BNL-xxx-yyy
MC-zzz

A Cost-Effective Design for a Neutrino Factory

J.S. Berg,[1] S.A. Bogacz,[2] S. Caspi,[3] J. Cobb,[4] R.C. Fernow,[1] J.C. Gallardo,[1]
S. Kahn,[1] H. Kirk,[1] D. Neuffer,[5] R. Palmer,[1] K. Paul,[6] H. Witte,[4] and M. Zisman[3]

[1]*Brookhaven National Laboratory, Upton, NY 11973, USA*

[2]*Jefferson Laboratory, 12000 Jefferson Avenue, Newport News, VA 23606, USA*

[3]*Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*

[4]*Oxford University, Oxford OX1 3RH, UK*

[5]*Fermi National Accelerator Laboratory, Batavia, IL 60510, USA*

[6]*Muons Inc., Batavia, IL 60510, USA*

(Dated: April 7, 2005)

Abstract

There have been active efforts in the U.S., Europe and Japan on the design of a Neutrino Factory. This type of facility produces intense beams of neutrinos from the decay of muons in a high energy storage ring. In the U.S. a second, detailed Feasibility Study (FS2) [1] for a Neutrino Factory was completed in 2001. Since that report was published, new ideas in bunching, cooling and acceleration of muon beams have been developed. We have incorporated these ideas into a new facility design, which we designate as Study 2B (ST2B), that should lead to significant cost savings over the FS2 design.

I. INTRODUCTION

A Neutrino Factory [2, 3, 4] facility offers an exciting option for the long-term neutrino physics program. In the U.S. there has been a significant investment in developing the concepts and technologies required for such an accelerator complex.

New accelerator technologies offer the possibility of building, in the not-too-distant future, an accelerator complex to produce and capture more than 10^{20} muons per year [3]. It has been proposed to build a Neutrino Factory by accelerating the muons from this intense source to energies of several tens of GeV, injecting them into a storage ring having long straight sections, and exploiting the intense neutrino beams that are produced by muons decaying in the straight sections. The decays

$$\mu^- \to e^- \nu_\mu \bar{\nu}_e \; , \qquad \mu^+ \to e^+ \bar{\nu}_\mu \nu_e \tag{1}$$

offer exciting possibilities to pursue the study of neutrino oscillations and neutrino interactions with exquisite precision.

A Neutrino Factory requires an intense multi-GeV proton source capable of producing a primary proton beam with a beam power of 1–2 MW or more on target. This is the same proton source required in the medium term for Neutrino Superbeams; hence, there is a natural evolution from Superbeam experiments to Neutrino Factory experiments.

The physics case for a Neutrino Factory will depend upon results from the next round of planned neutrino oscillation experiments [5]. If the unknown mixing angle θ_{13} is small, such that $\sin^2 2\theta_{13} < O(10^{-2})$, or if there is a surprise and three-flavor mixing does not completely describe the observed phenomenology, then answers to some or all of the most important neutrino oscillation questions will require a Neutrino Factory. If $\sin^2 2\theta_{13}$ is large, just below the present upper limit, and if there are no experimental surprises, the physics case for a Neutrino Factory will depend on the values of the oscillation parameters, the achievable sensitivity that will be demonstrated by the first generation of ν_e appearance experiments, and the nature of the second generation of basic physics questions that will emerge from the first round of results. In either case (large or small θ_{13}), in about a decade the neutrino community may need to insert a Neutrino Factory into the global neutrino plan. The option to do this in the next 10 years will depend upon the accelerator R&D that is done during the intervening period.

In the U.S., the *Neutrino Factory and Muon Collider Collaboration* (referred to herein as the Muon Collaboration, or MC) [6] is a collaboration of 130 scientists and engineers engaged in carrying out the accelerator R&D that is needed before a Neutrino Factory could be inserted into the global plan. Much technical progress has been made over the last few years, and the required key accelerator experiments are now in the process of being proposed and approved. In addition to the U.S. effort, there are active Neutrino Factory R&D groups in Europe [7, 8] and Japan [9], and much of the R&D is performed and organized as an international endeavor. Thus, because a Neutrino Factory is potentially the key facility for the long-term neutrino program, Neutrino Factory R&D is an important part of the *present* global neutrino program. Indeed, the key R&D experiments are seeking funding now, and will need to be supported if Neutrino Factories are to be an option for the future.

In this article we describe an updated Neutrino Factory design that demonstrates significant progress toward cost reduction for this ambitious facility. The paper is organized as follows. Section II describes in some detail the Neutrino Factory design concept. The design of the front end of the facility is described in Section III, the design of the accelerator chain in described in Section IV. In Section V we discuss in some details the assumptions and algorith, used to make estimate of the cost of a Neutrino Factory and finally, some open design issues are mentioned in Section VI.

Much of the work described in this paper was performed as part of the year-long Study of the Physics of Neutrinos, organized by the American Physical Society [5].

II. MACHINE CONCEPT

In this Section we describe the basic machine concepts that are used to create a Neutrino Factory facility [1, 5, 10]. This facility is a *secondary beam* machine; that is, a production beam is used to create the secondary beam that eventually provides the neutrino flux for the detector. For a Neutrino Factory, the production beam is a high intensity proton beam of moderate energy (beams of 2–50 GeV have been considered by various groups) that impinges on a target, typically a high-Z material (e.g. Hg). The collisions between the proton beam and the target nuclei produce a secondary pion beam that quickly decays (26.0 ns) into a longer-lived (2.2 μs) muon beam. The remainder of the Neutrino Factory is used to condition the muon beam (see Section III), accelerate it rapidly to the desired final energy of a few tens of GeV (see Section IV), and store it in a decay ring having a long straight section oriented such that decay neutrinos produced there will hit a detector located thousands of kilometers from the source.

Our present concept of a Neutrino Factory is based in part on the most recent Feasibility Study (Study-II, referred to herein as FS2) [1] that was carried out jointly by BNL and the U.S. *Neutrino Factory and Muon Collider Collaboration*. Two Feasibility Studies [1, 10] have demonstrated technical feasibility (provided the challenging component specifications are met), established a cost baseline, and established the expected range of physics performance. It is worth noting that the Neutrino Factory design we envision could fit comfortably on the site of an existing laboratory, such as BNL or FNAL.

The main ingredients of a Neutrino Factory include:

- **Proton Driver:** 1–4 MW of protons on target from, for example, an upgraded AGS; a new booster at Fermilab (or elsewhere) would perform equivalently.

- **Target and Capture:** A high-power target immersed in a 20 T superconducting solenoidal field to capture pions produced in proton-nucleus interactions. The high magnetic field at the target is smoothly tapered down to a much lower value, 1.75 T, which is then maintained through the bunching and phase rotation sections of the Neutrino Factory.

- **Bunching and Phase Rotation:** We first accomplish the bunching with rf cavities of modest gradient, whose frequencies change as we proceed down the beam line. After bunching the beam, another set of rf cavities, with higher gradients and again having decreasing frequencies as we proceed down the beam line, is used to rotate the beam in longitudinal phase space to reduce its energy spread.

- **Cooling:** A solenoidal focusing channel, with high-gradient 201.25 MHz rf cavities and LiH absorbers, cools the transverse normalized rms emittance from 17 mm·rad to about 7 mm·rad. This takes place at a central muon momentum of 220 MeV/c.

- **Acceleration:** A superconducting linac with solenoidal focusing is used to raise the muon beam energy to 1.5 GeV, followed by a Recirculating Linear Accelerator (RLA), arranged in a *dogbone* geometry, to provide a 5 GeV muon beam. Thereafter, a pair of cascaded Fixed-Field, Alternating Gradient (FFAG) rings, using a doublet lattice of combined-function magnets, is used to reach 20 GeV. Additional FFAG stages could be added to reach a higher beam energy, if the physics requires this.

- **Storage Ring:** We employ a compact racetrack-shaped superconducting storage ring in which $\approx 35\%$ of the stored muons decay toward a detector located some 3000 km from the ring. Muons survive for roughly 500 turns.

In the remainder of this paper we describe in detail the new design of the Neutrino Factory front-end for performing the required beam manipulations prior to acceleration and describe our new ideas for accelerating the muon beam using FFAGs.

III. FRONT END DESIGN

The front end of the neutrino factory (the part of the facility between the target and the first linear accelerator) represented a large fraction, about 40 %, of the total facility costs in FS2 [1]. However, several recent developments lead to the idea that a new design for the front end may be possible that is significantly less expensive:

- A new approach to bunching and phase rotation using the concept of adiabatic rf bunching [11, 12, 13, 14, 15] eliminates the very expensive induction linacs used in FS2.

- For a moderate cost, the transverse acceptance of the accelerator chain could be doubled from its FS2 value.

- This diminished the demands on the transverse ionization cooling section and allowed the design of a simplified cooling section with fewer components and reduced magnetic field strength.

We denote as *Study 2B* (ST2B) the simulations that have been made of the performance of this new front end, together with the new scheme for acceleration. The Monte Carlo simulations were performed with the code ICOOL [16]. The concept of the adiabatic buncher is compared with the system used in FS2 in Fig. 1. The longitudinal phase space after the target is the same in both cases. Initially, there is a small spread in time, but a very large spread in energy. The target is followed by a drift space in both cases, where a strong

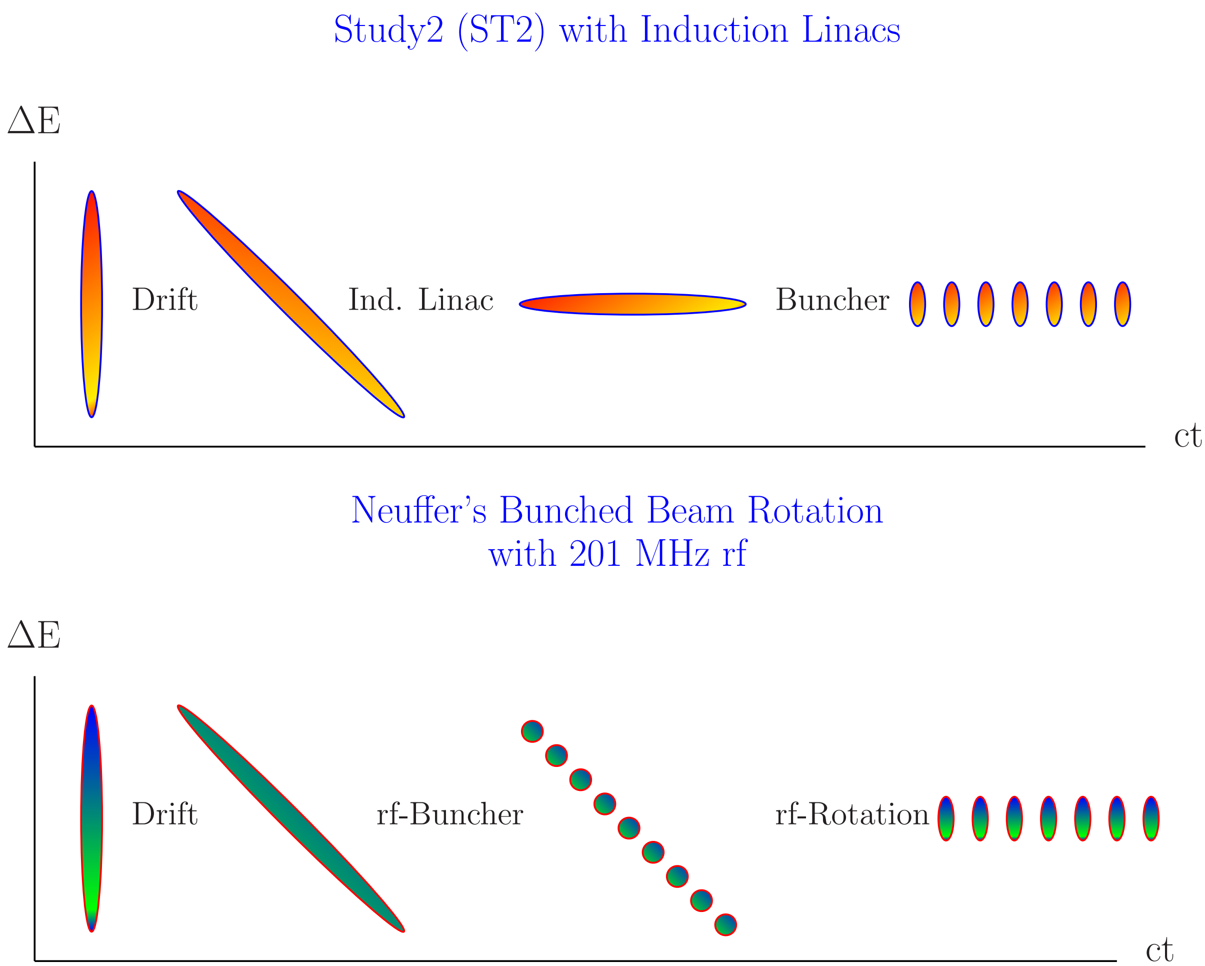

FIG. 1: (Color) Comparison of the buncher concept used here (bottom) with the bunching system used in FS2 (top).

correlation develops between time and energy. In FS2, the energy spread in the correlated beam was first flattened using a series of induction linacs. The induction linacs did an excellent job, reducing the final rms energy spread to 4.4%. The beam was then sent through a series of rf cavities for bunching, which increased the energy spread to ≈ 8%. In the new scheme, the correlated beam is first adiabatically bunched using a series of rf cavities with decreasing frequencies and increasing gradients. The beam is then phase rotated with a second string of rf cavities with decreasing frequencies and constant gradient. The final rms energy spread in the new design is 10.5%. This spread is acceptable for the new cooling channel. The overall layout of the new front-end design is shown schematically in Fig. 2.

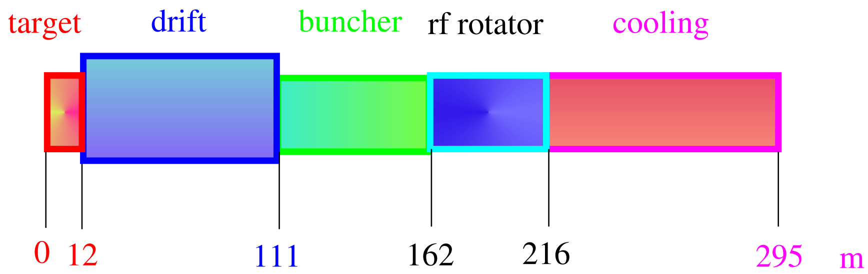

FIG. 2: (Color) Overall layout of the front-end.

