Main Injector Rookie Book

Table of Contents

Main Injector

Terry Asher wrote the Main Injector Rookie Book.
Steve Baginski and Terry Asher made all the Main Injector drawings.

Figures and Illustrations

The full-page illustrations listed below are either embedded in the text or added to the end of each chapter.

Chapter 1: Modes of Operation

Chapter 2: Magnets and the Lattice

The following maps (2-20 through 2-25) show the ring-wide distribution of the magnets, along with the associated CAMAC ramp cards and power supply connections:

Chapter 3: Power Supplies

Chapter 4: Water in the Main Injector

Chapter 5: Vacuum

Chapter 7: Beam Transport Lines

Main Injector

Ancient History

The Main Injector accelerator replaced the venerable Main Ring, once the most powerful accelerator in the world. The Main Ring, originally designed in the 1960's as a 200 GeV machine, was housed in a four-mile circular tunnel. Two "upstream" machines, the Linac and Booster, accelerated protons to 8 GeV, which was the minimum energy acceptable to Main Ring. After acceleration to 200 GeV in the Main Ring, the beam was extracted, simultaneously, to several experiments external to the ring. Many successive improvements brought the energy of the Main Ring—and the experiments—to 500 GeV.

Notes:

Introduction and History

In the early 1980s, the Tevatron, an accelerator using superconducting magnets, was built and installed in the space directly below the Main Ring. The Main Ring was now relegated to handing over its protons to an even more powerful accelerator. Beam was transferred from the Main Ring to the Tevatron at 150 GeV, and the experiments received protons from the Tevatron at 800 GeV. The placement of both accelerators in the same tunnel saved millions of dollars in civil construction costs.

In the mid-1980s, the business of experimental particle physics took a step forward with the development of colliding beams, in which particles (at Fermilab, protons) would circulate in one direction and collide with corresponding antiparticles (antiprotons) traveling in the opposite direction. The center-of-mass energy released in collisions was far greater than anything that extraction to stationary targets could provide.

Antiprotons, unlike protons, are not yet found in gas bottles, but are created when high-energy protons strike an appropriate target. The Main Ring acquired an additional role to provide 120 GeV protons to the antiproton production target. The antiprotons were created at an energy of 8 GeV; they were then stored in the Accumulator Ring of the Antiproton Source. The 8 GeV antiprotons could then be injected into the Main Ring, accelerated to 150 GeV, and sent on to the Tevatron.

In time, it was realized that having the Main Ring and the Tevatron in the same tunnel was more of a liability than an advantage. The construction of two gigantic particle detectors internal to the Tevatron forced a redesign of the Main Ring, since overpasses had to be built to carry beam up over the experiments. However, the vertical distortion impaired Main Ring performance, even as the losses leaking from Main Ring continued to degrade the data being taken by the detectors. Moreover, stray magnetic fields from the Tevatron interfered with those of the Main Ring and hurt the beam quality even further.

Main Injector

It was eventually decided that the functions performed by the Main Ring would be better served by constructing a new accelerator outside the four-mile tunnel. This machine, to be called the Main Injector, would be a smaller, faster, and more efficient version of the Main Ring, and it would not stand in the way of the Tevatron. It would do everything Main Ring had done: accepting 8 GeV particles from the Booster or the Antiproton Source and accelerating them to 150 GeV, and accelerating 120 GeV protons at a rapid rate for antiproton production. In addition, the original role of the Main Ring, that of extracting beam directly to the experiments, would be restored, this time at 120 GeV.

Perhaps out of fear that building an accelerator with such a multitude of functions would not be sufficiently challenging, it was decided to add a second machine to the Main Injector tunnel. It would be known as the Recycler, because it would salvage scarce antiprotons from the Tevatron once data taking was complete. It would be built of non-powered permanent magnets, and would store the antiprotons at 8 GeV.

Intention

The intention of this book is to describe the Main Injector and Recycler and at a level which will be operationally useful.

The first chapter, "Modes of Operation," will outline the geography, names, and places that pertain to the Main Injector. It will also describe, in a superficial way, the specific roles that Main Injector and Recycler will play: acceptance of 8 GeV beam from the Booster or Antiproton Source, acceleration of 120 GeV beam for antiproton production, extraction to the Fixed Target experiments, 150 GeV extraction into the Tevatron, and deceleration of antiprotons destined for the Recycler. The chapter will be an introduction to basic concepts and will provide context for the more detailed information to follow.

Chapter 2, "Magnets and the Lattice," describes the many kinds of magnets used in the Main Injector ring and how they work together to get the beam to circulate. Chapter 3 deals with the power supplies for the magnets. Chapter 4 describes the water systems used for cooling the magnets, and Chapter 5 is about the vacuum systems, which clear the air molecules from the path that the beam must travel.

There will be more chapters released in the near future. Chapter 6, "Beam Diagnostics," will explain how we can look at the beam and see what it is actually doing. Chapter 7, "Beam Transport Lines," describes how beam is transferred from one accelerator to another. Incomplete versions of these two chapters are available in the Main Control Room copies of the rookie book, but are not deemed ready for mass release.

Future chapters will tackle the "RF" (radio frequency) systems, which accelerate the beam and perform complex manipulations of the beam during certain modes of operation. The chapter on "Controls" will examine the communication of command and data information between the Main Injector and the Main Control Room. A chapter on "Utilities" will be concerned with the "mundane" aspects of the Main Injector infrastructure such as electrical power, ventilation, and water drainage, the failure of which can bring down an accelerator just as easily as the most complex instrumentation.

What To Look For

It is helpful, when reading this book, to recognize that the many complicated systems described are all directed toward a single goal, to maximize the production of the rare, exotic, and transient particles that are the domain of high-energy physics. From the standpoint of accelerator performance, that goal can usually be interpreted as the need to optimize intensity and/or luminosity. Of course, there are issues of beam quality as well.

Intensity is defined as the number of particles reaching the desired target. In Fixed Target mode, where beam is extracted to stationary targets, intensity is usually the bottom line.

In Collider Mode, where proton and antiproton beams collide with each other, luminosity is usually the most important goal. Luminosity is a measure of the collision rate, which depends not only on intensity but also on beam size and quality.

Experience

I hope that this book will assist the reader in attaining these goals. However, it cannot be read in isolation. For those responsible for running the Main Injector, it is not a substitute for experience. The material in the book is only a tiny fraction of what is required from a competent operator or user of the machine. When devices are discussed, find them in the field. Go to the service buildings and take advantage of access time in the tunnels. Learn your way around the applications pages and seize opportunities to troubleshoot and tune the accelerators. This book is only a first step.

Chapter 1: Modes of Operation

> A batch, a batch,
> A Booster full of buckets
>
> Old English nursery rhyme

Geography

The Main Injector ring is located in the southwest corner of the Fermilab site, adjacent to the Tevatron on one side and the Prairie Path bicycle trail on the opposite side. It is roughly elliptical in shape (Fig. 1-1). Beam transport lines connect the Main Injector not only to the Tevatron, but also to the Booster, Antiproton Source, Switchyard, and the Recycler Ring.

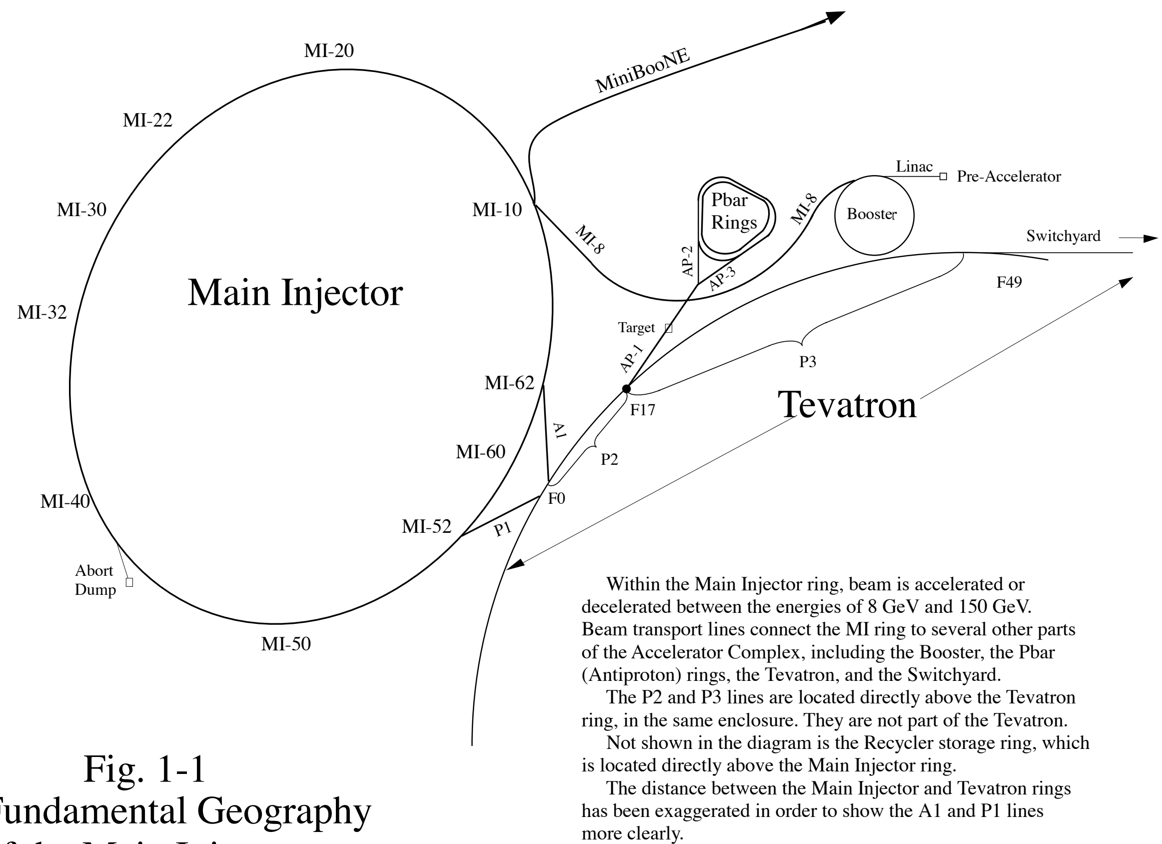

Fig. 1-1
Fundamental Geography
of the Main Injector

(The Recycler, not shown in Fig. 1-1, is located directly above the Main Injector in a shared tunnel.)

Main Injector

The Main Injector has such a central role in linking these areas that the other accelerators would be useless without it.

The Main Injector can accelerate or decelerate particles between the energies of 8 GeV and 150 GeV. The sources of the particles and their final destinations are quite variable and depend on the "mode of operation," that is, what the Main Injector is being used for at the time. This chapter will describe the modes of operation, emphasizing how the accelerators link together and the path the beam takes during each type of operation. Most of the technical details will be deferred to later chapters so that readers have a chance to grasp the big picture first.

Main Injector, like all of the accelerators, has a numbering scheme that defines locations around the ring. In principle, you can locate every component in the tunnel through this convention. Main Injector can be thought of as having six major "mileposts" at more or less equal intervals around the ring: MI-10, MI-20, MI-30, MI-40, MI-50, and MI-60. MI-10 is the point where protons enter the ring from the Booster. The protons travel counterclockwise, and so the numbers increase in a counterclockwise direction. MI-60 is the region adjacent to the Tevatron.

At each of these major locations there is a service building that houses equipment related to the accelerator components downstairs. Service buildings also provide access to the tunnel.

In between the major mileposts, you find locations by intermediate numbers; for example, the service buildings on either side of MI-60 are MI-52 and MI-62.

The beam transport lines, which connect the different accelerators to each other, have their own names and their own set of location numbers. For example, locations in the MI-8 line, which transports 8 GeV protons from Booster to the Main Injector, are predictably assigned numbers beginning with "8." Other beam lines are named for the type of particle transported: "P" for protons, or "A" for antiprotons. Protons leaving the Main Injector use the P1 line, originating at MI-52. (Do not be confused by the fact that antiprotons sometimes use the P1 line as well.) Antiprotons use the A1 line, originating at MI-62, to leave the Main Injector. The P1 and A1 lines both terminate at the F0 location of the Tevatron, which runs parallel to MI-60. Devices in the P1 line are given numbers beginning with a “7,” and device names in the A1 line begin with a “9.”

Two of the beam lines that deserve special mention are the P2 and P3 lines. Main Ring, like the Tevatron, was divided into six sectors, A through F. Although most of Main Ring has been disassembled and cannibalized, F Sector remains more or less intact. (Sometimes this old section of the Main Ring is referred to as the Main Ring Remnant.) It is used to link Main Injector to Switchyard in one mode, and to the Antiproton Source in another. Locations are designated F10 through F49. The boundary between P2 and P3 is at F17, which is where the beam line to the Antiproton Source branches off.

Main Injector accepts 8 GeV protons from the Booster, 8 GeV antiprotons from the Antiproton Source, or 8 GeV antiprotons from the Recycler. It can accelerate protons to 120 GeV for antiproton production, or for extraction to the Switchyard; it accelerates protons and antiprotons to 150 GeV during Collider Mode. When the Recycler is fully commissioned, it will accept 150 GeV antiprotons from the Tevatron and decelerate them to 8 GeV before injecting them into the Recycler to be stored. Finally, it is implicated in two neutrino experiments, MiniBooNE and NuMI.

One at a time, please!

The MI-8 Line

Linac accelerates protons to 400 MeV, and Booster accelerates them to 8 GeV. The purpose of the MI-8 line is to transport the protons from Booster to the Main Injector ring (Fig. 1-2). Extraction devices along the western perimeter of Booster kick the beam up and out of that machine. Permanent magnets (for the most part) guide the beam through the line. Generally, the magnets bend the beam to the right, avoiding the Antiproton Rings and eventually bringing the beam line parallel to the Main Injector ring. Injection devices at MI-10 place it onto the orbit it follows as it circulates around the Main Injector ring. The magnets in Main Injector keep the beam on that orbit.

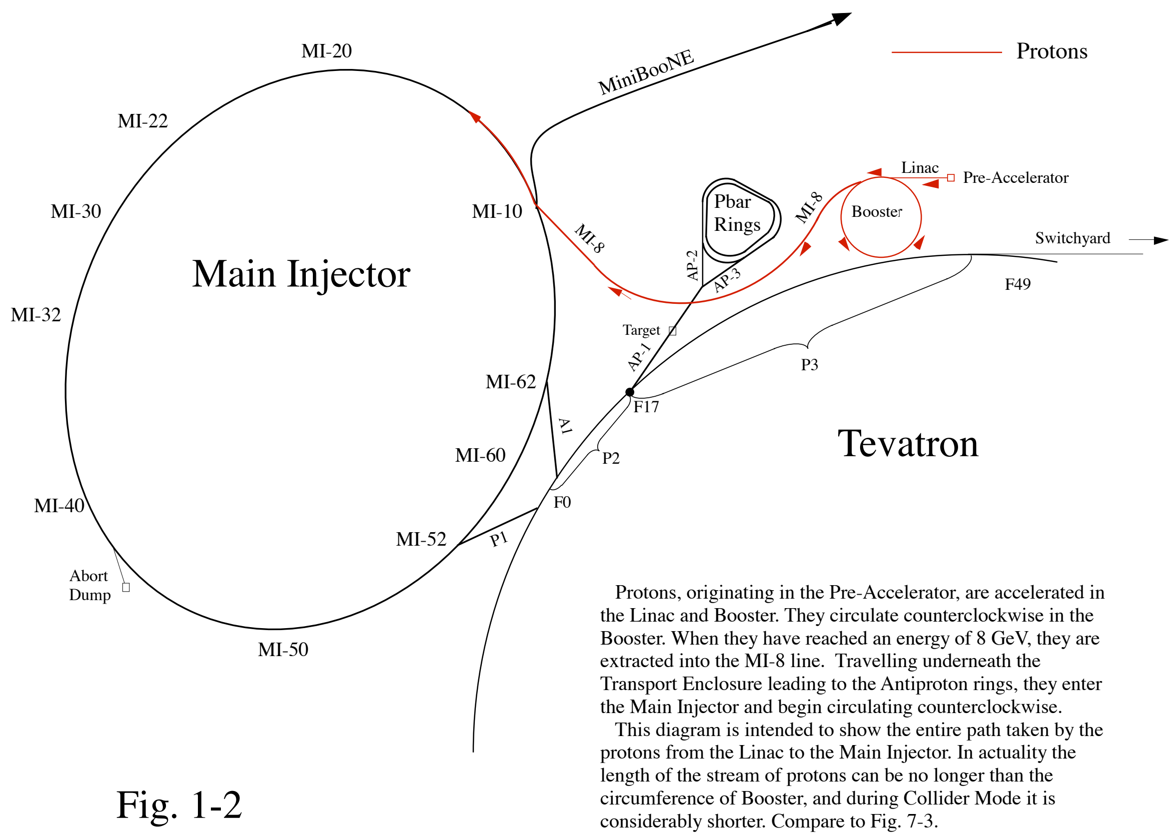

Fig. 1-2
Protons into Main Injector

Once the beam is circulating and stable, it is accelerated to the desired energy by RF (radio frequency) systems located at MI-60. The details of RF will be discussed in a later chapter, but to understand the various modes of operation it is necessary to understand batches and bunches.

A batch consists of the beam that Booster accelerates and extracts in one cycle. A full batch has a length equal to the circumference of Booster, about 1,545 feet.

The beam does not form a continuous stream, but is bunched. This is because the protons congregate around a certain phase of the RF wave—called a bucket—and nowhere else (again, the reasons for this are explained in a later chapter). When beam is extracted from Booster, the RF wave is slightly over 18 feet long, and there is one bunch per bucket:

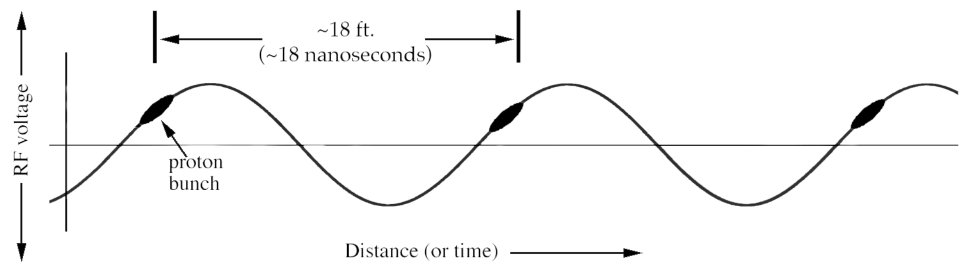

Because the wavelength of the RF wave is usually much smaller than the circumference of the accelerator, there are a number of bunches circulating in a machine at any given moment. The Booster, needs 84 bunches to fill all of the available RF slots around the circumference of the machine. The total sum of the bunches (or available buckets) is what constitutes a batch.

Main Injector Ramps

As soon as stable circulating beam is established, the acceleration process can begin. The RF adds energy to the beam. The magnets provide the constraining force that keeps the protons on the correct orbit, so that the current in the magnets increases to match the magnetic field to the energy of the protons (Fig. 1-3(a)). (It is the synchronous match of the magnets to the beam energy that entitles us to call our circular accelerators “synchrotrons.”)

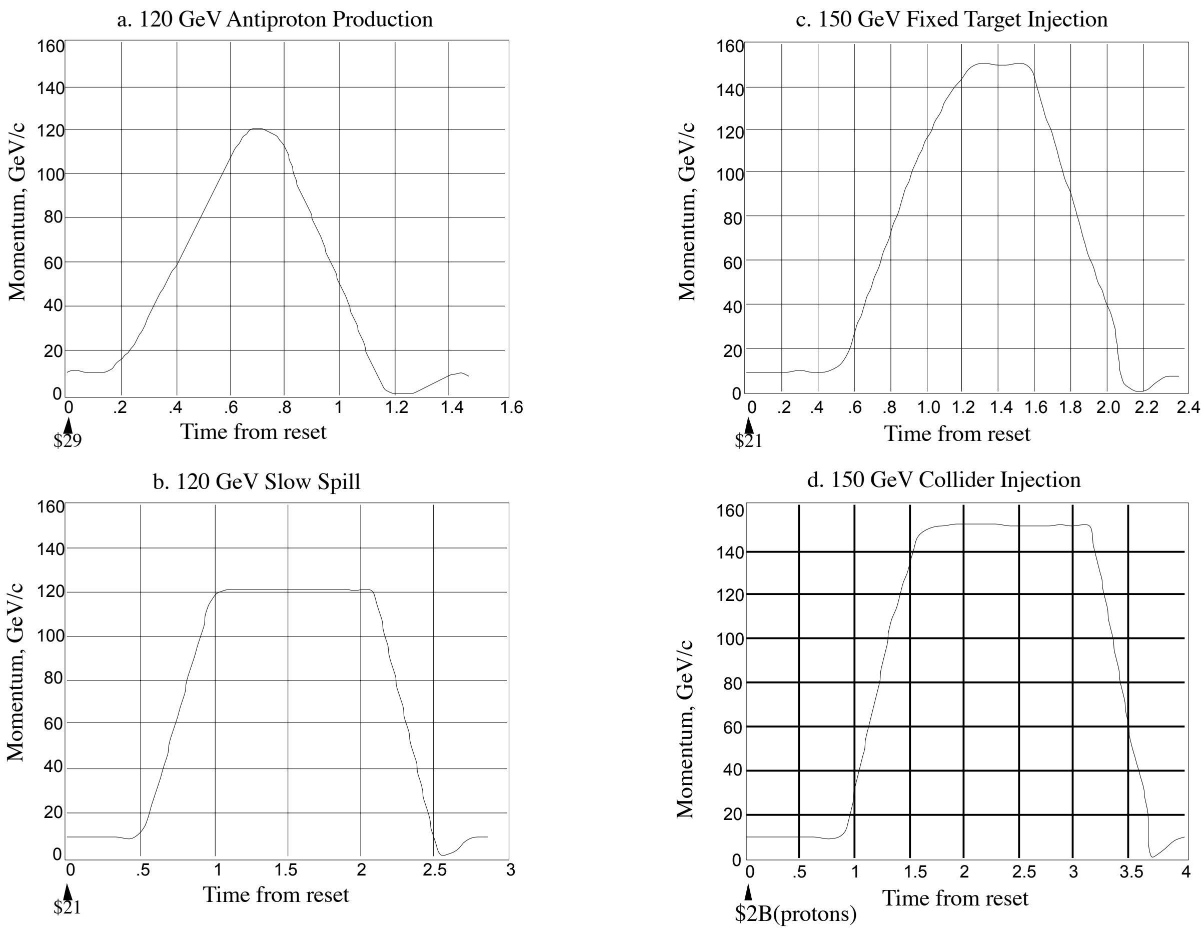

Fig. 1-3
Some Main Injector Ramps

Antiproton Production

Providing beam to the antiproton production target is one of the simplest tasks required of the Main Injector (Fig. 1-4). In this mode, a single batch of protons is accepted from the Booster at 8 GeV and accelerated to 120 GeV. The 120 GeV protons are extracted to the target, which yields 8 GeV antiprotons. (Scientists unable to pronounce the word “antiproton” use the synonym “Pbar,” which will occasionally be used in this book to simplify graphics.)

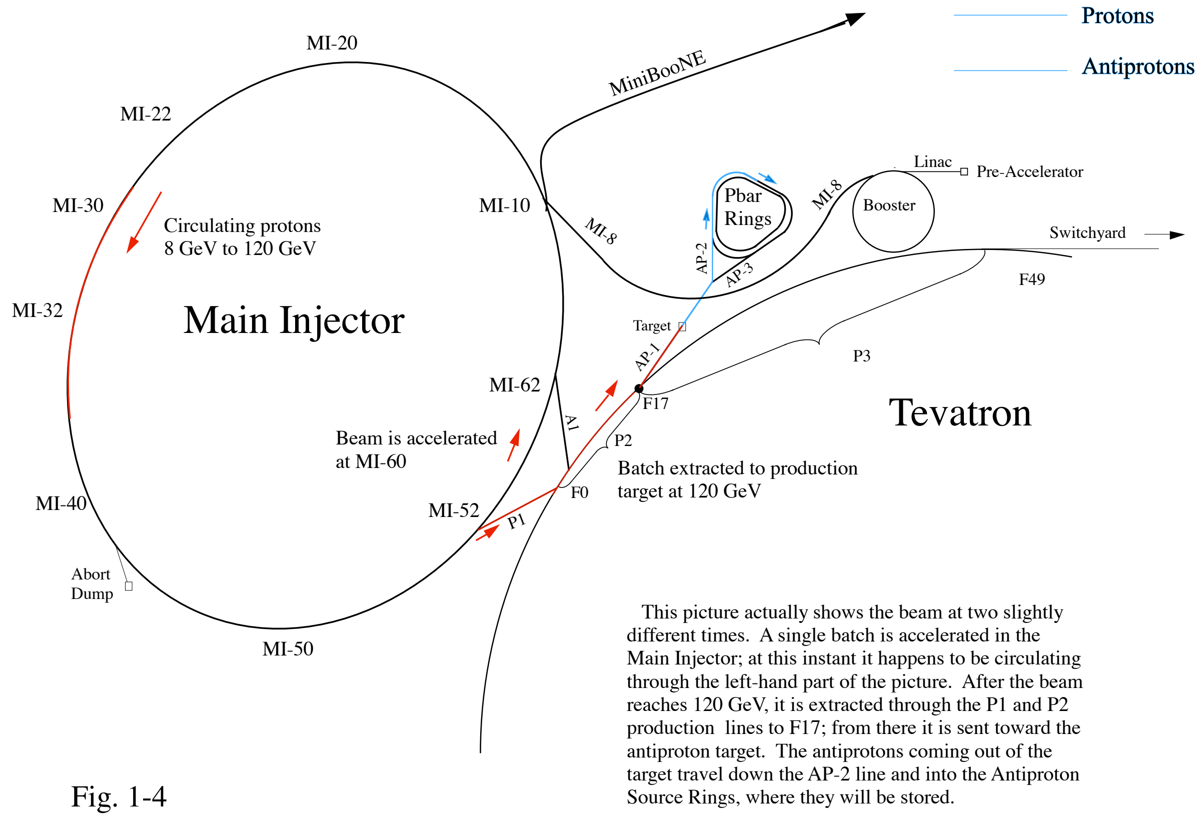

Fig. 1-4
Antiproton Production

Each of the modes of operation is initiated by a unique clock event. In the case of antiproton production, it is Event $29 (the "$" means that the number is hexadecimal). The clock events that trigger Main Injector cycles are designated by $2x, where x is a number 0 through F. These clock events come from the Time Line Generator (TLG).

When Event $29 occurs, it is normally accompanied by Event $14, which instructs the Booster to prepare one batch of beam. It requires exactly a thirtieth of a second for beam to be accelerated through Linac and Booster. It requires another couple of microseconds to transfer the beam out of Booster, through the 8 GeV line, and onto the Main Injector orbit.

The beginning of the acceleration is not linear; using a parabola softens what would otherwise be an abrupt change in the magnetic field. After the parabola, the rate of change does become linear for a while. As the final energy is approached, an inverted parabola eases the beam into "flattop." Flattop is a short period at the final energy during which the beam continues to circulate. In the case of the $29 cycle, flattop lasts for about 40 msec. Some modes of operation have distinct and complex tasks that must be completed during flattop.

When it is time for the beam to leave and make antiprotons, extraction devices in the Main Injector deflect the entire batch out of the ring and into the P1 line originating at MI-52. From there, the protons find their way into the P2 line; extraction from the P2 line into the AP-1 line takes place at F17. The AP-1 line carries them to the target.

The entire 120 GeV antiproton production cycle takes about 2 seconds. To make as many antiprotons as possible, the target is continuously bombarded at the fastest possible rate for hours on end.

Fixed Target Modes

During Fixed Target operations, protons are accelerated to the desired energy and then extracted to a stationary target external to the ring. Extraction takes place from the Main Injector at 120 GeV (Fig. 1-5). The target can be anything from a sliver of metal to a flask of liquid hydrogen.

Extraction is initiated by clock event $21. It starts, as usual, with 8 GeV beam from the Booster. But this time, six full batches—all that the Main Injector can comfortably hold—are loaded in quick succession. When the Main Injector is full, the six batches are accelerated to 120 GeV.

At flattop, a process called resonant extraction is initiated. Resonant extraction is a rather complicated process and the details will appear in a later chapter. Basically the beam "spills" out of the machine over an enormously long time—say, two seconds—as opposed to being kicked out in a few microseconds as is the case with other modes of operation. The flattop time is extended accordingly (Fig. 1-3(b)).

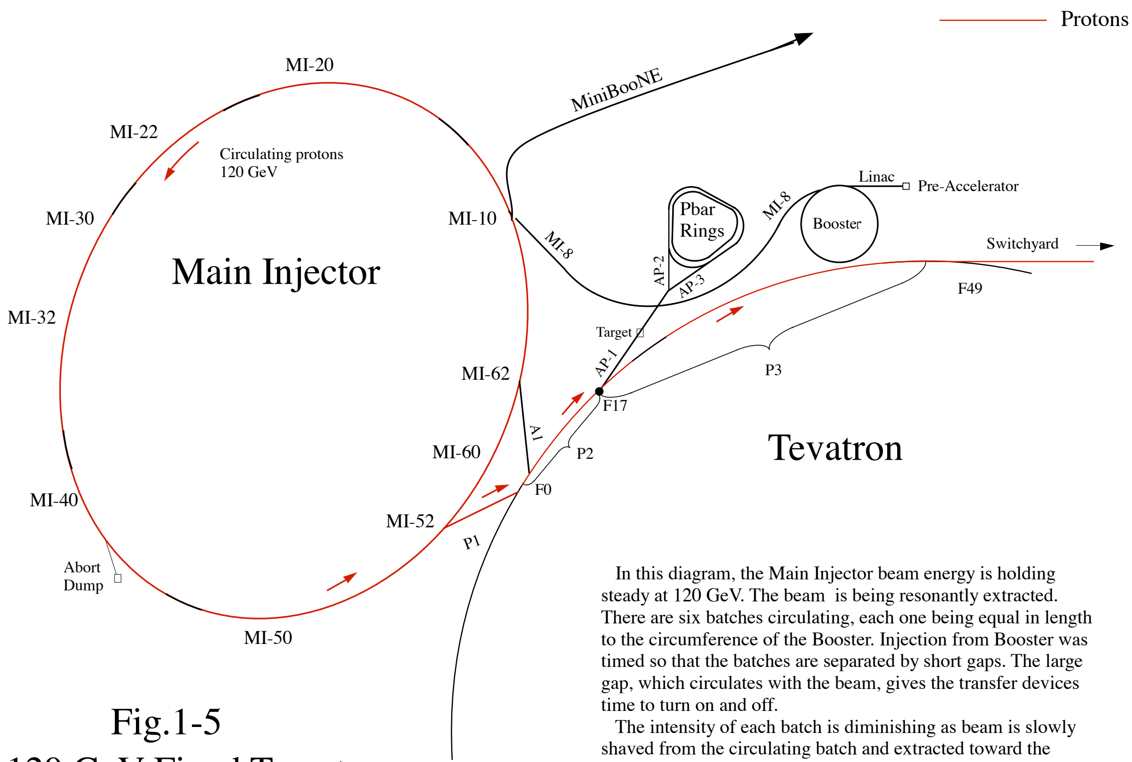

In this diagram, the Main Injector beam energy is holding steady at 120 GeV. The beam is being resonantly extracted. There are six batches circulating, each one being equal in length to the circumference of the Booster. Injection from Booster was timed so that the batches are separated by short gaps. The large gap, which circulates with the beam, gives the transfer devices time to turn on and off.

The intensity of each batch is diminishing as beam is slowly shaved from the circulating batch and extracted toward the Switchyard through the P1, P2 and P3 lines. It takes about a second to extract all of the beam from the machine.

Fig.1-5
120 GeV Fixed Target

For the Fixed Target experiments, beam is deflected into the P1 line—the same line used for antiproton production—and likewise enters the P2 line. This time the extraction devices at F17 are turned off, and the protons travel the length of the P3 line into Switchyard.

Spill over a couple of seconds is considered “slow spill.” There is also “fast spill,” usually required by neutrino experiments, which occurs over a period of a few milliseconds. Keep in mind that fast spill is still much slower than single turn extraction.

One of the experiments that will use fast spill is NuMI (for Neutrinos from Main Injector). This experiment will be a very long baseline version of MiniBooNE, also designed to look for neutrino oscillations. Protons, at 120 GeV, will be extracted from Main Injector at MI-60 at a sharp downward angle. The neutrinos from the target, moving in the same direction as their proton parents, are measured before they disappear into the ground. Tracing a cord through the Earth, they travel 425 miles before encountering a once-abandoned iron mine in northern Minnesota. Here an experiment is being set up to measure what changes, if any, have taken place over the 425-mile baseline.

(Once upon a time, the sole purpose of the Tevatron was to provide beam for the Fixed Target experiments. At this time, it is not expected that this mode will ever be revived, but the following paragraph is retained for historical perspective):

[If higher energies are needed for fixed target experiments, the Tevatron can be called into service. The Tevatron can comfortably hold up to twelve batches of protons, but limitations on timing restrict the number of batches that can be transferred. The Main Injector is first filled with five batches, and accelerates them to 150 GeV—the minimum energy at which the Tevatron magnets can sustain a high quality field—and extracts them to the TeV in a single turn (Fig. 1-3(c)). The same P1 line discussed earlier is used for Tevatron injection, but the protons are switched over to the Tevatron at F0 before they have a chance to enter P2 or P3. Main Injector then drops back to 8 GeV and loads up another six batches to finish the job. When the 11 batches have been loaded, the Tevatron accelerates the protons to some energy between 800 GeV and 1 TeV. Spill from the Tevatron takes about 40 seconds out of a total of 80 seconds in a cycle; when the Tevatron returns to 150 GeV, the Main Injector will be ready to load the next 11 batches.]

MiniBooNE

MiniBooNE (Fig. 1-6), a direct descendant of an experiment done in Los Alamos during the 1990's, is studying the phenomenon of "neutrino oscillations." It requires large numbers of protons with an energy of 8 GeV. Beam is extracted from the Booster and sent down the MI-8 line, but just as it

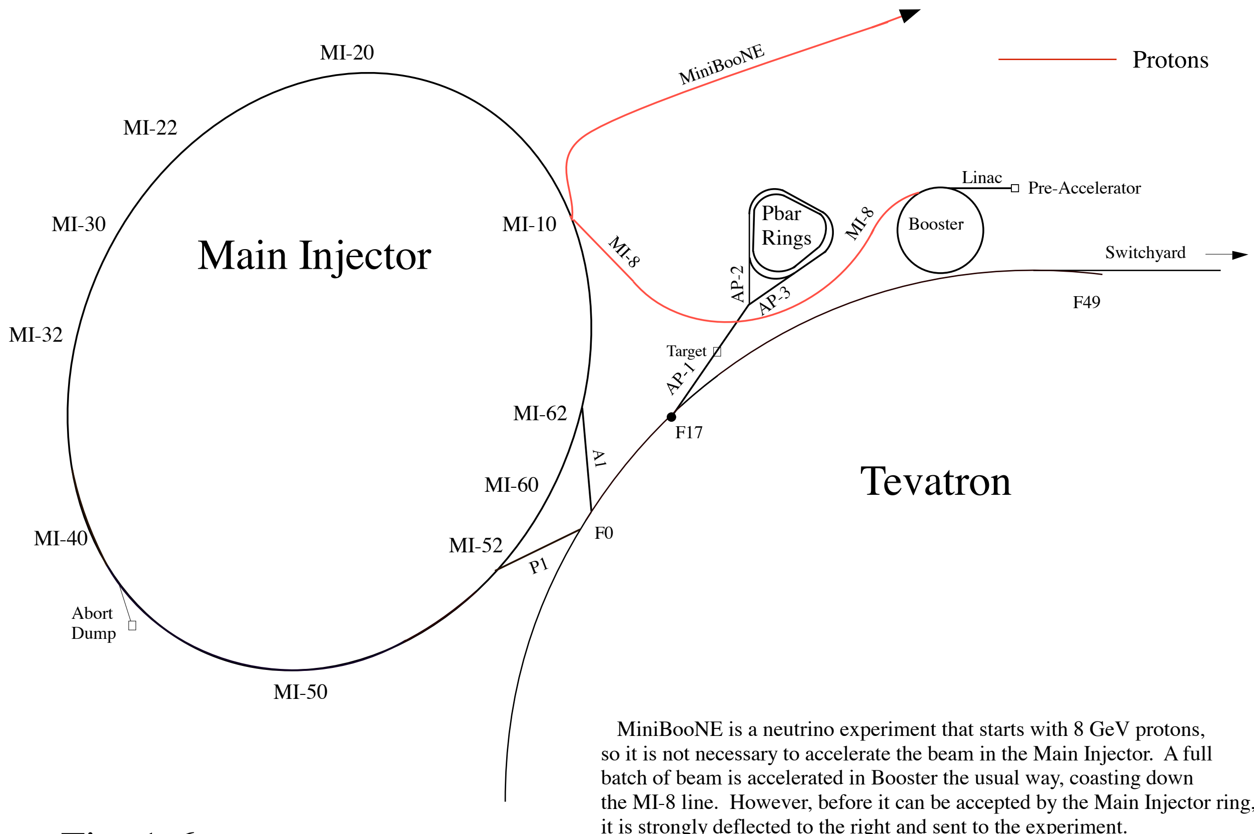

Fig. 1-6
MiniBooNE

reaches the Main Injector ring, it is deflected north toward the experiment.

Collider Operations

Collider Mode is the most complex scenario of all. Rather than striking a palpable, distinct target, protons circulating clockwise in the Tevatron collide with antiprotons that are circulating counterclockwise. The collisions take place in the Tevatron at an energy of 1 TeV (hence the name, although the current operational energy is actually only 980 GeV). Two large experiments hang like parasites around the Tevatron beam pipe at those points where the particles collide.

Main Injector must play several roles during Collider Mode. In addition to supplying 120 GeV protons for antiproton production, it must also feed the Tevatron protons and antiprotons at 150 GeV. To make matters worse, the experiments need "superbunches" more intense than any individual bunch than can be accelerated by Booster. A process called coalescing has been developed to create those superbunches; coalescing takes place at flattop in the Main Injector.

When protons are to be loaded, accelerated, and coalesced in the Main Injector, the sequence is initiated by Clock Event $2B (Fig.1-3(d)); for antiprotons, the event is $2A.

The following is a possible sequence of steps taken during a shot (the scientific term for loading the protons and antiprotons):

(1) One batch (84 bunches) of protons is accelerated to 8 GeV in the Booster.

(2) Only 7 (or so) bunches of the 84 are extracted to Main Injector; the remaining 77 are sent to the Booster dump. This group of bunches is a short batch or a partial batch, depending on who is doing the describing.

(3) The 7 bunches are accelerated to 150 GeV (Fig. 1-7).

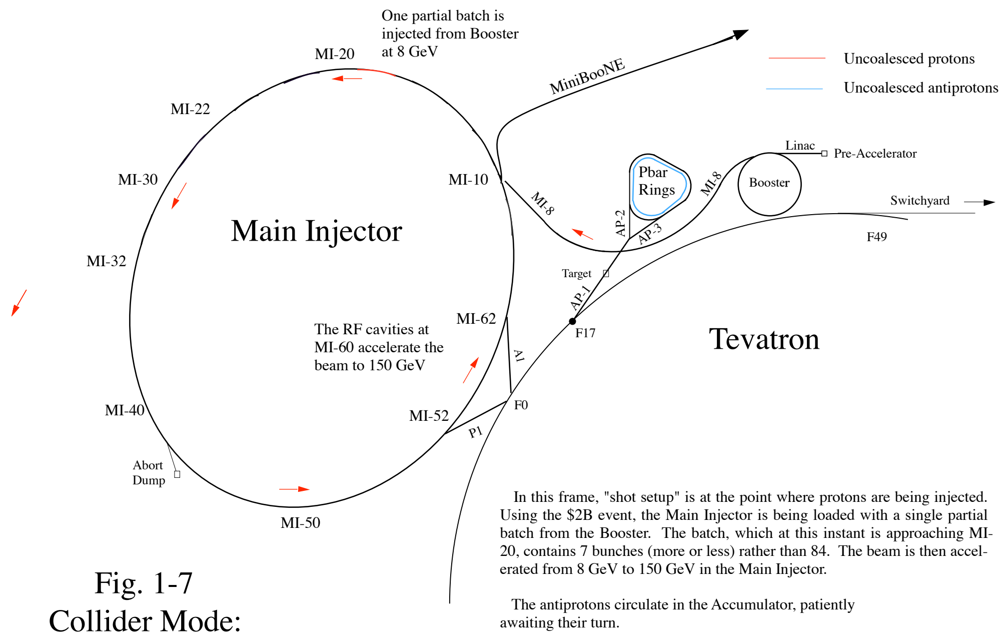

Fig. 1-7
Collider Mode:
Protons into Main Injector

(4) At flattop, the bunches are coalesced; that is, pushed together to form one narrow, high intensity bunch (Fig. 1-8).

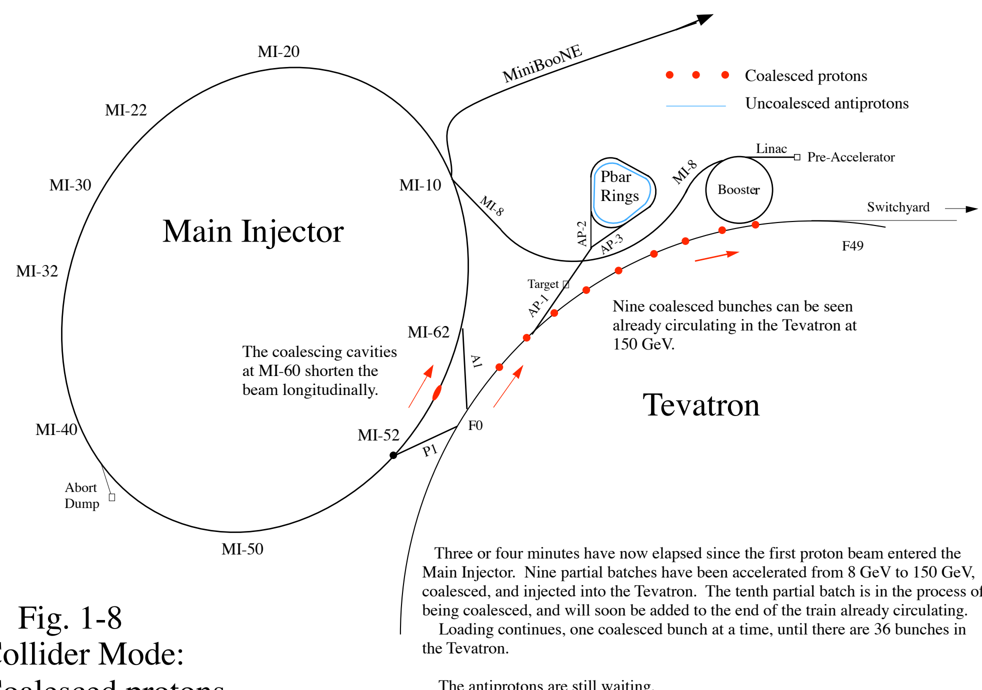

Fig. 1-8
Collider Mode:
Coalesced protons
into Tevatron

(5) The coalesced bunch is injected into the Tevatron in a single turn, passing through the P1 line on the way (Fig. 1-8).

(6) The above steps are repeated until there are 36 coalesced proton bunches in the Tevatron.

(7) The 8 GeV antiprotons created from the production target have been patiently waiting, circulating in the Antiproton Source (more specifically, the Accumulator). The Accumulator RF systems create four groups of antiprotons, each group consisting of several bunches. They are extracted from the Accumulator via the AP-3 and AP-1 lines, and into the P2 line at F17 (Fig. 1-9). Note that once they are in the AP-1 line, they are traveling backwards from the route taken by the 120 GeV protons. They enter the Main Injector from the P1 line at MI-52 and begin to circulate clockwise, which is the opposite direction that protons would be taking.

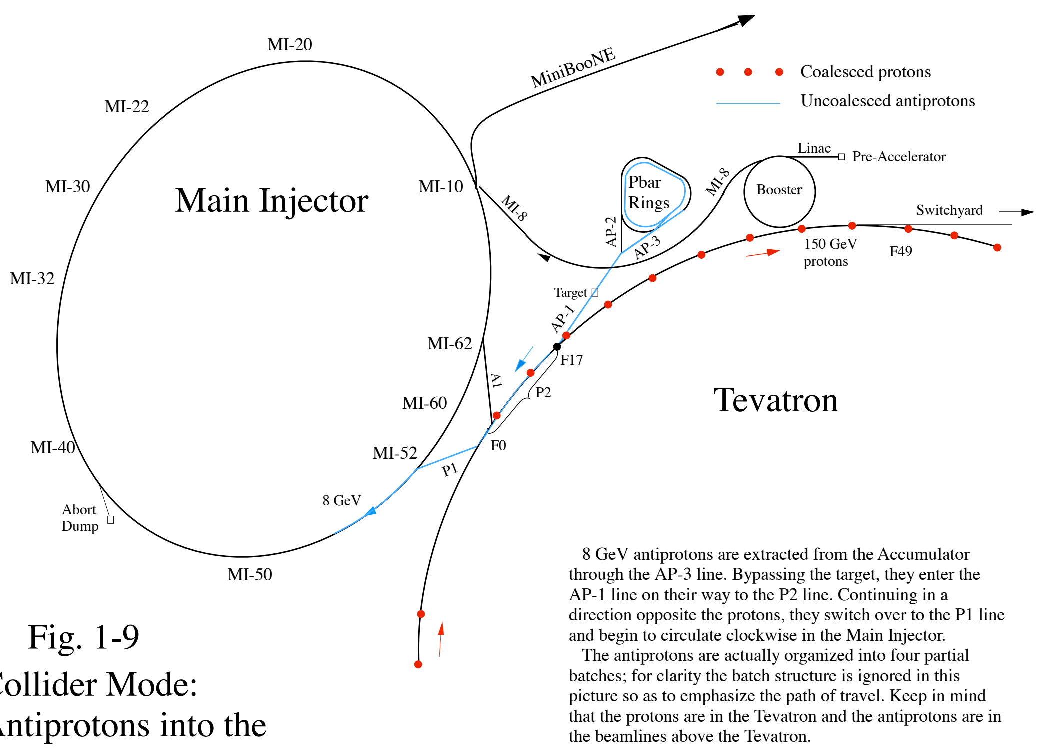

Fig. 1-9
Collider Mode:
Antiprotons into the
Main Injector

(8) The antiprotons are accelerated to 150 GeV (Fig. 1-10).

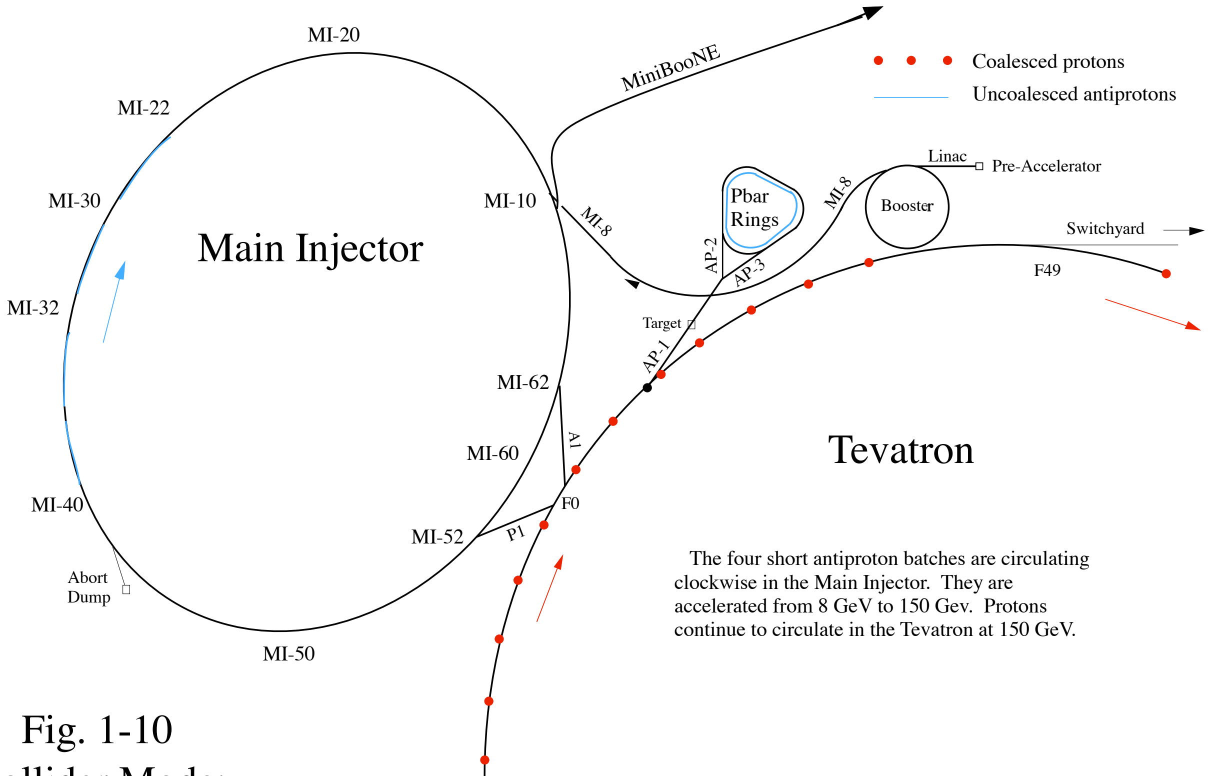

Fig. 1-10
Collider Mode:
Antiprotons Accelerated
in Main Injector

(9) The four groups of antiprotons are coalesced; four coalesced bunches are now circulating at 150 GeV.

(10)The four coalesced bunches are extracted to the Tevatron by way of the A1 line, originating at MI-62, and they enter the Tevatron at F0. (Notice that there is a gap in the string of proton bunches that allows room for the antiprotons.) The antiprotons are traveling counterclockwise in the Tevatron, opposite the direction taken by the protons. Both types of

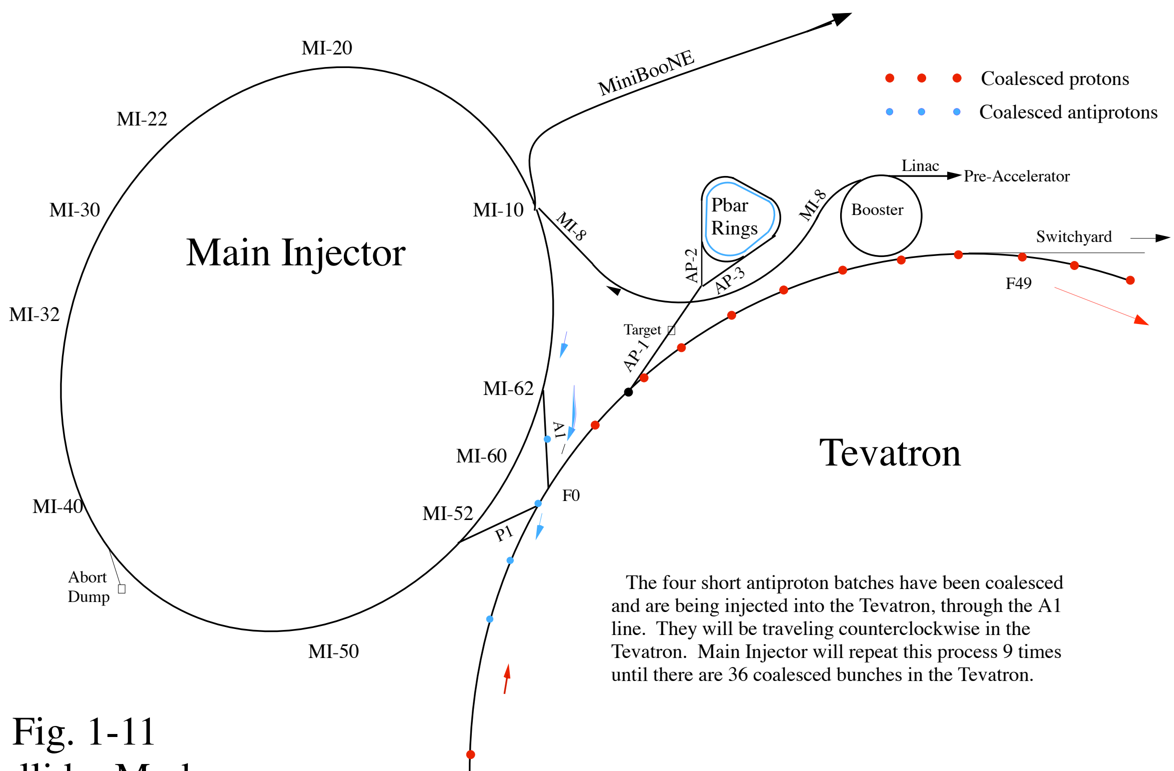

Fig. 1-11
Collider Mode:
Coalesced Antiprotons
into Tevatron

particles are now circulating in the Tevatron at 150 GeV (Fig. 1-11).

(11)Main Injector drops back to 8 GeV for another four groups of antiproton bunches. The process repeats until nine shipments have been delivered. There are now 36 coalesced bunches of antiprotons and 36 coalesced bunches of protons circulating in the Tevatron.

(12)The Tevatron ramps to 1 TeV with all of the coalesced bunches. A 36X36 store is established, and the two particle detectors in the Tevatron ring begin to take data. The Main Injector is now free to return to stacking. A store is intended to last for at least several hours (Fig.1-12).

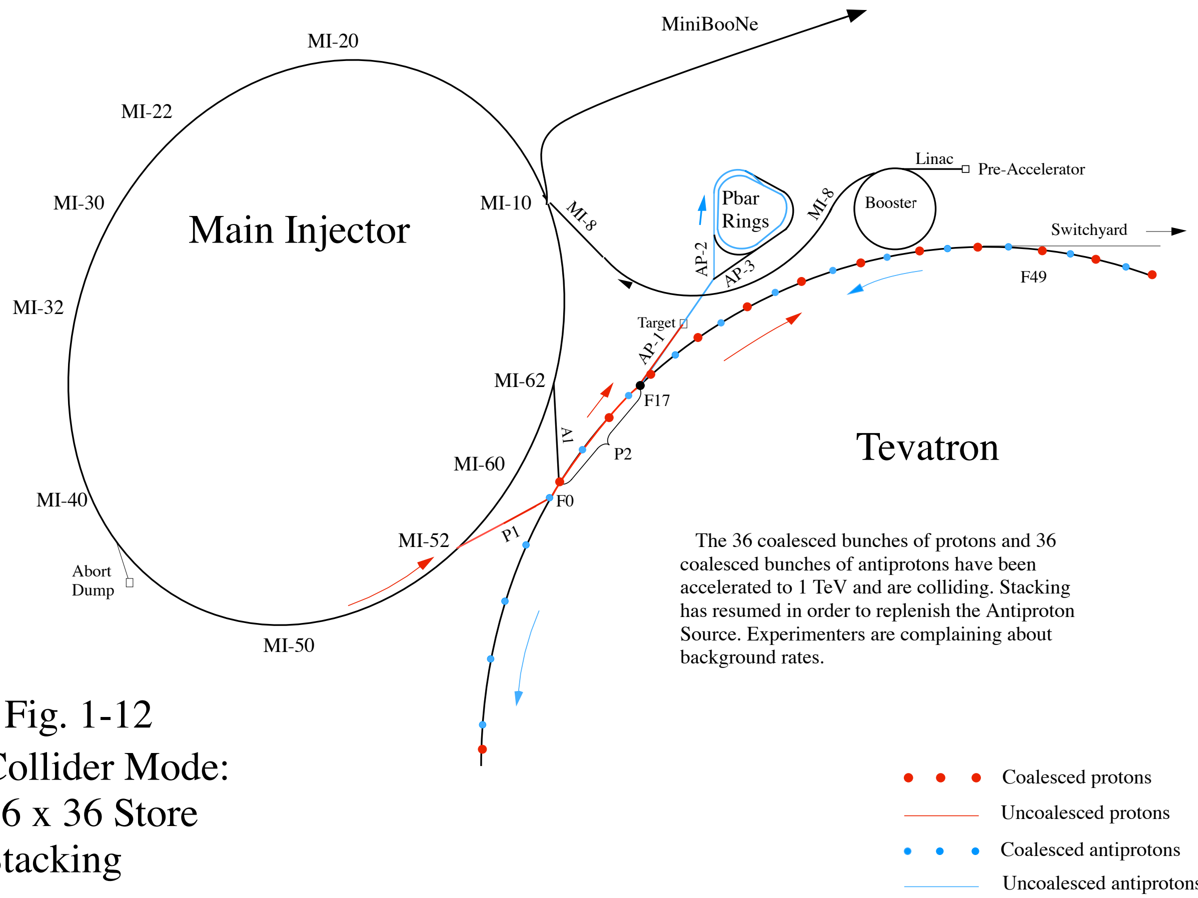

Fig. 1-12
Collider Mode:
36 x 36 Store
Stacking

Implicit in all of these maneuvers is the need for all of the coalesced bunches to be in the right place with respect to each other. This is accomplished by a technique called cogging, which in a rather abstract way resembles the cogs of intermeshing gears. Some cogging takes place in Main Injector and some occurs in the Tevatron. Cogging is one of the RF manipulations of the beam and will be discussed in a later chapter.

Main Injector

Role of the Recycler

The purpose of the Recycler is to store excess antiprotons. The number of antiprotons available has always been an important limiting factor in producing the high luminosities desired for stores in the Tevatron. They are difficult, or at least time-consuming, to produce—and if enough are created they are difficult to store because large numbers of antiprotons in the Accumulator tend to develop instabilities and can be suddenly lost.

The Recycler will store antiprotons from two sources: They can be siphoned off from large stacks, or they can be scavenged at the end of Tevatron stores. Both methods require the Main Injector as a middleman.

The Recycler ring is located directly above the Main Injector ring. Antiprotons are transferred from Main Injector to the Recycler through a beam line at MI-32. Transfer from the Recycler to Main Injector takes place at MI-22.

If the antiprotons are to come from the Accumulator, they will arrive in the same way as they would during a shot—through the AP-3 and AP-1 lines, into the P2 line at F17, and back through the P1 line, entering Main Injector at MI-52. Now, rather than being accelerated, they are transferred into the Recycler at MI-32.

Retrieving antiprotons from a store is a bit more involved, since the antiprotons are at 1 TeV. Here is a possible sequence of events, stepwise:

(1) The store has been circulating in the Tevatron for several hours. As particles are gradually lost and the beam size slowly grows, the luminosity degrades. At some point, a decision is made to terminate the store and load a fresh one.

(2) Scrapers, basically large chunks of metal, are slowly moved into the proton beam until only the antiprotons are left (the proton and antiproton orbits follow different paths in the Tevatron).

(3) The antiprotons are decelerated from 1 TeV to 150 GeV, using the Tevatron RF systems.

(4) A group of coalesced bunches of antiprotons is transferred from the Tevatron to Main Injector.

(5) While the antiprotons are still circulating at 150 GeV, they are "decoalesced"; that is, decomposed back into 7 or so bunches. If this is not done they will probably not survive the trip down.

(6) The antiprotons are decelerated to 8 GeV.

(7) The antiprotons are transferred to the Recycler via the MI-32 line.

(8) The Main Injector is ramped back to 150 GeV and accepts a new group of coalesced bunches; decoalescing and deceleration follow.

(9) This continues until all of the antiprotons from the store are now in the Recycler.

The specifics of this plan may very well change as operational experience is gained. Be alert.

The Recycler has its own RF system to facilitate transfers to and from the Main Injector. It also, unlike Main Injector, has its own cooling systems. Cooling, in this sense, means reducing the random energy of the particles. Details of the Recycler RF and cooling systems will be made available in the Recycler Rookie Book.

Mixed Modes

There are those who wish to complicate the simple scenarios described in the preceding sections. For example, it might be possible to perform fixed target extraction and antiproton production simultaneously by accelerating six batches to 120 GeV, kicking one batch out toward the antiproton target, and resonantly extracting the remaining five. There would have to be an increase in the flattop time in order to accommodate both operations, but in the end it might be more efficient. History shows that at Fermilab, people want to do everything at once.

Two modes that are not compatible are Fixed Target extraction from the Tevatron and Collider Mode. Because of the many tunnel components that need to be changed, at least a few weeks are required to complete the switchover.

The details that have been so conveniently ignored over the last few pages will be addressed in the following chapters. Many different kinds of magnets are necessary to direct the beam where it is supposed to go; most of the magnets require power supplies and cooling water. There are vacuum systems to clear the beam pipe of the air molecules that would disrupt passage of the beam. The RF systems must not only accelerate two different kinds of particles, but also perform the complex manipulations of bunch rotation, cogging, and coalescing. Diagnostics must be present to analyze the beam. Control systems must ferry information to and from the thousands of individual components necessary to make the Main Injector function, and provide the precise timing required to keep them synchronized. Much more has to be said about the beam transport lines that connect one accelerator to another.

The next chapter, “Magnets and the Lattice,” will discuss in some detail how the various kinds of magnets maintain circulating beam in the Main Injector.

Chapter 2: Magnets and the Lattice

Epoxy on both your houses!

- Romeo and Juliet

This chapter deals with the magnets that are used to maintain circulating beam in the Main Injector. For all of the effort that goes into building and maintaining the magnets in an accelerator, they have a rather limited role: to constrain the motion of the particles so that they are prevented from leaving the machine. This chapter has an even more limited role, dealing primarily with beam already circulating in the accelerator. Transfers into and out of the Main Injector will be discussed in the chapter on Beam Transfer Lines.

Forces

The basic mathematical equation describing how a charged particle interacts with electric and magnetic fields is:

$$\vec{\mathbf{F}} = \mathbf{q}\vec{\mathbf{E}} + \mathbf{q}\vec{\mathbf{v}} \times \vec{\mathbf{B}}$$

where $\vec{\mathbf{F}}$ is the force on the particle, q is the charge of the particle, $\vec{\mathbf{E}}$ is the electric field, $\vec{\mathbf{v}}$ is the velocity of the particle, and $\vec{\mathbf{B}}$ is the magnetic field.

The quantities with the arrows—that is, everything except for the charge—are vectors. With a vector, the direction of the field or force must be taken into account, in addition to the strength or magnitude. The × symbol represents a form of vector multiplication known as a cross product.

The first term of the equation, $q\vec{\mathbf{E}}$ represents the force created by the electric field. The electric field is crucial to the operation of any accelerator, because it is the only means available for changing the energy of the beam. The direction of the accelerating force is the same as the direction of the electric field. The energy-changing electric field belongs to the realm of the RF systems; it will be dealt with in excruciating detail in a later chapter. For now, strike $\mathbf{q\vec{E}}$ from the equation.

The remaining term, $\mathbf{q\vec{v}} \times \mathbf{\vec{B}}$, describes the influence of the magnetic field. The charge, q, is constant. Although actually measured in coulombs, for our purposes it can be reduced to (+1) for protons and (-1) for antiprotons.

The speed of the particle, surprisingly enough, changes only slightly as a particle is accelerated from 8 GeV to 150 GeV. This odd behavior is due to relativistic effects and will be explained later, but for now it can be assumed that the magnitude of $\mathbf{\vec{v}}$ is constant. Normally measured in, say, meters per second, for our purposes the velocity can be reduced to a +1 for protons and -1 for antiprotons. The difference in "polarity" is because the velocity vector requires a direction, and the frame of reference used here is that of the proton direction in the Main Injector.

The magnetic field $\mathbf{\vec{B}}$ also has a direction. Often the strength and direction of the field are visualized as "field lines" connecting one pole of a magnet to another pole. The magnitude of $\mathbf{\vec{B}}$ can be expressed in units of gauss, or in Tesla. One Tesla is equivalent to 10^4 gauss.

The cross-product, $\mathbf{q\vec{v}} \times \mathbf{\vec{B}}$, means that only the components of $\mathbf{\vec{v}}$ and $\mathbf{\vec{B}}$ which are perpendicular to each other will generate a force. That is, if $\mathbf{\vec{B}}$ is in exactly the same direction as $\mathbf{\vec{v}}$, there is no magnetic force on the particle. Conversely, the maximum force is produced when the magnetic field is exactly perpendicular to the direction of motion. Moreover, the force $\mathbf{\vec{F}}$ is itself perpendicular to both the direction of motion and the direction of the field.

There is an interesting consequence of $\mathbf{\vec{F}} = \mathbf{q\vec{v}} \times \mathbf{\vec{B}}$ that, as it has turned out, has been the single most important driving force behind the construction of the large accelerators built over the last two decades. If a proton (positive charge) passes through a magnet, it will see force of:

$$\mathbf{\vec{F}} = (+1)(+1) \times \mathbf{\vec{B}}$$

If an antiproton (negative charge) passes through the same magnet in the opposite direction, it will see a force of:

$$\vec{\mathbf{F}} = (-1)(-1)\times\vec{\mathbf{B}},$$

its velocity being negative because it is going "backwards" relative to the protons. The result is that $\mathbf{q}\vec{\mathbf{v}}$, and therefore $\vec{\mathbf{F}}$, is identical for both particles. As long as the particles are traveling in opposite directions, the same magnets can be used for both kinds of particles. This fact has spawned an entire industry of antiparticle sources, including Fermilab's Antiproton Source, and will be an indispensable part of the role that the Main Injector and the Recycler will play during Collider Mode. Remember from "Modes of Operation" that the protons travel in a counterclockwise direction, and the antiprotons travel clockwise.

Types of motion

The two fundamental types of particle motion are longitudinal and transverse. Longitudinal motion of a particle is in a "forward" or "backward" direction with respect to the desired path of the beam. Transverse motion is perpendicular with respect to the longitudinal, and can be horizontal, vertical, or some combination of the two. When particle motion is superimposed on coordinate axes, the x-axis represents the horizontal, the y-axis the vertical, and the z-axis the longitudinal.

In accelerators, then, the forward motion of the particles is longitudinal. The magnetic field is always set up so that it is perpendicular to the primary direction of motion, the field lines are usually set up to be either horizontal or vertical. The magnetic force itself is perpendicular to the direction of motion and to the field lines. If the field lines are vertical, the force will bend the beam horizontally, and if they are horizontal, the beam will be bent vertically.

That means that no magnet, no matter how it is designed, can change the longitudinal motion of a particle; magnets are only capable of influencing transverse motion. Magnets are excellent devices for bending and focusing the beam, but useless for changing its energy.

A magnetic field is created whenever an electric current flows in a conductor. The stronger the current, the stronger the magnetic field. It is necessary to use electromagnets whenever the beam changes energy, although the magnets themselves do not change the energy, their bending strength must still track the energy in order to constrain the beam. Usually (but not always), the conductor is arranged as a coil; the number of turns in the coil, along with the amount of current, is proportional to the strength of the field.

There are also permanent magnets, like those pinning down shopping lists on a refrigerator door, which do not explicitly depend on flowing current. The magnetic field comes from the unbalanced "spin" of the electrons in certain materials. Most of the magnets in the 8 GeV line and the Recycler are permanent magnets; the beam energy in both cases is held constant at 8 GeV. The magnetic material used is strontium ferrite.

The disadvantage of electromagnets is that they require a great deal of infrastructure, including large amounts of electrical power and elaborate water systems for keeping them cool. Permanent magnets do not require the same infrastructure, but there is the disadvantage that once the field is fixed it is difficult to change.

Main Dipoles

Dipoles are magnets in which the entire beam passing through them is pushed in the same direction. The two poles are the familiar "north" and "south" poles, although those terms are rarely used at Fermilab.

Dipoles in the Main Injector come in several sizes.

- The main dipoles described in this section weigh several tons apiece.
- There are also corrector dipoles that weigh a few hundred pounds apiece (but these magnets won't be described until later in this chapter).
- There are permanent magnet dipoles used in the MI-8 line and the Recycler.
- Finally, there are specialized dipoles in the beam transfer lines, which will be covered in another chapter.

Main Injector

Fig. 2-1 shows the ring-wide distribution of the main dipoles; each small box represents one of the main dipoles. These dipoles are responsible for bending the beam around the curvature of the ring. The “A” and “B” dipoles, shown in dark blue, are 6 meters long. The “C” and “D” dipoles, in light blue, are 4 meters long.

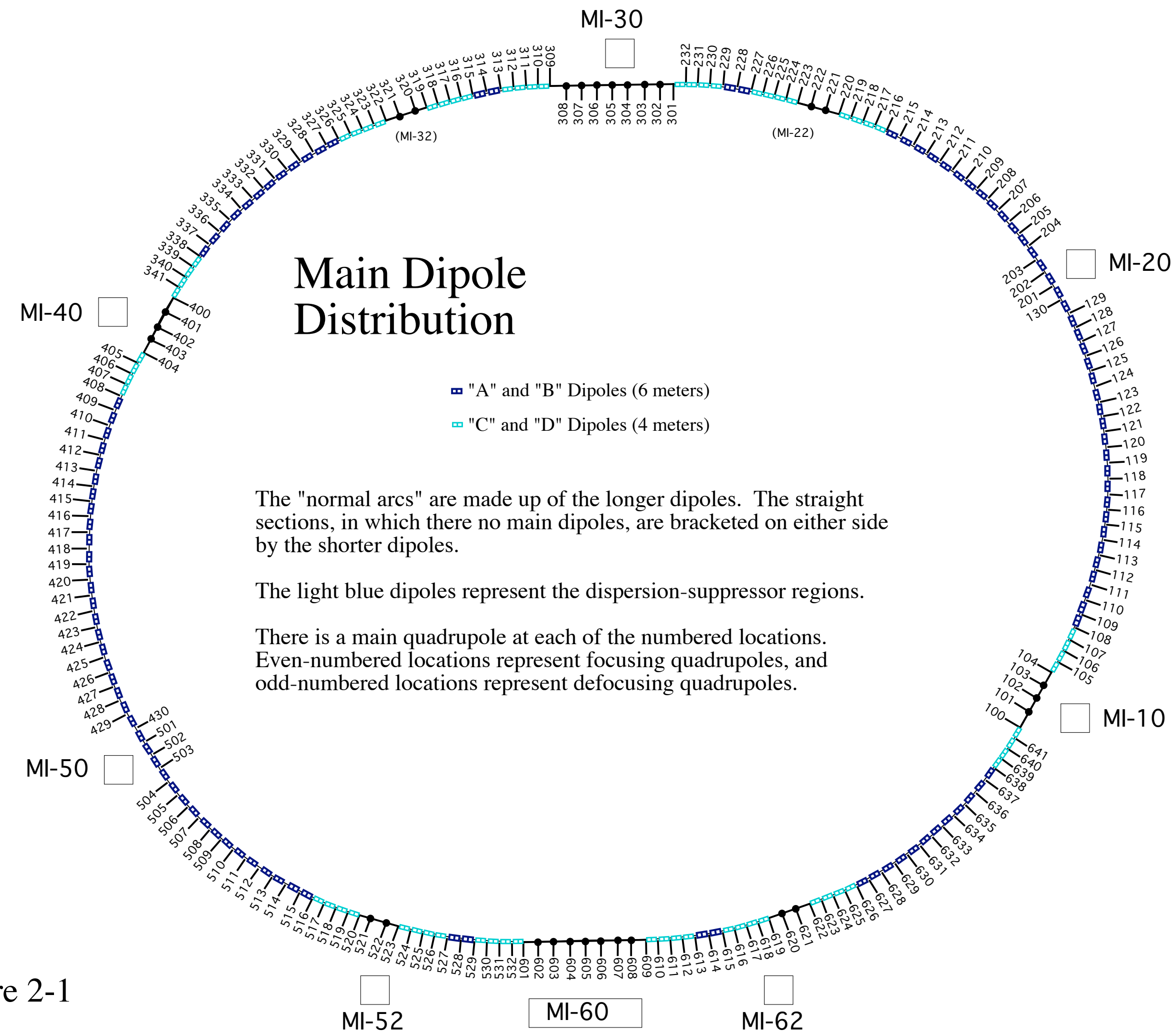

Figure 2-1

The scheme for numbering locations in the Main Injector is based on the direction that the protons travel. Notice in the diagram that the numbers increase in the counterclockwise direction, beginning at MI-10; MI-10 was chosen as the starting point because that is where protons first enter the machine from the MI-8 line. The first series of numbers, from MI-10 to MI-20, begins with "100." A series of numbers in the 200's begins at MI-20, with a new series being launched at each of the major service buildings.

Notice in Fig. 2-1 that wherever the main dipoles are present, there is one pair between each numbered location. "A" and "B" magnets are always paired together, as are the "C" and "D" magnets. Within each pair they are arranged alphabetically, also in a counterclockwise direction. (The convention in this book is that side-view diagrams will be drawn from the perspective of an observer standing in the tunnels. Since the magnets line the outside wall of the tunnel, this convention will soon lead to some awkward pictures: numbers and letters will increase from right to left. It will be uncomfortable for those trained in the English language or number lines, but get used to it. Think like an Egyptian.) There will be a more detailed description of the numbering system after the many different kinds of magnets have been introduced.

The regions made up of "A" and "B" dipoles are called the "normal" arcs. Where there are no main dipoles, there is no curvature; these regions are the straight sections. Altogether, there are 8 straight sections in the Main Injector. They are used for a variety of specialized functions, usually involving beam transfer. The "C" and "D" magnets are found in the dispersion suppressor regions, which act as bridges between the normal arcs and the straight sections. A more complete explanation of these different parts of the ring will appear later in the chapter.

All four types of magnets, although they differ in other respects, are identical in cross-section; a typical slice of a main dipole magnet is shown in Fig. 2-2. The beam travels inside the elliptical beam pipe (sometimes referred to as the beam tube) that runs along the central axis of the magnet. Obviously, it is inside the beam pipe where the strength and direction of the magnetic field are the most important.

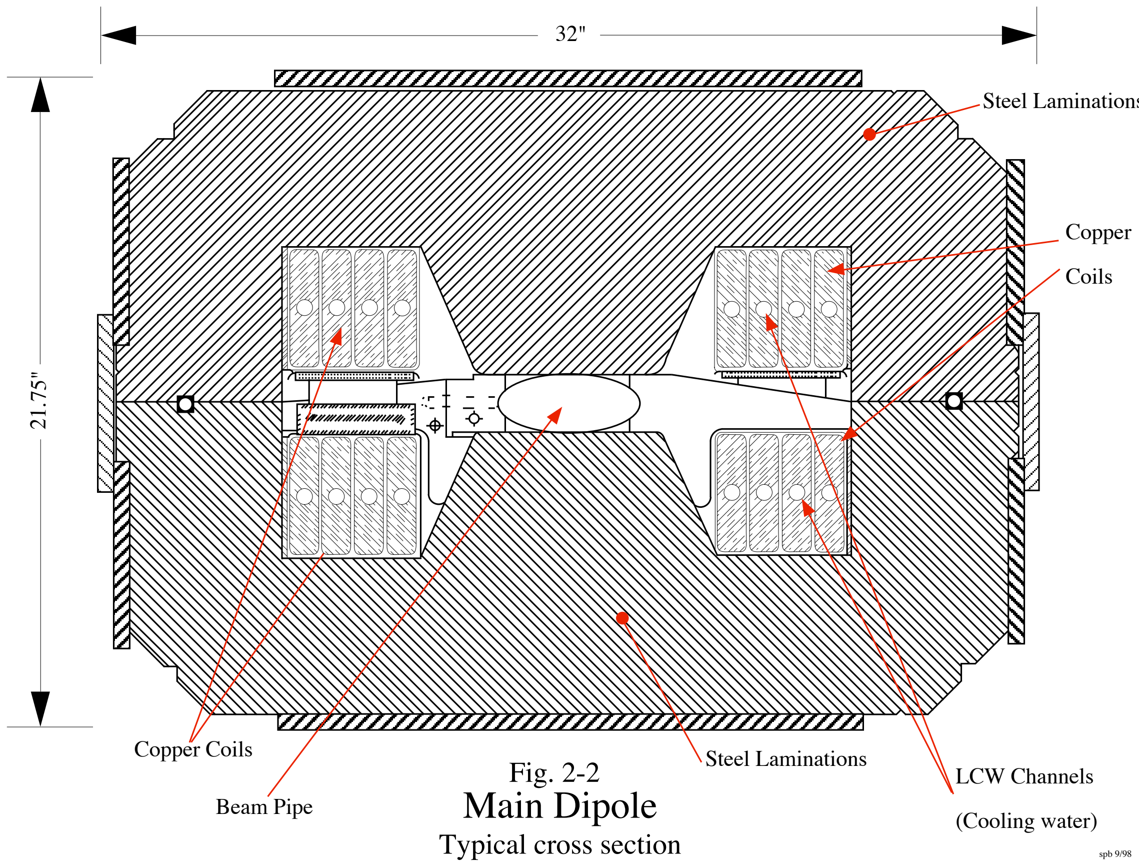

Fig. 2-2
Main Dipole
Typical cross section

Current flowing in a conductor creates the magnetic field, and, as with most electromagnets, the conductor is wound into a coil. A coil produces the effect of having several adjacent conductors working together to create the magnetic field. Fig. 2-3 shows what the coils of a pair of main dipole magnets would look like if all of the surrounding material were removed; the view of the upper and lower coils has been exploded in order to clarify the upcoming description of current flow.

The magnetic field near a current-carrying conductor can be visualized by grasping the conductor with your right hand (in your imagination, of course—to actually do so is probably a violation of LOTO). If your thumb is pointing in the direction of the current, your fingers will naturally “curl around” the conductor in the same direction as the magnetic field. In the figure below the current flow in a main dipole is schematically represented. The perspective for a main dipole would be looking downstream in the proton direction.

Fig. 2-3

Main Dipole Coil Current

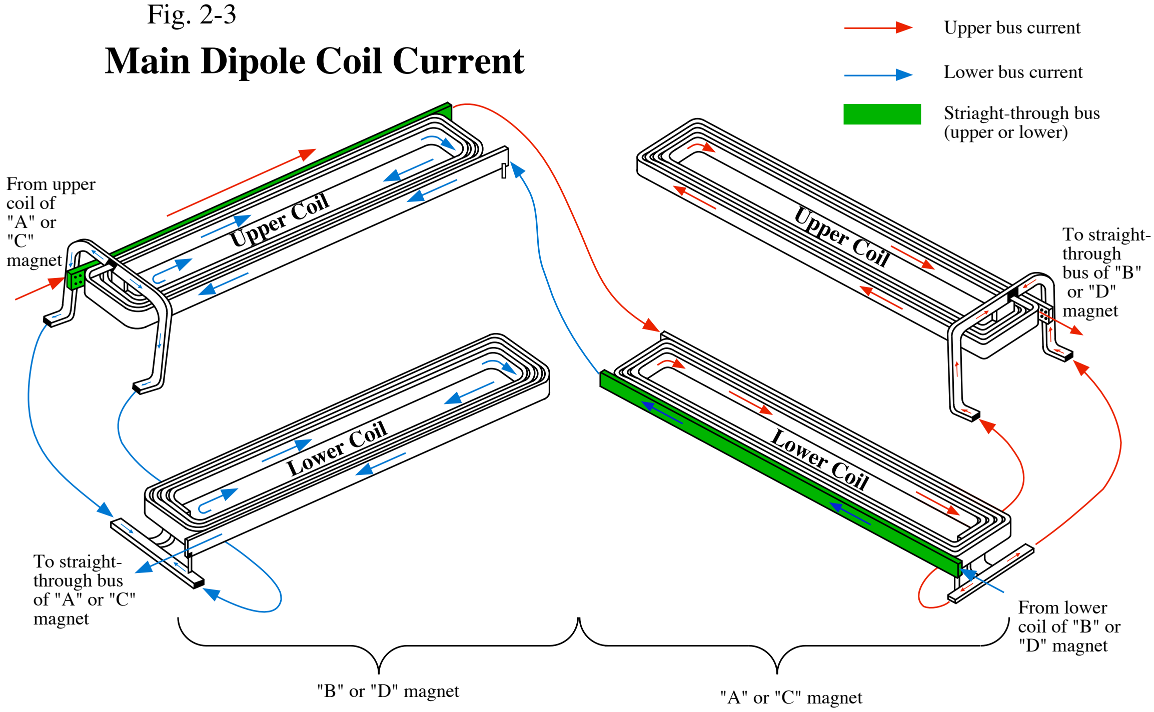

Current flow in a pair of main dipole magnets. Only the coils are shown; the view has been exploded for clarity.

Although the magnetic field "wraps around" the copper coils, it is uniformly pointing downward over the region of the beampipe. Disregarding the effects of the surrounding material, the strength of the field is proportional to the current in the coil. (Here, and throughout this book, conventional current—in which current flows from positive to negative—is represented instead of electron flow, because that is what countless generations of engineers have been brainwashed into using.)

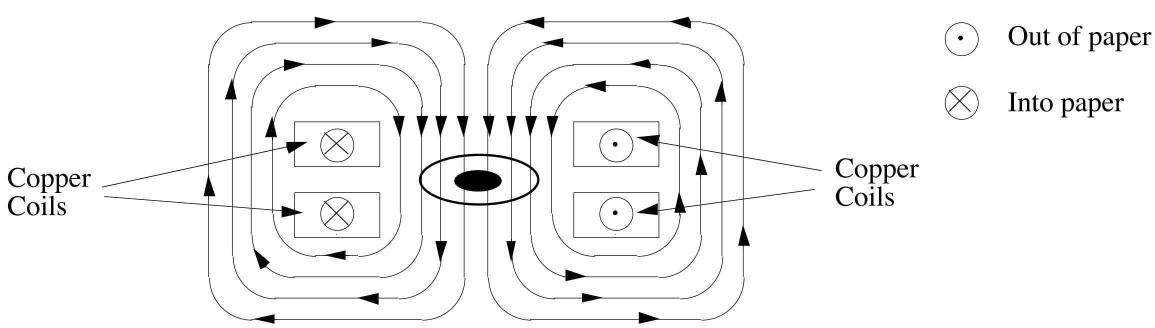

The direction of the force is determined from the right-hand rule, which itself is derived from $q\vec{v} \times \vec{B}$. (A quick-thinking graduate student caught making a gesture to his professor, originally developed the right-hand rule.) Hands are difficult to draw in software, but the directions of the vectors can be shown on three coordinate axes. The force on a proton, with its positive

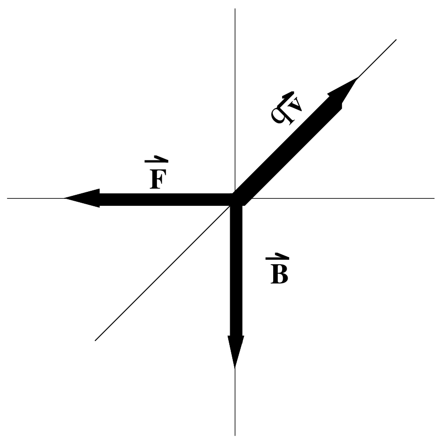

charge, is shown in the figure below:

Here, $\vec{v}$ is pointing "forward" in the proton direction. $\vec{B}$ is pointing down, as in the picture on the previous page. $\vec{F}$ is pointing to the left, so that the protons experience a force pointed towards the inside of the ring. This is how it should be, since it is this force that balances that irritating Newtonian tendency for the protons to move in a straight line and leave the beam pipe.

The force on antiprotons traveling backwards is the same; even though $\vec{v}$ is in the opposite direction, the force is in the same direction because q is negative:

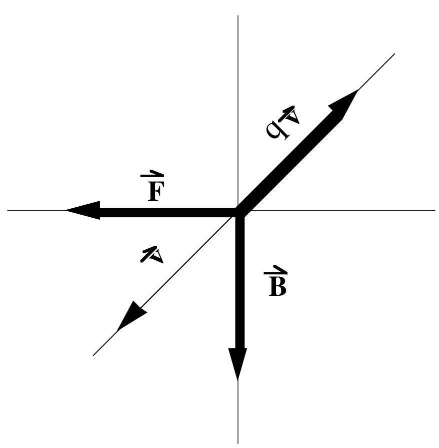

Returning to Fig. 2-2, recall that the ellipse in the center represents the beam pipe. The current-carrying coil is represented by the four sets of tall rectangles, 1” wide and 4” tall. The coil, made of copper, needs to be capable of carrying a peak current of 9400 amps. The current flows in one direction on the right side of the magnet and the other direction when it returns on the left. As the magnetic field “wraps around” the two sets of conductors, the field lines all point in the same direction in the region of the beam pipe.

The number of turns in a coil, and the amount of current in a coil, determines what the field strength will be. An “upper” bus and a “lower” bus (“upper” and “lower” referring to the relative locations of the conductors as they enter the magnets) power the coils. (The term bus refers to the entire length of the conductor, including the coils, the power supplies, and all of the connections between them.)

The upper and lower busses are not actually distinct entities, but “turn around” inside the power supplies at MI-60, as shown in Fig. 2-4. Since current in the upper bus is flowing in the opposite direction from that in the lower bus, the inductive load created by the magnets is balanced. (For those having trouble with that concept, an emergency section on inductance is located at the end of the chapter.) In addition, since current flowing in a loop also creates an external magnetic field that is perpendicular to the loop, having two loops of opposite polarity means that the two fields will cancel each other.

The electrical connections to the Main Injector magnets are a bit more complicated than with some other magnets. To trace the flow of current through an “A” and “B” pair—a “C” and “D” pair would be identical, except for the length—review Fig. 2-3. “A” and “B” magnets are ordered in the proton direction (counterclockwise). As mentioned earlier, when standing in the tunnel and looking at the magnets they are “read” from right to left. “A” will be on the right, “B” on the left.

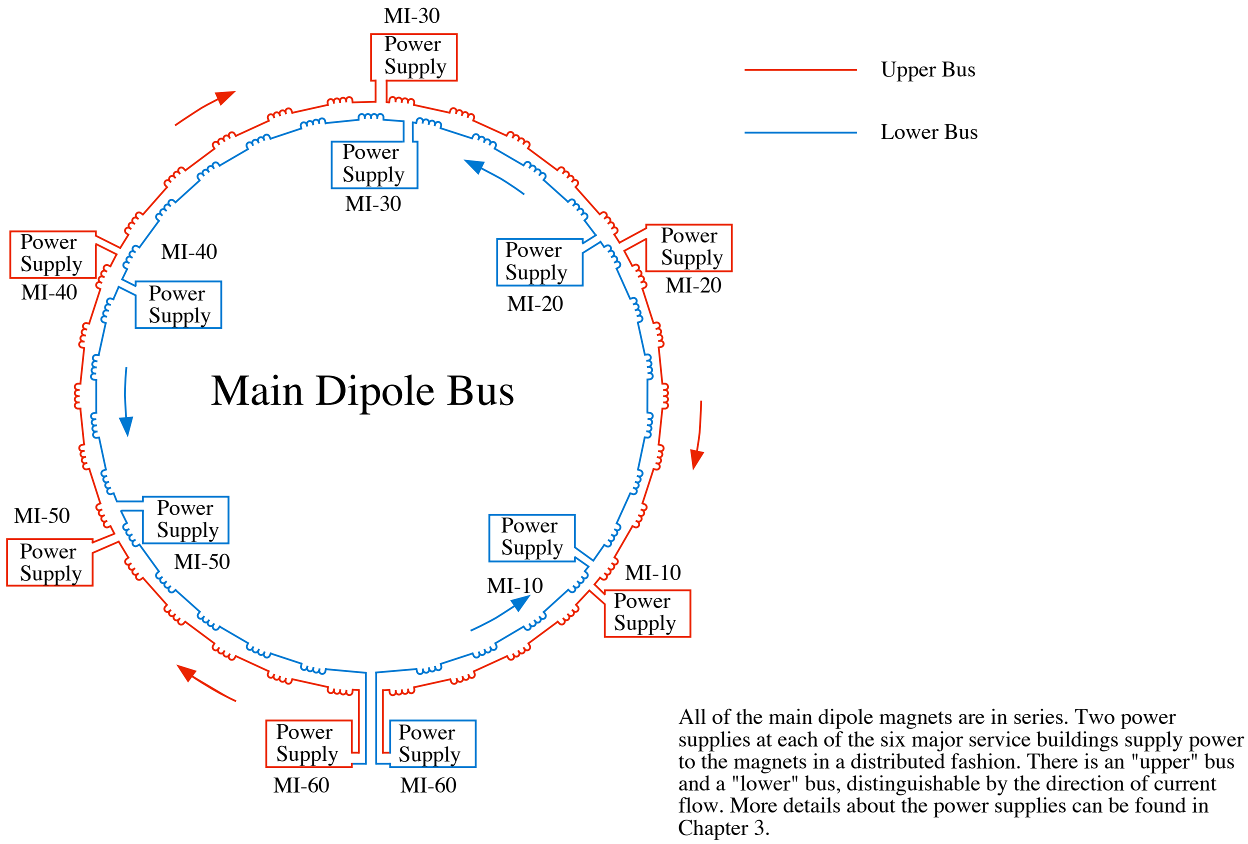

Fig 2-4

The top and bottom coils of each magnet each have four turns, but the main dipole magnets are designed so that there are “straight-through” conductors that only constitute half of one turn. Each bus powers a straight-through section in one magnet and both the top and bottom coils in the other.

Starting with the upper bus, current enters the “B” magnet at the downstream side, on the left. It first powers the straight-through bus at the top coil. Notice that this copper bar is not electrically connected to the rest of the coil, but because of the direction of its current, it contributes to the magnetic field as if it were. The bar emerges at the upstream side of the “B” magnet, where it is connected to the lower coil of the “A” magnet. The lower coil makes 3 1/2 turns, each turn nesting inside the previous one. At the upstream end of the “A” magnet, a braided copper jumper carries the current to the upper coil. The jumper temporarily divides the current into two branches so that there is no net magnetic field at the end of the magnet. Current enters the top coil on the inside turn and completes four full turns, working its way to the outside. Notice that the upper bus enters the “B” magnet at the top and leaves the “A” magnet at the top, making it easily identifiable in the tunnel.

Current from the lower bus enters the straight-through bus of the lower coil of the “A” magnet, adding the final half turn needed to complete the coil. From there it goes to the upper coil of the “B” magnet, powering 3 1/2 turns as it works its way to the inside. The external jumper carries the current to the inside turn of the lower coil. The current works its way to the outside of the coil, and when the four turns are complete it moves on to the next pair.

To summarize: The upper bus powers 7 1/2 turns of the “A” magnet and half a turn of the “B” magnet. The lower bus powers 7 1/2 turns of the “B” magnet and half a turn of the “A” magnet. This arrangement means that the voltage across the pair of magnets is minimized, since the inductive load is shared between the two busses. Moreover, placing the return bus inside the magnets reduces the amount of copper that has to be purchased, as well as reducing the total amount of electrical resistance that the power supplies must overcome.

Look again at Fig. 2-2, and compare it to Fig 2-3. The cross-section is representative of any main dipole; for the following discussion it can be assumed that it is an “A” dipole, looking downstream in the proton direction (counterclockwise in the tunnel). Protons in the elliptical beam pipe are moving into the paper. For the most part, current in the coils is coming out of the paper on the right-hand side, turning around at the reader’s nose, and going into the paper on the left-hand side. The two half-turns that do not turn around are the straight-through bar at the lower left, powered by the lower bus, and the final half-turn at the upper right, which connects to the upstream pair of magnets.

Although the current is continuous in a given bus, the individual turns within the magnet must be electrically insulated from each other. A turn-to-turn short would diminish the magnetic field for at least some of the length of the magnet. Voltages between turns on a given bus can be large, creating the possibility of arcing across the gap, destroying the insulation, and creating a short. The risk is even higher with the straight-through bus and adjacent coils, because the voltage between them can be as high as 2,000 volts. Several layers of fiberglass, impregnated with epoxy, are used to insulate the turns. Extra insulation surrounds the straight-through bus.

The holes in the coil are for the low conductivity water (LCW) used to cool the magnets. The water systems will be discussed in a later chapter.

Laminations

The hatched region in Fig. 2-2 represents the steel laminations surrounding the coils. The steel is of a specific high permeability type that concentrates the magnetic field. The laminations are 1.5 mm thick and are stamped out by the thousands in dies. When building a magnet, they are stacked and placed in a press. The reason a solid piece of steel is not used is that large eddy currents would develop in the presence of the magnetic field. The eddy currents would not only dissipate energy, but would change the character of the field. Small eddy currents still occur within the thin laminations, but they are not nearly as disruptive.

Of course, the laminations must be electrically isolated from each other or there is no point in using them. The ends of the magnets are constructed separately from the center section. Laminations for the end packs of the dipoles are coated in epoxy and heated until the epoxy cures. The epoxy provides cohesion for holding the laminations together as well as electrical insulation between the laminations. Laminations in the center section of the magnet are pre-coated with a plastic shield, which serves as electrical insulation; cohesion is maintained by welding a plate to the laminations while they are still in the press. Should the magnet need repair, it will be easier to remove the welded plate and separate the laminations than to deal with laminations that are glued together.

The reason for building the end packs separately is that in any magnet of finite length, the magnetic field is relatively uniform in the center but becomes geometrically complex toward the ends. The steel laminations, which are the most important factor in determining the shape of the field, must be assembled most carefully in the end packs.

The steel should not be thought of as a continuous entity, even within the laminations, because it is actually made up of innumerable microscopic crystals. Each individual crystal—also known as a magnetic domain—is a permanent magnet with its own magnetic field, but in a magnet that has not been powered the crystals are randomly oriented so that there is no net field. When current begins to flow in the coil, the field created by the coil begins to align the domains and the steel itself begins to contribute to the field.

Although the steel contributes much in the way of field strength and shape, it also introduces some problems. One is hysteresis, or the tendency for magnetic domains to remain as they are. Fig. 2-5 shows the hysteresis curve for a main dipole magnet. The horizontal axis is the "magnetizing force," initially supplied by the coil current, and the vertical axis is the component of the field created by the steel. (The magnetizing force is measured in units known as oersteds. Don't worry excessively about how oersteds are defined, but be aware that the magnetizing field initially comes from the coil current. As the magnetic domains are aligned with the field, they also begin to contribute to the magnetizing force.) The magnetizing force can be either positive or negative, depending on the polarity of the field. Point "A," at the center of the diagram, represents a magnet that has never been energized. There is no current in the coil, no magnetizing force, and no field. As current of the proper polarity energizes the coil, a field develops along with the magnetizing force, and the path from "A" to "B" is traced. This path would be taken when ramping to flattop. Notice that the curve begins to flatten as the magnetizing force reaches large values. This is because nearly all of the magnetic domains have been aligned: a phenomenon known as saturation. In the Main Injector, the effects of saturation begin to appear around 120 GeV.

Also notice that the path is not linear, and that the magnetic field in the dipoles is not always directly proportional to the current.

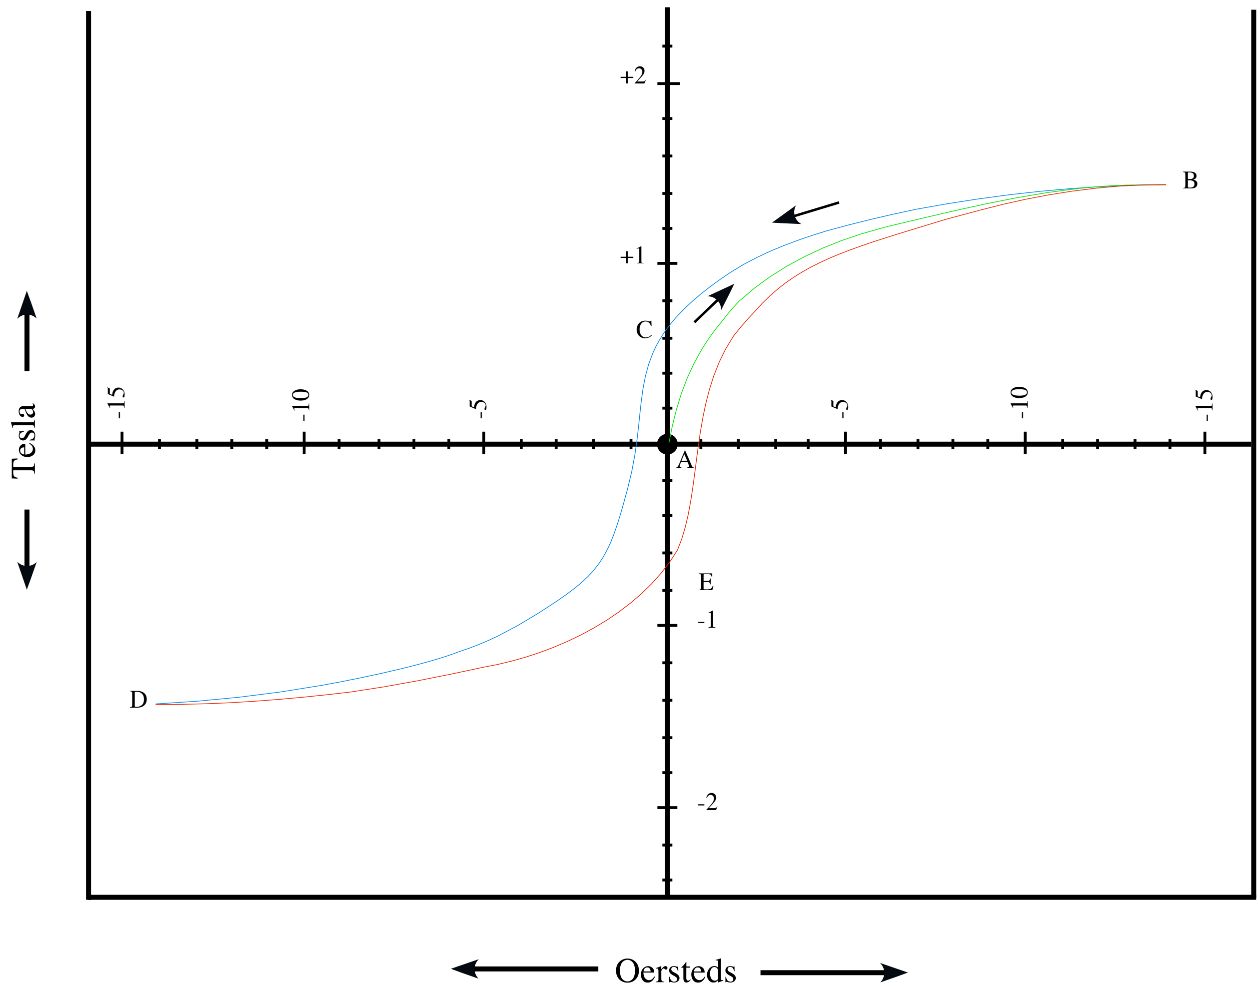

Fig. 2-5
Hysteresis in a Main Dipole Magnet

When the current in the coil is reduced, as when the ramp returns toward the 8 GeV level, the magnetic domains are already aligned; this “fossil” field remains considerably stronger than what would be expected based on the coil current alone. In fact, at point “C,” there is still a significant field left over even when the magnetizing force is zero. Hysteresis is this lack of ability to retrace the magnetization curve due to the history of the magnetic domains.

Hysteresis, saturation, and nonlinearity all have potentially adverse effects on the performance of the Main Injector dipoles. Strategies for dealing with these phenomena are implemented through the software that controls the power supplies.

The steel is grounded and obviously has to be insulated from the coils. A breach of the insulation between the coil and the steel would create a coil-to-ground short. Coil-to-coil and coil-to-ground shorts were all too frequent in the Main Ring, creating many hours of downtime for every failure. It is hoped that the Main Injector will be more robust.

Main Quadrupoles

The dipoles perform the necessary function of bending the beam around the curvature of the ring, but that is not sufficient to keep the beam inside the beam pipe. The beam consists of billions and billions of particles; in addition to their longitudinal motions, each one has its own unique transverse motion. Without some kind of focusing, the particles would spread out and within a few dozen feet would be lost from the machine.

Focusing the beam is done with quadrupole magnets. Quadrupoles, as you might guess, have four poles. The main quadrupoles, which are similar in cross-sectional size to the main dipoles but somewhat shorter, will be described in this section. There are also smaller quadrupoles, which will be described later.

The main quadrupoles in the Main Injector come in three lengths: 84", 100", and 116," but are virtually identical in cross-section (Fig. 2-6).

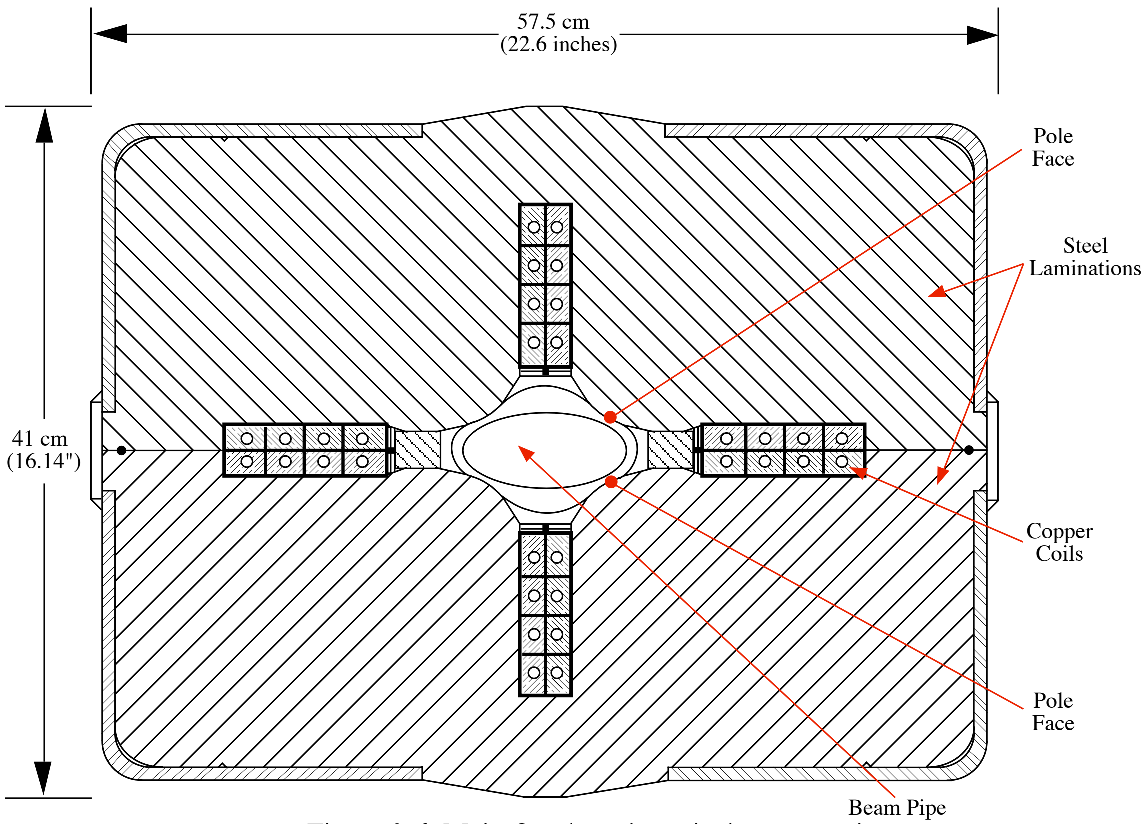

Figure 2-6 Main Quadrupole typical cross-section

As with the large dipoles, the beam pipe is the elliptical tube in the center. Current passes through the copper coils, square in cross-section, and cooling water passes through the holes in the coils. The steel laminations shape and concentrate the magnetic field.

Quadrupole Fields

The direction of current in a quadrupole alternates between adjacent poles. For the purpose of discussion, the direction of current in the two side coils can be arbitrarily chosen to be coming out of the paper. In the top and bottom coils current is going into the paper. The direction of the fields can be visualized, as for a dipole, by grasping the coil with the right hand:

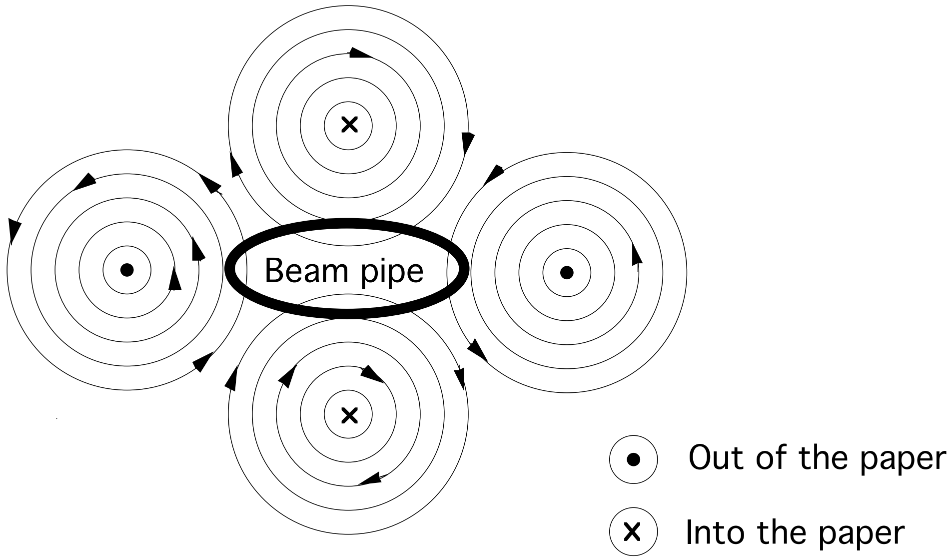

Another look at Fig. 2-6 shows that the iron laminations, also known as pole faces, are designed to capture and shape the field into the four regions between the coils. The predominant directions can be summarized as in the picture on the next page:

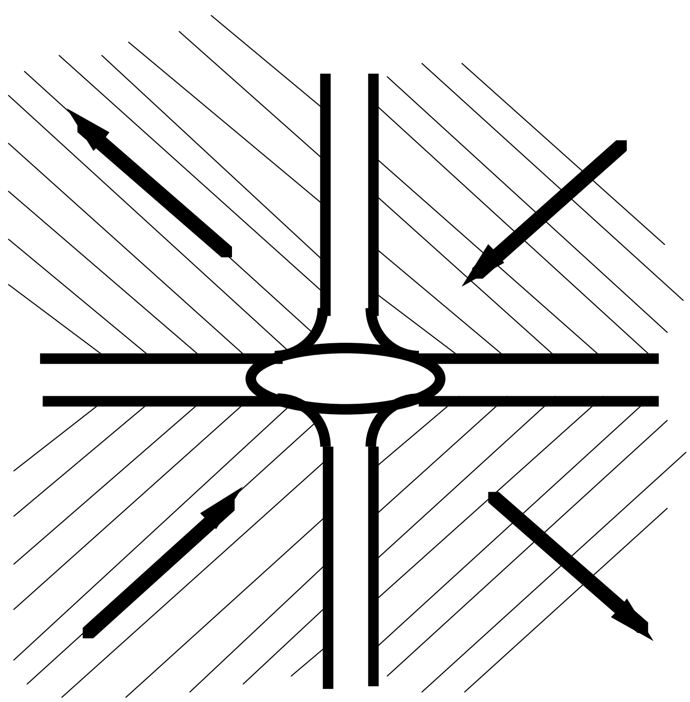

The field directions shown above can be resolved into horizontal and vertical components, noticing that the fields are consistently to the left on the top of the magnet, to the right on the bottom, up on the left and down on the right:

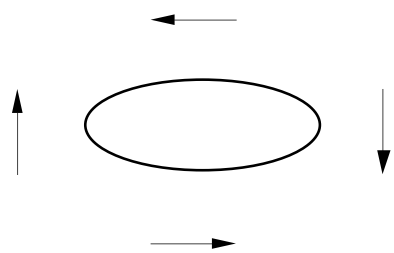

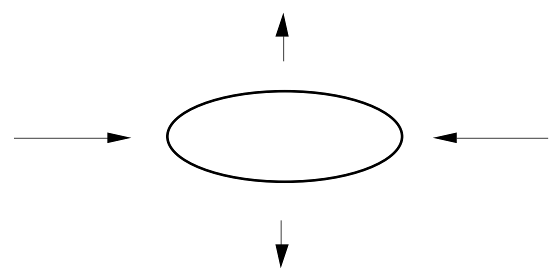

Now apply the right-hand rule to determine the direction of the forces:

The particles on the left will be pushed to the right, and particles on the right will be pushed to the left. A particle at the center will experience no correction (nor should it, since it is at the desired position). This quadrupole is horizontally focusing; horizontally focusing quads are also known as "F" quads.

If the field strength is carefully accounted for along each point on the horizontal axis, it turns out that the restoring force is linearly proportional to the horizontal distance from the center. A particle twice as far from the center will receive twice the correction.

This sounds great, but note that there are also vertical forces present that point away from the center, "defocusing" the beam vertically. These vertical forces are also linearly proportional to the vertical distance from the center.

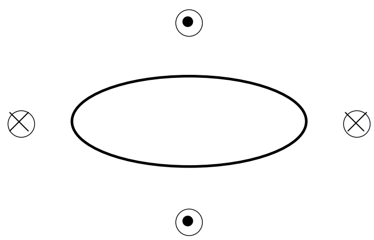

To keep the beam focused vertically, the next quadrupole in the sequence is connected with the opposite polarity.

The conscientious reader will take out a piece of paper and determine for herself that this quadrupole is vertically focusing. There is a restoring force vertically and a defocusing effect horizontally. Unfortunately, perhaps, this kind of quadrupole is usually referred to by its horizontally defocusing effects and is called a defocusing or “D” quad (not to be confused with the “D” dipoles).

The focusing and defocusing quadrupoles of a given type are identical in construction, but each is connected to a dedicated bus and set of power supplies (Fig. 2-7). Unlike the main dipoles, both leads to the magnet are connected to the same end, and there is no straight-through bus.

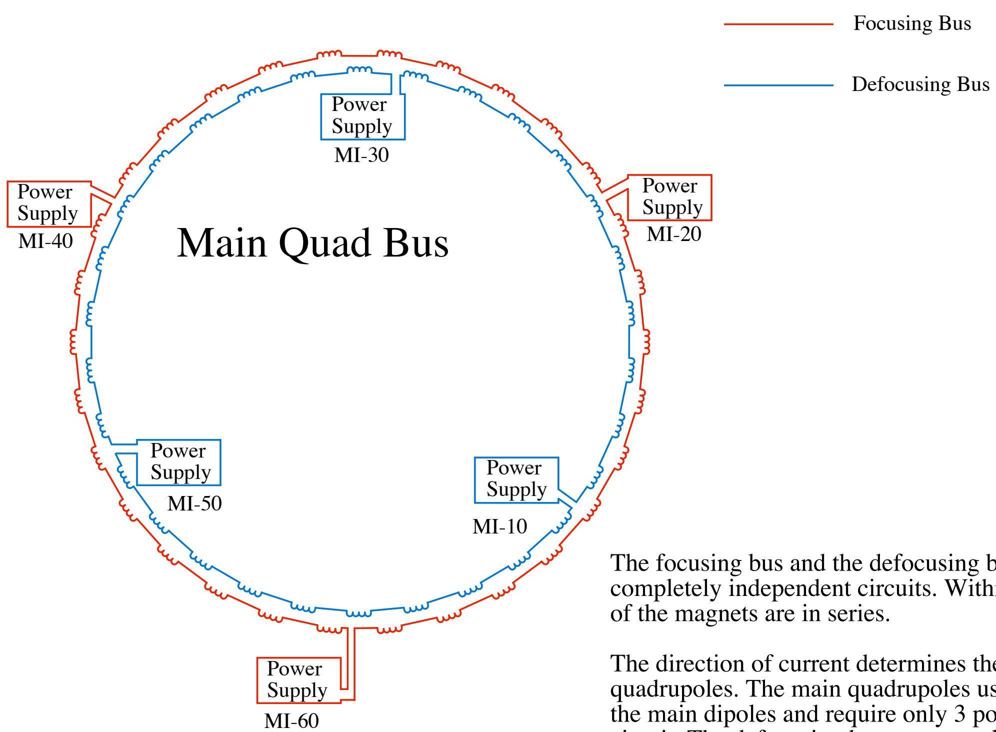

The focusing bus and the defocusing bus are completely independent circuits. Within each circuit all of the magnets are in series.

The direction of current determines the polarity of the quadrupoles. The main quadrupoles use less power than the main dipoles and require only 3 power supplies per circuit. The defocusing bus runs at a slightly lower current than the focusing bus.

Fig. 2-7

Ring-wide, the current in the focusing bus flows clockwise, while current in the defocusing bus flows counterclockwise. It will be seen shortly that the two currents are not quite the same, but to a first approximation the external magnetic fields are cancelled as they are for the dipoles. The "F" quads are connected to adjacent "F" quads (skipping over the "D" quads between them) and the "D" quads are likewise connected to each other.

The Main Injector Lattice

The previous sections in this chapter have described the main quadrupoles and dipoles. To understand the big picture requires some knowledge as to how these magnets are combined to form the lattice. The lattice is intended not as an analogy to the leafy vegetable, but rather to the regular and predictable pattern of atoms in a crystal. The main dipoles and quadrupoles are the "skeleton" around which the other components in the accelerator are arranged. The main quadrupoles actually define the lattice.

As discussed in the last section, there is "F" quads (horizontally focusing) and "D" quads (vertically focusing). The trick now is to space the quads at a distance that minimizes the defocusing effect in both planes. For example, in the "F" quads, the beam distribution looks like this:

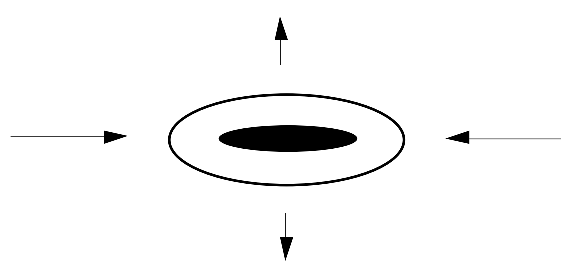

The beam is wide horizontally and narrow vertically. Since the horizontal and vertical forces are proportional to the distance of the particles from the center, the horizontal forces will predominate.

A single quadrupole magnet acts much as an optical lens would, except that focusing only takes place in one plane:

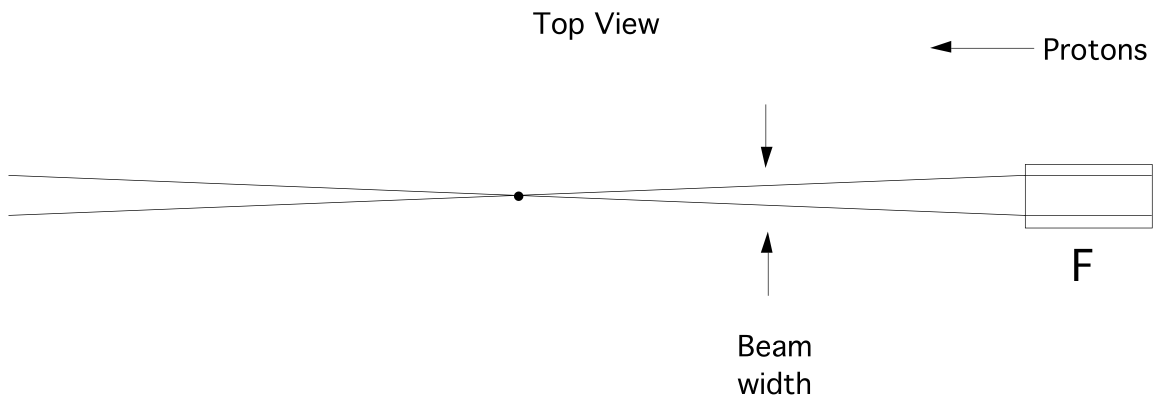

For any given quadrupole, there is only one place at which the beam size is at a minimum, beyond that point it begins to diverge again. In the lattice, the defocusing quads are placed where the horizontal beam size is the smallest, and vertical beam size is the largest; the focusing quads are placed where the horizontal beam size is the largest and vertical beam size is the smallest.

The vertical beam size is one half-cell out of phase with the horizontal:

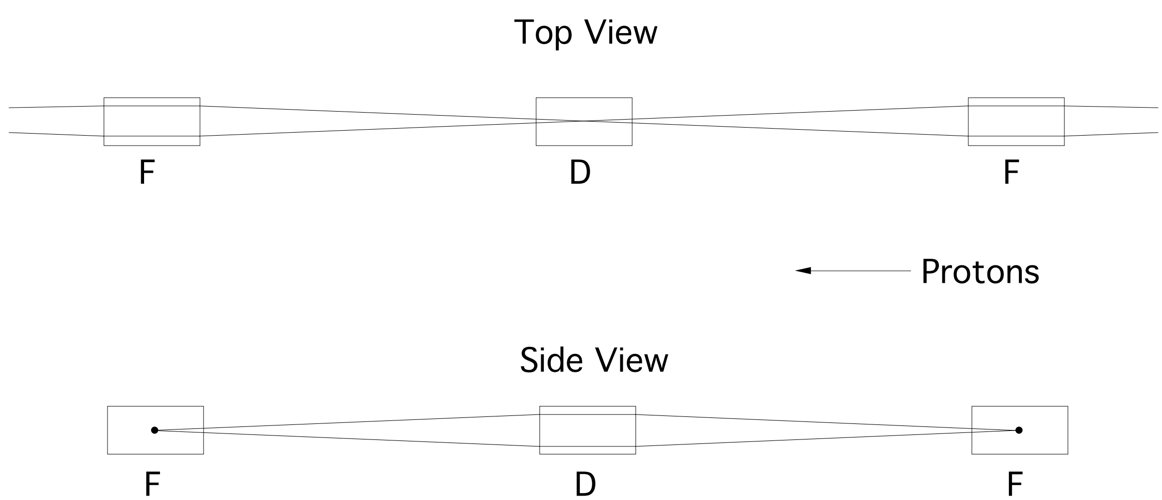

Again, for both the horizontal and vertical planes, the beam is intercepted and focused where it is the widest. In the Main Injector, the "F" quads and "D" quads alternate, creating what is known as a "FODO" lattice. The "O" designates the drift space between the quads. The drift space is often occupied by dipoles for bending the beam, but other devices or even unadorned beam pipe are valid options.

A typical section of the Main Injector FODO lattice is shown below:

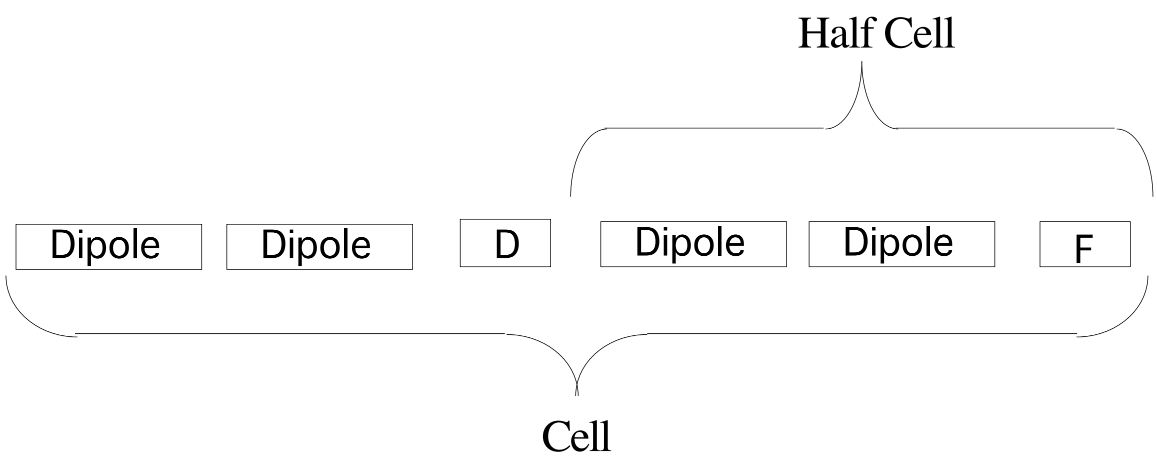

The length of the lattice shown is called a cell. The cell is the basic, repeatable unit of the lattice, in this case the FODO lattice. In this particular cell, each drift space contains two large dipoles. In the arcs between the straight sections, each cell contains six magnets: two quadrupoles and four dipoles. Where the main dipoles are present they are paired "A" and "B" or "C" and "D," as described earlier. There are no main dipoles in the straight sections (described below), but the number of quadrupoles per cell remains constant.

A half-cell is the distance that includes one quadrupole and the adjacent drift space. A half-cell is not repetitive because the polarity of the quadrupole changes in the next half-cell.

The alert reader will deduce that the longer dipoles will bend the beam in a tighter curve, and that the longer quadrupoles will focus the beam more tightly.

Since there is a main quadrupole at the beginning of every half-cell, and since they alternate between vertically and horizontally focusing quads, the naming convention has been adopted that every focusing location is assigned an even number, and every defocusing location is given an odd number. The other components in the neighborhood of the quadrupole incorporate the number as well, so any component with an even number can usually be assumed to be at a focusing location, and vice-versa.

Straight Sections

The Main Injector ring, as shown in Fig. 2-1, has a shape resembling that of a battered egg. The unevenness is due to the presence of straight sections at 8 locations. The straight sections provide space for specialized functions, usually for beam transfer to or from the Main Injector. Straight sections can be purchased in short, medium, or large sizes.

Even with the odd shape shown in Fig. 2-1, the symmetry of the lattice begins to emerge. The short straight sections at MI-52 and MI-62 (designed for beam transfers to and from the Tevatron, Antiproton Source and Fixed Target experiments) are mirrored by identical short straights at MI-22 and MI-32. The latter two sites were originally included only to preserve the symmetry of the ring, but have now been pressed into service as points of transfer to and from the Recycler. MI-10, a medium-length straight section, is the point where 8 GeV beam enters the machine; its counterpart at MI-40 has been chosen as the location of the abort line. The long straight section at MI-60—a very busy place with multiple beam transfers, the RF accelerating systems, and the extraction line to NuMI—is mirrored by MI-30, which is virtually empty.

It is the location of the straight sections that determines the distribution and strength of the main dipoles and quads. A careful look at Figure 2-1 reveals the two sizes of the magnets and where they are located. The "normal" arcs are comprised of the longer "A" and "B" magnets, shown as dark blue in the diagram. The shorter "C" and "D" magnets, in light blue, act as bridges between the arcs and the straight sections. When traveling through the dipoles the beam is essentially tracing part of a circle.

This arrangement of long and short magnets was made necessary by a phenomenon called dispersion. Dispersion is the tendency for low energy particles to be bent more than high-energy particles as they trace a curve through a dipole. The straight sections are designed as "zero dispersion" regions in which particles are not sorted out by their differences in energy. The "C" and "D" dipoles, and the quadrupoles associated with them, help to ease the beam into and out of the zero dispersion straight sections. The two cells on either side of each straight section are known as "dispersion suppressor" cells:

Fig. 2-8 represents the three basic types of straight section in the Main

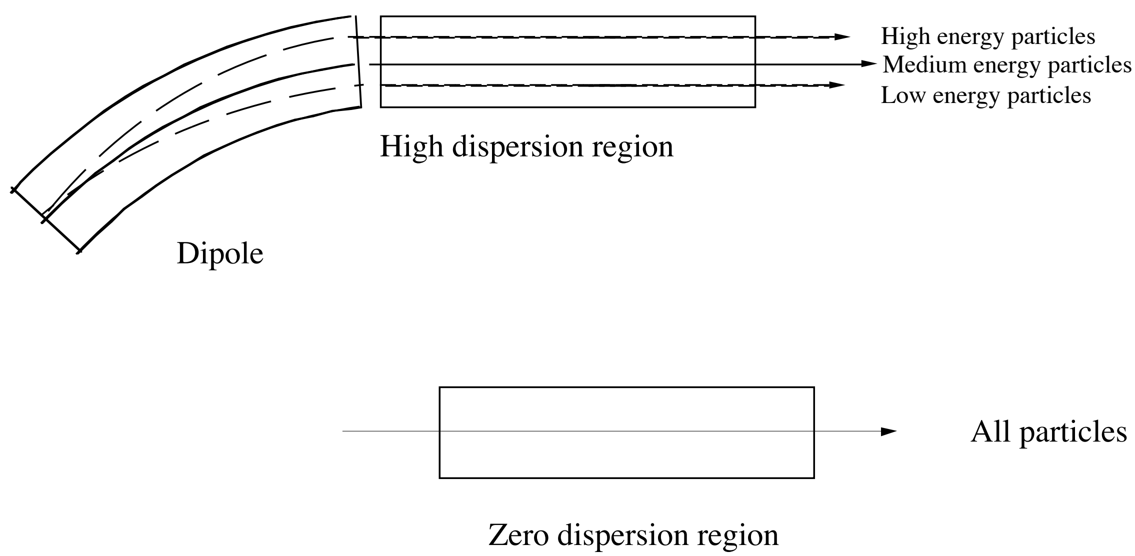

Injector. Fig. 2-8(a) represents the four short straight sections at MI-22, MI-32, MI-52, and MI-62; Fig. 2-8(b) the medium straight sections at MI-10 and MI-40, and Fig. 2-8(c) the long straight sections at MI-30 and MI-60. In each diagram, a part of the normal arc, with its "A" and "B" dipoles, can be seen to either side. Between each pair of dipoles in the normal arcs is an 84" quad.

The two dispersion-suppressor cells on either side of the straight section are made up of “C” and “D” dipoles, with 116” quads sandwiched between each pair. The boundary to either side of the dispersion-suppressor region is interfaced with a 100” quad. Within the straight sections themselves, the FODO lattice continues in the form of 84” quads.

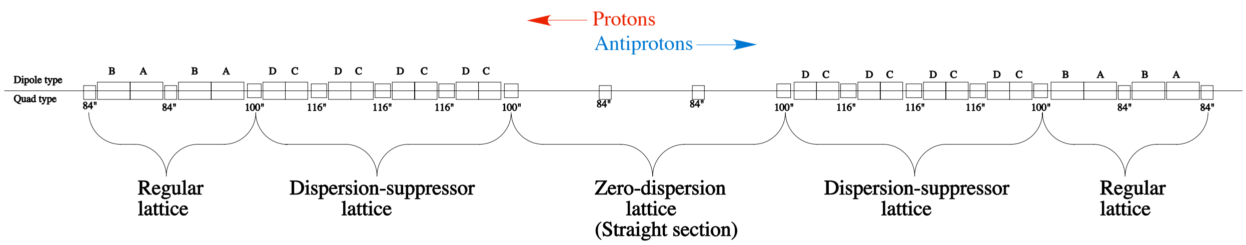

Fig. 2-8(a) Short Straight Sections
MI-22, MI-32, MI-52, MI-62

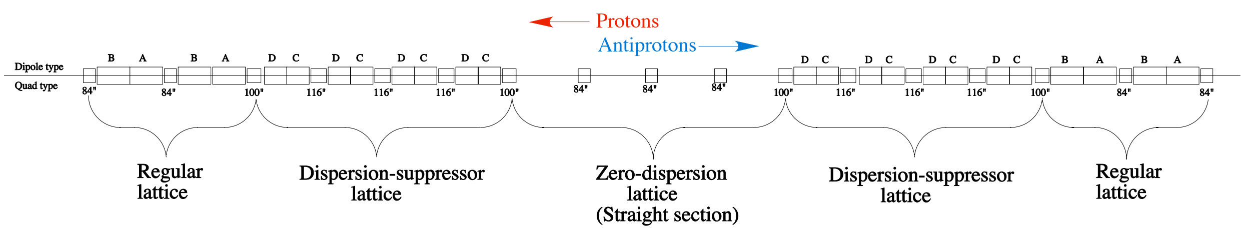

Fig. 2-8(b) Medium Straight Sections
MI-10, MI-40

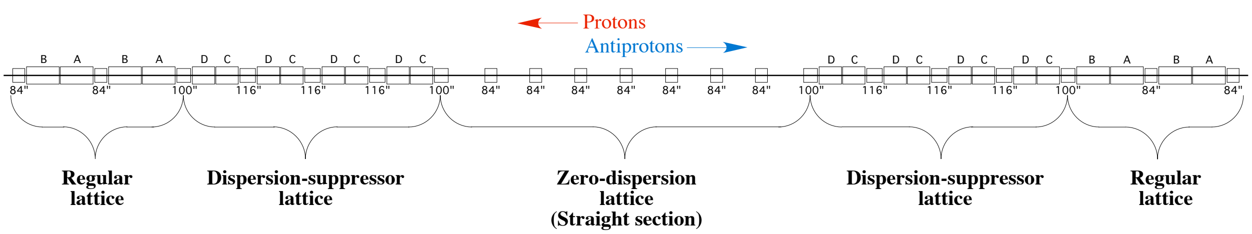

Fig. 2-8(c) Long Straight Sections
MI-30, MI-60

The various quadrupole lengths are necessary to match the changing bend field as the beam enters the zero dispersion areas. The curvature of the normal arcs and the dispersion-suppressor arcs is about the same, although the “C” and “D” dipoles are shorter than the “A” and “B” dipoles, there are more of them over any given length because the cells are shorter. However, the overall focusing strength is greater because all of the quadrupoles in the suppressor arcs are longer than in the normal arcs, and they are closer to each other, and because the cells are shorter. The fact that the ratio of focusing strength to bending strength is larger is what makes the suppressor cells work.

Figs. 2-8 only illustrates the basic lattice of the straight sections; in reality, much of the unused space is filled with magnets and other devices utilized for beam transfer, but the pattern of main dipoles and main quadrupoles does not change.

Just as each straight section has a symmetrically placed partner across the ring, it has an internal symmetry with respect to its own center. (Remember that magnetically there is no difference between magnets of a pair.) Circulating antiprotons are going to see the same sequence of magnetic fields as the protons.

Beta Functions and Tunes

Mathematically, the periodic widening and narrowing of the beam is described by a beta function. The beta function is only one of several lattice functions; dispersion is another. Fig. 2-9 shows the beta functions and dispersion through a typical straight section.

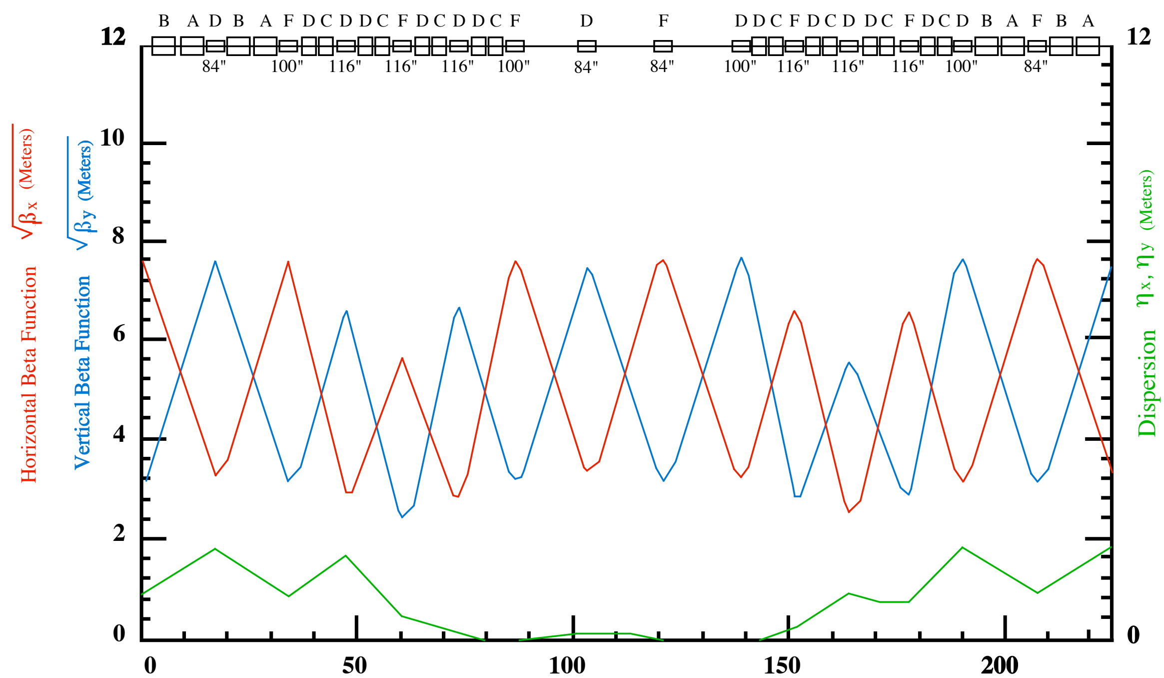

Main Injector lattice functions in the vicinity of a typical straight section; the red and blue traces represent the horizontal and vertical beta functions and the green trace represents the dispersion. The beta functions show that the beam is widest horizontally at the focusing quads and widest vertically at the defocusing quads. The dispersion goes to zero in the straight sections.

The particular lattice shown in the diagram is that of MI-32 or MI-62, but all of the straight sections are similar.

Fig. 2-9 Main Injector Lattice Functions

The location of the straight section itself, occupying the center third or so of the diagram, can be recognized by the zero dispersion.

The beta functions in this diagram for both the vertical and horizontal planes are superimposed on each other; remember that the horizontally the beam is widest at the “F” quads and the vertically it is widest at the “D” quads.

Beam width is intuitively a straightforward concept, but quantitatively it is rather complex. Notice that the vertical axis is labeled with $\sqrt{\beta_x}$ and $\sqrt{\beta_y}$, which are actually the square roots of the horizontal and vertical beta functions. These terms require about 20 years of graduate school to understand; let it suffice for now to say that the bigger the beta function, the wider the beam.

It would be logical to assume that the individual particles oscillate in exactly the same manner as the beta function. That assumption would be incorrect. Within the beam envelope, defined by the square root of the beta function, individual particles undergo an oscillation that is about four times slower than that of the beta function. The slower motions are known as betatron oscillations. Getting confused yet? The apparent discrepancy comes from the fact that the individual particles are not obliged to all start out with the same position and angle of deflection; the billions and billions of particles in each bunch each have their own trajectory. The beam envelope of the beta function defines the outer limits of all their motions.

Any given particle, as it moves through the quadrupoles of the lattice, is unlikely to be exactly on the desired orbit; it will be somewhat off-center horizontally and vertically. In either plane, as it receives a restoring force from a quadrupole, it receives a push toward the desired orbit position. It will then continue to drift past the desired position until it is on the opposite side of where it started. Its next encounter with a quadrupole of the proper polarity will push it back toward the desired position. It continues oscillating back and forth around the desired orbit as long as it remains circulating in the machine. This progression, which resembles a bumpy sine wave as the particle moves through the lattice, is referred to as phase advance; the Main Injector lattice is designed with a phase advance of about 90 degrees per cell. A particle returns to its original phase after traversing about four cells. If the field in the quadrupoles is strengthened, the oscillations are faster and take place over a shorter distance. The total number of betatron oscillations in a single revolution of the beam is called the tune, defined separately for the horizontal and vertical planes. Increasing the tune is equivalent to increasing the amount of current in the quadrupoles.

There are 104 cells altogether in the Main Injector ring: 104/4 = 26, so one would expect the tunes in Main Injector to be 26. In actuality, the Main Injector was designed for a horizontal tune of 26.425 and a vertical tune of 25.415. (The fact that the vertical tune is lower than the horizontal tune means that the current in the vertical bus is lower than the current in the horizontal bus.) If an exact integer tune were selected, a given particle would see the same tiny field errors once every revolution; these errors would quickly add up and the motion of the particle would become unstable. The same is true of half-integer tunes, say, 26.5, where the particle would see the same errors every other revolution. In fact, several easily visualized fractions, known as resonances, must be avoided if beam is to stay in the accelerator. Applications programs are available to adjust the current in each of the two quad busses to optimize the tunes; more about these in the chapter on power supplies.

During resonant extraction to fixed target experiments, the tunes are deliberately pushed toward a resonance in order to cause them to leave the machine in a controlled way. Perhaps someday there will be a chapter on resonant extraction.

A more general and mathematical treatment of tunes can be found in the "Concepts" book. Of particular relevance here are the chapters on "Beam Characteristics" and "The Lattice."

Corrector Dipoles

The main dipole magnets carry the lion's share of the work at keeping the beam on the correct orbit within the beam pipe. They are all in series with each other, tied to the same set of power supplies, and have identical currents at any given instant. However, no accelerator is perfect, and it is inevitable that numerous local adjustments will be needed due to the small differences in the construction of the individual magnets. The correction dipoles are puny compared to their multi-ton counterparts; the horizontal correctors weigh only 321 pounds apiece. The advantage to using them is that each one is individually controllable. They come in horizontal and vertical flavors.

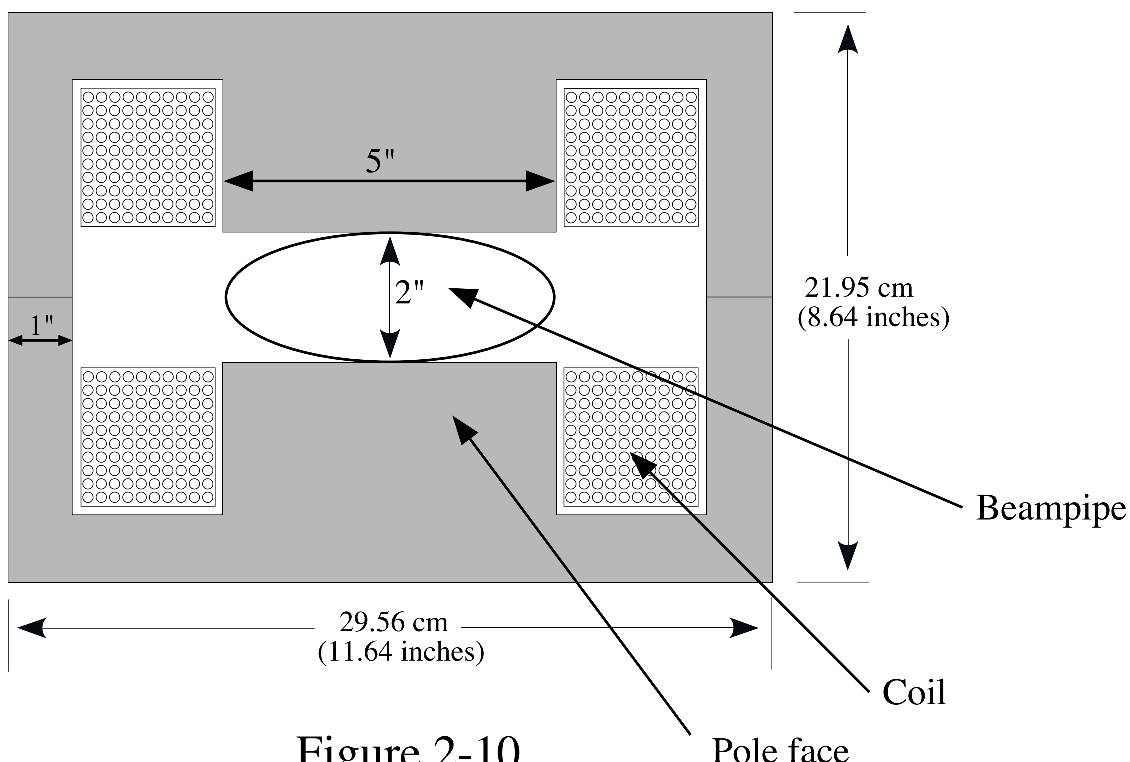

Figure 2-10
Horizontal Corrector Dipole

The horizontal correctors are just under 9" high and 12" wide, and are 17" long (Fig. 2-10). Like the large dipoles, the field originates in the coil current; unlike the large dipoles, the coils are formed from 400 turns of relatively thin (#10) wire. The top and bottom halves of the magnet are fabricated separately and must be bolted together around the beam pipe in the tunnel.

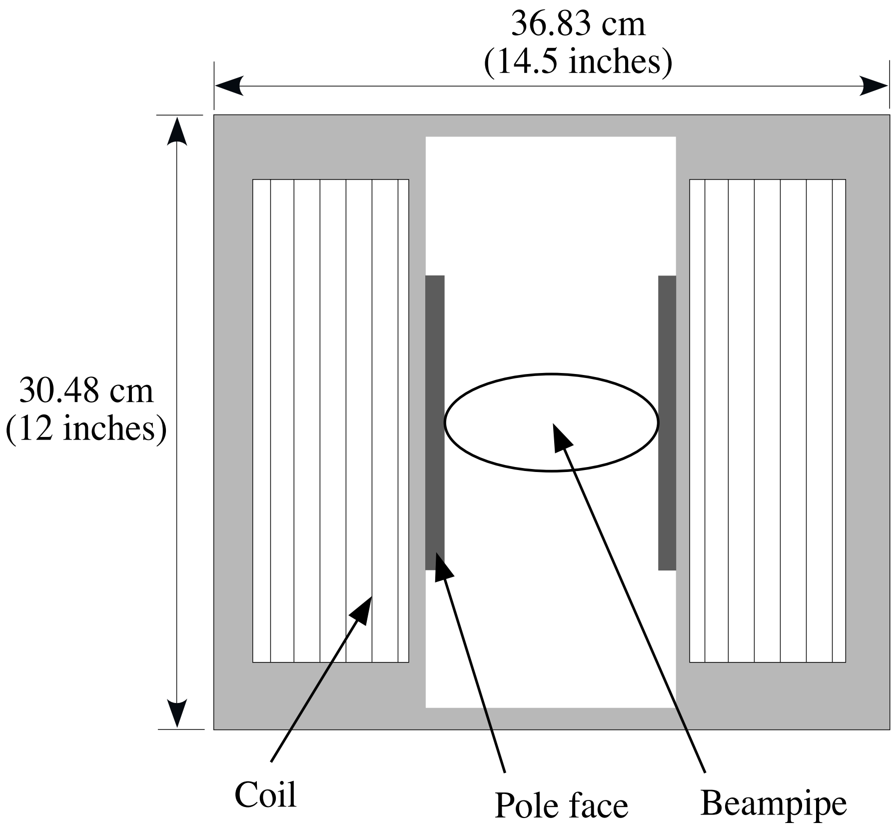

Figure 2-11
Vertical Corrector Dipole

The vertical corrector dipoles, since the field itself is horizontal, require that the left and right halves be built separately and then bolted together downstairs (Fig. 2-11).

Normally, there is a horizontal corrector dipole in front of each main focusing quadrupole and a vertical corrector in front of each main defocusing quadrupole, although there are occasional exceptions in the straight sections. Remember that beam is widest horizontally at the focusing locations and widest vertically at the defocusing locations. Changing the current in, say, a horizontal magnet will have the greatest effect on beam losses at the next horizontal magnet, where the lever arm is the greatest; the chances of scraping beam against the beam pipe is also the highest because of the beam size.

In practice, most changes to the corrector magnets come in the form of 3-bumps. A 3-bump is typed into a parameter page and knobbed. The upstream magnet changes the position of the beam at the point where losses are high, and then the magnet at the point of interest bends the beam in the opposite direction to get it back to its original position. Finally, the downstream magnet bends the beam to adjust the angle. If done properly, the bump is completely "local" and does not change the orbit at any other point in the ring. This miracle usually works because the software on board the parameter page knows the phase advance between the correctors, and can calculate the needed currents.

Skew Quadrupoles

There are 16 skew quadrupoles in the Main Injector (Figs. 2-12, 2-19). Their purpose is to control the amount of coupling between the vertical and horizontal tunes. When the tunes are coupled, any attempt to change the horizontal tune will also affect the vertical tune, and vice-versa. Reducing the coupling allows the tunes to be adjusted more independently.

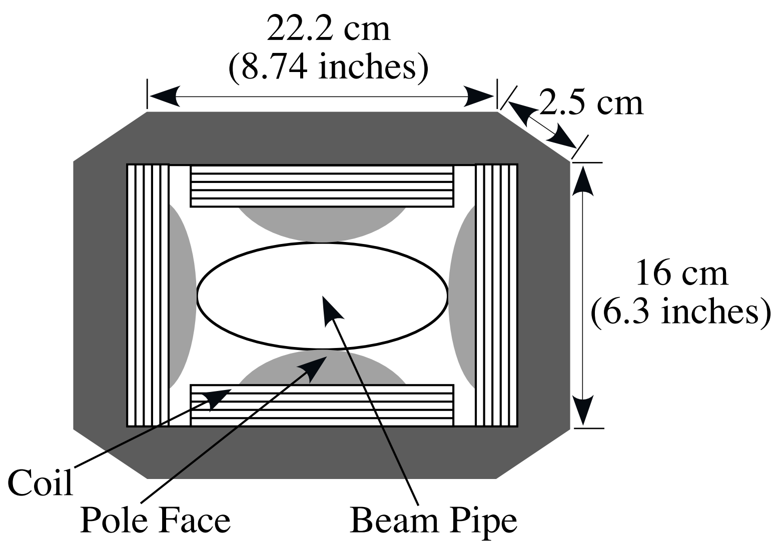

Figure 2-12
Skew quad

Unlike the other quads, both large and small, the pole faces have been rotated by 45° so that they are parallel to the horizontal and vertical planes; the skew quads look "normal" and the "normal" quads look skewed. They are distributed evenly between focusing and defocusing locations (Fig. 2-20), and clustered into four groups of four.

Trim Quadrupoles

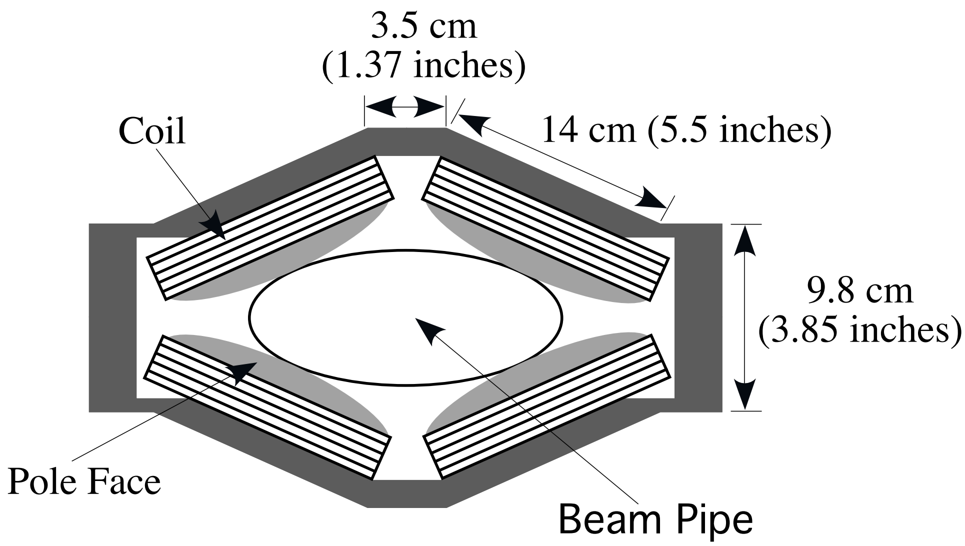

Figure 2-13
Trim Quad

There are also 16 trim quadrupoles in the Main Injector (Figs. 2-13, 2-21). Unlike the magnets already discussed, they all focus in the horizontal plane only. They are used during resonant extraction, which is a process that affects the beam horizontally. They can also be used as harmonic correctors, available for adjusting specific characteristics of the betatron oscillations, such as phase advance.

The trim quadrupoles are recycled from the Main Ring. The poles are oriented at 45˚ to the horizontal and vertical axes, as with the large quadrupoles. They are, along with the skew quads described above, among the smallest magnets to be found in the Main Injector. Like the horizontal correction dipoles and the skew quads, the top and bottom halves must be bolted together around the beam pipe. Like the skew quads, they are clustered into four groups of four.

The role of the trim quadrupoles in manipulating the beam will be discussed in the hypothetical chapter on resonant extraction.

Sextupoles

Sextupoles, which by definition have six poles, are used to control a phenomenon called chromaticity. The word for "chromaticity" is related to the word for "color," and is used in optics to describe lenses for different wavelengths of light that have different focal lengths, as in the picture below.

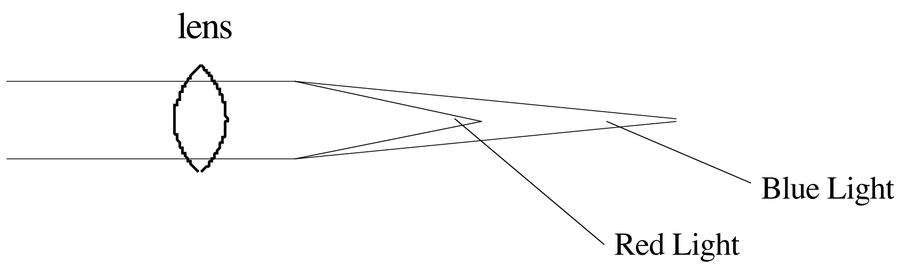

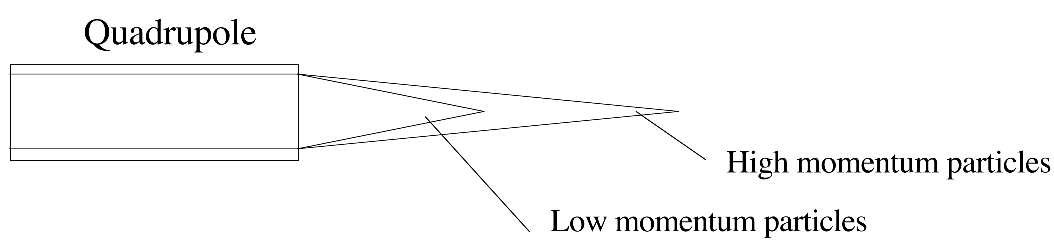

The separation of colors occurs because blue light has more energy than red light and is therefore more difficult to bend.

The concept of chromaticity in accelerators is intended as an analogy to what happens in optical systems. Just as light comes in a variety of energies, a group of particles in an accelerator has a range of energies. (Since accelerator physicists often speak of momentum rather than energy—momentum is the kinetic energy of a particle added to its rest mass—momentum will be used rather than energy in following discussion. The momentum of particles traveling at relativistic speeds is expressed in units of GeV/c; quantitatively, the numbers come out as the kinetic energy of the particle (in eV) plus the rest mass of the particle. The rest mass of a proton or antiproton is about 938 MeV.) Because the RF concentrates the beam in bunches, the range in momentum is fairly small: dp/p, the range in momentum spread divided by the total momentum, tends to be near 10^{-3}. Nevertheless, if an accelerator has significant chromaticity, even that level of momentum spread will create problems with the tune. Higher momentum particles, being harder to bend, will have longer focal lengths as they go through the quads, and will have a lower tune because they undergo fewer oscillations per revolution. A certain tune spread is acceptable, but if the tune spread gets to be too large some of the particles will begin to approach a resonance and they will be lost from the machine.

The equation defining chromaticity is:

$$\frac{\Delta\nu}{\nu} = \xi \frac{\Delta p}{p},$$

where ν is the tune, ξ is the chromaticity, and p is the momentum. In other words, the tune spread is proportional to the momentum spread; the chromaticity is the constant of proportionality.

