Tevatron Rookie Book

Table of Contents

INTRODUCTION

The Tevatron is currently the highest energy accelerator in the world. It derives its name from the phrase "tera electron volt," or TeV, which means one trillion electron volts. Although normal operations in the past were either at 800 GeV for Fixed Target experiments of 900 GeV for collider experiments, the TeV has attained an energy as high as 1.012 TeV during accelerator studies. The energy for Run II is 980 GeV. This accelerator is housed in a tunnel that has a radius of 1 km, roughly four miles in circumference. The TeV, as it is often referred to is a superconducting magnet synchrotron. All dipoles, quadrupoles, and correction element magnets are cooled to about 4.6 Kelvin with liquid helium, where upon they become superconducting.

The Tevatron has one distinct mode of operations, colliding beams. In colliding beams mode, the TeV is loaded with 150 GeV protons and antiprotons, which are ramped to 980 GeV and then made to collide with the detectors at CDF and D0.

The book is set up in the following way: The first chapter will introduce you to the geography of the Tevatron. Chapter 2 will show you the various magnets the TeV uses along with some important definitions. Chapters 3, 4, 5, 6, and 7 are the technical support chapters. They outline the vital systems that keep the accelerator operational. You will learn the intricacies of the RF system in chapter 8. Instrumentation follows next. One of the most important chapters for an operator is chapter 10: Operating the Tevatron. The next two chapters delve into the theory of colliding beams and the implementation of a store. The last chapter we have left to deal with the future of the TeV. For those who want to investigate a bit further, there are some advanced topics in the appendices.

Darren Crawford, 2006

1. Geography

Outside of the Tunnel

Fig. 1.1 Aerial view of the Tevatron ring

An aerial picture of the Tevatron ring shows several buildings around its circumference. There is a pattern as to how the buildings are arranged. The Tevatron ring is divided into 6 equal sectors, labeled A through F. At the beginning of each sector resides the zero building followed by four evenly spaced service buildings. The zero buildings contain the electronics specific to the specialized functions of the zero locations in the tunnel. Service buildings are labeled 1 through 4, i.e. F1, F2, F3, and F4. Each service building contains diagnostic equipment electronics for beam position monitors and beam loss monitors, a quench protection monitor microprocessor and its associated heater firing units, vacuum electronics, etc. You may be thinking, “What is all that stuff just mentioned?” In upcoming chapters you will learn the function of each of the devices mentioned in this chapter. You may not know how they work right now but at least you know their location. Behind each service building, sitting on the berm, is a satellite refrigerator building that is used for maintaining the flow and temperature of the liquid helium. Scattered around the ring other buildings are found at B48, E17, F17, F23, and F27. Buildings B48 and F17 will be described in the miscellaneous section. The remaining buildings pertain to the Main Ring remnant and Pbar transfer lines.

Zero Buildings

The zero buildings contain power supplies and electronics specific for the equipment in the long straight sections. Every zero building contains a room with helium compressors that maintain pressure and flow around the ring. A 386 based microprocessor that controls a sector's refrigerators and compressors can also be found at zero buildings.

Starting east of the high rise is the A0 service building which houses a magnet drop, helium compressors, and a staging area to prepare magnets for installation. The magnet mover can often be found next to the MVA (major vehicle access) gate.

The B0 service building is on the opposite side of the berm from CDF, the Collider Detector at Fermilab. It contains the power supplies for the low β magnets, separators, and spark counters. (If you don't understand the terminology yet, don't worry you will as you read on.)

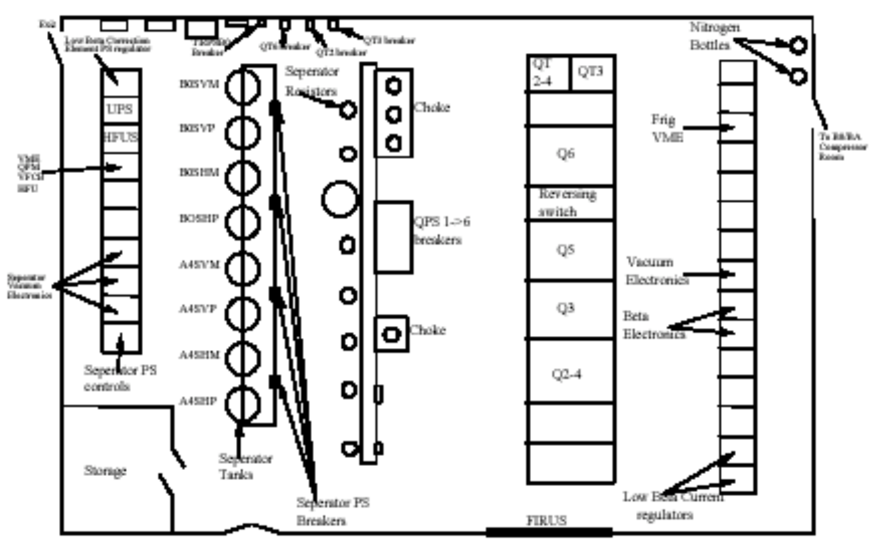

Fig. 1.2: B0 service building layout

Inside the C0 service building you will find the old proton abort line electronic diagnostics equipment. The new detector building is located on the opposite side of the berm from the service building. The SyncLite data acquisition system resides in the electronics room.

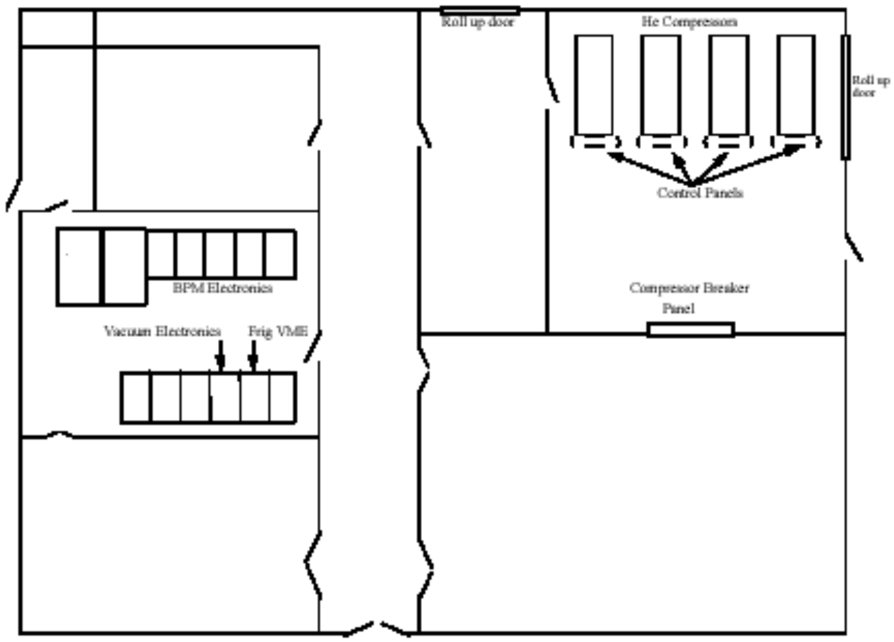

Fig. 1.3: C0 service building layout.

The D0 service building is located across the berm from the D0 detector building and next to the magnet storage building. This building is similar to the B0 service building in its contents and layout.

Picture of Service Building Fig. 1.3.1

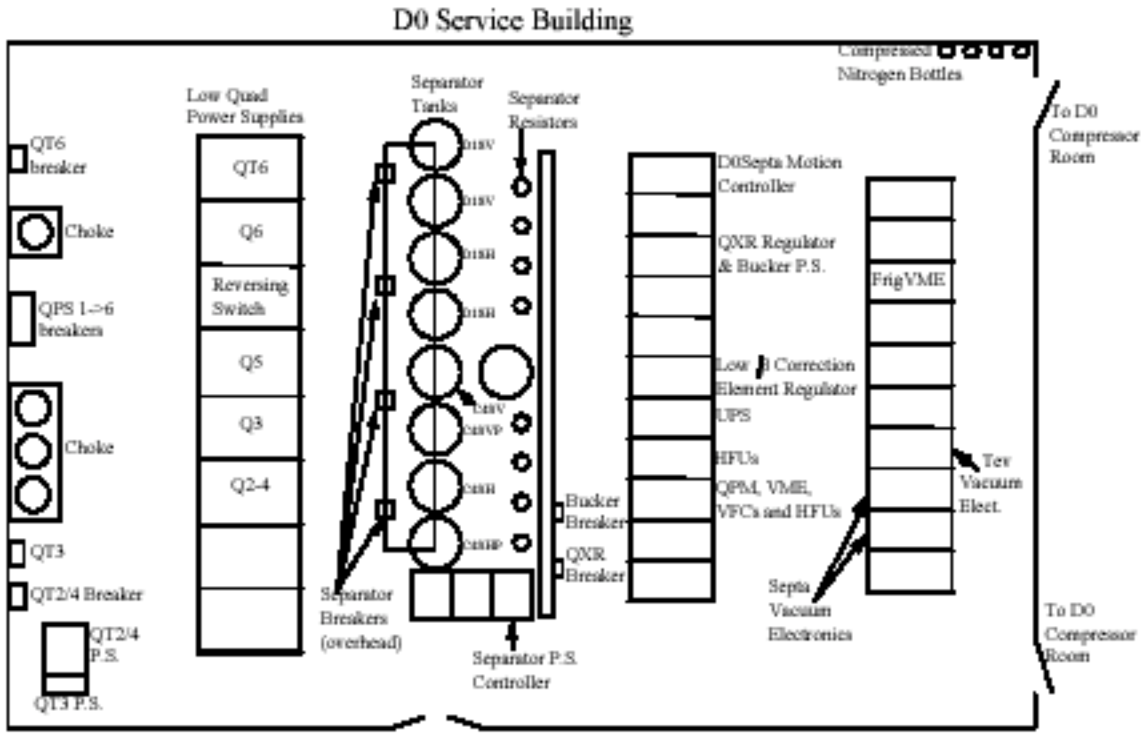

Fig. 1.4: D0 service building. In the picture, the left side door is to the electronics room and the middle door is to the compressor room. The diagram above shows the location of various electronics

The E0 service building contains the electronics that are used for scraping away the protons and antiprotons during a shot setup. Also, the Schottky detector electronics for measuring the horizontal and vertical tune of the protons and antiprotons can be found in the electronics room. Behind the building is an MVA. In this case MVA means minor vehicle access.

The final zero building is F0, which is also known as the TeV RF building. The building houses all of the equipment used to accelerate the beam from 150 GeV to 980 GeV. Inside you will find the anode power supply, 8 RF modulators, 8 power amplifiers, 8 transmission lines to the cavities, and 8 cavity water stations. All of the Tevatron damper electronics are located in F0. The LLRF VXI crate is located in the MI-60 control room, which is connected to F0 via a hallway. Outside the building, facing north, is the TeV anode power supply manual disconnect.

Service Buildings

Every sector has 4 service buildings labeled 1 through 4. Service buildings are mostly repetitive when it comes to their contents. The buildings contain a heat exchanger for the LCW, an air compressor that holds open vacuum valves and powers the Johnson Controllers, and an electronics room. Inside each electronics room you will find a Quench Protection Monitor with its associated Heater Firing Units and Voltage to Frequency Converters, Beam Position Monitor and Beam Loss Monitor electronics, a vacuum crate and ion pump power supplies, an Uninterruptible Power Supply system, and dipole correction element power supplies. Higher order correction element power supplies are found in some of the service buildings. A tunnel fan can be found outside of the 1 and 4 service buildings, facing the berm. The fan at an even numbered building forces air into the tunnel and at the odd number pulls air out. Just remember the phrase “odd man out.”

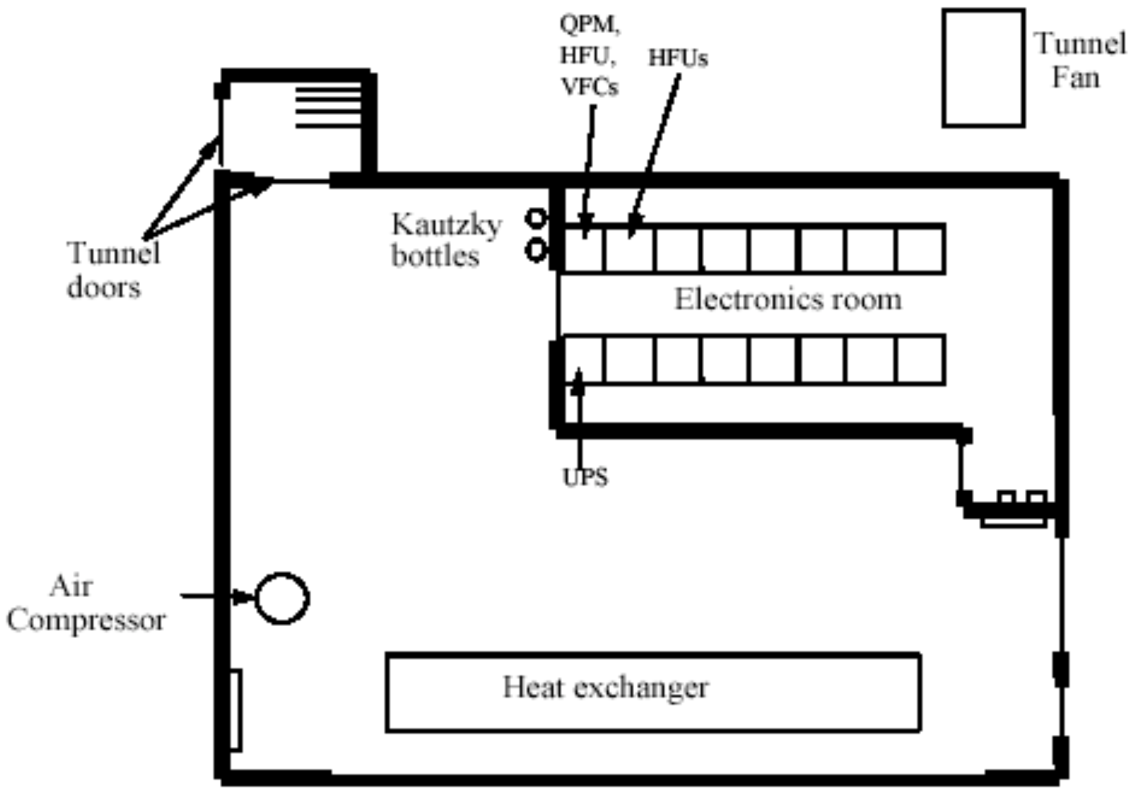

Fig. 1.5: Tevatron service buildings 1 and 4

The 2 and 3 numbered service buildings contain the power supplies that provide current to the TeV bus. Outside of these buildings you will find the manual disconnect, Vacuum Circuit Breaker cabinet, and transformer for the Tevatron power supply. On the opposite side of the building resides the 1/4-Ω dump resistor that is used to dissipate the power in the TeV bus. Because of the power supplies at these locations there will be extra equipment inside the building. The actual power supply resides outside of the electronics room and contains the commutating SCR (Silicon Controlled Rectifier) modules for rectifying the three-phase voltage signal from the transformer. Next to the power supply is a choke that is one part of the passive filter system. The remainder of the passive filter system is in the Filter/Dump cabinet. In the cabinet you will find the shunt and series SCR modules, a 1900 μf capacitor bank, the dump switch power module, and the knife switches used for placing the supply in or out of the circuit. Inside the electronics room you will also find equipment related to the power supply such as the CVT (Constant Voltage Transformer), TeV emergency off button, SPU (Standby Power Unit), QBS (Quench Bypass Switch) controller, safety coordinator, AC controller, and SCR firing unit.

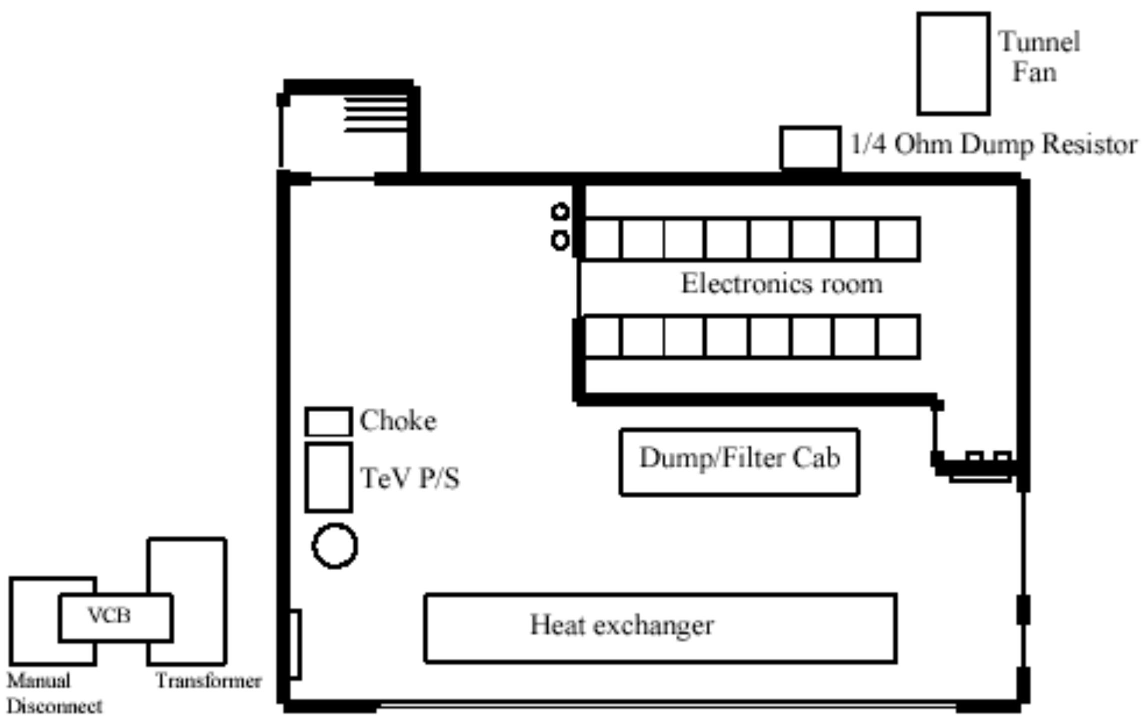

Fig. 1.6: Tevatron service buildings 2 and 3

Refrigerator Buildings

There are 24 refrigerator buildings ring wide, each located on top of the berm, directly behind a service building. A helium transfer line, also on the berm, runs the circumference of the ring connecting each refrigerator building with CHL. Next to the transfer line is the 3” discharge line, which is smaller in radius than the transfer line. Located over the transfer and discharge line is a blue cylindrical object called a heat exchanger.

Inside a refrigerator building you will find a wet and dry expander engine, valve box, bayonet can, suction header, device I/O crate, and cold compressor. Outside, next to the door is the crash button, which valves off the discharge line and thus bypasses high-pressure warm helium around the building. On the roof of the building reside the relief valves for the nitrogen and helium suction headers. The reliefs vent those gases during rapid expansion.

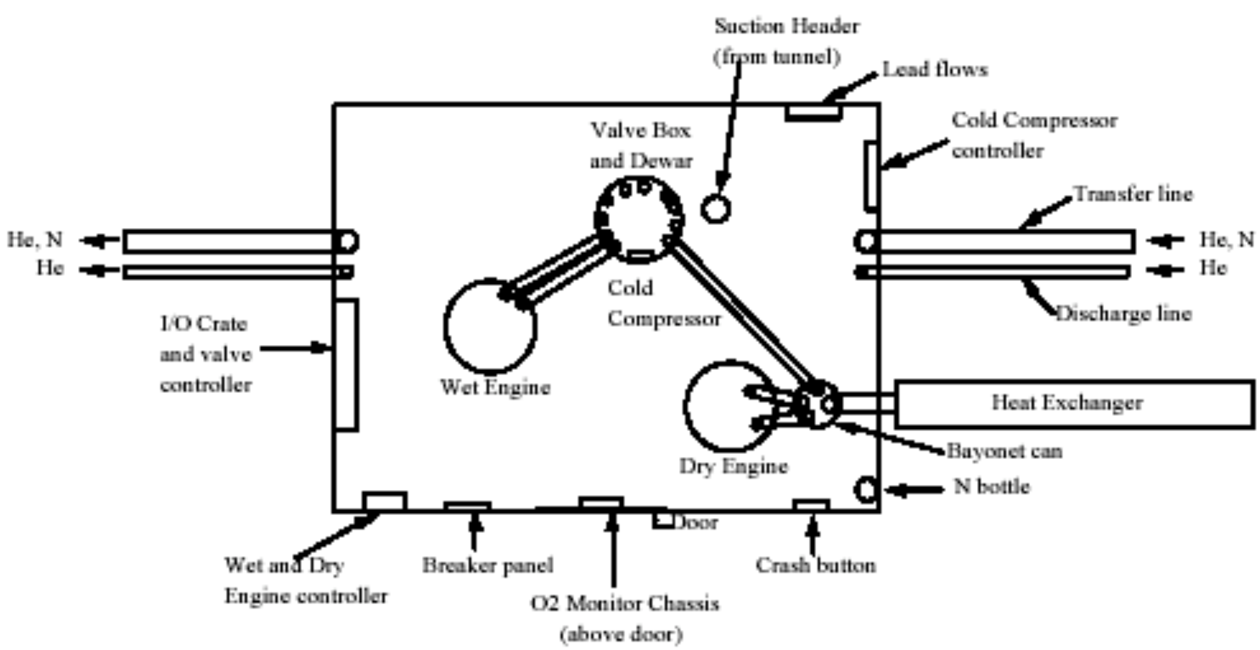

Fig. 1.7: Top picture is a typical refrigerator building. Bottom picture is a frig layout

In the Tunnel

The Tevatron ring is broken up into 6 symmetric sectors which are designated A, B, C, D, E, and F. Each sector starts with a section called the “zero” location and then the rest of the sector is broken up into four repetitive areas called “houses”, labeled 1 through 4. At the zero locations of each sector are long straight sections with specialized functions. A0 contains the proton/antiproton beam abort for collider. B0, C0, and D0 contain the collider detectors for top and bottom quark physics, supersymmetry studies, etc. Also located at C0 is the old proton abort line. The E0 straight section contains the Schottky detector electronics for determining the tune of the beam. F0 contains the 8 accelerating RF cavities as well as the P1 and A1 injection lines from Main Injector.

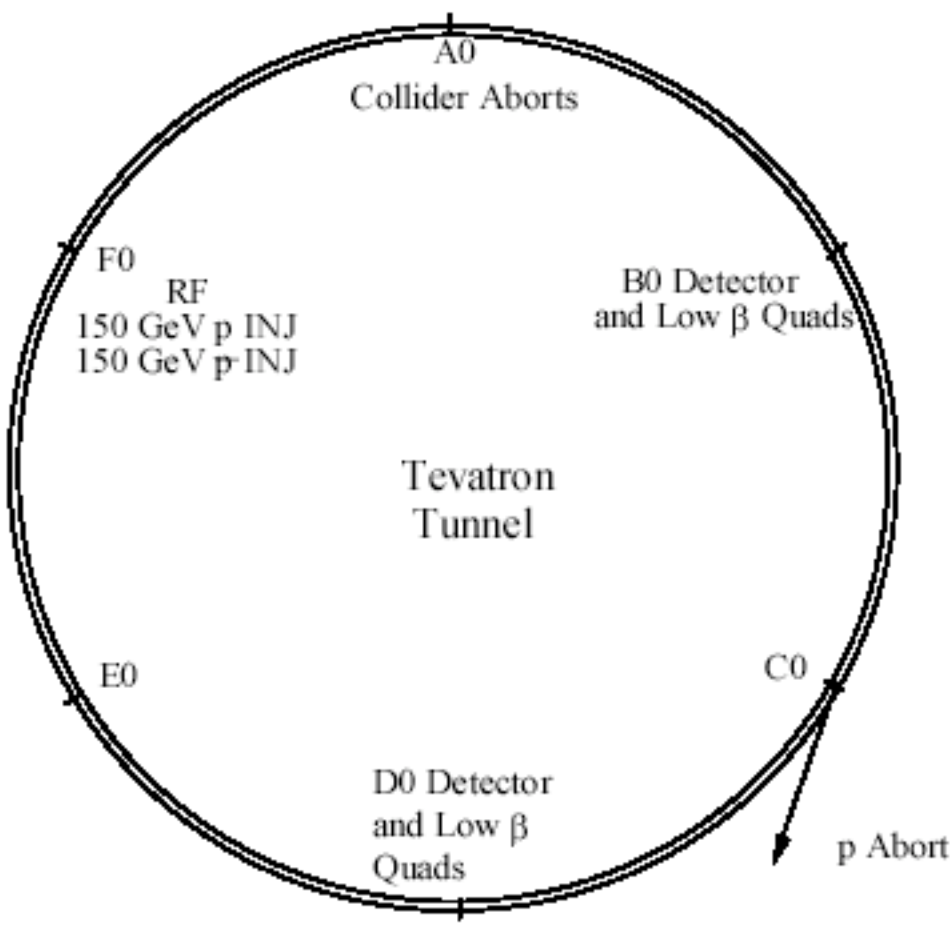

Fig. 1.8: Tevatron Tunnel

Houses consist of a number of repeating series of magnets called cells. Each cell has 10 magnets, 2 quadrupoles and 8 dipoles. You will learn more about magnets in chapter 2. Houses 1 and 2 have 4 1/2 cells each and houses 3 and 4 both contain 4 cells. This yields 17 cells for one sector. A cell starts with a quadrupole followed by 4 dipoles followed by another quadrupole and 4 more dipoles.

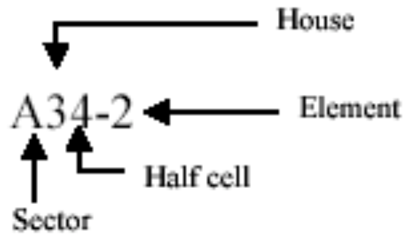

Example of A34-2, sector, house, half cell, and element

The numbering scheme for magnets in the tunnel is that a quadrupole is the first element of the half cell. If you were given a magnet number of A21-1 then you would know that you are in A sector, house 2, half cell 1, and element number 1. The diagram below should help clarify the numbering scheme. When you enter the tunnel from a service building you will be at the beginning of the 5 location for that house, so at the bottom of the A3 stairwell is the A35-1 magnet, etc.

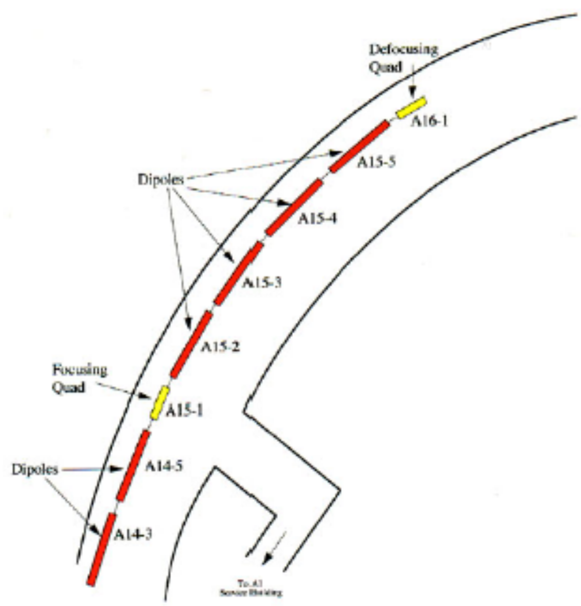

Fig. 1.9 Example of the magnet numbering scheme

To determine whether the quadrupole is focusing or defocusing depends upon what part of the sector you are in. In the 1 house all of the even numbered quads are defocusing and the focusing are odd numbered. The 2, 3, and 4 houses have the numbering scheme reversed. Reference the above diagram. The A15-1 quad is a focusing magnet. The box below shows the TeV ring numbering scheme.

Sector Numbering Scheme

Sector Numbering Scheme							
1 house		2 house		3 house		4 house	
Hor	Vert	Hor	Vert	Hor	Vert	Hor	Vert
11			21	32		42	
	12	22	23	34	33	44	43
13	14	24	25	36	35	46	45
15	16	26	27	38	37	48	47
17	18	28			39		49

Fig. 1.10: Magnet numbering scheme for each sector

2. Magnets and Lattice

Introduction to Superconductivity

Superconductivity is the phenomenon whereby a metal, alloy, ceramic, etc., when cooled to a low temperature, becomes a perfect conductor of electricity. As the temperature in a metal decreases the electrical resistance decreases until at a certain temperature, known as the critical temperature (T_c), the electrical resistance abruptly drops to zero. Critical temperatures are usually a few degrees above absolute zero. To describe superconductivity fully the concept of current density must be introduced. Current density, J, is the current per unit area carried through a conductor. For example, average house wiring is rated at 10^7 A/m^2, which is also the typical current density of conventional electromagnets with water-cooled copper windings.

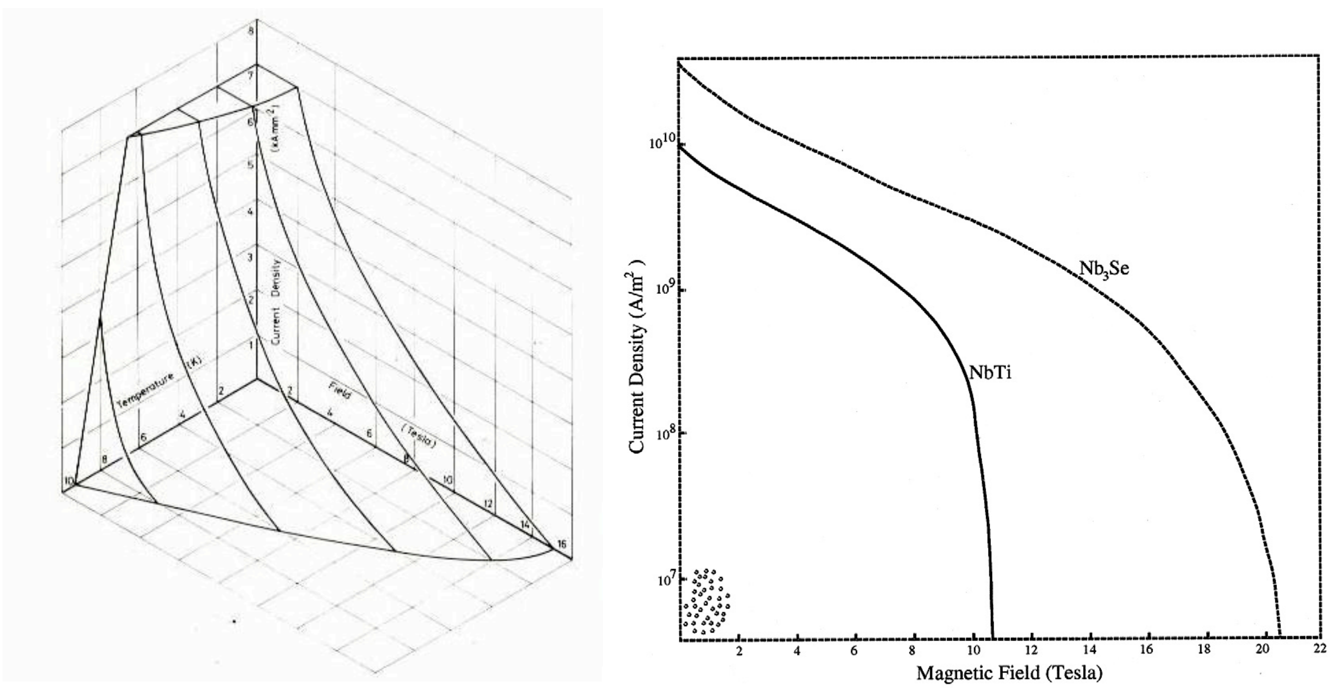

Fig. 2.1: The left plot shows the phase plot of NbTi. Below the surface NbTi is superconducting. The right plot is a comparison of 2 superconductors with conventional electromagnets.

Along with the critical temperature, the critical current density (J_c) and the critical magnetic field (B_c) help describe the characteristics of a superconductor. Similar to the phase plots of thermodynamics, the 3D plot of B, T, and J shows the critical surface of a material, where superconductivity exists below the surface and normal resistivity everywhere above it. The tests on the superconducting cable for the Tevatron magnets gave an average J_c of $1.8x10^9$ A/m^2 at 5 T and 4.2 K. At an operating point of 4.6 K, the Magnet Test Facility determined the average magnet current, I_{ave}, in the NbTi superconductor to be 4600 A, or 1.045 TeV.

To keep the NbTi bus at superconducting temperatures magnets must be placed in special helium vessels or cryostats. These are vacuum insulated containers that have an intermediate liquid-cooled nitrogen shield placed between room temperature and the low temperature region.

The NbTi cable has a keystone shape with dimensions 0.044 to 0.055-inch thickness X 0.308-inch width. Twenty-three strands of 0.0268-inch diameter wire are twisted flat to make up the cable. Each strand has 2050 filaments of NbTi alloy in a copper matrix that average 8.7 μm in diameter. The cable is insulated with 1 mil thick double wrapped Kapton and then spirally wrapped in a layer of fiberglass tape.

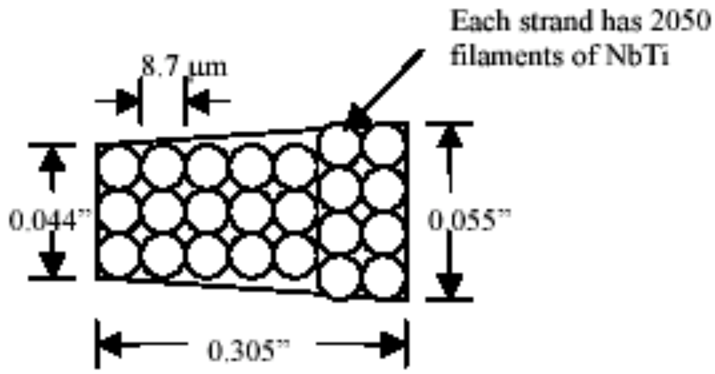

Fig. 2.2: A cross section of the NbTi bus.

Dipoles

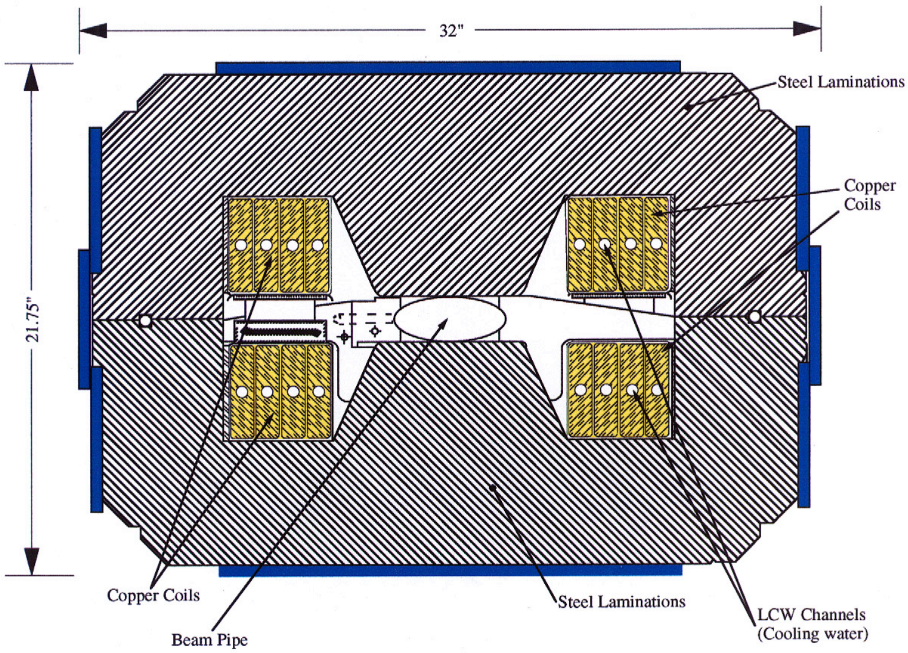

Fig. 2.3: Main Injector dipole cross section

We all know the cross sections of conventional dipoles for the Main Injector, Switchyard, etc. They have 2 sets of coils, top and bottom for a horizontal dipole or left and right for a vertical dipole.

The TeV magnets are much different in the fact that they are cryogenically cooled to become superconducting and have several vacuum chambers within. The center of the magnet contains the beam pipe, which, of course, is under vacuum of the order 10^{-9} torr. Surrounding the beam pipe and following the length of the magnet is single-phase helium at 4.6 K. The single-phase helium keeps the NbTi coil superconducting. Stainless steel collars clamp the magnet coil in place and keep it from distorting during ramping, which can be 4400 amps or more. Around the collared assembly is a two-phase (liquid and gas) helium jacket, which returns the cryogens along the length of the magnet in the opposite direction of the single-phase. This counterflow allows for heat exchanging to occur at the surface of the single-phase tube. Outside of the two-phase jacket are two concentric insulating vacuum spaces. Next is a liquid nitrogen jacket and finally an outer insulating vacuum space, which intercepts heat flow from room temperature. Superinsulation (aluminized Mylar) surrounds the outer insulating vacuum tube as an extra heat radiation shield.

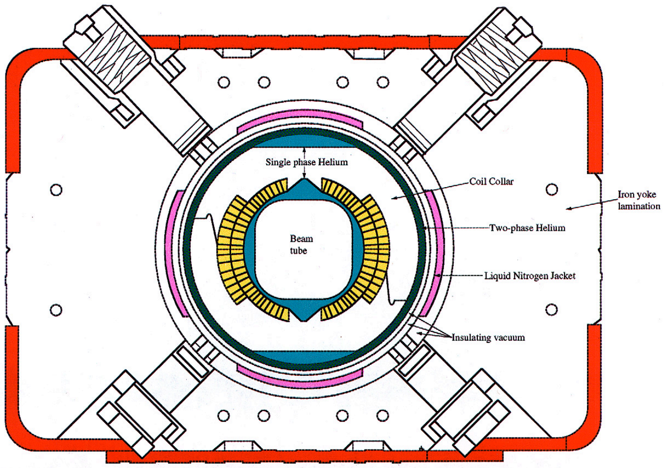

Fig. 2.4: Cross Section of a Tevatron Dipole

The entire magnet assembly is vacuum tight. It is held in a laminated iron yolk, which contributes roughly 18% to the total magnetic field. The assembly is precisely adjusted to within 1 mil of center with G10 suspension blocks and preloaded suspension cartridges that allow for thermal contraction and expansion. The 21-foot dipoles have nine sets of suspension points.

As you may have noticed, the coils are not configured like those of the Main Injector dipoles. The conventional dipole coil structure would not produce a uniform dipole field. Instead, the B field would have components in both transverse directions. To produce a perfect dipole field the windings have to take on the shape of two intersecting ellipses.

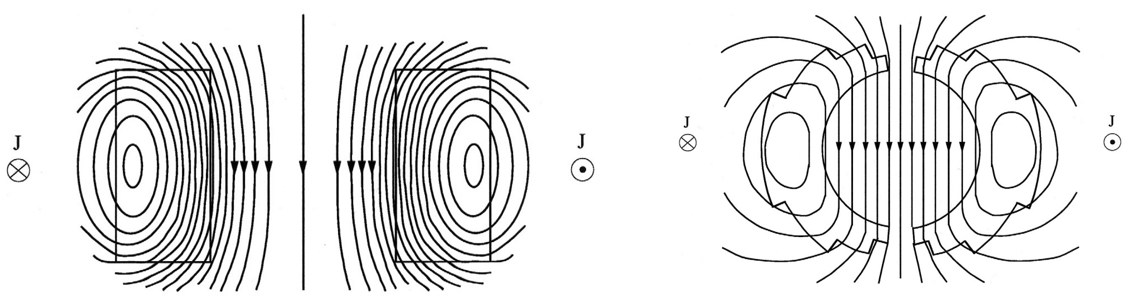

Fig. 2.5: The left diagram shows the magnetic field for a normal conductor. The right diagram shows the geometry for the NbTi superconductor that creates the dipole field

The dipole magnets have a magnetic length of 6.12 m and a total length of 6.4 m. Each magnet produces a deflection angle of 8.1 mrad.

Quadrupoles

Quadrupoles, of course, focus or defocus the beam with respect to the horizontal plane. The geometry for the superconducting quad coils are that of 2 sets of intersecting ellipses rotated 90 degrees apart.

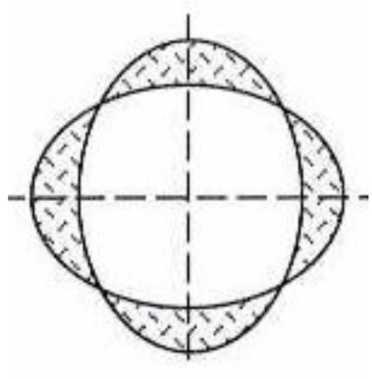

Fig. 2.6 Geometry of the superconducting bus that creates the quadrupole field

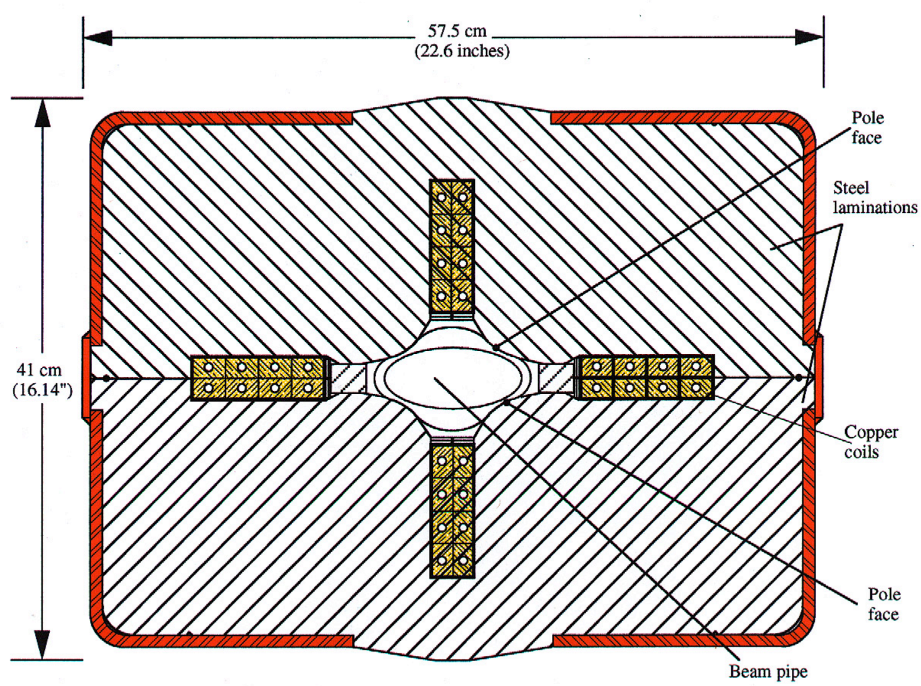

Fig. 2.7: Main Injector Quadrupole Magnet Cross Section

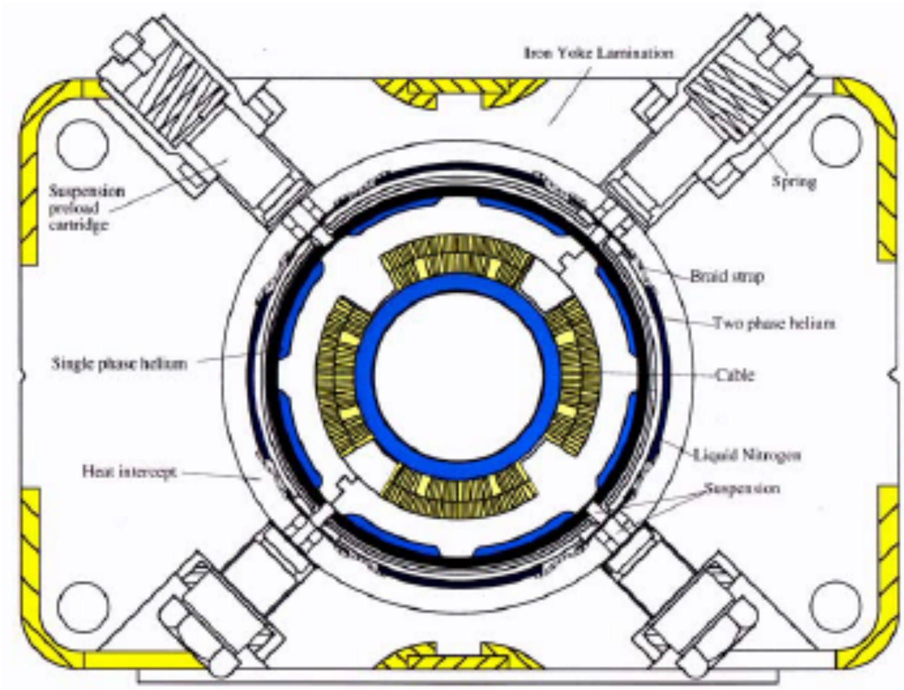

Fig. 2.8: Tevatron Quadrupole Magnet Cross Section

Cell and Lattice

As stated in chapter 1, a house consists of a repetitive series of magnets called cells. Houses 1 and 2 have 4 ½ of these repetitive cells and houses 3 and 4 have 4 cells each that yields 17 cells for each sector. Since each sector bends the beam 60 degrees, a cell deflects the beam by 3.5 degrees. Each cell has 10 magnets, 8 dipoles and 2 quads. A cell begins with a quadrupole followed by a mini-straight section, where the correction coil spool piece resides. Next, there are 4 dipoles followed by another quadrupole, correction coil spool piece, and 4 more dipoles. Reference the figure below. This repetitive cycle is broken in 3 places, at 0, 17, and 48 locations. The cell described above is one of the repetitive segments of the FODO lattice, which is a focusing quad, 4 dipoles, a defocusing quad, and 4 dipoles.

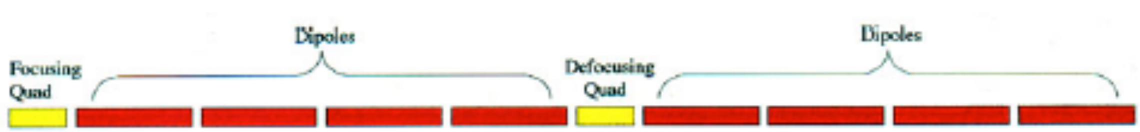

Fig. 2.9: The FODO lattice arrangement for the TeV

Each magnet is a four-pole device, with two leads on each end. The inductance of each magnet is concentrated either on the upper bus (TC type dipole, F quadrupole) or on the lower bus (TB type dipole, D quadrupole). The inductance of a typical “half cell”, that is the inductance of either the upper or lower bus through the cell, is about 0.18 H. This yields an inductance for the entire ring of 36 H. The inductive stored energy at 1 TeV (4440 A) is

$$\frac{1}{2}LI^2 = \frac{1}{2}(36H)(4440A)^2 = 3.50x10^8\,Joules.$$

That’s right, 350 MJ of energy.

Correction Elements

The previous sections have shown how different magnet design is for the TeV compared to the Main Injector. The same holds true for the correction and adjustment magnets. Two factors were considered in their design and use. First, error fields were no longer just a ramping phenomenon. Inevitable deviations in the superconductor from the ideal configuration produce significant field distortions that are independent of ramping. Therefore, certain corrections are required at all field levels. Second, the beam pipe is relatively inaccessible. It is buried throughout most of the accelerator inside an essentially continuous cryostat. Magnets cannot be added or shifted with ease.

Correction magnets are those required to correct field imperfections and alignment errors of the main quadrupoles and dipoles. Adjustment magnets are those required to tune the optics of the accelerator depending on the operating conditions. Often the same magnet performs both functions. Thus, dipole steering magnets are necessary to compensate for alignment errors so that the beam is in the center of the aperture and also necessary to bump the beam away from a restriction like an injection magnet. The correction and adjustment magnets are superconducting coils located within the main quadrupole cryostats. They are located immediately downstream of the quadrupole in a spool piece, a stainless steal tube that also houses the main magnet bus, cryogenic flows, and the vacuum connections for the insulation and beam tube vacuum. Most spool pieces contain 2 packages of 3 concentrically wound correctors with 6 pairs of leads coming from cryogenic temperature to room temperature.

