

UNIVERSITY OF CALIFORNIA

Los Angeles

Successful Beam-beam Tuneshift Compensation

A dissertation submitted in partial satisfaction

of the requirements for the degree

Doctor of Philosophy in Physics

by

Kip Aaron Bishofberger

2005

© Copyright by
Kip Aaron Bishofberger
2005

The dissertation of Kip Aaron Bishofberger is approved.

Harold Fetterman

Claudio Pellegrini

Alfred Wong

James Rosenzweig, Committee Chair

University of California, Los Angeles

2005

To my parents, for their belief in me.

Table of contents

Chapter 1:	Introduction	1
1.1.	Fermilab overview	2
1.1.1.	Basic Tevatron operations	4
1.1.2.	Getting there	7
1.2.	Tevatron Overview	9
1.2.1.	Establishing collisions	10
1.2.2.	Luminosity	12
1.2.3.	Detectors	14
1.2.4.	Pushing the limits — beam-beam effects	16
1.3.	Tevatron Electron Lens overview	20
1.4.	This thesis project	23
Chapter 2:	Particle motion in a circular accelerator	25
2.1.	Matrix formulation of linear dynamics	26
2.2	Building a machine lattice	31
2.2.1	Drift sections	31
2.2.2	Quadrupole fields	33
2.2.3	Equation of betatron motion	36
2.2.4	Transverse bunch properties: tune and emittance	41

2.3	The real world: imperfections	46
2.3.1	Effects of off-momentum particles	46
2.3.2	Resonant motion	52
2.3.3	Focusing errors	57
2.4	Beam-beam interactions	59
2.4.1	Linear weak-strong beam-beam focusing	60
2.4.2	Amplitude-dependent tuneshift	64
Chapter 3:	Requirements for tuneshift compensation	67
3.1	Specific goals of beam-beam compensation	69
3.2	Tuneshift mechanics using an electron beam	72
3.2.1	Derivation of antiproton tuneshift	73
3.2.2	Minimum interaction dimensions	76
3.2.4	Anode potential	78
3.2.5	Two tunes, two lenses	80
Chapter 4:	Electron-beam production and propagation	83
4.1	Generation of an electron beam	84
4.1.1	Electrons from a hot cathode	85
4.1.2	Space-charge limited flow	90
4.2	Cyclotron motion in a solenoid	94
4.2.1	General motion in a solenoid with space charge	95
4.2.2	Combining electrons and antiprotons	98
4.2.3	Gun and collector solenoids	101
4.3	Longitudinal Dynamics	105
4.3.1	Overcoming the space-charge potential	105

4.3.2	Beam-pipe acceptance	109
4.3.3	Electron-beam speed	114
4.3.4	Electron charge density	117
Chapter 5:	The Tevatron Electron Lens apparatus	122
5.1	Magnetic subsystem	124
5.1.1	Solenoids	124
5.1.2	Dipole correctors	129
5.1.3	Peripheral subsystems	133
5.1.4	Straightness of field lines	135
5.2	Electron gun	140
5.2.1	Mechanical design of T _E L electron guns	141
5.2.2	Gun electrical properties	147
5.3	Beam pipe and collector	150
5.3.1	Beam-pipe limitations on the electron beam	151
5.3.2	Electrodes along the beam pipe	157
5.3.3	Collector design	165
5.3.4	Properties of a depressed collector	168
5.4	Electrical subsystem	171
5.4.1	Basic model of electrical layout	171
5.4.2	Anode modulation	177
Chapter 6:	T_EL integration	180
6.1	Instabilities	180
6.1.1	Transverse-Mode Coupling Instability I	181
6.1.2	Transverse-Mode Coupling Instability II	185

Chapter 7:	Successful results of T _{FL} operation	191
7.1	Tuneshift measurement	192
7.1.1	Comparing tuneshift with theory	194
7.1.2	Translational alignment	199
7.1.3	Additional tuneshift dependencies	206
7.2	Successful operation	209
7.2.1	Controlling particle losses	210
7.2.2	Comparison of electron guns	214
7.2.3	Successful reduction of emittance growth	218
7.3	Other observations	222
7.3.1	Gentle bunch collimation	222
7.3.2	Effects of fluctuations	226
7.4	Conclusions	232
References	236

List of figures

1.1	Simplified cartoon of bunch spacing in the Tevatron.	5
1.2	Cute cartoon of the accelerators involved at Fermilab.	7
1.3	Illustration of loading antiprotons into the Tevatron.	11
1.4	Cartoon of proton and antiproton bunches colliding.	12
1.5	Cross-section views of the CDF and DØ detectors.	15
2.1	The reference frame attached to the orbit of the moving bunches.	29
2.2	Two examples of a particle propagating through optical elements.	32
2.3	Sketch of a particle passing through a series of FODO cells.	37
2.4	Illustration of three particles traveling through a FODO lattice.	42
2.5	Phase space of a particle circling a ring 500 times.	43
2.6	Tune plots depicting sum resonances.	57
2.7	Plot of Equation 2.30, normalized to the proton bunch size σ_r .	62
2.8	Diagram illustrating the tuneshift experienced by antiprotons.	65
3.1	Examples of schemes for beam-beam compensation.	70
3.2	Cartoon of the interaction of the electron beam and the antiprotons.	73
3.3	Cartoon of ramping the T _{EL} current between bunches.	78

4.1	Graph of Equation 3.2.	86
4.2	Cartoon showing how individual atoms create a conduction band.	87
4.3	Correlation between lattice spacing and work function.	88
4.4	Motion of a particle in magnetic and electric fields.	98
4.5	Field lines bending from the main solenoid to the collector.	104
4.6	Sketch of the potential energy that an electron beam passes through.	106
4.7	Plot of Equation 3.14.	107
4.8	Plot of $V_{s-c}^{\max}/V_{\text{cat}}$ as a function of r_e/r_{pipe} .	110
4.9	Electron-beam current versus the space-charge potential.	112
4.10	Beam-pipe acceptance for different sized electron beams.	113
4.11	Relationship between the beam speed and the total current.	115
4.12	Charge density at the center of an electron beam.	119
4.13	Dependence of charge density on cathode voltage.	120
5.1	A drawing of the solenoid magnets in the TEL.	125
5.2	Field strengths produced by the gun and collector solenoids.	127
5.3	Sketch of the placement and action of the dipole correctors.	130
5.4	Longitudinal strength of the dipole correctors.	131
5.5	Drawing of the gimble and photograph of the cart.	136
5.6	Cartoon of the technique used to measure magnetic field lines.	137
5.7	Drawing of one horizontal and vertical field line at 4 T.	138
5.8	Changes in central field lines with magnetic-field strength.	139
5.9	Comparison of several “parallel” field lines near the central axis.	140
5.10	Cross-section of the flattop electron gun.	142
5.11	The current density profile of the electron beam.	143

5.12	Total current expended by the cathode.	145
5.13	Cross-section of the more recent Gaussian gun.	146
5.14	Current-density profile typical of the electron guns.	147
5.15	Current production for different cathode temperatures.	149
5.16	Perveance of each of the guns.	150
5.17	Cross-section of the T _{FL} beam pipe.	151
5.18	Dependence of collector current on the anode voltage.	153
5.19	Plot of charge density on anode voltage.	155
5.20	Sketch of the electrodes in the T _{FL} .	158
5.21	Pictures of different diagnostics.	159
5.22	Example of optimal BPM positions.	162
5.23	Position dependence on length of electron beam.	164
5.24	Scaled drawing of the collector cross-section.	165
5.25	Measurement of collector acceptance.	170
5.26	Oversimplified schematic of the T _{FL} circuitry.	172
5.27	Realistic circuit of the T _{FL} .	174
5.28	Oscilloscope trace illustrating a typical pulse.	176
5.29	Signals produced by the anode modulator circuit.	179
6.1	Sketch of head-tail coupling, leading to TMCI.	184
6.2	Minimum field strength in the T _{FL} solenoid to prevent TMCI.	188
6.3	Magnetic field strength versus the synchrotron tune.	189
6.4	Minimum magnetic field strength versus T _{FL} tuneshift.	190
6.5	Minimum field strength versus the horizontal tune.	191

7.1	Horizontal Schottky spectra with and without the T _{EL} .	194
7.2	Observed horizontal tunes shift as it depends on current.	197
7.3	Illustration of how displacements can cause tunes shifts.	201
7.4	Measured tunes shift as the electron-beam path is displaced.	202
7.5	Measurement of tunes shift at various angles of the electron beam.	205
7.6	Two examples of Schottky spectra.	206
7.7	Tunes shift caused by a delayed electron pulse.	208
7.8	Antiproton tunes shift as a function of cathode voltage.	209
7.9	Typical bunch lifetime of a proton bunch interacting with the T _{EL} .	212
7.10	Analysis of how bunch losses compare with displacement.	214
7.11	Scans of working points with the T _{EL} .	216
7.12	Evolution of three bunch sizes during collisions.	222
7.13	Evolution of bunches without the T _{EL} .	223
7.14	Display of the bunches passing through the beam pulse.	224
7.15	Graphs of the T _{EL} acting as a “gentle collimator.”	226
7.16	Multiple traces of varying current amplitudes.	228
7.17	Effects of large current fluctuations on particles.	229
7.18	Schottky power versus the T _{EL} trigger delay.	230
7.19	Trace the T _{EL} agitating two antiproton bunches.	232
7.20	Antiproton-bunch data after the T _{EL} interfered with the store.	233

List of tables

1.1	Typical parameters associated with Run II.	3
2.1	Tevatron-related quantities.	27
2.2	Antiproton and proton parameters during recent stores.	28
5.1	Quantities associated with the TEL.	123
5.2	Attributes of the dipole correctors.	133
7.1	Growth of the rms vertical bunch size.	221

Acknowledgments

This thesis project would have been impossible without the collaboration of many people, to whom I owe great gratitude. As my advisor at UCLA, James Rosenzweig offered me the opportunity to pursue this exciting, profound, and unique experiment on the Tevatron at Fermilab. The experience and wisdom I have gained from him cannot be overstressed.

At Fermilab, Vladimir Shiltsev became my mentor and provided counsel, encouragement, support, and opportunities. When I forged my own path, he promoted my endeavors. When I discovered something unexpected, he believed it and explored the phenomenon with me. His ability to immediately and directly attack any problem demonstrates a penetrating knowledge of accelerator physics, though he successfully balances this with a warm sense of modesty, humor, and congeniality. I am indebted to him, not only for extending every opportunity, but also for offering his most sincere friendship.

Richard Temkin at MIT first opened my eyes to the field of accelerator physics. I feel that my current and future achievements stem directly from his guidance. He gave me opportunities that an undergraduate could only dream about, and since then, he has never ceased in advising and encouraging me in this field. His influence on me has been

more than he realizes. Additionally, Claudio Pellegrini at UCLA accepted me to work at PBPL. I am indebted to him for his sagelike wisdom and trust; without him, my development in this field would have been much briefer.

I cannot possibly list everyone that has assisted me scientifically or personally; the list would be far too long. In particular, the research presented in this dissertation required the dedication and diligence of several key people. Xiao-Long Zhang and Nikolay Solyak were fully devoted to the development of the Tevatron Electron Lens over different spans of time. Dan Wolff and Bob Hively covered many of the details necessary for the project's success. Howie Pfeffer and especially Greg Saewert deserve honorable mention; they invested much more time than necessary in order to turn a marginally operational machine into a superb one.

Upon arrival at Fermilab, I spent my time working with Helen Edwards at the AØ photoinjector facility. Though my time there was short-lived, I enjoyed every moment working with her and Don Edwards. Additionally, I attained a number of strong friends there, starting with Nick Barov and Jamie Santucci, and extending to a number of physicists, engineers, and students.

No life is complete without some social outlets, and I am very appreciative of the people which made my life at Fermilab an absolute pleasure. Especially noteworthy are two groups: the pick-up ultimate league, who can catch my long hucks, and my volleyball team, who can hit my quick sets. Matt Worcester, Dave LaFontese, and the rest of my fellow physics students at UCLA will always be cherished. And I thank Linda Coney for standing by me for these last years, helping me survive this strenuous and stressful process.

Lastly but most importantly, my parents were with me in the beginning, middle, and end, and their belief in me has never wavered. I honor and cherish their confidence, patience, encouragement, and unconditional love.

Vita

- | | |
|------|--|
| 1973 | Born Sacramento, CA |
| 1996 | Massachusetts Institute of Technology
Bachelor of Science, Physics |
| 1996 | University of California, Los Angeles
Masters of Science, Physics |
| 1996 | University of California, Los Angeles
Doctor of Philosophy, Physics |

ABSTRACT OF THE DISSERTATION

Successful Beam-beam Tuneshift Compensation

by

Kip Aaron Bishofberger

Doctor of Philosophy in Physics

University of California, Los Angeles, 2005

Professor James Rosenzweig, Chair

The performance of synchrotron colliders has been limited by the beam-beam limit, a maximum tuneshift that colliding bunches could sustain. Due to bunch-to-bunch tune variation and intra-bunch tune spread, larger tuneshifts produce severe emittance growth. Breaking through this constraint has been viewed as impossible for several decades.

A device known as the Tevatron Electron Lens (TEL) has been designed, constructed, and tested in the Tevatron at Fermilab. This device pulses an electron beam that travels for two meters with the antiproton bunches in the Tevatron. The peak beam current is typically 2 A, and the beam energy can range from 4 keV to 12 keV. The bunches interact with the beam's electromagnetic field, causing their tunes to shift opposite to the beam-beam limit. By tailoring the electron-beam current for each bunch, the tuneshift can be individualized to compensate for the adverse bunch-to-bunch variation. Additionally, shaping the electron-beam profile shifts the tune of each antiproton within each bunch selectively depending on their amplitude, compensating for the intra-bunch tune spread.

The maximum antiproton tuneshift is 0.0097 per interaction point. The bunch-to-bunch tune variation is approximately 0.007 and the intro-bunch tune spread is about 0.004. Experiments with the T_{EL} yielded tuneshifts with proton bunches equal to 0.0089, which is equivalent to antiproton tuneshifts of 0.0112. This value is more than necessary to successfully compensate for the beam-beam interactions. Additionally, the use of the T_{EL} has reduced antiproton losses and emittance growth, evidence that such compensation is indeed beneficial to synchrotron operation.

This dissertation introduces the physics of ultra-relativistic synchrotrons and low-energy electron beams, with emphasis placed on the limits of the Tevatron and the needs of a tuneshift-compensation device. A detailed analysis of the T_{EL} is given, comparing theoretical models to experimental data whenever possible. Finally, results of Tevatron operations with inclusion of the T_{EL} are presented and analyzed. It is shown that the T_{EL} provides a way to shatter the previously inescapable beam-beam limit.

Chapter 1:

Introduction

Effects of electromagnetic interaction between colliding beams seem to place rather principle restrictions on the achievable luminosity.

–V. Auslander et al. [13]

High-energy physics has been exploring our understanding of the physical world in the most extreme conditions, ultra-high densities and ultra-high temperatures, in order to explore the origins of our universe and the structure of matter beyond the standard model of quarks and neutrinos. The Tevatron at Fermilab, the LHC at CERN (currently under construction), and the focus on a possible Linear Collider is testament that high-energy colliders are by far the most useful tool for experimental particle physics. In addition, accelerators have a variety of other applications (light sources, FEL's, medical tools, solid-state and plasma physics experiments) as well as being a vital driver for technology development.

However, the most immediate goal for all of the recent and largest accelerators in the world is to provide a large number of high-energy collisions of two opposing beams and measure the decay products of each collision. The *luminosity* \mathcal{L} , defined as the number of these collisions per unit time, is the overall driving force behind Fermilab's Tevatron and this thesis project.

The accelerator projects at Fermilab consistently attempt to increase the production of luminosity, and several of these projects are currently being developed. These projects enhance certain parameters or attack certain limitations. This dissertation covers the recent construction and successful testing of the Tevatron Electron Lens, or TEL, an electrostatic focusing machine that compensates the *beam-beam tunes shift*. The elimination of the barrier set by this phenomenon is expected to boost luminosity production in the Tevatron and future colliders.

1.1. Fermilab overview

The Tevatron was first commissioned at its full energy of nearly 1 TeV (or 1×10^{12} eV) in 1986[1], but the first collection of collision data (referred to as Run I) lasted from 1993 to 1996. On either side of these dates, Fermilab performed various fixed-target experiments, enhanced the Tevatron's performance, and finished many other construction projects. From the spring of 2001 through the current printing of this dissertation, the Tevatron has been operating in what is known as Run II, which is expected to last several more years, until after the LHC at CERN is commissioned around 2009.

<i>Run II parameter</i>	<i>value</i>	<i>units</i>
collision energy	978.5	GeV
peak luminosity	92	$10^{30} \text{ cm}^{-2} \text{ sec}^{-1}$
integrated luminosity	14	$\text{pb}^{-1}/\text{week}$
luminosity lifetime	8	hours
store hours per week	100	hrs/week
store duration	28	hr
number of bunches	36	
protons per bunch	240	10^9
antiprotons per bunch	30–40	10^9
typical horizontal tune ν_x	20.585	
typical vertical tune ν_y	20.575	
beta at collision point (β^*)	35	cm
hourglass form factor	0.65	
overall antiproton efficiency	77	%
avg. antiproton stacking rate	9	$10^{10}/\text{hr}$
stack size at injection	140	10^{10}
injection energy	150	GeV
acceleration time	84	sec

Table 1.1. List of some typical parameters associated with Run II performance in the Tevatron. Average values over recent weeks are provided. Most of these are described in this dissertation; the others can be found in various references. Some of these numbers are likely to evolve as Run II progresses[2,3].

1.1.1. Basic Tevatron operations

In order to maximize the rate of proton-antiproton collisions, the Tevatron is loaded with as many protons and antiprotons as is manageable, and they are squeezed into a small area at the collision points (the proton-bunch rms radius is typically $30\text{ }\mu\text{m}$ at the collision points compared to generally over 0.5 mm elsewhere). Table 1.1 lists a number of parameters associated with a typical *store*, the term given to the period of time (usually twenty to thirty hours) that a specific set of particles remain revolving around the Tevatron. Most of these terms are defined in Chapter 2, and while some of them are fixed by the physical construction of the Tevatron, others can be adjusted for optimal performance.

The protons and antiprotons in the Tevatron are required to travel in small bunches. Each store operates with thirty-six bunches of each species which counter-rotate in one beam pipe. Most proton bunches contain approximately 25×10^{10} individual protons; decent antiproton bunches typically have 2 to 8×10^{10} particles. Each bunch is traveling with nearly 1 TeV of kinetic energy, and therefore they are moving within 0.00005% (5×10^{-7}) of the speed of light. Strong magnetic fields (over four tesla) are needed to keep bending these ultra-relativistic bunches around the Tevatron ring and to keep the bunches from spreading transversely.

The thirty-six bunches of either species are not evenly spaced around the ring. Three bunch trains are established with gaps between them, known as *abort gaps*. As the next section describes, these gaps are necessary in order to load the bunches and to dump the store after it is no longer fruitful. In Figure 1.1, the trains of both species are drawn in a highly stylized fashion. Since the two species are revolving in opposite directions,




Figure 1.1. Simplified cartoon of bunch spacing in the Tevatron. All of the bunches cohabit the one beam pipe, and the antiproton orbit and the proton orbit twist around each other. At CDF and DØ, the double-helix orbits are designed to collapse onto each other, which allow the protons and antiprotons to collide head-on.

Figure 1.1 only indicates one particular instant during a revolution. The bunches labeled P1 and A1 are situated at the same place in the ring. Meanwhile, P25 and A13 are colliding at CDF, one of the two colliders, and P13 and A25 are colliding at DØ, the other collider. As the two species continue to rotate according to the arrows, another pair of bunches will collide at each collider.

Figure 1.1 is sketched for easy comprehension, but in reality the two species cohabit the same beam pipe. During collisions, the orbits twist around each other, like a long circular DNA strand, for the majority of the ring. Only at the two collision points CDF and DØ do the orbits intersect, allowing *head-on* collisions between pairs of bunches.

Physicists commonly refer to the action of two bunches passing through each other as a *collision*. Even though there are very many particles in each bunch, the cross-section of

each particle is so small that typically only a few particles will physically (inelastically) collide. For each collision, there are an average of about two such *events*. The other particles continue through the *interaction point* and are able to orbit around the Tevatron and participate in another collision. As the revolution period of the Tevatron is $21\ \mu\text{sec}$ and there are two interaction points, each bunch participates in 95,000 head-on collisions each second.

The different particles in a bunch have slightly different trajectories, quantified as the *emittance*, which causes the bunch to spread out over a short time. Magnets known as *quadrupoles* continually focus the bunches transversely (horizontally and vertically), and large copper RF cavities prevent the bunches from spreading longitudinally. The magnets are discussed in detail in Chapter 2; the RF cavities are fed 53.1 MHz RF, which induces an oscillating longitudinal electric field along the path of the bunches. This frequency is exactly 1113 times the revolution frequency (47.713 kHz), establishing 1113 *buckets* where an antiproton or proton is able to circulate with stability.

Most of the magnets in the Tevatron use superconducting wires, which are cryogenically cooled by liquid helium at a temperature of 4.2 K. If the bunches are not adequately contained, a large number of particles can crash into the beam pipe in one area, depositing enough heat into the magnet wire, causing the magnet to quench. Quenches stress the affected magnets (and many people), cause the immediate loss of the entire store, and require hours for recovery. Careful operation, plenty of diagnostics, and numerous automated feedback systems all maintain the Tevatron and limit the frequency of these setbacks.




Figure 1.2. Cute cartoon of the accelerators involved at Fermilab. The Main Injector is used both to slam protons into the target for antiproton production and to inject proton and antiproton bunches into the Tevatron. Also shown are several fixed-target lines not being used during Run II.

1.1.2. Getting there

Protons and antiprotons start their lives as negatively charged hydrogen molecules. These are launched from a -750 -kV source and accelerated down a linear accelerator to 400 MeV, illustrated in Figure 1.2. The bunches pass through a foil, stripping the electrons off of the protons and get accelerated in the Booster ring to 8 GeV. During typical operation, the Booster accelerates protons a couple times a second.

The Main Injector is the next link in the chain for proton bunches. Typically seven bunches are injected at one time; then the 53 -MHz RF system accelerates them to 150 GeV. After this, different cavities are turned on that resonate at 2.5 MHz and 5.0 MHz, creating a large bucket that engulfs all seven bunches. Through a procedure known as *coalescing*, these bunches accumulate in the center of the bucket, when the 53 -

MHz RF cavities are switched back on. This generates high intensity bunches suitable for *injection* into the Tevatron. After all thirty-six bunches are injected, it is time to add the antiprotons.

Antiprotons are difficult to produce — and more difficult to collect into a tightly grouped, low-emittance bunch. While one Tevatron store is colliding, antiprotons are simultaneously being produced on a target, collected by the Debuncher, and cooled in the Accumulator, a long process referred to as *stacking*. When the Tevatron store is terminated, enough antiprotons will be usually ready for injection to start a new, full store. In this mode, the Main Injector again accelerates several bunches (albeit to 120 GeV) but instead sends them to the Antiproton Source target. This target is a high-cross-section metal, often nickel or copper, that sprays out a number of different particles, including 8-GeV antiprotons (at the unfortunately low rate of ten to twenty per million incident protons). These antiprotons travel around the Debuncher ring, where they are gathered, cooled, and finally transferred to the concentric Accumulator ring.

As more antiprotons are created, they are added to the *stack* in the Accumulator. Both the Debuncher and Accumulator employ *stochastic cooling*. This process uses sensitive beam-position monitors (BPMs) to pick up slight changes in the average transverse position of the passing beam; this signal is quickly sent downstream via a laser beam cutting a diameter across the ring and applied to kickers to correct the displacement. Each time this occurs, the average transverse movement is decreased slightly. After a few turns, the average transverse motion (emittance) has not changed much, but over the course of hours of repeated small corrections, the particles in the beam will have significantly less emittance.

After hours of continuous stacking, enough antiprotons are gathered and sufficiently cooled that they can be extracted into the Main Injector. The Accumulator contains an RF-cavity system that can provide four buckets in which to capture antiprotons, but the strength of the fields is kept very small, so that only a few antiprotons have just the right amount of momentum to be captured. These particles are purposefully accelerated slightly and drawn to a larger-radius orbit, while the majority of the antiprotons slide in and out of the buckets with incompatible momenta. Therefore an “outer track” of four bunches is created alongside the remaining continuous (DC) beam. These four bunches are kicked into the Main Injector, accelerated to 150 GeV, and injected into the Tevatron. Repeating this process nine times with increased Accumulator-RF strength generates the thirty-six antiproton bunches in the Tevatron[7].

Recently, another ring known as the Recycler has been commissioned; it can be used to *stash* portions of the antiproton stack. This is helpful, as the stacking rate is significantly larger at small stack sizes. Stashing the first portion allows the stacking process to run more efficiently.

1.2. Tevatron Overview

The Tevatron itself requires many systems beyond magnets and cavities for it to function well; a few of them are used in various contexts in this dissertation. Kickers are the electromagnetic devices that steer a few or all of the bunches at a time when loading or dumping a store. BPMs track the transverse position of the bunches at 216 locations around the ring. Beam loss monitors measure the rate at which particles are being

lost. Schottky detectors measure an important parameter known as the *tune*, defined in Chapter 2.

Many systems depend on precise timing with respect to the circulating bunches. A serial communications bus is distributed through the Tevatron that contains many specific timed signals along with a generalized beam sync. The time jitter on this sync with respect to the actual bunches is less than 1 nsec. The subject of this dissertation, the T_{EL}, uses this sync as a trigger. Diagnostics and time-dependent data expressed in the following chapters feed from this beam sync or derivative signals.

1.2.1. Establishing collisions

When the first four bunches of each antiproton train are injected into the gaps between proton-bunch trains, the Tevatron appears as illustrated in Figure 1.3*a*. Separate RF cavities control the energy (and therefore the speed) of the protons and antiprotons, which allows the antiprotons to precess with respect to the protons, a technique known as *cogging*. After the antiprotons have clogged a desired amount, shown in Figure 1.3*b*, the gap is now clear to add another set of four bunches. In Figure 1.3*c*, these antiproton bunches are added, and all of the antiprotons are again clogged until it reaches Figure 1.3*d*. Another batch of four bunches per train is added (Figure 1.3*e*), and the antiprotons are clogged one final time so that they are synchronized longitudinally to collide exactly at the centers of the CDF and DØ detectors (Figure 1.3*f*). During this time, devices known as *electrostatic separators* keep the paths of the proton and antiproton orbits transversely separated, lowering the space-charge interference between the two species[4].

After all of the bunches are loaded, the Tevatron is ramped from 150 GeV to 980 GeV, the nominal operational energy of the Tevatron (at the absolute maximum energy,

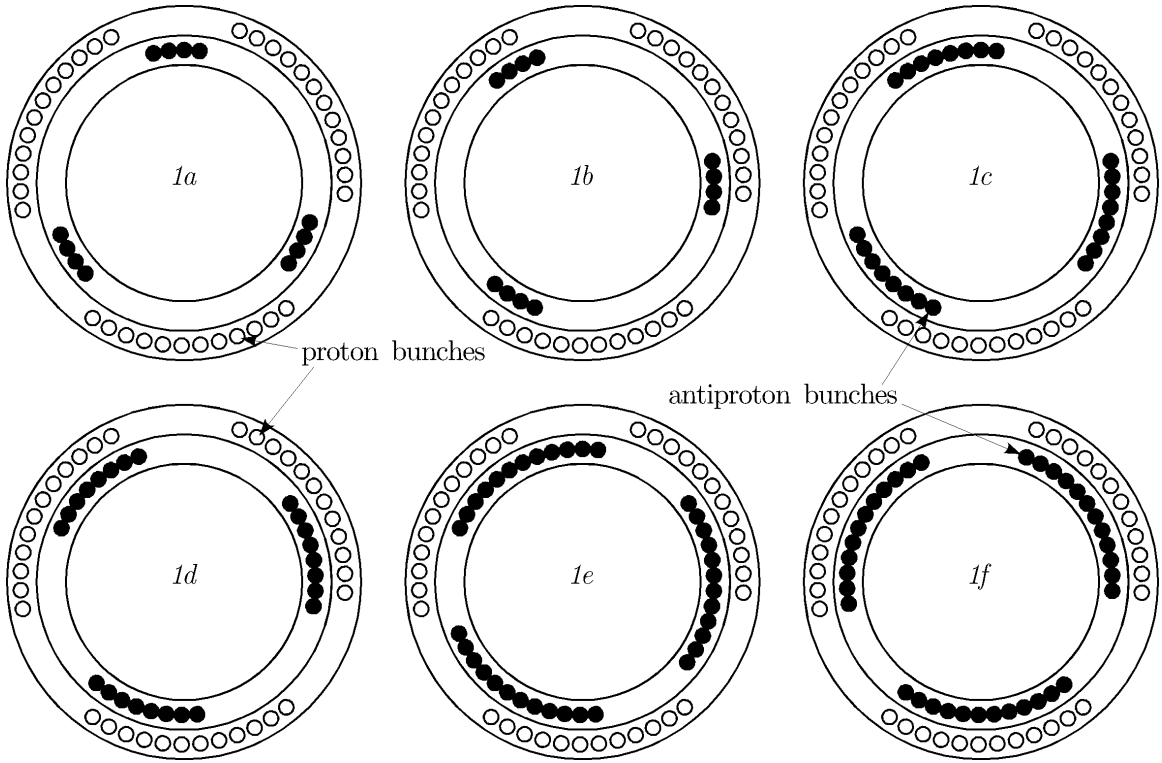


Figure 1.3. Illustration of loading antiprotons into the Tevatron. The hollow circles represent protons, which are loaded first. In *1a*, the first four antiproton bunches (filled circles) in each train have been injected into the abort gap. These bunches are advanced forward (*cogging*) in *1b*, before another batch of four bunches is injected in *1c*. The antiprotons are again advanced (*1d*), the final batch is injected (*1e*), and everything is advanced so that proper bunch can collide at CDF and DØ.

1.050 TeV, the superconducting magnets were more susceptible to quenches[1]). Again, the RF cavities accelerate the particles, while the dipoles, quadrupoles, and other systems are simultaneously ramped. In spite of the enormity of such an order, ramping typically succeeds with only a few percent of the particles being lost.

Flanking either side of the two detectors are *low-beta quadrupoles*, which have stronger fields than the other quadrupoles around the Tevatron. By increasing the




Figure 1.4. Cartoon of proton and antiproton bunches colliding. While this illustration is not to scale, it demonstrates the reality that bunches next to the collision point tend to pass each other quite closely, an issue discussed in Section 2.4.

strength of these magnets, the transverse size of the bunches is decreased at the collision points (the proton bunch radius becomes approximately $30\,\mu\text{m}$ and the antiproton bunch radius becomes about $25\,\mu\text{m}$). During this step, known as a *low-beta squeeze*, the protons and antiprotons are still on disjoint paths; no head-on collisions have yet occurred. The last step is to adjust the electrostatic separators so that the orbits cross at the two collision points only. Figure 1.4 sketches how the trains of bunches appear when collisions are taking place.

1.2.2. Luminosity

The term *luminosity* is often used for several different, but related, notions. In general, the goal of the interaction points is to provide a large and consistent number of collisions per unit time \dot{n} , which can be written as

$$\dot{n} = \sigma \mathcal{L}_i$$

in order to split the quantum mechanical cross-section σ and the accelerator-related instantaneous luminosity \mathcal{L}_i . The cross-section describes the natural structure of the

particles, while the instantaneous luminosity folds in all of the features related to the two colliding bunches. It is the latter that can be affected by accelerator parameters and is the subject of this discussion.

The instantaneous luminosity for the Tevatron is therefore defined to be the number of protons and the number of antiprotons passing through the same area together. If the two bunches collide perfectly head-on, then an arbitrary distribution of particles would be the integral over all space of the density of both species multiplied together[5]:

$$\mathcal{L}_i = 2\beta c \int_{\text{all space}} \rho_p(\vec{x}) \rho_{\bar{p}}(\vec{x}) dV, \quad 1.1$$

where $2\beta c$ represents the speed of the passing particles as seen by the detector, and $\rho_p(\vec{x})$ and $\rho_{\bar{p}}(\vec{x})$ describe the volume particle densities of the protons and antiprotons.

So far, the luminosity does not integrate over the time that the bunches pass through each other. It has been shown that bunches tend to stabilize in a Gaussian distribution[6], so that for each species,

$$\rho(\vec{x}, t) = \frac{N}{(2\pi)^{\frac{3}{2}} \sigma_x \sigma_y \sigma_z} \exp \left[\frac{x^2}{2\sigma_x^2} + \frac{y^2}{2\sigma_y^2} + \frac{(z \pm ct)^2}{2\sigma_z^2} \right]$$

applies to each of the colliding bunches. In this context, N is the total number of particles in the bunch, $\sigma_{x,y,z}$ are the rms bunch radii in three dimensions, and $\pm c$ is the velocity of the counter-propagating bunches. Integrating \mathcal{L}_i over time yields the luminosity from a single bunch crossing:

$$\mathcal{L}_b = \int_{-\infty}^{\infty} \mathcal{L}_i(t) dt = \frac{N_p N_{\bar{p}}}{4\pi \sigma_x \sigma_y}, \quad 1.2$$

and the time-average luminosity is \mathcal{L}_b times the bunch collision frequency.

1.2.3. Detectors

An individual proton or antiproton is constructed from three quarks, some gluons, and some other, more mystical, subatomic particles. A collision between an opposing proton and antiproton is actually the interaction of only one quark or gluon within each of them. Due to the large amount of total energy within each of the quarks and the kinetic energy of the bunches, many types of particles can be produced, most of which are very unstable and quickly decompose into a large number of simpler, more stable particles.

Scores of these various final particles are usually generated from a single collision, and large, immensely complicated detectors surround both of the interaction points, CDF and DØ. In Figure 1.5, a cross-section of each detector is shown. Each detector wraps around the horizontal axis, and is symmetric on either side of the interaction point. The two detectors are glued together in Figure 1.5 to display discernable differences between them; however, the general layout of each detector shows significant similarities.

Both detectors employ several methods to track particles out from the interaction point [8,10]. Resolution of tracks is high enough to see exactly where the final decay particles originated from. When an unstable intermediary, such as the top quark (discovered in Run I) or the Higgs boson (hopefully discovered in Run II), is created, it will travel a distance as short as $100\text{ }\mu\text{m}$ before decaying into a bevy of more stable particles. Meticulous analysis of tracks reveals the small offset and offers clues to the properties of these rare particles.

The fanciest piece of each detector is the silicon vertex detector, which employs almost a million separate photodetector-amplifier circuits crowded microns away from




Figure 1.5. Cross-section views of the CDF and DØ detectors show some of their different components. Each detector is symmetric around the central axis and surrounds the interaction point. In order to compare differences and similarities, only half of each detector is displayed, sandwiched together around a common interaction point.

each other[11,13]. Each photodetector is extremely sensitive to almost any passing particle, and hundreds will typically fire each time a pair of bunches collide. Because of the huge amount of data that is generated by each collision, state-of-the-art logic circuits (*triggers*) discriminate between “interesting” events and not-interesting events. Only the interesting ones, that appear to contain a sought particle (like a top quark or Higgs boson or some other rare particle), are saved to hard disk to be analyzed later.

Several constraints are forced on the bunches because of the detectors. Since the tracks must be measured accurately, the collision point must remain stationary within a few microns. Also, the time of each event is measured to determine speeds and delays, so

the collision time must be known precisely. In general, the Tevatron has a decently stable orbit path and timing system and allows adequate measurements of these rare particles.

1.2.4. Pushing the limits — beam-beam effects

Over the past half century, the two variables maximized in every new collider are the collision energy and the average luminosity. New technology enables accelerator upgrades that can increase those two attributes. The collision energy of the Tevatron is essentially limited; changing it would require an unrealistically huge amount of effort. The luminosity, on the other hand, depends on the adjustable quantities in Equation 1.2.

Independently, each variable can easily be improved. For example, more protons can be poured into each bunch, directly inflating the luminosity. This section illustrates how larger numbers of protons will harm the antiproton beam size and lifetime, in turn lowering the overall performance of the store. Operators are “trained people that vigilantly balance various aspects of the store with the intent of maximizing the average luminosity over the entire store[14].”

An unfortunate consequence of employing circular colliders is the disruption each bunch feels due to the electromagnetic fields generated by other, nearby bunches. It is possible to separate these effects into two categories: the head-on collisions at CDF and DØ, and long-range interactions when two bunches pass by each other, often called *parasitic* crossings. The head-on collisions are stronger than the parasitic ones. However, each bunch passes through only two head-on collisions, whereas it experiences seventy parasitic crossings. Tevatron stores have eight times the number of protons per bunch than antiprotons, so that the fields created by the protons are much larger, and the

antiproton bunches end up being much more afflicted. The forces that an antiproton bunch witnesses is qualitatively sketched here, but Chapter 2 canvasses the interaction more completely.

In the case of the head-on collisions, the antiprotons pass through the proton bunch and feel a transverse focusing force that varies with radial position. Antiprotons that happen to be axially centered in the bunch feel no force, antiprotons a small distance away feel a very strong force, and antiprotons on the outskirts of the bunch feel a weaker force. Plotted in Section 2.4, the nonlinear force is typically split into two portions: a linear part, expressed by $F_0(r) = -k_0 r$, accounts for the particles near the center of the bunch, and a nonlinear part, which asymptotically approaches $F_1(r) = -k_1/r$.

Both the linear and nonlinear aspects of the focusing force make maintaining the antiproton bunch more difficult. The linear force shifts a parameter known as the *tune* of the particles, but it shifts all of the particles the same amount. The nonlinear force shifts the tune different amounts for different particles, creating a tune spread within a given antiproton bunch. This latter consequence is shown in Chapter 2 to be detrimental to the survival of the bunch, but has been nearly impossible to remedy without sacrificing the luminosity of the machine.

Indeed, a several references over the past several decades have claimed that this *beam-beam limit* puts a clamp on future circular-collider performance [15,16,17,18]. John Rees said in 1986: “And finally, [they]...came up against what may fairly be characterized as the Fundamental Limit, the beam-beam limit of the colliding-beam storage ring. This phenomenon has proved to be insurmountable and continues to place the basic limit on the performance of colliding-beam storage rings. The theorem deserves to

be capitalized[20].” In 2004, Lyn Evans reaffirmed the finality of this constraint: “...the beam-beam interaction will always be a Fundamental Limit[21].”

The parasitic crossings are significantly different, since the antiproton bunch passes by a proton bunch at a significant distance away, usually eight to twelve times the rms bunch size. Unfortunately, the fields of the proton bunch, following the $-k_1/r$ -dependence, are still significant at that range, and the antiprotons feel a force at each of the seventy crossings. Mostly, the antiprotons are equally attracted to the charge of the protons, generating a small deflection of the entire bunch that can be compensated by various elements in the Tevatron. But the position-dependent force again shifts the tunes of the antiprotons, again causing havoc.

The orbits in the Tevatron have numerous other constraints placed on them, causing the transverse separation between proton and antiproton bunches to vary from one crossing to another. Figure 1.1 displays each bunch train with an abort gap between each one. Because of these gaps, the places where protons and antiprotons cross is different for each of the antiproton bunches. Hence, each bunch feels a unique net focusing effect and tunes shift.

Most of the bunches are tuneshifted nearly the same amount, since the transverse separations do not vary excessively around the ring. However, in order to draw the bunches together at both of the head-on collision points, the parasitic crossings on either side of each IP has a much smaller separation of 4.5–5 times the rms bunch size. The majority of the antiproton bunches are significantly affected by passing a proton bunch at each of these locations.

The first and last antiproton bunches in each train do not witness a proton bunch at one of these close-range crossings, so their tunes are not shifted like those of the “internal” bunches. As Chapter 2 elucidates, it is this difference in tuneshifts between bunches that proves problematic.

The intra-bunch tune spread and the bunch-to-bunch tune differences have different causes and attributes, but the final effect of both is that the antiprotons have a range of average tunes. This total range, known as the *tune footprint*, can be moved around through the use of various magnets. However, when the tune of any particle falls near one of many *resonances*, the particle can increase its transverse motion. This behavior expands the beam size $\sigma_{x,y}$ and leads to particle losses. Both of these lower the luminosity, and excessive particle losses can cause a quench of the Tevatron. The tune footprint and resonances are also detailed in Chapter 2.

John Rees and Lyn Evans perceived the beam-beam limit very fatalistically due to the fact that it is nearly impossible to correct either of these tune-expanding causes: the area of the tune footprint seemingly cannot be decreased without giving up luminosity. In order to repair the intra-bunch tune spread, a strong nonlinear radial electric field would need to be created. Gauss’ law mandates that there needs to be some sort of charged object along the axis to generate this[35]. Any metal wire placed in the path of the antiprotons would intercept particles and destroy the very bunch that it was trying to remedy. A strong azimuthal magnetic field could also generate radial forces, but due to Faraday’s law, a current needs to flow along the axis, creating the same problem.

The bunch-to-bunch tune variation presents its own complications. Changing the tune of a bunch requires a strong electric or magnetic field; the superconducting *quadrupole* magnets that establish the tune of the Tevatron put out several Tesla fields. In

order to change the tune of just one particle, these magnets would need to be pulsed at very high frequencies, but the natural inductance of the coils resists change in current. The current sources feeding the magnets would need unrealistically high levels of power to effectively change the fields, and all of this power translates into heat, which would immediately quench the superconducting wire.

Pulsing electrodes with voltages is easier than pulsing magnets, but the force an antiproton feels from 1 T of magnetic field would require an electric field of 300 MV/m [6]. Such huge voltages are impractical to generate or pulse at high rates, and the electrodes would be prone to spark [22]. It is due to these difficulties that the tune footprint has been perceived as immutable.

1.3. Tevatron Electron Lens overview

The Tevatron Electron Lens, or TEL, is a prototype machine that breaks through the beam-beam limit imposed on circular colliders. By introducing an electron beam into the midst of the antiproton bunches, a radial electric field is created that shifts the tune of the particles. The electron beam current can be easily modulated, allowing the tuneshift to change for each bunch. Proper shaping of the transverse electron density can change the tuneshift of individual particles in the bunch, allowing a compression of the intra-beam tune spread.

Electron beams are easy to generate, and the kinetic energy of the electrons can be chosen by convenience. Since electrons have the opposite charge of protons, the passing antiprotons are affected in the opposite direction as they were by the protons. Therefore,

by mimicking the shape and strength of the protons, the electron beam can “undo” the ailments antiprotons feel due to the protons. The goal of the TEL is to compensate for the beam-beam interactions, namely *tuneshift compensation*.

Historically, the concept of compensating for the beam-beam interactions has been proposed several times. In 1969, an idea was proposed to build an electron-positron machine using four simultaneously colliding bunches, two of each species[23]. The machine, known as DCI in Orsay, France, succeeded in producing the complicated arrangement and observed a 30% increase of the effective space-charge forces[24]. However, no significant increase in the luminosity was demonstrated, which was blamed on high-order collective instabilities[25].

Regardless, any machine employing four distinct bunches is unappealing, and Figure 1.5 indicates that the space around the collision points is already extremely crowded. Section 3.2 justifies that tuneshift compensation does not need to take place at the interaction point; it can be implemented anywhere around the ring (albeit certain regions prove better than others). An unpublished document in 1976 floated an idea of using an electron beam to compensate for the beam-beam tuneshift at CESR at Cornell University[26]. However, the electron beam crosses the particle orbit, limiting the amount of time that the bunches sense the fields. In order to impart a significant tuneshift, the electron beam required extremely high current densities[27].

A scheme for beam-beam compensation was also proposed for the SSC, with an appearance very close to that of the TEL[28]. Since the SSC was intended for proton-proton collisions, an electron beam could be used to compensate the tuneshift in both of the proton beams. The brief proposal imagined that the electron-proton interaction

region would act like another collision point. The electron beam size was identical to that of the protons, and an imaging and feedback system would constantly correct the transverse position. Concerns about the correction system, possible instabilities, and current densities would have needed to be resolved.

In 1997, Vladimir Shiltsev and David Finley proposed the device known as the T_{EL} for beam-beam compensation in the Tevatron[30]. While the design evolved slightly over the following years[31], it is the first proposal to combat the beam-beam limit while remaining straightforward, realizeable, and unobtrusive. Also, the fact that it is relatively inexpensive allowed it to progress as an “experiment,” without interfering with the general Tevatron program.

In its simplest form, the T_{EL} consists of an electron gun, collector, and a set of solenoid magnets that guide the electron beam from the gun, into the Tevatron beam pipe, along the path of the antiprotons for two meters, and out to the collector. The electron beam diameter is significantly larger than that of the antiprotons, so that perfect alignment schemes are not necessary. The kinetic energy of the electron beam is nominally 10 keV, and the gun puts out a couple amperes. The whole device only spans 2.5 m of the Tevatron circumference, and can be positioned wherever there might be an unoccupied space. Chapter 3 will show that the T_{EL} is more effective when certain Tevatron parameters meet optimal conditions, so only a few places around the ring are ideal.

Control of the electron gun is crucial to successful compensation. Being a thermionic gun, the cathode is set to a negative voltage, and the anode is pulsed to different voltages. The current drawn off of the cathode is proportional to $V_{\text{anode}}^{3/2}$, where V_{anode} is the voltage

of the anode with respect to the cathode. By adjusting this voltage between the passage of each antiproton bunch, the current and radial forces is adjusted for each bunch. This allows each bunch to be tune-shifted independently, allowing the bunch-to-bunch variation to be minimized. Additionally, another electrode in the gun can change the radial charge distribution in the electron beam simply by adjusting its voltage. This electrode, affectionately called the *profiler*, allows the central and outlying antiprotons to witness different focusing forces. Thus, the nonlinear tune spread is able to be diminished.

Of course, the previous discussion ignores many details, each of which are addressed in subsequent chapters. After the design of the T_{FL} was finalized, the pieces were fabricated and assembled in a testing area. Measurements of different properties and alterations were performed, and then it was installed in the Tevatron. Dedicated studies produced successful tuneshifts in a short amount of time, and after some more perseverance, compensation of the tuneshift improved the preservation of antiproton bunches.

1.4. This thesis project

This thesis describes the T_{FL}, a novel device designed to compensate for the adverse beam-beam interactions. In general, the tuneshifts associated with colliding beams had been believed to erect a limitation on the performance of circular colliders. The goal of the T_{FL} is to remove this barrier for future machines, and this dissertation evinces that initial experiments with the T_{FL} are successful.

More specifically, the T_{FL} is designed to produce an electron beam that interacts with passing antiproton bunches in order to change the tune of that bunch. This interaction is intended to compensate for the tuneshift that a given antiproton bunch feels

due to its collisions with proton bunches. Currently, the T_{EL} has succeeded in changing the tune of an entire bunch, known as *linear* compensation, with the intention of compensating for the bunch-to-bunch tune spread.

Additionally, the T_{EL} is able to produce an electron beam that generates nonlinear radial forces. With a shape similar to that of a proton bunch, the T_{EL} is able to vary the tuneshift for antiprotons within each specific bunch. By achieving this, the spread of tunes within each of the bunches is decreased, preserving the emittance of the antiproton bunch. Results are presented that indicate that the T_{EL} is also successful in this goal.

Neither of these results have ever been witnessed before the T_{EL} was created. Such a device paves a road to future colliders that are not hampered by the beam-beam limit. The conclusions of this dissertation find no significant hurdles in adapting the T_{EL}, or other T_{EL}-like devices, to future accelerators or using them to significantly increase the utility of any previous collider design.

In this dissertation, a brief introduction to general accelerator physics is given in Chapter 2, emphasizing aspects relevant to the Tevatron and the T_{EL}, notably beam-beam effects. Chapter 3 quantifies specific requirements that the T_{EL} must satisfy if it is to achieve the goals of successful beam-beam compensation. A theoretical analysis of electron-beam generation and propagation is given in Chapter 4, which concludes that the T_{EL} can provide the necessary quality of beam. A more physical description of the T_{EL}, including measurements of many beam properties, is provided in Chapter 5. The final two chapters address the operation of the T_{EL} on the Tevatron performance. In particular, Chapter 6 tests several complex interactions between the electron beam and proton/antiproton bunches. Lastly, verification of tuneshift and lifetime improvement is

displayed in Chapter 7, along with comparisons to theory, interpretations of the data, and predictions of future TEL research.

Chapter 2:

Particle motion in a circular accelerator

The subject, however various and important, has already been so frequently, so ably, and so successfully discussed, that it is now grown familiar to the reader, and difficult to the writer.
—Gibbon: *The decline and fall of the Roman empire*

An understanding of basic accelerator physics is necessary to follow the motivation of the TEL and its effects on the Tevatron bunches. This chapter develops accelerator physics from fundamentals; however, an introductory, grandiose portrayal of general accelerator theory is avoided for the sake of brevity. This dissertation does not do justice to the field that numerous books already have [6,32,33]. Instead, this chapter focuses on the background behind subjects that are relevant in subsequent chapters.

Much of the following discussion applies to numerous accelerators, but attention will be devoted only to the Tevatron. Chapter 1 briefly described the rest of the accelerator

chain at Fermilab, and other accelerator facilities have different operating characteristics. Additionally, the T_{EL}, which uses a low-energy electron beam immersed in a solenoidal magnetic field, is not discussed here. Chapter 3 devotes its effort toward understanding that type of beam.

First, this chapter describes how a beam propagates through a circular accelerator, using strong focusing. A development of tunes and resonances follows. An analysis of the beam-beam interactions is last, which illustrates the performance issues that the T_{EL} is designed to solve.

2.1. Matrix formulation of linear dynamics

The Tevatron consists of many systems and components. Thousands of electrical, magnetic, and mechanical devices, each powered by hardware and software, are vital to the sustainability of each store. Table 2.1 lists a selection of parameters involved in Tevatron operation. Some values are immutable since the construction of the Tevatron (the circumference, for example), while others are typical amounts that are changed often between or during stores.

In its simplest form, the Tevatron ring is simply a vacuum pipe bent into a large circle. Large, superconducting dipole magnets surround the pipe, which bend the path of the charged particles. The dipoles are 6.1 m long and bend the particles only a small amount, about 8.12 mrad maximum. Within the magnet, the radius of curvature is determined from $F = \gamma m a$ [35]:

$$p = eB\rho,$$

<i>Tevatron parameter</i>	<i>value</i>	<i>units</i>
Tevatron circumference	2π	km
number of dipoles	774	
dipole magnet length	6.12	m
dipole field strength	4.3	T
number of quadrupoles	216	
quadrupole magnet length	1.678	m
quadrupole focal length	26.1	m
cell type	FODO	
phase advance per cell	68	deg
standard half-cell length	29.7	m
bend magnets per cell	8	
β_{\max} in cells	~ 110	m
maximum dispersion D_x	6.2	m
revolution period	20.96	μsec
RF frequency	53.1	MHz
harmonic number	1113	
transition gamma γ_{tr}	18.7	

Table 2.1. List of Tevatron-related quantities, many of which are essentially fixed by the hardware installation in 1986. Some values are adjustable, but rarely are these significantly altered. This chapter defines most of these, but some are discussed in other chapters or references.

<i>operational parameter</i>	<i>value</i>	<i>units</i>
proton emittance (95 % norm.)	18	π mm-mrad
antiproton emittance (95 % norm.)	12	π mm-mrad
antiproton beam-beam tuneshift per IP	9.7	10^{-3}
proton beam-beam tuneshift per IP	3.8	10^{-3}
number of parasitic crossings	70	
transverse coupling	4	10^{-3}
bunch spacing	396	nsec
bunch length	37	cm
longitudinal emittance	3	eV-sec
momentum spread ($\Delta p/p$)	9	10^{-5}
synchrotron tune	2.7	10^{-4}

Table 2.2. List of antiproton and proton parameters typical during recent stores. Some of these parameters are adjusted regularly by operations [1,3,34], and others are being improved as Run II proceeds.

where p is the particle's momentum in MKS units, B is the magnetic field, and ρ is the radius of curvature (literature often uses the more convenient unit eV/c for momentum, so that $p = cB\rho$). As Chapter 1 explains, the antiprotons and protons in the Tevatron are traveling in opposite directions with the same momentum p . Therefore, both species rotate at the same radius and cohabit the same beam pipe.

The dipole magnets generate a field of 4.3 T during collisions. A particle passing through one magnet of length $L_{\text{dipole}} = 6.12\text{ m}$ bends an angle of

$$\theta \approx \frac{L_{\text{dipole}}}{\rho} = \frac{eBL_{\text{dipole}}}{p}, \quad 2.1$$




Figure 2.1. The reference frame used throughout this dissertation is attached to the orbit of the moving bunches. The discussion will refer to dimensions in the (inertial) laboratory frame, where the phenomena described are static situations.

using the small-angle approximation, so that θ is merely 8.12 mrad for the Tevatron. The sum of all of the dipoles must complete a full revolution, so

$$\sum_{i=1}^N \frac{eBL_{\text{dipole}}}{p} = 2\pi,$$

where $N = 774$ is the total number of dipoles around the ring.

The fields of the dipoles are quite uniform throughout a large region of the beam pipe. Particles that are not exactly on the closed orbit are still follow Equation 2.1. In order to examine the dynamics of these particles, it is convenient to ignore the common effects of the dipoles. A frame of reference is introduced that follows the beam around the ring. The origin is on the closed orbit, the longitudinal axis \hat{z} points tangentially along the orbit, and the two transverse axes \hat{x} and \hat{y} point horizontally and vertically respectively. Figure 2.1 sketches the frame of reference centered on a bunch and follows it around the ring.

By tracking a “perfect” particle as it traces out the closed orbit, describing the dynamics of all of the real particles around it becomes much easier. First, the relativistic

forward motion is removed from any calculation (though the effective mass γm must be retained). Second, the dipole magnets can be effectively ignored when quantifying the behavior of a bunch as it passes through these regions. This feature is justified in the next section and embraced in the rest of this chapter.

The tremendous forward momentum of the particles in the Tevatron is much larger than their transverse momentum. Transverse dynamics therefore almost never affect longitudinal dynamics, leading to uncoupled motion and a much simpler analysis. Often called the paraxial approximation, it also can be proclaimed as limiting the angle between particle trajectories and the longitudinal axis \hat{z} to very small values. Additionally, this accurate assumption allows the approximation $\tan \theta \approx \sin \theta \approx \theta$.

As the next section describes, the components on the Tevatron beamline affect particle motion in the horizontal and vertical motion largely independent of the other axis. This decoupled motion is another convenient and reliable simplification involved in understanding the dynamics. In general, a particle's motion is described fully by a six-dimensional vector $(x, x', y, y', z, \delta)$, where δ is defined as the particle's momentum over that of the ideal particle, $\delta p/p_0$. The z coordinate refers to the particle's position again with respect to the reference frame, and is equal to $-\beta c \Delta t$. A particle in front of the centroid of the bunch is said to be at a positive z -coordinate and arrives at a location in the Tevatron earlier than the centroid. If the centroid reaches a particular location at a time t_0 , the early bunch arrives at a negative time Δt .

Since the motion in each direction is largely decoupled, attention will be put on only one pair of factors, such as x and x' . The vertical dimension follows similar treatment to the horizontal, while the longitudinal follows distinctly different mechanics.

2.2 Building a machine lattice

When particles pass through most magnetic structures, they are kicked transversely with respect to the reference frame depending on their position. This process of kicking particles and then allowing them to drift is analogous to light passing through optical lenses. Each element, including drifts, can be described by a matrix that acts on the state vector of each particle. Therefore, passing a particle through a large system of elements can be reduced to the multiplication of a series of matrices.

In this manner, a whole ring can be thought as a series of simple optical elements that can be amalgamated into one transfer matrix. To a first-order approximation, most optical elements act in each dimension independently, and it is convenient to mathematically analyze different elements under this approximation. The current discussion tracks a particle in one transverse dimension, where the transfer functions describing different optical elements are simply 2×2 matrices. In Section 2.3, features such as coupling and momentum spread are shown to link motion in different dimensions.

2.2.1 Drift sections

The simplest element to analyze is a simple drift section, where a particle is allowed to travel without interference. In the top portion of Figure 2.2, an uninterrupted particle starts in a state $\vec{x}_0 = (x_0, x'_0)$ and moves along the length L_{drift} so that its final position is simply $\vec{x}_1 = (x_0 + Lx'_0, x'_0)$. This can be rewritten in matrix notation as:

$$\vec{x}_1 = \begin{pmatrix} 1 & L_{\text{drift}} \\ 0 & 1 \end{pmatrix} \vec{x}_0. \tag{2.2}$$




Figure 2.2. Two examples of a particle propagating through optical elements. The top illustration is a drift section of length L_{drift} . The bottom shows a focusing quadrupole bending the particle toward the axis. In both cases, (x_0, x'_0) represents the particle's initial trajectory and (x_1, x'_1) represents its final.

One of the conveniences of using a frame of reference which tracks the closed orbit is that it bends around dipole magnets the same amount as the particles do. Because everything bends in a nearly identical manner, the motion of a particle through the dipole with respect to the closed orbit is almost same as if they were traveling simply through a straight drift section. This means that each long dipole length can be approximated as a drift. In reality particles that travel at an angle through a given dipole sense a slightly longer length of field, and thus are bent more than particles that are traveling straight through. Additionally, if the ends of the dipoles are not perfectly parallel (for example, a *sector* magnet), particles passing through at different transverse positions will be affected differently. In general, these effects are known as *path-length focusing*, and contribute a

small correction to the drift-length assumption. The next section describes much stronger forms of focusing, so that the dipole contribution is either ignored or added to the larger contributions.

2.2.2 Quadrupole fields

The description of dipole magnetic fields is easy to write down, as the field strength is a constant throughout the region of interest and pointed vertically in order to bend the charged particles horizontally. For more sophisticated fields, it is helpful to start with the general formula for magnetic fields.

Any two-dimensional field must satisfy Laplace's equation $\nabla^2\Psi = 0$, and the solutions that converge at $r = 0$ can be written by the well-known multipole expansion [36],

$$\Psi = a_0 + \sum_{n=1}^{\infty} (a_n r^n \cos n\phi + b_n r^n \sin n\phi) , \quad 2.3$$

where a_n and b_n are constants determined by the specific geometry in question. A dipole directed vertically is simply Equation 2.3 when b_1 is the only non-vanishing term. The magnetic field becomes

$$\vec{B}_{\text{dipole}} = -\nabla\Psi_{\text{dipole}} = -\nabla(b_1 r \sin \phi) = -\frac{d}{dy}(b_1 y)\hat{y} = -b_1\hat{y} .$$

Since the magnetic field does not depend on radius, any particle passing through the dipole will be bent identically.

A more complicated magnet generates a position-dependent field. For example, if b_2 is the only nonzero constant, the field is

$$\begin{aligned}
\vec{B}_{\text{quad}} &= -\nabla \Psi_{\text{quad}} \\
&= -\nabla(b_2 r^2 \sin 2\phi) \\
&= -\frac{d}{dx}(2b_2 xy)\hat{x} - \frac{d}{dy}(2b_2 xy)\hat{y} \\
&= -2b_2 y\hat{x} - 2b_2 x\hat{y}.
\end{aligned}$$

This field produces a force on ultrarelativistic particles given by:

$$\vec{F} = qv\hat{z} \times \vec{B}_{\text{quad}} = (2b_2 ce)x\hat{x} - (2b_2 ce)y\hat{y}, \quad 2.4$$

which illustrates the separation of the two transverse dynamics: a particle with a specific offset x_0 in the horizontal directions is accelerated by $(2b_2 ce)x_0/\gamma m_p$ kick horizontally. Integrating this acceleration over the length of the quadrupole L_{quad} , the horizontal trajectory is altered by

$$\Delta x' = \frac{\Delta p_x}{p_z} = \frac{\gamma m_p \int a_x dt}{p_z} = \frac{2b_2 e L_{\text{quad}}}{p_z} x_0. \quad 2.5$$

A drawing of an example particle passing through a quadrupole is shown in the bottom half of Figure 2.2.

If the initial path is parallel to the axis, a quadrupole with negative b_2 will pull the particle toward the axis. A focal length, identical to that in light optics, can be introduced to described the strength of a quadrupole. The transfer matrix can be written[37,6]:

$$\vec{x}_1 = \begin{pmatrix} 1 & 0 \\ -\frac{1}{f} & 1 \end{pmatrix} \vec{x}_0, \quad 2.6$$

where the focal length f is equivalent to $-x_0/\Delta x'$, and an effective spring constant k , used in the following section, is

$$k = \frac{1}{cp_z} \frac{dF_x}{dx} = \frac{2b_2e}{p_z}.$$

Because the focal length of the quadrupole is commonly much larger than the length of the physical magnet, the thin-lens approximation can be employed[37]. This simplification is suggested by the dashed lines in Figure 2.2.

The linear field strength sets the focal length constant for particles at any incident position and transverse momentum. If b_2 is negative, the particle is focused in the horizontal direction. But the same b_2 in the vertical plane generates a force away from the axis, making the particle defocus. By reversing the current powering the magnet, b_2 will become positive, and the magnet will defocus in the horizontal plane and focus in the vertical. Accelerator terminology has developed the convention of referring to quadrupoles by their action in the horizontal dimension. Another convention, adopted in this dissertation but less consistently in other literature[38], uses a positive value for the focal length to imply focusing and a negative value to connote defocusing.

It is useful to see if any magnetic device is able to provide simultaneous horizontal and vertical focusing. To do this, a loop can be traced around the central beam axis, where the force \vec{F} acting on particles is pointing inward at every point. Since $\vec{F} = qv\vec{z} \times \vec{B}_{\text{quad}}$, B_ϕ must be positive around the loop. Integrating and using Ampère's law, $\int \vec{B} \cdot d\vec{l} = \mu_0 I_{\text{inside}}$ indicates that this device would need to generate a longitudinal current somehow immersed in the bunch path.

Three notable exceptions can be mentioned. First, a magnetic solenoid can introduce angular momentum in the bunch. This, in turn, produces a $\vec{v}_\phi \times \vec{B}_z$ force inward. A

second-order effect, the force is very weak on high-energy bunches and is analyzed more in Section 3.2. Second, plasma focusing produces focusing in both transverse directions, but it is mostly caused by radial electric fields. The electrostatic action of the TEL can be thought of as a (highly non-neutral) plasma lens [40], and this effect is calculated in Section 3.4. Last, the TEL successfully propagates a longitudinal current along the bunch path, producing an azimuthal magnetic field. The contribution of this effect is also quantified in Section 3.4.

2.2.3 Equation of betatron motion

Dipoles and quadrupoles alone cannot provide significant focusing simultaneously in both transverse planes, but focusing is crucial to the stability of bunches in any synchrotron, due to the need to contain orbits with transverse momentum errors[6]. Sections 2.3 and 2.4 provide examples of mechanisms through which particles can gain transverse momentum until they collide with the sides of the beam-pipe aperture and are lost. Overall focusing in both transverse dimensions *simultaneously* is imperative to preserve the stability of a bunch. It was shown that alternating focusing and defocusing quadrupoles, separated by drifts, can provide overall focusing in both directions and therefore the stability needed for bunch preservation[41]. This concept ushered in the era of large, high-energy synchrotrons such as the Tevatron[42].

The combination of a focusing magnet, drift, defocusing magnet, and another drift can be expressed by the matrix

$$M_{\text{FODO}} \equiv M_0 M_D M_0 M_F = \begin{pmatrix} 1 & L_{\text{drift}} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ +1/f & 1 \end{pmatrix} \begin{pmatrix} 1 & L_{\text{drift}} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -1/f & 1 \end{pmatrix},$$




Figure 2.3. Sketch of a particle passing through a series of FODO cells, viewed from above. The deviations from the axis is greatly exaggerated, and different particles will exhibit different paths. In the vertical plane, the roles of the focusing and defocusing magnets are reversed.

and is typically called a FODO cell. The stability of this cell can be proven under the condition that $L_{\text{drift}} \leq 2f$, though that proof is skipped in this discussion [38,32,6]. The majority of the Tevatron consists of FODO cells, with parameters described in Table 2.1. Near the interaction regions, large quadrupoles are used that effectively squeeze the bunches into a smaller area. Section 1.3 examines the increase in luminosity due to this compression. These regions are smoothly matched into the rest of the ring.

As a particle travels around the ring, it is continually passing through quadrupoles that have an average effect of kicking it toward the axis. A typical particle path through a series of FODO cells is drawn in Figure 2.3. This particle exercises oscillations around the closed orbit. At any point along its trip, it feels a focusing force $k(z)$. This function is nonzero inside quadrupoles and zero through drifts. Newton's law for such a linear system becomes

$$x''(z) + k(z)x(z) = 0, \tag{2.7}$$

where $k(z)$ is periodic when z completes each revolution.

This differential equation is known as Hill's equation[43] and is simply that of an undamped harmonic oscillator with a varying spring constant. The solution therefore can be written like a sinusoid, but with an amplitude and phase advance that can change with time:

$$x(z) = Aw_x(z) \cos(\phi(z) + \phi_0) . \quad 2.8$$

In this context, each particle can have its own initial amplitude A and phase ϕ_0 (corresponding to some initial conditions x_0 and x'_0), and the machine itself possesses its own amplitude and phase modifier that depends on the longitudinal position. Therefore $w_x(z)$ and $\phi(z)$ are affiliated with “machine parameters,” since they are defined by the lattice itself, regardless of the size, shape, and status of an injected beam.

It is important to note that Equation 2.8 is able to describe the path of any particle exactly, as it might appear in Figure 2.3. This description is deteriorated slightly when multi-particle phenomena is taken into account, such as collisions, space charge, and wake fields. Nevertheless, the amplitude function $w_x(z)$ and the phase advance $\phi(z)$ can be solved for any given machine lattice. These functions can be better understood by substituting Equation 2.8 into the differential equation given in Equation 2.7. Separating orthogonal terms yields two results,

$$2w'(z)\phi'(z) + w(z)\phi''(z) = \frac{1}{w(z)}(w^2(z)\phi'(z))' = 0 , \quad \text{and} \quad 2.9$$

$$w''(z) - w(z)(\phi'(z))^2 + k(z)w(z) = 0 . \quad 2.10$$

Simplifying Equation 2.9 reveals that $w^2(z)\phi'(z)$ must be constant for all z . If that constant is defined as c , it can be inserted into Equation 2.10, giving:

$$w''(z) + k(z)w(z) = \frac{c^2}{w^3(z)} . \quad 2.11$$

This equation is generally only solvable numerically, but directly relates the focusing strengths around the ring, $k(z)$, to the amplitude function $w(z)$. Additionally, integrating Equation 2.9 reveals that the phase advance between two points around the ring is

$$\Delta\phi_{z_0 \rightarrow z_1} = \int_{z_0}^{z_1} \frac{c \, dz}{w^2(z)}. \quad 2.12$$

Equations 2.11 and 2.12 reveal that both $w^2(z)$ and $\phi(z)$ scale linearly with the arbitrary constant c . In order to eliminate this unnecessary factor, a new function $\beta(z)$ is defined as:

$$\beta(z) \equiv \frac{w^2(z)}{c},$$

which is commonly called the beta function and allows $\Delta\phi$ to be rewritten as

$$\Delta\phi_{z_0 \rightarrow z_1} = \int_{z_0}^{z_1} \frac{dz}{\beta(z)}. \quad 2.13$$

Incorporating the constant c into the arbitrary constant A , Equation 2.8 can be rewritten as

$$x(z) = A\sqrt{\beta_x(z)} \cos(\phi(z) + \phi_0). \quad 2.14$$

At this point, two other parameters can be defined from the beta function for convenience:

$$\alpha(z) \equiv -\frac{\beta'(z)}{2}, \quad \text{and}$$

$$\gamma(z) \equiv \frac{1 + \alpha^2(z)}{\beta(z)}.$$

The three parameters $\beta(z)$, $\alpha(z)$, and $\gamma(z)$ are called Twiss parameters, first used by Courant and Snyder to describe the mapping of particles around a FODO-based accelerator [44].

The beta function, and the other Twiss parameters, have already been defined as machine parameters, which means they repeat in value after particles have traversed one ring revolution. The matrix representation which represents the entire ring of circumference C can be written as $\vec{x}(C) = M_{\text{rev}}\vec{x}(0)$, or

$$\begin{pmatrix} x(C) \\ x'(C) \end{pmatrix} = \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix} \begin{pmatrix} x(0) \\ x'(0) \end{pmatrix}, \quad 2.15$$

where $\beta_x(C) = \beta_x(0)$. Inserting Equation 2.8 and dividing out the common factors, Equation 2.15 becomes

$$\begin{pmatrix} \cos(\Delta\phi_{\text{rev}}) \\ -\frac{1}{\beta(0)} \sin(\Delta\phi_{\text{rev}}) \end{pmatrix} = \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix} \begin{pmatrix} \cos(\phi_0) \\ -\frac{1}{\beta(0)} \sin(\phi_0) \end{pmatrix},$$

with $\Delta\phi_{\text{rev}} \equiv \phi(C) - \phi(0)$. This system appears to contain two equations in four unknowns. However, this relation must hold for particles at any initial phase ϕ_0 , so solving the matrix for two cases (such as the orthogonal solutions $\phi_0 = 0$ and $\phi_0 = \pi/2$) defines each matrix entry exactly. After some algebra and substitutions, the solution eventually becomes

$$M_{\text{rev}} = \begin{pmatrix} \cos \Delta\phi_{\text{rev}} + \alpha \sin \Delta\phi_{\text{rev}} & \beta \sin \Delta\phi_{\text{rev}} \\ -\gamma \sin \Delta\phi_{\text{rev}} & \cos \Delta\phi_{\text{rev}} - \alpha \sin \Delta\phi_{\text{rev}} \end{pmatrix}. \quad 2.16$$

Every discussion of the Twiss parameters should mention what, on an intuitive level, the beta function is and is not. It is not the radius of a beam, though it is closely related and has units of length. It actually describes the wavelength of the particles at that particular position z , assuming they were able to freely oscillate around the closed orbit. The beta function in the Tevatron typically ranges from about 30 m inside a defocusing

lens to about 110 m inside a focusing lens, though it is crimped down to 35 cm at the collision points. It is not a coincidence that the FODO cell length is on the order of the beta function size [6]. However, the bunch size can be interpreted from Equation 2.14, because the average radius scales with $\sqrt{\beta}$. The next section demonstrates that A is generally much smaller than unity, so that the beam radius is significantly smaller than the large wavelength dimensions.

2.2.4 Transverse bunch properties: tune and emittance

The Twiss parameters in a specific area are not determined by the optics only in that area; instead, the matrix representation of the entire ring must be assembled in lattice form, multiplied together, and set equal to Equation 2.16. Then the beta function and phase advance can be derived. Repetition, such as the FODO cells defined in the previous section, simplifies this process, as does computer modeling of the lattice. When the Twiss parameters are determined at one longitudinal position, they can be tracked through the optics and plotted around the entire ring.

For example, Figure 2.4 shows a series of FODO cells similar to Figure 2.3, but this time an envelope is plotted in addition to several particle paths. The envelope equation is proportional to $\sqrt{\beta(z)}$. Two of the particles drawn are specifically chosen to have an amplitude of unity, such that, if the particle's phase passes through zero, it will reach the envelope. However, the two particles have different phase advances ϕ_0 , so they touch the envelope at different places. A third particle is identical to the first, except it has a smaller amplitude. Even though it moves synchronously with the first particle, it never reaches the outside of the beam envelope.




Figure 2.4. Illustration of three particles traveling through a FODO lattice. The envelope, proportional to $\sqrt{\beta(z)}$, is shown in grey. The first two particles have the maximum amplitude but different phases, so they brush against the envelope at different positions. The third has a much smaller amplitude, and deviates little from the central orbit.

Each particle witnesses the same phase advance as it progresses around the ring. In the Tevatron, the particles advance over twenty full betatron oscillations per revolution. This value ν is defined as the tune (or betatron tune), and can be produced from Equation 2.13:

$$\nu_x \equiv \frac{\Delta\phi_{\text{rev}}}{2\pi} = \frac{1}{2\pi} \oint_{\text{turn}} \frac{dz}{\beta_x(z)}$$

in the horizontal plane, and ν_y is analogous in the vertical plane.

The horizontal and vertical tunes are utilized extensively throughout this dissertation. The Tevatron possesses complications, discussed in Section 2.4, that depend on its tunes, and the TFL is shown in later chapters to mitigate some of these issues. The tunes in the Tevatron are easily adjusted independently, but are nominally set to 20.585 in the horizontal plane and 20.575 in the vertical plane. These numbers together define a *working point* in the two-dimensional tune space, but often the integer is ignored, revealing the *fractional* tune (for example, 0.585 and 0.575). Section 2.4 explains the motivation

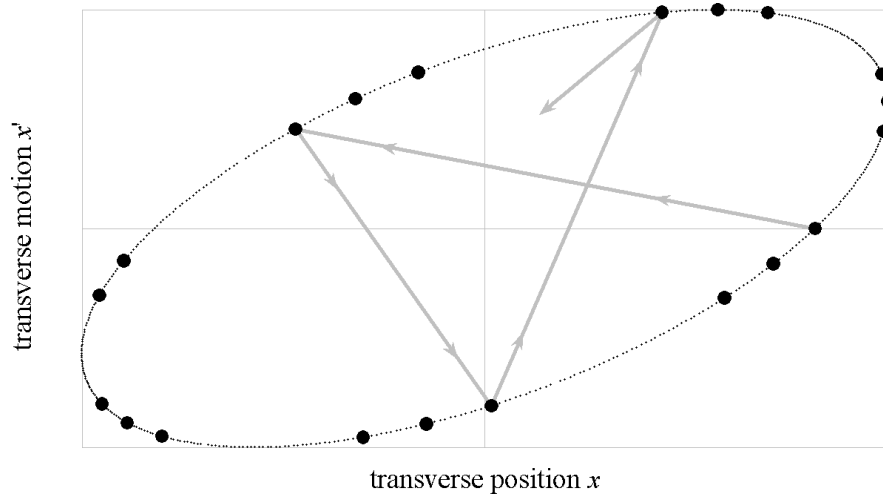


Figure 2.5. Phase space of a particle circling a ring 500 times, assuming the phase advance is 0.575. If the particle starts at position #1, then on the next passage it will appear at position #2, and so forth.

behind the choice of these values and the motivation to adjust the tunes during any store. Because the fractional tune plays a vital role in the collider’s performance, a reference to the “tune” commonly ignores the integer in favor of simply the fractional tune.

In a computer simulation, a single particle was tracked for 500 turns around a ring that possessed a fractional tune of 0.585. At a specific location, the particle’s horizontal displacement x and its derivative x' was recorded and plotted in Figure 2.5. Each point is shown, while the first twenty turns are drawn with larger dots. At each turn, the particle is at a new position and direction, but those points fill in an ellipse in phase space. If a similar analysis of a particle was performed at a different position along the ring, a different ellipse is generated. While the orientation and aspect ratio are different at each position, the area of each ellipse is identical.

The preservation of the area inside the ellipses reveals an underlying constant of particle motion. Figure 2.5 started at position #1, with position and derivative chosen

seemingly independently. However, the general equation for the particle's motion, expressed in Equation 2.8, gives only one variable, A , for the amplitude of the particle's motion (the initial phase advance ϕ_0 defines where on the ellipse the particle begins, not the size or shape of the ellipse). Therefore, A can be rewritten in terms of $x(z)$ and $x'(z)$. Since A is a constant of the motion, so must the expression involving $x(z)$ and $x'(z)$.

Squaring Equation 2.8 and its derivative permits the elimination of the sinusoidal term; in effect,

$$x^2 + (\beta_x x' + \alpha_x x)^2 = \beta_x A^2 \quad 2.17$$

at any point z around the ring. The area of the ellipse in Figure 2.5 is πA^2 and is termed the *emittance* ϵ . Using this definition, Equation 2.17 can be simplified to:

$$\frac{\epsilon}{\pi} = \gamma x^2 + 2\alpha x x' + \beta x'^2.$$

So far, this discussion has only analyzed one particle. However, the same derivations can be applied to a bunch incorporating a large number of particles; in this case, each particle travels around its own ellipse similar to Figure 2.5. Since the orientation and aspect ratio are determined by the beta function and its derivative, each particle will trace out concentric but identically shaped ellipses. Assuming the bunch is well-collimated, whichever particle traces the largest ellipse defines the total size of the entire bunch, and the emittance of that bunch is simply the area of that ellipse. The beam radius r at a specific location is defined to be the largest transverse position that this particle can ever take. Figure 2.5 indicates that this is $r_x \equiv A_x \sqrt{\beta_x} = \sqrt{\epsilon_x \beta_x / \pi}$ in the horizontal plane and similarly for r_y in the vertical plane.

For the Tevatron and most synchrotrons, the particle distribution is not nearly as well-defined as the previous paragraph purports. Due to diffusion and randomized processes, any bunch will, over time, develop a Gaussian profile[32]. The distribution can be modeled with a density function

$$n(x) = \frac{N}{\sqrt{2\pi} \sigma_x} e^{-x^2/2\sigma_x^2},$$

assuming N is the total number of particles and σ_x is defined as the rms beam size. Since there is no definite “maximum amplitude,” the beam emittance needs to be defined at some specific radius so that its ellipse circumscribes a large portion of the beam.

In general, the emittance can be defined according to:

$$\epsilon_x = -\frac{2\pi\sigma_x^2}{\beta_x} \ln(1 - f),$$

where f is the fraction of particles inside the designated ellipse. Physicists involved in hadron colliders habitually set f to the 95 % level[6], so that the definition of emittance becomes:

$$\epsilon_x \equiv \frac{6\pi\sigma_x^2}{\beta_x}. \quad 2.18$$

The emittance, under the conditions described so far, is a constant of motion for the bunch. However, when the Tevatron is ramped from injection energy (150 GeV) to its normal colliding level (978 GeV), the emittance changes significantly. Theoretically, this is because the two axes in Figure 2.5, x and x' , are not conjugate variables. Instead, a quantity known as the *normalized emittance* ϵ^N is defined in terms of the area of the traced ellipse in this space. More specifically, ϵ^N is defined as:

$$\epsilon_x^N \equiv \frac{1}{m_{\bar{p}}c} \oint_{\text{ellipse}} dp_x dx = \gamma_{\bar{p}} \beta_{\bar{p}} \oint_{\text{ellipse}} dx' dx = \gamma_{\bar{p}} \beta_{\bar{p}} \epsilon_x$$

for antiprotons, where $\gamma_{\bar{p}}$ and $\beta_{\bar{p}}$ are the relativistic factors. The unnormalized and normalized emittances have units identical to the units of $[x \cdot x']$: since x' represents an angle, the units for emittance are length times angle. In addition, the constant π is often preserved to better separate each contributor to the total area. As an example, Table 2.2 lists the typical normalized antiproton emittance as 12π mm-mrad. Subsequent chapters in this dissertation use the 95 % normalized emittance exclusively.

2.3 The real world: imperfections

The discussion thus far has analyzed an idyllic situation. Deflecting the particles through either bending dipoles or focusing quadrupoles is grounded in Equation 2.1, which simply gives the angle of deflection in terms of the particle's momentum p_z and an ideal magnetic field B . In reality, both of these quantities have errors which need to be understood and compensated for successful bunch propagation. Off-momentum particles will be deflected differently than particles with the design energy. Magnets can be constructed with intrinsically nonideal field patterns, and they can be incorrectly positioned or rotated.

2.3.1 Effects of off-momentum particles

Longitudinal dynamics plays a crucial role in all accelerators (even the name *accelerator* refers to longitudinal action), but is not the main thrust of this dissertation. However, some of the consequences push themselves onto the T_{EL}'s performance, and

therefore it is important to understand particle behavior in this dimension. The following discussion highlights relevant results, while more complete analyses can be found in many other sources [6,33,32].

Injection, acceleration, and collisions depend on a series of large copper cavities stationed along the Tevatron beamline. These RF cavities resonate at 53.1 MHz, and an intricate feedback system adjusts the phase and amplitude of this sinusoid in order to maintain the bunches and perform various functions. The excited mode of the cavity includes an on-axis electric field; it is this field that accelerates the bunches. Often compared to a surfer riding ahead of the crest of an ocean wave, the antiproton and proton bunches gain momentum as they traverse the cavities.

During acceleration, the amplitude and phase of the cavities are adjusted to transfer a large amount of forward momentum per passage. During collisions, however, the bunches pass through at a time when the electric field is nearly zero. Each turn, a noticeable percentage of kinetic energy is lost through several mechanisms (in the Tevatron, up to 50 keV is spent by a process called *beam loading* of the cavities [32,38], another 100 eV goes into wake fields around the ring, and a final 10 eV is lost to synchrotron radiation [3]). Therefore, the cavities must be phased such that the bunches regain the same energy that they lost on the previous turn.

However, another important feature of the sinusoidal electric field is the longitudinal focusing that it can provide. A particle with a nonideal energy behaves differently through each magnetic element. For example, if a specific particle had a higher energy than its peers, it would exhibit the following characteristics: a higher speed, resulting in an early arrival at any position; a lower amount of bending through the dipole magnets, resulting

in a larger circumference to travel and a late arrival; and lower focusing through the quadrupoles, resulting in a different tune. The first two characteristics have opposite effects, and for low-energy machines, the change in speed typically outweighs the change in path length, so high-energy particles precede low-energy particles.

In accelerators where the particles are ultrarelativistic, the change in speed is negligible. In these machines, low-energy particles arrive before high-energy machines. Many machines that accelerate particles pass through a transition point where the ratio between the revolution period and momentum (known as the *slip factor*, η) changes sign. The energy at which this occurs is called the *transition energy*, and for the Tevatron, it is calculated to be 17.5 GeV[45]. Since the Tevatron ramps from 150 GeV to 980 GeV, the Tevatron always operates *above transition*.

A more severe impact on the revolution period in the Tevatron stems from the change in quadrupole focusing on particle momentum. A particle with excessive momentum does not get bent as much and therefore traces out a quicker path. The amount of this effect depends on particle amplitude and lattice parameters, making it complicated to quantify. However, it is the largest contributor to the connection between an inherent momentum spread and a gradual elongation of the Tevatron bunches[3]. Counteracting this effect is the sinusoidal behavior of the RF in the cavities. The phase of the RF is adjusted such that the center of the bunch receives enough kinetic energy to remain at 980 GeV, but more energetic particles, arriving early, gain slightly less energy, and less energetic particles gain slightly more. This process constantly attempts to maintain the bunch's energy spread and longitudinal length.

Similar to the case of transverse motion, this undamped combination of displacement and restoring force sets up longitudinal oscillatory behavior. The *synchrotron tune*, ν_s ,

is defined as the number of longitudinal oscillations per revolution, but typically this number is much less than one. During collisions, the Tevatron's synchrotron tune is approximately 7.2×10^{-4} .

Particles with different energy levels trace out different paths. When they pass through quadrupoles, they are not necessarily centered anymore, and therefore they can receive a dipole kick in addition to focusing forces. The numerous bending and focusing magnets around the Tevatron create a defined function of this transverse deviation, known as the *dispersion function*, $D_x(z)$, and is equal to $\delta x_0(z)/(\delta p/p_0)$, where $\delta x_0(z)$ is the amount of transverse offset for the closed orbit and $\delta p/p_0$ is the particle's momentum difference divided by the design momentum. The dispersion function is another machine parameter (albeit related to the strength of focusing) that is independent of the momentum spread of the injected bunches.

Because the central orbit for off-momentum particles can change, the apparent beam radius σ_x can be larger than the formula given in Equation 2.18. The transverse position of any particle can be separated into the portion related to its betatron oscillations and the portion related to the orbit shift from its momentum. In other words,

$$x(z) = x_{\text{betatron}}(z) + \frac{\delta p}{p_0} D_x(z), \quad 2.19$$

and the rms beam radius is the quadrature sum of these terms, or [33]

$$\sigma_x(z) = \sqrt{\frac{\beta_x(z)\epsilon_x^N}{6\pi\beta_p\gamma_p} + \left(\frac{\delta p}{p_0} D_x(z)\right)^2}. \quad 2.20$$

The dispersion at the collision points is designed to be zero, minimizing the bunch radius.

The radius then simplifies to:

$$\sigma_{x,y}(z) = \sqrt{\frac{\beta_{x,y}(z)\epsilon_{x,y}^N}{6\pi\beta_{\bar{p}}\gamma_{\bar{p}}}}. \quad 2.21$$

The maximum horizontal dispersion around the Tevatron ring is approximately 6 m and usually positive, and the vertical dispersion varies around ± 0.4 m, since there is no significant bending in the vertical plane[46] (the major source of vertical dispersion is through *coupling*, defined in the following discussion[3]).

Additionally, the dependence of bending with momentum suggests that the focusing produced by each quadrupole is a function of momentum. Another parameter, named the *chromaticity*, relates a change in momentum to the consequential change in tune, or [6]

$$\xi \equiv \frac{d\nu}{\delta p/p_0} = -\frac{1}{4\pi} \oint_{\text{turn}} \beta(z)k(z)dz \quad 2.22$$

in both transverse directions, where $k(z)$ is the focusing strength used in Equation 2.7.

Without compensation, the negative natural chromaticity will induce a transverse instability in the bunch[50]. Additional problems associated with a large tune spread are analyzed in the next section, so limiting the chromaticity is advantageous. In order to influence it, additional magnets known as *sextupoles* are added to the lattice. These magnets have six poles and produce a magnetic field given by Equation 2.3 when a_3 or b_3 are the only nonzero terms. The field strength varies quadratically with transverse position. This is shown to lead to a focusing gradient that is dependent on the transverse position of a passing bunch.

The magnetic vector potential given by $\Psi_{\text{sex}} = b_3 r^3 \sin 3\phi$ generates a magnetic field:

$$\begin{aligned}
\vec{B}_{\text{sex}} &= -\nabla \Psi_{\text{sex}} \\
&= -\nabla (b_3 r^3 \sin 3\phi) \\
&= -\nabla (3b_3 x^2 y - b_3 x^3) \\
&= -6b_3 xy \hat{x} + 3b_3 (y^2 - x^2) \hat{y},
\end{aligned}$$

from which the force on a relativistic particle can be calculated similarly to Equation 2.4.

