The Dissertation Committee for Robert Miles Zwaska
certifies that this is the approved version of the following dissertation:

Accelerator Systems and Instrumentation for the NuMI Neutrino Beam

Committee:

Sacha Kopp, Supervisor

Duane Dicus

Karol Lang

Eric Prebys

Jack Ritchie

Christopher Sneden

Accelerator Systems and Instrumentation for the NuMI Neutrino Beam

by

Robert Miles Zwaska, B.S.

Dissertation

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

December 2005

For now we see in a mirror dimly, but then face to face.

Now I know in part, then I shall understand fully.

Acknowledgments

This document represents the culmination of a long period in my life when I learned some and worked much more. Much of that work would not have been possible without the help of people around me.

First of all, I must thank my adviser, Prof. Sacha Kopp. Not only did he provide support and guidance for almost six years, but he also took a personal interest in my progress. Through many late nights of fortran, kumacs, powerpoint, latex, and occasionally drink – we worked on (my apologies to his family). Particularly during my last six months, I must acknowledge the exceptional devotion he placed into ensuring the quality of this document and my eventual completion.

While at Texas, I gained much from the fellow graduate students in my group: Žarko "The Magnificent" Pavloviç and Dharma "Sky Pilot" Indurthy, each of whom excel in quiet contemplation. Similarly, but more noisily, a small army of undergraduates came and went; among them: Ryan Keisler, Chieze Ibeneche, and James Hall. Our engineer, Marek Proga, was invaluable as he could be diligent and efficient with either a smile or scowl upon his face.

While at Fermilab, I profited greatly from collaboration with my supervisor, Eric Prebys, and Booster tsar, Bill Pellico. I did nearly all my accelerator work with Bill; his experience and devotion to the task were crucial in the ultimate success of cogging. I would also like to thank my many MINOS colleagues who provided guidance; among them: Mark Messier, Alberto Marchionni, Jim Hylen, Debbie Harris, and Al Erwin.

At both Texas and Fermilab, I have had many fellow travelers. They could always provide commiseration and were willing to talk shop without too much embarrassment. I owe thanks to many, but trying to name a few: Trisha Vahle, Todd Tinsley, Rustem Ospanov, Tom Osiecki, Mike Kordosky, Phil Janisiewicz, Jay Hanssen, and Abinidab Dieter.

Ultimately my greatest gratitude goes to my family. My brothers, Michael and David, can play the roles of good friends, while also grimly resigned to the certainty of our permanent relationship. I must save my last and best appreciation for my parents: Robert Frank and Catherine. They raised me to be who I am. During my graduate experience they were always supportive. Particularly toward the end, I appreciate the restraint they exercised over their own anticipation, while still providing the supportive rock of a welcome family.

Robert Miles Zwaska

The University of Texas at Austin
December 2005

Accelerator Systems and Instrumentation for the NuMI Neutrino Beam

Publication No. ______________

Robert Miles Zwaska, Ph.D.
The University of Texas at Austin, 2005

Supervisor: Sacha Kopp

The Neutrinos at the Main Injector (NuMI) neutrino beam facility began operating at the Fermi National Accelerator Laboratory in 2005. NuMI produces an intense, muon-neutrino beam to a number of experiments. Foremost of these experiments is MINOS – the Main Injector Neutrino Oscillation Search – that uses two neutrino detectors in the beam, one at Fermilab and one in northern Minnesota, to investigate the phenomenon of neutrino oscillations.

NuMI is a conventional, horn-focused neutrino beam. It is designed to accept a 400 kW, 120 GeV proton beam from the Fermilab Main Injector accelerator. The proton beam is steered onto a target, producing a secondary beam of mesons which are focused into a long evacuated volume where they decay to muons and neutrinos. Pulsed toroidal magnets (horns) focus an adjustable meson momentum range. Design of the beamline and its components

is challenged by the 400 kW average proton beam power.

To achieve such high proton power, the Fermilab Main Injector (MI) must store and accelerate $\sim 4{\times}10^{13}$ protons per acceleration cycle. This requires the MI to be loaded with 6 or more batches of protons from the 8 GeV Booster accelerator. Such multiple-batch injection involves a synchronization of the two machines not previously required by the Fermilab accelerators. In this dissertation, we investigate timing errors that can arise between the two accelerators, and a feedback system which enables multiple Booster transfers into the Main Injector without significant loss of beam. Using this method of synchronous transfer, the Main Injector has delivered as many as $3{\times}10^{13}$ protons per pulse to the NuMI beam.

The instrumentation to asses the quality of the neutrino beam includes arrays of radiation-tolerant ionization chambers downstream of the decay volume. These arrays detect the remnant hadrons and tertiary muons produced with the neutrinos. This thesis discusses measurements using the arrays, including diagnostics of potential beam errors and initial alignment of the beamline elements. In addition, we describe theoretical calculations and experimental tests which demonstrate that the ionization chambers can withstand the intense particle fluxes in the beamline without signal loss due to space-charge build up in the detectors.

Contents

Chapter 1

The NuMI Beam

Accelerator-based neutrino beams are used as controlled, high-energy, and high-flux sources of neutrinos. Other sources of neutrinos (solar, atmospheric, supernova, reactor, geo) depend on knowledge of the underlying physics and details of the source; details which often cannot be measured. Accelerator beams, however, involve better control of the underlying processes, and thus of the neutrino beam.

An accelerator-based neutrino beam was first used to discover neutrino flavor [1]; that is, that muon-neutrinos interact differently than electron-neutrinos. Accelerator-based neutrino beams can be predominantly of the muon type, and can be delivered quickly ($\sim \mu$s), such that background rejection is not as crucial. Furthermore, accelerator-based neutrino beams offer higher-energy neutrinos (up to 600 GeV) and can be configured to produce well-controlled spectra of neutrinos (e.g., [2]).

Several high-intensity accelerator-based neutrino beams are being designed with long baselines (> 100 km) to investigate the phenomena of neutrino mixing and oscillations. While neutrino mixing was initially observed with solar [3, 4, 5, 6, 7, 8, 9] and atmospheric [10, 11, 12, 13] neutrinos, accelerator-based sources are necessary to more precisely measure the fundamental parameters of neutrino mixings and masses, as well as to definitively eliminate alternative theories (e.g., [14, 15, 16]). Accelerator-based beams allow direct measurement of the neutrino flux near the source, followed by another measurement after the beam traverses the baseline. Changes in the neutrino flux through energy-dependent appearance and disappearance of neutrino types provide discrimination between theories, and precision in the measurements of fundamental parameters.

1.1 NuMI Overview

NuMI (Neutrinos at the Main Injector) is a conventional, two-horn focused, accelerator neutrino beam intended for short and long baseline neutrino experiments. The neutrino beam provides an intense, high-energy (2-20 GeV) flux of muon-neutrinos, with low contamination of electron-neutrinos.

NuMI is designed to use the 120 GeV/c proton beam provided by the Fermilab Main Injector at an average power of up to 400 kW (instantaneous power of $\sim$ 1 TW). The primary beamline bends the beam downward, to achieve a 58 mrad vertical angle that directs the beam toward the Soudan mine in northern Minnesota, as shown in Fig. 1.1. The primary proton beamline (§1.2) includes instrumentation to measure the intensity, trajectory, and size of the proton beam.

The NuMI beamline converts the intense proton beam into a focused neutrino beam; the major components are illustrated in Fig. 1.2. The protons are steered into a solid carbon target to produce mesons through hadronic interactions. The production of mesons from the target is not directly measured

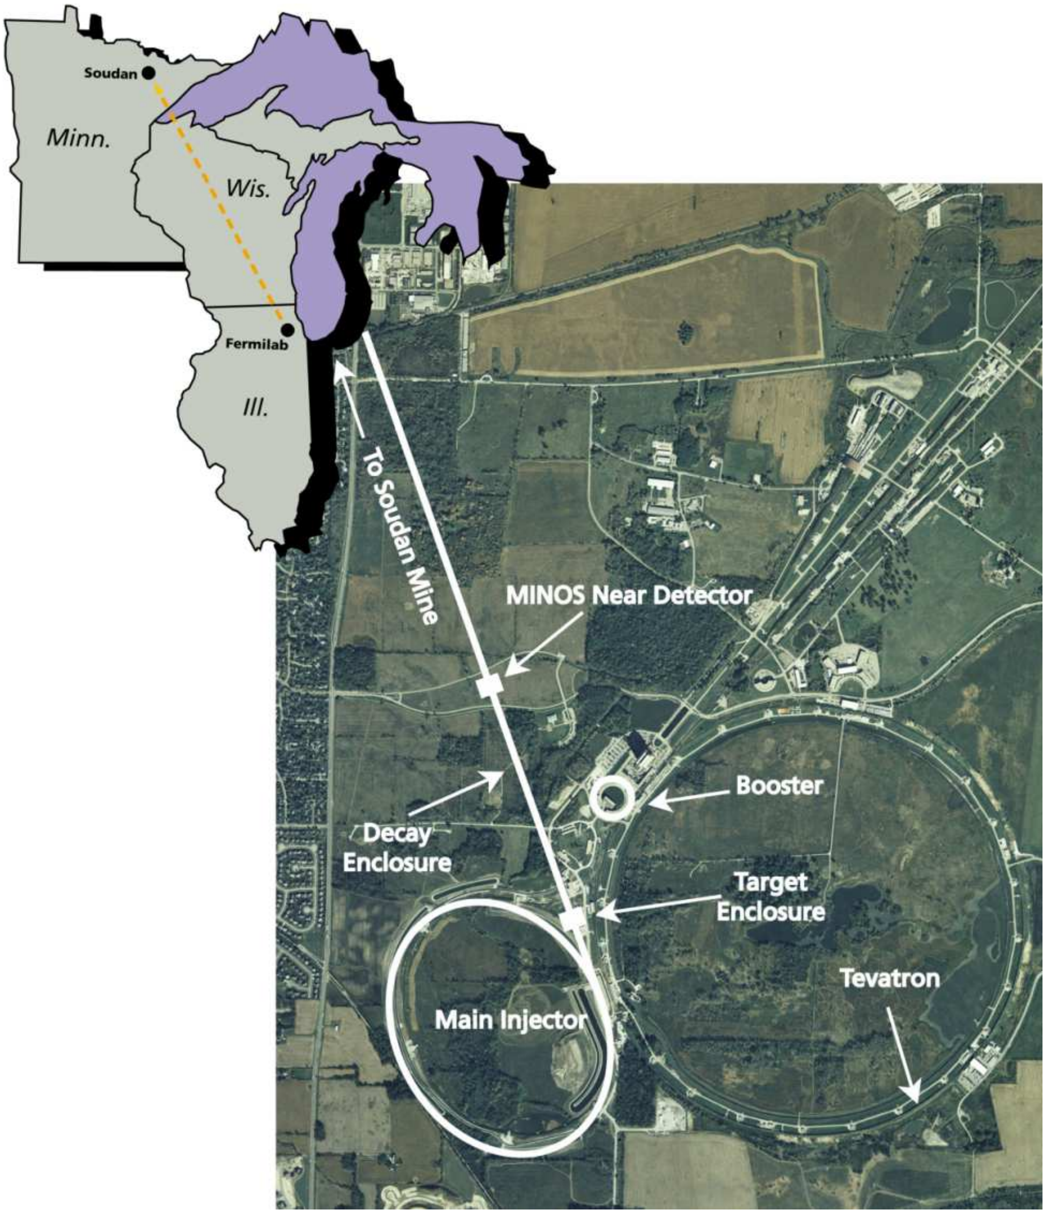

Figure 1.1: Maps of the Upper Midwest and Fermilab, with the NuMI beam superimposed. The NuMI beam is created underground and at a 58 mrad vertical angle, downstream of a Main Injector extraction region. The neutrino beam travels largely unimpeded through the earth, passing through the Soudan mine in Minnesota, and exiting the earth in Canada.

in the experiment, but is estimated through the use of standard Monte Carlo codes. Such estimates use input from separate particle production experiments (among others: [17, 18, 19]). The most relevant mesons for the production of neutrino beams are pions (π^+, π^-) and kaons (K^+, K^-, K^0); those mesons are produced in the greatest numbers and are likely to decay into neutrinos.

After production, the charged mesons are focused by a pair of pulsed, toroidal magnets called horns. A horn is a coaxial, transmission-line magnet with an air gap between the conductors. The inner conductors of the NuMI horns have parabolic profiles, such that they allow momentum selection of the pions through focusing (see §1.4).

Subsequent to the horns is a cylindrical evacuated volume into which forward-focused mesons are conducted. In the volume, the mesons decay into a muon and a muon-neutrino – establishing the neutrino beam. An additional focusing system, known as the Hadron Hose [20], was proposed for inclusion within the decay volume, but not ultimately built. The Hose would have further focused the mesons in the decay volume to reduce interactions with the wall.

The Hadron Monitor detector rests at the end of the decay pipe and measures the remnant beam of uninteracted protons and undecayed mesons; the Hadron Monitor is further described in Ch. 6. Immediately following the Hadron Monitor is the Hadron Absorber – a mass of steel, concrete, and aluminum that stops the hadrons. Muons penetrate the absorber and continue into unexcavated rock, where they range out over a distance proportional to their energy (up to $\sim$ 170 m). Before the start of the rock, and in two excavations into the rock, Muon Monitor detectors are installed to measure the spatial distributions of three energy ranges of the muon beam; the Muon Monitors are described further in Ch. 6.

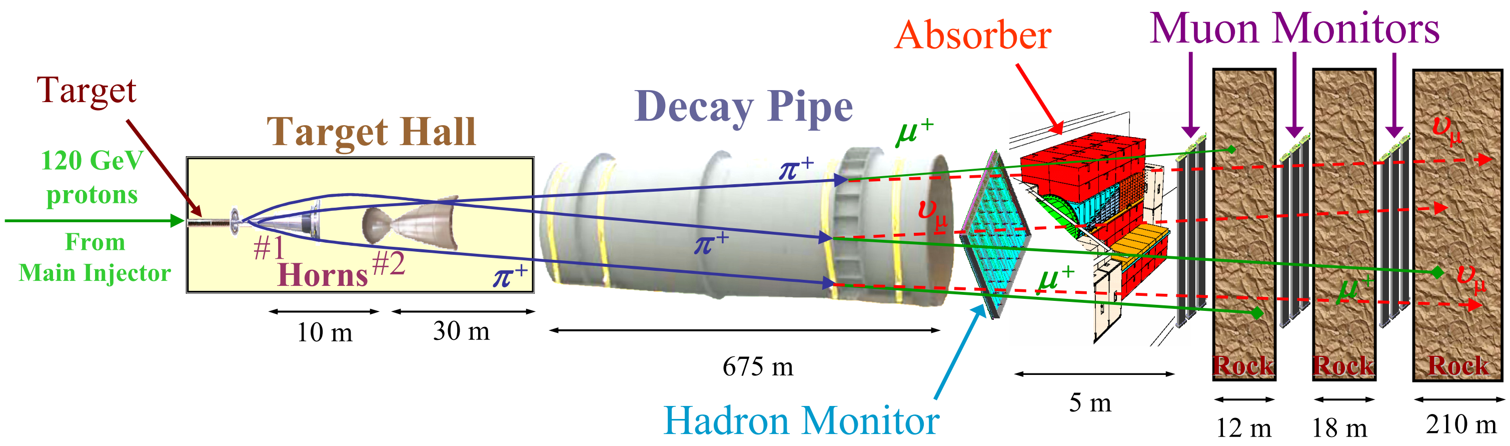

Figure 1.2: Illustration of the major NuMI beam components. 120 GeV/c momentum protons are delivered by the Fermilab Main Injector. The protons interact in a carbon target to produce mesons – predominantly pions. The pions are focused by a pair of pulsed "horn" magnets into the decay pipe. The pions decay to muons and muon-neutrinos. The Hadron Monitor measured the spatial distribution of the uninteracted protons and undecayed pions, after which they are stopped in the Hadron Absorber. The muons penetrate the absorber and some distance of rock, their spatial distributions being measured at three Muon Monitor stations. The neutrino beam is largely unaffected by the rock; the neutrinos propagate in straight lines to the detectors.

1000 m downstream of the target, and 300 m downstream of the end of the decay volume sits the MINOS (Main Injector Neutrino Oscillation Search) near detector. It measures the neutrino beam shortly after production in the NuMI line, and presumably before any oscillations occur. 735 km from the target is the MINOS far detector at the Soudan Underground Laboratory, in Soudan, Minnesota. The far detector measures the neutrino beam again, after the hypothesized neutrino oscillations shall have forced some portion of the muon-neutrinos to disappear.

1.2 The Primary Proton Beam

The NuMI primary beam extracts and transfers 120 GeV/c protons from the Main Injector to the meson production target. The transfer line is composed of main dipole and quadrupole magnets to steer and focus the beam onto the target. Redundant instrumentation along the beamline measure the intensity, position, and size of the beam (for further information, see App. B). A "Beam Permit System" monitors the operation of the beam components, and prevents beam whenever it might have been delivered incorrectly [21].

Extraction from the Main Injector is performed in a single turn through a fast kicker and lambertson magnet system. The kicker magnets are ferrite-loaded transmission line magnets, such that the beam pipe lies within a gap in the ferrite and between the conductors. The kicker is designed with a fast ($<$ 1.6 μs) risetime, setting the minimum abort gap length. The beam is kicked horizontally, such that after one-quarter of a betatron oscillation it arrives at the opposite aperture of a lambertson magnet than the circulating beam passes through. Now in the magnetic field, the beam is bent vertically, out of the Main Injector[1].

The transfer line consists of a strong focusing lattice that directs the beam to the target. The proton beam must be over-bent vertically to avoid having an extended tunnel at the level of a local aquifer. Furthermore, the tunnel had to be small during that portion and could not contain many components. The large vertical bends and restricted locations of the magnets lead to a line with vertical dispersion as large as 6 m. The magnet apertures and beam pipe were chosen to be very large so as to allow a maximum acceptance of 500π mm·mrad; preventing beam halo or accidents from causing any significant loss in the primary line. The currents through each of the main dipole and quadrupole magnets are independently adjustable. Fig. 1.3 shows calculations of beam size throughout the NuMI primary line, as well as measurements at several locations.

Trim dipole magnets in the transfer line are used to fine-tune the beam position[2]. Trim magnets can be adjusted manually to change the beam position at a location, or to move the beam across the target (see Ch. 11 & App. B). During operation, trim magnets are controlled by a program called AUTOTUNE that adjusts each magnet so as to maintain a desired trajectory.

Current-transforming toroids measure the intensity of the proton beam at two locations: immediately after extraction, and before the target. Their calibration was performed with a direct current source and empirically they are within 0.4% of the Main Injector direct-current current-transformer (DCCT), which provides a measurement of the beam intensity while still circulating in the Main Injector. The other beamline instrumentation described below can

[1]Two of the lambertson magnets are actually oriented at a small angle from horizontal, such that they apply a small horizontal bend to counteract quadrupole bending.

[2]Quadrupole trim magnet are not present, as the targeting of the proton beam does not require great precision in the lattice functions.

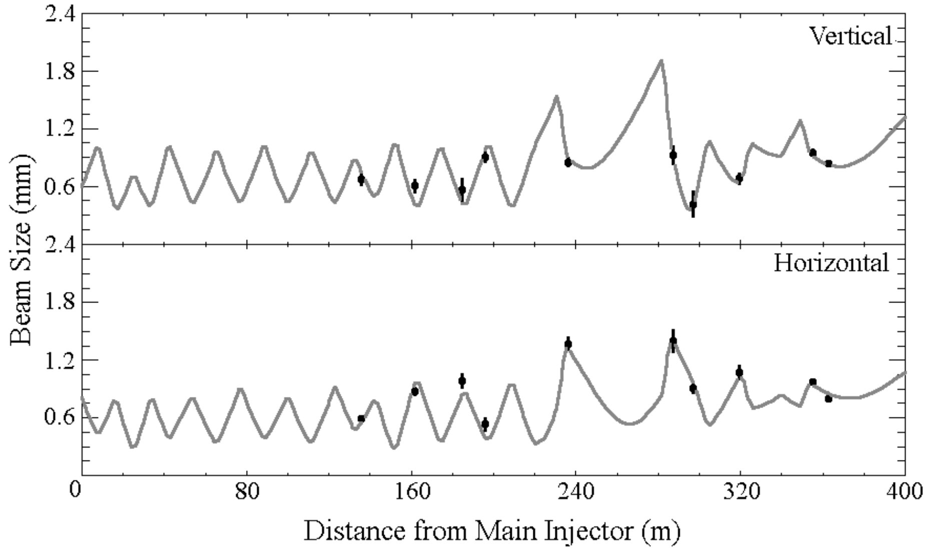

Figure 1.3: Measured beam widths (points with error bars) superimposed on lattice calculations of beam width along the NuMI primary line (solid lines). The widths are measured by Profile Monitors at each station – the last two points in each plot correspond to the two stations just upstream of the target. The beam width waists in both planes at the target. (courtesy Ming-Jen Yang)

also give intensity measurements, but they are less precise.

Beam Position Monitors (BPMs) are electromagnetic detectors which measure the beam's horizontal or vertical position at that location. Each BPM consists of two electrodes (replacing the beam pipe) that couple to the beam capacitively. Comparison of the charge induced on the plates provides the position measurement of the beam; summation also provides an intensity measurement [22].

Profile Monitors (PMs) are interceptive detectors that measure the beam's horizontal and vertical beam position and width. Each PM is composed of two sets of segmented titanium foils, oriented orthogonally to each other and transverse to the beam axis. The foils act as secondary-emission monitors (SEMs); such that the electric charge measured on each, from electron ejection, is proportional to the number of incident protons. The profile information is all recorded and can be plotted; the data are reduced to intensity, position, and width information for the beam, as well as some values for beam halo [23, 24, 25, 26]. App. B discusses the relative calibration and precision of the BPMs and PMs, as well as their ability to extrapolate beam position into the target hall.

The NuMI Beam Permit System (BPS) [21] monitors all of the primary beam magnets to ensure that they reach the prescribed current by the time of the beam spill. If any of the magnets fail to ramp correctly, the beam will be extracted to the Main Injector dump instead. Additionally, the BPS monitors several quantities from the target hall and Main Injector that indicate or affect beam quality.

1.3 Meson Production Target

NuMI uses a carbon target of almost 2 nuclear interaction lengths to maximize meson production, while surviving the bombardment of the intense proton beam. A schematic of the NuMI target assembly is shown in Fig. 1.4.

The target body is segmented into 47 individual 2 cm length graphite plates. Each plate is 6.4 mm wide and $\sim$ 18 cm tall and has rounded edges to minimize the effects of strain. Graphite can sustain very high temperatures and is relatively strong. Tests showed that graphite could survive the instantaneous stress of beam heating [27]. The graphite is $\sim$ 20% porous and has measured densities of 1.78 – 1.82 g/cm^3; we have used 1.81 g/cm^3 as the density in Monte Carlo calculations.

The target segments are contained between and sintered to a pair of stainless steel cooling tubes. The tubes circulate chilled water to remove heat

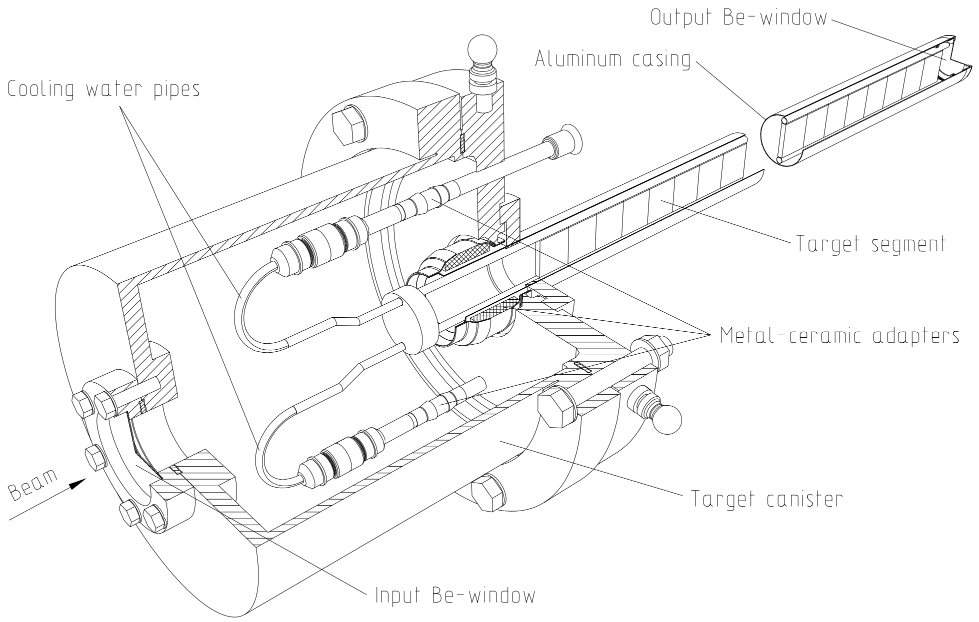

Figure 1.4: Schematic of the NuMI target assembly. The target is the length of forty-seven 2 cm carbon segments along the axis of the assembly. The target is cooled by water pipes along its length, all of which is electrically isolated from the enclosing canister. The interior of the canister is either evacuated or over-pressured with helium. The horizontal fin is not shown, but would be at about the position of the disk where the water pipes come together at the head of the target. (from [27])

deposited in the target, up to 16 kW at the design proton power. The tubes are 6 mm diameter; they sit out of the beam such that their centers are displaced 9 mm vertically from the beam spot.

The body of the target, including the graphite segments and cooling tubes, are electrically isolated from the target canister so that it can act as a "Budal Monitor" [28]. The proton beam will strip off electrons (δ rays) as it passes through the target, producing a current that is somewhat proportional to the proton flux. An electrical connection through a feedthrough brings the signal from the target outside for measurement.

The target canister consists of a large barrel section, and narrow snout section. The barrel section is upstream and contains all of the feedthroughs for water, helium, vacuum and instrumentation. The snout section is downstream and is designed so that some portion of the target may extend into the horn, within the inner conductor. The snout is an aluminum cylinder of 3 cm diameter. The canister can be either evacuated or filled with helium. Initial tests of the Budal monitor [27] suggested better signal with vacuum; however, the target's potential for water leaks requires a helium overpressure (see §10.3).

1.3.1 Meson Production from the Target

The spectrum of mesons created from the target is the first step in calculating the neutrino flux and interaction rates in the detectors. Various parametrized forms for the resultant yields of pions and kaons have been produced (e.g., see [29]); however, detailed Monte Carlo simulations of particle interactions are necessary to calculate a more precise flux. These models can be inputs or components of tracking Monte Carlos that calculate neutrino fluxes. While the Monte Carlo models generally have better correspondence to measured data, there still is significant uncertainty, which we explore below (and more so in [30]). In this section we provide meson yields using the FLUKA model of 2001a [31], MARS V14 [32], and the GEANT 3 implementation of FLUKA (GFLUKA) [33].

The predicted positive pion (π^+) yield is shown in Fig. 1.5. Depending on the model, 3.7 – 4.1 π^+ are created per proton ($p_\pi > 500$ MeV/c). The yields in momentum bins are within 30% of each other until 60 GeV/c.

As discussed in the next section, π^+ mesons will be focused, while π^-

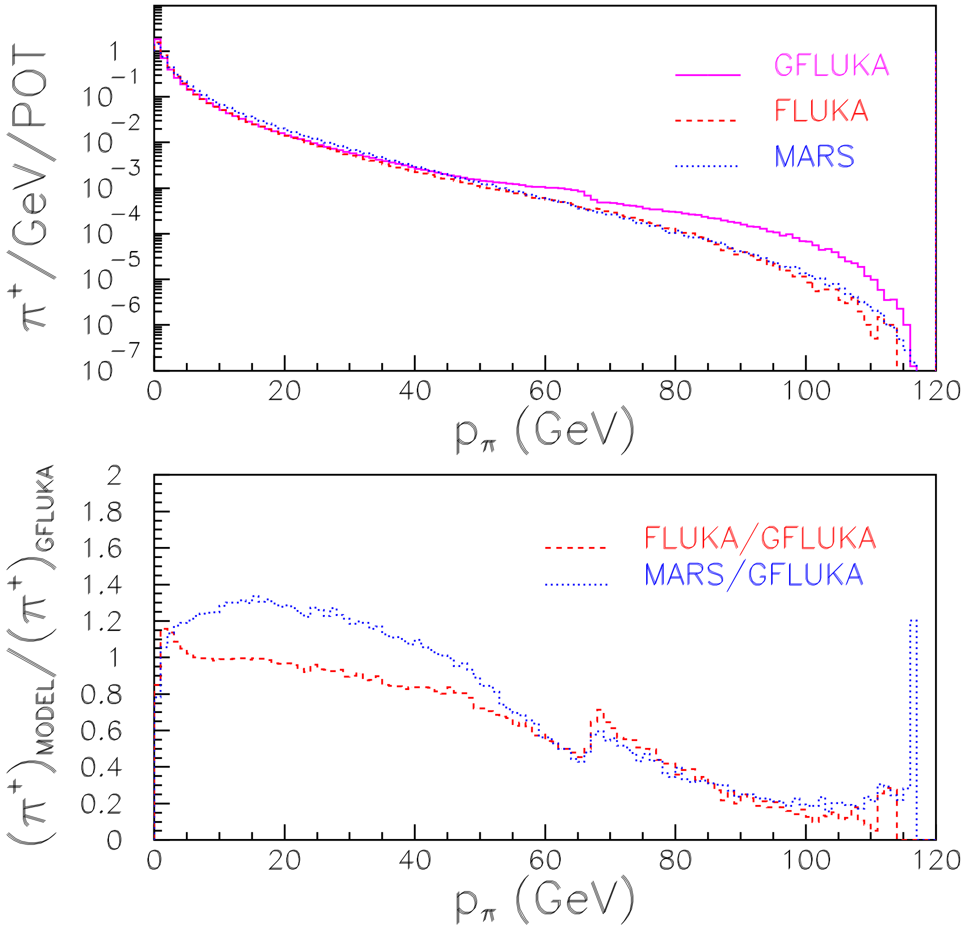

Figure 1.5: Yield of positive pions from the target, binned by pion energy, and as calculated by three different Monte Carlo models. The above plot shows the yield per proton, the below plot shows ratios of two hadron production models to the default one in the Monte Carlo.

mesons are defocused. The focusing preferentially selects pions within a particular band of momentum and transverse momentum (p_t). The transverse momentum range of interest is between 100 and 500 MeV/c, where the meson density is large[3]. The range of total momentum focused varies with beam configuration (see §1.4.2), and can be centered about 8, 18, or 25 GeV/c.

We have also explored the variation of hadron production with non-target material and with a detailed model of the target [34]. The variations of yield with non-target material is generally not as significant as the hadron production simulations themselves. The uncertainty in the neutrino flux and extrapolation between detectors may affect the ultimate reach of the MINOS experiment. While present hadron interaction codes limit the precision with which the neutrino flux may be calculated, we expect that the predictions for hadron production will be better constrained by data (e.g., see [35]).

1.4 Meson Focusing with Magnetic Horns

NuMI uses magnetic horns to focus the charged mesons and substantially increase the neutrino flux at the detectors. Originally implemented at CERN [36], horns are pulsed transmission-line magnets that create a focusing toroidal field. In the CERN implementations, as many as three horns were used along the length of the beamline – the later ones as "reflectors" to refocus the over-focused mesons, creating a type of strong focusing. In principle a greater number of horns produce greater neutrino flux; though material interactions will eventually absorb some portion of the mesons. In §1.5.1 we discuss an alternative focusing method for the length of the decay region.

[3]The beam flux is maximal in the forward direction (p_t=0), but the density of pions in transverse momentum reaches a model-dependent maximum around 300 MeV/c because of the growing solid angle.

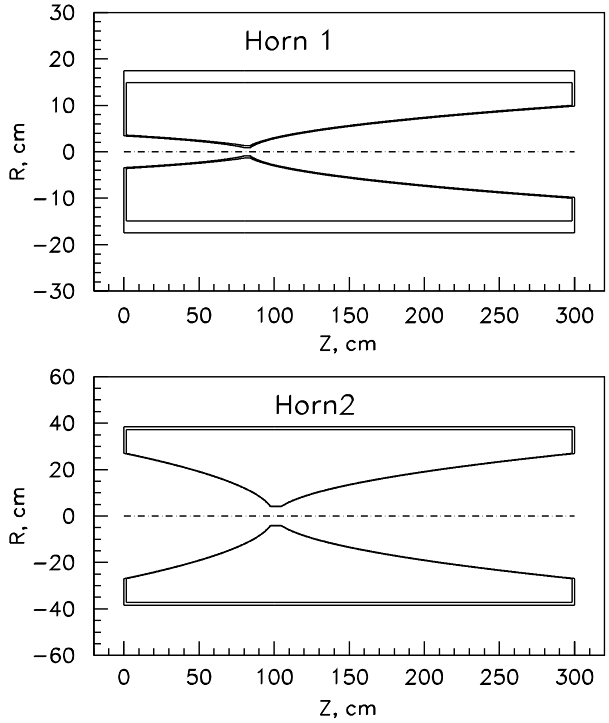

Figure 1.6: Profiles of the conductors of the two NuMI horns (from [37]). The pulsed current of up to 200 kA travels along the inner conductor and returns along the outer, producing a focusing magnetic field. The inner conductors have parabolic profiles.

Fig. 1.6 shows profiles of the NuMI horns [37]. Electric current travels along the inner conductor (closer to the axis) and returns along the outer. The current is the same along the length of the inner conductor, so a magnetic field of $B(R) \propto R^{-1}$ is created between the conductors. Only a negligible magnetic field is induced in the free space inside the inner conductor.

Particles pass from left-to-right through the horns in Fig. 1.6. Those produced parallelly will pass through the neck region, where there is no material and no field, and will continue parallel to the beam axis. Those mesons produced at an angle will encounter the horns at a nonzero radius (R). They then pass through the upstream parabolic inner conductor[4], and become sub-

[4]The passage of beam particles through the solid conductors of the horn will necessarily

ject to to the magnetic field. The field bends the mesons through some angle – mesons with very high transverse momentum exit the horn through the outer conductor and do not become part of the beam; those mesons that do constitute the beam exit the inner conductor on its downstream parabolic surface.

The mesons exiting the downstream face of the horn are, ideally, focused parallel with the beamline. Some number of underfocused and overfocused particles will not fully contribute to the beam, striking the beam enclosure before decaying to neutrinos. The contributions of these particles were explored further in [38]. In the case of NuMI, a second horn further focuses the beam $\sim$ 10 m after the first horn; the second horn also focuses some particles that had passed through the neck of the first horn[5].

The paraboloid shape of the horn inner conductor acts as a lens (in the thin-lens approximation); the focal length will be shown to be proportional to the secondary particle's momentum, so momentum-selective focusing is also possible. If we assume that a particle approaches the horn off-axis by a distance x it will a experience a magnetic field of:

$$B = \frac{\mu_0 I}{2\pi x} \tag{1.1}$$

Where I is the current in the horn and μ_0 is the permeability of free space. The focused particle will experience this field for a distance, ℓ, that depends quadratically on x:

$$\ell = ax^2 \tag{1.2}$$

Where a is the parabolic factor, and we have assumes that the particles are

involve some beam absorption and scattering. Accordingly, horns are unsuitable for multiple passes of particles, as in an accelerator ring. Additionally, they are made to present as little scattering, absorbing, and heating as possible; the NuMI horn conductors are fashioned of few mm thick aluminum.

[5]The second horn can be relocated an additional 10-20 m downstream of Horn 1 to focus higher energy mesons.

	Horn 1		Horn 2	
	Upstream	Downstream	Upstream	Downstream
Parabolicity (cm^{-1})	7.05	2.18	.135	.272
Focal Length	$0.119p$	$0.381p$	$6.14p$	$3.05p$
(m/(GeV/c))	$0.091p$		$2.03p$	

Table 1.1: The parabolic parameters for the shapes of the NuMI horn inner conductors, and the resulting focal lengths. Each horn's upstream and downstream portion have different parabolic factors. The focal lengths were calculated using a thin lens approximation; similarly the focal lengths of the two portions of each horn were summed harmonically to attain an aggregate focal length. All focal lengths are proportional to the momentum of the focused particle.

parallel to the beam axis and the deflection angle is small. The lateral, transverse momentum kick is found by integrating the magnetic field along the path:

$$\begin{aligned} \Delta p_{\mathrm{T}} &= \int_{\ell(x)} d\ell\, B(x) \\ &= \frac{0.3\mu_0 I}{2\pi} a x. \end{aligned} \tag{1.3}$$

The transverse kick is proportional to horn current, distance from the axis, and the parabolicity. Thus, the horn acts like a linear lens as the return force is proportional to the distance from the axis. As such, we can define a focal length for the horns:

$$\begin{aligned} f = 1/k &= \frac{x}{\Delta p_{\mathrm{T}}/p} \\ &= \frac{2\pi p}{0.3\mu_0 I a} \\ &= \left(1.67 \times 10^4 \,\frac{\mathrm{kA}}{\mathrm{GeV/c}}\right) \frac{p}{Ia} \end{aligned} \tag{1.4}$$

While, the NuMI horns have parabolic profiles, the parabolic factor, a, differs for the two halves of the horn (see Fig. 1.6). The values for the horns are enumerated in Table 1.1. We calculate the focal lengths for each half of each horn, as well as a focal length for the combined horn; calculated assuming that they are thin and immediately next to each other. In each case, the focal length grows linearly with the momentum of the pion.

In this section, we have treated horns as thin lenses for the purposes of illustration. However, the horns are not thin lenses by any reasonable approximation: the 1 m target sits inside the 3 m Horn 1 and does not provide a point source; the horns are extended objects with lengths approaching the calculated focal lengths; the bend of the particles within the horn is not negligible (Eqn. (1.2) becomes a poor approximation); and the material of the horn conductors provides substantial scattering and interaction with the beam. As such, useful optics calculations for neutrino spectra require tracking Monte Carlos that can simulate all of the above effects.

1.4.1 Neutrino Beam Spectra

The NuMI beam consists primarily of muon-neutrinos from the decays of pions:

$$\pi^+ \longrightarrow \mu^+ \nu_\mu \tag{1.5}$$

Additionally, kaon and muon decay create neutrinos of other species. The anticipated spectra of neutrino interactions at the detectors are calculated using a tracking Monte Carlo. The primary beam simulation for NuMI is known as GNUMI, as it is the GEANT [33] implementation of the NuMI beamline (see [37] for some simulation). The tracking portion of the simulation is constructed such that external models of hadron production (discussed in §1.3.1 and [30]) can be used for the target region. In this section, we will first discuss the major features of pion decay, and then then the resulting neutrino spectra produced by the Monte Carlo simulation.

The decay of relativistic pions results in a beam of neutrinos oriented in the direction of the pion beam. The neutrino energy in the lab depends on the pion's initial energy and on the decay angle with respect to the pion's direction (see [39] for details):

$$E_\nu \approx \frac{0.43 E_\pi}{1 + \gamma^2\theta^2} \tag{1.6}$$

θ is the angle the emitted neutrino makes with the pion trajectory, and γ is the pion's relativistic boost. We take from this equation that the neutrino will carry at most 43% of the pion's energy, and that the energy rapidly decreases if the they are not colinear. The muon will take the balance of the pion's energy.

The neutrino flux is concentrated in the forward direction:

$$\Phi \propto \frac{\gamma^2}{l^2(1 + \gamma^2\theta^2)^2} \tag{1.7}$$

l is the distance from the decay vertex. The neutrino flux thus falls off as the inverse-square of detection distance, is greater for high energy pions, and is concentrated in a forward cone whose width varies inversely with the boost.

The pions are produced at some angle from the target such that the characteristic angle varies inversely with pion momentum. As described above, the horn focusing will change these angles depending on the focusing details. For detailed calculation of the flux, the tracking simulation is used to produce mesons, focus them through the horns, and then let them decay at some point (or be lost in the shielding, if poorly focused). The decay is chosen to be in the direction of the detectors, and corresponding neutrino energies and fluxes for the detectors locations are produced. The interaction rates are the neutrino fluxes multiplies by the interaction cross sections on iron; the cross section is generally proportional to neutrino energy, above $E_\nu \sim 1$ GeV.

1.4.2 Variable Energy Beam

The NuMI target and horns system is designed to allow production of different energy neutrino beams with the same energy of incident protons. These different energy beams can be used to explore different ranges of neutrino oscillation parameters; additionally, the different beams are useful in analyzing the measurements of the Beam Monitoring System and the MINOS near detector.

There have been two modes of reconfiguration that have been extensively studied: movement of both horns and target[6] [27] to best focus high energy pions, and movement of only the target upstream of the horn system [38]. The first mode was considered in the design stages; it requires several weeks for reconfiguration and recommissioning. Three configurations of target and horn spacings were defined as the low-energy (LE), medium-energy (ME), and high-energy (HE) beams. The higher-energy beams generally produced a greater number of total neutrino interactions as the cross section is higher. Ultimately, NuMI was at first configured as the LE beam to best optimize the neutrino oscillation search.

The second mode of energy variation involves moving only the target to a different position upstream of the horns (see Fig. 1.7). The target is mounted on a remote-positioning module, so it can be moved in only a few minutes and without personnel entering the beam hall. In principle, any location of the target upstream of its nominal position is possible (within the ~ 2.5 m range set by the gantry); however, three nominal positions have been chosen. Two are established by moving the target 100 or 250 cm upstream of the nominal LE position, resulting in higher energy beams. These target positions are similar to the the medium- and high-energy beams, though have somewhat

[6]Differently designed targets are also considered for different energies of beam.

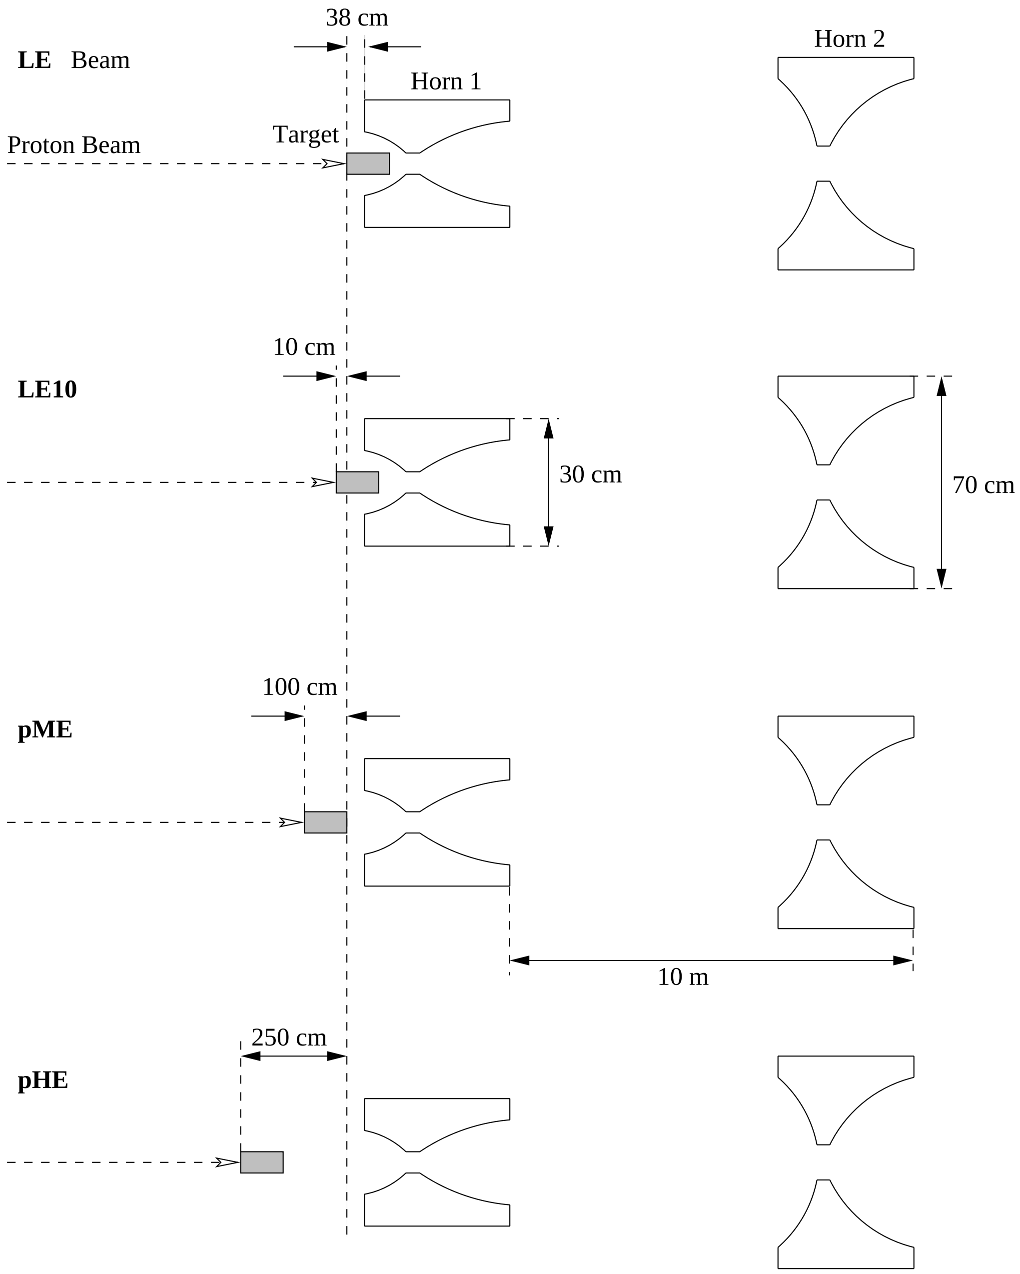

Figure 1.7: Four potential configurations of the target and horns to produce different neutrino beams. In the LE beam, the target sits substantially inside of the inner conductor of the first horn. The other beams are produced by moving the target upstream of the horns, increasing the average energy of the well focused pions. The pME and pHE beams, 100 and 250 cm upstream, produce "partial" medium- and high-energy beams.

reduced flux as Horn 2 is not relocated; as such, they are called semi, pseudo, or partial medium- and high-energy beams. More standardly, they are now as pME and pHE (instead of ME and HE).

The final configuration has the target 10 cm upstream of the low-energy position, but also a lower horn current of 185 kA. This beam was developed due to the target issues (see §10.3), and is designed to have the same peak energy as the LE beam, though with a few percent less flux. It has become known as the LE10 beam, and is the standard operating mode for NuMI.

The neutrino interaction spectra at the far detector are shown in Fig. 1.8 for the LE10, pME, and pHE beams. The peak energies increase as the target is moved upstream of the horns. The neutrino flux increases marginally, but the interaction rate increases substantially as the cross section grows with neutrino energy. NuMI has operated in each of these modes for some time. Some of the beam monitoring results from these runs are discussed in Ch. 10. The MINOS near detector has collected a substantial number of neutrino events in each run which are being used for calibration and analysis.

The neutrino beam includes anti-neutrinos and electron-flavor neutrinos, as shown in Fig. 1.9. The muon anti-neutrinos are produced mostly through the decays of negative pions. At low-energies these pions are defocused, leading to the large deficit as compared to the muon-neutrinos. At higher energies the focusing has no effect, but negative pions are produced at only $\sim 60\%$ of the rate of positive pions, and the interaction cross section is substantially smaller. Electron-neutrinos and electron anti-neutrinos are produced by certain types of kaon decay and the decays of muons. The electron-neutrinos provide an unwanted background to any searches for an electron appearance. In the NuMI LE beam, the electron contamination is $\sim 0.4\%$ at the peak.

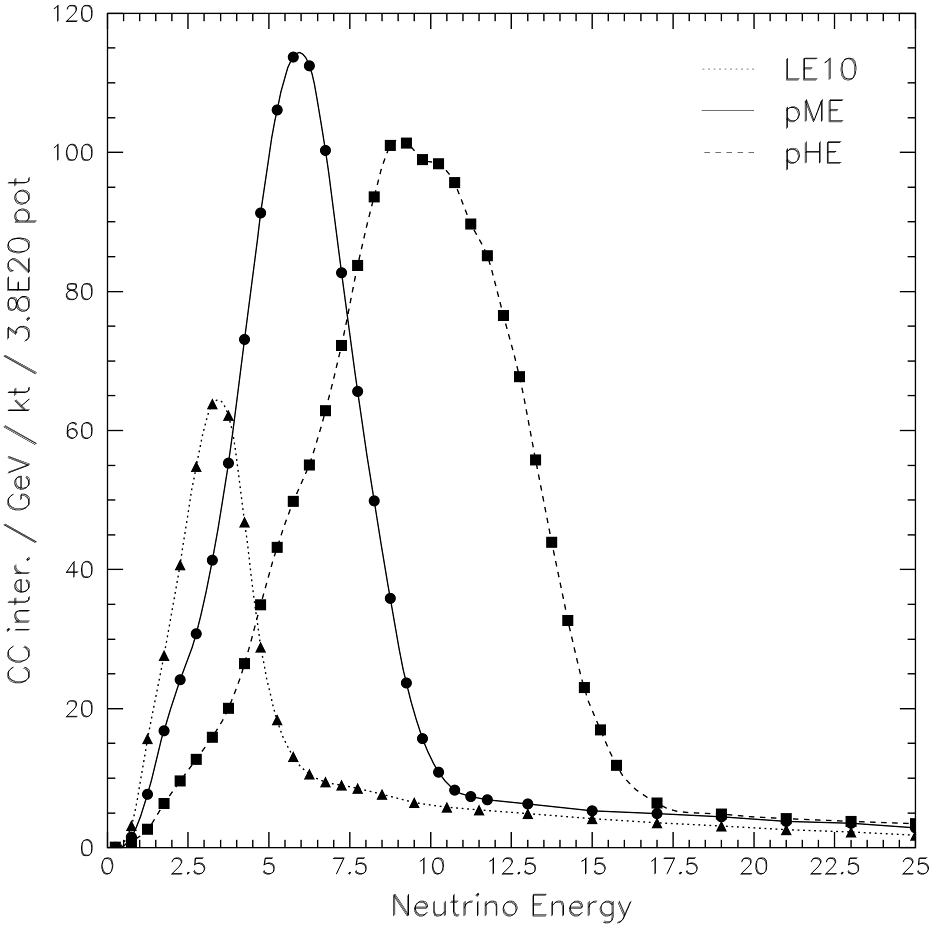

Figure 1.8: Calculated neutrino interaction rates at the MINOS far detector for three of the beam configurations shown in Fig. 1.7.

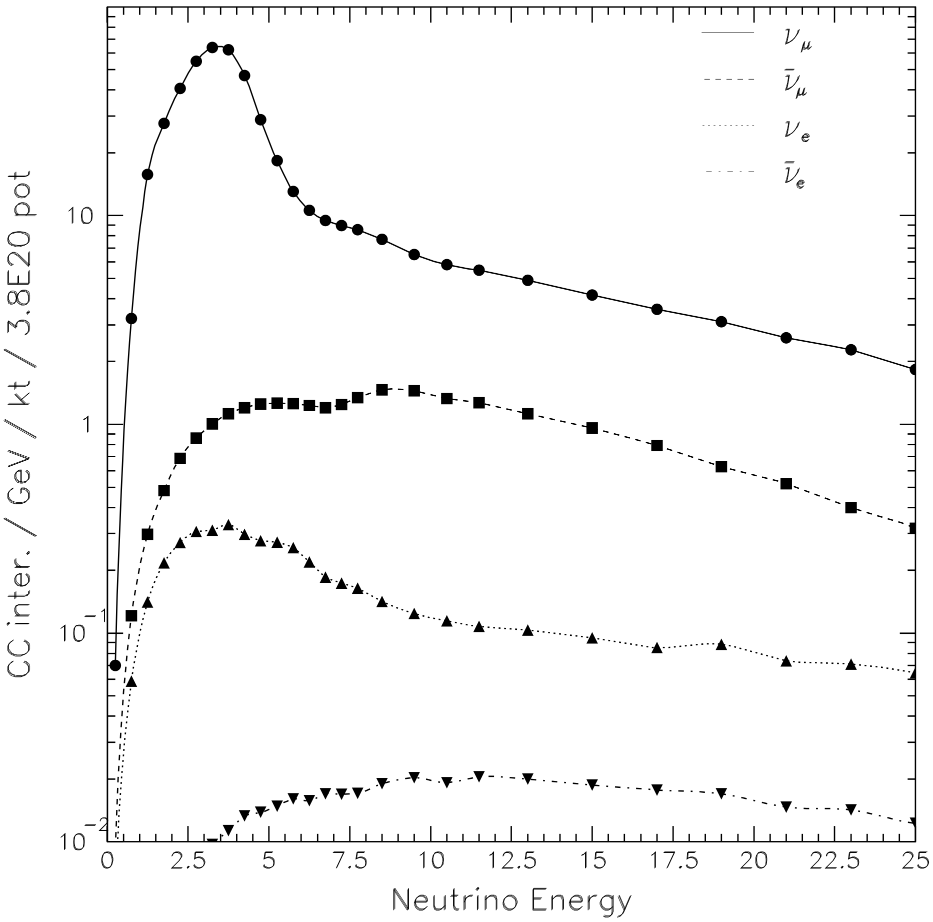

Figure 1.9: Calculated neutrino interactions in the MINOS far detector of different neutrino species (shown is LE beam without neutrino oscillations). The muon anti-neutrinos come from negative mesons that survive the horn defocusing. The electron-neutrinos are created from muon and kaon decay.

1.4.3 Limitations of Flux Calculations

Calculation of the neutrino fluxes at the near and far detectors requires accurate knowledge of particle tracking through the horns, and also the particle yields from the target. As noted in §1.3.1, such yields and their dependence on particle momentum and transverse momentum have substantial uncertainties.

Fig. 1.10 shows the calculated charged-current neutrino interaction rate for the MINOS near detector (1 km from the meson production target) for four different models of particle production from the target and with the LE focusing. The spectra were calculated using the four models, but with the same simulation of horn focusing. The beams all have a similar shape, but the difference in meson yields and angular distributions lead to neutrino rate predictions varying by $\pm$ 10% in the peak and $\pm$ 50 % in the high-energy tail of the spectra.

The neutrino spectrum measured in the far detector (in the absence of oscillations) will be similar in shape to that of the near detector. The similarity of the two spectra is quantified in the far-over-near (F/N) ratio, as shown in Fig. 1.11; it is simply the ratio of the flux spectra calculated for the two detectors[7]. The F/N ratio does vary with energy, but the variation with hadron production model is only a few percent.

The F/N ratio is expected to vary with neutrino energy. Pion decays produce an extended source of neutrinos (instead of a point source); higher energy pions will survive longer and decay closer to the near detector, so it will see relatively higher flux. This first effect is indicated by the black line superimposed on Fig. 1.11, which is calculated by assuming that a pencil beam

[7]Due to limitations in the focusing capabilities of a two horn system, the F/N ratio does not follow the ideal curve shown in Fig. 1.11; it is calculated for perfectly focused pions traveling parallel to and coincident with the beam axis down the decay volume.

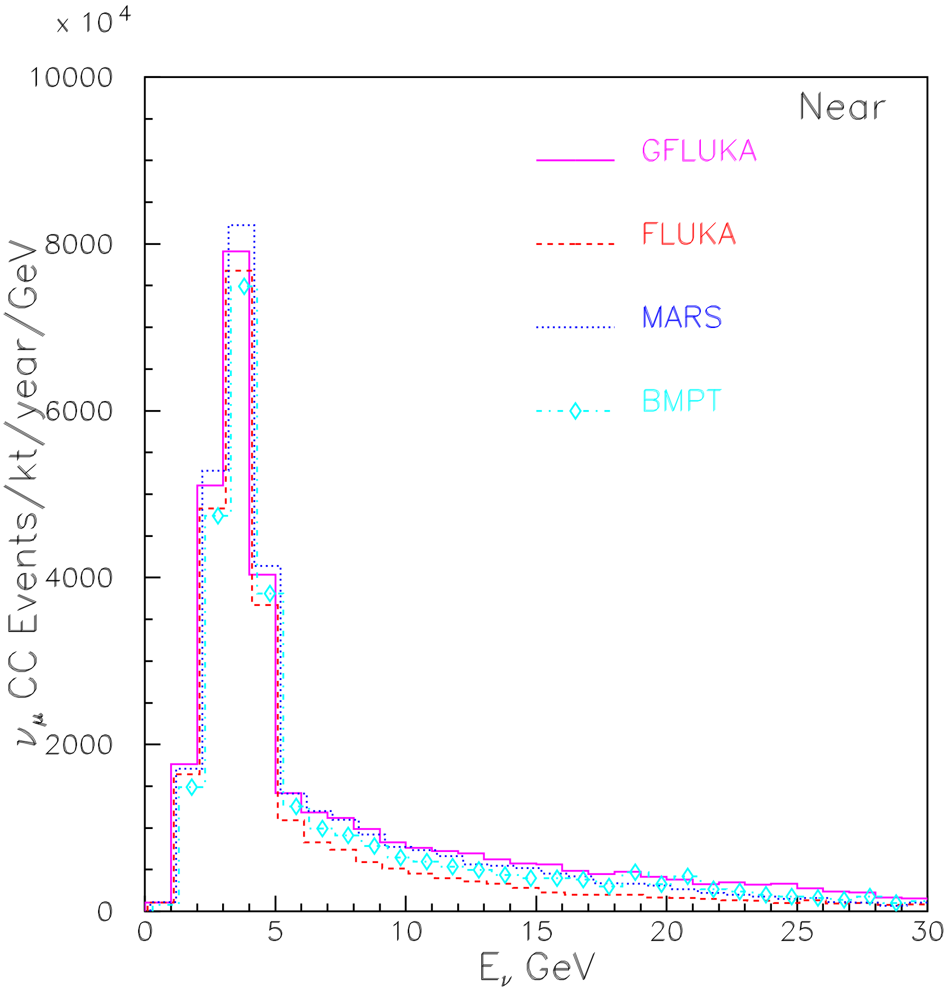

Figure 1.10: Calculated neutrino interaction rate in the MINOS near detector for the low-energy (LE) beam. In this case, 1 "year" corresponds to 3.8×10^{20} protons delivered to the target. The four curves plotted correspond to the rates predicted as a result of four different models of hadron production in the target; all used the same tracking algorithm thereafter. The predictions have a spread of $\pm$ 10% in the peak, and $\pm$ 50% in the high-energy tail.

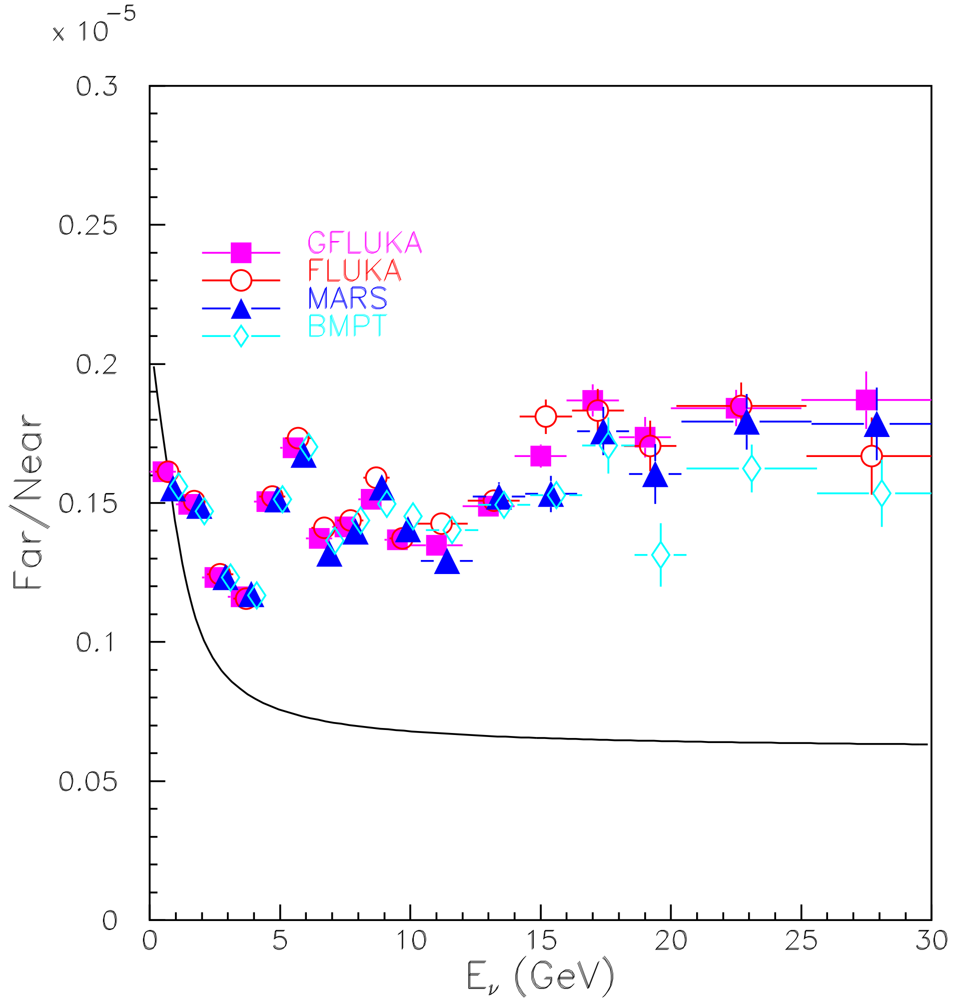

Figure 1.11: The ratio of predicted fluxes in the MINOS detectors as a function of energy (shown in the LE beam without neutrino oscillations). The plotted curve is a simple calculation in the (unrealistic) case where pions are focused parallel to and directly on the beam axis. The Monte Carlo predictions consistently lie above the curve as pions will generally have smaller decay angles to the far detector than near. Note that while the model predictions for interaction rate varied considerably in Fig. 1.10, their predictions for the ratio lie within a few percent in the peak energies.

of pions were colinear and coincident with the beam axis. The second effect is that the pions generally have a smaller decay angle to the far detector than near detector – producing relatively higher flux and neutrino energies in the far detector. The different angles arise because the pions are created and focused at small distances from the beam axis, but while subsequently drifting in the decay region the pion angle will generally produce a larger radial distance. This distance is highly correlated with the angle. At the time of decay the angles to the detectors must include that radial offset, such that $\phi_r \approx r/L$. For the far detector $L = 735$ km, so the angle is near zero; for the near detector, however, L is on the order of 500 m, so with a radial offset of as much as 1 m, an additional angle of 2 mrad could be required. For a 10 GeV pion, this could reduce the neutrino energy at the near detector by 6%, and the solid-angle flux by 12%.

1.5 Decay Volume and Beam Stop

The NuMI decay volume lies downstream of the target station and is the location where the bulk of pion decays occurs. The decay volume is a 675 m long, 2 m diameter, evacuated steel pipe. The pipe is encased in concrete for radiation protection, and must be water-cooled to conduct away the energy deposited by beam particles.

The pions entering the decay volume fly straight line paths until their decay. As mentioned above, these pions could be further focused within the decay volume to prevent wall interactions – a device to do this, known as the Hadron Hose, is described below.

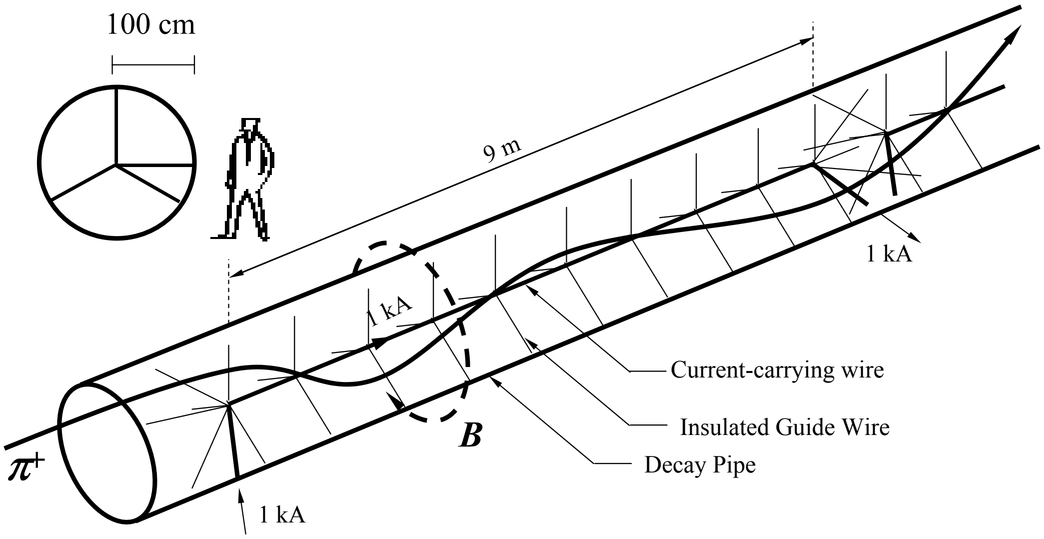

Figure 1.12: Schematic of a Hadron Hose segment in the NuMI decay pipe. Each segment consists of a 1 kA current-carrying wire that must enter from the edge of the pipe, travel down the pipe for 9 m (with supports every meter), and exit at the edge of the pipe. Positive particles are attracted to the hose wire and spiral about it.

1.5.1 The Hadron Hose

The Hadron Hose was proposed as a pulsed magnetic focusing device that extended for the entire length of the decay volume. The focusing field would be produced with a 1 kA pulsed current carried by a wire, broken into segments along the length of the decay pipe. The toroidal magnetic field would focus positive particles, forcing them to spiral about the wire. The focusing of the pions would increase the neutrino rates at the detectors and make the spectra more similar in shape.

The hose would physically consist of an aluminum wire 2.4 mm in diameter, broken into segments of $\sim$ 9 m, as shown in Fig. 1.12. Each segment would be connected in parallel via transformers to a transmission line, inducing the 1kA current. The hose wire was to be supported and positioned every

1 m by a set of guide wires. Misalignment of the hose wire by more than 2 mm along its length could lead to neutrino flux loss and distortion [40].

The hose's magnetic field varies inversely with the radial distance, and particles experience the field in the whole of the decay volume. As such, the hose does not act like a linear lens, having a much stronger force on particles near it, and forcing a spiraling motion. The focusing force causes some portion of the pions to spiral around the wire, instead of coming to a stop in the wall of the decay pipe. Additionally, the spiraling motion reduces the correlation between transverse position and angle, making the decay angles to the two detectors more similar. As a result, the shape of the neutrino spectra are more similar, reducing the potential uncertainties from such effects [41].

The effect of the hose on calculated neutrino interaction rates is shown in Fig. 1.13 for the low-energy beams. The hose focusing increases the neutrino flux, particularly in the high-energy tail. Additionally, the spectral shapes are more similar.

For higher-energy beams, the flux increase is not as substantial, but the spectral shapes are made much more similar. The angle-position correlation of drifting pions results in a significant deficit in the near detector spectrum at the falling edge of the energy peak. This deficit causes the predicted ratio of fluxes to change by a factor of three, and back, over the range of a few GeV. Failure to precisely model the spectra could results in a false-positive signal. The hose, by reducing the correlations, makes the spectra much more similar in shape, reducing fluctuations in the ratio by a factor of ten.

Ultimately, the hose was not included in the NuMI beam. The research and development toward an operational hose is recorded in [20].

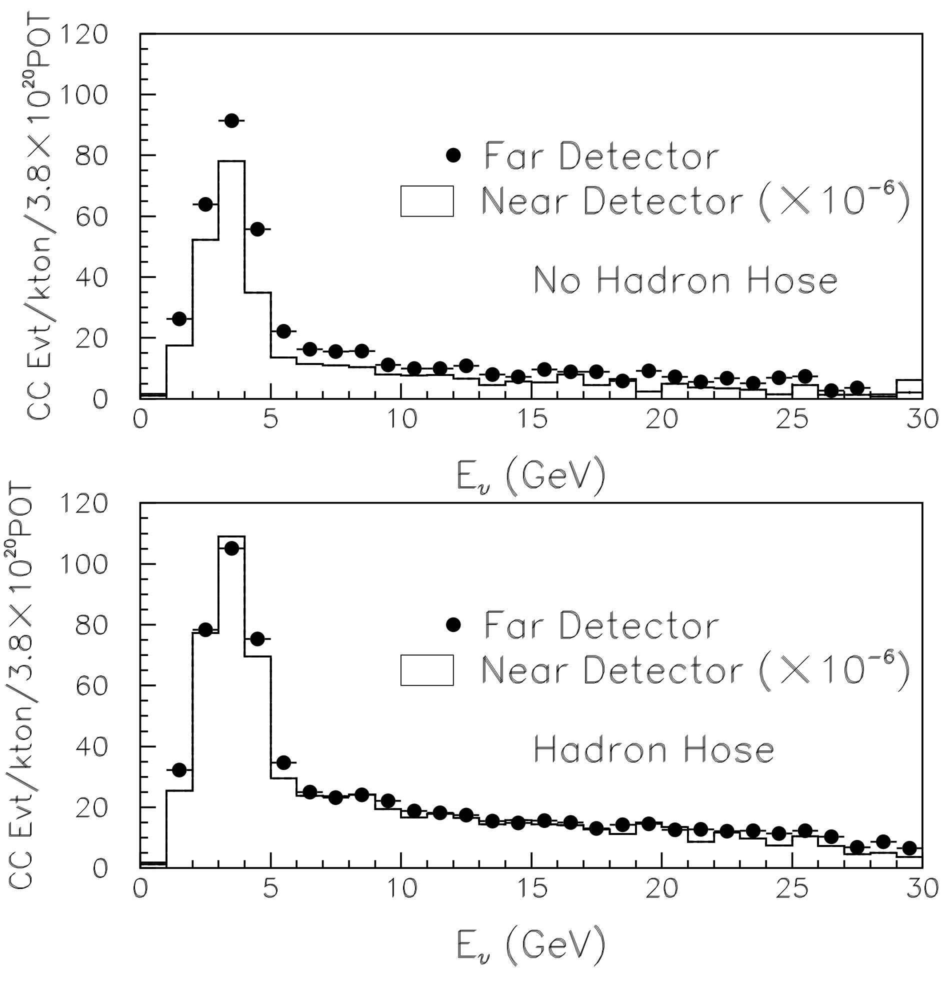

Figure 1.13: Calculated effects of the Hadron Hose on the low-energy (LE) beam spectra. The relative rates of the peak and tail are more comparable, and the peak energies closer as well. With the hose, the interaction rate is $\sim$ 20% higher in the peak and 50% higher in the tail.

1.5.2 Hadron Absorber

At the end of the NuMI decay volume is a large beam dump known as the Hadron Absorber. The Hadron Absorber stops the protons that do not interact in the target, as well as all hadrons that are produced but do not stop or decay before reaching the end. It is designed to stop beam at a rate of 160 kW in the steady-state, and to be able to survive the full 400 kW beam power for up to an hour (in case of emergency situations). To fully stop the high energy hadrons and absorb the beam power, it is constructed with a water-cooled aluminum core. The core is surrounded by shielding steel, and further surrounded by concrete for neutron protection.

1.5.3 Beam Monitors

The NuMI Beam Monitoring System measures the remnant hadron and tertiary muon beams (produced with the neutrinos) at the end of the decay volume. The Beam Monitors provide spill-to-spill information on target integrity, proton beam placement, and muon production; they can also be used in special runs for alignment purposes and other systematic crosschecks (e.g., horn scaling). The Beam Monitors are composed of a Hadron Monitor that sits upstream of the absorber and measures the spatial distribution of the hadron beam; and three sets of Muon Monitors that sit downstream of the absorber and in two other stations in the rock to measure the spatial distribution of the muon beam. The system is the subject of Chs. 6-11 of this dissertation.

1.6 MINOS

The NuMI beam's initial purpose is to provide a neutrino beam for a precision measurement of the neutrino oscillation parameters associated with muon-neutrino disappearance, and to eliminate alternative models of neutrino disappearance. The The Main Injector Neutrino Oscillation Search (MINOS) experiment will measure the neutrinos in the beam to make such a measurement. MINOS is composed of two neutrino detectors, one at Fermilab and one in the Soudan mine in Minnesota. The detectors are designed to be similar in composition and operation so as to reduce uncertainties in comparison. To observe oscillations, MINOS will measure the disappearance of muon-neutrinos over its 735 km baseline.

In the case of muon-neutrino disappearance, the dominant process is expected to be transitions from muon to tau flavor neutrinos, with negligible transitions to electron-neutrinos. The transition arises from the inference that muon-neutrinos are composed of a mixture of neutrino mas states. The mass states have different masses and their phases advance at different rates, thus the mixture changes. The different mixture will then have the possibility to interact as either a muon or tau neutrino. In this context, and in light of the experimental precision, neutrino oscillations are considered as a two-flavor approximation with the transition probability:

$$P_{\nu_\mu \longrightarrow \nu_\tau} = \sin^2 2\theta_{23}\ \sin^2 \left[1.27\, \frac{(\Delta m_{23}^2 \,/\, (\mathrm{eV/c^2})^2\,) \times (L \,/\, \mathrm{km})}{E_\nu \,/\, \mathrm{GeV}} \right] \tag{1.8}$$

$\sin^2 2\theta_{23}$ is the mixing fraction of μ and τ flavor neutrinos[8]; Δm_{23}^2 is the difference between the squares of the masses of the neutrino mass states[9]; L is

[8] θ_{23} is also the mixing angle in the unitary transformation matrix between mass and interaction flavor states. The mixing fraction is a more empirical concept. Currently, maximal mixing is favored by atmospheric neutrino observations: $\sin^2 2\theta_{23} \approx 1$.

[9] From atmospheric neutrino measurements, Δm_{23}^2 is on the order of 0.002 eV/c^2.

the baseline length from neutrino production to interaction[10]; E_ν is the energy of muon-neutrino in the lab frame. The results of the MINOS experiment will be arrived at by measuring the energy-dependent disappearance of the muon-neutrinos.

The MINOS detectors measure charged-current interactions of the muon-neutrinos in the NuMI beam:

$$\nu_\mu + \mathrm{N} \longrightarrow \mu^- + \mathrm{N}^* \tag{1.9}$$

Where N is a nucleon, and N^* becomes a hadronic shower. The detectors act simultaneously as spectrometers and sampling calorimeters to measure the momentum of the produced muon and the energy of the hadronic shower.

The detectors are constructed of consecutive planes of steel and segmented scintillator. The steel acts as a neutrino target, and the scintillator strips sample some portion of the deposited ionization energy. Energetic particles will leave tracks that can be reconstructed and the length used as a measure of the muon's initial energy. Additionally, the steel is magnetized by an electric coil, such that the muons are forced to bend. The bending of the the muon track provides charge identification and another measure of the muon momentum; this measure is particularly useful for energetic muons that might leave the detector. The total detector mass is 980 tons for the near detector, and 5400 tons for the far detector.

Fig. 1.14 shows simulated data corresponding to the level of statistics likely used in MINOS's first analysis. The number of neutrinos simulated is for 10^{20} protons delivered on target, which should be achieved by the end of 2005. 10^{20} protons will provide several million measured neutrino interactions

[10]The NuMI baseline is the 735 km distance from the NuMI decay volume to the Soudan Underground Laboratory.

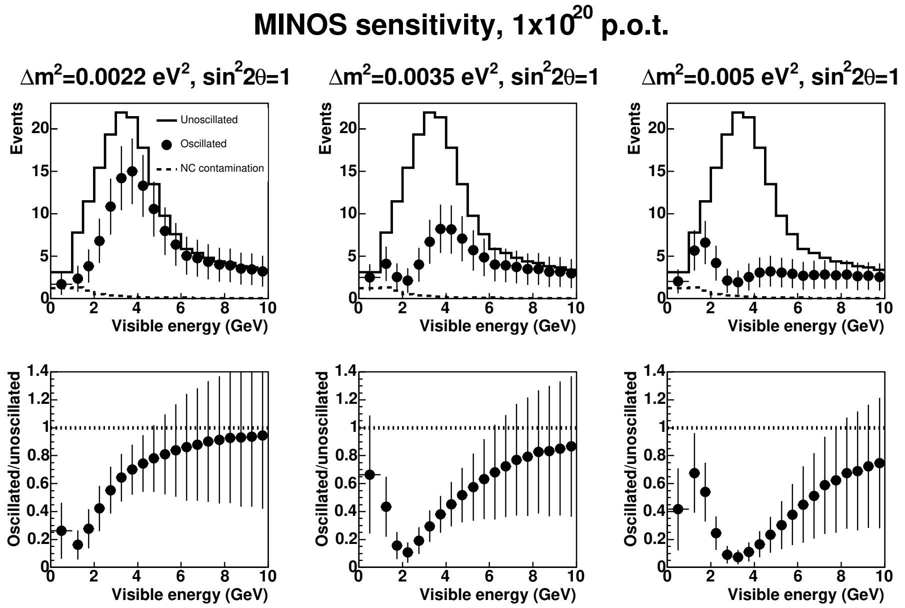

Figure 1.14: Predicted sensitivity of the MINOS experiment to neutrino oscillations when 10^{20} protons have been delivered to the NuMI target. The solid lines in the top plots show an extrapolated neutrino interaction spectrum at the far detector, using near detector data and beam simulation. The points with error bars show simulated data for different oscillation parameters. The ratio of each pair of spectra are shown in the lower set of plots.

in the near detector, and several hundred in the far detector (in absence of oscillations). The precision of the mass measurement should be comparable or greater than that of previous experiments. For higher mass differences the sinusoidal variation with E^{-1} will also be apparent.

Some eventual sensitivities expected in MINOS are shown in Fig. 1.15. The total number of protons to be delivered to MINOS has not yet been determined, but the experiment can improve its measurements up to at least $25{\times}10^{20}$ protons. The eventual statistical precision will allow measurement of neutrino disappearance to a high significance and provide discrimination between neutrino oscillations and alternative models of neutrino disappearance.

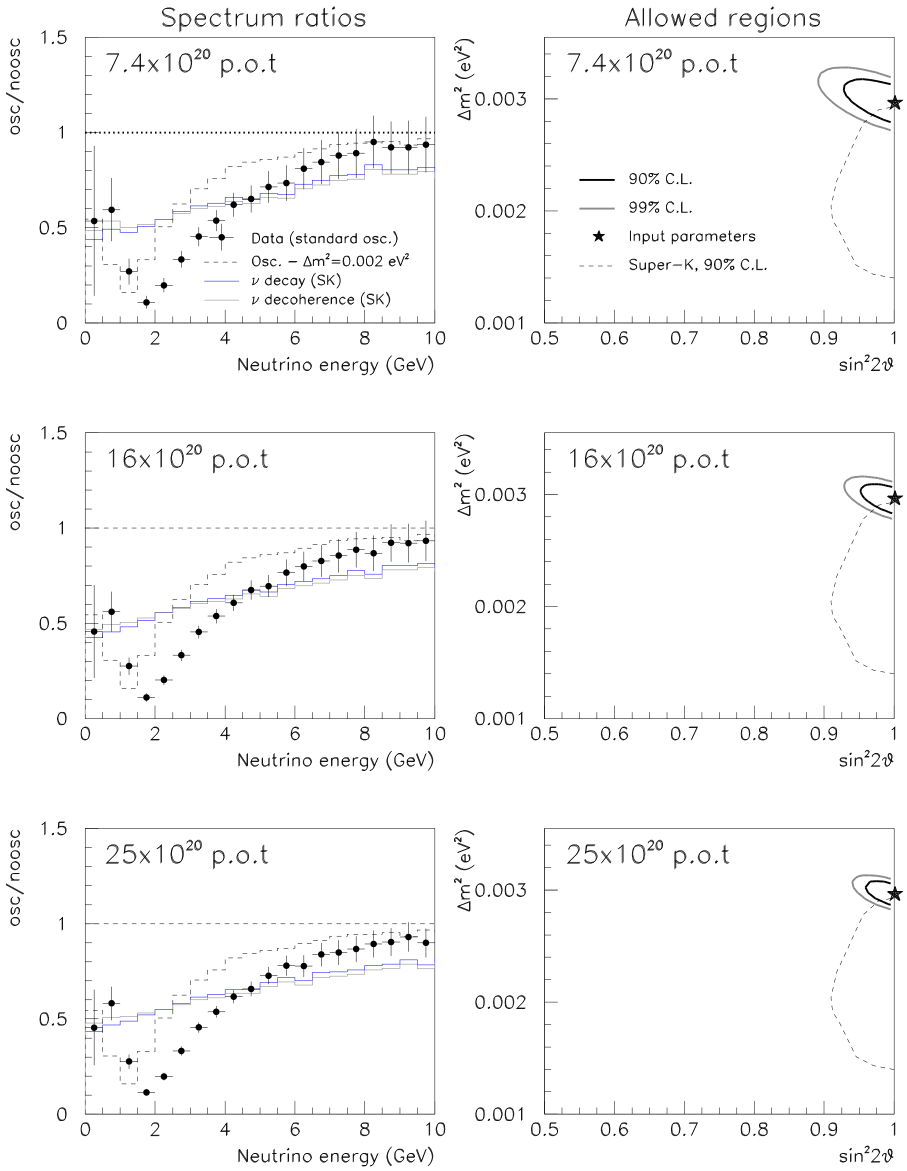

Figure 1.15: Simulated sensitivity of MINOS to oscillations, and discrimination from alternative models using a value of $\Delta m^2_{23} = 2.95\times10^{-3}$ eV^2. The left plots show ratios of the simulated spectra to the case of no oscillations for different levels of statistics. Also shown are the ratios expected for a different set of parameters and two alternative models of neutrino disappearance. On the right is plotted the 90% and 99% confidence levels achieved with the simulated MINOS data (solid lines), and the 90% confidence level achieved by SuperK with their zenith angle analysis (dashed line).

Additionally, in the case of neutrino oscillations, MINOS will make a precision measurement of the mass parameter, Δm^2_{23}.

In the future, the MINOS experiment is anticipated to perform other oscillation analyses than muon-neutrino disappearance. Neutral-current interactions will be measured to establish whether neutrinos are changing into other flavors. Also, electron-neutrino interactions will be measured to investigate oscillations through appearance. Furthermore, additional experiments are possible using the NuMI beam. These other measurements will require further understanding of the NuMI beam composition.

1.7 Outline of This Dissertation

This dissertation is concerned with maximizing the potential of the NuMI beam for oscillation measurements. The efforts detailed are split logically into two parts: (1) efforts to synchronize the Booster and Main Injector accelerators at Fermilab to produce high-intensity proton beams for NuMI; and (2) design and use of the secondary Beam Monitoring system in commissioning and operating the NuMI beam.

Chs. 2 – 5 make up the the accelerator synchronization portion of the dissertation. Known more colloquially as "Booster Cogging", this system allows the stacking of multiple batches of 8 GeV protons from the Booster into the Main Injector, so that they can be simultaneously accelerated to 120 GeV. Ch. 2 provides background on the Fermilab accelerator complex, and Ch. 3 discusses the Booster accelerator in detail (App. A also describes the accelerator physics concepts in use). Ch. 4 is concerned with analysis of the Booster acceleration cycle, possible perturbations to it, and measurements of variation. Ch. 5 details the implementation of a system to counteract variation in the

Booster, and properly transfer beam to the Main Injector.

Chs. 6 – 11 constitute the neutrino Beam Monitoring part of the dissertation. Ch. 6 describes the utility and design of the NuMI Beam Monitoring system. Ch. 7 reviews the principles and data governing the operation of ionization detectors. Ch. 8 provides simulations of ion pulse formation in ionization chambers, as well as data from a prototype beam test (also: App. C includes calculations of pulse formation, App. D contains the details of the numerical simulation, and App. E provides additional simulation results). Ch. 9 documents the initial performance of the Beam Monitoring detectors in the NuMI beam (App. F provides additional measurements). Ch. 10 describes the use of the Beam Monitors to monitor the neutrino beam. Ch. 11 details the use of the Beam Monitors and other instrumentation to align the proton beam with neutrino beam components (App. B provides additional information on the proton beam and instrumentation).

Chapter 2

The Fermilab Accelerators

Fermilab's particle accelerators consist of a linear section, six rings, several transfer lines, and other components. The accelerators accelerate, and in some cases generate, the high-energy particles used in the experiments. Different experiments place different demands on the accelerators – for MINOS the demand is for the maximum number possible of 120 GeV protons, within certain limits on beam parameters.

NuMI uses 120 GeV protons from the Main Injector (MI) accelerator. The protons are extracted in a single turn, delivering pulses of $\sim$ 8.6 μs in duration. NuMI is designed for as many as 4×10^{13} protons per pulse.

As shown in Fig. 2.1, the MI is fed batches of protons from the 8 GeV Booster accelerator. Because of its larger circumference, the MI can accept six batches of Booster beam, provided these batches can be injected with some longitudinal separation[1].

The Main Injector accelerates beam from 8 GeV to 120 GeV in $\sim$ 1 s, whereas the Booster is a rapid-cycling synchrotron, which can deliver 4-5$\times10^{12}$

[1]The MI circumference is actually 7 times that of the Booster; one Booster batch's space must be kept empty in the MI to allow the extraction kickers to fire.

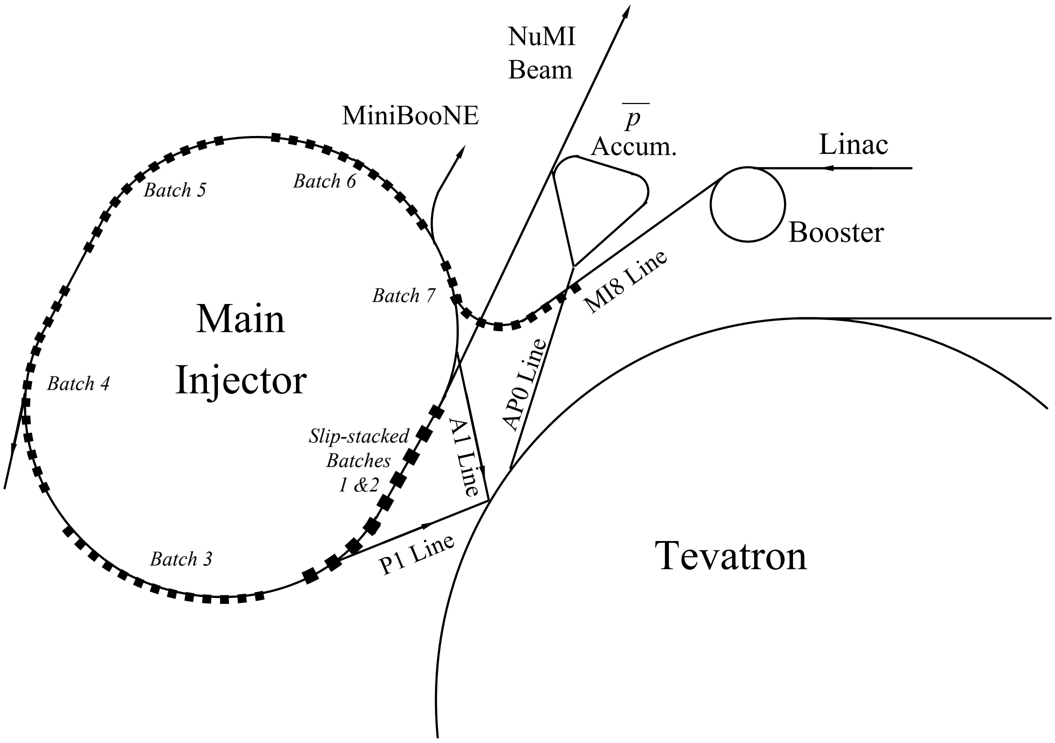

Figure 2.1: The Fermilab accelerators. Those used in generation of the NuMI proton beam are the Linac, Booster, and Main Injector.

protons every 67 ms. To make efficient use of the MI acceleration time, the MI accelerates beam both for extraction to the NuMI target and for the production of antiprotons. At MI flat top, a set of kicker magnets extracts ~ 84 bunches of beam for antiproton production, and another set of kickers delivers ~ 420 bunches to the NuMI beamline.

A method of accelerating more than six Booster batches has been developed and is called slip-stacking [42]. In this technique two Booster batches are injected into the MI, coalesced longitudinally into one double batch, and five additional 8 GeV batches injected. At present, this scheme is utilized for delivery of intense double batches for antiproton production, but will be developed to produce more intense NuMI beams in the future. A timeline of

Accelerator	Particle	Extraction Energy	N_p	Length
Pre-acc	H^-	750 keV	10^{13}	7 m
Linac	H^-	400 MeV	10^{13}	$\sim$ 300 m
Booster	p	8 GeV	5×10^{12}	475 m
Main Injector	p & $\bar{p}$	120-150 GeV	3×10^{13}	3320 m
Tevatron	p & $\bar{p}$	980 GeV	10^{13}	6290 m

Table 2.1: Fermilab accelerators, the particles accelerated, the maximum energy produced in each machine, a typical number for the number of protons that can be accelerated at once, and the length or circumference of the machine. All but the Tevatron are used to prepare NuMI protons.

this slip-stacking process followed by injection of beam for NuMI is shown in Fig. 2.2.

Table 2.1 contains some basic information about the accelerators used to accelerate beams[2]. All of the accelerators, except for the Tevatron, are used to generate the NuMI beam. Each machine depends on the previous for beam; since the later machines are longer than the previous machines several injections of beam can be used to fill a later machine. Additionally, the operation of the other accelerators and experiments depend on some of the same machines, such that influences are felt across the entire chain.

Chs. 4 – 5 of this dissertation describes a method of synchronizing Booster batches for transferal into the MI. Such is essential for delivering beam pulses of six Booster batches to NuMI. Because the Booster and Main Injector operate independently, a feedback system was developed to measure the location of beam in the MI and Booster, and to time the arrival of new Booster batches into the MI.

This chapter provides an overview of all of Fermilab's accelerators. Ch.

[2]One accelerator in the Antiproton Source does not provide any overall acceleration to the beam. This machine, the Debuncher, uses RF acceleration to longitudinally focus a secondary beam; the momentum spread is reduced at the expense of bunching, increasing the number of collected particles.

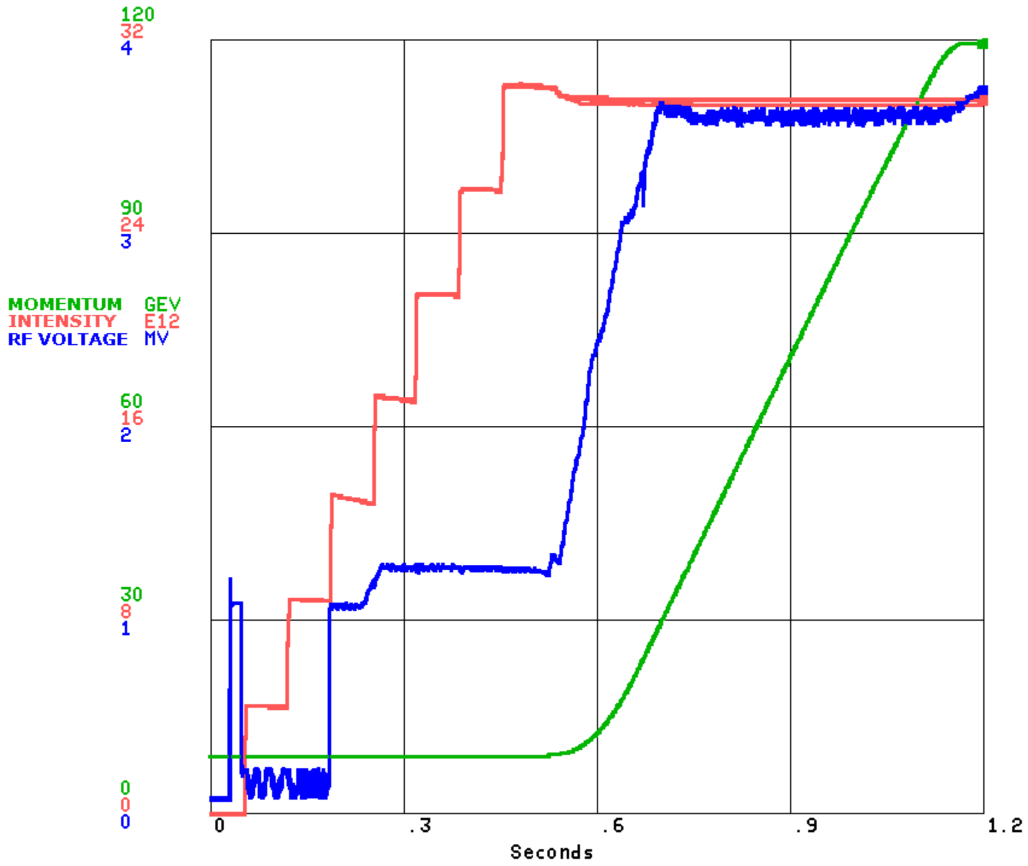

Figure 2.2: Injection of the Booster beam in the Main Injector for several MI Cycles. Overlaid on top of each other are three traces: the Main Injector charge (red) that increases seven times in seven separate injections from the Booster; the MI beam momentum (green) that is defined by the magnet currents, it ramps from 8 to 120 GeV/c; and the RF voltage (blue) that is low initially for matching to the Booster longitudinal phase space, but is increased to its maximum during the start of acceleration.

3 discusses the Booster accelerator in particular, since it will be used to provide this feedback. Ch. 4 describes timing errors that arise between the Booster and the Main Injector. Ch. 5 describes the forced synchronization that overcomes these errors. For reference, App. A defines many of the accelerator physics concepts used in these chapters.

2.1 Preaccelerator

The preaccelerator starts with hydrogen gas and provides an intense, DC H^- beam to the Linac. The preaccelerator has two primary components: an ion source that produces the H^- ions from the gas, and a Cockroft-Walton accelerator that accelerates the ions to 750 keV. Fermilab has two preaccelerators such that one can be serviced without significantly affecting operation. An outline of the preaccelerator with the ion source, potentials, and other equipment is reproduced in Fig. 2.3.

Ion Source

The Ion Source produces H^- beam pulses at 50 mA for 60 μs, at a repetition rate of 15 Hz [43, 44, 45]. The H^- ions are generated in a magnetron plasma source (shown in Fig. 2.4). Electrons are culled from a low work function surface via field or thermionic emission and forced to gyrate in a magnetic field, effectively confining the electrons near the surface. Hydrogen gas is fed into the gyrating electrons, ionizing the gas and creating a dense plasma. Hydrogen ions are attracted to the low work function cathode and create negative hydrogen ions via a sputtering-like process.

The ions are extracted from the source with a 20 kV electric potential and are passed through a 90° bending magnet. The magnetic field removes free

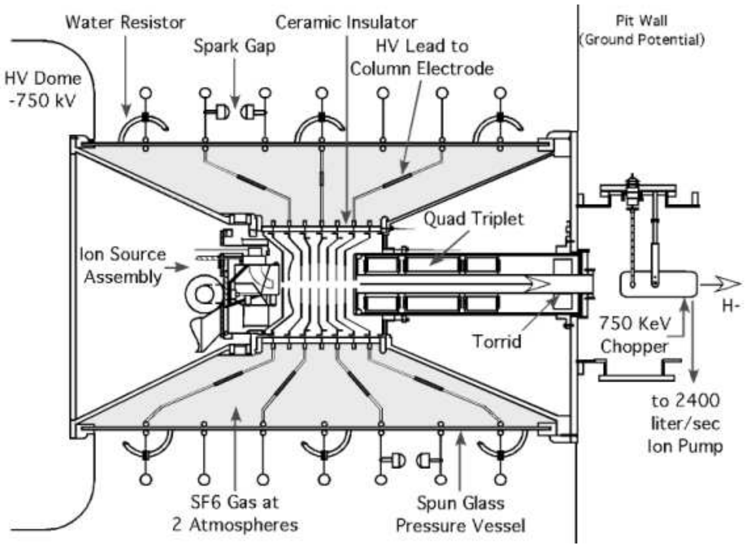

Figure 2.3: Diagram of the Fermilab Preaccelerator. The dome is kept at +750 keV from ground by the Cockroft-Walton. The ion source within the dome produces a 20 keV H^- beam that is focused and transported down a column to ground, accelerating the beam to 750 keV. (from [43])

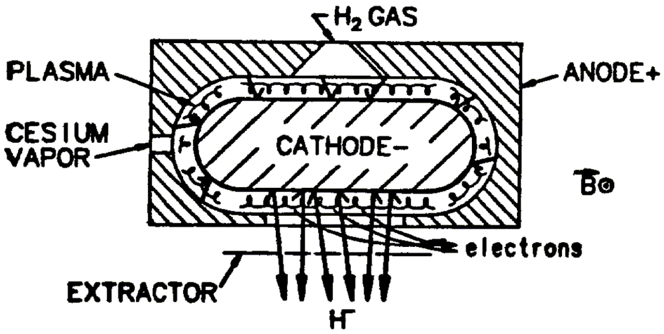

Figure 2.4: Schematic of a magnetron ion source. A dense plasma is maintained near a low work function cathode by electric and magnetic fields. Plasma protons bombarding the cathode can become negative ions. Negative particles leave the plasma via an electric field, electrons are filtered out of the beam by a magnetic field. (from [43])

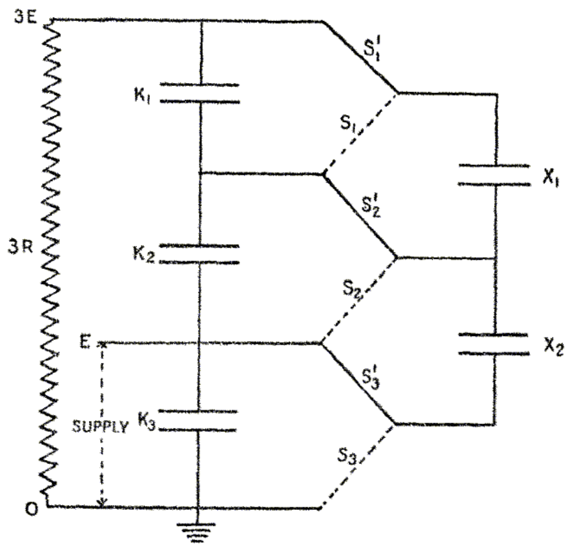

Figure 2.5: Cockroft-Walton voltage multiplier network. An electric current source provides charge to the lowest stages. Current flows up the network through switches, S_x (diodes in more modern systems). The maximum potential difference available is the applied voltage multiplied by the number of stages. (from [46])

electrons that are extracted along with the H^- as well as focusing the beam such that it has a mostly round profile when passing through the accelerating column; the normalized beam emittance is ~ 1 mm·mrad (90%). Some cesium and neutral hydrogen escape from the plasma (P $\sim$ 0.4 torr)

Cockroft-Walton

The Cockroft-Walton is a DC accelerator that provides the -750 kV potential that accelerates the H^- ions from the ion source. It achieves the 750 kV through a voltage multiplier network first discovered by Cockroft and Walton [46] and shown schematically in Fig. 2.5.

The Fermilab Cockroft-Walton accelerates particles by maintaining a "dome" at -750 kV from ground. The dome is connected to ground via an evacuated, glass cylinder. The cylinder is surrounded by field-shaping electrodes to maintain a uniform electric field. H^- ions are created within the dome by the ion source described above, and transported to the accelerating column. Passing through the column they acquire the 750 keV of kinetic energy. On exiting the column the ions are transported to the next accelerator by a magnetic bending and focusing transfer line.

2.2 Linac

The Fermilab Linac[3] accelerates the H^- beam from 750 keV to 400 MeV for injection into the Booster. It is composed of two different linear accelerators, two additional RF cavities, a magnetic focusing lattice, and numerous pieces of instrumentation. The two linear accelerators are in series and serve to provide the bulk of the acceleration. The two cavities are before and after the accelerator; the first cavity bunches the DC beam from the preaccelerator, and the the second debunches the beam for injection to the Booster.

	Initial Energy	Final Energy	Frequency
Low-Energy Linac	750 keV	116 MeV	201.24 MHz
High-Energy Linac	116 MeV	401 MeV	805 MHz

Table 2.2: Parameters for the two components of the Fermilab Linac. The two linear accelerators are connected in series between the preaccelerator and Booster.

Some of the parameters of the linacs are collected in Table 2.2. The Low-Energy Linac (LEL) was initially the only linear accelerator, which brought beam to 200 MeV for injection into the Booster. In the 1990s the High-Energy Linac (HEL) was added to upgrade the linac to 400 MeV, allowing more beam

[3]Short for Linear Accelerator

to be accelerated in the Booster. As part of the Linac Upgrade part of the LEL was removed and replaced by HEL components, such that the transition between machines occurs at 116 MeV.

Buncher

The Buncher is a single RF cavity that introduces 200 MHz structure to the DC beam produced by the preaccelerator, dramatically increasing the portion of the beam that is successfully accelerated. Without bunching some of the beam will be at the wrong phase with respect to the main RF, not be accelerated, and lost at some point along the beamline.

A single RF kick is delivered to the 750 keV beam at 202.24 MHz. This causes some to speed up, and some to slow down. A subsequent drift region allows these pieces of the beam to bunch together. By correctly choosing the bunching voltage, phase, and drift region length the beam is matched longitudinally into the main Linac. Use of the Buncher allows $\sim 75\%$ capture efficiency, whereas only $\sim 25\%$ would be captured from the laminar beam.

Low-Energy Linac

The Low-Energy Linac (LEL) is an Alvarez-type drift-tube linac (DTL) [47]. It has several RF "tanks". Each tank acts as an RF resonator and is supplied by a single vacuum-tube power amplifier. Within the tanks there are "drift tubes" placed periodically; the beam passes through their centers and is shielded from oppositely directed electric field.

The cavities are operated at roughly 200 MHz. When the beam reaches 116 MeV it has been compressed longitudinally to a small portion of the waveform, which will allow the conversion to higher frequency RF.

High-Energy Linac

The High-Energy Linac (HEL) is a side-coupled linac (SCL) [48]. A SCL is made up of a string of cavities that are lowest-order resonators of the supplied RF frequency. The accelerating cavities are coupled to each other by auxiliary cavities that act to slow down the phase velocity of the RF wave, and allow adjacent cavities to be at opposite phases with respect to each other.

As mentioned above, the beam has shrunk significantly on entering the HEL, so 800 MHz is used to accelerate the beam – exactly four times that of the LEL. The beam is still 200 MHz, however, so three out of four RF buckets are empty of beam. The higher frequency provides several advantages: smaller cavity size, higher gradients, and the freedom to use Klystron RF power supplies [49, 50] instead of vacuum-tube power amplifiers.

Debuncher

The Debuncher does the opposite job as the buncher. The Debuncher cavity sits on the transfer line between the Linac and Booster. The beam has drifted for some distance, allowing its longitudinal position within the RF buckets to become correlated with momentum. The Debuncher cavity provides a kick that reduces the momentum of the fast portion of the beam , and increases momentum of the slow portion. The debunching process thus reduces momentum spread at the expense of lost bunching. The bunching, however, would serve no purpose as the Booster operated at a much lower RF frequency than the Linac.

2.3 Booster

Because of its central role in the upcoming chapters, the Booster is described more fully in Ch. 3. For the purposes of this section we note that the Booster is a synchrotron that captures the debunched beam from the Linac after converting the H^- ions to protons and accelerates them to 8 GeV [51, 52]. The Booster cycle lasts 33 ms and the Booster magnets cycle at 15 Hz – only some portion of which actually accelerate beam.

The Booster beam is extracted to the Main Injector in a single turn by a series of fast kicker magnets. The Booster RF frequency at extraction is 52.8114 MHz ("53 MHz") and must be matched precisely to the Main Injector frequency and phase as the beam will be injected directly into the Main Injector RF buckets.

2.4 Main Injector

The Main Injector (MI) [53] accepts beam from the Booster, Recycler, and Antiproton Source. The MI delivers beam to NuMI, Recycler, Antiproton Source, Tevatron, and Switchyard. Most frequently, it accepts 8 GeV protons from Booster and accelerates them to 120 GeV for antiproton or neutrino production. The Main Injector is optimized for its central role by being designed to accelerate beam as quickly as possible: the magnets ramp as quickly as 240 GeV/s and the RF can supply up to 4 MV/turn. The Main Injector lattice is a separated-function strong-focusing lattice; the quadrupoles are reused from the Fermilab Main Ring [54] and the dipoles were made newly.

The Main Injector accepts protons from the Booster and increases their energy 15 times (to 120 GeV). The MI is exactly seven times the circumfer-

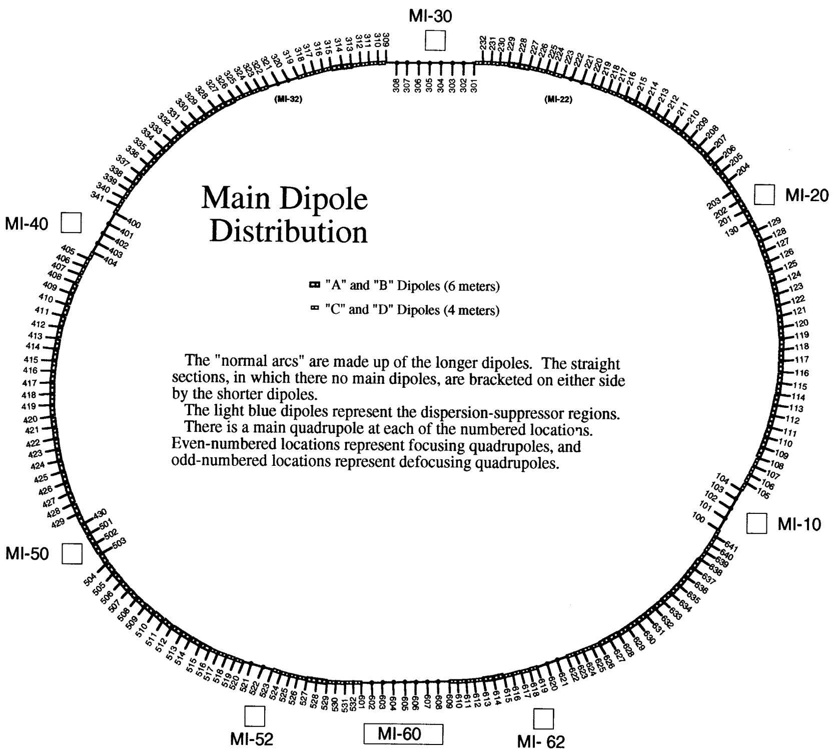

Figure 2.6: Overview of the Main Injector. The majority of the circumference is filled with bending and focusing magnets. The straight sections (usually numbered MI-XX) serve the purposes of injection, extraction, and acceleration.

ence of the Booster, allowing a larger bending radius so conventional magnets can provide the bending field. The circumference also allows the MI to accept multiple injections from the Booster and accelerate them simultaneously. Stacking is the general term used in accelerators for a process that increases the beam current in the machine by adding more particles. In the MI there are several methods of stacking, all of which combine multiple Booster batches into one beam; doing so requires precise synchronization of the Booster and MI, which will be the subject of the subsequent chapters of this thesis.

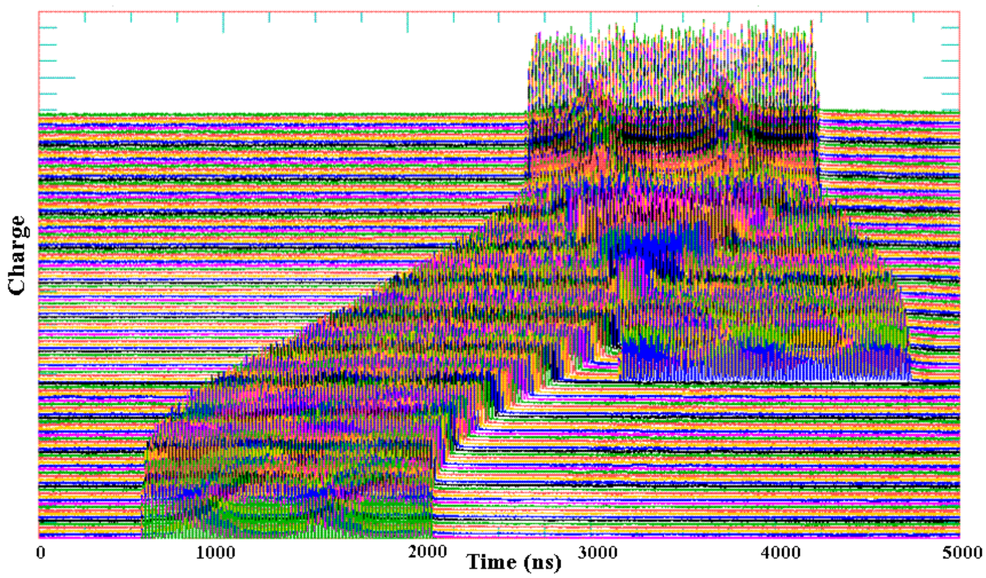

Figure 2.7: Data from the slip-stacking process. The two Booster batches are merged together in time (longitudinal position) over a period of $\sim$ 130 ms. Shown here is a series of scope traces during the process; the traces are separated by $\sim$ 1 ms.

Slip-Stacking

There are primarily two ways to increase the beam current in the MI: boxcar-stacking and slip-stacking. Boxcar-stacking is the simpler and more common of the two, practiced at many accelerator complexes and necessary for other types of stacking. In boxcar-stacking, a very fast kicker magnet is pulsed to lay beam from the Booster into the MI such that successive pulses of beam from are adjacent to each other longitudinally (azimuthally). Boxcar-stacking thus increases the total current in the ring while not increasing the instantaneous current – i.e. the number of particles in an RF bucket (measurements from this process in the MI are shown in Fig. 2.2.

Slip-stacking uses boxcar stacking to lay in two or more batches adjacent to each other, but then slightly accelerates the rear batch and decelerates the front batch. The two batches of beam pass through each other, until they are coincident; then, a much larger RF voltage is turned on to capture each pair of particle bunches into one larger bucket. Slip-stacking does increase the number of particles within a single RF bucket. In the process, the longitudinal emittance of each bunch is necessarily increased to more than the sum of the two precursor bunches. The procedure was initially proposed and demonstrated at CERN [55]. The slip-stacking in the Main Injector is the first operational use of slip-stacking [42]; measurements from the process are shown in Fig. 2.7. Currently, slip-stacking is used in the MI to enhance the proton current on the antiproton production target; expanded use of slip-stacking for the NuMI beam is under study.

53 MHz Acceleration

The primary RF system in the Main Injector operates at 53 MHz and can supply up to $\sim$ 4 MV of maximum voltage (V_0) per turn (revolution). The frequency actually changes from 52.8114 MHz to 53.103 MHz as the protons speed up somewhat; ferrite tuners are used to change the resonant frequency for the RF cavities.

A typical acceleration curve for NuMI operation is shown in Fig. 2.2. The cycle starts $\sim$ 15 ms before any beam is accelerated. The MI takes seven injections of Booster beam at a reduced RF voltage. The voltage must be at this lower level to properly match the beam longitudinally - otherwise the RF bucket would be too large in momentum acceptance and emittance dilution would occur. The first two batches are used for slip-stacking; these batches are bunch-rtoated in the Booster so that they have especially ong bunches, and reqire a low RF voltage. The fluctuation seen in the RF corresponds to beam-induced fields in the cavities that are damped by the beam-loading compensation system [56]. Acceleration consists of two parabolic regions and a linear region where the lattice magnets are ramped near the limit of the power supplies.

2.5 Antiproton Source & Recycler

Antiproton production and storage is a major component of Fermilab's collider program. While the NuMI beam does not use antiprotons, they must coexist and share some resources. Three rings are used for antiproton production: the Debuncher, Accumulator and Recycler. The Debuncher and Accumulator share a single tunnel and make up the antiproton source. The Recycler is a new storage ring for antiprotons.

Antiprotons are created by the interactions of 120 GeV protons with a solid target – similar to creation of the NuMI beam, only different secondary particles are collected. The protons are steered onto an Inconel target which produces secondaries. The secondaries are focused by a lithium lens, which is a mass of lithium with a current parallel to and centered on the beam axis. The current is distributed uniformly, leading to lens-like focusing similar to a horn.

The secondaries are directed into the Debuncher. The Debuncher is an accelerator with the same circumference as the Booster, which sets the length of the proton pulse at one Booster batch. The secondary beam has a very broad momentum spectrum, but a narrow width in time for each bunch as the beam is "rotated" in phase space before hitting the target. The Debuncher then applies RF that reduces the momentum spread of the beam at the expense of the bunching. After a few ms all unstable particles decay and only antiprotons reamin. The Debuncher's magnets are set to collect negative particles of 8.9 GeV/c momentum, which are all antiprotons. The Debuncher collects on the order of one antiproton for every 100,000 protons delivered to the target.

Once the beam is debunched it is transferred to the Accumulator, an 8 GeV storage ring. Multiple beam batches are collected into a "core". Each new batch is injected at a high dispersion region where it is transversely separated from the lower momentum core. The newly injected batch is then decelerated and merged with the core through stochastic cooling [57]. The maximum number of antiprotons that can be stored in the Accumulator is $\sim 2\times10^{12}$, which must be accumulated over $\sim$ 20 hours.

