Ann. Rev. Nucl. Part. Sci. 1985. 35 : 605–60

THE TEVATRON ENERGY DOUBLER: A Superconducting Accelerator

Helen T. Edwards

Fermi National Accelerator Laboratory,[1] Batavia, Illinois 60510

CONTENTS

1. INTRODUCTION

1.1 *The Facility*

The Tevatron project is a massive upgrade of the original Fermilab accelerator complex, which became operational in 1972 and operated from

[1] Operated by Universities Research Association, Inc. under contract with the US Department of Energy. The US Government has the right to retain a nonexclusive royalty-free license in and to any copyright covering this paper.

1973 until mid-1982 at typically 400 GeV and 2×10^{13} protons per pulse (ppp). The three parts of the Tevatron project when complete will allow for both fixed-target and collider hadron physics using primary beam energies in the 800–1000-GeV range. The Fermilab facility should thus be able to maintain its position at the forefront of high-energy physics research facilities (1). This paper describes the superconducting accelerator phase of the program. Because the new synchrotron provides protons with twice the energy of the "old" Main Ring, it is frequently called the Energy Doubler. It first accelerated a beam in July 1983 and is presently operating at 800 GeV with intensities above 1×10^{13} ppp.

During Main Ring operation from 1972 to 1982, maximum accelerator energies of 500 GeV and intensities of 3×10^{13} ppp were obtained (not simultaneously). The normal operating energy was 400 GeV. The accelerator chain consisted of a 200-MeV H^- Linac, an 8-GeV Booster Accelerator, and the Main Ring. The beam was resonantly extracted and split for distribution to numerous targets and secondary beam lines for fixed-target physics (2). The new project consists of the following additions or upgrades: (*a*) The Energy Doubler (3), which is the first large accelerator to use a superconducting magnet system for the main guide field. This technical achievement is the key to the Tevatron and for that matter to future higher energy proton accelerators. Excitation should be sufficient for 800–1000-GeV beam energies. (*b*) Beam lines to handle the increased energy capability of the accelerator for fixed target physics experiments. (*c*) The $\bar{p}$ Source—two 8-GeV rings the size of the Booster that accumulate antiprotons to be used for $\bar{p}$-p collisions in the Energy Doubler for colliding-beam experiments.

The Energy Doubler uses the Main Ring operating at 150-GeV peak excitation as an injector. The Doubler beam is located beneath that of the Main Ring in the same tunnel. The Main Ring is also operated at 120-GeV excitation in order to accelerate and extract protons to the $\bar{p}$ production target for the Source.

Operation with the Main Ring at 400 GeV typically was at a cycle time of 7 to 16 s. The slow 16-s limit was imposed to limit power demand during the daytime. The guide field ramp was such as to allow 1 to $1\frac{1}{4}$ s of 400-GeV flattop excitation during which the beam was slowly extracted from the ring to the various experimental areas. In addition, neutrino experiments required a few fast pulses of beam (~ 1 ms) with intensities up to $\sim 10^{13}$ ppp.

Cycle time with the superconducting magnets is of the order of once per minute, with slow extraction over 20 s. The beam is split in the switchyard area using electrostatic septa and directed to a number of primary targets. This very long time for slow extraction gives a 1/3 duty factor and provides for very good data-taking rates in experiments requiring low intensity. A number of pulses of fast resonantly extracted beam are produced during the excitation flattop along with the slow spill. These pulses of beam typically contain up to 2×10^{12} particles and last about 2 ms. The pulsed beam is diverted around the splitting septa and goes only to the neutrino production target. Though the long cycle time with limited intensity per cycle is not so favorable for neutrino experiments, the increase in the neutrino flux and interaction cross section with primary energy compensates for the slow repetition rate. Thus, in the Energy Doubler with a 60-s cycle time at 1000 GeV, one can expect the same average neutrino event rate as with the Main Ring at 400 GeV and a 10-s cycle time, assuming equivalent numbers of protons targeted per cycle for neutrino production.

The third part of the Tevatron project is the construction of the $\bar{p}$ Source rings and the adaptation of the Energy Doubler to $\bar{p}$-p storage operation (4, 5). In order for the superconducting ring to do colliding-beam physics, protons and antiprotons will be accelerated and stored in it while circulating in opposite directions. Initial operation calls for three bunches of beam in each direction. Collisions will take place at each of the six long straight sections. Major detectors will be constructed at two of the collision points.

The antiprotons are produced by extracting and targeting 120-GeV beam from the Main Ring. Resultant 8-GeV antiprotons are collected first in the debuncher source ring at an approximately 3-s cycle rate and then are transferred to the accumulator source ring where they are allowed to build up for a number of hours. During the accumulation time the $\bar{p}$ beam phase space is reduced or "cooled" so that the required emittance can be obtained for efficient injection into and acceleration in the Main Ring. Three equally spaced bunches of p and $\bar{p}$ are to be injected into the Energy Doubler and accelerated to full field, where they will be stored and collide for a number of hours.

Table 1 lists the general parameters of the Tevatron for both fixed-target and collider operation. Figure 1 illustrates the present operating cycle for the Main Ring and the Energy Doubler during fixed-target operation, with the Main Ring also being used for $\bar{p}$ production. Figure 2 is a schematic of the accelerator complex, and Figure 3 illustrates the beam lines available for fixed-target operation.

1.2 *Historical Motivation*

The original impetus for the Energy Doubler in the early 1970s was just what the name implies; to double the magnetic field and consequently the beam energy without constructing a new tunnel and support facilities. The physics justification was basically just higher energy to explore new regions. The Main Ring and its service buildings were originally laid out with the

Table 1 Tevatron parameters

General	
Accelerator radius	1 km
Peak beam energy	800–1000 GeV
Injection energy	150 GeV
Bend magnetic field at 1000 GeV	44 kG at 4400 amp
Beam emittance (95% normalized)	24π mm mr
Fixed-target	
Intensity	$\sim 2 \times 10^{13}$ protons/cycle
Acceleration rate	50 GeV s^{-1}
Cycle time	60 s
Slow spill duration (flattop time)	20 s
Fast spill	5 pulses at 2×10^{12} protons/cycle expected
Collider	
Intensity per bunch	0.6×10^{11} protons (antiprotons) expected
Number of bunches	3p, $3\bar{p}$
Luminosity	1×10^{30} cm^{-2} s^{-1}
Storage time between fills	$\sim$4 hr
Interaction point amplitude function (β)	1 meter (x, y)

idea that there would someday be a superconducting ring in the same tunnel, and the service buildings had room for the additional electronics of the superconducting ring (6).

Between 1973 and 1978 total power usage on site doubled, and the cost of power per kilowatt hour also doubled. It was clear that power costs would continue to rise and, in fact, over the ten-year period from 1974 to 1984 they have increased sixfold. The Main Ring magnet system alone at 400-GeV operation was by far the largest single user: over 50% of the total usage. The use of superconducting magnets could substantially reduce usage; the total power requirements for all Main Ring and Energy Doubler systems together at either 400- or 800-GeV operation uses half of what the Main Ring required when running with a 70% utilization factor at 400 GeV. As a bonus, the flattop duty factor relative to the whole magnetic cycle has increased from about 1/10 to 1/3, though obviously the number of protons accelerated per second has decreased. Details of cryogenic usage, manpower requirements, and equipment complexity all enter into the cost-benefit analysis, but the fact remains that the superconducting ring does operate at twice the energy with much better flattop duty factor for substantially less power than the old Main Ring.

It became apparent, as research and development of superconducting

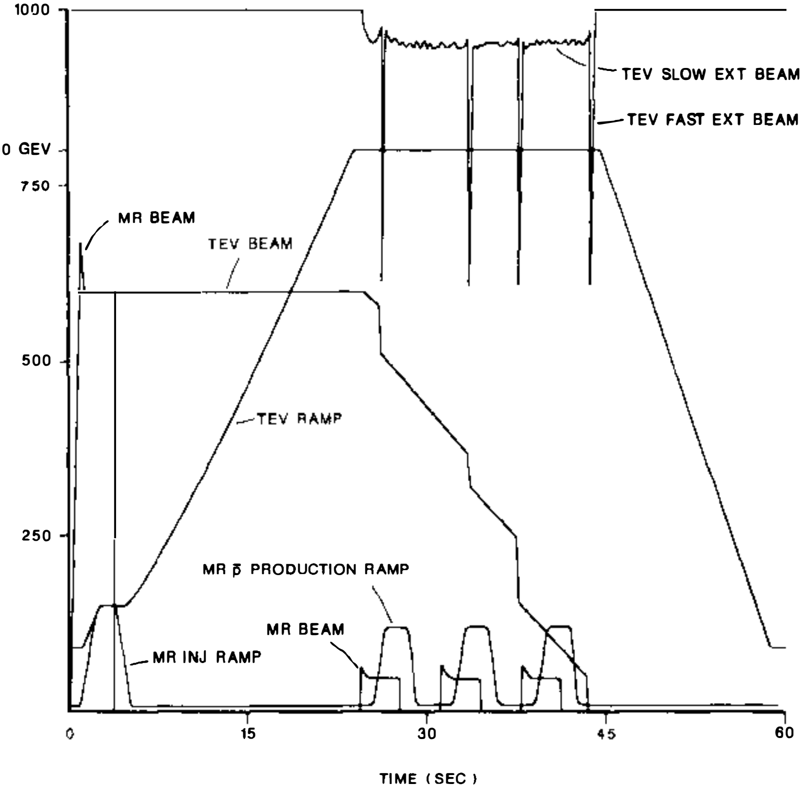

Figure 1 Main Ring and Energy Doubler cycle for 800-GeV fixed-target operation with three additional Main Ring p̄ production ramps included in the super cycle. (Vertical scale for beam intensity not shown.)

magnets proceeded, that collider physics was indeed the exciting new field with its great potential to reach high center-of-mass energies. The Main Ring would be extremely expensive to operate above 200 GeV for storage, and it probably would not make a good storage ring in any case. The tremendous advantage of superconducting ring magnets for beam storage at high excitation for hours with no additional operational cost became apparent. In fact, though there was talk of storage in the Main Ring or construction of other conventional rings for various colliding-beam scenarios, the only thing that really made sense was the development of the p̄ Source and the use of the Energy Doubler as a storage ring for p-p̄ collision experiments at 2-TeV center-of-mass energy.

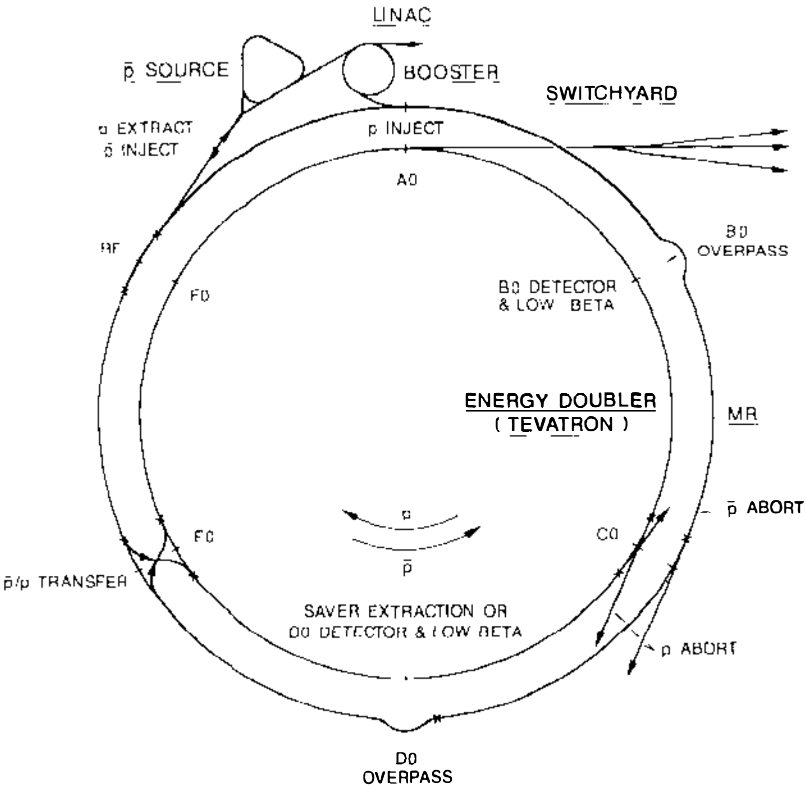

Figure 2 Tevatron Accelerator including the new p̄ Source rings and the Energy Doubler.

1.3 *Design Issues*

The design of a superconducting-magnet accelerator presented many real engineering challenges and raised several accelerator physics questions.

The magnets had to have a small cross section in order to fit under the Main Ring, and the superconducting cable could not be cryostatically stabilized as in the large bubble chamber magnets of the 1970s. This meant that the magnets had to be protected from their own stored energy in the event they stopped superconducting or "quenched." The magnet coils had to be thermally insulated from room temperature and cooled to 4.6 K with sufficient refrigeration to maintain their low temperature in the face of cryostat heat leaks, ramping eddy currents, and field hysteresis energy losses. The big question was how to wind and clamp the coils so they would not move under ramping field forces, and so that they had highly linear and reproducible magnetic fields throughout their excitation cycle.

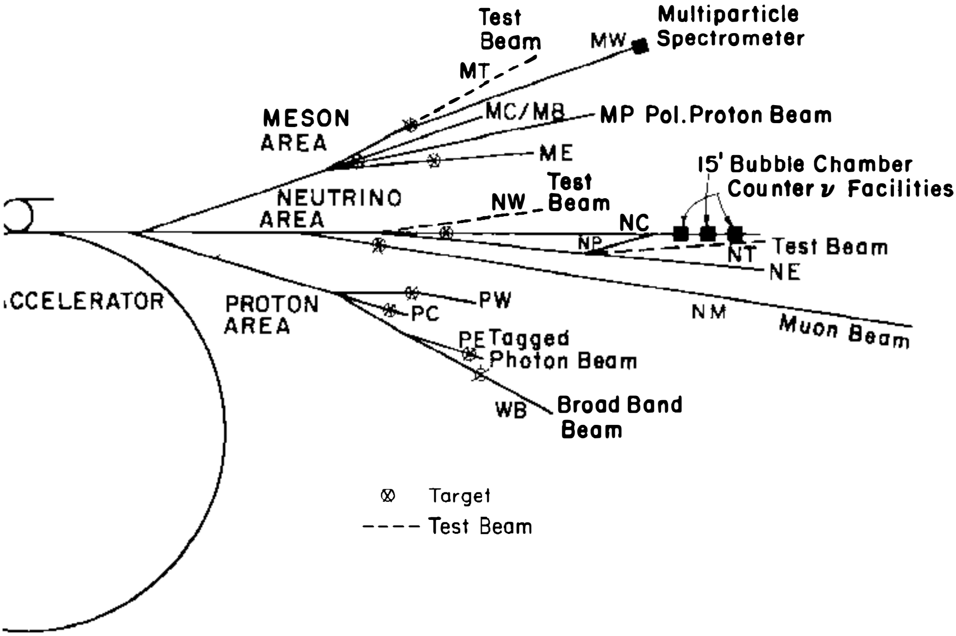

Figure 3 Expected layout of fixed-target physics beam lines upon completion of construction (1987).

The major accelerator physics questions were threefold. First, how good did the magnetic fields have to be over what region of the magnet aperture? Second, could beam losses be kept low enough so as not to quench the magnets during normal operation? Third, what sort of accelerator adjustment and control, instrumentation, and diagnostics would be required to aid in smooth commissioning and operation of an accelerator in which it was proposed to challenge the laws of nature that require an environment a few degrees above the absolute zero of temperature.

The first question tends to be one of the hardest to quantify in accelerator design, especially for proton accelerators where the memory of the particles is very long in the absence of synchrotron radiation. Designers really do not know what field quality is adequate to make a storage ring with long lifetime, for instance. An equally difficult problem is assessing the importance of various nonoptimized operational conditions, how often and to what extent they are likely to occur, and to what extent the design should allow for their occurrence. It is obvious, for instance, that the "good" magnetic field region should be as large as the beam size, but what about beam steering? In principle one can steer the beam down the center of the beam pipe within the resolution of the detectors and power supplies, but how do you get it there to begin with? How much will it drift with time? How much do you want to move it to explore its optical properties, etc? Taken in total, what is the difference between a good accelerator and a bad one? Why do some just run and others require continuous attention?

For the Energy Doubler, addressing these imponderables was cut short by evaluating the demands of the resonant extraction process. This process is precisely defined in terms of the amount of magnetic aperture it requires and the percentage of beam particles that under optimum conditions will hit the extraction septum. Thus, for the Energy Doubler beam-tracking studies were performed in detail for the few hundred turns required for extraction particle trajectories. Inside this required aperture, field quality was made as good as was reproducibly achievable by the manufacturing techniques and consistent with small resonant bands and realistically sized correction magnets. The losses from fast extraction were modeled by tracking particles after their interaction with the electrostatic septa. Particles were tracked either until they were extracted as halo or until they struck magnets or collimators and produced cascade showers. Energy deposition of the cascades in the superconducting magnets was calculated and could be compared with expected superconductor tolerance levels, and then allowed extraction beam intensities could be estimated.

The third question of control and diagnostics was confronted by realizing that the superconducting magnet system would not be forgiving, that magnets would quench from beam or refrigeration problems, that quench recovery would not be rapid, that cycle times would be long, and that a massive amount of hardware would be required for the refrigeration and power supply quench protection systems. Thus large amounts of distributed control and data processing capability were envisaged, along with diagnostics capable of recording the sequence of events during an acceleration cycle and leading up to a quench. The quench is prohibited, if possible, by "aborting" or removing the beam from the ring within the time duration of one turn, if for instance beam loss levels approach those that might produce a quench.

Finally, the problems of installation, commissioning, and operation had to be considered. The question of whether or not the vacuum system would ever become leak tight caused serious concern during initial installation. Reliability of the plant was not expected to be high during commissioning; accordingly, beam instrumentation and diagnostics were designed to provide information as efficiently as possible. Operating experience and operational reliability can now begin to be analyzed and the design and implementation judged in retrospect. It is possible to speculate upon changes and differences in approach for future accelerators.