The horizontal trajectory changes by an amount:

$$\Delta x' = \frac{3b_3 e L_{\text{sex}}}{p_z} (x^2 - y^2) \quad 2.23$$

in the horizontal plane, and the associated focusing strength becomes

$$k_{\text{sex}} = \frac{1}{cp_z} \frac{dF_x}{dx} = \frac{6b_3 e}{p_z} x.$$

The linear dependence of the focusing strength on position makes the sextupoles uniquely useful. Placed in positions where the dispersion is high, they can focus particles an amount dependent on the particles' energies. For example, particles with the correct energy are centered at $x = 0$, so that they receive no additional focusing regardless of the strength of the magnetic field. Particles with higher energy, however, pass through at larger horizontal positions, as expressed by Equation 2.19. Without sextupole magnets, these particles would have had lower tunes due to the lesser quadrupole focusing strengths contributing to Equation 2.22. But by incorporating sextupole fields so that $k_{\text{sex}} > 0$, these higher energy particles end up receiving additional focusing. Likewise, lower-energy particles, otherwise possessing a larger tune, pass through the sextupole at negative horizontal positions and feel a defocusing force, which drives down their tune.

In this manner, the sextupole magnets help decrease the tune spread of the protons and antiprotons in the Tevatron and are routinely used during stores[47]. However, Equation 2.23 shows the effect of introducing a nonlinear magnetic element; the action in one transverse plane is now dependent on the position in the other. Transverse kicks, dependent on sextupole settings, need to be included in the closed orbit calculation. In fact, the original equation of motion must be modified to include this nonlinear driving force. Equation 2.7 is rewritten as:

$$\begin{aligned} x''(z) + k_x(z)x(z) &= \frac{3b_3e}{p_z}(x^2 - y^2)\Theta(z) \quad \text{and} \\ y''(z) + k_y(z)y(z) &= -\frac{6b_3e}{p_z}xy\Theta(z), \end{aligned} \tag{2.24}$$

where $\Theta(z)$ is unity for the longitudinal region of the sextupole in question. This is the first example of the motion appearing as coupled harmonic motion, which implies that coupled resonances are possible. A full analysis of resonant motion requires a new section.

2.3.2 Resonant motion

The general vector potential for any magnet is given by Equation 2.3. Physical magnets are never perfect, and their imperfections will result in small nonzero values in each of the a_n and b_n terms[48]. A magnet with a pole slightly out of alignment can produce higher-order and lower-order multipoles in addition to the mode it was designed to generate. Positioning a magnet off-center produces lower-order multipoles as well. Rotating them induces *coupling* between the two planes. No matter how carefully each magnetic element is manufactured and aligned, small errors always exist, and therefore

every accelerator will have unexpected magnetic multipoles that can influence particle behavior.

Such motion can be hazardous due to the simple argument that the transverse motion mimics an undamped harmonic oscillator. The horizontal motion, for example, is

$$x''(z) + k_x(z)x(z) = g_x(x, y, z), \quad 2.25$$

such that $g_x(x, y, z)$ is the horizontal driving force due to all of the nonlinear or imperfect magnet fields. These are all fields not directly included in the closed orbit or linear betatron oscillations, the motion of undriven particle behavior as described in Section 2.2. In particular, the amplitude of the betatron oscillations can easily increase without bound if the particle has a tune near specific resonant frequencies associated with the driving function $g_x(x, y, z)$.

A full quantitative analysis of the resulting motion is well-documented in numerous sources[6,32]. Typically, this includes a Floquet coordinate transformation and a Hamiltonian rewriting of the equations of motion. This dissertation employs a simpler Fourier analysis to justify resonant behavior. While less rigorous, this straightforward approach provides an intuitive approach and yields the same qualitative results.

It is beneficial to first look at motion confined to one transverse direction. In this case, any driving magnetic fields $\Delta B(x, z)$ can be written as an infinite sum given by

$$\Delta B(x, z) = \sum_{i=0}^{\infty} c_i(z)x^i.$$

This general formula describes an arbitrary transverse field in terms of the orthogonal functions x^i , and it is expected that no $c_i(z)$ is identically zero for all z . In this case, the driving function becomes

$$g(x, z) = \frac{e}{p_z} \sum_{i=0}^{\infty} c_i(z) x^i. \quad 2.26$$

Since each $c_i(z)$ repeats with period equal to the ring's circumference, a Fourier transform can separate it into its harmonic constituents. Defining $\theta \equiv \phi/\nu_x$, each term in $g(x, z)$ can be expanded via:

$$c_i(z) = \sum_{k=0}^{\infty} d_{i,k} \cos k\theta, \quad 2.27$$

where k can be any nonnegative integer and $d_{i,k}$ is a constant. For example, the c_0 term contains terms oscillating at all harmonics $k\theta$ of the revolution frequency. However, the assumption by Equation 2.14 is that the particle oscillates with a frequency ϕ . Resonant behavior appears if the natural frequency is equal to a driving frequency. Therefore trouble occurs if $\phi = k\theta = k\phi/\nu_x$, or $\nu_x = k$. Setting the tune to any integer value will allow amplitude growth through resonant driving.

The next term in Equation 2.26 is c_1x (ignoring the constant e/p_z). Substituting the general solution to particle motion, Equation 2.14, and expanding c_1 through Equation 2.27 produces terms proportional to $\cos k\theta \cos \phi$. These factors represent two beating frequencies, so that c_1x can be written as

$$c_1(z)x(z) = A\sqrt{\beta} \sum_{k=0}^{\infty} \left(d_{1,k} \cos(k\theta + \phi) + d_{1,k} \cos(k\theta - \phi) \right).$$

Setting the second cosine term equal to that of the natural frequency yields $\phi = k\theta - \phi$, or $\nu_x = k/2$. This says that setting the tune equal to any half integer allows the particle to resonant with the driving function.

In general, the $c_m x^m$ term can always be expanded via Equation 2.27 and will produce a term proportional to $\cos k\theta \cos^m \phi$. Trigonometric identities can again expand this into separate beating frequencies, including $k\theta - m\phi$. Setting this frequency equal to ϕ yields the general result that any rational tune, that is, any tune that satisfies $m\nu_x = k$ for integers $m > 0$ and $k \geq 0$ has the potential to resonate with terms in the driving force*.

So far, the discussion has only looked at driving functions in one dimension. In reality, the particle can pass through the lattice at a different x and y position, and therefore be subject to a driving force dependent on both directions. At first, the analysis seems daunting, but following the previous technique, the driving force can be expanded in both transverse dimensions, and individual Fourier frequencies can again be set equal to the natural frequency to find resonance relations.

The general form for a driving function in both transverse planes can be written as

$$g(x, y, z) = \frac{e}{p_z} \sum_{i=0}^{\infty} \sum_{j=0}^{\infty} c_{i,j}(z) x^i y^j.$$

Each individual $c_{m,n}(z)$ term can be written as a sum of harmonics mimicking Equation 2.27; each term within that series contains the oscillatory factors

$$\cos k\theta_x \cos^m \phi_x \cos^n \phi_y = \cos(k\theta_x - m\phi_x - n\phi_y) + \cos(k\theta_x - m\phi_x + n\phi_y) + \dots$$

for driving the particle in the horizontal plane, where $\theta_x = \phi_x/\nu_x$ and k , m , and n are nonnegative integers. Setting the cosine argument equal to that of the particle's free motion reveals resonances that obey the relation:

$$m\nu_x \pm n\nu_y = k \tag{2.28}$$

* The integer m is silently incremented in this result, and again in Equation 2.28. Here, visual elegance eclipses irksome rigor.

and $m \neq 0$. This derivation is calculated for horizontal driving forces, but an identical result can be generated for vertical oscillations, with $n \neq 0$. The sum $p \equiv m + n$ is known as the *order* of the resonance.

Of course, for any rational horizontal and vertical tunes, there exists some integers m , n and k such that Equation 2.28 is satisfied. It may therefore seem that the ability to avoid resonances is infinitely impossible. However, the strength of the driving function falls rapidly as the order p is increased. Typically, the Tevatron is run such that ν_x and ν_y avoid resonances of order $p \leq 12$. The specific tunes chosen, called a *working point*, is often adjusted and optimized during many stores. Ideally, the working point is just that: a point in tune space. Section 2.4 shows that the realistic spread in tunes during collisions produces a large footprint, all of which is moved around in order to avoid as many resonances as possible.

The modal analysis used in this discussion justifies the ability and propensity of particles near certain tunes to begin resonant behavior, leading to particle losses and emittance growth. However, a more rigorous approach can predict the width of each resonant line and the amplitude of its oscillations. This treatment is reproduced in numerous other sources[6,32,38]. One interesting consequence of that derivation is that resonances such that both m and n are nonnegative, termed *sum resonances*, can increase without bound. The others are called *difference resonances* and tend to limit their amplitudes. The left side of Figure 2.6 draws all of the sum resonances (plus the $\nu_x = \nu_y$ line) through the fourth order. Numbers list the order of each cluster of lines. As usual, the tune space refers to the fractional tune; the square looks identical around the Tevatron working point.

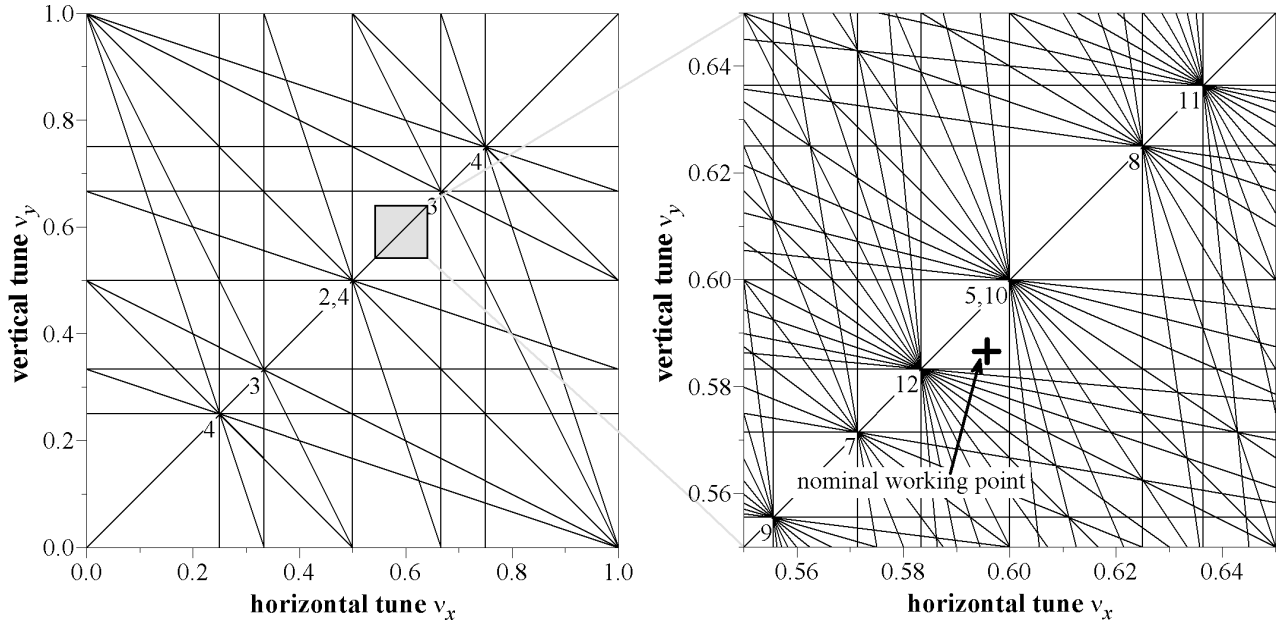


Figure 2.6. Tune plots depicting sum resonances. The left side covers the entire tune landscape and shows resonance lines through the fourth order (enumerated). The right side zooms in on a small section and includes resonances through the twelfth order.

The right side of Figure 2.6 is an expansion of a small region of tune space, but with sum resonance lines through the twelfth order drawn. A nominal working point for the Tevatron is shown, though it can be adjusted to find a region away from strong resonances. Often some of the stronger difference resonances are shown, though that is skipped here for simplicity. In practice, it is difficult to predict exactly how strong each resonance line will be for a specific machine. Calculations from lattice symmetries and tracking simulations combine to give reasonable estimates, and, as always, experimental evidence has the final voice.

2.3.3 Focusing errors

The inclusion of unexpected fields implies that the observed tune of the machine

will not necessarily equal the targeted value. If there is a single quadrupole error in an otherwise ideal lattice, the new revolution matrix would be

$$M'_{\text{rev}} = M_{\text{rev}} \begin{pmatrix} 1 & 0 \\ -\frac{1}{f} & 1 \end{pmatrix}, \quad 2.29$$

where M_{rev} is the expected revolution matrix given by Equation 2.16. Assuming that the additional gradient is small and does not excite resonant behavior, M'_{rev} would simply be a new revolution matrix of the same form as Equation 2.16, and each new Twiss parameter can be derived in terms of the old and the focal length f of the magnet error.

In particular, the trace of both sides of Equation 2.29 yields

$$\cos 2\pi\nu' = \cos 2\pi\nu - \frac{\beta}{2f} \sin 2\pi\nu$$

in each transverse dimension, with the substitution $\Delta\phi_{\text{rev}} = 2\pi\nu$. The change in tune can be defined through the substitution $\nu' = \nu + \Delta\nu$, and an expansion of the left-hand side produces the linear relation

$$\Delta\nu = \frac{1}{4\pi} \cdot \frac{\beta}{f}. \quad 2.30$$

This relation pertains to any focusing or defocusing lens added to an already stable lattice and applies to either transverse plane. Equation 2.30 is useful in both the beam-beam interaction at the collision points and the TFL's interaction region.

The linear result of Equation 2.30 implies that multiple focusing errors can be summed together. Indeed, it can be shown that a distribution of small errors generates the total tunes shift [32]

$$\Delta\nu_{\text{total}} = \frac{1}{4\pi} \sum_k \frac{\beta_k}{f_k}, \quad 2.31$$

where β_k refers to the beta function at the location of the k -th focusing error and in the relevant dimension.

This chapter has so far skirted the question of independently adjusting the horizontal and vertical tunes during a store. Section 2.2 implies that the phase advance through a FODO cell is nearly equal in both planes. This is true, due to the equal focal lengths of both quadrupoles in each cell, and only the path-length focusing of the dipoles can contribute differently to each dimension. A way around the near symmetry is suggested by the beta-function dependence in Equation 2.31. If the focusing magnet is increased in strength slightly, its focal length will decrease by some amount Δf . If the defocusing magnet is decreased in strength such that its focal length is increased by the same Δf , the tunes shifts due to both “errors” do not cancel. The larger beta function at the focusing magnet means that the total horizontal tunes shift will be slightly positive.

However, in the vertical plane, the roles of the quadrupoles are reversed, so that the vertically focusing magnet is decreased in strength, lowering the tune. Since this is where the vertical beta function is at its peak, the total vertical tunes shift is slightly negative. Through subtle manipulation of the FODO lattice, it is possible to independently change the horizontal and vertical tunes in any machine. For the Tevatron, this process is utilized often during stores, and it has proved invaluable in studying the effects of the TEL on the Tevatron operation, as explored in Chapter 6.

2.4 Beam-beam interactions

The discussion so far has been limited to linear motion and small higher-order perturbations around that motion. Nonlinearities in the lattice can create resonance

lines, but they also change the tune of particles transversing them. The chromaticity is an example of such a tuneshift dependent on each particle's energy.

The spread of tunes due to lattice imperfections is generally very small (typically less than ± 0.001 in the Tevatron); however, the interaction between the proton and antiproton bunches at the two collision points, CDF and DØ, generates a strong focusing force, altering the tunes of the bunches significantly (approximately $+0.02$). In addition, the focusing force is strongly amplitude-dependent. This has the effect of spreading the tunes between small-amplitude particles and large-amplitude particles within the same bunch.

As expressed in Table 2.2, the proton bunches have several times the number of particles as the antiproton bunches. The force that the antiprotons feel directly relates to the number of protons, while the force that the protons feel follows the number of antiprotons. Due to the significantly larger number of protons, it is the antiproton bunches that suffer the largest tuneshift and tune spread. If the total footprint is small enough, then it is possible to move it to a working point region away from any strong resonances. If the footprint is too large, then some antiprotons will find themselves on or near perilous resonance lines.

2.4.1 Linear weak-strong beam-beam focusing

It is vital to understand the forces that antiprotons feel when they collide with a proton bunch. This analysis is much simpler if the proton bunch is assumed to be unaffected by the antiprotons — a reasonable assumption, since the number of antiprotons is significantly less than that of the protons, and therefore so are the associated fields. This

assumption is called a *weak-strong* approximation, as the protons are nearly impervious to the antiprotons. The forces acting on single antiprotons are analyzed.

The proton bunch typically exhibits an approximately Gaussian transverse profile, so that its charge density is

$$\rho_p(r, z) = \frac{\lambda_p(z)}{2\pi\sigma_r^2} e^{-r^2/2\sigma_r^2}, \quad 2.32$$

where σ_r is the rms bunch size at the interaction point and $\lambda_p(z)$ is the longitudinal linear charge density. The longitudinal bunch shape is typically Gaussian or more complex, but it will prove irrelevant under the thin-lens approximation identical to that employed in Section 2.2.

An antiproton passing through the proton bunch will witness a force given by:

$$\vec{F} = -e(\vec{E} + \vec{v} \times \vec{B}) = -e(E + cB)\hat{r}$$

for a round proton bunch. Since the bunch is long in its own rest frame, the electric field is approximately radial, and Gauss' law can be applied to Equation 2.32 to get the electric field,

$$\vec{E}_p = \frac{\lambda_p}{2\pi\epsilon_0 r} (1 - e^{-r^2/2\sigma_r^2}) \hat{r},$$

and the magnetic field likewise follows from Ampère's law;

$$\vec{B}_p = \frac{\mu_0 c \lambda_p}{2\pi r} (1 - e^{-r^2/2\sigma_r^2}) \hat{\phi}.$$

Paralleling the derivation of magnetic focusing elements in Section 2.2, the total kick $\Delta r'$ becomes

$$\Delta r' = \frac{\Delta p_r}{p_z} = -\frac{e^2 N_p}{2\pi\epsilon_0 c p_z r} (1 - e^{-r^2/2\sigma_r^2}), \quad 2.33$$

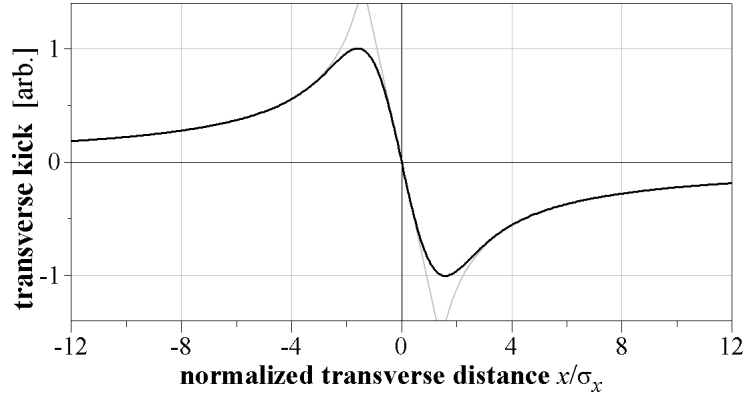


Figure 2.7. Plot of Equation 2.30, normalized to the proton bunch's rms bunch size σ_r and the maximum deflection angle. Two limits are also drawn, the $1/r$ dependence at large separation and the linear dependence near the center.

where the additional factor of $1/2$ stems from the fact that the test antiproton needs only to travel half of the length of the proton bunch, since over the same time, the proton bunch has also passed by the antiproton (the fields, the geometry, and the time are all measured in the laboratory's frame of reference).

In Figure 2.7, the transverse deflection given in Equation 2.33 is plotted as a function of the radial position of the incident antiproton. The vertical axis is scaled to the maximum deflection, and the horizontal axis is scaled to the rms beam size σ_x of the proton bunch. Antiprotons are expected to be focused by the proton bunch, so the kick is negative for an antiproton passing at a positive horizontal position, and the kick is positive for a negative horizontal position.

Two important limits should be noted in Equation 2.33. For large x , the exponential is nearly zero and the kick follows a $1/x$ dependence. This imitates the force a particle experiences due to a line charge. In Figure 2.7, a $1/x$ asymptotic line is also drawn to illustrate the rate at which the two functions converge. If the antiproton passes by the

proton bunch at $3\sigma_x$, the exponential term affects the kick by only 1 % from a simple $1/x$ dependence.

The second limit is also shown in Figure 2.7. At small x , the exponential term can be expanded, cancelling the constant and producing a linear relationship between separation and deflection. In this case,

$$\Delta x'(x \rightarrow 0) = -\frac{e^2 N_p}{4\pi\epsilon_0 c p_z \sigma_x^2} x. \quad 2.34$$

This line is drawn in Figure 2.7 and brings back memories of the linearly focusing quadrupole defined in Equation 2.5. Therefore, small-amplitude antiprotons feel a linear focusing force in both the horizontal and vertical planes at each collision point, and the tune will be increased, according to Equation 2.30, by an amount

$$\Delta\nu_0 = \frac{1}{4\pi} \frac{\beta^*}{f_{\text{IP}}} = \frac{e^2 N_p \beta^*}{16\pi^2 \epsilon_0 c p_z \sigma_x^2} = \frac{3R_p N_p}{2\epsilon_p^N}, \quad 2.35$$

where $R_p \equiv e^2/4\pi\epsilon_0 m_p c^2 = 1.53 \times 10^{-18}$ m is defined as the classical proton radius and ϵ_p^N is the proton bunch's normalized emittance, defined in Equation 2.18. It is interesting to note that the final formula does not depend on the machine lattice at all. A similar derivation for the tunes shift caused by an asymmetric proton bunch ($\sigma_x \neq \sigma_y$) gives[51]:

$$\Delta\nu_{0x,y} = \frac{R_p N_p \beta^*}{2\pi\gamma_{\text{p}} \sigma_{x,y} (\sigma_x + \sigma_y)}$$

in each transverse dimension.

The last two equations correspond to antiprotons with betatron amplitudes small enough to lie well within the linear approximation shown in Figure 2.7. A specific machine will often refer to this tunes shift as the zero-amplitude tunes shift or as the beam-beam parameter (usually with the letter ξ).

2.4.2 Amplitude-dependent tuneshift

Equation 2.35 sets the maximum tuneshift an antiproton can feel due to its interaction with a proton bunch. Core antiprotons are shifted by this amount, but larger-amplitude particles will be shifted significantly less. Computation of the average tuneshift for a particle with a specific amplitude has been performed in other publications; only the final conclusions are reiterated in this discussion [52,53].

In general, an antiproton traces out an ellipse in phase space, as drawn in Figure 2.5. By mapping normal phase-space coordinates to what are called *action variable* coordinates, it is possible to find a convolution integral that predicts the average tuneshift of the antiproton. This integral simplifies to

$$\frac{\Delta\nu}{\Delta\nu_0} = \frac{2}{\pi a} \int_0^{2\pi} \frac{1}{a^2} \left(1 - e^{-(a^2/2)\cos^2\phi}\right) d\phi,$$

where $\Delta\nu/\Delta\nu_0$ is the ratio of the antiproton's tuneshift to that of a zero-amplitude antiproton and a is the antiproton's amplitude normalized to the beam size of the proton bunch (that is, $a \equiv A_{\bar{p}}\sqrt{\beta_{\bar{p}}}/\sigma_p = A_{\bar{p}}/\epsilon_p$ using the general solution given by Equation 2.14). Solving this integral employs the modified Bessel function I_0 [54]:

$$\frac{\Delta\nu}{\Delta\nu_0} = \frac{4}{a^2} \left(1 - I_0\left(\frac{a^2}{4}\right) e^{-a^2/4}\right). \quad 2.36$$

Figure 2.8 plots this function, which is normalized to unity at $a = 0$. If the betatron amplitude of the antiproton is nearly zero, then the tuneshift generated by each collision point is equal to the maximum tuneshift calculated in Equation 2.35. Another antiproton with a larger betatron amplitude undergoes a smaller tuneshift. An antiproton with an

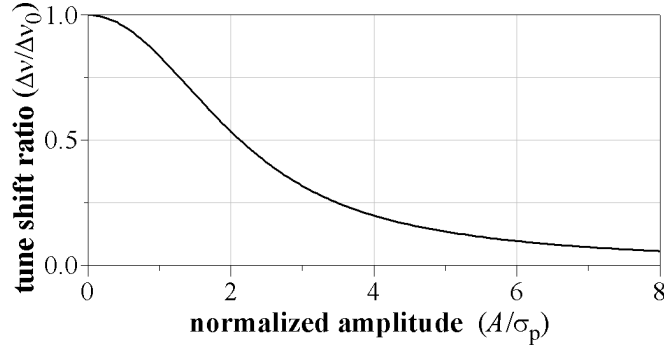


Figure 2.8. Diagram illustrating the tuneshift experienced by antiprotons with different betatron amplitudes. As $a \rightarrow 0$, the tuneshift approaches the zero-amplitude shift of Equation 2.35. The horizontal axis is scaled to the rms proton bunch radius σ_p .

amplitude equal to five times the rms proton size, for example, is only tune-shifted 13 % of the maximum.

Particularly interesting is the fact that a large portion of the antiprotons are smeared in tune space. Using parameters listed in Table 2.2, the antiproton bunch radius is about 20 % less than that of the proton bunch, so 68 % of the antiprotons' tunes are shifted at least 80 % of the maximum tuneshift. This leaves the other 32 %, still a large percentage, at tunes less than 80 % of the maximum.

During normal Tevatron operation, the maximum tuneshift can be calculated from Equation 2.35. The antiprotons typically experience a +0.010 maximum tuneshift per interaction point, or +0.020 total in each transverse dimension. The proton bunches are also affected by the antiprotons, and their tunes get increased also. However, due to the much smaller number of antiprotons, the protons are shifted about five times less, which means that they reside around a different tune region. The total footprint, as it would be drawn into Figure 2.6, must include both species, and is dramatically enlarged by these

head-on interactions. A discussion of mitigating this large footprint launches the next chapter.

Chapter 3:

Requirements for tuneshift compensation

Things should be made as simple as possible — but no simpler.
—Albert Einstein

The last chapter is an overview of standard accelerator physics, including the linear tuneshift of the colliding proton and antiproton bunches. However, Section 2.4 outlined the complications of beam-beam interactions: at the interaction points, the betatron tunes of outlying particles are shifted significantly less than that of the core particles, and, away from the interaction points, long-range fields cause additional tuneshifts unique for each bunch.

The spread in tune within each bunch and the bunches' differing tuneshifts both create a major difficulty in finding a stable Tevatron working point. Figure 3.1 presents

the antiproton tune footprint of a typical store, which crosses a number of resonance lines. This large footprint can be moved to different working points in order to minimize the emittance growth and particle losses, and in fact this is commonly done during many Tevatron stores. However, increasing the number of protons enlarges the antiproton-footprint size even more, which renders the possibility of finding a decent working point nearly impossible. This constraint on the available luminosity in any circular collider is referred to as the beam-beam limit and has impeded the functionality of colliders for decades.

Chapter 1 introduces the basic concept of the Tevatron Electron Lens, or TEL, an experimental device that is able to counteract the ill beam-beam effects of circular colliders. By breaking through this limit, future colliders will be able to take advantage of larger numbers of particles without suffering the consequences of large tune footprints.

This chapter elucidates the minimum requirements that the TEL must achieve in order to provide useful tuneshift compensation. By analyzing the focusing strength of an electron beam, a constraint on the current and interaction length is set. Practical considerations, such as adjusting the current and passing the current through the beam pipe, are also discussed.

A theoretical understanding of the TEL is developed in Chapter 4, and a full physical description of the TEL layout is saved for Chapter 5. However, it should be noted that most of the requirements are somewhat flexible; for example, a longer interaction region implies less current is needed. Historically, competing parameters were juggled as the optimization process matured. In this chapter, however, the final values listed in Table 5.1 are simply shown to satisfy the needs of beam-beam compensation.

3.1 Specific goals of beam-beam compensation

The overall mission of any beam-beam compensation device is to decrease the size of the tune footprint. However, outlining the specific goals of such a device adds focus to the discussion of the TEL.

- 1. Linear tuneshift compensation.** The first and simplest method to decrease the size of the antiproton tune footprint is to adjust the nominal tune for each bunch independently. As Section 3.2 discusses, an electron beam with a flat charge-density profile will provide a linear focusing force, which shifts the tune of each particle in a specific bunch the same amount. The spread of tunes within each bunch is therefore not affected, but (ideally) all 36 bunches can lie on top of each other in the tune plane. Figure 3.1 illustrates different examples of such a footprint. In order to shift each bunch's tune differently, the electron-beam current is modulated between bunches. This feature is discussed in Section 3.2.
- 2. Nonlinear tuneshift compensation.** The tune spread within each bunch, as derived in Section 2.4, is related to the betatron amplitude of each particle. The core (small amplitude) antiprotons feel a large tuneshift from the protons, but outlying (large amplitude) particles feel less tuneshift. In total, each bunch creates a large footprint. This issue can be remedied if the electron-beam profile is shaped to mimic that of the proton bunches. Complete control over the electron beam profile is a daunting request, and a full exploration of nonlinear compensation is not reported in this dissertation. Nevertheless, numerous experiments related to changing the beam profile are examined in subsequent chapters, and an example of nonlinear compensation is shown in Figure 3.1.

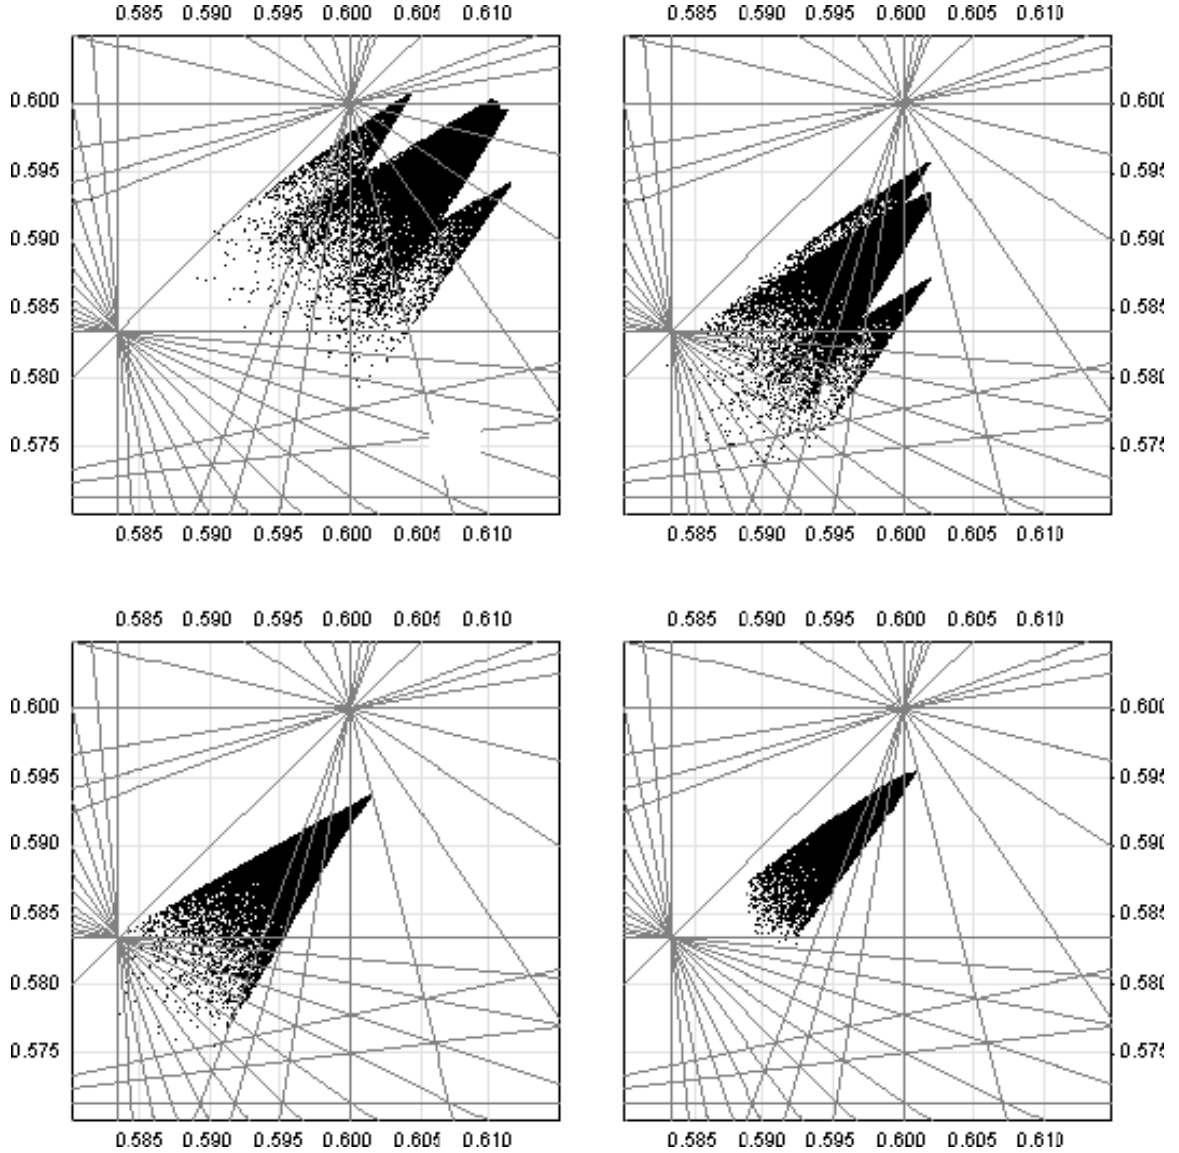


Figure 3.1. Examples of schemes for beam-beam compensation. The upper-left situation shows the original unperturbed antiproton footprint. The upper-right diagram shows the footprint if linear compensation is applied to each of the bunches, but the horizontal and vertical shifts are not independently controlled. The lower-left case allows independent control of horizontal and vertical tunes, so that all 36 bunches lie on top of each other. In the lower-right situation, nonlinear compensation is added to the linear, creating a dramatically small footprint. These simulations of the Tevatron used tracking codes and the insertion of one or two T_FLs at different places around the ring.

3. Turn-by-turn and long-term stability. As Section 2.3 mentions, focusing elements with fluctuating strengths encourage emittance growth and therefore a reduction in luminosity. The T_{FL} is identical to other focusing elements in this regard. In general, any emittance growth generated by the T_{FL} needs to be significantly smaller than the overall emittance growth observed in the Tevatron. The details of how emittance growth are caused by the T_{FL} are presented in Section 6.2, but the conclusion is that the turn-to-turn current stability ought to be better than approximately 0.5 %. Over longer time scales, the limit on current variations can be more lax; even 720 Hz and 60 Hz structure in the current is slow compared to the 47.7 kHz revolution frequency, and a percent of irregularity is not maleficent.

4. No adverse effects on the Tevatron. Adding a unique and experimental device such as the T_{FL} to the Tevatron during Run II can be risky. It is important that the T_{FL} can be turned on or off (including failures, such as solenoid quenches) with a minimum of harm to the Tevatron. T_{FL} magnets, instabilities, misalignment, and turn-on procedures are all examples of indirect methods that the T_{FL} can disturb a coasting store. Many of these effects are analyzed and tested in Chapter 6.

The T_{FL} is envisioned as a device that improves the functionality of the Tevatron. Turned off, the T_{FL} must be transparent to the operation of the Tevatron. Turned on, the T_{FL} should not restrict the manner that the Tevatron behaves. Instead, the T_{FL} is intended to simply broaden the range of operational parameters that the Tevatron runs by. A likely scenario is that the time-average luminosity remains more constant, that is, has a longer lifetime, over the course of a store. An estimate of this increase is

roughly 10–15 %, an appreciable enhancement of the Run II program[3]*

The remainder of this chapter derives the tuneshift that the T_{EL} can deliver, and derives basic minimums of the T_{EL} parameters. Additionally, the T_{EL} needs to be operated correctly in order to shrink the tune footprint as drawn in Figure 3.1. A description of the pulsing technique is introduced, along with a description of how two T_{EL}s can be used to maximize the effectiveness of tuneshift compensation.

3.2 Tuneshift mechanics using an electron beam

The linear tuneshift ν_0 of an antiproton bunch passing through a proton bunch at each IP is found in Section 2.4 to be:

$$\nu_0 = \frac{3R_p N_p}{2\epsilon_p^N},$$

where $R_p \equiv e^2/4\pi\epsilon_0 m_p c^2 = 1.53 \times 10^{-18}$ is the classical proton radius, N_p is the number of protons in the bunch, and ϵ_p^N is the emittance of the proton bunch. In a similar manner, the tuneshift of an antiproton passing through an electron beam can be calculated. The picture looks like Figure 3.2, where the electron beam enters the Tevatron pipe, gets steered onto the antiproton orbit, interacts with the antiprotons for about two meters, and turns back out of the pipe. The proton orbit, not drawn, passes through on an orbit anti-parallel to the antiprotons. However, this orbit is significantly far away from the

* A pragmatic measurement of utility has units of dollars: The budget for Run II is about \$100 million/year, so a 10 % increase in luminosity can be valued at \$10 million/year. If the T_{EL} delivers this yield for four years, its worth totals \$40 million. Of course, this type of measurement ignores the main notion behind basic scientific research.




Figure 3.2. Cartoon of the interaction of the electron beam and the antiprotons. The protons pass significantly far away from the electron beam that they feel little effect. It is assumed for now that the electron beam profile is flat and that its radius is larger than that of the antiproton bunch.

electron beam, and for the protons, the electric and magnetic forces oppose each other. In total, the proton bunches are not discernably affected by the electron beam.

First, the focusing strength and tunes shift of an electron beam is formulated. This relation is then used to set lower bounds on specific T_{EL} parameters. Additionally, nominal values for T_{EL} operation are justified. Chapter 4 continues with a more elaborate description of the electron beam, and Chapter 5 details the apparatus.

3.2.1 Derivation of antiproton tunes shift

An antiproton passing through a beam of electrons, as sketched in Figure 3.2, senses a radial electric field and an azimuthal magnetic field generated by the beam. Given a

uniform electron charge density ρ , the beam current is $I = \rho v(\pi r_e^2) = \rho \beta_e c(\pi r_e^2)$, and Gauss' law can be applied inside and outside the beam to derive the electric field[35]:

$$E_e(r) = \begin{cases} \left(\frac{I}{2\pi\epsilon_0\beta_e c} \right) \frac{r}{r_e^2} & \text{for } r < r_e, \\ \left(\frac{I}{2\pi\epsilon_0\beta_e c} \right) \frac{1}{r} & \text{for } r > r_e; \end{cases} \quad 3.1$$

where r_e is the radius of the electron beam. Ampere's law gives a similar result for the magnetic field:

$$B_e(r) = \begin{cases} \left(\frac{\mu_0 I}{2\pi} \right) \frac{r}{r_e^2} & \text{for } r < r_e, \\ \left(\frac{\mu_0 I}{2\pi} \right) \frac{1}{r} & \text{for } r > r_e. \end{cases} \quad 3.2$$

Of course, these formulas are largest at $r = r_e$. For a nominal electron beam in the TEL, $I = 2$ A, $r_e = 1.6$ mm, and $\beta_e = 0.2$; at these parameters, the peak electric field is $E_e(r = r_e) = 500$ kV/m and the peak magnetic field is $B_e(r = r_e) = 3.33 \times 10^{-4}$ T.

Paralleling the derivation in Section 2.4, the tuneshift can be calculated from the radial force $\vec{F} = e(E_r + cB)\hat{r}$. Assuming an antiproton has a betatron amplitude smaller than the beam radius, it will feel a radial momentum kick

$$\begin{aligned} \Delta p_r &= F \Delta t \\ &= \frac{eL_{\text{TEL}}}{c} \left(\frac{I_e r}{2\pi\epsilon_0\beta_e c r_e^2} + \frac{\mu_0 c I_e r}{2\pi r_e^2} \right) \\ &= \frac{(1 + \beta_e)}{\beta_e} \cdot \frac{e I_e L_{\text{TEL}} r}{2\pi\epsilon_0 c^2 r_e^2}. \end{aligned} \quad 3.3$$

This momentum kick is outward (that is, defocusing), and since the TEL is constructed such that the current of the electron beam opposes that of the antiproton bunch, the

magnetic kick supplements the electric kick (had the currents traveled in the same direction, the two terms would subtract, producing a $(1 - \beta_e)$ factor instead of the $(1 + \beta_e)$ seen here).

Since the antiproton's longitudinal momentum p_z is simply $\gamma_{\bar{p}} m_{\bar{p}} c$, the focal length of the lens becomes

$$f = \frac{r}{\Delta p_r / p_z} = \frac{\beta_e}{(1 + \beta_e)} \cdot \frac{ecr_e^2 \gamma_{\bar{p}}}{2R_p I_e L_{\text{T\overline{E}L}}},$$

where R_p is again the classical proton radius. For a uniform current profile of the electron beam, the focal length is independent of the incoming particle's position.

The tuneshift of any focusing element is calculated in Equation 2.30, so that the shift due to the electron beam is

$$\Delta\nu_x = \frac{1}{4\pi} \cdot \frac{\beta_x}{f} = \frac{\beta_x}{4\pi} \cdot \frac{(1 + \beta_e)}{\beta_e} \cdot \frac{R_p I_e L}{ecr_e^2 \gamma_{\bar{p}}} \quad 3.4$$

in the horizontal direction and similarly for $\Delta\nu_y$ in the vertical direction. This formula, the keystone of the T $\overline{\text{E}}$ L experiment, relates the tuneshift of the antiprotons to the various parameters of the T $\overline{\text{E}}$ L.

If the electron beam profile is azimuthally symmetric, then the focal lengths in the horizontal and the vertical directions will be identical. The tuneshift in Equation 3.4 is linear with the beta function; therefore the vertical tuneshift and the horizontal tuneshift follow the ratio

$$\frac{\Delta\nu_x}{\beta_x} = \frac{\Delta\nu_y}{\beta_y}. \quad 3.5$$

Since the beta functions are determined by the lattice, this relationship denies the possibility to independently control the horizontal and vertical tuneshifts. The T $\overline{\text{E}}$ L was

installed in the Tevatron at a position where $\beta_x \approx 101.67\text{ m}$ and $\beta_y \approx 30.89\text{ m}$, so that $\Delta\nu_x/\Delta\nu_y \approx 3.3$. The horizontal tuneshift, being larger than the vertical, is easier to measure. The results shown in Chapter 7 tend to emphasize horizontal values, but for every observed horizontal tuneshift, an appropriate amount of vertical tuneshift was also witnessed.

3.2.2 Minimum interaction dimensions

The rest of this chapter attempts to justify decisions made about the physical parameters of the original TEL. When the design of the TEL was being developed, numerous issues were identified and evaluated, which limited different parameters of the TEL to ranges. The choice of specific values was based on limiting the complexity of the overall design. For example, a longer interaction length would lower the needed electron-beam current, but it increases the complexity of the magnets. Using copper instead of superconducting wires in the solenoid eliminates the need of liquid helium, but requires much more space in the Tevatron tunnel. In this discussion and subsequent chapters, various issues are described, and the choices made during the design process are justified.

