

**Fermilab**

TM-1398  
1790.000

DESIGN AND OPERATION OF THE QUENCH PROTECTION SYSTEM  
FOR THE FERMILAB TEVATRON\*

P. S. Martin

May 1986

\*Submitted to the 1984 U.S. Summer School on High-Energy Accelerators, Fermilab, Batavia, Illinois, August 13-24, 1984.

DESIGN AND OPERATION OF THE QUENCH PROTECTION SYSTEM  
FOR THE FERMILAB TEVATRON

Philip S. Martin  
Fermi National Accelerator Laboratory, Batavia, IL 60510

TABLE OF CONTENTS

1. Introduction.....	2
2. Quench Characteristics of Superconductors.....	2
Quench Threshold.....	4
Quench Propagation--Theory.....	4
Quench Propagation--Experimental Results.....	5
3. Temperature Increase in Adiabatic Limit.....	7
Comparison with Experimental Results.....	8
4. Design of the Fermilab Quench Protection System.....	8
Layout of the Tevatron Superconducting Accelerator.....	8
Quench Protection System.....	12
Quench Detection System.....	14
5. Operational Experience at the Tevatron.....	17
Failures During Operation.....	17
Other Aspects of Operating a Superconducting Accelerator.....	21
6. References.....	24

## 1. INTRODUCTION

The operation of a superconducting accelerator, in addition to cryogenic requirements, introduces a new complexity not present in a conventional accelerator. A method is required to protect the magnets from possible overheating or overvoltage conditions in the event that some magnets quench, that is, are elevated in temperature such that they no longer are superconducting. The development of that system is the topic of this lecture.

There are two very important ingredients to any quench protection system. First, the system must be designed with sufficient integrity to remain functional even under abnormal circumstances. The magnets must be protected during power failures, for example. Quenches involving a large number of components can also be hazardous due to the redistribution of voltages during the quench. Some of the system integrity can be achieved through redundancy. Frequent testing of critical elements of the system also assures the overall integrity. Second, the quench protection system must protect against damage from quenches regardless of their location or the excitation current at the time. It is not sufficient to protect just the magnet coils; the leads between magnets must either be fully stabilized or the quench protection system must protect them.

The next section presents a brief discussion of the basic properties of superconductors and the phenomenon of quench propagation. A more complete discussion can be found<sup>1,2,3</sup> elsewhere.

## 2. QUENCH CHARACTERISTICS OF SUPERCONDUCTOR

For a given superconductor, there is a surface in the space defined by temperature, magnetic field and current density which marks the boundary between the superconducting and normal states. In accelerator magnets, in which large quantities of superconductor are required, cost constraints force the magnet designers to minimize the amount of superconductor and therefore the magnets operate as close to the boundary as possible. It is often convenient to reduce the surface in three dimensional space to a set of curves of current density versus temperature, with each curve representing the critical current density as a function of temperature at a fixed magnetic field. For any particular magnet design there is a transfer function relating the excitation current and the resulting magnetic field. This reduces the set of curves to a single curve, (Figure 1), which gives the maximum operating current as a function of temperature. At some particular operating current density, there is similarly some critical temperature above which the superconductor becomes resistive. The superconductor used in the Tevatron was specified to have a critical current density of 1800 A/mm<sup>2</sup> at 5 T and 4.2 K, for the individual strands. There is some degradation in the cabling and winding process. In recent years, there has been substantial improvement in the quality of superconductor, and current

densities in excess of  $2500 \text{ A/mm}^2$  are now obtainable, although industry has not yet produced superconductor with that current density on a large scale.

The transition from the superconducting state to the normal state might occur through the elevation of the conductor to a temperature above the critical temperature, perhaps due to beam losses. Or it might occur at a particular temperature, if the current and field are increased beyond the critical current  $I_C$ . The magnetic field seen by the superconductor is not the same as that in the beam aperture. The conductors see an entire range of fields from zero up to a peak field which might exceed that in the beam pipe. Therefore, the concepts of critical current and critical temperature are local properties.

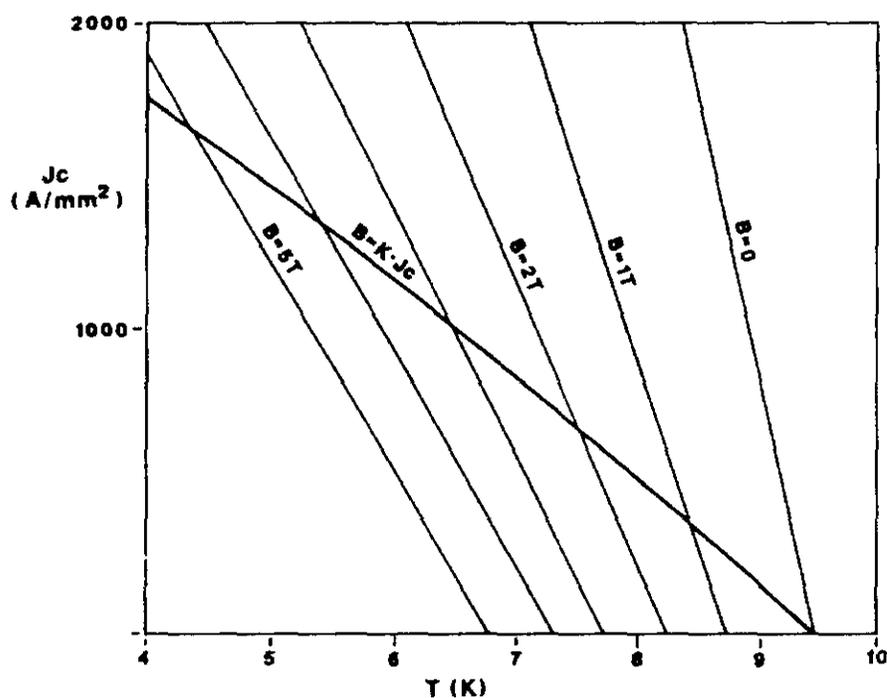


Figure 1. Curves of critical current density vs. temperature for fixed magnetic fields and for a magnetic field proportional to  $J_C$ .

The Fermilab magnets are constructed of superconductor consisting of 8 micron filaments of niobium-titanium imbedded in a copper matrix. Twenty-three strands, .068 cm in diameter and containing 2100 filaments each, are formed into a keystone cable<sup>4</sup>. While much of the discussion here applies specifically to this cable, it can be easily generalized to apply to magnets composed of other types of superconductor and other sizes of cable.

The copper matrix serves two important functions. First, it stabilizes the superconductor by providing an alternate current path, allowing small regions of the superconductor to become resistive for short periods of time and then recover as they are cooled back to the superconducting state by heat exchange with the environment. Second, in the event a larger section of superconductor becomes normal and cannot recover, it prevents overheating. Niobium-titanium, like other superconductors, is a rather poor conductor once it has exceeded the critical temperature, and the ohmic heating in the absence of the copper would quickly destroy the magnet.