2. SPECIAL DESIGN FEATURES

2.1 *Overview*

The main special design features of a superconducting accelerator are, of course, the magnets themselves and the cryogenic system to cool them. The magnets are packaged in their own cryostats, which provide a reservoir for the helium cooling and thermal isolation between liquid helium temperature and room temperature. Vacuum systems are required not only for the tube in which the beam circulates but also for thermal insulation in the cryostat. In the Energy Doubler, the magnet cryostats themselves also provide a transfer line and heat exchanger system for the helium cooling as well as a nitrogen heat shield to reduce the refrigeration requirements at helium temperatures (7).

Second only to these requirements are the questions of protecting the magnets when they quench and of determining what level of beam intensity can be tolerated without quenching. The power supply and quench protection system for the Energy Doubler main bend and quad magnet string is a large distributed "active" system in which resistive voltage development is monitored continuously across >200 separate segments of the ring in order to divert current and extract electrical energy from the magnets in the event of a quench. Regarding beam loss, arduous calculations are necessary to convince oneself that a superconducting accelerator is able to operate without quenching.

The system required to control and monitor the refrigeration, vacuum, power supply and quench protection systems, and special beam diagnostics must acquire and process far more data than its predecessors, even before considering the requirements of the more conventional accelerator systems. In order to obtain sufficient support, distributed process control and monitoring have been provided in the 30 remote service buildings around the ring.

2.2 *Magnet Cross Section and Accelerator Layout*

The cross section of the magnet-cryostat illustrated in Figure 4 shows the concentric design. Starting with the inside is the vacuum space for the beam. This is followed by the single-phase (liquid helium) cryostat, which contains the magnet coil clamped by stainless steel collars in a highly reproducible, accurate configuration that does not distort during magnetic excitation. Though most of the single-phase volume is filled by the collared coil assembly, small annular regions between the beam tube and coil, and between the collars and the outer single-phase pipe, allow helium to flow along the length of the magnet. The next concentric pipe encloses a volume for the two-phase (liquid and gas) helium, which returns along the length of the magnet in the opposite direction and thereby allows for counterflow heat exchange at the surface of the single-phase tube. Outside of the helium container is an insulating vacuum space and then two concentric pipes of only slightly different diameter; these contain nitrogen liquid and intercept heat flow from room temperature to liquid helium temperature (4.6 K). The insulating vacuum region between the nitrogen shield and the room temperature outer cryostat tube contains superinsulation (aluminized Mylar) as an additional radiation shield. This whole magnet-cryostat assembly is vacuum tight. It is held in a laminated iron yoke that contributes $\sim 18\%$ to the total magnetic field. The assembly is precisely (~ 1 mil) adjusted relative to center with G10 suspension blocks and preloaded suspension cartridges that allow for contraction and expansion during the thermal cycle. The dipole (21 ft in total length) requires nine sets of suspension points.

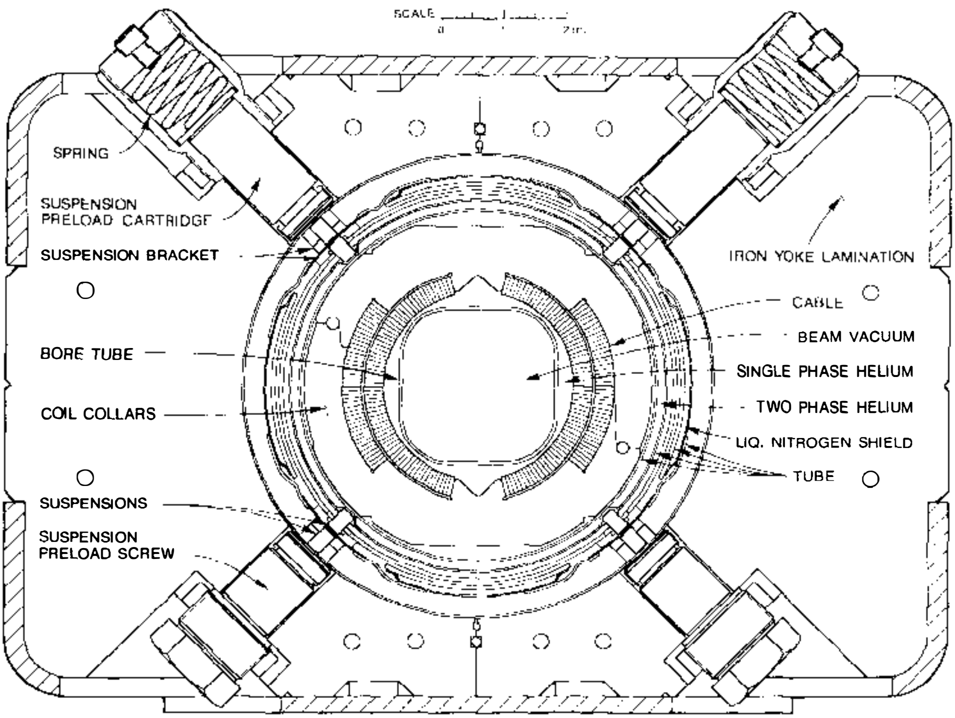

Figure 4 Cross section of the Energy Doubler dipole magnet showing the collared coil assembly, the cryostat, and the warm iron yoke.

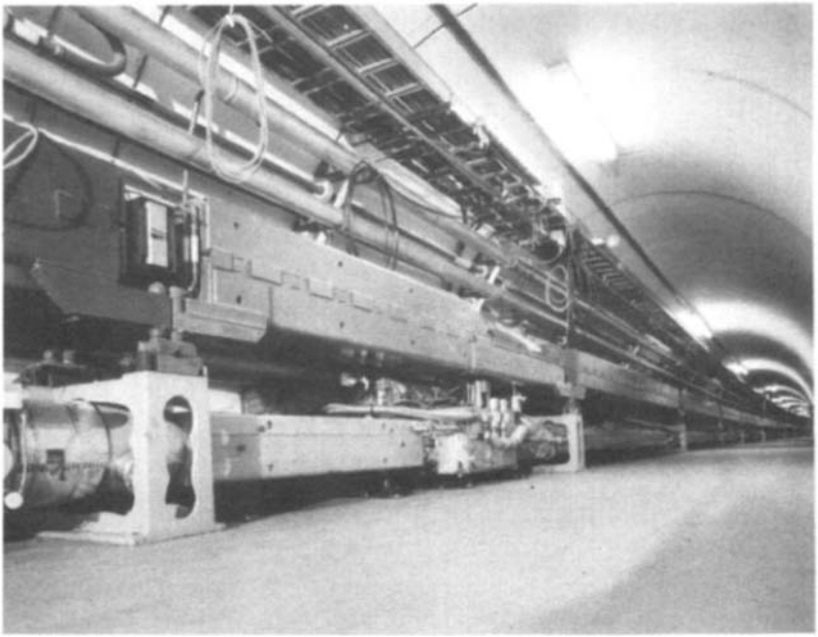

Figure 5 The Main Ring tunnel with the superconducting Energy Doubler magnets installed underneath the conventional Main Ring magnets.

The magnet layout for the Energy Doubler follows closely that of the Main Ring, constrained as it is by the requirement that the Doubler fit under the Main Ring magnets with their center lines separated by $25\frac{1}{2}$ inches (Figure 5). Both rings have six sectors and six long straight sections. In addition, each sector has two additional drift regions of 6 and 12 meters where kickers, special magnets, and detectors for injection, extraction, etc can be placed.

Each sector has 31 half-cells as well as cells with quadrupole doublets at the ends of the straight sections. A standard half-cell includes one quadrupole, one correction coil unit ("spool piece"), and four bend magnets. The sector is divided into four magnet string units or cryo sections, which typically contain eight half-cells. These 1/4 sector units are sometimes called "houses" because cryogens, power, vacuum, control, and monitoring for each unit are provided by a service building located above the tunnel. There are 24 houses for the support of the six ring sectors and an additional six for the support of the six straight sections.

2.3 *Magnet*

Superconducting accelerator magnet design was discussed in a prior review article (8). Here we only review some of the more basic features of the Energy Doubler magnets.

The desired cos θ current distribution for the coil is approximated by two layers of NbTi Rutherford-style cable with the inner and outer layers extending to $\pm 72°$ and $\pm 36°$ respectively. The 23-strand cable is keystone shaped in cross section to form a self-supporting semicircular Roman arch. Cable dimensions are 0.044 to 0.055 keystone thickness $\times$ 0.308 inch width. Twenty-three strands of 0.0268-inch diameter wire with a copper-to-superconductor-area ratio of 1.8 : 1.0 are twisted flat to make up the cable. Each strand has 2050 filaments of NbTi alloy (53.5 to 46.5% by weight) that average 8.7 μm in diameter. The small filaments are necessary for prevention of flux jumping in the superconductor (9), and also play a role in minimizing persistent current effects. To avoid inductive loops within the strands themselves, which would lead to interfilament coupling and large ac losses, the strands are twisted twice per inch. The cable is insulated with 1-mil thick double-lapped Kapton. Outside of the Kapton wrap is a layer of spirally wrapped fiberglass tape impregnated with epoxy that flows during the coil-molding process to provide a reasonably solid coil package. Since a typical cable in the magnet is subject to very large Lorentz forces ($\sim$100 lb per linear inch), laminated stainless steel collars are press-fitted around the insulated coil to prevent any motion that could cause quenching. The magnets must operate to ± 500 volts to ground under normal operation and 2 kV during failure mode, when the coils may be surrounded by up to 180-K helium gas at 40 pounds per square inch absolute (psia), i.e. relative to vacuum. The magnets are hi-potted to 2 kV at 4-atm room temperature helium prior to installation. (Note that while helium liquid is a good insulator, helium gas has a low breakdown threshold.)

Tests on superconducting cable for the Doubler magnets give an average J_c of 1800 amp mm^{-2} at 5 T and 4.2 K. Taking into account the magnet geometry, field, and operating point of 4.6 K, one finds the average magnet current that this J_c would allow is 4.6 kA (10). All magnets prior to their installation in the ring were measured at the Magnet Test Facility under two different excitation sequences: "quench" and "cycle." In the first test, magnets were ramped at 200 A s^{-1} until a quench occurred. In the second test repetitive ramps approximating the accelerator cycle were used and the flattop current increased until the quench occurred. The results from these tests for the magnets installed in the ring (see Figure 6) indicate that ring excitations of 900–950 GeV should be possible with only a few magnet replacements, whereas to reach 1000 GeV will probably require either very slow ramps or cryogenic modifications, such as "cold compressors" to lower the operational pressure (temperature).

Magnetic field quality is given by the multipole coefficients in the expansion

$$B_y + iB_x = B_0 \sum_{n=0}^{\infty} (b_n + ia_n)(x+iy)^n,$$

where the pole number is $2(n+1)$ and b_n (a_n) is the normal (skew) multipole coefficient, and b_0 is unity. The multipoles allowed by dipole symmetry, $b_2, b_4, b_6, \ldots$, are designed to be small and would be zero for a pure cos θ winding. The Energy Doubler magnets were designed assuming the conductor could be placed accurately enough ($\sim$1 mil) (8) to give high field multipoles a few times 10^{-4} at a reference radius of one inch. Generally speaking the measured values meet those expectations (11–13). An important feature of the magnet is the ability to null the normal and skew quadrupole moments by off-centering the collared coil in the external iron with the suspension cartridges. In this way the b_1 and a_1 rms widths for collared coils of 1.9×10^{-4} and 2.9×10^{-4} were both reduced to $\sim 0.7 \times 10^{-4}$ for the finished magnets.

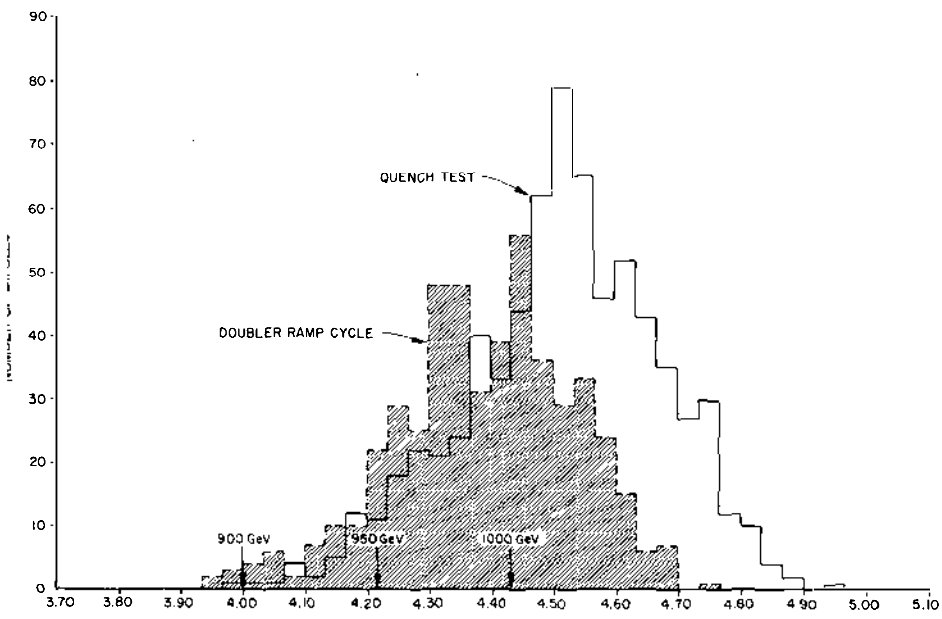

Figure 6 Quench current distribution for dipoles installed in the Energy Doubler.

Multipole coefficients that result from the position of conductors are expected to be independent of current as long as the coil assembly is rigid. "Persistent currents" in the superconducting filaments, resulting from eddy current shielding effects (Meisner effect), produce a hysteresis behavior in the b_0, b_2, b_4 coefficients that is dependent on whether the field is increasing or decreasing as well as on the size of the filament, the $J_c(B)$ of the filament, and the coil radius (14). This results in magnetization that persists even when the power supply current is reduced to zero (Figure 7). For the Energy Doubler, the persistent dipole and sextupole fields at one inch are about ~ 6 gauss; this means that the sextupole moment (b_2) is -5×10^{-4} at injection whereas the average high field value is 1×10^{-4}. The standard deviation from the mean is 3.6×10^{-4}, independent of excitation (15).

The filament-induced magnetization exhibits hysteresis that in turn leads to cyclic energy losses as the magnets are ramped from minimum to maximum excitation and back. Additional eddy currents induced between

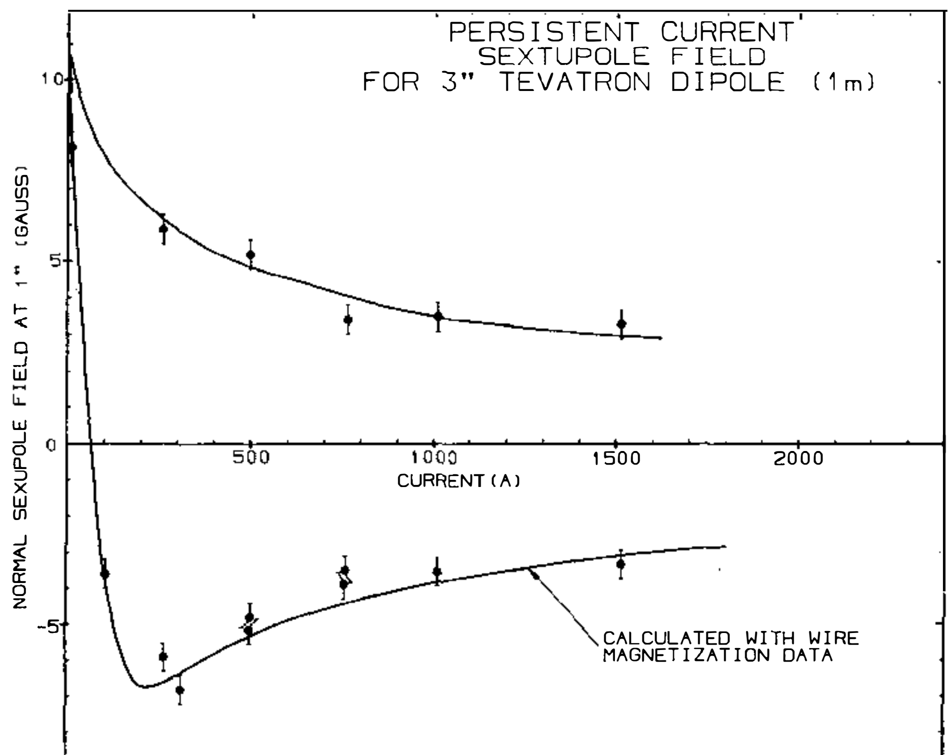

Figure 7 Persistent current sextupole field of an Energy Doubler dipole.

strands and in the Cu matrix surrounding the filaments also cause ac losses that are proportional to both dB/dt and B_{max}-B_{min}. For a typical 60-s cycle, the total ac loss is ~500 joules per dipole, which results in an average power consumption of 8.4 W (13). Superconductor hysteresis is responsible for about 50% of this value.

Average values of the low order multipole coefficients (b_1, b_2, b_3, a_1, etc) produce tune shifts, chromaticity (momentum-dependent tune), amplitude dependence of tune, and horizontal-to-vertical coupling in the accelerator that can be compensated for with distributed correction elements. The rms fluctuations of the multipoles in the magnets will produce resonant bands and make the accelerator unstable at certain tunes. These also can be compensated by corrections, but their effect can be minimized if the magnets are installed in an order that cancels the magnetic imperfections. This "shuffling procedure" (16) was performed on groups of ~32 bend magnets (1/24th of the ring) and was designed to minimize sextupole and skew octupole resonance driving terms. The driving term for the $3\nu_x$ resonance was reduced by over a factor of 30; more typically other terms were reduced by factors of 2 to 3 (17). Third-integer resonance (full) widths of $\Delta\nu = 0.003$ to 0.004 would have been expected at injection (for $0.15\pi \times 10^{-6}$ m emittance) from the measured b_2 rms value (18). Half-integer full widths of $\Delta\nu = 0.006$ were expected from the b_1 rms of 0.7 and $\Delta\nu = 0.007$ was measured at 500 GeV (19).

2.4 *Cryogenic System*

The Energy Doubler helium refrigeration system is the world's largest. It is a hybrid system and contains a large helium liquefier, a nitrogen reliquefier, a distribution system around the ring for liquid helium and nitrogen (transfer line), as well as header pipes for room temperature high- and low-pressure gases, and a distributed satellite refrigerator system that supplies the 24 individual magnet strings with helium and nitrogen cooling (20). These components (Figure 8) can provide a total of 24 kW of cooling at 4.7 K for the magnets as well as 600 liters per hour (ℓ h^{-1}) of liquid helium for power lead cooling and 1000 ℓ h^{-1} of nitrogen for the magnet heat shields.

The large helium liquefier can supply over 4500 ℓ h^{-1} of liquid and uses 3.6 MW of power. It consists of two 2000-hp, three-stage compressors (and a spare) that operate at 12 atm and supply 1200 g s^{-1} of helium gas to a heat exchanger "cold box" that operates with three turbine expanders. The first stage of cooling is provided by a liquid nitrogen counter flow heat exchanger that uses 0.64 liters of nitrogen per liter of helium produced.

The nitrogen reliquefier, which will become operational in 1985, can produce 85 tons of liquid per day and will use 3.0 MW. For operation without the reliquefier, liquid nitrogen must be delivered by 5000-gal commercial tankers (5000 gal = 17 tons). During full operation one tanker must arrive on the average of every four hours.

The distribution system around the ring is 6.5 km long and consists of a "transfer line" for liquid nitrogen and liquid helium and header pipes for high- and low-pressure helium and low-pressure nitrogen. The transfer line is made of four nested stainless steel pipes containing (from the inside out) supercritical helium at 5 K and 3 atm, an insulating vacuum space, 3-atm subcooled nitrogen that acts both as a radiation shield for the helium and as a nitrogen supply for the satellite refrigerators and magnets, and, finally, an

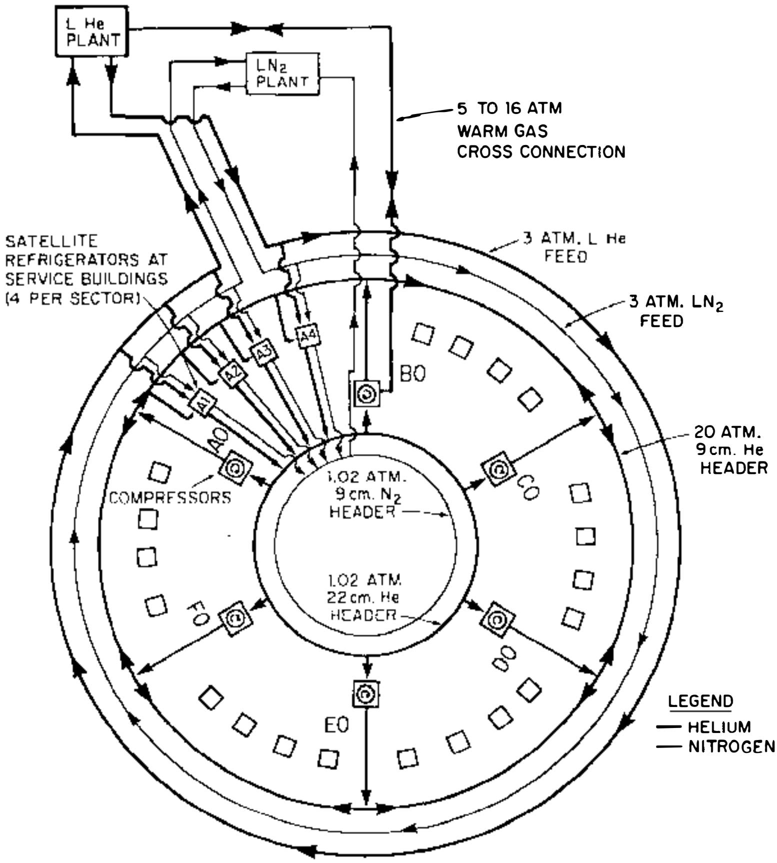

Figure 8 Layout of the helium and nitrogen refrigeration system, comprised of the central helium and nitrogen plant, the transfer line and header distribution system, and the satellite refrigerators and compressors.

outside vacuum insulation space. This transfer line is mounted on the ground above the tunnel and interrupted at each service building for branch taps to the satellite refrigerators and magnet strings. The transfer line can distribute 5000 ℓ h^{-1} of helium and 3000 ℓ h^{-1} of nitrogen. Estimated heat leak for the total line is very small: of the order of 3.6 kW (80 ℓ h^{-1}) at nitrogen temperature and 240 W at helium temperature.

Adjacent to the transfer line is a high-pressure (20 atm) helium header that supplies gas from the satellite compressors located at the straight-section service buildings (zero buildings) to the satellite refrigerators located at the sector service buildings (1–4 buildings). In addition, two low-pressure collection headers are located in the tunnel adjacent to the magnets. These headers collect helium and nitrogen from the satellite refrigerators and magnet system. Helium is vented through valves into the header from the magnets only during cooldown or during a quench. These safety valves are mounted on every magnet. Nitrogen from the magnet shield is vented at the end of every cryo string during normal operation. It also can be vented at every spool piece if pressure gets too high.

In principle no large quantities of helium or nitrogen should be vented into the tunnel except under extreme quench conditions (e.g. the simultaneous quenching of every magnet in a sector) or from the rupture of a magnet cryostat. Standard safety procedures do not allow personnel in the tunnel when the magnets are excited, and personnel entering the tunnel (whenever there is liquid in the magnets) are required to carry oxygen monitors and air escape breathing devices.

The satellite refrigerator system consists of 31 screw compressors, each of which requires 350 hp and can deliver 58 g s^{-1} at 20 atm, and 24 refrigerators, each of which contains a cold box, a liquid reciprocating expansion engine, a standby gas expansion engine, and a pair of subcoolers. The refrigerators are located directly above the tunnel and feed helium and nitrogen directly into the magnet string below. The flow is split and goes upstream and downstream through the magnets for typically four half-cells (16 bends, 4 quads, 4 spools) in each direction. At the ends of the string helium flow is controlled by Joule-Thompson valves that connect the single-phase region with the two-phase return region of the cryostat. The two-phase helium is returned along the length of the magnets to the shell side of the refrigerator heat exchanger. Nitrogen for the magnet shields is supplied at the refrigerator feed point and makes a single pass to the ends of the strings where it is discharged into the nitrogen header as 92-K gas.

Under normal (satellite) operation, the satellite refrigerator system uses about 7.0 MW for compressor power, 1500 ℓ h^{-1} (32 tons per day) of nitrogen (including magnet shields), and 4000 ℓ h^{-1} of helium from the Central Helium Liquefier. The Central Helium Liquefier in turn uses

$\sim$2600 ℓ h^{-1} (56 tons per day) of nitrogen and 3.6 MW of power (including all plant usage). This configuration results in 17 kW of refrigeration at 4.6 K and 600 ℓ h^{-1} of helium for lead cooling. As much as 24 kW is available under ideal conditions.

Refrigeration at 100% efficiency would require 66 watts at room temperature for every watt at 4.5 K, and 234 watts at room temperature for every liter per hour of liquid helium produced. In addition 100% efficient nitrogen production would require 161 watts per liter per hour. Typically nitrogen is produced at 30% efficiency and so requires 540 W $ℓ^{-1}$ h. With the above numbers, the Energy Doubler helium cryogenic system is of the order of 10–14% efficient (21).

Helium inventory in the system is 35,000 liters and includes 24,000 ℓ in the magnets and 11,000 ℓ in the transfer line plus liquid storage. Most of the magnet inventory is lost in less than one hour when a site power failure shuts down all compressors. Typical loss rates are 40,000 to 60,000 liters per month, of which $\sim$60% is estimated to be steady-state leakage, $\sim$15% is due to magnet quenches, and the remaining 25% is due to transient small leaks as well as maintenance decontamination and repair.

The cryogenic system has a broad spectrum of operating conditions. In "satellite mode," the central plant supplies large amounts of cold helium to the magnet strings and thus to the return side of the satellite heat exchangers. This excess flow imbalance in the exchangers results in 1000 W of refrigeration without use of the gas expansion engine and with minimal liquid nitrogen (12 ℓ h^{-1}) to do initial cooling of the warm high-pressure helium gas at the input to the cold box. At the other end of the spectrum in "stand-alone mode" no cold helium is used but the gas expansion engines must operate and over 60 ℓ h^{-1} of liquid nitrogen is required in the first heat exchanger. Of the order of 450–500 W of refrigeration plus 25 ℓ h^{-1} of liquid results from this mode and is sufficient to handle only the static heat load of the magnets. A variety of intermediate cases are possible depending on the availability of helium from the central plant.

The magnet strings are cooled by the force-fed single-phase subcooled liquid helium driven by the wet expansion engines. The single-phase liquid is continuously cooled along the length of the string by the return two-phase counter flow heat exchanger. This cooling balances the heat deposited in the string by the static heat load (thermal radiation and conduction through the magnet supports) and the ramping heat load (hysteresis and eddy currents in the magnets). The two-phase return helium changes from mostly liquid at the far end of the string to mostly gas near the refrigerator return, and it is this evaporation that supplies the required cooling. The flow control valves at the ends of the magnet strings are adjusted so that only a small amount of liquid (10%) remains in the end of the two-phase system independent of heat load. Typical flow in the magnet string is 21 g s^{-1}.

Operating temperature of the superconducting magnets is of critical importance as the allowable peak excitation varies by ~14% per degree Kelvin (22). In the Energy Doubler, the operating temperature is set by the requirement that all parts of the refrigerator system operate at positive pressure with respect to the atmosphere so as to reduce the possibility of contamination of the helium gas by leakage of air into the system (air freezes and blocks the cold piping). Because of this, the two-phase region of the magnets operates at 5 pounds per square inch "gauge" pressure (psig), i.e. above atmospheric pressure, and the helium boils at 4.5 K. A ΔT of ~0.4 K is required to carry the heat from the magnet conductor through the insulation, to the single-phase liquid, from the liquid to the stainless single-phase/two-phase tube, and through the tube to the two-phase flow. Thus the magnets are probably operating at ~4.9 K. Lowering the temperature at fast cycle rates more than 0.2 K may be difficult, but 0.4 K may be possible with very slow ramps.

Static heat leak measurements on the magnet strings give 430 ± 60 W for eight half-cells (32 dipoles, 8 quads, 8 spools) (23). The heat load due to ramping is an additional 270 W (500 J × 32 dipoles at 60-s cycle time), making a total of 700 W heat load per refrigerator. The measured static load (~9.5 W per bend magnet) (24) is typically 50–100% greater than expected from calculations.

2.5 *Vacuum System*

The vacuum system for the Energy Doubler consists of three separate subsystems with different characteristics and requirements (25). The cryostat insulating vacuum system is the most complex and is completely isolated from the high-vacuum cold beam tube system inside the magnets. The straight sections and other noncryogenic regions have warm beam tube, bakeable, conventional vacuum systems.

In total the three systems employ of the order of 700 vacuum gauges (Pirani, cold cathode, and ion gauges), 200 ion pumps, 50 turbo roughing pump stations, and 250 electropneumatic valves, all with computer control and readout. The system is divided into 30 individual units corresponding to the 24 cryogenic refrigeration loops and the six warm long straight sections.

All in all there are about 1300 cryogenic interfaces between magnets or between magnets and spool pieces (correction coil modules) or other special purpose modules. A magnet-to-magnet interface (Figure 9) consists of a beam tube seal, two liquid helium connections, one liquid nitrogen connection, and a large external room temperature insulating vacuum seal.

Solder splice joints between the pairs of high-current superconducting leads of the magnets must be made up and insulated inside the single-phase connection. Each of the cryogenic seals must be able to be verified at room temperature with sufficient sensitivity to assure that it will not leak liquid helium. By far the most time-consuming aspect of installation is the interface connection and leak checking. During initial installation each interface took on the average one man-week; subsequent work has been done in about half this time.

The leak-checking procedure is done in three steps. First each device,

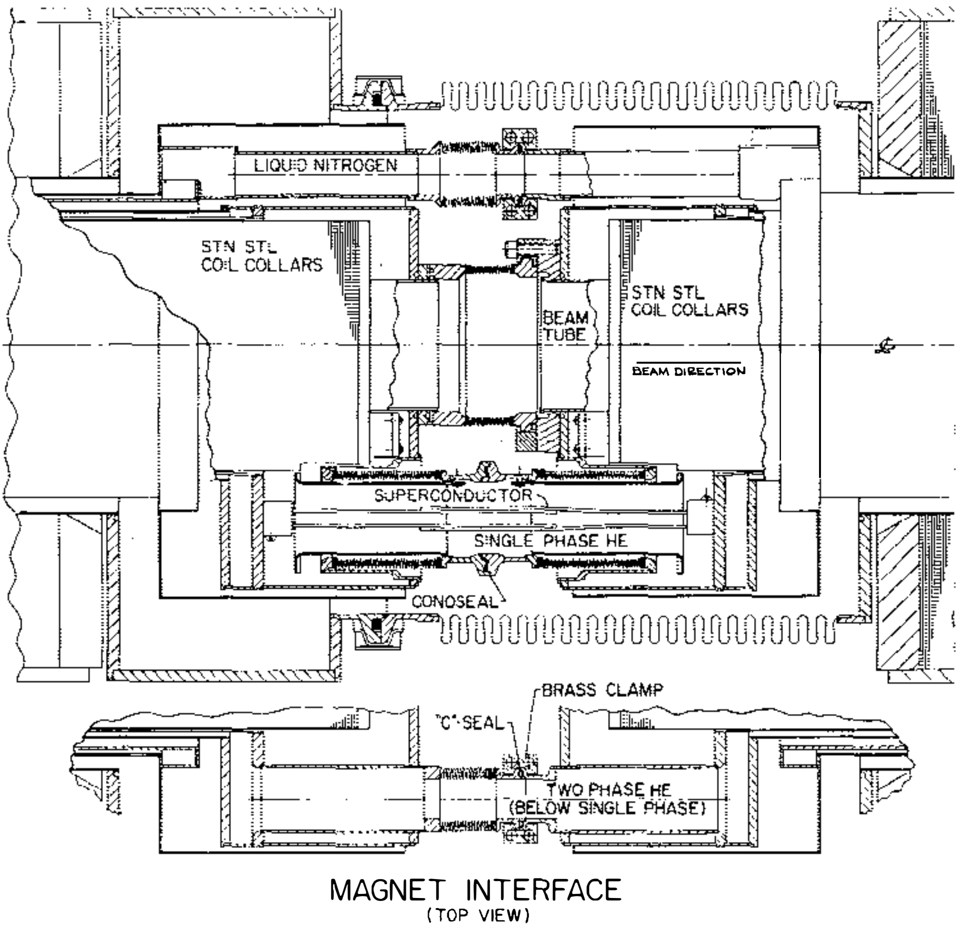

Figure 9 Magnet-to-magnet connecting interface with beam tube, superconductor, single-phase helium two-phase helium, nitrogen, and external room temperature bellows connections.

during its manufacture, is checked for leaks after each major assembly and welding procedure. Finished assemblies are checked for leaks and all devices are cooled to at least liquid nitrogen temperatures (all magnets and spools are cooled to helium temperatures) and rechecked. To pass requires that no leaks be detected, with leak detector sensitivities of better than 2×10^{-10} atm cm^3 s^{-1}. The second step is to install, connect, and test a complete half-cell (4 bends and 1 quad) in the tunnel. At this stage spool pieces are left out so that the ends of the half-cell can be capped off, with each cryogenic circuit brought through vacuum feedthroughs in the temporary end caps. The beam tube connections are first checked by pumping down the beam tube and checking the connections with helium in the usual fashion. (At this stage the large room temperature bellows interface has not yet been closed.) Next, the rest of the interface connections are made, the end caps installed, and the insulating vacuum space of the cryostat pumped down with leak detectors on pumpout ports adjacent to each interface region. After about three hours, leak checking can begin at 2×10^{-10} atm cm^3 s^{-1}. The leak detector output signals are connected to multichannel chart recorders for comparison of responses and permanent record. Each cryogenic circuit (single-phase, two-phase, and nitrogen) is pressurized in turn to 2 atm (absolute) of helium gas and held for 15 min. If a leak is detected it can be located by using the time and magnitude response of the various detectors on the chart recorder. If any leaks are detected, seals are replaced and the string is rechecked. The bore tube vacuum is also checked when the single-phase system is pressurized. This procedure is repeated for each of the eight or nine half-cells in a house (or quarter sector). When all half-cells have been certified, the interconnecting spool pieces are installed and a third and final leak check is performed with a leak detector now on each half-cell. Since there are insulating vacuum barriers in each spool piece, leaks can still be isolated.

The static heat leak of the cryostat-magnet system is due to thermal radiation and heat conduction through magnet supports and other structural elements as well as through residual gas in the insulating space of the cryostat. We find that for our geometry the static heat load doubles at a pressure of 2×10^{-5} torr of helium. We require a pressure below 10^{-5} for operation (which reads 3×10^{-6} on nitrogen-calibrated cold cathode gauges). Typically, vacuum reads 10^{-7} torr (nitrogen-calibrated). The insulating vacuum is pumped with turbo molecular pumps. Additional mobile pump stations can be added at the location of leaks to permit continued operation of the accelerator.

The beam tube vacuum is divided, by isolation valves, into 30 major sections. These consist of 24 250-m cryogenic sections and six 50-m warm straight sections. In addition there are some 23 short warm sections.

The pressure in the cold beam tube is expected to be very low if helium leaks are absent. Pressures of 5×10^{-11} torr cold (5×10^{-10} torr as measured warm) are typical. The cold beam tube provides an economical practical way of obtaining a very high vacuum (required for long beam storage times) over the major fraction of the ring circumference. A warm beam tube inside the superconducting magnets would have required considerable complexity and larger magnet coil diameter to accommodate insulation, in situ high-temperature baking, and conventional pumping. Beam current thresholds for pressure instability cascades induced by beam ionization of residual gas molecules have been estimated to be more than a factor of 30 above expected operating currents (25). In the cascade process the ions are accelerated to the vacuum wall and desorb more molecules.

Helium leaks to the beam tube are extremely difficult to detect or pump without warming the system up because the gas will condense on the beam tube surface and only slowly migrate along the pipe. Because of this, the cryostat was designed so that the beam tube seam weld is the only weld between the helium space and bore tube vacuum.

The warm sections of the beam tube vacuum contain not only pipe but also devices such as rf cavities, kicker magnets, electrostatic septa, magnets with vacuum tubes, and "Lambertson septum magnets," the last of which have their steel laminations exposed directly to the vacuum. Because the Lambertson magnets have the potential for outgassing from large virtual leaks, great care is taken in cleaning and baking them during assembly. Their outgassing rate has thus been reduced to a few times 10^{-8} torr ℓ s^{-1} per meter of exposed laminations.

The fact that most of the circumference of the ring (93%) is cold and has very good vacuum means that the pressure in the warm regions (7%) need not be particularly low. With 5×10^{-11} torr in the cold regions and 10^{-8} torr in the warm regions, the reduction of luminosity during storage due to Coulomb and nuclear scatterings of the beam from residual gas is expected to be $\sim$23% after 20 hours (4). The warm region gas contributes about one-half of this reduction.

The vacuum systems, though large and containing numerous flanges, seals, pumps, valves, and pressure-monitoring devices, have been remarkably trouble free and reliable. This must, to a large extent, be due to the substantial cryo-pumping ability of the refrigerated surfaces.

2.6 *Power Supply and Quench Protection System*

This system (26) has two opposing goals; namely, powering the main bend and quadrupole magnet string and protecting these same magnets from the stored energy in the magnetic field should any fraction of the superconductor in the whole ring become normal (resistive) for any reason. In addition to the power supplies, the power system includes dump resistors to dissipate rapidly the 350 MJ of stored energy. The quench protection system consists of detection units, heater-firing units to distribute the normal region in quenched magnets and thus reduce the energy dissipated in any section of conductor, and current bypass units that direct current around the quenched magnets. Because this system requires rapid detection of quenches and action of subsidiary electrical components in order to save the magnets from self-destruction, it is called "active" as opposed to one that might require little or no external action, i.e. "passive" (diode protection).

In order to understand the protection system it is necessary first to consider what happens in the superconducting cable when it becomes normal and current transfers from the NbTi to the copper (27). The copper, which now conducts most of the current, is generally insufficient to prevent further heating, and the cable will melt unless some means is found to remove the current. The rate at which the cable temperature rises is difficult to calculate because of the nonlinear behavior of the parameters (specific heat, resistivity, thermal conductivity, etc) that describe the cable constituents at cryogenic temperatures. Instead, the cable temperature, after the initiation of a quench, has been measured at discrete constant currents (I) as a function of $\int I^2\,dt$. In the Energy Doubler the peak temperature limit allowed was set at 460 K in order to protect the numerous soft-solder splices and the silver-tin coating on cable strands. At its maximum operating current of 4.4 kA, the cable reaches 460 K after 7×10^6 A^2 s (10^6 A^2 s = 1 MIIT), i.e. there is less than one-half second available for removing the magnet current to prevent permanent damage. At a current of 1 kA, the thermal damage threshold is reached after 12 MIIT.

A schematic of the Energy Doubler magnet power supply system is shown in Figure 10. 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. Each power supply is capable of ramping to 4500 A at 1 kV. Eleven of the power supplies operate in a voltage regulation mode to ramp the current up and down, with either the A2 or A3 supply acting as the system current regulator. Since the resistance of the system is small, the current regulator can provide the required voltage during flattop; hence only this power supply must be capable of conducting continuous flattop current.

The magnets and their interconnection leads are continuously monitored for a resistive voltage component. Once detected, the power supplies are turned off and 0.25-ohm resistive loads are inserted in their place by thyristor switches and a backup dc breaker. The resulting exponential current decay (12-s time constant) is too slow to protect the "normal zone" in the magnet that has quenched. To ensure removal of the current from a quenching element, the magnet circuit is configured into quench protection units, as shown in Figure 11. Eight dipoles and two quadrupoles (one cell) form a typical unit with half of the magnet coils connected into each of the upper and lower bus circuits. A "safety lead" connects the superconducting bus to a room temperature bypass circuit at the ends of each unit. This lead cannot carry steady-state operating currents, but can carry the decaying magnet current around a quenching section of magnets during the dump. Thyristors in each of two redundant external bypass circuits are triggered to allow current to pass as soon as the resistive voltage in the unit exceeds the inductive voltage arising from the decaying current. The fate of the cable depends on the outcome of a race between the cable temperature (function of $\int I^2\, dt$) and the current bypass rate, which is determined by the total resistance of the normal zone. [Longitudinal and transverse normal zone velocities are proportional to I^2 for currents between 1 and 4 kA (28).]

The problem is further complicated in the Energy Doubler by the requirement to protect intermagnet connections, which are not fully stabilized because of geometric restrictions. Quenches originating at these

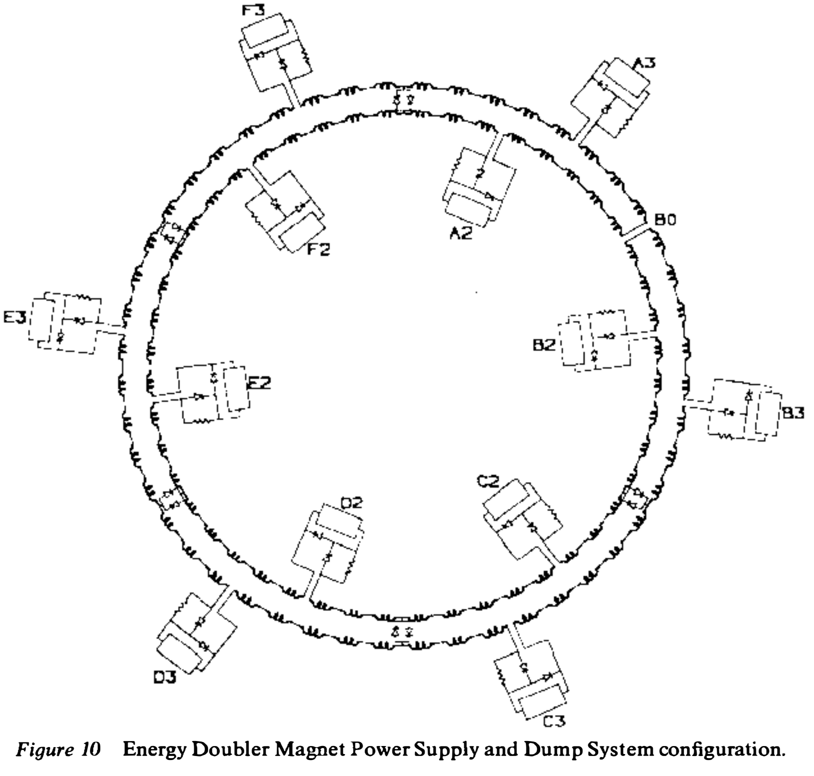

Figure 10 Energy Doubler Magnet Power Supply and Dump System configuration.

locations can only propagate longitudinally and are detected much later than quenches that start within a magnet coil. This uncertainty in quench resistance growth is overcome with resistive heater strips in each dipole. When a quench is detected in a protection unit, energy stored in capacitor banks (heater-firing units) is dumped into the dipole heaters of the protection unit. The resulting rapid resistance growth drives the current into the bypass thyristors and also yields a uniform voltage distribution within the bypassed unit.

The system is controlled by a network of 25 special purpose microcomputers [24 Quench Protection Monitors (QPMs) and one Tevatron Excitation Control and Regulation (TECAR) processor] connected by a serial communication link and redundant hardwire loop cable. The TECAR processor, in addition to controlling the power supply program-

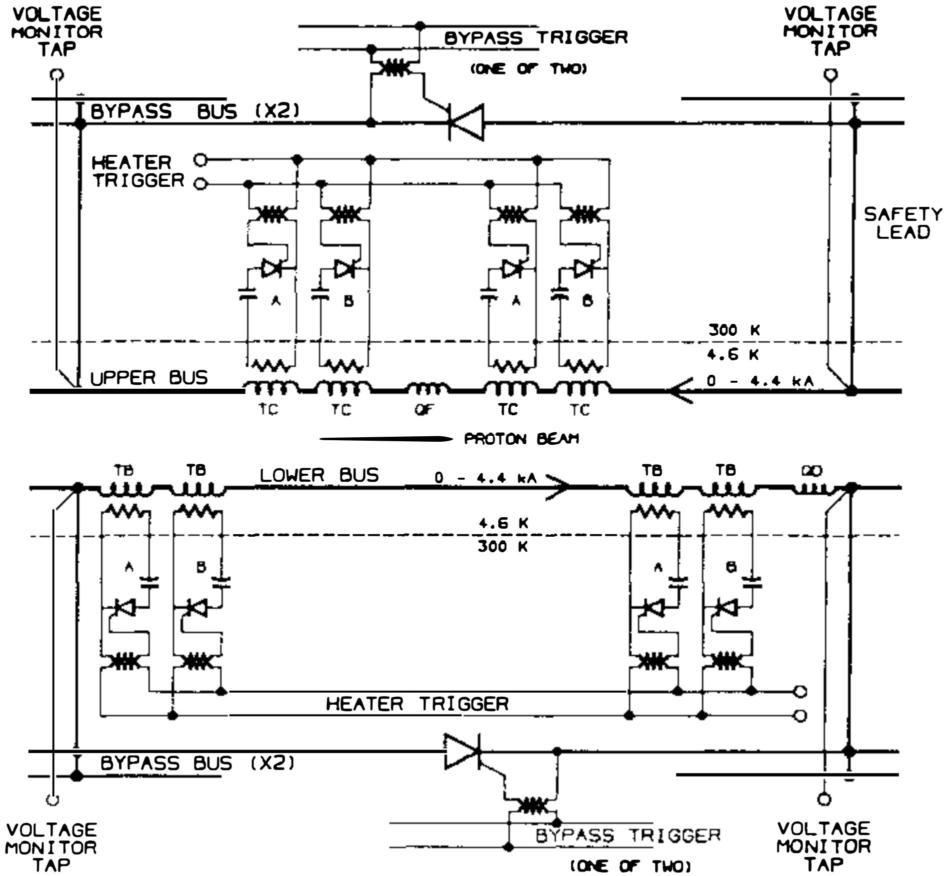

Figure 11 Energy Doubler Quench Protection System for one cell of magnets.

ming and current regulation, coordinates quench protection actions. The system is fail-safe in the sense that element failure results in a shutdown in a manner that protects the magnets from damage. Critical elements and actions are made redundant, and there is continual or frequent verification of the redundancy.

The QPM samples the voltage across each cell (a total of about 200 points) at a rate of 60 Hz; it then calculates the average dI/dt for the cells using the inductance values stored in the memory for each cell. For each cell, the "resistive voltage," equal to the measured voltage minus the voltage due to $L\,dI/dt$, is compared against the quench limits. If the resistive voltage is outside the limits ($\sim 1/2$ V), the protective actions described above are taken. This method of determining dI/dt was selected because it had less electronic noise than did methods that differentiated the current signal from a transductor, and because it avoids heavy reliance upon a communication system. Further checks are made, however, between the "magnet" dI/dt and the dI/dt derived from a transductor.

The QPM also communicates to the refrigerator microprocessor that a quench has occurred and which cell(s) are involved, so that the quench recovery cooldown procedures can be initiated. The recovery time can be as little as 20 minutes if the cooldown begins promptly. Even short delays result in hot gas returning to the refrigerator, which makes the cooldown much more difficult. The refrigeration status is also checked by the QPM so that the current can be removed from the magnets before they quench from insufficient cooling.

In addition to its protective role, the QPM also maintains a record of the event. This is accomplished by continuously storing data in a "circular buffer" six seconds long and then freezing the buffer when an event occurs. The data record includes the cell voltages, the calculated resistive voltages, dI/dt, the current, voltages-to-ground, equipment status, and the voltages across the vapor-cooled current leads (which the QPM also monitors and protects). This data record is essential for properly understanding the system behavior.

The QPMs normally monitor the differential voltages across each cell. In the event of repeated quenches or other anomalous behavior, a special QPM can be installed to monitor the voltages within a cell, typically on every dipole. Faults can usually be localized to a particular device or interface.

2.7 *Control System*

The Accelerator Central Control System (29, 30) for the whole accelerator complex, including the superconducting ring, was redesigned and partially implemented in parallel with the design and construction of the Energy

Doubler. Further implementation is ongoing, along with the p̄ Source construction. The new system thus has diverse requirements.

1. It is an upgrading and conversion of the original Linac, Booster, Main Ring, and Switchyard central controls. It has to be compatible with and adaptable to the interfaces of existing equipment. The conversion has had to take place over a time scale of years, with minimal interference to operation.

2. The Energy Doubler Accelerator requires remotely distributed processors operating independently of the central system. Over 500 microprocessors are in the Doubler system and there are now 1000 in the entire complex.

3. The control system must also be suitable for control of the new p̄ Source rings: the debuncher and accumulator.

4. The system must support not only fixed-target operation as in the past but also simultaneous or sequenced operation for (*a*) p̄ production, stacking, and accumulation in conjunction with fixed-target physics or beam storage; (*b*) injection and storage of protons in the Energy Doubler; and (*c*) unstacking, injection, and storage of antiprotons in the Energy Doubler.

5. The system must provide console (operator station) support to remote facilities such as the Central Helium Liquefier, the Colliding Detector Facility at B0 (and eventually D0), and the RF Building at F0.

The control system can be seen as two separate parts: the central "Host" system and the remotely distributed processors and interfaces for equipment in service buildings around the rings. Communication between the central and remote parts (Figures 12*a*,*b*) takes place via a serial CAMAC link for the Energy Doubler. General specifications for the system include the following features.

2.7.1. CONSOLES All consoles are identical and are able to control and monitor all parts of the accelerator.

2.7.2 DATA RETRIEVAL Selective retrieval is required so that only data required at a specific time is brought back to the central system and only as often as required. Thus instead of an indiscriminate central data pool where all monitor points are updated at a specified repetition rate, as in the old system, now a Data Pool Manager allows for discriminate data retrieval. This selective system was necessary because possible control/readback points increase from 6000 in the old accelerator to over 100,000 with the new Tevatron systems.

Data required at any time include those necessary for applications programs requested by the operators, by data-logged devices, and by alarms scans, though alarms consolidation is done as much as possible at

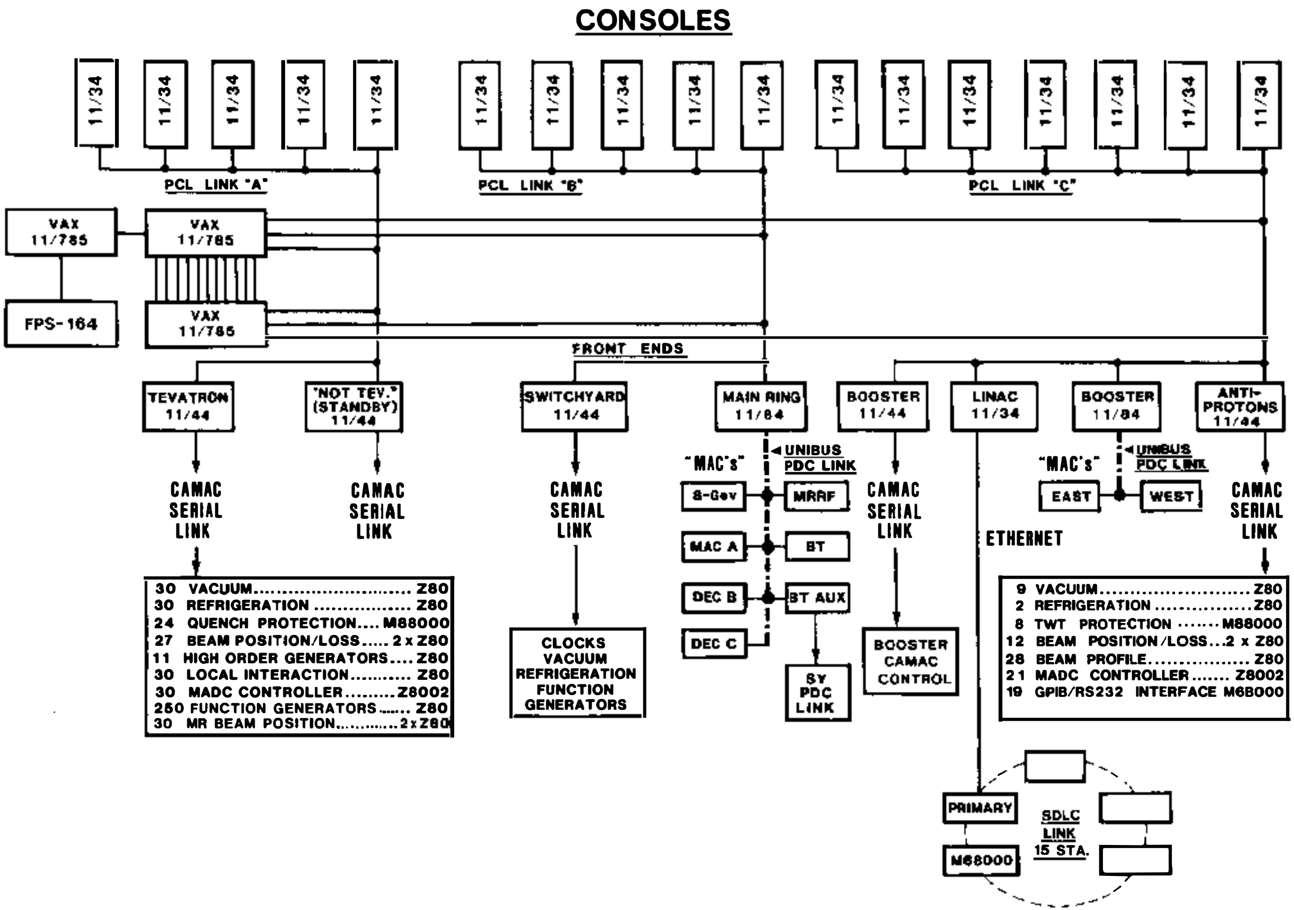

Figure 12a Tevatron Accelerator Control System: Central Host System including console computers, central VAX computers, and front-end

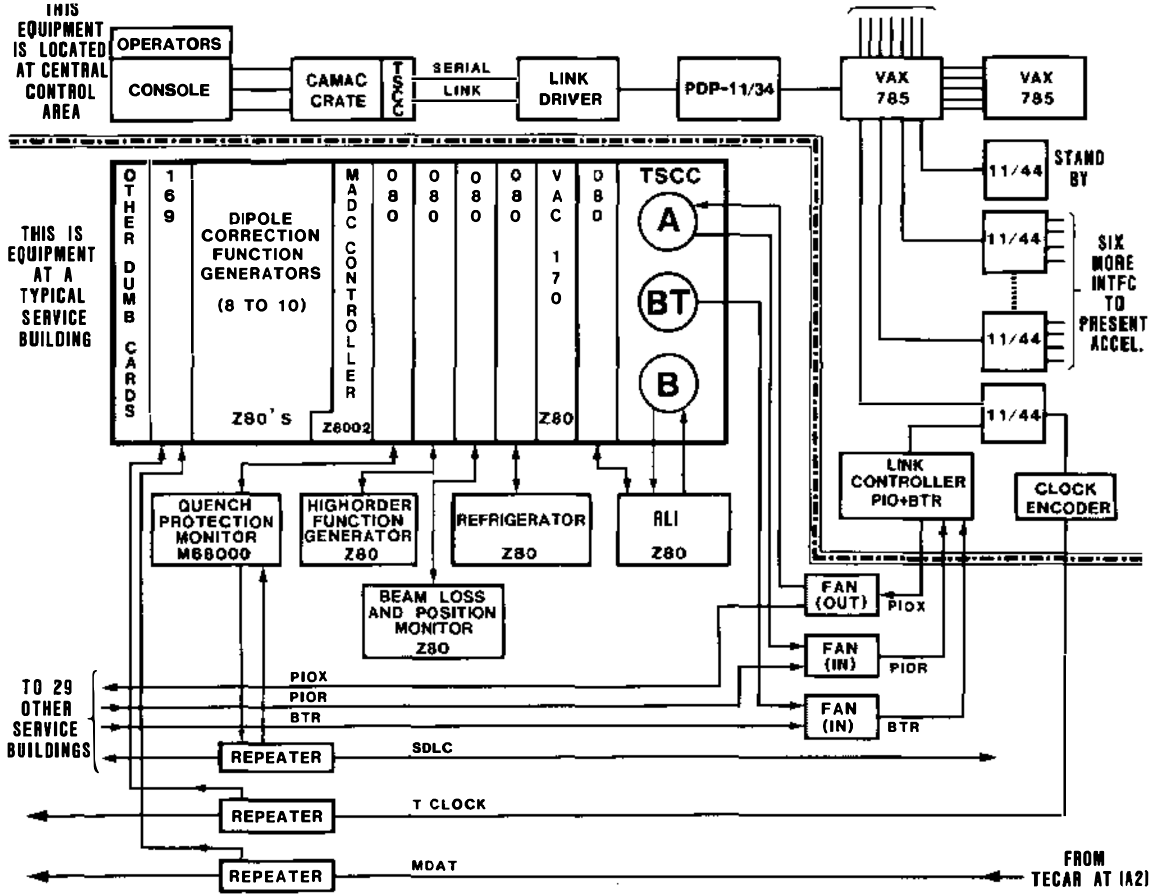

Figure 12b Tevatron Accelerator Control System: remote CAMAC crate with interface to the link and front-end computers on the one hand, and to stand-alone multibus microprocessors and accelerator components on the other hand.

the remote systems. Only small quantities of the data generated and used by the remote processors are ever brought back to the central system.

2.7.3 REMOTE PROCESSORS The remote processors operate separately from the central system either as quasi stand-alone units (refrigerator control processors) or in constant communication with each other (quench protection monitors). These remote systems act as servos in real time and in conjunction with the power supply ramp cycle. Thus each of the refrigeration processors manage 15 control loops and monitor of the order of 100 points at each satellite refrigerator location. The correction coil processors have full adjustment as a function of ramp excitation energy and cycle time ($V_{out} = E^*[f(E)+g(t)]$, where E is the ramp excitation and f, g are arbitrary line segment functions). Additional distributed microcomputers have been provided for vacuum scanning, beam position/intensity/loss monitoring, quadrupole excitation control for resonant extraction, and fast-time plotting of ADC channels. The central computer does *no* real-time control or precise real-time data collecting and thus avoids placing excessive demands on the links.

2.7.4 SNAPSHOT A number of the remote processors have been implemented with circular buffer storage that can save data recorded at time intervals preceding or following beam aborts or quenches. Any part of these data can be retrieved by the central system on operator demand for analysis. No attempt is made to retrieve all buffer data, as the total amount of the remote computer memory is about two million bytes.

2.7.5 LINKS Aside from the three control links from the central control computers to the remote CAMAC crates and multibus microprocessor units, there are five other links around the Energy Doubler ring.

1. Time clock—this link has a real-time clock signal along with "events" encoded on the link.
2. Tevatron beam sync—this is a beam sync clock to track the position of beam bunches in the ring. It allows for precise kicker timing or gate triggering relative to the bunches.
3. Magnet data clock—this link carries information on the Main Ring and Energy Doubler guide-field excitation encoded at a 720-Hz rate.
4. SDLC—this link carries power supply and quench protection information only. It is an integral part of the magnet safety system.
5. Abort link—this link is a fail-safe link. When interrupted it causes the beam to be extracted in one turn from the ring to an abort dump. This link can be deactivated by any one of a number of systems from any of the service buildings. For example, if beam losses above a given threshold are detected or if a quench is detected, the abort link is deactivated. The time

sequence in which different inputs pull on the link during an abort is recorded.

In the Central System there are two VAX 11/785 mainframe processors linked via DEC HSC units, one designated the "Operational VAX" and the other the "Development VAX." The Operational VAX supports a large central data base that contains addressing information for all "control points" in the accelerator system. It also supports certain "central applications programs" not run from the consoles. The Development VAX is used to support software development for computers throughout the control system. Examples of "Central Applications" are logging data for latter analysis, sending alarm messages on to the console computers, and gathering data from the beam position processors as input to the orbit-smoothing program.

Seven "Front End" PDP-11/44-style computers are used to interface the "Central Host" (Consoles and VAXs) to the various discrete accelerator components. A "Front End" computer attempts to hide from the Host computer the individual differences of the various subsystems. This allows the "Host" to treat all requests for receiving or sending data identically. Individual software drivers, such as that for CAMAC and those for communicating with the "smart" subsystems, are located in the appropriate Front End. The Front End also has the task of scanning for any alarms and then alerting the Host to devices that are at values outside their nominal limits. This includes the unsolicited messages generated by the "smart" subsystems as well as LAMs ("look-at-me" signals), which can be generated by "dumb" CAMAC modules such as power supply controllers.

There are 17 identical control consoles. Each console is supported by a DEC DP-11/34 with RSX/11 operating system, and up to four simultaneous user tasks. Each console is interfaced via a CAMAC crate to the PDP-11 and can be an arbitrary distance from it. Console equipment includes a keyboard, track ball, interrupt button, two Hitachi color monitors, a Lexidata precision color graphics monitor, and a Tektronix 613 storage scope. All PDP-11s and the two VAXs are interconnected with DEC PCL (Parallel Communications Link) busses. Three such busses are used since the fixed-time slice architecture does not permit a very high multiplicity on any one bus.

2.8 *Beam Loss and Resonant Extraction*

As discussed in the cryogenic section, the maximum operating field of the superconducting magnets is critically dependent on the temperature of the coils. Loss of even a small fraction of the circulating beam into the magnets will result in cascade showers, energy deposition, and heating in the superconductor. This heating may be sufficient to cause the magnets to quench even at low excitation; at high excitation where there is little ΔT margin, the peak accelerator energy may have to be reduced for reliable operation. In a system where quench recovery typically takes an hour, one must try to understand and reduce the effects of beam loss.

Beam loss in the accelerator fits into one of two categories: accidential and unavoidable. By unavoidable we mean loss produced during standard operation of the accelerator; by accidental we mean loss that can be minimized by being careful and by reducing failures or mistakes. The resonant extraction process of the beam from the accelerator to the fixed-target areas is an example of a standard operational mode that produces unavoidable losses. In fact, the fast resonant extracted beam pulses, which produce energy deposition over a short time interval (2 ms), must be carefully analyzed to determine if they are feasible at all.

Calibration studies with simple geometries have been carried out to determine how sensitive the magnets are as a function of excitation and time duration of energy deposition (3, 31). These studies consist of measuring the beam level at which a magnet quenches and comparing that to the calculated deposition per incident particle. In this way a design guideline of less than 4 mJ g^{-1} at 800 GeV was established for millisecond-like beam pulses. Modeling of the extraction process then allowed us to determine what intensity is possible and how to design straight-section components that maximize this intensity.

Resonant extraction is accomplished by producing a half-integer stop-band with quadrupoles distributed on the 39th harmonic of the accelerator circumference ($\nu = 19\frac{1}{2}$). Octupoles with the same phase relationship are used to provide a nonlinear amplitude-dependent tune so that large amplitude oscillations are unstable. As the quadrupole strength is slowly increased, the stable region shrinks down on the beam and particles stream out along specific trajectories (separatrices) of the phase space plot. The oscillation amplitude of particles grows as they progress away from the stable region; when this growth becomes of the order of 1 cm every two turns, it is possible to "split" the beam with an electrostatic septum. Particles that have "jumped" from the field-free region to the high-field gradient region are deflected sufficiently then to jump a magnetic septum on the other side of the accelerator ($9\frac{3}{4}$ oscillations away) and be deflected into the extraction beam line.

Statistically it is not possible for all particles to jump the electrostatic septum without any hitting the 2-mil wires that make up the anode plane. Of the order of $1\frac{1}{2}\%$ of the beam hits the wires, and it is this beam and its resultant interaction particles that must be tracked.

Monte Carlo analysis using CASIM (32, 33) and similar programs is done in four steps. First the elastically and inelastically scattered particles from the septa are calculated and their energy and direction determined. The particles are then tracked through a geometry that models the accelerator configuration. Inelastics need only be tracked through the first few downstream accelerator magnets. Elastics must be tracked through the machine until they hit a magnet or are extracted. This tracking is followed for up to two turns. Once the locations of hits in the superconducting magnets are recorded, energy deposition in regions with a large "hit probability" can be calculated. Finally, the expected response of loss monitors in the region is also calculated.

Results show that of the particles that hit the septa wires, 20% are inelastics deposited locally downstream of the septa. Because of the magnetic shielding design (resulting from the calculations) in the straight section, only 7% of these actually enter the superconducting magnets. The other 80% of the beam hitting the septa are elastically scattered; 10% are lost in the half-turn between the septa and the extraction channel; 5% are lost in the extraction channel; and 65% leave the machine via the extraction channel on the first or third turn.

It is the $1\frac{1}{2}$% inelastics and the 10% elastics lost in the half-ring, or less than 2×10^{-3} of the total beam intensity, that must be considered in detail. Calculations show that beam should hit in about a dozen locations, in good agreement with what is observed. Exact predictions are sensitive to the local closed-orbit beam position relative to the magnet bore tube and vary by about a factor of 3 for 1-mm changes. Particles hitting the magnets deposit about 40% of their energy in the superconducting coils.

Comparison of calculations with actual measurements of beam quench levels predicts that 4–15 mJ g^{-1} of energy is deposited locally in the coil to produce the quench when of the order of 2.5×10^{12} protons are extracted in a fast pulse. Considering the overall complexity of the calculations, the good consistency of the results with early calibration studies leads one to believe that these calculations are a viable design tool and can be used for further shielding and scraper design improvements (34).

3. ACCELERATOR LATTICE AND CONVENTIONAL SYSTEMS

In the previous section, aspects of the accelerator that are specifically associated with or strongly impacted by the superconducting nature of the magnets were discussed. In this section we describe other systems of the Tevatron, which, though impacted by the superconducting design, are basically required in any accelerator. The lattice, correction coil configuration, diagnostics, and radiofrequency acceleration system are some of these. The accelerator parameters are given in Tables 1 and 2 and the general layout was illustrated in Figure 2.

3.1 *Lattice*

As already described, the lattice for the Energy Doubler follows that of the Main Ring closely. The $15\frac{1}{2}$ normal cells per sector follow the pattern of focussing quad, 4 bends, defocussing quad, 4 bends. The bend magnets have a magnetic length of 6.12 m and an overall length of 6.40 m; the quadrupoles have a magnetic length of 1.68 m and a focal length of 26 meters. Adjacent to each quad is a 2.16-m space for the correction coil spool piece. Spacing between quads (half-cell length) is 30 meters and the betatron phase advance per cell is 68°. One of the half-cells is exceptional in that two of the bending magnets are omitted as in the Main Ring to provide space for special warm components.

Generally the fact that the Energy Doubler must follow the Main Ring geometry so closely is very constraining, especially in the straight-section regions for injection, extraction, minimization of the beam loss on the superconducting magnets, and in the design of the low-β regions for collisions. A number of compromises in design have been made, which would not have been necessary if the lattice had been designed independent of the Main Ring. Remember that beam energies of 150–1000 GeV instead of 8–400 GeV must still be handled by conventional iron magnets and electrostatic devices in spaces designed for the lower energy. Deviations from the Main Ring lattice do take place at the six 50-meter straight sections (A0–F0). Figure 13 illustrates the layout of these regions.

Two of the straight sections, E0 and F0, where injection and rf acceleration take place, are optically very similar to the Main Ring; the only difference is that the insertion quadrupole doublet on either side of the straight section is made up of two quadrupoles instead of four. At C0 where the beam is extracted to a dump when aborted by a one-turn kicker, superconducting bending magnets have been shortened on either side of the straight section, and conventional extraction magnets with equivalent $\int B\, dL$ have been substituted in the center of the straight section. Not only does this make space for the abort kicker, but it provides a certain amount of shielding for the superconducting magnets at the downstream end of the straight section.

The straight section A0 contains the extraction magnets to divert the beam into the switchyard for fixed-target operation. In both A0 and D0 the lattice quadrupoles on either side of the straight section have been reconfigured so as to produce a large horizontal amplitude function (β), at the upstream end of the straight section, suitable for fast extraction. This

Table 2 The accelerator parameters

General

Number of dipoles	774
Number of quadrupoles (without low-β)	216
Number of sectors	6
Good field region	± 2 cm horizontal ± 1.5 cm vertical
Tune (ν)	19.41 horizontal 19.38 vertical
Natural chromaticity (ξ)	-22

Lattice

Lattice for sector arc: 2(F, 4B, D, 4B) F, $S_1$2B, D, 4B, 12(F, 4B, D, 4B)
Warm length $S_1 = 12.5$ m
β_{max} 100 m β_{min} 29 m
η_{max} 6 m η_{min} 1.2 m

Lattice for normal straight section:
F, S_2, 3B, F, D, LS, F, D, 4B, D, 4B
Warm length long straight $LS = 51$ m
Warm length $S_2 = 6$ m
β_{max} 110 m, β_{min} 60 m in straight section

Lattice for high-β straight section:
F, S_2, 3B, D, F, LS, D, F, 4B, D, 4B
β_{max} 243 m, β_{min} 34 m in straight section

Lattice for low-β straight section: F, S_2, 3B, F, D, F, D, F, CS, D, F, D, F, D, 4B, D, 4B
Warm length collider straight $CS = 13$ m
β_{max} 900 m, β^*_{min} 1 m

Radio frequency

Longitudinal phase space:
0.3 eV-s fixed-target, 3.0 eV-s collider
rf frequency 53 MHz
Harmonic number 1113
rf voltage, 8 cavities at 300 kV each

Power supply

System stored energy 350 MJ
Total inductance of ring 36 h
Dipole inductance 0.045 h
Total resistance of ring:
90 m ohm (ring) + 90 m ohm (filter chokes)
Power supplies, 12 at 1 kV, 4500 A
Volts to ground (normal) 500 V peak
Volts to ground (fault conditions) 1.5 kV maximum
Volts bus-to-bus 1–1.5 kV peak

Injection

Single turn (single bunch for collider)

Abort

Single-turn extraction

Extraction

Half-integer resonant

lattice change does not affect the optics elsewhere in the machine but magnifies the horizontal beam size at the position of the electrostatic septa (D0) and magnetic Lambertson septa (A0). This magnification is crucial for obtaining efficient extraction and yet maintaining a reasonably small magnetic aperture in the superconducting arcs of the ring.

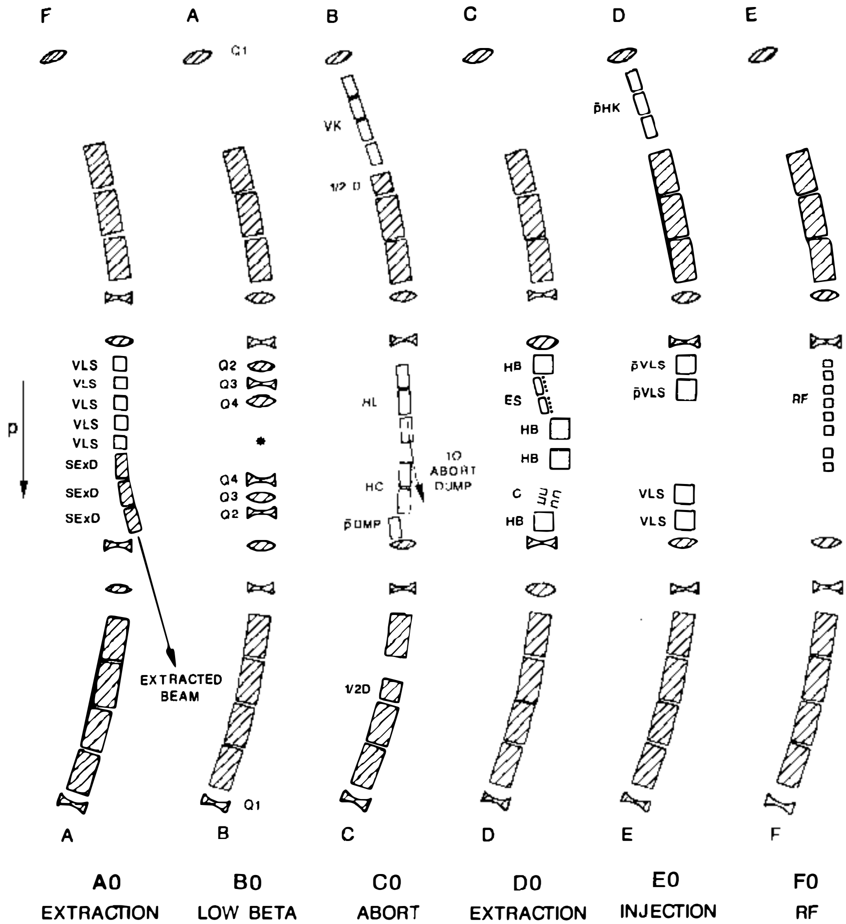

Figure 13 Straight-section layout A0–F0. Cross hatching indicates superconducting elements. VLS = Vertical bending Lambertson Septum magnet; SExD = Skew Extraction Dipole; VK = Vertical Kicker; 1/2D = half-length Doubler dipole; HL = Horizontal Lambertson septum magnet; HC = Horizontal C magnet; HB = Horizontal Bend; ES = Electrostatic Septum; C = Collimator; and HK = Horizontal Kicker.

The straight section D0 will be reconfigured depending on whether fixed-target or collider operation is underway. For the fixed-target operation, a high-β lattice configuration identical to that at A0 is used. In addition to the standard superconducting magnets on either side of the straight section, a chevron arrangement of conventional magnets is installed in the straight section, with the extraction electrostatic septa located after the first of these magnets. This configuration was devised in order to protect the downstream superconducting magnets from beam spray off the septa, as discussed in a prior section. For collider physics, straight sections B0 and D0 will be used for large solid angle detectors. In the case of D0, the extraction bends and electrostatic septa will be removed and the low-β quads and experimental detector installed. Changeover time will be less than two weeks.

The B0 straight section has been modified in order to produce a low-β^* of 1 m at the collision point. The quadrupole layout is such that the optics can be adjusted between the standard fixed-target optics (like E0 and F0) and the low-β optics by adjusting four quadrupole circuits (35).

In addition to the standard Collins quadrupole doublets at either end of the straight section, which are powered in series with the whole superconducting ring, triplets have been added inboard. Pairs of elements of the triplet (either side of the collision point) are powered in series by separate supplies. In addition, the quads a half-cell removed from the straight section need to be separately powered as a pair. During normal ramping these latter must track the normal excitation ramp and the other three supplies are off. During the process of squeezing the β^* from 70 m to 1 m with the energy held constant, all four of these supplies must simultaneously follow prescribed excitation curves in order not to perturb unduly the optics outside of the interaction straight-section region. Magnetic gradients of 25 kG per inch are required in the low-β quads. These quads have a 3-inch inside coil diameter, a copper-to-superconductor ratio of 1.3:1 and a 20-μm filament size.

In this design we have placed a minimum of reliance on the operation of the special low-β quads during fixed-target physics (one circuit). We have also allowed for continuous tuning from the fixed-target lattice to the low-β lattice, whereas in reality injection and acceleration for storage could probably take place at an intermediate β^* value. We have also chosen to minimize the number of special quads and power supplies for economy of high-current cryogenic power leads and separate high-current supplies and quench protection. Because of this simplicity the amplitude function, $\beta(z)$, and dispersion function, $\eta(z)$, $[\Delta x(z, \Delta p/p) = \eta(z)\Delta p/p]$ are not completely matched to fixed-target values around the ring. In particular, the dispersion is $\eta = 0.2$ m at the collision point and 60% larger in the arcs than for fixed-target operation.

Part of the preparation for collider operation is a reconfiguration of the Main Ring at B0 and D0 (36). Because the Main Ring will be accelerating and extracting beam for $\bar{p}$ production during the time the Energy Doubler is in storage mode, it is desirable not to have the Main Ring beam going through the center region of the detectors. Otherwise there would be difficulties in Main Ring acceleration with the detector magnetic fields turned on as well as background and radiation damage in the detector, produced by the Main Ring beam. The B0 overpass allows the Main Ring to bypass the detector completely; the D0 overpass bypasses the central region and argon calorimeter of the D0 detector.

The Main Ring is arched vertically away from the Energy Doubler in these regions in so-called overpasses; this allows for a distance between beams of 21 feet (B0) and $6\frac{1}{4}$ feet (D0). In order to generate the necessary deflection, vertical bending magnets will be inserted in the Main Ring lattice 250 meters either side of B0 and 105 meters either side of D0 straight-section centers. These magnets and the two adjacent horizontal bends, along with similar down bends near the straight sections, are operated at twice the field of the regular Main Ring magnets and limit the Main Ring to a maximum of 200-GeV operation. Separate tunnel enclosures for the Main Ring near B0 are also required.

3.2 *Correction and Adjustment Magnets*

Correction magnets are those required to correct field imperfections and alignment errors of the main quadrupoles and bend magnets. Adjustment magnets are those required to tune the optics depending on the desired operating conditions. Often the same "correction" magnets perform both functions. Thus, dipole steering magnets are necessary to compensate for alignment errors and put the beam in the center of the aperture, but they are also necessary to bump the beam away from the center where this is required to avoid certain restrictions like injection or extraction magnets. Chromaticity sextupole circuits adjust the variation of the tune of the machine (number of betatron oscillations per turn) with momentum, $\Delta\nu(\Delta p/p) = \xi\Delta p/p$. These circuits are necessary as adjustments, but they are also necessary as corrections to the dipole magnet persistent current sextupole error at low excitation.

In the Energy Doubler most of the correction magnets are superconducting and are designed with strengths sufficient for use at full excitation (37). Their power supplies can be programmed to different values continuously throughout the acceleration cycles. The correction magnets are located in "spool pieces" adjacent to each of the more than 200 quadrupole magnets in the lattice. Most spool pieces contain two "packages" of three concentrically wound correctors with six pairs of leads coming from cryo temperature to room temperature. In principle all correctors (over 800) could be powered independently but, in fact, only the dipole steering elements need independent control—the rest are connected in series combinations to produce "correction coil families" powered by single supplies. Originally, putting the correctors inside the quadrupole magnet was considered; however, it was expedient not to complicate the quadrupole design by having a separate spool with its quench protection safety leads and correction elements.

The first correction package (DSQ) contains a dipole (horizontal or vertical), normal sextupole, and normal quadrupole for steering, chromaticity adjustment, and tune adjustment. The second package (OSQ) contains an octupole, a sextupole (normal or skew), and a quadrupole (normal or skew). Not all of these units are connected and powered. (See Table 3.)

Power supplies for the correction elements are of two types: one for single elements (dipoles), or possibly two to three elements in series, with ± 15 V, ± 50 A output, and a current stability and ripple limit of 0.1% of full scale; a second high-precision supply (38) used for 4 to 90 elements in series strings. These precision supplies are rated at ± 500 V, ± 50 A with a stability of better than 2.5 ppm per °C and output voltage ripple less than 30 mV peak to peak. Reference voltage to these supplies is provided by an analog-smoothed high-precision 16-bit D/A (39). Every effort to minimize ripple has been made in order to minimize intensity modulation of the slow extracted beam. This requirement of $d\nu/dt < 0.0001\ \mathrm{s}^{-1}$ is expected to be more demanding than the requirements of storage.

Desired excitation of the circuits typically is derived from formulas of the form given below for the tune chromaticity adjustment circuits at the focussing and defocussing quadrupoles (40):

$$
\begin{aligned}
I_{\mathrm{QF}} &= E(39\Delta\nu_x + 11\Delta\nu_y - 3.3b_1) \\
I_{\mathrm{QD}} &= -E(12\Delta\nu_x + 43\Delta\nu_y + 3.3b_1) \\
I_{\mathrm{SF}} &= E(0.17\Delta\xi_x + 0.05\Delta\xi_y - 3.3b_2) \\
I_{\mathrm{SD}} &= -E(0.09\Delta\xi_x + 0.30\Delta\xi_y + 5.1b_2)
\end{aligned}
$$

where I is in amps, E is in TeV, and $\Delta\nu_{x,y} = \nu_{x,y} - 19.4$, $\Delta\xi_{x,y} = \xi_{x,y} + 22.5$. Here 19.4 is the natural tune of the accelerator and -22.5 is the natural chromaticity. Coefficients b_1, b_2 are the average bend magnet multipole coefficients in units of 10^{-4} at one inch.

Table 3 Correction coil circuits and their uses

	Number of circuits/ number of elements per circuit	Strength each element (kG-in. at 1 in., 50 A)
Steering dipole		
Horizontal and vertical	224/1	181
Quad tune adjustment		
Horizontal and vertical	2/90	75
Chromaticity sextupole		
Horizontal and vertical	2/90	57
Skew quad coupling adjustment		
0 harmonic ($\nu_x - \nu_y = 0$)	1/48	75
Horizontal-to-vertical coupling compensation		
39th harmonic quad		
Horizontal sin and cos terms for 39th harmonic ($2\nu_x = 39$) 1/2-integer stop-band adjustment for resonant extraction	2/4	75
39th harmonic octupole		
Horizontal sin and cos terms 1/2-integer resonant extraction adjustment	2/16	33
0 harmonic octupole		
Horizontal and vertical average octupole compensation	1/24	33
Extraction adjustment	1/12	—
Circuits below not powered		
39th harmonic quad		
Vertical sin and cos terms allow for simultaneous $2\nu_x$, $2\nu_y = 39$ compensation	8/2	75
39th harmonic skew quad		
Coupling resonance sin and cos terms compensation for $\nu_x + \nu_y = 39$	4/2	75
58th harmonic normal sextupole		
Sin and cos terms for $3\nu_x$, $\nu_x + 2\nu_t = 58$	14/2	44
1/3-integer resonance compensation	4/1	—
58th harmonic skew sextupole		
Sin and cos terms for $3\nu_y$, $2\nu_x + \nu_y = 58$	14/2	44
1/3-integer resonance compensation	4/1	
0th harmonic skew quadrupole	12/1	75

3.3 *Beam Position and Beam Loss Monitors*

Beam diagnostics in the Energy Doubler are still under development and rapid evolution, especially for storage operation. However, considerable effort has been taken with recording and displaying data from the more basic diagnostic devices so that information is available for analyzing causes of beam-induced quenches, aborts, and losses. The fact that recovery from quenches can take an hour or more makes it imperative that reasons for erratic beam behavior be found with as few beam pulses as possible, and at best with just one.

The basic diagnostic devices are the beam position and loss monitor systems (41). In general large amounts of memory have been incorporated into the electronics associated with the devices so that information related to first-turn injection and to 1000 consecutive turns, or to 500 samples at 15-ms intervals prior to an abort, can be automatically recorded, as well as data at other desired times. The Beam Position Monitor (BPM) system consists of about 200 position detectors and associated rf, analog, and digital circuitry. The detectors are a pair of directional strip line plates 20 cm long and can independently measure the beam in two directions. A detector is mounted at the end of a superconducting quadrupole inside the cold beam tube. Signals from the detectors go to the nearest service building for processing. They first pass through a 53-MHz resonant filter to an amplitude-to-phase modulation unit that produces an analog voltage proportional to the ratio of signals from the two plates. The resonant filter is necessary for isolated single bunches in collider operation. In addition to the AM/PM processing, the two-plate signals are added to provide intensity information. Position and intensity signals are then processed and digitized. There are of the order of 8 to 10 detectors per service building, with parallel analog processing for each detector. The digitized position and intensity information is read into a microprocessor, which can store either the single-turn information or average information over a selected time period in order to find the equilibrium orbit independent of coherent betatron oscillations.

The position system has a least count resolution of 0.15 mm. The dynamic range of the system is dependent on the number of rf buckets around the ring circumference filled with beam; its minimum threshold is 5×10^9 protons in less than 50 consecutive buckets, its maximum is 3×10^{13} protons in a full ring (~ 1000 buckets). This wide dynamic range is required in order to avoid quenches during initial turn on, as well as for collider operation.

Data from the beam position monitors are used in various ways. One can plot beam position as a function of location around the ring for one-turn information and for closed-orbit information. This allows one to make corrections with the dipole correctors over the whole accelerator cycle. On any two detectors one can obtain position information on 1000 consecutive turns. This allows one to observe and correct injection errors in either the transverse or longitudinal phase space by observing coherent betatron and synchrotron oscillations. By using a fast Fourier transform routine with these data, one can obtain coherent tune measurements either at injection time or later in the cycle by inducing a coherent oscillation. By looking at the turn-by-turn information during times of erratic operation, one can see if coherent oscillations have been induced by some unknown mechanism. If two detectors in the same plane (horizontal or vertical) with $\sim 90°$ phase advance between them are used for the turn-by-turn information, it is possible to plot the phase space (x, x') mapped by the coherent oscillation. In Figure 14 a closed-orbit position plot around the ring taken during beam extraction time shows a large deviation from center line to avoid injection devices at E0. The small horizontal deviations in D and F sectors and in the D0 and A0 straight sections are made so as to reduce beam loss from particles scattered off the electrostatic septum during extraction. A turn-by-turn plot at injection time of a horizontal detector (Figure 15) illustrates coherent betatron and synchrotron motion with and without the use of the fast beam damper. Data are displayed for both 128 turns and 1024 turns.

Beam loss monitor (BLM) detectors are located near each quadrupole in the ring. Their electronics are packaged along with the BPM electronics.

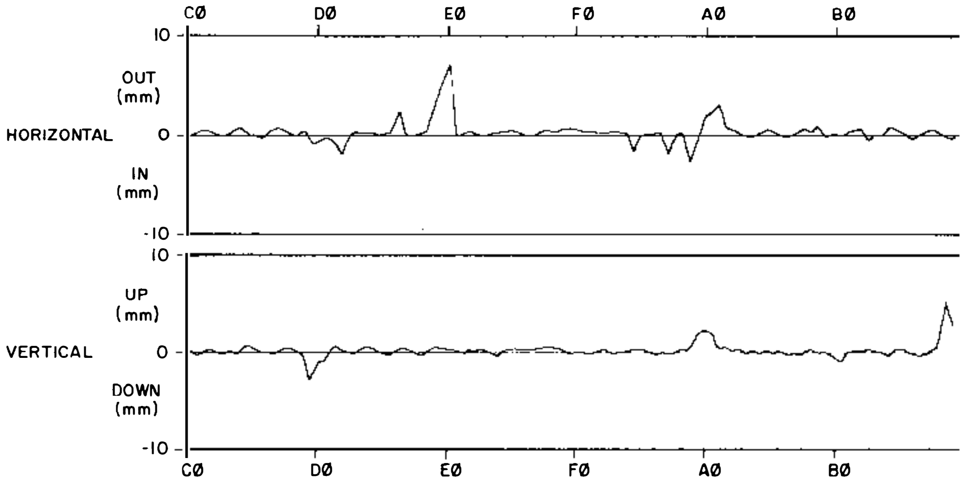

Figure 14 Horizontal and vertical closed orbit measured by the beam detector system at 800 GeV during resonant extraction.

Additional BLMs are installed in straight-section regions associated with extraction or abort. The BLM detector consists of a sealed glass argon ion chamber with a sensitivity of 2×10^{-8} coulombs per rad (C rad^{-1}) and linearity to 100 rad of instantaneous loss. The electronics for each detector consists of a 4-decade integrating logarithmic amplifier with an integrating time of 1/16 second to try to match the quench properties of the superconducting magnets so that the output voltage is related to the probability of quenching the magnets. Outputs of the loss monitors are continuously checked so as to abort the beam if the abort threshold is exceeded. Signals to abort the beam can be generated within 200 μs. Though abort triggers can be generated by numerous other signals, such as the kicker not being ready or the beam position being out of tolerance, the loss monitor system is by far the most used for preventing beam-induced quenches. Figure 16 illustrates loss patterns obtained during a quench induced by too much fast extracted beam; it also indicates abort thresholds. Loss readings are higher in the straight section because the detectors are not shielded as well by the magnets. Loss points in sectors D, E, F are typical of those caused by extraction beam spray.

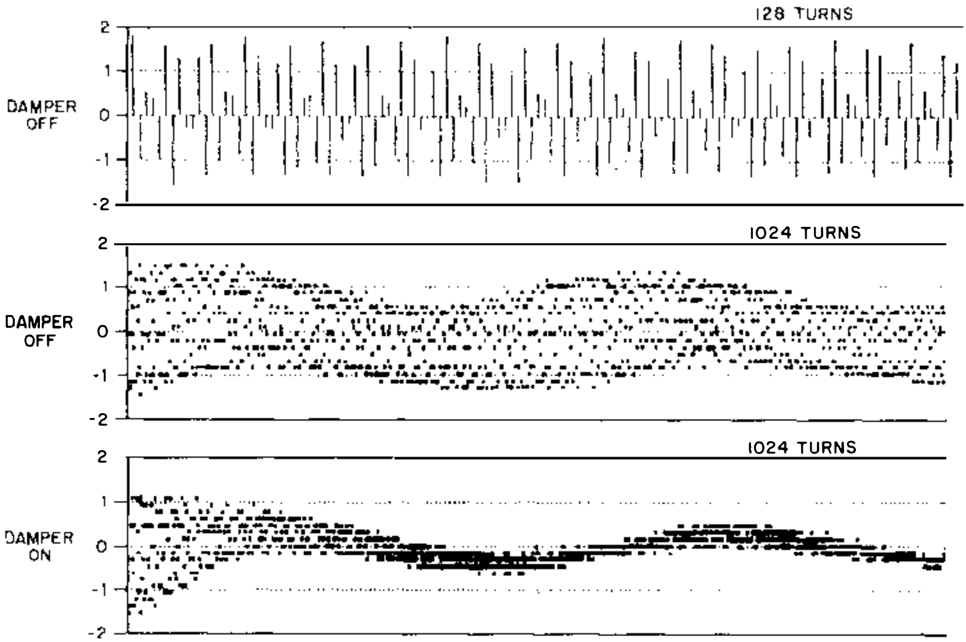

Figure 15 Turn-by-turn plot illustrating coherent betatron and synchrotron oscillations and showing betatron oscillations being damped by the beam damper system.

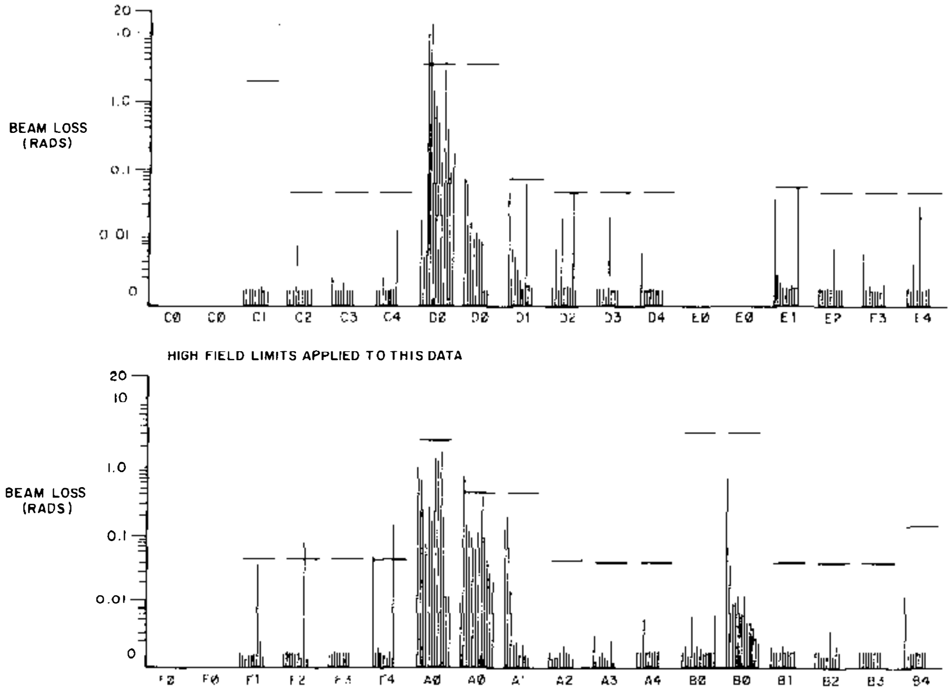

Figure 16 Beam loss reading during a quench induced by fast extraction losses.

3.4 *Beam Damper System*

The beam damper system (42) is used to inhibit beam blow up caused by electromagnetic coupling between the beam and the vacuum chamber wall or other devices installed in the ring. The damper can work on any collective betatron motion that the beam in individual rf buckets undergoes as a whole. The damper system is also a diagnostic device that can measure betatron motion of individual buckets, can damp or antidamp individual buckets selectively, or can antidamp small amplitude oscillations and damp larger amplitudes; this allows for continuous tune measurements on individual buckets during acceleration and extraction. We have just begun to explore the potential of this device.

The damper consists of two separate parts. The first is the low-level measurement and storage electronics, which also incorporates analysis and display electronics; the second is the high-level power system and deflector plates that act on the beam. The system is similar to that used in the Main Ring but it is considerably more powerful. The principle is very simple. The position of the beam in each rf bucket is measured as it goes by pickup electrodes. On the next turn it passes damper electrodes after $(n+\frac{1}{4})$ betatron oscillations and it receives a transverse deflection proportional to its measured displacement from center. Thus, after successive turns a coherent oscillation can be reduced to zero.

The low-level electronics consists of flash ADCs followed by a specialized "pipe line" computer for computing the beam position and multiplying it by the proper gain transfer function. The flash ADCs and computer are clocked at the 53-MHz frequency of the rf system, and the delay in the electronics is just equal to the time of flight of the beam around the ring so that the correct beam bucket is acted on.

The high-level electronics for each plate consists of a 24-tube (4CW800F tetrode) distributed amplifier feeding 50-ohm lines to the deflector strip line. Output to each plate is zero to $+1400$ V, with a 10–90% response of $3\frac{1}{2}$ ns. The gap between the plates is $2\frac{1}{2}$ inches and the length is $\lambda/4$ (of 53 MHz). The filling direction is opposite to the beam direction and a pulse length of the order of 12 ns is required.

3.5 *Flying Wire*

The flying wire (43) is a device to measure the profile of the beam and thus deduce the beam emittance if the amplitude function (β) is known or to deduce relative amplitude functions at different locations with simultaneous measurements. The principle is very simple: A thin wire is quickly rotated through the beam in either the horizontal or vertical direction and losses from the beam hitting the wire are picked up in a photomultiplier counter. (Alternatively, charge depletion from the wire can be measured directly.) The signals are proportional to the beam intensity as a function of wire position.

This very simple concept has some interesting engineering problems associated with it. Namely, the wire must go fast enough so that it does not burn up or unduly scatter the beam (carbon or beryllium 4-mil wires are used). The data-digitizing rate must be fast enough to obtain sufficient data points over the beam profile, and the wire must move in a reasonably continuous linear fashion when in the beam and stop after one transversal. Typically, velocities of 5–10 m s^{-1} and a digitizing rate of 47 kHz (revolution frequency of the beam) are used. Figure 17 illustrates wire profile measurement at $\beta = 100$ m, 1 m for 800-GeV beam.

3.6 *Radiofrequency Accelerating System*

The rf accelerating system (44) consists of eight 53-MHz cavities driven by pulse amplifiers. The cavities each produce a peak accelerating voltage of 1/3 MV and are coaxial, two-gap resonators, 12 inches in diameter and 160° total electrical drift tube length, or 108 inches total length. The cavities are

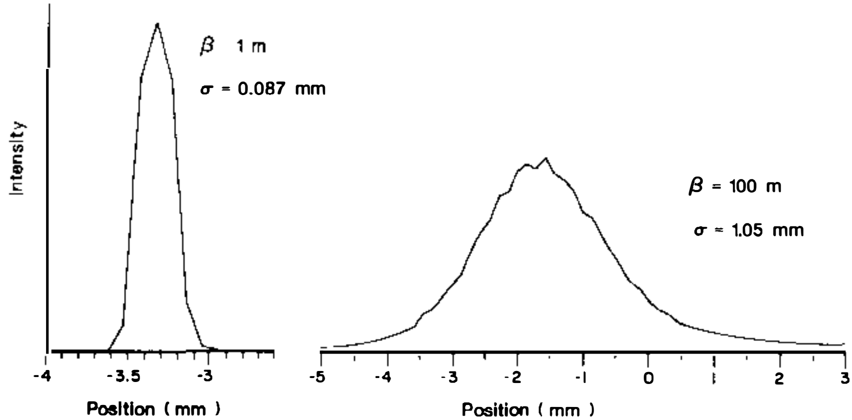

Figure 17 Beam profiles measured by the flying wire system at the collider low-β point and at a standard-β point.

excited from their center point in a $\lambda/4$ resonant mode. The small diameter unfolded structure was based on the geometrical requirement that the cavities fit between the Main Ring cavities and the floor. They can be mounted end to end with 180° electrical spacing if desired. The pulse amplifiers for the cavities are 200-kW units and produce ± 27 kV peak to peak and 10 amps rms. Fifty-ohm coaxial lines connect between the pulse amplifiers and the cavities.

The cavities must be frequency modulated by only 2 kHz during acceleration from 150 to 1000 GeV to compensate for the change in the protons' velocity. Radial position changes of 2 cm average can be made by frequency changes of 1 kHz. Because the required modulation is small, the normal frequency program is a completely dead-reckoned digitally generated modulation and runs off the control system Magnet Data Clock.

Though this system will run and accelerate beam without any feedback at low beam intensities, coherent phase oscillations are generally damped with a single-phase feedback loop. The phase detector itself must be capable of obtaining information from only one bucket (out of 1113 rf buckets) in the collider mode of operation.

The cavities have a 7-kHz bandwidth (-3 dB points) and are tuned to resonance (1 kHz per °C) by controlled changes in their operating temperature during the ramp cycle. In addition, it is necessary to compensate for the cavity heating that occurs during the long cycle time with the rf on for 40 out of 60 s. This is done by heating the water going to the cavities in proportion to the power not used when they are not at maximum power.

All eight cavities are used for acceleration of protons during fixed-target operation. However, the cavities have been positioned in pairs spaced by $3\lambda/4$ and phased in time by 90° so as to be in phase for beam going in one direction but 180° out of phase for beam traveling in the opposite direction. During collider operation, they operate as orthogonal sets of four for acceleration of protons and antiprotons independently. The colliding point of the two beams can be moved circumferentially around the accelerator and frozen at a particular point by independent adjustment of frequency of the two systems and by adjustment of the phase between the systems.

4. OPERATIONS

The Energy Doubler was commissioned between June and September of 1983 (19). Beam was first accelerated in July to 512 GeV. Fixed-target operation at 400 GeV began in October 1983 and the energy was increased to 800 GeV in February 1984. At this writing there have been three physics runs over a total of about 12 months duration, so one can begin to make some generalizations about the operational behavior of the accelerator. The three runs were at 400, 800, and 800 GeV with fast extraction. Intensity to date has not exceeded 1.1×10^{13} ppp. Single-beam storage studies have been carried out (45).

The rapid commissioning was successful probably because of the many step-by-step tests of components and subsystems that were made over the years prior to the final beam test. Though the major systems were not particularly reliable, they were all checked out and working. Interruptions to beam studies were at the level of trips of the ramp caused by refrigerator instabilities or spurious quenches caused by the quench protection system rather than any basic flaws.

As an accelerator the Energy Doubler was much more stable and reproducible than anyone had expected. (There had been concern that the coils would move during cooldown, quenches, or ramping.) For instance, steering adjustments once made did not have to be corrected from day to day. In addition, settings worked well over wide energy intervals and did not have to be changed appreciably as a function of excitation. The natural tune and chromaticity, though not exactly what was expected from the magnetic measurements, could be explained by differences in the average multipole coefficients a_1, b_1, b_2 of about one unit (1×10^{-4} at one inch). The momentum aperture was reassuringly large. Nevertheless continuous study and effort will be necessary to carry out all the detailed measurements required for a sufficient understanding of the accelerator (46) to permit both high-intensity, fixed-target operation and $\bar{p}p$ storage.

Operational reliability is a big issue in the Tevatron facility. There are now four accelerators in the chain, followed by the switchyard beam lines.

Both the Energy Doubler and the switchyard have superconducting magnets and the complications that go with them. It is obvious that the success of the Tevatron, as well as the credibility of larger, higher energy proton accelerators, lies in the ability to obtain high operational reliability. Improvements over the present reliability will have to be made in the Tevatron.

For initial operation with the Energy Doubler, actual uptime for high-energy physics has been at best of the order of 60% of the scheduled time. This amounts to about 300 hours per month as compared with 400 hours per month (or 80% actual to scheduled time) for Main Ring operation prior to Energy Doubler installation. At present about 50% of all downtime is due to the Energy Doubler systems. The major causes are the cryogenic system, the power supply and quench protection system, and quenches of the magnets themselves. These three categories account for 75% of the Energy Doubler downtime in a ratio of 2:1:1. As a reference, the Main Ring power supply system has about the same failure time as the Energy Doubler power supply and quench protection system and is the largest cause of downtime in the rest of the accelerator.

Cryogenic system downtime is of three types: trips due to unstable operation or oscillations in the system that indicate the magnets may be too warm; failures of expansion engines; and failure or reduced capacity of the central plant as a result of contamination, filter clogging, or inadequate supply of nitrogen or helium.

In the original design of the cryogenic system, the central helium liquefier played the role of a backup system that could be used to supplement the satellite refrigerators' capacity and provide continued operation if a refrigerator were down. In reality the central liquefier is required for ramped operation in addition to the above roles. This essential use of the central liquefier came about because the heat leak in the magnets is larger than originally expected. The increased heat leak resulted from a last-minute modification to the magnet-cryostat anchors, which prevent the collared coil assembly from rotating in the external iron.

Quenches of the superconducting magnets are obviously very disruptive to the accelerator operation. Not only does it take time to recover from a quench but the refrigeration and quench protection systems are exercised with each quench. During operation in the first quarter of 1985, there have been on the average of nine quenches per week. High-field quenches require about 70 minutes for recovery, injection-field quenches take about 20 minutes. The reasons for quenching can be put into three categories: typically 20% are directly due to failures of the quench protection system (like heaters firing without cause) or cooling problems in the cryogenic system; 25% are beam induced from failures of devices like kickers or correction coils; and the remaining 55% are beam related and have to do with tuning, unstable operation, or high intensity. The third kind of quench happens most often at high field, whereas beam quenches from device failures happen most often at injection. There are about equal numbers of high- and low-field quenches.

The expense of operating the superconducting accelerator can be assessed from the power and cryogens used. Figure 18 illustrates the power savings of the superconducting ring. The shaded area illustrates the power required to operate the main accelerator prior to the Energy Doubler operation versus that necessary to operate the Main Ring–Doubler and Central Helium Liquefier after Doubler turn on. There is of the order of a factor of two savings or approximately 15 GW hours per month (10 million dollars per year) for continuous operation. Cryogen usage (Figure 19) runs about 5.4 million dollars; thus if the program is operating 11 months of the year and the refrigeration operation is not cut back during the one month off, a savings of about $3\frac{1}{3}$ million dollars could be realized. Of course, this saving is offset by the additional operating cost of the cryogenic systems from manpower and materials that would not be needed in a conventional system. The main advantages are that power usage per se has been substantially reduced and that higher energy beam and collider physics is possible.

Cryogen usage is substantially greater than anticipated. Nitrogen usage is more than 50% higher than expected. This excess is due in part to the higher heat leak of the magnets and in part to a preference for operation in the most stable mode rather than the most economical mode. The large helium leakage rate is something of a mystery and probably is distributed over many parts of the plant. For instance, leakage of valve stems is a suspected cause and checking all possibilities is an arduous task. Work must be done to reduce both nitrogen usage and helium leakage as the operation stabilizes.

5. HINDSIGHT AND A LOOK TO THE FUTURE

It works! This is by far the most significant thing that can be said in retrospect. One tends to forget that it was not obvious that the Energy Doubler would work and that years of effort with many failures and setbacks went into the magnet and cryogenic development.

There are no major flaws in the system that we know of to date, but one should keep in mind that it is very much a prototype accelerator. It is as much an accelerator research tool as a high-energy physics tool. It will take time for the performance to become as dependable as expected from conventional accelerators.

So, first of all it works. Now if we were starting over today what would we have done differently? The most natural question to ask is whether the

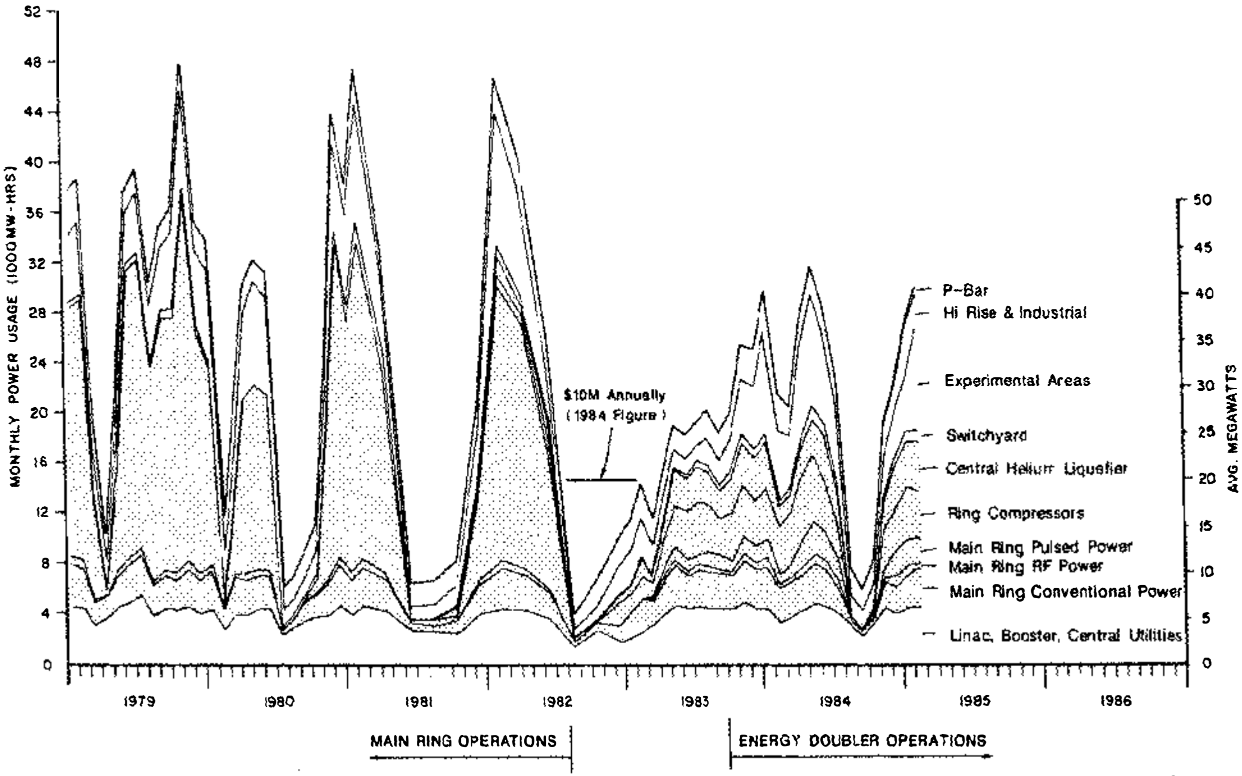

Figure 18 Power usage before and after Energy Doubler installation. Shaded area indicates Main Ring usage prior to Energy Doubler installation versus Main Ring and Energy Doubler usage after installation.

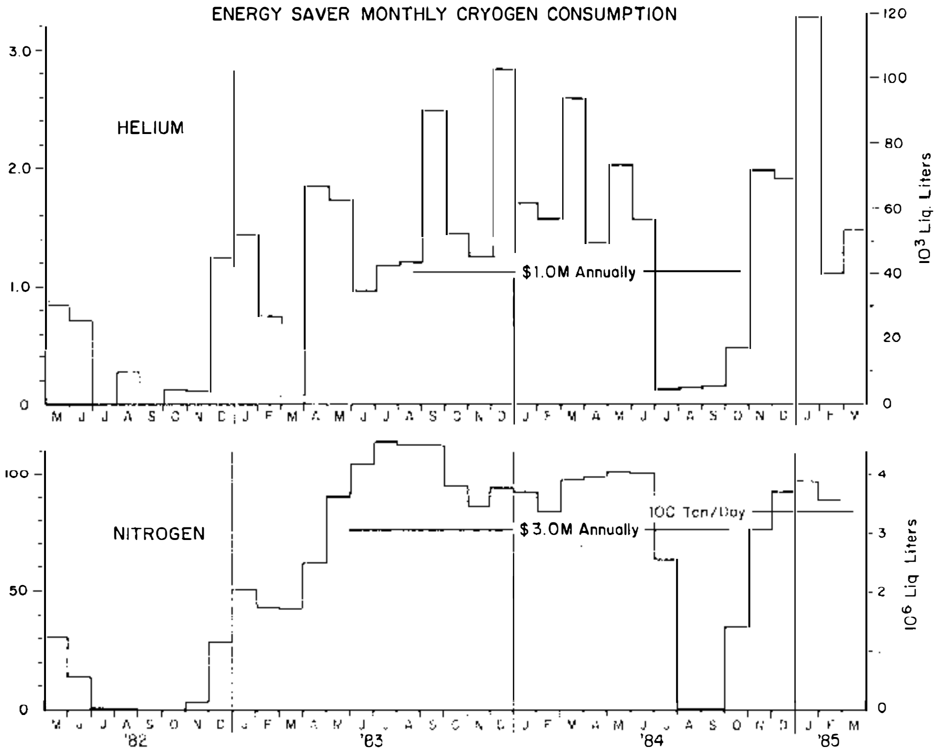

Figure 19 Energy Doubler helium and nitrogen usage.

Energy Doubler should have been built in the same tunnel as the Main Ring. Even at this point, one can only speculate as to the outcome if the same tunnel had not been used. It is quite possible that the whole project would not have happened at all. After all, with a risky research and development effort, one does not necessarily want to spend a lot on conventional construction in order to test the new idea. It is expedient having all the conventional utilities and buildings available and in good working order. By installing the Energy Doubler in the Main Ring tunnel and turning off all physics research during that time interval, one guaranteed total laboratory commitment to the project. That commitment would not have been available were the physics programs running simultaneously. On the other hand, the three long shutdowns for Doubler installation, $\bar{p}$ injection extraction tunnel modifications, and colliding detector building construction have been very disruptive to physics research. Of the order of $2\frac{3}{4}$ years of operation will have been lost during the six years (1981–1986) of construction, installation, and commissioning of the Tevatron. But physics at higher energy has been made possible sooner.

Many design choices for the Energy Doubler magnets were precluded by the requirement that they fit below the Main Ring. This boundary condition may actually have helped focus the effort. Possibly the esthetically pleasing warm iron design with short suspension blocks and high heat leak might have been modified had there been more space. The compact cryostat cross section leaves little room for the cryogen flow and thus limits the length of the magnet strings serviced by one refrigerator.

Independent of the magnet geometry, better heat transfer from the coils to the two-phase system would have helped obtain higher energy and reduce ramp turn-on transients (resulting from expansion of the helium liquid with heating). Increasing the capacity of the satellite refrigerators to match better the magnet heat load or obtaining redundancy in the Central Helium Liquefier is the single, most important change that could presently be made to the Energy Doubler in order to improve its reliability.

The overriding issue for the next generation of large superconducting accelerators-colliders must be reliability. Some real assessment of the probable success in meeting operating goals must be made, and large prototype systems tested for fatal flaws. The reliability issue must be treated with a care it never has received in past accelerators.

A second difficult issue is that of magnet aperture. By this one does not mean the physical aperture of the beam tube but rather the usable dynamic aperture within which the beam will survive. There are at least two levels here: over what region will the beam behave reasonably linearly so one can analyze it and so that it will store with very long lifetimes; and over what region will it survive long enough so that it can be crudely measured and adjusted? As indicated in the introduction, evaluation of the problem is very controversial. Computer simulations are not necessarily sufficient to answer these questions and must be checked against measured behavior in existing accelerators such as the $Sp\bar{p}S$ and the Energy Doubler.

The field quality in the magnets is affected by the choice of coil radius in two ways. First, as the radius is reduced conductor placement errors become more important and the random variations of the multipole coefficients of the field will be increased. It is believed that these variations in the b_n, a_n coefficients scale like $1/r_c^{n+\frac{1}{2}}$ all else being equal, where r_c is the coil radius. This possibly simplistic approximation comes from assuming dominance of azimuthal errors that scale like the square root of the number of turns in the coil and consequently like the $1/r_c^{1/2}$ for constant cable thickness (47). The second effect is produced by "persistent currents"; it changes the average multipole coefficients at low excitation, as well as producing some random behavior. The average multipole coefficients of the dc residual field scale like D/r_c^{n+1}, where D is the superconductor filament diameter. [They also scale like $J_c(B)$ and $\Delta r_c \varepsilon$, where Δr_c is the coil thickness and ε is the fractional area of the coil that is composed of superconductor (24).] The effect is dominant in the $b_0, b_2, b_4, \ldots$ coefficients; however, random behavior can appear in both normal and skew coefficients from variations in filament diameter, critical current $J_c(B)$, and temperature of the coils. Because persistent currents produce a field that is only slowly dependent on magnet excitation (Figure 7), the effect on the beam at injection varies inversely with the injection field.

Persistent currents lead to large chromaticity, ξ, in the accelerator if uncompensated [$\Delta\nu(\Delta p/p) = \xi\Delta p/p$, $\xi \simeq b_2\langle\eta\beta\rangle$]. Typically, η is unchanged as the radius of the accelerator is increased and β increases proportionally to $\sqrt{R}$, where R is the accelerator radius. Uncorrected chromaticities at injection 50 to 100 times larger than in the Energy Doubler can result in 20-TeV accelerators at 1 TeV injection, with coil radii reduced from 3.8 cm (the Energy Doubler coil radius) to 2.5 or 2.0 cm. Fortunately it appears that filament diameters of one half to one quarter that of the conductor used in the Energy Doubler (8 μm) may be possible. In any case, sextupole corrections will probably have to be made locally in each bend magnet (otherwise second-order terms may become important and misalignment between the correction and the bends will produce random quadrupole errors). Two such correction schemes are passive and require no external control. One relies upon a shorted superconducting sextupole coil driven by the unwanted sextupole flux (48). Another method relies on strips of superconductor mounted on the outside of the beam pipe (49). This scheme uses the magnetization induced by persistent currents in the passive strips to cancel the sextupole fields and does not involve any external (transport) current. Higher order, e.g. decapole, corrections may be required also.

The fundamental risk from the random multipole errors is the deterioration from rational linear behavior of the accelerator beam (50). This deterioration is roughly proportional to $\sqrt{L}x^{n-1}\langle a_n^2, b_n^2\rangle^{1/2}$ in the approximation that the amplitude function β is proportional to $R^{1/2}$. Here L is the length of a bend magnet (or the coherent length over which multipole coefficients remain constant if shorter than a magnet length); x is the aperture dimension of interest, either horizontal or vertical; and $\langle a_n^2, b_n^2\rangle^{1/2}$ are the rms fluctuations of the multipole coefficients. The above expression can be written $\sqrt{L}r_c^{-3/2}(x/r_c)^{n-1}$ using the expected scaling of the multipoles. Thus, for sextupole-dominated effects the aperture is expected to be about five times smaller for a 2-cm coil radius than for the 3.8-cm Energy Doubler coil radius.

Recently there have been pleasant surprises in superconductor development. Not only does it appear possible to produce conductor with 4-μm filaments as mentioned, but the current-carrying capacity of the conductor has already improved markedly since production of Energy Doubler magnets. At present (1985) a J_c of 2150 amp mm^{-2} at 5 T and 4.2 K is typical in short sample tests of Energy Doubler–type cable, an improvement of 20% since Doubler production. Moreover, there are preliminary indications that both small filaments and high current density can be achieved simultaneously, and strand with 2900 amp mm^{-2} has been produced.

Cryogenic requirements and system design will probably be quite different in a Superconducting Super Collider (SSC) from those for the Energy Doubler (51, 52). A large fraction of the magnet heat leak comes from synchrotron radiation heating by the beam and some comes from static or ramped heating loads. These must be substantially reduced. Static loads can be reduced through longer magnet supports; excitation ramping will be slow and only occasional. Total system cryogenic capacity may need to be only three times that required for the Energy Doubler. Heat leak estimates, possibly optimistic, are as little as 1–2 W at 5 K and 2 W at 10 K per (12-meter) dipole, as compared with 9 W per 6-meter Energy Doubler dipole.

The refrigerator system will probably consist of stand-alone refrigerators located around the ring at each of the dozen or so sectors. Each refrigerator might have a number of compressors, so there are backup spares. In addition, each refrigerator could be sized for 150% of nominal load so that if any one goes down its cooling load can be handled by its two neighbors.

It is hoped that transfer lines and most headers can be eliminated. By having more space for the magnet cryostat, the magnets themselves will serve as their own transfer lines. By designing the magnet cryostats to withstand high pressures, venting of helium during quenches and power failures can be substantially reduced.

The magnet quench protection system can potentially be "passive" for a slowly ramping collider in which radiation damage is not likely to be a problem for semiconductors (53). In such a system diodes are put across the coils of the individual magnets (or the coil is split in fractions and additional diodes are used). These conduct when a quench occurs. $L\,di/dt$ voltage across the magnet must be limited to less than the forward voltage drop across the diode, and the quench propagation velocity must be sufficient to allow the magnet to absorb its own energy without being destructively overheated. This scheme favors magnets with low stored energy or small coil radius. The quenches must still be detected in order to dump the energy from the rest of the magnet string, but heaters and very fast detection response have been eliminated. (There is a spectrum of possible solutions to the quench protection that range between this one and that used in the Energy Doubler.) The passive diode protection scheme does require a large number of diodes mounted in the magnet cryostats. Of the order of 25,000 would be required in a 20-TeV collider. Even if the probable number of times each one is used in quenching is less than one, there still are serious reliability questions with this scheme. (Magnets would have to be warmed up to change shorted diodes.) The HERA accelerator will probably try this method first, and its experience will be extremely valuable.

ACKNOWLEDGMENTS

The successful completion and operation of the first large superconducting accelerator are a result of the dedicated work and perseverance of the people at Fermilab. Special acknowledgment should be given to the following individuals for supplying material used in this article: D. Bogert, D. Edwards, E. Fisk, W. Fowler, M. Harrison, H. Jostlein, K. Koepke, P. Martin, C. Rode, and G. Tool.

Literature Cited

1. Bjorken, J. D. "Fixed-target physics at Fermilab," p. 144; and Lederman, L. M. "Physics at the Tevatron," p. 33; both in *Am. Phys. Soc., Proc. Santa Fe Meet., Oct. 31–Nov. 3, 1984*. Philadelphia/Singapore: World Sci. (1985)
2. Sanford, J. R. *Ann. Rev. Nucl. Sci.* 26: 151–98 (1976)
3. Design Report Superconducting Accelerator. *Fermilab Rep.* (May 1979)
4. Design Report Tevatron I Project. *Fermilab Rep.* (Sept. 1984)
5. Peoples, J. *IEEE Trans. Nucl. Sci.* NS-30(4): 1970 (1983)
6. Wilson, R. R. *Phys. Today*, Oct. 23 pp. (1977)
7. Vander Arend, P. C., Fowler, W. B. *IEEE Trans. Nucl. Sci.* NS-20(3): 119 (1973)
8. Palmer, R., Tollestrup, A. V. *Ann. Rev. Nucl. Part. Sci.* 34: 247 (1984)
9. Wilson, M. N. *Superconducting Magnets*, pp. 131 ff. Oxford: Clarendon (1983)
10. Turkot, F., Cooper, W. E., Hanft, R., McInturff, A. *IEEE Trans. Nucl. Sci.* NS-30: 3387 (1983)
11. Hanft, R., Brown, B. C., Cooper, W. E., Gross, D. A., Michelotti, L. et al. *IEEE Trans. Nucl. Sci.* NS-30: 3381 (1983)
12. Cooper, W. E., Fisk, H. E., Gross, D. A., Lundy, R. A., Schmidt, E. E., et al. *IEEE Trans. Magn.* MAG-19: 1372 (1983)
13. McInturff, A. D., Gross, D. *IEEE Trans. Nucl. Sci.* NS-28: 3211 (1981); Wake, M., Gross, D., Yamada, R., Blatchley, D. *IEEE Trans. Magn.* MAG-15: 141 (1979); McInturff, A. D., Carson, J., Engler, N., Fisk, H., Hanft, R., et al. *IEEE Trans. Magn.* MAG-21: 478 (1985)
14. Brown, B. C., Fisk, H. E., Hanft, R. *IEEE Trans. Magn.* MAG-21: 979 (1985)
15. Edwards, D. A. In *Proc. Int. Conf. High Energy Accelerators, 12th, Fermilab*, ed. F. T. Cole, R. Donaldson. Batavia: Fermilab. 90 pp. (1983). This reference gives a summary of dipole multipole data
16. Michelotti, L. P., Ohnuma, S. *IEEE Trans. Nucl. Sci.* NS-30: 2472 (1983)
17. Gelfand, N. M. In Accelerator Physics Issues for a Superconducting Super Collider, 12–17 Dec., ed. M. Tigner. Ann Arbor: Univ. Mich. 124 pp. (1983)
18. Collins, T. L., Edwards, D. A. *Fermilab Tech. Memo TM-614* (1975)
19. Edwards, H. T. See Ref. 15, p. 1
20. Rode, C. H. See Ref. 15, p. 529
21. Rode, C. H., Theilacker, J. C. "What % of Carnot is Realistic?" *Fermilab Intern. Memo* (1983)
22. McInturff, A. D. Private communication
23. Martin, P., et al. *Fermilab Rep. TM-1134*. 57 pp. (1982)
24. Kuchnir, M. "Second SPTF Test." *Fermilab Intern. Memo* (1981)
25. Bartelson, C. L., et al. *J. Vac. Sci. Technol. A* 1(2): 187 (1983)
26. Tool, G., Flora, R., Martin, P., Wolff, D.

IEEE Trans. Nucl. Sci. NS-30:2889 (1983)
27. Koepke, K., Martin, P. *Fermilab Intern. Memo UPC-155* (1982)
28. This is an experimental observation; private communication, K. Koepke
29. Bogert, D., Segler, S. In *Europhysics Conf.—Computing in Accelerator Design and Operations, Berlin*, ed. W. Burse, R. Zelazny. Berlin: Springer-Verlag. 338 pp. (1983)
30. Bonner, M. K. In *Research and Development*, ed. R. R. Jones. Barrington, Ill. 106 pp. (1984)
31. Cox, B., Mazur, P. O., VanGinneken, A. *Fermilab Tech. Memo TM-828-A* (1978)
32. VanGinneken, A. *Fermilab Rep. FN-272* (1975)
33. VanGinneken, A., et al. *Fermilab Publ. 85/64* (1985)
34. Drozhdin, A. I., Harrison, M., Mokhov, N. V. *Fermilab Rep. FN-418* (1985)
35. Johnson, D. E. *Fermilab Tech. Memo TM-1106* (1982); Johnson, D. E., Koepke, K., Willeke, F., et al. In *Part. Accel. Conf., Vancouver, IEEE Trans. Nucl. Sci.* In press (1985)
36. Gerig, R., et al. See Ref. 35
37. Johnson, M., McInturff, A., Raja, R., Mantsch, P. M. See Ref. 15, p. 536
38. Pfeffer, H., Yarema, R. J. *IEEE Trans. Nucl. Sci.* NS-30:2870 (1983)
39. Rotolo, C. J. *IEEE Trans. Nucl. Sci.* NS-30:2959 (1983)
40. Goodwin, R. W., Johnson, R. P. See Ref. 15, p. 594
41. Shafer, R. E., Gerig, R., Baumbaugh, A., Wegner, C. See Ref. 15, p. 609
42. Crisp, J., et al. See Ref. 35
43. Elseth, J., et al. See Ref. 35
44. Kerns, Q., Meisner, K., et al. See Ref. 35
45. Johnson, R. See Ref. 17, p. 45
46. Edwards, D. A., Johnson, R. P., Willeke, F. *Fermilab Publ. 85/59. Part. Accel.* Submitted (1985)
47. Edwards, D. A. Private communication
48. Gilbert, W., Borden, A., Hassenzahl, W., Moritz, G., Taylor, C. *IEEE Trans. Magn.* MAG-21:486 (1985)
49. Brown, B. C., Fisk, H. E. In *Proc. Summer Study on the Design and Utilization of the Superconducting Super Collider*, ed. R. Donaldson, J. Monfin. New York: Am. Phys. Soc. (1984)
50. Collins, T. *Fermilab Publ. 84/114*. 18 pp. (1984)
51. The Superconducting Super Collider. A Reference Design Study for the US Dept. Energy. Berkeley, Calif. (1984)
52. Rode, C. H. In *10th Int. Cryogenic Eng. Conf., Helsinki*, ed. H. Collan, P. Berglund, M. Knusius. Guildford, Surrey: Butterworth (1984)
53. Koepke, K., Tool, G. Workshop on SSC Commissioning and Operation. Unpublished (1985)