The Recycler is a new ring that accepts beam from the Accumulator. The Recycler is in the same tunnel as the Main Injector and is thus seven times the circumference of the Booster or Accumulator. The Recycler cools its beam through use of both stochastic cooling and electron cooling [58]. It is anticipated that the Recycler will be able to store 6×10^{12} antiprotons.

Antiprotons are extracted from the Accumulator or Recycler in 2.5 MHz coalesced buckets. The beam is then transferred to the MI where it is accelerated to 150 GeV and sent to the Tevatron.

The primary interaction of NuMI and the antiproton source is the combined slip-stacking/NuMI Main Injector cycle. Both beamlines have their beam accelerated together to 120 GeV; the slip-stacked beam is extracted to the antiproton source; then the rest of the beam is extracted to NuMI. The repetition rate of this process is set by the time required to stochastically cool each new batch of antiprotons injected into into the Accumulator core. The maximum rate the MI can deliver beam is about 0.5 Hz, but the antiproton cooling can require a time as long as 5 s; this time increases with the number of antiprotons in the core.

2.6 Timing

The timing of the Fermilab accelerator complex is controlled by TCLK – which sends outs signals at a maximum rate of 10 MHz around the site. Most devices that require only μs timing run off of TCLK and a timer, but each machine has its own additional timing process, and individual systems can require their dedicated timing.

The TCLK signals are decoded as hex bytes (e.g. \$1D), each with a particular meaning. Overall scheduling of clock events is set with a master clock and programmed with a Time-Line Generator (TLG). Events associated with accelerator operation are inserted via modules that are collections of events separated by specified time intervals; an example module is shown in Fig. 2.8. Timing of particular devices are typically set by at TCLK signal plus a timer good to μs. As the generation of signals is serial, during congestion a signal may be delayed by several 10 MHz ticks. TCLK is thus used only for events that do not require timing better than the μs level. Typically, TCLK events arm devices well ahead of their actual triggering, or signify the start of a slow event.

Additionally, each machine has its own timing system, usually focused on extraction; e.g. BES is the Booster Extraction Synchronization system. Whereas the TCLK provides a general-use, but "slow" system, these systems provide fast but proprietary communication. For examples of Fermilab timing systems see [59, 60].

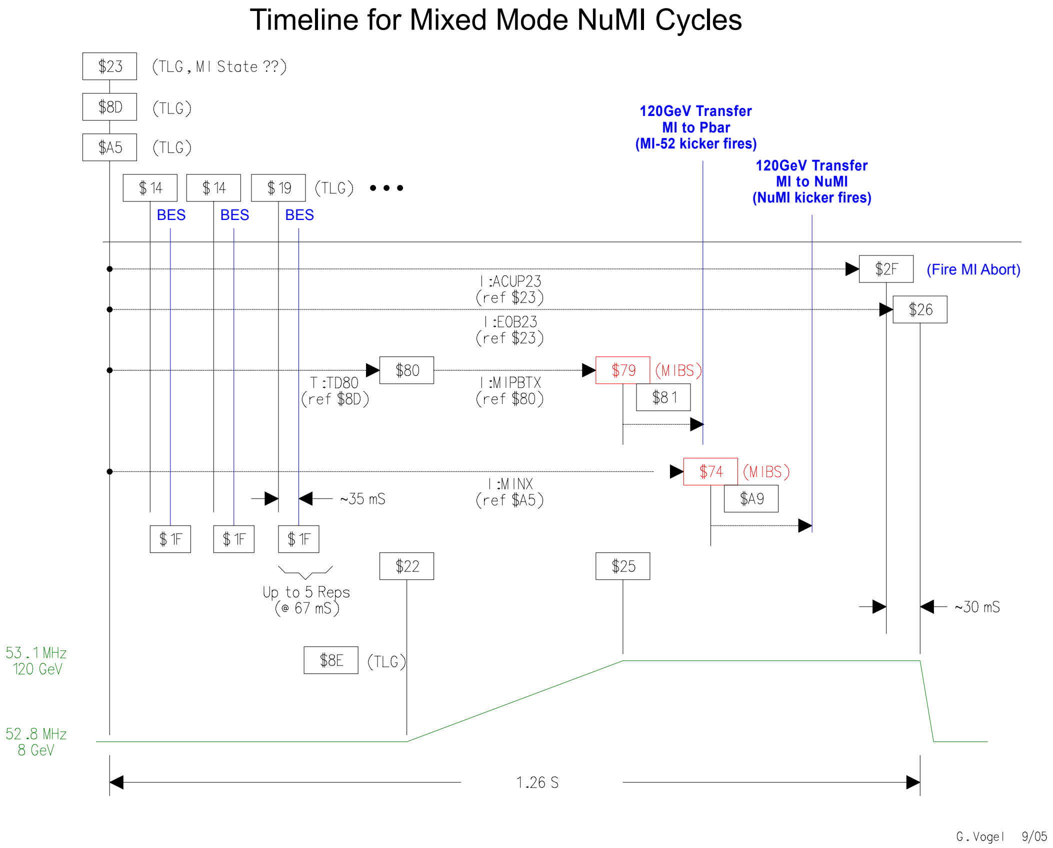

Figure 2.8: Schematic showing some of the timing structure of a Main Injector cycle that provides slip-stacked protons for antiprotons production and five Booster batches for NuMI. Booster events are \$14 for the start of a cycle to be slip-stacked, \$19 for the start of NuMI-destined cycle, and \$1F at the end of every Booster cycle; the diagram shows only one \$19, but typically there are four more. \$23 is the MI cycle that occurs previous to the Booster cycles, along with a number of other resets. The antiproton extraction is previous to the NuMI extraction by $\sim$ 1 ms, and there is another extraction to the MI dump at the very end in case there is any beam left in the machine. (Plot from Greg Vogel)

Chapter 3

The Fermilab Booster

The Fermilab Booster is an intermediate accelerator in the Fermilab accelerator chain. It has a strong focusing lattice made up of combined function gradient magnets. The magnets cycle at 15 Hz, setting the maximum rate; other components limit the accelerator's operational rate to more like 7 Hz. The Booster accepts 400 MeV H^- ions, converts them to protons, and accelerates them to 8 GeV for passage to the next accelerator or experiment.

The Fermilab neutrino program has called upon the Booster to accelerate far greater quantities of protons than in the past. Previously, the Fermilab fixed-target program was based on 200, 400, or 800 GeV protons from the Main Ring or Tevatron. The throughput of protons was limited by the Main Ring or Tevatron ramp time and the length of the spill time. The shorter, 1.5 second, ramp time of the Main Injector allows it to accept and accelerate almost an order of magnitude more protons than in the past. Most of the current program of Booster upgrades is focused on efforts to improve the rate of proton production; the limits on proton throughput will be discussed at the end of this section.

This chapter provides a review of the Booster accelerator and systems that are relevant for the forthcoming chapters:

- The magnet lattice, its properties, and how it is powered.
- Multiple-turn injection whereby a long-pulse of H^- ions is converted to a shorter pulse of protons.
- The Boosters RF acceleration system, its capabilities, and the process of adiabatic capture of the laminar beam into RF buckets.
- The low-level RF (LLRF) system that maintains acceleration stability, the associated dampers, and the systems used when accelerating through transition.
- Correction magnet systems used to fine-tune the orbit, including dipoles, quadrupoles, skew quadrupoles, and sextupoles.
- Beam instrumentation: position monitors, loss monitors, and the ionization profile monitor.
- The extraction system, including kickers, septa, the notch, phase-lock, and timing.
- The Booster performance in terms of proton throughput, and what limitations there have been.

3.1 Magnet Lattice

The Booster magnet lattice has 24-fold periodicity made up of 96 gradient magnets. The layout of the the Booster is shown in Fig. 3.1. The lattice also contains magnetic elements for injecting and extracting beam; these elements have distortive effects on the lattice that will be discussed below.

The cross-sections of the Booster magnets are shown in Fig. 3.2. The magnet design is taken from the Cornell synchrotron. Each magnet is composed chiefly of water-cooled copper coils and an iron yoke with shaped poles. The tilted pole shape give the magnetic field horizontal and vertical gradients, as described in §A.1. There is no beam pipe within the magnets; instead the entire magnet is within a vacuum vessel, such that the pole pieces define the available physical aperture [51].

The layout of a cell is shown in Fig. 3.3. Two F and Two D magnets are arranged in a modified FODO lattice (described in §A.1) – the FOFDOOD. It is so called as the F and D magnets are adjacent to each other, and there are two straight sections of unequal length – the longer earning the "OO". The long straight section allow major devices such as RF cavities, kicker magnets, and injection or extraction areas, and collimators. Instrumentation and correction packages (discussed below) are on both the short and long straight sections. In Booster terminology Lxx or Sxx refers to the long or short straight section of period xx.

The calculated lattice functions for the ideal Booster lattice is also shown in Fig. 3.3. The maximum β functions are 32 m in the horizontal and 21 m in the vertical; the maximum horizontal dispersion is 3.2 m, while vertical dispersion is always $\ll 1$ m. The beam is wide and dispersive in the short sections, while tall in the long sections. As the beam is larger vertically in the D-magnets they have larger vertical apertures and more current-carrying coils, as shown in Fig. 3.2. As discussed below, injection and extraction systems distort the envelope functions about the ring, and correction magnets are used to fine-tune the beam.

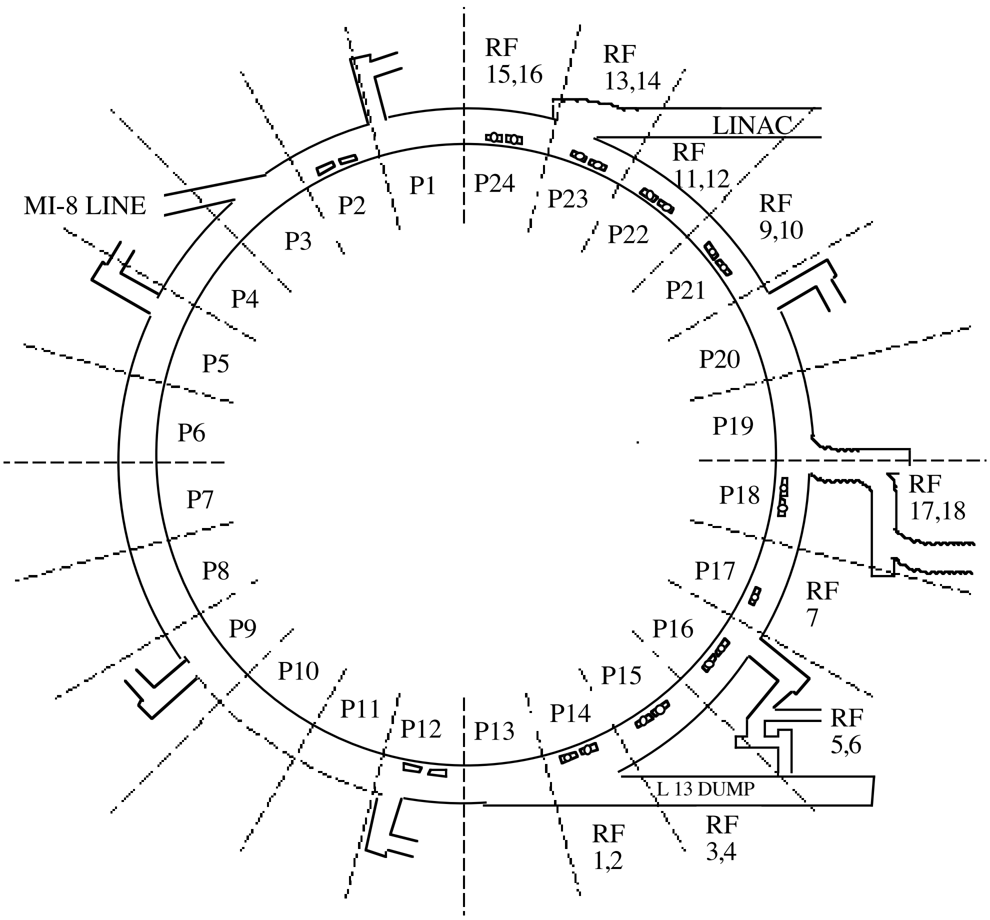

Figure 3.1: Footprint of the Booster accelerator. 400 MeV H^- ions enter from the LINAC. 8 GeV protons leave the Booster toward the L13 DUMP or the MI-8 LINE. The Booster is broken into 24 periods (Pxx), each of which consists of 4 gradient magnets, a 6 m long straight section, and a 1.2 m short straight section. Instrumentation, RF, and correction packages are contained in the straight sections, the RF cavities are shown as small rectangles on the right half of the ring. (from [52])

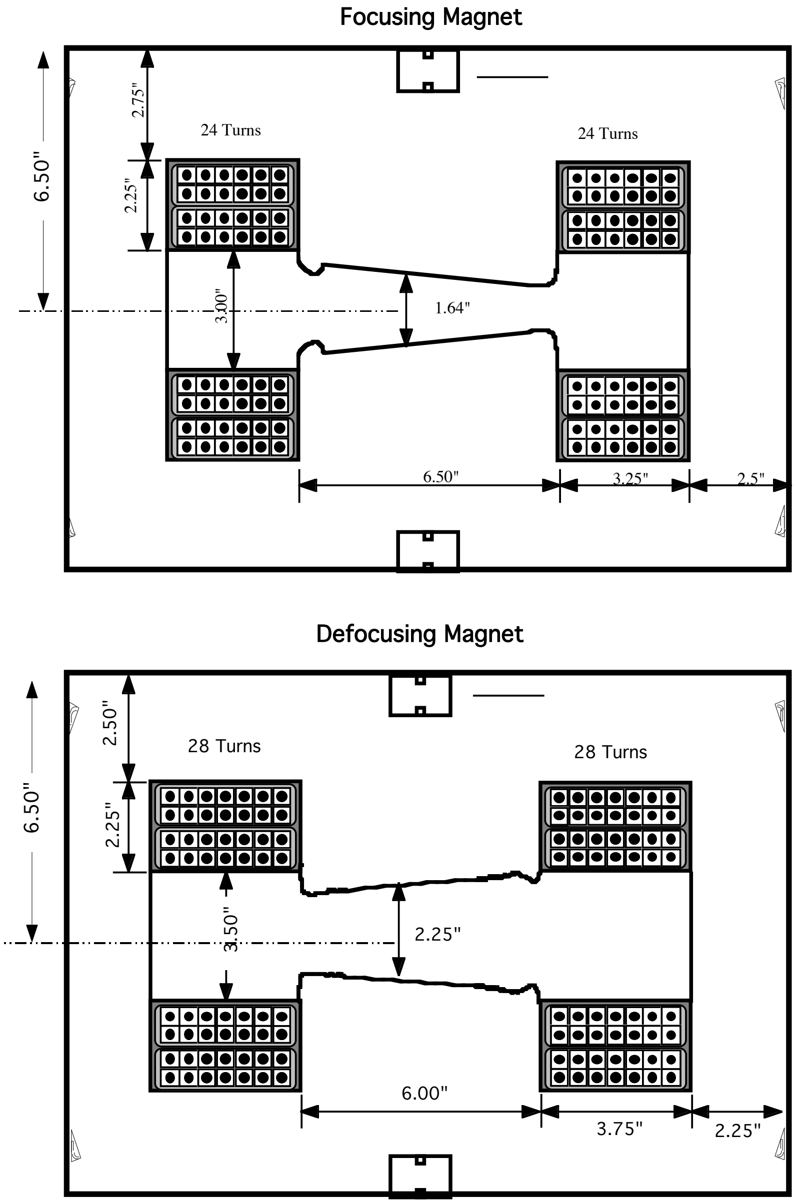

Figure 3.2: Cross-sections of the Booster combined function magnets. A central vacuum region is surrounded by an iron yoke and water-cooled copper coils. The slopes in the poles of the yokes determines the focusing of the magnet. D-magnets (defocusing) have a larger vertical aperture as the beam is larger when passing through them. (from [52])

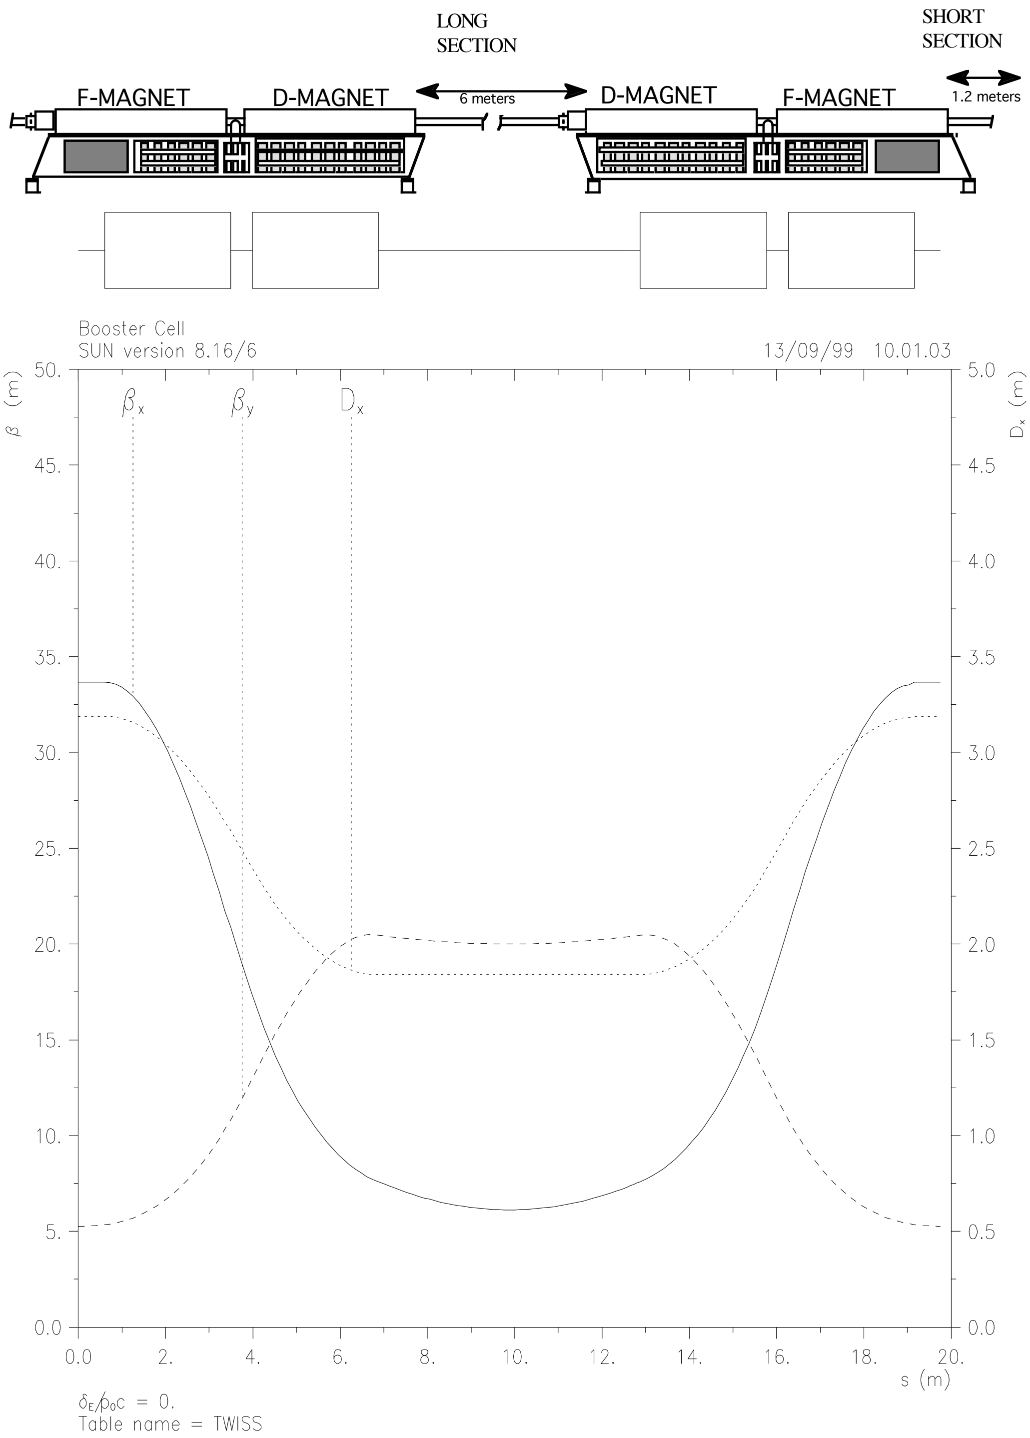

Figure 3.3: Calculated β and dispersion functions for a cell of the ideal Booster lattice. This cell is repeated 24 times around the circumference of the Booster.

The nominal tunes of the Booster are 6.6 horizontally, and 6.7 vertically; however, measurements indicate that the tune varies by ± .1 during acceleration and there is a calculated space-charge tune shift of as much as 0.6 at injection. The $\gamma_{\rm tr}$ of the lattice is 5.45, though early in the cycle the $\gamma_{\rm tr}$ might be a little different due to DC magnetic distortions.

3.1.1 Magnet Power System

The Booster's 96 primary bending and focusing magnets are powered by a system known as the Gradient Magnet Power Supply (GMPS). This system feeds AC and DC electrical current into the Booster magnet systems. The Booster magnets are accompanied by large capacitor banks and inductors mounted on each magnet's girder (visible in Fig. 3.3). The capacitors, inductors, and Booster magnets form an LC circuit that allows the AC component of the energy to be stored with a $Q \sim 20$, reducing power consumption.

The AC power is locked to the wallsocket power to reduce power disruptions. It had been found that if a power supply is not synchronized to the wallsocket frequency its maximum voltage will change with the relative phase. To maintain energy consistency GMPS was synchronized with the wallsocket frequency and varies with it[1]. The locking is such that the actual GMPS frequency is $1/4^{\rm th}$ of the wallsocket frequency[2] such that it is nominally 15 Hz, but changes by as much as 10 mHz.

The currents in the magnets follow a sinusoid with a DC offset. Beam Injection and extraction occur at the extremes of the sinusoid such that the acceleration is near zero at those times. The momentum varies directly with

[1] As we discuss below, this synchronization affects the Booster beam in other ways.

[2] the Booster magnets were designed so that they could also run at 7.5 or 30 Hz with suitable modification of the capacitors and inductor

the current such that:

$$p(t) = p_0 - p_1 \cos(2\pi f t) \tag{3.1}$$

Where $p_0 = (p_i + p_e)/2 = 4.9223$ GeV/c, $p_1 = (p_e - p_i)/2 = 3.9666$ GeV/c, $f = 15$Hz, and t is the time in the cycle; p_i and p_e are the injection and extraction momentums, respectively. The acceleration period is thus close to 33 ms, varying by where exactly on the sinusoid beam is injected and extracted – generally beam is injected shortly before the minimum current, and extracted shortly after the maximum current.

3.1.2 Injection & Extraction Chicanes

A series of four dipole magnets with equal bending power can be arranged such that the first and last bend in one direction, while the middle two bend in the opposite direction to produce an offset in position without distorting the rest of the orbit (e.g. see 3.4). Such a configuration of magnets is known as a chicane, double-dogleg, or 4-bump and is used for a variety of purposes in different accelerators. In the Booster, five sets of chicanes are used to create offsets in the orbit for injection and extraction. While a chicane's closed-orbit effects to the lattice are local, they have distortive second-order effects: the dipoles lead to vertical edge focusing that alters the β and dispersion functions around the entire ring.

ORBUMP is a system of four pulsed magnets used for injection, located in the Long-1 straight section (L1). Shown in Fig. 3.4, they are all pulsed dipoles designed to turn off in a few μs. The dipoles are used in the merging of protons and H^- beams and in deflecting them into the stripping foil. However, the magnets' edge focusing and slewing time cause distortions in the Booster orbit during and immediately after injection. Additionally, the magnets generate a great amount of heat; power dissipation limits their operation to a sustained rate of 7.5 Hz. In the near future the system will gain a new configuration [61] and new magnets [62] which will reduce the negative effects of the system and allow it to operate at 15 Hz.

There are two extraction regions located at L3 and L13, a schematic of which is shown in Fig. 3.12. Each region contains two chicanes: BEXBumps and the Dogleg system. The BEXBumps are pulsed dipoles that turn off at the end of acceleration and keep the beam away from the extraction septum until extraction is desired. The Dogleg is a DC chicane that moves beam away from the septum: the deflection away from the septum becomes smaller as the beam accelerates and becomes stiffer. The Dogleg allows the septum to actually reach into where the beam would otherwise circulate, simplifying extraction. As the beam accelerates its width decreases due to adiabatic damping, §A.3, allowing it to become closer to the septum without intersecting it. The Doglegs were found to cause severe distortions to the Booster's lattice at injection, increasing dispersion by as much as a factor of two in some locations [63]. Efforts have been underway in the Booster to minimize the negative effects of the Doglegs.

3.2 Multiple-Turn Injection

Multiple-turn injection is a tool used at most modern proton synchrotron facilities to inject additional charge into already occupied phase space. Typically, the proton density in a portion of phase-space cannot be increased – more beam cannot be injected into the beamline without deflecting some beam. Multiple-turn injection of protons involves bending [64] H^- ions to be colinear and coincident with circulating protons, and then stripped of their electrons [64].

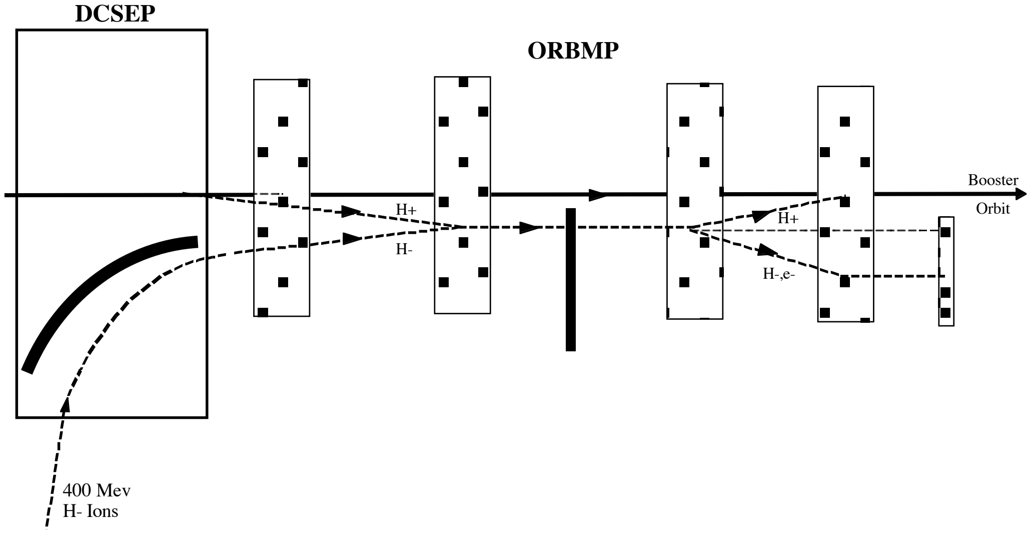

Figure 3.4: Injection diagram of the Booster. In a long straight section there is a septum magnet (DCSEP), a set of four pulsed dipoles (ORBMP), and a thin carbon foil (vertical black line). The 400 MeV H^- ions are bent into the beamline in the field region of the septum magnet – the circulating beam is in a field-free region. The dipoles of the ORBMP system bring together the ions and circulating protons into the same phase space (i.e. position and angle). The ions are stripped in the foil such that only protons and electrons exit. After several circumferences the H^- beam stops and the ORBMP magnets are turned off quickly so that the circulating beam does not intersect the foil. (from [52])

The current scheme for multiple-turn injection is shown in Fig. 3.4. The proton and H^- beams are brought together by a chicane of pulsed dipoles, then they pass through a stripping foil [65]. After the H^- injection the dipoles are rapidly turned off, taking the proton beam out of the stripping foil. Currently, the Booster uses a graphite stripping foil of $\sim 550\ \mu\text{g/cm}^2$ thickness, and has typically been in the range of 400-600 $\mu\text{g/cm}^2$. Stripping foils can be switched remotely using a machine with multiple foils.

The Booster can inject up to the full length of the 60 μs Linac beam, but usually takes much less. Depending on Linac current, the Booster takes between 12 and 15 turns to fill completely. More charge is counter-productive because either the excess beam is lost during capture or at transition. In both cases the efficiency is very low and beam power lost in the Booster is excessive – limiting the total throughput. In tests, as many as 8×10^{12} protons were injected into the Booster, but only 6.4×10^{12} survived to extraction. More typically, 5.5-6×10^{12} are injected and 4.5-5×10^{12} are extracted.

3.3 RF Acceleration

RF acceleration in the Booster is provided by 19 RF stations, each providing $\sim$ 50 kV of voltage. The Booster can accelerate beam reliably with as few as 16 stations, but efficiency suffers. The RF stations are individually tuned to be in phase with the beam, and are then divided into two groups: A & B; a balancing system tries to keep an equal amount of voltage on each group.

The RF power ($\sim V_0$) at each point of the cycle is programmable, and can be tuned for best operation. An example RF curve is shown in Fig. 3.5. The total RF is kept at is maximum until the end of the cycle where not as much is needed. The synchronous phase is slightly larger at those times, but it also stabilizes the beam by allowing a slightly longer bunch. Fig. 3.5 also shows that the synchronous phase grows to as much as 50° at maximum acceleration; with more beam this angle can become even larger and the RF bucket shrinks, leading to beam loss. The synchrotron frequency can be as much as 40 kHz early, but settles to a value of $\sim$ 2.5 kHz after transition.

The frequency of the RF also changes substantially during acceleration from 37.94 MHz to 52.8114 MHz. The resonant frequency of each cavity and the RF frequency of the driver must be adjusted in a predetermined way. The cavities are tuned via ferrite tuners: a magnet field induced by a large

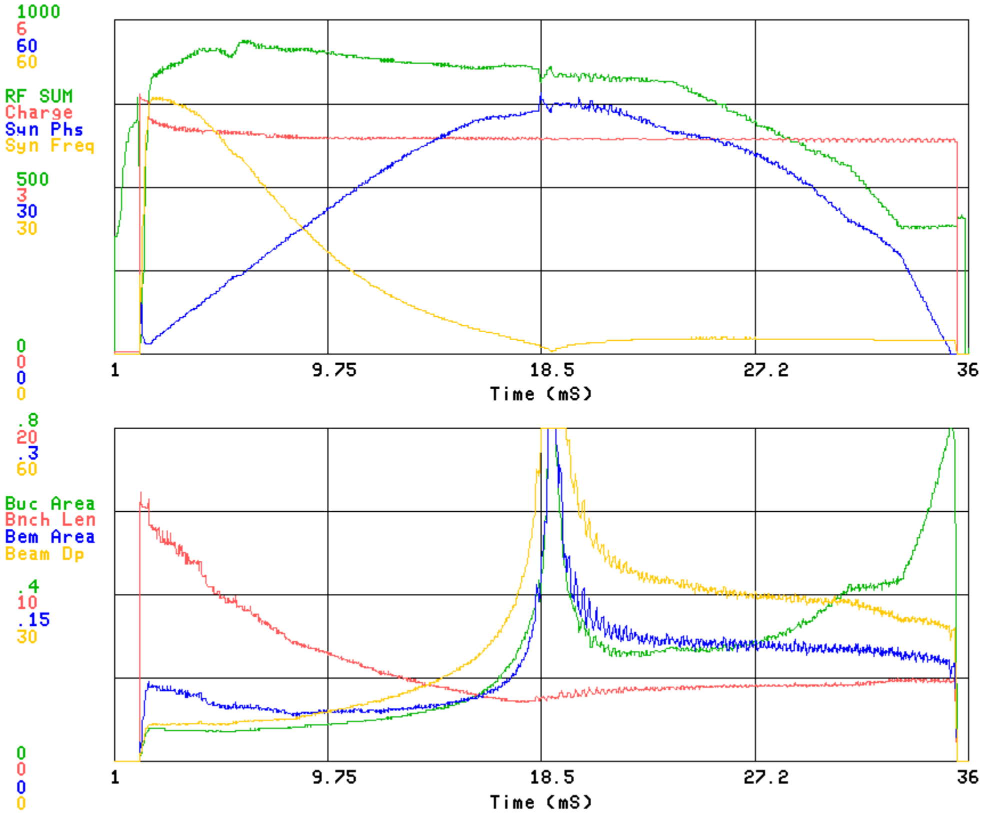

Figure 3.5: Parameters of the Booster beam during the acceleration cycle; quantities plotted are: the measured maximum RF voltage per turn (V_0) in kV, the measured charge in the machine in units of 10^{12} protons, the measured synchrotron phase in degrees, the calculated synchrotron frequency in kHz, the calculated bucket area in eV·s, the measured bunch length in ns, the beam area in eV·s, and the momentum width of the beam (δp) in MeV.

electric current changes how near the ferrite is to saturation, and thus the permeability of the ferrite. The phase velocity of RF waves through the ferrites is thus altered, changing the resonant frequency of the cavity. The Q of each cavity is about 400 including frequency-dependent losses to the ferrite, so the RF drivers must continually provide power. The phase of the RF system is controlled by the low-level RF system (LLRF), discussed below in §3.4.

3.3.1 RF Capture

The beam injected into the Booster is partially debunched 200 MHz beam, not the initial 38 MHz of the RF system; so the beam must be captured into RF buckets. This is accomplished by ramping up the RF power applied to the beam over a few hundred μs. As the power increases the synchrotron frequency increases, the RF bucket area increases, and beam is swept into the buckets with an efficiency of 96-98%. The initial RF frequency of the beam is thus imposed upon it and has no relation to the beam's energy.

The RF drivers and Q of the cavities prevent the RF from being ramped that quickly, so a procedure known as paraphrasing is used to rapidly increase the effective V_0. Paraphrasing uses the two circuits A & B mentioned earlier. They each make up one half of the RF power and there phases are individually tuned such that they are 180° apart. The phase can be adjusted rapidly enough allowing V_0 to be adjusted more quickly.

3.4 RF Feedback

The horizontal aperture of the Booster, ~ 2.5", is much smaller than its radius, so the beam energy must be controlled precisely during acceleration to keep it in that aperture. This control is provided by the low-level RF (LLRF) system. The LLRF system has many tasks, but its central goal is providing an RF phase at which to stably accelerate the beam (in contrast to the "high-level" RF that provides the power to the beam)..

A diagram of the LLRF system is shown in Fig. 3.6. The LLRF has a number of preset "curves" that describe standard settings for different times of the cycle, including RF frequency and synchronous phase. These are just used as a baseline, however, the actual values are found dynamically. As the beam accelerates its position is read back through BPMs, as described below in §3.6, one of which is part of the radial loop: it measures the radial[3] position (RPOS) of the beam and uses it for RF feedback[4].

The measured RPOS is compared to a predetermined radial offset (ROFF) curve. The ROFF curve presents the desired position of the beam at each moment of the cycle; this curve is tuned to maximize efficiency of the accelerator. The measured RPOS, following the ROFF curve, for several Booster cycles is shown in Fig. 3.7. In a synchrotron, radial position acts as a proxy for momentum: a change in the horizontal centroid of the circulating beam represents a deviation of the momentum from the first-order approximation in Eqn. (3.1); the position deviation is proportional to the dispersion and momentum deviation (see §A.1. Thus, by changing the rate of acceleration the momentum deviation and radial position of the beam can be changed. The LLRF does this by adjusting the phase of the RF that must be applied to beam. The RF phase is thus the actual synchronous phase plus some amount to push the beam toward its preferred radial position. As the synchronous phase is

[3]This is actually just a horizontal position, but is customarily known as radial in the LLRF context as it concerns the circumference of the machine.

[4]The particular BPM used as RPOS has changed over time. It is always a BPM in one of the long sections, and has most recently been L18.

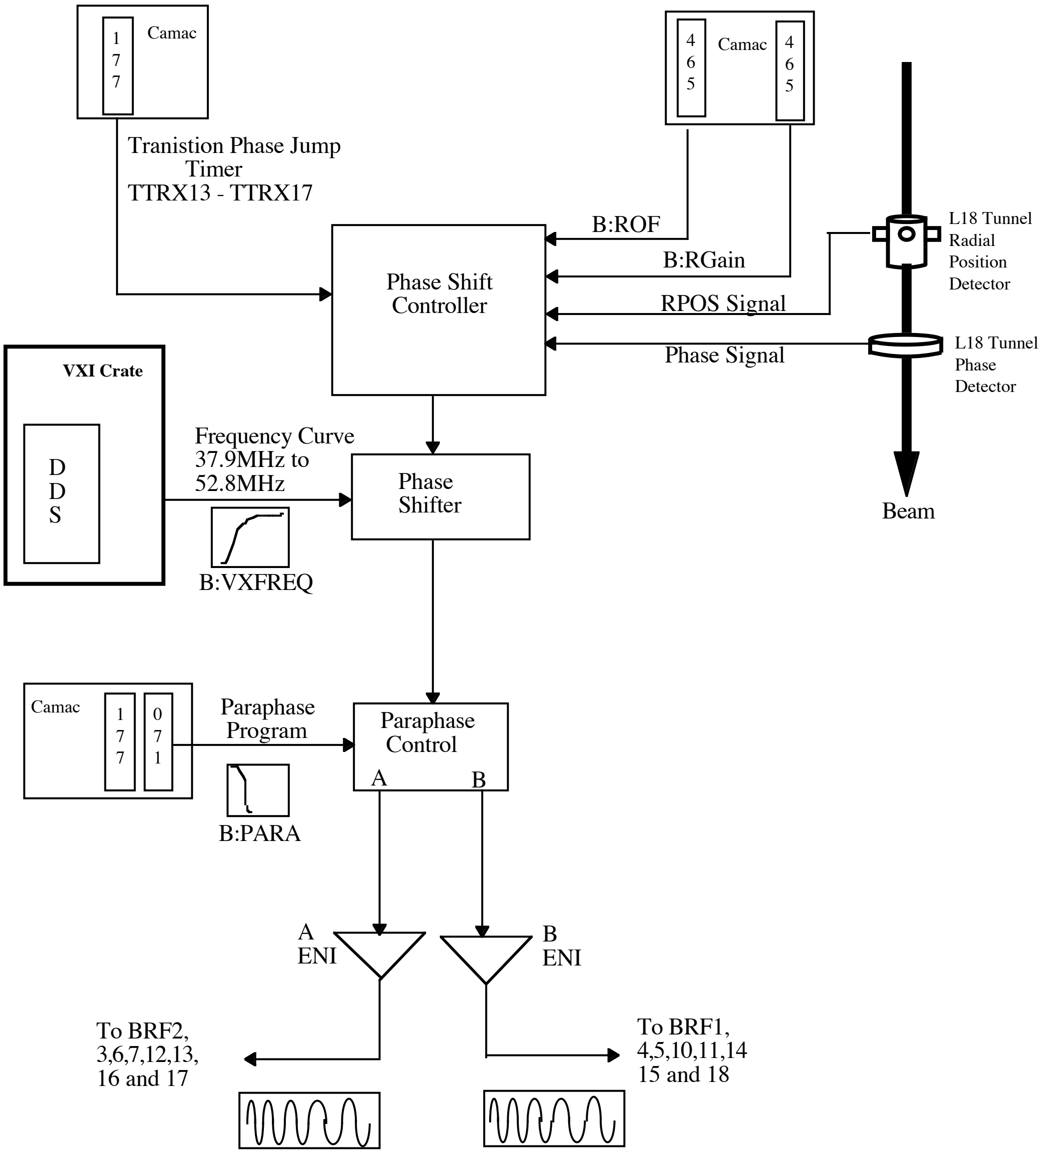

Figure 3.6: Simplified Booster LLRF diagram. The desired radial offset (ROF), radial feedback gain (RGain), frequency curve, and paraphase program are predetermined curves, tuned for efficiency. The radial position (RPOS) and beam phase are measured from the beam and used to adjust the RF applied to the beam. The paraphase splits the RF power into two circuits of RF cavities that can be separately phased. (from [52])

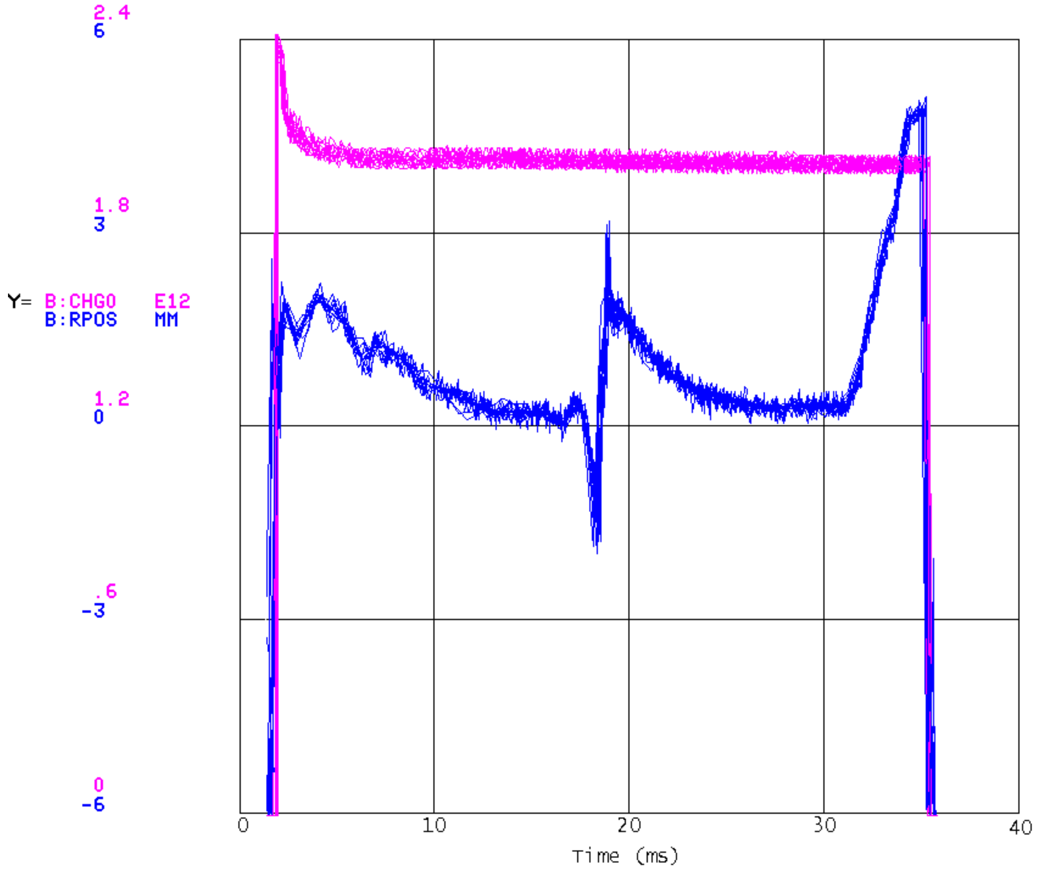

Figure 3.7: Plot of measured radial position (RPOS) throughout the Booster cycle, for multiple batches of Booster beam. The Booster LLRF maintains RPOS at a preprogrammed trajectory. In absence of synchronization with the Main Injector, all cycles follow closely the same trajectory, as shown here.

not exactly predetermined, the LLRF uses a PID loop to approach the proper phase. By changing the phase consistently the LLRF also effectively changes the RF frequency of the beam through out the cycle.

3.4.1 Transition Crossing

The LLRF must be aware of transition crossing as it effects where on the RF waveform the synchronous phase is, and also which direction the phase should be adjusted to get the desired change in momentum. The precise time of transition cannot be measured well dynamically, and is instead set by a tunable timer. When that timer is reached the RF system must change the RF phase as rapidly as possible to reach the other side of the waveform – a jump of $180° - 2\phi_S$. This jump induces oscillations into the beam and dilutes the longitudinal emittance.

To improve the transition crossing, the beam is typically forced to a rapid change in radial position at that time (visible in Fig. 3.7). This is a result of second-order focusing effects in the lattice known as chromaticity: different radial positions receive a different amount of focusing around the ring and thus have different tunes and γ_{tr}. Changing the radial position takes advantage of this and minimizes the time the beam spends near transition. Another strategy is to reduce V_0 through transition so the synchronous phase is larger and the phase jump smaller. A third method for dealing with transition: a gamma-t system (currently unused, but has been implemented at other accelerators). This is a system of pulsed quadrupoles that rapidly changes the focusing around the ring, and thus the γ_{tr} – similar to, but more substantial than the radial swing described above.

3.4.2 RF Dampers

Another feedback system to the RF is composed of dampers that deal with oscillations of the beam that are not of the natural synchrotron or RF frequencies. Typically, these are harmonics of higher-order modes (HOMs) in the RF cavities or rational sums of the revolution and betatron or synchrotron frequencies. Only limited feedback can be applied directly through the main RF system due to the limited bandwidth of the RF cavities. The Booster has several dampers for specific modes of oscillation, but continues to have issues with other modes. One such mode is a quadrupole mode associated with the transition jump that is essentially proton bunches fitting poorly into their RF buckets and tumbling. The oscillation is at twice the synchrotron frequency and eventually damps out through nonlinear oscillations – increasing the longitudinal emittance.

3.5 Orbit Control

The Booster orbit varies throughout the cycle due to misalignments and varying strengths of the magnets. An example of such motion is shown in Fig. 3.8. Whereas the radial motion discussed in the previous section is correlated around the entire ring and controlled by the RF (see Fig. 3.7), this motion varies from location to location and must have other correction methods. Currently, the improvements in the lattice are gained by regular re-alignment of the Booster, tweaking of the chicanes discussed above and the use of corrector magnets.

3.5.1 Correction Magnets

Correction magnets are auxiliary magnets, smaller and less powerful than the main bending magnets. Each sector of the Booster contains a pair of correction packages. The packages are a set of four air-core overlaid magnets: vertical and horizontal dipoles, a quadrupole, and a skew quadrupole. These allow fine tuning of the beam about the ring. The dipoles are used for moving the centroid of the beam, the quadrupoles for adjusting the focusing, and the skew quadrupoles for the coupling between planes. The quadrupoles were important

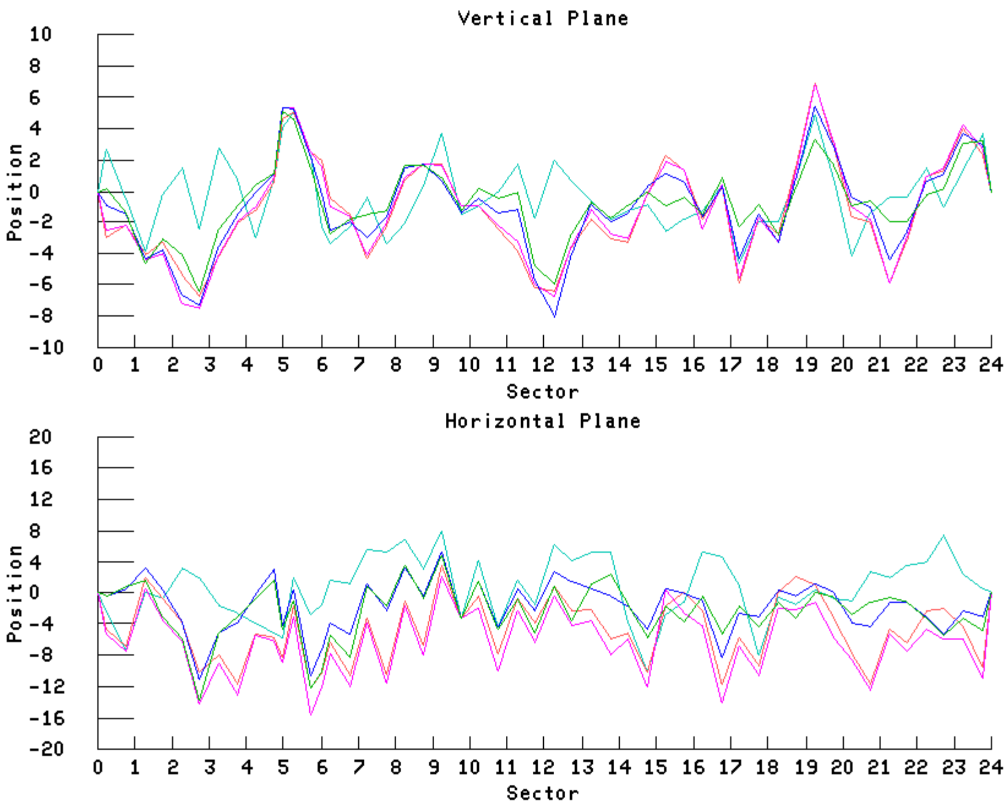

Figure 3.8: Example of beam motion during the Booster cycle. Plotted is the measured position of the beam at 48 locations (two for each period) around the circumference of the Booster. The five traces are from different times in the cycle; approximately: 2, 10, 16, 24, and 33 ms. Note that the change in positions at different locations are not always correlated.

in mitigating the distortive effects of the doglegs before modification (see §3.1). The corrector magnets were recently upgraded with a programmable control system that allow them to have different currents throughout the cycle. This system has had only limited use, however, as the magnets are not powerful enough to substantially affect the beam later in the cycle (the orbits in fig. 3.8 are with correction from the corrector packages) and cannot slew current quickly enough. A new corrector system consisting of new magnets, power supplies, and controllers is planned [66].

Another set of auxiliary magnets are sextupoles. Twenty-four sextupoles about the ring are split into two circuits: the SEXTL and SEXTR, depending on whether they are located in the long or short regions, respectively. The sextupoles substantially affect the chromaticities which are important for transition-crossing, the operation of the LLRF, and stability of the beam. Octupole magnets were also installed in the Booster, but were not substantially used and were later removed to install other components.

3.6 Instrumentation

The Booster has a large number of instruments for measuring the properties of the accelerating beam. Below, we discuss in some detail the Booster Beam Position Monitors (BPMs), Beam Loss Monitors (BLMs) and Ionization Profile Monitor (IPM). Other instrumentation that is not discussed below: toroids for beam intensity, resistive wall monitor for longitudinal measurements, total loss monitors, velocity monitor, phase monitor, and flying beam.

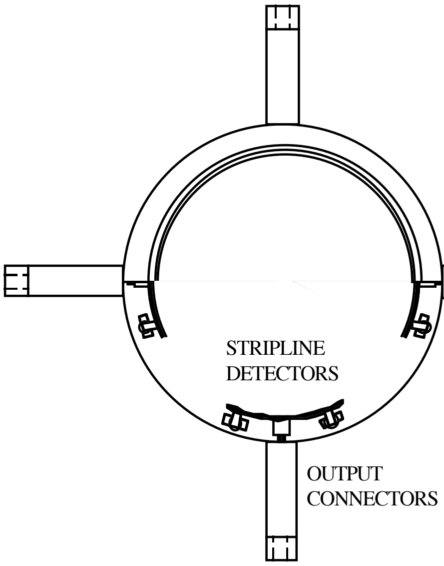

Figure 3.9: Schematic of a Booster BPM, split to show the beam pipe on the top half. Each BPM has four stripline electrodes, opposite pairs of which are used to make horizontal or vertical position measurements. (from [52])

3.6.1 Beam Position Monitors

As alluded to above, the Booster Beam Position Monitors (BPMs) are used to steer the beam and to control energy feedback. The Booster has 48 BPMs around the ring – one at each short and long straight section; more BPMs are present in the transfer lines. Each BPM provides vertical and horizontal centroid measurements, but provides no information on beam width or other moments.

A schematic of a Booster BPM is shown in Fig. 3.9. The BPM is of the stripline type which couples capacitively and inductively to the beam. Four striplines make up each BPM and are matched to be at about the same radius from the center as the walls of the beam pipe to reduce impedance-matching issues. When the beam is in the central region the beam induces a voltage on the striplines – generally inducing more on a stripline it is closer too.

The BPMs must operate at all the different frequencies of the Booster, and thus have no natural resonant frequency. The signals on each stripline are brought out and measured in pairs – the top and bottom for vertical measurement, and the left and right for horizontal measurement. The standard measurement is to divide the difference of measurements by the sum; the resulting number (between -1 and 1) is to first order where the beam is in units of the stripline spacing. Some corrections for nonlinearity are applied at larger distances.

The Booster electronics digitize and store position information for the entire cycle, but the network backbone cannot handle this data in a sustained manner. Typically, The BPMs are all read out ~ 25 times during a specific cycle. The orbits can be displayed individually, Fig. 3.8 is such a display with five time snapshots chosen. The one BPM outfitted as RPOS is measured out continuously into the low-level and a ADC that allows it to be plotted over time for every cycle.

3.6.2 Beam Loss Monitors

Beam Loss Monitors (BLMs) are sealed ionization chambers that respond to the passage of high-energy particles by producing a current – the current increases with increased charged particle flux. The BLMs are placed outside of the accelerator, usually on a magnet or beam pipe and measure only particles and showers associated with beam loss.

The Booster has ~ 66 loss monitors; two per period, six for each of two extraction regions, and six extra for the collimation region. There are more BLMs on the injection and extraction lines, but they are not part of the Booster BLM system. The signal from each BLM is directed into an integrator where the charge is held on an RC circuit. The integrated charge for each BLM, for

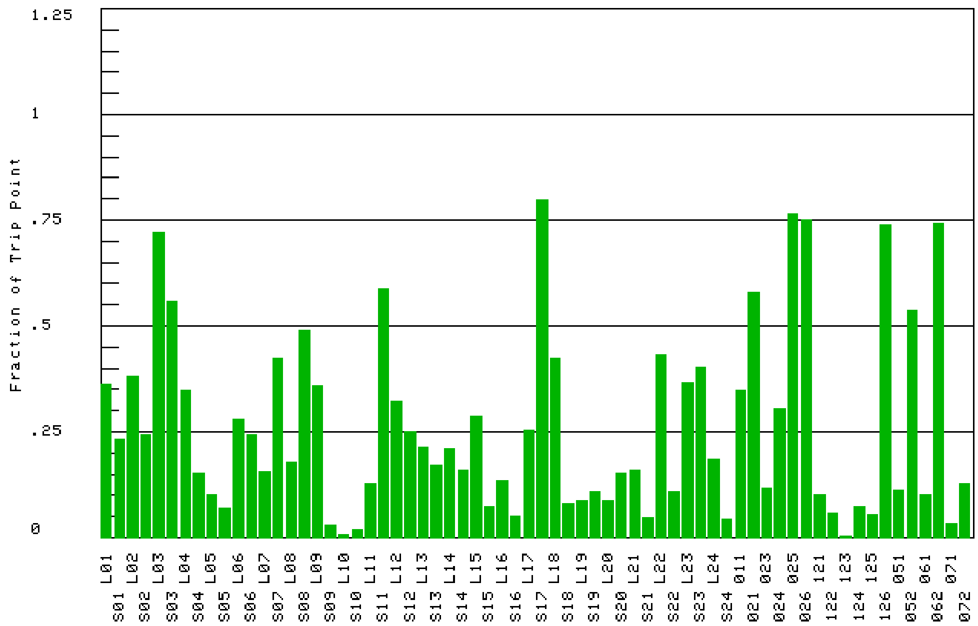

Figure 3.10: Measured beam loss around the Booster. There are two BLMs per period and additional ones at the extraction regions and collimators. Plotted is the sum of losses from all cycles in a 100 s period; each loss is normalized to the administratively set loss limit for that location.

each cycle can be plotted as a function of time. This is convenient for locating beam loss and understanding the beam loss mechanism by when it happens in the cycle and how quickly.

The BLMs in the Booster are used not only to locate beam loss, but also to limit Booster operation. The charge of every BLM at the end of the cycle is recorded as a measure of the total beam loss in that area. A running sum of losses is maintained for the last 100 s. Each BPM has a limit associated with its running sum. The limit is an administratively set level arrived at by correlating the measured beam loss during cycles with the measured residual activation at that site, and deciding upon what is the maximum residual radiation level allowable. If the beam loss at a particular location exceeds its limit the beam is temporarily disabled until the running sum falls below the limit [67]. Frequent bean inhibitions will cause operators to adjust the tuning of the Booster at that location or to reduce the rate at which the Booster is operating.

3.6.3 Ionization Profile Monitor

The ionization profile monitor (IPM) is a device designed to measure the profile of the proton beam while it is being accelerated. It is the only device that provides a measurement of beam width during acceleration. There is one IPM composed of two parts for vertical and horizontal measurements.

The IPM depends on ionization of the residual gas in the vacuum vessel; Booster vacuum is on the order of 10^{-7} torr. The ions then drift in an applied electric field to a micro-channel plate (MCP) where the charge is amplified and read out of the channels individually to create a profile. The profile is corrected for the effects of the beam's electromagnetic field and a beam centroid and width is produced [68]. The beam with can also be divided by the square root of the boost to produce an emittance.

Fig. 3.11 shows some typical beam width measurements from the IPM. The beam widths generally shrink during the cycle due to adiabatic damping. The vertical beam width grows rapidly during injection due to lattice issues at low field. The horizontal beam width has contributions from both horizontal emittance and dispersion; near transition the $\delta p/p$ grows and so does the beam width temporarily. Also, fast oscillations in the beam can also be measured as beam width, because the IPM's resolution is only a few μs. After transition the transverse emittances are noticeably larger; there is some longitudinal emittance dilution and oscillations at transition that leaks to the horizontal plane through synchro-beta coupling, and then through transverse coupling to the vertical plane.

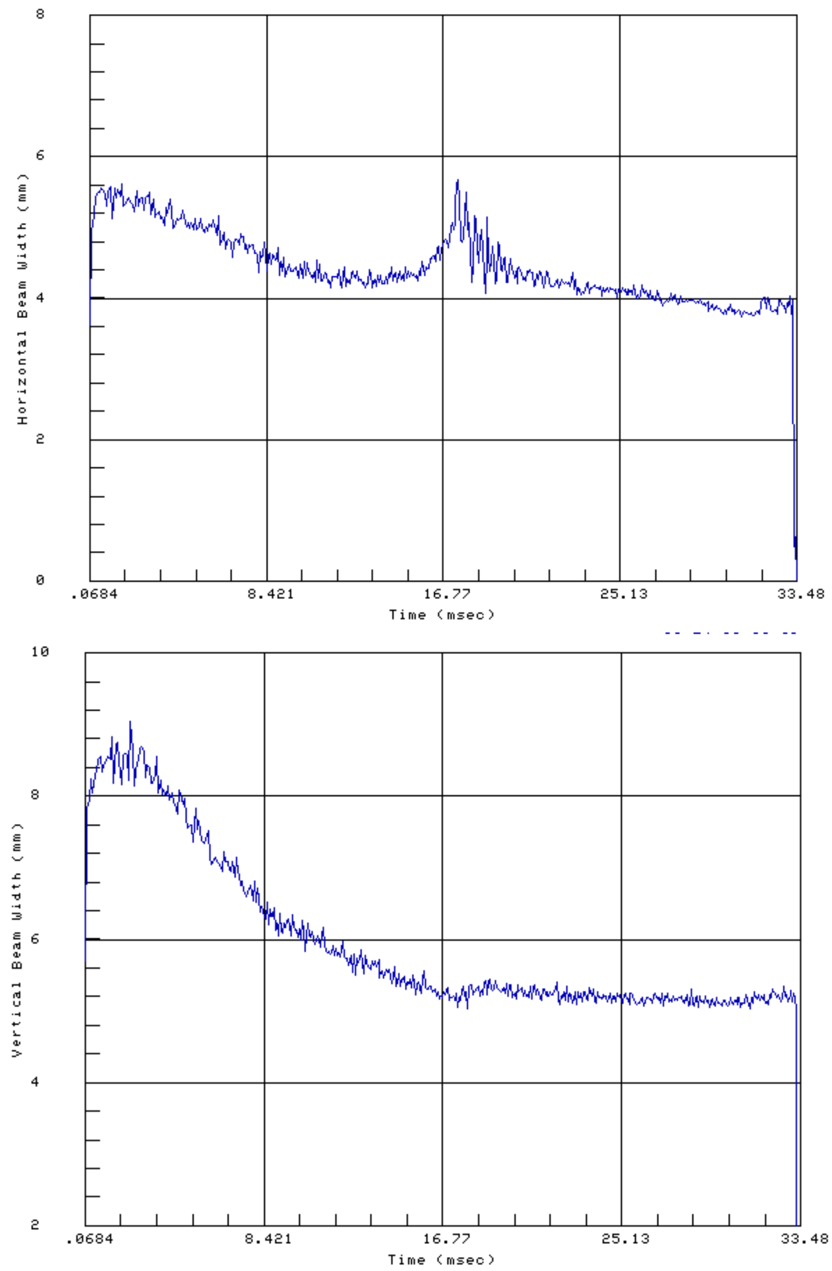

Figure 3.11: Beam width measured by the IPM during the Booster cycle. The horizontal beam width grows near transition, because the $\delta p/p$ is large and dispersion increases the beam width.

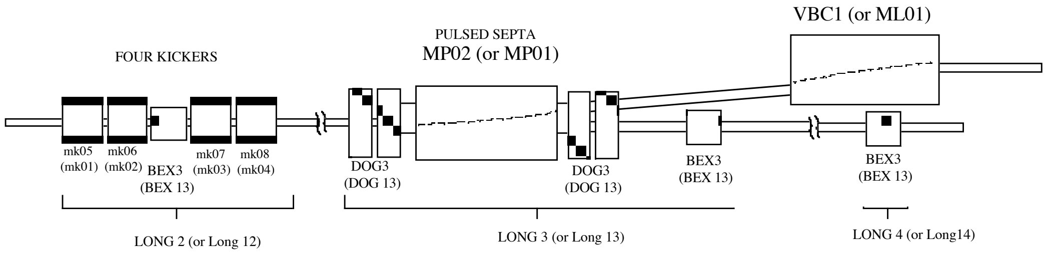

Figure 3.12: Layout of the Booster extraction regions (both are similar). Kicker magnets in the previous period give a vertical kick to the beam that becomes a vertical offset at the septum magnet. A set of four DC dipoles make up a double-dogleg about the septum. The septum magnet is pulsed and bends the beam vertically into the extraction channel. (from [52])

3.7 Extraction

The Booster has two nearly identical extraction regions, at L3 and L13, As shown in Fig. 3.1, the L3 extraction heads towards the Main Injector and MiniBooNE, while L13 heads toward the Booster dump.

The Booster extraction components are shown in Fig. 3.12. As discussed above, there are two four-dipole chicanes that steer the beam in the extraction region. The first is called the DC Dogleg, which steers the beam away from the extraction septum, but less so at high momentum than low. The second, the pulsed BEXBump system, holds the beam away from the septum until immediately before extraction. A set of four kicker magnets in the previous long section induce a large betatron oscillation into the beam that translates into a vertical offset at the septum magnet of ~ 20 mm, though that number is tuned substantially[5]. The vertical offset moves it to a position on a septum magnet such that it passes through an area of magnetic field that bends the beam upward, away from the Booster, and into the extraction channel.

[5]With a phase advance of about 97° the beam is actually slightly angled toward the center of the beamline – away from extraction. However, the vertical bend induced by the septum is much greater than the oscillation.

3.7.1 Kickers

The kicker are fast pulsed vertical magnets in the straight sections. The risetime of the pulse is $\sim$ 30 ns. To achieve that fast risetime the magnet must have only 1 conductor (i.e. have low inductance), be loaded with ferrite to sharpen the pulse, have short length ($\sim$ 1.2 m) to minimize filling time, and be powered by a high-voltage ($\sim$ 50 kV) pulse-forming network (PFN). Extraction is achieved by the pulsing of four kickers, so there are a total of eight in the Booster tunnel. If one kicker fails, a single magnet of the other set can be used to compensate – the betatron oscillation of a single kicker is not enough to cause the beam to fall outside of the available vertical aperture at extraction.

3.7.2 Septum Magnets

A picture of a Booster septum magnet is shown in Fig. 3.13. The bottom, larger aperture is where the beam circulates during acceleration. The top, smaller aperture is where the beam is placed by the kickers. Between them is the septum that is a copper plate that carries current; with current return along the top a horizontally-oriented magnetic field is produced above the septum that will bend the beam vertically. The bottom aperture is known as the field-free region, but the magnetic field is not zero so the magnet is pulsed over $\sim$ 100 μs so that the circulating beam is not affected by it. If beam only received a partial kick it will intersect the septum and be lost. Additionally, any large amplitude beam that could revolve in the large aperture of the Booster will be collimated by the small extraction channel, resulting in beam loss at extraction.

Figure 3.13: View inside a Booster extraction septum magnet. The circulating beam travels through the large horizontal region where there is no magnetic field. The extracted beam exits through the smaller aperture above the copper plate where the magnetic field is pulsed.

3.7.3 Phase Lock

The process of phase lock makes possible synchronous RF transfer between the Booster and Main Injector. The beam in the Booster is bunched into 84 53 MHz buckets and will be accepted by the Main Injector into its RF buckets. If the charge from the Booster enters the Main Injector at the wrong phase with respect to its RF buckets, then the beam will oscillate longitudinally, eventually being damped out and resulting in emittance dilution. Additionally, if the frequency is incorrect then only a few bunches can have the correct phase upon entering the Main Injector. The phase lock system matches the Booster RF phase to the Main Injector RF; in the process, the frequencies are also matched as a result of the phase-matching process. The original Booster phase lock system, still in use today, was designed for transfers to the Main Ring [69], but transfers to the Main Injector are governed by the same principles.

Phase lock begins about 2 ms before Booster extraction. Its precise trigger is based on frequency: when the Booster's RF frequency is $\sim$ 7 kHz below the reference frequency sent by the MI. Then, the Booster LLRF uses a phase comparator between the MI and Booster RF frequencies, resulting in an error signal. Phase lock then sends a signal to the LLRF indicating a change in desired radial position of the beam (RPOS). This changes the revolution frequency of the beam and is manipulated until both the frequency and phase of the beam are matched that of the MI.

The frequency of the Booster beam is then held fixed as the magnets reach their maximum value. The Booster's gradient magnets vary sinusoidally, so there is no time of constant magnetic field and beam energy. Instead, the $\pm$ 1 ms around the maximum is used as a pseudo-flattop for extraction. The frequency is fixed by the phase lock and LLRF and the beams radial position is allowed to vary to maintain the frequency in the slightly varying magnetic field. The frequency is held constant not the energy, so a variation in maximum magnetic field leads to a variation in the extracted beam energy. This variation persists into the Main Injector as the MI will only maintain the frequency until acceleration – i.e. it performs no radial feedback until the acceleration begins.

3.7.4 Notch

Some of the Booster's beam must be removed to allow clean extraction at full energy – the area of removed beam is known as the notch. As mentioned above, the risetime of the extraction kickers is ~ 30 ns while the bunch spacing is only 19 ns; additionally, any beam receiving a partial kick will intersect the septum plate of the extraction magnet. Most of the beam intersecting the septum plate will be lost and shower in that area, leading to residual activation. In larger machines with higher beam energy a similar technique is used to extract the beam quickly. In those machines it is typically known as an abort gap and is created by not putting any beam into that location in the first place. The distinction between the two might be that the notch is shorter and is created by removing beam that had been injected.

The Booster ran for many years losing 1-2 bunches of 8 GeV protons on its septa during extraction. The radioactivation was higher in those areas than the rest of the ring. The prospect for losses became much greater with the increase in the number of protons used for antiproton production and for the NuMI and MiniBooNE beamlines. All-in-all, the increase in extracted protons was to be about a factor of ten. The resulting radiation would have resulted in damage to accelerator components, and the residual activation would have limited the time that workers could approach the components.

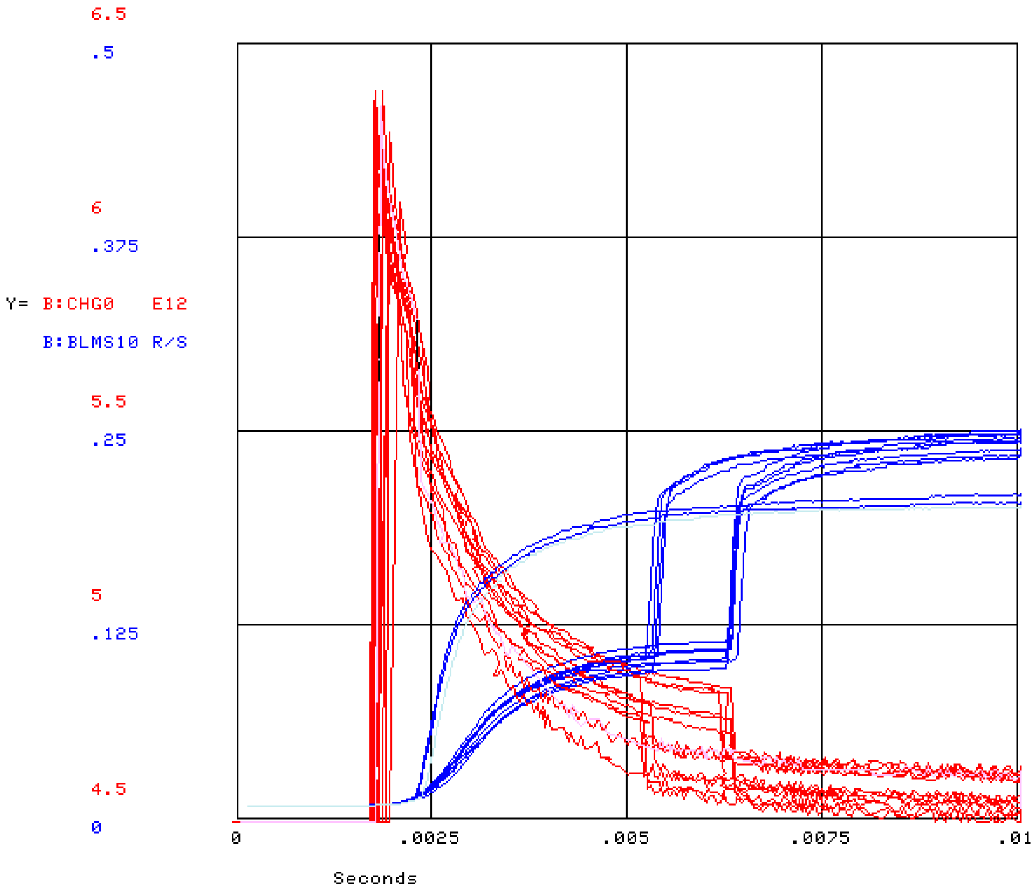

Figure 3.14: Notch creation within the Booster. Plotted are charge and loss traces from the first 10 ms of a Booster cycle. After charge is injected there is a slow loss due to Booster capture and collimation; the sharp loss is the creation of the notch which removes $\sim 4\%$ of the beam. The notch is shown being created at several different times in the cycle. The loss trace shows losses measured at a downstream location, caused by the notch creation.

A remedy to the above concerns was to remove the charge from the buckets of beam that would be swept over the septum. A system was developed which used a single extraction kicker to kick the beam vertically shortly after injection. As the beam is less stiff at that momentum, the kick from a single kicker is sufficient to put that portion of the beam entirely outside of the Booster's aperture. The pulse to this kicker is is tuned to remove ~ 3 bunches of Booster beam, while leaving the adjacent bunches as little disturbed as possible. The three bunches are still lost in the Booster, but at a lower energy of 400-600 MeV, instead of 8 GeV.

An example of the notching process as measured at the Booster is shown in Fig. 3.14. In that figure, the notcher is fired at either 2, 5.5, or 6.5 ms into the cycle. The loss of beam particles is immediate, and causes particle showers downstream. Fig. 3.15 shows the notch as measured in the MI by a resistive wall monitor.

The amount of radiation damage and residual activation is approximately proportional to the kinetic energy of the lost particles. Therefore, even though notching causes losses in the Booster, the effects of those losses can be reduced to 0.4 GeV $\div$ 8 GeV $= 5\%$, by forcing the loss at 400 MeV. Additionally, the location of the loss can be chosen: the notching magnet (notcher) was originally located in a region where there were relatively few components to damage or service; later the notcher was moved to a location so the beam could be lost into a system of collimators that were designed to absorb beam losses.

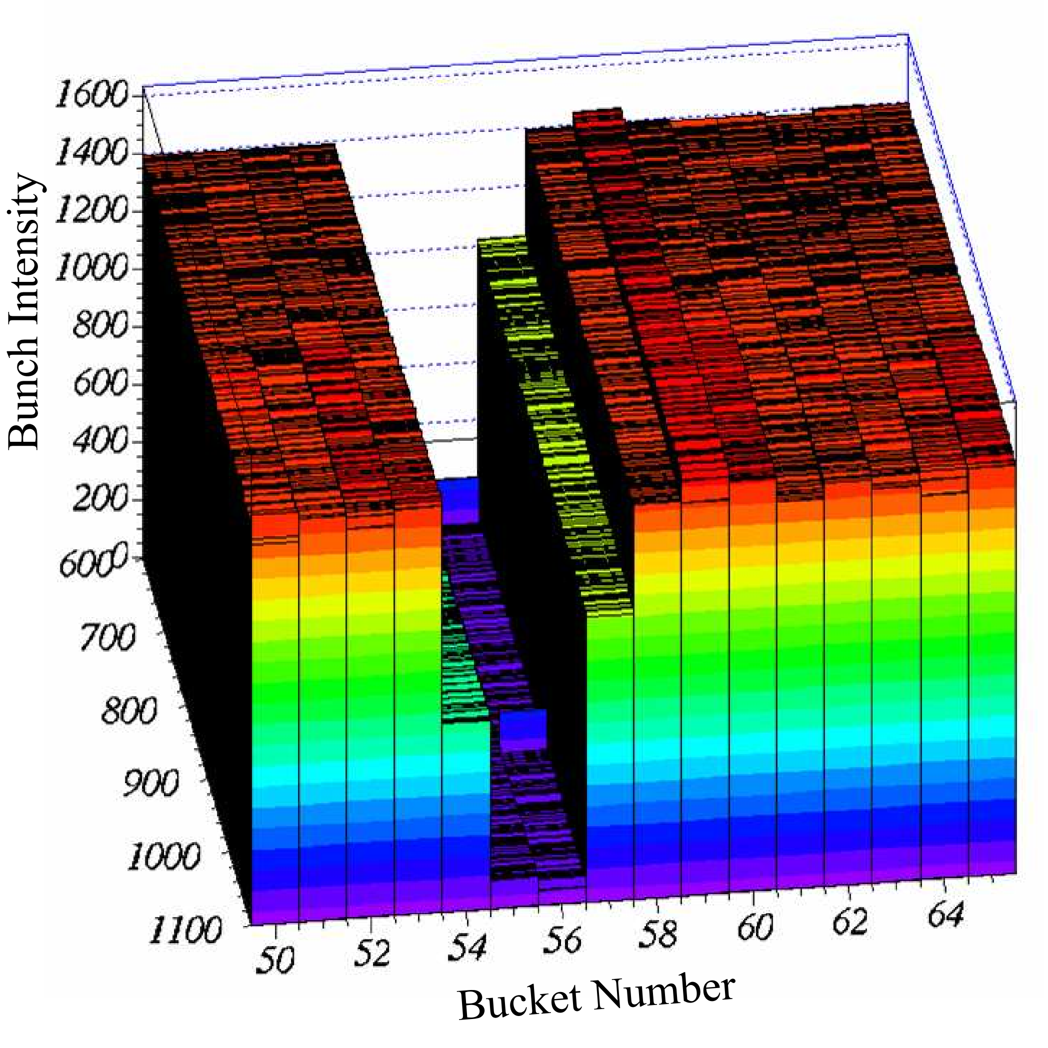

Figure 3.15: Measured beam charge of 16 RF buckets in the Main Injector. The notch is evident at buckets 54-57. This notch was created 5 ms into the Booster cycle. Two bunches are reduced in intensity by $\sim$ 95%, and two satellite bunches are partially reduced in charge. Figure courtesy: Phil Adamson.

3.7.5 Timing

The Booster Extraction Sync (BES) is the ultimate signal that triggers the Booster extraction kicker magnets to transfer the beam from the machine. How that signal is generated depends on several other systems. We will discuss extraction timing in terms of level of complexity: asynchronous extraction (e.g. to MiniBooNE), RF synchronous transfer to the MI, single batch to MI, and multi-batch to MI.

Common to all extraction methods is the Booster revolution marker (BREV). BREV is a signal generated by the Booster once every revolution of beam. It must scale with the speed of the beam, so it is actually generated by a "divide by 84 module" (÷84) which is fed the RF directly from the drivers and emits a TTL pulse every 84 counts[6]. The BREV marker corresponds to a particular bucket around the ring, usually the notch, or some offset from the notch. The BREV is used in some fashion for most types of extraction.

Asynchronous extraction from the Booster is used when there is no downstream accelerator to accept the beam; the Booster regularly performs this type of extraction to MiniBooNE and to the Booster beam dump. Approximately 400 μs after the start of the cycle the notcher magnet is fired by a timer and the BREV counter reset, so that BREV will be synchronized with the notch for the whole of the cycle. For extraction, a TCLK event around the time of the energy maximum[7]; this event signals the Booster extraction

[6]The Booster's harmonic number is 84, meaning that there are 84 buckets about the circumference.

[7]The Booster has en effective flattop of $\pm$ 100 μs. This is roughly the time that the beam can be held at the same frequency without significant changes in beam position due to the magnet currents changing. Because of the TCLK system's asynchronous nature and some of the wallsocket effects described in Ch. 4, the TCLK extraction signal can only be within $\pm$ 20 μs of the maximum energy. However, this variability is far less than the effective flattop time.

system to extract the beam. The extraction system pulses all of the slower magnets like the septum, waits a specified period of time, then waits for a BREV marker, and then pulses the kicker magnets to extract the beam[8]. As the BREV marker signifies the notch, the extraction is synchronized with it and no beam is lost on the septum magnet.

RF synchronous transfer to the Main Injector is a system whereby the beam is delivered to a particular location in the Main Injector without regard for the notch. This system is used regularly for proton injection to the Tevatron; it also was previously the system for doing multiple-batch injection and some of the details are still shared. In this case, beam must be delivered to a particular location in the Main Injector, but the Booster notch is not synchronized with that position[9]. Instead, the notch is not created at all; the BREV signal is created at the marker during injection, but is not used for the extraction. Instead the Main Injector sends a signal, known as the OAA[10], which the booster extraction should be synchronized to. As the beam nears extraction, phase lock is invoked so that the beam's frequency and phase match that of the MI. At the TCLK extraction signal the Booster disregards its BREV and instead waits for an OAA signal which it uses to trigger its kickers. 2-3 bunches of beam are lost on the septum magnet during the extraction. This type of transfer is is acceptable only if it occurs infrequently, such as proton acceleration for the Tevatron. Such losses, however, are not acceptable for proton delivery to NuMI.

[8] Generally, every device will have a specified delay associated with it to account for signal propagation times. We will neglect mentioning them unless they have specific relevance; however, setting and coordinating all the timing offsets is one of the more effort-intensive tasks associated with the extraction systems.

[9] The natural state of the Booster is for this synchronization to not exist. The next chapter will discuss the reasons for this.

[10] OAA stands for "offset AA", where AA is a hexadecimal signal corresponding to the Main Injector revolution marker.

Single-batch operation of the Main Injector does not require beam to be delivered to a particular location of the MI, so the notch can be used. In this case, the notch is created $\sim 400\ \mu$s after injection and a BREV reset is performed. Near extraction phase lock is performed, but the MI does not immediately provide an OAA marker. Instead, the Booster sends its BREV marker to the MI; the MI LLRF then uses the BREV marker to reset its OAA marker so that they are synchronized[11]. The new OAA is then sent back to the Booster and the Booster uses it to extract as described above for synchronous transfer. As the the OAA has been defined to be synchronous with the notch, no beam is lost on the septum during extraction.

Multiple-batch operation of the Main Injector is the concern of Chs. 4 - 5 where we implement a system that allows extraction using the notch to be used for all the Booster batches in a batch train. Originally, however, the multiple-batch running consisted of a single-batch transfer as above, followed by 1-6 RF synchronous transfers that did not use the notch. In this case the first batch was cleanly extracted, but every subsequent Booster extraction resulted in beam loss at the septum magnet. Such an unsynchronized method was only used for beam experiments and commissioning, but the radiation effects were still measurable in the Booster. Sustained operation for NuMI and antiproton production is only possible with clean extraction from the Booster, as is described in the next two chapters.

[11]As far as the MI LLRF is concerned the Booster beam is still being delivered to a particular location within the MI. Essentially, the MI redefines the positions around the ring to match the Booster beam. This redefinition only works when there is no beam within the MI that would otherwise define unique positions.

3.8 Performance

Precise limits have been placed on beam loss in the Booster as the demand for protons has increased. These limits define the effective metrics of Booster performance and determine the total number of protons deliverable to the experiments. The primary limits on Booster throughput are set by two measurements of loss: beam loss monitors (BLMs), and the Booster "watt-meter".

The operation of the BLMs was described above in §3.6.2. Here, we note that the BLM permit system limits losses at particular locations within the Booster. If the loss measured at a location is too high the beam will be inhibited until the 100 s running average falls below the acceptable level. Additionally, these trips often instigate operator tuning to reduce the loss at that location.

The Booster watt-meter uses beam loss information to calculate the total energy loss. The Booster charge is measured throughout the cycle and the derivative is taken as the loss at that time. The rate of loss is multiplied by the energy of the beam at that time and integrated to determine a total energy loss during the cycle:

$$P_l = \int_0^{33\text{ ms}} dt\, \frac{dQ}{dt} E(t) \tag{3.2}$$

The energy loss of all Booster cycles is kept in a 5 minute running sum. If the running sum exceeds a specific value then beam is inhibited until the running sum is within limits; as above, this sometimes causes operator intervention to improve the Booster tuning.

The watt-meter is a measure of the total power loss around the ring, as opposed the location specific BLMs. The watt-meter limit is set in the area of 1 W/m; this standard was developed from experience at various laboratories and is only a rule-of-thumb. The initial Booster limit was set at 400 W. However, the level has been increased to accommodate for the "controlled" losses into the collimator system. Today, the Booster typically runs around 450 W on the watt-meter while providing $\sim$ 30 kW of 8 GeV protons. Of note, is that the Booster efficiency[12] is 85-90%, but the energy loss is only 1-2%. This is accomplished as most of the beam loss occurs with creation of the notch and through collimation, when the beam energy is low ($<$ 1 GeV).

In the case of extraction losses, they occur at one location and are not included in the measurement of the watt-meter, as the device has no way to differentiate the zero charge of successfully extracted beam from the zero charge of lost beam. Therefore, the loss monitors are the only limit on these extraction losses. As mentioned above, the extraction region experiences significant radiation from beam loss, so the loss monitors must be set carefully to prevent inordinate radioactivation, while allowing beam operations.

The figure of merit for evaluating the quality of extraction in this dissertation will be the loss monitors around the extraction region. The loss caused by beam being lost on the septum magnet is roughly 20$\times$ that of cleanly extracted beam[13]; enduring this loss would limit Booster throughput of protons to $< 10^{16}$ / hour. That rate is insufficient even for only antiproton production, much less NuMI or MiniBooNE. Instead, the system described in the next two chapters allows the use of the notch in multi-batch mode so that extraction losses are not limiting the machine performance.

[12]Booster efficiency is defined as the ration of charge extracted to charge injected.

[13]This factor is reduced at high intensities where other effects increase the beam size, leading to collimation losses on the septum. In that case the ratio between notched and unnotched loss can be as low as 5, so the notch is not as dominant. However, that type of beam-quality loss at high intensity can potentially be reduced; as such, maintaining the notch remains significant.

Chapter 4

Longitudinal Evolution of the Booster Beam

4.1 Motivation

Extraction losses in the Booster can limit the number of protons delivered by the accelerators; the losses can be controlled only through extraction of the beam with a notch (see discussion is §3.8). The notch consists of 3 Booster buckets empty of beam. When synchronized with extraction the notch allows the kickers to reach full current without giving some bunches partial kicks; those partially kicked bunches would be lost on the septum of the extraction magnet. The beam must be extracted at exactly the time that the notch is at the kicker position on the circumference of the Booster. However, if the beam is destined for the Main Injector there might already be circulation beam in the MI, and the new beam must be in the correct position relative to the MI beam.

As discussed below, such dual synchronization between the Booster and

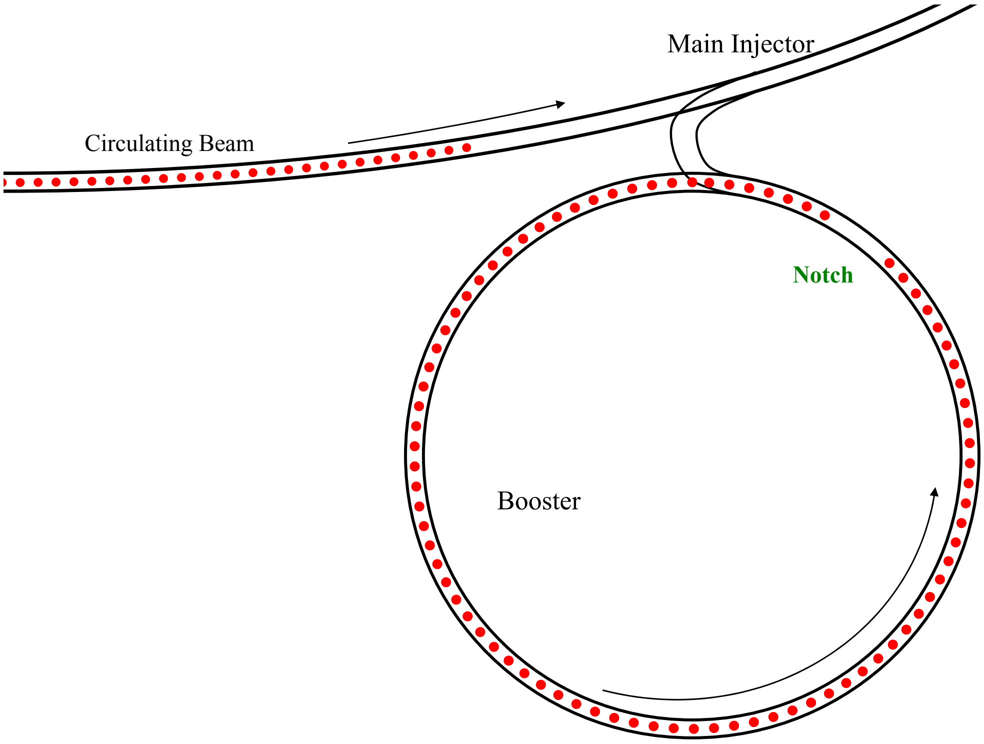

Figure 4.1: Illustration of the beam transfer situation. Beam can be revolving in both the Booster and Main Injector. To transfer from the Booster to MI, extraction must by synchronized with the notch and place the beam in the correct position within the MI.

Main Injector does not intrinsically exist. The primary culprit is the fact that the Booster RF frequency spans the range from 37 to 53 MHz. Even though the Booster's cycle is only 33 ms long, the beam orbits $\sim$ 20,000 times in that period; $\sim 1.6{\times}10^6$ RF buckets pass a given point. Furthermore, the difference between the RF frequencies of the Booster and MI means that the Booster beam travels $\sim$ 100,000 fewer buckets during that time period. A change of 2 in the number of buckets is enough to make the Booster miss the notch completely, meaning that having a fully repeatable timing structure between the Booster and MI would require consistency in all systems affected by the RF frequency to the 99.999% level.

This chapter will describe the longitudinal progress of the Booster beam, and sources of variation in that progress. Ch. 5 will describe the design and implementation of a system to force that synchronization. This chapter will provide: calculations of the Booster beam's longitudinal properties (§4.2); description of a hardware system to measure the beam's longitudinal progress (§4.3); initial measurements of Booster beam motion (§4.4); and an analysis of sources of variation in longitudinal beam motion, as well as methods to mitigate them where possible. Such partial mitigation reduces the demands on the synchronization methods of the next chapter.

4.2 Calculations

The slip factor is typically defined as the relative change in revolution period per relative change in momentum:

$$\eta = \frac{\Delta\tau/\tau}{\Delta p/p} \tag{4.1}$$

This can be interpreted as the slippage of the beam in revolutions, per revolution, and per proportional change in momentum. We will use a similar factor, but modified to be more useful for our applications:

$$\tilde{\eta} \equiv f_{\mathrm{MI}} - f_{\mathrm{B}} \tag{4.2}$$

Where f_{MI} and f_{B} are the Main Injector and Booster RF frequencies. We will call $\tilde{\eta}$ the slip rate. $\tilde{\eta}$ has units of buckets / sec, and describes the rate at which the Booster beam is falling behind the coasting Main Injector beam. The Main Injector frequency is fixed, as its beam is not accelerating during the period of Booster injections. The Main Injector injection frequency is $f_{\mathrm{MI}} = 52.8114$ MHz. The Booster frequency will start out much lower than the MI frequency

($f_{\mathrm{B}}(t=0) \sim 37.94$ MHz), but equal it at the end of a Booster cycle. The slip rate can also be expressed in terms of velocities and the Booster parameters:

$$\tilde{\eta} = f_{\mathrm{MI}} \left(1 - \frac{\beta_{\mathrm{B}}}{\beta_{\mathrm{MI}}} \frac{C_{\mathrm{MI}}}{7C_{\mathrm{B}}} \right) \tag{4.3}$$

Where β_{MI} and β_{B} are the relativistic velocities of the beam in the Main Injector and Booster. C_{MI} and C_{B} are the circumferences of the MI and Booster. The speed of the Main Injector beam is constant: $\beta_{\mathrm{MI}} = 0.994475$. We can express β_{B} in terms of momentum:

$$\beta_{\mathrm{B}} = \left[1 + \left(\frac{p}{mc} \right)^{-2} \right]^{-1/2} \tag{4.4}$$

Where p is the beam momentum, m is the mass of the proton, and c is the speed of light. During the cycle of the Booster, the momentum varies as an offset sinusoid (see §3.1):

$$p(t) = p_0 - p_1 \cos(2\pi f t) \tag{4.5}$$

Where $p_0 = (p_i + p_e)/2 = 4.9223$ GeV/c, $p_1 = (p_e - p_i)/2 = 3.9666$ GeV/c, $f = 15$Hz, and t is the time in the cycle; p_i and p_e are the injection and extraction momentums, respectively. Combining the above we can express the slip rate as a function of time during the Booster acceleration cycle:

$$\tilde{\eta}(t) = f_{\mathrm{MI}} \left[1 - \frac{1}{\beta_{\mathrm{MI}}} \frac{C_{\mathrm{MI}}}{7C_{\mathrm{B}}} \frac{1}{\sqrt{1 + \frac{m^2c^2}{[p_0 - p_1\cos(2\pi f t)]^2}}} \right] \tag{4.6}$$

The rightmost quotient within the brackets is equal to β_{MI} at t = $33.\bar{3}$ ms, so $\tilde{\eta}$ goes to zero at extraction. The Booster momentum and frequency during the acceleration cycle are plotted in Fig. 4.2 (slip rate is 52.8114 MHz - f_{B}).

We now define the total slippage as the integrated number of buckets

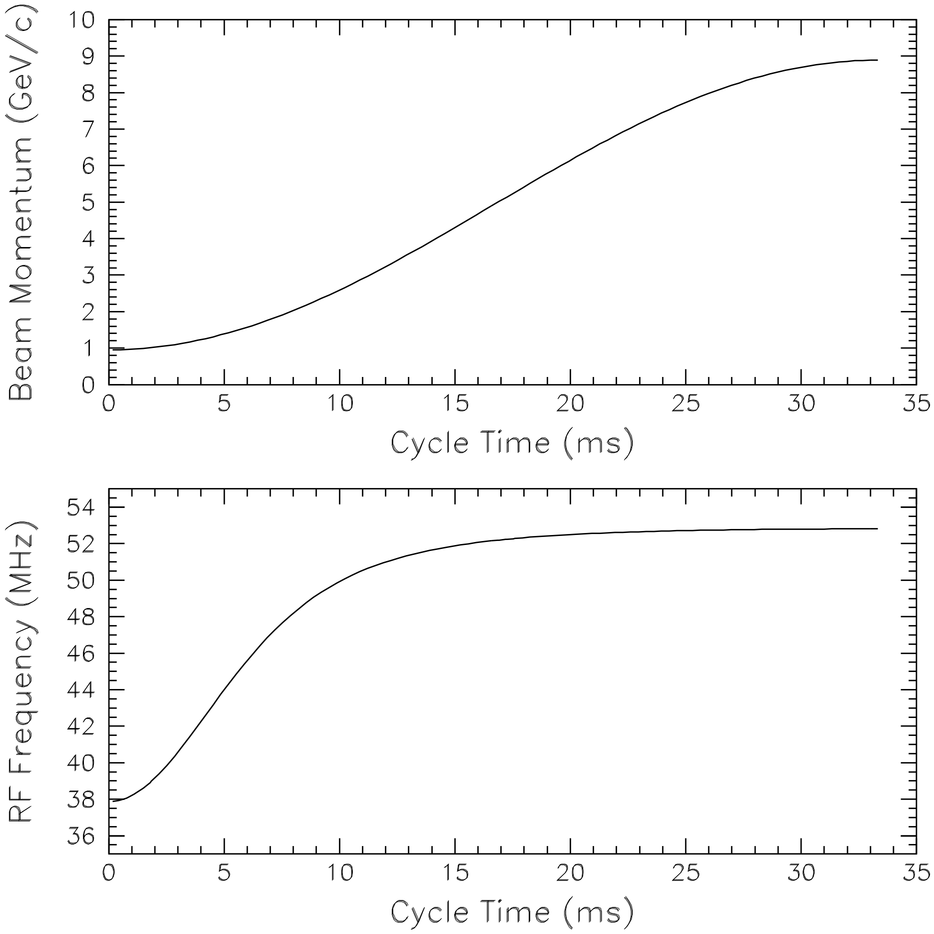

Figure 4.2: Idealized Booster beam momentum and RF frequency during the acceleration cycle. The slip rate, defined as $f_{\mathrm{MI}} - f_{\mathrm{B}}$ starts at ~ 15 MHz and rapidly decreases, reaching zero at the end of the cycle.

slipped, up to a certain time:

$$\begin{aligned} S_0(t) &= \int_0^t dt'\, \tilde{\eta}(t') \\ &= f_{\mathrm{MI}} t - \int_0^t dt'\, f_{\mathrm{B}}(t') \end{aligned} \tag{4.7}$$

This number is how much further Main Injector beam has traveled than Booster beam during the time period ($d = S_0 \times 5.645$ m/bucket). The value of S_0 is plotted for the Booster cycle in Fig. 4.3. The Main Injector beam passes $\sim$ 103,000 more buckets in the time of Booster cycle. If all Booster acceleration cycles slipped by precisely the same amount, the slip could be anticipated in designing the Booster extraction system. Unfortunately, $\tilde{\eta}$ suffers variation cycle-to-cycle, requiring feedback to correct it. A 0.001% change in $\tilde{\eta}$, for example, will results in a 1 bucket change in the total slippage. In the next sections we examine potential changes to $\tilde{\eta}$ and the effects upon slippage.

4.2.1 Cycle Variations

We consider variations in the Booster cycle that can change the total slippage amount. We take $p_0(t)$ for the Booster, as in Eqn. (4.5) and add variation:

$$p(t) = p_0(t) + \delta p(t) \tag{4.8}$$

Then, we find the relative change in slippage due to the change:

$$\begin{aligned} \delta S(t) &= S(t) - S_0(t) \\ &= \int_0^t dt'\, [\tilde{\eta}(t') - \tilde{\eta}_0(t')] \\ &\approx \int_0^t dt'\, \frac{d\tilde{\eta}}{dp}(t') \delta p(t') \end{aligned} \tag{4.9}$$

where $\delta S(t)$ is the cycle-to-cycle variation of the nominal slippage, $S_0(t)$. Below, we will consider the variation of constants that affect $p(t)$, so for a general

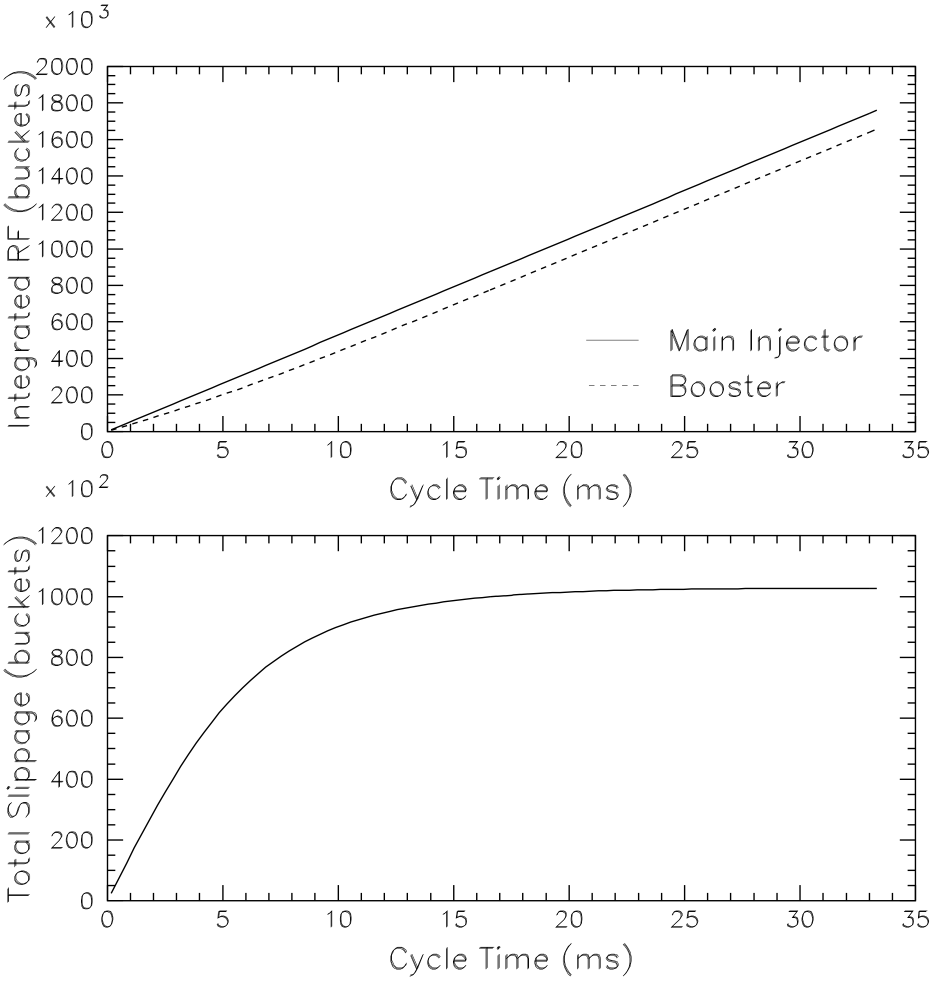

Figure 4.3: Top: integrated RF in the MI and Booster during a Booster cycle; essentially, the number of buckets is a measure of how far the beam has traveled. Bottom: the difference between the two above, or how much further the MI beam has traveled. This is also the result of integrating $\tilde{\eta}(t)$.

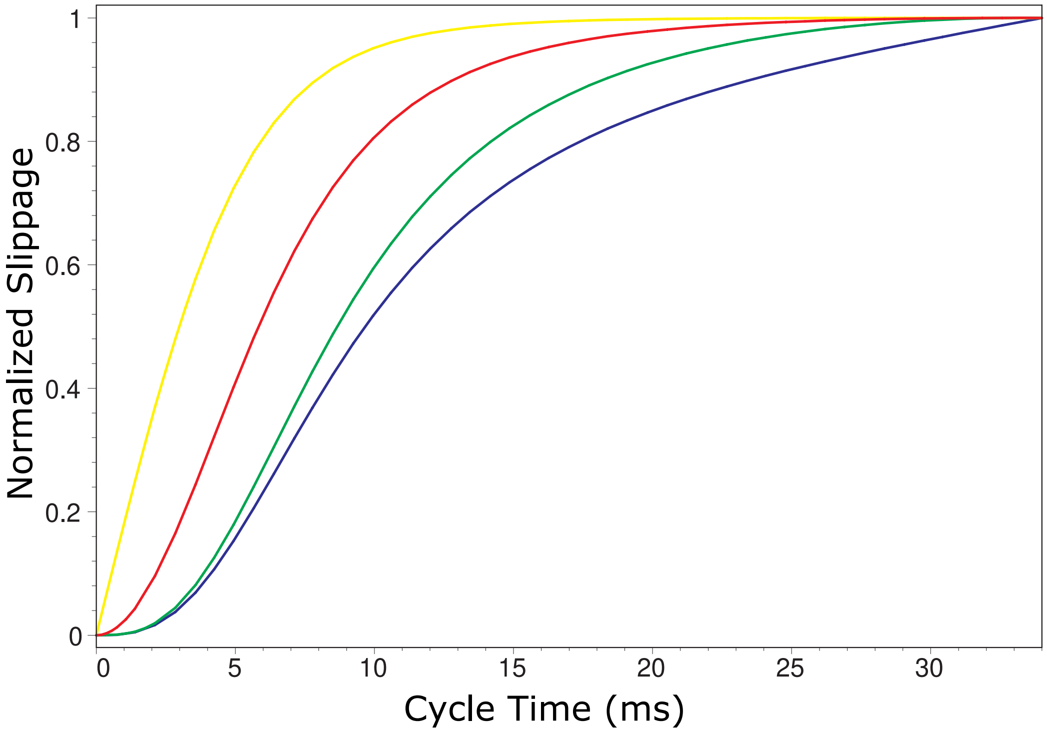

Figure 4.4: Calculated slippage due to various changes to the Booster's momentum parametrization, $p(t)$. From left to right they are: injection momentum (δp_i), timing offset (δt), magnet frequency change (δf), and extraction momentum (δp_e).

variable ξ we can describe the relative slippage as:

$$\frac{\delta S(t)}{\delta \xi} = \int_0^t dt' \, \frac{\partial \tilde{\eta}}{\partial \xi}(t') \tag{4.10}$$

The variables ξ, of whose effects on $p(t)$ we consider, are: a time offset δt in the start time of the Booster cycle, a Booster magnet frequency change δf, a change in the minimum magnet current δp_i, and a change in the maximum magnet current δp_e. Any parametrization of δp is possible, but we will consider these changes, particularly, in later sections.

Time Offsets

We consider first a variation in the momentum ramp, $p(t)$, due to a timing error, δt, at the start time of the Booster cycle:

$$p(t) = p_0 - p_1 \cos[2\pi f(t - \delta t)] \tag{4.11}$$

Then, we apply Eqn. (4.10):

$$\begin{aligned} \frac{\delta S_t(t)}{\delta t} &= \int_0^t dt' \, \frac{\partial \tilde{\eta}}{\partial t'}(t') \\ &= \tilde{\eta}(t) - \tilde{\eta}(0) \\ &= f_{\mathrm{B}}(t) - f_{\mathrm{B}}(0) \end{aligned} \tag{4.12}$$

The shape of this function is plotted in Fig. 4.4 as the red curve. It rises from 0 to a value of $\sim$ 15 MHz. This means that a 1 μs change in the start time results in a 15 bucket change in slippage by the end of the cycle.

Magnet Frequency Changes

We now consider a variation in $p(t)$ which arises from a variation δf in the magnet frequency:

$$p(t) = p_0 - p_1 \cos[2\pi(f - \delta f)t] \tag{4.13}$$

Applying Eqn. (4.10):

$$\begin{aligned} \frac{\delta S_f(t)}{\delta f} &= \int_0^t dt' \, \frac{\partial \tilde{\eta}}{\partial f}(t') \\ &= \int_0^t dt' \, \frac{f_{\mathrm{MI}}}{\beta_{\mathrm{MI}}} \frac{C_{\mathrm{MI}}}{7C_{\mathrm{B}}} \frac{1}{\gamma_{\mathrm{B}}^3(t')} \frac{p_1}{mc} (2\pi t') \sin(2\pi f t') \end{aligned} \tag{4.14}$$

γ_{B}(t) is the relativistic booster of the booster beam during the cycle. The integral was performed numerically and plotted in Fig. 4.4 as the green curve. The integrated change at the end of the cycle is $\sim$ 6 buckets of relative slippage for a 1 mHz change in the Booster frequency (out of 15 Hz).

Magnet Current Changes

We rewrite $p(t)$ in terms of p_i and p_e:

$$p(t) = \frac{p_i + p_e}{2} - \frac{p_e - p_i}{2}\cos(2\pi f t) \tag{4.15}$$

Again, we consider variations in $p(t)$ arising from δp_i and δp_e. Applying Eqn. (4.10) for each case:

$$\begin{aligned} \frac{\delta S_{p_i}(t)}{\delta p_i} &= \int_0^t dt' \, \frac{\partial \tilde{\eta}}{\partial p_i}(t') \\ &= \int_0^t dt' \, \frac{f_{\mathrm{MI}}}{\beta_{\mathrm{MI}}} \frac{C_{\mathrm{MI}}}{7C_{\mathrm{B}}} \frac{1}{\gamma_{\mathrm{B}}^3(t')} \frac{1}{mc} \cos^2(\pi f t') \end{aligned} \tag{4.16}$$

$$\begin{aligned} \frac{\delta S_{p_e}(t)}{\delta p_e} &= \int_0^t dt' \, \frac{\partial \tilde{\eta}}{\partial p_e}(t') \\ &= \int_0^t dt' \, \frac{f_{\mathrm{MI}}}{\beta_{\mathrm{MI}}} \frac{C_{\mathrm{MI}}}{7C_{\mathrm{B}}} \frac{1}{\gamma_{\mathrm{B}}^3(t')} \frac{1}{mc} \sin^2(\pi f t') \end{aligned} \tag{4.17}$$

These integrals are also performed numerically, and the results plotted in Fig. 4.4 as the yellow curve (p_i) and the blue curve (p_e). S_{p_i} grows to the value of ~ 100 buckets for a part-per-thousand change in p_i. S_{p_e} grows to the value of ~ 70 buckets for a part-per-thousand change in p_e. While the magnitudes of these are similar, note that $p_e \approx 8.4\, p_i$, so a relative change in p_e is a greater absolute change in magnet currents. The injection current has such a greater effect as $\tilde{\eta}$ is large at only the start of the cycle.

4.2.2 Radial Offset

In Ch. 5 we will adjust $\tilde{\eta}$ (providing feedback) by varying the radial position of the beam. Here we provide a calculation of the effect.

Assuming that the magnet lattice is the same, changing the beam's momentum from the ideal momentum at any time during the cycle changes the beam horizontal position around the circumference of the Booster according to dispersion (see §A.1). A (relatively) decelerated beam will trace a smaller circumference, while an accelerated one will trace a larger circumference. The changes in circumference and speed can change the revolution time, and thus the slip rate.

We start by calculating the variation of Eqn. (4.3) with momentum:

$$\frac{d\tilde{\eta}}{dp} = -\frac{f_{\mathrm{MI}}}{\beta_{\mathrm{MI}}}\frac{C_{\mathrm{MI}}}{7C_{\mathrm{B}}}\left[\frac{1}{C_{\mathrm{B}}}\frac{\partial\beta_{\mathrm{B}}}{\partial p} - \frac{\beta_{\mathrm{B}}}{C_{\mathrm{B}}^2}\frac{\partial C_{\mathrm{B}}}{\partial p}\right] \tag{4.18}$$

We use then find a form of $\partial\beta/\partial p$ in terms of γ:

$$\begin{aligned}\frac{\partial\beta}{\partial p} &= \left[1+\left(\frac{p}{mc}\right)^{-2}\right]^{-3/2}\left(\frac{p}{mc}\right)^{-2}\frac{1}{p}\\ &= \frac{\beta^3 m^2 c^2}{p^3}\\ &= \frac{1}{\gamma^3 mc}\end{aligned} \tag{4.19}$$

We then use momentum compaction in terms of the transition energy (see Eqns. (A.12) & (A.26)) to similarly find $\partial C_{\mathrm{B}}/\partial p$:

$$\begin{aligned}\frac{\partial C_{\mathrm{B}}}{\partial p} &= \frac{C}{p}\frac{1}{\gamma_{\mathrm{tr}}^2}\\ -\frac{\beta_{\mathrm{B}}}{C_{\mathrm{B}}^2}\frac{\partial C_{\mathrm{B}}}{\partial p} &= -\frac{1}{C_{\mathrm{B}}\gamma_{\mathrm{B}}\gamma_{\mathrm{tr}}^2 mc}\end{aligned} \tag{4.20}$$

We then get an expression for the derivative of the slip rate:

$$\begin{aligned}\frac{d\tilde{\eta}}{dp} &= \frac{f_{\mathrm{MI}}}{\beta_{\mathrm{MI}}}\frac{C_{\mathrm{MI}}}{7C_{\mathrm{B}}}\frac{1}{\gamma_{\mathrm{B}} mc}\left[\frac{1}{\gamma_{\mathrm{B}}^2} - \frac{1}{\gamma_{\mathrm{tr}}^2}\right]\\ &= \frac{56.6\,\mathrm{kHz/(MeV/c)}}{\gamma_{\mathrm{B}}}\left[\frac{1}{\gamma_{\mathrm{B}}^2} - 0.0337\right]\end{aligned} \tag{4.21}$$

In the above equation only γ_{B} varies with time during the Booster cycle and $C_{\mathrm{MI}} \approx 7C_{\mathrm{B}}$.

As expected, the slip rate derivative equals zero when at transition; that is, the beam is isochronous: to first-order, all particles have the same revolution frequency independent of momentum. Additionally, $d\tilde{\eta}/dp$ changes sign as it goes through transition. This reflects the fact that while the higher momentum beam has a higher velocity, when above transition the change in circumference more than compensates and the revolution frequency is lower.

For convenience, we usually like to think of the slip rate changing with radial (horizontal) position within the machine. Generally, the radial derivative is location specific, being dependent on the dispersion:

$$\begin{aligned} \frac{d\tilde{\eta}}{dx}(z) &= \frac{p}{D(z)}\frac{d\tilde{\eta}}{dp} \\ &= \frac{f_{\mathrm{MI}}}{\beta_{\mathrm{MI}}}\frac{C_{\mathrm{MI}}}{7C_{\mathrm{B}}}\frac{\beta_{\mathrm{B}}}{D(z)}\left[\frac{1}{\gamma_{\mathrm{B}}^2}-\frac{1}{\gamma_{\mathrm{tr}}^2}\right] \\ &= \frac{53.1\,\mathrm{MHz}}{D(z)}\beta_{\mathrm{B}}\left[\frac{1}{\gamma_{\mathrm{B}}^2}-0.0337\right] \end{aligned} \tag{4.22}$$

For further convenience, we average $D(z)$ around the ring and define radial position as $r = C/2\pi$; then:

$$\frac{d\tilde{\eta}}{dr} = \left(21\,\frac{\mathrm{buckets}}{\mathrm{mm}\cdot\mathrm{ms}}\right)\beta_{\mathrm{B}}\left[\frac{1}{\gamma_{\mathrm{B}}^2}-0.0337\right] \tag{4.23}$$

The value of $d\tilde{\eta}/dr$ during the Booster cycle[1] is plotted in Fig. 4.5. The relative slip rate is given in units of buckets / mm / ms; meaning that with a value of 1, a 1 mm shift in average radial position will cause the beam to slip at a rate of 1 RF bucket per ms compared to what it would have otherwise.