Fig. 2.10: Correction element magnet geometry

There are 26 various configurations of the spool types. The most common will be outlined. The following table will show all of the various packages and spool types. The first corrector package is DSQ I, which contains a horizontal dipole, normal sextupole, and normal quadrupole. The DSQ I is thus used for horizontal steering, chromaticity adjustment, and tune adjustment. The DSQ II package differs in that the dipole is vertical instead of horizontal. The second most common corrector package is the OSQ of which there are 3 varieties, I, II, and III. OSQ I contains an octupole, skew sextupole, and skew quadrupole, whereas OSQ II contains all of the magnets in the normal configuration. OSQ III is different in that only the quadrupole is skewed.

The table on the next page shows the outline of various spool types. The number of poles, P, for each of the 3 coils is given. Thus, a 2P is a dipole, a 4P is a quad, and so on. The S prefix indicates that the coil is skewed. For example, a S-2P is a horizontal dipole rotated 90 degrees, which makes it a vertical dipole.

Spool Type	Upstream Inner Middle Outer	Downstream Inner Middle Outer	Safety Leads	Notes
A	DSQ I 2P 6P 4P	None	No	43" spool length, only one corrector package
B	DSQ II S-2P 6P 4P	None	Yes	43" spool length, only one corrector package
BA	DSQ II S-2P 6P 4P	Bypass	Yes	72" B spool
C	DSQ I 2P 6P 4P	OSQ I 8P S-6P S-4P	No	72" spool
CA	DSQ I 2P 6P 4P	OSQ I 8P S-6P S-4P	Yes	72" spool
CB	DSQ I 2P 6P 4P	None	No	72" spool
D	DSQ II S-2P 6P 4P	OSQ I 8P S-6P S-4P	Yes	72" spool
DA	DSQ II S-2P 6P 4P	None	Yes	72" spool
DR	DSQ II S-2P 6P 4P	OSQ I 8P S-6P S-4P	Yes	72" D spool with recooler for lower temperature
E	DSQ I 2P 6P 4P	OSQ II 8P 6P 4P	No	72" spool
F	DSQ II S-2P 6P 4P	OSQ II 8P 6P 4P	Yes	72" spool
FR	DSQ II S-2P 6P 4P	OSQ II 8P 6P 4P	Yes	72" F spool with recooler for lower temperature
G	DSQ I 2P 6P 4P	OSQ III 8P 6P S-4P	No	72" spool
H	DDQ 2P S-2P S-4P	None	No	5000 A Power Lead 49.91" spool length
HA	DDQ 2P S-2P S-4P	None	No	5000 A Power Lead 49.91" spool length
HH	DDQ 2P S-2P S-4P	None	No	High Temp 5000 A power lead 49.91" spool length
I	30" DD 2P S-2P	Bypass	No	84" spool length
IA	30" DD 2P S-2P	Bypass	No	125" spool length
J	T6 Quad	DSQ I 2P 6P 4P	Yes	5000 A Power Lead Low β quad trim, 72" spool
K	T6 Quad	DSQ II S-2P 6P 4P	Yes	5000 A Power Lead Low β quad trim, 72" spool
L	30" DD 2P S-2P	None	Yes	5000 A Power Lead 56.149" spool length
M	S5 Quad	DSQ I 2P 6P 4P	Yes	2000 A Power Lead Bartelson quad, 72" spool
N	S5 Quad	DSQ II S-2P 6P 4P	Yes	2000 A Power Lead Bartelson quad, 72" spool
P	DDQ 2P S-2P S-4P	None	Yes	5000 A Power Lead 56.149" spool length
R	None Drift Space	None	No	5000 A Power Lead 36" spool length
S	DDQ 2P S-2P S-4P	None	No	49.91" spool length

Fig. 2.11: Spool Types

The Table below is the spool map for the Tevatron. It indicates the type of spool located downstream of each main quadrupole.

Quad		Sector					
D	F	A	B	C	D	E	F
	11	H	P(B10), R	H	S	H	H
12		D	L	F	L	D	D
	13	C	J	C	J	C	C
14		D	N	F	N	B	D
	15	A	A	A	A	A	A
16		D	N	F	N	D	D
	17	C	M	C	M	C	C
18		DR	DR	F	FR	D	F
	19	C	E	E	C	E	E
21		B	B	B	B	B	B
	22	C	C	G	C	C	G
23		D	F	F	D	F	D
	24	C	E	E	C	E	E
25		B	B	B	B	B	B
	26	C	C	G	C	C	G
27		D	FR	D	DR	FR	F
	28	C	E	E	C	E	E
29		B	B	B	B	B	B
	32	C	C	G	C	C	G
33		D	F	D	D	F	D
	34	C	E	E	C	E	E
35		B	B	B	B	B	B
	36	C	C	G	C	C	G
37		D	F	F	D	F	F
	38	C	E	E	C	C	E
39		B	B	B	B	B	B
	42	C	C	C	C	C	C
43		N	D	N	D	D	D
	44	M	C	M	C	C	C
45		B	B	B	B	B	B
	46	M	C	M	C	C	C
47		K	DR	K	B	FR	FR
	48	L, A, H	A	L, A	A	A	A
49		P	H	S	H	H	H

Fig. 2.12: Spool type and Location

3. Power Supplies

More power! -Tim Allen

So how does the power supply work? The acronym SCR (silicon controlled rectifier) is heard quite often in the MCR and it is a vital component to a power supply. The SCRs are semiconductors that behave like diodes except that they can be gated on. Diodes, when forward biased, will conduct until the voltage drop across it is zero. For an SCR, it can conduct when it is forward biased but only after receiving a gate pulse. This allows control over how much of a sine wave will pass through the SCR. Refer to the diagram below.

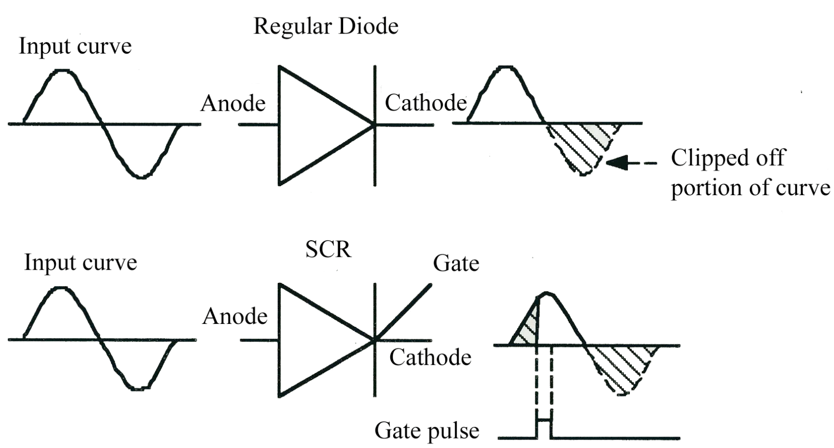

Fig. 3.1: Diagram showing the difference between a diode and SCR

In a Tevatron power supply, the input from the utility yard is in the form of 4 conductors each coming from the X and Y transformer secondary side. These inputs go to 2 banks of SCR modules that are coupled together. Each module contains 4 pairs of SCRs; one pair connected to each leg of the Wye transformer and one bypass SCR connected to the neutral terminal of the transformer. This configuration results in a 12 phase full wave rectification.

A 12 phase system is employed so that the resulting ripple at 720 Hz can be easily filtered out. The filtering is accomplished with a 0.8 mH choke (inductor) located next to the supply. The harmonic filtering network is located in the dump/filter cabinet.

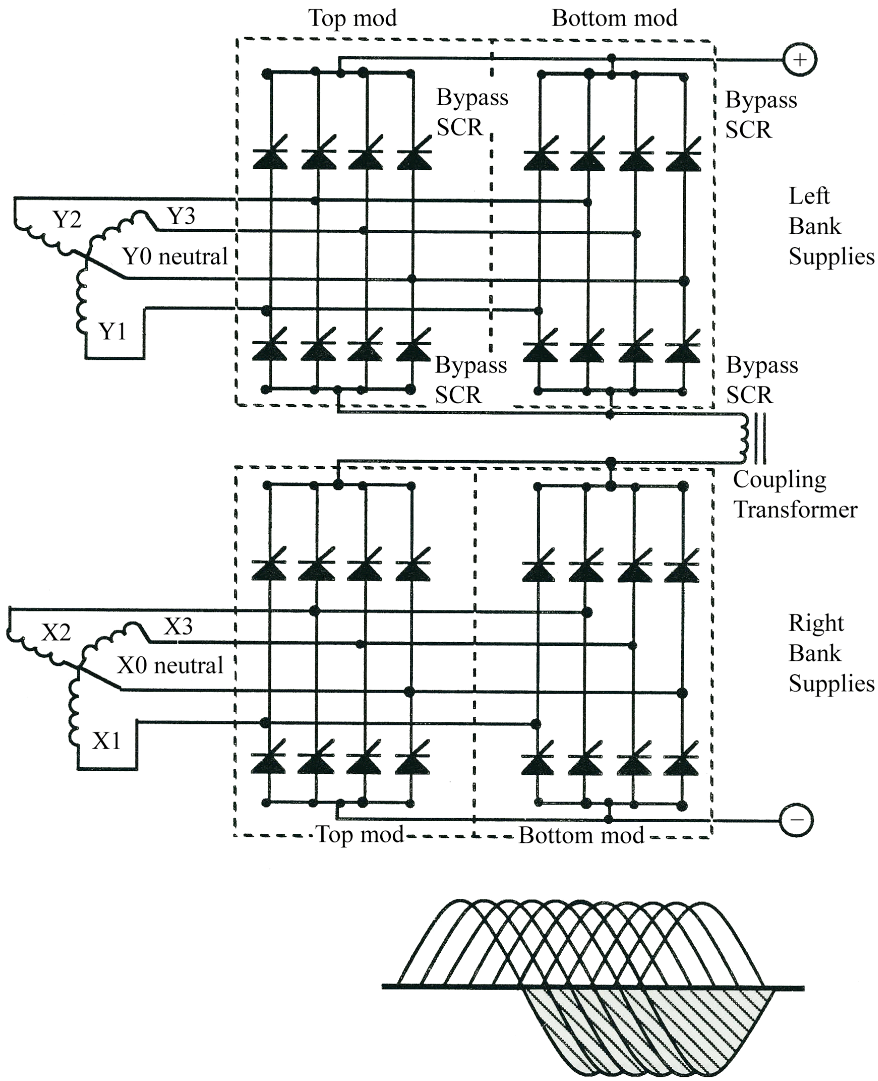

Fig. 3.2: Power Supply Schematic

The power supplies are turned on (or commutated) by providing a firing voltage to the gates of the SCRs in the SCR modules. Each SCR requires 14 V to conduct and that is supplied by the CVT (constant voltage transformer), located in the electronics rooms behind the racks. The CVT voltage itself is gated by the SCR Unit, which gets inputs from the PT (potential transformer) in the yard and TECAR. TECAR determines the firing angle of the individual SCRs and sends the command to commute, which, in turn, determines the output of the power supply.

The PT is a 13.8 kV to 120 V transformer located inside the 13.8 kV cubicle. Its purpose is to provide a phase reference for the SCR Unit by knowing the incoming 13.8 kV phase.

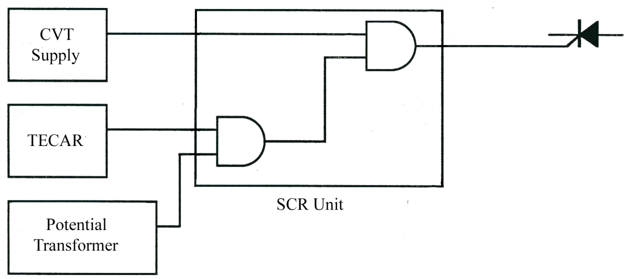

Fig. 3.3: SCR Unit Inputs

The 774 dipoles, 216 quadrupoles, and 12 power supplies form a single series circuit with “upper” and “lower” busses connected at a “fold” in the B0 straight section. All of the power supplies are capable of ramping to 4500 A at 1 kV. Eleven power supplies operate in a voltage regulation mode to ramp the current up and down. The A2 power supply acts as the current regulator and the A3 power supply is its back up.

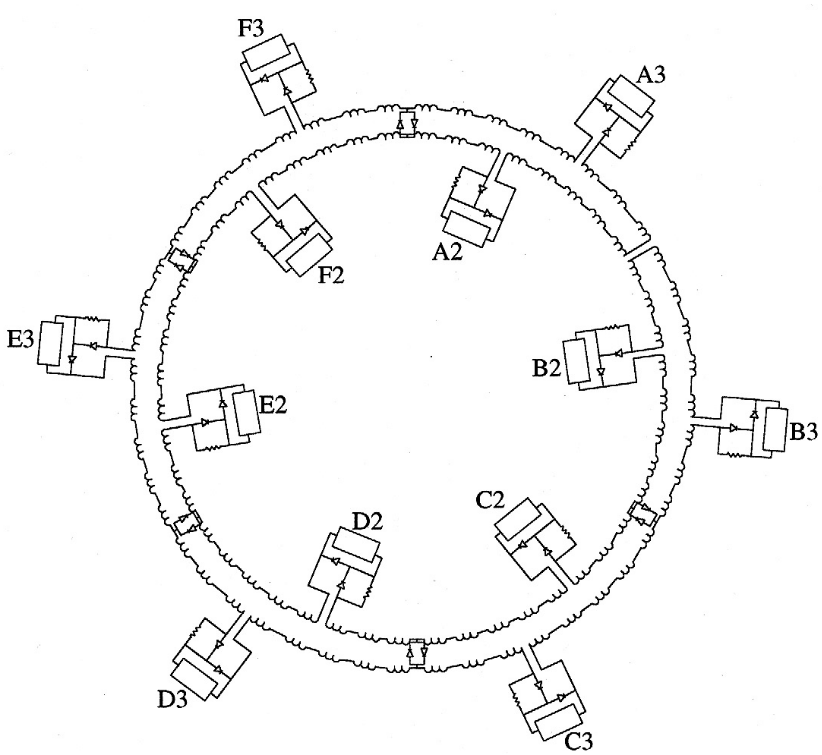

Fig. 3.4: Tevatron power supply distribution around the ring

Since the resistance of the entire Tevatron circuit is small, the current regulator can provide the required voltage during flattop. Only this supply must be capable of conducting continuous flattop current. Collider mode requires 6 power supplies with the remainder held in reserve.

The TeV power supplies are distributed around the ring at each of the “2” and “3” houses. Each power supply circuit consists of a DC breaker, series SCR, dump resistor, shunt SCR, passive filter, and (of course) the power supply. Reference the diagram below to see the layout of the circuit.

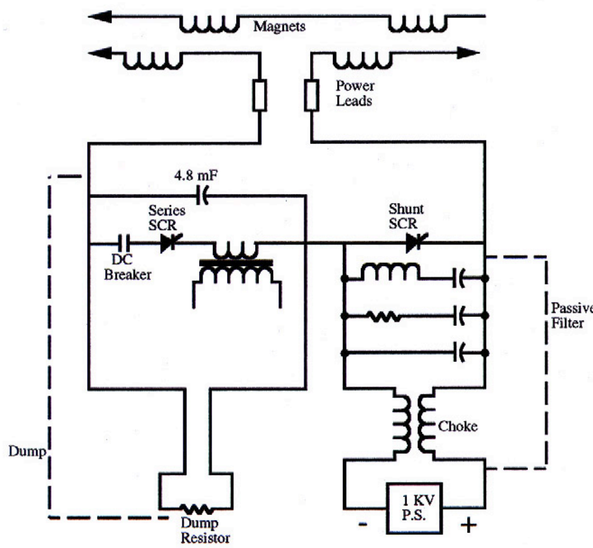

Fig. 3.5: Power Supply Circuit

There are four modes of operation for the TeV PS circuit.

1.) The PS is providing current to the bus
2.) The dump resistor is dissipating current from the bus
3.) The PS is not required due to the TeV being at flattop
4.) Current is being dissipated through the TeV bus

When the power supply is providing current to the TeV bus the DC breaker is closed, the series SCR is conducting, and the shunt SCR is open. This configuration allows the current that comes through the power leads, up from the tunnel, to bypass the 0.25 W dump resistor, enter the passive filter network and PS, and then back to the magnets. You may ask “But what about the 4.8 mF capacitor in parallel with the series SCR and DC breaker?” Not to worry. When the series SCR is conducting the capacitor is seen as an open circuit.

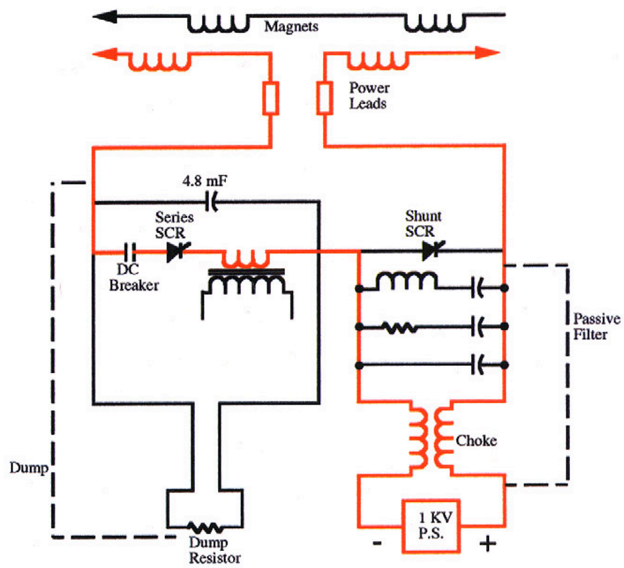

Fig. 3.6: Power supply circuit showing current path during ramp

In the case that we quench or have to dump the current in the TeV, the dump resistor will need to be brought into the circuit, which is done by commutating the series SCR, so that it does not conduct, and gating the shunt SCR, so that it does conduct. The DC breaker is opened as a backup to firing the series SCR. A 4.8 mF capacitor is added so that there is an RC circuit, which allows for a soft dump of the current. In order to open the series SCR a 1900 mF capacitor is discharged through the commutating transformer and the resulting current flow through the SCR goes to zero, which turns off the SCR.

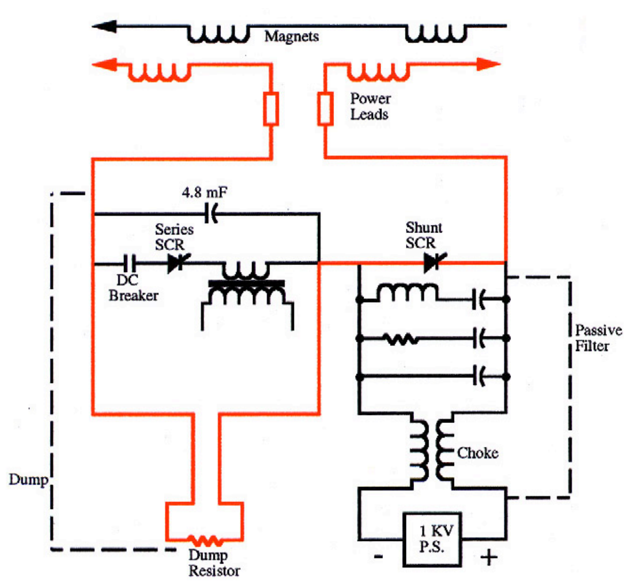

Fig. 3.7: Power supply in dump mode

Another mode the SCRs can be in is where both the series and shunt SCRs are on. This occurs when the power supply is not required due to the Tevatron being at a flattop. At this point only the regulating power supply, located at A2, provides current to the bus since the voltage required is small. At flattop or during any constant energy segment of the ramp the regulator is the only power supply on.

When the bypass is fired all of the series and shunt SCRs around the ring are gated on, making them conduct. The current in the Tevatron then bleeds away due to the resistive warm bus segments around the ring, usually in the zero locations and PS locations.

Fig. 3.8: The top picture shows the disconnect, VCB, and transformer outside of the service building. The picture below left is the power supply with the choke next to it. The picture below right is that of the dump/filter cabinet.

TECAR

A brief description of TECAR (Tevatron Excitation, Controller And Regulator) was given earlier in this chapter. Now a full explanation of this processor's responsibility will be given. This microprocessor controls the ramp waveforms of the main dipoles and quadrupoles and supervises the de-excitation of the power supplies during a quench. The VME crate that houses the TECAR associated electronics is found in the A2 service building electronics room.

During ramping or excitation, TECAR is responsible for the generation of the current waveform and developing the voltage waveforms for all the power supplies. The process for each supply playing out a waveform is achieved by TECAR providing the appropriate firing angles to the SCRs. Once flattop is achieved TECAR phases out the ramping supplies, except for the regulator, by gating the shunt SCRs.

In the event of a quench or some other fault, the microprocessor takes control of shutting down the power supplies and gathers diagnostic data. When a quench protection monitor (QPM) microprocessor communicates to TECAR that a quench has occurred, TECAR, via the other QPMs, performs a ring wide shutdown of the power supplies, activates the dump switches, and triggers the quench bypass switch (QBS) controllers.

A fiber optic link that connects TECAR with all the QPMs and a redundant hardware loop allows for continual cross checking of the current in all the power supplies and its derivatives, along with resistance measurements. TECAR broadcasts two 100 kHz pulse trains, one positive and the other negative, on the hardware loop and can inhibit either one or both of the pulses. The effect of inhibiting the pulses will cause the power supplies to turn off, and the dump switches to open. If any QPM loses communication with TECAR, the QPM will clamp both pulses.

Quench Protection System

Before discussing the quench protection system we must understand the dynamics of a quench and why it occurs. For any given superconductor there is a mathematical surface defined by temperature, magnetic field, and current density that mark a boundary between the superconducting and normal regime (Refer to figure 2.1). Tevatron magnets operate close to this boundary because cost constraints minimized the amount of superconductor to be used. The superconductor in the Tevatron was specified to have a Jc of 1800 A/mm at 5 T and 4.2K for the individual strands.

The transition from the superconducting state to the normal state can occur from the temperature in the conductor going beyond the critical temperature (T>Tc), for instance, due to beam loss. Another cause could be that the current or field is increased beyond the critical values. This is the beginning of a quench.

Once a normal resistive zone has formed it will propagate outward at a velocity that is dependent upon the current and field. A quench in a Tevatron magnet will expand with a velocity of

$$v = 0.36\ I^2\ (1 + 0.077\ B^2)$$

where the velocity is in units of meters per second, I is the current in kA, and B is the magnetic field in Tesla. The above formula is applicable from a current of 1 kA to 4 kA.

About 10% of the superconducting cable is open area between the strands, which is filled with liquid helium. When a quench propagates through a magnet the helium in the normal zone vaporizes and in turn the expanding gas displaces the liquid. The current transfers from the NbTi to the copper within the strands. If the cable reaches 460 K the numerous solder splices and silver-tin coating on the strands will begin to melt. At 4400 amps there is less than ½ second for removing the magnet current and preventing permanent damage.

The quench protection system protects the main bend and quadrupole magnets from the 350 MJ of stored energy in the ring should any portion of the superconductor become resistive. The major component of the quench protection system is the Quench Protection Monitor (QPM), which is responsible for monitoring the voltages across the half-cells of a given house. It determines whether a quench is occurring and, if so, fires the appropriate Heater Firing Units (HFUs) so that the magnet string becomes fully and uniformly resistive. Also, the QPM must communicate to TECAR that a quench has occurred, which in turn, through the other QPMs around the ring, turns off the power supplies, activates the dump switches, and triggers the QBS controllers. The refrigerator microprocessor receives the appropriate information from the QPM as to which cells have quenched so that the cool-down can begin immediately.

The half-cell voltage monitoring is accomplished by using Voltage-to-Frequency Converters (VFCs). The cables from the converters to the magnets in the tunnel are ~200 m long and any given cable serves as the positive input for one VFC and the negative input for the adjacent VFC. The VFC output is read by the QPM scalers, which are auto-zeroed periodically from the MCR to remove any voltage drift over time.

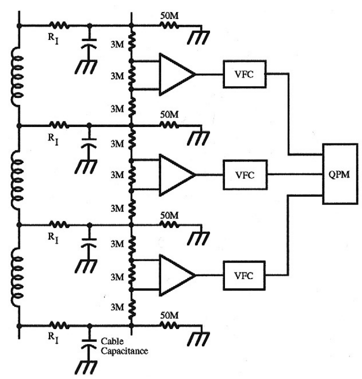

Fig. 3.9: VFC connections to each half cell

The detection of a quench is based upon the difference between the measured half-cell voltage and the expected inductive voltage. The cell voltage can be written as

$$V_{cell} = -L\ dI/dt + V_{resistive}$$

where L is the inductance of the half-cell and $V_{resistive}$ is, hopefully, zero during normal operation. With the resistive voltage at zero the cell voltage is purely inductive. The above equation can be rewritten as

$$V_{resistive} = V_{cell} + L\ dI/dt$$

The polarity of a VFC measurement is negative so that when Vcell is fully inductive Vresistive will be zero. During a quench the cell voltage becomes more negative and thus the resistive voltage becomes a negative value. The threshold limit is -0.5 V before the QPM detects a quench. Also a possibility is a malfunction, where a VFC card fails and Vresistive rises above 3.0 V. The QPM proceeds as if this were an actual quench. Since Vresistive is of opposite polarity to the actual quench, this event is called an antiquench.

The VFCs are located under the QPM in the Tevatron service building electronics rooms. There are three types of VFCs: 1) 10 V, 2) 200 mV, and 3) 200 V. The 200 Volt VFCs are used for measuring the resistive voltage across the half cell. The 200 mV VFCs are connected to the power leads and the 10 V VFCs are used for voltage to ground measurement on the bus.

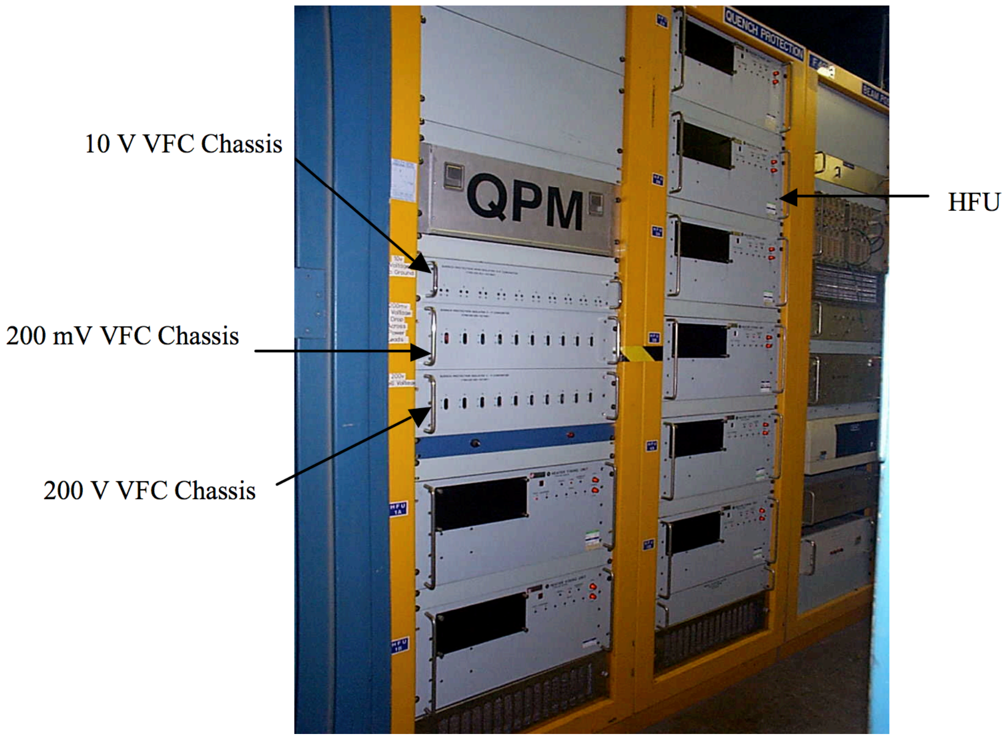

Fig. 3.10: QPM and HFU racks located in an electronics room

The HFUs contain capacitors that remain charged until a quench is detected within their cell. When fired by a command from the QPM, each HFU deposits 650 J of energy into the heater strip located between the superconductor and coil collar in the magnet. Testing the discharge time constants of the capacitors is a procedure performed in the MCR, which ensures that sufficient energy is being deposited into the magnets. When the HFU energy is dumped into the dipole heaters it causes a uniform voltage distribution throughout the magnet.

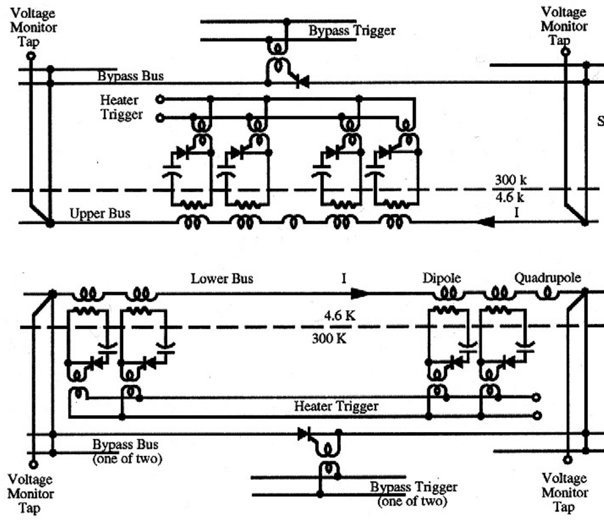

Fig. 3.11: Cell schematic showing an HFU circuit, voltage tap connections for VFCs, and a quench bypass system

The quench bypass circuits consist of two independently controlled thyristors (A and B) and a self-firing circuit that will turn on the QBS once the voltage across the cell reaches 200 V. They are connected to the magnets via the safety leads. The QBSs are semiconductors and sensitive to radiation so holes have been bored into the Tevatron tunnel and the QBS placed within for shielding.

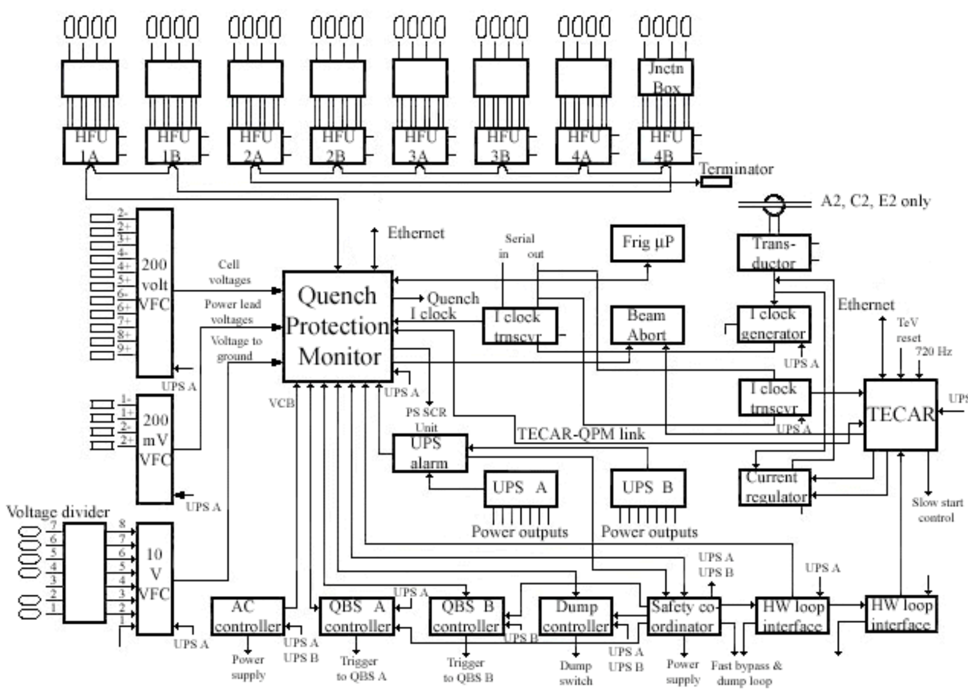

Fig. 3.12: Quench protection system schematic

Master Substation and TeV Power

There are 6 feeders that provide the necessary power to the Tevatron buildings and accelerator systems. Feeder 23 provides the pulsed power for the power supplies at the 2 and 3 service buildings and the low b magnets. This, of course, is racked out for any access into Tevatron (Transfer Hall, A-E, or F-sector). When Feeder 23 is unavailable, whether due to maintenance or otherwise, the pulsed power can be back fed from the Kautz Road Substation. Because of the load that the KRSS feeders are already providing for Main Injector and the transfer lines to the Antiproton Source this mode greatly reduces the maximum possible repetition rate of beam events in those machines.

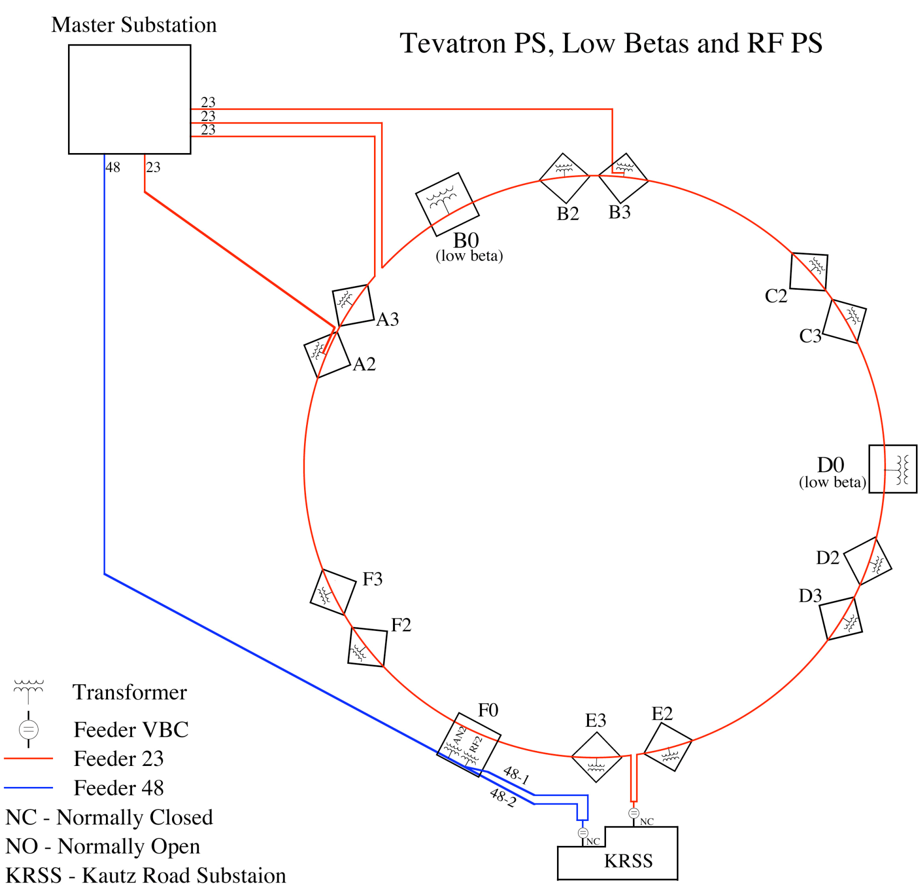

Fig. 3.13: Pulsed power distribution for the 2 and 3 service buildings, low b power supplies, and RF systems

Feeder 45 provides conventional power to all service buildings including zero buildings, and E4R. Conventional refers to the power for lighting, relay racks, and air conditioning. Critical systems, like the QPMs, CVTs, fire detection, etc., are each on an uninterruptable power supply or battery backup. At F0, this feeder provides all of the power to the RF systems except for the anode power supply and the water heaters. Feeder 45 also has the capability of being back fed from KRSS. The anode power supply and the water heaters for the TeV RF are on Feeder 48. The normal source of power for this feeder is MSS and the back feed is KRSS.

CDF collision hall, D0 collision hall, and the C4 pump house power are fed from Feeder 49. MSS provides the power to the feeder and it can be back fed from Feeder 45.

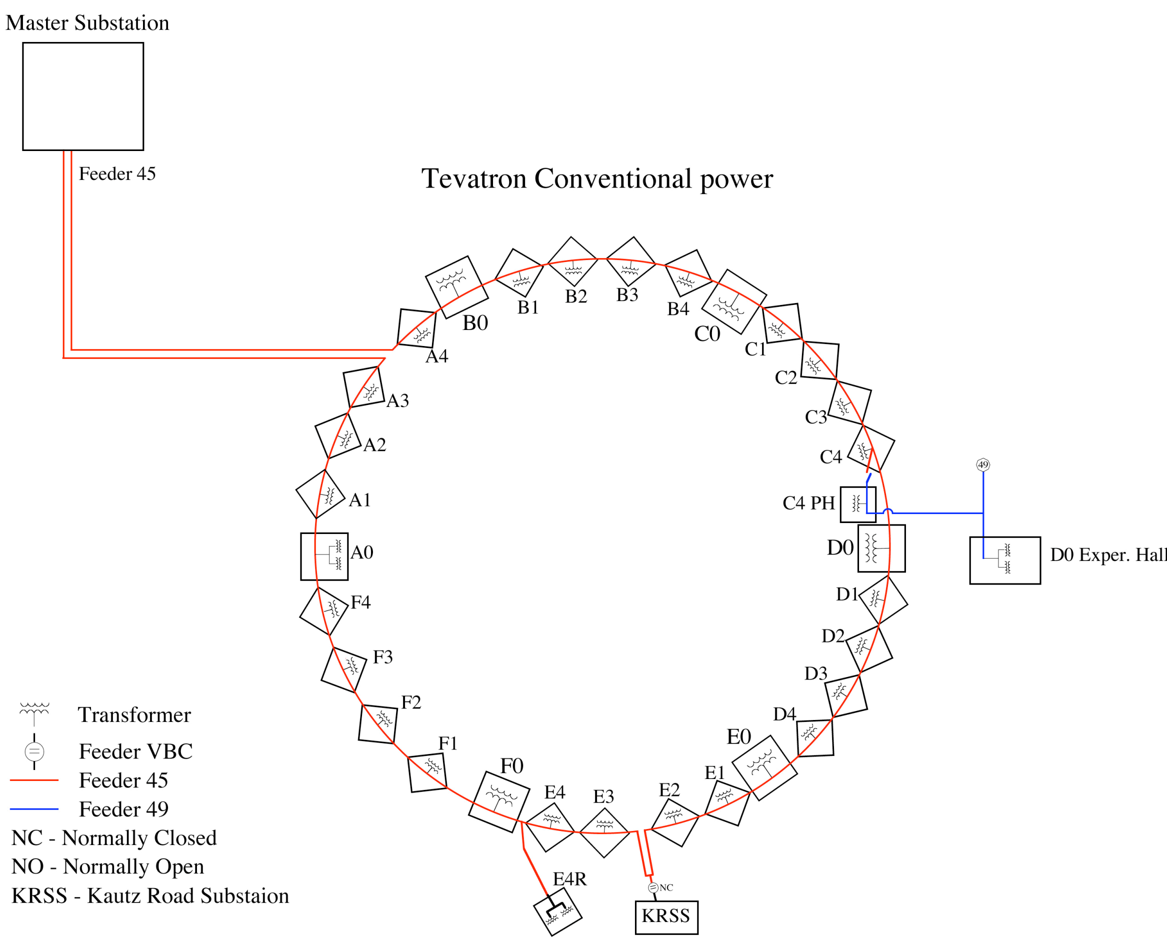

Fig. 3.14: Conventional power distribution for service buildings and collision halls

The Mycom compressors at each of the zero buildings used for maintaining helium pressures in the discharge and suction lines are on two separate feeders, 46A and 46B. Feeder 46A provides power to A0, C0, E0, and EA compressors. This circuit is fed from MSS and can be back fed from KRSS. Feeder 46B, which powers B0, BA, D0, DA, and F0, is fed from KRSS and back fed from MSS. Placing the feeders at separate substations was done so that if one feeder goes down the entire helium inventory would not be lost. Compressors from one feeder can partially maintain the pressure distribution around the ring.

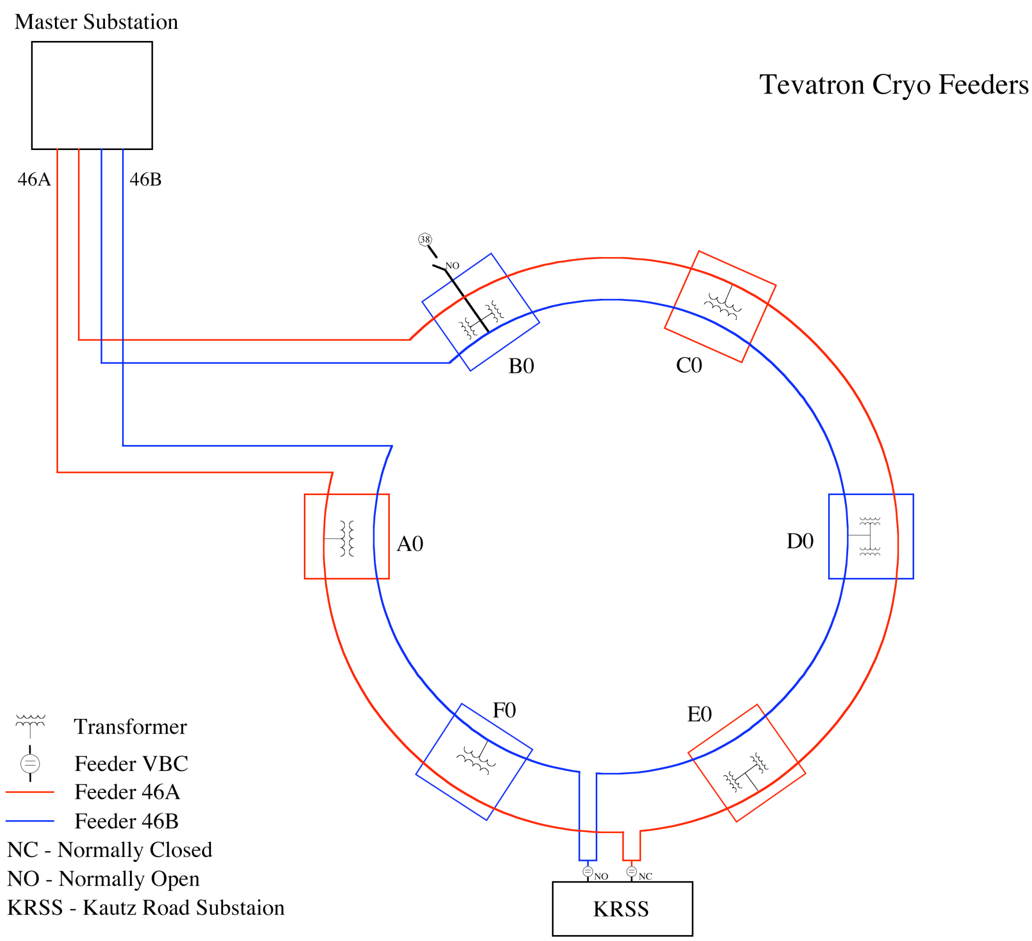

Fig. 3.15: Cryogenic compressor feeder distribution

LCW

The heat exchange system at each service building serves to cool Tevatron correction element power supplies, the warm bus, and the main power supplies. From F1 to F4, also known as the Main Ring Remnant, the heat exchange system provides additional cooling to remnant magnets, their power supplies, and the associated correction element power supplies. To accomplish all of this, the LCW heat exchanges with pond water in heat exchangers located at each service building.

The pond pumps are locates in the cement pit on the pond side of Ring road. These are used to circulate pond water, at a regulated flow, between the heat exchangers in the service buildings and the ponds outside. The desired temperature for the LCW is reached by diverting a portion of regulated LCW through the heat exchanger and mixing it back into the regular LCW flow.

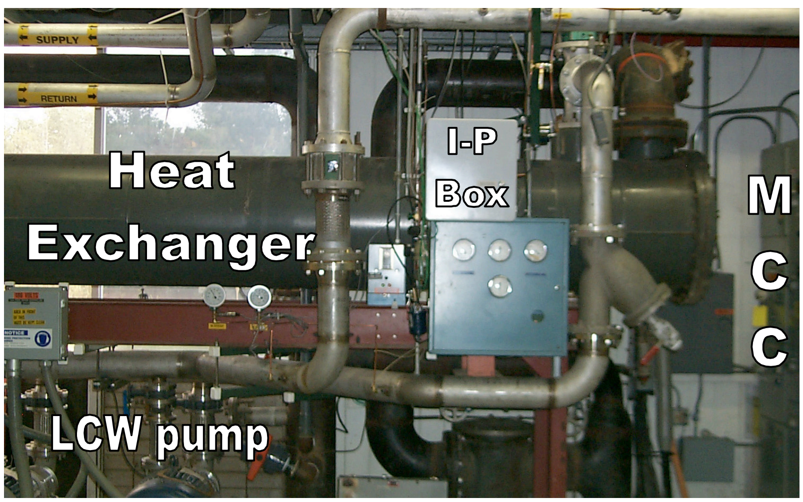

Fig. 3.16: LCW Equipment in a service building

LCW pumps are used to circulate the LCW through the heat exchangers, and out to the various heat loads. The service building LCW pumps alone accomplish the flow of LCW around the ring. Although the system is designed to be self-sufficient, one should still be aware that when a pump trips off it puts a burden on the two adjacent houses. The long straight sections have additional water requirements due to extra magnets in the tunnel and power supplies in the zero service buildings. Cooling problems can arise if an LCW pump is off at either the 1 or 4 building adjacent to the zero building. If necessary, the pond and LCW pumps can be run in local (or hand) at the Motor Control Center (MCC) of each house. These wall breakers are located opposite the service building door, behind the heat exchanger.

To maintain some form of consistency in the operating characteristics of the magnets and other heat loads, the supply temperature of the LCW is regulated between 90 degrees F to 100 degrees F, depending on the ring wide heat load. The Love 2600 controller regulates the temperature at the 1 and 4 buildings, with the 3 buildings serving as a back up. There is no regulation at the 2 buildings. The Love 2600 controller operates the LCW valve on top of the heat exchanger. So where's the Love. The electronics reside in a chassis mounted in the 3-bay-rack, opposite the electronics room.

Fig. 3.17: Love 2600 Controller

The controller operates a Current-to-Pressure transducer, which is supplied 20 psi air from the air compressor adjacent to the heat exchanger. The I-P converts the 4-20 mA control signal sent by the Lover controller into a 3-15 psi signal received by the pneumatic valve positioner/actuator (3=valve closed, no cooling; 15=valve full open, maximum cooling). A gray box in front of the heat exchanger contains the I-P. The key to access the I-P box is located in the Crew Chief's cabinet.

Under normal operation, all Love controllers are configured to operate off a remote set point. This set point can be modified via ACNET consoles. If the set point falls out of an acceptable range (65-115 degrees F), the controller will revert to its local set point. When the remote set point regains to an acceptable level the controller will again operate based on the remote set point.

The TeV LCW temperature control system also allows for local control of the valve in case of temperature regulation problems. There are 3 approaches one may take to valve manipulation; 1. adjusting the local set point, 2. manually setting the valve position with the Love Controller, or 3. bypassing the Love Controller/I-P entirely and manually controlling the amount of air to send to the valve.

Local Adjustment of the set point is accomplished by pressing the index key on the Love Controller, using the up or down arrows to change the set point, and pressing enter to lock it in.

The manual valve position control options are more crude ways of controlling the temperature, because the temperature regulation is completely bypassed, and no compensation for pond water temperature will be done, except for manual intervention.

Manually setting the valve position with the controller is accomplished by pressing the Auto/Manual Key on the Love Controller until the "MAN" lights up, using the up or down arrows to change the percentage of controller output (MAX = 20 mA), and pressing enter to lock it in.