A sextupole magnet (Fig. 2-14) compensates for the chromaticity by applying different forces to particles of different momenta, so that the variation

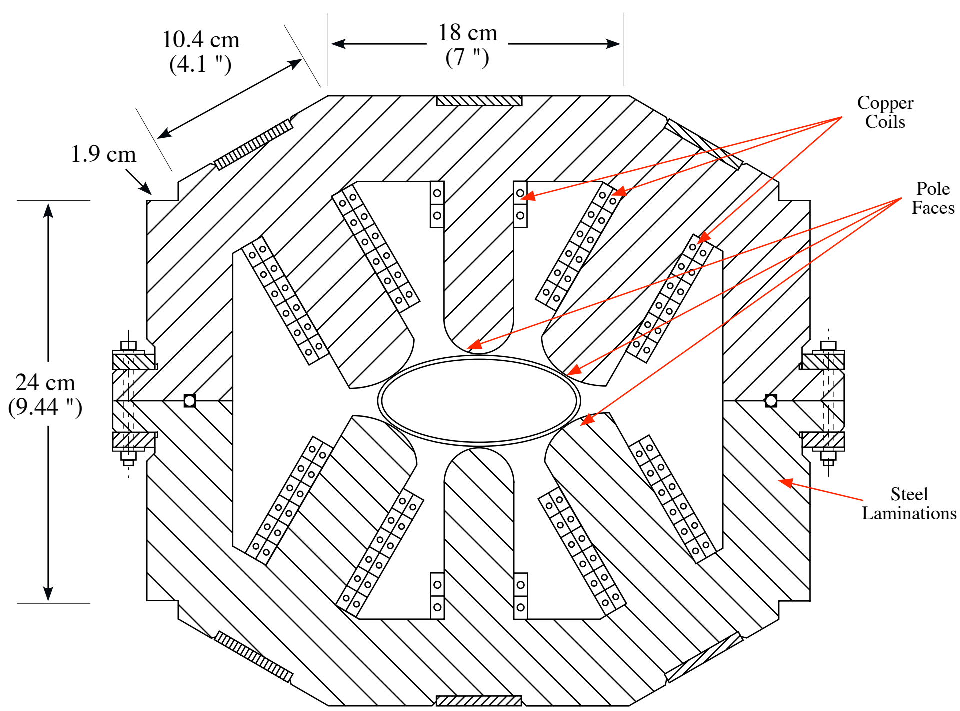

Figure 2-14 Cross-Section of Main Injector Sextupole

in focal length, and therefore the variation in tune, can be controlled.

A copper coil accompanies each pole face, the coils appearing as pairs in cross-section. As with the other magnets, the current in the coils magnetizes the steel pole faces; the steel concentrates and shapes the field.

The current flowing in a horizontally focusing sextupole looks something

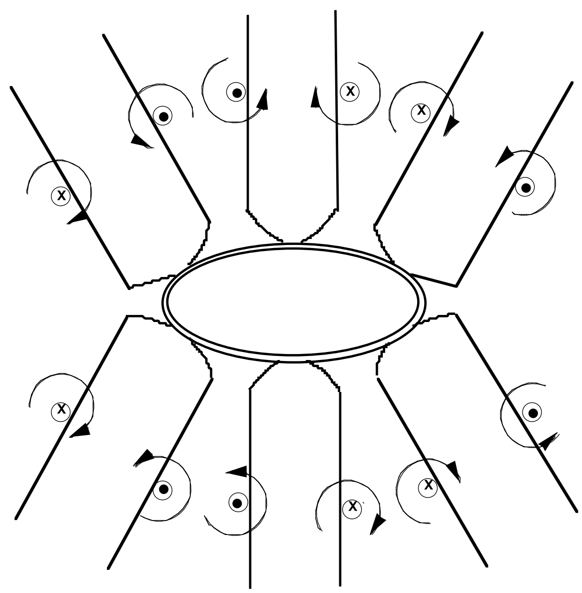

like this in cross-section below.

The direction of each field in each pole face is either toward the beam pipe or away from it, and can be summarized by the picture below.

When the right-hand rule is applied and the restoring force is plotted as a function of the distance from the center, a parabola results:

Algebraically, a parabola is expressed as $y = ax^2$, where "a" is a constant. The x^2 term is the reason that sextupoles are considered “second-order”

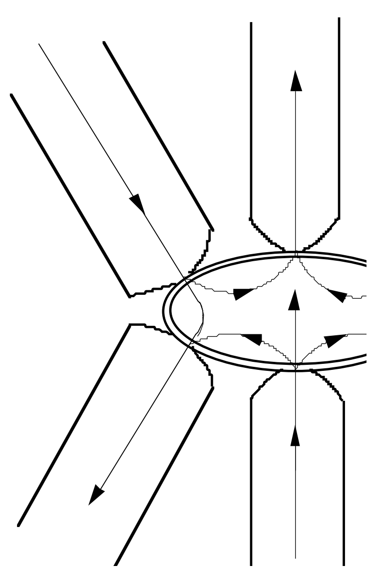

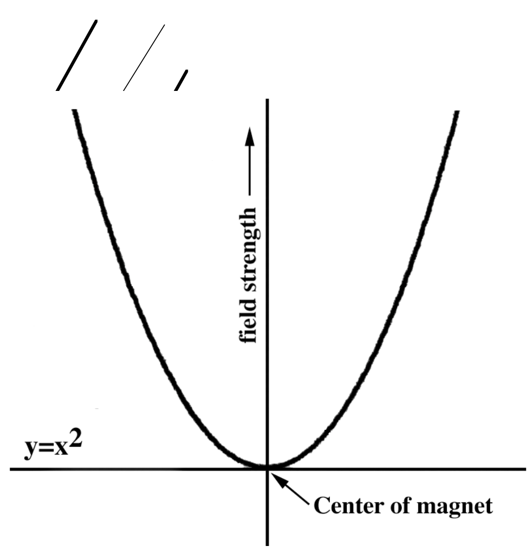

components. The parabola is always positive, except at the center, where it is zero.

The result is at first startling to those used to the symmetry of quadrupoles, because, being always positive, the horizontal forces are all in the same direction. Moreover, because the top and bottom coils combine to produce a field that opposes that of the "side" poles, the field vanishes at the center of the magnet.

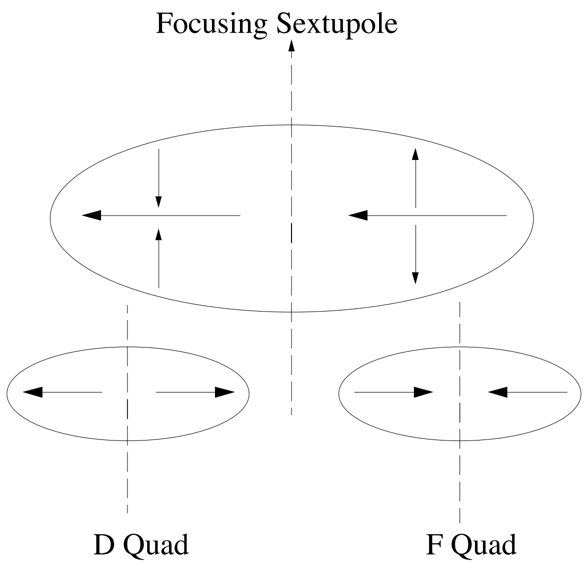

So how does this unidirectional parabolic field differentially focus particles of different energies? If a vertical line is drawn through the center of a horizontally focusing sextupole, the fields of the outside half are in the same direction as those of a horizontally focusing quad, and those of the inside half are in the same direction as those of a horizontally defocusing quad, as shown on the left.

Because of the dispersion created by the dipoles, higher energy particles will tend to be toward the outside of the beam path, and lower energy particles toward the inside (except, obviously, in the zero-dispersion regions.) Fortunately, it is the higher energy particles that need the extra focusing, because of their longer focal lengths; the low energy particles have a focal length that is too short, and being on the outside, the defocusing force undoes a portion of the quad focusing.

There are no sextupoles located in the straight sections or the dispersion-suppressor cells, because without dispersion they are of no use (see Figs. 2-22, 2-23).

Taken in combination, these effects narrow the tune spread to an acceptable range. Notice that the momentum spread of the beam has not been changed; it has just been managed so that beam is not lost from the machine.

In addition to the chromaticity sextupoles described above, there were trim sextupoles and skew sextupoles in the Main Ring. These have been salvaged and the option remains to use them in the Main Injector as harmonic correctors if deemed necessary.

The safe sextupoles are not used in the Main Injector.

Octupoles

A cross-section of a Main Injector octupole—yes, you can figure out the number of poles—is shown in Fig. 2-15. Octupoles are like sextupoles insofar as they selectively focus particles, and that they use the dispersion of the beam to identify those particles; they are like the quadrupoles in that the direction of the force changes when crossing the center of the magnet. They are unique in being able to preferentially focus particles with high amplitudes of betatron motion. Amplitude is strictly a transverse motion and is not the same as the longitudinal momentum spread dealt with by the sextupoles. (I agree, this is becoming more difficult to visualize with every step.) Octupoles are "third-order" components, and the force obeys equations such as $F=ax^3$, as in the graph on the next page:

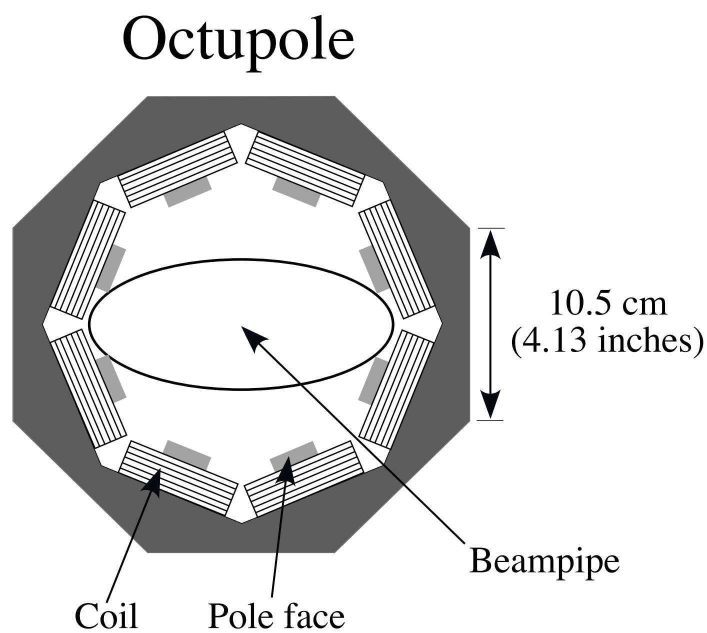

Figure 2-15

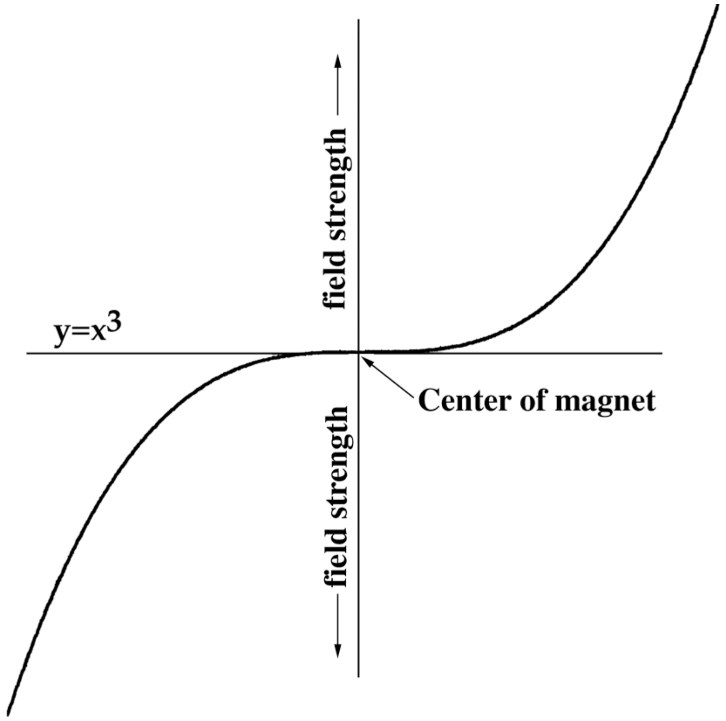

In a cubic equation like this, the force increases very sharply with increasing distance from the center. It is clear from the curve that a particle farther from the center will see a much larger restoring force, raising its tune.

There are vertical and horizontal octupoles in the Main Injector, and they have a relatively modest role in maintaining circulating beam. Part of their function is to compensate for octupole components generated by the shape of the laminations in the 84" quadrupoles recycled from the Main Ring. (The laminations had been deliberately trimmed in ancient times in order to push the Main Ring beam to higher energies.)

The horizontal octupoles greatly outnumber the vertical octupoles, because they are used during resonant extraction. Their role will be discussed in that chapter.

Permanent Magnets

As far as the Main Injector is concerned, permanent magnets are only used in the MI-8 line and beam transfer to and from the Recycler. If the reader is already saturated with the study of magnetic fields at this point, he/she may skip this section, and the loss of continuity will not be noticed until Chapter 7 is encountered. It would, however, be wise to read the "Names and Locations" section at the end of this chapter.

The extensive use of permanent magnets is new to accelerator design. In the Main Injector, permanent magnets make up most of the MI-8 line; the Recycler, of course, is built almost entirely of permanent magnets. Because their magnetic field strength cannot be changed at a whim from the control room, permanent magnets are only used where the beam energy does not change.

In recent years, an industry has developed for the production of strontium ferrite magnets. Strontium ferrite is the material of choice for a wide variety of applications, from magnets on refrigerator doors to inductive pickups in automobiles. Fermilab contracted with industrial suppliers of strontium ferrite to make tens of thousands of ferrite bricks. In the factory, the strontium ferrite is pulverized and compressed into a brick. Each brick spends several days slowly creeping through a 300-foot kiln until the ferrite has been baked into a ceramic.

Fermilab places the bricks in an Accumulator-style dipole that magnetizes them; the dipole is pulsed for about 3 seconds and the brick retains a field. (Although hysteresis is a nuisance in electromagnets, it is actually responsible for the "permanence" of a permanent magnet. Remember that once the magnetic domains are aligned, there is a self-sustaining field that keeps them aligned.) The bricks have a dull gray color and are predictably heavy.

The process of assembling the bricks is surprisingly empirical and, at first glance, imprecise. They are stacked around a section of beam pipe and, if field measurements do not meet specifications, they can be restacked. The assembly of bricks is enclosed in a long steel box. The ferrite bricks play the same role as the copper coils in the electromagnets, creating the magnetic field; the steel box is analogous to the steel laminations of the magnets, strengthening the field. Steel pole faces internal to the casing strengthen and shape the field.

The magnetic field is susceptible to changes in temperature, so temperature compensation shims are sandwiched between the bricks as necessary. The shims are made of an iron-nickel alloy, probably melted down from a meteorite; the alloy has a Curie point (demagnetization temperature) of 50˚C. The weakening of the field in the shims, as the magnet gets warmer, compensates for the field changes in the ferrite bricks. After assembly, the entire magnet is "frozen" to 0˚C for 24 hours to standardize the hysteresis due to thermal cycling. As the magnet warms back up to room temperature, the field is monitored in order to verify the effectiveness of the temperature compensation shims.

Adding a geometrically complex end shim at either end can compensate for higher order components of the field. A variety of end shims are kept on hand, and a computer matches the optimal pole face to the measured field. (This is analogous to the role that the end packs play in the main dipoles.)

Permanent magnets come in several varieties. Horizontal dipoles are used in the MI-8 line. However, permanent magnet quadrupoles are used sparingly, because gradient magnets, similar in concept to the combined-function gradient magnets in the Booster, do most of the focusing.

All of the dipole magnets are identical in cross-section (Fig. 2-16). In the MI-8 line, the horizontally bending dipoles are referred to as PDD magnets, the “P” standing for “permanent” and the “DD” standing for “double dipole.” “Double dipole” comes from the fact that there are two layers of ferrite brick in the magnet. Stacking the bricks in two layers allows for a shorter magnet length than would otherwise be the case.

The strength of the field is inversely proportional to the distance between the two steel pole faces; aluminum spacers allow the field to be weakened or strengthened as necessary.

Notice that the field lines in Fig. 2-16 point up in the region of the beam pipe, which is opposite the direction of the field in the Main Injector dipoles. The beam in the MI-8 line is normally bent to the right,

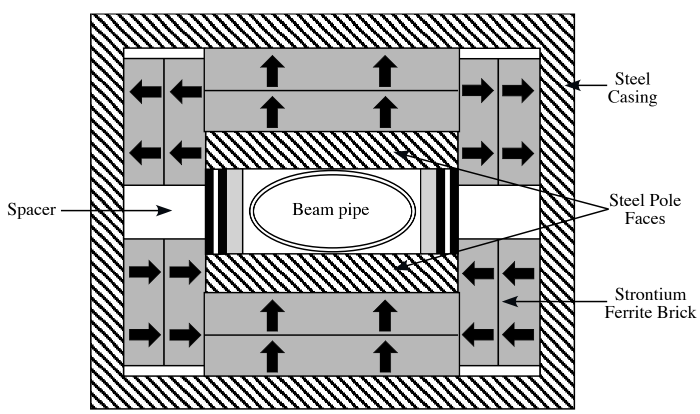

Figure 2-16
PDD Dipole

PDD stands for "Permanent Double Dipole." The ellipse in the center is the beam pipe. The gray rectangles represent the permanently magnetized strontium ferrite bricks; the arrows represent the direction of the magnetic field. The "Double" designation refers to the fact that the bricks are stacked two layers deep. The steel in the casing and the pole faces are of a low carbon type that can retain magnetic flux.

You can find PDD magnets in the MI-8 Line. Refer to Figs. 7-1, 7-2, and 7-4 to see how they are distributed in the line.

rather than to the left as in the Main Injector ring. (The MI-8 line does have a short stretch, called the Reverse Bending section, in which beam needs to be bent to the left; this is accomplished by turning the PDD magnets upside down.)

A cross-section of a typical permanent quadrupole magnet (PQP) is shown in Fig. 2-17. A focusing PQP can be converted into a defocusing quad by turning it 180˚ around its vertical axis, or by rotating it 90˚ with respect to its longitudinal axis. Minor adjustments to the field are accomplished by adding steel washers (developed by elite units of the Ace Hardware Research Division) to the corners of the magnets. The only PQP magnets used in the Main Injector are in the Reverse Bending section.

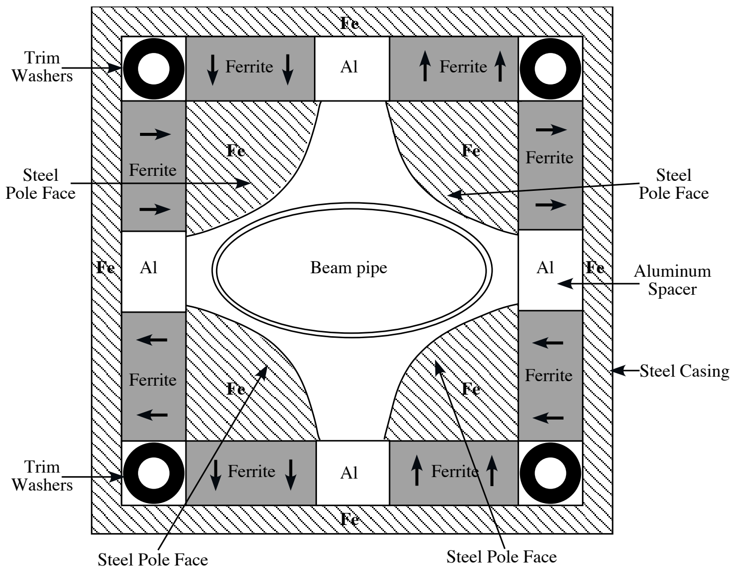

Figure 2-17
PQP Quadrupole

The PQP permanent quadrupoles are used in the Reverse Bending Section of the MI-8 Line. Permanently magnetized strontium ferrite bricks drive the magnetic field, which is shaped and reinforced by the iron pole faces. Trim washers are added for fine adjustments to the fields.

See Figs. 7-1 and 7-3 for the distribution of the PQP magnets in the MI-8 line.

A majority of magnets in the MI-8 line and the Recycler are the PGD (permanent gradient dipole) magnets (Fig. 2-18). The PGD magnets are similar in principle to the combined-function gradient magnets of the Booster. They are built to have both dipole and quadrupole field components. The shape of the pole faces, which are slanted in a linear fashion, generates the quadrupole components of the field.

More details on the use of permanent magnets in the Main Injector can be found in the “Beam Transport Lines” chapter.

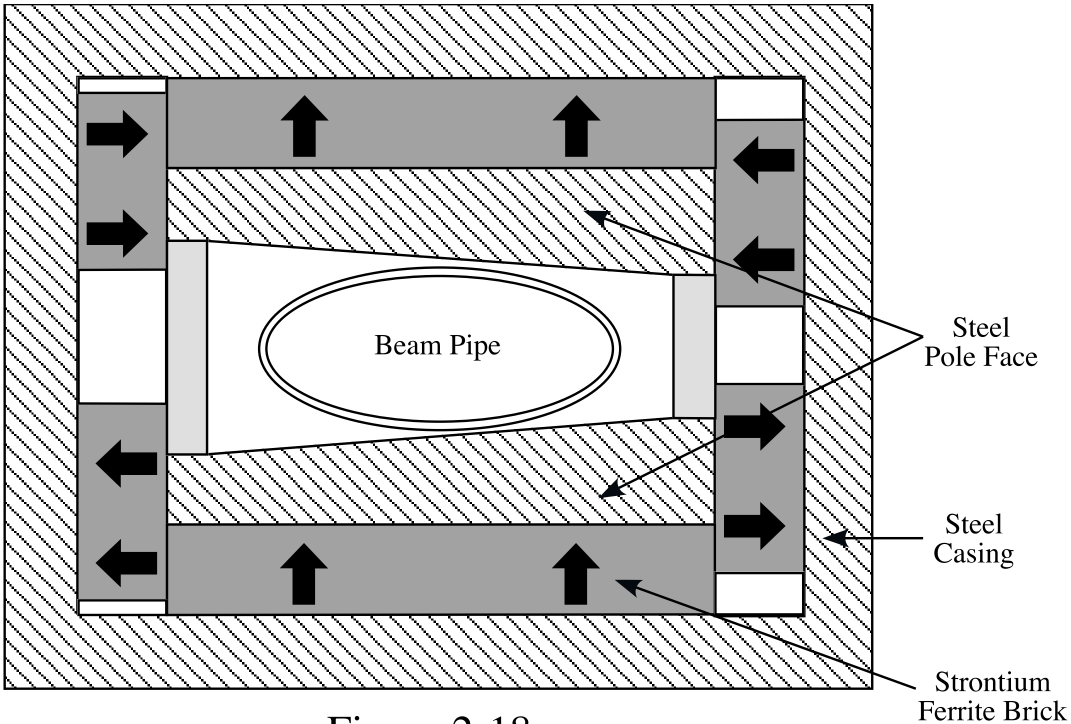

Figure 2-18
PGD Gradient Magnet

PGD stands for "Permanent Gradient Dipole." Like the Booster gradient magnets, PGD magnets can be thought of as having both dipole and quadrupole components. The slope of the iron pole faces above and below the elliptical beam pipe creates the gradient.

PGD magnets are used in the MI-8 line. See Figs. 7-1, 7-4, and 7-5 for their distribution.

Names and Locations

It is now possible to assemble an integrated view of the distribution and naming of the magnets responsible for maintaining circulating beam in the Main Injector.

The naming system is anchored around the main quadrupoles. Each quadrupole is assigned a number, and each device near the quadrupole includes that number as part of its name.

Since beam first enters the machine at MI-10, the numbering scheme is defined to start at the beginning of that straight section. The first series of numbers, from MI-10 to MI-20, begins with the “100” location at the first quadrupole (proton direction) in the straight section. Between MI-10 and MI-20 the number increments with each main quadrupole. The “100” location is occupied by a focusing quadrupole; the convention has been adopted that focusing locations be assigned even numbers, with the defocusing locations getting the leftovers.

When the MI-20 Service Building is reached, the “200” series begins. This series begins with “201” because there is a defocusing quadrupole at that location.

The “300” series begins with the straight section at MI-30. Again, whenever a straight section is present at a service building, the new series of numbers begins at the first quadrupole of the straight section. The rest of the pattern should be clear (review Fig. 2-1). Unfortunately, for those who have to remember such things, the number of locations is not consistent from section to section; it can be 31, 32, or 41.

Magnets and other components near the large quadrupoles are assigned a name based on the quadrupole number. Fig. 2-19 shows the layout of the

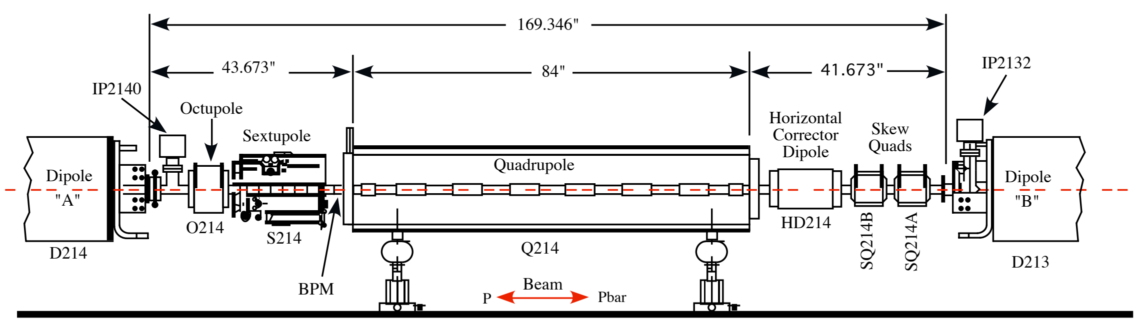

The correction magnets at any given location are clustered around the main quadrupole, and take on the number of that quadrupole. The name of the quadrupole itself starts with "Q." The sextupole and octupole magnets, which are generally downstream of the main quadrupole, are prefaced with "S" and "O," respectively. The corrector dipoles are prefaced with "HD" or "VD." The main dipoles are designated by "D." The two main dipoles downstream of the main quadrupole take the number of the main quad, but main dipoles upstream take the preceding number.

If there is more than one corrector of the same type, such as the two skew quads at this location, they are designated "A and "B."

Ion pumps have a 0, 1, or 2 linked to the location number. In this picture, the first Ion pump at Location 214 is IP2140; IP2141 is between the "A" and "B" dipole, to the left of the picture. IP2132, at the right of the picture, is the last pump belonging to location 213.

Compare the correctors in this picture to Figs. 20-25. Different combinations of correctors are found at different locations.

Fig. 2-19

The 214 Girder: Typical Magnet Locations and Naming Conventions.

magnets at a selected location.

This particular combination of magnets can be found at the 214, 336, 514, or 636 locations. Suppose that it represents the location at 214. The quadrupole itself is given the name “Q214.” The two large dipoles downstream (proton direction) also belong to this location, and are given the names “D2141” and “D2142,” in that order.

The magnets immediately upstream of the large quadrupole, up to the previous main dipole, also belong to this location. The main dipoles upstream belong to the “213” location. If an ion pump (responsible for maintaining vacuum) is upstream of the main quadrupole, it takes the name of the previous location.

Preceding (proton direction) every large quadrupole is a correction dipole: a horizontally bending dipole at the focusing locations, or a vertically bending dipole at the defocusing locations. The horizontal dipole at the 214 location is named “H214,” as the vertical corrector at the previous location is named “V213.”

The skew quads, 16 in all, are grouped into four clusters (Fig. 2-20). At the “x14” and “x36” locations, there are two skew quads on the same girder. The skew quads at location 214 are named “SQ214A” and “SQ214B,” in that order (proton direction). As mentioned earlier, the skew quads are divided equally between the horizontal and vertical locations. The other two skew quads in this cluster are SQ213 and SQ215.

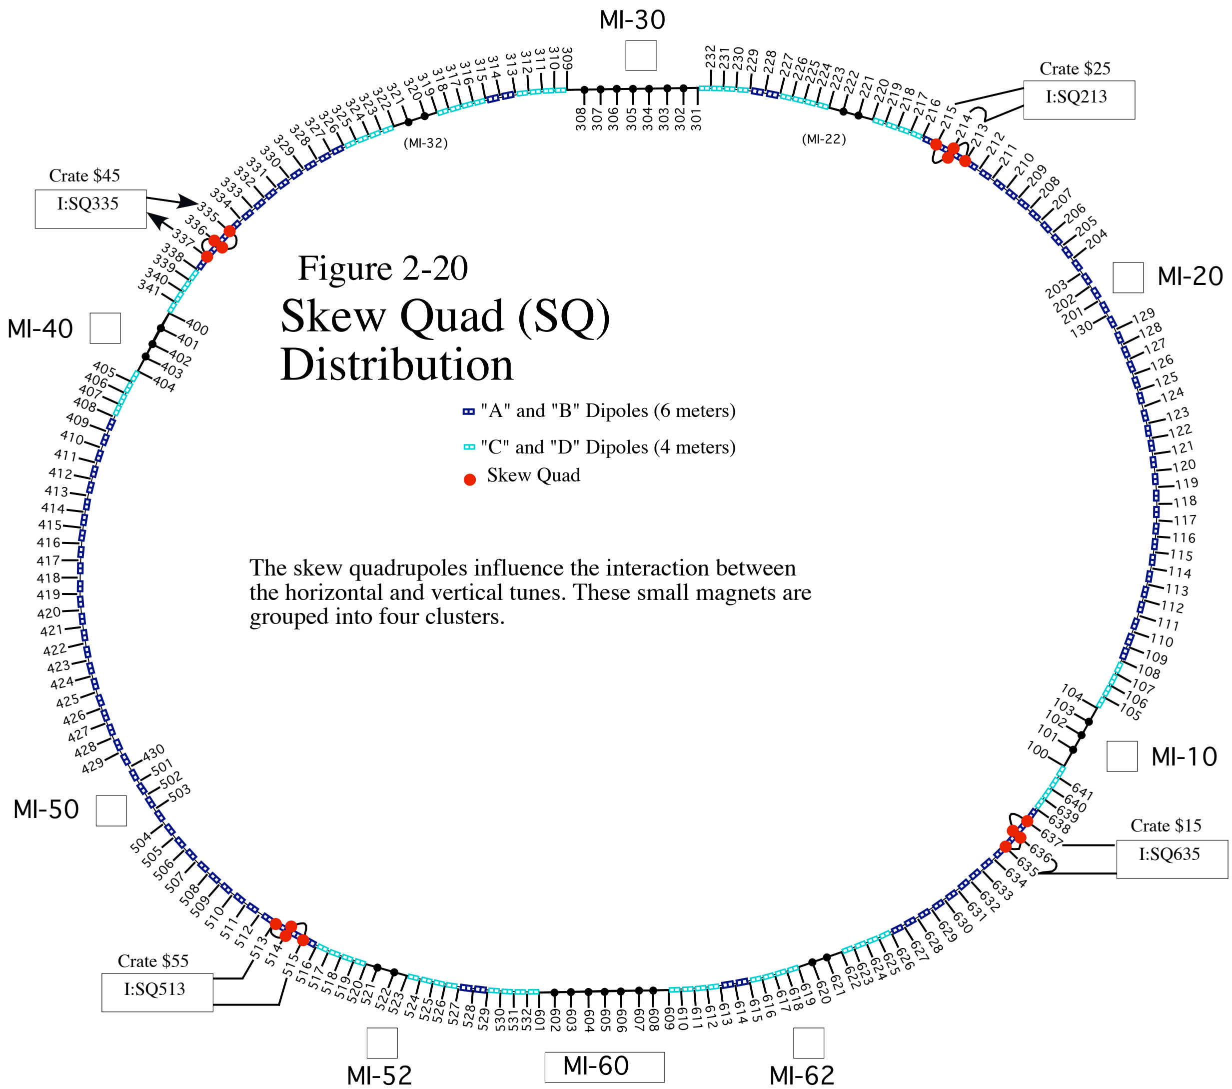

Figure 2-20
Skew Quad (SQ) Distribution

The skew quadrupoles influence the interaction between the horizontal and vertical tunes. These small magnets are grouped into four clusters.

The 16 trim quads are also arranged into four groups (Fig. 2-21). Notice that all of them are found at horizontal locations. The process of resonant extraction requires large changes in the horizontal tune, not the vertical tune. There are no trim quads at the 214 location, but where they are found, they are given names beginning with “QC” (for “quad corrector”), followed by the location number.

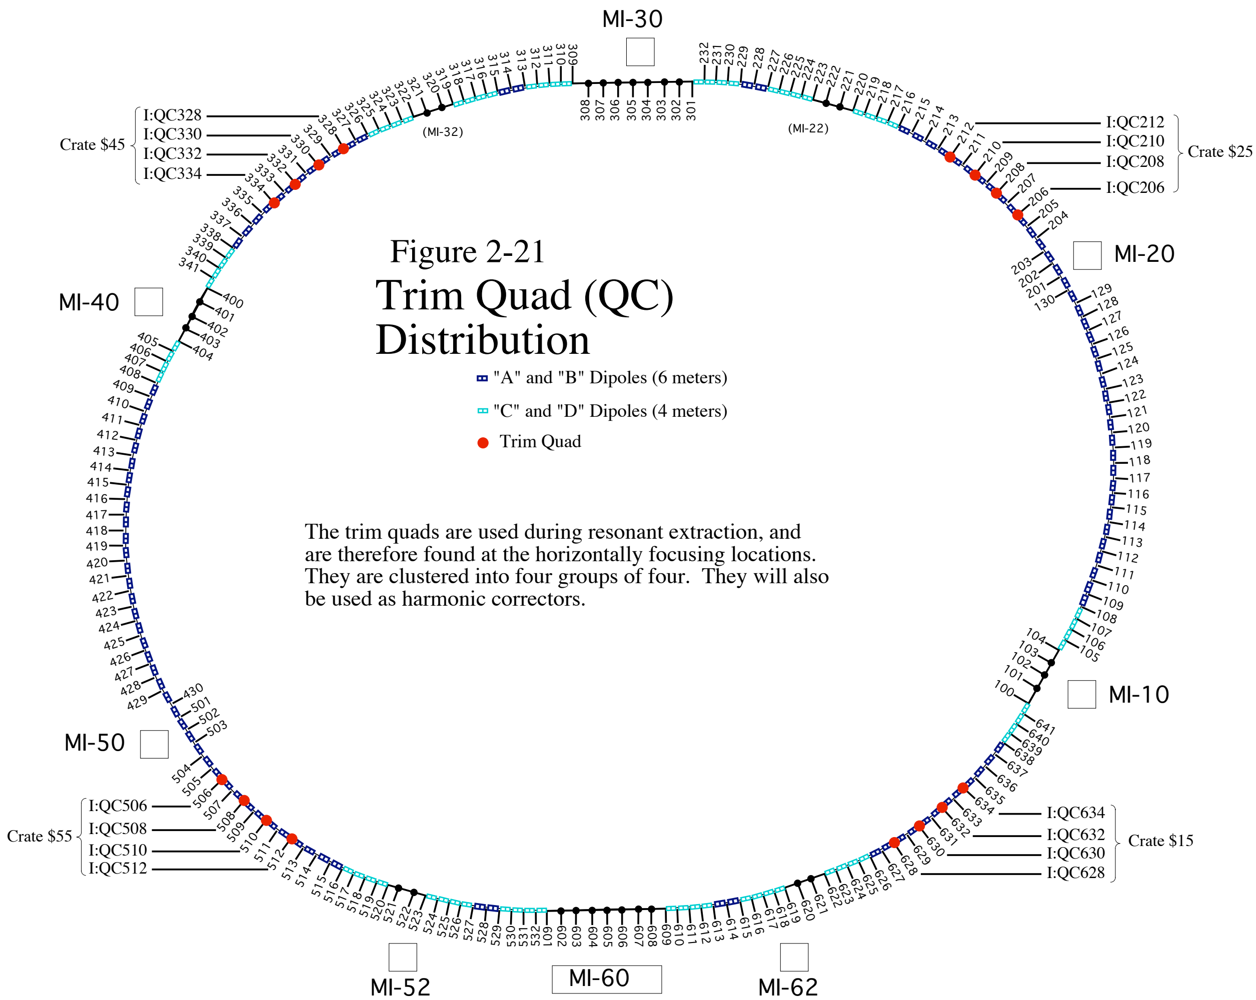

Figure 2-21
Trim Quad (QC) Distribution

The sextupoles (Figs. 2-22, 2-23) are distributed at each location in the normal arcs: horizontal sextupoles at the focusing locations and vertical sextupoles at the defocusing locations. Remember that sextupoles use dispersion to compensate for variation in the momenta of the particles, so they would be of little use in the low- and zero-dispersion regions.

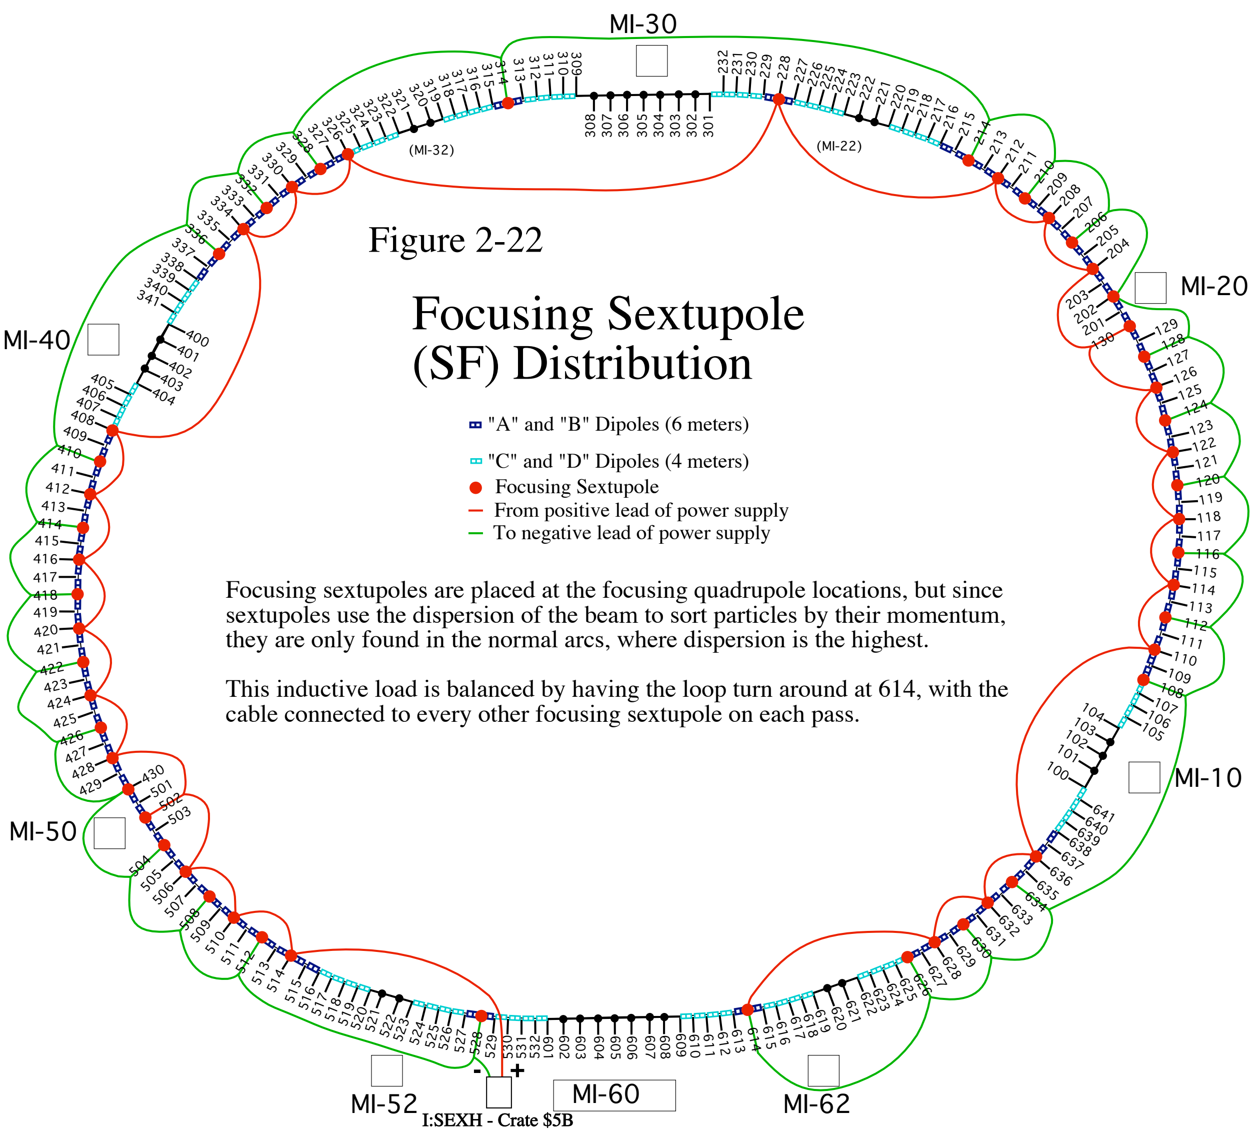

The sextupole at location 214 is named “S214.”

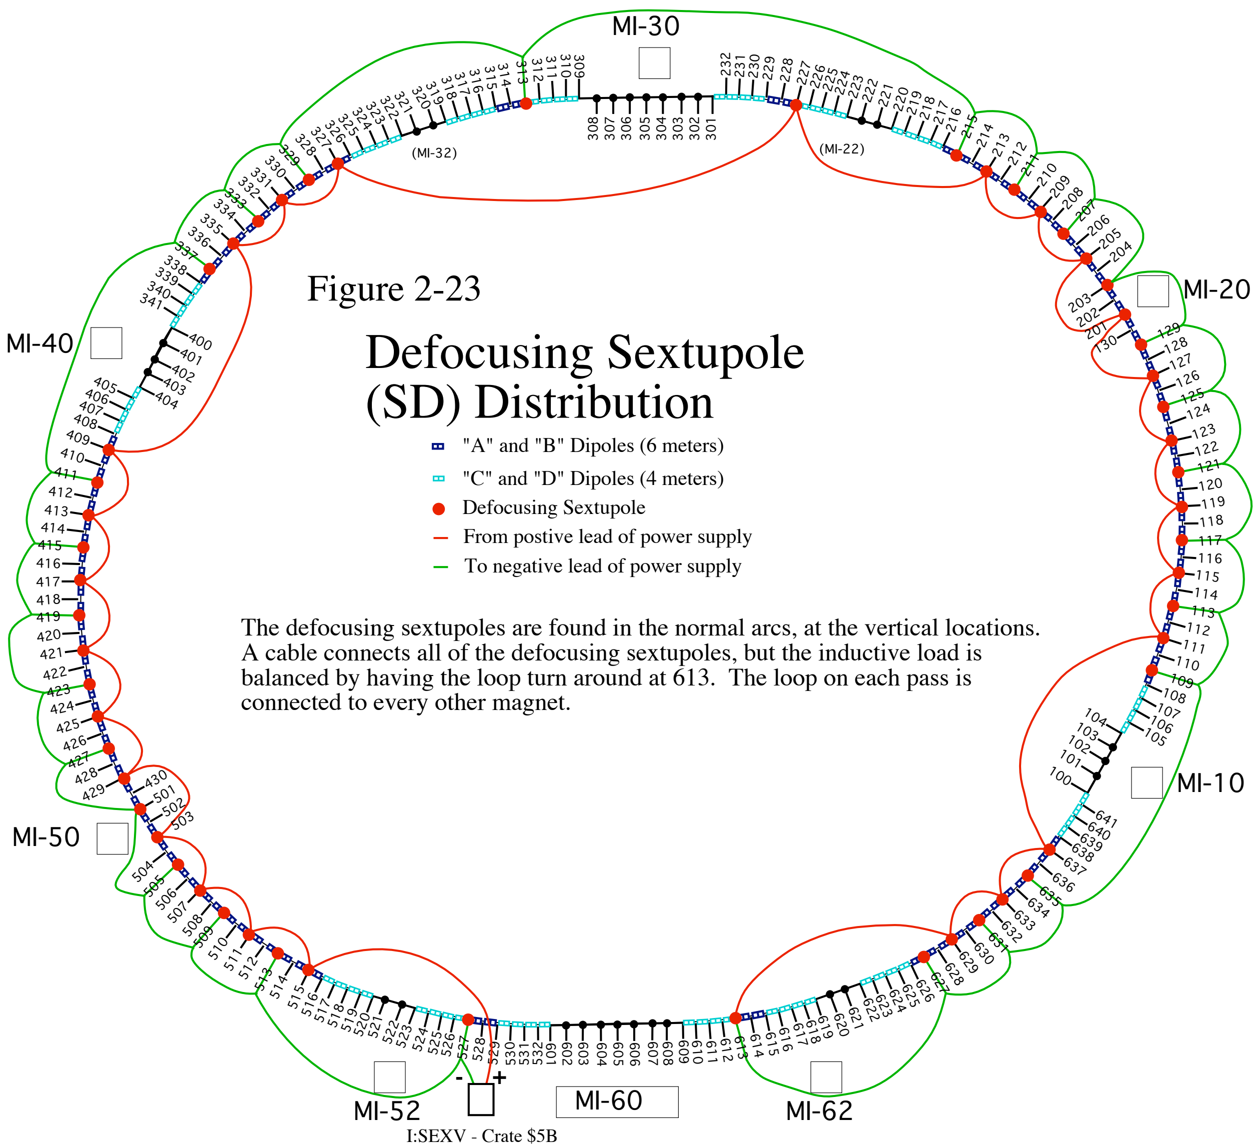

Figure 2-23

Defocusing Sextupole (SD) Distribution

The defocusing sextupoles are found in the normal arcs, at the vertical locations. A cable connects all of the defocusing sextupoles, but the inductive load is balanced by having the loop turn around at 613. The loop on each pass is connected to every other magnet.

The octupoles (Figs 2-24, 2-25) are used to stabilize circulating beam, but are also used for resonant extraction. There are only 16 octupoles at the vertical locations, but 50 at the horizontal locations. (The horizontal octupoles are present wherever the horizontal sextupoles are found, except at 108, 326, 408, and 626; the vertical octupoles are grouped into four clusters of four

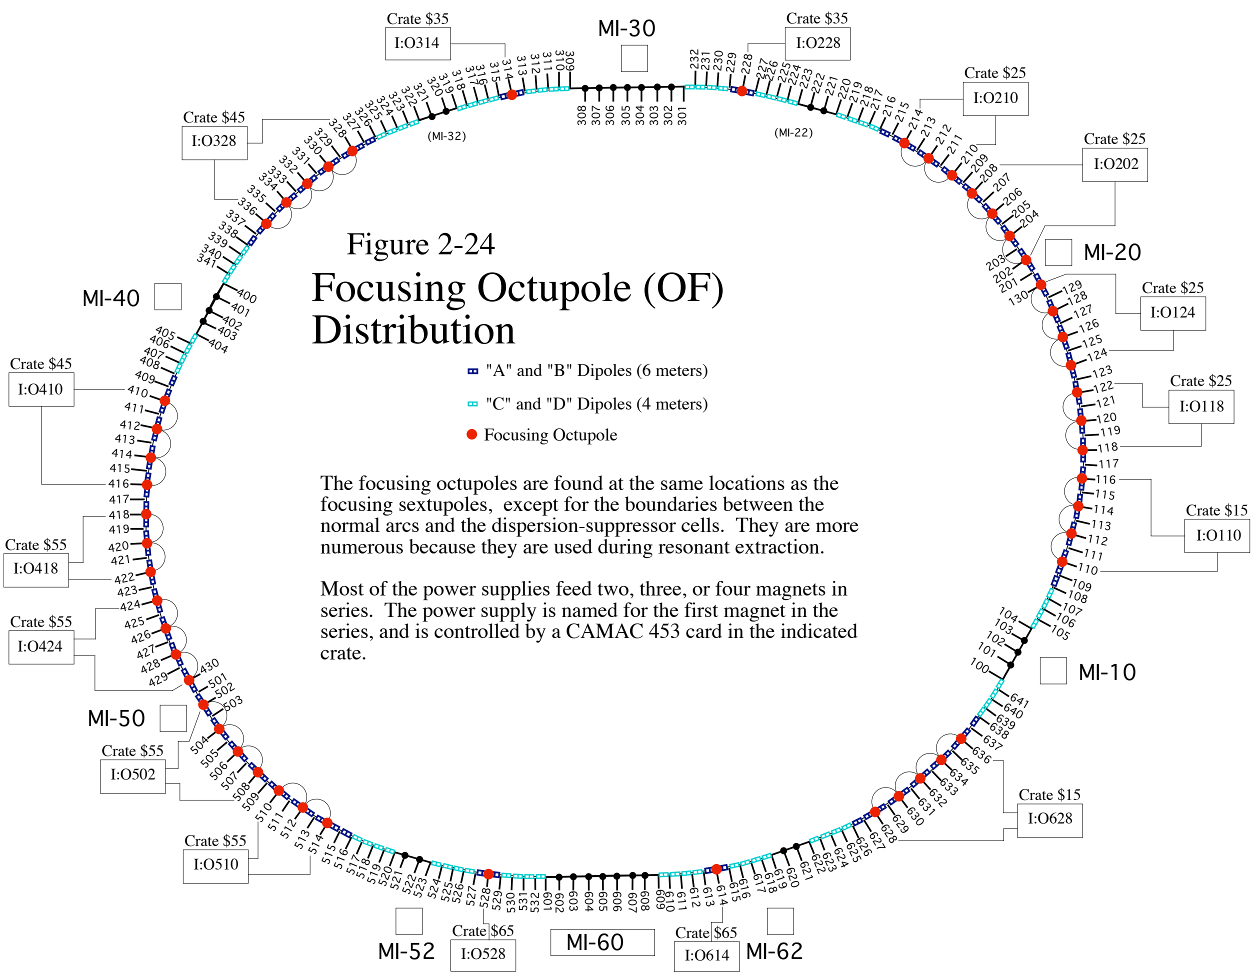

Figure 2-24
Focusing Octupole (OF) Distribution

magnets.)

The preponderance of the horizontal octupoles is due to their importance during resonant extraction. The octupoles are named with an “O” in front of the location name; at this location, the octupole is called O214.

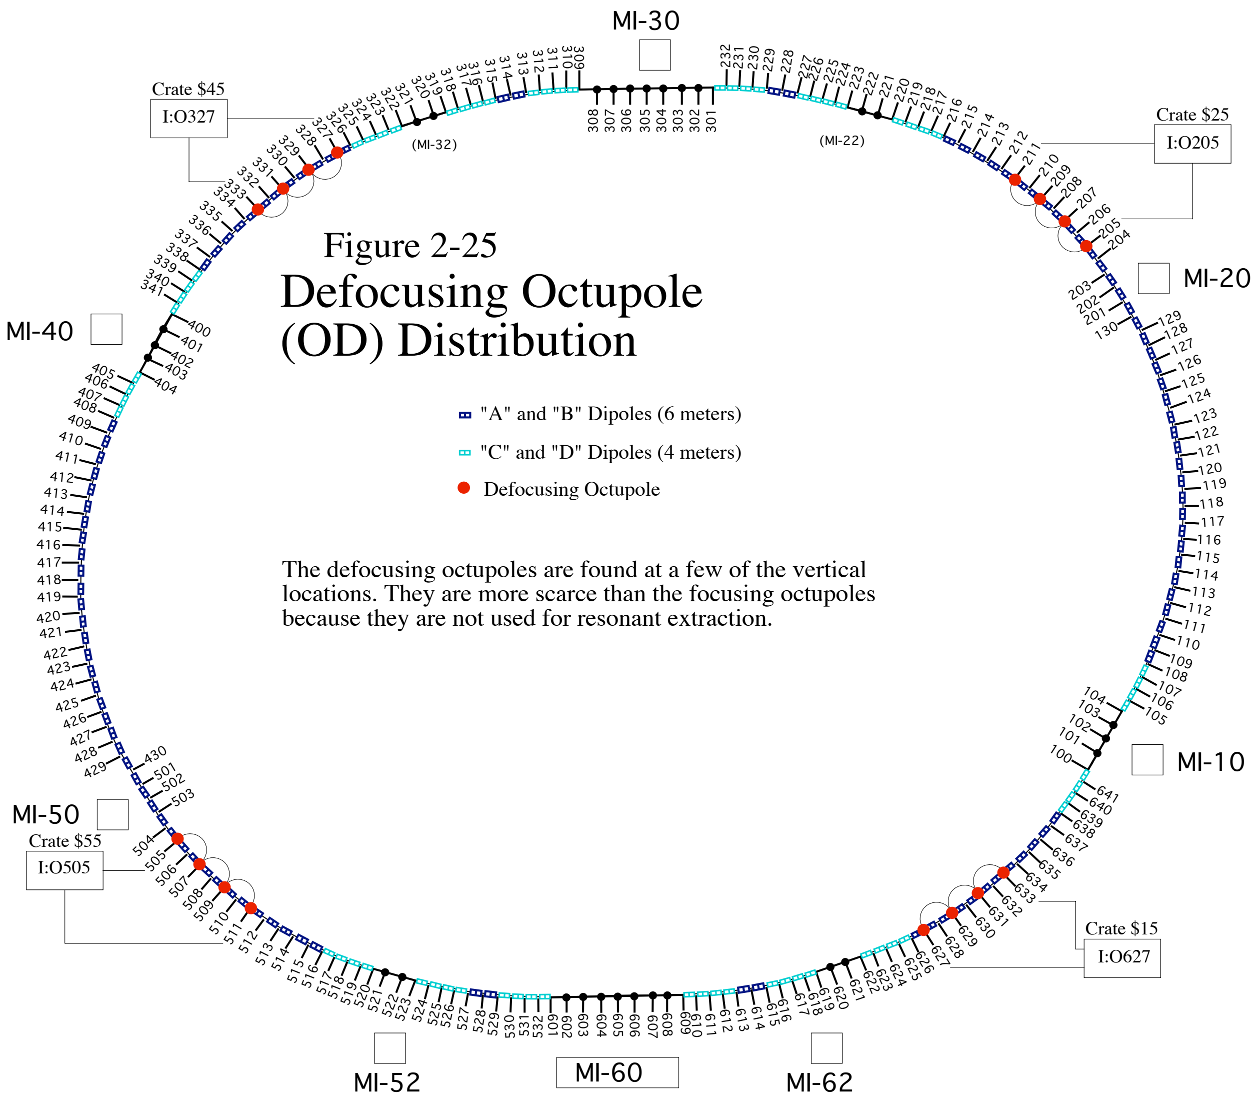

More detail on the trim quads and extraction octupoles will be discussed in the chapter on resonant extraction.

Remember, whether or not a magnet is focusing or defocusing does not need to be explicitly designated in its name, since the “evenness” or “oddness” of the location number encodes that information.

The next chapter, “Power Supplies,” is the story of how the magnets get their current.

What The Heck Is Inductance, Anyway?

A recent Gallup/CNN poll revealed that an astonishing 36% of the general population suffers from an inadequate understanding of magnetic inductance. That figure is even higher among physicists and engineers. Since the concept is central to understanding many aspects of the Main Injector, or of any accelerator, the basics will be presented here.

Inductance

Inductance can be thought of as resistance to a change in a magnetic field.

When current begins to flow in a conductor, a magnetic field begins to form around the conductor. The field forms in such a way to create a voltage that opposes the increasing current in the conductor. The polarity of the induced voltage is the opposite of the polarity of the voltage created by the changing current and the magnetic field. The induced voltage limits the rate at which current can be increased or decreased in a conductor.

As soon as the current reaches a stable value the magnetic field is consequently at a stable value, and the induced voltage vanishes.

The symbol used to represent inductance is, inexplicably, L. It is measured in units called henrys, probably after some guy named Henry. A henry is an unrealistically large unit for most everyday applications; the largest inductance in any Main Injector magnet is about 2 millihenrys. Line several hundred up in a big ring, however, and the total inductance will command some respect.

The inductive voltage across a system is given by:

$$V = L\frac{di}{dt}$$

where di/dt is the rate of change of the current, usually expressed as amps per second. The faster the current is changing, the greater the induced voltage.

The actual inductance, L, depends on the geometry of the magnet and is more complicated to compute. One factor is the number of turns. A straight conductor (one "turn") has the least inductance; inductance increases with the number of turns in a coil. That is the main reason why the coils in the Main Injector dipoles are large; if the same amount of current can be carried with fewer turns in the coil, the magnets can come to full strength faster because there is less inductance and a weaker opposing voltage. That means the Main Injector can be ramped more frequently, creating a higher stacking rate and ultimately higher luminosity for stores in the Tevatron. Wider is better.

The large dipoles in the Main Injector are in series. The copper bus emerges from one pair of magnets, passes behind the quadrupole, and re-enters the next pair, continuing in this fashion all the way around the ring. The inductance adds up with every magnet.

Pairing the "A" and "B" magnets, and the "C" and "D" magnets, distributes the inductance more evenly. The inductive voltage still opposes any change in current, but each power supply only sees about half of the opposing voltage that it would see otherwise.

The upper and lower bus are actually in series, but in the power supplies at MI-60 the upper bus turns around and becomes the lower bus, and the lower bus turns around and becomes the upper bus (Fig. 2-4). The inductive voltage is thus balanced in such a way that the voltages necessary for powering the ring are kept to a minimum.

Beneficial Aspects

There are beneficial aspects to inductance as well. Transformers use the inductance of the primary and secondary windings to step up or step down voltages. The fact that inductance resists sudden changes in current is employed when chokes are designed to smooth noisy power supply outputs.

The inductance of the magnets in the tunnel blocks ripple from the power supplies so that resistors can attenuate it.

Resonant Systems

The inductance of every component has to be taken into account whenever designing an electrical system, but it is especially crucial in resonant systems. (Actually, every electrical system is a resonant system of some sort, but some devices are built specifically to exploit the phenomenon.) Inductance, along with capacitance and resistance, determines the frequency of the oscillations in a resonant system. Kickers, which transfer beam from one accelerator to another, and RF cavities, which accelerate the beam, are examples of resonant systems used in the Main Injector. More detail about kickers can be found in Chapter 7, and RF cavities will be described in the RF chapter.

Chapter 3: Power Supplies

This chapter will discuss how electrical power is supplied to the magnets in the tunnel and to the equipment in the service buildings. It will also cover the regulation of the magnet power supplies, especially that of the major busses. The chapter is primarily concerned with power to the magnets in the ring; power to devices in the beam transport lines and to general service building utilities will be covered in later chapters.

The biggest consumers of electrical energy in the Main Injector are the large dipoles and quadrupoles. All of the dipoles are on a single bus, in series (Fig. 2-4). All of the focusing quads are on a second bus and the defocusing quads on a third (Fig. 2-7). The power delivered to these three sets of magnets is called pulsed power. Because of the large amounts of power required by these magnets, and the rapidly changing loads that they create, pulsed power is usually treated separately from other applications. The rest of the distributed power—to the smaller magnets, service building electronics, water systems, light bulbs, etc.—can be considered conventional or house power.

Commonwealth Edison and the Substations

Commonwealth Edison, or ComEd, is the public utility company that supplies power to most of the residential and commercial users in northern Illinois. Fermilab purchases all of its power from Com Ed. Fermilab's utter dependence on Com Ed becomes evident when a substation fire or an aggressive cherry tree interrupts power to the site, and sometimes weeks are required to recover from the damage. To be fair, Fermilab's rapidly changing inductive loads, especially the Tevatron and the Main Injector, can be a major headache for Com Ed as well.

ComEd generates power from several nuclear power plants distributed around northern Illinois. There are also coal-fired plants that can be called into service during times of heavy demand, such as on hot summer days when 15 million air conditioners are operating simultaneously. Generated power is placed on the grid, a large network for the distribution of power. The transmission lines on the grid are usually maintained at 345 KV or higher. (Power lost as heat is proportional to current; since power is the product of voltage and current, a given amount of power is most efficiently transmitted with a high voltage and a low current.) The transmission lines entering the Fermilab site are at a voltage of 345 KV.

The 345 KV is 3-phase at 60 Hz; that is, there are three separate power lines, running parallel to each other, each being 120° out of phase with respect to the other two. The phases are usually designated “A”, “B”, and “C”:

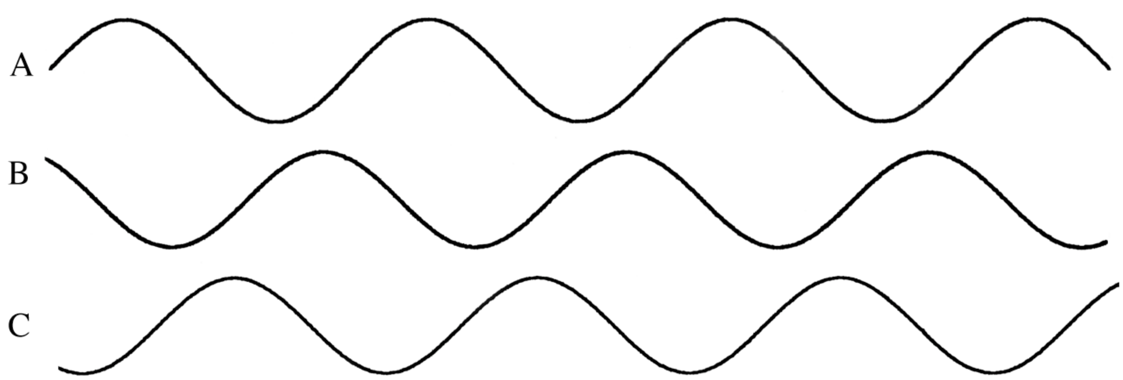

Wherever power is to be drawn from the grid, a substation is present. Transformers at the substations step down the 345 KV, a voltage generally too high for most household appliances, to a more manageable level. At Fermilab, the substation transformers step the voltage down to 13.8 KV.

There are three substations on the Fermilab site. The first is the Master Substation, located about half a mile north of the High Rise. The Master Substation supplies power to many of the accelerators as well as to most of the buildings on site. The second is Giese Rd., a small substation just to the west of the Antiproton Source. The Kautz Rd. Substation (KRS) is located just outside the Tevatron Ring near the MI-50 Service Building. It was built specifically for the Main Injector and will be the one of greatest interest in this book. However, Giese Rd. and the Master Substation can “back-feed” some power to the Main Injector when necessary, and Kautz Rd. can back-feed power to systems normally fed through the Master Substation.

Fig. 3-1, located below, provides a general schematic of power distribution related to the Kautz Rd. Substation. It should be referred to frequently when reading the next few paragraphs.

Fig. 3-1
Kautz Road Substation

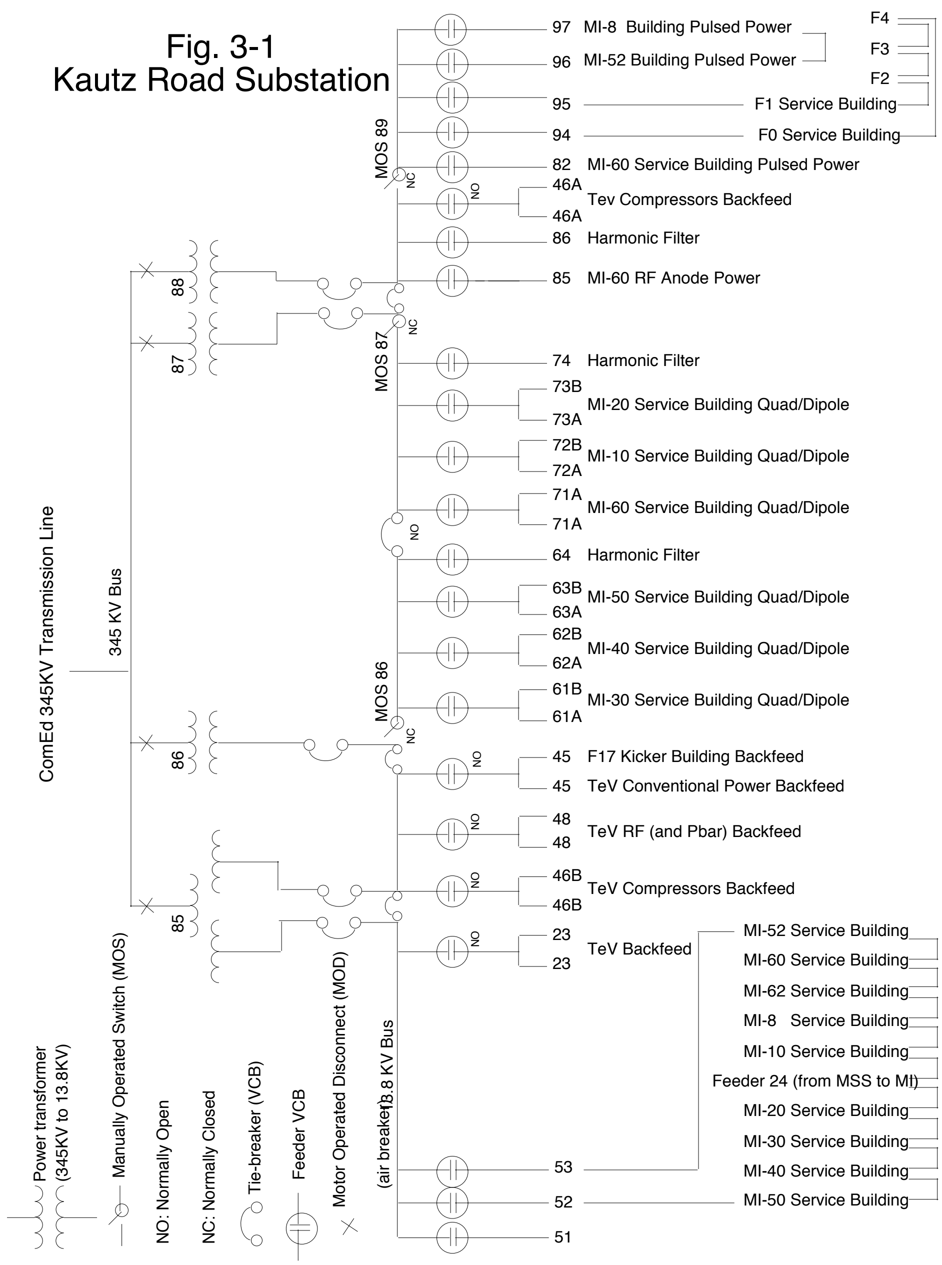

Main Injector

The 345 KV lines to the Kautz Rd. Substation run parallel to Butterfield Road near the southern border of the laboratory until they strike northwest to meet the substation. Inside the substation fence the three phases connect to gigantic transformers, designated T-85, T-86, and T-88. They all step the 345 KV down to 13.8 KV. (The numerically alert will notice that T-87 is missing from the sequence. It was to be installed with the others, but it failed on the manufacturer's test stand and is still awaiting shipment. As of this writing (October 1999), it appears that the substation will be short one transformer for another year or so. T-88 has temporarily taken over much of the work that T-87 was to have done. The discussion to follow describes the current temporary configuration and will have to be revised when the replacement transformer arrives.)

From time to time, T-85, T-86, and T-88 need to be isolated for maintenance or repairs. On the primary (345 KV) side, air breakers, also known as Motor Operated Disconnects (MODs), at the top of each transformer, are used to isolate Com Ed's 345 KV from the transformers. A MOD operates by rotating a segment of the line and breaking the path of the current.

On the secondary (13.8 KV) side, vacuum circuit breakers (VCBs) are used to isolate the transformers from the load. VCBs are breakers that open and close in an evacuated "jar;" under vacuum, there is less likelihood of arcing. These main VCBs, rated to withstand 3000 amps, are also known as tiebreakers. (The analogy is obviously patterned after the duty of the Vice President to resolve a deadlock in the Senate.) There are also tiebreakers for isolating segments of the 13.8 KV bus.

The 13.8 KV from the transformers is distributed to the service buildings through individual feeders. (A feeder is a large underground cable capable of carrying a correspondingly large amount of power.) Unlike the Main Ring feeders, which had been buried directly in the ground and were susceptible to ground faulting, the Main Injector feeders are protected in underground concrete enclosures. They are often named after the last digit of the transformer from which they originate (e.g., feeders originating from T-86 are assigned numbers beginning with "6," and those from T-88 sometimes begin with an "8").

T-85 has two secondaries, one that supplies conventional power to the Main Injector service buildings and another that can be used for back-feed to the Master Substation. Feeders 52 and 53 carry conventional power from T-85 to the Main Injector ring. T-86 and T-87 were designed to provide pulsed power to the main dipoles and quadrupoles in the Main Injector ring; since T-88 has taken over T-87's load, many of its feeders begin with "7." The once and future purpose of T-88 is to supply pulsed power to beamlines, including the P1, P2, and P3 lines. Except for the P3 line, which has not yet been commissioned, T-88 still performs that function.

The individual feeders can also be isolated from their loads by VCBs. These feeder breakers, required to carry less current, are rated to 1200 amps. In addition to the tiebreakers, the 13.8 KV bus can be broken into segments by manually operated switches (MOS). These knife switches at Kautz Rd. are opened during an access. The reader should be able to determine from Fig. 3-1 that if all three MOS switches (MOS 86, MOS 87, and MOS 89) are open, pulsed power is removed from the bus while still leaving conventional power available to the buildings.

The procedure for opening the MOS switches is popularly known as racking out. In the Main Injector tunnel, the pulsed power is carried on exposed copper busses. Since this situation would represent a serious electrical safety hazard to anyone present, it is vital to confirm that the feeders cannot be powered. Each of the MOS switches has a dedicated "Open" and "Closed" key; the keys are controlled from the MCR. In addition, the VCBs for each pulsed power feeder are opened.

Remember that three-phase is present at every point discussed so far. Each MOS switch actually consists of three knife switches, one for each phase. A window has been provided for confirmation that all three knife switches have opened.

It should be noted that the MOS switches were not designed to be opened under load; all pulsed power must be removed from the appropriate devices before operating the switch.

Some of the feeders are paired and feed a system from opposite directions. For example, Feeder 52 starts its rounds at MI-50, while Feeder 53 starts at MI-52. The service a building disconnects can be configured so that the load is shifted from one feeder to the other.

Most of the feeders isolated by MOS 89 are designated with a “9.” (Remember that there is no Transformer 89; these feeders are powered from T-88.) Most of these feeders are dedicated to F Sector pulsed power, but 96 and 97 serve MI-8, MI-40, MI-52, and MI-62. Note that these locations are where most of the beam transport lines reside.

Pulsed Power

Remember that T-86 and T-88 provide all of the pulsed power for the main dipoles and quadrupoles. The feeders from T-86 travel underground from the substation until they reach MI-50. Two of the feeders, 63A and 63B, supply MI-50 with power. The other T-86 feeders turn to the left. 62A and 62B supply pulsed power to MI-40, and 61A and 61B continue on to MI-30. Similarly, the feeders from T-88 branch to the right and feed MI-60, MI-10, and MI-20. Feeders 71A and 71B go to MI-60, feeders 72A and 72B to MI-10, and feeders 73A and 73B to MI-20. When T-87 is installed, it will take over Feeders 71, 72, and 73.

Remember that in addition to the MOS switches, which isolate the feeders from the transformers, the individual feeder breakers are opened prior to an access. The breakers are operated from a console at the Kautz Rd. Substation.

Because of the transient spikes created when the feeder breakers are opened and closed, they are manipulated one at a time in a specific order. The order is executed by programmed logic controllers (PLCs), which act as software relays. PLCs have a variety of uses in the Main Injector; more detail will be forthcoming in the chapter on controls.

One feeder from each of the pulsed power transformers is connected to a harmonic filter. Each of the transformers has a dedicated harmonic filter; the filters are all located in the northwest corner of the substation. Their purpose is to suppress frequencies near the 720 Hz and 1440 Hz (and other unwanted frequencies) that may be on the feeders. The 720 Hz and 1440 Hz noise, as will be explained later, originates with the power supplies in the service buildings and feeds backwards toward the pulsed power transformers. These high frequencies are added to the voltage already present on the feeders and increase the voltage stress on the system. The harmonic filters, which consist of large stacks of inductors and capacitors, greatly enhance the lifetime of the pulsed power transformers by screening off the higher frequencies.

It may seem odd that the filters are not directly in line with the feeders, but remember that the filters, magnets, and power supplies are all part of the same interconnected resonant system. The filter feeders (no, these aren't clams) are designed to work in parallel with the other feeders.

A large resonant system such as the Main Injector power supply/magnet network can generate side effects that would not be intuitive to the untrained observer. One of these is that a 200 Hz standing wave appears across the magnets. If not compensated for, the varying impedance would cause every magnet to have a current that is out of phase with that of its neighbors. The compensation comes in the form of a mode-damping resistor placed in parallel with each magnet:

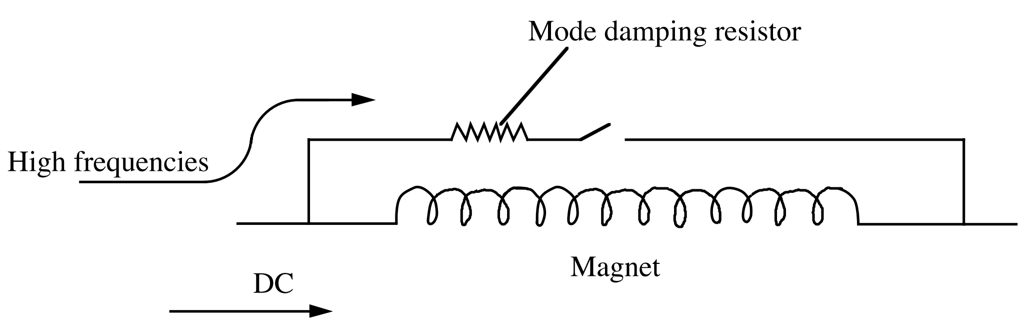

Exactly how this resistor works to block the standing wave is beyond the scope of this author . . . I mean, book. That is, it is left as an exercise for the reader. However, it can be said that higher frequencies are blocked by the high inductance of the magnet but pass readily through the resistor, where some of the energy is dissipated as heat.

The Service Building Utility Yard

The term “Utility Yard” as used here refers to the area outside each service building where you find the transformers and other high voltage equipment.

The way in which pulsed power is processed in order to create current in the magnets can be exemplified by the supplies at MI-20 (refer to Fig. 3-2 for a schematic overview). The process is similar at all of the major service buildings.

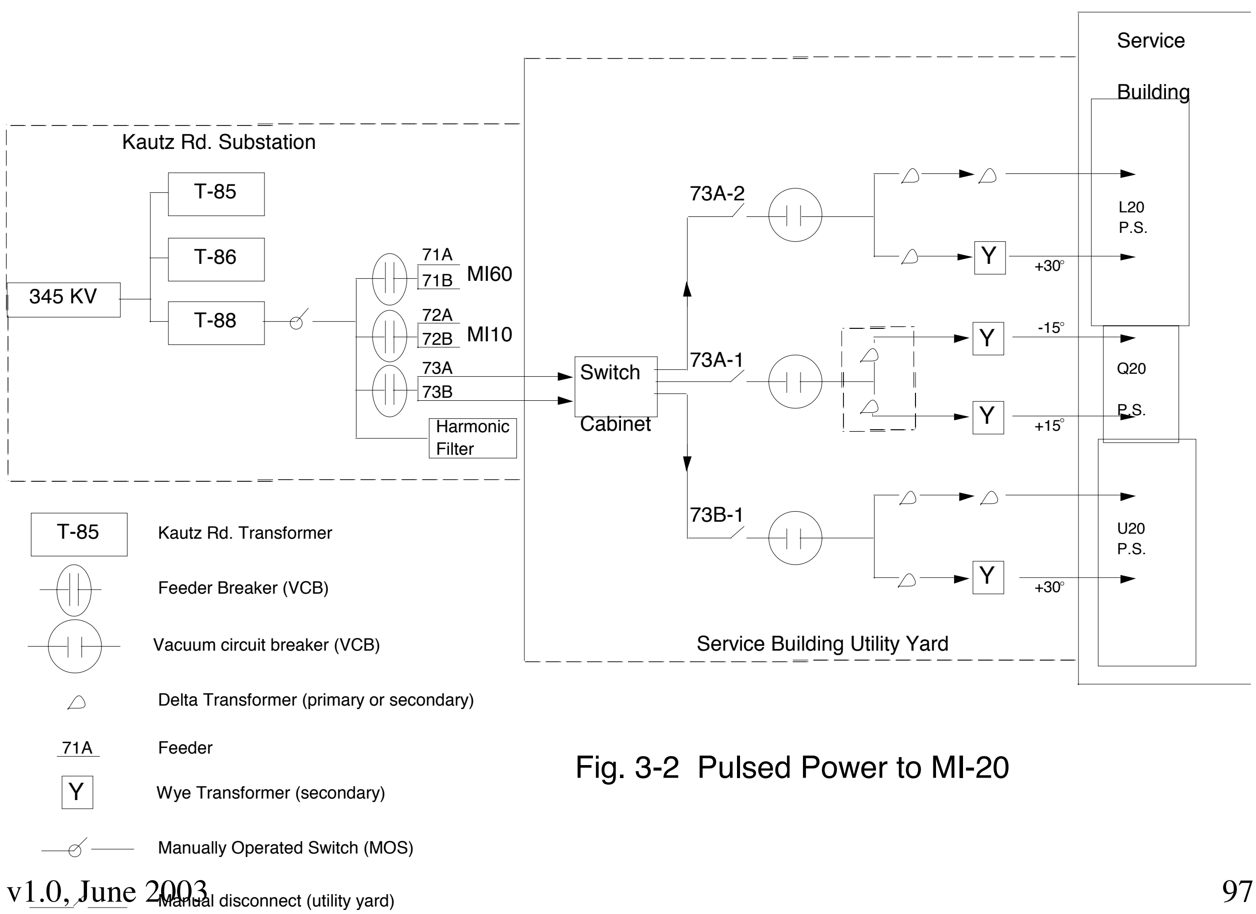

Fig. 3-2 Pulsed Power to MI-20

There are two dedicated pulsed power feeders leading to each service building. Remember that the feeders (73A and 73B) are energized to 13.8 KV. The first surface manifestation of power is the switch cabinet, a large box in the Service Building Utility Yard. Unlike the Tevatron and some other machines, these manual disconnect switches are not visible from the outside. These particular disconnects are used by FESS to reconfigure the feeders, not to isolate the power supplies.