A. Target and Decay Region

The first ≈12 m is used to capture pions produced in the target. The field here drops adiabatically from 20 T over the target to 1.75 T. At the same time, the radius of the beam pipe increases from 7.5 cm at the target up to 25 cm. Next comes ≈100 m for the pions to decay into muons and for the energy-time correlation to develop. The adiabatic bunching occupies the next ≈50 m and the phase rotation takes place in ≈50 m following that. Lastly, the channel has ≈80 m of ionization cooling. The total length of the new front end is 295 m. The longitudinal component of the magnetic field on-axis is shown for the full front-end in

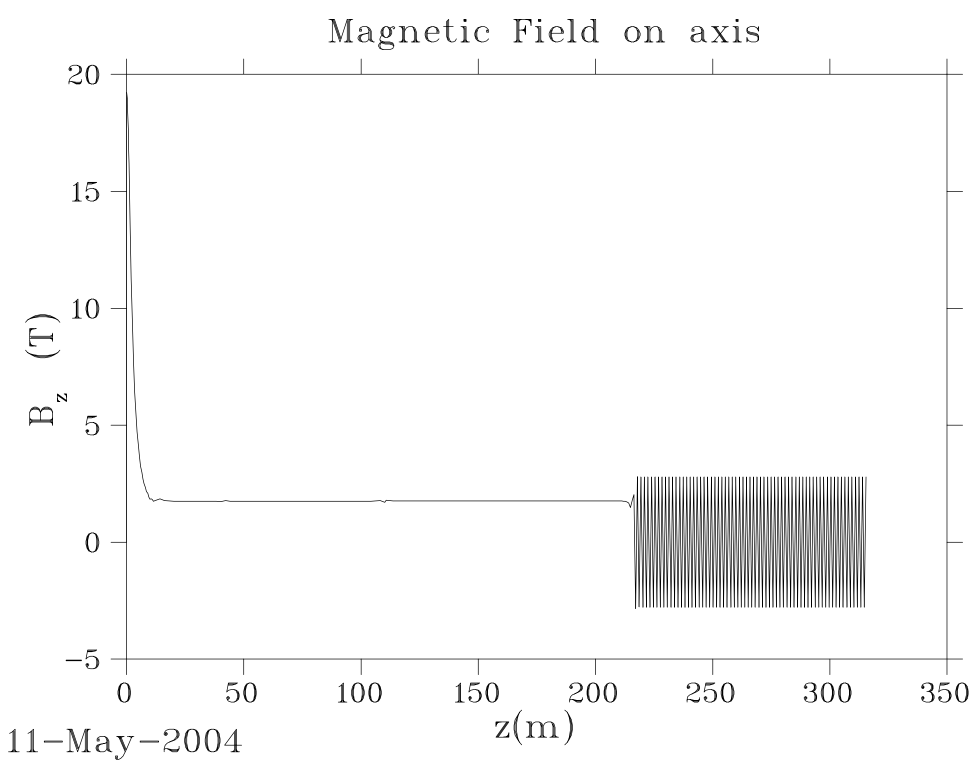

FIG. 3: Longitudinal field component B_z on-axis along the ST2B front-end.

Fig. 3. The field falls very rapidly in the collection region to a value of 1.75 T. It keeps this value with very little ripple over the decay, buncher and rotator regions. After a short matching section, the 1.75 T field is changed adiabatically to the alternating field used in the cooling section.

The beam distributions used in the simulations were generated using MARS [17]. The distribution is calculated for a 24 GeV proton beam interacting with a Hg jet [18]. The jet is incident at an angle of 100 mrad to the solenoid axis, whereas the beam is incident at an angle of 67 mrad to the solenoid axis. An independent study showed that the resulting 33 mrad crossing angle gives near-peak acceptance for the produced pions. An examination of the distribution of particles that were propagated to the end of the front-end channel shows that they have a peak initial longitudinal momentum of ≈300 MeV/c with a long high-energy tail, and a peak initial transverse momentum ≈180 MeV/c.

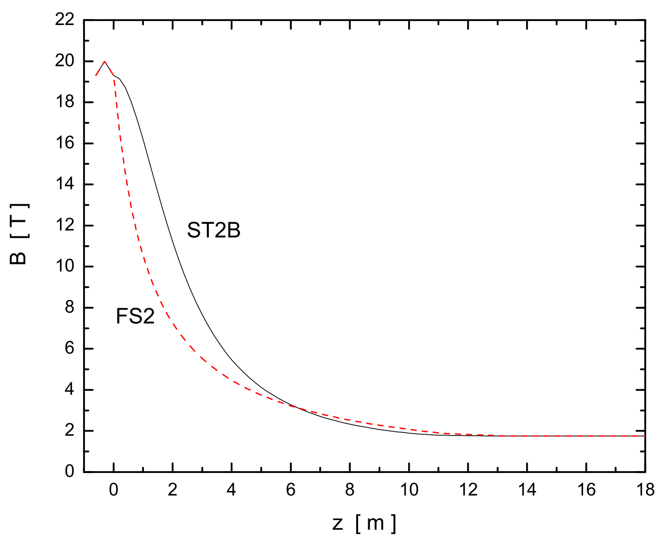

FIG. 4: (Color) Comparison of the capture region magnetic field used in the present simulation (ST2B) with that used in FS2.

We used an improved axial field profile in the capture region that increased the final number of muons per proton in the accelerator acceptance by ≈10%. The new axial field profile (marked ST2B) is compared in Fig. 4 with the profile used in FS2. Figure 5 shows the actual coil configuration in the collection region. The end of the 60 cm long target region is defined as $z = 0$. The three small-radius coils near $z = 0$ are Cu coils, while the others are superconducting. The left axis shows the error field on-axis compared with the desired field profile. We see that the peak error field is ≈ 0.07 T.

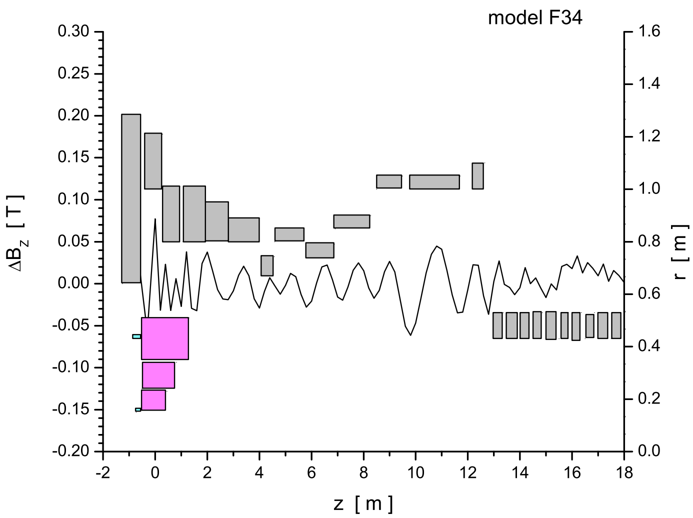

FIG. 5: (Color) Actual coil configuration in the collection region. The left axis shows the error field on-axis compared with the optimal capture field profile, denoted ST2B in Fig. 4.

Figure 6 shows a MARS calculation of the absorbed radiation dose in the collection region. The peak energy deposition dose in the superconducting coils, as illustrated in Fig. 6, is ≈ 1 Mgy/yr for a 1 MW beam running for a Snowmass year of 2×10^7 s. Assuming a lifetime dose for the insulation of 100 Mgy, there should be no problem with radiation damage in the coils.

B. Bunching and Phase Rotation Region

Two cells of the buncher lattice are shown schematically in Fig. 7. Most of the 75 cm cell length is occupied by the 50 cm long rf cavity. The cavity iris is covered with a Be window. The limiting radial aperture in the cell is determined by the 25 cm radius of the window. The 50 cm long solenoid was placed outside the rf cavity in order to decrease the magnetic field ripple on the axis and minimize beam losses from momentum stop bands. The buncher section contains 27 cavities with 13 discrete frequencies and gradients varying from 5–10 MV/m.

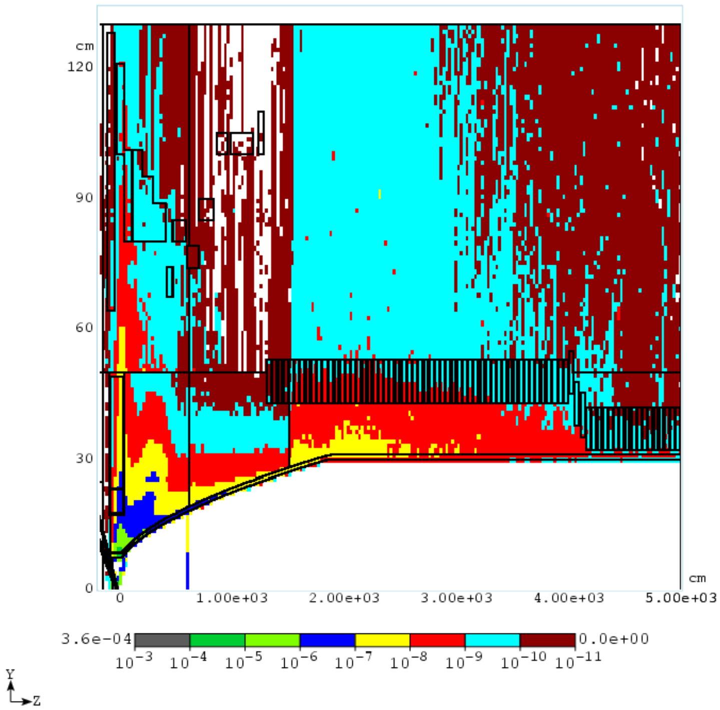

FIG. 6: (Color) MARS calculation of the absorbed annual radiation dose in the collection region.

The frequencies decrease from 333 to 234 MHz in the buncher region. The cavities are not equally spaced. Fewer cavities are used at the beginning where the required gradients are small. Figure 8 shows the correlated longitudinal phase space and the bunching produced by the buncher.

The rotator cell is very similar to the buncher cell. The major difference is the use of tapered Be windows on the cavities because of the higher rf gradient. There are 72 cavities in the rotator region, with 15 different frequencies. The frequencies decrease from 232 to 201 MHz in this part of the front end. All cavities have a gradient of 12.5 MV/m. The energy spread rms in the beam is reduced from 46 to 27 MeV.

C. Cooling Region

The cooling channel was designed to have a relatively constant with position transverse beta function with a magnitude of about 80 cm. One cell of the channel is shown in Fig. 9.

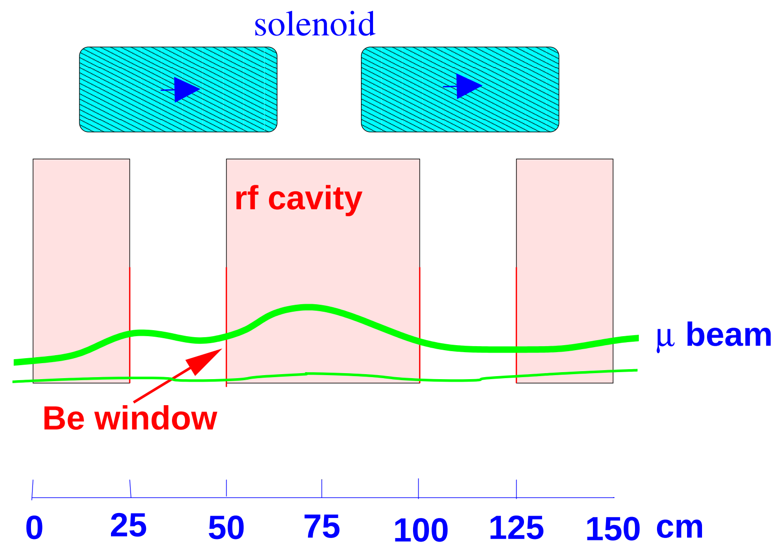

FIG. 7: (Color) Schematic of two cells of the buncher or phase rotator section.

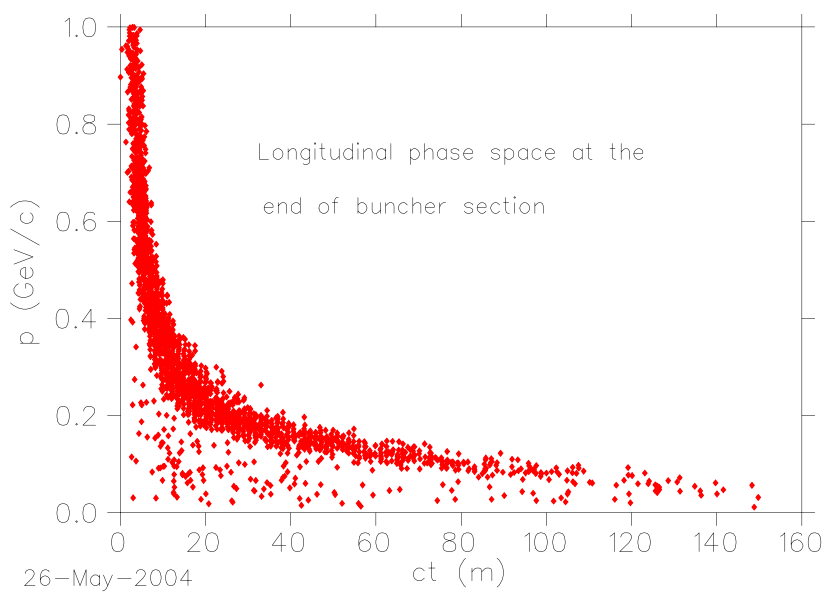

FIG. 8: (Color) Longitudinal phase space after the buncher section.

Most of the 150 cm cell length is taken up by the 50 cm long rf cavities. The cavities have a frequency of 201.25 MHz and a gradient of 15.25 MV/m. A novel aspect of this design comes from using the windows on the rf cavity as the cooling absorbers. This is possible because the near constant β function does not significantly increase the emittance heating at the window location. The window consists of a 1 cm thickness of LiH with a 75 μm thick layer of Be on the side facing the rf cavity field and a 25 μm thick layer of Be on the opposite side (The Be will, in turn, have a thin coating of TiN to prevent multipactoring [19].) The

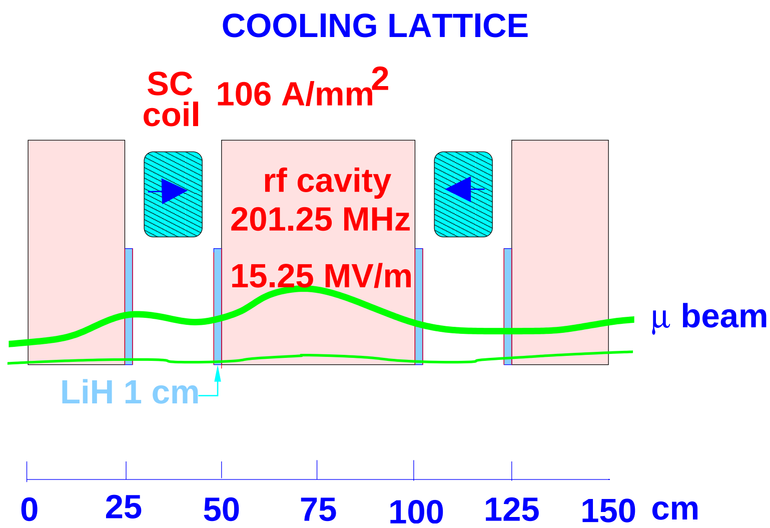

FIG. 9: (Color) Schematic of one cell of the cooling section.

alternating 2.8 T solenoidal field is produced with one solenoid per half cell, located between the rf cavities.

Figure 10 shows the longitudinal phase space at the end of the cooling section. The reduction in normalized transverse emittance along the cooling channel is shown in the left plot of Fig. 11 and the right plot shows the normalized longitudinal emittance. The channel

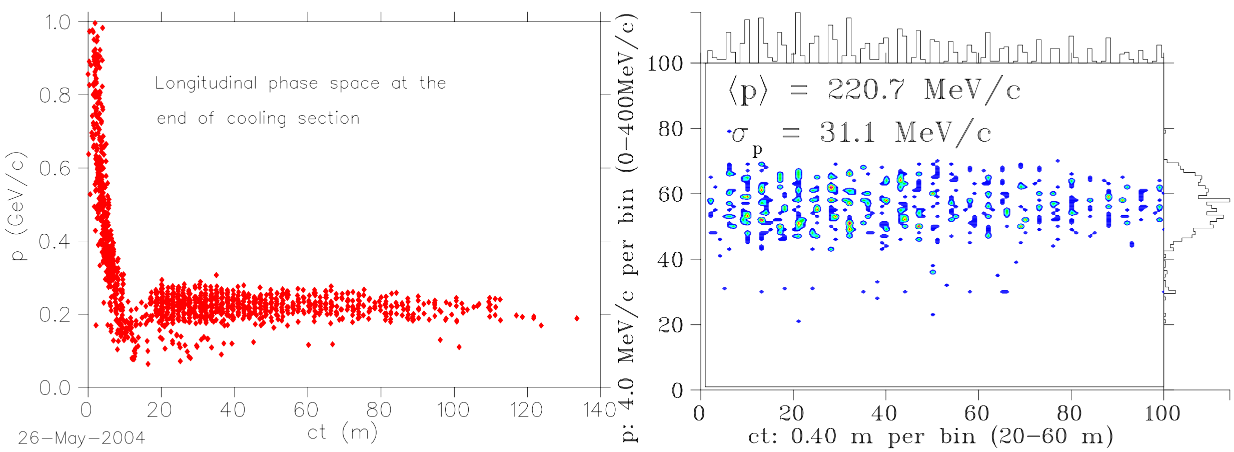

FIG. 10: (Color) Longitudinal phase space at the end of the channel.

produces a final value of $\epsilon_T = 7.1$ mm rad, that is, more than a factor of two reduction from the initial value. The equilibrium value for a LiH absorber with an 80 cm β function is about $\epsilon_T^{\text{equil.}} \approx 5.5$ mm rad. Figure 12 shows the muons per incident proton on target that fit into the accelerator transverse normalized acceptance of $A_T = 30$ mm rad and normalized

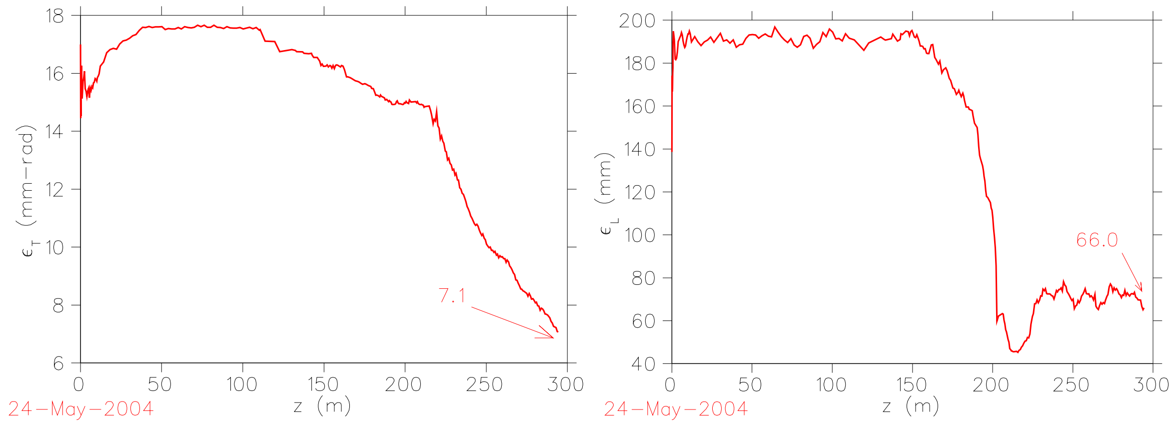

FIG. 11: (Color) Normalized transverse emittance (left) and longitudinal emittance (right) along the front-end for a momentum cut $0.1 \leq p \leq 0.3$ GeV/c.

longitudinal acceptance of $A_L = 150$ mm. The 80-m-long cooling channel raises this quantity by about a factor of 1.7. The current best value is 0.170 ± 0.006 muons per incident proton. This is the same value obtained in FS2. Thus, we have achieved the identical performance at the entrance to the accelerator as FS2, but with a significantly simpler, shorter, and presumably less expensive channel design. In addition, unlike FS2, this channel transmits both signs of muons produced at the target. With appropriate modifications to the transport line going into the storage ring, this design could deliver both (time tagged) neutrinos and antineutrinos to the detector. The beam at the end of the cooling section consists of a train of bunches with a varying population of muons in each one; this is shown in Fig. 13 (one sign).

Figure 14 depicts the longitudinal phase space of a superposition of all bunches projected onto a single period ($T \approx 5$ ns) and Fig. 15 shows a few interleaved μ^+ and μ^- bunches exiting the cooling section.

D. Heating of Absorber Windows

There are some unresolved issues with the absorber windows that will require further R&D. To minimize multiple scattering we have assumed the windows are made from LiH. In order to protect the LiH from the environment and to provide a high conductivity surface to close off the rf cavity, we have assumed the LiH is encased in a thin layer of Be. Assuming that the Be can be deposited or bonded to the LiH, there is the question of what happens

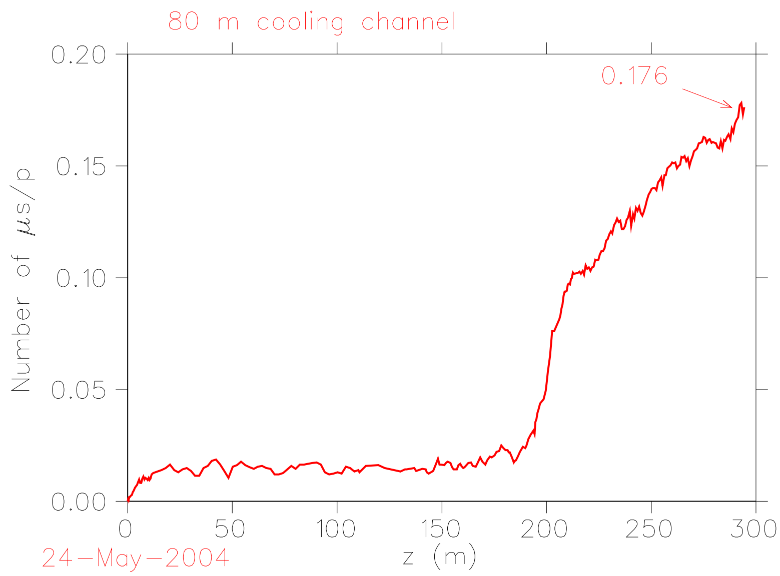

FIG. 12: (Color) The muons per incident proton on target into the accelerator transverse normalized acceptance of $A_T = 30$ mm rad and normalized longitudinal acceptance of $A_L = 150$ mm for a momentum cut $0.1 \leq p \leq 0.3$ GeV/c.

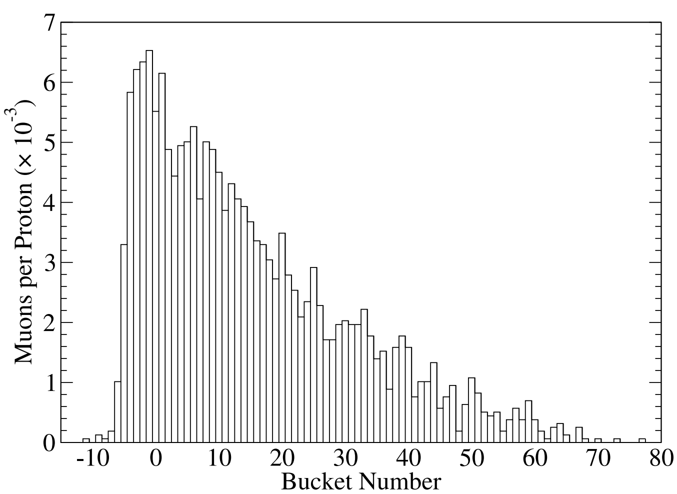

FIG. 13: Bunch structure of the beam delivered to the accelerator transverse normalized acceptance of $A_T = 30$ mm rad and normalized longitudinal acceptance of $A_L = 150$ mm for a momentum cut $0.1 \leq p \leq 0.3$ GeV/c.

when the window is heated by energy loss of the muon beam and by the power deposited by the rf cavity. If the heating becomes high enough, melting and differential stresses leading to buckling are possible. In addition the window could suffer degradation from radiation

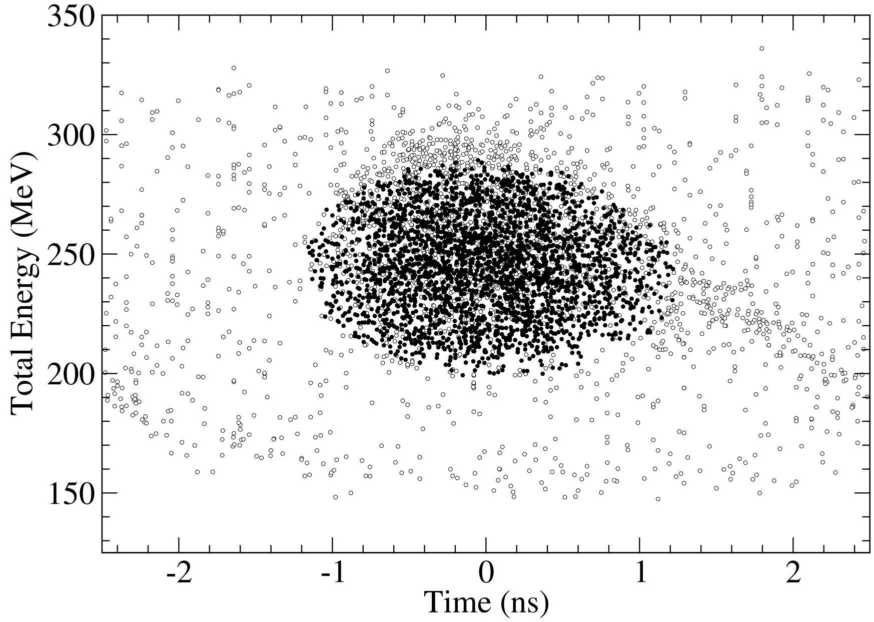

FIG. 14: Longitudinal phase space of overlayed bunches in the train at the end of the cooling section. The open circles are all the particles that reach the end of the channel and the filled circles are particles within the accelerator transverse normalized acceptance of $A_T = 30$ mm rad and normalized longitudinal acceptance of $A_L = 150$ mm for a momentum cut $0.1 \leq p \leq 0.3$ GeV/c.

damage.

Approximately 1.1×10^{14} muons of each charge enter the start of the cooling channel each second. This produces a total power deposition of ≈ 58 W distributed along the beam path. Most of the energy deposition takes place in the LiH. We assume that cooling is provided by a heat sink at the outer edge of the window. If we ignore any longitudinal heat conduction from the LiH to the Be layer, the LiH reaches a maximum temperature of 270°C in steady state. This is safely below its melting temperature of 690°C. The rf heating occurs in a skin depth on the side of the window facing the cavity. The skin depth for Be at 201 MHz is approximately 9 μm. The rf power deposited on the window of a pillbox rf cavity is

$$P = \frac{\pi^2}{2} \frac{d}{\lambda} E_0^2 \frac{b^2}{Z_0} \left\{ J_1^2(\alpha) - J_0(\alpha) J_2(\alpha) \right\} \tag{2}$$

where, $d =$skin depth, $\lambda =$rf wavelength, $E_0 =$peak rf gradient, $b =$window radius, $a =$ radius of rf cavity (pillbox), $Z_0 =$impedance of free space, and J_0, J_1, J_2 are Bessel functions with argument $\alpha = 2.405 \times \frac{b}{a}$. This gives a total rf power of ≈ 220 W in each window. Rough calculations predict that the temperature at the center of the 75μm thick Be layer should

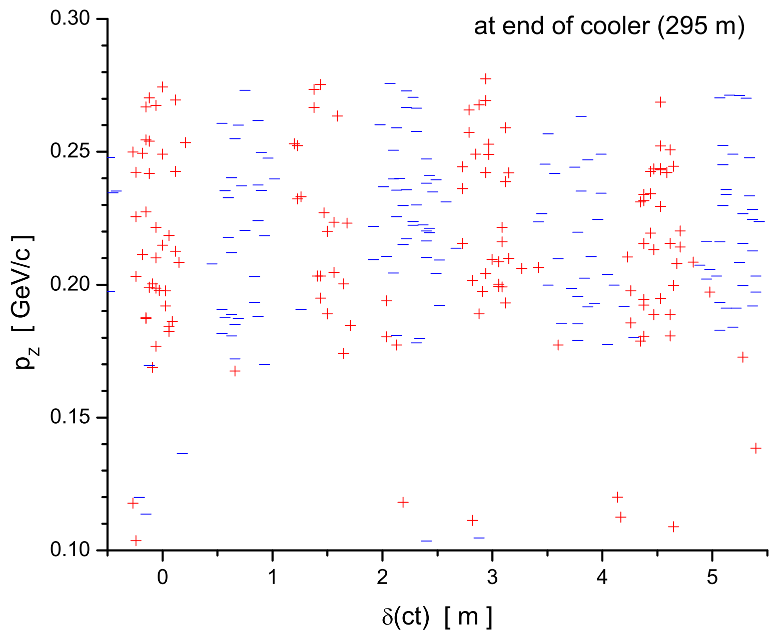

FIG. 15: (Color) A sample of the train of interleaved μ^+ (red) and μ^- (blue) bunches exiting the cooling section.

be less than 700 °C. This is safely below its melting temperature of 1275°C.

Although melting will not be a problem, buckling and delamination of the Be layer is a potential deleterious outcome. More accurate finite element thermal studies need to be done of the composite LiH-Be system. If this window design proves to be no feasible, a number of alternative window designs have been considered, as discussed in Section VI.

IV. ACCELERATION DESIGN

The acceleration system takes the beam from the end of the cooling channel and accelerates it to the energy required for the decay ring. Table I gives the design parameters of the acceleration system. Acceptance is defined such that if $A_\perp$ is the transverse acceptance and $\beta_\perp$ is the beta function, then the maximum particle displacement (of the particles we transmit) from the reference orbit is $\sqrt{\beta_\perp A_\perp mc/p}$, where p is the particle's total momentum, m is the particle's rest mass, and c is the speed of light. The acceleration system is able to accelerate bunch trains of both signs simultaneously.

TABLE I: Acceleration system design parameters.

Injection momentum (MeV/c)	273
Initial kinetic energy (MeV)	187
Final total energy (GeV)	20
Normalized transverse acceptance (mm)	30
Normalized emittance, rms (mm-rad)	3.84
Normalized longitudinal acceptance, $\Delta p L_b / m_\mu c$ (mm)	150
Total energy spread, ΔE (MeV)	± 45.8
Total time-of-flight (ns)	± 1.16
Energy spread, rms (MeV)	19.8
Time-of-flight, rms (ns)	0.501
Bunching frequency (MHz)	201.25
Maximum muons per bunch	1.1×10^{11}
Muons per bunch train per sign	3.0×10^{12}
Bunches in train	89
Average repetition rate (Hz)	15
Minimum time between pulses (ms)	20
Average beam power (per charge) (kW)	144

To reduce costs, the RLA acceleration systems from FS2 [1] will be replaced, as much as possible, by Fixed-Field Alternating Gradient (FFAG) accelerators.

A. Matching from Cooling to Acceleration Linac

As indicated in Table I, we inject into the linac at a momentum of 273 MeV/c, which is still higher than the average momentum in the cooling channel. We deal with this by designing a matching section from the cooling channel to the linac in which sufficient acceleration will occur to reach the required momentum for the linac. That matching section will consist of some cavities similar to those in the cooling channel, but with thinner windows, and some superconducting cells similar to the accelerating linac. Figure 16 shows a layout of the matching section. The current design for the matching section has about 15% loss; initial

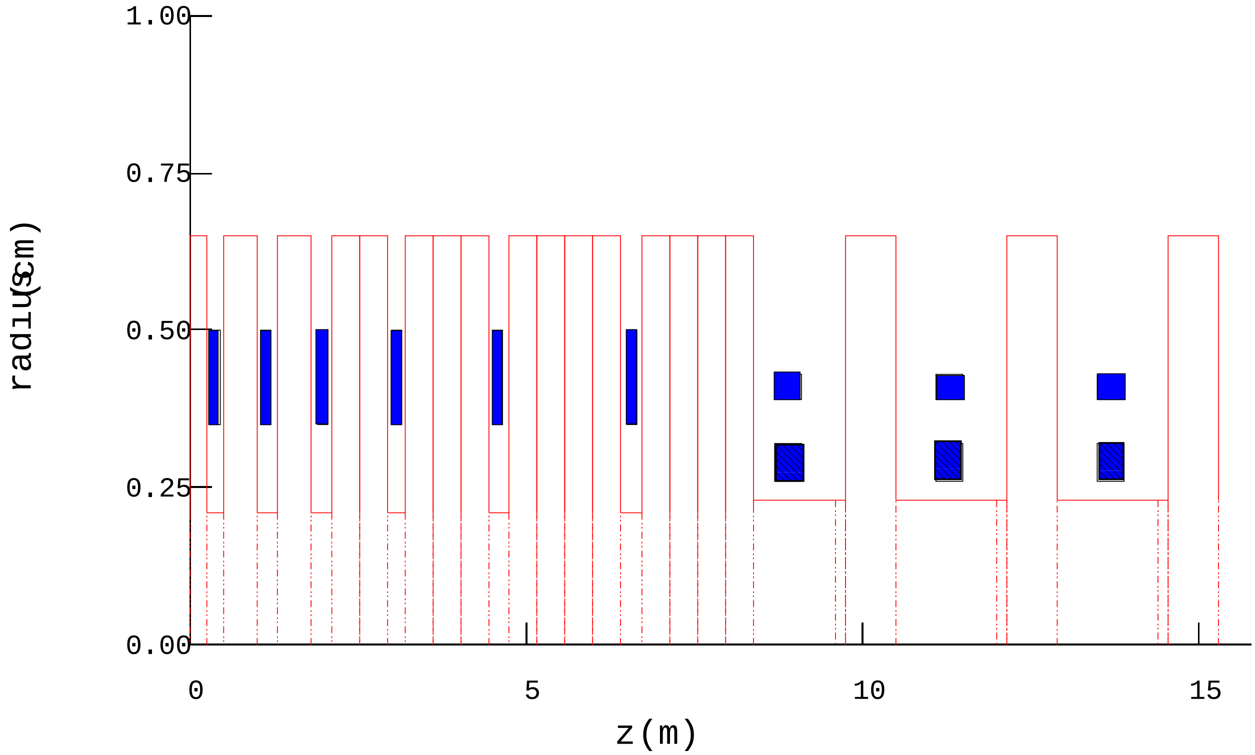

FIG. 16: Matching section from cooling to linac. Black squares are solenoids, red lines are rf cavity walls. The last three cavities are superconducting, the remaining ones are room temperature.

studies indicate that this may be due to performing the matching at low instead of high amplitudes. Initial attempts at performing the longitudinal match at high amplitudes have eliminated the losses, but we have not yet done the matching for the transverse plane as well.

B. Low Energy Acceleration

Based on cost considerations, we have chosen not to use FFAGs below 5 GeV total energy. Therefore, we must provide alternative acceleration up to that point. Similarly to what was adopted in FS2, we use a linac from the lowest energies to 1.5 GeV, followed by a recirculating linear accelerator (RLA).

The linac turns out to be strongly constrained by the transverse acceptance. In FS2, there were three types of cryomodules, containing one, two, and four two-cell cavities, respectively. With our larger acceptance, the cryomodules from FS2 would require the beam to have a momentum of at least 420 MeV/c, 672 MeV/c, and 1783 MeV/c, respectively. Note that

TABLE II: Linac cryomodule structure. Numbers are lengths in m.

Cryostat I		Cryostat II		Cryostat III	
Drift	0.45	Drift	0.70	Drift	0.70
Solenoid	1.00	Solenoid	1.00	Solenoid	1.00
Drift	0.50	Drift	1.00	Drift	1.00
Cavity	0.75	Cavity	1.50	Cavity	1.50
Drift	0.30	Drift	0.80	Drift	1.50
Total	3.00	Total	5.00	Cavity	1.50
				Drift	0.80
				Total	8.00

the lowest momentum is much higher than the average momentum in the cooling channel, which is about 220 MeV/c. Thus, we need to make adjustments to the FS2 design to be able to accelerate this larger beam.

In particular, to increase the acceptance, we must reduce the lengths of the cryomodules. We first construct a very short cryomodule by using a single one-cell cavity instead of the two-cell cavities in the FS2 cryomodules. Not only does this shorten the cavity itself, it also eliminates one of the input couplers. We also eliminate some of the drift space in the cryomodule. This is possible since we intend to run the cavities with up to 0.1 T on them [20]; this is acceptable provided the cavities are cooled down before the magnets are powered. The field profile of the solenoids shown in FS2 indicates that the iron shield on the solenoids is sufficient to bring the field down to that level even immediately adjacent to the solenoid shield. Together, these changes permit a total length for the first module type of 3 m. Table II shows the dimensions of the cryostats we will use and Fig. 17 depicts all three of them.

Table III summarizes parameters for the linac. The phase of the cavities in the linac will be varied approximately linearly with length from about -73° at the beginning of the linac to 0° at the end, as shown in Fig. 18.

TABLE III: Linac cryomodule parameters.

	Cryo I	Cryo II	Cryo III
Length (m)	3.00	5.00	8.00
Number of modules	12	18	22
Cells per cavity	1	2	2
Cavities per module	1	1	2
Maximum energy gain per cavity (MeV)	11.2	22.4	22.4
Cavity rf frequency (MHz)	201.25	201.25	201.25
Solenoid length (m)	1	1	1
Max Solenoid field (T)	1.5	1.8	4.0

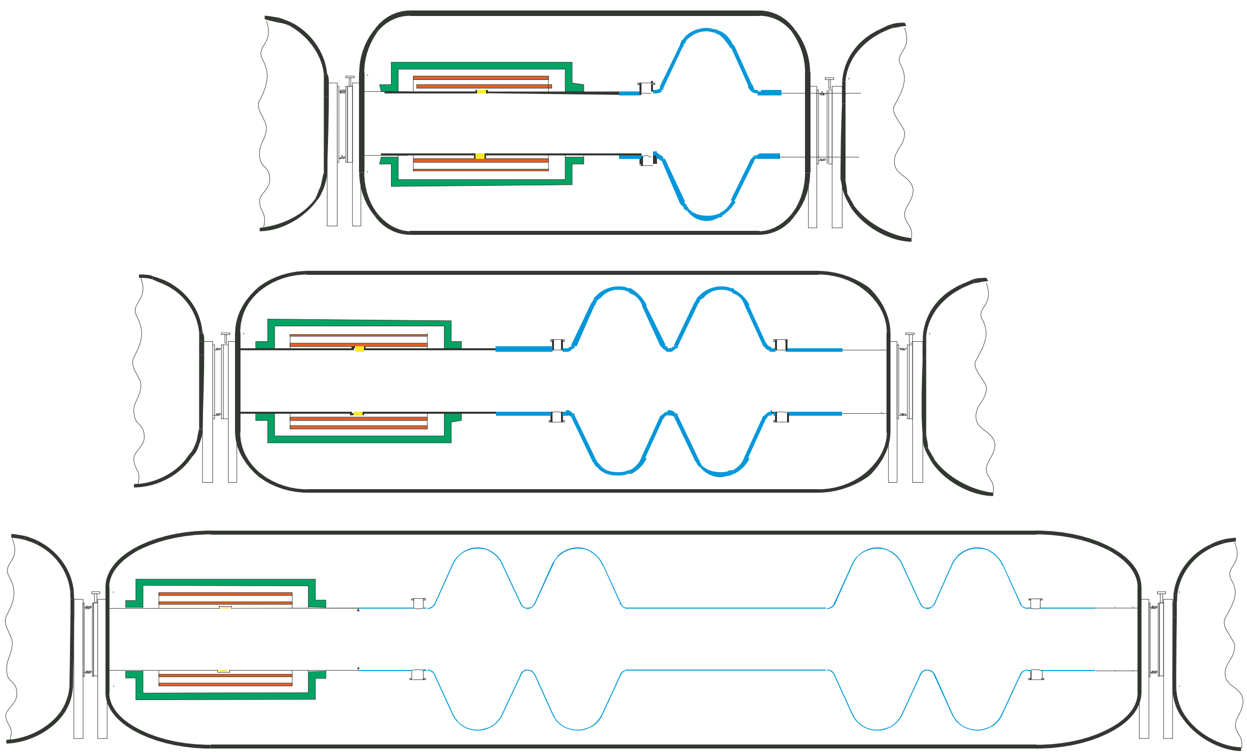

FIG. 17: (Color) Layouts of superconducting linac pre-accelerator cryomodules. Blue lines are the SC walls of the cavities and solenoid coils are indicated in red with the iron shielding in green. The dimensions of the cryomodules are shown in Table II, and Table III summarizes parameters for the linac.

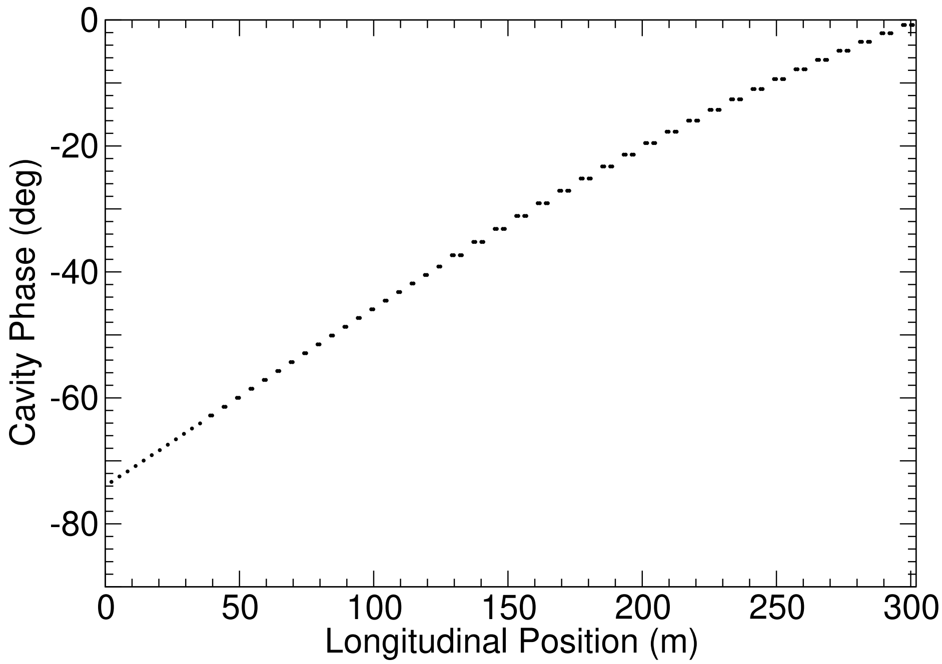

FIG. 18: Cavity phase vs. position along the linac.

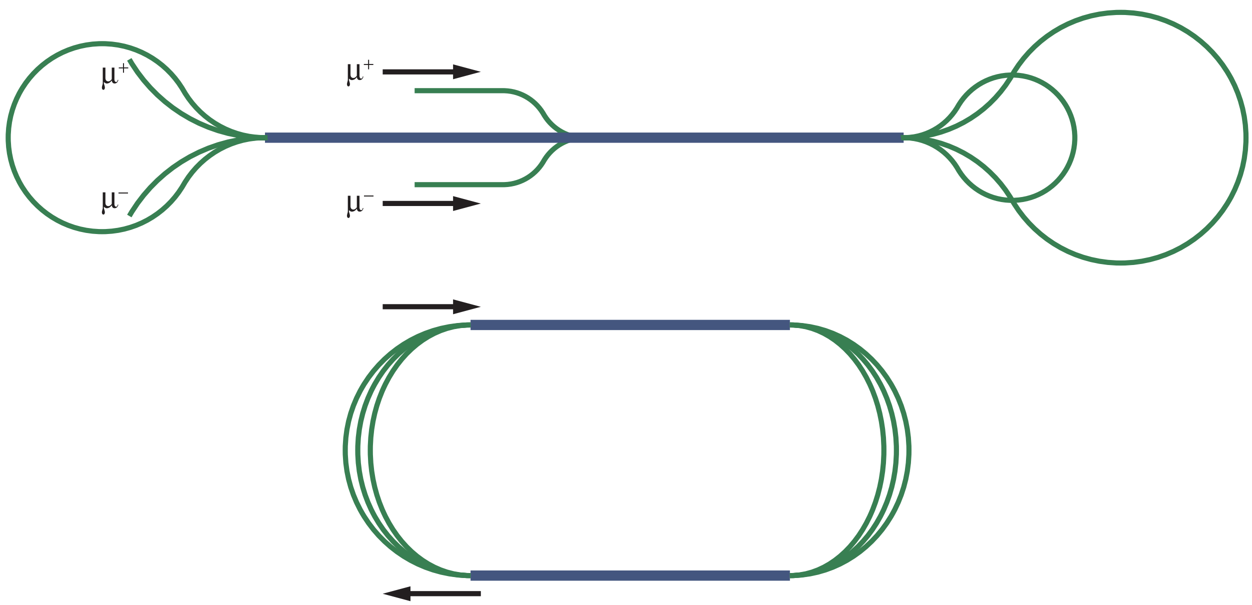

FIG. 19: (Color) Dogbone (top) and racetrack (bottom) layout for the RLA.

C. RLAs

Compared with FS2, we are injecting into the RLA at a lower energy and are accelerating over a much smaller energy range. This will make it more difficult to have a large number of turns in the RLA. To mitigate this, we choose a dogbone layout for the RLA [21]. For

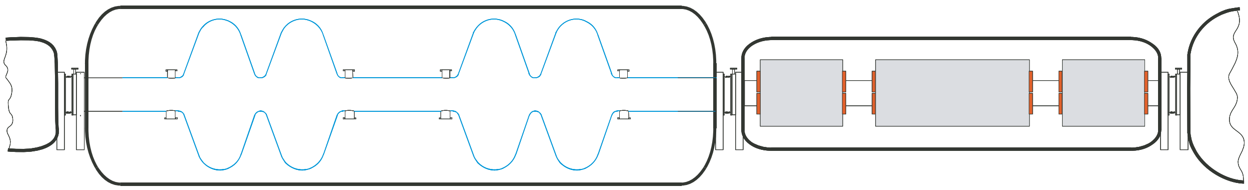

FIG. 20: (Color) Dogbone linac cell.

a given amount of installed rf, the dogbone layout has twice the energy separation of the racetrack layout at the spreaders and recombiners, making the switchyard much easier and allowing more passes through the linac.

One disadvantage of the dogbone layout is that, because of the longer linac and the very low injection energy, there is a significant phase shift of the reference particle with respect to the cavity phases along the length of the linac in the first pass (or the last pass, depending on which energy the cavities are phased for). To reduce this effect, we inject into the center of the linac (as shown in Fig. 19).

In the dogbone RLA, we have just over 1 GeV of linac, and we make three and a half passes through that linac to accelerate from a total energy of 1.5 GeV to 5 GeV. The linac with consist of 11 m long cells with two 2-cell cavities per cell, and quadrupole triplet focusing, as shown in Fig. 20. The cavities are the 30 cm aperture cavities assumed in FS2, as opposed to the 46 cm aperture cavities that were used in the linac that accelerated up to 1.5 GeV; this should allow a larger gradient (17 MV/m rather than 15 MV/m).

The arcs will also use quadrupole triplet focusing, with a 90° phase advance per cell in both planes, in order to cancel some chromatic effects. Both the quadrupoles and the dipoles in the arc and linac lattices will have 1 T maximum field at the coils and be superconducting.

Since the dogbone arc changes its direction of bend twice in each arc, dispersion matching must be handled carefully. This is done by having a 90° phase advance per cell, and removing the dipoles from two consecutive cells. This will cause the dispersion to switch to the other sign as desired, as shown in Fig. 21.

D. FFAGs

Once we reach 5 GeV, it appears to be more cost-effective to use FFAGs rather than RLAs. This conclusion is based on applying a procedure for producing minimum-cost FFAG designs

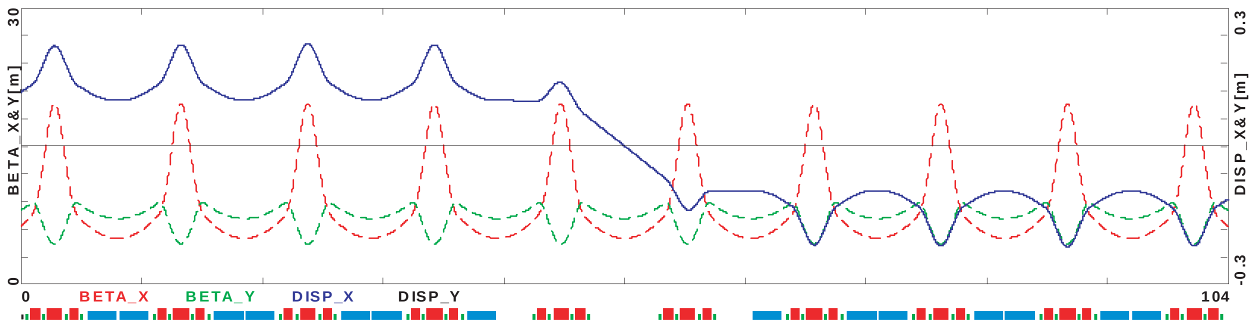

FIG. 21: (Color) A section of the dogbone arc where the bend changes direction, showing the dispersion (solid) and beta functions (dashed).

TABLE IV: Parameters for FFAG lattices.

Cavity rf frequency (MHz)	201.25			
Maximum energy gain per cavity (MeV)	12.7			
rf drift length (m)	2			
Drift length between quadrupoles (m)	0.5			
Initial total energy (GeV)	5		10	
Final total energy (GeV)	10		20	
Number of cells	65		82	
Number of rf cavities	49		56	
Decay loss (%)	5.1		6.5	
Circumference (m)	286		400	
Magnet type	Defocusing	Focusing	Defocusing	Focusing
Magnet length (cm)	77	113	97	141
Field on reference orbit (T)	4.4	-0.8	5.6	-1.1
Field gradient (T/m)	-19	16	-27	24
Magnet aperture radius (cm)	10	16	7	13
Maximum field at the coils (T)	5.7	3.5	7.1	4.3

[22], and comparing the resulting costs to the costs from FS2. Based on this optimization procedure, FFAG designs for accelerating from 5 to 10 GeV and from 10 to 20 GeV are given in Table IV. The lattices consist exclusively of combined-function doublet cells with a

TABLE V: Parameters for FFAG injection and extracting kickers. The specifications here are for FFAG designs that are slightly different from those presented here.

Energy (GeV)	5	10	10	20
Type	Inject	Extract	Inject	Extract
Length (m)	1.5	1.5	1.5	1.5
Kick field (T)	0.37	0.51	0.78	0.58
Maximum field at the coils (T)	3.6	2.6	4.2	5.6
Vertical aperture (cm)	10	10	7.6	7.6
Horizontal aperture (cm)	25	25	19.5	19.5
Current (kA)	44	60	71	53
Supply voltage (kV)	±58	±60	±52	±48
Rise/fall time (ns)	640	950	875	1270
Pulse length (ns)	300	300	300	300
Stored energy (J)	850	1620	2280	1260

drift length sufficient to hold a single-cell 201.25 MHz superconducting rf cavity (similar to the double-cell cavities from FS2). The 2 m length of the drift is needed to keep the fields on the cavity under 0.1 T [20].

With the beam intensity given in Table I, and both signs of muons, about 9% of the stored energy will be extracted from the cavities in the 5–10 GeV FFAG, and about 16% will be extracted in the 10–20 GeV. While this may seem substantial, it is easily handled. To keep the average voltage sufficient to accelerate over the desired range, one need only to add two cavities to the 5–10 GeV FFAG and three cavities to the 10–20 GeV FFAG. The most important effect is a differential acceleration between the head and tail of the bunch train, which is only about 0.3% for both cases. This should be at least partially correctable by a phase offset between the cavity and the bunch train and, in any case, is substantially smaller than the energy spread in a single bunch.

One of the biggest challenges for the FFAGs is injection and extraction. Table V gives the parameters required for injection and extraction kickers. The stored energy in the kicker is high, but is similar to that found in induction linac cells. The rise times and voltages are reasonable. These parameters assume that injection occurs from the inside of the FFAG.

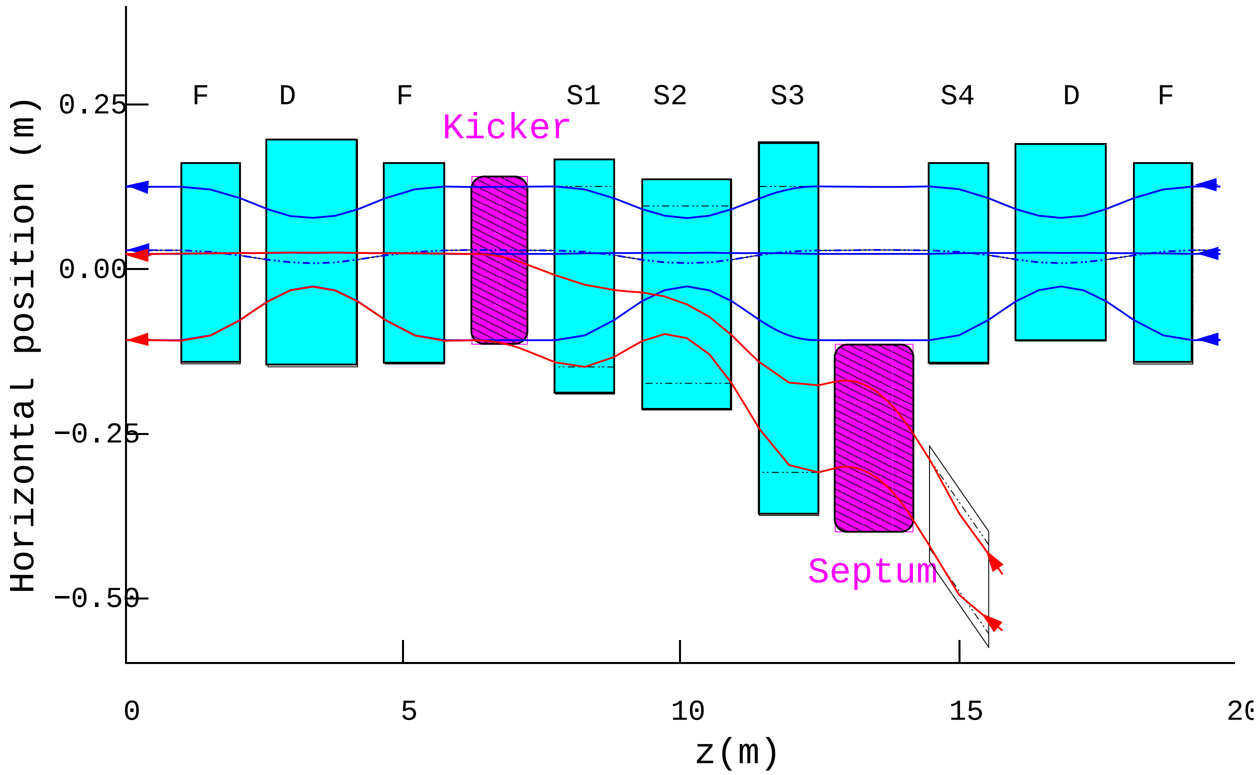

FIG. 22: (Color) Injection into the FFAG.

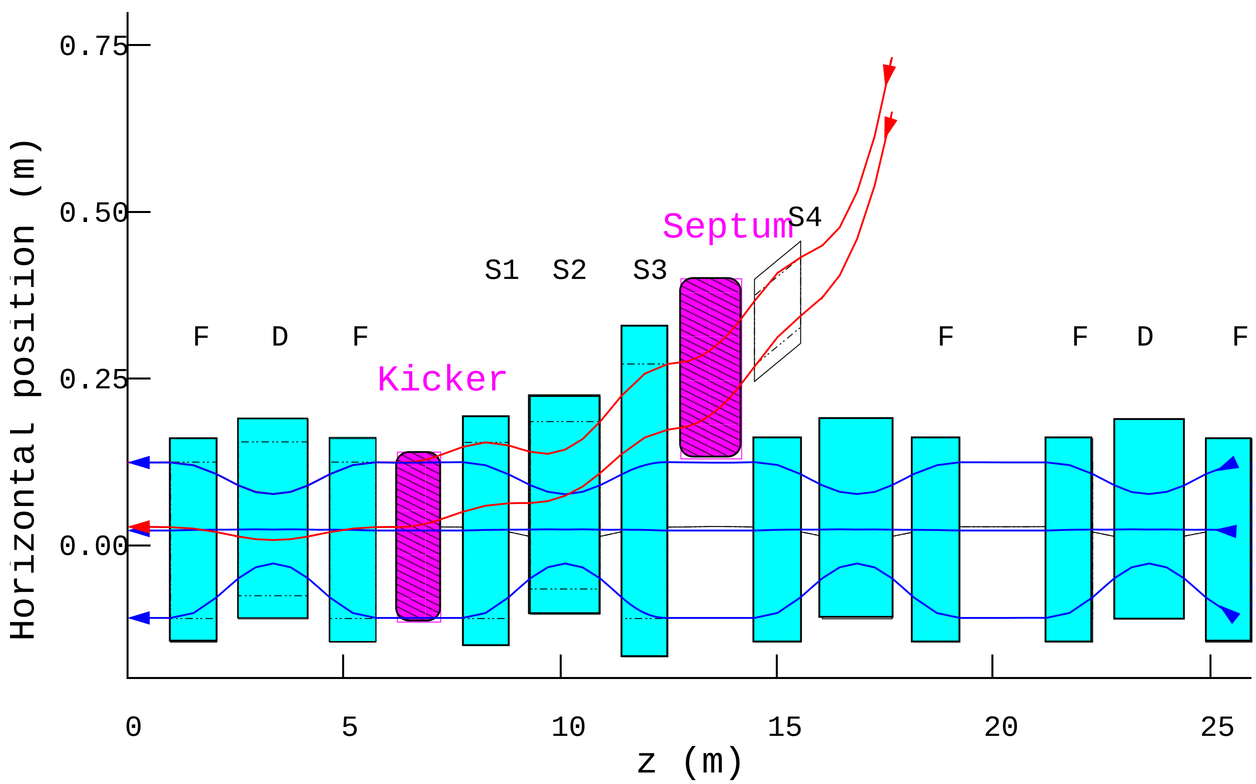

FIG. 23: (Color) Extraction out of the FFAG.

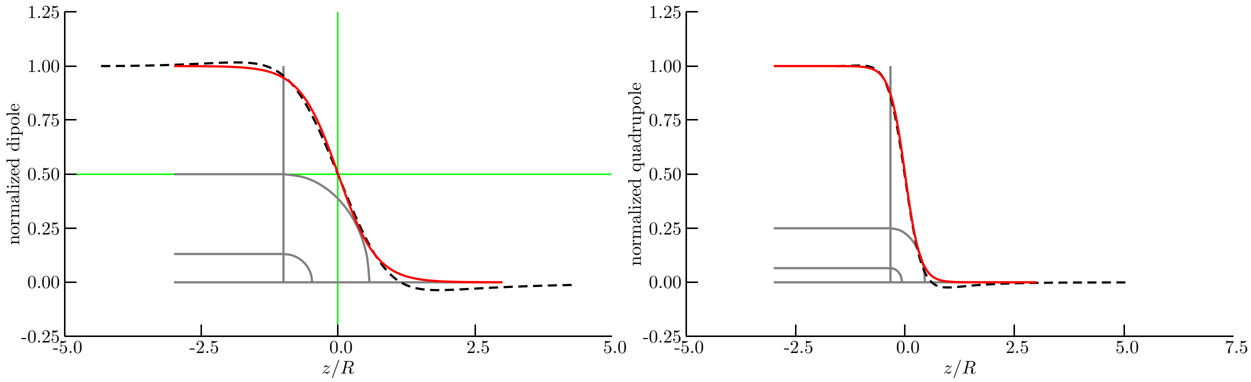

FIG. 24: (Color) Relative dipole field (left) and quadrupole field (right) near the magnet end. The dashed line is the field from TOSCA, while the solid line is our model.

This is preferred since the beam will be near the inside of the FFAG at the lowest energies. Figures 22 and 23 show the injection and extraction layout. The magnets near the kickers and septum must be modified to accommodate the injection and extraction systems, but their effects will be kept as close as possible to those of the other cells in the FFAG lattice to minimize the driving of resonances.

E. FFAG Tracking Results

Initial experience with FFAG lattices having a linear midplane field profile has shown them to have a good dynamic aperture at fixed energies. We are careful to avoid single-cell linear resonances to prevent beam loss. However, since the tune is not constant, the single-cell tune will pass through many nonlinear resonances. Nonlinearities in the magnetic field due to end effects are capable of driving those nonlinear resonances, and we must be sure that there is no beam loss and minimal emittance growth because of this. Furthermore, there is the potential to weakly drive multi-cell linear resonances because the changing energy makes subsequent cells appear slightly different from each other. These effects can be studied through tracking.

ICOOL [16] is used for tracking for several reasons. It will allow for a fairly arbitrary end-field description, it will attempt to make that description consistent with Maxwell's equations, and it will track accurately when the lattice acceptances, beam sizes, and energy spread are all large.

We begin by constructing a simple model of both a quadrupole and dipole $\cos\theta$-type

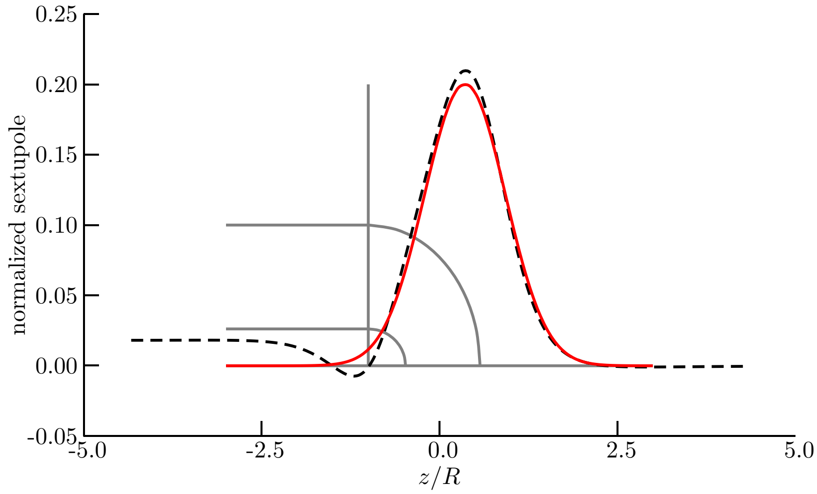

FIG. 25: (Color) Peak magnitude of the sextupole end field at radius R (the magnet aperture), divided by the dipole field. The dashed line is the field from TOSCA, while the solid line is our model.

magnet, without iron, using TOSCA [23]. At the end of the magnet, the field does not immediately drop to zero, but falls gradually, as shown in Fig. 24. The end-field falloff in the dipole and quadrupole generates nonlinear fields, which ICOOL calculates. In addition, there are higher-order multipoles generated by breaking the magnet symmetry at the ends where the coils form closed loops. We use TOSCA to compute the sextupole components that arise from this effect, as shown in Fig. 25, and include them in our computation.

The TOSCA computation is done without iron, which leads to the overshoot in the field values in Figs. 24–25. Iron in the magnet will likely eliminate that overshoot. Thus, we approximate the fields from TOSCA using functions without the overshoot. Fitting roughly to the TOSCA results, the fields are approximated by

$$
\begin{gathered}
B_0(z) = \frac{1}{2}B_{00}\left(1+\tanh\frac{z}{0.7R}\right), \qquad B_1(z) = \frac{1}{2}B_{10}\left(1+\tanh\frac{z}{0.35R}\right) \\
B_2(z) = -0.2B_{00}\exp\left[-\frac{1}{2}\left(\frac{z-0.36R}{0.57R}\right)^2\right],
\end{gathered}
\tag{3}
$$

where R is the magnet aperture radius, $B_0(z)$ is the dipole field, B_{00} is the dipole field in the center of the magnet, B_1 is the quadrupole field, B_{10} is the quadrupole field in the center of the magnet, and B_2 is the maximum magnitude of the sextupole field at the radius R. These fitted functions are shown in their corresponding plots in Figs. 24–25.

Our tracking studies were not done using the lattices described here, but were based on

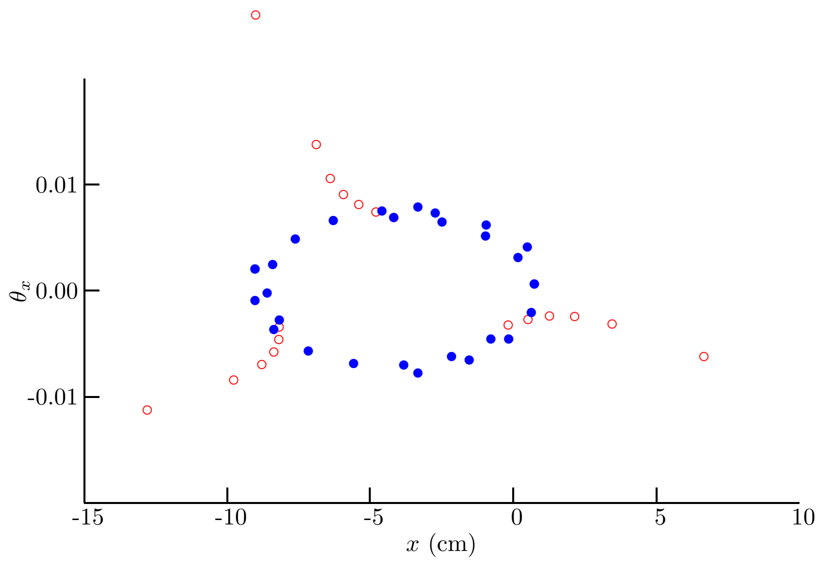

FIG. 26: (Color) Horizontal phase space from tracking at 5.1 GeV/c at the outer edge of the acceptance. Open circles are without the body sextupole fields and show a third-order resonance; filled circles are with the body sextupole fields.

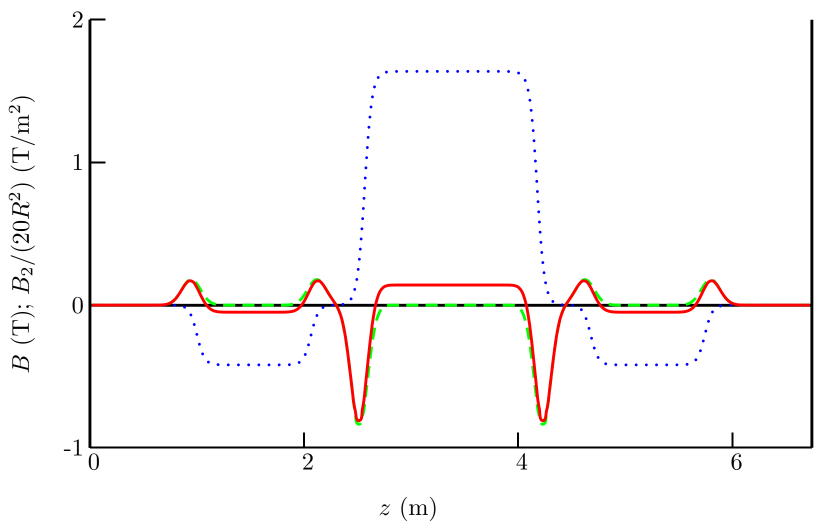

FIG. 27: (Color) Sextupole field components in the lattice used for tracking. The dotted line is the dipole field, the dashed line is $B_2/(20R^2)$ with zero body sextupole field, and the solid line is with sufficient body sextupole field to eliminate the third-order resonance.

earlier designs of a 5–10 GeV FFAG that used a triplet lattice. We expect qualitatively similar results for the lattices described above.

Injecting particles at the outer edge of the acceptance and tracking through several cells indicated a large third-order resonance at around 5.1 GeV/c, as shown in Fig. 26. This

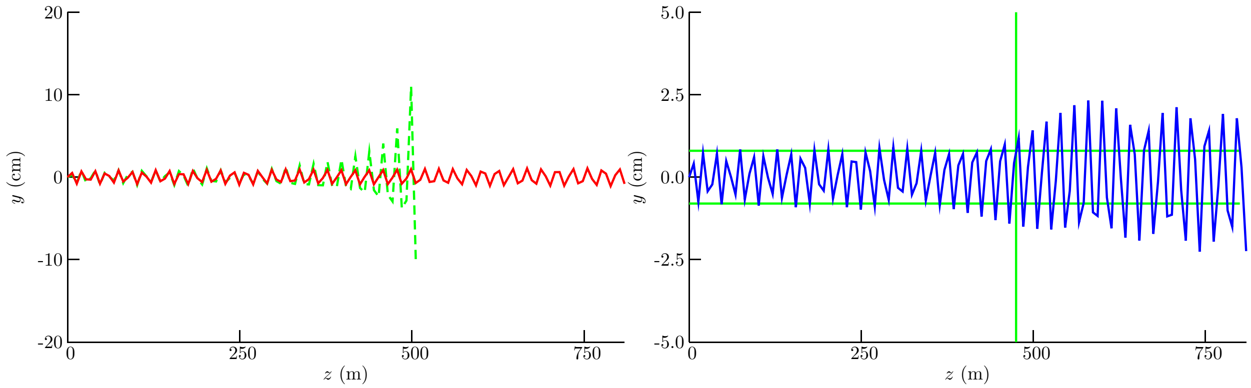

FIG. 28: (Color) Tracking of a particle at the edge of the acceptance with uniform acceleration. On the left the dashed line is without any body sextupole, and the solid line is with the corrected body sextupole. On the right a smaller integrated sextupole correction is used (40% instead of 68%), and significant emittance growth is observed.

resonance is presumably being driven by the sextupole fields at the magnet ends. With some experimentation, it was found that if the integrated body sextupole was set to 68% of the integrated end sextupoles, (see Fig. 27), the resonance was eliminated (also shown in Fig. 26). When acceleration is included, one sees particle loss when accelerating through the resonance if there is no body sextupole correcting the end sextupoles, while there appears to be almost none with the body correction included (see Fig. 28). If the body correction is only partially included, there is significant emittance growth, as seen in Fig. 28. With these sextupole corrections, we can uniformly accelerate over the entire 5–10 GeV energy range without losing a high-amplitude particle or having its amplitude grow by a large amount.

When tracking with rf is considered, the longitudinal dynamics behavior is complex [24]. If one begins with an upright ellipse, there is considerable emittance growth if only the 201.25 MHz rf is used (see Fig. 29). Adding a third-harmonic rf considerably reduces the emittance growth, as shown in Fig. 29. The amount of third-harmonic rf required is substantial and that, combined with space considerations, makes this alternative unattractive. An alternative that includes tilting the initial ellipse in phase space, which also reduces the emittance growth, is being studied.

Figure 30 shows a compact potential layout for all the acceleration systems described here.

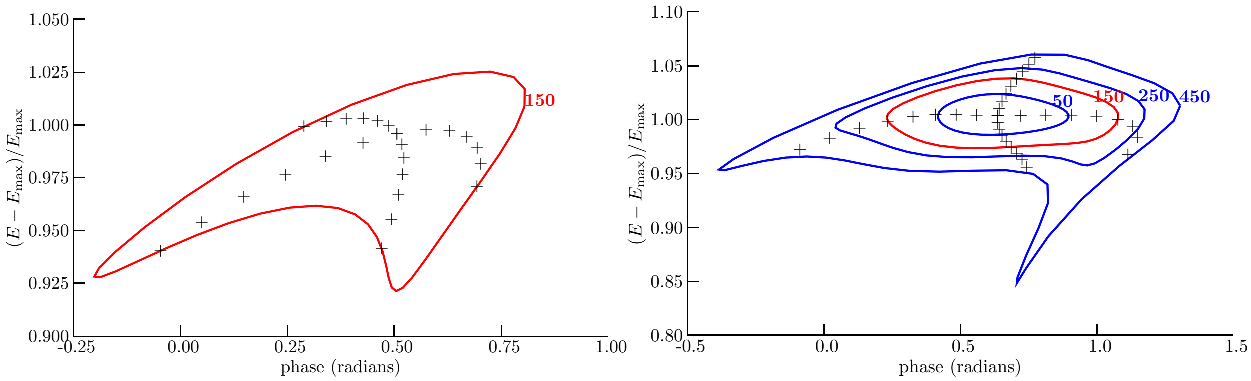

FIG. 29: (Color) Longitudinal tracking starting from an upright ellipse for the 5–10 GeV FFAG. On the left with only 201.25 MHz rf. On the right with third-harmonic rf having voltage equal to 2/9 of the fundamental rf voltage. Curves are labeled with their corresponding acceptance. Crosses for both cases started out as horizontal and vertical lines in phase space.

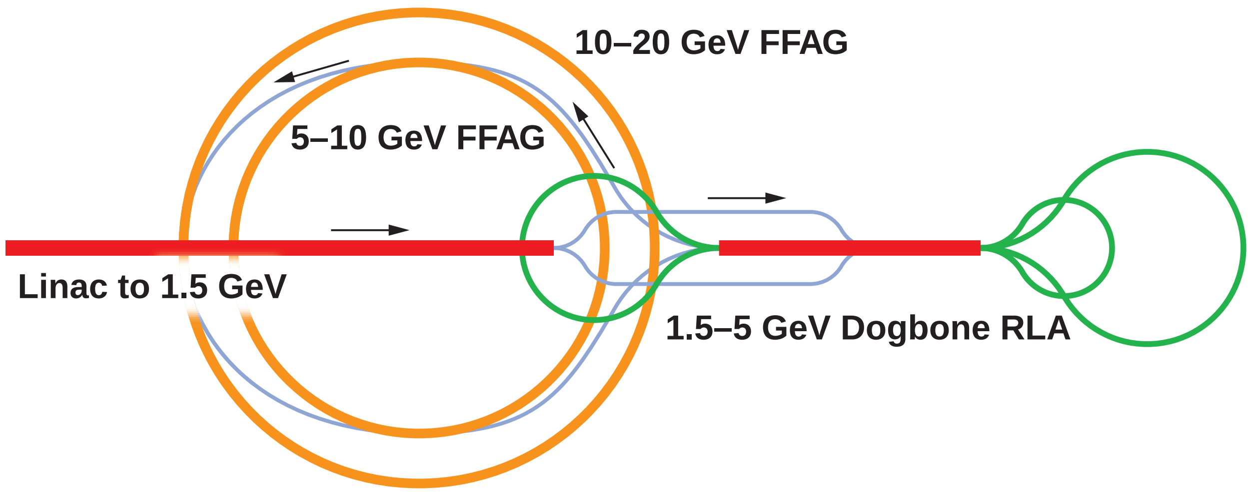

FIG. 30: (Color) Potential layout for the acceleration systems.

F. Design of Combined-Function Superconducting Magnet for FFAGs

An initial design of a superconducting combined-function (dipole–quadrupole) magnet has been developed [25]. The design is for the defocusing magnet from the triplet design used above for tracking studies. The magnets needed for the FFAGs described earlier in a previous sub-section could be designed similarly.

The parameters of this QD combined-function magnet are shown in Table VI, L is the length of the reference orbit inside the magnet, r is the radius of curvature of the reference orbit, R is the radius of the magnet bore,

TABLE VI: Parameters of the QD cell: L_0 is the length of the long drift between the QF magnets; L_q is the length of the short drift between QF and QD magnets; X_0 is the displacement of the center of the magnet from the reference orbit; B_0 is the vertical magnetic field at the reference orbit, and B_1 is the derivative of the vertical magnetic field at the reference orbit.

Initial energy E_{min} (GeV)	10
Final energy E_{max} (GeV)	20
Long drift L_0 (m)	2
Short drift L_q (m)	0.5
Type of magnet	QD
Length of reference orbit L (m)	1.762
Radius of curvature r (m)	18.4
Displacement X_0 (mm)	1.148
Radius of the magnet bore R (cm)	10.3756
Vertical magnetic field B_0 (T)	2.7192
Gradient B_1 (T/m)	-15.495

The magnet design is based on a cosine-theta configuration with two double layers for each function. The quadrupole coil is located within the dipole coil and both coils are assembled using key-and-bladder technology. All coils are made with the same Nb–Ti cable capable of generating the operating dipole field and gradient with about the same current of 1800 A. A single power supply is thus possible with a bit of fine tuning. The maximum central dipole field and gradient at short sample are 4.1 T and 26 T/m, as compared with the requirements of 2.7 T and 15.4 T/m, respectively. At this early design stage, excess margin is left for safety and perhaps a field-rise in the magnet end region. The maximum azimuthal forces required for magnet pre-stress are of the order of 1 MN/m (assuming maximum safety). The conductor strand size and cable parameters common to both dipole and quadrupole are listed in Table VII.

The initial cross sections of both dipole and quadrupole were designed to give less than 0.01 units of systematic multipole errors at a radius of 70 mm. It is straightforward to readjust the design to cancel the end-field multipoles as proposed in sub-section IV F. developed, as shown in Figs. 31, and 32.

TABLE VII: Nb–Ti conductor for dipole and quadrupole coils.

Strand diameter (mm)	0.6477
Cable width, bare (mm)	6.4
Cable thickness, insulated (mm)	1.35
Keystone angle (deg.)	0.6814
Conductor type	Nb–Ti
Cu:SC ratio	1.8:1
Current density (at 5 T, 4.2 K) (A/mm^2)	2850
Number of strands	20

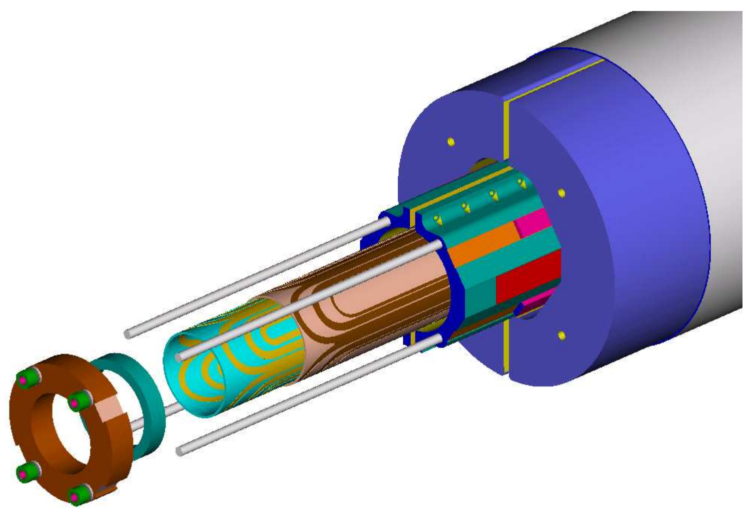

FIG. 31: (Color) Exploded view showing the two quadrupole layers (dipole coils not shown).

An alternative concept would be to use a single dipole like design with laterally displaced poles (see Fig. 33) as discussed in Refs. [26, 27, 28].

V. COST ESTIMATE: ASSUMPTIONS AND ALGORITHM

A. Estimation of ST2B costs based on FS2

The following analysis is intended only as a status report on the collaboration's ongoing efforts to reduce the cost of a Neutrino Factory. At this time major efforts have been directed at the three major components: Phase Rotation, Cooling, and the higher energy part of the

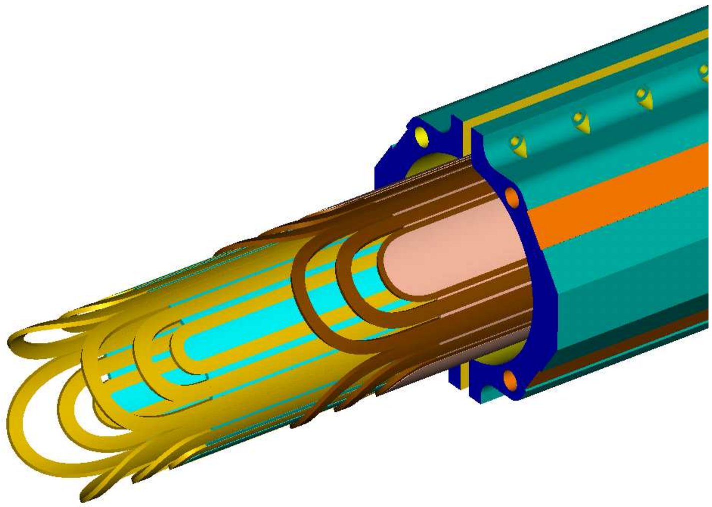

FIG. 32: (Color) Close-up of quadrupole coil return-end windings.

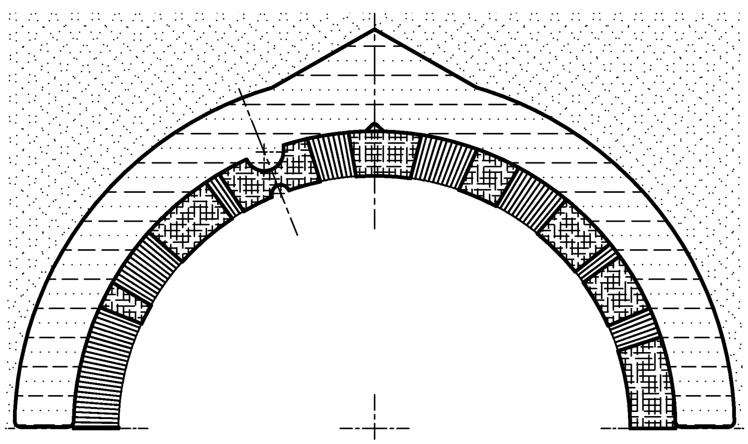

FIG. 33: Expanded view of the cross section of the superconducting combined function magnet used in the 50 GeV proton beamline for the J-PARC neutrino experiment [26].

Muon Acceleration. Nevertheless, a comparison with FS2 suggests substantial progress has been achieved.

Starting from a FS2 Work Breakdown, we derive element costs per unit length, integral rf voltage, or net acceleration; these costs are then applied to the ST2B parameters after scaling for magnetic fields, radii, stored energy, rf gradient, etc.

1. *Estimated Study 2B subsystem costs and comparison with FS2*

Estimates of ST2B subsystem costs are obtained by scaling from those assumed for FS2.

The Proton Driver component cost estimate is taken, without modification from FS2. This estimate was for an upgrade of the AGS to 1 MW beam power, and is thus completely site dependent. However, since it would already be required for an assumed preceeding Super-Beam experiment, it would not be part of the cost to upgrade to a Neutrino Factory. The remaining components are largely site independent.

The ST2B target and capture component are essentially the same as that in FS2, but the transport region in which the field is tapered down is shorter by 5.5 m because it tapers from 20 T to 1.75 T. The savings are estimated by subtracting the cost of 5.5 m of a 1.25 T transport, whose cost per meter is taken from the drift in FS2.

The first 18 m of drift, in either study, would be more expensive than later beam transports because of required radiation shielding in the early part. We therefor treat the first 18 m of drift separately from the subsequent transport. Costs for the first 18 m are taken from FS2, with only a correction for the higher specified solenoid fields (1.75 T vs. 1.25 T). For this correction we increase the magnet, power supply and cryogenic costs using the second scaling formula from Mike Green [29], which gives magnet costs $\propto (B\,R^2\,L)^{0.577}$. The subsequent drift will require less shielding and will thus be less expensive. In FS2, there is no further simple drift from which to scale this cost. Instead, we take the costs from the magnets, power supplies and cryogenics included in the induction linac section of FS2. These, which were at 1.25 T, are corrected for the higher 1.75 T specified fields.

The buncher and phase rotator in ST2B are quite different from the induction system in FS2. The focusing now consists of an essentially continuous solenoid at 1.75 T, as in the drift, but with a radius (65 cm), sufficient large to contain the rf cavities. For the cost of this solenoid we again scale from the FS2 induction linac transport magnets, corrected, again using Green's second formula, but now correcting for both the higher field and the greater radius. This estimate is again conservative, because again it is scaled from the more difficult transport solenoids inside the induction linacs of FS2. Instead of induction linacs, the ST2B uses a sequence of rf cavities at varying frequencies in the range of 300-200 MHz). The cost of these cavities and needed rf power supplies are scaled from the FS2 costs of cavities used in the cooling (201 MHz). These costs are scaled for the different average accelerating gradients as follow: cavity cost per GeV is inversely proportional to the gradient; power supply cost per GeV is proportional to the gradient.

The rf in ST2B cooling is essentially identical to that in FS2, so the costs per GeV are taken to be the same. The focusing lattice, however, is quite different: a simple alternating solenoid FOFO lattice in ST2B, instead of the more complicated, and tapered, super FOFO lattices in FS2. We estimate the new system cost by scaling in the lattice from the first FS2 cooling lattice, now using Green's first scaling formula that depends on the total stored energy (cost $\propto U^{0.662}$), with $U = B^2\, R^2\, L$.

The cryogenic cost is also scaled with the magnet costs, but now scaled from cryogenic cost in the FS2 phase rotation, rather than that in the cooling, because the cryogenic costs in FS2 cooling includes cooling of absorbers.

The matching section is required to match the beta function vs. momentum in the cooling channel to that in the pre-accelerator. In the FS2 case, the beta function vs. momentum in the cooling lattice had a highly non-linear character, with low beta functions at the upper and lower momentum limits and a maximum beta function in the center. In contrast, the beta functions in the pre-accelerator are approximately linear in momentum. Matching was achieved in 18.15 m of a modified 1.65 m cell cooling lattice. In ST2B, the match is simpler and less expensive because: **FIXME**

- the beta functions both before and after the match have similar linear momentum dependence

- the match is only between 0.8 to 2.7 m beta functions, compared with 0.2 to 2.7 m in FS2

- because the lattice on which it will be based has the relatively lower stored energy per unit length (189/1039)

The pre-acceleration cost is scaled from FS2 by energy gain, ΔE for rf cavities, power supplies and cryogenics, and with length (L) for magnets, vacuum and conventional equipment.

ST2B requires a 1.5 to 5 GeV RLA, whose cost is scaled from the 2.5 to 20 GeV RLA in FS2. The number of passes is 3.5, compared with 4 for FS2. A dog bone design is favored for ease of the switchyard design, whereas FS2 used a race track, but costs per unit length, or per energy gain are probably very similar, and assumed here to be the same. The arcs are assumed to have the same average bending field as the larger energy FS2 arc, and the cost per length is taken to be the same. The special and transport magnet costs were scaled with the final RLA energies.

The costs for all technical and conventional costs are taken from a cost algorithm discussed in the acceleration section. This algorithm when applied to the FS2 RLA gives a higher cost than that stated in the report. The algorithm thus appears to be conservative when compared with FS2. Injection/Extraction kickers are assumed to be driven by typical induction linac power sources, and will contain similar amounts of magnetic materials. The costs are taken from the FS2 induction linacs that has the same pulsed energy. Transfer line lengths are taken from Palmer's report [30] on Injection/Extraction at TRIUMF and include lines for both signs. Cost per m of these transport lines is taken from RLA arcs (magnets, PS, and vacuum: $23.9 + 4.1 = 28.0k\$/m$).

Storage ring costs are taken, without modification, from FS2. In that case, sited at BNL, there was a constraint that no part of the downward tilted ring should fall below the nearby water table. This constraint forced the construction of the ring in an artificial hill and also required unusually high, and not cost optimized, bending fields to keep the ring small. The cost at another location, without this constraint, would be most likely less.

A summary of ST2B Costs, and comparison with FS2is presented in Table VIII as well as percentage of the cost reductions.

VI. ALTERNATIVE DESIGN CHOICES

A major simulation effort will continue to focus on iterating the front-end channel design to optimize it for cost and performance. Some of the options that have already been studied briefly or that might be developed in future studies include:

- Be absorbers in place of LiH absorbers
- Different rf frequency
- Shorter buncher rotator
- Gas-filled cavities
- Shorter bunch train
- Quadrupole-based cooling channel

TABLE VIII: Costs comparison FS2 and the present study (ST2B) in arbitrary units.

System	FS2	ST2B	Cost Change
	M$	M$	%
Proton Driver	168	168	100
Target, capture	92	90	98
Drift Channel	6	26	433
Phase Rotation	307	80	26
Mini-cooling and magnetic field flip	11	0	
Buncher and Matching	76	45	59
Cooling	310	182	59
Matching to pre-accelerator	57	34	60
Pre-accelerator linac	133	70	53
Acceleration	356	261	73
Storage Ring	83	83	100
Totals without controls or site utilities	1599	1039	65

If a problem develops with heating in the absorber windows a number of alternative designs have been investigated. The thermal conductivity to the heat sink on the outer edge of the window can be improved by breaking up the LiH into several pieces, separated by layers of high conductivity Be. Using a total thickness that gives the same total energy loss as the original window results in only $\approx$ 3% loss in the accepted muon flux. Other window possibilities that gave reasonable muon fluxes include (1) using a thin Be layer on pure lithium and (2) separating the LiH and Be with a region containing helium gas. Another possibility would be to use thicker Be windows as the end plates of the cavities, with a thickness chosen to make the total energy loss the same as that in the baseline foil + LiH absorber case. This would eliminate the need for thin Be windows. Cooling would be a bit less effective because of the greater multiple scattering in Be absorbers. An initial evaluation [31] of a Be-only scenario showed somewhat less capture into the acceleration channel acceptance ($\approx$15%). A scenario in which Be absorbers are initially installed and then upgraded later to more efficient LiH absorbers is, of course, also possible.

The baseline scenario requires a roughly 110 m drift, a 51 m buncher and a 52 m high-voltage phase rotation section. These parameters have not been optimized. For comparison, we considered an example having only a 26 m phase rotation section [31]. The alternative phase rotation section would be significantly less expensive, since it is not only shorter but provides about 200 MV less of high-gradient rf voltage. Initial evaluations indicate only small decreases in captured muons ($\approx$10%).

The baseline case generates μ^+ and μ^- bunch trains that are about 100 m long. These bunch trains are matched to the FS2 scenario requirements; in particular, they fit within the circumferences of the presently envisioned accelerators and storage ring. However, other scenarios might make use of smaller circumference ring coolers, accelerators, and storage rings, and thus require shorter bunch trains. For example, a scenario with a 20 m drift, 20 m buncher and 20 m phase rotator has been explored [31]. This produces a roughly 20 m long bunch train. Although this shorter system would be much less expensive than the present roughly 200 m long system, an initial evaluation showed that the total number of captured muons was substantially reduced (by about 50%). (On the other hand, a longer system, capturing longer bunch trains, might produce more muons, at a small incremental cost.)

Both FS2 and our present scenario use 201.25 MHz rf as the baseline final operating frequency, because of the availability of rf components at that frequency and because it is a plausible optimum for large-aperture and high-gradient operation. Other baseline frequencies could be considered, e.g., scenarios at 50, 100, 300 or 400 MHz. Lower frequencies (larger bunches) may be desirable if the accelerator longitudinal motion requires larger phase-space buckets.

Muons Inc. has an STTR grant to explore the use of hydrogen gas–filled rf cavities for muon cooling [32]. This approach simplifies the cooling channel design by integrating the energy-loss material into the rf system. Moreover, it may be more effective in permitting high-gradient operation of the cavities. Such cavities could also be used in the cooling and phase-rotation (and possibly buncher) sections; an exploration with cost-performance optimization is planned.

The transport and cooling system in the front-end scenario considered here uses high-field solenoids for focusing. A cooling system with similar performance parameters using large-aperture quadrupoles has also been examined [33], though a cost-performance comparison has not yet been made.

Further effort will be given to beam dynamics studies in the FFAG rings and storage ring, including realistic errors. Work on optimizing the optics design will be done. Assessment of field-error effects on the beam transport will be made to define acceptance criteria for the magnets. This will require use of sophisticated tracking codes, such as COSY [34], that permit rigorous treatment of field errors and fringe-field effects. In many ways, the storage ring is one of the most straightforward portions of a Neutrino Factory complex. However, beam dynamics is an issue here as the muon beam must circulate for many hundreds of turns. Use of a tracking code such as COSY is required to assess fringe field and large aperture effects. Suitable magnet designs are needed, with the main technical issue being the relatively high radiation environment. Another lattice issue that must be studied is polarization measurement. In the initial implementation of a Neutrino Factory it is expected that polarization will not be considered, but its residual value will nonetheless be important in analyzing the experiment.

As should be clear from the design descriptions, the muon-based Neutrino Factory is a demanding project. The machine makes use of novel components and techniques that are, in some cases, at or beyond the state of the art. For this reason, it is critical that R&D efforts to study these matters be carried out.

Successful construction of a muon storage ring to provide a copious source of neutrinos requires many novel approaches to be developed and demonstrated; a high-luminosity Muon Collider, which might someday follow, would require an even greater extension of the present state of accelerator design. Thus, reaching the desired facility performance requires an extensive R&D program. Each of the major systems has significant issues that must be addressed by R&D activities [5]. Component specifications need to be verified. For example, the cooling channel assumes a normal conducting rf (NCRF) cavity gradient of 15 MV/m at 201.25 MHz, and the acceleration section demands similar performance from superconducting rf (SCRF) cavities at this frequency. In both cases, the requirements are beyond the performance reached to date for cavities in this frequency range. Development and testing of efficient high-power rf sources at a frequency near 200 MHz are also needed. The ability of the target to withstand a proton beam power of up to 4 MW must be confirmed. Finally, an ionization cooling experiment should be undertaken to validate the implementation and performance of the cooling channel, and to confirm that our simulations of the cooling process are accurate.

VII. SUMMARY

A new type of facility has been proposed that could have a tremendous impact on future neutrino experiments—the Neutrino Factory. In contrast to conventional muon neutrino beams, the Neutrino Factory would provide a source of electron neutrinos (ν_e) and antineutrinos ($\bar{\nu}_e$), with very low systematic uncertainties on the associated beam fluxes and spectra. The experimental signature for $\nu_e \rightarrow \nu_\mu$ transitions is extremely clean, with very low background rates. Hence, Neutrino Factories would enable very sensitive oscillation measurements to be made.

An impressive Neutrino Factory R&D effort has been ongoing in the U.S. and elsewhere over the last few years, and significant progress has been made towards optimizing the design, developing and testing the required accelerator components, and significantly reducing the cost.

The novel accelerator facility described here represents a significant improvement over previous designs. New ideas in bunching, phase rotation and ionization cooling have been incorporated into the design of the front end. The non-scaling FFAG acceleration concept has been further developed and used for accelerating the muons up to the 20 GeV design energy. Taken together, these new designs should give a facility that produces the required machine performance at a substantialy reduced cost.

Acknowledgments

We would like to thank S. Geer and Y. Fukui for helpful discussions. This research was supported by the U.S. Department of Energy under Contracts No. DE-AC02-98CH10886, No. DE-AC02-76CH03000, and No. DE-AC03-76SF00098.

[1] S. Ozaki, R. Palmer, M. Zisman, and J. Gallardo, eds., Tech. Rep., BNL-52623 (2001), `http://www.cap.bnl.gov/mumu/studyii/FS2-report.html`.

[2] S. Geer, Phys. Rev. **D57**, 6989 (1998), *ibid.* **59**, 039903E (1999).

[3] M. M. Alsharo´a *et al.*, Phys. Rev. ST Accel. Beams **6**, 081001 (2003).

[4] A. Blondel *et al.*, CERN 2004-002 ECFA/CERN.

[5] *APS Multi-Divisional Study of the Physics of Neutrinos*, `http://www.interactions.org/neutrinostudy/`, sponsored by the American Physical Society Divisions of: Nuclear Physics, Particles and Fields, Astrophysics, Physics of Beams (2004).

[6] The Neutrino Factory and Muon Collider Collaboration WEB page, `http://www.cap.bnl.gov/mumu/`.

[7] `http://hepunx.rl.ac.uk/neutrino-factory/`.

[8] `http://muonstoragerings.web.cern.ch/muonstoragerings/`.

[9] `http://www-prism.kek.jp/nufactj/index.html`.

[10] N. Holtkamp and D. Finley, eds., Tech. Rep. Fermilab-Pub-00/108-E, Fermilab (2000), `http://www.fnal.gov/projects/muon_collider/nu-factory/nu-factory.html`.

[11] D. Neuffer, *Exploration of the high-frequency buncher concept*, MUC-NOTE-269 (2003), all MUC-NOTE papers are available from `http://www-mucool.fnal.gov/notes/noteSelMin.html`.

[12] D. Neuffer, *Beam dynamics problems of the muon collaboration: ν-factory and $\mu^+ - \mu^-$ colliders*, MUC-NOTE-266 (2003).

[13] D. Neuffer, *High-frequency buncher and phase rotation for the muon source*, MUC-NOTE-181 (2000).

[14] D. Neuffer and A. Van Ginneken, Proceedings of the 2001 Particle Accelerator Conference (2001), `http://accelconf.web.cern.ch/Accel/Conf/p01/PAPERS/TPPH162.pdf`.

[15] A. Van Gineeken, Tech. Rep., Fermilab (2001), MUC-NOTE-220.

[16] R. Fernow, in *Proceedings of the 1999 Particle Accelerator Conference*, edited by A. Luccio and W. MacKay (1999), p. 3020, latest version available at `http://pubweb.bnl.gov/people/fernow/icool/readme.html`.

[17] N. Mokhov, in *Proceedings of the 2001 Particle Accelerator Conference* (2001), p. 745, see also `http://www-ap.fnal.gov/MARS/`.

[18] K. McDonald, in *Proceedings of the 2001 Particle Accelerator Conference* (2001), p. 1583, also, H.G. Kirk *et al.*, ibid. p. 1535 and Chapter 3 in [1]. All Particle Accelerator Conference papers can be obtained from `http://accelconf.web.cern.ch/accelconf/`.

[19] A. Chao and M. Tigner, eds., *Handbook of Accelerator Physics and Engineering* (World Sci-

entific, 1999).

[20] M. Ono *et al.*, *Magnetic field effects on superconducting cavity*, 9$^{\text{th}}$ Workshop on RF Superconductivity (1999), Los Alamos, NM, 2000), Los Alamos National Laboratory report, LA-13782-C.

[21] J.S. Berg, C. Johnstone, and D. Summers, in *Proceedings of the 2001 Particle Accelerator Conference*, edited by P. Lucas and S. Webber (2001), p. 3323, D.J. Summers, Snowmass 2001, hep-ex/0208010.

[22] J. S. Berg, *Recent results from optimization studies of linear non-scaling FFAGs for muon acceleration*, MUC-CONF-ACCELERATION-309, to appear in the proceedings of FFAG 04, Tuskuba, Japan, 13–16 October 2004 (2004), `http://www-mucool.fnal.gov/notes/noteSelMin.html`.

[23] *Vector Fields Inc., computer program OPERA-3d*.

[24] S. Koscielniak and C. Johnstone, in *Proceedings of the 2003 Particle Accelerator Conference*, edited by J. Chew, P. Lucas, and S. Webber (2003), p. 1831.

[25] S. Caspi and R. Hafalia, *A Combined Function Superconducting Magnet for Fixed-Field Muon Acceleration in an Alternating Gradient Ring: First-Cut*, LBNL Report SC-MAG-839 (2004).

[26] T. Nakamoto *et al.*, IEEE Trans. Appl. Supercond. **14**, 616 (2004).

[27] T. Ogitsu *et al.*, IEEE Trans. Appl. Supercond. **14**, 604 (2004).

[28] T. Nakamoto *et al.*, *Development of Superconducting Combined Function Magnets for the Proton Transport Line for the J-PARC Neutrino Experiment*, to appear in the Proceedings of the 2005 Particle Accelerator Conference (2005).

[29] M. A. Green, R. Byrns, S.J.St. Lorant, Advances in Cryo. Eng. **37**, 637 (1992), LBNL-30824; see also Phys. Rev. **D66**, Review of Particle Physics, 010001-217 (2002).

[30] J. S. Berg, R. Fernow, and R. B. Palmer, FFAG Workshop, Vancouver, Canada (2004), see `http://www.triumf.ca/ffag2004/`.

[31] D. Neuffer (2004), presentation at APS Study Workshop, ANL; `http://www.cap.bnl.gov/mumu/study2a/notes/neuffer.pdf`.

[32] R. Johnson *et al.*, in *Proceedings of Particle Accelerator Conference* (2003), p. 1792, Muon Inc., MUC-NOTE-247 (2002); R. Johnson *et al.*, AIP Conf. Proc. **671**, 328 (2003).

[33] C. Johnstone *et al.*, Nucl. Instrum & Meth. **A519**, 472 (2004).

[34] K. Makino and M. Berz, Nucl. Instrum. & Meth. **A427**, 338 (1999).