USING PROJECT X AS A PROTON DRIVER FOR MUON COLLIDERS AND NEUTRINO FACTORIES*

Charles Ankenbrandt and Rolland P. Johnson
Muons, Inc., 552 N. Batavia Ave., Batavia, IL, 60510 USA

Abstract

Muon colliders and neutrino factories impose very demanding requirements on the proton accelerator systems that are used to make the muons. Various advanced concepts to satisfy those needs have been developed. Fermilab is proposing Project X, a major intensity upgrade featuring a powerful 8-GeV H$^-$ linac. This paper describes a way to use the Project X linac to meet the needs of these major facilities that are based on muon storage rings. The concept makes use of one or two 8-GeV proton storage rings, followed by an external bunch combiner if necessary, to create the desired proton bunch structures at the pion production target. The concept is compared with other approaches.

INTRODUCTION

Fermilab is proposing a major intensity upgrade known as Project X [1]. The centerpiece of the proposal is a superconducting 8-GeV H$^-$ ion linac based on ILC technology. Initially the linac will supply protons to experiments that use 8-GeV protons directly and to the Main Injector for further acceleration for neutrino physics. A possible layout of the linac in relation to other accelerators on the Fermilab site is shown in Figure 1.

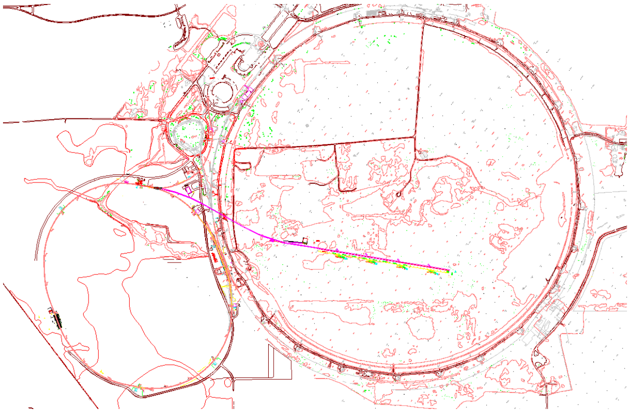

Fig. 1: A possible linac location on the Fermilab site.

Eventually the proton beam may be used to produce muons for programs based on muon storage rings, i.e., neutrino factories and muon colliders. Anticipating that eventuality, the Fermilab Directorate recently asked a group of Fermilab scientists [2] for advice on whether the Project X linac could serve as a driver for major muon facilities and on how best to use it for that purpose.

Lively discussions resulted in the following advice: “The three most important conclusions are as follows:

“1) If muon colliders and neutrino factories are separately designed and optimized, the front ends tend to diverge somewhat because muon colliders need luminosity whereas neutrino factories need flux. Nevertheless, there is considerable overlap between the proton beam power needs of energy-frontier muon colliders and those of neutrino factories based on muon storage rings. In many ways, muon colliders are somewhat more demanding on their front ends than neutrino factories, so any facility that meets the beam-power needs of the former is likely to meet the needs of the latter.

“2) Several muon collider design efforts have generated parameter sets that call for proton beam power of several megawatts. The most common requests fall in the ballpark of 3 to 4 MW; however, most designs are optimistic and none have been fully vetted, so it is advisable to provide considerable performance contingency. The required proton beam power is not likely to be a strong function of the center-of-mass energy of the collider.

“3) Several alternatives have been examined including synchrotron-based ones. The most promising front end is based on the Project X 8-GeV H$^-$ linac upgraded to about 3 MW, with a further upgrade path to ~10 MW held in reserve. One or more 8-GeV storage rings will be needed to provide stripping and accumulation, formation of the appropriate number of bunches, and bunch shortening. Of course an appropriate multi-megawatt target station will also be necessary.

“There are two main recommendations:

“1) The performance requirements on the aforementioned 8-GeV storage ring(s) are severe. Accordingly, a design study should be initiated. The main goals should be to establish design concepts and explore potential limitations due to beam instabilities.

“2) Planning should be initiated for an appropriately located muon test area that can evolve into a facility capable of handling several megawatts of proton beam power.”

This paper is intended to provide some insight into the considerations underlying that advice. In particular, the paper first examines the proton beam requirements of these muon facilities and then describes a way to meet those needs via a combination of one or two proton storage rings, possibly followed by a “trombone” and a “funnel” to combine bunches on the way to the pion production target. Finally, other alternatives are briefly considered.

---

*Supported in part by DOE STTR grants DE-FG02-05ER86252 and DE-FG02-08ER86350

[#]chuck@muonsinc.gov

PROTON BEAM REQUIREMENTS

The International Scoping Study (ISS) [3] has provided a well-defined set of major parameters specifying the requirements on the proton driver for a neutrino factory based on a muon storage ring. These specifications reflect a current consensus that was reached after some evolutionary development from previous neutrino factory design studies [4]. Meanwhile, the Neutrino Factory and Muon Collider Collaboration (NFMCC) and the Muon Collider Task Force (MCTF), working together, have considered various parameter sets for a multi-TeV muon collider, but no consensus has yet been reached. However, the so-called front ends for neutrino factories and muon colliders share a common design. (The front end starts at the pion production target and includes systems for collection and decay of the pions as well as capture and bunch rotation of the muons.) Accordingly, some of the requirements on the proton beam are the same for neutrino factories and muon colliders. In particular, the preferred proton energy is 10 ± 5 GeV, corresponding to the location of a broad peak in a plot of the production rate of useful muons normalized to beam power as a function of proton beam energy[5]. Figure 2 shows those results.

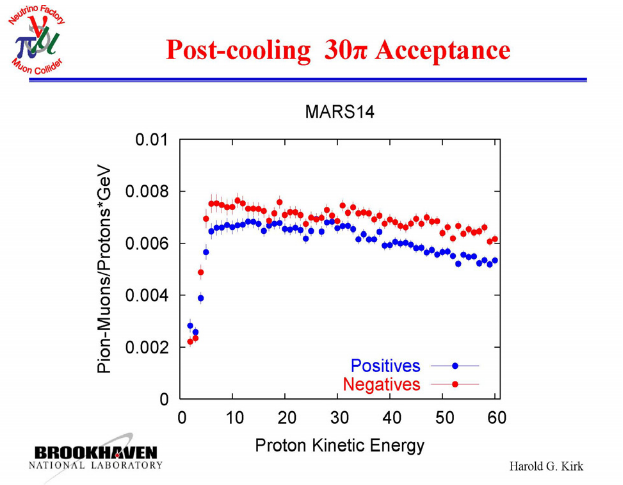

Fig. 2: Study of optimum proton energy for a NF.

Also, the front-end processing of the muons in longitudinal phase space requires a very short rms proton bunch length, specified to be 2 ± 1 ns at the pion production target. In the following two subsections, the requirements specific to neutrino factories and muon colliders, respectively, will be examined.

Neutrino factories

The ISS specifies an average proton beam power of 4 MW and a repetition rate of 50 Hz for the proton accelerator for a neutrino factory. (However, the operating time per year is projected to be only 10^7 seconds, so the integrated beam power amounts to 1.26 megawatt-years per calendar year.) Each accelerator cycle is to deliver 3 or 5 bunches at a time, resulting in 150 or 250 proton bunches per second striking the pion production target. The three or five bunches per cycle must be spaced in time in specific ways so as not to violate constraints imposed by the target and by beam loading in the muon acceleration systems. For the details, see the ISS document [3].

Muon Colliders

For muon colliders at the energy frontier, the situation is more complicated and less settled. The NFMCC and the MCTF have considered many sets of major parameters, reflecting differences of opinion about the most promising approaches to various design challenges. Recent attention has focused on three sets of major parameters, the so-called Low, Medium, and High Emittance parameter sets shown in Table 1. These parameter sets have appeared in many documents; the version shown here is adapted from a recent presentation by Jansson [6].

Table 1: Three parameter sets for a muon collider

	Low ε	Med ε	High ε	
CM Energy	1.5	1.5	1.5	TeV
Luminosity	2.7	1	1	10^{34} cm^2/s
N_μ/bunch	0.1	1	2	10^{12}
No. bunches	10	1	1	/charge
Ring circumf.	2.3	3	8.1	Km
$\beta^* = \sigma_z$	5	10	10	Mm
Dp/p (rms)	1	0.1	0.1	%
Ring depth	35	13	135	M
Mu survival	30	4	7	%
ε_T	2.1	12	25	π μm
ε_L	370,000	72,000	72,000	π μm
PD Rep rate	65	24	12	Hz
PD Power	≈4	≈6	≈4	MW

Regarding proton drivers, the first thing to notice about Table 1 is the last row: all three parameter sets call for proton driver beam power of about 4 to 6 MW. (The parameter sets have evolved a little since the advice quoted in the introduction above was written, resulting in a need for somewhat higher proton beam power.)

For ease of comparison, all three parameter sets were generated with the same CM energy of 1.5 TeV. However, it is worth noting that parameter sets that have been generated for higher CM energies in the multi-TeV range do not call for very different amounts of proton beam power. This can be understood as follows: The designers typically want similar rates for pointlike processes, whose cross sections scale inversely with the square of the CM energy. That means that the desired luminosity tends to go up with the square of the CM energy. In the formula for luminosity, one factor of energy comes from the adiabatic shrinkage of transverse beam sizes with energy. Another factor can come from the fact that, for a given longitudinal emittance of a muon bunch, the bunch length can go down as the energy goes up because there is a maximum tolerable fractional momentum spread dp/p. If the bunch length can be reduced, so also can the lattice beta function at the interaction point. Thus the same muon bunch intensities can provide the desired rise of luminosity with energy. (Admittedly, there may also be a psychological factor at work: designers of multi-TeV colliders may fear being deemed too greedy if they ask for very high luminosity as well as very high CM energy.)

Note that the low emittance parameter set has high repetition rates, more bunches in collision, and low muon bunch intensities, and *vice versa* for the high emittance set. Basically, those who lean toward the low-emittance set are concerned that large muon bunch intensities may not be feasible, whereas those who prefer the high-emittance set worry that very low muon emittances may not be achievable. The medium-emittance parameter set is a possible compromise.

Whether by accident or by design, the three parameter sets have similar values for the ratio of the intensity to the transverse emittance of the muon bunches. This ratio is close to the brightness corresponding to the expected maximum tolerable value of the beam-beam tune shift in the collider ring. With some plausible additional assumptions, it can be shown that under those circumstances the proton beam power should come out about the same, independent of whether the muon bunch intensities and emittances are separately low, medium, or high. That is, the luminosity can be written as a product of beam-beam tune shift and muon beam power, and it is plausible that the muon beam power is proportional to proton beam power, so if the luminosity and the beam-beam tune shift are the same, then the proton beam power is the same.

Conversely, if it is not possible to produce or to accelerate such bright muon bunches, the luminosity will decrease unless the proton beam power is raised, for example by delivering more proton bunches per second to the pion production target. There are also significant uncertainties in the estimates of muon survival probabilities. Those are some of the main reasons why considerable performance contingency is called for in the advice quoted in the introduction.

Cooling Schemes and New Technology

Figure 3 shows the emittance evolution of the beams in the three cooling schemes, where individual components of the cooling channels are indicated [7]. The beam cooling is based on the principle of ionization cooling, the only process that can provide the needed emittance reduction in the short muon lifetime [8]. The high emittance option uses a large helical structure called a "Guggenheim" based on large focusing solenoids interspersed with large aperture RF cavities. The solenoids are tilted to provide the bending and dispersion needed for emittance exchange, with wedge absorbers to achieve 6-dimensional cooling. For its final stage of cooling, the high emittance scheme uses very strong solenoids that may be enabled by newer superconductors developed for high temperatures but used here at LHe temperature.

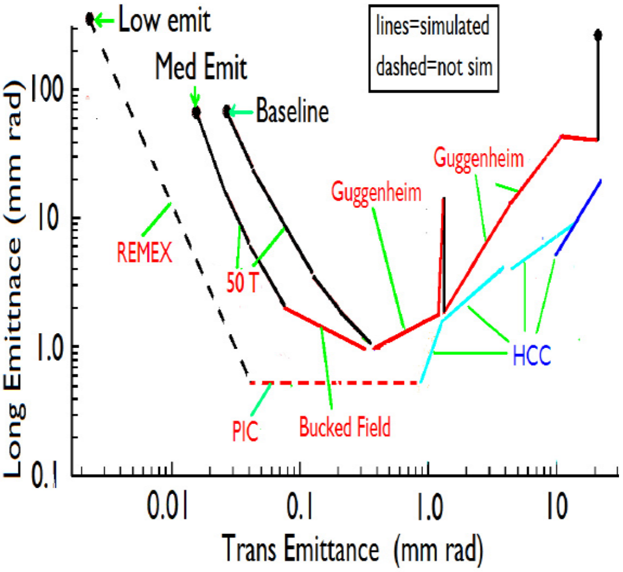

Fig. 3: Fernow-Neuffer plot of emittance evolution for the low, medium, and high emittance schemes.

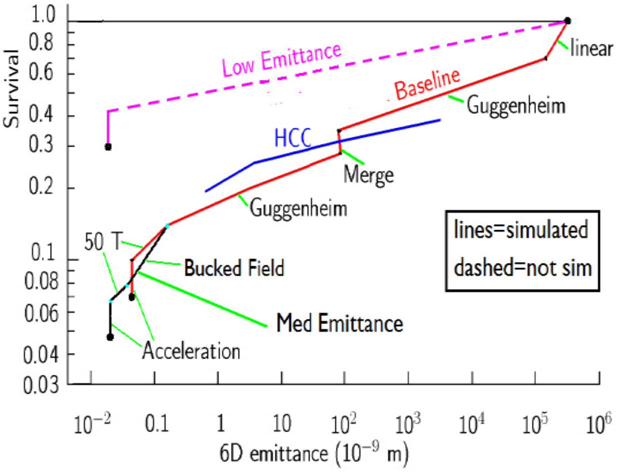

Fig. 4: The fraction of muons that survive after successive processes for the three schemes.

The low emittance option uses Helical Cooling Channels (HCC) with a continuous absorber and a special magnet system incorporating solenoid and helical dipole and helical quadrupole fields [9]. To get to even lower emittance, a technique called parametric resonance cooling (PIC) is required [10]. Reverse Emittance Exchange (REMEX) [10], the final stage of the low emittance option, involves the use of wedge energy absorbers to transfer the emittance from the transverse dimensions to the longitudinal one. That would allow the luminosity to be increased up to the point that the bunch length becomes comparable to the betatron wavelength at the interaction point in the collider. While the theoretical and analytical work has indicated that these last two techniques should work, complete simulations using plausible engineering designs are yet to be made.

Figure 4 shows some initial estimates of the muon survival rates for the cooling schemes shown in the previous figure [7]. Each of the three schemes depends on technology that is yet to be fully realized. For example, vacuum RF cavities that are required for the Guggenhim channel and perhaps other later cooling segments have been shown to be adversely affected by the strong magnetic fields needed for effective beam cooling. Mitigation methods that are being studied include bucked fields to reduce the fields at the cavities and methods to suppress the limitations due to dark currents (magnetic insulation) [11].

RF cavities in the HCC segments are assumed to be filled with dense hydrogen gas in operation. hile such cavities have been shown to work well in strong magnetic fields [12], they are yet to be tested in the conditions of a beam of intense ionizing radiation. The technology of the PIC and REMEX cooling segments is under development as well, where the required magnetic aberration control, which has been solved analytically, is yet to be demonstrated in a complete simulation.

Recent progress in muon beam cooling and beam manipulation concepts has been very encouraging. In the next few years even more new ideas can be expected; conversely, some ideas may not live up to expectations.

In the Tevatron proton-antiproton collider, the intensity of the cooled antiproton beam has usually been the most important limitation to luminosity. It is quite likely that will be the case for a muon collider as well, since in many ways the problem is compounded in that both the $\mu^+$and μ^- beams will have to be collected and cooled. One hopes that sufficient resources will be applied to develop production, capture, cooling, acceleration, and collision ideas to the point that the current estimates of required proton driver beam power are realistic. However, it is very likely that those estimates are optimistic, especially for the first years of operation.

MEETING THE BEAM REQUIREMENTS

An important goal is to design a proton driver facility that can meet the needs of both a neutrino factory and a muon collider. Two major considerations dominate the proton driver design: the proton beam power and the required bunch structures.

The required beam power is at least 4 MW at 8 GeV, with considerable performance contingency. Eventually a decision must be made whether to design *a priori* for more than 4 MW or to incorporate the potential for a future intensity upgrade in the design.

Regarding proton bunch structures, the neutrino factory and muon collider parameter sets call for widely varying numbers of bunches per second delivered to the pion production target. In the neutrino factory case, the bunches are to come in bursts of three or five. Accordingly, another important design goal is to develop a bunching strategy that allows considerable flexibility.

The design concept that has resulted is straightforward: two 8-GeV storage rings, followed by a system to deliver bunches simultaneously at the pion production target if necessary. The storage rings must have very large apertures and relatively small circumferences to allow mitigation of space-charge effects while creating bunches of relatively small longitudinal emittance. Figure 5 is an example based on the use of the existing Fermilab rings, now used for antiproton production, which may be available in the era of neutrino factories and muon colliders.

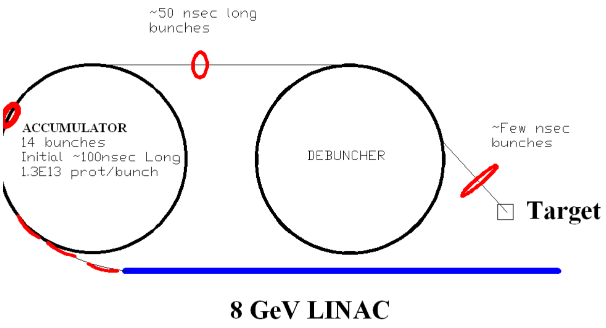

Fig. 5: Proposed use of the Fermilab antiproton accumulator and debuncher rings to form intense, short proton bunches for muon production.

The Accumulator

The first storage ring will accumulate many turns of linac beam via charge-stripping of the H^- beam. The incoming beam from the linac will be chopped to allow clean injection into pre-existing RF buckets to form the desired number of bunches. Painting will be necessary in the 4-d transverse phase space and possibly also in the longitudinal. Very large transverse emittances must be created in order to control space-charge forces. For example, Figure 6 shows the Project X plan to paint three "squirts" from the linac into the acceptance of the Recycler to minimize the number of passes of the stripped protons through the stripping foil.

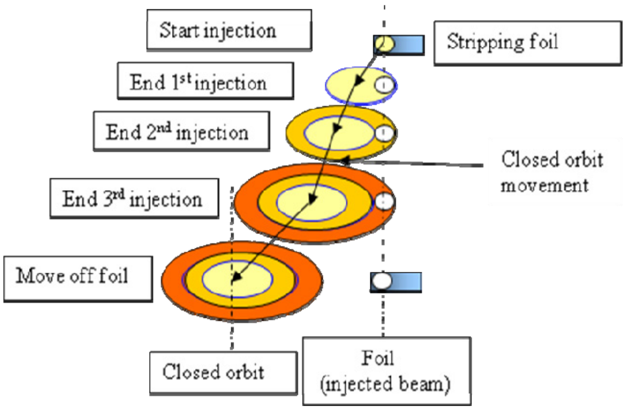

Fig. 6: Schematic of Project X painting scheme to strip many turns of H^- ions into the Recycler to minimize the number of passes of the protons through the stripping foil.

The Buncher

The second storage ring will be used to accept one or more bunches at a time from the Accumulator. Then a 90 degree bunch rotation in longitudinal phase space will be performed to shorten the bunches just before they are extracted. Of course, the momentum spread will become large, of order 5%, at that time, so the ring must have a large momentum acceptance. Also, the space-charge tune shift will be large when the beam is short.

The existing 8-GeV Fermilab Accumulator and Debuncher rings in the Antiproton Source are high-quality storage rings having the right energy and roughly the right circumferences. Furthermore, their apertures are large. They are, however, in a shallow tunnel, which probably obviates using them in their current location. They might, however, serve the purposes described here if they are relocated to a deeper tunnel.

The Combiner

The combiner is a set of transfer lines and kickers downstream of the rings that can allow more than one bunch to arrive simultaneously at the pion production target. The first major subsystem, the “trombone”, sends bunches on paths of different lengths. The second subsystem, the “funnel”, nestles the bunches side-by-side on convergent paths to the pion production target. The cartoon in Figure 7 illustrates the concept.

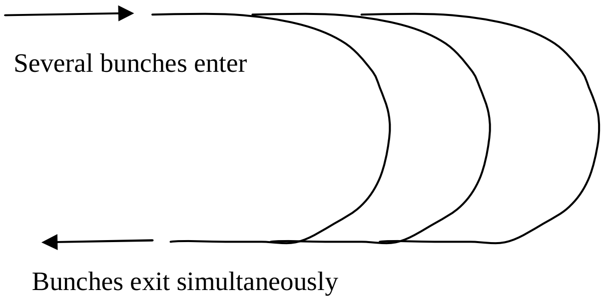

Fig. 7: The Combiner concept.

It is probable that any specific muon facility might not initially need all three of these major systems. If that is the case, then a phased approach would be possible. For example, facilities needed to drive a neutrino factory could be built first, followed by further elaboration for a muon collider. It would be important in that case to understand *a priori* how the facility might evolve.

Although these systems are conceptually simple, the performance requirements for the rings are beyond the present state of the art. Detailed design studies are clearly required. The studies should start by establishing first-order major parameters, layouts, and linear lattice designs for both rings. Then the performance limitations due to space charge, electron cloud, and other intensity-dependent effects should be addressed by theory and simulations. The Combiner systems also require careful design work. These studies should commence soon to determine viability, and if the problems are too daunting, then other alternatives will have to be considered.

OTHER ALTERNATIVES

Several alternative approaches for proton drivers for muon facilities have been suggested. The ISS [3] considered a variety of ideas to drive neutrino factories; those ideas will not be considered here. Instead, some ideas that have been suggested for implementation specifically at Fermilab will be discussed.

CW linac instead of pulsed linac

A group of people has been examining the possibility of using a CW superconducting linac instead of a pulsed one for this purpose. As the required beam power increases, this option looks more attractive. Since a paper has been submitted to the LINAC2008 conference [13], the discussion will not be repeated here.

Linac plus rapid-cycling synchrotron

Another idea is to use an 8-GeV rapid-cycling synchrotron instead of an 8-GeV linac. This suggestion seems to be motivated mainly by the expectation that it might be less expensive. However, an 8-GeV synchrotron for this purpose should have a high injection energy (~2 GeV) and a very large aperture to mitigate space-charge effects. Since the low-energy part of a linac is the most expensive part and since rapid-cycling, large-aperture synchrotrons are also very expensive, any cost savings would be modest at best.

Many subsystems in rapid-cycling synchrotrons are technically more difficult than the corresponding storage ring subsystems. Among these are the beam pipe, the RF systems, the magnets and their power supply. However, the biggest technical risk in these rings is excessive beam losses. In that regard, it is important to realize that there are beam losses in synchrotrons that do not occur in storage rings. Among these are the uncaptured beam that is lost at the start of the magnet ramp, losses if transition must be crossed, and various resonant conditions that occur at particular energies. The storage rings considered here have other advantages regarding the reduction of beam losses: the beam does not stay in the rings for very long, and beam collimation is easier and more effective in storage rings because the energy is fixed.

Finally, storage rings provide more flexibility regarding the formation of a variable number of bunches per cycle and the provision of a variable repetition rate to the target.

Use of Main Injector with Project X

Another idea is to use the Main Injector, with the Project X linac as the injector, to accelerate beyond 8 GeV. There are many variants on this idea.

In the original Project X concept, the Recycler is used to accumulate 3 linac “squirts”, which are then transferred to the Main Injector for further acceleration. In this way, the Main Injector is expected to deliver 2 MW of beam power at 120 GeV with a 1.4 sec cycle time. However, 120 GeV is far from the optimal value for production of the low energy pions that are most likely to produce useful muons, and 0.7 Hz is far from the ideal repetition rate, so the idea is to raise the repetition rate and reduce the energy, perhaps to 40 GeV. Less beam power is likely to be available with faster ramps going to lower energies; how much less power is a matter of current lively debate and would depend on what upgrades are implemented. According to Figure 2, the rate of production of useful muons per unit proton beam power is ~10% lower at 40 GeV compared to 8 GeV.

If the Project X linac is designed to deliver three times more protons per cycle than originally proposed, then the beam could be injected directly into the Main Injector, bypassing the Recycler. There has not been much attention paid to the possibility of accelerating even more protons per cycle in the Main Injector; the initial Project X concept already calls for more than a factor of five beyond the original Main Injector design intensity. Beam losses at injection and transition and beam instabilities might limit the intensity. So it is possible that the maximum beam power available from the Main Injector at ~40 GeV would be considerably less than 2 MW.

At any rate, the concept is to rebunch the beam in a storage ring at ~40 GeV, rotate or compress to short bunches in the same ring or another, and then extract to the target. Since this scenario can furnish fewer proton bunches per second but with more pions produced per bunch, it has been supported mostly by those who prefer the high emittance, low repetition rate muon collider scenario.

SUMMARY

Factors influencing the design choices for a proton driver based on Project X for muon storage ring facilities have been examined. For either the muon collider or the neutrino factory, 8-GeV beam power of at least 4 MW may be needed. The preferred system, consisting of two 8-GeV storage rings plus an external bunch-combiner, may be able to provide at least 4 MW of beam power with considerable flexibility in the way bunches are delivered to the target.

REFERENCES

[1] S. Nagaitsev, LINAC2008.

[2] Yu. Alexahin, C. Ankenbrandt, S. Geer, A. Jansson, D. Neuffer, M. Popovic, and V. Shiltsev, FERMILAB-TM-2418-AD-APC.

[3] J. S. Berg et al., “Accelerator design concept for future neutrino facilities”, http://www.hep.ph.ic.ac.uk/iss/reports/ISS-AcceleratorWG-final.pdf.

[4] Neutrino Factory Design Studies 1, 2, 2A, 2B are described in the previous ISS reference.

[5] H. Kirk et al., EPAC06.

[6] A. Jansson, Nufact08.

[7] R. B. Palmer, Low Emittance Collider Workshop, http://www.muonsinc.com/lemc2008/presentations/0804_LEMC_v1.pdf .

[8] R. P. Johnson, COOL07.

[9] Y. Derbenev and R. P. Johnson, Phys. Rev. Spec. Topics Accel. and Beams 8, 041002 (2005).

[10] Yaroslav Derbenev et al., COOL05.

[11] R. Fernow et al., https://mctf.fnal.gov/meetings/2008/02_21/fernow_lattice.ppt/view.

[12] P. Hanlet et al., EPAC06.

[13] C. Ankenbrandt et al., LINAC2008.