From the switch cabinets the power is initially split three ways, one branch going to each bus (upper dipole, lower dipole, and quad). At this point there are two final opportunities to isolate the 13.8 KV from the power supply transformers. The first is the manual disconnect switch. This switch can be opened if work needs to be done on an individual supply. The second is a VCB, similar to those described earlier. The VCBs are opened routinely before an access or a hipot, and always when the permit loop is dropped. There will be more about those italicized items later in the chapter.

When the VCBs and manual disconnects are closed, power reaches the power supply transformers, which are also in the Utility Yard. The voltage output on the secondary windings of the transformers is about 1 KV. There is some further processing of the AC power by the transformers, but that discussion should be deferred until the power supplies themselves are described

There are six transformers altogether: two each for the upper and lower dipole supplies, and two for the quadrupole supply. (There are, of course, two quadrupole busses in the Main Injector ring, focusing and defocusing, but the power supplies for each bus are only found at every other building. At “even” numbered buildings such as MI-20 the quadrupole supply is dedicated to the focusing bus, and at “odd” numbered buildings it is connected to the defocusing bus.) The quadrupole transformers are both in a single tank, so it appears that there is only one.

AC to DC

The power supplies are located inside the service buildings. Their purpose is to convert the 60 Hz AC from the transformers to a DC voltage. This process is known as rectification (literally, “to make right”). The sine wave originating on the transformers spends half its time at the “correct” polarity and half its time at the “wrong” polarity. The rectification techniques involve (1) reversing the “incorrect” polarity and (2) packing the rectified sine waves so closely together that the output is nearly indistinguishable from a flat DC voltage. This approach is used many times over for power supplies at Fermilab and elsewhere.

The power supplies themselves are inside the service building. The core of the power supplies is a bank of silicon-controlled rectifiers (SCRs). An SCR is a device that, like a diode, allows current to pass in only one direction. An SCR has the additional feature of being triggered; without the trigger, the SCR will not conduct current even if the polarity is correct. If a conducting SCR is reverse biased, that is, if the voltage is reversed, it not only stops conducting but also has to be re-triggered before it can conduct again. Triggering, also known as gating, allows for precise control of the amount of voltage transmitted.

If a sinusoidal voltage is sent through an SCR (assume that it is gated), the waveform will look like this:

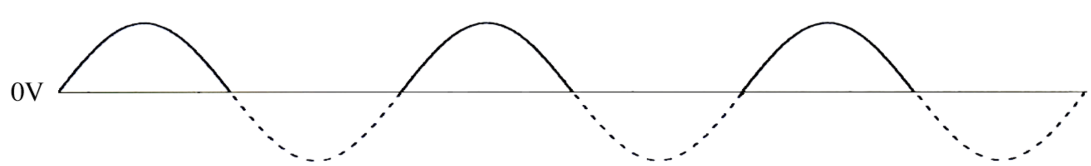

The SCR will only allow the positive portion of the waveform to pass. The dashed curve here indicates the portion of the waveform which was blocked; that is, where the voltage is negative. This type of rectifier is called a “half-wave” rectifier because only half of the available voltage survives the passage.

The specific type of power supply used for the dipoles is known as a full wave bridge rectifier (Fig. 3-3). In this simplified schematic, the transformer secondary is connected to a bridge consisting of four SCRs, and the load—in this case, a string of main dipoles or quadrupoles—is placed across the network.

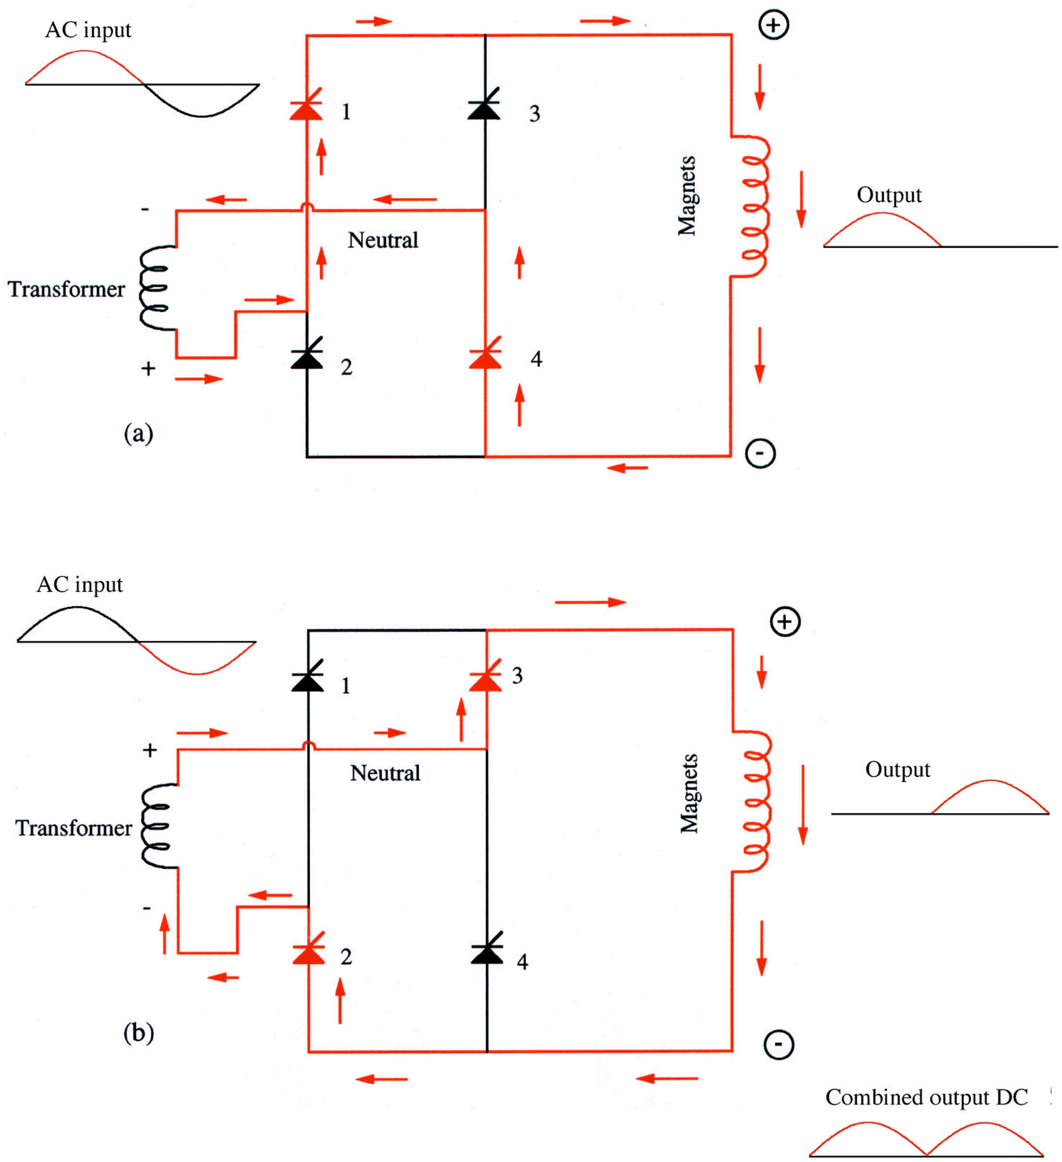

Fig. 3-3 Voltage Rectification using SCR's

This highly simplified diagram shows how the SCR's in the Main Injector power supplies convert an AC input to a DC output. See text for details.

Suppose that the current is to flow through the magnets as shown in Fig. 3-3(a). When the polarity of the voltage is positive (as in the picture above) the top SCR conducts and current flows through the magnets in the proper direction.

When the polarity reverses, as in Fig. 3-3(b), the bottom SCR conducts. The bridge is actually looking at the negative half of the waveform:

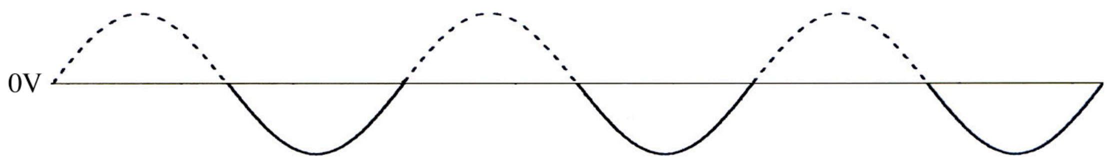

However, because of the way that the SCRs are connected, the direction of the current through the magnets is unchanged. The overall effect is to reverse the polarity of the negative portion of the wave. (The “full wave” designation comes from the fact that the entire wave is used.) The combined waveform is just beginning to resemble a steady voltage:

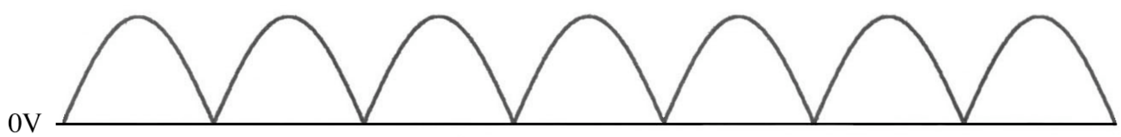

There is a new peak voltage every 180˚ ,but remember that we are only looking at one phase out of three. When all three phases are considered (requiring more SCRs), there is a new peak every 60˚:

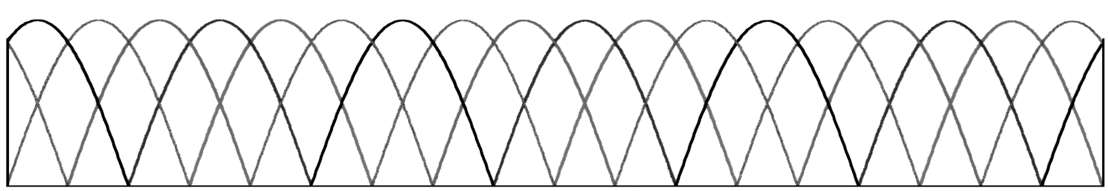

The voltage is now being refreshed at a 360 Hz rate, which is not quite good enough, sort of like driving across a cobblestone road. But the transformers have one more trick up their proverbial sleeve, the primary coil can be connected to two secondary coils instead of one, and the windings built so that the secondaries are shifted in phase with respect to each other. The dipole and quadrupole power supplies differ as to how this is done.

The two transformers on each dipole bus are of two different types. Both have a primary winding called a "delta," named for the Greek letter that it resembles. The secondary windings are "delta" in one transformer and a "wye" in the other. ("Wye" is the letter "Y" spelled phonetically, for ease of pronunciation. It wasn't my idea.) So there is one "delta-delta" and one "delta-wye" transformer for each dipole bus. The "delta" secondary retains the same phase as the primary, but the "wye" secondary shifts the phase by 30˚. The phase shift doubles the number of sinusoidal peaks available to the power supplies, that is, it goes into the primary coil as 3-phase and comes out of the secondary as 6-phase.

The transformers and power supply components for the quadrupole power supplies have been recycled from old Main Ring supplies. Both transformers for a given supply utilize an "extended delta-wye" configuration. In one transformer, the phase of one of the secondaries is shifted forward by 15˚, and in the other the phase is shifted backwards by 15˚. The resulting net shift of 30˚ is the same as what happens in the dipole supplies. The SCR network is similar, except that the individual SCRs are smaller. The chokes and capacitors in the filters are also less impressive than those in the dipole supplies, but the quadrupoles operate at a lower current, so they don't have to be as impressive.

Combining the peaks coming from the three-phase input, the inverted negative portion of the wave, and the 30˚ phase shift, there are now 12 sinusoidal peaks per cycle. Voltage is being refreshed at a 720 Hz rate:

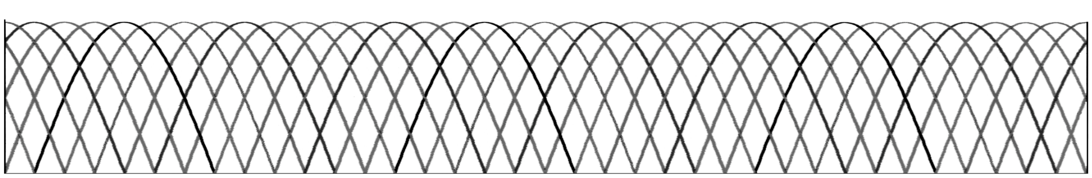

Of course, not all of the voltage is needed all the time; the SCRs are gated according to how much voltage is needed from the power supply at any given instant. The delay, measured along the sinusoidal wave, is called the firing angle:

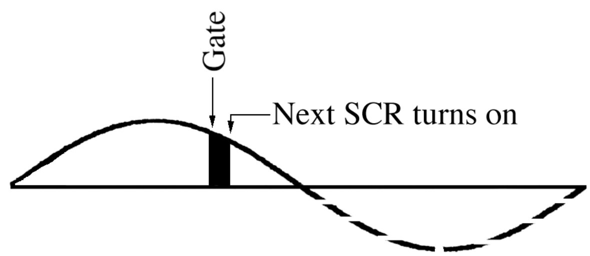

The shaded area represents the time during which the SCR is conducting current. In isolation, the SCR would continue to conduct until the wave naturally reached a zero value. However, soon after the voltage begins to droop, an adjacent SCR fires and begins to conduct current from the next up-and-coming peak. Since the next peak is now at a higher voltage than the first, the first SCR is reverse biased and shuts off.

Remember that Fig. 3-3 is a simplified version of the actual SCR banks. In reality, there is a set of three pairs of SCRs for each of the three phases; then that number is doubled to account for the two secondaries feeding each supply, for 36 SCRs per supply.

This is as much resolution as can be achieved with the existing SCRs. More SCRs could always be added, but it would be impractical; most of the remaining 720 Hz ripple can be smoothed over with the chokes and capacitor banks inside the power supplies. A small amount remains.

The 720 Hz that enters the tunnel is blocked by the inductance of the magnets. A portion of the 720 Hz ripple created by the supplies, even though it is on the secondary side of the transformers, is picked up by the primary coils of the transformers in the utility yard. A 1440 Hz ripple (the first harmonic of 720 Hz) is generated as well. It is these two frequencies (technically, sidebands of those two frequencies), which find their way back to the Kautz Rd.

transformers, that the harmonic filters at the substation are designed to suppress.

There are also bypass SCRs associated with each power supply. The bypass SCRs shunts current away from the load. They are turned on routinely during every ramp cycle, depending on the specific requirements of the ramp at that instant. They are also activated if an unexpected failure in the power supplies or magnets requires that current be removed from the system. The fast bypass loop, described below, monitors conditions and fires the bypass SCRs if necessary.

The power supplies can be connected or disconnected from the magnet bus using the knife switches. These "large" knife switches can also be used to bypass the local power supply, leaving the bus unbroken so that the magnets can still be powered from the other supplies. A third configuration is to open up the knife switches completely, breaking the bus. This last option is used when searching for a ground fault.

There are "small" knife switches as well; they will be described shortly.

A signal from the potential transformer (PT) is also one of the inputs to the SCR firing module. A PT is not a transformer that has yet to live up to its ability, but rather a device for measuring the phase of the incoming voltage. All of the clever rectification techniques described above are useless if the SCRs don't know what the phase is.

The Power Supply Link

The voltage level produced by the power supplies depends entirely on the timing of the SCR firing triggers. If all of the voltage is blocked, there is no output; if none of it is blocked, the supply produces its maximum output. The local power supply hardware, given a voltage request, can calculate and implement the firing angles. The firing angles need to be updated with every new sine wave, at a 720 Hz rate.

In order to get all of the power supplies to produce just the right amount of voltage for the current needed in the magnets, the scheme must be coordinated from a central location. A VME processor called MECAR, which resides at MI60N, transmits the voltage request to the power supplies via the power supply link (PSL). The link originates at MECAR; a cable launches the signal toward MI-10, counterclockwise around the ring. At each of the major service buildings, the signal is intercepted and decoded by CAMAC 269 cards, which then pass the information on to the local power supplies. These instructions are updated at 1440 Hz—twice as fast as the SCR firing angles are updated—with the hope that the increased resolution will carry over to 720 Hz. The local hardware knows how to convert the voltage requests into firing angles.

The only function of the CAMAC crate is to provide power and other basic services to the card; the link is independent of the PIOX, PIOR, and BTR links that CAMAC normally uses.

The health of the PSL link can be monitored from I17.

The scarcity of information about MECAR in this section will be rectified later in the chapter.

To the Magnets

Once the voltage has been produced, it must find its way to the magnets. The dipole busses emerge from the top of the power supply cabinets, surrounded by (but not touching) a metallic shield, and plunge downward into the tunnel near the door to the tunnel stairwell. They are bent horizontally when they reach the alcove downstairs.

The bus work is circular in cross-section when it leaves the power supply and enters the tunnel, but soon has to switch to the 1”X4” rectangle of the magnet coils (see Figs. 2-2, 2-3). The 1”X4” shape is generally retained when the bus has to pass behind the numerous quadrupoles, but reverts to a circular cross-section in the straight sections.

From a local perspective, it would seem that the upper and lower dipole supplies at each building are powering two separate busses, but in reality the busses are continuous because they turn around at MI-60 (see Fig. 2-4, which the reader has surely memorized by now).

The Quadrupole Bus

An important difference between the dipole and quadrupole busses is that there is no fold in either of the quad busses. The focusing and defocusing busses are separate circuits, with current flowing one way in the focusing bus and the other in the defocusing bus. The individual busses require less current than the dipole bus, so there are only three power supplies per bus. The focusing bus is powered from MI-20, MI-40, and MI-60; the defocusing bus is powered from MI-10, MI-30, and MI-50.

MECAR controls the quadrupole busses in addition to the dipole busses.

The quadrupole busses enter the tunnel in a fashion similar to the dipoles, but their cross-section is always circular. Quads of a given polarity are connected with straight lengths of bus; there can be as many as three sections of quad bus running parallel to each other at a given location, and, because of the way they leapfrog, it is not always obvious which is which. Fig. 4-2 illustrates the quadrupole bus connections (in the context of LCW flow). One reason the busywork alternates is to accommodate the curvature of the tunnel; it is possible to use a straight length of bus in most locations. The other, and more important, reason is to allow the inductances and fields from each bus to cancel each other out.

The Permit and Fast Bypass Loops

The permit and fast bypass loops are used to ensure that the power supplies are not allowed to operate under adverse or unsafe conditions. When the permit loop is dropped, power is removed from the system by opening the VCBs; when the fast bypass loop is dropped, the bypass SCRs are turned on so that no current can be passed on to the load.

The permit loop is used for all situations involving personnel safety, because opening the VCB is the surest way to remove power from the system in a hurry. The Main Injector Electrical Safety System (ESS), the power supply door interlocks, and the big red "Panic Button" near each supply are all inputs to the permit loop. Certain overcurrent trips and "sudden pressure" indications are tied to the permit loop as well.

However, opening the VCB can cause transient voltage spikes if there is power flowing through it, possibly causing more damage. Situations that do not require such a drastic response use the fast bypass loop; these include problems such as high temperatures, water pressure, and transformer oil levels. In addition, the bypass loop is dropped whenever there is an unexpected ground fault; it is unwise to open the VCBs because of the voltage spikes just mentioned. The power supplies are in bypass whenever power is not available to the magnets, as when the 13.8 KV is racked out, or the VCB is open.

When turning on the Main Injector supplies, the permit loop must first be made up, and then the VCBs closed, before the bypass loop can be made up. Once the loops are made up, MECAR can be told, via the applications page I2, to begin sending current to the magnets.

The permit/fast bypass control chassis is located at MI-60 South. The chassis has buttons to push to bring the loops up, but fortunately, the applications page I17 provides a software interface to the MCR.

Fault Detection

The current traveling through the main dipole and quadrupole circuits must navigate through dozens of miles of copper busywork. Over the entire path, the bus must be adequately isolated from any conducting material that could cause a fault, including other parts of the bus itself. A turn-to-turn short inside a magnet would, for example, prevent part of the magnet from being powered, something which would obviously be detrimental to the beam trajectory. A ground fault will draw current through components, and two simultaneous ground faults could draw enough power through the magnets to destroy them.

Faults can be generated in several ways: a coil-to-coil short can be caused by a breach of the magnet insulation, or a piece of copper tubing can be left wrapped around two busses after a maintenance day. In the Main Ring, high voltages on the bus would slowly cause copper to be electroplated inside the ceramic insulators that were supposed to isolate the LCW plumbing from the bus. It has been recently discovered in the Main Injector that copper is slowly absorbed by the thermoplastic hoses isolating the quadrupoles from the LCW headers, gradually creating a path for current flow from the bus to the grounded headers. Whatever the cause, it is important to continuously monitor the integrity of the three main busses in the Main Injector.

Ground faults and bus-to-bus shorts are detected through ground fault circuitry that is interlocked to the power supplies. In addition, a procedure called hipotting is employed, particularly following an access, to specifically look for potential faults. All of the main dipoles and quadrupoles are connected to a ground fault/hipot loop, which makes the hipot measurement straightforward.

Coil-to-coil shorts can be found by measuring the inductance across the magnets. This is a tedious and time-consuming process, since the magnets must be measured one at a time, and is not done as routinely as hipotting.

First, ground faults and hipotting.

The Ground Fault/Hipot Loop

The ground fault/hipot loop is used to search for faults on the Main Injector bus. It actually consists of two separate entities sharing the same cables. The cables run from building to building, surfacing in the knife switch cabinets where the magnet bus is connected to the power supplies. There are loops for the QF, QD, and Upper and Lower bend busses; they can be dealt with separately, or in combination. In the knife switch cabinets, the cables are connected to the bus through a 50K resistor and a small knife switch (in contrast to the large knife switches that connect the power supplies to the magnet bus.) The resistor limits the amount of current flowing through the bus in the event of a ground fault. There is resistor and switch on each side of each power supply. Having many different paths to ground—in this case, through numerous current-limiting resistors—is known as a distributed ground system.

The hipot power supply is located in the back racks of the Main Control Room. The hipot controller itself is located at MI-60 South, near the controller for the permit and fast bypass loops. The software interface with the controller can be found on page I17.

Fig. 3-4 shows the hipot cabling for the upper and lower bend busses. This picture should be compared to Fig. 2-4; also be aware that there is a

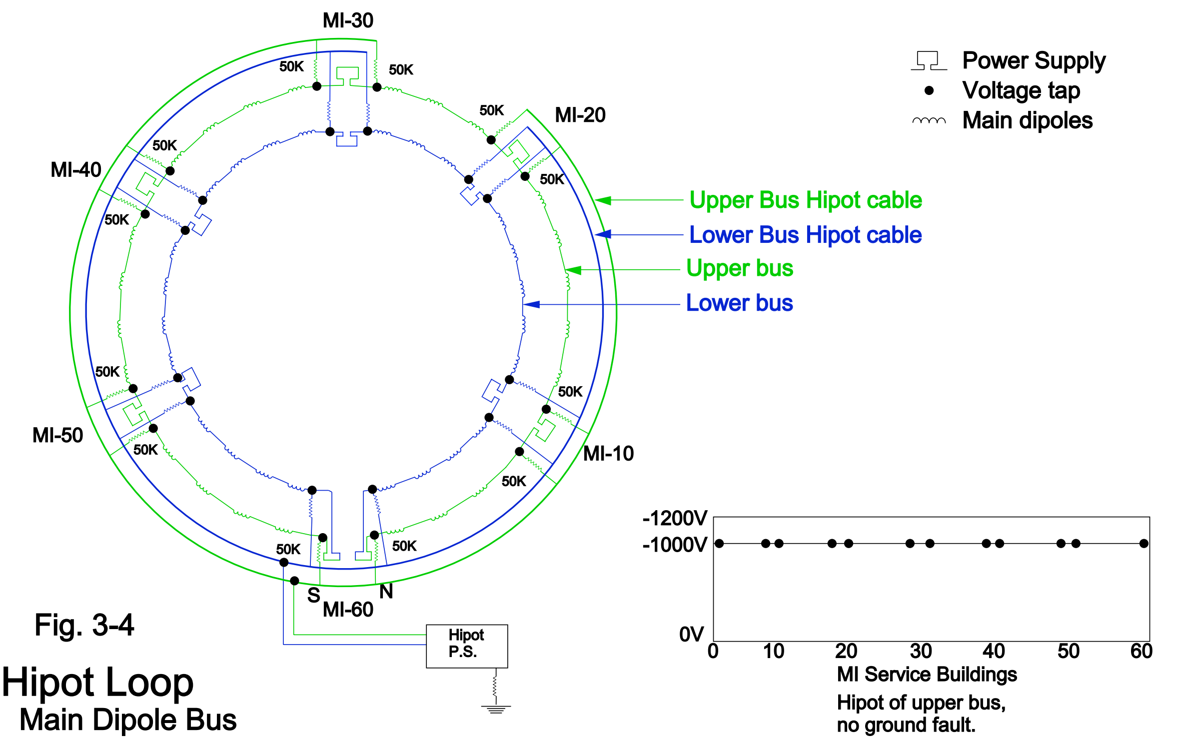

Fig. 3-4

Hipot Loop
Main Dipole Bus

The Hipot loop consists of cables connected to the magnets and their power supplies through 50K resistors. There are voltage taps on the bus side of the resistors. The Hipot power supply places -1000V on the loop. If there is no ground fault, no significant current flows through the resistors, and the voltage taps all see the same voltage as that produced by the Hipot supply.
The Main Dipole Hipot Loop, which includes the upper and lower busses, is shown. There are also Hipot loops for each of the two main quadrupole busses.

similar setup for the quadrupole busses.

All four busses share the ground fault/hipot electronics.

During normal operation, any current flowing to ground will return to the ground fault circuitry (via ground, of course) and will be measured by an ammeter. The ground fault detector circuit is used to catch faults if they occur when the machine is running. If a fault is found, the fast bypass loop is pulled.

To actively hunt for a ground fault, a switch connects the loop to the hipot power supply (the VCBs must be opened before making the switch). The hipotter provides a static voltage (up to 3 KV, although 1 KV is the norm) to the hipot cable. The hipotter is usually first used in the “Bus and Power Supply” configuration. In this configuration, the bypass SCRs are gated. If there is no ground fault, there will still be a small amount of current through the 50K resistors, but no appreciable voltage drop. (The more current flowing through a resistor, the greater the voltage drop.) Voltage taps on the bus to either side of the power supplies measure the voltage on the bus. If there is no voltage drop, the bus will be at the same voltage as the hipot loop. The voltage from the taps can be graphically displayed from page I17.

If, however, the bus is ground faulted at some point, current will be drawn through the resistors, and the voltage drop will be noted on the taps. Since the resistors are in parallel, a hard ground fault will mean that none of the taps will read any voltage.

If that happens, the next step is to go to the “Bus Only” configuration (Fig. 3-5 (a) and (b)). In this state, the bypass SCRs are not gated; a ground fault on one side of the power supply will pull down the voltage on that side, but not on the other. Since the situation is similar at the adjacent power supplies, the ground fault is now isolated to a region between the two power supplies.

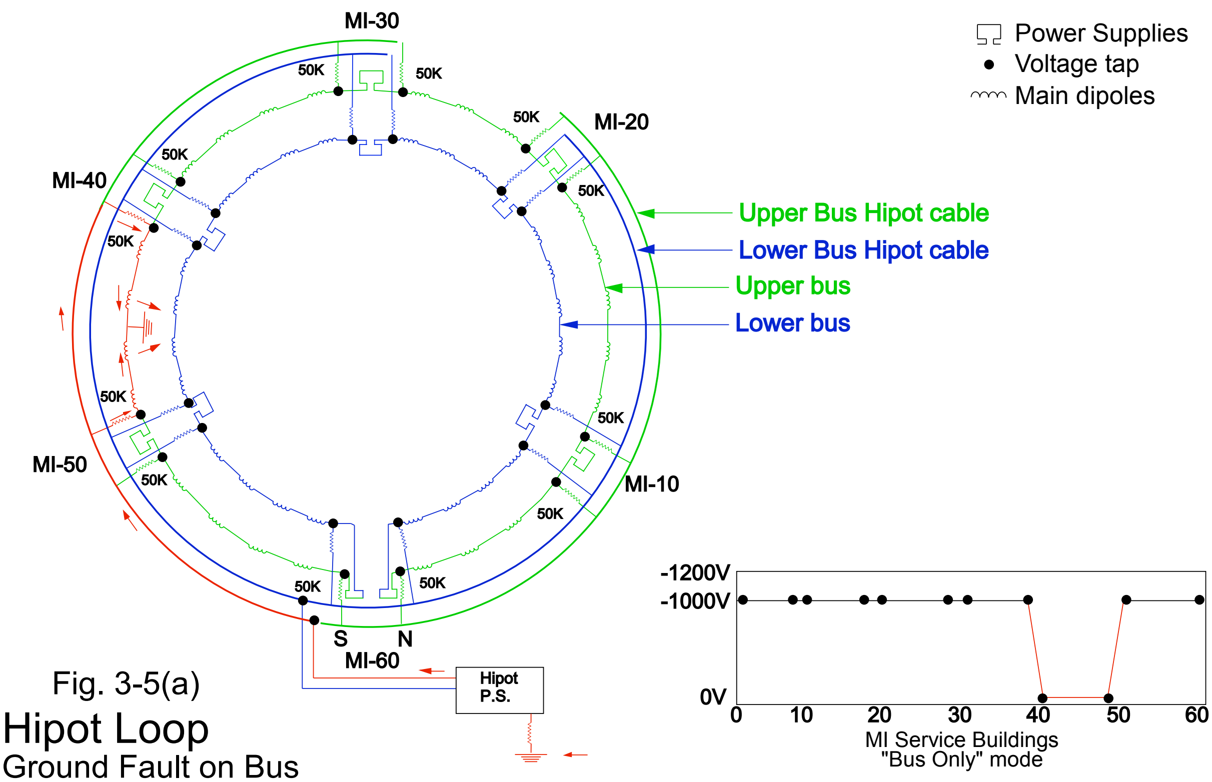

Fig. 3-5(a)
Hipot Loop
Ground Fault on Bus

A ground fault has developed on the upper bus. Current is flowing to ground through the 50K resistors to either side. The hipotter measures the current. In this case, enough current is flowing to cause all of the voltage to be dropped across the two resistors, as shown in the voltage plot at right.

This isolation technique may not work, however, if one of the transformer secondaries or some of the SCRs are shorted. To completely isolate the sector, both the large and small knife switches must be opened.

The other two hipot options look at coil-to-coil faults. In the P2 mode, the upper bus is grounded and the lower bus hipotted; any short between the busses is interpreted as ground current. In the P3 mode, the lower bus is grounded and upper bus hipotted. A bad result could mean that the busses are touching, or (worse) one of the magnets is internally shorted.

Gadgets for checking the direction of ground current in the tunnel include Groundhog and the Blacklister. Groundhog operates with a 45V supply while the Blacklister places up to 1 KV on the bus.

Inductance Measurements

Occasionally a magnet will develop a coil-to-coil short that does not produce a ground fault. If only a small part of a main dipole or quadrupole is shorted, the effect on the beam may be subtle enough for the cause to go undiscovered for months; much troubleshooting of poor beam quality may occur without finding the cause. Each type of main dipole or main quadrupole has a predicted inductance based on the geometry of its coils. Inductance probes, which have to be manually placed on the magnets one at a time, can measure the inductance. If the inductance of a magnet falls short of the ideal, it may mean that a short has disabled part of the coil, and the magnet will have to be replaced.

An inductance probe works by generating a high frequency signal and measuring how it is propagated through the magnet. A switch, usually located behind the magnet in the most awkward location imaginable, opens the mode-damping resistor so that it does not interfere with the measurement by attenuating the higher frequency components of the signal.

MECAR

MECAR stands for “Main Injector Excitation Controller And Regulator” (the superfluous words are included to make the acronym easier to pronounce, as Mee-kar). MECARs job is to orchestrate the current in the main dipoles and main quadrupoles; it physically inhabits a VME crate at the north end of MI-60. (There is a spare MECAR in the same rack. One is known as MECAR A, the other as MECAR B.) The software in MECAR controls the amount of current the magnets will see by regulating the voltage produced by the power supplies. Some of the issues are discussed in the following sections.

Momentum, Current, Voltage

Current in the dipoles is frequently expressed in terms of beam momentum rather than amperes, because the magnetic field is more or less proportional to current, and the beam has to be within a narrow range of momentum if it is to be constrained by a given magnetic field. The two concepts diverge somewhat in the Main Injector because the saturation of the magnets at high currents distorts the proportionality between current and momentum.

A note to purists about the following discussion: ask your average Joe standing in the grocery line about the relativistic properties of subatomic particles, and he will probably express his feelings in terms of energy, as measured in electron volts, without giving it a second thought. The MECAR tables, however, list the momentum of the beam in GeV/c. As mentioned in Chapter 2, relativistic momentum is the kinetic energy, in electron volts, plus the rest mass of the particle, which is also measured in electron volts. The rest mass of a proton or antiproton is about 938 MeV. MECAR lists the momentum at injection as 8.889 GeV/c. That is the same as a kinetic energy of 8 GeV. (All right, the numbers don't quite add up. This business of Momentum = Kinetic Energy + Rest Mass is really just an old folk tale, but one based on fact. It becomes increasingly accurate with increasing energy. 8 GeV just isn't quite relativistic enough to hide the difference.) Anyway, the units are almost arbitrary. Before the request for momentum can be put on the power supply link, momentum has to be converted to magnetic field strength, the field strength has to be converted to the equivalent current in the magnets (adjusted for hysteresis and saturation), current has to be converted to voltage (adjusted to compensate for the inductance of the magnets), and voltage has to be converted to a complex set of SCR firing angles (subject to a resolution of 720 Hz). In an attempt to mimic everyday banter, this text will carelessly bounce back and forth between energy and momentum, but the reader should know that there is a difference.

The applications page I2 is the user interface to MECAR. I2 includes time vs. momentum tables, which dictate the amount of current in the main dipoles, as well as time vs. tune tables, which ultimately determine the current in the main quadrupoles. As a convenience to the tuner, I2 also includes sextupole and octupole tables, but, as will be seen later in the chapter, these two types of magnets depend on CAMAC 453 cards rather than MECAR.

The various calculations are divided between I2's software and MECAR. First, I2 figures out what the magnet current needs to be, based on the desired beam momentum. Detailed conversion tables are then consulted that take saturation into account. (The relation between momentum and current is almost proportional up to 120 GeV or so, after which the increase of the field slows significantly compared to the increase in current.) Some compensation is made for hysteresis as well, making certain assumptions about previous ramps. When the calculations are complete, I2 sends a table of desired current to MECAR.

It is MECARs responsibility to calculate exactly what voltage is needed to produce the desired current. To do that, it needs to know the inductance of each segment of the ring between the power supplies. The di/dt can be very high, up to 10,000 amps/second, so the inductive impedance can be very large during those times. Up to 750 volts per power supply may be required to keep the ramp going.

Page I2 has an additional step to complete for the quadrupoles. The tunes are listed in the tables. The gradient field establishes the tune, which is measured in, say, kilogauss per meter. However, the beam momentum is changing during the ramp; the gradient has to change to match the beam momentum just to keep the tune constant during the cycle. So, the quadrupoles have "baseline" ramps to match the momentum. I2, when calculating the current, increases or decreases the current request based on whether the requested tune is above or below the baseline curve.

Regulation

The calculations of I2 and MECAR get the current close to where it should be, but reality soon takes its toll. In such a large system there are hundreds of variables, temperature being the most important, that can influence the final value of the current. There are several mechanisms available for controlling the outcome:

- The individual power supplies, like all decent power supplies, have internal feedback to ensure that they are producing voltage as requested. They compare their output with the voltage request arriving from the power supply link, and adjust the SCR firing angles appropriately.
- MECAR monitors the current in the busses and adjusts the voltage program accordingly, in real time. There are four transductors at MI-60: one for each quad bus, and one each for the upper and lower bend busses. (Only one of the bend bus transductors is used at any given time, since they are in series with each other.) MECAR reads the current on the three busses and implements a fast feedback loop to adjust the voltage as necessary. Any one of the 12 bend supplies can be selected as the bend bus regulator. Each of the quad busses also has a designated regulator. The fast feedback information is sent only to the regulators; the remaining supplies continue to play the basic pre-calculated waveform.
- If a ramp of a given type continues to deviate from its expected output during a cycle, MECAR remembers the error and applies a correction to the next cycle. This is called learning. These updates are semi-permanent and reduce the need for fast feedback. Updated voltage information is sent to all of the power supplies, not just the regulators.

The overall voltage request for a given cycle is called a profile. By "saving" a profile, the learning updates are assimilated by the profile, although they will be lost if MECAR is rebooted. Sometimes, however, the profiles are deliberately cleared of their updates. This is often done after MECAR learns in garbage (a universally accepted scientific term) when the ramp is faced with less-than-optimal conditions.

Ramps

The discussion in this section will emphasize the dipole (i.e., momentum) ramps. As mentioned earlier, the quadrupole ramps are generally proportional to the momentum ramps, with adjustments being made for changing tunes.

There are 8 GeV "ramps" where the beam momentum doesn't change, but all the other ramps share some major milestones that occur during each cycle.

Injection always takes place at 8 GeV. Occasionally the amount of time that beam is circulating at 8 GeV is referred to as the dwell time.

During parabola, the current in the dipoles begins to increase. As implied by the name, the increase starts out slowly, the change in current (to a first approximation) literally following a parabola. As the current increases, the RF acceleration systems are changing the beam energy to match the strength of the magnets.

Near the top of the ramp, an inverted parabola eases the transition to flattop. The current in the magnets, and therefore the beam energy, does not change during flattop. The two flattop values used for Main Injector are 120 GeV and 150 GeV. During flattop, beam is at its highest energy for the cycle and is usually extracted from the Main Injector. Flattop is also the time when some of the more complex RF manipulations take place, such as bunch rotation and coalescing, but even if nothing special is done, the RF must still be present at some level in order to maintain the bunch structure of the beam and to compensate for energy lost to synchrotron radiation. (Synchrotron radiation is a relativistic phenomenon and will be discussed in the RF chapter.)

After flattop, the ramp is driven into invert. The current in the magnets is taken back down to the 8 GeV value, and then some. Just because the energy is decreasing does not mean that invert is a passive process; the considerable inductance of the magnets would cause the current to linger for an unacceptably long time. It is not enough to shut off the voltage, the power supply voltage must actually be reversed during invert to force current out of the magnets.

The magnet current during invert is driven below the 8 GeV level in order to "reset" the hysteresis, ensuring that every cycle will begin with the same magnetizing force.

Finally, the magnet current bounces back to the 8 GeV level and the next cycle begins.

The total amount of time required for a complete cycle varies from about 2 seconds, in the case of antiproton production, to several seconds when injecting coalesced bunches into the Tevatron. Each type of ramp is designed to complete all of its necessary tasks as quickly as possible.

MECAR, then, must control three different ramps: those of the main dipoles, the main focusing quads, and the main defocusing quads. After making all of the appropriate calculations, the commands are sent to the power supplies via the power supply link. All phases of the ramp cycle (dwell time, parabola, linear ramp, flattop, and invert) are implemented at the power supplies by the timing of the SCR triggers.

The tables for the ramps can be found on the MECAR page, I2. There is a separate operational table for each type of reset ($20, $21, $29, $2A, $2B, $2D, or $2E). Backing every operational table is a set of files. The tables determine the ramp in terms of time vs. momentum.

It is worthwhile to analyze a specific ramp using numbers from an actual MECAR table. The following discussion follows the fate of the main dipole magnet current over the course of a cycle. The $21 cycle, which accelerates five or six batches to 150 GeV, can be used as an example. The numbers in this example have been taken from operational tables, so they can be considered realistic, but not immutable, since they tend to be adjusted over time. Remember that these tables are reset-dependant; for example, a $29 ramp will include many of the same basic features as the $21 but will differ in many of the details.

The ramp has a dwell time of .55 seconds, during which the beam momentum is held constant at 8.889 GeV/c. That is sufficient time for all of the batches to be accelerated in the Linac and Booster, and to establish a stable orbit in the Main Injector. On I2, this line is labeled INITI, for "initial."

Remember that there are twelve dipole power supplies altogether. At 8 GeV only the regulator is producing current.

At .55 seconds, the current begins to slowly increase. Between 0.55 seconds and 0.5855 seconds, the momentum changes from 8.889 GeV/c to 8.96 GeV/c. This is the first part of the parabola, and of course, the rate of change in a parabola is slowest in the beginning and speeds up over time. This line is labeled VPARAB.

Notice that there are three numbers to the right of the "momentum" column. These are the derivatives, or rates of change, of the momentum. "Pdot" is the first derivative, or the rate of change of momentum. Pdot is measured in units of GeV/c per second if looking at momentum, or amps per second if looking at current. At .5855 seconds, the rate of change is 6 GeV/c per second.

To the right of the "Pdot" column are the second and third derivatives of momentum, Pddot and Pdddot. When Pdddot has a nonzero value, it means either that the ramp waveform equation has added a cubic term, or that there is a problem with the "d" stroke on the keyboard.

The regulator is the first supply in the turn-on order. The turn-on order—the sequence in which the supplies are turned on—is flexible and can be configured from Page I2. As the demand for current increases, the regulator needs to be supplemented by some of the other supplies. The supplies are organized into tiers, with each tier comprised of a group of power supplies. Like the turn-on order, the tier structure can be modified, but it is important to distribute the supplies evenly around the ring to prevent the voltage to ground from becoming too high at any given point. For example, if the first tier consists of 6 supplies (including the regulator), the order might be Lower 60 (L60), U60, L20, U40, U20, and L40. Except for the regulator, which is already on, these supplies phase in simultaneously so that the current is getting a boost at regular intervals around the ring. When those supplies approach their maximum output, the next tier (which may consist of either three or six supplies) begins to turn on and fill in the gaps. The last power supplies to phase in will also be the first to phase out when the extra voltage is no longer needed.

As discussed in Chapter 2, the dipole busses have considerable inductance and will resist the change in current. The change in voltage has to be much steeper than the desired change in current. Since the inductive voltage is proportional to the change in current, Pdot is an indicator of how hard the supplies will have to turn on.

Remember that at the hardware level, any change in the voltage output has to be implemented by adjusting the SCR firing angles. At 8 GeV, the supplies can be on, but delaying the trigger gate may block their voltage output. By reducing the delay, more voltage is gated through and produces more current.

At line 3, PARAB, the momentum is 9.5 GeV/c and Pdot is 20 GeV/c per second. Clearly, things are picking up. By line 5, the beam momentum is 85 GeV/c and Pdot is 188 GeV/sec, but Pddot has a negative value (-185 GeV/c/sec^2); the rate of rise is beginning to slow down. At 2.3217 seconds, the momentum is finally at 150 GeV/c and Pdot is zero, in other words, flattop.

The ramp remains at flattop for nearly half a second, to 2.7292 seconds. That is sufficient time for transfer cogging (an RF process) and transfer into the Tevatron. Transfer currently takes place at 2.6 seconds; the Tevatron may have already quenched by the time the ramp starts down. The cleanup abort is normally scheduled to occur sometime between 2.6 seconds and the end of flattop.

Once beam is out of the machine, during invert, the momentum rapidly plunges toward its minimum value. The trip down is faster than the trip up because without beam there is no need to maintain any sort of field quality. The magnets only have to be at their minimum value for an instant to set the hysteresis, and they are brought back to their 8.889 GeV/c level before the next cycle begins.

As mentioned earlier, during invert the SCRs are configured so that the polarity of the applied voltage is reversed, not just eliminated, so that current will not linger in the magnets. After all, we may be in a hurry to move on to the next pulse. Pdot assumes large negative values during this phase.

Since this chapter deals with power supplies and their regulation, the use of the tune sub page on I2 as a tool for controlling the beam will be deferred to a hypothetical chapter on "Beam Tuning."

The sextupole and octupole tables are also kept on I2. It is important to recognize that these are not controlled by MECAR; the ramps have been placed on the same page for the convenience of the tuner.

MDAT Generation

MDAT stands for "machine data." The data carried by MDAT is broadcast widely on the MDAT link and can be listened to by any device. MECAR, since it is already full of information about the Main Injector ramps, encodes some of that information onto MDAT. The MDAT data frame is updated at a 720 Hz rate, which is of course the same rate that the SCRs in the power supplies update. Each type of data within the frame is assigned a number, and there is a dedicated parameter for each type as well. The data generated by MECAR is based on the main dipole bus; the devices that read MDAT use the information as a baseline for calculating their own ramps. Some MDAT signals generated by MECAR include:

MDAT30: Main Injector beam momentum, in GeV/c. Remember that there is not a strict linear relationship between energy and current in the main dipoles, because the magnetic domains in the dipoles saturate at higher energies. The beam momentum is calculated from the program current based on the saturation tables stored in I2, the applications page. It is MDAT30 that the correction elements watch when playing their energy ramps (see below). If MECAR and MDAT are alive and well, this signal will be present even if there is no power to the magnets.

- MDAT31: Program Pdot. “Pdot” is the rate at which the momentum is changing. Program Pdot is a value calculated by MECAR from the program momentum. It can be useful in determining if the slew rate for a particular ramp is too high, and is used in the sextupole ramp tables.
- MDAT40: Measured current, or what the current is actually doing. It is measured by a transductor; MECAR reads the information to use in its feedback algorithms and then formats the data for broadcast on the MDAT link.
- MDAT41: Measured Idot, calculated from measured current.

MDAT uses the same fiber-optic repeater chassis as the other links. More information on MDAT in general can be found in the Controls Rookie Book.

Corrector Element Power Supplies

The correction dipoles, trim quads, skew quads, and octupoles in any given region all draw their power from a single Corrector Power Supply (CPS). The sextupoles are powered separately, as explained below. There is one CPS at each of the major service buildings, with two at MI-60. (Those who remember Main Ring should consider the CPS analogous to the two bulk supplies for the old correctors. Those power supplies and magnets have been retained for use in the MI-8 line and the Recycler.) The CPS supplies, and therefore the correctors, are powered from house power and not pulsed power.

Fig. 3-6 is a representative layout of the modules at a typical service building, MI-30 and Fig. 3-7 is a simplified schematic representation of the CPS. The devices at MI-30 will be used as examples repeatedly in the next few paragraphs and diagrams, not only to heighten the sense of individual drama, but also to provide continuity of the descriptions in a specific context.

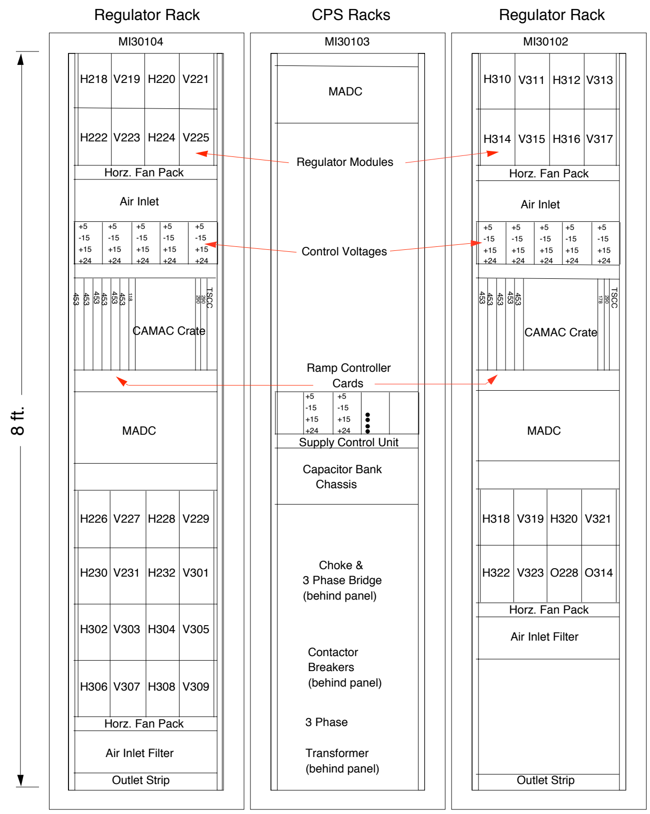

Fig. 3-6
Corrector Power Supplies at MI-30

The CPS first feeds power to a number of switching regulators (only one of the switching regulators connected to the CPS is included in Fig. 3-8). Each regulator is dedicated to a single magnet or a short string of magnets. The specific job of the CPS is to provide a constant 180V to each of the regulators, and the regulators then send the proper current to the magnets. A CPS is capable of powering up to 48 regulators, but in practice no more than 44 are actually connected.

The CPS uses 3-phase power. The three phases are first sent through a slow start circuit, that is, a small amount of current is applied at first to ensure that the phases are all matched up properly. When the circuit is energized, the slow start contactors are closed so that resistors limit the current. After a tenth of a second, the main contactors close. Since they have very little resistance compared to the slow-start resistors, most of the current then passes through them and on to the supply.

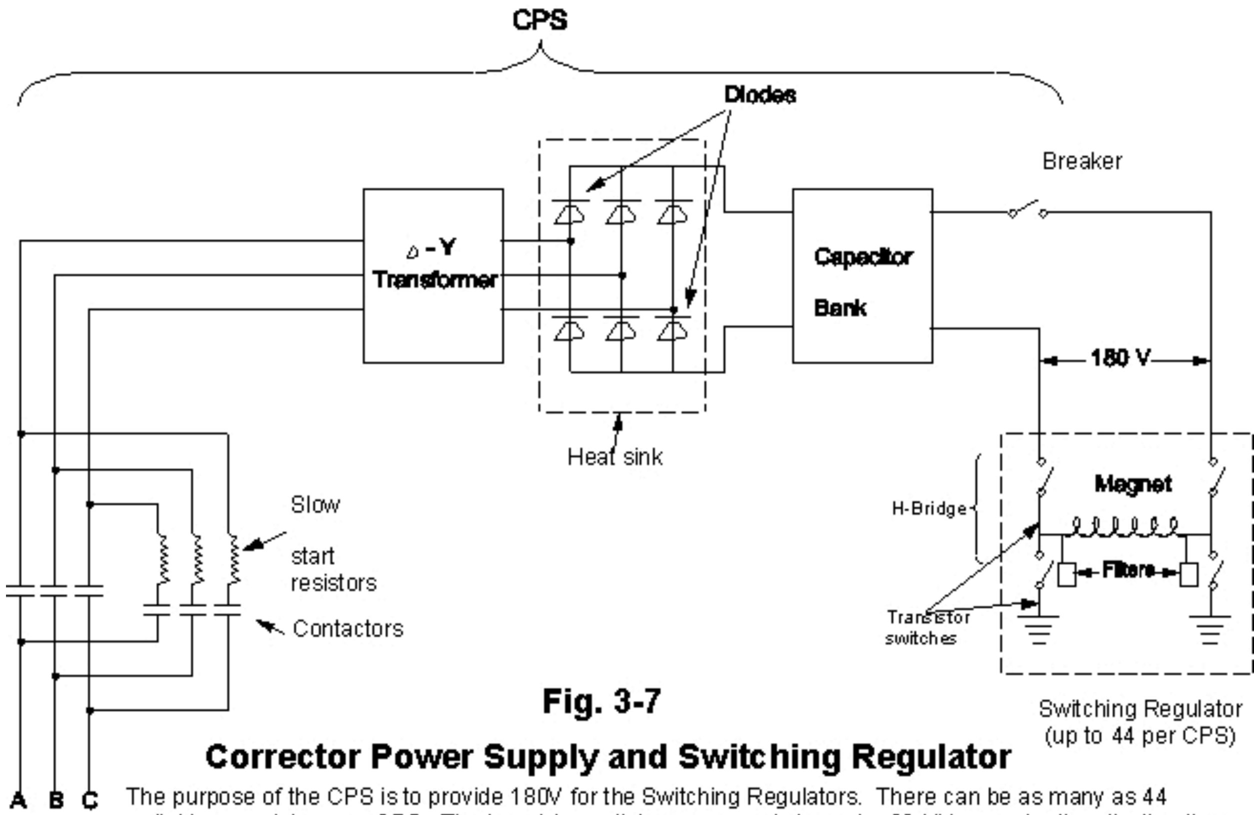

Fig. 3-7

Corrector Power Supply and Switching Regulator

The purpose of the CPS is to provide 180V for the Switching Regulators. There can be as many as 44 switching regulators per CPS. The transistor switches open and close at a 30 KHz, constantly adjusting the amount of current flowing through the magnet.

The 3-phase passes through a delta-wye transformer and is rectified by a bank of diodes. (Diodes are much like SCRs in that they only allow current to pass in one direction; unlike SCRs, they are not triggered.) The diode bank needs to be kept cool, so it sits on a heat sink.

The DC current from the diode rectifier charges a capacitor bank, which in turn provides the 180V needed by the regulators.

There is a dedicated cabinet for each supply; the manual disconnect for the supply is mounted on the back door of the cabinet, to ensure that the supply is racked out before accessing the cabinet.

The CPS is air-cooled by three fans at the top of the rack. Unlike many magnet power supplies elsewhere, LCW is not used.

The switching regulators are the silver-colored modules found in adjacent racks. The design of a regulator is based on an H-Bridge, in which the magnet or magnet string forms the horizontal bar of the “H.” The switches—

actually transistors—are found at the four poles. The switches can open and close at a 30 kHz rate, constantly adjusting the voltage applied to the load as needed. This setup allows bipolar current to flow in the magnets and updates quickly enough to keep up with the rapid changes demanded by the ramp rate.

The correction dipoles and trim quads each have a one-on-one relationship with their personal regulator. The skew quads, which as you remember from Chapter 2 are clustered into four groups of four, are assigned one regulator per cluster. The regulators assigned to the octupoles can be assigned between one and five magnets, depending on the location. The specific power supply configuration for each type of corrector is graphically displayed in a series of pictures at the end of Chapter 2.

On the front panel of each regulator is a Load Compensation Switch. The magnet loads vary from place to place, especially with the number of magnets in the string. The Load Compensation Switch anticipates what the total resistance and inductance of a string will be. The switch is not automatic; it is important to set the switch properly when replacing a regulator.

The ACNET parameter name for a given CPS is based on its location, I:CPS10 at MI-10, I:CPS20 at MI-20, etc. The parameters for the supplies at MI-60 are I:CPS60N and I:CPS60S, for "north" and "south." The ACNET names for the regulators are named after the magnets, e.g. I:H310. If there are several magnets in series using the same regulator, the supply is named for the first magnet in the string.

Ramp Controller Cards

A ramp waveform is a series of current values, played out over the course of a machine cycle and sent to an individual regulator. The waveforms for most of the correctors are stored and played from CAMAC 453 cards. The 453 card is one of a series of ramp controller cards widely used at Fermilab. The applications page I14 consolidates access to all of the ramp tables in the 400 series CAMAC cards, including those of the Main Injector, Booster, Tevatron, and the Antiproton Source.

Each 453 card is capable of storing the waveforms for up to four individual devices; each device can theoretically be assigned up to 48 waveforms, each waveform appropriate for a particular magnet.

One way in which the ramps are listed is by clock event. There can be up to 16 different clock events for each device, designated A through P. Any clock event can be assigned to any slot, but the list is usually standardized to minimize confusion. The ramp begins to play out its program when the 453 card detects the appropriate clock event, carried on TCLK. Most of the clock events used are those for Main Injector resets, e.g., a $29 for a stacking ramp, or a $21 for Fixed Target Tevatron injection. One clock event, a $26 (end of beam operations), is put into Slot "L" to help smooth the transition to the beginning of the next ramp.

Within each clock event slot, there are three ways of defining a ramp: (1) A ramp that is a function of time, or f(t); and two additional functions, g(M1) and h(M2), where M1 and M2 are selected MDAT signals. The MDAT signal used almost universally for correctors in the Main Injector is MDAT30, the Main Injector programmed momentum. MDAT30 is generated by the applications software of Page I2 and launched from MECAR. The f(t), g(M1), and h(M2) ramps can be superimposed if desired.

The f(t) ramp is the easiest to understand, and the least used. The ramp plays when the clock event is received, oblivious to any further instructions. In the "t" column, each slot is given a time interval. In the f(t) column are the current values the power supply is expected to reach during that time interval. The card calculates the rate of change needed and sends the appropriate signal to the regulator.

The 453 cards use MDAT functions most of the time. Specifically, MDAT30 is assigned to the M1 function so that the regulators produce a current that is directly related to the beam energy. In that way the waveform stored in the table is still appropriate even if the ramp timing is modified.

There are applications programs, such as those for smoothing the beam orbit, whose primary function is to adjust the 453 ramps. These programs will

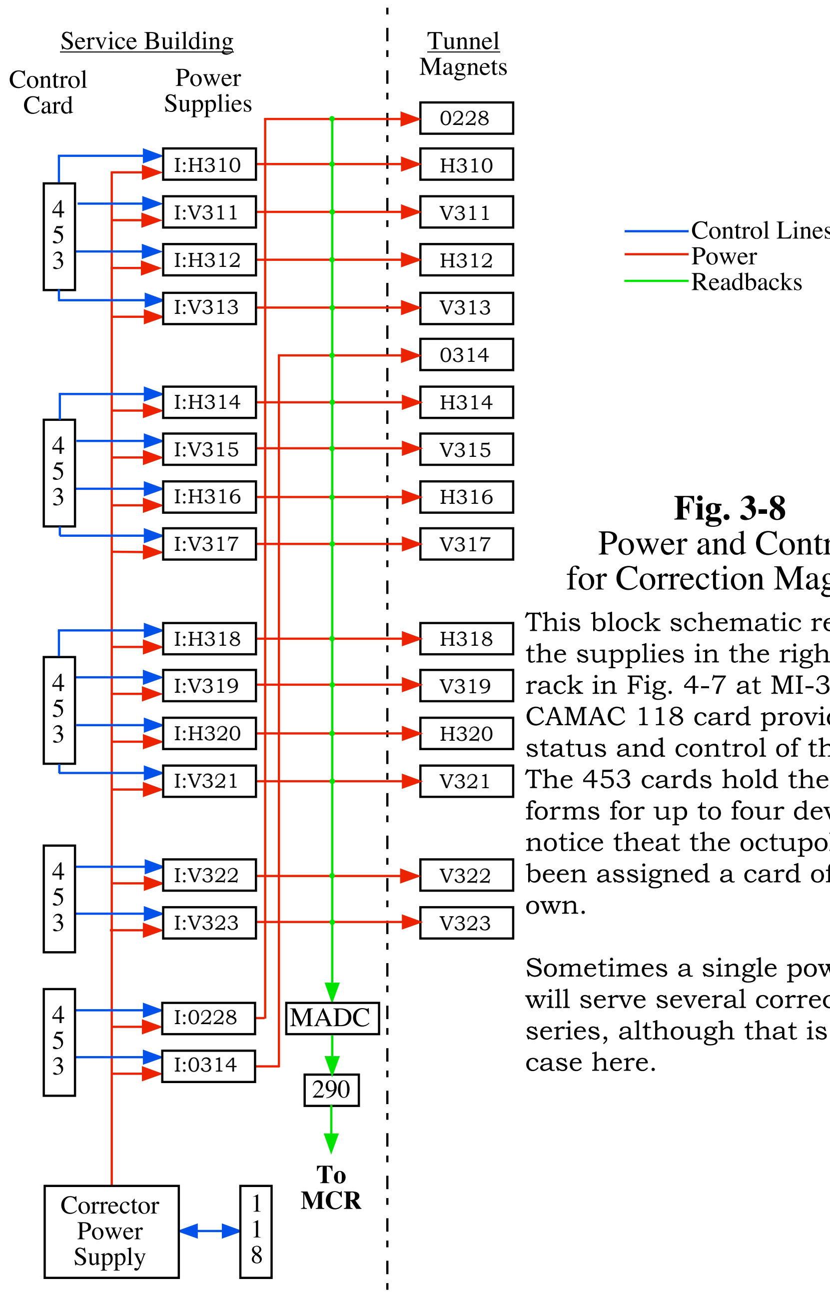

Fig. 3-8
Power and Control
for Correction Magnets

This block schematic represents the supplies in the rightmost rack in Fig. 4-7 at MI-30. A CAMAC 118 card provides digital status and control of the CPS. The 453 cards hold the wave-forms for up to four devices; notice theat the octupoles have been assigned a card of their own.

Sometimes a single power supply will serve several correctors in series, although that is not the case here.

be discussed in a mythical chapter called “Beam Tuning.”

Fig. 3-8 schematically shows the control lines and power distribution to the regulators in the right-most rack at MI-30. The 453 cards each talk to four power supplies, except that there aren’t quite enough dipoles to go around.

The last card with dipole tables only has two dipole correctors under its control. The final 453 card has two octupoles assigned to it. Although it would have been possible to put those two octupoles on the previous 453 card, the octupoles are always segregated from the other devices because they are ultimately controlled from a different source (page I2).

Often, as with the skew quadrupoles and most of the octupoles, a single regulator may drive several magnets in series. The two octupoles at MI-30, however, happen to be the only devices powered by their respective regulators (Fig. 2-24).

The digital interface to the CPS is a CAMAC 118 card. On and off commands are issued through this card, and the digital status information described below comes back to the MCR through the card.

Language Traps

Some of the corrector magnets, of course, are quadrupoles.

The 453 cards are labeled as "Quad Controllers." Here, "quad" does not mean that the cards necessarily control quadrupole magnets, but that they control four devices.

The regulators are called "Quadrant Switching Supplies." This does not mean that the regulator is necessarily powering four devices; the term refers to the four poles of the H-Bridge.

One can only imagine what electrical engineers are like on a golf course.

Interpreting Digital Status

When a CPS or regulator power supply trips off, the cause can usually be determined from the extended digital status (I:CPSxx) found on, for example, S53. The string "xx" refers to the house number.

Some of the status bits for the CPS are:

- Power Supply: This bit refers to the control voltages from the Control Unit in the middle of the CPS rack.
- Slow Start: This bit actually indicates that the main breakers of the CPS, which are supposed to close in a tenth of a second after the slow start breakers, actually do so. If they don't, current will continue to flow in the slow start resistors until they overheat. A Klixon on the resistors opens and toggles the bit.
- RMS Overcurrent: "RMS" stands for "root-mean-square," and is a standard method for measuring an average. If the average current produced by the CPS over the course of an entire cycle exceeds 120 amps, it will trip. The average current can be monitored through the parameter I:CPSxxI. An RMS overcurrent trip can often be fixed just by reducing the current in some of the regulators.
- Door Interlock: If the back door of the CPS rack is opened, the supply will trip. The door is also equipped with a manual disconnect which must be opened before entering the cabinet.
- Cooling Fans: The CPS is interlocked to the three fans at the top of the rack through motion sensors.
- Transformer/Choke: These components of the CPS are interlocked to the temperature through klixons.
- Reset: Clears appropriate interlocks on the CPS, and after a short delay sends resets to the regulators. The regulators are set sequentially in groups of eight. At MI30, for example, the CPS would first reset a group beginning with H218, followed by groups beginning with H226, H302, H310 and finally H318. The sequence can be watched on a string of LED's when resetting the CPS locally.

Some of the status bits for the regulator supplies (e.g. I:H310) are:

- Ramp Enabled: The ramp waveform generated by the 453 card can be enabled or disabled by toggling a bit on the 453 table page. It is also possible for the ramp to become disabled because of a failure on board the card.
- Tracking Error: A "tracking error" means that the actual current from the supply deviates from the ramp waveform playing from the 453 card. The tolerance allowed is about a quarter of an amp. Tracking errors can be diagnosed by plotting the actual current (e.g. I:H310) vs. the reference voltage (e.g. I:H310F). (In the Main Ring the 453 cards were updated at a 15 Hz rate, and tracking errors were common. The cards have now been upgraded to update at 720 Hz, which is expected to alleviate the problem.)
- Analog alarms are impractical for corrector supplies, because the amount of current can legitimately be anywhere within a large range. However, a digital tracking error will post an alarm if a regulator is not producing current or has other serious problems following the ramp program.
- DC Overcurrent: The regulator supplies are capable of producing up to ±20 amps. A current limit of about 19 amps is set in the hardware that will trip the supply if exceeded. This bit is keyed to the instantaneous value of the current rather than to an average. The problem can often be diagnosed by plotting the output of the supply to see if it crosses over the limit at any point.
- Ground Fault: Obvious.
- RMS Overcurrent: This is the same as for the CPS, an average current over an entire cycle. The limit for an individual regulator is 12 amps in the case of the dipoles. As with the CPS, an RMS overcurrent for an individual regulator supply can often be cleared by reducing the overall current that it produces over a cycle.
- Low Input Voltage: This means that the regulator is not seeing enough voltage from the CPS, which it reads from the capacitor bank. (When turning off the CPS locally, the LED for the input voltage on the regulator

supplies can be seen lingering for a few seconds as the capacitors bleed down.)

- External Permit: This is the permit issued by the CPS to the blocks of 8 regulators, causing them to turn on sequentially (see under CPS reset).
- Overtemp: This bit looks specifically at the filter resistors and heat sink inside the regulator chassis. Temperature sensors are used here instead of klixons.

Ramped Sextupoles

All of the horizontal sextupoles are in series with each other; the vertical sextupoles form a similar but independent loop (Figs. 2-22, 2-23). Unlike the other correctors, the ramped sextupoles only require one power supply per loop. The large power supplies, I:SEXHPS and I:SEXVPS, are located in the MI-52 Service Building. Power to the sextupole supplies comes from Feeders 96 and 97, so power is removed when MOS 89 is opened.

There are numerous sextupole magnets for each plane, so connecting one terminal of the power supply to every other magnet, and catching the other half with the return path balances the inductive load. For example, the positive lead of I:SEXHPS is connected to S514, but the cable from S514 skips over S512 and is instead connected to S510. The loop leapfrogs nearly all the way around the ring before reaching S614, where it turns around and connects to all of the horizontal sextupole magnets that were missed the first time.

Digital control of the power supplies is through PLCs, which in turn are controlled through the VME crate in the electronics room of the service building.

The sextupole ramp waveforms, I:SEXH and I:SEXV, are sent to the power supplies from a CAMAC 453 card in Crate $5B. The waveforms are not only a function of MDAT30 (Main Injector momentum), but also of MDAT31 (Pdot, or the rate of change of momentum).

Main Injector

As you may recall from Chapter 2, ramped sextupoles are used to control the chromaticity of the ring. For tuning convenience, the chromaticity tables are normally adjusted from I2, but that should not imply that they are directly controlled by MECAR; they are not. There are separate chromaticity tables for each type of Main Injector reset.

Notes:

Chapter 4: Water in the Main Injector

> The study of water is such a dry subject that nothing funny can be said about it; there is no such thing as aqueous humor.

The magnets, power supplies, and RF systems in the Main Injector use tremendous amounts of power, and the heat generated from these components must somehow be removed. The Low Conductivity Water (LCW) systems provide cooling to the devices that need it. At each of the major service buildings, about 2.5 MW of power is dissipated as heat; a total of up to 8,000 gallons a minute passes through the components to keep them cool.

"Low conductivity" is essential because the water circulates directly through electrical conductors, such as the copper coils inside the magnets. The high voltages applied to the magnets would cause ions dissolved in the water to migrate and create currents of their own, disrupting the performance of the magnets and possibly causing damage.

The LCW system can be appreciated at several levels. It is important to understand how the water is distributed to the magnets in the tunnel, as well as the power supplies upstairs. Pumping, cooling, and deionization of the LCW all occur in the service buildings. The water required to cool the LCW is circulated in ponds distributed around the Main Injector ring. Finally, there are controls to coordinate the system.

This chapter is concerned primarily with water systems in the ring, as well as the supporting role played by the Central Utility Building (CUB). Detailed discussion of systems specific to the beam lines (e.g. the abort, P1, and A1 lines) is deferred to the chapter on Beam Transport Lines.

LCW in the Tunnel

Two stainless steel pipes, running the length of the tunnel between service buildings, distribute LCW to the magnets (Fig. 4-1, 4-2). One of the 6" diameter pipes is known as the supply header because it supplies cool water to the magnets; the other pipe, which collects the warm water emerging from the magnets, is the return header. The supply pressure, the pressure of the water in the supply header, is about 185 psig. The return pressure is about 15 to 20 psig. (Incidentally, psig stands for "pounds per square inch gauge." Pressure gauges are often calibrated to read zero pounds at atmospheric pressure, so a pressure reading in psig should be interpreted to represent pounds per square inch above atmospheric pressure.) The difference between the supply and return pressures forces the water through the magnets.

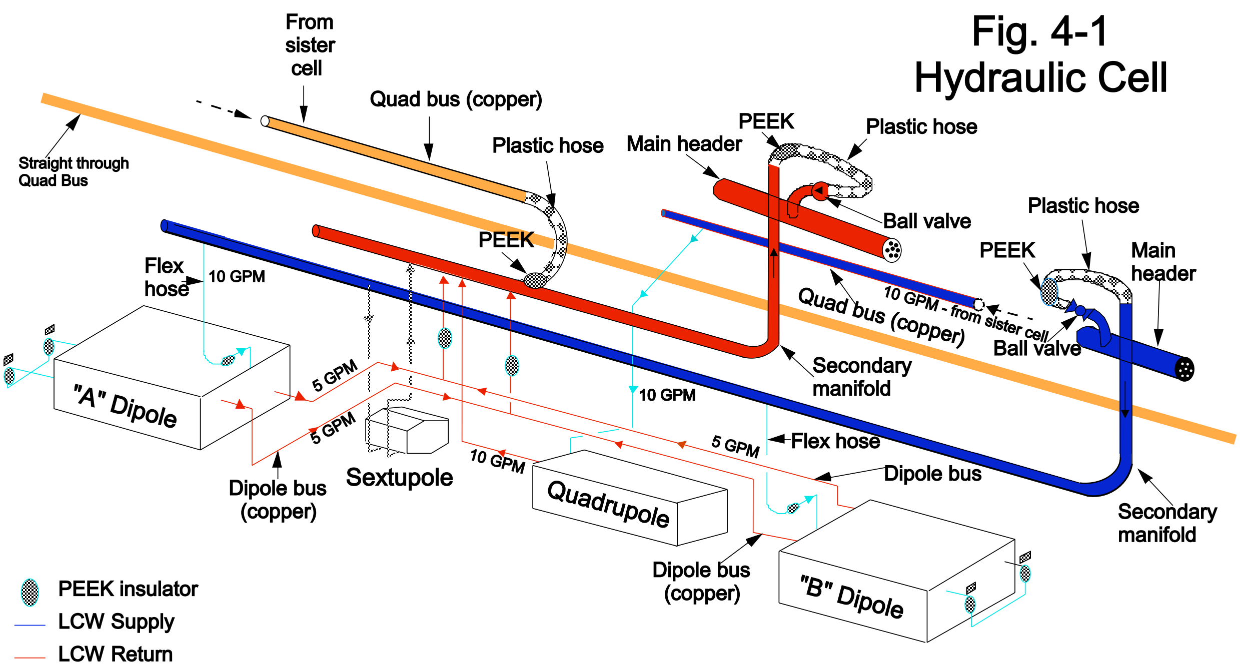

Secondary manifolds branch from the main LCW header at each main quadrupole location, supplying water to the main dipoles and main quadrupoles. The sextupoles are not water cooled at present. The ball valves on the main headers isolate the dipoles, but the quadrupoles must be isolated with valves on the quadrupole bus (not shown here). Compare to Fig. 4-2; the quadrupole shown here could be Q630. Her sisters would be Q628 (supply) and Q632 (return).

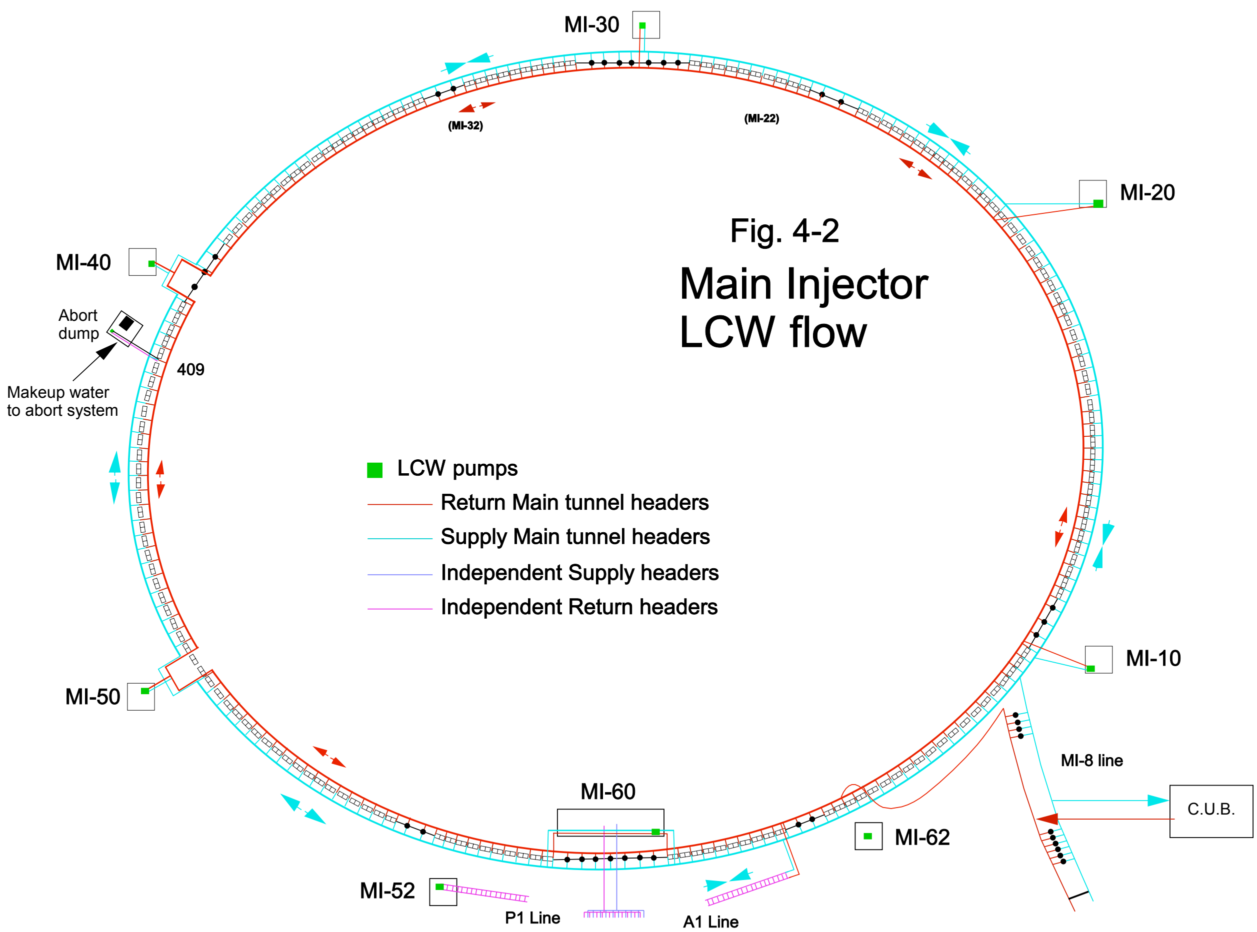

Because the LCW is going to pass through electrically active components, the stainless steel headers must be electrically isolated from the magnets. All of the hoses connecting the water pipes to the magnets are equipped with PEEK insulators to prevent the magnets from being grounded through the pipes. ("PEEK" stands for "polyether ether ketone." Unfortunately, it is also a completely unrelated trade name found on some of the LCW metering electronics.) PEEK is a relatively hard and inflexible substance used where sturdiness is required; plastic hoses of non-PEEK composition are also used where it is anticipated that thermal expansion and contraction of the pipes could cause problems. There are specialized cases where ceramic insulators are used as well. All of the insulators allow continuity of water flow, but block electrical current. The water itself will not conduct significant amounts of current because most of the ions have been removed.

