_{\vec{J}}(s, \vec{\phi}) = [\mathbf{u}_1, \mathbf{u}_2, \mathbf{n}]_{\vec{J}}$ called the *Invariant Frame Field* (IFF). In this moving frame the one-turn spin rotation is diagonalized to a single rotation about the third axis:

$$U_{\vec{J}}^{-1}(\mathcal{M}_0(\vec{z})) \mathcal{R}_0(\vec{z}) U_{\vec{J}}(\vec{z}) = R_3(2\pi\nu(\vec{J})), \quad (2.20)$$

where $R_3(\alpha)$ denotes a rotation through angle α about \mathbf{n} , and $\nu(\vec{J})$ is the *amplitude-dependent spin tune*.

Equation (2.20) is directly analogous to the orbital normal form

$$(\mathcal{T}A) \mathcal{M}_0 (\mathcal{T}A)^{-1} =_N \exp(: H_{\text{eff}}(\vec{J}) :),$$

with $U_{\vec{J}}$ playing the role of the normalizing transformation and $\nu(\vec{J})$ the analogue of the amplitude-dependent orbital tunes $\partial H_{\text{eff}}/\partial \tilde{J}_i$. In this geometric picture the spin motion forms an SU(2) (or SO(3)) fibre bundle over the nonlinear orbital tori, and the IFF defines a local coordinate system that brings the coupled spin-orbit map into SO(2) normal form. Note that the fiber bundle formalism is naturally capable of including the Stern-Gerlach-type feedback of spin on orbital motion. The simplification we make that comes from ignoring/decoupling these effects is that the principal bundle reduces to a product bundle, i.e. Cartesian product space.

2.2.5 Resonant Terms

The normal-form transformation of Eq. (2.14) assumes that no linear combination of the tunes satisfies a resonance condition. When a combination of integer coefficients $\vec{k} = (k_x, k_y, k_z)$ exists such that

$$\vec{k} \cdot \vec{Q}(\vec{J}) = m, \quad m \in \mathbb{Z},$$

the corresponding Fourier term cannot be eliminated by any canonical transformation. The effective interpolating Hamiltonian then acquires an irreducible angle-dependent component [24]:

$$H_{\text{eff}}(\vec{J}, \vec{\phi}) = H_0(\vec{J}) + \sum_{n=1}^{N/|k|} V_{n\vec{k}}(\vec{J}) \cos(n\vec{k} \cdot \vec{\phi}). \quad (2.21)$$

This term represents an angle-dependent deformation of the otherwise invariant torus, and its strength $V_{\vec{k}}(\vec{J})$ measures the amplitude of the resonant coupling.

The equations of motion in the resonant subspace become

$$\frac{d\tilde{J}_i}{dn} = -\frac{\partial H_{\text{eff}}}{\partial \tilde{\phi}_i} = \tilde{k}_i V_{\vec{k}}(\vec{J}) \sin(\vec{k} \cdot \vec{\phi}), \quad \frac{d\tilde{\phi}_i}{dn} = \frac{\partial H_{\text{eff}}}{\partial \tilde{J}_i}. \quad (2.22)$$

Hence the nonlinear invariants \tilde{J}_i are no longer constant but oscillate around their unperturbed values with a modulation period inversely proportional to the resonance strength $V_{\vec{k}}$. In the exact resonance ($\vec{k} \cdot \vec{Q} = m$), the \tilde{J}_i can experience large secular changes corresponding to slow drift through the resonance.

When a system parameter τ varies slowly (for instance, during an acceleration ramp), the normal-form transformation \mathcal{T}_τ that defines the coordinates $(\vec{J}, \vec{\phi})$ itself changes with τ . Hence the evolution of \tilde{J}_i acquires an additional geometric term due to the slow deformation of the canonical basis:

$$\frac{d\tilde{J}_i}{d\tau} = \left. \frac{\partial \tilde{J}_i}{\partial n} \right|_\tau + \left. \frac{\partial \tilde{J}_i}{\partial \tau} \right|_n.$$

The second term represents the parametric rotation of the normal-form frame in phase space, analogous to the adiabatic precession of the spin frame described by the vector field η . This geometric contribution couples slow machine variations to the orbital resonances and is responsible for irreversible transport of the nonlinear actions.

This mechanism underlies the well-known Froissart–Stora behavior of adiabatic spin flips and the analogous secular diffusion of orbital actions.

Thus, in the presence of resonant terms, the nonlinear invariant tori are deformed into *angle-dependent* invariant curves, and the invariants \tilde{J}_i acquire slow modulations driven by the resonant harmonics $V_{\vec{k}}(\vec{J})$.

2.3 Spin Resonance Overview

The rich resonance structure of hadron spin dynamics emerges from the periodic nature of particle motion in circular accelerators. As particles circulate around the ring, they experience time-varying electromagnetic fields that can drive spin transitions when certain frequency matching conditions are satisfied. Understanding these resonances, their locations, strengths, and mitigation strategies, is central to maintaining polarization during acceleration.

2.3.1 Spin Tune and Resonance Conditions

The fundamental parameter characterizing spin motion in a circular accelerator is the spin tune ν , defined as the number of spin precessions per orbital revolu-

tion. For a particle in circular orbit in a uniform magnetic field, we refer to the revolution azimuth along the path as:

$$\theta(s) \equiv \int_0^s \frac{q}{p} B_y(s) ds, \quad (2.23)$$

such that $\theta(L) = 2\pi$.

Building on this, we define the spin phase advance along the closed orbit, relative to the revolution azimuth, as:

$$\psi(s) = \int_0^s \frac{q}{p} G\gamma B_y(s) ds' = G\gamma \theta(s), \quad (2.24)$$

Such that the closed orbit spin tune is given by $\nu_0 = \psi(L)/2\pi = G\gamma$

where the integral is taken around the closed orbit, B_y is the magnetic field component perpendicular to the velocity, and p is the particle momentum.

For protons with $G = 1.793$, the closed-orbit spin tune grows linearly with energy, reaching $\nu_0 \approx 525$ at the top energy of RHIC (275 GeV per beam). This rapid growth with energy means that hadron beams cross many potential resonances during acceleration, each representing an opportunity for depolarization.

Resonances occur when the spin precession frequency becomes commensurate with orbital motion frequencies. The general resonance condition is:

$$\nu = k + m_x Q_x + m_y Q_y + m_z Q_z \quad (2.25)$$

where $k, m_x, m_y,$ and m_z are integers, and Q_x, Q_y, Q_z are the horizontal, vertical, and synchrotron tunes, respectively. The integers m_x, m_y, m_z determine the type and order of the resonance, while k describes an imperfection resonance or shifts the location in energy of an intrinsic resonance.

2.3.2 Intrinsic Resonances

Intrinsic resonances arise from the periodic focusing structure of the accelerator lattice and are characterized by $|m_x| + |m_y| + |m_z| > 0$ in Equation 2.25. The most important class consists of first-order vertical resonances with $m_y = 1$ and $m_x = m_z = 0$:

$$\nu = k \pm Q_y, \quad k \in \mathbb{Z}. \quad (2.26)$$

These resonances are driven during coherence between spin precession and the vertical betatron motion of particles as they oscillate around the closed orbit. The strength of an intrinsic resonance depends on the amplitude of betatron oscillations and the details of the lattice optics, particularly the vertical beta function and the vertical magnetic field components created by focusing elements.

With protons in flat rings generally having $\nu_0 = G\gamma \approx 1.793\gamma$, each of the \pm intrinsic resonances are spaced by approximately $\Delta\gamma m_p c^2 \approx 523 \text{ MeV}$. During a typical acceleration ramp from injection to top energy, a hadron beam potentially crosses hundreds of these resonances, with strength ϵ_k^\pm :

$$\epsilon_k^\pm = \frac{G\gamma + 1}{2\pi} \oint \sqrt{\frac{J_y \beta_y(s)}{2}} k(s) e^{-i(k\theta(s) \pm \phi_y(s))} ds \quad (2.27)$$

Fig. 2.1 shows the characteristic first-order intrinsic spin resonance spectrum of RHIC using the 2022 lattice optics without snakes.

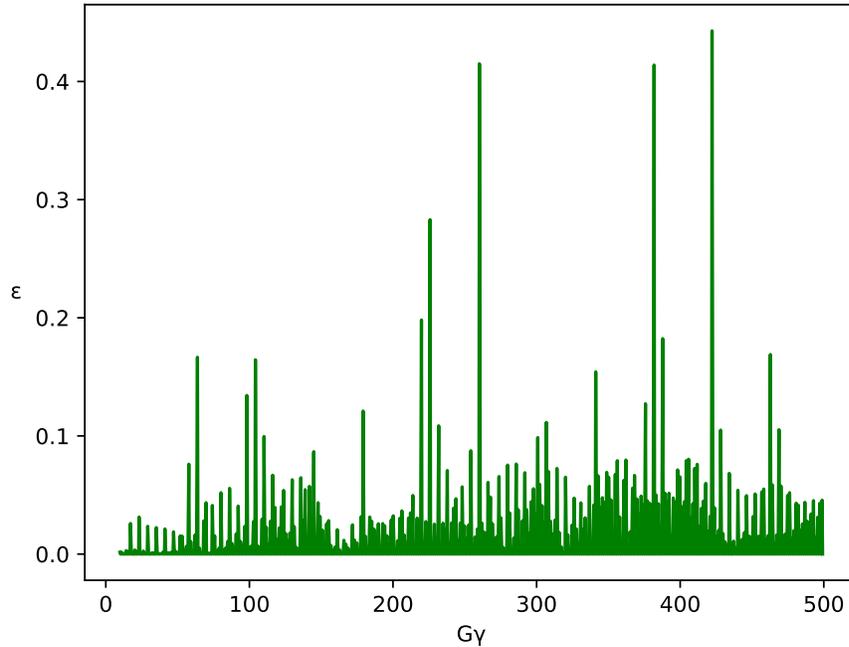


Figure 2.1: First-order resonance strength spectrum of RHIC ε_κ in the absence of snakes.

Higher-order intrinsic resonances with $|m_x| + |m_y| + |m_z| > 1$ also exist but are generally weaker because they are driven by nonlinearities in the spin motion arising from strong precession within an element or higher-order terms in the multipole expansion of the magnetic fields. However, in machines with strong coupling or nonlinearities, such as those with insertion devices or Siberian snakes, these higher-order resonances can become significant.

2.3.3 Imperfection Resonances

Imperfection resonances correspond to $m_x = m_y = m_z = 0$ in Equation 2.25, giving the simple condition:

$$\nu = k, \quad k \in \mathbb{Z}. \quad (2.28)$$

These resonances are driven during coherence between spin precession and beam circulation by magnetic field errors, magnet misalignments, and closed orbit distortions that create periodic vertical magnetic field components at fixed locations. Unlike intrinsic resonances, which are driven by betatron motion, imperfection resonances affect the entire beam regardless of beam size or emittance.

For protons, imperfection resonances similarly occur each 523 MeV increase in beam energy. For helions, they occur each 224 MeV increase in nucleon energy.

The strength of an imperfection resonance is proportional to the vertical magnetic field integrated around the closed orbit:

$$\epsilon_k = \frac{G\gamma + 1}{2\pi} \oint y_{\text{co}}(s) k(s) e^{-ik\theta(s)} ds \quad (2.29)$$

where $y_{\text{co}}(s)$ is the vertical component of the closed orbit relative to the center of quadrupoles, and $k(s) \equiv \frac{q}{p} \frac{\partial B_x}{\partial y}$ is the normalized quadrupole strength. This integral is directly proportional to the magnitude of vertical closed orbit distortion, making orbit correction a primary tool for controlling imperfection reso-

nance strengths.

2.3.4 Higher-Order Resonances

Beyond first-order intrinsic and imperfection resonances, higher-order spin-orbit resonances become critical in lattices with Siberian snakes. While snakes lock the closed-orbit spin tune at $\nu_0 = 1/2$, the amplitude-dependent spin tune $\nu(\vec{J})$ can deviate from $1/2$ for finite orbital actions \vec{J} and cross higher-order resonance conditions [25]

$$\nu(\vec{J}) = j_0 + \vec{j} \cdot \vec{Q}, \quad (j_0, \vec{j}) \in \mathbb{Z}^4. \quad (2.30)$$

A convenient description uses the Invariant Frame Field, $U_{\vec{j}(s, \phi)} = [\mathbf{u}_1, \mathbf{u}_2, \mathbf{n}]_{\vec{j}}$, in which the spin precesses by $\nu\theta$ about the ISF \mathbf{n} . The ISF projects the spin to the adiabatic invariant $J_S = \vec{S} \cdot \mathbf{n}$. Under a slow parameter change τ (e.g., energy ramp), the IFF varies according to the vector

$$\boldsymbol{\eta} \equiv \frac{1}{2}(\mathbf{u}_1 \times \partial_\tau \mathbf{u}_1 + \mathbf{u}_2 \times \partial_\tau \mathbf{u}_2 + \mathbf{n} \times \partial_\tau \mathbf{n}), \quad (2.31)$$

which satisfies $\partial_\tau \mathbf{n} = \boldsymbol{\eta} \times \mathbf{n}$. Components of $\boldsymbol{\eta}$ transverse to \mathbf{n} drive changes in J_S and determine higher-order resonance strengths via Fourier components [26] of $\boldsymbol{\eta} \cdot (\mathbf{u}_1 + i\mathbf{u}_2)$ along the azimuth $\theta(s)$:

$$\varepsilon_\kappa \propto \lim_{N \rightarrow \infty} \frac{1}{NL} \int_0^{NL} \boldsymbol{\eta} \cdot (\mathbf{u}_1 + i\mathbf{u}_2) e^{-i\kappa\theta} ds, \quad (2.32)$$

with $\kappa \equiv \vec{j} \cdot \vec{Q} \bmod 1$. Crossing isolated higher-order resonances during a ramp follows the Froissart–Stora law with spin-tune jumps of size $2\varepsilon_\kappa$ and depolarization set by $(\varepsilon_\kappa, \alpha)$ where α is the tune-ramp rate [26, 27].

Fig. 2.2 illustrate typical behavior extracted from the higher-order resonance analysis: sharp drops in P_{lim} and associated jumps in $\nu(\vec{J})$ as energy varies.

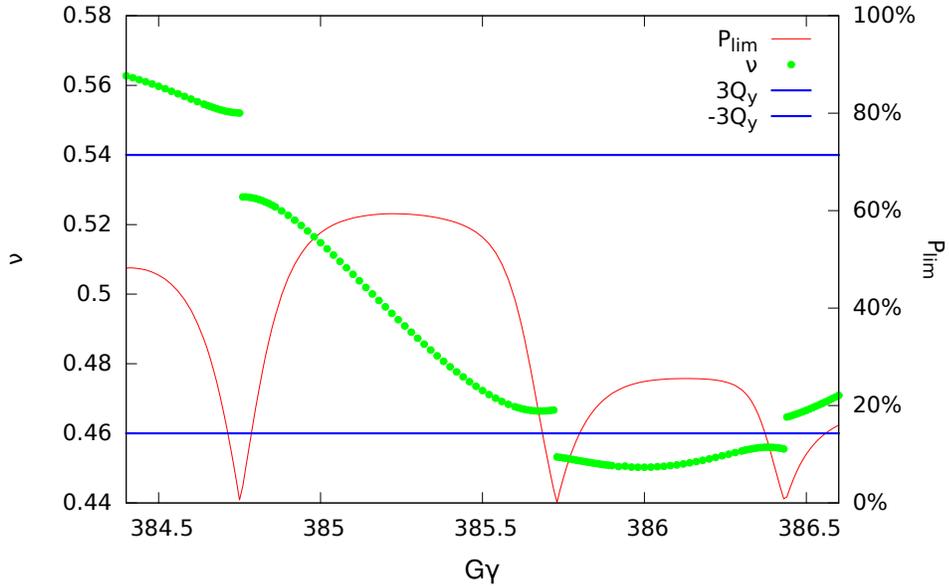


Figure 2.2: Jumps in the amplitude-dependent spin tune $\nu(\vec{J})$ across 3rd-order resonances; jump size = $2\varepsilon_\kappa$.

2.4 Field Modeling and Map Construction

Accurate modeling of electromagnetic fields is fundamental to reliable spin dynamics simulations. The challenge lies in representing the complex three-dimensional field distributions found in real accelerator magnets while maintaining computational efficiency and preserving the mathematical properties, particularly symplecticity, required for long-term tracking studies.

Traditional approaches to field modeling in accelerator physics often rely on thin-lens approximations or simplified field models that, while computationally convenient, can introduce systematic errors that accumulate over millions of turns. For hadron spin dynamics, where polarization survival depends sensitively on the details of field distributions and their higher-order multipole components, more sophisticated approaches are required.

2.4.1 Maxwell's Equations and Field Constraints

All electromagnetic fields must satisfy Maxwell's equations, which in the absence of time-varying fields reduce to:

$$\nabla \cdot \vec{B} = 0 \quad (2.33)$$

$$\nabla \times \vec{B} = \mu_0 \vec{J} \quad (2.34)$$

$$\nabla \times \vec{E} = 0 \quad (2.35)$$

$$\nabla \cdot \vec{E} = \frac{\rho}{\epsilon_0} \quad (2.36)$$

In most regions of interest for beam dynamics, away from current-carrying conductors and space charge, the source terms vanish, leading to the conditions $\nabla \cdot \vec{B} = 0$ and $\nabla \times \vec{B} = 0$. These constraints mean that magnetic fields can be derived from a scalar potential ϕ_m such that $\vec{B} = -\mu_0 \nabla \phi_m$, and the potential satisfies Laplace's equation $\nabla^2 \phi_m = 0$.

The divergence-free condition $\nabla \cdot \vec{B} = 0$ places strong constraints on the allowed field configurations. In particular, it couples the different field components and ensures that knowledge of one component (say, B_z) on a surface determines the other components throughout the volume, provided appropriate boundary conditions are specified.

These Maxwell constraints reflect fundamental physical principles that must be preserved in any realistic field model. Violations of these constraints, even small ones, can lead to unphysical effects such as artificial focusing or defocusing that accumulates over many turns and corrupts long-term tracking studies.

2.4.2 Generalized Gradient Field Equations

Following Venturini and Dragt [28], static fields can be described by scalar potentials

$$\vec{B} = -\nabla \psi_B, \quad \vec{E} = -\nabla \psi_E. \quad (2.37)$$

In cylindrical (ρ, θ, z) coordinates, decompose ψ as

$$\psi = \sum_{m=1}^{\infty} \psi_{m,s}(\rho, z) \sin(m\theta) + \sum_{m=0}^{\infty} \psi_{m,c}(\rho, z) \cos(m\theta). \quad (2.38)$$

Each harmonic expands in powers of ρ using the generalized gradients $C_{m,\alpha}(z)$ ($\alpha = c, s$):

$$\psi_{m,\alpha} = \sum_{n=0}^{\infty} \frac{(-1)^{n+1} m!}{4^n n! (n+m)!} \rho^{2n+m} C_{m,\alpha}^{[2n]}(z). \quad (2.39)$$

The resulting magnetic field components are

$$B_\rho = \sum_{m=1}^{\infty} \sum_{n=0}^{\infty} \frac{(-1)^n m! (2n+m)}{4^n n! (n+m)!} \rho^{2n+m-1} \times \quad (2.40)$$

$$[C_{m,s}^{[2n]}(z) \sin m\theta + C_{m,c}^{[2n]}(z) \cos m\theta] + \quad (2.41)$$

$$\sum_{n=1}^{\infty} \frac{(-1)^n 2n}{4^n n! n!} \rho^{2n-1} C_{0,c}^{[2n]}(z) \quad (2.42)$$

$$B_\theta = \sum_{m=1}^{\infty} \sum_{n=0}^{\infty} \frac{(-1)^n m! m}{4^n n! (n+m)!} \rho^{2n+m-1} \times \quad (2.43)$$

$$[C_{m,s}^{[2n]}(z) \cos m\theta - C_{m,c}^{[2n]}(z) \sin m\theta] \quad (2.44)$$

$$B_z = \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} \frac{(-1)^n m!}{4^n n! (n+m)!} \rho^{2n+m} \times \quad (2.45)$$

$$[C_{m,s}^{[2n+1]}(z) \sin m\theta + C_{m,c}^{[2n+1]}(z) \cos m\theta]. \quad (2.46)$$

Two properties are central:

- **Locality:** $\vec{B}(x, y, z)$ depends only on adjacent $C_{m,\alpha}^{[n]}$ near the evaluation z ; interpolation over $[z_i, z_{i+1}]$ uses the coefficients at the bracketing points.

- **Truncation:** Derivative cutoffs $N_{m,\alpha}$ bound Maxwell error with radius scaling $\rho^{m+N_{m,\alpha}}$.

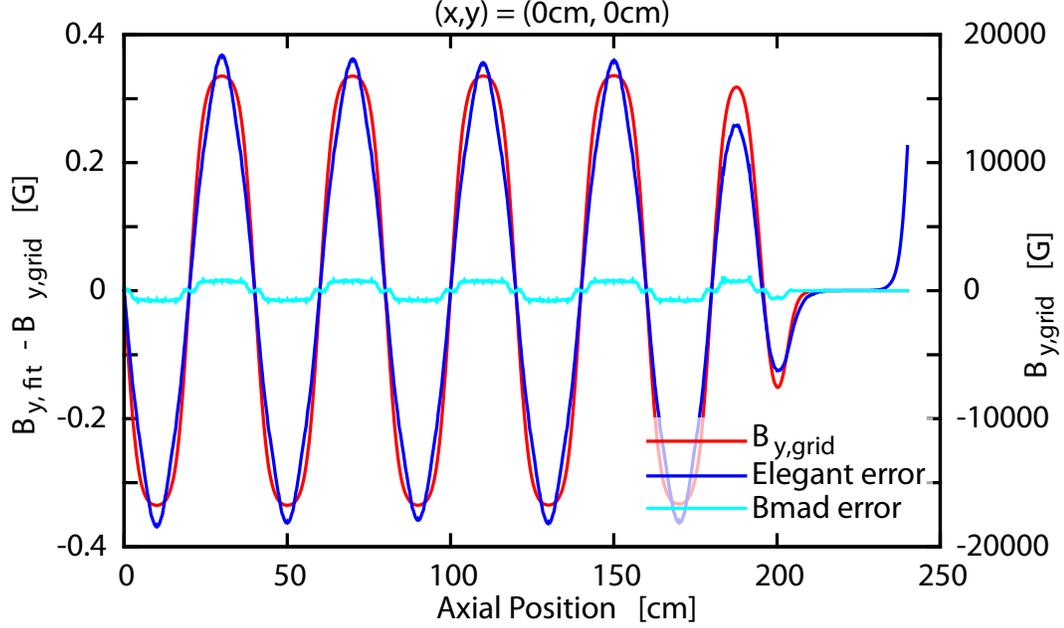


Figure 2.3: On-axis wiggler field agreement: grid value vs. GG fits using Elegant vs. Bmad algorithms (surface vs volume fits). In blue/cyan are the absolute errors of the fits (left axis), and in red is the actual value of the grid field points (right axis).

2.4.3 Generalized Gradient Fitting Algorithm

To use GGs in tracking, derivatives $C_{m,\alpha}^{[n]}(z)$ must be computed from field data. Two complementary approaches are relevant:

- **Venturini-Dragt (surface) method:** From $B_\rho(R, \theta, z)$ on a cylinder of radius R , decompose into A_m/B_m and obtain Fourier-Bessel transforms, e.g.

$$C_{m,s}^{[n]}(z) = \frac{i^n}{2^m m!} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} dk e^{ikz} \frac{k^{m+n-1}}{I_m'(kR)} \tilde{B}_m(R, k). \quad (2.47)$$

- **Sagan (volume) fitting** : Fit $C_{m,\alpha}^{[n]}(z_i)$ to nearby field points within $[z_i - \delta z, z_i + \delta z]$ and extrapolate within an interval using

$$C_{m,\alpha}^{[n]}(z) = \sum_{j=n}^{N_{m,\alpha}} \frac{(z - z_i)^{j-n}}{(j-n)!} C_{m,\alpha}^{[j]}(z_i). \quad (2.48)$$

Coefficients follow from minimizing a weighted merit function $M = \sum_k W_k \|\vec{B}_{\text{fit}}(\mathbf{r}_k) - \vec{B}_{\text{grid}}(\mathbf{r}_k)\|^2$ with core-emphasis weighting W_k .

The local method exploits GG locality, avoids nonlocal contamination from distant high-field regions, supports arbitrary grids, and affords smooth edge control at element ends. Derivative cutoffs $N_{m,\alpha}$ balance accuracy against truncation error.

Figs 2.3, and 2.4 show the advantage of the volume-fitting method used by Bmad [10] vs the surface-fitting method used by Elegant [29].

2.4.4 Vector Potentials and Symplectic Maps

For symplectic integration, it is often preferable to work with vector potentials \vec{A} rather than fields directly, since the Hamiltonian formulation naturally involves the vector potential through the canonical momentum $\vec{p} = \gamma m \vec{v} + q \vec{A}$. The magnetic field is related to the vector potential by $\vec{B} = \nabla \times \vec{A}$, but this relationship does not uniquely determine \vec{A} due to gauge freedom.

The gauge freedom can be exploited to construct vector potentials with desirable properties for numerical integration. A common choice is the Coulomb gauge $\nabla \cdot \vec{A} = 0$, which simplifies many calculations. For multipole fields, explicit expressions for the vector potential can be derived that automatically satisfy both the gauge condition and Maxwell's equations.

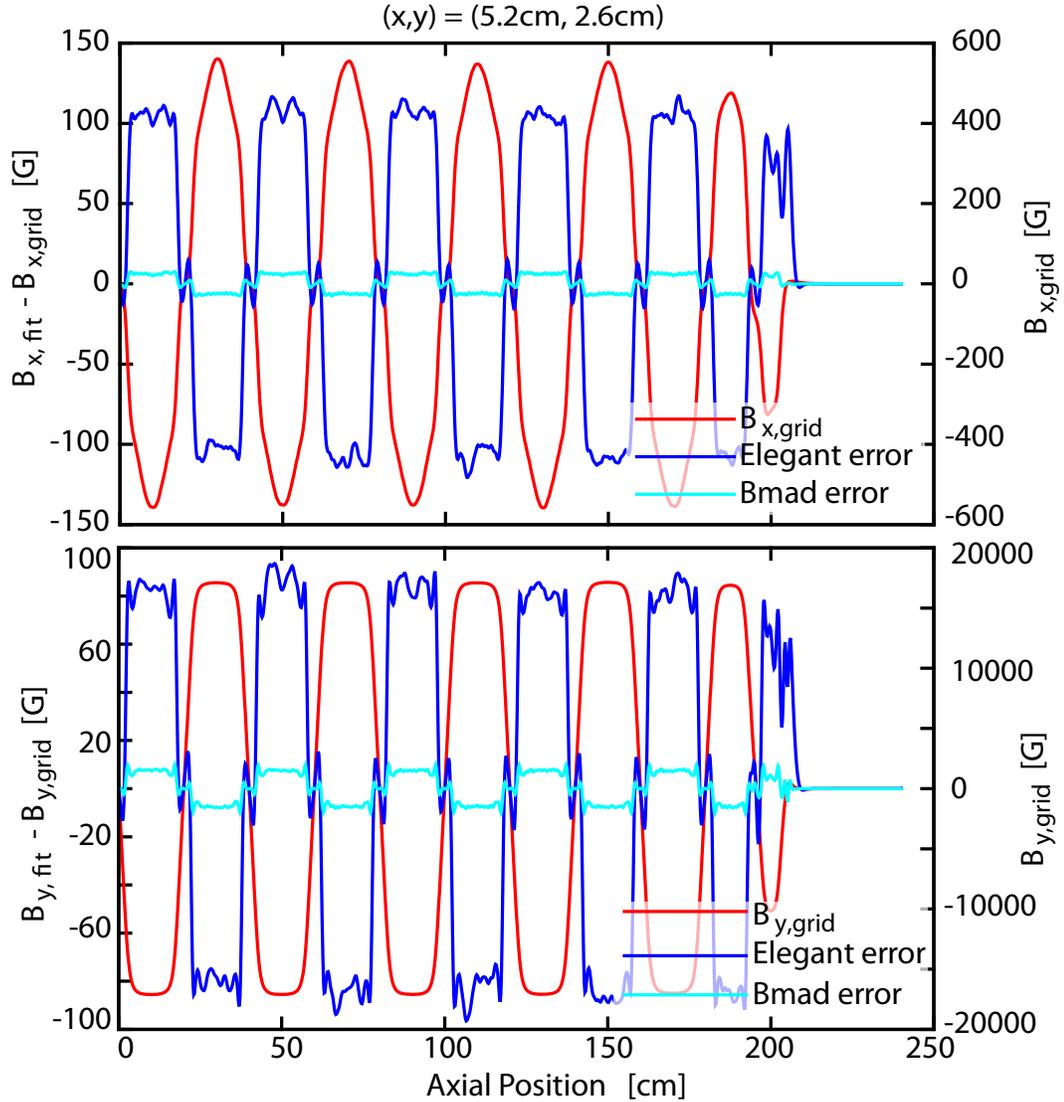


Figure 2.4: Off-axis wiggler field agreement at max grid radius: grid value vs. GG fits (surface vs volume fits). In blue/cyan are the absolute errors of the fits (left axis), and in red is the actual value of the grid field points (right axis).

Consider a pure dipole field $\vec{B} = B_0\hat{y}$ in the y -direction. A suitable vector potential is $\vec{A} = B_0z\hat{x}$, which gives $\vec{B} = \nabla \times \vec{A} = B_0\hat{y}$ and satisfies $\nabla \cdot \vec{A} = 0$. However, using a different gauge of $\vec{A} = -B_0x\hat{z}$ enables a closed-form solution of Hamilton's equations through such a field. For more complex multipole fields, the vector potential expressions become correspondingly more complicated but can be expanded systematically.

The advantage of working with vector potentials becomes apparent when constructing symplectic integration schemes. The Hamiltonian in Equation 2.3 involves the vector potential directly, and symplectic integrators can be designed to preserve the canonical structure exactly. This preservation ensures that long-term tracking studies maintain the correct phase space volume and do not exhibit artificial damping or growth.

2.4.5 From Local Models to Global Tracking

The field modeling techniques described above apply to individual accelerator elements, dipoles, quadrupoles, sextupoles, etc. To perform global tracking around an entire accelerator, these local models must be combined into a comprehensive description of the machine.

The standard approach is to divide the accelerator into a sequence of elements, each with its own field model and integration scheme. Particles are tracked through each element in sequence, with the output of one element serving as the input to the next. This element-by-element approach allows different integration methods to be used for different element types, optimizing computational efficiency while maintaining accuracy.

However, care must be taken to ensure that the overall tracking remains symplectic even when different integration schemes are used for different elements. This requirement places constraints on the choice of integration methods and the interfaces between elements. In particular, the transformation from the coordinate system of one element to the next must preserve the symplectic structure.

For nontrivial elements with complex field distributions, such as partial Siberian snakes or combined-function magnets, tracking may require explicit or implicit integration. The challenge is to choose the field description and tracking method to balance accuracy against computational cost while preserving physical and geometric properties.

The Julia ecosystem provides sophisticated tools for constructing these global tracking codes. Packages like `DifferentialEquations.jl` offer a wide variety of geometric integrators, while `BeamTracking.jl` has been developed to provide high-level interfaces to extremely performant low-level symplectic integration schemes. The combination of these tools enables the construction of efficient, accurate tracking codes that can handle extremely complex field distributions found in modern hadron accelerators. Automatic differentiation packages like `ForwardDiff.jl` and `GTPSA.jl`, which can perform arbitrary-order Taylor expansions of arbitrary smooth functions, enable the construction of symplectic polynomial maps known as truncated power series algebra (TPSA) maps that are necessary for normal form analysis.

2.5 The Big Picture

The mathematical foundations presented in this chapter form an interconnected framework that underpins all subsequent work in this dissertation. Each component, from the fundamental T-BMT equation through field modeling and symplectic integration, addresses a specific aspect of the broader challenge of maintaining spin polarization in hadron accelerators.

The T-BMT equation serves as the fundamental physical law governing spin

evolution, encoding the coupling between electromagnetic fields and particle spin that makes both polarization and depolarization possible. Its relativistic structure explains why hadron spin dynamics become increasingly challenging at higher energies, where the spin tune grows linearly with energy and resonances become densely packed.

Hamiltonian mechanics and symplectic integration provide the mathematical framework that ensures faithful long-term tracking. While the preservation of per-particle phase space volume can be violated by collective effects, beam dynamics in most high-energy synchrotrons is described to leading order by single-particle dynamics. Therefore, maintaining the symplectic nature of motion is essential for simulations that span over millions of turns. Violations of symplecticity lead to artificial resonances and spurious beam losses that can completely invalidate simulation predictions.

Spin resonance theory explains the fundamental mechanisms by which polarization is lost and defines the conditions under which it can be preserved. The distinction between intrinsic and imperfection resonances provides a conceptual framework for understanding different mitigation strategies: lattice design and snake placement for intrinsic resonances, orbit correction and field quality for imperfection resonances.

Field modeling bridges the gap between abstract mathematical descriptions and the complex electromagnetic environments found in real accelerators. Maxwell-consistent field models ensure that simulations capture the essential physics while avoiding unphysical artifacts that can arise from oversimplified representations.

Together, these mathematical tools form the conceptual scaffolding that supports the practical work described in subsequent chapters. The investigation of field realism in the AGS (Chapter 3) relies on the Maxwell-consistent field modeling techniques developed here. The analysis of snake configurations in RHIC and the HSR (Chapter 4) uses the spin resonance analysis and normal form methods to understand polarization survival mechanisms. The development of correction strategies (Chapter 4, Section 4.6) builds on these approaches to effectively leverage multi-fidelity information about the system.

Perhaps most importantly, the computational infrastructure requirements identified here, the need for symplectic integration, Maxwell-consistent fields, and efficient ensemble tracking, have driven the development of the Julia-based tools that enable all subsequent simulations. The performance requirements of million-turn tracking, combined with the mathematical constraints of symplectic integration, and the ability to rapidly and rigorously analyze beam dynamics define a unique computational niche that existing accelerator codes did not adequately address.

Looking forward, these mathematical foundations provide the basis for addressing even more challenging problems in future accelerator facilities. The EIC's requirements for multiple polarized species, each with different gyromagnetic ratios and resonance structures, will push these methods to their limits. The techniques developed here provide both the theoretical understanding and computational tools needed to meet these challenges, ensuring that the remarkable scientific potential of polarized hadron beams can be fully realized in the coming decades.

The path from fundamental equations to operating accelerators is indeed

long and complex, but the mathematical foundations presented in this chapter provide a solid starting point. By ensuring that our theoretical understanding is both rigorous and computationally tractable, we create the possibility for genuine progress in our quest for knowledge.

CHAPTER 3

MODELING MAGNETIC FIELDS AND SYMPLECTIC TRACKING IN THE AGS

The Alternating Gradient Synchrotron (AGS) presents a stringent test for long-term spin tracking: combined-function dipoles and two partial Siberian snakes create strong local coupling and optics distortions, especially near injection [1, 30–32]. Prior non-Maxwellian field models led to non-symplectic tracking with spurious emittance growth and artificial depolarization within $O(10^3)$ turns.

This chapter replaces those models with Maxwell-consistent field representations supported by symplectic integration through such fields. Symplectic tracking here does *not* use generalized gradients (GGs). Instead, it relies on the exact cylindrical map decomposition; this is only a few times faster than Runge–Kutta tracking through interpolated grid fields, but it preserves symplecticity and stability. Separately, GG maps are used for fast non-symplectic tracking and benchmarking, and they are $\sim 1000\times$ faster than grid maps, serving a different role than the symplectic model.

We proceed as follows. Section 3.1 concisely states the field-modeling problem in the AGS. Section 3.2 documents the failure modes of non-symplectic tracking and the Maxwellian symplectic solution. Section 3.2.6 presents validation and performance. A new Section 3.4 performs lattice matching to machine settings across 100 energies and highlights the stability of motion.

3.1 Field Modeling Challenges in the AGS

The Alternating Gradient Synchrotron occupies a unique position in the landscape of hadron accelerators. With a circumference of 807 meters and an injection energy of 2.3 GeV (corresponding to $\gamma = 2.45$ for protons), it accelerates beams to 23.8 GeV ($\gamma = 25.4$) before extraction to RHIC [30]. While this energy range encompasses only a subset of its full capabilities before saturation, it is still a critical regime for spin dynamics: the closed-orbit spin tune ν_0 grows from approximately 4.5 to 45.5, crossing dozens of depolarizing resonances [33–35].

The magnetic environment of the AGS reflects its history and the constraints of its physical footprint [36]. The main bending and focusing is provided by 240 combined-function dipole magnets arranged in 12 superperiods of 20 each. Unlike modern separated-function lattices where dipoles provide pure bending and quadrupoles provide pure focusing, the AGS dipoles accomplish both tasks simultaneously through pole face angles that create a superposed quadrupole gradient on top of the bending field. This is accomplished by using sections of a rectangular hyperbola where the positive and negative solutions are almost parallel, effectively producing an off-center quadrupole. Such an elegant solution minimizes the number of magnets and the required circumference but creates complex three-dimensional field distributions that resist common simplifying descriptions.

3.1.1 Combined-Function Dipole Magnets

The AGS main magnets are long, straight, C-shaped dipoles with midplane-symmetric hyperbolic pole tips [36], shown in Fig. 3.1. Their geometry intentionally combines dipole and quadrupole functions: a vertically dominant field B_y that bends the beam, and a built-in horizontal gradient $\partial B_y/\partial x$ that provides focusing. The hyperbolic profile of the pole tips ensures that the field is approximately linear in x across the good-field region, producing the desired alternating-gradient focusing of a strong-focusing synchrotron.

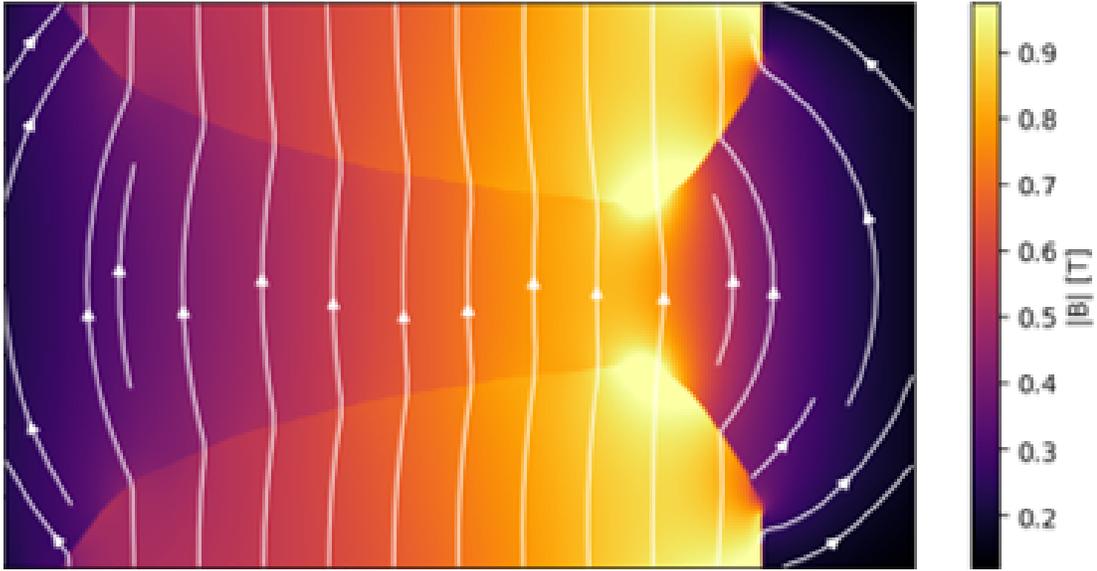


Figure 3.1: Simulated field of an AGS Type-B main magnet at 2650 A, calculated via 2D FEM with 1–5 mm grid spacing.

A simple geometric estimate of the focusing gradient follows from the local slope of the pole faces. For a given hyperbolic pole-face height $h(x) \propto \frac{1}{x_0-x}$ and pole-face angle $\theta(x) = \tan^{-1}(h'(x))$ measured from the midplane,

$$\frac{\partial B_y}{\partial x} = \frac{\partial B_x}{\partial y} \approx B_{y,0}(x) \frac{h'(x)}{h(x)} \propto \frac{h'(x)}{h^2(x)}, \quad (3.1)$$

where the vertical field at that point on the midplane is approximately inversely proportional to gap height $B_{y,0}(x) \propto 1/h(x)$. The gradient arises because the

field lines converge in regions where the gap narrows, and diverge where it widens. This geometric relation provides a first-order estimate of the effective quadrupole strength built into the dipole field.

Historically, the AGS has modeled these straight combined-function magnets as if they were *curved multipoles*—that is, magnets in which the total field remains constant along an ideal circular reference trajectory through the ring [37–39]. In this curved-multipole approximation, the dipole field $B_{y,0}$ is assumed constant along the arc, while the superimposed gradient $G = \partial B_y / \partial x$ provides the focusing. Although this treatment is convenient for optics calculations, it effectively replaces the true straight magnet geometry with an idealized curved body, thereby neglecting the natural longitudinal variation of field experienced by a particle as it bends through a straight gradient field.

In the physical magnet, the body field (neglecting fringes) is genuinely two-dimensional and obeys Laplace’s equation in a straight Cartesian geometry:

$$B_y(x, y) = B_0 + Gx, \tag{3.2}$$

$$B_x(x, y) = Gy, \tag{3.3}$$

$$B_s(x, y) = 0. \tag{3.4}$$

A reference particle that bends through this field must satisfy

$$B_y(s) = B_0 + G x_{\text{ref}}(s),$$

where $x_{\text{ref}}(s)$ is the horizontal displacement of the actual orbit through the magnet. Since $x_{\text{ref}}(s)$ varies along the path, the effective field seen by the particle is *not constant in s* . There is no trajectory through the straight magnet for which B_y remains uniform: the bending itself results from sampling the gradient across the aperture. The common curved-multipole approximation thus

omits the additional focusing and chromatic terms associated with this intrinsic s -dependence.

A more faithful description treats the AGS main magnets as true straight-field regions joined to the ring by entrance and exit rotations that match the global curvature of the lattice. In this representation, the reference orbit within each magnet is slightly sinusoidal, not circular, and the resulting focusing varies smoothly along the magnet length. This approach preserves the exact Laplacian structure of the magnetic field, avoids spurious longitudinal components, and maintains full symplecticity in tracking and map-based analyses.

From the standpoint of field classification, it is important to distinguish between the exact multipole solutions in a *curved* coordinate system and the ordinary multipoles in a *straight* geometry. McMillan's horizontally and vertically pure multipole families apply when the reference coordinate system itself follows a curved arc, and are relevant only if one insists on representing the magnet as an element with intrinsic curvature. By contrast, the AGS main magnets are physically straight and are most naturally described by the simple linear model above: a superposition of a uniform dipole field B_0 and a constant gradient G , satisfying the two-dimensional Laplace equation in Cartesian coordinates. In this straight-field picture, any additional focusing and nonlinear effects do not arise from any exotic multipole mixing, but from the fact that the closed orbit is necessarily non-zero and traverses the gradient field, so that the effective curvature and focusing become s -dependent through $x_{\text{ref}}(s)$ even though the underlying body field is uniform in s .

3.1.2 Partial Siberian Snakes and Local Coupling

The AGS contains two partial Siberian snakes installed in the early 2000s to enable polarized proton acceleration [5,6]. Unlike the full 180° snakes in RHIC, space and budget constraints limited the AGS snakes to 90° rotations each. The snakes are helical dipole magnets approximately 2.5 meters long, distinguished as a "cold snake" which is superconducting with field strength of ~ 3 T and a "warm snake" which is normal-conducting with a field strength of ~ 1.5 T. Their location is shown schematically in Fig. 3.2

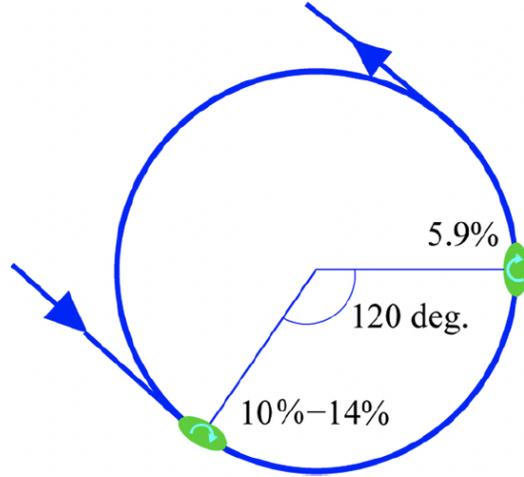


Figure 3.2: Locations of two partial snakes in the AGS ring.

Idealized spin rotation model. To a leading-order approximation, the snakes are treated as producing pure spin rotations of angles ψ_1 and ψ_2 about the longitudinal axis, with no distortion of the closed orbit. Between the snakes, the spin precesses about the vertical guide field by angles

$$a = \frac{2\pi}{3}G\gamma, \quad b = \frac{4\pi}{3}G\gamma,$$

corresponding to the two unequal arc lengths separating the cold and warm snakes. The one-turn spin transport operator can then be expressed as a se-

quence of rotations:

$$R_{\text{turn}} = R_z(\psi_2) R_y(b) R_z(\psi_1) R_y(a),$$

where R_y and R_z denote rotations about the vertical and longitudinal axes, respectively.

Using standard $SU(2)$ matrices, the trace of this composite rotation gives the closed-orbit spin tune ν_0 via

$$\cos(\pi\nu_0) = \frac{1}{2} \text{Tr}(R_{\text{turn}}) = \cos\frac{\psi_1}{2} \cos\frac{\psi_2}{2} \cos(\pi G\gamma) - \sin\frac{\psi_1}{2} \sin\frac{\psi_2}{2} \cos\left(\frac{\pi G\gamma}{3}\right),$$

so that

$$\nu_0(\gamma) = \frac{1}{\pi} \arccos\left[\cos\frac{\psi_1}{2} \cos\frac{\psi_2}{2} \cos(\pi G\gamma) - \sin\frac{\psi_1}{2} \sin\frac{\psi_2}{2} \cos\left(\frac{\pi G\gamma}{3}\right)\right].$$

In the limit $\psi_1 = \psi_2 = 0$, the usual snake-free relation $\nu_s = G\gamma$ is recovered. Two unequal partial snakes shift the tune away from this line, altering the energies of depolarizing resonances but not eliminating them entirely.

Energy dependence during ramp. In the real AGS lattice the snake currents are held approximately constant during the acceleration ramp, so the magnetic fields $B_{\text{snake},i}$ are nearly energy independent. The helical windings are designed such that the *integrated transverse field* through each snake vanishes to a very good approximation,

$$\left| \int \vec{B}_{\perp,i}(s) ds \right| \approx 0,$$

so that the net orbit bending angle from each snake is essentially zero at all energies. This makes the snakes as orbitally transparent as possible in the sense of total deflection. Nevertheless, the local curvature inside the snake is nonzero, and the closed-orbit excursion through the helical channel decreases roughly as

$$x_{\text{max},i}(\gamma) \propto \frac{q}{p} \propto \frac{1}{\beta\gamma},$$

so that the orbit wave in the snake, shown in Fig. 3.3, shrinks with increasing energy even though the net entrance–exit angle remains zero.

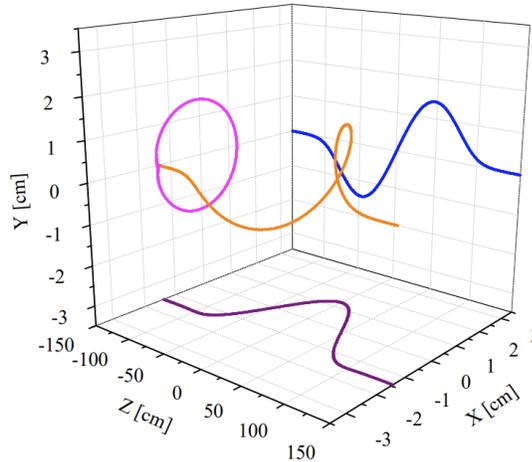


Figure 3.3: Simulated orbit through the warm snake at 2 GeV, from [1].

At the same time, the shrinking orbit excursion and the changing focusing and coupling inside the snake imply that the *orbital* dynamics remain highly energy dependent. The local beta functions, dispersion, and transverse coupling through the snakes all evolve during the ramp, and these changes must be taken into account in lattice design, matching, and orbit correction.

Orbital coupling and field realism. From a beam dynamics standpoint, snakes act as localized regions of intense transverse coupling between horizontal and vertical motion. The helical field geometry means that a horizontal displacement induces a vertical kick and vice versa, requiring careful local compensation to preserve emittance and optical matching into RHIC.

The field modeling challenge is acute for these magnets. Each snake consists of helical dipole windings wrapped around H-shaped iron laminations that pro-

duce a field rotating along the beam direction:

$$B_x(s) = B_{\text{snake}} \cos(k_{\text{snake}}s), \quad (3.5)$$

$$B_y(s) = B_{\text{snake}} \sin(k_{\text{snake}}s), \quad (3.6)$$

where $k_{\text{snake}} = \pi/(2L_{\text{snake}})$ is chosen to yield a nominal 90° spin rotation over the magnet length L_{snake} . In reality, solenoidal field components, helical multipole errors, and end effects introduce significant deviations from this idealized model, requiring detailed field maps for accurate spin-orbit tracking.

3.1.3 The Impact of Diagnostic Limitations

A final factor that amplifies the field modeling challenge is the limited diagnostic instrumentation available in the AGS. There are 72 beam position monitors (BPMs) distributed around the 807-meter circumference [40], but are currently only capable of measuring 100-turn averages of orbit distortions due to digital electronics limitations. This critical input for understanding imperfection resonance strengths is difficult to measure and also subject to large uncertainties due to the scarcity of beam loss monitors. On the other hand, there are up to 2 dedicated turn-by-turn BPMs, one of which is used for measuring tune through the coherent response of the beam to a single-turn fast-kicker. This can be used to capture and measure oscillations due to the mismatch of the beam during the injection process.

Emittance measurements rely primarily on ion profile monitors (IPMs) that infer beam size from the ionization of residual gas by the beam. These devices are highly sensitive to beam position, gap potential, and vacuum conditions, making quantitative emittance determination challenging. An existing electron-

collecting IPM (eIPM) system offers the possibility of turn-by-turn emittance measurements due to the fast response rate of electrons released by ionization events [41], however this is not being generally maintained. The eIPM system would enable us to observe beam instabilities in real-time.

Polarimetry in the AGS is challenging due to the severe trade-off between accuracy and rapidness. The primary polarimeter uses a carbon wire target and intercepts a small fraction of the beam, making continuous polarization monitoring during acceleration impossible [42]. Measurements are somewhat destructive and are typically made at fixed energies by collecting events from many cycles. This is approximately 25 acceleration ramps to achieve a relative uncertainty of 5%, about 2 mins long, leaving significant gaps in the energy dependence of polarization survival.

These diagnostic limitations place a heavy burden on simulation: we must rely on computational models to understand aspects of beam behavior that cannot be directly measured. This reliance makes the accuracy of field models paramount. Any systematic errors in field representation will propagate into predictions of orbit distortions, emittance evolution, and polarization survival, potentially leading to incorrect conclusions about machine performance and limiting the effectiveness of proposed corrections.

The combination of complex magnetic geometries, partial snake effects, and speed-limited diagnostics creates an environment where careful field modeling is not merely desirable but essential. The remainder of this chapter describes the evolution from inadequate models to Maxwell-consistent representations that finally provide reliable foundations for symplectic integration of spin-orbital dynamics studies in the AGS.

3.2 From Generalized Gradient Expansions to Symplectic Maps

The history of accelerator physics is replete with examples of approximate methods that work well in their intended domain but fail when pushed beyond their assumptions. Thin-lens approximations break down for long magnets or large transverse offsets. Linear transfer matrices fail in the presence of strong sextupoles or octupoles. Paraxial approximations fail for large-angle trajectories. The generalized gradient method for modeling combined-function magnets represents another entry in this catalog of approximations that seemed reasonable but ultimately proved inadequate for the demanding requirements of hadron spin dynamics.

3.2.1 Manifestations of Non-Symplectic Tracking

Modeling with non-Maxwell-consistent field maps induces emittance growth and instability on the $O(10^3)$ -turn timescale relevant to AGS injection. Emittance growth manifests as phase-space area increase in the absence of non-conservative processes, indicating broken symplecticity. Tracking at constant injection energy for 10^4 turns highlights stark differences between field maps and symplectic fits: the latter remain stable over the $\sim 10^5$ -turn ramp timescale.

In field-map models, large-amplitude particles experience tune drift, hitting the $3/4$ orbital resonance and causing emittance growth, most pronounced at injection [43]. Thin-lens or “effective matrix” approximations avoid some instabilities but neglect higher-order field effects such as chromaticity.

3.2.2 Maxwell Multipole Expansions

We adopt Maxwell-consistent solutions of vector potentials and fields using cylindrical multipole expansions that are tracked through using symplectic integration [44]. For a given azimuthal index m and longitudinal harmonics k_n the field components read

$$\begin{aligned}
 B_\rho &= \text{Re} \left[\sum_{n=-N/2}^{N/2-1} \frac{1}{2} b_m(n) e^{ik_n z} \cos(m\theta - \theta_{0m}) \left(I_{m-1}(k_n \rho) + I_{m+1}(k_n \rho) \right) \right] \\
 B_\theta &= \text{Re} \left[\sum_{n=-N/2}^{N/2-1} \frac{-1}{2} b_m(n) e^{ik_n z} \sin(m\theta - \theta_{0m}) \left(I_{m-1}(k_n \rho) - I_{m+1}(k_n \rho) \right) \right] \\
 B_z &= \text{Re} \left[\sum_{n=-N/2}^{N/2-1} i b_m(n) e^{ik_n z} \cos(m\theta - \theta_{0m}) I_m(k_n \rho) \right]
 \end{aligned}$$

With 300 terms the RMS deviation of the fit from field data is $< 1\%$ over a $2 \text{ cm} \times 2 \text{ cm}$ aperture (and $< 5\%$ over $10 \text{ cm} \times 10 \text{ cm}$). The first-order transfer matrix satisfies symplecticity to better than 10^{-12} . Note that realistic sextupole components from the fringes are included in the main magnet bodies to capture chromaticity; these substantially affect optics and must be modeled.

One vector potential that can generate these field components is:

$$\begin{aligned}
 A_\rho &= (1 - \delta_{m,0}) \text{Re} \left[\sum_{n=-N/2}^{N/2-1} \frac{-i\rho}{2m} b_m(n) e^{ik_n z} \cos(m\theta - \theta_{0m}) \left(I_{m-1}(k_n \rho) + I_{m+1}(k_n \rho) \right) \right] \\
 A_\theta &= \delta_{m,0} \text{Re} \left[\sum_{n=-N/2}^{N/2-1} \frac{i}{k_n} b_0(n) e^{ik_n z} I_1(k_n \rho) \right] \\
 A_z &= (1 - \delta_{m,0}) \text{Re} \left[\sum_{n=-N/2}^{N/2-1} \frac{-i\rho}{m} b_m(n) e^{ik_n z} \cos(m\theta - \theta_{0m}) I_m(k_n \rho) \right]
 \end{aligned}$$

3.2.3 Construction of Symplectic Maps

Given a Maxwellian field representation, the next step is to construct a symplectic map that advances particles from one location to another. For elements with simple field distributions, analytical expressions for the map can sometimes be derived. For complex elements, numerical integration is necessary, but the integration method must preserve symplecticity.

Paraxial symplectic formulation. Let (x, p_x, y, p_y, δ) denote the canonical variables with path length s as the independent variable, where $\delta \equiv (p - p_0)/p_0$ and momenta are normalized by the reference momentum p_0 . In current-free regions, the magnetic field can be expressed as $\vec{B} = \nabla \times \vec{A}$, where \vec{A} is the vector potential. For longitudinally uniform fields, one often chooses a purely longitudinal potential $\vec{A} = (0, 0, A_s(x, y))$, which yields a closed multipole expansion. In more general situations, $\vec{A}(x, y, s) = (A_x, A_y, A_s)$ may vary nonlinearly and explicitly with s .

To maintain a Hamiltonian description without implicit equations, the *paraxial approximation* assumes small transverse momenta $|p_{x,y}| \ll 1 + \delta$, leading to the quadratic, explicitly s -dependent Hamiltonian

$$H(x, y, p_x, p_y; s) = \frac{(p_x - a_x)^2 + (p_y - a_y)^2}{2(1 + \delta)} - a_s - H_0(\delta), \quad (3.7)$$

where $\vec{a} = \frac{q}{p_0 c} \vec{A}$.

Because Eq. (3.7) is already quadratic in the canonical momenta, it admits linear canonical transformations that remain symplectic to the same order as the paraxial expansion. In particular, the transverse momentum shifts may be

symplectically absorbed by the Lie transformations generated by

$$G_x(x, y, s) = \int a_x(x, y, s) dx, \quad G_y(x, y, s) = \int a_y(x, y, s) dy, \quad (3.8)$$

yielding

$$e^{iG_x(s)} : p_x \mapsto p_x - a_x(x, y, s), \quad x \mapsto x, \quad (3.9)$$

$$e^{iG_y(s)} : p_y \mapsto p_y - a_y(x, y, s), \quad y \mapsto y. \quad (3.10)$$

These transformations are symplectic for any fixed s because they represent canonical momentum shifts. When the potentials depend explicitly on s , the generating functions also acquire an explicit s -dependence, so that the transformed Hamiltonian gains an additional term

$$H' = e^{iG(s)} H + \frac{\partial G(s)}{\partial s}.$$

The correction $\partial_s G = \int \partial_s a_x dx + \int \partial_s a_y dy$ represents induction-like effects associated with longitudinal field variation. For smooth, quasi-static fields, these terms are of higher order than the paraxial truncation and can be neglected; in rapidly varying fields (wigglers, solenoids, or fringe regions) they can be included explicitly or treated through midstep evaluation.

The overall effect is that each kinetic term can be symplectically recast in unshifted form:

$$:\frac{(p_x - a_x)^2}{2(1+\delta)}: \mapsto e^{-i\int a_x dx} : \frac{p_x^2}{2(1+\delta)} : e^{i\int a_x dx}, \quad (3.11)$$

$$:\frac{(p_y - a_y)^2}{2(1+\delta)}: \mapsto e^{-i\int a_y dy} : \frac{p_y^2}{2(1+\delta)} : e^{i\int a_y dy}. \quad (3.12)$$

Noncommutativity between the x - and y -dependent transformations generates terms proportional to $B_z = \partial_x a_y - \partial_y a_x$, which again enter only at higher order than retained in the paraxial model.

Equations of motion. Including the longitudinal pair (z, δ) , Hamilton's equations generated by Eq. (3.7) are

$$\frac{dx}{ds} = \frac{\partial H}{\partial p_x} = \frac{p_x - a_x(x, y, s)}{1 + \delta}, \quad (3.13)$$

$$\frac{dy}{ds} = \frac{\partial H}{\partial p_y} = \frac{p_y - a_y(x, y, s)}{1 + \delta}, \quad (3.14)$$

$$\frac{dz}{ds} = \frac{\partial H}{\partial \delta} = -\frac{(p_x - a_x)^2 + (p_y - a_y)^2}{2(1 + \delta)^2}, \quad (3.15)$$

$$\frac{dp_x}{ds} = -\frac{\partial H}{\partial x} = -\frac{1}{1 + \delta} \left[(p_x - a_x) \partial_x a_x + (p_y - a_y) \partial_x a_y \right] + \partial_x a_s(x, y, s), \quad (3.16)$$

$$\frac{dp_y}{ds} = -\frac{\partial H}{\partial y} = -\frac{1}{1 + \delta} \left[(p_x - a_x) \partial_y a_x + (p_y - a_y) \partial_y a_y \right] + \partial_y a_s(x, y, s), \quad (3.17)$$

$$\frac{d\delta}{ds} = -\frac{\partial H}{\partial z} = 0, \quad (3.18)$$

where all derivatives of a_i are taken at fixed s , and the explicit s -dependence of the potentials enters parametrically through their arguments $a_i(x, y, s)$. In the absence of RF cavities or time-dependent potentials, H does not depend on z and δ remains constant. The longitudinal equation dz/ds encodes the path-length (or arrival-time) slip induced by transverse motion and energy offset.

Symplectic kick–drift–kick (KDK). Splitting $H = T + V$ with

$$T(p_x, p_y, \delta) = \frac{p_x^2 + p_y^2}{2(1 + \delta)}, \quad V(x, y, s, \delta) = -a_s(x, y, s), \quad (3.19)$$

yields a second-order, time-reversible symplectic map over a step Δs :

$$\text{(Half-kick)} \quad p_x \leftarrow p_x - \frac{\Delta s}{2} \partial_x a_s(x, y, s), \quad (3.20)$$

$$p_y \leftarrow p_y - \frac{\Delta s}{2} \partial_y a_s(x, y, s); \quad (3.21)$$

$$\text{(Drift)} \quad x \leftarrow x + \frac{\Delta s}{2} \frac{p_x - a_x}{1 + \delta}, \quad (3.22)$$

$$y \leftarrow y + \frac{\Delta s}{2} \frac{p_y - a_y}{1 + \delta}, \quad (3.23)$$

$$z \leftarrow z - \Delta s \frac{p_x^2 + p_y^2}{2(1 + \delta)^2}, \quad (3.24)$$

$$y \leftarrow y + \frac{\Delta s}{2} \frac{p_y - a_y}{1 + \delta}, \quad (3.25)$$

$$x \leftarrow x + \frac{\Delta s}{2} \frac{p_x}{1 + \delta}; \quad (3.26)$$

$$\text{(Half-kick)} \quad p_y \leftarrow p_y - \frac{\Delta s}{2} \partial_y a_s(x, y, s + \Delta s), \quad (3.27)$$

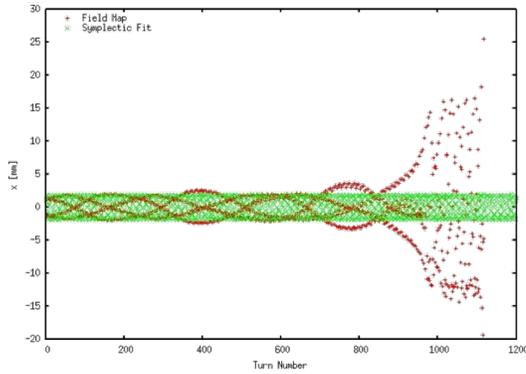
$$p_x \leftarrow p_x - \frac{\Delta s}{2} \partial_x a_s(x, y, s + \Delta s). \quad (3.28)$$

The vector potential and its derivatives are evaluated analytically from the cylindrical multipole expansion at the particle's local (x, y, s) . The scheme is symplectic by construction and time-reversible by symmetry, ensuring bounded phase-space errors over long propagation. Higher-order symplectic integrators also follow from symmetric Yoshida compositions [45] of the 2nd-order KDK map, but the basic integration steps are missing nonlinear details due to the approximations made. Interfaces between elements, including longitudinally varying edges or fringe fields, are implemented as explicit canonical maps to preserve global symplecticity.

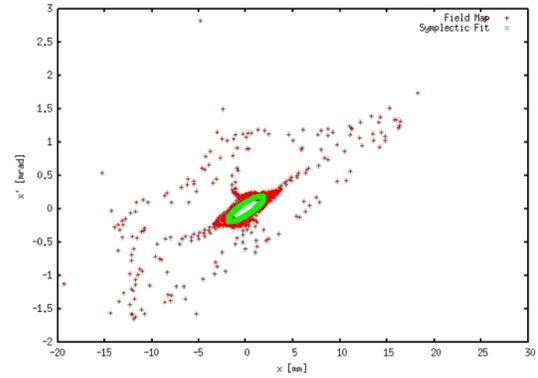
3.2.4 Validation and Performance

We validate the symplectic tracking with the cylindrical field map expansions against grid field tables from finite-element simulations of the AGS snakes and quantify performance improvements.

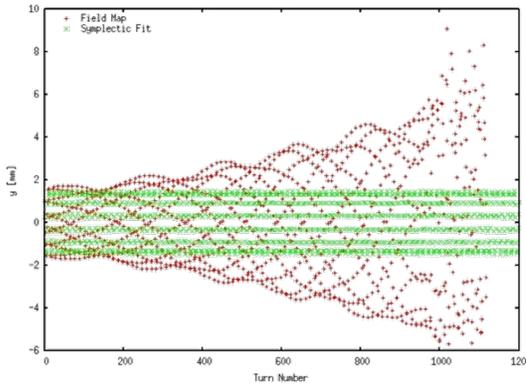
Field reconstruction and tracking through the ring with snakes using cylindrical expansions closely match the initial grid field integration results, as shown in Figs. 3.4a through 3.4d. However, the grid field integration quickly destabilizes the motion, whereas the symplectic integration remains stable as expected.



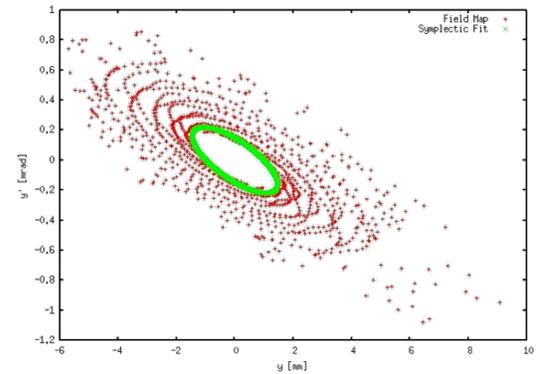
(a) Horizontal particle position vs time.



(b) Horizontal phase space.



(c) Vertical particle position vs time.



(d) Vertical phase space.

Figure 3.4: Stability of phase space motion with grid-table integration (red) vs symplectic tracking of cylindrical expansion (green)

Snake strengths computed from GG maps agree with grid tables and tracking-time comparisons show a $\sim 1000\times$ speedup using GG-derived Taylor maps versus grid-field Runge–Kutta tracking, and approximately $\sim 10\times$ speedup for 2nd order symplectic integration with 300 terms and the same number of steps, c.f. Tables 3.1 and 3.2

Table 3.1: AGS Snake Strength Calculation Results

Snake	GG map	symplectic	grid table
warm	5.32%	5.28%	5.28%
cold	9.77%	10.32%	10.31%

Table 3.2: AGS Snakes Tracking Times (seconds)

Snake	GG map	symplectic	grid table
warm	1.7×10^{-5}	7.6×10^{-3}	2.2×10^{-3}
cold	1.98×10^{-5}	7.9×10^{-3}	2.1×10^{-3}

3.3 Simulated Field Models

The Maxwell-consistent field expansions described in the previous section represent a substantial improvement over generalized gradient methods, restoring symplecticity and eliminating some systematic errors. However, they remain approximations to the true electromagnetic field distributions in real accelerator magnets. Multipole expansions assume that fields can be accurately represented by low-order polynomial terms near the reference trajectory, an assumption that breaks down in regions of strong nonlinearity, near magnet edges, or in elements with complex three-dimensional geometries.

The ultimate goal of field modeling for the EIC era is to incorporate fields derived directly from the actual magnet geometries: CAD models of pole faces, coil configurations, and iron return paths, fed into electromagnetic finite element codes like Opera, COMSOL, or ANSYS Maxwell, or custom meshing and solving libraries. The goal is to solve Maxwell’s equations numerically on fine meshes, producing field maps that capture every subtlety of the real design, possibly including saturation effects in iron, higher-order multipoles from finite pole widths, three-dimensional fringe field distributions, and even interference effects between nearby elements if problematic.

3.3.1 Field Modeling Workflow

There are many choices of finite-element field solvers that can be used for modeling fields in 3D geometries. The relevant modules solve the static and quasi-static Maxwell equations using finite element analysis on unstructured tetrahedral meshes. The typical workflow for modeling an accelerator magnet proceeds through several stages:

CAD geometry import. The magnet geometry is defined in a CAD system, including detailed pole face profiles, coil positions and cross-sections, iron yoke shapes, and any additional elements like correction coils or trim windings. Modern finite-element solvers can import geometry directly from standard CAD formats (STEP, IGES) or sometimes have their own interface for direct geometry construction.

Material property assignment. Each geometric region is assigned appropriate material properties: permeability curves for iron components (including saturation effects), conductor properties for coils, and vacuum for field regions. Grid solvers' material library includes standard magnet steels, but custom B-H curves can usually be imported from measurements or manufacturer specifications.

Mesh generation. The finite element mesh is generated adaptively, with finer elements near material boundaries and in regions of high field gradients. Mesh quality is critical. Poorly shaped elements or insufficient refinement lead to inaccurate field solutions. Modern meshing algorithms handle the complex geometries of accelerator magnets reasonably well, but manual refinement is often necessary in critical regions like pole tips or coil edges.

Excitation specification and solution. The current in each coil set is specified, along with boundary conditions at the outer boundaries of the solution domain (typically far from the magnet where fields approach zero). The finite-element solver then solves the nonlinear system of finite element equations, iterating if necessary to account for saturation effects in iron. The solution produces the vector potential \vec{A} at each mesh node, from which field components can be derived.

Field map export and post-processing. The computed fields are exported on a regular grid covering the region of interest for beam tracking. Typical grid spacings range from sub-millimeter in the core beam region to centimeter-scale in the outer regions. Post-processing includes extraction of integrated field quantities (dipole strength, gradient, etc.), computation of multipole components through harmonic decomposition, and visualization of field patterns to verify physical reasonableness.

3.3.2 From Field Maps to Vector Potentials

The output of finite-element solvers is naturally a field map: $\vec{B}(x, y, z)$ evaluated on a three-dimensional grid. However, symplectic integration requires the vector potential \vec{A} rather than the field itself. This requirement poses a nontrivial challenge because the transformation from \vec{B} to \vec{A} is not unique (gauge freedom) and requires solving a differential equation. While it is possible to directly solve for the vector potential rather than the field, this is not usually done during new magnet commissioning, so it would be beneficial to adapt a process for working indirectly through the B-fields.

Given \vec{B} , we need to find \vec{A} such that $\vec{B} = \nabla \times \vec{A}$. In the Coulomb gauge ($\nabla \cdot \vec{A} = 0$), this reduces to solving

$$\nabla^2 \vec{A} = -\nabla \times \vec{B}. \quad (3.29)$$

In current-free regions where $\nabla \times \vec{B} = 0$, this becomes Laplace's equation $\nabla^2 \vec{A} = 0$, which admits many solutions depending on boundary conditions.

One practical approach is to integrate the defining relation $\vec{B} = \nabla \times \vec{A}$ along a chosen direction. For example, in the gauge $A_z = 0$, the curl components read

$$B_x = -\frac{\partial A_y}{\partial z}, \quad B_y = \frac{\partial A_x}{\partial z}, \quad B_z = \frac{\partial A_y}{\partial x} - \frac{\partial A_x}{\partial y}. \quad (3.30)$$

Integrating the first two relations with respect to z from a reference plane z_0 gives

$$A_x(x, y, z) = A_x(x, y, z_0) + \int_{z_0}^z B_y(x, y, z') dz', \quad (3.31)$$

$$A_y(x, y, z) = A_y(x, y, z_0) - \int_{z_0}^z B_x(x, y, z') dz'. \quad (3.32)$$

The functions $A_x(x, y, z_0)$ and $A_y(x, y, z_0)$ act as integration constants (depending only on x and y) and are constrained by the remaining curl component,

$$\frac{\partial A_y}{\partial x} - \frac{\partial A_x}{\partial y} = B_z, \quad (3.33)$$

together with any additional gauge or boundary conditions.

An alternative approach, more robust for gridded data, is to perform a least-squares fit of multipole expansions to the field map data. The Maxwellian multipole forms discussed in Section 3.2 have known vector potentials that can be

derived analytically. By fitting multipole coefficients to match the finite-element field data, we simultaneously obtain both the field representation and its corresponding vector potential.

For the AGS partial snake magnets, preliminary fits using cylindrical and cartesian decompositions expansions achieved agreement with Opera field maps to better than 1% over transverse regions spanning ± 20 mm from the central axis. This accuracy is sufficient for most beam dynamics studies, though higher-order terms would be needed for particles at large amplitudes or for detailed studies of nonlinear resonances.

3.3.3 Symplectic Integration with Interpolated Fields

Once a vector potential representation is available, symplectic integration can proceed using standard methods. The Hamiltonian for a particle in a magnetic field is:

$$H = \sqrt{(\vec{p} - q\vec{A})^2 + m^2c^2} \quad (3.34)$$

Hamilton's equations give:

$$\frac{d\vec{x}}{dt} = \frac{\partial H}{\partial \vec{p}} = \frac{\vec{p} - q\vec{A}}{E/c^2} \quad (3.35)$$

$$\frac{d\vec{p}}{dt} = -\frac{\partial H}{\partial \vec{x}} = q(\nabla\vec{A}) \cdot \frac{\vec{p} - q\vec{A}}{E/c^2} \quad (3.36)$$

where $E = \sqrt{p^2c^2 + m^2c^4}$ is the energy. These equations are integrated using

a symplectic method like the Verlet algorithm or higher-order Yoshida variants.

The challenge is that \vec{A} is now known only on a discrete grid, requiring interpolation. Standard interpolation methods (trilinear, tricubic splines) are not exactly Maxwellian. However, if the interpolation error is small enough, the symplectic motion remains approximately Maxwellian over finite time intervals.

For practical AGS tracking with typical grid spacings of 1-2 mm and symplectic integrator step sizes chosen to resolve the field variation, the phase space volume drift is acceptably small. Ensemble tracking over 100,000 turns shows emittance growth less than 0.5%, comparable to the measurement uncertainty in real AGS emittance measurements and smaller than space-charge effects.

An alternative approach that preserves exact symplecticity is to fit the gridded vector potential data with analytical functions designed for symplectic integration. Polynomial splines with carefully chosen knot points can represent the vector potential while allowing analytical evaluation of derivatives.

3.3.4 Validation Against Measurements

The ultimate test of any field model is comparison with experimental measurements. For the AGS, direct field mapping of the combined-function magnets was performed during the original construction in the 1960s. Within the accuracy of available data, simulated fields show decent agreement with measurements. The integrated dipole field strength agrees to better than 0.5%, well within the uncertainty of the measurement technique. Gradients inferred from

field measurements during commissioning agree with simulation predictions to within 2-3%.

More telling validation comes from orbit response measurements. When corrector dipoles are energized, the resulting orbit distortion provides information about the beta function and phase advance around the ring. Comparing measured orbit responses to predictions from map-based models shows agreement typically at the 5-10% level, with discrepancies attributable to uncertainties in corrector strengths, BPM calibrations, and the inevitable mismatch between the ideal design and the as-built machine.

For spin dynamics, direct validation is more challenging due to the limited polarimetry in the AGS. Measured polarization transmission from injection to extraction typically ranges from 60% to 85% depending on beam intensity, bunch structure, and other operational parameters [33].

The most convincing validation may be indirect: the elimination of pathological behaviors seen with earlier models. Cylindrical decomposition simulations of snakes do not predict the dramatic injection losses, spurious depolarization, or orbit instabilities that plagued grid table [43] and generalized gradient models. Resonance strengths match theoretical estimates, orbit distortions evolve as expected from lattice optics, and spin dynamics show the signatures of partial snake effects predicted by analytical theory. In fact, while simulations using Opera-based field models of the partial snakes are entirely unstable at injection energy, the Maxwellian cylindrical expansions of the snakes are successfully able to predict spin resonance magnitudes and phases within less than 10%.

3.3.5 Future Directions

Looking forward to the EIC era, the goal is a fully automated pipeline from magnet CAD designs to symplectic tracking simulations. Engineers design magnets in SolidWorks or other CAD systems, optimizing geometry for field quality, mechanical constraints, and manufacturability. Those designs are automatically exported to a finite-element solver for field solution, with material properties and excitation currents drawn from a central database. The resulting field maps are processed into multipole fits or interpolated vector potentials, then integrated into tracking codes that simulate entire acceleration chains from source to collision point.

This integrated workflow enables rapid iteration on magnet designs. If tracking simulations reveal excessive polarization loss due to a particular resonance, engineers can adjust pole face shapes or add correction coils and immediately see the impact. If manufacturing tolerances are found to be critical for certain multipole components, field quality specifications can be tightened before fabrication begins.

The technological pieces for this vision are largely in place. CAD-to-field interfaces are mature and reliable. Field map processing and multipole fitting are computationally straightforward. Symplectic tracking codes with interpolated field support exist in both established packages and emerging frameworks. The challenge is integration: building the software infrastructure to connect these pieces seamlessly and establishing the workflows and standards that allow multiple groups to collaborate effectively.

In summary, finite-element solvers validates realism while cylindrical ex-

pansions enable practical, stable, symplectic long-term tracking used throughout this chapter.

3.4 Validation and Lattice Matching to Machine Settings

With symplectic field models in place, we fit the AGS lattice to match machine settings across the acceleration ramp, following the established methodology. Three optimizations are performed at 100 energy values between 2.5 GeV and 25 GeV: (i) compensate orbit defects due to the snakes locally before and after each snake using dipole correctors and backleg windings; (ii) compensate snake optical focusing effects locally via seven dedicated compensation quadrupoles; and (iii) fit the tune quadrupole currents to match the measured tunes along the user ramp, fixed above γ -transition at $Q_x, Q_y = (8.72, 8.985)$. This is shown in Fig. 3.5.

Interpolating between all 100 settings during the ramp keeps the beam centered and reproduces the measured tunes throughout acceleration. The matching problem is under-determined: the number of truly independent knobs, thin quadrupoles, compensation quadrupoles, backleg windings on the main magnets, and correctors, is fewer than the desired local Twiss targets across the snakes. Nevertheless, there are sufficient controls to obtain a closed lattice with reasonable optics and minimal beta beating in practice. By relaxing exact control of the optics, the independent knobs provide a degenerate set of approximate solutions.

Polarization tracking is performed only *above* γ -transition and does not use the optics-matching process. The normalized fields used for tracking are static

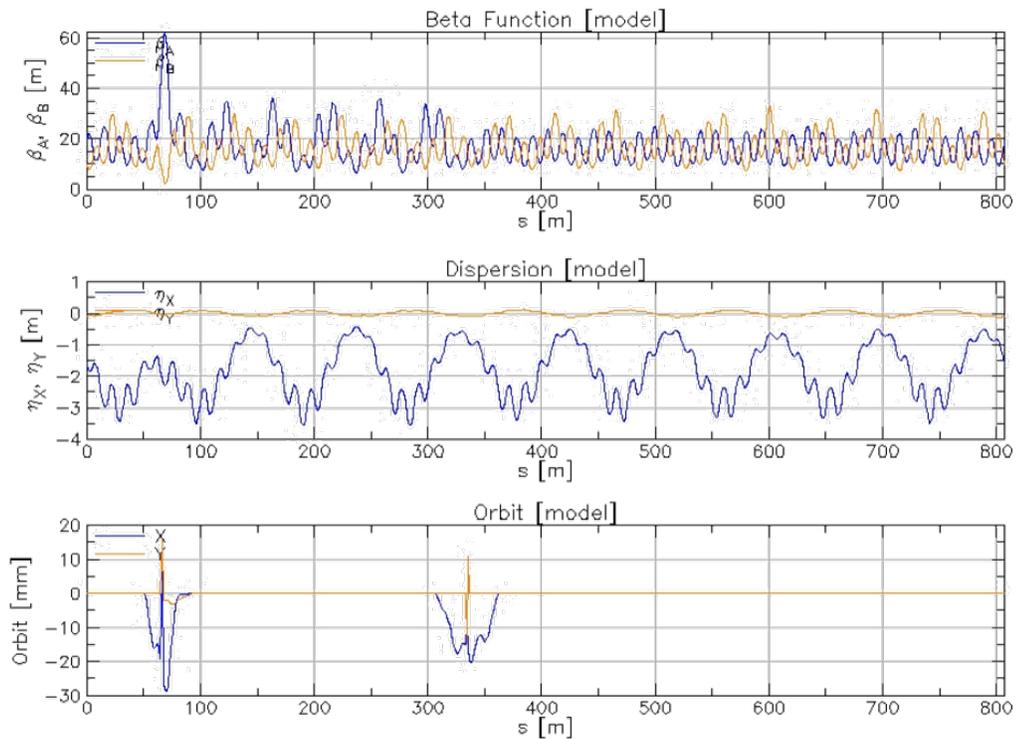


Figure 3.5: Optics and orbit at 2.5 GeV showing closed-orbit bumps around both snakes after local compensation.

lattice fields, while the partial snakes use static unnormalized fields; the optics fits serve solely to reproduce measured tunes and close the lattice. As seen in Fig. 3.6, tracking a single proton with normalized action of 1.7 mm mrad from 10 to 23 GeV yields 97% polarization transmission through the AGS cycle. This amplitude sits well within the core of the beam and is representative of the beam motion.

Finally, the matched lattice was used to predict appropriate corrections for partial-snake resonances with betatron coupling.

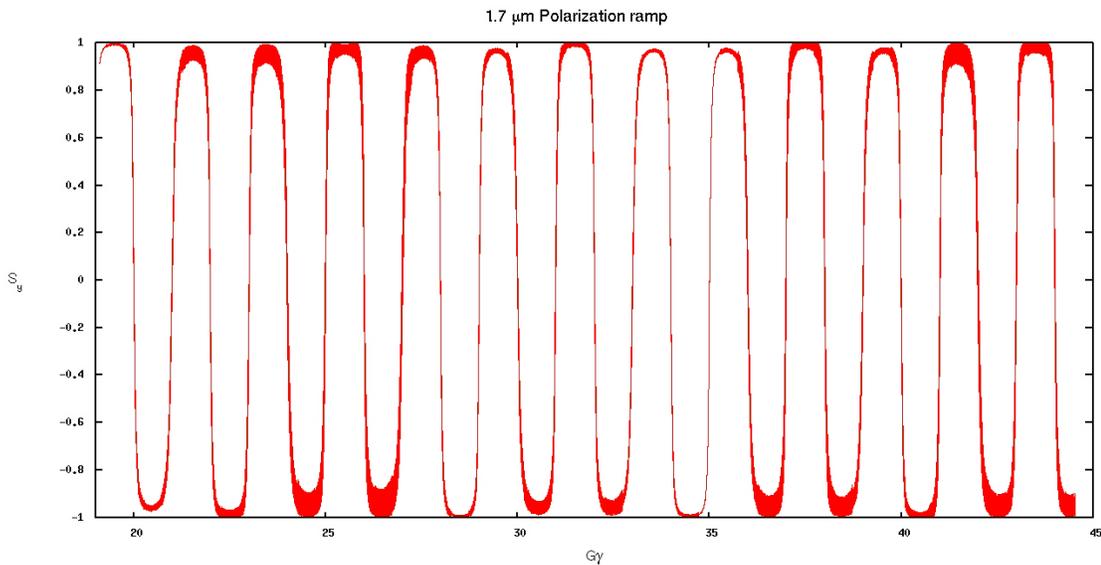


Figure 3.6: Polarization transmission through the AGS cycle: **97%** achieved with symplectic snake models and optics compensation.

3.5 The Big Picture

The journey through field modeling challenges in the AGS recounted in this chapter exemplifies a broader theme that runs throughout this dissertation: the indispensability of mathematical rigor in computational accelerator physics. It is tempting to view field modeling as a frivolous matter, a box to be checked on the path to addressing “real” physics questions about spin dynamics and polarization optimization. This chapter demonstrates that such a view is fundamentally mistaken.

Field modeling is not preparatory to the physics. It is part of the physics. The requirement that electromagnetic fields satisfy Maxwell’s equations is not a technicality but a fundamental constraint that shapes the possible dynamics. The preservation of symplecticity in numerical integration is not an optional performance enhancement but a necessary condition for meaningful long-term

tracking. When these mathematical principles are violated, the resulting simulations do not merely lose accuracy. They lose physical meaning.

The consequences of inadequate field models extended beyond numerical artifacts. Early simulations using grid tables and generalized gradient model predicted particle losses that, if real, would have made beam operation in the AGS impossible. These predictions might have led to unnecessary hardware modifications, wasted beam time investigating non-existent problems, or even abandonment of the snakes altogether. The transition to Maxwell-consistent symplectic models eliminated these false positives and revealed the actual, more manageable challenges facing AGS polarization.

Looking forward to subsequent chapters, the field modeling framework developed here provides the foundation for all beam dynamics studies. Chapter 4 relies on accurate field models to predict polarization survival in RHIC and the HSR under various snake configurations, then uses these models to design and evaluate resonance correction schemes based on orbit bumps and beta matching. Chapter 5 couples machine learning algorithms to symplectic tracking engines built on the Maxwellian field representations.

The broader lesson extends beyond the AGS to all future accelerator projects. The EIC's demanding polarization requirements, complex magnet geometries, and need for multiple species with different gyromagnetic ratios will push field modeling to its limits. The symplectic framework established here provides a template for handling these challenges without repeating the false starts and systematic errors that plagued earlier efforts.

Perhaps most fundamentally, this chapter illustrates that computational ac-

celerator physics is not merely "running codes." It is physics, requiring the same rigor in formulation, the same care in approximation, and the same scrutiny of assumptions that characterize theoretical work. The electromagnetic fields in accelerator magnets are as real as the quantum fields in particle physics, and they deserve the same respect for their mathematical structure and physical content. When we respect that structure by ensuring our models inherently satisfy Maxwell's equations and our integrators exactly preserve symplecticity, we are rewarded with simulations that capture reality rather than artifacts.

CHAPTER 4

SPIN DYNAMICS IN RHIC AND THE HADRON STORAGE RING

This chapter examines spin polarization transport in high-energy hadron storage rings, focusing on two complementary facilities that bracket the evolution of polarized hadron beam physics: the Relativistic Heavy Ion Collider (RHIC), which has successfully delivered polarized proton collisions since 2001, and the future Hadron Storage Ring (HSR) [7] for the Electron-Ion Collider (EIC) [46], which will extend polarized beam capabilities to higher luminosities and multiple hadron species. The unifying theme is how Siberian snake configurations govern polarization survival during acceleration, and how resonance structures differ across machine conditions.

We begin with unresolved questions about RHIC's operational resilience following magnet damage [47], move to predictive modeling of the HSR under ideal and realistic conditions, and culminate in the development of novel snake configurations that offer new solutions for helions with challenging spin dynamics. Throughout, the central concern remains polarization survival during energy ramps, where beams must traverse dense forests of depolarizing resonances while maintaining the spin coherence required for physics measurements.

The work presented here draws extensively on the mathematical foundations developed in Chapter 2, particularly spin resonance theory and normal form analysis. The insights gained from modeling these facilities reveal both the remarkable capabilities of Siberian snake technology and its fundamental limitations, motivating the search for new approaches that can extend polarized beam physics to increasingly challenging parameter regimes.

4.1 Early Work: RHIC Snake Re-Optimization

The Relativistic Heavy Ion Collider (RHIC) represents one of the most successful large-scale implementations of *Siberian snake* technology to date, routinely achieving proton beam polarizations of 50–60% at collision energies of 255 GeV (corresponding to $\gamma \approx 270$). This success relies on the installation of two full Siberian snakes in each ring, located diametrically opposite one another and with horizontal spin rotation axes oriented at $\pm 45^\circ$ relative to the longitudinal beam direction. This configuration fixes the closed-orbit spin tune to $\nu_0 = 1/2$, thereby eliminating all first-order imperfection and intrinsic resonances when the orbital tunes are chosen appropriately. RHIC operates at a working point of approximately $(Q_x, Q_y) = (28.695, 29.685)$, which avoids snake and intrinsic resonances to first order.

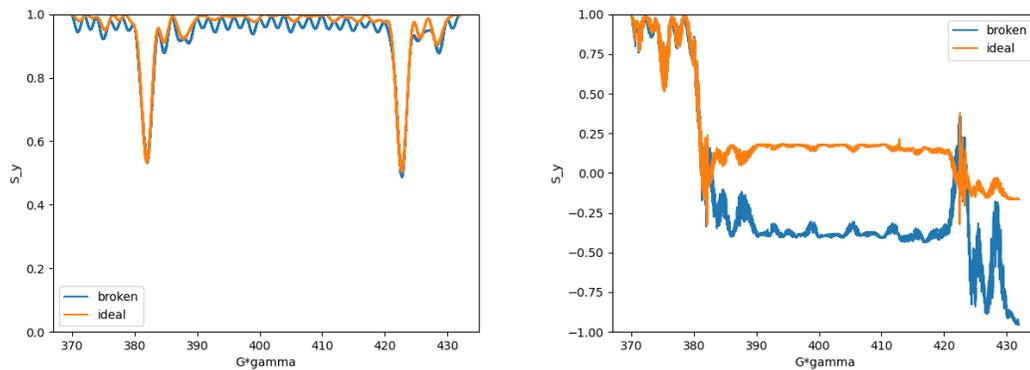
4.1.1 The Snake Damage Event and Operational Response

In December 2021, a series of power-supply failures [47] in one of the Blue ring's Siberian Snakes forced it to operate as a *partial Siberian Snake* with a reduced precession angle and a slightly altered spin rotation axis of $\sim 49^\circ$. To maintain approximate symmetry, the undamaged snake was reconfigured to be complementary. Nevertheless, the closed orbit spin tune was slightly detuned from $1/2$. Unexpectedly, the time-averaged beam polarization during this run was slightly higher in the Blue ring, which operated with two partial snakes, than in the Yellow ring, which retained two fully functional snakes. This surprising robustness of polarization suggested that slightly asymmetric or partial snakes may, under certain conditions, be no more detrimental (and possibly even bene-

ficial) to polarization than perfectly symmetric configurations. This observation motivated a detailed series of spin tracking simulations to investigate the underlying mechanisms.

4.1.2 Simulation Studies and Key Findings

Comprehensive spin tracking simulations [48] of the RHIC Blue ring were carried out to examine polarization transmission in the “dangerous” resonance region $G\gamma \in [370, 432]$, encompassing the strongest intrinsic resonances near $G\gamma = 393 + Q_y$ and $G\gamma = 411 - Q_y$. A sequence of increasingly realistic lattice models was considered, beginning with an ideal lattice having a perfect closed orbit and nominal magnet strengths, followed by a lattice incorporating measured magnet-to-magnet field-strength variations, and finally a configuration including the six interaction-region (IR) closed-orbit bumps of ± 5 mm used during the energy ramp. Each of these models compared a pair of full snakes against a pair of partial snakes, shown in Figs 4.1a and 4.1b.



(a) Tracking a $\sim 1.4\sigma$ particle with $J_y = 5$ mm mrad (b) Tracking a $\sim 2.8\sigma$ particle with $J_y = 20$ mm mrad

Tracking with realistic emittances showed that, in the ideal lattice, both con-

configurations preserved full polarization through the ramp, and no depolarization was observed. Including field-strength variations slightly degraded polarization by tilting the invariant spin field and shifting the spin tune away from $1/2$. Adding the IR bumps further enhanced imperfection-type resonance terms but still did not reproduce the polarization losses observed experimentally. Significant depolarization appeared only after introducing finite RMS lattice misalignments, demonstrating that misalignments, rather than the partial-snake asymmetry itself, are the dominant source of polarization loss. For a 1.4σ particle ($J_y = 5\pi$ mm mrad), polarization remained essentially identical between full and partial-snake configurations, whereas a 2.8σ particle ($J_y = 20\pi$ mm mrad) exhibited a complete spin flip by the end of the ramp in the broken-snake case, highlighting the strong amplitude dependence of depolarization.

4.1.3 Physical Interpretation and Implications

The simulations reveal that partial or asymmetric snakes introduce a small additional resonance-driving term whose phase can partially cancel existing terms. Consequently, modest deviations from symmetry such as the 49° axis used during the 2021 run do not necessarily worsen polarization and may even yield near-ideal performance depending on phase relationships. In contrast, magnet misalignments and closed-orbit distortions have a far greater impact. They tilt the invariant spin field and shift the closed-orbit spin tune away from $\frac{1}{2}$, allowing higher-order spin-orbit resonances to be crossed during the energy ramp. This explains why experimentally observed polarization losses cannot be reproduced in simulations without including RMS misalignments [49].

The practical implication is that orbit control and alignment dominate polarization survival in RHIC. While adjustments to snake rotation angles or axes can offer incremental improvements, they cannot compensate for poor orbit quality. These results provide valuable guidance for future polarized-proton facilities. If RHIC, with only two snakes and nearly perfect symmetry, exhibits such sensitivity to orbit distortions, then machines employing more numerous or asymmetrically placed snakes will require even tighter orbit and alignment tolerances to maintain polarization through acceleration.

4.2 The EIC Hadron Storage Ring: Design and Polarization Challenge

The scientific mandate of the EIC requires the acceleration of novel, highly polarized hadron beams to unprecedented energies. The successful operation of the future EIC is contingent on maintaining high hadron beam polarization up to 275 GeV in the HSR. However, the HSR lattice features a symmetry-breaking interaction region, especially the vertical beta function [50], that excites strong, non-systematic spin resonances, posing a significant threat to polarization preservation.

This section presents a systematic investigation of two complementary strategies to ensure high polarization transmission in the HSR. The first method involves optimizing the vertical betatron phase advance between Siberian snakes to orchestrate a cancellation of depolarizing kicks across the ring. The second, more powerful approach involves optimizing the snake rotation axes to more strongly avoid resonance driving terms [51]. We introduce a novel, highly

symmetric “Doubly Lee-Courant” (DLC) scheme, which enforces a local π spin phase advance across every consecutive pair of snakes. Our analysis reveals a clear performance hierarchy, with the DLC configuration providing an exceptionally robust and energy-insensitive baseline for polarization preservation.

4.2.1 HSR Design Parameters and Snake Configuration

The baseline HSR design incorporates six full Siberian snakes with each snake producing a 180° spin rotation. This six-snake configuration, with $\sim \pi/3$ radians of revolution between each snake, ensures an energy-independent closed-orbit spin tune of $\nu_0 = 1/2$, providing greater redundancy than RHIC’s two snakes.

The HSR operates over a wide energy range, accelerating protons from injection at $\gamma \approx 25.4$ (kinetic energy 23 GeV) to top energy at $\gamma \approx 290$ (kinetic energy 275 GeV). The HSR must accommodate multiple hadron species, protons, deuterons, and helium-3, each with different gyromagnetic ratios and hence different spin dynamics. The fundamental parameter determining spin behavior is the anomalous magnetic moment $G = (g - 2)/2$, which varies significantly across species.

The species-dependent resonance density underscores these challenges. The diversity of hadron species required for the EIC physics program leads to dramatically different resonance behavior, and Table 4.1 summarizes the corresponding linear resonance spectrum.

As shown in the table, helium-3 encounters roughly 2.3 times as many intrinsic resonances as protons during the same energy ramp, with resonance

Table 4.1: Species-Dependent Linear Resonance Spectrum in the HSR

Species	G	max $ G\gamma $	No. of Resonances
p	1.7928	525	1575
D	-0.1430	21	63
${}^3\text{He}^{2+}$	-4.1842	819	2457

spacing $\Delta\gamma = 1/|G|$ of approximately 0.56 for protons and 0.24 for helium-3. This higher resonance density, combined with resonance strengths that scale as $\sim G\sqrt{\gamma}$, makes helium-3 polarization preservation particularly challenging at high energies. For this reason, our optimization studies focus primarily on helium-3, the most demanding case.

4.2.2 Symmetry Breaking and Resonance Driving

The efficacy of any snake matching scheme is determined by the structure of the accelerator lattice. RHIC, the EIC's predecessor, has an approximate 3-fold super-periodicity that plays a crucial role in passively suppressing 2/3 of spin resonances, simplifying the task of polarization preservation. The EIC HSR, in contrast, inherits the RHIC tunnel but incorporates at least one new, large, and vertically asymmetric interaction region (IR) that fundamentally breaks this 3-fold symmetry.

The total resonance strength ϵ can be expressed as a sum of contributions from all insertions (arcs, IRs) around the ring:

$$\epsilon \propto \sum_{j=1}^M \int_{\text{insertion } j} f(s) e^{i\chi_j(s)} ds = \sum_{j=1}^M \mathcal{F}_j e^{i\bar{\chi}_j}, \quad (4.1)$$

where $f(s)$ contains the lattice functions (β_y, k) and $\chi_j(s)$ is the total spin-orbital phase advance. In a symmetric lattice, the complex \mathcal{F}_j are identical for similar

sections, leading to constructive or destructive interference. The HSR's IR introduces very large and distinct \mathcal{F}_j terms that are not balanced by other insertions. This broken symmetry excites a dense spectrum of strong, first-order, intrinsic resonances, making polarization preservation qualitatively more difficult than in RHIC [52].

Figure 4.2 compares the intrinsic resonance spectrum of the more symmetric lattice of RHIC with the denser, less symmetric lattice of the HSR. The HSR exhibits both stronger dominant resonances due to its lower β^* and a higher resonance floor due to its symmetry-breaking IR. Despite RHIC having two low-beta interaction points compared to the single one in the HSR, its overall lattice symmetry is preserved to a much greater degree.

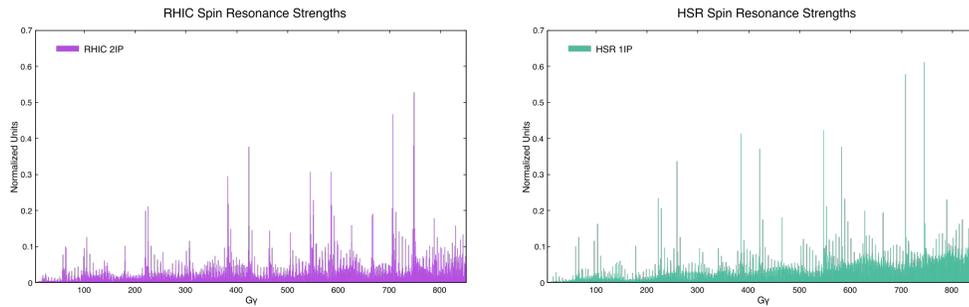


Figure 4.2: Schematic comparison of the intrinsic resonance spectrum of RHIC (left) versus the HSR (right). The HSR's broken symmetry from the complex interaction region creates both stronger systematic resonances and a higher non-systematic resonance floor, making polarization preservation significantly more challenging.

4.2.3 Baseline Tracking and the Need for Optimization

To demonstrate the inadequacy of polarization transmission for a nominal snake configuration, non-perturbative spin tracking simulations were per-

formed using `Bmad`. A beam of 49 helium-3 nuclei was initialized on a 4D phase-space torus with amplitudes corresponding to the beam periphery ($J_x = \sigma_x$, $J_y = 2\sigma_y$) to probe the depolarizing effects on large vertical-amplitude particles. The initial spin of each particle was aligned with its local invariant spin field (ISF). The beam was then accelerated at a realistic rate of $\dot{\gamma} = 0.75/\text{s}$ through the most problematic energy range ($700 \leq G\gamma \leq 800$), where significant polarization loss occurs. To measure polarization loss, we calculate the final value of P_{dyn} , the dynamic quantity describing the instantaneous average polarization after ramping relative to the maximum time-averaged polarization at a fixed energy $P_{\text{lim}} \equiv |\langle \vec{n} \rangle_\phi|$. The relationship of interest is $P_{\text{dyn}} P_{\text{lim}} \equiv |\langle \mathbf{S} \rangle_i|$.

Figure 4.3 shows the result of this baseline tracking with a nominal $\pm 15^\circ$ snake axis configuration. The polarization rapidly deteriorates as the beam crosses strongly resonant regions, particularly around $G\gamma \approx 705$ and $G\gamma \approx 745$, confirming that this baseline setup is insufficient to preserve polarization in the face of the HSR's strong non-systematic resonances. Each colored trace represents the vertical spin component of each particle in the bunch, and the dramatic divergence indicates severe depolarization. This motivated the development of two complementary optimization strategies: Method I (betatron phase advance optimization) and Method II (snake axis optimization).

4.3 Method I: Optimization of Betatron Phase Advance

The first strategy investigated for restoring polarization transmission is the manipulation of the accelerator's orbital dynamics, specifically the vertical betatron phase advance Φ_y between the Siberian snakes [53]. This approach treats

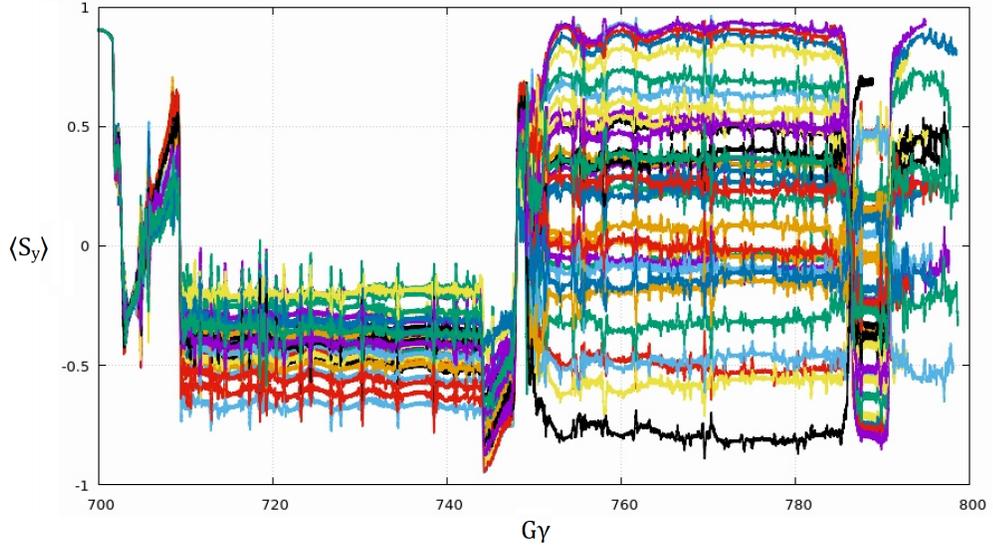


Figure 4.3: Baseline spin tracking of uncooled, round, helion beams of $3 \mu\text{m}$ emittance in the injection HSR lattice with $\pm 15^\circ$ snakes. Each colored trace represents the vertical spin component of each particle in the bunch. The ensemble polarization is lost rapidly as the beam is accelerated through strongly resonant energies around $G\gamma \approx 705$ and $G\gamma \approx 745$.

the snake configuration as fixed (e.g., the $\pm 15^\circ$ axes) and instead adjusts the lattice optics to force the depolarizing spin kicks from different sections of the ring to cancel one another.

4.3.1 Theoretical Basis for Phase Advance Optimization

The core idea is to treat the total first-order resonance strength, represented by the complex-valued spin-orbit coupling integral, as a vector sum in the complex plane. The ring is composed of $2N$ sections between the snakes, and the total integral is the sum of the contributions from each section. A Siberian snake fundamentally alters the spin precession that follows it, transforming the spin basis and inverting the subsequent vertical spin precession. Consequently, the phase of each contribution depends on the accumulated spin phase Ψ and betatron

phase Φ_y up to that point, including the discrete jumps from the snakes.

The snake match condition is achieved when the sum of these complex vectors is zero. Since the snake axes φ_i and arc precession angles Ψ_i are fixed in this method, the only free parameters are the betatron phase advances $\Phi_{y,i}$ between the snakes. By adjusting the quadrupole settings in the arcs and straight sections, we can change the values of $\Phi_{y,i}$, effectively “rotating” each phasor in the complex plane until their sum cancels. It should be noted that it is not physical to artificially change the phases without also making changes to the beta function, which implies that both the magnitude and angle of the complex \mathcal{F}_j are varying. At first, we only consider an artificial change of the phases $\Phi_{y,i}$ while neglecting changes to the beta function. Nevertheless, the HSR lattice is not symmetric, so a numerical optimization will still be required to find a set of $\{\Phi_{y,i}\}$ that satisfies this condition over a specific energy range.

4.3.2 Baseline Behavior and Tune Scans

Before optimization, we diagnose the lattice by performing a “tune scan,” where we compute the maximum time-averaged polarization P_{lim} for large-amplitude particles as a function of the vertical betatron tune Q_y at a fixed energy using Bmad’s stroboscopic averaging program [54]. Fig. 4.4 shows such a scan for the nominal HSR lattice at $G\gamma \approx 704$. The yellow/bright areas indicate high polarization ($P_{\text{lim}} \approx 1$), while the dark bands indicate strong depolarization. These bands are higher-order resonance doublets, which arise from the condition $\nu(J) = k \pm mQ_y$. The splitting of the doublets is a direct measure of the deviation of the amplitude-dependent spin tune (ADST) from the central tune.

The wider the split, the stronger the depolarization and the more difficult it is to find a safe working point.

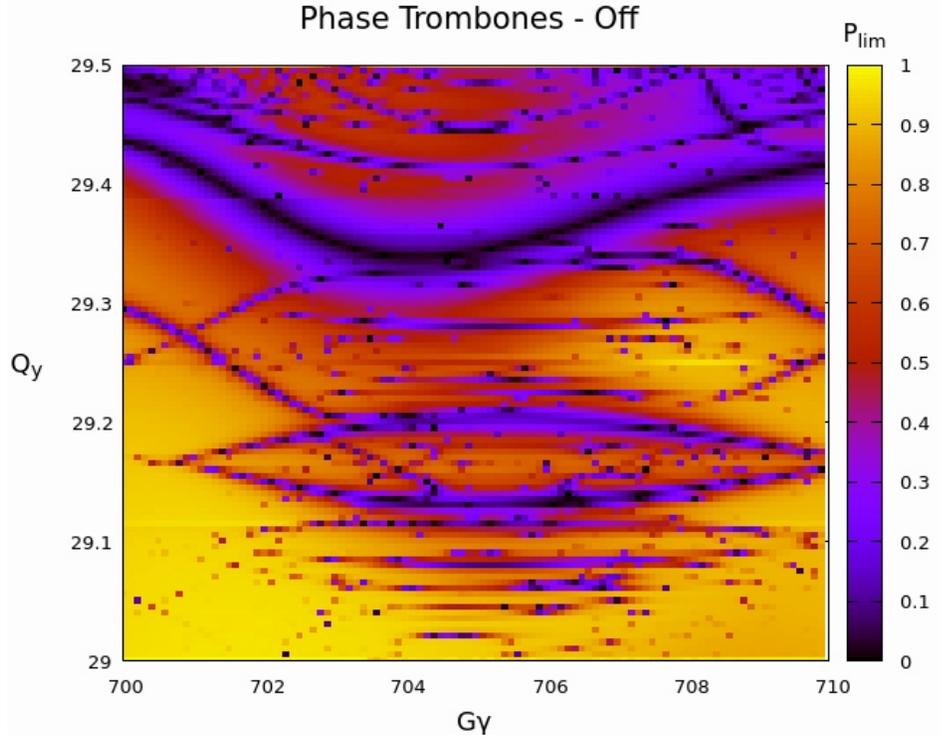


Figure 4.4: Tune scan of the limiting polarization P_{lim} for the baseline HSR lattice at $G\gamma \approx 704$. The dark bands indicate strong higher-order resonance doublets whose splitting reveals the magnitude of the ADST spread. The nominal working tune is marked by the crosshairs. The wide doublet splitting indicates poor polarization preservation with the baseline configuration.

4.3.3 Optimization Results

Our optimization procedure uses the Bmad-based optimizer Tao [55] to vary the six artificial phase advances in the lattice, one per arc. The optimization target is to minimize the spin-orbit coupling integral I^\pm as a proxy for the ADST spread coefficient $|\nu_{1,y}|$, which corresponds to closing the resonance doublets seen in the tune scan. Figure 4.5 shows the tune scan after optimizing the phase advances

at $G\gamma \approx 704$ using idealized “phase trombone” matrix elements. The optimization has successfully squeezed the doublet lines together, creating a much larger region of stable polarization transmission compared to the baseline (Fig. 4.4).

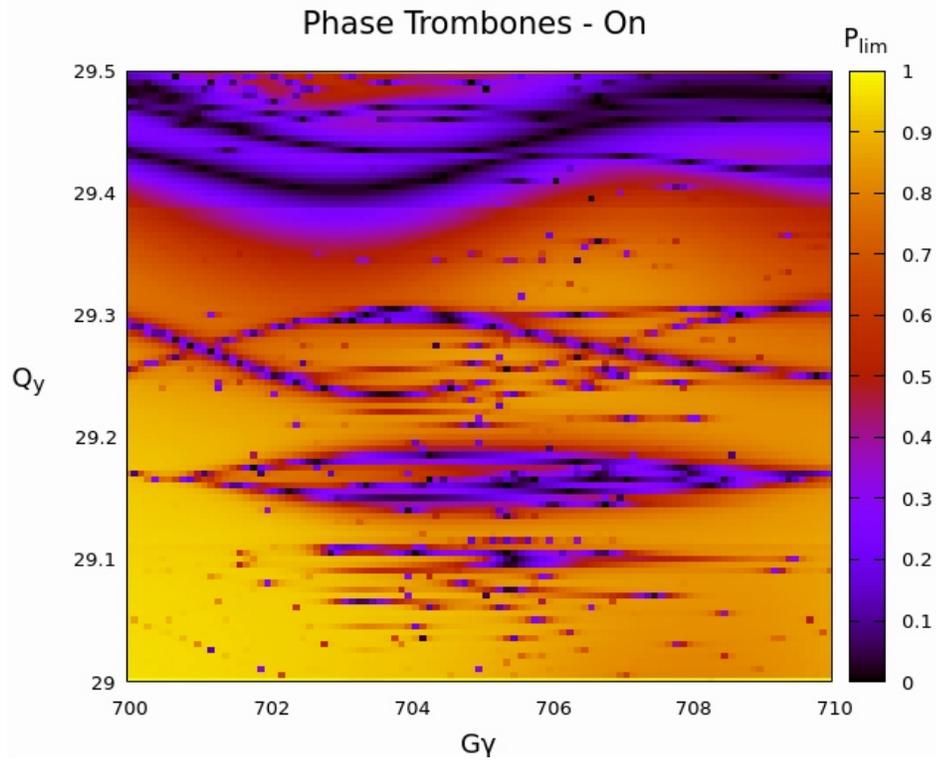


Figure 4.5: Tune scan for the HSR lattice after optimizing the inter-snake betatron phase advances at $G\gamma \approx 704$. The resonance doublets are significantly narrowed compared to the baseline (Fig. 4.4), creating a much larger area of high polarization (bright yellow regions). This demonstrates the effectiveness of phase advance optimization in suppressing ADST spread.

This improved lattice, however, is optimized for a single depolarizing energy region. Since the spin-orbital phasing is energy-dependent, a solution that works at one energy may not work at another. To achieve high polarization throughout the acceleration ramp, the phase advances must be dynamically adjusted. By identifying several key energy points corresponding to the strongest resonances and finding an optimized set of phase advances for each, we create a continuous “phase ramp” through interpolation. Figure 4.6 shows the result

of tracking through the full energy range using such an interpolated ramp. Polarization transmission is successfully maintained with $P_{\text{dyn}} \geq 99\%$, a dramatic improvement over the baseline (Figure 4.3).

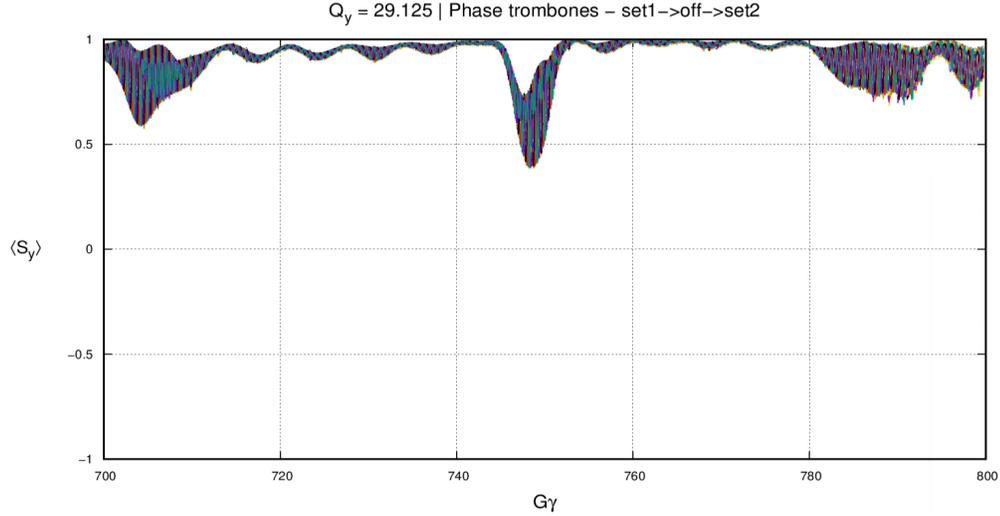


Figure 4.6: Spin tracking with an energy-dependent interpolation of optimized phase advances. Each curve shows the vertical spin component of individual particles. Polarization is successfully preserved through the challenging energy range with $P_{\text{dyn}} > 99\%$, demonstrating that carefully orchestrated phase advance variations can navigate the dense resonance structure.

A more practical solution is to find a single, fixed set of phase advances that provides adequate performance over the entire energy range. Using scalarized multi-objective optimization, simultaneously targeting the minimization of I^\pm at several key energies, we had trouble converging to an adequate solution. Since I^\pm was being used as a proxy for $|\nu_{1,y}|$, we instead leveraged Bmad/PTC to directly minimize $|\nu_{1,y}|$ at multiple energies and found a robust, constant-phase solution at the expense of temporarily shifting the vertical tune away from the desired working point. Fig. 4.7 shows the result: this single setting also preserves polarization to a high degree, with $P_{\text{dyn}} > 96\%$.

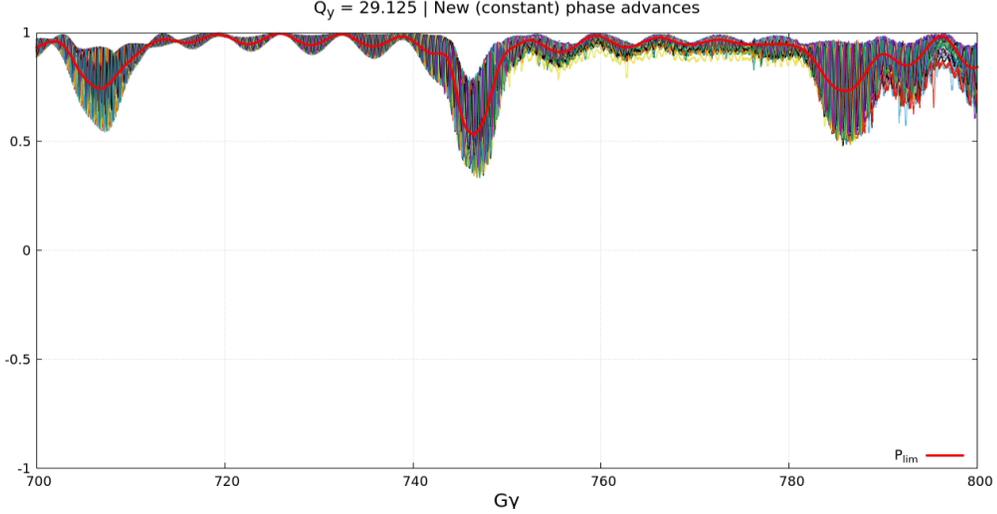


Figure 4.7: Spin tracking using a single, constant set of optimized phase advances, found via multi-objective optimization targeting multiple energies simultaneously. This provides a robust solution without requiring dynamic lattice changes during the ramp, achieving $P_{\text{dyn}} > 96\%$ while being operationally simpler than the dynamic solution.

4.3.4 Practical Implementation and Limitations

The idealized phase adjustments must be realized by adjusting the currents in the actual HSR quadrupoles. This involves a complex lattice matching procedure, as changing quadrupole strengths affects not only the phase advance but also the beta functions and dispersion, which must be matched at centers of the straight sections. Furthermore, adjusting quadrupoles to control Φ_y inevitably causes a correlated change in the horizontal phase advance Φ_x , which is undesirable and could hypothetically excite new coupling resonances. To achieve this control, the phase advance in each arc is made adjustable as a function of an additional current parameter applied to the arc quadrupoles. Each arc is independently matched to its IRs by solving for unique and smooth optics configurations that satisfy $(\beta_x, \beta_y, \alpha_x, \alpha_y, \eta_x, \eta'_x)$ at both IPs. The optics match is performed sequentially for a range of magnet strengths $\pm 10\%$ around their

nominal value, assuming a linear relation between current and magnetic field (unsaturated magnets). Different sets of matching quadrupoles are used for positive and negative changes in arc phase advance, and the resulting matched quadrupole currents are saved and cubically interpolated to provide a piecewise smooth modulation of the phase advances in each arc. After developing a realistic matching procedure using the IR quadrupoles as a function of the arc quadrupole currents, it appeared that we re-optimized the lattice by targeting the minimization of $|\nu_{1,y}|$ using Bmad/PTC. Final tracking simulations with this fully-matched, realistic lattice and a single, constant set of optimized quadrupole settings confirm the viability of the approach, preserving $\sim 97\%$ polarization through the most difficult part of the energy ramp, as shown in Figure 4.8.

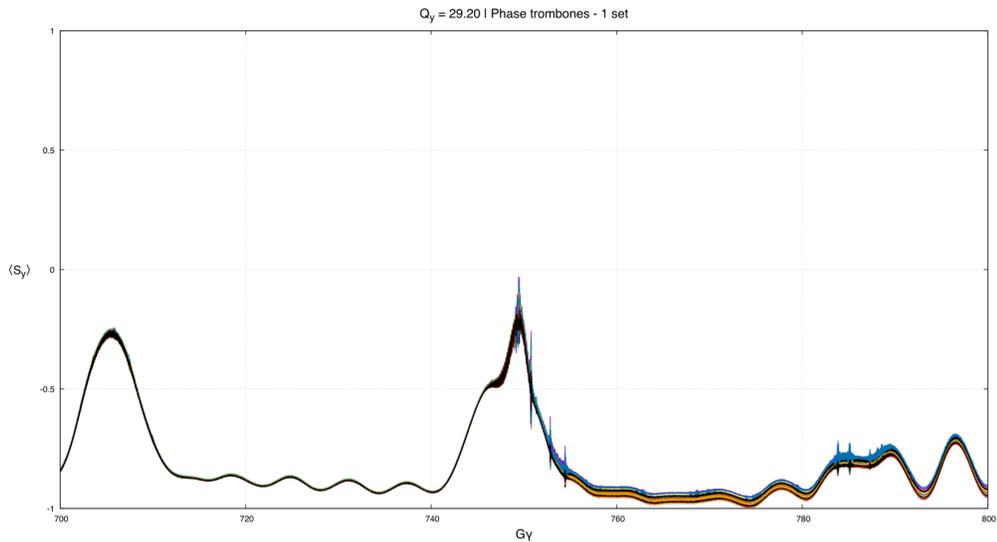


Figure 4.8: Final tracking result with a single, optimized set of quadrupole currents in the fully matched lattice with realistic optics constraints. The solution achieves $P_{\text{dyn}} > 97\%$ while being implementable with actual machine hardware, demonstrating that Method I provides a viable path to high polarization transmission.

Figure 4.9 shows the tune scan after this realistic optimization. While the target vertical resonance doublets are narrowed at the strongly resonant energy,

the tune space shows lower average P_{lim} due to the breaking of approximate lattice symmetry when matching realistic optics. Nonetheless, a safe working point of $Q_y = 29.20$, close to the original working point $Q_y = 29.21$, can be identified with no depolarization across a reasonable energy range.

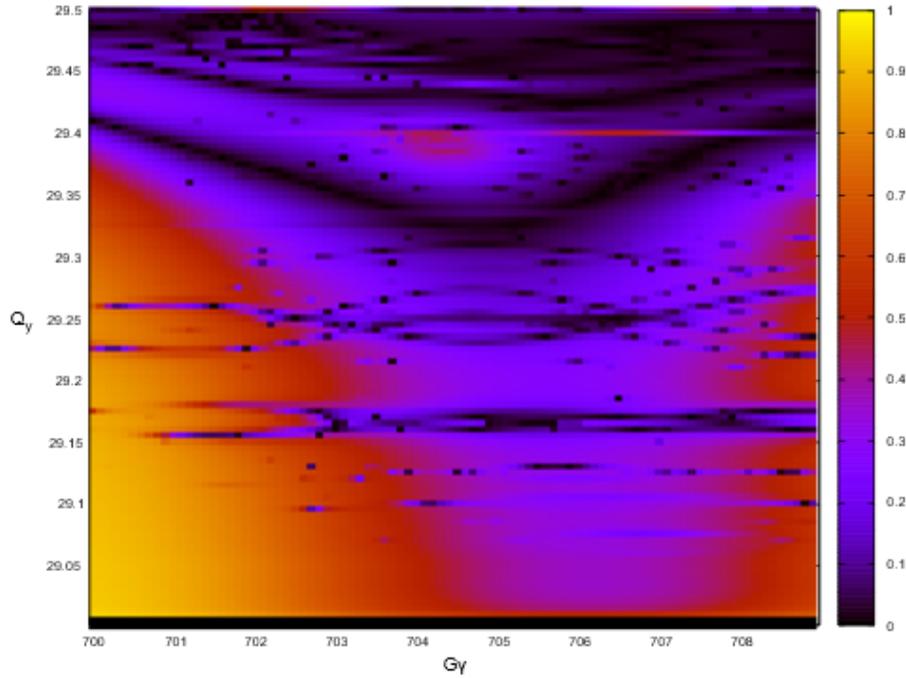


Figure 4.9: Tune scan after optimization using realistic quadrupole matching with actual IR optics constraints. The vertical resonance doublets are narrowed at the target energy, though the tune space is more complex than with idealized phase trombones. A safe working point (marked) can be identified that provides high polarization transmission.

In summary, Method I achieves $> 97\%$ polarization transmission for a 4D distribution of helium-3 through optimized betatron phase advance matching using realistic lattice control.

The primary limitation of the betatron phase optimization method is its nature as a numerical, energy-dependent compensation. It does not eliminate the source of the depolarizing kicks but rather carefully arranges for them to cancel at specific energies. While effective, the solution can be sensitive to lattice er-

rors and, for superconducting machines like the HSR, requires either complex, dynamic control of quadrupole currents (challenging due to quench risks) or a carefully crafted compromise solution that may not be optimal at all energies. This motivates the search for a more powerful solution that suppresses the resonance driving terms at their source through snake axis optimization.

4.4 Method II: Optimization of Snake Rotation Axes

The second, more fundamental strategy for achieving polarization transmission involves manipulating the spin dynamics directly by optimizing the rotation axes of the Siberian snakes [51]. Instead of compensating for depolarizing kicks using the lattice optics, this method aims to suppress the resonance driving terms at their source by imposing a higher degree of symmetry on the spin transport itself.

4.4.1 Exploration of the General Snake Axis Space

For a ring with $2N$ snakes, the condition to fix the closed-orbit spin tune at $\nu_0 = 1/2$ imposes only a single constraint on the $2N$ snake rotation axes $\{\varphi_i\}$. This leaves a vast $2N - 1$ dimensional parameter space of possible configurations. A practical consideration often limits this space: snakes whose axes deviate significantly from the longitudinal direction cause larger orbit excursions into regions of poor field quality, and may even lead to scraping against the vacuum pipe. However, an orbitally-optimal solution is not necessarily optimal for polarization.

We performed a statistical survey over the 5-dimensional snake axis parameter space for the HSR over 49 equidistant beam energies in the $G\gamma \in (700, 800)$ range. We considered the maximum ADST deviation $\max_{G\gamma}(\Delta\nu)$ to quantify the doublet splitting as well as $\min_{G\gamma}(P_{\text{lim}})$ to quantify the residual resonance driving terms. Sampling almost 250,000 unique snake configurations revealed that the nominal $\pm 15^\circ$ scheme is extremely inadequate: it is far from the best-performing configuration. Many other solutions offer significantly improved spin stability, motivating a search for underlying principles that identify these superior configurations.

4.4.2 The Statistical Advantage of Lee-Courant Schemes

A known class of symmetric configurations, which we refer to as Lee-Courant (L-C) schemes, imposes additional constraints on the snake axes [8]. In a ring with $2N$ snakes, N odd, an L-C scheme constrains the spin transport across pairs of snakes to be a pure π rotation. For example, in the HSR with six snakes, this requires the one-turn map sections from snake 1 to 3, from 3 to 5, and from 5 back to 1 to each correspond to a spin phase advance of π . This imposes $N = 3$ independent constraints on the six snake axes, reducing the dimensionality of the solution space from five to three.

Figure 4.10 compares the distribution of $\min_{G\gamma}(P_{\text{lim}})$ from our statistical scan for the general set of solutions versus the subset that satisfies the L-C conditions. The result is unambiguous: the L-C schemes show a strong statistical preference for high polarization. The peak of the L-C distribution is shifted significantly toward higher P_{lim} values compared to the general distribution. This

demonstrates that imposing additional symmetries on the spin transport system systematically improves polarization preservation.

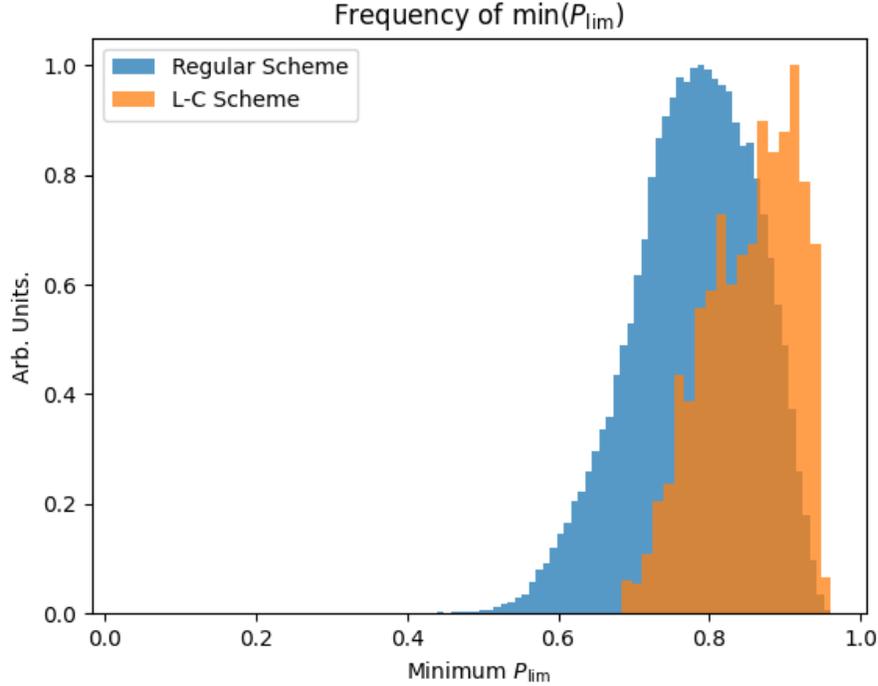


Figure 4.10: Histogram of $\min(P_{\text{lim}})$ over the energy range $G\gamma \in (700, 800)$ for a random sampling of general snake axis configurations (blue) versus the subset satisfying the Lee-Courant conditions (orange). The L-C schemes are statistically far more likely to have higher equilibrium polarization, demonstrating the power of symmetry-based design principles.

4.4.3 The Doubly Lee-Courant (DLC) Scheme

Building on this principle, we introduce and analyze a novel, more restrictive arrangement which we term the “Doubly Lee-Courant” (DLC) scheme [9]. The DLC scheme elevates the concept of local spin cancellation to its logical conclusion by enforcing a π spin phase advance across *every consecutive pair* of snakes, not just across one exclusive set of pairs as in standard L-C schemes.

In this highly symmetric configuration, each snake simultaneously participates in fixing the spin phase advance across two overlapping segments which begin at its upstream *and* end at its downstream neighbor. This imposes the highest possible degree of local symmetry on the spin motion. To our knowledge, this is the first systematic investigation of the DLC class of snake configurations.

For a ring with six snakes, the DLC condition requires that the combined spin transport across any two consecutive snakes and the arc following them is equivalent to a pure vertical rotation of π (or $-\pi$). Using the SU(2) formalism, this stringent condition simplifies to a simple arithmetic relationship between consecutive snake axes:

$$\varphi_{k+1} - \varphi_k = \pm \frac{\pi}{2} \pmod{2\pi}. \quad (4.2)$$

This condition imposes $2N - 1$ constraints on the $2N$ snake axes, leaving only one continuous and one discrete degree of freedom: an overall rotation of the entire snake axis pattern, and an overall reflection of all axes. For the HSR ($2N = 6$), one solution is the alternating $\pm 45^\circ$ snake axes. The general solution is an alternating pair of axes separated by 90° .

Fig. 4.11 illustrates the DLC configuration schematically. Snakes are represented by diamond shapes around the ring. Complete segments begin at dashed lines of one color and end at the second snake of the same color, going clockwise. **The key feature is that each snake simultaneously participates in two overlapping color schemes**, i.e., cancellation in two overlapping segments. This maximal overlap is what distinguishes DLC from standard L-C schemes and provides its superior performance.

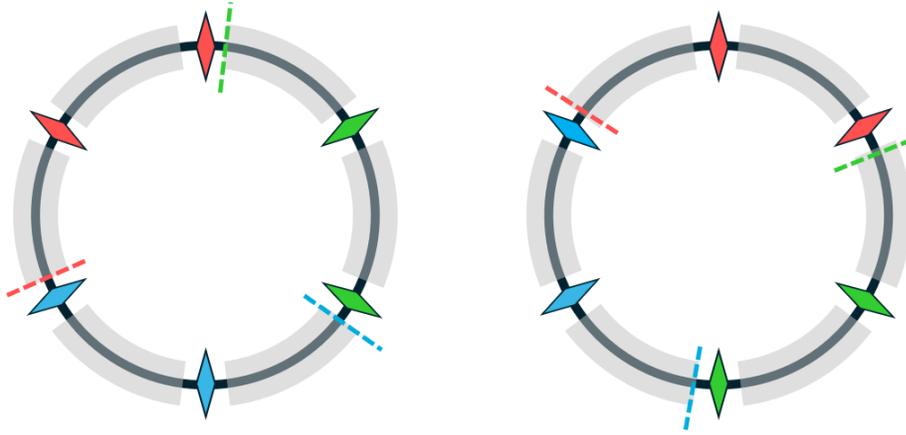


Figure 4.11: **Schematic of the Doubly Lee-Courant (DLC) snake configuration.** Snakes are shown as diamond shapes around the ring. Complete segments (each enforcing a π spin phase advance) are indicated by colored dashed lines, with each segment spanning two consecutive snakes. The defining feature of DLC is that each snake simultaneously participates in two overlapping segments (note how each diamond is part of two different color schemes), providing maximal local spin cancellation.

4.4.4 Performance Results and the DLC Discovery

Comparing the distribution of $\min(P_{\text{lim}})$ for a refined DLC set of snake configurations with the previous general and L-C sets shows that DLC schemes yield a categorically stronger statistical preference for high polarization. Figure 4.12 shows this comparison: the DLC distribution (green) is dramatically shifted toward higher polarization values compared to both general (blue) and L-C (orange) schemes. The median P_{lim} for DLC configurations is near unity, indicating that most DLC solutions provide excellent polarization preservation.

Similarly, evaluating the distribution of maximum ADST spread $\max(\Delta\nu)$ for all three schemes reveals the success hierarchy of General $<$ L-C $<$ DLC in suppressing resonance driving terms. Figure 4.13 shows that DLC schemes consistently achieve the smallest ADST spread, with most configurations having $\max(\Delta\nu) < 0.01$, compared to typical values of 0.02 – 0.04 for L-C schemes and

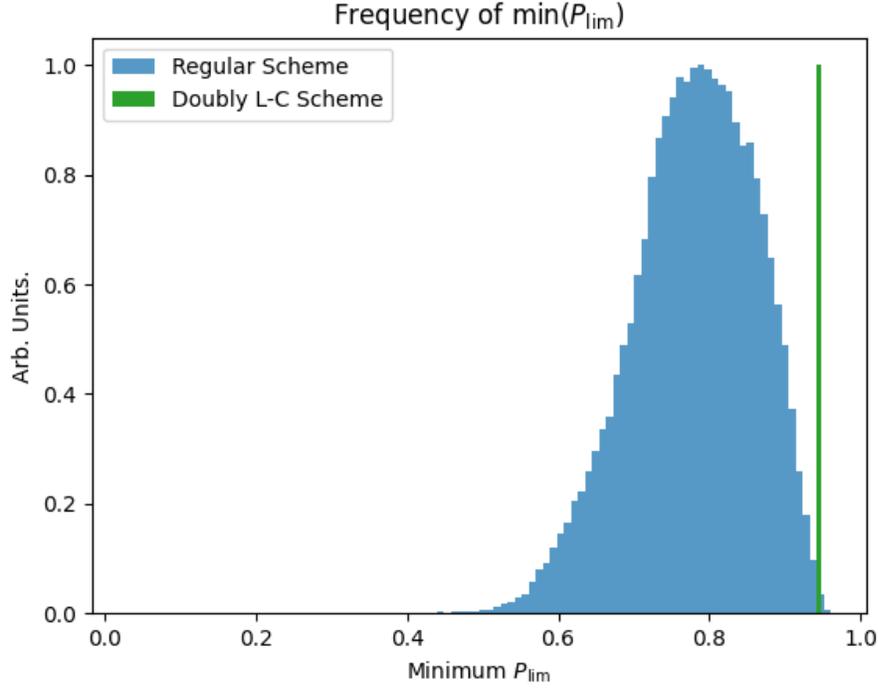


Figure 4.12: Histogram of $\min(P_{\text{lim}})$ for general snake axis configurations (blue), Lee-Courant schemes (orange), and Doubly Lee-Courant schemes (green). The DLC schemes show statistically superior equilibrium polarization, with the vast majority achieving $P_{\text{lim}} > 0.95$. This demonstrates the power of maximal local spin symmetry.

0.04 – 0.08 for general schemes. This factor of 4-8 reduction in ADST spread directly translates to much narrower resonance doublets and hence larger safe operating regions in tune space.

Table 4.2: Snake Axis Schemes for a $2N$ -Snake Ring

Scheme	Constraints	Free Parameters
General	1	$2N - 1$
Lee-Courant (L-C)	N	N
Doubly L-C (DLC)	$2N - 1$	1

The superior performance of the optimal DLC scheme is confirmed with non-perturbative tracking simulations using B_{mad} . Figures 4.14 and 4.15 compares the polarization survival for two different DLC configurations with $\sigma_\delta =$

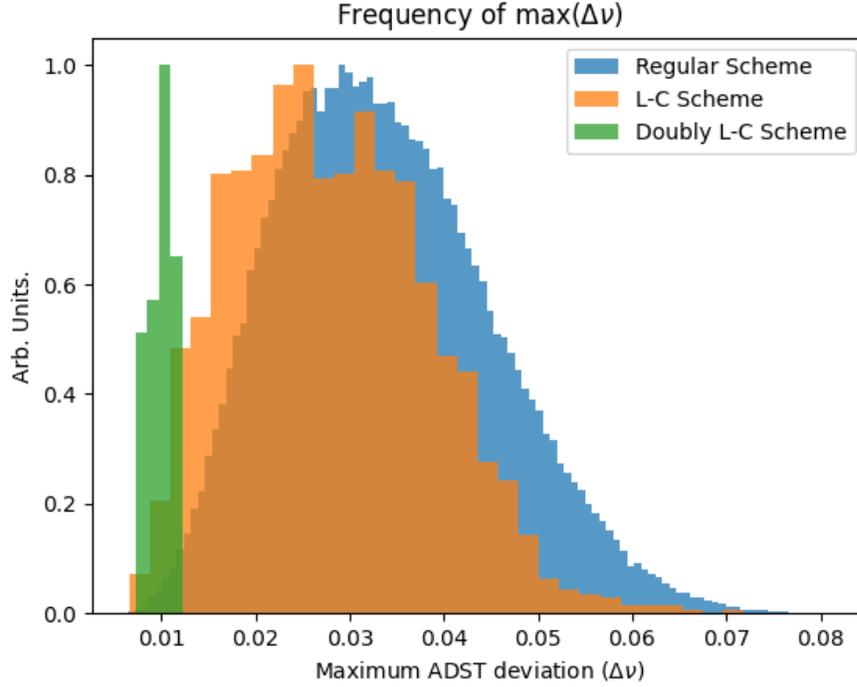


Figure 4.13: Histogram of $\max(\Delta\nu)$ (maximum ADST spread) for general configurations (blue), Lee-Courant schemes (orange), and Doubly Lee-Courant schemes (green). Better performance is indicated by smaller ADST spread. DLC schemes achieve a factor of 4-8 reduction compared to general configurations, demonstrating superior resonance suppression.

0.06% energy spread in the beam: the standard $\pm 45^\circ$ solution and the optimal DLC configuration found through systematic optimization of the single free parameter. The optimal DLC configuration provides exceptionally stable polarization, navigating the entire energy range with minimal loss and dramatically outperforming the baseline configuration shown in Figure 4.3.

The key to this performance is its energy insensitivity. Because the DLC scheme enforces a local cancellation of spin precession that is independent of the arc precession angle ($G\gamma$), it is inherently robust over a wide energy range, unlike other schemes that rely on global, energy-dependent cancellations. As seen in Figure 4.15, the spin vectors of all particles remain tightly grouped through-

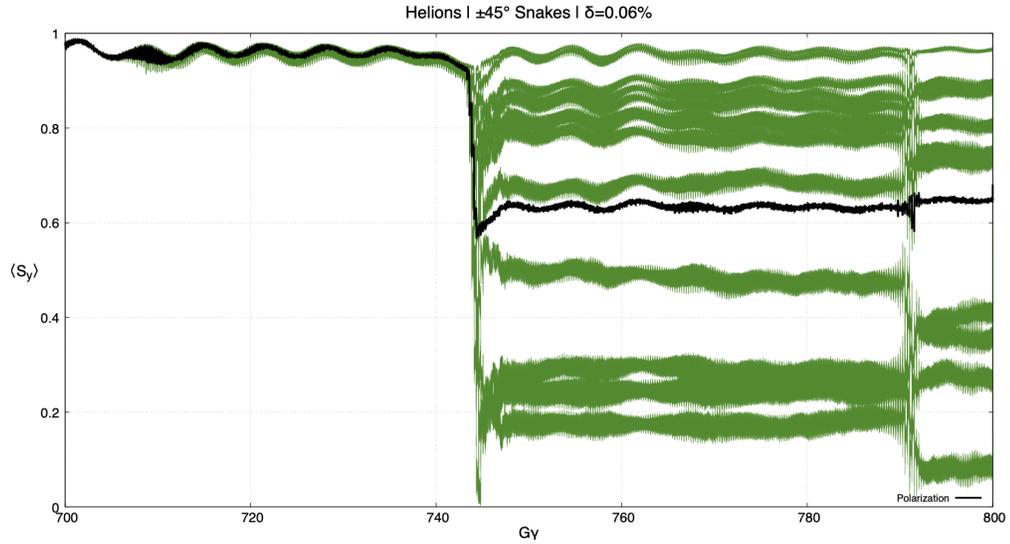


Figure 4.14: Spin tracking result for the standard DLC scheme with $\pm 45^\circ$ snake axes. The configuration shows significantly improved polarization preservation compared to the baseline (Figure 4.3), though some polarization loss still occurs around $G\gamma \approx 745$.

out the ramp, indicating that the DLC symmetry successfully suppresses resonance driving terms at their source.

Taken together, these results establish that the Doubly Lee-Courant scheme achieves near-perfect polarization transmission for helium-3, approaching 100% under ideal conditions and maintaining high transmission even with realistic orbit errors, constituting a qualitative advance in polarization preservation for high- G species. The improvement from the baseline (Fig. 4.3: significant depolarization) to the optimal DLC (Fig. 4.15: near-perfect transmission) demonstrates the transformative power of symmetry-based snake configuration design.

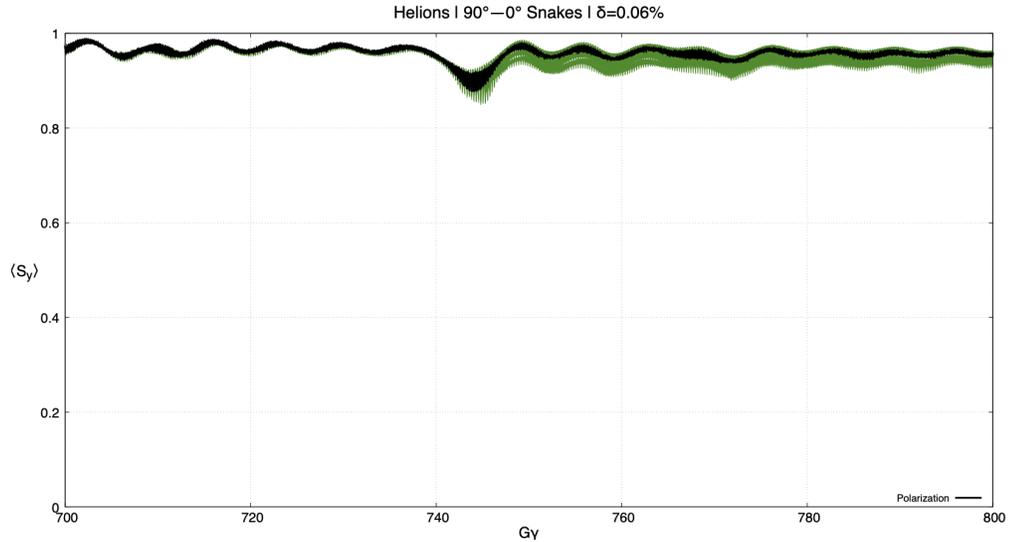


Figure 4.15: **Spin tracking result for the optimal DLC scheme** with maximal P_{dyn} found by optimizing the single free parameter. The configuration provides near-perfect polarization transmission throughout the energy range, with all particle spins remaining tightly aligned. This represents the best achievable performance through snake axis optimization alone.

4.5 Comparative Analysis and Path Forward

A comparative analysis of the two strategies, betatron phase optimization (Method I) and snake axis optimization (Method II), reveals their distinct strengths and underlying philosophies.

Two contrasts frame this comparison. In terms of robustness, Method I is a form of *numerical compensation*: it relies on fine-tuning the orbital phase to cancel depolarizing effects that are still fundamentally present, which makes the solution inherently energy-dependent and potentially sensitive to machine errors that can disrupt the precise phase relationships. Method II, particularly the DLC scheme, is a *symmetry-driven* solution. By modifying the fundamental spin transport to be locally antisymmetric at a granular level between the arcs, it is intrinsically more robust against variations in energy and is predicted to be

less sensitive to lattice imperfections.

From an operational standpoint, practicality also differs. Method I primarily uses existing hardware (quadrupoles). However, it requires precise—and, in superconducting machines, potentially dynamic—control over many magnet circuits during the acceleration ramp, which poses significant operational challenges. Method II offers a “set-and-forget” solution that is robust over the whole energy range. The primary challenge is that ideal DLC configurations often require snake axes far from the longitudinal direction, leading to larger orbit excursions within the snake magnets, which must be accommodated by the beam aperture and may require specialized snake designs.

The two methods are not mutually exclusive but are, in fact, complementary. The most effective strategy for the EIC HSR is a hybrid approach. The DLC scheme should be implemented as a robust, symmetric baseline, which will passively eliminate the vast majority of the resonance strength. Then, betatron phase advance matching can be used as a fine-tuning tool, applying small, static corrections to compensate for any residual effects from lattice imperfections or the symmetry-breaking IRs.

4.5.1 Implications for EIC Commissioning

The results presented here inform the EIC commissioning strategy in several concrete ways. Helium-3 commissioning effectively requires the DLC scheme: the baseline configuration is inadequate at top energy, even when pre-cooling the beam, and implementing the Doubly Lee-Courant configuration offers a decisive improvement, trading increased engineering complexity for dramatically

better physics performance. Diagnostic capabilities remain paramount as well. Both the RHIC experience (4.1) and HSR simulations emphasize that orbit quality dominates polarization preservation, motivating high-resolution BPMs, frequent polarimeters to track polarization evolution during ramps, and sophisticated orbit feedback systems. More broadly, the DLC discovery provides a new toolkit: by embedding local π spin phase advances across every consecutive snake pair, the framework elevates snake placement from empirical optimization to symmetry-based design for future high-energy polarized-beam facilities.

4.6 Discussion: Closed-Orbit and Sextupole Feed-Down

In RHIC, most corrector magnets are themselves superconducting. Nonetheless, they operate at substantially lower currents than main dipoles and quadrupoles and are therefore more tolerant of rapid, cycle-to-cycle adjustments without quench. This makes correctors the preferred actuators for dynamic, energy-dependent corrections, whereas superconducting quadrupoles typically cannot be ramped quickly during an acceleration cycle. In that direction, closed-orbit manipulation may prove more practical for spin corrections.

Closed-orbit bumps affect spin dynamics through two qualitatively distinct mechanisms. Orbit-flattening bumps used to correct closed-orbit errors reduce vertical field integrals seen by the beam and directly mitigate *imperfection* spin resonances (those driven by a non-zero vertical closed orbit). In contrast, deliberately steering through sextupoles generates effective quadrupole fields via sextupole feed-down, modifying the local linear optics; these feed-down bumps can be tuned to cancel *intrinsic* resonances that arise from the coupling of spin

to betatron motion.

Among orbit-flattening strategies, certain families of Fourier harmonics of the reference orbit have an outsized influence on spin dynamics, specifically those coherent with strongly systematic spin resonances. In rings whose arcs are built from repeating FODO cells, coherence between spin phase advances cell-to-cell produces strong systematic resonances with characteristic harmonic content. For the present case with arc periodicity 3 and approximately 27 effective FODO cells per arc, the dominant systematic content occurs near harmonic $h \approx 3 \times 27 = 81$. Orbit-control that suppresses (or avoids exciting) these coherent harmonics yields disproportionate gains in polarization retention.

4.6.1 Sextupole Feed-Down

Sextupole magnets, designed for chromatic correction, generate fields $B_y = k_2(x^2 - y^2)/2$ where $k_2 = \partial^2 B_y / \partial x^2$. When the closed orbit passes off-axis through a sextupole at $x = x_{\text{CO}}$, the Taylor expansion yields

$$B_y(x_{\text{CO}} + \Delta x) \approx k_2 \frac{x_{\text{CO}}^2}{2} + k_2 x_{\text{CO}} \Delta x + \dots \quad (4.3)$$

The second term, proportional to $k_2 x_{\text{CO}}$, acts as an effective quadrupole field, the feed-down. By steering the closed orbit through selected sextupoles, effective quadrupole perturbations can be introduced at chosen locations to destructively interfere with intrinsic resonance driving terms. This suggests that we can leverage orbit control for correcting both imperfection *and* intrinsic spin resonances.

4.6.2 Operational Guidance

In practice, a robust workflow is: (i) flatten the orbit to suppress imperfection resonances, emphasizing suppression of coherent orbit harmonics near $h \approx 81$; (ii) introduce targeted feed-down bumps only as needed to cancel residual intrinsic resonances; and (iii) verify dispersion and emittance at critical apertures using the local relation above.

4.6.3 Chapter Summary

This chapter has presented a comprehensive investigation of spin dynamics in hadron storage rings, from the operational lessons of RHIC to the challenges of the future HSR. The journey from RHIC observations through HSR predictions to the Doubly Lee-Courant innovation illustrates the integration of rigorous physical modeling, innovative optimization strategies, and advanced computational methods. Taken together, the contributions form a coherent progression from diagnosis to design. It first demonstrated that orbit quality, rather than snake fine-tuning, governed RHIC's resilience to a damaged snake. Building on that insight, Method I was developed (betatron phase-advance optimization) and achieved $> 97\%$ helium-3 transmission in realistic lattices. The central advance was the discovery and characterization of the Doubly Lee-Courant (DLC) configuration, which attains near-perfect polarization transmission through maximal local spin symmetry. This, in turn, established a symmetry-based framework for snake placement that supersedes purely numerical searches and culminated in a hybrid operating recommendation: deploy DLC as a robust baseline and apply modest phase-advance adjustments

for fine correction. The Doubly Lee-Courant scheme stands as this chapter's, and this dissertation's, most prominent contribution to polarized beam physics. By enforcing a π spin phase advance across every consecutive pair of snakes, the DLC configuration provides energy-insensitive polarization preservation that is qualitatively superior to all previous approaches.

CHAPTER 5

MACHINE LEARNING FOR INJECTION OPTIMIZATION

The preservation of spin polarization in hadron accelerators extends beyond the challenges of acceleration and storage. It begins with the quality of beam injection. Poor injection efficiency directly impacts polarization through mechanisms ranging from mismatch-induced betatron oscillations to scraping losses that preferentially remove particles with certain spin orientations. For machines like RHIC and the future HSR, the injector chain represents a critical bottleneck where careful optimization can yield substantial improvements in final polarized luminosity.

Yet optimizing injection in existing facilities presents a challenge distinct from those addressed in previous chapters. Unlike the symplectic tracking of Chapter 3 or the depolarization correction strategies of Chapter 4, injection optimization must contend with sparse, noisy, and sometimes unreliable diagnostics. The AGS and its injectors lack the high-quality turn-by-turn BPM measurements and emittance measurement systems available in modern storage rings. Instead, operators must make tuning decisions based on $O(100)$ -turn averaged orbit data, orbit-dependent beam-size data from IPMs, and integrated current transformer readings. This imperfect diagnostic environment makes traditional model-based optimization difficult to validate and motivates the application of machine learning methods that can adaptively learn from limited data.

This chapter explores the intersection of hadron spin physics and modern machine learning, demonstrating how data-driven optimization can improve injection performance even when physics models are incomplete or diagnostic information is limited. The work progresses through three conceptual stages.

The narrative begins by establishing the diagnostic limitations that define the optimization challenge, then proceeds by applying Bayesian optimization to improve transmission efficiency using only scalar intensity measurements. Ultimately, there is a discussion on how to develop reinforcement learning frameworks that incorporate richer state information from multiwire profile monitors to enable more sophisticated, polarization-aware optimization.

The broader picture connects injection quality to the central theme of polarization preservation. Just as snake configurations mitigate resonances during energy ramps and optics corrections suppress spin-orbit coupling, injection optimization ensures that beams enter the acceleration cycle with minimal emittance, optimal trajectories, and maximum transmission. This foundation is essential: no amount of careful acceleration can recover polarization lost to poor injection. The machine learning frameworks developed here provide operational intelligence that complements the physics-based modeling of earlier chapters, enabling practical adaptation to the inevitable imperfections of actual machines.

5.1 Injector Chain Instrumentation

Before exploring machine learning solutions, it is essential to understand the diagnostic environment that motivates their application. The HSR injector chain, comprises the Optically-Pumped Polarized Ion Source (OPPIS), the 200 MeV Linac, the Linac-to-Booster (LtB) transfer line, the 1.5 GeV Booster synchrotron, the Booster-to-AGS (BtA) transfer line, the AGS, and the AGS-to-RHIC (AtR) transfer line. This accelerator chain currently suffers from instrumentation limitations that would be considered unacceptable in a modern facility but are the

reality of machines designed decades ago and incrementally upgraded.

5.1.1 Beam Position Monitors

The most fundamental diagnostic tool in any accelerator is the beam position monitor. In modern rings like RHIC, turn-by-turn BPMs provide position measurements of *select* bunches for each revolution in chunks of 1024 turns, enabling detailed orbit reconstruction, tune measurements via Fourier analysis, and observation of injection transients. The temporal resolution of turn-by-turn BPMs allows operators to distinguish between closed orbit distortions, coherent betatron oscillations, and injection mismatch.

In contrast, with only one or two turn-by-turn BPMs dedicated for destructive fast-kicker-based tune measurement systems, the AGS employs 72 dual-plane, multi-turn-averaged BPMs that integrate signals over hundreds of revolutions before reporting position [40]. While this averaging improves signal-to-noise ratios for low-intensity, stationary beams, it fundamentally obscures fast dynamics. Injection oscillations, which carry critical information about trajectory matching, are averaged away. Coherent modes that might indicate optics errors or space charge effects are invisible. Especially during injection, which is a dynamic process, the closed orbit measurement becomes ambiguous when averaged over changing beam distributions. To alleviate this, the turn-by-turn tune-meter system must be retrofitted and re-calibrated to non-destructively acquire transient signals at injection which is unsustainable in the long term. Transfer lines between machines are particularly under-instrumented, with the only position data coming from 2-3 destructive multiwire screens (harps).

This sparse, temporally-averaged position information makes traditional orbit correction challenging. Furthermore, standard algorithms like MICADO or SVD-based orbit correction are most effective for addressing steady-state orbital distortions rather than transient distortions such as injection mismatch. Measured orbit responses to corrector (steering) magnets become unreliable during the dynamic injection process especially since BPMs average over changing orbits or when spatial gaps require extrapolation.

5.1.2 Emittance and Profile Measurements

Beam emittance, the phase space area occupied by particles, is a fundamental quality metric that directly impacts luminosity and polarization lifetime. The AGS relies primarily on ion profile monitors (IPMs) where the beam ionizes residual gas and the resulting ions are extracted to micro-channel plate detectors. While conceptually straightforward, these devices suffer from intensity-, voltage-, and orbit-dependent systematics. For example, at low intensities ion statistics are poor and profiles are noise dominated, whereas at high intensities, space charge modifies the extraction fields, distorting apparent profiles.

Calibration of IPMs is non-trivial. The relationship between measured the spatial distribution of ion counts and actual beam density depends on extraction gap voltage, gas composition, and beam distance away from the microchannels [56]. Absolute emittance measurements require careful analysis that is rarely performed during routine operations where quick results are needed. Instead, operators use qualitative comparisons, whether a profile is "broader" or

“narrower” than previous measurements, as proxies for emittance quality.

To be more accurate, trustworthy emittance measurements can be made in the LtB and BtA transfer lines via upstream quadrupole scans using multi-wire profiles [57]. The downside to this is that it introduces uncertainty about whether observed emittance growth occurred during injection or acceleration.

5.1.3 Closed-Orbit Distortions

Perhaps most striking from an accelerator physics perspective is the AGS’s tolerance for large closed-orbit distortions. In well-designed modern machines, closed orbits are corrected to sub-millimeter RMS deviations, ensuring that beams experience only the intended magnetic fields and minimizing higher-order multipole effects.

The AGS routinely operates with closed-orbit distortions reaching centimeter scales. At injection energies (typically 1.5 GeV kinetic energy for protons), orbit excursions of 5–10 mm are not uncommon. While these distortions could trigger alarms in facilities like RHIC or the LHC, they have become an operational norm in the AGS, partly because incrementally correcting them yields negative results. This indicates that operators have empirically found a local maximum for performance across the closed-orbit landscape, while it is unlikely but conceivable that this is the global maximum due to the 25A limits of corrector strength. Furthermore, it ought to be stated that some closed-orbit distortions are deliberately set for centering of the orbit through the partial snake magnets which generate a large helical trajectory within them up to 2cm at injection.

This acceptance of large orbits has interesting implications. On one hand, it suggests that the AGS lattice and apertures are robust enough to tolerate such deviations without serious losses. On the other hand, it raises questions about systematic effects: are these orbits truly optimal, or merely "good enough"? Do they generate sextupole feed-down that happens to partially correct other errors, creating an accidental compensation that would be disrupted by naive orbit flattening?

The difficulty in answering these questions definitively reflects the diagnostic limitations. Without detailed turn-by-turn orbit data, comprehensive field quality measurements, and accurate emittance monitors, it is difficult to establish ground truth against which to validate corrections. This is precisely the environment where machine learning approaches offer value: they can optimize measurable objectives (transmission, downstream intensity) without requiring complete understanding of the underlying physics.

5.1.4 Implications for Optimization Strategies

The diagnostic limitations described above fundamentally shape the optimization problem. Traditional accelerator physics approaches assume that one can measure the closed orbit with high precision at many locations, determine emittances and beam distributions accurately, construct and validate response matrices for correctors, and compare measured beam parameters against model predictions to identify discrepancies.

In the AGS injector chain, none of these assumptions hold reliably. Orbits are poorly measured. Emittances are qualitative at best. Response matrices are

nontrivial due to the coupling induced by snakes and the nonzero orbit through sextupolar fringe fields of main magnets. Most models fail to capture the dominating nonlinear effects, misalignments, and systematic errors present in a machine that has evolved over decades of incremental modifications.

This situation motivates a shift from model-based to data-driven optimization. Rather than attempting to build perfect physics models and validate them against imperfect data, machine learning approaches directly optimize measurable outcomes: transmission efficiency, downstream beam intensity, or ultimately polarization transmission. These methods treat the accelerator as a complex function to be optimized, without requiring explicit models of all internal dynamics.

The key insight is that even limited diagnostics provide significant information about performance. A current transformer measuring injected intensity is a noisy scalar, but it unambiguously indicates whether transmission improved or degraded. Ion profile monitors may have systematic errors, but relative changes in profile widths still indicate whether matching improved, so long as the vertical orbit remains fixed. By leveraging such partial information intelligently, machine learning can gradually explore parameter space and identify improved operating points.

Importantly, this data-driven approach complements rather than replaces physics understanding. The choice of which parameters to vary (corrector magnets rather than main dipoles), the ranges over which to vary them (avoiding aperture limits), and the structure of the search space (symmetries, periodicity) all benefit from accelerator physics expertise. Machine learning provides the optimization engine; physics provides the constraints and priors that make the

search tractable.

The following sections describe two progressively more sophisticated applications of machine learning to the AGS injection problem. Section 5.2 presents results from a Bayesian optimization of the BtA transfer line using only downstream intensity as an objective [58]. Section 5.3 discusses the hypothetical application a reinforcement learning frameworks that incorporates multiwire profile information as state variables, enabling more nuanced optimization strategies.

5.2 Bayesian Optimization for Booster-to-AGS Injection

This first application of machine learning to AGS injection focuses on the Booster-to-AGS transfer, where beam extracted from the Booster at 1.5 GeV kinetic energy must be matched into the AGS injection lattice. This process involves approximately 20 magnets, dipoles for trajectory steering, quadrupoles for matching, and correctors for fine-tuning, that must be set to values that maximize transmission while preserving emittance and beam quality.

Traditionally, this tuning is performed by experienced operators who adjust magnets iteratively based on downstream intensity measurements. The process is time-consuming, requires significant operator expertise, and produces results that depend on individual operator strategies. Moreover, the parameter space is high-dimensional (20+ magnets) and strongly coupled (changing one magnet affects the optimal settings of others), making systematic exploration difficult.

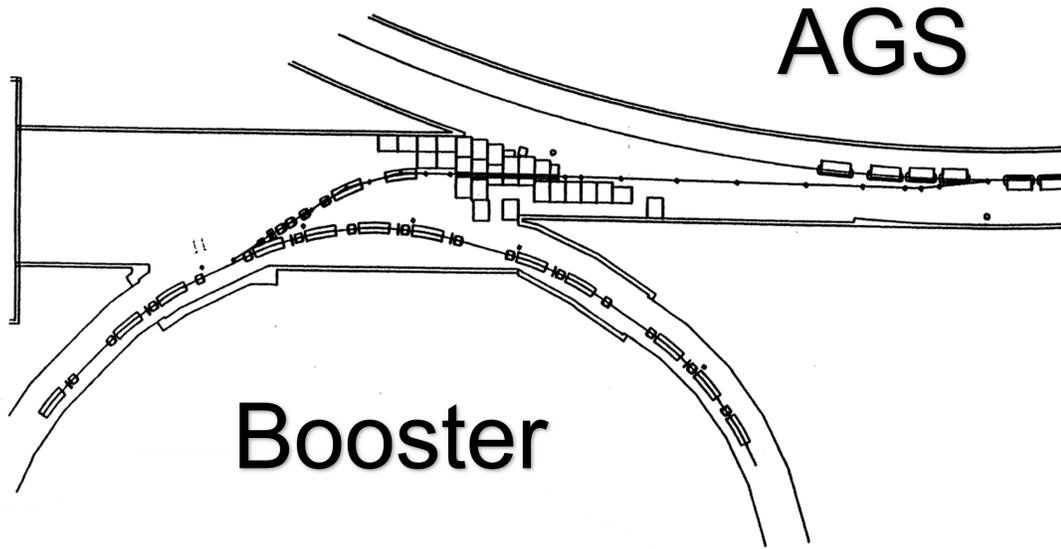


Figure 5.1: Schematic diagram of the BtA line and its connections with the Booster and the AGS.

5.2.1 Bayesian Optimization Framework

Bayesian optimization (BO) provides a principled approach to optimizing expensive black-box functions in high-dimensional spaces. The core idea is to build a probabilistic surrogate model of the objective function that is cheap to evaluate, use this easy-to-optimize surrogate model to decide where to sample next, and update the surrogate model iteratively as new data arrive.

For the injection optimization problem, the objective function maps magnet settings $\vec{x} = (x_1, x_2, \dots, x_N)$ to transmission efficiency $T(\vec{x})$, measured by the ratio of intensity detected in the AGS after injection to the intensity of the Booster at the end of its ramp. Diagnostic measurements arrive each AGS injection cycles, effectively each 5–10 seconds, providing a natural timescale for optimization iterations. The transfer line beam dynamics are dominated by linear

motion, dipole steering and first-order focusing, which simplifies the optimization landscape compared to the nonlinear ring dynamics encountered during acceleration. This function poses several challenges for traditional optimization. Evaluations are expensive: each one requires ramping magnets, injecting beam, and waiting for measurement systems to stabilize, a process taking up to 10 seconds per sample. Measurements are noisy, since shot-to-shot variations in source performance, timing jitter, and diagnostic uncertainty cause repeated measurements at the same \vec{x} to yield different transmission values. Derivatives are unknown, because the relationship between magnet settings and transmission involves complex transient beam dynamics and orbit information, making gradient estimation unreliable. Finally, the dimensionality is high: with over 20 parameters, exhaustive grid searches are infeasible (even a coarse 10-point grid per parameter requires 10^{20} evaluations).

Bayesian optimization addresses these challenges by treating $T(\vec{x})$ as a realization of a Gaussian process (GP), a probability distribution over functions characterized by a mean function $\mu(\vec{x})$ and covariance kernel $k(\vec{x}, \vec{x}')$. The GP prior encodes assumptions about function smoothness and typical lengthscales of variation, which are refined as data accumulate [59].

After observing n samples $\{(\vec{x}_i, T_i)\}_{i=1}^n$, the posterior distribution over $T(\vec{x})$ at any new point \vec{x}_* is Gaussian with mean and variance:

$$\mu_n(\vec{x}_*) = \vec{k}_*^\top (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \vec{T} \quad (5.1)$$

$$\sigma_n^2(\vec{x}_*) = k(\vec{x}_*, \vec{x}_*) - \vec{k}_*^\top (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \vec{k}_* \quad (5.2)$$

where $\vec{T} = (T_1, \dots, T_n)^\top$ is the vector of observed transmissions, \mathbf{K} is the $n \times n$

covariance matrix of sampled points with $K_{ij} = k(\vec{x}_i, \vec{x}_j)$, \vec{k}_* is the covariance vector of the new point with respect to the sampled points with $k_{*,i} = k(\vec{x}_*, \vec{x}_i)$, and σ^2 is the prior-variance of the measurement noise (assumed to be independent of \vec{x} in this case).

The mean $\mu_n(\vec{x}_*)$ represents the model's best estimate of transmission at \vec{x}_* , while the variance $\sigma_n^2(\vec{x}_*)$ quantifies uncertainty. Note that the predictive variance σ_n^2 is independent of the measured outputs \vec{T} , generically reflecting that regions far from observed data have high uncertainty while regions near many observations have lower uncertainty since the covariance kernel comprises radial basis functions (e.g. Gaussians).

5.2.2 Acquisition Functions and Exploration-Exploitation

The key to effective Bayesian optimization is the *acquisition function*, which uses the GP posterior to decide where to sample next. An acquisition function can have the ability to balance *exploration*, by sampling regions of high uncertainty to learn more about the objective, and *exploitation*, by sampling regions with high predicted transmission to find the optimum [60]. The three most common acquisition functions are described in the following section.

We define Z_ξ to be a biased Z-score of a normal distribution centered at $T_{\max} + \xi$, with standard deviation determined by the posterior uncertainty $\sigma_n(\vec{x})$:

$$Z_\xi \equiv \frac{\mu_n(\vec{x}) - T_{\max} - \xi}{\sigma_n(\vec{x})},$$

where $\xi \geq 0$ is a parameter offering an explicit trade-off between exploration (large ξ) and exploitation (small ξ).

Probability of improvement (PI) measures the probability that $T(\vec{x})$ exceeds the current best plus a bias, ξ :

$$\alpha_{\text{PI}}(\vec{x}) = \mathbb{P} [T(\vec{x}) > T_{\text{max}} + \xi] = \Phi(Z_\xi), \quad (5.3)$$

where Φ is the standard normal CDF.

Expected improvement (EI) quantifies the expected amount by which $T(\vec{x})$ exceeds the current best observed value T_{max} plus a bias, ξ :

$$\alpha_{\text{EI}}(\vec{x}) = \mathbb{E} [\max(0, T(\vec{x}) - T_{\text{max}})] \quad (5.4)$$

$$= \sigma_n(\vec{x}) [Z_\xi \Phi(Z) + \phi(Z_\xi)] \quad (5.5)$$

$$= (\mu_n(\vec{x}) - T_{\text{max}} - \xi) \Phi(Z_\xi) + \sigma_n(\vec{x}) \phi(Z_\xi) \quad (5.6)$$

where ϕ is the standard normal PDF. Beyond the ξ bias hyperparameter, the EI acquisition function also implicitly balances exploitation (through $\mu_n - T_{\text{max}} - \xi$) with exploration (through σ_n).

Upper confidence bound (UCB) provides an alternative direct trade-off between mean and uncertainty:

$$\alpha_{\text{UCB}}(\vec{x}) = \mu_n(\vec{x}) + \kappa \sigma_n(\vec{x}) \quad (5.7)$$

where κ is a tunable hyperparameter controlling exploration (large κ) versus exploitation (small κ).

For the AGS injection optimization, UCB was chosen. While EI is often preferred for its robust performance in noisy environments, practical experience has shown that the acquisition function favors dangerous exploration of parameter space even when ξ was tuned to be exploitative (stick to high-performance neighborhoods). UCB, on the other hand, was more responsive to variations of κ in terms of safety, more diligently avoiding regions with complete beam loss and providing faster convergence in the BtA tuning application.

The optimization loop begins by initializing with a small number of random samples (typically 5–10) to establish the GP prior. It then computes the acquisition function $\alpha(\vec{x})$ over the constrained parameter space using the current GP posterior and selects the next sample point \vec{x}_{n+1} by optimized the acquisition function $\alpha_{\text{UCB}}(\vec{x})$, as follows: $\vec{x}_{n+1} = \arg \max_{\vec{x}} \alpha(\vec{x})$. After measuring transmission T_{n+1} at \vec{x}_{n+1} by injecting beam with those magnet settings, the GP posterior is updated with the new observation (\vec{x}_{n+1}, T_{n+1}) . This process repeats until convergence or the evaluation budget is exhausted.

5.2.3 Kernel Selection and Dimensionality Reduction

The choice of covariance kernel $k(\vec{x}, \vec{x}')$ encodes assumptions about the function being optimized. For smooth functions, a Gaussian kernel is common:

$$k_{\text{SE}}(\vec{x}, \vec{x}') = \sigma_f^2 \exp\left(-\frac{1}{2} \sum_{i=1}^N \frac{(x_i - x'_i)^2}{\ell_i^2}\right) \quad (5.8)$$

where σ_f^2 is the signal variance and ℓ_i are length-scale parameters for each dimension. The length-scales determine how rapidly the function varies along

each axis; smaller ℓ_i implies more rapid variation. These parameters can either be fitted to collected data during optimization, or, as is usually done, they are fixed to unity after linearly rescaling the components of \vec{x} .

For injection optimization, the 20+ dimensional parameter space poses strong demands of machine time, requiring thousands of observations over hours of machine time. Moreover, the true effective dimensionality is likely much lower than 20, since many magnets have correlated effects or minimal individual impact on transmission.

To address this, dimensionality reduction was applied by choosing magnets based on their phase advances rather than via principal component analysis. While there are only 4 independent corrector magnets, there are at least 15 independent quadrupole magnets to choose from. Specifically, quadrupoles were selected to be as far from $\pi/2$ phase advance from each other as possible in the nominal BtA line model, ensuring that their individual effects on the beam are approximately orthogonal and can be adjusted independently. While none of the available quadrupoles were ideally placed near $\pi/2$ apart (which would provide maximum orthogonality), this phase-advance-based selection reduced the effective parameter space from 15 quadrupoles to 3–5 key quadrupoles, dramatically reducing the dimensionality while retaining physically meaningful degrees of freedom.

This reduction has both computational and physical benefits. Computationally, optimizing in a low-dimensional space (3–5 parameters) rather than 20+ reduces the number of samples needed to adequately explore the space and speeds up GP computations. Physically, selecting quadrupoles based on phase advance ensures that control parameters affect different components of the 4D

beam distribution, avoiding redundancy and preventing the optimizer from exploring pathological configurations that violate basic optics principles.

5.2.4 Experimental Results

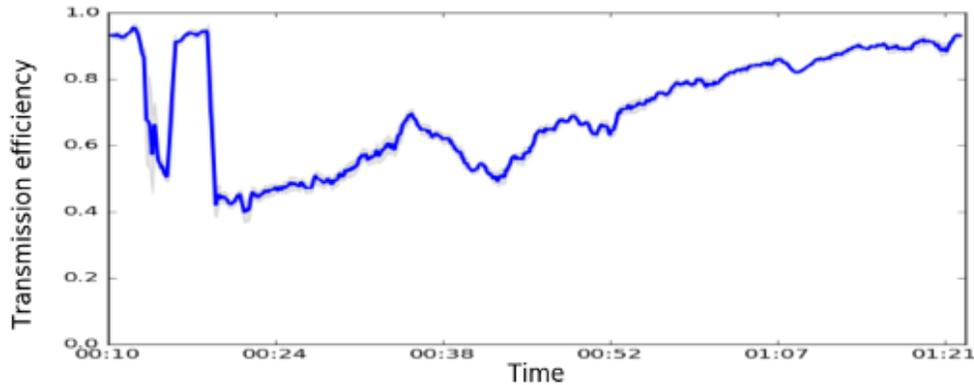


Figure 5.2: BO results of the ratio of AGS injection intensity to Booster extraction intensity after spoiling the initial settings, while varying two horizontal and two vertical steering magnets.

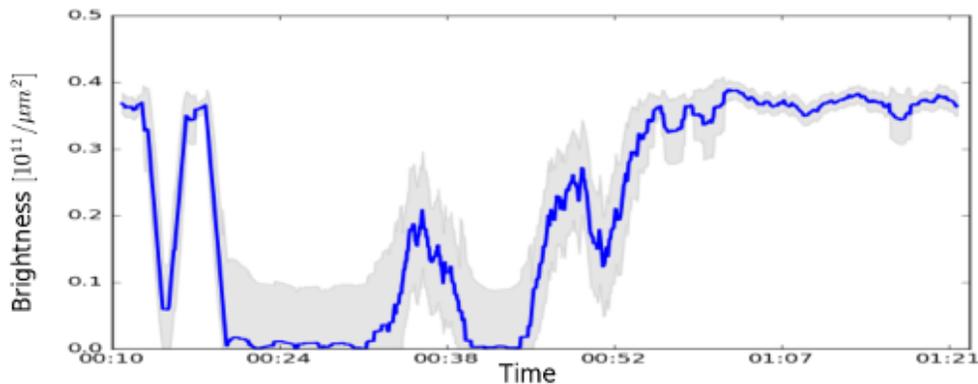


Figure 5.3: Beam brightness is a secondary target not included in optimization, but shows strong recovery following spoilage of the settings.

To test the ability of this approach, we started with a well-tuned beamline that was producing 90%+ proton transmission efficiency and spoiled 4 steering magnets (2 horizontal and 2 vertical). We then launched the Xopt BO toolkit and were capable of recovering the initial transmission efficiency after about 1

hr, as shown above. The reason for the extremely slow convergence in this low-dimensional space is because extra care was taken to obtain clean beam size signals from IPMs, to increase the signal-to-noise ratio through averaging over three cycles, and to accommodate slow-changing magnets between optimization steps.

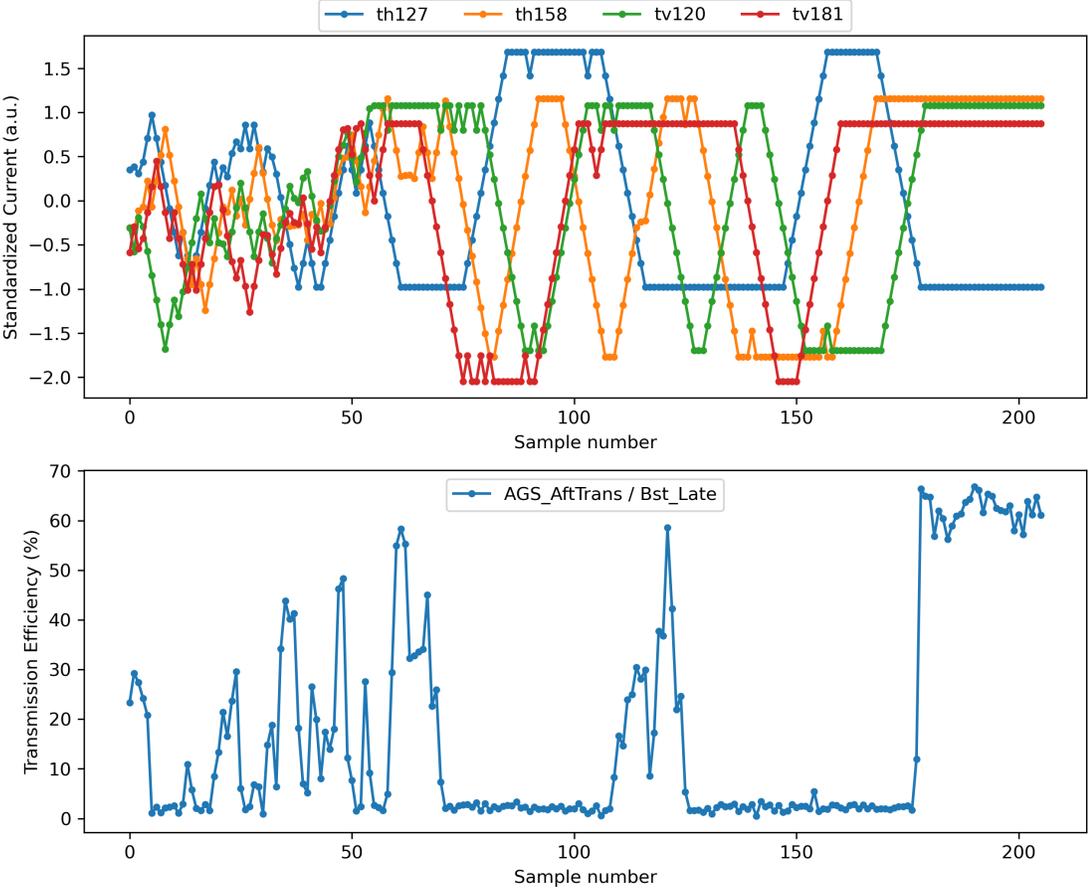


Figure 5.4: BO results of the ratio of AGS extraction intensity to Booster extraction intensity after spoiling the initial settings, while varying two horizontal and two vertical steering magnets.

We perform another test to mimic real-world operations after spoiling the transfer line steering where we did not explicitly average multiple cycles and did not collect IPM data on beam size due to poor SNR. Here we used the EI acquisition function and optimized the ratio of AGS extraction to Booster ex-

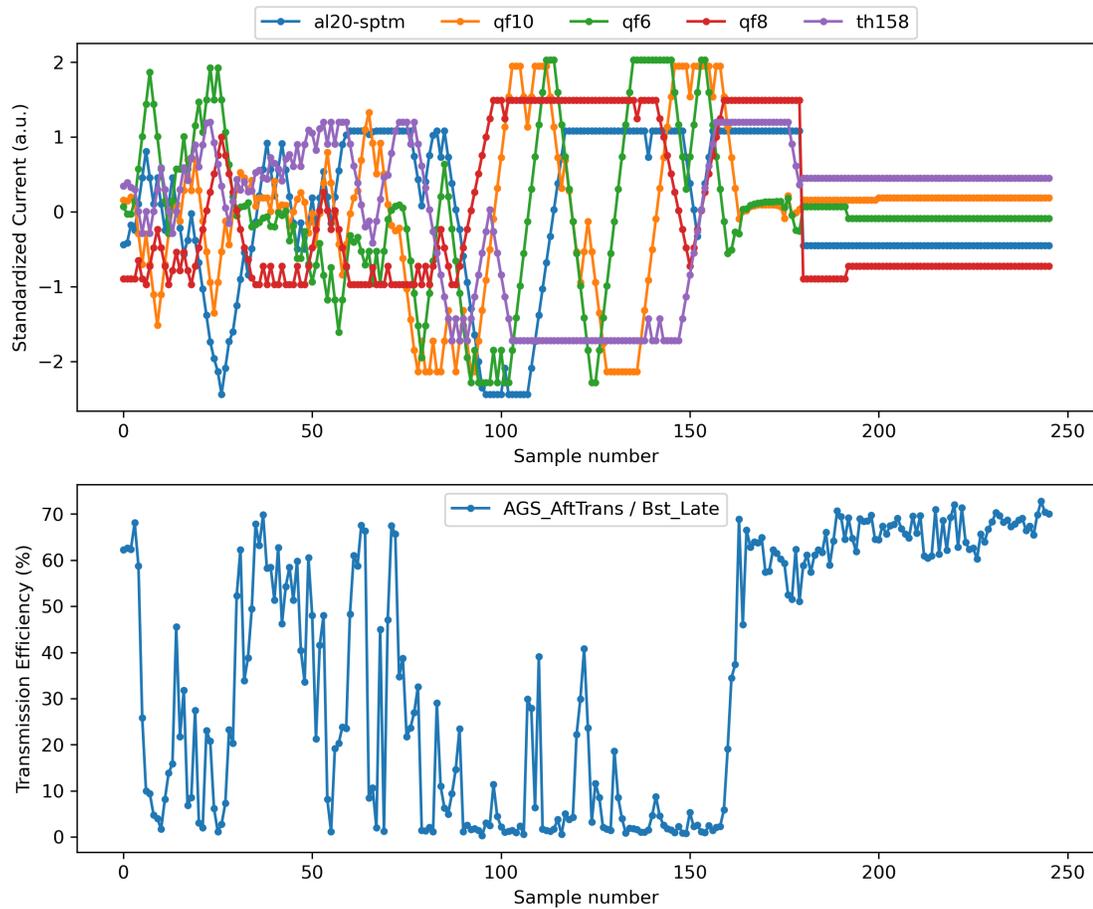


Figure 5.5: BO results while varying three focusing quadrupoles and two horizontal steering magnets to compensate unintentional quadrupole-induced steering (horizontal optics matching).

traction intensity (as opposed to AGS injection to Booster extraction intensity) by varying the 2D steering and achieved much faster convergence results in under 20 mins, which are shown in Fig. 5.4.

After aligning the beam with steering magnets, we focused on horizontal optics matching because of a nonzero horizontal dispersion, using 3 horizontally-focusing quadrupoles as well as 2 horizontal steering magnets and without spoiling any quadrupole settings. While it would be ideal to ignore the steering magnets after the previous alignment, any nonzero orbit through a quadrupole

kicks the beam from feed-down in the field, and this would have to be actively compensated while manipulating quadrupole strengths via two horizontal steering magnets to control position and angle (x, p_x) . In fact, since we do not assume the orbit is zero in either plane through quadrupoles, the vertical feed-down kicks would also have to be compensated, but this effect is much smaller since the horizontally-focusing quadrupoles are placed near the vertical beam-waist: β_y is small, implying a smaller vertical kick from the horizontal quadrupoles. Furthermore, to precisely target horizontal phase space with quadrupoles, we would need 4 linearly independent quadrupoles to constrain the first-order, uncoupled optics. This includes the envelope (β_x, α_x) and the dispersion (η_x, η'_x) . However, relaxing these strong matching requirements slightly allows us to use just 3 quadrupoles.

Again, we focus on the AGS extraction intensity to bypass any consistent and deleterious transient effects. To choose the most suitable quadrupoles, we considered quadrupoles that were as far from $\pi/2$ horizontal phase advance from each other as possible in the ideal BtA line model. None of the quadrupoles were close to $\pi/2$ apart as would be desirable, potentially contributing to a weak result. Fig. 5.5 shows that we were only able to achieve a marginal improvement of about 5% in injection performance compared to the prior result with just steering. Also note that the final converged values of the magnet currents approximately coincide with the initial values before optimization, indicating either that the system is over-constrained and needs more free parameters, or that the optics matching is already strong.

Finally, in a real-world test of the framework, we launched the optimizer to address the former injection efficiency target at a time when the injection was

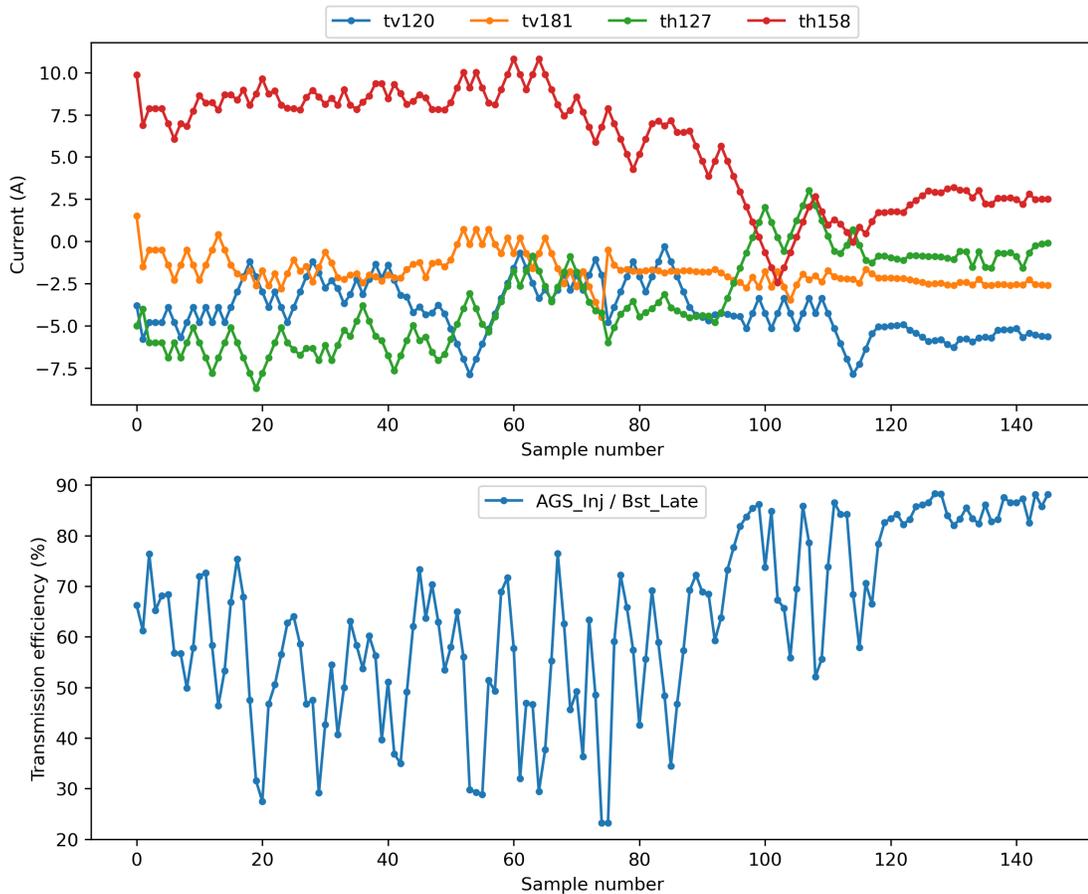


Figure 5.6: Real-world BO results of the ratio of AGS injection intensity to Booster extraction intensity starting from an untuned machine (about 65% transmission), while varying two horizontal and two vertical steering magnets. Convergence in approximately 15 minutes.

poorly tuned, starting at an average 65% transmission efficiency, and performed a pure steering optimization in 4D. In this attempt, we again did not explicitly average multiple cycles and did not collect IPM data on beam size due to a poor signal, achieving impressive and much faster convergence results in approximately 15 minutes as shown above. The final transmission efficiency reached an average 87%, a relative improvement of 34%!

5.2.5 Limitations and Insights

Despite its successes, the Bayesian optimization approach revealed important limitations that motivated subsequent developments. A central shortcoming is the lack of state awareness: because the method optimizes a scalar objective (transmission) without encoding information about *why* losses occur, it cannot distinguish losses due to aperture scraping from those originating in mismatched emittance or timing jitter; any configuration that yields higher transmission is preferred, irrespective of mechanism. Closely related is the absence of trajectory shaping. Intermediate diagnostics, such as multiwire profiles and partial BPM data, are not incorporated, so two settings that achieve identical downstream transmission but follow very different trajectories through sensitive regions are indistinguishable to the optimizer. The framework is also single-objective in practice. It prioritizes transmission without jointly considering correlated quantities such as emittance growth or orbit distortion, which means a configuration that maximizes transmission by barely threading tight apertures may be chosen over a slightly less intense but substantially safer operating point. Finally, the approach does not generalize knowledge across tasks: each optimization is independent, and insights from proton injection do not transfer to gold injection despite shared hardware and similar beam physics, so the algorithm learns specific parameter sets rather than broadly useful operating principles. These limitations point toward more sophisticated frameworks that incorporate partial state information, learn across related tasks, and optimize multi-objective criteria. The next section develops reinforcement learning approaches that begin to address these shortcomings.

Practical insights from operations. In operational tests, the choice of acquisition function materially affected convergence and safety. While EI performed robustly in noisy environments, UCB without a trust region converged faster and better avoided low-performance regions during BtA tuning, which is desirable for operations safety. In addition, explicit multi-shot averaging to denoise IPM signals slowed convergence and did not improve final performance; omitting explicit averaging enabled the 15-minute convergence shown in Fig. 5.6. Furthermore, avoiding averaging allows the GP to learn the true measurement uncertainty as a function of parameter-space: regions where the beam is barely threading a tight aperture will be sensitive to shot-to-shot noise and jitter, and will record a larger measurement uncertainty and be disfavored by the BO process.

Computational constraints of GP surrogates. GP regression scales cubically with the number of observations ($O(n^3)$). For online optimization where each evaluation costs several seconds, sample counts remain modest, but repeated campaigns can accumulate data. Practical mitigations include dimensionality reduction (PCA over magnets), sparse/approximate GP variants, and periodic data pruning to bound n . On the other hand, if the parameter space is large enough and upstream conditions are held approximately constant, then training a neural network-based predictive mean function would be an effective fast prior for any freshly deployed BO process, avoiding retraining from scratch while simultaneously providing more weight to the real-time measurements.

5.3 Reinforcement Learning with Multiwire Profiles

While Bayesian optimization successfully improved transmission efficiency, its limitation to scalar objectives motivated exploration of reinforcement learning (RL) frameworks capable of incorporating richer diagnostic information. Specifically, multiwire profile monitors in the injection lines provide 2D beam profile distributions at several locations, data that contains information about matching quality, trajectory angles, and halo development that is invisible to downstream intensity measurements alone.

Reinforcement learning offers a natural framework for this problem [61]. Unlike Bayesian optimization, which seeks the parameter configuration maximizing a fixed objective, RL optimizes a *policy*, a mapping from observable states to corrective actions. This policy can incorporate partial measurements from multiwire profiles, adapt to different beam conditions, and balance short-term actions such as correcting intermediate mismatches against long-term goals (maximizing downstream transmission).

5.3.1 Markov Decision Process Formulation

The injection optimization problem can be formulated as a Markov Decision Process (MDP), a mathematical framework for sequential decision-making under uncertainty [62]. An MDP is defined by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$ comprising the state space \mathcal{S} representing system observables, the action space \mathcal{A} representing system controls, the transition dynamics $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ representing the probability of likely next states given a state-action pair, the reward function

$\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, and the discount factor $\gamma \in [0, 1]$ that determines a trade-off between short- and long-term rewards.

For AGS injection, these components are instantiated as follows:

State space \mathcal{S} : A state $s \in \mathcal{S}$ comprises summary features from 3–4 multi-wire monitors—transverse centroids and standard deviations $(\mu_x, \mu_y, \sigma_x, \sigma_y)$ per monitor—together with current magnet settings (or, more practically, deviations from a reference configuration), measured transmission efficiency at available current transformers, and possibly historical information such as previous profiles and recent trends over beam pulses.

Importantly, the state space is not high-dimensional: by reducing beam profiles to centroids and standard deviations rather than using full pixel arrays, the state vector is kept to tens of scalar values rather than hundreds or thousands. This compact representation captures the essential trajectory and matching information needed without requiring extremely large datasets for training.

Action space \mathcal{A} : The action space consists of proposed changes to magnet settings:

$$a = (\Delta x_1, \Delta x_2, \dots, \Delta x_N) \quad (5.9)$$

Actions are continuous but bounded to respect magnet operational limits. The action space might be further restricted to the subspace spanned by dominant PCA modes identified in the Bayesian optimization work, reducing dimensionality and incorporating prior knowledge.

Transition dynamics \mathcal{P} : The transition from state s to next state s' after applying action a is determined by the accelerator’s response, beam dynamics,

magnet transfer functions, and diagnostic measurements. Crucially, these dynamics $\mathcal{P}(s'|s, a)$ are *unknown* and must be learned from data. The RL agent does not require an explicit model of \mathcal{P} ; it learns optimal behavior directly from experience.

Reward function \mathcal{R} : The reward $r(s, a)$ quantifies the immediate benefit of taking action a in state s . For injection optimization, the reward combines multiple objectives:

$$r(s, a) = w_1 \cdot T + w_2 \cdot Q_{\text{profile}} - w_3 \cdot \|\Delta\vec{x}\|^2 - w_4 \cdot \text{penalty}(s, a) \quad (5.10)$$

where T is transmission efficiency (if measurable at this state, otherwise estimated), Q_{profile} is a profile quality metric (narrowness, Gaussianity, absence of halo), $\|\Delta\vec{x}\|^2$ penalizes large magnet changes to encourage smooth adjustments, and $\text{penalty}(s, a)$ enforces hard constraints (beam loss, aperture violations, and magnet limit violations).

The weights w_i balance these objectives and are tuned based on operational priorities.

Discount factor γ : Typical choices are $\gamma \in [0.95, 0.99]$, favoring long-term cumulative reward (downstream transmission) while still responding to immediate feedback (early-stage profile quality).

5.3.2 Policy Representation and Deep Reinforcement Learning

The goal of reinforcement learning is to find a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that maximizes expected cumulative discounted reward,

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right], \quad (5.11)$$

and, with continuous action variables such as injection optimization, this policy is naturally parameterized by a neural network $\pi_{\theta}(a|s)$ whose weights θ are trained from data. The network ingests compact state summaries and outputs either distributions over actions or direct action proposals in a deterministic variant.

Two families of algorithms are commonly employed. Value-based methods learn an action–value function $Q^{\pi}(s, a)$ that estimates the return from taking action a in state s and then following π , with the implied greedy policy $\pi(s) = \arg \max_a Q(s, a)$; however, direct maximization over continuous actions is nontrivial, which limits practicality unless the action space is coarsely discretized. Policy-gradient methods instead optimize the parameters of π_{θ} directly via gradient ascent on expected return. Representative algorithms include Proximal Policy Optimization (PPO), Trust Region Policy Optimization (TRPO), and Soft Actor–Critic (SAC), each designed for stability and sample efficiency in continuous control.

For the AGS injection problem, Soft Actor–Critic (SAC) is particularly well matched because it combines an off-policy critic $Q_{\phi}(s, a)$ with a stochastic actor $\pi_{\theta}(a|s)$, enabling efficient reuse of data and a principled exploration mechanism [63]. The critic supplies low-variance targets that guide policy improvement, while the entropy-regularized objective maintains exploration and ro-

bustness, properties that are desirable when optimizing under noisy diagnostics and hardware constraints.

5.3.3 Sim-to-Real Transfer Challenges

The reinforcement learning framework described in this section and the following subsections has not yet been implemented on the real machine. All results, methods, and discussions presented here are hypothetical and based on simulation studies only. Unlike the Bayesian optimization work described in Section 5.2, which was successfully deployed and tested operationally at BNL, the RL approach remains a conceptual prototype. Future work would be required to validate these ideas experimentally.

Simulating injection and transmission efficiency in a predictive way poses severe physical and diagnostic challenges. Accurate modeling of beam loss requires knowledge of the *absolute injected beam position* and angle relative to the true vacuum-pipe apertures throughout the entire transfer line and in the first several hundred turns of the AGS. While BPMs do measure the absolute orbit, their mechanical and electrical centers introduce millimeter-scale uncertainties, and their electronics can only currently process signals representing the *average* beam position over hundreds of turns. As a result, they obscure the most critical beam-loss time scales, the first few turns after injection, where the majority of transmission loss occurs. Furthermore, the AGS is equipped with only a single turn-by-turn BPM, leaving the early-turn dynamics largely unobserved. Under such diagnostic limitations, even small unmodeled offsets can drastically alter predicted losses: a centimeter-scale rms beam must clear an elliptical aperture

of roughly 17 cm width and 8 cm height, so millimeter-scale trajectory errors can determine whether the beam transmits cleanly or scrapes.

These limitations imply that a policy pre-trained purely in simulation is likely to internalize incorrect causal relationships between control actions and transmission outcomes. Rather than learning physically meaningful strategies, it would effectively overfit to a simulator whose absolute orbit reference is uncertain. In practice, improving predictive fidelity quickly devolves into a model-fitting exercise: determining the absolute closed orbit and its offsets throughout the ring to millimeter accuracy. While this is an ambitious and worthwhile sub-goal in its own right, as it would enable more accurate modeling of orbit-dependent spin resonance compensation and aperture effects, it is distinct from the higher-level control objective. Without a verified absolute orbit reference, simulation pre-training risks stunting performance, and transfer to the real machine would require extensive online fine-tuning to recover meaningful behavior.

Consequently, the proper first step in any future effort will be to *validate our knowledge of the absolute beam position through the BtA transfer line*. This can be achieved by experimentally benchmarking the predicted multiwire profiles and beam-loss patterns against measurements, thereby patching the BtA and AGS reference systems into a consistent global orbit model. Only after establishing this alignment can simulation-based reinforcement learning or other model-dependent approaches be meaningfully calibrated for predictive control.

5.4 Summary

This chapter demonstrated operational ML success for injection optimization. In the BtA line, Bayesian optimization recovered transmission from about 65% to 90%+ within roughly 15 minutes using 2D steering with 4 parameters. Reinforcement learning concepts were outlined as future, proposed extensions.

These results support a path toward autonomous operations: Bayesian optimization provides fast, safe, and reproducible tuning under diagnostic limitations. The next chapter turns to slow extraction to NSRL, where similar optimization and modeling principles apply.

CHAPTER 6

SLOW EXTRACTION TO THE NSRL

The challenges of preserving beam quality in hadron accelerators extend beyond injection and acceleration to the final stage: extraction and delivery to experimental areas. This chapter examines slow extraction from the Booster synchrotron to the NASA Space Radiation Laboratory (NSRL), demonstrating how resonance analysis, symplectic tracking, and data-driven control methods developed for spin dynamics apply equally to controlled beam extraction.

The central result of this chapter is experimental validation of a BNL-original extraction mechanism [64]: the Booster's constant-optics slow extraction scheme achieves one to two orders of magnitude suppression (factor of 20–100) of dispersion compared to the periodic dispersion function, enabling the uniform dose delivery required for NSRL's radiobiology and materials science program. This suppression arises from maintaining fixed Twiss parameters during extraction while ramping the reference momentum, causing particles of all momenta to converge onto common unstable trajectories regardless of their dispersive closed orbits.

Validation proceeds through three stages that connect analysis, computation, and measurement. First-order normal form analysis around the dispersive elliptical fixed point, using Kobayashi's Hamiltonian formalism [65], predicts the separatrix geometry and explains why extracted particles' transverse positions are nearly independent of momentum offset. Symplectic particle tracking through the full s -dependent Booster lattice reproduces the predicted dispersion suppression, with no detectable residual dispersion within statistical error. Finally, D6 flag measurements with two different beam momenta ($\Delta\delta = 0.93\%$)

yield a transverse displacement $\Delta x \approx 0.7 \pm 0.5$ mm, confirming $D_{\text{eff}} = 8 \pm 5$ cm—one to two orders of magnitude reduction from the periodic value of $\eta_x = 2.93$ m.

This progression from theory to simulation to experimental confirmation exemplifies the dissertation’s methodology: rigorous analytical frameworks validated through high-fidelity computation and verified through machine measurements. The resonant normal form approach parallels Chapter 4’s spin resonance analysis, the symplectic tracking extends Chapter 3’s methods, and future control strategies connect to Chapter 5’s data-driven optimization.

NSRL serves a critical mission: simulating space radiation environments for biological samples, electronics, and materials. Unlike instantaneous high-flux irradiation, radiobiology studies require prolonged exposure at moderate dose rates mimicking months of cosmic ray exposure compressed into hours. This demands extraction with stable, uniform spill characteristics, enabled by the dispersion suppression mechanism validated in this chapter.

The chapter proceeds in three sections. Section 6.1 develops the normal form theory of constant-optics extraction, deriving the Kobayashi Hamiltonian and separatrix structure. Section 6.2 presents simulation results and experimental measurements demonstrating one to two orders of magnitude dispersion suppression. Section 6.3 discusses future directions for spill quality control using RF channeling and feed-forward compensation, techniques paralleling the injection optimization of earlier chapters.

6.1 Extraction Mechanics

Slow extraction differs fundamentally from fast extraction schemes used in many fixed-target experiments. Fast extraction, typically accomplished using fast-rising kicker magnets that deflect the entire beam into an extraction channel within a single turn or less, delivers beam in short pulses on microsecond to millisecond timescales. While conceptually simple and well-suited for particle physics experiments requiring coincidence measurements, fast extraction cannot achieve the prolonged, low-intensity irradiation required for radiobiology studies where samples must accumulate dose over timescales comparable to cellular response times.

Slow extraction instead relies on gradually driving particles across a resonance separatrix over thousands to millions of turns, enabling extraction periods ranging from milliseconds to many seconds depending on the application. The technique most commonly employed is third-integer resonance extraction, where the horizontal betatron tune is brought near $Q_x = m/3$ for some integer $3 \nmid m$. At this tune, particles executing betatron oscillations complete exactly three oscillations per turn, creating a resonance condition that fundamentally modifies the phase space topology.

6.1.1 First-Order Resonant Normal Form à la Kobayashi

To understand the dispersion suppression mechanism in the Booster's constant optics extraction, we employ first-order normal form analysis around the dispersive, stable fixed point of the third-integer resonance. This rigorous analyt-

ical framework, developed by Kobayashi [65, 66], provides both the separatrix geometry and the extracted dynamics.

Elliptical Fixed Point and Separatrix Structure

The power of normal form transformations of truncated power series is enhanced when the map is close to the identity. In the vicinity of $Q_x \approx m/3$, typically $m = 13$ in the Booster giving $Q_x \approx 4.333$, we transform to the co-moving reference frame that is the composition of 3 one-turn maps. The transformed tune becomes $\omega \equiv 6\pi\Delta Q_x$, where $3\Delta Q_x = 3Q_x - m$.

After normalizing phase space coordinates into (\tilde{x}, \tilde{p}_x) via a first-order normal form transformation (using Twiss parameters) at the extraction point, we write the phasor basis coordinates:

$$z = \sqrt{2J}e^{i\phi_x}, \quad 2J = z\bar{z} = \tilde{x}^2 + \tilde{p}_x^2, \quad \phi_x = \text{phase advance} \quad (6.1)$$

While the RF cavities are not entirely off during extraction, the debunched beam coasts with approximately constant energy. Thus, individual particle tunes Q_x are parametrized through chromaticity by the fractional momentum deviation $\delta = \frac{p-p_0}{p_0}$ between particle momentum p and the reference momentum $p_0 = eB\rho$ of the dipole fields. The amplitude-dependent tune shift is negligible in the Booster since the highest multipoles are sextupoles whose effective tune shift has only nonlinear dependence on amplitude. Then to first order, $\mathbf{Q}(\delta) = \mathbf{Q}_0 + \delta\boldsymbol{\xi} + \mathcal{O}(\delta^2)$ and $\omega = 2\pi(3Q_{0,x} - m) + 6\pi\delta\xi_x$.

Following Kobayashi's formalism, we write the transformed Hamiltonian to

first order in sextupole strength:

$$H = \omega J - \frac{1}{8} (S_v z^3 + \bar{S}_v \bar{z}^3) \quad (6.2)$$

where S_v is the virtual sextupole (resonance strength/driving term):

$$S_v = \frac{1}{2} \sum_j [k_2 l \beta_j^3 e^{3i\phi_x}]_j \quad (6.3)$$

Here, k_{2j} , l_j , β_j , $\phi_{x,j}$ are the normalized strength, length, beta function, and phase advance at each sextupole location. This sum defines an effective coherent strength that governs the shape and orientation of the separatrix, given by the effective phase $3\psi \equiv \arg(S_v)$.

When $H = \frac{8\omega^3}{27|S_v|^2}$, the Hamiltonian factors into a product of three lines forming an equilateral triangular separatrix in the normalized (\tilde{x}, \tilde{p}_x) plane. For $S_v \in \mathbb{R}^+$, the lines defining the separatrix are:

$$\tilde{x} = -\frac{2\omega}{3|S_v|} \quad (6.4)$$

$$\tilde{x} \pm \sqrt{3}\tilde{p}_x = \frac{4\omega}{3|S_v|} \quad (6.5)$$

For general $S_v \in \mathbb{C}$, the separatrix is obtained by rotating these equations by ψ .

Transforming back to physical coordinates (x, p_x) via the Twiss transformation:

$$\begin{pmatrix} x \\ p_x \end{pmatrix} = \begin{pmatrix} \sqrt{\beta_x} & 0 \\ -\frac{\alpha_x}{\sqrt{\beta_x}} & \frac{1}{\sqrt{\beta_x}} \end{pmatrix} \begin{pmatrix} \tilde{x} \\ \tilde{p}_x \end{pmatrix} \quad (6.6)$$

maps the equilateral triangle in normalized phase space to a distorted triangle in real phase space. Critically, because β_x and α_x are constant during the constant-optics spill, this mapping does not rotate the separatrix as the tune changes, keeping its unstable arms fixed in physical space.

For a fixed S_v , the area of stable beam motion within the separatrix equals $2\pi J$ for a particle on the separatrix (by Liouville's theorem). Thus the border of stability is determined by:

$$2\pi J = \frac{8\omega^2}{3\sqrt{3}|S_v|^2} \quad (6.7)$$

Hyperbolic Fixed Point and Extraction Dynamics

As the separatrix is approached and crossed with decreasing ω (driven by the momentum ramp through chromaticity), particles arrive on unstable trajectories approaching the *hyperbolic fixed points*. Such motion is characterized by growth and decay rates rather than tunes, but still approximately respects hyperbolic invariants [67].

Translating the Hamiltonian to the rightmost vertex $x \rightarrow x + \frac{4\omega}{3|S_v|}$ and transforming the intersecting separatrix legs to become perpendicular $\tilde{x} \pm \sqrt{3}\tilde{p} \rightarrow \hat{x}, \hat{p}$, we find (when $S_v \in \mathbb{R}^+$):

$$\tilde{H} \approx \sqrt{3}\hat{x}\hat{p} \left(\omega + S_v(\hat{x} - \hat{p}) \sqrt[4]{\frac{3}{64}} \right) + \frac{8\omega^3}{27|S_v|^2} \quad (6.8)$$

When the beam's phase-space trajectory is carried across one arm of the triangular separatrix, it is funneled into the neighborhood of the hyperbolic fixed point on that arm. The local (linearized) flow is generated by:

$$\tilde{H} \approx \sqrt{3}\omega\hat{x}\hat{p} \quad \implies \quad \hat{x}(t), \hat{p}(t) \approx e^{\pm\sqrt{3}\omega t} \quad (6.9)$$

so that an initial perturbation grows/decays at the rate $\sqrt{3}\omega$ to leading order in S_v . In a perfectly constant- ω lattice, the turn-count $T = 3t$ to reach the thin magnetic septum at \hat{x}_f is:

$$T_f \sim \frac{\sqrt{3}}{\omega} \log(\hat{x}_f/\hat{x}_0) \quad (6.10)$$

When ω itself is ramped, so that the separatrix edge sweeps inward by $\Delta\omega$ each turn, a larger $\Delta\omega$ per turn drives earlier crossing and larger exit amplitude. However, as long as the adiabaticity condition holds:

$$|\dot{\omega}| \ll \sqrt{3} \omega^2 \quad (6.11)$$

the hyperbolic escape remains adiabatic, and the instantaneous growth factor $e^{\sqrt{3}\omega t}$ is essentially the same as in the constant- ω case.

Extracted Beam Distribution and Dispersion Suppression

The dominant contribution to the extracted beam width comes from the distance between the closed-orbit and the septum position, determining the step sizes before extraction, and secondarily by the range of amplitudes at which particles cross their separatrices, effectively reflecting the initial beam amplitude distribution, since instability occurs when the separatrix narrows below a certain fraction of a particle's betatron ellipse.

Critically, the separatrix position at crossing is largely momentum-independent. While dispersion $\eta(s)$ exists at the septum, the final escape position does not follow the usual $\Delta x = \eta\delta$ relationship. Instead, particles exit along a common unstable manifold defined only by J and δ , their betatron amplitude and their momentum offset relative to $eB\rho$, not their absolute momentum. This geometric effect, particles of all momenta converging onto the same unstable trajectory at different times, results in a notably flat transverse beam profile with dramatically suppressed dispersion.

Consider two particles with the same betatron amplitude J but with different momenta $\langle p \rangle \pm p$, where $p > 0$, living around different dispersive orbits

$\sim (x_0, p_{x0}) \pm \frac{p}{\langle p \rangle} (\eta_x, \eta'_x)$. At the start of extraction t_0 , the $+p$ particle will witness a smaller separatrix than the $-p$ particle since chromaticity is negative and the tune is larger than the resonant tune. At some time t_1 , the reference momentum is $p_1 = eB_1\rho$ and the $+p$ particle orbiting around (x_1, p_{x1}) hits its separatrix, having $\omega_+ = |S_v| \sqrt{\frac{3\pi}{4} J \sqrt{3}}$ and $\delta_+ = \frac{1}{2\pi\xi_x} (\omega_+ - 2\pi(3Q_{x,0} - m))$. Note that $\delta_+ \equiv \frac{\langle p \rangle + p - p_1}{p_1}$ is independent of p in the separatrix equation. Some time later at t_2 , the reference momentum is $p_2 = eB_2\rho$ and the $-p$ particle is extracted. Since $J_+ = J_- = J$, we see that $\delta_- = \delta_+$. But $\delta_- \equiv \frac{\langle p \rangle - p - p_2}{p_2}$, so we find that $\frac{B_2}{B_1} = \frac{\langle p \rangle - p}{\langle p \rangle + p}$, implying that the initial dispersion (momentum-position correlations) gets transformed into correlation between momentum and extraction time. Furthermore, since dispersion is evaluated relative to the reference momentum, both particles are circulating around the same final closed-orbit position once they are extracted, and will follow the same subsequent trajectory between the $D3$ and the $D6$ septa.

Figure 6.1 compares the predicted separatrix from first-order normal form theory with symplectic tracking through the exact s -dependent lattice Hamiltonian, showing remarkable agreement. Also shown is the extracted bunch distribution, demonstrating that particles indeed extract along the predicted separatrix arms regardless of momentum offset.

This analysis reveals why the constant optics extraction scheme is so effective for NSRL: by maintaining fixed β_x and α_x throughout the spill while ramping the reference momentum (and thus ω through chromaticity), the separatrix remains geometrically fixed in real space. Particles of all momenta converge onto the same unstable arms, producing an extracted beam whose transverse position is nearly independent of momentum, exactly what is needed for uniform dose delivery in radiobiology experiments.

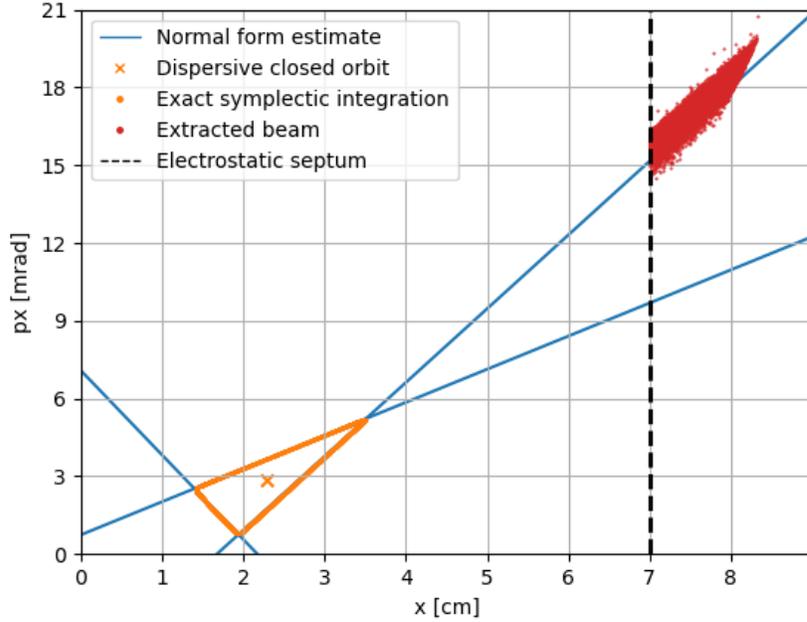


Figure 6.1: Comparison of normal form separatrix prediction for a particle with $\delta = 0.35\%$ and $J = 8.45\mu\text{m}$ with symplectic tracking through the exact s -dependent lattice Hamiltonian shows remarkable agreement. Also shown is the extracted bunch distribution. The lattice parameters are $\beta_x = 8.02\text{ m}$, $\alpha_x = -1.31$, $Q_x = 4.364$, $\xi_x = -7.6$, and $S_v = 9.27 + 8.63i$. This validates both the normal form analysis and the constant-optics extraction mechanism that enables dispersion suppression.

6.1.2 Lattice Elements and Extraction Channel

Implementing third-integer extraction in the Booster relies on several specialized elements that work in concert. The sextupole system comprises four families: one family of 24 horizontal chromatic sextupoles, one family of 24 vertical chromatic sextupoles, and two families of two drive sextupoles each. During chromatic slow extraction, the chromatic sextupole families are turned off to utilize the ring's natural chromaticity, while the drive families are powered to deliberately excite the third-integer resonance: strong enough to form a well-defined separatrix yet weak enough to permit controlled extraction over prac-

tical timescales. The distribution and relative strengths of the drive sextupoles determine not only the resonance strength but also the orientation of the stable and unstable regions in phase space; by tuning the two drive families, operators rotate the resonance pattern so that one unstable separatrix leg aligns with the extraction septum, thereby maximizing extraction efficiency.

The extraction channel begins with a thin magnetic septum, 0.76 mm thick and located in the D3 cell, positioned at the edge of the circulating beam's aperture to provide the initial separation of extracted particles from the circulating beam. This element must sit precisely at the extraction amplitude threshold: if it is too close to the beam core, circulating particles are intercepted and losses increase, whereas if it is too far, particles that have become unstable and are taking larger steps can hit the beam pipe on the other side of the septum. The optimum setting depends on the closed-orbit, septum configuration, beam emittance, and sextupole resonance strength, and has been established empirically during commissioning.

Additional bending is provided by a thick ejection magnetic septum with 15.2 mm thickness and 2.5 m length that imparts a 155 mrad kick. This septum supplies the majority of the extraction angle, steering the beam into the dedicated NSRL transport channel; its field must be set to match both the extraction energy and the desired trajectory.

A distinctive feature of the NSRL extraction line is a thin stripping foil—typically aluminum with thickness 25–100 μm —inserted in the channel. For heavy ions the foil strips residual electrons, and for all species the associated multiple Coulomb scattering broadens the angular distribution to produce uniform irradiation patterns. Beams extracted via third-integer resonance oth-

erwise have a narrow, highly directional angular distribution; the foil converts this into the broader spread required for radiobiology studies that demand uniform irradiation over extended areas with minimal angular bias relative to the beam's rest frame.

6.1.3 The Extraction Cycle and Spill Structure

A typical NSRL extraction cycle begins with injection and acceleration. Beam is injected into the Booster from the upstream Linac at 200 MeV kinetic energy, and the process—sharing the same challenges discussed in Chapter 5—must deliver beam with minimal emittance and good matching to the Booster lattice. The beam is then accelerated over approximately 200 ms to the extraction energy, typically 1.0 GeV depending on the ion species and experimental requirements.

Once flat-top energy is reached, RF is largely turned off and the Booster reference momentum is ramped slowly downward by decreasing the main power supply for dipoles and quadrupoles. Due to the large negative chromaticity, as the beam's effective momentum deviation increases, the effective tune decreases from its pre-extraction value (typically $Q_x \approx 4.36$) toward the extraction tune ($Q_x \approx 4.333$). This ramp must be slow enough to avoid an adiabatic resonance crossing, which would preserve invariants and prevent extraction, yet fast enough to complete within a reasonable flat-top duration. Typical tune ramp rates are 0.01 per second, implying extraction periods of a few seconds.

As the tune approaches resonance, the separatrix amplitude shrinks according to Equation 6.7, causing particles in the beam's amplitude distribution tail to begin crossing into unstable regions. These particles spiral outward on succes-

sive turns until reaching the thin magnetic septum and entering the extraction channel.

Once extraction begins, the system ideally reaches a quasi-steady state where the tune ramp speed, beam emittance, and resonance strength combine to produce a constant extraction rate. In this regime, the extracted particle flux dN/dt can be approximated by:

$$\frac{dN}{dt} \approx N_{\text{circ}} \cdot \int \rho(A_{\text{sep}}(\delta), \delta) \cdot \frac{dA_{\text{sep}}(\delta)}{dt} d\delta \quad (6.12)$$

where N_{circ} is the number of circulating particles, $\rho(A, \delta)$ is the amplitude density distribution, and $dA_{\text{sep}}(\delta)/dt$ is the rate at which the separatrix moves inward for a given momentum deviation δ . This expression reveals that uniform spill requires either a uniform amplitude distribution $\rho(A, \delta)$ or a dynamically adjusted tune ramp rate that compensates for variations in ρ .

In practice, power-supply ripple in the main dipoles and quadrupoles causes instantaneous tune modulation, creating fluctuations in the extraction rate. These non-uniformities pose challenges for experiments requiring precise dose control, motivating the spill-quality control techniques discussed in Section 6.3.

As the separatrix amplitude continues to decrease, eventually even particles in the beam core cross the threshold and are extracted. Once the circulating intensity drops below a threshold and the extraction cycle is concerned complete, any remaining beam is dumped (less than a few percent), and the extraction tune is jumped back to the nominal value in preparation for the next cycle.

6.2 Simulation and Validation

Accurate modeling of slow extraction requires integrating several computational elements: a realistic lattice model including magnetic field nonlinearities, symplectic tracking to preserve phase space structure over millions of turns, extraction threshold detection to identify when particles reach the septum, and normal form analysis to predict separatrix geometry from the lattice parameters. This section describes each component and presents validation results comparing tracking simulations against normal form predictions.

6.2.1 Booster Lattice Model

The Booster synchrotron uses separate-function magnets with residual multipole errors. Unlike combined-function machines, the Booster's dipoles and quadrupoles are distinct elements, though they share a common power supply arrangement. The main power supply, which is ramped during the acceleration cycle, provides the only supply for the dipoles and is also wound with 5 turns around the quadrupoles. The quadrupoles also have a dedicated power supply for tune control which only uses 1 winding and is not ramped. This means the optics are not *exactly* constant during extraction but approximately so.

The Booster lattice consists of 24 superperiods, each containing separate-function dipoles, quadrupoles, sextupoles, and drift spaces. At the extraction energy of 1.0 GeV for protons ($\gamma \approx 2.07$), the machine operates with tunes $Q_x \approx 4.6$, $Q_y \approx 4.7$ during acceleration, with Q_x adjusted to $Q_x = 4.333$ (corresponding to the $\nu_x = 13/3$ third-integer resonance) for extraction.

The lattice model was constructed using both Bmad and the new Beam-lines.jl package that is part of the SciBmad project in Julia. Neither Bmad nor SciBmad represent magnet bodies as transfer maps for tracking. Rather, their magnetic fields are used to symplectically integrate the respective Hamiltonians of each element in order. This ensures conservation of phase space volume (Liouville’s theorem) over arbitrary numbers of turns, critical for extraction studies where particles must be tracked for 10^5 – 10^6 turns as they slowly spiral outward, and even small symplecticity violations would accumulate to produce artificial emittance growth or damping.

The model incorporates several refinements beyond simple thin-lens approximations. In Bmad, the fringe fields of bends are represented using a second-order transfer map rather than direct field integration (no other elements include fringe fields by default) and these edge effects introduce nonlinear focusing that can influence trajectories, particularly for large-amplitude particles near extraction thresholds. In addition, the beam’s intrinsic momentum spread ($\Delta p/p \sim 10^{-3}$) was modeled explicitly through chromatic dependence of quadrupole strengths so that particles at different momenta experience distinct tunes and effective resonance strengths; this momentum dependence governs the relative extraction rates across the distribution and therefore impacts spill uniformity.

The complete lattice model consisted of approximately 600 elements (24 cells \times 25 elements per cell). This representation allowed turn-by-turn tracking of particle ensembles through the lattice in approximately 1–2 μ s per particle-turn on modern CPUs, enabling simulation of full extraction cycles with 10^4 particles over 10^5 turns in several hours of computation time.

6.2.2 Model Simulation and Normal Form Validation

The analytical predictions from Section 6.1.1 provide the theoretical foundation for understanding dispersion suppression, but require validation through detailed particle tracking to confirm their applicability to the real Booster lattice with its s -dependent fields, fringe effects, and higher-order multipoles.

The simulation framework tracks particles through the full Booster lattice using a third-order resonance model with chromaticity, synchrotron motion turned off (coasting beam assumption), and a time-dependent ramp on the magnetic rigidity. The main power supplies are ramped slowly (adiabatically) downward with primary windings around dipoles and secondary windings on quadrupoles which preserves the linear optics (beta functions and phase advance), ensuring approximately constant optics throughout the spill, which is the key to the dispersion suppression mechanism.

The Booster lattice is initialized at a working point of $\mathbf{Q}_0 = (4.36, 4.82)$ to avoid interference with the $Q_y = \{19/4, 24/5\}$ resonances, and a reference kinetic energy of 1 GeV/u for protons. A de-bunched Gaussian beam with normalized emittances of $\epsilon_{x,y,\text{norm}} = 3 \mu\text{m}$ and a momentum spread of $\sigma_\delta = 0.1\%$ is matched to the periodic Twiss parameters and populated along the reference closed orbit.

Two pairs of drive sextupoles are used to stimulate the third-integer horizontal resonance with minimal effect on chromaticity, while the two families of chromaticity-correcting sextupoles are turned off to facilitate chromatic slow extraction. With natural chromaticities $\xi = (-7.6, -2.1)$, this setup creates an unstable separatrix parametrized by δ . The separatrix orientation in (x, p_x) phase space remains approximately fixed as the reference momentum is slowly swept

through the resonance, thanks to the constant optics. This is a key distinction from conventional extraction schemes where tune changes would rotate the phase space structure and separatrix arms.

10,000 particles are initialized as a full 5D Gaussian distribution. As the tune decreases due to chromaticity and the energy ramp, particles slowly drift into the unstable region. After crossing the separatrix, particles become unstable and quickly spiral out away from their stable dispersive fixed point along the hyperbolic trajectories predicted by Equation 6.9. Those that cross over the thin magnetic septum are extracted, while those that collide with the septum cross-section or length are lost.

The normal form analysis from Section 6.1.1 was performed on the Booster one-turn map using `Bmad/PTC` to extract the virtual sextupole strength S_v and predict the separatrix geometry. The calculation yields $S_v = 9.27 + 8.63i$ at the nominal extraction working point, corresponding to a separatrix phase $\psi = \arg(S_v)/3 \approx 21.5^\circ$ and strength $|S_v| = 12.7$. These values were used to generate the predicted separatrix curves shown in Figure 6.1, which agree with particle tracking to within 1% in amplitude, remarkable given that the normal form truncates at first order.

6.2.3 Dispersion Suppression in the Extracted Beam

The most remarkable feature revealed by the simulations is the dramatic suppression of dispersion in the extracted beam: a central result that makes this extraction scheme suitable for NSRL's stringent uniformity requirements. A key diagnostic tracks the final transverse position of each extracted particle as a

function of their relative momentum $p/\langle p \rangle - 1$ at the D3 thin magnetic septum.

Figure 6.2 shows the extracted particle horizontal position versus momentum deviation relative to the bunch average. The plot reveals remarkably weak dependence on p : particles spanning a momentum spread of $\pm 0.3\%$ emerge with horizontal positions varying independently of momentum, the lack of correlation corresponding to a lack of dispersion. This is to be compared with the periodic dispersion function at the D3 septum location, which is $\eta_x = 2.93$ m, a suppression factor of approximately 10.

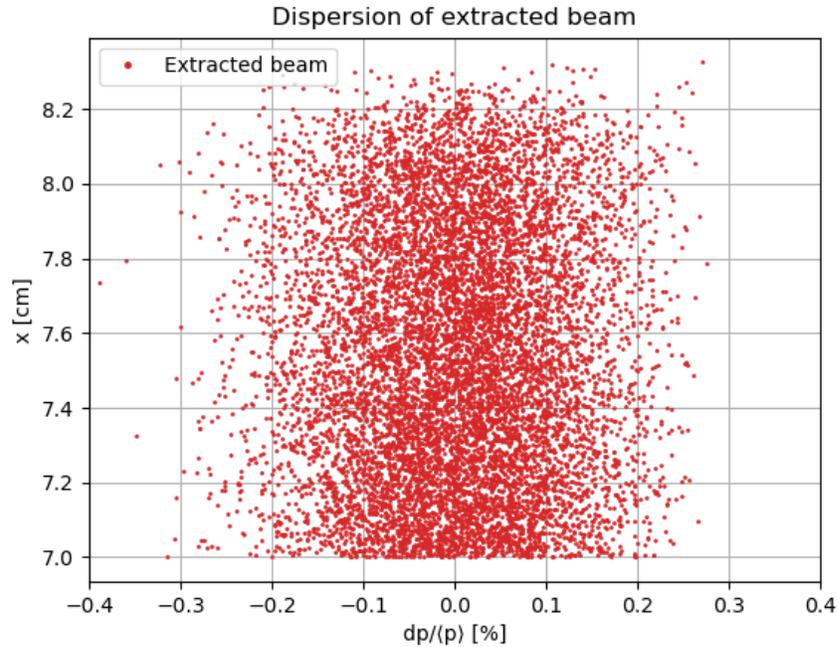


Figure 6.2: Extracted particles' horizontal position vs momentum deviation relative to bunch average at the D3 thin magnetic septum shows almost zero correlation, indicating dramatically suppressed beam dispersion. The near-vertical scatter demonstrates that particles of all momenta exit along common unstable trajectories defined by betatron amplitude, not momentum. This validates the normal form prediction that the constant-optics extraction scheme produces dispersion-free beams despite non-zero periodic dispersion at the septum.