4.3 Timing Measurement System

A system was implemented in the past [70] to first measure slippage in the Booster and test forms of radial manipulation. This dissertation will build off

[1] Plotted is the negative of Eqn. (4.23) as the sign convention is arbitrary. We will mostly be interested in the values after transition, so they were chosen to be positive

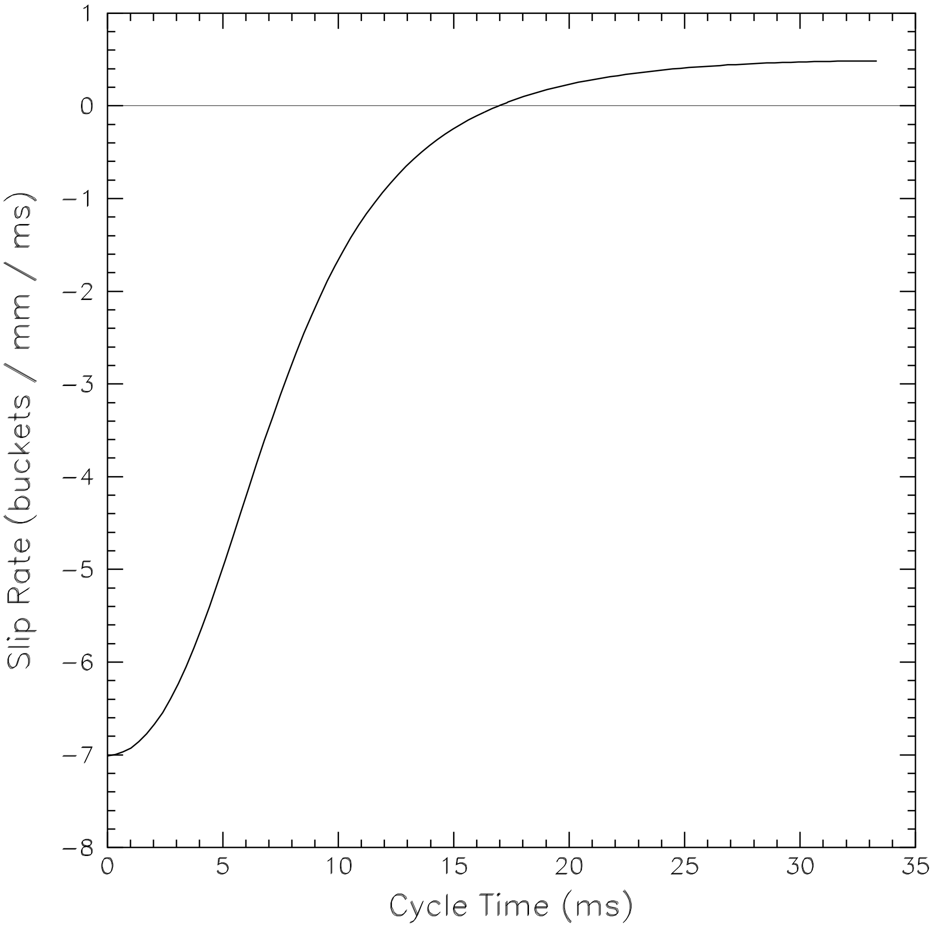

Figure 4.5: Calculations of $d\tilde{\eta}/dr$ during the Booster cycle. Note that the curve crosses zero at transition.

that system to fine tune the counting method and develop the manipulation (in Ch. 5) into a working system. In this section, we describe the hardware platform and algorithms for making useful longitudinal measurements.

4.3.1 Hardware

The hardware platform was chosen early on to be the "Booster Generic DSP Board". The DSP (Digital Signal Processor) board is Fermilab designed and operates inside a VXI crate. Each board has 1 DSP and 4 FPGAs and on-board shared memories. An FPGA (field-programmable gate array) is a programmable integrated circuit capable of performing logic on signals (see [71]). The front panel has 4 digital inputs, 4 digital outputs, 3 analog inputs, 3 analog outputs, 2 RF inputs, as well as two inputs for the proprietary Fermilab signals of TCLK and MDAT. The backplane allows a Motorola PPC 603 processor in the VXI crate to read and write to the shared memory. The PPC operates using VXWorks, enabling network communication to Booster front-end and UNIX systems, allowing integration into the Fermilab ACNET system and data storage.

The on-boards DSP is a Analog Devices SHARC processor, model number ADSP-21062. It is a 40 MHZ DSP that can access shared memory, accept and transmit data to the FPGA, and allow hardware interrupt-driven code. The DSP is programmed using assembler language and compiled with a PC, the executable code is loaded onto the DSP over the network. We use the DSP primarily for memory access and calculations.

The DSP board contains 4 Altera FPGAs, but only one will be used for the processing described in this dissertation (the others provide for the wiring of the board and are essentially static). The logic is defined by a program provided on-board via an EEPROM (electrically erasable programmable read-only memory). The FPGA allows great flexibility in design as its logic functions can be reconfigured by reprogramming the EEPROM. Additionally, the FPGA can have multiple inputs into different parts of itself and perform logic in parallel, allowing logic-type computation much more quickly than possible with a DSP or microprocessor. We use the FPGAs for fast counting of the RF clocks, various internal timers, and for the reception and transmission of various external signals.

The Altera FPGA is of the Flex10k type, has $\sim$ 30,000 logic gates, and can operate in excess of 100 MHz. The FPGA is programmed in AHDL (Altera hardware description language) on a PC where the EEPROM is burned. The EEPROM must be installed physically and the system rebooted to reprogram the FPGA. The FPGA is connected to the front panel where it may send & receive the digital signals; it is also connect to the RF inputs which include comparators and generate digital versions. The FPGA is connected to the DSP in such a way that it provides two separate hardware interrupts.

4.3.2 Tracking Algorithm

The goal of this first algorithm is only to measure the longitudinal progress of the Booster beam, and variations thereof[2]. We will us the Main Injector RF as a standard clock, as the MI RF is fixed by a digital source during injection. We measure the distance (in RF buckets) between an MI revolution marker and a Booster revolution marker. A change in that difference changes the relative locations of the notch and circulating beam in the MI. The differences are stored as arrays and can be compared between cycles.

[2]As a reminder: we are interested in variations of slippage of as small as 1 in 100,000.

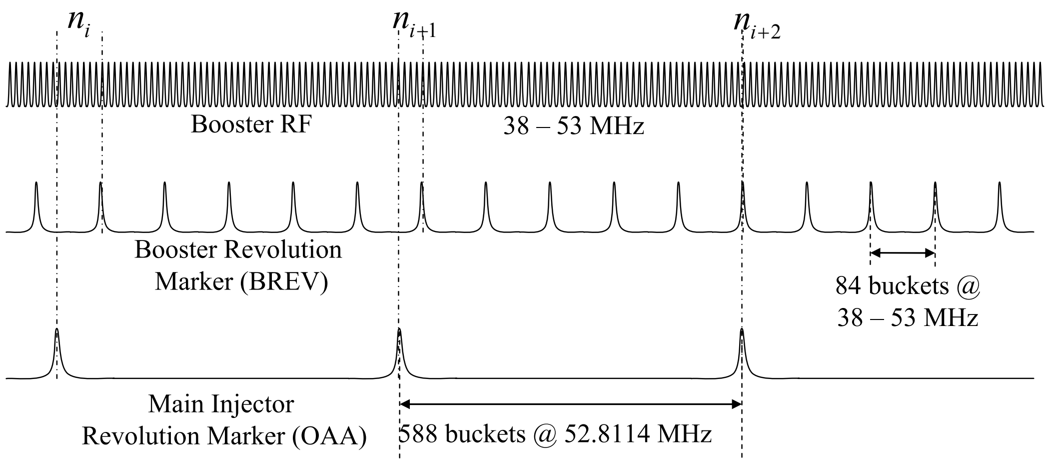

Figure 4.6: Schematic of the signals fed into the counting system. The Booster RF is read directly from the RF drivers. The Booster revolution marker (BREV) pulses once every 84 Booster RF pulses. The Main Injector revolution marker (OAA) pulses once every 588 MI RF pulses. As shown in Fig. 4.7, the OAA starts a counter, which increments every Booster RF pulse and stops at the first subsequent BREV. The counts are stored in an array.

The outline of the counting scheme is shown in Fig. 4.6, and the hardware logic layout in Fig. 4.7. The FPGA receives 3 digital signals: a TTL version of the Booster RF, a Booster revolution marker, and a Main Injector revolution marker. At this point, the precise timing of the revolution markers are unimportant, as long as they recur every 84 and 588 buckets; later, we shall decide upon concrete timing definitions for the revolution markers. For now, we consider the Booster revolution marker to be an externally provided signal that corresponds to the notch in the Booster: the BREV signal. We consider the MI revolution marker to be the signal sent by the MI LLRF that specifies the desired position of the Booster beam in the MI: the OAA signal.

As shown in the figures, the OAA is used to reset and start a counter. The counter then counts Booster RF buckets until it receives a BREV. The OAAs are spaced at 11.13 μs, so there will be about 3000 counts generated during the Booster cycle. The count is never more than 84 buckets as the

BREVs are spaced by 84 buckets[3].

The count is sent to the DSP along with an interrupt informing the DSP that data is available to be read. The DSP takes the value and stores it in shared memory in an array. The DSP also recalls a value from another array (with the same index) that was generated on a previous Booster cycle. The DSP then finds the difference between the two counts – if the Booster cycles were identical then the difference should be zero. The difference is always in the range of [-83,+83], but values separated by 84 are equivalent. Thus, the DSP performs a correction step that verifies the difference is similar to the previous difference[4]; the resulting number is known as the turns-corrected error (TCE). The TCE is stored to memory and is also output via a DAC so that the TCE can be plotted on an oscilloscope during a Booster cycle.

4.4 Timing Measurements

This section describes some of the first measurements made with an updated counting system in 2003. These measurements were significant as they set the scale of slippage variations to be corrected: it had been known that the Booster beam slipped and that the notch position was inconsistent after a specific time interval; however the magnitude of the effect was unknown.

Fig. 4.8 shows the 3000 values of the count array for a single cycle of Booster beam. Plots for subsequent Booster cycles are superficially similar to Fig. 4.8 as they are dominated by the Booster's slower frequency: at injection the Booster's velocity is only $\sim 5/7$ that of the MI, so it falls behind two

[3]Note that the range of possible values would change during the cycles if we counted MI RF buckets, instead of the Booster RF buckets

[4]While the count values vary quickly, the differences on subsequent turns cannot vary substantially without extreme changes in the Booster.

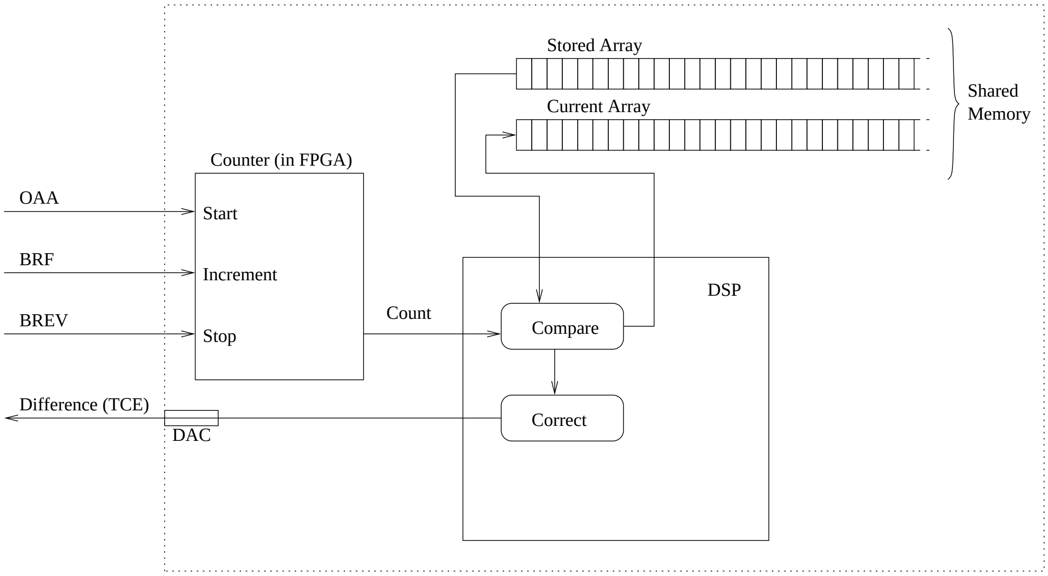

Figure 4.7: Overview diagram of the counting process. The FPGA prepares counts of Booster RF cycles with the MI OAA as a start and the BREV as a stop. The count is stored to a shared memory array by the DSP and compared to a previously recorded value, stored in another array. The difference is calculated and corrected to be continuous with the previous values (un-mod 84). The difference (TCE) is output via a DAC. The arrays start at the beginning of the Booster cycle and have $\sim$ 3000 entries each.

Booster circumferences every time a count is taken (every MI revolution). At then end of the cycle, as the Booster frequency approaches that of the MI the curve levels off and locks to the MI.

Fig. 4.9 shows the results of taking the difference of each pair of consecutive counts in a single cycle (the data are nominalized to fit int the range of [0,83]). We expect the Booster to slip between counts a distance defined by the current slip rate and the Main Injector revolution period:

$$S_\tau \approx \tilde{\eta}(t) \times \tau_{\mathrm{MI}} \tag{4.24}$$

The slip rate does not change rapidly during the cycle, so it is relatively constant over a count period. The data in the figure provide a measurement of the slip rate throughout the cycle. The granularity prevents precise measurements

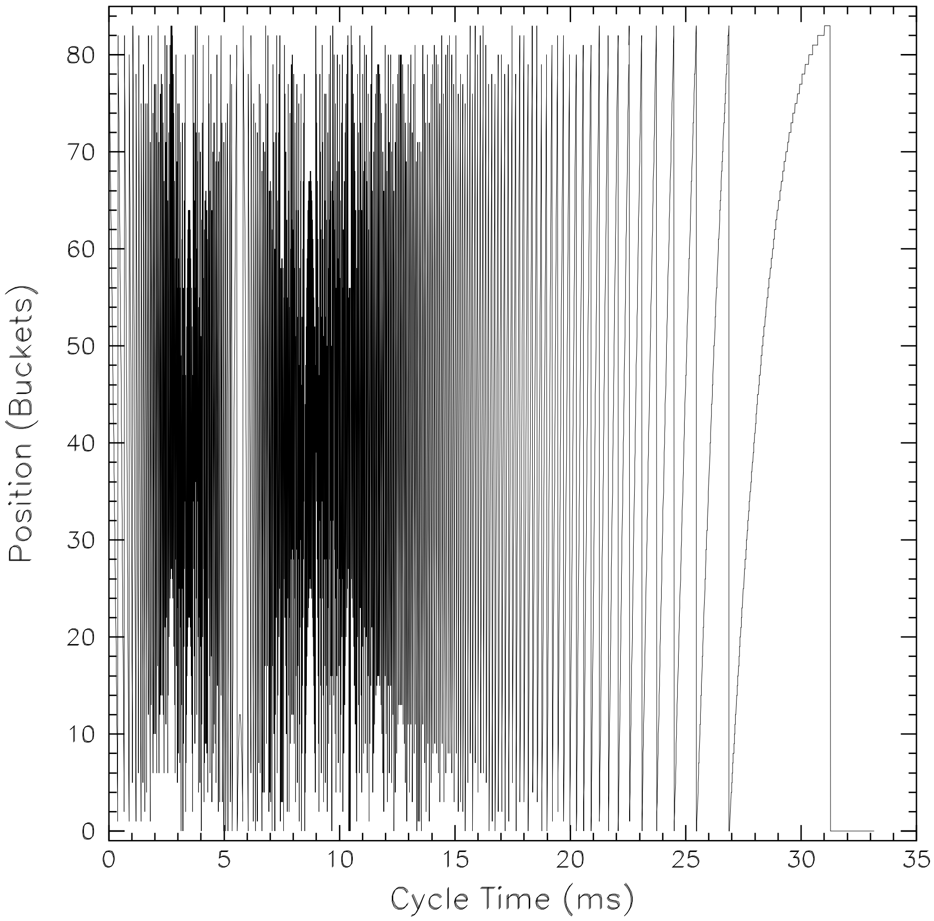

Figure 4.8: Measured position of the Booster beam with respect to the MI beam throughout one Booster cycle. The Booster beam falls behind the MI beam by about 1200 revolutions, causing the wrapping. At the very end of the cycle the Booster frequency approaches the MI and they stop changing.

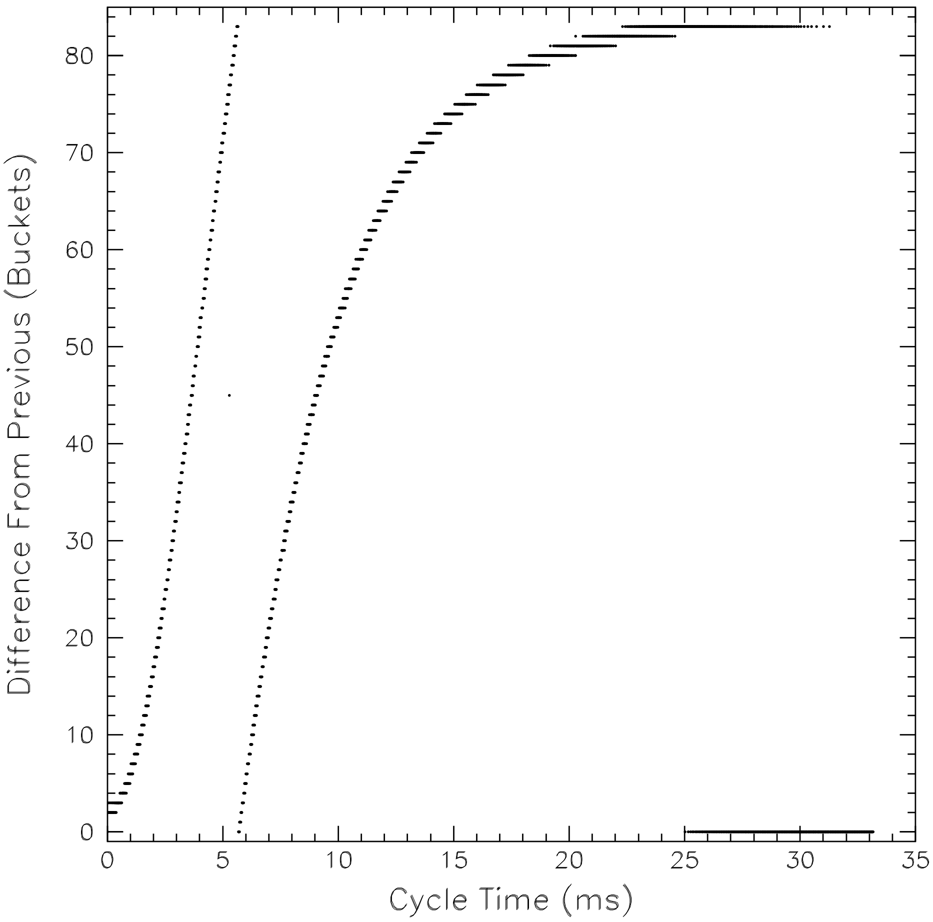

Figure 4.9: Measured difference (mod 84) between successive position measurements. This number represents how many Booster buckets fewer the Booster beam traveled than the MI beam. All of these numbers are “negative”, but because of the modulus they are all brought to the range of [0,83]. The data points at < 5.5 ms slip by more than a Booster circumference.

over short times scales, but they conform to our expectations from Fig. 4.2 and simple simulations of the counting.

Finally, the measurement we are interested in is the relative slippage, or TCE. As described above, the TCE is generated by finding the difference between the measured position at a given time in the cycle, and and that from a previous cycle at the same time in that cycle. The results are shown in Fig. 4.10 for ten seprate Booster cycles. The slippage starts at zero in each case, but then evolves quickly, until leveling off after transition.

Referring back to our calculations in §4.2 and the resulting curves in Fig. 4.4, we see consistent behavior in the measured curves with our predicted ones. The bulk of the slippage will always occur at the start of the cycle as that is where the slip rate is its greatest and most sensitive.

While the shapes of the curves look reasonable, we had no *prima facie* expectation of what the magnitude of the variations should be; in fact, the measured radial position for all of these cycles follow the same trajectory (shown in Fig. 3.7). The maximum excursion of nearly 200 buckets corresponds to a difference of more than two circumferences[5]. These values of slippage are much larger than the ~ 1 RF bucket accuracy required for successful beam extraction using the notch.

4.5 Sources of Deviation

Having measured the variation, we now turn to the potential sources for it. Recall the general form for relative slippage Eqn. (4.10): any variation in the

[5]The concept of relative slippage can be elusive in its precise meaning. Here, however, it means that one Booster cycle went through 180 more RF cycles in the same amount of time as another cycle. The faster beam traveled a further distance of 1000 m during the acceleration cycle.

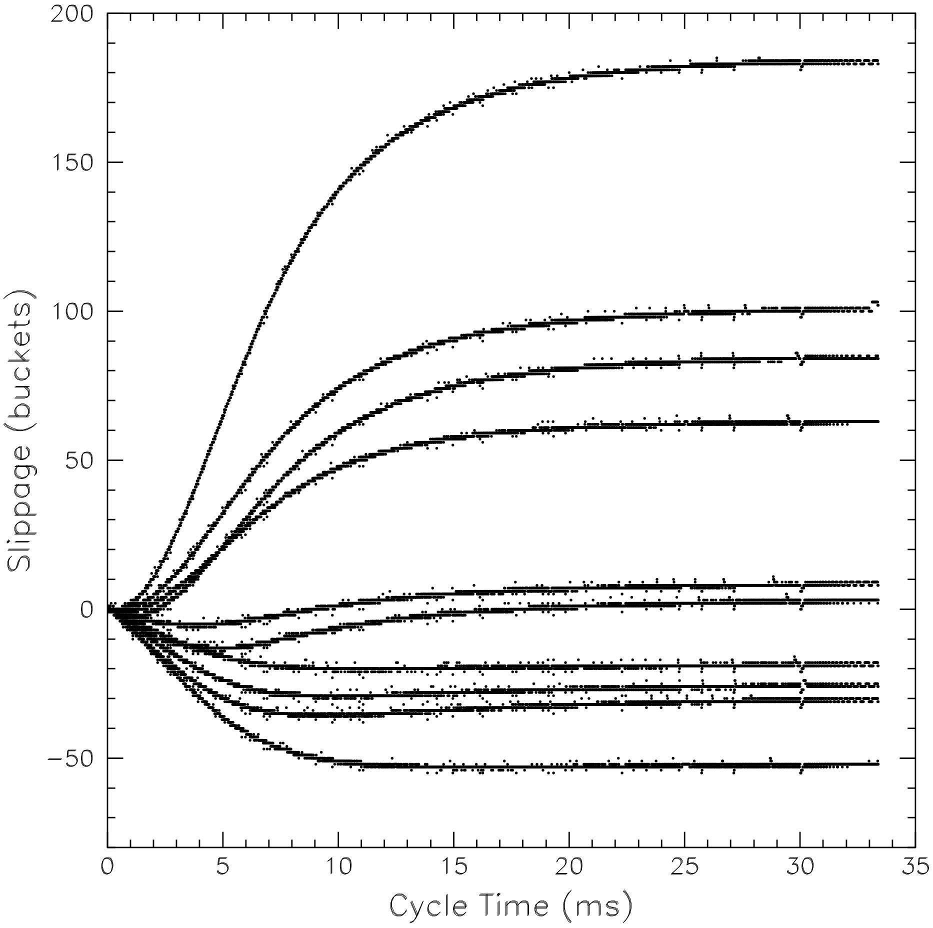

Figure 4.10: Initial measurements of slippage in the Booster (data from July 15, 2003). The slippage is measured as growing to almost 200 buckets, meaning that notch created early in the cycle would move from an expected position by as much as 2 $^1/_2$ circumferences. Some artifacts from miscounting in the prototype system cause the noisy appearance.

Booster RF frequency leads to a change in the slip rate and total slippage. We also showed that any deviation in magnet current (beam momentum) or the beam's radial (horizontal) position induces a change in slippage.

One source of variation could be the radial position. A change of position on the orders of $\sim 20\ \mu$m would be enough to cause a 1 bucket shift by the end of the cycle. Unfortunately, we lack the ability to measure position that accurately. Additionally, as discussed in §3.5, the Booster's orbit changes significantly around the ring during the acceleration cycle. The most consistent part of the orbit is that in the region of RPOS detector as the LLRF keeps the radial position their to its preprogrammed values. However, the precision of the BPM limits that control[6]. We take the stability of the Booster orbit as an irreducible error in this case.

On the other hand, we can consider some specific variations in the Booster's momentum (i.e. magnet currents), and partially correct them. The below sections analyze the issues of timing, magnet current variation, and magnet frequency variation.

4.5.1 Timing

One of the first issues to become clear was that of timing. The evidence, shown in Fig. 4.11, was that we measured significant slippage when there was no beam. No BPM issues or magnet current issues could cause this variation as there was no feedback. Instead, the LLRF generates default frequency curves from a digital source whose variation should be less than 1 Hz.

Further investigation showed that there were a number of issues that could lead to timing errors at the 10 μs level and lead to slippage with the form

[6]Furthermore, the Booster BPMs' response has a weak intensity dependence. Even a small deviation, though, can lead to a several bucket offset.

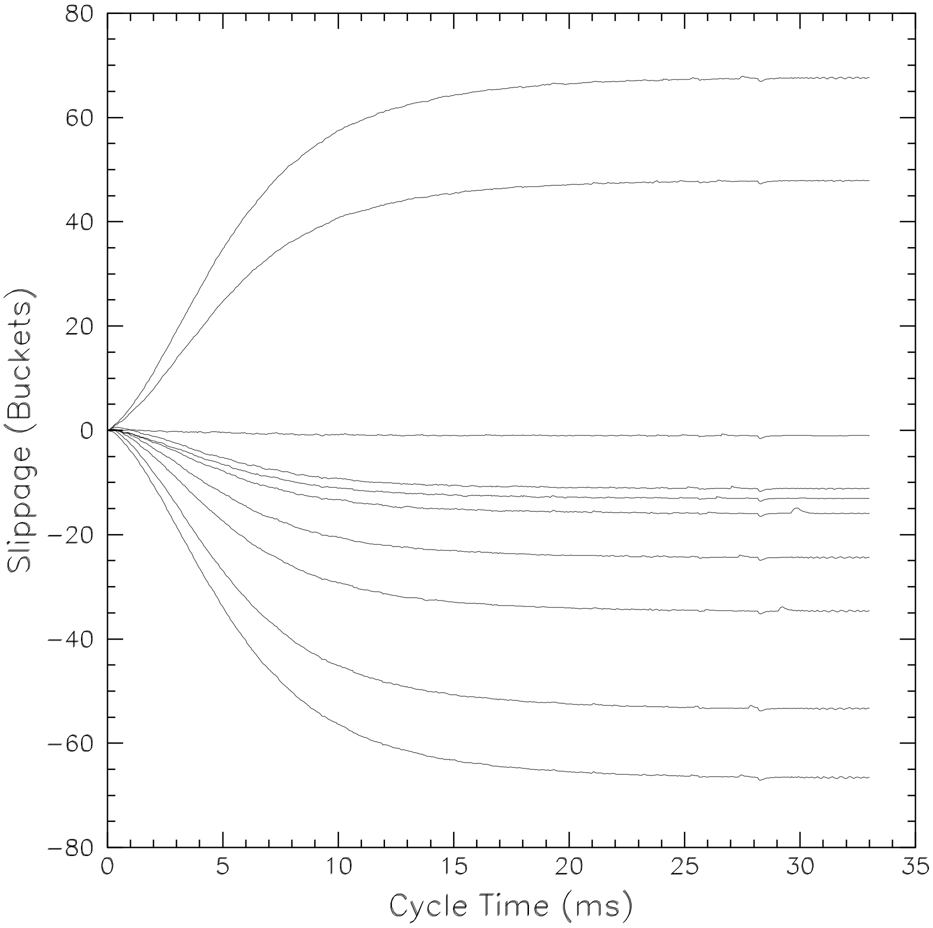

Figure 4.11: Measured slippage on Booster cycles without beam. The default Booster RF curves can generate slippage because of timing errors. Compare to Fig. 4.4. The total slippage range is consistent with timing errors of ~ 10 μs. The datapoints have been smoothed for clarity via an exponential moving average.

of Eqn. (4.12). First was the issue of what, precisely, started the Booster curves and the measurement program; these issues could produce timing errors of $\sim$ 10 μs. Second was the issue of revolution markers, which had no correlation with the start of the Booster cycle, and could produce timing errors of $\sim$ 11 μs.

The errors and attempts to rectify them are discussed further in the below sections. Here we note that with substantial rewiring of the triggers for the Booster RF and slippage measurement, the trigger errors were mostly eliminated, and through the addition of several extra counters and book-keeping in the FPGA and DSP we eliminated the problem caused by the revolution markers.

The fruits of the efforts are shown in Figs. 4.12 & 4.13. Fig. 4.12 shows the measured slippage without beam once all the timing errors had been corrected. The measured slippage was at the level of 1 RF bucket, or 20 ns – demonstrating that the timing relative to the RF and the revolution markers had been corrected.

Fig. 4.13 shows the measured slippage with beam and with all of the timing corrections. In this case, the beam frequency is controlled by the LLRF as described elsewhere, so slippage can still occur from other variations in radial position or magnet currents. The slippage measured is less than that in Fig. 4.10 by $\sim$ 60%, but is still far too much for notched extraction. Also, the shape of the slippage curves appears more complicated than in Fig. 4.10; assumedly this is because the dominating, simple timing curve variation has been removed and we are left with other, more subtle, variations.

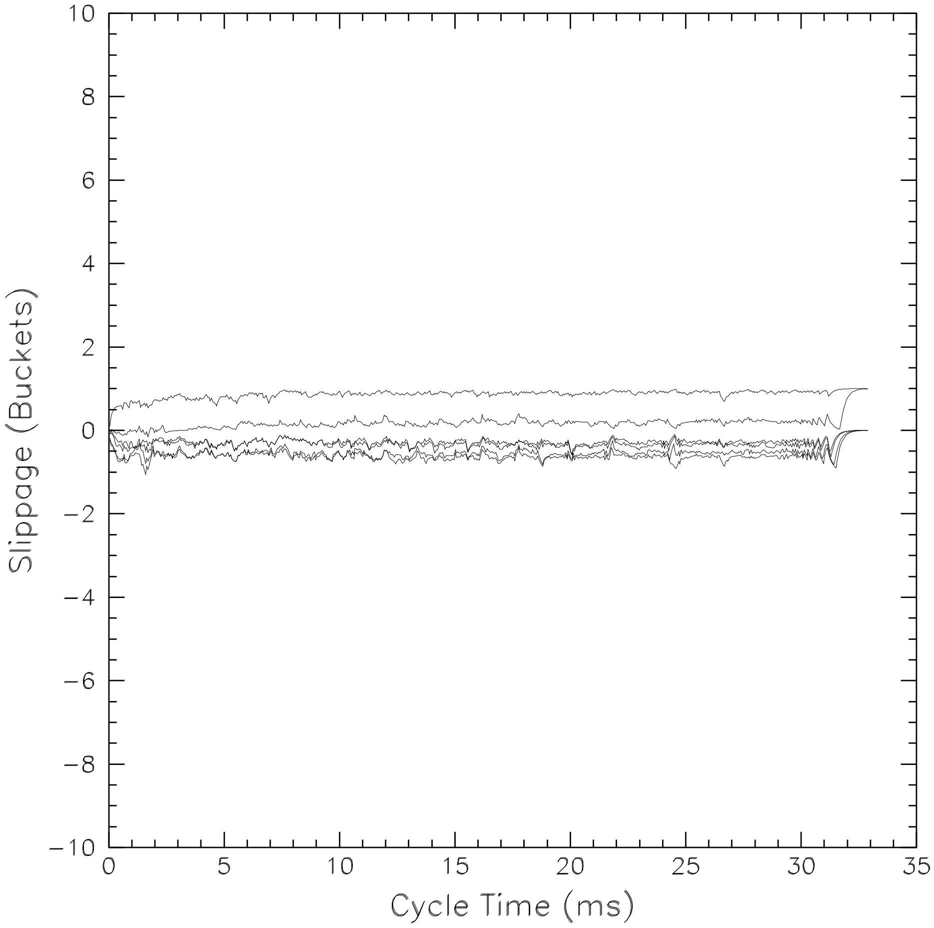

Figure 4.12: Measured slippage on Booster cycles without beam, but with correction to the various timing errors. Compare to Fig. 4.11. The remaining variation is on the order of 1 bucket, suggesting that timing is now consistent to the 20 ns level. The data points have been smoothed for clarity with an exponential moving average.

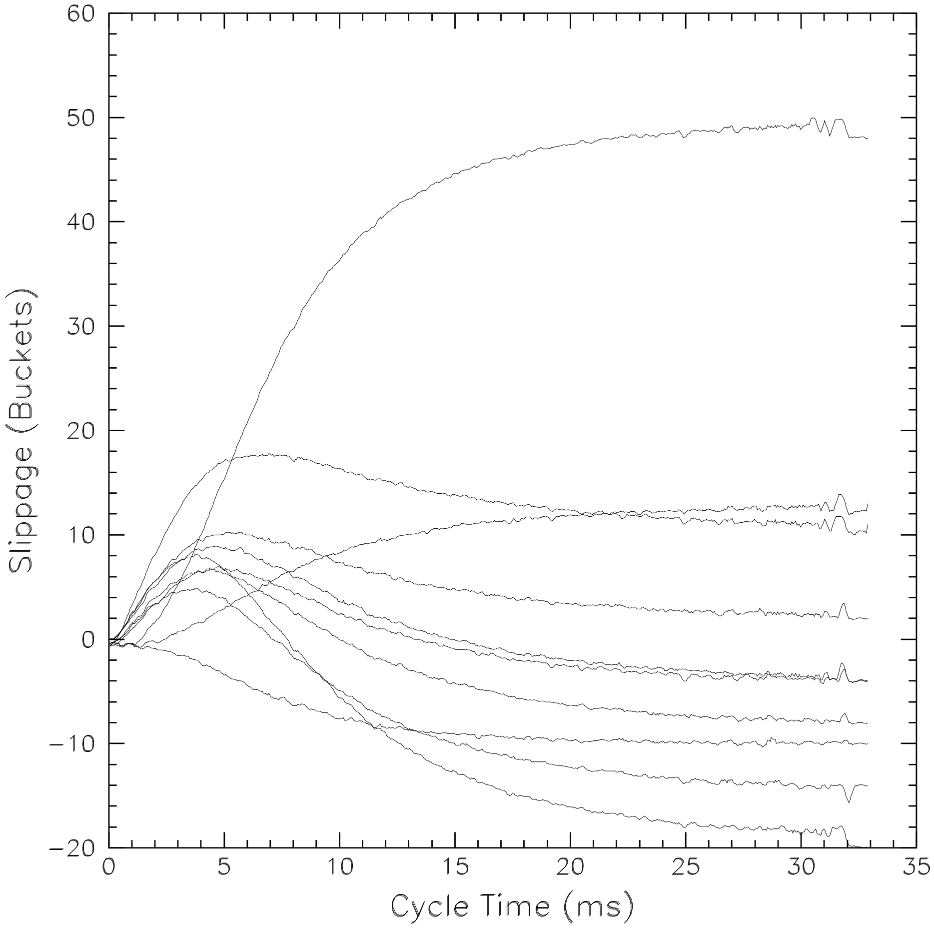

Figure 4.13: Measured slippage on Booster cycles with beam, but with the timing corrections – compare to Fig. 4.10. The range of slippage has been reduced by 60%, but still amounts to ~ 1 circumference due to other variations.

Cycle Trigger

The issue of the cycle trigger has several components that deal with how the beam injection is triggered, how the default RF curves are triggered, and how the measurement system is triggered with relation to those two and the Booster magnets.

We started with the realization that all of these systems had been triggered with a TCLK Booster reset that occurred ~ 2 ms before injection to the Booster. This was necessary for some systems,such as RF drivers and pulsed magnets, so that their power supplies could be charged before beam arrived. However, the TCLK signal is not well timed with the Booster cycle and is generated asynchronously such that variations > 1 μs are possible.

As will be discussed below, the wallsocket AC frequency varies, causing the Booster frequency to vary. The TCLK system uses the two previous magnet current minimums (measured by a BDot zero crossing) to extrapolate to the current Booster cycle. However, that extrapolation occurs over a lever arm of 66 ms – a reasonable variation of 1 mHz in frequency would lead to 4 μs variation is timing, which could result in a slippage of 60 RF buckets.

Therefore, it was decided that the slippage measurement and RF systems should be triggered on a more relevant signal. The Booster BDot zero-crossing was used with some success, but we found that the measurement system and RF curves also required some time before beam injection for initialization. A system using a BDot level crossing was chosen to trigger the LLRF and measurement system ~ 80 μs before the magnet current minimum.

Revolution Markers

The issue of revolution marker was discovered when going from a test system's OAA to the real one from the MI. We had been using our own OAA, generated by a $\div$ 588 module, and our own MI RF source. These were triggered at the start of the measurement program. However, simply by changing to an MI OAA we saw a large amount of timing variation like in Fig. 4.11.

The source of this variation was the asynchronous relationship of the MI RF with the Booster magnet's resonant cycle. As shown in Fig. 4.14, the measurement cycle was triggered on the magnet minimum, but would not start counting until the next OAA came by – a variable period of up to 11 μs. This time offset persists throughout the cycle and is equivalent to mistiming the initial triggering of the measurement cycle.

The solution, as also shown in Fig. 4.14, was to generate an internal timer (dubbed the fOAA) that is synchronized with the start of the measurement program, and repeats at the same rate as the OAA (once every 588 MI RF cycles). However, the OAA still represented the intended position of the beam, so the interval between the fOAA and OAA was measured and is applied at various times in the cogging program discussed in the next chapter.

A similar internal marker was established in lieu of the Booster revolution marker (BREV), but is not as significant as the BREVs are more closely spaced. Additionally, the BREV orientation with respect to the OAA is not fixed, so variation in its initiation only led to a slightly more noisy measurement, not a systematically slipping one.

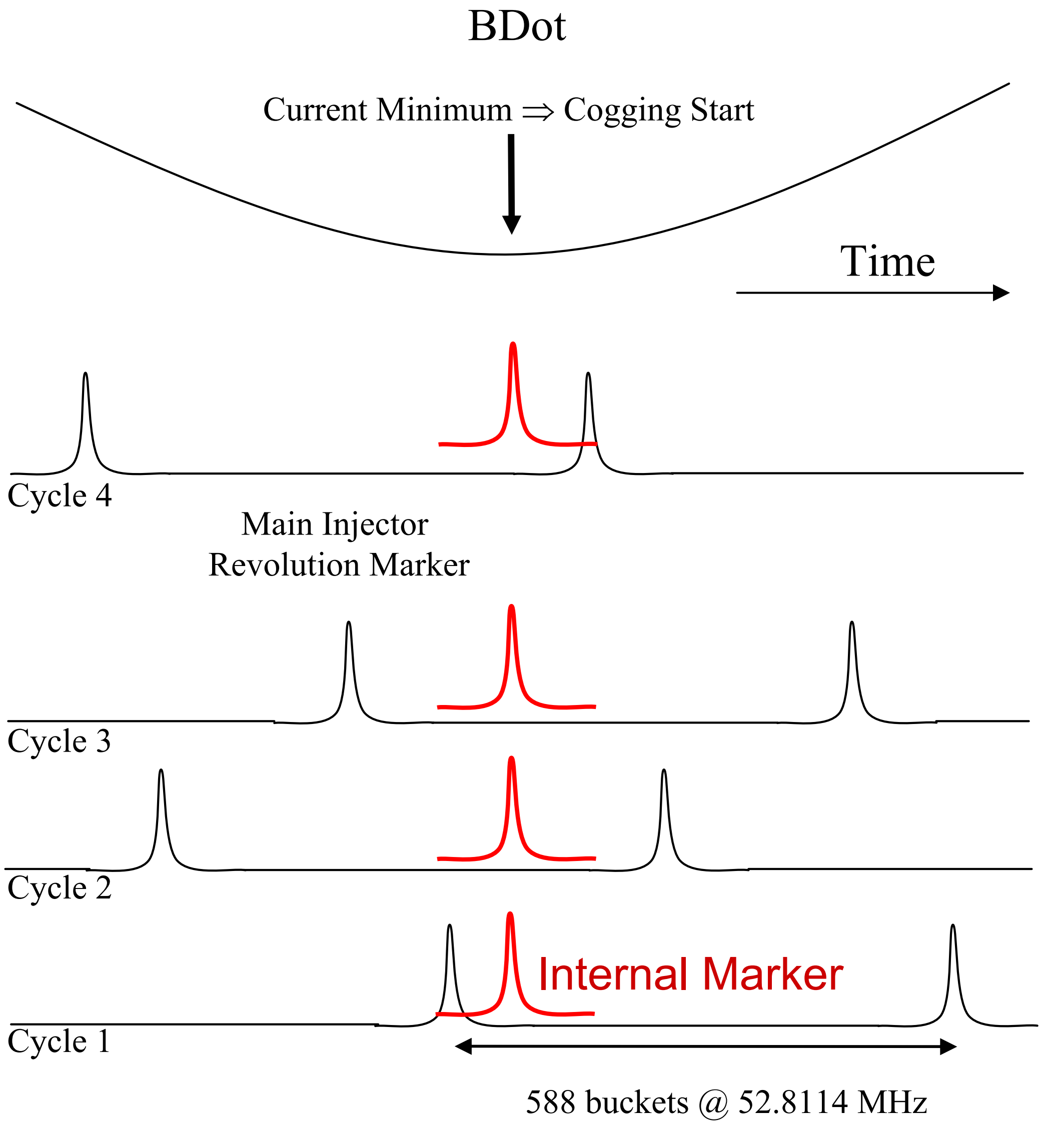

Figure 4.14: Illustration of the timing issue caused by the asynchronous OAA and the remedy. The slippage measurement system starts with the magnet current minimum (BDot), but the OAA signal from the MI can can at any time in relation to BDot. The offset causes a timing slippage of up to 160 buckets (2 revolutions). To remedy this effect, an internal MI revolution marker is kept (fOAA) that is synchronized to the cycle start.

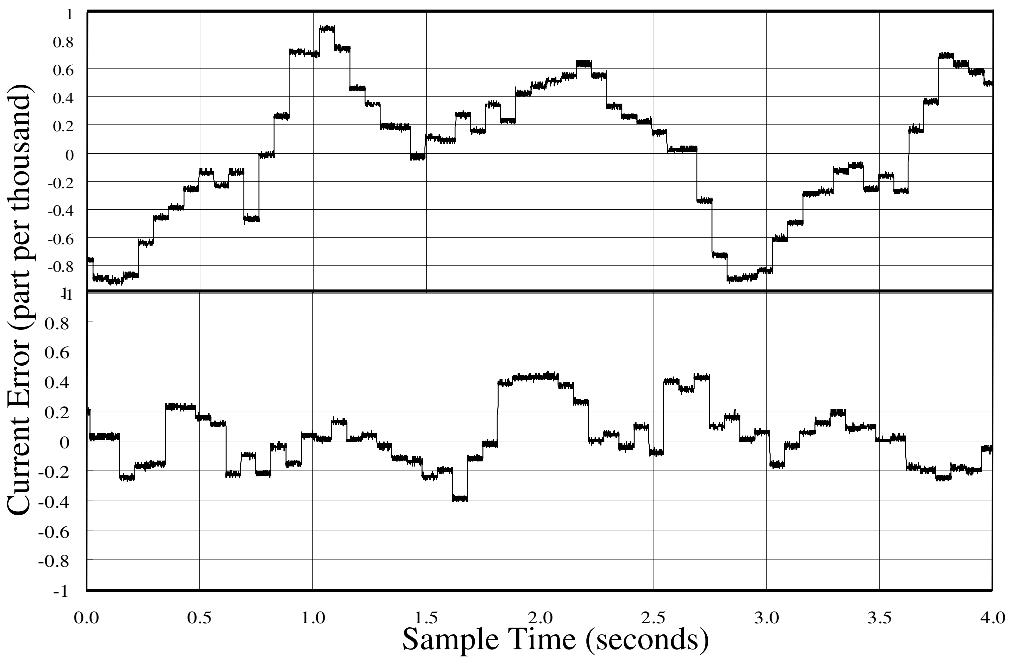

Figure 4.15: Booster minimum magnet current variation measurements. Above: naturally occurring error signal. Below: error signal with feed-forward compensation. The measured variation is reduced by about a factor of two.

4.5.2 Magnet Currents

Measurements of slippage and magnet currents showed that the pulsing of the RF anode supplies was causing the supply voltage to the gradient magnets to sag, causing the injection current to vary pulse-to-pulse.

The gradient magnets have a feedback system through GMPS, but it takes a few pulses to correct the minimum magnet current. In this situation each subsequent cycle has a consistent magnet offset that led to slippage early in the cycle. Measured variation of the injection magnet current is shown in Fig. 4.15. The spikes and troughs were typically correlated with the beginning and of a bunch train.

To combat the variation, a feedforward magnet current compensation system was implemented. The compensation consists of a pulse generated on

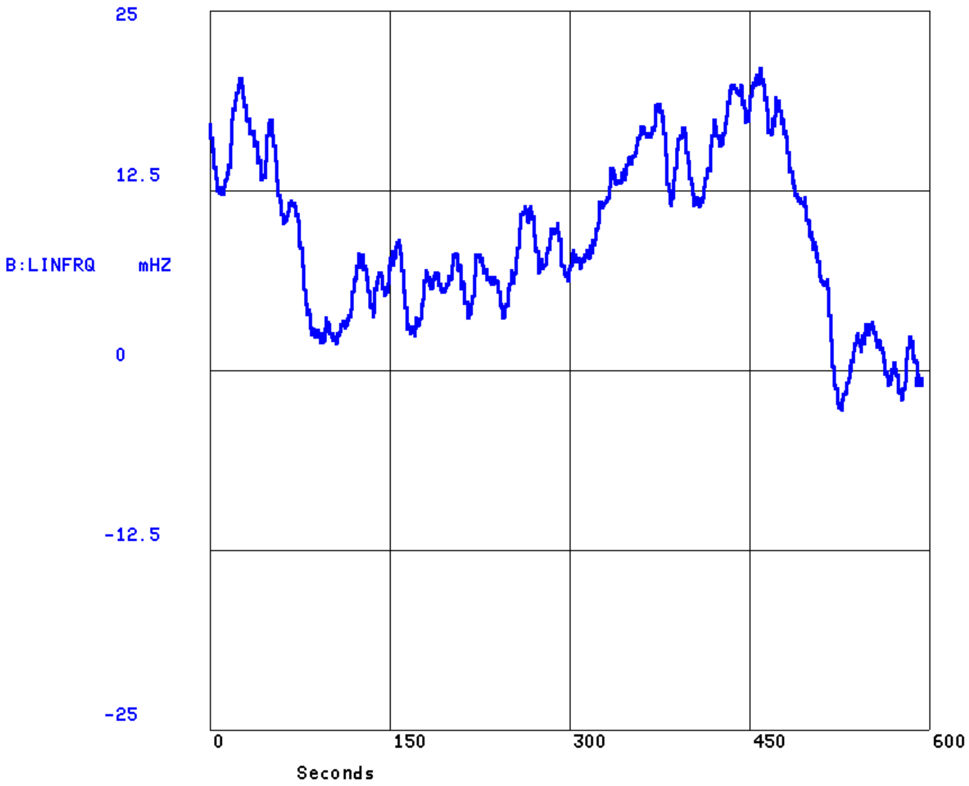

Figure 4.16: Measured line frequency variations from the nominal 60 Hz. The Booster frequency is locked at $1/4^{\text{th}}$ that, so its variations are a quarter of those plotted.

each cycle modulating the magnet current regulation. The pulse is a half-sine wave offsets to the magnet current throughout the cycle, increasing the magnet current at injection[7].

This simple system reduced the variation of the measured minimum current by more than a factor of two, improving the quality of the prediction used to create the notch. The effect of the compensation on current variation is shown in figure 4.15. The effect on measured slippage was less clear: the magnitude of slippage was not reduced substantially, but early swings in the

[7]The maximum magnets currents were found to vary by a similar amount to the minimum currents. However, the relative change was a factor of 10 smaller. Additionally, the slippage expected from such a variation is less. As we could only reduce the injection current variation by a factor of 2, and the extraction current effect was $\sim 15\times$ smaller, we ignored it.

slippage were reduced. A more consistent shape of the slippage curve is useful for the correction system of Ch. 5.

4.5.3 Magnet Frequency

Another potential variation to the Booster momentum curve is variation in the Booster magnet frequency. The Booster's resonant frequency does not change, but the supplied AC wallsocket frequency does change, which also changes the actual frequency of the Booster magnets[8]. The measured variation in the wallsocket frequency over 600 s is shown in Fig. 4.16. The variation over periods of multiple seconds is on the order of 20 mHz, but the fast variation in less than a second is on the order of a few mHz, and is further suppressed by the Q of the Booster's resonant circuit which is ~ 20.

We had no unobtrusive way of forcing the Booster's magnet oscillation frequency, so errors from the variation are considered irreducible. Similarly, there are most likely other sources of variation that could not be easily identified, some of which will remain undiscovered, and perhaps also be irreducible.

[8]Changing the frequency of the input power affects the resonant gain of the system, which could change the minimum and maximum magnet currents. However, the GMPS system provides feedback to correct this variation.

Chapter 5

Booster Synchronization with the Main Injector

5.1 Introduction

In the previous chapter, we established that the distance traveled by the Booster beam during the acceleration period can vary by as much as $\pm$ 200 RF buckets. After refinements in timing and control of the Booster, this variation could be made as small as $\pm$ 70 buckets. However, we require a consistency of about $\pm$ 1 bucket for proper extraction of the Booster beam. This chapter describes a system to force the Booster beam to the correct longitudinal position for extraction.

In accelerators, it is common to use the RF system to radially manipulate the beam to adjust its forward progress, and prepare the beam for transfer to another ring (e.g., see [72, 73, 74]). This procedure is known as "cogging", a designation that calls to mind correcting the orientation of two sets of gears so that their cogs mesh. Typically, the beam is moved radially a distance of

~ 1 mm while at flat top[1] over a period of a few tens of ms. The Fermilab Booster, however, does not have a flat top, because its magnets are part of a resonant circuit and follow a sinusoidal curve. The virtual flattop of the cosine is only a few hundred μs long – insufficient to move the beam even 1 RF bucket's length.

The Booster requires a system that can correct up to 80 bucket's worth of slippage before the time of phase lock to the Main Injector. From the profiles of slippage in Fig. 4.13, we observe that the majority of slippage occurs in the first 10 ms of the cycle. A fast feedback system could try to fix the frequency progression of the Booster by altering the beam's radial position at those times. However, the beam size is comparatively large early in the cycle; moving the bream would bring some portion of it to the limit of the dynamic aperture, where it would be lost. Instead, we will allow the slippage to develop early in the cycle, and correct it later in the cycle. Later in the acceleration cycle the beam has become smaller due to adiabatic damping (§A.2), and can be moved about without exceeding the aperture. Additionally, we will delay creation of the notch by ~ 4 ms so that it can be placed in anticipation of slippage, based on a prediction from measurements of the first few ms slippage.

The cogging system described in this chapter was originally proposed and studied in the 1990s [70], and was further studied and implemented by this author and others [75, 76]. The system builds off of the measurements system described in the previous chapter – we add input and output signals to the various involved systems and develop additional logic to interpret the measurements and control the signals. Ultimately, the cogging system must control the creation of the notch, the radial position during much of the Booster

[1] Flat top is when the accelerator reaches its maximum energy and the magnets stay at a constant current – creating a "flat top" in the momentum profile vs time.

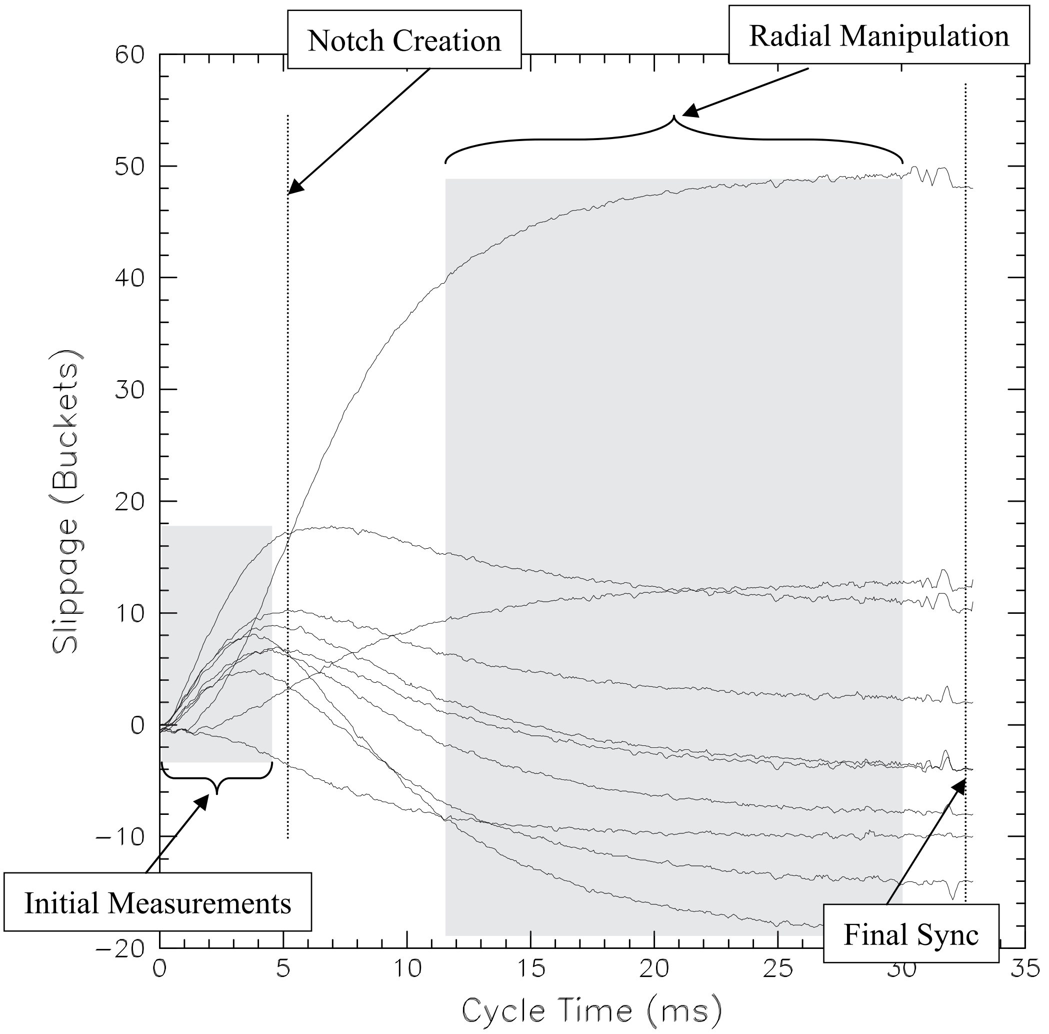

Figure 5.1: Steps of the synchronization process. Slippage is measured for the first 4 ms, after which the notch is created anticipating slippage. Once the beam has shrunk in lateral size, radial manipulations correct the remaining slippage; the last 3 ms are reserved for phase lock. Signaling between accelerators sets the final synchronization at the very end of the cycle.

cycle, and the timing of the Booster extraction and MI injection kickers; the various stages are outlined in Fig. 5.1.

5.2 Predictive Notch Placement

In this section, we describe how delaying the creation of the notch 4 ms into the acceleration cycle[2] can, with measurement of slippage during those 5 ms, significantly reduce the final slippage error. As shown previously, much of the slippage occurs in the early part of the cycle; particularly, much of the cycle-to-cycle variation occurs there as well. The similar shapes of the slippage curves (see Fig. 5.1) allow extrapolation of the final value of slippage from early values, with some precision. By measuring the slippage during the first few ms of the cycle we can extrapolate to the end of the cycle, and place the notch anticipating that further slippage.

Of course, if the notch were made at the end of the cycle it could be placed in precisely the right location. However, that would negate the advantage of notching (over unnotched extraction) as the same amount of beam would be lost at the same energy – leading to the same amount of radioactivation as extraction without a notch [3]. Instead, we balance the decreased energy lost from creating the notch earlier with the reduced predictive power. Additionally, the current notching system of a single kicker magnet cannot kick the beam strongly enough to remove it after $\sim$ 6 ms into the cycle. The time of kick is chosen as late as possible such that the kicker can fully remove the beam; this time is $\sim$ 5 ms after beam is injected, at which time its kinetic

[2]In other circumstances, the notch is made $\sim$ 1 ms into the cycle – after the beam has been fully bunched

[3]Potentially, there could still be some advantage to making such a notch as the loss could be located away from the extraction region. However, that gain is marginal as no location of the Booster can afford such radiation damage and activation.

energy has increased to $\sim$ 700 MeV. As was shown in Fig. 3.15, some beam remains in the extraction notch because of the limited notch kicker charging voltage.

5.2.1 Procedure for Predictive Notch Creation

The logic for choosing where the notch will go must be performed in real time, using the DSP. The time is limited as a slippage measurement occurs every 11 μs and must be processed by the DSP before the next cycle starts. With a few μs of calculation and a 40 MHz DSP, only $\sim$ 200 instructions can be executed as part of the DSP's calculations. To accommodate this, running sums are kept of the slippage measurements; also the extrapolation logic is a simple two (averaged) point extrapolation to the end of the cycle, using a tunable scale factor:

$$P = G(s_2 - s_1) + s_1 \tag{5.1}$$

Where P is the predicted position at the end of the cycle, G is the tunable scale factor, and s_1 and s_2 are the two running averages:

$$\begin{aligned} s_1 &= \sum_{i=210}^{300} s_i \\ s_2 &= \sum_{i=300}^{390} s_i \end{aligned} \tag{5.2}$$

where s_i are the individual slippage measurements; the limits of the summation are revolution numbers whose values were tuned. The running sums act as averages; at least 30 points are necessary to allow fine prediction to the end of the cycle. The scale factor, G, is approximately 6 buckets / bucket.

Once the predicted slippage at the end of the cycle is calculated, the notch is made at the corresponding buckets in the opposite direction of the

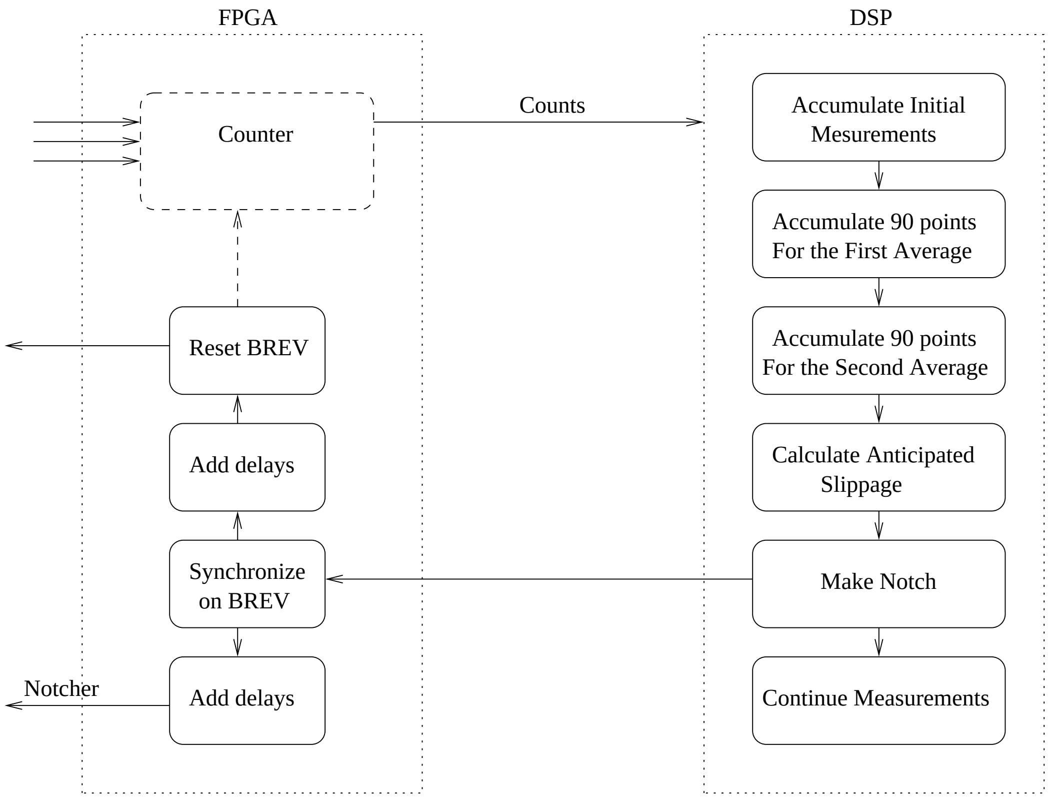

Figure 5.2: Flowchart of the notching process. Counts are generated and processed as in Fig. 4.7. After an initial number of slippage measurements, two sets of ninety consecutive slippage measurements are averaged and used to extrapolate a final slippage. The notch is then made which involves firing the Notcher and resetting the Booster Revolution Marker (BREV), each of which requires synchronization with the previous BREV and different sets of delays.

current Booster revolution marker. The implementation of the logic in the hardware is shown in Fig. 5.2. An initial, arbitrary Booster revolution marker (BREV) is established – as there is no notch, there is no distinguishable characteristic of the beam to reference it to. The counts are created and turned into relative slippage measurement as described in §4.3.2.

The DSP accumulates some number of initial measurements that are not included in the measurement. These are discarded as they are from when the LLRF's radial feedback is still weak and the beam's RF frequency is being set more by the preset program. Then, the DSP accumulates two sets of running sums, which are effectively averages. Then the total anticipated slippage is calculated as described above. The calculation takes less than two μs, and the value of anticipated slippage is sent to the FPGA.

The FPGA then must create signals that trigger the notcher to create the notch, and to internal and external counters that maintain a BREV. The FPGA waits for a BREV, and then adds a number of delays to each signal (measured in RF counts) and creates the pulses. Each signal must wait the number of buckets of the predicted slippage (making the notch later). Each must include an offset due to the OAA – fOAA difference (see §4.5.1); this offset must be measured in MI RF buckets, but be applied as Booster RF buckets. Each must also include an offset for the desired slippage count at the end of the cycle; recall that we are striving for consistency, so this number must be consistent with the previously measured cycle used as a baseline. The notch creation signal must then have a tunable offset applied that accounts for signal propagation time.

The notcher kicker magnet then immediately makes the notch. The internal BREV marker is reset, and the FPGA must allow it time to reset in case of another reset. Additionally, the reset is provided the external BREV counters that might be aware of the notch position, particularly the BES module that will control the timing of the extraction kickers.

5.2.2 Results

The results of the predictive notching are shown in Fig. 5.3. Generally, the prediction is accurate enough to reduce the range of final slippages to $\pm$ 10 buckets at extraction. Tuning of the average ranges and extrapolation scale factor is rarely necessary once proper values have been found. Necessary for this level of error reduction is the GMPS feed-forward compensation described in §4.5.2, and the learning described in §5.5. The injection magnet current variation caused slippage of a different shape than typical, reducing the effectiveness of the prediction (i.e. it had a different optimal scale factor). The learning reduces long-term variations, like magnet frequency variation, that would also have different slippage shapes.

We tested more complicated predictive algorithms, such as adding a third point and trying to analyze the shape of the slippage. However, these gave only minimally superior results to a well-tuned two-point extrapolation like Eqn. (5.1). Furthermore, they all required more processing time in the DSP. When calculations are too lengthy they can cause the subsequent measurement to be missed; which translates into a timing error for the remainder of the cycle, carrying with it artificial slippage that will need correcting. The two-point system was retained as adequate; the remainder of the slippage is corrected by the RF manipulations of §5.3.

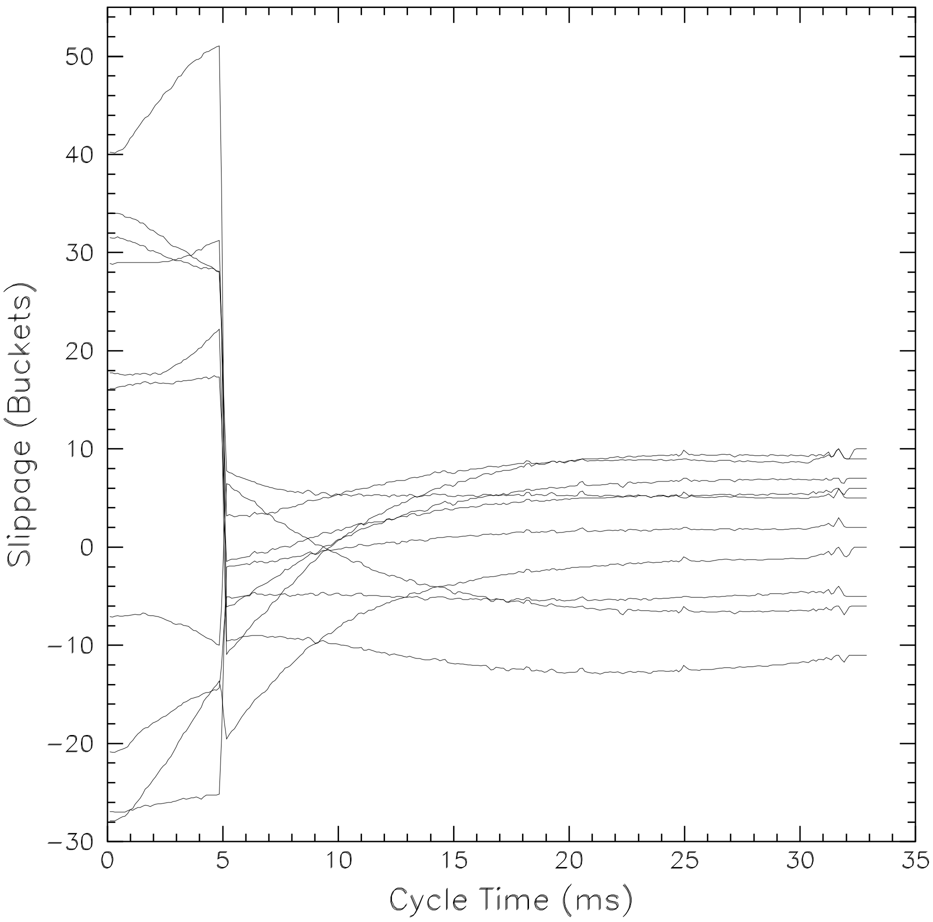

Figure 5.3: Example of ten predictively notched cycles. The initial slippage offset is arbitrary as there is no notch at that point. The slippage develops and is measured for ~ 5 ms; the notch is then made including the current slippage and anticipated slippage. The spread in final value of slippage is reduced from that shown in Fig. 5.1.

5.2.3 Issues with Later Notching

Delayed creation of the notch markedly reduces the inconsistency in the final position of the notch and is an integral part of the cogging system; however, certain compromises must be made to achieve these results. This section discusses the issues of increased beam power lost, increased requirements of the notching system, and restrictions placed on proposed alternative systems for notch creation.

Power Loss

As mentioned above, one of the expected effects of later notching is the greater amount of beam energy lost. The amount of radiation damage and residual activation is roughly proportional to the beam energy lost – not the number of particles. Thus, delaying the notch to 5 ms involves kicking out a beam with $\sim$ 80% higher kinetic energy. Fortunately, the notching losses are deposited primarily in the collimator sections which is designed to absorb significant beam losses. However, the collimator region is a rather radioactive area of the Booster; further activation (e.g. from higher intensity, higher rate beam) may not be acceptable. Loss measurements in the collimator region are dominated by collimation losses and the effect of the notch is not directly apparent.

Notcher Power

Delaying the notch creation to 5 ms after the start of the cycle results in a stiffer beam that requires a greater magnetic field to kick out. At 5 ms, the beam momentum has increased 45 % and a correspondingly larger kick is required to fully create the notch.

The kicker magnet that creates the notch is powered by the discharge

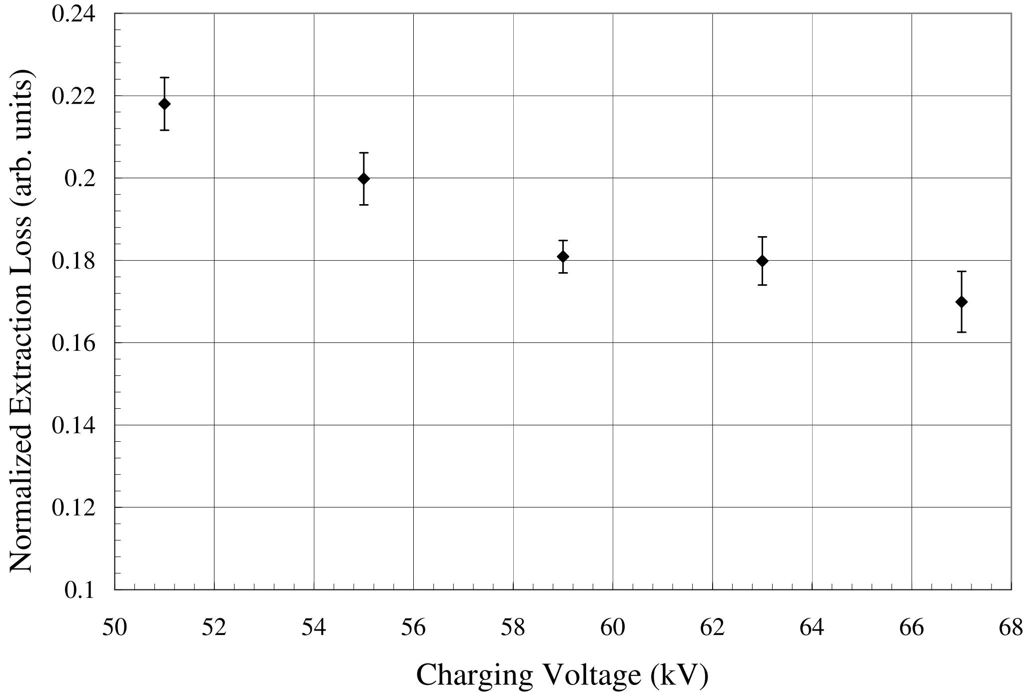

Figure 5.4: Extraction loss per extracted proton as a function of notcher charging voltage. Losses are significantly higher when the charging voltage drops below $\sim$ 58 kV. The error bars reflect cycle-to-cycle variations.

of a high voltage pulse-forming network (PFN). For notching at 400 MeV, a charging voltage of 43 kV was typically sufficient, suggesting that this implementation of cogging requires as much as 70 kV. The existing supply could not operate reliably in excess of 57 kV.

Fig. 3.15 showed the results of creating a notch 5 ms into the Booster cycle at a charging voltage of only 53 kV. The notch has an effect four buckets wide – with two bunches only reduced by $\sim$ half, and the two central bunches being reduced by $\sim$ 95%. A higher charging voltage achieves a more square and deeper notch.

Fig. 5.4 shows measurements of extraction losses with cogging as a function of charging voltage with a power supply newly purchased to address the problems of the delayed notch. The new supply capable of a maximum of 70 kV, but is currently run at $\sim$ 63 kV as its behavior is unstable at higher charging voltages. The Booster department is considering replacement of the notcher magnet with an electrostatic kicking system that should be capable of equal or greater kicks.

Alternative Notching Methods

The Proton Source Department had studied several schemes to create the notch previous to injection of beam into the Booster – removing the notching loss altogether [77, 78, 79]. These schemes involve creating a notch near the begininning of the Linac where the energy involved will be far less. However, to cog cycles the notch must be places anticipating further slippage, which is impossible that early in the cycle.

Predicitive notch placement generally precludes Linac notching for cogged cycles. However, the Linac notching could still be applied on the first cycle of a train, which is not cogged and used as a baseline, as well as any other uncogged cycles used in the Booster.

5.3 Radial Manipulations

After the placement of the notch we will correct any slippage using radial manipulation of the beam through the Booster LLRF. As explored in §4.2.2, a change in the radial position of the beam (RPOS) will involve a change of momentum and a change in circumference. At any time in the cycle, other than transition, these changes will change the revolution frequency of the beam and induce slippage. We will intentionally use the Booster LLRF to change the momentum of the beam, changing its radial position, and correcting any slippage as the Booster cycle progresses.

In this section: we discuss the available horizontal aperture in which the beam can be moved; the implementation of the radial manipulation logic in the DSP of the measurement system; various manipulation algorithms with their results; and some of the issues encountered when implementing the radial manipulations system for cogging.

5.3.1 Available Horizontal Aperture

Changing RPOS at the very beginning of the Booster cycle is not possible as the beam fills the horizontal aperture of the Booster – any change in the beam's centroid would involve pushing some of the beam's tails out of the aperture where they would be lost. Instead of losing that beam, we will wait until the beam has shrunk through adiabatic damping. The fine-tuning of the beam's longitudinal progress will occur after transition when the beam has shrunk significantly and the nominal slippage has mostly developed.

In anticipation of our post-transition cogging, we performed a series of manual manipulations on the beam to gauge the range of allowable radial change. For these tests, we used a study cycle with an already tuned RPOS trajectory that we treated as nominal. Fig. 5.5 shows the radial trajectories followed in two of the tests. We found that at low intensity ($\sim 10^{12}$ protons) the beam could be moved $\pm$ 8 mm after transition without beam loss, or other apparent negative effect.

The second test shown in Fig. 5.5 involved placing as much beam into the Booster as it would hold. Necessarily, there was substantial beam loss throughout the cycle; however, we only saw marginal beam loss at the most extreme RPOS settings. A closer look with the Ionization Profile Monitor

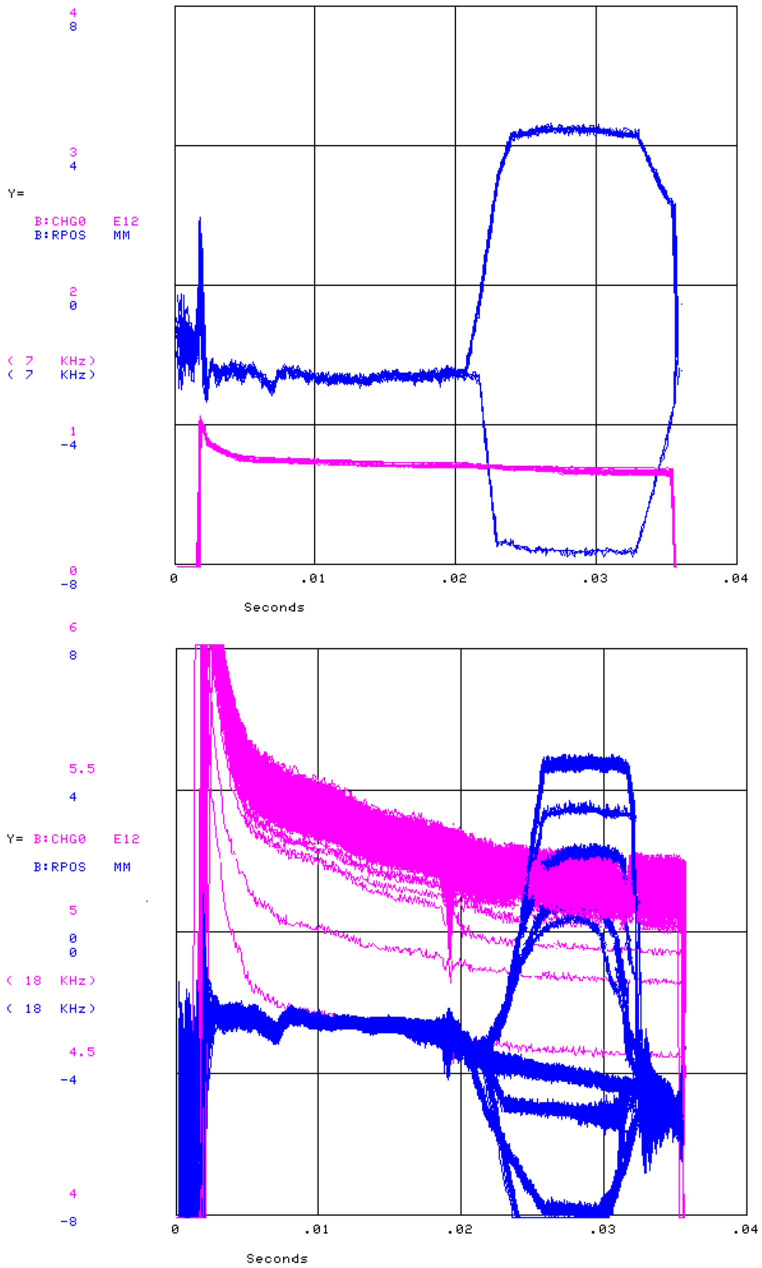

Figure 5.5: Exploration of horizontal aperture with manual radial manipulations. Each shows the the beam intensity as a function of time, and the radial position (RPOS) as it was programmed to go. The upper is a low-intensity test with high efficiency; the lower is a very high-intensity test with poor efficiency throughout. While the intrinsic loss dominates the lower plot, greater loss was experienced at the greatest radial excursions.

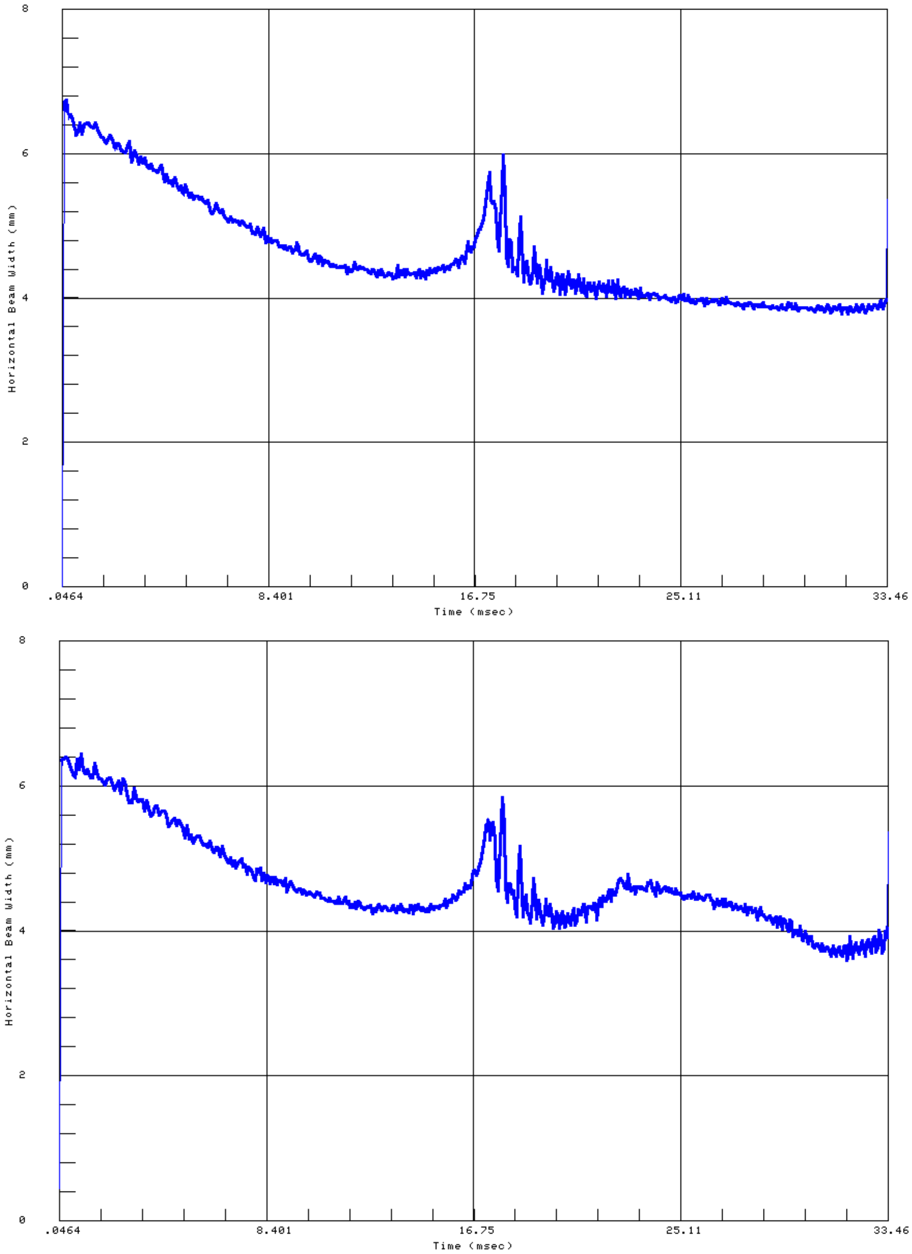

Figure 5.6: Measurements of the horizontal beam width during the Booster cycle. These two correspond to the high-intensity cycles of Fig. 5.5. The upper is for the nominal radial position throughout the cycle, and the lower is for an 8 mm manual radial offset. While the radial offset resulted in little additional beam loss, it did contribute to growth of the beam width.

(Fig. 5.6), however, showed that the beam width increased at precisely the time that the greater radial offsets were applied.

We eventually chose $\pm$ 5 mm as being the maximum radial excursion we could withstand, though we would ultimately need less of an offset.

5.3.2 Procedure for Radial Manipulations

The logic for the radial manipulations is performed entirely within the DSP, given the measurements described previously. The DSP effects radial motion by outputting a value to a DAC, which is input to the LLRF as an offset to the programmed radial offset.

Fig. 5.7 shows the basic outline for the DSP logic. For each measurement, the DSP reviews its output to the LLRF and modifies it accordingly. There is a wide variety possible of radial manipulation algorithms, several of which are discussed below, but their inputs and outputs are essentially the same. The DSP will analyze the current error measurement in terms of previous measurements, the time in the acceleration, the current RPOS offset, or the results of a prediction akin to that used for the notch. The DSP will then choose to increase, decrease, or maintain the level of RPOS offset, and output that value to the LLRF.

The RPOS value is a 16-bit word, translated by the DAC into a $\pm$ 10 V scale. The LLRF interprets the voltage as a further offset to the programmed ROFF curve: the LLRF will attempt to maintain the radial position of the beam as the programmed curve *plus* the modified value from the cogging system. The LLRF system has intrinsic bandwidth, limiting the speed of feedback to ~ 5 μs. However, the synchrotron frequency of the beam is usually in the range of 1-10 kHz, so a fast offset shift to the LLRF may result in

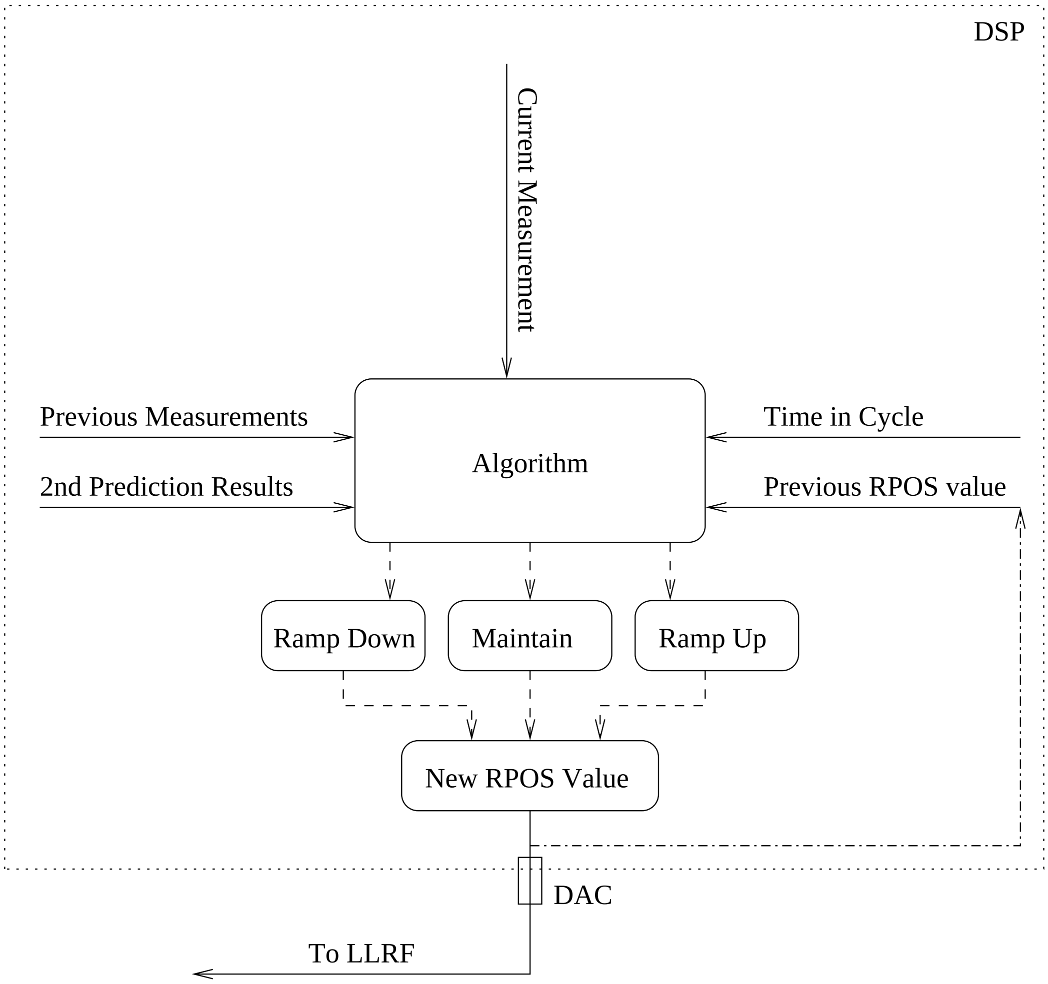

Figure 5.7: Flowchart of the logic for radial manipulations. The logic takes place entirely within the DSP. After each measurement, an algorithm is employed that uses the measurement, previous measurements, the results of a prediction, the time in the cycle, and the current RPOS offset to decide whether to maintain the RPOS offset, or ramp it up or down. The RPOS offset is converted to a voltage signal that is interpreted by the Booster LLRF.

jerking the beam longitudinally, and increasing its emittance. As such, the radial manipulations of the cogging system are fed through a 10 kHz low-pass filter before reaching the LLRF.

5.3.3 Algorithms

The digital nature of the cogging systems allows great flexibility in designing synchronization algorithms, such that the system can apply non-traditional methods of feedback and feedforward manipulation. During the design and testing phases of the cogging system, several different algorithms were tested. Here, we describe the three notable iterations: simple proportional feedback; "flat" feedback where the amplitude of the feedback offset is kept large to decrease the damping time; and the ultimate system which employed portions of flat and proportional feedback, as well as a pre-transition flat feedforward and a gain doubling.

Proportional Feedback

The initial manipulation method implemented in [70] was to initiate a radial offset to the LLRF after transition that would move the beam slippage toward zero. The radial offset was proportional to the size of the error and an adjustable gain. The DSP generated a new offset with each measurement, such that the offsets decreased as the slippage error decreased. Fig. 5.8 shows the results of such proportional feedback with a gain setting of ~ 0.15 mm / bucket. The rate of slippage induced is proportional to the offset and the cogging power as defined in Eqn. (4.22). The dispersion at the BPM used for radial positions measurements is 1.7 m. Roughly accounting for the variation in $d\tilde{\eta}/dx$, the damping rate results in an e-folding time of ~ 9 ms (halving

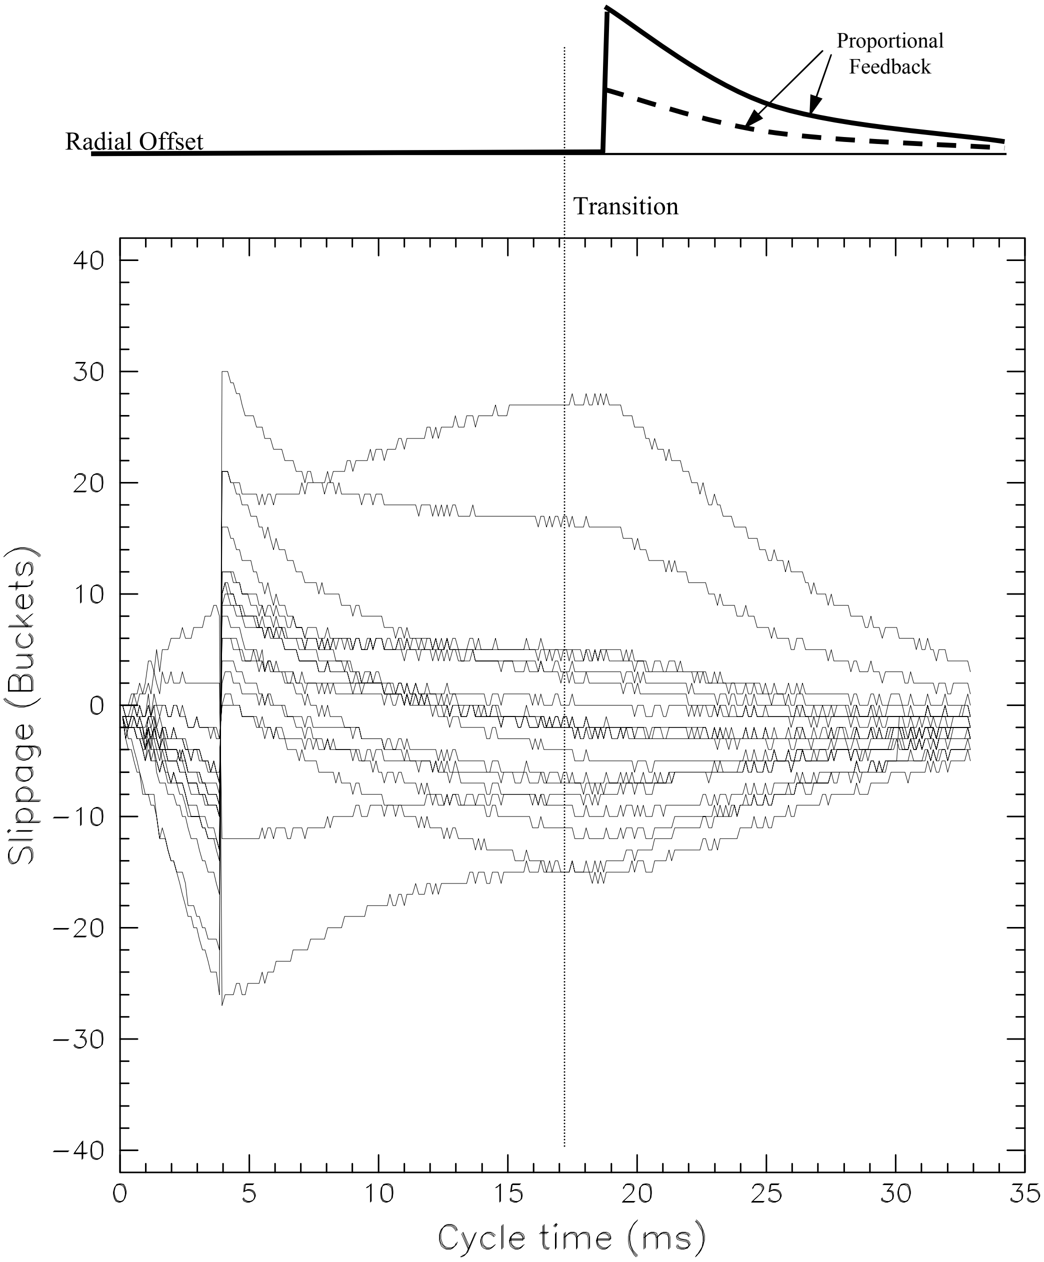

Figure 5.8: Schematic of the proportional feedback algorithm, and measured data from cycles subjected to it. The notch is performed at 4 ms; the radial feedback starts after transition, and the induced change in RPOS is proportional to the slippage error. For the largest offset at transition (shown) of 27 buckets, the cogging offset to RPOS is $\sim$ 4 mm, resulting in a slippage of $\sim$ 2.4 buckets / ms, and an e-folding time of $\sim$ 9 ms. The final slippage values are spread over a range of $\pm$ 5 buckets.

time of $\sim$ 6 ms).

In the figure, the largest offset is $\sim$ 4 mm for the 27 bucket error, and it reaches an ultimate error of 3-4 buckets at the end of the cycle. Generally, the final range of errors was $\pm$ 5 buckets. This error is too large for synchronization. A higher feedback proportionality would results in radial offsets at 20 ms that would result in either beam loss or emittance growth.

Flat Feedback

A method of "flat" feedback was devised where the DSP chose a radial offset to correct an error, and then maintain that offset until the slippage error had been reduced to a small value. Once that small value had been reached proportional feedback would again be applied, but at a higher gain than possible before. Additionally, logic was implemented that would ramp the RPOS offset over the course of 0.5 ms, instead of changing it immediately. We had found that rapid changes in the RPOS offset could cause beam loss or emittance growth.

Fig. 5.9 shows the results of a flat cogging algorithm. In this iteration, offsets larger than 10 buckets at transition were given a 4 mm offset; those with an error of 5-10 buckets received a 2 mm offset; and those with smaller errors went into proportional offset with a proportionality constant twice that of the algorithm used in Fig. 5.8. Those in flat feedback had their offsets maintained until the error fell to two buckets, at which time they also entered proportional feedback.

The final range of offsets was $\pm$ 1 bucket. We believe the nonzero range of final errors is caused by cycle variation effects like the variability in the maximum magnet currents (see §4.2.1). With an error of 1 bucket, the proportional feedback was 0.3 mm; the induced slippage from such an

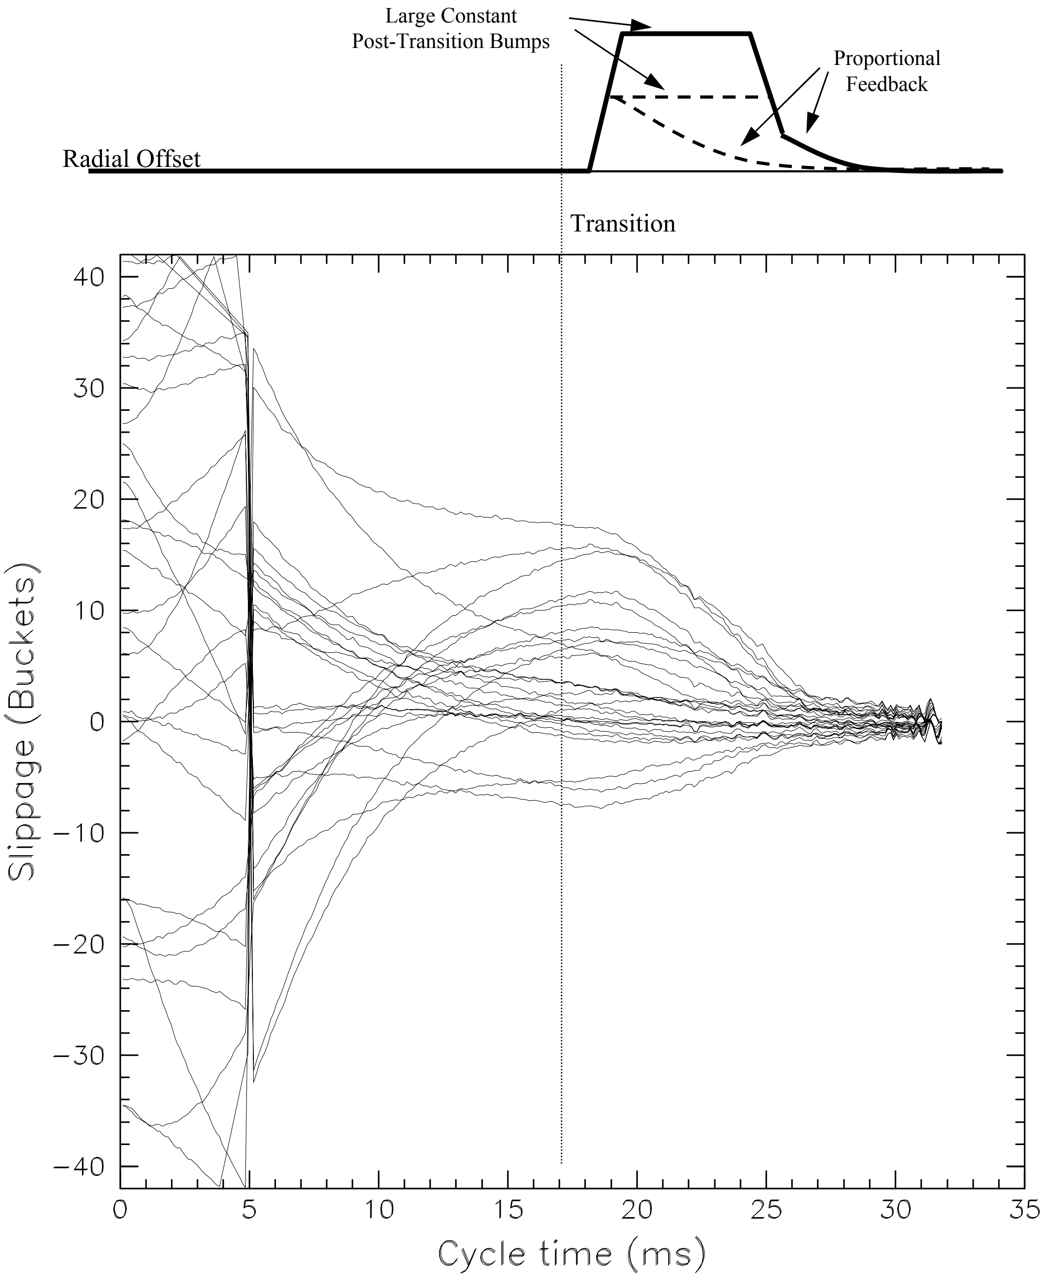

Figure 5.9: Schematic of the flat feedback algorithm and along with measured data. The notch is performed at 5 ms; the radial feedback starts after transition. The feedback magnitude is greater for greater errors, but stays large until the error is small, at which point proportional feedback takes over, but with a larger proportionality constant. The largest RPOS bump here is 4 mm, and the values approach a final range of ± 1 bucket.

offset could be entirely canceled by a 2 parts in 10,000 variation in the final magnet current, such that the slippage error would remain at 1 bucket. These variations were eventually reduced by the GMPS compensation (§4.5.2), and further dealt by implementation of a gain doubling in the ultimate algorithm, whereby errors of 1 bucket were subject to twice the proportional feedback they would otherwise receive.

While the flat feedback could reliably reduce slippage errors to a small range that could be accommodated, it required regular excursions of $\pm$ 4 mm in the radial feedback ($\sim$ 10% of the cycles received the larger feedback). These excursions could be withstood by careful tuning of the Booster, but imposed a narrow tuning range on RPOS. We were also uncertain of whether such a range of movement could be sustained at higher intensity and still not increase beam loss or result in emittance dilution.

Final Cogging Algorithm

The final radial manipulation algorithm is shown in Fig. 5.10, along with the data from such an algorithm. Fig. 5.11 shows the actual RPOS measurements taken during the above manipulations. The cogging system is continually tweaked by Booster operations, but the radial manipulation algorithm has been mostly unchanged since being implemented in normal operations. Data from such operations are shown in Fig. 5.10. The substantial change from previous algorithms is the implementation of small pre-transition bumps to correct slippage error.

Pre-transition bumps had been previously avoided as the beam is still slipping at that point, and is considerably larger in horizontal size. However, the cogging power is somewhat larger before transition (see Fig. 4.5). Using

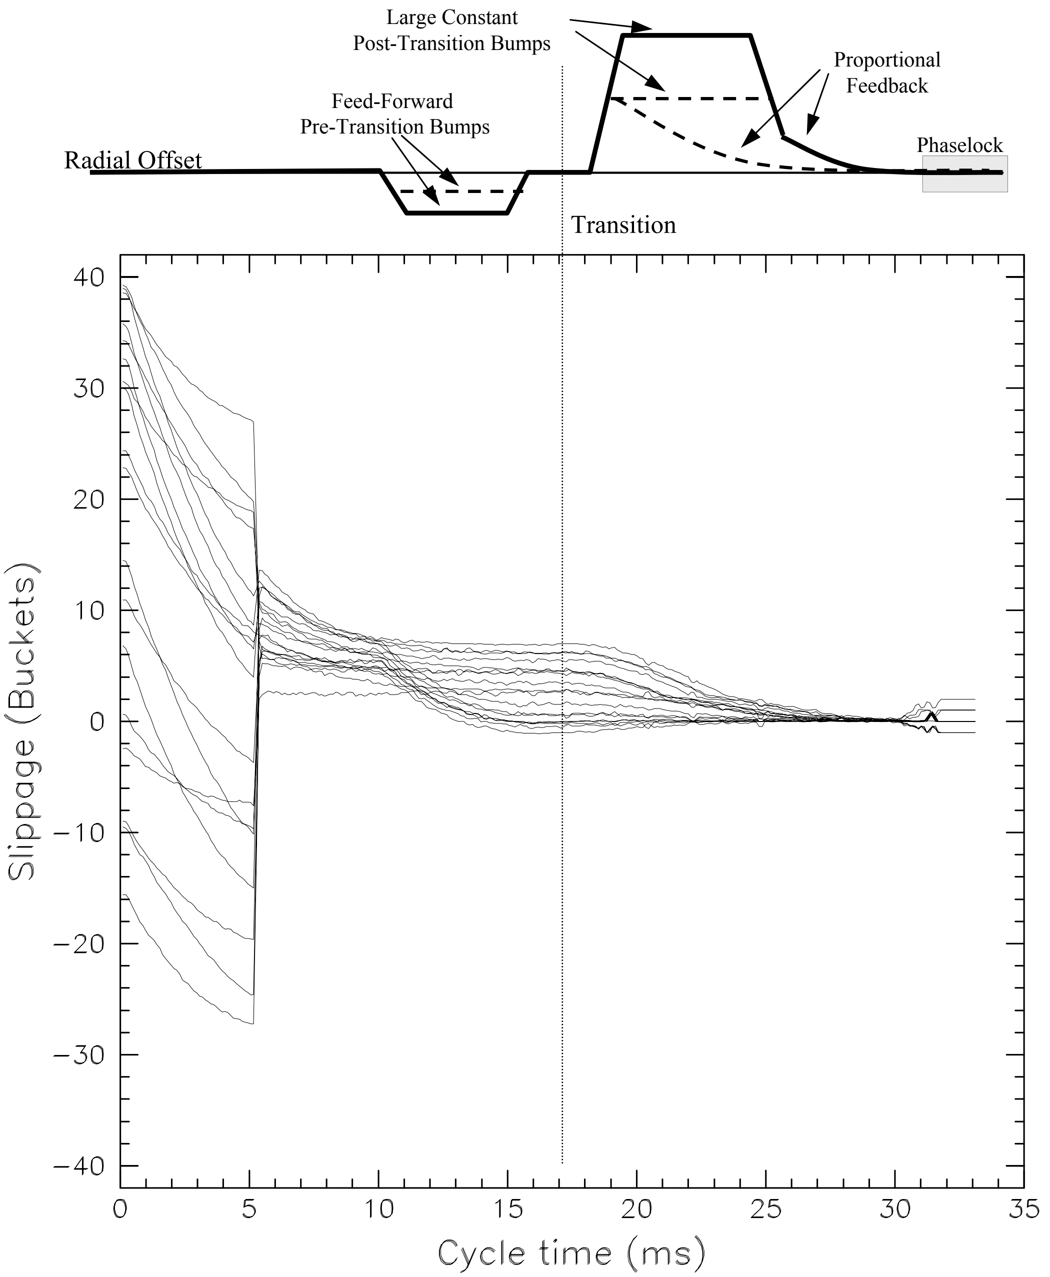

Figure 5.10: Schematic of the full cogging algorithm along with measured data. The notch is performed at 5 ms; followed by pre-transition radial bumps, if necessary; after transition there are larger radial bumps, if necessary; or proportional feedback for smaller errors. The largest cogging bump in RPOS shown is ± 1.7 mm (it can be as large as ± 3.4 mm for large errors, but rarely is). The relative slippage is entirely corrected by 29 ms, however phase lock can pull the beam 1 or 2 buckets away in the final few ms.

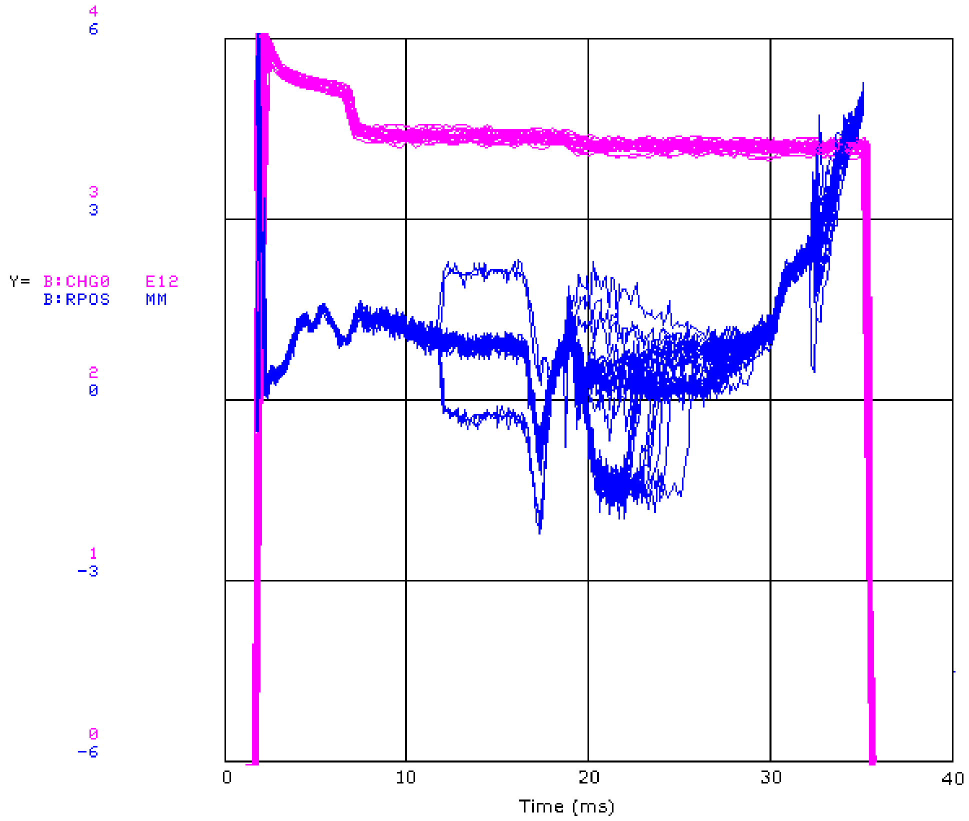

Figure 5.11: Measured radial positions on a series of cogged Booster cycles. The bands of pre-transition cogging are clearly separated, and the post-transition cogging shows more variation. These data can be compared to the uncogged measurements of Fig. 3.7.

a prediction system similar to that for placement of the notch (see §5.2), small constant-amplitude pre-transition bumps were implemented to reduce the range of errors at transition, and thus reduce the maximum excursion necessary.

In the final system, there are also two levels of pre-transition cogging: for predicted errors larger than 15 buckets they receive a 2 mm bump for $\sim$ 7 ms, and for predicted errors of 4-15 buckets they receive a 1 mm bump. Predicted errors smaller than 4 buckets receive no pre-transition feedforward bump. The larger bump is only required $\sim 3\%$ of the time – usually corresponding to unusual tuning conditions in the Booster. Examples of beam cycles with the smaller or no pre-transition bump are visible at $\sim$ 10 ms into the cycle in Fig. 5.10.

Reducing the error at transition by the pre-transition bumps also reduced the amount of post-transition feedback necessary. The size of the two levels of flat offset are 1.7 and 3.4 mm, and the larger offset is used less than 1% of the time.

The slippage error is reliably reduced to zero towards the end of the Booster cycle. However, as shown in Fig. 5.10 and discussed below, the RF synchronization of phase lock (§3.7.3 [69]) occasionally shifts the slippage by up to two buckets. With the phase lock complication, about 2% of the cycles conclude with an error of 2 buckets, and another 6% have an error of 1 bucket.

5.3.4 Issues with Radial Manipulation

While the cogging system is mostly successful, and has been operating adequately since the start of 2005, there are several issues that could confound further improvement in the Booster. Here we discuss the usage of horizontal aperture and the effects of the phase lock system on ultimate cogging performance.

Aperture Usage

The Booster continues to improve its performance, accelerating greater numbers of protons with greater frequency, and is anticipated to further improve in the next few years [80]. Intensity increases in the Booster imply a larger beam, and a correspondingly greater usage of horizontal aperture. Cogging at these higher intensities may start to result in some beam loss.

The cogging system currently uses about $\pm$ 2 mm previous to transition, and $\pm$ 3.4 mm after transition[4], though typical variation is only half that. Further reductions in the radial offsets used are not trivial, but could come from more precise control of the Booster's magnet cycle (e.g., through better GMPS compensation) or through some gross use of the new Booster correction system [66].

Phase Lock

As shown in Fig. 5.10, the phase lock system for synchronous transfer occasionally undermines the synchronization performed by the cogging. The phase lock system (see §3.7.3) aligns the RF phase of the Booster and Main Injector, and is necessary for the beam to be captured at all by the RF in the MI. Cogging is necessarily disabled as the phase lock system engages.

Phase lock is a legacy system originally designed for synchronous transfer for the Booster to the Main Ring. It is composed of analog electronics

[4]These position variations are measured at the RPOS detector where dispersion is calculated as being 1.7 m. So these radial offsets correspond to $\Delta p/p$ of about $\pm$ 0.1% and 0.2%.

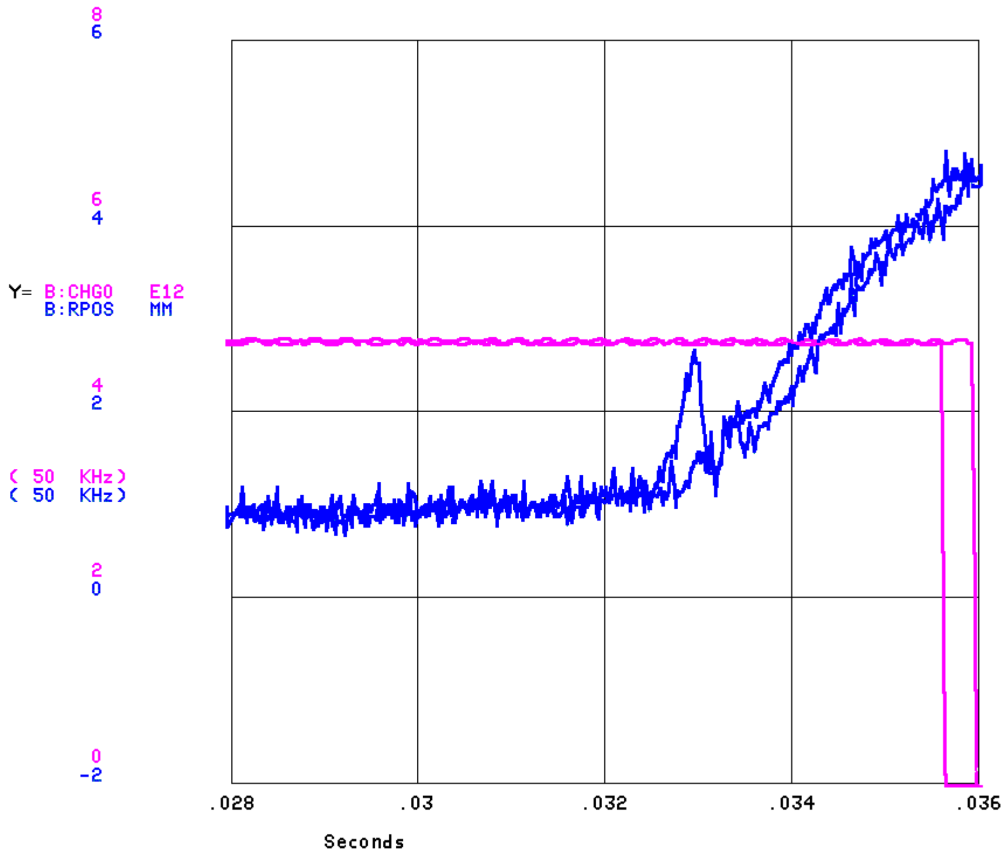

Figure 5.12: Example of the radial movement caused by phase lock. Shown are the charge and radial position traces of the last quarter of two subsequent Booster cycles. The 1 ms, 2 mm bump at 33 ms is an example if the kind of position variation that can be caused by phase lock. This bump is sufficient to change the slippage of the beam by ~ 1 bucket.

and has very available adjustments. The action of the phase lock system is demonstrated in Fig. 5.12. Several second before extraction the phase lock system imposes a radial bump on the beam that matched its frequency and phase to that of the MI RF.

In principle, phase lock should never have to effect a phase shift of more than 180°. However, for no (as yet) understood reason, it occasionally moves the beam by as much as 2 entire buckets ($\sim 720°$). No adjustments to the phase lock operation have been able to reduce the variation. Additionally, the large and rapid radial adjustments performed by phase lock can contribute to beam loss.

The long-term plan for the Booster includes a digital replacement of the phase lock system. Such a system should be more robust, as well as better understood. Potentially, cogging information could also be provided to the new phase lock system so it would perform a fine-tuned notch synchronization, as well as RF phase synchronization.

5.4 Final Synchronization

Once cogging has reduced the beam's slippage, it must then be transferred to the Main Injector. The cogging program controls extraction timing in combination with the Main Injector, providing the Booster extraction system and the MI with information as to the position of the notch and the quality of the cogging. Implementation of an "override" has allowed clean extraction even with a final slippage error of 2 buckets.

The default extraction method for cogged cycles is to extract on the OAA marker, like the RF synchronous transfer described in §3.7.5. If the slippage error is reduced to zero, then the notch should be synchronized with

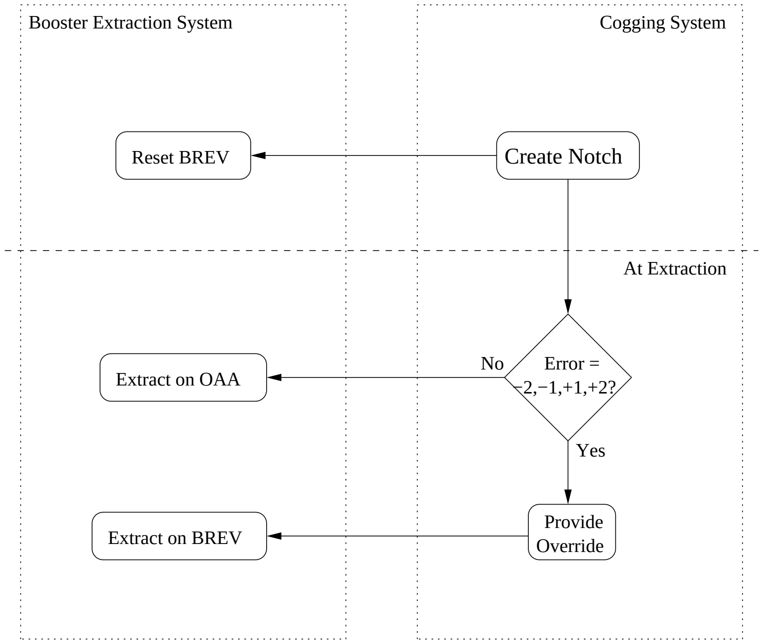

Figure 5.13: Cogging signaling for extraction. When the cogging system creates the notch early in the acceleration cycle it also resets the Booster revolution marker (BREV), so that it corresponds to the notch. At the time of extraction, cogging checks if the error is small, but nonzero. If the error is large or zero then the beam is extracted on the OAA from the MI; otherwise, if the error is small, but nonzero it is extracted on the BREV.

the OAA and extraction should occur on the notch and to the correct bucket in the Main Injector. However, when the slippage error is nonzero then the notch will miss the extraction point, some beam will be lost, but the remaining beam will still end up in the correct location in the MI.

As it became clear that the phase lock issue could not be quickly resolved, we implemented a system of overriding the OAA extraction. The logic and signaling is illustrated in Fig. 5.13: when the notch is created a signal concurrently resets the BREV so that it can be used to extract the beam on the notch. At the very end of the cycle, the cogging system checks the slippage error to see if it has varied from zero, but is still small. If the error is nonzero, a override signal is generated that forces the Booster to extract on the BREV instead of the OAA. The beam is thus always extracted on the notch, but its position in the MI can vary a small bit.

The override tolerance was originally set at $\pm$ 1 bucket, so that the beam in the MI would never be out of position by more than one bucket. However, the NuMI cycle eventually resulted in a batch spacing of 86 buckets, so that there was enough room to override extraction on $\pm$ 2 buckets, and cover the entire range of slippage error caused by phase lock. The precise position of each batch can then vary by $\pm$ 2 buckets in the MI, as illustrated in Fig. 5.14. In the unlikely case that a -2 cycle follows a +2 cycle, a small amount of beam will be lost in the MI; this coincidence has not been directly observed, and is estimated to occur on less than 1 in every 1000 transfers. Otherwise, the $\pm$ 2 bucket variation in batch position does not affect MI operation.

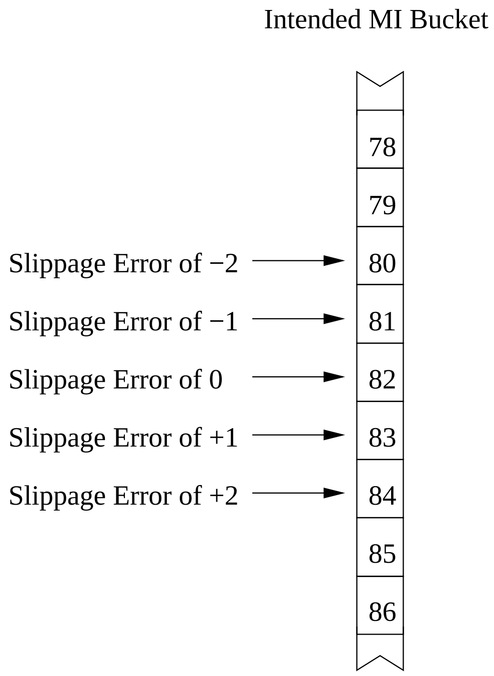

Figure 5.14: Example of the override procedure. If the Main Injector requests beam be delivered to bucket number 82, and the final slippage error is $\leq$ 2 buckets, then cogging fires the kickers offset by that error, and the beam position in the MI is off by 1 or 2 buckets.

5.5 The Operational System

In operation, the cogging system works for all types of Booster cycles transferred to the MI, without intervention. To do so, the first Booster batch in a multi-batch train is uncogged. Instead, the first batch is transferred to MI like in single-batch operation as described in §3.7.5. The beam is notched at the very beginning of the cycle, and the MI revolution marker is reset upon Booster extraction, such that no loss occurs.

The first batch is measured and recoded by the cogging system and is used as a slippage baseline for subsequent, cogged cycles. Taking a new baseline eliminates sources of slippage that are slowly varying, such as those resulting from magnet frequency variation and tuning. Using the first batch as a baseline requires that the Main Injector be at its injection frequency before the first batch[5], and also that cogging system be able to adapt to an arbitrary final slippage[6].

Each subsequent batch is cogged as described earlier in this chapter. Typically, seven batches make up a combined-cycle to the MI: two for stacking and five for NuMI. However, this operation is often interlaced with dedicated NuMI cycles with six Booster batches. There is also a combined cycle for slow-extraction that must have three batches. The cogging systems is agnostic to these different types of cycles as long as the first Booster batch is of the stacking or NuMI type.

The remainder of this section describes the performance of the cogging system since implementation. Particularly, we discuss the reduction of extraction losses, the induced radial motion of the beam, and the effects upon beam quality. In each case, we discuss the possible implications of higher intensity beam.

5.5.1 Reduction of Extraction Losses

The cogging system, along with the override extraction, reduces extraction losses by $\sim 90\%$ – near the maximum possible. As mentioned above, the cogging system reliably reduced the slippage error toward the end of the cycle; however, phase lock can move the beam 1 or two buckets in the final few ms such that only 90% of the batches finish with an error of 0. With the

[5] At one time, it was discovered that the MI changed its RF frequency before beam was injected, but while it was being accelerated in the Booster. As the subsequent Booster cycles did not have a changing MI RF, it manifested itself as a slippage that was difficult to remove.

[6] The first batch is not constrained in any way as to what its final slippage will be. Therefore, subsequent batches will be cogged to the same arbitrary slippage, so there must be offsets introduced to the notch creation and extraction to compensate.

implementation of the override system, the extraction is moved such that it is coincident with the notch on more than 99% of batches.

The measured extraction losses define the figure of merit for extraction. 98-99% show the reduced loss characteristic of notched extraction. The remaining 1-2% are either the victims of unusual Booster operation (such as widely varying wallsocket voltage or frequency) or the accumulation of fractional bucket errors from the various timings.

The precise values of loss vary with Booster tuning and intensity, but from a period of typical running: for beam of intensity 4.5×10^{12} protons, notched extraction resulted in a loss of 1.0 Rad/s, and unnotched resulted in a loss of 8.5 Rad/s. In this period, $\sim$ 98% of the cogged cycles showed reduced losses, for a total reduction of integrated extraction losses of 86%, where a reduction of 88% would result from perfect cogging.

Extension to Higher Beam Intensity

The ability to cog beam is based on the linear optics of the Booster and should not be dependent on beam intensity. However, the proportional reduction of beam losses will vary with intensity, as the unnotched losses do.

As described in §3.7.2, the aperture of the extraction septum magnet is very narrow compared with that of the area for circulating beam in the Booster. As such, any large amplitude beam that was created during the acceleration cycle may be collimated off in the septum magnet. The process of large amplitude beam growth is not well described in the Booster, but it is known to occur nonlinearly with intensity.

As an example of extraction loss nonlinearity: in the example above, beam extracted (properly) on the notch results in a measured loss of 1.0 Rad/s, when its intensity of 4.5×10^{12}. At a lower intensity of 3.5×10^{12} the loss might only be 0.3 Rad/s. In such cases cogging has reduced losses by up to 95%.

The growth of this lost halo beam is probably related to instabilities and resonances in the Booster. New damper systems are intended to reduce some of these instabilities. Additionally, the losses are greater and more variable on batches that not cogged, but are phase locked to the MI. Probably, the phase lock bump is happening quickly enough to generate some large amplitude beam; eventual replacement of the phase lock system may ameliorate this effect.

The end result is that extraction losses can be anticipated to become worse and more of an issue with higher Booster batch intensities. However, these increased losses will be caused by large amplitude beam around the whole of the Booster, and not just in the notch. As such, reduction of extraction losses will become more important, but the reduction from cogging will become relatively smaller.

5.5.2 Beam Motion

The implemented cogging system swings the beam before and after transition, reducing the tuning range available. The pre-transition radial offsets occur between 10 and 16 ms in the cycle, and are usually 1 mm in size (as measured at RPOS), and occasionally 2 mm. After transition, typical bumps are 1.7 mm, sometimes 3.4 mm; they start shortly after transition, and are mostly damped out by the time of phase lock.

The role of the pre-transition bump is reduce that necessary after transition. The algorithm could be tuned to reduce one type of motion, at the expense of the other.

Operational experience has suggested that the Booster runs slightly less efficiently with cogging than without. Particularly, high-intensity MiniBooNE cycles (which are not cogged) are often tuned to a somewhat different default RPOS trajectory than the cogged cycles. This indicates that the optimal RPOS trajectory is usually not in the center of the safe range for RPOS variation.

At higher beam intensity, cogging should not require radial movement in greater magnitude or extent than at lower intensities. However, the beam will be larger and a greater amount of the aperture used. As the Booster achieves higher intensity, the balance of cogging periods and magnitudes may have to be examined for more optimal values.

5.5.3 Beam Quality

The possible effects of cogging's radial manipulations on beam quality was a concern from the outset. Rapid radial shifts in position have the capacity to dilute emittance space longitudinally. Furthermore, beam may be directed into a part of the Booster's dynamic aperture that is nonlinear, or non-optimal in some other way, potentially leading to transverse emittance growth. For examples of beam size growth see §5.3.1.

Mindful of the above possibilities, cogging was optimized to reduce the radial manipulations within the ultimate cogging system. Additionally, radial manipulations were ramped over multiple synchrotron periods, where possible, and the radial offset signal further passed through a 10 kHz low-pass filter preceding the LLRF.

For the most part, no variation in beam quality has been noticed with the cogged beam. The slip-stacking in the Main Injector project had difficulty

Beam Type	**Cogged?**	**Intensity** (10^{12})	**Beam Size** (mm)
Switchyard	No	0.4	3.84 ± 0.05
Switchyard	Yes	0.4	3.93 ± 0.06
MiniBooNE	No	2.8	4.71 ± 0.03
Stacking	No	4.2	4.43 ± 0.05
Stacking	Yes	4.2	4.00 ± 0.08

Table 5.1: Extracted beam size for various types of beam, at different intensities from the Booster. The Switchyard and stacking cycles are phase locked with the MI before extraction. No beam size growth is evident with cogging.

with the longitudinal emittances of the extracted beam from the Booster, but that was eventually traced to unbalanced RF during a bunch rotation process.

To be more concrete, we studied the extracted beam's size after leaving the Booster. Table 5.1 shows data collected from the first multiwire profile detector downstream of the Booster extraction region. The multiwire was used to measure the horizontal beam width of the beam. During the period of time of the study, there were three types of beam being extracted: low intensity batches intended for the MI and ultimately the Switchyard; moderate intensity MiniBooNE batches, and high intensity batched intended for slip-stacking. During this study the cogging could be enabled and disabled on various cycles. We were able to interlace different types of cycles over a period of $\sim$ 30 minutes, hopefully nullifying any slow variation in the tuning quality.

The measured data give a null result for the impact of cogging on beam quality. In fact, cogged beam had a smaller beam size in the case of high intensity batches. Furthermore, the moderate intensity MiniBooNE beam showed the largest beam size by far. While unexpected, we believe this to be caused by the Booster's tuning being optimized for high-intensity, such that moderate intensity beam was less well accelerated. The conclusion from this study was that cogging has no consistent negative effects on the beam and that proper

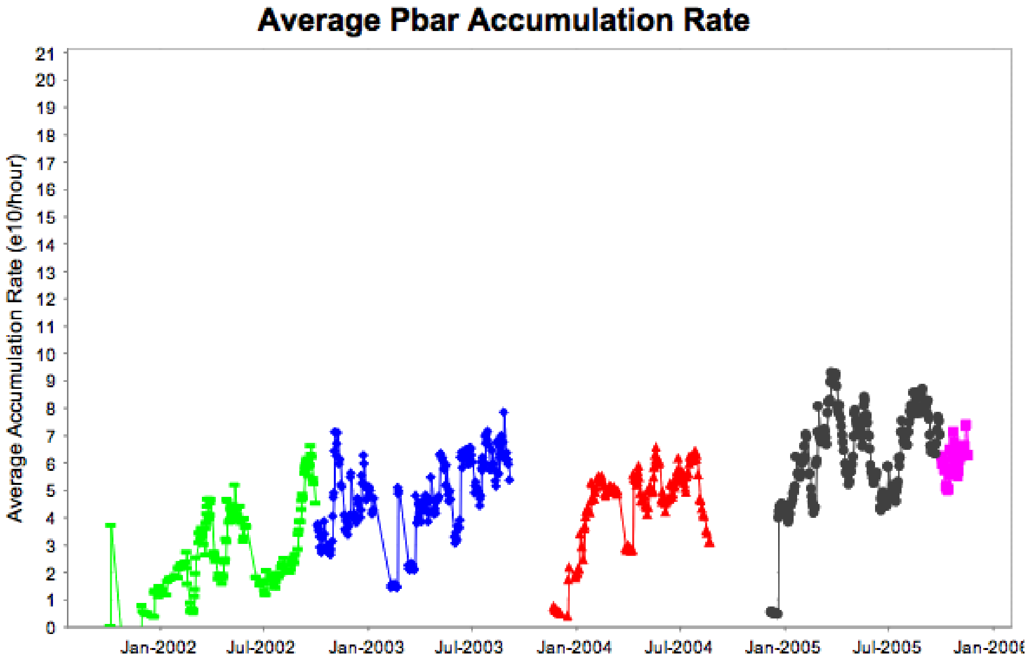

Figure 5.15: Antiproton production rate since the start of Fermilab's Run II. The implementation of slip-stacking with cogging in Jan. 2005 contributed to the higher accumulation rate.

tuning of the machine is likely to have the most significant effect on any operation.

5.6 Conclusion

The ultimate cogging system was implemented in the end of 2004 and has operated in support of the NuMI beam and slip-stacking for antiproton production. Since implementation, over 30 million Booster pulses have been cogged. The system works close to its ideal specification in terms of reducing losses, while requiring a later notch and radial motion of the beam during the cycle.

Fig. 5.15 shows the improvement of antiproton production rate since the start of Run II. Many sources of improvement have contributed to the increase; these include the storage of antiprotons in the Recycler and delivery of more protons per hour to the antiproton target. The increased number of protons has arisen through the process of slip-stacking, requiring multi-batch MI operation with cogging. Slip-stacking has allowed a maximum increase of $\sim 60\%$ in the number of protons to the antiproton production target. In the long term, it is hoped that the slip-stacking process can be improved to double the number of protons delivered to the antiproton production target – cogging will be required for the duration of this run.

The NuMI beam has been able to ramp up to a sustained rate of near 200 kW of beam on target, as shown in Fig. 5.16. This beam power is the highest for a high-energy neutrino beam[7], besting the previous WANF beam at CERN which ran at ~ 150 kW. Booster cycles typically exceed 4×10^{12} protons per batch. Cogging has been used to inject up to seven batches into the the Main Injector, as many as six for delivery to the NuMI line. Improvements in the rate of antiproton cooling have allowed the Main Injector combined cycles to occur more rapidly, and thus increase the NuMI average beam power. Further increases have occurred with the inclusion of dedicated NuMI cycles in the spacing between combined-cycles, and during times when the antiproton source is not stacking.

Fig. 5.17 shows the integrated number of protons delivered to NuMI. A total of 9×10^{19} have been collected so far, leading to millions on neutrino events in the MINOS near detector, and some number in the far. The first oscillation results should be possible when the number of protons exceeds 10^{20}, and improvements will be possible up to $\sim 25\times$ that number, to be delivered over the next several years.

In the very long term, alternatives to cogging become more attractive.

[7]The sub-GeV LAMPF beam at LANL ran at a higher power of ~ 700 kW

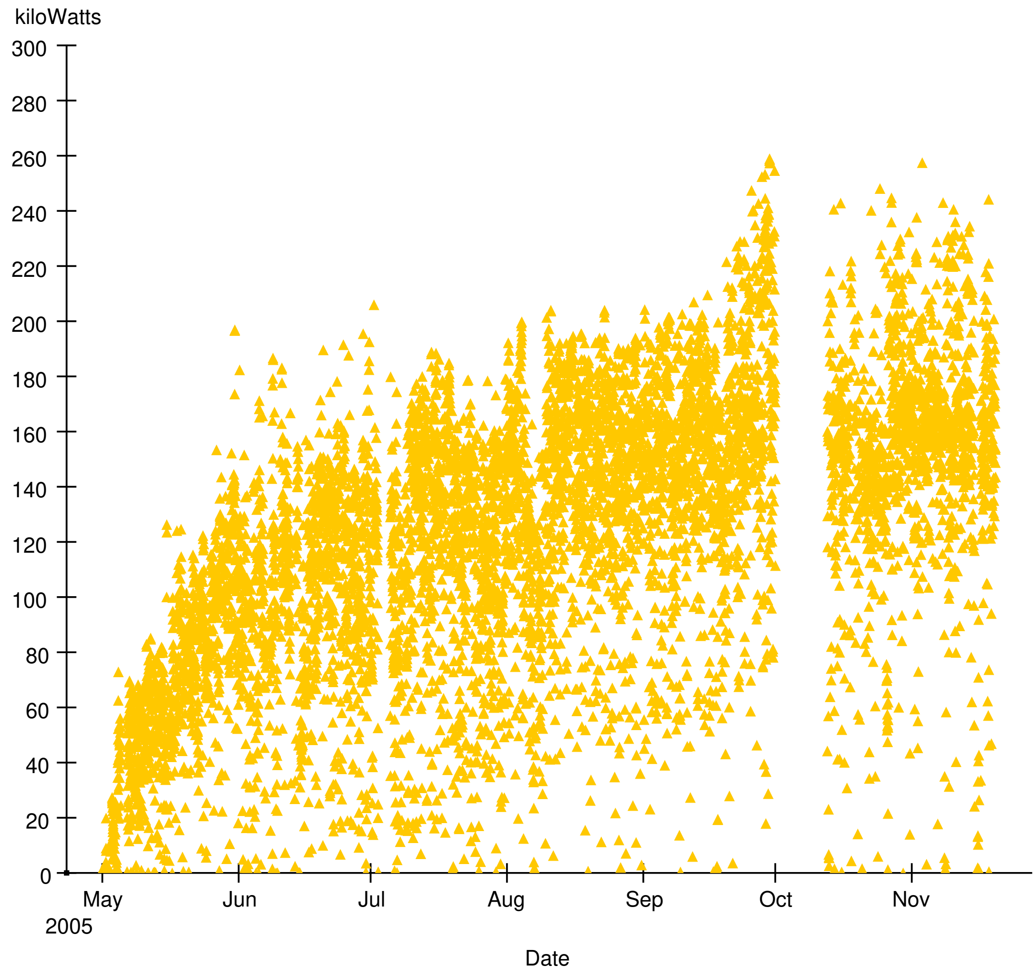

Figure 5.16: Average proton power delivered to the NuMI target over the length of the run. Each point is averaged over 30 minutes. The power increased quickly as the Booster batch sizes reached their maximum, and the Main Injector reached its multi-batch capacity. Remaining improvements come from reducing the cycle time, which is limited by antiproton production.

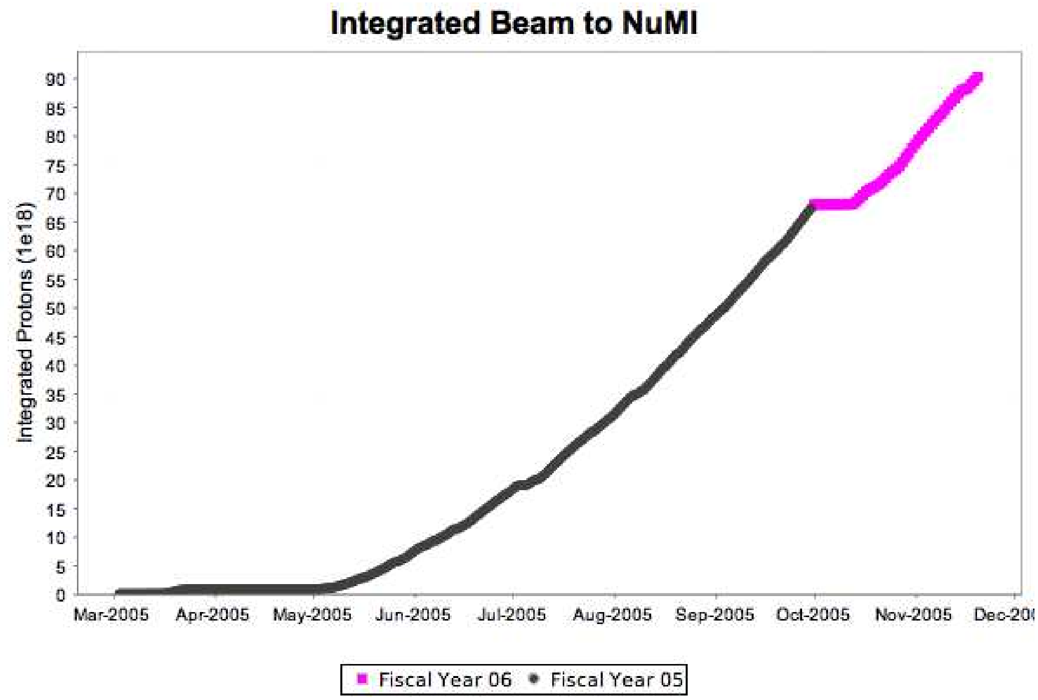

Figure 5.17: The integrated number of protons to NuMI since the start if its run. Cogging has been operational from the very start as part of the multi-batch procedure.

Some plans call for operating the Booster at 10-15 Hz, at the same or higher batch intensities. The increased throughput of 2-3 times as many protons will exceed extraction loss limits by simple extrapolation. Furthermore, the notching loss starts to become prohibitive, providing a preference for the alternative notch creation schemes. Either additional systems must be used to reduce the non-cogging extraction losses; or a different type of extraction that does not require cogging is necessary. Two such possibilities are the use of an intermediate storage ring or accelerator between the Booster and MI that could accommodate longitudinal slippage [81], or the development of an extraction kicker system that could reach full deflection within 10 ns – the inter-bunch spacing.

Chapter 6

Neutrino Beam Monitoring

6.1 Introduction

Experimental use of a neutrino beam requires knowledge of its fluence, energy, and flavor composition. Full measurement of the neutrino spectrum requires a neutrino detector; however, due to the small interaction rates of neutrinos, the MINOS near detector requires hours or weeks to accumulate enough data [82, 83] to observe changes in the beam. A less direct, but much quicker way to verify the beam quality is to observe the secondary and tertiary particles produced in association with the neutrinos, which can be detected on every beam pulse.

The NuMI Beam Monitors are a series of detector arrays located at the downstream end of the decay volume. The arrays consist of a number of ionization chambers that act as charged-particle counters. The Beam Monitors measure the fluence and spatial distribution of remnant hadrons in the beam and of the tertiary muons produced concomitantly with neutrinos in pion and kaon decays. Because of the large fluxes of these charged particles, the

Beam Monitors can obtain, in a single spill, a reasonably accurate measure of the neutrino beam intensity, direction, and consistency. Furthermore, these instruments permit rapid diagnostics to be performed in the event of a failure of the upstream neutrino beam systems.

This chapter reviews the design of Beam Monitor detectors built for the NuMI beam, and the measurements of which they are capable. Comparisons of the NuMI system are made to those installed in previous beamlines. Several aspects of the chamber design are reviewed. The Beam Monitors must operate in extremely high particle fluxes, so particular attention was applied such that their response behaves linearly.

Ch. 7 reviews literature data on helium gas necessary to describe ion chamber performance. Ch. 8 describes the results of a theoretical model used to extrapolate the expected ion chamber performance with the intense fluxes anticipated in NuMI. Ch. 9 explores the chambers' performance in the NuMI beam. Ch. 10 provides several diagnostic measurements for which the chambers have been used in NuMI beam operation. Ch. 11 contains an analysis in which the Beam Monitors were used to align the components in the target hall using the beam itself.

6.2 Secondary Beam Monitoring in Previous Experiments

The first neutrino beam, used at Brookhaven National Laboratory (BNL), did not explicitly use beam monitors (e.g., see [1]). That experiment based its neutrino flux calculations on estimates of hadron production yields from the targets. During the experiment they verified that beam was being produced using indirect methods, such as measuring the rate of loss of protons in a circulating beam. They subsequently used emulsion detectors placed in slots in the steel shielding in front of the detector to analyze the muon spectra, and thus provided a check on the muon-neutrino fluxes [84]. The CERN neutrino beam utilized a spectrometer downstream of its target station to measure secondary particles production *in situ* [85]. Thus, they verified the properties of the focusing horns; however, the spectrometer diverted the secondary beam and could not be used during neutrino running. A similar system, which does not impact the beam and can run continuously, is used in the Fermilab MiniBooNE beam [86]; it is limited to measuring certain wide-angle decays of kaons.

The "dichromatic", narrow-band beams of Fermilab [2, 87] and the CERN SPS [88] used dipole and quadrupole magnets to select a small momentum band of secondary pions and kaons from the target. Although such beams greatly reduce the overall neutrino yield per proton, they offer a precise flux measurement because of the known π/K momentum selection and because of the possibility of placing instrumentation along both the secondary π/K and tertiary μ beams [89, 90, 88, 91]. Such monitors were placed in the beam throughout the run and were essential for cross-section measurements of neutrinos on various target materials.

The wide-band, horn-focused beams do not lend themselves well to continuous secondary beam monitoring, so instead rely on measurements of the tertiary muon flux. The first such continuous muon beam monitoring system was built at the CERN neutrino beam run in 1967 [92, 93]. The system consisted of arrays of ionization chambers and scintillator paddles in the muon filter that measured profiles of the muons at various depths within the filter. At this beamline the filter was an arrangement of steel walls that stopped the muons via ionization energy loss. Variously called a muon shield, dump, or ranger, the filter is designed to cause all muons to slow and decay before reaching the neutrino detector[1]. The muon rates were used to verify the horn focusing and steer the proton beam. The results were also used for calculations of neutrino flux to the experiment [93, 94].

The recent oscillation experiments depend more on knowledge of the spectral shape and flavor composition of the beam, and less on the absolute rate of the beam. In their West Area Neutrino Facility (WANF) beam, CERN added secondary-emission foils downstream of their target and ionization chambers at the upstream end of their decay pipe to ensure proper targeting of the proton beam [95]. Likewise, the Institute of High Energy Physics-Serpokov (IHEP) used ionization chambers placed in their muon shield [96, 97]. BNL added "pion monitor" ionization chambers to its decay region [98]. The K2K beam in Japan used temporary Cerenkov counters and current transformers in the secondary beam [99], and also used several additional muon and neutrino detectors to better understand the tertiary beam composition [100]. The latest CERN beam, CNGS, will use arrays of ionization chambers for its muon monitors [101].

6.3 The NuMI Beam Monitoring System

The NuMI Beam Monitoring System consists of a single Hadron Monitor and a set of three Muon Monitors(see Fig. 6.1) . The primary goals of the system are to measure beam characteristics, on a pulse-to-pulse basis, that are used as verification for flux predictions at the MINOS far detector, and to assist in

[1]In MINOS, the filter consists of the unexcavated rock through which the muons must travel. As discussed below, small "alcoves" were excavated in portions of the rock; at which places the NuMI Muon Monitors reside and measure the muon fluences.

commissioning and aligning the beamline components.

MINOS relies on its Near Detector for an absolute measurement of its neutrino flux, significantly reducing the systematic errors involved with predicting the neutrino flux. The limitation on experimental precision then falls to our ability to extrapolate near detector measurements to the far detector. Beamline simulation shows that various changes in the beamline (such as component positions, beam position, or magnet currents) can affect the ratio of neutrinos between the two detectors. Such changes in the beamline may be detectable in the Near Detector; however it would take weeks to accumulate sufficient statistics.

The NuMI Beam Monitors provide spill-to-spill quality control measurements. They provide effective quality control as every muon-neutrino is accompanied by a muon. The monitors gain sensitivity to the energy spectrum of the muon beam by sampling the distribution at several locations in the shielding . The beam monitors' measurement ability is enhanced by configuring the NuMI beam for different neutrino beam energies [38]. Such energy variation studies are important, as measurement of the absolute flux of neutrinos in the peak of the low-energy distribution is not possible with the muon monitors – some of the corresponding muons range out in the upstream shielding. Past experience from CERN shows that it is desirable to have fast (i.e. single spill) measurements from muon monitors to diagnose problems such as misalignments [102] or non-ideal horn optics [103].

The Hadron Monitor, located at the end of the decay pipe, monitors the integrity of the target and baffle system. The flux at the Hadron Monitor is dominated by protons and hadrons from the target [83], so is relatively insensitive to the details of horn focusing. As such, it is used to locate the target-baffle system (see Ch. 11), and on a spill-to-spill basis to ensure that

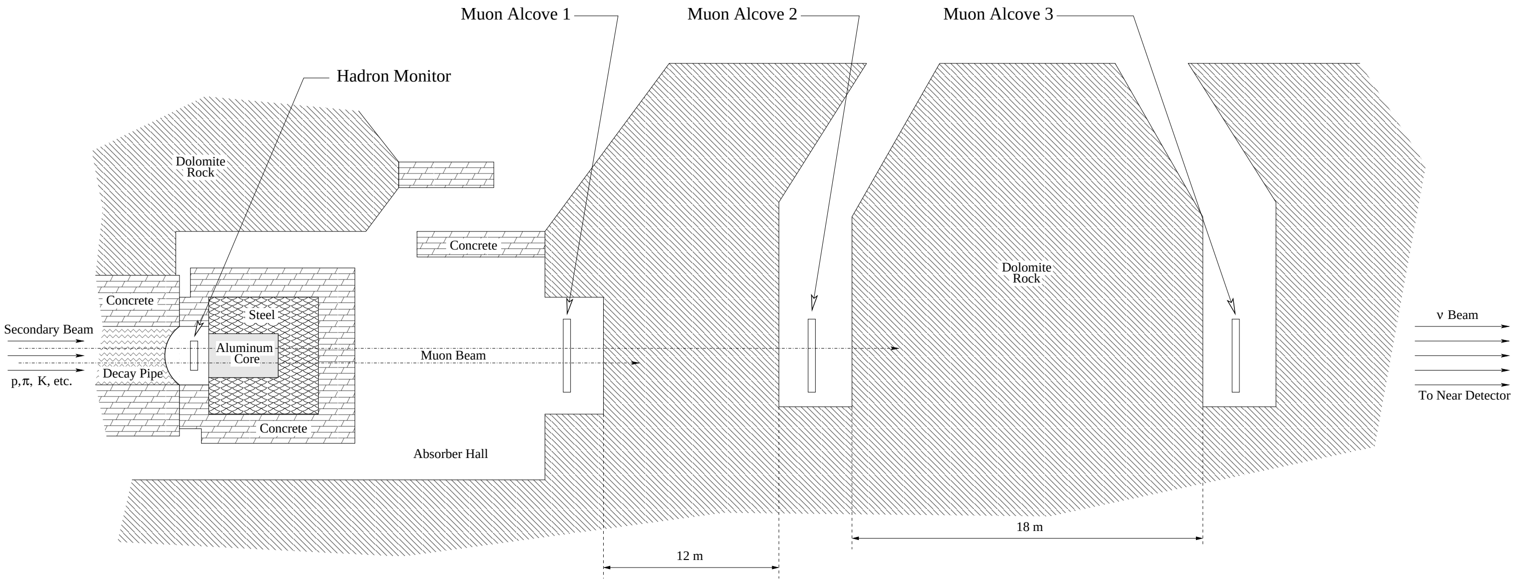

Figure 6.1: Bird's-eye view of the downstream areas of the NuMI beamline. The beam enters the area through the decay pipe consisting of remnant protons, secondary hadrons, decay muons, and neutrinos. The hadron beam's spatial distribution is measured at the Hadron Monitor and then stopped in the Hadron Absorber. The higher-energy muons penetrate the absorber and some distance of rock; along the way their spatial distributions are measured by the Muon Monitors. After 60 m of rock only the neutrino portion of the beam remains.

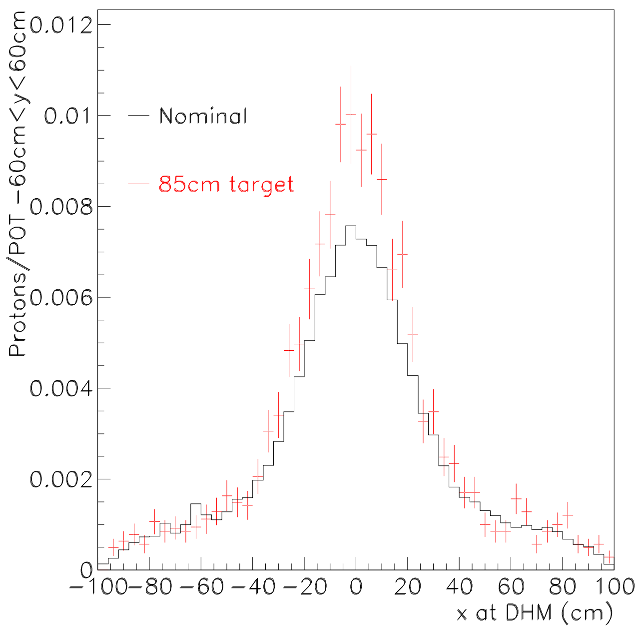

Figure 6.2: Simulated profile of proton fluence at the Hadron Monitor for the nominal beam case, and for a case where the last 10% of the target is missing (as in damaged).

no failure of the target has occurred. Experience from BNL ([104], Ch. 2) and our own experience operating NuMI (see Ch. 10) have shown that having an *in situ* monitor of the target's integrity can be quite useful, especially in light of the higher proton beam power available at Fermilab.

To appraise the Hadron Monitor's capability as a target monitor, various target failure scenarios were simulated. Fig. 6.2 shows the simulated transverse distribution of protons arriving at the Hadron Monitor during normal running, and for a case where the last 10% of the target is destroyed or displaced. The loss of target material reduces the attenuation and scattering experienced by the proton beam, resulting a more intense and narrower distribution at the Hadron Monitor.

The Hadron Absorber removes most of the protons and hadrons from

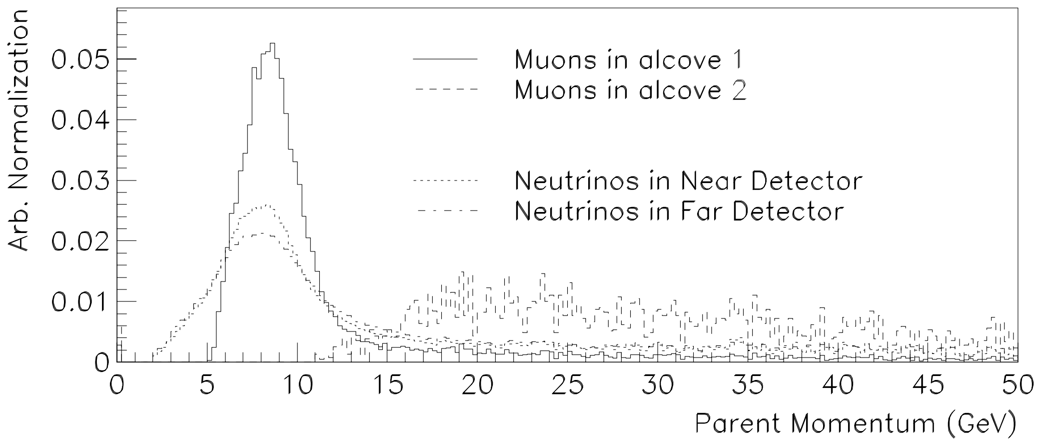

Figure 6.3: Momentum distributions of parent particles that produce particle interactions in the first two muon monitors and the neutrino detectors in the low-energy beam. The cutoff for muons penetrating to alcove 1 is ~ 5 GeV, and 12 GeV to alcove 2. The muons measured in alcove 1 come from parent hadrons which contribute strongly to the peak of the neutrino spectrum in the MINOS detectors, while alcove 2 measures the high-energy tail.

the beam. The absorber consists of a water cooled aluminum core, surrounded by steel, which is surrounded by concrete. The beam muons must pass through about 2.5 m of Aluminum, 2.3 m of steel, and 0.9 m of concrete to reach the first Muon Monitoring station.

The first Muon Monitoring station, located behind the absorber, measures the spatial distribution of the muon flux penetrating the absorber. As shown in Fig. 6.3, the muons that survive to that point are those with 5 GeV or greater momentum. Fig. 6.3 shows the momentum distribution of the parent particles that decay into the muons and neutrinos, weighted by the probability of either the muon to arrive at the monitoring station or the probability of the neutrino to interact in the MINOS near detector. The peak parent momentum is slightly higher for the first alcove than for the MINOS near detector. Therefore, alcove 1 measurements correspond most closely to the neutrinos of interest.

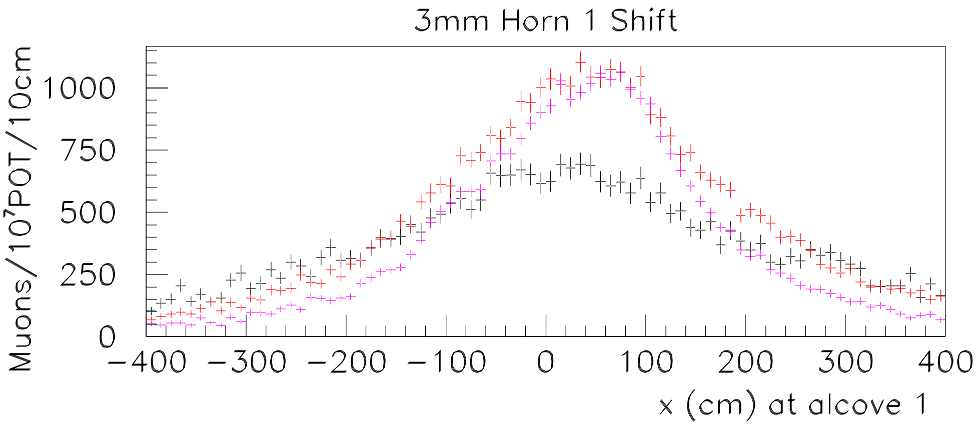

Figure 6.4: Simulated horizontal profile of muons at alcove 1 in the case where the first focusing horn is displaced 3 mm to the right of its nominal location. The three curves are for the low-energy (blue), pME (red), and pHE (magenta) beams. The center of the muon distributions are consequently displaced horizontally, allowing the use of the alcoves in assuring the horn focusing. Also, the sensitivity to misalignments is generally greater for higher-energy beams. (from [38])

Considering the distribution in Fig. 6.3, we might have wished to have our first monitor sampling below the peak in the neutrino parent momentum distribution. In fact, previous beam monitors in other experiments have sampled down to muons of initial energy 1.2 GeV [84]. Such a measurement, however, is not possible in the NuMI beam. Previous beamlines used lower energy primary proton beams such that the hadrons ranged out quickly. A muon monitor placed in the NuMI shielding, sensitive to lower energy muons, would be dominated by the remains of the hadronic shower, as has been shown previously in other experiments [105, 106].

Muon alcove 1 is located ~ 8 m behind the absorber. Fig. 6.4 shows simulated muon profiles at the alcove when the first focusing horn is misaligned. The alcove can be used to measure the shifted muon profile in a single spill, and provide notice that such a displacement occurred. Muon distributions similarly shift for dislocations of the primary proton beam (see §10.2.2). The array is subject to a neutron background emanating from the absorber; the signal from which is usually much less than the muon signal [107].

Muon alcove 2 follows approximately 12 m of unexcavated rock downstream of alcove 1. This rock further attenuates the muons such that ~ 12 GeV of initial momentum is necessary to penetrate the rock. The distribution of parent particle momenta for these muons is also shown on Fig. 6.3. The muons measured here correspond to the high-energy tail of the low-energy neutrino spectrum, or at about the peak of the medium-energy beam.

Muon alcove 3 follows another 18 m of unexcavated rock. The cutoff here for muon energy is about 24 GeV, so these muons correspond to only the highest-energy portion of the neutrino beam above the focusing peak, even for the medium- and high-energy beams. The utility of alcoves 2 and 3 is mostly limited to situations where higher energy pions are focused down the decay pipe, whether intentionally by a change of beam energy, or unintentionally through a missteering of the proton beam.

6.4 Expected Particle Rates & Distributions

Table 6.1 shows the calculated charged particle fluence at each detector, along with predictions for the amount of charge collected on a pixel and the amount of radiation each monitor is exposed to. Of note is that the Hadron Monitor is exposed to two orders of magnitude more charged particle fluence, and three orders of magnitude more radiation[2]. The Hadron Monitor construction needed to be different from the Muon Monitors to accommodate higher signals

[2]The additional order of magnitude in radiation is due to the contribution of neutrons. Different Monte Carlo models gave very different estimates of neutron fluence, but it was always higher at the Hadron Monitor than at the Muon Monitors. Additionally, protons and other hadrons cause more radiation damage than do muons.

	Fluence (/cm^2/10^{12}ppp)	**Charge** (pC/pixel)	**Dose** (MRad/year)
Hadron Monitor	6.8×10^7	25,000	2,000
Muon Alcove 1	6.5×10^5	720	14
Muon Alcove 2	8.5×10^4	95	1.2
Muon Alcove 3	2.5×10^4	28	0.5

Table 6.1: Predicted fluence, collected charge, and radiation dose at the beam monitors in the low-energy beam. The maximum fluence is calculated for the low-energy beam and evaluated at the center of each monitor. The charge collected per spill is calculated using the fluence, an area of 58 cm^2, a gap of 1 or 3 mm, ionization of 1.6 mm^{-1}, and a beam intensity of 2.5×10^{13} protons. The dose is calculated for the center of each monitor; the Hadron Monitor calculation includes an upper bound for the neutron contribution, which can cause the estimate to vary by an order of magnitude. Monte Carlo calculations taken from [82, 83], radiation estimates from [108].

and withstand greater radiation damage. The Muon Monitors were designed to withstand a marginal amount of radiation and tuned to more precisely measure the smaller signals.

Fig. 6.5 shows the calculated radial distribution of the secondary beam at the Hadron Monitor for three beam conditions. The first condition is the nominal beam profile; the half-width-half-maximum of the distribution is $\sim$ 20 cm, as expected from multiple scattering in the NuMI target. The distribution flattens by 50 cm away from the center. The signal at that point starts to be dominated by electromagnetic particles from showers; furthermore, the effect of horn focusing is almost immeasurable in the Hadron Monitors, so coverage outside of 50 cm is of limited use.

The third distribution shown in Fig. 6.5 is that of the proton beam at the Hadron Monitor when the target is removed. Such direct transport of the proton beam to the absorber occurred during the commissioning period; at which time the Hadron Monitor verified proton transport through the target

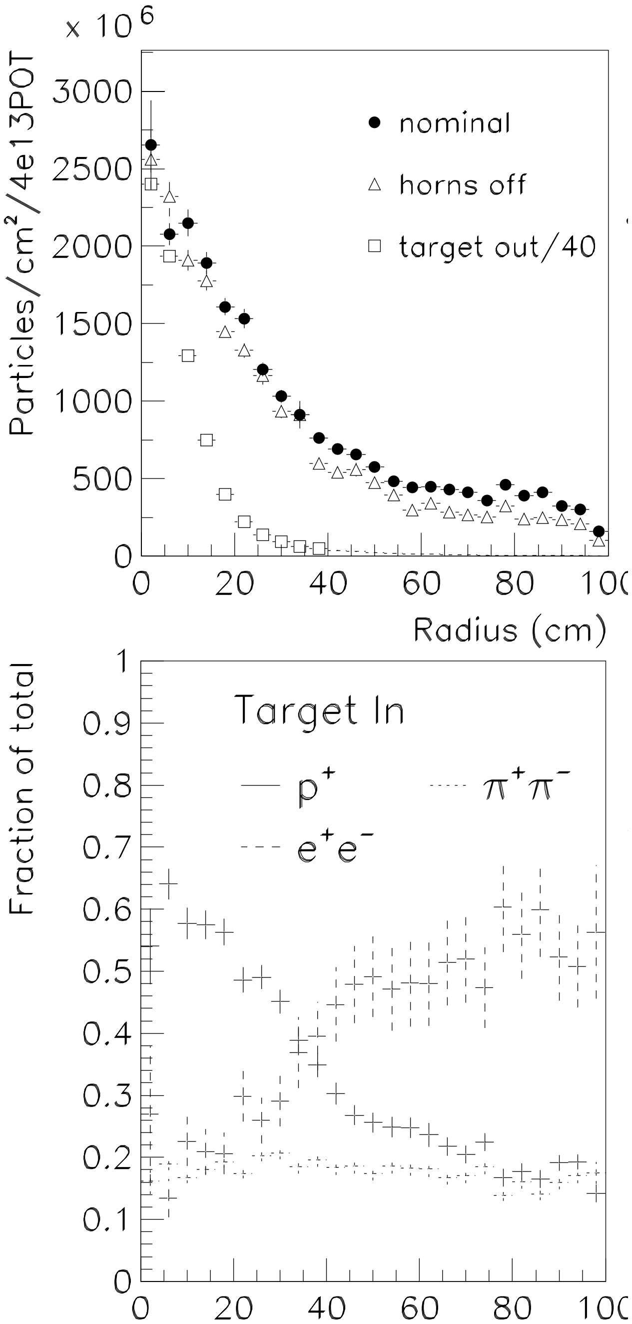

Figure 6.5: Above: Monte Carlo simulated radial distributions of particle fluence in the Hadron Monitor for three beam conditions: nominal low-energy running, low-energy-running without horn focusing, and the proton beam transported to the monitor without an intervening target. Below: simulated radial distribution of fluence composition for nominal low-energy running. (from [82])

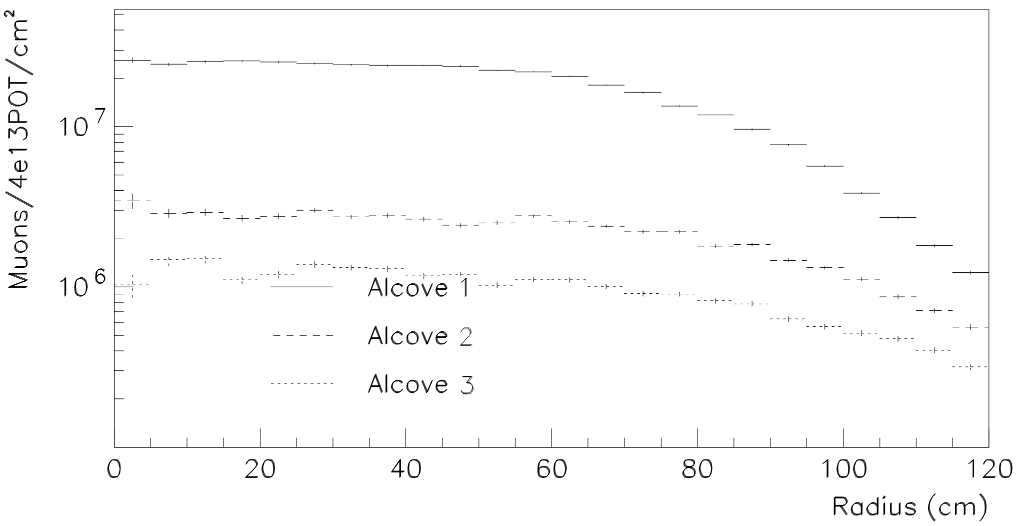

Figure 6.6: Simulated radial distributions of the muon fluence at each of the muon alcoves in the low-energy beam. (from [82])

hall, and was used for the beam-based alignment of target hall components (Ch. 11). The direct proton beam is significantly narrower, such that the beam half-width is only $\sim$ 10 cm. To measure this beam well we need pixel spacings of less than 20 cm and coverage gaps of less than 5 cm.

Fig. 6.6 shows the radial distributions of muon fluence in the Muon alcoves for nominal low-energy running. The muon beam is very broad by the time it reaches the Muon Monitors, its extent being set largely by the decay pipe size and multiple scattering in the intervening material[3]. The great breadth of the beam allows the Muon Monitors to operate with significantly less granulation than the Hadron Monitor.

Table 6.2 shows simulation predictions of beam size and intensity at the monitoring stations for three different beam configurations. The low-energy is the design beam which will likely be used for the majority of neutrino oscillation studies. The pME and pHE beams are higher energy neutrino

[3]The simulated details of these distributions will change once the Hadron Absorber is fully described in the Monte Carlo simulation.

		LE	pME	pHE
Hadron	fluence (10^7/cm^2/10^{12}ppp)	6.8	6.8	6.8
Monitor	rms (cm)	20.	20.	20.
Muon	fluence (10^5/cm^2/10^{12}ppp)	6.5	10.0	9.0
Alcove 1	rms (cm)	190	130	110
Muon	fluence (10^5/cm^2/10^{12}ppp)	0.9	5.0	7.2
Alcove 2	rms (cm)	250	140	110
Muon	fluence (10^5/cm^2/10^{12}ppp)	0.35	0.5	2.3
Alcove 3	rms (cm)	190	250	120

Table 6.2: Predicted maximum particle fluence and beam size at the monitoring stations for different types of beam. **LE** is the low-energy, nominal design, beam – now replaced by the LE10 beam. **pME** and **pHE** are the partial-medium and -high energy beams described in [38].

beams produced by moving the target back to improve the focusing of high-energy mesons [38]. The rates in the later alcoves are predicted to increase with higher-energy beam.

6.4.1 Radiation Levels

Table 6.1 enumerates the significant amounts of radiation the monitors will be exposed to. The radiation has the potential to damage materials and create residual activation. Damage from radiation can manifest directly, as in causing the dielectric in a cable to fail, shorting the high voltage. Such radiation can also damage the monitors indirectly by vaporizing organic materials that can then be deposited on the electrodes and reducing their efficiency. To mitigate the effects of radiation the monitors were designed with materials that could survive the radiation, eschewing materials such as Teflon and other radiation-vulnerable plastics for seals and dielectrics, as well as potentially volatile organic materials like G-10 circuit board within the gas vessels. The design is made up of almost entirely four types of components: metal, ceramic, polyimide[4], and PEEK plastic. Radiation tests of chamber components were conducted at the University of Texas 1 MW fission reactor; doses of up to 12 GRad were shown to have no detrimental impact.

The issue of residual activation is important for the Hadron Monitor[5]. The Hadron Monitor is designed to be removable, and possibly replaced in the event of failure. It is estimated that after 1 year of running the Monitor will be radioactivated to $\sim$ 700 Rad/hr on contact [109], which is too high for any substantial human proximity. To minimize activation, the Hadron monitor was constructed out of Aluminum, instead of steel, wherever possible. The resulting radionuclides from irradiation of aluminum generally have shorter half-lives than those of steel.. Additionally, the dose from an extended object depends strongly on its mass, so material was removed wherever possible to minimize the mass.

6.5 Ionization Chambers

The Muon and Hadron Monitors are arrays of ionization chambers. Each charged particle passing through the chamber ionizes a gas subject to an electric field; the ions and electrons then drift to electrodes and the current measured is proportional to particle flux. Details of ionization chamber operation are discussed in Chs. 7 and 8. Here we discuss the particular parameters of the ionization chambers and some of the reasons for those choices.

The chamber chosen for the NuMI Beam Monitors is of the parallel-plate type, where two planar electrodes are oriented parallel to each other and

[4]Commonly known by the trade name Kapton.

[5]The Muon Monitors will suffer almost no residual activation as a result of the particle fluence being markedly lower and that muons rarely interact in such a way to produce radioisotopes.

an electric bias applied across the separation. This geometry is useful for its simplicity, for the uniformity of the electric field , and for the possibility of small electrode separation. The uniform electric field controls the drift of the ions and forestalls ionization amplification (see §7.5). A small gap is desired as smaller ionization chambers generally have better collection efficiency at the very high particle fluences anticipated.

Ceramic chambers were chosen in place of the circuit-board G-10 chambers used elsewhere [110]. The ceramic is more radiation tolerant; it is furthermore stiffer, allowing smaller electrode separation. A chamber with a large transverse electrode size and small separation was chosen to maximize both signal strength and collection efficiency.

The chambers in Fig. 6.7 have a substrate of 99.7% pure alumina ceramic cut to 4" squares. Silver-platinum electrodes are deposited via photolithography. The high-voltage plate has a solid electrode across its entire face. The signal plate consists of the 3"×3" signal electrode and a 1 cm guard ring that surrounds the sensitive electrode. The guard ring, sense electrode, and bias electrode make their electrical connections in the corners.

The plates are separated by precision ceramic washers of either 1 or 3 mm width – for the Hadron or Muon Monitor, respectively. While these washers varied by only 0.0005" in thickness, the plates' flatness varied by 0.002" (Muon Monitors) or 0.0005" (Hadron Monitor).

The ionization medium for these chambers was chosen to be pure helium gas, continuously flowed, at atmospheric pressure. Helium has low density and low ionization per traversing particle, mitigating the effect of space-charge buildup within the chamber. A pure, inert gas is desired for chamber longevity. The choice of atmospheric pressure was for simplicity in the design of the gas system. Lower gas densities might have been preferable in the Hadron Monitor

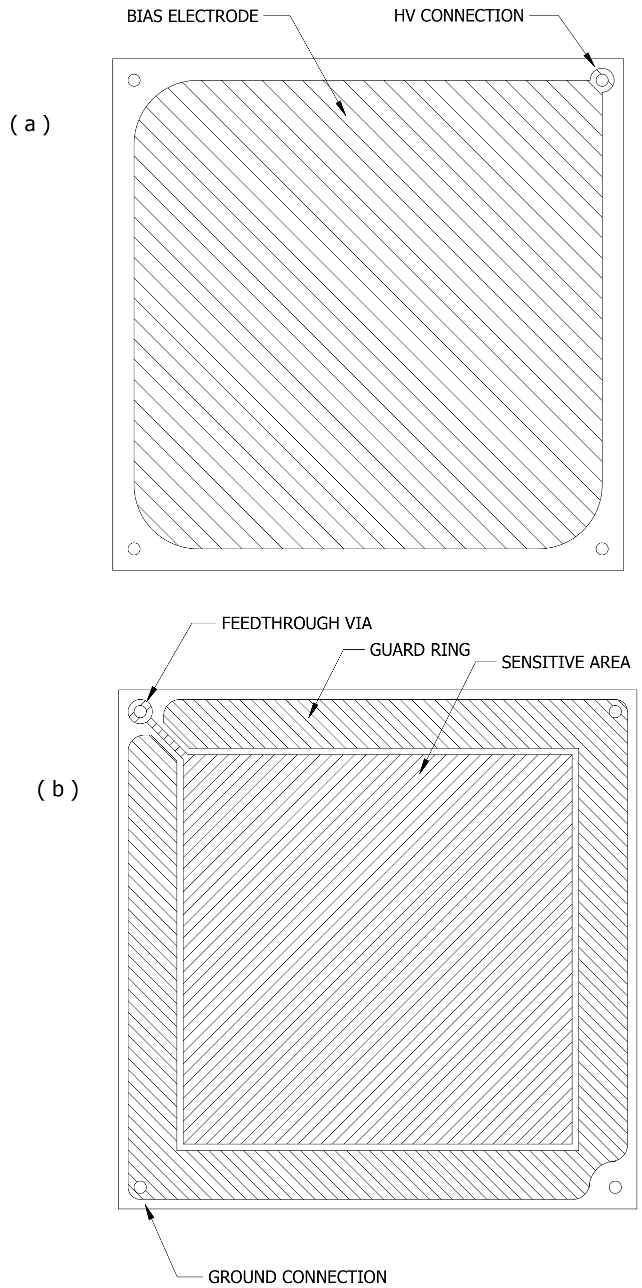

Figure 6.7: The ceramic plates that make up each vessel. Both have silver-platinum electrodes facing the the ionization volume. Plate a is the high voltage plate with a single electrode that connects to the bias supply through a corner post. b is the signal plate that has two electrodes: the sense pad that defines the active volume and a grounded guard-ring electrode around it. Both make electrical connections in the corners.

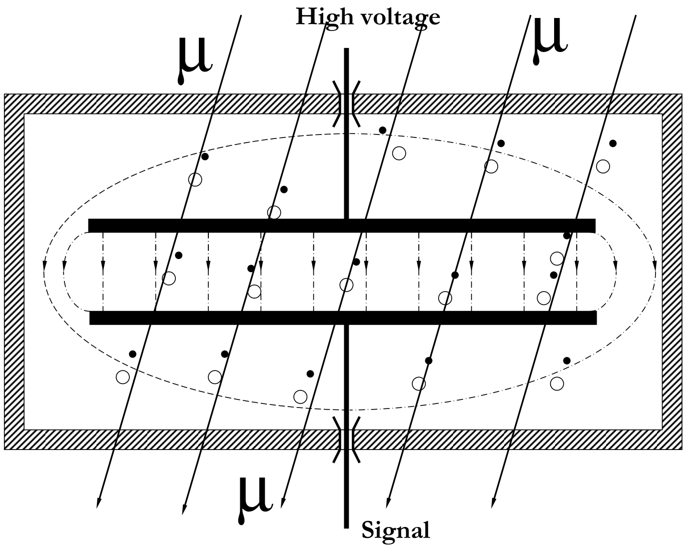

Figure 6.8: Schematic of a stray ionization process. Some electric field within the chamber (in this case caused by an unshielded high-voltage cable) causes charge to drift to the signal cable – if unshielded. As the volume of gas outside of the chamber can be large, the signal from stray ionization can be significant if not properly avoided.

to reduce space charge effects, but would have required a more leak-resistant system to operate at negative pressure, as compared to atmosphere.

6.5.1 Stray Ionization

Only a small portion of the helium within each detector vessel is part of the active volume. However, all of the helium volume is ionized by the beam. The produced charge may drift to part of the wire that carries the signal from the detector's sense pad to the electronics. Such long conductors, if they collect stray ionization, will distort any spatial profile measured by the chambers.

A configuration resulting in stray ionization is shown in Fig. 6.8. In the case shown, both of the cables providing high voltage and removing the signal charge are unshielded within the gas vessel. Without an intervening ground, there are lines of electric field from the HV wire to signal wire. Ions and electrons will drift along these lines and increase the measured signal. The additional ionization will come from areas that may have been irradiated to a different extent than the chamber; also the drift will be over a larger distance, so will be less linear with intensity.

In order to better understand the mechanisms for stray ionization and to optimally design electrical connections to the ion chambers, we performed several bench tests with different conditions [111]. These tests were used to determine configurations which would not collect charge from outside the active volume. A series of ionization chambers were mounted inside a single gas vessel filled with either helium or air. These chambers had electrical connections consisting of either bare wire, insulated wire, or coaxial cables in which an outer ground shield covered the cable insulation. Similarly, several electrical feedthroughs were tested. α sources were mounted inside and outside the ionization chambers, and the relative currents from signal (inside) and stray (outside) ionization noted. It was found that all connections and feedthroughs required an insulator around the contacts, and further a grounded conductor outside the insulator.

6.6 Hadron Monitor

The Hadron Monitor was designed as a 7×7 grid of 1 mm ionization chambers in a single gas vessel, oriented transversely to the beam axis. The chambers are spaced 4.5" center-to-center, for a total active extent of 30" horizontally and vertically. The vessel is an aluminum box with a cover that is sealed to the box with a solder-wire gasket. Each high voltage and signal channel has

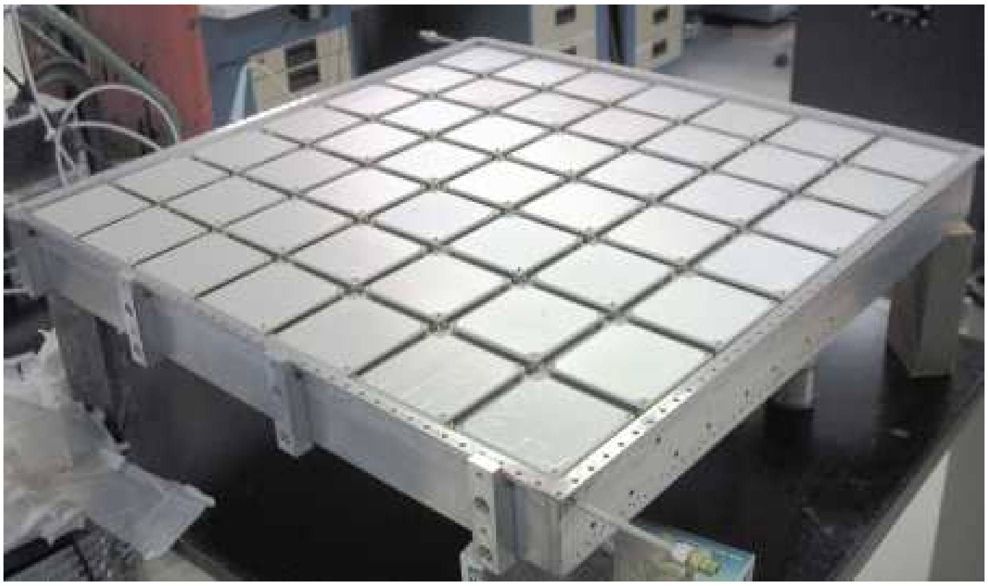

Figure 6.9: View of the Hadron Monitor interior. The 49 chamber array measures about 1 m square.

its own ceramic feedthrough[6], custom high-radiation cable, and coaxial cables back to the equipment racks outside of the radiation area; this separation allows individual channels to be disconnected in the case of pixel failure. The following sections describe the design, construction, and installation of the Hadron Monitor

6.6.1 Construction

The Hadron Monitor interior is shown in Fig. 6.9. The vessel is entirely made of aluminum and was as thin as possible to minimize radioactivated mass. The ceramic feedthroughs were first installed on the blank aluminum box which had the 98 holes in it. The ceramic feedthroughs are standard high-vacuum 10 kV feedthroughs with a central stainless conductor brazed to the ceramic

[6]Each channel was provided its own feedthrough to avoid the challenge of running cable within the gas volume while avoiding stray ionization and gas poisoning. See §6.6.2 for details on the cabling.

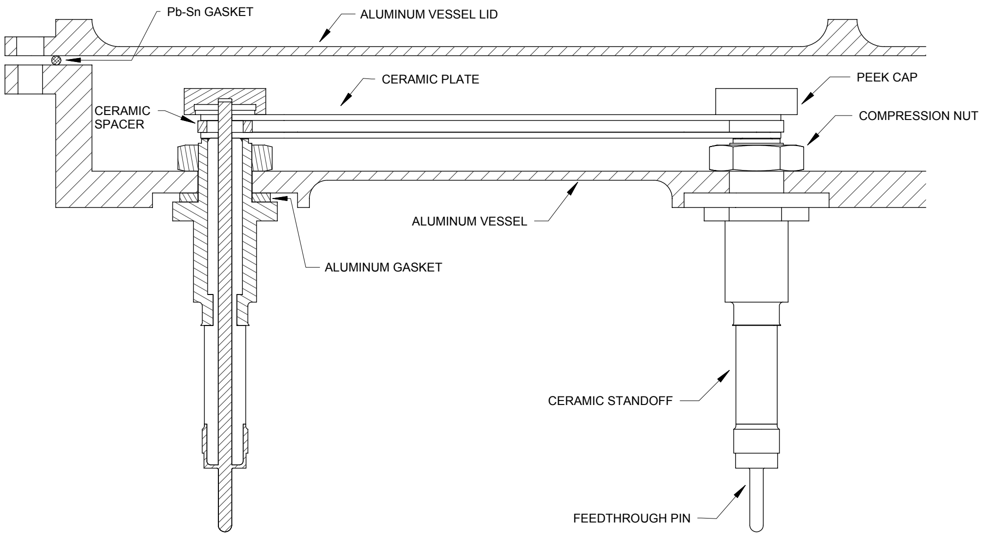

Figure 6.10: Profile of the Hadron Monitor feedthrough and chamber assembly. The feedthrough is a modified vacuum high-voltage feedthrough that is gasketed to the aluminum vessel. The chamber plates mount directly on the feedthroughs. A peek cap is attached to the top of the feedthrough lead to avoid collecting stray ionization.

and an external stainless jacket for welding. Each feedthrough was fitted with a male threaded fixture which was used to compress an a aluminum gasket at the aluminum vessel wall (see Fig. 6.10). The ceramic plates are supported at two corners by these feedthrough pins, and at their other two corners by aluminum standoffs.

The cover to the Hadron Monitor is an aluminum lid, sealed to the vessel with a lead-tin solder gasket. A back cover of aluminum sheet metal was fastened over the feedthroughs and cables for protection.

6.6.2 Cabling

Because the Hadron Monitor is buried inside the absorber shielding, where radiation levels are quite high, particular attention was given to its cabling. Calculations show [108] that the radiation levels fall off by a factor of ~ 3 at 1 m radius, and another factor of 10 through the concrete shielding. Our expectation was that Kapton-insulated cable would survive these radiation levels. For an additional safety margin, a redundant cable was designed for the span behind the Hadron Monitor, where the radiation levels are highest: a coaxial Kapton cable was stripped of its outer ground braid for its last $\sim$ 1 m and a ceramic tube slid over for extra insulation. The ceramic has an aluminum sleeve around it, acting as the exterior ground; each sleeve was soldered to its associated braiding.

Outside of the absorber shielding the coaxial Kapton cable transitions to standard RG-58 coaxial and SHV cables for the run to the electronics. The transitions were soldered to circuit boards inside shielded stainless boxes. The boxes were potted with RTV[7] silicone rubber to prevent breakdown and spurious ionization. The RG-58 were run for about 100' through the absorber hall to the electronics racks outside the radiation area. Inside the hall, the cable trays were routed to avoid higher-radiation areas where feasible.

6.6.3 Installation & Alignment

The Hadron Monitor sits directly in front of Hadron Absorber and is only accessible by a 6" $\times$ 40" slot in the side of the absorber. The Hadron Monitor is installed by sliding it on rails into position, and removed the same way – if ever necessary. The alignment of the Hadron Monitor is defined by the

[7]Room-Temperature Vulcanizing

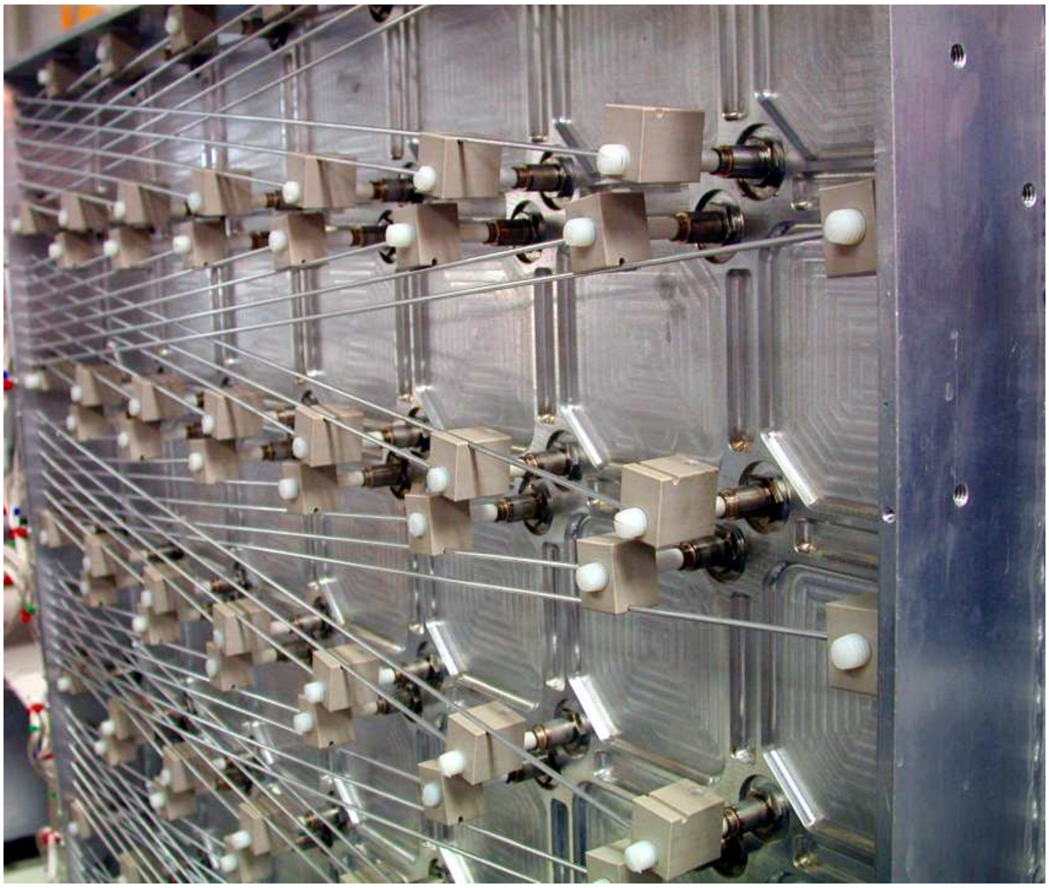

Figure 6.11: The backside of the Hadron Monitor showing the hard-line cables and the couplings to the electrical feedthroughs. The cable is a central Kapton insulated wire surrounded by ceramic cylinders sheathed in a solid aluminum tube. The couplings are machined PEEK cylinders with a polyethylene nut to allow access to the electrical connection. HV feedthroughs are grouped together to reduce electric fields at the signal feedthroughs.

location of the rail and the stop on that rail. Alignment of these components was performed six months previous to installation [112].

While the Beam Absorber was under construction a lower rail for supporting the Hadron Monitor, an upper rail for keeping it vertical, and a stop for defining the horizontal position were installed. Survey markers in the absorber hall were used for alignment and expected to be accurate to a precision of $\sim$ 0.5".

A length of stainless-steel wire rope was attached to the Hadron Monitor to allow removal, if necessary. The wire rope, gas lines, and Kapton cables were bundled together at the edge of the shielding. A 6" thick steel door was then closed over the installation slot, sealing off the interior except for a few inch hole in the door that allowed the bundle of cables and pipes to penetrate.

6.7 Muon Monitor

The Muon Monitors are composed of three 9×9 arrays of 3 mm ionization chambers. The arrays are each made up of nine "tubes" in which there is a tray with nine chambers mounted. The chambers are spaced 10" center-to-center (horizontally and vertically) for a total array coverage of 2×2 m^2. Each tube is an independent helium volume with electrical feedthroughs and gas connections. For design details, see [113] and [114].

6.7.1 Construction

The layout of a Muon Monitor tray is shown in Fig. 6.12. The tray is an machined aluminum U-channel 90" in length, 5.5" in width, and 1.5" in height. Each of the nine chambers is mounted on four standoffs that are each fastened

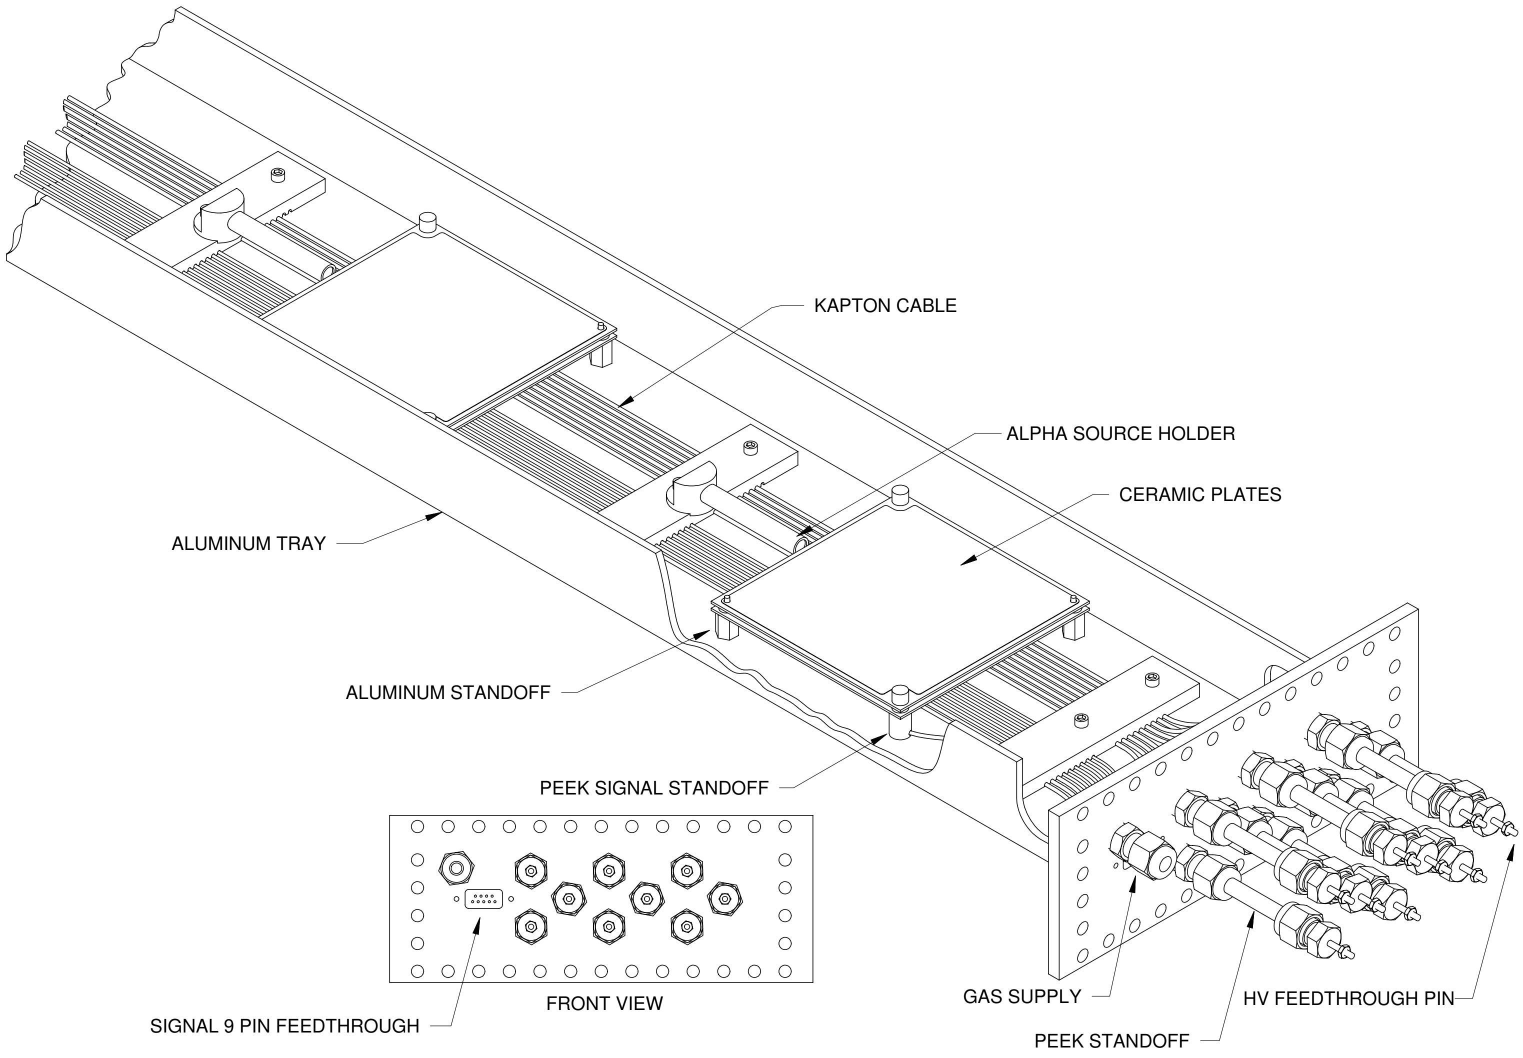

Figure 6.12: View of a portion of a Muon Monitor “tray”. The chambers are mounted to an aluminum channel and the cables run down the length to the flange.

to the tube. The two standoffs that carry the ground are solid aluminum with pins that are soldered to the ground corners of the signal plates. The other two standoffs are like coaxial cables, in that they carry the electrical connection within a dielectric.

The coaxial cables are routed on the tray toward one end, underneath the other chambers. The cables are held in place between each tube by an aluminum comb that is fastened to the tray. Mounted on each comb is a holder for a 1 μCi Americium-241 source. Each source is situated directly adjacent to the gap of each chamber, providing illumination of the active volume by α-particles. These sources produce a small ionization current that was intended to be used for calibration. However, the character of the ionization was so different than muons and the signal so small that they were unsuitable for calibration [115].

All of the cables are terminated in an end-plate that was bolted to the tray, the exterior of which is shown in Fig. 6.13. The high-voltage cables fed into custom feedthroughs (see Fig. 6.14). The signal cables were all soldered onto the pins of a 9-pin D-type ceramic vacuum feedthrough. A PEEK collar fit over the cables and pins to insulate them from the gas.

The assembled tray and end-plate were inserted into a "tube", which was aluminum box-channel 6" $\times$ 2" with 1/16" side-walls and 90" in length with flanges welded onto the ends. One side of the tray was fixed to the end-plate, and was thus fixed to the tube upon sealing. The opposite end was fastened to the tube via an internal plate. The tubes were flushed and sealed with dry nitrogen for transportation and storage until installation.

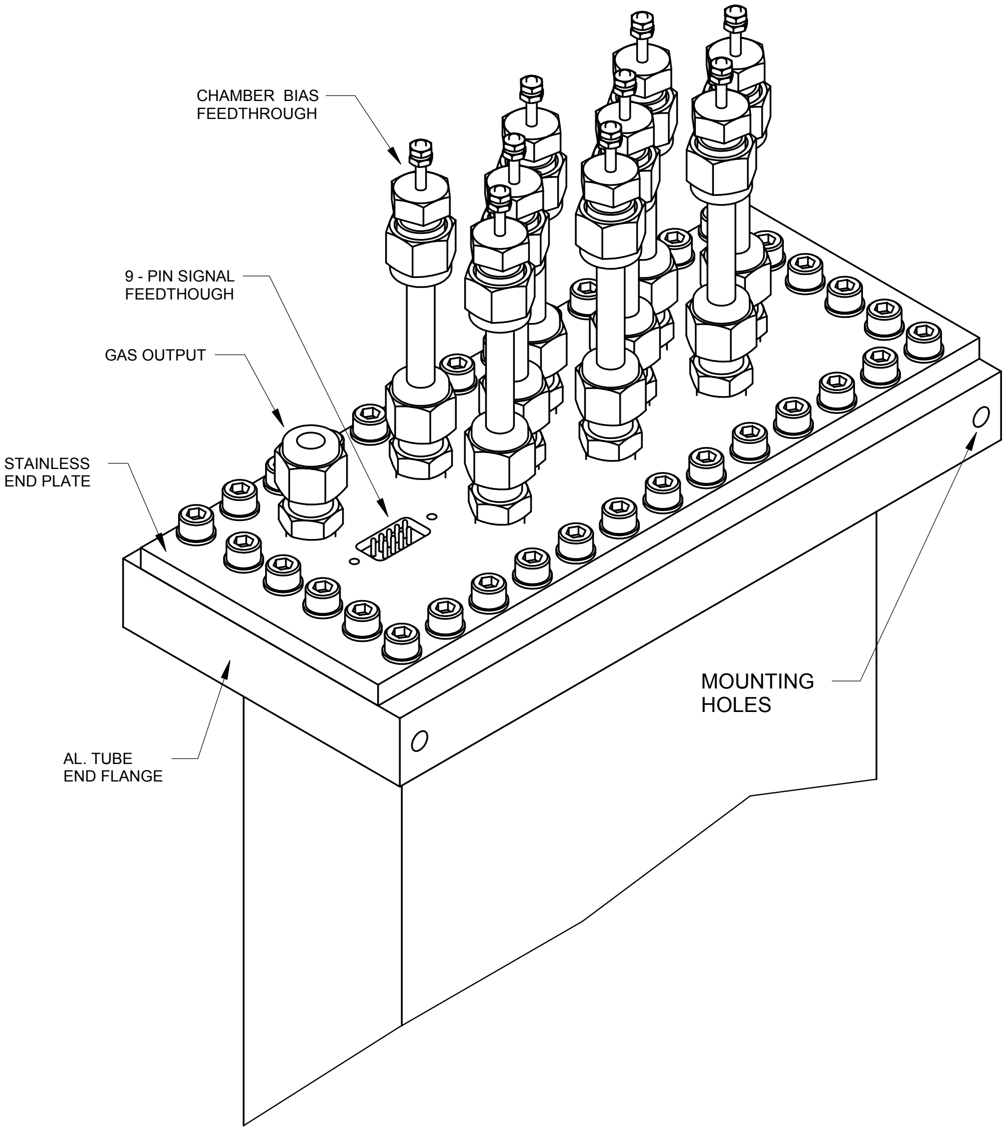

Figure 6.13: View of an end-plate for the muon tubes. The nine high-voltage channels each have their own feedthrough. The nine signal channels are routed through a single metal-ceramic nine-pin connector. The gas connection is made with a compression fitting. The seal of the plate to the tube is made by compressing an Aluminum wire gasket. The opposite plate on each tube has only a single compression fitting for a gas connection.

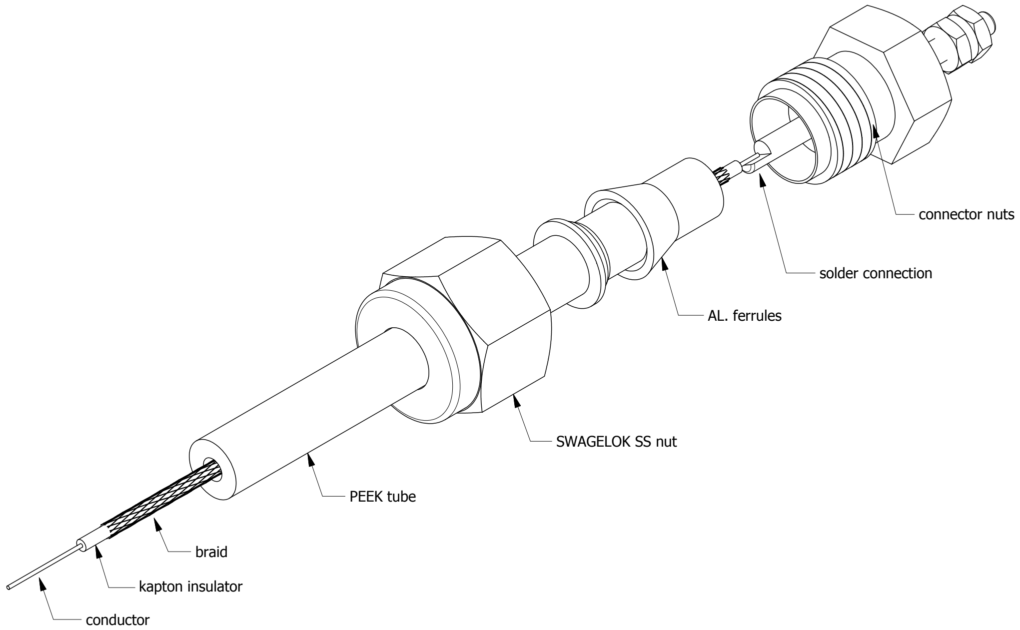

Figure 6.14: The high-voltage feedthrough used in the Muon Monitors. It is composed of compression fittings and a PEEK cylinder as the dielectric. One compression fitting (not shown) is welded to the end-plate and makes a gas-tight seal with the PEEK. The coaxial Kapton wire is fed into the cylinder; the end is stripped and the center conductor is soldered to a pin welded to the outer compression-fitting cap. This cap provides the other gas seal on the PEEK. The external electrical connections are made by nut compression to a G-10 circuit board.

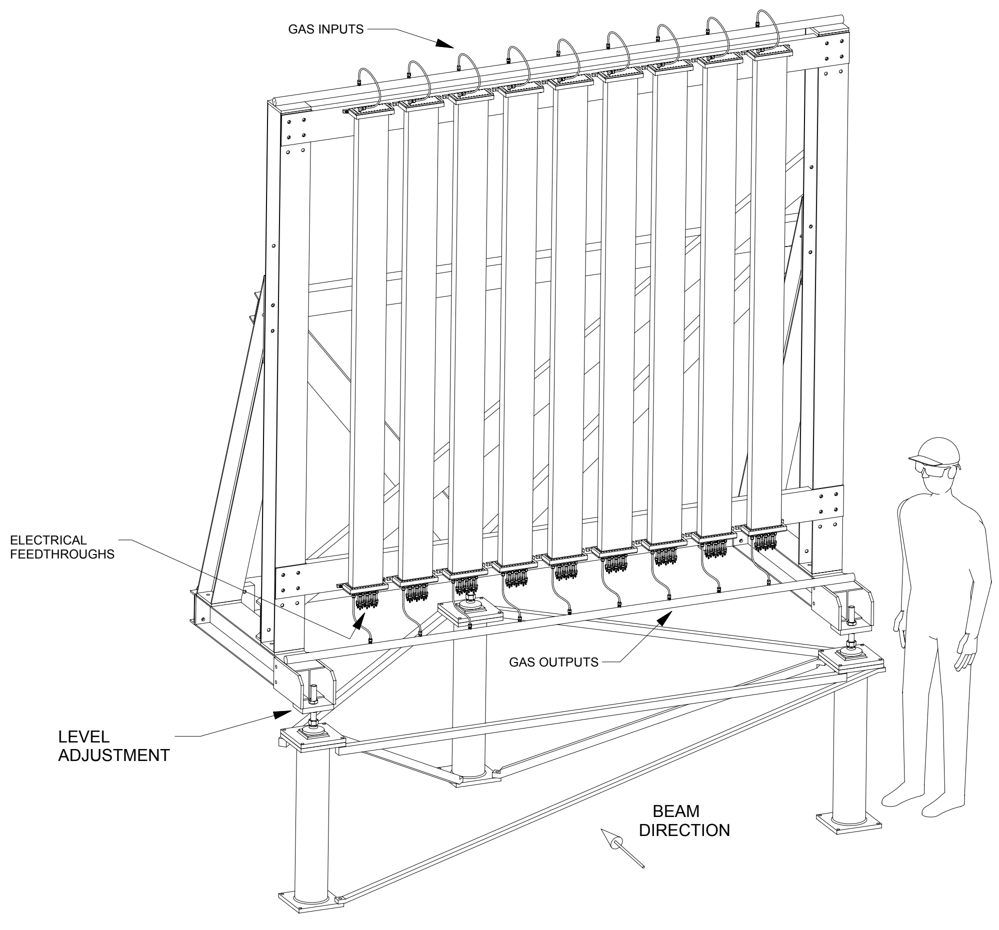

Figure 6.15: One of the three Muon Monitor arrays. There are nine ionization chambers in each of nine tubes.

6.7.2 Installation & Alignment

The muon tubes were mounted vertically on an aluminum support structure as shown in Fig. 6.15. Each tube was fastened via screws at the top and bottom to the support structure. Above and below the tubes are two 1" stainless steel gas distribution manifolds connected to each chamber with 1/4" stainless-steel tubing.

A potted nine-pin cable was attached to each tube, bring the signals to junction boxes. There, the signals were individually coupled to the conductors of 50-pin twisted flat cable - two cables for each alcove. This cable ran to the electronics, outside of the radiation area. The junction boxes were each potted; alcove 1 was potted with RTV, alcoves 2 and 3 were potted with wax.

The Muon Monitor positions and orientations were measured with respect to survey markers in the absorber hall and muon alcoves. These survey markers had only been surveyed to an accuracy of ~ 6 mm with respect to the network defined by the markers in the target hall.

The most important part of alignment, aligning the muon tubes on stand, was accomplished by adjusting the monitors so that the lower end-plates were within 1 mm of each other vertically. Also, the tubes were spaced horizontally at 10" to a precision of ~ 2 mm.

The support structures were each aligned to the survey marks in the hall. Using laser lines and levels, tape measures, right angles, and plumb bobs the results of Table 6.3 were obtained.

	Alcove 1	**Alcove 2**	**Alcove 3**
x (mm)	0.8	0.0	0.0
y (mm)	2.5	0.0	-1.3
yaw (mrad)	-7.0	2.4	0.0
pitch (mrad)	1.3	0.0	0.0
roll (mrad)	0.0	0.0	0.0

Table 6.3: Measured positions and orientation angles of the three Muon Monitor arrays with respect to the survey points. The uncertainty in the positions is $\pm$ 3 mm, and in the angles is $\pm$ 1 mrad. The correspondence between survey points in the upstream and downstream portions of the beamline is $\sim$ 6 mm.

6.8 Gas Delivery System

As mentioned above, the beam monitors operate with a continuous flow of helium gas. A diagram of the gas delivery system is shown in Fig. 6.16.

The helium source is an 8-pack of ultra-high purity (UHP) helium cylinders. The cylinders are changed regularly – typically once every three or four weeks. The cylinders connect to a stainless steel high-pressure manifold. A two-stage regulator controls the output pressure that is delivered from the cylinder – set to $\sim$ 60 psi.

After the regulator, a 1/4" stainless line of about 1000' length travels down the MINOS vertical shaft ($\sim$ 330') and up the absorber access tunnel to the equipment racks outside of the beam enclosure.

The input pressure to the rack is measured via a transducer to a precision of $\sim$ 1 psi. The line then splits into four – one for each beam monitoring station. Each line leads to a manual needle valve that controls the flow to each station. The flow and pressure are measured on the low-pressure lines via flowmeters and transducers that readout into the data acquisition[8]. Each line is fitted with a check valve: the Hadron Monitor is limited to 1 psi and

[8]The precision of these transducers is $\sim$ 1 torr; of these flowmeters: $\sim$ 0.1 L/hr.

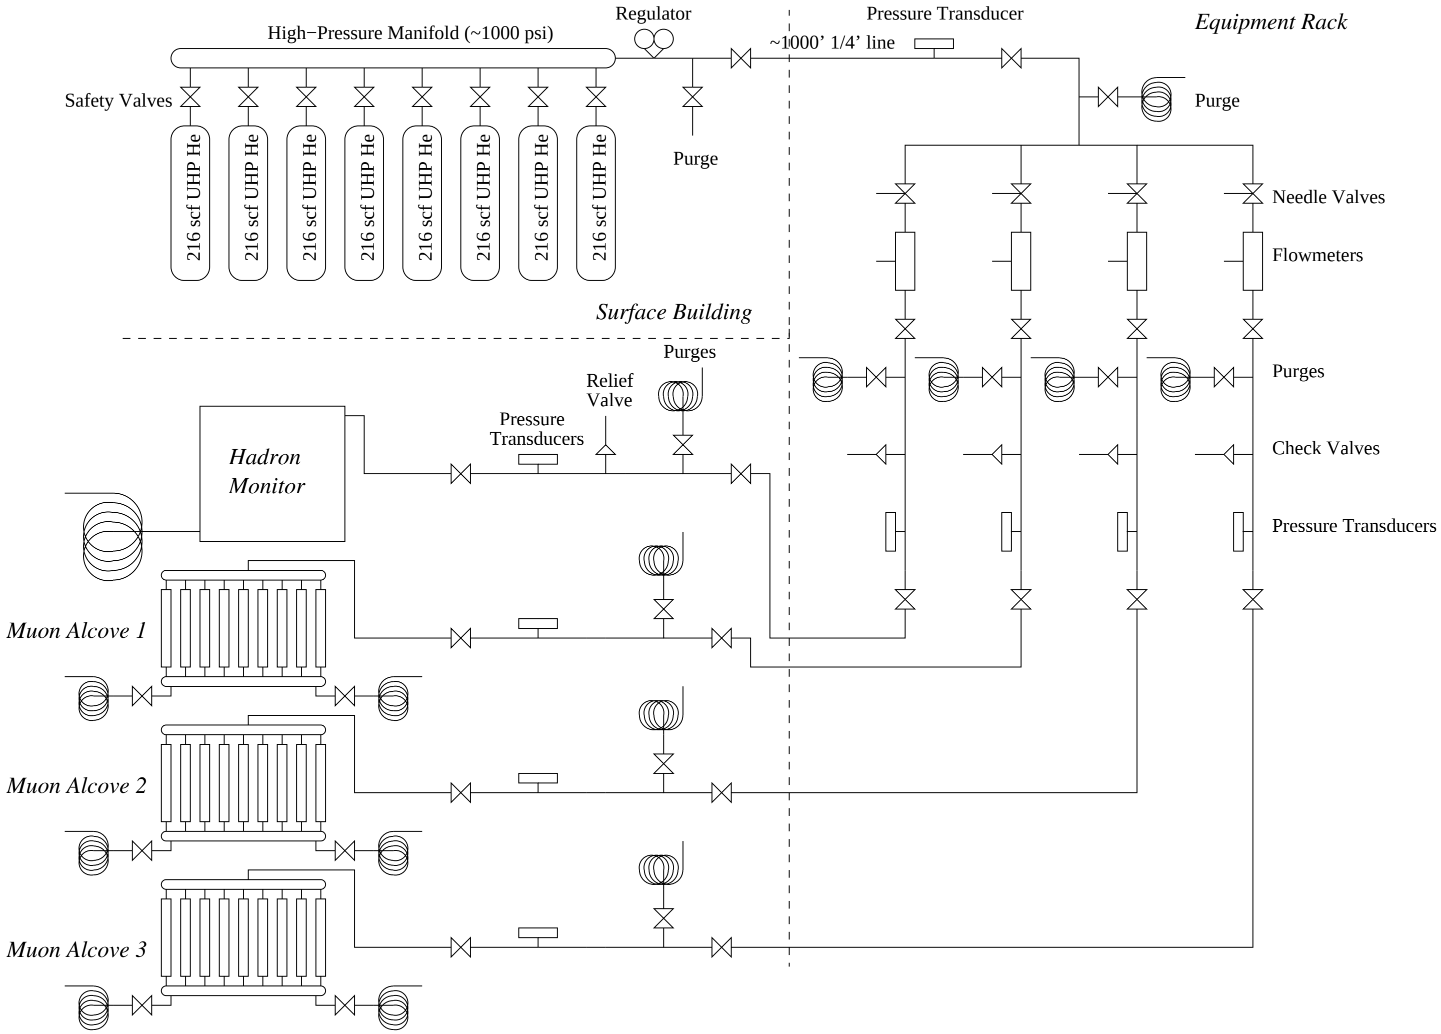

Figure 6.16: Diagram the Beam Monitoring gas system. It is broken into three parts: the cylinders above ground at the service building, the distribution and instrumentation at the equipment racks just outside of the radiation area, and final parts in the beam tunnel. The configuration shown is as installed – some components in the beamline were later removed. Detailed descriptions are in the text.

the Muon Monitors to 3 psi.

At each of the Hadron and Monitors there is another set of isolation valves, purge, and pressure transducer. The Hadron Monitor also has a relief valve which cracks at 1/3 psi. As anticipated, the transducers at the detectors failed due to radiation within one month of operation. Before this failure, the transducers at the equipment rack were calibrated to the detector transducers; these are used for long-term monitoring. There is also an RTD[9] attached to each of the detectors for temperature corrections.

To establish the operating flow of the monitors we measured the gas purity of the exhaust from each monitor as a function of flowrate. An Illinois Instruments Oxygen Analyzer was placed in line with the exhaust of each monitor. The oxygen analyzer has a significant impedance, so each test resulted in a larger overpressure than the flowrate would otherwise induce. The flowrate was adjusted and the analyzer allowed to come to equilibrium at a value.

The results of the flow tests are shown in Fig. 6.17. The Hadron Monitor oxygen level reaches a minimum at a flow of 40 l/hr. The minimum exists because the overpressure is inducing a leak at the check valve and perhaps other locations. The high impurity level is probably due to a known leak at one of the feedthroughs. An operating flow of 25 l/hr. was chosen for the Hadron Monitor and 10 l/hr. for each of the Muon Monitors.

6.9 High-Voltage System

The high-voltage for the Beam Monitors is generated with 16 “Droege” high-voltage (HV) supplies. The supplies were modified to output one-tenth of their

[9]Resistive Temperature Device. These are thought to be more radiation resistant than thermocouples.

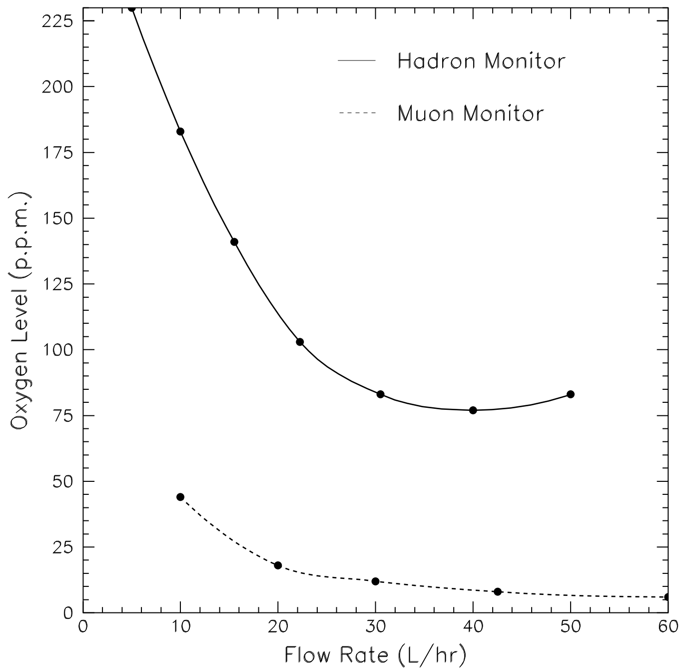

Figure 6.17: Measured Oxygen level at the outlet of the monitors as a function of flow rate. The upswing in the Hadron Monitor flow is thought to be caused by the relief valve, or other leaks, opening up with the increased pressure necessary to induce the flow.

nominal bias to allow better precision [116]. The HV supplies are located in the equipment rack outside of the radiation area. Each supply has a trip setting and is set to 0.1 mA.

The HV supplies are controlled with standard Fermilab CAMAC quad-DAC voltage controllers. The settings for the supplies can be adjusted through the Fermilab controls network, ACNET; the measured voltages are similarly recorded by the system.

The layout of the muon HV system is shown in Fig. 6.18. Each HV channel supplies HV to 27 chambers. The HV channel is split in the equipment rack to three cables – one each for a muon tube. The cable is terminated at the tube end-plate by soldering onto the filter and fan-out circuit board on each tube. Each circuit board has a 1 kΩ input resistor to reduce ground loop effects. After the input resistor the signals are fanned out to a 10 Hz low-pass filter and then to the HV feedthroughs on the muon end-plate. The filters are designed with large capacitance to minimize the voltage drop over the resistor due to current flow. The capacitors are in the muon radiation field, so they have ceramic dielectrics. Should one chamber fail such that it is shorted to ground the remaining chambers will remain biased, except that the short will act like a voltage divider.

The layout of the Hadron Monitor HV system is shown in Fig. 6.19. The filtering and fanout is entirely accomplished at the equipment rack. A large filter box takes HV input from seven supplies and outputs to 49 RG-58 cables. This allows individual channels to be disconnected at the rack in case of failure. Also, the capacitance of these filters is even larger to accommodate the greater current draw of the Hadron Monitor without significant voltage drop over the filter resistors.

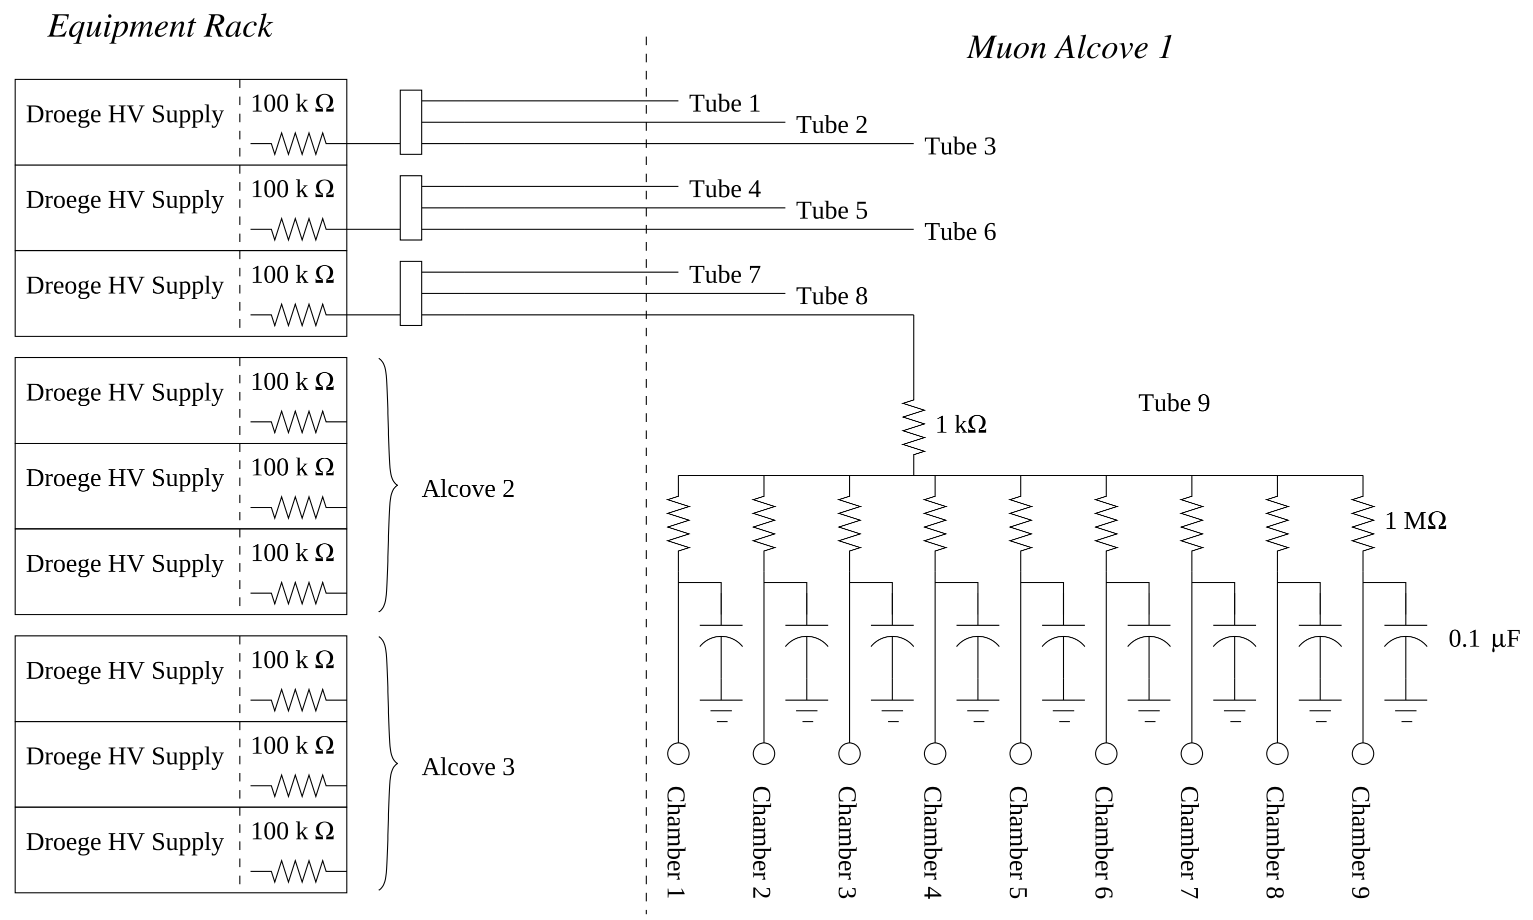

Figure 6.18: Layout of the High-Voltage system for the Muon Monitor. Nine HV supplies are used. Each supply has a 100 kΩ internal resistance. The output of each supply is fanned out to three RG-58 cables that run to the alcoves. Each cable runs to a single muon tube; at the tube the cable is soldered to a filter board attached to the feedthroughs. Each chamber has an independent 10 Hz filter.

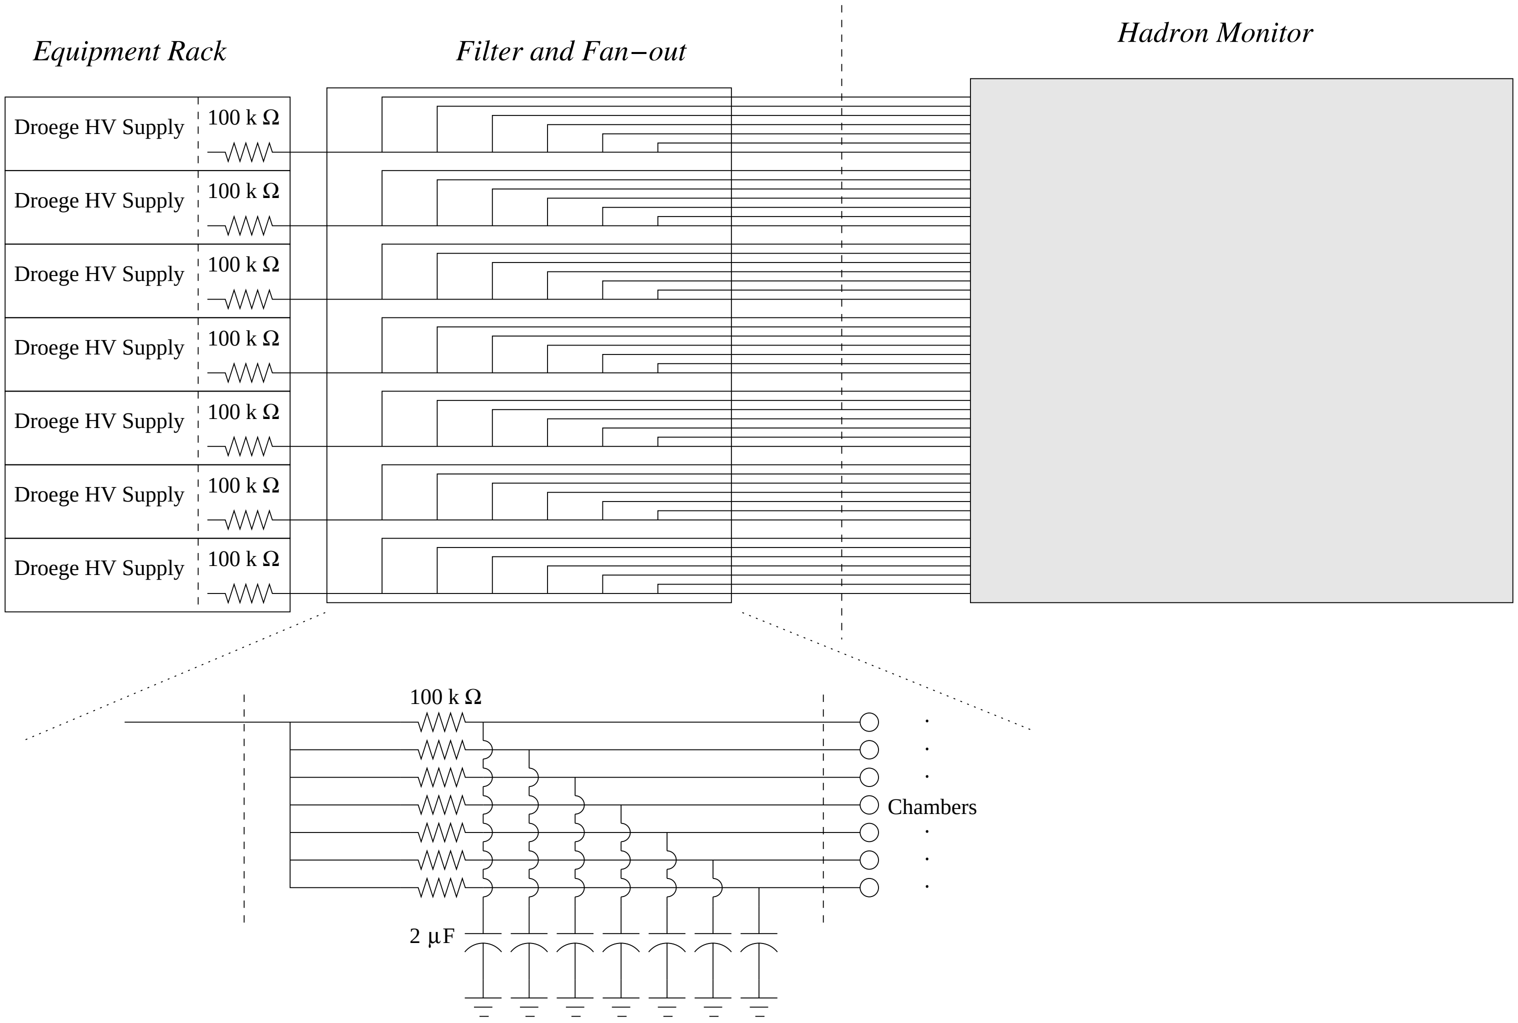

Figure 6.19: Layout of the High-Voltage system for the Hadron Monitor. Seven HV supplies are used. Each supply has a 100 kΩ internal resistance. The output of each supply is connected to a filter and fan-out box. Within the box, each HV channel is fanned out to seven lines, each independently filtered at 5 Hz. Forty-nine RG-58 lines exit the box and run to the transition box at the Hadron Monitor.

Chapter 7

Principles of Ionization Chambers

The NuMI Secondary Beam Monitoring System consists of Ionization Chambers (or Ion Chambers). The ion chambers are designed to measure the fluence of charged particles. The charged particles ionize helium gas within the chambers; the liberated ions and electrons are collected on a set of electrodes; measurements of the current are proportional to the fluence.

Maintaining the chambers' linear response to particle fluences is essential in using them as diagnostic tools, but is challenged by the intense fluences in the NuMI beamline. Such high fluences lead to the buildup of "space charge" of slow moving ions in the ion chamber, which cause nontrivial signal collection behavior at the highest intensities. We therefore developed a theoretical description of the charge flow in the ion chamber; the goal of this description is to establish the effective limits of operation of the chambers for which space charge does not significantly degrade their performance.

In order to develop the theoretical description, we found it necessary to

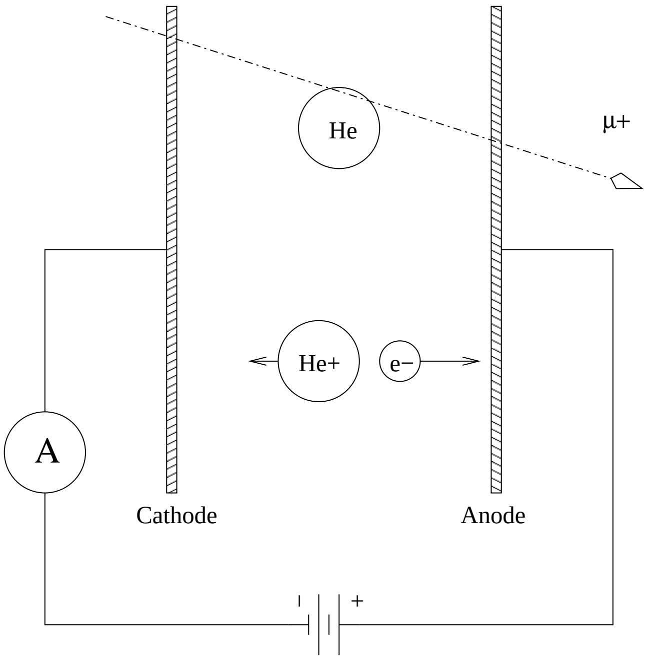

Figure 7.1: Schematic of an ion chamber. A bias is applied across two electrodes, separated by an ionization medium – in this case helium. An incident charged particle ionizes a helium atom into an He^{+} ion and an electron, known collectively as an ion-pair. The ion and electron drift to their respective electrodes, producing a current that is measured externally.

review the data on helium gas properties relevant to ion chambers, and the limitations on the state of knowledge of such parameters. This chapter introduces the basic operation of a parallel-plate ionization chamber with helium gas, and provides a survey of the relevant helium gas properties. The helium data will be used in Ch. 8 to create a simulation of pulse formation with the chamber that is compared to data from a beam test, and in Ch. 9 to interpret the performance of the ion chambers *in situ* in the NuMI beam. As will be shown in Ch. 8, there is a complicated interplay between charge recombination and gas multiplication in the chamber, which becomes highly nontrivial as the space charge build-up accrues with higher beam intensities.

7.1 Ion Chamber Introduction

A parallel-plate ionization chamber is illustrated in Fig. 7.1: an electric current is induced by the ionization of helium gas confined between two electrodes at different electric potentials. The ideal ion chamber will result in a current measurement that is equal to the amount of charge ionized by energy deposition from an energetic particle. For high-energy particles, the ionization per length is nearly uniform for different species and energies of particles, so simply counting the charge deposited gives a measure of the number of traversing particles. A complication is that different particle species may ionize different amounts of charge per particle. §7.2.1 describes the process through which the high-energy charged particles produce charge in the chamber.

The electric field causes the ions and electrons to drift to the electrodes, as described in §7.3 and §7.4. Fig. 7.2 is a conceptual plot of the ion chamber response variation with applied electric field. A low electric field results in slower charge drift such that some charge will not be collected on the electrode due to charges recombining in the gas (§7.6), or electron attachment on impurities (§7.7). As discussed in Ch. 8, and exemplified by the dashed curve in Fig. 7.2, the processes of recombination and attachment are exacerbated at higher ionization rate, causing larger electric fields to be required for full charge collection.

At very high electric field the drifting electrons will acquire greater velocities, resulting in charge multiplication (§7.5), where electrons gain enough kinetic energy from the electric field to ionize neutral atoms (or impurities). This effect is represented in Fig. 7.2 by the rise in collected charge at high fields. The response is much more sensitive to variation in temperature and pressure, so is not desired for operation.

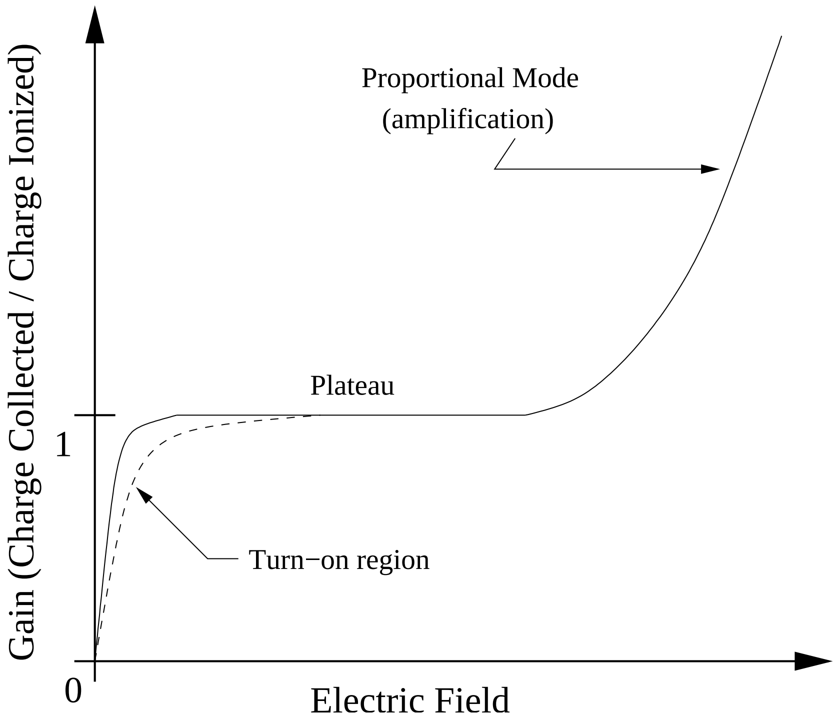

Figure 7.2: The response of an idealized ion chamber; the "gain" is plotted as a function of the applied electric field, where gain is defined as the ratio of charge collected on the electrodes to that ionized in the gas by the incident particles. The ion chamber is intended to be operated on the "plateau", where all of the ionized charge is collected at the electrodes, and variations in field have no effect. At low field, below the plateau, some portion of the ionized charge is not collected on the electrodes due to charge recombination. Such charge loss effects are dependent on the rate of ionization – the dashed line shows a greater turn-on region associated with higher ionization rates allowing more opportunity for recombination. At high field the ion chamber enters proportional mode in which the electrons gain enough energy from the electric field to ionize more atoms.

7.2 Ionization by Charged Particles

7.2.1 Description of Ionization

It is necessary to know how much ionization is induced per track length traversed in the gas volume. These ionizations result in a pair of electron and positive ion, but are commonly known as an "ion pair". The total number of ion pairs, s_t, created per cm of track length is inferred (see Wilkinson [117] pp. 8-11, and Knoll [118] p. 130) from the energy lost by the particle traversing the medium (dE/dx) and the average energy necessary to create an ion pair in the medium(w):

$$s_t = \frac{dE}{dx}(\chi, \xi, N, T) \Big/ w(\chi, \xi, N, T) \tag{7.1}$$

where T is the kinetic energy, N is the number density of the medium, χ refers to the species of the incident particles, and ξ describes the composition of the medium.

7.2.2 Energy Loss by Charged Particles

The values of dE/dx are described well by the Bethe-Bloch equation for particles of intermediate momenta. Taking the equation from [119] (p. 2):

$$-\left(\frac{1}{\rho}\right)\frac{dE}{dx} = Kz^2\frac{Z}{A}\frac{1}{\beta^2}\left[\frac{1}{2}\ln\frac{2m_ec^2\beta^2\gamma^2T_{max}}{I^2} - \beta^2 - \frac{\delta}{2}\right] \tag{7.2}$$

For an explanation of all the variables and constants see [119] (pp. 1-4). In this case dE/dx has the units of MeV/cm. The curves in Fig. 7.3 cover the useful range of the Bethe-Bloch equation for several materials, of atomic number Z, and incident particles of charge z.

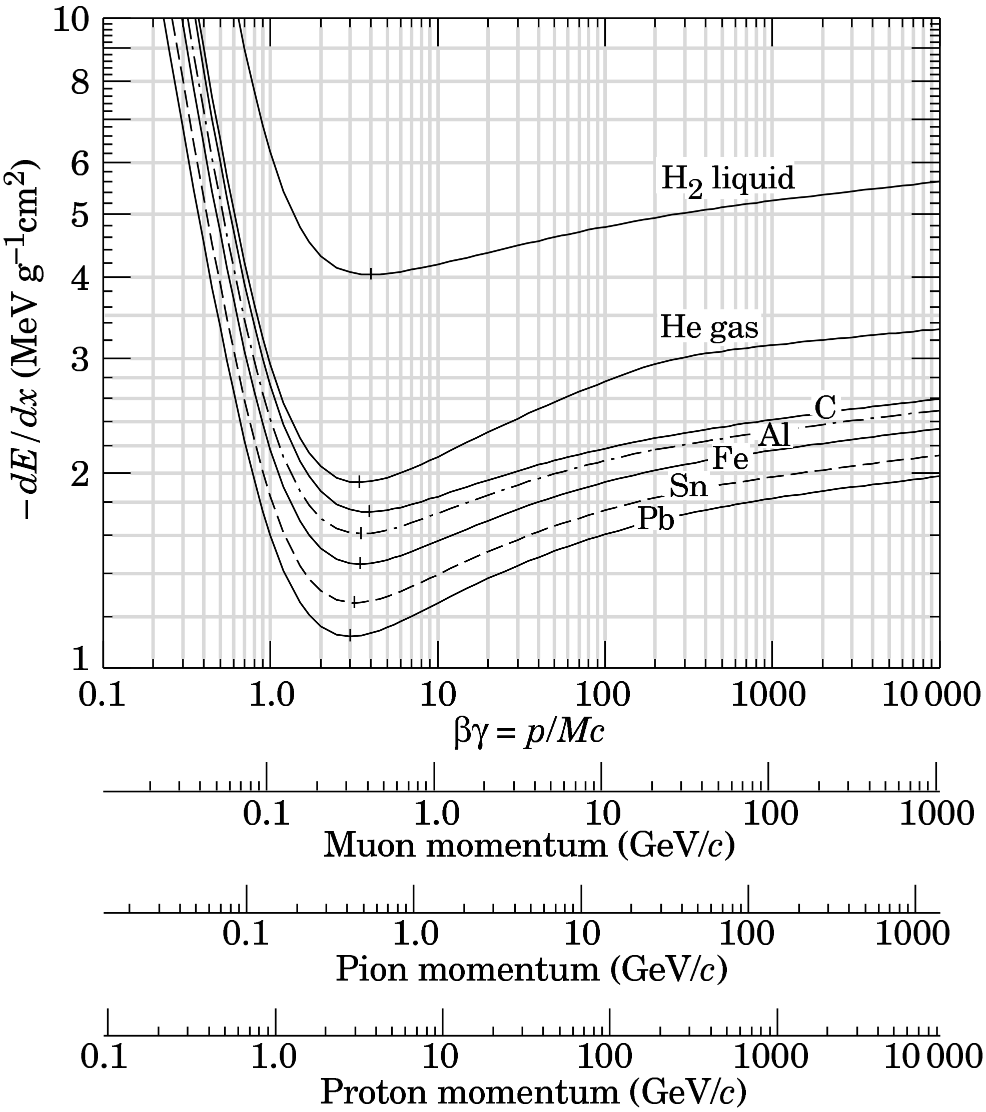

Figure 7.3: dE/dx curves due to ionization for various singly charged particles, taken from [119] (p. 4). This curve also applies for electrons as far as ionization losses are concerned, but this data is not an accurate measure of dE/dx if bremsstrahlung radiation is significant. Energy loss for electrons in helium is dominated by bremsstrahlung above 240 MeV, i.e. $\beta\gamma \approx 480$. The similar quantity for muons is $E_{\mu c} \approx 2300$ GeV, and higher for other particles.

Particle	Energy	w (eV)	Reference
p	3.6 MeV	45.2 ± 0.9	[121], cited in [120] (p. 20)
α	5 MeV	46.0	[122], cited in [120] (p. 23)
α	5 MeV	42.7	[123], cited in [120] (p. 23)
e^-	variable; β decay	42.3	[124], cited in [120] (p. 28)
e^-	variable; β decay	40.3	[125], cited in [120] (p. 28)
e^-	variable; γ irradiation	42.3	[126], cited in [120] (p. 28)
α	4.979 MeV	30.0	[127], cited in [117] (p. 20)
α	5.3 MeV	31.0	[128], cited in [129] (p. 227)

Table 7.1: Values of w as measured for several different sources in helium. We assume that charged mesons and muons will have values similar to that for protons.

7.2.3 Energy Required to Produce an Ion Pair

The variable w (in Eqn. (7.1)) describes the average energy necessary to create an ion pair. It is much higher than the ionization potential (24.6 eV for helium) as kinetic energy and excitations make up a significant portion of the energy transferred from the incident particle to the medium [120].

The sensitivity of w on impurity level (known as the Jesse Effect [123]) makes accurate measurement of this quantity difficult. The criterion used by the International Commission on Radiation Units and Measurements (ICRU) is to choose the measurements from literature with the highest values ([120] p. 23). Some of the cited values of w for helium are listed in Table 7.1. The values vary from 30 to 46 eV, but modern measurements suggest a higher range of 40 to 46 eV (probably due to trace impurities and higher pressures in older measurements). The value most relevant for our measurements will be $\sim$ 31 eV ([130], p. 42), because of impurity levels within the chambers.

7.3 Ion Mobility

An ionized helium atom will drift toward the cathode under the influence of the electric field. Because of collisions with other atoms, the speed by which an ion proceeds toward the cathode will be much less than its thermal (instantaneous) velocity.

The relationship between the ion drift velocity and electric field is:

$$v_I = \left(\frac{760\text{ mm Hg}}{P}\right)\left(\frac{T}{291\text{ K}}\right)\mu_I \times X \tag{7.3}$$

Here, v is the drift velocity of the ion, μ is the mobility[1], and X is the electric field. Typical units for v, μ, and X are cm/s, cm^2/(volt·s), and volt/cm.

Early experiments and calculations of helium's mobility were extensive (see [117], p. 28), but initially led to confusion (see [131], pp. 74-85) and many values of mobility that are unsuitable for application to the case of ionization chambers operated at atmospheric pressure. Experiments had assumed that the drifting ion was He^+; as was only shown by later investigations [132, 133] molecular ions (such as He_2^+) form at higher gas pressures via the collisional process: $He^+ + 2He \longrightarrow He_2^+ + He$, and will drift with a different mobility than the atomic ion (He^+). Some mobility data of the two ions are shown in Fig. 7.4 as a function of X/P.

The NuMI ionization chambers will run at atmospheric pressure, so the higher value of around 20 cm^2/(volt·s) can be used for the mobility.

7.4 Electron Drift Velocity

Some of the data on electron velocity are plotted in Fig. 7.5. For the lowest fields the drift velocity is proportional to the field. Here the electron's kinetic

[1]Denoted as K in some sources.

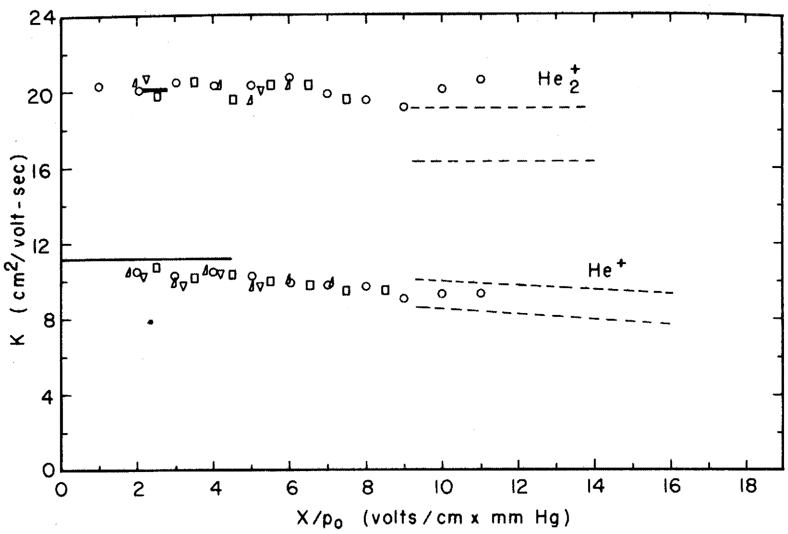

Figure 7.4: Loeb's summary of the data on He^+ and He_2^+ mobilities in helium gas [131], p. 102. The points are the measurements of Chanin [134], the dashes are the bounds on the measurements of Hornbeck [135], the short line represents Mitchell's measurements [136], and the long line is Massey's calculation [137].

energy is still less than thermal energy and a mobility could be defined as in Eqn. (7.3). Above a few V/cm the velocity curve starts to turn as the electron energy increases, and kinetic theory predicts a form of (e.g. see [117] pp. 31-5):

$$v = \kappa \times \left(\frac{X}{P}\right)^{\frac{1}{2}} \tag{7.4}$$

At the higher energy region the relation becomes, once again, linear. The prudent approach for implementing the electron drift velocity is to simply fit some function to the data, assuming a linear region below the lowest data. The function used in the simulations of Ch. 8 is shown as the line on Fig. 7.5.

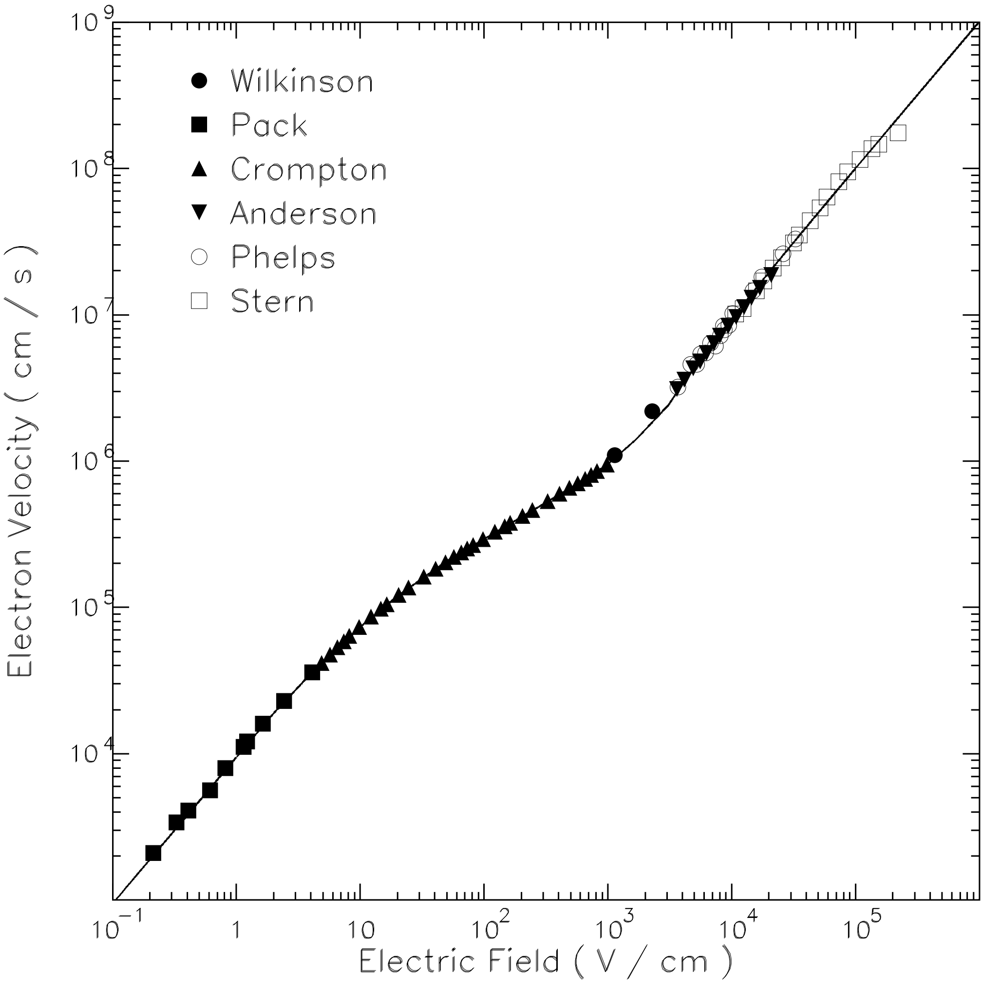

Figure 7.5: Collection of various sources of electron drift velocities in helium. Most of the data were taken from Dutton ([138], [139], [140], [141], and [142] as in [143], pp. 586 - 587, 608 - 609). Dutton's data is quoted in terms of E/N, so a value of $N = 2.69\times10^{19}$ cm^{-3} is used to convert to X. Wilkinson's data have been used to fill in the gap [117]; a value of $P = 760$ mm Hg was used in interpreting his data. The line shows the velocity used in the simulations of Ch. 8.

7.5 Gas Multiplication

As electrons drift through a gas under the influence of an electric field they gain kinetic energy. If the kinetic energy is high enough the electrons will ionize atoms through inelastic collisions. This process of charge multiplication continues until the charge is collected on the electrode. The amount of charge collected on the anode is related to the initial ionization by a proportionality constant

$$Q_{collected} = \mathcal{M} \times Q_{ionized} \tag{7.5}$$

This constant is related to the average number of electron-ion pairs liberated per length by a drifting electron

$$\mathcal{M} = \mathrm{e}^{\alpha x} \tag{7.6}$$

Where x is the distance traveled, and α is the number of ionizations per cm of path, known as the first Townsend Coefficient.

The first Townsend coefficient is usually defined effectively for a given gas, but ideally it should only depend on X/P (or X/N). A popular parametrization of α is that due to von Engel ([144] p. 138):

$$\frac{\alpha}{P} = A \exp\left[\frac{-B}{X/P}\right] \tag{7.7}$$

This form is only an empirical approximation. Von Engel states that, for helium gas, $A = 3$ ion pairs / (cm $\cdot$ torr), and $B = 25$ V / (cm $\cdot$ torr), and the relation for α/P above is valid over a range of $3 < X/P < 10$ V / (cm $\cdot$ torr). For the range of $20 < X/P < 150$, these values are $A = 3$ and $B = 34$. His parametrizations are reproduced in Fig. 7.6.

Impurities significantly affect the First Townsend Coefficient, particularly in helium. Penning ionization ([143], p. 706) further complicates the situation when X/P is high enough to excite a helium atom, but not ionize it. While experimental measurements of α for helium have been brought into agreement with theoretical predictions through the use of very high purity (p.p.b.) gas systems and low pressures [143], such data and calculations are not of great relevance for our chambers, which operate at moderate (p.p.m.) impurity levels. Further, our chambers operate at atmospheric pressure, where gas interactions become more complicated. Therefore, we will use the model of von Engel which was fit to older experimental data, taken with less pure helium gas and higher pressures.

Helium at atmospheric pressure can show significant deviation (see [131], pp. 700-4) from the idealized predictions for α due to excited states of helium which can cause further ionization via: $\mathrm{He}^* + \mathrm{He} \longrightarrow \mathrm{He}_2^+ + \mathrm{e}^-$ and $\mathrm{He}^m + \mathrm{He}^m \longrightarrow \mathrm{He}^+ + \mathrm{He} + \mathrm{e}^-$. Some calculations of total excitation in helium are reproduced in Fig. 7.6 as the open symbols. These data set an upper bound to the enhancement of the first Townsend Coefficient due to metastable excitations

Experimental work suggests that the effective values of α in our chambers lie somewhere between von Engel's parametrization and the excitation data. Our impurity levels are typically on the order of 1-100 p.p.m. and we use helium at atmospheric pressure. However, the values vary widely with impurity levels that are difficult precisely control and measure.

7.6 Recombination

In the NuMI ionization chambers there are high densities of electrons and ions present that must be collected on the electrodes. Some of the charge ionized will not be collected on the electrode, however, due to recombination of ionas

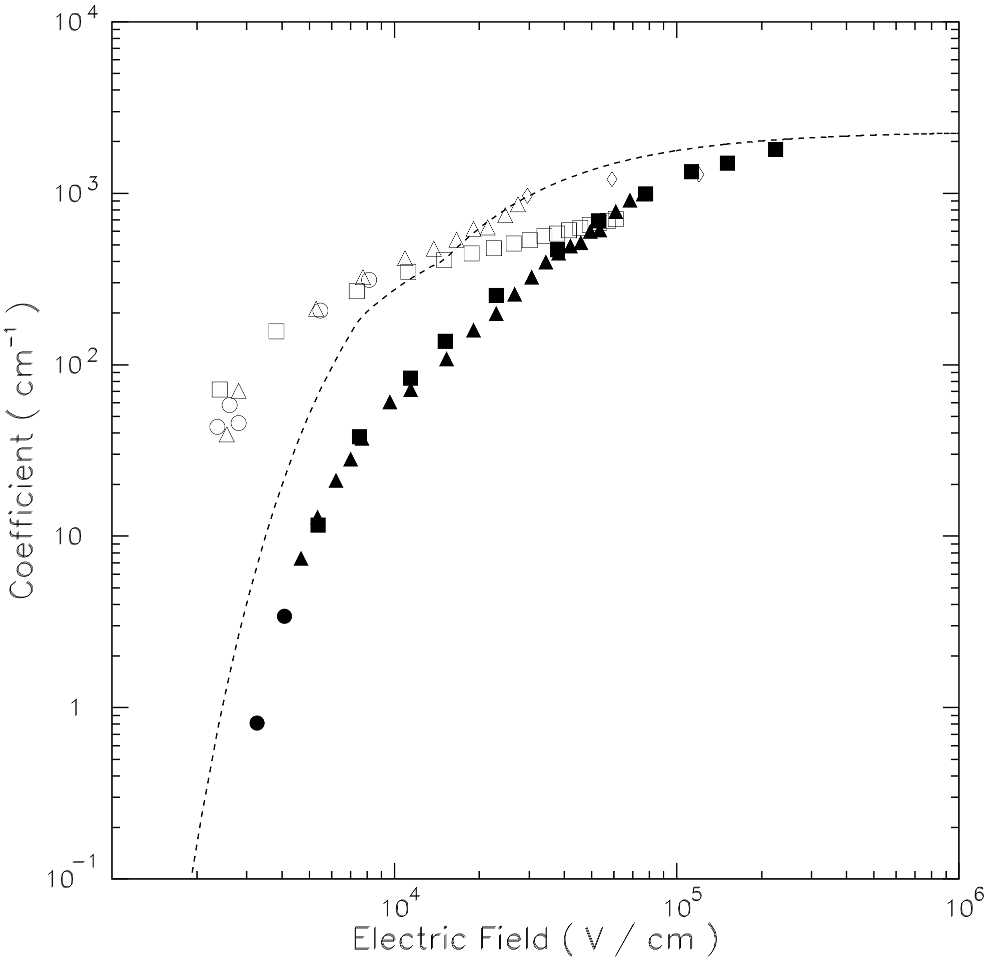

Figure 7.6: Data on ionization and excitation by electrons in helium. The line is von Engel's fit for multiplication ([144], p. 504) cited in Brown ([145], p. 140). The filled symbols are Dunlop's [146], Chanin's [147], and Davies's [148] data on multiplication as cited in Dutton ([143], pp. 711, 722); these data are probably correct for direct ionization. The open symbols are Corrigan's [149], Bortnik's [150], Hughes' [151], and Itoh's [152] calculations for total excitation in helium as cited in Dutton ([143], pp. 694-5, 704-5).

and electrons. A generalized form of two-body recombination can be written as

$$\frac{dn_-}{dt} = \frac{dn_+}{dt} = -r_2 n_- n_+ \tag{7.8}$$

where we have used Dutton's convention with r as the recombination coefficient. n_- and n_+ are the densities of the electrons and ions. A typical unit for r_2 is $\mathrm{cm}^3/(\mathrm{ions{\cdot}s})$. There are two dominant contributions for helium[2]:

Dissociative Recombination results when molecular ions in a gas capture an electron and dissociate into two neutral atoms: $\mathrm{He}_2^+ + \mathrm{e}^- \longrightarrow 2\mathrm{He}$. The value of Biondi and Brown [153] has been cited in Brown ([145] p. 195) and Loeb ([131] p. 560) as "probably exact". It is $r_2 = 1.7\times10^{-8}$ $\mathrm{cm}^3/(\mathrm{ions{\cdot}s})$. We expect dissociative recombination to be the dominant mechanism in our gas.

Three Body Recombination can occur in higher pressure gases where multiple particles can be in close proximity for a reaction: $\mathrm{He}_2^+ + \mathrm{e}^- + \mathrm{He} \longrightarrow 3\mathrm{He}$ or $\mathrm{He}^+ + \mathrm{e}^- + \mathrm{He} \longrightarrow 2\mathrm{He}$. As the process now depends on the number density of the gas it will have a different reaction rate

$$\frac{dn_-}{dt} = \frac{dn_+}{dt} = -r_3 n_- n_+ N \tag{7.9}$$

Where now r_3 is the three body recombination coefficient with typical units of $\mathrm{cm}^6/(\mathrm{ions{\cdot}s})$ and N is the number density of neutral atoms. At atmospheric pressure $N = 2.688 \times 10^{19}$ cm^{-3}. Massey gives a value for $r = r_3 N = 3.8 \times 10^{-9}$ At N.T.P. [3] ([154], pp. 634-5). Dutton describes some scaling with pressure ([143], pp. 733-4), such that $r_3 N = 7 \times 10^{-9}$.

[2]Radiative recombination: $\mathrm{He}^+ + \mathrm{e}^- \longrightarrow \mathrm{He} + \mathrm{h}\nu$ is found to be small ([145] pp. 192-3).
[3] 18° C and 760 Torr

Using the Dutton value for three-body recombination and the previous values for dissociative recombination we arrive at a total value for the effective recombination coefficient of 2.4×10^{-8} cm^3/(ions·s). From the above numbers and statements in Dutton there is at least an uncertainty of a factor of two in this coefficient.

7.7 Attachment

Trace impurities of electronegative gases can attach electrons to themselves. The primary method of attachment is radiative electron capture

$$\mathrm{X} + \mathrm{e}^- \longrightarrow \mathrm{X}^- + \mathrm{h}\nu \tag{7.10}$$

The likelihood of attachment can be a strong function of electron energy. A common ($\sim$ p.p.m.) contaminant in a commercial gas cylinder that has a high electron affinity is molecular Oxygen (O_2). Fig. 7.7 shows the probability h of attachment per atomic collision as a function of electron energy.

To estimate the magnitude of attachment losses we follow the method of Wilkinson ([117], pp. 39-42). If x is the partial pressure of the Oxygen in units of p.p.m., the mean free path of an electron is $6/\mathrm{x} \times 10^{-2}$ cm. The instantaneous velocity of electrons is $10^7\sqrt{\eta}$ cm/s, where η is as defined as the ratio of electron kinetic energy to the kinetic energy of the gas (see Fig. 7.8). The effective rate of charge loss in the gas is:

$$\xi = -\frac{1}{N}\frac{dN}{dl} = \frac{1.5 \times 10^5 x h \sqrt{\eta}}{v_d} \tag{7.11}$$

where v_d is the electron drift velocity in cm/s. The fractional loss δ across the path is:

$$\delta = 1 - \exp(-\xi l) \approx \xi l \qquad (\xi l \ll 1) \tag{7.12}$$

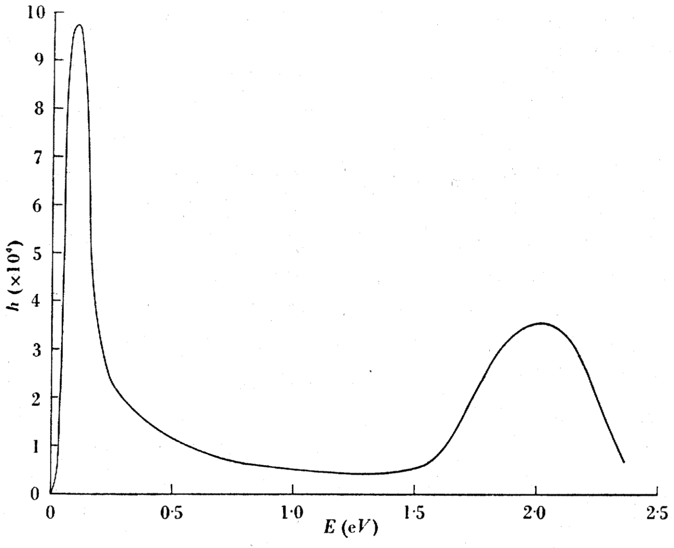

Figure 7.7: The probability of attachment (h) for electrons colliding with O_2 atoms as a function of the electron energy. Compiled from the data of Bradbury [155], Healy [156], and Bloche [157] as cited in [117], p. 41.

For a 5 mm chamber at 500 V, we set $v_d = 8 \times 10^5$ cm/s (see Fig. 7.5), $\eta = 65$ (see Fig. 7.8), and h $\approx 10^{-4}$. Requiring a maximum attachment of 1% over 5 mm, we need $x < 130$ p.p.m. Oxygen.

Another approach is to consider a region in the ion chamber where the electric field is near zero due to space-charge effects. Here, the region will persist for an amount of time and the drift velocity will be zero. Substituting:

$$\frac{dN}{dt} = v_d \frac{dN}{dl} \tag{7.13}$$

So that using the same variables we can restate Eqn. (7.11)

$$\xi_t = -\frac{1}{N}\frac{dN}{dt} = 1.5 \times 10^5 x h \sqrt{\eta} \tag{7.14}$$

And, Eqn. (7.12)

$$\delta = 1 - \exp(-\xi_t t) \approx \xi_t t \qquad (\xi_t t \ll 1) \tag{7.15}$$

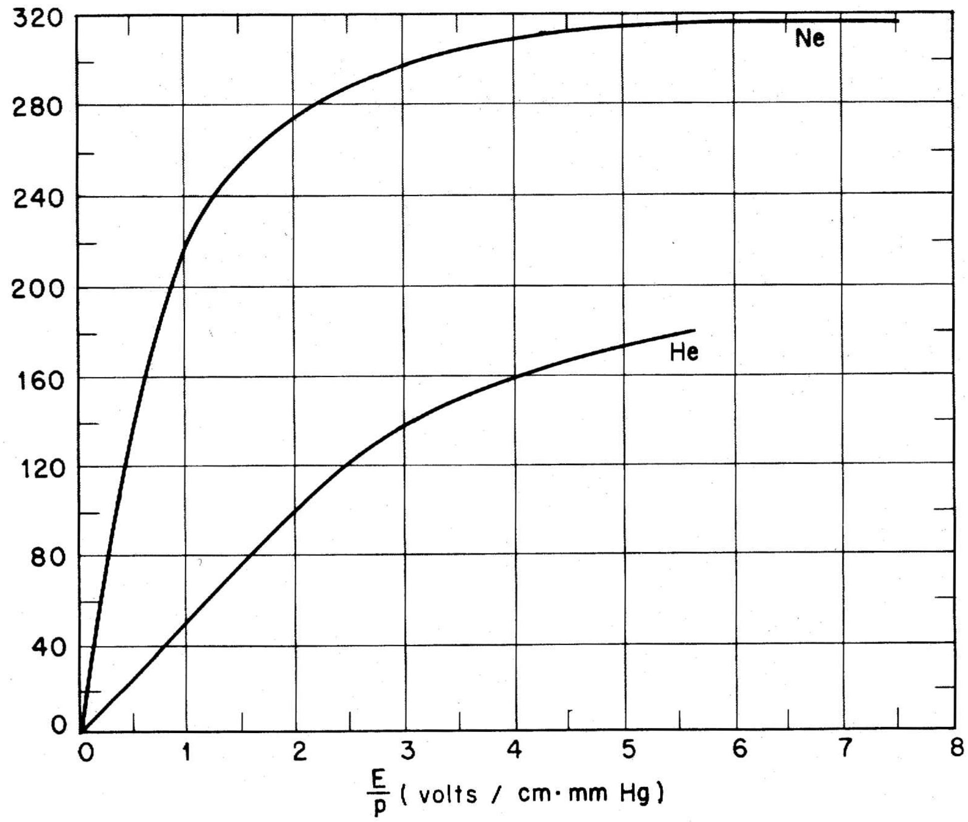

Figure 7.8: η (the ratio of average electron kinetic energy to that of the gas) as a function of E/P for helium and neon. Taken from [156] as cited in [129], p. 16.

We now consider a dead zone that will persist some number of μs. Considering the unknown electron temperature in a dead region, we consider two cases of $\eta = \{1, 4\}$. This corresponds to electron energies of $\epsilon = \{0.0235, .094\}$ eV, for which $h = \{10^{-4}, 10^{-3}\}$, providing a spread of ξ_t by a factor of 20. Requiring an attachment of at most 1% we need $xt < \{670, 33\}$ p.p.m.·μs. So, depending on the electron temperature, dead zones persisting for, say, 3 μs will require O_2 levels less than $\{220, 10\}$ p.p.m. In the Muon Monitors, the impurity levels are below these thresholds, but as discussed in §6.8 the Hadron Monitor O_2

level is around 70 p.p.m. so it may experience attachment effects.

7.8 Summary

The present chapter reviewed relevant data for charge creation and flow in helium gas. Such data are necessary to understand, in detail, the performance of our ion chambers.

Particularly important is to possess such data at a variety of electric fields. The electric field applied to the NuMI chambers is typically low and nominally uniform; however, the space charge created by slow moving ions also creates significant electric fields. As will be shown in Ch. 8, the total electric field within the NuMI chambers becomes distorted and spans several orders of magnitude – from near zero to several times the applied electric field.

Chapter 8

Pulse Formation in an Ionization Chamber

The ionization chambers for the NuMI beam monitors will be exposed to intense particle fluxes, $10^6 - 10^9/\mathrm{cm}^2/\mu\mathrm{s}$. As is well known, ion chambers exposed to such fluxes can exhibit distinctly nonlinear response due to the build-up of space charge from slow moving ions. Such nonlinear behavior, if present in our chambers, severely degrades their utility as beam monitors.

In order to understand the limitations to the ion chambers' useful range of application, the data from the previous chapter are utilized to develop an accurate description of the collected charge from an ionization chamber resulting from a particle beam pulse. Two limiting cases occur: in the first, the beam's duration is much longer than the time for an ion to traverse the chamber. In the second, the beam pulse may be approximated as instantaneous with respect to the ion collection time. App. C studies several limiting cases of chamber performance which may be described analytically. While not directly applicable to the case of the NuMI beam, they constitute important inputs and checks of the numerical simulations below.

In the case of the NuMI beam, neither approximation is valid. The beam pulse of 8.6 μs duration is comparable to the time for an ion to drift across the chamber gap. This necessitates a numerical simulation of the ion chamber. The present chapter will present results of such a numerical simulation obtained from a computer code described in App. D. The results of these numerical simulations are compared to a beam test performed at the Fermilab Booster RDF [158, 159]. As will be shown, the simulations and RDF data demonstrate that the ion chamber has a range of effective linear operation up to $3\times10^{10}/\text{cm}^2/\mu\text{s}$, approximately a factor of 10-20 greater than is required the NuMI beam.

8.1 Numerical Simulation

For the purposes of our calculations, we restrict ourselves to the case of a parallel-plate ion chamber where the ionization is uniform across the surface of the plates. This restriction is convenient as it simplifies the problem to 1-dimension; it is also close to reality as the ionization rate does not change over transverse length scales smaller than the gap size.

Even with the simplification of a 1-dimensional system the problem requires the solution of a set of coupled partial differential equations:

$$\frac{d\rho_+}{dt} = \Phi - v_+\frac{\partial\rho_+}{\partial x} - \rho_+\frac{\partial v_+}{\partial x} + \alpha\rho_- - r_e\rho_-\rho_+ + D\frac{\partial^2\rho_+}{\partial x^2} \tag{8.1a}$$

$$\frac{d\rho_-}{dt} = \Phi - v_-\frac{\partial\rho_-}{\partial x} - \rho_-\frac{\partial v_-}{\partial x} + \alpha\rho_- - r_e\rho_-\rho_+ + D\frac{\partial^2\rho_-}{\partial x^2} \tag{8.1b}$$

$$\frac{\partial E}{\partial x} = \frac{e}{\epsilon_0}(\rho_+ - \rho_-) \tag{8.1c}$$

$$\int_0^d dxE = V_{app} \tag{8.1d}$$

The dynamical variables here are the densities of ions, $\rho_+(x,t)$, and electrons, $\rho_-(x,t)$; the electric field, E, is determined entirely by the ion densities in Eqn. (8.1c) and the boundary condition in Eqn. (8.1d) where V_{app} is the applied bias voltage[1] to the chamber of width d. The beam ionization, $\Phi(x,t)$, is an arbitrary input; generally, we consider it to be uniform in space and either a square-wave pulse or constant in time. The second and third terms in Eqns. (8.1a) and (8.1b) arise from the continuity equation; the velocities, $v_+(E)$ (§7.3) and $v_-(E)$ (§7.4), provide coupling through their dependence[2] on E. Multiplication (§7.5) is provided through the First Townsend Coefficient, $\alpha(E)$, which has a very nontrivial dependence on electric field. Recombination (§7.6) is provided through an effective recombination coefficient, (r_e), that is constant, but the term introduces more coupling. Diffusion requires the coefficient of diffusion, D, and involves second derivatives.

There is clearly no general solution for the above system. We constructed a numerical simulation of ion chamber dynamics to investigate the response behavior of the ion chambers which extends beyond the approximations examined in the previous section. Using the data on Helium reviewed in Ch. 7, we can fully simulate the equations in Eqn. (8.1)[3]. The simulation treats the chamber as uniform in the dimensions transverse to the electric field, as in the above calculations. The dimension parallel to the field is simulated through finite-element cells, each containing a uniform density[4] of electron or ion charge.

[1] We generally assume that the electrodes can exchange collected charge to and from sinks of charge to maintain the potential difference; we examine this assumption in §D.5.

[2] v_+ is linearly related with E, while v_- varies nonlinearly.

[3] Diffusion is only simulated incidentally, as discussed in §D.4. Diffusion's effect on chamber performance is negligible.

[4] To allow exact implementation of average velocity and charge conservation, densities are used instead of quantities.

The finite-element analysis undertaken utilized sufficiently small cell sizes and time increments such that no advanced methods were used for solving field equations or performing propagation. App. D describes detailed steps of the numerical simulation and execution. The remainder of this chapter sketches the technique and results.

The numerical simulation is used to visualize the ion chamber dynamics and to calculates the response over a range of possible physical parameters. In this section we concentrate on the 1.56 μs beam pulses relevant for the beam-test results of §8.2.

8.1.1 Field and Charge Distribution Evolution

We start by considering the charge and electric field distributions within the chamber during and after the ionization period. For this exercise we only consider charge transport and space charge – multiplication and recombination are not in effect.

Fig. 8.1 shows the charge and field distributions for a 1 mm chamber biased at 200 V and ionized at a rate of 10^{10} ionizations / cm^3 / μs for 1.56 μs. This rate is corresponds tp the start of space charge becoming appreciable within the chamber, and corresponds to a fluence of $\sim 10^{10}$ charged particles per square centimeter[5]. The electron density quickly gains a mostly linear distribution, as described in §C.1. The ion density is much larger due to its lower drift velocity; it begins to form a linear distribution, but the ionization ends before it is fully formed. The net charge distribution $(\rho_+ - \rho_-)$ is approximately that of the ions. The space charge is enough to cause some distortion

[5]The ratio between flux and ionization rate depends on the average energy to create an ion pair, w, as discussed in §7.2.3. In the simulations the question of a precise value for w is avoided by considering only the ionization rate.

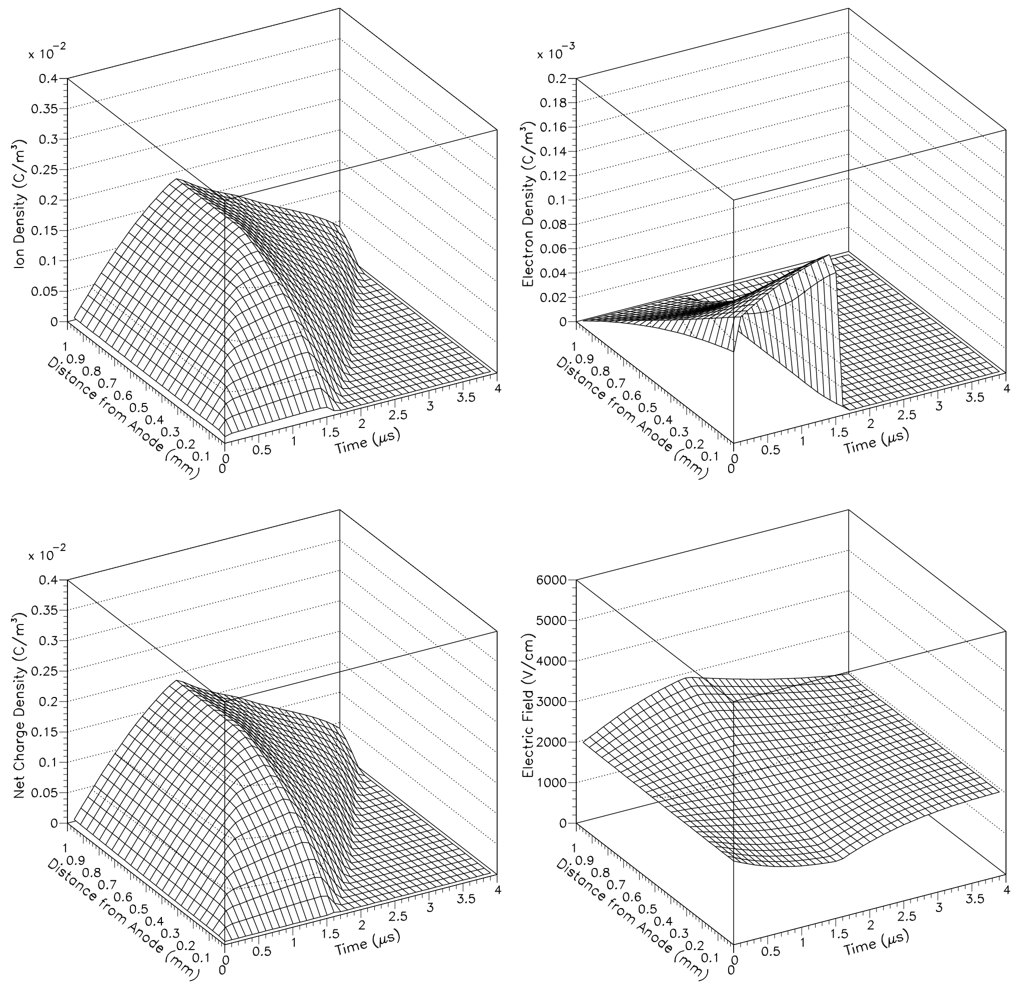

Figure 8.1: Field and charge distributions for a 1 mm chamber ionized at a rate of 10^{10} ionizations / cm^3 / μs for 1.56 μs and biased at 200 V. The horizontal axes on each plot consist of the location in the chamber and the time after the beam pulse begins: the pulse lasts for first 1.56 μs; the next 2.5 μs shows the charge leaving the chamber. Upper left: positive ion density in the chamber builds up during the pulse and then leaks out after. Upper right: electron density in the chamber builds up and leaves the chamber immediately upon the end of the pulse. Lower left: net charge density – the difference of ion and electron densities – is almost equal to the ion density in this case. Lower right: the electric field within the chamber, determined from the applied field of 2000 V/cm and the space charge, is slightly distorted from a uniform field.

of the electric field such that the field at the anode reaches almost 150% of the applied electric field. The chamber is mostly clear of charge by 4 μs, i.e. 2.5 μs after the end of the beam pulse. Such is consistent with the time expected for He_2^+ ions to drift 1 mm in a field of 2000 V/cm.

An ionization rate of 10^{11} ionizations / cm^3 / μs results in the effects shown in Fig. 8.2. The linear distribution exists only briefly for the electrons. Instead, the excess ion density becomes enough to fully screen the electric field and create a dead region. The dead region is fully formed by 0.5 μs; the ion density and electric field distributions then approach that of Eqns. (C.13) and (C.14), and plotted in Fig. C.3. The ion density increases linearly within the dead zone as ionization conitinues, buts stays finite because of the limited ionization period. The electron density is almost identical to the ion density within the dead region, and very small without. The net charge density is that of the ions outside of the dead zone and zero inside. After ionization ends the chamber transitions to the situation of Eqn. (C.21), and illustrated in Fig. C.4; the dead zone contracts as ions leak out of it. The electron density remains equal to the ion density within the dead zone; the net density is dominated by a ridge at the border of the dead zone where $x^{-1/2}$ behavior of the ion density causes a peak, before it is canceled by the electron density. The jagged peaks in the plot are caused by aliasing the 10,000 cells in the simulation to 40 for the visualization; the peak moves smoothly within the simulation. The plots are truncated at 4 μs, but the dead zone persists for a few more μs afterward. Once the dead zone is depleted the remaining charge leaves the chamber in $\sim$ 2.5 μs.

An intensity of $3{\times}10^{11}$ ionizations / cm^3 / μs brings us to Fig. 8.3. The behavior of the chamber is similar to that of Fig. 8.2, but the ion and electrons densities in the dead zone area each three times higher, the dead zone takes

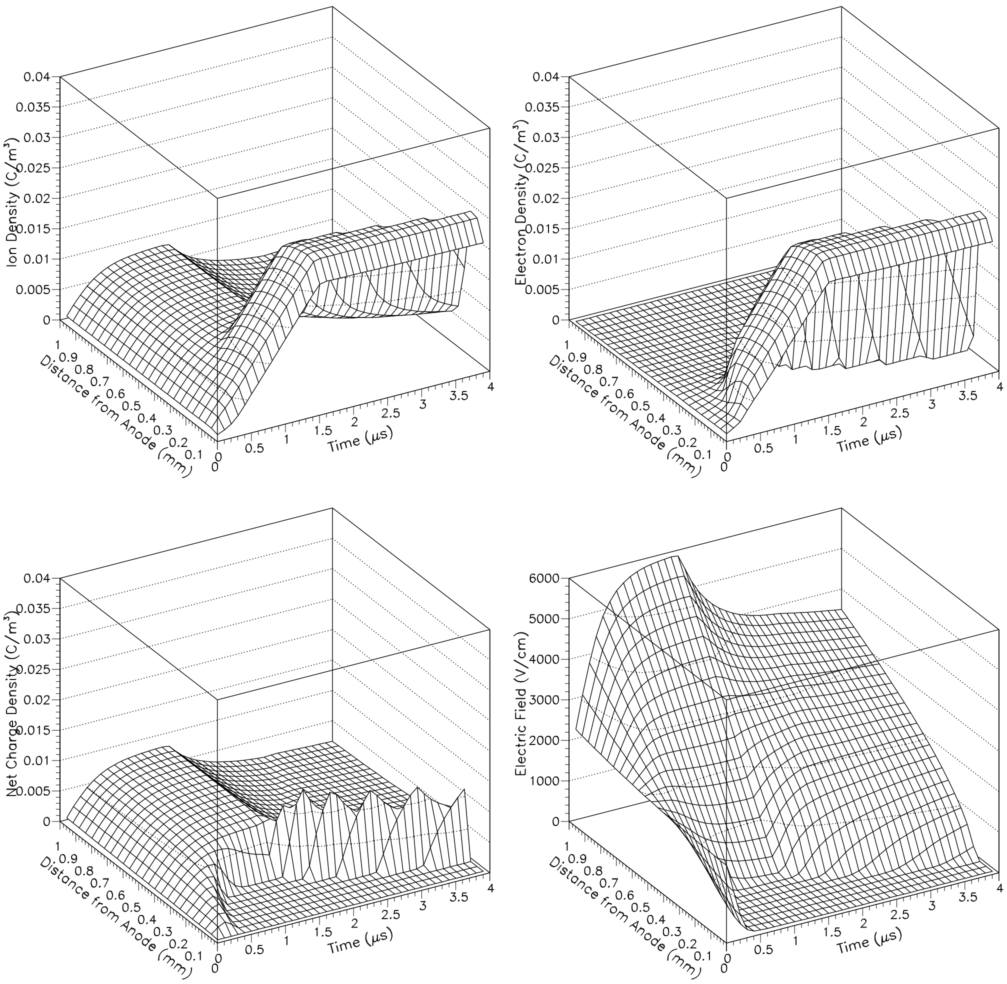

Figure 8.2: Same distributions as in Fig. 8.1, except for a chamber ionized at the higher rate of 10^{11} ionizations / cm^3 / μs. The excess charge distorts the electric field substantially, such that a region of no electric field is formed. Electron and ion densities increase quickly within the dead zone from ionization. Charge leaks out of the dead zone after ionization ends. The ridge in the net charge is aliased by the plotting.

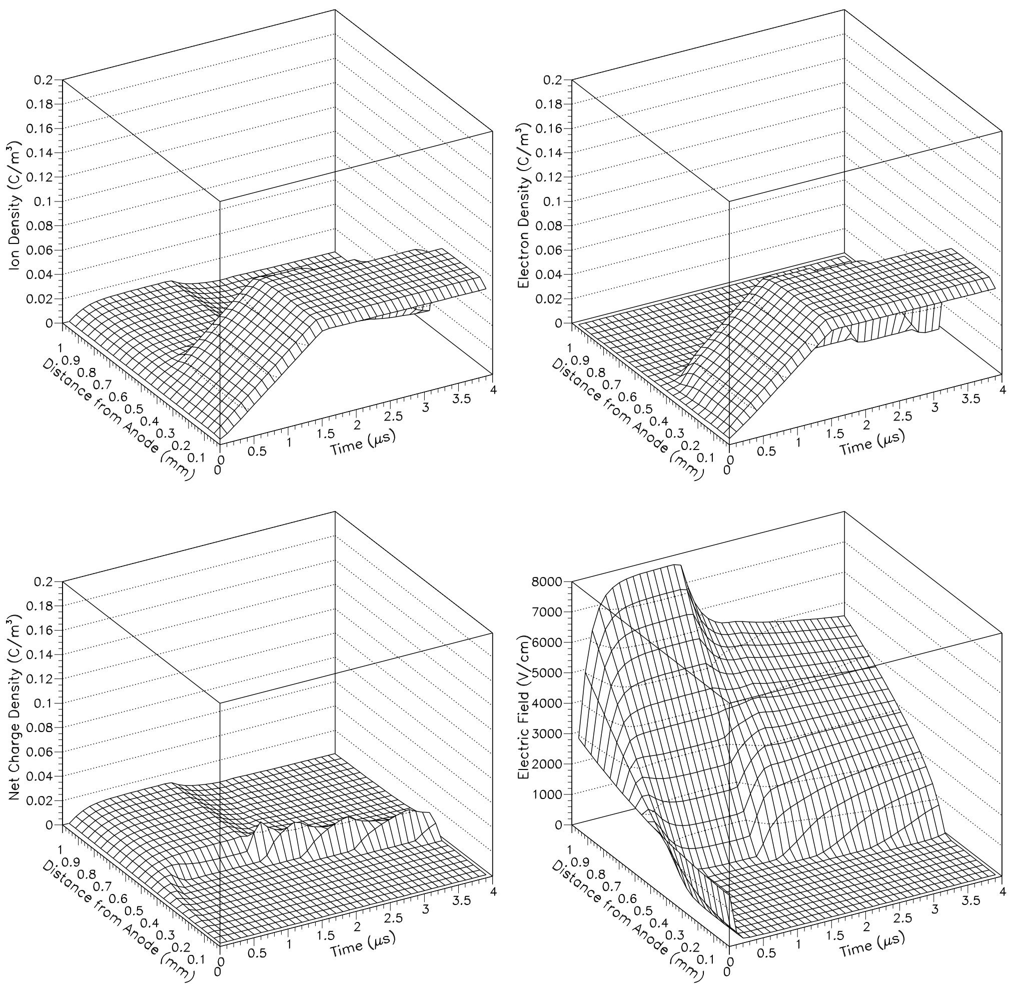

Figure 8.3: Same as Fig. 8.1, except ionized at a rate of 3×10^{11} ionizations / cm^3 / μs. The dead zone extent, charge build-up, maximum electric field, and leakage time are all larger than in Fig. 8.2.

up half of the chamber, and it will evidently last much longer. The higher densities sequestered in the dead zone for a greater amount of time are what causes the recombination loss to scale quickly with ionization rate. Note also that the maximum electric field now reaches four times the applied electric field – which will contribute to gas multiplication.

8.1.2 Space Charge-Enhanced Multiplication

From the preceding calculations and discussion we expect that multiplication will be an issue for ion chamber response even when the applied electric field would not cause significant multiplication. Fig. 8.4 shows a chamber with the same configuration as Fig. 8.3, except that multiplication is enabled within the simulation.

The multiplication results in an increase of ion density near the anode during ionization. The electron density is somewhat higher as well, but is still negligible compared to the density in the the dead zone. For the purposes of space-charge screening, ionization through multiplication is equivalent to beam ionization, except for the location of the produced charge. A higher ionization rate causes further compression of the active region and a higher peak electric field. Note that multiplication ends as ionization ends – while elevated peak electric fields persist, there are no electrons in the active region to cause multiplication.

8.1.3 Chamber Response: Linearity & Plateau

This section employs the full simulation, including charge drift, multiplication, and recombination. The results are derived for a 1.56 μs beam spill, of relevance for the test beam data discussed later in this chapter. Fig. 8.5 shows

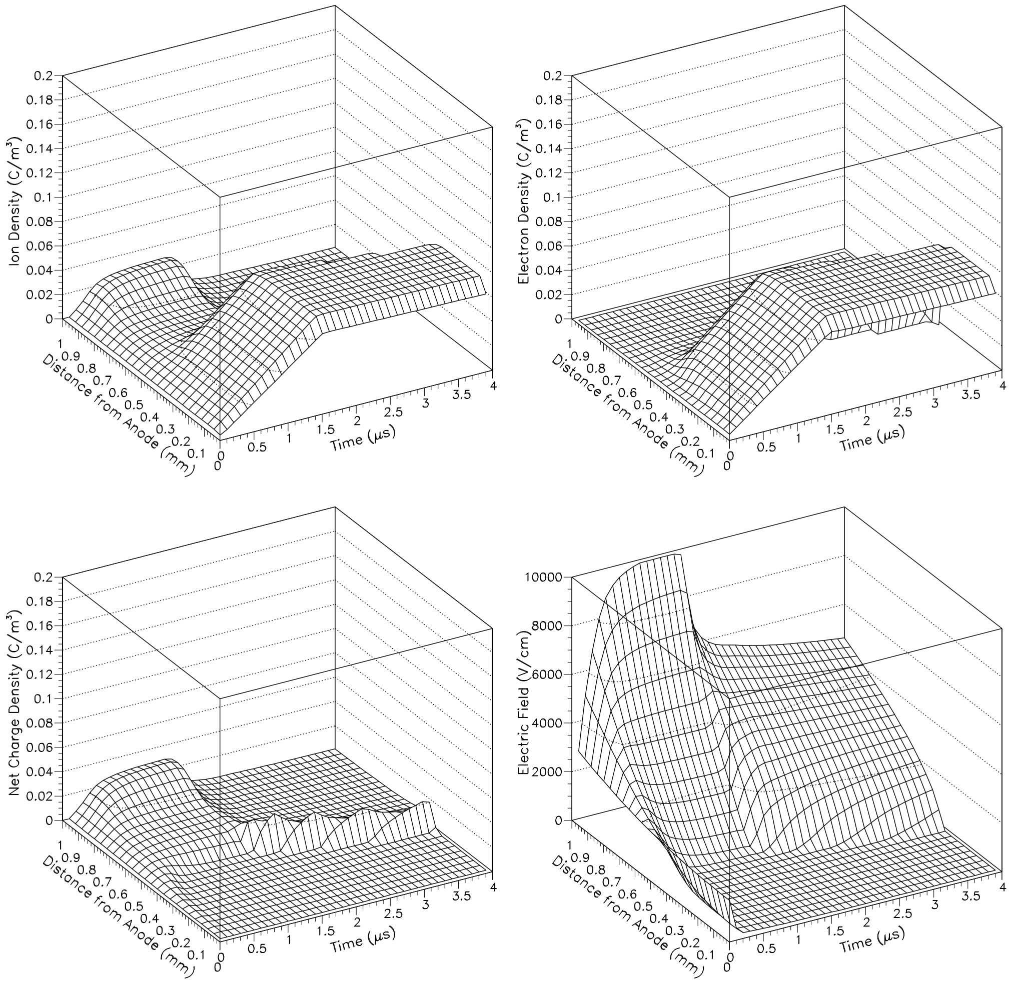

Figure 8.4: Same as in Fig. 8.1, except ionized at a rate of 3×10^{11} ionizations / cm^3 / μs, and the effects of multiplication are included. The dynamics are similar to that of Fig. 8.3, except for an increase in ion density near the cathode during ionization and the higher peak electric field.

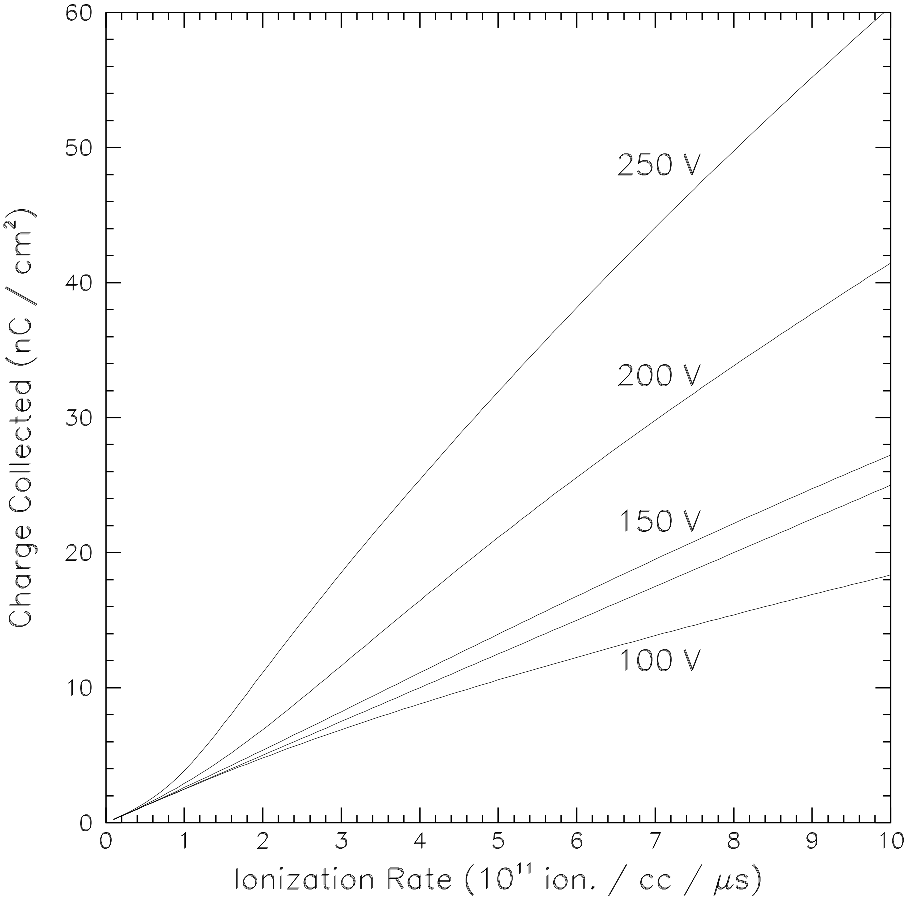

Figure 8.5: Results of simulation for a 1 mm chamber at various bias voltages. The total charge collected is plotted as a function of the ionization rate; a straight line is drawn where the total charge collected would be equal to that ionized within the chamber by the beam. Nominal parameters for multiplication and recombination are used.

the results of simulations at four different bias voltages where the ionization rate was varied up to 10^{12} ionizations / cm^3 / μs. We find that at the lower intensities all the bias voltages result in roughly linear response, with near total collection efficiency. At high rates the 100 V bias shows decided loss due to recombination. The behavior of the higher biases, however, is more complicated. An increase of collected charge is caused by space charge induced gas multiplication. At higher intensities the long duration of sequestered charge in the chamber results in some charge loss due to recombination which evidently counteracts the amplification due to space charge.

The situation is shown more clearly in Fig. 8.6 where the charge collected is normalized by the charge ionized by the beam – resulting in a "gain". While the amplification due to the applied field from 250 V bias should only be a few percent, the space charge increases this to 150% at 3.5×10^{11} ionizations / cm^3 / μs. We can also compare the recombination loss at 100 V to the analytical prediction plotted in Fig. C.5. We see that both depart from full collection efficiency at about the same rate of 10^{11}, but that the simulation loss does not drop off as quickly; this result is because the analytic calculation does not consider charge that was never captured in the dead zone.

Fig. 8.7 shows the response as a function of bias voltage for several intensities. The behavior at biases less than 100 V is similar to that calculated for Fig. C.5, but the loss is smaller in magnitude as discussed above. The lowest intensity curve at 10^{10} gives a good baseline for behavior – it has no loss down to 20 V and is only starting to experience multiplication at 225 V. Increasing intensity to 4×10^{10} or 10^{11} results is some loss below 100 V, but drastically affects multiplication, increasing it by a factor of 10 at 300 V. At the highest intensity simulation, 10^{12}, both effects are enhanced, but recombination is starting to dominate, such that gain is decreasing with intensity above 250 V.

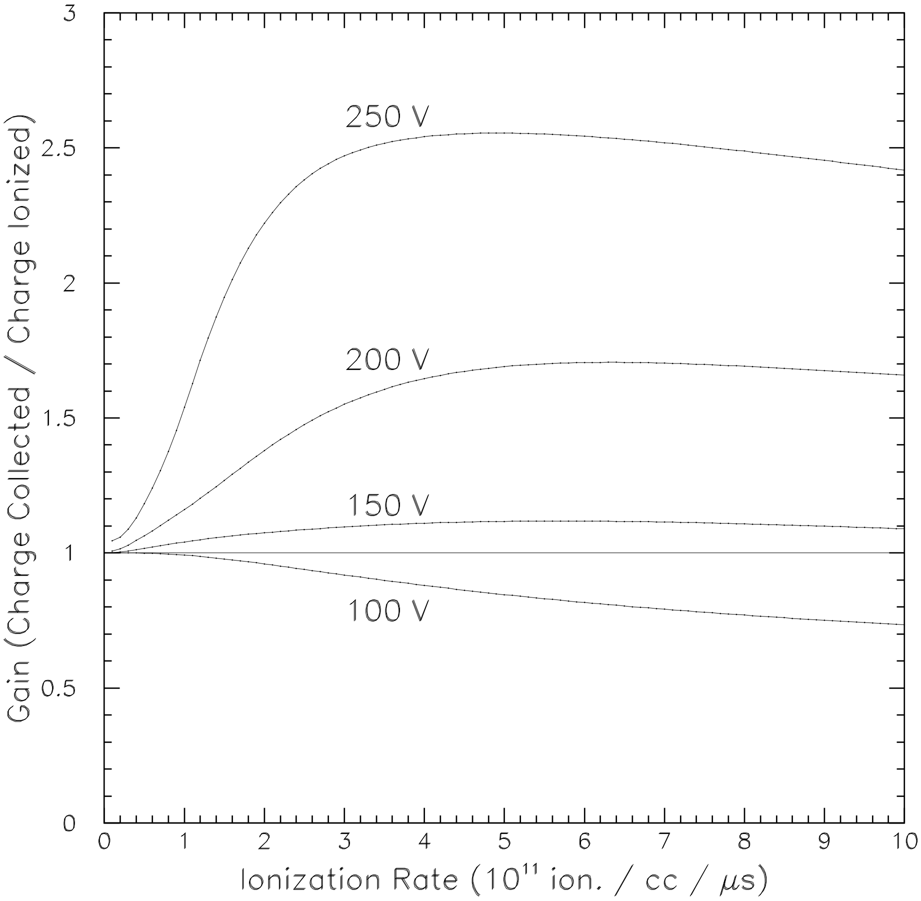

Figure 8.6: Results of simulation for a 1 mm chamber at various bias voltages. The "gain" is plotted as a function of the ionization rate. Nominal parameters for multiplication and recombination are used. At low intensity with no space-charge effects the multiplication at 250 V is only a few percent, but is more than a factor of two at high ionization rates.

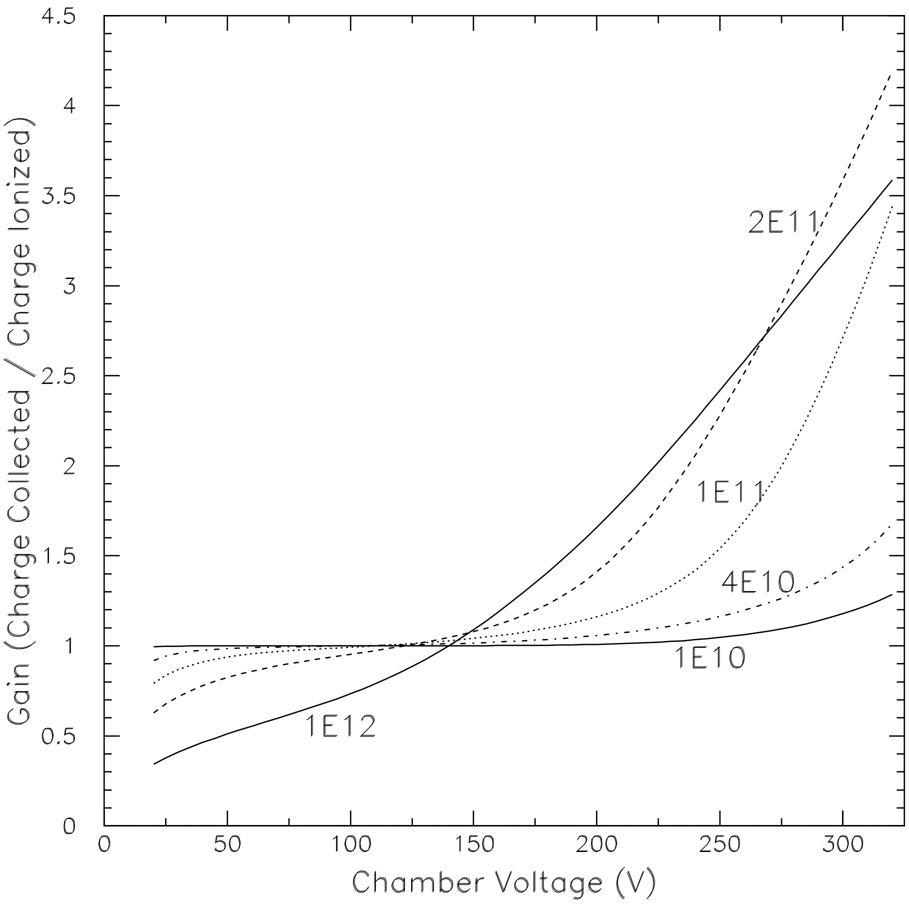

Figure 8.7: Results of simulation for a 1 mm chamber at various ionization rates (in units of ionizations / cm^3 / μs). The "gain" is plotted as a function of the bias voltage. Nominal parameters for multiplication and recombination are used. Note that increased intensity increases recombination (decreases signal) at low biases, and increases multiplication (and signal) at high biases. Several curves cross each other at approximately the same voltage of 130 V – response would be mostly linear at this bias.

Qualitatively, we notice a point near 130 V where all of the curves converge. At this "crossing point" the response will be most linear with intensity.

Simulations of 2 mm chambers were also carried out and are discussed in App. E.

8.2 Comparison to Beam Test Data

We have performed beam tests of 1 mm and 2 mm gap ionization chambers in helium and helium-hydrogen gases at the 8 GeV Fermilab Booster accelerator. The chambers were exposed to proton fluences of $2-100\times10^{10}$ per 1.56 μsec spill – to greater intensities and shorter pulse durations than expected in the NuMI beam. Details of the beam test can be found in [158, 159].

We took three intensity scans with helium and four with helium-hydrogen flowing through the chambers in which the beam intensity was varied from 1.5×10^{10} to 1×10^{12} protons per spill. The chambers' bias voltages were set to 100, 150, 200, or 250 V. In Fig. 8.8 we plot the output signal for the 1 mm and 2 mm chambers versus the 8 GeV proton beam intensity (measured with a beam toroid). The beam spot size was about 5 cm^2.

The data for the 1 mm chamber operated at 200 V were fit over the range of $0-4\times10^{10}$ protons/cm^2/spill (see Fig. 8.9). The fits indicate the chamber response is reasonably linear with respect to ionization rate.

Recombination results in the curves of Figs. 8.9 and 8.10 falling below a straight line fit. Gas amplification is also present in these data, as evidenced particularly by the scans taken at 250 V. The interplay between amplification gain and recombination loss caused by space charge is perhaps best seen in Fig. 8.10. The figure shows the collected charge divided by the beam intensity as a function of beam intensity. If neither amplification nor recombination

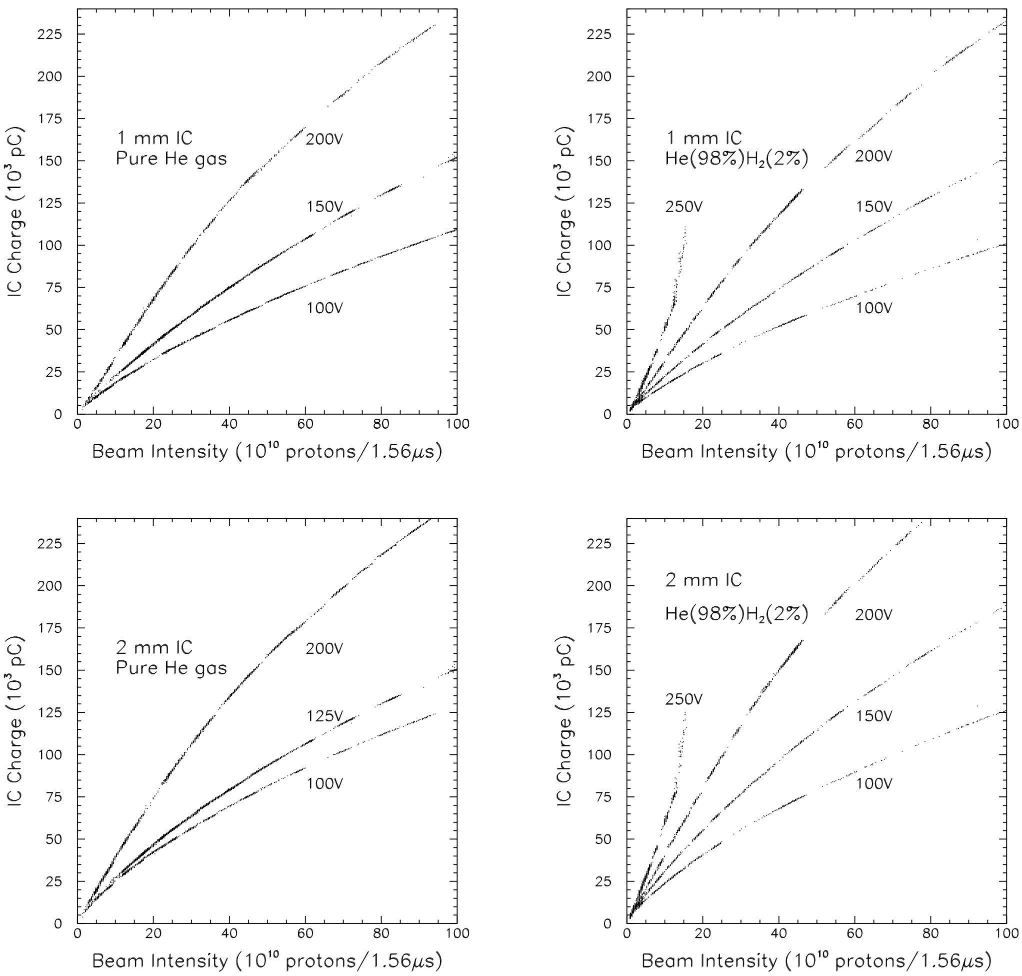

Figure 8.8: Scans of the 1 mm and 2 mm chambers versus beam intensity in He and HeH_2 gas. The beam intensity is that measured by the toroid, corrected for the fraction outside the active area of the chamber.

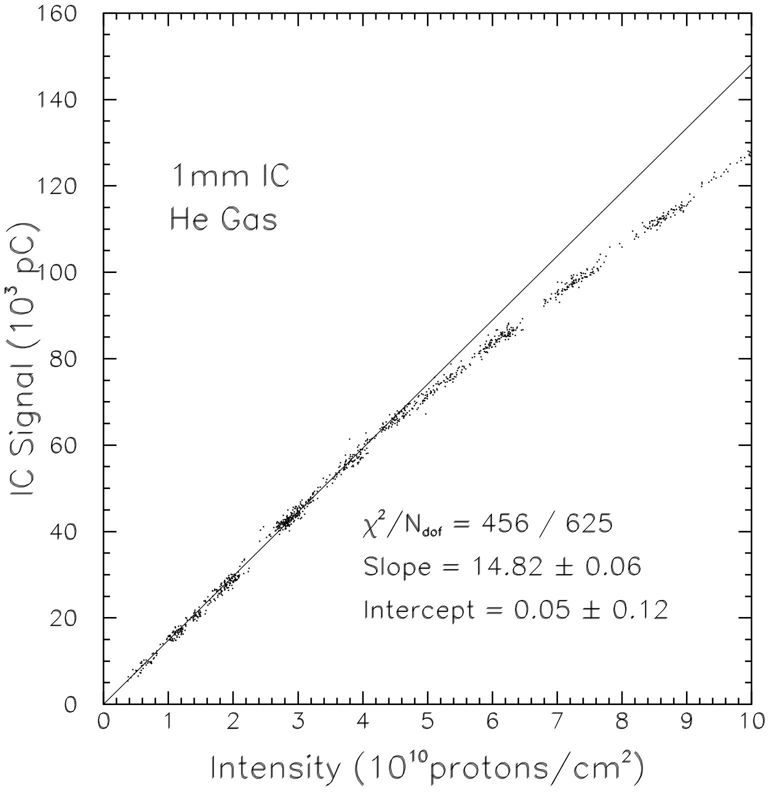

Figure 8.9: Scan of the 1 mm chamber operated at 200 V versus beam intensity in He gas, showing only the low intensity region. Also shown is a linear fit using the data below $4 \times 10^{10}/\mathrm{cm}^2$.

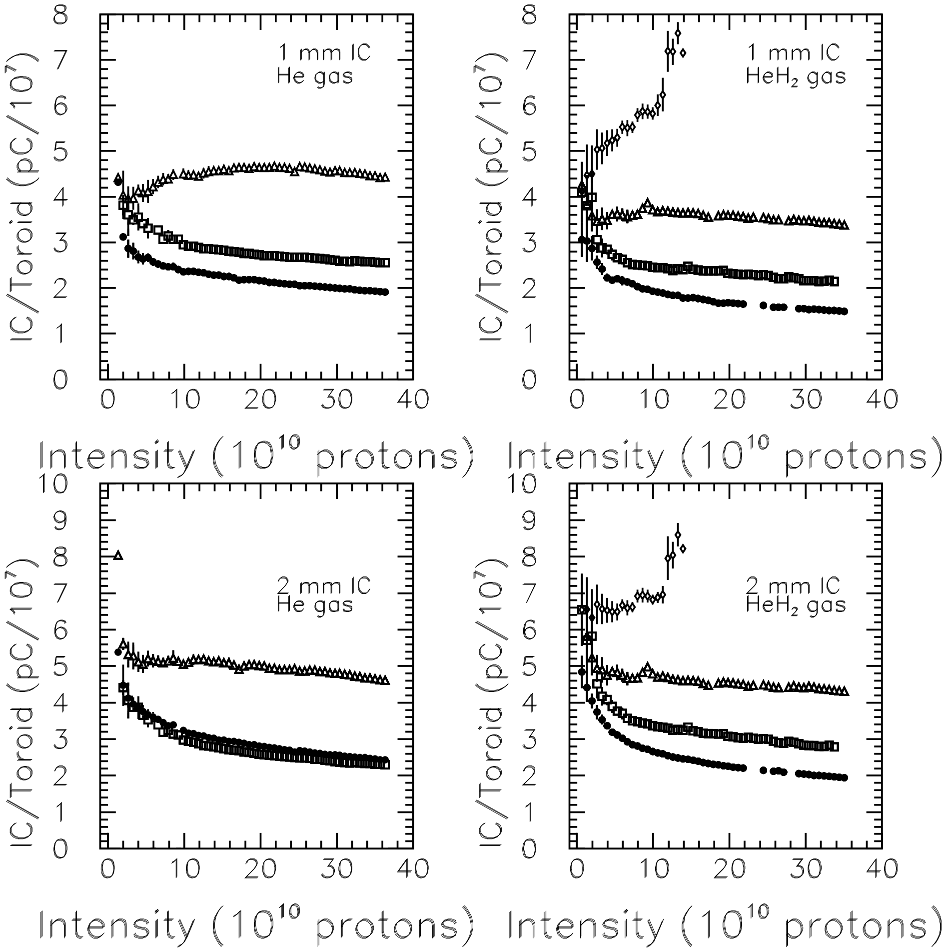

Figure 8.10: Scans of the 1 mm and 2 mm chambers versus beam intensity in He and HeH_2 gas as in Fig. 8.8. Plotted is the charge collected per proton (in pC/10^7 protons) as a function of beam intensity. The points are for applied bias potential of 100 V (filled circles), 150 V (open squares), 200 V (open triangles), and 250 V (open diamonds).

are present, all curves should be horizontal lines at the value of the primary ionization in the gas. Recombination losses result in a dip below this value as the beam intensity increases. Amplification results in a rise in this quantity. As may be seen, the initial recombination losses are moderated at intensities of $\sim 5 \times 10^{10}$ protons/spill, and in the 1 mm chamber at ≥ 200 V or the 2 mm chamber at 250 V amplification can actually overtake recombination. These observations confirm the expectation of such behavior from the previous section.

The interplay of gas amplification and charge recombination is a significant modification to earlier theoretical investigations [144, 160], which considered only recombination effects. Both recombination and amplification occur as a result of the space charge from the ions. The ionization chamber has an effective range of linear operation that is protracted by this interplay. However, while the response of a unity gain ionization chamber follows a relatively simple scaling with temporal variations of gas density (pressure and temperature variations), the space-charge dominated ion chamber will have a more complicated dependence as the beam intensity increases. The results of this chapter show that the ion chamber must be operated in the middle of the voltage plateau in order to avoid space charge effects at higher intensities.

Chapter 9

Ionization Chamber Performance in the NuMI Beam

In this chapter, we present data from the NuMI beam which demonstrates the *in situ* performance of the ionization chambers and their operation as the Beam Monitoring system. §9.1 catalogs fluence distributions measured for various beam conditions. §9.2 examines the variance of signal response with bias voltage. Finally, §9.3 examines the signal linearity with intensity.

9.1 Measured Particle Distributions

The Beam Monitoring system measures the two-dimensional particle fluence at each station. While the beam is operating, profiles are calculated and displayed in the control room for each beam pulse. Here, we recreate a few demonstrative distributions and profiles for different beam situations. The data in this section are from the first few months of running when the Hadron Monitor was biased at 190 V and the Muon Monitors at 300 V.

9.1.1 Proton Beam Without Target

During commissioning, and at a few later times, the proton beam was transported through the target hall with the target removed. In this state the proton beam interacts only with the air and vacuum windows, such that most of it is deposited in the Hadron Absorber. The Hadron Monitor then measures the proton beam directly, and the Muon Monitors only a small number of particles produced by interactions in the absorber. Target-out data are also used in App. B to study the pointing accuracy of the primary beam system.

Fig. 9.1 shows the two-dimensional charge distributions measured at the monitoring stations for this beam configuration[1]. With an expected ionization of 1.6 electrons per proton in the Hadron Monitor, we expect 2.6×10^5 pC to be deposited in the Hadron Monitor for each 10^{12} protons. If all this charge were deposited in the central pixel we would measure an ionization density of 4.4×10^4 pC/cm^2/10^{12}ppp. The data correspond roughly to this number considering some charge is deposited on the adjacent pixels and in the gaps between pixels.

The signal in alcove 1 is substantially smaller than the Hadron Monitor. The signal is made up of neutrons created in the absorber and of muons from decays of the few pions created in the upstream material. Neutrons do not directly ionize the chamber gas, but are expected to produce some signal through nuclear collisions in the chamber gas [107]. The signals in alcove 2 and 3 are smaller due to the higher momentum threshold on muons necessary to reach there.

[1]Here the charge is presented as an areal density. The area of a pixel is 58 cm^2 and the beam intensity was 0.37×10^{12} protons.

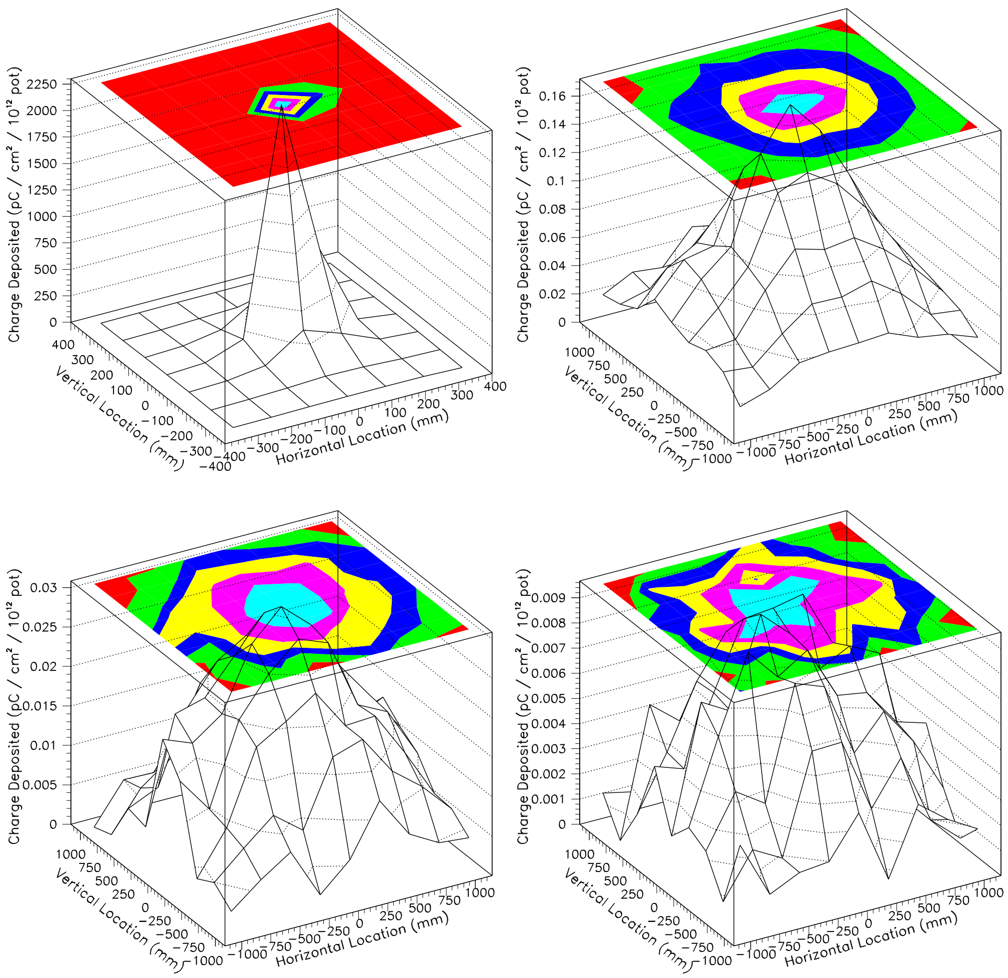

Figure 9.1: Measured charge distributions at the monitoring stations with the proton beam, but no target, such that the beam hits the absorber. The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right. Alcoves 2 and 3 show noise as their signal is barely larger than background.

9.1.2 Bare Target Beam

A bare target beam is one where the proton beam is properly steered onto the target, but the horn focusing is inactive. In this case, most of the protons interact in the target to produce pions, but the low-energy pions are not focused and stray into the shielding. The muon (and neutrino) beam has higher average energy than the focused low-energy beam, but much lower rate.

The charge distributions at the stations are shown in Fig. 9.2. The distribution in the Hadron Monitor is significantly wider and lower in amplitude; caused by the substantial scattering and absorption in the target material. From Coulomb scattering alone, one expects that the proton beam acquires a divergence of:

$$\hat{\theta} = \frac{13.6\,\mathrm{mrad}\cdot\mathrm{GeV}}{120\,\mathrm{GeV}} \times \sqrt{4} = 0.24\,\mathrm{mrad} \tag{9.1}$$

resulting in 18 cm RMS beam size at the Hadron Monitor. Nuclear scattering effects provide an additional divergence. The central peak in Fig. 9.2 is mostly the remnant proton beam. The flatter portion of the distribution is secondary particles emitted at wider angles as well as neutrons from the absorber.

The distribution in alcove 1 has a feature that reduces the peak muon flux in a vertical band. This feature was not predicted by precursory Monte Carlo simulations. Our conjecture is that this feature is a result of incomplete simulation of the Hadron Absorber[2] or a "shadow" cast by the target. The shadow would possibly result from pions with smaller angles being forced to re-interact at a greater rate in the target.

[2]Initial simulations of the muon flux were performed before the Hadron Absorber was constructed. Plans and drawings were used for the absorber geometry. During assembly, however, it is known that small (few inch) vertically oriented gaps exist on either side of the aluminum core. A more detailed simulation geometry using as-built information is in preparation.

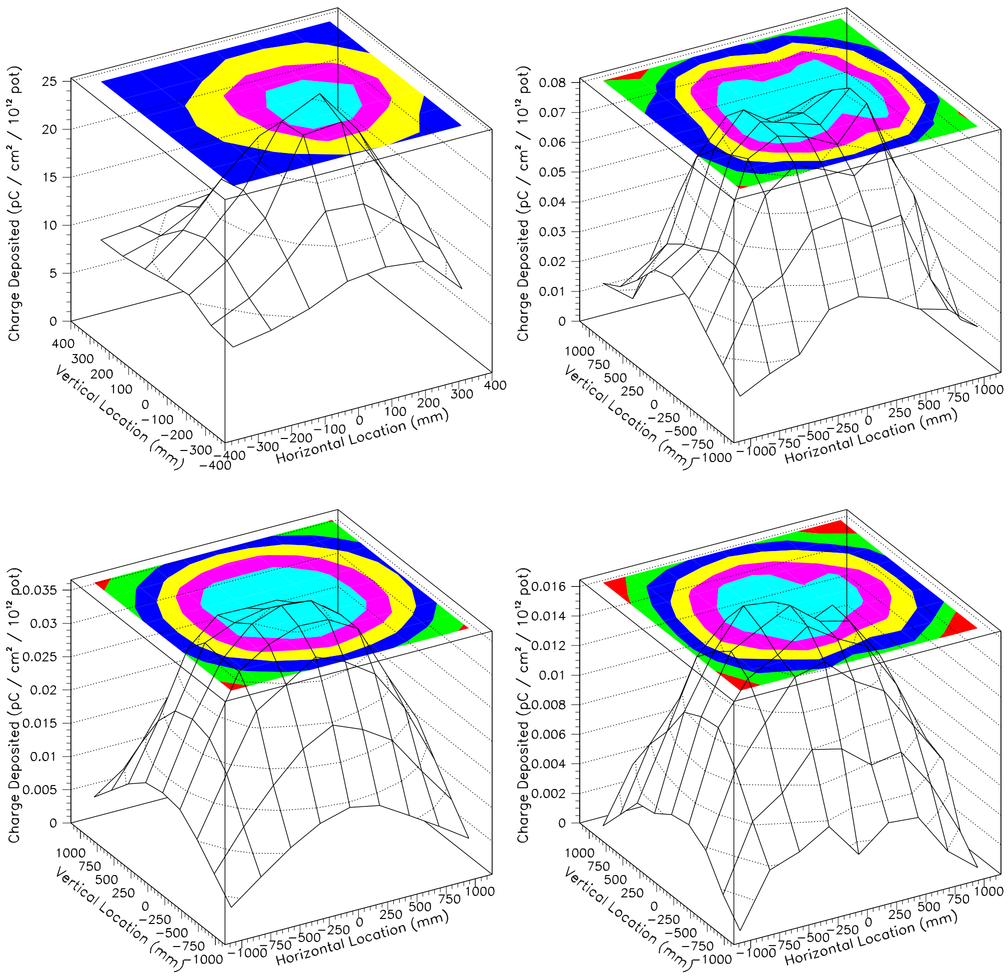

Figure 9.2: Measured charge distributions at the monitoring stations with the bare target beam, such that the proton beam hits the target, but there is no horn focusing. The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right.

9.1.3 Target Positions with Horn Focusing

Fig. 9.3 shows the hadron and muon distributions with the beam centered on target with the horn focusing. The target in the partial Medium Energy (pME) position (see §1.4.2). The Hadron Monitor signal is substantially the same as the bare target beam. The muon alcoves show similar distributions; only the rates are higher and the distributions are somewhat wider.

Fig. 9.4 shows the distributions when the proton beam is targeted between the baffle and target, to beam-right (the target and baffle geometry are discussed in App. B). Some portion of the beam is incident on each of these devices as the beam size is comparable to the distance between them ($\sigma_{\mathrm{beam}} \approx 1.0$ mm; $gap \approx 2.3$ mm). These plots can also be compared with Fig. 9.1: the proton beam without a target. The Hadron Monitor distribution is narrow, but is still scattered and absorbed some by the material. There is also enough material to create a significant number of pions that are focused and decay to muons that reach the monitors.

Fig. 9.5 shows the distributions when the proton beam is targeted on the right ($x > 0$) side of the baffle. In this case the beam is 2.2 mm from the edge of the baffle, so the Hadron Monitor signal has a very wide distribution from the bulk of the beam that impacts the baffle, but still has a small peaked portion from the 1% or so that evades the baffle. Muon alcove 1 shows a very pronounced shadow of the baffle. In this case the shadow has at least two causes: high-energy pions more easily escape from the left side of the baffle; also, the creation point of the pions is now off center in the horn focusing, so pions will be overfocused to the left. This overfocusing is also evident in alcove 2, and to a lesser extent alcove 3.

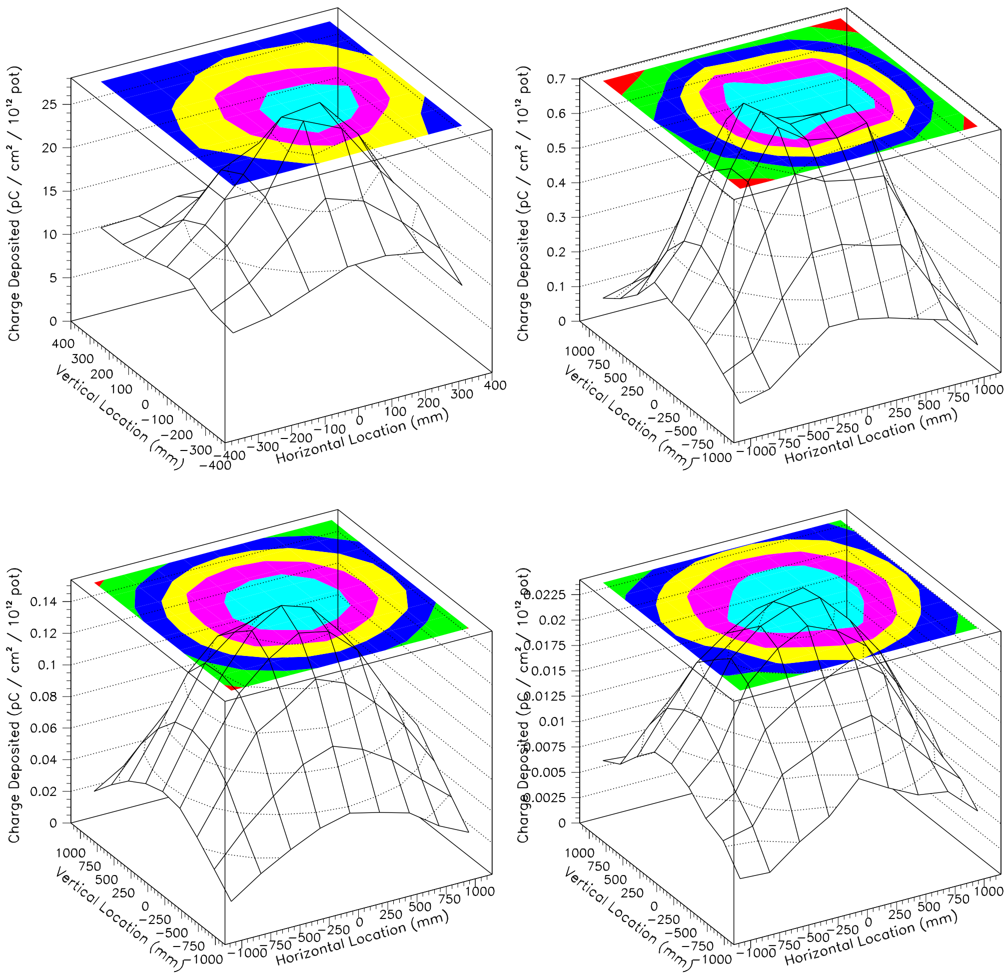

Figure 9.3: Measured charge distributions at the monitoring stations with the beam centered on the target in the pME position and with horn focusing. The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right.

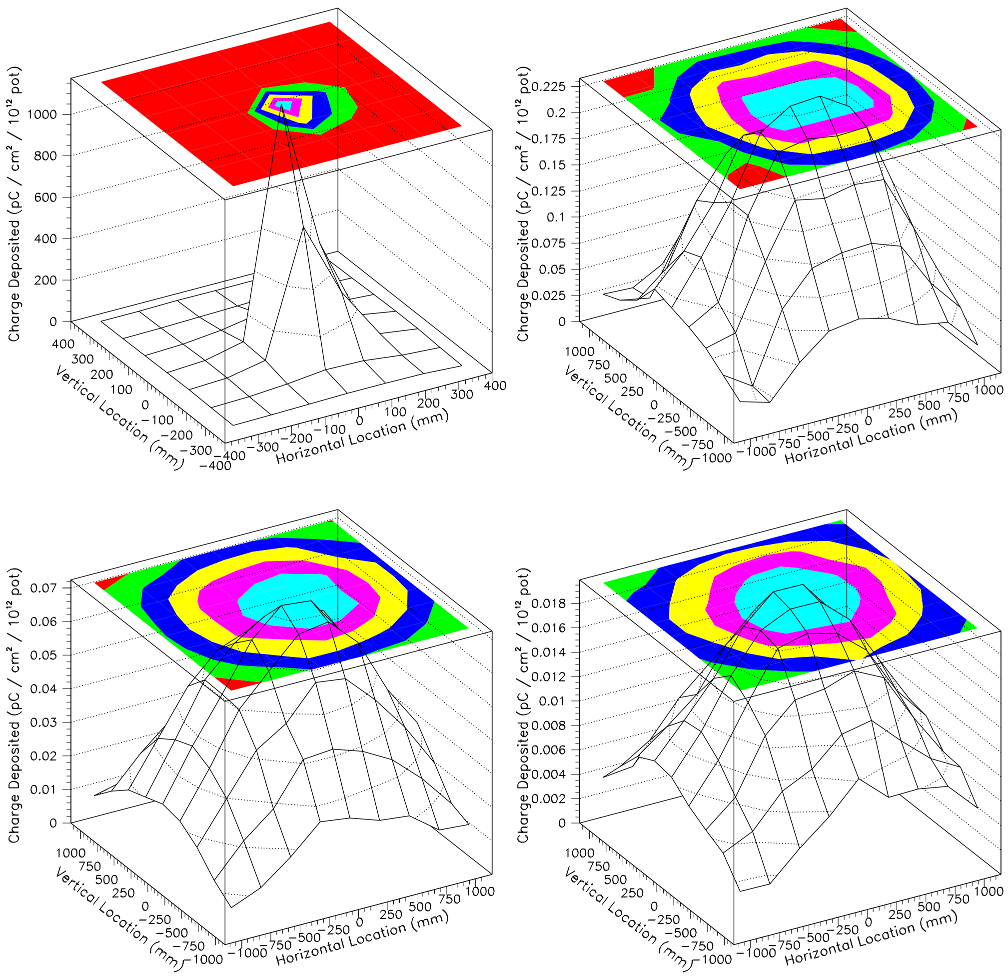

Figure 9.4: Measured charge distributions at the monitoring stations with the proton beam in between the baffle and target in the pME position; this is with horn focusing. The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right.

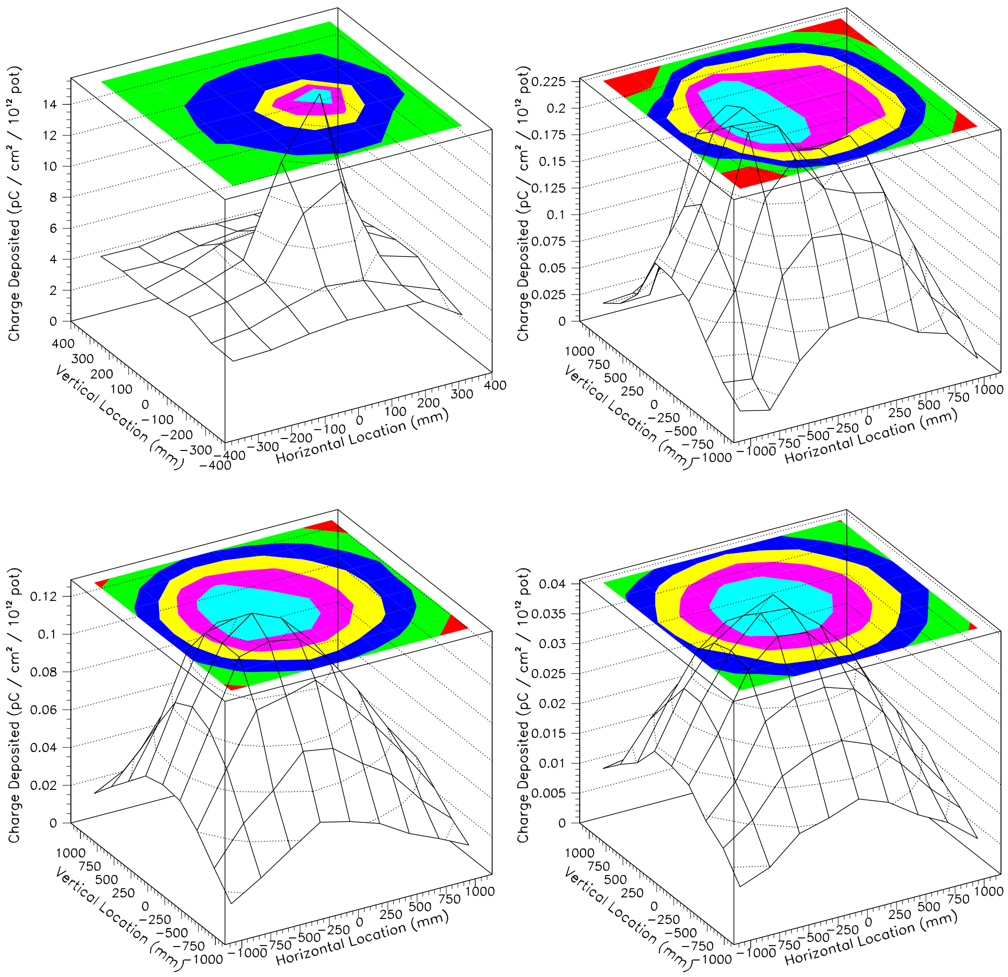

Figure 9.5: Measured charge distributions at the monitoring stations with the proton beam hit the baffle in the pME position, 2.2 mm from edge (; this is with horn focusing. The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right.

9.1.4 High-Energy Target Position

Fig. 9.7 shows the monitor distributions when the target is in the high energy position, such that the pHE beam is generated.

9.1.5 Standard Beam Configuration

Fig. 9.7 shows the monitor distributions for what has become that standard beam configuration for long-term running, namely the LE10 beam.

9.1.6 Different Energy Beams

The NuMI beamline has been operated in four different energy beams: the design Low Energy (LE), partial Medium Energy (pME), and partial High Energy (pHE) beams as well as the, now standard, LE10 (production of these beams is discussed in §1.4.2). Each of these beams produces different profiles in the Muon Monitors, while the Hadron Monitor distribution is largely the same. Fig. 9.8 shows the horizontal and vertical muon profiles for the four beams. The shapes of the distributions are largely the same, with only the amplitudes changing significantly.

9.2 Signal Response to Bias Voltage

The Beam Monitoring ionization chambers were designed such that their performance would be linear at the intensities in the NuMI beam. Here, we evaluate the chambers' performance *in situ* and set the operating bias for the chambers. We consider the central pixel of the Hadron Monitor and the central pixel of Muon Monitor alcove 1. At the date of writing, the maximum ionization measured in a plateau curve at the Hadron Monitor is 3.3×10^{10}

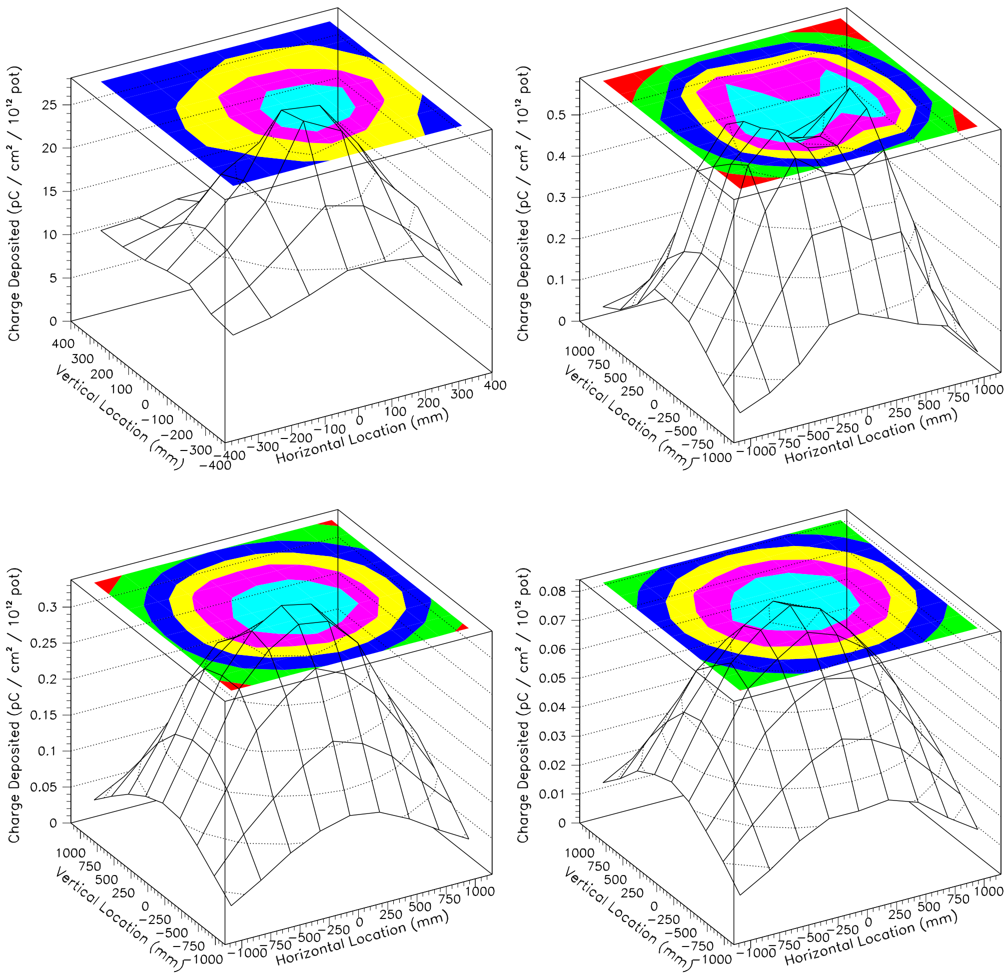

Figure 9.6: Measured charge distributions at the monitoring stations with the beam in the partial High Energy (pHE) configuration. The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right.

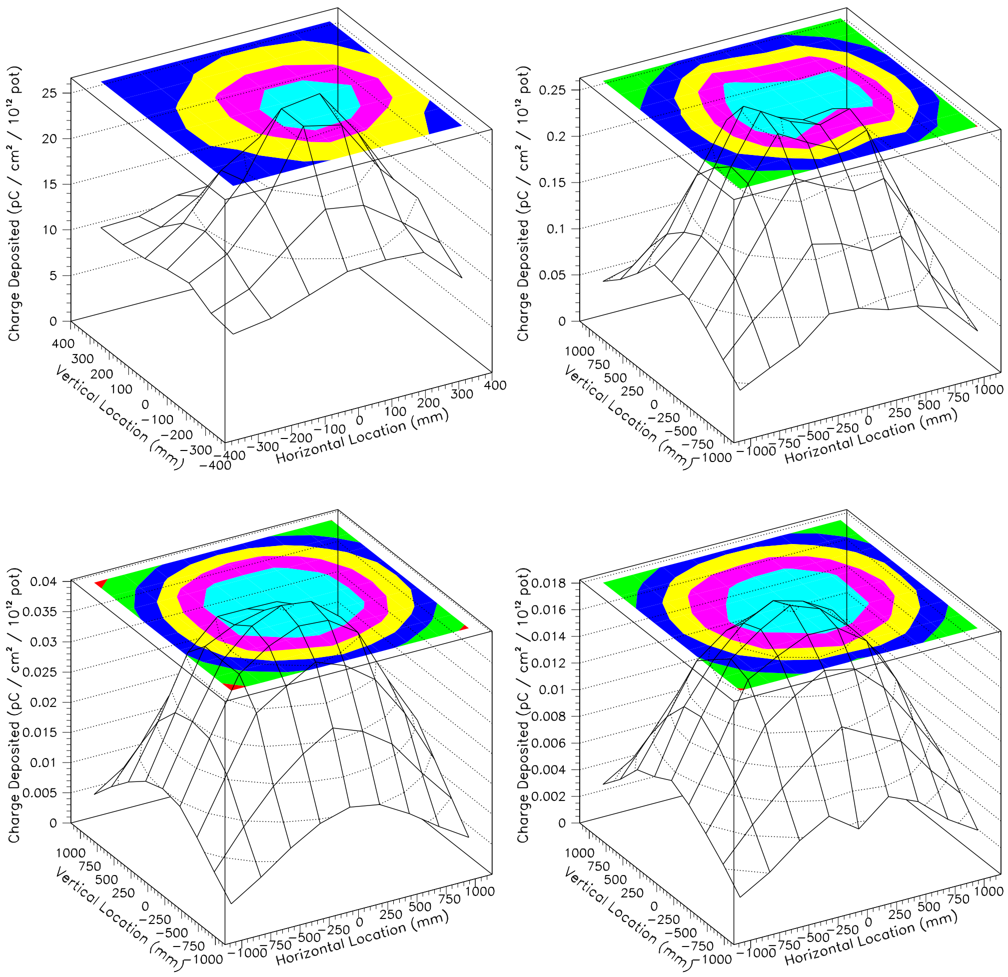

Figure 9.7: Measured charge distributions at the monitoring stations with the beam in the standard running configuration (LE10). The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right.

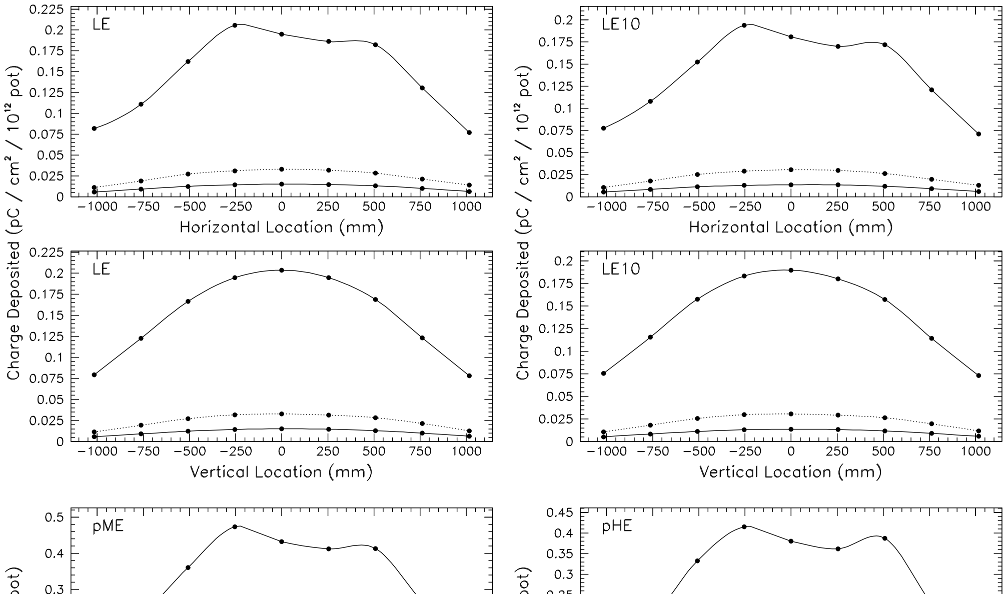

Figure 9.8: Profiles of the measured charge distributions in the muon monitors for different energy beams. The measured 2-D distributions are projected and averaged both horizontally and vertically. Solid lines are for alcove 1, dashed for alcove 2, and dotted for alcove 3.

ionizations / cm^3 in a 8.0 μs spill, for a rate of 4.2×10^9 ionizations / cm^3 / μs; the maximum ionization rate observed in the Muon Monitors is 1.8×10^7 ionizations / cm^3 / μs (taken in the medium-energy beam).

Fig. 9.9 shows the plateau curves taken with the Hadron Monitor central pixel. The bias necessary to reach full collection efficiency increases with intensity (i.e. the plateau is depleted). The loss of plateau is greater than would expected from the simulations shown in Ch. 8. The early depletion likely result from an incorrect value for the recombination coefficient in the simulation[3], or the higher than anticipated oxygen level in the Hadron Monitor ($\approx$ 80 p.p.m).

There is only limited evidence of space charge-enhanced multiplication. However, at the highest ionization rate shown we do not expect formation of a dead zone, only marginal deformation of the electric field (see §C.3). The curve for the highest ionization rate does lie above the others in the 130-180 V region. At higher biases it actually lies belows the others, suggesting that the substantial recombination is limiting the gains from space-charge.

Extrapolating from the plateau depletion at this intensity we expect good linearity to probably another factor of two increase in intensity. Based on this data we have operated the Hadron Monitor at 130 V (see Fig. 9.11).

Fig. 9.10 shows the plateau curves for the Muon Monitor alcove 1 central pixel. It reaches complete collection efficiency at $\sim$ 15 V. Based on this meager depletion, we do not expect any problems until far beyond the highest muon intensity reached in the NuMI beam. Based on this data, we have operated the muon monitors at 300 V.

One note of concern, however, is the slight slope to the plateau curve. The signal rises $\approx$ 7% between 20 and 300 V of bias. In the past, such slopes

[3]The recombination coefficient, r, has an uncertainty of at least a factor of two

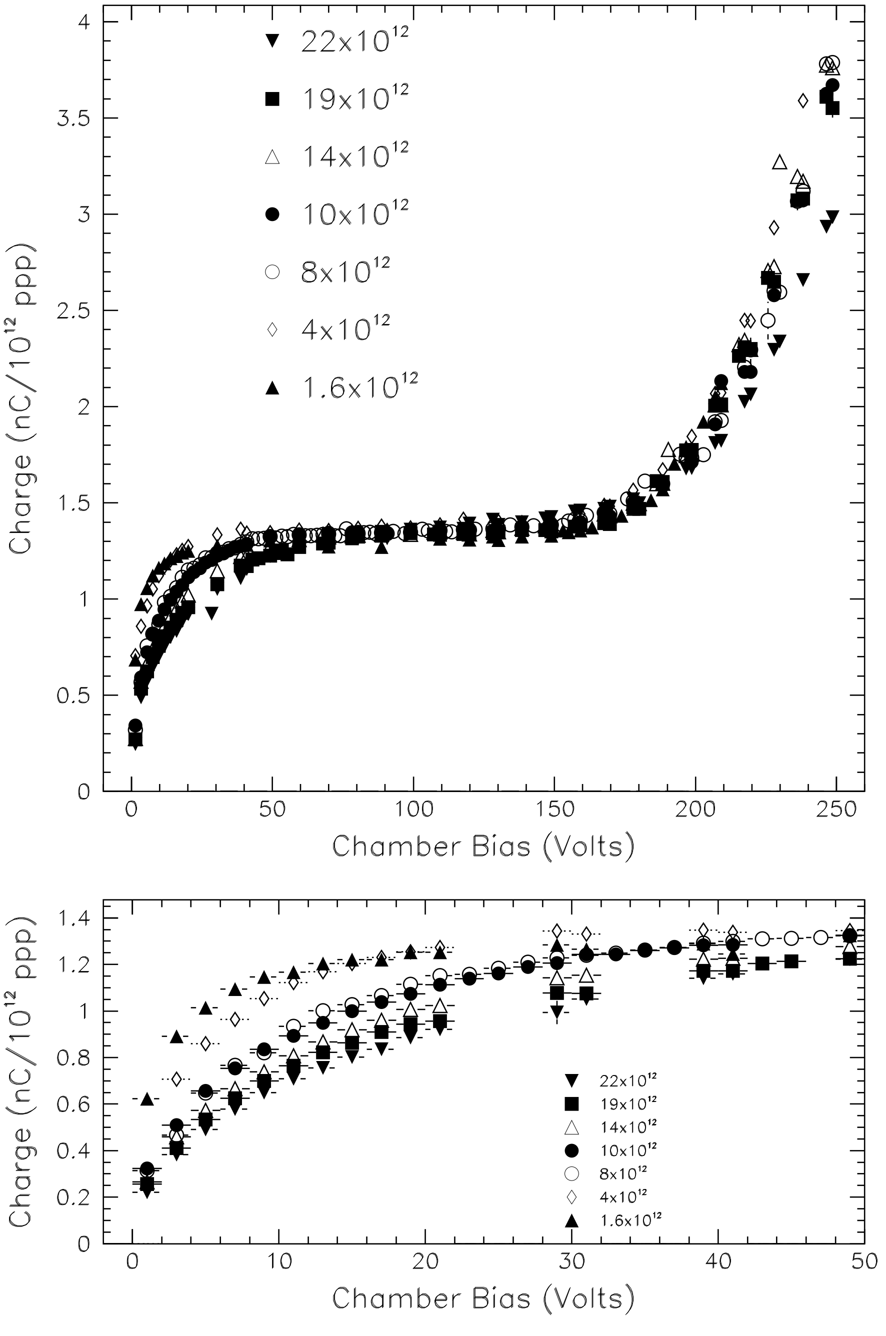

Figure 9.9: Plateau curves of the Hadron Monitor central pixel for several different proton beam intensities. The lower plot is zoomed in on the low voltage region.

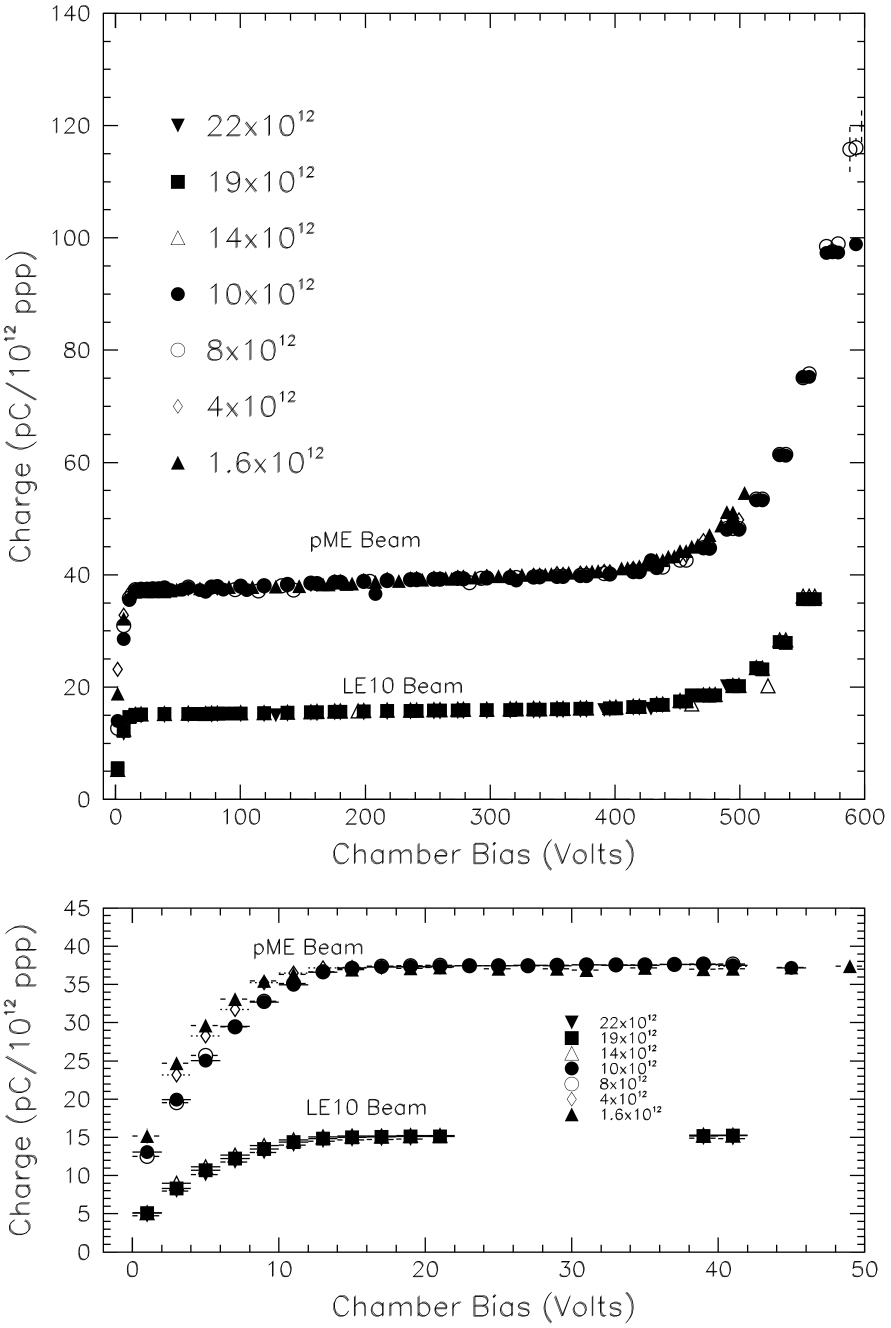

Figure 9.10: Plateau curves of the central pixel from Muon Alcove 1 for several different proton beam intensities. The lower plot is zoomed in on the low voltage region. The upper curve is from intensities met with the pME beam, the highest proton intensities achieved are in the lower curve and were achieved with the LE10 beam.

have been a sign that some stray charge outside of the active area is being collected on the electrode [111]. Efforts were made in construction to limit this behavior, but apparently some persists. We expect that the excess signal will scale mostly linearly with intensity, as the ionization throughout the vessel volume is linear with muon rate. However, if the volume is large and has a weak electric field space-charge effects could cause a nonlinearity. Further study with other chambers in the array and other beam intensities will be needed to evaluate this issue.

9.3 Response Linearity

Performing a direct measurement of chamber linearity *in situ* in the NuMI beam would require reducing the proton beam intensity from its maximal value, reducing the number of neutrinos sent to the experiment. In the future a scan may be permitted where intensity is reduced temporarily. Furthermore, different intensity proton beams have different characteristics (e.g., size and divergence) which affect the secondary and tertiary production nontrivially.

In place of a direct scan, we use the data accumulated over several months that happens to fill in much of the desired range of intensities. However, the beam conditions are also evolving over this time so the flux at the monitoring station fluctuates. We can make some limits on beam parameters using the primary beam instrumentation such as the SEMs, but some variability in the signal response remains.

The linearity curve for the Hadron Monitor central pixel is shown in Fig. 9.12. Notwithstanding the afore-mentioned changes in beam condition, the response can be said to be mostly linear for each of the of the bias voltages.

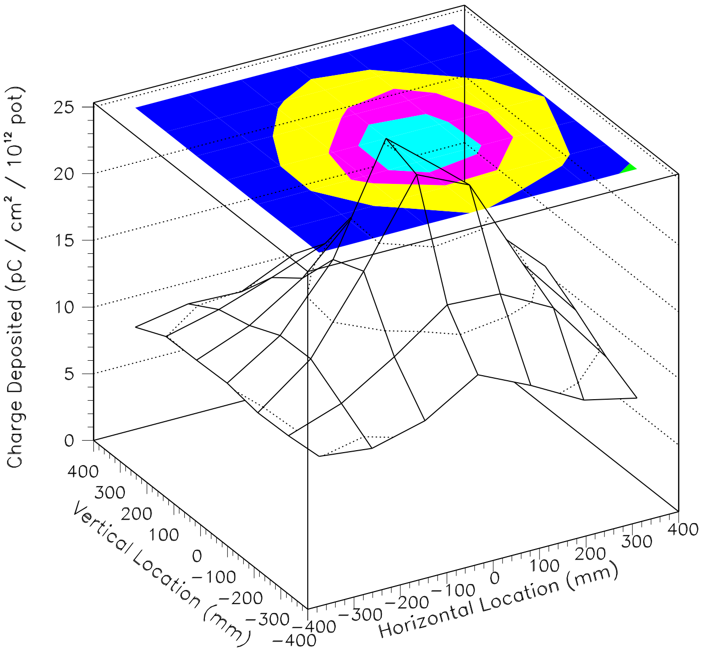

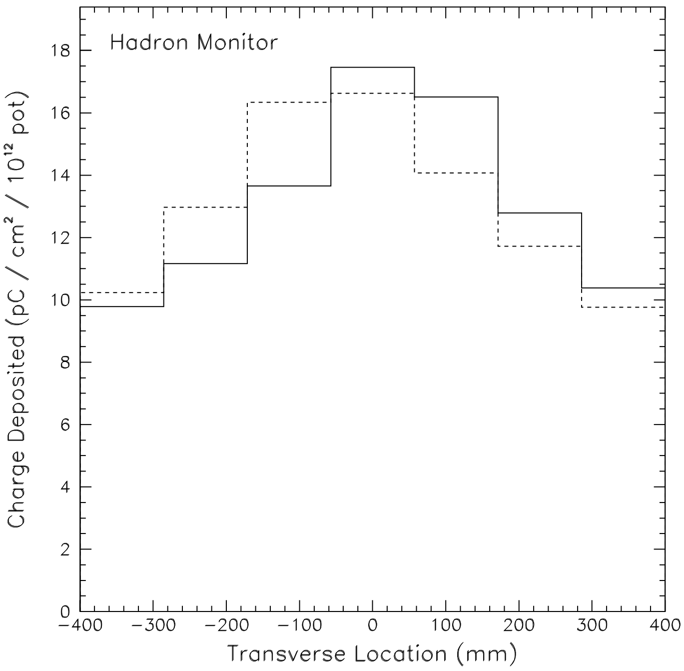

Figure 9.11: Measured charge distribution and profiles in the Hadron Monitor when set at 130 V, the established operating point. The signal is ≈ 10% smaller and the variability in the flat regions at the edges is marginally smaller. The solid line is the horizontal profile, the dashed is the vertical.

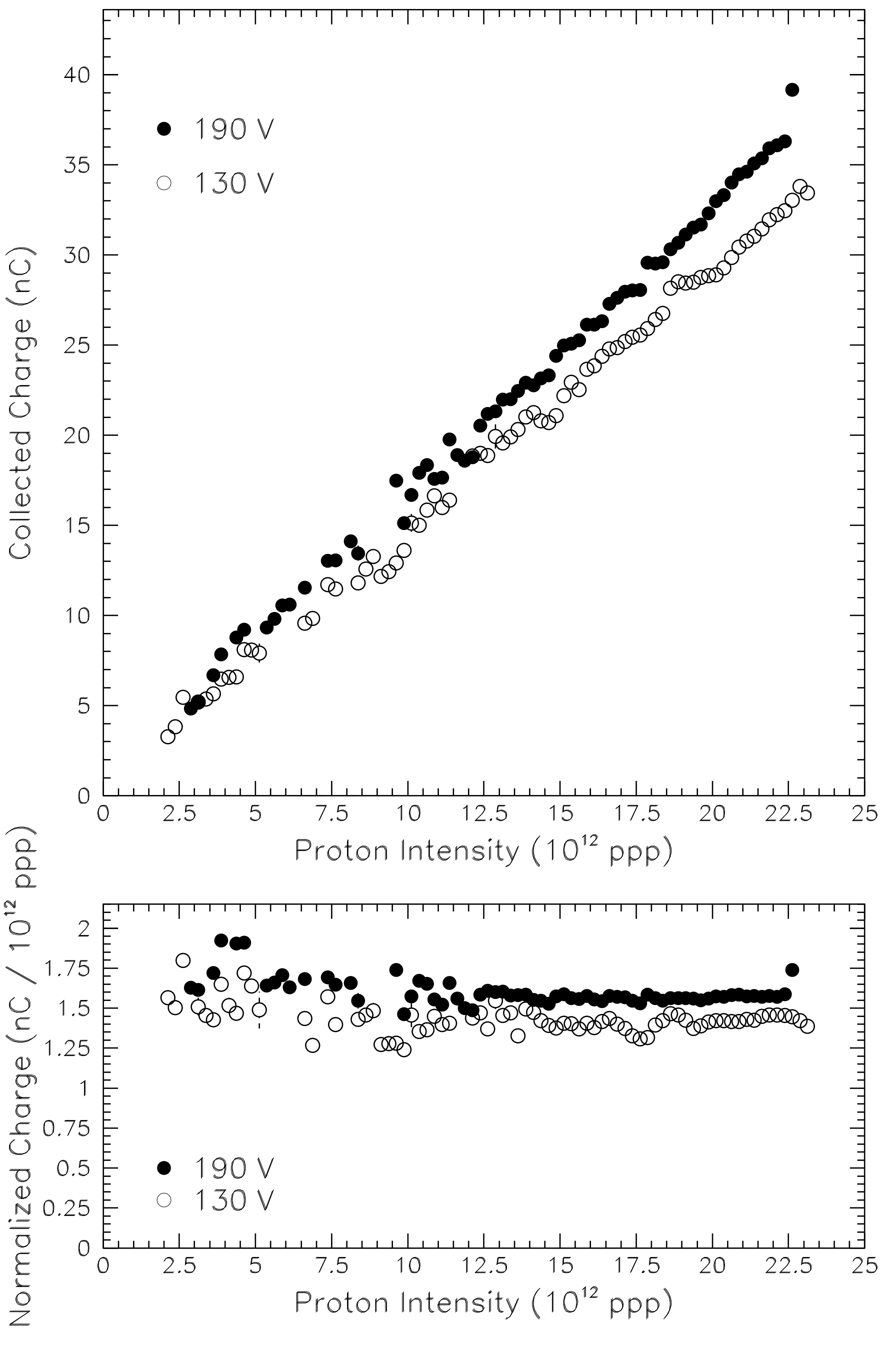

Figure 9.12: Hadron monitor central pixel linearities with beam intensity. The data for the two operating biases is shown. These data have not been taken in dedicated studies and are from months of running over which variation in the beam takes place.

The 130V data has slightly lower response than the 190 V data, as expected:

$$Gain(190\,\mathrm{V}) \approx 1.64\ \mathrm{nC}/10^{12}\ \mathrm{ppp}$$

$$Gain(130\,\mathrm{V}) \approx 1.46\ \mathrm{nC}/10^{12}\ \mathrm{ppp}$$

The ratio of these two slopes is 1.12, consistent with expectations form the plateau curves of Fig. 9.9. The short scale ripples are changes in beam quality.

The normalized data suggest a possible slope in the response, but it is comparable to the amount of signal variation. The situation should become clearer with even higher intensity data, and possibly with a dedicated beam study.

The linearity for the central pixel of Muon Alcove 1 is shown in Fig. 9.13. The muon signal sustains the same temporal variability as the Hadron Monitor signal, further suggesting that the variation is because of beam conditions and not chamber response. As above with the Hadron Monitor, we call the response mostly linear.

We have to question also whether in a dedicated study the beam conditions can be kept the same between low- and high-intensity beams. The results of these linearity data suggest that the plateau curve will be the more useful tool for analyzing chamber performance.

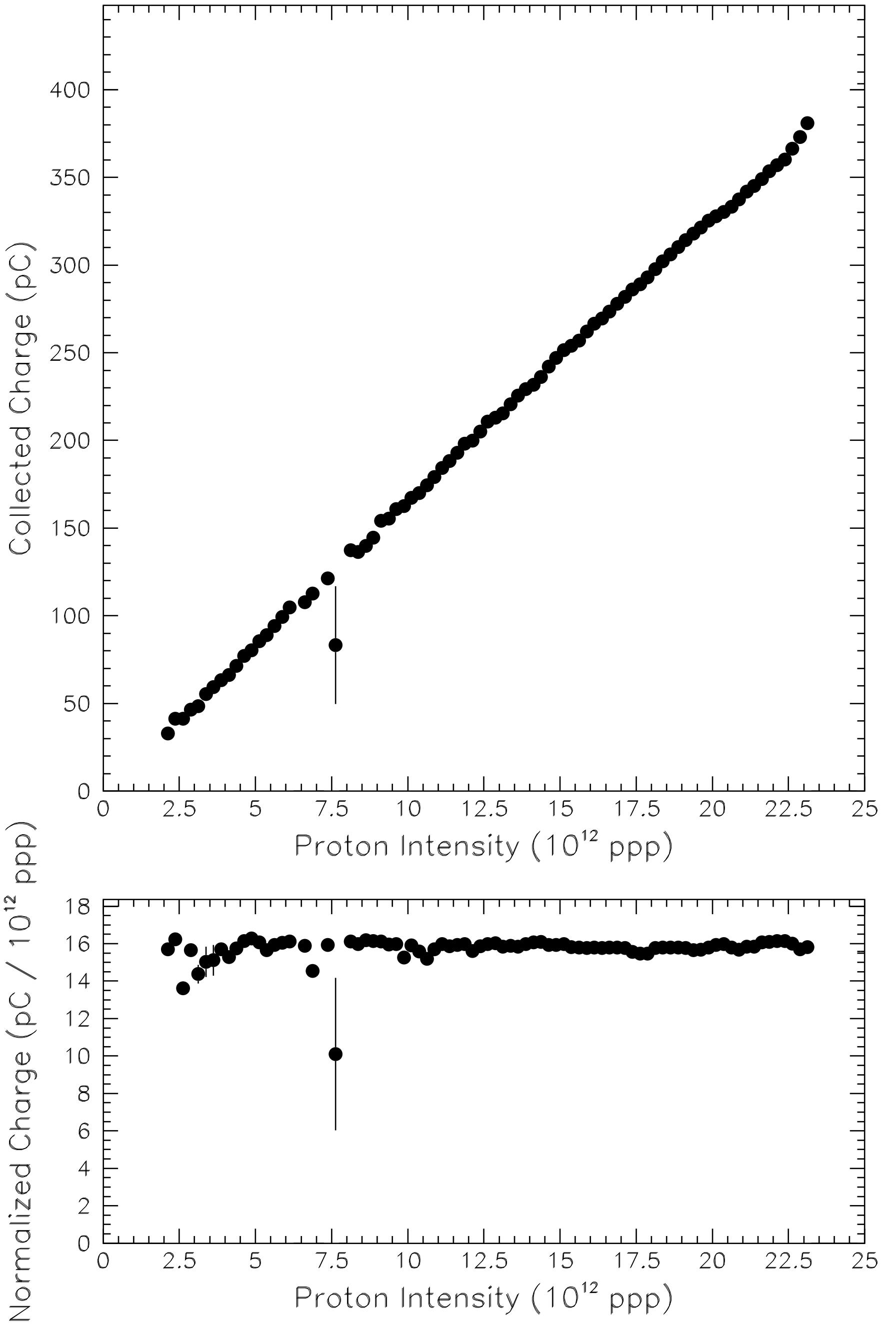

Figure 9.13: Muon Monitor alcove 1 central pixel linearity with intensity. These data have not been taken in dedicated studies and are from months of running over which variation in the beam takes place.

Chapter 10

Beam Monitoring Measurements

As discussed in Ch. 6, the secondary and tertiary beam monitors have an important roles in: (a) confirmation of the neutrino flux predictions from the beam Monte Carlo, and (b) monitoring the quality of the neutrino beam. At the time of this writing, the NuMI has only recently commenced operation, hence long term experience with the monitors is lacking. However, a number of studies were undertaken to demonstrate the monitors' sensitivity to the horn optics, proton beam steering on the target, etc. The studies described in this chapter confirm the measurement capabilities of the monitors. Additionally, this chapter describes how the monitors were used as a diagnostic tool during the failure and replacement of the NuMI target.

10.1 Crosschecks of Neutrino Flux

As discussed in Ch. 6, the location of the target relative to the horns selects the particle momentum focused by the horns. Likewise, the horns' current changes their effective focal length. By varying both of these parameters independently, an effective check on the beam flux can be made. We conducted such a test during the early commissioning of the beamline, although the study awaits more effective simulation tools for interpretation.

Fig. 10.1 shows the results of scanning the horn current from 0 kA to 200 kA. The total charge in each detector is measured as a function of the horn current and target position. The Hadron Monitor shows only a marginal dependence on horn current and beam energy. The charge from the remnant proton beam and electromagnetic showers dominates the signal. The fact that there is no clear correspondence with beam energy suggests that much of the signal at the Hadron Monitor is not from muons.

Muon alcove 1 shows strong dependence on horn current and beam energy. The LE, LE10, and pME beams show consistent gains with higher current, but the pHE beam shows a maximum at ~ 160 kA. Above this current the pHE is overfocusing low-energy pions to the extent that the muon flux at alcove 1 is actually reduced. The difference between the LE and LE10 curves provides some measurement of the effect of the 10 cm target displacement on muon rates, and hence the neutrino flux. Such must be confirmed with a Monte Carlo simulation.

Alcove 2 shows little effect from horn focusing in the LE and LE10 positions, but the pME and pHE beams show enhanced flux. Alcove 3 is only responsive in the pHE beam, in accordance with our expectations (§6.4).

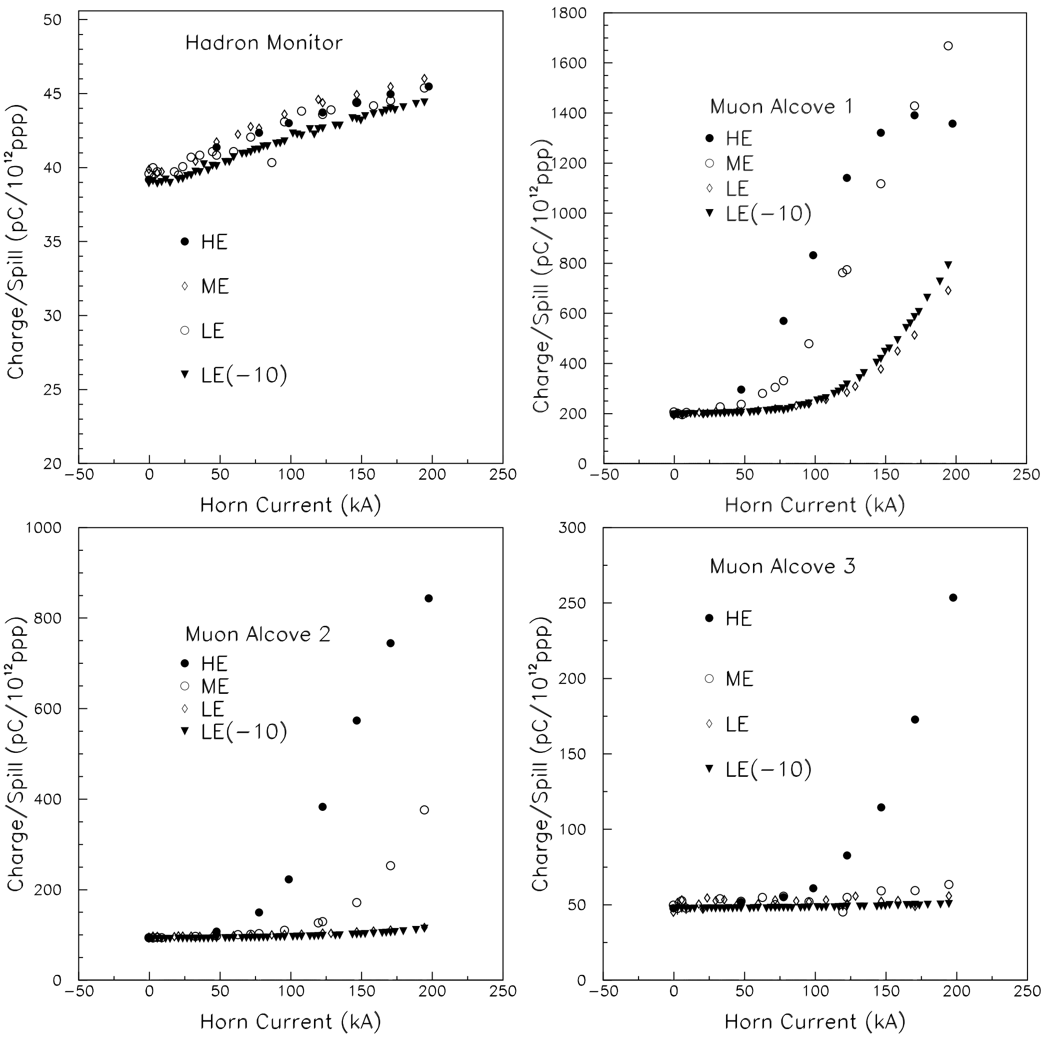

Figure 10.1: Horn current scans in the LE, LE10, pME (labeled as ME), and pHE (labeled as HE) beams. The beams differ in target position, being 0, 10, 100, and 250 cm behind the nominal LE position. In each of these target positions the horn current was varied. The Hadron Monitor shows a marginal increase in flux from horn focusing. The alcoves all show increased flux with increased horn current and beam energy - except for alcove 1 where the muon flux actually decreases in the HE beam with the highest currents.

10.2 Crosschecks of Proton Beam Stability

The Muon Monitors are quite sensitive to changes in the proton beam position at the target. Thus, this system can be used as a redundant check of the proton beam steering, complementing the extrapolation provided by the primary beam instrumentation. We performed several scans early in the NuMI commissioning process which can be used to calibrate any subsequent excursions detected by the Muon Monitors. The results of this study have been confirmed with two accident spills during subsequent running, one in March and one in November.

10.2.1 Study of Muon Yield

Results from the muon alcoves are shown in Fig. 10.2 for two scans of the proton beam across the target while located in the low-energy position (with and without horn focusing). The alcove 1 signal comes from both the decay of pions and from particles emanating from the Hadron Absorber. Its peaks in intensity, without horn focusing, correspond to when the proton beam penetrates in the region between the baffle and target and strikes the absorber. The opposite is true with horn focusing, in which case muons from focused pions dominate the signal.

Alcove 2 is sensitive to higher energy muons, so the center of the target stands out promptly without the horns. The baffle's greater lengths provides for more reinteractions of high-energy pions, so it produces fewer muons here. With horn focusing far more muons are produced off the baffle than the target; this is a result of the baffle being $\sim$ 2m upstream of the target. Alcove 3, being sensitive to only the highest-energy muons, sees very little signal. The flux of higher energy muons (alcoves 2 & 3) drops noticeably when the proton beam

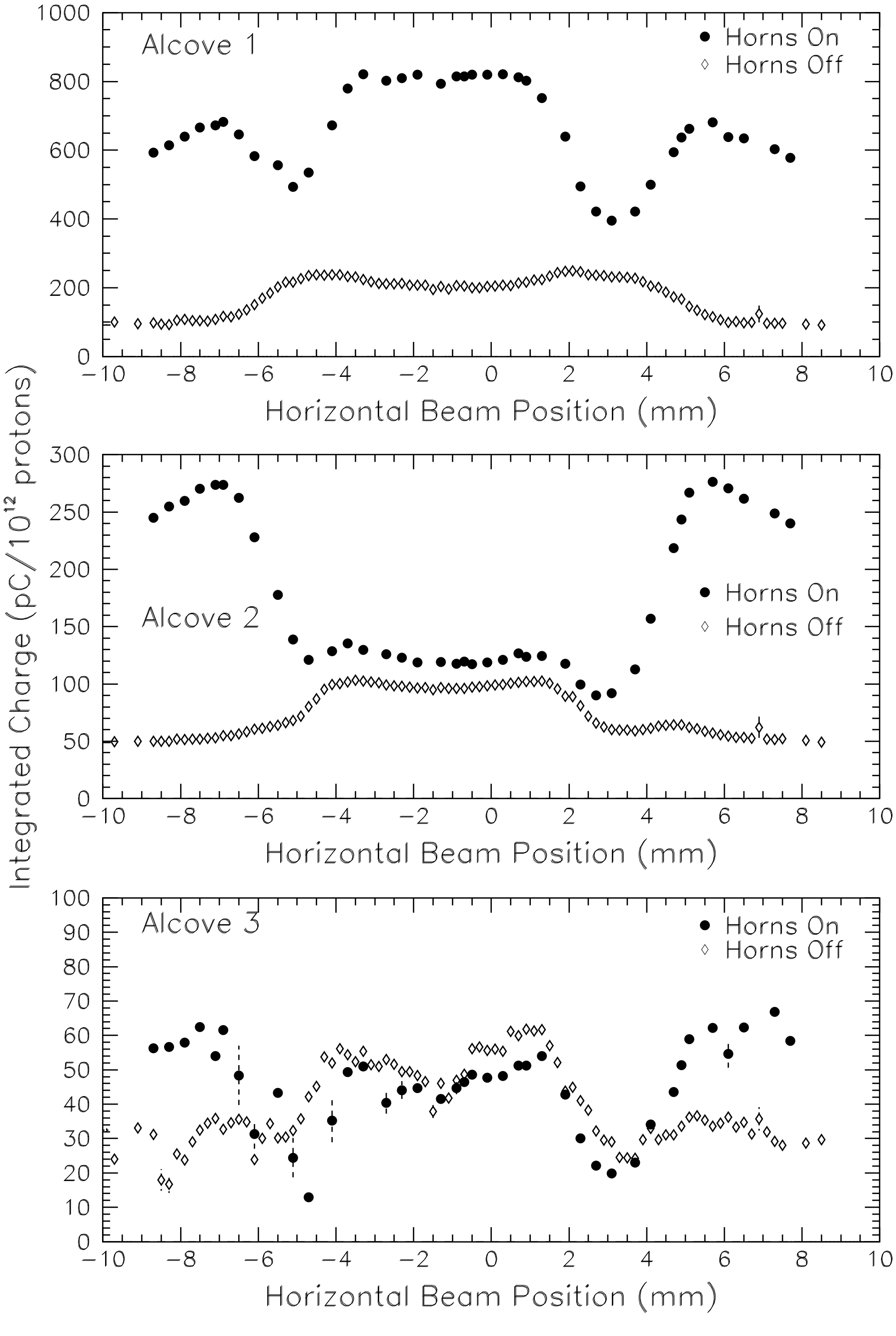

Figure 10.2: Muon Monitor results from horizontal target scans in the low-energy position. The target is centered at -.95 mm. These scans are different from that in Fig. 11.8 which was used for alignment; beam intensity is higher here. The target and baffle features produce different peaks and valleys in the different alcoves, depending on horn focusing. Note that the alcove 2 flux with horns is much higher when the beam is targeted on the baffle instead of the target.

strikes the center of the target, instead of the edges; this is a result of pions reinteracting in the target: when near the edge, high-energy pions are more likely to escape out the side of the target before interacting again. Such can be a useful monitoring crosscheck of the quality and position of the proton beam.

The Muon Monitors become more sensitive to target-baffle features with the horn focusing and higher-energy beams.. Examples of this are shown in Figs. 10.3 (pME beam) and 10.4 (pHE beam). The horn-off data is similar to that of the low-energy scan. With the horns pulsing, the alcove 1 signal clearly shows the target in the center and has small bumps corresponding to the baffle edges. Alcove 2 shows the same situation, but with a greater signal from the baffle pions. Alcove 3 is a more extreme example of this with the baffle muons producing almost as much signal as the target muons.

For NuMI, the Muon Monitors play a secondary role in alignment (see Ch. 11). However, other experiments have used their Muon Monitors more extensively to align the target hall components (see [102, 95]). In that case the proton beam and components were moved until the muon flux is maximized; such a method cannot be used for the NuMI alignment. As can be seen in Figs. 10.2, 10.3, and 10.4, centering the proton beam on the target often does not maximize muon flux. There is often a local minimum in the flux at the center where pions reinteractions in the target deplete the number of medium- and high-energy pions. In fact, the long-term monitoring of the proton beam stability on the target can be accomplished by detecting a rise in muon yields due to missteering of the beam.

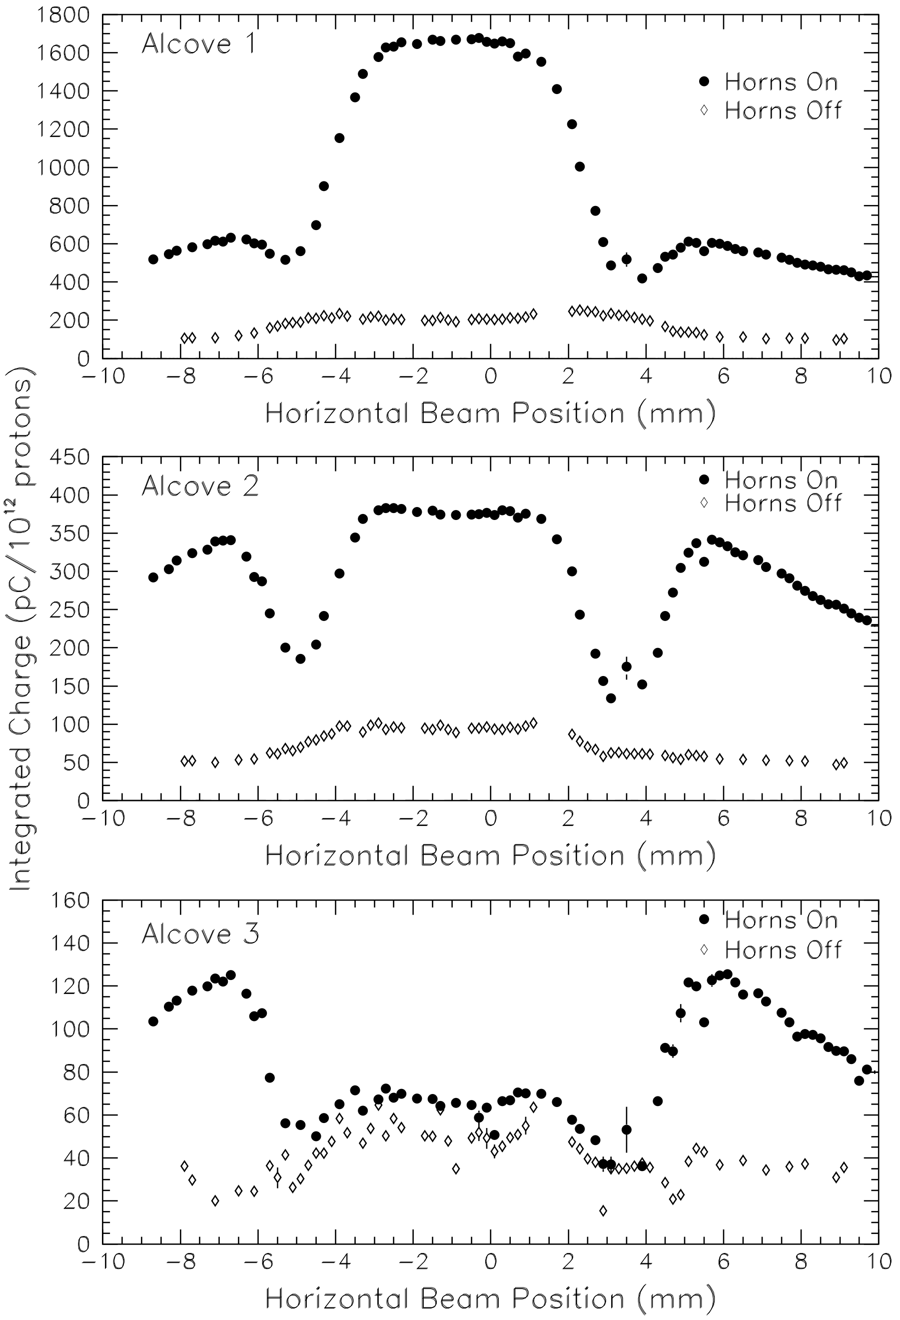

Figure 10.3: Muon Monitor results from horizontal target scans with the target in the medium-energy position; with and without horns pulsing. The target is centered at -.95 mm.

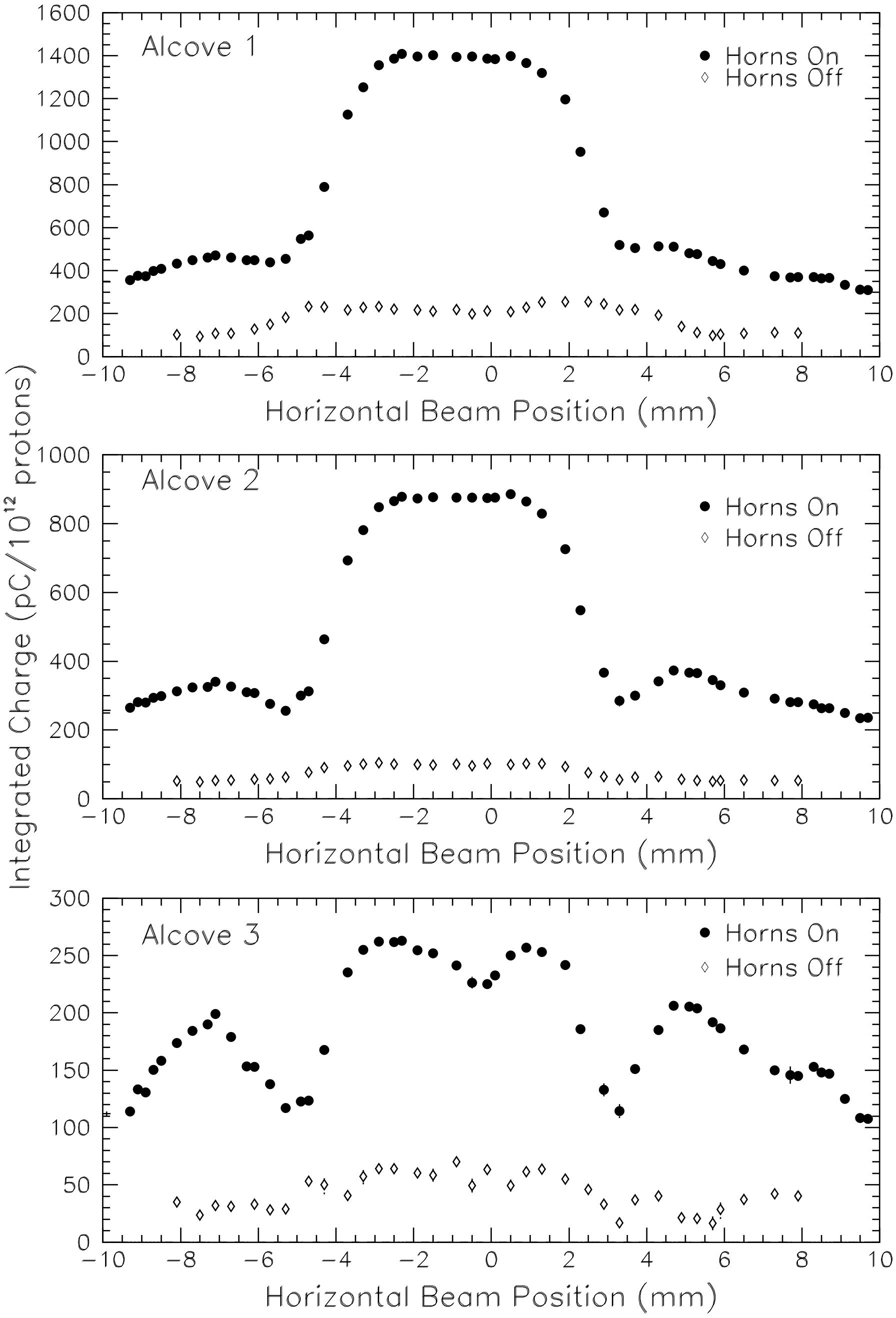

Figure 10.4: Muon Monitor results from horizontal target scans with the target in the high-energy position; with and without horns pulsing. The details of the target and baffle are more apparent than in Fig 10.2, particularly in alcove 3, with the horns and with the target in a higher energy position. Interestingly, the maximum muon fluxes in alcove 3 occur $\lesssim 1$ mm into either the target or baffle edge. The target is centered at -.95 mm.

10.2.2 Muon Monitor Centroids

The centroid of the distributions measured at the Muon Monitors is capable of giving significant amounts of information, though interpretation of details will benefit from further Monte Carlo simulation of the beamline.

Fig. 10.5 shows the horizontal centroid position of the muon distributions measured during horizontal scans of the proton beam across the target. The horn focusing acts like a lens for various energies of muons and can overfocus or underfocus them[1]. The focusing results in something like an image of the proton beam to be projected via muons at the alcoves. The centroid positions measured in the alcoves thus are correlated (positively or negatively) with the beam position on target, the proportion being dependent on the focusing. As each alcove samples a different portion of the muon energy spectrum, the alcoves' centroid correlation will vary differently.

The correlations in Fig. 10.5 suggest that for the LE beam the pions are underfocused for all energies measurable in the Muon Monitor. The ME beam is underfocused in alcoves 2 and 3, but alcove 1 appears slightly overfocused. In the HE beam, the muons reaching alcoves 1 and 2 were clearly overfocused, while alcove 3 is underfocused. From these distributions, it is clear that the Muon Monitors can detect future excursions of the proton beam off target center.

[1]In some cases the muons are not focused at all, because they pass through the neck of the horn. However, they still acquire an angle if steered away from the center of the target. The angle results from more secondaries escaping the nearer side of the target.

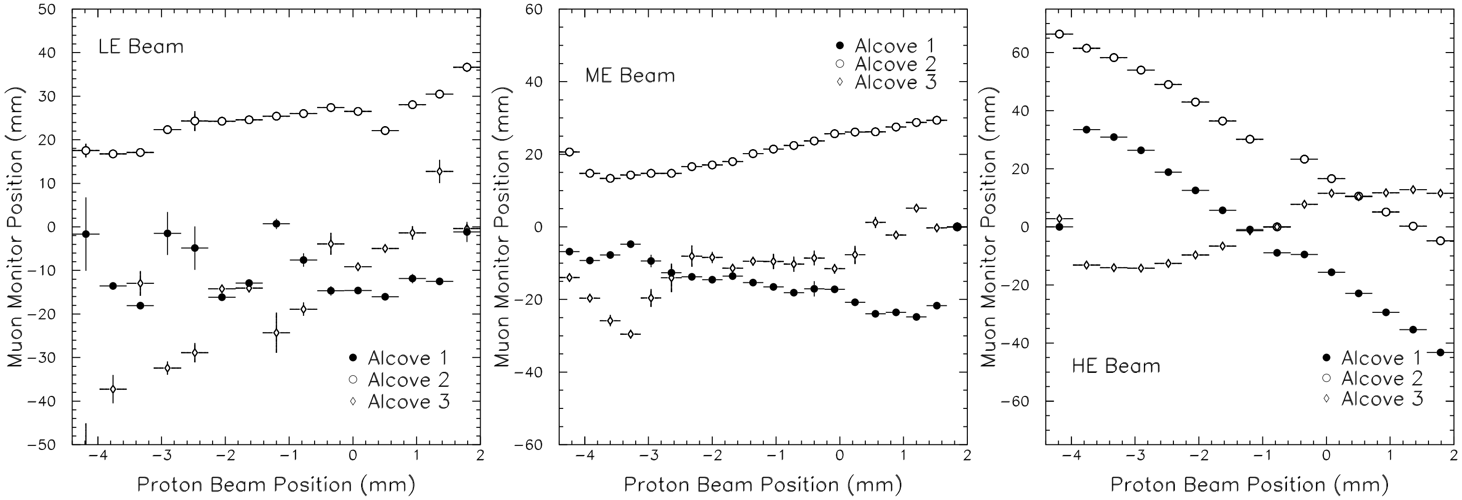

Figure 10.5: Muon Monitor centroids during target scans in the various beams. When the centroid position is correlated with the proton beam location there are two possible explanations: the proton beam is nearer edge of the target and more particles escape out that side, or underfocusing results in an image of the proton beam that scales with its position. Anti-correlation suggests overfocusing where the image is inverted. Everything is underfocused in the LE beam; the pME beam is slightly overfocused in alcove 1 muons, and underfocused in alcoves 2 and 3; the pHE is clearly overfocused in alcove 1 and 2 muons, but still underfocused for alcove 3.

10.3 Target Integrity

On March 23, 2005 water was found in the oil of the vacuum pump keeping the target canister evacuated. Hadron Monitor data from a subsequent scan of the proton beam across the target showed that there was an appreciable amount of water within the target vessel. Following attempts to remove the water from the canister while still in the beamline, beam scans (across the target) were performed periodically and the data from the Hadron Monitor used to evaluate the success of such attempts.

Ultimately, the target was removed from the beamline. The target was drained of water and back-pressured with helium with the intent being that the helium pressure would minimize any leak into the the target vessel. Upon replacement, the target was scanned several times using the Hadron Monitor to verify that back-pressure system operated as planned.

This section describes some of the measurements made with the Beam Monitoring System to diagnose the extent of the target leak.

10.3.1 Measurements with Compromised Target

The target cooling water leak was discovered shortly after it began – though the particular event and time remain unknown. The first horizontal beam scan across the target (shown in Fig. 10.6) on March 23 showed reduction in the proton beam transmission through the target-baffle gap. The reduction of beam reaching the Hadron Monitor was $\sim 60\%$, indicating ~ 1.9 interaction lengths of water filling the target's snout. However, the signal measured while the beam was centered on the target was not significantly reduced. We thus concluded that all of the water intersected by the beam in these scans must be in the downward sloping snout of the target (see Fig. 1.4).

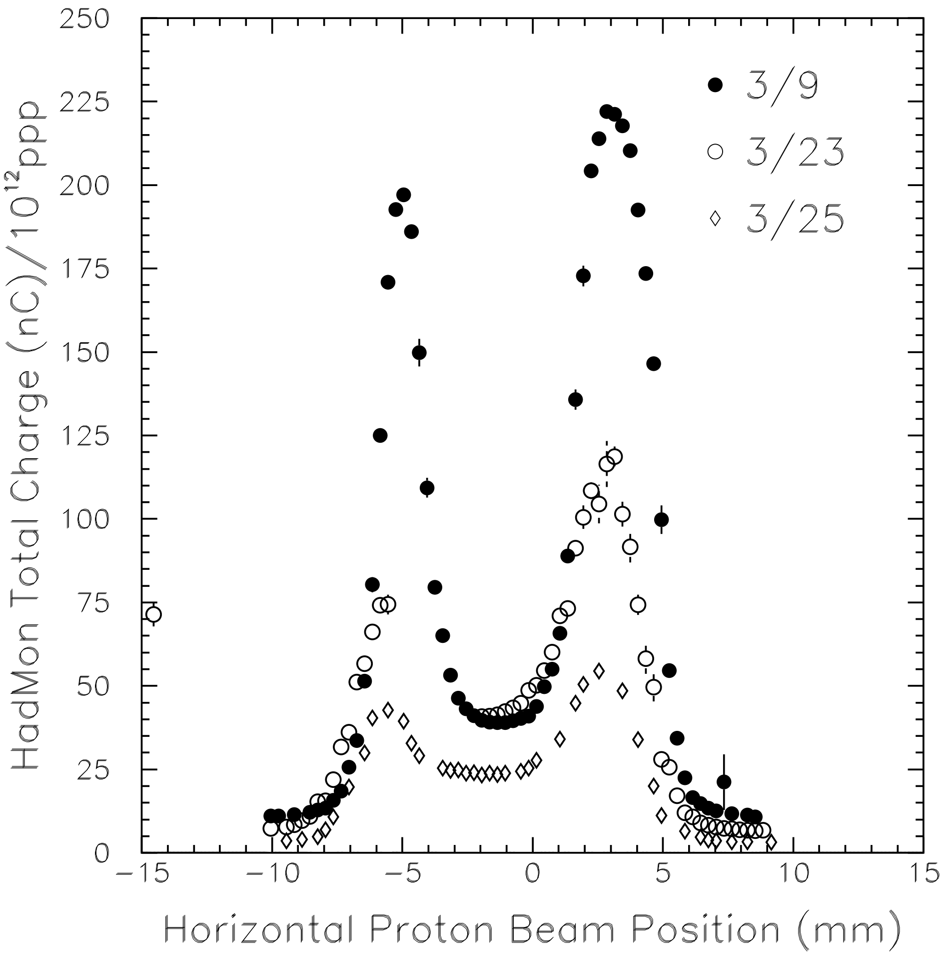

Figure 10.6: Hadron Monitor normalized intensity of horizontal target scans at several dates during the target incident. The March 9 data show a typical horizontal scan, presumably before any damage. The March 23 and 25 data show reduced signal in the Hadron Monitor from water in the target canister. The March 23 data shows reduction in the gap between target and baffle, but not on the target itself – suggesting that there is some amount of water in the target snout, but not in the barrel upstream of the target. The March 25 data shows that the target canister must have been almost entirely full of water up to the beam centerline.

A later scan on March 25 (Fig. 10.6 showed greater reduction in transmission, as well as a 40% reduction on the center of the target. This new information implied that the water had filled the target canister completely up the entrance point of the beam[2].

The beam width at the Hadron Monitor also showed the expected variation with water in the target, as shown in Fig. 10.7. The water's effect as a scatterer is more apparent. In fact, it can be seen that the beam width measured on the target center was somewhat larger in the March 23 data. This scattering is caused by ~ 8 cm of space downstream of the last target fin where there is nominally no material. Such is expected, by adding 0.25 radiation lengths of water to the 4 radiation lengths of carbon in the the target, which would be expected to scale the scattered proton beam size by $\sqrt{4.25/4} = 1.03$.

Vertical scans of the proton beam across the target gave another measure of the water level. Several are shown in Figs. 10.8 & 10.9. The depth of water the beam has to traverse increases ~ 17 mm for each mm of vertical displacement; however, the beam has its own width so the effect is not as simple. The kink seen around the center of the target was suggestive at one point, but is more likely some effect of surface tension around the horizontal fin or other internal features of the target canister.

10.3.2 Target Recovery

Ultimately, the target assembly was removed from the beamline and examined directly. The water inside the target canister was level with the start of the snout, as expected; however, no obvious leak was found. Testing suggested

[2]There are about 30 cm of distance inside the target canister upstream of the first target segment where there is nominally no material

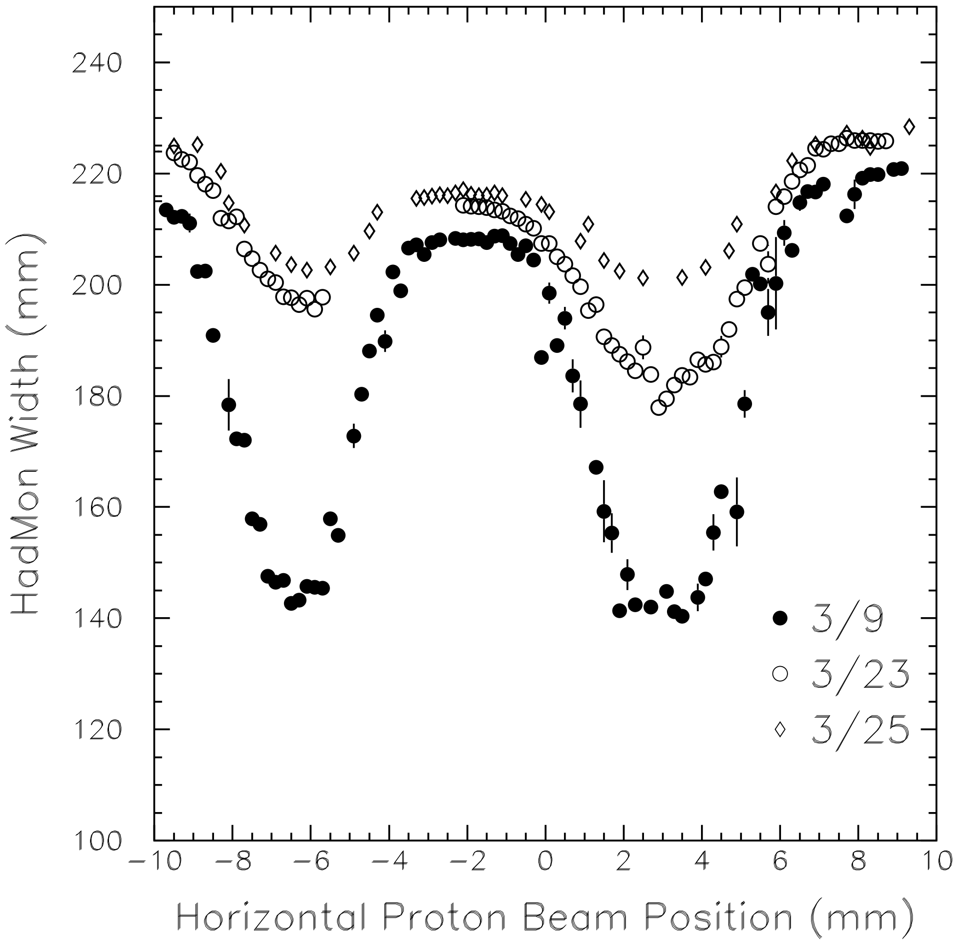

Figure 10.7: Hadron Monitor measured beam width from the same horizontal target scans as Fig. 10.6. The beam width measured at the Hadron Monitor increases with the amount of water within the target canister.

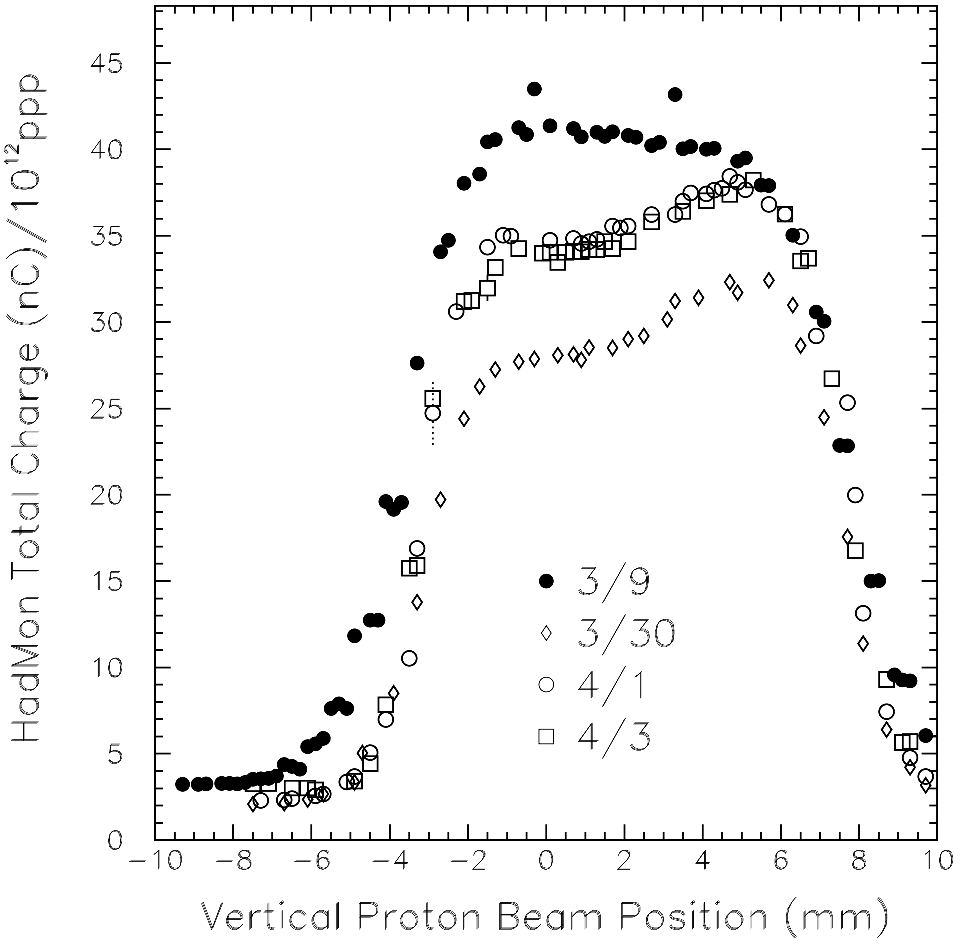

Figure 10.8: Hadron Monitor normalized intensity of vertical target scans during the target incident period. The pre-incident scans shows a negative slope across the target face due to the horizontal fin's position. The target canister is angled at the same slope of the beamline: 58 mrad. As such the depth of water encountered by the proton beam is higher at lower positions at the target – leading to the positive slope in measured signal across the target face. The March 30 data are from when the target canister was completely full with water. The April scans are after an attempt drain water from the target without removing it; after these attempts the water filled the target snout up to some level, either where the snout opens into the larger barrel or the leak location on the cooling pipe.

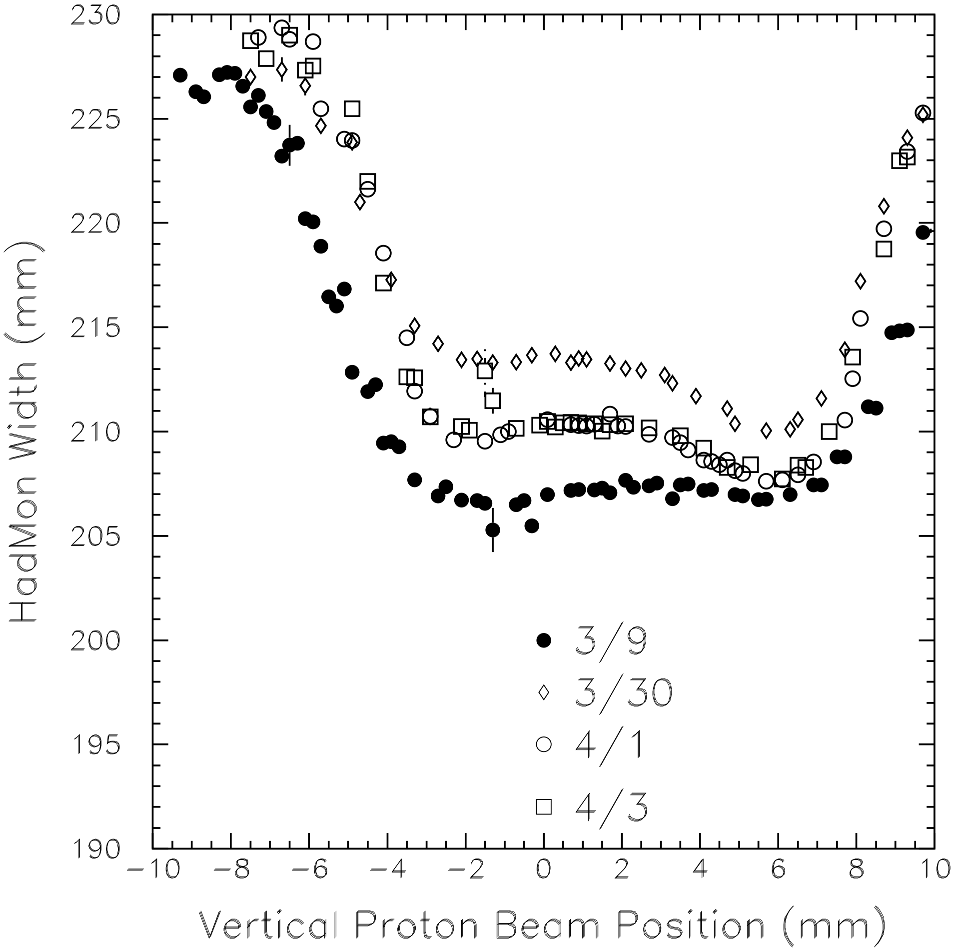

Figure 10.9: Hadron Monitor beam width measurements from the same vertical scans as Fig. 10.8.

that a helium backpressure of 18 psi against the cooling water pressure should be sufficient to keep small leak at bay.

Horizontal scans were again used to verify target integrity after the replacement. The Hadron Monitor measurements are shown in Figs. 10.10 & 10.11. The transmission through the center of the target went back to the levels measured before the target incident. The transmission through the target-baffle gap increased because the primary beam optics had been changed to deliver a more round beam spot to the target, with lesser divergence.

10.3.3 Long-Term Monitoring

The water leak is prevented by an over-pressure of helium in the target canister. As the leak may enlarge and water re-enter the canister, we have established a regular program of target scans to monitor target integrity.

Horizontal scans of the proton beam across the target are performed on an at least monthly basis to get good measurements of the transmission through the target-baffle gap. Additional scans are performed after any major change of operating mode, or when unusual events are observed with the helium overpressure or cooling water systems. As of writing, no further leaks had been observed in the target.

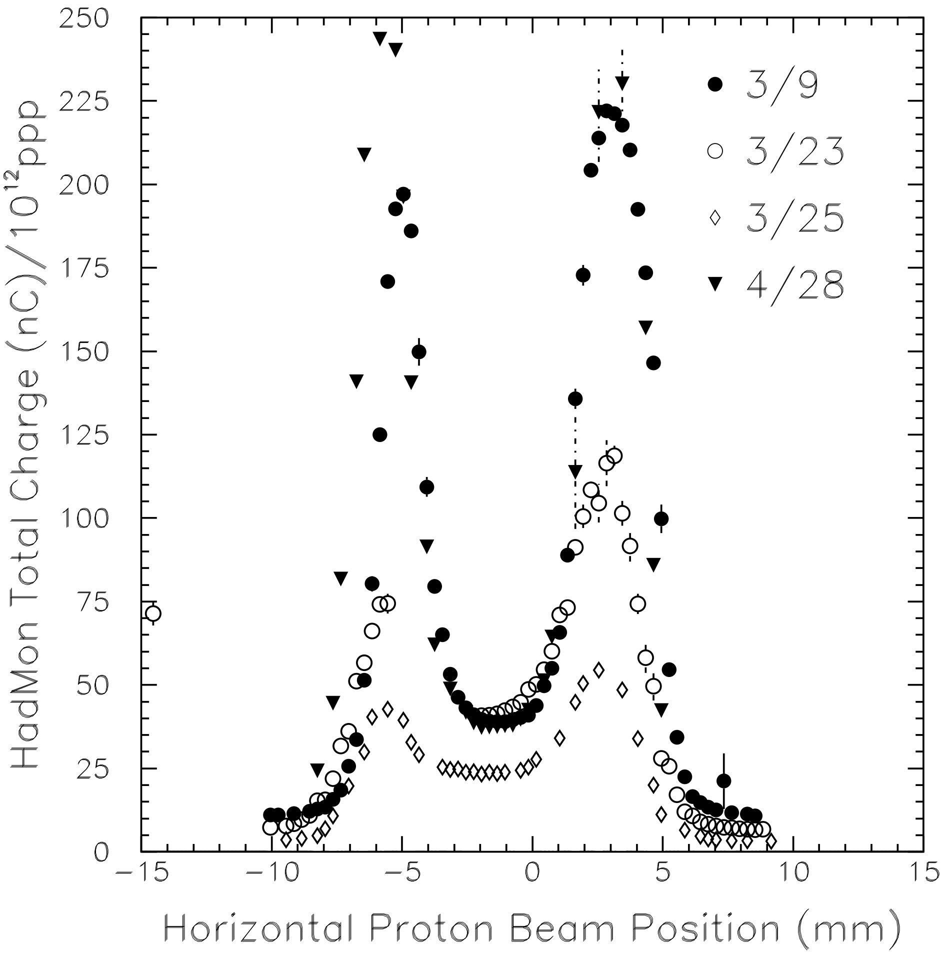

Figure 10.10: Hadron Monitor normalized intensity as measured during horizontal target scans before, during, and after the target incident. The somewhat higher signal after target replacement has been attributed to better target and proton beam angling, as well as the optics change that occurred just before the target incident.

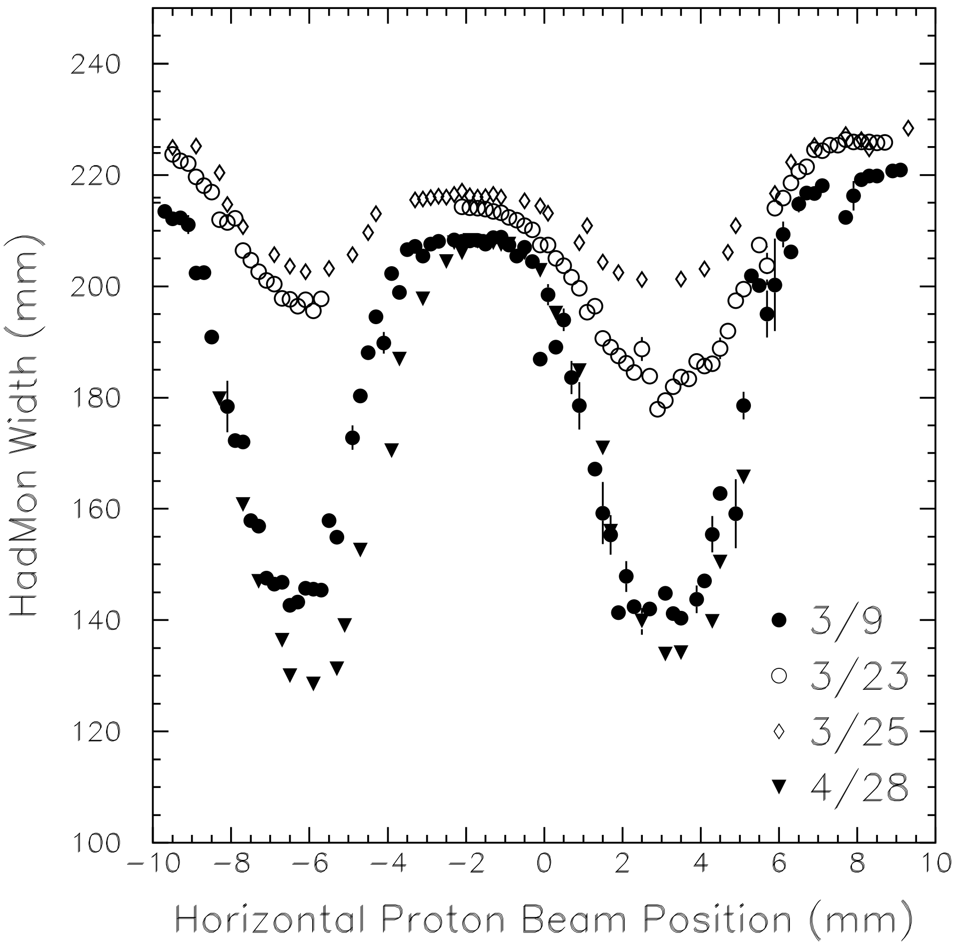

Figure 10.11: Hadron Monitor beam width measurements as measured during horizontal target scans before, during, and after the target incident. Again, better proton beam penetration was achieved after the target replacement by improving the target angle and proton beam angle during scans, as well as changing the proton beam spot through correcting the primary beam optics.

Chapter 11

Beam Based Alignment of Target Hall Components

11.1 Introduction

As discussed in Ch. 6, the beam monitors assist in the commissioning of the beamline through alignment of the target hall components (target, baffle, & horns). The neutrino beam, and particularly the ratio of beam spectra at the the near and far detectors, is sensitive to misalignments, as has been shown in previous experiments [102]. In this chapter we describe the process of beam-based alignment, where the proton beam itself is used to locate the positions and angles of these components with the assistance of the downstream instrumentation.

The low-energy neutrino beam is particularly sensitive to misalignments because there are so many higher energy hadrons that will be better focused for most any misalignment. An off-center beam on the target results in stiffer pions exiting the target; an off-center target results in greater focusing through

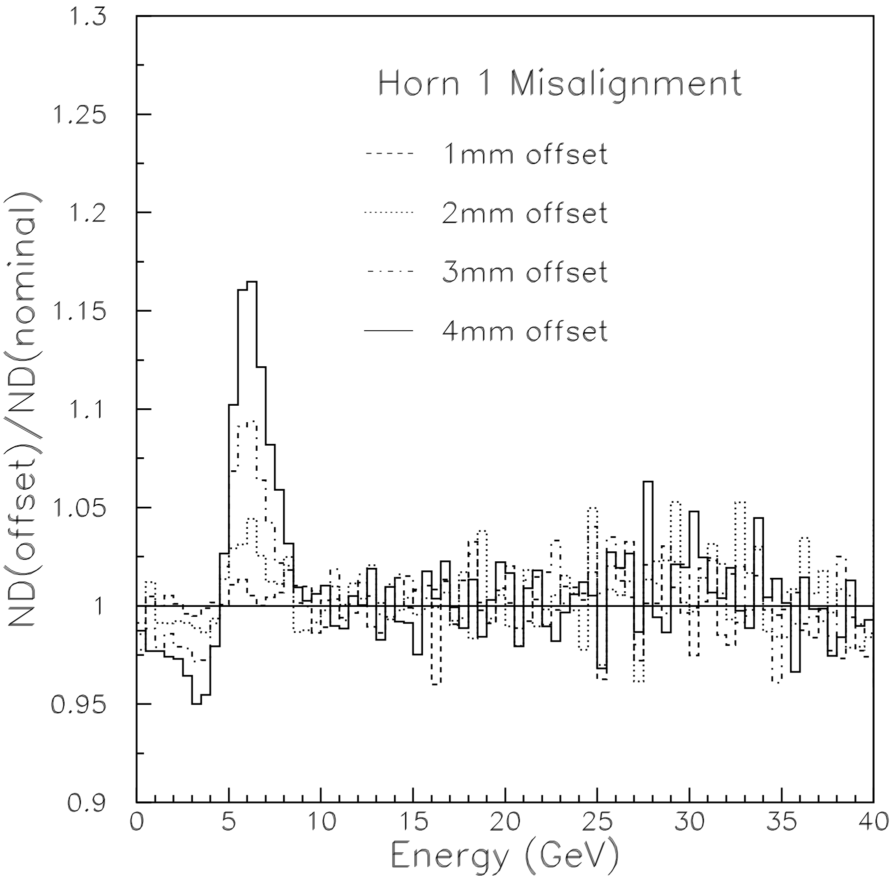

Figure 11.1: Results of a Monte Carlo simulation showing the effects of transverse displacements of Horn 1 on the neutrino beam spectrum at the MINOS near detector. Plotted is the ratio of neutrino flux with the offsets compared to the nominal flux. Each displacement reduces the number of neutrinos in the peak and increases the flux just after the peak, raising the average energy of the beam. (Simulation courtesy Z. Pavlovich.)

the horns, as do off-center horns. The effect of transverse offsets of Horn 1 on the near detector neutrino spectrum are shown in Fig. 11.1; greater displacements increasingly reduce the peak flux and increase the flux just above the peak, raising the average energy of the beam.

The secondary beam monitors play an important role in alignment of the target hall components. The basic method of beam-based alignment is to move the proton beam spot laterally along the target hall components, the effect of whose features upon the beam is measured by the downstream instrumentation. As mentioned in Ch. 1 and further described in App. B, instrumentation available in the primary proton beam line can accurately determine the position of the proton beam as it traverses the target hall. The downstream instrumentation then measure: (a) the attenuation of the proton beam as it passes through target and horn material; (b) particle showers created when the proton beam is incident on the target or horns; and (c) the divergence of the proton beam as it passes through target or horn material. Fig. B.2 in App. B is a layout of the beamline devices relevant to this procedure.

The particular attraction of beam-based alignment is that it provides direct measurements of the components relative to the beam. This avoids uncertainties associated with using the intermediate step of an optical survey network and reduces targeting errors. BBA also avoids the error noted in Ch. 10 of aligning the components based on maximal neutrino yield. To our knowledge, this kind of beam-based alignment has been performed previously only by Argonne National Laboratory (ANL), which temporarily placed a series of glass plates before and after the target and horns to observe the location of the radiation-blackening [161].

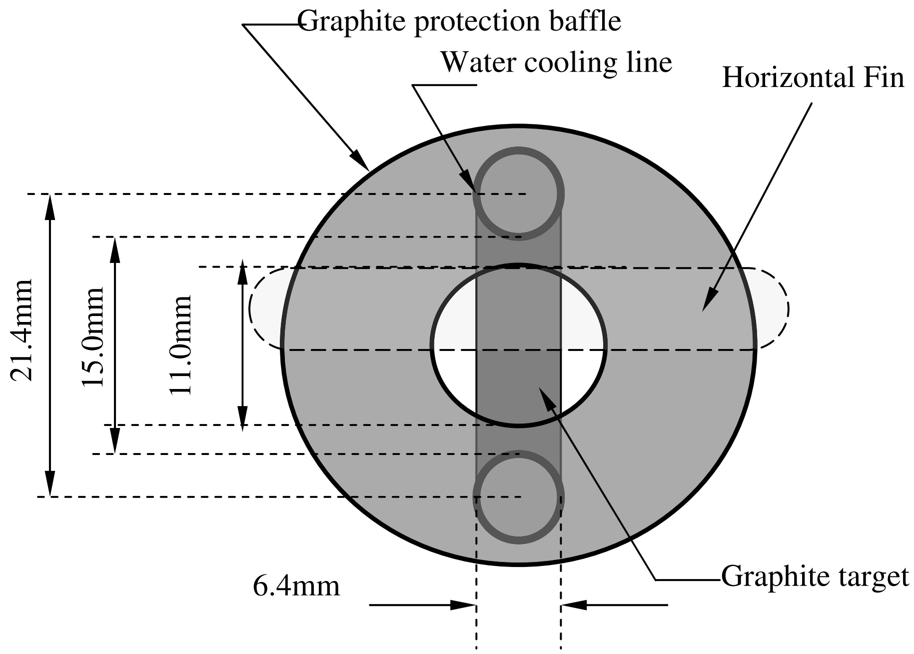

Figure 11.2: Beam's-eye view of the target-baffle system. The beam first passes the baffle, then a single horizontal graphite target fin, and then the target itself (47 fins). For the surveys in this section the proton beam is scanned across the target in a vertical or horizontal direction. The beam is attenuated and scattered by the material in its path, changing its intensity and angular spread as detected by the Hadron Monitor.

11.1.1 Target-Baffle System

A beam's eye view of the target-baffle system [162] is shown in Fig. 11.2. The target and baffle are mounted on a module that allows manipulation of the target and baffle in all three coordinates as well as pitch and yaw. However, the target and baffle are fixed with respect to each other and move as a unit.

The baffle is a graphite rod with a 11 mm diameter clear aperture. The baffle is 150 cm long, corresponding to 3.1 nuclear interaction lengths and 6.3 radiation lengths. When the beam is incident on the baffle material (and not the aperture) we expect only 4.5% of the proton beam to survive, and it to have acquired a RMS angle of 0.30 mrad due to multiple scattering. Both effects can be measured using the Hadron Monitor.

The target is positioned 68 cm after the baffle. The main part of the target consists of forty-seven 2 cm long graphite segments, spaced by 0.3 mm. The segments have a width of 6.4 mm and a height of 15 mm (all oriented vertically). The segments contribute 2.0 nuclear interaction lengths and 4.0 radiation lengths, such that 13.5% of the proton beam will penetrate the target and acquire a RMS angle of 0.24 mrad due to multiple scattering. The beam can be steered such that its center intersects neither the target nor baffle. This allows determination of the horizontal target position.

The vertical position of the target is not as crucial. To determine the vertical target position there is a 48^{th} fin, oriented horizontally, 15 cm upstream of the main target. This fin, called the "horizontal fin", is made of the same graphite material as the main target. It is centered 2.26 mm above the vertical center of the main target. The horizontal fin contributes 0.08 radiation lengths, measurable at the Hadron Monitor as an increase of the proton beam RMS scattering angle increases by 0.002 mrad.

11.1.2 Horns and Crosshairs

To allow beam-based alignment of the horns, a system of "crosshairs" was designed [162, 163]. Using these crosshairs and the Horn 1 neck, the locations of the upstream and downstream ends of each horn can be independently measured, giving a measurement of both position and angle of the horn. Figs. 11.3 and 11.4 show the crosshairs on the horns.

The crosshairs are three aluminum bars attached to the downstream face of Horn 1, and both faces of Horn 2. Each crosshair consists of a main spar oriented vertically and providing a width of 1 mm in the horizontal direction;

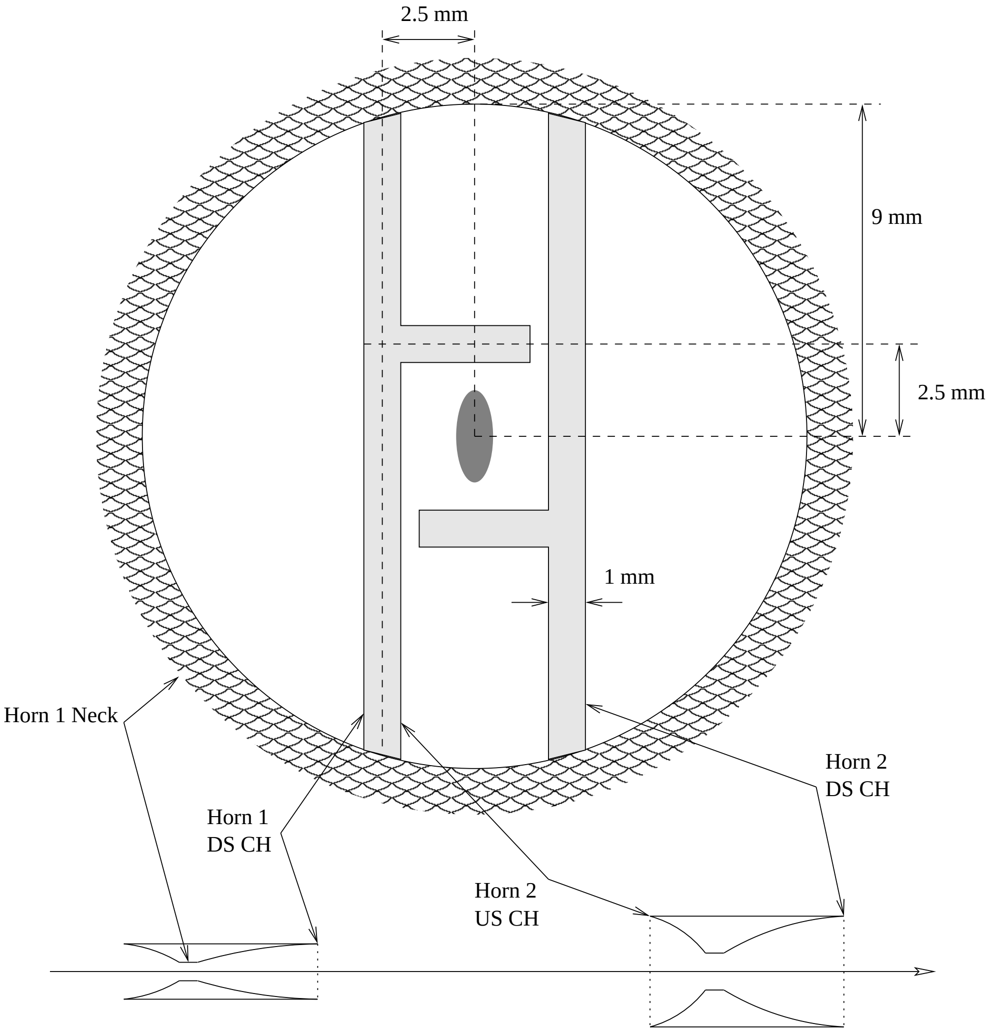

Figure 11.3: Beam's eye view of the Horn 1 neck and crosshairs. The proton beam is scanned horizontally and vertically over the neck and crosshair features (the ellipse in the center is the approximate 1 σ beam spot for these studies). The particle spray from interactions in the material produce signals in two loss monitors, one just downstream of each horn; the scattering of the beam also appears as larger beam size in the Hadron Monitor. There are three crosshairs that are either above and to the left of the beam (Horn 1 downstream and Horn 2 upstream), or below and to the right (Horn 2 downstream). Each crosshair consists of a vertical spar over the entire opening of the horn and a small horizontal nub; they establish the horizontal and vertical positions, respectively, of that end of the horn.

there is an additional horizontal nub that is 1 mm in width and extends 3.5 mm back toward the center of the beamline. In the beam's-eye view the Horn 1 downstream and Horn 2 upstream crosshairs overlap.

A pair of ionization chambers, one downstream of each horn (one is shown in Fig. 11.4), detects the scattered particles created when the proton beam is incident on the crosshairs. Additionally the Hadron Monitor can detect the increased divergence of the proton beam after penetrating the crosshairs.

11.1.3 Proton Beam Extrapolation

Extrapolation to the target was done with a simple linear model:

$$x(z) = x_{\mathrm{TGT}} + (x_{\mathrm{TGT}} - x_{121}) \times \left(\frac{z - z_{\mathrm{TGT}}}{z_{\mathrm{TGT}} - z_{121}} \right) \tag{11.1}$$

x_{121} and x_{TGT} are the measured positions at the 121 and TGT instrumentation stations (see Fig. B.2 & App. B); z_{121} and z_{TGT} are the positions of the instruments along the beamline; $x(z)$ is the extrapolated position at the longitudinal location z. The extrapolated resolutions from profile monitor (PM) measurements are estimated to be $\sim$ 32 μm for the baffle, target, and Horn 1; and 54 μm at Horn 2. For the BPM measurements the estimates are 260 μm and 450 μm at the two locations (resolutions are somewhat better for target scans where intensity is higher).

The PMs are used for the beam-based alignment measurements because of their better resolution at low intensity and fewer uncertainties in the translation from measured positions to nominal positions as defined by optical survey[1].

[1]During normal beam operations, the BPMs are used to steer the proton beam. Therefore, the target and horn positions are also given at the end of this chapter in BPM coordinates.

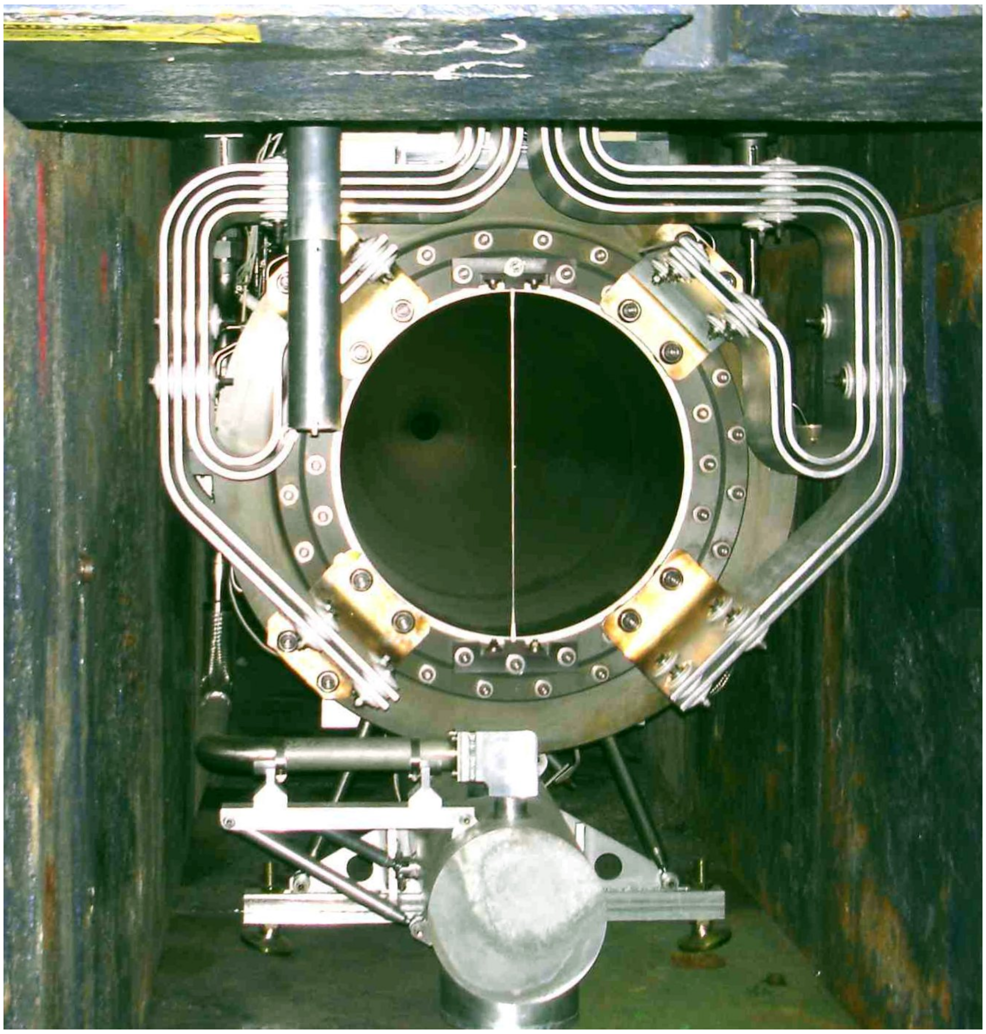

Figure 11.4: Picture of Horn 2, as installed, from the downstream end. The crosshair is the vertical bar across the horn aperture. The horizontal nub on the crosshair is just visible at the center of the horn aperture. The loss monitor ion chamber is the cylinder extending from above on the left (beam right).

11.2 Alignment of Target and Baffle

The target and baffle horizontal and vertical positions can be established by a single horizontal and a single vertical scan of the target. The angles of the target and baffle cannot be well measured, but the common angle of the two components can be considered degenerate with their differential offsets.

The scans used to determine the final position of the target were performed on April 25, and had a more-or-less nominal beam profile of 0.9 x 0.9 mm^2 sigmas. The scans from a prior installation of the target, March 3, are also shown. During the March 3 scans the primary proton beam vertical sigma was $\sim$ 1.4 mm, and horizontal sigma was $\sim$ 0.7 mm.

11.2.1 Vertical Target-Baffle Measurements

The vertical scan of the baffle is shown in Fig. 11.5. The charge collected in the Hadron Monitor is reduced as it passes into the baffle by absorption and scattering in the baffle material. The signal is fit to a constant multiplied by two error functions of oppositely signed arguments and the offsets allowed to vary.

The upper scan in Fig. 11.5 was performed on March 5, 2005 and has the vertical position of the baffle as +1.18 mm. The lower scan is from April 25, 2005 – after the target replacement – and results in the vertical position of the baffle being measured as +0.21 mm. This one mm vertical shift in baffle position was unexpected, but is consistent with other measurements and can be accommodated.

The vertical position of the target itself can be established by inferring the presence of the horizontal fin of the target. The Hadron Monitor RMS beam size and central pixel intensity are sensitive to the presence of the single

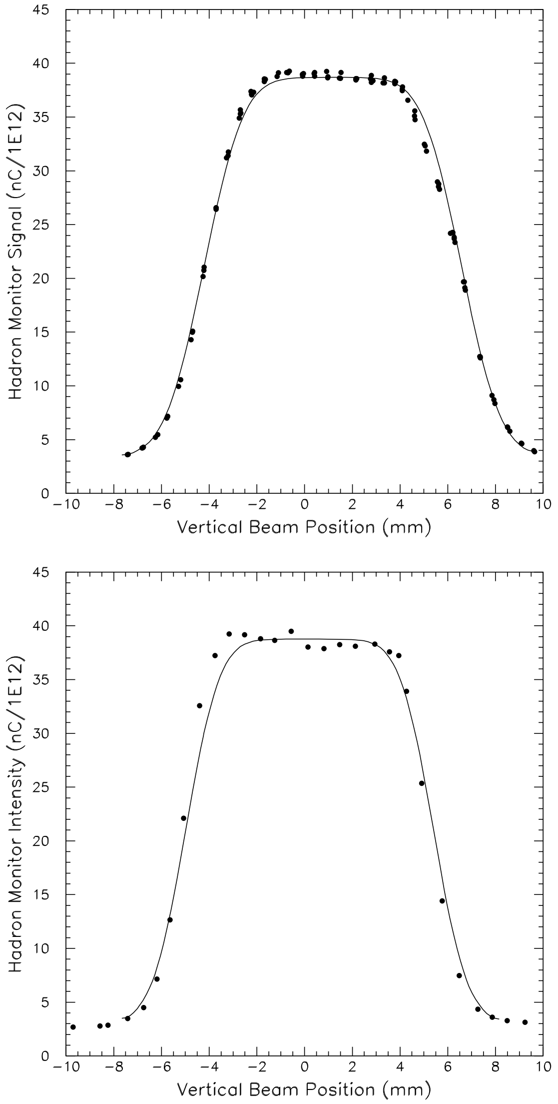

Figure 11.5: Two vertical scans of the target and baffle. The former being performed on March 5, 2005, the second performed on April 25, 2005. The data in each plot are the total amounts of charge collected on the Hadron Monitor, normalized by proton beam intensity, and plotted as a function of beam position at the target as extrapolated from profile monitor position measurements. The line is a fit to the data based on a simple model of proton absorption in the target and baffle. The fits show a 1 mm vertical shift of the baffle after re-installation.

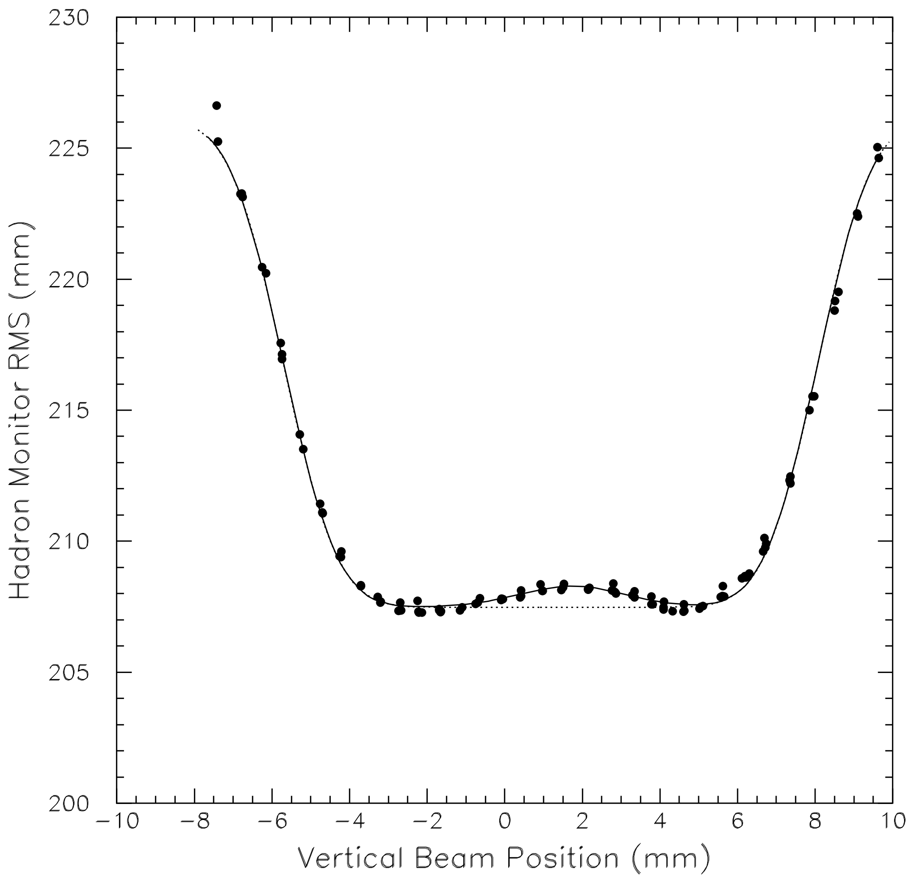

Figure 11.6: Information from the Hadron Monitor used to find the vertical target position. This plot corresponds to the upper scan in Figure 11.5, in which the beam is scanned vertically along the horizontal center of the target. Plotted is the vertical RMS of the measured distribution in the Hadron Monitor. The position of the horizontal fin is indicated by an increase in RMS due to scattering. The solid line is a fit to a simple model of absorption, with the dashed line showing the fit without the effect of the horizontal fin.

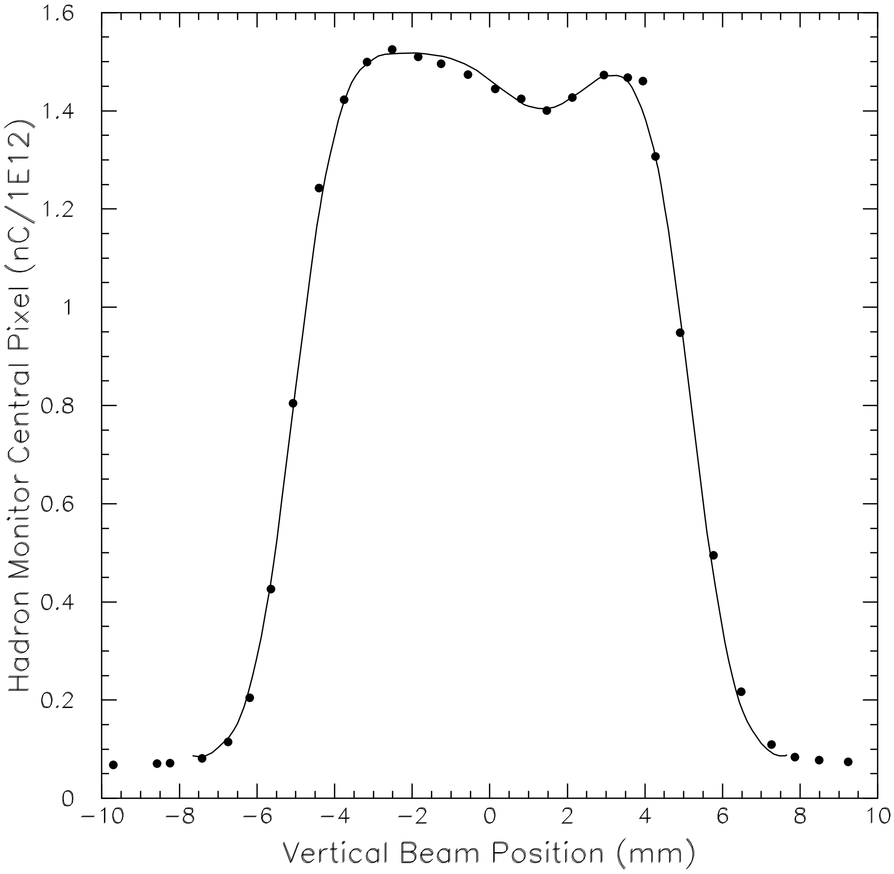

Figure 11.7: This plot corresponds to the lower scan in Figure 11.5. Plotted is the the charge collected on the center pixel of the Hadron Monitor, normalized by proton beam intensity. The position of the horizontal fin is indicated by a decrease in the charge on the center pixel due to scattering and absorption. The black line is a fit to the data considering the baffle and the fin. This scan also shows a vertical displacement as compared to 11.6 that occurred after the target re-installation.

fin, as the RMS responds directly to scattering; the central pixel responds to absorption the same as total intensity, but is more sensitive to scattering. Figs. 11.6 and 11.7 show these quantities as measured from the same vertical beam scans as before. In the March 5 scan the RMS of the distribution makes the best measurement of the fin position. The data are fit to a flat background plus two error functions[2] for the baffle and a gaussian for the fin, the value for the centroid of the gaussian is +1.75 mm. The response in RMS to the horizontal fin is not expected to be gaussian, but serves well enough as the beam width is comparable to the fin width.

For the April 25 scan, the central pixel turned out to be a better measure of the fin position, as shown in Fig. 11.7. The data are fit to a constant background minus two error functions for the baffle, and a gaussian from the fin, the value for the centroid of the gaussian is +1.36 mm. Thus, the target also sees a vertical change of position after re-installation, though only half as much as the baffle sees. The baffle position is measured as +0.07 mm (compared to 0.21 mm above). This discrepancy is likely from the quality of the fit and indicates a fitting uncertainty of about $\pm$ 0.1 mm.

11.2.2 Horizontal Target-Baffle Measurements

All the information used to find the horizontal positions of the target and baffle are shown in Fig. 11.8. During horizontal scans of the proton beam, the Hadron Monitor intensity varies widely with the amount of material traversed. The high contrast of baffle-gap-target-gap-baffle (Fig. 11.2) provides significant

[2]The increase in RMS from the baffle is not expected to be exactly an error function as the RMS is related to angle spread which will be related to the square root of material traversed and the Hadron Monitor RMS is poorly measured for very large widths, such the 230 mm is about the maximum value; however, the error functions fit well enough for these purposes.

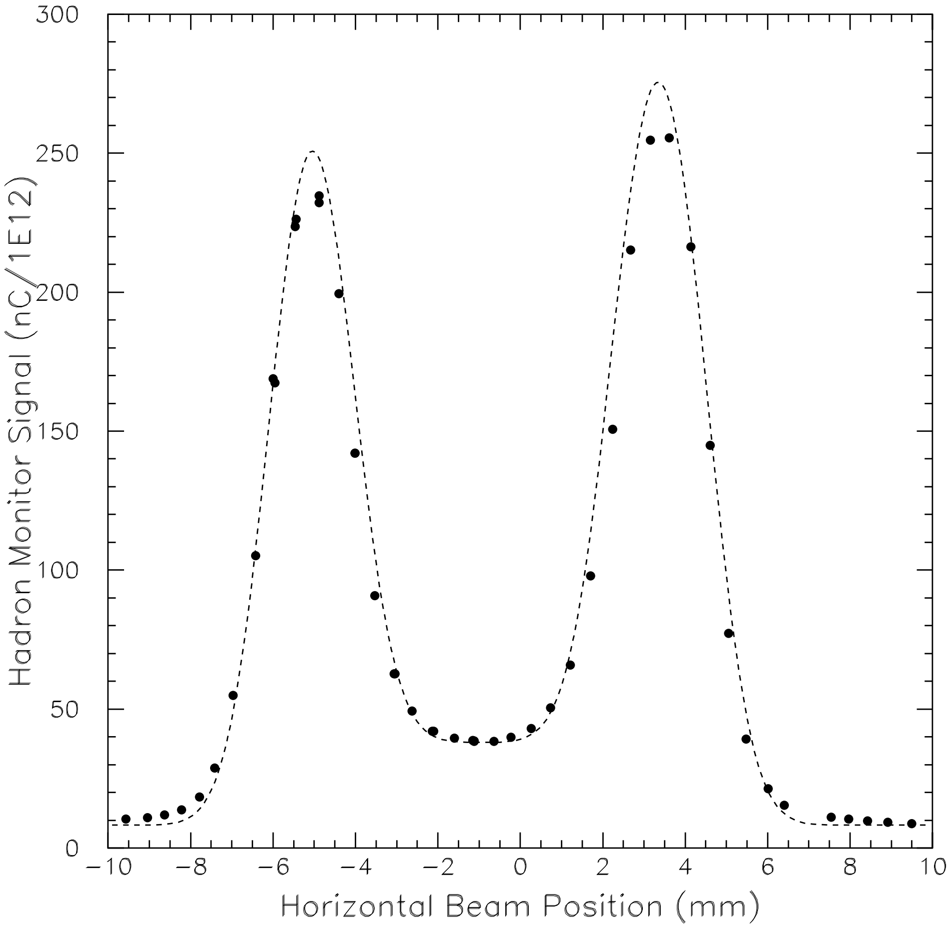

Figure 11.8: Data from a single horizontal scan that establishes the horizontal positions of the target and baffle. Plotted is the total Hadron Monitor charge normalized by proton intensity. The edges of the target are found by fitting to the central dip in charge; the edges of the baffle are found by fitting to the outer dips in charge.

modulation of the areal density, and thus in proton flux measured at the Hadron Monitor. The data is fit to a constant multiplied by two sums of two error functions. The target and baffle centers are fitted independently. The fitted baffle center is at -0.75 mm, the target center at -0.95 mm. The 0.2 mm offset between target and baffle results in the unequal signal amplitude of the two peaks corresponding to the gaps.

The fit to the baffle width gives 10.7 mm, suggesting that the beam was off-center by as much as 1.3 mm vertically, or at an angle up to 200 μrad. The proton beam position was -0.5 mm vertically, which with a 0.2 mm baffle offset, suggests the beam was off-center by only 0.7 mm, limiting the angle to 140 μrad.

The fitted reduction of signal in the Hadron Monitor is about 88.5% for the target, or 2.2 e-foldings, and 97.5% for the baffle, 3.7 e-foldings. The ratio between e-foldings is 1.7, compared to the anticipated absorption ratio of $150/96 = 1.56$. The Hadron Monitor intercepts a solid angle of $(1.06\ \mathrm{mrad})^2$, so scattering contributes to the further signal reduction in the baffle.

11.2.3 Other Indications

The Hadron Monitor is the preferred way to determine the target and baffle positions, but there are a few other corroborating measurements: the Muon Monitors and the Budal Monitors. The muon fluxes depend (in a nontrivial way) on how much material the proton beam traverses – affecting the muon flux measured. Details of the muons response during target scans are give in §10.2. For the purposes of beam-based alignment we note that the target scans also have identifiable features in the Muon Monitors that provide a cross-check on the Hadron Monitor measurements.

The Budal Monitors [28], another set of ionstrumentation, are intended to give direct electrical signals from the target that correspond to the amount of beam passing through; results of this system were mixed. A Budal Monitor is made by electrically isolating the target (or some part of the target) and measuring the charge ejected from it by the passing proton beam. The NuMI target has two Budal Monitors: one being the main span of the target itself and the other being the horizontal fin. The horizontal position of the target can be found by performing a horizontal scan and reading out the main span of the target. This technique produces measurements consistent with the target scan, but with less contrast. To find the vertical position, the intention was to perform a vertical scan and read out the horizontal fin signal. Unfortunately, this signal was dominated by the spray off of the baffle and was not usable to find the vertical position of the target. This was likely in part due to the fin's short length as compared to the target and baffle, not allowing a proper signal to develop.

11.3 Alignment of Horns

The horizontal and vertical angles and positions of the horns are determined with a set of scans performed with the target removed and the horns off. Three scans (one vertical, and two horizontal) are used to to find the positions of the Horn 1 neck and the long spars and nubs of three crosshairs (see §11.1.2). All of the scans for horn alignment were performed prior to target re-installation and have the anisotropic beamspot of 0.7 mm horizontal sigma and 1.4 mm vertical sigma. The beamspot has an impact as feature can appear sharper or narrower.

In these analyses the proton beam was always projected to the position of the feature in question in the profile monitor coordinate system. However, as a single detector may be used for multiple purposes a place of extrapolation much be chosen for plotting. In the plots show here the beam is extrapolated to: the Horn 1 downstream crosshair for the Horn 1 loss monitor; the Horn 2 upstream crosshair for the Horn 2 loss monitor; and the Horn 2 downstream crosshair for the Hadron Monitor RMS. The features on the plots visible in different devices will not always line up due to the different points of extrapolation.

11.3.1 Horizontal Horn Measurements

A typical horizontal scan is shown in Fig. 11.9. The data from the three detectors are plotted: "Horn 1 LM" is the loss monitor just downstream of Horn 1 and is sensitive to particle spray from the Horn 1 neck and the downstream crosshair on Horn 1. "Horn 2 LM" is the loss monitor just downstream of Horn 2 and is sensitive to the Horn 1 neck and both crosshairs on Horn 2. The signal from the upstream Horn 2 crosshair is about eight times as strong as that from the downstream crosshair, leading to difficulty in differentiating the signals[3]. "HadMon RMS" is the vertical RMS about the centroid of the beam at the Hadron Monitor and is sensitive to all material, so is only useful where crosshairs do not significantly overlap.

[3]From the crosshair lengths, we calculate that the material presents 1.5% (4.6%) of a nuclear interaction length and 7% (20%) of a radiation length for a 6 mm (18 mm) long crosshair. Because of the 1 mm width, the entire beam will never be incident on the crosshair material. So, there is an upper bound of 1.5% (4.6%) of the beam attenuated in the crosshairs, which is not measurable in the Hadron Monitor; however, most of that attenuation will occur through inelastic collisions producing charge particles at larger angles which can be measured by the loss monitors. The electromagnetic scattering induces an angular spread of as much as 0.027 (0.048) mrad, which corresponds to 20 (35) mm at the Hadron Monitor. The increase in angular spread will depend on the pre-existing divergence without the material.

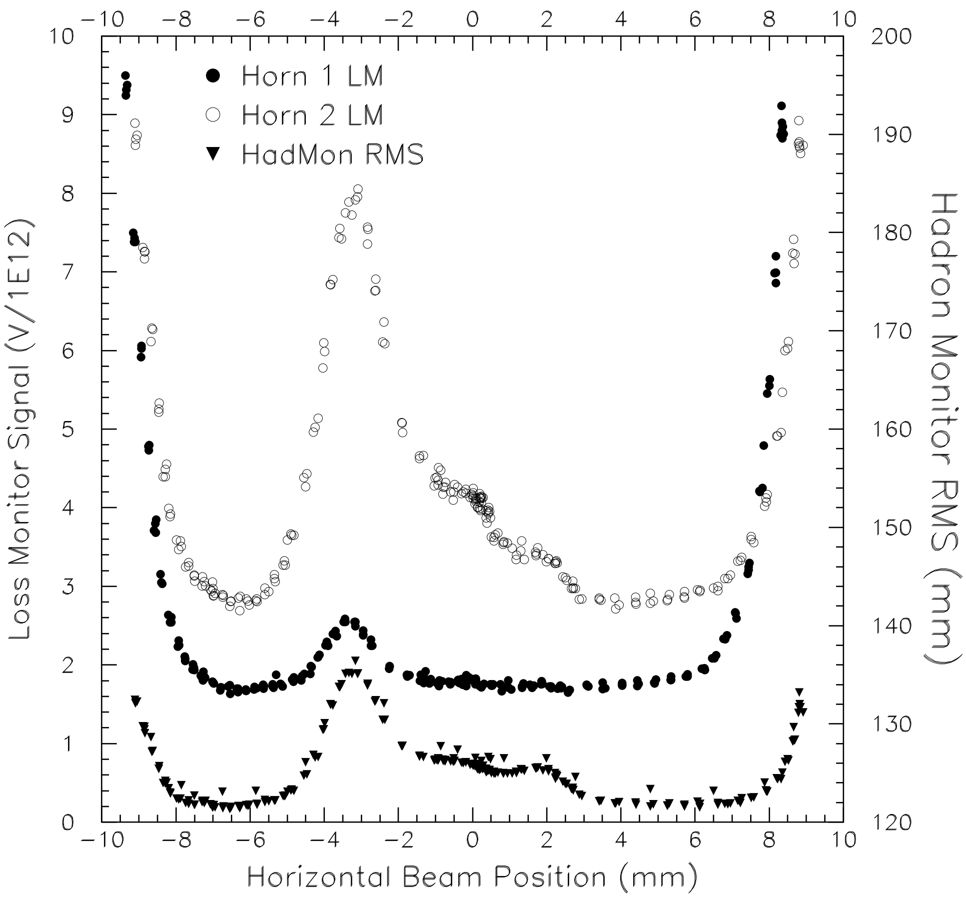

Figure 11.9: Summary of measurements made to establish the horizontal positions and angles of the horns. Shown are the signals in the two loss monitors, one downstream of each horn, and the vertical RMS of the distribution in the Hadron Monitor, about the centroid, during a horizontal scan of the proton beam across the system.

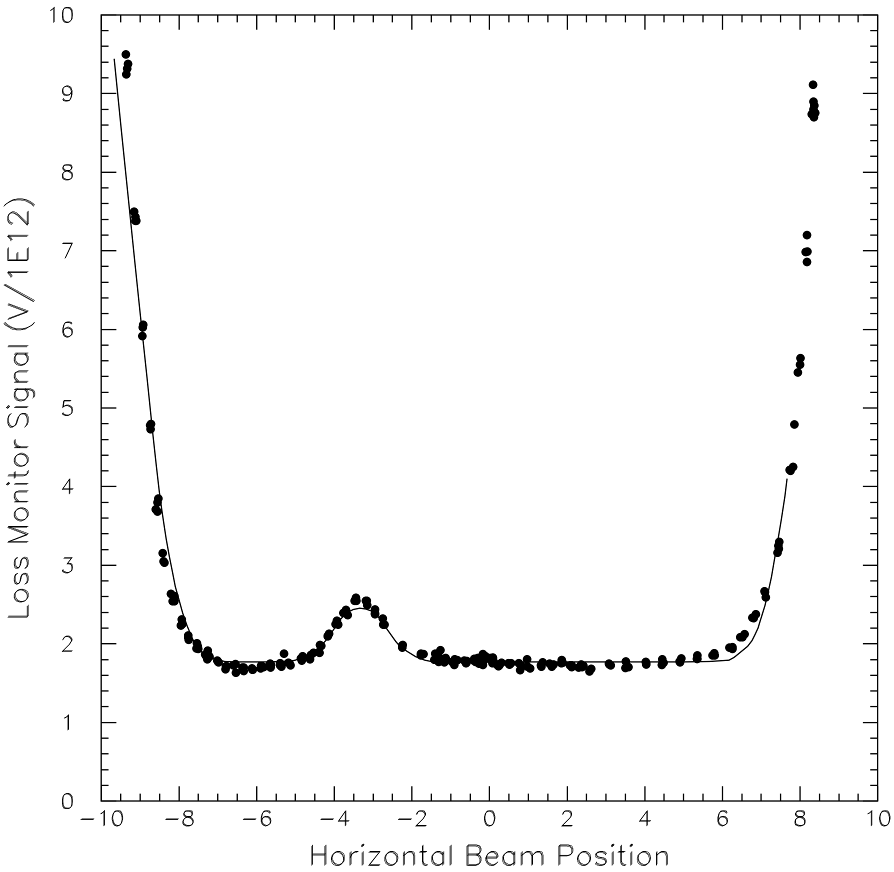

Figure 11.10: The data from Fig. 11.9 for the Horn 1 loss monitor alone. Superimposed on the data is the fit used to determine the center of the horn neck from the outer peaks, and the crosshair position from the central peak.

The Horn 1 loss monitor data are shown alone on Fig. 11.10 along with the fit used to determine the horizontal position of the horn neck and crosshair. The fit is a constant background plus two error functions for the horn neck, and a gaussian for the crosshair. The fit results in a position of -0.46 mm for the horn neck and -3.36 mm for the downstream crosshair.

The Horn 2 crosshair signals are more difficult to separate. The signal induced by the upstream crosshair is eight times as strong, so even the small amount of beam crossing the nub causes interference. The problem can be seen in the three-humped structure of the Horn 2 LM signal in Fig. 11.11. A similar scan is show as the open circles in Fig. 11.11, it is displaced a fraction of a millimeter downward from the other scan. This distribution is fit to a linear background (to approximate the nub) and a gaussian. The gaussian gives the upstream crosshair position as -3.35 mm.

A second horizontal scan, with the beam displaced a further 4 mm downward, is used to separate the downstream crosshair signal. This scan (the filled circles in Fig. 11.11) is well clear of the upstream nub, allowing the downstream crosshair to be resolved and fit to a linear background and a gaussian. The fit gives the downstream crosshair as +1.33 mm.

11.3.2 Vertical Horn Measurements

The data from a single vertical scan of the proton beam across the horn system are shown in Fig. 11.12. Here the beam scans vertically along the centerline of the horn neck – avoiding the crosshair spars as much as possible. The signal peaks represent the horizontal crosshair nubs; the peaks are wider and less intense than in the horizontal scans because of the smaller amount of material in the nub and because of the wider vertical beam size.

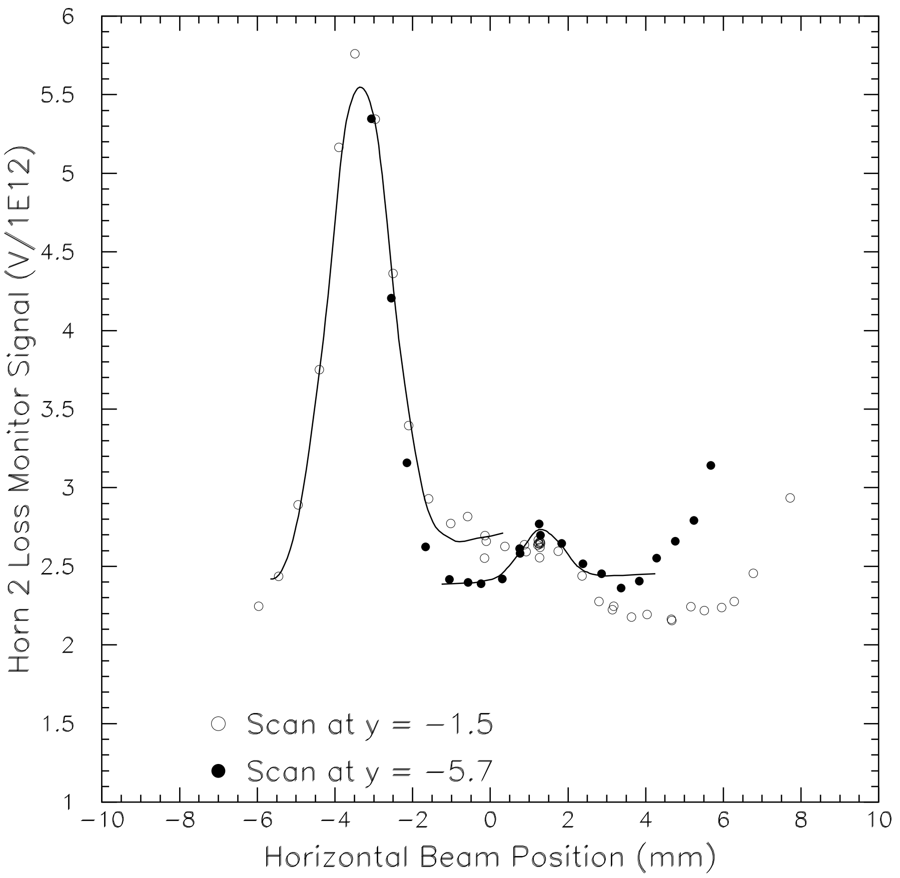

Figure 11.11: Data from two separate horizontal scans of Horn 2 that establish the positions of the two crosshairs. The two scans were done at different vertical positions, as indicated. In the higher scan the signal from the upstream nub interferes with the signal from the downstream crosshair; the second scan, displaced downward, avoids this problem.

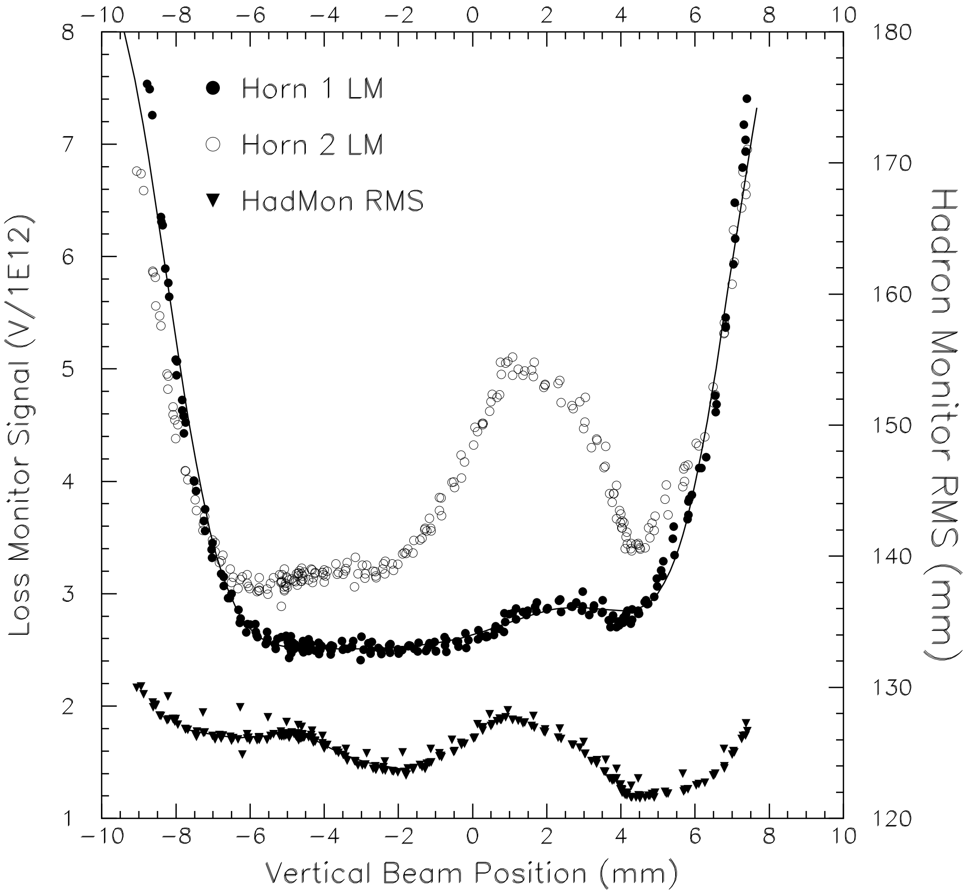

Figure 11.12: Summary of measurements made to establish the vertical positions and angles of the horns. The signals are the same as in Fig. 11.9. The Horn 1 neck and crosshair nub positions are measured as in Fig. 11.10. The Horn 2 upstream nub position is measured from the loss monitor signal, as in Fig. 11.11. The Horn 2 downstream nub position must be measured by the change in measured RMS at the Hadron Monitor.

The Horn 1 downstream crosshair nub and neck positions are found by a fit the Horn 1 loss monitor data. The neck, as before, is clearly visibly, but the nub provides only a weak bump. Regardless, the data is well fit by two error functions for the neck and a gaussian for the nub. The fit value for the nub position is +2.78 mm; other scans give positions of +2.02 and +2.04 mm, so an average is used. The fit value for the neck position is -0.64 mm.

The Horn 2 upstream crosshair nub is found by a fit to the peak in the Horn 2 loss monitor data. The fit is done to a linear background and a gaussian over a range excluding most of the neck signal. Two scans were fit giving positions of +1.48 and +1.57 mm.

The Horn 2 downstream crosshair nub is not resolvable in the Horn 2 loss monitor data. The expected position is around -3 mm, and a slight bump is perhaps visible there in the loss monitor data, but is dominated by background. However, the Hadron Monitor data can be used here as there are no other features competing with the nub. A fit is performed to the Hadron Monitor RMS with a linear background and a gaussian over a 6 mm range about the peak in the RMS. The fitted value for the position of the nub is -4.74 mm.

11.4 Final Results

The position measurements in the previous sections are used to find the positions (and some angles) of the target hall components. Some of the measured positions are of features that are not expected to be in the middle of the beam-zero (such as the horizontal fin and crosshairs). The raw data for measured positions are shown in Table 11.1.

In some case fits from multiple scans and different measures were used

Feature	**Scan**	**Fit Value**	**Orth Pos**	**Fig #**
Baffle Neck	Horz	-0.75 mm	-0.5 mm	11.8
Baffle Neck	Vert	+0.21	-1.1	11.5
Baffle Neck	Vert	+0.07	-1.1	11.7
Target	Horz	-0.95	-0.5	11.8
Target	Vert	+1.36	-1.1	11.5
Horn 1 Neck	Horz	-0.46	-0.3	11.10
Horn 1 Neck	Vert	-0.64	-0.2	11.12
Horn 1 Downstream Crosshair	Horz	-3.36	-1.2	11.10
Horn 1 Downstream Crosshair	Horz	-3.32	-0.3	
Horn 1 Downstream Nub	Vert	2.78	-0.2	11.12
Horn 1 Downstream Nub	Vert	2.02	-0.4	
Horn 1 Downstream Nub	Vert	2.04	-1.4	
Horn 2 Upstream Crosshair	Horz	-3.35	-1.5	11.11
Horn 2 Upstream Nub	Vert	+1.48	-0.1	
Horn 2 Upstream Nub	Vert	+1.57	-1.1	
Horn 2 Downstream Crosshair	Horz	+1.33	-5.7	11.11
Horn 2 Downstream Nub	Vert	-4.74	+0.1	11.12

Table 11.1: Summary of raw measurements used to establish the positions of the target and horn, and the positions and angles of the horns. The measurements are broken down by feature and type of scan; in some cases multiple measurements of the same feature were used. The Profile Monitor (PM) coordinated system is used here. "Orth Pos" refers to the scan position in the opposite coordinate; e.g. the horizontal scan used to establish the position of the baffle neck occurred at the y position of -0.5 mm. The figure number is also provided for fits shown in the previous sections.

Device	Direction	Coord	Offset (mm)	Angle (mrad)
Baffle	Horz	PM	-0.75	
Baffle	Horz	BPM	-1.21	
Baffle	Vert	PM	+0.14	
Baffle	Vert	BPM	+1.12	
Target	Horz	PM	-0.95	
Target	Horz	BPM	-1.41	
Target	Vert	PM	-0.90	
Target	Vert	BPM	+0.13	
Horn 1	Horz	PM	-0.65	-0.18
Horn 1	Horz	BPM	-1.24	-0.18
Horn 1	Vert	PM	-0.33	+0.20
Horn 1	Vert	BPM	+0.81	+0.26
Horn 2	Horz	PM	-1.01	-0.11
Horn 2	Horz	BPM	-1.82	-0.18
Horn 2	Vert	PM	-1.61	-0.42
Horn 2	Vert	BPM	+0.08	-0.43

Table 11.2: Positions and angles of target hall components calculated from the positions measured in Table 11.1. The positions in the profile monitor (PM) coordinate system are calculated directly from the fitted values, the beam position monitor (BPM) coordinate system includes the difference in extrapolated position measured in each run. There is only one position measurement in each direction for the target and baffle, so there is no angle information.

to find a feature. For example, the baffle vertical position was fit to both total Hadron Monitor intensity, and the intensity measured on the central pixel. In the case of the Horn 1 downstream nub, independent measurements differed substantially ($\sim$ 0.7 mm). This sets a reference for the some of the systematic uncertainty involved in the beam-based alignment procedure. Within a certain coordinate system, we estimate an uncertainty $\pm$ 0.3 mm for the target, baffle and Horn 1; and $\pm$0.5 mm for Horn 2.

Using the knowledge of feature offsets from component centers and averaging the numerous results we get the offsets listed in Table 11.2. Additionally, for the horns where we have measurements separated by 2-3 m, we can calculate the angles of the components with respect to the beamline nominal. We estimate uncertainties of $\pm$ 0.2 mrad for the angles of the horns. The Horn 2 vertical measurements (particularly the angle) may be biased by the large offset measured in the Horn 2 downstream nub, which could not be seen in the loss monitors – only in the Hadron Monitor RMS.

All of the previous results were in the profile monitor (PM) coordinate system, but as the BPMs are used to steer the beam we translate them into BPM coordinates. Here, the translation is done on a scan-by-scan basis. While the resolution of the BPMs is poor at these intensities, an average offset is measured for the entire run. A preferred method would be to use comparisons that are robust up to normal operating intensities, but such comparisons are not yet available.

The PM and BPM coordinate systems are clearly diverging in Table 11.2 such that there is a 1.7 mm difference in the vertical direction at Horn 2. This is another uncertainty in the measurements of the target hall components. Generally, it has been our bias that the offsets from optical survey necessary for making PM measurements are better understood than those for BPM measurements. So, we feel that the PM measurements better correspond to reality. However, the BPMs are used to steer the beam during typical running, so we must understand the offsets in their coordinate system, as we do in the next section.

11.5 Interpretation of Results

The position deviations arrived at in §11.4 range up to 2 mm; larger than might be desired. To quantify any effect, we use simulation to find the beam changes due to any offset or angle. Additionally, in some cases remediative

Device	Dir.	Offset	Effect	Angle	Effect
Baffle	Horz	-1.2 mm	2.5%	-0.1 mrad	$<$ 0.1%
Baffle	Vert	+1.1	2.2	-0.7	$<$ 0.1
Target	Horz	-1.4	2.5	-0.1	$<$ 0.1
Target	Vert	+0.1	$<$ 0.1	-0.7	0.3
Horn 1	Horz	-1.2	1.1	-0.2	0.3
Horn 1	Vert	+0.8	1.4	+0.3	0.4
Horn 2	Horz	-1.8	1.2	-0.2	$<$ 0.1
Horn 2	Vert	+0.1	$<$ 0.1	-0.4	$<$ 0.1

Table 11.3: Tabulation of effects on the Far-to-Near ratio due to misalignments if the beam were steered at (x,y) = (0,0). The offsets are those from 11.2 for BPMs, except for the target-baffle angle where a representative angle is chosen as that that would give the different offsets measured, instead of an actual physical offset between the centers of the two. There is an approximate budget of 2% for effects on the ratio, so some of these offsets are intolerable. See Table 11.4 for the aligned version.

action can be taken to correct offsets. We expect that in the long term the MINOS experiment may accumulate enough data such that statistical errors in flux measurements will approach 2% (2500 events) in its most populated energy bin; as such, we adopt a budget of $< 2\%$ for systematic errors in the far-to-near extrapolation from alignment uncertainties.

A Monte Carlo parametrization was used to understand the effect of alignment tolerances ([162], §4.6). In particular, the the effect on the ratio of fluxes measured at the near and far detectors (far-to-near ratio) was considered. The far-to-near ratio will vary with the focusing as the average angle of the pions to the detectors will change (differently for the two detectors); generally, as more higher-energy pions survive to decay to the detectors they have a lower ratio of far-to-near as they decay closer to the near detector.

The effects to the far-to-near ratio were fit and parametrized to the offsets with a power law in [162]. In Table 11.3 we use these parametrizations to arrive at estimated fluctuations induced in the far-to-near ratio from the

Device	Dir.	Offset	Effect	Angle	Effect
Baffle	Horz	0.0 mm	$< 0.1\%$	-0.1 mrad	$< 0.1\%$
Baffle	Vert	+0.1	< 0.1	-0.7	< 0.1
Target	Horz	-0.2	0.4	-0.1	< 0.1
Target	Vert	-0.9	< 0.1	-0.7	0.3
Horn 1	Horz	-0.0	< 0.1	-0.2	0.3
Horn 1	Vert	-0.2	< 0.1	+0.3	0.4
Horn 2	Horz	-0.6	0.2	-0.2	< 0.1
Horn 2	Vert	-0.9	0.4	-0.4	< 0.1

Table 11.4: Tabulation of effects on the Far-to-Near ratio due to misalignments if the beam were steered at (x,y) = (-1.2,+1.0) – as it is actually, unlike the case of Table 11.3. None of the uncertainties due to offsets or angles approach the alignment budget of 2% effect on the ratio.

offsets measured in this section. The BPM system is used to steer the beam, so we analyze the data in the BPM coordinate system and consider the case of the beam being steered toward "0". While angles were not measured for the baffle and target, we estimate an upper bound by considering the difference in baffle and target offsets to be due to a common angle in their mounting. There are several large effects in the ratio here, mostly resulting from horizontal offsets, particularly those of the target and baffle which exceed the 2% budget. This is not surprising as a 1.1 mm targeting error is large on a target of width 6.4 mm.

In reality, the beam is not steered at the 0 position. The target scans clearly show a horizontal offset of the target and a vertical offset of the baffle. Using this knowledge, the beam has been steered at the x = -1.2 mm and y = +1.0 mm position. We recalculated the offsets from the beam location and show those in Table 11.4 along with the recalculated effects of the far-to-near ratio. In this case, all effects contribute less than 0.5%, ensuring that the budget is not approached.

Should a second attempt at beam based alignment be made in the future several aspects of the present analysis could be optimized:

- Many of the scans were performed with improper optics in the primary proton line. The new optics setting results in a rounder, smaller beam spot that can be better extrapolated throughout the target hall.

- The correspondence between BPMs and PMs is still not established robustly at all intensities. Additionally, the extrapolation to the Hadron Monitor is not fully understood (see §B.3).

- Simulation of these scans in the detectors is not complete. The contrast is great enough that feature positions can be measured, but good simulation is desirable to understand the profile details in the detectors. Additionally, functions derived from simulation with variable parameters should be used for the fittings, instead of the *ad hoc* functions used in this chapter.

Addressing the above concerns will yield more reliable results. Additionally, we can prescribe several modifications to any future beam based alignment procedure:

- Scans should be operated using as many profile monitors as possible. These can be used to constrain the primary beam optics.

- Detailed knowledge of the primary beam optics should be used to calculate more complicated magnet variations to move the beam parallelly across the aperture. The existing bumps cause the beam to acquire small angles (i.e., are not entirely parallel) and induce some motion in the opposite plane because of quadrupole steering.

- Operation of the PMs, BPMs, and toroids in a high precision mode is very valuable. After the initial beam based alignment, the dynamic range of these instruments was altered in such a way that reduces the precision at low beam intensity. The high precision for low-intensity running must be reestablished, if not improved.

- Multiple scans should be performed for each component feature, but the scans should have different positions or attitudes. For example: horizontal scans with a 0.1 mrad difference in proton targeting will allow better measurements of component angles, as well as provide limits on the precision of proton beam extrapolation from the primary beam instrumentation.

- Alignment scans should be performed with a granularity of $\ll 0.5$ mm per step, and multiple data points per step. Subsequent scans intended for monitoring target integrity have usually been inadequate for alignment purposes. Use of an automated scanning program that can run in parallel with normal antiproton stacking should be investigated.

In the future, we hope beam-based alignment can be revisited in a more comprehensive manner. For the present, beam-based alignment has been used only for choosing the position at which to target the proton beam. When the measured offsets are better understood, the components may be repositioned with respect to each other and angled to reduce any possible effect on the neutrino flux. Such alignment may be required if the run of MINOS is extended and the statistical uncertainties become smaller.

Appendix A

Glossary of Accelerator Physics Terms

The primary tool of the modern accelerator complex is the synchrotron. A synchrotron uses magnets to bend a beam of particles in a repeating, somewhat circular path while also focusing the beam and using radiofrequency (RF) acceleration. This appendix reviews magnetic focusing and RF acceleration, and describe their integration as the synchrotron.

A.1 Particle Beam Optics

High-energy charged particle beams are directed with magnets. Dipole magnetic fields bends the beam, changing its direction. In a constant, uniform dipole magnetic field a charged particle will be bent transversely to the magnetic field and the particle's motion according to the Lorentz force law. If the particle's motion is entirely perpendicular to the magnetic field it will be bent into a circle. This behavior is called gyration or cyclotron motion; the radius of gyration is given by the equation:

$$r = \frac{p}{qB} = \left(3.34\,\frac{\mathrm{m}\cdot\mathrm{T}}{\mathrm{GeV/c}}\right)\frac{p}{B} \tag{A.1}$$

Where B is the magnetic field, p is the particle momentum, and q is the particle charge – we will be concerned with singly-charged particles and so will often neglect q. For example, in the case of the Booster, $p = 8.9$ GeV/c and $B \approx 0.7$ T, so the bending radius is 44 m. Any momentum a particle has parallel to the magnetic field will be unchanged; in this case the particle follows a spiral path along a magnetic field line.

Weak Focusing

A beam is made up of a large number of particles that have some small variation in position and direction; maintaining the beam size requires focusing the beam. Any particle momentum parallel to the magnetic field will continue, so the beam would diverge indefinitely in a uniform field. The beam motion parallel to the beam, however, does not diverge indefinitely. If all the particles have the same momentum they will have the same radii of gyration and cyclotron frequencies. For example, for particles starting at the same point with a small angle between their motion of α, the maximum distance between them will be $2r\alpha$. This is form of focusing intrinsic to radial motion in a dipole field; it is a component of "weak focusing" that was used in early accelerators.

Weak focusing is accomplished by introducing a gradient in the magnetic field to redirect the vertically diverging particles toward the center orbit. A profile of such a magnet is shown in Fig. A.1 in the center row – right; the pole faces of the magnet are sloped to produce gradients in the magnetic field. If a positive particle beam is directed into the page it will be bent to the right by the dominant dipole portion of the field. There will also be a horizontal

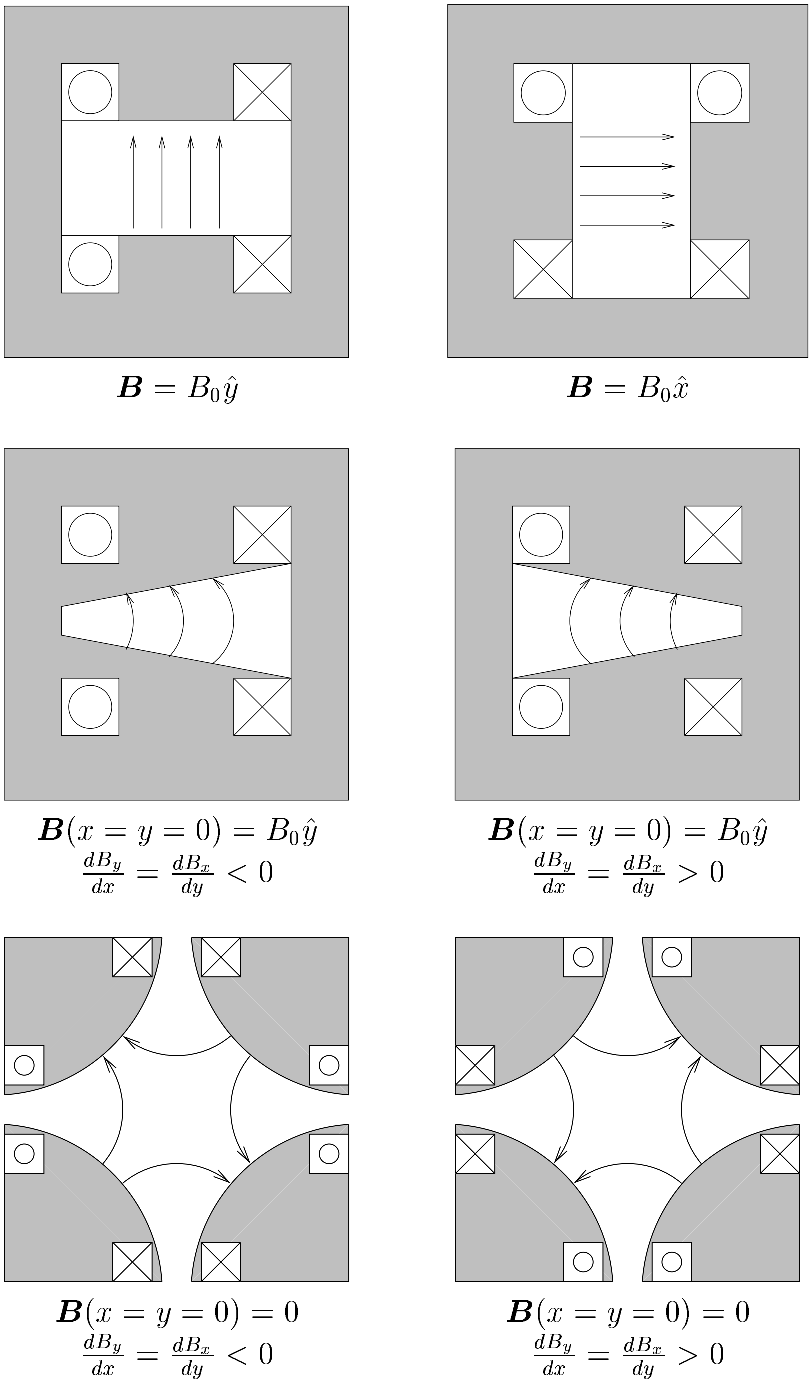

Figure A.1: Common types of magnets. The top row are dipole bending magnets; horizontal and vertical. The middle are combined function magnets that bend horizontally and focus; the left being F, the right D. The bottom row are quadrupole, pure focusing magnets; the left being F, the right being D.

component of the magnetic field directed to the right above the centerline, and to the left below the center. This field gradient produces a return force toward the central orbit. With the magnet face rotated around the center of gyration it produces a stable vertical oscillation around the central orbit. Note however that the implicit horizontal focusing force is reduced as there is also a horizontal gradient in the vertical field. Any strengthening of the vertical focusing reduces the horizontal focusing – thus the nomenclature of "weak" focusing. The constraints of weak focusing limited beam intensities until the development of strong focusing.

Strong Focusing

The ideal method of focusing would be to have a strong return force toward the central orbit. This would require the gradients:

$$\frac{dB_y}{dx} < 0 \qquad \frac{dB_x}{dy} > 0 \tag{A.2}$$

The quadrupole magnet, as shown in Fig. A.1, can be used as a focusing lens. The magnet has four poles generated by currents electrical currents outside of the beam path. The very center of the magnet has zero field with gradients in each direction. These gradients, however, must satisfy Ampere's law, so one of them is in the opposite direction required to focus. The magnet focuses in one plane and defocuses in the other. Nevertheless, magnets focusing in the horizontal plane (F-type) can be used in combination with those focusing in the vertical plane (D-type) to produce a net focusing. The combined focusing from these magnets is stronger than the weak focusing above in that it can compress a given particle distribution to a smaller transverse size. This method of focusing is known as alternating-gradient or strong focusing.

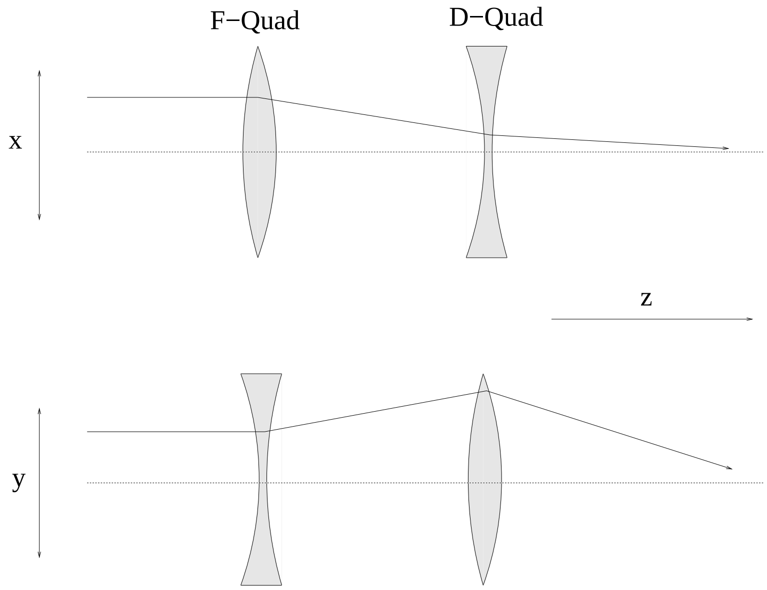

Figure A.2: Demonstration of strong focusing. A beam initially off the central orbit horizontally and vertically travels through an F type quad and a D type quad. Horizontally, the beam is first focused then defocused; vertically, the beam is defocused then focused. However, the angle of focusing is greater than the defocusing in both cases.

The dynamics of horizontal and vertical motion of the beam about a central orbit are, to first order, uncoupled and can be treated individually.

The basic arrangement of magnets for strong focusing is illustrated in Fig. A.2. The beam travels through one magnet of each type – separated by some distance. The beam is both focused and defocused in each direction; i.e. it is kicked toward the center and away from the center[1]. However, the kicks are of uneven magnitude – the kick toward the center is always greater. This is accomplished by keeping the separation between the magnets less than the focal length of the magnets. The beam that is first focused is nearer to the center when it receives the defocusing kick so it is smaller; similarly, the beam

[1]Magnets are shown here using the metaphor of thin lenses. The magnets do actually act much like cylindrical lenses, particularly if the magnets are short. In crowded rings where the magnets take up much of the space the dynamics need to include the finite length of magnets. However, the concept is much the same as that described here.

that is defocused first is further from the center so its focusing kick is larger.

A continuing pattern of these magnets is know as a FODO lattice – where F and D signify the magnets and O signifies a non-focusing region. Other forms of strong focusing lattices exist, but are based on the same principle as the FODO lattice. These lattices can be used in transfer lines where the O is usually a drift region and the focusing magnets are quadrupoles. They also can be used in rings where either the focusing magnets are quadrupoles and the Os are dipoles, or the focusing magnets can be combined function magnets that both bend and focus. At Fermilab, the Booster is an example of a combined function machine, and the Main Injector is an example of a separated function machine.

Beam Emittance & Envelope

Above, we considered a few particles off the central orbit as examples of how focusing works. A better tool for keeping track of the particles within a beam is emittance. Emittance is a measure of the six-dimensional phase space occupied by the beam particles. Generally, the dynamics of horizontal, vertical, and longitudinal phase space can be separated. Horizontal and vertical act similarly, but are mostly decoupled; longitudinal will be discussed in §A.2.

Transverse emittance, ϵ, is measured as the phase space area occupied by the beam[2] The phase space variables used are commonly beam location and divergence – x and dx/dz. Typically, the beam particles will occupy some ellipse in this space. When traveling along a transfer line or properly matched in a ring the moments of the beam at any point in the line are defined by the

[2]The precise definition of emittance varies – usually changing to account for some fraction of the beam. Here we will adopt the definition such that σ_x in Eqn. (A.3) corresponds to 68% of the beam.

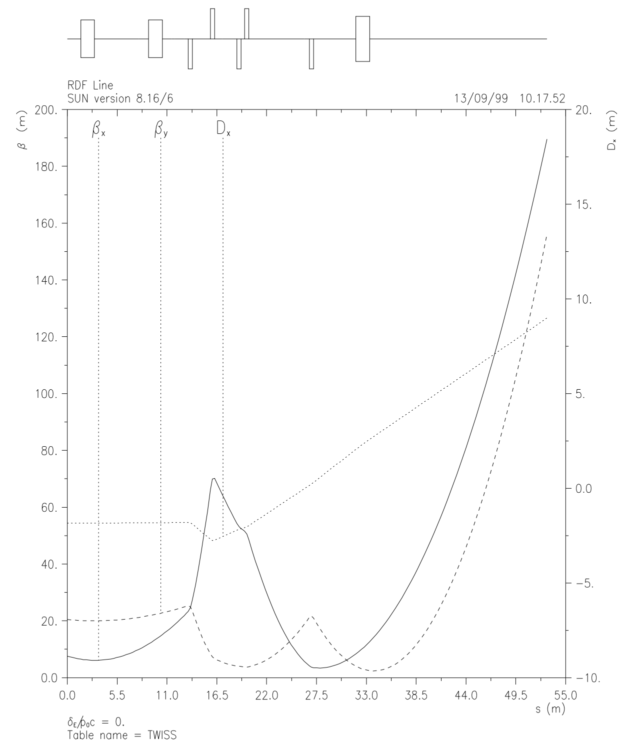

Figure A.3: Lattice functions for an extraction line from the Booster. The β functions and horizontal dispersion are plotted down the length of the line. The top schematic shows magnets along the line: bends span the line, F-quads extend above, and D-quads below. Note that the beam blows up toward the end of the line in the absence of focusing.

beam emittance and Courant-Snyder parameters[3] that are properties of the lattice: $\alpha(z)$, $\beta(z)$, and $\gamma(z)$.

Of the Courant-Snyder parameters β is the most often used as it defines the beam size in the machine:

$$\sigma_x = \langle x^2(z) \rangle^{1/2} = (\beta(z)\epsilon_x)^{1/2} \tag{A.3}$$

The parameter γ similarly gives the beam divergence and α describes the correlation between between the position and angle. The horizontal and vertical β functions of a Booster extraction line are shown in Fig. A.3. Note how the functions are bent at each quadrupole, and how a net focusing effect is

[3]Also known as Twiss parameters.

achieved.

The β function can also be used to describe single-particle dynamics along with another quantity: the phase advance, ψ.

$$x(z) = (\beta(z)\epsilon_0)^{1/2}\sin(\psi(z) - \psi_0) \tag{A.4}$$

where ϵ_0 as a single particle emittance. The phase advance is another quantity of the lattice and goes as the inverse of the β function:

$$\frac{d\psi}{dz} = \frac{1}{\beta(z)} \tag{A.5}$$

The phase advance is a measure of a particle's oscillation within the envelope defined by the β function. This last relation can be understood by considering $1/\beta$ to be a measure of focusing – greater focusing results in smaller beam (smaller β). Also, greater focusing increases the rate at which particles oscillate around the central orbit. Just as there are different β functions for horizontal and vertical motion, phase advances at different rates vertically and horizontally.

Dispersion

The above discussion considers the role of transverse phase space in determining beam size. However, momentum deviation also plays a role in transverse dynamics as each magnetic field acts as a spectrometer. Different momentum beams will have different central orbits throughout the machine, primarily from different bending angles in dipole fields; the off-momentum beam then is focused back toward the central orbit by the quadrupole fields[4].

[4]Generally, off-momentum beam will also receive a different amount of focusing. In machines as described so far, the focusing would be less with greater momentum. However, in actual machines the variation of focusing with momentum is controlled with sextupole fields. Regardless, this difference in focusing is a negligible, second-order effect for the purposes of this thesis.

In a uniform dipole magnetic field a difference in momentum, $\Delta p/p$, results in a proportional change in the radius of gyration, $\Delta r/r$. However, the focusing of quadrupole fields will try to focus back the beam toward the center. At each point of the lattice, there will be a dispersion function, $D(z)$, that describes the position variation of the beam with momentum along the beamline. The dispersion function is a property of the lattice – as well as the beam parameters in the case of transfer lines. For a single particle, or beam centroid:

$$\Delta x(z) = D_x(z)\frac{\Delta p}{p} \tag{A.6}$$

The spread in momentum of a beam will also contribute to the beam width through dispersion in quadrature with the beam width from emittance:

$$\sigma_x(z) = \left[\beta(z)\epsilon_x + \left(D_x(z)\frac{\delta p}{p}\right)^2\right]^{1/2} \tag{A.7}$$

The horizontal dispersion for the Booster extraction line is also shown in Fig. A.3. Usually horizontal dispersion will be more significant as that is the plane where most of the bending is done. However, in some cases like the NuMI line where there is significant vertical bending vertical dispersion must be considered as well (see Fig. 1.3).

A.2 Particle Acceleration

Charged particles are accelerated with electric fields. Electrostatic fields can be used for acceleration up to $\sim$ 10 million volts, but become impractical for higher energies. All high-energy particle accelerators use high-frequency alternating electric fields to accelerate particles. As the frequencies are usually in the radio or microwave range this is referred to as radiofrequency or RF acceleration.

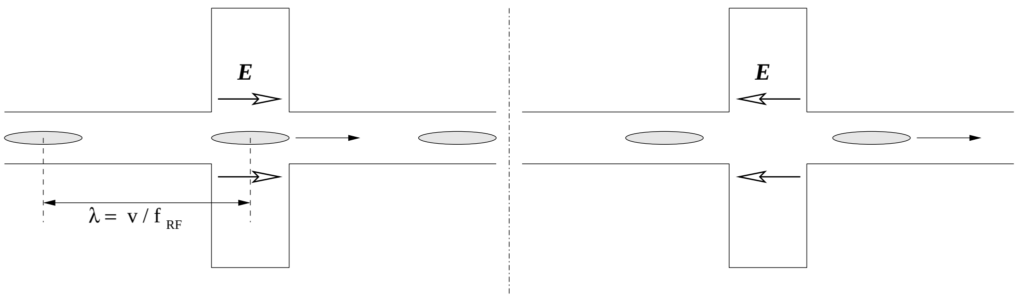

Figure A.4: Conceptual diagram of RF acceleration. A standing RF wave exists in a resonator through which the beam passes. The electric field of the RF wave is longitudinally oriented and varies sinusoidally with time. By bunching the beam and spacing it appropriately the beam can be in the cavity only with the positive electric field, and will thus be accelerated.

RF Acceleration

The principle behind RF acceleration is illustrated in Fig. A.4. RF energy is supplied to a resonator, or cavity, where it is stored and passed on to the beam. Holes are cut in the cavity to allow beam particles to enter and exit. The RF wavelength and cavity extent must be larger than the beam pipe so that the energy is confined to the resonator. As a beam particle passes through the cavity it acquires a kick corresponding to the potential:

$$\begin{aligned} V_{acc} &= -\int_{-d/2}^{d/2} dx\, E_0 \sin(\omega t - \Phi) &&= -\int_{-d/2}^{d/2} dx\, E_0 \sin(kx - \Phi) \\ &= E_0 \frac{2}{k} \sin\Phi \sin\left(\frac{kd}{2}\right) &&= E_0 d \sin\Phi \left[\frac{\sin(kd/2)}{kd/2}\right] \\ &= V_0 \sin\Phi \end{aligned} \tag{A.8}$$

Where we have defined V_0 as the product of maximum field, accelerating gap, and the transit time factor. The transit time factor is a measure of how much the field changes while the beam crosses the gap; it essentially decreases the maximum voltage. Note that the k in the transit time factor is the wave number for the beam and is thus related to the RF frequency as $k = 2\pi f/\beta c$.

For beam of a given momentum the transit time is fixed and the amount of acceleration is controlled by the amount of energy in the cavity and the phase, Φ, between the particle and the RF.

Further RF acceleration is achieved by passing the beam through a series of such cavities. This is accomplished by either a long string of cavities (a linear accelerator) or by bending the beam around and passing them through the same cavities more than once (e.g. the synchrotron – §A.3). To properly accelerate the beam the cavities must be spaced appropriately and have the appropriate RF phase and amplitude. In some cases feedback systems must be used to ensure proper acceleration.

RF Bunching

Above, we see how an individual particle can be accelerated by an RF cavity, but we must accelerate a distribution of particles. As the accelerating voltage varies with time, sometimes the cavity accelerates and sometimes the cavity decelerates. The key is to have beam only pass through the cavity while it is accelerating – this is accomplished by delivering the beam in bunches (as shown in Fig. A.4) instead of as a continuous beam.

A DC beam can be bunched by passing it through a single RF cavity followed by a drift region. The cavity will accelerate part of the beam, decelerate part, and leave the mean momentum unchanged. The drift region allows the faster beam to catch up with the rest, and the slower beam to fall back. Once the beam is bunched it can be accelerated together as described in the next section. A small portion of the beam is lost through this process, but as it usually happens at low energy it is not a significant concern.

In a linear accelerator it is often convenient to supply RF power of

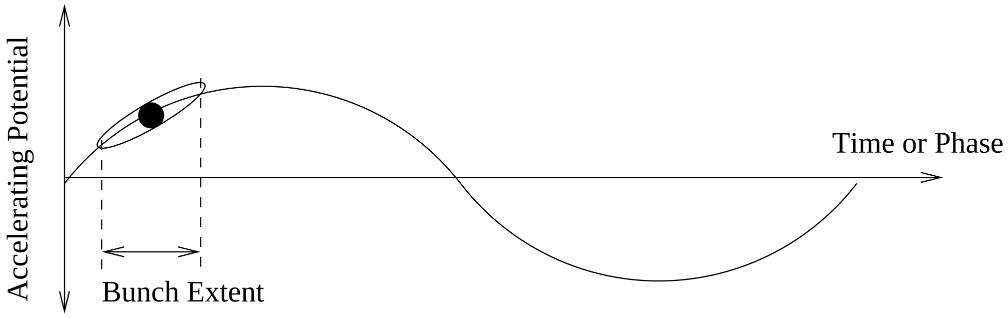

Figure A.5: Schematic diagram of RF focusing. The bunch of particles arrives at the cavity over some time. In a linear channel, the head particles get a smaller acceleration than the tail particles. Given a subsequent drift the tails particles catch up with the head particles.

the same frequency to different parts of the machine. As the beam speeds up, the distance between bunches and thus resonators must increase. This is usually done by modifying the phase velocity of the RF waves in an accelerating structure along its length. As will be discussed, synchrotrons keep the distance between bunches fixed, so the frequency of the RF must change.

Longitudinal Focusing

Even after being bunched, the particle beam will have a finite extent in time. Upon passing through a cavity different parts of the beam will receive different accelerations, as illustrated in Fig. A.5. One apparent solution might be to choose a phasing of $\Phi \approx 90°$ such that the beam lies on the top of the waveform and receives as uniform a kick as possible. However, this is bound to fail except for very narrow bunches[5] as the tail of the bunch will always get less kick and will fall inexorably behind.

The beam is accelerated off-peak (as shown in Fig. A.5) in order to

[5]Some electron machines which do have very narrow bunches do actually accelerate "on peak", but in proton machines the longitudinal emittance is always too great to do so.

maintain a stable bunch throughout acceleration. This acceleration producd type of strong focusing akin to that discussed above for transverse motion. By combining off-peak acceleration and drift periods the beam is locked into a focused beam – slow beam lagging behind gets a greater kick making it catch up with the rest of the beam until it overtakes the centroid and receives less of a kick afterward. This method can also can also be used to maintain longitudinal focusing while the beam is not accelerating, or even decelerating; the phase between the beam and RF need only be adjusted so that the average acceleration is zero.

A.3 Synchrotrons

The synchrotron is the marriage of a strong-focusing ring and RF acceleration. As the beam is accelerated the magnet currents are increased to keep the beam on the same orbit; the RF frequency is also adjusted to compensate for the quicker beam. At Fermilab and elsewhere a chain of synchrotrons is used to elevate a particle beam to ever higher energies.

Each synchrotron starts by having beam of protons injected to it, already with a relativistic velocity. The beam is focused transversely and longitudinally from the start, i.e. the magnets and RF are energized. The magnet currents are then increased, changing the stable momentum of the central orbit; concurrently the RF phase is adjusted to start accelerating the beam such that its rate of acceleration matches the rate of change of momentum required to stay on the central orbit. Often, a feedback system is used to maintain this synchronicity.

This section describes some of the relevant features of synchrotrons for this thesis. Of concern will be the basic beam dynamics within the machine;

the application and effects of the RF acceleration; and the interaction of longitudinal focusing with a strong focusing lattice and relativistic effects.

Betatron Oscillations and Tunes

We have discussed strong-focusing above with reference to it being used in linear channels and rings – there are a few specifics of operating in a ring that deserve bearing out. First, we consider the issue of Betatron oscillations. Above, we could use the β function to define the size of a beam or the oscillation of a particle about the center of the beam while traveling down a line. In the case of a transfer line the β function is defined by the lattice and the properties of the beam (i.e. injected distribution). If the next beam is injected off-center, then it has different β functions down the line associated with its new central orbit.

The case of magnet rings is different as multiple revolutions requires that a beam follow a repeating orbit. This defines a single closed orbit for the beam[6] about which the beam centroid will oscillate. These oscillations are called betatron oscillations and have specific frequencies.

Consider again the case of a particle beam gyrating in a uniform dipole magnetic field. The frequency of the gyration is known as the cyclotron frequency[7]. This is also the frequency that a particle will oscillate around a central orbit in the field. To an observer at a particular location, however, the particle might not appear to be oscillating as the particle will always return to the same position at a point in revolution. This aliasing where the oscillation appears stationary is only an issue for the cases where the frequencies of

[6]Different momentum beams and beam particles will have different closed orbits. These orbits are discussed throughout the next several sections.

[7]The early accelerators lent their name of cyclotron to this motion.

oscillation and sampling are related by an integer. Also note, that the vertical motion has no periodic behavior, so its frequency is essentially zero.

In the case of a focusing ring – whether weak or strong – there will both horizontal and vertical oscillations; also, they will generally be different from the revolution frequency. In the case of weak focusing a return force is pressed upon the vertical motion at the expense of a horizontal return force. There will be frequencies for each oscillations, but they generally will be unequal, and both will be less than the revolution frequency. That is, it will take more than one revolution to execute a full oscillation. These oscillations were first measured in betatron accelerators, and are thus called betatron oscillations. The frequencies of oscillation can be expressed in terms of the revolution frequency:

$$\omega_x = \nu_x \omega_c \quad \omega_y = \nu_y \omega_c \tag{A.9}$$

There ν_x and ν_y are called the betatron tunes and are related to each other in a weak machine by:

$$\nu_x^2 + \nu_y^2 = 1 \tag{A.10}$$

Strong focusing dramatically increases the return force, and thus has higher oscillations frequencies – generally greater than the revolution frequency. To evaluate the betatron frequencies we go back to the concept of phase advance as defined in Eqn. (A.5). The phase of the oscillation advances more quickly with greater focusing. We can calculate the phase advance of one revolution:

$$\begin{aligned} \Psi_i &= \oint_C dz \, \frac{d\psi_i}{dz}(z) \\ &= \oint_C \frac{dz}{\beta_i(z)} \end{aligned} \tag{A.11}$$

Where i can be x or y. The tune is thus related back to the β function as $\nu_i = \Psi_i / 2\pi$.

Measurement of the tune is complicated by the aliasing mentioned above. On successive revolutions the particle or beam position at a particular location will be different (unlike the cyclotron). A collection of such measurements can be used to measure the tune via frequency analysis. However, with strong focusing the length of oscillation will generally be less than the circumference – several oscillations will occur over the course of a revolution. As such, only the fractional part of the tune is measured. This part is more important anyhow as it described the frequency of re-occurrence of position and angle at a location. To prevent resonant conditions that can damage the beam, rational numbers are avoided as tunes and as ratios between the horizontal and vertical tunes. In a synchrotron all of the bending and focusing magnets are ramped simultaneously so that the focusing remains constant during the cycle; this keeps the tunes constant and avoids crossing rational resonances.

Momentum Compaction

Another effect of the strong focusing lattice is to reduce the beam spread due to momentum spread, i.e. dispersion. If we consider again the cyclotron case of a uniform dipole magnetic field, we expect the radius of gyration to change proportionally to any change in momentum according to the magnetic rigidity – Eqn. (A.1). Again, this radius is only in the plane perpendicular to the magnetic field – any momentum parallel to the magnetic field will cause the beam to diverge indefinitely. In the weak focusing case the beam will be confined vertically, but again at the cost of an enlarged beam horizontally. Similarly, we expect a strong-focusing synchrotron to compress the beam in both planes – leading to the term "momentum compaction". We will define momentum compaction, α – not the Courant-Snyder parameter, for a ring to be the double ratio of circumference (or radius of gyration) change to momentum change:

$$\alpha = \frac{\Delta C/C}{\Delta p/p} \tag{A.12}$$

Apparently for a cyclotron α=1, for weak focusing $\alpha > 1$, and for strong focusing $\alpha < 1$.

Dispersion, as defined in Eqn. (A.6), will allow us to evaluate momentum compaction for a synchrotron. Dispersion is primarily horizontal in a ring, so the vertical portion can be neglected. To first order the circumference can be evaluated as:

$$C = \oint_C dz \left[1 + \frac{\Delta x(z)}{\rho(z)}\right] \tag{A.13}$$

Where $\rho(z)$ is the local bending radius of the dipoles. Substituting, the change of circumference can be calculated as:

$$\begin{aligned} \Delta C &= \frac{\Delta p}{p} \oint_C dz \frac{D_x(z)}{\rho(z)} \\ &\approx \frac{\Delta p}{p} 2\pi \langle D_x \rangle \end{aligned} \tag{A.14}$$

The approximation is the case where the dispersion does not vary much over the ring. This is close to true for round machines like the Booster, but can be far from the truth for other, less round machines[8]. We can then calculate the momentum compaction factor:

$$\begin{aligned} \alpha &= \frac{1}{C} \oint_C dz \frac{D_x(z)}{\rho(z)} \\ &\approx 2\pi \frac{\langle D_x \rangle}{C} \end{aligned} \tag{A.15}$$

Momentum compaction will be used in some of the discussion below as well as in subsequent chapters to consider the effect of off-momentum beam.

[8]In fact, a ring can be designed to have zero or even negative change in circumference with momentum change. To do so requires more focusing, but can have significant advantages in terms of avoiding isochronous acceleration [164].

Synchronous Phase

As discussed above: the acceleration from an RF cavity depends on both the maximum accelerating voltage and the phase with respect to the beam, as defined in Eqn. (A.8). A synchrotron will have a number of RF cavities around it circumference, N, such that for each revolution (turn) the beam will gain the energy:

$$\Delta E = NV_{acc} = NV_0 \sin(\Phi) \tag{A.16}$$

Amounting to an average acceleration rate:

$$\begin{aligned} \frac{dE}{dt} &= \frac{\Delta E}{\tau} \\ &= \frac{\beta c N V_0 \sin(\Phi)}{C} \end{aligned} \tag{A.17}$$

By adjusting V_0 and Φ the acceleration can be set to a wide range of values – even for negative acceleration.

The acceleration rate is set by the ramping of the magnets in the synchrotron. The magnets currents are increased in unison such that the momentum of the closed orbit increases. This involves both bending and focusing magnets as we will usually want to keep the tunes (focusing) constant. The rate and timing of the ramp are arbitrary as long as the RF can keep pace. The RF voltage and phase are modulated to provide enough power to the beam to keep it on its central orbit. Generally, the RF voltage can only be adjusted slowly and is kept constant over short periods of time. The RF phase can be varied more quickly and is used to fine tune the beam momentum.

The synchronous phase, Φ_s, is the phase that must be maintained between the applied RF beam charge in order to keep the beam centered in the magnets while the currents are increased[9]. In fast-ramping synchrotrons the

[9]Conventions vary, but usually $0°$ corresponds to zero acceleration.

phase can be quite large, $> 60°$, and the synchronization will not be perfect, possibly to the point that the beam will fall outside of the available aperture. To ensure proper acceleration, all fast-ramping machines include feedback systems that further modulate the RF phase to keep the beam on the desired orbit. The system that applies the synchronous phase and feedback is known as Low-Level Radio Frequency (LLRF).

We should also note that synchrotrons must change the RF frequency, sometimes substantially, during the ramp of the machine. In linear accelerators the frequency is kept constant while the bunch spacing increased with velocities. In synchrotrons, however, a fixed number of bunches can be fit around the circumference of the ring. The bunch spacing must remain fixed, so the frequency must be changed. As beam accelerated the frequency is increased asymptotically toward the limit of a light-speed traveling particle:

$$f_{RF} \underset{\beta \to 1}{\longrightarrow} \frac{hc}{C} \tag{A.18}$$

Where h is the harmonic number: the number of RF buckets around the circumference of the machine.

Synchrotron Oscillations

While accelerating, a beam particle might have different momentum or longitudinal position from the center of the beam. This particle can still be accelerated as long as it falls into an area of longitudinal phase space known as the RF bucket; the particle will then oscillate about the central part of the beam in a process known as a synchrotron oscillation, like in the longitudinal focusing illustrated in Fig. A.5. A particle at the center of the bucket is known as the synchronous particle, as its phase with respect to the RF is always the synchronous phase.

The extent of the RF bucket is defined by the RF voltage available and and the synchronous phase necessary to keep up with the beam. Lesser synchronous phase allows a greater range of longitudinal positions (or particle phases) to be accelerated. Greater RF voltage allows a greater range of particle momentum to be accelerated. For a given magnet ramp, increasing the RF voltage reduces the synchronous phase required, so the RF bucket's phase extent is larger as well.

The synchrotron frequency, ω_s, is the rate at which particles oscillate about the synchronous particle in phase space. ω_s depends on a number of factors including RF voltage, synchronous phase, and momentum compaction. The synchrotron frequency is typically less than the revolution frequency, in contrast to the betatron frequency which is typically greater.

The betatron oscillations described above are linear as long as the magnetic fields are pure dipole and quadrupole, however the synchrotron oscillations are necessarily nonlinear due to the sinusoidal nature of the accelerating voltage. In principle, if the RF voltage were very large and the beam was very small in proportion to the RF wavelength, then the behavior of the particles could be somewhat linear. The reality for proton machines is that the beam size is a significant fraction of the RF wavelength[10]. Any perturbations in the original beam will oscillate at different rates about the synchronous particle – eventually filling a larger phase space volume, increasing the longitudinal emittance of the beam.

[10] In synchrotrons where the beam is captured into RF buckets from a laminar beam, the beam initially takes up the entire RF waveform. The Booster is a machine of this type.

Adiabatic Damping

As the beam accelerates its transverse phase space is compressed due to a process known as adiabatic damping; as a result, the beam emittance is reduced, leading to smaller beams size. This effect arises as the RF acceleration is purely in the longitudinal direction. As the total momentum of the particle is increased the transverse momentum remains constant, such that the angular divergence of the particles is decreased. The beam emittance is thus inversely related to the momentum of the beam:

$$\epsilon_i \propto \frac{1}{\beta\gamma} \tag{A.19}$$

The beam size will decrease as well. At any particular position around the ring the beam size will scale inversely with the root of momentum[11]:

$$\sigma_i \propto (\beta\gamma)^{-1/2} \tag{A.20}$$

In Fig. 3.9 we showed beam width measurements during the Booster cycle, which shrink as the beam energy grows.

Transition

The combination of momentum compaction and relativistic effects produce an effect known as transition. When in transition the beam is isochronous, in that the revolution time of a particle is independent of its momentum. Transition has several effects in a synchrotron, most significantly affecting longitudinal focusing.

Consider again the case of the cyclotron where a uniform dipole magnetic field turns particles in circular orbits. The momentum compaction, α,

[11] Occasionally, the beam size will be dominated by momentum spread and dispersion, so the scaling will not strictly follow this rule.

for this type of field is 1, in that any change in momentum produces a proportional change in circumference. The revolution time for a non-relativistic beam will be independent of the momentum:

$$\tau_C = \frac{C}{v} = \frac{C_0(1+\Delta p/p_0)}{\frac{p_0}{m}(1+\Delta p/p_0)} = \frac{C_0}{p_0/m} \tag{A.21}$$

$$\frac{d\tau_C}{dp} = 0 \tag{A.22}$$

Where the default circumference, C_0, is determined by the default momentum, p_0, and the magnetic rigidity – Eqn. (A.1). Cyclotron motion is thus isochronous at non-relativistic momenta as all particles will gyrate at the same frequency in the magnetic field. Relativistic effects will cause deviations as the beam will not be able to go arbitrarily fast as the circumference increases.

$$\tau_R = \frac{C}{p/\gamma m} = \gamma\tau_C \tag{A.23}$$

$$\frac{d\tau_R}{dp} = \tau_C\frac{d\gamma}{dp} = \frac{\beta^2\gamma}{p}\tau_C \tag{A.24}$$

In this case of a cyclotron, the beam is isochronous at low momentum, but the rotation frequency decreases as the particles become relativistic.

For most synchrotrons the momentum compaction is significantly less than $\alpha = 1$. In these cases the revolution time for a non-relativistic particle decreases with the momentum; the circumference increases less quickly than the momentum and velocity. As the beam accelerates and becomes relativistic, the speed of the particles increases more slowly than the momentum until eventually it takes higher energy particles longer to make a revolution. The rate of change in the revolution period:

$$\begin{aligned}\frac{d\tau_S}{dp} &= \frac{\tau_S}{C}\frac{dC}{dp} - \frac{\tau_S}{\beta}\frac{d\beta}{dp} \\ &= \frac{\tau_s}{p}\left(\alpha - \frac{1}{\gamma^2}\right)\end{aligned} \tag{A.25}$$

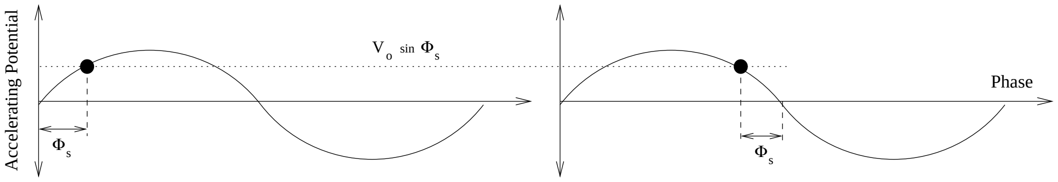

Figure A.6: RF acceleration before and after transition. Before transition higher momentum beam arrives earlier at a particular location, so the accelerating voltage should be increasing with time, as on the left. After transition lower momentum beam arrives earlier, so the accelerating voltage should be decreasing, as on the right.

For small γ, $d\tau_S/dp$ is negative as $\alpha << 1$. There is a specific beam energy where the beam becomes isochronous[12]. The isochronous period is usually very short in a synchrotron as the momentum is ramping; the event in time is called transition and occurs at a boost defined solely by the momentum compaction of the lattice:

$$\gamma_{\text{tr}} = (\alpha)^{-1/2} \tag{A.26}$$

γ_{tr} is known as the transition gamma, the corresponding energy being known as the transition energy.

Transition's most significant effect is reversing the direction of longitudinal focusing. As discussed earlier, longitudinal focusing works by providing a return force to the center of the batch in the form of an energy kick. In linear accelerators and in synchrotrons below transition higher-momentum particles will reach a cavity first and should receive a lesser acceleration, as shown in the left plot of Fig. A.6. After transition, however, slower particles revolve more quickly and will reach the cavity before the more energetic ones, thus the slope of the kick must be reversed, such that the synchronous phase lies on an entirely different quadrant of the RF waveform, as shown in the left plot

[12]The beam is only isochronous to first order. There is still revolution time deviation with momentum, but it is second-order.

of Fig. A.6. The angle between the applied RF and the beam must change by $\pi - 2\Phi_s$. To make this change the machine's LLRF must be aware of the time of transition and be able to rapidly shift the phase of the RF – the jump will also be limited by the bandwidth of the RF driver and cavity.

A secondary, yet significant, effect of transition is the reduction of longitudinal focusing near transition. Longitudinal focusing has two components: the near-linear variation of accelerating voltage with phase, and the variation of revolution time with particle momentum. Near transition the beam is isochronous and the focusing breaks down. The beam stays on the same part of the RF waveform for successive terms. Before the phase jump the particles at the head of the beam will receive significantly less acceleration than the tail of the beam. The phase jump must be timed precisely to undo this momentum stretching on the opposite slope, before focusing takes over and leads to emittance dilution. During the above process, the momentum spread of the beam becomes significantly larger and the bunch length smaller – the beam thus becomes wider due to dispersion (see Fig. 3.9 for this effect in the Booster). For the above reasons, transition and the regime close to it are avoided as much as possible; for example: the momentum can be ramped more quickly through transition, the central orbit can be moved to take advantage of momentum compaction changes, or the lattice can be quickly changed to affect the momentum compaction.

Appendix B

Proton Beam Positioning

The proton beam position at the target is calculated from position measurements at two upstream instrumentation stations. Each station measures the horizontal and vertical position of the beam with capacitive beam position monitors (BPMs) and segmented foil secondary emission profile monitors (PMs). These two types of detectors give redundant measurements that each have their own advantages. Corrector dipoles upstream of the instrumentation permit fine-tuning of the beam's trajectory.

The standard for the NuMI beamline is that the z-coordinate is measured along the beam's path. For convenience then, the y-coordinate is up and the x-coordinate is to the right (when looking downstream). This results in a left-handed coordinate system.

B.1 Proton Beam Preparation

The proton beam is transported from the Main Injector to the target by the primary beamline (see §1.2). A series of dipole corrector magnets fine-tune

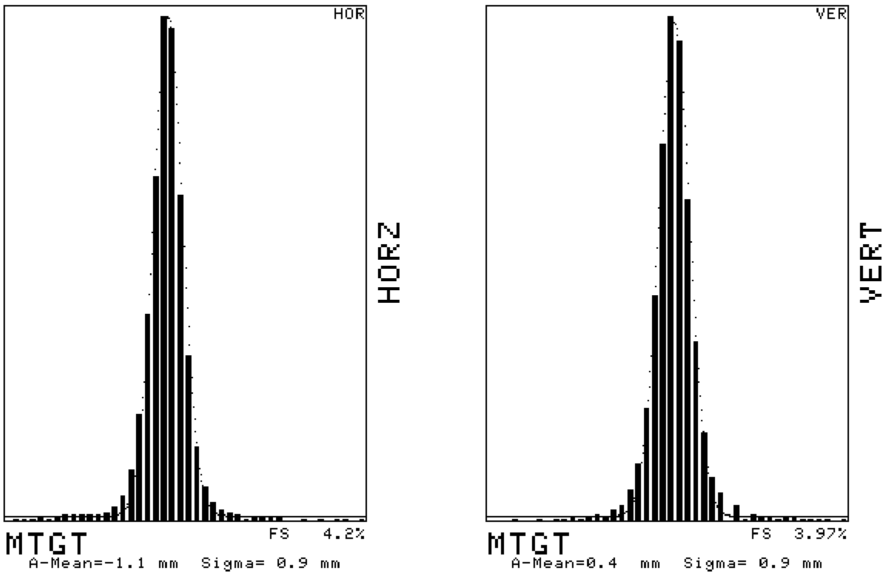

Figure B.1: Typical proton beam profiles as measured in the target SEM profile monitor. This display is from the online system where only rudimentary pedestal subtraction has been performed; both horizontal and vertical profiles are shown. The offline calculations of centroid and width have another digit of precision.

the position of the beam to better than 1 mm average along the beamline, and 0.25 mm at the target.

The nominal beam spot is supposed to be round and to have a waist in focusing at the target; the nominal RMS width is 1.0 mm, but varies with beam intensity and conditions. A typical profile as measured by the target profile monitor is shown in Fig. B.1. During the early commissioning of the beamline, one quadrupole magnet was operating with the incorrect current, changing the beam spot at the target. In this period the beam was measured as being very narrow and converging horizontally while approaching the target, and wider and diverging vertically. The horizontal beam size probably waisted horizontally between the two horns. The result of this is that the beam had an aspect ratio of 2:1 during these scans, being narrower in the horizontal. The estimated typical beamspot at the target was 0.7 x 1.4 mm^2.

To scan the beam laterally across the target, a pair trim magnets are varied upstream to move the beam parallelly across the target; angle manipulations are also possible, but not as useful. The precise ratio to change the dipole currents were determined using the TRANSPORT code [165], taking into account the intervening main dipole and quadrupole magnets. These ratios are not entirely robust: if the position of the beam in a quadrupole magnet changes it will acquire a different bend from that magnet, sometimes inducing an angle in the beam, sometimes causing motion in the opposite direction.

B.2 Proton Beam Position Measurement and Extrapolation

The proton beam position is measured independently by BPMs and PMs at the two instrumentation stations, shown in Fig. B.2, upstream of the target pile. These two systems (BPMs and PMs) have different absolute offsets from the beam center, defining two different coordinate systems for the beam.

The Profile Monitors (PMs) are segmented secondary-emission monitors (SEMs) [25]. The PMs detect the passage of charged particles by the expulsion of electrons from Titanium foils through the process of secondary emission. The foils are segmented and individually instrumented, allowing the beam profile to be recorded. From that profile we extract the beam central position and width in the form of a fitted gaussian sigma. The expected precision of the PMs at these stations was < 25 μm [24], and measured in the beam to be 18 μm [25]. The PM foils were optically referenced to exterior monuments used to survey the PM in place in the beamline.

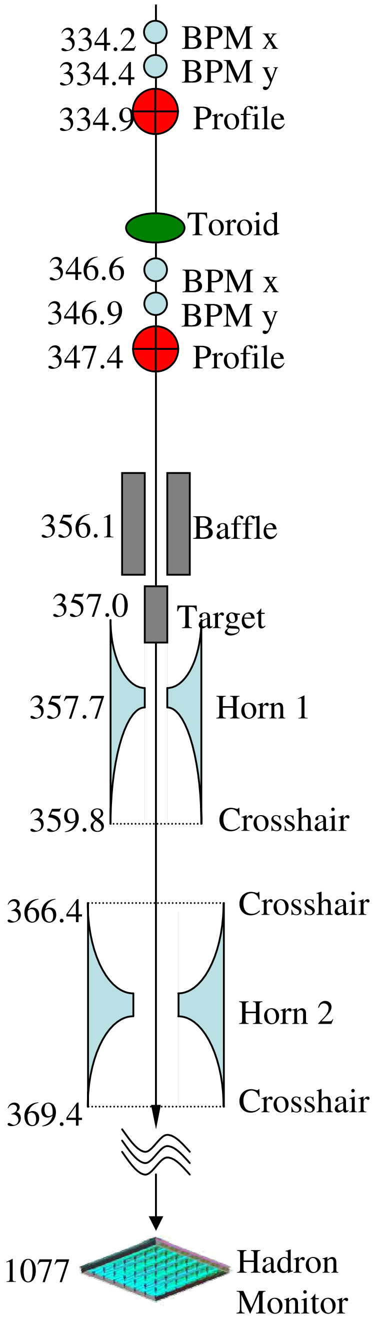

Figure B.2: Schematic layout of the target hall components and instrumentation used in this section. The numbers give the position along the beamline of the device in meters. Two sets of BPMs and profile monitors give the proton beam and angle used to extrapolate into the target hall. The target hall contains the target, baffle, and horns with their crosshairs and loss monitors. The Hadron Monitor is located far downstream, after the decay pipe.

The Beam Position Monitors (BPMs) are electrically isolated plates on the walls of the beam pipe vacuum chamber that capacitively couple to the charge of the proton beam. Two plates, of specifically chosen geometry, are on opposite sides of the beam pipe; the beam will induce more image charge on one side than the other. Comparison of the signal on the two electrodes gives beam position information; the difference in signals divided by the sum gives the position to first-order. The target BPMs are expected to have resolutions of 50 μm at high intensities, but at lower intensity the resolution is more like 150 μm over a range of $\pm$ 6 mm ([162], §4.1.3.3). The BPMs' electrical centers are measured on the bench and referenced to external monuments. There is an additional electrical offset that is dependent upon the cabling details and must be measured *in situ*. This offset changes the signal measured on each plate by a constant amount (like a pedestal); the effect of this offset is inversely proportional to intensity, so is more important for some of the target alignment scans of Ch. 11. The BPMs were optically surveyed in place (like the PMs), and these various offsets summed.

The correlation between the last two BPMs immediately upstream of the target is shown in Fig. B.3. The data shown are from a period of steady running when the beam was not adjusted. The spread in positions is the natural variation during running. The angle has a similar variation pulse-to-pulse; as these BPMs are separated by 12 m, the angle causes some spread in the correlation. The spread is shown in the residuals, but is convolved with the precision of the instruments. The relatively small residual demonstrates that the angle does not change appreciably pulse-to-pulse.

The correlation between the last BPM and PM is shown in Fig. B.4. Here, the two instruments are at nearly the same location and the measurements are not be affected by the beam angle.

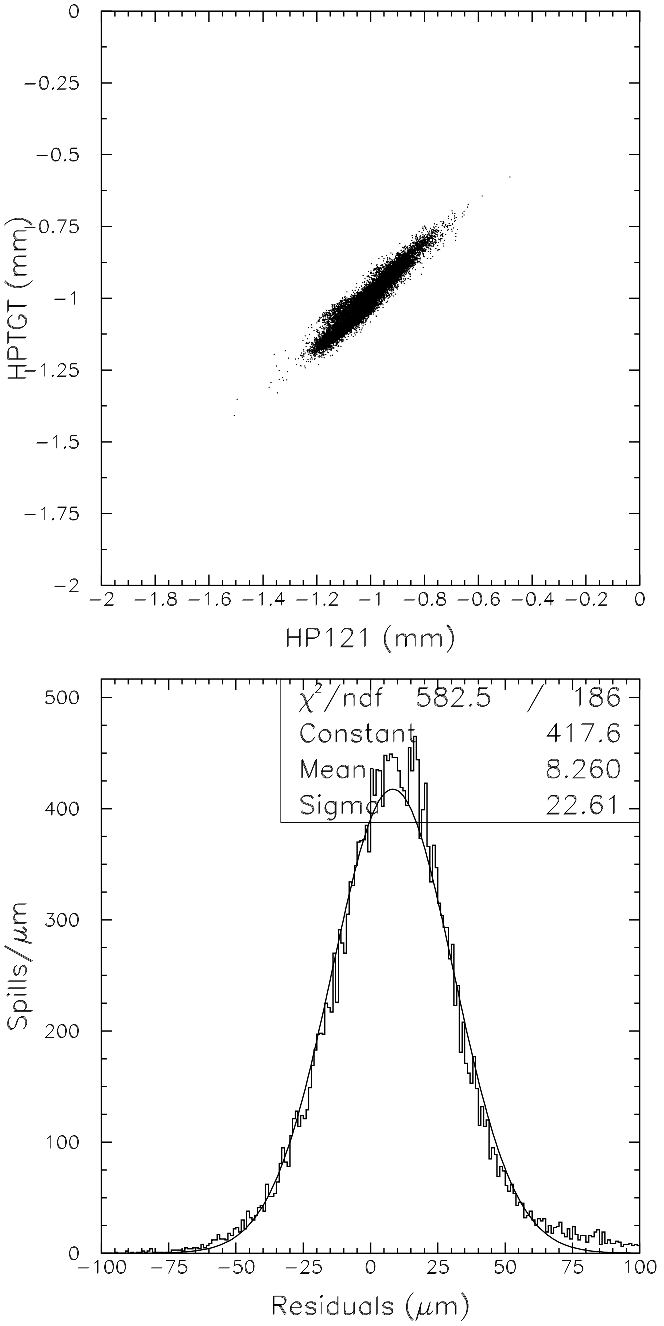

Figure B.3: Scatter plot and residuals of the measurements of horizontal position at the two BPMs upstream of the target. This running is taken from stable conditions, so the angle of the proton beams through the BPMs is only caused by variability in the beam. The residual width is thus a convolution of measurement precisions and the position difference caused by the variable angle.

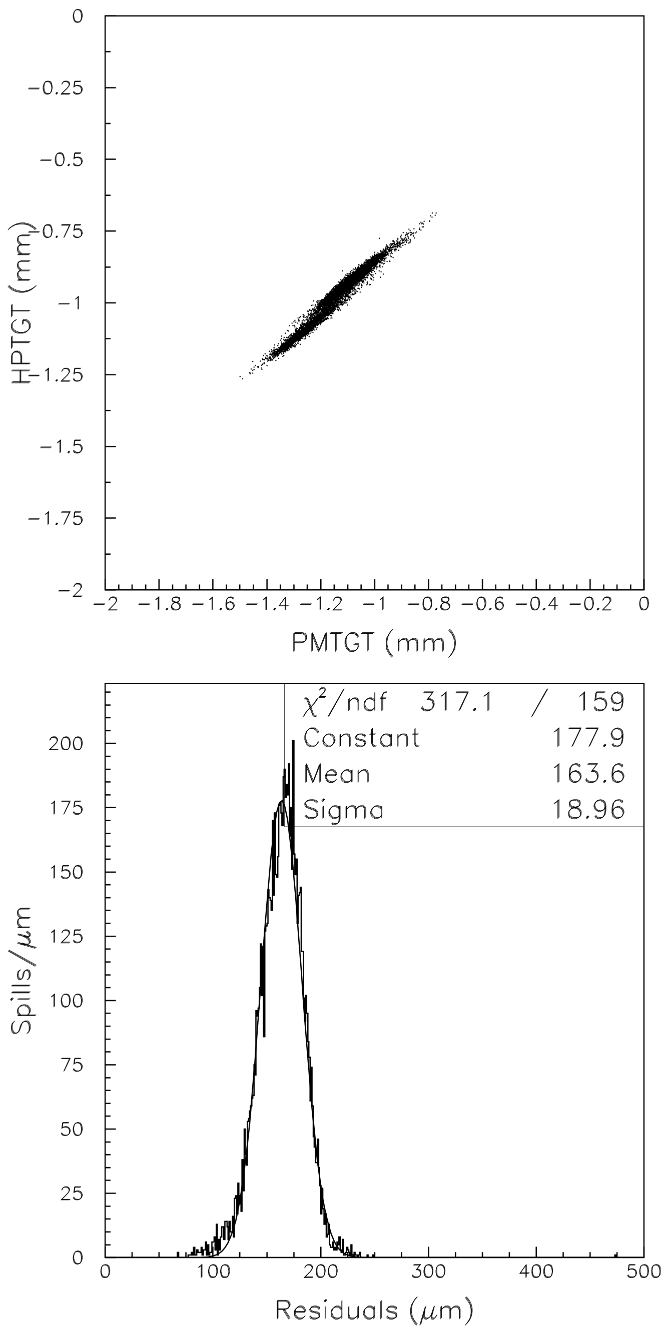

Figure B.4: Scatter plot and residuals of the measurements of horizontal position at the BPM and PM just upstream of the target. This running is taken from stable conditions, so the difference between the measurements is caused only by their calibrations and precision. The centroid of the residual is the offset of the calibrations, and the width is the convolution of the devices' precisions.

B.3 Proton Beam Extrapolation to the Hadron Monitor

For the alignment scans in §10 we use the BPMs and profile monitors to extrapolate proton beam position to locations within the target hall. A cross-check of this procedure is to directly transport the proton beam through the target hall and decay volume to the Hadron Monitor. A special run was performed when the target was removed where the beam was angled such that the position at the neck of Horn 1 remained constant. The range of angles available allowed the proton beam to be scanned horizontally across most of the face of the Hadron Monitor. The results of this scan are shown in Fig. B.5.

The BPMs were used in this analysis as the large transverse positions at the instrumentation stations included positions where dead foils on the PMs interfered with the position measurement. Overall, the results are the same for both: the proton beam can be well extrapolated to the center of the Hadron Monitor, but there is an unresolved scaling that results in poor correspondence at large angles. The offset of 41 mm at the Hadron Monitor is not large in terms of angle ($\sim 50\ \mu$rad), and is also consistent with the measured offsets of the BPM coordinate system as compared to the PMs – the intercept is closer to zero in the PM coordinate system.

The slope of 89-90% exists in both the BPM and PM extrapolations as compared to the measured Hadron Monitor position. This number has been cross-checked several ways: counting strips in the SEMs to verify position measurements; counting cells in the Hadron Monitor to check its position calculation. The source of this variation remains unknown.

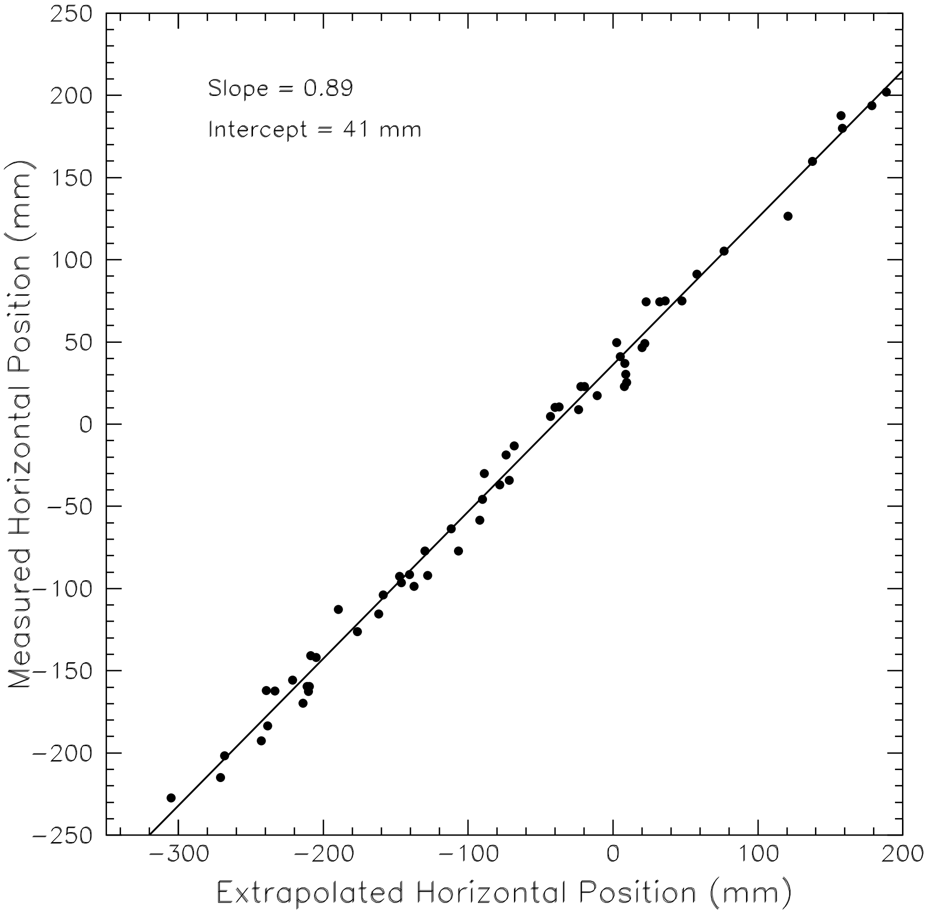

Figure B.5: Measured position at the Hadron Monitor vs the positions extrapolated from the BPMs upstream of the target. The deviation of the slope from unity is not understood. The intercept is consistent with our understanding of BPM and Hadron Monitor positions.

Appendix C

Analytical Approaches to Pulse Formation

As discussed in Ch. 8, the flow of charge in an ionization chamber is described by a set of coupled partial differential equations:

$$\frac{d\rho_+}{dt} = \Phi - v_+\frac{\partial\rho_+}{\partial x} - \rho_+\frac{\partial v_+}{\partial x} + \alpha\rho_- - r_e\rho_-\rho_+ + D\frac{\partial^2\rho_+}{\partial x^2} \tag{C.1a}$$

$$\frac{d\rho_-}{dt} = \Phi - v_-\frac{\partial\rho_-}{\partial x} - \rho_-\frac{\partial v_-}{\partial x} + \alpha\rho_- - r_e\rho_-\rho_+ + D\frac{\partial^2\rho_-}{\partial x^2} \tag{C.1b}$$

$$\frac{\partial E}{\partial x} = \frac{e}{\epsilon_0}(\rho_+ - \rho_-) \tag{C.1c}$$

$$\int_0^d dx E = V_{app} \tag{C.1d}$$

While there is no general solution for the above system, the below sections provide solutions and calculations for certain limiting cases. These cases are useful in understanding the results of the numerical simulations presented in Ch. 8, and provide a crosscheck on the simulation details.

C.1 First-Order Recombination Losses

We first consider the case where ionization in the chamber is insufficient to affect the applied electric field, but the drift velocity is low enough to allow recombination. This is the limit of low applied field and steady-state, low-rate beam ionization, reducing Eqn. (8.1) to:

$$0 = \frac{d\rho_+}{dt} = \Phi - v_+ \frac{\partial \rho_+}{\partial x} \tag{C.2a}$$

$$0 = \frac{d\rho_-}{dt} = \Phi - v_- \frac{\partial \rho_-}{\partial x} \tag{C.2b}$$

$$\frac{\partial E}{\partial x} = 0 \tag{C.2c}$$

$$\int_0^d dx E = V_{app} \tag{C.2d}$$

Here Φ, v_+, and v_- are constants. We apply the additional boundary conditions that $\rho_+(x=0)=0$ and $\rho_-(x=d)=0$ to find the solutions:

$$E = V_{app}/d \tag{C.3a}$$

$$\rho_+ = \frac{\Phi}{\mu E} x \tag{C.3b}$$

$$\rho_- = \frac{\Phi}{v_-(E)}(d-x) \tag{C.3c}$$

These densities are plotted in Fig. C.1; clearly the electron densities are much less than the ion densities.

We can now define a recombination density:

$$\begin{aligned} \mathcal{R}(x) &= r_e \rho_-(x)\rho_+(x) \\ &= \frac{r_e \Phi^2}{\mu E v_-(E)}(d-x)x \end{aligned} \tag{C.4}$$

$\mathcal{R}$ is the rate of charge loss per volume per time. We can then define the collection efficiency as the ratio of charge collected on the electrodes to that

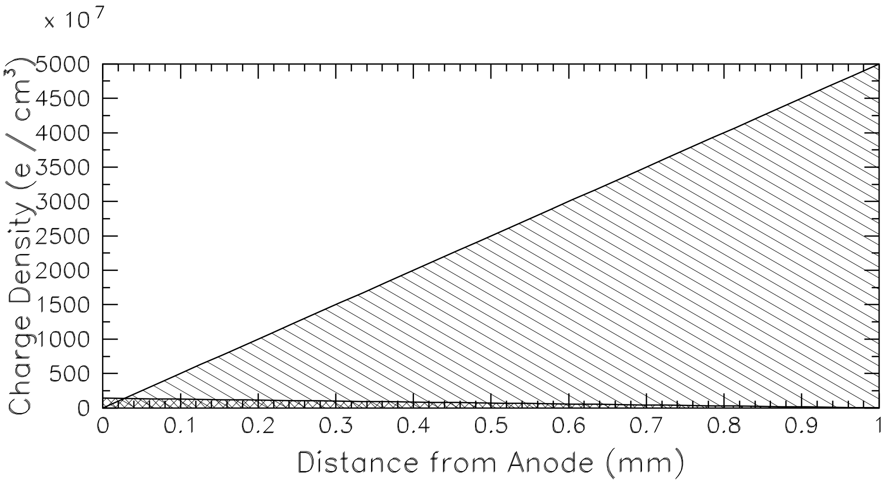

Figure C.1: Calculated charge distributions using the method of §C.1. This is for a 1 mm chamber, biased at 100 V with steady-state ionization rate of 10^{10} ionizations / cm^3 / μs. The upper line is for the ion density, the lower for electrons.

originally ionized by the beam:

$$\begin{aligned}\epsilon &\equiv \frac{1}{\Phi d}\left[\Phi d - \int_0^d dx\mathcal{R}\right] \\ &= 1 - \frac{r_e \Phi d^2}{6\mu E v_-(E)}\end{aligned} \tag{C.5}$$

The efficiency is plotted as functions of voltage and intensity in Fig. C.2 for a 1 mm chamber[1]. Note that at a typical applied voltage of 1000 V/cm there is no appreciable loss of efficiency until an ionization rate of 10^{19} ionizations / cm^3 / μs; this is seven orders of magnitude greater than the maximum in the NuMI beam or beam tests.

However, the validity of Eqn. (C.2) fails at these high ionization rates.

[1]The value of effective recombination coefficient r_e used is 2.4×10^{-8}.

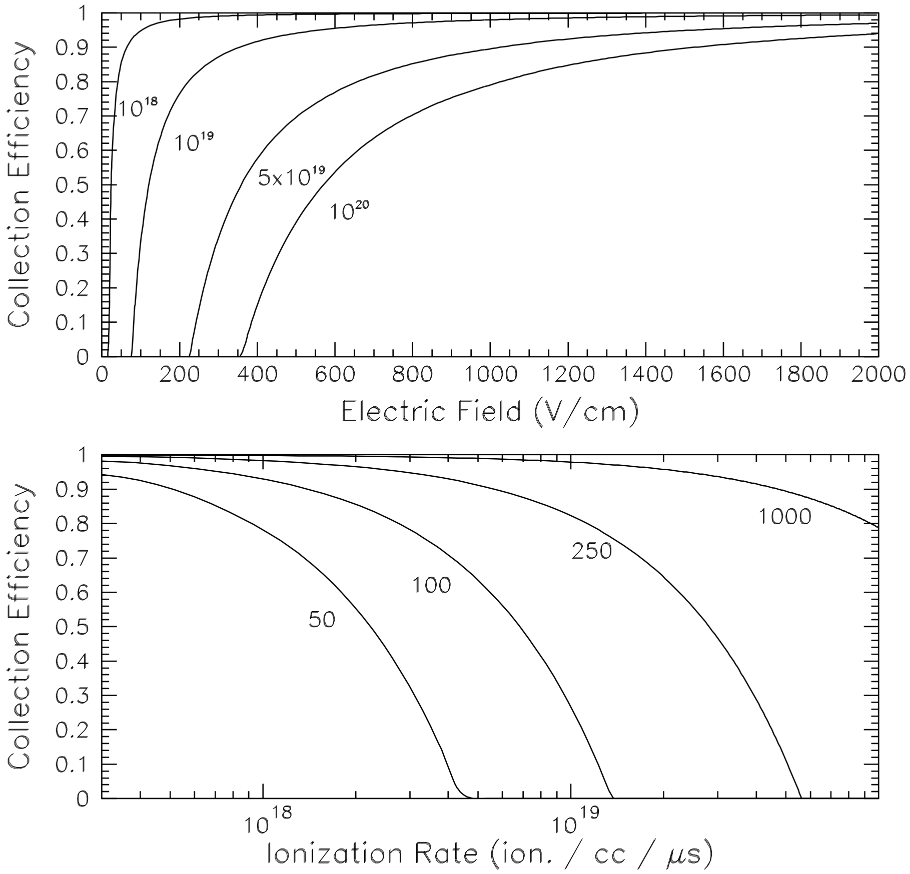

Figure C.2: Calculated collection efficiencies using the method of §C.1. These plots are for a 1 mm chamber. The upper plot is versus the applied electric field; the curves are for different (very high) ionization rates indicated by the numbers in units of ionizations / cm^3 / μs. The lower plot is versus the ionization rate, with the different lines corresponding to different applied electric fields in units of V/cm.

This can be seen by considering the effect of ion density on the electric field:

$$\begin{aligned}\frac{\partial E}{\partial x} &= \frac{e\rho_+}{\epsilon_0} \\ &= \frac{e\Phi d}{\mu V_{app}\epsilon_0}x\end{aligned} \tag{C.6}$$

Applying the boundary condition Eqn. (8.1d), we get the electric field solution:

$$\begin{aligned}E &= E_0 + \frac{e\Phi\mu V_{app}d^2}{6\epsilon_0}x \\ E_0 &= \frac{V_{app}}{d} - \frac{e\Phi d^3}{6\mu V_{app}\epsilon_o}\end{aligned} \tag{C.7}$$

If $E_0 < 0$ then at some point $E < 0$, which is unphysical as it would cause infinite accumulation of charge at some point [2]. The ionization rate at which $E_0 = 0$ is:

$$\Phi = \frac{V_{app}^2}{d^4}\frac{6\mu\epsilon_0}{e} \tag{C.8}$$

For a typical 1 mm chamber biased at 100 V this rate is 6.6×10^9 ionizations / cm^3 / μs; far below the rates quoted above, and even below the ionization rate in the Hadron Monitor. The Muon Monitors, however, see a lower rate, so this approximation is of some use there.

C.2 Steady-State Solution with Space Charge

As the previous approximation is not valid for the intensities of the Hadron Monitor, we expand the approach to include the space charge of the ions. This approach still does not include multiplication and considers the charge loss to recombination to be a correction. Additionally, the electron contribution to space charge is ignored as it is much smaller than the ion contribution[3].

[2]This situation corresponds to total space-charge screening of the anode

[3]The smaller electron density is caused by its much faster velocity. In some gases this factor can be more than a thousand; in Helium, however, the electron velocity is only about fifty times greater than the ion velocity

We start by considering the simplified form of Eqn. (8.1):

$$\frac{d\rho_+}{dt} = \Phi - v_+ \frac{\partial \rho_+}{\partial x} - \rho_+ \frac{\partial v_+}{\partial x} \tag{C.9a}$$

$$\frac{d\rho_-}{dt} = \Phi - v_- \frac{\partial \rho_-}{\partial x} - \rho_- \frac{\partial v_-}{\partial x} \tag{C.9b}$$

$$\frac{\partial E}{\partial x} = \frac{e}{\epsilon_0}\rho_+ \tag{C.9c}$$

$$\int_0^d dx E = V_{app} \tag{C.9d}$$

As ρ_- is decoupled, we ignore it for finding the field solution. A general solution is found in [160], but we will consider only the limit where the electric field approaches zero at the anode; thereby establishing the range of usefulness for this solution. Using $v_+ = \mu E$, we can find the solutions for ion density and electric field:

$$E = \frac{2V}{d}\frac{x}{d} \tag{C.10a}$$

$$\rho_+ = \frac{2V}{d^2}\frac{\epsilon_0}{e} \tag{C.10b}$$

The electric field increases linearly across the chamber and the ion density is constant as the spreading of charge due to the increasing velocity is exactly offset by ionization. Note that the electric field adjacent to the cathode is twice the applied applied field (V_{app}/d).

The ionization rate to which Eqn. (C.10) corresponds can be found by the steady-state flux of ions into the cathode $\rho v_+(x = 0)$:

$$\Phi_{crit} = \mu \frac{4V^2}{d^4}\frac{\epsilon_0}{e} \tag{C.11}$$

For the typical case of a 1 mm chamber with 100 V bias the critical intensity is 4.4×10^9 ionizations / cm^3 / μs. As expected, this is even lower than that in §C.1. As the Hadron Monitor ionization rates exceed this rate, we will have to consider more effects to model the detector response.

C.3 High-Intensity Space Charge Effects

When the critical intensity in Eqn. (C.11) is exceeded the electric field in part of the chamber will fall to a very small value such that there is almost no drift of ions or electrons. This region is commonly known as a "dead zone". For steady-state beam conditions, charge ionized in the dead zone is essentially lost as there is no method for it to escape before recombination (e.g., see [160]). As the NuMI beam is pulsed, however, we will need to consider the lifetime of such a dead zone.

With a dead zone present, the ionization chamber essentially compresses to a smaller volume. Part of the volume adjacent to the anode will have no electric field (i.e. be dead), the other part will have the linearly growing field as in §C.2. The length of the reduced chamber is then:

$$x_0 = \left[\frac{\epsilon_0}{e}\frac{4\mu V^2}{\Phi}\right]^{1/4} \tag{C.12}$$

If we assume that all charge ionized in the dead zone is lost[4] then the efficiency is simply $\epsilon = x_0/d$. The field solution is:

$$E = \begin{Bmatrix} 0 & x < (d - x_0) \\ \dfrac{2V}{x_0}\dfrac{x-(d-x_0)}{x_0} & (d-x_0) < x \end{Bmatrix} \tag{C.13}$$

And the ion density solution in the steady-state is:

$$\rho_+ = \begin{Bmatrix} \infty & x < (d - x_0) \\ \dfrac{2V}{x_0^2}\dfrac{\epsilon_0}{e} & (d-x_0) < x \end{Bmatrix} \tag{C.14}$$

Additionally, we know that $v_- = v_+$ in the dead zone, while v_- is still small outside the dead zone. The electric field and ion density are plotted in Fig.

[4]The charge lost in the area outside the dead zone is comparatively small as the ion and electron densities are still comparable to that in §C.1

C.3 for the typical chamber of 1 mm and 100 V at an ionization rate of 10^{10} ionizations / cm^3 / μs. Note again that the maximum electric field (at the cathode located at $x = d$) can be several times the applied electric field as the active volume becomes more compressed.

C.3.1 Pulsed Beam Structure

The steady-state efficiency of $\epsilon = x_0/d$, calculated in the dead zone approximation of above, falls to 82% with an ionization rate of 10^{10} ionizations / cm^3 / μs, which is well within the NuMI range of intensities for the Hadron Monitor. However, the NuMI beam (and test beams) are pulsed with 1.5 or 9 μs beam spills. The time is short enough that we can reasonably expect some charge to escape from the dead zone before recombining.

After the beam pulse ends, charge leaks out of the dead zone as the charge in the active area is no longer being replenished by the beam ionization. We approach this problem by finding the quasistatic solution to the post-ionization situation. We assume the existence of a dead zone of extent $x_1 = d - x_0$ from which charge is leaking. In the active area continuity requires:

$$\rho_+ v_+ = A = \rho \mu E \tag{C.15}$$

Where A is a constant. We substitute for ρ_+:

$$\mu E E' \frac{\epsilon_0}{e} = A \tag{C.16}$$

and integrate to find:

$$E^2 \mu \frac{\epsilon_0}{e} = Ax + B \tag{C.17}$$

where B is a constant of integration. Then we solve for E:

$$E = \left[\frac{e}{\mu \epsilon_0} (Ax + B) \right]^{1/2} \tag{C.18}$$

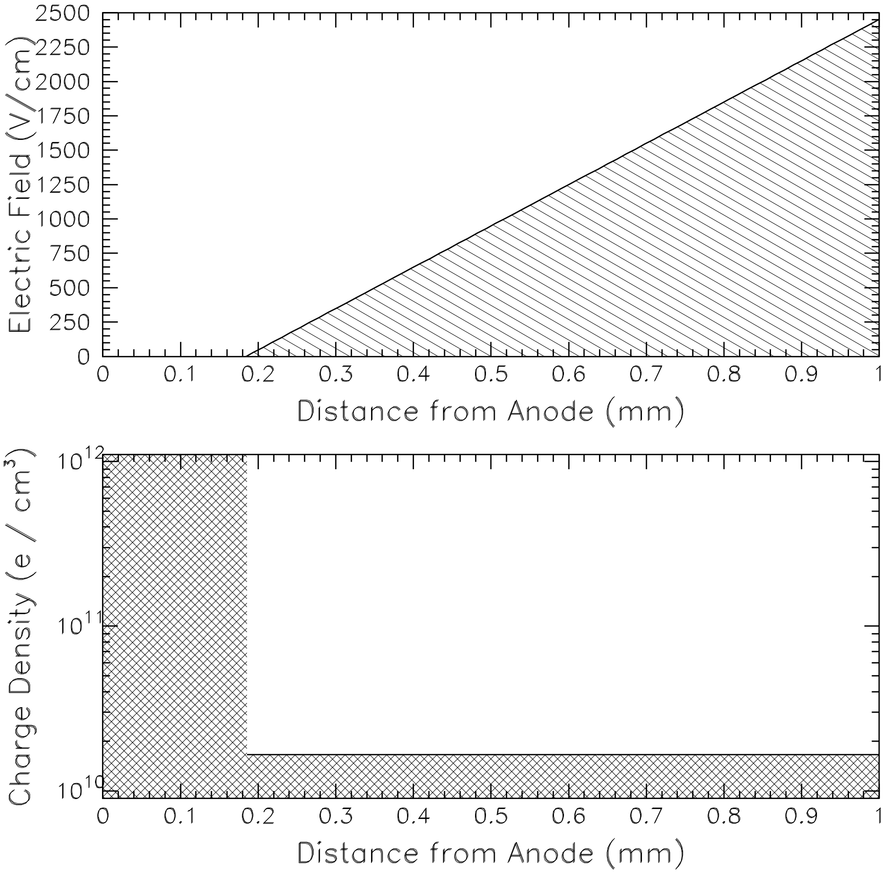

Figure C.3: Steady state field and ion density distributions, calculated with the approximation of §C.3 for a 1 mm chamber with 100 V of bias and an ionization rate of 10^{10} ionizations cc μs. The dead zone extends up to x=0.184 mm; inside the dead zone the ion density approaches infinity.

We require the boundary condition of $E(x_1) = 0$ at the edge of the dead zone, which gives us $B = -Ax_1$:

$$E = \left[\frac{eA}{\mu\epsilon_0}(x - x_1)\right]^{1/2} \tag{C.19}$$

We then apply the other boundary condition Eqn. (8.1d) to find A:

$$A = \frac{9}{4}V^2(d - x_1)^{-3}\frac{\mu\epsilon_0}{e} \tag{C.20}$$

Then we have the solution for the electric field and (by differentiation) the ion density:

$$E = \frac{3}{2}\frac{V}{x_0}\left[\frac{x - x_1}{x_0}\right]^{1/2} \tag{C.21a}$$

$$\rho_+ = \frac{3}{4}\frac{V\epsilon_0}{x_0^2 e}\left[\frac{x - x_1}{x_0}\right]^{-1/2} \tag{C.21b}$$

Then, the charge current out of the chamber is:

$$J = \frac{9}{4}\mu\epsilon_0\frac{V^2}{x_0^3} \tag{C.22}$$

This current will be depleting the charge out of the dead zone. If we assume an initial, uniform, ion density in the dead zone we can find the rate at which x_0 changes and calculate how long the dead zone persists after the end of ionization. We start by considering the growth of the active region due to leakage:

$$\frac{dx_0}{dt} = \frac{J}{\rho_0 e} = \frac{9}{4}\frac{\mu\epsilon_0}{\rho_0 e}\frac{V^2}{x_0^3(t)} \tag{C.23}$$

where x_0 is now allowed to change with time; we also use the accumulated density in the dead zone $\rho_0 = \Phi\tau_p$, where τ_p is the duration of the beam pulse. Integrating to find a solution, we get:

$$\begin{aligned} x_0(t) &= \left[\frac{9\mu\epsilon_0}{\rho_0 e}V^2 t + x_0^4(0)\right]^{1/4} \\ x_0(t) &= x_0(0)\left[1 + \frac{9}{4}\frac{t}{\tau_p}\right]^{1/4} \end{aligned} \tag{C.24}$$

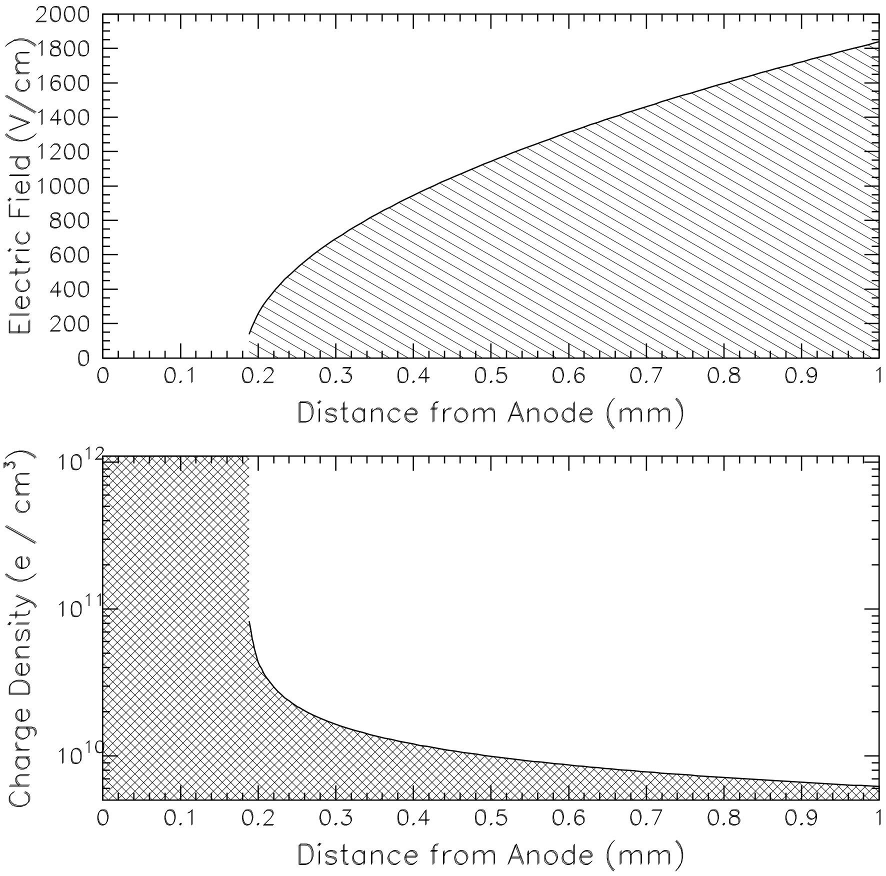

Figure C.4: Quasi-equilibrium post-ionization field and ion density distributions for a 1 mm chamber with 100 V of bias and an ionization rate of 10^{10} ionizations / cm^3 / μs, as described in Eqn. (C.21). The dead zone extending up to x = 0.184 mm was established during the beam ionization, the rest of the density and field are established by ions leaking out of the dead zone.

In these equations t is the time elapsed after the end of the pulse.

Next, we define the depletion time (τ_d) as the time required for x_0 to equal d, such that there is no longer total screening. Solving for x_0(t) = d:

$$\tau_d = \left(\frac{\rho_0 e}{9\mu\epsilon_0 V^2}\right)\left[d^4 - x_0^4(0)\right] \tag{C.25}$$

Consider again the typical case of 100 V over 1 mm at 10^{10} ion./ cm^3 / μs; if we assume that the dead zone forms very quickly[5] and consider a 1.5 μs beam pulse: the dead zone extent (x_1) would be 0.184 mm with a density of 1.5 $\times$ 10^{10} e / cm^3 for a total areal charge density of 2.8 $\times$ 10^8 e / cm^2; the initial current[6] out of the dead zone would be 4.6 $\times$ 10^8 e / cm^2 / μs , the current decreases as the dead zone is eroded, and the final time to erosion is 0.84 μs.

C.3.2 Recombination in the Dead Zone

We now consider recombination in the dead zone. As has been discussed, recombination losses are very small in the active regions of the detector for the intensities relevant to the NuMI beam. However, the electron densities in the dead zone are several orders of magnitude higher than in the active region of the detector. We will consider a model where all of the recombination loss occurs within the dead zone; we assume that the dead zone has reached its equilibrium extent immediately, so the amount of charge sequestered there will be somewhat overstated – subsequently the recombination losses will be an upper bound on what really might be expected.

We again use the recombination density rate in the dead zone $\mathcal{R} = r_e\rho_0^2$

[5]Actually, the dead zone takes on the order of 1 μs to form as enough electrons must move away from the cathode to create a net positive space-charge that produces the screening.

[6]Note that there must be a complementary current of electrons to the anode, as well as escaping ions. To enable this there must be a small electric field persisting in the dead zone – in this case it would be $\sim$ 3 V/cm.

and now define the total areal loss:

$$\begin{aligned}\mathcal{L} &\equiv \iint dt\, dx\, \mathcal{R} \\ &= r_e \rho_0^2 \int_0^{\tau_d} dt\, (d - x_0(t)) \\ &= r_e \rho_0^2 \left[d\tau_d - \frac{16}{45} x_0(0) \tau_p \left[\left(1 + \frac{9}{4}\frac{\tau_d}{\tau_p}\right)^{5/4} - 1 \right]\right]\end{aligned} \tag{C.26}$$

The total collection efficiency is then $\epsilon = 1 - \mathcal{L}/\Phi d\tau_p$. This equation, though ungainly, gives us an analytical model of charge loss in the chamber[7]. It has a complicated structure depending separately on the ionization rate and the length of the ionization pulse. For our typical chamber of 100 V over 1 mm, ionized at 10^{10} ion. cm^3 μs for 1.5 μs the charge loss amounts to only 3.7×10^4 ions / cm^2, for an efficiency of 99.996%. The charge loss, however, scales approximately as the third power of intensity; this is in accordance with our expectation that the loss in the dead zone is much greater than in the active regions.

Figs. C.5 and C.6 shows calculated collection efficiencies for a 1 mm chamber at different biases and ionization rates. Fig. C.5 is for a 1.5 μs beam pulse, Fig. C.6 is for 9 μs beam pulse[8]. Comparing to the calculations of §C.1 and Fig. C.2 we see that including the dead zone in our model reduces the turn-on ionization rate of recombination effects by $\sim$ seven orders of magnitude; this rate is in the upper range of that expected in the NuMI beam. The calculation of this section overstates the effect of recombination as it does not consider the time necessary to create the dead zone and treats the loss to recombination to be linear, so we need to consider the effects to a greater extent. Furthermore,

[7] Note that we treat recombination as small and assume the density changes are not significant enough to reduce the loss rate to recombination. Thus, Eqn. (C.26) overstates the loss to recombination at higher loss rates.

[8] Note that the total ionization in Fig. C.6 is six times that of Fig. C.5 for a given ionization rate.

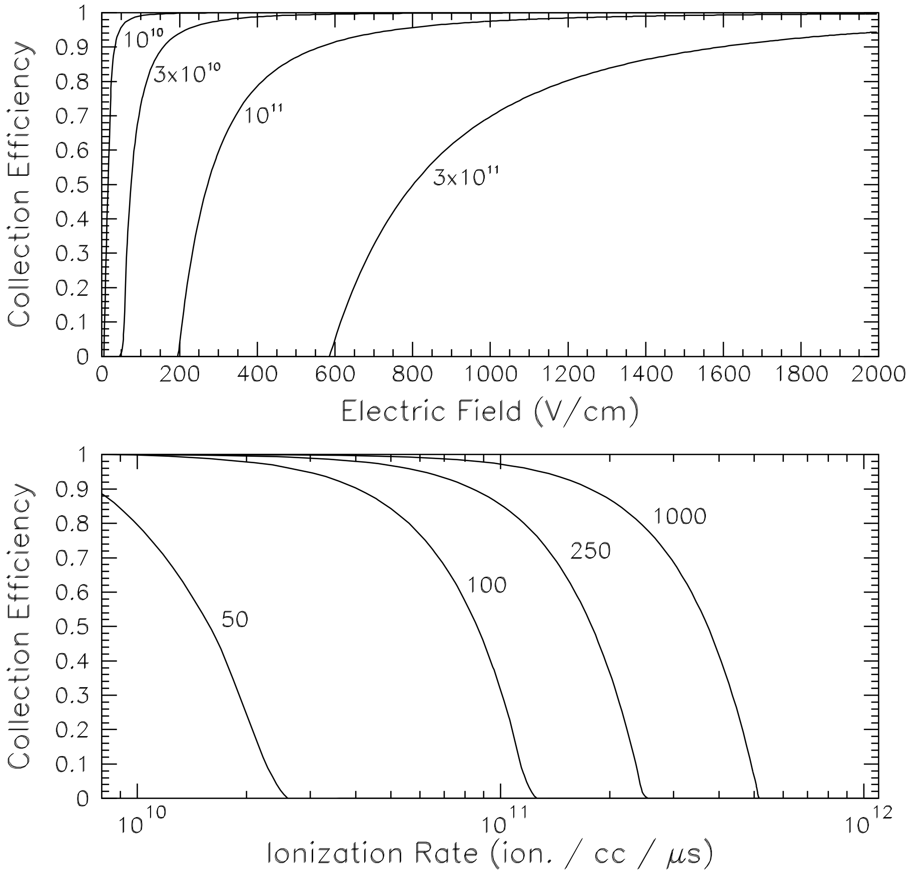

Figure C.5: Calculated collection efficiencies using the method of §C.3 for a pulsed beam that induces creation of a dead zone and subsequent recombination. These plots are for a 1 mm chamber with a 1.5 μs pulse. The upper plot is versus the applied electric field; the curves are for different ionization rates indicated by the numbers in units of ionizations / cm^3 / μs. The lower plot is versus the ionization rate, with the different lines corresponding to different applied electric fields in units of V/cm.

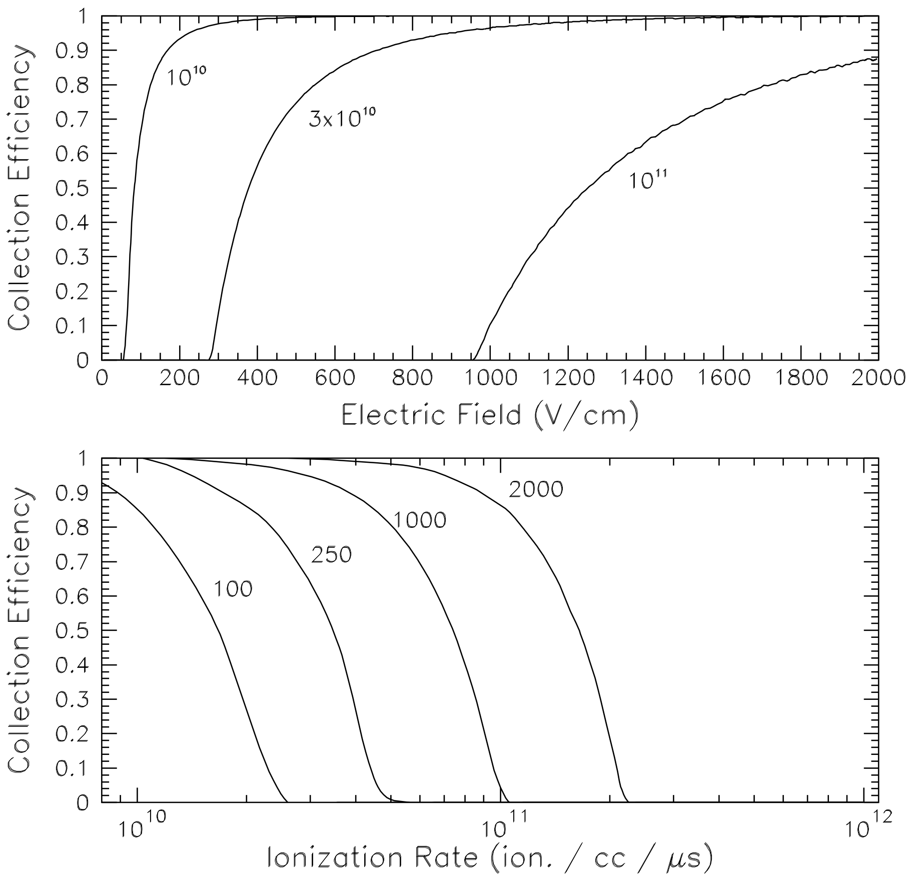

Figure C.6: Calculated collection efficiencies using the method of §C.3. These plots are for a 1 mm chamber with a 9 μs pulse. The upper plot is versus the applied electric field; the curves are for different ionization rates indicated by the numbers in units of ionizations / cm^3 / μs. The lower plot is versus the ionization rate, with the different lines corresponding to different applied electric fields in units of V/cm.

as has been noted above, the electric field at the cathode becomes several times as high as the applied electric field doe to compression of the active volume by space-charge screening; the higher electric field will cause gas multiplication which needs to be considered as well.

Appendix D

Numerical Ion Chamber Simulation

As discussed in Ch. 8, the flow of charge in an ionization chamber is described by a set of coupled partial differential equations:

$$\frac{d\rho_+}{dt} = \Phi - v_+\frac{\partial\rho_+}{\partial x} - \rho_+\frac{\partial v_+}{\partial x} + \alpha\rho_- - r_e\rho_-\rho_+ + D\frac{\partial^2\rho_+}{\partial x^2} \tag{D.1a}$$

$$\frac{d\rho_-}{dt} = \Phi - v_-\frac{\partial\rho_-}{\partial x} - \rho_-\frac{\partial v_-}{\partial x} + \alpha\rho_- - r_e\rho_-\rho_+ + D\frac{\partial^2\rho_-}{\partial x^2} \tag{D.1b}$$

$$\frac{\partial E}{\partial x} = \frac{e}{\epsilon_0}(\rho_+ - \rho_-) \tag{D.1c}$$

$$\int_0^d dx E = V_{app} \tag{D.1d}$$

App. C discussed analytic solutions to these equations under certain limiting cases. The situation provided by the NuMI beam where the beam-pulse duration and ion collection times are comparable requires numerical solution of these equations.

This appendix describes a computer code which performs a finite-element simulation of charge propagation in an ionization chamber. The computer code deposits charge in the chamber from the beam; drifts the charge to the electrodes under the influence of an electric field; and subjects the charge to recombination, amplification, and diffusion. The charge propagates through 1 μm cells along the chamber gap, and the field configurations are recalculated following 1 ns time increments.

The program was implemented in FORTRAN [166] using CERNLIB [167] routines and the program PAW [168] for visualization and analysis. All internal units are SI, except that cells are considered to be 1 cm^2 in area orthogonal to the electric field.

D.1 Program Outline

An flowchart of the program organization is shown in Fig. D.1. The steps are explained below and in the following sections.

Program Initiation

The input files are read and the HBOOK[1] output ntuples are setup along with the logfile.

There are two input files: `input.dat` and `repeater.dat`. `input.dat` gives default input parameters for simulation and physical parameters. If there are to be multiple runs then the `repeater.dat` file is used to specify parameters to be changed and how so.

There are three output ntuples. The first, `100`, is for the final run summary, giving information used to calculate response. The next, `101`, is a "snapshot" ntuple that provides information at a given time interval during the run; it contains a series of numbers that describe the state of the chamber at that time. The final ntuple, `102`, is a more detailed snapshot ntuple that

[1]Part of the CERNLIB package.

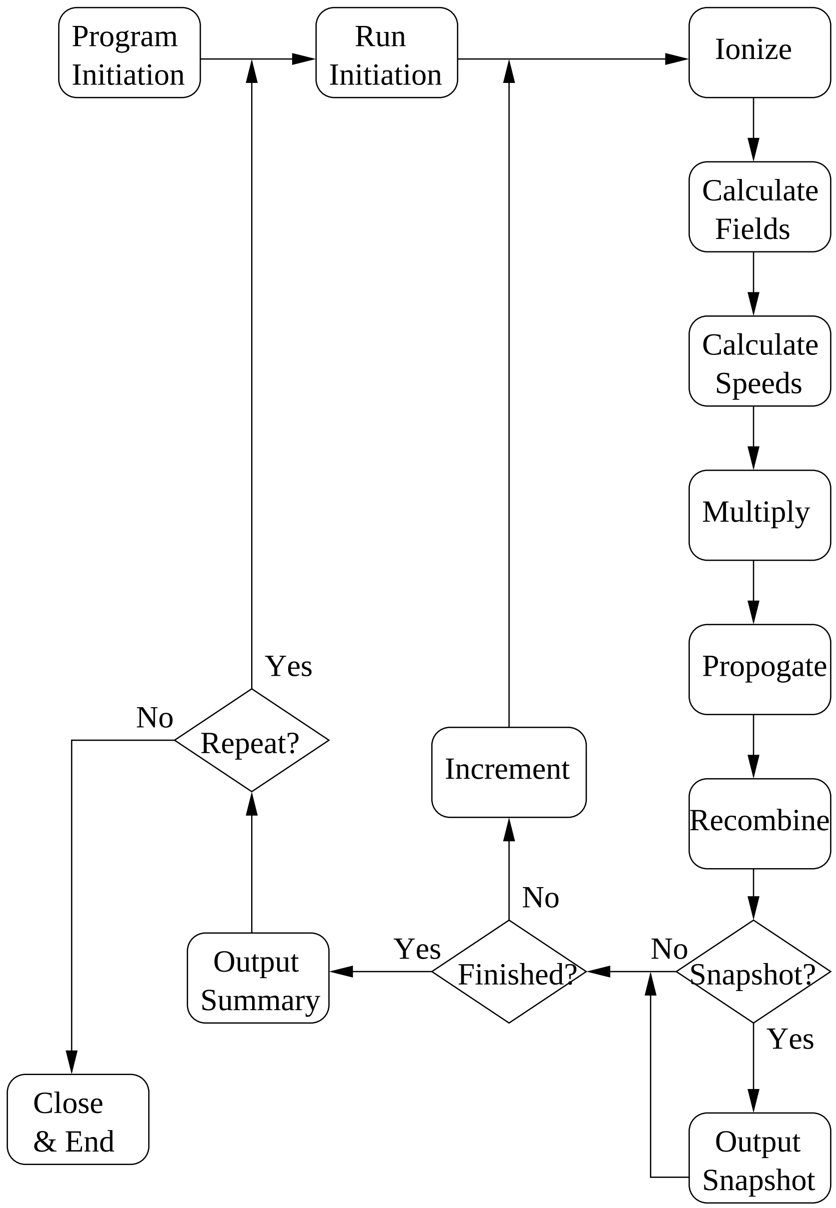

Figure D.1: High-level flowchart of the ionization chamber simulation.

has arrays of the ion and electron densities throughout the chamber, as well as the electric field at each point in the chamber. The three give, in turn 0-, 1-, and 2- dimensional data about the run. The more detailed snapshots are used only for single runs and allow visualization of the chamber dynamics; whereas the summary ntuple is generated on every run.

The logfile provides basic information about each run for runtime information and debugging. It tells the method of exit for each run – success, (type of) spark, or time. Also given are a few of the input parameters for that run: bias voltage, ionization intensity, and chamber gap. As a check, the "time" of the simulation is also given in simulated microseconds.

Run Initiation

When the program is run in repeated mode it will simulate some large number of chambers with different configurations of the input parameters; Run Initiation initializes the variables used and run-specific tables for each set of input parameters used. The electron drift velocity and Townsend coefficient are nonlinear functions of the electric field (see §7.4 & §7.5). The functions are not very complicated, but can consume a disproportionate amount of computing resources. In a typical run, each of these variables is calculated ~ 1 billion times. To optimize running, a hash lookup table is generated for each of these two variables at the start of each run[2]. Typically, 10,000 calculations of each value would be stored in a lookup table; when the variable is determined during the run it is found by linear interpolation between the two adjacent points on the lookup table.

[2]The tables are generated at the start of the run, instead of the start of the program, because the physical parameters might differ between runs

Ionize

Ionization is performed over a preset time period. In each step of this period the electron and ion density of each cell are increased uniformly. The total ionization is kept as a running sum throughout the execution.

Calculate Fields

The electric field at each point is a superposition of the field from the charge on the electrodes and the space-charge in the chamber. Importantly, the charge on the electrodes is affected by the space-charge, in the form of image charges. By using knowledge of the 1-dimensional nature of the problem and assuming that the potential between the two electrodes is fixed, we created an $\mathcal{O}(\mathrm{n})$ algorithm to calculate the electric field at each cell, the induced charge on the electrodes, and the image current. These are illustrated in detail in §D.2. Additionally, a possible exit condition here is if more that 0.5% of the cells have an electric field more than $100\times$ that of the applied field; the run is terminated, registered as a spark in the logfile, and, does not earn an entry in the summary ntuple.

Calculate Speeds

The drift velocities of the ions and electrons are then calculated for each cell. The speeds are calculated before propagation as the electron drift velocities are necessary for the multiplication calculation. The ion velocities are calculated through a simple multiplication by the mobility. The electron velocities are determined by using the lookup table as described above.

Multiply

The first Townsend coefficient is determined from a lookup table as described above. As the the time and distance steps are so small the coefficient is applied linearly, instead of exponentially:

$$\rho_{\pm}(x, t+\delta t) = \rho_{\pm}(x,t) + \rho_{-}(x,t)\alpha(E)v_i(E)\delta t \tag{D.2}$$

There is a running sum of the Townsend coefficient across the chamber. If the coefficient satisfies the condition:

$$\int_0^d dx\, \alpha(E) > 20 \tag{D.3}$$

then it is said to have satisfied the Raether condition [169], which is another criterion for breakdown within the chamber. The run is then ended, listed as a spark in the logfile, and not entered into the summary ntuple.

Propagate

Propagation of charge is achieved by essentially picking the charge up out of a cell and depositing it a certain number of cells downstream – this ends up being two cells as the calculated number of cells is never exactly integral. The process is described in more detail in §D.4, and results in a sort of diffusion that is managed to be about that of natural diffusion.

Recombine

Recombination was implemented in several ways. The most useful type is performed according to the process described in §7.5, except it is linearized due to the small steps:

$$\rho_{\pm}(x, t+\delta t) = \rho_{\pm}(x,t) - r_e\rho_{+}(x,t)\rho_{-}(x,t)\delta t \tag{D.4}$$

The above method is superior to steady-state studies where all the charge within the dead region was considered to have been lost [160]. In such a method, one defines a minimum electric field, below which the ion and electron densities were lowered by the same amount such that one of them equaled zero (to conserve charge). We found this method to not only overstate loss, but to also introduce instability into the simulation at the interface between the dead zone and active region.

Snapshot? & Output Snapshot

If a particular number of time steps have occurred since the last snapshot, the program outputs variables about the state of the system to ntuple `101` and dumps the electric field and ion and electron densities to ntuple `102`. The distributions are grouped into super-cells over which the cells are averaged with equal weight. The parameters for output frequency and detail are specified by the input.

Finished?

A number of completion criteria are checked to determine if the program must continue running. The criteria involve some calculation, so are only checked at the same frequency as the Snapshot Output (usually 100 or 1000 time steps). The run is defined as coming to a proper conclusion if ionization has ended and the amount of charge remaining in the chamber is several orders of magnitude smaller than the amount collected. The run is declared a spark if the charge collected is more than 100× as much as that ionized. The run is canceled if it has lasted 100× longer than it would nominally take the charge to cross the chamber after ionization.

Increment

During ionization the time is incremented by the defined time step and the cycle repeated. However, we have found that the simulation is more stable after the end of ionization and can withstand larger time steps. Thus, there are parameters that allow a larger time step after ionization, and one that scales the time step as the execution continues.

Output Summary

When a run finishes an output summary is written to ntuple `100` as long as the run was not declared a spark or to have gone to long. The exceptions are to avoid the `nans` that usually accompany any extreme behavior.

Repeat?

If multiple runs with different parameters were requested by the input the repeat function scans through the space specified in the input. The parameters to be modified can be changes linearly or exponentially and include physical parameters as well as simulation parameters like cell size.

Close & End

At the end of the program there is only simple housekeeping to properly close the logfile and ntuples.

D.2 Space Charge Computation

As has been seen earlier, the electric field can be significantly modified by the presence of net space charge within the chamber. The electric field is the

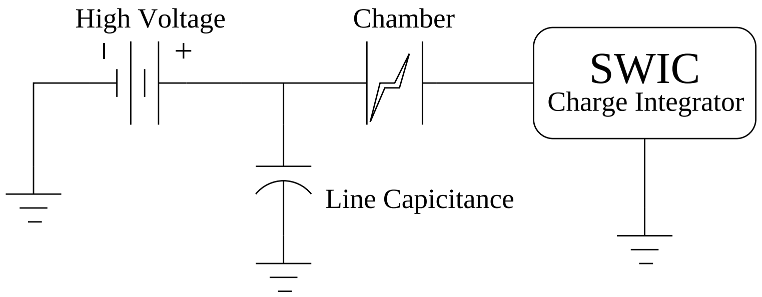

Figure D.2: Circuit diagram of the chamber from the point of view of image charge. The space charge within the chamber can cause the potential between the electrodes to vary. Charge is transferred from the ground and high voltage to compensate.

solution of Eqns. (8.1c) and (8.1d); however, instead of solving this at each time step, we use a less computation-intensive approach.

The electric field has three components: the applied electric field, the space-charge field from the charges in the chamber, and the image-charge field caused by the space charge. The image charge can be thought of as being the image of the charges in the electrodes. More usefully for our application we consider the image charge to be a result of the field of the space charge – the space charge creates a potential difference between the two electrodes that differs from Eqn. (8.1d); this difference is compensated by charge flowing between the charge reservoirs and the electrodes balancing the potential (see Fig. D.2).

We start by finding the electric field of the space-charge by itself; the ions and electrons in the chamber are considered to be infinite flat sheets of charge. We define the quantities $Q_{\downarrow}(x)$ and $Q_{\uparrow}(x)$ as all of the charge below

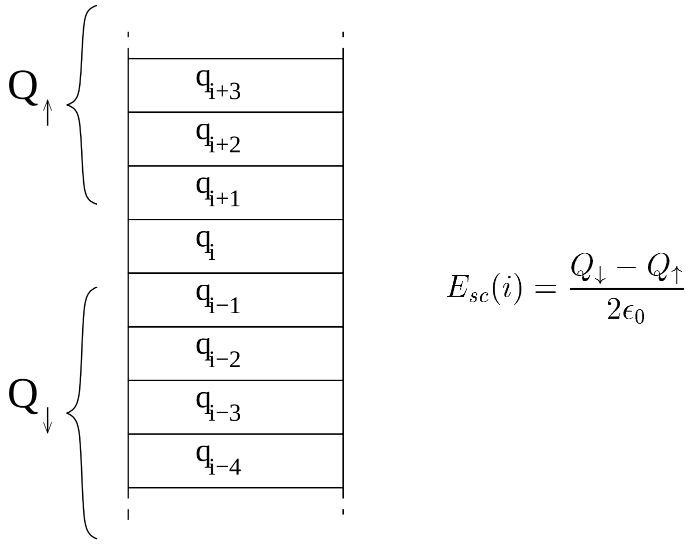

Figure D.3: Illustration of space-charge field calculation method. The space-charge portion can be determined solely from running sums of the charge above and below the cell.

or above that point in the chamber:

$$\begin{aligned} Q_\uparrow(x) &= \int_x^d dx' \, \frac{e}{\epsilon_0}[\rho_+(x') - \rho_-(x')] \\ Q_\downarrow(x) &= \int_0^x dx' \, \frac{e}{\epsilon_0}[\rho_+(x') - \rho_-(x')] \end{aligned} \tag{D.5}$$

Then the electric field is found as treating them as sheets of charge:

$$E_{sc}(x) = \frac{Q_\downarrow(x) - Q_\uparrow(x)}{2\epsilon_0} \tag{D.6}$$

To minimize the necessary computation, this process is not repeated for every cell. Instead, the Qs are calculated for only the first cell; each successive pair can be calculated by adding or subtracting the charge in that layer (see Fig. D.3).

As the electric field is calculated for each cell, a running sum is also maintained allowing a potential difference to be calculated as an integration of the electric field to that point:

$$V_{sc}(x) = \int_0^x dx'\, E_{sc}(x') \tag{D.7}$$

The voltage from integrating the electric field gives the potential difference at the two electrodes due to the space charge. As the nominal potential is maintained by exterior sources, there must be compensating charge present on the two electrodes to offset the space-charge potential. The image charge present on the electrodes is then:

$$\sigma_{sc} = -\frac{\epsilon_0 V_{sc}}{d} \tag{D.8}$$

Generally V_{sc} will be negative, so the image-charge on the chamber will be positive (i.e. in the same direction as the flow of charge).

D.3 Image Current

Any change in the image charge on the chamber results in an image current that is measurable as a current in the line to ground. Thus, the measured signal is not only the ionization current through the electrode, but also the image current from charge moving through the chamber. Here, we briefly consider the image current from charge motion. Usually, the electronics we use measure the total charge over a time window much longer than the ionization period, so the question is moot.

Consider an isolated charged particle between two grounded, conducting planes[3]. The particle induces charge on each of the planes to maintain the

[3]To accommodate the applied electric field in an ion chamber we simply use superposition.

potential – this charge depends on the particle position:

$$q_1 = -q\left(\frac{d-x}{d}\right)$$
$$q_2 = -q\left(\frac{x}{d}\right) \tag{D.9}$$

Where q is the charge of the particle and the electrodes are (q_1) at $x = 0$ and (q_2) at $x = d$. As the particle moves it will create a current:

$$I_{img} = \frac{dq_1}{dt} = -\frac{dq_2}{dt} = q\frac{v}{d} \tag{D.10}$$

Where v is the velocity of the particle and I_{img} is defined as being in the same direction as the particle's current.

When the charge is at the electrode ($x \sim d$) the induced charge is $-q$. Thus, as the charge is collected on the electrode, the collection current is exactly canceled by the oppositely directed current. The total charge flow caused by a particle created at position a will then be:

$$Q = \int_a^d dx\,\frac{q}{d} = \frac{q(d-a)}{d} \tag{D.11}$$

For an electron-ion pair where the ion travels toward $x = d$ and the electron toward $x = 0$:

$$Q = \frac{e(d-a)}{d} + \frac{(-e)(0-a)}{d} = e \tag{D.12}$$

As expected, transport of an ion pair through the chamber produces an integrated current equal to the electron charge. The current has a fast component due to the electron drift, and a slow component due to the ion. The magnitude of the currents depends on the drift velocities, and the integrated currents depend on the ionization location within the chamber.

Next, consider the case of recombination as shown in Fig. D.4. The current from a single ionization is truncated by the recombination. Consider

Figure D.4: Illustration of charge flow in the case of recombination.

the ionization on the left: the electron reaches the anode, but the ion is recombined shortly after ionization – the contribution to the current is less than one charge. However, the recombination requires an electron from another ionization that is similarly truncated. The sum of these two currents gives one charge – from two ionizations.

Within the simulation image charge is calculated as described above at each time step. Image charge is calculated simply as the change in image charge at each step. For simulating the current we consider the sum of the flux into the electrode and the image current as what will be measured.

D.4 Charge Transport

Generally, the charge current density varies with not only the charge density and velocity, but also the gradients of those quantities (see Eqn. (8.1)). To avoid calculating these gradients and to strictly conserve charge we move cells of charge individually according to the velocity at each location – naturally allowing stretching and compression of the charge distribution. In detail, this is accomplished by calculation of the spatial step necessary in units of cell widths, and splitting that into integral and fractional parts:

$$\begin{aligned} \Delta_\pm &= \frac{v_\pm\,\delta t}{w_{cell}} \\ &= n_\pm + f_\pm \end{aligned} \tag{D.13}$$

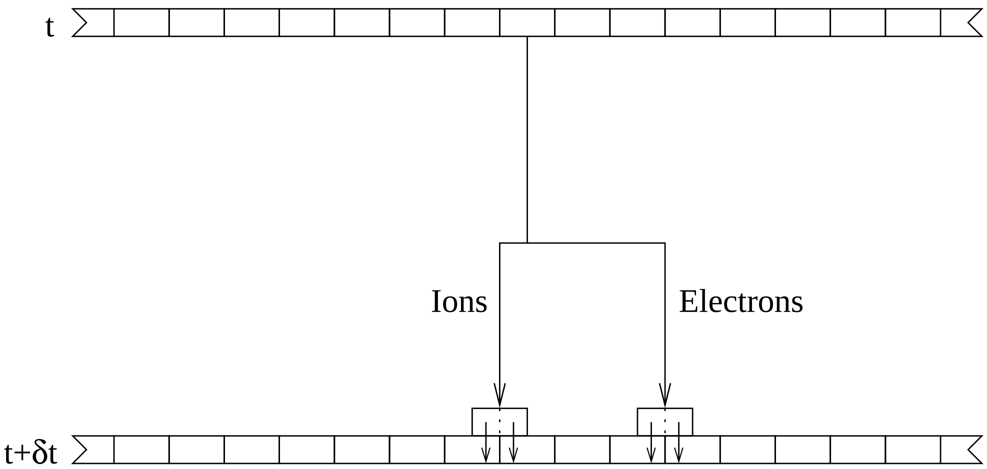

Figure D.5: Illustration of charge transport in cells. The drift velocity of ions and electrons are calculated for every cell. The charge is redistributed among two cells such that the mean speed is maintained.

w_{cell} is the individual cell width; $n_\pm$ and $f_\pm$ are the integral and fractional portions of the step. The charge is then split up fractionally among two downstream cells to strictly conserve the average velocity. This process is illustrated in Fig. D.5 and expressed quantitatively:

$$\begin{aligned} \delta\rho_\pm(i+n_\pm, t+\delta t) &= (1-f_\pm)\rho_\pm(i,t) \\ \delta\rho_\pm(i+n_\pm+1, t+\delta t) &= f_\pm\rho_\pm(i,t) \end{aligned} \tag{D.14}$$

Implicit Diffusion

The splitting of charge from cells results in that charge spreading out over time – making a kind of diffusion. For the ions, the step is typically less than a cell's width. In this case the charge from a single cell will develop a Poisson distribution with the average distance corresponding to the velocity

$$\overline{x} = vt = vN\,\delta t = Nw_{cell} \tag{D.15}$$

Where N is the number of steps taken. The spread of the charge can then be equated to the form for a diffusion coefficient:

$$\begin{aligned} \sigma_x &= \sqrt{N w_{cell}^2} \\ &= \sqrt{v t w_{cell}} \\ &= \sqrt{2Dt} \end{aligned} \tag{D.16}$$

to get an effective (or implicit) diffusion coefficient:

$$D = \frac{v w_{cell}}{2} \tag{D.17}$$

For typical simulation parameters of a 1 mm gap, 100 V bias, and a cell width of 1 μm this gives $D = 1$ cm^2/s for ions, which is roughly twice that naturally expected. This value, however, varies with the local electric field reaching a maximum when one-half cell width is traversed each time step. The maximum diffusion coefficient is then:

$$D_{max} = \frac{w_{cell}^2}{4\,\delta t} \tag{D.18}$$

To limit the maximum diffusion for electrons we must manage the time step such that the maximum diffusion does not become too large. For typical run parameters of 1 μm cell width and 100 ps time step[4] the maximum diffusion rate that electrons might experience is 25 cm^2/s, somewhat smaller than the 160 cm^2/s expected naturally for electrons.

[4]The simulation was found to be unstable at higher intensity with time steps of 1 ns or larger. This is thought to be because some electrons would jump too large of distances when the electric field had a steep gradient. This could reverse the electric field resulting in "hiccups" whereby charge was ejected quickly from the dead zone. To avoid this, all simulations were performed with time increments of 200 ps or less for higher intensity or precision runs.

D.5 Others Issues

The simulation described above is limited by the quality of physical data available and the level of detail simulated. Here we discuss a few areas of detail that were investigated, but never fully simulated: a Second Townsend Coefficient that increases the signal, High Voltage sag that reduces the signal, and the problem of electron acceleration which also reduces signal in this case.

Second Townsend Coefficient

The Second Townsend Coefficient can describe a number of phenomena that cause electrons to be emitted from the cathode in an ion chamber. In our case, we might have some electrons emitted as Helium ions impact the cathode. When the Helium ion is neutralized there is 22.5 or 24.6 eV of excess energy generated. Sometimes, a portion of this energy will be carried away by an electron ejected from the cathode. In our case, we expect this process to be suppressed as we have He_2^+ ions that will dissociate upon neutralization – providing two bodies to carry away the energy. While no directly relevant measurements exist for our system, we could expect as many as 0.15 electrons per ion impact. The electrons emitted would also have a potentially greater effect as they must traverse the entire active region, including the peak field region.

A mechanism for Second Townsend effects was implemented in the simulation and some simulations carried out. The results gave greater signal that the previous simulations, but there was not any significant difference. To avoid adding another arbitrary degree of freedom to the problem, we consider the Townsend effect to have only a multiplicative affect on w and perhaps shift the multiplication parametrization to lower voltages. As we know neither to

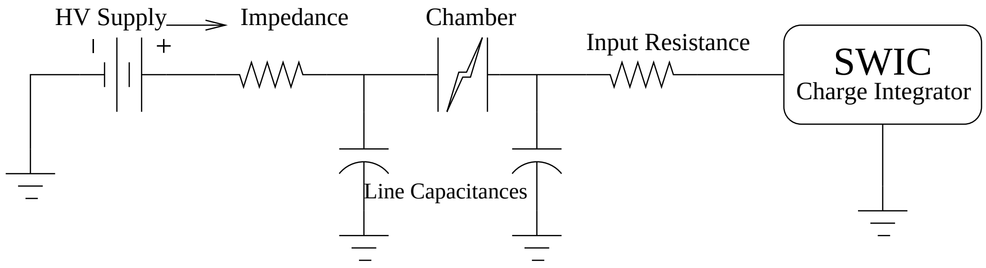

Figure D.6: Circuit diagram for of the chamber system from the perspective of HV sag. Each electrode of the chamber is connected by an effective RC circuit to its reservoir of charge. If too much charge is deposited too quickly on one of the electrodes the electrode potential will vary from that of the reservoir.

any precision the situation is not changed.

High Voltage Sag

When considering the quantity of image and ionization current involved we realized that the method by which charge on the plates is replenished is significant. As shown in Fig. D.6, each electrode of the chamber is separated from the charge reservoirs by an RC circuit. In the case of the beam test, there was an RC filter on the HV supply to reduce ripple providing a capacitance of 47.6 nF and a resistance of 1 MΩ (21 Hz); on the signal side there was only a 10kΩ resistance and the capacitance from the coaxial cables for 3.8 nF (26 kHz). In the beam test as much as 100 nC was deposited in a single beam spill; reducing the potential of the HV side of only 2 V, but on the signal side by 26 V – with decay constants longer than the spill length. Thus, it is possible that the applied voltage was effectively reduced by 30 V.

Some steps were taken to simulate this effect, but were abandoned. Instead, care was taken with the design of the final system to avoid sag. The input RC filter was designed to have a larger capacitance and smaller resistance, additionally, the input resistor was removed from the readout electronics allowing an almost direct connection to ground.

Electron Acceleration

In the simulation above we used the First Townsend Coefficient as a direct coefficient to multiplication. However, multiplication requires that electrons have enough energy to ionize the gas. Generally, measurements are performed when the electron energy distribution has reached its quasi-equilibrium value. To generate the 24.6 eV necessary for ionization the electron must travel, at minimum, through 24.6 V of potential drop. Usually this is just a small correction; however, in this simulation a large portion of the multiplication occurs near the cathode where the electrons are still gaining energy. The effect of this is probably a correction such that we should have an effectively smaller Townsend coefficient. As our knowledge of this coefficient is so poor the effect is negligible.

Appendix E

Simulation Results for 2 mm Ionization Chambers

We consider the simulation of 2 mm chambers like those used in the Booster beam test. The simulation is the same as described in Ch. 8 and App. D. The eventual beam monitoring system uses no 2 mm chambers, but simulation is still worthwhile as it provides comparison to the experimental beam test results

Field and Charge Distribution Evolution

A 2 mm chamber simulation is shown in Fig. E.1 with 10^{11} ion. / cm^3 / μs, for 1.5 μs, biased at at 200 V. The dead zone extends to over half of the chamber. The half of the 2 mm chamber near the cathode has almost the same charge distributions and field as the 1 mm chamber in Fig. 8.2, while the other half has no electric field and high charge densities. If the chamber were instead biased at 400 V, to maintain an applied field of 2000 V/cm, the potential difference would be compressed to a region $\sim$ 1 mm in size, and consequently

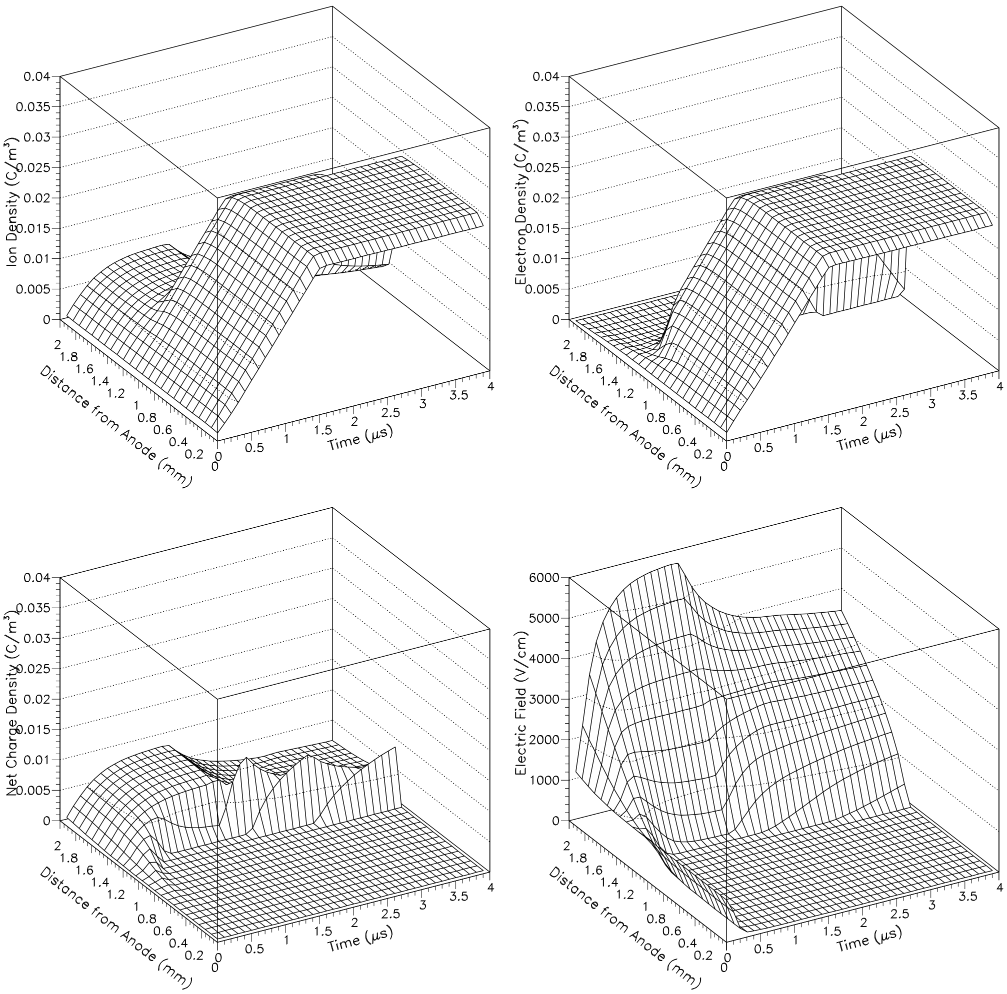

Figure E.1: Same as in Fig. 8.1, except for a 2 mm chamber ionized at a rate of 10^{11} ionizations / cm^3 / biased at 200 V. The behavior of the half of the chamber near the cathode is almost identical to the 1 mm chamber in Fig. 8.2. The dead region persists much longer in the half of the chamber near the anode, resulting in greater recombination.

the field distributions would be almost equivalent to a 1 mm chamber biased at 400 V. With multiplication considered, such a chamber would spark.

Chamber Response: Linearity & Plateau

Fig. E.2 show the linearity of simulated response for 2 mm chambers. The recombination is evidently much greater an effect than with the 1 mm chamber; however, multiplication is clearly occurring at 200 & 250 V (1000 & 1250 V/cm) whereas they might not be expected to reach that range from the 1 mm results. However, as shown in §8.1.1 and Fig. E.1, when sufficiently compressed the size of the active region is independent of the full size of the chamber. Thus, the active region in a 2 mm chamber at 250 V is the same as that as the 1 mm chamber at 250 V and leads to just as much recombination. Some difference in the response remains as the extra 1 mm contributes to twice as much total ionization for normalization and greater recombination losses.

The plateau for the 2 mm chamber also shows some of this variation in Fig. E.4. The recombination losses are much worse at low field and multiplication is caused above 150 V, whereas the low intensity pulses show almost no multiplication up to 300 V. The crossing point here has been shifted to ~ 190 V, an increase of 46% for a 100% increase in chamber gap. Also, the highest intensity pulse of 10^{12} does not meet the crossing point, indicating that the recombination is dominating the multiplication – its gain is always less than that of the 2×10^{11} pulse. Further increases in ionization rate only move the curve to the right.

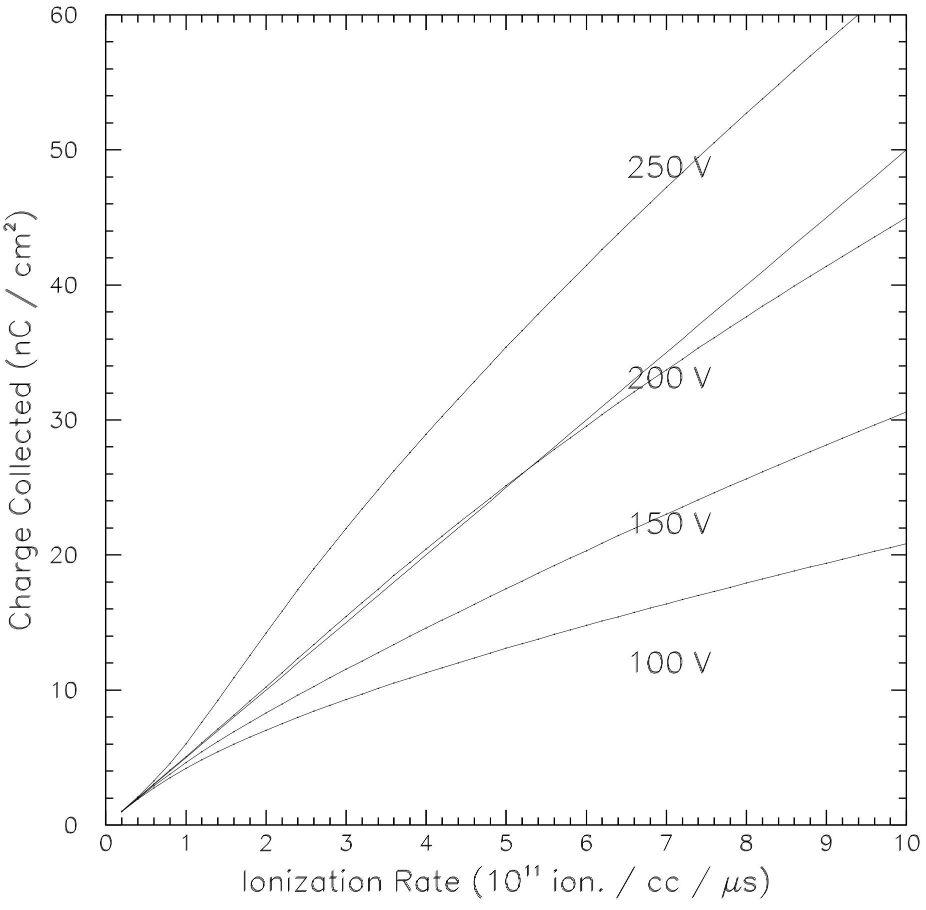

Figure E.2: Results of simulation for a 2 mm chamber at various bias voltages. The total charge collected is plotted as a function of the ionization rate; a straight line is drawn where the total charge collected would be equal to that ionized within the chamber by the beam. Nominal parameters for multiplication and recombination are used.

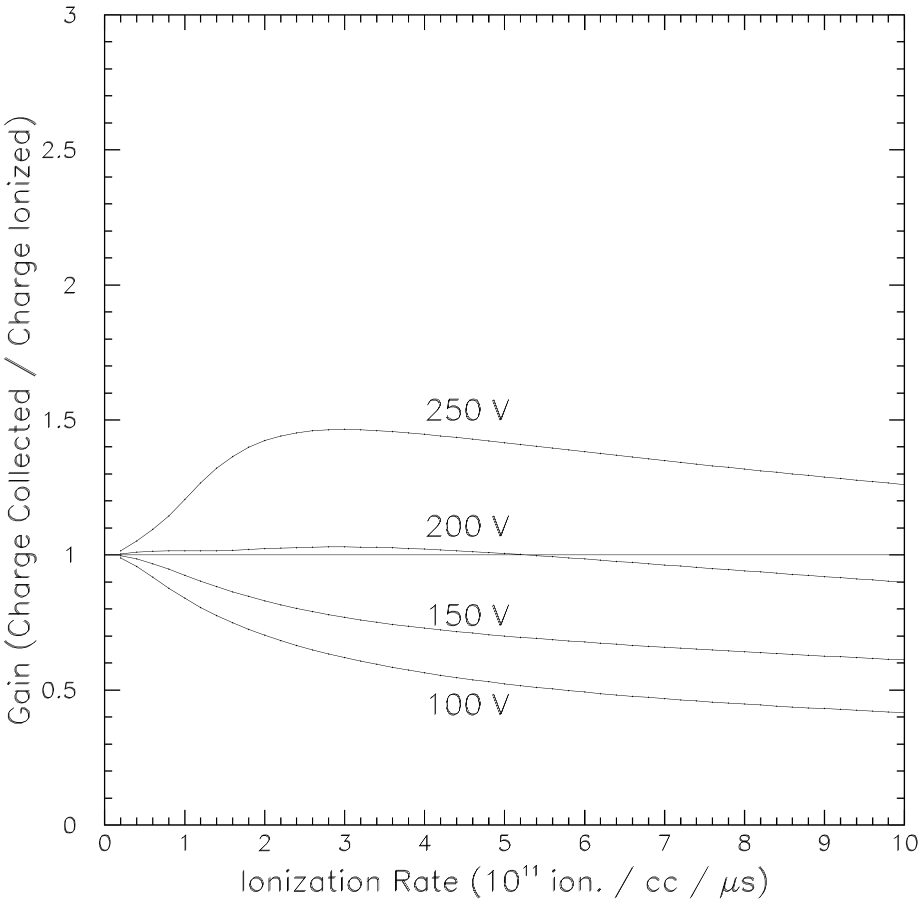

Figure E.3: Results of simulation for a 2 mm chamber at various bias voltages. The "gain" is plotted as a function of the ionization rate. Nominal parameters for multiplication and recombination are used. While 250 V of bias results in the same nominal electric field as 125 V on a 1 mm chamber the response behavior does not scale in the same way.

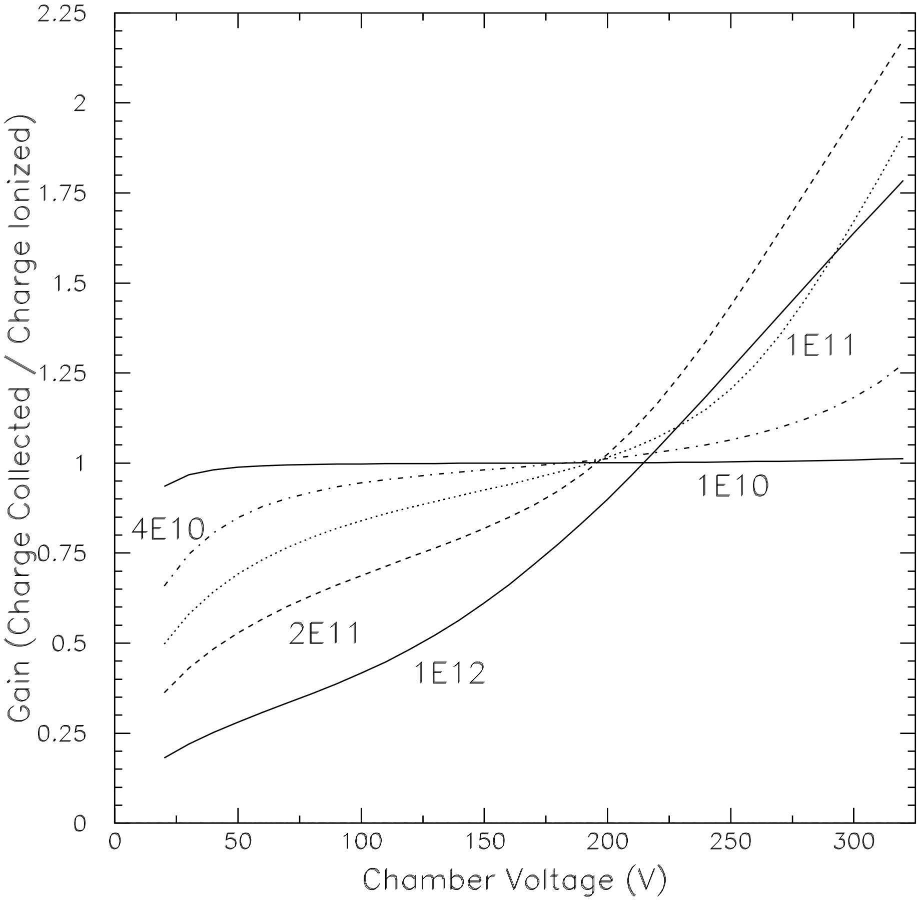

Figure E.4: Results of simulation for a 2 mm chamber at various ionization rates (in units of ionizations / cm^3 / μs). The "gain" is plotted as a function of the bias voltage. Nominal parameters for multiplication and recombination are used. Note that increased intensity increases recombination (decreases signal) at low biases, and increases multiplication (and signal) at high intensity. Several curves cross each other at approximately the same voltage of 190 V – the highest intensity curve suffers too much recombination to approach the crossing point.

Appendix F

Monitor Distributions & Projections

In this appendix we provide the 1-dimensional profiles of the 2-dimensional distributions in Ch. 9. The data are averaged along the seven or nine chambers in a column or row. We also supply duplicates of the distributions, but plotted in the "lego" form that allows closer examination of adjacent bin heights.

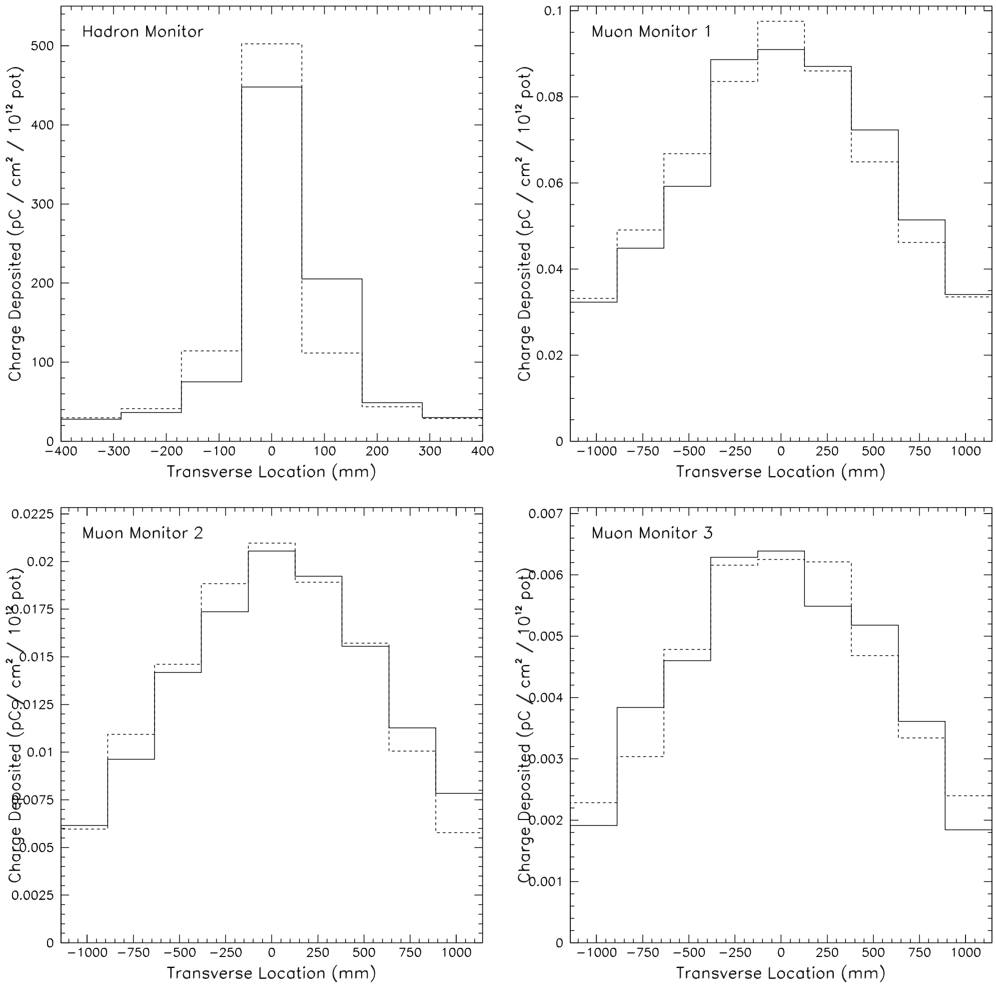

Figure F.1: Measured charge profiles at the monitoring stations with the proton beam, but no target, such that the beam hits the absorber. The solid line is the horizontal profile, the dashed is the vertical. Alcoves 2 and 3 show noise as their signal is barely larger than background. This data corresponds to that of Fig. 9.1.

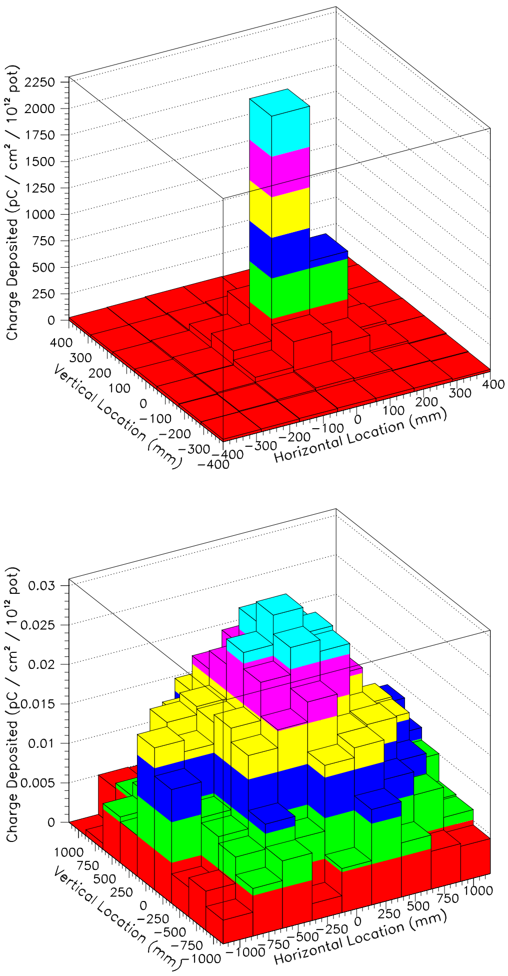

Figure F.2: Measured charge distributions at the monitoring stations with the proton beam, but no target, such that the beam hits the absorber. The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right. Alcoves 2 and 3 show noise as their signal is barely larger than background. This data corresponds to that of Fig. 9.1.

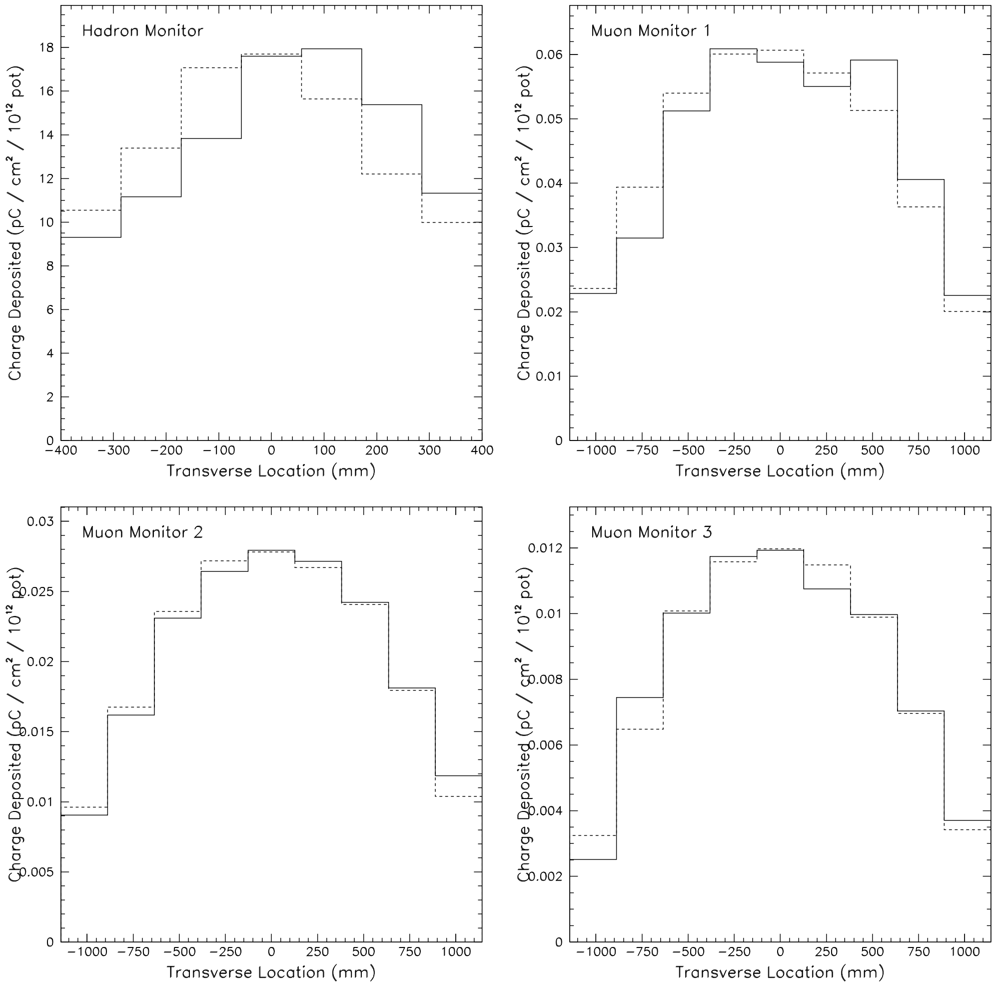

Figure F.3: Measured charge profiles at the monitoring stations with the bare target beam, such that the proton beam hits the target, but there is no horn focusing. The solid line is the horizontal profile, the dashed is the vertical. This data corresponds to that of Fig. 9.2.

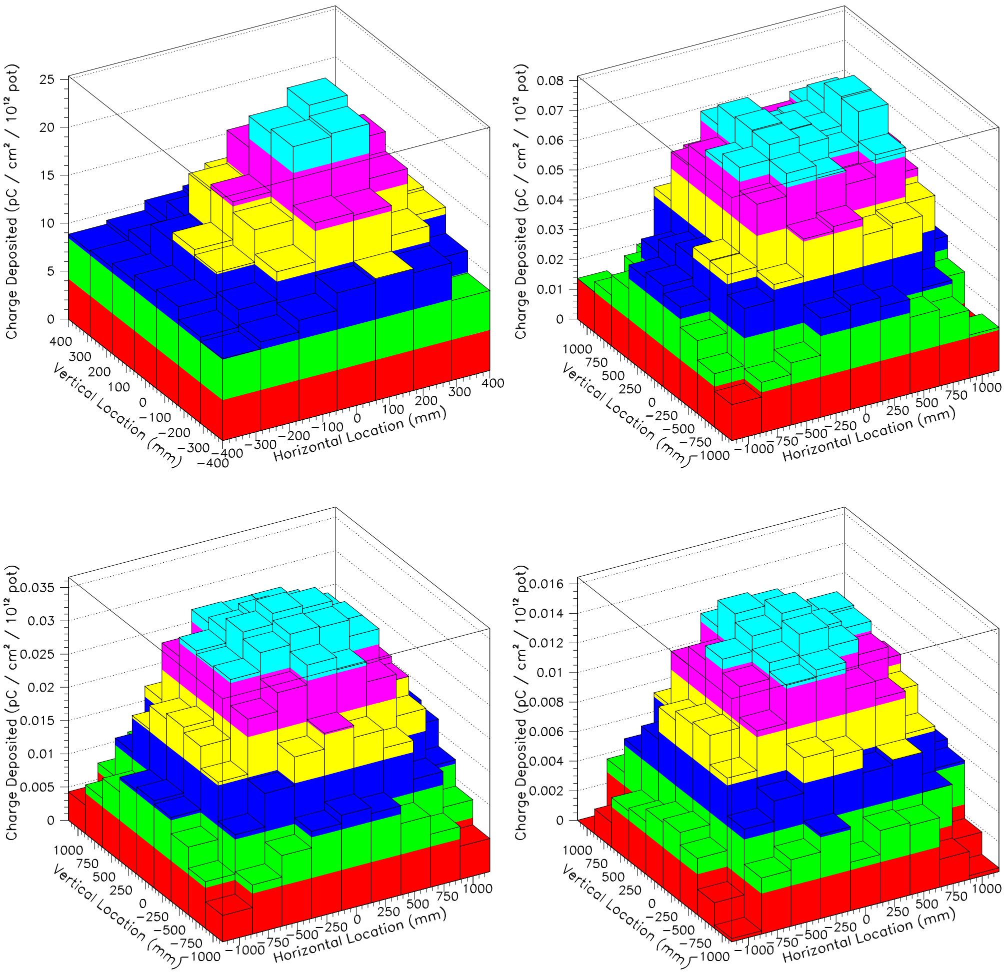

Figure F.4: Measured charge distributions at the monitoring stations with the bare target beam, such that the proton beam hits the target, but there is no horn focusing. The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right. This data corresponds to that of Fig. 9.2.

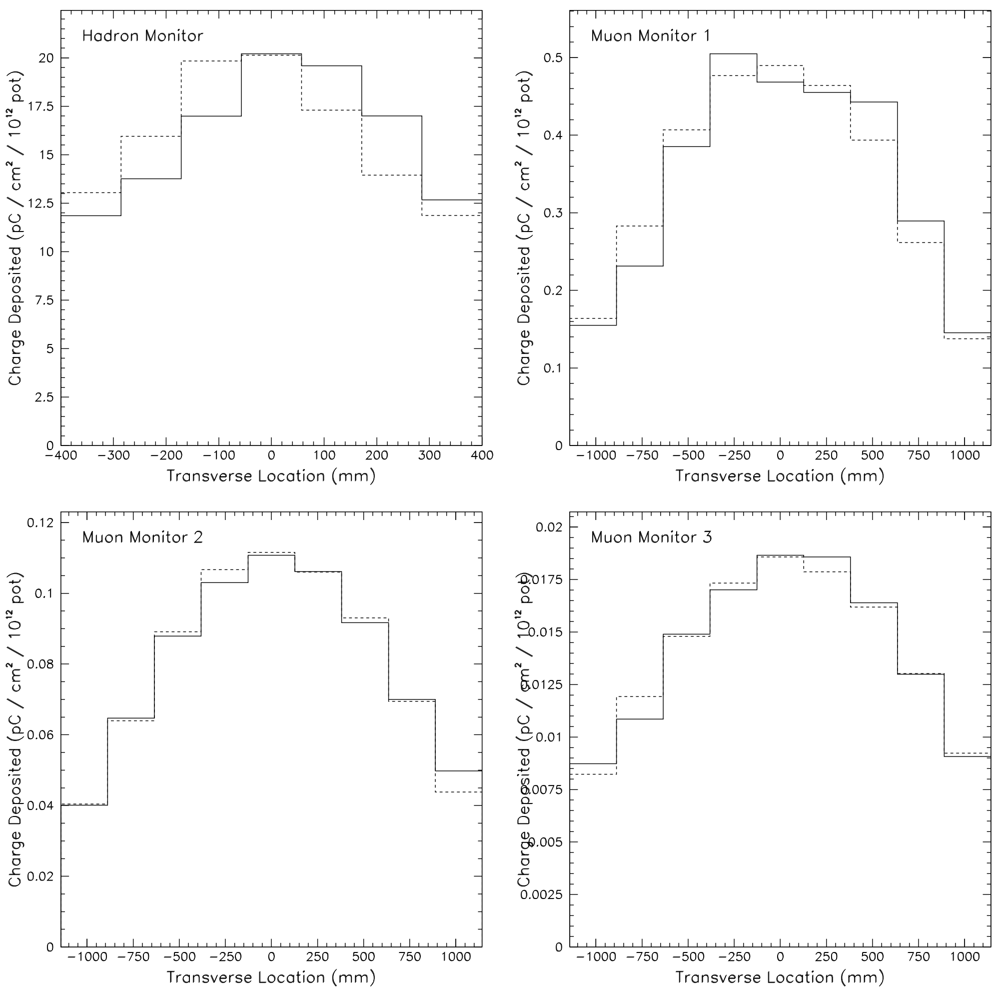

Figure F.5: Measured charge profiles at the monitoring stations with the beam centered on the target in the pME position and with horn focusing. The solid line is the horizontal profile, the dashed is the vertical. This data corresponds to that of Fig. 9.3.

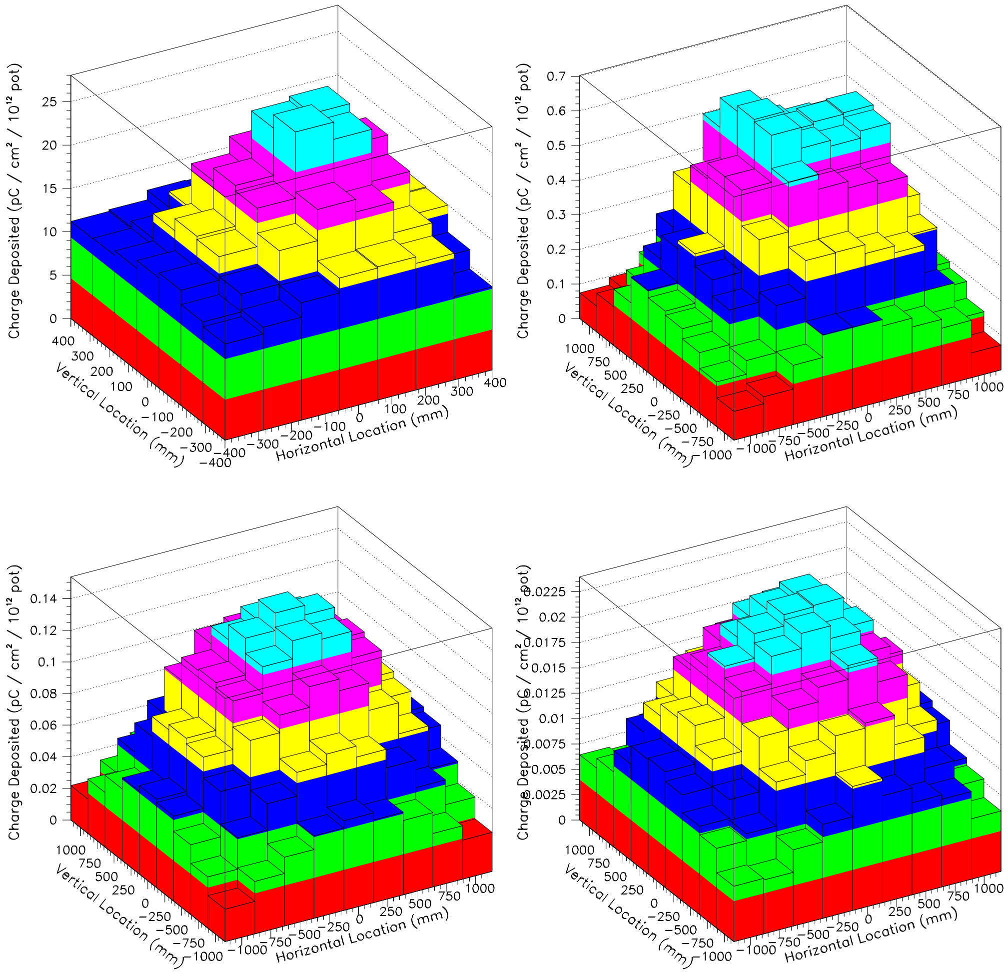

Figure F.6: Measured charge distributions at the monitoring stations with the beam centered on the target in the pME position and with horn focusing. The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right. This data corresponds to that of Fig. 9.3.

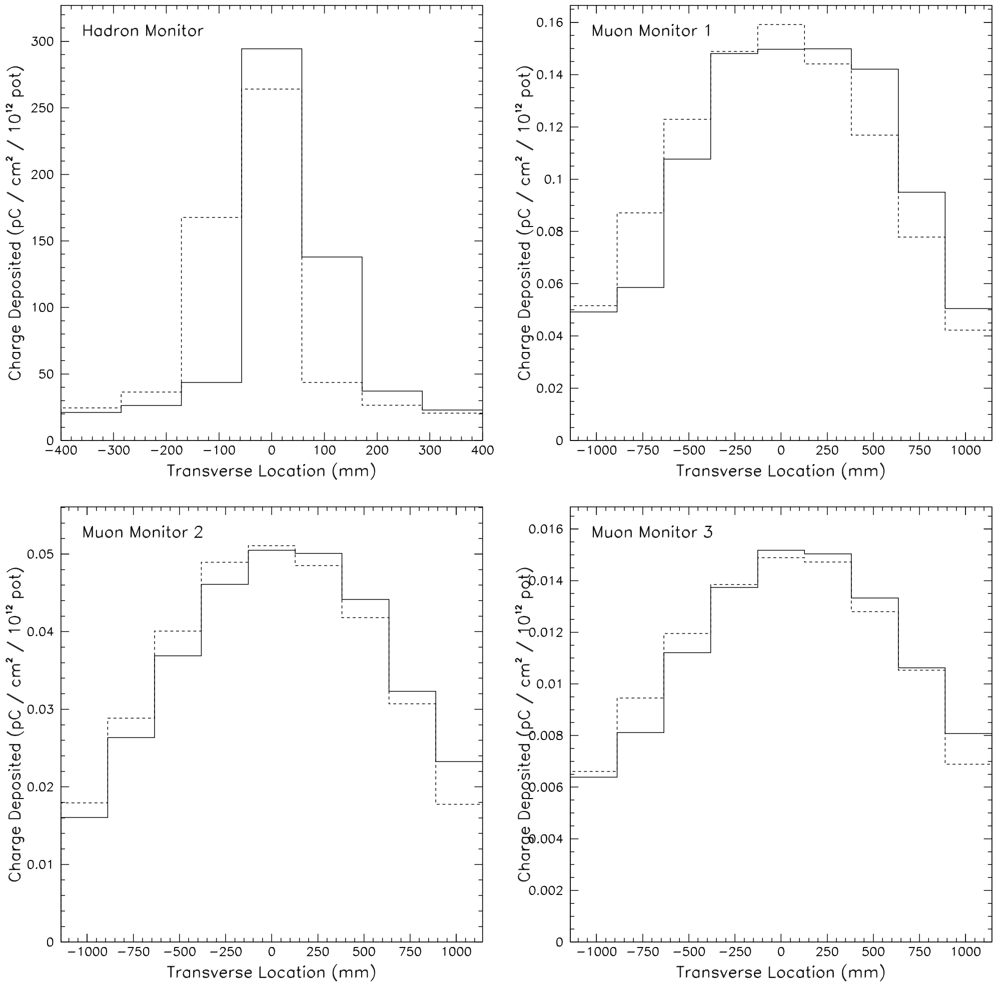

Figure F.7: Measured charge profiles at the monitoring stations with the proton beam in between the baffle and target in the pME position; this is with horn focusing. The solid line is the horizontal profile, the dashed is the vertical. This data corresponds to that of Fig. 9.4.

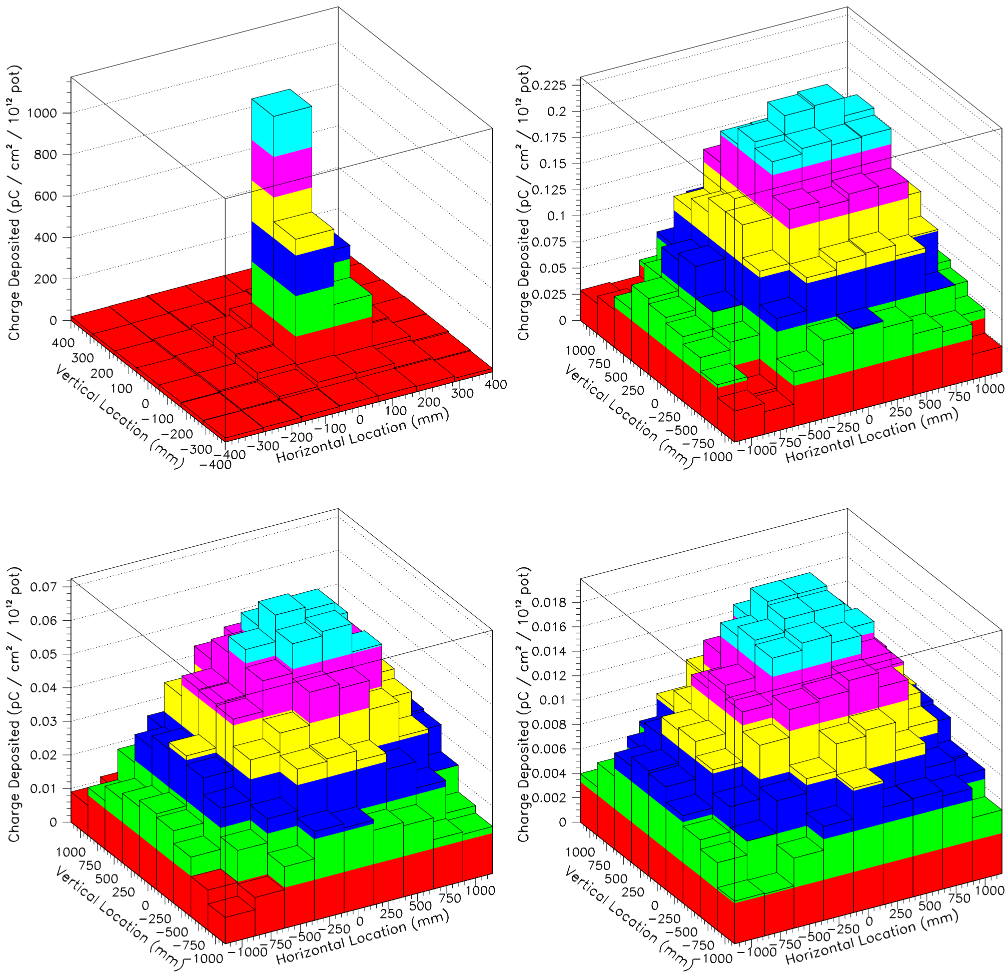

Figure F.8: Measured charge distributions at the monitoring stations with the proton beam in between the baffle and target in the pME position; this is with horn focusing. The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right. This data corresponds to that of Fig. 9.4.

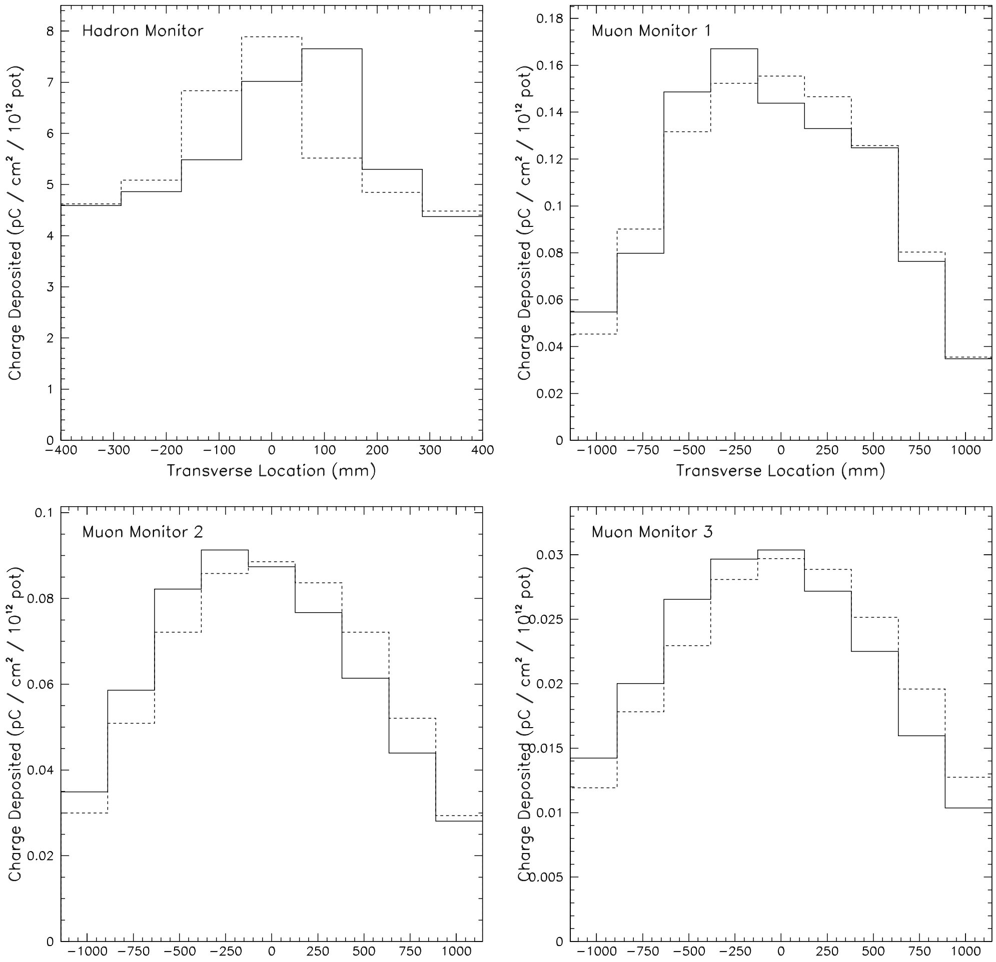

Figure F.9: Measured charge profiles at the monitoring stations with the proton beam hit the baffle in the pME position, 2.2 mm from edge; this is with horn focusing. The solid line is the horizontal profile, the dashed is the vertical. This data corresponds to that of Fig. 9.5.

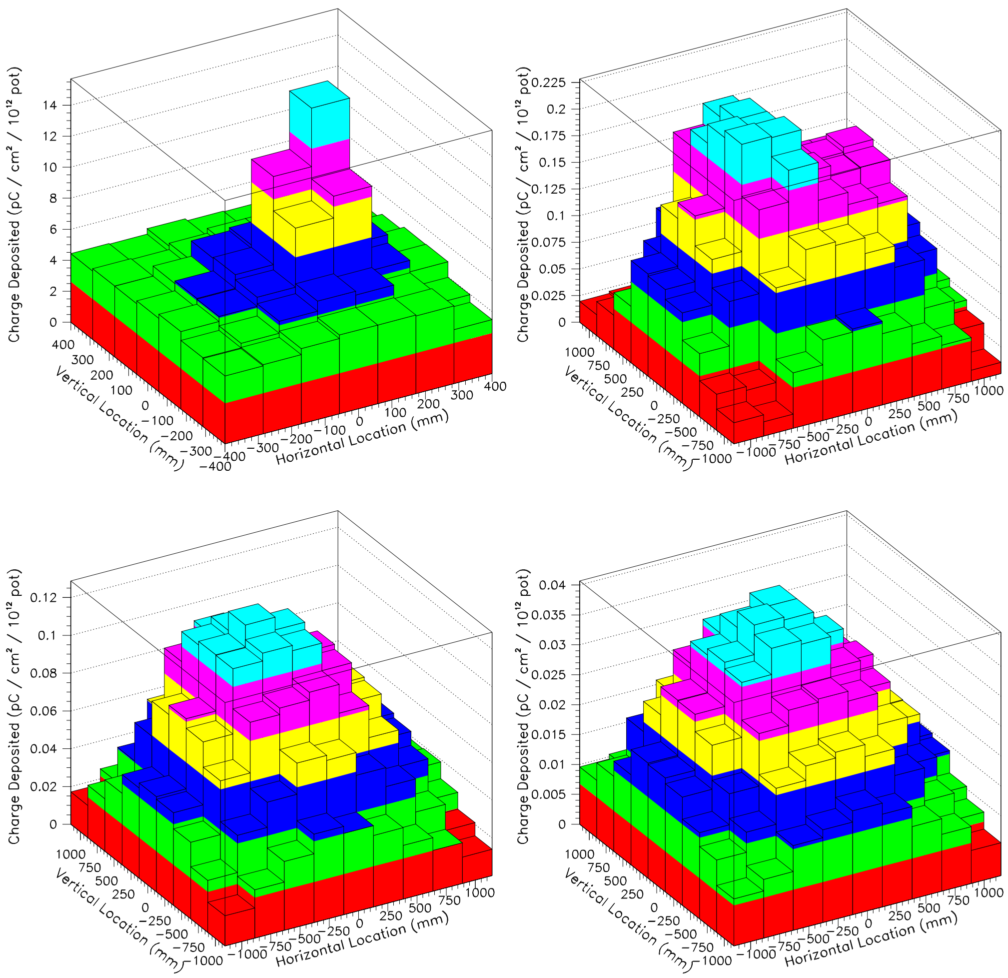

Figure F.10: Measured charge distributions at the monitoring stations with the proton beam hit the baffle in the pME position, 2.2 mm from edge; this is with horn focusing. The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right. This data corresponds to that of Fig. 9.5.

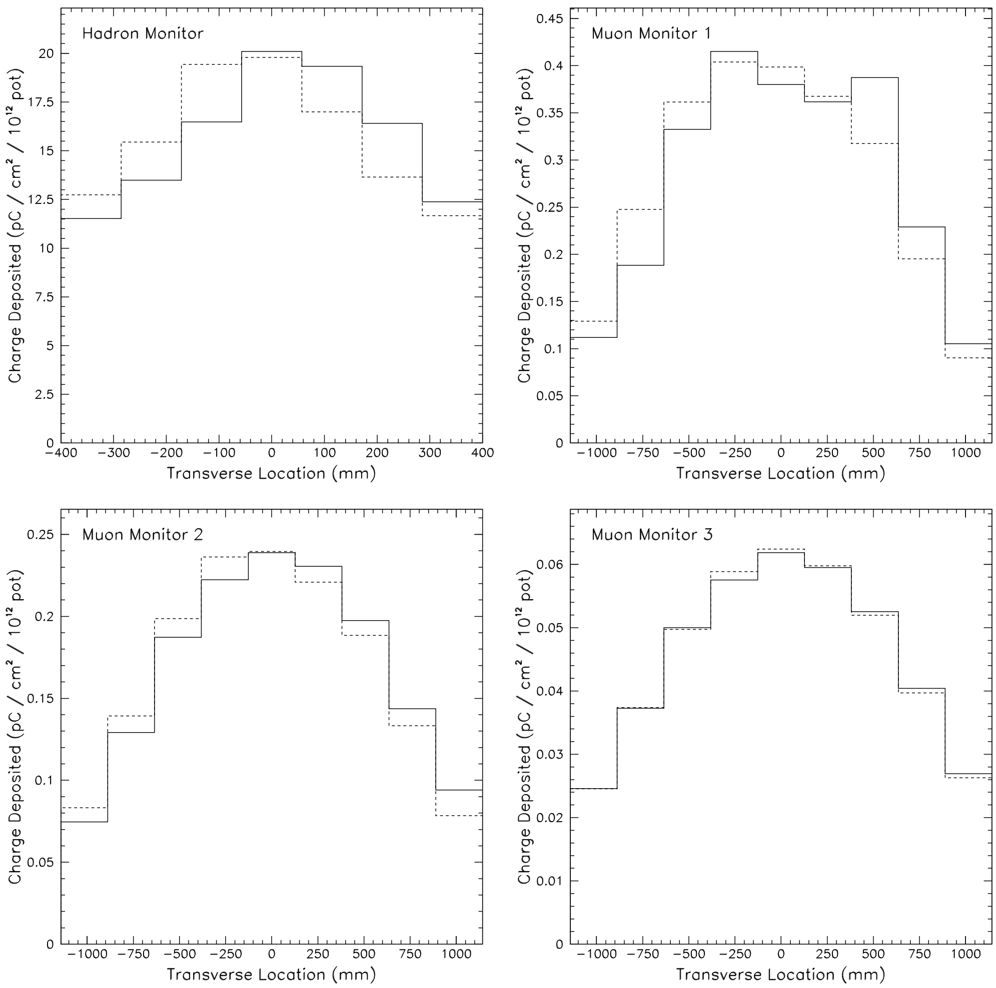

Figure F.11: Measured charge profiles at the monitoring stations with the beam in the partial High Energy (pHE) configuration. The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right. This data corresponds to that of Fig. 9.6.

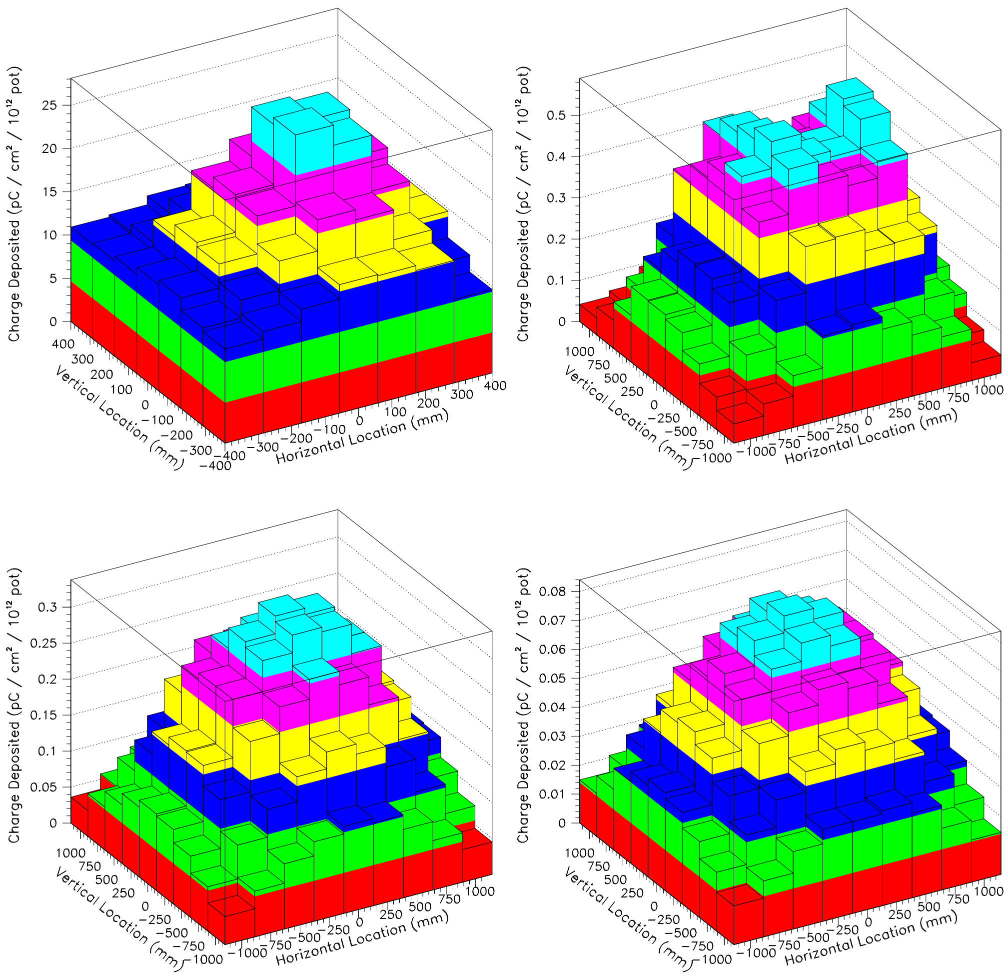

Figure F.12: Measured charge distributions at the monitoring stations with the beam in the partial High Energy (pHE) configuration. The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right. This data corresponds to that of Fig. 9.6.

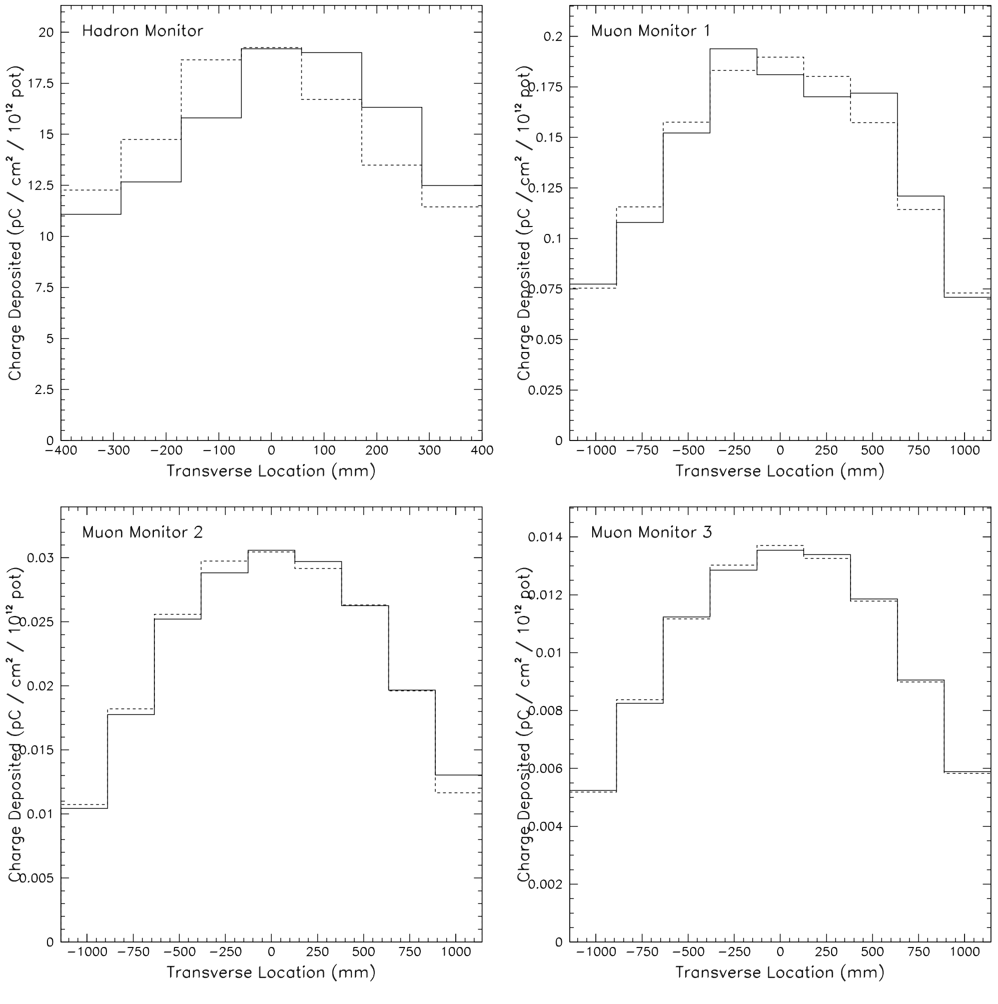

Figure F.13: Measured charge profiles at the monitoring stations with the beam in the standard running configuration (LE10). The solid line is the horizontal profile, the dashed is the vertical. This data corresponds to that of Fig. 9.7.

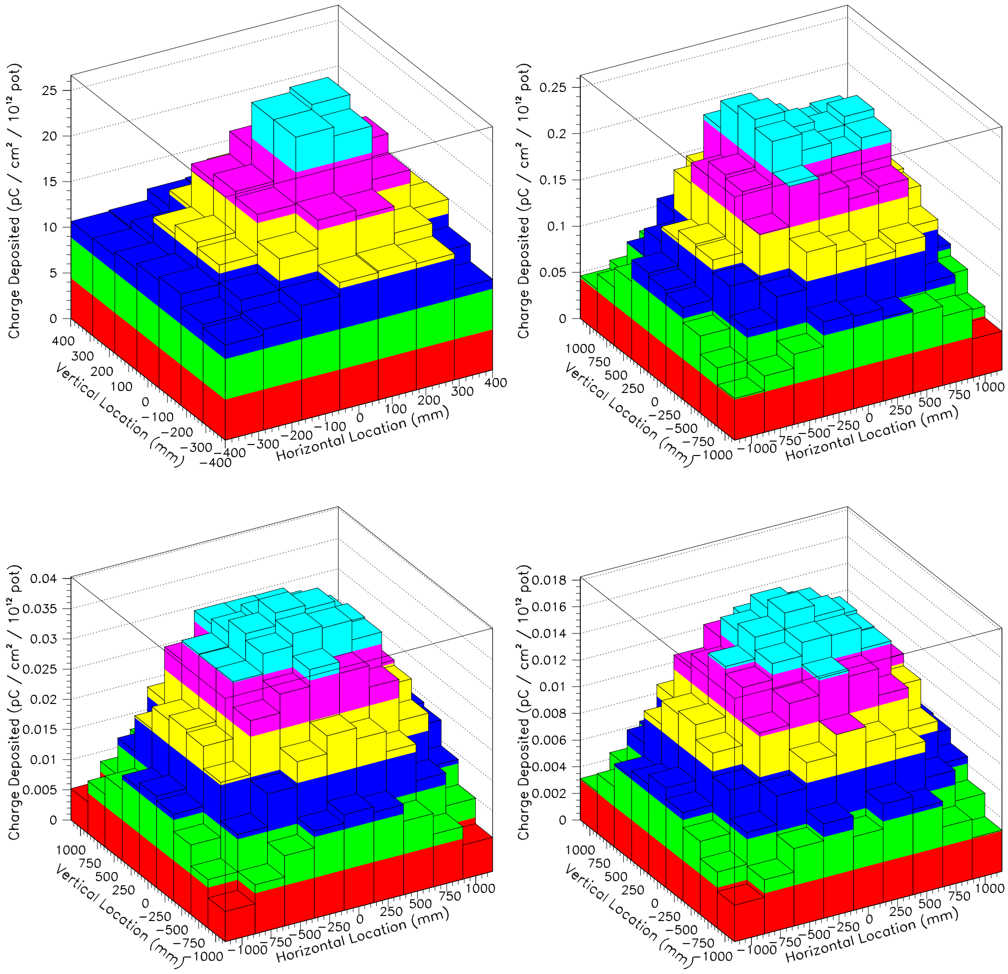

Figure F.14: Measured charge distributions at the monitoring stations with the beam in the standard running configuration (LE10). The upper-left is the Hadron Monitor; alcoves 1, 2 & 3 are the upper-right, lower-left, and lower-right. This data corresponds to that of Fig. 9.7.

Bibliography

[1] G. Danby *et al.*, "Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos", *Phys. Rev. Lett.* **9** (1962) pp. 36–44.

[2] P. Limon *et al.*, "A Sign-Selected Dichromatic Neutrino Beam", *Nucl. Inst. Meth.* **116** (1974) p. 317.

[3] B. T. Cleveland *et al.*, "Measurement of the solar electron neutrino flux with the Homestake chlorine detector", *Astrophys. J.* **496** (1998) pp. 505–526.

[4] J. N. Abdurashitov *et al.*, "Solar neutrino results from SAGE", *Phys. Atom. Nucl.* **63** (2000) pp. 943–950.

[5] M. Cribier, "Results of the whole GALLEX experiment", *Nucl. Phys. Proc. Suppl.* **70** (1999) pp. 284–291.

[6] S. Fukuda *et al.*, "Solar B-8 and he p neutrino measurements from 1258 days of Super-Kamiokande data", *Phys. Rev. Lett.* **86** (2001) pp. 5651–5655, `hep-ex/0103032`.

[7] S. Fukuda *et al.*, "Constraints on neutrino oscillations using 1258 days

of Super-Kamiokande solar neutrino data", *Phys. Rev. Lett.* **86** (2001) pp. 5656–5660, `hep-ex/0103033`.

[8] Q. R. Ahmad *et al.*, "Measurement of the charged current interactions produced by B-8 solar neutrinos at the Sudbury Neutrino Observatory", *Phys. Rev. Lett.* **87** (2001) p. 071301, `nucl-ex/0106015`.

[9] J. Hosaka *et al.*, "Solar neutrino measurements in Super-Kamiokande-I", `hep-ex/0508053`.

[10] Y. Fukuda *et al.*, "Evidence for oscillation of atmospheric neutrinos", *Phys. Rev. Lett.* **81** (1998) pp. 1562–1567.

[11] W. W. M. Allison *et al.*, "The atmospheric neutrino flavor ratio from a 3.9 fiducial kiloton-year exposure of Soudan 2", *Phys. Lett.* **B449** (1999) pp. 137–144, `hep-ex/9901024`.

[12] M. Ambrosio *et al.*, "Matter effects in upward-going muons and sterile neutrino oscillations", *Phys. Lett.* **B517** (2001) pp. 59–66, `hep-ex/0106049`.

[13] Y. Ashie *et al.*, "A measurement of atmospheric neutrino oscillation parameters by Super-Kamiokande I", *Phys. Rev.* **D71** (2005) p. 112005, `hep-ex/0501064`.

[14] V. D. Barger, J. G. Learned, S. Pakvasa, and T. J. Weiler, "Neutrino decay as an explanation of atmospheric neutrino observations", *Phys. Rev. Lett.* **82** (1999) pp. 2640–2643, `astro-ph/9810121`.

[15] R. Barbieri, P. Creminelli, and A. Strumia, "Neutrino oscillations from large extra dimensions", *Nucl. Phys.* **B585** (2000) pp. 28–44, `hep-ph/0002199`.

[16] E. Lisi, A. Marrone, and D. Montanino, "Probing possible decoherence effects in atmospheric neutrino oscillations", *Phys. Rev. Lett.* **85** (2000) pp. 1166–1169, `hep-ph/0002053`.

[17] H. W. Atherton *et al.*, "Precise Measurements of Particle Production by 400-GeV/c Protons on Beryllium Targets", CERN-80-07.

[18] D. S. Barton *et al.*, "Experimental Study of the A-Dependence of Inclusive Hadron Fragmentation", *Phys. Rev.* **D27** (1983) p. 2580.

[19] G. Ambrosini *et al.*, "Measurement of charged particle production from 450-GeV/c protons on beryllium", *Eur. Phys. J.* **C10** (1999) pp. 605–627.

[20] J. Hylen *et al.*, "The Hadron Hose: Continuous Toroidal Focusing for Conventional Neutrino Beams", *Nucl. Instrum. Meth.* **A498** (2001) pp. 29–51, `hep-ex/0210051`.

[21] A. Godley, "Description of the ACNET Devices that Control the Main Injector Inputs to the NuMI Beam Permit System", Fermilab Beams Document 1527, 2005.

[22] P. Prieto, "NuMI BPM Electrical Offset, BPM Polarity measurement and Intensity Scaling Calculation", Fermilab Beams Document 1534, 2005.

[23] D. Indurthy, S. Kopp, M. Proga, and Z. Pavlovich, "Profile monitor SEM's for the NuMI beam at Fermilab", `physics/0405021`.

[24] S. Kopp *et al.*, "Beam Test of a Segmented Foil SEM Grid", *Nucl. Instrum. Meth.* **A554** (2005) pp. 138–146, `physics/0507211`.

[25] S. Kopp *et al.*, "Segmented Foil SEM Grids at Fermilab", `physics/0506194`, Proceedings of the Particle Accelerator Conference (PAC05), Fermilab-Conf-05-092-AD.

[26] S. E. Kopp and Z. Pavlovic, "Studies of the Precision Proton Beam Profile Monitor SEMs", MINOS Internal Note 1068, 2005.

[27] A. Abramov *et al.*, "Dynamic Stress Calculations for ME and LE Targets and Results of Prototyping for the LE Target", MINOS Internal Note 675, 2000.

[28] K. Budal, "Charge Transport from Targets in Proton beams as a Means of Monitoring", CERN Yellow Report 67-17, 1967.

[29] M. Bonesini, A. Marchionni, F. Pietropaolo, and T. Tabarelli de Fatis, "On Particle Production for High Energy Neutrino Beams", *Eur. Phys. J.* **C20** (2001) pp. 13–27, `hep-ph/0101163`.

[30] S. Kopp, M. Kostin, and R. Zwaska, "Hadron Production Models Revisited", MINOS Internal Note 768, 2001.

[31] A. Fasso, A. Ferrari, P. R. Sala, and J. Ranft, "FLUKA: Status and Prospects for Hadronic Applications", Prepared for International Conference on Advance Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications (MC 2000), Losbon, Portugal, 23-26 Oct 2000.

[32] N. V. Mokhov, "The MARS Code System User's Guide Version 13(95)", FERMILAB-FN-0628.

[33] R. Brun *et al.*, "GEANT Detector Description and Simulation Tool", CERN Program Library Long Writeup W5013, October, 1994.

[34] S. Kopp, M. Kostin, and R. Zwaska, "Effect of Inert Material in the NuMI Target", MINOS Internal Note 1098, 2005.

[35] H. Meyer, "Physics of the MIPP Experiment", *Nucl. Phys. Proc. Suppl.* **142** (2005) pp. 453–458.

[36] S. van der Meer and B. de Raad, "Proposal for an Enhanced Neutrino Beam", CERN Report 61-3, 1961.

[37] A. G. Abramov *et al.*, "Beam Optics and Target Conceptual Designs for the NuMI Project", *Nucl. Instrum. Meth.* **A485** (2002) pp. 209–227.

[38] M. Kostin, S. Kopp, M. Messier, D. Harris, J. Hylen, and A. Para, "Proposal for Continuously-Variable Beam Energy", MINOS internal note 783, October, 2001.

[39] C. A. Ramm, "Neutrino Spectra from the Two-Body Decay of Relativistic Parents", in *The 1963 NPS Seminars: the Neutrino Experiment*, pp. 111–141. 1965, CERN Yellow Report 63-37.

[40] M. Messier *et al.*, "Neutrino Fluxes, Hadron Production, and the Hadronic Hose", MINOS Internal Note 700, 2000.

[41] L. Buckley-Geer *et al.*, "Effect of the hadronic Hose on MINOS Physics", MINOS Internal Note 701, 2000.

[42] K. Seiya *et al.*, "Status of Slip-Stacking at Fermilab Main Injector", To appear in the Proceedings of the Particle Accelerator Conference (PAC 05), 2005.

[43] C. W. Schmidt, "Review of negative hydrogen ion sources", Presented at 1990 Linear Accelerator Conf., Albuquerque, NM, Sep 10-14, 1990.

[44] C. W. Schmidt and C. D. Curtis, "A 50-mA Negative Hydrogen-Ion Source", *IEEE Trans. Nucl. Sci.* **NS-26** (1979), no. 3,.

[45] D. P. Moehs, "Studies on a Magnetron Source", *AIP Conf. Proc.* **639** (2002) pp. 115–120.

[46] J. D. Cockroft and E. T. S. Walton, "Experiments with High Velocity Positive Ions I: Further Developments in the Method of Obtaining High Velocity Ions", *Royal Sociey of London Proceedings Series A* **136** (1932) pp. 619–630.

[47] L. W. Alvarez, "The Design of a Proton Linear Accelerator", *Physical Review* **70** (1946) pp. 799–800.

[48] E. A. Knapp, B. C. Knapp, and J. M. Potter, "Standing Wave High Energy Linear Accelerator Structures", *Review of Scientific Instruments* **39** (1968), no. 7, pp. 979–991.

[49] E. L. Ginzton, W. W. Hansen, and W. R. Kennedy, "A Linear Electron Accelerator", *Review of Scientific Instruments* **19** (1948) pp. 89–108.

[50] G. Caryotakis, "High-Power Klystrons: Theory and Practice at the Stanford Linear Accelerator Center. Part I: Theory and Design", 2004. SLAC-PUB-10620.

[51] E. L. Hubbard, "Booster Synchrotron", 1973. FERMILAB-TM-0405.

[52] B. Worthel *et al.*, "Booster Rookie Book", Fermilab Beams Document 1022, 1998.

[53] C. M. Bhat, "The Main Injector Facility: A New Tool for High Energy Physics at Fermilab", *Int. J. Mod. Phys.* (2001) pp. 1187–1189.

[54] H. Hinterberger *et al.*, "Quadrupole Magnets of the NAL Main Accelerator", Proceedings of the 1971 Particle Accelerator Conference (PAC71).

[55] D. Boussard and Y. Mizumachi, "Production of Beams with High Line Density by Azimuthal Combination of Bunches in a Synchrotron", *IEEE Trans. Nucl. Sci.* **26** (1979) pp. 3626–3625.

[56] J. Dey and I. Kourbanis, "53-MHz Beam Loading Compensation for Slip-Stacking in the Fermilab Main Injector", To appear in the proceedings of the Particle Accelerator Conference (PAC 05), 2005.

[57] S. van der Meer, "Stochastic cooling and the Accumulation of Anti-Protons", *Rev. Mod. Phys.* **57** (1985) pp. 689–697.

[58] G. I. Budker *et al.*, "Experiments on Electron Cooling", *IEEE Trans. Nucl. Sci.* **22** (1975) pp. 2093–2097.

[59] W. Blokland and J. Steimel, "Main injector synchronous timing system", To be published in the proceedings of 8th Beam Instrumentation Workshop (BIW 98), Stanford, CA, 4-7 May 1998.

[60] S. Saritepe and G. Annala, "Tevatron injection timing", FERMILAB-TM-1792.

[61] M. Popovic, "A Proposed H- Injection System for the Fermilab Booster", Fermilab Beams Document 1784, 2005.

[62] J. R. Lackey *et al.*, "New Pulsed Orbit Bump Magnets for the Fermilab Booster Synchrotron", To appear in the Proceedings of Particle Accelerator Conference (PAC 05), Knoxville, Tennessee, 16-20 May 2005.

[63] W. Chou, A. Drozhdin, P. Lucas, and F. Ostiguy, "Fermilab Booster Modeling and Space Charge Study", presented at Particle Accelerator Conference (PAC 03), Portland, Oregon, 12-16 May 2003.

[64] G. I. Budker and G. I. Dimov, "On the Charge Exchange Injection of Protons into Ring Accelerators", International Conference 1962, CONF-114, USAEC TID-4500, pp. 1372-1377.

[65] C. Hojvat *et al.*, "The Multiturn Charge Exchange Injection System for the Fermilab Booster Accelerator", *IEEE Trans. Nucl. Sci.* **26** (1979) pp. 3149–3151.

[66] V. S. Kashikhin *et al.*, "A New Correction Magnet Package for the Fermilab Booster Synchrotron", To appear in the Proceedings of Particle Accelerator Conference (PAC 05), Knoxville, Tennessee, 16-20 May 2005.

[67] P. Kasper, "Booster Radiation Issues", Fermilab Beams Document 865, 2003.

[68] J. Amundson *et al.*, "Calibration of the FNAL Booster Ionization Profle Monitor", *Phys. Rev. ST Accel. Beams* **6** (2003) p. 102801.

[69] L. C. Teng, "Matching and Synchronization During Transfer from Booster to Main Ring at 8-GeV", 1971. FERMILAB-TM-0281.

[70] W. A. Pellico and R. C. Webber, "RF Cogging in the FNAL Booster Accelerator", Proceedings of IEEE Particle Accelerator Conference (PAC 99), New York, 1999.

[71] M. J. S. Smith, *Application-Specific Integrated Circuits.* Addison-Wesley, 1997. available online: http://www.edacafe.com/books/ASIC/ASICs.php.

[72] K. Meisner, "RF Synchronous Transfer into Specific Buckets between Fermilab Main Ring and Tevatron Accelerators", *Proceedings of the 1993 Particle Accelerator Conference* (1993) p. 2522.

[73] M. T. Ronan *et al.*, "Timing and RF Synchronization for Filling PEP/SPEAR with the SLC Damping Rings", Presented at IEEE Particle Accelerator Conf., Chicago, Ill., Mar 20-23, 1989.

[74] J. Brennan *et al.*, "Synthesizer Controlled Beam Transfer from the AGS to RHIC", Presented at IEEE Particle Accelerator Conference (PAC2001), Chicago, Illinois, 18-22 Jun 2001.

[75] R. Zwaska, S. Kopp, W. Pellico, and R. Webber, "Synchronization of the Fermilab Booster and Main Injector for Multiple Batch Injection", `hep-ph/0101163`, Proceecedings of 2004 European Particle Accelerator Conference (EPAC04), Lucerne, Switzerland.

[76] R. Zwaska, S. Kopp, and W. Pellico, "Cycle-to-Cycle Extraction Synchronization of the Fermilab Booster for Multiple Batch Injection to the Main Injector", To appear in the proceedings of Particle Accelerator Conference (PAC 05), Knoxville, Tennessee, 2005.

[77] R. Tomlin, "Ion beam notcher using a laser", Prepared for IEEE Particle Accelerator Conference (PAC 2001), Chicago, Illinois, 18-22 Jun 2001.

[78] D. Moehs, “Sub microsecond notching of a negative hydrogen beam at low energy utilizing a magnetron ion source with a split extractor”, *AIP Conf. Proc.* **763** (2005) pp. 189–193.

[79] D. P. Moehs, “Sub-microsecond beam notching at low energy”, Presented at 11th International Conference on Ion Sources (ICIS05), Caen, France, 12-16 Sep 2005.

[80] E. Prebys *et al.*, “Increasing the Intensity of the Fermilab Booster”, presented at Particle Accelerator Conference (PAC 03), Portland, Oregon, 12-16 May 2003.

[81] D. McGinnis, “A 2 MegaWatt Multi-Stage Proton Accumulator”, Fermilab Beams Document 1782, 2005.

[82] D. Harris *et al.*, “Muon Monitoring Requirements for the NuMI Beamline”, MINOS Internal Note 756, 2001.

[83] D. Harris *et al.*, “Downstream Hadron Monitor Requirements for the NuMI Beamline”, MINOS Internal Note 785, 2001.

[84] R. Burns *et al.*, “Determination of the neutrino flux”, in *CERN Informal Conference on Experimental Neutrino Physics*, pp. 97–109. 1965, CERN Yellow Report 65-32.

[85] G. Plass, K. M. Vahlbruck, and H. W. Wachsmuth, “The determination of the spectra of high-energy neutrino parents”, in *CERN Informal Conference on Experimental Neutrino Physics*, pp. 121–126. 1965, CERN Yellow Report 65-32.

[86] T. Hart, "Kaon Monitoring Using the MiniBooNE Little Muon Counter", 2005. Proceedings of the Particle Accelerator Conference (PAC05).

[87] B. C. Barish *et al.*, "Exploratory Study of High-Energy Neutrino Interactions", *Phys. Rev. Lett.* **31** (1973) pp. 565–568.

[88] H. Wachsmuth, "Neutrino and Muon Fluxes in the CERN 400 GeV Proton Beam Dump Experiments", CERN/EP Report 79-125, 1979.

[89] R. Blair *et al.*, "Monitoring and Calibration System for Neutrino Flux Measurement in a High-Energy Dichromatic Beam", *Nucl. Instrum. Meth.* **A226** (1984) p. 281.

[90] B. C. Barish *et al.*, "Charged Current Neutrino and Anti-Neutrino Cross-Section Results from the CITFR Experiment", Presented at the 3rd Int. Conf. on New Results in High Energy Physics, Nashville, tenn., Mar 6-8, 1979, FERMILAB-conf-78/46-EXP.

[91] E. H. M. Heijne, "Muon flux measurement with silicon detectors in the CERN neutrino beams", CERN Yellow Report 83-06, 1983.

[92] J. B. M. Pattison, "The 1967 Neutrino installation", in *CERN Neutrino Meeting*, pp. 13–32. 1969, CERN Yellow Report 69-28.

[93] D. Bloess *et al.*, "Determination of the Neutrino Spectrum in the CERN 1967 Neutrino Experiment", *Nucl. Inst. Meth.* **91** (1971) pp. 605–612.

[94] H. W. Wachsmuth, "The neutrino spectrum for the CERN 1967 neutrino experiment", in *CERN Neutrino Meeting*, pp. 33–46. 1969, CERN Yellow Report 69-28.

[95] P. Astier *et al.*, "Prediction of neutrino fluxes in the NOMAD experiment", *Nucl. Instrum. Meth.* **A515** (2003) pp. 800–828, `hep-ex/0306022`.

[96] V. B. Anikeev *et al.*, "Total Cross-Section Measurements for Muon-Neutrino, Anti- Muon-Neutrino Interactions in 3-GeV - 30-GeV Energy Range with IHEP-JINR Neutrino Detector", *Z. Phys.* **C70** (1996) pp. 39–46.

[97] A. P. Bugorsky *et al.*, "Muon Flux Measuring System for Neutrino Experiments at the IHEP Accelerator", *Nucl. Inst. Meth.* **146** (1977) pp. 367–371.

[98] C. Y. Chi *et al.*, "Segmented Ionization Chambers for High Intensity Beam Measurements", *Nucl. Instrum. Meth.* **A281** (1989) pp. 448–452.

[99] J. E. Hill, "Beam monitoring and modeling for the K2K long-baseline neutrino oscillation experiment", *Int. J. Mod. Phys.* **A16S1B** (2001) pp. 758–760.

[100] **K2K** Collaboration, Y. Oyama, "Current status of the K2K experiment", 2001.

[101] G. Acquistapace *et al.*, "The CERN neutrino beam to Gran Sasso (NGS): Conceptual technical design", 1998. CERN98-02.

[102] L. Casagrande *et al.*, "The alignment of the CERN West Area neutrino facility", CERN Yellow Report 96-06, 1996.

[103] J. C. Dusseux *et al.*, "The CERN Magnetic Horn (1971) and its Remote-Handling System", CERN Yellow Report 72-11, 1972.

[104] A. K. Mann, *Neutrino Interactions with Electrons and Protons.* American Institute of Physics Press, 1993.

[105] E. D. Zimmerman, “BooNE Decay Region and Absorbers”, Presented at the 2nd International Worksop on Neutrino Beams and Instrumentation (NBI 2002), CERN, Geneva, Switzerland, March, 2002.

[106] A. Auchincloss, *A Measurement of the Total Cross Section for Neutrino Nucleon Interactions.* PhD thesis, Columbia University, 1987.

[107] R. Keisler, “Neutron Induced Ionization in the NuMI Beam Monitors”, University of Texas undergraduate thesis, 2005.

[108] D. Harris and N. Grossman, “Radiation Dose Calculations for Monitoring System”, MINOS Internal Note 844, 2002.

[109] M. Diwan *et al.*, “The Hadron Monitor”, MINOS Internal Note 846, 2002.

[110] D. Naples *et al.*, “Pad Ionization Chamber Design”, MINOS Internal Note 769, 2001.

[111] E. Fuller *et al.*, “Studies of Cross Talk Between Ion Chambers”, MINOS Internal Note 823, 2002.

[112] D. Indurthy and S. Kopp, “Installation and Survey of the Hadron Monitor Support Structure”, MINOS Internal Note 1039, 2004.

[113] D. Naples, A. Erwin, and D. Northacker, “Preliminary Design of a Radiation-Hard Pad Ionization Chamber Array”, MINOS Internal Note 850, 2002.

[114] M. Diwan *et al.*, "The Muon Monitor", MINOS Internal Note 847, 2002.

[115] D. Indurthy *et al.*, "Calibration of the NuMI Beam Monitors", MINOS Internal Note 1084, 2005.

[116] A. R. Erwin, D. Glover, and H. Ping, "Modification of Droege HV Supplies for Ethernet Control of the Beam Monitoring Ion Chambers", MINOS Internal Note 942, 2003.

[117] D. H. Wilkinson, *Ionization Chambers and Counters.* Cambridge University Press, 1950.

[118] G. F. Knoll, *Radiation Detection and Measurement.* John Wiley & Sons, Inc., New York, Third Ed., 2000.

[119] S. Eidelman *et al.*, "Review of Particle Physics", *Physics Letters B* **592** (2004) *Passage of Particles through Matter.*

[120] ICRU, "Average Energy Required to Produce an Ion Pair", Tech. Rep. 31, International Commission on Radiation Units and Measurements, Wahington D. C., 1979.

[121] J. E. Parks *et al.*, "Ionization of the noble gases by protons: Jesse effects as a function of pressure", *Journal of Chemical Physics* **57** (1972) p. 5467.

[122] T. E. Bortner and G. S. Hurst, "Energy per ion pair for 5 Mev alpha particles in helium", *Physical Review* **90** (1953) p. 160.

[123] W. P. Jesse and J. Sadauskis, "Alpha-particle ionization in pure gases and the average energy to make an ion pair", *Physical Review* **90** (1953) p. 1120.

[124] W. P. Jesse and J. Sadauskis, "Ionization in pure gases and the average energy to make an ion pair for alpha and beta particles", *Physical Review* (1955) p. 1668.

[125] J. Weiss and W. Bernstein, "Energy required to produce one ion pair for several gases", *Physical Review* **98** (1956) p. 1828.

[126] R. Cooper and R. M. Mooring, "Ionization current measurements and the energy required to form an ion pair in γ-ray irradiated gases", *Austrian Journal of Chemistry* **21** (1968) p. 2417.

[127] G. Stetter, "Die Ionisation einzelner Alpha-Strahien in Verschidenen Gasen", *Zeitschrift für Physik* **120** (1943) p. 639.

[128] L. H. Gray *Proceedings of the Cambridge Philisphical Society* **40** (1944) p. 72.

[129] B. B. Rossi and H. H. Staub, *Ionization Chambers and Counters.* McGraw-Hill Book Company, Inc., New York, 1949.

[130] L. G. Christophorou, *Atomic and Molecular Radiation Physics.* Wiley–Interscience, 1971.

[131] L. B. Loeb, *Basic Processes of Gaseous Electronics.* University of California Press, Berkeley, California, 1955.

[132] J. A. Hornbeck and J. P. Molnar, "Mass Spectrometric Studies of Molecular Ions in the Noble Gases", *Physical Review* **84** (1951) p. 621.

[133] A. V. Phelps and S. C. Brown, “Positive Ions in the Afterglow of a Low Pressure Helium Discharge”, *Physical Review* **86** (1952) p. 102.

[134] L. M. Chanin and M. A. Biondi, “Mobilites of Atomic and Molecular Ions in the Noble Gases”, *Physical Review* **94** (1954) p. 910.

[135] J. A. Hornbeck, “The Mobilities of Molecular and Atomic Rare Gas Ions in the Present Gases: Helium, Neon, and Argon”, *Physical Review* **80** (1950) p. 297.

[136] J. H. Mitchell and K. E. W. Ridler, “The Speed of Positive Ions in Nitrogen”, *Proceedings of the Royal Society* **A146** (1934) p. 911.

[137] H. S. W. Massey and C. B. Mohr, “Free Paths and Transport Phenomena in Gases and the Quantum Theory of Collisions; Part II: The Determination of the Laws of Force between Atoms and Melecules”, *Proceedings of the Royal Society* **A144** (1934) p. 188.

[138] J. L. Pack and A. V. Phelps, “Drift Velocities of Slow Electrons in Helium, Neon, Argon, Hydrogen, and Nitrogen”, *Physical Review* **121** (1961) p. 798.

[139] R. W. Crompton *et al.*, “The Momentum Transfer Cross Section for Electrons in Helium”, *Australian Journal Of Physics* **20** (1967) p. 369.

[140] J. Anderson, “Hall Effects and Electron Drift Velocities in the Plasma and the Positive Column”, *Physics of Fluids* **7** (1964) p. 1517.

[141] A. V. Phelps, J. L. Pack, and L. S. Frost, “Drift Velocity of Electrons in Helium”, *Physical Review* **117** (1960) p. 470.

[142] Stern, "Drift Velocities of Electrons in Helium at High E/p", in *Proceedings of the Sixth International Conference on Ionization Phenomena in Gases*, P. Hubert and E. Cremieu-Alcan, eds., Vol. 1, p. 331. Paris, 1963.

[143] J. Dutton, "A Survey of Electron Swarm Data", *Journal of Physical and Chemical Reference Data* **4** (1975) p. 577.

[144] A. von Engel, *Handbuch der Physik*, Vol. 21. Springer Verlag, Berlin, 1956.

[145] S. Brown, *Basic Data of Plasma Physics: The Fundamental Data on Electrical Discharges in Gases.* AIP Press, New York, 1994.

[146] S. H. Dunlop *Nature* **164** (1949) p. 452.

[147] L. M. Chanin and G. D. Rork, "Experimental Determination of the First Townsend Coefficient in Helium", *Physical Review* **133** (1964) p. 1005.

[148] D. K. Davies *et al.*, "Primary Ionization Coefficient of Helium", *Proceedings of the Physical Society* **80** (1962) p. 898.

[149] S. J. B. Corrigan and A. von Engel, "The Excitation of Helium by Electrons of Low Energy", *Proceedings of the Physical Society of London* **72** (1958) p. 786.

[150] I. M. Bortnik, "Investigation of the Positive Corona Discharge in Helium", *Soviet Physics, Technical Physics* **9** (1965) p. 1683.

[151] M. H. Hughes, "Electron Energy Distribution Function and Transport Coefficients in Helium and Neon", *Journal of Physics B* **3** (1970) p. 1544.

[152] T. Itoh *et al.*, "Monte Carlo Calculations of Motion of Electrons in Helium", *Journal of the Physical Society of Japan* **15** (1960) p. 1675.

[153] M. A. Biondi and S. C. Brown, "Measurement of Electron-Ion Recombination", *Physical Review* **76** (1949) p. 1697.

[154] H. S. W. Massey and E. H. S. Burhop, *Electronic and Ionic Impact Phenomena.* Clarendon Press, Oxford, 1952.

[155] N. E. Bradbury, "Electron Attachment and Negative Ion Formation In Oxygen", *Physical Review* **44** (1933) p. 883.

[156] R. H. Healy and J. W. Reed, *The Behaviour of Slow Electrons in Gases.* Amalgamated Wireless, Ltd., Australasia, 1941.

[157] F. Bloche and N. E. Bradbury, "On the Mechanism of Unimolecular Electron Capture", *Physical Review* **48** (1935) p. 689.

[158] R. Zwaska *et al.*, "Beam Tests of NuMI Ionization Chambers at the Fermilab Booster Accelerator", MINOS Internal Note 810, 2001.

[159] R. M. Zwaska *et al.*, "Beam tests of ionization chambers for the NuMI neutrino beam", *IEEE Trans. Nucl. Sci.* **50** (2003) pp. 1129–1135, `hep-ex/0212011`.

[160] S. Palestini *et al.*, "Space charge in ionization detectors and the NA48 electromagnetic calorimeter", *Nucl. Instrum. Meth.* **A421** (1999) pp. 75–89.

[161] S. J. Barish *et al.*, "Study of Neutrino Interactions in Hydrogen and Deuterium. 1. Description of the experiment and Study of the Reaction $\nu + d \to \mu^- + p + p_s$", *Phys. Rev.* **D16** (1977) p. 3103.

[162] Fermilab, *NuMI Technical Design Handbook*, 2002-4. URL `http://www-numi.fnal.gov/numwork/tdh/tdh_index.html`

[163] A. Para, "Locating the Neutrino Beam Elements", MINOS Internal Note 796, 2001.

[164] Y. Mori *et al.*, "Syncrotron Design Issues of the Japanese Hadron Project", Presented at 1996 European Particle Accelerator Conf.

[165] D. C. Carey, K. L. Brown, and F. Rothacker, "Third order TRANSPORT with MAD input: A computer program for designing charged particle beam transport systems", 1998. SLAC-R-0530.

[166] L. P. Meissner, "Fortran Programming language: recent developments and a view of the future", *SIGPLAN Fortran Forum* **1** (1982), no. 1, pp. 3–8.

[167] J. Shiers, "CERNLIB-CERN Program Library Short Writeups", CERN, Geneva, 1996.

[168] R. Bock *et al.*, "PAW – toward a physics analysis workstation", *Comp. Phys. Commun.* **45** (1987) pp. 181–190.

[169] H. Raether, "The development of electron avalanche in a spark channel (from observations in a cloud chamber)", *Zeitschrift fur Physik* **112** (1939) p. 464.

Vita

Robert Miles Zwaska was born in Milwaukee, Wisconsin on August 5, 1977 and grew up in the village of Brown Deer, Wisconsin under the direction of his parents: Robert Frank Zwaska and Catherine Zwaska. He attended Marquette University High School in Milwaukee, Wisconsin, graduated in 1995, and continued on to the University of Notre Dame. He earned the degree of Bachelor of Science from the University of Notre Dame in the spring 1999. He started as a graduate student in the Department of Physics at the University of Texas at Austin in the fall of 1999, and in the spring of 2000 began working for Professor Sacha Kopp in the High Energy Physics Laboratory. While at the University of Texas, he earned the Continuing University Tuition Fellowship in 2002 and travel grants to attend the 2001 Tamura International School, the 2001 and 2005 Particle Accelerator Conferences, and the United States Particle Accelerator School in the summer of 2001, winter of 2002, and winter of 2004.

Permanent Address: 416 N. River St.
Batavia, IL 60510

This dissertation was typed by the author.