The magnets in the tunnel that dissipate enough heat to require water-cooling are the main dipoles and main quadrupoles. The other, smaller magnets are air-cooled, although provisions have been made to supply water to the sextupoles should it become necessary. The cooling units are organized into hydraulic cells, which consist of a main quadrupole and the two main dipoles to either side. The boundaries of the hydraulic cells do not correspond to those of the cells of the FODO lattice.

At each hydraulic cell, LCW is distributed to the individual magnets through secondary manifolds. There are supply and return manifolds at each cell. Ball valves are present which can be used to isolate the main headers from their respective secondary manifolds in the event of a leak or scheduled maintenance.

A line carrying 10 gallons per minute (GPM) runs to each of the two main dipoles from the supply manifold; PEEK insulators electrically isolate the magnets from the manifold. The water enters the magnets through a connection at the top, on the side facing the main quadrupole. Inside the magnet, the flow splits into two branches. Water passes through the holes in the copper coils (Fig. 2-2); in order to cool the upper and lower coils as well as the straight-through bus, the flow turns around at the opposite end of the magnets and makes a second pass through. The jumpers are provided with PEEK insulators so that the magnet is not electrically shorted.

The 1” X 4” main dipole bus which runs behind the quadrupole must also be cooled, so water returning from the magnets passes through the bus on its way back to the return manifold.

Water to the quadrupoles is handled differently for historical reasons. Many of the quadrupoles have been recycled from the Main Ring, where the LCW was carried inside the copper bus. For consistency, the new quadrupoles have been built in the same manner. The quadrupole bus is a hollow copper pipe; the copper carries the power and the LCW flows inside the pipe. At the secondary manifold, the quad bus is connected to either the supply or return manifold through a plastic hose. Consider the focusing bus (Fig. 4-3). At a given location, the bus may be connected to, say, the supply manifold. Since this particular length of bus is connected to two focusing quadrupoles, the flow is divided between the two magnets. For example, the supply header at Q628 also supplies Q630. The quad bus at the next focusing location (in either direction) is connected to the return manifold. For example, Q630 is connected to the return manifold, which also collects warm water from Q632. Two hydraulic cells tied together via the quadrupole bus are called "sister cells."

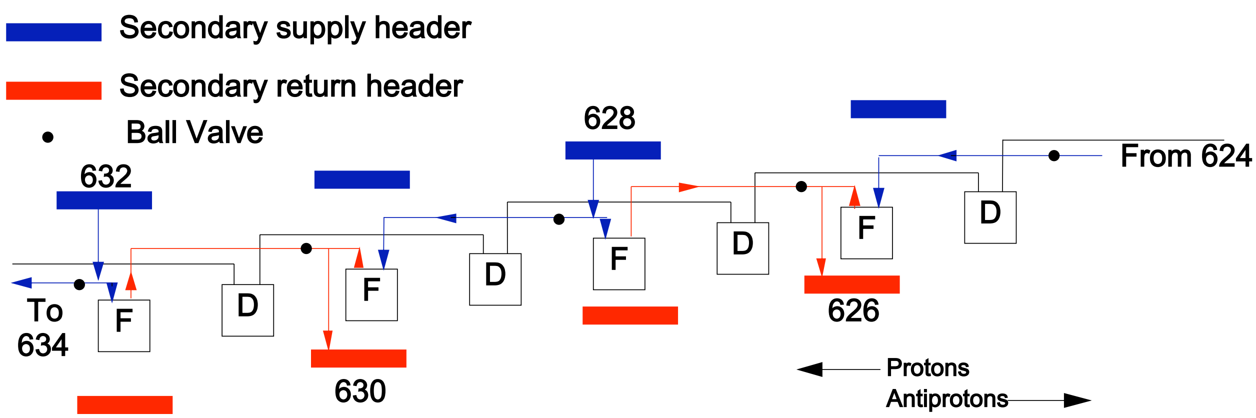

At each main quadrupole location, the magnet is alternately connected to either the supply header or the return header. Since every magnet needs access to both a supply and a return header, each secondary header must be connected to two quadrupoles (compare to Fig. 4-1). Also, since the water travels through the buss work, the two " sister" magnets must have the same polarity. Ball valves are used to isolate the quadrupoles from the main header.

This picture, for the sake of clarity, only shows connections to focusing quads.

Fig. 4-3, Sister Cells

The defocusing bus follows the same pattern, but is one half-cell "out of phase" with the focusing bus.

A careful look at Figs. 4-1 and 4-3 reveals that closing the ball valves to the secondary headers is insufficient to isolate the quadrupole magnets, because LCW can still flow in from a sister cell. If a leak develops at a quad, there are small (3/4") ball valves upstream and downstream, on the quadrupole bus itself, which must also be closed. The upstream valve is almost a half-cell away, while the downstream valve is nearly a full cell away.

The valves are adjacent to breaks in the bus that have been jumpered with braided copper cabling to allow for thermal expansion of the pipe.

Since the 6" supply and return headers are present around the entire circumference of the ring, and since LCW passes from the supply header through the magnets to the return header at several hundred locations, the ring-wide view can be thought of as a large capillary network with the components being cooled in parallel. In Fig. 4-2, each "crossover" in the ring from the supply to the return header represents a secondary manifold. At the service buildings, warm return water from the magnets is drawn upstairs, where it is cooled, deionized and pressurized before being sent back into the tunnel as supply water.

Notice that water is sent out from the service buildings in both the upstream and downstream directions, and returns from both directions. At some point along the route, as water is drawn from the supply header, the pressure from adjacent service buildings will be equal. To ensure that the water will not completely stagnate, there are flow restrictors on the headers for adjusting the differential pressures. They are located in the service buildings, just before the headers enter the tunnel, and are set up to induce a slow net flow in a clockwise direction.

Occasionally, the headers branch out to cover specialized functions. LCW supply water is shunted to CUB from a point between 640 and 641, eventually coming back to the return line at 624. A short branch at MI-40 supplies the Abort Lambertsons. At 409, water is siphoned off for the quasi-independent LCW system of the abort dump. The abort water heat exchanges with the ring LCW, and makeup water is supplied from the main headers, but the abort LCW is a closed-loop system with its own deionizer. Between 618 and 619 there is a branch from the main header that supplies water to devices in the A1 (or A150) line.

In addition, there are three LCW systems completely independent from the rest. At MI-52, there is a closed-loop system that supplies LCW to the devices in the P1 (or P150) line. Water is pumped from the MI-52 Service

Building; the headers run along the inside wall of the tunnel before crossing over to the magnets at 523.

More detail on the abort, P1, and A1 LCW systems will be included in the "Beam Transport Lines" chapter.

The other two independent systems are at MI-60. (This would be a good time to begin pondering Fig. 4-7, but don't sweat the details yet.) One system known as the "RF" or "95°" system, supplies water to the RF electronics upstairs, the RF amplifiers downstairs, and to the anode supplies. The other, known as the "Cavity" (or quite inaccurately, the "55°" system) provides water to the RF cavities downstairs. The RF and Cavity LCW systems will be dealt with in a separate section.

LCW in the Service Buildings

Each service building houses equipment dedicated to maintaining the LCW under its regional sphere of influence (Fig. 4-4). The layout is similar from house to house (except at MI-60, which will be discussed later).

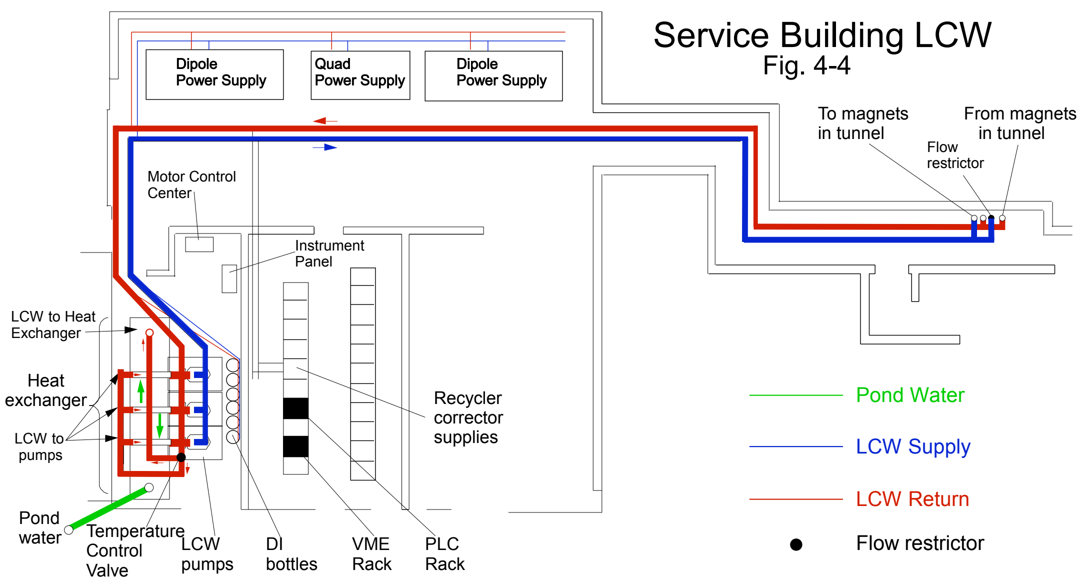

At the service buildings, cool LCW is pumped to the magnets in the tunnel and the power supplies upstairs. The main headers are suspended from the ceiling.

Only the external plumbing of the heat exchanger is shown here; this diagram should be compared to the cross-section in the text. The warm LCW returns to the Temperature Control Valve (TCV), to the right of the exchanger in this diagram. (The header is not connected to any of the structures shown beneath it.) The valve splits the LCW into two branches. The branch in the center is heat exchanged with pond water from outside (the pond water is represented somewhat schematically). The other branch bypasses the exchanger. The two flows are then combined (not shown). Finally, the combined flow is divided into three branches, which pass under the exchanger to the LCW pumps.

The hardware includes LCW pumps, a heat exchanger, and deionization tanks.

In addition to the figures in this book, the graphics on page I56 will be helpful in understanding the details discussed for the remainder of this chapter.

LCW Pumps

The pumps are the heart of the LCW system, providing the differential pressure necessary to keep the water circulating in the ring and through the magnets. Each of the 100 horsepower pumps is capable of receiving LCW at a pressure of 15 or 20 psig from the return header and pressurizing it to about 185 psig as it enters the supply header. All three pumps are normally running at any given time; about a thousand gallons of water is pumped to the magnets in the tunnel from each service building every minute.

Low-pressure water in the return line before it enters the pump is called suction, and the high-pressure water leaving the pump is referred to as discharge.

The motors for the LCW pumps are powered and controlled from panels in the Motor Control Center (MCC). The pumps each have a dedicated panel on the MCC from which they can be turned on and off. For remote control of the pumps, the MCC is the interface with the controls system.

The supply and return headers work their way out of the pump room, past the main power supplies; the penetrations into the tunnel are located next to the door leading into the tunnel. Notice in Fig. 4-4 that the supply and return headers each split into two branches just before entering the tunnel; one runs upstream of the service building and the other downstream (compare to Fig. 4-3). The flow restrictors are located on one of the branches of the supply header.

Cooling for the main dipole and main quadrupole power supplies is tapped from the main headers that branch of the LCW runs behind the power supply cabinets. There are manually operated supply and return valves for isolating each of the three supplies. In addition, a small branch taps off the main headers to feed the Recycler corrector supplies in the electronics room.

Heat Exchangers

The return flow to the pumps takes a somewhat convoluted path as it approaches the heat exchanger. The heat exchangers, one in each building, cool the water returning from the magnets. Some of the warm LCW passes through a pipe that runs the length of the exchanger, while cool water from outdoor ponds flows past the pipe. Although the two bodies of water do not actually meet, heat is transferred from the warm LCW to the cooler pond water, as in the diagram on the next page.

Before encountering the heat exchanger, water returning from the tunnel and power supplies enters the Temperature Control Valve (TCV), also known as the 3-way valve. The TCV splits the flow between a branch that passes through the heat exchanger, and a branch that does not. The ratio of cooled to warm water determines the temperature of the water that will be sent to the magnets. The valve is given a setting temperature, usually around 90˚ F, and continuously adjusts the ratio in order to maintain that temperature. The readback is set up so that "% open" refers to the percentage of the return flow sent to the heat exchanger.

In Fig. 4-4, in which only the external plumbing of the heat exchanger is represented, the return flow can be seen splitting into the two branches. The branch seen directly over the heat exchanger enters the exchanger at the far end and flows back through the inner pipe. The other branch bypasses the exchanger altogether. The two flows merge at the point indicated (the inner pipe leaving the heat exchanger is hidden in this view). Finally, the combined flow splits one more time to feed all three of the LCW pumps.

The cool pond water from the pond pumps comes up through the floor, passes through a large filter, and enters the heat exchanger at the top. It flows the length of the exchanger, first passing over the LCW pipe and then, on the return path, underneath. It leaves the bottom of the heat exchanger and returns underground to the outside sprayers. A more global view of the pond pumps is forthcoming.

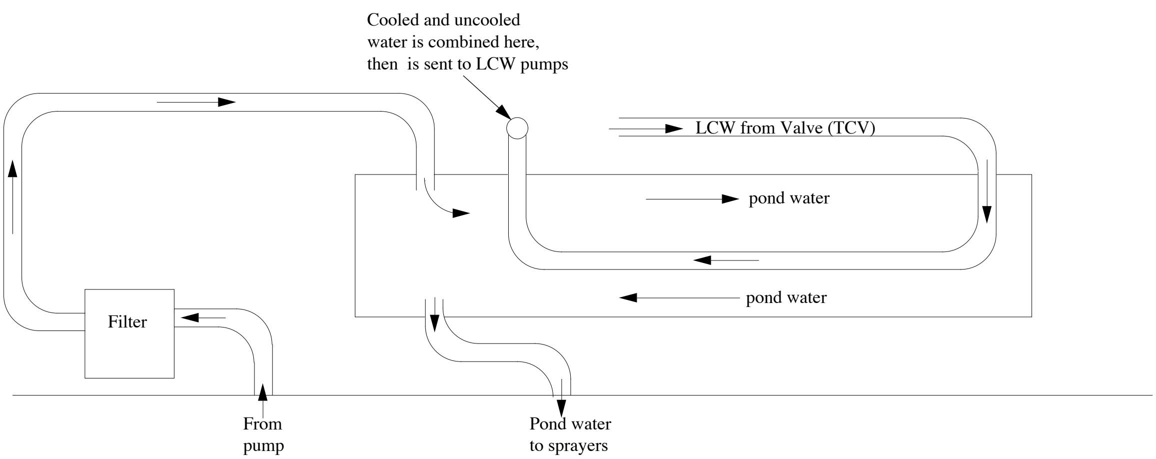

Cross-Section of Heat Exchnager

Deionization

A small portion of the discharge from the LCW pumps is sent through a series of deionizer (DI) bottles. The bottles are basically just fancy (and very thorough) water softeners that remove dissolved ions from the water. The DI bottles are filled with tiny resin particles that have been treated to retain an electric charge. These “mixed bed” deionizers are comprised of negatively and positively charged resins. Ions in the water, such as calcium, magnesium, sodium, sulfate, bicarbonate, etc., are captured by the particles. When the resins become saturated with ions, they must be flushed with caustic substances in order to recharge them.

Non-ionic nuisances, such as copper oxides, are immune to capture by the resins.

The optimal flow through recharged deionizer bottles is about 40 GPM, but that rate slows down as the resins become saturated. The bottles need to maintain a flow of at least 5 GPM in order to function properly.

To protect the deionizers, there is a regulator upstream of the bottles which steps the pressure down from the 185 psig of the supply line to about 50 psig.

The conductivity of the water is usually measured in units of its antithesis, resistivity. Resistivity is expressed in ohms per centimeter. The greater the resistivity, the cleaner the water. LCW in the Main Injector should be in the neighborhood of at least 8 Megohms (millions of ohms) per cm, and can be as high as 18 Megohms per cm.

As a backup, there are filters at the beginning and end of the series of bottles that clean the water and offer some protection to the rest of the system in the event that a bottle goes bad.

Ponds and Pond Pumps

The cooling ponds ultimately dissipate heat from the magnets and power supplies. These elongated ponds are adjacent to the berm, sometimes to the outside of the ring and sometimes to the inside (Fig. 4-5). Designated A through H, they have a total of about 19 acres for evaporative cooling. It turns out that the rate of evaporation is the most important factor limiting the ramp rate of the ring.

The pond water is not low conductivity, it would be rather difficult to maintain ion-free water in an open outdoor environment.

There are two pond pumps at the downstream end of each pond that pump water out of the pond and into the service buildings to be heat exchanged with the LCW. Usually, only one pump is running at a time at any given service building. After absorbing heat from the LCW, pond water is usually pumped to the upstream side of the adjacent pond through sprayers. Water in a spray will cool faster because of the increased surface area of the droplets. (The spraying nozzles are removed in the winter to prevent the roads from icing over on windy days.) The pond water before it is heat exchanged is defined as supply water; after the heat exchanger it is referred to as return water.

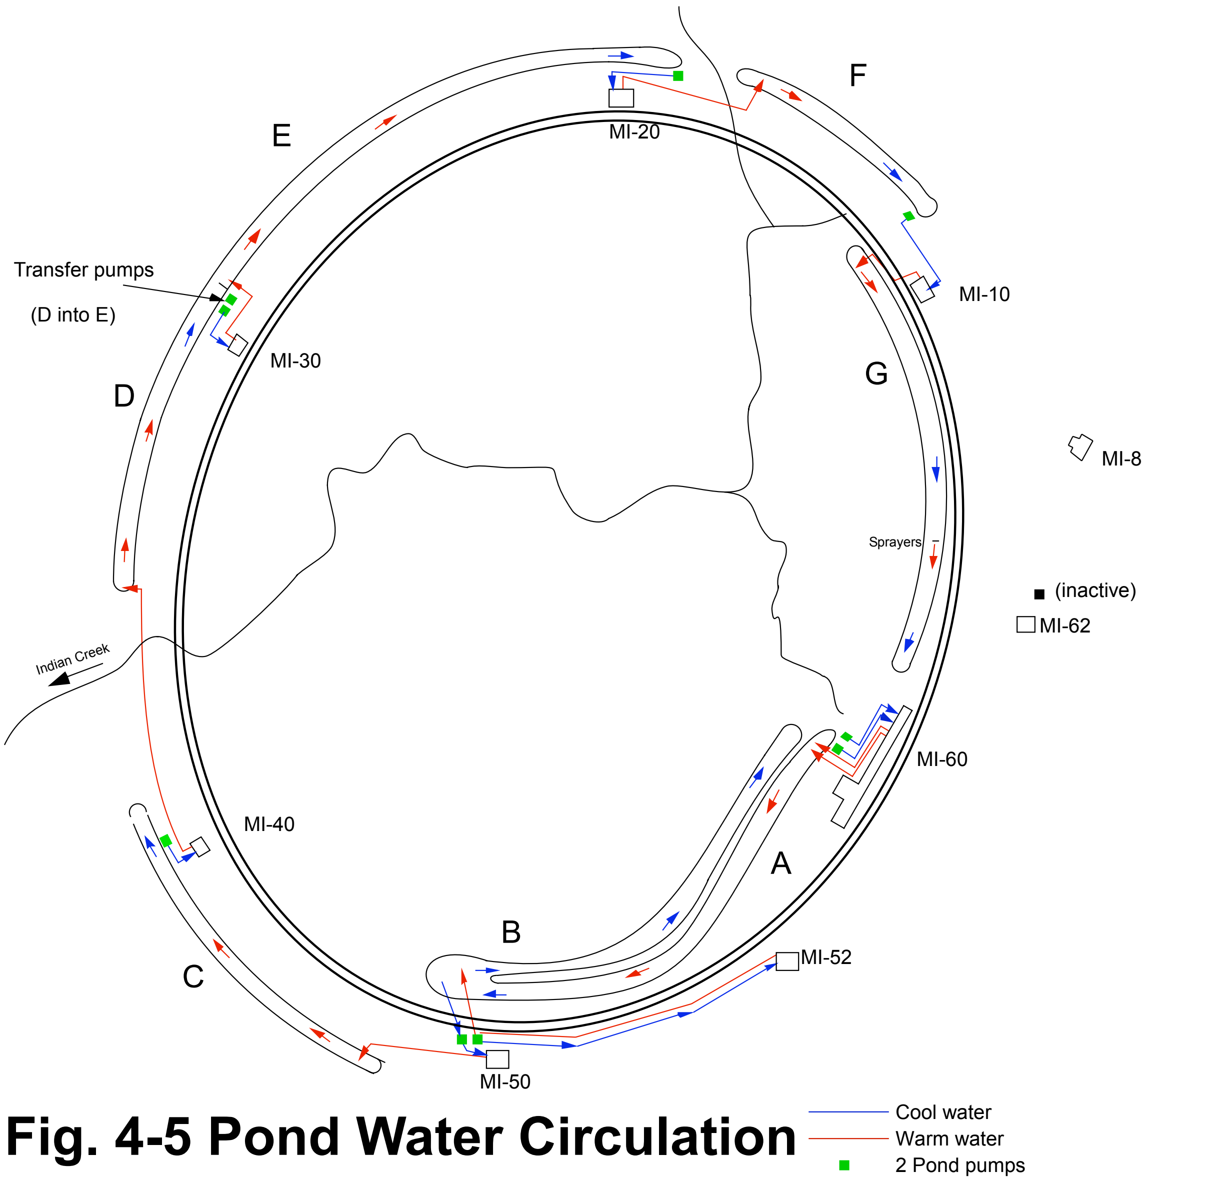

Fig. 4-5 Pond Water Circulation

Cool water is drawn from the ponds and used in the heat exchangers to cool the LCW returning from the magnets. The warm water from the heat exchangers is returned to the ponds via sprayers, and allowed to cool by evaporation. Each green square represents a pair of pumps.

The inactive pumps near MI-62 will probably be used when the A1 and NuMI lines are commissioned.

The location of the pumps is designed so that there is (in general) a constant clockwise flow around the ring. After warm water is discharged at the upstream end of a pond, it slowly drifts downstream and cools by evaporation. At the downstream end of the pond, it is cool enough to be used as supply water for the next service building.

Sometimes there is a significant distance between ponds, as between C and D (so that Indian Creek can run between them), and sometimes only a dam separates the ponds. The A and B ponds are continuous, as are G and H, but are given distinct names because there are pumping stations midway through.

The MI-52 LCW system, which is independent of the ring system, sends its warm pond water to the point where the "A" pond turns around, behind the MI-50 Service Building. It draws cool water from a point upstream of the discharge.

In addition to the usual pumps, there is a "lift pump" at MI-30. Rather than being tied into a heat exchanger, its sole purpose is to transfer water from one pond to another in order to maintain the proper levels.

The pond pumps are housed in concrete pits but can be susceptible to flooding. Sump pumps in the pits prevent the pond pumps from being submerged. When the pit fills to a certain level, the sump pump discharges the water back into the pond.

Local LCW Controls

The local LCW controls at a service building must be able to read the various pressure, temperature, flow, and conductivity values, as well as pump and valve status. They must also be able to execute commands such as turning the pumps on or off, and regulating water temperature through the TCV. Finally, all of this information must be shuttled back and forth between the service building and the Main Control Room (see Fig. 4-6 on the next page).

Most of the instrumentation reports to modules in the Instrument Panel, located in the pump room. The equipment in the Instrument Panel is more of a collection than an integrated system, and the individual units are usually named after the manufacturer. Conductivity, measured at the DI bottles, reports back to the "Thornton" meters; each of the Thornton modules can handle two conductivity sensors. A PEEK meter (brand name only, not chemical composition) measures flow through the DI bottles, while the other flow meters report to the "Ronan X54."

Fig. 4-6
Service Building LCW Controls

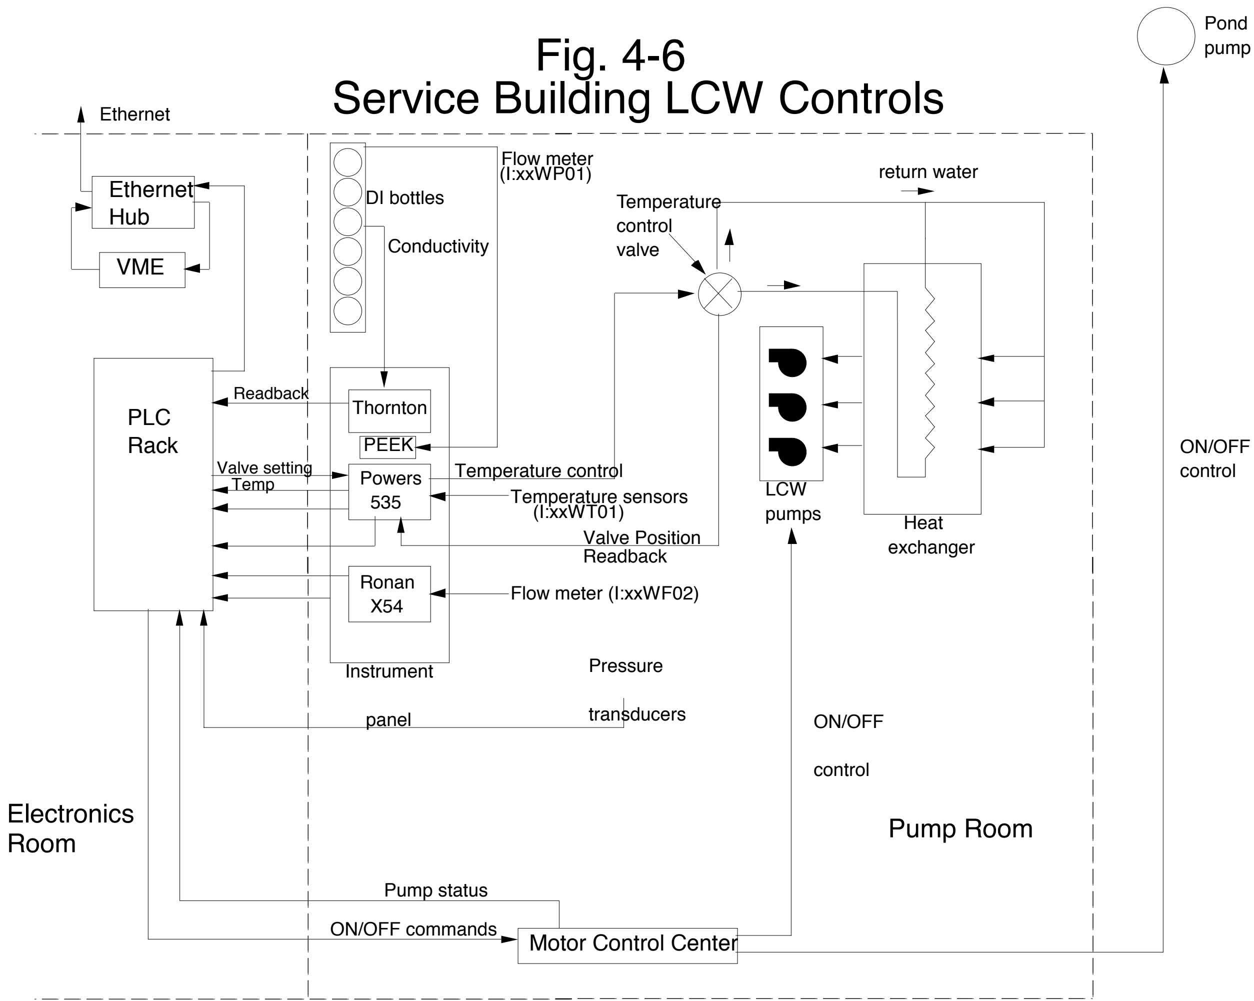

The most sophisticated of the modules in the Instrument Panel is the "Powers 535," which handles temperature control. Not only does it read the temperature sensors and the TCV position, it also runs the algorithm that maintains the temperature of the LCW at the desired level by adjusting the TCV.

Overseeing the Instrument Panel and other controls functions are the PLCs, or Programmable Logic Controllers. These are Six-trak PLCs, a different type than those used for the power supplies. They are housed in racks in the Electronics Room. Six-trak PLCs, unlike those for the power supplies, are capable of reading analog signals. The PLCs read the analog readbacks for conductivity, temperature, and flow rates from the Instrument Panel; the pressure transducers and pump status from the MCC are read directly by the PLCs. In the other direction, digital control of the pond and LCW pumps is routed through the PLCs.

A VME crate, the same one that consolidates data from the power supply PLCs at the service building, is also the interface between the Six-trak PLCs and the rest of the world. Since this is the Age of Networking, the PLCs and the VME, which are three feet away from each other, must talk to each other through the local Ethernet hub. The VME, after organizing the information, starts the data on its journey toward the Main Control Room via the ring-wide Ethernet. For the rest of the story, read the hypothetical Controls chapter.

Water at MI-60

The LCW systems at MI-60 (Fig. 4-7) are more complicated than those at the service buildings. They not only have to supply water to the magnets, but to the RF systems as well.

Fig. 4-7 MI-60 LCW

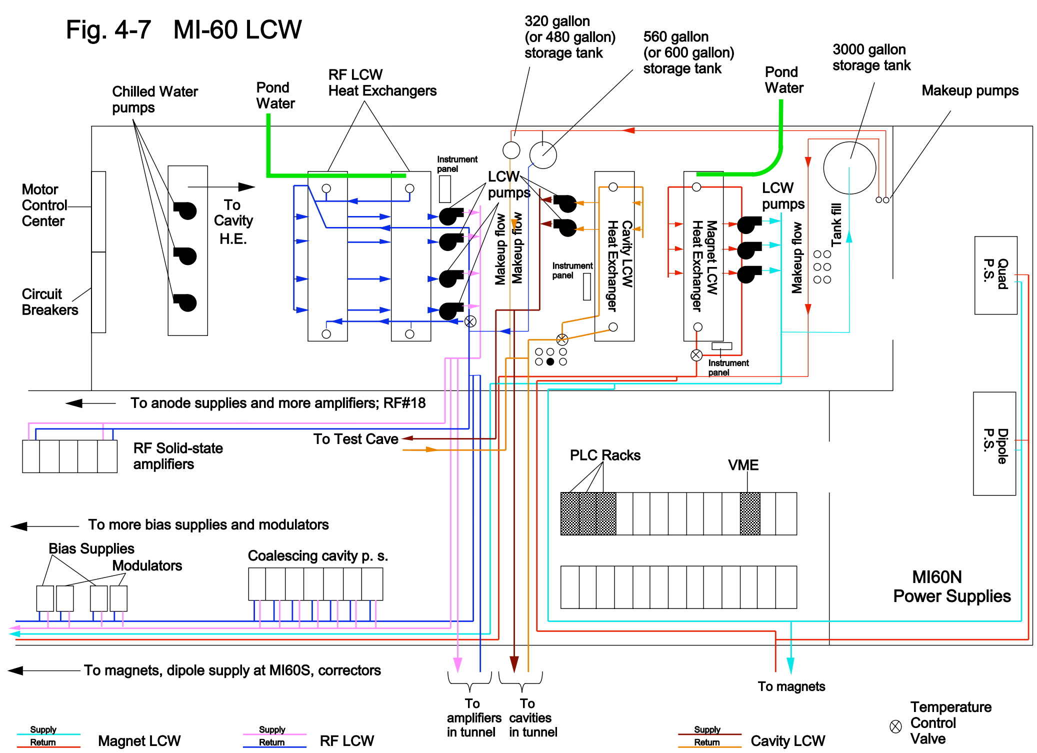

This map is believed to be topologically correct, although numerous details have been omitted in a futile attempt to achieve clarity.

All of the systems are gathered together in the Pump Room at the north end of the building. There are actually four distinct water systems in the Pump Room, going from south to north (left to right in Fig. 4-7). They are the Industrial Chilled Water system, the RF system, the Cavity system, and the Magnet system.

Industrial Chilled Water

The Industrial Chilled Water (ICW) is not actually a low conductivity system. It supplies water to the fire sprinkler heads and air conditioners throughout the building. The ICW pumps and pipes are painted red everywhere in the building for quick identification during an emergency. Its importance to the LCW is that it is heat exchanged with the LCW in the Cavity system (see below).

RF System

“RF” here is understood to mean the RF power supplies rather than the cavities themselves. The RF water is divided into three major branches. One branch provides water to the coalescing cavity supplies as well as to the bias supplies and modulators for 17 of the 18 accelerating stations. The second branch, running along the east wall of the building, goes to the solid-state MOSFET amplifiers, the anode supplies, and the bias supply and modulator for the remaining accelerating station (#18). The third branch plunges into the tunnel in order to cool the power amplifiers on top of the cavities.

The RF LCW, which is responsible for cooling a large heat load, requires two heat exchangers and four LCW pumps. Heat is exchanged with pond water as with the other service buildings.

Being a stand-alone system, the RF LCW has its very own 560-gallon storage tank for replacing lost water.

Cavity System

Since the RF cavities represent a relatively small heat load, only one heat exchanger and two LCW pumps are needed. As mentioned above, the Cavity LCW is heat exchanged with the ICW rather than with pond water. The Cavity system, also stand-alone, is provided with a 320-gallon storage tank. In addition to cooling the cavities in the tunnel, there is a branch that feeds the Test Cave area upstairs.

Notice on the diagram that the 320-gallon tank is physically located nearest the RF system, and the 560-gallon tank is located nearest the Cavity system. Both of these storage tanks get their makeup flow from the 3000-gallon tank of the magnet system, to be described next.

Magnet System

The magnet LCW is similar in most respects to the LCW setup at the other service buildings with one heat exchanger and three LCW pumps. Inside the Pump Room it splits into a north and a south branch. Upon reaching the east wall, each of the two branches splits again, with one secondary branch entering the tunnel to service the magnets and the other going to the main power supplies.

In the tunnel, the branches join the main LCW headers at the upstream and downstream ends of the long straight section. A relatively small crossover supplies the main quadrupoles that are scattered among the cavities. Remember that there are no main dipoles in the straight section, so the heat load is minimal.

What is different about the magnet system at MI-60 is the 3000-gallon storage tank (not to be confused with the 3000-gallon tank at CUB, which has a similar name). Since most of the magnets in the ring are connected to a common header, this tank is used to make up water for the entire ring. The water is added to the suction side because of its lower pressure. The tank, in turn, is filled with LCW from CUB when it runs low. In addition to the magnet system, the 3000-gallon tank at MI-60 is used to fill the storage tanks for the RF and Cavity systems, and the (otherwise) independent system at MI-52.

Nitrogen gas, obtained from bottles, is used to control the head pressure of the 3000-gallon tank. That is, there is nitrogen at the top of the tank exerting a pressure that is transmitted throughout the system; since makeup flow is added to the suction side, the regulated pressure of the nitrogen controls the suction pressure of the water. It is important to keep the suction pressure comfortably above atmospheric pressure in order to keep the pumps from cavitating. Cavitation occurs when the pressure of the system drops below the vapor pressure of the water, and the water begins to boil. The dramatic increase in volume and pressure and the highly energetic “bubbles” bouncing around inside the pipe can cause a great deal of damage. There is a line from the tank to the return header that must be open to maintain continuity of pressure whenever the pumps are running.

There are two small makeup pumps associated with the MI-60 system. They are turned on only when makeup flow is needed. One of the pumps sends makeup water to the Magnet system. The LCW is first pumped through the DI bottles; the flow can then go either to the return line or back to the makeup tank. The other pump sends makeup water to the RF and Cavity systems.

Makeup water is configured more or less the same way for the RF and Cavity systems, and at MI-52.

An advantage of using nitrogen, as opposed to ordinary air, is that oxygen is excluded. The presence of dissolved oxygen in the water would lead to the creation of copper oxides in the magnets and busywork. In other accelerators, copper oxides have had a shameful history of clogging up the plumbing and causing much downtime due to overheated magnets.

Each of the three LCW systems at MI-60 is supported by a technical infrastructure similar to that found at the other the service buildings: Instrument panels, Temperature Control Valves, DI bottles, etc. Three PLC racks and the VME crate are found in the gallery to the east.

CUB

So far in this chapter, the Main Injector LCW systems have been treated as “stand-alone;” that is, as if all of the pumping, cooling, and deionization takes place locally at the service buildings. It is indeed possible to operate all of the Main Injector water systems without outside help, except for occasional makeup flow. However, CUB plays a major role in deionizing the LCW, and is the primary source of makeup water to the ring.

Fig. 4-8 represents the layout of the Central Utility Building. CUB supplies water to much of the accelerator complex; those functions having to do with the Main Injector and the Tevatron are shown in color. The two

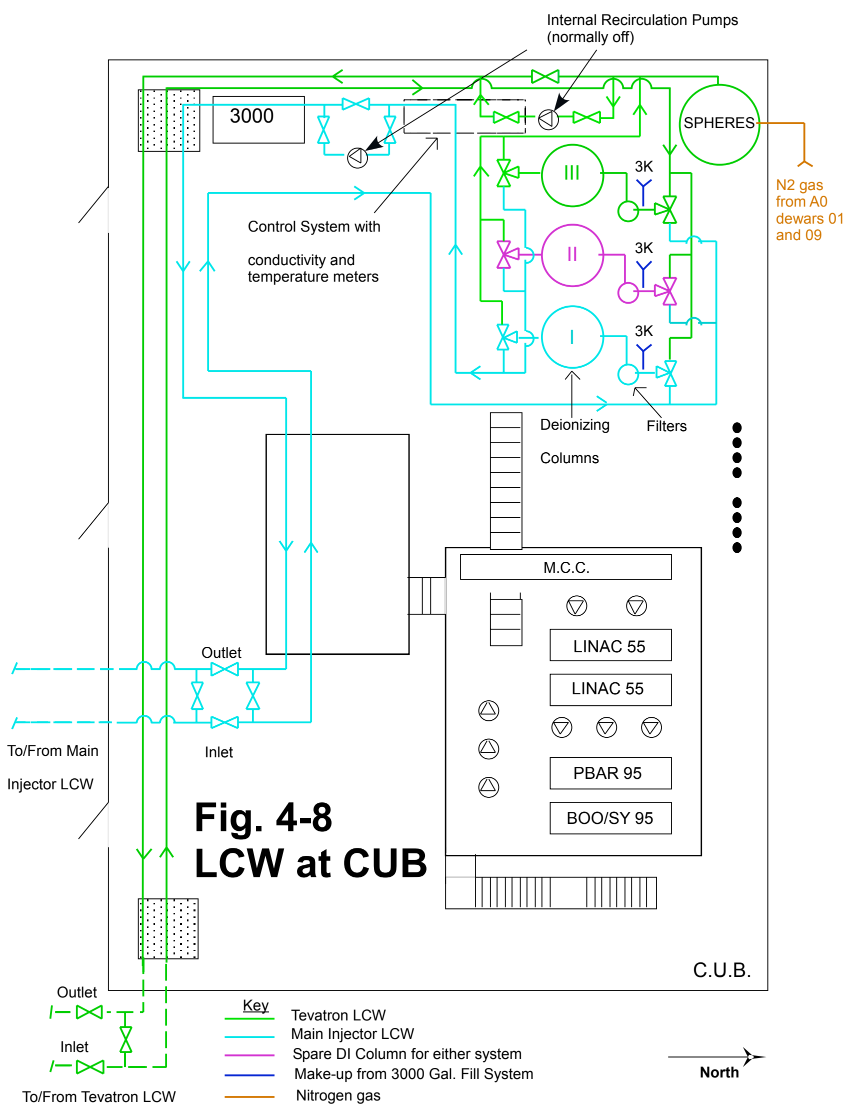

Fig. 4-8
LCW at CUB

systems are interrelated and will be considered as a single entity here.

The heart of the system (or perhaps more appropriately, the kidneys) consists of the three large deionizing tanks (or columns) in the northwest corner, labeled I, II, and III. Normally, the Main Injector uses Column I and the Tevatron uses Column III, while Column II can be used as a spare for either. If necessary, the columns can be reconfigured so that any column can cleanse the water from either machine.

The primary reservoir of LCW at CUB for the Main Injector and the Tevatron is contained in the spheres. Certainly among the most unusual water containers on the planet, the three large spheres are stacked vertically in the northwest corner of CUB. Like the magnet storage tank at MI-60, they can hold about 3000 gallons. Also, like the 3000-gallon tank, the head pressure at the top of the spheres is maintained by nitrogen gas at the top of the water. Boiling off some of the liquid nitrogen from the Dewars at A0 derives the nitrogen.

The spheres are the primary source of makeup flow for the Main Injector as well as the Tevatron. Water can also be drawn directly from the 3000 gallon tank at CUB, not to be confused with the other two 3000 gallon systems already discussed.

The supply and return lines to the Main Injector exit CUB from a point near the center of the south wall. Traveling underground, they skirt the Antiproton Source and penetrate the wall of the MI-8 enclosure at location 809 (look again at Fig. 4-3). There is a branch from 837 leading up to the MI-8 building, used for cooling the bulk supplies for the correction elements in the MI-8 line. A small branch from 809 feeds the powered magnets downstream of the shielding wall, but the magnets of the MI-8 line upstream of the wall are cooled separately, from the Booster LCW header originating at Long 4. The powered quads at the end of the MI-8 line are supplied from the same branch that services the quads downstream of the shielding wall.

There are no pumps to force the water back to CUB, so the supply water, which is already at high pressure) is tapped from the header at a point between

640 and 641. However, rather than entering the return line at 640, the water coming back from CUB is diverted to 624. The reason for this is to try to put enough separation between the point of entry and the point of departure. Otherwise, the water might end up looping exclusively between 640 and CUB. The LCW pumps in the service buildings, the flow restrictors, and the makeup flow maintain the overall clockwise flow in the ring at 624.

Entry and exit to CUB is controlled by the inlet and outlet valves, which are normally open. Once inside, the water passes through Column I, and, refreshed and not under so much pressure, starts back to the Main Injector ring.

Sometimes, especially when the columns need to be regenerated, closing the inlet and outlet valves and opening the bypass valve can set up a pattern of internal circulation. The setup is similar for both the Main Injector and the Tevatron. A small pump, used only for internal circulation, can be used to maintain flow through the columns.

If water needs to be made up and CUB is down, domestic water can be used. Domestic water is the "ordinary" water used site wide for bathrooms and drinking fountains, and of course is of insufficient purity to use in the magnets. The water, drawn from the domestic line at MI-60, is heavily filtered before being allowed to join the LCW; one of the "filters" is UV light, intended to destroy bacteria that could potentially eat through the stainless steel pipes. Domestic water is also drawn at CUB during normal operations, and processed in a similar way.

Parameter Names

The LCW parameters can be found on I58. Of course, all Main Injector parameter names are of the type I:xxxxxx. On I58, each "character place" is more or less standardized to represent some particular facet of the LCW.

Going from left to right:

- The first two characters give the location of the device: “10” for MI-10, etc. At MI-60, “60” is assigned to the magnet system, “6R” to the RF system, and “6C” to the Cavity system.
- The third character in an LCW parameter name is usually “W,” for “water”.
- The fourth character describes the type of device; for example, “L” means “liquid level,” e.g. I:60WL01 refers to the water level in the 3000 gallon storage tank for the magnet system at MI-60, and I:6RWL01 refers to the liquid level in the 560 gallon storage tank for the RF system. “V” stands for valve, “P” for pressure, “F” for flow, “T” for temperature, and “R” for resistivity. For digital control parameters of the pumps, such as I:60WS01, the “S” stands for “start,” or, I suppose, “stop.”
- The last two digits in an LCW parameter name are numeric, specifying which device of a particular type is meant. A serious attempt has been made to ensure that diagrams, parameter names, and labels on the equipment all correlate with one another.

By themselves, the parameter names and descriptor text often do not reveal much about the specific functions of the devices that they represent. However, the names and numbers are standardized from building to building (except for MI-60), and most functionality can be determined from the graphics called up from I56. At MI-60, the three systems are more complicated than the rest—primarily because of the storage tanks and associated makeup flow—and the conventions used at the other service buildings are largely abandoned.

Specific parameter naming schemes for the major service buildings MI-10 through MI-50 are described below, along with some of the ways that parameters can be interpreted and used; MI-60 parameters are unique enough to require a separate discussion. Since the first two characters in an LCW parameter usually refer to the service building or RF system, and since the third character is usually "W," names in the following section will frequently be abbreviated to the last three characters.

Pumps

"S", as mentioned earlier, designates pumps. I:xxWS01, S02, and S03 display the digital status of the three LCW pumps at most service buildings; all three pumps are normally on. S04 and S05 represent the pond pumps; one of the pumps is usually on. A sixth pump parameter, I:xxWPND, is the "Pond Pump Digital Status."

The pumps can also be controlled through these parameters, although other applications (such as I56) are used more frequently.

In the MI-60 Magnet system, the LCW pumps are still named I:60WS01, S02, and S03, but S04 and S05 refer to the two small makeup pumps next to the 3000 gallon storage tank. The pond pumps are designated S06 and S07.

The RF system at MI-60, because of the large heat load, requires two heat exchangers and four LCW pumps. The pumps are designated I:6RWS01, S02, S03, and S04. There is one pump for handling makeup flow, designated S05, and two pond pumps named S06 and S07.

In the Cavity system, I:6CWS01 is the makeup pump. S02 and S03 are the LCW pumps. There are no pond pumps as such, since the LCW heat exchanges with the ICW.

Valves

In non-MI-60 buildings, V01 and V03 represent the valves on the supply headers as they enter the tunnel; V02 and V04 are the valves on the return headers. These four parameters are digital readbacks, showing whether the valves are open or closed.

I:xxWW01 is the position of the TCV, measured in "% open." The percentage refers to the amount of return water sent through the heat exchanger.

In the MI-60 Magnet system, V01 and V02 refer to valves upstream of the makeup pumps S04 and S05, respectively. V03 is a valve on the makeup line that adds water to the return line, near the TCV. V04 is on the line that connects the supply header to the DI bottles. The TCV is designated I:60WW03.

In the RF system, there are six valves related to handling the flow of makeup water. The 560 gallon tank supplies much of the makeup water to the system, but that tank, in turn, gets its water from the 3000 gallon tank of the MI-60 Magnet system.

V01 opens to let water from the 560-gallon tank flow to the DI bottles.

V02 is the gate (water-gate) that allows water to come in from the 3000-gallon flow line. If V03 is open, the makeup water flows into the 560-gallon storage tank, and if V04 is open, the water goes directly to the DI bottles.

V05 and V06 are valves in the DI loop. V05 is in the supply line and V06 is in the return line.

In the Cavity system, I:6CWV01 is on the makeup line from CUB; V04 lets the makeup water into the storage tank. V02 isolates the storage tank from the makeup pump, and V03 lets the water from the 3000-gallon tank enter the Cavity system on the supply side of the DI loop. V05 is on the supply side of the DI circuit, and V06 is on the return line of the circuit. Since V06 can isolate the storage tank from the return line, it should normally be left open, otherwise the return header will not feel the head pressure from the storage tank.

Flow Readbacks

Flow readbacks are given in gallons per minute (GPM).

I:xxWF01, “DI processing,” measures the amount of flow going through the DI bottles. This is the value read at the PEEK meter on the Instrument Panel.

I:xxWF02, “Gallery Supply,” is the total amount of flow going to the magnet strings in the tunnel. The PLCs use F02 to monitor the flow and shut off the LCW pumps if it exceeds certain limits. If the upper limit is exceeded it is likely that there is a major leak downstairs, and one might as well not send lots of pressurized water in that direction. The upper and lower limits are set through the parameters I:xxWF2U and I:xxWF2L. The flow meter hardware can be found hanging from the supply header at the entrance to the corridor leading to the tunnel.

In the MI-60 Magnet system, F01 measures flow into the 3000 gallon storage tank; F02 is the flow out of the DI bottles; F03 and F04 measure flow to the two branches of the supply headers as they go into the tunnel. Trip limits (I:60WF3L, etc.) can be set for F03 and F04.

In the RF system, I:6RWF01 measures the flow coming into the RF system from the 3000-gallon tank. F02 reads the flow going from the DI bottles into the 560-gallon tank, and F03 reads the amount of flow from the DI bottles to the return header (i.e. the water actually being added to the system). F04 measures the total amount of flow headed for the solid-state MOSFET amplifiers and the anode supplies. F05 represents the flow to the modulators, bias supplies, and coalescing supplies in the gallery; F06 reads the flow to the amplifiers in the tunnel.

In the Cavity system, F01 measures makeup flow from the 3000-gallon tank, F02 is the flow from the DI bottles to the 320-gallon (Cavity) storage tank, and F03 is the flow from the DI tanks to the return header. F04 is flow in the supply line leading to the Test Cave, while F05 is flow in the supply line to the cavities in the tunnel.

Pressure Readbacks

Remember that discharge pressure is read from the output side of the pump, while suction pressure is read from the input side. (Discharge is the same as supply pressure and suction is the same as return pressure.) Of course, if a pump is on and working properly, the discharge pressure will be significantly higher than the suction pressure.

As mentioned earlier, pressure gauges are read directly by the PLCs, and do not go through the Instrument Panel.

I:xxWP01 represents the pressure of the LCW return from the tunnel, measured just upstream of the TCV.

P02 is the pressure in the supply header, before it branches and enters the tunnel.

P03, P04, and P05 measure the discharge pressure from the three LCW pumps S01, S02, and S03 respectively. P06, P07, and P08 represent the suction pressure for those same pumps. The PLCs monitor upper and lower limits on the suction pressure, which are set through the parameters I:xxWP6U, 6L, 7U, 7L, 8U, and 8L. If the suction pressure is too low, the pumps will cavitate, and the PLCs will shut them down.

P09 is the pressure of the pond water at the inlet to the pond pumps, and P10 is the output pressure.

P11 is the water pressure in the DI line, measured between the regulator and the bottles.

The MI-60 Magnet system has about 30 pressure gauges.

P01 measures pressure on the makeup line leading to the RF and Cavity systems. There is no P02 or P03.

P04 is the pressure of the pond water as it leaves the heat exchanger. P05 is the pressure of the LCW as it leaves the heat exchanger. P06 is at the pond strainer outlet, and P07 is at the pond strainer inlet.

The next block of pressure gauges, from P08 to P16, is dedicated to the LCW pumps. LCW pumps in the Magnet system are equipped with 3 pressure gauges, on the suction side there is one at the strainer inlet; the second, “suction pressure,” is at the inlet of the pump. The difference between the two represents the pressure drop across the strainer. The third is on the discharge side. The pumps are interlocked to the suction pressures.

P17 is the pressure of the LCW as it leaves the heat exchanger. P18 measures the water pressure in the outlet line from the 3000-gallon tank, and P19 measures the head pressure inside the tank. P20 measures the discharge pressure from S05 (the makeup pump to the RF and Cavity systems) and P21 does the same for S04 (the makeup pump for the Main Injector ring). P23 measures the pressure of the LCW going into the 3000-gallon tank. P24 reads the pressure coming out of the DI bottles, and P25 measures it coming out. P26 is at the check valve downstream of the DI bottles; P27 is at the strainer.

On the lines running to the magnets and power supplies, P30 and P31 monitor supply and return pressures on the south branch , while P32 and P33 monitor the south branch.

The RF system also has numerous pressure readbacks. The first ten (numerically) monitor the pond pumps and heat exchangers. There is a strainer for the pond water coming in to the heat exchangers; it is bracketed by P04 upstream and P01 downstream. A significant pressure drop across the strainer might indicate that it is getting clogged. P05 reads the pressure of the pond water entering Heat Exchanger #1 and P07 reads it as it leaves. P08 and P10 do the same for Heat Exchanger #2. P03 reads the pressure after the two lines have merged, as the water enters the sprayer of the pond.

P06 and P09 read the pressures of the LCW coming out of the two heat exchangers. Jumping ahead (numerically), P23 and P24 measure the pressure of the LCW going into the heat exchangers.

The block of gauges from P11 to P22 monitors the pressures at the four LCW pumps. Each pump has three gauges. P11, P14, P17, and P20 are “pre-valve” readings, the manually operated valve being underneath the heat exchangers. P12, P15, P18, and P21 monitor the suction. P13, P16, P19, and P22 measure the pressure in the discharge lines.

The many ways that LCW can be moved around the DI bottles and storage tanks can be appreciated by sampling the related pressure gauges. P25 measures the pressure on the output branch of the 560-gallon tank that empties into the return header, adding to the volume of water that is circulating among the components. P27 reads the head pressure inside the 520-gallon tank. P28 measures the pressure on the return side of the DI loop, and P29 measures it on the supply side, upstream of the regulator. P31 is read downstream of the regulator, and also downstream of the point at which flow merges with any makeup flow. The readback for P32 is taken at the outlet of the DI bottles.

Of course, all of this water has to be used somewhere. P36, and P34 further downstream, measure the pressure of the supply header servicing the solid-state amplifiers and the anode supplies in the gallery. P33 and P35 read the pressure on the associated return header. P41, and P38 further downstream, measure the supply pressure to the bias supplies, modulators, and coalescing cavity supplies; P39 and P37 read return pressures. P42 measures the supply pressure to the amplifiers in the tunnel and P40 measures the return pressure.

Exploration of the pressure gauges in the Cavity system is left as an exercise for the reader, because this is becoming way too tedious for the author. The Cavity system is actually considerably simpler than the RF system. The important differences to remember are that pressures are measured on the ICW system rather than pond pumps, and that there is a branch leading to the Test Cave.

Resistivity

On a parameter page, resistivity is measured in units of “Megohms,” by which is really meant Megohms per centimeter.

There are three critical measurement locations at a typical service building. I:xxR01 measures resistivity in the return line from the tunnel. It is here that resistivity is likely to be at its worst, because the water has had a chance to pick up dissolved material as it moves through the headers and magnets. R02 measures the resistivity going into the DI bottles, and R03 measures it coming out. A comparison of R02 and R03 is a measure of how effectively the DI bottles are working.

In the MI-60 Magnet system, I:60WR01 is measured as it leaves the 3000 gallon storage tank; R02 is read as LCW is being diverted from the supply header toward the DI bottles, upstream of the DI filters; R03 is downstream of the DI filters; and R04 samples the resistivity on the return line just upstream of the TCV (i.e. it is the equivalent of R01 in the other service buildings.

In the RF system, I:6RWR01 is the resistivity coming out of the 560-gallon storage tank; R02 is at the inlet to the DI bottles; R03 is at the outlet of the DI bottles; and R04 is on the return line just upstream of the TCV. In other words, the naming convention is the same as that in the Magnet system.

In the Cavity system, I:6CWR01 is the resistivity of the LCW coming out of the 320-gallon storage tank; R02 measures the return water after being scrubbed in the DI bottles; R03 is supply water on its way to the DI bottles, and R04 is water returning from the tunnel or the Test Cave. In other words, the naming convention for R02 and R03 has been reversed from that of the Magnet and RF systems.

Temperature

Since Fermilab is the world's leading institution for the study of the fundamental nature of matter and energy, temperature is measured using the Fahrenheit scale.

I:xxWT01 is the temperature of the LCW in the supply header before it encounters the magnets, power supplies, or DI bottles. At T01 the water has not passed through any heat loads, so the reading at that location can be interpreted as a measure of the overall health of the heat exchange process. T01 is therefore used as the set point controlling the position of the TCV, which regulates the amount of water going to the heat exchanger. If T01 is too high, the TCV sends more water through the exchanger, and vice-versa. Being a controllable parameter, typing in a new D/A value can change the set point, but it is usually set near 90˚.

The PLCs also use T01 as an interlock to protect the magnets from undue thermal stress. The readback is approximately the temperature that the magnets will see. The main threat is to the epoxy and insulation surrounding the magnet coils; too cold and it may become brittle and shatter, or too warm and it could become soft and melt. Either extreme could compromise the integrity of the coils. The upper and lower limits are set through the parameters I:xxWT1U and I:xxWT1L. Typical limits are 40˚F for the lower limit and 150˚ for the upper limit. If the limits are exceeded, the LCW pumps at the service building trip off.

T02 is the temperature of the water going into the heat exchanger, while T03 is the temperature coming out. Comparing the two is a measure of how effectively the heat exchanger is performing. For example, on a hot day the pond water may be too warm to allow the heat from the LCW to be dissipated efficiently, and T02 and T03 will not differ by much.

The return flow, of course, is divided between that which flows through the heat exchanger and that which doesn't. T04 looks at the temperature of the water after the two branches have merged, and is a measure of the LCW temperature just upstream of the pumps.

At the MI-60 Magnet system, T01 measures the temperature of the pond water as it emerges from the heat exchanger; T02 is the temperature of the LCW at the inlet of the heat exchanger; T03 is the pond water temperature as it enters the heat exchanger; and T05 is the LCW temperature after the water has been cooled by the heat exchanger.

T04 measures the temperature of the supply water before it goes into the magnets in the tunnel. It performs the same function as T01 does at the other service buildings, providing feedback to the TCV for temperature regulation and protecting the magnets from temperature extremes. The low trip limit for T04 is set from I:60WT4L, and the high trip limit is set through I:60WT4H.

In the RF system, there are two heat exchangers, but the temperature readbacks monitor lines that are common to both. T01 is just upstream of the pumps, as they pull water out of the exchangers and into the pond. T03 measures the temperature of the LCW as it leaves the heat exchangers and T04 measures it as it is going in.

T05 measures the temperature of the water in the supply header before it goes to the power supplies and amplifiers. This reading provides feedback for the TCV.

In the Cavity system, I:6CWT01 measures the temperature of the supply water headed toward the cavities or the test cave; the reader can probably guess that it is used in the feedback for the TCV. T02 and T03 monitor the ICW temperature going in to and out of the heat exchanger, respectively. T04 measures the return water temperature, just downstream of the TCV, as it enters the heat exchanger.

Liquid Level Readbacks

The storage tanks for the various LCW systems are at MI-60, MI-52, and CUB. There are two types of parameters for reading the amount of water in the tank. One of the parameters displays the water level of the tank in inches, which is the way that the instrumentation hardware actually sees it. The second parameter uses a scale factor to convert the value to gallons, which is intuitively easier for the user.

In the MI-60 magnet system, I:60WL01 represents the height of the water column of the 3000-gallon tank in inches. I:60WVOL converts inches into gallons. A trip indication is generated if the liquid level is out of bounds. I:60WL1L is the lower trip limit, I:60WL1H is the high trip limit, and I:60WL1I is the “low inner trip limit.”

The liquid level in the 560-gallon tank of the RF system is read from I:6RWL01; for the 380-gallon tank it is I:6CWL01, and at MI-52 it is I:52WL01. None of these storage tanks have corresponding parameters for volume.

The "leak" parameters work by measuring the change in the makeup tank levels. In the MI-60 Magnet system, I:LEAK is the almost real-time leak rate, averaging the rate over the previous minute. I:LEAKF (fast leak) is averaged over the previous 5 minutes, and I:LEAKS (slow leak) over the previous ten minutes. Equivalent parameters exist for the RF and Cavity systems (e.g. I:RLEAK, I:CLEAK). A negative leak value means that the tank is being filled.

There is a special measure of makeup flow to the RF and Cavity systems known as the TOTalizer, donated by the Arnold Schwarzenegger Foundation. Makeup flow to the RF system is I:6RWTOT, and makeup flow to the Cavity system is I:6CWTOT. The TOTalizer integrates the amount of makeup water transferred to those systems over a period of time (say, midnight to midnight). If the amount is excessive, as with a leak, an alarm is generated.

Vibration Sensors

These are named I:xxWUxx, where the first two characters designate the house number and the last two numbers correlate to the pump number. The vibrations are measured in Hz. At the time of this writing, none of these devices have actually been implemented, but parameters can be found on some pages.

CUB Parameters

The most important CUB parameters to know might be:

- T:CWVOL, which is the volume of the LCW in the spheres, measured in gallons.

- T:CWL01, the height of the LCW in the spheres, measured in inches. Being spheres, you might want to revert to T:CWVOL if you want to know how much water you have, unless you really like geometry problems.
- I:CWR101, which measures the resistivity coming out of Column I (the column normally used for Main Injector).
- I:CWV28 and I:CWV29, which are the outlet and inlet valves (respectively) to CUB from the Main Injector.
- I:CWV22, makeup water to the spheres.
- I:CWV419 and I:CWS03. When the pump S03 is running and the valve V419 is open, makeup water is pumped from the spheres to Column I, eventually to make its way back to the Main Injector.

Pump Interlocks

Much of the information in this section has already been covered in this chapter, but this section consolidates material relevant to the interlocks.

Traditionally, status for the RF water systems in the Main Ring was monitored through FIRUS. For some reason unknown to the author, this tradition has been extended to include the Main Injector. There are analog readbacks through ACNET, but there are no PLC trip limits or trip flags read through ACNET, nor is there any digital control of the pumps or valves. (On I25, there is the Watchdog trip, which inhibits the RF drive if the RF system conductivity or temperature limits are exceeded, but the pumps are not affected.) The discussion below does not include the RF or Cavity systems.

The trip limits alluded to in the description of parameter names are usually enforced through interlocks to the pump motors. A temperature, pressure, or flow rate that is out of bounds will cause the pumps at to trip off. The trip limits can be modified from Page I56 or from a parameter page.

Trips can be global, or restricted to a particular house or even a single pump, depending on what is being protected. Global trips turn off all of the pumps ring-wide (except for the independent systems such as at MI-52). They include:

- I:60WL01, the liquid level in the 3000 gallon tank, and
- I:60WV03, which connects the tank to the return header. Remember that V03 needs to be open in order to maintain the suction pressure.

House trips shut off all of the pumps at a given service building. They include:

- F02, the flow rate in the supply header going to the magnets,
- T01, the temperature of the water going to the magnets, and
- V14, which is actually a summation of V01 through V04, the valves on the supply and return lines just before they enter the tunnel.

The equivalent parameters for the magnet system at MI-60 are:

- F34, which is the summation of F03 and F04;
- T04 (which has the same function as T01 at the other houses), and
- There is a house trip if both pond pumps trip off. More about that below.

Single pump trips are generated whenever the return pressure (P06, P07 or P08) strays out of limits. The pump is turned off primarily to prevent cavitation. When a single pump trips, a possible consequence is that the local power supplies will also trip.

The LCW pumps are also interlocked to the pond pumps. Normally, only one of the two pond pumps needs to be running at each house. If the one pump trips off, the other is automatically turned on. In order to give the second pond pump time to turn on, there is a 15-second delay before the LCW pumps respond to a trip. A failure of both pond pumps causes a house trip of the LCW pumps, so that hot water is not pumped to the magnets.

At MI-10 through MI-50, the pond pumps are designated S04 and S05, so the interlock summation is S45. For the MI-60 Magnet system, pumps S06 and S07 are summed into S67. Pond pumps are interlocked to high temperature, not to the water, but to a temperature sensor internal to the motor. They are also interlocked to their own output pressure (P10), and to the differential pressure across the pump.

Turning on Pumps

On Page I56, on the “Global” sub page, there is an “Auto-on” option that performs the necessary steps to turn on the LCW pumps ring-wide. (As of this writing, there is no option for turning on the pond pumps globally.) This is what it does:

- Sends a reset to all of the trip flags. This checks that the pond pumps are on and the inlet and outlet valves to each building are open. Some of the trip flags, such as flows and pressures, will have bad status because the pumps are off. To avoid a viscous circle, there is a delay of 15 seconds before those trips re-assert themselves in order to give the pumps time to turn on.
- V03 and V04, the DI loop valves in the magnet system, are opened. Remember that it is these valves, especially V03, which maintain the proper pressures in the ring.
- Pumps 1 and 2 at each house are turned on.
- After several seconds, Pump 3 at each house is turned on.

The MI-52 system does not yet have an automatic startup sequence. To turn on the pumps manually, a similar sequence is followed.

PLC Logic

The turn-on sequence, of course, is controlled by the PLCs. The PLCs progress through a series of operational states, as shown in Fig. 4-9.

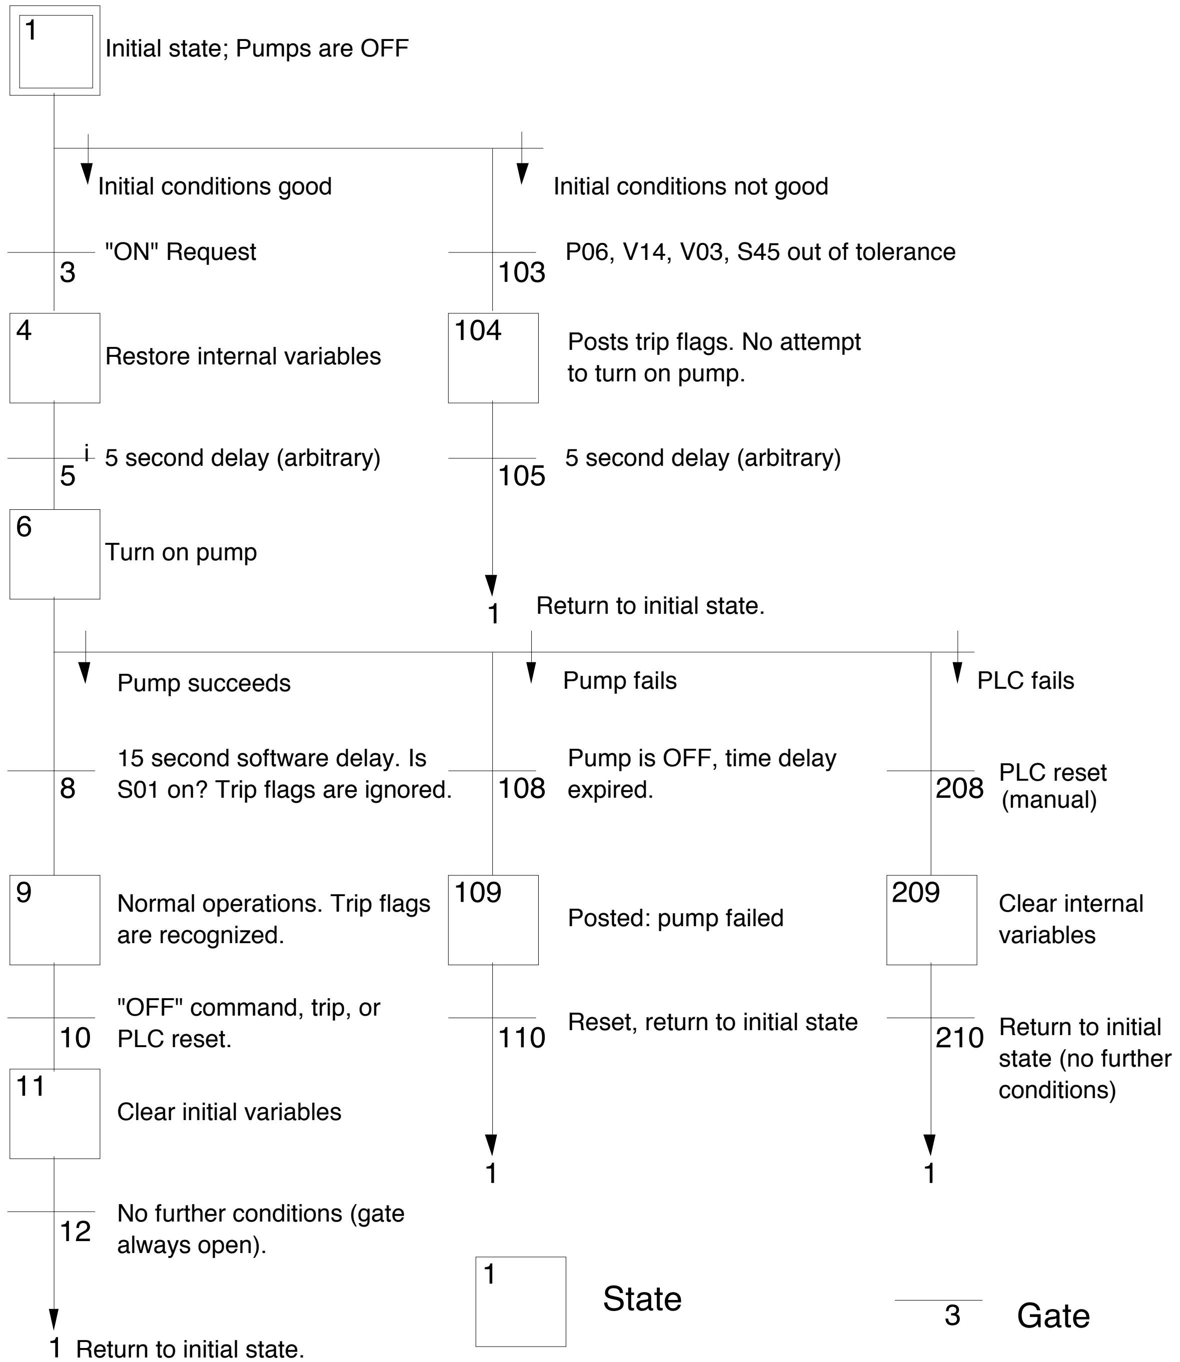

Fig. 4-9 Typical PLC State Logic for LCW pump

The PLC state is usually read from I56, where states 1, 2, and 3 refer to the LCW pump number. The operational state can also be read as a parameter, for example, I:10WST1 represents LCW Pump #1 at MI10.

In Fig. 4-9, the boxes represent operational states, while the horizontal lines are gates leading to the next state. A successful sequence would begin at State 1 (the initialized state) and, passing through Gate 8, ends at State 9. If the machine is operational, most of the pumps are in State 9. If an "off" command is sent, or a trip detected, the logic continues through State 11 and returns to State 1. The same is true of a PLC reset, which reinitializes the logic and shuts the pump off.

If the pump fails to come on during the 15-second delay at Gate 8, it passes instead through Gate 108 and ends up in State 109. Usually pumps that are off are in State 109.

If an "on" command is sent, but the suction pressure, inlet and outlet valves, or pond pumps are not in the proper state, the logic is routed through Gate 103 and back to the initial state. No attempt is made to turn the pump on.

This chapter has been about filling pipes with water. The next, "Vacuum," will be about how to keep pipes empty.

Notes:

Chapter 5: Vacuum

A particle accelerator abhors a vacuum.
- Aristotle

The beam tube in the Main Injector, as with any accelerator, must be kept under vacuum. Whatever air molecules may be in the pipe, act as obstacles to the circulating protons and antiprotons. Sometimes the particles are disturbed to the point where the beam quality is degraded, and sometimes they are scattered out of the beam pipe altogether.

Unlike the Antiproton Source and the Tevatron, in which the particles must circulate for hours, days, or weeks, beam in Main Injector is never in the machine for more than a few seconds. Although a “few seconds” may mean hundreds of thousands of revolutions, the vacuum requirements are not quite as stringent as with accelerators in which beam is stored.

The quality of vacuum is frequently measured in units called torr. Atmospheric pressure is traditionally measured in inches or millimeters of mercury as measured by a barometer. One torr represents one millimeter of atmospheric pressure. For comparison, pressure at sea level is 760 torr (except during hurricane season). Vacuum in the Main Injector is in the neighborhood of 2×10^{-8} torr. Another unit of pressure is the micron. One torr is equivalent to 1,000 microns.

Vacuum Pumps

Different types of vacuum pumps are effective at different pressures. The beam pipe in the Main Injector, if starting from atmospheric pressure, is pumped down in three stages, using roughing pumps, turbo molecular pumps, and ion pumps, in that order. Of these, only the ion pumps are permanently affixed to the ring, and are responsible for maintaining the high vacuum in the pipe. The roughing pumps and turbo pumps, combined as a unit, are on portable carts. They are moved to locations where the vacuum is especially poor, as during initial pump down.

Roughing Pumps and Turbo Pumps

A roughing pump uses a mechanical piston to remove air from the beam tube. It is the only type of pump effective at atmospheric pressure. It is capable of pumping down to about 10 torr.

The turbo pumps are multi-tiered “fans” which drive the air molecules out of the pipe. They are turned on when the roughing pump has removed what it can. The blades spin at a rate of several tens of thousands of RPM; the ends of the blades are actually moving faster than the molecules they are trying to hit. The turbo will pump the beam pipe down to approximately 10^{-5} torr.

There are about ten carts with the roughing/turbo combination available for responding to vacuum problems in the Main Injector. If everything is running well, none of them are needed. The units are connected as needed at the pump out ports at the downstream side of some of the main quadrupoles.

Pumping Ions

Any pumping scheme that expels air outside the system, such as the roughing and turbo pumps, will have some problems with “back-streaming,” or the tendency for air to slip in from the outside. Ion pumps avoid this problem because, as they take over from the turbo pumps at about 10^{-6} torr, there are so few molecules left that they can just be swept under the rug. They are literally buried in the pump itself.

Inside an ion pump, 5200V is placed across a stainless steel anode and a titanium cathode. Free electrons in the electric field will be accelerated and strike the air molecules, freeing even more electrons. The resulting positive ions will also pick up energy from the electric field and be pulled toward the cathode. The more chemically reactive molecules, such as oxygen (and to a lesser extent, nitrogen) will combine chemically with the titanium. Other molecules, such as the chemically inert argon, may still have enough energy to bury themselves in the metal of the cathode.

The high voltage for the ion pumps comes from individual power supplies in the service buildings. Red cables (a safety standard for high voltage) connect the power supplies to the pumps. It is important not to use wet dirty scissors when severing the cable.

The cathode and anode assemblage is encapsulated in permanent magnets, which bathe the components in a magnetic field of about one-kilogauss. The magnetic field causes all of the charged particles to spiral, increasing their path length and their chances for ionization.

The pumping capacity of an ion pump is measured in units of liters/second (L/S). It may be difficult to interpret that unit literally, but it is proportional to the number of ions removed from the beam pipe. The vast majority of ion pumps in the Main Injector ring are 30 L/S. There are a few locations that need powerful pumps, such as the regions connected to beam transfer lines and near the RF cavities. The ion pumps used in these critical regions have a capacity of either 300 L/S or 600 L/S.

The density of ion pumps in the tunnel is high, three per half-cell in the normal arcs, and two per half-cell in the dispersion-suppressor regions. They are named by adding a digit to the location number. For example, at 629, a location in a normal arc, IP6290 is the pump between the Q629 and the "A" dipole just downstream, IP6291 is between the "A" dipole and the "B" dipole, and IP6292 is downstream of the "B" dipole (IP6292 looks as if it belongs to the downstream location, but it doesn't). Fig. 2-19 illustrates ion pump nomenclature used around a main quadrupole.

The dispersion-suppressor cells are shorter and only require two pumps. For example, at 639 there is no ion pump between the quad and the "C" dipole; IP6391 is between the "C" and "D" dipoles, and IP6392 is downstream of the "D" dipole.

In the straight sections, the ion pumps are given letter designations. For example, at Q100, where the MI-8 line meets the Main Injector ring, the ion pumps are designated IP100A through IP100G.

More details on the various beam line vacuum systems will be included in the chapter on Beam Transfer Lines.

Under normal circumstances, the ion pumps are sufficient to maintain vacuum in the Main Injector ring. If, however, there is an air leak, the pump will respond by drawing more current as it impounds the air molecules. If vacuum becomes worse than, say, 10^{-5} torr, or if the pump is shorted, the titanium cathode will overheat and begin to outgas its store of trapped molecules. At that point, using the ion pump to improve vacuum has become a losing proposition, and it trips off. If the leak is bad enough, what will often happen is that several pumps will trip off sequentially as they try to take over the work of those who have fallen before them. Sometimes it is possible to turn the pumps back on after they have had a few minutes to cool, but if not, an access has to be made to fix the leak, and the turbo pumps brought in to recover the vacuum.

Sector Valves

[Author's note: This section describes, among other things, how the sector valves are interlocked to the ion pumps. That is the once and future plan. However, during the commissioning phase, when the new pipe was laden with trapped gases and the beam was not yet properly tuned, the valves would close so frequently that commissioning became impractical. Since that time, the valves have been interlocked only to the Pirani, which are less sensitive. At some point in the future the ion pumps will be used again for interlocking.]

The beam pipe in the Main Injector ring is divided into 18 vacuum sectors (see Fig. 5-1 on the next page). Sector valves can close to create a barrier to bad vacuum, thereby isolating vacuum problems to relatively short sections of the ring.

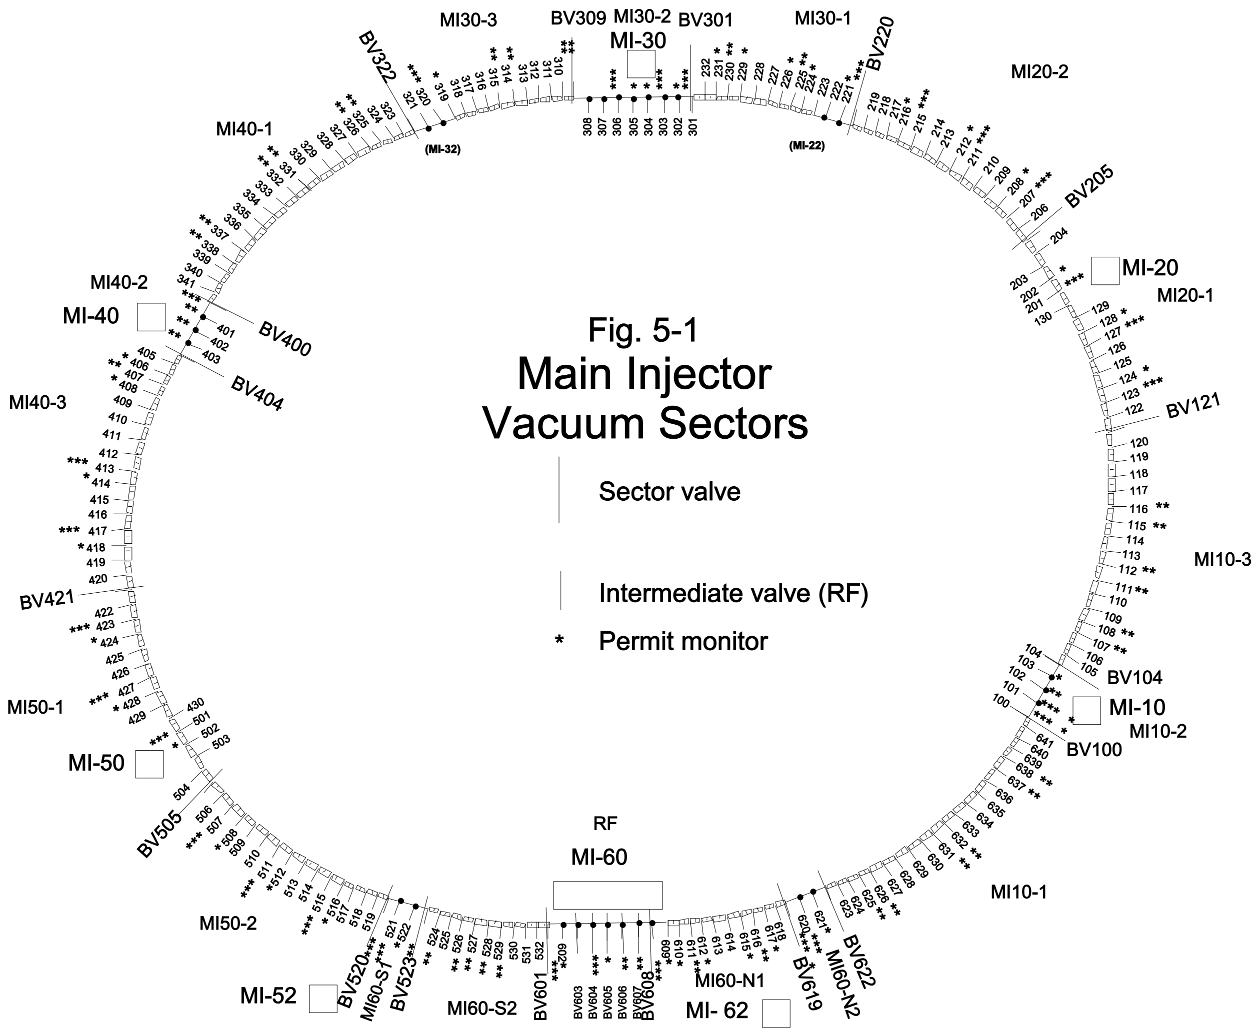

Fig. 5-1
Main Injector
Vacuum Sectors

(Unfortunately, the valves are also effective barriers to the beam; they are often referred to as beam valves.) The status of the valves can be read through the controls system. The beam valves are designated “BV” followed by the location number.

The length of the sectors varies. The beam valves are close together where vacuum activity is particularly intense, or where there are connections to external beamlines. Sector valves bound most of the straight sections. The RF Sector has several beam valves internal to the sector, called intermediate valves, for isolating short groups of RF cavities.

On the vacuum applications page, the sectors are usually named for the service building where its equipment resides, and numbered consecutively if there is more than one sector associated with a service building, e.g. MI10-1, MI10-2, and MI10-3. Locally, however, inside the service building, the sectors are labeled for the first location (proton direction) in the sector; the same three sectors at MI-10 are labeled #622, #100, and #104.

The sector valves can be closed manually, but they are also interlocked to selected readbacks from the ion pumps. Except for MI-60 and some of the very short sectors, each sector has three sets of ion pumps that can be used to interlock the valves. Each set includes four pumps; at least two of them need to be on in order to grant a permit (see Fig. 5-2).

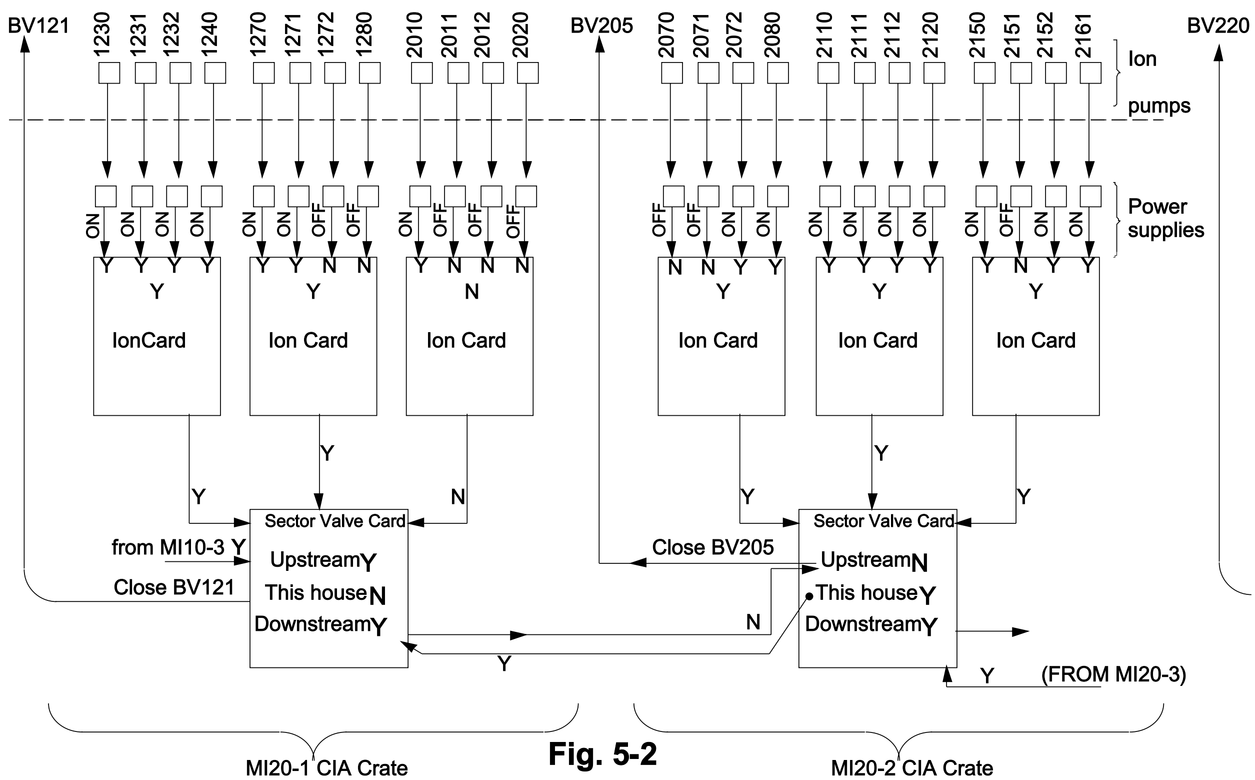

Fig. 5-2
Ion Pump Permits

In this diagram, a vacuum problem has begun to develop near BV205; a majority of the pumps in the downstream block have tripped off. The block has sent a "NO" vote to the permit and for M220-1, so BV121 and BV205 are ordered to close. Compare to Fig. 5-1.

Readers should ask themselves, if either 2072 or 2080 trip off, will BV220 also close?

If a beam valve between two sectors is to be opened, a permit must be obtained from both the upstream and downstream sectors. For example, in order to open BV205, permits must be granted from both MI20-1 and MI20-2.

Two beam valves, BV301 and BV309, are used by the safety system to control circulating beam in the Main Injector ring. The valves will close if the electrical permit is present and the radiation permit is not. (If the electrical permit is not present, it is impossible to circulate beam, and there is no need to close the valves.)

The sector valves are held open by air pressure. Air compressors at MI-20, MI-40, and MI-60S maintain pressure of 80 to 100 psig in copper tubes. The tubes enter the tunnel penetrations and run parallel to the cable trays on the ceiling; the pressurized air can be tapped from the tubes at intervals, and, of course, where there are sector valves. A failure of the air compressor to maintain pressure will eventually close the valves, although each valve is provided with a ballast tank. The valve will also shut in the event of a power failure. Unlike the old Main Ring, the air compressor is dedicated to the vacuum system, and is not shared with the LCW equipment.

The pressurized air can be (and usually is) sent through a unit that cools and dehumidifies the air before it enters the tunnel. The unit can be manually valved out if necessary.

Vacuum Instrumentation

There are two methods for measuring vacuum in the Main Injector. One is with a Pirani (Pe-RAH-nee) gauge. A Pirani (which have occasionally been known to skeletonize a cow in less than three minutes) works by heating a wire with a constant current source and measuring its resistance. The resistance of the wire increases with temperature. Air molecules near the wire carry away the heat. The better the vacuum, the fewer the air molecules and the hotter the wire, so the applications software as can interpret the resistance of the wire pressure. Pirani gauges read back measurements from normal atmospheric pressure down to 10^{-3} torr; they are most accurate at the lower end of that range.

One alias for a Pirani gauge is “convection” gauge, because heat is carried away by the convection of the air. There are two Pirani gauges in every sector, including the "sub sectors" in the RF sector. They are found near the bottom of the pump out port and can be recognized by their blue caps.

The second method of measuring vacuum is through the ion pumps themselves. The arc of an ion pump creates a current that can be read by the power supply. Since the current is proportional to the number of molecules ionized, it can be interpreted by applications software as pressure; the vacuum readback on an ion pump is derived from the arc current. The readback is valid for the range in which the ion pumps are active, from about 10^{-4} torr down to at least 10^{-9} torr.

Vacuum Controls

The applications program that interfaces with the vacuum system is called, inexplicably, "MI VACUUM." It is found on Page I55 (no relation to the Stevenson Expressway). The program also interfaces with the Tevatron, Switchyard, and Booster vacuum systems.

The individual sub pages on I55, 8GeV, MI10-1, MI-10-2, etc., select which of the 18 sectors will be viewed. You can find the Pirani gauges readbacks from the and ion pumps on the "Main Page"; the status of the beam valves and their permit status is also here.

Locally, at the service buildings, the CIA crate controls the vacuum (Controls Interface Adapter). There are several kinds of cards in the CIA crate. As a rule, there is one ion card for every 6 ion pumps. The ion card gathers the readbacks from its little domain of ion pumps; the first four channels on each card can be used to establish the beam valve permits. (There are usually two ion pumps monitored by each ion card that are not represented on the front panel.) In most sectors, three cards, that is, 12 ion pumps, are tied into the beam valve permit. The three groups of pumps are spaced apart to more accurately sample the length of the sector. If three pumps within one of the selected groups fail, the permit is withdrawn and the beam valve closes.

The sector valve card is what actually decides whether a beam valve should close or not. It directly controls only the upstream beam valve for the sector. For example, Sector MI20-1 is bounded by BV121 and BV205. The sector valve card for MI20-1 is responsible for BV121, but the sector card for MI20-2 is responsible for BV205. In the situation represented by Fig. 5-2, IP2011, IP2012, and IP2020 have all tripped off. With the majority of the pumps off, the ion card representing those pumps, interpret the information as a vacuum problem that could spread to other sectors. It sends a "NO" bit to the sector card for MI20-1, which orders BV121 closed. It also sends a status bit to the "Upstream Permit" bit of the sector valve card for MI20-2, which responds by closing BV205.

Locally, on the CIA crate, the sector valve card displays the permit status and the open/closed status of each of the two valves. There is also a "request" bit that indicates that a valve has been asked to open but has not yet complied. The same information is available on I55 under the "Sector Valve" sub page.

Alarms are generated when any three consecutive ion pumps are off, or if 20% of the total number of pumps in a sector trip. The alarms are independent of the valve permits, but provide a warning that vacuum problems are developing.

The CIA crate interfaces to the outside world using an Arcnet loop, not through CAMAC as with some of the other machines. There is a 186-bearing card in the rightmost slot that functions as the interface. All of the CIA crates in the Main Injector talk to a dedicated VME front end, called MIVAC; MIVAC is located at MI60 south, Rack 123. The Arcnet loop uses the CATV cables (the same system that carries TV images) to relay data back and forth from the CIA crates to MIVAC. The VME, in turn, is linked to the general controls system through Ethernet.

The VME uses a MOOC software platform, so many of the controls related error messages will read back as MOOC errors.

Notes:

Disclaimer

Chapter 6, "Beam Diagnostics," is not anywhere near completion. It is being released now because it contains some basic information about multiwires, BPMs, and BLMs. There are no diagrams. System experts have not recently evaluated it.

Future versions of this chapter will include information on flying wires and "remote sensing" instrumentation such as DCCTs, stripline detectors, and bunch length monitors. There will also be material on the timing of diagnostics, and, of course, pictures.

Chapter 6: Beam Diagnostics

> I'm beginning to wonder whether the physicists are pulling some kind of elaborate scam here. I'm starting to wonder if they don't sit around their $23-million atomic accelerators all day, drinking frozen daiquiris and shrieking, "There goes one now!" and then laughing themselves sick. Maybe it's time we laypersons asked some hard questions about this idea that all matter consists of tiny invisible particles whizzing around.
>
> \- Dave Barry

Dave Barry is wrong of course, our accelerators cost a lot more than $23 million. But how do we know that there are tiny invisible particles whizzing around the Main Injector? And, almost as importantly, how do we know that they are doing what they are supposed to? How wants a daiquiri? This chapter will attempt to answer those questions.

There are a few basic kinds of techniques used to look at beam. One requires that an object, such as a wire, be put directly into the beam. This method needs to be implemented with great care, because if done improperly the beam can be destroyed or significantly degraded. Devices called multiwires are used in "one-pass" beam lines, where the beam will only strike each set of multiwires once. Multiwires are inappropriate for looking at circulating beam, since the wires and the beam would be destroyed in very short order.

The solution adopted for using wires in a circulating beam is called a flying wire. The flying wires are very thin and move through the beam very quickly, so that disruption to the beam is minimized.

Both of these wire methods are destructive to the beam to some extent, so there is always a trade-off: Use the wires and degrade the beam, or don't use the wires and remain ignorant about what the beam is doing.

The second type of approach is less destructive. Protons and antiprotons have an electric charge, and can be "remotely" sensed by devices designed to detect beam electromagnetically. (This method is not entirely risk-free. There are image charges created in the instrumentation that can, in turn, influence the beam. There is no such thing as a free lunch.) Devices that detect the beam electromagnetically include the beam position monitors (BPMs), fast bunch integrators (FBIs), bunch length monitors, DC Current Transformers (DCCTs), SBDs (Sampled Bunch Displays), and resistive wall monitors.

A third, and indirect, way of "seeing" the beam is through the beam loss monitors (BLMs). When beam strikes an object, a spray of secondary particles is produced which can be detected outside of the beam pipe. (This is the reason no one is allowed in the tunnel when beam is present.) It is preferable that the losses not occur in the first place, but when they do, they provide valuable information about the location of the beam.

All of the methods mentioned above could be used to create a picture of what the beam is doing. Now, one at a time.

Multiwires

The multiwires are used in the beam lines: the MI-8 line, the abort line, and the P1, P2, P3, and A1 lines. This section describes the general principles behind the design and operation of multiwires; the specific layout in each line will be covered in Chapter 7.

When beam strikes a metal wire, electrons are dislodged. If the wire is part of a circuit, current will flow into the wire from elsewhere in the circuit to replace the missing electrons, and the charge imbalance can be measured. The greater the intensity of the beam striking a particular wire, the greater the charge displaced:

[Future Picture]

Horizontal and vertical sets of wires are strung in a G-10 board (G-10 is the name of a green fiberglass-reinforced epoxy commonly used at Fermilab). The wires are normally placed 1 mm apart. Sometimes the board is referred to as a paddle, because it rotates around the vertical shaft. The assembly is suspended in an evacuated metal can in the beam pipe. The horizontal set of wires measures the vertical profile of the beam, and vice-versa.

Because of the way that the board is rotated into the beam, only one plane can be measured at a time. The wires can move from the “out” position to either the horizontal or vertical position, and back again. Since the wires scatter some of the beam, it is usually preferable that the wires be out of the beam unless the beam profiles are actually being measured. Also, if the wires are in an intermediate position, the beam may hit the G10 board and not get through at all.

In addition to the position readbacks available in the control room, the position of the wires can be determined in the tunnel by looking at the top of the can. There is a pin visible from outside the can that is linked to the assembly:

[Future Picture]

The limit switches to either side of the "out" position provide the readbacks to the software read in the control room.

The cables that connect the multiwires to the electronics upstairs have a known capacitance, and cables that can store a charge:

[Future Picture]

When the time comes to read the beam intensity, the switches are closed and the charge accumulated by the wires is read by the electronics.

The result is a picture that looks something like this:

[Future Picture]

Each column in the histogram represents the beam intensity as measured by one wire. The horizontal and vertical planes are read separately.

The multiwires provide information about the shape of the beam that is not available from a simple position readback. For example, if the beam is too wide or too narrow, there could be a problem with one of the quadrupoles upstream; or, if the beam is scraping some object, the "shadow" can often be seen in the Multiwire plot. Once nominal profiles have been established, hardcopies or save files can be used to compare to current conditions and to diagnose problems.

Multiwire Controls

Some of the multiwires interface to the controls system through CAMAC. There are two types of CAMAC cards used by the multiwires. One is the 184 card, which is the motor controller card. It talks to the stepping motor power supply, which in turn sends streams of pulses to the motor downstairs that rotates the paddle. Each pulse rotates the motor through an incremental angle.

The other CAMAC card is the 192 card, which digitizes the data from the cables and makes it available to the CAMAC links. The 192 card has a small microprocessor on board which can make a few simple calculations on the data, including determination of the beam centroid (average center position) and the beam sigma (related to the width of the beam).

Multiwires not connected to CAMAC are operated through SWIC controllers similar to those used in Switchyard. They are linked through an ARCNET loop that feeds the data into the ARCNET card of a VME crate; the data returns to the computer room via Ethernet.

It is important that the wires sample the data at the exact moment that beam is present. The Multiwire electronics listens to the beam sync clocks to synchronize the sample time with the arrival of beam. (Remember that the beam sync clocks also trigger the kickers responsible for transferring beam.) Specific clock events used in each line are discussed in the chapter on beam lines.

Beam Position Monitors

The Beam Position Monitors, or BPMs, measure the average position of the beam at numerous locations in the beam lines and around the Main Injector ring. Unlike multiwires, which must be put directly into the beam, BPMs detect the beam electromagnetically. This property allows them to be used for measuring circulating beam.

The Main Injector BPMs are located just downstream (proton direction) of each quadrupole in the ring, and at strategic locations in each of the beam lines. The physical construction of the detectors differs from those of the former Main Ring in that each one can measure the beam position in both planes. Normally, only one plane is selected for readback, but where deemed necessary both planes can be read back simultaneously.

Structurally, the BPM detectors have been cut from the inner surface of the beam pipe itself. They consist of four strips of steel, each connected to an output at the upstream end. The strips are arranged more or less diagonally around the beam pipe, as in this longitudinal view:

[Future Picture]

The principle used here is basically that of a stripline detector; several variations of this type of device will soon be encountered. As our tiny and invisible but charged particles whiz by the stripline, an image charge is created in the metal. "Image charge" is just a way of saying that as a proton passes next to the strip, electrons attracted to the positive charge will be drawn to the surface of the stripline, forming an "image" of the proton; likewise, an antiproton creates a positive "image" as the electrons are repulsed. Since the particle bunches are traveling in a direction, the image charge becomes an image current traveling in the opposite direction of the beam. It the image current which is sensed at each of the four outputs:

[Future Picture]

What makes the BPMs work as position detectors is that the strength of the signal is proportional to the proximity of the beam to the stripline. The beam position is calculated by comparing the relative strengths of the four signals. (Those who understood how BPMs worked in the Main Ring are at a slight disadvantage now, because there were two dedicated plates for measuring the horizontal position, and two for the vertical. In the Main Injector, information for the two planes must be extracted from all four detectors.) Suppose that the beam is left of center. The two striplines on the left will each see a stronger signal than the two on the right.

BPM Electronics

The signal developed from the image current is first sent to an RF module, located upstairs in the service building. There is a dedicated RF module for each of the BPMs. The term “RF” may be a bit misleading in this case, because the modules have nothing to do with generating the RF for the acceleration systems. They can be better thought of as RF detectors or receivers, tuned to 53 MHz because that is the rate at which bunches are passing through the detector. They compare the signal strengths from the BPMs and convert them into DC analog signals palatable to the next stage in the processing, the Analog Box.

The Analog Box—another slightly misleading name—is a dedicated MADC that digitizes the analog signals from the RF modules and delivers the data to the microprocessor.

BPM Controls

In the Main Injector, the BPM microprocessors are actually called MBP (Main Injector Beam Position) microprocessors, because the Tevatron had already copyrighted the name “BPM”. The MBP is capable of organizing the BPM data coming from the Analog Box in several different ways. Data is written into a variety of buffers. Keep in mind that the microprocessor at a given service building only handles data for the BPMs under its control, and only sees part of the picture. When a BPM application program is called up in the control room, the Main Injector front end gathers data from each of the service buildings and integrates it into a single picture, or frame. A frame represents data taken at a given moment in time.

The most fundamental variation in the way data is taken and stored is between flash frames and snapshot frames. Flash frames only look at one turn of beam, that is, a single sample of beam all the way around the ring once, or once through a beam line. They are usually used to look at either the first or last turn of beam in the machine, at the moment of injection or extraction. Flash frames require exquisite timing, which is orchestrated through the beam sync clocks. The sample of beam must be measured at each BPM location as it moves around the ring at nearly the speed of light. There is a beam sync event for the injection or extraction kickers; a delay set by a CAMAC 279 card waits for a specified number of Main Injector revolutions before triggering the BPMs. When the delay is complete, the flash trigger is fanned out simultaneously to all of the service buildings. Local delays are implemented at the service buildings to ensure that the trigger arrives precisely when the sample of beam is passing through the BPM. The flash data is then stored into a dedicated flash buffer. The data in the buffer is replaced with new data each time a new injection or extraction event is issued.

The snapshot frames do not need to be timed quite as stringently. The BPM signals are averaged over several turns—8 turns is typical—before being assembled into a frame. After this flurry of activity, there is a wait of several milliseconds before data for the next frame is taken:

[Future Picture]

The circular buffer for the snapshot frames (that is, a buffer that continuously writes new data over the oldest data) holds up to 512 entries. It "freezes" at the end of a beam cycle, or when an abort occurs, saving the last second or so of data. In that way, the history of a problem can be analyzed as it developed.

There are two other buffers that take their data from the circular buffer; these are snapshot frames that have been earmarked for specific times. One is the display buffer, for which one time in the machine cycle can be selected and displayed for every pulse. New data is written into the display frame once a cycle.

The other is the profile buffer, which selects a series of sample times within a cycle. It holds up to 128 entries, all taken from the snapshot buffer, and are spaced at equal intervals chosen by the user. It also is refreshed once a cycle.

Several additional application programs use the BPMs. They will be covered in the prophesied chapter on beam tuning.

In the beam lines the MBP microprocessors are not used. Any given pulse of beam only passes through once. Data is sent back to the appropriate front ends via the MADC chassis, and the applications software assembles the flash frame. There is no need for the relatively complex overhead required to generate the display, snapshot, and profile frames.

Beam Loss Monitors

When all else has failed, and beam is so far off course that none of the other beam detectors will work, there are always the beam loss monitors. (The BLMs are useful in less drastic situations as well.) The BLMs are placed outside the beam pipe in order to detect the spray of secondary particles created when beam strikes an object, such as the inside of a beam pipe.

Externally, a BLM presents itself as a metal cylinder, usually a few inches long, with cables connected at either end. Inside the cylinder is a sealed glass tube filled with argon. The tube also houses electrodes: the anode runs along the axis of the tube, and the cathode is a metal cylinder just inside the surface of the glass. There is a potential of about 2 KV across the electrodes. When a particle passes through the glass tube, some of the argon is ionized. The ions drift toward their favorite electrode and a current, proportional to the number of particles passing through, is produced.

The red cables connected to the BLMs provide the high voltage to the electrodes; the power supply for these is upstairs. The green cable connected on the other side of the cylinder carries the detected signal to the electronics, which is also upstairs.

There are two integrators in the electronics; the signal from the BLM is split and sent to both integrators. One of the integrators is “lossy” and provides something resembling a real-time signal; the other actually integrates.

The microprocessor that coordinates BLM data is actually the BPM microprocessor. One consequence of this union is that the times for sampling the losses correspond to the BPM timers: e.g. the BLM Display timer will be the same as the BPM Display timer.

As with the BPMs, BLMs in the beam lines return data to the controls system through the MADCs instead of MBP microprocessors.

The following chapter is correct, but is missing some information and pictures. It will be updated at a future date.

Chapter 7: Beam Transport Lines

I'd like to be
Under the sea
In an octupoles' garden
In the shade

Main Ringo

The purpose of the beam lines is to transport protons and antiprotons from one accelerator to another, or to the Fixed Target experiments. The beam lines to be discussed in this chapter include the MI-8 line, which delivers protons from the Booster to the Main Injector; the abort line, in which protons are kicked out of the Main Injector to a beam absorber; the P1 line, which serves as a connection between the Main Injector and the Tevatron or other beam lines; the P2 line, which transports beam towards the Antiproton Source and the Fixed Target experiments; the A1 line, which transfers antiprotons from the Main Injector to the Tevatron; and NuMI, which sends protons to an external target and ultimately neutrinos to Minnesota. (These beamlines are sometimes given alternate names—the P1 line is sometimes referred to as the P150 line, and the A1 line is also known as the A150 line, because transported beam in those lines is often at 150 GeV.) The P3 line will not be discussed in detail in this book, because it is remote from the Main Injector ring and should more properly be thought of as a dedicated beam line to the Fixed Target experiments. Sometimes, the P2 and P3 lines are collectively named the Main Ring Remnant, since most of the magnets were once part of F Sector in the Main Ring.

All of the beam lines have some things in common: for one, any given pulse of beam only passes through them once. This simplifies their design, because there is no need for any of the devices specifically needed for circulating beam, such as sextupoles, octupoles, or RF. It also makes possible the use of beam-intrusive instrumentation, such as multiwires, which can survive one hit of beam each cycle but not the tens of thousands of hits they would have to sustain if they were in the circulating beam.

In a general way, all of the beam lines also share a common process for transferring beam from one machine to another. Kickers—very fast but relatively weak magnets—deflect the beam in order to position it with respect to a septum magnet. A magnetic septum has one aperture with a magnetic field, and another aperture that is allegedly field-free. It can be thought of as a way to place a beam pipe and a magnet next to each other when there is not enough room for them to coexist as separate components. The kicker determines which aperture the beam enters. Normally (but not always), the kicked beam enters the field region of the septum and is bent into the beam transport line, while the circulating beam passes through the field-free region. The specific type of magnetic septum found almost universally in the Main Injector beam lines is called a Lambertson magnet, named for its inventor.

When beam is being injected into an accelerator from a beam line, the sequence is reversed—a Lambertson magnet bends the beam onto the appropriate trajectory, and then a kicker corrects the angle so that everything is on the proper orbit.

The beam sync clock initiates the very precise timing required by the kickers. Beam sync events are responsible for injection timing as well as extraction timing.

Kickers

Magnet Construction

Kicker magnets are unique because they have to be very fast. Beam that is being injected into or extracted from a machine has to be deflected without affecting any beam that is already circulating. Beginning with zero current, a kicker magnet has to achieve full current quickly—usually within a microsecond or less—and remain at a steady current until all the beam has passed through. Often, the current must then be removed from the kicker within a very short time to prevent subsequent beam from being kicked.

Some of the kickers, such as those in the abort line and P1 line, must be prepared to deflect beam at a variety of energies. In contrast, the injection kicker at the end of the MI-8 line and the kickers for transfer in and out of the Recycler only need to deal with 8 GeV beam, and the kicker at MI-62 for extracting antiprotons to the Tevatron always operates at 150 GeV. For all of the extraction kickers, the object of the game is to get the beam into the field region of a Lambertson magnet; injection kickers generally accept beam from a Lambertson magnet.

Since the inductance of a magnet limits the speed at which the field can change, kickers are built with a single “turn;” there is no coil as such. The current only makes one pass through the magnet. The field generated by the current is amplified by sections of ferrite surrounding the beam pipe. Unfortunately, the absence of a coil also means that a kicker magnet is relatively weak, so the beam trajectory must be adjusted to minimize the amount of deflection required.

For the following discussion, refer to Fig. 7-1, which is a block diagram representing a generic kicker power and control scheme. Variations on this theme will be discussed as the tour of the various beam lines progresses.

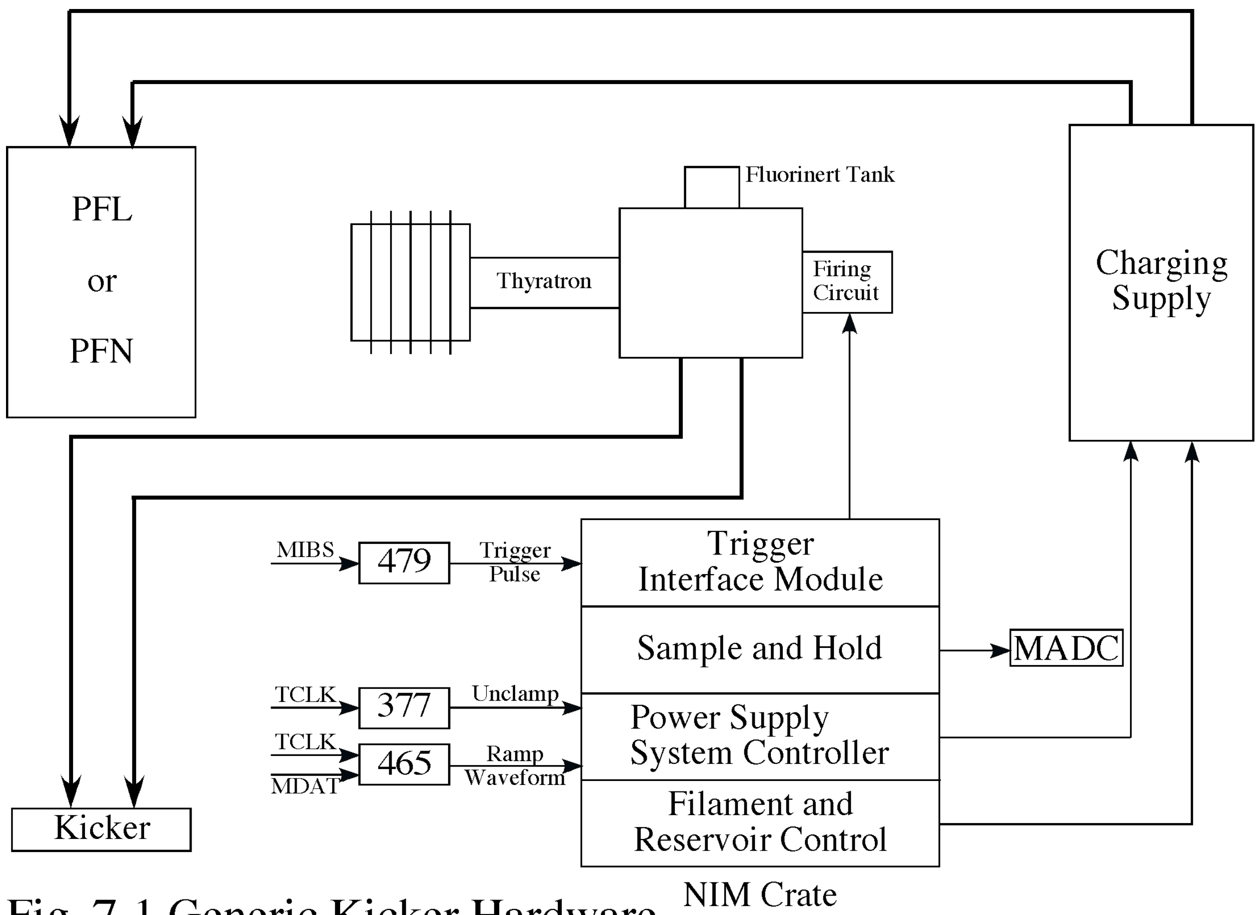

Fig. 7-1 Generic Kicker Hardware

PFNs and PFLs

The current in a kicker magnet is at a very high level for a very short time. Energy is dissipated in the magnet faster than it can be supplied, so a mechanism is required to store energy and then release it at the proper time. A charging supply feeds charge into a resonant system and the energy is stored in the capacitance and inductance of the system. In a Pulse Forming Network (PFN) the capacitors and inductors exist as discrete components; in contrast, a Pulse Forming Line (PFL) uses the “natural” inductance and capacitance of a large (RG-220) cable. After a PFN or PFL is fully charged, a trigger event allows the energy to be discharged into the system at the proper time.

However, if the energy in a capacitor or inductor were to be discharged directly into a magnet, the current would start out high and decay exponentially—the strength of the magnetic field would decline significantly over the interval of time that the beam is present. Instead, a waveform is required that is flat through the duration of the beam pulse. A PFN gets around the problem of exponential decay by modularizing the capacitor/inductor units—the greater the number of modules the closer the current waveform will approximate a square wave. The ideal square waveform is flat on top, with nearly instantaneous rise and fall times. In contrast, a PFL is designed so that the inductance and capacitance of a cable acts as a continuum of infinitesimal “modules;" here, the waveform can also be made flat, but for a shorter period of time. PFLs are adequate for the duration of a Booster batch (2.2 microseconds). PFN's are found at locations where several consecutive batches needed to be kicked, as during Fixed Target mode.

The charging supply is usually a Spellman or Glassman module. The module itself usually runs around 3 KV or so, but the charge that it delivers to the PFN or cable can produce voltages in the tens of kilovolts. The supply is controlled via a NIM-based “Power Supply System Control” module. Of course, the system must be charged to a specific voltage, so a reference voltage is sent to the control module via a CAMAC card. If the beam energy is constant, a CAMAC 118 or 119 card will do—these cards establish a DC reference. If the kicker voltage must track changing beam energy, a C465 card is required. The kickers at MI-10 use C118 cards, those at MI-62 use C119 cards, and those in the abort line and at MI-52 use C465 cards.

To save wear and tear on the hardware, the supply is “clamped” down until an unclamp event is sent to the controller. The unclamp event is sent from a C377 card and is usually referenced to a Main Injector reset-specific event (say, a $2B). When the controller receives the event, the supply begins to charge, tracking its reference.

Thyratrons

It takes longer to charge up a PFN or PFL than it does to discharge into the kicker magnet. A Thyratron is a massive switch—actually, a tube—that isolates the PFN or PFL from the kicker magnet until it is time to fire. Generally, hydrogen, or its isotope deuterium, serves as the charge carrier. When the Thyratron is triggered, and the charge carrier begins to conduct, the system "avalanches" as charge is suddenly transferred from the PFN or PFL into the magnet.

Thyratrons are equipped with a reservoir, which can be used to adjust the sensitivity of the Thyratron tube. If the reservoir is set too high, the Thyratron may discharge spontaneously before receiving a trigger, and the system will not be ready when the beam arrives. If the reservoir is set too low, the Thyratron may not fire at all.

Triggers

Ideally, the trigger pulse to fire the Thyratron should be set just early enough so that beam arrives immediately after the current has reached full value. The timing must be accurate to within a few nanoseconds. The Trigger Interface Module, in the same NIM crate as the Power Supply Control Module, generates the pulse that triggers the Thyratron and fires the kicker.

In most cases, the ultimate source for timing the trigger comes from the beam sync clock. (The MIBS clock is ultimately generated by a collaboration of the RF systems and the Time Line Generator, and correlates events by counting RF buckets.) Locally, a C279 or C479 decoder card is "armed" by the beam sync transfer event, and sends the pulse to the trigger module after a predetermined delay counted in RF buckets. (In contrast, note that the timing of the unclamp event does not require a great deal of precision, as long as the kicker is fully charged by the time it fires. A delay from a TCLK event is close enough. The trigger event, on the other hand, requires the precision of the beam sync clock.)

Finally, a TCLK "reflection" event is created simultaneously with the broadcast of the beam sync event. The reflection event comes in handy for such things as beam diagnostics or sequencer steps, since hardware capable of decoding TCLK can already be found at nearly every location. For example, the Tevatron sequencer listens to reflection events to confirm that a beam transfer attempt has just taken place.

More information about beam sync clocks can be found in the Controls Rookie Book, or in the mythical Controls chapter of this book.

Lambertsons

The Lambertson magnets, according to local Fermilab definition, have a field-free aperture and an aperture with a magnetic field, in close proximity to each other. The Lambertson design shown in Fig. 7-2 is used in the P1, A1, NuMI, and abort lines, as well as injection into the Tevatron. The lattices for all four lines are similar, except that the Lambertson in the abort line is “wired” upside-down so that beam is extracted downward; the other beam lines bend beam up from the Main Injector ring.

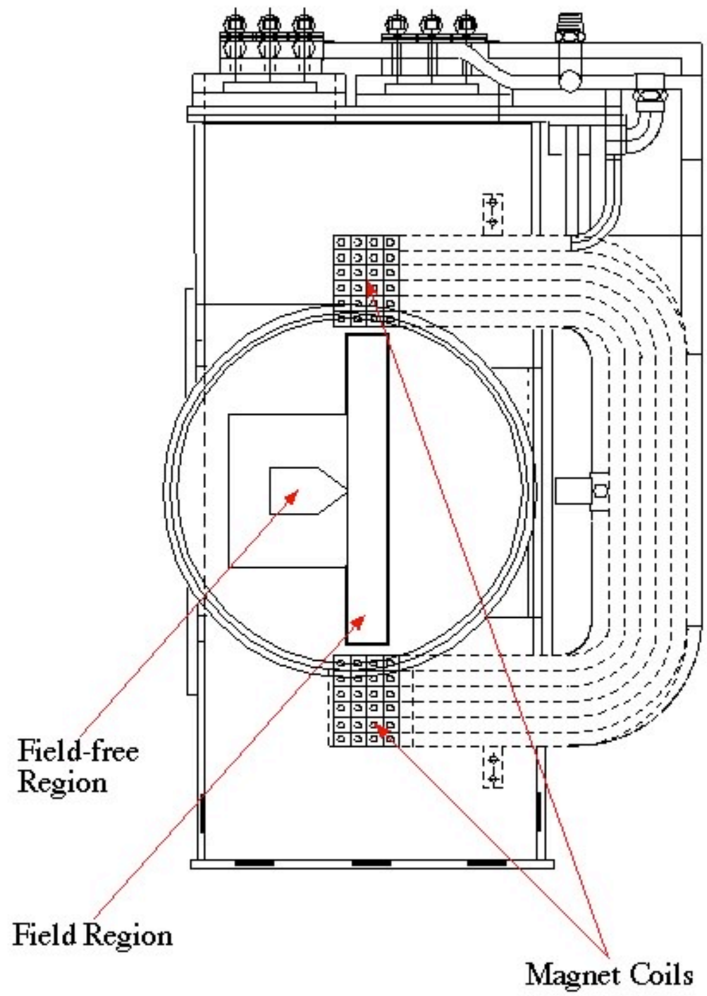

Lambertson magnets have a field region and a field-free aperture. In the beam transfer lines, circulating beam passed through the field-free region, Extracted or injected beam passes through the field region, where it is deflected. This particular design is used for the abort line, P1 line, NuMI, and the Tevatron injection.

Fig. 7-2 Lambertson Magnet

Another difference among the various lines is that the abort line and the P1 line must be able to extract beam at a variety of energies between 8 GeV and 150 GeV, while NuMI and the A1 line extract at a constant energy (120 GeV and 150 GeV, respectively). At 120 or 150 GeV, three of the relatively powerful Lambertsons are needed to produce a vertical deflection large enough (at least 3.5 inches) to clear the Main Injector beam pipe and cleanly enter the extraction pipe. Extraction kickers displace the beam horizontally into the field region of the Lambertson—but remember that in the Main Injector, the main quadrupoles are closer together than, say, in the Tevatron. This high packing ratio, or relative density of quadrupoles, helps circulating beam by keeping its size small, but it also leaves less room between them for the extraction devices. Only two of the three Lambertsons will fit between a pair of quadrupoles, the remaining one being placed upstream (Fig. 2-8 shows the basic lattice types that the extraction devices must fit into).

This arrangement works at 150 GeV, where the beam size is small and the beam is hard to bend. However, at 8 GeV, the beam has a larger cross-section and must be deliberately steered further out horizontally in order to avoid the septum of the upstream Lambertson (peek ahead to Fig. 7-22). To immediately deflect the beam vertically would guarantee that a great deal of it would be scraped off in the quadrupole, with its tighter vertical aperture. To avoid those losses, the role of the upstream Lambertson has to be minimized or eliminated at 8 GeV. The abort and P1 lines use different strategies to deal with this problem; those strategies will be described in later sections.

A BPM horizontal display frame of the ring reveals the locations of most of the Lambertsons, since the beam is bumped around the septum between the field and field-free apertures.

New and Used Magnets

The beam lines are a haven for magnet refugees arriving from decommissioned accelerators and the older beam lines. Design tolerances are often less stringent for straight-through beam as opposed to circulating beam, and, where feasible, money can be saved by recycling old components. Other magnets are new, but based on older designs.

Main Ring Large Dipoles

The large dipoles taken from the old Main Ring performed the same general function as the main dipoles in the Main Injector. The major differences between the two types are that the coils are smaller in cross-section, requiring more turns, and that the beam pipe through the Main Ring magnets is rectangular rather than elliptical.

There were three major types of large dipoles in the Main Ring: B1, B2, and B3. Most were either B1 or B2 dipoles. The B1 magnets had a larger horizontal aperture (1.5”X5”), while the B2 magnets had a relatively larger vertical aperture (2”X4”). This is because of where they were located in the lattice—there would be two B1 dipoles to either side of the horizontally focusing quadrupoles, and two B2 dipoles would be placed to either side of the defocusing quads. (Remember that beam is biggest horizontally at the focusing quads, and vice-versa.) That pattern is still intact through the P3 line, which of course is part of the Main Ring remnant. In most of the other beam lines where large dipoles are needed, only the B2 magnets are used—the magnets are often rotated, and the B2 magnets provide the best overall aperture in both planes.

The different vertical gap sizes of the B1 and B2 dipoles means that in order to maintain the same overall field strength, the B1 dipoles must have 12 turns while the B2 dipoles need 16 turns.

There were several B3 dipoles in the Main Ring, used primarily in the overpasses that bent the beam up and over the detectors at CDF and D0. The B3 magnets have an aperture of 3”X5”. Two of these industrial-strength dipoles have been incorporated into the MI-8 line; the other is at the end of the P2 line, where the AP-1 line branches off toward Pbar.

EPB Magnets

"EPB" stands for "Extracted Proton Beam." These magnets were originally used in the Fixed Target beam lines. The design of the dipoles is similar conceptually to the large dipoles described above, but smaller and shorter. In cross-section, they are 12.5" high and 16" wide and either five feet or ten feet long.

Newly built quadrupoles based on the EPB model are the "3Q120" (10 foot) and "3Q60" (5 foot) magnets; the design of the manifolds and steel laminations has been modified from the original.

Some superfluous technical information on EPB magnets can still be found in the Switchyard Rookie Book.

C-Magnets

C-magnets are one degree of separation beyond the Lambertsons. The beam pipe is just far enough away from the magnet to allow the two to be separate components. The "C" designation refers to the shape of the magnet, which surrounds the beam tube on three sides. C-magnets are found just downstream of the Lambertsons in several of the beam lines; extracted beam passes through the magnet, while circulating beam goes through the external pipe.

Correction Dipoles

Horizontal and vertical trim dipoles have been liberally interspersed among the other beam line components. Some of these correctors are recycled from the Main Ring, and others are of the newer Main Injector type.

Many of the corrector dipoles are from LEP (the now decommissioned Large Electron-Positron collider at CERN). These ugly magnets with the lead-based green paint were sold to Fermilab at a dollar apiece, and are probably overpriced. The dipole field is generated by a single coil and is shaped by two plates; beam passes between the plates. Their advantage: they are cheap and available. Their primary disadvantage: they have high inductances and cannot be used for applications that require rapid changes in the field strength. They are used primarily in the MI-8 line and A1 line.

More Quadrupoles

The SQ quads were originally built for the Antiproton Source and the old 8 GeV line. A number of those magnets from the 8 GeV line have now been recycled for use in the MI-8 line.

Some of the 84" quads from the Main Ring have been relocated to various Main Injector beam lines.

The MI-8 Line

The MI-8 line (Fig. 7-3) transports 8 GeV protons from the Booster to the Main Injector ring. The path is somewhat convoluted for two reasons.

Fig. 7-3
MI-8 Line

Vertically, the Main Injector ring is about 11 feet below the level of the Booster ring (Fig. 7-4). Horizontally, the line must avoid the pre-existing Antiproton Source, although it must still pass under the Transport Enclosure.

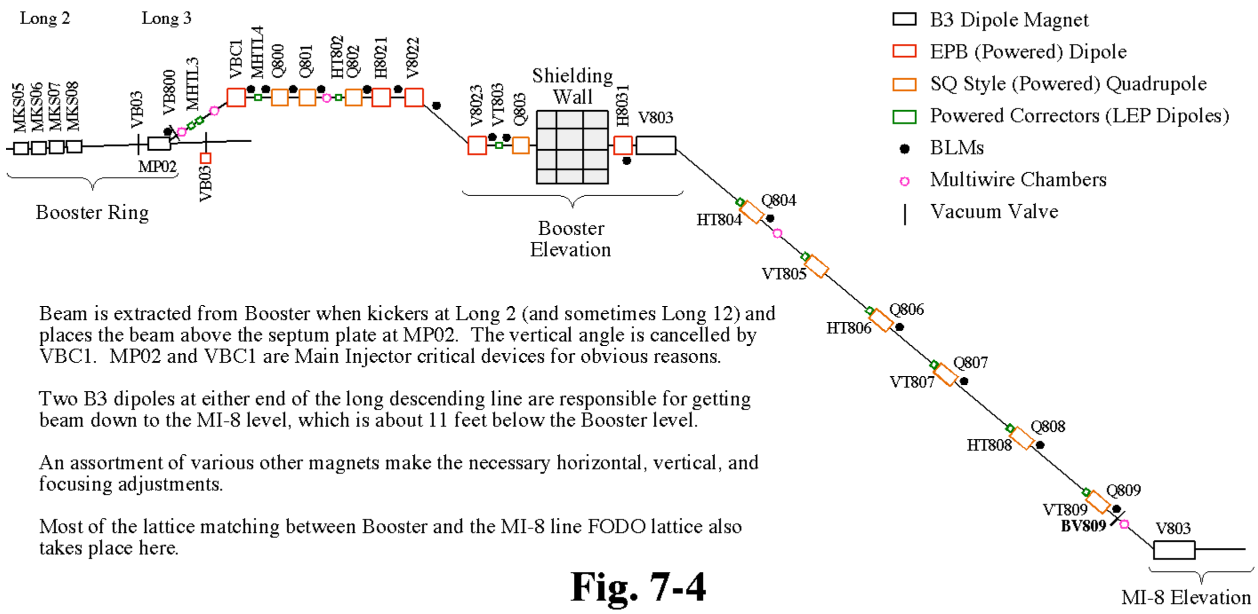

Fig. 7-4
Vertical Profile of Upstream of MI-8 Line

Most of the MI-8 line is composed of permanent magnets. Originally, the line was designed to use powered magnets, but it became a testing ground for the permanent magnet concept when the Recycler was in the earliest stages of development. Large powered magnets are still used at the beginning and end of the line, and LEP dipole correctors are used throughout the line.

The details of extraction from the Booster are included in the Booster Rookie Book. Essentially, in its simplest form, the four kickers (MKS05, MKS06, MKS07, and MKS08) at Long 2 give the beam an upward kick so that extracted beam passes over the septum plate of MP02 at Long 3. The beam leaves Booster at a tangent to the Booster ring horizontally and at an upward angle vertically. There are kickers at Long 12 as well; their purpose is to send beam to the Long 13 dump or to the Radiation Damage Facility (RDF). To create partial batches, used for coalesced beam in Collider mode, a few bunches are sent to the Main Injector and the rest are sent to the Long 13 dump.

Numbering in the line begins with “800” at the point at which the beam leaves the Booster, and ends at “852,” where it enters the Main Injector. In its entirety (Fig. 7-3), the MI-8 line looks rather complicated, but in reality there are only five different kinds of lattice sections to worry about. The stretch through the middle is your basic FODO lattice, and the two at either end are the matching lattices. Since the Booster, MI-8 line, and Main Injector ring each have a unique lattice, the transition from one to the next has to be done in a controlled fashion. There is a Booster matching lattice at the beginning of the line and a Main Injector matching lattice at the end. Background information on lattice functions can be found in Chapter 2.

Most of the FODO lattice in the MI-8 line consists of two alternating lattice types—the beam is gradually being bent to the right through this region—but just after the Booster matching section is a short length of reverse bending, in which beam is bent to the left so that the line is steered away from the Pbar rings. In order of appearance:

The Booster Matching Lattice

The region between 800 and 809 (Fig. 7-4) is dedicated to making the transition between the Booster lattice and the MI-8 FODO lattice, but there are also important things happening vertically and horizontally.

Although most of the MI-8 line consists of permanent magnets, the magnets in the Booster matching section are powered. The quadrupoles are SQ-style electromagnets similar to those found in the Pbar beam lines; there is one quadrupole at each numbered location, and the even/focusing, odd/defocusing convention is applied here as in the Main Injector ring. The purpose of the matching quadrupoles is to convert the lattice functions (dispersion as well as the amplitude and phase of the beta functions) in Booster to that of the FODO lattice of the MI-8 line. Within the matching section, the amplitude of the beta function varies from cell to cell until the FODO lattice is reached.

Within this section of the line, the vertical path dominates the scene. The vertical angle out of the Booster from MP02 is cancelled by a 5-foot EPB magnet known as VBC1. The short, level section following VBC1 contains three of the matching lattice quads (Q800, Q801, and Q802), and a 10-foot EPB (H8021) for steering the beam to the left of the AP-4 dump. A dogleg (the picture explains the term best) consisting of the two EPB magnets V8022 and V8023 brings the beam back down to the original elevation it had in the Booster.

(At this point the beam pipe passes through a shielding wall. Why is there a shielding wall in the middle of an enclosure? Once upon a time, when antiprotons were much more difficult to produce, the Antiproton Source was tuned by sending protons, from Booster, into the rings. The line connecting Booster and the Debuncher was known as the AP-4 line. Later, when Collider operations required that partial batches be sent to the Main Ring, the upstream end of the AP-4 line was reworked to send the excess beam to a beam absorber (called a dump in those less sensitive days). The shielded area around the absorber became known as the AP-4 dump. Now that the excess beam is being sent to the Long 13 dump, there is a big pile of concrete blocks surrounding a section of the MI-8 line for no apparent reason.)

After the shielding wall, a B3 dipole recycled from the Main Ring starts the beam down the long slope to the Main Injector elevation 11 feet below. At the end of the slope, another B3 dipole straightens the beam out at the Main Injector level. The two B3 dipoles, which are in series, are collectively known as V803. Although the MI-8 line continues on for a considerable distance, V803 is the last major vertical bend in the line. However, at the bottom of the slide the beam is still 35 mm higher than circulating beam in the Main Injector—a fact that will be explained in due time.

The first horizontal priority in the MI-8 line is to prevent the beam from hitting the enclosure wall and to steer it to the left of the shielding wall. That task is performed by a 10-foot EPB named H8021. After the shielding wall, another 10-foot EPB, H8031, gives the beam another push to the left to position it correctly for the trip down the slope. H8031 is the last major horizontally bending powered magnet unique to the MI-8 line. The bulk of the horizontal bending from that point forward is done by permanent magnets.

LEP correctors precede most of the powered quadrupoles. Since beam in the MI-8 line is always at the same energy, the large inductance of the magnets doesn't matter. The practice of including one powered corrector at every half-cell continues down the rest of the line.

Reverse Bending Section

Background information about permanent magnets can be found in Chapter 2.

At 809, the beam has reached the bottom of the slope and, vertically, can virtually coast into the Main Injector. Horizontally, the beam will soon be bent to the right until its trajectory is parallel to that of the Main Injector at MI-10. However, the dispersion match is not quite complete, and besides, if the rightward bend is started now the beam will intersect the domain of the Antiproton Rings.

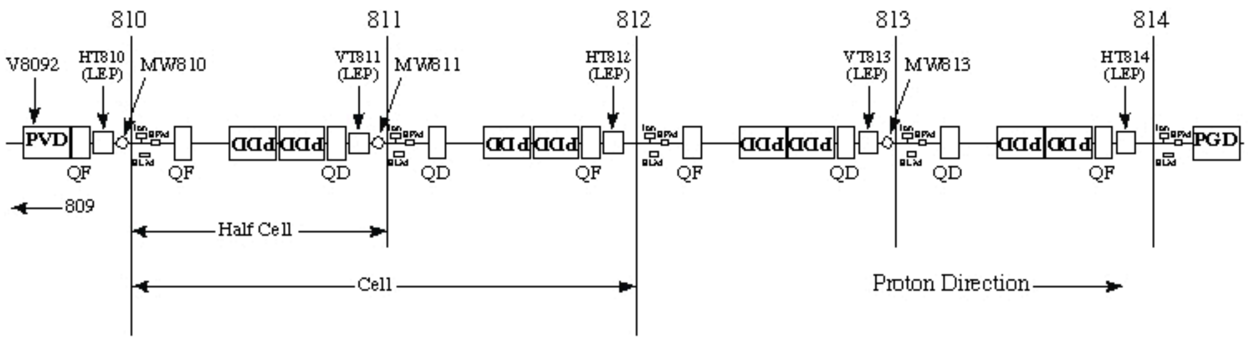

The reverse bending section begins at the bottom of the vertical slope at 809. The permanent dipoles have been turned upside-down so that they bend to the left instead of the right. Focusing is done with short permanent quadrupoles. The three trim magnets, from LEP, are powered. The section ends at 814.

The purpose of this section is to steer cleat of the Antiproton source rings (see Fig. 7-4).

Fig. 7-5
Reverse Bending Section, MI-8 Line

The distance from 809 to 814 (see Figs. 7-3, 7-5) has been designated as the reverse bending section; beam is bent to the left in order to steer clear of the rings and the surrounding roads. The reverse bending is done by turning the PDD dipole magnets “upside-down;” because most of the MI-8 line bends to the right, the convention—and the labeling—is that the right-bending configuration is “right-side-up.” Right side up in this case means that the magnetic field is pointing up as well (Fig. 2-16).

Focusing in the reverse-bending section is done by two-foot long permanent quadrupoles (Fig. 2-17). The entire reverse-bending section is only two cells long. The focusing and defocusing quadrupoles are identical in construction, but the defocusing quads have been rotated end-to-end with respect to the focusing quads (i.e. they are “backwards”). The section begins with a focusing quad at 810 and ends with a focusing quad at 814. By the end of the reverse bending section, the dispersion match is complete and the trajectory is set up for the long run to the Main Injector.

The MI-8 FODO Lattice

The long stretch from 814 through 846 consists of two types of cells: the full-arc cells (Figs. 7-3, 7-6) and the dispersion-suppressor cells (Figs. 7-3, 7-7). Beam is consistently bent to the right in both types of cells, as the beam is

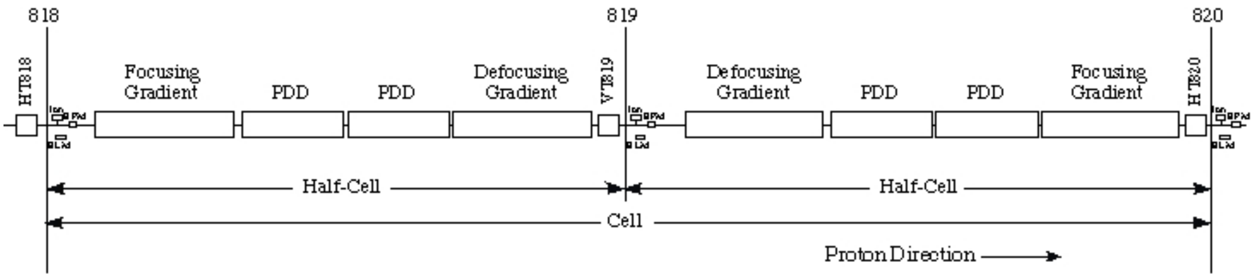

The full arc-cells begin at 818, 820, 830, 832, 834, 844 and 846. The gradient magnets bend and focus the beam.

Fig. 7-6
Full Arc Cell

brought to an angle nearly tangent to the Main Injector at MI-10.

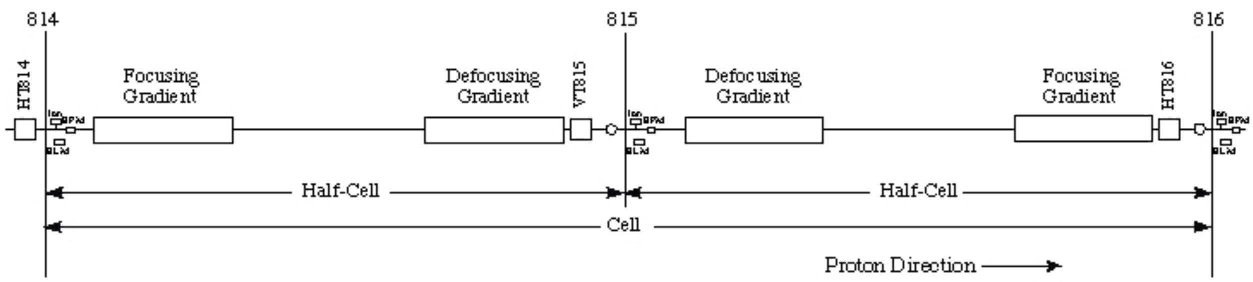

The dispersion suppressor cells begin at 814, 816, 822, 824, 826, 828, 836, 838, 840, and 842.

○ - Multiwire
IP - Ion Pump
BLM - Beam Loss Monitor
BPM - Beam Position Monitor

Fig. 7-7
"Missing Dipole" Dispersion-Suppressor Cell

The permanent magnets in the dispersion-suppressor cells are all PGD gradient dipoles, while the "full-arc" cells are a combination of PDD "double dipoles" and PGD gradient dipoles.

The PDD magnets are shorter than the PGD magnets, but since the ferrite bricks are stacked two layers deep in the PDD magnets, the two types have the same bending power. Yet another look at Fig. 7-3 shows that the curvature is greater in the full-arc regions. Since this part of the line is populated primarily by permanent magnets, the only control of the beam in this region is provided by the LEP dipole correctors at each half-cell boundary.

The Main Injector Matching Lattice

Location 847, about a hundred meters from the finish line, marks the beginning of the transition to the Main Injector lattice. Powered quadrupoles once again appear at each numbered location, and the oscillations of the beta functions once again change in amplitude. The PGD gradient magnets disappear as the powered SQ-style quads take over the task of focusing. Q852 is the final magnet belonging to the MI-8 line, and the lattice has now been matched to that of the Main Injector (again look at Fig. 7-3).

As of this writing, MiniBooNE magnets have been added to this last stretch of the line. The MiniBooNE line branches out of the Main Injector tunnel before encountering the ring. MiniBooNE extraction will be covered in a future section of this chapter.

Injection into the Injector

Protons at the end of the MI-8 line (Fig. 7-8) are necessarily approaching the Main Injector at a horizontal angle of about 35 milliradians (mrad); it will be seen momentarily that this angle is defined by the bending strength of the Lambertson magnet.

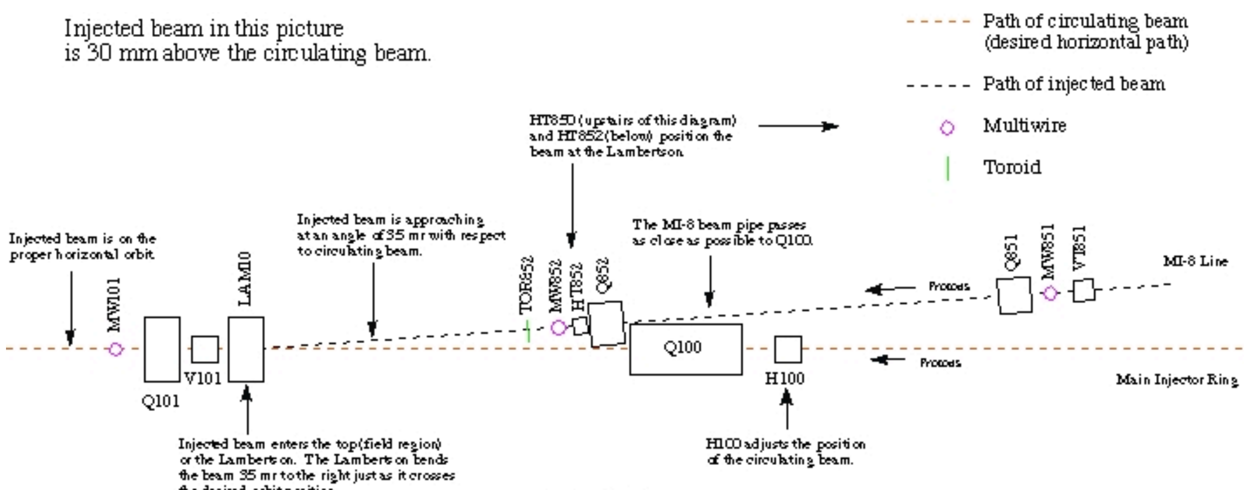

Fig. 7-8
Horizontal Closure in Main Injector
(Top View)

The final bulky component in the MI-8 line, Q852, is placed just downstream of Q100 (in the Main Injector ring) so that the beam line can be set at the 35 mrad angle. Remember from earlier in the line that beam has deliberately been placed about 35 mm above the desired vertical center for circulating beam.

The process of placing injected beam on the proper orbit in an accelerator is called closure. When the orbit is closed, the beam is at the correct position, and angle of trajectory, in both planes. In the Main Injector, the horizontal orbit is closed with a Lambertson magnet, and the vertical orbit is closed with kickers.

The injection Lambertson (LAM10) is located just upstream of Q101, a defocusing quadrupole (V101 is between the two, but doesn't take up much space). The horizontally bending Lambertson has a field region on top and a field-free aperture on the bottom; it was originally built as a spare for the 8 GeV line in the Main Ring. Injected beam was designed to be 35 mm higher than circulating beam and will pass through the top aperture (Fig. 7-9). Beam may already be circulating in the Main Injector, and, if so, will pass through the field-free aperture on the bottom.

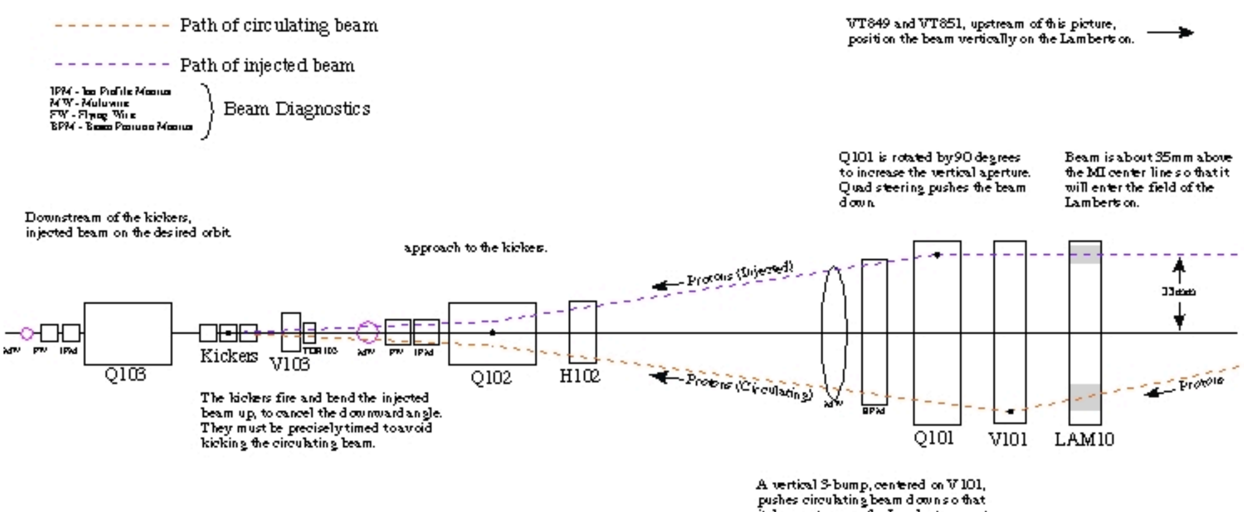

Fig 7-9
Vertical Closure in Main Injector
(Side View)
(Vertical scale grossly exaggerated)

The injected beam in the Lambertson is bent 35 mrad to the right, just enough to offset the angle coming in from the MI-8 line. Since the Lambertson is located at the point where injected beam crosses the closed orbit horizontally, the position and angle are now correct and injected beam is closed in the horizontal plane.

This scheme will only work if the beam actually enters the Lambertson magnet with the correct position and angle of approach. HT850 and HT852, two corrector dipoles in the MI-8 line, are both used to ensure that those two conditions are met.

To close the beam vertically (Fig. 7-9), the beam is deflected onto a downward slope, and the kickers are placed where the beam will cross the desired vertical position. Because of the large vertical separation between the injected beam and circulating beam required at the Lambertson, the kickers are placed as close as possible to where the vertical betatron phase advance is 90˚—the "natural" oscillation of the beam will do much to bring it to the desired vertical position.

Specifically, remember from Chapter 2 that in the Main Injector there is an approximate 90˚ phase advance over the distance between two adjacent "F" quads or two adjacent "D" quads. Beam coming out of the Lambertson is vertically high at Q101 but will have been pushed down by the time it reaches the center of the beampipe at Q103, the next defocusing quad. The kickers are therefore placed as close to Q103 as possible. It is a little more complicated than that, because quad steering—the dipole bending effect resulting from the beam being off-center at Q101—pushes the beam down faster than desired. The correction dipole VT849, in the MI-8 line, can be used to adjust the angle at the Lambertson so that the position is correct at the kickers. The Lambertson is also rolled (rotated) so there will be a slight upward angle imparted to the injected beam.

The orbit of the circulating beam must also be modified, because at 8 GeV the beam is large enough vertically to scrape against the septum plate of the Lambertson. A 3-bump using V641, V101, and V103 pushes the beam down, with the low point being at V101. (A vertical BPM display frame clearly shows this bump.) The Lambertson is close enough to V101 to make it less likely that circulating beam will strike the septum plate

However, because injected beam is high and circulating beam is low, the vertical separation of the injected and circulating beam at Q101 is too large to be accommodated by the aperture of a main quadrupole as normally configured. Q101 has been rolled 90˚ so that the long transverse axis of the aperture is oriented vertically instead of horizontally. Moreover, although most of the other 84" quads from the Main Ring were upgraded with an elliptical beam pipe, Q101 and Q102 were left with the original "star" aperture—giving the beam extra room. (The cross-section of a main quadrupole in Fig. 2-6 shows the elliptical beam pipe surrounded by the four-lobed star aperture. With the elliptical beam pipe removed and the magnet turned on its side, as with Q101, there is much more vertical aperture.) The bus connections are such that it is still a defocusing quad. The two beam trajectories are closer together at Q102, and the beam is wider horizontally; therefore the star aperture is used at Q102 but the magnet has not been rolled.

When the Lambertson closes the beam horizontally, there is the luxury of the injected and circulating beam being in two distinct places—but all of the beam passes through the kickers. The only thing separating injected beam from circulating beam is time, and not very much of that. Timing constraints are very stringent. To kick the injected beam and not the circulating beam, the kickers must have as fast a rise time, and fall time, as possible. The flattop for the kicker waveform is 1.6 microseconds long, the length of a Booster batch.

Fig. 7-10, on the next page, illustrates some of the timing and control for the injection kickers. The trigger for the kickers comes from the Booster Extraction Sync (BES), which is itself derived from TCLK. BES, which predates the beam sync clocks, is ultimately referenced to Booster $10 or $12 (beam or pre-pulse) events. A parameter called B:MIEXTR counts down the approximate time before beam is to be extracted. As of this writing, B:MIEXTR has been given a value of 35,450 microseconds. (Remember that although acceleration time in Booster is about 33.333 milliseconds, the Linac and Booster devices require about 2 milliseconds to prepare for beam after the reset has been issued. The exact value of the delay is then determined empirically, after much anguish.) A NIM module in the Booster LLRF room adds the delay to the clock event and then issues the BES.

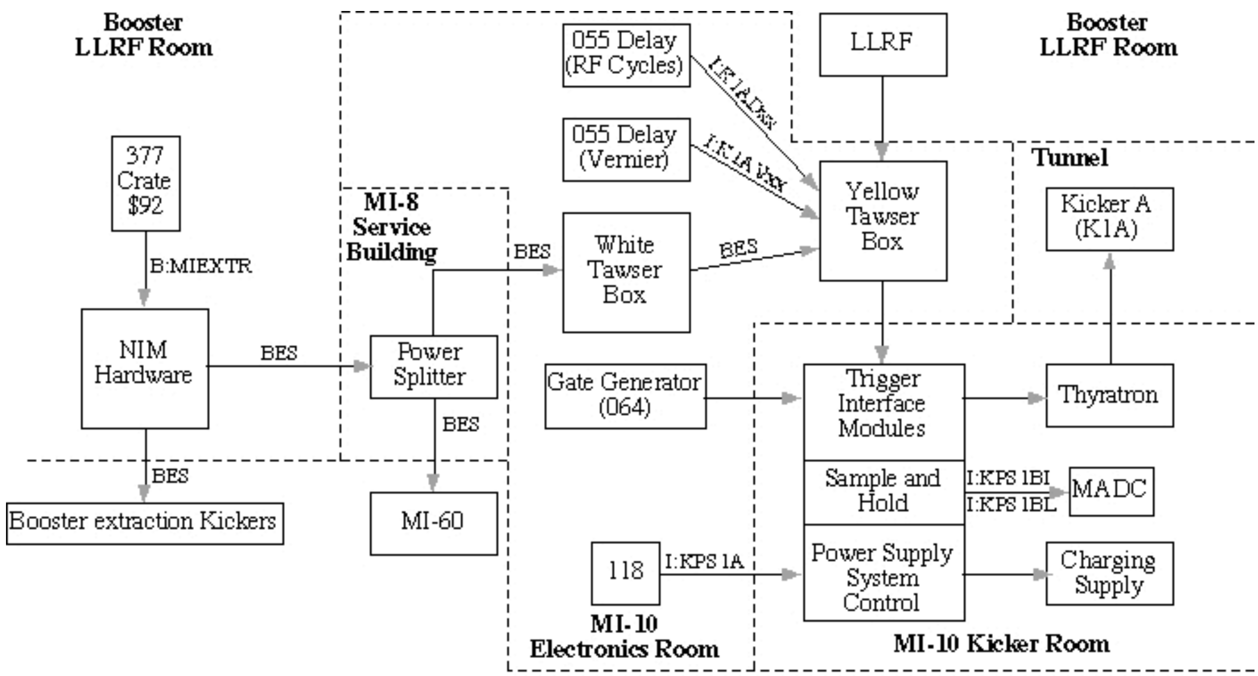

Fig. 7-10, Injection Kicker Timing and Control

The injection kickers in the Main Injector get their timing from the Booster Extraction Sync (BES), which predates the other beam sync clocks. BES is initiated by Booster beam and pre-pulse events ($10s and $12s) and generated by NIM hardware in the Booster LLRF room. The BES signal is fanned out to the Booster extraction kickers, MI-60 and MI-10. The white Tawser box in the electronics room at MI-10 distributes the signal to three yellow Tawser boxer, one of which is shown in this diagram. The yellow Tawser boxes add the RFC and vernier delays to the BES time, using the LLRF as a reference. Finally, the combined signal tells the trigger interface modules exactly when to fire the Thyratron.

A CAMAC 118 card sets the DC level for the kicker voltage. A DC level is sufficient because the beam is always at 8 GeV.

For simplicity, this diagram only shows timing and control for the first kicker, K1A.

BES is used to trigger the Booster extraction kickers as well as the injection kickers at the end of the MI-8 line. A common source for the trigger increases the likelihood that the Main Injector kickers will actually have current in them when the beam passes through. BES is fanned out to the electronics room at MI-10, where it becomes an input to the first Tawser box. The Tawser boxes are named after Stan Tawser, of Tawser box fame.

There is a “White” Tawser box in the electronics room that fans BES out to the three “Yellow” boxes—one for each kicker. The yellow boxes combine the BES, LLRF, and two delays (measured in RF cycles) from nearby 055 cards to create a trigger timing pulse. The pulse is shipped next door to the NIM hardware in the kicker room that triggers the Thyratron.

The delays for the triggers can be found from B4, the Booster Beam Turns Control page (kicker times to LX), and on B19, the parameter page dedicated to Booster extraction. The main type of delay is measured in RF cycles (RFCs); the other, shorter delay is the vernier and is measured in nanoseconds. Both delays are set from 055 cards in Crate $10. It is sometimes difficult to make sense of the numbers, since there is a lot of fine tuning done to compensate for cable propagation delays and kicker rise times.

The injection kicker power supplies reside in the kicker room at MI-10. The pulse is shaped by a PFL, the cables being wrapped around large spools.

MI-8 Power Supplies

The larger powered magnets in the MI-8 line are found at the beginning of the line (800 to 809), and at the end (847 to 852). There are also the small powered LEP dipole correctors at each half-cell. The power supplies for the upstream segment of the line are found in the Booster West Gallery and the Booster West Tower. The corrector dipole supplies for the remainder of the line are located in the MI-8 Service Building. Figs. 7-11 and 7-12 are intended to serve as a complete list of power supplies for the MI-8 line, excluding devices which are actually part of the ring proper. (The magnets themselves have names that may or may not be similar to the ACNET names.) The small box at the upper left corner of each power supply name in Fig. 7-11 (next page) represents the type of CAMAC card controlling the supply; to the right of the box is the type of power supply.

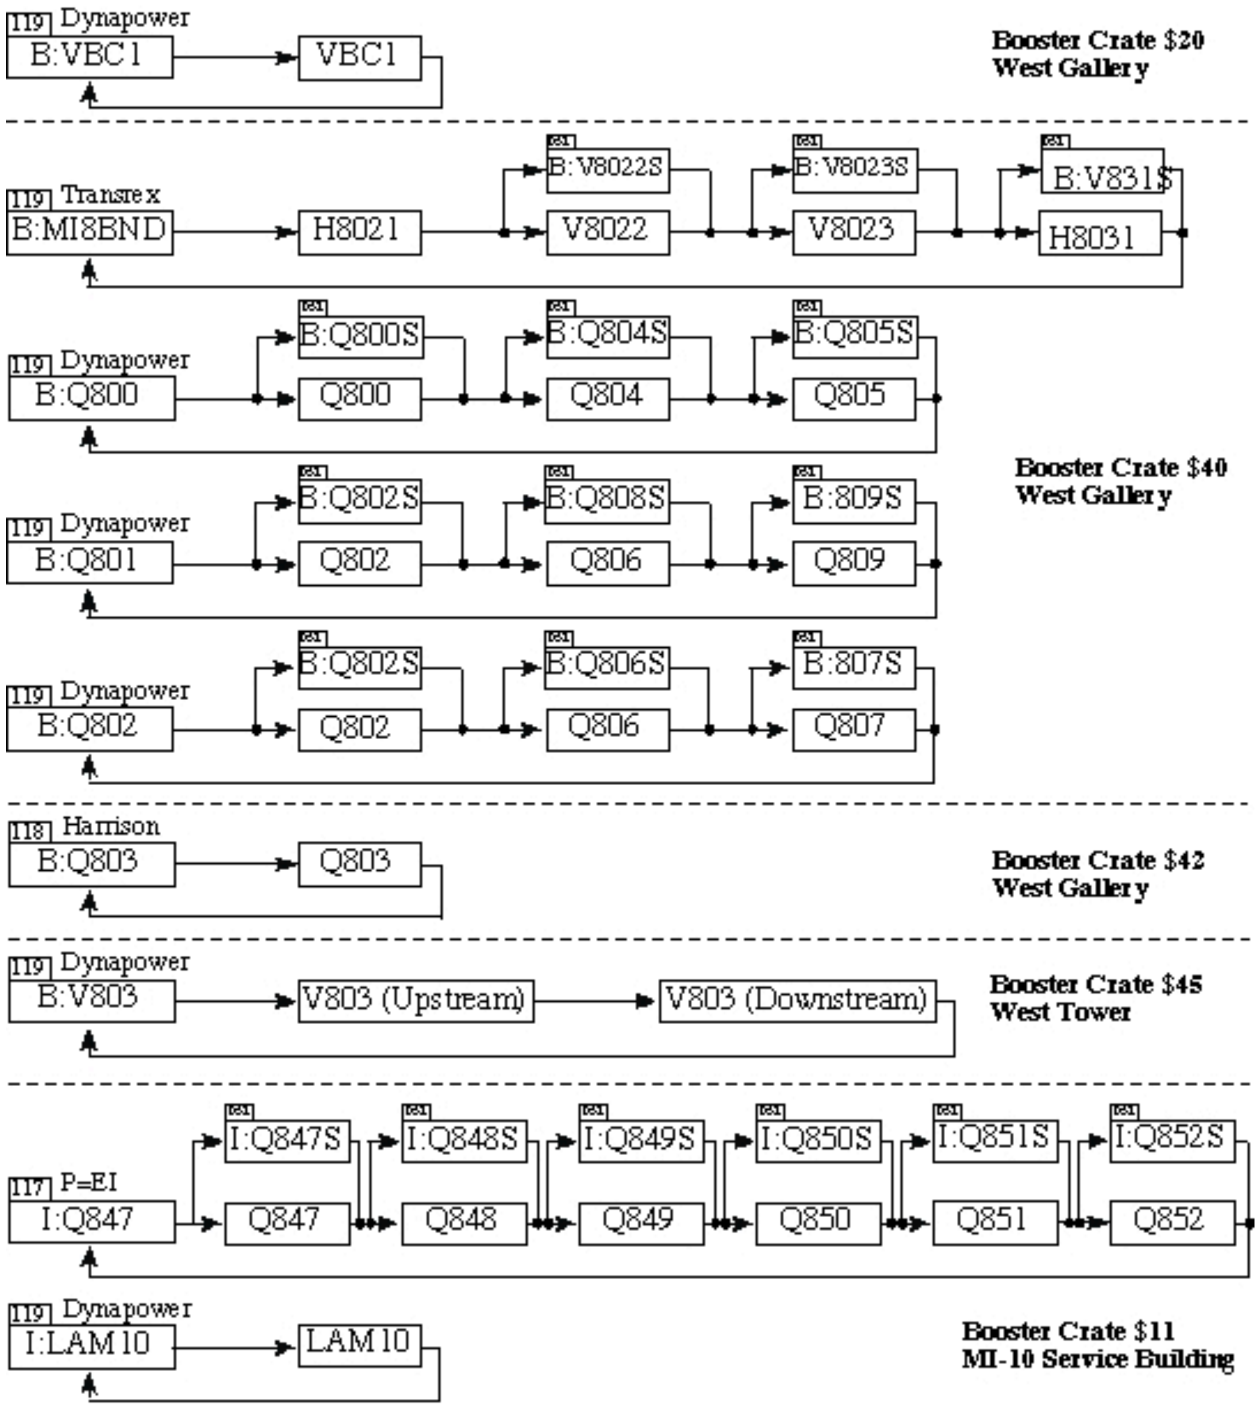

Fig. 7-11 MI-8 Line Power Supplies

Several of the supplies for devices in the MI-8 Line power strings of magnets. Power to individual magnets in a string is customized by diverting some of the current through shunts. The "S" at the end of a parameter designates a shunt.

The number at the upper left of the power supply name represents the CAMAC control card.

Power supplies for the LEP dipole correctors in the MI-s line are listed in Figure 7-12.

VBC1, the first EPB dipole of the line (and a critical device), uses a dedicated Dynapower supply located in the Booster West Gallery. V803, which powers the B3 dipoles at either end of the vertical drop to the Main Injector elevation, uses a Dynapower supply located in the Booster West Tower. These two devices are controlled through CAMAC 118 or 119 cards—cards in the 117,118, and 119 series are capable of sending a power supply an analog reference voltage as well as providing digital control and digital readbacks. Analog readbacks return through MADCs.

Other supplies make extensive use of power supply shunts, where a single power supply drives several devices in series. All of the devices in the series would carry the same current, but some of the current destined for any individual magnet can be bypassed, or shunted, around the magnet, thereby reducing its bending strength. The power supply in each case is controlled through CAMAC 119 cards, while the shunt current—an analog signal—is controlled and read back through 052 cards. ACNET names for the shunts consist of the magnet name followed by an "S;" for example, I:Q804S shunts current around the magnet Q804, which is in series with other magnets that may need to run at different currents. The current in the magnets themselves is calculated by subtracting the shunt current from the primary current; the calculation is done by a disembodied "open-access" front end called BOOSTR that runs on CFSS (one of the central VAXs). The parameter name for the calculated current ends with "I," e.g. "Q804I."

In the upstream MI-8 line, four power supplies use shunts. B:MI8BND powers four EPB dipoles: H8021, V8022, V8023, and H8031; the last three of these have shunts. There are three quadrupole supplies in the Booster matching lattice, each supply powering three magnets: B:Q800 powers Q800, Q804, and Q805; B:Q801 powers Q801, Q808, and Q809; and B:Q802 powers Q806 and Q807.

The LEP magnets, distributed throughout the line, use CPS bulk supplies and regulators identical to those used in the Main Injector ring (Fig. 7-12). I:CPSBW, in the Booster West Tower, powers I:HT800D through

I:HT812D. Although 453 cards control the regulators, the correctors run DC—the ramp tables are not populated, and even if they were, the ramps are disabled. The "D" at the end of the parameter name stands for "DC."

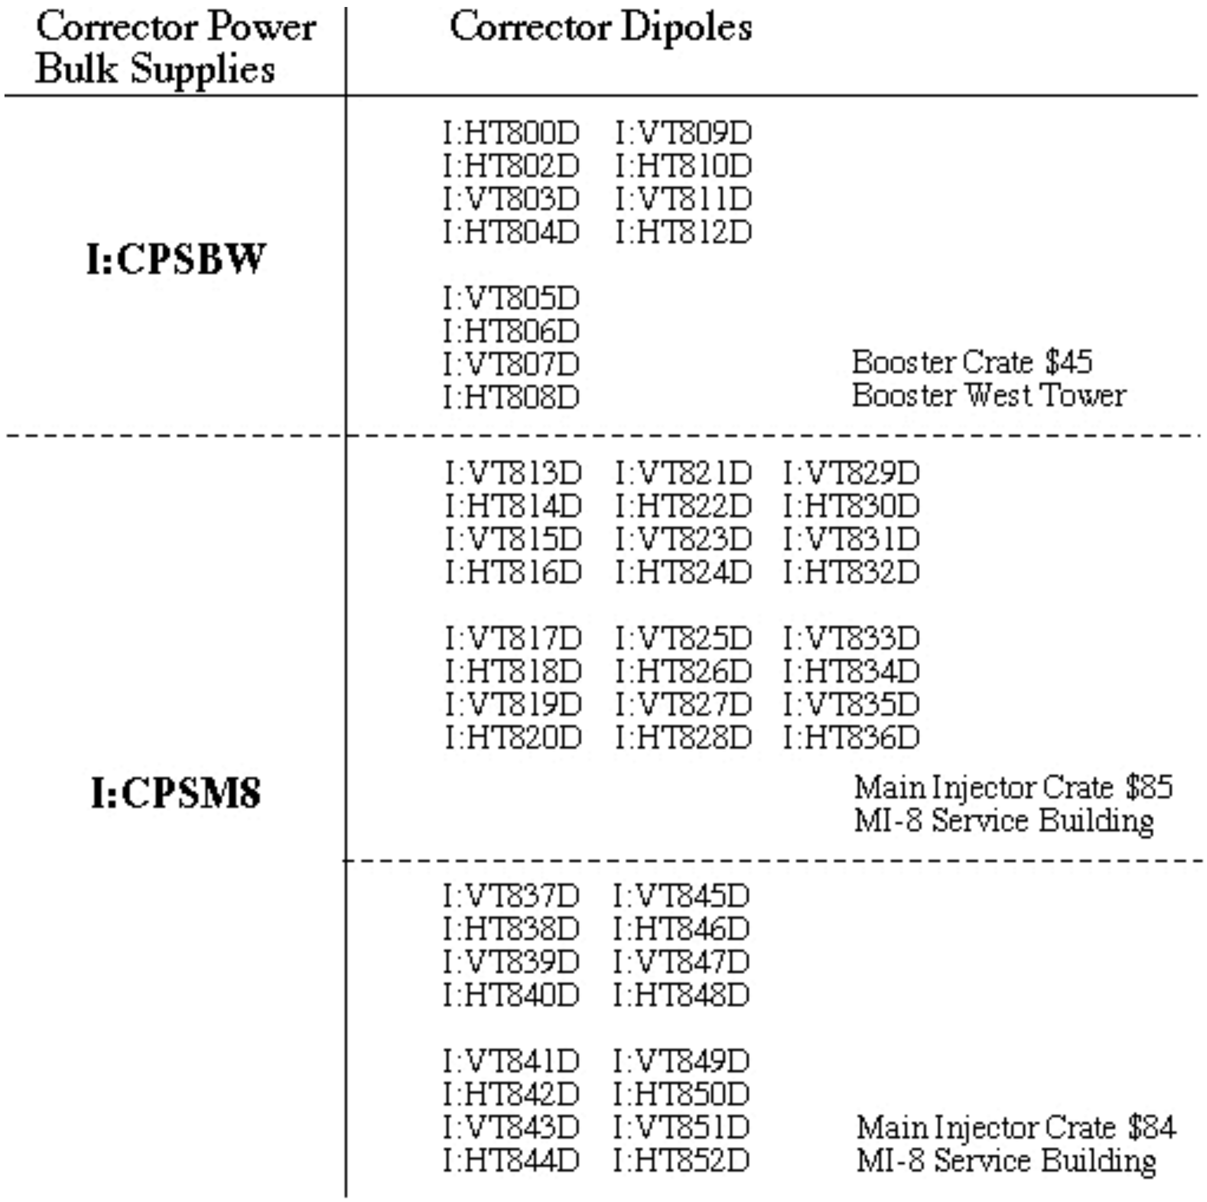

Fig. 7-12 MI-8 Line Dipole Corrector Supplies

All correctors in the MI-8 line are from LEP

The "D" at the end of each parameter name designates a DC value.
Each cluster of four devices represents an individual 453 card.

Compare to figures 3-7 and 7-3.

All of the other correctors, from I:VT813D to I:HT852D, are powered by I:CPSM8 in the MI-8 Service Building.

The quadrupoles of the Main Injector matching lattice, from Q847 to Q852, are all powered in series from a single P=EI supply, I:Q847, located in the MI-10 Service Building. There are shunts for all of the individual magnets.

The Lambertson magnet, LAM10, is powered from a Dynapower supply located in the MI-10 Service Building. The kicker equipment is located in a separate room at MI-10.

The correctors V641, V101, and V103, which create a three-bump at the Lambertson, are part of the Main Injector ring and are powered from I:CPS10.

MI-8 Vacuum

Background information on vacuum systems can be found in Chapter 5.

The MI-8 line can be thought of as being divided into four vacuum sectors (see Fig. 7-3). The first sector stretches from the Booster ring to 809. There are actually three valves that can be used to isolate Booster vacuum from the first sector (Fig. 7-4). MP02 is bracketed by two valves collectively known as VB-03. The two valves are designed to open and close in concert; if MP02 is isolated from the rest of Booster, the MI-8 line is isolated from Booster as well. The third valve, VB-800, is part of the MI-8 line immediately downstream of MP02. Obviously, in order to isolate the MI-8 line from Booster while still allowing circulating beam, VB-03 must be open and VB-800 must be closed. VB-800 can assume a secret identity—on Booster applications pages, it is known as VBMI8.

The next beam valve is BV809, just upstream of V803 at the bottom of the slope.

The MI-8 FODO lattice stretches from 809 to 847. It is broken into two sectors, with a beam valve at 828 (Fig. 7-3). The next beam valve is near the downstream end of 846; it isolates the lattice matching section from the FODO lattice.

The final beam valve in the MI-8 line is between Q852 and the Lambertson.

As in the Main Injector ring, vacuum in the MI-8 line is maintained with ion pumps. The power supplies for the pumps are located in the MI-8 Service Building.

Vacuum in the MI-8 line is ideally in the 10^{-8} range, but tends to be slightly worse in the regions near the numerous multiwires. The multiwires use a G-10 (fiberglass) support; the mesh of microscopically thin fibers tends to trap air, which can then continue to outgas for several months following a pump down.

MI-8 Beam Diagnostics

Background information on beam diagnostics can be found in Chapter 6.

The MI-8 line is, of course, the point where beam first enters the Main Injector. The precision with which the beam is injected has a significant effect on its quality later on, so the MI-8 line is crammed with diagnostics—BPMs, BLMs, multiwires, flying wires, toroids, and ion profile monitors.

BLMs and BPMs in the MI-8 line (Fig. 7-13 on the next page) are generally located at every half-cell boundary, although the stretch from MP02 through 803 is more densely instrumented in order to ensure efficient extraction from the Booster. The locations of many of the BPMs and BLMs can be found in Figs. 7-4 through 7-7.

The BLM hardware in the MI-8 Line, unlike its counterpart in the Main Injector ring, does not use microprocessors to organize the data—cards inside the BLM chassis integrate the losses from each ion chamber and then send the value directly to an MADC. The timing to set the interval of integration comes from nearby 377 cards. CAMAC 190 or 290 cards transmit the MADC values back to the MCR. BLM hardware is located in the Booster West Gallery (B:LMVBCU), the Booster West Tower (B:LM800 through I:LM822), and the MI-8 Service Building (I:LM823 through I:LM852). The BLMs in circulating beam

(I:LM100, etc.), being part of the Main Injector ring, are handled through the MBP microprocessors.

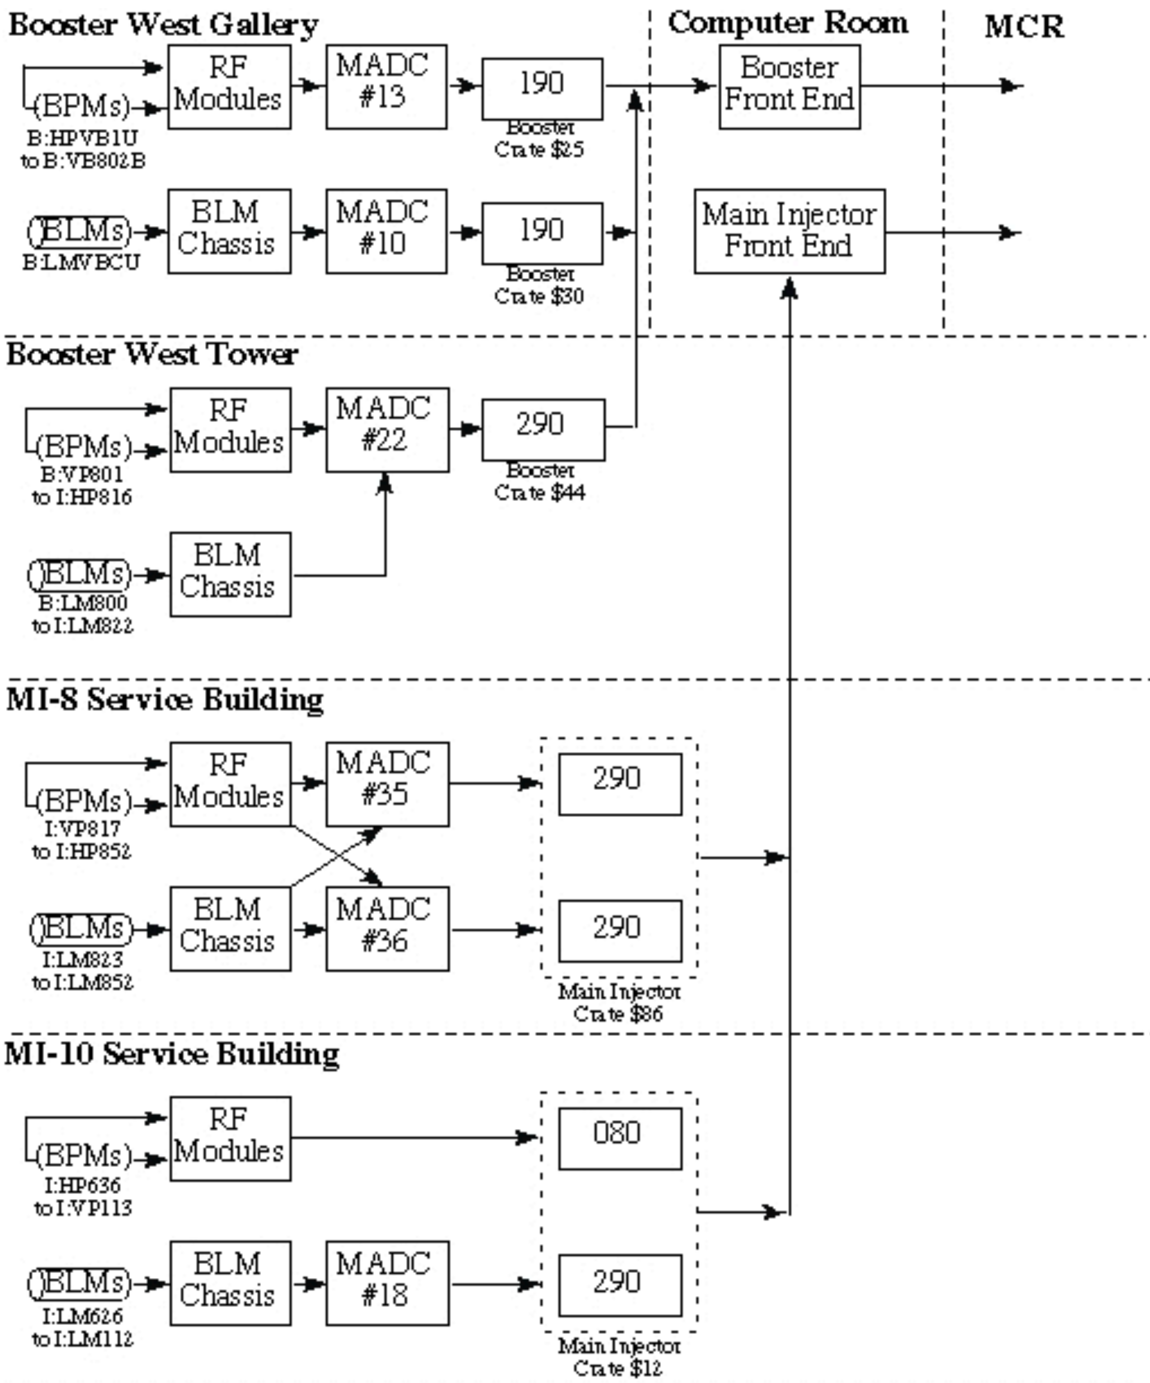

Fig. 7-13 MI-8 BPM and BLM Readbacks

The BPMs, like the BLMs do not use the MBP microprocessors. However, calculations of beam position and intensity are done on board the individual BPM modules and passed on to the MADCs. The ACNET parameter names for the calculated values are HPxxx and VPxxx for the beam positions, and HIxxx and VIxxx for the beam intensities (xxx representing the location). The BPM applications page I39 assembles those readbacks into a coherent picture.

The reason that the MBP microprocessors are not used in the beam line is that there is no need for the averaging of orbit positions for display frames, nor is there a need to collect the multiple snapshot frames necessary for constructing profile buffers. One flash frame is sufficient for a single pass.

The two toroids, TOR800 and TOR852, are obviously located at the endpoints of the line. Readbacks are through the MADCs. The TOR800 electronics is located in the Booster West Gallery, and the TOR852 electronics is at the MI-10 Service Building.

The MI-8 multiwires (Fig. 7-14 located on the next page) are distributed over a large distance, and consequently require the services of several different segments of the controls system.

Multiwires 800 through 830 communicate with the controls system through CAMAC 184 and 192 cards. MW800, 802, and 804 report to Booster Crate $31 in the Booster West Gallery; MW810, 811, 813, and 814 talk to Booster Crate $45 in the Booster West Tower; and MW825, 826, and 830 are connected through Main Injector Crate $86 in the MI-8 Service Building.

Multiwires 836, 839, 840, 851, 101, 102, 103, and 104 are operated through SWIC controllers similar to those found in Switchyard. 836 and 839 are found in the MI-8 Service Building, and the rest are located in the MI-10 Service Building. There is an ARCNET loop at each house that ties the SWIC controllers to the local VME crate, so that the data can be sent back via Ethernet toward the MCR.

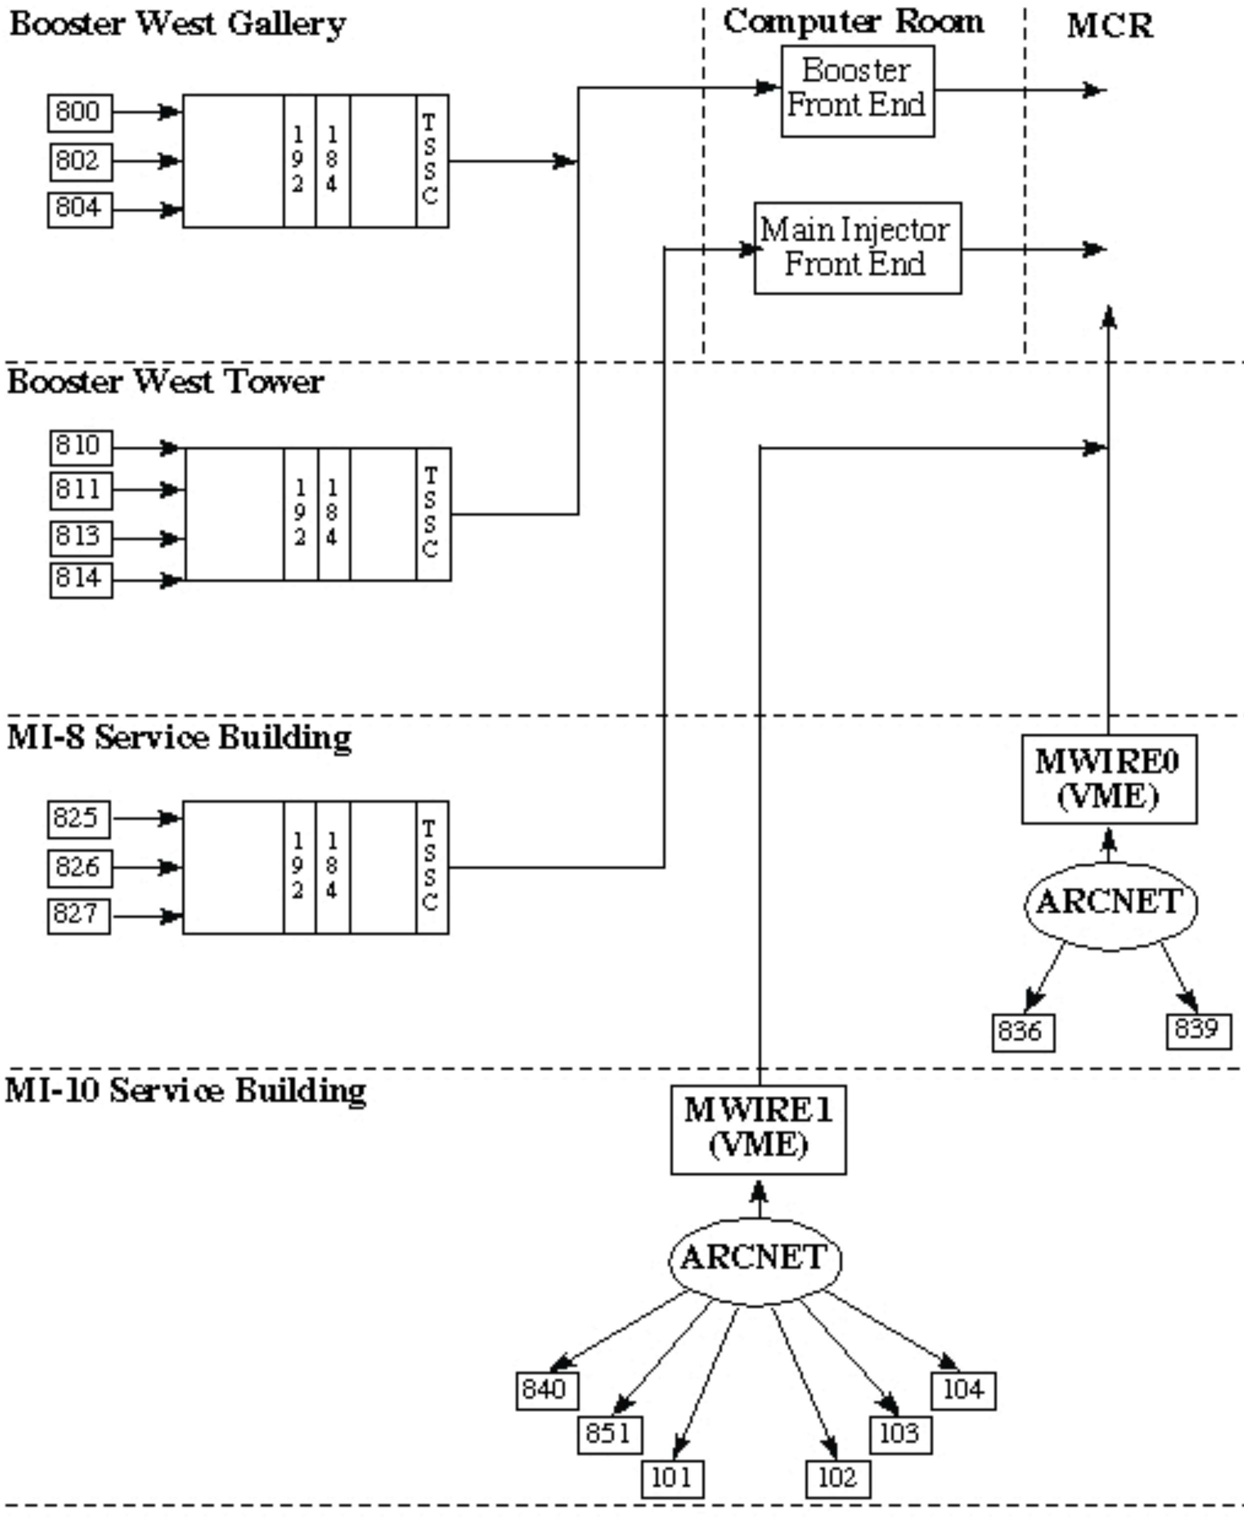

Fig. 7-14 MI-8 Multiwire

Multiwires 101 through 104 must only be used when beam is aborted on the first turn, since they are in the path of circulating beam.

The horizontal Flying Wire and the Ion Profile Monitor are located downstream of Q102; their vertical counterparts are just downstream of Q103. They have been placed in locations where injected beam can be compared to circulating beam. Both of these types of instrumentation are controlled through Mac's. These computers have a built-in Ethernet adapter, so data can be loaded directly onto the local Ethernet hub.

Timing for the MI-8 line diagnostics is coordinated through the Booster Extraction Sync.

The Abort Line

Protons circulating in the Main Injector are normally sent to one of several destinations—the Tevatron, the Fixed Target experiments, the Antiproton Source, or NuMI. Sometimes the beam needs to be disposed of without being sent to a user. This can happen routinely, as during machine studies, or unexpectedly, such as when an abnormal situation is detected. The beam has enough energy to create radioactive isotopes from the metals in the magnets and the beam pipe; in order to minimize residual radiation from aborted beam, the beam is sent to a beam absorber (also known as the beam abort or the abort dump.) The beam still activates the materials in the absorber, but losses are localized and easily isolated.

In the Main Injector (and the Recycler), beam is aborted from the straight section at MI-40 (Fig. 7-15). This section will deal with several aspects of the abort line. There are kickers to deflect the beam horizontally. Lambertsons are needed to bend beam down vertically. There are several dipole and quadrupole magnets in the line to steer the beam and keep it focused on its way to the absorber, and they must be prepared to abort the beam at any energy between 8 GeV and 150 GeV. The magnets and the beam absorber require water-cooling. Finally, some consideration will be given to the set of logical conditions used to determine if the beam will be aborted.

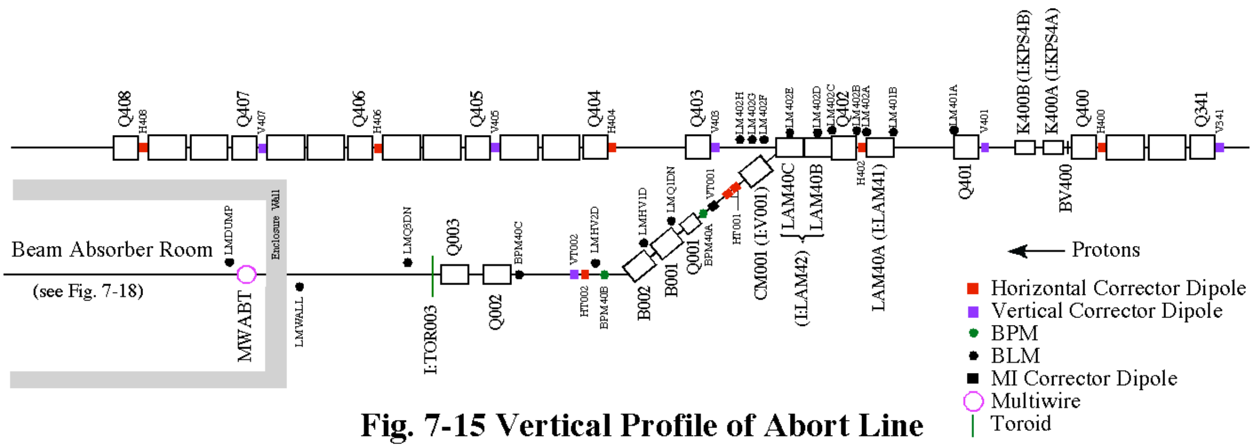

Fig. 7-15 Vertical Profile of Abort Line

Vertically, the beam is pushed down, out of the ring, Horizontally, it is bent to the right. The abort line passes through the enclosure wall and on to the abort dump. The vertical scale of the abort line is exaggerated.

The ACNET names differ from the magnet names and are given in parentheses—they normally indicate controllable power supplies. All of the magnets downstream of I:V001 are powered by the main quadrupole busses and are not individually controllable. B001, B002, and Q002 are powered from the QF bus; Q001 and Q003 are on the AD bus.

Incidentally, antiprotons are not aborted from the Main Injector. They are normally circulating only during Collider mode, and their numbers are not sufficient to justify the expense of a second absorber.

Abort Kickers

The two abort kickers are recycled from the Main Ring, where they served the same purpose. They are located just downstream of Q400, at the beginning of the straight section at MI-40. Background information on kickers can be found earlier in this chapter.

Two conductors flank the elliptical beam pipe inside the kickers. Viewing the magnet from the perspective of an approaching proton, the pulse of current first enters the conductor on the left side of the magnet and returns through the one on the right. For one brief shining moment (about 10 microseconds) the current creates an upward-pointing field in the beam pipe; protons experience a force which pushes them to the right. As with other kickers, the field is intensified by surrounding ferrite.

The charge is stored for the abort kickers in a PFN instead of a PFL. The PFN consists of 14 modules made up of capacitors and inductors. The kicker has to be able to carry the current required to abort the beam for the entire 10 microseconds of the Main Injector revolution (at least during Fixed Target mode). Such a long pulse would require a prohibitive amount of cable for a PFL.

In the kicker room at MI-40, there are three high voltage cabinets for each of the two kickers—one for the PFN, one for the Thyratron, and one for a pulsed transformer that steps down the voltage (for higher current) just before being sent to the kickers. The NIM and CAMAC controls are similar to those of the other kickers (Fig. 7-16).

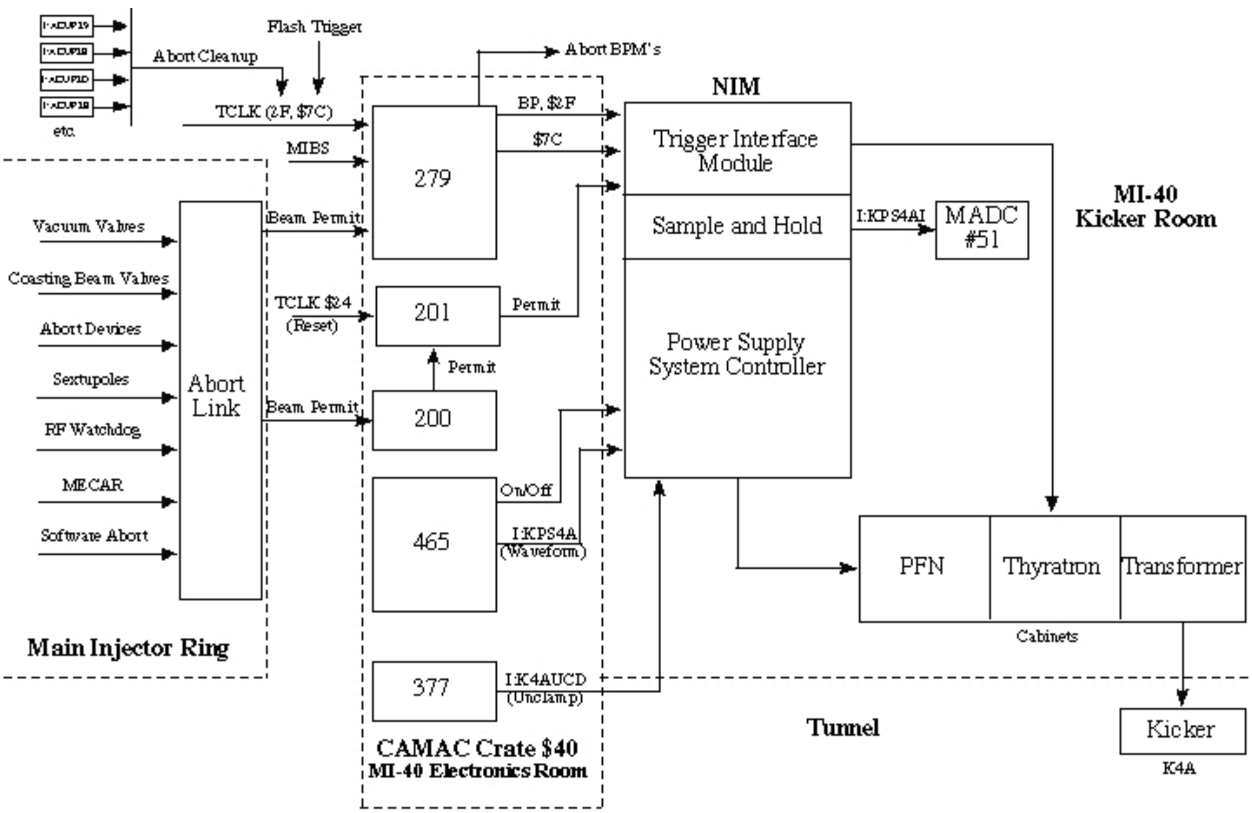

Fig. 7-16 Abort Kicker Control

This diagram shows a few of the permits and timers that control the upstream abort kicker, K4A. A diagram for the second kicker, K4B, would be similar.

The kickers must be able to abort the beam at a moment's notice, at any energy, so the charge stored in the PFN must be matched to the beam energy at all times. The ACNET parameters for the charging waveforms are I:KPS4A and I:KPS4B. The waveforms, played from CAMAC 465 cards, are reset-specific and track the Main Injector program momentum (M30). Digital control is also implemented through the C465. The unclamp parameters, stored in a 377 card, are I:K4AUCD and I:K4BUCD; they are issued with each machine reset.

Each kicker has its own Trigger Interface Module, which in turn gets its timing pulse from a CAMAC 279 beam sync decoder module. (To simplify the following discussion on timing, only Kicker A parameter names will be mentioned. Kicker B uses similar names.)

There are three main inputs into the 279 module: TCLK, MIBS, and the beam permit:

- TCLK encodes the $2F, or abort cleanup pulse. The cleanup pulse is reset dependant and is set for a time in the cycle when the beam is no longer needed. For example, on a $29 cycle the cleanup pulse occurs after the scheduled extraction time to Pbar but before the end of flattop. If for some reason extraction does not happen, the beam is still aborted in a controlled fashion. The $2F itself is generated from a 377 module in the MAC Room; the $2F is a summation of the I:ACUPxx series, where xx stands for a Main Injector reset event.
- MIBS controls the precise time that the kicker begins to fire, which must occur during the gap in the beam. When the 279 card intercepts the $2F from TCLK, it waits for the $AA marker on MIBS and then counts down the correct number of buckets to the gap before sending the timing pulse to the trigger module.
- If the beam permit is removed (i.e. the abort permit is dropped), the trigger is fired. This can happen at any time in the cycle. The cleanup pulse is irrelevant in this case, but the 279 card must still wait for the proper bucket in order to insure that the beam is aborted cleanly. More on the beam permit is on the way.

The $2F and the loss of the beam permit are summed into a single MIBS delay time, I:K4APFD (i.e., if the permit is lost before the cleanup is scheduled, the beam gets aborted anyway).

Another way of aborting the beam is to use the TCLK $7C event. The $7C, or flash trigger, is unusual in that the TCLK event itself is generated by another 279 module, which is in the MAC room. This option is used when aborting the beam on the first or last turn. Since the flash trigger and the abort time are synched to the same source it is possible to get a coordinated picture of a single turn. The delay time between the $7C and the abort trigger is given by the parameter I:K4AKD1.

The choice between I:K4APFD and I:K4AKD1 is made by turning them on or off from a parameter page.

The BPMs in the abort line get their trigger pulse from the same 279 card.

Abort Lambertsons

The “standard” Lambertsons developed for the P1 and A1 lines are also used in the abort line—except that in the abort line, the magnets are connected with the opposite polarity so that the beam is bent down instead of up (Fig. 7-15). There is one Lambertson magnet upstream of Q402 (LAM40A) and two downstream of Q402 (LAM40B and LAM40C).

The rightward kick from the kickers places the aborted beam in the field region, and of course the circulating beam passes through the field-free notch. It is important that any aborted beam be pushed cleanly over to the field region, but it is perhaps even more important that the circulating beam have a clear path through the Lambertsons (after all, aborted beam is the exception, not the rule). This is achieved partly by offsetting the main horizontal quadrupoles (Q400, Q402, and Q404) to produce a three-bump from the quad steering. Q400 is offset slightly to the left (that is, to the inside of the ring), so that beam sees, on average, a push to the left. Therefore, at the Lambertsons 90° in phase space away, circulating beam is further to the left and fits more comfortably in the field-free notch. Q402 is offset to the right, so the beam sees a push to the right. Q404 is offset to the left, completing the 3-bump. The main quadrupoles, of course, roughly track the energy of the beam.

More Magnets

Once beam is clear of the Lambertsons, there are several more magnets for steering beam toward the absorber. The first is a "C" magnet called CM001. The notch in the magnet that surrounds the beam pipe is relatively shallow; beam coming out of the Lambertsons is on a downward slope, so the circulating beam passes through the upper beam pipe while aborted beam enters the magnet. CM001 bends the beam down further. The ACNET parameter for CM001 is called I:V001.

Following the "C" magnet are two B2 magnets, called B002 and B003. They have been recycled from the Main Ring, and have been rotated 40° in order to provide both horizontal and vertical bending. Horizontally, the beam is bent to the right; vertically, since the beam is now close to the proper vertical position, it is bent up so that the downward slope is almost cancelled.

There are also several recycled Main Ring dipole correctors used in the abort line. HT001 is actually a doublet just downstream of CM001; VT001 is just upstream of Q001; and HT002 and VT002 are between B002 and Q002.

So much for the bending magnets. For focusing, there are three quadrupoles. Q001, which with its "odd" designation can be recognized as a defocusing magnet, is just after the "C" magnet. It is a 52" recycled quad from the Main Ring. Q002 and Q003 form a doublet further downstream, and are recycled 84" quads.

Power Supplies and Ramps

Block diagrams of power supplies and their associated magnets can be found on Fig. 7-17.

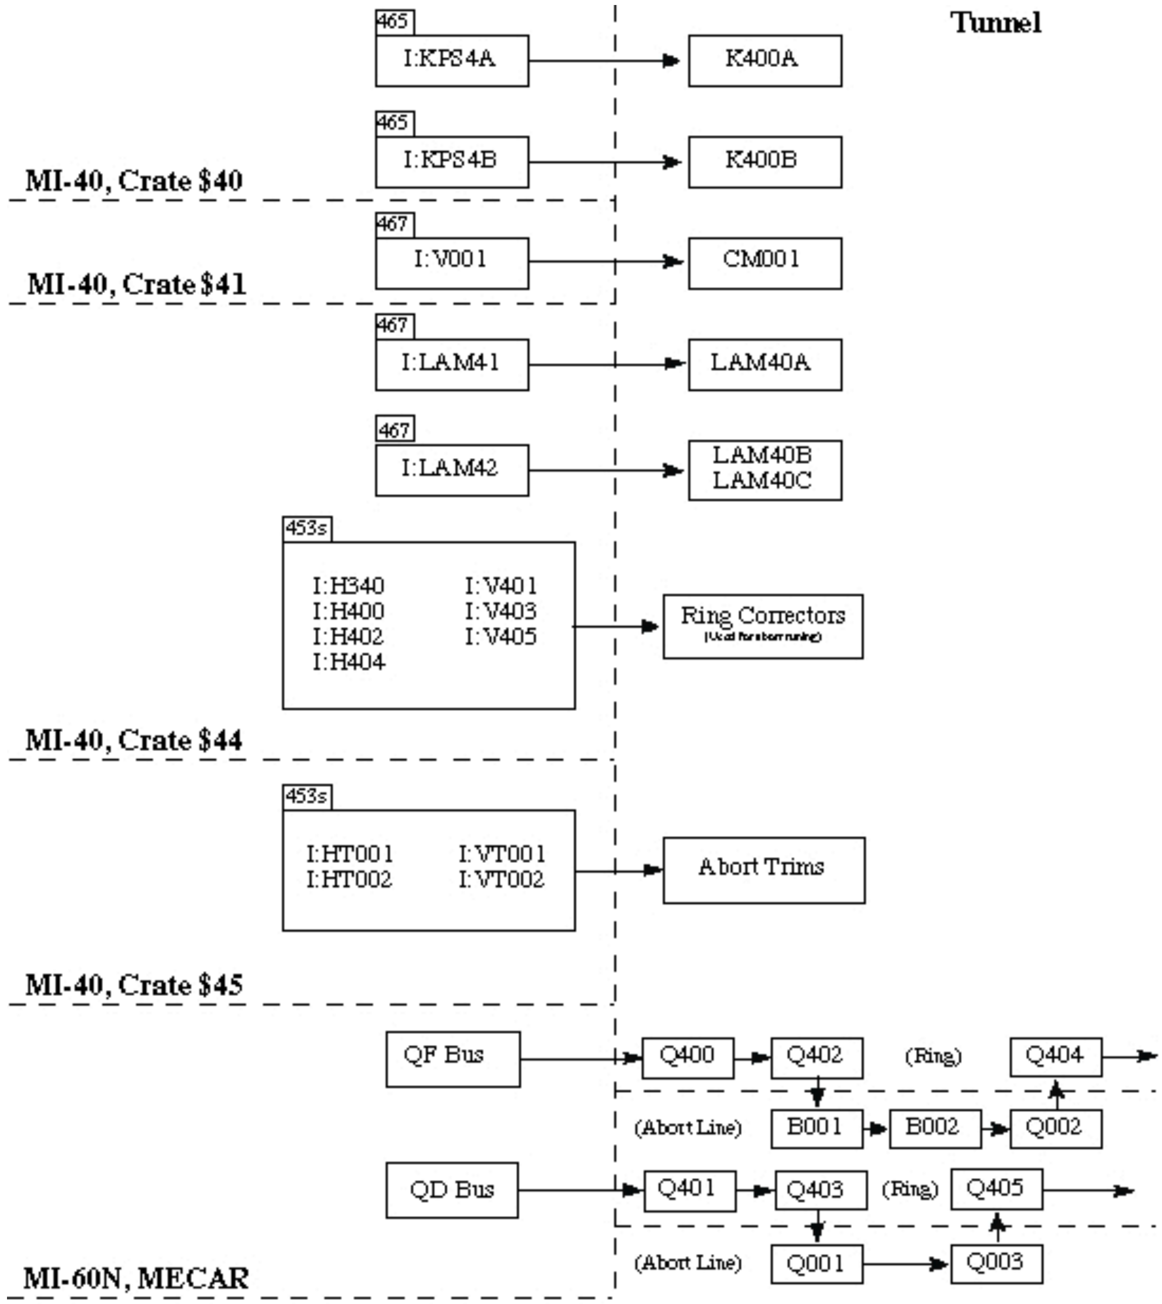

Fig. 7-17 Abort Line Power Supplies
Compare to Fig. 7-15

There are two power supplies for the Lambertsons—I:LAM41 and I:ILAM42. I:LAM41 controls LAM40A, the Lambertson upstream of Q402.

I:LAM42 powers LAM40B and LAM40C, which are the two magnets downstream of Q402. As mentioned earlier, I:V001 powers the "C" magnet. The three Lambertsons and V001 form a quartet of downward bending magnets, but they are not always ramped in concert. Since LAM40A lowers the beam going through Q402, quad steering will tend to push the beam back up. This dilemma can be postponed by letting the LAM42 magnets do most of the bending early in the ramp. Then, at higher energies, LAM41 and V001 can be brought into play. LAM41 and V001 are ramped more or less symmetrically so that the angle of descent through LAM42 remains constant. By the time the beam energy is at 150 GeV, all four magnets are sharing the load equally.

Further ahead in this chapter (Fig. 22(a, b, and c)) there are diagrams of beam as it passes through the P1 line Lambertsons. They might be useful for visualizing what is happening in the abort Lambertsons, but remember that beam is being pushed down in the abort line, and up in the P1 line.

The Lambertsons and "C" magnet are controlled through CAMAC 467 cards in Crate $41.

Several magnets are powered through the main quad busses, and therefore have no ACNET parameters or independent control. B001, B002, and Q002 are on the QF bus, while Q001 and Q003 are powered from the QD bus.

The Absorber Room

The beam pipe of the abort line, which has been deflected downward and to the right from the Main Injector ring, passes through the enclosure wall near 406. Eventually it enters the Absorber Room (Fig. 7-18).

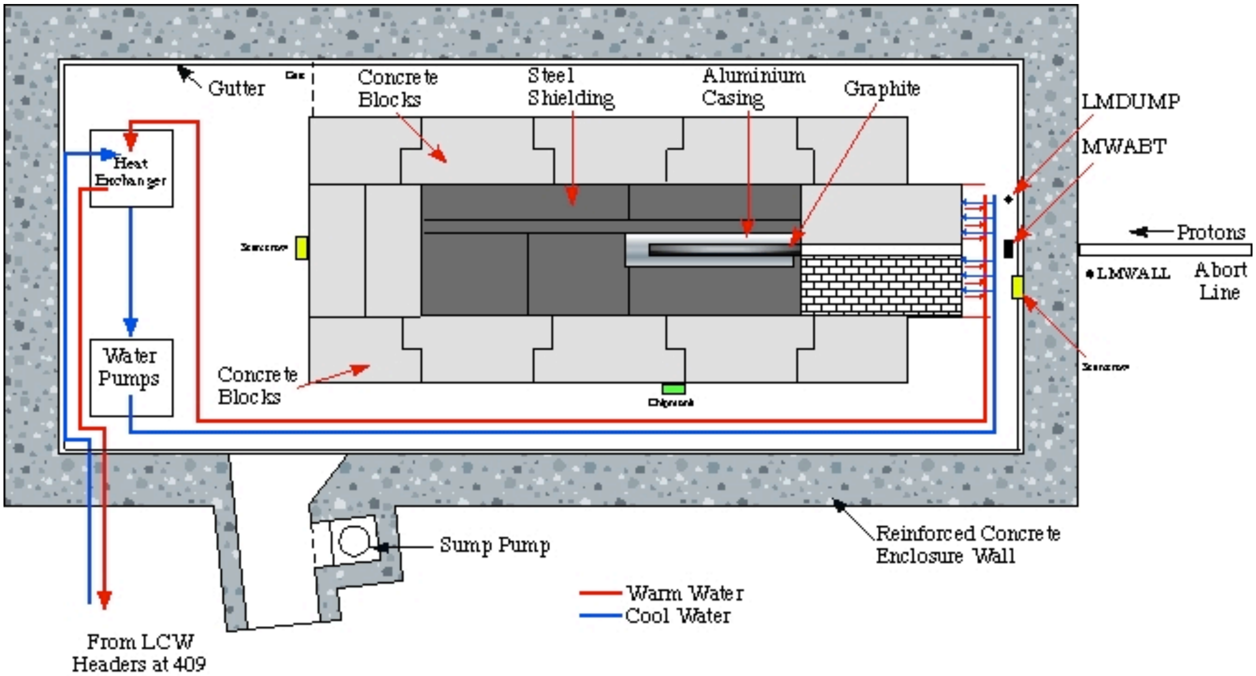

Fig. 7-18 Beam Absorber Room (Top View)
The aborted beam, arriving from the right, is stopped by a series of materials. First is the graphite core, which absorbs much of the thermal energy and slows the secondary neutrons as well. The graphite is contained in an aluminum sheath that is surrounded by steel and then concrete.

The steel, which conducts heat readily, contains numerous channels for water cooling. The closed loop system prevents activated water from entering the main lines. Compare Fig. 4-3 and 7-15.

The purpose of the absorber is to shield the outside world from exposure to the high-energy beam, which creates radioactive isotopes in the materials it encounters. It consists of several concentric layers.

Graphite: Graphite, similar to that found in pencil "lead," is a form of pure carbon (although in pencils, the graphite is mixed with clay to make it harder). Unlike the chemically identical diamond, in which the carbon atoms are polymerized in three dimensions, the carbon atoms in graphite are only polymerized in two dimensions; graphite is therefore much softer than diamond.

Why graphite? One reason is its high temperature of sublimation (it doesn't melt) — about 3825˚C. During periods of tune-up in the Main Injector, the absorber may be required to take a thousand or more hits an hour of 120 GeV beam, which represents a significant thermal load.

The other major reason for using graphite is its ability to slow down energetic neutrons. Fast neutrons produced when the protons hit the target collide with the relatively light carbon nuclei and lose much of their energy; otherwise, they might penetrate the walls of the absorber room. "Fast" neutrons are converted to "thermal" neutrons.

And, of course, graphite is cheaper than diamond.

The graphite absorber is shaped as a long cylinder. The beam impacts the cylinder along its longitudinal axis and is therefore attenuated by several meters of graphite. The graphite is encased in an aluminum sheath.

Steel: The second line of defense against secondary particles is the steel surrounding the graphite core. Steel is a relatively inexpensive, yet relatively dense material that can stop many secondary particles.

Embedded in the steel are numerous water channels for carrying off the heat radiating from the graphite. The outer dimensions of the steel block, looking head-on, are 2' by 2'. The water system is described in more detail below.

Concrete: Finally, the graphite/steel block is surrounded by large concrete blocks. The blocks are a cheap way of providing a lot of additional shielding.

Water Cooling: (Background information on LCW systems in general can be found in Chapter 4.) As mentioned above, channels used for water-cooling permeate the steel block. The water, although heat exchanged with the LCW from the main headers, is a closed loop system. This is because the water itself can be activated by energetic secondary particles passing through the steel. There are short-lived isotopes that can decay in a matter of minutes, but there is also tritium (H^3), with a half-life of over 12 years.

A simplified drawing of the abort water system is included with Fig. 7-18. There are 3" supply and return headers that branch off from the main LCW headers at 409 (which is also the location of the stairwell entrance to the absorber room). The LCW makes a single pass through a small heat exchanger and returns to the main header. Flow is maintained by the differential pressure between the supply and return headers.

Within the closed loop system, warm water returns from the channels in the absorber, passes through the secondary loop of the exchanger, and is cooled by the LCW. Circulation is maintained by two small pumps.

A 21-gallon storage tank receives the makeup flow from the LCW supply line. As is usually the case, head pressure on the tank is maintained with nitrogen gas. Makeup water from the storage tank to the abort system enters the return line just upstream of the pumps. There is also a "recirculation" pump that siphons excess water from the high-pressure side of the abort system and transfers it to the tank.

The configuration of the abort cooling system is somewhat simpler than most of the other water systems in the Main Injector because there is no need to deionize the water (it doesn't pass through any electrical components), nor is there a need to precisely control the temperature of the absorber—it just needs to be kept reasonably cool.

Abort Link

The beam permit is based on the abort link, which originates in a CAMAC 201 module in the MI-40 electronics room. The signal launched from the 201 module is repeated from building to building by the CAMAC 200 abort concentrator modules. The purpose of the concentrator modules is to accept local inputs that are able to interrupt the signal and break the link, which causes the abort kickers to fire. Refer to the Controls Rookie Book for more details on abort links in general.

Compared to the Tevatron, the Main Injector 200 modules are sparsely populated with inputs. All of the inputs can be found on I67. At all houses, a closing vacuum valve will pull down the permit—there is no sense in pounding a valve with beam. In addition, at MI-30 there are two beam valves (BV301 and BV309) specifically designated to stop beam in the unlikely event that all else fails—these are also tied in to the abort. At MI-40 the inputs are from the power supplies for the abort line itself—for the kickers, Lambertsons, and C-magnets. At MI-52 the inputs come from the vertical and horizontal sextupoles, at MI-60S it is the RF watchdog (that also inhibits the anode supply); and at MI-60N it is MECAR that pulls the abort if the main power supplies are becoming too much of a problem.

Finally, at MI-10 there is the ever-mysterious “Software Abort.” Through this input, any ACNET parameter can become an input to the abort system with the proper modifications to the database; many of the loss monitors in the abort line itself can pull the abort. Of course, in a case like this the abort is already in progress and its real purpose is to inhibit the next pulse of beam, and to make operators aware of the problem.

When an abort is reset, TCLK event $24 is broadcast on the clock. When the C201 module intercepts the $24, it attempts to restart the abort link.

TCLK event $27, issued by the BSSB, is an announcement to hardware that “Beam has been aborted.” By creating a time stamp, the $27 creates a record of the exact time of the abort that can be useful in diagnosis.

The P1, P2, and P3 lines each have their own permit systems, to be described in the next sections. Since there is no way—or need—to abort beam from these single-pass lines, the permit system can only prevent unwanted beam from entering those lines.

Abort Diagnostics

Beam diagnostics for the abort line include a multiwire (MWABT), a toroid (TOR003), and a dense array of loss monitors. TOR003 is upstream of the absorber room, and MWABT is located just inside the wall where the beam pipe enters the absorber room. There are also BPMs, but they are seldom if ever used. Fig. 7-19 shows the abort system located in the MI-40 service

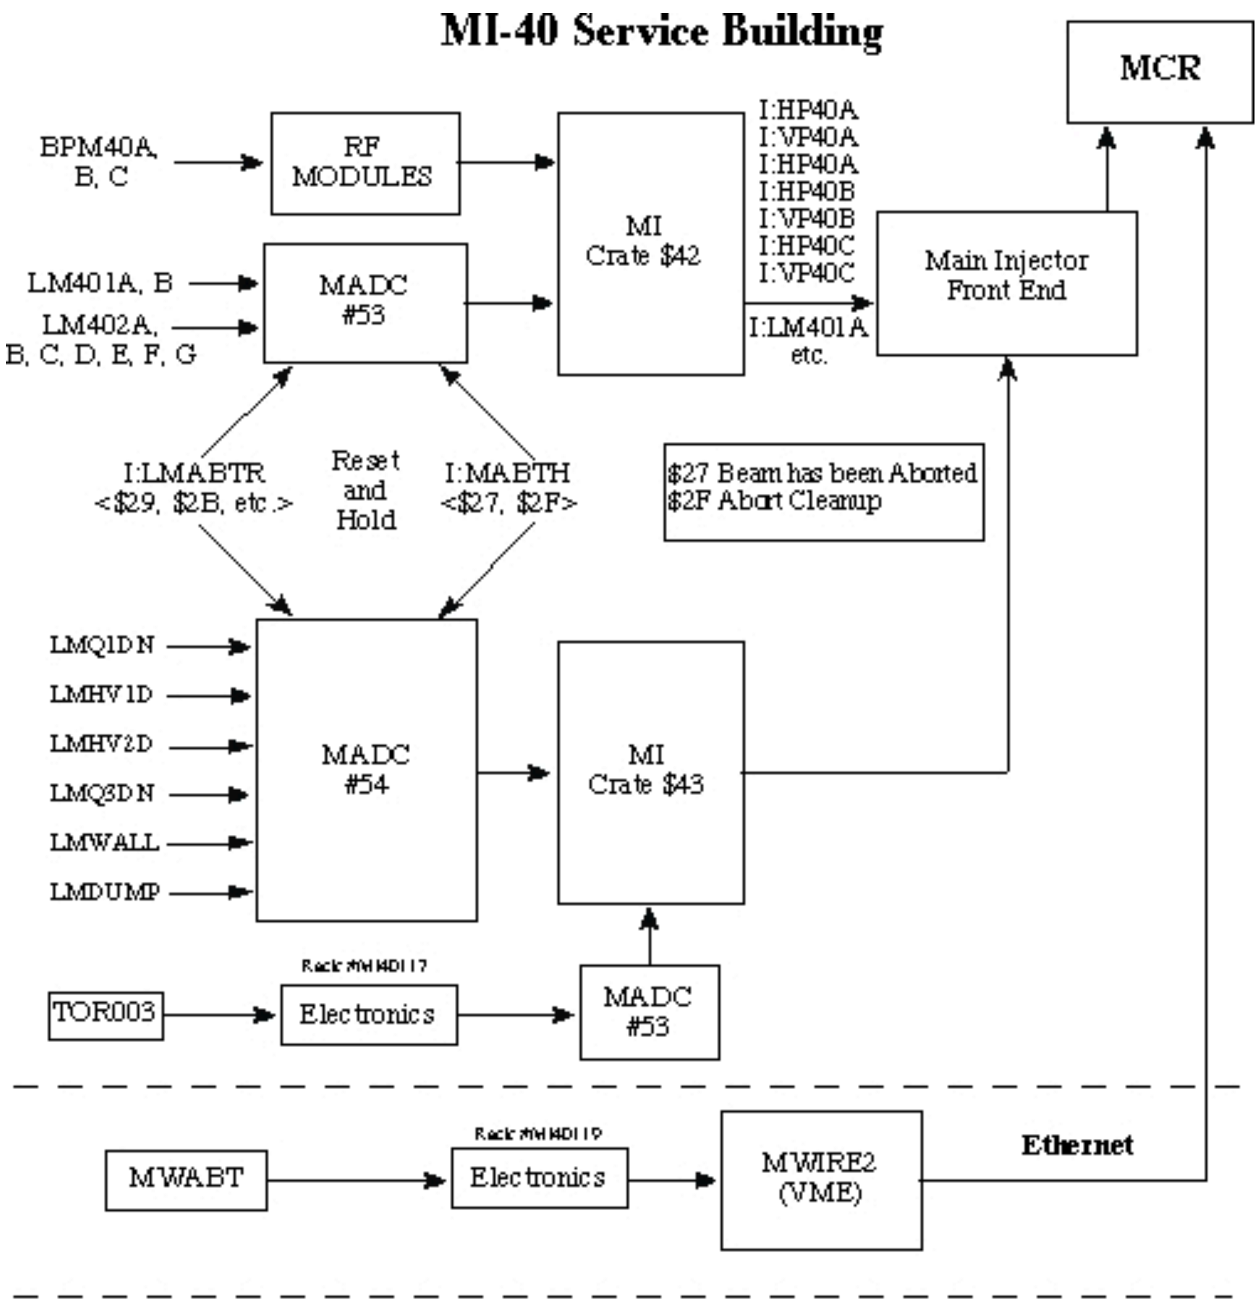

Fig. 7-19 Abort Diagnostics

building.

I:LMABTR, the abort reset trigger, is referenced to the Main Injector machine resets ($29, $2B, etc.) with a delay of about half a millisecond.

I:LMABTH is the "hold" event that tells the BLMs to latch onto the integrated value; it is referenced to $27 and $2F events with a delay of 5 milliseconds (enough time for the loss monitors to integrate the charge). The latched values can be read on a parameter page, but, in practice, real-time plots of critical monitors (such as I:LMQ3DN) are used while tuning.

The P1 Line

The P1 line connects the Main Injector to the F0 Lambertsons (Fig. 7-20). It is used by protons extracted from the Main Injector and by antiprotons entering the Main Injector. Operationally, it is probably the most complex of all the beam lines, because it must transport beam in both directions and at several different energies. It also sends protons toward two distinct destinations—the Tevatron and the P2 line—each of which requires customized beam trajectories and optics (Figs. 1-4, 1-5, 1-6, 1-8, and 1-9). In addition, the Main Injector, Tevatron, and P2 line are all at different elevations.

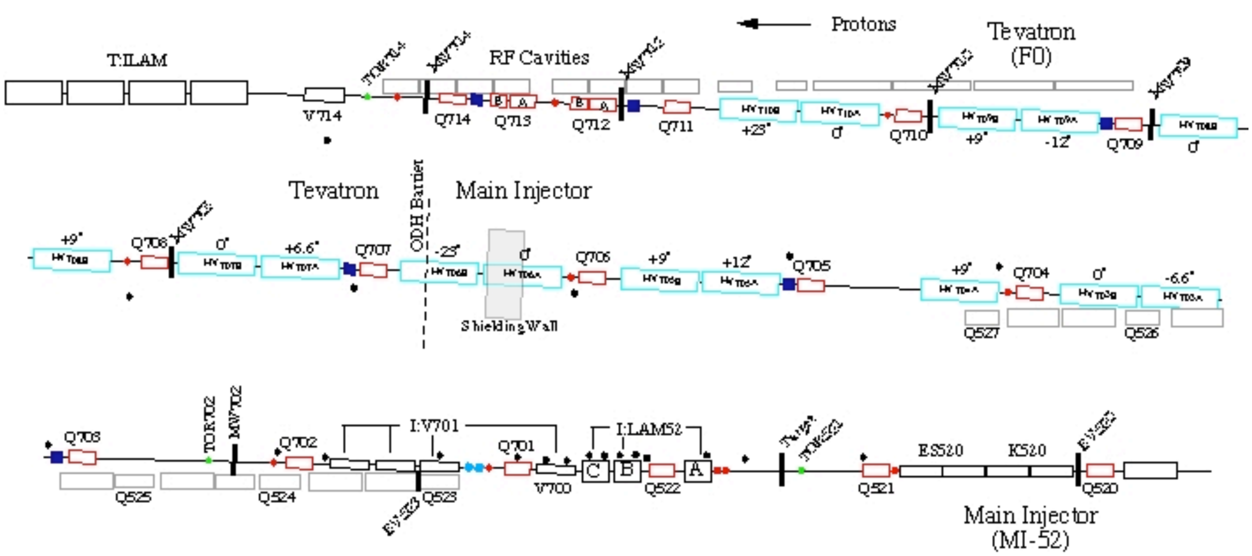

If you are a proton, read this map from right to left and from bottom to top. If you are an antiproton, do just the opposite. Beam is kicked horizontally by K520 (two single turn magnets) and then vertically upward (for the most part) by LAM52. Collectively, the B2 magnets bend horizontally, but many are rolled to create a vertical component to the beam path. The amount of roll, in degrees, is indicated above the magnets. (Compare to figure 7-23.)

B2 Dipole
Vertical Corrector, Main Ring Style
Horizontal Corrector, MR Style
Quadrupole
Vertical Corrector, Main Injector Style
BLM
Not Part of P1 Line

Fig. 7-20 Vertical Profile of the P1 Line

In this book, the beginning of the P1 line is considered to be at 520, where kickers initiate the extraction process. The end point of the line is the Injection Lambertson at F0 (ILAM), where the beam must choose whether to enter the Tevatron or the P2 line. P1 line locations are designated by numbers in an "I:700" series (not to be confused with numbers in the AP-2 line, which are in a D:700 series).

There are nearly a dozen different operations that use the P1 line to transfer beam from one machine to another. Extracted beam may be headed for the Tevatron at 150 GeV. 120 GeV protons in the line may be used for stacking, or resonantly extracted to Switchyard. The P1 line also sends 8 GeV protons to the Antiproton Source for tune-up; those can be sent to either the Debuncher or the Accumulator.

In the other direction, the line carries 8 GeV antiprotons from the source to the Main Injector, to be accelerated there or deposited in the Recycler. It will also accept 150 GeV antiprotons freshly decelerated in the Tevatron, sending them to the Main Injector to be further decelerated to 8 GeV and stored in the Recycler.

To complicate things further, most of the power supplies used by the P1 line are also shared with the A1 line, because the two lines are, for the most part, symmetrical.

All of these demands on the design of the P1 line can sometimes create complications for the tuner. The approach in this book will be to first describe the line in a general way, and then to go back and discuss the specific adaptations required by different modes of operation.

Several of the beam lines (P1, A1, the abort line and NuMI) are very similar to each other in the way the components are set up in the lattice to extract beam. To minimize redundancy, some of the details will be discussed in excruciating detail for the P1 line so they can be dealt with less harshly in the other sections.

P1 Line Overview

The basic pattern of kickers, Lambertsons, and C-magnets used to extract beam from the Main Injector at MI-52 is similar to that of the abort line, except that beam is deflected up instead of down (Figs. 7-15, 7-20). Vertically, the beam needs to be moved up to the level of the Tevatron, which is 2.133 meters higher, by the time it reaches F0. T:ILAM, if powered, straightens out the upwardly rising beam so that it continues on into the Tevatron. If those Lambertsons are not powered, beam continues upward into the P2 line.

Considering the horizontal component, be aware that if the beam were simply allowed to escape at a tangent to the Main Injector at 520, it would quickly reach the Tevatron upstream of ILAM without any assistance from kickers or magnets. Although the kickers kick to the outside of the Main Injector ring, the P1 magnets as a whole continue to bend beam to the left until it is tangent to the circle of the Tevatron at F0.

A majority of the magnets in the P1 line are “rolled” so that they provide both horizontal and vertical changes to the beam trajectory.

The MI-52 extraction kickers are located at the beginning of the straight section, immediately downstream of Q520. There are two single-turn magnets; they kick the beam to the right, toward the outside of the ring. The parameter for the kicker power supply is I:KPS5S, for “Kicker Power Supply at MI-52, Short.” The “Short” designation is a historical remnant from the days when there was a “Long” kicker as well. “Short” means up to 84 bunches, or a complete Booster batch, in length. (The kickers are immediately followed by the extraction septa, which are only used during resonant extraction. Except for the origin of the deflection, resonantly extracted beam follows the same path through the P1 line as kicked beam.)

The kicked beam soon finds itself in the field region of I:LAM52, the extraction Lambertsons (Fig. 7-2, 7-22 a, b, & c). There are three Lambertson magnets—LAM52A, LAM52B, and LAM52C, with LAM52A being upstream of

Q522 and LAM52B and LAM52C downstream of Q522. Collectively, they begin to bend the beam up toward the Tevatron. The Lambertsons are rolled by 10˚ or so, which helps with the horizontal angle as well. Fig. 2-8(a) shows the short straight-section lattice at MI-52 that must accommodate the extraction devices.

Following the Lambertsons, there are four C-magnets that, on average, deflect the beam down—they are powered in series by I:V701. The first, called V700, is upstream of the first quad in the line (Q701), and the other three, known collectively as V701, are downstream of Q701.
V700 is only used during 120 GeV or 150 GeV extraction and the magnet bends the beam up. By the time the beam reaches Q701, it has risen above the Main Injector centerline by 21.8 cm and has an upward deflection of 28.5 mrad. The magnets of V701, which are powered for all extracted beam, push down on the beam and reduce the upward angle. (A bit more on the reason for this is on the way.)

The large bending magnets in the P1 line are recycled B2 dipoles from the Ring formerly known as Main. One of these, H703B, is independently powered, by I:H703. The task of H703B is to adjust the horizontal angle of the beam. The remaining 14 are powered in series from a common power supply (I:HV703). Walking past the line leaves the impression that an earthquake has recently struck, because many of the magnets have been rolled to produce both horizontal and vertical components to the field (Fig. 7-20 shows the roll angles for the B2 dipoles, in degrees).

Four of the dipoles are not rolled; i.e. they bend the beam only in the horizontal plane. Whether rolled or not, all of the magnets in the I:HV703 series bend beam to the left, bringing it parallel to the Tevatron beam pipe.

The remaining ten of the B2 dipoles each have at least some vertical component to their bending strength. They can be grouped into four "families." The dipoles in the first family are rolled counterclockwise by 9˚. (The convention here is that a counterclockwise roll, viewed from upstream to downstream, is positive, and that the vertical component of the force pushes the beam up.) There are two pairs in this family, and they are responsible for much of the remaining upward deflection in the line.

The remaining three families include one pair of magnets each, with the roll of one in the opposite direction of the other: +/-6.6˚, +/-12˚, and +/-23˚. These rolls are designed to cancel out the vertical dispersion in the line.

At the end of the line, a C-magnet called V714 gives the beam a final push upward to ensure that it will enter the field region of the Lambertsons. V714 pushes even harder if beam is meant to go on up to the P2 line. When all is said and done, there is an upward angle of 24 mrad as the beam arrives at the injection Lambertsons. If the Lambertsons are powered at the time, they cancel the angle, and horizontally level beam continues into the Tevatron. (The beam still has a horizontal angle that must be removed by the F17 kicker in the Tevatron.) If they are not powered, the beam continues to rise toward the P2 line, which is 64.6 cm higher than the Tevatron. In that case, the horizontal and vertical injection angles are eventually cancelled in the P2 line (next section).

In a manner similar to the MI-8 line, the lattice in the middle of the P1 line is flanked on either side with a matching lattice to the adjacent machines. Upstream, of course, is the match to the Main Injector. Downstream, the lattice has to match that of either the Tevatron or the P2 line, which are not identical. The string of quadrupoles from Q703 to Q709 establishes the central FODO lattice in the P1 line; they are powered in series by I:Q703 (Fig. 7-23). Q701 and Q702, which are independently powered, match the P1 FODO lattice to that of the Main Injector. The last five quads, Q710 through Q714, are all individually powered and match the P1 FODO lattice to either the Tevatron or the P2 line. The matching lattice is changed from pulse to pulse by varying the current in the magnets— in particular, the ramp waveforms of Q713 and Q714 change the most.

Of course, there are corrector dipoles in the line, usually located downstream of the appropriate quad. All of the horizontal dipoles are recycled from the Main Ring, while most of the vertical dipoles are of the newer Main

Injector style. The vertical correction at 701, however, uses a doublet of MR-style dipoles (VT701-1 and VT701-2). LEP magnets are not used in the P1 line because their high inductances make for a sluggish response to the constantly changing energy requirements.

Now for the excruciating details....

Kickers

As mentioned earlier, there are two kicker magnets downstream of Q520. The pulse is generated by KPS5S, located at MI-52. The pulse has a short flattop and is used to kick single batches or partial batches. It uses a pulse forming line (PFL); that is a long cable, to store its charge. The charge is fed to the cable from a Glassman high voltage supply.

Like the abort line, the P1 line uses TCLK, MIBS, and reflection events for controlling transfers. The biggest difference is that the P1 line has so many of them. Table 7-1, located at the very end of this chapter, attempts to summarize these. The kickers, magnets, and beam diagnostics use the events, plus delays, as triggers. For example, the flash frame delay triggers the BPMs so that they sample the beam at the instant it passes through the line, using many of the same events that allowed the kickers to fire. Every mode of operation has its own unique and complicated set of timers.

As usual, a TCLK event initiates the events associated with charging the supply (Fig. 7-21). For example, during stacking a single batch is injected and accelerated on the $29 cycle. Event $80, the “Stacking Reset” trigger, is actually issued by the TLG one 15 Hz tick (67 ms) earlier than the $29. Devices specifically needed for stacking, such as the kickers and the P1 line magnets, will be referenced directly or indirectly to the $80 event. First, there is the unclamp delay (I:K5SUCD) which, once initiated by the $80, tells KPS5S to begin charging its PFL.

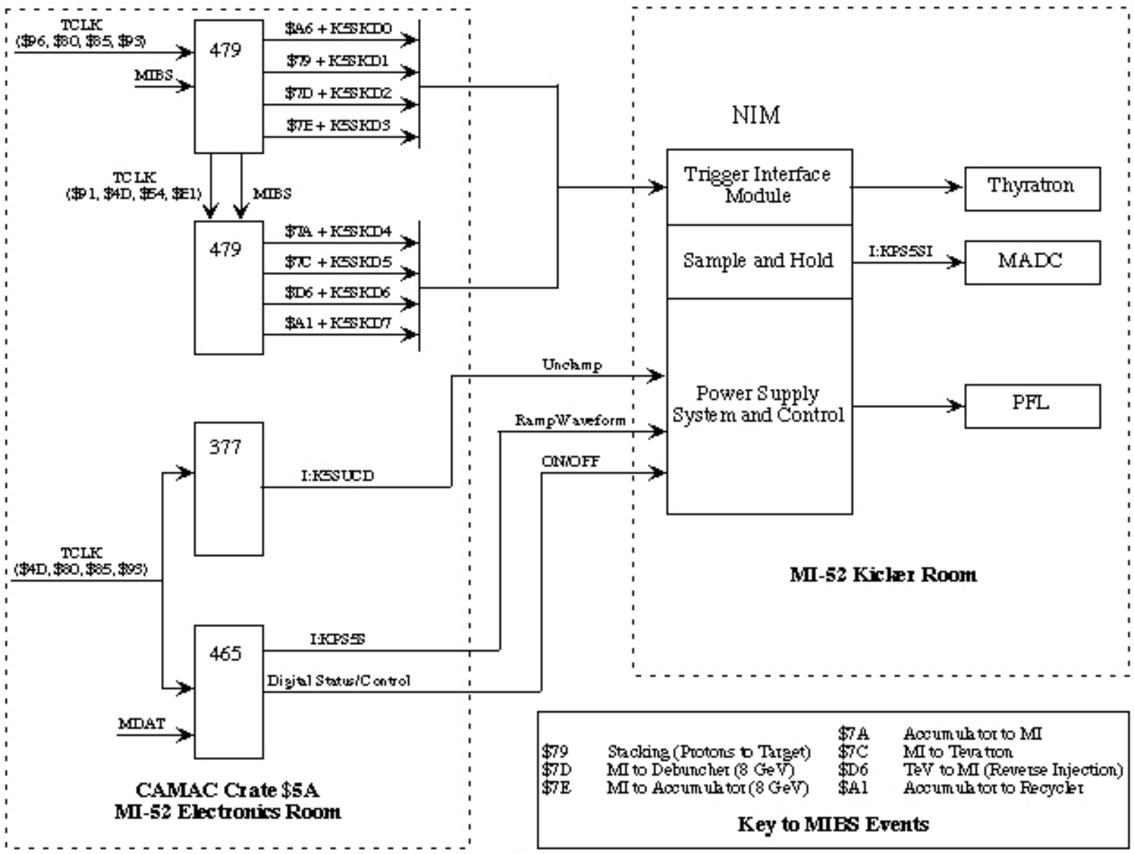

Fig. 7-21 MI-52 Kicker Timing

KPS5S also begins to charge up when it sees events $4D, $85 and $93 (Fig. 7-21, Table 7-1). The $85 and $93 events will transfer 8 GeV beam to the Debuncher and Accumulator, respectively; beam on the $80 event will be at 120 GeV, and the $4D beam will be at 150 GeV. The kickers will charge to a level appropriate for the beam energy. The ramp tables for the kicker waveforms track the MDAT signal for Main Injector momentum.

Once the PFL is fully charged, another trigger is needed to fire the kicker—that is, to unload all of that stored energy into the kicker magnets. Continuing with the story of the stacking pulse, on board a CAMAC 377 card in Slot 3, Crate $94 (home of MIBS) there is a timer called I:MIPBTX. This timer, with a value (as of this writing) of .905555 seconds, initiates the extraction of beam. It begins counting down the interval as soon as it decodes an Event $80 from TCLK.

When the time is up—and it had better coincide with flattop on the $29 cycle—the 377 issues a pulse which, after some processing by other cards in Crate $94, is transformed into the beam sync Event $79 and broadcast on MIBS. Event $79 is, in turn, decoded by a 479 card in Crate $5A at MI-52. The 479 card now knows that it must choose a moment in the very near future to send the timing pulse to the Trigger Interface Module; that moment must be just as the leading edge of the batch approaches the kickers, taking into account the rise time of the kickers.

Once armed, the 479 card begins its own countdown. The countdown delay, set from the parameter I:K5SKD1, is given as 24.9 "MREV," or "Main Injector Revolutions." (It is actually counting in "RF buckets," in groups of seven, with one revolution equal to 588 buckets.) Since the bucket offset between the location of the batch and Event $79 is fixed, the delay can be set so that the kicker fires at the correct bucket.

Mode of Operation	TCLK	Beam Sync	BS Reflection	BS Delay	Flash Delay	Ext Kicker	Inj Kicker	Beam Lines
Booster to Main Injector	$2x	BES			I:BPFTBM	MKS02	MI 10	MI-8
Tevatron Fixed Target	$4D	$78	$58	I:MITPX	I:BPFTBM	MI52	F17K	P1
TeV Collider, Forward Protons	$4D	$7C	$5C	I:MITCPX	I:BPFT7C	MI62	F17K	P1
Main Injector to TeV Antiprotons	$40	$7B	$5B	I:MITPBX	I:BPFT7B	MI62	E48	A1
TeV to MI Antiprotons	$34	$D6	$5F	I:MITBX	I:BPFTD6	F17K	MI52	P1
Reverse Protons, TeV to MI	$5D	$D8	$55	I:TMIPX	I:BPFTTM	E48	MI62	A1
STACKING	$80	$79	$81	I:MIPBTX	I:BPFTMD	MI52	D:IKIK	P1, P2, AP1, AP2
MI Protons to Debuncher	$85	$7D	$86	I:MIDX		MI52	D:IKIK	P1, P2
Accumulator to MI Antiprotons	$91	$7A	$94	I:AMIPBX	I:BPFTAM	D:EKIK	MI52	A3, AP1, P2, P1
Reverse Protons to Accumulator	$93	$7E	$99	I:MIAX	I:BPFT7E	MI52	D:EKIK	P1, P3
Recycler to Accumulator Protons	$96	$A6	$97	R:RAX		RR30	D:EKIK	RR30, P1, P2, AP1, AP3
Accumulator to Recycler Antiprotons	$E1	$A1	$98	R:ARX		D:EKIK	RR30	AP3, AP1, P2, P1, RR30
Main Injector to Recycler Protons	$E2	$A2	$F2	R:MIRPX	I:BPFTA2			RR20
Recycler to MI Protons	$E3	$A3	$F3	R:RMIPX	I:BPFTA3			RR30
Recycler to MI Antiprotons	$E4	$A7	$F7	R:RMIPBX	R:BPxAFD			RR20
MI to Switchyard	$30	$75	$39	I:MIBLX	None	Septa	None	P1, P2, P3
MI to NuMI	$A5	$74	$A9	I:MINX				NuMI

Table 7-1 Beam Line Timing

The parameters for the kicker delay times are listed in Table 7-1.

The tables for the kicker waveforms are stored in a CAMAC 465 card in Crate $5B in the MI-52 Service Building. The C377 and C479 cards for the unclamp and trigger signals are located in the same crate.

Lambertsons and C-Magnets

The extraction Lambertsons (I:LAM52) and the C-magnets (I:V701) differ from the other magnets in the P1 line in two ways: (1) Their power supplies are not shared with magnets in the A1 line, and (2) Different subsets of each magnet string are powered, depending on the energy of the extracted beam. 8 GeV beam calls for one set of magnets, while 120 or 150 GeV calls for another. When a change from one configuration to another is required, the ramp waveform goes to zero current, and "load transfer switches" connect the proper magnets to the power supplies. The switch defaults to the 120/150 GeV state.

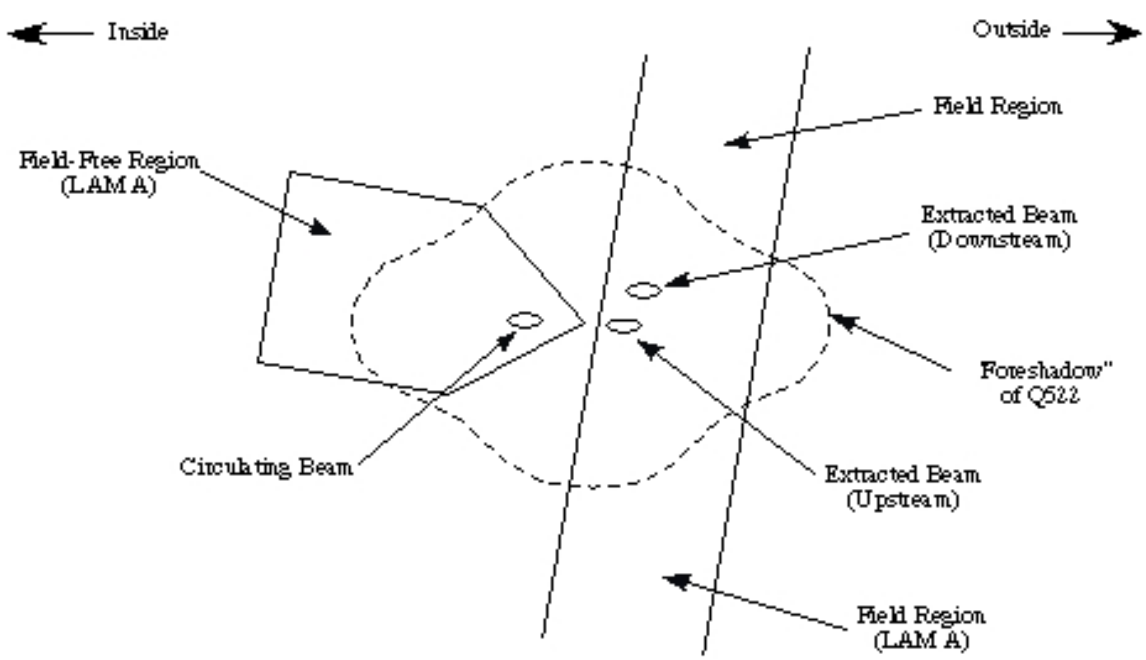

Fig. 7-22(a) 150 GeV Beam at LAM A

The extraction kickers at 520 deflect the beam into the field region of LAM A, which in turn bends it up and slightly outward.

In the case of the Lambertsons, remember that there are actually three magnets—LAM52A, LAM52B, and LAM52C. During 120 or 150 GeV extraction, all three are powered. But during 8 GeV extraction, only LAM52B and LAM52C are powered. Why is this switch necessary? Let me tell you. At 120 or 150 GeV, the strength of three Lambertsons is required to give the beam enough of a deflection to clear the beam pipe by the time it reaches Q701. In fact, there is an additional C-magnet, V700, which is used to help push the beam up even further. But there is only enough room for two Lambertsons (B and C) between Q522 and Q523. The remaining Lambertson (A) is located

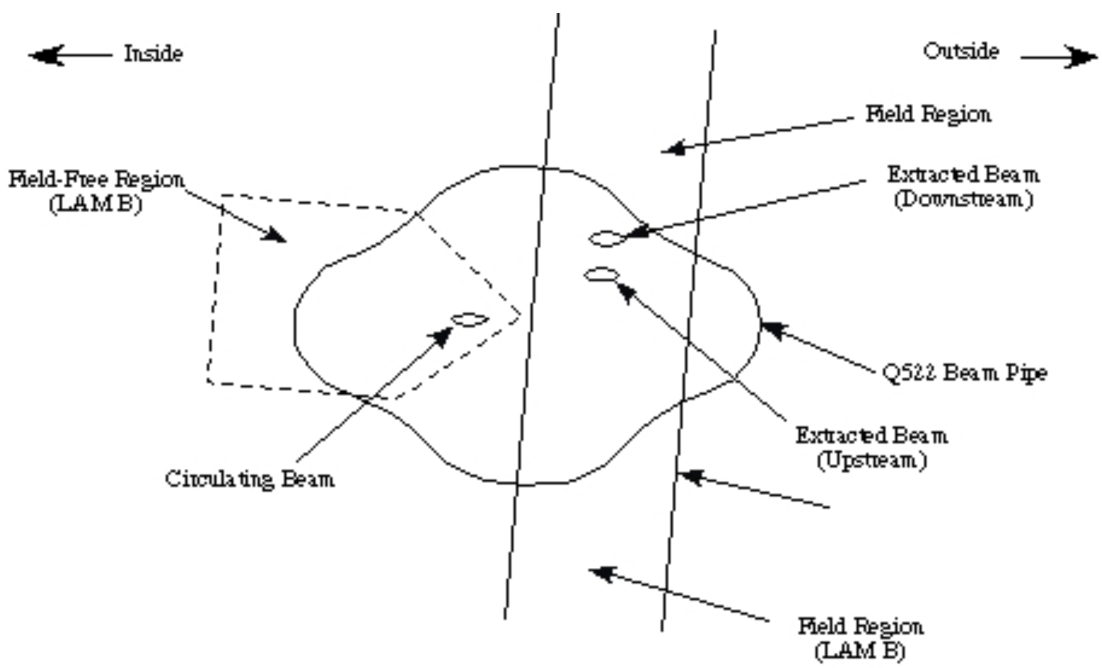

Fig. 7-22(b) 150 GeV Beam at Q522

At Q522, the extracted beam continues to move on its upward trajectory, and is helped slightly by the vertically defocusing field of the quadrupole. Compare to Fig. 7-20.

upstream of Q522 (Figs. 7-20, 7-22(a) and (b)).

8 GeV beam, however, is not only easier to deflect, it is also bigger (Fig. 7-22(c)). In order to go cleanly through the field region at LAM52, the beam is kicked horizontally and steered to the outside. That means that beam will also be to the outside as it enters Q522, which is just downstream of LAM52. If 8 GeV beam were to be bent vertically up at LAM52A, as 150 GeV beam is, it

would scrape on the upper surface of the beam pipe at Q522. So, at 8 GeV, LAM52A and V700 are disconnected from their respective power supplies.

When an 8 GeV clock event is detected, the ramp current goes to zero, the Load Transfer Switch removes LAM52A from the circuit, and current is sent only to LAM52B and LAM52C. Near the end of the 8 GeV cycle, current again goes to zero, the switch adds LAM52A to the circuit, and all three Lambertsons are set up to ramp on the following cycle if needed.

In the case of I:V701, remember that the first C-magnet, V700, is distinct from the following three C-magnets because it bends the beam up rather than down. V700 behaves like LAM52A—during 120 or 150 GeV extraction, all four magnets are powered; with an 8 GeV event, the ramp goes to zero, V700 is removed from the circuit, and the remaining three are powered.

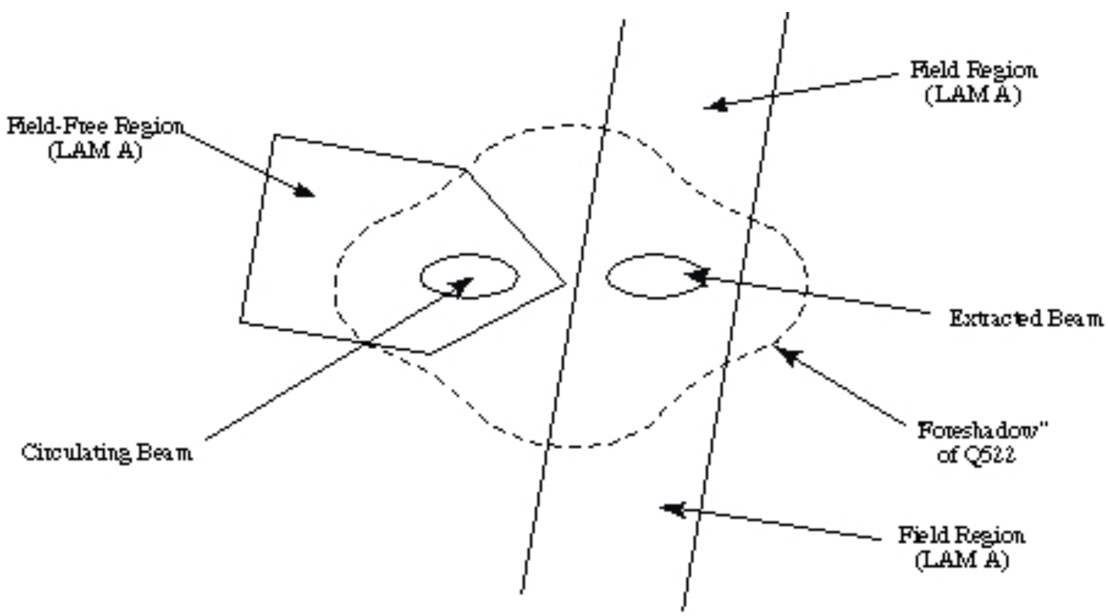

Fig. 7-22(c) 8 GeV Beam at LAM A, P1 Line

LAM A stays off so that extracted beam does not scrape on the beam pipe at Q522

At 8 GeV, if LAMA were to be turned on, the large and easily swayed beam would scrape the top of the beam pipe. The extraction Lambertsons are in series for all 120 or 150 GeV extraction. However, when beam is to be extracted at 8 GeV, as during a $2D cycle, I:LAM52 is ramped down to zero amps and a transfer switch LAMA out of the circuit (Fig. 7-23). Then, since the beam remains vertically centered through Q522, LAMB and LAMC can be ramped up to their 8 GeV value.

V700 is disconnected from the I:V701 supply at the same time as LAMA is disconnected from I:LAM52.

The load transfer switches are triggered by the parameter I:LS8ON, which itself is referenced to 8 GeV reset events ($93, $85, $E1, etc.). When I:LSSON is active, the 8 GeV subset of magnets is connected. I:LS8OFF, referenced to the same events but with a longer delay, turns the switches off, and connects the magnets needed for 120/150 GeV operation. The switch from 8 GeV to 120/150 GeV is sometimes referred to as "S1," to distinguish it from the switch "S2" that transfers power between the P1 line and A1 line magnets (to be discussed later).

The ramps for I:LAM52 and I:V701 are stored in CAMAC 468 cards located in Crate $5B at the MI-52 Service Building. A 377 card in the same crate generates I:LS8ON and I:LS8OFF. The 468 ramps used by the 8 GeV events are set up to go to zero current at those times when the load transfer switches are potentially active.

The Rest of the Magnets

The power supplies for the rest of the magnets are found at the north end of the F0 Service Building. As mentioned earlier, H703, Q701, Q702, Q710, Q711, Q712, Q713, and Q714 are independently powered and controlled, whereas I:HV703 and I:Q703 power many magnets in series (Fig. 7-23). Their ramp tables are stored in 468 cards, but from a practical standpoint they are controlled from the applications page I68. Making seemingly innocent modifications (such as for the current values during extraction) actually requires changes at several levels of the ramp to assure continuity from cycle to cycle; I68 calculates these changes automatically.

The P1 (and A1) corrector power supplies are merely regulators that share the Corrector Power Supplies for the ring magnets (Fig. 7-29). The upstream trims, VT701 to VT709, use CPS6S, and those downstream use CPS6N.

A second load transfer switch, S2, allows the power supplies to switch back and forth between the P1 and A1 line supplies; I:LAM52 and I:V701 are exempt because the equivalent devices in the A1 line only need to run at 150 GeV.

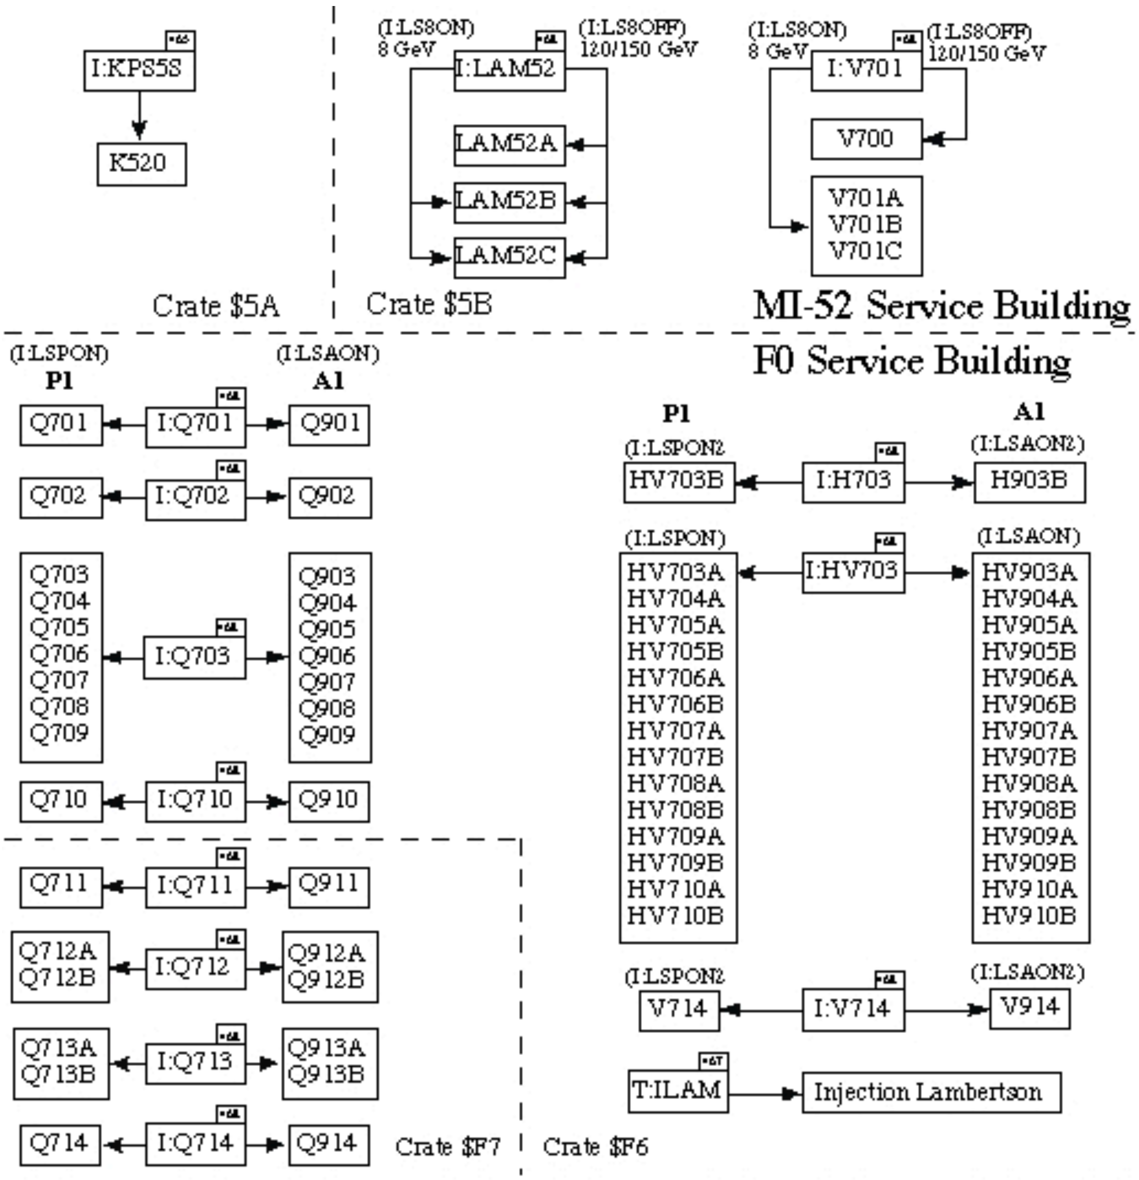

Most of the major power supplies in the P1 line lead a double life for one reason of another; the modes are changes via load transfer switches. At the beginning of the line, I:LAM52 and I:V701 are disconnected and reconnected to certain magnets, depending on the beam energy. Most of the remaining elements in the line must periodically share power supplies with corresponding magnets in the A1 line. I:Q703 is on when beam is to be injected into the Tevatron, and off when beam is destined for the P2 line.

The timer parameters (in parenthesis) tell the transfer switcher when to connect to the indicated magnets.

Fig. 7-23 P1 Line Power Supplies

Beam Diagnostics

P1 line diagnostics (Fig. 7-24) include BPMs, BLMs, multiwires, and toroids.

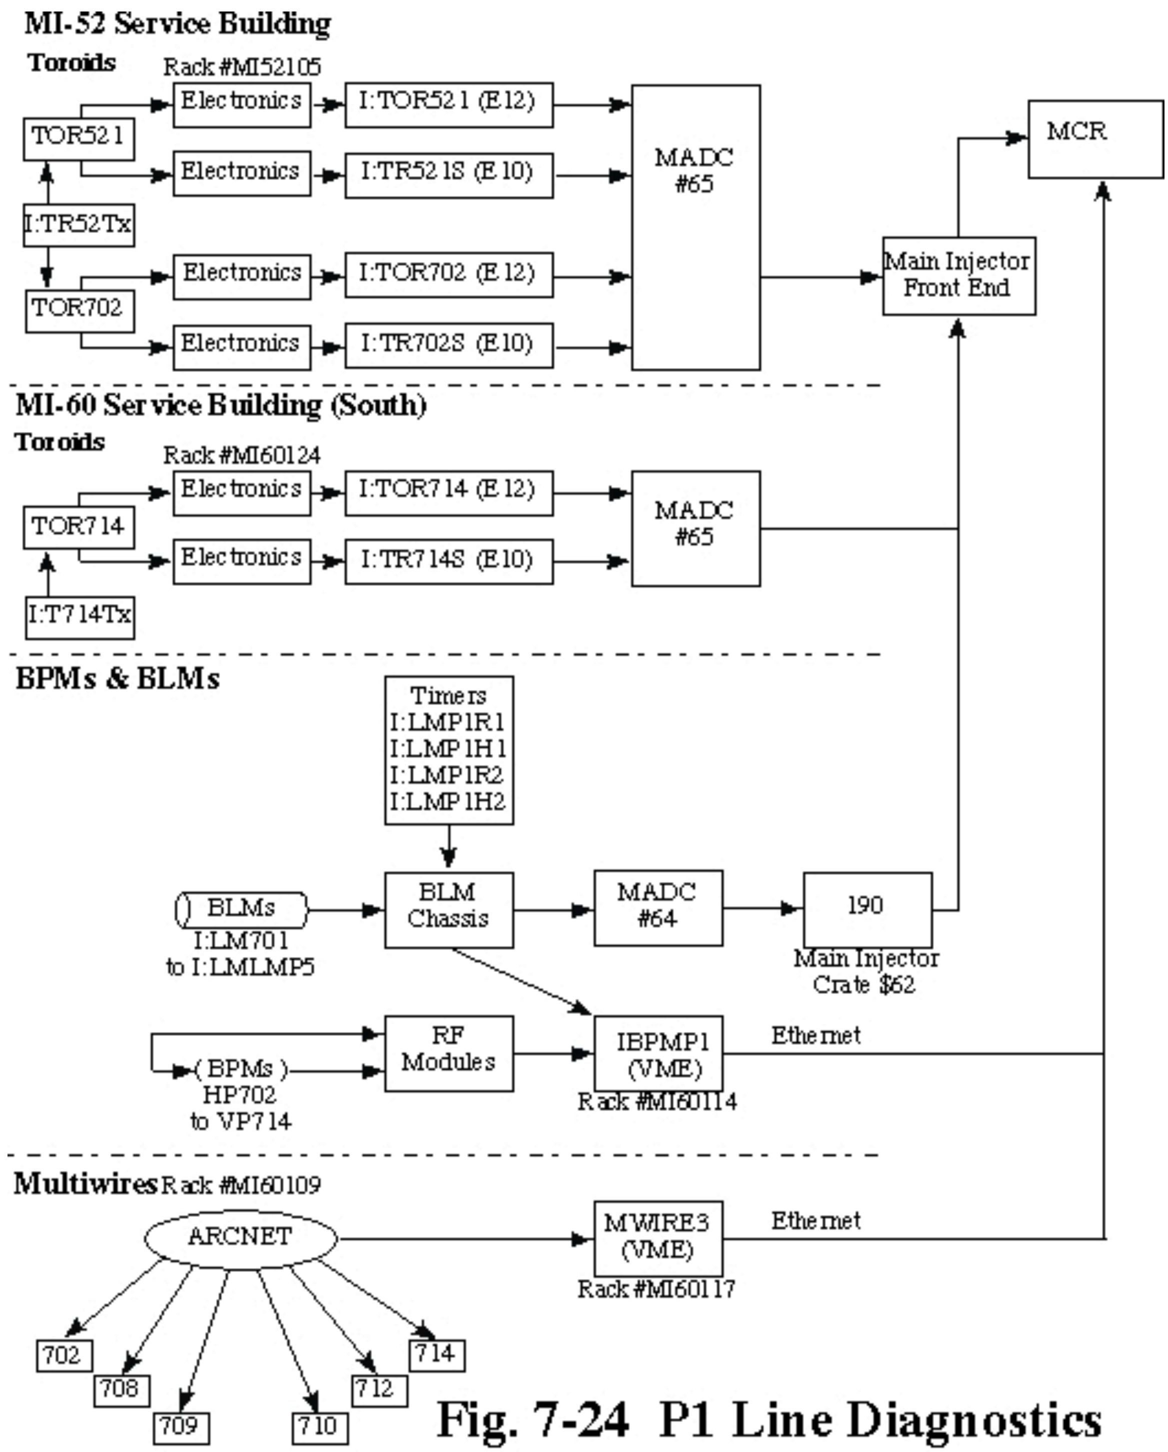

Fig. 7-24 P1 Line Diagnostics

Being a single-pass line, the only BPM option on I39 is the flash frame; choosing the cycle to display is a matter of choosing the proper beam sync trigger. Each beam sync clock event has a flash delay associated with it; the delays are all stored in 279 cards in Switchyard Crate $09 in the MAC room.

The delays can be found on I39 and are also listed in Table 7-1 (the last of the pictures at the end of this chapter). I39 can be used to view the P2, P3, and A1 lines as well.

The "Timing" window on I39 offers several options for setting the flash frame delay times:

(1) BPTRIGS opens a window that includes most of the beam sync events and flash delays associated with beam transfer to and from the Main Injector. For example, if looking at reverse proton tune-up to the Accumulator, I:BPFT7E—the Beam Position Flash Trigger for MIBS 7E—would be enabled.

All of the events in the BPTRIGS series are ORed together. Enabling an event and launching an SA will display a continuous update of positions in the beam lines. If more than one extraction event is selected, however, the events will overwrite each other. Since the selections are not console-dependant, altercations in the control room may develop as one person's data obliterates another's.

(2) BP1EXx, the P1 BPM beam sync triggers, are designed to prevent this conflict. They are triggered independently of BPTRIGS:

I:BP1EXT ($7D, $7E)
I:BP1EX1 ($7C)
I:BP1EX2 ($79)
I:BP1EX3 ($78)

Time delays in MREV can be chosen by the user. This option allows studiers—or automated programs—to collect flash data on particular events without interfering with the BPTRIGS frames.

(3) BPDELAY encompasses the parameters I:BPFTD1 and I:BPFTD2, which are user-controlled delays that can time the flash frame with respect to any TCLK event (as opposed to a beam sync event). In practice, BPDELAY is used more for diagnosing problems with circulating beam than it is for the beam lines.

BPMs in the P1 line interface with the controls system through IBPMP1, a VME crate at MI-60 south. The node can be rebooted through the parameter I:VRST61.

There are about two-dozen loss monitors in the P1 line. I:LM701 through I:LM714 are located near the quads Q701 through Q714, respectively. The density of BLMs is highest at the upstream and downstream ends of the line: at V701 (LMPC1U, LMPC1M, and LMPCD) and V714 (LMPC2U, LMPC2M, and LMPC2D—the "C1" and "C2" designations come from the fact that V701 and V714 are C-magnets). In addition, the Tevatron Injection Lambertson is instrumented with I:LMLMP1 through LMLMP5 ("Loss monitor" and "Lambertson" use the same abbreviation in these parameter names.)

The window of integration for the loss monitors in the P1 line is currently 5 milliseconds. The timing does not have to be incredibly precise; the only requirements are that the window be open when the beam passes through the line, and that the interval is short enough to prevent noise and background radiation from overwhelming the signal. The trigger and hold times are referenced to the TCLK events or reflections of the beam sync event (e.g. $80, $86, $5C, $94, $99). The reflection event appears only when the associated beam sync event is issued and a beam transfer is expected.

The reset time is set by I:LMP1R1 or I:LMP1R2; the hold time is set by I:LMP1H1 or I:LMPH2.

There are six multiwires in the P1 line: MW702, MW708, MW709, MW710, MW712, and MW714. They are controlled from page I41. The wires are controlled through a card in the MWIRE3 VME, located in the electronics racks at MI60S.

Since Main Injector acceleration cycles that use the P1 line come in many different varieties, the usual "Start, Sample, Stop" clock events are inadequate to set the timing for the multiwires. Instead, a 377 card at MI60S has been co-opted for that role. Fortunately, each cycle only uses the P1 line once, so each 377 timer can be assigned a delay referenced to the appropriate clock events. These timers can be found in the "External Events" box, to be enabled or disabled as needed.

The two toroids in the P1 line are TOR702 and TOR714. TOR521, which measures circulating beam at the point of extraction, can also be useful. Obviously, a comparison of the three is a measure of P1 line transmission efficiency.

Beam in the P1 line can be high intensity, such as during stacking, or low intensity, as during Pbar transfers. Therefore, the raw signal from each toroid is sent to two amplifiers: low gain amplifiers for I:TOR521, I:TOR702, and I:TOR714 that produce traces calibrated in units of E12, and high gain amplifiers for I:TR521S, I:TR702S, and I:TR714S, calibrated in units of E10 (the "S" stands for "small").

The triggers for TOR521 and TOR702 are based on I:TR52Tx, where x is 0-7; each value of x represents a TCLK event. (Like the BLMs, toroids can get away with the less precise timing of TCLK because all they need is an integration window that includes the beam pulse. The TCLK event is usually a beam sync reflection.) The series I:T714Tx triggers TOR714. Times for TOR714 are set about 23 microseconds later than the others, accounting not only for the time of flight of the beam down the line, but also for propagation delays of the TCLK events from the MAC room.

The parameter page I34 lists many timers related to beam diagnostics, including those for the toroids; most toroid timers can be found in the "8 GeV" sub pages.

LCW

Most of the magnets in the P1 line–-beginning with the C-magnets—are cooled by the stand-alone LCW system at the MI-52 Service Building. The setup is similar in many respects to a "regular" service building, details of which can be found in Chapter 4. Exceptions to the norm are discussed below; be sure to consult the graphics on page I56.

The system includes a storage tank; head pressure on the tank ultimately determines the pressure in the return line (V02 must remain open in order for the pressure to be transmitted). Water can be added to the system from the makeup tank by opening V03 and turning on the small makeup pump S06. Makeup water is also obtainable from the Main Injector system. Makeup water is added, as always, to the return side.

The supply pressure at MI-52, about 250 PSI, is higher than that in the rest of the ring.

The pond pumps are located a significant distance from MI-52, behind the MI-50 Service Building.

Vacuum

Vacuum is maintained with ion pumps. BV701 isolates the P1 line from the Main Injector. BV714 and BVF11 isolate the TeV Injection Lambertson.

P1 Line Permits

I:LAM52 and I:V701 are critical devices for F Sector and the Pre-Target Enclosure, as well as being the failure mode (backup) devices for the Tevatron CDC, in the unlikely event that T:ILAM should fail to trip off. There are also radiation monitors that will trip the critical devices. The critical device parameter name is I:P1INJ.

There is also a P1 line permit system, analogous to an abort system but only designed to inhibit beam. It is completely distinct from both the safety system and the Main Injector abort. All of the major power supplies for the P1 line magnets, as well as BV701, are tied in to it. When the permit is removed, the BSSB will instruct the Linac pulse shifter not to accelerate beam if it is destined to pass through the P1 line on its way to the Tevatron, Pbar, or Switchyard.

The P2 Line

The P2 line is the section of the Main Ring remnant that begins at the Injection Lambertson (ILAM) and ends with the B3 dipole at F17 (Fig. 7-25).

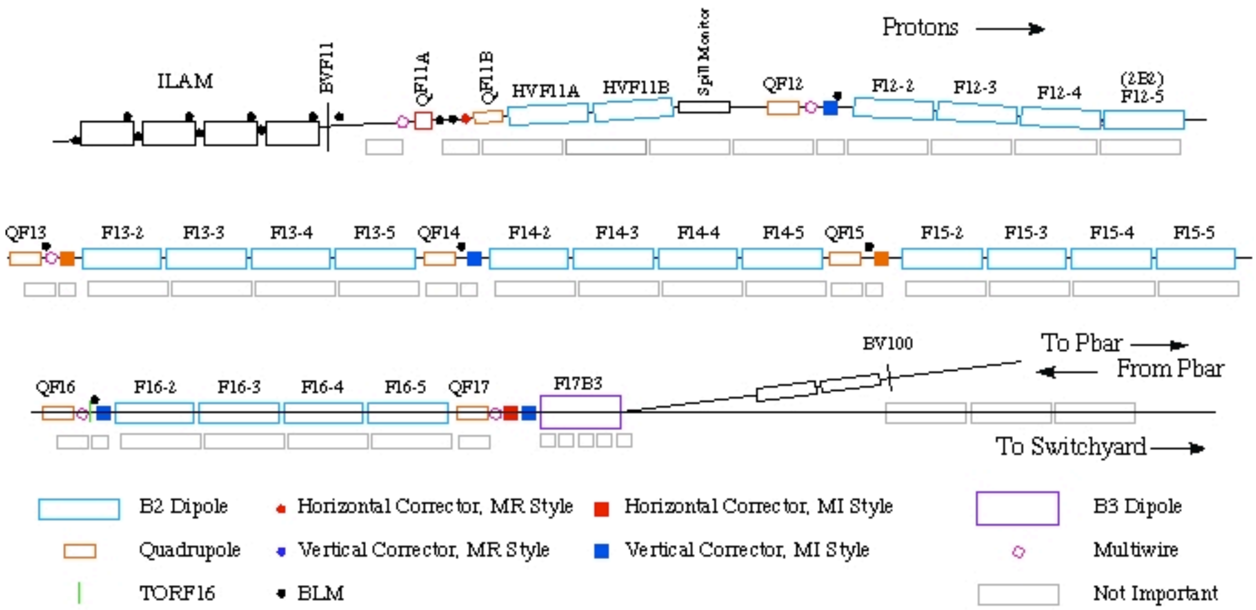

Fig. 7-25 Vertical Profile of the P2 Line

The P2 Line connects the P1 Line to F17B3 (compare to fig. 7-20, 7-26). If T:ILAM is not pulsed, the beam continues to rise until it reaches the two magnets at F11 (powered by HVF11). Those two magnets are skewed to provide a downward vertical deflection that cancels the upward angle of the beam. Skewed B2 dipoles at F12 then deflect the beam down to the level of the old Main Ring and finally remove the downward vertical angle. From QF13 to F17B3, the beam essentially traces part of a circle through the old Main Ring.

If F17B3 is powered, beam will be bent up into the AP-1 Line. If not, beam continues into the P3 Line.

The P2 Line can run at either 8 GeV or 120 GeV. At 8 GeV, it is probably being used for reverse protons, or for antiprotons arriving from the Accumulator. At 120 GeV, protons are being sent to the Pbar target or to the Fixed Target experiments.

It is used by 120 GeV protons headed toward the Pbar target, 8 GeV antiprotons from the Accumulator headed toward the Main Injector, or 8 GeV protons headed for the Antiproton Source for tune-up purposes (Figs. 1-4, 1-5, and 1-9). In the future, it will also be a bridge for 120 GeV beam to the P3 line and the Fixed Target experiments (Fig. 1-6).

The numbering/naming system has been inherited from the Main Ring, where the seven main quadrupoles defined the locations F11-1 through F17-1. (Unlike the Main Injector, the Main Ring had no consistent convention assigning even and odd numbers to focusing and defocusing quads. In fact, from F12 through F17, it turns out to be just the opposite.) Then there were four dipoles following each quad—for example, F14-2, F14-3, F14-4, and F14-5 followed F14-1. The quads have been renamed (e.g. F14-1 is now QF14), but most of the dipoles stay the same.

The original Main Ring lattice from F13-1 through F17-1 has remained (more or less) intact. The region through F11 and F12 has been reworked extensively to match the lattice and geometry of the P1 line.

Remember from the last section that beam in the P1 line is moving upward as it approaches ILAM. Then, if ILAM is pulsed, the beam is deflected downward so that it is horizontal, and it continues into the Tevatron. If ILAM is not pulsed, beam continues upward into the P2 line.

Most of the P2 line is 64.6 cm (25”) above the Tevatron, as inherited from the Main Ring. However, the geometry coming out of the P1 line is such that beam overshoots slightly. The first task of the P2 line is to bend the beam down until it is at the correct level. The B2 dipoles HVF11A and HVF11B, both powered by I:HVF11, straighten the beam to the horizontal, although it is still too high. The four dipoles at F12 then act as a dogleg to bend the beam down and straighten it horizontally at the correct level.

The dipoles at F11 and F12 are, for the most part, rotated. This is because the Main Ring was, of course, a ring, and the beam has to conform to the horizontal curvature even as it is bent vertically.

The last dipole in the F12 string is a double-strength B2 dipole (called a B2B); it serves to complete the vertical dogleg as well as provide bending for the horizontal curvature.

There are two major quadrupoles upstream of the HV11 dipoles—QF11A and QF11B. They are individually controllable.

The region from F13 to F17-1 is nothing more than the original Main Ring lattice, serving here to transport beam to and from F17. This is a FODO lattice much like that of the Main Injector, except that there are four dipoles for every quadrupole instead of two. There are five 84” quads between F12 and F17B3. The dipoles, from the Main Ring, are of the "B2" style. (Unlike the original Main Ring lattice, that used both B1 and B2 dipoles, all of the dipoles in the P2 line are the magnets with the larger vertical aperture.)

There is a “cradle” attached to each main quadrupole that can hold correction elements. In the P2 line, the original dipole correctors have been replaced by Main Injector style dipoles, except at F11. (As with the P1 line, LEP dipoles lack the agility required to track the ramps used in the P2 line.) The higher-order correctors such as the sextupoles, octupoles, and skew quads have been removed because this is a single-pass line.

At the end of the line is F17B3, a B3 dipole from the Main Ring. If the magnet is not powered, beam sails on into the P3 line and presumably on to Switchyard. F17B3 lies on its “side” and is primarily a vertical bend; it deflects protons into the AP-1 line. 8 GeV antiprotons approaching from the AP-1 line, such as during a shot, are bent onto the proper path in the P2 line. The magnet is therefore powered to either an 8 GeV or 120 GeV level.

Power Supplies

Magnets near the beginning of the P2 line, where tuning is the most critical, tend to have individual power supplies (Fig. 7-26). Further down the line there is a tendency to clump them all together. QF11A and QF11B each have a dedicated supply (I:QF11A and I:QF11B); HV11A and HV11B are both powered by I:HV11. All of the remaining main dipoles are powered from I:HVF12, and all of the remaining quads from I:QF12.

I:QF11A and I:QF11B are 75 KW Spang supplies located at the north end of F0. I:HVF11, and I:QF12 (which must power six quads in series), are 500 KW PEI supplies, also found at the north end of F0. The latter two supplies produce enough power to require their own dump circuitry.

Power Supplies	Cards	Magnets	Location
I:QF11A (75KW Spang)	468	QF11A	F0 North
I:QF11B (75KW Spang)	468	QF11B	
I:HVF11 (75KW Spang)	468	HVF11A HVF11B	Main Injector Crate $F7
I:QF12 (500KW P=E1)	468	QF12 (F12-1) QF15 (F15-1) QF13 (F13-1) QF16 (F16-1) QF14 (F14-1) QF17 (F17-1)	
I:HVF12 (Main Ring P.S.)	A468	F12-2 F13-2 F14-2 F15-2 F16-2 F12-3 F13-3 F14-3 F15-3 F16-3 F12-4 F13-4 F14-4 F15-4 F16-4 F12-5 F13-5 F14-5 F15-5 F16-5	F1 Service Building Main Injector Crate $1F
I:CPSF1	453s	I:VTF11 I:VTF14 I:HTF11 I:HTF15 I:VTF12 I:VTF15 I:HTF13 I:HTF17 I:VTF17	Main Injector Crate $5F
I:F17B3 (Main Ring P.S.)	468	F17B3	F2 Service Building Tevatron Crate $F6

Fig. 7-26 P2 Line Power Supplies
Compare to Fig 7-25

I:HVF12, which must power many large dipoles in series, is a modified Main Ring power supply located in the F1 service building.

The dipole corrector supplies, I:VTF11 through I:VTF17, are powered from I:CPSF1 and controlled from three 453 cards.

The supply for F17B3 is located in the F2 service building. I:F17B3, like I:HVF12, is a modified Main Ring supply.

LCW

The Tevatron LCW pump and the pond pumps at F1 are responsible for cooling the magnets in the P2 line and the I:HVF12 power supply. The cooling is shared with the TeV HOPS and corrector power supplies.

Vacuum

The P2 line can be isolated from the upstream world by BVF11, which is just downstream of ILAM (Fig. 7-25). BV100 can be closed to isolate the AP-1 line (indeed, this valve is sometimes inadvertently left closed after accesses). Although F17B3 is considered the end of the P2 line, the vacuum sector actually extends through F19; a third valve at F19 isolates the P2 line from the P3 line (not shown on Fig. 7-25).

The vacuum system for the P2 line has been inherited directly from the Main Ring; in fact, parameter names begin with M: instead of I:.

A roughing pump and a diffusion pump are connected to a pump down line at F15. The roughing pump is upstairs, between the motor control center breakers and the air compressor; the diffusion pump is downstairs at F15. (The concept of a diffusion pump has not yet been introduced in this book—basically, it heats and vaporizes oil; the hot droplets absorb air molecules and carry them away from the beam pipe. The diffusion pump does its work over approximately the same pressure range as a turbo molecular pump.)

Finally, there are ion pumps between each magnet; the supplies for the ion pumps are upstairs at the F2 Service Building. An LED on the front of each supply indicates whether the supply is on or off. Vacuum is monitored by thermocouple (TC) gauges. A CAMAC 145 card in the electronics room is the interface to the controls system.

More detail on the vacuum system can be found in the old Main Ring Rookie Book.

Beam Diagnostics

Beam diagnostics for the P2 line, and their controls links, are shown in Fig. 7-27.

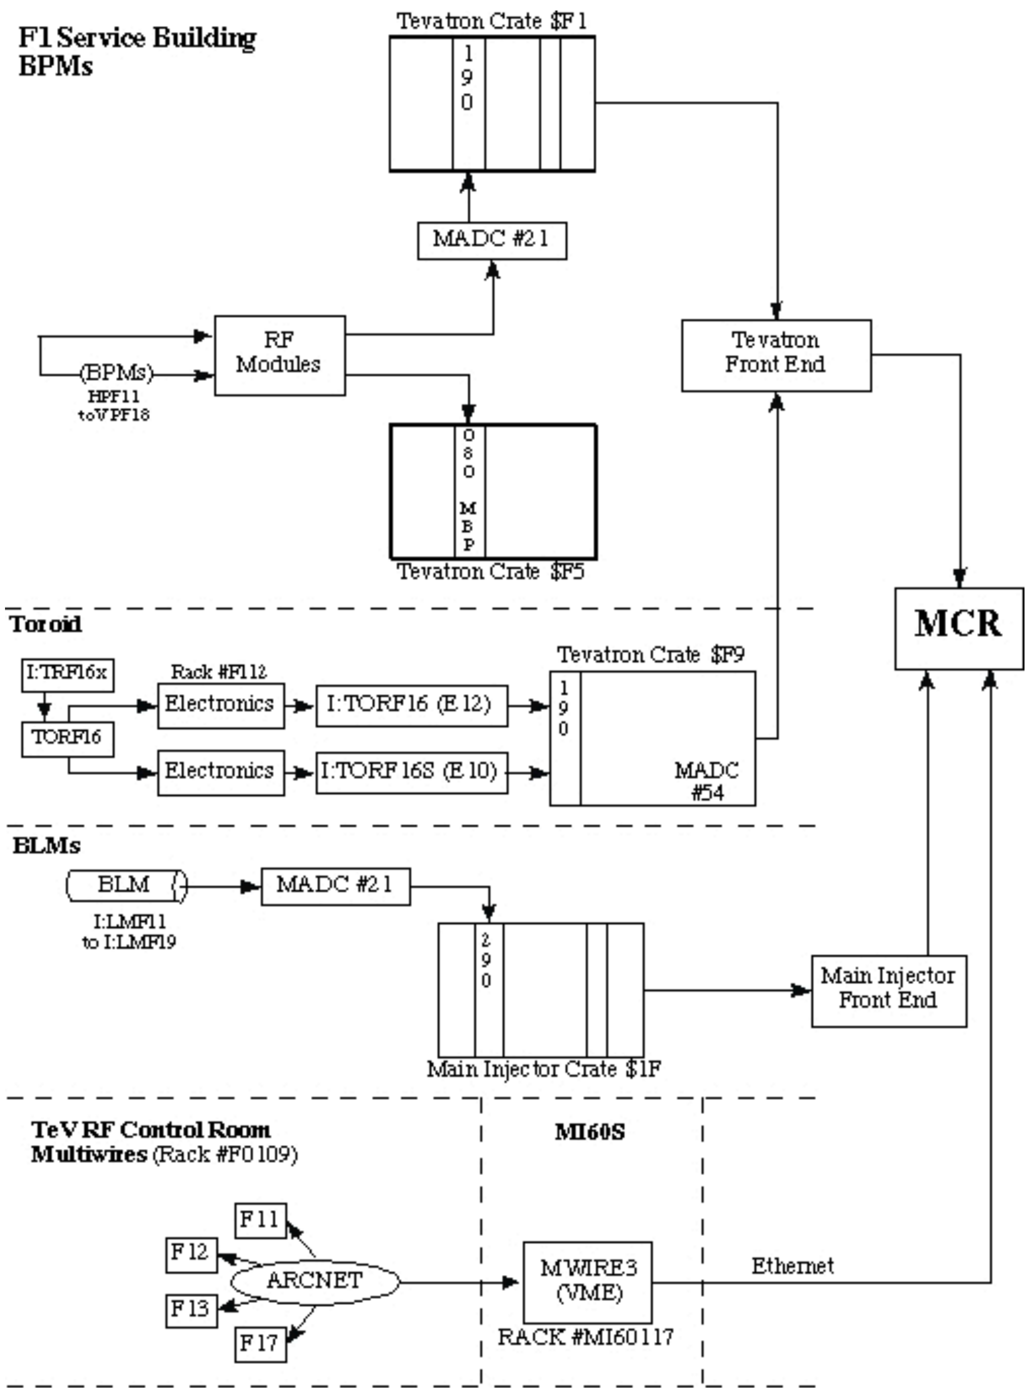

Fig. 7-27 P2 Line Diagnostics

The BPMs in the P2 line can be viewed on I39, either by themselves or in conjunction with the P1 or AP-1 lines. They rely on the same timers as the P1 line. The same is true of the BLMs.

There are five multiwires in the P2 line: (1) upstream of QF11A; (2) downstream of QF12; (3) downstream of QF13; (4) downstream of QF16; and (5) between QF17 and F17B3. The last one in particular is essential for establishing positions going into the AP-1 line. Like the P1 multiwires, the P2 multiwires are controlled from a card in the MWIRE3 crate.

There is one toroid in the P2 line, at F16 (Fig. 7-25). Like the P1 line toroids, the raw signal is sent to two amplifiers. I:TORF16 is scaled in units of E12, and I:TRF16S in units of E10.

The A1 Line

The A1 line has but one purpose: to extract 150 GeV antiprotons from the Main Injector and get them into the Tevatron (Figs. 7-28, 1-11).

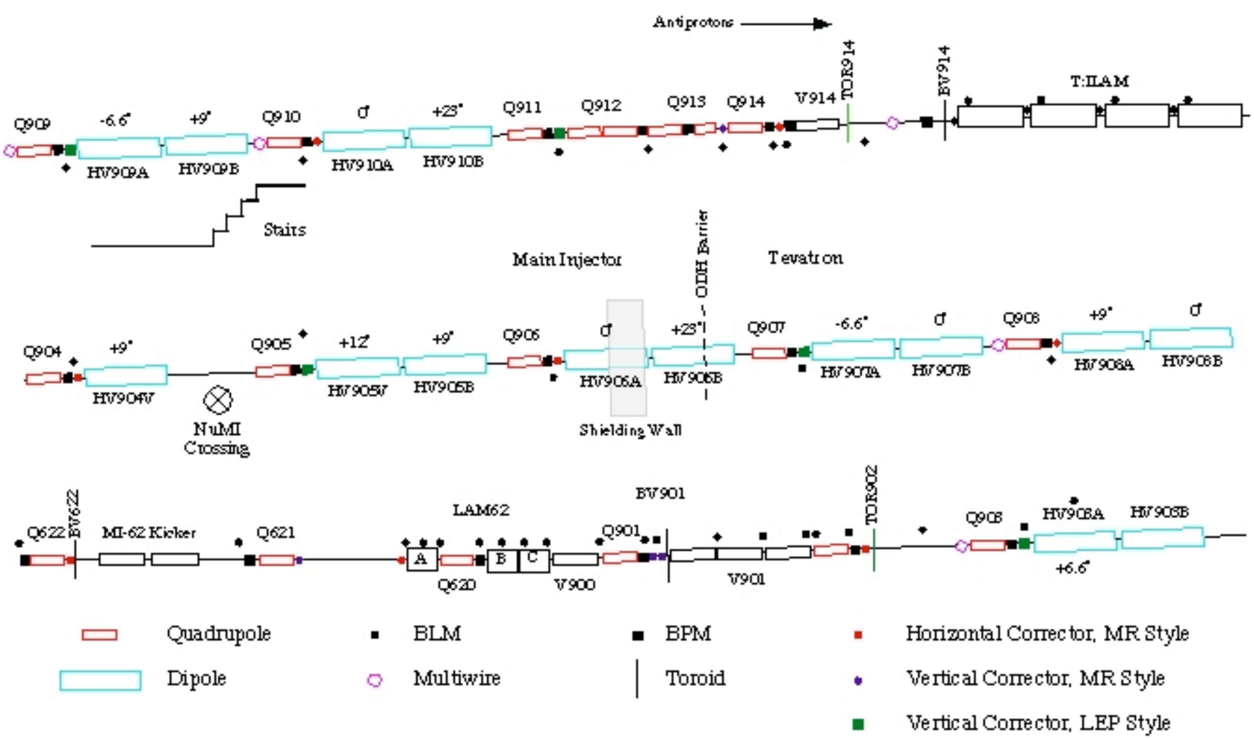

Fig. 7-28 Vertical Profile of the A1 Line

In many respects the A1 line is a mirror image of the P1 line, although the design is simplified by the fact that the beam has only one destination (the Tevatron) and one energy (150 GeV). (This simplicity is somewhat disrupted by the process of reverse injection, to be described below. But the sole purpose of reverse injection is to tune the A1 line for antiprotons at 150 GeV.)

Device names in the A1 line are brought to you by the number "9" (I:, that is).

During shots to the Tevatron, antiprotons are extracted from the Accumulator and sent down the AP3/AP1 lines. They enter the P2 line at F17B3 and transfer to the P1 line at the Tevatron Injection Lambertsons, which have to be off. After they enter the Main Injector from the P1 line, they are accelerated to 150 GeV and coalesced. The kickers at MI-62 fire to extract the beam into the A1 line. The A1 line, which looks remarkably like the P1 line, transports the beam to the Injection Lambertsons (which now have to be on) and on into the Tevatron. The E48 kicker cancels out the horizontal injection angle in the Tevatron. The Main Injector tasks in this process are initiated by a $2A event, and Tevatron events by a $40.

During shot setup, it may take several passes of beam to tune the A1 line. Rather than consume rare antiprotons during this time, reverse injection is used to send protons, going backwards through the line, to accomplish the same purpose. First, protons are accelerated and coalesced in the Main Injector and injected into the Tevatron via the P1 line. $2B and $4D events are used to load the protons, just as if setting up for a proton store. The protons circulate in the Tevatron for a short while; then the E48 kicker is fired to extract the beam. The beam enters the Injection Lambertsons, which bend the beam down into the A1 line. At the end of the line, the kickers at MI-62 are used to close the beam in Main Injector (instead of extracting the beam, as they would for forward antiprotons). A clean pass through the A1 line, and closure in the Main Injector, bodes well for a subsequent beam of antiprotons going the other way.

A $2A/$5D combination is used to orchestrate reverse injection. The $2A ramps the Main Injector to 150 GeV, and the $5D references such things as the transfer switches, kicker resets, and beam diagnostics. A completely new beam sync event is created as well—MIBS $D8. I:TMIPX sets the delay from TCLK $5D to generate MIBS $D8. The reflection event is TCLK $55.

The extraction kickers are just downstream (in the antiproton direction) of Q622. The power supply for the kickers, I:KPS62, is located in the MI-62 Service Building (Fig. 7-29). It always charges to a 150 GeV level. The unclamp parameter is K6AUCD. The beam sync delay timer for antiproton transfers, I:MITPBX, is referenced to TCLK $40, the Tevatron event dedicated to antiproton injection. At the end of the delay, MIBS Event $7B is issued, which sends the timing pulse to the Trigger Interface Module. The reflection event for MIBS $7B is TCLK $5B.

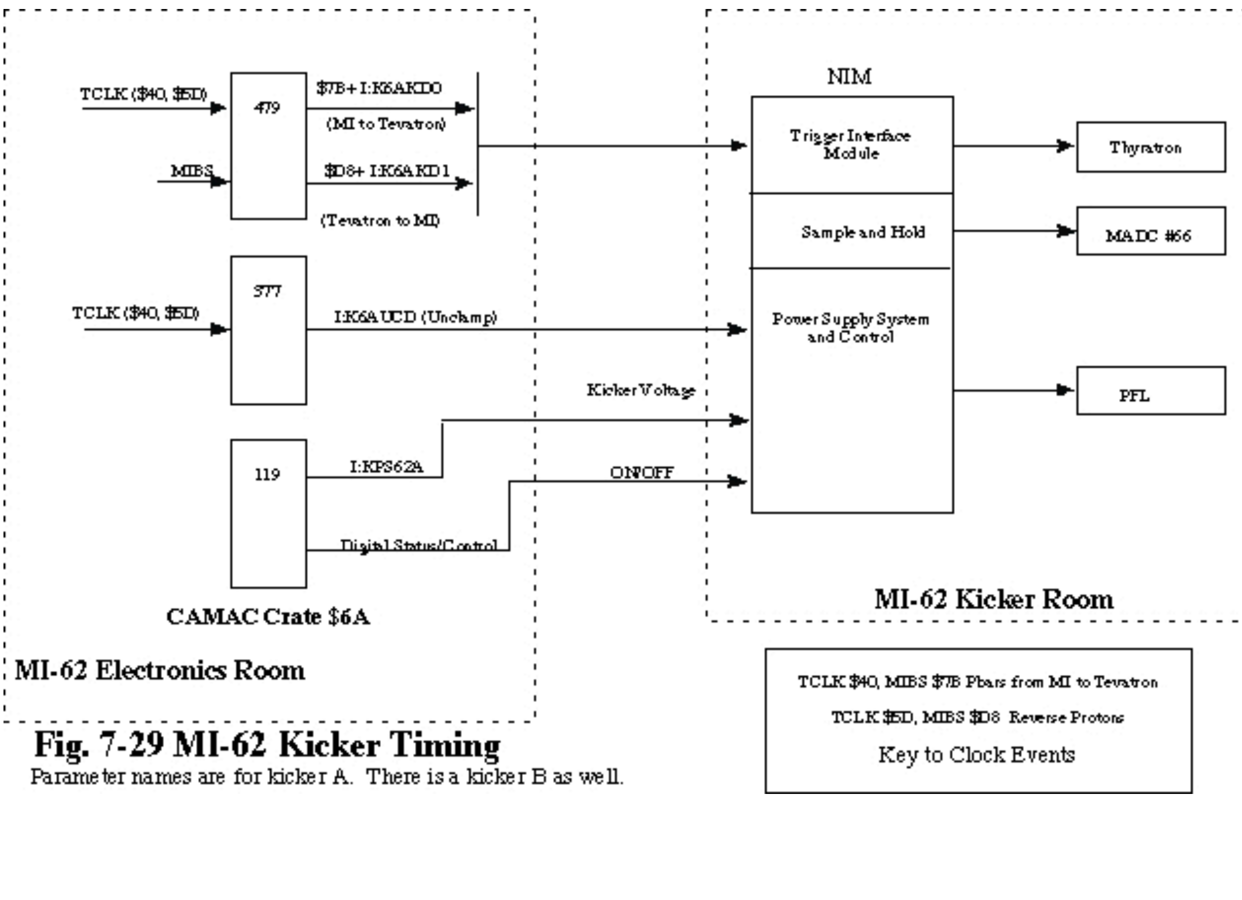

Fig. 7-29 MI-62 Kicker Timing
Parameter names are for kicker A. There is a kicker B as well.

The MI-62 Lambertsons form the usual 3-magnet cluster, this time around Q620. They bend beam up and out from the Main Injector. All three are powered.

The C-magnets, V900 and V901, are analogous to V700 and V701; V900 bends the beam up and the three magnets of V901 reduce the upward angle. (After the C-magnets, as in the P1 line, a series of 84" quads and rolled B2 dipoles takes the beam up and out toward T:ILAM. (The second would-be dipole between Q904 and Q905 is missing, in order to let the NuMI line pass through.) Toward the end of the line, Q911 through Q914 form a matching lattice to the Tevatron, which happens to be where the beam is headed. Then, as with 150 GeV protons, the Tevatron Injection Lambertsons are used to cancel the upward motion of the beam. Finally, the beam is closed horizontally in the Tevatron with the E48 kicker.

Power Supplies

The supplies for LAM62 and V901 are P=EI supplies. The extraction Lambertson is powered by I:LAM62. Miraculously, there are no transfer switches involved, because of the constant energy of extracted beam. I:V901 is actually powered by two P=EI supplies, I:V901A and I:V901B.

All of the magnets downstream of V901 are powered from the same supplies as their counterparts in the P1 line (Fig. 7-23). The alert reader will recognize a dilemma here—during shots, the P1 line magnets are needed to get particles in to the Main Injector at 8 GeV, but seconds later, the A1 line magnets are needed to transfer 150 GeV beam to the Tevatron. This is where a set of transfer switches is needed.

The switch between the two beam lines is sent from two parameters: I:LSPON and I:LSAON, in the case of quadrupoles, and I:LSPON2 and I:LSAON2 for the dipoles. For example, I:Q701 powers Q701 in the P1 line and Q901 in the A1 line. During shot setup, these parameters are referenced to TCLK $40 (forward antiprotons) or $5D (reverse protons), which are issued simultaneously with the $2A event.

The default state is for the P1 devices to be powered. As of this writing, antiprotons are sent to the Main Injector through the P1 line at about 1.6 seconds after the $2A. Then the 465 tables ramp down to zero output; I:LSAON and I:LSAON2 are issued 2.2 seconds after the $2A. At that point, the transfer switch physically disconnects the P1 line magnets from their power supplies and connects the supplies to the A1 line magnets. Beam begins accelerating around 4.5 seconds, is coalesced at flattop, and injected into the Tevatron via the A1 line at 6.7 seconds. At the end of the ramp, around 7.8 seconds, the tables again go to zero, I:LSPON and I:LSPON2 are issued, and the P1 magnets are reconnected, ready for the next cycle.

From the CAMAC perspective, the 468 cards servicing the P1 magnets also cover the corresponding magnets in the A1 line—Ramp "E" is usually referenced to a $40 and $5D to handle Pbar injection and reverse injection. As with the P1 line devices, page I68 is normally used to tune the line (except for closure).

Since the beam in the A1 line is always at 150 GeV, it is possible to use the high-inductance LEP dipoles; they are found at all of the vertical locations except for 901 and 913, which use Main Ring style dipoles. The horizontal correctors are all of the Main Ring style. The A1 corrector regulators are tied into CPS6N, as are some of the P1 correctors (Fig. 7-30).

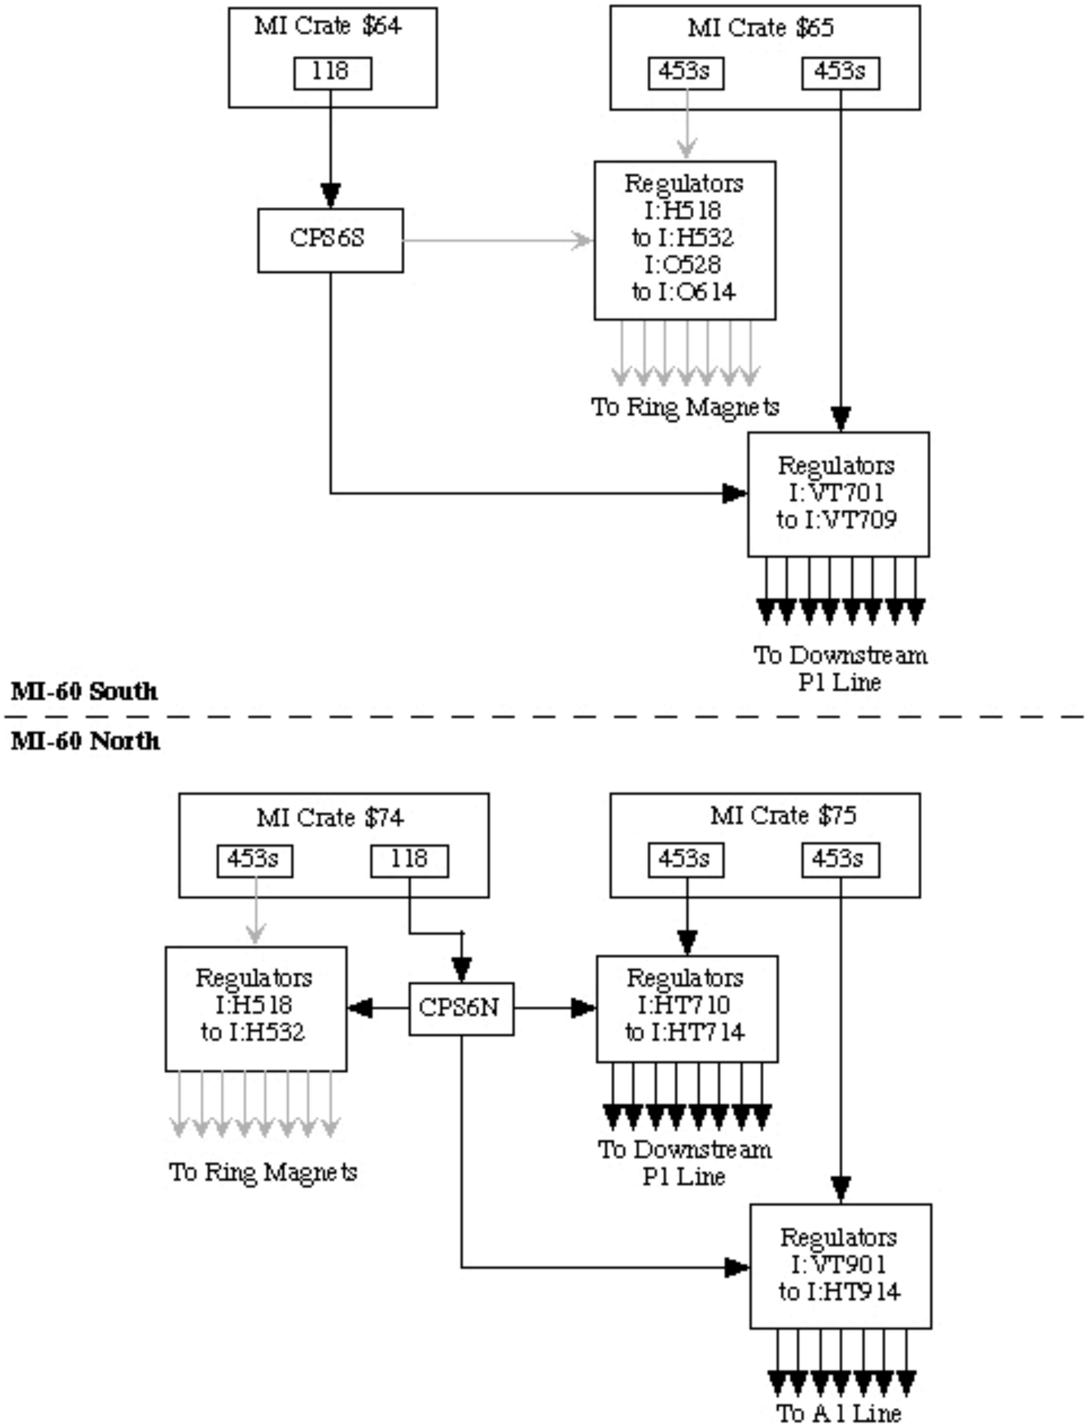

Fig. 7-30 P1 and A1 Corrector Dipole Power Supplies

The dipole correctors in the P1 and A1 lines use the same bulk power supplies as the ring correctors

LCW

Although the A1 line components are nearly identical to those in the P1 line, they do not require quite as much cooling. That is because they are only used during Pbar transfers and reverse injection, and at a relatively low rep rate. A stand-alone system like the one at MI-52 is not required. The LCW branches off of the ring headers at the C-magnets and begins feeding the A1 line components at Q903.

At the MI-62 service building, there is a branch coming up from the tunnel to feed the power supplies for V901 and LAM62.

Vacuum

The extraction region into the A1 line is bounded by three valves — BV622 and BV619 isolate the short straight section, while BV901 isolates the beam line just downstream of the Lambertson magnets. The downstream end of the beam line is isolated from the injection Lambertsons by BV914. Ion pump and Penning gauge status can be read under "Beam Lines/A150" on page I55. The vacuum hardware can be found at MI60N, under the name "150 GeV Line." Included are a CIA crate and ion pump power supplies.

Beam Diagnostics

The A1 line beam diagnostics (Fig. 7-31) are nearly symmetrical with those in the P1 line, but somewhat simplified by having only two modes of operation: antiproton extraction to the Tevatron, and reverse injection ($40 and $5D, respectively). They include:

- Multiwires (MW902 to MW914), which are controlled through an ARCNET loop originating from the VME MWIRE3.
- BPMs (HP902 to VP914), which report through the VME IBPMA1.

- BLMs (LM901 to LMAC2D), which can be read back through MADC channels; readbacks are also sent to IBPMA1 so that an integrated picture of losses and beam positions can be assembled.
- Toroids at 902 and 914, which are read through MADCs.

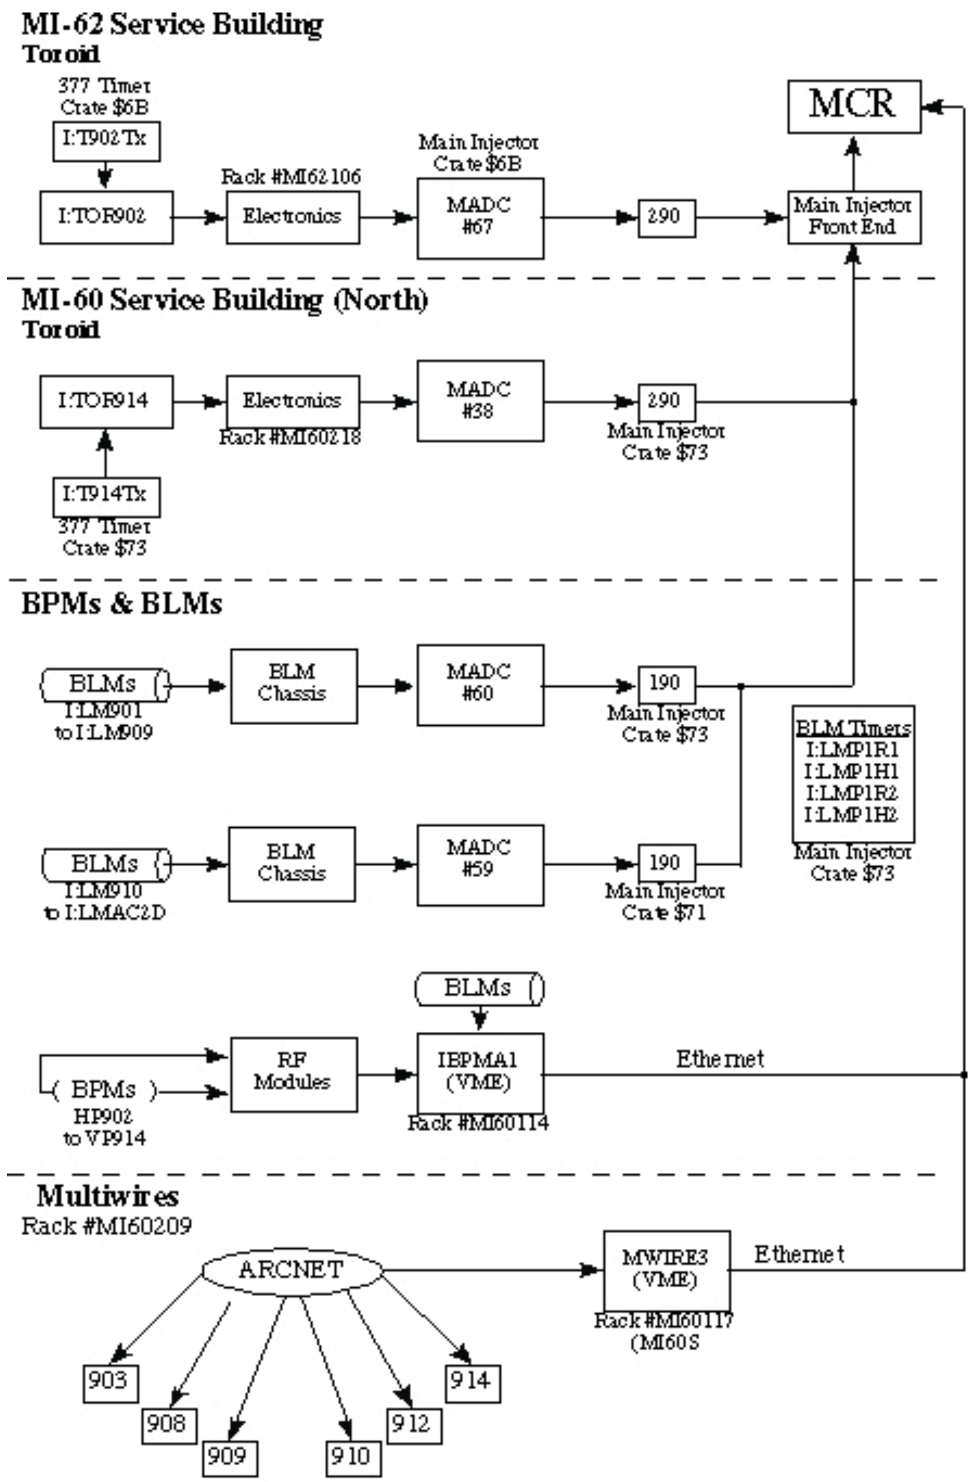

Fig. 7-31 A1 Line Diagnostics

The Recycler

Antiprotons from the Recycler enter the Main Injector at MI-22 and are extracted at MI-32; permanent magnet Lambertsons at 321 and 222 bends the beam appropriately. There is a single kicker at 304, in the MI-30 long straight section, that either kicks the beam into the field region of one of the Lambertsons or kicks the beam onto a closed orbit in the Main Injector, as needed:

[Embedded schematic of transfer]

There are, of course, also kickers and Lambertsons in the Recycler ring itself. For reasons of economy, the extraction kicker for antiprotons is also the abort kicker.

[Clock events]

NuMI

The initial extraction to NuMI is virtually identical to the standard used by the P1 and A1 lines. Extraction is from the MI-60 straight section—there is a stretch of space available just downstream of the coalescing cavities. There is a kicker just downstream of Q606, while the three Lambertsons are clustered around Q608 and Q609. Although the beam is eventually destined to plunge deep into the earth, where no beam has gone before, C-magnets raise the beam up over the Main Injector magnets so that it can be transported toward the NuMI stub.

Notes: