Fermi National Accelerator Laboratory

FERMILAB-FN-650

Design Concepts for a Novel Fermilab Superbooster

Charles M. Ankenbrandt

Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510

November 1996

Operated by Universities Research Association Inc. under Contract No. DE-AC02-76CH03000 with the United States Department of Energy

Disclaimer

Distribution

DESIGN CONCEPTS

for a

NOVEL FERMILAB SUPERBOOSTER

Charles M. Ankenbrandt
Fermi National Accelerator Laboratory*
P.O. Box 500, Batavia, IL 60510

October, 1996

Abstract

Design concepts for a novel high-intensity, high-brightness proton accelerator complex are described. In particular, a two-ring replacement for the Fermilab Booster that would outperform that machine by about an order of magnitude in beam intensity is proposed. Charge stripping of H^- ions is used to transfer protons from the first synchrotron to the second. The transfer is initiated by a laser beam that strips the first electron from H^- ions after they have been accelerated in the first ring. The resulting beam of neutral hydrogen atoms streams across to the second ring, where the second electron is removed by a stripping foil as in a conventional charge-stripping multi-turn injection scheme. By modulating and spatially shaping the laser beam, selectivity in both transverse and longitudinal phase spaces can be achieved, affording considerable flexibility in tailoring phase-space distributions to the diverse needs of various downstream machines or users. The enhancement of future physics programs at Fermilab by application of these capabilities is discussed.

TABLE OF CONTENTS

INTRODUCTION

Recently an idea for a new type of proton accelerator complex was conceived by the author. The novel feature is the use of charge stripping of H^- ions to transfer protons from one circular accelerator to another. The transfer is initiated by a laser beam that strips the first electron from circulating H^- ions that have been accelerated in the first ring. The resulting beam of neutral hydrogen atoms streams across to the second ring, where the second electron is removed by a stripping foil. The multi-turn charge-stripping injection into the second ring then proceeds fairly conventionally, followed by further acceleration in that ring.

The primary advantage of this scheme is that the laser beam can be used to select portions of the phase space of the H^- beam in order to populate desired regions of the phase space in the second ring. Since it is a multi-turn process, ions that do not satisfy the selection criteria on one turn may fall into the acceptable phase space on subsequent passes. By modulation and spatial shaping of the laser beam, selectivity in both longitudinal and transverse phase spaces can be achieved, affording considerable flexibility in tailoring phase-space distributions to the needs of downstream machines or experimenters.

Note that, unlike conventional charge-stripping injection from a linac, the method proposed here makes use of the two electrons on each ion for two separate purposes: removing the first one selects beam to be transferred in the first ring, and removing the second one allows overlaying beam into the same phase space in the second ring.

The utility of this idea depends not only on whether the laser transfer process itself is feasible, but also on whether two accelerators that use the process can outperform conventional facilities at reasonable cost. This paper elaborates upon the laser transfer idea by describing design concepts for a complete functional replacement for the Fermilab Booster that would outperform that machine by about an order of magnitude in beam intensity while also providing considerably more flexibility. In discussions within the Accelerator Division, the proposed complex has come to be called the SuperBooster.

The SuperBooster proposal is somewhat similar in spirit to two other papers that appeared many years ago. The first paper proposed a Fermilab Booster upgrade called the B^4 Booster.[1] The second paper presented more general ideas about boosters for a hadron-hadron collider.[2] Both articles proposed accelerating H^- ions in the first ring. However, those papers did not include the idea of using a laser for the beam transfer. In the first paper, extraction of H^- from the first ring was to be by resonant extraction. The second paper proposed simultaneous displacement of the beam in both rings onto a stripping foil in a shared straight section.

The status of development of the SuperBooster design is reported. No insoluble problems have been identified, and plausible approaches to design issues have been found, but some topics require more detailed design work to prove feasibility. Routine aspects of accelerator design have not been addressed.

OVERVIEW

A brief overview of the proposed complex will facilitate understanding the rest of this document. The complex consists of two rapid-cycling synchrotrons, both about the size of the Fermilab Booster, operating in series in the same tunnel. The first ring accelerates H^- ions from the

Table 1. Major SuperBooster parameters.

Major parameters	**First Ring**	**2nd Ring**	
Bending radius	53.46	52.83	m
Dipole packing factor	0.70	0.70	
Average radius	76.37	75.47	m
Circumference	479.85	474.20	m
Ramp rep. rate	15	15	Hz
Transition gamma	5	12	
Harmonic number	85	84	
Max E gain/turn	138	455	keV
Injection Parameters			
Kinetic energy	0.40	2.00	GeV
Beta	0.71	0.95	
Mag. field	0.60	1.76	kg
Rev. freq.	0.45	0.60	MHz
rf freq.	37.97	50.36	MHz
Extraction Parameters			
Kinetic energy	2	8	GeV
Beta	0.95	0.99	
Mag. field	1.74	5.61	kg
Rev. freq.	0.592	0.629	MHz
rf freq.	50.360	52.813	MHz
Beam Parameters			
Protons per bunch	5.88×10^{11}	5.95×10^{11}	
Total no. protons	5×10^{13}	5×10^{13}	
Norm. 95% emittance	$200\ \pi$	$40\ \pi$	mm-mr
Laslett tune shift	-0.40	-0.42	

present Linac energy of 400 MeV to 2 GeV; the second one accelerates protons from 2 GeV to 8 GeV to match the output energy of the Fermilab Booster. Table 1 shows major parameters of the two rings. The parameter table is presented here to illustrate current thinking; the numbers do not represent final design choices.

In order to avoid stripping the H^- ions, the first ring must have good vacuum (about 10^{-9} torr or less) and low magnetic field (B<1.7 kg). It must also have large acceptance to accommodate beams having normalized 95% transverse emittances of about 200 π mm-mr. The large beam size serves two essential purposes: it allows injection of many turns by stacking in phase space, and it raises the Laslett incoherent space-charge tune shift limit to the desired value of about 5×10^{13} protons per cycle.

The rf bucket spacing in both rings (5.645 m) matches that of the Fermilab Booster, Main Ring, Main Injector, and Tevatron; the rf frequency at 8 GeV is 52.813 MHz. The two rings have incommensurate harmonic numbers (85 and 84), and correspondingly different circumferences, so that the beam in the first ring cogs rapidly with respect to the second, thereby allowing different bunches in the first ring to be overlaid into the same rf bucket in the second ring. The number of rf cavities required is modest in both rings because of the small energy gain in the first ring and the small frequency swing in the second. This is a considerable advantage not only because rf systems are expensive but also because the cavities present high impedances to the beam, resulting in instabilities at high intensity.

The injection methods into the two rings have already been mentioned; establishing their feasibility is really the crux of this proposal. The phase space stacking into the first ring and the charge-stripping transfer into the second ring will be discussed at length in two following sections. Other sections will examine major subsystems of the two rings. Finally the question of whether the existing Booster has a role to play in the context of considerably higher performance will be addressed. But first the desired performance of the SuperBooster and the factors limiting its performance will be discussed.

SUPERBOOSTER PERFORMANCE

Desiderata

Various ideas for future physics programs at Fermilab are now under active consideration. Naturally, the plans for the near future rely on small and reasonable extrapolations of the present performance of the Fermilab Booster. However, some of the ideas for the physics of the next millennium require much better performance of the proton source, and the discussion of the present Booster will make it obvious that large performance improvements are not to be expected from it. It is not too early to start thinking about a Booster replacement, because it inevitably takes a long time from conception to commissioning of such a complex. Before choosing performance specifications for a Booster replacement, the needs of the different programs ought to be reviewed.

Main Injector Performance

The Main Injector design calls for the Booster to deliver $5x10^{12}$ protons per batch. That is only about 20% above the intensity that the Booster can deliver to the Main Ring now, and since the Main Injector has a much larger aperture than the Main Ring, the Booster should be able to deliver the design intensity. A standard antiproton production cycle uses one batch; a standard Main Injector fixed-target cycle uses six batches totaling $3x10^{13}$ protons.

High beam intensity is obviously desirable both for antiproton production and for the Main Injector fixed-target program. In discussions about SuperBooster performance specifications in the context of the Main Injector, Division Head Finley, alluding to conversations with Director Peoples, suggested an intensity goal of $5x10^{13}$ protons per SuperBooster cycle.[3]

Weiren Chou recently presented an analysis of Main Injector intensity limitations at a TeV33 meeting. To make a long story short, he felt that a factor of five above the design intensity might be achieved with suitable modifications, but that a factor of ten seemed very difficult.[4]

TeV33

After extensive discussions within the Accelerator Division in connection with TeV33, a plausible path to a luminosity of 10^{33} $cm^{-2}sec^{-1}$ has begun to emerge. The plans are described in a paper resulting from Snowmass 1996.[5] It appears that the proton beams needed to achieve the luminosity goal might be supplied by a more economical means than replacing the Booster; instead, the plans creatively circumvent the limitations of the Booster by taking advantage of the capabilities of downstream machines. In particular, slip-stacking[6,7,8] at injection into the Main Injector is planned to increase the number of protons on target for pbar production, and electron cooling of both protons and antiprotons in the Recycler can be used to create the bright and intense bunches needed in the Tevatron Collider.

Nevertheless, it is instructive to point out how the capabilities of a machine like the proposed SuperBooster could be used to enhance the performance of the Tevatron Collider. First, the ability to overlay bunches could be used to compress a SuperBooster batch into 42 or even 28 bunches, and then two or three of these shorter batches could be loaded end-to-end into the Main Injector to make a very intense 84-bucket batch for pbar production.

Second, the process of merging multiple proton bunches into one for the collider could be implemented by laser stripping at 2 GeV in the SuperBooster rather than by rf coalescing at 150 GeV in the Main Injector. Furthermore, overlaying several bunches with the aid of the laser could create an intense bunch having the same longitudinal emittance as the original bunches. However, the antiproton bunches would have to be prepared some other way. Upgraded stochastic cooling might be used to reduce the longitudinal emittance of the antiproton bunches. (If a proton-proton collider supersedes the present antiproton-proton collider in the future, then presumably both beams would be processed the same way. Smaller longitudinal emittances would then enable shorter bunch lengths, allowing use of smaller beta functions at the interaction point.)

Muon Collider

The demanding specifications for the so-called proton driver for a Muon Collider have been evolving for the past year or so. Present plans call for a total of 10^{14} 8-GeV protons per cycle at a repetition rate of 15 Hz.[9] The bunch structure at the pion-production target is to consist of four bunches with rms lengths of 3 nsec apiece. (There is an alternative specification that uses 30 GeV protons instead.)

For the muon collider, the initial complement of 85 bunches in the first SuperBooster ring might be transformed into four clumps of beam (this is deliberately vague because there are various possibilities), which can subsequently be manipulated at 8 GeV to the four tightly bunched distributions desired at the pion production target.

For the Muon Collider, the macroscopic duty factor can approach 100%; that is, every cycle would be a beam-on cycle. In contrast, when a proton booster feeds a larger proton synchrotron such as the Main Injector in series, the duty cycle of the booster is typically of order 10% because the beam in the booster is off while the protons are accelerated in the latter machine. The combination of very high beam intensity and 100% duty factor implies that even less beam loss can be tolerated for the stand-alone applications than for the booster applications.

SuperBooster Fixed-Target Experiments

An example of an experiment that could use the SuperBooster output directly is an initiative called BOONE. The LSND (Liquid Scintillator Neutrino Detector) collaboration, which has published evidence for oscillations of muon anti-neutrinos to electron anti-neutrinos based on data from Los Alamos[10], is working on the design of a similar experiment at higher energy at the Fermilab Booster that would be capable of collecting almost two orders of magnitude more events.[11] The advantage of using relatively low proton energies for neutrino production is that fewer kaons, and hence relatively few electron anti-neutrinos, are produced. For this or any similar beam-dump experiment that is not rate-limited, the SuperBooster would produce more events than the Booster in proportion to their respective intensities. Like the muon collider, an experiment such as this could use a duty factor approaching 100%.

Really Large Hadron Collider

The Really Large (about 50 TeV x 50 TeV) Hadron Collider, sometimes informally called the Pipetron, is a proposed proton-proton collider at Fermilab.[12] While a variety of beam parameters for such a machine have been suggested, it is reasonable to expect that the proton beam requirements will evolve toward those of the SSC, with appropriate changes to account for a larger circumference and higher design luminosity.

In order to limit the number of events per crossing, the SSC design called for a very large number of bunches, with bunch spacing similar to the normal 53 MHz bunch spacing at Fermilab. To achieve high luminosity while limiting the beam power and synchrotron radiation losses, the design specified a high beam brightness N_p/ε_n. The maximum usable value of brightness is limited by beam-beam effects in the collider.

The beam-beam limited value of brightness is somewhat higher than the maximum value of brightness for individual bunches that the present Fermilab injector chain can produce directly. This should not be surprising. In any straightforward collider filling scenario where the individual bunches are formed in the first booster and preserve their identity throughout the injector chain to the collider, there is a relationship between the space-charge tune shift at low energy in the first booster and the beam-beam tune shift in the collider.[13] Both tune shifts are proportional to the brightness. The beam-beam tune shift is independent of energy, whereas the space-charge tune shift has the kinematic dependence $1/\beta\gamma^2$. The upshot is that the booster whose space-charge tune shift determines the limiting beam brightness (ordinarily the first booster in the chain) must have high injection energy and small circumference in order to create bright enough bunches to meet the collider requirement.

In the plans for TeV33, the brightness limit of the Booster is circumvented by using electron cooling in the Recycler to create beams of the desired brightness. The relatively slow electron cooling times[14], measured in minutes, are tolerable in the case of TeV33 because not very many Main Injector cycles are needed to fill the Tevatron Collider. Note, however, that the Pipetron circumference of 500 km is about 150 times larger than the Main Injector and Recycler circumferences of 3.3 km. Thus it would take about 300 Main Injector cycles to fill both Pipetron rings using straightforward "boxcar" filling. Slow electron cooling of each Main Injector batch in the Recycler would then lead to excessive Pipetron fill times.

In the SuperBooster, beams of the desired brightness would be created by shrinking the transverse emittances during the beam transfer from the first to the second ring. The 2-GeV injection energy into the second ring is more than high enough to make the space-charge tune shift acceptable for beams of the desired brightness. In this context the SuperBooster is comparable in performance to a 2-GeV Linac upgrade.

Early versions of the SSC design called for variable bunch spacing in the collider. This flexibility was eventually given up in the SSC design, apparently because it was deemed too complicated to implement conventionally. However, the laser transfer process would allow virtually any set of buckets to be populated for the Pipetron.

It is worth noting in this context that the ability to vary the number of bunches would also be useful for a muon collider. At the design luminosity, the present plans call for two bunches of positive muons and two of negative muons, rather than one of each polarity, in order to avoid excessive rf beam loading in the muon accelerators. In the early days, before the design luminosity is routinely achievable, a single more intense bunch of each polarity could be used instead, thereby increasing the luminosity. Conversely, if performance improves to the point that the design muon bunch intensities are routinely exceeded, it would be easy to increase the number of proton bunches on target if necessary.

Present Linac Performance and Upgrade Potential

In a recent paper edited by M. Popovic, the Fermilab Linac staff addressed present and future performance of the Linac.[15] The present limitations on beam current and pulse length were analyzed for the linac subsystems, and straightforward upgrade paths were identified. Their results are as follows. As of August, 1995, the Linac was able to deliver a beam current of 46 mA and a pulse length of 35 μsec for the high-energy physics program. This corresponds to 1.0×10^{13} protons per cycle. They assert that "with existing hardware" the Linac can deliver a 48 mA, 80 μsec pulse, corresponding to 2.4×10^{13} protons. After discussing upgrade possibilities, they conclude that "with small modifications" a 55 mA, 110 μsec pulse could be delivered, corresponding to 3.8×10^{13} protons. A pulse length of 110 μsec corresponds to about 50 Booster turns at 400 MeV.

Choice of Design Intensity

A review of the requirements of the various proposed programs is in order as a prelude to choosing the design intensity of the SuperBooster.

An intensity of 5×10^{13} from the SuperBooster would imply an order-of-magnitude improvement in Main Injector capabilities. However, a theoretical analysis of the intensity limits of the Main Injector suggests that it might be futile to inject more than about 2.5×10^{13} per 84 bunches. Either of these numbers would be more than adequate to support the future needs of the Tevatron Collider.

Really large hadron colliders typically require a very large number of very bright bunches delivered rapidly. However, the intensity required of the injectors is not as large as that needed by some of the other proposed programs. The ability of the SuperBooster to shrink transverse emittances is quite important for this application.

The muon collider is the most demanding of the high-energy facilities; not only does it need the highest intensity per batch, 10^{14}, but it also requires a special bunch structure: four short bunches. The ability of the

SuperBooster to combine bunches without emittance growth is very valuable in this regard; but still the specifications are challenging. A complex designed to deliver 5×10^{13} might gain a factor of two to reach the muon collider design luminosity in either of two ways: by producing both polarities of muons with each proton bunch, or by a Linac upgrade to about 1 GeV to double the SuperBooster space-charge limit.

A class of beam dump experiments that directly uses the SuperBooster output would be able to take advantage of the highest intensities that it could deliver.

Pending the examination of various SuperBooster intensity limitations, an intensity of 5×10^{13} is tentatively adopted as a design goal. That intensity is a reasonable compromise among the demanding requirements of the various future physics programs, and the ability to deliver that many protons is probably within the range of an aggressive Linac upgrade. We shall see that normalized transverse emittances of about 200 π mm-mr are required in the first ring to make the space-charge tune shift barely tolerable at that intensity; such emittances can be reduced to 40 π mm-mr during the laser transfer to the second ring in order to fit into the Main Injector acceptance. The higher injection energy of the second ring makes the tune shift tolerable there when the emittances are reduced by a factor of five.

Factors Limiting Performance

Very high-intensity, medium-energy synchrotrons have been frequently proposed but seldom built. Demanding performance specifications like those of the SuperBooster are typical of the design of proposed new spallation neutron sources as well as advanced hadron facilities such as KAON at TRIUMF and LAMPF II at Los Alamos.[16] Achieving such exacting specifications is a challenging task. The large number of design reports that have been written and workshops that have been held typically focus on the following impediments to high intensity.

Laslett tune shift

The Laslett incoherent space-charge tune shift[17] limits the beam brightness at low energy. The limitation on beam intensity can be raised by increasing the injection energy and by making the transverse emittances larger. Of course physical and dynamic apertures must be large enough to accommodate the large emittances.

A useful approximation for the space-charge tune shift $\Delta\nu_{sc}$ at the center of a round Gaussian beam is

$$\Delta\nu_{sc} = -\frac{3 r_p N_{tot}}{2\varepsilon_n \beta \gamma^2 B}$$

In this expression $r_p = 1.535 \times 10^{-18}$ m is the electromagnetic "radius" of the proton, N_{tot} is the total number of protons in the ring, ε_n is the 95% normalized transverse emittance, β and γ are the usual Lorentz kinematical factors, and B is the bunching factor, defined as the ratio of the average beam current to the peak current. Note that B is always less than or equal to one.

Results on the evolution of bunch lengths during acceleration in both rings are presented below in the section entitled "RF/RAMP CONSIDERATIONS FOR BOTH RINGS"; they can be used to calculate the bunching factor. They show, for example, that for an intensity of 5×10^{13} and normalized transverse emittances of 200 π mm-mr, the maximum tune shift in the first ring is about 0.4.

The kinematical quantity $\beta\gamma^2$ in the denominator of the tune shift formula increases by a factor of 6.4 as the beam accelerates from 400 MeV to 2 GeV in the first ring. Accordingly, making the beam brighter by about the same factor during the beam transfer from the first to the second ring, either by longitudinal recombination or by transverse compression or a combination of the two, will not violate the space charge limit in the second ring.

Collective instabilities

Collective instabilities are a major concern. Thresholds can be raised by aggressively reducing the impedances that the beam sees. A smooth beam pipe is important in this regard. The following paragraphs on transition and rf systems are relevant to collective instabilities. Instabilities can be counteracted by various means, one of the most effective being active damping systems. A comprehensive discussion of instabilities is beyond the scope of this paper; Mills has recently addressed the subject in the context of a proton driver for a muon collider.[18]

Transition

Several deleterious effects happen around transition for intense beams.[19] To make a long story short, the best (and ultimately, at very high intensity, the only) solution is to move transition out of the operating range. Elegant lattice solutions for high and/or imaginary transition gamma are available.[20]

Rf systems

Providing enough rf voltage and power to rapidly accelerate intense beams is difficult and expensive. Medium-energy synchrotrons have a large frequency swing, which exacerbates the problem. Furthermore, the cavities are a source of high impedances that drive instabilities. Higher-order modes can be damped passively at the cavities, and active feedback to damp the longitudinal modes of oscillation of the beam can be implemented. Other things being equal, rings having fewer cavities should have higher instability thresholds.

Radiation protection

Beam losses raise safety concerns ranging from worst-case scenarios for short-term accidents to routine long-term activation of components. For these reasons, designs for very high-intensity synchrotrons have typically aimed to keep losses below about the 1% level or less. Oftentimes, attempts are made to control where the losses occur in order to protect sensitive components.

As SuperBooster design concepts are described in the following sections, it will become apparent that the SuperBooster has the potential to circumvent the intensity limitations.

MULTI-TURN INJECTION INTO THE FIRST RING

H^- stripping has become the preferred method of multi-turn injection from a linac into a proton synchrotron. However, the SuperBooster uses stripping instead to transfer beam from the first ring to the second one. Hence multi-turn injection into the first ring must be accomplished by other means.

Liouville's Theorem

Liouville's theorem applies to all the injection methods that might be used instead of stripping; in this context, it means that incoming beam cannot be superposed on previously injected beam in the same phase space volume. It is worth noting first of all that there is considerable room in the six-dimensional phase-space acceptance of the first ring to inject many turns without violating Liouville's theorem. The transverse admittances of the ring must be quite large anyway to accommodate the large emittances that are necessary in order to reduce the space-charge tune shift to manageable proportions. The normalized transverse emittances of the Linac beam are about 5 π mm-mr, 40 times smaller than the design admittances of 200 π mm-mr. Thus in principle there is room for 40x40=1600 turns in the transverse phase spaces without violating Liouville or his theorem. The inherent longitudinal emittance of the Linac beam is also much smaller than the longitudinal acceptance of the first ring. The relevant question is whether a method can be designed to populate the available phase space volume efficiently.

Required Number of Turns

The present H^- ion source and Linac combination can produce a beam current of about 50 mA. Injecting one hundred turns (225 microseconds) is a nice round goal; at 50 mA, that would correspond to 7×10^{13} protons, enough to end up with more than 5×10^{13}, even after allowing for modest beam losses. (Recall that, according to the Linac staff, an intensity of 3.8×10^{13} is within reach after minor modifications. In the context of the SuperBooster, raising the linac beam current would be a very worthwhile improvement because it would allow a corresponding reduction in the number of turns to be injected.)

Transverse Stacking

There are at least two ways to implement stacking in transverse phase spaces: programmed closed-orbit motion and resonant injection. To accomplish the first method, a local closed-orbit bump at the injection septum is programmed to move the closed orbit as the injected beam streams in. Different turns then occupy different regions of transverse phase space. The injection process usually begins with the closed orbit near the injection septum, so that the first turns are injected near the middle of the phase space. The closed orbit then moves away from the septum toward the design orbit, so that the early turns do not strike the septum when they come back around, and particles injected later populate the periphery of the phase space.

The second method, resonant injection, can be thought of as a time-reversed version of the familiar process of resonant extraction.[21] To be specific, consider third-integer resonant injection. With the fractional betatron tune near one third or two thirds, beam can be injected on an in-streaming separatrix. The tune and the sextupole strength can be adjusted and programmed to capture the injected particles into the growing stable region of phase space.

It is fortuitous that there are two transverse degrees of freedom and two ways to stack transversely. Accordingly, one method could be used in the horizontal plane and the other one simultaneously in the vertical. The resulting transverse distributions will probably be considerably flatter than Gaussian, resulting in less tune spread across the bunch, which may ameliorate space-charge effects.

The Injection Straight Section

The layout of the injection straight section must accommodate all of the injection stacking methods that will be used. To allow stacking in both transverse phase spaces, the injected beam ought to enter in a corner of the physical aperture, i.e. displaced both horizontally and vertically from the design closed orbit. A thin electrostatic septum should be the last element in the injected beam line. (The circulating beam will then necessarily not contain particles having both horizontal and vertical emittances simultaneously very large; this may help to reduce subsequent losses.) As has been seen, there should be a fast-acting programmed closed-orbit bump across the injection region. Also, it will soon become clear that dispersion at the injection straight section may be useful to facilitate momentum stacking. It is worth noting that the magnetic fields used for injection are limited to about 7 kg to avoid appreciable stripping. Finally, the layout ought to incorporate a method to gracefully dispose of lost beam in a dump.

Longitudinal Stacking

Room is also available in the longitudinal phase space. The longitudinal emittance of the Linac beam is inherently considerably smaller than the bucket area normally used to accelerate the beam. Suppose for simplicity that the incoming beam is adiabatically captured after passing through a debuncher, as is now the case in the Fermilab Booster. In that case the small emittance of the individual Linac rf bunches is deliberately allowed to filament during the capture process at the Booster rf frequency. Even so, there will probably still be some room for momentum stacking.

One way to implement momentum stacking is as follows. (To simplify the discussion of this paragraph, any deliberate offsets of the beam in the transverse phase spaces are ignored.) At injection time in a rapid-cycling synchrotron, the guide field $B(t)$ is near its minimum value and its variation with time is approximately quadratic. Suppose we cause the momentum of the injected beam to vary with time in a corresponding way, perhaps by programming the debuncher phase; suppose further that we start the injection process when the guide field is at its minimum value and that there is dispersion in the injection straight section. Then, since the momentum $p(t)$ tracks the guide field $B(t)$, at any time the beam is coming in on the instantaneous design orbit for that momentum, but previously injected beam of lower momentum is pulled away from the injection septum by the rising guide field. Of course the bucket area must grow large enough at the end of adiabatic capture to contain the entire distribution. That is not a severe requirement: the guide field changes by only 0.1% in 485 microseconds in the first ring proposed here, so that the momentum spread will not be excessive.

Variations on a Theme

Multi-turn injection is a complicated process involving many parameters. The choice of betatron tunes and their possible variation during injection must be specified. Lattice functions and the closed-orbit program must be chosen. Perhaps linear transverse coupling can be deliberately introduced to facilitate injection. The linac beam might be scraped before injection to reduce losses in the ring. The conception and evaluation of further variations on these ideas is left as an exercise for the reader. Optimizing the design of multi-turn injection requires extensive numerical modeling that is beyond the scope of this paper.

STRIPPING OF CIRCULATING NEGATIVE HYDROGEN IONS

H^- ions can be stripped by various means as they circulate. The uncontrolled loss of particles to stripping must be held below about 1% in a practical high-intensity accelerator. Low magnetic fields and good vacuum are required in order to reduce stripping to manageable proportions. However, it is not difficult to meet these constraints in a rapid-cycling low energy machine; furthermore, the design can take advantage of the ease of stripping to facilitate disposal of beam destined to be lost.

Field Stripping

H^- ions can be stripped as they pass through a magnetic field. In the rest frame of the ion, the magnetic field transforms into an electric field that exerts opposite forces on the proton and the electrons. The electric field in the rest frame is proportional to the product of the magnetic field strength B and the ion momentum p; hence the proper lifetime τ of the ion in its rest frame is a function only of the product pB: $\tau = f(pB)$.

The ion lifetime is a very steep function of this equivalent electric field.[22] In a rapid-cycling booster, the ion spends only a few milliseconds near the maximum momentum p_{max}; a one-second ion lifetime at the maximum momentum would then assure losses of less than 1% due to field stripping. (A detailed quantitative design should integrate the stripping probability over the acceleration cycle after parameters are well-defined; such a calculation would also take time dilation into account. Here an approximate treatment is appropriate.) To assure an ion lifetime greater than a second, the product pB must be kept below about 5 kg GeV/c.[22]

$$pB < k = 5 \text{ kg GeV/c}.$$

This inequality can be combined with the relation $p = eB\rho$ to establish the maximum achievable momentum in a ring of harmonic number h=85 and dipole fill factor $f_d = \rho/R = 0.7$. The result is:

$$p^2_{max} = ekf_d h s_B / 2\pi; \qquad p_{max} = 2.83 \text{ GeV/c},$$

where $R = h s_B / 2\pi$ has been used and s_B is the bunch spacing. The corresponding maximum magnetic field is only 1.76 kg. The corresponding kinetic energy is 2.04 GeV; for design purposes, this has been rounded off to 2 GeV. In a real engineering design, the energy would probably be pushed as high as possible to maximize the space-charge limit of the second ring.

Although it may seem wasteful to use a Booster-sized ring to accelerate the ions to only 2 GeV, the tunnel is in some sense "free", as it is needed anyway for the second ring. The component costs of the accelerator are not a strong function of its circumference; low-field magnets are inexpensive to build and to operate.

Gas Stripping

H^- ions can also be stripped by encounters with the atoms of the residual gas. To hold gas stripping to tolerable levels, the pressure must be kept low; but the high injection energy and short cycle time of the proposed ring help to make the vacuum requirements not too stringent. A back-of-the-envelope estimate for the required pressure can be obtained as follows by scaling from the known foil thicknesses required to strip ions during the normal charge-stripping injection process.

It is known (e.g. from experience with the Fermilab Booster) that at 400 MeV, a carbon foil 300 $\mu g/cm^2$ thick will convert 99% of the negative hydrogen ions to protons; the cross sections for stripping fall off with energy as $1/\beta^2$. According to Hojvat and Webber, the cross section for removing the first electron is about three times larger than the cross section for removing the second one.[23] It follows that the "interaction length" for removing the first electron is about 20 $\mu g/cm^2$ of carbon. Assuming that the molecular makeup of the residual gas is equivalent to carbon, a thickness of 0.2 $\mu g/cm^2$ will strip 1% of the ions at 400 MeV.

At what pressure is the residual gas equivalent to 0.2 $\mu g/cm^2$ of carbon? At close to the speed of light, the ions travel 10^9 cm during the

1/30 second acceleration time. Suppose that the species found in the residual gas have densities of about 10^{-3} gm/cm^3 at NTP; then the ions would traverse 10^6 g/cm^2 in 1/30 sec at NTP. The density at atmospheric pressure must then be reduced by a factor of 5×10^{12} to make it equivalent to 0.2 μg/cm^2 of carbon. In other words, the vacuum pressure must be less than about 1.5×10^{-10} torr. This estimate is pessimistic for several reasons: first, the stripping cross section declines somewhat with energy; second, the ion velocity is lower than the speed of light; third and most importantly, the residual gas at that pressure is probably largely hydrogen, which has a low stripping cross section.

More optimistic results could have been obtained by scaling from the pressures required in other H^- accelerator designs. For example, the proposed "Progetto ADROTERAPIA" synchrotron accelerates H^- from 11 MeV to 250 MeV in 0.15 sec and extracts during 0.2 sec. That design estimates losses of 0.75% at a vacuum pressure of 10^{-10} torr, with a gas content of 95% hydrogen and 5% oxygen.[24] Extrapolating that result to shorter cycle time and higher energy implies that 10^{-9} torr would be good enough for the SuperBooster.

The next iteration of the design process should include cross section estimates for the expected molecular makeup of the residual gas and should integrate over the acceleration cycle to account for the slow dependence of the stripping cross sections on energy.

<u>Stripping by Intrabeam Scattering</u>

Colestock has pointed out that the rate of stripping by intrabeam scattering ought to be estimated.[25] Intrabeam scattering refers to Coulomb scattering of beam particles by other particles in the same beam. One way to estimate the magnitude of the effect is to scale from gas scattering. For the parameters considered here, the ion density within a bunch is similar to the density of residual gas molecules; however, a typical velocity of relative motion within a bunch is three orders of magnitude smaller than the velocity of the ions through the residual gas molecules. Stripping by intrabeam scattering is then negligibly small compared to stripping by the residual gas unless the ion-ion stripping cross section at low velocities is 3 orders of magnitude larger than the cross section for stripping on the residual gas. Such a large cross section is highly unlikely because the first electron is bound by about half an electron volt, and it will then take a fairly close and energetic Coulomb encounter to dislodge it. Colestock has arrived at the same tentative conclusion via a different train of thought.[26]

Stripping Collimators

The susceptibility of H^- ions to stripping is an inconvenience, but it can also be used to advantage in the design of the first ring. Recall that activation of components due to inevitable beam losses is a major concern in high-intensity, moderate-energy synchrotrons. The fact that H^- ions strip easily suggests a strategy for gracefully disposing of beam particles that are likely to be lost. Thin stripping collimators upstream of short dipoles can be used to define the limiting physical aperture; the dipoles are

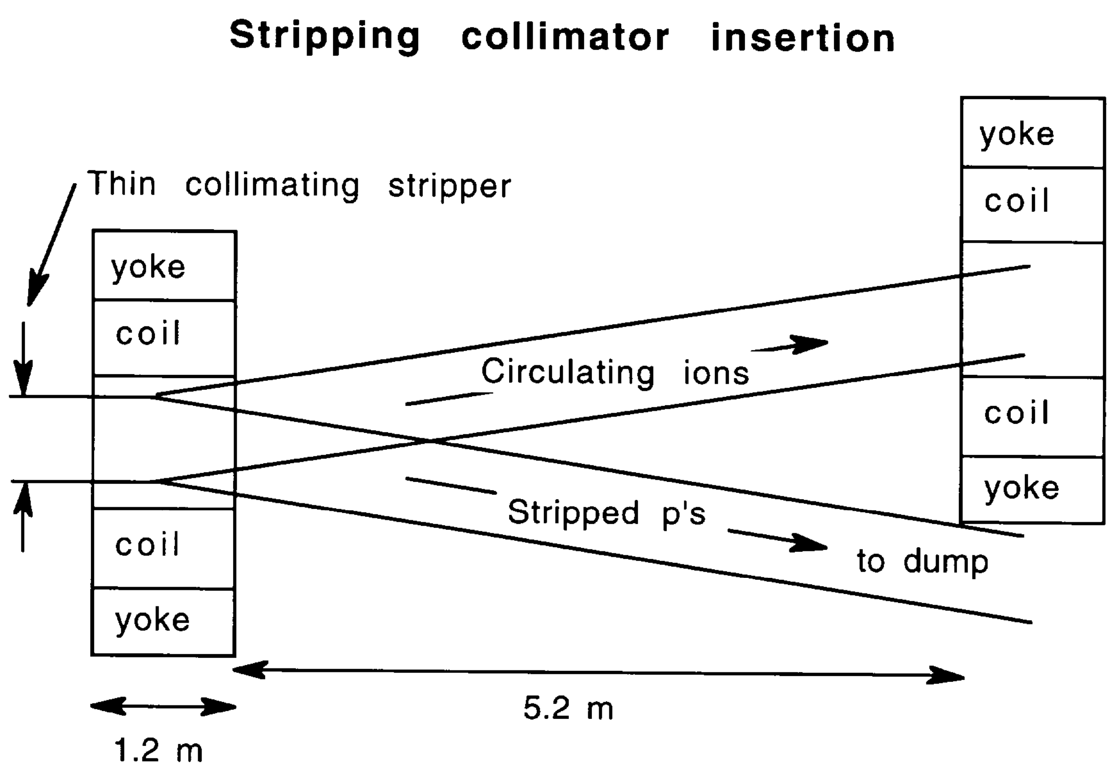

Figure 1. The schematic layout of a stripping collimator insertion. The purpose is to dispose gracefully of halo ions destined to be lost.

followed by straight sections as shown in Figure 1. A few of these insertions can readily be incorporated in the lattice. The stripping collimator converts the halo ions to protons, which then bend outward in the dipole and leave the ring in the straight section. The protons are directed to a dump buried in the outside wall of the tunnel. The dump entrances can be shielded by remotely moving blocks before people enter the accelerator enclosure.

BEAM TRANSFER FROM FIRST RING TO SECOND RING

The crux of the present proposal lies in the way that beam is transferred from the first ring to the second ring. Recall that a laser beam strips the first electrons from the circulating H^- ions, and the resulting neutral atoms stream across to a stripping foil in the second ring. The laser allows flexible selection of parts of the beam in the first ring, and the foil allows superposition of the injected beam in the phase space of the circulating beam in the second ring. This method allows reduction of the transverse emittances and/or overlaying of bunches in the second ring.

Layout of the Laser-Induced Beam Transfer Area

Figure 2 shows schematically a possible layout of the area for laser-induced beam transfer. For clarity in this conceptual sketch, focusing magnets are not shown. The view in the horizontal plane shows how Lambertson magnets can be used to bring the beams close together. The horizontal bends in the Lambertsons are part of the normal lattice of the H^- ring and must ramp to track the momentum of the ions. The view in the vertical plane shows the relative arrangement of laser beam, mirror, and stripping foil to accomplish the transfer. The vertical orbit-bump magnets pulse on for the duration of the beam transfer; while they are on, the proton orbit is aligned with the H^- ion orbit extrapolated from the photon-ion interaction region. The end of the orbit-bump pulse serves to pull the proton beam off the stripping foil after the transfer is complete.

It is important not to strip the ions while they are bending, as that would spray beam around uncontrollably. Three measures serve to confine the photon-ion interactions to the straight part of the ion orbit: using a slight vertical crossing angle between the photon beam and the ion beam, modulating the photon beam at the rf frequency, and limiting the length of the interaction region so that only one ion bunch at a time occupies it.

The crossing angle also solves the problem of getting the photon beam around the stripping foil. For head-on collisions, the foil would have to

Laser-Induced Beam Transfer Insertion

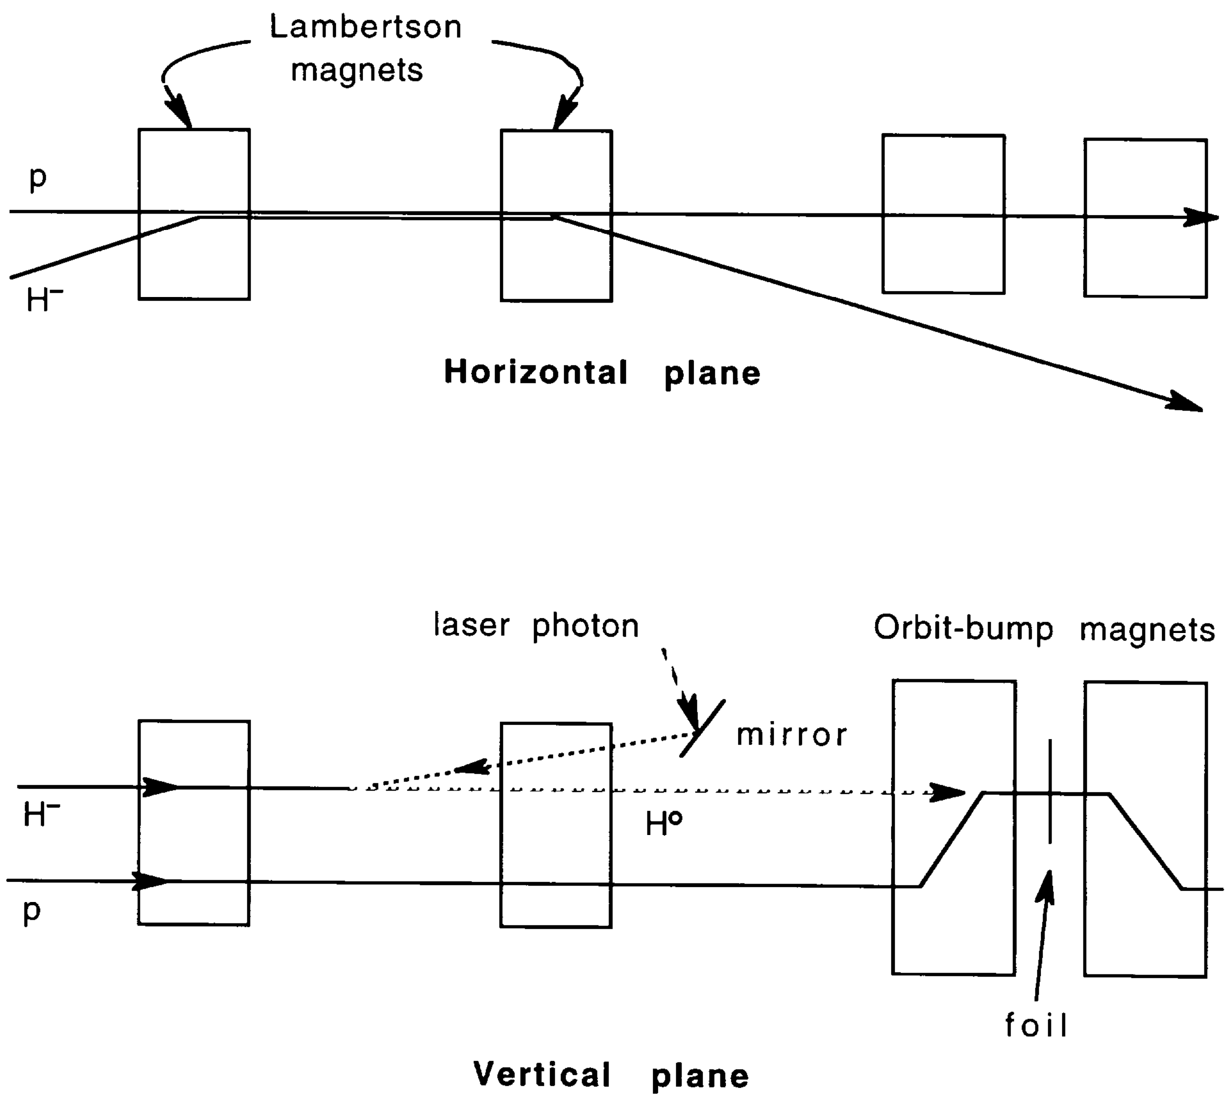

Figure 2. The layout of the insertion for laser-induced beam transfer. The Lambertson magnets bend the ion beam horizontally throughout the cycle; the orbit-bump magnets deflect the proton closed orbit vertically onto the foil during the transfer. The vertical view shows the relative arrangement of laser beam, mirror, and stripping foil to accomplish the transfer.

either serve as a mirror or transmit the laser beam transparently, both of which seem problematical.

Factors Limiting the Duration of the Transfer

Several factors may limit the number of turns that can be used to accomplish the beam transfer. Multiple scattering of the circulating proton beam on the stripping foil will cause emittance growth. The guide fields of the two rings will diverge as one rises and the other one falls. Also, there may be practical limits to the duration of the laser pulse. However, it will not take very many turns to compress the transverse and/or longitudinal phase spaces by the combined factor of six to seven that is allowed by the respective space charge limits of the two rings.

The change of rms transverse emittance ε_{rms} due to multiple Coulomb scattering on the foil is given by

$$\Delta\varepsilon_{rms} = \beta_L \theta_{rms}^2 / 2,$$

where β_L is the lattice function at the foil and θ_{rms}, the rms projected scattering angle, is given approximately by

$$\theta_{rms} = \frac{15 MeV / c}{p\beta} \sqrt{L / L_{rad}}$$

To convert to 95% normalized emittance, the rms emittance must be multiplied by $6\pi\beta\gamma$. At 2 GeV, a carbon foil 540 $\mu g/cm^2$ thick will strip 99% of the hydrogen atoms. Using that thickness and assuming for example that the beta function at the foil is 10 m, the change of normalized 95% emittance is only 0.035 π per pass through the foil. At that rate, hundreds of passes are tolerable, especially for those applications that do not require small emittances from the second ring. Having relatively small beta functions at the foil in the second ring would help to minimize the effect of multiple scattering.

Obviously, the beam optics and the laser optics of the transfer process need careful design integration. Notice that the interaction rate between the photons and the ions is governed by a luminosity calculation analogous to a bunched-beam colliding-beam interaction region with a crossing angle, with all the complexity implied by that. One significant difference is that the ion beam attenuates significantly as it traverses the photon beam, whereas a high-energy luminosity calculation is implicitly a thin-target situation. Extensive numerical simulations are necessary to calculate the luminosity-like integral over the photon-ion interaction region and to optimize the proton distributions in the second ring for various applications.

Other Beam Transfer Methods

Although laser-induced beam transfer is the raison d'être and the most interesting aspect of the SuperBooster, the layout of the rings should also accommodate other transfer methods. One method worth including is to have a simultaneous orbit-bump onto a stripping foil in both rings as shown in Figure 4. This would allow transverse emittance reduction without a laser, but would not provide the bunch-by-bunch selectivity that a gated laser beam would. The orbit-bump magnets for the H^- beam must be weak enough to avoid appreciable field stripping of 2-GeV ions. The availability of a simple alternative method that also accomplishes transverse emittance reduction should be reassuring to readers not totally convinced about the feasibility of the laser transfer process. It might also be important operationally to have another transfer method to cover periods when the laser system needs maintenance.

A routine single-turn kick transfer system should also be implemented; this could double as a beam abort system for the first ring. The beam abort could serve to clean up ions remaining after the laser extraction process. Extraction from the second ring would also be by a single-turn kick; this system could occupy a space in the second ring analogous to the location of the single-turn kick/beam abort system for the first ring.

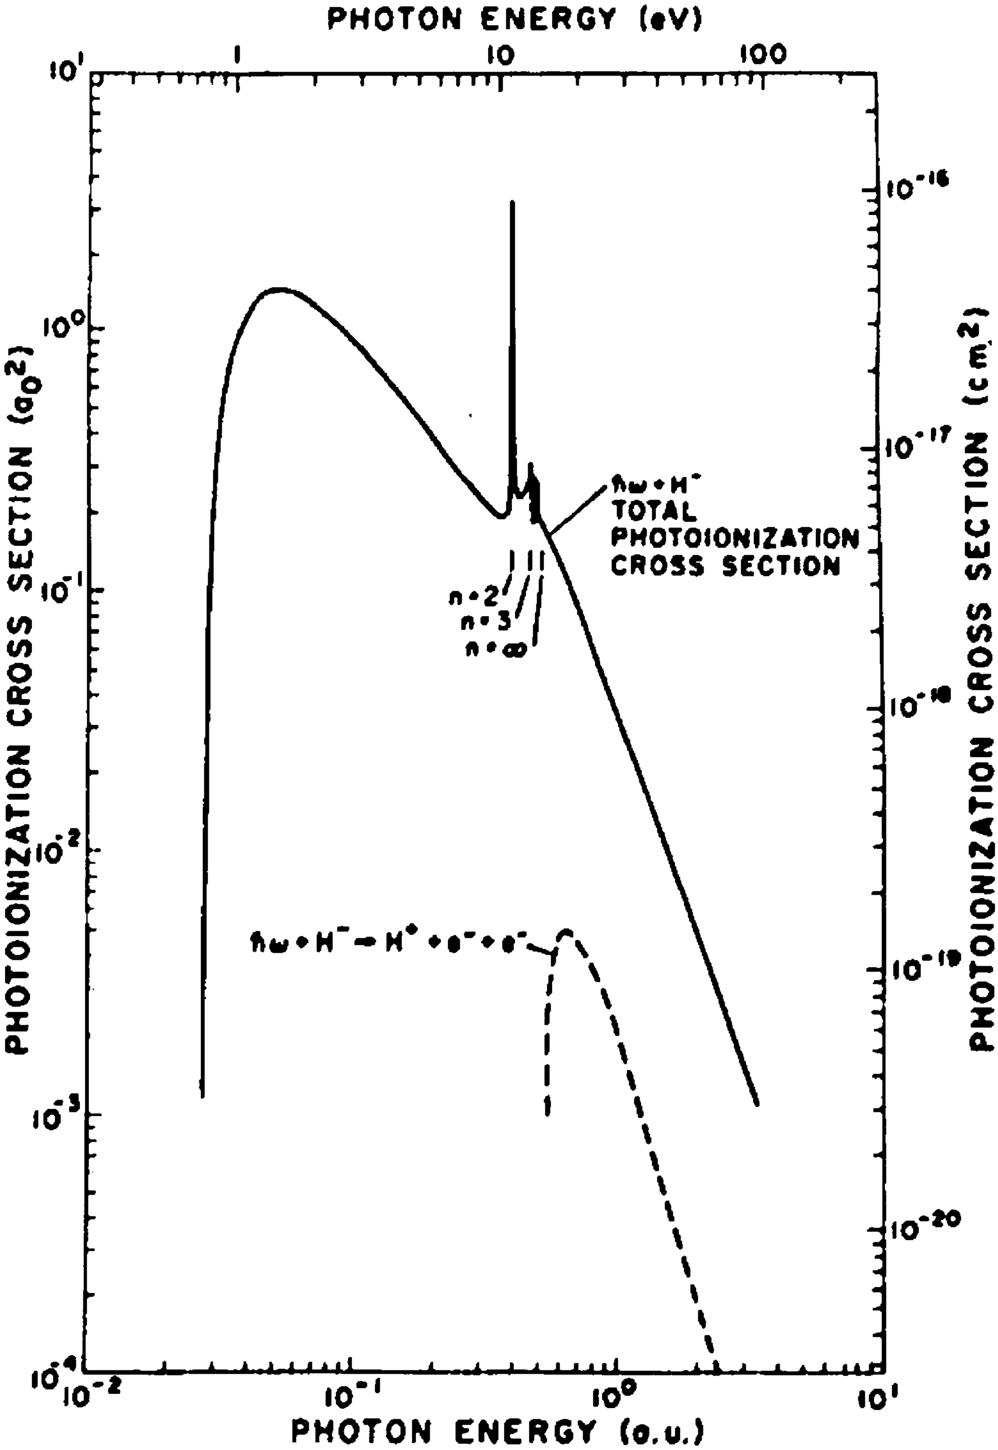

Figure 3. The cross section for photo-detachment of electrons from H- ions as a function of the photon energy. The SuperBooster operating point is on the peak at 1.4 eV on the left side of the curve.

parameters such as the bunch lengths and the emittances of both beams. The optimal conditions may well be different for the different intended uses of the proton beam. However, simple estimates can set the scale for what is required of the laser system. Imagine, then, a very short bunch of N_γ photons with uniform intensity over an area A. Such a photon bunch will be one interaction length thick for the ions if

$$N_\gamma/A = 1/\sigma = 2.5\times10^{16} \text{ photons/cm}^2,$$

where σ is the aforementioned 40 Mb photodetachment cross section. The photon energy per unit area is then given by

$$E_\gamma N_\gamma/A = 1 \text{ mJ/cm}^2.$$

That is a convenient number to remember. The total photon energy $E_{t,\gamma}$ required during a cycle to strip N_B = 85 bunches having area A_B with 99% probability (4.6 interaction lengths) is then given by

$$E_{t,\gamma} = 4.6 N_B A_B (E_\gamma N_\gamma/A) = (391 \text{ mJ/cm}^2)\, A_B.$$

(In deriving this expression it has been assumed that the photons are used only once and not "recycled"; in reality a photon recirculation system, i.e. an optical cavity, could greatly reduce the laser power requirement.)

The area A_B of the beam depends on the normalized horizontal and vertical emittances and beta functions in the interaction region: $A_B = \pi r_x r_y$, with

$$r_x = (\varepsilon_{n,x}\beta_{L,x}/\beta\gamma)^{1/2}$$

and similarly for y. A low-beta insert for the ions at the interaction region is advisable, but a compromise is in order: the ion beam must not be so divergent as to miss the foil. Taking for example betas of 5 m and normalized emittances of 200 π mm-mr in both planes at 2 GeV yields a beam radius of about 1.8 cm and an area of about 10 cm^2. So a laser system capable of delivering about 4 joules at a repetition rate of 15 Hz is required, corresponding to 60 watts. This simple estimate is presented only for illustration; as noted, it assumes that the photons are sent through the interaction region only once.

Optical components for manipulating the photon beam are of course also necessary. The availability of suitable components can be a challenge, particularly in a high-power beam. The crucial idea in the critical area of photon "logic", i.e. of modulating and gating the light beam to create the desired photon time distribution, is to do it at low power, then amplify the resulting photon pulse train. This not only greatly reduces the laser power requirement, but also avoids the problem of having to dispose of a lot of unneeded photons.

A carbon monoxide (CO) gas laser is capable of meeting the requirements outlined above. It produces a narrow band of wavelengths around 5 microns. The physical structure of a CO laser is just like the highly-developed industrial workhorse carbon dioxide (CO_2) gas laser, which produces light at 10 microns. A high-power CO laser needs liquid nitrogen cooling. Its efficiency, defined as the ratio of power in photons divided by wall plug power, is high; numbers in the literature range from 10% to 35%. The desired repetition rate of 15 Hz and pulse length up to hundreds of microseconds are typical of normal operation.

Another alternative would be to use a CO_2 gas laser, which produces wavelengths around 10 microns, together with a frequency-doubling crystal to shift the wavelength from 10 microns to 5 microns. Since CO_2 lasers are more commonly used industrially than CO lasers, the requirements could probably be met with off-the-shelf components, whereas the CO solution may require some development. The choice between CO and CO_2 may depend not only on the laser itself, but on whether the ancillary components for beam manipulations (lenses, shutters, beam splitters, etc.) are more readily available at 10 microns than at 5.

In this section the critical aspects of the laser transfer process have been discussed, and some first-order design concepts have been described.

Obviously, the beam optics and the laser optics of the transfer process need careful design integration. Notice that the interaction rate between the photons and the ions is governed by a luminosity calculation analogous to a bunched-beam colliding-beam interaction region with a crossing angle, with all the complexity implied by that. One significant difference is that the ion beam attenuates significantly as it traverses the photon beam, whereas a high-energy luminosity calculation is implicitly a thin-target situation. Extensive numerical simulations are necessary to calculate the luminosity-like integral over the photon-ion interaction region and to optimize the proton distributions in the second ring for various applications.

Other Beam Transfer Methods

Although laser-induced beam transfer is the raison d'être and the most interesting aspect of the SuperBooster, the layout of the rings should also accommodate other transfer methods. One method worth including is to have a simultaneous orbit-bump onto a stripping foil in both rings as shown in Figure 4. This would allow transverse emittance reduction without a laser, but would not provide the bunch-by-bunch selectivity that a gated laser beam would. The orbit-bump magnets for the H^- beam must be weak enough to avoid appreciable field stripping of 2-GeV ions. The availability of a simple alternative method that also accomplishes transverse emittance reduction should be reassuring to readers not totally convinced about the feasibility of the laser transfer process. It might also be important operationally to have another transfer method to cover periods when the laser system needs maintenance.

A routine single-turn kick transfer system should also be implemented; this could double as a beam abort system for the first ring. The beam abort could serve to clean up ions remaining after the laser extraction process. Extraction from the second ring would also be by a single-turn kick; this system could occupy a space in the second ring analogous to the location of the single-turn kick/beam abort system for the first ring.

Beam Transfer Insertion: Simultaneous Orbit-Bump in Both Rings onto Stripping Foil

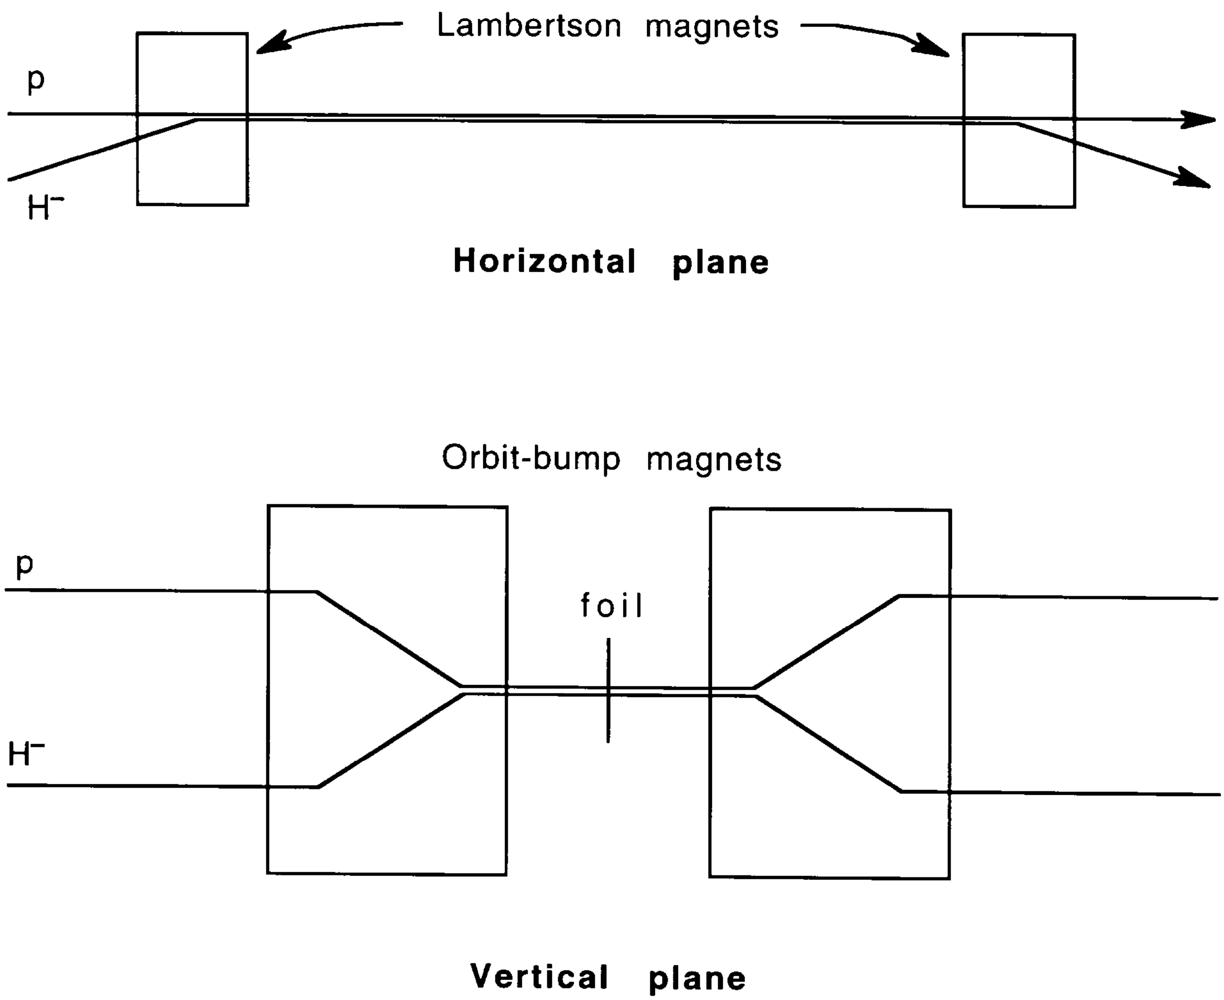

Figure 4. Layout of a region for beam transfer by simultaneously bumping the beams in both rings onto a stripping foil.

RF/RAMP CONSIDERATIONS FOR BOTH RINGS

It is convenient to discuss the rf and ramp requirements for both rings together. The most important conclusion that will emerge is that the use of two rings rather than one leads to considerable savings on the rf systems. That might seem surprising at first; it would seem that the total number of cavities in a two-ring system should be about the same as in a one-ring system because each of the two rings requires only about half as much accelerating voltage as the single ring. But the actual situation is much better than that. The first ring requires only a modest number of cavities because its energy range is modest, from 400 MeV to 2 GeV. The second ring also needs only a modest number of cavities because its rf frequency range is modest, so that much higher cavity voltages can be used. According to Wildman[29], the cavity voltages in the first ring should be like those of the Booster, but the cavity voltages in the second ring should be about four to five times higher like those of the Main Ring. Having fewer cavities means not only less expense but also better performance because the cavities are an important source of the high impedances that drive instabilities.

Magnet Excitation

The main dipole and quadrupole magnets of the ring lattice provide the guide fields for the synchrotron beams. As is customary with rapid-cycling synchrotrons, the magnets comprise the inductive part of LC circuits driven at their resonant frequency, which matches the repetition rate f_{rep} of the synchrotron. The power supply circuit provides a DC bias current as well as a sinusoidal AC current. Furthermore, the magnets run well below saturation. Ideally, then, the magnet currents, the field strengths B(t), and the beam momenta p(t) all follow sinusoidal excitation curves or ramps of the form:

$$p(t) = (p_{max} + p_{min})/2 - [(p_{max} - p_{min})/2] \cos 2\pi f_{rep} t$$

In this expression injection occurs near time t=0 when $p = p_{min}$, and extraction occurs near $t = 0.5/f_{rep}$ when $p = p_{max}$. Of course the ramps in the two rings will be oscillating out of phase, so that the field in the second ring will be at its minimum value when the field in the first ring reaches its maximum value, and vice versa.

The time rate of change of the beam momentum determines the required energy gain per turn ΔE and hence the accelerating voltage which the rf system must supply, according to:

$$\Delta E = 2\pi\ R\ (dp/dt)/c,$$

where R is the average radius of the machine, and the rate of change of beam momentum is found by differentiating the expression for p(t):

$$dp/dt = \pi\ f_{rep}\ (p_{max} - p_{min})\ \sin 2\pi\ f_{rep}\ t$$

It is worth noting that it is possible to use more complicated magnet waveforms in order to reduce the maximum value of dp/dt and hence the peak required rf voltage. One possibility is to add a second harmonic component to the AC current oscillating at the fundamental frequency f_{rep}. Another possibility is a switched system that ramps up relatively slowly and ramps down faster.[30] Clearly a tradeoff between higher power supply costs and higher rf costs is involved. Such power supply complications are probably not cost-effective for the SuperBooster because the rings need only a handful of cavities apiece; henceforth a simple sinusoidal ramp is assumed.

The intriguing possibility that stored magnetic energy might be made to oscillate from the magnets of one ring to the other has not been examined in detail.

Given the time dependence of the magnetic field, the time available for beam transfers can be estimated. To set the scale, the time for the field to change by 0.1% has been calculated. (The experience with the Fermilab Booster is that no serious adverse effects result from a mismatch of that magnitude between the guide field and the beam momentum at injection time.) The results are that it takes 485 microseconds or 216 turns for the field to rise 0.1% above its minimum value in the first ring, and it takes 453 microseconds or 272 turns (at higher velocity and a slightly smaller circumference) in the second ring. It takes 828 microseconds or 490 turns for the field to fall 0.1% below its maximum value in the first ring. Note that these numbers can be doubled by working on both sides of the minimum or maximum values. The upshot is that the time window for multi-turn injection into the first ring is of order 400 turns, and it is of order 500 turns for transfer from the first ring to the second. Of course other factors may place more severe limits on the time available.

Rf Requirements and Longitudinal Beam Parameters

Table 2. Rf and longitudinal phase space parameters.

Rf & Long. p.s. params.	**First Ring**	**2nd Ring**	
Average radius	76.37	75.47	m
harmonic no.	85	84	
gamma t	5	12	
No. protons	5×10^{13}	5×10^{13}	
Norm. emittances	200 π	40 π	mm-mr
Bunch area	0.03	0.06	eV-sec
Bucket area at inj.	0.04	0.08	eV-sec
Max. tune shift	0.4	0.42	
Max. ring voltage	330	600	kV
Voltage/cavity	55	200	kV
No. rf cavities	6	3	

Table 2 shows a number of parameters related to the rf systems and the longitudinal phase space of the beam in the two rings. Some of these parameters are used for the calculated results presented here. The transition gammas are assumed to be 5 and 12, respectively, comfortably beyond the maximum energies of the two rings. (The results are only weakly dependent on the values of gamma-t, so lattice refinements will not qualitatively change the conclusions drawn here as long as gamma-t is safely above the extraction energy.) The harmonic numbers are taken to be 85 and 84. The rf frequency ranges from 37.9 to 50.3 MHz in the first ring and from 50.3 to 52.8 MHz in the second ring. The longitudinal emittance is assumed to be 0.03 eV-sec per bunch in the first ring and 0.06 eV-sec in the second. For stability below transition, the synchronous phase must be below 90 degrees. In fact, the maximum value in the first ring is only about 25 degrees; most of the voltage is needed to provide bucket area. The synchronous phase in the second ring reaches a maximum value of 52 degrees about halfway through the cycle. The bucket area is assumed to be 0.04 eV-sec immediately after adiabatic capture in the first ring and 0.08 eV-sec at injection into the second ring to allow for emittance growth in the first ring.

Figure 5 shows the time dependence during the accelerating cycle of four rf-related parameters for the first ring. The curve labeled E in the figure shows the variation of total energy corresponding to the specified sinusoidal momentum dependence. The curve labeled Fs shows the synchrotron frequency; its maximum value of 21 kHz is high enough to be somewhat concerned about synchrotron-betatron coupling, implying that the rf cavities ought to be in dispersion-free straight sections. The curve labeled Vrf shows the total ring voltage; the curve labeled A shows the bucket area rising monotonically from its value at injection time. In fact the rf voltage curves are numerically adjusted by hand to produce monotonically rising bucket areas throughout the acceleration cycles in both rings.

Figure 6 shows, for the first ring, the time dependence during the accelerating cycle of four parameters related to the beam and its evolution in longitudinal phase space. Δt is the full bunch length and $\Delta p/p$ is the half-width of the fractional momentum spread of the beam. (All widths given relate to emittances containing 95% of the beam.) Note the precipitous drop in bunch length early in the cycle, even though the acceleration rate is fairly gradual. The curve labeled $\Delta \nu$ represents the magnitude of the Laslett incoherent space-charge tune shift; the early rise is caused by the bunch length reduction. The curve labeled Z/n represents the limiting longitudinal impedance calculated according to the naive Keil-Schnell criterion. The latter results depend on beam intensities and emittances; Table 2 shows the values used to calculate the curves. The

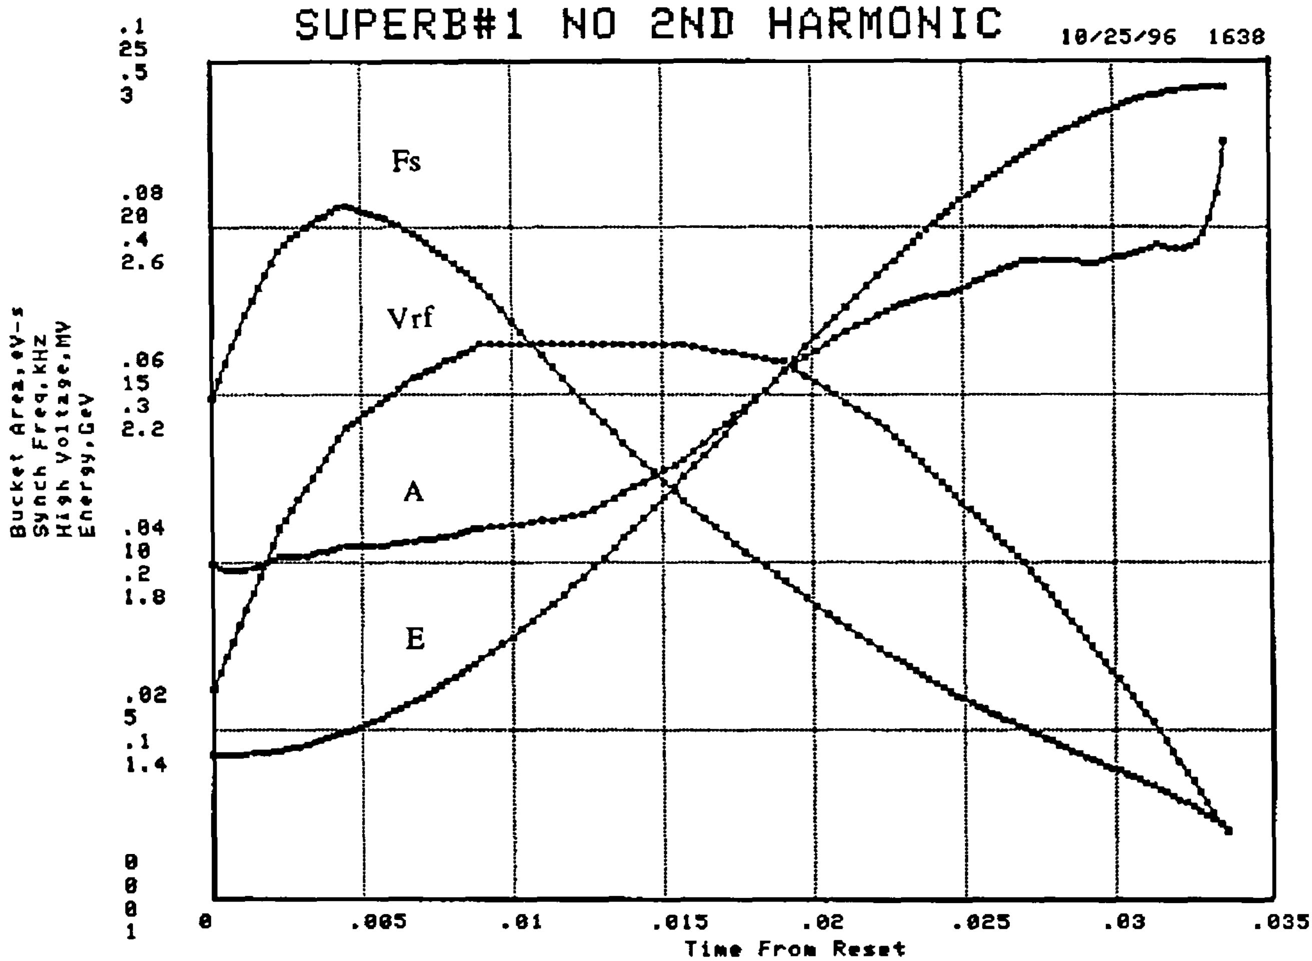

Figure 5. The time dependence of four rf-related parameters for the first ring. See text for details.

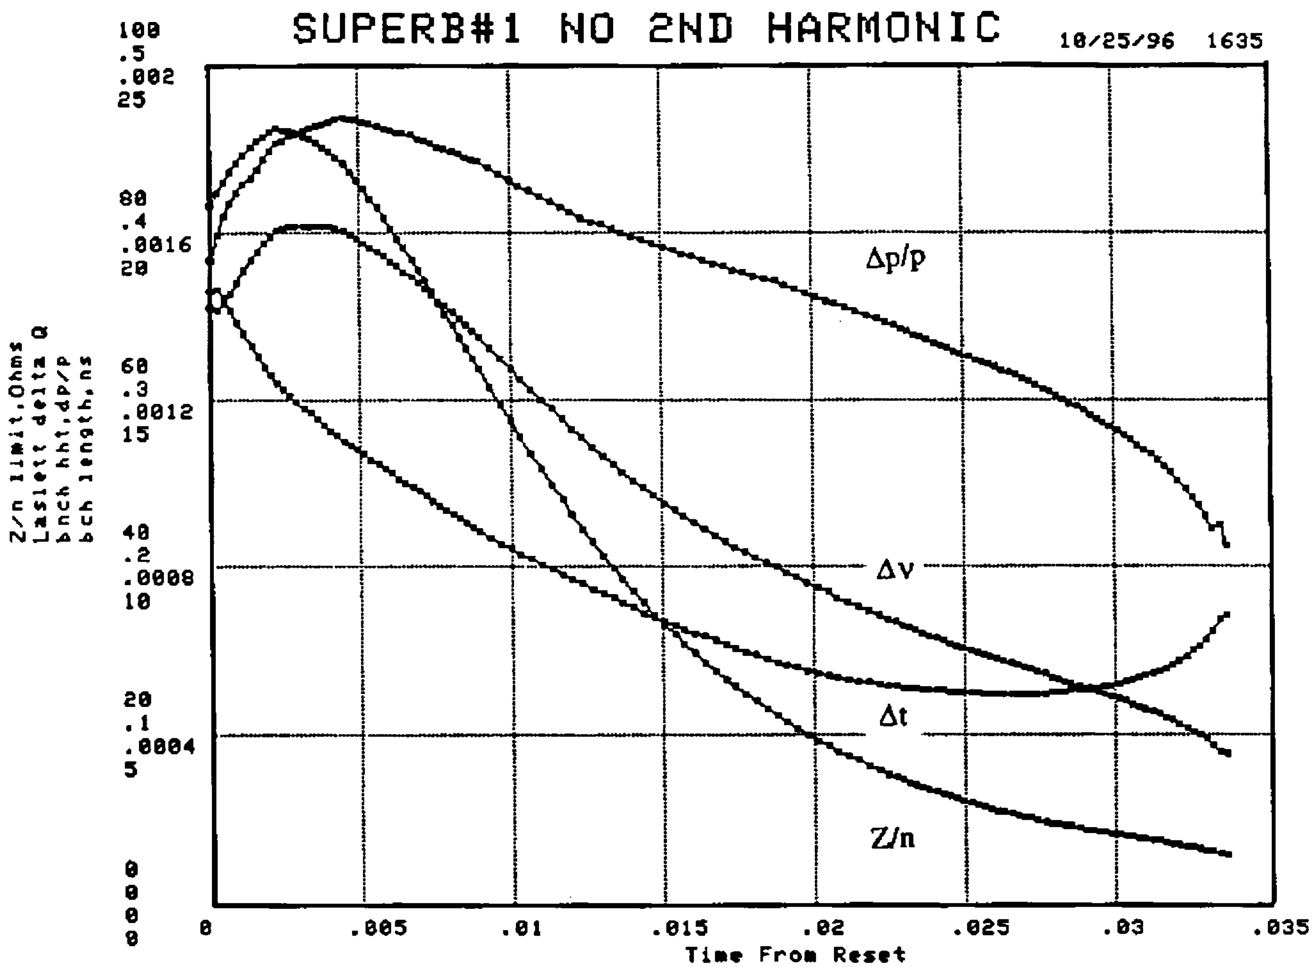

Figure 6. The time dependence of four parameters related to the beam and its evolution in longitudinal phase space for the first ring. See text for details.

space-charge tune shift rises to a value of about 0.4 shortly after injection for the design intensity of $5x10^{13}$; the normalized transverse emittance enclosing 95% of a Gaussian distribution was chosen to be 200 π mm-mr to make the tune shift tolerable, albeit barely. The Z/n limit looks safe everywhere; impedances should be kept small to avoid instabilities near extraction time. Although careful design of modern rings can hold the value of Z/n to about 1 Ohm, cavalierly designed rings can be much worse. Kickers are often the offending elements.

Figures 7 and 8 show the corresponding results for the second ring. The space-charge tune shift rises to a value of about 0.42 for the assumed transverse emittance of 40 π mm-mr. (For this calculation, a transverse emittance reduction by a factor of five during the laser beam transfer is assumed, as would be appropriate for a cycle feeding the Main Injector.) The fact that the tune shifts in the two rings are about the same is by design; the ratio of the kinematic factors $1/\beta\gamma^{-2}$ at injection in the two rings is about the same as the ratio of the normalized emittances for the case presented.

The maximum ring voltages in the two rings are 330 kV and 600 kV. For the first ring, cavities like those in the Booster can provide about 55 kV apiece, so that six cavities are needed. For the second ring with its much smaller rf frequency swing, each cavity can provide about 200 kV as in the Main Ring[31], and as a result only about three are needed.

Although the required ring voltages are manageable, the SuperBooster does place some severe requirements on the rf systems; the results presented so far in this section have ignored longitudinal intensity-dependent effects. The high beam intensities lead to high rf power requirements as well as concerns about intensity-dependent effects such as beam loading and longitudinal space-charge effects. Griffin has examined these aspects and concluded that the requirements can probably be met.[32]

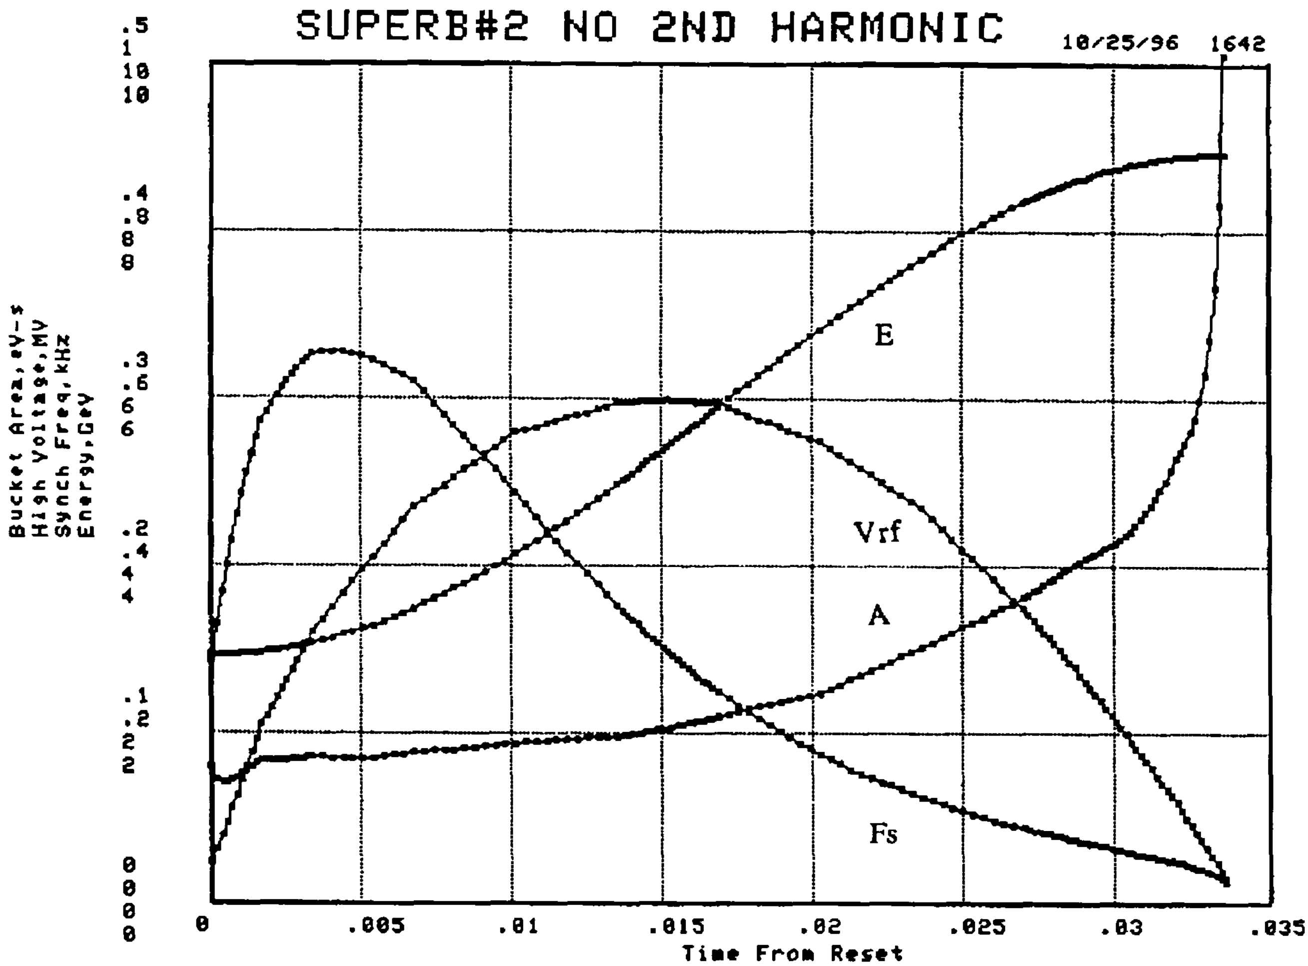

Figure 7. The time dependence of four rf-related parameters for the second ring. See text for details.

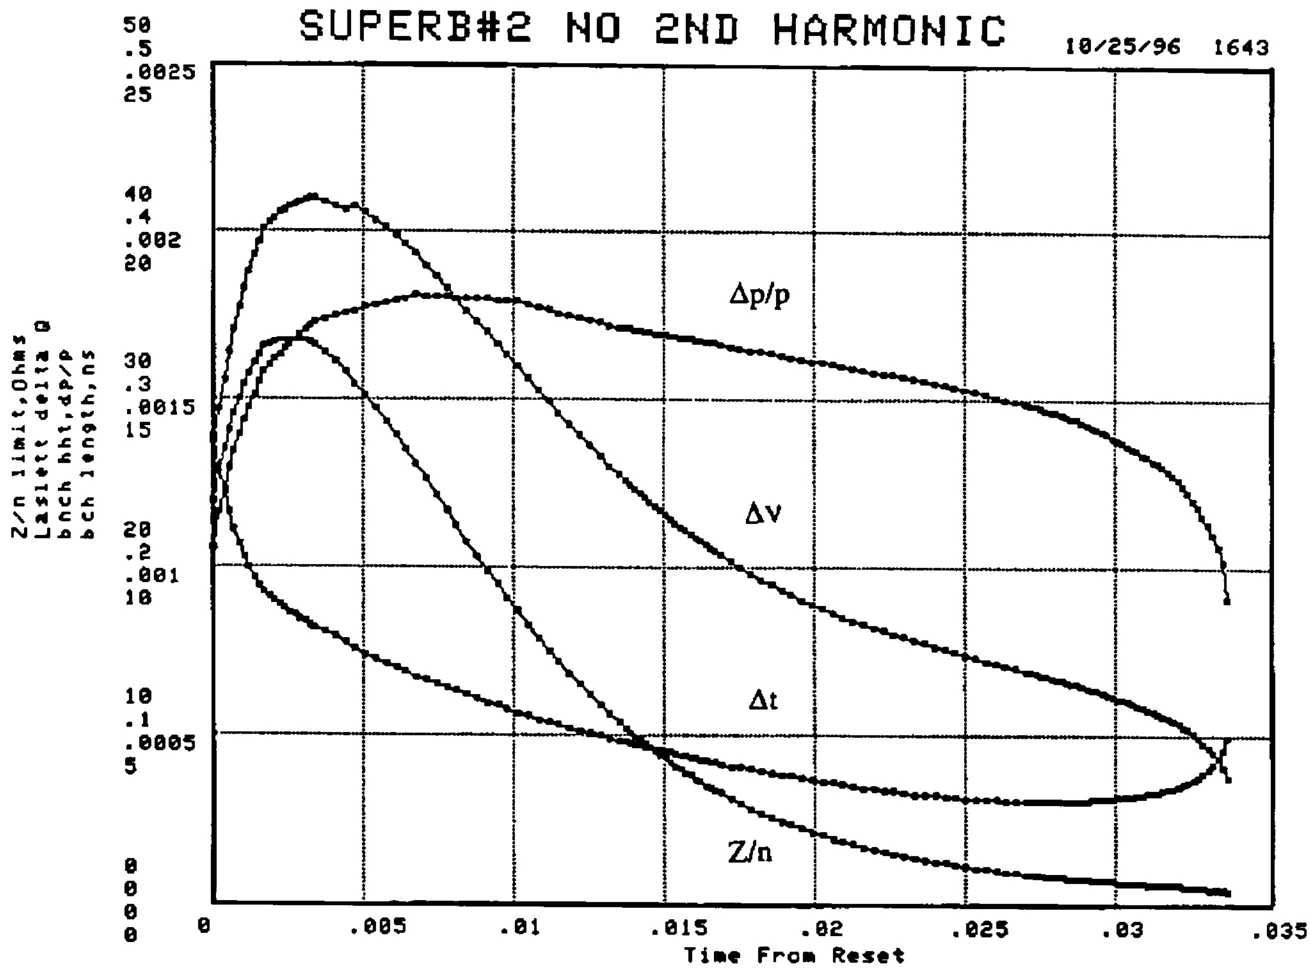

Figure 8. The time dependence of four parameters related to the beam and its evolution in longitudinal phase space for the second ring. See text for details.

The SuperBooster also requires unusually large physical apertures through the cavities. Griffin has pointed out that such cavities are feasible as long as the physical apertures are small compared to the rf wavelengths.[33] The SuperBooster easily satisfies that requirement.

BEAM PIPE DESIGN FOR BOTH RINGS

Beam pipe design is a major issue in the design of rapid-cycling, high-intensity proton synchrotrons. Use of a conventional beam pipe through the dipole magnets is often precluded because eddy currents cause resistive heating as well as induced sextupole fields, either or both of which may be too large to tolerate. At the same time, in order to reduce the impedances that can make intense beams unstable, the pipe ought to look smooth and highly conductive to the beam. These conflicting requirements create interesting design challenges.

Designs for existing and proposed high-intensity, rapid-cycling machines take a variety of interesting albeit complicated approaches to the beam pipe problem. For example, ceramic pipes with intricate conductor patterns to reduce impedances were proposed for the LAMPF II design at Los Alamos.[34] The space-charge impedance of the beam is lower if the wall is close to the beam; the ISIS machine uses a wire cage enclosing the beam inside a ceramic vacuum pipe. The same approach is proposed for the IPNS Upgrade at Argonne.[35] Very thin-walled metal tubes with external ribs have been proposed for other machines.

In the Fermilab Booster, the eddy current problem is circumvented by having no beam pipe at all through the magnets(!); the vacuum enclosure is an external "skin" on the magnets. The price paid for this approach is high impedances that can cause instabilities at high intensity, because there is no smooth conductive wall next to the beam in the dipoles.

A significant advantage of the SuperBooster proposal is that a rather simple beam pipe solution seems applicable to both rings. The problem is tractable because the repetition rate is a moderate 15 Hz and because the maximum magnetic fields are small: only about 1.7 kg in the first ring and about 5.4 kg in the second ring if it has the same dipole fill factor as the first ring. Scaling arguments indicate that both of the rings in the present proposal can adopt the elegant solution used in the AGS Booster at Brookhaven.[36] That machine has an almost conventional thin metallic beam pipe. The departure from convention is that eddy current correction windings are incorporated on the beam pipe. These windings are driven by transformer coupling to the flux in the magnet backleg, which assures that the correction windings automatically track the eddy currents generated in the beam pipe by the varying magnetic field. The beam pipe in the high-energy ring may need to be water-cooled. The Muon Collider Snowmass book proposes the same solution for its proton driver.[37]

LAYOUT AND LATTICE CONSIDERATIONS FOR BOTH RINGS

General Requirements

Any machine intended to replace the Fermilab Booster ought to have the same beam energy and rf structure, because the designs of downstream machines including the Antiproton Source and the Main Injector are based on those Booster beam parameters. Thus the second ring should have the same energy, the same bunch spacing, and about the same circumference as the Booster; it should be capable of producing a batch of 83 or 84 8-GeV bunches at an rf frequency of 52.8 MHz. Although the first ring has lower peak energy, the necessity to avoid field stripping of the H^- ions favors a large circumference as well as a high dipole packing fraction for it. Since the two rings must be close together at the beam transfer location(s), it is natural to put them both in the same tunnel. (Since the two rings will never store beam very long and since neither is much good without the other, occupying the same tunnel is not a big operational disadvantage.)

To facilitate rapid cogging of beam in one ring with respect to the other, the two rings should have incommensurate harmonic numbers, say 85 and 84, 89 and 84, or 83 and 84. In this report harmonic numbers of 85 and 84 are chosen for the first and second rings, respectively.

Straight Section Allocation

Straight sections of various types are needed in the two rings. These will determine the best shape for the rings.

The rf cavities ought to be in one or more dispersion-free straight sections in order to avoid synchrotron-betatron coupling effects. The simplest way to provide space for the cavities is to intersperse them among quadrupoles in one or two straight sections having normal quadrupole spacing.

A long straight section is needed in the first ring for injection from the linac. This ought to be a transversely matched insert devoid of quadrupoles in the middle. As noted previously, dispersion would facilitate momentum-stacking in the injection straight section.

The two rings must have collinear straight sections for the laser-induced beam transfer process. A sketch of how this might be accomplished has already been presented. This should be a matched dispersion-free insert in each ring. A variable-beta insert would be ideal to help optimize the laser stripping and allow matching of the laser-selected beam into the second ring.

There should be side-by-side straight sections for the beam transfer that uses simultaneous orbit bumps onto a stripping foil in the two rings. This might conceivably be integrated with the laser-induced transfer area.

In the first ring, there should be a few short straight sections downstream of short dipoles to dispose of halo beam converted to protons by the stripping collimators. To protect the injection system from excessive losses, a system like this ought to be either integrated into the injection straight section or located just upstream of it.

Finally, there should be provision for extraction from both rings by a single-turn kick to a septum magnet.

Aperture Requirements

Recall that the first ring must have large transverse admittances in order to alleviate space-charge problems. If the maximum beta function in a normal cell is 20 m and the normalized 95% emittance is 200 π mm-mr, the corresponding beam half-size is 63 mm at injection time. In other words, the 95% contour has a 5-inch diameter; there is good reason to keep the beam tightly focused! Although this seems large by Fermilab standards, the dipole field strength is low, only 1.7 kg, so the overall magnet cross sections in the first ring need not be too unwieldy.

The second ring can have somewhat smaller physical apertures than the first one. One reason is that the beam adiabatically shrinks during acceleration in the first ring, so that the geometrical acceptance can be smaller by a factor of p_{max}/p_{linac}, where p_{max} is the momentum at transfer to the second ring. The physical aperture can be smaller by the square root of that ratio. This is another advantage of using two rings. If the machine will only be used to feed the Main Injector, then the aperture of the second ring could be reduced further to match the 40 π mm-mr normalized admittance of the Main injector. In that case, larger emittances from the first ring would always have to be shrunk transversely during the beam transfer process.

Further Lattice Optical Properties

Transition avoidance is an essential requirement for high-intensity rings like these. Any reasonable lattice would naturally have transition well above the maximum energy (gamma = 3.13) of the first ring. However, a normal lattice would have transition in the operating range of the second ring. A so-called flexible momentum compaction lattice can be used to raise transition well above the peak energy of the second ring.[20] Such a lattice can have beta functions that are almost the same as those of the corresponding FODO lattice.

For the first ring, separated-function FODO arcs are probably ideal. A FODO lattice has several virtues including simplicity, and separated-function magnets are preferred for several reasons. First of all, a combined-function magnet has higher field strengths off-axis than on the design orbit, which would cause field stripping unless the peak energy were reduced. Secondly, separate quadrupoles will be needed anyway for matching, straight-section inserts, and so forth; and special short dipoles are needed. The main virtues of a combined-function lattice, simplicity and economy, are compromised when a number of different special quadrupoles and dipoles are also needed. Finally, a separated-function lattice provides adjustability of the betatron tunes via separate control of dipole and quadrupole buses.

THE FERMILAB BOOSTER

Present Booster Performance

"The design intensity in the main accelerator is $5x10^{13}$ protons per pulse and is to be achieved by injecting 13 pulses from the booster synchrotron. This requires that $3.5x10^{12}$ protons be accelerated in each booster cycle." [38]

Although the Fermilab Booster was apparently designed for a beam intensity of about $3.5x10^{12}$ protons per cycle, it originally fell far short of that goal. (The Main Ring also fell short of the transmission of 110% implied by the quote above, but that is another story.) The Booster performance has been enhanced by a number of major upgrades stretching over many years. Conversion to H^- charge-exchange injection allowed overlaying many more protons into the limited admittance of the ring.[39] The Linac Upgrade doubled the injection energy, thereby roughly doubling the space-charge limit.[40] Aggressive implementation of beam damper systems has allowed control of the most damaging beam instabilities.[41] As a result, the Booster can now reliably deliver up to $4.2x10^{12}$ protons per cycle to the Main Ring at 8 GeV.

The Main Injector design calls for $5x10^{12}$ protons per cycle from the Booster. Since the Main Injector aperture will accept larger-emittance beams from the Booster than the Main Ring does, the required improvement of roughly 20% in intensity is well within reach. Furthermore, as has been discussed in connection with TeV33, even the demanding requirements of that program can plausibly be met by taking advantage of the capabilities of the Main Injector and the Recycler. The Main Injector momentum acceptance is big enough to allow slip-stacking to merge two Booster batches for antiproton production. Also, the Recycler, assuming that it performs as expected, will allow emittance reduction by electron cooling for protons as well as antiprotons. Thus for any process that need not be very fast (the cooling times are measured in minutes), such as delivering protons to the Tevatron Collider, the Booster-Recycler combination is capable of delivering not only intense beams but also bright beams of small emittance.

Various proposed physics programs for the next millennium beyond TeV33 (discussed above in connection with the desired performance of the SuperBooster) call for much better performance from the proton source. It is appropriate to examine how the present Booster measures up against the demanding design criteria for a very high performance source. A caveat is in order: any such discussion necessarily includes subjective and debatable aspects.

Booster Performance Limitations

Magnets

The combined-function magnets that provide the Booster guide field have several problems. First, there is no beam pipe through them; the vacuum enclosure is provided by an external stainless-steel "skin". As a result, the beam environment includes exposed magnet laminations as well as abrupt changes in the walls at the ends of the magnets. The resulting large impedances can drive a variety of instabilities for intense beams. Furthermore, the vacuum pressure is rather high in the magnets, typically 10^{-8} torr; and they take a long time to pump down after being let up to atmospheric pressure.

Another problem is that the dynamic aperture is rather small at low energy, probably due to large field errors. The designed non-normalized admittance of the Booster is 90 π mm-mr horizontal by 40 π mm-mr vertical, but the usable non-normalized aperture is considerably smaller, about 25 π by 16 π. In the vertical plane, the dominant limiting apertures may well be the extraction septa hanging into the aperture, and these can be circumvented by dogleg magnets. However, in the horizontal plane, there are no physical apertures to explain the observed small admittance. Raising the linac energy alleviated this problem somewhat because of adiabatic damping; at higher energy a larger normalized emittance fits in a given geometrical acceptance. A current project to implement ramped doglegs around the extraction septa should open up the vertical aperture so that the physical aperture of the guide-field magnets constitutes the limiting vertical aperture.

Oral tradition from the early days (and much of Booster lore is unwritten) carries the rumor that there was a design error in the original sextupole-correcting end-packs that was corrected partway through the magnet production cycle, and that as a result the installed magnets have a variety of different end-packs. If this is true, it is probably not beneficial to the beam.

Rf cavities

The Booster rf cavities were designed and built for the original operating energy range from 200 MeV to 8 GeV, corresponding to a frequency swing from 30 MHz to 53 MHz. In order to tune over this frequency range, the cavities have considerable ferrite, which absorbs considerable rf power. As a result the gap voltage is limited to 30 kV. Consequently, it takes a large number of cavities (presently 17) to generate the required ring voltage. The many cavities present large impedances to the beam; some of the higher-order modes in the cavities drive longitudinal coupled-bunch oscillations. Some of the dominant coupled-bunch modes are suppressed by active damping systems at present operating intensities.

Now that the linac has been upgraded to 400 MeV, the operating rf frequency range is smaller: 37 MHz to 53 MHz. New or modified cavities designed for the smaller frequency range could presumably generate somewhat more voltage, and the number of cavities could be reduced.

The physical aperture through the rf cavities is rather small, only about 6 cm in diameter. Griffin has pointed out that it might be feasible to enlarge the aperture through the cavities; of course this would be a major project.[42]

Lattice

The main problem with the lattice is that the transition energy lies within the operating range. Also, there are no dispersion-free straight sections; the dispersion at the rf cavities produces undesirable synchrotron-betatron coupling effects. Furthermore, a simple separated-function FODO lattice would be preferable to the combined-function FOFDOOD lattice of the Booster.

Extraction

The extraction systems from the Booster consist of a vertical single-turn kick followed by a current-sheet septum. The septum hangs into the physical aperture and constitutes the limiting physical aperture vertically for the circulating beam. The rise time of the extraction kickers is not as fast as the time between bunches, so that at least one bunch is lost on the septum at extraction time. Because of both of these kinds of beam losses, the septum magnet is probably the most radioactive device in the Booster. It is noteworthy that the proposed SuperBooster provides straightforward ways to avoid these problems associated with extraction.

Tunnel

The Booster tunnel is not deep enough; too much radiation resulting from beam losses can reach the surface. The problem is most acute under the Booster Towers. Present beam intensities do not lead to intolerable problems, but the present level of shielding would not support much higher integrated intensities unless fractional losses per cycle were significantly reduced. Of course, activation of components would also be a problem if routine running of significantly higher intensities were not accompanied by lower fractional beam losses. Unfortunately, the tendency in the Booster is for fractional losses to increase with the beam intensity.

Unsuitability as the Second Ring

The proton ring for the SuperBooster is a fairly conventional synchrotron, except for the requirements imposed by the novel injection method. For present purposes little need be said about its design beyond the special considerations already discussed. However, it is appropriate to consider the possibility that the existing Booster in its present tunnel might serve as the second ring, just as in the original B^4 Booster proposal.[1]

(It should be clear for a variety of reasons that the Booster would not be an appropriate choice for the first ring.)

The question to be addressed is whether higher injection energy can overcome the limitations of the Booster. It is true that the normalized admittance and the space charge limit would be increased considerably by injecting at 2 GeV. But not all the problems already discussed would be solved thereby. The transition energy would still lie in the operating range; the transition jump system[43] of the existing Booster is not enough to circumvent all transition-related problems at very high intensity. The beam would still see the large impedances of the magnets and the rf cavities, and magnetic-field errors would be substantial. The amount of ferrite in the rf cavities is much too large for the limited frequency swing of the high-energy ring.

It is undeniably appealing to avoid both the cost of a new tunnel and the considerable bureaucratic effort and time required for approval of new civil construction. However, there is no free space in the existing tunnel for another ring alongside the Booster; substantial relocation of Booster components would be necessary no matter where in the tunnel another ring was located. Also, the existing tunnel is very nearly circular, and it would be difficult or impossible to accommodate the lengthy collinear straight sections needed by the laser beam transfer system. The required extensive modifications to the Booster would have a major operational impact on the high-energy-physics program. Finally, the existing tunnel is not deep enough to provide adequate shielding against the much higher beam intensities that are a goal of the present proposal.

Recycling the Booster

The Booster is part of the original complement of accelerators at Fermilab, along with the Linac and the Main Ring. The Main Ring is being replaced by the Main Injector. The Linac has undergone major changes, including conversion from proton to H^- operation and replacement of the downstream end in order to double the energy. The Booster has also undergone aggressive improvements, but many of its remaining limitations can be neither circumvented nor controverted. Perhaps it is time to park it in the back yard and use it for spare parts. What might be reusable for a new SuperBooster?

Various parts of the rf systems might be reusable, such as bias supplies, power amplifiers, the low-level system, and so forth. Even the rf cavities themselves might be reworked to increase the beam aperture and reduce the amount of ferrite.

Large parts of the gradient magnet power supply system might be reused. The capacitors and chokes that allow resonant operation of the Booster represent a significant investment that might be recouped to serve the same purpose in the new rings.

Much of the injection and extraction systems might be recycled, including beam components and power supplies. Depending on where the new machine is located, most of the existing 400-MeV and 8-GeV beam transfer lines might remain intact, or their components might be reused.

Depending again on where the new machine is located, the Booster service galleries might serve the same purpose for the new machines. Otherwise, they could be converted to valuable work space for people.

It is conceivable (but probably a longshot) that the existing gradient magnets could be converted, by cutting off the pole tips and adding beam pipe, to large-aperture separated-function dipoles suitable for the first ring.

Reusing the existing tunnel is probably also a longshot. Its shielding would have to be augmented. It is not the right shape for the proposed new rings. Furthermore, using the existing tunnel would entail a much longer interruption to the high-energy physics program than would a new tunnel.

The replacement of the Main Ring by the Main Injector is an example of how well the process can be managed when the new ring uses some components from the old ring but occupies a new tunnel.

SITING CONSIDERATIONS

There are several possibly appropriate locations for the new rings. The best solution, albeit probably the most expensive, is a new "green-field" site appropriately located near the linac and the potential users of the output beam. It should be far enough downstream of the linac to allow room for a future linac upgrade to one GeV, as might be needed for a muon collider program.

If the construction of the proposed machines were to occur at some hypothetical future time when antiprotons were no longer used at Fermilab, then the Antiproton Source tunnel might be a suitable location. However, its shielding would probably need to be augmented.

Another location worth considering is a new tunnel well below the existing Booster tunnel, with enough vertical separation to allow occupancy while the Booster is running. The advantages would be the ability to reuse the service galleries, minimal environmental impact (and hence a less bureaucratic approval process), a logical location relative to existing accelerators, and presumably little interference with existing utilities.

SUMMARY AND CONCLUSIONS

Design concepts for a two-ring SuperBooster have been presented. The centerpiece of the design is a flexible method for laser-induced transfer of beam from the first ring to the second one. A plausible case has been made that the design concepts are viable, and no insurmountable obstacles have been identified. To advance to an airtight case for technical feasibility, further work is required on the design of the beam transfers. The multi-turn injection into the first ring requires mainly detailed numerical simulations in order to prove feasibility and to optimize parameters. The laser transfer system also needs numerical simulations, but in addition some hands-on prototyping is in order. Of course a considerable amount of conventional accelerator physics and engineering design work is also required. Given the potential of the SuperBooster for very high performance to serve the varied needs of future physics at Fermilab, it would seem advisable to take the necessary steps to advance the design.

ACKNOWLEDGMENTS

The Proton Source Department of the Accelerator Division held a series of SuperBooster meetings in the Summer of 1995. The development of the ideas in this report has profited from these and other discussions with members of that department. In particular, Milorad Popovic and Chuck Schmidt contributed heavily in the area of Linac performance and Linac upgrade possibilities, and Carol Johnstone and Ray Tomlin to laser requirements and possibilities.

The set of demanding specifications for the so-called proton driver for a muon collider was one of the original stimuli for the ideas embodied in the SuperBooster design. A muon collider working group led by Bob Noble has been meeting regularly, and one of the regular topics of discussion is how to implement the proton driver. Discussions in that forum of both the specific SuperBooster ideas and the challenges of high-performance synchrotrons in general have advanced the concepts embodied in this report. Jim Griffin, Carol Johnstone, Fred Mills, and Milorad Popovic as well as Jim Norem of Argonne National Laboratory are among those who have contributed most regularly and heavily to these discussions.

Pat Colestock and Alvin Tollestrup have provided encouragement both directly and via interesting discussions.

REFERENCES

*Operated by the Universities Research Association, Inc., under contract with the US Department of Energy

[1]C. Ankenbrandt, D. Johnson and R. P. Johnson, "The B^4 Booster", 1977 Fermilab Summer Study, Colliding Beam Physics at Fermilab 2, 369.

[2]C. Ankenbrandt, "Emittance Reduction by H^- Charge Exchange for Hadron-hadron Colliding Beams", Proceedings of the 1982 Summer Study on Elementary Particle Physics and Future Facilities, Snowmass, Colorado (1982) p. 341.

[3]D. Finley, private communication.

[4]W. Chou, semi-private communication.

[5]P.P. Bagley, F.M. Bieniosek, P. Colestock, A. Gerasimov, G.P. Jackson, J.A. Johnstone, J.A. MacLachlan, J.P. Marriner, C. S. Mishra, S.M. Pruss, S. Shukla, R.H. Siemann, W. Wan, "Summary of the TeV33 Working Group", Proceedings of Snowmass 96: New Directions for High Energy Physics, Snowmass, Colorado, June 25-July 12, 1996.

[6]D. Boussard and Y. Mizumachi, "Production of Beams with High Line-Density by Azimuthal Combination of Bunches in a Synchrotron", CERN-SPS/ARF/79-11; 1979 Particle Accelerator Conference.

[7]C. Ankenbrandt, "Slip-Stacking: A New Method of Momentum Stacking", Fermilab FN-352 (1981).

[8]Shekhar Shukla, John Marriner and James Griffin, "Slip Stacking in the Fermilab Main Injector", Proceedings of Snowmass 96: New Directions for High Energy Physics, Snowmass, Colorado, June 25-July 12, 1996.

[9]R. Noble, private communication.

[10]C. Athanassopoulos et al., "Candidate Events in a Search for Muon Antineutrino to Electron Antineutrino Oscillations," Phys. Rev. Lett. 75, 2650 (1995). More recent results were presented by W. Louis at a Fermilab Joint Experimental and Theoretical Physics Seminar on Oct. 4, 1996.

[11]F. Federspiel and H. White, "Letter of Intent to Perform a Neutrino Experiment Using the Fermilab 8 GeV Booster" (1994); Fred Federspiel, Geoff Mills, Rex Tayloe, and Hywel White, "A BOONE Primer", on the World Wide Web at http://nu1.lampf.lanl.gov/~boone/.

[12]G.W. Foster and Ernest Malamud, "Low-cost Hadron Colliders at Fermilab: A Discussion Paper", Fermilab TM-1976, 1996. See also the Pipetron World Wide Web site at http://www-ap.fnal.gov/PIPE/.

[13]C. Ankenbrandt, "The Case for a Small SSC Booster", Proc. DPF Snowmass SSC Summer Study (1984).

[14]G. Jackson, Editor, "Recycler Ring Conceptual Design Study", Fermilab TM-1936 (1995); S. Holmes, "Run II Handbook" (unpublished draft), 1995.

[15]Linac Staff, M. Popovic, Editor, "Present, Near Future, and Future Performance of the Fermilab Linac", Fermilab-Pub-96/046, 1996.

[16]For an overview of the designs of many such machines, see for example "Proceedings of the International Workshop on Hadron Facility Technology", Los Alamos Report LA-11130-C, compiled by Henry A. Thiessen (1987).

[17]L.J. Laslett, "On Intensity Limitations Imposed by Transverse Space-Charge Effects in Circular Particle Accelerators," Proceedings of the 1963 Summer Study on Storage Rings, Accelerators, and Experimentation at Super-High Energies", BNL-7534, p. 325.

[18]F. E. Mills, "Problems in the Design of a Proton Driver for a μ-μ Collider", Fermilab TM-1983, 1996.

[19] For a relatively recent review, see e.g. the articles in the section on Transition Crossing Instabilities in "Proceedings of the Fermilab III Instabilities Workshop", Fermilab TM-1696 (1990), edited by S. Peggs and M. Harvey.

[20] S.Y. Lee, K.Y. Ng and D. Trbojevic, "Minimizing Dispersion in Flexible-Momentum-Compaction Lattices", Phys. Rev. E 48, 3040 (1993).

[21] See, for example, M. Tomizawa, Y. Arakaki, K. Chida, S. Watanabe, T. Watanabe, T. Katayama, M. Yoshizawa, and A. Noda, "Injection Method Using the Third Order Resonance at TARN II", Proceedings of the 1993 Particle Accelerator Conference (IEEE, Piscataway, NJ, 1993) pp. 41-43.

[22] G. M. Stinson, W. C. Olsen, W. J. McDonald, P. Ford, D. Axen, and E. W. Blackmore, "Electric Dissociation of H^- Ions by Magnetic Fields", NIM 74, 333 (1969); Andrew J. Jason, Daniel W. Hudgings, and Olin B. van Dyck, "Neutralization of H^- Beams by Magnetic Stripping", IEEE Trans. Nucl. Sci. NS-28, 2704 (1981)]; Leonard B. Scherk, "An improved value for the electron affinity of the negative hydrogen ion", Can. J. Phys. 57, 558 (1979).

[23] Robert C. Webber and Carlos Hojvat, "Measurement of the Electron Loss Cross Sections for Negative Hydrogen Ions on Carbon at 200 MeV", IEEE Trans. Nucl. Sci. NS-26, 4012 (1979).

[24] G. Arduini, R.L. Martin, S. Rossi, and M. Silari, "The H^-/Light Ion Synchrotron for Hadron Therapy. I. Vacuum and Magnetic Field Constraints", TERA Note 93/12 (1993).

[25] P. Colestock, private communication, 1995.

[26] P. Colestock, private communication, 1996.

[27] Carol Johnstone and Ray Tomlin contributed heavily to the development of the ideas in this section via both verbal and unpublished written

communications. One of the documents that emerged from that process to become a Fermilab internal note is "Laser Stripping of Relativistic H- Ions with Practical Considerations", Fermilab TM-1957 (1995) by R. Tomlin.

[28]John T. Broad & William P. Reinhardt, "One- and two-electron photoejection from H^-: a multichannel J-matrix calculation", Phys. Rev. A 14, 2159 (1976).

[29]D. Wildman, private communication, 1995.

[30]W.F. Praeg, IEEE Trans. Nucl. Sci. NS-30, 4 (1983).

[31]D. Wildman, private communication, 1995.

[32]J.E. Griffin, "Aspects of Operation of the Fermilab Booster RF System at Very High Intensity", Fermilab TM-1968 (1996).

[33]J.E. Griffin, private communication.

[34]"The Physics and a Plan for a 45 GeV Facility that Extends the Higher-Intensity Capability in Nuclear and Particle Physics," Los Alamos National Laboratory Report LA-10720-MS (1986). See also L. Walling, D. Neuffer, H. A. Thiessen, and G. Spalek, "Transmission-Line Impedance Measurements of Beam Pipes for an Advanced Hadron Facility", in "Proceedings of the International Workshop on Hadron Facility Technology", Los Alamos Report LA-11130-C, compiled by Henry A. Thiessen (1987), p. 279.

[35]"IPNS Upgrade: A Feasibility Study", Argonne National Laboratory Report ANL-95/13, April, 1995.

[36]G.T. Danby and J.W. Jackson, Part. Accel. 27, 33 (1990).

[37]The $\mu^+\mu^-$ Collider Collaboration, "Muon Muon Collider: A Feasibility Study", J. Gallardo, Editor. BNL-52503, Fermilab Conf 96/092, LBNL-38946 (1996).

[38]The National Accelerator Laboratory Design Report, Second Printing, July 1968, page 9-10.

[39]C. Ankenbrandt, C. Curtis, C. Hojvat, R. P. Johnson, C. Owen, C. Schmidt, L. Teng, and R. C. Webber, "H^- Charge Exchange Injection Systems", Proc. 11th Int. Conf. High-Energy Accelerators, CERN (1980) p. 260.

[40]"Fermilab Linac Upgrade: Conceptual Design", Fermilab internal document, unpublished, July, 1989.

[41]C. Ankenbrandt, E. Higgins, Jr. and R. P. Johnson, "Suppression of Transverse Instabilities by Fast Feedback in the Fermilab Booster", IEEE Trans. Nucl. Sci. NS-24, 1698 (1977); James M. Steimel Jr. and Dave McGinnis, "Damping in the Fermilab Booster", Proceedings of the 1993 Particle Accelerator Conference (IEEE, Piscataway, NJ, 1993) pp. 2100-2102; James M. Steimel Jr., "Fast Digital Dampers for the Fermilab Booster", Proceedings of the 1995 Particle Accelerator Conference; James Steimel, Jim Crisp, Hengjie Ma, John Marriner, and Dave McGinnis, "Overview of Coupled Bunch Active Damper Systems at FNAL", Fermilab Conf - 96/102.

[42]J.E. Griffin, private communication.

[43]W. Merz, C. Ankenbrandt, and K. Koepke, "Transition Jump System for the Fermilab Booster", Proc. 1987 Particle Accelerator Conf., Washington, DC (1987).