### Quench Threshold

Imagine a superconducting magnet, operating at some current and temperature  $T_b$  less than the critical temperature  $T_c$  (at that current and magnetic field). The superconductor will make a transition to the normal (resistive) state if the temperature is raised above  $T_c$ . The energy required to raise the temperature from  $T_b$  to  $T_c$  is determined by the enthalpy of the constituents. For metals at low temperatures, the specific heat varies like  $T$  to the third power (plus a linear term which is comparable in magnitude at 5 K). For copper, niobium and titanium, the specific heats are in the range 1 to 5 mJ/cm<sup>3</sup>K. Cables like the Fermilab cable contain about 10% open area between the strands. That space is filled with liquid helium, whose specific heat is far larger than the specific heat of the conductor itself. The details of how the helium participates is not clear; further, it may well depend upon the process of interest. For example, fast energy deposition, due to beam loss, is certainly different from slow processes such as eddy current heating. In the case of quench propagation, which we will be discussing in the next section, it seems reasonable to assume that the helium in the normal region has been vaporized. It then expands and displaces the liquid helium surrounding the still superconducting cable. Even under these assumptions, the enthalpy of the gaseous helium is an order of magnitude higher than that of the cable constituents. Experiments have demonstrated<sup>5</sup> that the energy required to initiate a quench varies roughly like  $1/I$  over the range .1 to .9  $I_c$ . At .9  $I_c$ , the energy required is about 4 mJ.

### Quench Propagation-Theory

If some small region of superconducting cable has been elevated in temperature so that it becomes resistive, the region will expand (or contract) according to the balance between the power generated by the resistive section and the heat absorbed by the liquid helium.

The heat equation in one dimension is

$$\frac{d}{dx} \left( k \frac{dT}{dx} \right) - c \frac{dT}{dt} - \frac{hP}{A} (T - T_b) + J^2 \rho = 0$$

where  $k$  is the thermal conductivity and  $c$  the specific heat of the superconductor,  $h$  the coefficient of heat transfer,  $P$  the cooled perimeter,  $A$  the area of the superconductor,  $T_b$  the helium (bath) temperature, and  $J^2 \rho$  is the power density in the normal zone. By transforming to a coordinate system moving with velocity  $v$ , one obtains an equation for the motion of the normal zone. The solution to this equation generally is obtained after making a number of assumptions. The thermal conductivity and specific heat are taken to be constants, and the temperature of the normal region is assumed to reach a thermal equilibrium. While these assumptions are oversimplifications, they allow a solution to the equation. That solution takes the form

$$v = \frac{J}{c} \sqrt{\frac{\rho k}{(T_c - T_b)}}$$

The formula predicts a velocity which grows slightly faster than linearly with current, due to the current-dependence of  $T_c$ . The  $T_c$  term also introduces a B-field dependence. As mentioned earlier, the specific heat of the helium in the Fermilab type cable dominates that of the metallic constituents. Thus, knowledge of details of how the helium participates is necessary if one wishes to calculate quench velocities, either analytically or numerically. In the latter case, one need not make such simplifications as constant specific heats and thermal conductivities.

#### Quench Propagation--Experimental Results

The design of the Fermilab quench protection system relies in part on a set of measurements<sup>6,7</sup> made on a short (17 cm) sample of superconductor mounted in a "hairpin" fixture (Figure 2). A heater was located at one end, and two voltage taps were installed for monitoring the resistance of the sample. A typical voltage signal, after inducing a quench with the heater, is shown in Figure 3. There are two distinct parts to the voltage growth. The first part, the very nearly linear rise in voltage, is due to the propagation along the sample. Measurements at a variety of currents and with different magnetic fields yielded a formula for the quench velocity

$$v = .36 I^2 (1 + .077 B^2)$$

where  $I$  is the current in kiloamperes,  $B$  the field in Tesla, and the velocity is given in meters/sec. This formulation holds for currents from 1 to 4 kA.

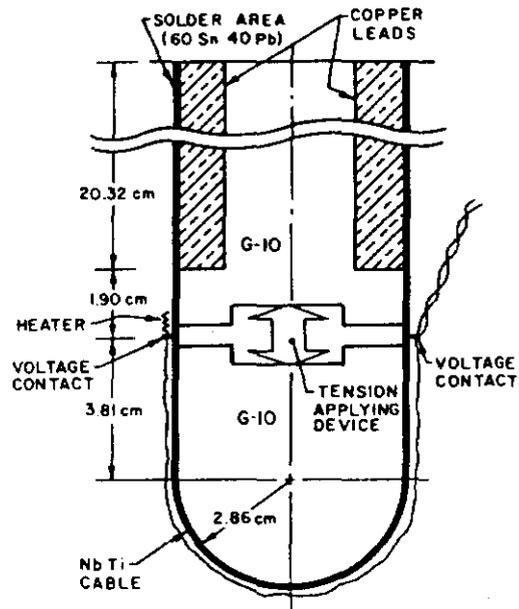


Figure 2. Hairpin fixture used for studying quench velocity and conductor heating.

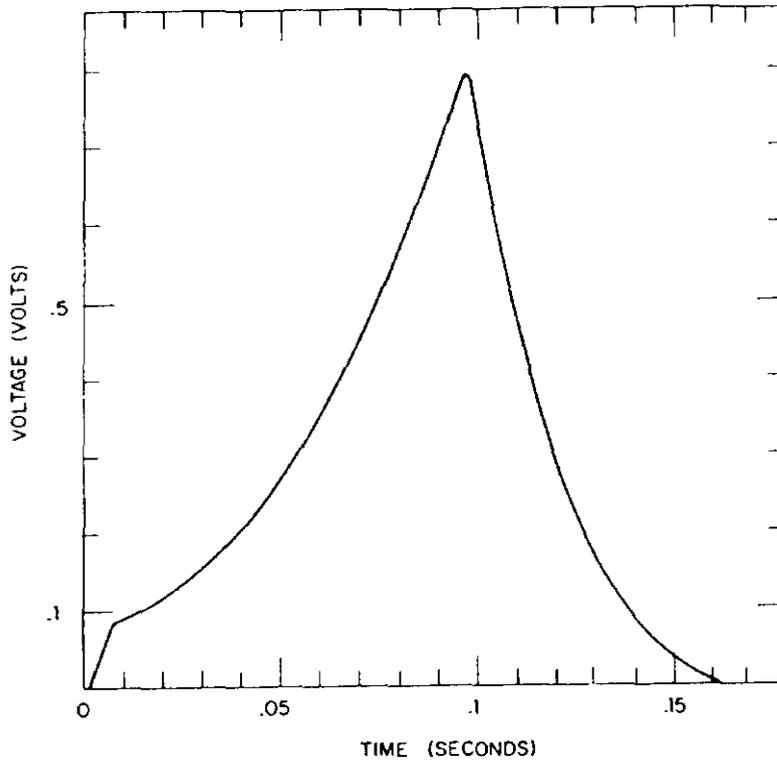


Figure 3. Typical voltage pattern observed in hairpin fixture. The power supply was turned off at .095 sec.

### 3. TEMPERATURE INCREASE IN THE ADIABATIC LIMIT

Once a section of conductor has quenched, the Joule heating will start increasing its temperature. If we neglect any heat transfer, the heat equation (setting  $k$  and  $h$  to zero) simply becomes

$$c \frac{dT}{dt} = J^2 \rho$$

Let us work in terms of a unit volume of conductor. Then

$$J^2 \rho = I^2 \rho / \left( A^2 \frac{r}{r+1} \right)$$

where  $r$  is the ratio (by volume) of copper to superconductor and  $\rho$  is the resistivity of the copper matrix. Taking the resistivity to the other side and integrating, one obtains

$$A^2 \frac{r}{r+1} \int_{T_0}^T \frac{c}{\rho} dT = \int_0^t I^2 dt$$

The integral on the right is often referred to as the MIITs (mega-amp-squared-seconds). In the adiabatic limit, the number of MIITs

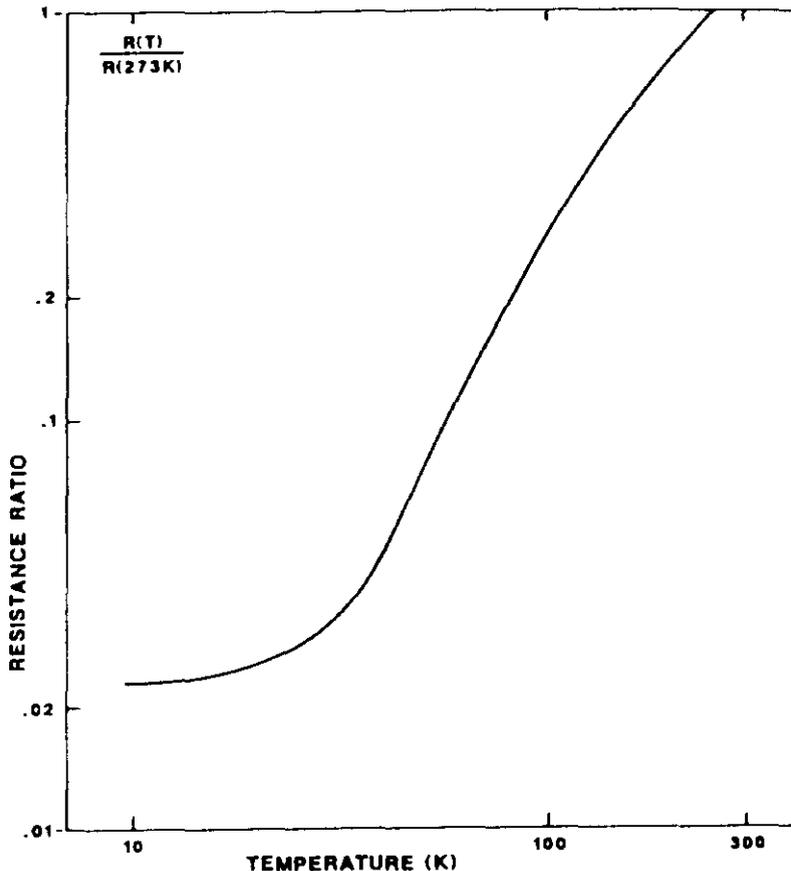


Figure 4. Resistance ratio vs. temperature for samples of Fermilab cable (from Reference 8).

determines the peak temperature in the superconductor. Since we have neglected heat transfer, this is a worst-case calculation. The peak temperature will be somewhat less than that predicted by this formula.

#### Comparison With Experimental Results

Measurements<sup>8</sup> of the resistance vs. temperature for samples of superconductor are shown in Figure 4. The residual resistivity ratio (RRR), that is the ratio of resistivity at 273 K to that at 4 K, is approximately 40 for cables from several different sources. The residual resistivity is a result of impurities in the copper. Above about 40 K the resistance of the sample can be used as a thermometer. The second part of the voltage growth in Figure 3 is due to the increase in temperature (and resistance) of the conductor. This data has been replotted in Figure 5 for several different currents. Changing variables, instead of voltage vs. time, we plot temperature vs. MIITs, Figure 6.

The Fermilab magnets are interconnected using solder, and half of the strands are coated with solder to achieve a balance between proper current-sharing and eddy currents. Therefore, the melting point of solder--450 K--has been selected as the maximum temperature to be allowed during a quench. From Figure 5 we see that this corresponds to roughly 7 MIITs at 4 kA. That is, within one-half second after a quench begins, the current must be reduced to essentially zero. The selection of 450 K is obviously somewhat arbitrary. Other limits worth noting are (i) insulation (Kapton, Mylar) degradation begins around 600 K, and (ii) superconductor degradation occurs above 800 K. Damage to materials like Kapton, however, depends on duration of the temperature level and upon presence of oxygen in the environment.

#### 4. DESIGN OF THE FERMILAB QUENCH PROTECTION SYSTEM

Quench protection implies primarily the removal of the current from a quenching magnet before it overheats. To accomplish this requires two activities: (i) detecting the quench before it's too late, and (ii) taking the necessary action. First let us discuss the configuration of the Tevatron magnets.

##### Layout of the Tevatron Superconducting Accelerator

The magnets in the Tevatron are configured as shown in Figure 7. The magnets and power supplies form a single series circuit consisting of an "upper" and "lower" bus connected at the B0 straight section. Each power supply is capable of ramping to 4500 A at 1 kV. One supply (A2) is the current regulating supply, and is capable of dc operation. The other eleven supplies operate in voltage regulation mode. The uniform spacing of the power supplies minimizes the peak voltage to ground. A .25 ohm "dump" resistor is also located at each power supply. When a quench or other fault is

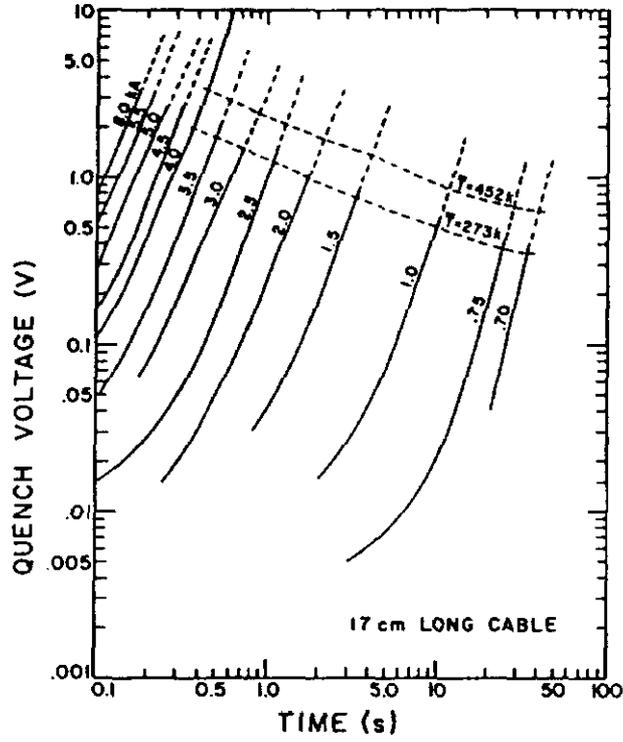


Figure 5. Voltage vs. time from hairpin measurements at different currents.

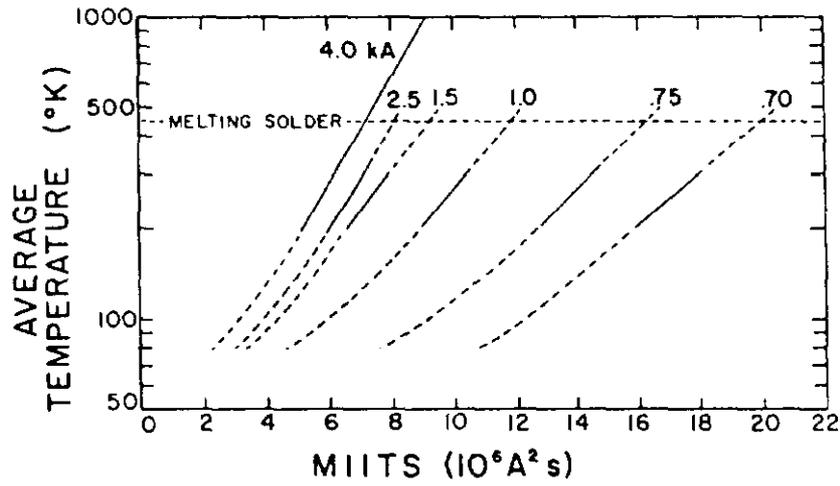


Figure 6. Temperature vs. MIITs.

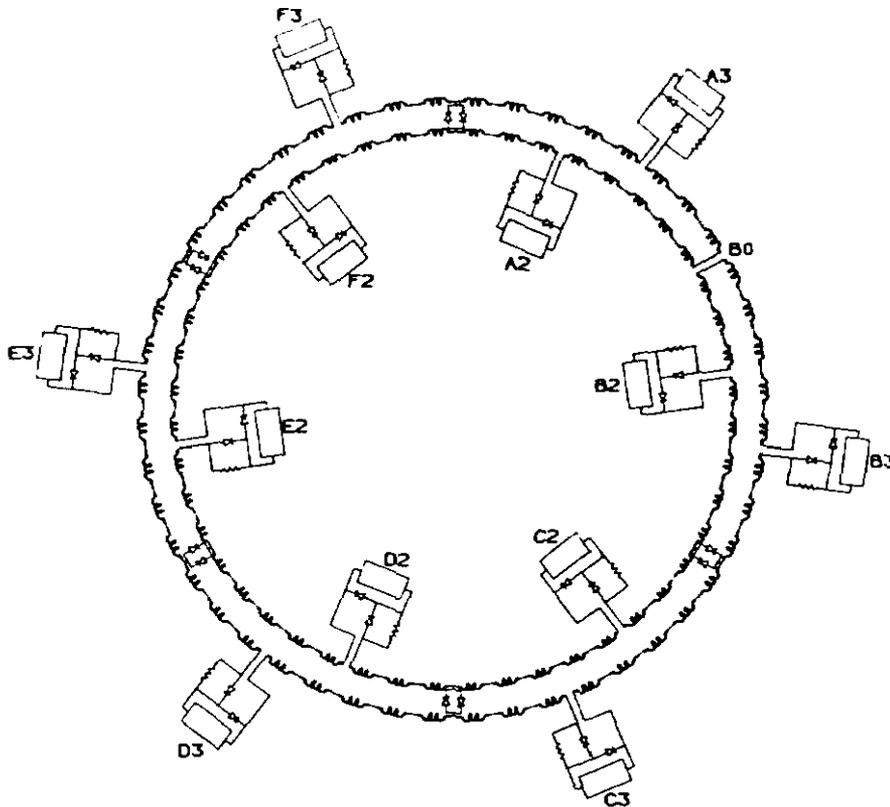


Figure 7. Schematic of Fermilab Tevatron

detected, the power supplies are turned off and the switches which normally conduct current around the dump resistors are opened. The current in the bus then decays with a 12 second L/R time constant. The dump switch is an SCR backed up by a dc contactor. A schematic of a power supply-dump rack is shown in Figure 8.

The normal operating conditions of the Tevatron permit voltages up to 1 kV with a single ground fault. To maintain this during a dump requires an inductive cancellation of the voltage provided by the dump resistors. When magnets quench, their inductance is removed from the circuit, and if many magnets quench, substantial redistributions of voltage can occur. At each straight section (other than B0) the upper and lower busses are connected with a bipolar thyristor switch which is closed just prior to opening the dump switches. This decouples the sectors from one another during the dump, and prevents potentially destructive voltages from arising should many magnets quench simultaneously.

Figure 9 shows the layout of a typical lattice cell of the Tevatron. It consists of eight dipoles and two quadrupoles arranged

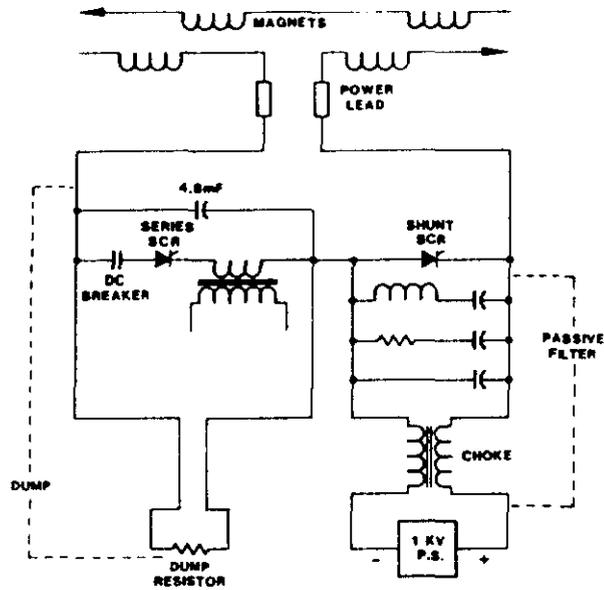
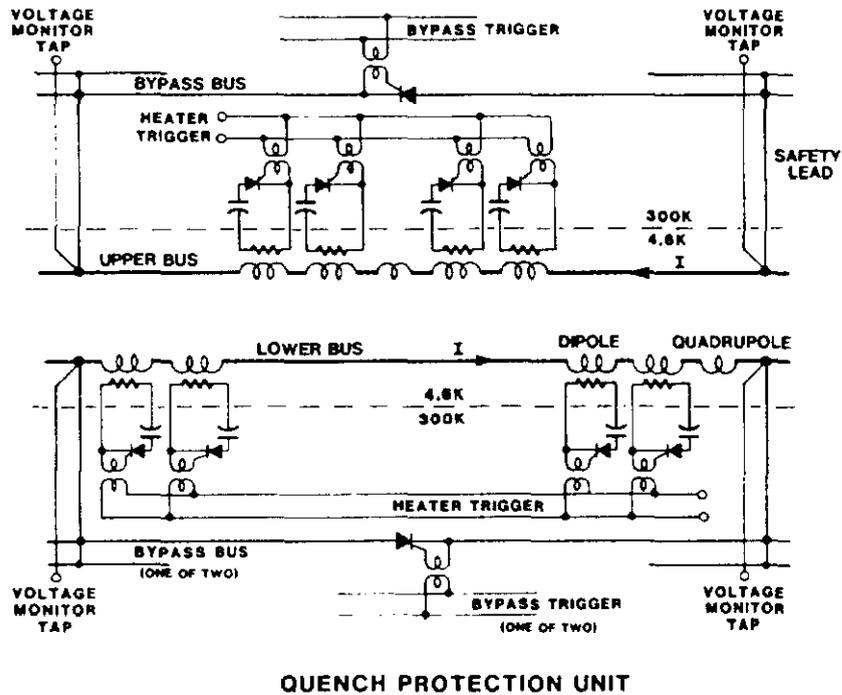


Figure 8. Power supply/dump rack schematic



QUENCH PROTECTION UNIT

Figure 9. Typical cell layout

in an "alternating bus" configuration. 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 dipole, D quadrupole). The inductance of a typical "half-cell", that is, of one bus of one cell, is about .2 H. The nominal ramp rate of the Tevatron is about 250 A/s, so the voltages observed are on the order of 50 V.

### Quench Protection System

The Fermilab quench protection system is an active one; it requires prompt detection of a quench and active components to remove the current from the quenching magnets. The heart of the system is the Quench Protection Monitor (QPM), a microprocessor which monitors the voltages across the cells at a 60 Hz rate. If a quench is detected, heater firing units (HFUs) are discharged into all the dipoles in the cell which is quenching. This drives a large fraction of the cell normal, resulting in a large resistance. Rather than absorbing the stored energy of the half-cell (2 MJ at full current) in the area of the quench origin, that energy will be absorbed by all the coils in the half-cell. At the same time, quench bypass switches (QBSs) are gated on to allow current to bypass the cell. The switches and cables which carry the bypass current are at room temperature and are connected to the superconducting bus through "safety leads" located in the spool pieces on either end of the cell.

The rate at which current is bypassed depends on the quench resistance and the inductance of the half-cell. Shortly after the heaters are fired, the dump resistors are switched into the circuit. This reverse-biases the QBSs, and no current can be bypassed until the resistive voltage reaches the voltage due to the dump, about 50 V. The current in the cell then begins to decay according to the equation

$$L \frac{dI}{dt} + I R(t) = V_{qbs}$$

where the resistance and SCR drop in the bypass circuit determine the term on the right. The time-dependence of  $R$  is shown to emphasize that this is not a simple LR circuit with an exponential decay.  $R$  is increasing very rapidly with time, as the heater-induced quench propagates and the coil heats up. The maximum  $dI/dt$  exceeds 10 kA/s which implies voltages of over 2 kV. Because the resistive and inductive voltages overlap in space, there is a cancellation, and only modest voltages are observed within a quenching cell (on the order of 200 V across an element.) The term  $V_{qbs}$  is small and also time dependent. The resistance of the bypass circuit is on the order of .010 ohm. That small resistance is useful; it allows one to determine the bypass and magnet currents as a function of time, and from that, the MIITs. The magnet current during a quench is shown in Figure 10.

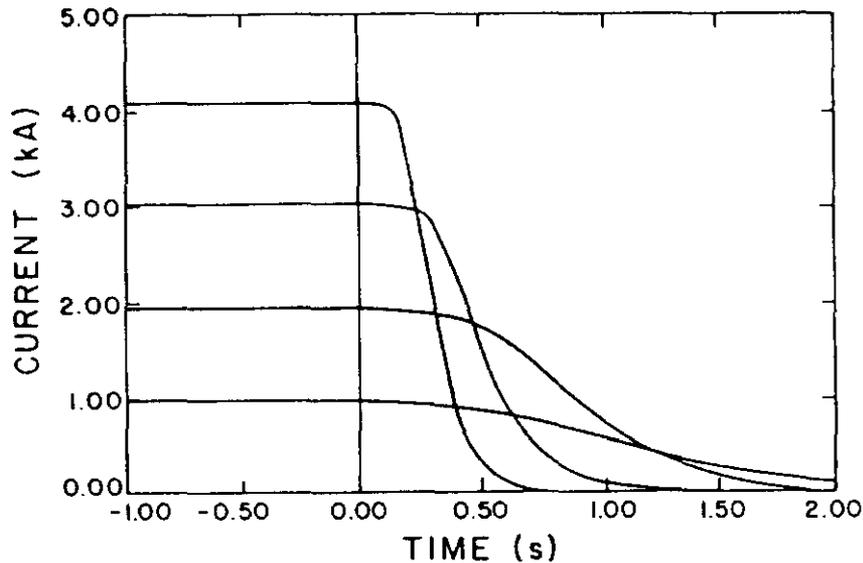


Figure 10. Magnet current for quenches at different currents.

The critical elements of the quench protection system are powered from an Uninterruptible Power Supply, which uses storage batteries as a buffer between the AC line power and its output. The UPS continues to supply power to the QPM, HFU logic and QBS controllers during a power outage.

The heater firing units use capacitors as a storage device; each unit deposits 650 J into the heater strip in the magnet. The heater strip is a thin stainless steel strip located between the superconductor and the collars, electrically isolated from both, of course. In addition, the capacitor charging circuit is transformer coupled, so the heater strip is floating. This avoids having the heater strips being a potential ground fault for the coils. The insulation between the heater and the coil results in a time delay before the quench is actually initiated. There is an additional time delay due to the discharge RC time constant of the heater circuit. However, this is not a serious problem, since the energy stored was determined by the requirement that the heaters work at low currents, where the time scales are relatively long. At high current, only a small fraction of the total energy is required to induce the quench. The HFUs are checked as part of a start-up procedure whenever the Tevatron has been off for more than four hours. The test includes a measurement of the discharge time constant, which insures that the energy is being deposited into the magnets.

The quench bypass circuits consist of two independently controlled thyristors. In addition, there is a self-firing circuit that will turn the QBS on once the voltage across the cell reaches 200 V. The "safety leads" which connect the superconducting bus to the room-temperature QBS, as well as the QBS itself, are not

intended to carry current for extended periods of time. Like the superconductor itself, the safety lead and QBS have a MIIT limit, but a much higher one. It is the protection of these elements that dictated the L/R dump time constant of twelve seconds. The QBSs, like all semiconductors, are radiation-sensitive. Consequently, they have been shielded by placing them in holes bored in the tunnel wall. Each of the QBS circuits is checked separately during the start-up check out.

The "straight section shorts" -- the devices for clamping the bus to bus voltage in each of the straight sections, as mentioned earlier -- are very similar to a normal QBS. The only differences are (i) the self-firing circuit triggers at 1 kV and (ii) the device is activated only when both QBS circuits are triggered. The latter requirement prevents the straight section shorts from conducting during the tests of the normal QBS system. The 1 kV self-firing level was dictated by the fact that there is only one current-regulating power supply. When the current is first being brought up, the full voltage of that supply can appear across the closest straight section. Also, it is desirable to be able to run the Tevatron with one or more power supplies out of the circuit, which raises the bus to bus voltage in the straight sections.

### Quench Detection System

The major component of the Fermilab quench detection system is the Quench Protection Monitors (QPMs). The QPM is responsible for monitoring the voltages across the half-cells, determining if a quench is occurring, and if so, firing the appropriate HFUs. It must also communicate to TECAR, the Tevatron Excitation Controller and Regulator, that a quench has occurred. TECAR in turn, through the other QPMs, turns off the power supplies, activates the dump switches, and triggers the QBS controllers. The QPM also communicates to the refrigerator microprocessor which cells have quenched so that cool-down can begin promptly.

The half-cell voltage monitoring is accomplished using voltage-to-frequency converters (VFCs) located in the service buildings near the QPMs. Some care is required in matching resistances and cable capacitances (Fig. 11) to minimize sensitivity to common mode V and dV/dt. The cables to the magnets in the tunnel are approximately 200 m in length, and one cable generally serves as the positive input for one VFC and the negative input for its neighbor. This requires a small correction in software. The 50 M $\Omega$  resistors which tie one end to ground result in a quench signal if one of the cables is open (e.g. disconnected). The scalars in the QPM, which read the VFC output, can be "auto-zeroed" from the Main Control Room, removing the effect of drift in the VFCs.

The ability of the QPM to turn off the power supplies is critical to the overall protection of the system. As noted above, the primary means is via TECAR. This communication occurs over a dedicated link, whose function is transmitting data between TECAR, the twenty-four QPMs and the twelve power supplies around the ring.

The power supplies receive the appropriate current or voltage program, and QPMs receive I and  $dI/dt$  information. The link communication is backed up by a second control loop which broadcasts two 100 kHz pulse trains, one positive and one negative. TECAR can clamp either one or both polarities, inhibiting further propagation of the pulses, which in turn causes, in the first case, the power supplies to turn off, and in the second, the dump switches to open. If the QPM-TECAR communication fails, the QPM will clamp both pulses.

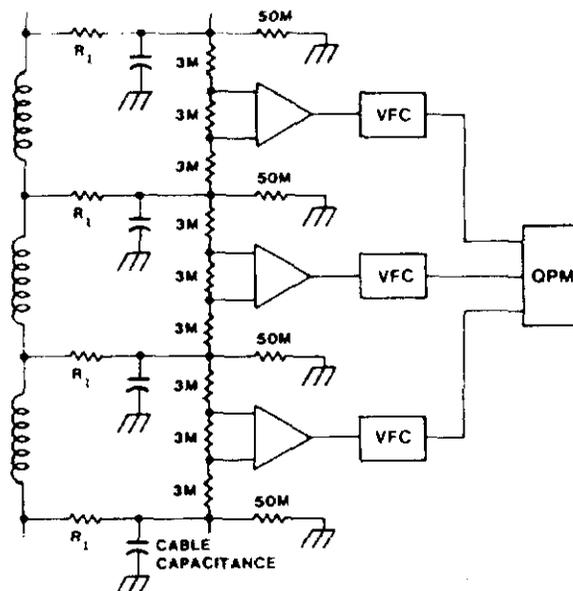


Figure 11. Schematic of voltage monitoring circuit.

The determination that a quench has occurred is based upon what we call the "resistive voltage", the difference between the measured voltage across a half-cell and the expected inductive voltage due to  $L dI/dt$ . When the resistive voltage exceeds the tolerance, then that half-cell is considered to be quenching. What should the tolerance be? Tests on a string of magnets at the B-12 facility at Fermilab determined that once the HFUs have been fired, a certain time, or MIITs, is required before the current decays away. Those data are shown in Figure 12. From the earlier hairpin measurements, we know that 7 MIITs is the limit at 4 kA. The difference determines how many MIITs can occur prior to detection. The hairpin data for voltage vs. time can be extrapolated to longer pieces of cable by a summation procedure, giving the voltage as a function of time for an arbitrarily long piece. The combination of that calculation with the MIITs allowable specifies that the quench must be detected at a level of 0.5 V at 4 kA. This detection level is required in order to protect against quenches which start in the single conductor in low field regions outside the coil. A higher threshold could be used if one only had to protect the conductor inside the coils, where the

propagation is faster, and where turn-to-turn propagation amplifies the voltage growth. At lower currents, the tolerance increases somewhat, but not very dramatically. The required level at 1 kA is approximately 3 V. The 0.5 V level is used as the quench detection tolerance in the Fermilab QPM system if the current is greater than 600 A. Below that level, we use a higher tolerance (10 V) which permits special procedures (such as testing the QBS systems) to be performed. The time scales for reaching 3 V at 600 A are so long that a quench must propagate into a magnet coil, regardless of where the quench originates. Once it reaches a coil, many turns contribute to the resistive voltage, and the 10 V level is safe.

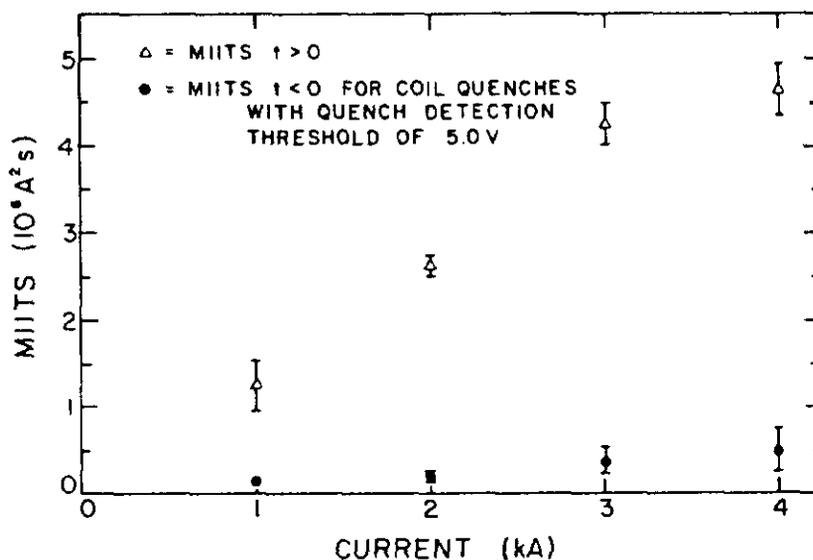


Figure 12. MIITs after firing heaters (open triangles).

In the early development of the quench protection system, the  $dI/dt$  signal used in calculating the resistive voltage was generated by analog differentiation of the current signal from a transducer. That method has a number of drawbacks. First, the differentiation process is noisy; as it became clear that 0.5 V was the required tolerance, that noise became a concern. Second, the transducer  $dI/dt$  signal was only derived in six locations around the ring, once in each sector. To transmit that signal around to the other QPMs would imply that the quench detection algorithm would require a failsafe transmission system. Failure of the transmission system, or of the transducer  $dI/dt$  hardware, would result in a large section of the ring being quenched needlessly. The quench detection algorithm was therefore changed to utilize a "relative"  $dI/dt$ . The QPM monitors the voltages across each half-cell within its domain (either 8 or 10 half-cells, depending on the QPM location.) Since the upper bus and lower bus half-cells are far apart electrically and may see transients arriving at different times, the QPM treats them as logically distinct. For each (upper and lower) bus, the relative

$dI/dt$  is determined by dividing the sum of the measured voltages by the sum of the inductances. This change was made after considerable deliberation. There was concern over the possibility of quenches growing simultaneously in the different half-cells, and therefore not being detected. However, examination of quenches showed that once the quenches have propagated into the coils, the differences soon exceed 0.5 V. Further, the different half-cells are in different cryogenic circuits, and each circuit has a temperature gradient along it. Earlier papers<sup>9,10</sup> give a more detailed account of the voltage monitoring circuitry, voltage and  $dV/dt$  common mode problems, transient response and other subtleties which enter into the quench protection algorithm. Figure 13 shows a block diagram of the QPM system, showing the interconnections within the system and to the other systems in the accelerator.

In addition to its active role in quench protection, the QPM has a passive role in the form of data gathering. Each QPM maintains a circular buffer containing analog data and status information; this buffer is stopped whenever something "abnormal" happens, e.g. quench or other event causing the system to deexcite. This circular buffer, which stores data from one second before the event to five seconds after, at a 60 Hz rate, has been invaluable, especially in resolving cause and effect questions.

## 5. OPERATIONAL EXPERIENCE

At the time of this writing, there has now been considerable experience in operating the Tevatron, primarily in the "Fixed Target" mode, but also in the "Collider" mode, in which the Tevatron stores beam for many hours. The accelerator has operated with increasing reliability, but not without some problems along the way.

### Failures During Operation

There has been only one instance in which the quench protection system failed and as a result the magnet system was not completely protected. The event was in fact a double fault, one in data transmission and one in software. As the Tevatron was first commissioned, a number of problems were encountered with the power supply turn-on sequence. As a result, one section of QPM software essentially said: it's OK if there is some status wrong, provided the current is less than 20 A. This situation, combined with a faulty current transmission which indicated a large but negative current, resulted in the failure described below. The situation has been remedied by (i) taking the absolute value of the current and requiring it to be below 20 A, and (ii) requiring the currents measured separately in the six sectors to agree within 50 A, and (iii) only masking off certain "safe" status bits. Additional checks on  $dI/dt$  are also made which further guarantee the system integrity.

The failure occurred when an injection kicker presumably misfired, resulting in beam loss in three separate cells, one each in



E1, E2, and E3. The quenches in those three cells were detected and the heaters were fired, but the power supplies were not turned off, nor were the QBS gates fired. The self-firing circuits in the QBSs fired, allowing current to bypass the quenched magnets. The power supplies kept ramping for several minutes, overheating the safety leads until the insulation failed and a ground fault tripped the power supplies. Five spool pieces were damaged and had to be removed from the tunnel for repair. It is ironic that the only quench protection failure to date was a failure mode to which passive quench protection methods are also susceptible, and it underscores the requirement for a failsafe quench detection system (and control of the power supplies and dump switches) even if the magnets themselves are passively protected.

Near the end of the first fixed target run (mid-1984), problems began developing in the magnets themselves. As noted earlier, the dipoles come in two types, the TB and TC. There are slight differences in their construction, and the TC magnets as a result contain a section of cable roughly one foot long from the point it leaves the collared coil assembly to the interface to the next magnet. The Lorentz force from the fringe field at the end of the magnet produces flexing of the cable as the current is ramped up and down. Individual strands began breaking, and the ends of the broken strands were very likely to create ground faults or bus-to-bus shorts. A single ground fault is a rather benign event; the power supplies are turned off, but the dumps remain out of the circuit. The current decays away very slowly (several hundred seconds). If there is a quench, the dump resistors are switched into the circuit, and the current comes down quickly. The voltage-to-ground distribution is shifted, and larger than normal voltages generally appear, but the system can tolerate these voltages and again, nothing terrible happens, although the faulty component will have to be located and replaced.

The voltages normally seen by the magnets are less than 500 V. The voltage on any particular component depends on its location in the ring. Each magnet has a capacitance to ground of about 60 nanofarads, and when a ground fault develops, the charge stored in the capacitance of the adjacent magnets is discharged into the fault. Consequently, ground faults which develop suddenly at high field generally result in a quench signal due to the equal but opposite currents (and  $dI/dt$ ) into the ground fault.

The fact that ground faults cause quenches (a quench-like signal, that is) turns out to be useful. Ground faults are often difficult to locate, particularly when the joints between magnets cannot be opened up very easily. For sparking-type discharges, ground faults can be located to within a few components by looking for differences in arrival time at different locations, even with the magnets superconducting. High-impedance faults are most easily identified by warming the magnets up to 20K or so, so that the voltage drop of the leakage current across the now-resistive magnet can be sensed. In both cases, tunnel access is required to make the measurements. The situation is obviously much more difficult when

the ground fault only occurs at high current. Safety considerations preclude tunnel access when the system is powered to high currents.

Quenches, on the other hand, are relatively easy to locate. The initial quench location is determined by the QPM to be on the upper or lower bus of one cell. By installing a special QPM, known as PDQ, which performs monitoring functions only, it is possible to subdivide the cell into much smaller increments. Every dipole and every spool piece has voltage taps which may be monitored. In a typical cell, the only device which doesn't have a voltage tap is the quadrupole, and there are a few special devices which don't have taps. Thus PDQ can be installed to look across every one or two devices. By stopping its circular buffer at the same time the quench occurs, the quenching component can be identified. (For quenches which are very close to the interface between magnets, there was originally some confusion in identification, since PDQ doesn't really look across one device, but from the point the voltage tap is soldered to the superconductor in one magnet to the similar point in the next magnet. As the geometry was better understood, it became possible to identify the location with better accuracy. Quenches which occur in leads, as opposed to coils, have a slower growth, and can therefore be distinguished. If the quench propagates across the interface, one can also determine which end of the magnet is quenching.) The utility of PDQ was first established during the initial Tevatron commissioning. A magnet with a turn to turn short and three poorly soldered splices had to be located.

Bus-to-bus shorts, or multiple ground faults, are another matter. The current can take an alternate path, and the quench protection is basically defeated. As a result, such failures are generally more catastrophic; melted superconductor and cryostats with holes burned in them are the most common result. In this case, locating the fault is very easy, but the insulating vacuum and superinsulation are filled with helium. Leak-checking after replacing the faulty component becomes difficult to impossible, due to the contamination. A total of five TC dipoles failed during operation between March 26 and July 9, 1984. Two were single ground faults, three involved bus-to-bus shorts. During the 1984 summer shutdown, all TC magnets were repaired by opening the cryostats and tying the two leads together to prevent motion. Broken strands were found in many dipoles; if more than six were broken, a new section of cable was spliced in place. This was done to five dipoles in the course of the repair.

The problems with the TC dipoles all occurred at the upstream end; at the downstream end, where the voltage taps and heater firing connections are made, the leads were tied together already. However, in 1985, one TC dipole in which the leads had not been tied on the downstream end, also failed in a similar manner.

The strand-breakage itself is an interesting situation. Since the cable is in low field, particularly right at the interface where the breakage was occurring, its critical current is much higher. As a result, the increase in current density in the remaining strands

does not exceed the critical current if only a few strands are broken, so the system can operate under those conditions. But should a quench occur anywhere in the cell, the heaters will fire and the quench will propagate to the region of broken strands. The temperature increase in the cable is related to the MIITs divided by the area squared. In the bus-to-bus short failures, the short typically appeared well after the quench, and may well have been due to overheating from this effect. The subsequent melting of superconductor may have been inevitable and may have occurred without the shorts.

A TB dipole failed during an 800 GeV quench three days before the scheduled end of the 1985 fixed target run. The failure occurred in the leads at the upstream end of the magnet, about 1.5 inches from where they exit the coil. A bus-to-bus short and cryostat rupture resulted. There was evidence of a deep score mark in the conductor at the point of the failure, which was inside a G-10 block securing the leads. The damage had to occur during construction, since the area was covered by G-10. What is most puzzling is the timing: a few days earlier, in an effort to locate the weakest magnets in each sector, that area had been ramped to 900 GeV without quenching. The cell containing the magnet that later failed quenched at 910 GeV without damage, and operation continued at 800 GeV for several days.

Finally, there was a failure in January, 1985, in which a power supply transformer shorted from primary to secondary, raising the bus potential to 13.8 kV, far in excess of its rating. The damage was limited to F-sector, where several devices had to be replaced.

The point of this rather lengthy expose on magnet failures can best be summarized as follows: the ends of the magnets, in addition to being difficult and costly in terms of cryostat construction, have also been the source of virtually all of the magnet failures to date. Designers of future superconducting magnets should pay attention to the electrical details of the ends, as well as to simplifying the cryostat.

#### Other Aspects of Operating a Superconducting Accelerator

Quenches were the largest source of unscheduled downtime during 1985. Almost ninety percent of the quenches were due to beam loss; about half of the remainder can be attributed to the Quench Protection System. The QPM-related quenches arise from spontaneous HFU discharges, faulty or drifting VFCs, or QPM failures, (in which case all HFUs in the building are fired, so that the magnets are protected.) The HFU problem has been the most aggravating, and that system is in the process of a redesign. Another QPM-related problem is the inability to distinguish real quenches from "false" quenches. The latter can occur in association with power supply regulation problems which place fast transients on the bus.

Repeated quenches in the same cell pose a special problem. They occur at certain regions of the ring: near the injection, extraction and abort channels where beam loss is most likely. Kicker timing problems, for example, can cause local beam loss. The problem with

repeated quenches is that the safety leads which bypass current heat up, and the cooldown time is much longer than the time required for the cryogenic system to recover. To shorten the cooldown, vapor cooling was installed on those spools in regions most likely to quench. Running at 800 GeV, the problem has not been too severe; there were a few occasions when, after three quenches, the accelerator was shut down to avoid potential overheating should a subsequent quench occur. The overheating problem becomes more severe as the energy of the machine is increased, and vapor cooling is being installed on all leads. The vapor cooling decreases the cooldown time by roughly a factor of two, and some future problems are anticipated. It is desirable to limit the flow of the cooling in order to avoid forming large iceballs. Even with adjustable flow (e.g. refrigeration type flow controllers, as presently used on the main power leads and the already-modified safety leads) the question remains--how does one insure that there is flow?

With the beginning of commissioning the Tevatron as a Collider, it became evident that the beam lifetime was not limited by the slow losses due to the quality of the vacuum, or to collisions at the anticipated luminosity, but rather, the lifetime was limited by sudden loss of the entire beam as one device or another caused an abort. The QPM system was among those causing aborts. The beam must be aborted any time the system is deexcited. Changes were made in the QPM response with the primary objective of avoiding the need to abort. Many of these changes have had the further benefit of easing the operation for fixed target physics as well. The most notable example is the response to the refrigeration status becoming "bad". The original response was to deexcite immediately, thereby aborting the beam. Now, the response is to inhibit the next ramp from occurring. In fixed target mode, this allows the beam to be extracted to the experimental areas and allows the refrigeration system an extra chance to recover. If it doesn't recover, the accelerator idles at 400 A until it does. This sequence is operationally much easier than turning off and back on again. In the Collider mode, the scarcity of antiprotons dictates that the machine continue operation without interruption. There is little difference if the beam is lost due to a quench or due to deexcitation. The time required for quench recovery is sufficiently short that it is better to give the refrigerator a chance to recover. A similar response is taken for (i) a single HFU channel showing faulty status or (ii) the circular buffer being stopped.

#### SUMMARY

This paper has discussed only the quench protection system for the main Tevatron system. In addition, there are four independent low- $\beta$  quadrupole circuits in B0, each with their own QPM, and three systems in Switchyard. These "auxillary" systems required developing

somewhat different quench protection algorithms and hardware, but the QPMs themselves are basically the same as the 24 in the Tevatron.

The QPM system is indeed a complicated system. It must monitor and protect the superconducting magnets without introducing potential ground faults. It must be intimately linked to the power supply system. It must have sufficient redundancy that single-device failures do not jeopardize the system's ability to protect the magnets. It must have sufficient sensitivity to detect quenches under many conditions, from low current to high, inside coils and in leads, while ramping or running dc; and yet have sufficient immunity to noise and transient conditions to distinguish them from real quenches. Many people at Fermilab have struggled to accomplish these goals. The QPM system is now a fairly stable system, as opposed to the situation just two years ago when the software was still rapidly evolving. With the improvements to the HFU system, and streamlining of the checkout procedure, it should prove more and more "user-friendly."

## REFERENCES

1. M. Wilson, Superconducting Magnets, Clarendon Press Oxford (1983).
2. R. Palmer and A. Tollestrup, *Ann. Rev. Nucl. Part. Sci.* 34, 247 (1984).
3. A. Tollestrup, *AIP Conf. Proc.* No. 87, 699 (1982)
4. A Report on the Design of the Fermilab Superconducting Accelerator, May, 1979
5. W. Sampson, M. Garber and A. Ghosh, *IEEE Trans. on Nucl. Sci.*, Vol. NS-28, No. 3, 3245 (1981).
6. W. Fowler, M. Kuchnir, R. Flora and R. Remsbottom, *IEEE Trans. on Magnetics*, Vol. MAG-17, No. 1, 925 (1981).
7. K. Koepke, P. Martin, M. Kuchnir, *IEEE Trans. on Magnetics*, Vol. MAG-19, 696 (1983).
8. M. Kuchnir and J. Tague, Fermilab TM-679 (unpublished).
9. G. Tool, R. Flora, P. Martin and D. Wolff, *IEEE Trans. on Nucl. Sci.*, Vol. NS-30, No. 4, 2889 (1983).
10. P. Martin, R. Flora, G. Tool and D. Wolff, Fermilab TM-1134 (unpublished)