The original TEL was designed to shift the antiproton tune to cover the spread between individual bunches, as illustrated in Figure 3.1. This variation is theoretically calculated to be about 0.007 and measured at less than 0.006. The tuneshift due to an electron beam has been shown in Equation 3.4. If the interaction length L is 2.05 m and the energy of the electrons is 10 keV (giving a speed of about $\beta_e = 0.2$), the necessary current of the electron beam can be calculated. A beam of current 2 A and radius 1.6 mm gives a horizontal tuneshift of 0.00458.

The calculation in the previous discussion assumed that the electron beam radius is larger than the antiproton bunch radius. The latter is unfortunately not a specific number, as the antiproton distribution approximates a Gaussian profile. In Section 2.2, the radius containing 95% of the particles was shown to be

$$\sigma_{x,95\%} = \sqrt{\beta_x \epsilon_x} = \sqrt{\beta_x \frac{\epsilon_x^N}{\beta_p \gamma_p}}.$$

At the T_{EL}, 95% of the antiprotons are within a radius of 1.2 mm, and over 99% are inside 1.5 mm. Results presented in Section 7.2 underscore the negative impact that the edge of a uniform electron beam have on large-amplitude particles. The T_{EL} was built to have a flexible beam size through the use of magnetic compression, but its nominal radius is shown in Section 5.1 to be 1.6 mm. If the focusing strength of the T_{EL} is to be preserved, increasing the beam radius through manipulation of the magnetic compression requires more total current. The necessary current scales as the square of the beam radius, which is shown in Section 4.1 to be challenging.

The electron beam must be able to change currents between neighboring bunches, which are spaced every 396 nsec. Electrons located at the gun cathode when the first bunch passes must be able to travel the entire distance from the gun through the interaction region, and into the collector before the second bunch arrives (theoretically, only the actual interaction region needs to be refreshed, but this idea is realistically difficult to implement and gains little). If the entire path length from gun to collector is 2.5 m and, if the speed of the electrons is $0.1c$, then the total transit time is $t = d/v = 84$ nsec, significantly less than the time between bunches. Figure 3.3 illustrates a simple method for ramping the T_{EL} current between the passage of two antiproton bunches.




Figure 3.3. Cartoon of ramping the T_EL current between adjacent antiproton bunches. The ramp shape mimics the real shape as observed in Figure 5.28. The events along the time axis indicate when the head of an antiproton bunch first enters one end of the electron beam and when the tail finally exits the other; this duration is much longer than just the length of the antiproton bunch. It is important that the electron-beam current reaches the desired level and is stable there early enough to ensure that the entire interaction length has a defined charge density.

It is important that the T_EL reach its “set” current long enough that the entire two-meter length has the same charge density; if the current level is varying as the antiproton bunch passes, then the tunes shift will be partially dependent on the previous current level. The antiprotons, moving at essentially c , take less than 8 nsec to traverse the length of the T_EL, less than the 84 nsec maximum time the electron beam requires to cover the same distance. The total is 92 nsec, which is the entire length of time that the electron-beam current needs to be steady if it is to provide the full focusing force that was initially desired.

3.2.4 Anode potential

The voltage applied to the anode directly relates to the amount of current that

the electron gun supplies. The complexities associated with the cathode voltage, as the preceding sections discussed, determine whether the current is able to pass through the TEL's main solenoid. Under normal operating conditions, the cathode voltage is sufficient, and the anode voltage becomes the means to select and change the TEL's focusing strength.

This relation is derived in Section 4.1, where it is proved that for any gun geometry, the current that is drawn from the cathode surface is related to the anode voltage by the relation

$$I_e = PV_{\text{anode}}^{3/2}. \quad 3.6$$

The constant P has units of perts, equivalent to the units expressed for the beam-pipe acceptance. In this relation and the rest of this dissertation, V_{anode} refers to the voltage on the anode with respect to the voltage set on the cathode. This voltage is ideally either zero, providing no current, or positive, even though the voltage applied to the anode is typically always below that of the grounded beam pipe.

Section 4.1 discusses typical gun perveances. Relevant geometries typically yield perveances around two to four microperts, and two guns were designed, tested, and utilized in the TEL. The first gun, known as the *flattop* gun, measured a perveance of $4.35 \mu\text{P}$, and produced an even distribution of current over a large region. The second gun produced a *Gaussian*-like distribution and possessed a perveance of $2.1 \mu\text{P}$. In order to provide a desired current of two amperes, the anode of each gun needs to be pulsed to 5.4 kV for the flattop gun and 9.7 kV for the Gaussian gun. Operation of these guns is discussed further in Section 5.2.

Pulses at these voltage levels are not insurmountable. However, the ability to produce a staircase-like pattern, as sketched in Figure 3.3, with a fast risetime and repeatable to voltages well under 1 %, has proved challenging. Section 5.4 details the complexities associated with anode modulation, but in general the ability to pulse the anode to high voltages has diminished over time. A pulsed voltage of around 5 kV has been sustainable throughout the life of the TEL, but significantly higher voltages have only been obtained early during the experiment.

If the bunches are separated by 396 nsec and it takes as much as 84 nsec for the electron beam to complete its trip from the gun down the length of the main solenoid, the risetime of the anode modulator must be less than 312 nsec. This includes being able to shut off completely when the TEL needs to be off.

As described in Figure 3.3, the anode needs to follow an arbitrary staircase-like waveform. This requires the modulation circuit to be a high-bandwidth amplifier instead of a simpler pulse generator. The modulation circuit should ideally reach to DC in frequency, since the desired waveform sits on the positive side of the cathode voltage. Since the cathode voltage is adjustable, the “zero level” of the modulator must float on top of this voltage while not affecting it. Section 5.4 addresses the ability of the TEL modulator to conform to these constraints.

3.2.5 Two tunes, two lenses

For any azimuthally symmetric electron beam, the horizontal and vertical tunes will follow Equation 3.5. The ability to independently adjust the horizontal and vertical shifts for each bunch does not exist. While the tune footprint can be decreased somewhat,

it is not nearly as small as if the bunches were gathered to a single working point. This can be seen in the upper-right frame of Figure 3.1, where the size of the footprint has improved over that of the original (upper-left), but it is not as impressive as if all of the bunches resided at the same tune (lower-left).

By installing a second T_{FL}, this handicap can be alleviated; however, the beta-function ratio needs to be different at the two positions. If this condition is met, then by adjusting the current levels of each T_{FL} independently, specific horizontal and vertical tuneshifts can be chosen for each bunch.

In general, the beta functions where the first T_{FL} is located can be defined as β_{x1} and β_{y1} , and it produces tuneshifts $\Delta\nu_{x1}$ and $\Delta\nu_{y1}$ for a current I_1 . Likewise, the second T_{FL} can have equivalent parameters β_{x2} , β_{y2} , $\Delta\nu_{x2}$, $\Delta\nu_{y2}$, and I_2 . Any specific bunch could require specific tuneshifts $\Delta\nu_{xT}$ and $\Delta\nu_{yT}$, which can be produced by

$$\begin{aligned}\Delta\nu_{xT} &= \Delta\nu_{x1} + \Delta\nu_{x2} = c_1\beta_{x1}I_1 + c_2\beta_{x2}I_2 , \\ \Delta\nu_{yT} &= \Delta\nu_{y1} + \Delta\nu_{y2} = c_1\beta_{y1}I_1 + c_2\beta_{y2}I_2 ;\end{aligned}$$

where c_1 and c_2 fold in the other parameters associated with Equation 3.4. In this scenario, the two T_{FL}s, set to I_1 and I_2 , produce the desired amount of tuneshift.

Reversing these equations enables a quick determination of the necessary currents. To provide the horizontal and vertical tuneshifts $\Delta\nu_{x1}$ and $\Delta\nu_{y1}$, the two T_{FL}s need to be set to the currents I_1 and I_2 given by

$$I_1 = \frac{\beta_{y2}\Delta\nu_{xT} - \beta_{x2}\Delta\nu_{yT}}{c_1(\beta_{x1}\beta_{y2} - \beta_{x2}\beta_{y1})} \quad \text{and} \quad I_2 = \frac{\beta_{y1}\Delta\nu_{xT} - \beta_{x1}\Delta\nu_{yT}}{c_2(\beta_{x2}\beta_{y1} - \beta_{x1}\beta_{y2})}. \quad 3.7$$

If the two TELs are identical, then $c_1 = c_2$ and the ratio of I_1 to I_2 depends solely on the beta functions at the two locations. Obviously, only positive currents are allowed, which implies that the desired tunes shifts need to be above some thresholds. These minimums can be calculated by solving Equation 3.7 for $\Delta\nu_{xT}$ and $\Delta\nu_{yT}$, and physically they can be verified by independent adjustments of the two currents I_1 and I_2 .

For the proof-of-principle experiment of this dissertation, however, only one TEL has been built and operated. As this dissertation is being written, a second TEL is being assembled and tested. The second TEL is positioned in the Tevatron in sector AØ, where $\beta_{x2} \approx 49$ m and $\beta_{y2} \approx 140$ m, which is almost the reverse of $\beta_{x1} \approx 102$ m and $\beta_{y2} \approx 31$ m located at the F49 sector. Placing it where the beta-function ratio is significantly different is of primary importance, and it is hoped that, in the near future, the two TELs will be operated together to explore the possibility of full linear compensation of the horizontal and vertical tunes shifts, expressed in Figure 3.1.

Chapter 4:

Electron-beam production and propagation

*In all realms of life, it takes courage to stretch your limits,
express your power, and fulfill your potential.*

—Suze Orman

The goal of Chapter 2 was to develop a generalized treatment of particle behavior in a collider, which provided a basis for many of the conclusions reached in Chapter 3. Likewise, this chapter analyzes the motion of low-energy electron beams immersed in solenoidal fields, as this theory is fundamental to an explanation of the TEL mechanics.

This chapter is divided into three distinct parts. The first analyzes electron guns and the generation of an electron beam in general. Then a description of how a solenoidal field contains a charged electron beam is developed. Finally, the loss of kinetic energy for an electron beam passing through a beam pipe is quantified in the last portion. Each

phenomenon is handled analytically for a generic beam, but direct connections to the TEL are made, elucidating some of the requirements described in Chapter 5.

4.1 Generation of an electron beam

Throughout the past century, a number of fundamentally different techniques for generating electron beams have been invented. Each style has its strengths and weaknesses, and early in the TEL program, a list of the necessary characteristics of the TEL electron beam was compared to the various schemes. As Section 1.3 states, the gun is expected to supply up to two amperes of current at roughly 10 keV of energy.

As Chapter 3 details, the current needs to be adjustable, responsive, and repeatable, and the transverse current profile ought to be malleable. Additional factors are also considered, such as very tight space requirements, high reliability, radiation hardness, and, like every project, low cost. The best option for sustained large currents at low energy levels is the well-understood thermionic emission gun[33]. Fortunately, this technology also fits best with all of the other preferences.

At its most unadulterated level, a thermionic gun is simply one very hot metal surface which is negatively charged, the *cathode*, facing another surface which has a hole in it, the *anode*. The apparatus sits in a solenoid field, and if the electric field between the cathode and anode increased enough, electrons on the cathode surface will stream off of it toward the anode. These electrons, however, become trapped by the solenoidal magnetic fields and are directed through the hole, where they proudly become known as a beam.

In order to understand the general behavior of thermionic guns more diligently, the principles behind cathode emission must first be explored.

4.1.1 Electrons from a hot cathode

The first step to understanding a thermionic electron gun is to describe how electrons are expelled from the cathode surface. A simple model for describing the bulk metal treats the conduction electrons as an electron gas obeying Fermi-Dirac statistics. Since fermions are not permitted to fill already occupied states, a metal at temperature $T \approx 0$ K is at its lowest energy, filling each state with an electron until the supply of electrons is exhausted. The highest occupied state in a cold metal defines the Fermi energy level E_F . At a nonzero temperature, the likelihood that a state at energy E is filled is given by the Fermi-Dirac distribution:

$$P(E) = \frac{1}{e^{(E-E_F)/k_B T} + 1}, \quad 4.1$$

where $k_B = 8.617 \times 10^{-5}$ eV/deg is Boltzmann's constant. Figure 4.1 plots this distribution for several temperatures. It is interesting to note that at room temperature, the probability distribution deviates only slightly from the distribution at zero temperature. However, a temperature of above 1000 K increases the number of occupied higher-energy states. The T_{PL}'s cathode is thus understandably heated to between 1100 K and 1200 K [40].

Each atom in a macroscopic piece of metal can be thought of as an individual quantum well with discrete energy levels. With infinite distance between each atom, each well can be treated independently. However, bringing the atoms closer together brings the wavefunctions together; there are redundant energy levels, and the energies for these levels begin to split. With a huge number of atoms, the degeneracy is large enough to create a band of states instead of discrete levels.

This effect is illustrated in Figure 4.2. At large interspatial distances, each atom acts like a distinct harmonic oscillator. Zero potential is set at the top of each atom, so




Figure 4.1. Graph of Equation 3.2. The probability distribution of particles filling the lowest energy shells in a metal at absolute zero and at room temperature are not very different. By the time the metal reaches 1000 K or above, however, a significant portion of the electrons are substantially energized; even a small percentage of a macroscopic cathode will provide a myriad of conduction electrons.

that free electrons have positive energy. When the atoms approach each other, isolated states become connected and create a large degeneracy. This spreads out the energy levels for the myriad of interconnected (conducting) states. Now the higher-energy electrons can move freely about the metal, but are still trapped by the surface of the material. This explains a major difference between the energy levels in conductors and the overly simplified Fermi-Dirac distribution; however, within the conduction band itself, Fermi-Dirac statistics apply, and it is valid to use this theory to describe metal behavior.

Another important effect that is observed in Figure 4.2 is the fact that even the most energetic electron requires a specific amount of energy gain before it can escape from the conductor. This difference is called the work function and is a property of each metal. The work function, typically between two and five electron-volts, depends in complicated




Figure 4.2. Cartoon showing how individual atoms, shown by separate quantum wells, will create a material with a conduction band as the interspatial distance decreases. Electrons in a warm metal will fill some of the conduction band, but they require additional energy (the work function) to escape the entire lattice.

ways on the lattice spacing, atomic size, and filled orbits. The work function for several metals is pictured in Figure 4.3. Exposing different crystal faces will somewhat alter this value, and providing a transition layer of a different material can significantly change it. Section 5.2 discusses the choices made in the TEL's cathode.

The goal of this discussion is to determine the number of electrons that will escape the conductor's surface at a given high temperature. Answering this requires figuring out the number of electrons that reach a surface region A with high enough momentum $p_z > p_z^{\min}$. Heisenberg's uncertainty principle says that the smallest volume that phase space can be divided is $d\tau = dx dy dz dp_x dp_y dp_z = h^3$, where h is Planck's constant and only two electrons (of opposite spin) can occupy this region. The Fermi-Dirac distribution is the most likely scenario of randomly distributing fermions among the available cells, so that the number of electrons in a specific cell is

$$\frac{2 d\tau}{h^3} \cdot \frac{1}{e^{(E-E_F)/k_B T} + 1}.$$

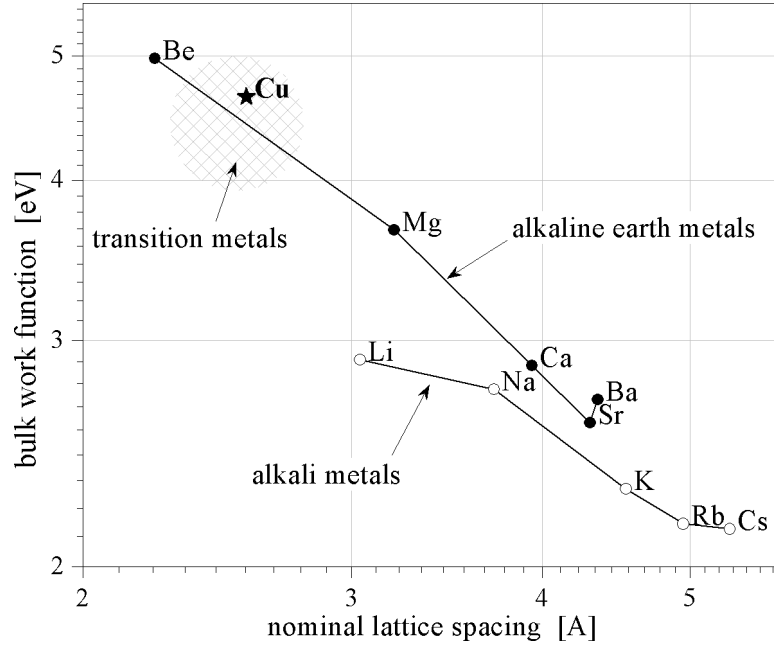


Figure 4.3. Correlation between lattice spacing and work function. Different elements will form different crystalline structures, and different faces will yield different atomic planar densities; hence the lattice spacing is not perfectly comparable. The alkali metals show a strong logarithmic connection between spacing and the work function, as do the alkaline earth metals. The transition metals tend to have similar behavior; most of them huddle in the region around copper.

The Fermi energy is the boundary between filled and empty levels at low temperatures and is given by

$$E_F = \left(\frac{3}{8\pi} \right)^{\frac{2}{3}} \frac{h^2 n^{\frac{2}{3}}}{2m_e},$$

where n is the density of free (conducting) electrons.

The number of electrons that reach the surface A during a time dt is in the volume $A dz = A (p_z/m_e) dt$. The current density J can therefore be written as

$$J = \frac{2e}{m_e h^3} \iiint_{p_z^{\min}}^{\infty} \frac{p_z}{e^{(E-E_F)/k_B T} + 1} dp_z dp_y dp_x,$$

where the energy $E = (p_x^2 + p_y^2 + p_z^2)/2m_e$. It is important to note that the limit p_z^{\min} is the minimum momentum to escape the confines of the material; that is, $p_z^{\min} = \sqrt{2m_e(E_F + W)}$.

Analyzing the integrand will elucidate certain features. The work function of copper is 4.65 eV and its Fermi energy is 7.00 eV, so at room temperature ($T \approx 300$ K) the exponent in the denominator for the lowest energy electrons, for which $E = E_F + W$, is about 180. In other words, the current density at this low temperature is minuscule. As the temperature rises to above 1000 K, the exponent decreases to about 27, allowing significantly more electrons to escape.

At temperatures significantly above absolute zero, the exponential term is much larger than unity, leading to a useful simplification of the integrand (equivalently, the distribution at moderate temperatures appears Maxwellian). In this case, pulling the constant term $e^{(E_F/k_B T)}$ and evaluating gives

$$J = \frac{2e}{m_e h^3} e^{E_F/k_B T} \int_{p_z^{\min}}^{\infty} p_z e^{-(p_x^2 + p_y^2 + p_z^2)/2m_e k_B T} dp_z dp_y dp_x ,$$

which can be separated and quickly processed:

$$J = \frac{4\pi m_e e k_B^2}{h^3} T^2 e^{E_F/k_B T} e^{-(p_z^{\min})^2/2m_e k_B T} .$$

Substituting p_z^{\min} yields

$$J = \frac{4\pi m_e e k_B^2}{h^3} T^2 e^{-W/k_B T} . \tag{4.2}$$

Equation 4.2 is called the Richardson-Dushman equation, and illustrates the maximum available current density from any cathode surface. It does not mean, however, that this many electrons will continually pour off of the conductor. In fact, as soon as one electron

does manage to escape, a net positive charge deepens the potential well that the second electron must overcome.

If a strong electric field is applied between the hot cathode and a nearby anode, all electrons overcoming the potential barrier will be pulled away, and the current density will obey the Richardson-Dushman law. This situation is known as *temperature-limited* flow and varies greatly with small changes in the cathode temperature. In practice, controlling the temperature of a cathode is very difficult, yet the current needs to be maintained to rather tight tolerances ($\Delta T/T \approx 0.1\%$ for the TEL, for example). Hence, temperature-limited emission is not a good regime for the operation of most electron guns.

4.1.2 Space-charge limited flow

However, at more moderate potential differences, a subset of electrons succeed in being pulled to the anode and the rest are "reflected" back into the conductor. This can be understood quantitatively by realizing that in steady-state flow, the current density $J(z)$ must be a constant value throughout the region between the cathode and anode. If this were not the case, excess charge would quickly build up in one spot — obviously this situation could not continue forever.

In the planar electron gun, Poisson's equation can be simplified to

$$\frac{d^2V(z)}{dz^2} = -\frac{\rho}{\epsilon_0}, \quad 4.3$$

and the definition of current is given by

$$J = \rho v = \rho \sqrt{\frac{2eV(z)}{m_e}}. \quad 4.4$$

Solving Equation 4.4 for ρ and inserting it into Equation 4.3 gives a straightforward single-variable differential equation:

$$\frac{d^2 V(z)}{dz^2} = \frac{J}{\epsilon_0} \sqrt{\frac{m_e}{2e}} V(z)^{-1/2}.$$

Integrating this twice with appropriate boundary conditions and solving for current density gives [55]

$$J = \frac{4\epsilon_0}{9} \sqrt{\frac{2e}{m_e}} \frac{V^{3/2}}{d^2} = (2.335 \times 10^{-6}) \frac{V^{3/2}}{d^2}, \quad 4.5$$

where d is the separation between cathode and anode. If d is in centimeters, J will be in amperes per square centimeter. The constant in the equation is called the *perveance* and is a heavily discussed parameter in this thesis.

Equation 4.5 describes the current flow in a planar diode gun as a function of the applied voltage difference and the plate's difference (note that for electrons, ρ is negative and thus J points from anode to cathode, even though the beam is moving away from the cathode — this assumption applies to the following discussions also). When the voltage is turned high enough, the current begins to saturate at the Richardson-Dushman limit. The mechanism that regulates the current flow below this limit is the high space charge near the cathode surface; therefore this regime is called *space-charge-limited* flow, and is the basis for a vast assortment of electron and ion-based sources.

Other geometries of electron guns can be solved through the same approach. In the cylindrical case, Poisson's equation still reduces to a one-dimensional equation:

$$\frac{1}{r} \frac{d}{dr} \left(\frac{r dV}{dr} \right) = -\frac{\rho}{\epsilon_0}.$$

Solving the resulting differential equation requires a series expansion, and the final equation requires use of an awkward series[55]:

$$\beta \equiv u - \frac{2u^2}{5} + \frac{11u^3}{120} - \dots,$$

where $u \equiv \ln(r_{\text{anode}}/r_{\text{cathode}})$. In this case, the current per unit length of the coaxial gun is found to be

$$\frac{I}{l} = (14.66 \times 10^{-6}) \frac{V^{3/2}}{r_{\text{anode}} \beta^2}.$$

It should be noticed that $\beta(u)$ starts at zero at $u = 1$, crosses $\beta = 1$ at around $r_{\text{anode}}/r_{\text{cathode}} = 11$, reaches a maximum near $r_{\text{anode}}/r_{\text{cathode}} = 40$, and approaches unity again. As long as $r_{\text{anode}}/r_{\text{cathode}} > 7$, then β is within 10 % of unity.

For spherical electrodes, a similar effect is found. Employing u from the previous paragraph, a function[55]

$$\alpha \equiv u - 0.3u^2 + 0.075u^3 - 0.00143u^4 + 0.00216u^5 - \dots$$

plugs into the total current from a spherical cathode:

$$I = (29.34 \times 10^{-6}) \frac{V^{3/2}}{\alpha^2}.$$

Ignoring edge effects, the total current out of a cathode stretching over a portion of a sphere is just the fraction represented by the area of the cathode over that of the full sphere.

In fact, it can be shown through dimensional analysis that any geometry will still provide a $J \sim V^{3/2}$ relationship. In general, Poisson's equation is

$$\nabla^2 V = -\frac{\rho}{\epsilon_0},$$

the current definition is

$$\vec{J} = \rho \vec{v},$$

and energy conservation maintains

$$\frac{1}{2}mv^2 = eV.$$

In the general case, the electron's velocity will not necessarily move along the potential gradient, so \vec{v} and \vec{J} must be vector quantities.

Let the potential be increased by some factor k . Then the space-charge density ρ will increase by k , and the velocity v will increase by $k^{1/2}$. Hence, J will increase by $k^{3/2}$, and all that remains is the constant of proportionality. The constant is known as the *perveance* and is only dependent on the geometry of the electron gun. Hence, the generalized relation is

$$I = PV^{3/2} \tag{4.6}$$

and is commonly called the Child-Langmuir law. The perveance P is usually expressed in units called *pervs*, where $1 \text{ perv} \equiv 1 \text{ ampere/volt}^{3/2}$. A couple micropervs are typical for fist-sized electron guns.

Obviously the above examples are useless for producing a beam that persists past the anode. In order to create a beam, a hole is drilled in the anode and the assembly is submerged in a solenoidal magnetic field. The latter directs the electrons along the magnetic paths, following the theory of the next section, which will often force the electrons to pass through the hole with kinetic energy equal to the voltage difference between anode and cathode. Not only is the beam size (and density profile) frozen by the magnetic field, but the anode shape can be chosen so no electrons ever touch it. Adjusting

or pulsing the anode voltage is much easier when there is no current that is drawn from the electrode.

However, the hole size for most guns is typically big enough to significantly shift the electric field lines such that analytic analysis of a given electron gun becomes impractical or impossible. Computer simulations fill that void and allow adjusting of electrodes' positions and shapes in order to optimize the gun's perveance, density profile, etc. The specific designs of the TFL electron guns is reserved for Section 5.2.

4.2 Cyclotron motion in a solenoid

Once the electrons are launched from the cathode, they act as free charged particles that are pushed around by various electromagnetic forces. Without any other force, nonrelativistic electrons impart an electric field on each other, causing the beam to spread apart. significantly. A relativistic beam of electrons also generates a magnetic field that pulls the electrons together, which mitigates the space-charge repulsion.

If the electron beam is ultra-relativistic, the two forces nearly cancel, allowing the beam to propagate without significant expansion. The TFL accelerates the electrons only to a speed of approximately $0.2c$, so that the electric space-charge force is much stronger than the magnetic current-induced force. Immersing the electron beam in a solenoidal field is shown to contain the beam to a specific radius.

4.2.1 General motion in a solenoid with space charge

The electric and magnetic fields produced by a uniform electron beam is given in Section 3.2. If there is no solenoid, then a single electron would feel a total force given by

$$\vec{F} = e(\vec{E}_e + \vec{v} \times \vec{B}_e) = e(E_e - \beta_e c B_e) \hat{r}.$$

Combining Equations 3.1 and 3.2 yields the simple relation:

$$B_e = (\beta_e/c) E_e \tag{4.7}$$

at any radius, so that the total force can be written as

$$\vec{F} = e(1 - \beta_e^2) E_e \hat{r} = \frac{e E_e}{\gamma^2} \hat{r}, \tag{4.8}$$

where γ is the electron's relativistic factor. This conclusion is applicable to any particle beam, and the $1/\gamma^2$ -dependence of the total force F_{free} verifies that the force approaches zero as the particle speed nears the speed of light. For the electron beam in the TEL, however, $\beta_e \approx 0.2$, so $\gamma \approx 1.02$. In this case, the strength of the magnetic force is only 4% of the strength of the space-charge force; the beam will expand quickly if unattended.

If a solenoid surrounds the beam, the total force acting on a particle is more complicated. By keeping the derivation general, the solutions will be applicable to both the electrons and the antiprotons.

Any round beam possesses cylindrical symmetry, so it is convenient to write the generalized motion of a relativistic particle in cylindrical coordinates:

$$\vec{v} = \dot{r} \hat{r} + r \dot{\theta} \hat{\theta} + \beta c \hat{z},$$

where βc is the particle's speed. In the laboratory frame of reference, a specific particle witnesses the following fields:

$$\begin{aligned}\vec{E}_{\text{self}} &= E\hat{r}, \\ \vec{B}_{\text{self}} &= (\beta/c)E\hat{\theta}, \quad \text{and} \\ \vec{B}_{\text{solenoid}} &= B\hat{z};\end{aligned}\tag{4.9}$$

where Equation 4.7 was substituted into \vec{B}_{self} . The relativistic equation of motion becomes

$$\dot{\vec{p}} = \gamma m \dot{\vec{v}} = e(\vec{E}_{\text{self}} + \vec{v} \times (\vec{B}_{\text{self}} + \vec{B}_{\text{solenoid}})).$$

Breaking this vector equation into orthogonal components yields the coupled differential equations

$$\begin{aligned}-\omega_c \dot{r} &= r\ddot{\theta} + \dot{r}\dot{\theta} \\ \ddot{r} &= \frac{\omega_c E}{\gamma^2 B} + \omega_c r\dot{\theta},\end{aligned}\tag{4.10}$$

where $\omega_c \equiv eB/\gamma m$ is known as the *cyclotron frequency*. The time-derivative of the first can be equated to the second, yielding the relation $\ddot{r} = \omega_c \dot{r}$. This equation describes a simple harmonic oscillator, which can be solved by

$$r(t) = r_0 + r_1 \cos(\omega_c t + \phi_r).$$

By choosing specific initial conditions, some interesting results can be gleaned. For example, a particle that sets

$$\begin{aligned}r(t=0) &= r_0 \quad \text{and} \\ \dot{r}(t=0) &= 0\end{aligned}$$

as boundary conditions produce the specific radial solution

$$r(t) = r_0 + \frac{mE}{\gamma e B^2} (1 - \cos \omega_c t). \quad 4.11$$

The radial position of the particle begins at r_0 and oscillates to a larger value. This suggests that by providing a solenoid around the beam, each particle is confined to a specific range of radii; the beam is unable to expand without bound.

Under the assumption that $r_0 \gg mE/\gamma e B^2$, the particle's radius does not change significantly during its passage through the solenoid, which facilitates a straightforward solution for the azimuthal motion. Equation 4.10 can be simplified to

$$\ddot{r} = \frac{\omega_c E}{\gamma^2 B} + \omega_c r_0 \dot{\theta},$$

which in turn gives a general solution for the azimuthal motion:

$$\theta(t) = \theta_0 + \theta_1 \sin(\omega_c t + \phi_\theta) + \Omega t.$$

If the coordinate axes are chosen so that $\theta_0 = 0$ and the radial motion is defined by Equation 4.11, then the constants can be easily determined. The specific motion is found to be

$$\theta(t) = \frac{mE}{\gamma e B^2 r_0} \sin \omega_c t - \frac{E}{\gamma^2 B r_0} t. \quad 4.12$$

Equations 4.11 and 4.12 define the motion of a particle in the transverse plane as it flies along the solenoidal field with momentum $\gamma m \beta c$. These parametric equations, illustrated in Figure 4.4, describe a cycloidal path slowly rotating around the beam axis, which is closely related to the Cartesian equivalent of perpendicular uniform electric and




Figure 4.4. Motion of a particle in a solenoidal magnetic field and radial electric field. This motion is found to be nonrelativistic and does not affect movement along the longitudinal solenoidal field. The longitudinal motion, however, increases the relativistic mass, which in turn influences the cycloidal radius.

magnetic fields (indeed, that connection is enforced when $r \approx r_0$ is assumed; movement along \hat{r} is ignored, so the divergence of the electric field is disregarded). This motion is in addition to the particle's motion along the magnetic field lines. Assuming the cycloidal motion is nonrelativistic, the amount of forward momentum that is converted to transverse momentum is negligible (that is, $p_{\perp} \ll p_z$), so the forward momentum remains essentially unaffected.

4.2.2 Combining electrons and antiprotons

The previous analysis included the forces that the self-generated electric and magnetic fields that the electron beam produces and that the individual electrons witness. However, during the act of beam-beam compensation, an antiproton bunch also passes through the electron beam, producing its own radial electric field and azimuthal magnetic

field. These forces are not included in the list given by Equation 4.9, yet they have the potential to alter the equations of motion.

Using the formalism of the last section, it is relatively straightforward to append these additional forces. In Section 2.4, the electric and magnetic fields due to the Gaussian bunches are found to be:

$$\begin{aligned}\vec{E}_{\bar{p}} &= \frac{\lambda_{\bar{p}}}{2\pi\epsilon_0 r} (1 - e^{-r^2/2\sigma_r^2}) \hat{r} \quad \text{and} \\ \vec{B}_{\bar{p}} &= \frac{\mu_0 c \lambda_{\bar{p}}}{2\pi r} (1 - e^{-r^2/2\sigma_r^2}) \hat{\theta}.\end{aligned}\tag{4.13}$$

In this case, $E_{\bar{p}} = cB_{\bar{p}}$. The force due to the electric field and the magnetic field acts on antiprotons in opposite directions (and as the previous discussion indicated, they differ in magnitude by merely $1/\gamma_{\bar{p}}^2$). Since the electron beam flows in the opposite direction, the magnetic force adds to the electric force.

The equation of motion is now

$$\begin{aligned}\dot{\vec{p}}_e &= e(\vec{E}_e + \vec{E}_{\bar{p}} + \vec{v} \times (\vec{B}_e + \vec{B}_{\bar{p}} + \vec{B}_{\text{solenoid}})) \\ &= e((1 - \beta_e^2)\vec{E}_e + (1 + \beta_e)\vec{E}_{\bar{p}} + \vec{v} \times B_{\text{solenoid}}).\end{aligned}$$

By defining a new “effective” electric field $E_{\text{eff}} \equiv (1 - \beta_e^2)E_e + (1 + \beta_e)E_{\bar{p}}$, the equations of motion for electrons become

$$\begin{aligned}r_e(t) &= r_0 + \frac{\gamma_e m_e E_{\text{eff}}}{e B^2} (1 - \cos \omega_c t) \\ \theta_e(t) &= \frac{\gamma_e m_e E_{\text{eff}}}{e B^2 r_0} \sin \omega_c t - \frac{E_{\text{eff}}}{B r_e} t\end{aligned}\tag{4.14}$$

given the same approximations as in the previous section.

The point of this exercise is to confirm that the electron beam passes through the TEL beam pipe without being adversely affected by the numerous forces that the individual electrons experience. Quantifying the maximum amount of movement under worst-case conditions is therefore important. For a 2-A beam with a radius of 1.6 mm and moving at a speed of $0.2c$, the peak electric field is $E_e(r_e) = 500 \text{ kV/m}$.

The peak electric field due to the antiproton bunch is more obscure, as it follows the relation shown in Equation 4.13. While solvable, the peak electric field is very brief due to the Gaussian nature of the linear antiproton charge density. Since the equations of motion are linear with the electric field, it makes more sense to compute the peak field averaged over the length of the bunch; the electrons at the end of the interaction length will have responded the same amount as if the antiproton distribution were evenly distributed over the entire bunch length. For the typical antiproton-bunch parameters listed in Table 2.2, the average electric field is 144 kV/m .

Sifting this together yields an effective electric field of 650 kV/m . The main solenoid, discussed more extensively in Section 5.1, is typically operated at $B_{\text{solenoid}} = 3.5 \text{ T}$. In this field, the electrons have an angular frequency of $\omega = 6.03 \times 10^{11} \text{ rad/sec}$, implying that they make nearly 3300 full revolutions along the length of the TEL. However, the cyclotron radius is $0.31 \mu\text{m}$, which is much smaller than r_e . This substantiates the remainder of the derivation, so that the electrons follow Equation 4.14. The precession frequency $\Omega = E_{\text{eff}}/Br_e$ is high enough to allow the electron beam to twist over one radian through the two-meter length.

4.2.3 Gun and collector solenoids

The behavior of the electrons inside the main solenoid has been described extensively. The gun and collector solenoids are independently adjustable, which allows them to run at field strengths significantly lower than that of the main solenoid. By doing this, the electron beam in the gun gets compressed when it enters the main solenoid, which enhances the space-charge density.

Section 4.1 derives the maximum density of electrons that a thermionic cathode can deliver; in general, typical oxide-coated cathode can produce a maximum current density of about 10 A/cm^2 [13]. For linear tuneshift compensation, Section 3.2 indicates a beam with current 2.0 A and radius 1.6 mm suffices. Unfortunately, the corresponding current density is 25 A/cm^2 , which is significantly larger than oxide cathodes can reliably deliver.

Hence, the use of a weaker gun solenoid is vital in order to provide the current density needed for full beam-beam compensation. The gun solenoid is typically operated at approximately 3.8 kG and is situated at a right angle to the main solenoid. The lower field is advantageous because the magnetic field lines that pass through both solenoids are close together in the main solenoid but more spread out in the gun solenoid.

The electrons follow the cyclotron orbits defined in Equation 4.14, which means that, other than the azimuthal precession around the beam, each electron circles a specific field line. By means of magnetic compression, the density of the electron beam at the cathode can be decreased while maintaining the strong space-charge density inside the main solenoid. As long as the field lines themselves pass through both magnets and the electrons' energy is high enough to overcome the potential energy of compression, the beam will continue from one solenoid to the next.

The density of field lines is just a description of the strength of the solenoid. Therefore, the area of the beam scales inversely with the magnetic field strength:

$$\frac{B_{\text{main}}}{B_{\text{gun}}} = \frac{A_{\text{gun}}}{A_{\text{main}}} = \frac{r_{\text{gun}}^2}{r_{\text{main}}^2} . \quad 4.15$$

The gun solenoid, in practice, operates at about 3.8 kG, which increases the electron-beam radius at the cathode to 4.9 mm and decreases the current density to 3.4 A/cm², well below the limit of oxide cathodes. In addition, the weaker field allows the gun solenoid to be constructed from copper wires, eliminating the complications involved in superconducting magnets.

The electron beam is collected after being bent back out of the Tevatron beam pipe. Another solenoid is needed at this end to steer the field lines, and thus the beam, into the collector. Section 5.3 discusses the collector more completely, but a normal-conducting solenoid almost identical to the gun solenoid is used.

It is important to reanalyze the electron-beam behavior in the gun and collector solenoids to ensure their feasibility. In this situation, Equation 4.14 is still valid, and in fact only the self-induced radial electric field contributes to the effective electric field E_{eff} . No antiproton bunches are nearby to exert electric or magnetic forces. The maximum radial electric field inside the gun solenoid is approximately 160 kV/m at $r = r_{\text{e}}^{\text{gun}}$. With this electric field and a solenoidal strength of 3.8 kG, the cyclotron frequency is 6.6×10^{10} rad/sec and the orbit radius is about 6.5 μm . The electrons spiral many times within the gun solenoid, but the radius is less than 1 % of the beam radius, so the overall beam shape is unaffected.

The bends themselves are more difficult, for several reasons. First, the field strength depends on the fringe fields of the two solenoids, which is much more difficult to calculate

than inside a solenoid with known parameters. Second, the field is not cylindrically symmetric, making analytic computations of the field lines much more challenging. Also, the fringe fields of a solenoid depend strongly on the amount and position of any ferrous material near the exits of the solenoids. Each of the solenoids use ferrous material to help confine and smooth the magnetic field, so a small change in their parameters (for example, the ferrous material is slightly thicker or its permeability μ is slightly different) can significantly impact the trace of the field lines. Related to this, any ferromagnetic material in the vicinity of the bends will change the trajectories significantly, and may even spell disaster for the electron beam.

Simulations of the bends were performed to confirm that the field lines between the two solenoids connect and that the electron beam can successfully pass from one magnet to the next [80,82]. Figure 4.5 illustrates the results of one such simulation. The region of the bend between the main solenoid and the collector solenoid is shown. The horizontal dotted line represents the axis of the main solenoid, and the vertical dotted line represents the axis of the collector solenoid. The small vertical line illustrates where the main solenoid ends; its length extends to the left.

The three black curves trace three field lines: one in the center of the main solenoid, and the other two are at a horizontal displacement of ± 5 mm, which purposefully maps out a width larger than the electron beam. These three lines spread out and then recombine as they enter the collector solenoid. Without any adjustments of the magnetic field, the lines enter the collector solenoid displaced about 4 cm from the axis of the solenoid. Section 5.1 describes the correcting dipole magnets within the main solenoid in greater detail, but these magnets have the ability to offset the electron beam through

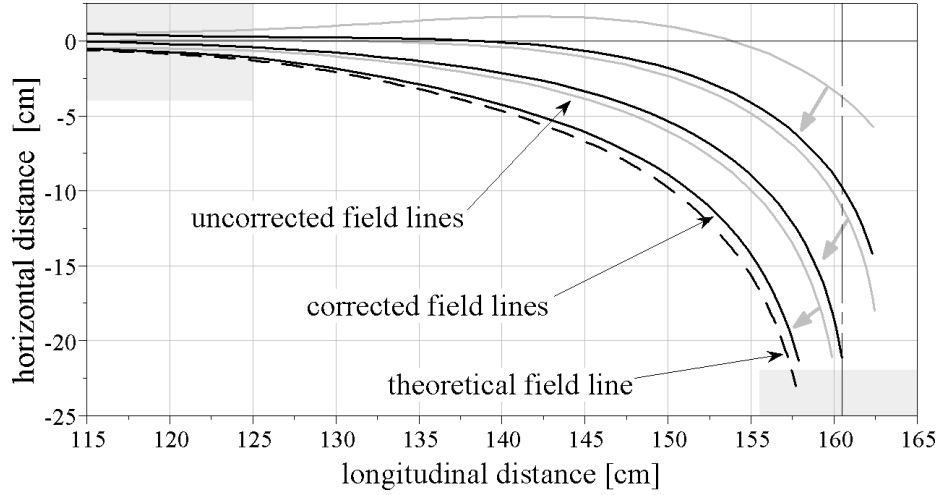


Figure 4.5. Field lines bending from the main solenoid (upper-left shaded region) to the collector solenoid (lower-right region with axis). During these tests, the main solenoid was set to 65 kG and the collector solenoid was 4 kG. The shaded lines miss the axis of the collector, but 4 kG in the downstream corrector brings these lines to the black, centered ones. A theoretical curve (dashed line) is also included.

the length of the main solenoid or upon the beam's exit. By adjusting these correcting magnets, the offset can be eliminated. The three shaded lines in the right drawing assume 0.4 T is set in the last correcting magnet. The use of this small perturbation successfully brings the field lines to the center of the collector solenoid.

In the right drawing of Figure 4.5, the dashed line is a fourth line starting at a horizontal position of -5.1 mm but without the final corrector's help. This field line reaches the axis of the collector solenoid, indicating that a slight displacement in the main solenoid results in a strong displacement in the collector. It was anticipated in Section 3.3 that the electron beam would need significant steering in order to encompass the antiproton bunches. Figure 4.5 verifies that by adjusting the final correctors, the electron beam can easily be steered into the collector.

4.3 Longitudinal Dynamics

The magnetic subsystem has specific requirements outlined in the preceding section. However, the electron beam also depends on several electrical premises as well. The two most important voltages affecting the performance of the TEL is the cathode-to-ground voltage V_{cat} and the anode-to-cathode voltage V_{anode} . The former establishes the energy of the electron beam traveling through the TEL, and the latter is pulsed to produce the desired amount of current from the cathode. The current pulse or waveform that is shown in Figure 3.3 is produced by a similar voltage waveform applied to the anode. In general, tunes shift compensation desires a specific maximum current, so the electrical requirements are steered toward providing that level.

There are two constraints on the cathode voltage. The first is that the beam needs significant energy to overcome its space-charge potential; the second is that the beam needs to move fast enough to “clear out” of the Tevatron beam pipe between bunches. The anode voltage needs to be large enough to generate the current required for tunes shift compensation, and it should be adjustable for each bunch in the train repetitively.

4.3.1 Overcoming the space-charge potential

It is tempting to think that the anode accelerates the electrons from the cathode, and, having gained this kinetic energy, they happily travel down the pipe without regard to the beam pipe. However, the same electric field that accelerates the electrons from the cathode toward the anode also pulls back on the electrons, possibly even *decelerating* them, when the electrons exit the anode region.




Figure 4.6. Sketch of the potential energy that an electron beam passes through. Electrons, being negatively charged, see a negative voltage as positive potential energy, so a region of higher energy corresponds to a more negative voltage.

Energy is conserved, and the electrons' initial potential energy while on the cathode is the kinetic energy that they have on their journey through the TEL. The anode voltage generates the current pulse, but it is the voltage on the cathode with respect to ground that defines the kinetic energy of the beam. In Figure 4.6, a cartoon describes the potential energy of the different regions through which the beam passes. The initial hurdle of overcoming the work function at the cathode surface is described in Section 4.2, but afterward the beam gains significant kinetic energy while traveling through the anode tube. After it leaves the anode, it passes along the grounded beam pipe, and its kinetic energy is defined by the cathode's voltage.

An undiscussed complication is the space-charge potential that particles within the electron beam sense. Bringing the electrons together (first at the cathode and again near the main solenoid) requires work; the kinetic energy of the individual particles is reduced to provide this potential energy.

A naïve but approximately correct approach to solving for the space-charge potential $V_{sc}(r)$ within the beam is to derive it from the electric field inside and outside the

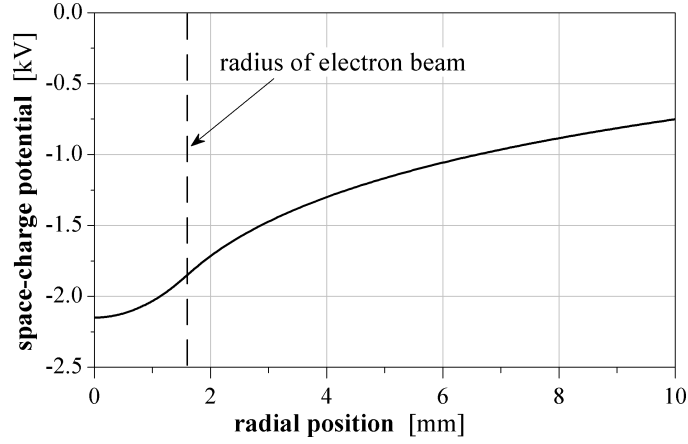


Figure 4.7. Plot of Equation 3.14 assuming $I_e = 2$ A, $\beta_e = 0.2$, $r_{\text{pipe}} = 35$ mm, and $r_e = 1.6$ mm. At $r = r_{\text{pipe}}$, the potential equals zero, and creates a negative well inside the pipe. This function assumes that the particles are traveling at the same speed.

beam as expressed in Equation 4.13. Solving for $V_{\text{s-c}}(r) = -\int_r^{r_{\text{pipe}}} E(r)dr$ in this case, where r_{pipe} is the radius of the beam pipe, yields the potential within the beam. Setting $V_{\text{s-c}}(r_{\text{pipe}}) = 0$ gives

$$V_{\text{s-c}}(r) = \begin{cases} -\frac{I}{4\pi\epsilon_0\beta_e c} \left[1 + 2 \ln \frac{r_{\text{pipe}}}{r_e} - \left(\frac{r}{r_e} \right)^2 \right] & \text{for } 0 \leq r \leq r_e; \\ -\frac{I}{2\pi\epsilon_0\beta_e c} \ln \frac{r_{\text{pipe}}}{r} & \text{for } r_e < r \leq r_{\text{pipe}}. \end{cases} \quad 4.16$$

This solution indicates a quadratic relation between the potential and radial position within the beam and a logarithmic relation outside. At the grounded pipe wall, the potential is zero. This curve is plotted in Figure 4.7 using typical TEL parameters, including a relativistic beta β_e of 0.2, which corresponds to 10 keV of kinetic energy.

Several important conclusions can be drawn from Equation 4.16. The space-charge potential for the entire beam is a significant portion of the total kinetic energy. For

example, in Figure 4.7, the beam's potential energy is 1.85 keV at $r = r_e$ and 2.15 keV at the center. Both of these values are a significant portion of the beam's total energy, 10 keV, which implies that the actual kinetic energy, and therefore the speed, is lower than originally assumed.

Second, the center of the beam requires 300 eV more potential energy than the edge, so the center of the electron beam is moving somewhat slower than the edges. This creates a larger charge density in the center, affecting the linearity of the electric fields. It also means that the center of the beam requires a larger cathode voltage to propagate than the edge; if the cathode voltage is set too low, the outside of the beam will pass through the TEL with less in the center. Beam profile measurements at low cathode voltages are displayed in Section 5.2. These profiles indeed have decreased charge densities in the center.

Assuming that a uniform current density leaves the cathode and all of the electrons are able to travel the length of the TEL, the current density will remain constant throughout the beam. If the central portion of the beam is traveling slower, however, then the charge density will be enhanced, requiring more potential energy for the central electrons. This increase in the charge density is not accounted for in the simple analysis of Equation 4.16. Instead, the difference in potential between the center and the edge of the electron beam is expected to be larger than predicted by Equation 4.16. At low or medium currents, the error is negligible, but if the current reaches a level such that the space-charge potential is a large fraction of the cathode voltage, the discrepancy becomes severe.

4.3.2 Beam-pipe acceptance

The circular relationship between the charge density and the potential can be analyzed more rigorously through Poisson's equation, $\nabla^2 V_{s-c}(r) = -\rho(r)/\epsilon_0$. The radially dependent charge density is $\rho(r) = J/v(r)$, where the current density J is independent of radius and the particle velocity $v(r)$ is tied to the kinetic energy, $eV_{\text{cat}} - eV_{s-c}(r)$. Using cylindrical symmetry, Poisson's equation is reduced to the one-dimensional differential equation:

$$\frac{1}{r} \frac{d}{dr} \left(r \frac{dV_{s-c}(r)}{dr} \right) = -\frac{J}{\epsilon_0} \sqrt{\frac{\gamma m}{2e(V_{\text{cat}} - V_{s-c}(r))}}. \quad 4.17$$

This differential equation cannot be solved analytically, but a limit can be found on the amount of sustainable current before the space-charge potential of the beam center is too large to allow the current to continue[85,86]. This limit is defined by the maximum potential at the beam center, $V_{s-c}^{\text{max}}(r=0)$, and is equal to the cathode voltage V_{cat} in the simplified analysis of Equation 4.16. Any realistic beam radius is less than the beam-pipe radius, for which a converging series solution is found to be

$$V_{s-c}^{\text{max}}(r=0) = \frac{2}{3} \left(1 + \frac{1 + 6 \ln r_{\text{pipe}}/r_e}{4(1 + 2 \ln r_{\text{pipe}}/r_e)^2} + \dots \right) V_{\text{cat}}. \quad 4.18$$

The truncation of the series to the first two terms yields an error of less than 5% when $r_e \rightarrow r_{\text{pipe}}$ and less than 0.02% for the TEL's beam ratio, $r_e/r_{\text{pipe}} = 0.0457$. This function is plotted in Figure 4.8.

As expected, the energy provided by the cathode cannot be completely converted into space-charge energy as Equation 4.16 hoped. For a large beam that fills the entire pipe, $r_e \rightarrow r_{\text{pipe}}$ and $V_{s-c}^{\text{max}}(r=0)$ approaches $0.833 V_{\text{cat}}$. Stating this slightly differently,

immaterial to the understanding of the T_{EL}.

Equation 4.19 indicates a complicated relationship between the cathode voltage V_{cat} and the beam current I_e . However, for a specific beam radius, Equation 4.18 provides a fixed value for the ratio $V_{\text{s-c}}^{\text{max}}(0)/V_{\text{cat}}$ for the maximum current. Inserting Equation 4.18 into Equation 4.19 gives a simple power relationship between the maximum allowed current and the cathode voltage for a specific beam radius. In Figure 4.10, three curves are shown, each with a slightly different T_{EL} beam radius. For a radius of 1.6 mm, the maximum current that the pipe will accept is given by the simple equation,

$$I_e^{\text{max}} = (3.75 \mu\text{P}) V_{\text{cat}}^{3/2}, \quad 4.21$$

where the derived numerical constant $3.75 \mu\text{P}$ is known as the acceptance of the beam pipe and has the same units as the perveance defined in Section 4.2; that is, $1 \mu\text{P} \equiv 10^{-6} \text{ A/V}^{3/2}$. The interpretation is essentially the same, too: given a particular voltage potential of the beam, a specific amount of current is allowed to flow.

If the electron-beam radius is increased to 2.0 mm, the acceptance rises to $4.01 \mu\text{P}$. However, the current density decreases to 68 % of the original beam, severely limiting the radial electric field and thus the tunes shift. Reducing the beam size lowers the acceptance to $3.45 \mu\text{P}$, but the current density is 63 % larger than that of the 1.6-mm-radius beam. As Section 3.2 describes, a small-radius electron beam has more difficulty surrounding an entire antiproton bunch, though such a beam will have a stronger effect on the bunch due to its increased space-charge field.




Figure 4.11. Relationship between the speed of the beam center and the total current pushed through the beam pipe, assuming $V_{\text{cat}} = 10$ kV. With negligible current, both the center and edge of the beam move at nearly $0.2c$. As the current is increased to the maximum, the center slows to less than $0.12c$.

For beams of radius 1.2 mm, 1.6 mm, and 2.0 mm, this relation is also plotted in Figure 4.11. Each curve indicates a higher speed than the corresponding curve for the beam center. Also, as the current decreases, the difference between center and edge disappears. At the maximum current for each beam size, the edge of the beam travels at about $0.14c$.

Figure 4.11 covers the entire range of allowable current for T_{FL}-like beams, but the electron guns employed by the T_{FL} are only able to operate between zero and roughly two amperes, and most of their operation is less than 1.5 A. In this range, the curves in Figure 4.11 are fairly linear. For comparisons between theory and experiment in

Chapter 6, it is reasonable to approximate the central electrons' velocities as:

$$\begin{aligned}
r_e = 1.2 \text{ mm} : \beta_e(0) &\approx 0.1960 - 0.0133I_e ; \\
r_e = 1.6 \text{ mm} : \beta_e(0) &\approx 0.1960 - 0.0122I_e ; \\
r_e = 2.0 \text{ mm} : \beta_e(0) &\approx 0.1960 - 0.0113I_e .
\end{aligned}
\tag{4.25}$$

The errors between these linear approximations and all of the numerical values are less than 0.3 %.

The relations reached in Equation 4.25 are important for two reasons. First, the tuneshift as expressed in Equation 4.10 depends on the electron velocity; the above results provide a more detailed description of β_e than suggested in Section 3.2. Second, the T_{EL} gun is required to adjust the current between antiproton bunches. This means that after a new level of current begins to flow, the electron speed must be high enough to pass down the entire interaction length before the next antiproton bunch arrives. If the current does not travel fast enough, the “new” bunch will feel a tuneshift corresponding partially to the “old” current level, causing unnecessary controls complications. Figure 3.3 sketches the idea of giving the beam enough time to travel down the main solenoid.

Section 3.2 calculates the necessary time for the electron beam to pass through the entire interaction region. As a “worst-case scenario,” an electron speed corresponding to $\beta_e = 0.1$ was shown to require 84 nsec to travel the entire length of the T_{EL}. As Figure 4.11 suggests, the T_{EL} is realistically producing electrons moving significantly faster than this, so that an antiproton bunch will only witness a level of electron-beam current intended for that one bunch. In this manner, independent control of each bunch's tuneshift is ensured.

4.3.4 Electron charge density

The preceding section analyzed the decrease in electron speed as the current is increased. This decrease in speed enhances the otherwise linear connection between the charge density and current, because $\rho_e = I/(\pi r^2 \beta_e c)$. By substituting this relation into Equation 4.10, the tunes shift becomes:

$$\Delta\nu_x = \frac{R_p \beta_x \rho_e L_{\text{TEL}}}{4\gamma_{\text{p}}} \cdot (1 + \beta_e) \quad 4.26$$

in the horizontal plane and similarly in the vertical plane. This equation indicates an almost linear relationship between tunes shift and charge density. However, the $(1 + \beta_e)$ term folds in a small nonlinearity, which can be explored by substituting Equation 4.25 into Equation 4.26: at $I_e = 1$ A, the difference in tunes shift between assuming a fixed β_e of 0.196 and the calculated value is 1 %. At $I_e = 3$ A, this difference increases to 3.8 %. These discrepancies decrease as the beam is enlarged, since β_e is less mutable.

Modeling the charge density is important, since the fields scale with the density and it can be measured experimentally through the use of BPM signals. It is thus a useful stepping stone in comparing theory to experiment. Inserting the relations plotted in Figure 4.11 into Equation 4.26 yields the connections between charge density and current for differently sized beams; this is plotted in Figure 4.12. For each beam radius, the charge density scales fairly linearly with low currents. However, as the current pushes toward its upper limit, the density increases more rapidly. This is anticipated by the decrease in electron speed at higher currents.

A significant difference in charge density can be seen between electron beams of various radii. The 1.2 mm beam must pack the same number of electrons into a much




Figure 4.12. Charge density at the center of an electron beam. The density bends upward due to the slowing of the electrons immersed in the large current. A decrease in the beam radius drastically increases the density, but it also increases the dependence on accurate alignment.

smaller area than the 2.0 mm beam with the same current. The charge density, therefore, is much higher. The slope of the plotted functions at low currents are proportional to $1/r_e^2$.

It is important to note that Figure 4.12 follows the center of the electron beam. The charge density around the perimeter of the beam can also be easily calculated, and is somewhat lower (and more linear) than the central portion. The BPM plates detect the total linear charge density, and integrating ρ_e over the cross-sectional area of the beam would yield this total. However, this measurement is not critical to the tunes shift. The large majority of antiprotons have small beta functions, so they only interact with a small region of the electron beam around the center. Large-amplitude particles will witness the density outside of the center, though even they also often pass through the center of the electron beam. As a result, the charge density at the edge of the beam does not play a large role in shifting antiproton tunes. Hence, the comparison between

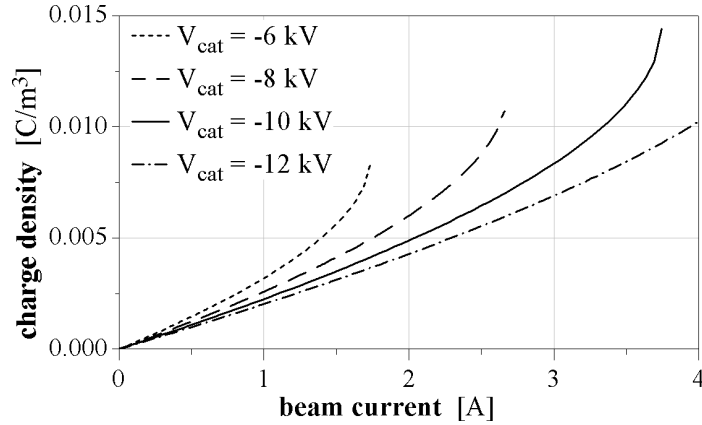


Figure 4.13. Dependence of charge density in the center of the electron beam on cathode voltage. By setting the cathode voltage larger (more negative), the electrons have larger total energy and are not as influenced by the space-charge potential. A beam with only 6 keV of total energy, for example, cannot sustain much more than one ampere.

experiment and theory in Chapters 6 and 7 uses the center of the electron beam for its analysis.

The previous discussion has incorporated different beam sizes in its theory. Figure 4.12 proves that a smaller electron beam is an effective way to provide more tunes shift for the same current. However, as Section 3.2 points out, a small beam has more difficulty surrounding an entire antiproton bunch, and it also is more susceptible to misalignment. Another method for adjusting the tunes shift capabilities of the T_{EL} is to change the cathode voltage. As the voltage is decreased in amplitude (toward zero, since the voltage is negative). Adjusting V_{cat} in Equation 4.24 provides an understanding of how the charge density can be altered. In Figure 4.13, the cathode voltage ranges between -6 kV and -12 kV.

A significant difference can be seen between different voltage levels. This is directly related to the amount of kinetic energy the electrons begin with, as Figure 4.6 illustrates. Without much space charge, the velocity follows $\beta_e = \sqrt{2eV_{\text{cat}}/\gamma_e m_e c^2}$, so the density at low currents is expected to have different slopes proportional to $V_{\text{cat}}^{-1/2}$. Higher currents, as usual, require more energy to be earmarked for the beam's space-charge potential, which slows down the electrons (seen in Equation 4.23) and increases their density. As Figure 4.13 indicates, the divergence from a linear relationship occurs at significantly different currents depending on the cathode voltage. For example, at $V_{\text{cat}} = -6 \text{ kV}$, the beam is unable to sustain even two amperes. However, at $V_{\text{cat}} = -12 \text{ kV}$, the density is moderately linear above four amperes. The cathode voltage sets the total energy available; when an abundance of energy is available, the limit is not sensed. When it is in short supply, the space-charge potential easily consumes a large portion and leaves very little left for motion, a commonly observed phenomenon[73,56,57].

Chapter 5:

The Tevatron Electron Lens apparatus

*All my study be to no effect?
You are deceived: for what I mean to do
See here in bloody lines I have set down;
And what is written shall be executed.
– Shakespeare: Titus Andronicus*

So far, a need for beam-beam compensation has been described, and an analytical description of the necessary beam parameters has been developed. A detailed look at the physical machinery is important, as it elucidates numerous features and issues. Chapter 1 outlined the main subsystems, which are expounded in this chapter. Table 5.1 lists many of the different TEL parameters, including a nominal setting and an adjustable range in order to conduct experiments.

The TEL can be divided into three main subsystems. The magnets comprise the most important one for the success of the TEL, as they are needed to guide the electron

T_{EL} parameter	nominal	range	units
electron beam energy	10	5–12	keV
peak beam current	1	0–3.5	A
main solenoid field	35	25–40	kG
gun/collector solenoid field	3.7	3.0–4.2	kG
cathode radius	5		mm
beam radius in main solenoid	1.6	1.1–2.0	mm
electron density	3.3	2.1–6.9	$10^{10}/\text{cm}^3$
interaction length	2.05		m
horizontal beta function	101.67		m
vertical beta function	30.89		m
anode modulator rise time	>200		nsec
modulator pulse width	~300		nsec
current stability	0.1		%

Table 5.1. Table of a few quantities associated with the T_{EL} as it is typically run for tuneshift compensation. In anticipation of bunch-by-bunch compensation, the anode modulator can produce staircase waveforms as sketched in Figure 3.3.

beam along the orbit of the antiprotons. Their fields must be steerable yet remain straight to ensure that the antiproton bunch and the electron beam are collinear through the whole interaction range. Additionally, the magnets must turn the electrons into the collector regardless of how they were steered through the interaction region.

Two electron guns were developed to appease two different goals. The first was designed to create a flat charge density over a wide range, while the second generated a rounded, Gaussian-like distribution. Both were optimized to produce the largest amount of current as possible. A beam pipe and collector were also designed such that the beam could be diagnosed and collected without complications.

The third subsystem is the electrical circuit that furnishes current and power for the TEL. The largest amount of effort was devoted to the anode pulse modulator. Also, the TEL uses recirculating current, necessitating more complex circuitry between the gun and collector. Other details complicate the electrical layout of the TEL, discussed in Section 5.4.

As Section 3.2 mentions, a second TEL is currently being fabricated. Due to the success of the design of the first TEL, the second uses a very similar design, adding very few significant changes. This fact affirms the notion that the first TEL's design is quality.

5.1 Magnetic subsystem

The three solenoids in the TEL are oriented as shown in Figure 5.1. The gun solenoid sits in the lower-left corner perpendicular to the long Tevatron beam pipe, the main solenoid surrounds the beam pipe, and the collector solenoid resides in the lower right. Electrons, germinating from the electron gun, follow these field lines as shown in Section 4.2.

5.1.1 Solenoids

The solenoids were manufactured at IHEP in Protvino, Russia and tested at Fermilab. The main solenoid uses NbTi wire intertwined with copper wire, rated for 550 A at a temperature of 4.2 K; the wire itself measures 1.44 mm by 4.64 mm cross-section [81]. A 4.85 cm-thick, low-carbon steel shield wraps around the coils, which enhances the field




Figure 5.1. A drawing of the solenoid magnets in the TEL. The electrons are contained by the magnetic field throughout their journey from the electron gun, along the interaction length, and finally into the collector. By adjusting the relative strengths of the solenoids, the beam size, position, and curvature can be altered. This is a view from above the TEL.

strength, keeps the field lines compressed near the solenoid's ends, improves the homogeneity throughout the interaction region, and reduces stray fields. The specific design of the solenoid is outside the scope of this dissertation, but it involved careful simulations and sophisticated experience [81,82].

The coils of the main solenoid are immersed in a liquid helium bath, and the total weight of this *cold mass* is 1350 kg. Due to hysteresis effects and eddy currents in steel, a small amount of heat is generated whenever the current in the superconductor is changed. Quantifying the heat generated from a current ramping rate is difficult, but the manufacturers specify that 10 A/sec is a maximum rate at which the current is changed. In practice, the main solenoid is rarely powered up or down, but usual ramp rates lie under 4 A/sec. Quenches are avoided whenever possible.

While the nominal operating strength of the main solenoid is 3.5 T, it is designed to handle 6.5 T. However, superconducting magnets often require a “training period,” where they quench at lower strengths before attaining their rated capacity. During initial operation with the main solenoid, it successfully reached 6.6 T before quenching. The current was pushed to 6.7 T again, without quenching. After some time, it was decided to put it in operation without further conditioning, and it has rarely been pushed beyond 4.0 T. In the occasional times that other Tevatron magnets quench, the main solenoid will also sometimes quench. In addition, occasionally during operation the main solenoid has quenched on its own. However, the number of times this has happened is extremely small, and it is not considered a liability to the performance of the T_FL.

The field measured inside the main solenoid scales very linearly with current; less than one percent of variation was observed. The measured field strength was 36.7 G/A, which compares well with predictions[82]. The main solenoid does not contain a closed current loop; when energized, the current flows out of its current leads and through external power supplies. The following sections describe the advantage that this has for quench protection, as the solenoid can contain up to 1 MJ of energy when it reaches its maximum rating of 6.5 T. Also described are superconducting dipole correctors contained in the solenoid’s assembly.

The gun and collector solenoids use copper windings, produce much weaker fields, and require less support systems. Water-cooled, these magnets can support 340 A across their 0.19- Ω load. The measured inductance of the 391 turns of wire is roughly 18 mH. The bore of each magnet has a diameter of 24.0 cm and a length of 30.0 cm, enough to contain the electron gun and the entrance to the collector.

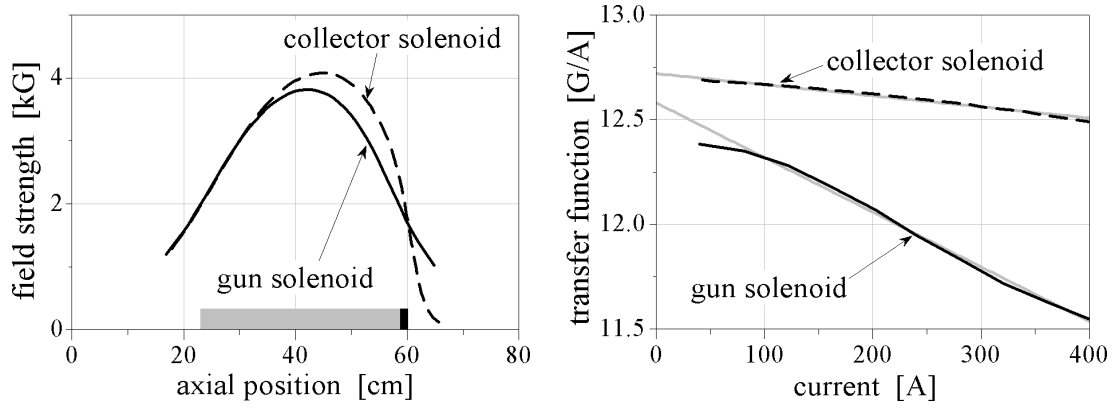


Figure 5.2. Field strengths produced by the gun and collector solenoids. The left side shows the strength through each solenoid with 325 A of current. The grey rectangle represents the length of the gun solenoid, and the black represents the additional iron plate of the collector solenoid. The right plot shows the measured field strength in black with fitted approximations in grey.

A small design difference between the gun and collector solenoids is useful. The collector solenoid has an additional iron plate on its back end, which deliberately reduces the field strength outside of the solenoid (in the region of the collector itself, as seen in Figure 5.1). Measurements of the field strength in each magnet is shown on the left side of Figure 5.2. The horizontal axis indicates distance from the center of the Tevatron beam pipe ($-x$), while the vertical shows the strength of the field at a current of 325 A in each magnet, close to the typical operating currents.

A grey bar at the bottom of Figure 5.2 extends the length of both magnets' coils. The iron plate added to the collector solenoid is indicated by the small black region. The effect of this plate is apparent: the strength of the field beyond it, inside the collector region, is significantly reduced from the equivalent position along the gun solenoid. This reduction of field strength is intended to encourage the electron beam to spread before colliding with the collector walls, an effect which is elaborated in Section 5.3.

The extra iron has three additional effects. First, the inclusion of more iron increases the maximum field generated in the collector solenoid for the same current level. This can be seen by the higher attainable fields in the collector-solenoid data than the gun-solenoid data. Second, the iron plate breaks the symmetry of the field.

The last effect of the iron plate is that it decreases the nonlinear aspect of the magnetic field. Saturation effects in the gun solenoid cause the field-to-current ratio, also known as the *transfer function*, to sag at high currents, a characteristic which is less observed in the collector solenoid. The right side of Figure 5.2 illustrates both the higher peak fields and the improved linearity of the collector solenoid over a broad range of currents. In this plot, the transfer functions B/I for both magnets are measured at different currents. As expected, both solenoids exhibit decreasing transfer functions, with the gun solenoid deteriorating faster than the collector solenoid. Lines were fitted to the data with weight given to the measurements around 325 A, since the corresponding 3.8–4.0 kG fields are typical settings during operation.

A linear fit corresponds to a quadratic fit between field strength and current. For the two solenoids, these fits are

$$B_{\text{gun}} = (12.58)I - (0.00261)I^2 \quad \text{and}$$

$$B_{\text{col}} = (12.72)I - (0.00053)I^2,$$

where B is measured in Gauss and I is in Amps. These polynomials describe the peak magnetic fields in both solenoids quite well over the range shown in Figure 5.2, and are very accurate over the relevant range of fields around 3.8 kG.

5.1.2 Dipole correctors

Proton and antiproton orbits are observed to drift transversely over time during a given store, and often the orbits are deliberately altered either during or between stores. Therefore, the T_{EL} needs the ability to orient the electron-beam path such that it coincides with the antiproton path. Without any bending magnets, the antiproton orbit is a straight line through the T_{EL}. At the *upstream* end of the T_{EL}, that is, the place where the electron beam first enters the Tevatron beam pipe and is intended to contact the antiproton orbit, the orbit has a horizontal and vertical position with respect to the center of the beam pipe. The orbit extends in a straight line through the T_{EL} at specific horizontal and vertical angles. Therefore the T_{EL} needs the ability to adjust to four degrees of freedom: the upstream position and the angle, both in the horizontal and vertical directions.

The T_{EL} design included four dipole corrector magnets in order to accomplish this steering. Two of these correctors, one oriented horizontally and one vertically, are located at the upstream end of the main solenoid; their goal is to adjust the upstream transverse position of the electron beam to equal that of the antiproton orbit. Figure 5.3 illustrates an example of how these dipole correctors work, though the transverse offsets in the illustration are exaggerated. In reality, the length of the main solenoid is two meters and offsets are only several millimeters.

The black line represents the antiproton orbit while the grey line is the electron beam. Uncorrected, the electron beam follows the straight, dashed line, which is not collinear with the antiproton orbit. The short upstream correctors are then activated to “push” the electron beam over to where the antiproton orbit is located. The upstream




Figure 5.3. Sketch of the placement and action of the dipole correctors, with an extremely exaggerated transverse scale. Without activating the correctors, the electron beam would follow the dashed grey path. By using them, the electron beam can overlay the antiproton orbit.

horizontal and vertical correctors therefore have the ability to account for the upstream position of the antiproton orbit. Section 5.3 describes beam-position monitors (BPMs), which are diagnostics intentionally located at the end of the upstream correctors. These BPMs are able to detect the transverse position of the antiprotons, protons, and electrons, and can verify that the upstream position of the electron beam equals that of the antiproton bunches.

Two other correctors extend nearly the length of the main solenoid. These *long correctors* have the ability to angle the electron beam along their entire length. Once the upstream correctors are set, the long correctors are adjusted so that the electron beam coincides with the antiproton orbit, as drawn in Figure 5.3. Another set of BPMs are situated at the downstream end of the long correctors in order to confirm that the two species finish at identical transverse positions.

The electron beam can end at a variety of positions, yet it must be able to pass into the collector. To accomplish this, a third set of correctors are located downstream

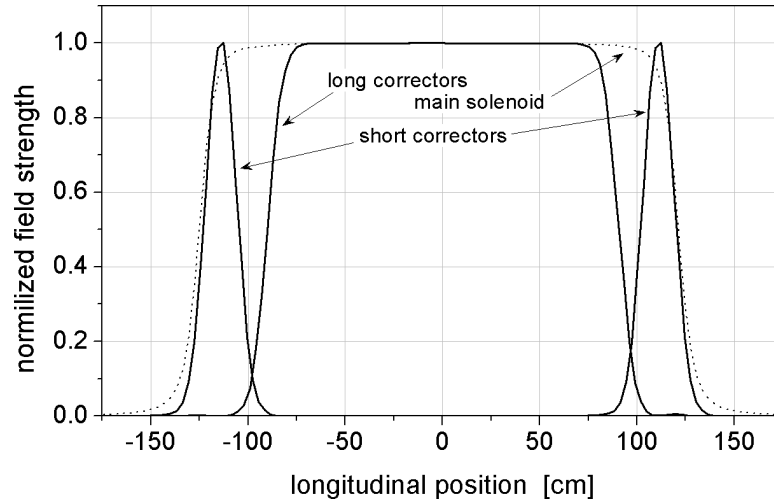


Figure 5.4. Longitudinal strength of the dipole correctors inside of the main solenoid. The solenoid field strength B_z is also included, and the longitudinal position is referenced to the center of the main solenoid.

of the long correctors in order to steer the beam back into a position where it will successfully enter the collector. These correctors, identical to the upstream correctors, often are adjusted simultaneously with either the upstream or the long correctors, but in the opposite direction; in this sense, they “undo” the changes made by the other correctors.

The specific location, and length, of each of these correctors is shown in Figure 5.4. The dashed line illustrates the strength of the main solenoid on axis as a function of longitudinal position. It is at a maximum nearly from -100 cm to $+100$ cm and rapidly falls to almost zero at -150 cm and $+150$ cm.

The solid lines in Figure 5.4 represent the measured strength of each set of dipole correctors. The upstream short correctors peaks around -115 cm, the long corrector extends from -75 cm to $+75$ cm, and the downstream short corrector appears at $+115$ cm.

Since the strength of each corrector and the main solenoid can be arbitrarily set, their magnitudes are all normalized to 1.0 in Figure 5.4. In the actual measurements, the solenoid was set to 6.5 T, the short correctors were 0.8 T, and the long was 0.2 T. No significant difference was observed between horizontal and vertical correctors.

Without the solenoidal field, the dipole correctors would act as dipole magnets were described in Section 2.2. There, each dipole bends the bunches. In this context, however, the electron beam spirals around the solenoidal field lines as derived in Section 4.2. The dipole correctors add a small perturbation to the nominally longitudinal solenoidal field. By superposition, the vector field of the correctors gets added to the vector field of the main solenoid. Since the former is a uniform field pointing transversely and the latter is a uniform field pointing longitudinally, the net result is a field that points at an angle represented by the sum of the two vectors. The electron beam, dutifully following the field lines, tracks the resultant field. Beyond the region of the corrector, the field lines and the electron beam again point longitudinally, but from this new position.

The strength of the correctors is listed in Table 5.2. However, a more useful measure is the amount that the beam gets shifted transversely, but this depends on the field of the corrector, the length of the corrector, and the field of the solenoid. More specifically, geometry of the two fields says that the total horizontal deflection Δx can be derived from

$$\Delta x = \int \frac{B_{\text{corrector}}(z)}{B_{\text{solenoid}}(z)} dz,$$

where the two field strengths are functions of the longitudinal position z , and the integral covers the pertinent length shown in Figure 5.4. A similar expression can be written for the vertical corrector, assuming the appropriate corrector field is used.

<i>corrector name</i>	<i>coil length</i> [cm]	<i>field strength</i> [G/A]	<i>displacement</i> [kG-mm/A]
upstream horizontal	27	39	5.88
upstream vertical	27	41	6.59
long horizontal	196	20	34.6
long vertical	196	20	37.7
downstream horizontal	27	39	5.75
downstream vertical	27	41	6.35

Table 5.2. Attributes of the dipole correctors. The field strengths were measured in Protvino, where the solenoid was constructed [81]. The displacement strengths were measured at Fermilab and agree with BPM observations.

In Table 5.2, the total deflection power of each corrector is given in units of kG-mm/A, which is $\Delta x \cdot B_{\text{main}}$ per unit current. Dividing these numbers by the main solenoid strength yields a valid transverse displacement for a known amount of current. Separate measurements of electron-beam deflection using BPM readings verified these calibration numbers [83].

5.1.3 Peripheral subsystems

The main solenoid and its correctors are superconducting. If the solenoid quenches at 6.5 Tesla, nearly 1 MJ of stored energy is released over a mere two seconds [84]. The current in each magnet loops through external power supplies, allowing external quench detection circuits and loads to absorb most of that energy. Simulations of quenches suggests that roughly 90 % of the total energy can be dissipated in external resistive loads, with the remaining 10 % being dissipated in the solenoid itself. In these simulations, the hottest point in the coil rises to about 270 K.

The dipole correctors only contain up to 1.3 kJ of energy, and dissipating this energy within the magnet is not worrisome. However, heat in one region could cause a quench in the main solenoid. Therefore, the correctors are also connected to quench protection circuits and loads.

Each monitor was originally designed to observe the voltage across its magnet and the time-derivative of the current, which were compared to an assigned limiting voltage:

$$\left| L_{\text{magnet}} \frac{dI(t)}{dt} - V(t) \right| < V_{\text{limit}} , \quad 5.x$$

where L_{magnet} is the inductance of the magnet. If the above condition ever became invalid, the magnet was assumed to have begun to quench, and crowbar resistors would be immediately bridged across the magnet to absorb its energy. The reason that Equation 5.x works is that an ideal inductor generates a voltage difference given by $L dI/dt$, so the left side of the equation ideally is zero. If the magnet quenches, the current suddenly decreases without changing voltage; the left-hand side increases dramatically.

However, the inductance L of a large (0.5 Henry) solenoid is typically not constant at low frequencies (1-10 Hz), due to iron-saturation effects and eddy currents. The overly simplified model expressed in Equation 5.x led to occasional false quench detections. A more sophisticated model using higher-order effects of both $V(t)$ and $I(t)$ was adopted and is now in use. The quench protection monitor tests the relation,

$$\left| L_{\text{magnet}} \left(\frac{dI}{dt} + \kappa_1 \frac{d^2 I}{dt^2} \right) - \left(V + \kappa_2 \frac{dV}{dt} \right) \right| < V_{\text{error}} .$$

The addition of κ_1 and κ_2 offers the ability to better mimic the physical behavior of the magnets over a range of frequencies. As expected, the number of false quench detections has decreased dramatically since this change.

The power supplies for each of the solenoids and correctors need to be able to sustain each magnet's full current. The main solenoid in normal operation requires a full kilo-Ampère; large cables send this current from the power supply to the solenoid itself. The short dipole correctors employ 200-A supplies, while the long correctors use 50-A supplies. Since the correctors might need to be energized in either direction, each supply is fed through a reversing-switch box. This box is able to swap the leads, effectively turning the unipolar supplies into bipolar supplies.

The current ramp rates for each of the superconducting magnets is limited, and all of the settings are done remotely through computer control. The reversing-switch circuits automatically handle ramping the current through zero and switching polarity properly. Scanning the electron beam transversely becomes feasible.

5.1.4 Straightness of field lines

This chapter so far has assumed that the main solenoid's magnetic field is extremely straight through its length. If there is any significant bending, then the electron beam, which religiously follows the field lines, would not interact properly with the antiproton bunches. Worse yet, the beam would impart nonlinear forces and kicks on the bunch, as Section 6.3 discusses.

If the electron beam has a diameter of about 3 mm inside the main solenoid, the centroid of the electron beam should not deviate from a line by significantly less than that. This provides a comfortable column within which the antiproton bunches may pass and be effectively focused. A rough estimate of this requirement is 10% of the electron




Figure 5.5. Drawing of the gimble and photograph of the cart. The gimble measured only an inch across and had very little mechanical resistance. The cart was pulled by a long track, and it rolled inside an aluminum pipe temporarily thrust into the solenoid bore.

beam diameter, or 0.3 mm. While the magnet was designed and built to be as uniform as possible, it was imperative to measure the magnetic field lines.

Observing the field lines requires some clever techniques. A small iron rod was centered in a non-magnetic gimbal and mounted on a small cart, as illustrated in Figure 5.5. The cart was dragged through the solenoid, and the solenoid magnetized the rod, which constantly attempted to align itself along the field lines (a magnetized ferromagnet feels a torque $\vec{\tau} = \vec{M} \times \vec{B}$ attempting to align it along the field lines[36]). A small mirror that was attached perpendicularly to the rod (actually surrounding the rod) reflected a laser beam from one end of the solenoid back down the same direction.

The resolution of this measurement can be calculated from the fact that the deviation of the field lines is simply the integral of the perpendicular magnetic field:

$$x_{\text{field line}} = \int_{x_0}^x \frac{B_{\perp}}{B_{\parallel}} dx'.$$

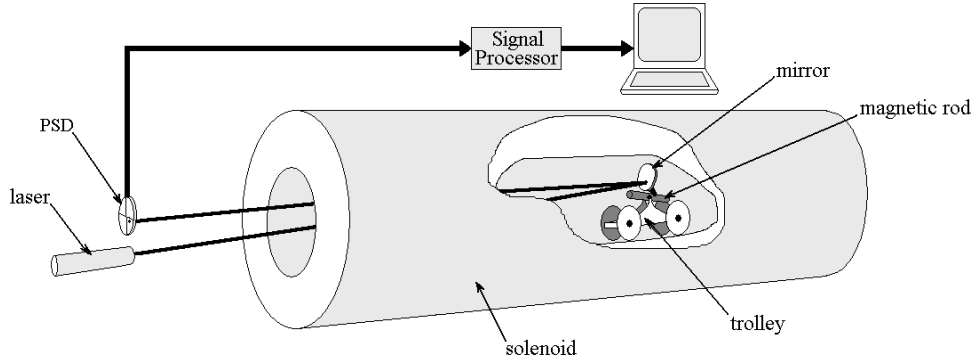


Figure 5.6. Simplified cartoon illustrating the technique used to measure the magnetic-field lines in the main solenoid. In reality, the laser beam reflected off of the mirror nearly on top of the incident beam, and a one-way mirror deflected the beam onto the light-position sensor.

Turning this formula around yields what angles of the rod are necessary to measure (these angles are very much within the small-angle approximation), namely,

$$\theta_{\text{rod}} \approx \frac{B_{\perp}}{B_{\parallel}} = \frac{\Delta x}{\Delta z}.$$

A maximum deflection from the bounds described earlier would be $\Delta x < 0.3 \text{ mm}$ over half of the solenoid length, or 1 m. This corresponds to angles of 3×10^{-4} , so the magnetic rod needs to be sensitive to angles less than that.

The reflected laser beam (returning at twice the angle of the magnetic rod) struck a two-dimensional light-position sensor. This large CCD target contains some processing electronics that reports where on its surface the laser light was incident. In this manner, minute angles could be observed, and Figure 5.6 depicts how the entire apparatus was operated.

A difficulty with sensitive magnetic field measurements is a systematic offset or tilt brought by the measuring device. In this case, the laser was positioned by hand such




Figure 5.7. Drawing of one horizontal and vertical field line at 4 T. A non-parallel laser beam generates an offset and angle artifact in the data, which has been normalized out of this plot. Each circle is a data point.

that its reflection successfully tracked on the CCD target, but there is no feasible manner to ensure that the laser beam was within microns of being parallel to the solenoid or its fields. The analyzation process therefore needed to include a removal of these arbitrary artifacts.

The field lines in the center of the solenoid are shown in Figure 5.7. The field does not bend more than $200\ \mu\text{m}$ in the horizontal direction and only about $45\ \mu\text{m}$ in the vertical. Therefore the electron beam is able to surround the antiproton bunches through the entire solenoid length. The $200\ \mu\text{m}$ variation is conveniently small and allows the option to experiment with beam-beam compensation at different electron-beam sizes.

Figure 5.8 shows how the straightness of the field lines change as the solenoid's field is ramped up or down. The deviation is rarely more than $20\ \mu\text{m}$, so changing the solenoid strength should not impair the ability of the TEL to work its magic on the antiproton tune (obviously, a different solenoid strength does shift the electron-beam path horizontally




Figure 5.8. Changes in central field lines with magnetic-field strength. Above 1 T, the magnetic field lines do not shift more than $5\text{ }\mu\text{m}$. At 1 T, the deviations reach $100\text{ }\mu\text{m}$. The 100 data points per trace are not drawn in order to improve clarity.

and will affect the charge density, but adjusting the correctors and the beam current (following Equation 5.1) negates these issues).

The cautious reader will note that dragging the cart down the axis of the solenoid is not the same as following a specific field line; Figures 5.7 and 5.8 assume that different field lines follow parallel trajectories, and therefore measuring the field angles for numerous field lines still accurately portrays the trajectory of one line. In order to test this hypothesis, the cart's track was offset by a small distance and data was again taken. In Figure 5.9, data was compared for five different transverse positions. Both horizontal and vertical field lines were adequately parallel to the central axis, so the solenoid would adequately carry the entire electron beam in a mostly straight path while preserving the beam diameter. If the solenoid showed significant deviation from its ideal model, calculating the expected tunes from the TEL would be much more difficult. Fortunately, the field generated by the solenoid is similar enough to a straight, ideal solenoidal field to avoid such complications.




Figure 5.9. Comparison of several “parallel” field lines near the central axis. The central field line and the constant offset (either 0.5 or 1.0 mm) was subtracted from each field line. The field lines are parallel to within $6\text{ }\mu\text{m}$ of the central line.

5.2 Electron gun

The physics behind electron guns is theoretically developed in Section 4.1. In that discussion, attention is spent on attaining high current densities while maintaining a desired current profile. In order to supply a beam with a flat density distribution, the size of the hole in the anode should be significantly smaller than the longitudinal distance between anode and cathode[36]. By scaling both dimensions equally, a large gun can produce more total current, but the anode requires pulses of higher voltages.

In the TEL, a compromise between gun size and anode voltage was reached; the anode is pulsed up to around 7 kV above the cathode potential, a number detailed in Section 5.4. Meanwhile, the gun’s emitting cathode has a radius of 5 mm, and the anode’s inner radius approximately 7 mm. Using Equation 4.5, a 1 cm cathode-to-anode

distance yields 1.4 A/cm^2 , far less than the potential yield that a more creative design could generate. In addition, the hole diameter is much larger than the cathode-to-anode distance creating a beam intense on its edges but weak or nonexistent in its center. This type of profile is not desirable.

Instead, the electron guns use more complicated geometries to preserve their desired profiles while producing as much current as possible for a given anode voltage, that is, a high perveance. In order to predict and optimize the guns' geometries, they were forged initially in simulation. The first gun was specifically designed to produce a wide, uniform density profile, and the second produces a tapered profile mimicking the proton-bunch distribution.

5.2.1 Mechanical design of TEL electron guns

The layout of the first gun is illustrated in Figure 5.10. It is called the *flattop gun* since it was designed to maintain a uniform current density across the width of the beam. In Figure 5.10, the horizontal axis is the central axis of the gun; the heated cathode is in the lower left and the anode in the upper right. Wrapping the diagram around the central axis, the anode appears more like a short tube than a plate. The electric field from the anode draws electrons off of the rounded surface of the cathode. A strong longitudinal magnetic field oriented horizontally forces these electrons to stream to the right from the cathode.

Section 4.1 models a planar gun as having a cathode plate and an anode plate, and the hole in the anode is small with respect to the distance between the plates. The previous discussion suggests that the TEL guns need to have a much larger hole in order




Figure 5.10. Cross-section of the flattop electron gun, the first electron gun designed and built for the T_{FL}. All three grey electrodes are symmetric around the longitudinal axis.

to provide suitable levels of current. However, the electric field from the anode will not be uniform against a planar cathode surface; the field is strongest around the edge of the cathode and drops near the center. Figure 5.10 exemplifies this problem, since the diameter of the hole is larger than the separation distance.

In order to counteract the nonuniform current that would result from the nonuniform field distribution, the cathode is spherically convex. This feature concentrates the electric field nearer the center, enhancing current flow from the center. This technique is common in high-current guns and can provide significantly flatter profiles than a simple flat cathode can.

One important addition to this gun was the use of a secondary electrode circling the cathode. This non-heated electrode can be imagined as a grid or control electrode. When




Figure 5.11. The current density profile of the electron beam. With the profiler set to the cathode voltage, electrons are easily emitted from the entire surface. The widest trace indicates that the current density was strongest at the edge. But applying a negative voltage to the profiler, the emitting area is effectively reduced.

not in use, it is set to the same potential as the cathode and simply extends the apparent surface of the cathode for the electric field lines, without actually giving off current, since it is not heated. However, if it is set to a voltage negative with respect to the cathode, it steals electric field lines from the outer edges of the cathode and effectively suppresses emission from that region. Figure 5.11 illustrates the effect of changing the voltage on this electrode, which is named the *profiler*.

If the profiler is set to zero volts with respect to the cathode, then the entire cathode/profiler surface senses electric field lines from the anode. For reasons explained in the next section, the field is not perfectly uniform; significantly stronger fields are felt

on the edges of the cathode, which translates into a higher current density than in the center. Nevertheless, the profiler is not suppressing these fields.

When the profiler is set to a voltage more negative than the cathode, the electric field lines are attracted to it, leaving less strength on the ring of the cathode. Even at -300 V, the profiler successfully reduces the current emission from the edges, while the current density in the center is nearly unchanged. Not only has the peak current density decreased dramatically, but the radius of the beam has shrunk approximately 1 mm.

As the profiler is set more negatively, the influence it has on the cathode's current gets stronger. An additional -300 V tends to reduce the beam radius by roughly 1 mm, which is the same as 0.33 mm difference in the beam radius inside the main solenoid. Even though the radius of the beam is significantly reduced, the density of the electron beam at the center remains relatively constant.

The total current can be easily measured as a function of profiler voltage. In Figure 5.12, the total current decreases significantly when the magnitude of the profiler voltage is increased. This confirms the notion that emissions from the edges of the cathode are suppressed by the influence of the profiler.

Many results in the following chapters discuss results using this flattop gun, which proved reliable for linear tuneshift compensation. Often it was run with -300 V applied to the profiler in order to limit the edges of the profile. However, it was observed that the steep sides of the profile provided a significant nonlinear focusing effect on large-amplitude particles, leading to significant losses. This phenomenon is analyzed in Section 7.2, but a second electron gun, with a more rounded profile, was also constructed. Adjustments




Figure 5.12. Total current expended by the cathode. It can be seen that the profiler can easily cut off some of the current, but must be set to a large (negative) voltage to suppress most of the current.

of the positions of its electrodes produced a profile that better imitated the profile of the proton bunches, since Section 3.1 argues that the ideal compensation scheme requires a radially dependent focusing force similar to the beam-beam effects at each IP.

Since proton bunches stabilize in a Gaussian distribution, the current profile of this second gun was designed to appear similar to a Gaussian curve[68]. This gun, known therefore as the *Gaussian gun*, is drawn in Figure 5.13, which includes simulations of the paths of electrons streaming from the surface. In comparison to the flattop-gun design, the most apparent changes are the profiler's position and the addition of a fourth electrode. The fourth electrode, called the *collar electrode*, has been electrically connected to the cathode, extending the surface similar to how the profiler behaved in the flattop gun. Not heated, the collar electrode does not produce any current, but smooths out the electric field lines that otherwise would concentrate on the edge of the cathode.




Figure 5.13. Cross-section of the more recent Gaussian gun. The profiler is shaped different, and a new *collar* electrode surrounds the cathode. Simulated paths of electrons are drawn in (black), along with equipotential lines (grey).

The profiler has been positioned much more prominently in the Gaussian gun than in the flattop gun. Its position now shields the edge of the cathode from the electric field, while the center of the cathode still feels the full field strength. In fact, the cathode is also curved more in the Gaussian gun, so that the electric field is further enhanced at the center of the cathode.

In all, the electric field at the center of the Gaussian gun's cathode is nearly as strong as it is in the flattop gun's cathode. The field strength near the edges, however, is significantly reduced. In Figure 5.13, field equipotential lines are drawn in grey. These lines are tightly clustered near the center of the cathode, whereas they are spaced apart near the rim of the cathode. Since the electric field is the gradient of the potential, the simulation agrees that the current profile will smoothly decrease at larger radii.




Figure 5.14. Current-density profile typical of the Gaussian electron gun (solid points) and the flattop gun (open points). The Gaussian gun puts out much less total current, but clearly it better mimics the profile of the proton bunches at collisions.

Measurements of the current profile confirm the simulations. Figure 2 display the radial profile of the Gaussian gun, and the profile, shown with solid points, has a rounded appearance much closer to the Gaussian distribution of the bunches in the Tevatron. As a reference, a measured profile of the flattop gun, under identical conditions, is shown with open points. It is immediately obvious that the edges of the beam are hugely suppressed in the Gaussian gun, providing a smooth distribution of space charge.

5.2.2 Gun electrical properties

The cathode itself is tungsten metal and impregnated with calcium oxide[13]. By heating the cathode in a vacuum, the oxide slowly leaks out of the surface, providing an atomic layer that lowers its effective work function. It has been found that the work function with the calcium oxide is about 1.1 eV, significantly lower than the 4.65 eV

work function of raw copper. The Richardson-Dushman law (Equation 4.3) predicts an exponential relationship between the maximum possible current and the work function, and indeed, before the oxide layer has been “activated,” very little current is observed. As the cathode is heated, that current begins to rise until finally the limit set by Richardson-Dushman surpasses the Child-Langmuir limit, and the gun then operates in the space-charge-limited regime. If the gun is exposed to air, the oxide layer is ruined (the tungsten will oxidize, creating a nasty layer with a high work function) and another processing is required.

Figure 5.15 illustrates the transition between space-charge and temperature-limited regions of operation. The resistive filament behind the cathode was set to different temperatures, and, after the cathode’s temperature re-equilibrated, the total current was measured. At high enough temperatures, Child-Langmuir fixed the current by the anode’s voltage and the gun perveance. At low temperatures, Richardson-Dushman limited the current following a $T^2 e^{-W/T}$ dependence. As the first gun was used extensively, it was noticed that the required temperature to reach the space-charge-limited regime increased somewhat. While this issue never limited its functionality, the most likely cause is the slow elimination of the oxide; as less oxide remains, it requires higher temperatures to replenish the supply.

Unfortunately, this solution circles back to the problem, since a higher temperature boils off the oxide layer faster [73]. Toward the end of experimentation with the flattop gun, it was feared that the cathode might be nearing its expiration date. After installation of the Gaussian gun, the filament’s temperature was more carefully regulated to barely more than necessary in order to prolong the cathode’s life as long as possible. No significant degradation has been seen.




Figure 5.15. Current production at a constant anode-cathode voltage difference for different cathode temperatures. At lower temperatures, the gun operates in the temperature-limited regime expressed by Equation 4.2, while above a certain temperature, the gun follows Equation 4.6 of the space-charge-limited regime. Because the temperature is very difficult to measure directly, the power dissipated in the cathode filament is shown (the strong Stephen-Boltzmann function makes an accurate calculation of temperature impossible). A bolometer check of the cathode revealed that the temperature is around 1000 K to 1100 K at the knee in this function.

One of the most important goals of these designs was to produce the largest current for a given anode voltage. This means the perveance of the guns was maximized while maintaining a decent current profile. The electrode shapes and positions shown in Figures 5.10 and 5.13 were adjusted in many ways before settling on the displayed schemes. The simulation code SuperSAM predicted a perveance of $4.7 \mu\text{pervs}$, and the results of measurements are shown in Figure 5.16. The data very well matches Child-Langmuir's relation, though the values for the perveance is somewhat different between experiment and theory.




Figure 5.16. Perveance of each of the guns. The Gaussian gun produces significantly less current than the flattop for the same anode voltage; this was a known side effect of reducing the edges of the beam so that the profile, shown in Figure 5.14, would be much smoother. Unfortunately the anode modulator used to pulse the Gaussian gun can not generate the same amplitude as the modulator used for the flattop gun.

5.3 Beam pipe and collector

As Figure 5.1 illustrates, the electron beam streaming out of the gun travels through a beam pipe until it finally reaches the collector. The beam pipe through which the electron beam travels preserves a constant inner diameter along nearly the entire length of the electron-beam path. This starts just downstream of the anode, continues around the first bend, through the main solenoid, around the second bend, and into the collector solenoid.

The inner pipe diameter in the main solenoid needed to be as large as the physical aperture of the Tevatron, as the the T_{EL} was not intended to ever inhibit the performance




Figure 5.17. Cross-section of the TEL beam pipe. A number of items need to be squeezed between the minimum radius of the Tevatron beam pipe and the maximum radius of the bore of the main solenoid's cryostat.

of the Tevatron. The dotted line in Figure 5.17 outlines the typical aperture found around the Tevatron ring. The circumscribed circle of radius 35 mm is the inner radius of the pipe through the length the TEL. The assurance that the pipe is larger than or equal to that of most of the Tevatron has the effect that, if the TEL is turned off, the Tevatron performs exactly as it did before the TEL was installed.

5.3.1 Beam-pipe limitations on the electron beam

A theoretical discussion of the space-charge potential is developed in Section 4.3. That analysis is crucial to the comprehension of measurements of beam current and charge density in the TEL, as it derives the acceptance of a beam pipe such as the TEL. In particular, a beam pipe of radius 35 mm supporting a 1.6-mm radius beam is calculated in Equation 4.21 to be $3.75 \mu\text{P}$. This value is dependent on the beam size and therefore the magnetic field strengths of the gun and main solenoids. The importance of the beam size is preserved in Equation 4.19 and demonstrated in Figure 4.10.

The acceptance sets the maximum current that the beam pipe will allow, given a set cathode voltage and beam size. If the electron gun attempts to send more current, the space-charge potential uses up all of the beam's available energy, denying the electrons the ability to move forward. This traffic jam extends back into the gun and shields the cathode from the electric field of the anode[73]. Increasing the anode-pulse amplitude further has no effect on the T_{EL} current. This argument assumes a steady-state situation; that is, the rise time of the anode pulse must be slow enough that the beam current can reach the main solenoid and relay back information about space-charge issues. The distance from the gun cathode to the entrance of the main solenoid is less than one meter, so electrons traveling at a speed of $0.2c$ require 15 nsec to traverse the distance. An additional 3 nsec is necessary for the electric fields to propagate back to the cathode surface.

The modulator circuit, expounded by Section 5.4, generates pulses with characteristic rise times of about 300 nsec, significantly larger than the 18 nsec total delay time for the current to begin, reach the main solenoid, and send back information about the space-charge effects. Because the rise time is much larger than the feedback delay time, any limits to the current due to the beam pipe's finite acceptance are broadcast back to the fields in the gun, deterring additional production of current, before the anode has increased significantly in voltage. If, however, the rise time was close to or less than this delay time, this feedback loop could enter an unstable, oscillating mode detrimental to the performance of the T_{EL}[66].

During normal operation of the T_{EL}, the cathode potential was set to a large negative voltage (nominally -10 kV) so that the acceptance of the beam pipe would not limit




Figure 5.18. Dependence of collector current on the anode voltage. A cathode voltage near zero produces large charge densities, which rob the electrons' energy and denies them the ability to pass through the beam pipe.

the beam current. In selective studies, the cathode was intentionally decreased in magnitude while the anode was still pulsed. Measurements of the beam current are shown in Figure 5.18.

When the cathode voltage is set to a large negative level, such as -12 kV, there is enough beam energy that the acceptance of the beam pipe does not limit the current. Instead, the current follows the anode-pulse voltage according to $I_e = PV_{\text{anode}}^{3/2}$, where $V_{\text{anode}}^{3/2}$ is the voltage of the anode with respect to that of the cathode, as discussed in Section 4.1. Over 3 A of current was produced without being significantly limited by the beam pipe.

As the cathode voltage is brought toward zero, however, the allowed current diminishes. Even at $V_{\text{cat}} = -10$ kV, the highest attainable anode voltage only produced 2.8 A of current. Further decrease in the cathode voltage brought down the maximum current.

At $V_{\text{cat}} = -4\text{ kV}$, a mere 1.5 A was produced at anode voltages even larger than those for other cathode-voltage settings (difficulties with the anode-modulation circuit prevented pulsing the anode to these impressive levels at the other cathode settings[67]).

At lower anode-voltage amplitudes, Figure 5.18 indicates a convergence of the cathode-current curves. Indeed, the current is independent of the pipe acceptance, contrasting with the constraints imposed at higher anode voltages. Instead, the current in this instance is solely determined by the perveance of the gun, identified in Figure 5.16.

A theoretically derived acceptance of $3.75\text{ }\mu\text{P}$ implies that at -4 kV , only 0.948 A may pass. However, the observed current for this low cathode voltage is seen to extend to at least 1.45 A. Similarly, a limit of 1.33 A exists for $V_{\text{cat}} = -5\text{ kV}$, while the measured current extends to 1.6 A and is continuing to rise. A couple reasons can explain this significant discrepancy between theory and experiment. The calculation presented in Section 4.3 assumes that the beam is centered in the beam pipe. However, an off-center beam can be shown to exhibit a higher acceptance[33]. In these measurements, an attempt to center the beam was not performed; instead, the TFL parameters followed their setting during tuneshift compensation.

In addition, Figure 5.18 displays data that was taken when the anode was being pulsed at the usual 47 kHz repetition rate with a long pulse duration of approximately $1\text{ }\mu\text{sec}$. The result is a relatively large duty cycle (4.8%). Residual gas molecules bombarded by the electron beam can result in positively charged ions. These ions are attracted to the negative potential well of the electron beam, and can accumulate over the time of each pulse. Between pulses, the potential well disappears, allowing the ions to scatter, but because the duty factor is rather large, many of the ions do not have

time to fully escape or recombine[51,68]. Due to this phenomenon, a net positive space charge can develop over time (usually just a few pulses), which counteracts the space-charge forces of the electron beam. More electrons can enter the beam pipe without being stopped by the space-charge potential, and therefore more current is allowed to flow[76].

The simple acceptance relation between cathode potential and beam current is convenient, but tuneshift compensation relates more directly to the space-charge density of the beam. A pair of BPM plates are described more fully in the following discussion, but essentially they are two electrodes that detect the electric field produced by a passing charged beam. In this case, adding the integrated waveforms from both plates in a BPM pair yields a measure of the total space charge passing through. In Figure 5.19, this amplitude is plotted again against anode-pulse amplitude for the same variety of cathode potentials. The vertical scale has arbitrary units, but the magnitude of the received signal scales linearly with the density of the beam's space charge.

Since the current obeys $I_e = \lambda_e \beta_e c$, where λ_e is the linear charge density, the space charge is expected to depend on the anode and cathode voltages similarly to the dependence of the current. However, the speed of the electrons depends on the current and cathode potential according to Equation 4.24. This complicated relation generates a more challenging connection between the beam current and charge density.

Nevertheless, it can be seen in Figure 5.19 that at low cathode potentials, the acceptance of the beam pipe limits the charge density; the curve for $V_{\text{cat}} = -4 \text{ kV}$ flattens at high anode amplitudes. However, even though the density is limited, the slower speed of the electrons creates more space-charge for essentially all of the parameter ranges accessible by the TEL. As the cathode potential is increased, two effects are observed.




Figure 5.19. Plot of charge density, as seen by BPM plates, on anode voltage. Due to the higher kinetic energy at $V_{\text{cat}} = -12$ kV, the charge density is significantly decreased, whereas the beam-pipe acceptance limit is observed when $V_{\text{cat}} = -12$ kV.

First, the speed increases, lowering the space-charge density for most of the space in Figure 5.19. This suggests that the largest tuneshifts from the T_{EL} will typically occur at lower cathode potentials.

However, the second phenomenon is that the beam's charge density does not saturate. At the highest levels of anode pulsing, a beam with more kinetic energy has a greater charge density than one with less kinetic energy, even though it is actually traveling faster. This is simply due to the fact that the electrons are given more total energy, so that they can continue traveling forward in spite of strong space-charge forces in the main solenoid.

The saturation effect at low cathode potentials is undesirable to typical T_{EL} operation; tuneshift beyond a certain amount become impossible. Even though the beam

current needs to be higher for moderate charge densities, the TEL was envisioned to run with a cathode potential of about -10 kV, which provides a smoothly, monotonically increasing level of space charge with anode voltage over the entire range of interest.

Both Figure 5.18 and Figure 5.19 do not show a quick transition point where the perveance-limited regime switches to the acceptance-limited regime. The main reason for the very gentle “rolling off” stems from the potential inside the beam itself as given by Equation 4.16. The potential at the very center of the beam is always lower than at the edge of the beam. Therefore, at moderate anode voltages, the edge of the beam may still have enough kinetic energy to continue moving, while the center of the beam has used it all to overcome the charge potential. The center then will begin to drop out, even though the edges keep traveling at full current density.

As the anode voltage is increased, a larger central portion of the beam is adversely affected, repelling the full current from entering the beam pipe. Eventually, the anode is pulsed at a voltage high enough that the full beam is space-charge limited. However, there exists a large range where some portion of the beam is limited by the pipe while the rest is not. The effect of this is to spread the transition point out into a rounded knee. This effect is readily apparent in Figure 5.18 and Figure 5.19. Not expressed in Figure 5.19 is the susceptibility of the acceptance limit to different parameters, such as cathode potential, residual gas density, pulse length, and corrector settings. The fact that these parameters easily change the charge density creates a difficult situation for a steady level of compensation. By avoiding operating the TEL under these conditions, these complications can be avoided. Instead, the TEL’s cathode potential is typically kept high enough so that the entire beam current can travel through the beam pipe; it




Figure 5.20. Sketch of the electrodes in the T_{EL} . While not to scale, a qualitative picture of the shape of each electrode is given.

is simply the anode voltage that determines the strength of tuneshift compensation from one bunch to the next.

5.3.2 Electrodes along the beam pipe

The cross-section of the beam pipe shown in Figure 5.17 indicates a space of radius 1.5 cm between the desired aperture and the high-vacuum wall. Along the length of the main solenoid in this region must fit the T_{EL} diagnostics. Figure 5.20 sketches the numerous electrodes involved in detecting and measuring the electron beam and the antiproton and proton bunches. Each of these electrodes is electrically isolated from the grounded beam pipe and is wired, through vacuum feedthroughs, to coaxial cables leading out of the Tevatron tunnel and to support electronics. Not drawn to scale, the two sets of BPMs reside near the beginning and the end of the main solenoid.

The wires shown in the center are two mechanically actuated forks. One is oriented vertically, the other horizontally, and each has a tungsten wire strung across the gap.




Figure 5.21. Pictures of different diagnostics. The left looks down into the gun bend area, where a pair of curved *elbow electrodes* and one *holy electrode* can be seen. The middle picture peers down the pipe at the two profile-measuring, retractable wires. The right looks at the support pipe with a diagonally cut BPM plate showing in the window.

Remotely operated motors are able to swing each fork into the middle of the beam pipe, where the electron beam is flowing. By adjusting the correctors, the beam can be swept across the fork, and the intercepted charge flows through the fork and again into cables that bring the signal out to be measured. The amount of current as a function of beam position yields data that can be converted into a profile of the current density. This conversion process uses a well-known technique often called an Abel inversion[58,60], and the resulting profiles are displayed in Section 5.2. The forks are always rotated out of the beam pipe before stores are injected in the Tevatron. In Figure 5.21, a picture is shown that looks down the beam pipe while both wires are inserted.

The *elbow electrodes* are curved imitating the path of the electron beam around each of the bends. Horizontally opposed, a high-voltage difference can be applied in order to generate a strong horizontal electric field. While this was tested in order to witness an $\vec{E} \times \vec{B}$ drift, the electrodes have not been used during normal TEL operation. Figure 5.21

shows a pair of curved elbow electrodes inside the bending area near the gun. Some mechanics involved in the operation of the wires can also be seen.

The *holy electrodes* are simply cylindrical electrodes that have a hole cut into one side. The electron beam passes through this hole as it enters and leaves the region of the antiproton orbit. These and the elbow electrodes were installed to assist with initial TEL commissioning. If the electron beam failed to pass through the solenoids and into the collector, observing which electrodes were absorbing current would indicate how to correct the guiding fields. Pleasantly, the electron beam had little difficulty propagating completely into the collector, and the utility of these electrodes diminished quickly. A picture of one can also be seen in Figure 5.21 next to the elbow electrode.

Next to the holy electrodes in Figure 5.20 are cylindrical electrodes intended for clearing out ions. Ions are created by electrons bombarding residual gas molecules floating in the beam-pipe vacuum. The once-neutral molecule can easily lose electrons, turning it positively charged and attracted to the electron beam's space charge. Past machines often witness instabilities associated with these ions interacting with the propagating beam[61,62], and a moderately high voltage on these electrodes successfully attracts these ions away from the electron beam. In actuality, the influence of ions has been small enough not to induce instabilities or other problems, and these electrodes are typically grounded with the others.

The beam-position monitors, or BPMs, are pairs of plates that produce signals corresponding to the transverse position of any passing charged particle. Discussed at length in other references[33], a BPM consists of two electrodes opposing each other and electrically connected to separate amplifiers or integrators. When any charged particle

passes through them, the electrical fields set up by the particle induce a charge on each plate. While this phenomenon occurs in any electrode, the two separate plates each induce an image charge linear with the distance that the particle is from the plate. By normalizing the amplified voltage on each plate to the total charge observed, the average position of any amount of charge can be measured. The position simply follows the linear relationship:

$$x = k \frac{V_A - V_B}{V_A + V_B}, \quad 5.4$$

where x is the transverse distance of the beam from the center of the beam pipe, V_A and V_B are the measured voltages on the two plates, and k is a constant empirically determined from calibration measurements. The BPMs in the T_{EL} have a calibration constant $k = 33.57$ mm when each V_i is the total integral of the bunch's charge profile, which in turn is the integral of the doublet current signal as charge is first pulled onto the electrode, then returned, as the bunch passes by.

Each BPM measures the beam's position in only one dimension and at only one longitudinal position. Therefore a pair of BPMs are needed on the "upstream" end of the main solenoid in order to record both the horizontal and vertical positions, and another pair of BPMs are needed on the "downstream" end. By pinpointing the electron beam's position at these two points, the trajectory of the electron beam is determined.

In the T_{EL}, a fast oscilloscope and a dedicated computer process the waveform signals and compute positions constantly during stores. The design of this system includes the ability to measure the position of antiproton bunches and proton bunches with the same BPM plates as the electron beam. In this manner, it becomes possible to confirm that the electron beam and the antiproton bunches are collinear within the main solenoid.

BPMs, while theoretically simple, are difficult to implement without complications. The first is the amount of noise in the system — the T_EL BPM system must average measurements over hundreds of turns in order to report positions with low error bars. Additionally, the oscilloscope cannot digitize the electron, antiproton, and proton signals at precisely the same time. Instead, the software accumulates and processes each species sequentially over several seconds. Similarly, the one oscilloscope can only monitor one BPM at a time; the software is able to automatically switch through all four BPMs. In this manner, the horizontal and vertical positions both upstream and downstream of all three species are reported during normal operation.

In Figure 5.22, the positions reported by the BPM system for one location and plane are shown over the course of five hours. In order to see the fluctuations between readings, the average value of each measurement was subtracted and a small offset added to better display the spread in each data set. Over the course of the five hours, the Tevatron orbit positions will change slightly, reflected in the long-term trends in the data. After accounting for this, the standard deviation for the data sets are $\pm 1.2 \mu\text{m}$ for the proton signal, $\pm 2.1 \mu\text{m}$ for the antiproton signal, and $\pm 8.9 \mu\text{m}$ for the electron signal[70].

An additional issue is the linearity of the BPM signals across a wide portion of the beam pipe. In order to improve this linearity, the BPMs are constructed slightly differently than Figure 5.20 implies. The plates are cut at an angle, which can be seen inside its support pipe in Figure 5.21. This design allows the measurement of position to be linear over a large range[63,64], which is important since the orbits of the antiproton and proton bunches are approximately six millimeters apart.

The T_EL BPM system has received a significant amount of attention. Several references have detailed the operation of the system[40,65], but one significant drawback




Figure 5.22. Example of optimal BPM positions. The average positions reported by the downstream horizontal BPM was 2.29 mm, -2.12 mm, and 2.33 mm for protons, antiprotons, and electrons respectively. This average value was subtracted out (with $\pm 20 \mu\text{m}$ added) in order to zoom in on the fluctuations. The upper black data set represents the proton-bunch position, the lower black set are the antiprotons, and the grey are the electrons.

has been the frequency-dependent response of the plates. The frequency spectrum of a passing antiproton bunch centers around 53 MHz, while the much longer electron beam encompasses a wide bandwidth around 2 MHz. This difference generates a systematic difference in the BPM calibration between the two species, and is mutable by a change in the electron-beam pulse width.

To showcase this effect, the anode was pulsed at varying lengths ranging from 20 nsec to 100 nsec, and the apparent position of the beam was measured and plotted in Figure 5.23. The BPM system reported over 1 mm difference in position over this range, which does not include the pulse length of a typical antiproton bunch. The anode modulator circuit is unable to produce shorter pulses. Because the magnetic correctors did not need to be adjusted to push the electron beam all the way to the collector, the actual beam position is believed to be reasonably stationary over this range. During beam-beam compensation studies, the electron beam has been centered on the bunch orbit by maxi-




Figure 5.23. Position dependence on length of electron beam. The fact that the BPM system calculates different positions depending on the length of the beam corroborates the frequency-dependent response of the plates.

mizing the tunes and minimizing particle losses, described in Chapter 7, even though the BPM system reports a significant offset in the position of the two species.

The calculation of position can be challenging, as each signal V_A and V_B always contains some level of noise, but subtracting the two signals yields a result much smaller than either of the two sources, while the noise level is increased. The signal-to-noise ratio can be difficult to overcome, which explains why typical position measurements averages over hundreds of turns. Due to this susceptibility to noise and the nanosecond-scale signals on the BPM leads, special attention needed to be placed on the cabling from the plates all the way to the oscilloscope. Fifty-ohm cable is attached to each BPM plate, drawn through coaxial feedthroughs in the vacuum wall, brought out of the Tevatron tunnel, and into the BPM electrical apparatus. While the outer conductor is grounded at several places along the route (such as the vacuum feedthrough and the signal switcher) reasonable preservation of the signals has been observed.

The other electrodes also have leads drawn through the vacuum wall and to various power supplies and equipment. All of these cables are also shielded in 50-ohm cabling and separated from pulsed power signals, such as the anode modulator pulses. This level of caution succeeded in preventing significant contamination of low-level signals by high-power level transients.

5.3.3 Collector design

A drawing of the collector appears in Figure 5.24, with a view of the expanding electron beam. Spreading out the beam has several uses, one being the distribution of the heat load. Locating the heating in one spot could melt the copper if the TEL is operated at full capacity. Instead, the magnetic field lines beyond the collector solenoid spread out, and the electrons, religiously following the field lines, are absorbed by a much larger area of the copper collector. Additionally, chilled water is piped into the collector, where it passes through ducts within the copper and extracts up to 50 kW of heat [76]. The intake pipe is shown in Figure 5.24 to illustrate the setup.

A more serious concern is the production of secondary electrons. Whenever an energetic particle impinges on a surface, it has the ability to produce two effects. First, it can inelastically collide with other electrons and impart enough kinetic energy for them to escape from the metal's surface. Second, it can elastically collide with an atom and reflect off of the surface.

The inelastically produced electrons generally have a small fraction of the incident electron; few are above 20 eV. The elastically reflected electron, on the other hand, often has most of its original energy. These electrons can travel backward through the TEL,

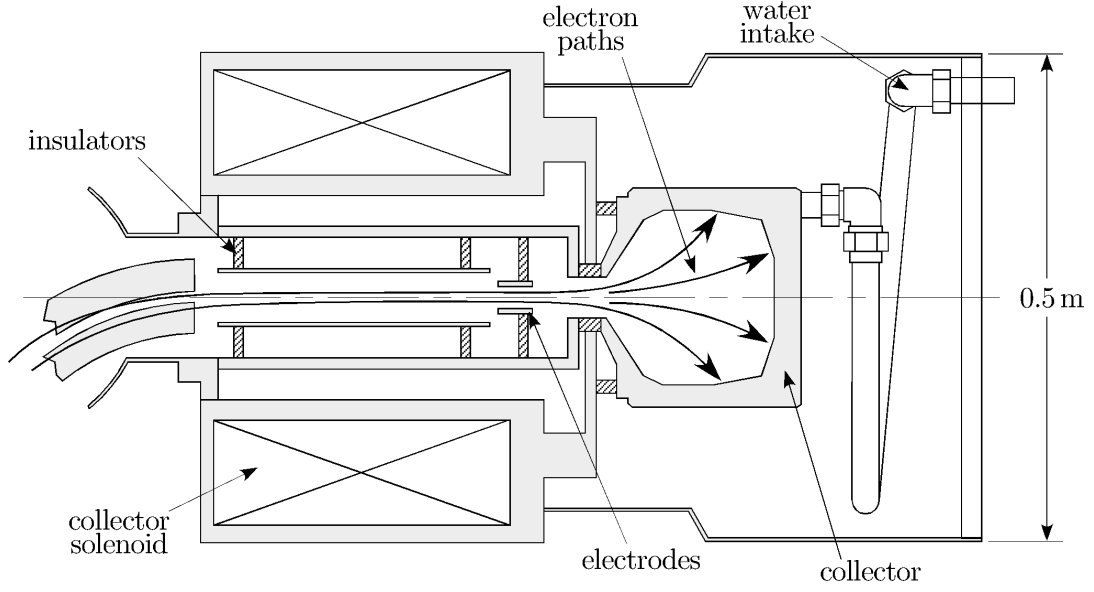


Figure 5.24. Scaled drawing of the collector cross-section. The collector itself is a water-cooled copper cavity that resides outside of the solenoids. This allows the electron beam to spread out, distributing the heat load and decreasing the production of secondary electrons.

creating a second, weaker, electron beam collinear with the first. The primary and secondary beams are able to interact adversely, generating what is typically known as a *two-stream instability*[71,72].

The mechanical design of the collector targets this issue. The secondary electron has a kinetic energy that is conserved as it attempts to pass into the collector solenoid. Its energy can be split into parts parallel and perpendicular to the field lines; that is,

$$v_{\text{tot}}^2 = v_{\parallel}^2 + v_{\perp}^2, \quad 5.5$$

where each portion is flexible total remains fixed. The electron, witnessing the magnetic field lines, takes on a spiraling orbit as derived in Section 4.2. As pointed out in that

discussion, the angular momentum of the field-and-particle system must be conserved, leading to the result that Br^2 is constant as the magnetic field B intensifies as the electron approaches the collector solenoid.

However, the transverse momentum p_\perp is related to the orbital radius by $p_\perp = eBr$, which means that v_\perp^2/B must also be constant, or as the magnetic field increases, so must the perpendicular velocity. Equation 5.5 can therefore be written

$$v_\parallel^2(s) = v_{\text{tot}}^2 - \left[\frac{B(s)}{B(0)} \right] v_\perp^2(0),$$

where s is the distance along the field lines starting at the collector surface. This relation holds for real values of velocity, which enforces the constraint

$$v_\parallel(0) > v_\perp(0) \sqrt{\frac{B(s)}{B(0)} - 1}. \quad 5.6$$

An electron must originate with enough parallel momentum to overcome the magnetic compression; if the electron does not have enough, it will run out of longitudinal momentum and be returned back to the collector surface.

The T_{FL}'s collector surface has a residual field on the order of 100 G, while the collector solenoid runs at 3.8 kG, implying that the longitudinal momentum needs to be over six times larger than the perpendicular momentum. If the electrons are emitted from the collector surface uniformly over all solid angles, only 1.2% of those electrons meet that constraint. More importantly, the number of electrons that can pass all the way into the 35 kG main solenoid is 0.14%.

The total number of secondary electrons is able to be larger than the number of primaries, but the kinetic energy of most of them is typically a few eV, which significantly

impairs their ability to pass from the collector into the solenoid region[73]. Backscattered electrons often have kinetic energy close to that of the primary, but the number of these that are germinated from copper is about 30 % of the number of primaries[74], and for an oblique incident angle, the percentage that point in the small solid angle satisfying Equation 5.6 is even smaller than the percentages listed above[75]. Under normal operation, the TEL is able to retain at least 99.7 % of the incident electron beam, as comparisons with the cathode and ground current (explained in the next section) determine.

5.3.4 Properties of a depressed collector

The collector is electrically isolated from the ground of the beam pipe. The electron beam is absorbed by the collector and produces a current in a lead that is brought back, via a floating power supply, to the cathode. By adjusting the power supply, the collector can be set to a voltage different than ground. If the power supply is set to zero, then the voltage on the collector would equal that of the collector. In this theoretical situation, electrons are born at the cathode potential and accelerated by the anode and the beam pipe according to Section 4.3, but as they approach the collector, they are slowed back down to zero energy.

Realistically, the collector must be at a voltage somewhat more positive than the cathode, so that the electrons still carry some forward momentum when they are received by the collector. Nevertheless, there are two advantages to keeping the voltage difference small. First, the heat load generated by the incident beam is directly related to its kinetic energy with respect to the collector: a 10 keV electron on a (-)6 kV collector

imparts only 4 keV of energy. When the TEL is operated at maximum capacity, the total power deposited in the collector could be considerable, and lowering the voltage difference would be useful.

The second advantage to a small voltage difference is related, since the power dumped into heat originates with a power supply. If the collector is grounded, the supply feeding the cathode would need to maintain -10 kV for as much as three amperes. It is this supply that feeds the heat load of the collector. Instead, a supply connecting the collector back to the cathode but with a smaller voltage difference needs to provide much less power while still maintaining the complete current loop. This technique of returning the current back to the cathode on a path electrically separate from ground is known as a *recirculating* beam, and is used in many high-power beam systems[33].

A side benefit of floating the cathode power supply is that the current between the cathode and ground is, ideally zero. Discussed in the next section, this current is nonzero only when electrons are able to get to the beam pipe, which means the beam is scraping along a side, a magnet is incorrectly set, or something has failed. A safety interlock can use this signal to confirm that the TEL is operating adequately.

The complex geometry of the collector, the changing beam size, and the non-grounded collector walls make a calculation of the collector acceptance difficult. However, it is important that the electron beam is able to pass into the collector just as Section 4.3 analyzed the beam passing into the beam pipe. A test of the collector consists of setting its voltage to nearly that of the cathode and measuring the current that it receives. As the voltage difference approaches zero, the current is expected to drop to nearly zero, implying the space-charge potential is too high for the current to continue moving forward.




Figure 5.25. Measurement of collector acceptance. As the collector voltage is adjusted with respect to the cathode voltage, the current admitted by the collector changes. At around 0.75 kV, the space-charge potential of the entire beam just equals to the kinetic energy.

At larger voltages, more and more current will be collected, until finally the amount will level off at the full available current.

The data from this experiment is shown in Figure 5.25. At zero voltage difference, a small amount of current is measured on the collector, implying that through energy spread or space charge forces, a quarter of the beam can still manage to reach the collector surface. However, as the voltage is raised, the amount of current reaching the collector increases. As the difference approaches 1 kV, all of the current is received.

Ideally, a sharp corner would appear where the space-charge potential just barely equals the full energy of the electron beam, and the acceptance could be calculated using the voltage and current at that point. However, due to the beam's energy spread, the finite rise time, and the complex geometry of the collector, this corner is smoothed out,

and the acceptance can only be approximated. Using the data point where the maximum current is truly witnessed, the acceptance of the collector appears to be $9.21\ \mu\text{P}$. Using the point somewhat lower in total current for the calculation, the collector functions at $13.61\ \mu\text{P}$. In reality, the collector has an approximate acceptance of around $10\ \mu\text{P}$, much larger than that of the beam pipe[76]. The collector has never been a limit on the T_{EL}'s performance, since its voltage can always be increased if necessary. Normal T_{EL} operation usually sets it at about 5 kV above the cathode voltage.

Two electrodes are also shown in Figure 5.24, which have primarily been used for monitoring the beam's passage into the collector. For example, the *scraper* electrode, closest to the collector, has its own current monitor, and if it reads something other than zero, the T_{EL} beam is, at least partially, running into it. Adjusting the downstream magnetic correctors such that this signal is eliminated assures optimal performance of the T_{EL}.

5.4 Electrical subsystem

While the various elements, such as the electron gun, the collector, and various electrodes, have been analyzed, the manner in which these elements are typically operated has not been described. Likewise, each of the probes used to measure various features of the beam must be understood before they can be trusted.

5.4.1 Basic model of electrical layout

A recirculating DC beam could be generated by the simple circuit illustrated in Figure 5.26. In this case, the cathode power supply does not need to produce any

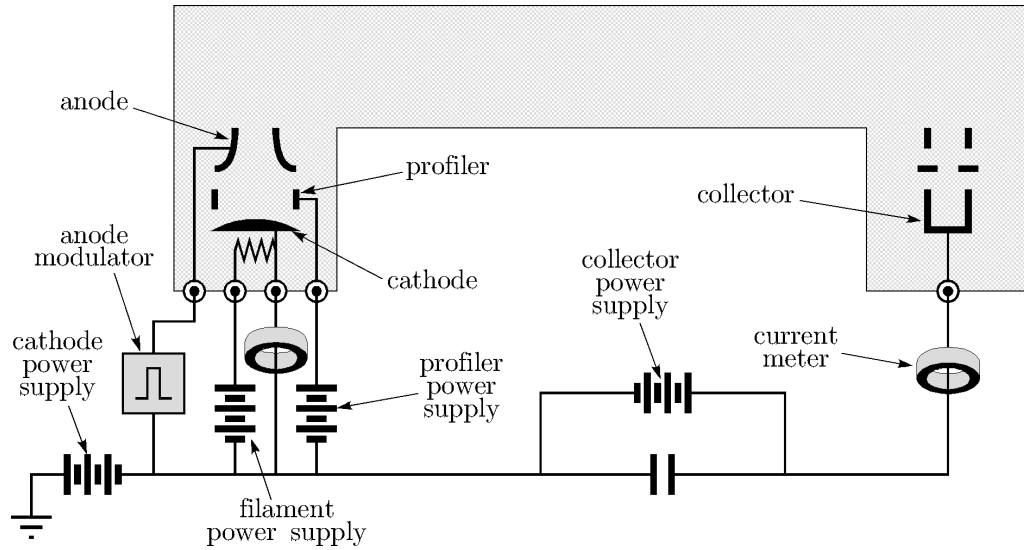


Figure 5.26. Oversimplified schematic of the TFL circuitry. In this picture, the power supplies are located close to the TFL, and the anode, filament, and profiler voltages are referenced to stable cathode voltage.

current. The beam current does flow through the collector power supply, but the voltage that this supply must support can be significantly less. The anode power supply also carries no current and simply “kicks” the current around the loop.

The actual TFL requires pulses of beam current with fast rise and fall times. In addition, the Tevatron tunnel, filled with radiation during operation, is a poor place for solid-state electronic equipment. The power supplies must be located above the tunnel, connected by cables over a hundred feet in length. Keeping the inductance of the current loop to a minimum and assuming the TFL will only be used in a pulsed application, the loop was configured with a large capacitor near the TFL, while the collector power supply that simply maintains the voltage difference is located upstairs.

Likewise, the cathode power supply is located upstairs, but clearly as current begins to flow from the cathode, the voltage will begin to drop. This simply relates to the total

capacitance between the cathode and ground. In direct measurements, the cathode-ground capacitance of the gun itself is only about 20 pF, but the cables attached to it increase that significantly. The process of changing the cathode voltage follows $dV/dt = I/C$ if C was a known value, but the cables act like a distributed capacitance, creating a confusing circuit to model. It was decided to add an additional capacitance between the cathode wire and ground to “make sure” the voltage remains steady, and monitors of the cathode voltage before and after the addition show marked improvement in its stability.

Another concern is the cathode filament connections. While an ideally designed cathode would be electrically isolated from the filament, constraints on the space allotment made this difficult. Instead, the filament is connected on one side to the cathode, and the other side has its own tap. The filament is therefore powered by 60 Hz which is electrically isolated through a high-voltage transformer. This allows the primary side to be referenced to ground and only the secondary is floating near the cathode voltage.

While some of this discussion may seem overly fastidious, it is important to understand the cathode’s connections when the question is asked how the current from the cathode is measured. A high-bandwidth, commercial current transformer encircles the wire attached to the cathode. The transformer’s signal is preserved over the long Heliax cable by only grounding it upstairs where an oscilloscope measures it. This prevents ground noise to corrupt the signal. Keeping the impedance at $50\ \Omega$ also reduces electrical coupling from other sources over the long propagation distance and reflection issues. Risetimes of 1 or 2 nsec and currents of a few milliamperes are visible. However, the low-frequency filament current passes through the current transformer, so the return cable was obliged to pass through it also; the two currents are always in opposition and therefore cancel.

This complexity is depicted in Figure 5.27. The majority of the electronics and power supplies are harbored outside of the tunnel due to the damaging radiation produced by the Tevatron operation. This means that long cables must span the sixty meters between the T_PL itself and the gallery where these supplies are situated. For the majority of these cables, high-voltage, shielded coax (RG-213) of $50\,\Omega$ impedance is used. In order to connect the capacitors, current meters, and various connections, a high-voltage enclosure was constructed that sits next to the electron gun and houses the components. This enclosure is not shown in Figure 5.27, but is detailed in other documents[77]. In order to keep cable issues at a minimum, the cables leading to the gun and collector from this box are minimized.

A simple brief analysis of the capacitances reveals one concern. The capacitance from the cathode conductor to ground provides the initial current when the anode draws current (hence increasing that capacitance holds its voltage better). If the cathode itself has about 20 pF to ground, the two-foot cable provides another 60 pF, and the additional capacitor is 1 nF, then only about 92 % of a sudden current pulse will be witnessed by the current transformer. Shortly after this moment, the additional cables (traveling upstairs) and capacitance on the other side of the recirculating capacitor contribute to the current flow (eventually replenishing all of the lost electrons). The effect on the current transformer's signal is similar to a low-pass filter, but the cathode's capacitance is low enough to keep the distortion below significant levels.

In order to confirm that the beam current passed all of the way into the collector, another current transformer monitored the current returning from the collector to the recirculating capacitor. A third transformer watched the scraper electrode's current

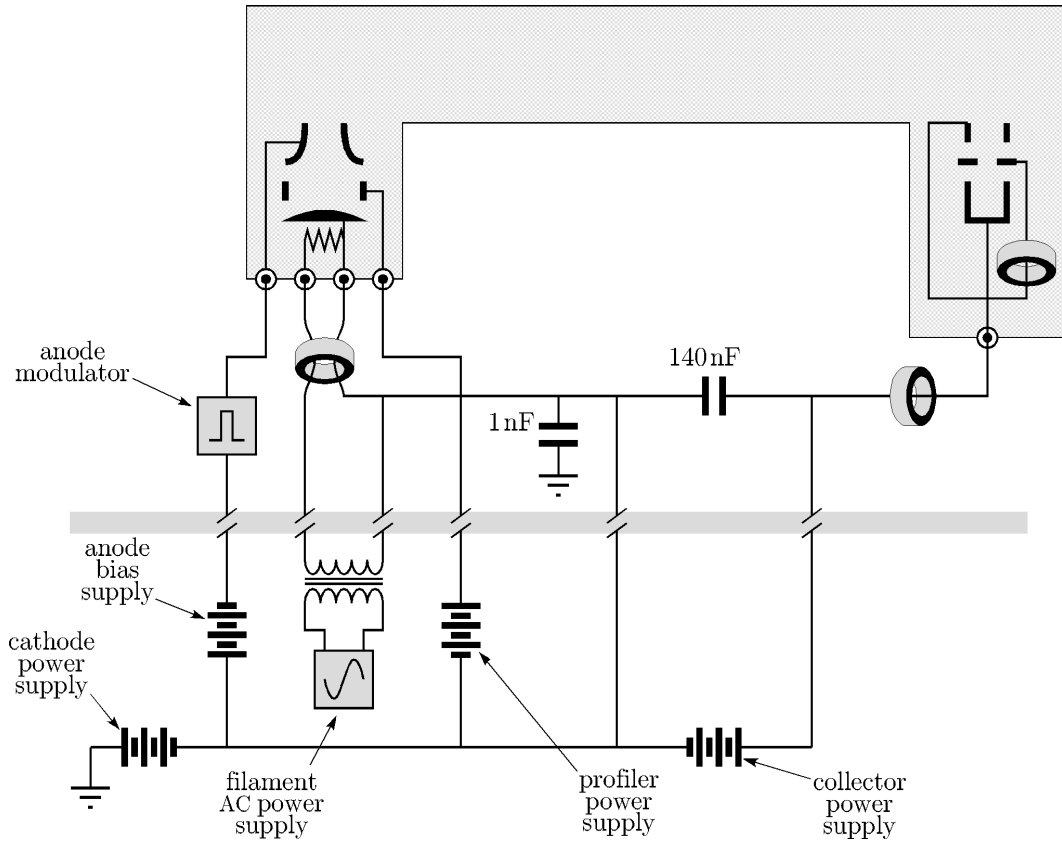


Figure 5.27. Realistic circuit of the TEL. The shaded barrier is actually about sixty meters between the Tevatron tunnel and the gallery where the supplies and electronics are stationed. The filament is connected to the cathode potential on one side, so the AC power supply needs to be isolated from ground via a transformer, and its signal needs to be “subtracted” from the cathode current meter.

(Figure 5.17 shows that they were placed beside the collector, before any cable runs), since this provided the narrowest aperture and the easiest way to adjust field strengths in order to steer the beam into the collector (finding a zero is less challenging than finding a maximum). The scraper fed into the collector cable, so that the collector current should be identical (though delayed) to the cathode current. A typical oscilloscope trace of these three currents is shown in Figure 5.28. As has been consistently observed, the peak




Figure 5.28. Oscilloscope trace illustrating a typical pulse of cathode, collector, and scraper currents. The cathode and collector currents peak current of at approximately 2.45 A. The time delay is simply the amount of time it takes for the electron beam to travel the entire path from cathode to collector, plus another 1.5 m of additional cabling to the high-voltage enclosure.

collector current is less than that of the cathode current. Section 6.1 explains how the electron beam is expected to lengthen longitudinally due to energy spread, which would lower the peak current at the collector.

A second reason for the discrepancy between the cathode and collector currents is the much larger capacitance between the collector and ground. It was found to be about 600 pF, much larger than that of the cathode. An additional 3 nF of capacitance was placed on the other side of the current transformer to attempt to offset this capacitance, but a low-pass filtering effect is still expected on the signal. Fast risetimes and low noise has been seen on the collector current transformers; however, the cathode current

remains the more dependable resource for observing the beam current. Most data in this dissertation use the cathode current readings.

It should be noted that a better design would be to include a vacuum-compatible current transformer (or one outside a ceramic break and covered by a continuous ground plane) slightly downstream of the anode. This would enable direct measurement of the beam current without complications, though juggling the bore size, low-frequency limit, and high-frequency limit, especially in a magnetic field, would require some effort. Another less-choice option is putting it right behind the cathode-filament connections around their leads — still some analysis of its feasibility would be needed. Both of these options are more trustworthy in their measurement of actual beam current, the former being the best, since it surrounds the beam itself. A similar approach would be useful for the collector current as well.

5.4.2 Anode modulation

The requirements for pulsing the TEL were described in Chapter 3. The beam current not only must reach an arbitrary new level within 396 nsec, it must be stable at the new level early enough to propagate down the interaction region. If the speed of the electrons is about $0.2c$ and the distance is 2 m for the interaction length plus another 1 m from the cathode, the “flushing” time is at a minimum 50 nsec. The anode voltage must be stabilized at its selected level, as Section 3.2 points out.

It is tempting to think that a significant overshoot (or undershoot) is acceptable, since the added “bump” in current can be compensated by dialing in a lower (or higher) chosen level. This works well for single-bunch operation (including this discussion), but

if tune-shift compensation is set for multiple bunches, adjusting the current level for one bunch should not affect the current levels for neighboring bunches.

Impressive solid-state technology is available for high-voltage, high-speed power switching. However, the waveform as shown in Figure 3.3 requires a high-bandwidth amplifier. These difficult constraints led to a large effort to produce an adequate amplifier. The current scheme, which has barely succeeded the T_{EL}'s requirements, centers around a tetrode RF tube. Several publications describe this circuit, but the high-power circuit is briefly mentioned here.

A commercial, low-level waveform generator is triggered by the Tevatron-based timing system mentioned in Section 1.2 and produces a small version of the desired pulse train. This is amplified by a 500 W wideband amplifier. The output of this device, illustrated in Figure 5.29, is transported via high-power Heliax into the Tevatron tunnel to the grid electrode of the RF tube. Meanwhile, about 6 kV provided through a large resistor, shown with the tube in Figure 2, to the top plate of the tube and, before any pulses arrive, the tube is left *on*, which means current pours from the top plate to the bottom collector. This maintains the plate's voltage at nearly the same low voltage as the bottom collector, while the large resistor dissipates the large current.

When a pulse arrives, the grid cuts off the current a certain amount, which allows the tube's plate to jump in voltage. The gun's anode is capacitively coupled to this plate, so the jump in voltage is transferred to the anode. The waveform arriving at the grid is magnified at the gun anode. In order to ensure that the anode stays at zero potential with respect to the cathode when beam current is not wanted, the anode is inductively coupled to the cathode.




Figure 5.29. Typical trace of the signals produced by the anode modulator circuit. The low-level pulse is sent to the wideband amplifier, which is located in the gallery. Then the high-power signal is sent into the tunnel, where the RF tube is located. The voltage is applied to the grid electrode of the tube, and the output voltage is sent to the gun anode.

This system has a few complications. The first is that the 6 kV power supply must be willing to drive a $1.5\text{ k}\Omega$ load up to 100% of the time. Second, the resistor (and its cooling water) must dissipate the same power. These specifications lead to rather large and unwieldy components. The RF tube also must be high-voltage and high-power, making it also quite expensive, and it requires filament heating and biasing. The entire amplifier circuit, therefore, has several adjustable parameters that will affect its gain, risetime, and repeatability.

While the overall gain can be achieved by turning everything up enough, the performance of the RF tube deteriorates rather quickly (weeks of operation) until it either

puts out negligent voltage or it fails catastrophically. Either way, replacing the RF tube is expensive and often impossible, so recent experiments with the lens use severely restricted gains. Even so, the beam current has slowly decreased significantly over a year, and several teams have recently been analyzing possible methods to improve or replace the RF-tube concept. One option that is showing promise is a “distributed amplifier” scheme. In this case the one RF tube is replaced by several smaller (and therefore more manageable) tubes. Delaying the low-level signal activates each tube in succession, but their outputs are delayed also, so each contributes to the final single waveform[78].

Chapter 6:

T_EL integration

*For it is written: I will destroy the wisdom of the wise;
the intelligence of the intelligent I will frustrate.
—1 Corinthians 1:19*

6.1 Instabilities

When the concept of beam-beam compensation was being proposed for the T_EL, a number of questions arose in regards to the introduction of a new and untested device that can affect and be affected by the Tevatron bunches[13]. Besides the hope of tunes shifting the bunches, the T_EL electron beam can react to the passage of bunches while affecting them at the same time. In general, this circular influence can lead to instabilities. Cal-

culations, simulations, and measurements of different effects have been performed, since any adverse situations will negate any benefits of tuneshift compensation.

6.1.1 Transverse-Mode Coupling Instability I

The electron beam interacts with the head and tail of each antiproton bunch. Fortunately the electron beam is “replenished” between consecutive bunches, so that the electron beam does not transmit forces from one bunch to the next. However, one portion of the electron beam does interact with the head and tail within a given bunch. This implies that if the head of the bunch was to affect the electron beam, the electron beam would be altered when the tail arrived. On each passage, this phenomenon could escalate, creating an instability for the antiproton bunches.

This description is similar to numerous head-tail interactions. In each case, the head of the bunch produces a wakefield that, through some medium, affects the tail. The kick is imparted on every turn, but the tail particles migrate to the head (through synchrotron oscillations) and cause other particles to be affected. The circular relationships between different portions of the bunch is what destabilizes the motion, leading to a transverse-mode coupling instability (TMCI). In this section, the simplest mode of this TMCI is considered. The antiproton is modeled as two macroparticles, each containing half of the total antiproton charge. Oscillations of the two macroparticles illustrate the dipole mode of this instability and can be solved analytically. The next section presents results from simulations of more general modes.

In most instances of TMCI, an offset in one plane generates a kick in the same plane. However, the T_EL solenoid generates motion in one transverse plane from motion

in the other. Therefore the analysis of this instability is somewhat more complicated than standard head-tail phenomena. The full treatment has been produced in several other publications and requires a development of the concept of wake functions. In this dissertaion, pedantic theory is replaced by relevant results.

A simple model of the coupling is illustrated in Figure 6.1. First, the electron beam interacts with the head of the antiproton bunch. However, the beam and the bunch cannot be perfectly colinear, so there is a small Δx offset between their centers. If a slice of the antiproton bunch near the head has total charge Q , then the electron beam feels a collective momentum kick

$$\Delta \mathbf{p}_e(t=0) = \frac{2eQ\Delta x}{(1+\beta_e)c\sigma_p^2} \hat{x} .$$

This momentum kick causes coherent gyromotion of the entire electron beam similar to Equations 3.3 and 3.5. In this situation, the motion will follow

$$x_e(t) = \frac{\Delta p_{x0}}{eB} \sin \omega t ,$$

$$y_e(t) = \frac{\Delta p_{x0}}{eB} (1 - \cos \omega t) ;$$

where ω is still defined as $eB/\gamma m$.

It is evident that a displacement in one direction results in movement in both directions as the rest of the antiproton bunch passes. Wake-function analysis describes the impact of the electron beam's fields on any subsequent slice of the antiproton bunch. If a particular slice is distance s behind the first and has the same total charge Q , then it




Figure 6.1. Sketch of head-tail coupling, leading to TMCI. The electron beam begins by being slightly off-center of the antiproton bunch. This offset generates a kick in the electron beam away from the antiprotons. However, the solenoid field translates that horizontal kick into horizontal and vertical displacement, so that the electron beam's impact on the tail of the antiproton bunch depends on the head.

will receive momentum kicks

$$\begin{aligned}\Delta \mathbf{p}_x(s) &= -\frac{eQ}{c} (W_d(s)\Delta x - W_s(s)\Delta y), \\ \Delta \mathbf{p}_y(s) &= -\frac{eQ}{c} (W_s(s)\Delta x + W_d(s)\Delta y); \end{aligned}\tag{6.1}$$

in which the direct wake function $W_d(s)$ and the skew wake function $W_s(s)$ are defined by

$$\begin{aligned}W_d(s) &= W \sin ks \theta(s), \\ W_s(s) &= W (1 - \cos ks) \theta(s); \end{aligned}$$

where $\theta(s)$ is the step function and W and k are

$$W = \frac{4LI_e}{(1 + \beta_e)\beta_e c r_e^2 \sigma_{\bar{p}}^2 B},$$

$$k = \frac{\omega}{(1 + \beta_e)c}.$$

In this situation, L is the interaction length of the T_{EL}, I_e is the beam current, r_e and $\sigma_{\bar{p}}$ are the electron beam and antiproton bunch radii, and B is the solenoid field strength.

Through a tedious analysis of the synchrotron motion, the impact of the momentum kicks described in Equation 6.1 can be solved. The direct wake effect can be mitigated if the gyrofrequency ω is fast enough to allow many rotations during the passage of the bunch. The skew wake is decreased by keeping the tune difference $|\nu_x - \nu_y|$ large.

Bounded solutions can be found in terms of the horizontal and vertical tune-shifts $\Delta\nu_{x,y}$ and the Tevatron's synchrotron tune ν_s . In this case, the minimum solenoid field B_{\min} is given by

$$B_{\min} \approx \frac{1.3 e N_{\bar{p}}}{r_e^2} \sqrt{\frac{\Delta\nu_x \Delta\nu_y}{\nu_s |\nu_x - \nu_y|}}, \quad 6.2$$

where $N_{\bar{p}}$ is the number of protons. For a realistic situation where $\Delta\nu_x = 0.01$, $\Delta\nu_y = 0.003$, $N_{\bar{p}} = 6 \times 10^{10}$, $r_e = 1.6$ mm, $\nu_s = 0.001$, and $|\nu_x - \nu_y| = 0.01$, the minimum magnetic field is approximately 0.96 T. This value is based on theoretical precision; a real machine could easily have other phenomena occurring that encourages the growth of the instability. In addition, the T_{EL} is designed to allow for different operating parameters. For example, a higher-current electron gun or more antiprotons per bunch can easily push the minimum magnetic field higher. Finally, this analysis only pertains to the dipole mode of TMCI

instabilities; the head and tail act as two macroparticles that oscillate against each other. A more general description of the antiproton bunch that incorporates multiple modes is pursued next.

6.1.2 Transverse-Mode Coupling Instability II

The preceding section derives the relation between the minimum stable magnetic field and various parameters, but only accounts for the dipole motion of the antiproton bunch. By splitting the bunch into a greater number of macroparticles, a numerical tracking algorithm observed the excitation of many different TMCI modes as a function of the relevant parameters found in Equation 6.2. In these simulations, between 128 and 2048 macroparticles were assembled following a Gaussian distribution in longitudinal position and momentum. The horizontal and vertical position of every macroparticle was followed as they rotated through their synchrotron orbits. A revolution around the Tevatron ring was modeled as a rotation through phase space in all three dimensions; the longitudinal position was advanced by the phase angle $2\pi\nu_s$, and the transverse positions were likewise rotated by $2\pi\nu_{xi}$ and $2\pi\nu_{yi}$. The transverse tunes were spread over a small range, typically 0.002, in order to mimic the tune spread of a real bunch.

After each revolution, the macroparticles were sorted from head to tail and allowed to collide with the electron beam. This interaction caused electron-beam motion as described by Equation 4.15, which in turn influenced subsequent antiproton macroparticles. Each macroparticle acted on and responded to the electron beam. Afterwards, the new transverse position and momentum were translated back into a betatron amplitude and phase for each direction, and the antiproton bunch was ready to be sent around the Tevatron again. The tracking program repeated this process for tens of thousands of turns, depending on the specific application.

The first analysis related the threshold magnetic field to the strength of the T_{EL} electron beam. In this situation, the tuneshift of the electron beam and the number of antiprotons were fixed, and the solenoid's field strength was chosen. After 10,000 turns, the macroparticle transverse-oscillation amplitude was tested. The instability would manifest itself as a significant increase in the oscillation amplitude over this many turns. If no significant increase was detected, the simulation's solenoid strength was lowered; if a large increase was observed, the solenoid strength was raised.

By repeating this method and iteratively adjusting the field strength, the threshold magnetic field could be estimated for a specific bunch population and T_{EL} tuneshift. Figure 6.2 plots the results of these simulations, where each data point represents the threshold magnetic field for TMCI. If the magnetic field is larger than the threshold, the instability was held in check by the rigidity of the electron beam. Below the threshold, and significant oscillations of the antiproton macroparticles was observed. As the figure indicates, three different antiproton numbers were tested at a range of T_{EL} tuneshifts.

As the strength of the T_{EL} or the number of antiprotons increases, the required magnetic field strength also increases. Three bunch populations were chosen, $N_{\bar{p}} = 1 \times 10^{10}$, 6×10^{10} , and 1×10^{11} , and the effective tuneshift in both the horizontal and vertical directions was adjusted between 0.002 and 0.020. The simulation was run for 10,000 turns, which is significantly less than a second of time in the Tevatron, but the minimum magnetic field is clearly dependent on both parameters. More specifically, B_{\min} follows roughly a linear dependence on $N_{\bar{p}}$ and $\Delta\nu$. In these simulations, the horizontal and vertical tuneshifts were identical, so that $\Delta\nu = \sqrt{\Delta\nu_x \Delta\nu_y}$. This linear relationship is predicted in Equation 6.2, so the simulation is in good agreement with the two-mode analytical model.




Figure 6.2. Minimum field strength in the T_{EL} solenoid to prevent TMCI. The simulation gives sensible results; the approximately linear dependence on the number of antiprotons and on the tuneshift is apparent. The staircase appearance of the $N_{\bar{p}} = 1 \times 10^{10}$ case stems from the limited number of iterations performed. The linear dependence relates well to the linear dependence of B_{\min} on $N_{\bar{p}}$ and the total tuneshift $\Delta\nu = \sqrt{\Delta\nu_x \Delta\nu_y}$ expressed in Equation 3.13. In these simulations, $\nu_x = 0.585$, $\nu_y = 0.575$, $\nu_s = 0.0012$, and $\sigma_{\bar{p}} = 0.7$ mm.

Figure 6.2 predicts that at the current operation of the T_{EL} at 35 kG, a tuneshift of $\Delta\nu = 0.020$ could act on antiproton bunches with $N_{\bar{p}} < 6 \times 10^{10}$ without exciting the TMCI. This is double the maximum tuneshift results discussed in Chapter 6, providing a large safety margin before exciting this instability.

The synchrotron tune ν_s is also expected to influence the onset of TMCI. In Figure 6.3, this parameter was varied and the minimum solenoid strength was again found. The two-mode model expressed in Equation 6.2 predicts a $\Delta\nu_s^{-1/2}$ dependence, and the data points are fitted to this curve. Again, the data matches the simple model quite well.




Figure 6.3. Minimum magnetic field strength versus the synchrotron tune. The data shown was generated with $\nu_x = 0.585$, $\nu_y = 0.575$, $\Delta\nu = -0.01$, $N_{\bar{p}} = 6 \times 10^{10}$, and $\sigma_{\bar{p}} = 0.7$ mm. A fitted function $B_{\min} = k\Delta\nu_s^{-1/2}$ produced a constant $k = 0.557$ and is also drawn.

The Tevatron’s current operating parameters include a synchrotron tune of about 0.0007, where the data and curve are approaching 30 kG. The operation of the T_{FL} at 35 kG is above this limit, though close enough to be warrant more careful simulations.

Combining the parameters from Figures 6.2 and 6.3, contours of the minimum solenoid strength can be plotted as a function of tuneshift and synchrotron tune. In Figure 6.4, the range of synchrotron tunes associated with Tevatron operation are analyzed more closely. As the synchrotron tune approaches zero, the required field increases to infinity. This limit makes sense, as a bunch that never reverses the head and tail, or reverses them too slowly, will drive the tail of the bunch at the betatron frequency via the skew wake function in Equation 6.1. This driving force is not damped by the reversal




Figure 6.4. Contours of minimum magnetic field strength for a range of T_{EL} tune-shifts and synchrotron tunes. Typical operation sets the tuneshift $\Delta\nu = \sqrt{\Delta\nu_x \Delta\nu_y}$ at less than 0.01 and the synchrotron tune ν_s greater than 0.0003.

of roles, and the tail will be driven into the beam pipe. Therefore the minimum magnetic field in Figure 6.4 explodes as ν_s approaches zero.

In the realistic case where the synchrotron tune is approximately 0.0003 or larger, the simulations predict a tuneshift above 0.015 is necessary to excite TMCI. Typical T_{EL} operation operates with a horizontal tuneshift ν_x no more than 0.01 and a vertical tuneshift ν_y less than 0.003, so the simulations suggest that operation of the T_{EL} is stable.

Equation 6.2 shows a $|\nu_x - \nu_y|^{-1/2}$ dependence for the minimum solenoid field. By adjusting one tune while holding the other fixed yields a strong influence on B_{\min} . The TMCI simulation was run while scanning the horizontal tune from 0.52 to 0.63; the vertical tune was held at 0.575. Figure 6.5 presents the minimum field for stability along




Figure 6.5. Minimum field strength as the horizontal tune is scanned across the vertical tune. In this case, $\nu_y = 0.575$, $\nu_s = 0.001$, $\Delta\nu = -0.01$, $N_{\bar{p}} = 6 \times 10^{10}$, and $\sigma_{\bar{p}} = 0.7$ mm.

this scan. It is clear that when the tunes are close ($|\nu_x - \nu_y|$ is less than approximately 0.015) the ability to suppress TMCI is extremely difficult. A fitted curve following the $|\nu_x - \nu_y|^{-1/2}$ dependence is drawn also, and this curve again connects the simulation results with the analytic theory behind Equation 6.2.

Farther away from the strong coupled-mode resonance, the simulation levels off, contrary to Equation 6.2. This is expected, as higher-order modes of TMCI take over. Lastly, as the horizontal tune approaches 0.5, the half-integer resonance forces the minimum field strength to increase. These phenomena are still significantly below the operational 35 kG that the T_EL functions at. As long as the horizontal and vertical tunes are not within about 0.01 of each other, the T_EL solenoid does not excite TMCI.

Chapter 7:

Successful results of T_{EL} operation

*Understanding is a three edged sword:
your side, their side, and the truth.
– Captain John Sheridan*

The measurements detailed in Chapter 5 focus on internal aspects of the T_{EL}. Chapter 6 proves that undesirable interactions between the electron beam and Tevatron bunches will not manifest. In this chapter, the successful results of T_{EL} are presented. Section 7.1 measures tunes shift with respect to several parameters, and for much of the discussion, the flattop gun was used, as its linear focusing yielded the largest and clearest shifts.

Because the edges of the flattop beam profile tended to interfere with long lifetimes, the Gaussian gun replaced the flattop gun. In Section 7.2, a comparison between the

two guns is given; it is with the Gaussian gun that the slowest emittance growths were observed.

Due to the time required to produce antiproton bunches, outlined in Section 1.1, many tests were performed on proton bunches. The dynamics are similar, but the electric and magnetic forces that the electron beam imparts on proton bunches is different. The electrostatic force is negative, since the protons are positively charged and attracted to the electron beam. The proton bunch travels in the same direction as the electron beam, so their associated currents are in opposite directions. Therefore, the force due to the magnetic field is outward, or positive, weakening the effect of the electric field. Hence, any tuneshift that the TEL performs on a test proton bunch is less than the same TEL acting on a corresponding antiproton bunch.

Of course, this can be quantified by copying the derivation of TEL-induced tuneshift given in Section 3.2. In particular, the factor $(1 + \beta_e)$ in Equation 3.4 adds the electric and magnetic forces together when acting on antiprotons, since both terms are defocusing and therefore of the same sign. Since they are of opposite sign when acting on proton bunches, the tuneshift is instead proportional to $(1 - \beta_e)$. If the electron beam travels at a speed $0.2c$, the effective antiproton tuneshift from a measured proton tuneshift can be found by multiplying the former by $-(1 + \beta_e)/(1 - \beta_e) = -1.5$.

7.1 Tuneshift measurement

The Tevatron has had finely tuned Schottky detectors that measure the average betatron tune of both protons and antiprotons. In Figure 7.1, two spectra are shown,




Figure 7.1. Horizontal Schottky spectra with and without the T_{EL} acting on a single proton bunch. The left set of peaks is crosstalk from the vertical tune, while the right set are the significantly changed horizontal tune.

each of which shows the horizontal tune spectra of a single proton bunch that was injected into the Tevatron.

In the lower trace, the electron beam in the T_{EL} is turned off, and the observed spectra is simply the horizontal signature of the tune. The peaks on the right side of the lower trace shows peaks, which are clearly centered around a tune of 0.5831 (the actual tune is 20.5831, but Section 2.2 explains that the integer is typically ignored). The peaks

on the left side actually bleed in from the vertical tune. Decreasing the crosstalk between the two dimensions has required serious effort.

The trigger of the T_{EL} was then timed to coincide with the bunch passage, the cathode voltage was set to -10 kV, and the collector-to-cathode voltage was +5 kV. So far, the anode was not being pulsed, so there was no beam current. The magnetic fields in the gun, main, and collector solenoids were 3.8, 35, and 3.8 kG respectively, corresponding to a beam radius of 1.6 mm. After adjusting the correctors to center the electron beam onto the proton bunch, the anode was pulsed so that the current slowly increased. During this process, the horizontal spectra slowly shifted to the right. The upper plot in Figure 7.1 shows the new spectra at a full current of 2.03 A. The fact that the new central peak is centered around 0.5868 is proof that the T_{EL} successfully tuneshifted the proton bunch. The lifetime of the bunch, defined later in subsequent sections, was measured to be about twenty hours. This is just an example of many successful attempts to clearly tuneshift proton and antiproton bunches.

Additional features in Figure 7.1 are worthy of mention. The area under each of the traces represents the amount of transverse energy being transferred to the particles, which leads to higher losses and emittance growth. This area is known as the total *Schottky power*, and it is significantly larger when the T_{EL} is operating on the bunch. This fact suggests that the T_{EL} is exciting betatron motion in the particles. Section 7.3 delves into the measured impact of T_{EL} noise.

7.1.1 Comparing tuneshift with theory


The expected tuneshift caused by an electron beam is calculated in Section 3.2. Since tuneshift is the immediate goal of the T_{EL}, the equation for tuneshift given in

Equation 3.2 is the first to experimentally verify. Over the course of testing the performance of the TEL, many measurements of tuneshift were performed while adjusting its different parameters.

The relation between tuneshift and electron-beam current is important and relatively straightforward to test. This connection was measured a number of times, a few of which are shown in Figure 7.2. Each time, the translational position of the TEL needed to be optimized in order to produce the maximum tuneshift. Chapter 5 demonstrates the complexities involved in aligning the electron beam well, and it describes the problems associated with trusting the beam-position monitors (BPMs). Instead, adjustments to the dipole correctors were made while maximizing the observed tuneshift. For a number of reasons, attaining the best alignment proved difficult.

The experimental data often appeared close to the values expected by theory. During some experiments, mostly during the first tests, the experimental tuneshift were seen to be much less than theoretical predictions. The cause for this large discrepancy is blamed on poor translational alignment. As experience developed over three years of experimentation, large differences were rarely seen.

In Figure 7.2, the data from four select experiments is shown. In all of these cases, the flattop gun was used with a beam radius of 1.6 mm. The experiment labeled as Experiment #4 is the earliest data plotted in Figure 7.2, but is in fact the fourth experimental session with the TEL after it was installed in the Tevatron. As the table shows, the Tevatron's horizontal tune was set to 0.5825 before the electron-beam current was turned on, and the TEL's cathode voltage was set to -7.5 kV. As the current was increased, the tune shifted by the shown amounts. While the amount of tuneshift was significant, the



Experiment	V_{cat}	ν_0	$\Delta\nu_{\text{max}}$
#4	-7.5 kV	0.5825	0.0062
#14	-7.5 kV	0.5642	0.0089
#18	-8 kV	0.5832	0.0067
#25	-4.7 kV	0.5832	0.0044

Figure 7.2. Observed horizontal tuneshift as it depends on current, over several separate experiments. Theoretical predictions, following Equation 4.3, are plotted as grey lines for each relevant cathode potential. Between each data point, the current was adjusted slowly to allow the bunch’s tune to change adiabatically. It is important to note that in each of these experiments, the gun and main solenoids had respective field strengths of 3.8 and 35 kG, so the electron beam had a radius of 1.6 mm.

amount was significantly below the expected amount, which is expressed by the grey line labeled Theory #4.

In Experiment #14, the same test was performed. First, the T_{EL} was aligned to the proton orbit, then the cathode voltage was set to -7.5 kV. Tuneshifts were recorded as the current was ramped up and down, with the highest attained shift of 0.0089, which corresponds to an antiproton tuneshift of 0.0112. While the base tune in Experiment #14 was 0.5642, significantly different than that of Experiment #4, the difference in results is not attributable to that. Instead, it is believed that the electron beam was much better aligned to the proton orbit, causing a stronger tuneshift for the same current. The tuneshift of 0.0089 on a proton bunch is the largest observed tuneshift during all of the operation of the T_{EL}, and the equivalent tuneshift of 0.0112 on an antiproton bunch achieves and surpasses the requirements of successful tuneshift compensation.

Additionally, Experiment #14 provides tuneshifts comparable with theoretical prediction, with significant departure only at its highest current levels. The model follows Equation 3.4, including the description of the electron-beam speed discussed in Section 4.3. This theory says that, up to the current accepted by the beam pipe, the current density and speed is uniform across the beam radius. In that scenario, the tuneshift curves upward as the larger charge density is enhanced by the reduction of electron speed.

In Figure 4.7, the center of the electron beam is suggested to be at a lower potential than the edge, and therefore the electrons at the center are somewhat slower and more dense than those on the edge. Since the majority of the protons in the affected bunch have small betatron amplitudes with respect to the electron-beam radius, they sense the

center of the electron beam more than the edge. This could explain why the data of Experiment #14 below 1.5 A is slightly higher than the prediction of the theory, which does not tackle the nonlinear potential inside the electron beam.

More interesting is the divergence of the data and the theory at large currents. Without consideration of the space-charge potential within the electron beam, the current should follow the Child-Langmuir law, given in Equation 4.6, until the acceptance of the beam pipe limits it. A plot of current versus anode voltage under this scenario would show a $3/2$ -power relation to a sharp corner, beyond which the current would remain constant. Instead, Figure 5.18 demonstrates a very gentle rolling-off of current.

This phenomenon, analyzed in Section 5.3, agrees with the idea that the center of the electron beam requires significantly higher energy than the edge. Because of this, the center is partially rejected even at moderate currents. Since the protons typically witness the center of the beam more than the edges, their tunes are not shifted as much as the uniform-density model predicts.

Other T_{EL} parameters were explored in other experiments. Two such cases are also shown in Figure 7.2. In Experiment #18, the cathode potential was set to -8 kV and tuneshift measurements from the base value of 0.5832 produced a curve less than, but close to, the corresponding theoretical prediction. Other than a small upswing at its highest current, the data exhibits similar features as Experiment #14. It is interesting to note that both of these examples yielded somewhat linear relationships between observed tuneshift and current, which is discussed in Section 7.4.

Experiment #25 displays results when the cathode voltage was set to -4.7 kV, significantly below its value during most of the other experiments. Due to the slowly-moving

electrons, a significantly higher charge-density yielded large tuneshifts even at low current levels. This experiment impressively matches the theoretical prediction. If more current was attainable during that experiment, it is expected to drift away from the theoretical model, since electron rejection is a significant problem at low cathode potentials. This prediction is justified in Figure 5.18, where, at currents as low as 1 A, the beam pipe is unable to accept the total available current when the cathode potential is -4 kV or -5 kV .

7.1.2 Translational alignment

Much of the discussion in this dissertation stems from the difficulty of producing a large but dense electron beam. The net result is that, in the interaction region, the electron-beam radius is significantly larger than the rms bunch radius. Nevertheless, it is small enough that a little error in position can cause the electron beam to miss the bunch at least partially. For example, if the head of the electron beam is centered on the bunch and it travels at a mere 1 mrad angle to the bunch's orbit, the centers will be 2 mm apart, which is larger than the beam radius.

Figure 7.3 sketches the effects of a displaced electron beam on a proton bunch. The left diagram illustrates the horizontal force exerted on protons in the ideal case, where the electron beam and the proton bunch are perfectly collinear. Protons initially to the right of the center receive a negative kick, while protons to the left receive a positive kick. A dashed line passing through the center of the proton bunch and tangent to the force represents the linear tuneshift, as has been analyzed in Section 3.2. The slope of this line correlates to the magnitude of the shift.




Figure 7.3. Illustration of how displacements can cause positive or negative tune shifts. A proton bunch centered on the electron beam feels a maximum positive tuneshift (left). A vertical displacement generates a lesser tuneshift (middle), but a horizontal displacement can create a negative tuneshift (right).

The middle diagram in Figure 7.3 illustrates the horizontal force protons feel when the electron beam is significantly displaced in the vertical direction. Even though the protons are mostly outside the volume of the electron beam, they still witness a focusing force in the horizontal direction. The magnitude of this force, however, sharply reduced, yielding a smaller tuneshift. For proton bunches, both scenarios exhibit positive tuneshifts.

A horizontally misplaced electron beam creates a more interesting situation, as shown on the right side of Figure 7.3. In this case, all of the protons feel an attractive force toward the electron beam (toward $-\hat{x}$), but the protons closest to the electron beam are more strongly pulled than the others; the bunch is defocused. This can be interpreted by the force diagram. The proton bunch in this case passes through the force field where




Figure 7.4. Measured tuneshift as the electron-beam path was displaced from the proton-bunch orbit. The horizontal tune was measured as the electron beam was steered vertically (solid circles), then horizontally (open circles). Gaps in the horizontal data were due to indecipherable Schottky spectra.

the slope is of opposite sign (the force scales with $1/r$ instead of with r). The linear tuneshift, represented by the dashed tangent line, is now negative.

Thus, the proton bunch is expected to demonstrate a positive horizontal tuneshift when it is well-aligned with the electron beam, or if it is simply vertically offset, but if it is significantly misaligned horizontally, a negative tuneshift should be observed. In Figure 7.4, this hypothesis is tested and proven correct. An electron beam, initially centered on a proton bunch, was steered vertically while the tune was monitored. After it was re-centered, it was steered in the horizontal direction. Throughout this time, the horizontal tune was measured and plotted.

The plateau at the center of the figure indicates positions where the flattop electron beam surrounded the proton bunch. With a current of 2.5 A and a cathode voltage

of -7 kV , the simplified theoretical model predicts a horizontal tuneshift of over 0.010, but Figure 7.2 predicts that the T_{EL} will only shift approximately 0.009. Indeed, the highest tuneshift in Figure 7.4 is 0.0088, with most of the plateau resting at about 0.0085.

The solid dots in Figure 7.4 show a sudden decrease in tuneshift as the electron beam is displaced more than its radius of 1.6 mm. As the separation distance is increased, the tune asymptotically approaches the unshifted base tune (which was 0.5644 during the experiment). Between each datum, the electron beam was moved approximately 0.57 mm; however, at positions where the proton bunch was scraping the edges of the electron beam, the forces acting on the bunch are highly nonlinear. The Schottky spectra under these circumstances are unable to measure a well-defined tune, and these displacement values are skipped.

The open-circle data shows the horizontal tuneshift while the beam was displaced horizontally. Each increment was about 0.29 mm, but large regions failed to produce a clear, distinct proton tune. Nevertheless, the separate regions reveal that the tuneshift does in fact become negative when the proton bunch passes outside the electron beam; Figure 7.3 predicted this effect. At large separations, the tune again approaches the base tune, but when the bunch is near the edge of the electron beam, the nonlinear space-charge forces prevent proper measurement of the proton tune.

An additional conclusion can be drawn from Figure 7.4. Within the plateau region, both vertical and horizontal changes suggest that there is a slight decrease in tuneshift in the exact center of the electron beam, with larger shifts observed closer to the beam edges. This corroborates with the understanding in Figure 7.2, where at high currents, the center of the electron beam begins to be rejected, decreasing the space-charge density

in this region. Due to this theoretical understanding, a lower tuneshift is expected at the center of the beam, which is what is observed in Figure 7.4.

The electron beam can be angled with respect to the closed orbit instead of simply offset. The effect of such a misalignment is less severe, since the effective tuneshift becomes an average of the derivative of the electric field over the relevant extent of the transverse range. For example, if the angle is small enough such that the head of a 1.6 mm-radius electron beam is horizontally displaced by +1 mm and the tail is displaced by -1 mm, then a small bunch would still exhibit the same horizontal tuneshift, since the derivative of the field is constant within the beam. If the angle was increased beyond the edge of the electron beam, then part of the interaction region would defocus the bunch while another part would focus it. The net effect for moderate angles is still a positive tuneshift for protons, though of a decreased magnitude.

Evidence supporting this claim is shown in Figure 7.5. This data was taken by more sophisticated manipulation of the dipole correctors within the main solenoid. Described in Section 5.1, it is possible to pivot the electron beam around its center by simultaneously changing the upstream, long, and downstream correctors. Figure 5.3 in particular illustrates two distinct electron-beam paths, one slightly rotated around the center of the main solenoid from the other.

In Figure 7.5, the correctors were adjusted to cover 2 mrad of total pivoting range. This experiment was performed during a different study period than the previous tests. At 3 A of electron-beam current and -7.5 kV on the cathode, a maximum of 0.0036 horizontal tuneshift was observed in the proton bunch during this test. Over a range of 1 mrad, the tuneshift remained at this maximum (the center of which was defined




Figure 7.5. Measurement of tuneshift at various angles of the electron beam. The beam path was angled around the center of the main solenoid. This required coordinated adjustments of the upstream correctors, long correctors, and downstream correctors.

as zero), but at larger angles, the shift began to decrease. The extent of this plateau corresponds to the head and tail regions of the interaction length each moving by one millimeter, a value significantly less than the diameter of the electron beam. A vertical displacement or angle would have the effect of decreasing the range of maximum tuneshift in the horizontal direction. This possibility is also supported by the fact that the maximum observed tuneshift is noticeably less than what is expected for the electron-beam current.

The Schottky detectors are quite sensitive to the alignment of the electron beam to the orbit of interacting bunches. Without the T_{EL}, the Schottky spectra of the protons or antiprotons appears as a small number of peaks clustered around the nominal tune. Separate detectors in the horizontal and vertical planes provide good separation of the signals in either direction. When the T_{EL} is optimally positioned on a bunch, the spectra generally looks like the left example in Figure 7.6. First, the set of peaks slides to a new




Figure 7.6. Two examples of Schottky spectra. The left data show the horizontal and vertical spectra of a well-aligned, one-Ampère electron beam acting on a proton bunch, while the right data indicate that the T_{FL} is poorly aligned to the bunch. While the tune peaks are still visible, the higher power (up to -40 dBm, instead of -53 dBm) indicates significant emittance growth. Worse alignment obscure the peaks altogether.

position, providing proof positive that the T_{FL} is indeed shifting the tune. Additionally, the height of the peaks and the surrounding baseline increases in magnitude, which implies that the particles are in general oscillating with larger betatron amplitudes. When the T_{FL} is centered correctly, the total Schottky power, represented by the area under the peaks and baseline, does not increase dramatically.

When the displacement is not optimized, however, the increase in total power is more significant. The right set of data in Figure 7.6 shows much more power being deposited into the proton bunch. The T_{FL} was unchanged between the two cases, except the beam was translated by about two millimeters. The tunes shift in both planes has

decreased, and the height of the peaks has risen. The horizontal data shows the most extreme difference. The pedestal on the left side sits at about -69 dBm, while on the right, it has become -62 dBm. The peaks reach to -53 dBm of power in the optimal case and -40 dBm in the suboptimal case.

If the alignment is worse, or if there is a higher level of fluctuations in the current, the pedestal can swallow the peaks altogether, so that a reading of a specific tune becomes impossible. This was the case during portions of the displacement experiment shown in Figure 7.4, for example. Additionally, the tune spectra produced by the Schottky detectors during collisions is significantly more complicated, due to the inclusion of antiprotons and protons and the complex distribution of particle behavior. The use of Schottky detectors to verify proper T_{EL} performance is not realistic.

7.1.3 Additional tuneshift dependencies

The previous discussion scrutinized the consequences of a bunch passing through the electron beam at different transverse positions. The bunch must also pass through the electron beam at a time when the beam current is maximized. Since the anode modulator is pulsed by a delay circuit after the Tevatron clock signal, it is possible to observe the tuneshift while adjusting the delay pulse over a wide range.

When the delay is very short, the electron beam fires before the bunch arrives at the location of the T_{EL}, resulting in zero tuneshift. If the delay time is increased slightly, the electron bunch will have fired too early, but the bunch will pass through the tail end of the pulse, generating a small amount of tuneshift. Increasing the delay further allows the bunch passes through the optimal time when the full electron-beam current fills the




Figure 7.7. Tuneshift caused by a delayed electron pulse. As the trigger timing was adjusted, the bunch witnessed either the head or the tail of the electron pulse. Again, changes were adiabatic to prevent bunch blowup.

interaction region. At larger delays, the bunch arrives early and witnesses just the head of the electron beam.

In short, sweeping the delay and recording the tuneshift ought to map the longitudinal profile of the electron beam. Figure 7.7 plots data under these conditions; the delay circuit was swept over a wide range of time while the tuneshift mapped out the shape of the electron pulse. Since a larger delay time correlates to the front of the pulse, the outline is horizontally reverse of the typical current pulse.

In order to compare the data to the pulse shape, a shaded rendering of the current profile is reversed and included in Figure 7.7. The data impressively copies the pulse shape, including the small bumps riding on top of the flat plateau. The fact that the bunch can discern these small changes is important in the discussion on fluctuations in Section 7.3.

The tuneshift generated by the TEL also depends on the cathode voltage, as Figure 7.2 suggests. An experiment was performed on the tune of an antiproton bunch




Figure 7.8. Antiproton tuneshift as a function of cathode voltage. During experimentation with antiproton bunches, the voltage applied to the cathode was varied, while the current remained at approximately 1.92 A and the profiler electrode was set to -150 V.

while changing the cathode voltage. The data produced in this experiment is plotted in Figure 7.8. The black data points show a decreasing tuneshift as cathode voltage is increased in magnitude. This is expected, since the tuneshift is linear with the charge density $\rho = I_e/v_e$. An electron beam with 13 keV has a large speed, so for constant current, the charge density is small, producing small tuneshift in the antiproton bunch. An electron beam with only 6 keV of energy has more charge density and generates a larger tuneshift.

It should be noted that the current was observed to decrease slightly at the lower cathode potentials. This current, also plotted in Figure 7.8, ranged from 1.75 A at the least cathode potential to 1.92 A at the largest. Additionally, the profiler electrode was set to -150 V with respect to the cathode potential. Shown in Figure 5.11, the effect of the profiler is to eliminate current emission from the edge of the cathode. This reduces

the radius of the beam and the total current output, but has negligible effect on the current density near the center. The effective reduction of the radius can be determined from Figure 5.12. Setting the profiler to -250 V from zero reduces the effective area by 4.1 % or the beam radius by 2.0 %.

Using the correlated cathode voltage, beam current, and beam radius, Equation 3.4 generates a theoretical prediction for the optimal tunes that the antiproton bunch can experience. The upper grey curve in Figure 7.8 plots this relation. As previous discussion have concluded, it is difficult to align the electron beam perfectly with the orbit, so a small reduction of the theoretical prediction is not unexpected. The lower grey curve is simply the original theory multiplied by 0.9. This curve fits the data impressively well at cathode voltages beyond -7 kV , while at -6 kV and -7 kV , the difference becomes significant.

An explanation for the data's departure from theory is similar to that of Figure 7.2; namely, when the cathode voltage is extremely low, the space-charge potential is large enough that some of the current in the center of the beam cannot penetrate the beam pipe. Since the majority of antiprotons in a bunch have small betatron amplitudes with respect to the radius of the electron beam, they only sense the charge density near the beam's center. With less current at the center, the electron beam cannot shift the tune of the bunch as much as Equation 3.4 predicts.

7.2 Successful operation

During many study periods, the shifting of the tune of proton or antiproton bunches has been quite successful. The results shown in Figure 7.2 is a sampling of many quality experiments. The dependability and alacrity with which TEL parameters could be

optimized increased with the accrument of experience. However, a number of attempts indicates that a significant number of particles in the bunches were being lost as the TEL intercepted them. The lifetime of these bunches was noticeably reduced, which has an obviously adverse effect on the eventual utility of the TEL.

7.2.1 Controlling particle losses

Typically the lifetime of proton bunches in the Tevatron, before collisions and without the TEL, is on the order of 150–300 hours. As the bunch traverses the Tevatron ring billions of times, several mechanisms contribute to a gradual growth in emittance of about 0.5π mm-mrad/hr. These include: residual-gas scattering, intra-beam scattering (IBS), and fluctuations in the ring elements. As the bunch size increases, particles gradually diffuse to larger oscillation amplitudes until they finally collide with some aperture restriction, usually one of many adjustable collimators inserted in the Tevatron beam pipe.

Each time a particle intercepts a collimator, it emits radiation due to the loss of its forward momentum. X-ray detectors, known as *beam-loss monitors*, are strategically positioned near the collimators such that they detect most of these pulses of radiation. The readouts from these monitors report a frequency at which X-ray pulses are registered, but reliable gating techniques allow losses from each individual proton and antiproton bunch to be monitored independently. Lifetimes of specific bunches can be observed via these detectors, and the results of certain actions can be observed relatively quickly.

Figure 7.9 presents the results of an experiment where the rate of losses, as indicated by the beam-loss monitors, was measured as the TEL current was changed. This test was




Figure 7.9. Typical bunch lifetime of a single proton bunch interacting with the TEL electron beam. This data is for the flattop gun, and it singlehandedly indicates the necessity of switching to a smoother beam profile, as discussed in Section 3.4. Shortly after this problem was noted, the Gaussian gun was installed.

performed on a single proton bunch, and the cathode voltage was -10 kV. The losses varied from about 250 Hz at low currents to 1 kHz at the highest currents. At zero current, the average loss rate was approximately 230 Hz over a large portion of an hour.

Several features are apparent in Figure 7.9. The losses data were converted into lifetimes, producing more tangible results. This translation was straightforward, since the lifetime $\tau = -cN/(dN/dt)$, where N was the current total number of particles and dN/dt was the loss rate measured by the beam-loss monitors. The constant c was determined from a calibration test. During a period of a couple hours, the bunch was allowed to proceed without any changes made to the TEL. After this amount of time, the number of particles had diminished enough to directly compute the lifetime. Since the number of particles and the average loss rate were also known, the constant could be derived.

The maximum current in this experiment was about 0.75 A, far less than the two or more Ampères that is intended for useful beam-beam compensation. Nevertheless,

the proton bunch's lifetime at the higher electron-beam currents is less than fifty hours, significantly less than the typical 175-hour lifetime without interference. While it was impossible to guarantee that the electron beam was perfectly centered on the proton orbit, adjustments of the beam's position yielded no improvement in the bunch's lifetime.

In order to predict the relationship between beam current and bunch lifetime, the sources of the losses must be understood. Without the T_{EL}, a nonzero loss rate can be measured. The mechanisms behind these losses, previously listed in this discussion, are unaffected by the action of the T_{EL}. Therefore, the loss rate is expected to have a constant level to which the effects of the T_{EL} add.

If the electron beam was perfectly uniform and had a radius much larger than the bunch size, then the losses due to the T_{EL} are expected to be negligible. It is the edges of the flattop beam profile that is expected to effectively eliminate particles of large betatron amplitude. If any particle that gains sufficient amplitude is assumed to be removed quickly, the loss rate due to the T_{EL} then corresponds directly to the average expansion of the particle betatron amplitude. In Section 3.3, the emittance growth rate was shown to follow the square of fluctuations in the electron-beam current, which is usually a constant percentage of the total current. Since the emittance is defined to be an average of the particles' amplitudes, it follows that the loss rate due to the T_{EL} should depend on the square of the current.

The measured losses data was fitted with a curve consisting of a constant (non-T_{EL}) term and a quadratic (T_{EL}) term. The curve plotted in Figure 7.9 is the inversion of this curve to represent lifetime. The data follows this trend very well, suggesting that the fluctuation level of the electron beam is a major contributor to particle losses during experimentation.




Figure 7.10. Analysis of how bunch losses compare with displacement of the electron beam with the intercepted bunch. As expected, the losses decrease drastically as the electron beam is far away from the bunch. However, a local minimum was expected at the ideal center position ($x = 0$), which is not witnessed.

A different experiment, shown in Figure 7.10, measures the loss rate with respect to the transverse position of the T_{FL} electron beam. If the beam and the proton orbit are perfectly aligned, the core protons would be tune-shifted but stable. Only the most extreme protons could sense the nonlinear edges of the flattop electron beam and be lost. At this position, the average losses are expected to be relatively small.

If, however, the electron beam is slightly displaced transversely, the core protons would pass near the edge of the electron beam, generating much larger losses. Finally, as the electron beam is positioned further from the proton orbit, fewer protons would sense much of its fields at all, again reducing losses. Hence, the shape of the losses was expected to resemble a hill with a narrow hole in its center.

While the results in Figure 7.10 indicated a significant increase in the loss rate when the electron beam was intercepting the proton bunches, there was no indication of a stable center position. The center position ($x = 0$) was determined by the position of the largest tuneshift. However, at that position, the loss rate also reached its maximum. The lack of a central stability region can be explained by an imperfect orientation either in the other direction or an angle.

7.2.2 Comparison of electron guns

The last discussion presented more examples of the severe constraints that the flat-top edges place on the performance of the TEL. The Gaussian gun was designed to possess much smoother edges so that particles at large betatron amplitudes would not feel strongly nonlinear space-charge forces. Section 5.2 describes the differences between the two guns. In order to quantify the differences between them, a scan of working points was performed with each of the guns. In this test, the Tevatron horizontal and vertical tunes were independently adjusted to cover approximately a 0.020 span in both dimensions. By adjusting the tunes in 0.002 increments, the loss rate was measured, recorded, converted to a lifetime, and plotted in Figure 7.11.

In order to simplify the understanding, both guns were set to currents such that the horizontal tuneshift was 0.004 and the vertical was 0.0013. The Tevatron is equipped with tune-adjustment parameters, which provided a convenient way to adjust the tune. Confirmation of the correct tune was sometimes possible with the Schottky detectors, but when the loss rate was high, an accurate measurement of the tune was difficult to determine.




Figure 7.11. Scans of working points with the T_{EL} shifting the Tevatron tune. Both of the flattop (top) and Gaussian (bottom) guns were set to push the tune of a single proton bunch by about 0.004 horizontally. Data points generating these greyscale maps were taken at 0.002 intervals in each direction.

Whenever the tunes were adjusted, a short amount of time was needed before the loss rate stabilized. Sometimes it reacted quickly, while other times required a longer period before a specific loss rate could be determined. The number of protons in the test bunch was measured throughout the experiment period, and a calibration test was performed as mentioned in the previous section. This allowed the loss-rate data to be changed into lifetimes as shown in Figure 7.11.

The shaded scale shown on the right side of the scans indicates the lifetime, in hours, witnessed at each data point. In order to more effectively convey the regions of high and low lifetime, a two-dimensional interpolation algorithm turned the individual data points into a smooth, shaded surface. Individual isocurves are also drawn at multiples of 20 hours.

Unfortunately, the regions covered by the two scans do not span the exact same tune space, but significant overlapping allows significant differences between the flattop and the Gaussian guns to become apparent. In the upper scan, the flattop gun usually produced poor lifetimes. This implies that the T_{EL} flattop gun tended to excite oscillations in at least some portion of the bunch particles, and recorded lifetimes were mostly less than ten hours.

However, along a diagonal region near the main diagonal ($\nu_x = \nu_y$), there is a relatively consistent pattern of lower losses. Along this strip, lifetimes as high as seventy hours were observed, almost as high as the lifetime of the bunch without the T_{EL} operational. This achievement suggests that fluctuations in the T_{EL} current can be small enough that the bunch lifetime is nearly acceptable. However, this is only observed in a very small portion of the tested tune space.

It is important to notice that if the losses incurred by the TEL operated by a scattering process or driven oscillations by TEL fluctuations (steering, current, ion instabilities, etc.), losses would be significant regardless of the bunch's tune. The particles being driven would be lost due to the mechanism independent of the Tevatron settings. Instead, the observation of low losses at particular tunes indicate that the losses are driven by a resonant phenomenon. While fluctuations can still contribute to the process, the particles must be excited at select resonant frequencies, as described in Section 3.4.

The large regions of low lifetime again supports the hypothesis that the flattop electron beam is adversely affecting protons. The outlying particles, witnessing strongly nonlinear focusing forces from the electron beam, do not survive as long as the core particles. Through the majority of the tested tune space, these particles are able to escape very quickly, and the gradual emittance growth of the core protons constantly feeds these losses. Only in a small working-point region do the outlying particles not escape so nimbly, slowing the rate at which particles are lost.

The second scan in Figure 7.11 shows the massive difference that the Gaussian gun had on the lifetime. The highest measured lifetimes were around 130 hours, almost indistinguishable from the bunch's lifetime without the TEL. Much larger regions of lifetimes over twenty hours can also be seen.

The fact that the highest lifetimes are nearly the same as the unperturbed proton bunch bolsters the idea that TEL fluctuations cannot, by themselves, remove particles from the bunch completely. Instead, it is believed that the fluctuations contribute to a gradual emittance growth, and because there are no strongly nonlinear edges to the electron beam, the protons are still stable at larger orbits. This explains why a much

larger percentage of the tested tune space offered moderate lifetimes over that of the flattop gun and the best lifetimes using the Gaussian gun significantly outlast those of the flattop.

It is interesting to note that the diagonal region of stability for the flattop gun does not appear in the Gaussian-gun case. It is known that the particular resonance in question ($\nu_x = \nu_y$) requires a minimum amount of coupling between the two degrees of freedom[68]. While relatively weak, this resonance has been observed in Tevatron studies, and the regions are avoided during operations[3]. In general, a coupling resonance is effective because motion in one transverse plane encourages motion in the other. It is possible that the flattop gun's highly nonlinear forces disrupts this correlation too much and disables this particular resonance, while the Gaussian gun is gentle enough that the protons are able to resonate. More experiments are necessary to test this conjecture.

7.2.3 Successful reduction of emittance growth

The previous section presented the exciting success of the T_{FL}, producing a strong tuneshift without compromising the lifetime significantly. However, the initial goals of the T_{FL}, stated in Section 3.1, were to compensate for the tuneshift and possibly the spread of the antiproton bunches incurred by the protons. The underlying issue is the fact that bunch lifetimes are already lowered by the beam-beam interactions, and that the T_{FL} should mitigate this malady by *increasing* the lifetime of the afflicted bunches. Can the T_{FL} improve the performance of the Tevatron during collisions?

When a batch of protons and antiprotons first collide, the emittances of the antiprotons tend to increase significantly. This initial emittance growth is studied in Section 2.4,

but after the first period of bunch expansion, lasting on the order of an hour, the bunches reach a more stable size and suffer lower losses. Since the Tevatron is typically operated such that stores last for a long time, it is difficult to verify that the T_{FL} might be improving the store's performance by measuring the lifetime over the latter portion of the store. Additionally, the period of time it takes to measure lifetime scales with the lifetime, so doing T_{FL} experiments at the end of a store, when emittance-growth lifetimes are at their longest, would take an extremely long time to produce a significant result.

Instead, the T_{FL} was operated during the first half hour of collisions, acting on a single bunch, while the bunch's size was monitored. The goal was to significantly decrease the growth of the particular bunch with respect to its sibling bunches (the equivalent bunches in the other two trains), or with respect to the same bunch in other, similar stores. This feat would be a successful instance of tuneshift compensation.

Several attempts were made to test this ability, each on a new store during the first short period. These tests, along with pertinent parameters, are shown in Table 7.1, where each row represents a different store (the designated store number is listed down the left column).

In the three stores listed without the T_{FL} in Table 7.1, the growth of bunch A33 is similar or larger than that of its siblings. In stores #2546 and #2549, it is still larger. However, in store #2540, the growth of bunch A33 is significantly less than that of the other two bunches. The size, as a function of time, is shown in Figure 7.12. The data starts just after collisions began and the T_{FL} was applied to A33. While the other two bunches significantly increased in size over the next thirty minutes, A33 did not increase as much.

<i>store number</i>	<i>duration</i>	<i>A9 growth</i>	<i>A21 growth</i>	<i>A33 growth</i>	
#2536	40 min	9.9	9.2	9.3	
#2538	35 min	1.9	1.7	2.8	
#2540	34 min	4.1	2.2	1.0	T _{EL} on
#2546	30 min	3.9	1.9	4.0	T _{EL} on
#2549	26 min	4.5	3.6	7.1	T _{EL} on
#2551	34 min	6.7	6.6	7.0	

Table 7.1: Growth of the rms vertical bunch size in the first portion of several stores. All of the growth numbers are in units of π mm-mrad. For the indicated stores, the T_{EL} was acting on bunch A33.

While the other stores are analyzed in the same way, with the results shown in Table 7.1, it is useful to visualize the bunch sizes in a similar store without the T_{EL}. Figure 7.13 plots the same parameters for the same bunches in the previous store, #2538. In this situation, the three sibling bunches all increase at nearly the same rate. The difference between the two consecutive stores is considerable, and the only intentional difference is the application of the T_{EL}.

After the initial half hour of each store listed in Table 7.1, the growth of each bunch decreased significantly; it is understood that after the antiproton bunches reach a certain size, their interaction with the proton bunches at collisions is much more benign. This activity was witnessed regardless of whether the T_{EL} was operated or not, and decreasing this already minimal growth rate would be very difficult.




Figure 7.12. Evolution of three bunch sizes during the first thirty minutes of collisions in store #2540. A line was fitted to each of the bunches to determine their average growth rate. Antiproton bunch A33 clearly grows at a much slower pace than its siblings.

The scope trace shown in Figure 7.14 shows the exact timing situation of the T_{EL} during store #2540. The lower trace shows as raw BPM-plate signal inside the T_{EL} . Four proton signals are clearly marked, with antiproton signals occurring roughly 150 nsec after each proton pulse.

The upper trace is the time-integral of the BPM plate signal. The large rise and fall is the electron beam pulse, and the short negative spikes are the proton bunches passing. The integrator has difficulty tracking the fast but small antiproton pulses, but following the same time delay as seen on the raw BPM signal, one antiproton pulse, namely that




Figure 7.13. Evolution of the same bunches as Figure 7.12, but for the previous store (#2538) and without the T_{EL} . A linear fit of each bunch size shows the average growth of each bunch to be nearly the same. The bunch size of A9 and A22 start growing quickly before leveling off.

of bunch A33, intercepts the electron beam shortly before the latter's falling edge.

7.3 Other observations

So far in this chapter, results directly related to expressed goals of the T_{EL} have been shown. A number of uses for the T_{EL} have become apparent, mostly involving shifting the tune of antiprotons or protons, increasing the lifetime of bunches in the Tevatron, or shaving unwanted particles out of the Tevatron ring.




Figure 7.14. Display of the bunches passing through the T_{EL} electron beam pulse. The upper trace shows an integrated BPM plate, the lower shows a raw BPM plate signal. The slowly rising and falling electron beam intercepts several proton and antiproton pulses, but only the antiproton pulses are transversely centered on the electron beam.

The interaction between the T_{EL} and the Tevatron possess a number of other interesting features, which are discussed here.

7.3.1 Gentle bunch collimation

The flattop gun is shown in Figure 7.11 to typically cause higher losses than the Gaussian gun. The major cause for this difference is the fact that the flattop beam profile, shown in Figure 5.11, has such sharp edges. Indeed, Section 3.2 shows that the ideal profile for linear compensation is an infinitely large flat distribution. Small-amplitude particles and large-amplitude particles would both feel the same focusing power, and the tunes shift would be independent on a particular particle's betatron amplitude.

Therefore the flattop gun was predicted to safely shift the tune of the core particles — those with small amplitudes. The particles with larger amplitudes would witness the sharp edges and would be driven out of the bunch. The remains, after some amount of time, would be a smaller, lower-emittance bunch. The T_{FL} would act as a “gentle” collimator, coaxing the outlying particles away from the bunch and eventually into the beam pipe wall.

A convenient way to measure this effect is to observe the bunch size as the T_{FL} trims away extraneous particles. In Figure 7.15, one bunch was monitored over a hundred minutes as the T_{FL} shaved the bunch size. The T_{FL} was first set to 1 A of peak current for the first 45 minutes. After a ten-minute respite, the current was increased to 2 A (these settings are shown above the two graphs). After about 85 minutes, the T_{FL} was purposefully missteered in order to observe a “blowup” the bunch sizes.

The upper plot shows the horizontal and vertical beam size measured many times during this process. Also shown is the longitudinal bunch size. While this parameter is largely ignored in this dissertation, coupling between the transverse and longitudinal oscillations suggests that the longitudinal behavior ought to mimic the transverse. Indeed, it does.

The lower plot shows the intensity of the bunch during this process. As shaving particles suggests that the intensity should always go down, it is expected that, under one set of conditions (like the T_{FL} staying at one current), the intensity should exponentially decay. The semi-log graph makes the different rates of decay stand out more clearly.

If the bunch was not being collimated, in other words, particles were driven out irrespective of their betatron amplitude, then the beam sizes would decrease exponentially until all particles are eliminated (and the sizes would be zero). Instead, there is




Figure 7.15. Graphs that monitor the TFL acting as a “gentle collimator.” Above both plots is a description of what was done. First the TFL was set to 1 A peak current. After 45 minutes, it was turned off, and after an additional 10 minutes, was set to 2 A. After 85 minutes, the correctors were changed so that the bunch passed through the edge of the beam. The upper graph shows the bunch sizes decrease as the bunch was slowly shaved, and the lower graph monitors the total number of particles remaining. The horizontal axis (time) is the same for both graphs.

a fast initial decreasing of size, but after about ten minutes, the rate of decrease drops significantly; this implies that the large-amplitude particles have been cleared, and the core is more stable inside the electron beam.

In addition, the increase of the TFL current to 2 A was expected to worsen the bunch-size lifetime, but the smaller bunch was well preserved for the remaining time that the TFL was on and centered. The stability of the bunch size is remarkable, suggesting that the flattop profile was ideal for the small bunch size.

However, it is known that particles inside of a bunch will diffuse through phase space; a collimated edge in a bunch will not last indefinitely[6]. It is expected that even when the bunch size is fixed, a slow attrition of particles will be witnessed. The lower graph in Figure 7.15 illustrates this fact. Again, after a short interval of faster losses, the rate decreases significantly to a smooth exponential (linear on the logarithmic vertical axis). When the T_EL current is doubled, the exponential decay is nearly unchanged. A direct analysis of particle diffusion could be performed from this data.

After the bunch was observed for a while, the T_EL was moved transversely so that the bunch intercepted the edge of the electron beam. As expected, the particles were suddenly feeling extremely nonlinear forces, causing emittance (and size) growth, shown by the bump in the upper plot, and heavy losses, shown by the fast decline of the lower plot.

7.3.2 Effects of fluctuations

Successful antiproton preservation is dependent on a very repeatable electron beam. The tolerable limit of shot-to-shot variations in the peak beam current is shown in Section 3.5 to be around 0.1 % of the total beam current. It is interesting and important to observe the effect that the level of fluctuations has on emittance growth in the antiproton bunches.

The modulator pulse circuit was modified to produce a random-amplitude pulse. This was established by setting an average pulse amplitude modulated by a noise generator. A trace of multiple pulses is shown in Figure 7.16, where the average current is about 500 mA and the rms variation was adjustable from 10 mA to about 150 mA. At different




Figure 7.16. Multiple traces of intentionally varying current amplitudes. First, the average current is established (about half of an Ampère); then, the rms fluctuation level is adjusted.

noise levels, the bunch size is observed long enough to record its growth (and converted into emittance growth). At the same time, the Schottky spectra was monitored, and the total power in the spectra was measured.

Figure 7.17 shows the data of this analysis. According to Section 3.2, the emittance growth is expected to increase as the square of the amplitude fluctuations. The measured data were fitted to an arbitrary quadratic, which they follow adequately. In addition, the Schottky power data, having units of decibels, should scale as the logarithm of the emittance growth, or linearly with the level of fluctuations.

As Figure 7.17 indicates, the power is a little more anemic than fully linear. However, the process by which the test was done allowed the bunch to grow somewhat before the




Figure 7.17. Effects of large current fluctuations on bunch particles. The emittance growth reasonably follows a quadratic curve. The total power measured in the Schottky spectra is also recorded. Ideally, this should be linear, but the effect of bunch growth accounts for the less-than-perfect relation.

final, noisiest data was taken. This means that the total bunch population had decreased before the final power was measured, which lowers the measured Schottky power. It is therefore not surprising that Schottky power does not track as well as the emittance growth; if it were possible to measure the effects of fluctuations without significantly disrupting the bunch, the results should follow theory better.

Section 4.3 explains the inimical complexities of directly measuring shot-to-shot variations in peak current (low-frequency changes, between 0.1 Hz and 1 kHz, are much easier, but they are of lesser consequence to Tevatron bunches). Though the maximum allowed fluctuations are stated in that discussion, a clever technique for estimating that limit uses the curve in Figure 7.17. The Tevatron, without the T_{FL}, has a typical emittance growth of 0.2π – 0.5π mm-mrad/hr. If the T_{FL} is allowed to only enhance the emittance growth




Figure 7.18. Measured Schottky power as the T_{EL} trigger delay is adjusted with respect to the intercepted bunch. The total Schottky power is plotted at each step, and should mimic the derivative (magnitude) of the pulse shape if timing jitter is significant.

by 0.1π mm-mrad/hr, added in quadrature to the Tevatron's inherent emittance growth, Figure 7.17 equates that limit to 3.5 mA rms variation, which is 0.17% of a nominal 2 A-peak current.

Another source of fluctuations is timing jitter. Because the peak of the current pulse is typically quite broad, the timing jitter has shown itself to be rather benign. Exploring this further, the electron beam was delayed various amounts with respect to an impending bunch so that the bunch intersected the rising or falling edge (following the description in Figure 7.7). Instead of waiting for the bunch's emittance to grow at each step, a measurement of the Schottky power was taken. Figure 7.18 displays the results of these measurements.

Although spectra power appears somewhat random, there is a discernible correlation between its peaks and the largest slopes of the pulse shape. For example, the maximum

observed power is 4 dB over the next highest peaks and coincides with the falling edge of the pulse. Along the rising edge, there are two somewhat steeper edges with a more gentle portion between 665 nsec and 675 nsec. The registered power is lower in this region than before or after.

In addition, the small bump in charge density at 645 nsec, observed also in other figures in different sections, can contribute to jitter-induced Schottky power, since a small change in timing still produces a variation in the charge density. Indeed, a large peak in Schottky power is centered directly on the small bump.

An interesting event occurred when a bad connector in the triggering circuit caused the T_{FL} to begin failing intermittently. Large variations in the current, as shown in Figure 7.19, impacted a couple bunches during a Tevatron collider store. Two antiproton bunches, A28 and A29, intercepted the T_{FL} electron beam. Over the course of about two minutes, those two particular bunches were essentially eliminated. The multiple black traces exhibited in Figure 7.19 are BPM signals after being integrated. The electron-beam current is fluctuating significantly, indicated by the large range in signal amplitudes. The BPMs also detect the passage of the proton bunches, shown as negative spikes, and the antiproton bunches, which are small positive spikes after each proton bunch. Even though two proton bunches pass during the electron-beam pulse, they are transversely separated from the electron beam. The effect of the electron beam on the proton bunches is therefore negligible.

However, the two antiproton bunches A28 and A29 also pass through the electron beam, and they are transversely aligned with the beam. Figure 7.20 illustrates the impact of the frenetic T_{FL} on the two bunches. The number of particles in each antiproton bunch




Figure 7.19. Trace of the T_{EL} agitating two antiproton bunches. While the individual signals are overlaid, the light grey is the current from the cathode and the black is the BPM plate signal (integrated once). The multiple traces typify the large pulse-to-pulse variations.

is shown; each of the bunch trains are grouped together. Most bunches contain 20×10^9 to 28×10^9 particles, but the two bunches in question only have a tenth of those amounts. Before the two-minute episode, these bunches also contained approximately 20×10^9 antiprotons apiece. The severe reduction in numbers illustrates the consequences of current oscillations.

In addition to measuring the bunch intensity, the horizontal and vertical emittances of each bunch were analyzed after the event. In Figure 7.21, the horizontal emittances of the unaffected bunches average less than 20π mm-mrad, while the vertical emittances average around 30π mm-mrad, with a systematic error of about 10% and a statistical error of less than 5% [87]. In each transverse dimension, the reported emittances for bunches




Figure 7.20. Antiproton-bunch intensities after the crazed T_{EL} interfered with the store. Bunches A28 and A29 are nearly completely eliminated, and A30 was also slightly affected.

A28 and A29 are at least twice the other values. However, due to the small number of particles in the two bunches, the Tevatron technique for measuring the emittance is quite inaccurate [88].

The subsequent bunch, A30, passed through the trailing end of the T_{EL} pulse in Figure 7.19. The number of particles in that bunch was reduced by roughly 20%, which provides enough remaining particles to provide a solid measurement of its transverse emittances. As Figure 7.21 illustrates, the detected horizontal and vertical emittances are both significantly higher than those of the remaining bunches.

Close to this time, the T_{EL} also interacted with bunches A4, A5, and A6. The time that the T_{EL} interacted with them was much shorter than the time that it toyed with A28 and A29; nevertheless, the number of antiprotons in each of those bunches was slightly diminished. This can be observed in Figure 7.20, where the intensity of those three bunches dips below the average of the bunches around them. More noticeable is the effect of the T_{EL} on the transverse emittances. In both the horizontal and vertical plots,




Figure 7.21. Horizontal and vertical emittance measurements after the current fluctuations nearly obliterated bunches A28 and A29. All of the touched bunches show significantly increased emittances.

the emittances of those three bunches is significantly higher than the other measurements. In general, Figure 7.21 shows relatively smooth emittances amongst all of the bunches except the ones that the TEL interacted with while having large current fluctuations.

7.4 Conclusions

The beginning of this dissertation quoted several references to the beam-beam limit as a Fundamental Limit to the progress of synchrotrons. The concept that bunches intrinsically possess a large tune spread which increases with higher collision rates seems

unpreventable. The resonance theory presented earlier indicates that accelerators must avoid putting particles on resonances. Large accelerators are unable to avoid having imperfect optics, and therefore these resonances will always exist.

The idea of using radial fields from an immersed charged beam is a very recent idea, and has, until now, never been attempted. It is very exciting to show in this dissertation the first successful attempt at tuneshift compensation. Over many tests and under different situations, the ability to tuneshift a chosen bunch through the use of an electron beam is proven. The theory matches the experiment reasonably well, with qualitative description given to explain the small differences. Throughout the experiment, the T_{EL} has responded to different experiments very similarly to how it was expected to behave. Measurements of the Tevatron-bunch parameters have also followed expectations.

More specifically, the tuneshift rises with current, as shown in Figure 7.2, very well. After gaining experience at alignment of the electron beam, the results appear very promising. The interesting results in Figure 7.4 follow the theoretical arguments that predict a reverse tuneshift when the electron beam is suitably misaligned.

The beam-loss monitors in the Tevatron often appear fickle. Nevertheless, by using them in the best manner while adjusting the T_{EL} current and plotting the results, the data resembles a theoretical curve in Figure 7.9 strikingly well.

Perhaps one of the most interesting data taken (and definitely the most laborious) appears in Figure 7.11. The mapping out of the bunch lifetime at numerous working points while the T_{EL} was operating proved to be very rewarding. The regions of high lifetime reside between known resonance lines, and the difference between the two electron guns is likewise very clear.

The TEL itself was thought out very carefully before being constructed, and the extra design has paid off. While officially called an “experiment,” the TEL has been extremely reliable during the operation of the Tevatron. Only a few difficulties remain to be worked out before the TEL can truly be an operational device.

First, the electron guns have a difficult time producing the desired current after many months of operation. This is believed to be caused by the slow atrophy of the oxides that are originally contained within the cathode’s interior. Months of heating and running allows this layer to dissipate, which in turns raises the work function and lowers the gun’s maximum output. In fact, the flattop gun, when it was first commissioned, was able to produce 3.5 A of peak current. By the time it was switched out for the Gaussian gun, it could only put out 0.8 A. This problem needs to be rectified if the TEL is to be depended upon for long-term use.

Another issue has been the anode modulator circuit. Several designs for such a system have been proposed, but the operation of the TEL requires some difficult high-power equipment which has succeeded for a while before one part or another breaks or deteriorates. Originally, the circuit was thought of as just one of many electrical subsystems that can be built rather straightforwardly. Unfortunately, it turns out to be a major project, requiring a large allocation of resources. Currently, another design is being built and tested which is hoped to be more reliable, while improving on the peak voltage and risetime constraints.

It is exciting to know that a second TEL is currently being tested. The option of having two TEL’s has been explained to be much more useful for compensation than just one. Building it is indicative of the worth of such a device. The second TEL employs

a few relatively minor changes, including a longer main solenoid and only forty-degree bends (instead of the ninety in the first T_{EL}).

Lastly, a few control and diagnostic issues need to be resolved. For unattended compensation, a control loop must be created, which can adjust different parameters of the T_{EL} given outputs from the diagnostics. While much of the elements are in place, no control loop has yet been attempted. Before that occurs, it is imperative that the BPMs provide reliable data, the tune-tracking systems are stable, and that the timing system stays functional. In addition, these diagnostics must respond quicker than the timescale of changes in the Tevatron. While the orbits drift rather slowly compared to the revolution frequency, they still can move over the range of seconds, and as of this writing, the diagnostics require more time to respond. The second T_{EL} is attempted to improve on all of its diagnostics.

The thesis project described in this dissertation has been personally extremely rewarding and beneficial. It is also an attempt to explore an idea that has never been experimentally tried before. The concept of beam-beam compensation has far-reaching implications in the realm of high-energy accelerators, and it is hoped that devices similar to the T_{EL} become standard in future accelerator endeavors. The beam-beam limit can be broken.

References

- [1] H. Edwards. The Tevatron energy doubler: a superconducting accelerator. *Annual Review of Particle Science*, 35:605-660. 1985.
- [2] V. Shiltsev *et al.* Status of Tevatron collider Run II and novel technologies for luminosity upgrades. *Proceedings of EPAC*, pp.239-243. 2004.
- [3] V. Shiltsev. Private communication.
- [4] D. Siergiej. Beam-beam interaction effects in the Fermilab collider. Ph. D. dissertation, University of New Mexico. 1995.
- [5] M. Month. Collider performance with ideal collisions. Brookhaven Note BNL-51909. 1985.
- [6] D. Edwards and M. Syphers. An introduction to the physics of high energy accelerators. Wiley. 1993.
- [7] Fermilab Accelerator Division. Antiproton source rookie book. Unpublished. 1999.
- [8] T. LeCompte and H. Diehl. The CDF and DØ upgrades for Run II. *Annual Review of Nuclear and Particle Science*, 50:71-177. 2000.
- [10] L. Coney. Diffractive W and Z boson production in $p\bar{p}$ collisions at $\sqrt{s} = 1800$ GeV. Ph. D. dissertation, University of Notre Dame. 2000.
- [11] A. Machiel. The CDF and DØ B physics upgrades. *Proceedings of Conference on B Physics and CP Violation*, pp.27-39. 1997.
- [13] M. Worcester. Inclusive search for anomalous high- p_T like-sign lepton pair production at the Fermilab Tevatron collider. Ph. D. dissertation, University of California, Los Angeles. 2004.
- [14] D. Newhart. Private communication.

- [15] A. Piwinski. Satellite resonances due to beam-beam interaction. *IEEE Transactions on Nuclear Science*, 24:1408-1414. 1977.
- [16] CESR logbook 67:125. 1986.
- [17] H. Burkhardt and R. Kleiss. Beam lifetimes in LEP. *Proceedings of EPAC*, pp.1353-55. 1994.
- [18] F. Willeke. HERA status and perspectives of future lepton-hadron colliders. *Proceedings of EPAC*, pp.20-24. 2002.
- [20] J. Rees. Colliding beam storage rings — a brief history. *SLAC Beam Line*, 9:2-8. 1986.
- [21] L. Evans. The challenge of the LHC. *CERN Courier*, 44:1. 2004.
- [22] D. Bollinger. Private communication.
- [23] J. Augustin *et al.* A multigev electron-positron colliding beam system with space charge compensation. *Proceedings of International Conference on High Energy Colliders*, 2:113. 1969.
- [24] J. LeDuff and M. Level. Experiences faisceau-faisceau sur DCI, laboratoire de l'Accélérateur linéaire. LAL-RT 80-03. 1980.
- [25] B. Podobedov and R. Siemann. Coherent beam-beam interaction with four colliding beams. *Physical Review E*, 52:3066-073. 1995.
- [26] R. Talman. Non-linear perturbation of a cyclic accelerator lattice. Unpublished. 1976.
- [27] R. Talman. Private communication.
- [28] E. Tsyganov *et al.* Compensation of the beam-beam effect in proton-proton colliders. SSC-519. 1993.

- [30] V. Shiltsev and D. Finley. “Electron compression” of beam-beam footprint in the Tevatron. Fermilab Note TM-2008. 1997.
- [31] V. Shiltsev *et al.* Considerations on compensation of beam-beam effects in the Tevatron with electron beams. *Physical Review Special Topics - Accelerators and Beams*, 2:071001. 1999.
- [32] S. Y. Lee. Accelerator physics. World Scientific. 1999.
- [33] S. Humphries. Charged particle beams. Humphries. 2002.
- [34] S. Eidelman *et al.* Review of particle physics. *Physics Letters B*, 592:22. 2004.
- [35] D. Griffiths. Introduction to electrodynamics. Prentice-Hall. 1991.
- [36] J.D. Jackson. Classical Electrodynamics. Wiley. 1999.
- [37] E. Hecht. Optics. Addison-Wesley. 1987.
- [38] U. Wienands. Physics of high-energy particle accelerators. Unpublished. Course given at USPAS, June 2001.
- [40] K. Bishofberger *et al.* Initial results from the Tevatron Electron Lens. *Proceedings of Advanced Accelerator Concepts*, 647:821-829. 2002
- [41] E. Courant, M. Livingston, and H. Snyder. The strong-focusing synchrotron — a new high energy accelerator. *Physical Review*, 88:1190-1196. 1952.
- [42] A. Sessler. The development of colliders. *Multi-GeV High-performance Accelerators and Related Technology*, pp.8-15. 1997.
- [43] G. Hill. On the part of the motion of the lunar perigee which is a function of the mean motions of the sun and moon. *Acta Mathematica*, 8:1-36. 1886.
- [44] E. Courant and H. Snyder. Theory of the alternate gradient synchrotron. *Annals of Physics*, 3:1-48. 1958.

- [45] Fermilab Accelerator Division. Accelerator concepts rookie book. Unpublished. 2004.
- [46] M. Syphers. Skew quadrupole tuning and vertical dispersion in the Tevatron. Fermilab Beams-doc 611. 2003.
- [47] Fermilab Accelerator Division. Tevatron rookie book. Unpublished. 2005.
- [48] Y. Papaphilippou *et al.* Deflections in magnet fringe fields. *Physica Review E*, 67:046502. 2003.
- [50] C. Pellegrini. On a new instability in electron-positron storage rings: the head-tail effect. *Nuovo Cimento*, 64A:447-472. 1969.
- [51] A. Chao and M. Tigner. Handbook of accelerator physics and engineering. World Scientific. 1999.
- [52] P. Bambade. Effets faisceau-faisceau dans les anneaux de stockage e^+e^- à haute énergie: grossissement résonant des dimensions verticales dans le cas de faisceaux plats. Ph.D. dissertation, Université de Paris Sud, Centre d'Orsay. 1984.
- [53] W. Herr. Tune shifts and spreads due to short and long range beam-beam interactions in the LHC. *CERN LHC Note 119*. 1990.
- [54] W. Herr. Beam-beam interactions. Unpublished. Course given at CERN Accelerator School, June 2003.
- [55] K. Spangenberg. Vacuum tubes. McGraw-Hill. 1948.
- [56] R. Pollock. Storage rings for nuclear physics. *Annual Review of Nuclear and Particle Science*, 41:357-388. 1991.
- [57] F. Zakaria. The future of freedom. Norton. 2003.

- [58] W. Kreyer, J. Hemsley, and M. Andrews. A simple numerical method for a modified Abel inversion in which the density can be approximated by elliptical symmetry. *Journal of Physics D*, 26:1836-1842. 1993.
- [60] C. Kim *et al.* Characterizations of symmetry and asymmetry high-density gas jets without Abel inversion. *Review of Scientific Instruments*, 75:2865-68. 2004.
- [61] P. Chou. The nature of transverse beam instabilities at injection in the Fermilab Main Ring. Ph.D. dissertation, Northwestern University. 1995.
- [62] Y. Miyahara. Stability of trapped ions in electron storage rings. *Proceedings of the Workshop on Beam Instabilities in Storage Rings*, pp.136-157. 1994.
- [63] H. Koziol. Beam diagnostics revisited. *Proceedings of EPAC*, pp.164-168. 1998.
- [64] S. Smith. Beam position monitor engineering. *Proceedings of Beam-Instrumentation Workshop*, pp.50-65. 1996.
- [65] X. Zhang *et al.* Operation of the beam diagnostics system for Tevatron Electron Lens. *Proceedings of Beam Instrumentation Workshop*, pp.483-490. 2002.
- [66] P. Horowitz and W. Hill. The art of electronics. Cambridge. 1989.
- [67] Fermilab beam-beam compensation logbook. September 27, 2001.
- [68] R. Davidson and H. Qin. Physics of intense charged particle beams in high energy accelerators. World Scientific. 2001.
- [70] K. Bishofberger. BPM memo #3. Unpublished. 2003.
- [71] L. Zhijing *et al.* Instability of relativistic electron beam with strong magnetic field. *Chinese Science Bulletin*, 45:18-23. 2000.
- [72] R. Zimmerman. Review of single bunch instabilities driven by an electron cloud. *Physical Review Special Topics - Accelerators and Beams*, 7:124801. 2004.

- [73] A. Gilmour. Principles of traveling wave tubes. Artech House. 1994.
- [74] E. Darlington. Backscattering of 10–100 keV electrons from thick targets. *Journal of Applied Physics*, 8:85-93. 1975.
- [75] S. Lai. Secondary and backscattered electrons. Lecture notes from a class given at the Swedish National Graduate School of Space Technology. Unpublished. 2003.
- [76] N. Solyak. Electron beam system of the Tevatron Electron Lens. *Proceedings of PAC*, pp.1420-22. 2001.
- [77] K. Bishofberger. Electrical memo #5. Unpublished. 2001.
- [78] G. Saewert. Private communication.
- [80] K. Bishofberger *et al.* Characterization of the Tevatron Electron Lens magnetic system. *Proceedings of PAC*, pp. 3406 – 08. 2001.
- [81] Fermilab Accelerator Division. Plans for Tevatron Run IIB. D. McGinnis and H. Montgomery editors. Unpublished. 2001.
- [82] L. Tkachenko *et al.* Tevatron Electron Lens magnetic system. *Proceedings of PAC*, pp. 3630 – 32. 2001.
- [83] K. Bishofberger. Magnet calibrations. Unpublished. 2001.
- [84] L. Tkachenko *et al.* Test results of the magnetic system for the Tevatron Electron Lens. *Proceedings of EPAC*, pp. 2433 – 35. 2002.
- [85] L. Smith and P. Hartman. The formation and maintenance of electron and ion beams. *Journal of Applied Physics*, 11:220-229. 1940.
- [86] C. Olson and J. Poukey. Force-neutral beams and limiting currents. *Physical Review A*, 9:2631-2634. 1974.
- [87] A. Hahn. Online calculation of Tevatron collider luminosity using accelerator instrumentation. *Proceedings of PAC*, pp. 2226 – 2228. 1997.
- [88] D. Crawford. Private communication.