This suppression is not an artifact of the simulation but rather a fundamen-

tal consequence of the constant-optics extraction mechanism described in Section 6.1.1. Because the Twiss parameters β_x and α_x remain fixed throughout the spill, the separatrix maintains a constant orientation and size in real (x, p_x) space throughout the tune ramp. Particles of all momenta converge onto the same unstable manifolds, the hyperbolic fixed point trajectories given by Equation 6.9, regardless of where their dispersive closed orbits lie. In the absence of RF energy bucketing, the extraction process itself transforms the correlation of $\langle x p \rangle_0$ into a correlation of $\langle x \delta \rangle_1 \propto \langle x p \rangle_0$, and decorrelating the position from the true momentum $\langle x p \rangle_1 \rightarrow 0$.

Within statistical error, we do not detect any residual momentum dependence in the simulated extracted beam distribution. The constant-optics extraction scheme suppresses dispersion by an order of magnitude compared to conventional extraction methods. This suppression is the key to NSRL's ability to deliver uniform dose distributions with minimal beam steering corrections, which is a practical validation of the normal form theory that will be confirmed experimentally in Section 6.2.4.

6.2.4 Experimental Validation of Dispersion Suppression

While the simulation results in Section 6.2 provide strong evidence for dispersion suppression in the constant-optics extraction scheme, experimental validation is essential to confirm that the predicted behavior occurs in the real machine. This section describes measurements performed at the Booster using the D6 extraction line flag to directly quantify the output beam dispersion and compare it to the periodic dispersion function, and demonstrates one to two orders

of magnitude suppression that validates both the normal form theory and the simulation predictions.

Measurement Methodology

To quantify the residual output dispersion, we exploited the Booster's ability to adjust the flattop beam momentum before initiating extraction. The radial position of the beam can be controlled via the radial loop in the RF system: by changing the RF frequency before disengaging the RF for coasting beam extraction, the beam's closed orbit radius (and thus its momentum) can be precisely adjusted.

Two radial offsets were chosen for comparison: $dR = -13.5$ mm and $dR = -6.5$ mm, corresponding to a relative momentum difference of $\Delta\delta = 0.93\%$ between the two conditions. This momentum difference was calculated from the Booster's transition gamma $\gamma_{tr} = 4.33$ and average radius $\rho = 13.866$ m using the standard relation:

$$\frac{\Delta p}{p} = \frac{1}{\gamma_{tr}^2 - \gamma^2} \frac{\Delta R}{\rho} \quad (6.13)$$

For each radial offset, the beam was allowed to de-bunch and coast, then the slow extraction was initiated with the standard third-integer resonance drive. A flag (a scintillating screen with calibrated position markers) was inserted into the D6 extraction vacuum chamber, located just over one FODO cell downstream from the initial D3 thin magnetic septum. This location was chosen to ensure that extracted particles had completed their hyperbolic escape from the separatrix and were traveling on stable trajectories in the extraction line, while still being close enough that beam envelope growth from natural emittance and

scattering remained minimal.

The flag width is precisely known to be $1.11'' = 2.819$ cm, with position markers labeled "L" and "R" at the left and right edges. By observing the beam spot position for the two momentum conditions and fitting Gaussians to the intensity distributions, the relative transverse displacement Δx can be measured and converted to an effective dispersion:

$$D_{\text{eff}} = \frac{\Delta x}{\Delta\delta} \quad (6.14)$$

Experimental Results and Comparison to Periodic Dispersion

Figure 6.3 shows overlaid images of the two beam spots on the D6 flag, corresponding to the two radial offsets. The blue spot corresponds to $dR = -13.5$ mm (lower momentum) and the grey spot to $dR = -6.5$ mm (higher momentum). The red "L" and "R" lines indicate the flag edges, providing the absolute position reference.

The measured difference between the spot centroids is approximately $\Delta x \sim 0.7 \pm 0.5$ mm. The uncertainty arises from several sources: the finite spot size ($\sigma_x \approx 3$ mm), shot-to-shot variations in extraction conditions, and systematic errors in the Gaussian fitting procedure. Nevertheless, the central value and its uncertainty bound allow us to quantify the effective dispersion:

$$D_{\text{eff}} \approx \frac{0.7 \pm 0.5 \text{ mm}}{0.0093} \approx 8 \pm 5 \text{ cm} \quad (6.15)$$

This measured effective dispersion should be compared to the periodic dispersion function at the D6 flag location. The Booster lattice model gives

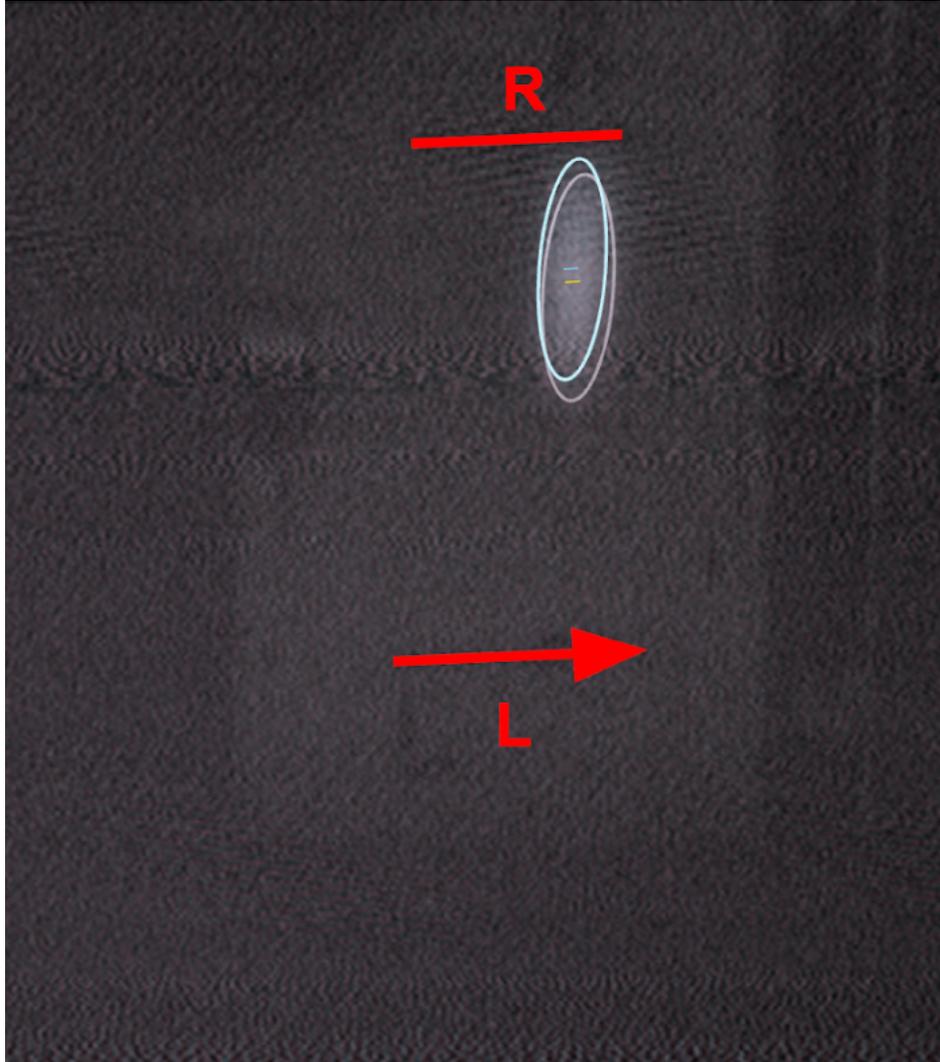


Figure 6.3: D6 flag showing two overlaid beam spots corresponding to radial offsets of $dR = -13.5$ mm (blue) and $dR = -6.5$ mm (grey). The red "L" and "R" lines are 1.11" apart, providing calibrated position reference. The measured horizontal displacement between the two spots is $\Delta x \approx 0.7 \pm 0.5$ mm despite a momentum difference of $\Delta\delta = 0.93\%$, demonstrating dramatic dispersion suppression compared to the periodic value of $\eta_x = 2.93$ m at this location.

$\eta_x(D6) = 2.93$ m at the nominal extraction working point. Thus, the suppression factor is:

$$\text{Suppression factor} = \frac{\eta_x}{D_{\text{eff}}} = \frac{2.93 \text{ m}}{0.08 \pm 0.05 \text{ m}} \approx 20\text{--}100 \quad (6.16)$$

We conclude that the dispersion is suppressed by one to two orders of mag-

nitude from the periodic value.

Validation of the Normal Form Theory

This experimental result provides strong validation of the theoretical predictions. The dramatic dispersion suppression occurs precisely because the Twiss parameters β_x and α_x remain approximately fixed during the spill while the reference momentum is ramped, keeping the separatrix geometrically fixed in real space and ensuring the separatrix arm does not mix with the momentum-time correlations, exactly as predicted by the Kobayashi normal-form analysis in Section 6.1.1. The measurement also confirms the hyperbolic fixed-point dynamics: extracted particles follow the escape trajectories given by Equation 6.9, which are defined by betatron amplitude rather than momentum and thus explain the near-independence of transverse position on δ despite non-zero periodic dispersion. Finally, within statistical error, the simulations exhibit no residual dispersion, consistent with the experimental finding and validating both the lattice model and the tracking framework developed for this study.

6.3 Spill Quality Control and Future Directions

While the experimental validation in Section 6.2.4 confirms the dispersion suppression mechanism, operational experience reveals that spill uniformity remains limited by power supply ripple and beam distribution effects. This section briefly outlines control strategies to improve spill quality, demonstrating connections to the machine learning and control themes developed in earlier chapters.

6.3.1 Power Supply Ripple and Mitigation Strategies

Power-supply ripple at 60–720 Hz (harmonics of 60 Hz AC) creates tune modulation $\Delta Q_x \sim 10^{-3}$ that couples to extraction-rate fluctuations. Two complementary mitigation strategies parallel techniques from earlier chapters. One is RF empty-bucket channeling, which modulates the RF voltage to create longitudinal density structure that “gates” extraction during the quiet portions of the ripple cycle; experiments at CERN indicate roughly a factor-of-two reduction in spill ripple, provided the RF phase is precisely synchronized to the power-supply ripple. The other is feed-forward compensation: measure the ripple in real time and drive fast quadrupole correctors to cancel the induced tune modulation. A newly installed AC quadrupole at the Booster enables such feed-forward control at the source, though practical deployment requires kHz-bandwidth actuation and sub-millisecond latency.

6.3.2 Machine Learning and Future Control Strategies

The parallels between extraction control and the injection optimization of Chapter 5 are striking: both involve high-dimensional parameter spaces, limited diagnostics, and model uncertainties. Bayesian optimization could autonomously tune closed-orbit bumps and tune-ramp parameters to maximize spill efficiency, while reinforcement learning approaches could discover effective feed-forward strategies for the AC quadrupole modulation based on ripple measurements. These applications remain for future work but illustrate the broader applicability of data-driven control methods across accelerator systems.

6.3.3 Summary and Connections to Dissertation Themes

This chapter has demonstrated that slow extraction to NSRL, though not directly addressing polarization, exemplifies the dissertation's unifying themes. The constant-optics extraction scheme achieves one to two orders of magnitude dispersion suppression through careful resonance engineering and is validated experimentally. The normal form analysis parallels the spin resonance theory of Chapter 4, the symplectic tracking extends the methods of Chapter 3, and the future control strategies connect to the machine learning approaches of Chapter 5.

Both spin preservation and extraction involve navigating resonance-dominated phase space where small changes produce large effects. The common philosophy of high-fidelity modeling to predict instabilities, robust diagnostics to verify machine state, and data-driven corrections to maintain desired behavior unifies these apparently disparate applications.

CHAPTER 7

CONCLUSIONS AND FUTURE DIRECTIONS

The preservation of polarized hadron beams demands a synthesis of first-principles theory, high-fidelity modeling, operational diagnostics, and data-driven control. This dissertation pursued that synthesis across the RHIC injector chain and the proposed EIC Hadron Storage Ring, developing tools that span from Maxwellian field representations to symmetry-based snake configurations and autonomous injector tuning. The results demonstrate that polarization control is tractable when physics realism, computational rigor, and operations-aware optimization are designed in concert.

7.1 Key Contributions

7.1.1 Mathematical and Computational Infrastructure

Chapter 2 established the mathematical backbone required for million-turn spin studies. By embedding the T-BMT equation within a symplectic Hamiltonian framework and coupling it to Maxwellian field descriptions, the dissertation provided a rigorous context in which subsequent models could operate. The integration of TPSA-based normal forms, invariant frame fields, and cylindrical multipole expansions enabled the faithful calculation of amplitude-dependent spin tunes and resonance strengths across the wide energy ranges relevant to RHIC and the HSR. These foundations also motivated the combined use of the `Bmad` toolkit and the emerging Julia-based `SciBmad` ecosystem, establishing a reproducible computational stack for polarized-beam studies.

7.1.2 Field Realism and Symplectic Tracking in the AGS

Chapter 3 confronted the long-standing shortcomings of non-Maxwellian field models in the AGS. By fitting cylindrical multipole expansions to detailed finite-element data and integrating them within symplectic kick-drift-kick schemes, the work eliminated artificial emittance growth and recovered stable polarization transport through the entire acceleration ramp. The resulting stable beam transmission for representative emittances demonstrates that realistic snake modeling, coupled with lattice matching across 100 measured energy points, tightens the gap between simulation and reality. The chapter articulated an end-to-end workflow that takes magnet CAD models through finite-element solvers into symplectic tracking, a template that can be replicated for future facilities.

7.1.3 Spin Dynamics from RHIC Operations to EIC Design

Chapter 4 bridged operational diagnostics and future-lattice design. Detailed tracking of the 2021 RHIC partial-snake incident showed that orbit distortions, rather than snake asymmetries, dominated depolarization, underscoring the primacy of alignment and closed-orbit control in existing machines. Building on that insight, two complementary strategies were developed for the HSR. Method I demonstrated that modest, optics-feasible adjustments to inter-snake betatron phase advance can recover $\geq 97\%$ polarization transmission for helium-3 across the most dangerous resonance band. Method II delivered the central discovery of this dissertation: the Doubly Lee–Courant (DLC) snake arrangement, which enforces local π spin phase advances across every consecutive snake pair. The optimal DLC scheme achieves near-perfect polar-

ization transport in realistic lattices with minimal energy dependence, promoting a symmetry-driven design principle that supersedes brute-force numerical searches.

7.1.4 Machine Learning for Injector Operations

Chapter 5 translated the theoretical and modeling advances into operational leverage. Bayesian optimization applied to the Booster-to-AGS (BtA) line recovered spoiled injection efficiencies from 65% to beyond 90% within minutes using only scalar intensity diagnostics, while accommodating the sparse BPM coverage and noisy profile monitors inherent to the injector complex. The framework demonstrated that carefully constrained Gaussian-process surrogates, informed by accelerator physics priors, can provide safe, reproducible tuning knobs for day-to-day operations. The chapter also laid out the conceptual architecture for reinforcement learning policies that incorporate multiwire profiles and steering constraints, pointing toward autonomous, polarization-aware injector tuning.

7.1.5 Extending the Toolkit to Controlled Extraction

Chapter 6 broadened the scope to slow extraction for NSRL, showing that the same analytical and computational machinery governs spill quality. First-order normal form analysis predicted that constant-optics third-integer extraction would suppress dispersion by one to two orders of magnitude, a result confirmed by symplectic tracking and by D6 flag measurements that observed an effective dispersion of 8 ± 5 cm against a periodic value of 2.93 m. This val-

idation links theoretical resonance control, high-fidelity tracking, and beamline diagnostics in a second operational context, demonstrating the generality of the dissertation's methodology.

7.2 Future Directions

7.2.1 Toward an End-to-End EIC Digital Twin

The tools developed here set the stage for integrated source-to-collision simulations that couple injector tuning, ramp optics, and collider snake design within a single workflow. Extending the symplectic, Maxwellian modeling chain to include space-charge, intrabeam scattering, and coherent beam-beam effects will enable quantitative forecasts of polarization margins for the full EIC program.

7.2.2 Implementing the Doubly Lee–Courant Baseline

Realizing the optimal DLC scheme in hardware will require detailed engineering studies of orbit excursions and snake field quality. Future work should evaluate error sensitivity under realistic alignment and field tolerances, and design commissioning procedures that explore the 1D space of DLC snake configurations. Comparative studies of hybrid strategies, using optimal DLC as a baseline supplemented by modest phase-advance trims or orbit-feed-down bumps, will quantify the operational latitude available when real machines deviate from design ideals.

7.2.3 Autonomous Data-Driven Control

The successful application of Bayesian optimization invites broader adoption of data-driven control. Future campaigns should integrate contextual priors (e.g., beam species, bunch intensity) and leverage transfer learning to warm-start optimizers from historical data. Reinforcement learning policies, trained on calibrated digital twins and refined online, can simultaneously minimize losses, preserve polarization, and respect hardware limits. Extending these approaches to spill regulation, for example by driving the Booster's AC quadrupole with feedback from fast spill monitors, would close the loop on slow extraction as well.

7.2.4 Collective Effects

Incorporating collective effects into the spin dynamics framework remains an open frontier. Future research should develop symplectic models of space-charge-induced spin diffusion, coherent beam-beam depolarization, and intra-beam scattering within the Maxwell-consistent paradigm established here. Coupling these effects to the existing tools will enable comprehensive studies of polarization preservation under high-intensity, high-luminosity conditions, informing both machine design and operational strategies for the EIC and beyond.

7.3 Closing Remarks

Polarized hadron beams offer unique insights into the spin structure of matter, but their preservation demands a confluence of theoretical rigor, computational fidelity, and operational savvy. This dissertation has advanced that confluence across multiple fronts, delivering tools and insights that bridge from first principles to real-world machines. As the field moves toward the ambitious goals of the EIC era, the methodologies and discoveries presented here will serve as a foundation for continued innovation in polarized beam dynamics and control.

BIBLIOGRAPHY

- [1] J. Takano et al. Optimization of AGS Polarized Proton Operation with the Warm Helical Snake. In *Proc. PAC'05*, number 21 in Particle Accelerator Conference, pages 1003–1005. JACoW Publishing, Geneva, Switzerland, 2005.
- [2] Y. S. Derbenev and A. M. Kondratenko. Acceleration of polarized particles to high energies in synchrotrons. *10th International Conference on High Energy Accelerators*, 10:70–75, 1977.
- [3] Y. S. Derbenev and A. M. Kondratenko. On the Possibilities to Obtain High-Energy Polarized Particles in Accelerators and Storage Rings. *AIP Conf. Proc.*, 51:292–306, 1979.
- [4] M. Bai et al. Polarized proton collisions at 205 gev at rhic. *Phys. Rev. Lett.*, 96:174801, May 2006.
- [5] J. Takano et al. Field Measurements in the AGS Warm Snake. In *Proc. EPAC'04*, number 9 in European Particle Accelerator Conference, pages 2116–2118. JACoW Publishing, Geneva, Switzerland, 9 2004.
- [6] E. Willen et al. Superconducting Helical Snake Magnet for the AGS. In *Proc. PAC'05*, number 21 in Particle Accelerator Conference, pages 2935–2937. JACoW Publishing, Geneva, Switzerland, 2005.
- [7] J. Berg et al. Lattice design for the hadron storage ring of the Electron-Ion Collider. In *Proc. IPAC'23*, number 14 in International Particle Accelerator Conference, pages 903–905. JACoW Publishing, Geneva, Switzerland, 9 2023.
- [8] Mathias Vogt et al. Bounds on the maximum attainable equilibrium spin polarization of protons at high energy in hera. Technical report, Deutsches Elektronen-Synchrotron (DESY), Hamburg (Germany), 2000.
- [9] E. Hamwi, J.P. Devlin, and G.H. Hoffstaetter. Polarization transmission in the electron-ion collider's hadron storage ring. *Phys. Rev. Accel. Beams*, 2025. To be published.
- [10] D. Sagan. Bmad: A relativistic charged particle simulation library. *Nucl. Instrum. Meth.*, A558(1):356–359, 2006. Proceedings of the 8th International Computational Accelerator Physics Conference.

- [11] V. Bargmann, L. Michel, and V. L. Telegdi. Precession of the polarization of particles moving in a homogeneous electromagnetic field. *Phys. Rev. Lett.*, 2:435–436, May 1959.
- [12] M. Conte et al. Beam optics of the Dirac particle with anomalous magnetic moment. *Part. Accel.*, 56:99–125, 1996.
- [13] K. Heinemann and D. P. Barber. The semiclassical Foldy-Wouthuysen transformation and the derivation of the Bloch equation for spin 1/2 polarized beams using Wigner functions. In *15th Advanced ICFA Beam Dynamics Workshop on Quantum Aspects of Beam Physics*, pages 695–701, 9 1998.
- [14] P. J. Mohr et al. CODATA recommended values of the fundamental physical constants: 2022. *Rev. Mod. Phys.*, 97:025002, Apr 2025.
- [15] Ronald D. Ruth. Single-particle dynamics in circular accelerators. *AIP Conference Proceedings*, 153(1):150–235, 02 1987.
- [16] S. Blanes, F. Casas, J. A. Oteo, and J. Ros. The Magnus expansion and some of its applications. *Physics Reports*, 470(5):151–238, 2009.
- [17] H. Goldstein, C.P. Poole, and J.L. Safko. *Classical Mechanics*. Addison-Wesley Longman, 2001.
- [18] N. J. Hitchin, G. B. Segal, and R. S. Ward. *Integrable Systems: Twistors, Loop Groups, and Riemann Surfaces*. Oxford University Press, 05 1999.
- [19] A. J. Dragt and J. M. Finn. Lie Series and Invariant Functions for Analytic Symplectic Maps. *J. Math. Phys.*, 17:2215–2227, 1976.
- [20] E. Forest, M. Berz, and J. Irwin. Normal form methods for complicated periodic systems. *Part. Accel.*, 24:91–107, 1989.
- [21] P.-L. Giscard et al. An exact formulation of the time-ordered exponential using path-sums. *J. Math. Phys.*, 56(5):053503, 05 2015.
- [22] D. P. Barber, J. A. Ellison, and K. Heinemann. Quasiperiodic spin-orbit motion and spin tunes in storage rings. *Phys. Rev. ST Accel. Beams*, 7:124002, Dec 2004.
- [23] Klaus Heinemann, Desmond P. Barber, James A. Ellison, and Mathias Vogt.

A detailed and unified treatment of spin-orbit systems using tools distilled from the theory of bundles, 2015.

- [24] Helmut Wiedemann. *Hamiltonian Resonance Theory*, pages 479–502. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.
- [25] Georg Heinz Hoffstaetter. *High energy polarized proton beams: a modern view*, volume 218. Springer Science & Business Media, 2006.
- [26] G. H. Hoffstaetter, M. Vogt, and D. P. Barber. Higher-order effects in polarized proton dynamics. *Physical Review Special Topics-Accelerators and Beams*, 2(11):114001, 1999.
- [27] E. Hamwi and G. Hoffstaetter. Higher-order spin depolarization analysis. In *Proc. IPAC'23*, number 14 in International Particle Accelerator Conference, pages 3133–3135. JACoW Publishing, Geneva, Switzerland, 9 2023.
- [28] M. Venturini and A. J. Dragt. Accurate computation of transfer maps from magnetic field data. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 427(1-2):387–392, 1999.
- [29] M. Borland. Elegant: A flexible sdds-compliant code for accelerator simulation. Technical report, Argonne National Lab., IL (US), 08 2000.
- [30] N. Tsoupas et al. Acceleration of polarized protons in the ags. Technical report, Brookhaven National Laboratory, 02 2010.
- [31] W. Meng and M. Tanaka. Three Dimensional Field Analysis for the AGS Combined Function Magnets. In *Proc. PAC'93*, pages 2907–2910. JACoW Publishing, Geneva, Switzerland, 1993.
- [32] F. Meot, Y. Dutheil, R. C. Gupta, H. Huang, N. Tsoupas, and J. Takano. AGS Snake Stories. In *Proc. IPAC'14*, pages 1220–1222. JACoW Publishing, Geneva, Switzerland, 2014.
- [33] V. Schoefer et al. Increasing the AGS Beam Polarization with 80 Tune Jumps. In *Proc. IPAC'12*, pages 1015–1019. JACoW Publishing, Geneva, Switzerland, 2012.
- [34] V. Schoefer. Using betatron coupling to suppress horizontal intrinsic spin resonances driven by partial snakes. *Phys. Rev. Accel. Beams*, 24:031001, Mar 2021.

- [35] V. Schoefer et al. Correction of horizontal partial snake resonances with pulsed skew quadrupoles at the Brookhaven AGS. In *Proc. IPAC'24*, number 15 in International Particle Accelerator Conference, pages 1000–1002. JACoW Publishing, Geneva, Switzerland, 7 2024.
- [36] Brookhaven National Laboratory. Ags design report. Technical report, Brookhaven National Laboratory, Upton, NY, USA, 1966.
- [37] R. E. Thern and E. Bleser. The dipole fields of the ags main magnets. Technical report, Brookhaven National Laboratory, 01 1996.
- [38] R. E. Thern and E. Bleser. The quadrupole and sextupole fields of the ags main magnets. Technical report, Brookhaven National Laboratory, 03 1996.
- [39] Y. Dutheil, F. Meot, and N. Tsoupas. A Model of the AGS Based on Stepwise Ray-Tracing Through the Measured Field Maps of the Main Magnets. In *Proc. IPAC'12*, pages 1395–1397. JACoW Publishing, Geneva, Switzerland, 2012.
- [40] P. Menga, L. Ahrens, and R. Schroeder. Ags bnl tests. Technical report, Brookhaven National Laboratory, 08 2007.
- [41] R. Connolly et al. RHIC-Style IPMs in the Brookhaven AGS. In *Proc. IBIC'14*, pages 39–41. JACoW Publishing, Geneva, Switzerland, 2014.
- [42] G. Bunce et al. The ags cni polarimeter. *AIP Conference Proceedings*, 667(1):77–80, 05 2003.
- [43] Y. Dutheil et al. Optimization of the pp AGS Zgoubi Model in the Low Energy Range. In *Proc. IPAC'14*, pages 3089–3091. JACoW Publishing, Geneva, Switzerland, 2014.
- [44] E. Hamwi, G. H. Hoffstaetter, V. Schoefer, and W. Lin. Comprehensive modeling of Siberian Snakes in BNL's AGS: symplectic tracking and optical compensation. In *Proc. IPAC'24*, number 15 in International Particle Accelerator Conference, pages 686–689. JACoW Publishing, Geneva, Switzerland, 7 2024.
- [45] Haruo Yoshida. Construction of higher order symplectic integrators. *Physics Letters A*, 150(5):262–268, 1990.

- [46] Ferdinand Willeke. Electron Ion Collider Conceptual Design Report 2021. Technical report, BNL, 2 2021.
- [47] F. Méot et al. Rhic run 22, 9 o'clock, a snake in the blue. Technical report, Brookhaven National Laboratory (BNL), Upton, NY (United States), 06 2022.
- [48] E. Hamwi and G. Hoffstaetter. Proton polarization in RHIC with partial Siberian snakes. In *Proc. IPAC'24*, number 15 in International Particle Accelerator Conference, pages 690–693. JACoW Publishing, Geneva, Switzerland, 7 2024.
- [49] François Méot. A possible origin of the 10 15 Technical report, Brookhaven National Laboratory, 07 2024.
- [50] H. Witte et al. The Interaction Region of the Electron-Ion Collider EIC. In *Proc. IPAC'21*, International Particle Accelerator Conference, pages 2574–2577. JACoW Publishing, Geneva, Switzerland, 5 2021.
- [51] Georg H Hoffstaetter. Optimal axes of siberian snakes for polarized proton acceleration. *Physical Review Special Topics—Accelerators and Beams*, 7(12):121001, 2004.
- [52] K. Hock et al. Simulations of polarized helions in the HSR. In *Proc. IPAC'24*, number 15 in International Particle Accelerator Conference, pages 1634–1636. JACoW Publishing, Geneva, Switzerland, 7 2024.
- [53] E. Hamwi et al. Phase variation for snake matching in the eic's hsr. In *Proc. NAPAC2025*, number 2025 in North American Particle Accelerator Conference, pages 547–549. JACoW Publishing, Geneva, Switzerland, 08 2025.
- [54] D. Sagan. A Superconvergent Algorithm for Invariant Spin Field Stroboscopic Calculations. In *Proc. IPAC'18*, pages 145–148. JACoW Publishing, Geneva, Switzerland, 2018.
- [55] D. Sagan and J. C. Smith. The tao accelerator simulation program. In *Proceedings of the 2005 Particle Accelerator Conference*, pages 4159–4161. IEEE, 2005.
- [56] C. Hall et al. Simulation of the ion profile monitors in the Brookhaven AGS. presented at IPAC'25, Taipei, Taiwan, Jun. 2025, paper MOPS082, unpublished.

- [57] B. Dhital et al. An automated quad scan based emittance measurement software. In *Proc. IPAC'24*, number 15 in International Particle Accelerator Conference, pages 3453–3456. JACoW Publishing, Geneva, Switzerland, 7 2024.
- [58] E. Hamwi et al. Application of bayesian optimization to bta injection at bnl. In *Proc. NAPAC2025*, number 2025 in North American Particle Accelerator Conference, pages 58–60. JACoW Publishing, Geneva, Switzerland, 08 2025.
- [59] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 11 2005.
- [60] J. T. Wilson, F. Hutter, and M. P. Deisenroth. Maximizing acquisition functions for bayesian optimization, 2018.
- [61] K. Rajput et al. Harnessing the power of gradient-based simulations for multi-objective optimization in particle accelerators. *Machine Learning: Science and Technology*, 6(2):025018, apr 2025.
- [62] C. Szepesvári and M. L. Littman. Generalized markov decision processes: Dynamic-programming and reinforcement-learning algorithms. In *Proceedings of International Conference of Machine Learning*, volume 96, 1996.
- [63] T. Haarnoja et al. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. Pmlr, 2018.
- [64] E. Hamwi et al. Minimizing dispersion through resonant extraction for bnl's nsrl. In *Proc. NAPAC2025*, number 2025 in North American Particle Accelerator Conference, pages 515–517. JACoW Publishing, Geneva, Switzerland, 08 2025.
- [65] Y. Kobayashi and H. Takahashi. Improvement Of The Emittance In The Resonant Beam Ejection. In *6th International Conference on High-Energy Accelerators*, pages 347–351, 1967.
- [66] S. T. Wang, V. Khachatryan, and P. Nishikawa. Lattice design and experimental studies of nonlinear resonance at the cornell electron storage ring. *Phys. Rev. Accel. Beams*, 26:104001, Oct 2023.
- [67] L. Bongini, A. Bazzani, G. Turchetti, and I. Hofmann. Analysis of a model

for resonant extraction of intense beams by normal forms and frequency map. *Phys. Rev. ST Accel. Beams*, 4:114201, Nov 2001.