To bypass the Love Controller/I-P, close the LCW valve using the MANUAL mode of the Love Controller if possible. Otherwise, turn the gradual switch all the way to the left (counterclockwise) before closing valves C and D and opening valves B and A.

Turning the gradual switch clockwise opens the LCW valve to the heat exchanger and provides more cooling. Once a manual valve position is selected using the gradual switch, the valve will not move if the system is undisturbed.

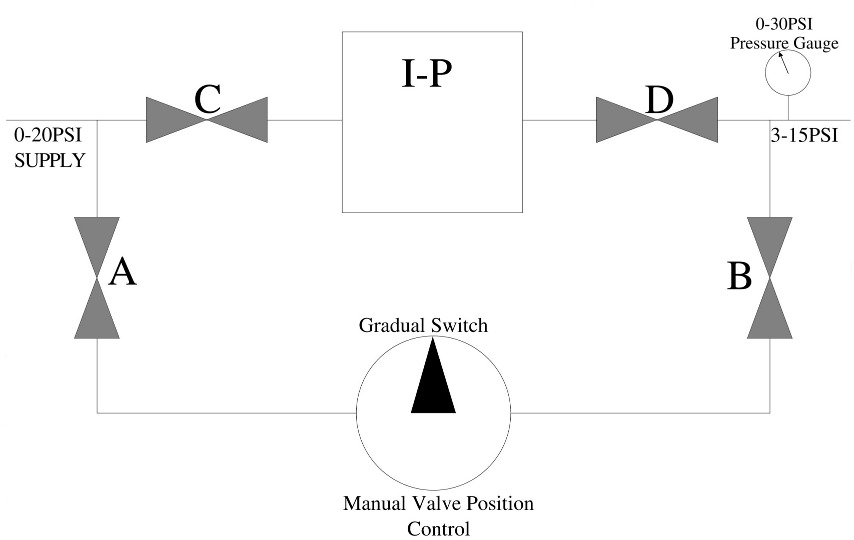

Fig. 3.18: Manual override schematic.

In the service building, the analog temperature gauges are mounted on the pipe near the enclosure door. In addition, the supply temperature is displayed on the top readout of the Love Controller. In the A-E sector #2 and 3 buildings, the pressure gauges are located behind the 3 Bay Rack. In all other buildings, the pressure gauges are located on the heat exchanger frame, above the pumps. In buildings with running pumps, the nominal supply pressure is about 180 psi and return pressure is around 30 psi with an average flow of 300 gpm.

4. Beam Manipulation Systems

What is a beam manipulation device? It is any component, other than the main dipoles and quadrupoles, that bends, focuses, separates, collimates, kicks, or in any other way intentionally affects the beam for some given purpose.

Dipole Correction Elements

The correction element dipoles are strong enough to steer beam at any energy. The correctors have a current range of +/-50 A. They are powered by a raw bulk supply per service building, which feeds one regulator per dipole. The regulators are driven by a programmed low level curve from a CAMAC 460 module, also known as a dipole function generator (DFG).

The C460 module is a programmable ramp controller whose functionality mirrors that of the C46x family and is capable of generating an analog waveform which is based on the time in the cycle, current (MDAT value), and change in current (MDAT value). The time portion of the waveform is updated at a 1 kHz rate, and the MDAT portion is updated at a 720 Hz rate. The C460 also contains digital control capabilities to turn on, turn off, and reset a power supply. It can also return sixteen status bits from the power supply.

Each DFG receives a digital number via the CAMAC 169 card in slot 1 of the DFG crate. The C169 card is a single wide CAMAC card that receives the serial information from the MDAT link, converts it to a parallel format, and daisy chains out to the C460 cards.

The regulators contain circuitry for protection of the magnets against ground faults, quenches, and protection of the regulator circuits against loss of water or over-temperature.

Collimators

Whenever proton and antiproton beams are injected into the TeV and ramped to 980 GeV they will always have a distribution of particles with some residing at lower and higher energies from the desired energy.

A finite fraction of the beam will move beyond the stable phase space Separatrix due to possible beam-gas interactions, intra-beam scattering, proton-pbar interactions at the interaction points inside the detectors, RF noise, ground motion, and resonances excited by magnet imperfections. These particles produce a beam halo, which can interact with the beam pipe to create electromagnetic and hadronic showers in the accelerator and detectors causing a higher background level at the interaction regions.

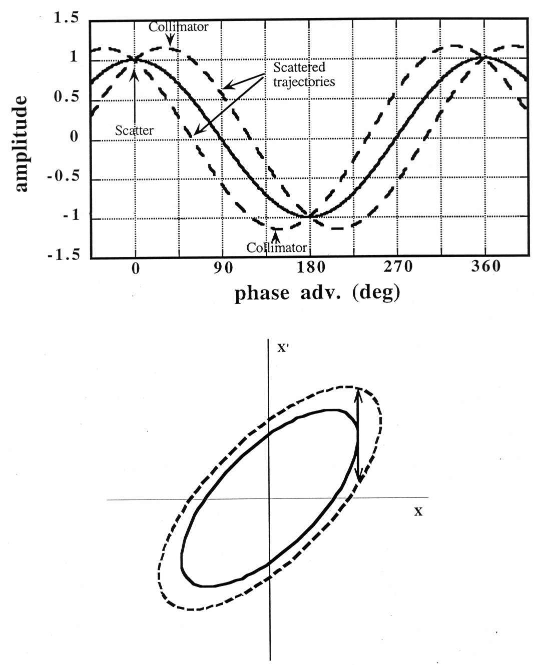

Fig. 4.1

The collimator system localizes most of the losses in the straight sections D17, D49, E0, F17, F48, F49, and A0. The system consists of horizontal and vertical primary collimators and a set of secondary collimators placed at a desired phase advance so it can intercept most of the out-scattered particles during the first turn after interaction with the primary collimators.

The primary collimator, often called the target, consists of a movable, narrow Tungsten target 5mm thick and the secondary collimator is a 1.5 m long stainless steel absorber. The target moves into the beam pipe about 5 s from the beam axis and the secondary collimator moves in about 7 or 8 s from the beam axis.

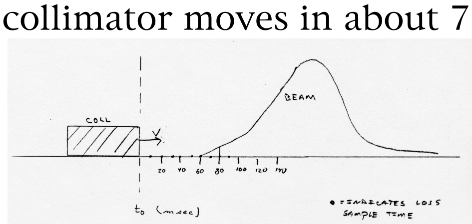

Fig. 4.2

Equipment and Control

The map below shows the location of the primary and secondary collimators around the ring. In Run Ib the typical scraping procedure took about 20 minutes to complete. The Run II system is automated. There are 4 sets of collimator systems for scraping away the halo, 2 are for proton removal and 2 are for pbar removal. A 5th collimator system at E0 is for proton removal after a store so that pbars can be recycled for future stores.

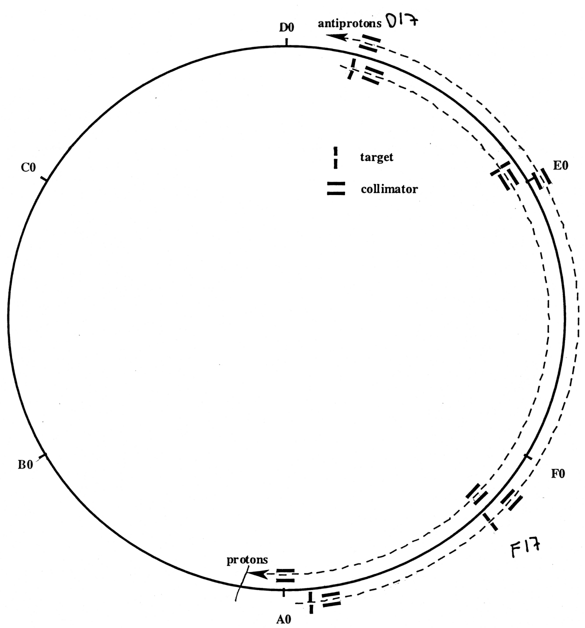

Fig. 4.3

A local processor running VXWORKS in a VME crate controls each collimator. The targets have a single motor for the vertical motion and a single motor for the horizontal motion. Secondary collimators have 2 motors in each transverse dimension to control upstream and downstream motion independently. The stepping motors are geared so that the collimator can be moved 1" in 13 seconds, which is the full distance from the out position to the beam axis. Position readbacks are provided by LVDTs (Linear Variable Differential Transformers); 4 per secondary collimator, 2 per target collimator. Limit switches protect the hardware from damage. Local feedback for the motion control, operating at 720Hz in the CPU, is provided by 4 standard TeV Balm's – 2 upstream and 2 downstream for redundancy. Stepping motors, loss monitors, and LVDTs are interfaced to the CPU via 3 IP's (Industrial Packs). Communication with ACNET is through Ethernet. A single VME crate can house multiple collimator systems.

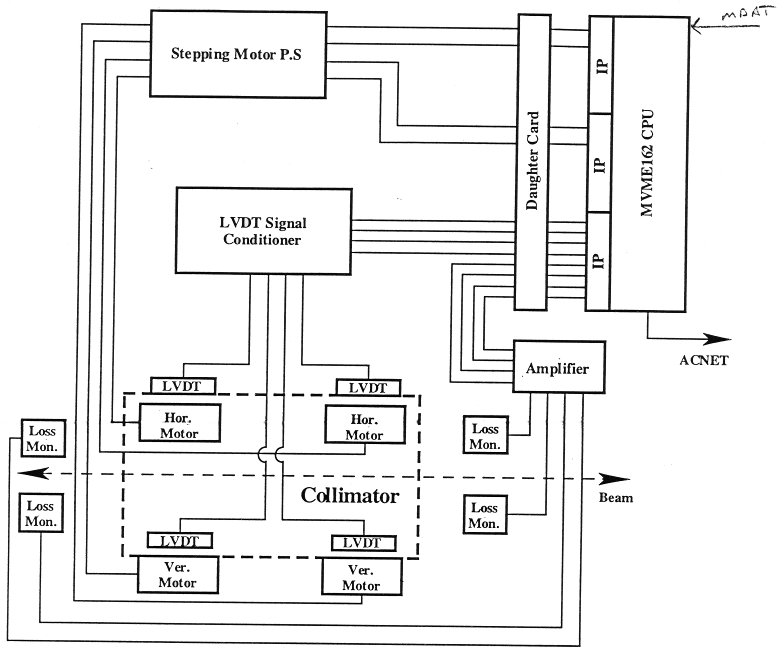

Fig. 4.4

An application program controls the beam halo scraping sequence. The program initiates motion for each collimator, waits for a completion status from that collimator, and then initiates the next collimator move. The application program communicates with an Open Access Client (OAC) that runs the algorithm for stepping the collimators in/out of the beam path. The collimator move commands are sent to the Collimator front end, which loads the commands into the local VMEs.

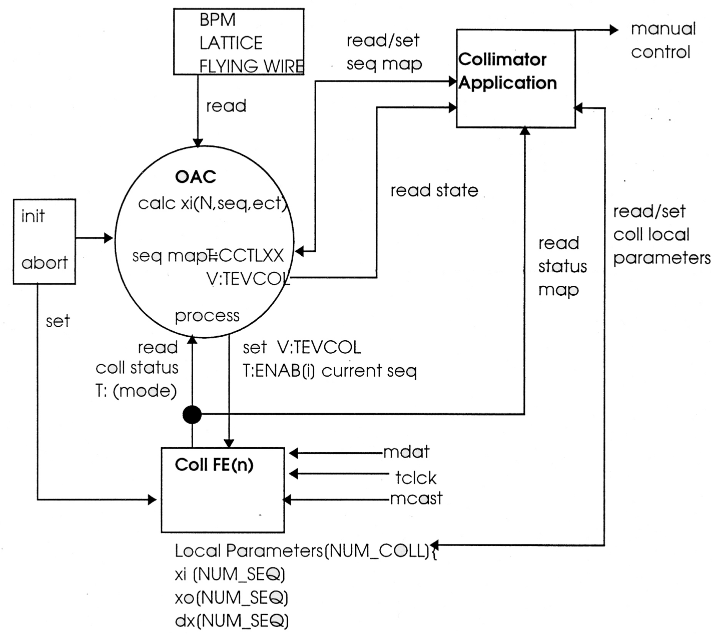

Fig. 4.5

Kickers

The Tevatron uses kickers for injection of protons and pbars and also for sending beam to the A0 abort. The proton injection kicker is located at F17 and the pbar injection kicker is at E48. Both are short batch kickers and have a rise time of 396 ns, which is the spacing between bunches. These kickers must maintain a 1.21 microsecond flattop and 2.61 microsecond fall time.

Beam-Beam Compensation (TEL)

In the Tevatron, the antiproton bunches suffer a tune shift due to their interactions with the more intense proton bunches. In multi-bunch operation, the tune shifts vary from antiproton bunch to antiproton bunch, leading to an effective spread in tune. An electron lens, consisting of a short, low energy, electron beam propagating along the axis of a solenoidal field, can induce a tune shift on the antiproton bunches opposite to that which they experience from the protons. Two such lenses could provide effective beam-beam tune shift compensation for both tune planes.

The main beam-beam concern for multi-bunch operation is that the bunches are not evenly spaced around the ring, different bunches within a train encounter the bunches in the opposing beam at different places around the ring. Proton and antiproton beams share the same vacuum pipe and, in addition to the two main interaction points at B0 and D0, there are many near misses. This causes differences between bunches in the train. Because of much higher intensity in the proton beam, the antiprotons suffer most from the beam-beam effects.

The tune footprints for most bunches are almost identical. However, pbar bunches 1 and 12, 13 and 24, 25 and 36 are shifted from the others because they do not see protons at the first crossing point upstream or downstream of the IPs, respectively.

Two electron beam setups for compensation of the beam-beam effects in the Tevatron (TEL- Tevatron Electron Lens) are located at F48 and A0. They provide the electron beams that collide with the antiproton beam. The electron beam is created on an electron gun cathode, transported through the interaction region in a strong solenoidal magnetic field, and absorbed in the collector. The electron charge is opposite to the proton charge. The electron beam can compensate for the electromagnetic force on antiprotons due to the proton beam. The proton beam has to be separated from the electron and antiproton beams in the device.

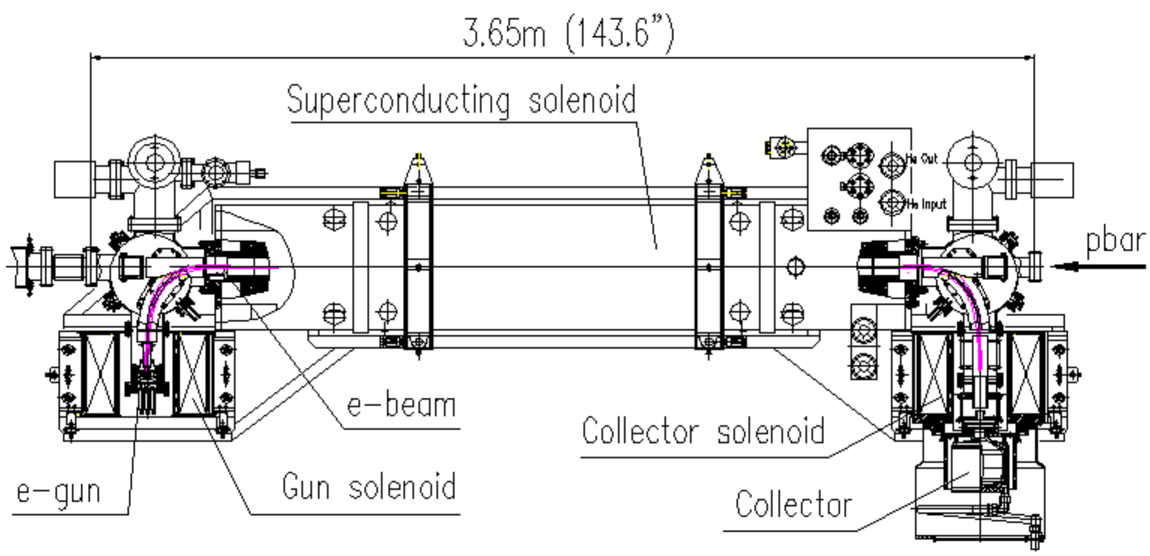

Fig 4.6

The TEL is equipped with 4 BPMs: one vertical and one horizontal at the beginning and at each end of the main solenoid. The BPMs measure transverse positions of electron, proton and antiproton beams passing through and thus, allow the electron beam to be centered on the antiproton or the proton one. 100 μm diameter tungsten wires, vertical and horizontal, can be introduced into the very middle of the interaction region for electron current profile measurements. They are remotely controlled and removed when high energy beams circulate in the machine. There are 10 HV electrodes around the electron beam trajectory that can be used for ion or secondary electron cleaning (though most of the time they are grounded).

The TEL vacuum under working conditions with 3 ion pumps with a total pumping speed of 300 l/s ranges from 4 to 10 e-8 Torr.

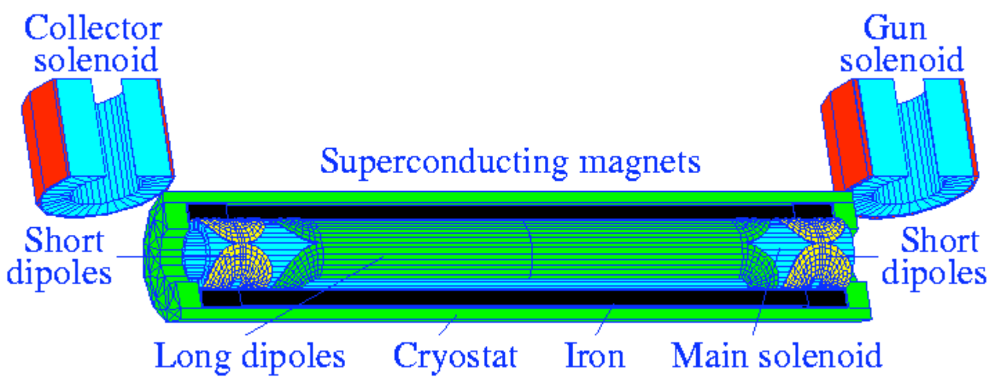

Fig. 4.7

The solenoid coil is constructed of a flat transposed cable consisting of 10 superconducting wires (NbTi filaments in copper matrix) each 0.85-mm diameter. The wire has 550 amps critical current at 4.2 K and 5 T. The dimensions of the bare cable are 1.44×4.64 mm2. Six steering dipoles are placed on the outer surface of superconducting solenoid coil. Four pairs of 250-mm long coils form (short) lateral vertical and horizontal dipoles at each end of the solenoid. Two pairs of 2-meter long coils are placed in the central region of the superconducting solenoid. All these dipoles are to correct the electron beam trajectory inside the magnetic system. The steering dipoles are wound of cable transposed from 8 wires of 0.3-mm diameter. The wire has 50 A critical current at 4.2 K and 5 T. Dimensions of bare cable are 0.45×1.48 mm2. The lateral dipole cable is made of superconducting wires only. The current in central dipoles is small, and the cable has three superconducting wires and five Cu wires. The central dipoles have one layer; lateral dipoles consist of two layers and an inter-layer spacer of 0.2-mm thickness.

The superconducting solenoid coil together with steering dipoles is enclosed in a magnetic shield made of low-carbon steel. The shield is 48.5-mm thick over the length of 270 mm and 38.5-mm thick in the central part over 1.96-m length. The yoke reduces currents in steering coils, improves homogeneity of magnetic field inside the solenoid aperture, compresses magnetic field lines at the ends of the coil block, and reduces stray fields.

The solenoid coil is not self-protected against resistive transition and fast quench detection, so removal of stored energy to the external dump resistor must be provided. The quenching lasts about 2 s. 90 % of the stored energy (about 1 MJ at 6.5 T) dissipates in the dump resistor and 10 % inside the cryostat, and the maximum temperature at the hottest point in the coil is about 270 K.

The energy stored in the dipoles is much smaller, about 1.3 kJ, and, in principle, one can allow all the energy to be dissipated in the coil if the quench is detected and the current is interrupted. In that case, the hot spot temperature will not exceed 120 K. However, to lower the risk of spreading the quench to the main solenoid, the scheme of quench protection with an external dump (as for the main solenoid) is also used in this case. The hot spot temperature does not exceed 43 K for lateral, and 29 K for central, dipoles.

Quench protection circuits for each superconducting coil compare the voltage across the coil with LdI/dt. If the difference exceeds 1 V, a signal is sent to high current IGBT switches to disconnect the coil from power supply and to dump the coil current into the resistive load. Mechanical current breakers are installed in series with the solid state switches for redundancy.

The gun and collector solenoids have almost identical design. Each is wound of 8.25×8.25 mm2 Cu conductor with a 5.5-mm diameter water hole. The solenoid has a 0.4 T nominal magnetic field, 0.19-Ohm electrical resistance, and 18-mH inductance. The coil has 250-mm inner diameter, 474-mm outer diameter, and 300-mm length. The solenoid coil consists of 17 pancakes (total number of turns 391), which are assembled on a common pipe of a 240-mm inner diameter. Water temperature rise in the coil is 300 C at 0.7 MPa pressure drop and nominal current of 340 A. About 100 A of operating current are needed, in the short steering superconducting dipole, in order for the electron beam to be transported along the center of the warm solenoid.

Electron beam shape and position correctors are set inside each of the conventional solenoids. The corrector consists of four coils, which can be commutated either as a quadrupole or as two dipoles (vertical and horizontal). Each coil layer is shaped with 0.74° inner and 40.04° outer angles, 112.5-mm inner radius and 8.6-mm thickness. The length of coil is equal to 298 mm. The dipole field is equal to 19 G/A; the quadrupole field is equal to 6 G/A/cm.

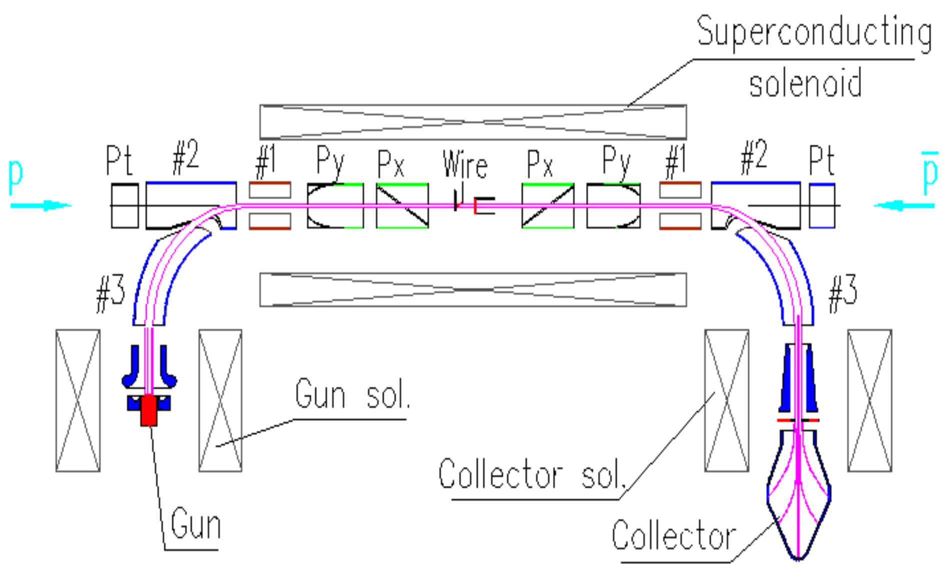

Fig. 4.8

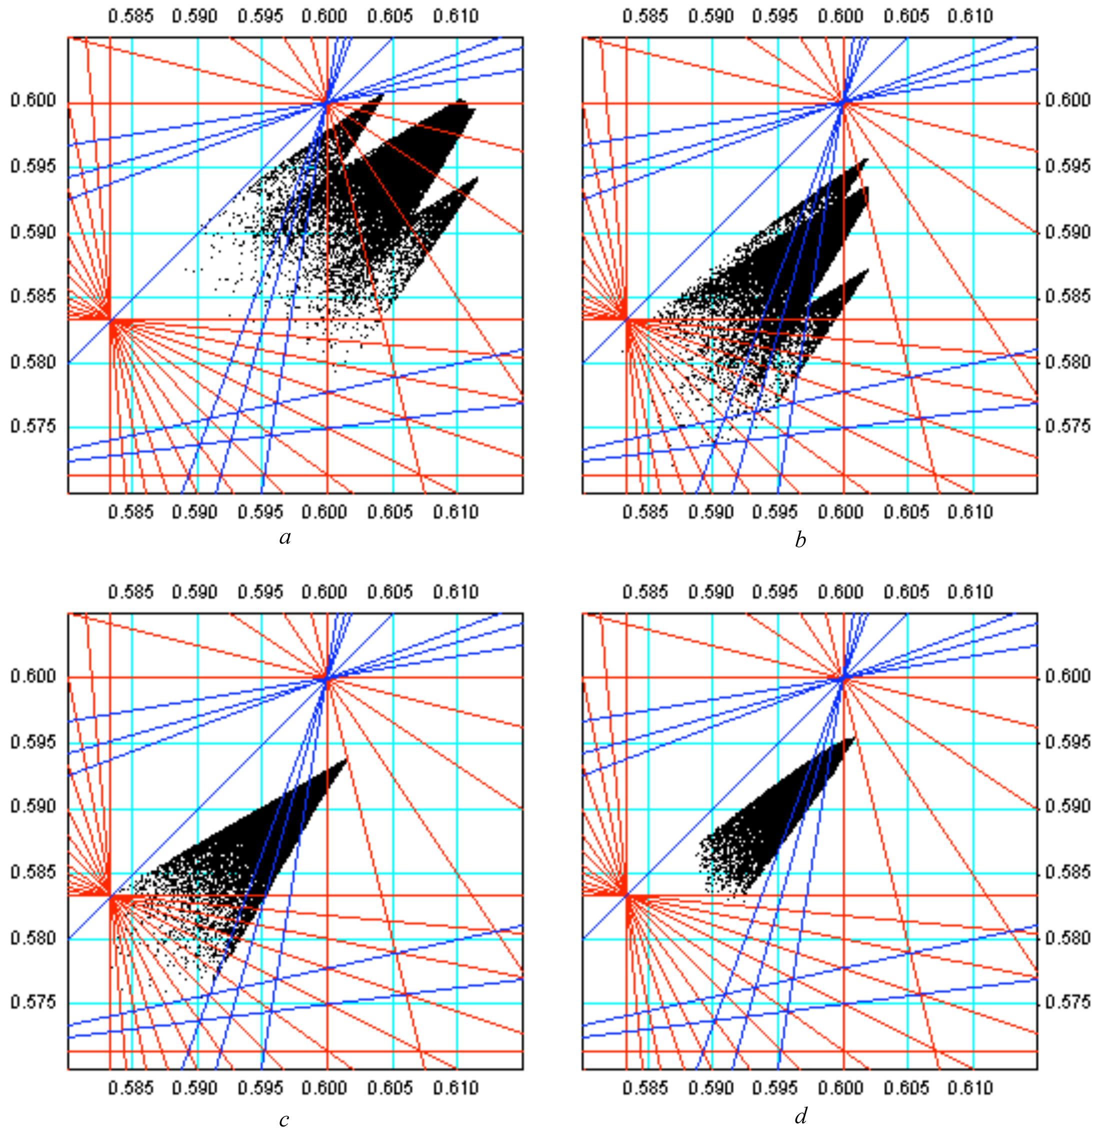

Fig. 4.9: Tune spread without beam-beam compensation during a store (a). Beam-beam compensation in both planes (d)

TEL Vacuum Interlocks

There are certain power supplies in the Tevatron Electron Lens (TEL) systems, both TEL1 and TEL2 that need to be disabled in the event of bad vacuum in the Tevatron. Two vacuum monitoring systems have been installed for both TEL machines to provide permit interlocking signals. Input signals to the interlock chassis are outputs of ion pump and vacuum gauge power supplies whose corresponding vacuum devices in the beam line are those that are located closest to the electron guns and collectors. Comparators in the interlock chassis monitor the ion pump power supply and ion gauge signals and disable the permit when vacuum is out of range. The permit signal is provided as both a TTL signal and a relay contact. The interlock chassis are connected to Ethernet for remote control and monitoring, and each interlock chassis utilizes an IRM chassis as an ACNET front end. Table 1 lists the hardware physical locations and associated Ethernet assignments. Each chassis has one ACNET parameter defined for the purpose of monitoring a trip condition and to provide the ability to reset it.

The interlock chassis have input connectors for as many as 5 vacuum devices, although presently 3 are used for TEL1 and 4 are used for TEL2. The interlock circuits contain window comparators that monitor the input vacuum signals and sum the comparator conditions to generate the permit signals. The upper and lower limit values of the four window comparators are settable remotely. The last saved limit values are stored permanently in the interlock chassis. Note that the hardware and firmware for both TEL systems is configured the same, likewise ACNET parameters L1VAC and L2VAC. Channel 3 is jumpered to channel 4 on the front of the TEL1 interlock chassis at the present time for this system to work properly.

Window comparators are used for checking vacuum limits because the vacuum monitor signals out of the ion pump and ion gauge power supplies have fault conditions at both high and low voltages. The vacuum signal voltage range is 0 to 10 V. When high voltage on the ion pump power supplies is turned off, the vacuum monitor output voltage goes up to about 9.8 V. Excellent vacuum is also a high voltage, but not as high as when the supply is off. Poor vacuum is a low voltage. At the low voltage end, the monitored voltage asymptotically approaches roughly 2.25 V for vacuum approaching 10-5 Torr and lower.

The hardware interlock circuitry is an assembly composed of two PCBs: a digital board and data acquisition board. The digital board communicates over Ethernet with an ACNET front end.

		TEL1	TEL2
I.R.M. Front End	node name:	67D	67B
	IP address:	131.225.128.110	131.225.128.4
	Chassis location:	ZG5-ZRR-57	TG9-ZRR-103
Vacuum Interlock Chassis	domain name:	Tel1vacintlk.fnal.gov	Tel2vacintlk.fnal.gov
	IP address:	131.225.139.126	131.225.139.2
	local port:	4524	4524
	VLAN:	3	3
	Chassis location:	Relay rack A01-04	Relay rack A01-03

Table 4.1: Ethernet assignments

ACNET parameters and chassis functionality

The ACNET parameters defined for monitoring trip conditions and providing a reset are T:L1VAC and T:L2VAC. These are defined in Table 2. These parameters also indicate which of the input channels was the first to trip out of range. Trip conditions are latched and must be manually reset by an operator before the interlocked TEL equipment will resume operation.

The remaining parameters provide the ability to set the trip limit values for the window comparators. These parameters are listed in Tables 3 and 4 that also shows which vacuum device is monitored by which input channel. Table 5 defines the properties for all the trip limit parameters. Their properties are all the same. Keep in mind though that for an upper trip limit, the input signal must go above the set voltage value to trip the permit. For a lower trip limit, the input must go below the set value.

There is one ACNET parameter for each trip limit. There are, therefore, two parameters associated with one input channel. The trip limit parameters display the vacuum of the device associated with it in the reading field. In the case of the readings, both limit parameters of a channel will display the same vacuum reading. The setting field displays the trip limit setting value and is settable. The trip limit settings are archived permanently in the interlock chassis and will independently return to the previous operating settings after a power outage or a power-on cycle.

The status field of each trip limit indicates whether it is tripped out of range high or low. In addition, if it was the first to trip in the chassis, its third status bit will indicate. All of the trip limit parameters have reset capability. Resetting any of the parameters will result in resetting all of the latches in the chassis.

The interlock chassis are always on. There is no way to remotely defeat the permit function. If one of the ion gauges is broken and is stuck with at either 0.0 or 10.24V, for example, about the only way to remedy the situation is to jumper two channels together at the input of the interlock chassis. It may be necessary to change limit values of the jumpered channel if this occurs.

Analog Properties	
Setting field:	none
Reading field:	none
Common units:	n/a
Primary units:	n/a
Digital Properties	
Status bit 1:	"." (grn) – when there is no trip
	"1" (rd) – Vacuum channel #1 was the first of the four to have tripped out of range
Status bit 2:	"." (grn) – when there is no trip
	"2" (rd) – Vacuum channel #2 was the first of the four to have tripped out of range
Status bit 3:	"." (grn) – when there is no trip
	"3" (rd) – Vacuum channel #3 was the first of the four to have tripped out of range
Status bit 4:	"." (grn) – when there is no trip
	"4" (rd) – Vacuum channel #4 was the first of the four to have tripped out of range
Reset Control Bit	This reset bit will reset all four vacuum channel comparator latches. (If the vacuum is still bad trip conditions will quickly reappear.)

Table 4.2: T:L1VAC and T:L2VAC. TEL vacuum interlock trip and reset parameters

ACNET Parameter	Trip Limit	Interlock Chassis Channel #	Vacuum Device
T:L1VM1U	upper	1	T:ATIG8A
T:L1VM1L	lower		
T:L1VM2U	upper	2	T:ATIG8B
T:L1VM2L	lower		
T:L1VM3U	upper	3	T:ATIP8M
T:L1VM3L	lower		
T:L1VM3U*	upper	4	T:ATIP8M*
T:L1VM3L*	lower		

*Redundant information. Channels 3 and 4 are jumpered together

Table 4.3: TEL1 vacuum interlock ACNET trip limit parameters

ACNET Parameter	Trip Limit	Interlock Chassis Channel #	Vacuum Device
T:L2VM1U	upper	1	T: A0IGTU
T:L2VM1L	lower		
T:L2VM2U	upper	2	T:A0IPTA
T:L2VM2L	lower		
T:L2VM3U	upper	3	T: A0IPTB
T:L2VM3L	lower		
T:L2VM4U	upper	4	T: A0IPTD
T:L2VM4L	lower		

Table 4.4: TEL2 vacuum interlock ACNET trip limit parameters

Analog Properties	
Setting field:	Setting field is the trip limit setting—either upper or lower.
Reading field:	Reading field is the vacuum reading of the
Reading common units:	Primary units scaled to 0 – 10.24 V. Displayed units are "V".
Reading primary units:	0x0000 – 0xFFFF 16-bit binary scaled to 0 – 10.24 V. Displayed units are "V".
Digital Properties	
Status bit 1:	"." (grn) – Vacuum voltage is not excessively high.
	"U" (rd) – Upper limit trip. Vacuum voltage is out of range--high.
Status bit 2:	"." (grn) – Vacuum voltage is not excessively low.
	"L" (rd) – Lower limit trip. Vacuum voltage is out of range--low.
Status bit 3:	" " (blank)
	"1" (rd) – This vacuum channel is the first to have tripped out of range.
Reset Control Bit	This reset bit will reset all four vacuum channel comparator latches. (If the vacuum is still bad trip conditions will quickly reappear.)

Table 4.5: ACNET parameter properties for all upper and lower trip limit parameters

Notes:

5. Cryogenics

Why Cryogenics?

The niobium-titanium alloy used in the Tevatron magnet coils is only superconducting when it is at a temperature of a few degrees above absolute zero. To complicate matters further, the threshold of superconductivity drops even lower when current is present in the coil, or if it is permeated by a strong magnetic field. It would be counterproductive to try to eliminate the latter two conditions, so the only way to achieve superconductivity in the magnets is to get them cold enough. (Review Chapter 2 for more detail on the conditions required for superconductivity.) For 800 GeV operations, the coil needs to be around 5K; at 1 TeV, where current and field strength are higher, the temperature of the coil needs to be close to 4K. To establish those temperatures, the magnet coils and the surrounding stainless steel collars are immersed in a slowly flowing stream of liquid helium. Helium is the only known substance that is a liquid at the desired temperatures. Oxygen, nitrogen and most other substances have long since frozen at 5K.

Liquid nitrogen, which liquefies at 77K, is also used in the magnets, but only as an outside layer of insulation for the helium. Even then, a layer of vacuum and reflective Mylar insulation separates the helium and nitrogen in order to keep the helium from being overheated by the nitrogen. (Again, review Chapter 2, which includes a cross-section of a Tevatron dipole.)

Heat energy always flows from warmer to cooler. That unfortunate principle of thermodynamics makes it difficult to create an environment that is several hundred degrees colder than its immediate surroundings. (The interior of a Tevatron magnet, in degrees Fahrenheit, is 400° below the coldest temperature ever recorded in the Chicago area.) But there are a few standard techniques that can be used to produce cryogenic liquids, including the judicious application of compression and expansion.

Compression of a gas heats the gas. Suppose that helium is compressed with a piston. Its temperature will rise—not only because some of the energy of the moving piston may be transferred to the helium, but also because the energy originally present in the gas has been pushed into a smaller volume. A small, hot volume of gas can have the same amount of heat energy as a large, cool volume. (Remember that the only temperature at which a substance has no heat energy is at absolute zero. Even an ice cube contains a lot of heat compared to what it might have at a lower temperature.) The heat energy can then be easily removed from the compressed gas as it cools to ambient temperature.

Once cooled, the compressed gas can be allowed to expand; now it will drop below ambient temperature, because its limited supply of heat energy is spread out over a larger volume.

It is possible to cool a gas even further if it is allowed to do work as it expands, as against a piston. The gas provides the work (energy) required to move the piston. Since the gas is expanding at the same time as it is losing energy, the gas is more effectively cooled because there are actually two mechanisms at work. Much of the refrigeration for the Tevatron ring uses piston-driven expansion engines.

Sometimes in the Tevatron, the helium does not do work as it expands and cools. An example of this process, which is less efficient than that of the expansion engines, occurs when the helium passes through a narrow aperture called a Joule-Thompson (JT) valve.

The last major principle of cooling to be mentioned here is that of heat exchange. Heat exchangers take advantage of the fact that heat flows from warm to cold. A gas or liquid, which is to be cooled, is brought into proximity (but not direct contact) with gas or liquid that is already cold. Usually the two substances flow in opposite directions.

By carefully designing the sequence of compressors, expanders, and heat exchangers, it is indeed possible to cool helium and nitrogen to the temperatures necessary for operation of the Tevatron.

There are two sources for the cryogenic liquids used by the Tevatron magnets. One is the ring system of refrigerator buildings and compressors. The ring system includes compressor buildings at each of the “zero” locations around the Tevatron ring, and also one at SSB. The compressors pressurize the helium coming back from the magnets and send it to the 24 refrigerator buildings distributed around the ring. The pipe transporting the helium is the discharge header, also known as the 3” header. The refrigerator buildings are located on the Tevatron berm, behind the “1” through “4” service buildings. Heat exchangers just outside the buildings, and expansion engines inside, liquefy the helium; each house is responsible for providing helium to upstream and downstream "strings" of magnets. The strings extend approximately halfway to the adjacent houses. Helium returning from the magnet strings is dumped into the suction header (also known as the 8" header). Now at low pressure, the helium gas is pulled in by the compressors and pressurized.

The second source of cold helium, and the only source of liquid nitrogen, is the Central Helium Liquefier (CHL). CHL is a large building off to the side of the Tevatron ring, filled with gigantic compressors, expansion turbines and Dewars, all dedicated to producing a constant supply of liquid helium and nitrogen to the transfer line. The transfer line originates at CHL and meets the Tevatron berm at a point between the A3 and A4 refrigerator buildings. From there it goes counterclockwise, on top of the Tevatron berm, servicing each of the 24 refrigerator buildings. The transfer line supplements the liquid helium being produced by the expansion engines, and also supplies the liquid nitrogen needed by the heat exchangers and magnets.

Since helium is constantly being fed into the Tevatron ring from CHL, sooner or later an equivalent amount of helium has to be returned to CHL. It returns in the form of excess high pressure gas from the 3" header.

All of the standard techniques described in the previous section—compression, expansion, and heat exchange—are put to use in the refrigerator buildings. The normal life cycle of the helium in the ring is summarized in the steps below. More details will follow later in the chapter. The numbers in the sequence correspond with those in Fig. 5.1:

1. Compressors around the ring take up helium from the suction header, pressurize it, and fill the discharge header. The helium, which has been heated by the compression, is heat exchanged with water in order to bring the temperature down to approximately room temperature.
2. Pressurized helium from the discharge header is again cooled, this time in a heat exchanger located outside the refrigerator building. The cold counterflow may be liquid nitrogen from CHL, cold helium returning from the magnet

strings, or even a portion of the pressurized helium that has been cooled in an expansion engine.

3. The cold pressurized helium is allowed to expand, cool further, and liquefy. This is usually done through an expansion engine (the wet engine), but a JT valve is available as a backup.
4. Heat exchanging with liquid helium from a Dewar cools the liquid further. The Dewar itself is filled with liquid helium from CHL and cold helium returning from the magnet strings.
5. Just before it enters the magnet string, in the feed can, the liquid helium is heat exchanged with the cold helium returning from the magnet string.
6. Splitting between the upstream and downstream strings, the liquid helium makes its first pass through the magnet strings, immersing the coil and the stainless steel collars. Here the helium is referred to as one-phase or single-phase, the "phase" in this case being liquid.
7. At the end of each string, the liquid helium passes through the magnet JT valve, where it expands, thus cooling further, and is partially converted to a gas. The magnet JT valves are located at the turnaround boxes at the ends of each string.
8. The expanded helium, now known as two-phase (liquid plus gas), is used as a counterflow to cool the one-phase further. The two-phase circuit is a shell just outside the one-phase. Flow is back toward the feed can.
9. Upon reaching the feed can, the two-phase helium is nearly boiled off and can now be considered a cold gas. It is used to heat exchange with the incoming gas in the feed can, cool the Dewar, and finally cool the pressurized helium in the heat exchanger.
10. Beam energies greater than 900 GeV require colder temperatures in order to maintain the magnets in a superconducting state. A cold compressor, located downstream of the Dewar, pumps down on the two-phase helium to lower its pressure, and thus its temperature. The colder two-phase lowers the temperature of the Dewar and the magnets appropriately. The discharge of the cold compressor passes through the heat exchanger; although it is

heated some by the compression, it is still cold enough to be effective as counterflow to the incoming helium.

11. After extracting as much heat as possible from the incoming gas in the heat exchanger, the low pressure helium enters the suction header, finds its way to the compressors, and is pressurized. The cycle starts over.
12. Throughout this process, CHL provides liquid nitrogen for use in the heat exchanger and the magnets, in addition to its role mentioned earlier of providing liquid helium.

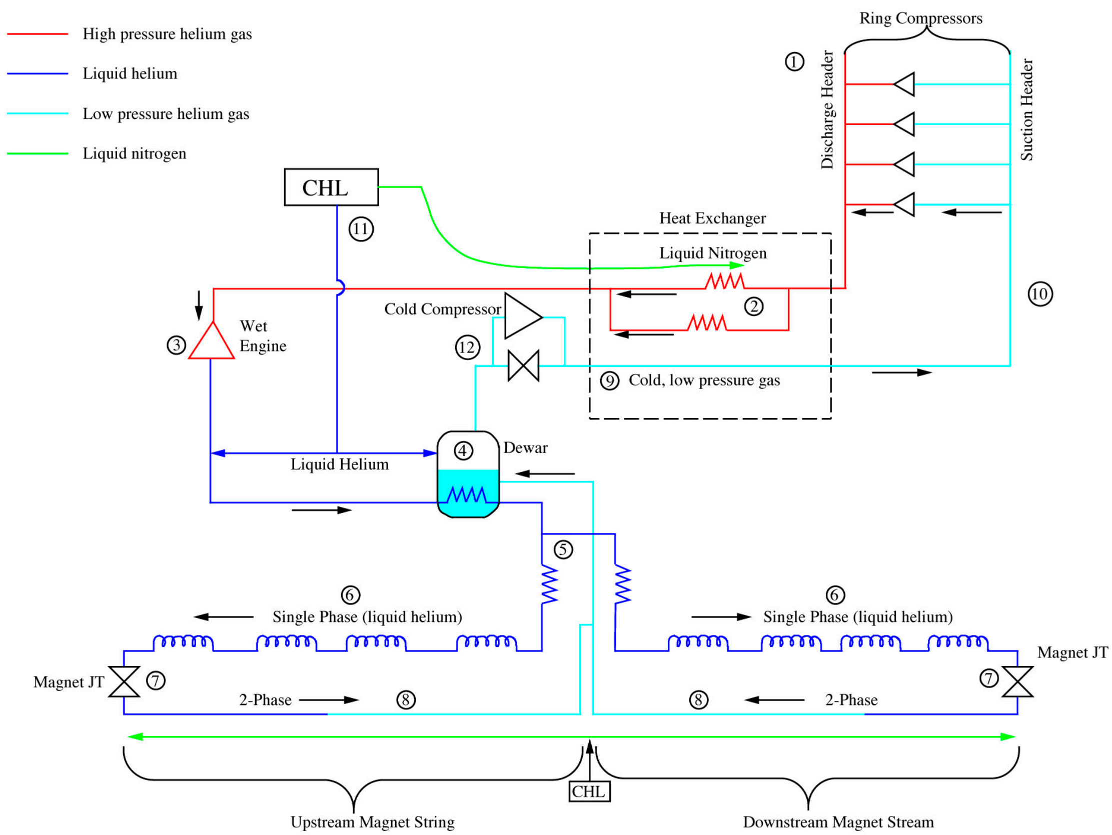

Fig. 5.1 Cryogenic Overview
See text for details

Instrumentation

Before describing the components of this scheme in more detail, it will be useful to look at some of the methods used to measure pressure and temperature.

Pressure is usually measured in psig. The psig unit needs to be distinguished from the slightly more familiar unit psi (pounds per square inch). The g stands for "gauge," i.e. the pressure transducer itself. Since pressure gauges are often calibrated to read zero pounds at atmospheric pressure, a measurement in psig represents the value above the standard atmospheric pressure of approximately 15 psi. This distinction is important in cryogenics, because if the psig of a helium circuit goes negative, it is "sub-atmospheric;" it is possible that freezable gases such as oxygen and nitrogen could leak in and clog up the plumbing. Parameter names for pressure gauge readbacks usually include "PI" (pressure indicator) as part of the name. Occasionally units of pressure are given in psi, or psia (the a standing for "absolute.")

There are three ways of measuring temperatures in the cryogenic systems. One is known as a VPT (vapor pressure transducer). Vapor pressure is actually measured in the vicinity of a liquid. In every liquid, there are at least a few molecules that have enough energy to break free of the liquid; these molecules are in a gaseous state and are responsible for the vapor pressure.
The higher the temperature, the greater number of free molecules can be found in the vapor. If the temperature reaches the boiling point of the liquid, the vapor pressure equals the pressure confining the liquid, and all of the liquid begins to convert into a gas.

The VPT's are actually liquid-filled bulbs fitted with pressure transducers. Although they read back in units of psig, the vapor pressure actually represents a temperature. The bulbs may be filled with nitrogen, hydrogen, neon, or helium—the choice depends on what temperature range is desired (of course, it would not work for the liquid to freeze solid in the environment it is supposed to be monitoring). Tables relating temperature to vapor pressure can be found in the parameter page HELP files. VPT parameter names usually include the phrase "TI" (for Temperature Indicator).

The second type of "thermometer" is the carbon resistor. These carbon resistors are simply precision versions of those used on any electronics board. Since the resistance varies with temperature, a constant current source upstairs is fed into the resistor; the voltage across the resistor can be converted into resistance, and the resistance can be scaled to temperature. (The scaling is not as obvious as one might think. In the cryogenic temperature range, the resistance of a carbon resistor is inversely proportional to temperature, and nonlinear at that. Fortunately, resistance values are normally converted to temperature values in the database before they are read on a parameter page.) Carbon resistors are at their best at the coldest temperatures (<10K), but very inaccurate at room temperature. Carbon resistors include “TR” (Temperature Resistor) in their name.

Knowing where each carbon resistor is located is essential when cooling a magnet string (Fig. 5.2 (a), (b), and (c)). To a first approximation, the resistors are named after the nearest quadrupole—using the E2 strings as an example, the resistor at the base of the feed can is in the spool piece next to the quadrupole at E25; it is named TRQ5. Going upstream, the numbers count back to TRQ1, in the turnaround box; downstream, TRQ9 is in the last spool piece in the string, at E29. In the case of the “1” houses, the numbers count up to TRQ0. (The downstream string at E1 continues to the spool piece at E21, so “0” really means “10.”)

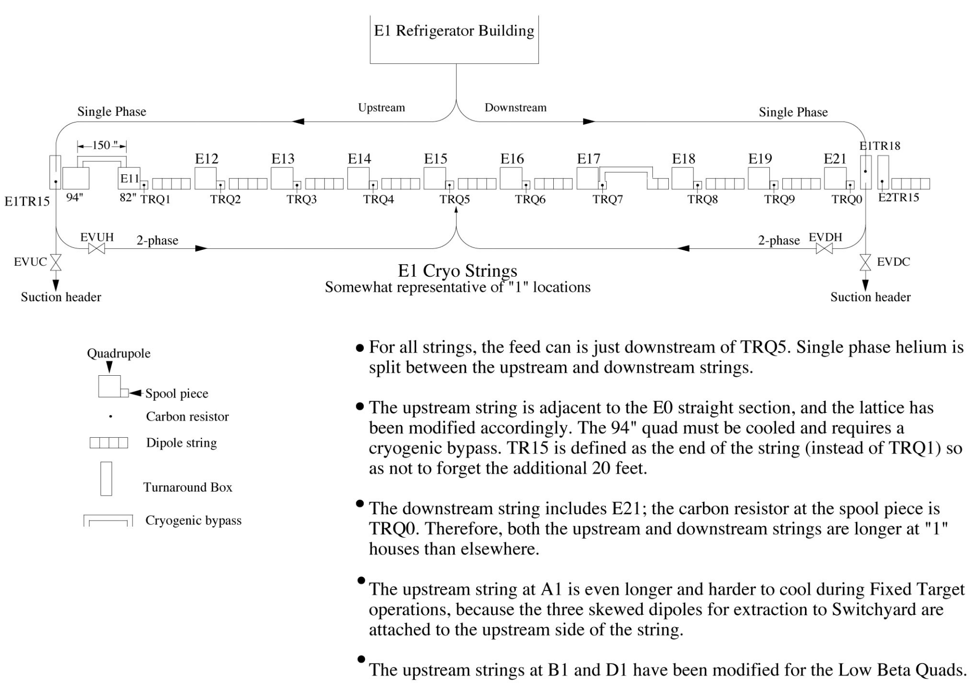

Fig. 5.2 (a) Cryogenic Layout of a "1" House

The turnaround boxes themselves are instrumented with TR15 (upstream) and TR18 (downstream). Sometimes TR15 or TR18 should be used to evaluate temperature at the end of the string. In the "1" houses, TRQ1 is separated from the turnaround box by two quads and a cryogenic bypass—or at B1 and D1, by low beta quads—so TR15 should be used instead. At the "4" houses, TRQ9 is separated from the turnaround box by a 94" quad. Another important carbon resistor is TR12, which measures the temperature of the helium entering the heat exchanger.

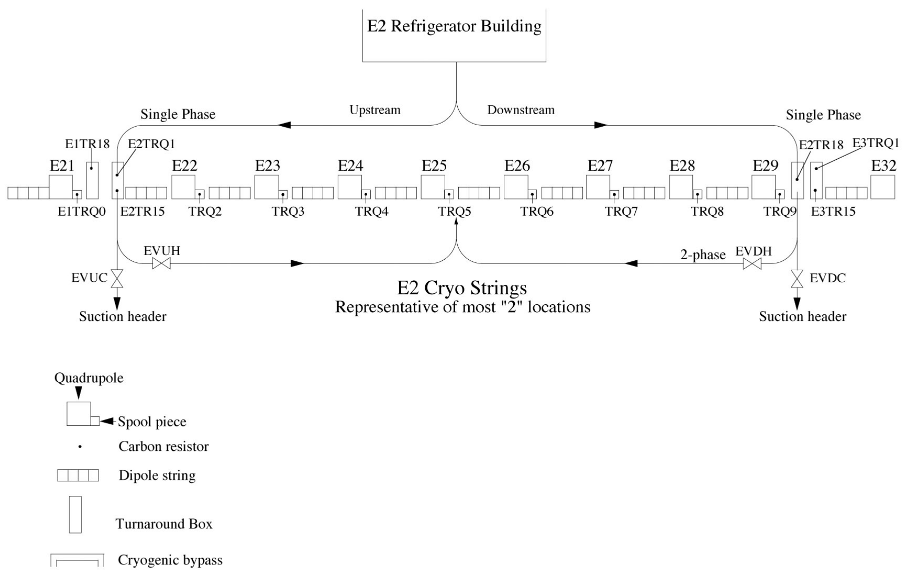

Fig. 5.2 (b) Cryogenic Layout of a "2" House

For measurements in the warmer regions of magnet strings, platinum resistors are used. They operate on the same principles as the carbon resistors, except that they are at their most accurate near room temperature. "TP" is a part of every platinum resistor's name.

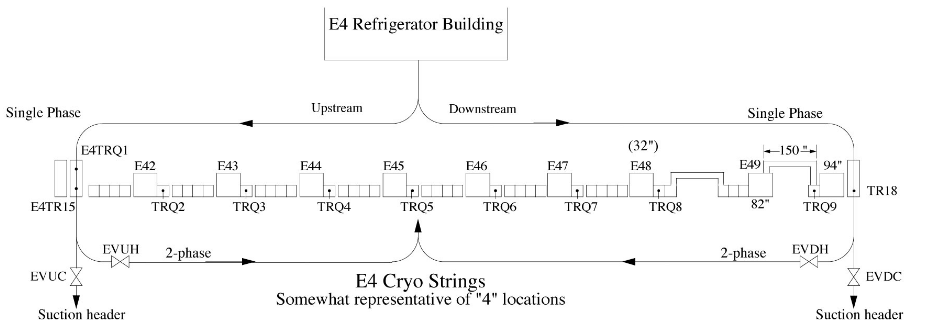

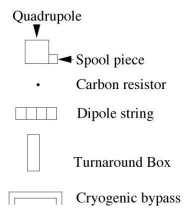

- The downstream string interfaces to the F0 straight sections; the lattice at the end is a mirror image of that at E1.
- There is an additional cryogenic bypass at the "48" location to accomodate the short straight section.
- The downstream strings at A4 and C4 have been modified for the Low Beta Quads.

Fig. 5.2 (c) Cryogenic Layout of a "4" House

Control Loops

There are a few generalized comments about control loops that should be made before moving on to the specific refrigerator components. The refrigerator system is a large, complex, and dynamic one, and the machinery must constantly be adapting to changing conditions. Most of the valves and engines are operated by feedback loops; they adjust their output, within minimum and maximum bounds, to try to match a measured value to a set point. A change in one component of the system may prompt a response from another component. The following discussion will emphasize the way that these variables interact in a steady-state operational mode, but necessary complications will be have to be introduced later.

Finite State Machines (FSMs) can take control of the loops. The FSMs are also used to control other devices, as explained later.

Ring Compressors

The ring compressors pull in low pressure helium exhaled by the magnet strings, and pressurize it so the expansion engines and JT valves in the refrigerator buildings can cool it. There are compressor rooms at all of the zero locations (A0, etc.) and one at SSB. Each compressor building normally houses four compressors (except for SSB, where there are only two.) If there are more than four, the additional compressors are grouped separately, at least from a naming and controls standpoint. For example, at B0 there are 8 compressors; four belong to the "B0" group and the other four belong to the "BA" group. Altogether there are 32 compressors; however, in Collider Mode one compressor at B0 is dedicated to the cryogenic solenoid at CDF and is unavailable as a ring compressor. All of the helium compressors are in parallel; that is, they all draw low pressure helium from the same set of suction headers, and, after pressurizing it, send it to a common 3" header.

By the time the helium has arrived in the suction header, it has completed its circuit of the expansion engines and magnet strings, and dropped to a pressure of about 1.2 psig.

The 1.2 psig helium from the suction header is pulled in to the compressors and leaves at a discharge pressure in the neighborhood of 280 psig. The high pressure gas enters the 3" header, which runs parallel to the transfer line on the Tevatron berm. The expansion engines in the refrigerator buildings use the compressed helium. It is important to note that even the ideal suction and discharge pressures are expected to vary somewhat from point to point around the ring. For example, there is a suction profile that requires different sections of the suction header to be at slightly different pressures.

There are valves in the B0 Compressor Room for shunting excess helium in the discharge line back to CHL, more about those later.

A few parameter names should be mentioned at this point. Suction pressure, which is measured at each zero building, takes the name T:xxPI1, where xx represents the location (i.e. B0PI1). Discharge pressure takes the name T:xxPI2.

Digital status for the individual compressors is organized under the parameter type T:xxHPyy, where yy is a number related to the compressor number at a given house. For some totally inexplicable reason, Compressor #1 is given the name T:xxHP30, #2 is T:xxHP40, #3 is T:xxHP50, and #4 is T:xxHP60. This strange scheme is also reflected in some of the other compressor parameter names that will be introduced later.

The ring compressors are of a screw-type design, not piston-driven as one might expect. Each stage consists of two large rapidly rotating screws that are roughly parallel to each other. Bridging the two screws is a slider (no relation to White Castle), which acts something like a zipper. If the slider is at one end of the stage, the screws are “unzipped;” they are misaligned enough that the helium gas is not forced through the system. If the slider is at the other end, the two screws mesh and the helium, trapped inside the spaces, is compressed as it passes through. The slider can also be at any intermediate position, is which case only some of the helium is compressed. The degree to which the slider causes the screws to mesh is called loading, and a compressor, which is processing the maximum possible quantity of helium, is said to be fully loaded.

At each individual compressor, the helium is compressed in two stages in order to keep pressure differentials at a minimum. The first (or “low”) stage takes helium from the suction header; the helium typically starts between 1 and 2 psig and is compressed to about 20 psig. The second (or “high”) stage takes the 20 psig helium and compresses it to the discharge header value of 280 to 300 psig.

The parameter name of the low stage slider for Compressor #1 at a given location is T:xxHV1L; the high stage slider is called T:xxHV1H. Compressor #2 parameters would be T:xxHV2L and T:xxHV2H, etc.

Regulation loops (Fig. 5.3) control the amount of loading for each stage. The low stage regulates from PI2—the discharge pressure—and in doing so determines how much helium is pulled from suction. For example, if the loop senses that the discharge pressure is too low, it sees to it that more helium is taken from suction; a greater quantity ends up in the discharge line and the pressure in the line increases. The set point for the loop is usually in the high 200’s (psig), but varies a little from compressor to compressor in order to keep the discharge profile healthy and stable.

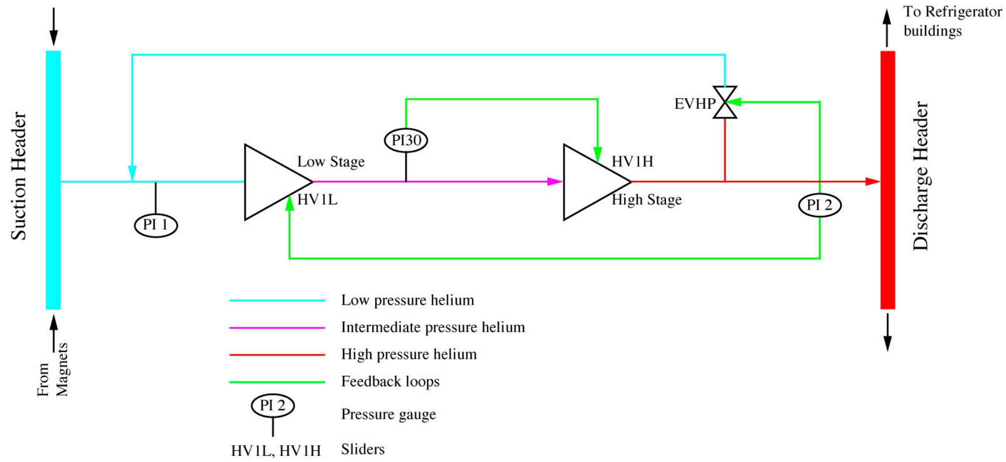

Fig 5.3 Compressor Loops

The high stage regulates from the interstage pressure. (The interstage pressure reading for Compressor #1 at a given location takes the name T:xxPI30.) A common set point for the interstage pressure is 20 psig. Remember that the low stage has already made up its mind as to how much helium is going to be compressed. The high stage takes that gas and compresses it. The harder the high stage works, the lower the interstage pressure. If the interstage pressure drops below the set point, the high stage senses that it is working too hard, and the slider "unmeshes" the screws a little. If the interstage pressure is too high, the slider moves to mesh the screws more tightly so that more helium is removed from the interstage area and compressed.

Compressors are somewhat messy machines—the screws have to be heavily lubricated with oil, and there is no way to prevent the oil from mixing with the helium. After compression, but before actually being allowed into the discharge line, the helium passes through separators and charcoal filters in order to remove the oil. Helium leaving the compressor is also quite hot, so it is also heat exchanged with chilled water.

T:xxEVHP, the compressor bypass valve, is common to all four compressors at a given location. It opens when the discharge pressure exceeds the set point; obviously, the set point is normally made higher than the output requested from any of the individual compressors. The excess high pressure gas is shunted back to the suction header.

Finally, there are two specialized valves at B0 and BA called the kickback valves. Remember that CHL is constantly sending liquid helium to the magnets via the transfer line. As the helium works its way through the refrigerator buildings and the magnets, it warms and vaporizes, and enters the suction lines as gas. To keep the Tevatron from exploding (just kidding, of course), kickback valves return an equivalent amount of helium to CHL from the discharge line; the suction pressure regulates the valves. These two valves will be examined in their full context later in the chapter when helium inventory is discussed.

Computer control of the compressors is an integral part of the refrigerator controls system, and will be addressed as part of that section.

From each compressor building, the discharge line connects to the 3" header located on top of the Tevatron berm. The header passes through each of the 24 refrigerator buildings, where the helium is withdrawn by the local cryogenic systems.

Because high pressure helium can present a hazard to personnel—as a mechanical energy hazard as well as from an ODH standpoint—the first valve encountered in the building is SV101, the Emergency Shutoff Valve. The valve and its motor control hang from the ceiling. Pushing the crash button on the outside wall of the building closes the valve. This action shuts down much of the equipment in the building and should only be taken judiciously.

Once past the SV101 valve, the next stop for the helium is the heat exchanger.

Heat Exchangers

The purpose of the heat exchangers is to cool the high pressure gas as much as possible before it is sent to the wet engine. The exchangers are the long tube-shaped structures attached to the sides of the refrigerator buildings. Internally, each heat exchanger actually consists of four interconnected exchangers.

As it enters the exchanger from inside the building, the high pressure line temporarily divides into two branches (Fig. 5.4). One branch is cooled by heat exchange with liquid nitrogen supplied by CHL, and the other branch is cooled by heat exchange with the cold low-pressure helium (alias "suction") returning from the magnets in the tunnel. The loops for both exchangers look at TI5, a nitrogen VPT. TI5 measures the temperature of the low-pressure helium at the point where the two branches combine again. The strategy: If the returning helium is cold enough, use it, if it isn't, use the liquid nitrogen, and if the gas is somewhere in between, use both.

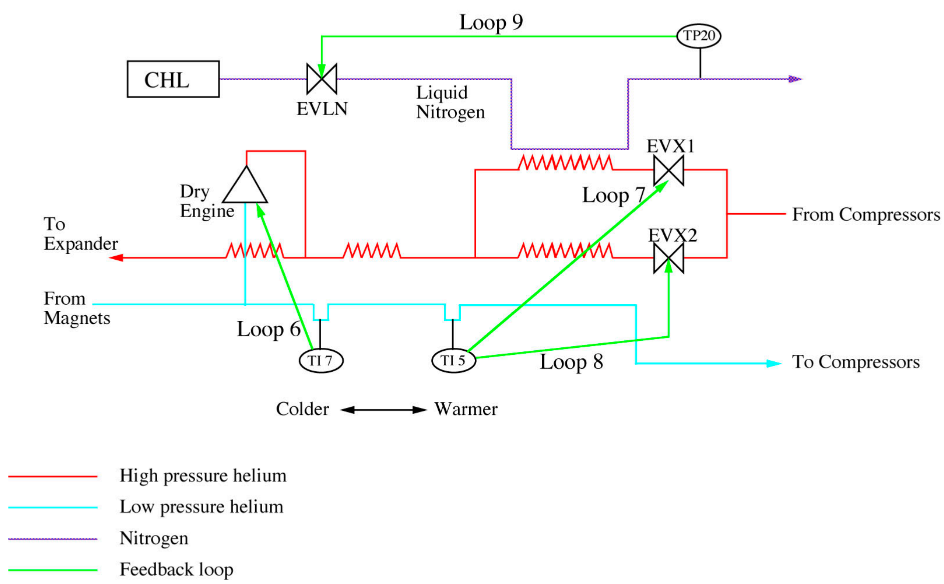

Fig. 5.4 Heat Exchanger Loops

The nitrogen exchanger is controlled by two loops—one that regulates the flow of nitrogen into the exchanger (Loop 9, EVLN), and one that controls the amount of high-pressure helium going through the nitrogen exchanger (Loop 7, EVX1). EVX1 adjusts the flow of helium through the exchanger to move TI5 toward its set point of 25 psig, which corresponds to a temperature of about 87K. Remember that the temperature of liquid nitrogen is about 77K—the valve opens up if TI5 is too warm, and closes down if it is too cold.

EVLN regulates to a platinum resistor called TP20, located in the nitrogen line downstream (nitrogen direction) of the exchanger. TP20 is assigned a set point near 200K. The purpose of EVLN's loop is to conserve liquid nitrogen; if the nitrogen is colder than 200K after heat exchanging with the helium, some of it is probably being wasted and the valve closes some.

Valve EVX2 (Loop 8) controls the flow of high pressure helium through the branch that exchanges with the cold helium from the magnets. Its loop is also regulated around TI5, but the feedback is opposite that of EVX1—if TI5 is too warm, EVX2 closes down to keep from warming the incoming helium. If TI5 is too cold, EVX2 opens in order to fully exploit the opportunity.

In practice, the loops are changed only during cooldown. During steady-state operation, returning helium is usually cold enough to carry the full load. EVX1 is locked nearly closed, and EVX2 is locked completely open. There is no need to conserve the low temperature of the returning helium, because this is the last point at which it will be used for heat exchange.

The heat exchanger is also equipped with an expansion engine called the dry engine. The dry engine, except for a smaller piston, is virtually indistinguishable mechanically from the wet engine. The only real difference between them is their location in the cooling sequence. The dry engine expands gas that is still too warm to liquefy.

The dry engine is seldom used any more, but when it is, it pulls high pressure helium out of the line, expands it, and dumps it back into the suction line. The relatively cold gas then joins the helium from the magnets and helps in the heat exchange. This cooling technique is normally used when cold helium from CHL is not available.

The loop (Loop 6) for the dry engine regulates to TI7, which is a hydrogen VPT. Notice in Fig. 5.4 that TI7 is downstream (cold helium direction) from the dry engine. The set point for TI7 is 10 psig, which corresponds to a temperature of about 22K. The harder the dry engine works, the colder TI7 becomes. More detail about expansion engines is forthcoming in the next section.

The dry engines are not normally used because cold helium is available from CHL. The loops regulating the CHL helium are best deferred until later, when the refrigeration system as a whole has been more fully explored.

Expansion

After the heat exchanger has cooled the high pressure gas as much as possible, it is time to call on the next technique—expansion (see Fig. 5.5). In a healthy, steady-state operation, the temperature of the gas just before it enters the expanders is about 6K. Expansion will reduce the temperature to a point at which the helium begins to condense into a liquid.

There are two devices, in parallel, which can implement the expansion. One is the wet engine. As an engine, the cold gas pushes against a piston as it expands. It is doing work, exactly as steam once did in locomotive engines. And, like the steam, energy is transferred from the gas to the piston, and the gas condenses. In fact, the activity of an expansion engine is monitored by attaching a small electric generator and reading its output. The helium gas is generating power, but only enough to rob it of its energy.

The engine expels a certain quantity of helium with every stroke, so the speed of the engine, in RPM, is the measure of its output. The wet engine normally runs between 300 and 1200 RPM.

The feedback loop (Loop 2) regulating the speed of the engine uses a pressure gauge downstream of the output; the faster the engine runs, the greater the quantity of helium pumped, and the higher the pressure of the output. There are several pressure gauges downstream of the engine, which can be used, but the preferred one is PI13. If PI13 is malfunctioning, PI14 or PI17 can be used. Since the latter two gauges are downstream of the feed can, the set point should account for the pound or so of pressure drop the helium experiences going through the feed can. The set point of PI13 during steady-state operation is 17 psig.

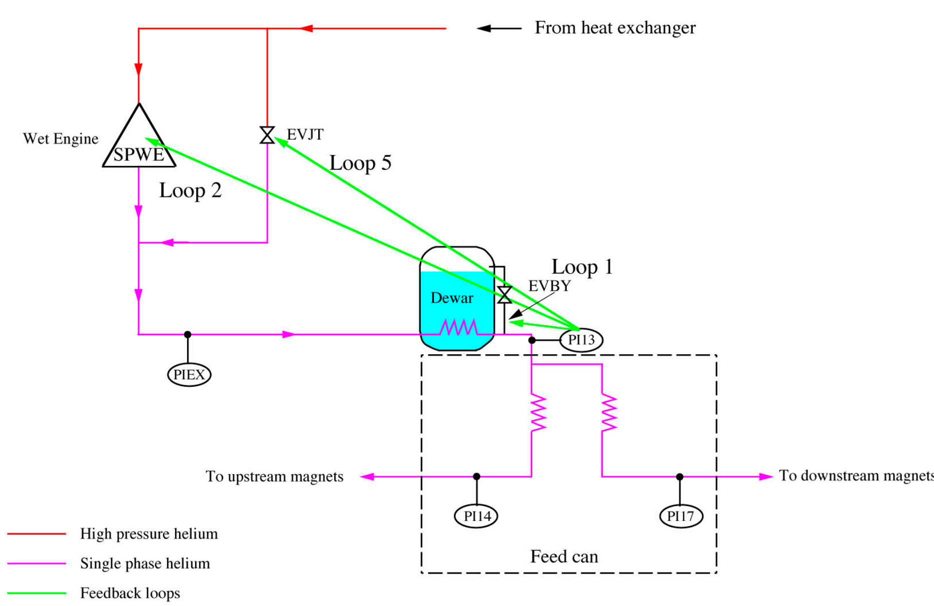

Fig. 5.5 Pressure Regulation Loops

The other expansion device is a Joule-Thompson (JT) valve, called EVJT. A JT valve forces the gas through a narrow aperture, after which it is allowed to expand. There is no mechanism for removing energy from the gas, so the cooling is less effective.

The wet engine is always the preferred method for liquefying the gas; the JT valve is strictly used as a backup. Not only is the quality of cooling inferior with the JT valve, but it also requires a greater quantity of compressed gas to achieve the same effect. The valve, like the wet engine, regulates to PI13 (Loop 5), but the set point is 12 or 15 psig; in the event of an unexpected failure of the wet engine, the JT valve acts as a safety net. The set point of the JT can be raised if wet engine recovery is not imminent.

There is one more pressure-related valve—EVBY, the bypass valve—which is used during cooldowns but not usually during steady-state operations. It also looks at PI13 (Loop 1), but is set to open if the pressure exceeds the set point of 18 psig. It acts as a safety valve to prevent over pressurizing the magnet strings, but normally the wet engine will slow down before the pressure gets that high. During cooldowns, EVBY is used to siphon off cold helium to the Dewar and heat exchanger – more about that later.

There are two additional opportunities to further cool the helium before it reaches the magnets. One is the Dewar, which during steady-state operation is full of liquid helium. The other is in the feed can, where the incoming helium is heat exchanged with the 2-phase helium returning from the magnets. The 2-phase is actually colder than the single phase.

Magnet Strings

The single-phase helium from the feed can splits to supply the upstream and downstream magnet strings. Inside each magnet is the cryostat, which is a long tube containing the beam pipe, the magnet coil, the helium and nitrogen "circuits," and the insulating vacuum. The cryostat is surrounded by the iron yoke. In most of the Tevatron magnets, the iron yoke is at room temperature; however, the low-beta quads are of a "cold-iron" design.

The single-phase helium enters the magnets and surrounds the coil and the stainless steel collar surrounding the coil. (The entire cryogenic system was built for this one moment.) The single-phase helium slowly moves through the magnet string until it reaches the turnaround box at the end. During normal operations, the helium passes through the magnet JT valve at the turnaround box. The liquid helium, after passing through the narrow aperture, expands and cools. It is then sent back into the magnets via the 2-phase circuit, which is concentrically to the outside of the single phase circuit. Because of the JT effect, it is a few tenths of a degree colder than the single phase helium; heat exchange takes place throughout the magnet string as it returns to the feed can.

The 2-phase helium consists of liquid and gaseous helium (hence the name). Immediately after passing through the JT valve, it has a relatively high ratio of liquid to gas. As it moves back along the 2-phase circuit, it absorbs heat from the liquid helium on the inside and from the environment on the outside. As with boiling water, it remains at the same temperature even as it is being converted to gas. It is preferred that all of the liquid—but no more—boils off by the time that the end of the string is reached. In that way the temperature remains cold, but no liquid helium is wasted.

The feedback loop controlling the JT valve uses a hybrid device called a DT (for differential temperature). The DT measures the difference between the pressure at the end of the 2-phase circuit and the temperature as measured by a VPT (DT = VPT - PI). Pressure is constant throughout the 2-phase circuit, so as long as there is a liquid/gas mixture the VPT should remain constant. If all of the liquid boils off, the temperature will begin to rise as it absorbs energy from its surroundings; in that case, the DT rises. It is even possible to “freeze out” a DT by producing a liquid below the boiling point.

The PI for the DT measurement is usually PI11, upstairs in the refrigerator building. There is a slight pressure drop going from the 2-phase circuit to PI11, so the DT usually has a set value of about 0.2 psid (d stands for differential).

Establishing the operating point for a JT valve is a fine and delicate art. Often, the value is between 50% and 60%, but it varies from string to string. The effectiveness of the JT depends on the balance between the amount of flow and the amount the valve is constricted, and that balance depends in turn on the temperature. Specifics will be discussed in the sections on cooling the strings.

The feedback loops (Loop 3 upstream, Loop 4 downstream) works by opening up the JT (increasing flow) if the DT rises above its set point, and decreasing flow by closing the JT when the DT shows that the helium is too cold. The limits are usually set to restrict changes to a few percent.

Nitrogen

Surrounding the helium layers is a vacuum break. Outside the vacuum layer is the liquid nitrogen circuit. The nitrogen is a free gift from CHL. At roughly 77K, it is an insulating layer between the helium and the balmy climate outside the magnet.

The nitrogen makes a single pass through the magnets and empties into a common nitrogen header behind the magnets. A pipe between A3 and A4 connects the header to CHL, where the nitrogen is re-liquefied.

The valves controlling the amount of nitrogen flowing through the magnets are EVUN (Loop 11, upstream string) and EVDN (Loop 12, downstream). They regulate to the temperatures measured by two platinum resistors, TP23 (upstream) and TP24 (downstream). The set point for each is 90K, slightly above liquid nitrogen temperature. The goal is to have converted all of the liquid nitrogen to a gas by the time it leaves the string, while still keeping the magnets cold.

Cryostat Vacuum

The single phase and 2-phase circuits are adjacent to each other and within a degree or so of the same temperature, but the nitrogen circuit is much warmer. A "layer" of vacuum separates the two layers of cryogens. The same vacuum extends to the outside of the nitrogen shields. This vacuum is completely distinct from the vacuum of the beam tube. It is also of a lower quality; beam tube vacuum is in the range of 10-10 torr, while the cryostat vacuum is closer to 10-7 or 10-8 torr. But it is sufficient to maintain temperatures in a cryogenically healthy magnet.

Within the vacuum layer are several layers of reflective Mylar insulation, which prevent infrared light from penetrating to the cryogenic levels. These elaborate layers of insulation work—the surface of the magnets is at room temperature, leaving no clue as to the near absolute-zero temperatures a few inches away.

Chapter 6 includes more detail on the cryostat vacuum.

Kautzky Valves

A quench, which vaporizes a great deal of helium, carries the risk of over pressurizing the magnets. Kautzky valves, in one of their roles, act as pressure relief valves.

Kautzky valves are of a very simple design—a plastic "poppet" is held in place by a control pressure of 30 psig, supplied from helium bottles upstairs in the service building. (Helium, as opposed to air, is used in order to ward off contamination in the magnet string.) One side of the poppet faces the suction header, and the other side is facing one of the cryogenic circuits—single phase, 2-phase, or nitrogen. If, during a quench, the pressure inside the magnets exceeds 30 psig, the poppet opens and allows the helium to enter the suction line.

There is one Kautzky valve for each of the dipole magnets dedicated to the single-phase circuit, since that is the circuit with the greatest potential for expansion of gas. At the spool piece, there is one Kautzky valve dedicated to each of the three circuits.

The single-phase Kautzky at the spool pieces can be remotely controlled, and during cooldowns they are sometimes opened and closed in a systematic fashion in order to expel warm helium from the magnet strings—more about that later.

The Dewar

The Dewars in the refrigerator buildings were installed as part of the upgrade to high energy (1 TeV) operations. They are filled with liquid helium; during normal operations, the helium is used not only to heat exchange with the exhaust of the wet engines, but also to provide a reservoir of cold helium for the shell side of the heat exchangers. But the real reason for the existence of the Dewar is to compensate for the heat load generated by the cold compressor, described below, during low-temperature operations.

The Dewar obtains helium from several different sources, including CHL, 2-phase returning from the magnets, and, during cooldowns, single phase helium from the wet engine exhaust via EVBY.

CHL

Cold helium arrives at the refrigerator building from the transfer line; the gatekeeper at the building is EVSH, a solenoid-operated valve. It is either open or closed—almost always open when the refrigerator is operational, and closed only during maintenance periods or when CHL is not producing liquid. There are no feedback loops controlling EVSH.

Downstream of EVSH are two regulated valves controlling the flow of CHL helium: EVLH and EVQH.

EVLH (Loop 10) adds liquid directly to the Dewar, and regulates to the amount of helium in the Dewar. The liquid level probe LL11, which expresses itself in units of "% full," often measures the quantity of helium. Sometimes it is measured by another type of probe, DP11, which measures the level in the less intuitive unit of "inches of water;" 4.6 inches represents a Dewar that is 100% full. The Dewar should be about 60% full, so EVLH is opened if the liquid level drops below that level.

EVQH (Loop 15) adds liquid helium to the exhaust of the wet engine (or JT valve) before it is heat exchanged with the helium in the Dewar. As such, it does not directly contribute to the Dewar level. But indirectly it does, because it means that the 2-phase helium will be colder (see below). EVQH is normally used only during quench recovery. In the original design, the feedback loop for EVQH looked at TI7, which measures temperature at one point inside the heat exchanger. In practice, it now often uses the liquid level of the Dewar, with the set point at 50% full (10% lower than the set point for EVLH, so they don't fight each other.)

2-Phase Helium

After the 2-phase helium has been used for heat exchange in the magnets and feed can, it enters the Dewar. The liquid helium settles at the bottom of the Dewar and the cold gas moves on to the cold compressor and the heat exchanger. It is important that no liquid helium gets inside the cold compressor, as explained in the next section.

Cold Compressors

During “normal” operations, the cold helium liquid in the Dewar evaporates and becomes the cold helium gas used as counterflow in the heat exchanger. During low-temperature operations, the helium from the Dewar first passes through the cold compressor, and then continues on through the heat exchanger. The cold compressors should not be confused in any way with the ring compressors!

Since the helium at this point is already quite cold, and since it has already been through the magnets, it may seem counterintuitive to put a compressor into the system knowing that it will heat the gas about to enter the heat exchanger. Indeed, nearly half of the heat load of the system comes from this compressor, and compensating for that heat presents a new engineering challenge. The advantage actually comes from lowering the pressure of the input of the compressor, that is, the 2-phase helium. Since the 2-phase circuit is continuous all the way back to the magnet JT valves, lowering the pressure lowers the temperature throughout the system. The drop of about 0.8K is enough to allow the magnets to sustain currents and fields equivalent to 1 TeV.

The compressors are of a turbine design, with the blades spinning at a minimum of 20,000 RPM (20 KRPM) and a maximum of 95 KRPM. The faster they spin, the lower the input pressure, and the greater the heat load on the system. (At those speeds, a tiny droplet of liquid could create enough centrifugal force to destroy the compressor.) Although most of the liquid of the 2-phase condenses in the Dewar, there are heaters at the input of the compressor to take care of any leftover mist.

The compressor speed regulates to PI11 (Loop 16), a pressure gauge located at the point where helium is taken from the compressor. The set point, when the compressors are fully operational, may be around 11 psia. Notice that the units are in absolute pressure and that the operating point is sub-atmospheric. All of the 2-phase circuits between the magnet JT’s and the cold compressor had to be “hardened,” to prevent leaks, before adding the cold compressors to the system.

The lower pressure in the 2-phase circuit creates a larger pressure differential with respect to the single phase, which increases the flow across the magnet JT valves. The valves are closed down by a few percent in order to reduce the flow to normal levels. The increased "JT effect" also helps to lower the temperature.

There is a valve, PVCB, which bypasses gas around the compressor. (The "P" means that the valve is pneumatically driven.) If left in remote, the valve closes automatically when the compressor turns on.

It should now be possible to explain why the Dewar is a necessary adjunct to the cold compressor. The exhaust from the compressor is warmer than the gas that would otherwise be returning through the heat exchanger. Incoming warm gas from the exchanger works its way through the wet engine, producing a lower quality of cooling than would otherwise be the case. The Dewar is an opportunity for the wet engine exhaust to heat exchange with a colder substance before it enters the magnets.

The cold compressors are not turned on until the refrigerator has already been cooled to "normal" cryogenic temperatures.

* * * * *

Finally, the exhaust, whether or not it has been through the cold compressor, passes through the heat exchanger and absorbs all of the heat it can from the incoming high pressure helium. (In an ironic twist of fate, a few hours earlier it was the incoming helium.) Now, at roughly room temperature, it enters the ring wide suction header and is taken in by the ring compressors, beginning the cycle again.

Power Leads

Another cryogenic component is required before the Tevatron is allowed to ramp—power lead cooling. The power leads are where the copper bus from the power supplies meets the superconducting alloy in the feed can. The copper, of course, is not superconducting, but it is required to carry several thousand amps to the feed can—there is considerable resistive heating at the leads. Leads at the "2" and "3" locations are associated with the feed cans, because there are power supplies at those buildings. At the "1" buildings they are at the upstream end, because they link to the copper bus of the straight-section bypass; at the "4" buildings they are at the downstream end for the same reason. Helium is tapped off of the single-phase circuit and passes through the superconducting part of the leads in order to cool them; the copper is cooled with LCW.

Current in the Tevatron, and therefore the cooling requirement, varies with operational demand. Sometimes the Tevatron is off, sometimes it is ramping, and sometimes it is in a 150 GeV store and sometimes a 1 TeV store. Cooling which would be appropriate at 1 TeV would freeze the water going through the leads during a 150 GeV store.

Control of the flow is implemented via three orifices upstairs in the refrigerator building, two of which are equipped with solenoids. The two with solenoid control can be interchanged, if necessary, with orifices that have different diameters and flow rates; the remaining orifice is fixed.

Helium flows through the fixed orifice at 15 SCFH (Standard Cubic Feet per Hour). It is always open.

The second solenoid-controlled orifice, carries either 40 or 60 SCFH of helium. Since the solenoid is open whenever the Tevatron is above 70 GeV or so, it is known as the low-energy solenoid.

The third orifice, with a flow of either 60 or 80 SCFH, has a solenoid that is opened whenever the TeV current is above 150 GeV. It is known as the high-energy solenoid. Actually, all three branches have to be open when the Tevatron is at higher energies.

A platinum resistor measures the temperature of the helium downstream of each lead. (There are “upstream” and “downstream” leads, although the terminology is strictly true only for the “2” and “3” houses; both power leads for the “1” houses are at the upstream end and both leads for the “4” houses are at the downstream end.) The parameter for the upstream lead temperature is TPUL, and that of the downstream lead is TPDL.

Using a slightly confusing convention, digital control for the solenoids uses TPUL and TPDL for parameter names, but assigns the low-energy solenoid to TPDL and the high-energy solenoid to TPUL. In reality, “low-energy” and “high-energy” have nothing to do with “upstream” and “downstream” when it comes to digital control of the leads.

Control of the lead flows is through parameter names (they are listed on page F14, subpages 33 and 34), the FSM, or the Sequencer.

The FSM automatically turns on the high-energy flows whenever the Tevatron energy exceeds 160 GeV (i.e., the Tevatron starts ramping). (It is preferred, however, not to risk a ramp dump by forgetting to turn them on manually.) They will also turn off the high-energy flows if the Tevatron has not been above 160 GeV for 5 minutes, in order to prevent the leads from freezing. (There are fans downstairs blowing on the leads to minimize the risk of freezing, but again, it is preferable not to take chances.)

Finally, the temperature readbacks for TPUL and TPDL are monitored by the FSM to make sure that the leads are not getting too hot. The limit is normally 297K. If the limit is exceeded, the FSM removes the ramp permit. (The ramp permit looks at several different measurements and is described in a section below.)

Safety Leads

At every other spool piece, safety leads are attached to the bus. The safety leads connect the bus to the QBS switches, and carry the current away from a quenched cell. If a particular cell quenches repeatedly, the leads could overheat. Cold helium is pulled from the magnets to prevent that from happening. Safety lead temperatures are not monitored from the Control Room, but there are flow meters in the tunnel that must be checked if the quenches become too frequent.

The Cooldown Sequence

This section will describe the cooldown sequence beginning at liquid nitrogen temperature, which is about 77K. Cooling from room temperature to nitrogen temperature is the task of specialists for whom the information in this chapter would be superfluous.

If the magnet string is sufficiently warm, four modes of cooling are required—cooldown, transition, fill, and operate. The control of the sequence can be done manually through the feedback loops, or by the Finite State Machine (FSM). Understanding the feedback loops is an important prerequisite for understanding what the FSM does; the emphasis will be on the loops for this discussion.

Cooldown

This stage (which has admittedly been given a totally ambiguous name—isn't it all cooldown?) refers to the first surge of cold helium passing through the magnet string. The magnets are too warm at this point to allow any of the helium to remain as a liquid; also, the single-phase circuit itself is likely to be full of warm helium. Cold helium enters from the feed can and makes one pass through the single phase circuit, displacing the warm gas. The helium being displaced is too warm for the magnet JT valve to be of any use; in fact, above 35K or so, JT expansion actually heats the gas. Instead, displaced helium is dumped directly into the suction header via the cooldown valves at the end of the string. If the refrigerator is strong enough to sustain the load, Kautzky valves can be used to speed up the process.

At first glance, the feedback loops for the cooldown valves (Loops 13 and 14) seem rather counterintuitive (i.e., they don't make sense). They look at TR12, a carbon resistor at the inlet of the wet engine. Because of its location, TR12 is a good indicator of heat exchanger performance.

What is keeping the heat exchanger cold? Remember that during normal operations, the counterflow of cold gas comes from returning 2-phase helium. However, if the magnet JT valves are closed and the cooldown valves are open, there will be no returning cold helium. This problem is solved by changing the set point of the wet engine to 20 psig, while leaving the set point of the bypass valve at 18 psig. The wet engine will speed up to its maximum of 1200 RPM as it tries to raise the pressure at PI13 to 20 psig, but the bypass valve will open at 18 psig and shunt the extra helium into the Dewar. That cold helium is what finds its way into the heat exchanger.

If the cooldown valves are open too far, the pressure will drop throughout the single phase circuit, and the bypass valve won't shunt enough cold helium into the Dewar. The heat exchanger, and therefore TR12, will consequently warm up. To preserve the quality of the cooldown wave, the feedback loops will close down on the cooldown valves to raise the pressure in the circuit.

The set point of the loop during this stage is usually set at either 10K or 20K, and produces a "10° wave" or a "20° wave." The 10° wave is slower, because the valves are not as far open, but it produces a better quality of cooling for this stage.

Often, the Kautzky valves are opened during cooldown mode. Any Kautzky downstream of the wave can be used. Since the Kautzky were originally designed as relief valves, the volume of helium purged from the string is enormous compared to the cooldown valves. They are only opened for a few seconds at a time; TR12 rises very rapidly during that time. They are then closed until TR12 indicates that it is safe to proceed. There is no loop control for the Kautzky valves; they are controlled either manually or through the Finite State Machine.

EVX1, which allows incoming helium to be exchanged with liquid nitrogen, is opened during cooldown mode to assist the heat exchanger. (It is normally closed to a trickle during steady-state operations.)

EVQH, usually closed during normal operation, is opened, and regulates to the liquid level of the Dewar. During cooldown phase the Dewar is not likely to fill, so EVQH remains open at its maximum value. Officially, EVLH (normally open) is closed, although in practice it is sometimes opened if needed. (EVLH would try to fill the Dewar directly, but it is not necessary for the Dewar to have liquid during this phase. It can, however, keep the heat exchanger colder.)

Transition

The cooldown phase is terminated when the cooldown wave reaches the end of the string. At that time, the magnet JT valves are opened to 100% and the cooldown valves are closed. It is important to properly assess when and where that happens—if the change is made too soon, helium going through the JT valve will be too warm and the counterflow will actually heat the single phase instead of cooling it. If it is done too late, the cold helium will be dumped into the suction header and wasted; also, since the cold gas has a higher density, flow through the system will increase, helium will leave faster, and less will be available to cool the heat exchanger.

The switch to transition mode should be made when the end of the string reaches 9K or so; if the temperature is being plotted, the change can be delayed until the drop in temperature begins to level off. The cooldown and magnet JT valves are located in the turnaround box at the end of the string, so it is important to look at a carbon resistor in or near the turnaround box (Fig. 5.2). The upstream measurement at “2,” “3,” and “4” buildings is read from TRQ1. The “1” houses must use TR15 for the upstream measurement because the lattice requires additional quadrupoles upstream of TRQ1. At B1 and D1 it is the low beta quads, and at A1, C1, E1, and F1 it is the 82” and 99” lattice-matching quads that separate TR15 and TRQ1.

Downstream, the “2,” “3,” and “4” houses use TRQ9 as the end of the string, except for A4 and C4, which use TR18 because of the low beta quads. The downstream “1” strings extend to the “21” location, so they use TRQ0.

The reason that TR15 and TR18 are not always used is that they can be susceptible to “cross-talk” from adjacent turnaround boxes. For example, if the downstream F2 string is warm, F3TR15 may read artificially high.

The magnet JT is locked open at 100% during transition because at the temperatures typical of this stage the JT effect is relatively weak, and flow is favored over restriction of the valve. The flow/restriction ratio will change as cooling progresses, but at this point establishing a reasonably cold counterflow in the 2-phase circuit, and lots of it, is the most important goal.

Fill Mode

To “fill” the magnets is to coax the helium to condense into a liquid. The change from transition to fill mode is implemented when the end of the magnet string reaches 5.3K; the magnet JT is changed from 100% open to 80% open, since the helium is now cold enough to take greater advantage of the JT effect. Also, the set point for the wet engine is returned to its operational value of 17 psig; the counterflow through the 2-phase circuit is sufficiently cold to allow the bypass valve to be closed. The reduced pressure in the circuits will also aid in the cooling, although the engine will continue to run fairly fast as long as the JT is open at 80%.

Initially, the DT will be high—around 6 or 8 psid. After what may seem like an eternity, liquid will begin to condense in the 2-phase circuit and the DT will begin to come down. The Dewar will begin to collect some liquid helium as well. The magnet JT valves should be closed down very gradually during the latter stages of fill mode.

Operate mode

At operating temperatures, the helium is cold enough to take full advantage of the JT effect. EVLH is opened to maintain liquid level in the Dewar, and EVQH is closed. EVX1 is reduced to a trickle, nitrogen no longer being needed by the heat exchanger. All of the loops can be restored to their operational steady-state values.

Low Temperature Operations

If high-energy beam is desired—which should be the case throughout the Collider run—the cold compressors must be used. The cold compressor is left off during most of the cooldown, because of the strain it puts on the heat exchanger. Only after a "normal" permit is established is the cold compressor turned on.

The purpose of the cold compressor is to lower the 2-phase pressure, as measured by PI11. The faster the compressor spins, the lower PI11 will be. The lower pressure makes the magnets colder, but there are two negative consequences: (1) the return gas flowing through the heat exchanger is warmer, and (2) there is greater flow through the magnet JT's, because of the increased differential pressure across the valves. For those reasons, the cold compressor is initially turned on to a speed of 20 KRPM; the set point of PI11 starts out relatively high. The cold compressor speeds up as the set point of PI11 is gradually lowered; at the same time, the JT valves are slowly closed down in order to reduce the flow. The 2-phase gradually becomes colder as the heat load from the compressor increases.

The sequence of cooling from "normal" permit temperatures to high-energy temperatures is coordinated through a Finite State Machine; it requires about half an hour to execute.

Refrigerator Ramp Permits

One of the tasks of the local controls system—specifically, the FSM—is to determine whether or not the system is healthy enough to allow the Tevatron to ramp. If it is, it issues a permit to the local QPM. The QPMs report to TECAR; there must be permits from all 24 refrigerator buildings before the Tevatron is allowed to ramp.

The permit is implemented via a ramp inhibit, not a ramp dump. For example, if a refrigerator parameter drifts out of tolerance at flattop, the Tevatron finishes its current ramp normally but locks at 90 GeV when it reaches that level. In Collider mode the refrigerator permit is completely ignored once a store is established. Given the investment required to set up collisions, it is better to risk a quench than to guarantee the death of a store.

There are two types of permits—a normal permit and a low-temperature permit. The normal ramp permit is required for any current up to 900 GeV. The low-temperature ramp permit, which includes everything required of the normal permit (and more), is necessary for operations above 900 GeV.

The ramp permits should not be confused with the refrigerator alarms on AEOLUS, which may or may not have the same limits for a given device.

The normal permit looks at parameters involving pressure, temperature, and DTs:

- Pressure: PI13, PI14, and PI17 (helium pressure upstream of the feed can, upstream string, and downstream string, respectively) are monitored to determine if they are within minimum and maximum limits. The FSM customizes the limits for each device, but a typical minimum value is 14 psig (remember that during normal operations, the wet engine set point is 17 psig). Common reasons for low pressure might be (1) the magnet JT valves are not close enough to their minimum value following a cooldown, (2) there is a leaky Kautzky valve, or (3) there is a problem with the wet engine or pressure transducer.
- Magnet temperature: Most of the carbon resistors, including those in the feed can and all of the TRQs, are monitored for temperature violations. A typical limit is 5K, but resistors in the feed can and turnaround boxes are allowed more leeway. If the cooldown sequence is complete and the temperatures are close to the limit, and the permit is still missing, opening the magnet JT slightly or turning on the low-energy lead flows might bring it back.
- Lead flow temperature: The limit for the lead flow temperatures is normally 297K before the ramp permit is removed. (This temperature permit, like the other permits listed here, is a subset of the refrigerator permit sent to the QPM—it should not be confused with the voltage monitoring of power leads done by the QPM itself. The penalty for a voltage violation is a ramp trip, not the ramp inhibit of the refrigerator permit.)

- DTs: Violation of a DT limit does not immediately revoke the ramp permit, because it usually takes a few minutes for the lack of subcooling to impact the temperature inside the magnets. Instead, when a DT rises above its set point, a "counter" is initiated; if the DT is slightly above the set point the numbers accumulate slowly, but if the DT is significantly high they accumulate quickly. When the integrated value reaches a value 1300, the permit is removed. The number is reset to zero as soon as the DT drops below the set point. One technique for deferring a ramp inhibit—to be applied judiciously—is to momentarily put the set point at an unrealistically high value in order to reset the counter.
- The low-temperature permit includes all of the requirements of the normal permit, and adds one. PI11, which measures the inlet pressure of the cold compressor, must be below the set value. PI11 is an indicator of pressure throughout the 2-phase circuit, and the magnet temperatures rise and fall with that pressure. The value of the set point is chosen to produce the desired temperature.

Refrigerator Controls

Control of the Tevatron cryogenic systems is implemented through VME crates distributed around the ring. There is one VME crate per ring sector, located at the zero buildings.

Each VME uses an ARCNET loop to control its little empire. (A strand of the 19-conductor cable was cut into six sections to create the loops.)

Locally, at the buildings, the commands and readbacks (i.e., I/O) are interfaced to the hardware through Intel 186 processors. There are actually two systems at the "1" through "4" houses: the device I/O crate and the thermometry I/O crate; each system has its own 186 card. A device crate controls the compressors, but they do not need a thermometry crate because there are no low temperatures there to measure.

The device I/O crates are located in the refrigerator buildings; the crate interfaces with the valve actuators, engine controllers, and the cold compressor. There are also digital input cards and A/D converters for monitoring the instrumentation (e.g. the pressure transducers), and a vacuum card. (The vacuum card monitors conditions in the equipment upstairs, such as the valve box and U-tubes. It should not be confused with the CIA crate in the service building, which spies on the vacuum in the beam pipe and magnet cryostats.)

The thermometry crates are responsible for providing the carbon and platinum resistors with the pulses of current necessary to measure the resistance, and therefore temperature. They are located in the Tevatron service buildings, on the back wall behind the QPM racks. In addition to their role of measuring temperature, the thermometry crates have a link to the QPMs.

All of the software, whether for thermometry or the device crate, gathers data at a 1 Hz rate. Altogether, there are about 700 readbacks from each house.

Loop Control

The reader who has successfully followed the preceding discussion of refrigerator operation already has a good intuitive understanding of how the loop controls work. A few points should be made anyway about F8, the loop control page. There are 16 loops normally used in refrigerator control. Remember that the loops and the Finite State Machines are two distinct entities; failure to recognize this difference is a common source of confusion for newcomers.

Itemizing some of the parameters on page F8:

- Input variable, or the parameter that the loop is trying to maintain. Typical input variables include temperature or pressure readbacks. The purpose of the loop is to maintain the parameter at the set value.
- Output variable, or the device responsible for maintaining the parameter at the set value. The device associated with each loop number is constant for all operational Tevatron refrigerator buildings and might as well be listed: (1) EVBY, (2) SPWE, (3) EVUH, (4) EVDH, (5) EVJT, (6) SPWE, (7) EVX1, (8) EVX2, (9) EVLN, (10) EVLH, (11) EVUN, (12) EVDN, (13) EVUC, (14) EVDC, (15) EVQH, and (16) SPCC. Details on each of these is scattered throughout the text above.
- Minimum and Maximum Positions: ("Position" can also mean things other than valve position, such as engine speed.) When

the loop is enabled and active, the device is free to roam within these limits in order to bring the input variable to its set value. If a position outside of the range is needed, the loop is helpless.

- Current Position: A readback, but it is also possible to type in a desired position.
- Enable/Disable: When disabled, the position is locked in place and will not respond to loop controls. However, it is still possible to type in a desired position, and the Finite State Machine can still change the position.
- Active/Inactive: This is a tricky one—it indirectly reflects what the Finite State Machine is doing. When the loop is "Inactive," the FSM is actually controlling the device, based on algorithms that have nothing to do with what is on the screen. Moreover, one can toggle the bit to "Active," but the FSM continues its work in secret.
- Remote/Local: This is controlled by the switch on the actuator card in the Device I/O crate.
- Sample Time: A measure of how often the loop checks conditions, and therefore how quickly the device changes. There is a wide range of response times. For example, EVBY must respond rapidly to pressure fluctuations, and has a sample time of 6 seconds. EVUN and EVDN, which control flow through the nitrogen shield, respond to slow temperature fluctuations and have a sample time of over 15 minutes.

Finite State Machines

The Finite State Machines (FSM) provides automatic control for various cryogenic systems, including quench recovery and cooldowns. Each house can run up to 32 FSMs; collectively, they can be thought of as a controls system in parallel with the loops. When called upon, the FSMs will inactivate selected loops and take over their functions until the task is completed. (The FSMs can be very efficient, as long as cryogenic conditions are predictable, but—just in case—human oversight should never be neglected.)

The FSMs can perform the following tasks:

- Quench response. Coordinating the quench response is the responsibility of the QPM, but the QPM makes a request that the FSM begin refrigerator recovery. A QPM failure will also trigger a frig recovery, unless the response is disabled (via page F3).

- Automatic cooldowns, which can also be initiated from page F3. A separate FSM is needed to turn on the cold compressors for low temperature operations.
- Control of lead flows; the FSM continuously monitors the lead flows, turning them on when the Tevatron begins to ramp, and turning the appropriate ones off if the TeV is off or set at 150 GeV.
- Monitoring conditions for a ramp permit. For example, if a temperature wanders out of limits, the FSM tells the QPM that the permit is being revoked, and the QPM will tell TECAR to hold the ramp at 90 GeV.

In the case of quench recovery and automatic cooldowns, the FSMs begin by evaluating current conditions—such as temperature, pressure, or liquid level in the Dewar—and then executing a predetermined set of algorithms to achieve the desired result. At each stage, called a state, the FSM monitors a specific set of conditions. When those conditions are met, the FSM advances to the next state (fortunately, there are only a finite number of states). Each state has a "library" of operations, actions, and timers it calls upon to do its job and evaluate its own performance. The FSMs, states, operations, actions, and timers are all explicitly listed on page F13, although it can be difficult to sort them all out from that page. For a more intuitive interpretation, live graphic displays for cooldowns can be launched from page F24, and displays for the refrigerator permits can be launched from F23.

The Consolidator

The "Fridge Consolidator," as it is affectionately known, is housed in a VME crate in the Computer Room. Its purpose is to gather data from around the ring and organize it for display applications. Programs dependent on the Consolidator include F3, F5, F16, and F17, along with any other application that aspires to present a ring wide view of the refrigerators. Failure of the Consolidator results in no data being returned through these programs.

6. Vacuum

The Tevatron vacuum system is divided into 3 types: cold beam tube; warm beam tube; and cryostat insulating vacuum. Each type is discussed below.

Cold Beam Tube Vacuum

The beam tube vacuum is completely separate from the cryostat vacuum. Beam tube vacuum is divided into 24 sections coinciding with the 24 cryogenics loops. Remotely controlled beam valves located at the turnaround boxes can isolated each vacuum section. The beam tube is cold (4.6 K) at all but the 6 long straight sections and each of the “17” and “48” locations.

The beam tube is pumped down initially by portable pumps, which are not a permanent part of the system. Four ion pumps per house (not including the zero houses since these are warm) maintain the vacuum after pump down. These ion pumps are at the “X2”, “X4”, “X6”, and “X8” locations, where X=1, 2, 3, 4, except the “4” house where the last ion pump is at the “49” location. There are also 2 nude ion gauges per house. The term nude refers to the ion gauge not having a glass bulb. These are at the “X3” and “X7” locations. The ion pumps have an ion current readback that is converted to pressure, although the ion gauges are more accurate. The typical working ranges of the ion pumps and ion gauges are 10-6 torr to 10-10 torr. In practice the beam tube, once cold, needs no external pumping anyway. The liquid helium temperature of the tube condenses out any gases, even helium. This is referred to as cryo pumping. The beam tube vacuum is required to be on the order of 10-8 torr, and is usually better.

Fig. 6.1: TeV Sector A vacuum

The beam valves at each end of a house are interlocked. They can only be opened when an ion pump permit exists at both houses on each side of the valve. The ion pump card is responsible for issuing an ion pump permit. An ion pump card has 6 inputs. Four of these are the 4 ion pumps per numbered house and the last two are the two ion gauges. Only the 4 ion pumps take part in the permit. The card issues a permit when at least 2 of the 4 ion pumps are on. Typically, the 24 numbered houses have only 1 ion pump card, while the zero houses can have up to 3 cards. Each of the 3 cards must issue a permit for that zero house to issue an ion pump permit. The ion pumps will trip off in 15 minutes if the vacuum exceeds 10-6 torr for that duration.

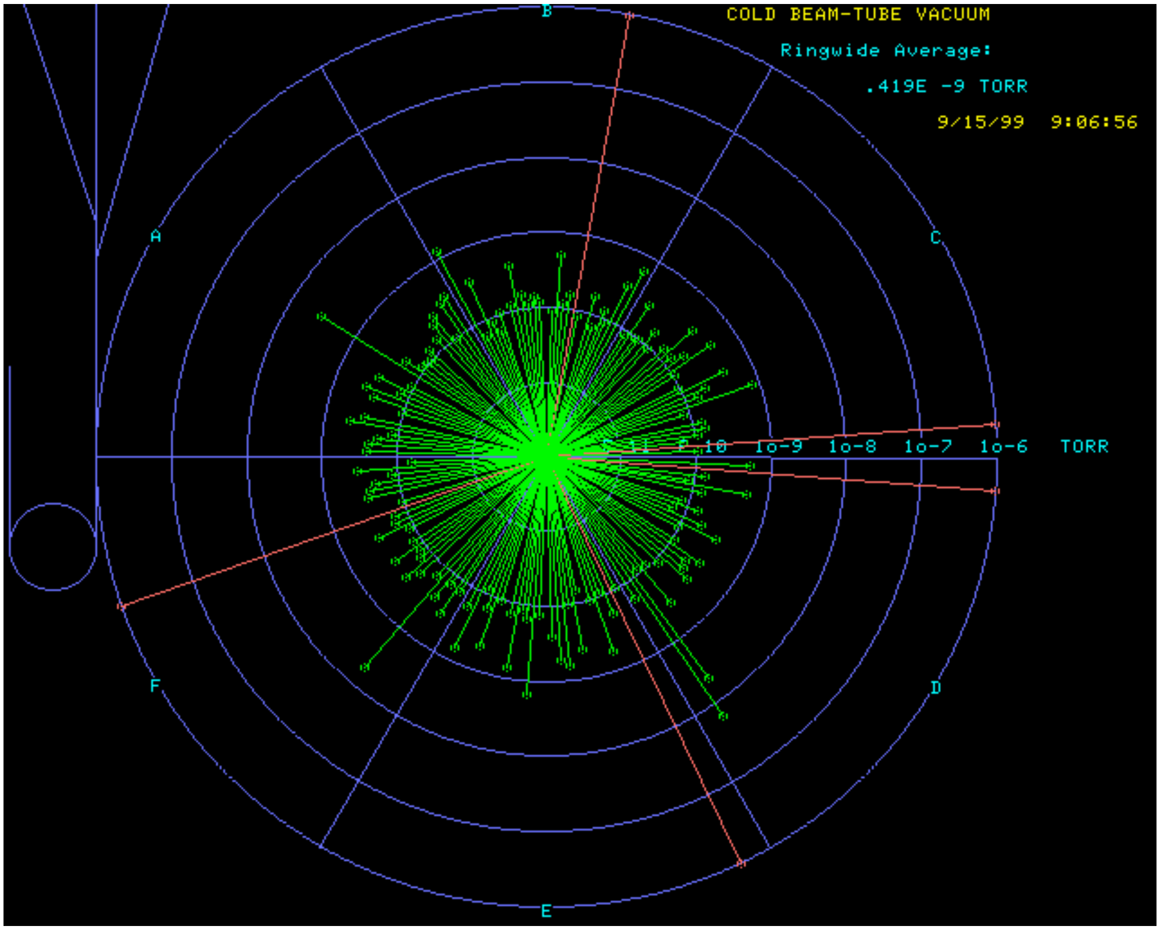

Fig 6.2: Cold Beam-Tube Vacuum

Warm Beam Tube Vacuum

The warm beam tube resides at room temperature and exists at the 6 long straight sections and each of the “17” and “48” locations. These sections generally contain specialty devices such as kickers, injection and extraction magnets, RF, etc. The vacuum equipment at these locations is an integral part of these devices and thus is different at each warm section. Each warm straight has an isolation valve at each end. Some of the long straights have intermediate valves to break up the section.

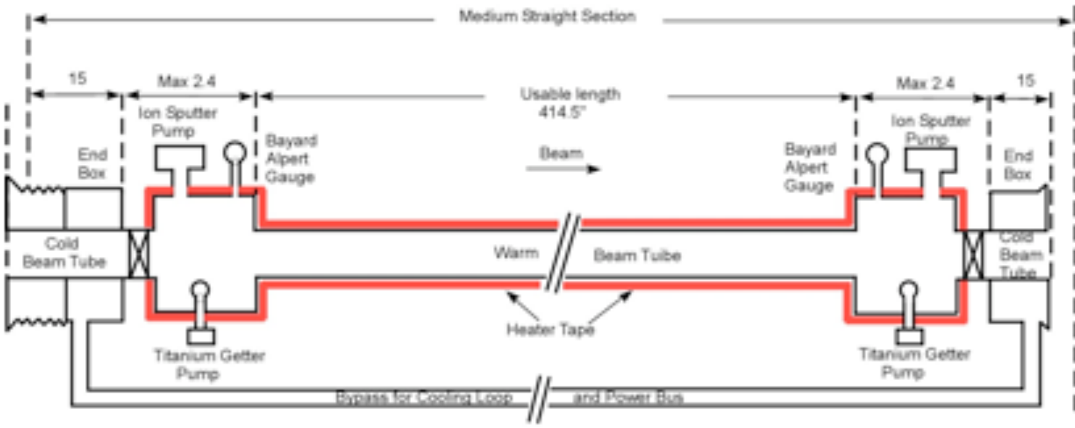

Fig. 6.3: Warm Beam Tube Vacuum

Cryostat Insulating Vacuum

The cryostat vacuum insulates the liquid nitrogen and liquid helium in the magnets. With a good vacuum, the main heat load becomes conduction through the magnet super-insulation and radiant heat across the vacuum. Radiant heat is a function of the temperature to the fourth power. This is why the liquid nitrogen shield exists; 774<<2734. The cryostat vacuum is required to be only better than 10-5 torr. It is usually much better, about 10-7 torr. The cryostat vacuum is broken into half-cells by permanent barriers at each quadrupole. This is to facilitate the isolation and elimination of vacuum leaks. The half-cell vacuum sections are each connected by remotely controlled cryostat valves. Once the vacuum is good in 2 adjacent half-cells the valve is normally left open.

Each numbered house has 2 pumping stations. These are located at the “X3” and “X7” locations, X=1, 2, 3, 4. The pumping stations are all identical. A station consists of a small turbo-molecular pump (100 l/s) and a rotary roughing pump (5 l/s). The pumping stations are connected to both sides of the barrier at this location (“X3” or “X7”) through two cryostat valves (upstream and downstream). The roughing pumps are remotely controlled while the turbo pumps simply come on automatically when the rougher has pumped the manifold down.

Two types of devices make pressure measurements in the cryostat vacuum: thermocouples and cold cathodes. Thermocouples are also called Pirani gauges and measure vacuum between atmospheric pressure and 10-3 torr. Cold cathodes measure from

10-3 to 10-8 torr. The cryostat valves are interlocked to the thermocouples. When there is a significant pressure differential between two half-cells or between the pumping manifold and either upstream or downstream half-cell, the appropriate cryostat valve will not open on command. It will hold a “request” to the valve and open when the differential pressure is low. The cold cathodes are interlocked to the thermocouples; the cold cathode will turn on when the adjacent thermocouple reads less than 1 micron (10-3 torr).

The vacuum system and the cryogenic system are invariably coupled. A magnet string can never be cooled down when the insulating vacuum is poor. On the other hand, if there are no leaks, a vacuum that is only fair will be cryo pumped by cooling the magnet string.

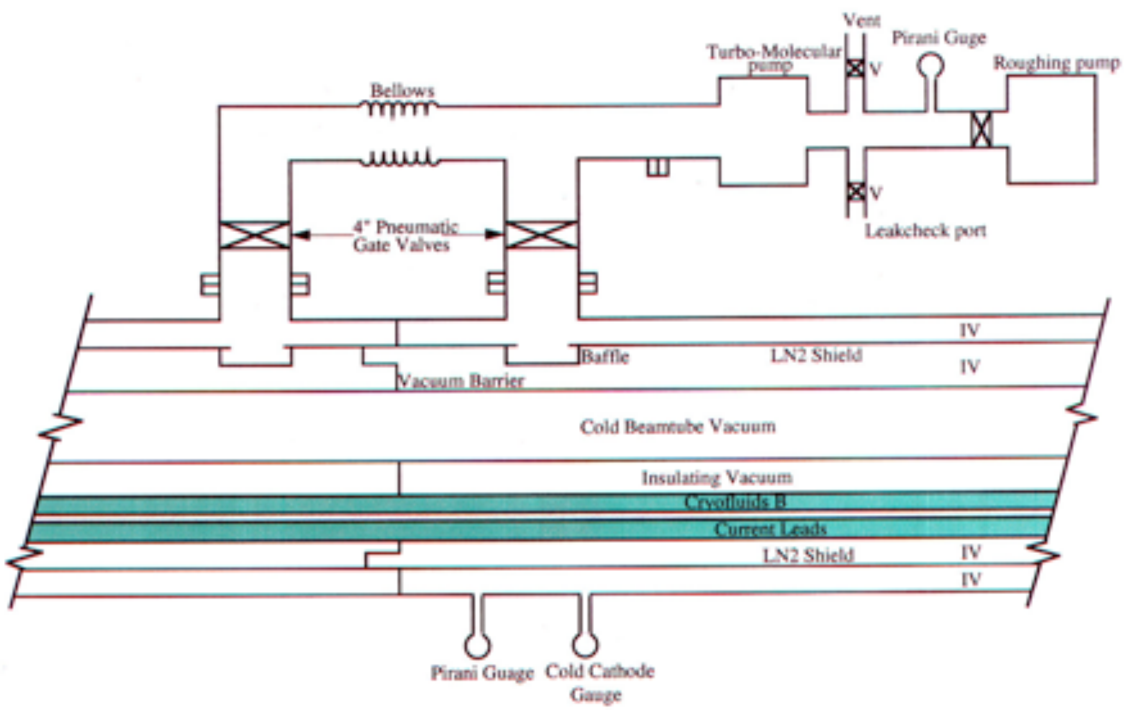

Fig. 6.4: Pump Station for Insulating Vacuum

Notes:

7. Controls

In order for Operations to monitor and control all of the hardware throughout the Tevatron a means of communication is required to bring information to the Main Control Room. The "links" are what provide this communication. There are several links associated with the TeV and a variety of transmission media, hardware, and software used for each. The most common link is the CAMAC link. Other links are the QPM, TVLLRF, Refrigerator, HLRF, and Abort.

CAMAC Link & Cards

CAMAC is an acronym for "Computer Automated Monitor and Control." Its hardware consists of a Multibus II front end (MB II FE), a 19-conductor cable, crates, repeaters, and specialized cards.

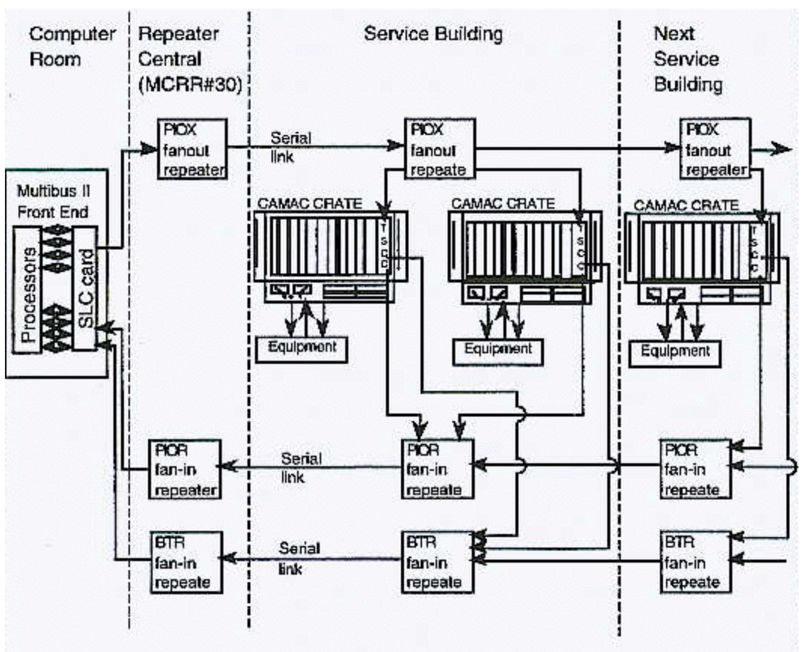

Fig. 7.1: CAMAC Hardware

The CAMAC link is generated by the Serial Link Controller card that resides in the MB II FE located in the computer room. Three links are associated with each FE. The first is PIOX, or the transmission link. Requests for information originate at the FE and are sent via the repeaters throughout the link, where the appropriate CAMAC crate decodes them. The second link is PIOR, also known as the receiving link, provides the replies of requests to the FE. The third link is BTR, or block transfer. The BTR link allows large amounts of data to be sent to the FE without interruption, such as for fast time plots.

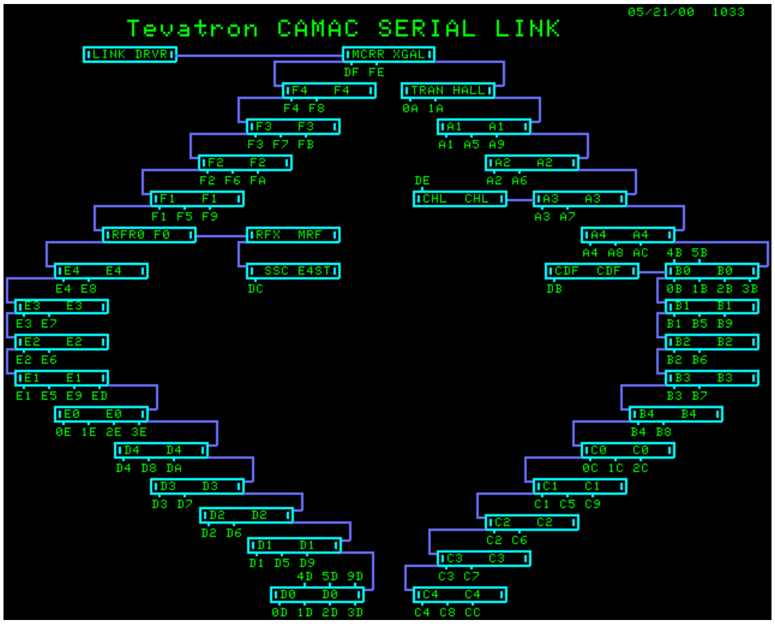

Fig. 7.2: CAMAC Link

From the SLC card the next step in the link is repeater central, located in Rack #30 in the MCR. From there the link divides with one branch going from F4 to D0 and the other from Transfer Gallery to C4. See the diagram below for the TeV link layout.

The link uses a 19-conductor heliax cable to propagate the signals. The cable can be seen fanning out from a large junction box attached to the ceiling of each service building. Since the Tevatron ring circumference is 4 miles, the electrical signal broadcast onto the link would normally have degraded by the time it reached the furthest service buildings. This is why repeaters are found at each service building. They boost the signal at regular intervals. Repeater cards are housed in half-high NIM crates.

Each CAMAC crate has a unique address, which is expressed in hexadecimal. Note the addresses in the picture of the serial link above. Within each crate is the Tevatron Serial Crate Controller (TSCC) card that resides in the 2 leftmost slots of the 25 slots available. The TSCC decodes the addresses within the transmissions on the link so that when it detects its own address the data is routed to the intended specialized card.

The CAMAC cards in each crate are the most extensive part of the CAMAC system. The cards can have a variety of functions. Some may be timing cards while others are ramp controllers. A listing and function of the most commonly used cards will follow later.

So how does the communication work? The front end has a 386 processor that uses an 8 bit parallel bus for data manipulation. In order to avoid burying miles of ribbon cable, a serial format to the data was a better choice and only one cable would be needed for the link. The SLC card converts the parallel data into the serial format. Along with the data an eight bit crate address, a sub address for the card, and two parity bits are added to the data stream. The TSCC card decodes the address, converts the serial format back to eight bits parallel format, and passes it to the appropriate card.

The transmission rate is 10 MHz and is coded as a square wave that alternates between 0 and 2 volts. The square wave is interpreted by the hardware as consisting of “cells” 100 nanoseconds long. If the voltage remains constant for the full lifetime of the cell, the bit is interpreted as a zero. It doesn’t matter if the voltage is high or low. If the voltage makes a transition from one state to another in the middle of the cell, then the bit is interpreted as a one.

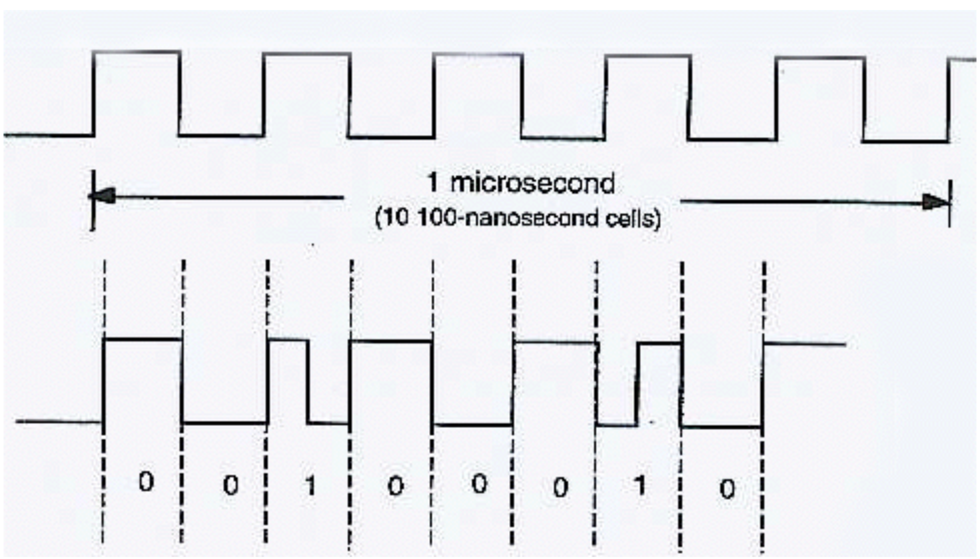

Fig. 7.3: CAMAC Data Manipulation

Abort Link

The Tevatron has kickers that remove beam either at the end of a machine cycle or during an unexpected occurrence. The abort loop monitors certain devices and if they fail then the beam is instantaneously removed from the machine. For beam to be permitted in the TeV the abort loop has to be continuous. If a failure occurs the loop is broken and the beam is kicked out of the machine.

In addition to the abort dump at C0, there is an internal dump located in the Transfer Hall, which is used during Collider operations. When antiprotons are in the Tevatron both protons and pbars are aborted into this dump. Since the intensities during Collider operation are relatively low this enclosure dump can be used.

The abort loop utilizes one of the 19 conductor cables. The link must have a 50 MHz signal present on it for the loop to be up. A CAMAC 201 module housed in a crate at the C0 service building generates the 50 MHz signal. The loop circles the ring where at every service building it encounters 200 module which can inhibit the transmission of the signal if a monitored device goes into a bad state. The 200 module is known as the abort concentrator and it accepts a maximum of 8 inputs. Behind the CAMAC crate that houses the 200 module is the abort patch panel that has the input cables for those devices to be monitored. Below the line of spigots for the abort inputs is another row of spigots labeled "current sources." The 200 module interprets the presence of current as "good" or "1". If the device input signal or current source is absent then the abort module will interrupt the 50 MHz signal and take down the loop. Device inputs can be jumpered using the current sources.

When an abort reset is sent from T67, the Abort Link Status page, the CAMAC 201 module is told to initiate the 50 MHz signal. The signal takes roughly 34 msec to traverse the ring if all is well. If the 201 module does not detect the signal within 100 msec then the card will cease transmitting and wait for the next reset to be issued.

The majority of the abort inputs will immediately remove beam when pulled. These are "Type 0" aborts. This group includes the dipoles, correction elements, low beta quads, and the Tevatron ramp itself via the A2 TECAR input. The QPMs will issue an abort in the event of quench or ramp dump. The BLMs generate an abort if losses are too high and vacuum crates will cause an abort if a beam valve closes. The 200 module in the MCR has inputs from the manual abort buttons and the safety system.

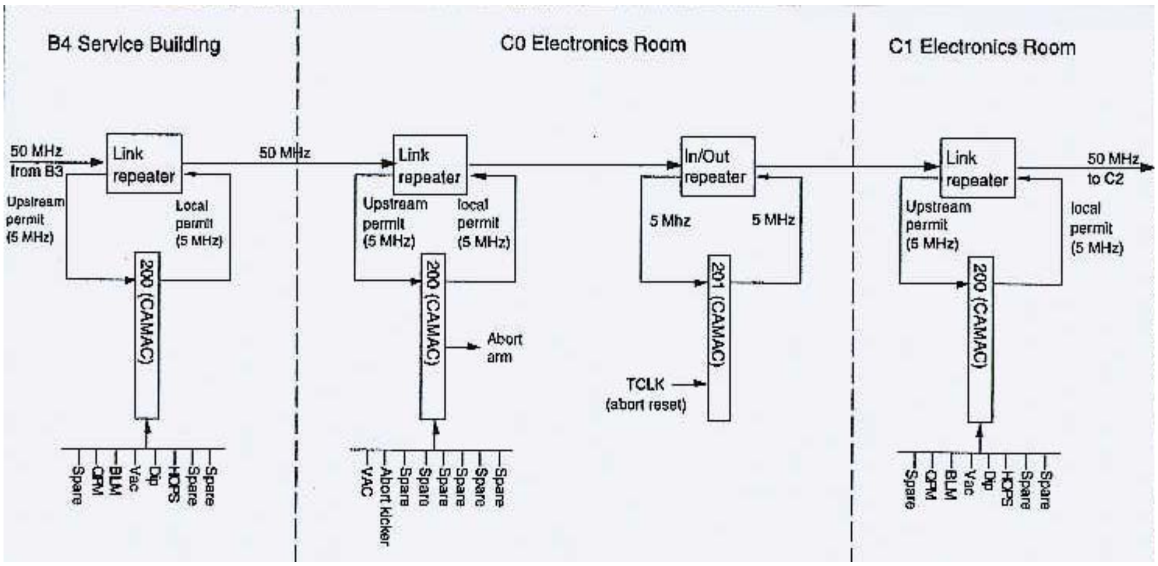

Fig. 7.4: Abort Link

Notes:

8. RF and Acceleration

The RF system for the TeV is somewhat standard compared to the Booster and Main Injector systems but there are a few differences. The beam injected into the Tevatron does not go through transition since it is already above the transition energy. There are no Ferrite bias supplies for the cavities because the frequency change from 150 GeV to 1 TeV is about 1 kHz. Kind of interesting, huh?

Introduction to RF

The Tevatron RF system is comprised of 4 major parts: 1) low level RF signal, 2) high level RF amplification, 3) transmission line, and 4) resonant cavity. Each of these systems will be explained in the following sections.

What is RF? What does RF mean? RF is an abbreviation for radio frequency. Radio frequencies are a form of electromagnetic waves. The frequencies used in the Tevatron RF cavities are specifically in the 53 MHz region. So how does a wave help in the acceleration of a proton, antiproton, or, for that matter, any charged particle? It has to do with the Lorentz force law,

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

where q is the charge of the particle, E is the value of the electric field that the particle is experiencing, v is the velocity of the particle, and B is the magnetic field the particle is going through. For accelerating a particle, we are only concerned with the first part of the equation because the cavities used have no magnetic field when the particle is present.

$$\vec{F} = q\vec{E}$$

The arrows in the equation tell us that the force exerted on a positive charged particle is entirely in the direction of the electric field. So what is the connection between RF and electric field? The answer is

where V is the electric potential, also known as voltage, and s is

$$E = -\frac{\partial V}{\partial s}$$

distance through which the voltage changes. For a synchrotron accelerator, like the Tevatron, the distance or displacement is in the longitudinal direction or z direction.

$$E_z = -\frac{\partial v}{\partial z}$$

The pieces are now starting to fall into place. One more piece should complete the picture. In order to create the voltage difference so that the electric field can be formed in the z direction, a gap must be placed within the cavity. The gap needs to have one end at ground potential, 0 volts, and the other end with the applied voltage.

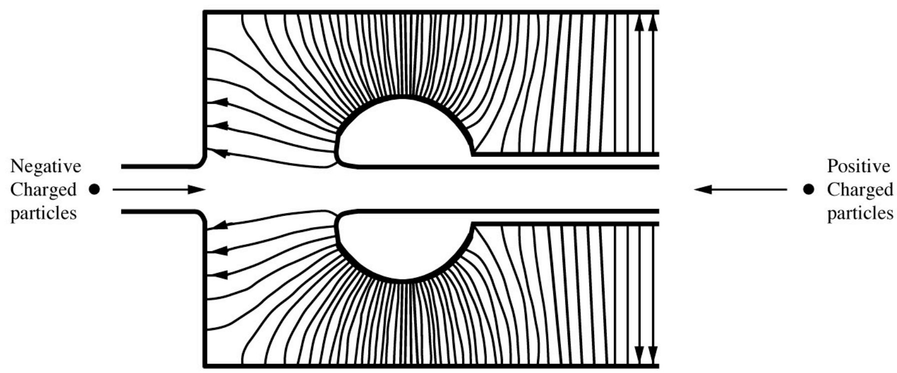

Fig. 8.1: RF Gap Fields

This side view of an RF gap shows the electric fields that are formed due to the voltage difference between the cavity shell (0 V) and inner electrode (where the RF voltage is applied).

The picture is now complete. If I take a power supply that applies to the cavity a voltage at a frequency synchronized to the beam then the charged particles will feel a force due to the electric field that is created by the voltage difference across the gap.

Low Level RF

If there is any magic box at Fermilab it is the LLRF VXI system. Booster, Main Injector, and Tevatron use direct digital synthesis (DDS) to create the waveforms output to the RF stations. In the most basic of terms, a DDS system can create any waveform by manipulating a cosine wave. The system carries a ¼ cosine table and can add, subtract, multiply, phase change, etc., that table to create the desired waveform.

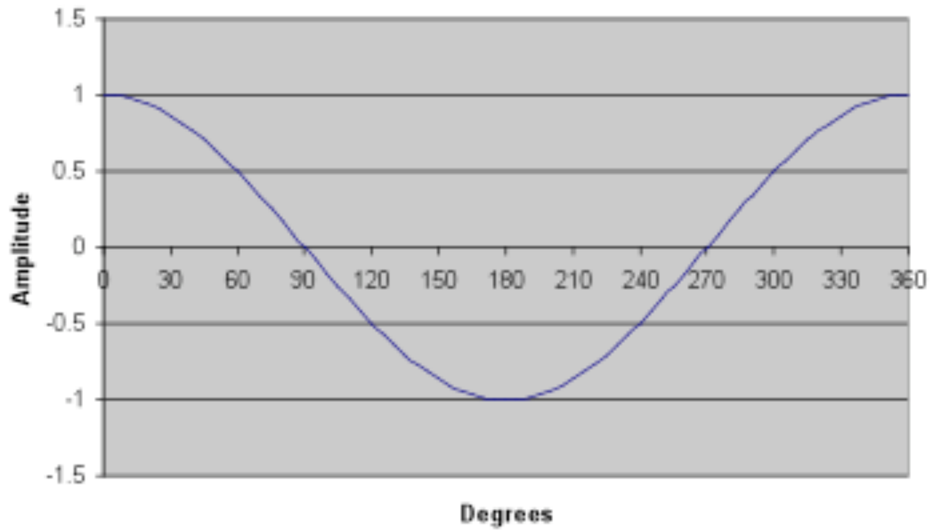

Fig. 8.2: The basic DDS system uses 1/4 of a cosine wave to create a desired waveform. A cosine is shown above

A DDS system can be thought of as a 3 block system. A clock signal and frequency signal, which is represented as a number, are input to a digital accumulator. The waveform map provides the cosine waveform to be manipulated and the output of that is sent to a digital-to-analog converter, which then provides the RF system with the appropriate LLRF signal. The whole system is often called a number controlled oscillator (NCO).

Fig. 8.3: DDS system block diagram

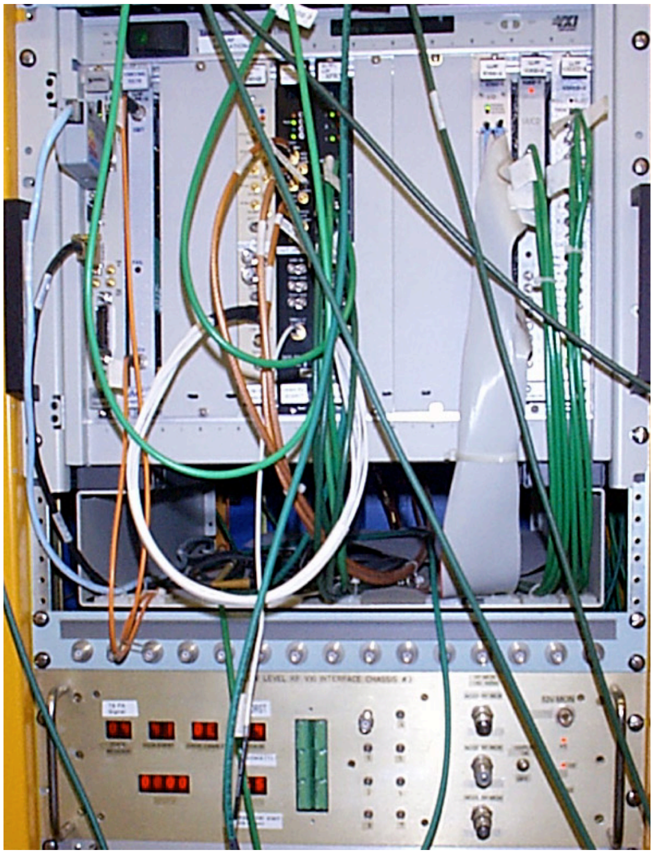
Fig. 8.4: The LLRF system housed at MI-60 consists of a VXI system and VXI Interface chassis.

So now that you know the basics of creating a low level signal what does the actual LLRF system look like? The VXI crate that houses the DDS system is located in the control room at MI-60.

The low level system is a combination of software and hardware. An introduction to the basic system diagram was just given but now the intricacies of the software and hardware will now be explained.

The digital signal processor (DSP) and its software create the RF bucket while controlling bucket area, radial beam position, and longitudinal bucket position. This must be done with very little RF phase noise so that the beam is not heated, which could cause beam loss. The software code must perform the frequency and phase calculations and ensure that they are smooth and accurate. The feed forward frequency is calculated from the real time programmed bend bus current, T:MDAT10, which is broadcast at 720 Hz. The frequency calculation is as follows. The frequency of revolution for a particle is

$$f_{rev} = v/2\pi r$$

where v is the velocity of the particle and r is the radius of the Tevatron, 1 km. The harmonic number of the TeV is h=1113. Multiplying the revolution frequency by the harmonic number gives the frequency of the RF.

$$f_{RF} = hv/2\pi r$$

If the particle was able to move at the speed of light (it cannot because it has a mass) it would take an infinite amount of energy to attain that velocity. The frequency of the particle at this infinite energy is

$$f_{\infty} = \frac{1113c}{2\pi r} = 53{,}105{,}071 Hz$$

where c is the speed of light, 3.0x108 m/s. The energy dependence of the RF frequency can be determined by using the relativistic energy

$$E = \gamma E_0 = \frac{E_0}{\sqrt{1 - v^2/c^2}}$$

where E_o is the rest mass of the particle moc2. The above equation can be rearranged to find the velocity

$$v = c\sqrt{1 - \left(\frac{E_0}{E}\right)^2}$$

Since E>> E_o a Taylor expansion† can be performed on the square root and the first 2 terms used.

$$v = c\left[1 - \frac{1}{2}\left(\frac{E_0}{E}\right)^2\right]$$

Substituting the velocity term into the frequency equation yields

$$f(E) = \frac{hc}{2\pi r} - \frac{hc}{4\pi r}\left(\frac{E_0}{E}\right)^2$$

where the first term is the frequency at infinite energy and the second term is df(E), the change in frequency with respect to energy.

---

† *There are 2 forms used in this document. The form used on this page is (1 + x)* ^{½} *= 1 + ½ x*^{2} *+ ... The form used on the next page is (1 + x)* $^{-½}$ *= 1 − ½ x*^{2} *+ ...*

$$f(E) = f_\infty - df(E)$$
$$f(E) = f_\infty\left(1 - \frac{1}{2}\left(\frac{E_0}{E}\right)^2\right)$$

Taking this equation and reapplying a Taylor expansion to it yields an almost final form.

$$f(E) = f_\infty\left(1 + \left(\frac{E_0}{E}\right)^2\right)^{-\frac{1}{2}}$$

Finally, one more piece of the puzzle is added. The conversion from GeV to amps is 1000 GeV=4440 A, and by knowing this fact the frequency with respect to current can be determined.

$$f(I) = f_\infty\left(1 + \left(\frac{4.44[GeV/A]E_0}{I}\right)^2\right)^{-\frac{1}{2}}$$
$$f(I) = f_\infty\left(1 + \frac{K}{I^2}\right)^{-\frac{1}{2}}$$

where K is a constant.

$$K = (4.44E_0)^2$$

So now you have it. The software for the LLRF system is written to calculate the appropriate frequency from the programmed ramp current, T:MDAT10 (T:IPROG). The feed forward frequency program resolution is 100 kHz and this gives a precise output to the HLRF system.

The frequency is one calculation a LLRF system DSP needs to perform but usually there is one other calculation that is needed, the synchronous phase fs. The TeV LLRF system was designed such that it doesn't need to calculate the synchronous phase.

Now, lets discuss the hardware of the LLRF system. Of course, already the DSP has been mentioned but the system contains much more than that. The LLRF system uses the VXI platform so that software and hardware can be incorporated together to generate, with great precision, the required output signals. Other than the

VXI, a custom interface chassis is also used to distribute the signals to the RF stations. Both of these components will be discussed below.

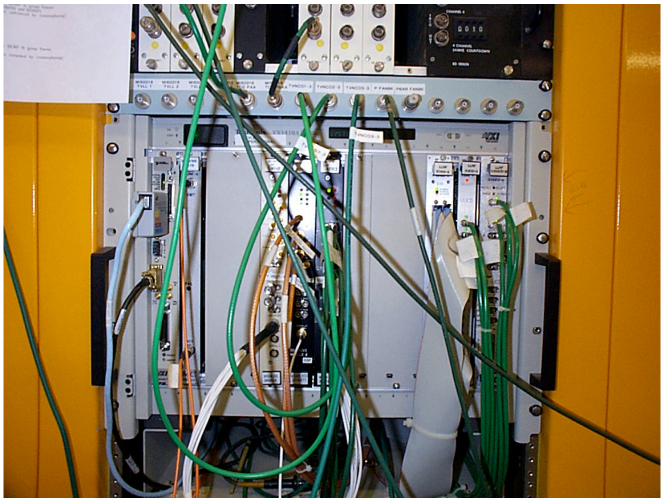

Fig. 8.5: The LLRF VXI crate and it associated cards.

The VXI system has 8 cards in the chassis. The card in slot 0 houses the VXI CPU, which configures all of the system's devices. An Ethernet port on the card provides all of the communication with the ACNET control system. Slot 1 has two fiber optic connections on the front of the card: transmit and receive. This card is used for reflective memory. All of the LLRF systems are linked by keeping copies of each other's code.

An IO100VXI digital I/O module, located in slot 10, generates TTL output signals for the machine states and drives the LED display on the VXI Interface chassis. Slots 11 and 12 house the VME and VXI Universal Clock Decoder modules, which receive and process TCLK, MDAT, and Beam Sync signals. The UCDs synchronize the LLRF system and provide ACNET timing resources for plots.

A 2 channel 200 kHz analog-to-digital converter in Slot 4 converts the beam position and phase detector signals and sends the data to the DDS. The beam transfer information is decoded in the XFR card in slot 6. Beam Sync, MDAT, and Transfer Sync information is processed and sent to the DDS.

Slot 5 houses the DDS card, where the DSP resides. The output of the DSP (frequency) goes to 3 NCOs, which in turn have their output sent to a phase modulator that amplifies the signal by +17 dBm. The LLRF output signals are carried on 3 gold colored cables. The top cable is for the beam sync system, the middle cable is for the proton RF stations, and the bottom cable is for the pbar RF stations.

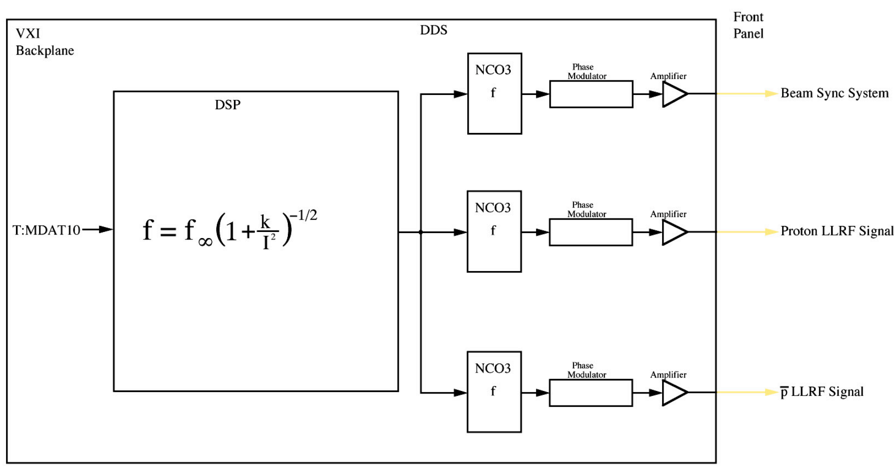

Fig. 8.6: DDS board schematic showing the frequency calculation and the output signals

The above mentioned cables are fed into the back of the VXI Interface chassis, where each signal is sent through a bandpass filter, amplifier, and power splitter. The LLRF signals are then sent to their designated RF station.

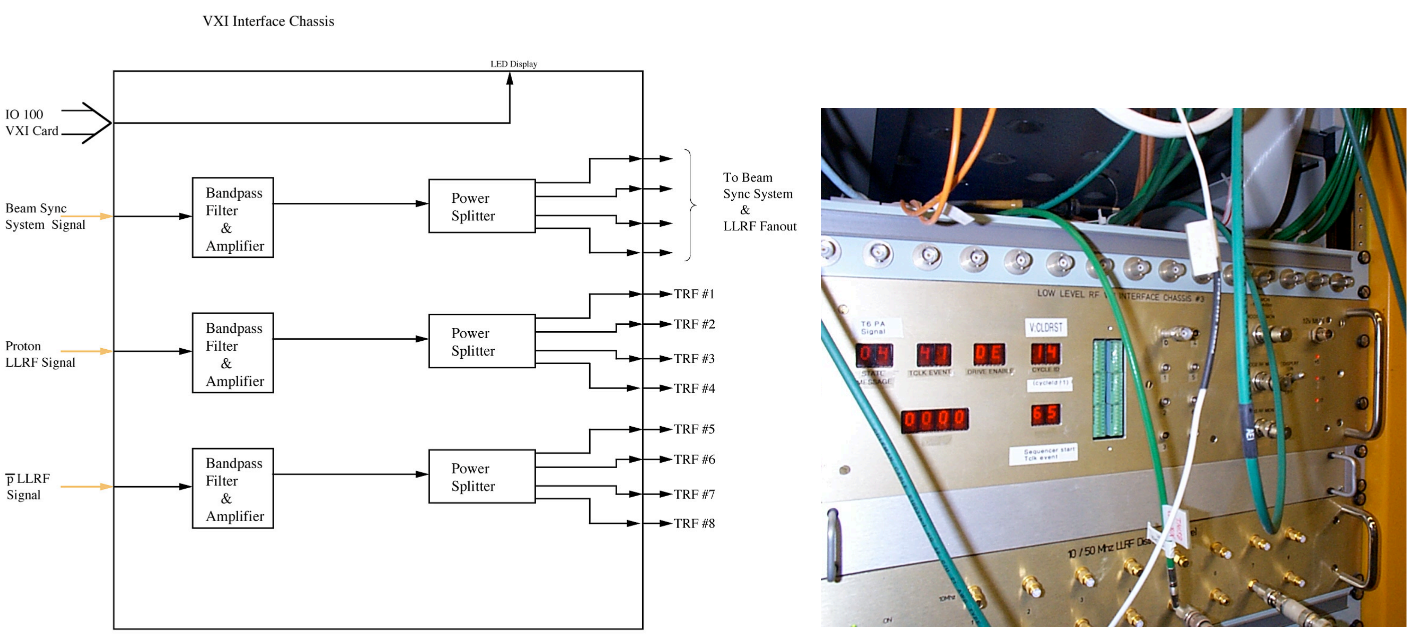

Fig. 8.7: VXI interface chassis board schematic and a picture of the chassis at MI-60

High Level RF

There are special requirements of the RF when counter-rotating beams of non-uniform charge are to be accelerated. For storage of protons and antiprotons there are restrictions on cavity spacing. Appropriately spacing and phasing of the RF fields in the individual cavities can optimize acceleration of the counter-rotating beams. The requirements of the RF system for p-pbar operation are:

The RF system must create a sufficient bucket area for simultaneous acceleration and storage of the protons and antiprotons.

The RF system must provide the capability for moving the bunch collision point longitudinally.

The system must allow for independent control of phase and amplitude, i.e. bucket size and location, of both protons and antiproton buckets.

The high level RF is where the amplification of the low level signal takes place. This section of the RF system is very similar to the Booster and Main Injector RF systems except that the Tevatron RF does not employ bias supplies. The high level system is comprised of an anode supply, modulator, power amplifier, transmission line, and resonant cavity.

Transmission Line

The transmission line is a means of transporting power from the RF source to the resonant cavity. In essence, it is a distributed capacitor. A coaxial transmission line consists of an inner conductor cylinder and grounded outer cylinder separated by a dielectric medium. In the case for the TeV RF stations the medium is air. An air blower at F0 forces air into the transmission line; the exits through a copper mesh near its connection to the RF cavity. The coaxial line has a 50 W impedance and is 9 3/16 inches in diameter. At 53.104 MHz the transmission line will have a power rating of 600 kW. One end of the transmission line is fixed at the cavity and the other end has an adjustable sliding section at the resonator in the RF gallery, which allows the line to be tuned to 3l/2. The resonator matches the PA output impedance into the 50 W transmission line.

Resonant Cavity

The cavity is a coaxial resonator that is 12 inches in diameter and 108.25 inches in length. As with all conventional RF cavities, it is made of copper and has ceramic RF windows made of 99% Al2O3. The impedance is roughly 70 W over the majority of its length, being lower in the center. Because the frequency range of the Tevatron is small, Df=2.271 kHz, the cavity is designed as a fixed-tuned two-gap structure of length ~bl/2.

The accelerating gap utilizes corona rings to minimize the corona ring gradient. At 200 kV peak gap voltage, the RF gradient is roughly 70 kV/cm. RF power is applied near the center of the cavity at a point where the 50 W transmission line is impedance matched at full load.

The Tevatron has 8 RF cavities that reside in the F0 straight section. In colliding beams mode they operate as 2 independent groups. Cavities 1, 3, 5, and 7 accelerate antiprotons while 2, 4, 6, and 8 accelerate protons.

Each cavity contains 2 quarter wave resonators with a drift tube separating the 2 acceleration gaps. Refer to the drawing below. The cavity is tuned to operate at 53.104 MHz and a peak voltage of 360 kV (180 per gap). The cavities are kept resonant by a temperature-controlled water system that circulates LCW through the upstream and downstream ends of the drift tube to maintain the 53.104 MHz center frequency.

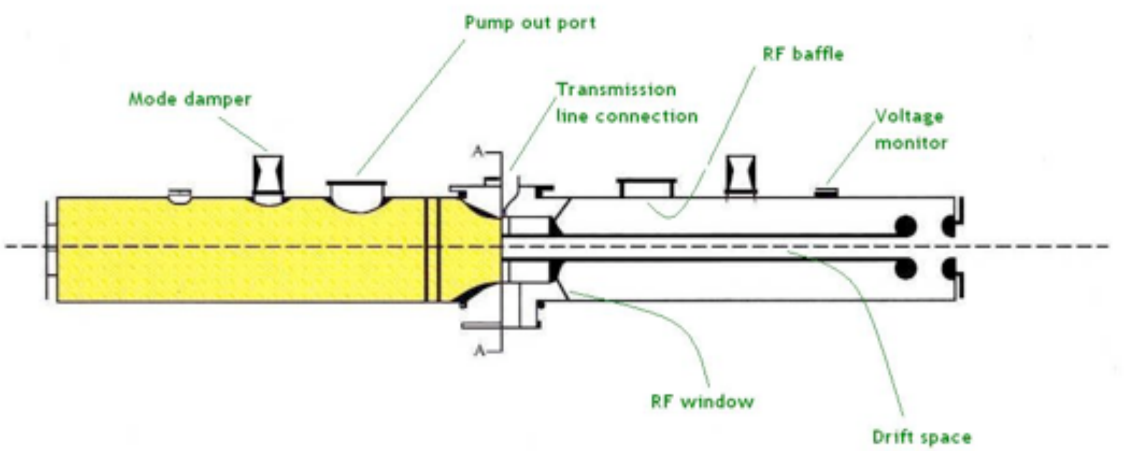

Fig. 8.8: Wave Resonator with Drift Tube

The high level RF system contains the typical devices used for particle acceleration. An anode power supply provides 35 kV to 8 modulators. The modulators pulse the high voltage to the anode of the PA (power amplifier) through the series tube in the modulator. The anode program, APGS, is input to the grid of the series tube. A 100 W ENI solid state amplifier amplifies the LLRF signal. This signal is sent to the solid state driver, which then drives the cathode of the PA.

Notes:

9. Beam Diagnostics

Beam Position Monitors

The Tevatron Beam Position Monitor (BPM) readout electronics and software system consists of 960 channels of electronics to process analog signals from 240 BPMs, front-end software, online and controls software, and the T39 ACNET interface. The system reads signals from both ends of the existing directional stripline pickups to provide simultaneous proton and antiproton position measurements.

BPM Pickups

The pickups in the Tevatron ring are part of the superconducting quadrupole assemblies. Each BPM is a pair of 50 W striplines 18 cm long, each subtending 110 degrees of arc, with a circular aperture of 7.0 cm diameter. Each BPM measures either the vertical or the horizontal coordinate, and there are approximately 240 BPMs situated around the Tevatron ring. The pickups are directional (26dB), and are read out on both ends. Special half-length BPMs are installed near the B0 and D0 interaction regions.

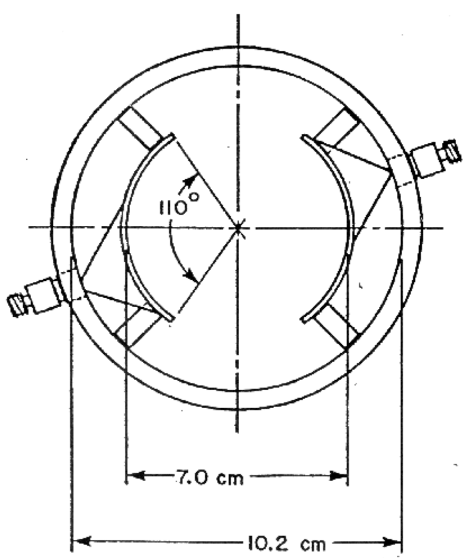

Key Specifications (Protons)*: Measurement Range: ±15mm Absolute Position Accuracy: < 1.0 mm Long Term Position Stability: < 0.05 mm Best Orbit Position Resolution: < 0.02mm Position Linearity: < 1.5% Relative Position Accuracy: < 5% Intensity Stability: < 2%
Key Specifications (Pbars)*: Measurement Range: ±15mm Absolute Position Accuracy: < 1.0 mm Long Term Position Stability: < 0.05 mm Best Orbit Position Resolution: < 0.05mm Position Linearity: < 1.5% Relative Position Accuracy: < 5% Intensity Stability: < 2% *All values are 3s

Fig. 9.1: Tevatron BPM and Specifications

Electronics and Signal Processing

The 53 MHz component of the BPM pickup response is used to measure beam positions. A schematic of the signal processing path is shown in Figure 9.2. Signals from the BPM pickups in the Tevatron tunnel are carried over foam RG-8 coaxial cables to the electronics system in one of 27 service buildings situated on the surface above the ring. Each VME subrack contains a Motorola processor module, a timing board providing clock and interrupt signals, front-end analog filter boards providing 53 MHz bandpass filtering and signal attenuation, and 8-channel 80 MHz digital signal receiver boards. All interconnections are made using double-shielded coax cables. A photograph of a completed and installed VME subrack is shown in Figure 9.3.

The digital receiver board is very similar to those used for the Recycler BPMs and identical to boards acquired for the Main Injector and P1 transfer line BPM upgrades. The digital signal receiver board consists of a 14-bit A/D converter, digital down-converter, FPGA (Field Programmable Gate Array), RAM, and VME interface. Signals are synchronously digitized at 74 MHz. Narrow-band (about 1 kHz) and wide-band (47 kHz) digital filters provide closed-orbit (CO) and turn-by-turn (TBT) measurements.

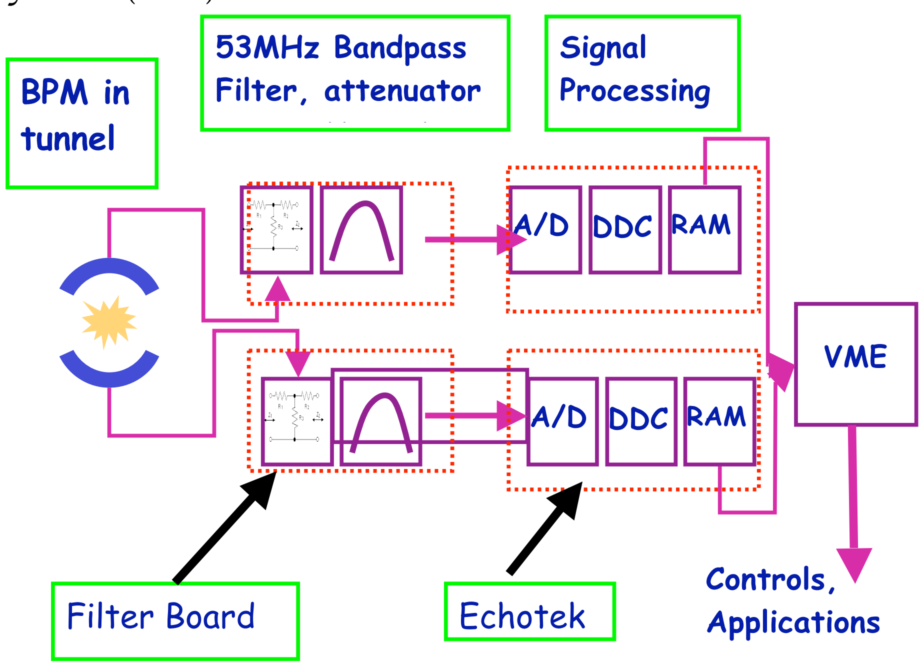

Fig. 9.2: Tevatron BPM signal processing path

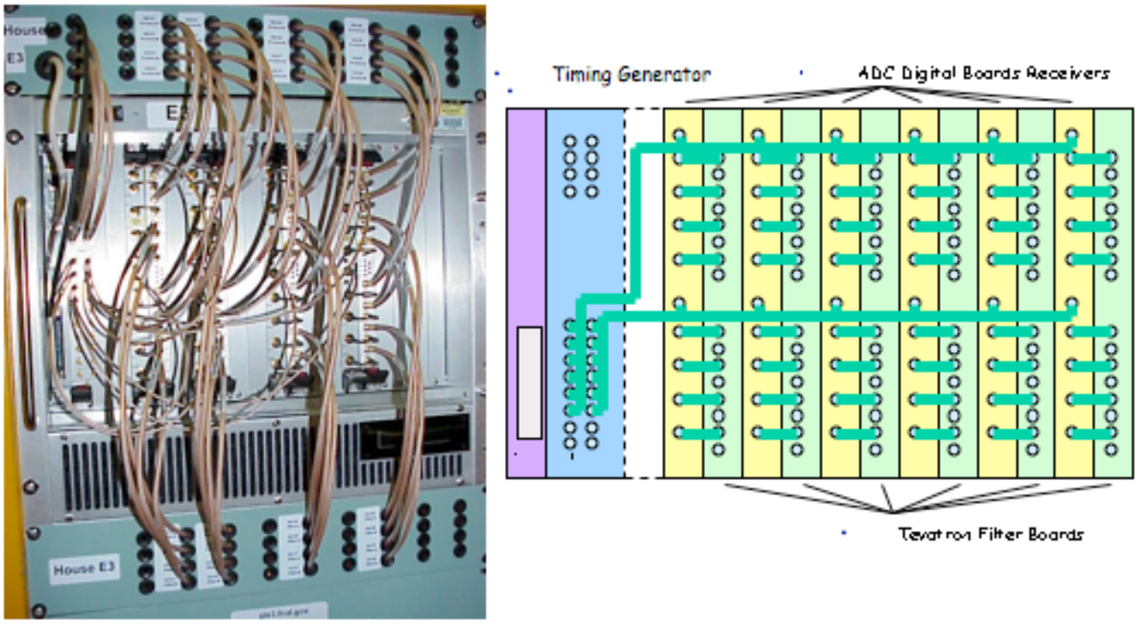

Fig. 9.3: Completed VME subrack in the E3 service building and crate layout

Closed Orbit

Transverse positions are computed using the following formula:

$$P = 26 * (|A| - |B|)/(|A| + |B|)$$

where A and B are the BPM response from the two plates and 26 is the scaling factor to convert from BPM response to position (in mm) for this pickup geometry.

An example of the performance of the system can be seen in Figure 9.4. In this figure the difference in proton positions at each BPM at 150 GeV for two stores are shown. Orbit oscillations of about 100 mm are seen.

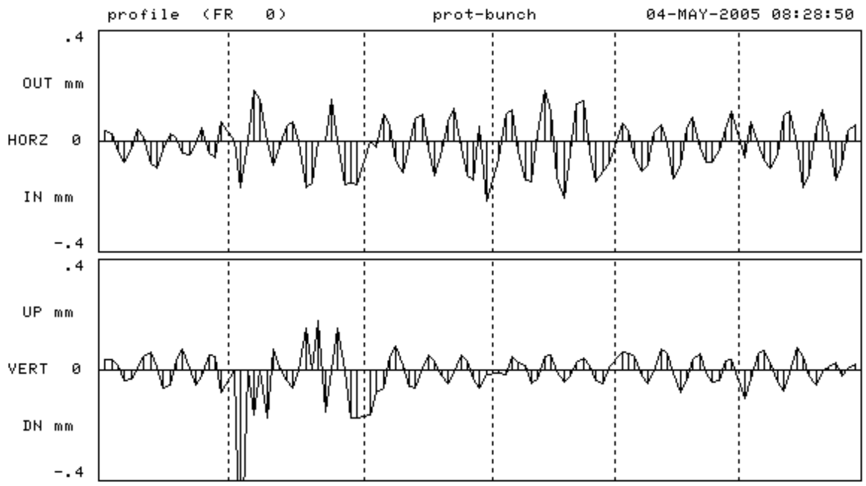

Fig. 9.4: Horizontal (top) and vertical (bottom) orbit differences between two stores, measured at 150 GeV. Vertical scale is a +- 400 mm

Antiproton Position Measurements

The antiproton beam positions are determined by a "deconvolution" technique that subtracts the proton signal contamination on the antiproton pickup. This subtraction is required because of the imperfect directionality of the pickups. The subtraction is implemented using the following formulas:

A'Pbar = APbar – aAP - bBP
B'Pbar = BPbar – cBP - dAP

The coefficients a, b, c, and d are determined empirically. The coefficients depend on the beam position in the pickup so it is important to determine these at the beginning of every store. The proton and antiproton positions at the beginning of a proton-pbar store are shown in Figure 9.5 (after deconvolution).

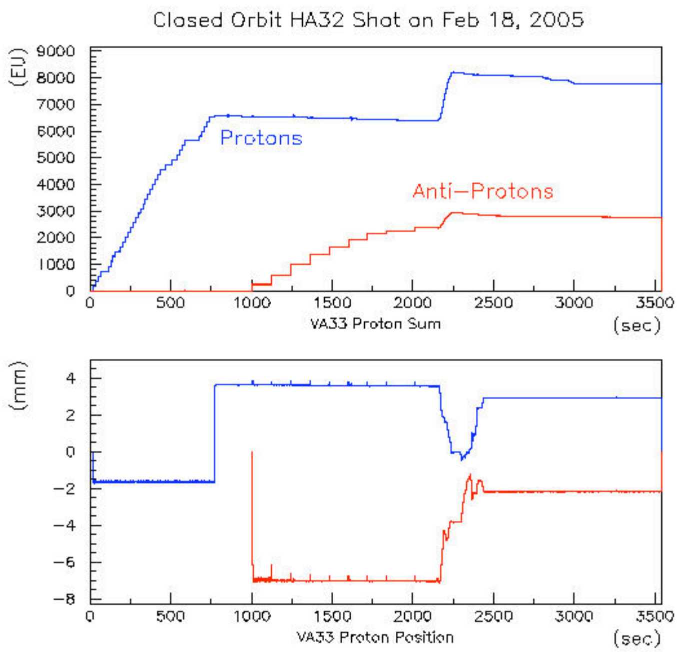

Fig. 9.5: Top left plot: proton and antiproton sum signals. Bottom left plot: Antiproton position (bottom curve) during injection and beginning of collisions.

Turn by Turn Measurements

The BPM system provides TBT measurements at beam injection and on request. The system provides 8192 turns of data. An example of TBT measurements on injection can be seen in Figure 9.6, showing synchrotron oscillations.

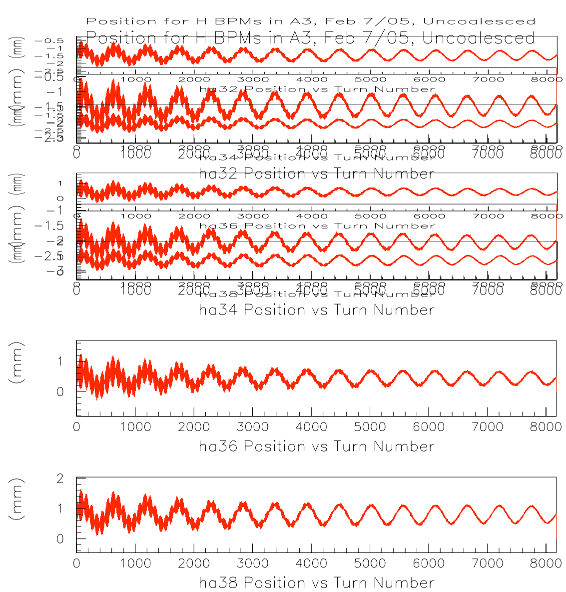

Fig. 9.6: Proton position at 4 locations during the first 8192 turns after injection into the Tevatron

Beam Loss Monitors

The BLM readout system is designed to perform several tasks: to provide a flexible and reliable abort system to protect Tevatron magnets; to provide loss monitor data during normal operations of the Tevatron; and to provide detailed diagnostic loss histories when an abort happens. Beam losses are detected using ion chambers.

The basic principle of operation of the new BLM system is to integrate for a short period of time, typically 21 μs, and digitize to 16 bits. There are two integrators per channel, running in a "Ping-Pong" mode, alternating between charge integration and digitization, so that no loss is missed. While one channel is integrating, the other is digitized, its integrator is reset, and the data are processed. The reset and processing time set a lower limit of 15 μs. The digital data are used to construct several numbers that are compared against thresholds to generate abort signals. These constructed data are sliding sums, which are a measure of the integrated loss over a variety of time scales from a single reading to the integrated loss over a period of up to 64k cycles. The abort signal is made in firmware by looking at these sums and thresholds as well as the number of channels requesting an abort.

The system uses a standard VME format crate. Besides the VME crate computer in Slot 1 that communicates data to the main control system, the BLM system includes five types of custom cards:

- Digitizer Cards (DC)
- Timing Card (TC)
- Control Card (CC)
- High Voltage card (HV)
- Abort Card (AC).

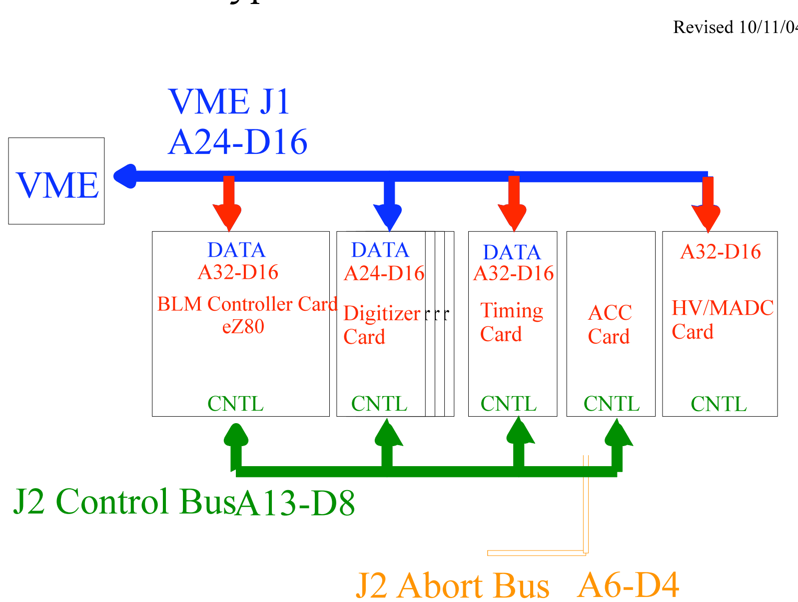

Fig. 9.7: Block diagram of a BLM crate

A custom J2 backplane is used for local system communication. A Control Bus using the user-defined pins on the J2 VME connector handles all of the critical BLM controls. This bus has 13 address lines and 8 data lines. The Controller Card is the only master on this bus, and the other cards are slaves. Also on the J2 connector is an Abort Bus where the AC is the master and the digitizer cards are the slaves.

The sliding sum time scales and corresponding buffers and abort channels are referred to either by the sum (or abort number) or intended time scale. These are as follows:

Number	Name	Typical Time Scale	Circular Buffer Depth
0	Immediate	20 μs	64kB
1	Fast	1 ms	16kB
2	Slow	50 ms	4kB
3	Very Slow	1 s	4kB

In this document, we include a summary description of each of the components followed by a description of the bus and communications protocol and detailed descriptions of the functions performed by each module including address maps.

Digitizer Card

The Digitizer Card (DC) integrates and digitizes the current from four loss monitor chambers each beam revolution. To avoid dead time between measurements, signals for each input are switched between the two channels of an integrator chip. Results are digitized from the two channels on alternate cycles and fed to on-board programmable logic devices.

The digitizer has a 16 bit resolution. Scaling is such that one digitizer count represents 15.26 fC (femto Coulombs) of charge in the integrator. The sensitivity of the BLM ion chamber is approximately 70 nC of charge per Rad.

The logic maintains three running sums per channel with programmable durations of up to 65,536 base clocks (1.4 seconds for the Tevatron) and compares the current measurement and the running sums to abort thresholds (4 thresholds in all). Each threshold can be set independently for each channel. There can be up to 15 digitizer cards in a crate.

The block diagram in Fig. 9.8 illustrates the signal processing for each channel. Note that the Sum registers will be read and the Threshold Registers written over the BLM Control Bus. The SRAM memory, which stores the integrator output values, can be read over the VME bus (J1) by the crate computer.

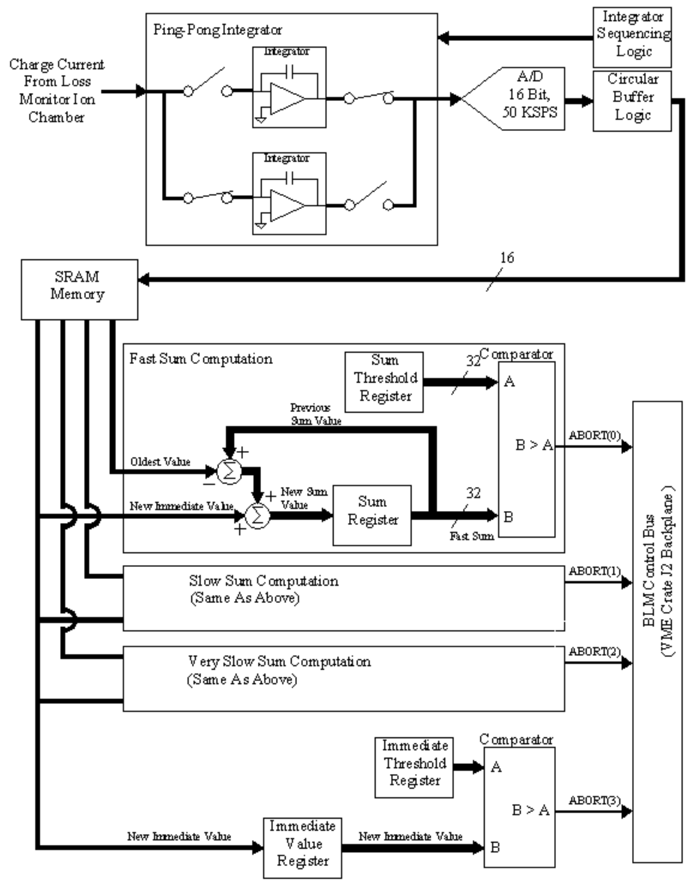

Fig. 9.8: Block diagram of the signal processing for one of the four channels on the Digitizer Card

Timing Card

The Timing Card (TC) receives accelerator system-wide timing information from three sources, the Tevatron Clock (TCLK), the Beam Sync Clock (BSYNC) and Machine state Data (MDAT).

The TC decodes BSYNC to generate the BLM system master clock, which it distributes on the BLM Control Bus. This will be generated from the AA marker with a 21ms period. The master clock signal is known as Make_Meas (“Make Measurement”).

The TC maintains a 64k circular buffer of timing information for each cycle including a 32-bit Unix time (seconds since 1970) and a 24-bit microsecond counter which is reset at one second intervals; this buffer is in parallel with the circular buffer of loss measurements in the digitizers. The master clock defines the integration interval of the digitizers and sets the threshold-comparison timing and abort-logic comparison timing. The TC also generates signals at appropriate intervals to cause the digitizers to latch the current values of the sliding sums and the Controller Card to read these sums with the latched timing information.

The TC decodes TCLK and sends a signal to freeze the data buffers in the Control card, Timing Card and Digitizers in the case of an abort. Other events from TCLK are used to signal the BLM system to collect and store synchronous ring-wide data samples for beam studies. The MDAT signal is decoded to determine the machine state and generate an interrupt to the Control Card causing it to load the appropriate abort thresholds and logic when the Tevatron machine state changes.

Control Card

To ensure that data communications and other tasks running on the VME crate computer do not impact the reliability of the BLM abort logic, the Control Card (CC) provides an independent dedicated processor that manages the setting of abort thresholds and other parameters used in the abort logic. The Control Card CPU uses a 24-bit address, 8-bit data, and 50 MHz microcontroller. The CC communicates with the other system cards over the dedicated custom J2 backplane keeping local communications separate from VME data transfers. The CC also maintains circular buffers that store the histories of the three running sums for each digitizer channel with time stamps provided by the TC. The histories will be at least 4096 time bins deep. The history can be read out via VME either on command from VME crate computer or saved in response to an accelerator control signal. The CC also stores abort thresholds for each of the sums for each channel for up to 256 machine states.

When a change in accelerator state is detected, the CC updates the thresholds in the digitizer cards as well as the abort masks and multiplicity requirements in the Abort Control Card.

HV Card

The High Voltage card is a double-wide 6U high VME module that can carry one to four high voltage modules that are independently controlled through the VME bus. It has an 8-bit switch selectable card number that sets the card address corresponding to VME address. A quad 12-bit DAC provides the program voltage for each of the high voltage modules for voltage output control. The combination multiplexer and 16-bit ADC reads the high voltage monitors for each module.

The card incorporates an FPGA (Field Programmable Gate Array) to interface with the VME bus, local timing and control. The FPGA receives all the VME bus control, address and data lines for read/write of data on the card. The FPGA can be programmed from the front panel through the Active Serial Program connector.

The program voltage circuitry consists of a 12-bit four-channel DAC device and an op-amp gain circuit. The DAC is controlled from the VME bus through an FPGA to select and send data to the selected channel. The DAC accepts straight binary which corresponds to a program voltage output of 0 to 10V. The high-voltage is linear in the setting value with a maximum value of 2250 V.

To read all the high voltage monitors, the card uses a 16-to-1 multiplexer and 16-bit unipolar input digitizer. The circuit operates in a circular mode, such that the digitizer continues digitizing all the monitor signals from each channel and storing the data into registers. The registers then can be read at anytime through the VME bus. The ADC has an input range of 0 to 3.33V with an output of straight binary. The timing and control is done through the FPGA.

The High Voltage module is a self-contain module with two SHV connectors, one for high voltage output, and the other for a high voltage return. A third connector is used as an I/O for input power, program voltage and high voltage monitoring. High voltage is produced by a dc to high voltage converter that is controlled by a programmed voltage input. The high voltage output is regulated and has low ripple output. The module also has three high voltage monitors for voltage output, voltage input return and current output monitoring. Shown below are the specifications for the module.

Abort Card

The four abort signals from each channel on each digitizer card are read by the Abort Card (AC) every integration interval. The aborts of a particular type are counted and compared to a programmable multiplicity requirement for that abort type. It is possible to mask channels off in the AC so they do not participate in the count. If the multiplicity for that integration interval equals or exceeds the threshold, a beam abort signal is generated. This logic is illustrated in Fig. 9.9. To accommodate the different operating conditions, the abort masks and multiplicity thresholds in the Abort Card can change depending on the machine state. We have also included a serial link on the Abort Card to allow a single point to receive information from all the BLM crates around the ring to be able to implement a ring-wide abort condition

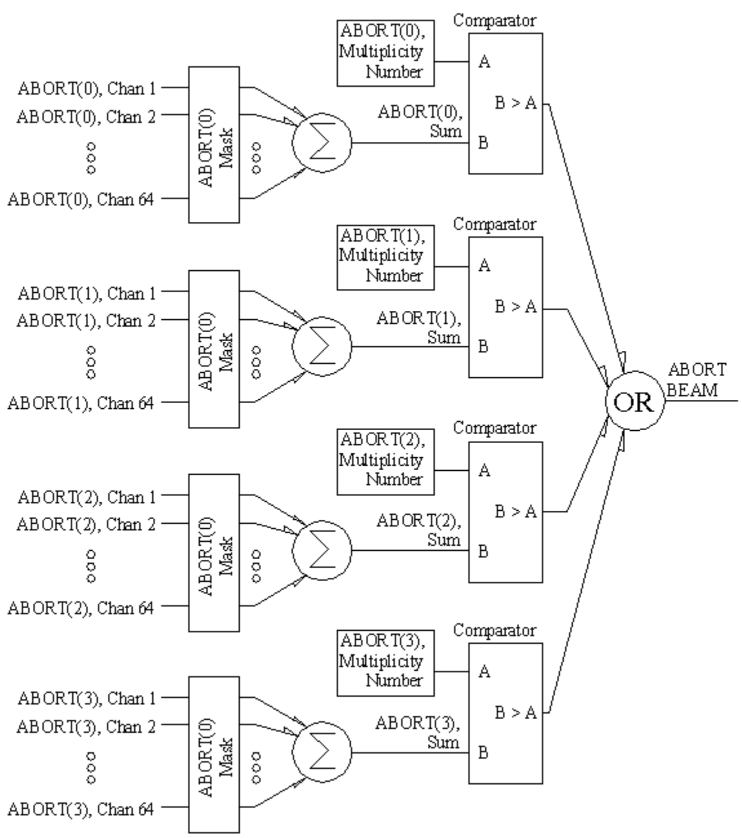

Fig. 9.9: Abort Card multiplicity logic

AC Block Diagram

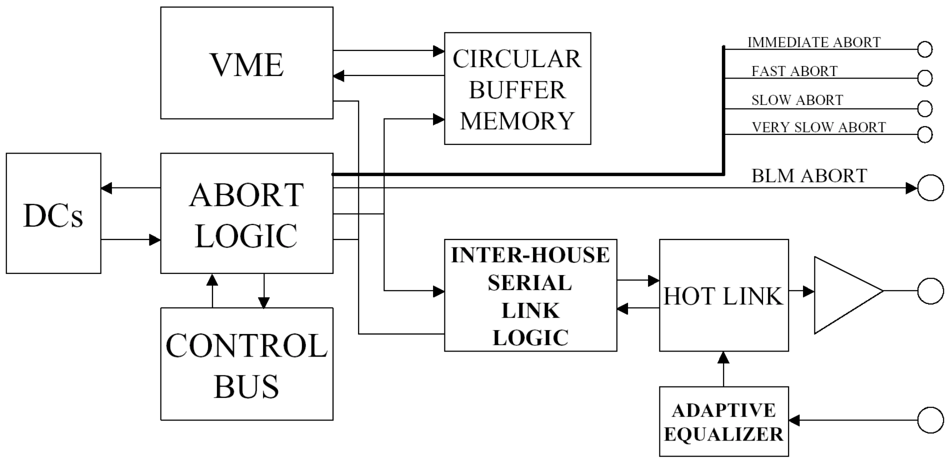

Fig. 9.10: Overview of Abort Card functions

Chassis

The chassis is an integrated 6Ux160mm VME crate, power supply and fan. In addition to the J1 backplane, each crate includes a custom J2 backplane that handles the BLM control bus with all lines bussed on the A and C rows for slots 4-21. Slots 1-3 will have no backplane connections on rows A and C. Row B includes the standard extensions for VME operation. The power supply blocks the rear of the backplane, so transition modules cannot be used in a BLM crate. The fan tray also provides an interface to allow slow control and monitoring via Ethernet.

Turn-by-Turn Buffers

In addition to the diagnostic buffers maintained by the Controller Card, which are circular buffers that are periodically overwritten, the BLM system has two linear buffers in the digitizer that are triggered and are not automatically overwritten. These buffers are Turn-By-Turn (TBT) and are each 8 kB. The TBT buffers are designed to allow the simultaneous sampling of beam position in the BPM and beam losses in the BLM.

The injection TBT buffer (ITBT) is designed to match the BPM injection TBT buffer and is triggered by ITBT Trig. The studies TBT buffer is designed to match the BPM TBT buffer used for beam studies. STBT_Trig triggers the STBT, both of these triggers are generated on the TC by clock events on TCLK or BSCLK. When either of these TBT buffers is triggered, the digitizer cards and timer card will set a bit in a status register indicating the TBT operations are in progress and that the TBT memory is not accessible from VME. Once the TBT operation completes, or if an abort happens, the digitizer cards and timer cards will reset the status bit, and VME will have access to the TBT memory.

Once it has been triggered, the ITBT will fill to its limit of 16k and stop. It will not be overwritten until another injection clock event happens. If another ITBT_Trig happens, prior to the completion of the ITBT operation, the ITBT pointer will be reset to 0, and 16kB of new TBT data will be written into the ITBT buffer. Only the first 8kB are protected from an STBT_Trig.

The STBT buffer once triggered, will fill to the limit of 16kB and stop. It will not be overwritten until another studies clock event happens. If another STBT_Trig happens, prior to the completion of the STBT operation, the STBT pointer will be reset to 8kB and 8kB of TBT data will be written into the STBT buffer.

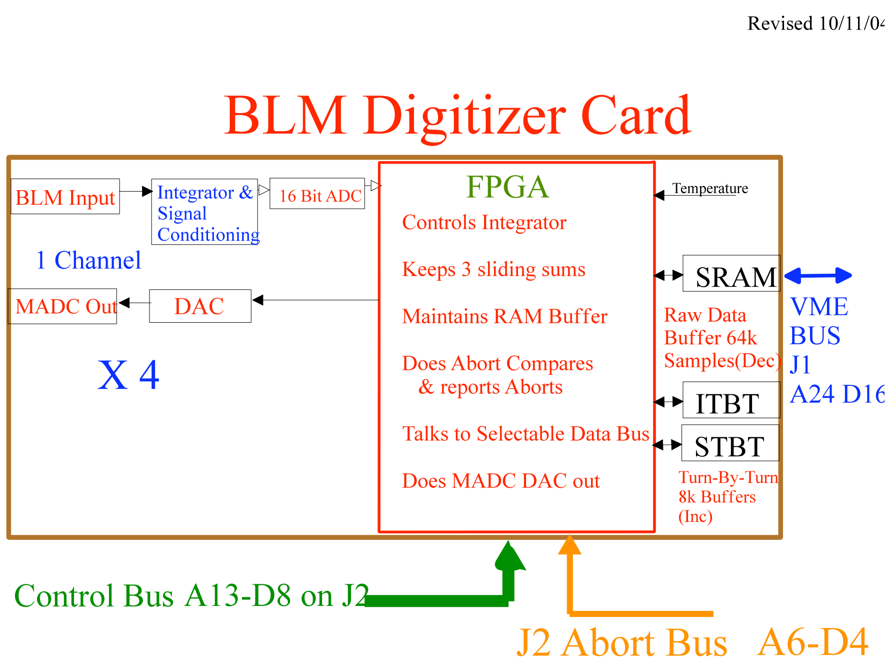

Fig. 9.11: Digitizer Card Functions

BLM Crate Normal Operations Sequence

Once the settings are loaded into the TC, DCs, and AC, the system is ready to run. The BLM operations are initiated by a clock event such as "Prepare for Beam" which will cause the TC to issue a Digitizer Card Reset (DC_Reset) on the control bus. The DC_Reset causes the DCs to zero all sliding sums and causes the DCs and the TC to set all circular buffer pointers to FFFF. This assures that all buffers are synchronized and ready to take data.

The primary clock for the BLM system, Make_Meas, is derived from the AA marker on the beam sync clock (typically 21 μs). Make_Meas is transmitted on the BLM control bus to all BLM cards. Optionally the Make_Meas signal can be created by dividing the AA marker or by dividing down an internal clock. The shortest allowable period for this signal is 15 microseconds due to the reset time needed by the integrators.

On the digitizer cards the Make_Meas signal defines the sample period, causing the integrators to switch between channels for each input and triggering the ADCs to digitize the charge for the channel not being integrated. After that, the sliding sums are updated and all abort comparisons are made. At this time the new ADC readings are written to a 64kB circular buffer, which is used for diagnostic purposes as well as the source of the sliding sums. The new ADC data may also be written to one of two turn-by-turn (TBT) dedicated studies buffers. The abort states are latched on the next Make_Meas. Thus the DC has the full sample period to do its conversions, make the sliding sums and do the abort compare with thresholds. The timing card stores real-time clock data on each cycle in a 64kB circular buffer that is synchronized with those of the digitizers.

On the AC, the Make_Meas signal causes the abort summing state machine to cycle through each BLM channel by putting the channel address on the abort bus and to read back from each channel the state of each of its abort requests. For each abort type, each channel has an abort mask bit that determines if that channel is allowed to request an abort of that type. A count is made for each of the four abort types of allowed AND requesting channels (i.e. those above threshold). If the number of channels requesting an abort for any of the four abort types equals or exceeds the abort multiplicity setting for that abort type, an abort request is transmitted from the card on a 50Ω TTL line driver.

The Make_Meas signal, therefore, causes the data to be taken and the abort logic to be updated every cycle. While a sliding sum might be the sum over 500 samples (10 ms) its abort threshold is compared every 21 μs.

During each 21μs cycle, the digitizer cards make and update the three sliding sums of samples. These sliding sums are compared every cycle to their abort limits. However, for diagnostic purposes, these sums are stored periodically in circular buffers on the Control Card. This process is controlled by the TC, which periodically generates 3 latch signals, one for each sliding sum. The latch signals cause the DCs to latch the appropriate sum and the TC to latch the time stamp and to interrupt the CC so that it knows the data is latched and ready to be read and stored in the appropriate circular buffer. The individual ADC readings are 16 bits; however, the sliding sums are 32 bit numbers. Therefore, the dynamic range of, for example, the 1 second sliding sum is almost 32 bits. These sliding sums are the total integrated loss over the sum interval, not just samples of losses spaced in time.

At any given time, the BLM has a variety of stored loss histories with different time resolutions: the 64kB raw measurement buffer provides 1.4 seconds of loss data with 21 μs resolution; the 16kB fast circular buffer provides 16 seconds of integrated loss data with 1 ms resolution, the 4kB slow circular buffer provides 200 seconds of integrated loss data with 50 ms resolution; and the 4kB very slow buffer provides 4096 seconds, over an hour, of integrated loss data with 1 second resolution. As one can see, in the event of an abort, there is a very detailed history of losses prior to the abort, which may be examined to aid in diagnosing the problem.

Sampled Bunch Display

The Tevatron SBD is used to provide information on the longitudinal parameters of coalesced beam bunches in the Tevatron. The quantities provided for each proton and antiproton bunch include the intensity, the longitudinal bunch profile, the timing of the bunch with respect to the low-level RF, the momentum spread and the longitudinal emittance. The system is capable of 2 Hz operation; it operates at 1 Hz.

The pick-up is a wide-band resistive wall-current monitor (RWCM) positioned where the proton and antiproton bunches are maximally separated (~200 ns), E48. The signal from the RWCM is brought from the tunnel and digitized in an oscilloscope located in a service building. The low level RF system provides the scope trigger. The data are read from the oscilloscope over Ethernet and processing is performed in LabView running on a Macintosh G5 computer in the Accelerator Division computer room. Accelerator parameters such as the beam energy and the RF voltage are read from the accelerator control system and the longitudinal quantities are returned via ACNET.

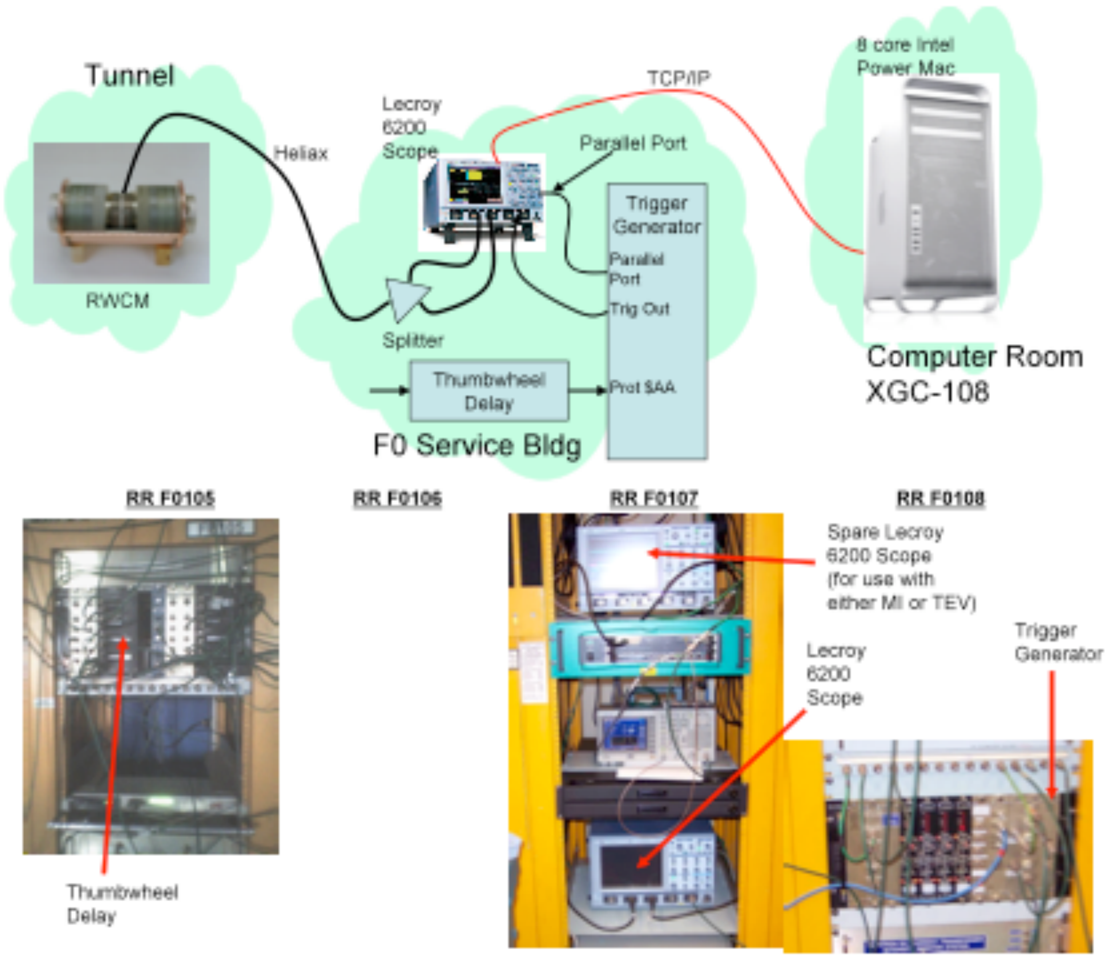

Fig. 9.12: SBD system components and connections

The RWCM has broad (>2 GHz) bandwidth with a 1.34 ohm resistance formed by 88 120 ohm resistors across the ceramic gap. A Lecroy 6200 oscilloscope is used as a digitizer. The RWCM output is brought to the F0 service building via 280 ft of 7/8” heliax cable and then split to provide two copies of the signal, just upstream of the oscilloscope. The signals are fed to two input channels with a gain ratio of ~8, the present ratio of proton to antiproton intensities. The split is positioned so that any reflections from one channel input will arrive at the other channel with 50 ns (2.5 buckets) delay.

Figure 9.13 Resistive wall current monitor cut-away

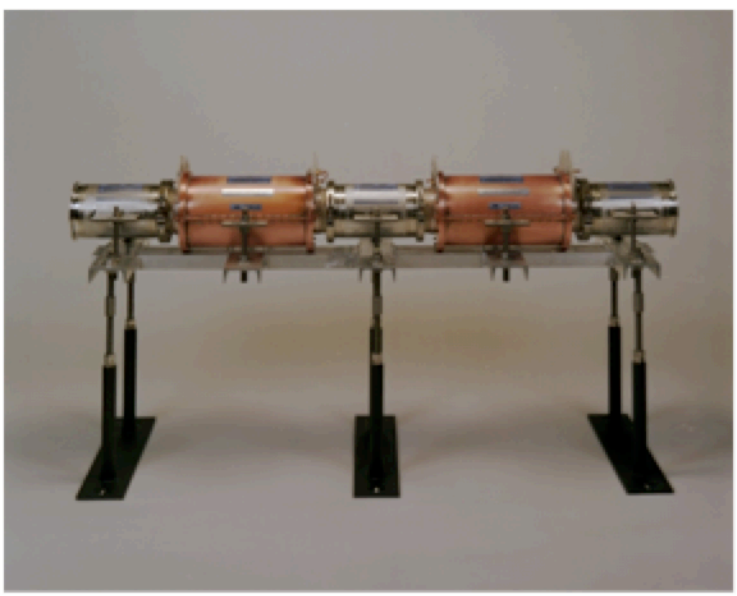

Fig. 9.14: Picture of resistive wall current assembly prior to installation at E48

The oscilloscope provides 8 bits of resolution. The high gain channel accommodates the antiproton signal and the proton signal is contained in the low gain channel. In practice, we synthesize the proton signal from both channels thus improving the resolution on the proton signal by ~ 8.

The signal is sampled at 5GS/s. To reduce the effect of digitizing noise, a set of 32 sweeps is taken and averaged by the scope. Each sweep covers 21 msecs, a full Tevatron period, and successive sweeps are taken every ~ 42 msec, triggered by the Low Level RF proton marker.

A second set of sweeps triggered by the antiproton marker is taken to obtain the antiproton RF timing. The data from an acquisition (200 Kbytes) are transferred via Ethernet to the Macintosh G5 for processing.

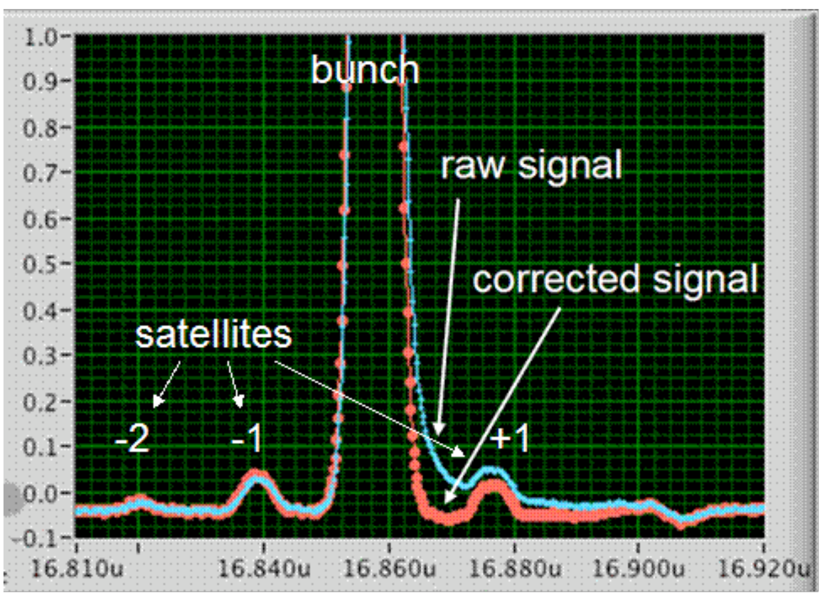

Fig. 9.15: A proton bunch signal (raw) and after correction. The feature at the far right is a 0.75 % reflection from one channel through the splitter to the other. Full height of the main bunch is ~5 amps.

Beam parameters are derived from the proton and antiproton pulse by summing the 36 sets of signals in the 5 buckets centered on the bucket containing the main bunch. The intensity measurement has a resolution of better than 0.5%; the centroid and the rms width measurements have a resolution of about 20 picoseconds.

The SBD calculated values are:

- T:SBDPIS[0 - 36] - Proton intensity
- T:SBDAIS[0 - 36] - Antiproton intensity
- T:SBDPWS[0 - 36] - Proton RMS bunch length

- T:SBDAWS[0 - 36] – Antiproton RMS bunch length
- T:SBDPMS[0 - 36] - Momentum spread for protons {rms}
- T:SBDAMS[0 - 36] - Momentum spread for antiprotons {rms}
- T:SBDPLS [0 - 36] - Longitudinal emittance for protons
- T:SBDALS [0 - 36] - Longitudinal emittance for antiprotons

Where 0 is the sum of all bunches; 1 - 36 are the individual bunches.

Fast Bunch Integrator

The Fast Bunch Integrator, aka FBI, uses the Resistive Wall Current Monitor just like the SBD. There are 3 RWCM in the TeV; one for the SBD, another for the FBI, and the last one for studies.

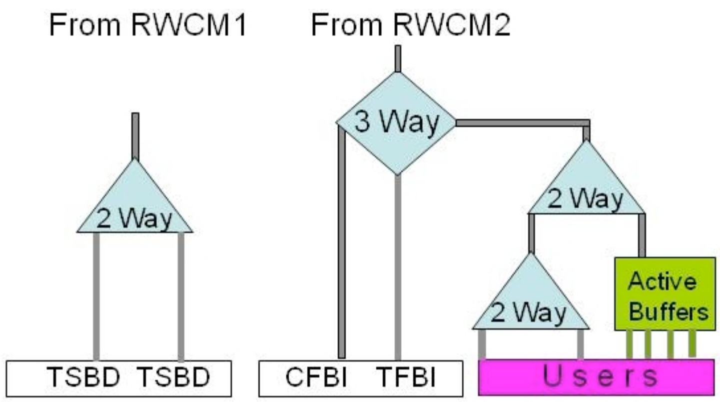

Fig. 9.16: Resistive Wall Current Monitor assignments.

The FBI returns intensity values for each of the 36 protons and antiproton bunches, a background value, and an intensity summation value. All intensity readbacks are arrayed devices. The digitizer cards are located in the F0 control room.

SyncLite and Abort Gap Monitoring

During operation of the Tevatron in colliding beam mode, a small amount of the beam diffuses out of the bunches and spreads around the ring. The presence of beam in the abort gap can have a serious effect on superconducting magnets and a devastating effect on the silicon detector of CDF. During an abort, the kicker magnets ramp up during the abort gap. Beam passing through the kickers while they are ramping sprays into magnets and into the silicon detector. There are 2 methods for directly measuring the beam in the abort gap using synchrotron light: CID camera, and photomultiplier tube.

Theory

A charged particle that undergoes transverse acceleration emits radiation in a cone around its velocity vector. This radiation is called synchrotron radiation after its first observation in a synchrotron. The Tevatron has 1113 RF buckets and typically contains 1013 protons in 36 bunches arranged in 3 trains of 12 (the antiproton intensity is ~1/10 the proton intensity).

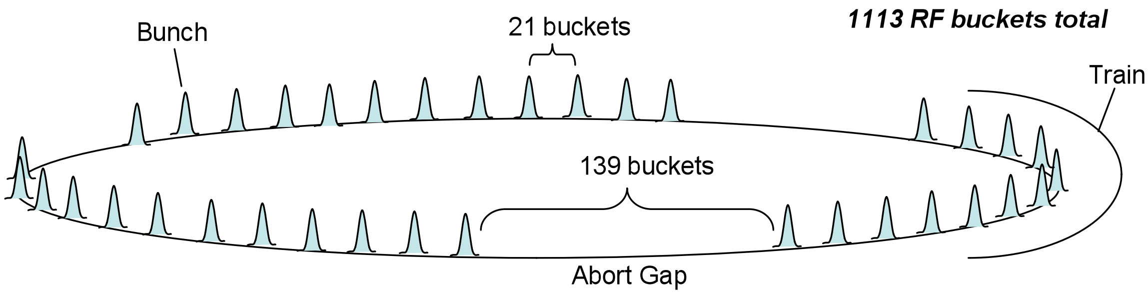

Figure 9.17: Bunch spacing in the Tevatron

Apparatus

The following two figures shows the optics of the synchrotron light apparatus, which is located near the short warm section at C11.

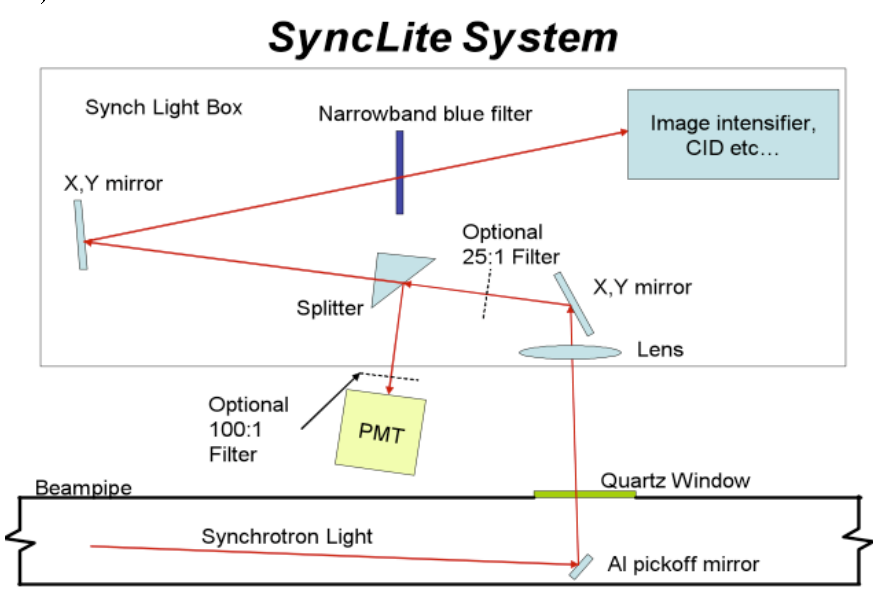

Fig. 9.18 A: Diagram of SyncLite system.

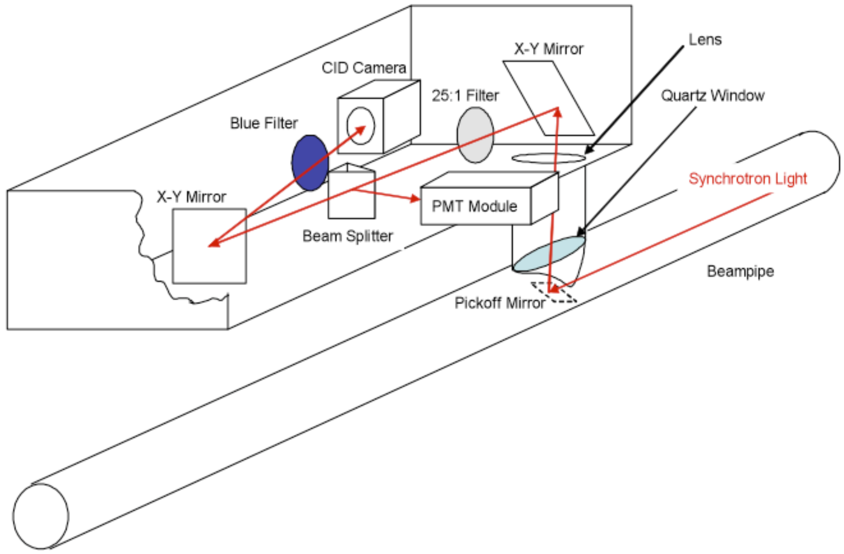

Fig. 9.18 B: Diagram of SyncLite system. The optics through the beam splitter is shared by both the camera and PMT systems. The top drawing is the logical diagram and the bottom drawing is the physical layout (100:1 filter is in PMT module).

The light is picked off by a mirror in the beam pipe and directed out a quartz window to the light box. Inside the light box, the light traverses a 1500mm focal length lens and another mirror before hitting the beam splitter. The synchrotron light is clearly visible on the wall of light box in Fig.

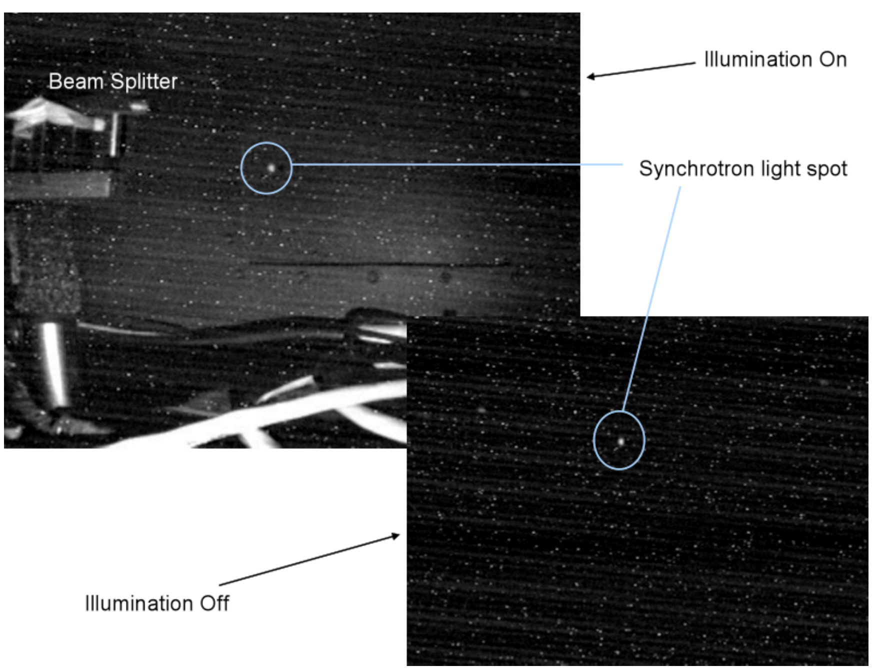

Fig. 9.19: Video stills from a CCD camera mounted in the light tight box showing synchrotron light impacting the side of the light tight box. The smaller white specs are radiation damaged pixels.

After the beam splitter, the PMT system and camera system follow separate paths. Just before the beam splitter there is a 4% neutral density filter that can be inserted into the light path to facilitate calibrations.

The observed number of photoelectrons is tabulated as follows:

	Optical Efficiency Through Beam splitter	Optical Efficiency after beam splitter	Wavelength acceptance	Photocathode Quantum efficiency	Photoelectrons /109 protons /bucket
PMT System	0.34	–	100 nm	0.14	0.1
CID System		0.92	20 nm	0.12	0.015

CID and Camera Version

Every pixel in a CID array can be individually addressed via electrical indexing of row and column electrodes. Unlike Charge Coupled Device (CCD) cameras, which transfer collected charge out of the pixel during readout (and hence erase the image stored on the sensor), charge does not transfer from site to site in the CID array. Instead, a displacement current proportional to the stored signal charge is read when charge "packets" are shifted between capacitors within individually selected pixels. The displacement current is amplified, converted to a voltage, and fed to the outside world as part of a composite video signal or digitized signal. Readout is non-destructive because the charge remains intact in the pixel after the signal level has been determined. To clear the array for new frame integration, the row and column electrodes in each pixel are momentarily switched to ground releasing, or "injecting" the charge into the substrate.

This principle of operation makes CID technology fundamentally different from other imaging techniques, giving rise to a number of technical advantages that can be used to solve imaging problems. For instance, the nondestructive readout capability of CID cameras makes it possible to introduce a high degree of exposure control to low-light viewing of static scenes. By suspending the charge injection, the user initiates "multiple-frame integration" (time-lapse exposure) and can view the image until the optimum exposure develops. Integration may proceed for milliseconds or up to several hours with the addition of sensor cooling, applied to retard accumulation of thermally-generated dark current.

CID sensors also offer wide spectral response, from 200 to 1100 nanometers, allowing capture of images produced by light sources ranging from UV to the near IR. And the PMOS structure reduces the effect of radiation on sensor operation, making CIDs less vulnerable to disruption in low-level radiation environments than NMOS devices (structure used in many CCDs). Radiation hardened CIDs are currently employed in nuclear power, industrial X-ray, scientific, and space applications.

Since each pixel in the CID array can be addressed individually, flexible readout and processing options are made possible. For example, "Progressive Scan" readout enables real-time processing by eliminating the delay required to combine odd and even fields (2:1 Interlace scanning). Instead, lines are read sequentially (1, 2, 3, 4, etc.) allowing an image processor to analyze the latest row of video information while readout continues to the next line. The 60 frames per second output of these cameras provide high-speed operation.

The SyncLite system functions by using a gated image intensifier to act as a fast shutter and amplifier for a generic CID (charge injection device) camera. This allows for the accumulation of many short-duration 'frames' during one 1/30 sec camera frame. The intensifier is operated at a gain of ~1000. The number of times the shutter is opened during a single camera frame is adjusted by the DAQ system based on the measured intensity. In the case of the abort gap, this is typically every 4th turn (~12 kHz). A LabView DAQ system collects the camera frames and fits horizontal and vertical projections of the beam profile to obtain the integrated intensity. Each abort gap measurement is the sum of 200 camera frames, or 8 x 104 abort gaps. Camera data of the abort gap are displayed in Figure 4 and Figure 5. The bump corresponds to a DC beam intensity around the ring of ~5 E9.

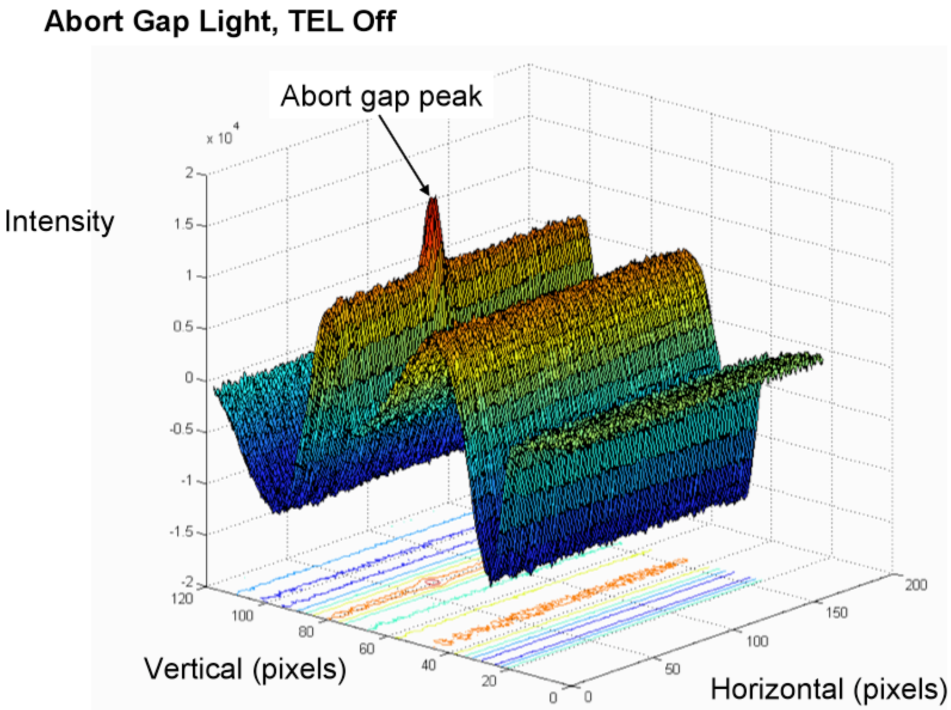

Fig. 9.20: Camera image in abort gap. This is after pixel by pixel background subtraction, but before horizontal line subtraction. The peak corresponds to a DC beam intensity of ~5 E9

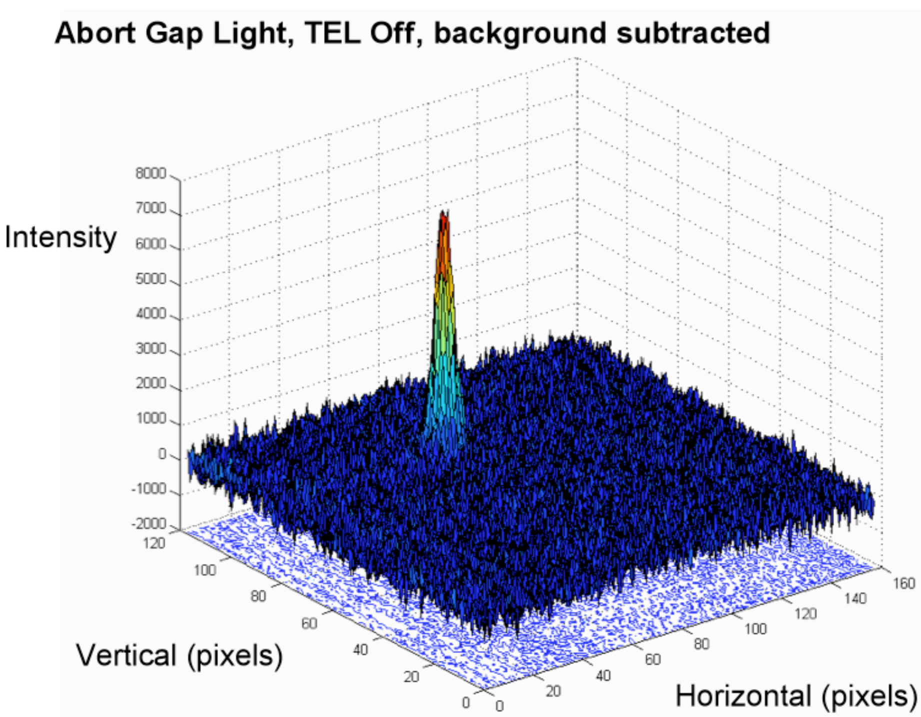

Fig. 9.21: This image is after both pixel and horizontal line subtraction. The peak corresponds to a DC beam intensity of ~5 E9

Photomultiplier Version

A 9-stage side window photomultiplier tube is attached to the SyncLite optics box and observes the light from the beam splitter. Between the beam splitter and the PMT, there is a 1% neutral density filter that can be moved in or out of the light beam.

Gated PMT

To avoid saturating the PMT with the light from the main bunches, there is a gating circuit attached to 2 of the dynodes of the PMT. The gating circuit holds the dynodes at a potential below the previous dynodes effectively turning off the tube (Figure 6). When the gate is on, the dynodes are pushed up to their nominal operating voltage.

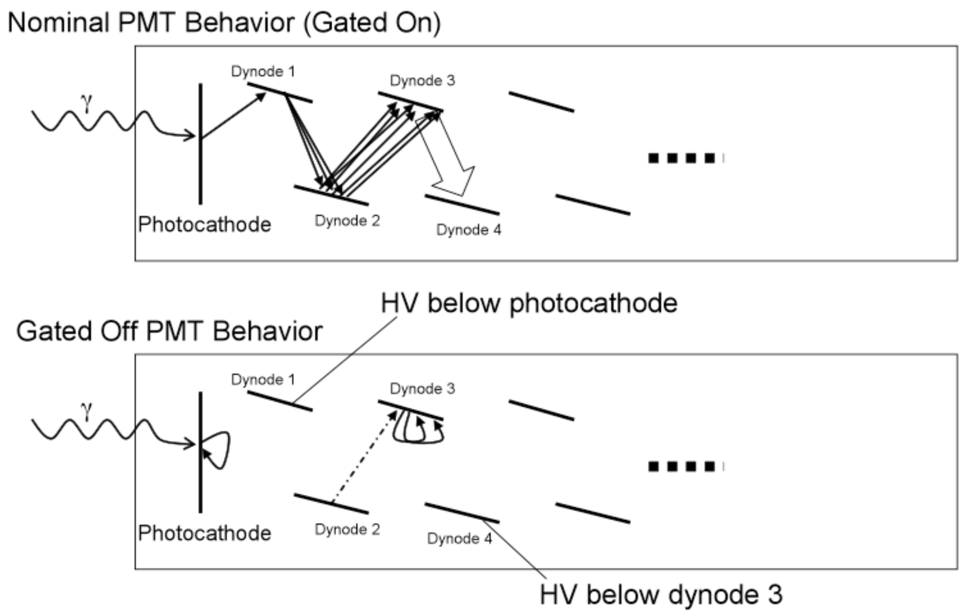

Fig. 9.21: PMT behavior for gated on and gated off modes. The gated on mode looks like an ordinary PMT, while the gated off mode turns around the field lines between several dynodes effectively shutting off the multiplication.

Data Acquisition

The DAQ system consists of an MVME board running VxWorks talking to a COMET 12-bit ADC board and a VRFT board for beam timing. The PMT anode signal is brought upstairs to a fast integrator that feeds the ADC board. The integration gate is typically 1.4 microseconds (2/3 of the abort gap). The DAQ program on the MVME performs a ~60ms readout cycle once every second (1000 samples of each abort gap).

Figures 9.22 and 9.23 shows the timing of the gating for a number of turns. Only one gate happens every turn, so over 3 turns each abort gap is sampled once.

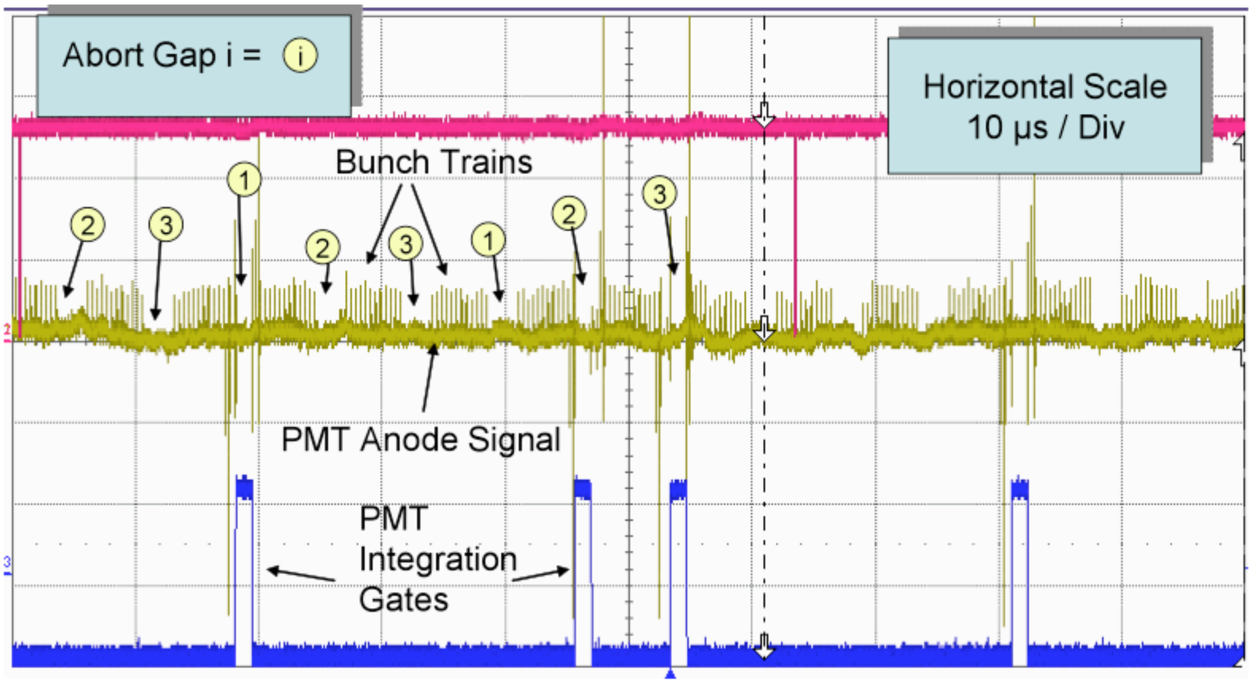

Fig. 9.22: Scope trace showing timing of PMT gates. The anode signal is into 50Ω and is displayed on a 2mV scale. The first gate from the left occurs at abort gap 1. The second gate is abort gap 2, one and one third turns later, and the third is abort gap 3. The spacing between gates is necessary to keep the duty cycle low enough for the gating circuit.

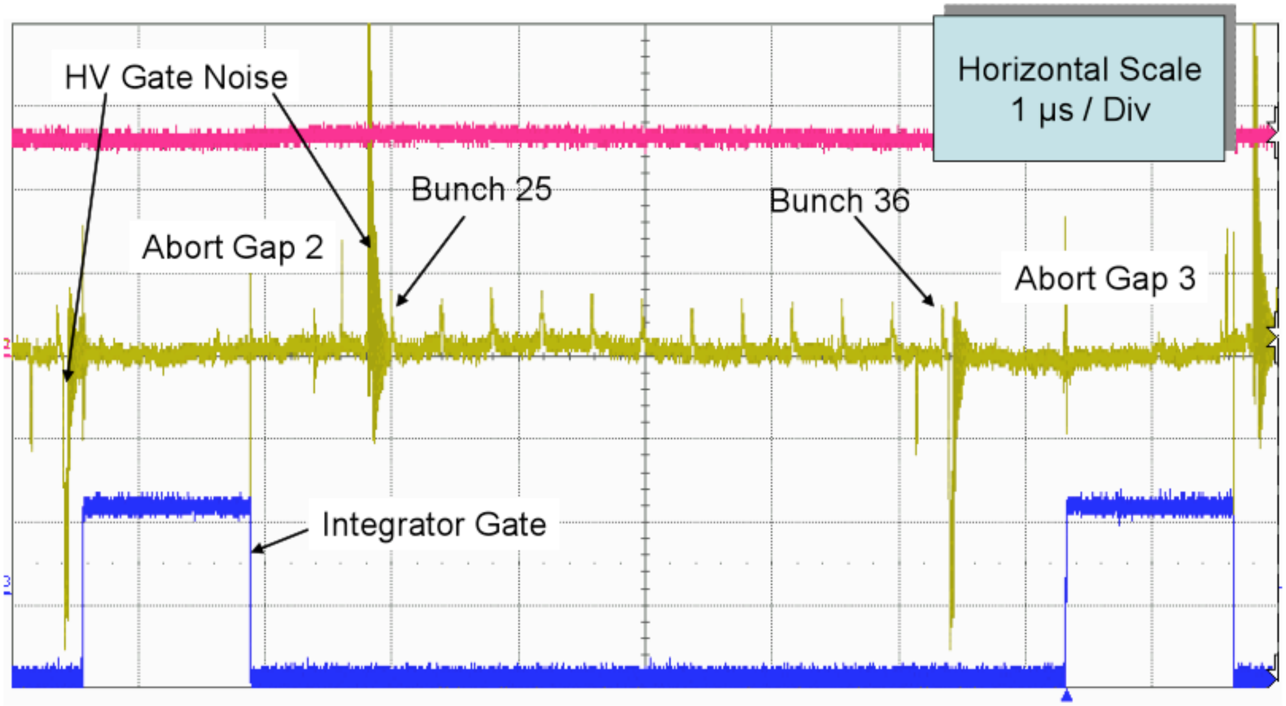

Fig. 9.23: Each abort gap has a different portion of it sampled, i.e. abort gap 1 is gated in the middle, abort gap 2 is gated at the beginning, and abort gap 3 is gated at the end

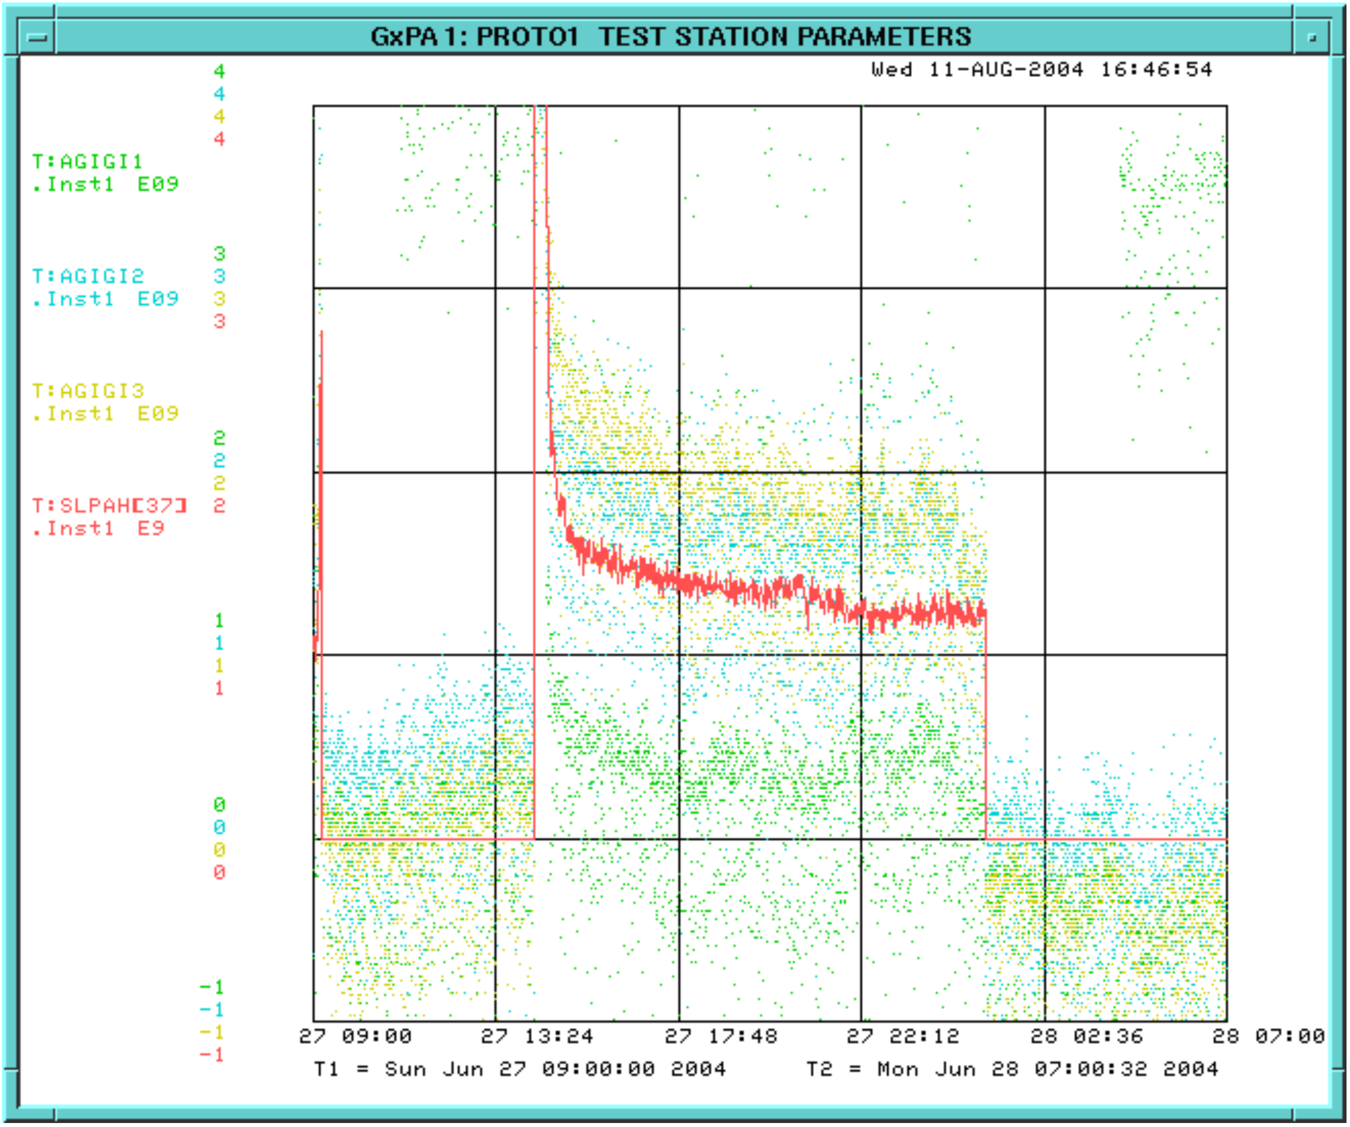

Fig. 9.24: Plots of both the PMT and the SyncLite system. The PMT system is the AGIGIx devices. The SyncLite system is the SLPAH device. One can see the shift in backgrounds for the PMT system before and after the store

10. Collider Theory

β Function

The β function is also known as the amplitude function. To understand the concept of the β function we must understand what is occurring as a particle traverses a FODO lattice. We all know that the real world does not follow ideal conditions. Accelerators are no exception. Magnets are not always constructed with perfect field configurations. For example, as magnets are ramped the laminations begin to heat up and the size of the magnet will increase, albeit a small change in volume. The field strength will change or an aberration in the field may become more apparent.

Suppose a proton is injected onto its ideal orbit in a circular accelerator made up of only dipoles and say that one dipole magnet in the ring has a momentary imperfection that causes a deflection in the plane perpendicular to the magnetic field lines. The resulting orbit would be another circle with the same radius but offset from the ideal orbit. The particle will oscillate about the ideal orbit and be considered stable because it remains in the accelerator.

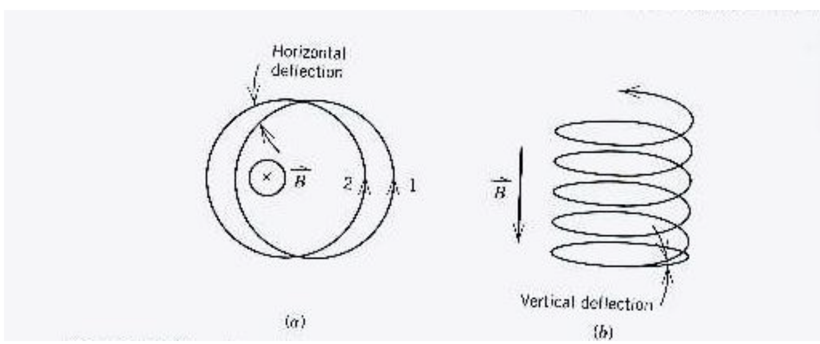

Fig. 10.1: The left diagram shows the deflection of a particle in the horizontal plane. The right figure shows the effect if the deflection has components in both transverse planes

Now consider a deflection from a magnet that has a component parallel to the magnetic field lines. The particles will begin to spiral out of the beam tube. This, of course is an unstable orbit.

The alternating gradient synchrotron was developed to provide strong focusing in both transverse planes. In the Tevatron we accomplish this with the use of quadrupoles arranged in a FODO lattice.

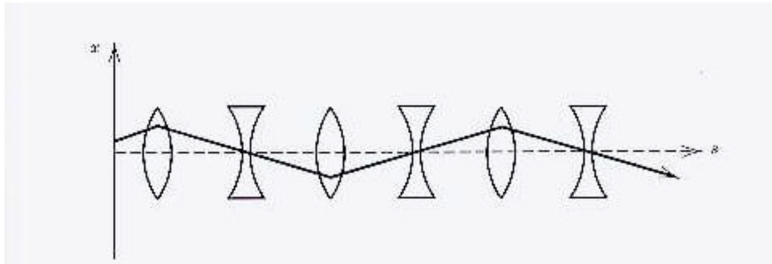

Fig. 10.2: Oscillatory motion due to strong focusing by a FODO lattice

In the above figure the motion of a particle is now periodic due to the placement of quadrupoles amongst the dipoles that keep the particle within the circumference. As with anything that is periodic, it can be compared with the solution for a simple harmonic oscillator.

$$x(s) = A\cos[\Psi(s) + \delta]$$

where A is the amplitude, Ψ(s) is the phase of the particle oscillation, and δ is the phase shift. After some manipulation via matrix applications the general form for the equation of motion of a particle traversing a FODO lattice is

$$x(s) = A\sqrt{\beta(s)}\cos[\Psi(s) + \delta]$$

where β(σ) is the amplitude function. The amplitude function is interpreted as the local wavelength of the oscillation divided by 2π. This function has units of length and is often quite a large value, on the order of meters, while the actual particle deviation from the ideal orbit is rather small.

The number of oscillations the β function goes through in one revolution of the accelerator is called the tune, ν.

Gaussian Distribution and Luminosity

In particle physics a colliding beams experiment has a great advantage over a fixed-target experiment due to the center-of-mass energy attainable for the creation of new particles.

In fixed-target the center-of-mass energy available for new secondary particle creation goes as the square root of the initial proton's energy, $E^{1/2}$. As the proton's energy is increased the gain in the secondary particles energy is small. However, in the colliding beam center-of-mass frame the proton and antiproton annihilate upon collision and give their total energy to the creation of new particles. So for a 980 GeV beam of protons and antiprotons the total available energy for new secondaries is 1.96 TeV.

There are disadvantages, too, in colliding two beams. The particles must be stable, although a muon collider has been considered due to the "long" lifetime of that particle. In the fixed-target experiments the proton collision rate with the target is high, whereas in the colliding beam experiment the collision rate is low.

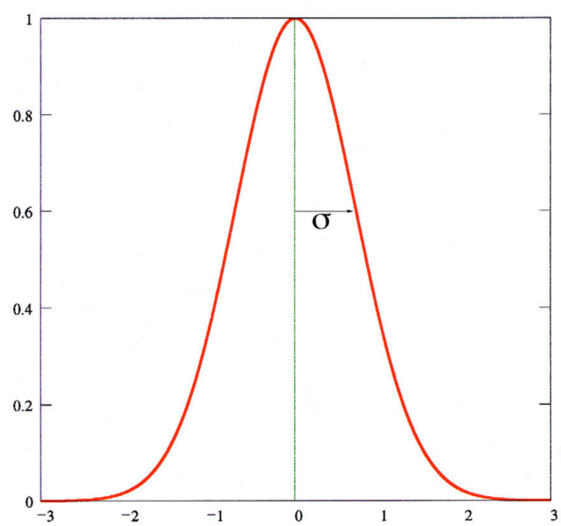

Fig. 10.3: Gaussian shape of the beam. The majority of protons/pbars reside at the correct energy and position. Some particles have higher/lower momentum and position. The vertical axis is the number of particles

The two beams as they are counter rotating in the accelerator ring have a Gaussian shape, hopefully. Refer to the picture above. Each particle has a probability of interacting with another particle traveling in the opposite direction. This is known as the interaction cross-section, σ_{int}. The rate of interaction within a detector is given by

$$R = \sigma_{\text{int}} L$$

where L is the luminosity. So what is luminosity? Lets define that now. The luminosity is a measure of how the particles in both bunches are interacting with each other. It is dependent upon the revolution frequency and the area that the beam occupies.

$$L = \frac{fnN_p N_{\bar{p}}}{A}$$

where N_p and N_{pbar} are the number of particles in each bunch, f is the revolution frequency, n is the number of bunches in either beam, and A is the cross-sectional area of the beams. Since the antiproton bunches and proton bunches can have different cross-sectional areas, A can be defined in terms of the width of the Gaussian shape, σ_p and σ_{pbar}. The luminosity in the Tevatron is defined as

$$L = \frac{f\, nN_p N_{\bar{p}}}{2\pi\left(\sigma_p^2 + \sigma_{\bar{p}}^2\right)} F\left(\frac{\sigma_l}{\beta^*}\right)$$

where f, n, N_p, and N_{pbar} are the same as defined above. The denominator contains σ_p and σ_{pbar}, which is the standard deviation of the beam spatially at the interaction point in the detector. This is just a measure of the width for the bunch. $F(\sigma_l/\beta^*)$ is a form factor (a percentage) dependent upon the bunch length, σ_l, and the beta function at the interaction point, β^*. Referring to the table on the next page, in Run II β^* is 35 cm and σ_l is 0.37m.

	RUN IB	RUN II with MI	RUN II with MI + Recycler	
Protons/bunch	2.32E+11	2.70E+11	2.70E+11	
Antiprotons/bunch	5.50E+10	3.00E+10	7.00E+10	
Total antiprotons	3.30E+11	1.30E+12	2.50E+12	
Pbar production rate	6.00E+10	1.70E+11	2.00E+11	pbar/hr
Proton emittance	23p	20p	20p	mm-mr
Antiproton emittance	13p	15p	15p	mm-mr
b*	0.35	0.35	0.35	mm-mr
Energy	900	1000	1000	GeV
Bunches	6	36	36	
Bunch length (rms)	0.6	0.43	0.38	mm-mr
Form Factor	0.59	0.7	0.7	
Typical Luminosity	1.60E+31	8.10E+31	2.00E+32	cm-2sec-1
Integrated Luminosity	3.2	16.3	41	pb-1/week
Bunch spacing	3500	396	396	nsec
Interactions/crossing	2.7	2.3	5.8	at 50 mb
Pbar tune shift	0.015	0.02	0.02	Horizontally
Proton tune shift	0.006	0.003	0.007	Horizontally

A high luminosity is what Operations strives for because it will yield a large interaction rate. By looking at the Luminosity equation above it can be seen that the luminosity increases as the intensity per bunch increases. Also, if the bunch cross-sectional area decreases then the luminosity increases. The average luminosity in Run Ib was 1.6×10^{31} $cm^{-2}sec^{-1}$. The luminosity goal for Run II is 5×10^{32} $cm^{-2}sec^{-1}$. This will be achieved due to the larger N_p from the Main Injector and an increased N_{pbar} from the Recycler.

The performance of a collider is determined by the integrating the luminosity over time. This yields units of cross section, which are units of barns (1 b = 10^{-24} cm^2).

Emittance

We are at a point where the topic of emittance can be discussed. As stated in the β function section, the solution to periodic motion through a FODO lattice is

$$x(s) = A\sqrt{\beta(s)}\cos\left[\Psi(s) + \delta\right]$$

By taking the derivative of the above equation, x'(s), and plotting its value against x(s) we obtain a phase space diagram.

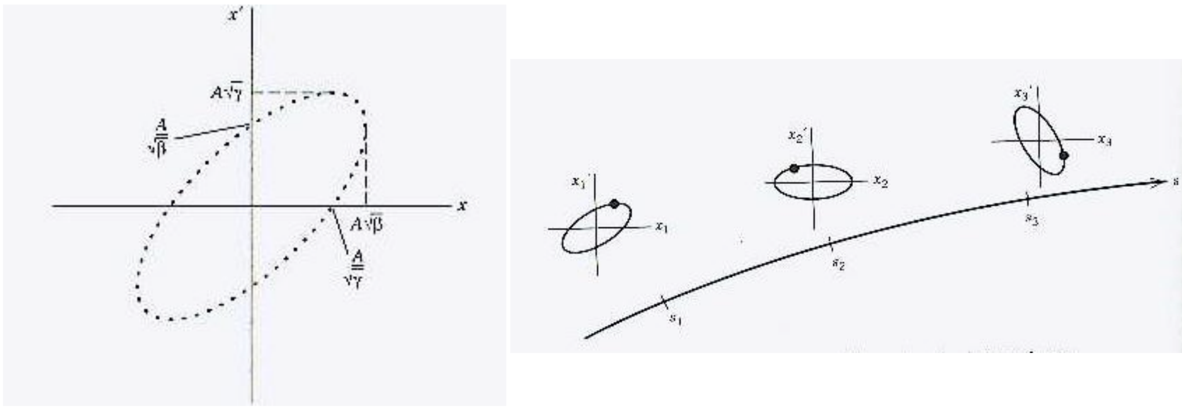

Fig. 10.4: The left figure shows the phase space diagram of a stable orbit. The right side figure shows the rotation of the ellipse throughout the accelerator. Notice the area of the ellipse is unchanged

At any point in the accelerator, the maximum value of x is $A\beta^{1/2}$. The area of the phase space remains the same but the ellipse rotates with respect to the position in the ring. The phase space occupied by the beam is called the emittance, ε. If all the particles were on the ideal orbit then the emittance would be zero because all of the particles would reside at one point on the phase space diagram. If the particles in the beam have a Gaussian distribution then the emittance is

$$\varepsilon = -2\pi\sigma^2 \frac{\ln\left(1 - {6\pi\sigma^2}/{\beta}\right)}{\beta}$$

where σ is the width of the Gaussian defined earlier in this chapter. The above equation gives the phase space that contains 95% of the beam. The units associated with emittance are mm-mr (mr = milliradians).

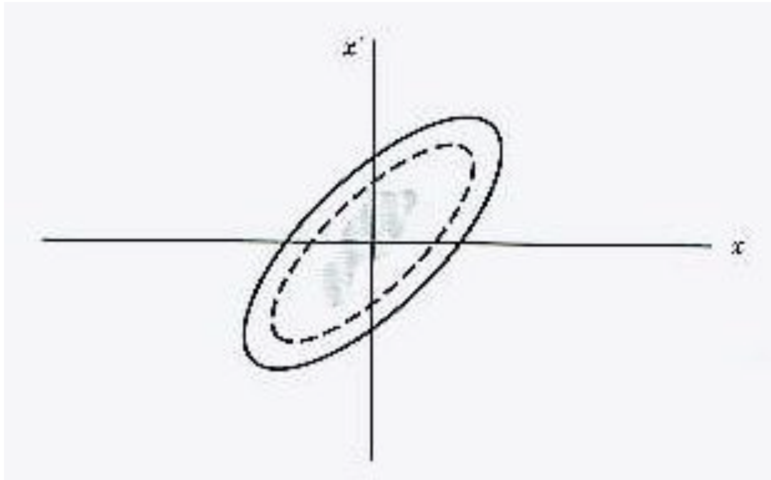

Fig. 10.5: Phase space and its relation to the beam within

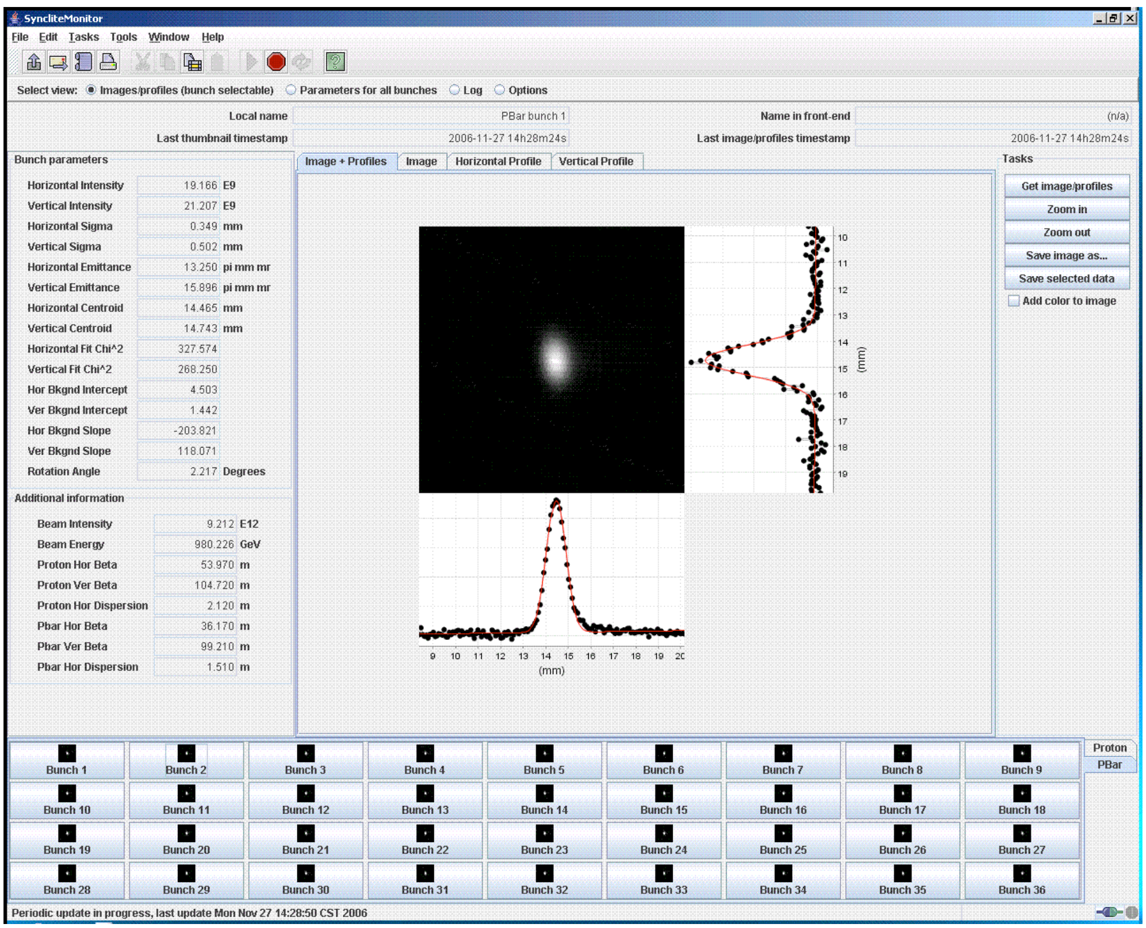

Fig. 10.6: SyncLite data for 1 antiproton bunch in the TeV at collisions. The Gaussian s is calculated along with the horizontal and vertical emittance

RF Theory

The previous section dealt with the motion in the transverse plane. Now the equations of motion in the longitudinal direction will be developed.

The progress of a particle through an accelerator can be charted via a phase space diagram of the longitudinal direction (z-axis). Let τ be the time of flight of the ideal particle passing through an RF station in one turn.

$$\tau = \frac{C}{v}$$

where C is the circumference and v is the velocity of the particle. The fractional change in τ is then

$$\frac{\Delta\tau}{\tau} = \frac{\Delta C}{C} - \frac{\Delta v}{v}$$

In relativistic terms

$$\frac{\Delta v}{v} = \frac{1}{\gamma^2}\frac{\Delta p}{p}$$

where p is the momentum and γ is $\dfrac{1}{\sqrt{1-\left(v/c\right)^2}}$.

The first term in the fractional time equation also depends on the momentum deviation. Of course, more than one particle is accelerated and statistically some will be slightly higher and lower in momentum, which implies there will be various orbits about the ideal orbit. A new parameter, γ_τ, is introduced.

$$\frac{\Delta C}{C} = \frac{1}{\gamma_t^2}\left(\frac{\Delta p}{p}\right)$$

The value of γ_τ is actually determined in the design of an accelerator. For the Tevatron, γ_τ is 18. Thus the expression of the fractional change in τ is

$$\frac{\Delta\tau}{\tau} = \frac{1}{\gamma_t^2}\left(\frac{\Delta p}{p}\right) - \frac{1}{\gamma^2}\left(\frac{\Delta p}{p}\right)$$

The term $\frac{\Delta\tau}{\tau}$ is called the slip factor, η.

$$\eta = \frac{1}{\gamma_t^2} - \frac{1}{\gamma^2}$$

Now you can see that when $\gamma = \gamma_\tau$ the sign of η transitions towards a positive number. This occurs at the transition energy. Luckily, for the TeV the beam injected is already above the transition energy. The longitudinal equations of motion can now be constructed.

Suppose a particle arrives at the η^{th} accelerating station with the energy and phase E_η and Φ_η. At the entrance to the $(n+1)^{th}$ cavity the energy and phase are $E_{\eta+1}$ and $\Phi_{\eta+1}$.

$$\frac{\Delta\tau}{\tau} = \eta\frac{\Delta p}{p}$$

The angular RF frequency, ω_{rf}, multiplied by the time of flight yields $2\pi h$, where h is the harmonic number of the TeV, 1113.

$$\frac{\Delta\phi}{2\pi h} = \eta\frac{\Delta p}{p} = \frac{\eta}{\beta^2}\frac{\Delta E}{E}$$

$$\phi_{n+1} = \phi_n + \frac{2\pi h\eta}{\beta^2}\frac{\Delta E}{E}$$

The above phase equation is one of our equations of motion in the longitudinal direction. The next equation deals with the energy of the ideal particle. Every time the particle traverses the RF cavity it gains energy,

$$(E_s)_{n+1} = (E_s)_n + eV\sin\phi_s$$

where e is the charge of an electron, V is the amplitude of the emf across the cavity's gap, and Φ_σ is the phase for arrival of the ideal particle, aka the synchronous phase. For any particle on any orbit the energy gain as it traverses the cavity is

$$E_{n+1} = E_n + eV\sin\phi_n$$

and so the change in energy between any particle and the ideal particle is

$$\Delta E_{n+1} = \Delta E + eV(\sin\phi_n - \sin\phi_s).$$

This is the second equation of motion. These equations transcribe orbits on a phase plot, ΔE vs. Φ, which show where particles in a beam are on stable and unstable orbits.

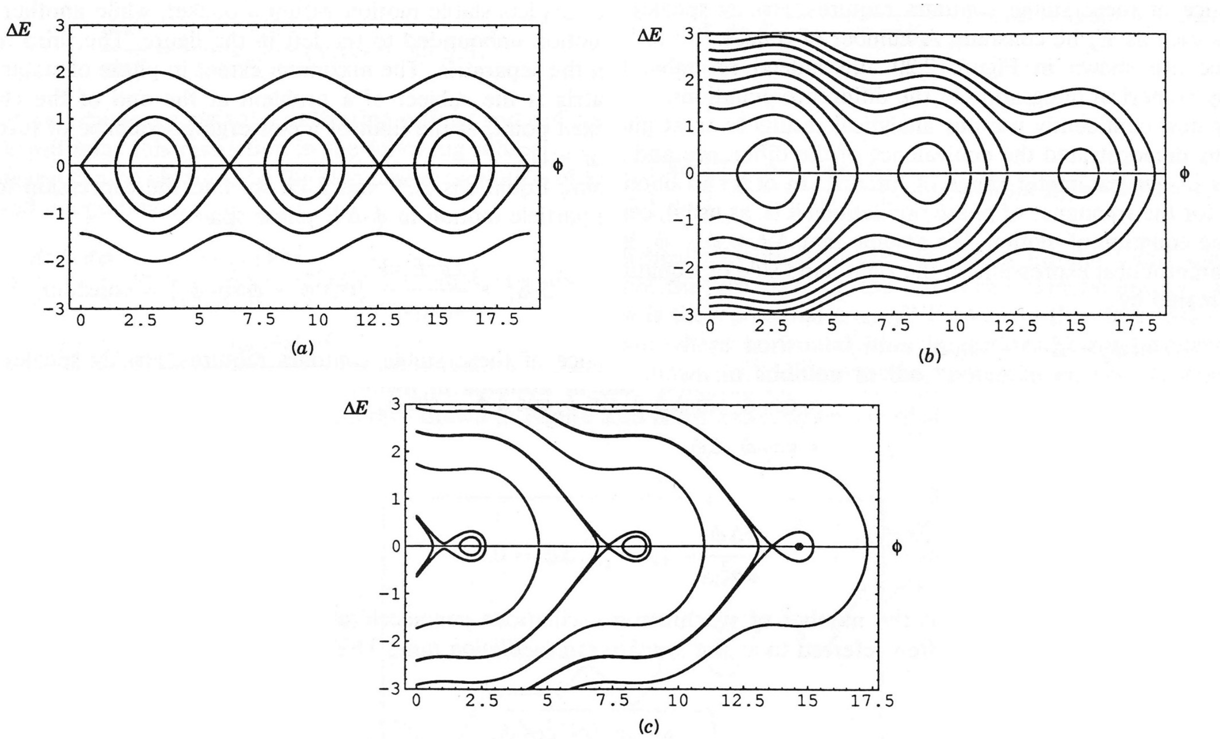

Fig. 10.7: Longitudinal phase space development. Particles are injected into a stationary bucket (a) and as the beam is accelerated (b) the phase space of the bucket shrinks until finally (c) the beam reaches the destination energy of 980 GeV

The above equations form a second order differential equation.

$$\frac{d^2\Delta\phi}{dn^2} - \frac{2\pi h\eta eV\cos\phi_s}{\beta^2 E_s}\Delta\phi = 0$$

From this equation the synchrotron frequency is found to be

$$\Omega_s = \sqrt{\frac{-h\eta eVc\cos\phi_s}{2\pi C^2 E_s}}$$

where c is the speed of light in a vacuum, C is the circumference of the TeV, and the other parameters have been previously defined. Notice from the equation that as the energy increases the synchrotron oscillations decrease. If we plug in the values for the TeV then we find

$$\Omega_s \approx 100Hz$$

$$;h = 1113$$

$$\eta = \left(\frac{1}{18^2}\right)$$

$$c = 3\times10^8 m/s$$

$$\phi_s = \pi$$

$$C = 2\pi\times10^3 m$$

$$E_s = 150\times10^9 eV$$

Φ_σ is π because the TeV is above transition.

11. Colliding Beams Mode

Chapters 1 through 8 have provided pieces to a puzzle, a large one at that. Chapter 9 brought a lot of those pieces together. For example, in chapter 2 you learned that a TeV dipole is part of a cell, which contains 8 dipoles and 2 quadrupoles. In turn, that cell is part of a FODO lattice, which repeats all the way around the ring except at certain locations. Before establishing any current in a superconducting magnet the cryogenics must bring the temperature of the NbTi bus to Tc. The satellite refrigerator building pulls helium gas from the discharge line, cools it via the heat exchanger and wet engine, and then sends it through the magnets. Once the temperature of the superconducting magnets reaches Tc, the power supplies at the 2 and 3 service buildings can supply current to the bus. The current plays out on a predetermined ramp that is loaded into TECAR from the console page C49 in the MCR. The ramp changes the field in the magnets so that the protons and antiprotons will feel the appropriate force to keep them in the beam pipe as the energy increases from 150 GeV to 980 GeV.

Lets not lose sight of what the final picture is in this mode. The results are p and pbar collisions that produce high quantities of rare, massive particles for study.

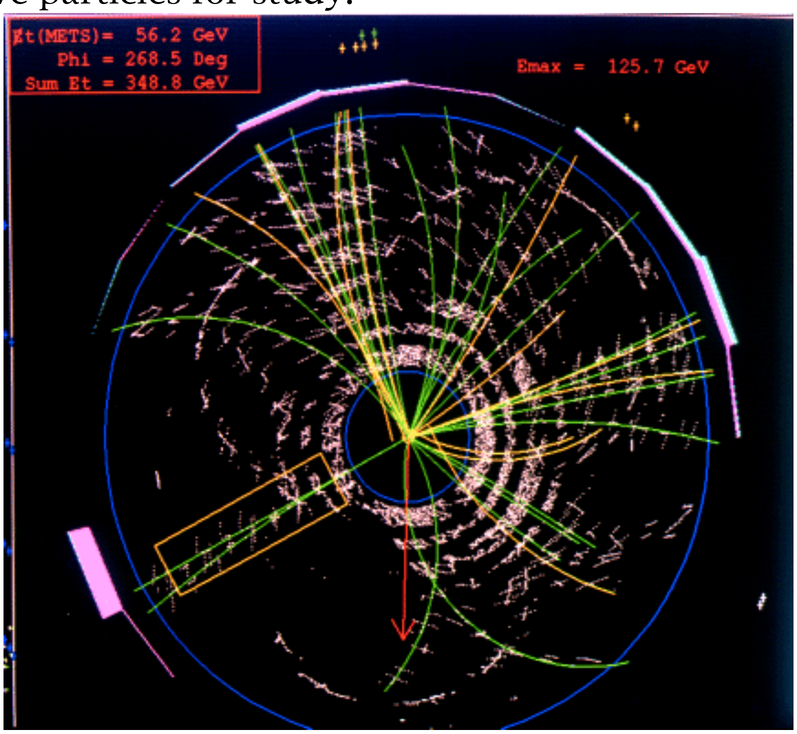

Fig. 11.1: A detected top-antitop quark event

A Shot Setup

In the collider run Ib a shot setup took about 2.5 hours to complete. Run II hopes to achieve an average shot setup of about ½ hour. The motivation is that this will increase the integrated luminosity by roughly 20%. How is this possible, you say? The plan is to automate most of the setup with sequencers.

There are 5 major steps to a shot setup:

1. TeV tune up
2. Inject proton and antiproton bunches
3. Accelerate to 980 GeV
4. Initiate a low β squeeze
5. Begin colliding and scrape away beam halo
6. Declare HEP and document store

Let's take a look at these steps individually.

1) TeV Tune Up

Before injecting protons and antiprotons destined for a store, the transfer lines into the TeV must be tuned up so that beam is placed onto the injection closed orbit with the least amount of losses so as to preserve emittance. To do this the sequencer aggregate **Decelerate, Goto Inj Porch** sets the TeV ramp at 150 GeV and sets the state device V:CLDRST to HEP shot setup mode. **Setup and Inject Protons** aggregate is next. The P1 line is tuned up first. The aggregate loads the TLG with a one-shot timeline (#192), which contains, among other things, the module Collider protons to Tevatron. Next, the TeV sequencer aggregate **Injection Closure** injects uncoalesced beam into the TeV via the event sequence of $2B $4D. Beam enters the Main Injector, accelerates to 150 GeV, is extracted to the P1 line, enters the Tevatron, and is aborted at A0.

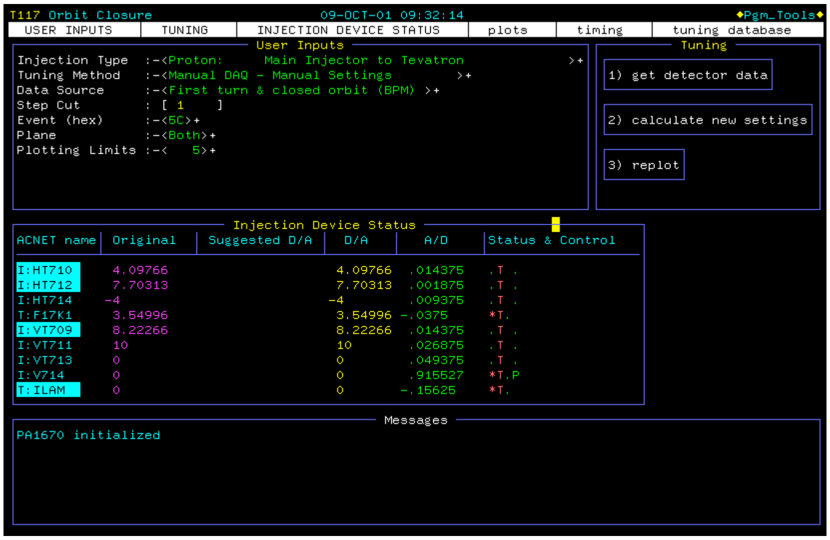

Fig. 11.2: T117 closure program for proton injection

To determine if the protons have been correctly injected onto the closed orbit, the Injection Closure program on T117 will take data on the first turn flash and the display when the $78 is broadcast. The program will calculate a change in current settings to trims I:HT710 and I:HT712 for the horizontal direction and trims I:VT709 and T:ILAM for the vertical direction, which will place the protons on the proper injection closed orbit.

Once the P1 line positions are set, the A1 line has to be tuned up. Since antiprotons are an expensive commodity, the line is tuned with reverse protons. The sequencer turns on the separators so that the injection helix is present but with the opposite polarity for mimicking the pbar helix. An event sequence of $2B $4D is initiated, which injects 150 GeV protons through the P1 line and into the TeV. Once the protons have settled onto the Pbar orbit an event sequence of $2A $5D is broadcast and the protons are extracted out through the F0 Lambertson. The protons exit the A1 line at MI-62 and go into the Main Injector. I52 is used to calculate the strengths of the I:HT902 and I:HT904 trims for the horizontal direction and I:VT901 and I:VT903 for the vertical direction so that the proper closed orbit can be attained.

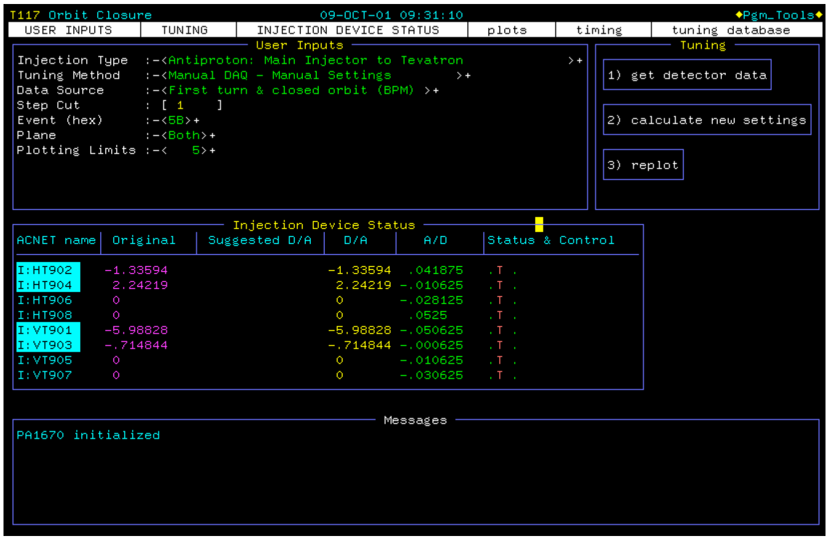

Fig. 11.3: T117 closure program for reverse injection tune up

The **Tune, Chrom, Coupling tuneup** aggregate is next. One uncoalesced bunch is injected. The chromaticity tracker is turned on and the tunes are manually adjusted to 0.590 horizontally and 0.571 vertically and the chromaticity is measured. The chromaticity needs to be set between 3-5 units in both planes.

Once the chromaticity is set, the coupling needs to be measured and adjusted. The tunes are manually adjusted to 0.586 horizontally and 0.578 vertically. Program C101 is called and the coupling is measured and adjustments are issued if desired. The tunes are then manually set to 0.583 horizontally and 0.576 vertically.

2) Inject Proton and Antiproton Bunches

With the Tevatron at 150 GeV, 36 coalesced proton bunches are injected via the P1 line. The proton bunches are injected one at a time with each bunch spaced 21 RF buckets (396 ns) apart. There are 3 “trains” of 12 bunches and between each train is an abort gap of 2.617 msec (139 buckets).

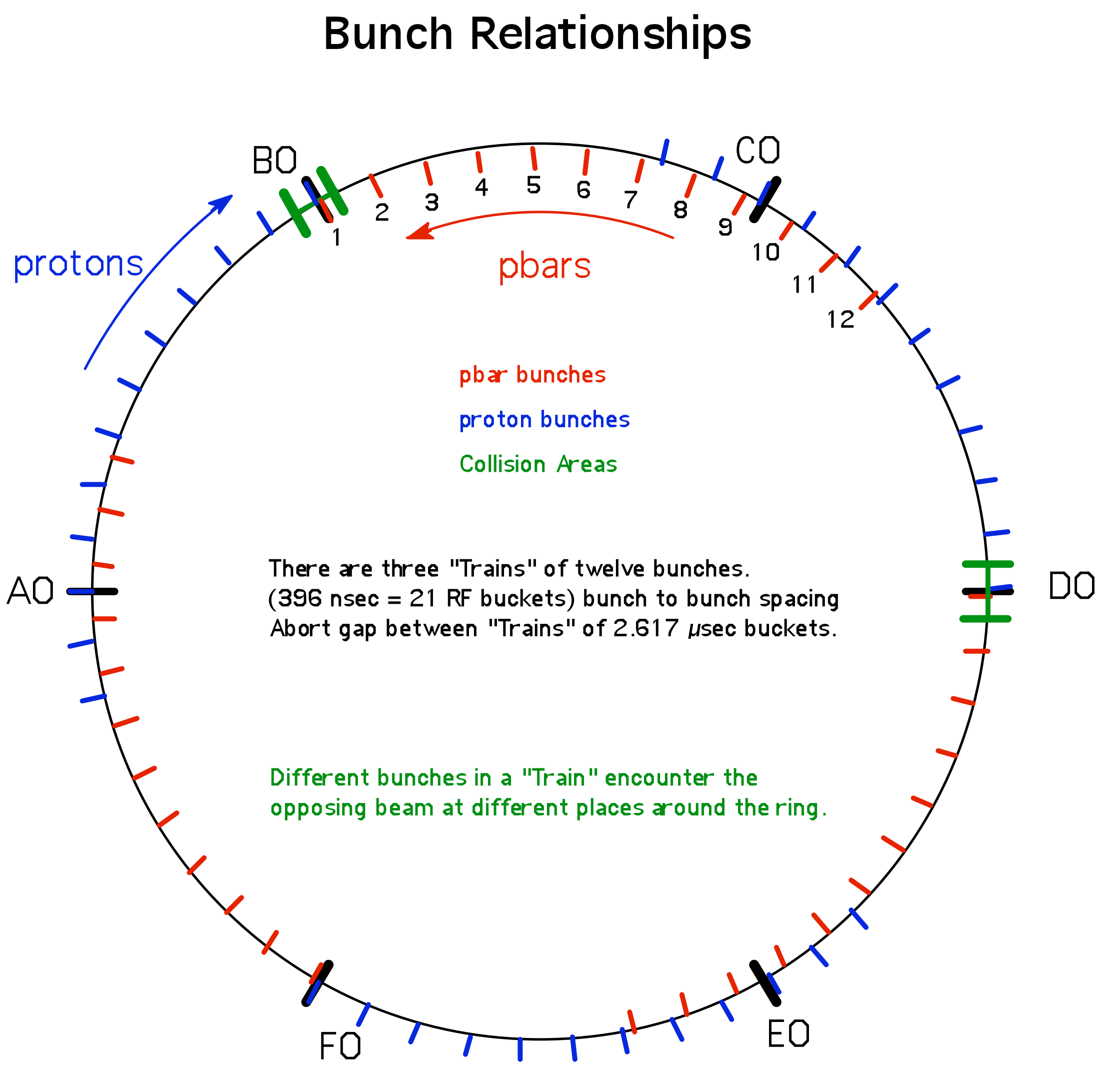

Fig. 11.4: Bunch and train spacing of the protons and antiprotons around the Tevatron.

The aggregate loads into the TLG timeline #113, #123, or #130 *Shot-RR to TeV-Proton Load-Plus XXX* (depending upon the current stacking and NuMI operational status), which contains the event sequence $2B $4D on a short supercycle so that injection of the bunches is done quickly. The proton qualifier is turned on so that bunches within a desired intensity window can be loaded. If any portion of the proton bunch is undesired, intensity, emittance, etc., then the TeV beam can be aborted and the aggregate restarted. Once the proton intensity is satisfactory, then the antiproton bunches from the Recycler can be loaded into the TeV.

To begin loading antiprotons, the **Open Helix** aggregate must be run first. This aggregate, among other things, uses a set of electrostatic separators to create a pair of non-intersecting helical closed orbits with protons on one strand and the Pbars ready to be placed on the other. The Pbar sequencer sets up the TLG to play out timeline #124, #125, or #126 Shot-RR to TeV-Pbar Load-Plus XXX. An event sequence of $2A $40 is initiated which causes 4 trains† of Recycler pbars with a 21 ns bucket spacing to be extracted to the Main Injector. The Pbars are ramped to 150 GeV, coalesced, and then injected onto the TeV helical orbit via the A1 line. Each group of 4 bunches are placed, in order, into the buckets using the following scheme: A01-A04, A13-A16, A25-A28, A05-A08, A17-A20, A29-A32, A09-A12, A21-A24, A33-A36. Each bunch is 396 ns apart.

Acceleration to 980 GeV

Before accelerating both counter-rotating beams to 980 GeV, the sequencer aggregate **prepare to ramp** removes the F0 Lambertson injection bump. The **Accelerate** aggregate then turns on the high energy lead flows, and enables a $63 that initiates the DFGs to ramp up.

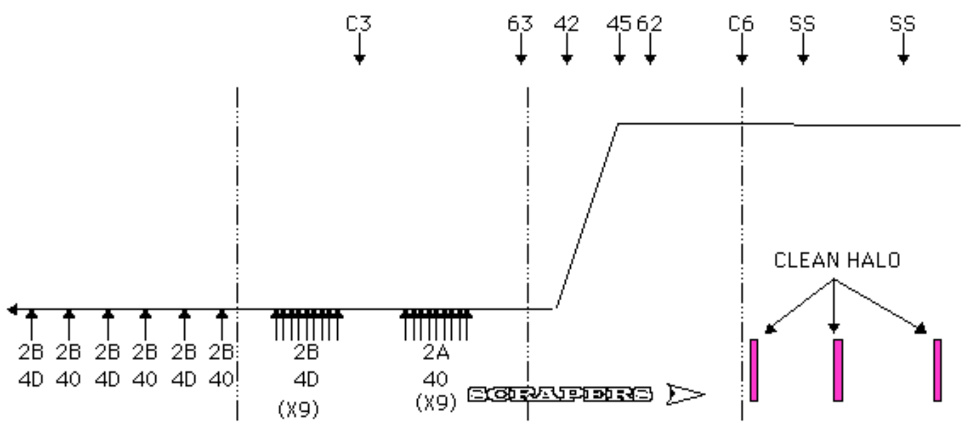

Fig. 11.5: Sequence of events used in loading and ramping to flattop

During the acceleration to 980 GeV the tunes, coupling, and chromaticity are controlled to minimize beam loss and emittance growth.

† There are 9 bunches of pbars in each train.

4) Initiate a low β squeeze and collide

After the beam has reached 980 GeV, the sequencer aggregate **Goto Low Beta** is next. The physics derivation of low beta is found in the Theory chapter.

In general, colliding beam accelerators are designed such that its constituent beams will only collide in regions monitored by detectors. In the case of the Tevatron, there are two interaction points: one in the CDF detector and one in the D0 detector. Once the interaction points have been established, the efficiency of the collider (measured primarily through luminosity) can be further raised by maximizing the number of interactions at the interaction points. This is done with the introduction of a low beta insert.

A low beta insert is a symmetric set of quadrupole magnets about the interaction point that minimize the beta function of the beam at the interaction point to a value β* for each plane. The beta function varies directly with beam size and, therefore, conversely to particle density within the beam. With very low values for β_x^* and β_y^* the beam will be much denser. As the denser beams pass through each other, the result is a higher rate of interaction.

Historically, there were ten pairs of quadrupoles in each low beta insert in the Tevatron. They can be split into two general groups of magnets: the Bartleson quadrupoles and the main quadrupoles. The main quads include Q1 – Q6 and Q8. The Bartleson quadrupoles include Q7, Q9 and Q0. Q1 and Q8 of the main quadrupole group are no longer in service.

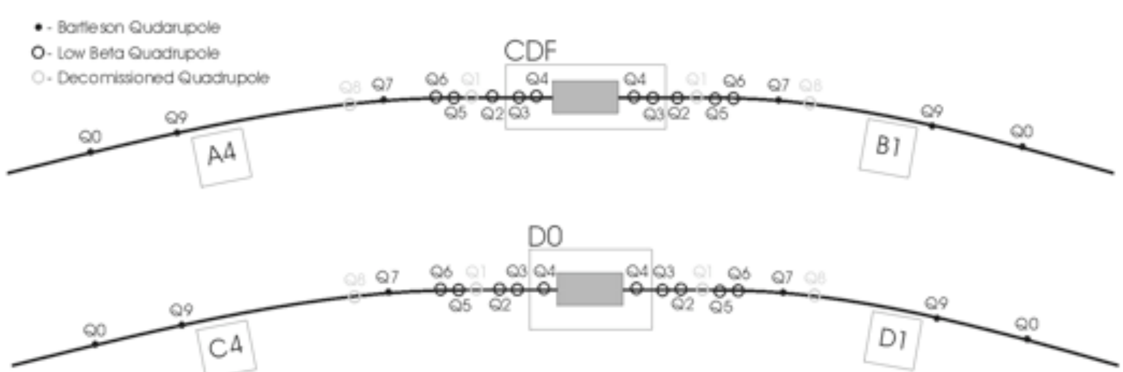

Fig. 11.6: This shows the layout of the low beta quadrupoles around each of the collision halls. Q1 and Q8 are no longer used at either collision hall

Because of the symmetric nature of a low beta insert, it is convenient to run each pair of the main quadrupoles off of a single supply. It follows that each magnet in the pair will have the same field at any given time. This is important because we want the beam to have identical qualities exiting the low beta insert as when it enters the low beta insert. As shown in figure 10.6, the low beta inserts around CDF and D0 have slightly different configurations with D0Q4 being the only low beta quad pair in the D0 collision hall whereas both B0Q3 and B0Q4 pairs are in the CDF collision hall. All other low beta quads are in the Tevatron enclosure.

The nomenclature for the supplies powering these quads includes the location and quad number. For example, Q3 around the CDF detector is powered by C:B0Q3. The supplies for each of these quadrupole pairs are located in the B0 and D0 service buildings along with their associated electronics: C468 ramp cards, current regulator chassis, HOLEC 7500 amp current transductors, and PLC electronics. It should be noted that the Q2 and Q4 pairs of quadrupoles are run off of one supply, the Q2-4 supply, at each collision hall. Because of this we only have readings for Q2, as C:B0Q2, and Q4 will be the same by design.

The Bartleson quadrupoles are in spool pieces on either side of the main low beta quadrupole string stretching from A43 to B17 around CDF and from C43 to D17 around D0. Their associated electronics are in the A4 & B1 and C4 & D1 service buildings respectively, shown again in figure 10.6.

There are also trim quadrupole and dipole supplies in the low beta strings. The trim quadrupoles do not have their own magnets with the exception of QT6, but use the main low beta quadrupoles, raising or lowering the current through one of each quadrupole pair by a small amount according to their ramps. In this way the trim quadrupoles act as verniers across one side of the main low beta quads. Nomenclature for the trim quadrupoles matches the trim with its associated low beta quad. For example C:B0QT3 acts as a vernier for C:B0Q3. Because Q2 and Q4 are powered by the same supply at each collision hall, it follows that the trims will follow the same convention. There is, for example, no C:B0QT4 since C:B0QT2 supplements C:B0Q2, which powers both Q2 and Q4 at CDF.

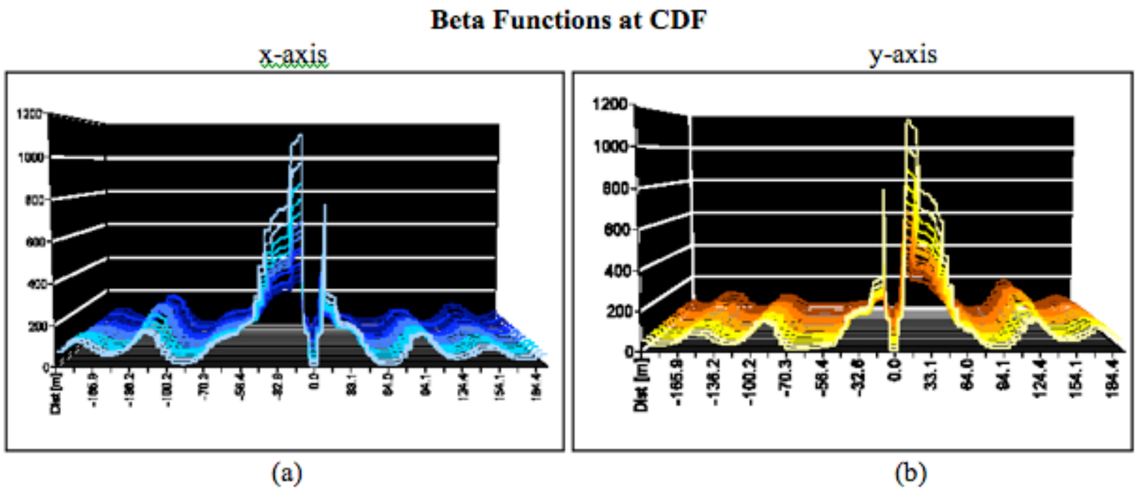

Fig. 11.7: These are the beta functions at CDF for both the x-axis through the low beta sequence as a function of distance from the interaction point. Steps of LBSEQ are plotted as progressively lighter traces. Some coupling between the beta functions can be seen. ß*, or the value for the beta function at the interaction point, is 28cm in each plane

Because of the design of the low beta insert, our beta function will get quite large before it finally arrives at its minimum, β*, at the interaction point as shown in figure 10.7. Since the beta function is directly proportional to beam size, the beam size will reach some maximum within the low beta insert before shrinking again at the collision point in the middle of our low beta insert.

The logistics of beta function manipulations makes other beam manipulations much more problematic. For this reason, most of the heavy manipulation of the beam, namely acceleration, is done prior to ramping the low beta insert to reach the desired β*. The act of ramping the low beta insert, commonly referred to as “going to low beta” or “squeezing”, is done in a sequence of steps coordinated using C:LBSEQ, the MDAT low beta sequence parameter. C:LBSEQ is broadcast on MDAT channel M13.

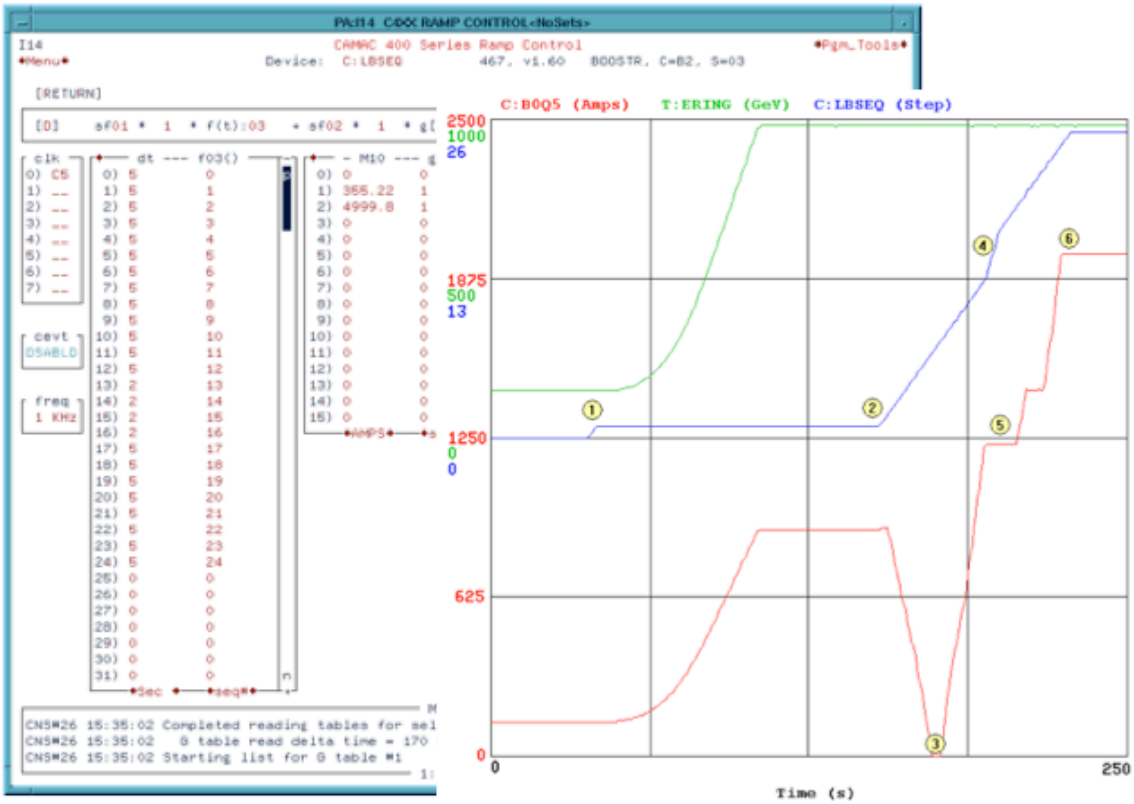

Fig. 11.8

1. When the $C4 is issued in the accelerate aggregate, LBSEQ goes to its initial value of 1.
2. Once the TeV has ramped to 980 GeV, the low beta squeeze is initiated with a $C5 and LBSEQ starts stepping through its f(t) ramp as seen on I14 (above).
3. B0Q5, for example, reverses polarity midway through its ramp.
4. The LBSEQ ramp is not constant in time, changing from 5 seconds between steps to 2 seconds briefly to allow a quicker transition from the injection lattice to collision lattice through an instability.
5. Between LBSEQ steps 10 and 16 the low beta quadrupoles do not ramp, rather, the feeddown sextupoles and separators ram to their collision lattice positions.
6. At LBSEQ step 25 the ramp to low beta is complete and the initiate collisions aggregate may be run

Once the $C5 is sent out from the sequencer, the C467 card in TeV crate $BF, located in the MCR back racks, ramps C:LBSEQ from step 1 to step 25 according to its f(t) ramp. The low beta quadrupoles, as well as feeddown sextupoles and separators, ramp with the changing C:LBSEQ step number according to their h(M13) tables. An example of the low beta sequence is shown in figure 10.8. The end result is β* goes from 1.7 m to 35 cm.

After the beam has reached low beta, a $C6 event collapses the separation bumps at the interaction points and brings both beams into collision at the detectors. The separator helix is phased so that the proton and antiproton beams only collide at the center of the detectors.

Scrape Away Beam Halo

When the beams are brought into collision and the luminosity begins to increase, the halo of protons and antiprotons needs to be scraped away in order to avoid both radiation damage to the detectors and numerous background events. The sequencer aggregate **Remove Halo** will begin to step in the scrapers while monitoring the losses from specified BLMs. Once the scraping algorithm determines the halo is no longer present the scrapers are moved out by 40 mils (1 mm).

CDF and D0 shift personnel are then informed that their detector systems can be turned on. An ACL script monitors the luminosity readbacks from both collision halls and completes the aggregate once a 10% increase in luminosity is detected.

Declare HEP and Document Store

Sequencer aggregate **HEP store** is started once CDF and D0 personnel report their luminosity high voltages are to 100%. The aggregate declares the start of a new store by setting states device HEP (V:CLDRST to 14). The time is set for the flying wires to periodically go through the beam to gather bunch structure information. Plots are also started to monitor the luminosity and losses at the detectors.

The final aggregate is document store, which collects ramp and squeeze values from the DFGs, gathers plot images, and sets the reference orbit for the orbit stabilization program.

Maintaining a Store

Orbit motion in the Tevatron has in the past contributed to lower integrated luminosity levels. Slow orbit drift on order of mm had on many occasions driven the beam into collimators causing high losses and even requiring the detectors to be turned off to prevent damage. In 2006, use of an application, which corrected the orbit periodically using two sets of correctors in each plane, became part of normal TeV operations. This system has managed to prevent the drastic orbit drift as well as keep the optics at collision fairly stable.

A study of the long-term orbit drifts showed that the dominant correctors necessary to correct the motion were the HA49 and HC49 in the horizontal plane and the VB11 and VD11 in the vertical plane as can be seen in Fig?. As a result these are the correctors used in this system.

Normally the system runs in a non-user slot on console 3. There is a boot script that automatically starts pa2119 in the non-user slot of console 3 if the console is rebooted. Once running it can be controlled and monitored via the devices listed on T55 <6> CALC VALUE.

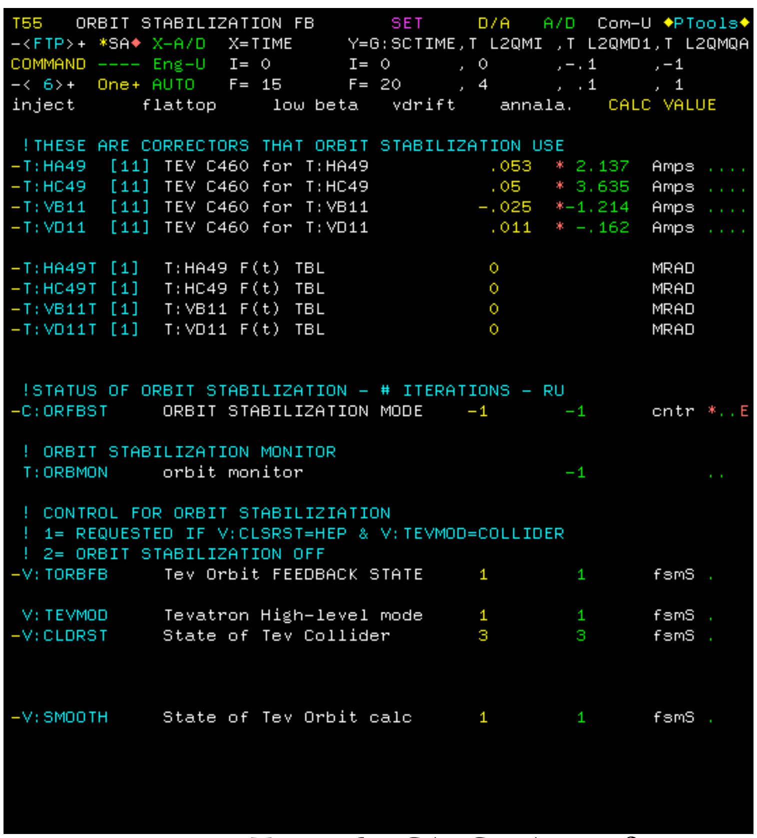

Fig. 11.9: Parameter page T55 <6> CALC VALUE for monitoring the orbit stabilization program

When the states device V:TORBFB is set to 1 it turns on orbit stabilization, likewise, state 2 turns off the program. The C:ORFBST device monitors how many corrections have been completed since the reference orbit has been updated. As well the correctors used by PA2119 are also listed.

PA2119 is currently setup to perform a single correction to an ideal orbit defined in the Tevatron Orbit subdirectory file 1, right after initiate collisions before halo removal. Then after HEP is declared the program re-initializes and begins correction to the "found" orbit every 30 seconds.

Occasionally, a bad BPM reading occurs so in-order to detect this, the following error checks are performed before any correction can be applied.

1. Each BPM reading is checked to see if it differs by more than 2 mm from the reference orbit. If it is greater then the reading is ignored and not used to calculate new correction settings
2. The calculated correction requires a change in corrector settings greater than .01 mrad then it is ignored
3. If the correction requires the total correction setting to exceed 0.12 mrad then it is also ignored since this is the corrector's current limit
4. If the total number of BPM readings which are "bad" exceeds" 10 then the whole correction pass is ignored

The program first writes corrections to the CAMAC time table for each corrector. Next an event is triggered which runs the time table and then the time table is copied to the H table. Finally the time table values are zeroed out.

On the following Page is an example of orbit and tune data taken during a store without and with the orbit stabilization on.

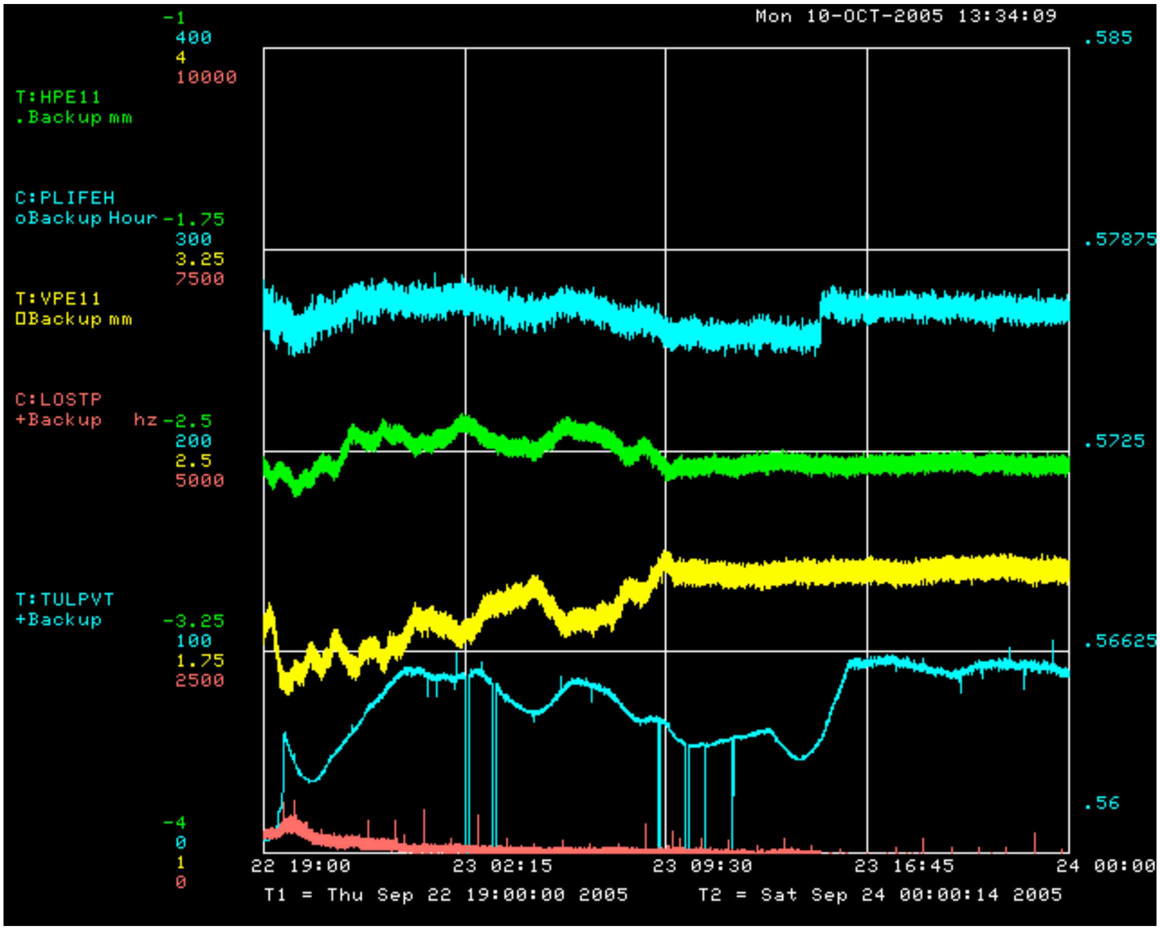

Fig. 11.10: Orbit and tune data from a store where the orbit stabilization program was turned on mid-store

Notes: