Accelerator Concepts

Table of Contents

Acknowledgements: In 1995 Cons Gattuso (with the help of Andrew Braun, Denton Morris, and Robin Spayde) wrote the original Concepts aid, version #1. Bruce Worthel updated the aid in 1997, version #2. Brent Evanger rewrote the Concepts book and turned it into a full-fledged Rookie Book in 2002, version #3. Pat Karns updated the Recycler information in 2006, version #3.1.

I. INTRODUCTION

A. Purpose of the book.

Learning about the various accelerators and subsystems found at the lab is a full time job. The intent of this book is to familiarize the new operator with some of the accelerator concepts that he or she will encounter again and again.

B. Characteristics of Fermilab accelerators

The Operations Department is responsible for the efficient running of a number of different accelerator systems: The Pre-accelerator, Linac, and Booster (collectively known as the Proton Source), Main Injector, Recycler, Tevatron, Debuncher and Accumulator. (These last two machines are referred to as the Antiproton Source). Operators are also responsible for operating the various transfer lines between the different accelerators as well as those between accelerators and experiments. In the next few pages, this Rookie Book will address the general characteristics of these machines.

Below, you will find a map of the FNAL site and a brief introductory description of each of the accelerators found here.

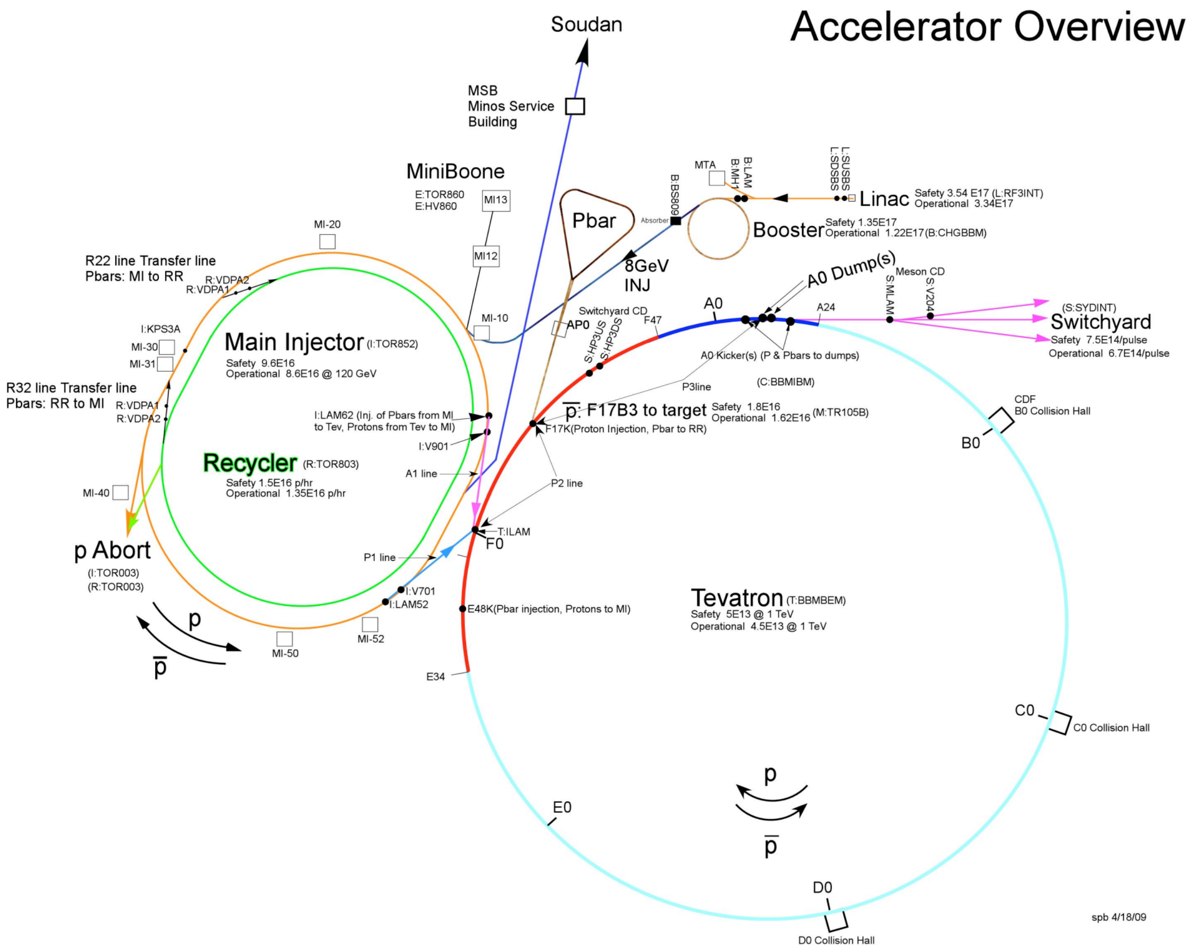

As an aid to understanding the terminology used to describe the beam energies reached in the various accelerators, it is useful to define the unit 'eV', or electron volt. One eV is the amount of kinetic energy given to a particle with the same charge as an electron crossing a potential difference of one volt. This unit is most useful for our purposes in much larger quantities; thus a series of semi-metric prefixes has been developed: KeV (Kilo-electron volt, 1E3 eV), MeV (Mega-electron volt, 1E9 eV), GeV (Giga-electron volt, 1E12 eV), TeV (Tera-electron volt, 1E15 eV). This last unit is the final designed beam energy of our largest accelerator: the TeV-atron (written simply as Tevatron).

1. Preacc

The Pre-accelerator, or "Preacc," is really the first accelerator. It is the source of the negatively charged hydrogen ions accelerated by the linear accelerator. The Preacc consists of the source housed in an electrically charged dome. The source converts hydrogen gas to ionized hydrogen gas (H^-). The dome is charged to a potential of -750 kV. The ionized gas is allowed to accelerate through a column from the charged dome to the grounded wall to an energy of 750 keV. The Preacc accelerates beam every 66 milliseconds (a 15 Hz repetition rate) whether beam is being requested or not.

There are actually two Preaccelerators at the beginning of our Linac (although only one is used at a time). They are called the H- source and the I- source. It should be pointed out that these names do not refer to the type of ion accelerated, as both machines are configured to deliver H- ions.

After beam exits the accelerating column, it travels through a transfer line called the 750 KeV line (referring to the transported beam's kinetic energy) and then enters the Linac.

2. Linac

The Linear Accelerator or "Linac" is the next level of acceleration for the negatively charged hydrogen ions. It takes the ions with an energy of 750 KeV and accelerates them to an energy of 400 MeV.

The Linac consists of two main sections, the low energy drift tube Linac and the high-energy side coupled cavity Linac. The drift tube Linac (DTL) makes up the first five RF stations. A large power amplifier tube (PA) powers each drift tube. These tubes amplify the 201 MHz RF signal used to drive the low energy cavities and accelerate the beam. The last 7 RF stations use Klystron amplifiers instead of the outdated tube technology of the low energy end. The Klystrons amplify an 805 MHz RF signal that is then fed into a series of side coupled cavity Linac (SCL) modules. The difference in these two resonating frequencies is a result of the differences in geometry between the DTL cavities and the SCL modules. In the DTL, every RF cycle is used to accelerate beam; in the SCL, only every fourth cycle is used (805 MHz = 4 x 201 MHz). Between the low energy DTL and the high energy SCL there is a 'transition section,' made up of the buncher and the vernier. These two additional klystron stations are used to ease the change in accelerating structure and RF frequency and improve the efficiency of the transfer. The Linac can accelerate beam once every 66 milliseconds (a 15 Hz repetition rate).

Beam in the DTL is focused by means of quadrupole magnets located inside the drift tubes, which in turn are located inside the RF cavities. The beam traveling through the SCL is focused by quadrupoles placed between the accelerating modules (outside of the accelerating cavities). After beam is accelerated in the Linac, the 400 MeV H- ions can be sent to one of three locations, depending on the need:

1) The 400 MeV line, a transfer line which connects the Linac to the Booster
2) The Linac dump, a large concrete block (the Linac actually has two separate dumps, one of them used only for studies)
3) The MTA beamline that leads to the MTA enclosure.

Linac modes of operation:

HEP

High Energy Physics, or HEP, is the mode used any time Linac sends beam to the Booster.

NTF

The Neutron Therapy Facility, or NTF, is a cancer treatment center located in the Linac gallery. NTF beam strikes a target with an H- ion beam from the Linac and directs the secondary neutrons toward cancerous tumors. The energy deposited by the neutrons bombarding the tumor kills the diseased cells. When NTF is treating patients (typically a couple of mornings a week), only a portion of the Linac is used to provide the primary H- ions. The DTL tanks 1 through 3 accelerate the beam to a kinetic energy of 66 MeV. The H- ions then coast through the tank 4 (without being accelerated) before special NTF-only magnets bend the beam 90 degrees to the left and out of the path of the Linac enclosure. This bend occurs between DTL tanks 4 and 5.

Studies

The Linac can run beam solely to the either of the two Linac dumps for various tune-up or diagnostic purposes.

MTA

The MTA enclosure houses muon cooling experimental apparatus and RF testing facilities that will use short pulses of Linac beam.

3. Booster

Booster is the next level of acceleration. It takes the 400 MeV negative hydrogen ions from the Linac and strips the electrons off, which leaves only the proton, and accelerating the protons to 8000 MeV, or 8 GeV. The Booster is the first circular accelerator, or synchrotron, in the chain of accelerators. It consists of a series of magnets arranged around a 75-meter radius circle, with 19 RF cavities interspersed. The accelerated proton beam in Booster can be directed to two different locations depending on the need. These locations are:

1) The MI-8 line (a transfer line from Booster to MI), which is located at Long 3 in the Booster tunnel.

2) The beam dump, also in the MI-8 line that uses 3 kickers to direct beam downward into an absorber.

Booster can accelerate beam once every 66 milliseconds (15 Hz).

Booster modes of operation:

HEP

High Energy Physics is a catchall meant to include all Booster cycles used by the Main Injector.

MiniBooNE

A follow up experiment to one run originally at Los Alamos, MiniBooNE (BooNE stands for Booster Neutrino Experiment) uses 8 GeV protons from the Booster to produce a neutrino beam aimed at their detector. The beam will leave Booster, head down the MI-8 line as if it were destined for the Main Injector, but then turn off at the last moment and head towards the MiniBooNE target.

Studies

Beam can be run solely to the Booster dump for various tune-up and diagnostic purposes.

4. Main Injector

The Main Injector (MI) is a circular synchrotron seven times the circumference of the Booster and slightly more than half the circumference of the Tevatron. The ring is divided up into 6 sections, or sectors, labeled MI-10 through MI-60. MI-60 is the region adjacent to the Tevatron. Main Injector has 18 accelerating cavities. It can accelerate 8 GeV protons from the Booster to either 120 GeV or 150 GeV, depending on their destination. When used to stack antiprotons or send beam to NuMI, the final energy is 120 GeV. When used to inject into the Tevatron, the final beam energy is 150 GeV. As well as accepting protons from Booster, the Main Injector can accept antiprotons from the Antiproton Source. The Main Injector can accelerate beam as fast as every 2.2 seconds.

As mentioned above, the Main Injector can provide beam to a number of different places at a number of different energies. The following list names main transfer lines associated with the Main Injector:

- ***The P1 line***. This line connects the MI to the Tevatron or to the P2 line, depending on the beam's final destination. The P1 line can accept 8 GeV protons or antiprotons coming to and from the Antiproton Source, 120 GeV protons going to the Antiproton Source target station or eventually to switchyard, or 150 GeV protons going to the Tevatron or, eventually, 'recycled' antiprotons coming back from the Tevatron to MI.
- ***The A1 line***. This line connects the MI to the Tevatron. The A1 line can accept 150 GeV antiprotons going into the Tevatron (or protons coming back from the Tevatron, usually for tune up).
- ***The MI abort***. The abort line is designed to accommodate protons only, and is located at MI-40. Beam of any kinetic energy can be sent to the abort.
- ***The P2 line***. The P2 line is not directly connected to the MI. It connects the P1 line to the AP1 line, the first of the transfer lines associated with the Antiproton Source. The P2 line can accept 8 GeV protons from the P1 line (going to AP1, usually for tune up or studies), 8 GeV antiprotons going to the MI (coming from AP1, usually for

shots to the Tevatron or Recycler), or 120 GeV protons en route to the target station (stacking).

- ***The P3 line***. Like the P2 line, the P3 line is not connected directly to the MI. The P3 line will accept 120 GeV protons traveling along the P1→P2 lines and guide them toward the switchyard and fixed target areas.
- ***The NuMI beamline.*** This beamline exits the Main Injector at MI-62 and travels downward to the NuMI Target Hall underneath MI-65.

Main Injector modes of operation:

Pbar Production

The most commonly used method is where two Booster batches are injected into Main Injector and then “slipped” together. This method is called “slip-stacking”, as the first injected batch is slowed down using the RF, and a second batch is injected and as it slips by the first batch they are merged together into one.

NuMI

Neutrinos from the Main Injector, or NuMI, is another neutrino experiment (like MiniBooNE) set to be running at the lab. For this mode of operation, up to 6 Booster batches are injected into MI and accelerated to 120 GeV. This proton beam is extracted near MI-62 and sent to the NuMI target hall. The resultant spray of secondary particles includes the particles that will eventually decay into the neutrinos the experiment wants. These neutrinos will be directed toward two experimental halls, one called the near detector (on site), and one appropriately called the far detector (in a mine in Soudan, MN).

Mixed-mode(NuMI + stacking)

The most typical running mode for Main Injector involves a beam cycle that send beam to both Pbar for pbar production and NuMI. In this mode two Booster batches for pbar production and 9 batches for NuMI and injected into Main Injector, accelerated then split to the two targets. This is done with a version of slip-stacking mentioned earlier. In this mode a stacking pulse plus four NuMI pulses are injected and slowed. Then another set of stacking plus four NuMI pulses are injected and slipped together with the first set. Once that is done a 9th NuMI batch is injected behind the now 5 slipped batches for a total of six batches. The first batch is extracted to Pbar and the remaining five are sent to NuMI.

Shot Setup

This mode relates to the act of extracting antiprotons from the Recycler, commonly referred to as a shot. Shot setup is the time before actually transferring the antiprotons when the various transfer lines are tuned up with protons to ensure efficient antiproton transmission. When loading the Tevatron with protons, 7 bunches are injected from Booster and accelerated to 150 GeV. A process called coalescing makes one bunch out of the 7 originals, and this coalesced bunch is extracted at MI-52 and travels down the P1 line and into the Tevatron. By repeating this process 36 times in a row you load the protons necessary for a

36x36 store. When loading antiprotons, 4 bunches are extracted from the Recycler, accelerated to 150 GeV in the MI, and extracted at MI-62, sent down the A1 line and into the Tevatron. This process is repeated 9 times to give a total of 36 antiproton bunches.

120 GeV Fixed Target

As the name implies, the final energy of this mode of MI operation is 120 GeV. One batch is injected and accelerated, then extracted at MI-52, traveling down the P1→P2→P3 lines and out to switchyard. Once in switchyard, the beam can be split (using electrostatic septa) and directed towards the various end users in the experimental beamlines.

Studies

Beam can be run solely to the Main Injector dump for various tune-up and diagnostic purposes.

5. Tevatron

The Tevatron is the largest of the Fermilab accelerators, with a circumference of approximately 4 miles. It is a circular synchrotron with eight accelerating cavities. The Tevatron can accept both protons and antiprotons from Main Injector and accelerate them from 150 GeV to 980 GeV. In Collider mode, the Tevatron can store beam for hours at a time. Because the Tevatron is a primarily storage ring, the length of time between acceleration cycles is widely variable.

The Tevatron is the only cryogenically cooled accelerator at Fermilab. The magnets used in the Tevatron are made up of a superconducting niobium/titanium alloy that needs to be kept extremely cold (~4 K) to remain a superconductor. The benefit of having superconducting magnets is the increased magnetic fields possible when high currents can be run through thin wires without fear of damage related to excessive resistive heating. This low operating temperature is responsible for the Tevatron's extensive cryogenic 'plumbing' and unique magnet protection systems.

The Tevatron is not a perfect circle either. The ring is divided into six sectors labeled A through F. Each sector has five service buildings, a "0" building and "1" through "4" buildings. Each "0" location contains a large straight section, and each such straight section has a special function. The A0 straight section is where the Tevatron tunnel connects to the Switchyard enclosure and is where the Main Ring remnant that is the P2 and P3 line connects to Switchyard. It is also the location of the beam abort for the Tevatron (the colliding beams abort). The CDF collision hall is located at the B0 straight section, while C0 is unused. The D0 experiment is named for the location it occupies in the tunnel, while E0's only claim to fame is that it was the site of the transfer line from the old Main Ring to the Tevatron. Perhaps the busiest section of tunnel at this laboratory is located at F0. This is where the Tevatron RF cavities are located, as well as the connection points of both the P1 and A1 transfer lines from Main Injector. The P2 transfer line also passes through, carrying beam to and from the Antiproton Source and protons on their way out to the Switchyard. Truly, F0 is the crossroads of the lab.

Tevatron modes of operation:

Collider

As mentioned above, the primary purpose of the Tevatron is to act as a storage ring where protons and antiprotons can collide with each other and produce interesting secondary particles. When operating in Collider mode, protons and antiprotons are injected at 150 GeV and then accelerated to 980 GeV. Once the final energy is reached, the two counter-rotating particle beams pass through each other for hours at a time (or until some component failure causes the beam to be lost). This stable situation of 980 GeV proton and antiproton collisions is called a **Store**. After the number of collisions per second (described by the ***luminosity*** of the store) drops too low to be useful for the experimenters, the store is ended and the Tevatron prepared for a new store.

Studies

This is a catchall referring to all sorts of operating modes used for specific tuning or diagnostic purposes. Two examples include proton-only stores (the protons are loaded and accelerated to 980 GeV but the antiprotons are not) and injection studies (the Tevatron stays at 150 GeV and protons are injected and sent to the abort repeatedly).

6. *Antiproton Source*

Target

The antiproton target station is not an accelerator. It is added here to maintain some continuity in the discussion of the antiproton source. The target station is found at the end of the AP1 transfer line, and is located beneath the AP0 service building. When we are stacking, 120 GeV protons coming from the MI through the P1→P2→AP1 lines strike a nickel alloy target. These high-energy protons striking the target produce a spray of all sorts of secondary particles. Using magnets to choose which momentum and charge we can collect 8 GeV antiprotons from this spray. These antiprotons are directed down the AP2 transfer line and into the Debuncher.

Debuncher

The Debuncher is one of the two synchrotrons that make up the Antiproton Source (commonly referred to as the Pbar Source). The Debuncher is a rounded triangular-shaped synchrotron with a mean radius of 90 meters. It can accept 8 GeV protons from Main Injector for beam studies, and 8 GeV antiprotons from the target station. Its primary purpose is to efficiently capture the high momentum spread antiprotons coming off of the target, using a RF manipulation called bunch rotation. There are also beam-cooling systems that act to make the beam more manageable. This so-called stochastic cooling is accomplished by picking up a signal from the circulating antiprotons on one side of the ring, amplifying the signal, and then applying that signal to the antiproton beam at another part of the ring. There are three cooling systems in the Debuncher: a momentum system, and two transverse systems (horizontal and vertical). The Debuncher does not 'accelerate' beam in the same sense as the other accelerators, but maintains the beam at a constant energy of 8 GeV. The antiproton beam can be transferred to the Accumulator via the D/A transfer line, located beneath the AP10 service building.

Accumulator

The Accumulator is the second synchrotron of the antiproton source. It is also a triangular-shaped synchrotron of radius 75 meters and is housed in the same tunnel as the Debuncher. It is the storage ring for the antiprotons; all of the antiprotons made are stored here at 8 GeV and cooled until transferred to the Recycler. The Accumulator has a number of different cooling systems: stacktail momentum, core momentum, and core transverse (horizontal and vertical). An 8 GeV antiproton beam can be extracted from the Accumulator and sent down the AP3 transfer line (which eventually meets up with the AP1 line) towards the MI.

Antiproton Source modes of operation

Pbar Production (Stacking)

120 GeV beam extracted from MI at MI-52 travels down the P1→P2→AP1 lines until striking the nickel target beneath the AP0 service building. Out of the spray of random secondary particles, 8 GeV antiprotons are taken down the AP2 line and into the Debuncher. The Debuncher cools the Pbar and then transfers them down the D/A line and into the Accumulator. The Accumulator further cools the beam and stores it until it's needed.

Shots

During shot setup, 4 bunches of antiprotons are extracted from the Accumulator and sent down the AP3→AP1→P2→P1 lines into MI, and into the Recycler.

Reverse Protons

This is a mode used during shot setup to tune up the transfer lines. This mode is also used during studies periods.

7. Recycler

The Recycler is an antiproton storage ring located along the ceiling of the Main Injector tunnel. The proposed purpose of the Recycler was to 'Recycle' the antiprotons from a Tevatron store, cooling them and storing them alongside those sent from the Antiproton Source. This was abandoned after early problems in Collider Run II. The Recycler now accepts transfers only from the Antiproton source and cools them further than the Pbar Accumulator is capable. The Recycler uses both a stochastic cooling system (like the Antiproton Source) and an electron cooling system.

Stochastic cooling is used to cool the beam in Recycler, but loses its effectiveness with higher intensities. Once above 200e10 antiprotons in the Recycler, Electron Cooling is required. Electron cooling works on the principle of momentum transfer between electrons and antiprotons. A highly concentrated, cool beam of electrons is driven at the same energy as the antiprotons, and laid overtop of the antiprotons. The resulting glancing collisions between electrons and antiprotons transfer some of the momentum from the "hot" antiprotons to the "cool" electrons. With enough electrons, a substantial longitudinal cooling force is produced by absorbing momenta from the antiprotons, allowing for more compact, brighter bunches to send to the Tevatron. These electrons are produced in a 5MV Pelletron, and guided through beamlines to a section of Recycler beam pipe in the RR-30 section. Once the electron beam has made its pass through the antiprotons, it is returned to the Pelletron to recover the charge.

The Recycler does not accelerate particles, but rather stores them at a constant kinetic energy of 8 GeV. There are three transfer lines associated with the Recycler:

1. *RR32 line*. Protons coming from the Main Injector (for tune-up or studies) travel through the RR32 line to the Recycler. Conversely, antiprotons heading out of the Recycler eventually destined for the Tevatron also use this line.
2. *RR22 line*. The RR22 line is the other transfer line connecting the Recycler to the Main Injector. This line is used for Pbar traveling into the Recycler, or protons traveling out (for tune up or studies).
3. *RR40 abort line*. The RR40 line is a proton-only abort line that directs tune-up protons to the same abort that the Main Injector uses.

Recycler modes of operation:

Stashing

The process of accepting pbars from the Pbar Source and cooling them to prepare for more transfers, and eventually once the stash is large enough to begin HEP shot setup.

Studies

8 GeV protons from the Booster are sent into the Main Injector first (but not accelerated) and then transferred to the Recycler. This is done as protons are easier to make, and less costly to lose for studies. This is typically done as a "check-out" of the Recycler systems after a Main Injector enclosure access.

Shot setup

After cooling the "stash", the antiprotons are "mined" into 9 "parcels". Each of these parcels is split into four bunches, which are extracted to the Tevatron, after acceleration in MI.

8. Switchyard 120 GeV

The Switchyard is not really an accelerator, but rather a complex crossroads where 120 GeV beam coming from the MI could have been directed to a number of final destinations in the three main fixed target beamlines: the Proton line, the Meson line, and the Neutrino line. Its configuration as of today (January 6, 2010) will only send beam to Meson.

Notes:

II. RF CONCEPTS

A. Methods of Acceleration

This chapter will deal with concepts related to Radio Frequency (RF) electromagnetic fields and their use in particle acceleration. This chapter will present an overview of RF acceleration including definitions of terms commonly used, a generalized view of RF amplification systems, and finally the considerations when transferring beam from one accelerator to another. But first, a little E&M background. This is the general equation of what forces an electric charge ‘feels’:

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

B. Why Use RF?

First of all, you cannot accelerate particles using a magnetic field. By looking at the above force equation and focusing on the magnetic contribution $(q\vec{v} \times \vec{B})$, it can be seen that if you take a charged particle at rest and place it in a magnetic field, the particle will experience no force because it’s velocity, v, is zero. Hence, the magnetic contribution to the accelerating force is zero. Even for a charged particle that is moving, the force created by moving through a static magnetic field is perpendicular to the direction of motion.

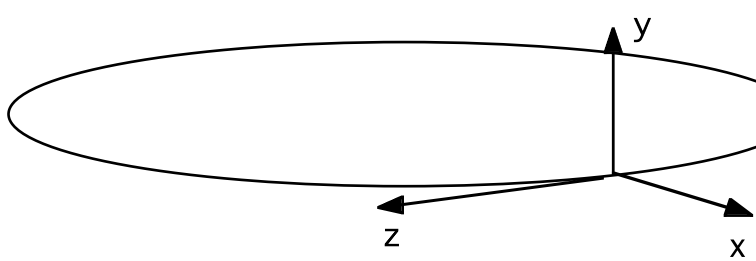

It is helpful to point out the coordinate system used in this chapter and many other discussions of particle trajectories. The coordinate system is a right-handed Cartesian system with the Z-axis pointing along the direction the beam is traveling, the X-axis pointing radially outward, and the Y-axis pointing vertically upward as in the diagram above. In the above example, if the magnetic field is in the Y-plane and the particle is traveling in the Z-plane, the force will be in the X-plane. To find out if this magnetic force will change the particle’s kinetic energy, we need to refer to the following equation:

$$U(r) = \int \vec{F} \bullet d\vec{r}$$

With the force in the X-plane and the particle traveling in the Z-plane, the dot product will be zero, and thus the energy gained by the particle (due to the magnetic field) is zero. This leaves only an electric field to accelerate a charged particle.

Gravity also acts upon the charged particles in question, so technically one could use gravity to accelerate them. In order to go from rest to a kinetic energy of 750 keV, an H- ion would have to be acted upon by the force of gravity over an extremely long distance.

The equation $U(r) = \int \vec{F} \bullet d\vec{r}$

becomes

$$U(h) = Fh$$

when the force, F, is in the h direction.

Setting U(h) = 750 keV, and F = mg (the force of gravity)

we get the equation 750 keV = mgh

where h is the height the particle needs to be dropped from.

Substituting in for some variables as seen below and solving for h, we find:

1eV = 1.602^{-19} J <u>thus</u> 750 keV=$1.204\text{x}10^{-13}$ J

m = $1.672\text{x}10^{-27}$ kg + 2*$9.11\text{x}10^{-31}$ kg = $1.6738\text{x}10^{-27}$ kg (mass of the proton and 2 electrons in the H- ion)

h ~ $7.34\text{x}10^{12}$ cm
(What does this mean? We would have to drop an H- ion from ~73 million km to yield a 750 keV kinetic energy—assuming a constant value for g.)

C. RF Cavities

So we are left with the fact that subjecting charged particles to electric fields is the only feasible way to increase their kinetic energy. The Cockcroft-Walton Preaccelerators at the beginning of the Linac create large static electric fields that the newly created H- ions are released into. By charging the Preacc dome to a potential of –750 kV and keeping the wall of the Preacc pit at ground potential, we produce an ion beam with a kinetic energy of 750 keV. To increase the final energy of the beam, we would have to increase the potential difference between the dome and the wall. The limiting factor in this approach is the phenomenon of electrical breakdown. If we have too large an electric field, a breakdown will occur. This breakdown of the field causes a lightning-bolt like arc and releases to ground the negative charges built up on the dome. What we want is some sort of structure inside that will efficiently subject the beam to moderate electric fields (which won't break down and cause sparking) many times in a row. In this way we can incrementally build up the particle's kinetic energy. What we want is an ***RF cavity***. An RF cavity is essentially a gap cut into the beam pipe across which an electric field is created. When the particle traverses the gap, it "feels" the electric field and gains kinetic energy. If we can place a number of these gaps in a row, we can build the particle's energy up to the desired level.

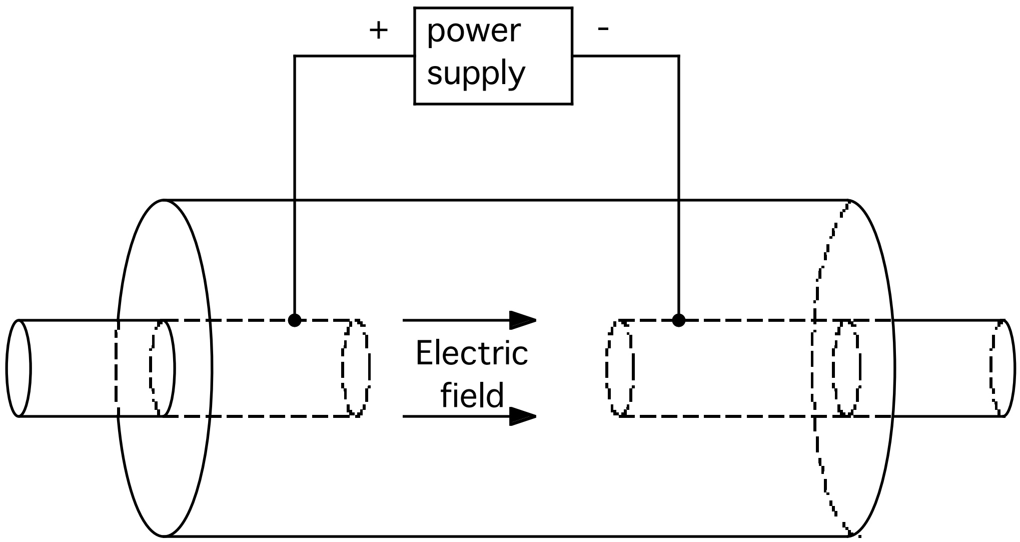

RF cavities are electrically resonant structures. They can be thought of as an R-L-C circuit (an R-L-C circuit is one with a resistor, R, an inductor, L, and a capacitor, C). Like the R-L-C circuit, an RF cavity has a specific natural resonant frequency. The cavity uses this resonant frequency to trade off the storing of its energy in the capacitive portion (in an electric field) and the storing of its energy in the inductive portion (in a magnetic field). Please note that in practice the natural resonant frequencies of most accelerating structures built lie within the **Radio Frequency** portion of the electromagnetic spectrum (hence the name RF). When driven with an AC power supply at the cavity's resonant frequency, an RF cavity is very efficient at storing the energy put into it. Designing the cavity appropriately allows the mode of this electromagnetic oscillation to be controlled. As noted above, electric fields and magnetic fields will be alternately produced within the structure. Choosing the appropriate cavity geometry allows one to dictate how those fields behave when they are formed within the cavity. For instance, in order to accelerate beam, we want the electric field to be produced along the axis of beam travel. Similarly, we want the magnetic field to form in such a way that it does not deflect the beam traveling through our cavity (see picture on next page).

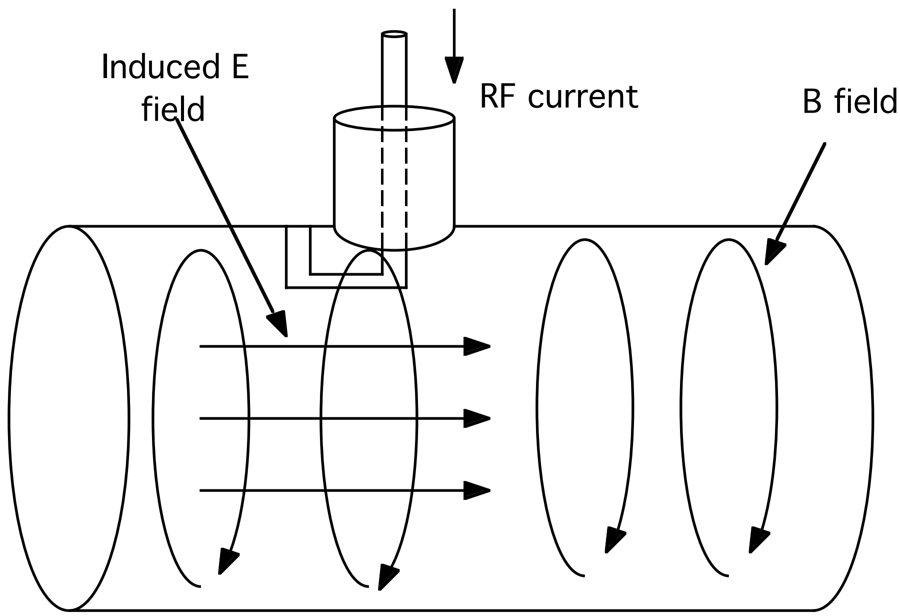

Notice how the magnetic field is in a direction that will not affect beam traveling along the lines of electric field.

During half of a period of oscillation, the electric fields will start off at a maximum value in one direction, shrink down to zero, then change signs and increase back to a maximum value in the opposite direction. During this same time, the magnetic field will start off at zero (all of the energy in the cavity is stored in the electric field at its maximum value). Then, as the electric field starts diminishing, the induced magnetic field increases to oppose the change until the electric field is zero and the magnetic field is at its maximum value (all cavity energy stored in the magnetic field). When the electric field changes direction after passing through zero, the magnetic field begins to decrease, etc. This is just like the mechanical example of the "mass on a spring" simple harmonic oscillator. Stretch the spring and you store energy in it. Release the spring and that stored energy converts itself to kinetic energy—the movement of the mass. The mass eventually passes through some maximum kinetic energy point (where the spring isn't stretched at all) and then slows down as it puts the energy back into the spring, and the cycle repeats. The same is true for the electric and magnetic fields oscillating inside an RF cavity.

The oscillating nature of these fields raises a good point. Imagine a positively charged particle entering the above (very long) cavity from the left-hand side. Given the field lines shown in the picture, this particle will feel a force to the right due to the electric field (its kinetic energy would be increased) and no effect from the magnetic field. Half a period of oscillation later, the electric fields would point towards the left-hand side of the page and the particle would feel a force to the left (its kinetic energy would be decreased). Because the particle would experience equal amounts of acceleration and deceleration, the net change in kinetic energy would be zero (see picture on next page).

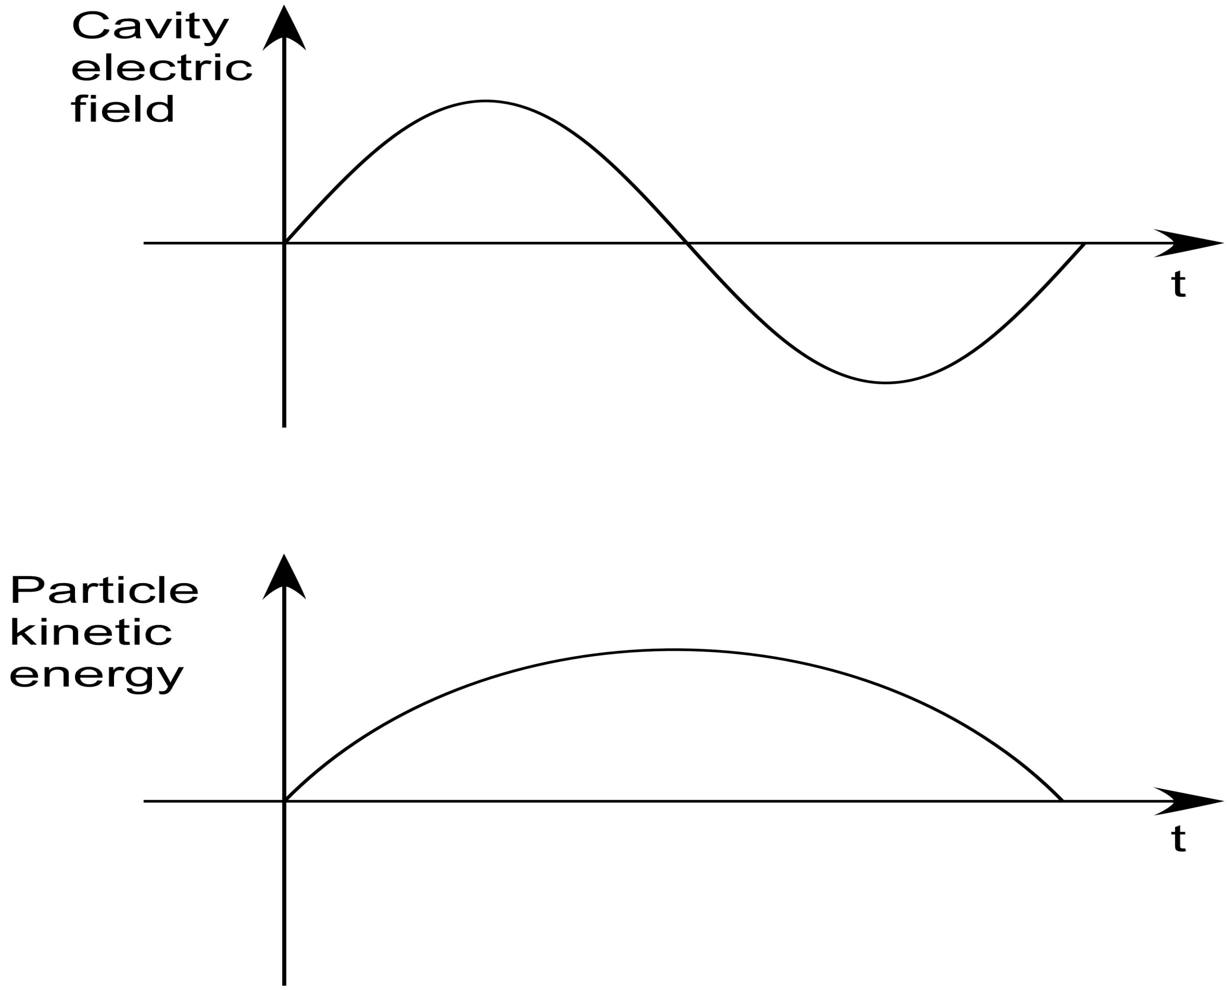

No kinetic energy gained. The sinusoidal curve on top represents the magnitude of the accelerating electric field in a RF cavity versus time. The other graph shows the change in particle kinetic energy versus time if the particle is present in the RF cavity for both accelerating and decelerating fields (net kinetic energy gain equals zero).

What we would like is a method of allowing the particles injected into our RF cavity to "feel" the electric fields that would accelerate them to the right, and not "feel" those pointing towards the left. If we did this we would be effectively breaking the RF cavity down into a series of smaller acceleration gaps and each time our particle traveled through one of these gaps its energy would be increased. There are essentially two methods of shielding particles from decelerating fields inside an RF cavity. One solution is called a ***drift tube***.

Drift Tubes

In order for the particle to receive a net increase in kinetic energy, it must feel a net force in the correct direction. In the graph above the overall change in kinetic energy is zero when the particle sent into the cavity is allowed to feel both an accelerating and a decelerating electric field. The particle needs protection (shielding) from the improper electric field.

We want to insert electrically grounded drift tubes, as seen in the picture below, inside the cavity so that when the electric field is 'wrong' the particle is hiding inside the drift tube. The particle will then emerge when the field is 'correct' again. While in the drift tube the particle is not accelerated, it simply 'drifts' (hence the name). When the particle emerges from the drift tube it experiences another accelerating force, and so on.

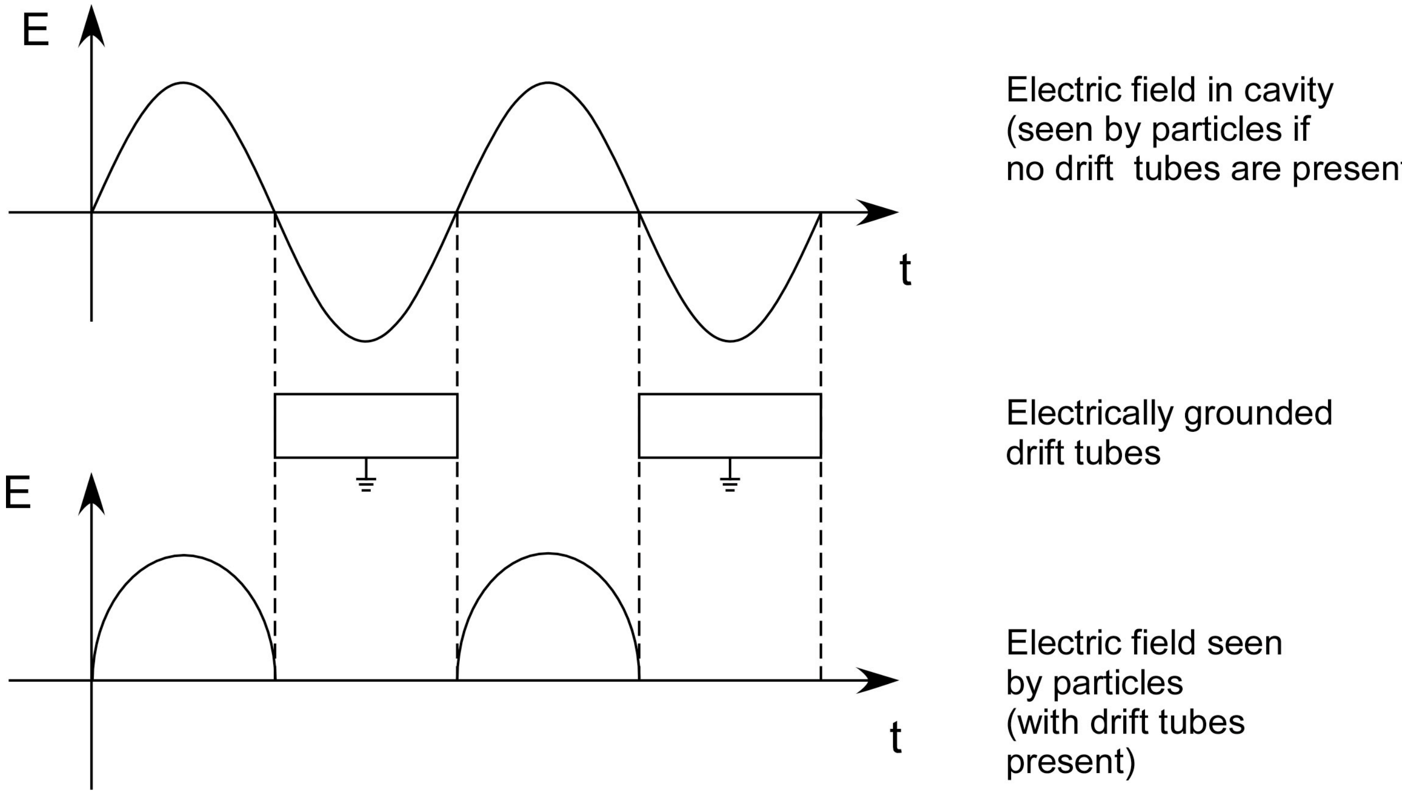

Electric fields with drift tubes. The graph on top represents the magnitude of the accelerating electric field in a RF cavity versus time. The bottom graph shows only that portion of the electric field seen by particles in a cavity equipped with drift tubes.

Using the above picture, if we inject a continuous stream of particles from the left-hand side, a time structure will be placed on the beam in the following way: particles that enter the cavity when the electric fields are correct get accelerated through the drift tube. Particles slightly behind in time will enter the cavity only to see an electric field in the wrong direction and will not be accelerated. Thus beam emerging from the right hand side of the cavity will be accelerated and formed into ***bunches***—discrete packets of particles—rather than a continuous stream. This bunching is a feature of all RF cavities. If there is no beam this empty area is called an RF ***bucket***. When a bucket actually has beam in it, we call it a ***bunch***.

Traveling through a drift tube equipped RF cavity allows the overall kinetic energy of the particles to be increased (as shown below)—the particles are shielded from those electric fields that would decelerate them.

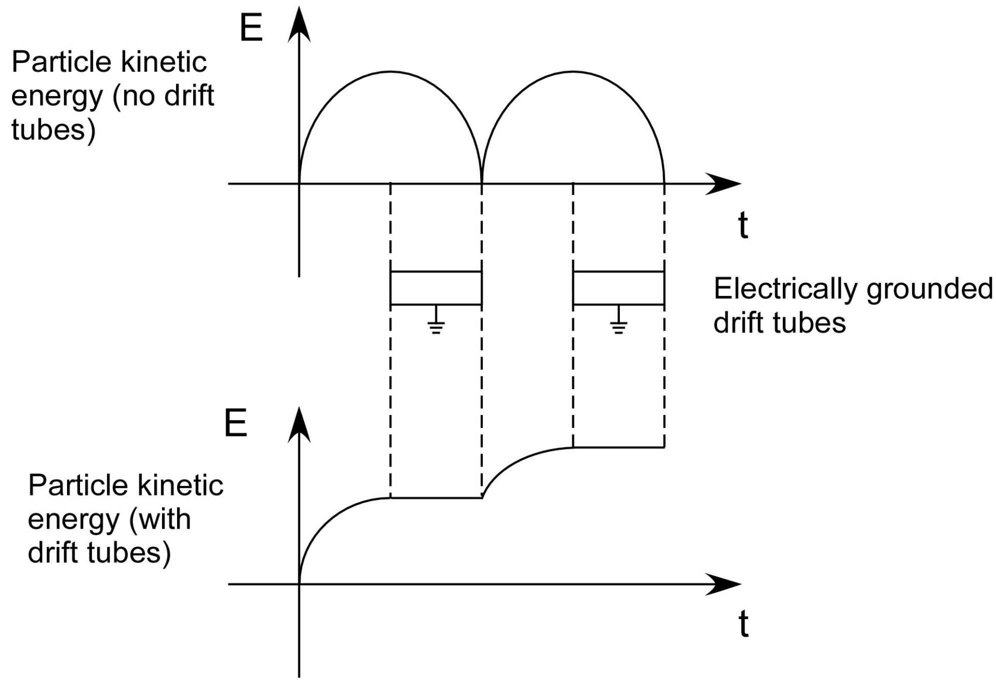

Kinetic energy gained. The top graph shows the change in particle kinetic energy versus time if the particle is present in the RF cavity for both accelerating and decelerating fields (net kinetic energy gain equals zero). The bottom graph shows a non-zero net increase in kinetic energy due to the shielding effect of the drift tubes.

The particles traveling through a drift tube equipped cavity operating with a fixed RF frequency must be shielded for a fixed amount of time during each RF cycle, namely one half of the RF's period. However, as the particle picks up more and more velocity, the time the particle takes to pass through each drift tube shortens. If all of the drift tubes within a constant frequency RF cavity were the same physical length, the particle will eventually arrive at the end of a drift tube while there is still a "wrong" electric field in the following gap. To combat this potential problem, the drift tubes can be made longer at the higher energy end of the tank giving the particle a longer drift space to traverse while the electric field is in the negative direction. As mentioned above, this method assumes that the RF frequency is fixed inside the tank. This is the way our low energy, or Drift Tube Linac (DTL), is arranged.

The picture below shows an exaggerated view of the increase in drift tube length found in the DTL. The arrows point in the beam direction — the left side is lower energy (slower) particles.

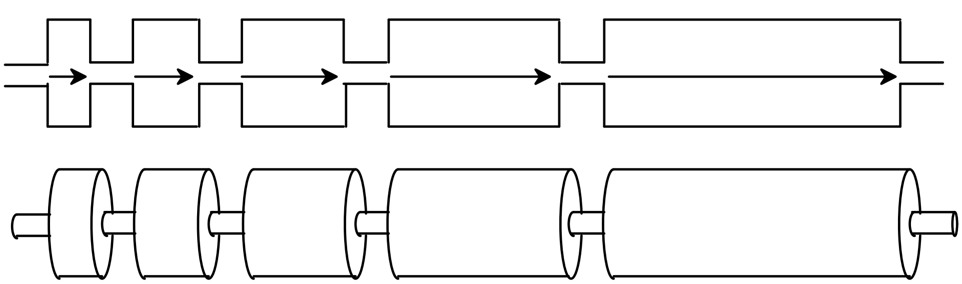

Constant Frequency with an Increasing Acceleration Gap Length

The drift tube solution to shielding particles from decelerating electric fields can only be used when the physical structure of the accelerator can be varied based on beam energy. For example, the energy of the beam (and the beam's velocity) increases as one moves further down the length of the Linac. The drift tubes found at the high-energy end are physically longer and would not work properly if moved to a different spot within the Linac. Each drift tube is custom made for it's specific location in the accelerator. It should also be pointed out that a DTL type cavity has many actual accelerating gaps and drift tubes within it (low energy Linac tank 1 has ~55).

The drift tube RF cavity is not suited to a circular accelerator. The beam traveling around a circular accelerator passes through each RF cavity thousands of times each second. Each successive pass the particles have a slightly different energy; and would require a physically different drift tube length. Obviously, radically changing the physical dimensions of a copper RF cavity thousands of times per second is impossible. We need to try something else.

Changing RF frequency

As the beam accelerates in a circular machine, it arrives at the RF cavity successively faster with each lap. Changing the RF frequency found inside the cavity as the particle speed increases makes the accelerating electric fields appear within the cavity more frequently. This method is used in the RF cavities of our circular accelerators where varying the size of the RF cavity would be impossible between successive beam revolutions.

In practice, one can rapidly vary the R-L-C electromagnetic properties of the cavity that will in turn vary the frequency at which the cavity resonates. The use of an increasing RF frequency timed such that the beam always 'feels' the correct force is one of the defining characteristics of a ***synchrotron***. A synchrotron is a circular machine where the RF frequency and the strength of the magnetic bending field are *synched* to the beam momentum and revolution frequency.

Phase

It is useful at this point to describe a term used to describe the status of the oscillating fields in the RF cavity at the time the beam arrives. The ***phase*** of the beam is essentially like the timing of the particles relative to the RF oscillations. When particles are injected into an RF cavity for acceleration they must have the right phase—they must arrive in the acceleration gap when the correct fields are present.

As noted above, if one injects a continuous stream of charged particles into a cavity only those with the correct phase get accelerated; the other particles are lost. If the beam is already formed into bunches, the phase of the beam is critically important. The beam transferred between the Booster and the Main Injector is bunched, so any beam arriving in the Main Injector's RF cavities too far out of phase will be lost.

The ***synchronous phase*** corresponds to the perfect proton with the perfect energy that arrives in the RF tank at the perfect time and receives the perfect incremental change in kinetic energy. It should not be surprising to find out that not all of our particles are perfect. Real bunches of particle beams made up of real collections of particles do not have one phase, but rather a distribution of phases centered (hopefully) on the synchronous phase.

Consider the picture below:

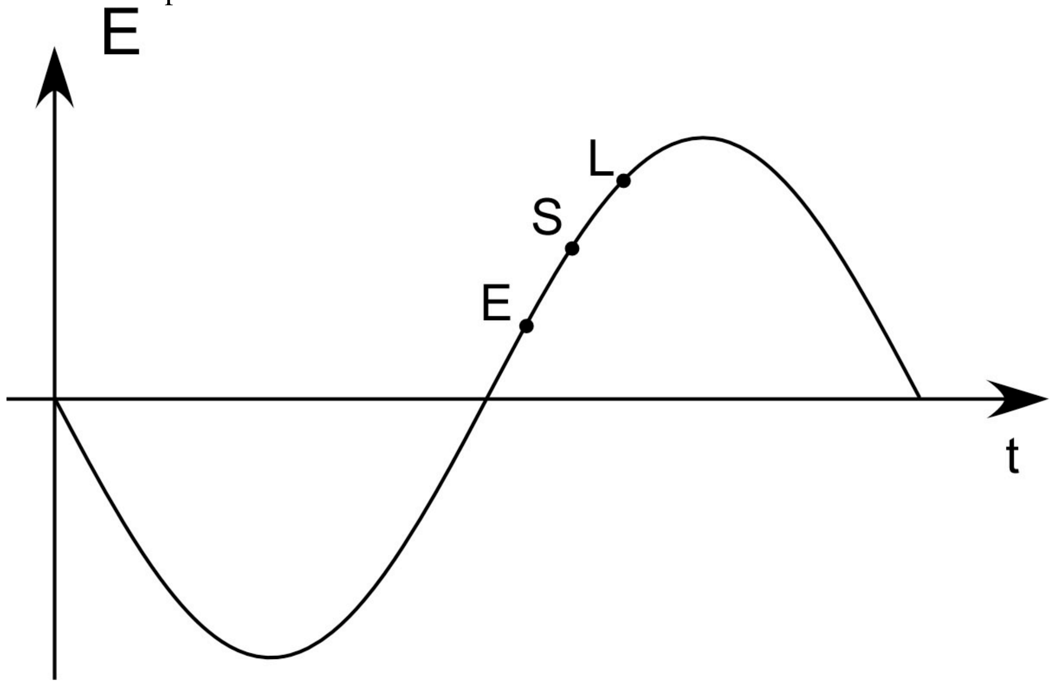

Acceleration. The sinusoidal curve represents the magnitude of the accelerating electric field in a RF cavity versus time. The labeled points represent various beam phases.

Point 'S' on the graph corresponds to the synchronous phase. Point 'E' indicates particles in a bunch that arrived in the RF cavity slightly earlier in time. Point 'L' indicates particles in the same bunch with a slightly later arrival time. Notice how the magnitude of the oscillating electric fields is not constant across the length of the bunch. The particles that arrived earlier in time—at point E—receive less of an increase in kinetic energy than those that arrived 'on time' at the synchronous phase, point S. Conversely, those particles lagging behind slightly—at point L—receive a larger boost in energy than the synchronous particles. This is seemingly a problem, until one considers why these particles are too fast or too slow. Those particles that arrive at the RF cavity too early do so because relative to the others in their bunch they have slightly too much energy. Those arriving too late have slightly less energetic. Thus, another benefit that arises from the uses of RF accelerating cavities is ***phase focusing***. Imperfect particles with imperfect energies are corrected slightly with each pass through the RF cavity. A particle at point E gets a smaller incremental change in energy than the others on this pass and will slowly move towards point S. Similarly, the slow particle at point L will arrive at an earlier time on the next pass and will move slowly towards point S as well. These initially off-energy particles will trade places with each other eventually and thus be caught in energy-versus-phase oscillation around the synchronous phase. These motions are known as ***synchrotron oscillations***. When undergoing this type of movement, individual particles inside the bunch are slowly moving from the rear of the pack to the front and back again. (For the Main Injector at 8 GeV the frequency of synchrotron oscillations is about 500 Hz while the revolution frequency of the beam traveling around the ring is about 90,000 Hz.) Motion like this taking place parallel to the direction of beam travel is called ***longitudinal***. Motion taking place perpendicular to the beam direction is called ***transverse***.

An alternate way to look at the information presented on the acceleration plot is to use a phase space diagram. In the above plot, we are comparing time to the magnitude of the electric field present inside the RF cavity when the particle arrives.

In the phase space plot below we compare "ΔE", the difference between a particle's energy and the energy of the synchronous particle, and "ΔΦ", the difference between a particle's phase and the phase of the synchronous particle (like the timing of the beam's arrival).

Like the example above, a particle at point 'E' is out of phase, gets a kick, and eventually winds up at point 'L'. In the phase space diagram, we see that a particle at point 'E' is initially out of phase with the synchronous particle at point 'S'. That error in phase becomes an error in energy as the particle follows the elliptical path to point 'L', where the energy error has again turned itself into a phase error. Particles following these elliptical orbits around the synchronous phase and energy are undergoing synchrotron oscillations.

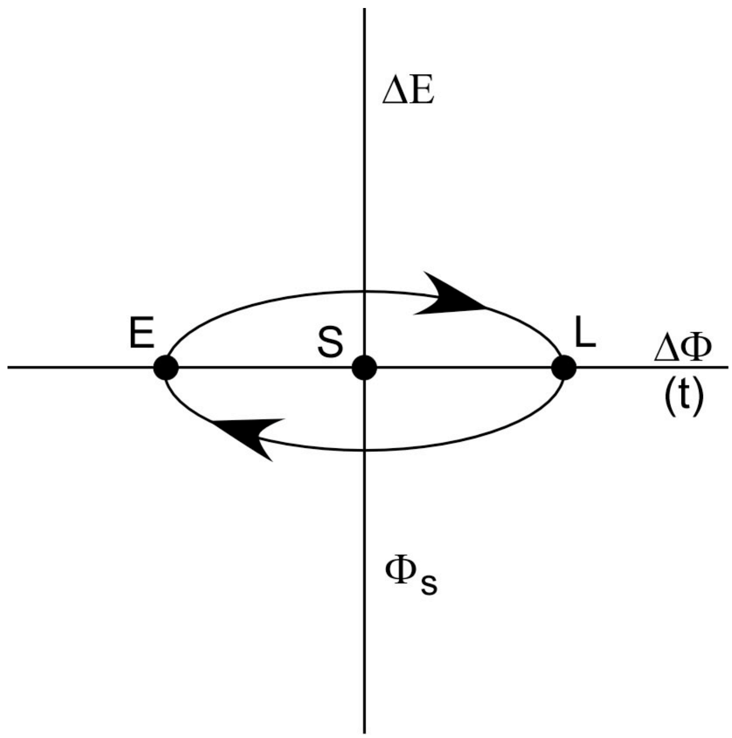

Phase Space Diagram. Here is a good way to view synchrotron oscillations—phase errors become energy errors become phase errors and so on. The labeled points represent various initial beam phases.

Did you now that protons from Booster are not accelerated immediately upon entering the Main Injector? How can this be when the Main Injector's RF cavities are on when the protons are injected? Shouldn't there be electric and magnetic fields present within the cavities? Even though there are fields present we can set things up so that the oscillating electric fields in the RF tanks are passing through their zero point when the beam arrives. Thus the protons passing through the tanks "feel" no net force and are not accelerated on average. Consider this picture:

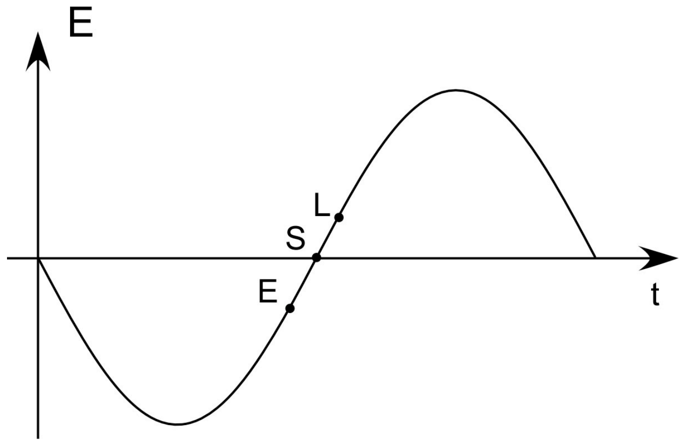

Injection. The sinusoidal curve represents the magnitude of the accelerating electric field in a RF cavity versus time. The labeled points represent various beam phases.

Just because the newly injected protons "feel" no net accelerating force, this does not mean that the RF does not act them upon. Notice how the particles at point 'E' (traveling too fast) are slightly decelerated while the particles at point 'L' (too slow) are slightly accelerated. Particles with the synchronous phase arrive at point 'S' and see no accelerating fields. The RF has actively captured and imposed a time structure via phase focusing on the beam (bunching), even though no net acceleration is taking place.

After the beam begins coasting around the Main Injector at a fixed 8 GeV, we can slowly change the RF phase with relation to the beam. We create a situation where electric fields *are* present in the cavities when the beam arrives, similar to the Acceleration picture from the previous page. Now the protons feel a net force and the acceleration process begins.

Another term used in reference to RF acceleration in circular machines is the ***harmonic number***. The harmonic number is essentially the number of RF buckets a circular accelerator contains. (This number is not used in the Linac.) Consequently, the harmonic number is also the maximum number of bunches of beam a machine can accelerate at once. The harmonic number of the Booster is 84; thus 84 bunches of beam represent what's called a ***batch***. A batch is simply one complete load (84 bunches) of beam from the Booster. We can inject this batch of beam into Main Injector and then accelerate it, or we can put additional batches of beam from Booster into Main Injector before accelerating. This is known as multi-batch injection and is used primarily when we are running NuMI + stacking cycles.

To reiterate this terminology:

- An RF structure is called a ***bucket***.
- If a bucket has particles in it, it is referred to as a ***bunch***.
- Eighty-four bunches equals one Booster ***batch***.

It is useful to look further into another few terms closely related to buckets and bunches. The 400 MeV chopper at the end of the Linac controls how much beam we inject into the Booster.

A chopper is a device used to deflect all or a portion of a line of beam. Beam traveling down the Linac is by default heading for the Linac dump unless we change its course and steer it toward the Booster. The amount of beam sent into the Booster during an acceleration cycle is measured in ***turns***. One turn is an amount of beam equal to one lap around Booster. Imagine a continuous line of H- ions traveling towards the Linac dump. If we want two turns worth of beam, we set the 400 MeV chopper to deflect Linac beam to Booster for approximately 4.44 μs.

It takes 2.22 μs for 400 MeV beam to complete one lap around Booster, so deflecting beam for twice that amount of time means that ‘newer’ beam will be layered in on top of ‘older’ beam. This is how we vary beam intensity in the Booster.

Unlike all of the other machines at the lab, the bunches coming from Linac do not have the same RF frequency as the Booster’s RF cavities. The Linac beam has a 201 MHz bunch structure, while the injection frequency of the Booster RF cavities is ~37.8 MHz. To allow this frequency mismatch to occur without large amounts of beam loss, a process called ***paraphasing*** takes place. When the Linac beam first comes into Booster, the Booster RF cavities are on and running, but they are out of phase with each other. Half of the tanks have one phase, and half of them are 180 degrees different. The particles injected see no net effect from the RF (neither acceleration nor deceleration). With no net restoring forces from the RF, the 201 MHz structure from the Linac begins to dissipate, and the beam becomes almost a DC line of charge. Once the all of the beam is fully injected (as many turns as we want), the cavities slowly come into phase with one another, forming new 37.8 MHz bunches, and the acceleration process goes on from there. This injection process allows us to use a variable number of turns per cycle, but it is also rather inefficient. Not all of the more or less DC beam is captured when the new buckets form, and thus a plot of Booster intensity vs. cycle time generally shows a characteristic fall off early in the cycle. These bunches, once formed in Booster, are the same bunches handed off from machine to machine the rest of the way down the line. This type of very efficient transfer between machines is called a ***bucket-to-bucket*** transfer. If we should want more intensity per bunch (or equivalently per batch), we vary the number of turns.

Below is a table describing most of the RF systems used here at Fermilab. Notice that not all of these systems accelerate charged particles.

Accelerator	**RF system**	**Harmonic Number**	**Purpose**
Linac	Accelerating RF	Not applicable	Acceleration
Booster	Accelerating RF	84	Acceleration
Main Injector	Accelerating RF	588	Acceleration
Tevatron	Accelerating RF	1113	Acceleration
Pbar Debuncher	DRF1	90	Bunch rotation
Pbar Debuncher	DRF2	4	preserve kicker gap
Pbar Debuncher	DRF3	4	Studies
Pbar Accumulator	ARF1	84	Sweep beam to core
Pbar Accumulator	ARF2	2	Ion clearing
Pbar Accumulator	ARF3	2	Studies
Pbar Accumulator	ARF4	4	Unstack beam from core

D. RF Manipulations

As mentioned above, RF systems are not only used to accelerate charged particles. They can also apply a structure to constant energy beams. When beam is first injected into an accelerator, for instance, the phase of the beam is such that no net acceleration occurs, but phase focusing still takes place. Therefore the bunched structure of the beam stays intact.

You should be aware of the following RF manipulations:

Bunch Rotation

Bunch rotation is a process used in the Main Injector while stacking (producing antiprotons). The beam from the Booster is accelerated to 120 GeV. Then, just before the beam is sent down the P1 line to eventually strike the target beneath AP0, the voltage in the accelerating RF cavities is drastically reduced for a moment. This allows the bunch to naturally 'rotate' in phase space.

> It should be noted that the RF voltage applied to each cavity does not actually change—the relative phases with respect to the beam of opposing pairs of cavities change. For instance, if two separate cavities are 180° out of phase with respect to each other, then the overall effect on the beam of both cavities combined will be zero. Conversely, if the cavities have the same phase relative to the beam, the overall effect is twice as much as a single cavity. By controlling the individual cavity phases in this manner, one can control the effective RF voltage as seen by the beam.

Initially the bunch has a small distribution of individual particle energies and a large distribution in individual particle phases. After the bunch rotates, this situation is reversed. Usually it would be a bad thing to have a larger than normal energy distribution within the beam, but in this case the large energy distribution is temporary. Once the protons hit the target 'all bets are off'—the secondary particles will have a much larger (random) spread in energies.

How then does bunch rotation help? Recall that the increased energy spread in the beam came with a smaller time spread. The 'energy spread information' from the proton bunch does not survive past the target, but the time spread does. This means that with bunch rotation, the clusters of secondary particles (including some antiprotons) coming off the target will be grouped closely in time (or phase). Therefore, with the appropriate phasing, the RF systems in the Debuncher can more efficiently capture this large energy spread (because of the target) beam.

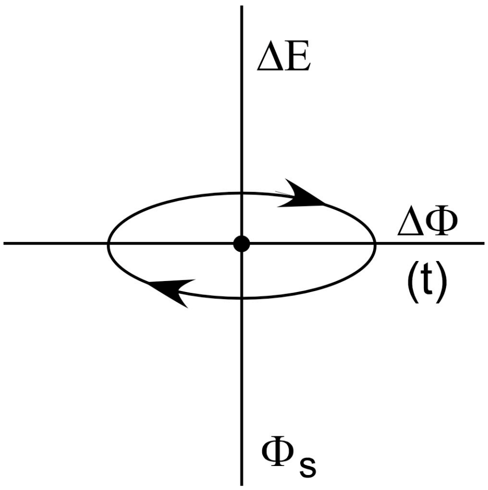

Bunch Rotation. This shows an RF manipulation performed on a bunch that swaps a large time spread and small energy spread (graph on left) for a large energy spread and a small time spread (graph on right).

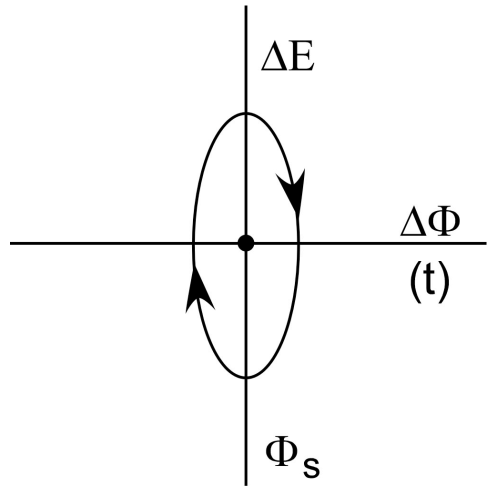

Coalescing

Coalescing is the process of using a special RF system to make one large bunch of beam out of many smaller bunches. This larger bunch is called (strangely enough) a coalesced bunch. The protons and antiprotons used in the Tevatron during collider mode are coalesced. In the case of the protons, 7 Booster bunches are injected into the Main Injector and accelerated to 150 GeV. Once at flattop, the voltage in the main acceleration RF cavities is nearly all the way down—this eliminates most of the restorative forces keeping the beam bunched. Then, special RF cavities make these 7 bunches converge on a single target RF bucket. When the main accelerating RF voltage is turned back up, the beam has been recaptured in the target bucket and we have a single coalesced bunch. The process works in the same way for antiprotons except that 4 separate coalesced Pbar bunches are made at the same time.

Cogging

Cogging is the process of adjusting the energy of beam in one machine relative to another in order to place bunches of beam into specific RF buckets. For example, when the MI and Tevatron are both at 150 GeV, they are also both at the same RF frequency. To complete an efficient bucket-to-bucket transfer between the two machines, these RF frequencies must be matched in phase so that particles sent from the Main Injector arrive in the Tevatron RF cavities when the correct electric fields are present. However, beam is not simply transferred from the MI into a random Tevatron bucket even if it is in phase. The particles are placed in a specific pattern around the ring through a process called ***transfer*** **cogging.** Imagine two large gears meshed together. The Main Injector gear has 588 teeth (RF buckets), and the Tevatron gear has 1113. Once these gears are synched up with each other, they are locked into position relative to each other as well, and particle transfers can occur between them.

We want to send protons in a given MI bucket into any Tevatron bucket. The solution is to change the RF frequency in the MI slightly, making the two machines out of phase with each other for a time. While the two ‘gears’ are out of phase with each other, they will rotate at different speeds, causing different sets of ‘teeth’ to come near to one another. If the MI frequency were changed back to its original value at the appropriate time, any MI bucket could line up with any Tevatron bucket. Another version of cogging occurs solely within the Tevatron. This type, called ***collision point cogging***, varies the RF frequency of the antiprotons relative to the protons. Collision point cogging allows us to specify exactly where the bunches of opposing particles pass through each other, namely at the center of the Collision Halls.

E. Generalized RF amplification system

This section will describe a 'typical' RF amplification system. Although there are subtle differences between the RF amplification systems used here at Fermilab, it will be helpful for the new operator to realize that there is a general scheme to the RF systems.

In order to drive an RF system, one needs a to produce a high power sinusoidally varying signal at the cavity's resonant frequency. This drive signal can be fed into the cavity and used to create the large electric fields that accelerate the beam. The diagram below shows the main pieces of the RF amplification system. On the next page you will find system descriptions.

RF oscillator/Low Level RF (LLRF)

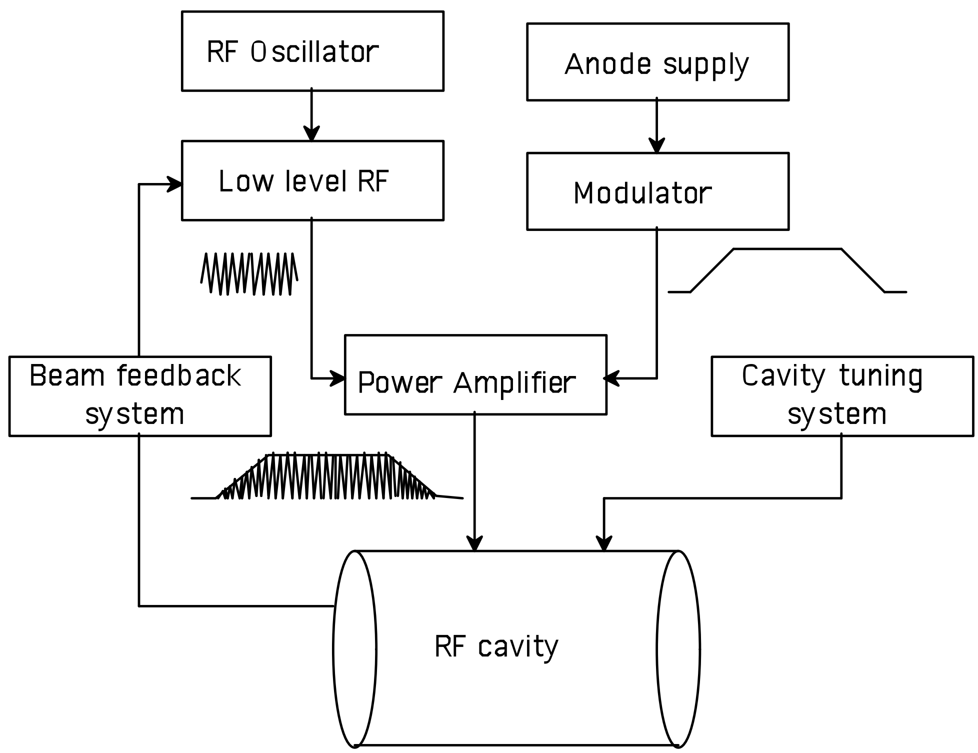

The oscillator is the low power source of the correct frequency signal that drives the cavity. The LLRF system is typically a computer that can adjust the phase of the oscillator's signal as needed.

Anode supply

The anode supply is a large raw power supply used as the source of the high voltage that modulates the LLRF signal.

Modulator

A modulator uses the anode power supply's high voltage to control the power amplifier.

Power Amplifier

The power amplifier, or PA, mates the anode supply's high voltage to the LLRF's correct frequency and phase signal. The PA produces an amplitude modulated high voltage signal.

RF Cavity

An RF cavity is the electromagnetically resonant structure in the tunnel that creates the strong electric fields needed to accelerate the beam. Applying this large amplitude signal at the correct frequency creates the field.

Beam Feedback system

A beam feedback system monitors the actual phase of the beam with respect to the RF. Based on this input, the LLRF adjusts the phase of the low-level signal accordingly.

Cavity Tuning system

This dedicated system controls the parameters of the RF cavity in order to achieve or change the desired resonant frequency. In practice this can be split into two categories: a cavity volume control system and a cavity tuning system.

In the volume control system, the precise volume of the RF cavity is maintained, keeping the resonant frequency constant. Why? Because even slight changes in the temperature of the cavity can cause enough thermal expansion or contraction to be noticeable. The temperature of any cavity's cooling water must be carefully maintained. In the low energy Linac a small copper slug, called a tuning slug, controls the volume by its position in or out of the cavity. Slightly varying the water temperature or the mechanical position of a tuning slug may keep a cavity resonating at a constant RF frequency, but these methods would be much too slow acting to allow for acceleration in a synchrotron.

In the cavity tuning system, the cavity's natural resonant frequency alters quickly. Those accelerators like Booster and the Main Injector that have to dramatically change the resonant frequency of their cavities throughout the acceleration cycle rely on devices called ferrite tuners.

Ferrite tuners connect electromagnetically to the RF cavity that controls the system's R-L-C parameters. Controlling the current flow through the ferrite tuners in turn controls the resonant frequency of the cavity. The bias supply (found standing beside the modulator at Booster and Main Injector RF stations) is responsible for this controlled current flow to the ferrite tuners.

When the RF frequency sweep necessary for acceleration is small enough, the RF cavity does not need a fast cavity tuning system. All RF cavities actually have a small distribution of possible frequencies at which they will resonate, rather than an exact single frequency. If the entire frequency sweep necessary for acceleration falls within this small range, as is the case with the Tevatron RF cavities, a system using ferrite tuners is not needed. This is the reason Tevatron RF stations do not feature a bias supply.

The table below shows the operating frequency and tuning method of various RF systems.

Accelerator	**Frequency Range (MHz)**	**Cavity volume control**	**Cavity tuning system**
Low energy Linac (DTL)	201.24	tuning slug/cavity temp	n/a (fixed frequency)
High energy Linac (SCL)	805	cavity temp	n/a (fixed frequency)
Booster	37.8 – 52.8	cavity temp	3 ferrite tuners per cavity
Main Injector	52.8 – 53.1	cavity temp	2 ferrite tuner per cavity
Tevatron	53.103 – 53.104	cavity temp	none (Δf small)

Notes:

III. MAGNETS

A. Deflecting Charged Particles

As we have seen in the last chapter, RF cavities accelerate charged particles. In order to reach higher and higher kinetic energies, one must allow the charged particles to travel through more acceleration gaps. Either you build a longer Linac *or* force the particles to travel through the same acceleration gap multiple times by building a circular machine. Using the same RF cavities over and over is efficient, but it also imposes some requirements. First of all, one has to be able to bend the particles around in a circle. Secondly, one has to be able to keep bending the particles around in the same radius circle as the particles gain energy.

Accelerators use time-varying magnetic fields to accomplish these tasks. Recall the E&M force equation:

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

A magnetic field cannot be used to change a charged particle's kinetic energy, but it can be used to deflect a particle from its original path.

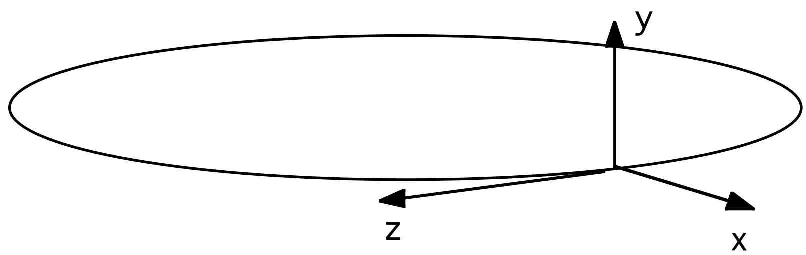

Using the right-handed coordinate system one can see that if we set up a magnetic field in the Y-direction, particles traveling in the Z-direction will "feel" a force in the X-direction. This is exactly what we need to keep our particles traveling in a circle. Only charged particles (where $q \neq 0$ in the force equation above) can be deflected by magnetic fields. Also, the force "felt" by a positively charged particle in a given magnetic field is equal and opposite to that "felt" by a negatively charged particle with the same energy. This fact allows protons and antiprotons in the Tevatron to circle in opposite directions—bent by the same bending magnets.

Throughout the rest of this chapter, we will examine some of the different types of magnets found in our accelerators. As it turns out, magnetic fields are used to do more than simply bend the paths of the charged particles into a circle.

B. Bend Magnets

The reason we bend the particles around in a circle is to efficiently use the same RF cavities over and over again. Each time a charged particle passes through a cavity, it is subjected to an electric field. The electric field increases the particle's kinetic energy a little bit with each pass. Slowly, over thousands of passes, the particle's kinetic energy is built up to the desired level. Magnetic fields force the particles into a circular orbit. When entering a region containing a uniform magnetic field, as shown below, a charged particle will follow a circular trajectory.

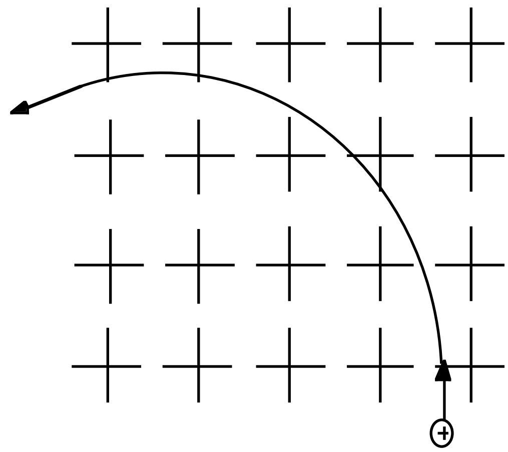

Bend Field. The crosses indicate a magnetic field pointing into the page. Given the positively charged particle initially traveling upwards, the right hand rule dictates a force to the left, deflecting the particle from its original path.

The radius of the circular trajectory seen above can be shown to be:

$$r = \frac{p}{qB}$$

Notice how this radius, *r,* is proportional to the particle's momentum, *p*, and inversely proportional to the strength of the magnetic field, *B* and the particle's charge, *q*. In order to keep charged particles confined to the same radius circle during acceleration (as *p* is increasing), the strength of the bending magnetic field *B* has to be increased. This is one of the characteristics of a **synchrotron**, a type of RF particle accelerator where the frequency of the accelerating RF and the strength of the magnetic bend field vary as needed to keep up with the beam's changing momentum and revolution frequency. The act of changing the magnetic field over time to compensate for the changing beam energy is commonly referred to as "ramping the magnets."

Before we go on, you might wonder how Fermilab creates these magnetic fields in the first place? First of all, there are two basic ways to create a magnetic field.

1. Use a magnetized piece of metal, like a large version of a magnet stuck on a refrigerator door. This type of magnet can also be called a **permanent magnet**, because after it is made, the strength of its magnetic field cannot be varied. Because the field cannot be changed, this type of magnet wouldn't work in an accelerator where the beam's momentum changes. This type of magnet would work, however, in a situation where the beam momentum is constant, or very nearly so. The Recycler and the MI-8 line are both made up of primarily permanent magnets, because the beam in the Recycler and the beam coming out of the Booster are both at a constant, fixed energy.
2. Use an **electromagnet**, which is basically a loop of wire that has an electrical current flowing through it. One of the byproducts of current flowing through a wire is the creation of a magnetic field. If one increases the amount of current flowing through the loop of wire, the strength of the magnetic field created increases proportionally. Because of this fact, electromagnets are used in machines where the beam energy changes, such as the Booster, Main Injector, and the Tevatron.

Synchrotrons are not made up of a single very strong electromagnet, but rather a series of individual bending magnets distributed in a ring with each magnet only required to deflect the beam a few degrees (not in a complete circle).

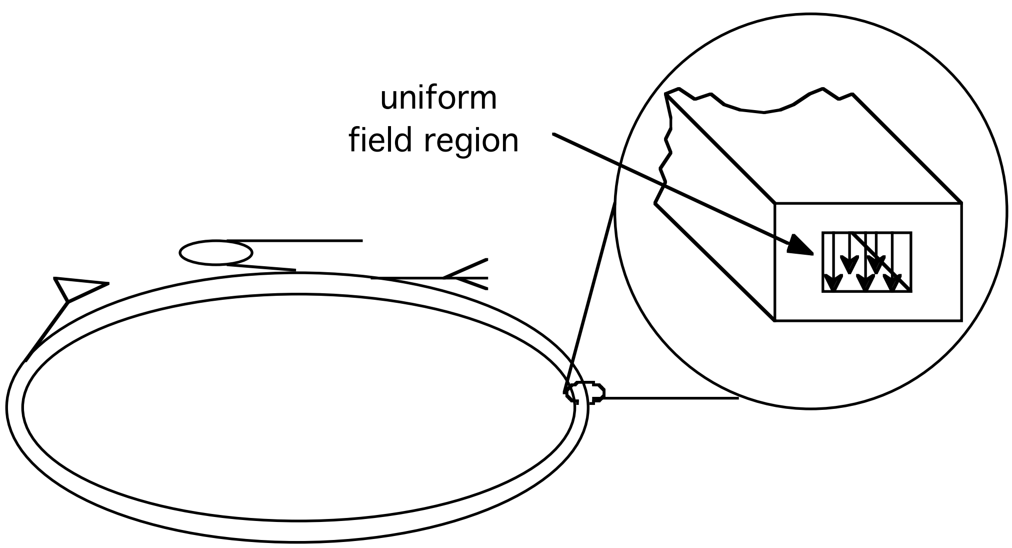

Below is a picture of a modern bending magnet like those used in the Main Injector. There are some interesting features worth pointing out. First of all, notice the copper "bus" used to make the current loop coils of the magnet (and used to connect the magnet to a power supply). This bus is actually a thick, high current carrying wire with a hole down the center.

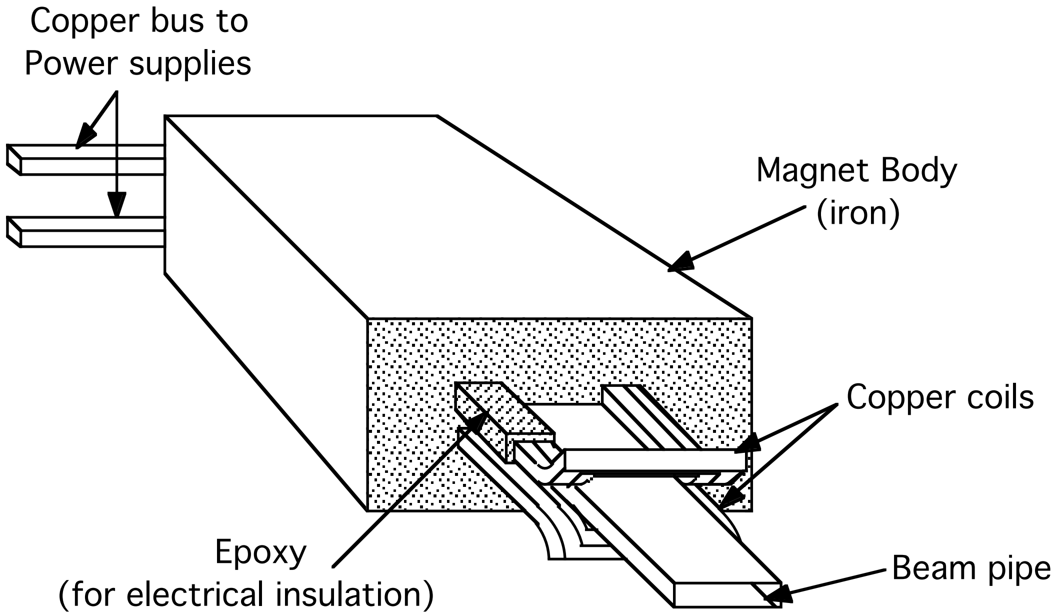

Specially treated cooling water (called LCW, which stands for Low Conductivity Water) flows through the passage inside the bus to carry away the large amounts of heat generated by the high current flow. Also notice that an iron magnet body, which concentrates and shapes the magnetic field, surrounds the coils. Finally, notice how the beam pipe inside the magnet is placed very close to the coils themselves. This maximizes the magnetic field "felt" by the beam passing through.

The bend magnet is called a ***dipole*** because of the two magnetic pole faces formed by the two opposing coils. The dipole generates a magnetic field in the Y-direction that in turn produces a radial force in the X-direction on the charged particle traveling in the beam pipe at the magnet's center. If the current in the coils of this electromagnet is increased, the strength of the bend field increases, allowing a higher energy beam to be bent in the same radius.

Recall the synchronous particle from the previous chapter. A synchronous particle is that particle whose energy and phase are perfect. With each lap around the accelerator, this particle arrives in the RF cavity at the correct time, and hence receives the perfect increase in kinetic energy. The increasing energy of the particle is exactly accounted for by this increase in magnetic bend field created by the ramping dipoles, and therefore the perfect particle is eternally bent along the same circular path. This perfect path, traveling down the center of the beampipe through the accelerator, is called the ***ideal orbit*** or the ***closed orbit***.

Not surprisingly, the vast majority of the particles we accelerate are not perfect. They have slight errors in energy and/or phase that cause them to undergo synchrotron oscillations. These small differences in momentum relative to the perfect particle cause the small differences in the amount that these non-perfect particles are deflected by the dipoles. Remember that the energy of the synchronous particle determines the strength of the magnetic field at any given time during the acceleration cycle. Therefore, if a particle's momentum is not the same as the synchronous particle, neither will its circular orbit have the same radius. Higher momentum particles will be bent slightly less than the perfect particle, and thus travel on an orbit somewhat radially to the outside of the ideal orbit. Similarly, particles with a slightly lower momentum, relative to the perfect particle, will travel on a path slightly to the inside of the ideal orbit. This distribution in orbit radius based on the natural momentum spread found in the beam is called ***dispersion***.

The imperfect collection of particles that we accelerate has other problems as well. We'd like to think that these particles are traveling solely in the Z-direction—that is, straight ahead—but in reality each non-perfect particle in the bunch has a nonzero transverse velocity. These particles are moving in the transverse plane for a variety of reasons. Remember that like-charges repel one another, so when you have beam in a bunch strong Coulomb forces exist that tend to make the particles want to separate from each other. If this separation is not resisted, the bunch will eventually fly apart and the particles will collide with the beam pipe. We need to have a method of focusing the particles in the transverse planes, much like how the RF focuses the particles in the longitudinal plane.

C. Focusing Magnets

To begin our discussion of focusing magnets, let's imagine an imperfect proton inside a beam bunch. As this proton is traveling through the accelerator, it finds itself naturally being repelled by the other protons within the bunch. Assume this repelling force causes our proton to begin to drift upwards in the Y-direction. If something isn't done to correct this particle's upward travel, it will hit the beam pipe and be lost. As we have seen, a deflecting magnetic field could be just the solution. If we were to arrange a magnetic field in the X-direction, perhaps by rolling a bending dipole onto it's side, we could make our particle feel a force in the Y-direction. This restoring force could resist the Coulomb repulsion and guide the wayward particle back towards the ideal orbit. While this type of simple solution might work for this particular particle, it would cause problems elsewhere in the bunch.

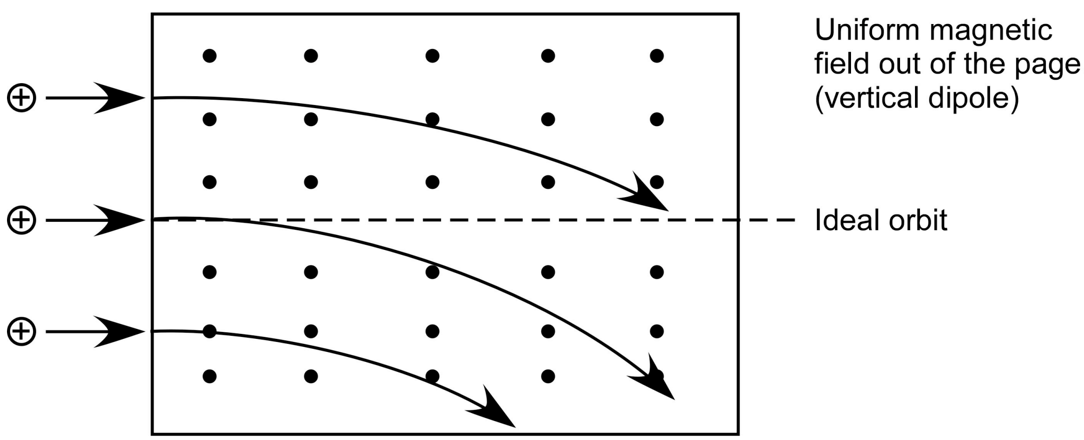

Attempt at vertical focusing. The three arrows at the left represent three different particles within a bunch. The top most is the imperfect particle in the example in that has drifted upwards. The center arrow is a perfect particle initially traveling on the design orbit. The bottom arrow is another imperfect particle that has drifted downwards. When all three particles enter the field created by the vertical dipole, they are deflected downwards.

We have perhaps corrected the path of the particle in question, but at the expense of all of the other protons. We want to create a magnetic field that corrects the orbit of the particles found to both sides of the design orbit, but leaves the perfect synchronous particle alone. Also, because particles found further from the ideal orbit need correspondingly stronger restorative forces, we want the strength of the focusing magnetic field to vary with distance from the ideal orbit.

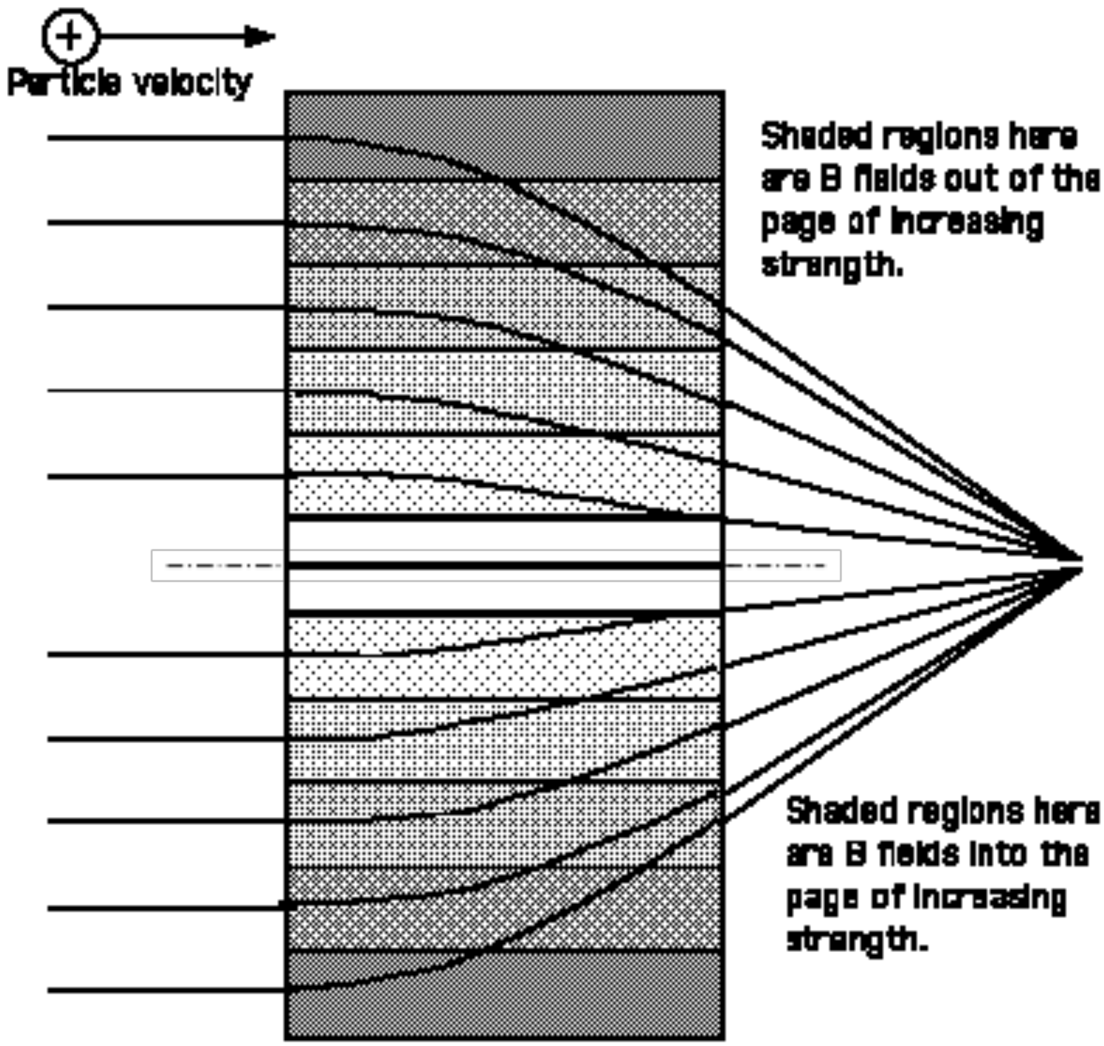

Magnetic focusing 'lens'. The diagram shows the type of magnetic focusing fields we would like to create. Particles found further from the design orbit travel through a stronger magnetic field and therefore feel a stronger restorative force.

As it turns out, a magnetic field distribution much like the focusing lens picture above can be produced using a special electromagnet called a ***quadrupole***, so named for it's four magnetic poles created by four current loops. A cross-sectional picture of a quadrupole is shown below from the particle's eye view.

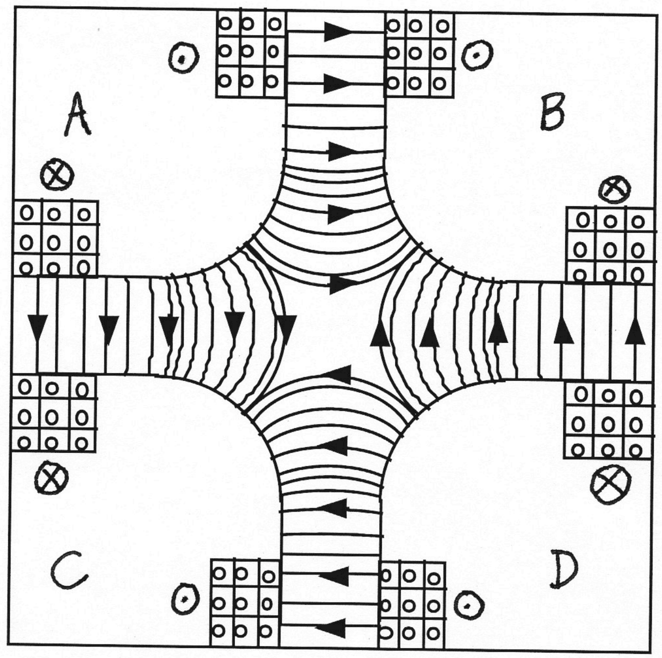

Quadrupole magnet. The four (usually iron) magnetic pole faces are labeled A, B, C, and D. The arrowed lines indicate the magnetic field. The direction of current flow, which is into or out of the page, is labeled near the square bus work. Notice the holes in the bus for cooling LCW. A given quadrupole can only focus the beam in one plane—while defocusing the beam in the other.

Looking at the above diagram, one can see that the quadrupole is made up of four separate loops of wire wrapped around four pieces of iron, called magnetic pole faces. The electric current flowing in four loops creates the magnetic field. The direction of current flow in these four loops is indicated. (A circle with a dot indicates 'out of the page' while a circle with an *X* indicates 'into the page'). The bus work, which are the heavy-duty wires forming the current loops, can be seen as small clusters of squares with holes in the middle. These holes are passages for cooling LCW to flow. The four pole faces concentrate and shape the magnetic field.

A given quadrupole magnet can only focus the beam in one plane. Applying the right hand rule to the lines of magnetic field shown in the picture, one can see that a positively charged particle traveling into the page will see no net force when passing through the center. When the same particle passes through the quadrupole off-center, net forces are felt. For instance, a positive particle traveling into the page between the A and B pole faces will feel a magnetic force downwards. This force acts to direct the particle back towards the ideal orbit, focusing in the vertical plane. If the same positive particle were to pass through this magnet off-center between the A and C pole faces, a force to the left, away from the ideal orbit, would be felt, defocusing in the horizontal plane. If we rotate the quadrupole above by 90 degrees, we get a magnet that behaves in the opposite way—it defocuses beam in the vertical plane and focuses in the horizontal plane. One can also change the quadrupole's orientation by reversing the flow of electrical current running through the magnet's bus work.

Quadrupoles are named for the effect they have in the horizontal plane—focusing or defocusing. Alternating the placement of focusing and defocusing quadrupoles in an accelerator produces a net focusing effect in both planes.

There is another type of quadrupole that is neither horizontal nor vertical known as a ***skew quadrupoles***, where the magnetic field is rotated or 'skewed' by 45 degrees. Skew quadrupoles control the coupling of the horizontal and vertical tunes. If the tunes are coupled, a change made to the tune in one plane affects the tune of the other plane as well, a situation that is not desirable.

Quadrupole focusing does not perfectly restore a wayward particle to the ideal orbit, much like phase focusing, as described in the RF chapter, it does not restore an out of phase particle to the synchronous phase. Instead the restorative forces created by the quadrupoles send particles, which otherwise would have been lost, into stable transverse oscillations about the ideal orbit. These side-to-side motions are called ***betatron oscillations***. The term betatron refers to the early accelerators in which these oscillations were first noted. Real particles traveling in real accelerators follow complicated paths indeed, moving back and forth longitudinally due to the RF and transversely in both the horizontal and vertical planes due to the quadrupoles.

The trajectory of any particle traveling along the sinusoidal path shown above can be described in terms of two variables: the particle's position, or distance from the ideal orbit, and the angle that the particle's trajectory deviates from the ideal orbit.

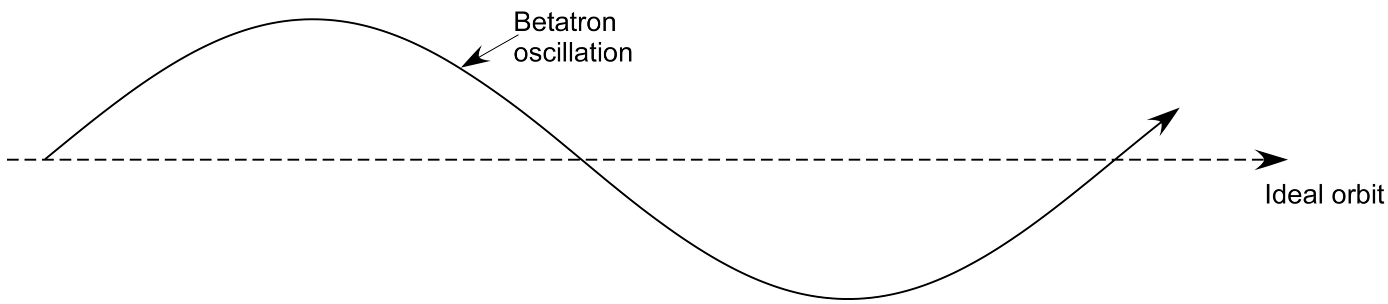

Betatron Oscillations. Quadrupole focusing forces particles into transverse oscillations about the ideal orbit. The particle's path is shown as the solid line above.

As the beam moves through the accelerator, the focusing and defocusing quadrupoles constantly correct individual particle trajectories. These restorative forces produce betatron oscillations about the ideal orbit in both the vertical and horizontal planes. The number of betatron oscillations that the beam undergoes in one complete trip around the accelerator is defined as the ***betatron tune***, or simply the ***tune***. An accelerator has largely independent tunes in both the horizontal and vertical planes. As will be discussed later in the Beam chapter, the tunes have to be controlled in order to avoid instabilities that could lead to beam loss.

Much like the case of longitudinal motion detailed in the RF chapter, the transverse motion of betatron oscillations can be shown as a phase space diagram. The variables are *x,* the particle's position at a given point, and *x'*, the angle between the particle's trajectory and the ideal orbit. There is some set of initial position and angle variables that newly injected particles can have and still stay in an accelerator.

In the graph below, the solid ellipse marks the stability boundary called the Separatrix. Particles whose initial *x* and *x'* coordinates do not fall within this closed boundary will be lost. For instance, one can inject a particle into a machine with such a large initial deviation from the ideal orbit, *x*, that it cannot be forced into stable oscillations (focused) by the quadrupoles. Such particles oscillate wildly until they eventually collide with the beam pipe or another accelerator component, causing losses and radioactivity in the tunnel.

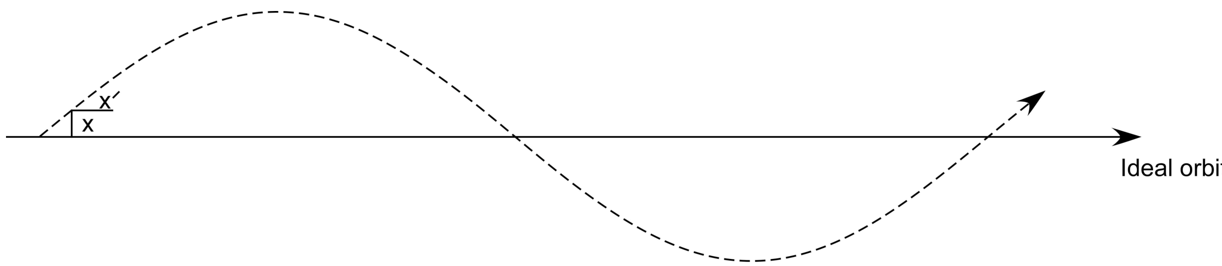

Position and angle. A particle's trajectory at any point can be described in terms of its distance from the design orbit, *x*, and the angle between its trajectory and the ideal orbit, *x'*.

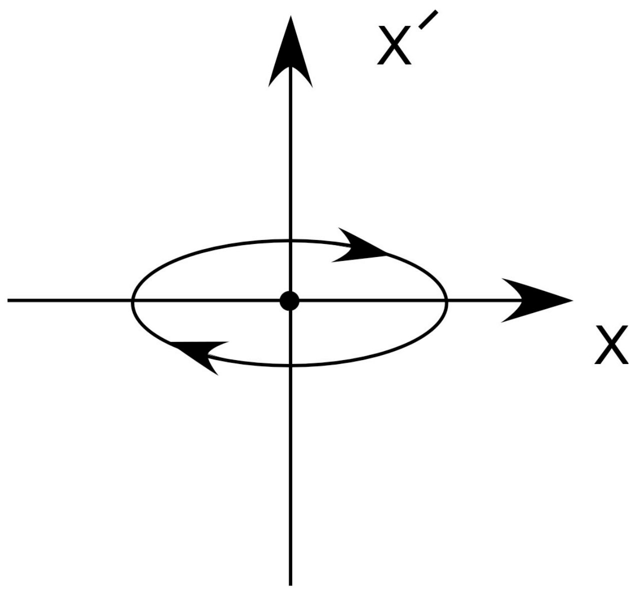

Phase space diagram. Here's another way to view betatron oscillations. As the particle travels along the ellipse its distance from the ideal orbit, *x*, increases and decreases, as does the angle of the particle's trajectory relative to the ideal orbit, *x'*. The ellipse shown is the Separatrix—the boundary between stable particles that stay in the machine and unstable particles that are lost.

The area contained within the phase space ellipse above is called the ***emittance***. Emittance is like the density of the tiny clouds of particles and is roughly akin to the cross-sectional area occupied by the beam. Emittances can be measured transversely in both the horizontal and vertical plane and longitudinally. Generally, lower is considered better as far as emittance goes. This is because a larger area on the above graph will mean that the actual beam size, the furthest excursion of the ellipse in the "x" direction, is larger and the probability of scraping some of the particles against accelerator components goes up.

D. Combined Function Magnets

Not all of the accelerators here at Fermilab use a system of separate magnets for the bending and the focusing of beam. Booster is a prime example. The main magnets used in Booster are known as combined function magnets. What this means is that instead of having bend magnets and quadrupole magnets around the ring, a different type of magnet is used to produce the combined field of a dipole and quadrupole, so the magnet both bends beam in an arc, as well as providing a focusing force. Below is a cross section of both a focusing and defocusing magnet.

The easiest way to see how this field bends beam and focuses it is by breaking the field lines into their x and y components. From there, one can compare the component lines with

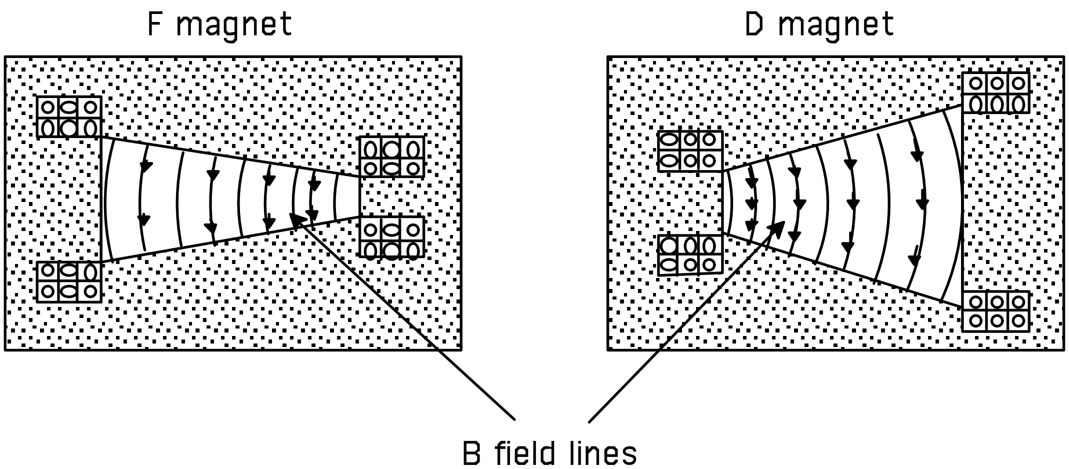

those of the bend field and quadrupole field shown in the previous sections, which we will leave as an exercise for the reader. Permanent (non-powered) gradient magnets are also used in the Recycler and the MI-8 line.

E. The Lattice

The placement of the bend magnets, quadrupole magnets, and drift spaces within an accelerator is not determined randomly. These components are placed according to a precise plan. This pattern of focusing, bending, and drifting is referred to as the accelerator's ***lattice***. Lattices are often periodic, with the smallest repeating unit called a ***cell***. The Main Injector, for instance, has two quadrupole and four dipole magnets making up its basic cell. Specifically, a Main Injector cell is made up of a focusing quadrupole magnet followed by two dipole magnets followed by a defocusing quadrupole magnet followed by two more dipole magnets. A shorter way of saying the above statement is to say that the Main Injector has a FODO lattice, pronounced 'foe-doe'. The 'F' and 'D' denote the types of quadrupoles while each 'O' denotes a drift space. In the Main Injector these spaces contain two bend magnets. Booster, on the other hand, has a FOFDOOD lattice, pronounced 'foff-dude', where the 'F' and 'D' represent focusing or defocusing combined function magnets. In the case of Booster, the 'O' denotes a short straight section of beam pipe and the 'OO' denotes a long straight section of beam pipe. These straight sections are where RF cavities, kickers, and other correction or diagnostic equipment are placed.

Accelerator	Cell	Cell Description
Booster	FOFDOOD	F, D are gradient magnets, O is a short straight, OO is a long straight
Main Injector	FODO	F, D are quadrupoles, O represents two dipoles
Tevatron	FODO	F, D are quadrupoles, O represents four dipoles

F. Correction Elements

It should not be surprising to note that real accelerators are not built perfectly. Magnets, RF cavities and other components may have slight construction errors in them, or they may be slightly misaligned. These small defects in the layout of a large machine could be enough to render an accelerator inoperable if they were not accounted for. That is where correction elements come in. Correction elements (or correctors, as they are also known) are scaled down versions of the larger magnets discussed above. Dipole correctors are used to finely tune the path the beam travels, and quadrupole correctors are used to tweak the tune of the machine.

The dipole corrector is a small dipole magnet powered by an individually controlled power supply. These correction elements impart a small bend to the particle beam in order to move it to the desired location. Correction elements of this type are usually found at regular intervals around the ring and throughout the beamlines between accelerators.

The quadrupole corrector plays the same role as the main quadrupole magnets mentioned above. They are used to control the tune or the coupling of the tunes in circular accelerators and to maintain beam focusing in Linac and the beam lines.

G. Higher Order Magnets

The equation below shows a magnetic field in its most general form, expanded via a Taylor series where the magnetic field *B* is shown as a function of *y,* the distance from the center of the magnet, or distance from the ideal orbit:

$$B(y) = b_0 + \frac{b_1 y}{1!} + \frac{b_2 y^2}{2!} + \frac{b_3 y^3}{3!} + \cdots \text{L}$$

The various terms of this expansion correspond to some of the single purpose magnets we've seen so far. Recall the uniform magnetic field generated by a dipole—represented by the constant b_0 term shown above. Because this term is constant, all beam traveling through a given dipole encounters the same magnetic field, regardless of the beam's transverse position.

The remaining terms in the expression are called 'higher order' components. The field generated by the quadrupole is the second term. Notice how the quadrupole field increases linearly with distance from the ideal orbit. Particles further away from the center of the quad see a correspondingly stronger magnetic field, while particles traveling through the center of the quad (where $y = 0$) see no net magnetic field at all.

The remaining two terms shown are non-linear additions to the magnetic field. Two types of non-linear magnets are commonly encountered in accelerators: ***sextupoles*** (corresponding to the y^2 term), and ***octupoles*** (the y^3 term). Again, beam traveling through the center of these magnets, where $y = 0$, sees no magnetic field. Particles traveling off-center through these magnets see exponentially increasing magnetic fields as the distance from the ideal orbit increases. The exact effects of nonlinear fields on the beam are subtle and difficult to see in terms of magnetic forces. These higher order fields allow us to control the coupling between two pairs of accelerator parameters that we do not want to affect each other. This is similar to how the skew quadrupole decouples the horizontal and vertical tunes to allow for independent control. The first higher order magnet we will look at is the sextupole.

Sextupole

The individual particles in a bunch do not all have the exact same momentum. Instead, the bunch has some bell curve like distribution of possible particle momenta. This is referred to as a momentum spread within the beam. The momentum spread is usually given as "Δp/p". Read as "delta p over p", this ratio is equal to the momentum spread divided by the momentum of the synchronous particle. While a non-zero momentum spread is unavoidable, we would like the value to be small. The value of this dimensionless unit is around 0.001 typically. Similarly, there is a spread of individual betatron tunes within a bunch. Not every particle is effected by the quadrupoles in the same way. There exists a coupling between the momentum spread and the tune spread within the beam we accelerate. The ratio of tune spread to momentum spread within an accelerator is called the ***chromaticity***.

The sextupole magnet is used to control the chromaticity of the accelerator in order to prevent subtle instabilities from causing beam loss.

A diagram of a sextupole and the resulting magnetic field lines is shown below:

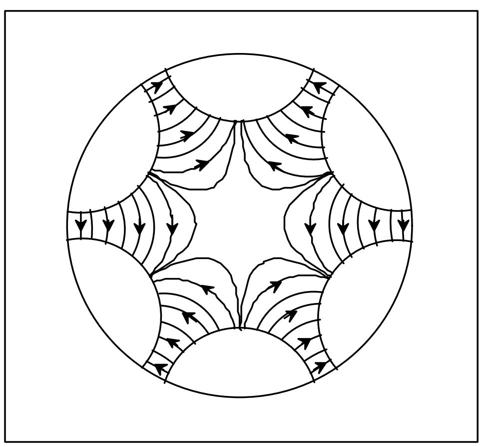

Sextupole magnet. The sextupole has six magnetic pole faces. The arrowed lines indicate magnetic fields. A sextupole controls the chromaticity in an accelerator.

There are three different types of sextupole fields: horizontal, vertical, and skew. The horizontal and vertical sextupoles control the chromaticity in the horizontal and vertical planes, respectively. The skew sextupole controls the coupling of the chromaticity between the horizontal and vertical planes.

Octupoles

Like the sextupole, the octupole controls the coupling between two seemingly independent accelerator parameters. In this case, the two parameters are the accelerator's tune and the amplitude of betatron oscillations the beam undergoes.

The octupole has eight pole faces and is shown below:

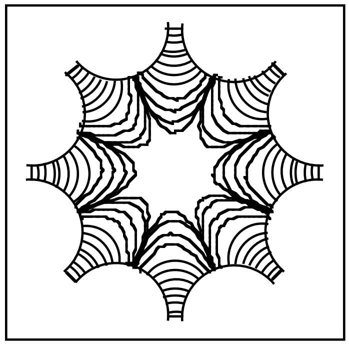

Octupole magnet. The octupole has eight magnetic pole faces. The lines indicate magnetic fields. It controls the coupling between betatron tune and betatron oscillation amplitude in an accelerator.

Octupole magnets are used in a process called ***resonant extraction***. Normally when a bunch is sent from one machine to another, it is an all or nothing proposition. During resonant extraction, natural instabilities are controlled in such a way that beam slowly dribbles out of the accelerator over a long (one or more seconds) time period. See the Beam chapter for a more complete discussion of resonant extraction.

H. Special Devices

There are other specialized electromagnetic devices that can be found in the various machines. This section will describe some of the most common of these accelerator components. Not all of the components are magnets.

Kickers

Kicker magnets, or simply kickers, are like little dipole magnets that can turn on very quickly. Kickers are used to alter the trajectories of particles, usually when injecting beam into or extracting beam out of a circular accelerator.

Physically, a kicker system is actually made of three components. The pulse forming network, or PFN, is a group of capacitors that are charged up by a high voltage supply. Cables from the PFN run down into the tunnel to the kicker magnet. The kicker magnet is formed from a single piece of large copper bus, unlike a dipole that is made from many windings of bus work. Like any current-carrying straight piece of wire, the copper bus in the kicker generates a magnetic field. Cables then lead from the kicker magnet back upstairs to a resistor, completing the circuit.

When it is time for the kicker to fire, a switch tube allows the PFN to discharge a large pulse of current through the kicker magnet. A magnetic field proportional to this high electrical current is generated quickly, within tens of nanoseconds. The current pulse travels through the kicker magnet and is then dissipated in the resistor back upstairs. A kicker is a so-called pulsed device due to the current pulses that power it.

The extremely rapid formation of a magnetic field within the kicker is quite unlike a dipole, which takes perhaps seconds to generate a stable field. Because the kicker is so fast acting, extracted beam from a circular machine can pass through the dormant kicker thousands of times during the acceleration process with no effect. However, when the time for extraction arrives, the kicker quickly turns on. This turn on time is usually synchronized with the beam so that as the field is forming no beam is passing through the kicker. With the magnetic field in place, the next time the beam enters the kicker it receives a ‘kick’ (hence the name). This small deflection changes the particle’s trajectory. When extracting beam, this change in trajectory is said to knock the beam off of the closed orbit.

When the beam to be extracted is knocked off of the closed orbit, the beam undergoes oscillations that can be taken advantage of by the appropriate placement of a Lambertson (see below). When injecting beam into an accelerator, kickers are used to place beam onto the closed orbit, delivering a kick that dampens out any irregular oscillations in the newly injected particle trajectories.

Chopper

A chopper is an electrostatic device that deflects a portion of a continuous stream of beam. Found at both the low energy and high-energy ends of the Linac, choppers decide how much beam will enter the Linac to be accelerated (750 keV chopper) and also decide how much of that newly accelerated beam will make it into the Booster (400 MeV chopper).

Physically, a chopper is a pair of conductive plates each connected to a high voltage power supply. Switch tubes control which (if either) of the plates are charged by it’s respective 28 kV power supply.

Look at the diagram of the Linac 750 keV chopper shown below, taken from the Linac rookie book:

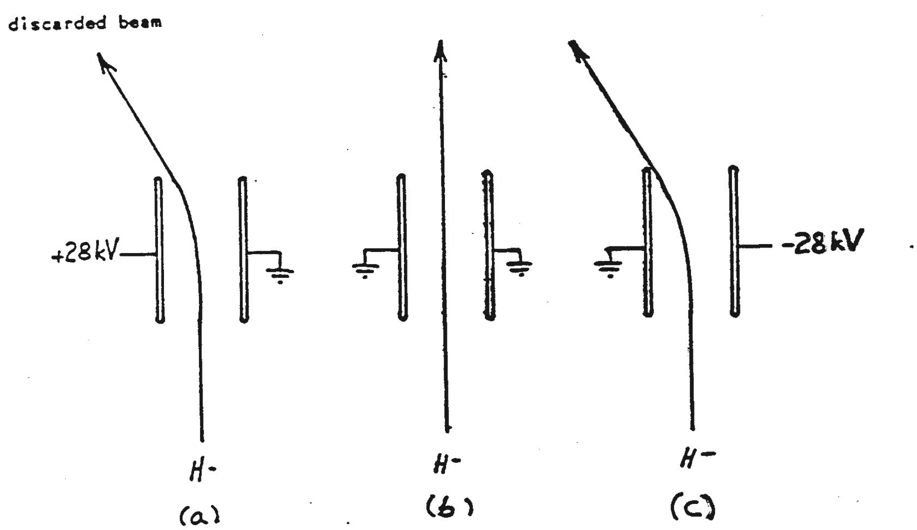

750 keV chopper – .The chopper is made up from a pair of conducting plates that straddle the beam path, allowing a portion of a more or less continuous stream of charged particles to be deflected.

In figure (a), the conductive plate on the left is connected to a 28 kV power supply that provides a deflecting electric field across the path of the beam from the Preacc. This electric field discards the unwanted H- ions into a carbon disk. When Linac requests beam in, the first switch tube (called the "ON" tube) fires and allows the left-hand plate to discharge to ground. The electric field is then gone, (b), and the beam passes through the chopper undeflected. After the appropriate amount of beam has been allowed into the Linac, the second switch tube (the "OFF" tube) fires and connects the right-hand plate to its power supply (c). The deflecting electric field returns and beam is again sent to the carbon disk. This beam chopping occurs 15 times every second. The 400 MeV chopper behaves in a similar fashion, sending beam either towards the Linac dump or down the 400 MeV line towards the Booster.

Lambertson

The Lambertson magnet is a special type of magnet used where beam is injected into or extracted out of an accelerator. Named for the FNAL employee, who designed them, Lambertson magnets are like modified dipoles. The unique characteristic of this type of magnet is the addition of a hole in the iron magnet body that contains no magnetic field. Beam coming up to a Lambertson has two choices: the standard dipole like aperture that will deflect the beam, or the so-called field-free region that will not deflect particles passing through. Other devices, called kickers, (described above) are located upstream of the Lambertson and force the beam into one of the two apertures depending on what we want to do. The picture below shows a typical cross sectional view.

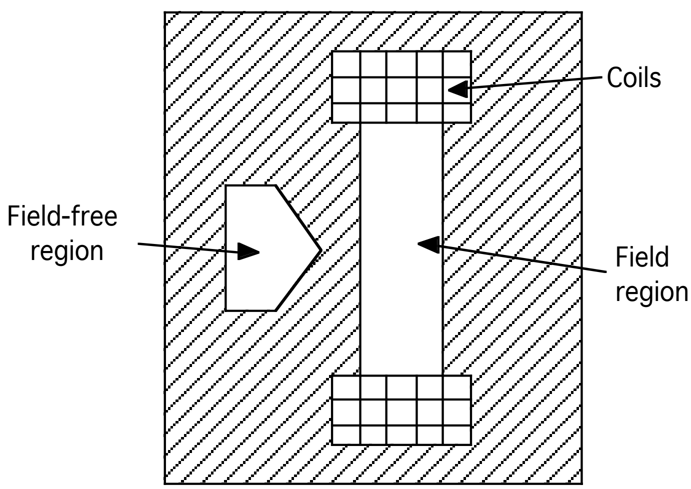

Lambertson. This magnet has two apertures: a magnetic field region (like a dipole), and a field free region. It's used as an injection or extraction device in circular accelerators or as a way to separate beams in a beam line.

Using the extraction process as an example, the goal is to get the beam circulating in an accelerator into an extraction beam line. Typically the injection-extraction process occurs in straight sections of the accelerator, where no bending force is required to keep the circulating beam circulating.

During the accelerating process, beam travels through the Lambertson magnet in the field free region. Once the extraction process has begun, the beam is given a ‘kick’ (using a kicker, appropriately enough) that distorts the orbit the beam travels on and so forces the beam to enter the field region. Once in the field region, the beam is bent out of the plane of the accelerator and directed into a beam transfer line.

For the injection process, the beam enters the Lambertson at some angle with respect to the plane of the accelerator. The dipole field in the Lambertson helps reduce this angle so that when beam leaves the Lambertson it is nearly parallel to the plane of the accelerator. Further downstream, an injection kicker moves the beam onto the closed orbit, ensuring that the beam stays in the machine. When the beam reaches the injection point again, it will be traveling through the field-free region of the Lambertson and will continue to do so until the beam is extracted at a different straight section.

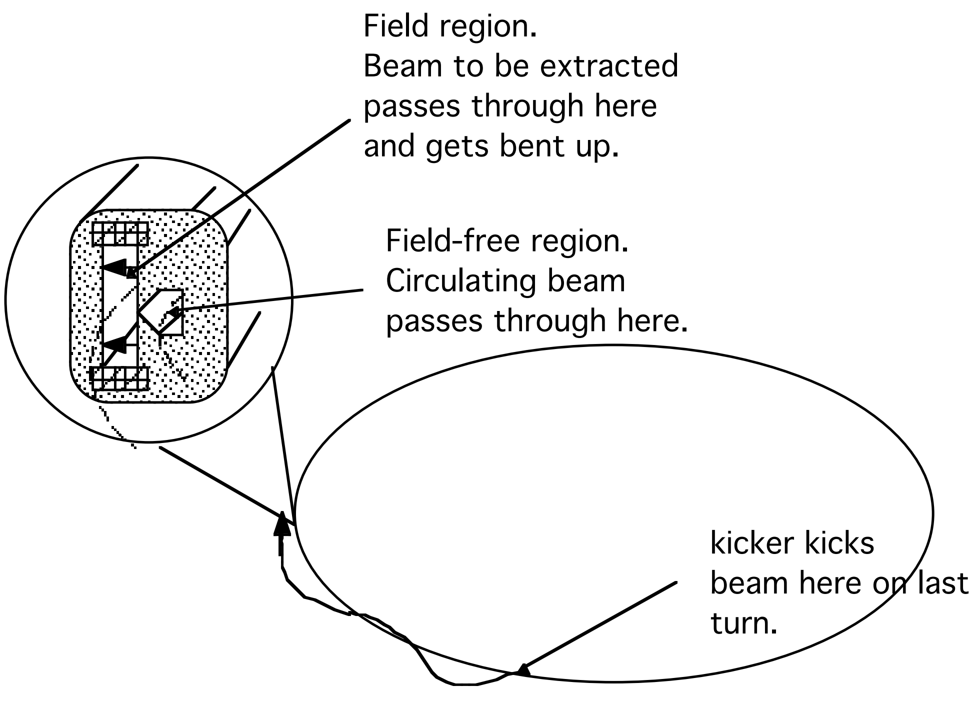

Beam Extraction. Beam initially traveling on the closed orbit is given a kick (by a kicker) and is sent into unstable oscillations. Appropriately placing a Lambertson downstream of the kicker allows the field region to be selected, bending the beam away from the plane of the accelerator. For injection, the same events happen in reverse.

Septa

Septa (plural of septum) provide a way of modifying beam trajectory using an electrostatic or magnetic field. The picture below shows an electrostatic septum like the kind used in Switchyard to split a beam of charged particles into two separate beams. Oncoming particles are directed towards the grounded wires found along the length of a septum, and the electric fields formed between the wires and the high voltage plates cause the stream of particles to be split into two separate streams. Usually the split is not accomplished with a single septum. Strings of individual septa are required to produce the appropriate splitting force. At a point downstream of the septa, after an adequate separation distance between the two beams is achieved, a Lambertson is located in the beamline. One of the two beams travels forward undeflected (through the field free region) while the other is bent into a new beamline. Switchyard and the external beamlines beyond switchyard are full of such “splitting stations”, allowing a single stream of particles from the Main Injector to be divided up, and then provide beam to a number of different experiments simultaneously.

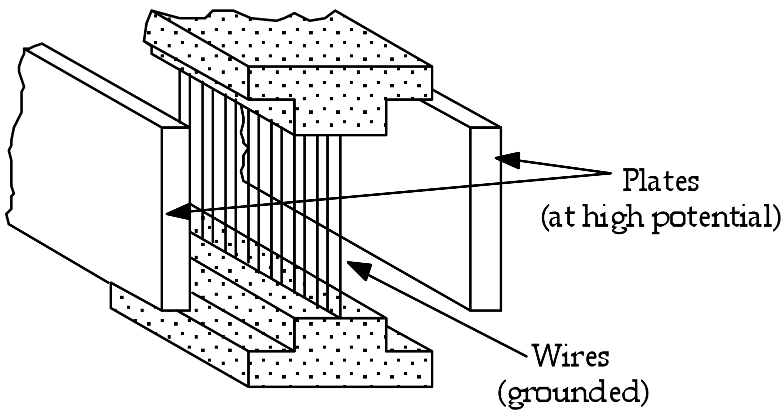

Electrostatic Septa. Beam travels between the high voltage plates and is split into two separate beams. The electric fields formed between the conductive plates and the grounded wires provide the force for this splitting effect.

Magnetic septa are devices that function like Lambertsons. A conducting plate separates the field and the field free regions of a magnetic septum rather than an iron magnet core as found in a Lambertson. In Booster, the extraction septa B:MP02 for the MI-8 line is of this design. MP02 is a pulsed device. This septum has its own PFNs and charging supply, much like kickers. While beam is being accelerated in the Booster it is traveling through the field free region—but in actuality both regions are field free. When it is time to extract beam from the Booster, current pulses are sent individually to the extraction kickers and the septum. When the beam crosses the conducting plate separating the two halves of the septum it encounters the newly formed magnetic field and is deflected.

There are a number of pulsed magnetic septa used for injection and extraction in the Antiproton Source as well (D:ISEP, D:ESEP, A:ISEP1, and A:ISEP2). These septa and the kickers that coordinate with them, make up the bulk of the “pulsed devices” that are turned off when we are not stacking.

Separators

Separators are devices somewhat like electrostatic septa without the grounded wires running down the centerline. Found in the Tevatron, separators are used in the colliding beams mode of operation. The electric field produced between the two conducting plates forces the orbits of the counter-rotating protons and antiprotons apart. Placing both horizontally and vertically oriented separators at strategic locations around the ring allows the formation of helical beam orbits. The protons and antiprotons spiral past each other everywhere in the machine without colliding. The effect of the separators is undone just before the particles enter a collision hall—allowing the beams to briefly pass through each other and collide. Upon exiting the collision hall both beams return to the helical orbit. This spiraling beam path is commonly referred to as the helix.

Ensuring that the protons and antiprotons collide only when we want them to is critical to the lifetime of a Tevatron store. This is because the proton bunches exert tremendous forces on the antiprotons bunches as they pass through each other. The protons are roughly 5 times more intense than the antiprotons. These extreme forces cause beam instabilities within the antiprotons, which can lead to beam loss. Without the helix, the combined effect of these bunches colliding everywhere in the ring all the time would knock the antiprotons out of the machine entirely. See the Beam chapter for more on beam instabilities.

IV. BEAM

A. Introduction

This chapter will look at a number of ideas relating to the streams of ions accelerated in the various machines at the lab. We will examine the calculations for beam energy and transit time, explore the concept of transition in circular synchrotrons, and look briefly at some sources of beam instabilities.

B. Beam Energy and Transit Times

Remember that the basic unit of energy used when dealing with ions traveling through accelerators is the 'eV', or electron volt. One eV is the amount of kinetic energy given to a particle with the same charge as an electron crossing a potential difference of one volt. This unit is most useful for our purposes in much larger quantities. A series of semi-metric prefixes has been developed: keV (kilo-electron volt, 1E3 eV), MeV (mega-electron volt, 1E6 eV), GeV (giga-electron volt, 1E9 eV), and TeV (tera-electron volt, 1E12 eV). This last unit is the final designed beam energy of our largest accelerator, the TeV-atron (written simply as Tevatron).

Below is a table summarizing the initial and final beam kinetic energies for all of the accelerators at the lab.

Accelerator	Initial kinetic energy	Final kinetic energy	Destination of beam
Preacc	~0 keV	750 keV	Linac
Linac	750 keV	66 MeV	NTF
		400 MeV	Booster, Linac dumps
Booster	400 MeV	8 GeV	Booster dump, RDF, MiniBooNE, Main Injector
Main Injector	8GeV	8 GeV for Recycler, Pbar studies	Recycler or Antiproton Source
		120 GeV for Pbar production	Antiproton source
		120 GeV for NuMI	NuMI target hall
		120 GeV for fixed target	Switchyard
		150 GeV for Tevatron injection	Tevatron
Tevatron	150 GeV	980 GeV	Stays in Tevatron (Collider)
Antiproton source	8 GeV	8 GeV	Main Injector
Recycler	8 GeV	8 GeV	Main Injector
Switchyard	120 GeV	120 GeV	External Beamlines

The goal of a particle accelerator is to increase the kinetic energy of a group of charged particles to the point where interesting things happen when said particles strike some sort of target. Usually, the interesting collision involves the creation of 'new' particles—the particles that come out of the collision that did not go in. When high-energy protons strike and destroy nickel nuclei, for example, antiprotons (among a whole host of other particles) can fly out of the interaction. The total energy available to create new secondary particles in such collisions is not limited to just the kinetic energy stored in the particles striking the target.

The total energy of a particle, E_T, is the sum of the particle's kinetic energy, T, and another quantity called the particle's "rest mass". The rest mass, E_0, of a particle is given by the famous equation: $E_0 = mc^2$, where *m* is the mass of the particle in question and *c* is the speed of light.

$$E_T = T + E_0 \qquad \text{(Eq. 1)}$$

where

E_T = total energy of particle (in eV)

T = kinetic energy of particle (in eV)

E_0 = particle "rest mass" (in eV; $E_0 = mc^2$)
mass conversion factor: 1.661E-27 kg = 1 Atomic Mass Unit = 931.5 MeV/c^2

Particle	**Rest Mass, E_0**
electron (e-)	0.511 MeV
proton (p)	938 MeV
neutron (n)	939 MeV
H- ion (H-)	939 MeV

Another way to find the total energy of a given particle is to use a definition common to the study of relativity called the "relativistic mass factor". This quantity, labeled γ (Greek letter gamma), relates the rest mass and the total energy directly as seen below:

$$E_T = \gamma E_0 \qquad \text{(Eq. 2)}$$

Where, by definition, γ is related to the particle velocity, *v*, and the speed of light, *c*, via:

$$\gamma = \sqrt{\frac{1}{1-\frac{v^2}{c^2}}} \qquad \text{(Eq. 3)}$$

Doing a bit of algebra using the three equations above, we find two results:

$$\gamma = \frac{T + E_0}{E_0} \qquad\qquad v = c\sqrt{1-\frac{1}{\gamma^2}}$$

Using these two results, we can solve for the relativistic velocity of an accelerated particle if we know that particle's kinetic energy and rest mass.

The table below shows this velocity (as a percentage of *c*, the speed of light) calculated for a number of different kinetic energies.

particle	kinetic energy	γ (gamma)	velocity (% c)
H-	750 keV	1.0008	4 %
H-	66 MeV	1.07	36 %
H-	400 MeV	1.43	71 %
p	8 GeV	9.53	99 %
p	120 GeV	128.93	99.99 %
p	150 GeV	160.91	99.999 %
p	980 GeV	1045.8	99.99995 %

Another important equation is the following:

$$(E_T)^2 = (pc)^2 + (E_0)^2$$

This shows the relationship between the total energy, E_T, the particle's relativistic momentum, *p*, and rest mass, E_0, which is relativistically invariant. That means that the above equation holds true no matter what reference frame you are in. Using this equation, one can solve for the particle momentum, as shown in the few steps of algebra below.

$$pc = \sqrt{(E_T)^2 - (E_0)^2}$$

$$pc = \sqrt{(T + E_0)^2 - (E_0)^2}$$

$$p = \frac{1}{c}\sqrt{T^2 + 2TE_0}$$

A particle's momentum is expressed in units of "eV/c", electron volts divided by *c*, the speed of light. The table on the next page shows relativistic momentum as a function of particle kinetic energy and rest mass.

particle	relativistic kinetic energy	relativistic momentum
H-	750 keV	37.5 MeV/c
H-	66 MeV	358 MeV/c
H-	400 MeV	954.5 MeV/c
p	8 GeV	8.89 GeV/c
p	120 GeV	120.9 GeV/c
p	150 GeV	150.9 GeV/c
p	980 GeV	980.9 GeV/c

Detectors

$$(E_T)^2 = (pc)^2 + (E_0)^2$$

The large particle detectors at CDF and D0 make use of the above equation to identify the flood of secondary particles resulting from colliding beams. In order to identify a given particle, one must figure out the particle's rest mass and charge. To see how the detectors accomplish this, let's take a quick conceptual tour.

Both the CDF and D0 detectors are basically composed of a tracking chamber surrounded by a calorimeter. The tracking chamber is situated inside a powerful solenoidal magnetic field. This field is aligned along the direction of the beampipe, and thus does not deflect the protons or antiprotons circling in the Tevatron. Secondary particles, produced in collisions and traveling perpendicular to the beampipe *are* deflected by this magnetic field. The solenoid bends the paths of these secondaries by an amount proportional to their momentum. The direction that the particle bends in the magnetic field is determined by its charge. While spiraling through the tracking chamber, the charged secondary particles leave an electronic footprint, telling experimenters about their charge and momentum. After the particles leave the tracking chamber, they encounter the calorimeter. The calorimeter is situated outside of the magnetic field generated by the solenoid. The purpose of the calorimeter is to stop the escaping particles in their tracks. The total energy of a particle, E_T, is deposited into the calorimeter and then measured.

Obtaining this information about momentum (from the tracking chamber) and total energy (from the calorimeter), experimenters can calculate each particle's rest mass. Because the charge is also known from the tracking chamber data, the particle can be classified.

Transit times

The amount of time it takes a particle to make one complete lap around an accelerator is also known as the **transit time**. We will look at two different ways to calculate the transit time for a given energy in a given machine.

The first method of calculating the transit time, *t*, of a particle is the most straightforward. If we know the particle's velocity and the length of the trip it has to take (the circumference of the accelerator in question), we can easily find the time needed to complete the trip. Consider the dimensional analysis below:

$$C = \frac{\text{meters}}{\text{lap}} \qquad v = \frac{\text{meters}}{\text{second}} \qquad t = \frac{C}{v} = \frac{\frac{\text{meters}}{\text{lap}}}{\frac{\text{meters}}{\text{second}}} = \frac{\text{seconds}}{\text{lap}}$$

Accelerator	Circumference (C)	kinetic energy	velocity (v)	Transit time C/v
Booster	475 m	400 MeV	0.71 c	2.2 µs
		8 GeV	0.99 c	1.6 µs
Main Injector	3319 m	8 GeV	0.99 c	11 µs
		120 GeV	0.9999 c	11 µs
		150 GeV	0.99999 c	11 µs
Tevatron	6283 m	150 GeV	0.99999 c	21 µs
		980 GeV	0.9999995 c	21 µs

The second method for calculating transit times is a bit subtler, and involves the RF frequency and the harmonic number. Consider the following equation:

$$h = \frac{\text{RF cycles}}{\text{lap}} \qquad f_{RF} = \frac{\text{RF cycles}}{\text{second}} \qquad t = \frac{h}{f_{RF}} = \frac{\frac{\text{RF cycles}}{\text{lap}}}{\frac{\text{RF cycles}}{\text{second}}} = \frac{\text{seconds}}{\text{lap}}$$

Accelerator	Harmonic number (h)	kinetic energy	RF frequency (f_{RF})	Transit time h/ f_{RF}
Booster	84	400 MeV	37.8 MHz	2.2 μs
		8 GeV	52.8 MHz	1.6 μs
Main Injector	588	8 GeV	52.8 MHz	11 μs
		120 GeV	53.1 MHz	11 μs
		150 GeV	53.103 MHz	11 μs
Tevatron	1113	150 GeV	53.103 MHz	21 μs
		980 GeV	53.104 MHz	21 μs

C. Transition

Imagine an imperfect proton being accelerated in a synchrotron. The proton completes lap after lap, each time passing through an RF cavity and gaining a small amount of kinetic energy. This increase in kinetic energy will speed the particle up slightly, and increase the particle's momentum slightly. All other factors being the same, speeding the particle up will decrease the particle's transit time. Conversely, boosting the particle's momentum will increase the radius of its orbit, and therefore increase the orbit circumference. This increases the transit time.

So we have two competing effects during acceleration: speed increases (making shorter transit times) and path length increases (making longer transit times). Which of these effects is most dramatic at any given point in a specific machine's acceleration cycle depends on many factors, such as the voltage and spacing between the RF cavities, the machine circumference, and the lattice arrangement.

At a certain point in the acceleration cycle, these two effects exactly cancel each other out. The incremental change in speed a particle receives from the RF is exactly accounted for by the incremental change in that particle's path length. At this point, called **transition** or the **transition energy**, all particles regardless of their energy errors have the same transit time. Points in the acceleration cycle before this time are said to be 'below the transition energy'. Points later in the acceleration cycle are said to be 'above the transition energy'. Transition itself is an unstable point where some beam loss is common, and a special RF manipulation has to be performed to avoid loosing the entire beam.

Recall the requirements for phase focusing, as discussed in the RF chapter. We assumed that higher energy particles would be traveling faster and therefore reach the RF cavity earlier than the synchronous particle. Because these fast particles were too energetic already, they were subjected to a proportionally smaller electric field, and therefore received a smaller kick.

Basically, we assumed that this acceleration was taking place below the transition energy. When the particles are above the transition energy, those arriving earlier in time do so not because they are traveling faster, but because they had a shorter path length to travel. Thus, above transition, particles arriving early need a larger boost than the synchronous particle. Strangely enough, above the transition energy, particles with more energy than the synchronous particle—which are traveling *faster* than their lower energy counterparts—actually take *longer* to complete one lap.

In order to maintain phase focusing, the RF phase has to be changed at the transition energy as shown on the next page.

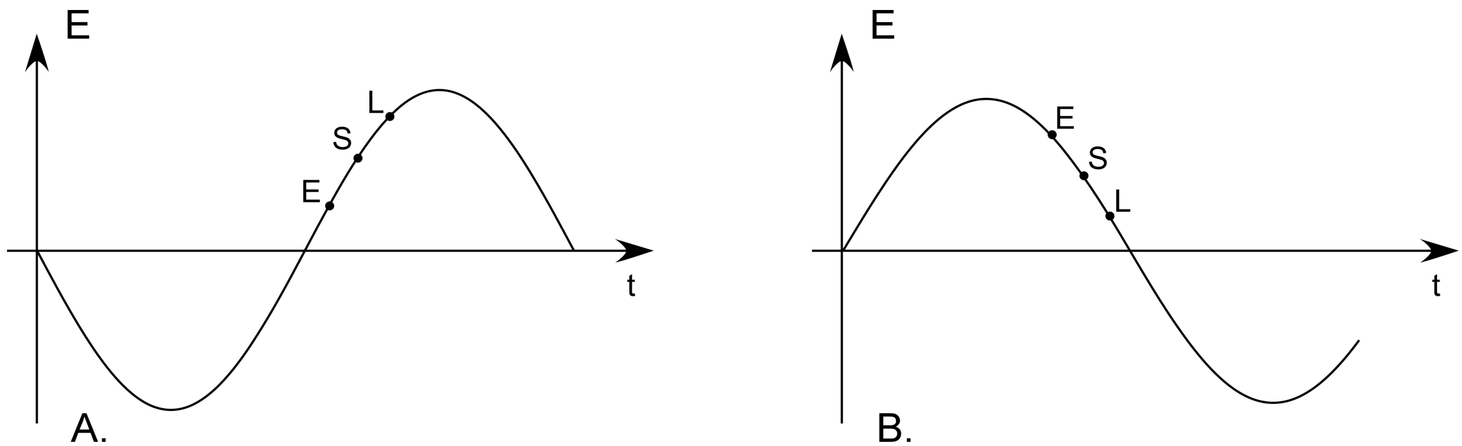

Transition. The moment in the acceleration cycle of a circular machine when the incremental change in a particle's speed is exactly canceled out by its incremental change in orbit path length is what we call transition. An RF phase change is needed at transition to avoid loosing beam. The picture on the left, labeled *A,* shows the acceleration phase below the transition energy. The picture on the right, labeled *B,* shows the new phase after transition. Notice how the magnitude of the electric field seen by the synchronous particle, labeled *S,* is the same above and below transition.

As the picture points out, the synchronous particle receives the same magnitude kick from the RF before and after transition. In fact, if all of our particles were perfectly synchronous, there would be no such thing as transition energy. The magnetic guide field would at all times be such that these synchronous particles would be bent along the ideal orbit. Beams that have no momentum spread (i.e., no off-momentum particles) do not have to deal with transition, because the length of the orbit's circumference never changes. Thus transit times are determined solely by speed increases.

It should also be pointed out that accelerators could be designed to avoid the problems associated with transition altogether. For instance, one can build an accelerator whose injection energy is higher than the transition energy. This is the case with the Debuncher, the Accumulator, and the Tevatron.

Another trick that was used in the Booster in the past (with limited success) was called the 'gamma-T jump'. Gamma-T, or γ_T, is the relativistic mass factor associated with the transition energy. In this transition loss avoidance scheme, a number of special quadrupoles were used to change the Booster lattice just before transition in the acceleration cycle. Changing the Booster lattice in this way also changed the value of the transition energy, so that the kinetic energy of the accelerated beam quickly 'jumped' over transition before any beam had a chance to fall out of the machine.

As noted above, each machine has its own transition energy, with the specific value determined by a number of factors such as the lattice and the voltage and spacing of RF cavities.

The table below summarizes these values. Transition is a concept that only applies to circular accelerators, so it is not a factor in the design and operation of the Linac.

machine	γ_T	transition kinetic energy	time in cycle
Linac	n/a	n/a	n/a
Booster	5.45	4.2 GeV	17 msec
MI	21.6	18.2 GeV	depends on **ramp rate**: $29 ~0.25 sec $2B ~0.67 sec
Tevatron	18.7	16.6 GeV	n/a

Ramp Rate. MI defines the transition time in terms of a specific RF frequency, thus the exact transition time during the acceleration cycle can vary if the magnet ramp rate is changed.

D. Beam Stability

Bunches of particles traveling around an accelerator are subjected to many different forces. The RF acts upon the beam longitudinally and the magnets act upon the beam transversely. Remember that imperfect particles are focused by both the RF and the quadrupoles, and forced into oscillations around the synchronous phase (longitudinal), and oscillations about the ideal orbit (transverse).

The particles themselves, being of the same charge, tend to want to repel one another within the bunch (longitudinally and transversely). Less than perfect vacuum can supply the beam with a variety of air molecule targets, causing scattering and beam loss. One major source of possible beam instability comes from the betatron tune. As we will see below, the exact value of the betatron tune in both planes is important. If the tunes stray too far from where they are supposed to be, beam loss can occur.

This term ***betatron tune*** has already been discussed in the magnet section. The tune is simply the number of betatron oscillations that a particle undergoes during one revolution in the accelerator. The horizontal and vertical planes have their own tunes. These two values are largely independent, but coupling between the two planes is common. The tunes are said to be strongly coupled when changes made to the horizontal tune also affect the vertical tune (or vice versa).

The tune of a circular accelerator is a design parameter determined by the location and number of quadrupole magnets. For example, during Collider operation the horizontal tune of the Tevatron is ~20.583 and the vertical tune is ~20.575. These values represent stable operating conditions. Using quadrupole trim magnets, the lattice can be altered slightly, and thus the tunes can be varied.

Tunes and stability

If the vertical tune of the accelerator were a simple integer of 1, then an individual particle would follow the same beam path on every revolution. At some point on the circumference of the accelerator, the beam displacement from the design orbit would be a positive maximum, whereas on the opposite side of the ring the orbit would have a negative maximum displacement. The picture below shows this condition.

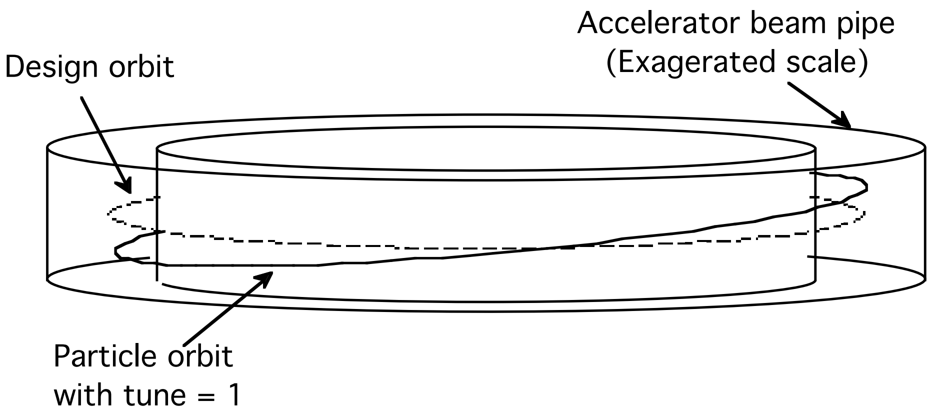

On the surface, it doesn't seem like the above particle orbit would be all that bad. The problem arises when we consider non-perfect magnetic fields. Imagine that we have a magnet somewhere in the ring whose field is not perfect. Imagine then that this small magnetic field error acts like a small vertical dipole, delivering a kick to the beam on each pass. Since we are assuming this error to be unique, the small vertical kick received by the beam is never corrected. Because the tune is equal to 1, the orbit passes through the dipole error in the same way on every pass. Each time the beam passes the location of the field error the distortion to the beam orbit is increased. Soon, the small field error has, over many revolutions, caused a large orbit change, as seen below. Eventually the beam will be lost in collisions with the beampipe wall.

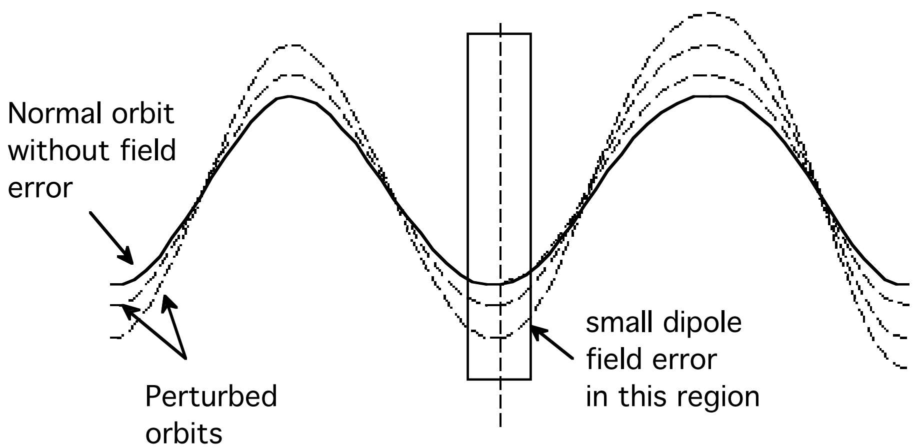

In reality, we do not simply have one imperfect magnet with a dipole field error. Every magnet in every accelerator can be considered to be slightly flawed. Even if all of the magnets were perfect, slight misalignments could generate these types of incorrect fields as well. Because of dipole field errors, the tune of any accelerator cannot be equal to 1, or in fact equal to any integer value.

If the tune were 0.5 (or any half-integer value), then the orbit position at some point would be a positive maximum during the first revolution and a negative maximum during the second revolution. The diagram below illustrates this condition.

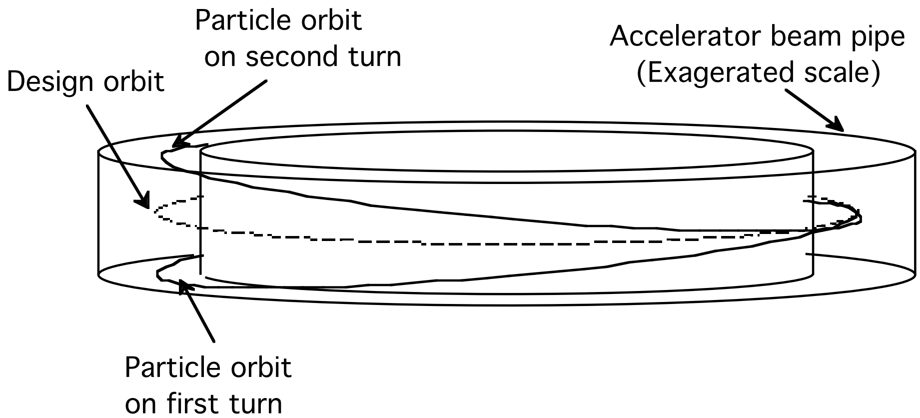

If a particle has a fractional tune value of 0.5, the same trajectory is traced out every other revolution in the accelerator. If there is a dipole error at some point, like in the example above, on one revolution the particle will see a kick to increase the oscillation. On the next pass through the region, however, the same kick would tend to reduce the oscillation. See below:

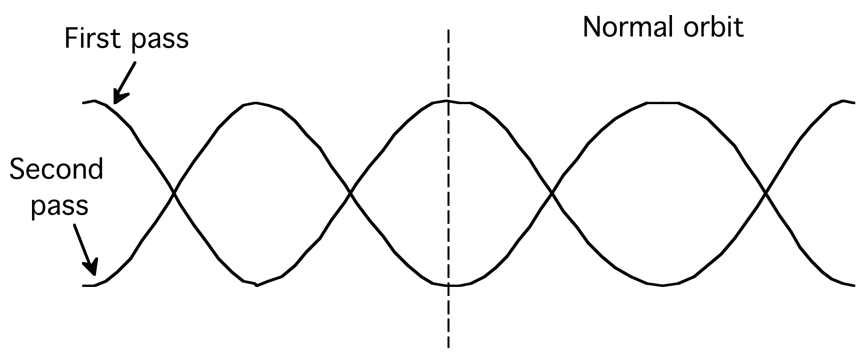

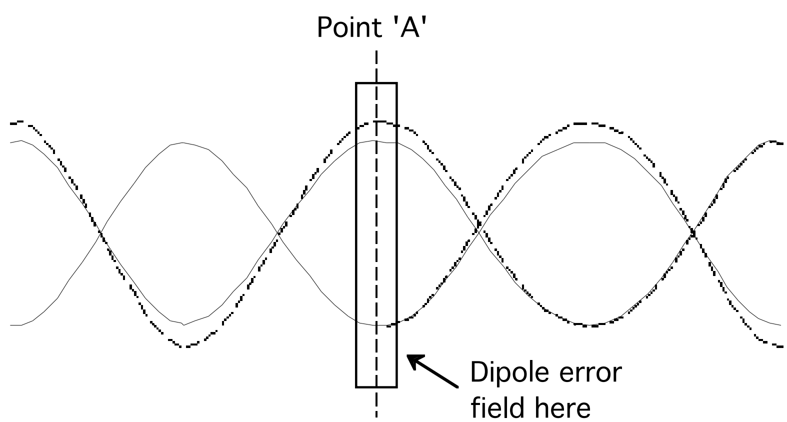

It seems that a particle with a 0.5 fractional tune is not affected by dipole field errors. Unfortunately dipoles are not the only magnets that can be imperfect. Consider a quadrupole field error, represented by a small defocusing quadrupole magnet. If the particle with a fractional tune of 0.5 passes through this quadrupole vertically off center, it will receive a kick in the vertical plane, causing the orbit to be distorted. On the second pass through the quadrupole, the particle will find itself on the other side of the magnet centerline, where it receives another kick in the vertical plane. From pass to pass, this quadrupole field error will keep distorting the orbit, increasing the amplitude of the betatron oscillations with each pass. The picture below shows this effect. The numbers represent successive passes through the region.

As with the dipole field error above, every quadrupole in every accelerator is slightly flawed, forcing us to avoid allowing particles to have half-integer tunes. Although it will not be discussed here, it can be shown that sextupole errors affect particles with 1/3 integer fractional tunes.

In general, the tunes that are driven by field errors are given by the equation,

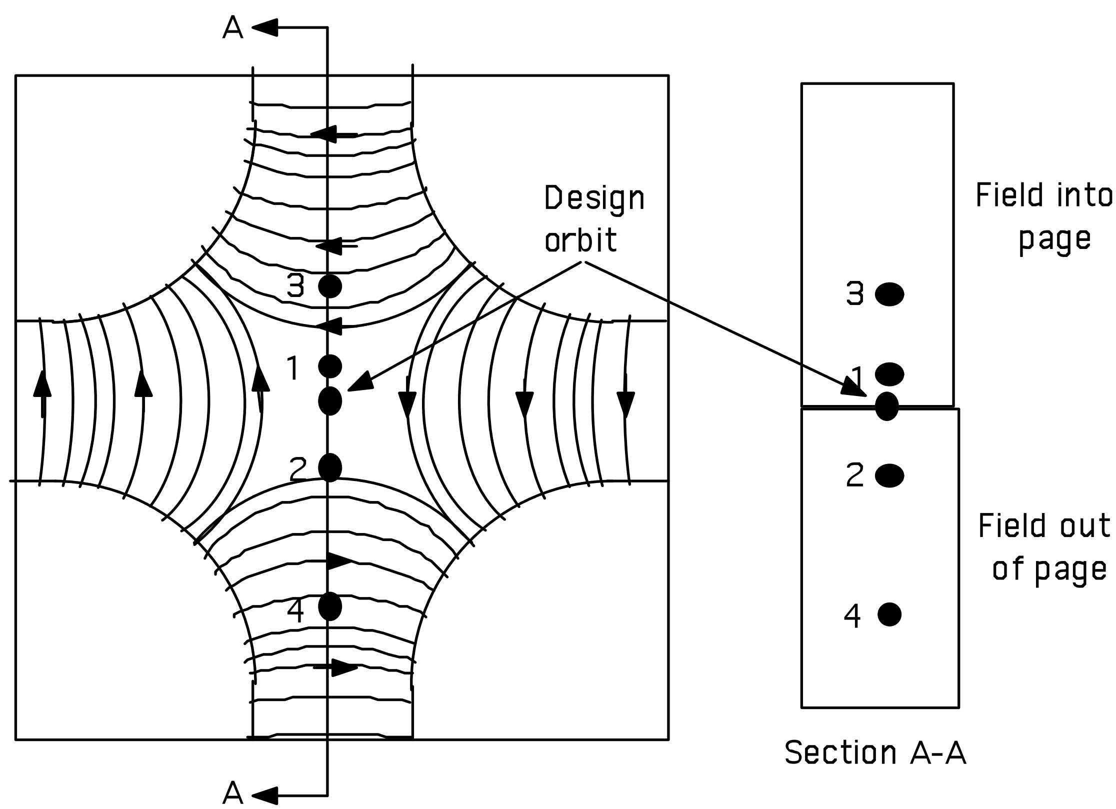

$$v = \frac{k}{n}$$

where v is the tune, k is an integer, and n, the order of the resonance, is an integer.

There is a degree of strength associated with the above equation in that the lower the order, the higher the strength of the resonance. So, for example, a first order resonance (n = 1) is stronger than a 15th order resonance. Dipole field errors are first order resonances. Quadrupole field errors, which are slightly less destructive to the beam, are second order resonances.

Sextupole field errors represent third order resonances, and so on.

Order (n)	Tune	Field error
1	integer	Dipole
2	1/2	Quadrupole
3	1/3	Sextupole
4	1/4	Octupole

One can plot the resonances to be avoided on a graph. Below is a representation of tunes up to third order for a particle in a circular accelerator with one transverse degree of freedom.

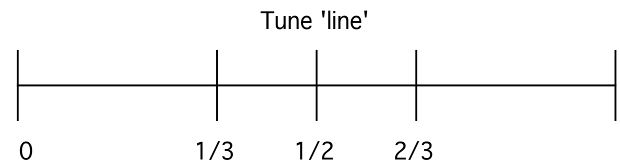

In order for beam to remain in the machine, we need to choose an operating tune that does not fall on any of the labeled points above. If we expand this concept to include the other transverse plane, we get

$$v_A + v_v = \frac{k}{n} + \frac{1}{m},$$

or

$$mv_A + nv_v = r\ .$$

Again, m, n and k are integers, and the order of the resonance is given by $|m| + |n|$. In the two-dimensional case, *m* and *n* can be positive, negative, or zero. A popular way of viewing the above equation graphically is to create a two dimensional version of the tune line, known as a tune diagram. The point on the tune diagram corresponding to the horizontal and vertical tune values is called the operating point.

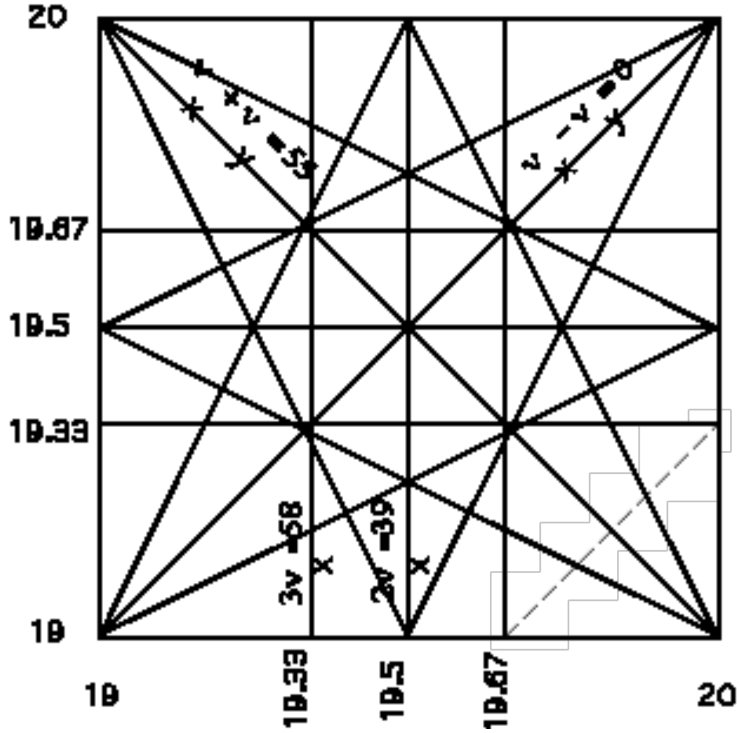

So far, we have talked about the tune of a single particle. In actuality, accelerators accommodate billions of particles in a single cycle. In that case, the ‘tune’ of the accelerator is actually the average of all of the individual particle’s tunes. The graphic below shows a sample distribution of tunes (along the x-axis) versus signal intensity (the y-axis). Notice how the beam has a distribution of individual tunes. The width of this tune spread is related to the chromaticity of the machine, and arises from the fact that the beam has a natural distribution of individual particle momenta as well.

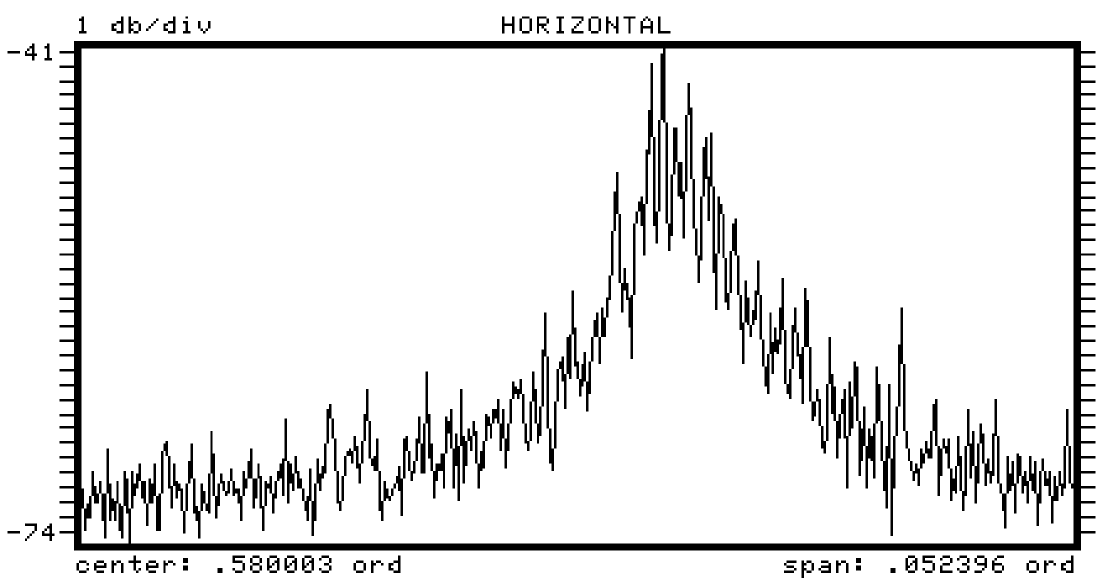

Because of this tune spread found in a many particle system, the operating point chosen on the tune diagram must be some distance away from all major resonance lines. The partial tune diagram below shows that some of the particles centered around the operating point are actually crossing over resonance lines, and therefore are unstable.

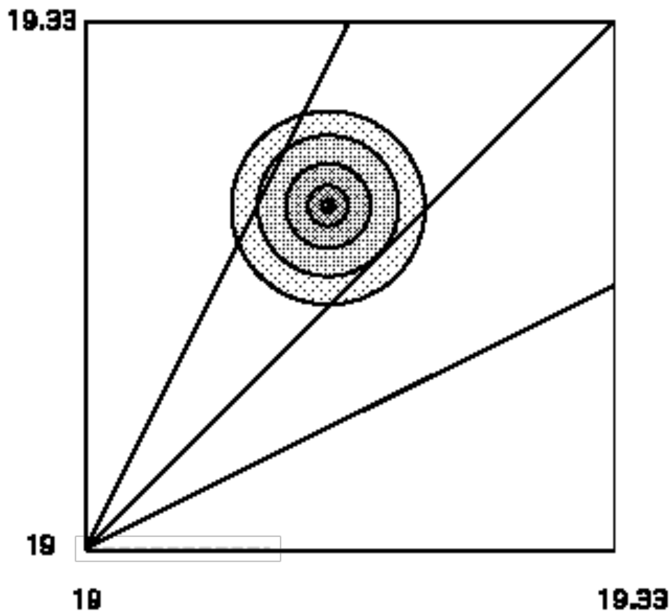

E. Resonant Extraction

In most cases, the operating point on a tune diagram is chosen so that beam does not cross a resonance line and fall out of the machine. In the case of resonant extraction, however, beam falling out of the machine is exactly what we want.

Resonant extraction allows an accelerator to slowly extract a more or less continuous stream of beam over the span of one or more seconds, taking advantage of a natural instability. Normally, when beam is extracted from an accelerator, it's all or nothing. A kicker fires, and the entire beam leaves the machine on the same turn (within microseconds). During resonant extraction, the operating tune is adjusted to make the beam distribution cross a resonance line in a controlled manner. With the MI 120 GeV fixed target program, the beam is slowly inched towards a half-integer resonance line. The betatron oscillation amplitude of the beam on the resonance line increased dramatically, until finally the orbit was so distorted that the beam crossed into the field region of the electrostatic septa. Once in the field region of the septa, the beam received a kick allowing it to be extracted at a Lambertson downstream.

This process allows a stored bunched beam to be slowly shaved away, each bunch loosing its particles over the span of thousands of turns. A special computer, called QXR, for Quadrupole eXtraction Regulator, calculates how much it should adjust the horizontal tune on the fly in order to maintain constant extracted beam intensity.

V. SUPPORT

A. Subsystems

This chapter is dedicated to three subsystems necessary for operation of all accelerators: vacuum, water, and power. Without an evacuated beam pipe particles would collide with air molecules and be lost. Without cooling water magnets would overheat and be damaged or destroyed. Without power nothing would happen at all.

B. Vacuum

There are two basic reasons why the beam pipe in an accelerator is under vacuum. First of all, the beam traveling inside the machine would be scattered and lost through collisions with air molecules if no vacuum were present. Secondly, without an evacuated beam pipe we would be unable to sustain the high electric fields required by various devices, RF cavities, separators, septa, etc. These high electric fields would ionize the air, forming a path to ground. To overcome these problems it is necessary to have evacuated regions wherever the beam is intended to travel and in locations where high electric fields are present.

Pressure

The air pressure in a vessel is a measure of the force per unit area exerted by the air upon the walls of the vessel. Using a form of the ideal gas law equation:

$$P = \frac{nRT}{V}$$

one can see that the pressure, P, in a vessel can be reduced by increasing the volume, V, of the vessel, decreasing the temperature, T, of the air in the vessel, or by reducing the number of particles, n, in the vessel. The most practical approach in the case of evacuating a beam pipe is the removal of air molecules—but as we will see later the vacuum pressure in the Tevatron also benefits from the extremely low temperature of the surrounding magnets and beam pipe.

Typical units of pressure are millimeters of mercury (mm Hg) with standard room pressure being around 760 mm Hg at sea level. In vacuum work, the more common units are “Torr” and “micron.” One Torr is equal to one mm of Hg, and one micron is equal to one μm of Hg. One atmosphere is thus equal to 760 Torr or 760,000 microns. A good beam pipe vacuum here at Fermilab would measure around 10^{-8} to 10^{-10} Torr. Here is a table of the different units.

1 atmosphere……….760 mm of Hg
1 atmosphere……………...14.7 psi
1 Torr………………….1 mm of Hg
1 micron…………………10^{-3} Torr

The basic components of any vacuum system are the airtight beam pipe (typically constructed out of stainless steel, aluminum, or copper) and a vacuum pump of one kind or another. Different vacuum pumps use different methods to remove air molecules. Each of these pumps works best over a limited range of air pressures and so evacuation of a beampipe is done in stages. One type of vacuum pump removes air molecules until the pressure is reduced to where the next type of pump will be effective. In the next few pages we will examine a couple of the more common types of vacuum pumps.

Roughing pumps

The first stage of pumping is done using an oil-roughing pump. This type of pumping is effective from one standard atmosphere to about 10^{-3} Torr. The pump consists of an oil-lubricated wheel with vanes on the perimeter. As the wheel rotates, the vanes trap air molecules from the vacuum vessel between the wheel and the pump wall and transport them to an exhaust port. Once there, the air molecules exit via a one-way valve and are expelled.

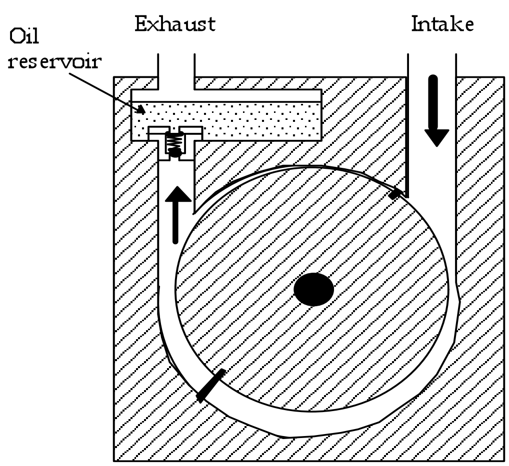

Roughing Pump. This pump has rotating vanes that trap air molecules and mechanically force them out the one-way exhaust valve.

As the pressure inside the beampipe drops, the number of particles randomly entering the intake port also drops. At an intake pressure of one micron or so, the roughing pump is no longer able to effectively remove more air molecules from the beampipe.

Turbomolecular (turbo) pumps

Turbo pumps, which work from 0.1 Torr to 10^{-6} Torr, usually help out roughing pumps. Affixed to the intake port of a roughing pump, turbo pumps work as turbochargers, increasing the number of air molecules present at the roughing pump's intake port. The turbo pump consists of different layers of vanes rotating at very high velocity (50,000-100,000 rpm). Like the roughing pump, the turbo works on the principle of momentum transfer. As the vanes spin and hit air molecules, they are driven towards the exhaust of the turbopump and into the intake of the roughing pump. Since the number of molecules at the intake of the roughing pump increases, so does the intake pressure (back up into the efficient operating region of the roughing pump). Due to the delicate nature of the spinning vanes, turbo pumps cannot be turned on until the roughing pump has lowered the pressure inside the beam pipe sufficiently.

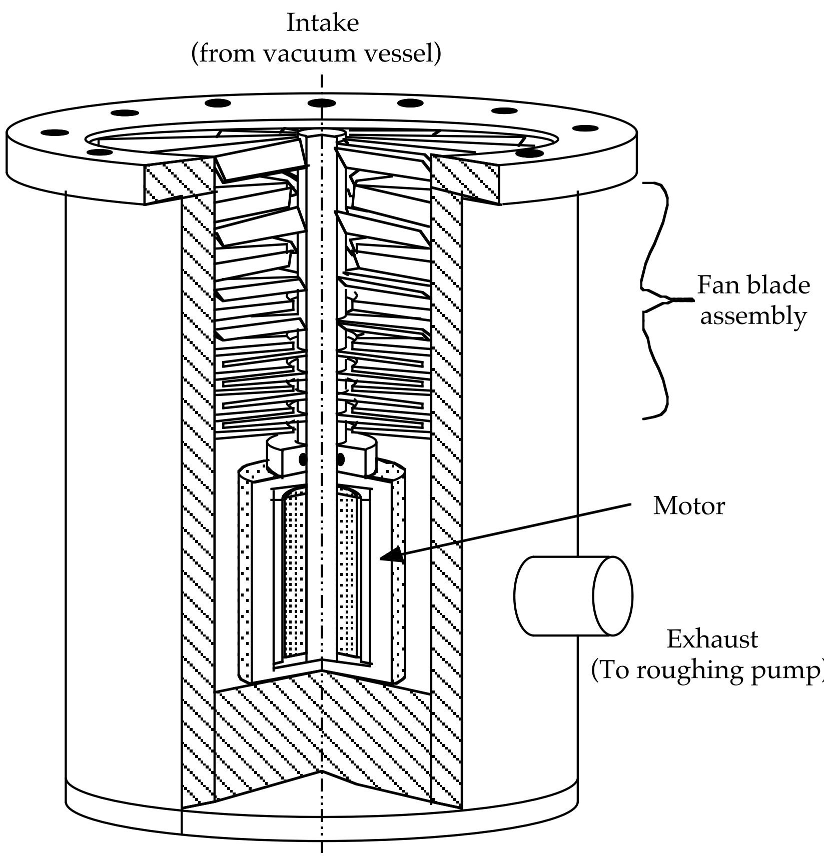

Turbo pump. This pump acts like a turbocharger for the roughing pump. Spinning vanes deflect molecules down toward the exhaust port and the connected roughing pump.

Ion pumps

Ion pumps, which are effective from 10^{-5} to 10^{-12} Torr are fundamentally different from the roughing and turbo pumps described above. Instead of mechanically removing the air molecules from inside the beampipe, an ion pump captures and traps air molecules. Physically, the ion pump consists of two parallel titanium plates with an electrically insulated titanium cell structure in the middle. Titanium is chosen because it readily interacts and bonds with ions. To complete the structure, a magnet is placed on either side of the outer titanium plates, creating a magnetic field perpendicular to the plates.

A high voltage (5 kV) is applied to the cell structure in the middle. This voltage creates an electric field between the plates and the middle cell. As neutral air molecules inside the beam pipe randomly enter the ion pump, the electric field ionizes them. Once ionized, the free electrons travel towards the positively charged middle cell while the ions move toward the outer plates. Both particles are then captured, the ions binding with Titanium atoms on the outer plates and the electrons collecting on the central anode. Removing molecules from the beampipe in this way decreases the air pressure.

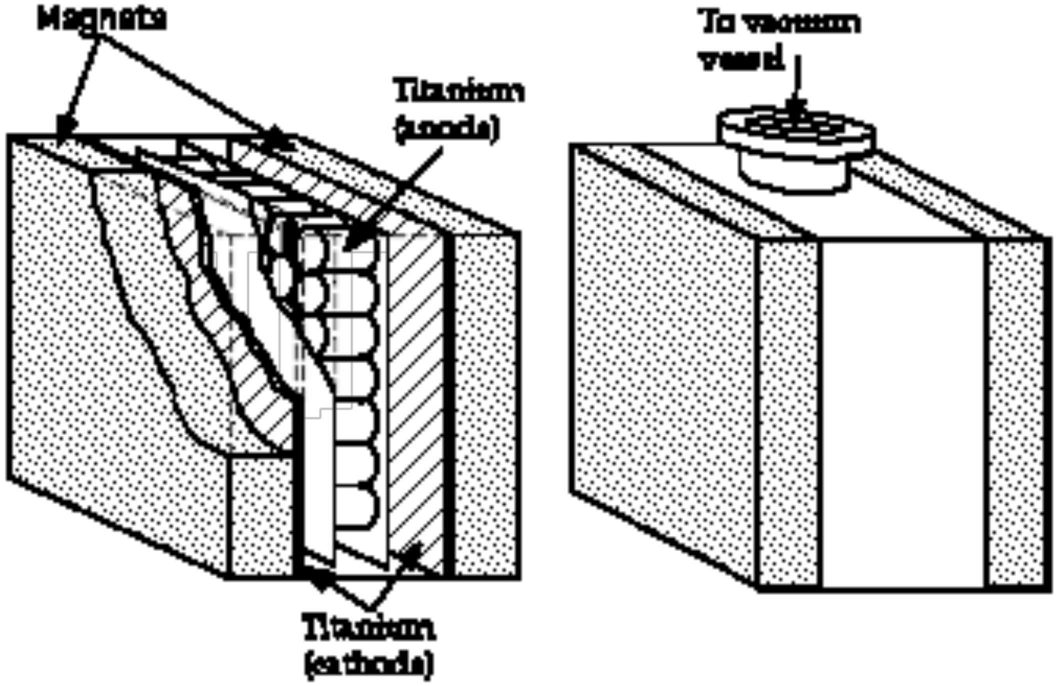

Ion Pumps. This pump works by ionizing air molecules then capturing the free electrons on the anode and the ionized molecules on the titanium plates. An electrical current proportional to the number of ions created can be measured, indicating the quality of the vacuum. Molecules trapped in this way no longer contribute to the vacuum pressure inside the beampipe.

As the newly charged particles begin to move to their respective sides they feel a magnetic force due to the noted magnetic field. This force causes the ions and electrons to spiral around inside the pump, effectively increasing the path length the ions travel before becoming trapped. A longer path length gives the charged particles more opportunity to collide with neutral air molecules just entering the pump. If such a collision occurs, the neutral molecule could be ionized and trapped as well, increasing the efficiency of the pump.

An electric current proportional to the number of electrons collected by the middle cell is used as a measure of vacuum quality. This is because the number of electrons collected is proportional to the number of ions created, which in turn is proportional to the number of air molecules inside the beampipe. If this electric current gets too high, the ion pump turns itself off as a precaution. The high voltage components are not made to work in (relatively) high-pressure environments where electrical sparking could damage the pump.

Ion pumps have a finite lifetime after which their surfaces are so contaminated with trapped molecules that they can no longer pump effectively. Typically this lifetime is 50,000 hours.

Sublimation pumps

Titanium sublimation pumps work on a principle similar to the ion pump. The sublimation pump consists of rods of titanium connected to an external power supply. When this power supply is turned on, small amounts of the titanium rods sublimate—they go directly from a solid to a vapor without boiling. The sublimated titanium vapor condenses on nearby surfaces like the inside of the beampipe, where it will readily interact with any air molecules that come into contact with it. Again, the air molecules are chemically trapped on the vessel wall, unable to contribute to the pressure.

Cryo pumping

According to the ideal gas law equation at the beginning of this chapter, we recall that the pressure inside a vacuum vessel can be lowered if the temperature of the air molecules within the vessel is reduced. As we have seen with the ion and sublimation pumps, one does not have to remove ions from the vacuum chamber to reduce the pressure. Simply trapping them in a state that they cannot contribute to the air pressure works as well.

In a machine like the Tevatron where sections of beampipe are just a few degrees above absolute zero, the sub-frigid temperatures similarly improve the quality of vacuum. Air molecules that would otherwise be free to collide with accelerated protons and antiprotons liquefy and condense on the inner surfaces of the beampipe in a process known as cryo pumping. As long as the surface remains cold, the molecules will remain trapped in liquid form, unable to bother the passing beam. If the surface heats up suddenly, like when the Tevatron magnets quench, these molecules are suddenly released back into gaseous form and cause what is known as a vacuum burst. The sudden influx of new particles causes a sharp increase in vacuum pressure, causing ion pumps near the quench to trip off.

Another common use of the cryo-pumping phenomenon is the cold trap. A cold trap is essentially a Dewar or thermos filled with liquid nitrogen that is in contact with the intake of a vacuum pump. The cold trap works to condense out any oil molecules that may have leaked from the pump, capturing them before they make it into the vacuum system.

Measuring Vacuum Quality

Various instruments can be attached to the beam pipe to measure the pressure of the air inside. Three such instruments used at Fermilab are thermocouple gauges, cold cathode gauges, and ion gauges.

Thermocouple gauges work on the principle of heat convection. A heated wire is exposed to an evacuated chamber. A temperature sensing circuit called a thermocouple is placed in contact with the heated wire where it constantly measures the wire temperature. Because the heated wire will naturally transfer some of its energy to the air molecules around it as the pressure around the wire is reduced, the thermocouple will detect a proportionally higher wire temperature. Thermocouples are usable from atmospheric pressure to about 10^{-3} Torr.

Cold cathode gauges are somewhat like miniature ion pumps. A high voltage wire is placed into the vacuum chamber. The amount of current flowing from this wire to the outer casing of the gauge is proportional to the amount of ionized gas in the system, which is a measure of the vacuum quality. Cold cathodes operate in the range from 10^{-3} to 10^{-8} Torr.

Ion gauges are much like cold cathode gauges. Instead of a high voltage wire, ion gauges have a filament and a collector. The filament is heated, liberating electrons, which are accelerated towards the collector. As they travel this short path, the electrons from the filament will ionize some of the neutral air molecules in the area. The collector picks up the electrons from the filament and the newly liberated electrons, and measures the total electrical current. The magnitude of this current is related to the pressure. Ion gauges are operable from 10^{-2} to 10^{-10} Torr.

Out gassing and bake outs

Molecules have a tendency to attach to the surface of the metal, and slowly break away when the surface is under reduced pressure. The surfaces of some metals are porous, providing small cavities where gases can be trapped at standard pressure only to slowly bleed into the vacuum once the pressure is reduced. The process of releasing gases from the surface of a material under reduced pressure is known as out gassing. Out gassing can be a source of seemingly spurious vacuum bursts, and is a problem to overcome when re-evacuating a component that has been ‘let up to air’ (i.e., brought from a vacuum up to atmospheric pressure).

In some cases, the outgassing is not caused by trapped gases being liberated from inside the metal, but rather by some contaminant on the surface evaporating. Oils and certain rubber components can out gas in this way; even the natural body oils from a person touching the inside of a beampipe can cause problems during pump down.

Out gassing is sometimes intentionally induced in high vacuum components to remove trapped gases or to boil off any oils collected on the inner surfaces. This procedure is known as a bake out. Special electric heating blankets (commonly called bake out blankets) are wrapped around the outside of the beam pipe or device to be baked. When turned on, these blankets heat the component in question. Raising the temperature of the vacuum vessel frees any gases trapped on the inner surface and forces the molecules back into the system where they can be pumped out. Due to the expected increase in the number of air molecules in the vacuum vessel during a bake out, extra portable vacuum pumps are usually added to the system. Once the system has been sufficiently baked, the special blankets are turned off, the vacuum valves connecting the extra pumping stations to the vessel are shut, and the ion pumps and sublimation pumps are turned on.

C. Low Conductivity Water (LCW)

As seen in a previous chapter, the large magnets used in building accelerators are commonly water-cooled. This cooling water runs through the bus work that forms the magnet windings. Because other portions of the water distribution system are electrically grounded, this water must not be a conductor. High currents flow through the bus work. If we used water from a pond or standard faucet for cooling, the impurities and free ions present would short out the magnet. What we need is a supply of purified water that is free of potentially conductive sediments. Low Conductivity Water (or LCW) is the answer.

The LCW flowing through the magnets is essentially confined to a closed loop. This loop includes a pass through the magnets in the tunnel, and a pass through a heat exchanger back upstairs in the service building. The purpose of the heat exchanger is to allow the hot closed loop LCW to cool off by coming into thermal contact with pond water. Physically the heat exchanger is like a large diameter pipe with a smaller diameter pipe coiled within it. The LCW to be cooled is sent through the smaller pipe, the pond water through the larger one. These two pipes are sealed so that the water within each does not mix. The pipes are also thin enough to allow the warmer LCW, coming back up from the hot magnets, to transfer its heat to the colder pond water coming in from outside. After the pass through the heat exchanger, the LCW is ready for another round of cooling hot magnets. Of course, as this is a continuous process, some water is running through all parts of the system at all times.

Separate pumps keep the LCW and the pond water flowing along their respective paths. The need for cooling (and therefore heat exchanging) is the reason there are so many cooling ponds around the various accelerators.

It should be noted that magnets down in the tunnel are not the only users of LCW cooling. Power supplies and RF equipment upstairs and RF cavities in the tunnel also utilize closed loop LCW cooling systems.

The Central Utility Building (or CUB), located behind the Booster pond, is the primary source of LCW. A number of machine specific LCW systems exist, each with their own carefully controlled temperatures and pressures. LCW starts out as ordinary tap water—special filters remove the free ions and impurities before allowing the water into the closed loop systems.

D. Power Distribution

The large electrical currents flowing through the magnets and other accelerator equipment have to come from somewhere. In our case the electricity comes from Commonwealth Edison, the local Power Company. Fermilab uses as much power as a not-so-small city, requiring an electrical supply delivered via very high voltage lines. Coming in from off site are 345 kV transmission lines ending at one of the electrical substations. Massive transformers at these substations convert the low current 345 kV into higher current 13.8 kV. A number of underground cables called feeders carry this power around the site and to the various service buildings. Outside the individual service buildings, yet more transformers convert the 13.8 kV into whatever voltage is required, usually very high current, such as 480 V for large power supplies.

The two main electrical substations found at the lab are the Master substation at the north end of the site and Kautz Road substation to the south. A variety of numbered feeder cables originate at both substations. For instance, feeder 23, the main power cable for the Tevatron, comes from the Master Substation. Feeders 86, 87, and 89 come from Kautz Road, powering the Main Injector and the P1, A1, P2, and AP1 line supplies. When preparing either the Main Injector or the Tevatron for access, the flow of power must be interrupted for safety reasons. Not only are the power supplies turned off, but a breaker is also opened at the appropriate substation, preventing current from flowing down the feeder cable. This process is called “switching off”.