BNL-75054-2005-JA
MUC-NOTES-327

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: November 9, 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 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 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

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

yield neutrinos that are directed along the line of the straight sections. This allows them to be observed at near and far detectors, and offers 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 1–4 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 NFMCC [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 several of the required key accelerator experiments are now 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. 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 performance improvements and cost reduction for this ambitious facility. The paper is organized as follows. Section II describes the Neutrino Factory design concept. The design of the front end of the facility is described in Section III and the accelerator chain is described in Section IV. In Section V we discuss the storage ring and the overall performance. Required R&D is described in Section VI. In Section VII we discuss the assumptions used to make the cost estimate for a Neutrino Factory and finally, we conclude with a summary in Section VIII.

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]; a schematic of the whole facility is shown in Fig. 1. The facility is a *quartenary beam* machine; that is, a primary proton beam is used to create first a secondary pion beam and subsequently, a tertiary muon beam that decays and eventually provides the neutrino flux for the detector. For a Neutrino Factory the primary 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

TABLE I: Summary of the main parameters

Proton energy (GeV)	24
Driver cycle rate (Hz)	2.5
Bunches per spill	6
Average bunch rate (Hz)	$2.5 \times 6 = 15$
Protons per bunch (10^{12})	1.6
Bunch length, rms (ns)	3
Proton power (MW)	1
Final muon energy (GeV)	20
Muons of each sign per proton after cooling	0.17
Muons of each sign per proton after acceleration	0.11
Muons of both signs per 10^7 sec decaying toward detector	2×10^{20}

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.

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. 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. NFMCC. 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. Figure 1 shows a schematic of the facility. A summary of parameters is given in Table I.

The main ingredients of a Neutrino Factory include:

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

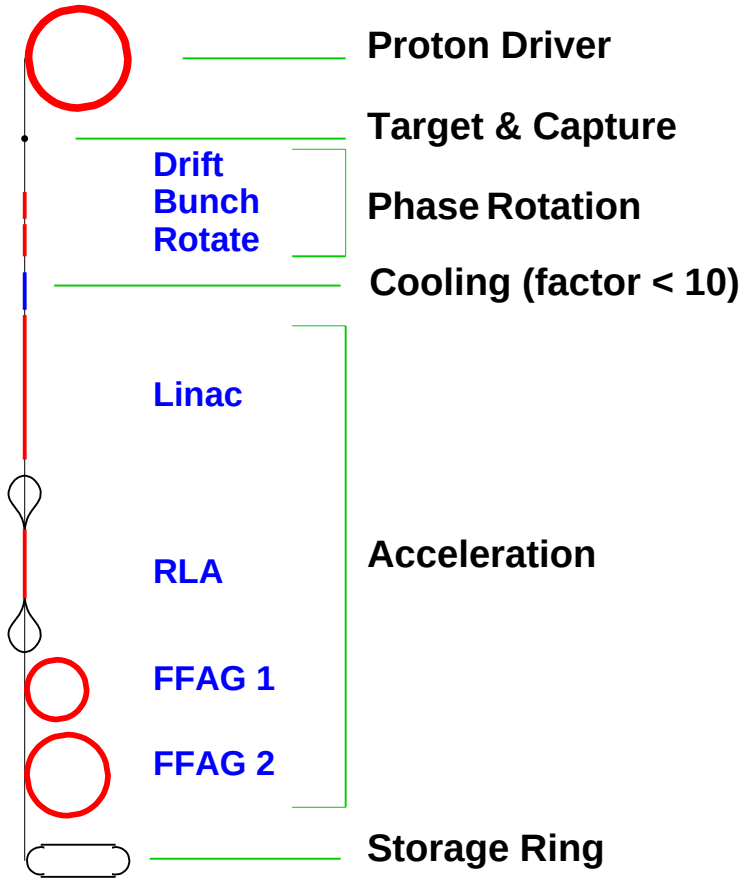

FIG. 1: (Color) Schematic of a Neutrino Factory Facility.

- **Target and Capture:** A high-power target is 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, with a triplet 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 while traveling toward detectors located nearby, and 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

Some front end parameters are given in Table II. 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 have led to a new design for the front end that has a crucial performance advantage and is also significantly less expensive. The new concepts are:

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

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

- The increased accelerator acceptance diminishes the demands on the transverse ionization cooling and allows the design of a simplified cooling section with fewer components and reduced magnetic field strength.

TABLE II: Summary of the front-end parameters

Solenoid capture magnetic field (T)	20
Length of taper (m)	12
Solenoid field at end of taper (T)	1.75
Length of drift (m)	99
Solenoid field in drift (T)	1.75
Length of buncher (m)	50
Solenoid field in buncher (T)	1.75
Maximum rf gradient in buncher (MV/m)	10
Length of phase rotator (m)	54
Solenoid field in phase rotator (T)	1.75
RF gradient in phase rotator (MV/m)	12.5
Length of cooler (m)	80
Maximum solenoid field in cooler (T)	3
RF gradient in cooler (MV/m)	15.25

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 [15]. The concept of the adiabatic buncher is compared with the system used in FS2 in Fig. 2. 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 correlation develops between time and energy. Figure 3 shows the longitudinal phase space after the long drift. In FS2 the energy spread in the correlated beam was first flattened (phase rotated) using a series of induction linacs. The induction linacs did an excellent job, reducing the final rms energy spread to 4.4%, but were expensive. The beam was then sent through a series of rf cavities for bunching, which increased the energy spread to $\approx 8\%$.

In the new scheme, the correlated beam is first adiabatically bunched using a series of rf cavities with decreasing frequencies and increasing gradients in such a way that the bunch centers remain at the rf zero crossings, even as their spacing increases because of their differing energies and velocities.

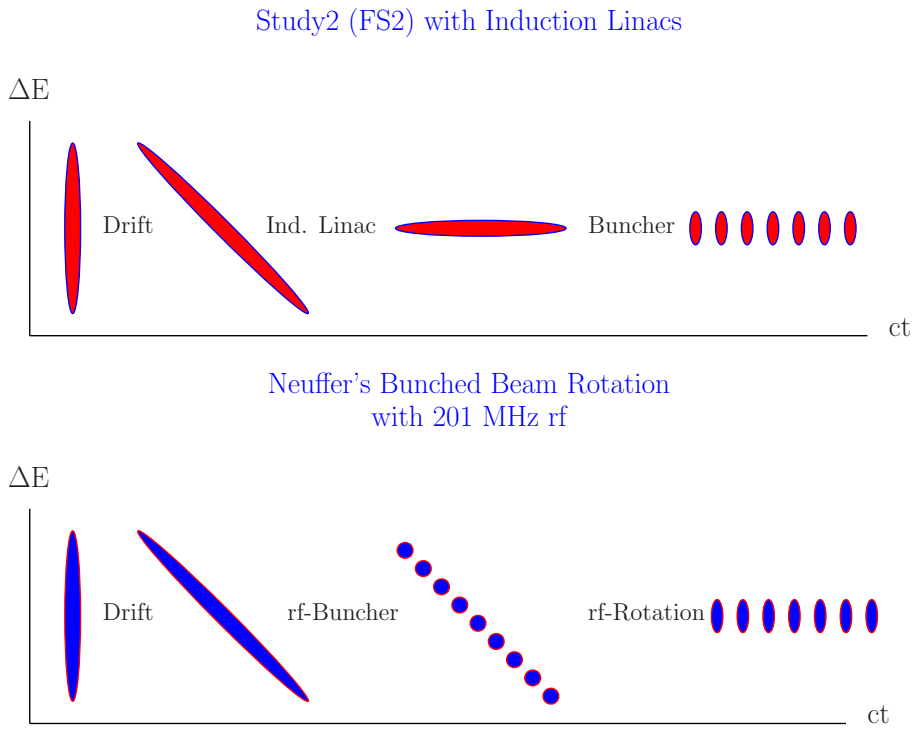

FIG. 2: (Color) Comparison of the buncher concept used here (bottom) with the bunching system used in FS2 (top).

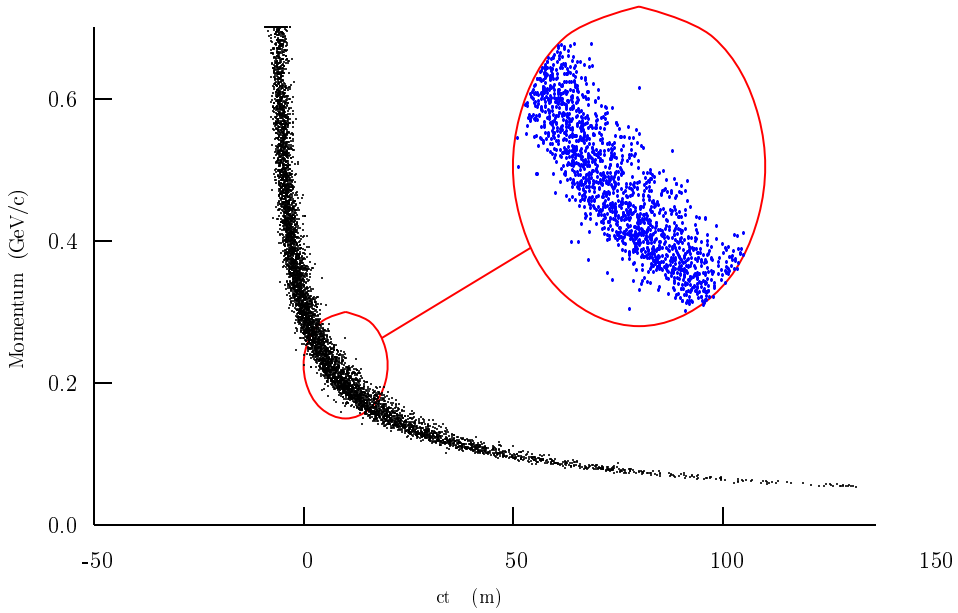

FIG. 3: (Color) Longitudinal phase space after the drift section.

The beam is then phase rotated with a second string of rf cavities with decreasing frequencies and constant gradient. In this case the frequencies are chosen with a slightly different

criterion than that in the bunching. They are chosen so that the high energy bunch centers see a decelerating rf field, while the low energy particles see an accelerating field. 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. 4.

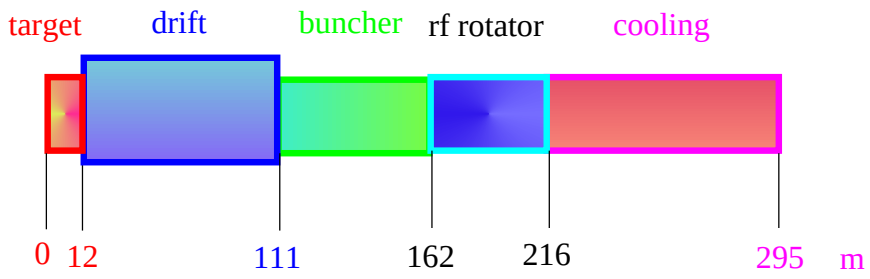

FIG. 4: (Color) Overall layout of the front-end.

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 99 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 and matching take place in 54 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. 5. 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 to the alternating field used in the cooling section.

A. Target and Decay Region

The beam distributions used in the simulations were generated using MARS [16]. The distribution was calculated for a 24 GeV proton beam interacting with a Hg jet [17]. The jet was incident at an angle of 100 mrad to the solenoid axis, whereas the beam was 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 capture region showed

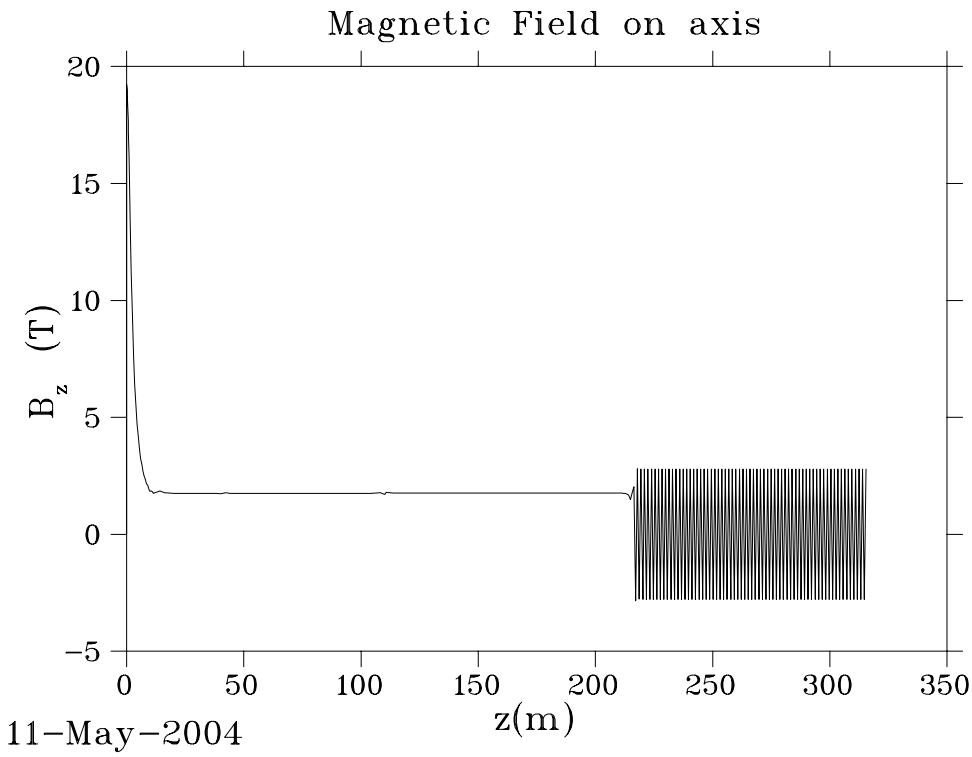

FIG. 5: Longitudinal field component B_z on-axis along the ST2B front-end.

that they have a peak initial longitudinal momentum of $\approx$ 300 MeV/c with a long high-energy tail, and a peak initial transverse momentum $\approx$ 180 MeV/c.

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 $\approx$ 10%. The new axial field profile (marked ST2B) is compared in Fig. 6 with the profile used in FS2. Figure 7 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 maximum error field is $\approx$ 0.07 T.

Figure 8 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. 8, is $\approx 0.5 \times 10^{-8}$ GeV/g per proton on target. This dose is $\approx$ 1 MGy/yr for a 1 MW beam running for a Snowmass year of 1×10^7 s. Assuming a lifetime dose for the insulation of 100 MGy, there should be no problem with radiation damage in the coils, even at a 4 MW intensity level.

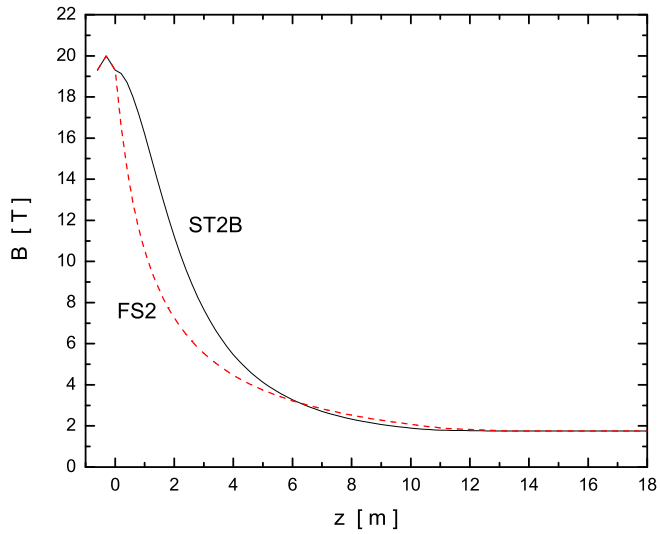

FIG. 6: (Color) Comparison of the capture region magnetic field used in the present simulation (ST2B) with that used in FS2.

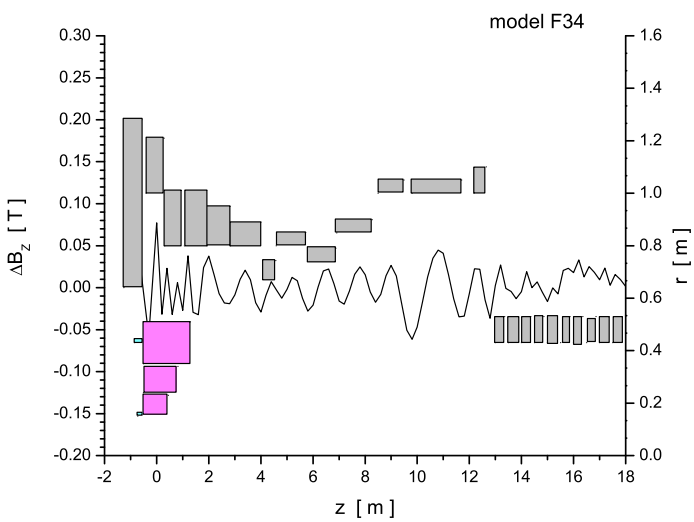

FIG. 7: (Color) Actual coil configuration in the collection region. The left axis shows the error field on-axis compared with the optimal capture field profile, denoted ST2B in Fig. 6.

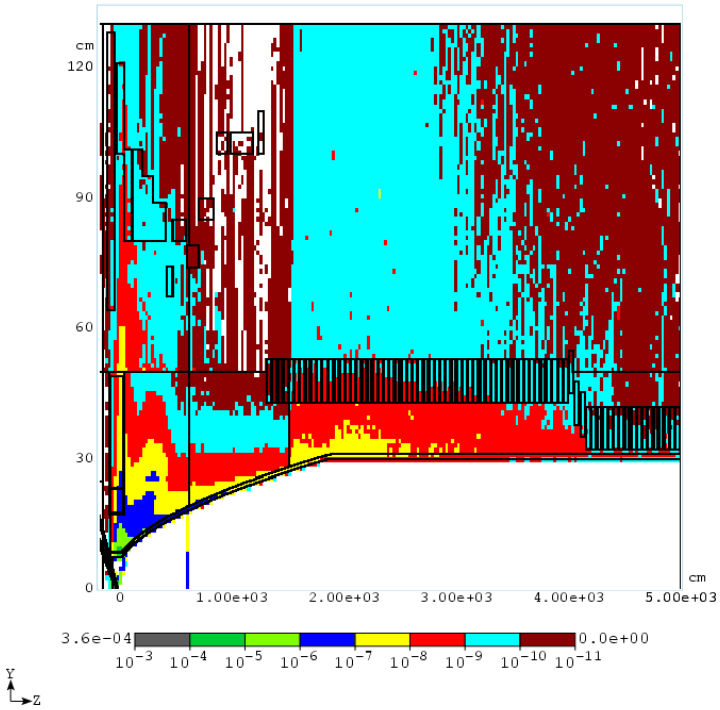

FIG. 8: (Color) MARS calculation of the absorbed annual radiation dose in the collection region for a 1 MW proton beam. The scale shown is in units of GeV/g per proton on target.

B. Bunching and Phase Rotation Region

A cell of the buncher lattice is shown schematically in Fig. 9. Most of the cell length is occupied by the 50 cm long rf cavities. The cavity iris is covered with a Be window. The window thicknesses varied from 200 to 395 microns, depending on the cavity gradient. The limiting radial aperture in the cell is determined by the 25 cm radius of the window. The 50 cm long solenoids were placed outside the rf cavity with periodicity of 75 cm, 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.

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 10 shows the correlated longitudinal phase space and the bunching produced

by the buncher.

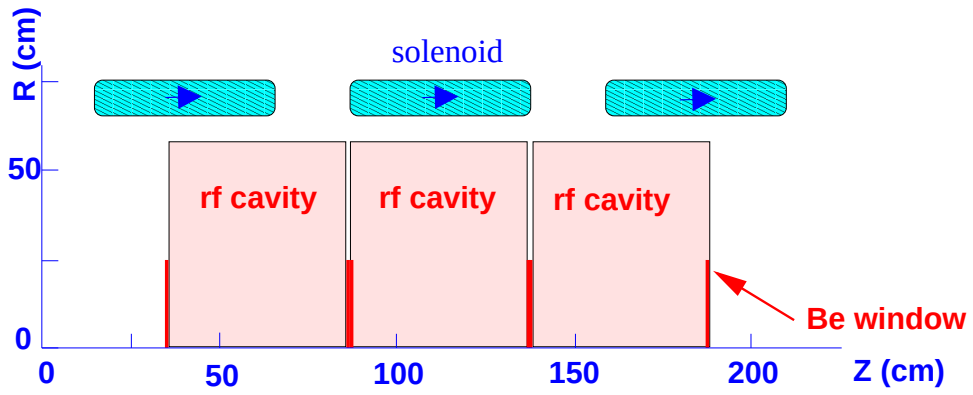

FIG. 9: (Color) Schematic of a cell at the beginning of the phase rotator section.

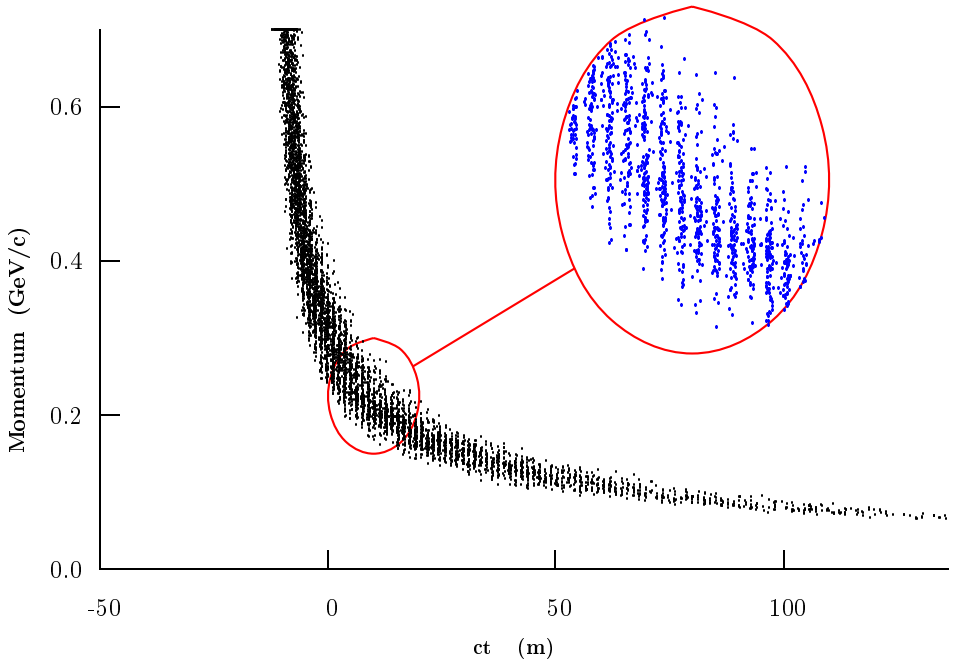

FIG. 10: (Color) Longitudinal phase space of one sign muons after the buncher section.

The phase 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. The tapered window had a thickness of 750 micron from the axis out to a radius of 20 cm and a thickness of 1.5 mm from 20 to 25 cm. 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. Figure 11 shows the longitudinal phase space after the phase rotator. The rms energy spread in the beam is reduced to 27 MeV. To study the

effect of a shorter phase rotator, we also considered an example having only a 26 m phase rotation section [18]. This alternative design would be significantly less expensive, since it is not only shorter but requires about 200 MV less high-gradient rf voltage. Initial evaluations indicate a small decrease in captured muons (≈10%).

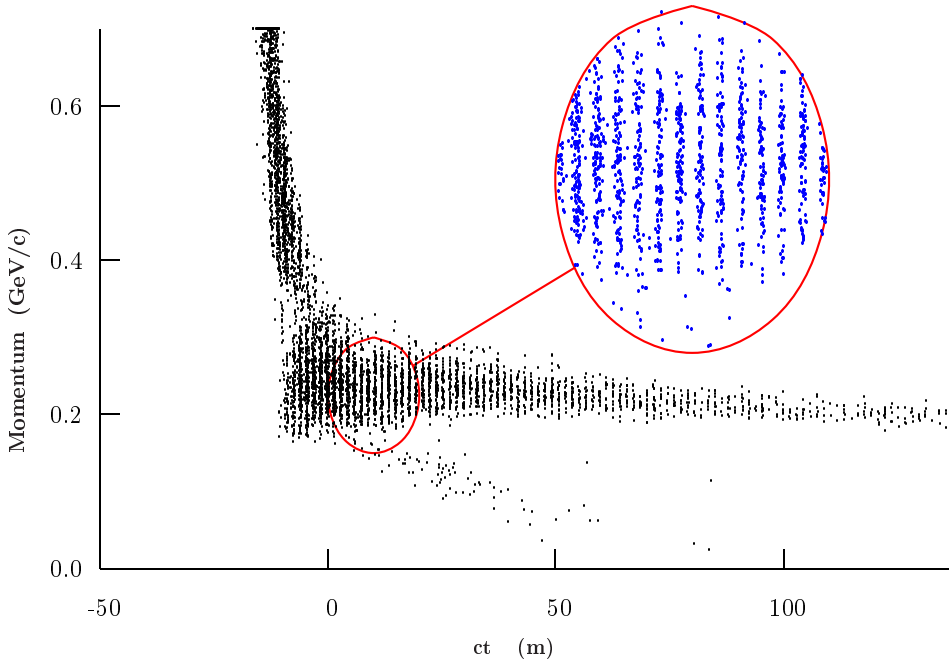

FIG. 11: (Color) Longitudinal phase space after the phase rotation section.

C. Cooling Region

The cooling channel was designed to have a transverse beta function that is relatively constant with position and has a magnitude of about 80 cm. One cell of the channel is shown in Fig. 12. Most of the 150 cm magnetic cell length is taken up by two 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 eliminates the need to place the absorbers at a low-β point to prevent emittance heating. The window consists of a 1 cm thickness of LiH with a 300 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]. A 1 cm space might be introduced between the Be rf window

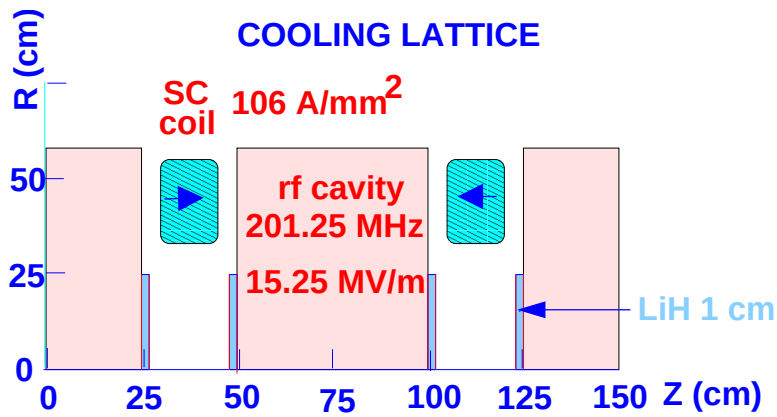

FIG. 12: (Color) Schematic of one cell of the cooling section.

and LiH absorber with flowing He gas to cool both. The alternating 2.8 T solenoidal field is produced with one solenoid per half cell, located between the rf cavities.

Figure 13 shows the longitudinal phase space of one muon sign at the end of the cooling section. The reduction in normalized transverse emittance (ϵ_T) along the cooling channel is shown in the left plot of Fig. 14 and the right plot shows the normalized longitudinal emittance (ϵ_L). The channel produces a final value of $\epsilon_T = 7.1$ mm rad, which is more

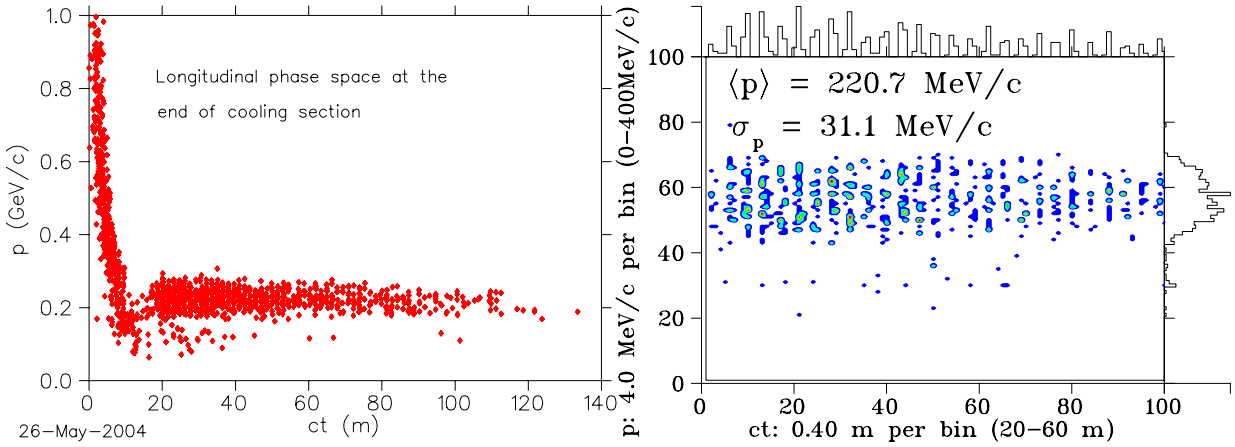

FIG. 13: (Color) Longitudinal phase space at the end of the channel.

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 5.5 mm rad. Figure 15 shows the muons per incident proton on target that fit into the accelerator transverse normalized acceptance of $A_T = 30$ mm rad and normalized 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

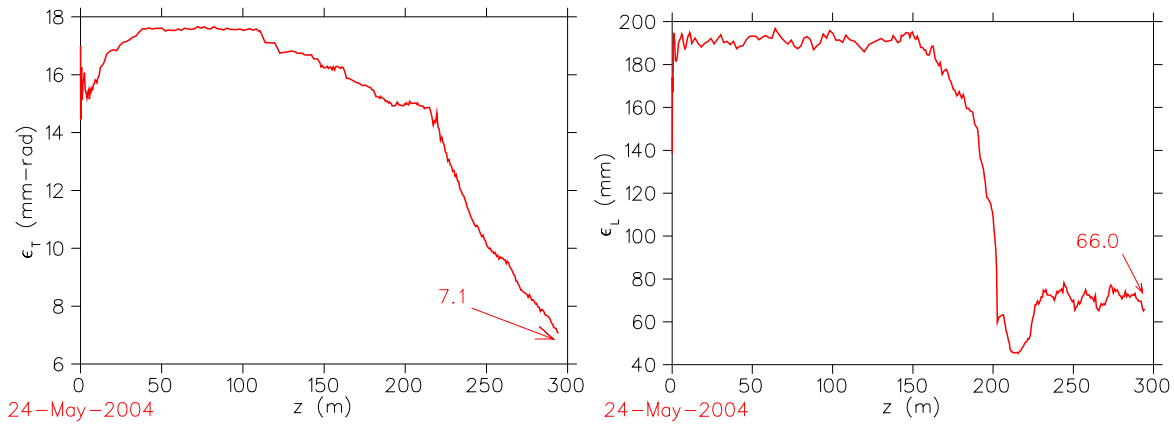

FIG. 14: (Color) Normalized transverse emittance (left) and longitudinal emittance (right) along the front-end for a momentum cut $0.1 \leq p \leq 0.3$ GeV/c.

0.176 ± 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 (see Sec. VII). 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 and the storage ring itself, 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. 16 for one sign.

Figure 17 depicts the longitudinal phase space of the superposition of all bunches projected onto a single period ($T \approx 5$ ns). We assume that particles outside the accelerator acceptance are intercepted by collimators located in the matching section, although this collimation system has not been designed yet. Approximately 60% of the beam leaving the cooling channel is intercepted by the collimators. The heat load on the collimator is 7.3 kW. Fig. 18 shows a few interleaved μ^+ and μ^- bunches exiting the cooling section. The opposite-sign bunches are mostly separated in time. There are a small number of wrong sign particles in the bunch train after cooling, but these will be cleanly separated by the dipoles in the subsequent accelerators and storage ring.

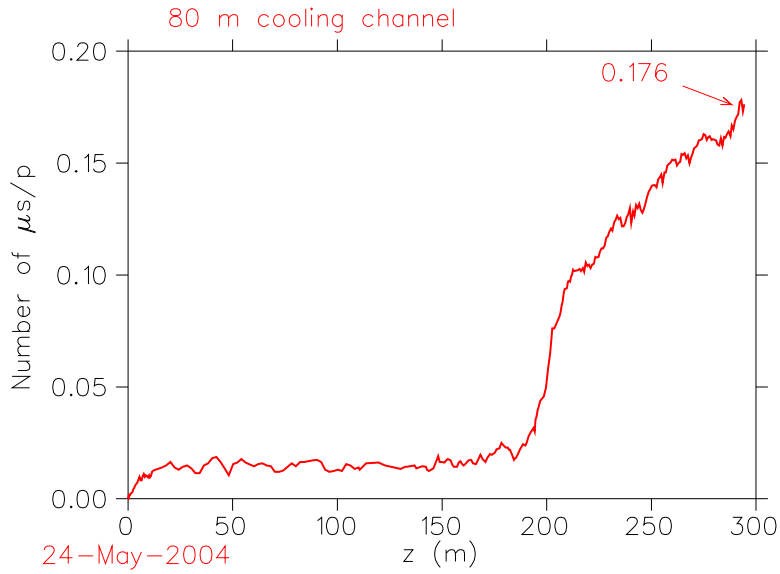

FIG. 15: (Color) The muons per incident proton on target into the accelerator normalized transverse acceptance of $A_T = 30$ mm rad and normalized longitudinal acceptance of $A_L = 150$ mm for a momentum cut $0.1 \leq p \leq 0.3$ GeV/c.

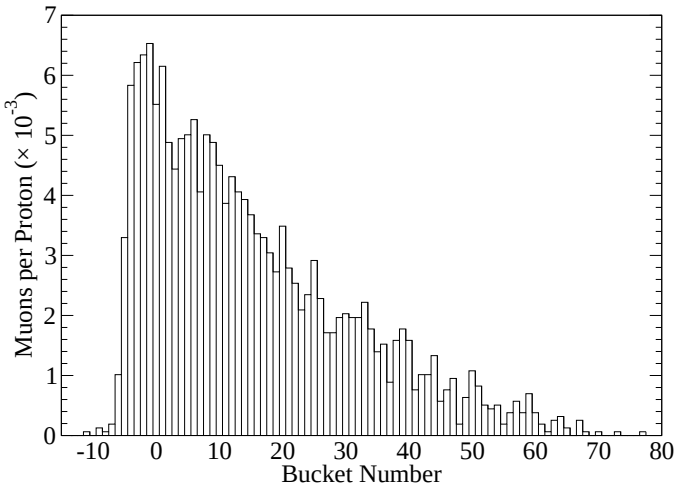

FIG. 16: Bunch structure of the beam delivered to the accelerator normalized transverse acceptance of $A_T = 30$ mm rad and normalized longitudinal acceptance of $A_L = 150$ mm for a momentum cut $0.1 \leq p \leq 0.3$ GeV/c.

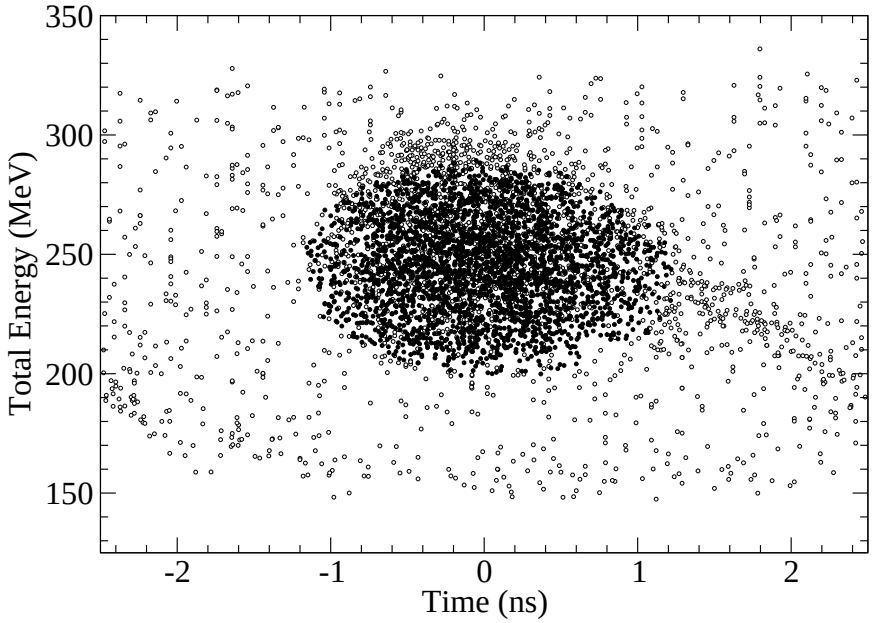

FIG. 17: Longitudinal phase space of overlayed bunches in the train at the end of the cooling section. The open circles are all the particles that reach the end of the channel and the filled circles are particles within the accelerator normalized transverse acceptance of $A_T = 30$ mm rad and normalized longitudinal acceptance of $A_L = 150$ mm for a momentum cut $0.1 \leq p \leq 0.3$ GeV/c.

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 bonded to the LiH, there is the question of what happens 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 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

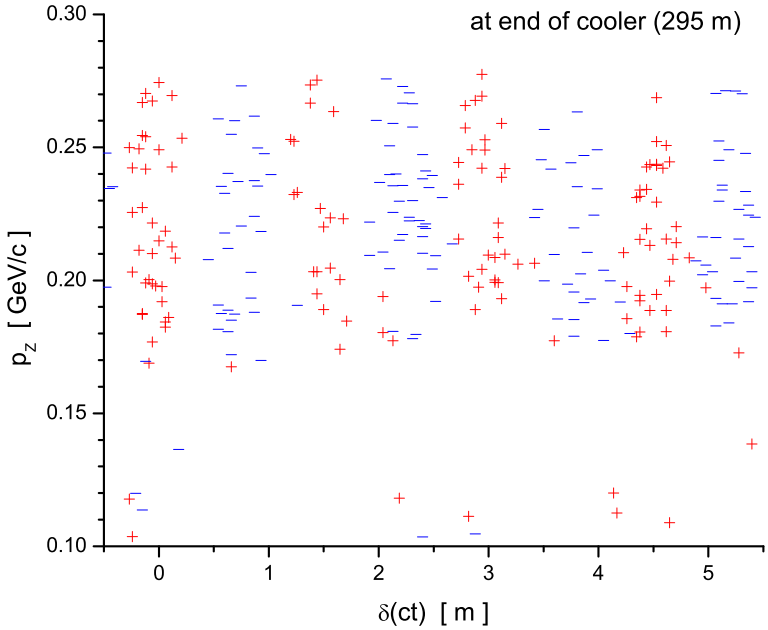

FIG. 18: (Color) A sample from the train of interleaved μ^+ (red) and μ^- (blue) bunches exiting the cooling section.

between the LiH and Be layers [20], the LiH reaches a maximum temperature of 310°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, λ = 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 $\approx$ 220 W in each window. Rough calculations predict that the temperature at the center of the 300 μm thick Be layer should be less than 175 °C. This is also 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. In case this window design does not prove to be feasible, a number of alternative absorber designs have been investigated. The window could be cooled

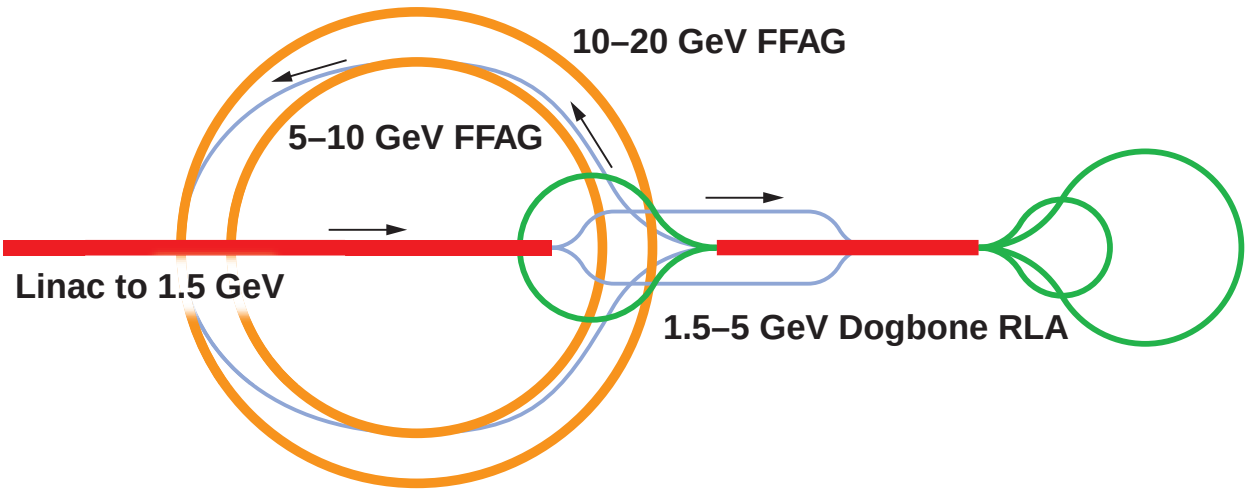

FIG. 19: (Color) Potential layout for the acceleration systems.

by flowing He gas between the Be window and the LiH. 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 possibilities that gave reasonable muon fluxes are a thin Be layer on pure lithium or thicker Be windows and no LiH, with a thickness chosen to make the total energy loss the same as that in the baseline LiH absorber case. Cooling would be a bit less effective because of the greater multiple scattering. An initial evaluation [18] of a Be-only scenario showed 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.

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. Figure 19 shows a compact potential layout for all the acceleration systems described here. It includes five sub-systems: a matching section, a linac, a Recirculating Linear Accelerator (RLA), and two Fixed-Field Alternating-Gradient (FFAG) circular accelerators.

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

TABLE III: 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 transverse emittance, rms (mm-rad)	3.84
Normalized longitudinal acceptance, $\Delta E \Delta t / m_\mu c$ (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 (each charge)	3.0×10^{12}
Bunches in train	89
Average repetition rate (Hz)	15
Minimum time between pulses (ms)	20
Average beam power at the end (each charge) (kW)	144

that accelerate a beam over a large energy range (generally at least a factor of 2) without varying the magnets' fields, allowing for very rapid acceleration. Since they are rings, the bunches make multiple passes through the RF cavities, reducing the RF voltage required to accelerate. The number of turns is not limited by the switchyard, as it is in an RLA. The original FFAG designs [21] ("scaling" FFAGs) used large, highly nonlinear magnets. For our design, we instead use so-called linear non-scaling FFAGs [22, 23]. These FFAGs use very linear magnets to maximize the dynamic aperture (necessary for our large-emittance beams), and the magnets generally have smaller apertures than those in a corresponding scaling FFAG design, bringing down the machine cost.

Table III 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 $\sqrt{\beta_\perp A_\perp mc/p}$,

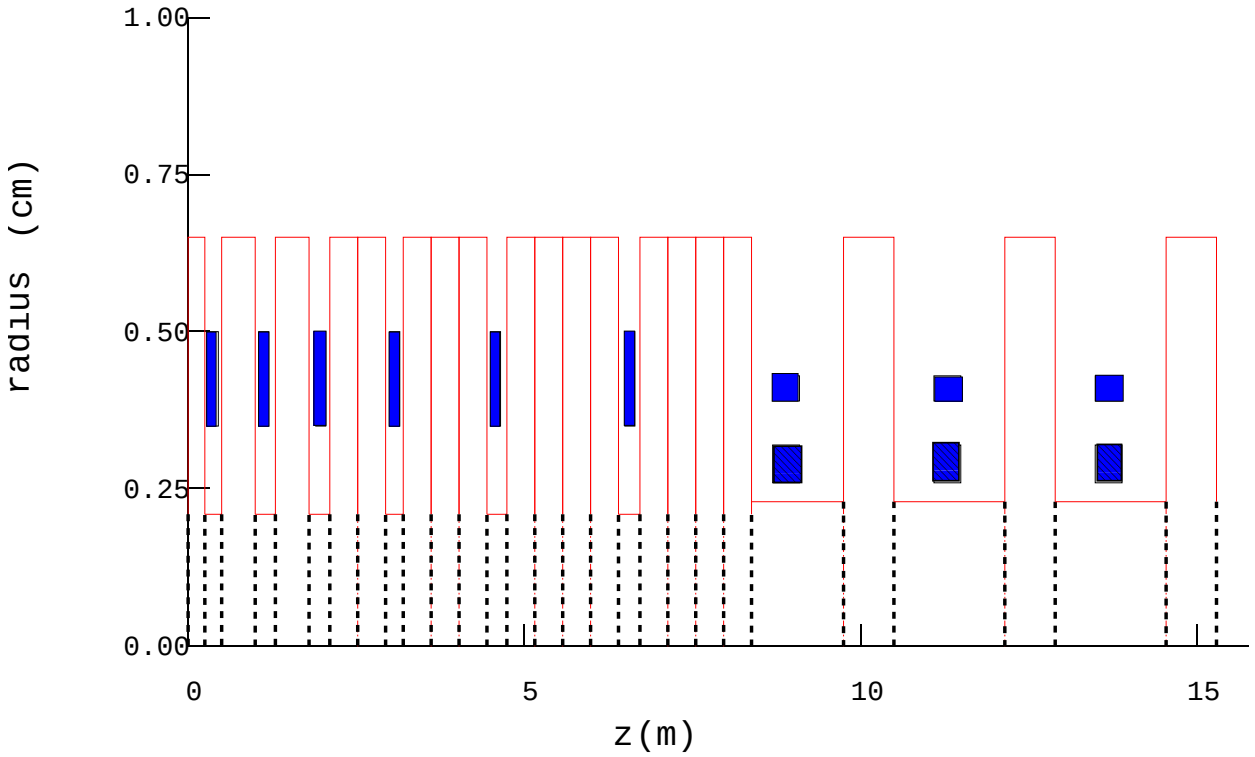

FIG. 20: (Color) Matching section from cooling to linac. Blue rectangles are solenoids, red lines are rf cavity walls and dashed black lines are the windows of the cavities. The last three cavities are superconducting, the remaining ones are room temperature.

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. The time of flight refers to the length of a single bunch.

A. Matching from Cooling to Acceleration Linac

The cooling section has a beta function of around 0.8 m, whereas the beginning of the acceleration linac has a beta function of around 2.7 m. A matching section is required to gradually change the beta functions from one section to the other so as to avoid emittance growth and/or particle loss. Furthermore, the reduced acceptance of the longer cells in the acceleration linac, as compared to the more compact cells of the cooling section, necessitates that the acceleration linac start at an energy above that of the cooling section (see Table III); the matching section will thus also begin to increase the beam energy after the cooling. That matching section will consist of six cells similar to those in the cooling channel, but with

increasing lengths and numbers of cavities per cell, and three superconducting cells similar to the accelerating linac, but made shorter by the use of shorter, higher field focusing solenoids. Figure 20 shows a layout of the matching section. The current design for the matching section has about 15% loss; initial 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 longitudinally, but we have not yet done the matching for the transverse plane as well.

B. Low Energy Acceleration

Based on preliminary 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 parameters are strongly constrained by the transverse acceptance. In FS2 there were three types of cryomodules, containing one, two, and four two-cell cavities, respectively. Because of our larger acceptance requirements, the cryomodule-dimensions 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 the momentum for the first stage of the linac is, already, 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 employ a very short cryomodule using a single one-cell cavity as opposed to the two-cell cavities used in all of 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 now consider it acceptable to run the cavities with up to 0.1 T on them [24], 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 reduce the total length for the first module type to only 3 m. Table IV shows the dimensions of the cryostats we will use and Fig. 21 depicts all three of them.

TABLE IV: 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

TABLE V: 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

Table V summarizes parameters for the linac. The phases 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. 22.

C. RLAs

Compared with FS2, we are injecting into the RLA at a lower energy and are accelerating over a much smaller energy range. These features 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 [25].

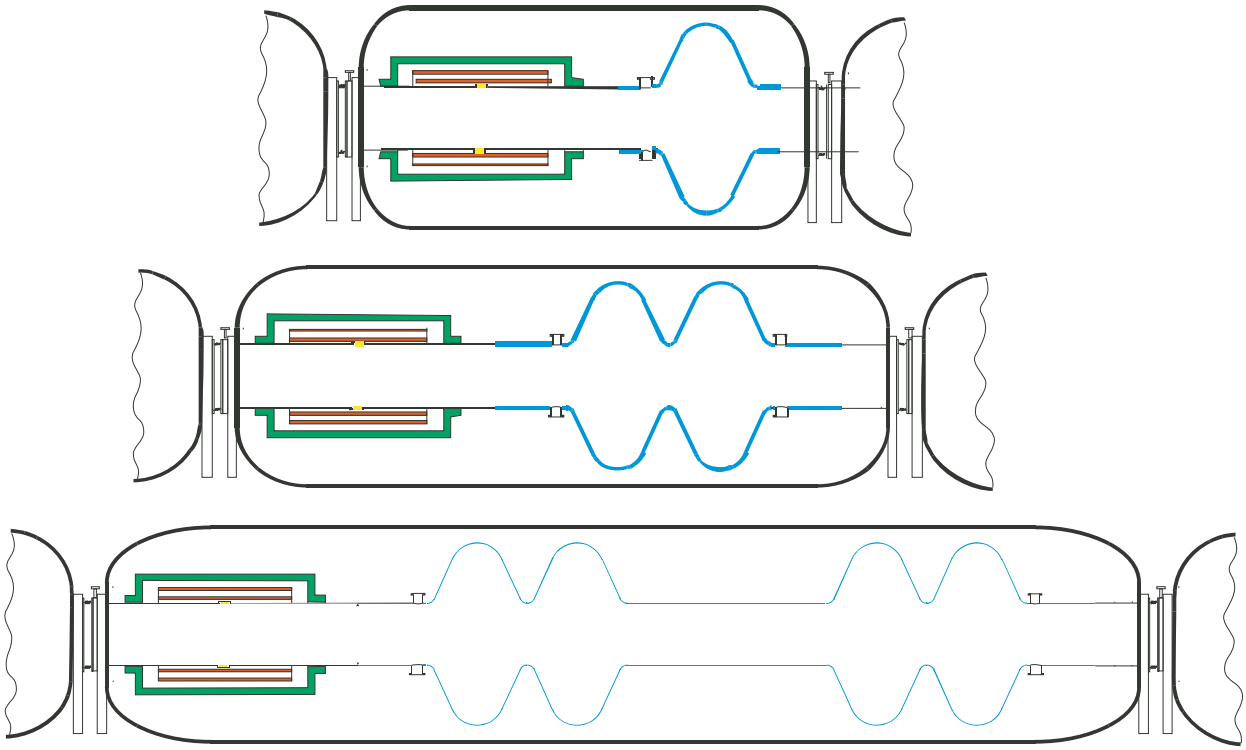

FIG. 21: (Color) Layouts of superconducting linac pre-accelerator cryomodules. Blue lines are the SC walls of the cavities, solenoid coils are indicated in red, and the iron shielding is in green. The dimensions of the cryomodules are shown in Table IV, and Table V summarizes parameters for the linac.

For a given amount of installed rf, the dogbone layout has twice the energy separation of the racetrack layout at the spreaders and recombiners (see Fig. 24), 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, relative to later passes. To reduce this effect, we inject into the center of the linac as shown in Fig. 23. This injection is accomplished with a chicane similar to that used for injection in FS2, but here, to inject both signs, there are two chicanes, one on either side of the linac (see Fig. 25). The start of the chicanes is the point at which the particles with differing charges are first separated. To avoid this point overlapping the earlier part of the linac, the chicanes are tilted slightly upwards.

In the dogbone RLA we have just over 1 GeV of linac, and we make three and a half

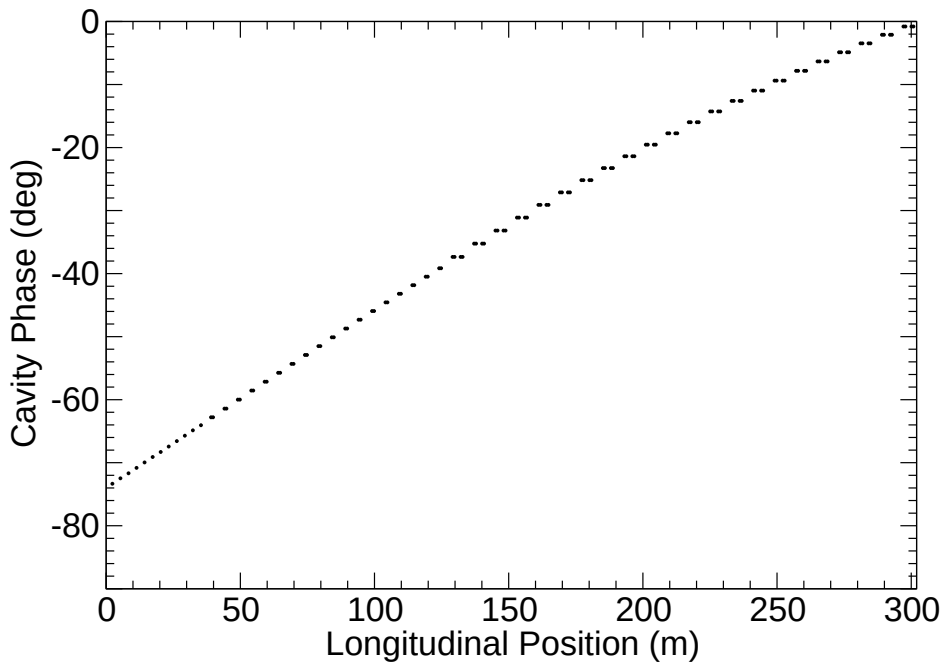

FIG. 22: Cavity phase vs. position along the linac.

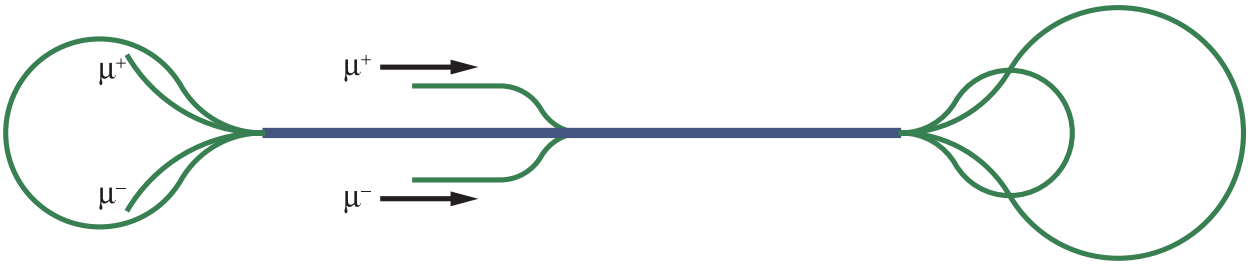

FIG. 23: (Color) Dogbone (top) and racetrack (bottom) layout for the RLA.

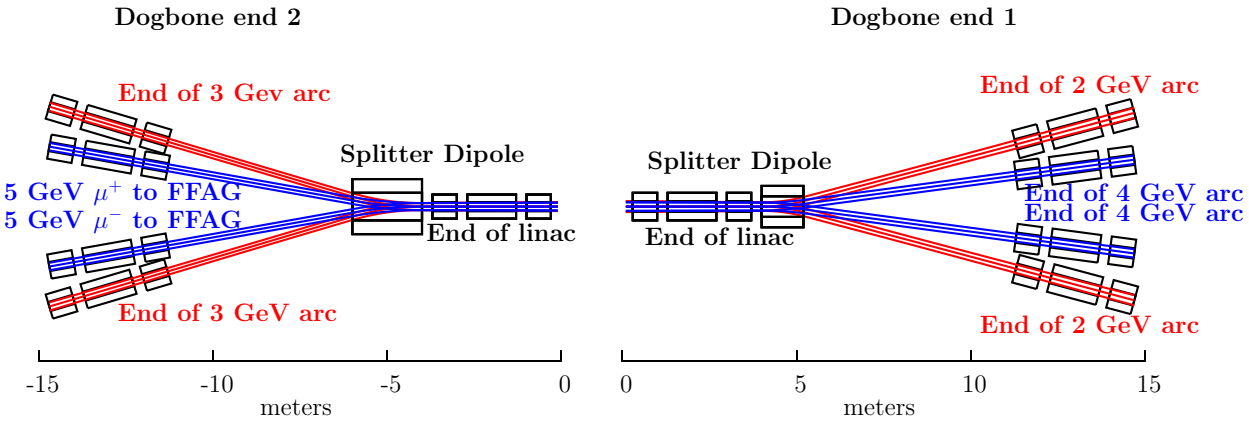

FIG. 24: (Color) Dogbone RLA Switch Yards.

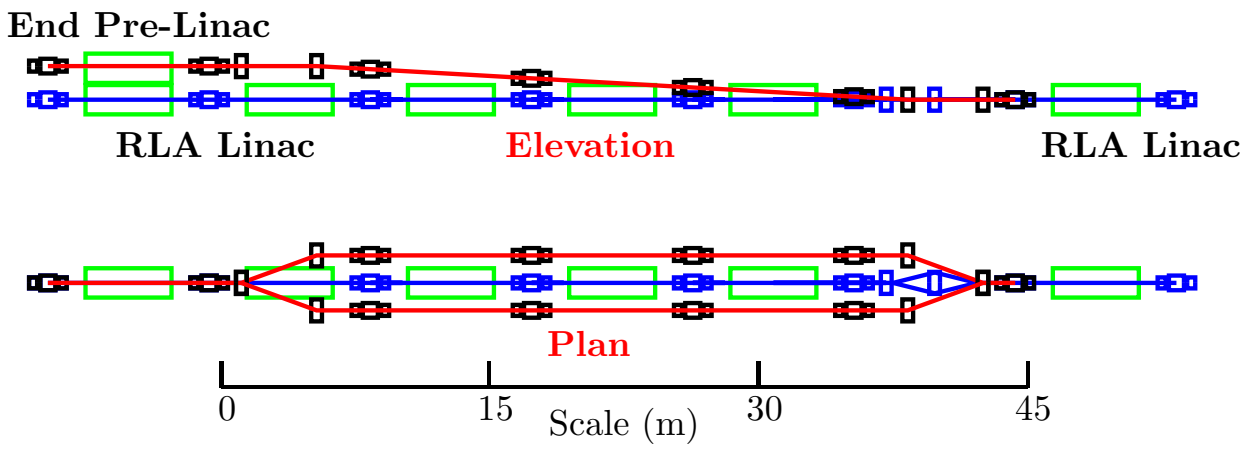

FIG. 25: (Color) Injection chicanes. Magnets in black, rf in green, RLA linac in blue, and injection lines in red.

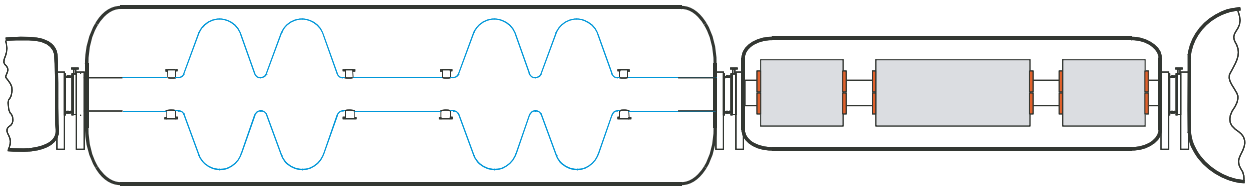

FIG. 26: (Color) Dogbone linac cell.

passes through that linac to accelerate from a total energy of 1.5 GeV to 5 GeV. The RLA linac consists of 11 m long cells with two 2-cell cavities per cell, and quadrupole triplet focusing, as shown in Fig. 26. 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 permit a somewhat higher 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 can be warm magnets.

Since the dogbone arc changes its direction of bend twice in each arc, dispersion matching must be handled carefully. This is done straightforwardly 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. 27. Matching of off-momentum particles is controlled using sextupoles.

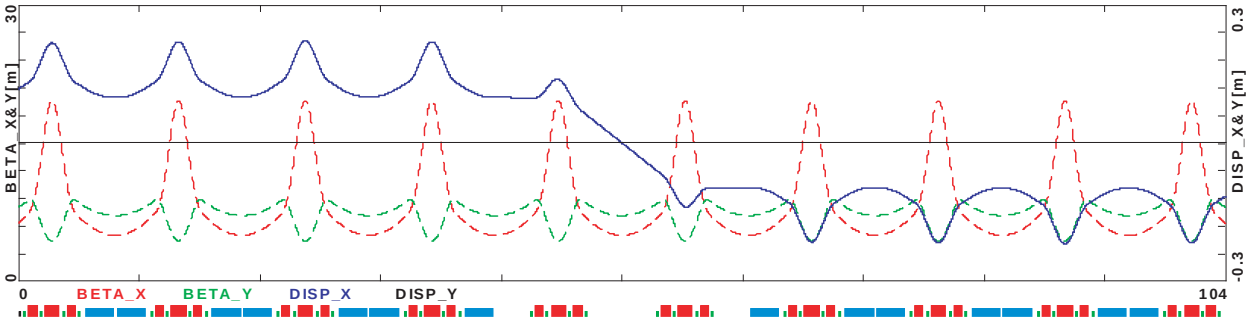

FIG. 27: (Color) A section of the dogbone arc where the bend changes direction, showing the dispersion (solid) and beta functions (dashed).

TABLE VI: Parameters for FFAG lattices. See Fig. 28 to understand the signs of the parameters.

Maximum energy gain per cavity (MeV)	7.5			
Stored energy per cavity (J)	368			
Cells without cavities	8			
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	90		105	
Magnet type	Defocusing	Focusing	Defocusing	Focusing
Magnet length (m)	1.612338	1.065600	1.762347	1.275747
Reference orbit radius of curvature (m)	15.2740	-59.6174	18.4002	- 70.9958
Magnet center offset from reference orbit (mm)	-1.573	7.667	1.148	8.745
Magnet aperture radius (cm)	14.0916	15.2628	10.3756	12.6256
Field on reference orbit (T)	1.63774	-0.41959	2.71917	-0.70474
Field gradient (T/m)	-9.1883	8.1768	-15.4948	12.5874

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 [26, 27] and comparing the resulting costs to those from FS2. FFAG designs for

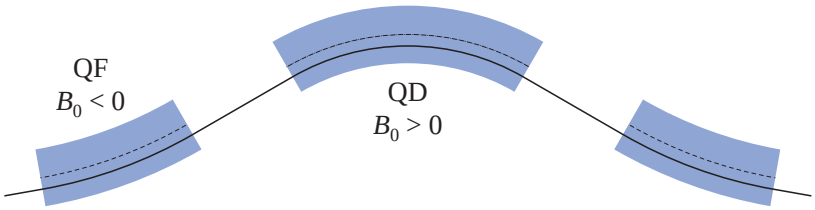

FIG. 28: (Color) Geometry of the Triplet Lattice. The "magnet center offset from reference orbit" listed in Table VI is positive for both magnets in this diagram.

accelerating from 5 to 10 GeV and from 10 to 20 GeV are given in Table VI. The lattices consist exclusively of combined-function triplet cells with a 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 [24].

With the 1 MW beam intensity given in Table III, and both signs of muons, about 16% of the stored energy will be extracted from the cavities in the 5–10 GeV FFAG, and about 27% 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, 7.5 MV, one need only to increase the initial voltage to 7.8 MV for the 5–10 GeV FFAG and to 8.1 MV for 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 1% 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 VII 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 also similar to those in induction linacs. These parameters assume that injection occurs from the inside of the FFAG. This is preferred since the beam will be near the inside of the FFAG at the lowest energies. Figures 29 and 30 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.

TABLE VII: Parameters for FFAG injection and extracting kickers.

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

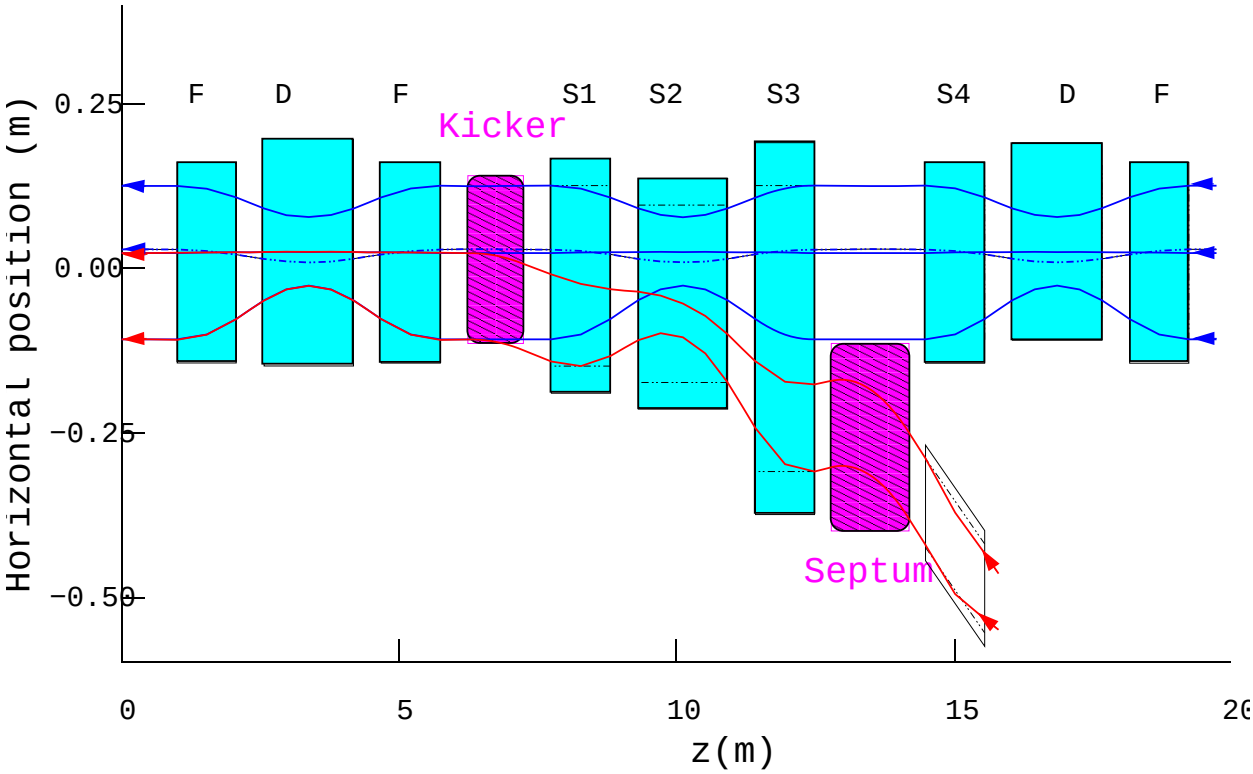

FIG. 29: (Color) Injection into the FFAG. S1 to S4 are special injection region dipoles.

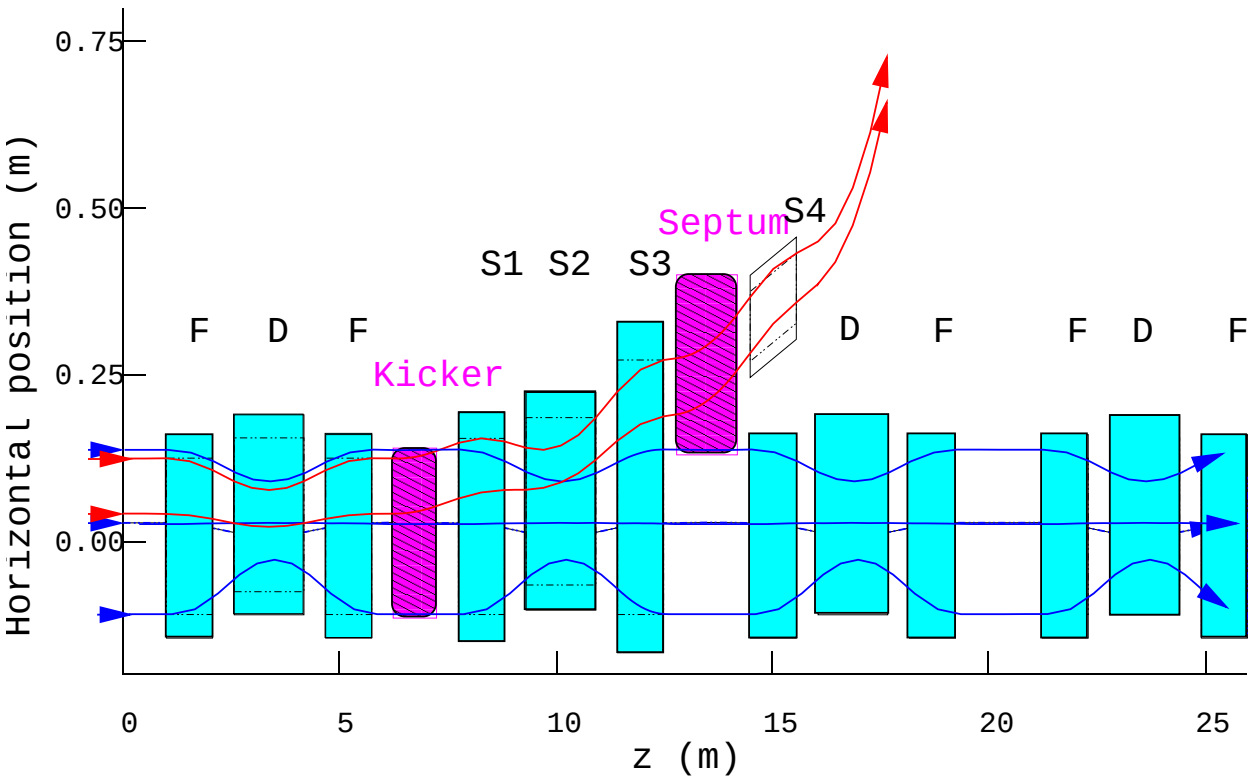

FIG. 30: (Color) Extraction out of the FFAG. S1 to S4 are special extraction region dipoles.

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 [15] is used for tracking for several reasons. It will allow for a fairly arbitrary end-field description, it forces that description to be consistent with Maxwell's equations, and it will track accurately even 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

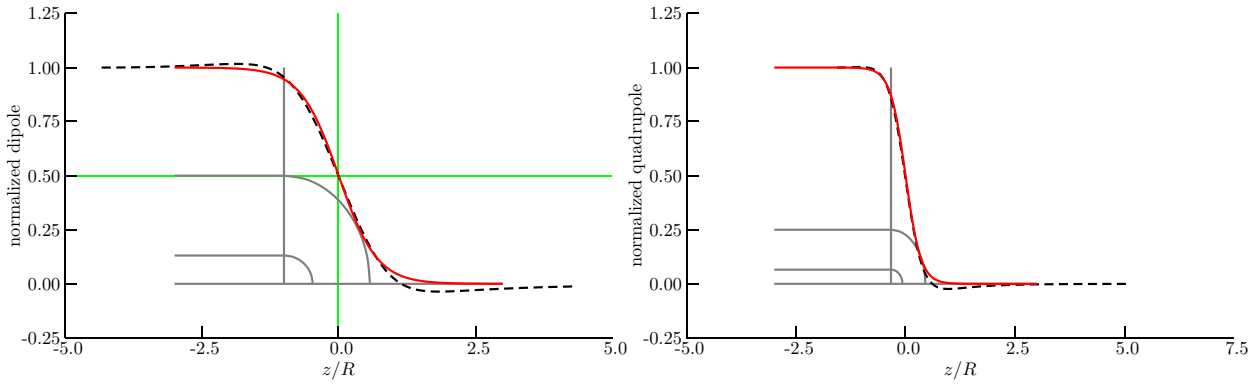

FIG. 31: (Color) Relative dipole field (left) and quadrupole field (right) near the magnet end. The dashed line is the field from TOSCA, while the solid line is our model.

magnet, without iron, using TOSCA [28]. At the end of the magnet, the field does not immediately drop to zero, but falls gradually, as shown in Fig. 31. The end-field falloff in a dipole or a 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. 32, and include them in our computation.

The TOSCA computation is done without iron, which leads to the overshoot in the field values in Figs. 31 and 32. 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

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

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 red in their corresponding plots in Figs. 31 and 32.

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. 33. This 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

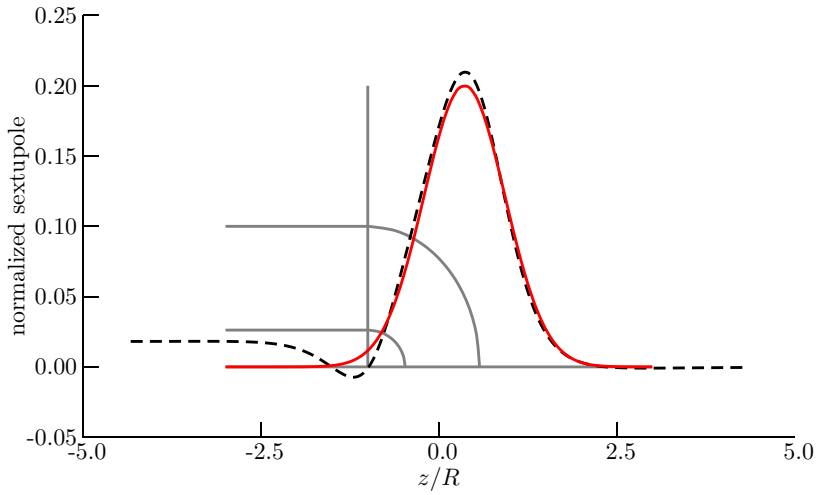

FIG. 32: (Color) Peak magnitude of the sextupole end field at radius R (the magnet aperture), divided by the dipole field. The dashed line is the field from TOSCA, while the solid line is our model.

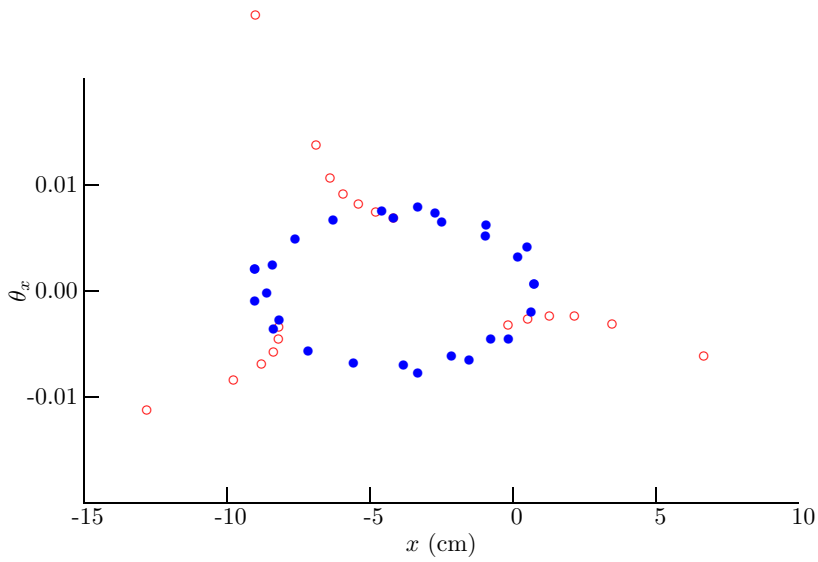

FIG. 33: (Color) Horizontal phase space from tracking at 5.1 GeV/c at the outer edge of the acceptance. Open circles are without the body sextupole fields and show a third-order resonance; filled circles are with the body sextupole fields.

integrated end sextupoles, (see Fig. 34), the resonance was eliminated (also shown in Fig. 33). 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 no loss with the body correction included (see Fig. 35). If the body correction is only partially

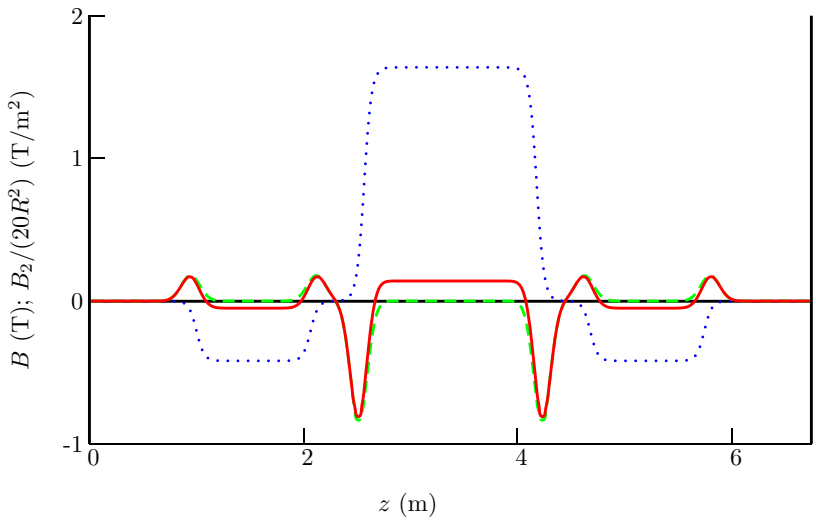

FIG. 34: (Color) Sextupole field components in the lattice used for tracking. The dotted line is the dipole field, the dashed line is $B_2/(20R^2)$ with zero body sextupole field, and the solid line is with sufficient body sextupole field to eliminate the third-order resonance.

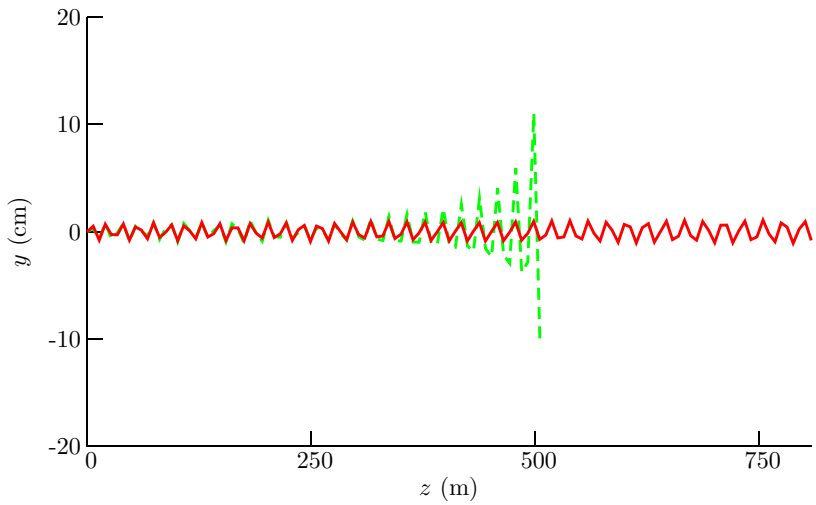

FIG. 35: (Color) Tracking of a particle at the edge of the acceptance with uniform acceleration. The dashed line is without any body sextupole, and the solid line is with the corrected body sextupole.

included, there is significant emittance growth. 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 is complex [29]. If one be-

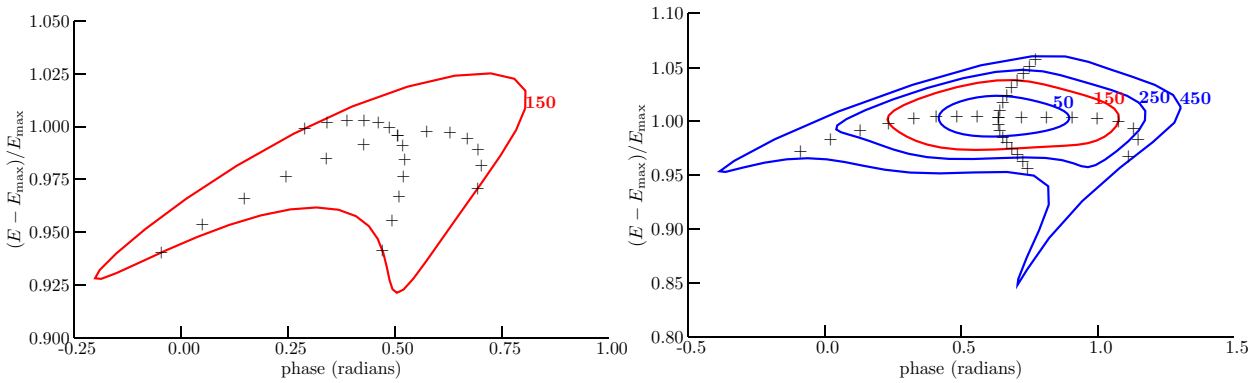

FIG. 36: (Color) Longitudinal tracking starting from an upright ellipse for the 5–10 GeV FFAG. On the left with only 201.25 MHz rf. On the right with third-harmonic rf having voltage equal to 2/9 of the fundamental rf voltage. Curves are labeled with their corresponding acceptance. Crosses for both cases started out as horizontal and vertical lines in phase space.

gins with an upright ellipse, there is considerable emittance growth if only the 201.25 MHz rf is used (see Fig. 36). Adding a third-harmonic rf considerably reduces the emittance growth, as shown in Fig. 36. 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.

F. Design of Combined-Function Superconducting Magnet for FFAGs

An initial design of a superconducting combined-function (dipole–quadrupole) magnet has been developed [30]. The work has been done for the defocusing magnet from the above design. The parameters of this QD combined-function magnet are shown in Table VIII.

The magnet design is based on a cosine-theta configuration with two double layers for each function. A cross section for one quadrant is shown in Fig. 37. 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

TABLE VIII: Parameters of the QD magnet: 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 (see Fig. 28); 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, $E_{\rm min}$ (GeV)	10
Final energy, $E_{\rm max}$ (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

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 IX.

TABLE IX: 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

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.

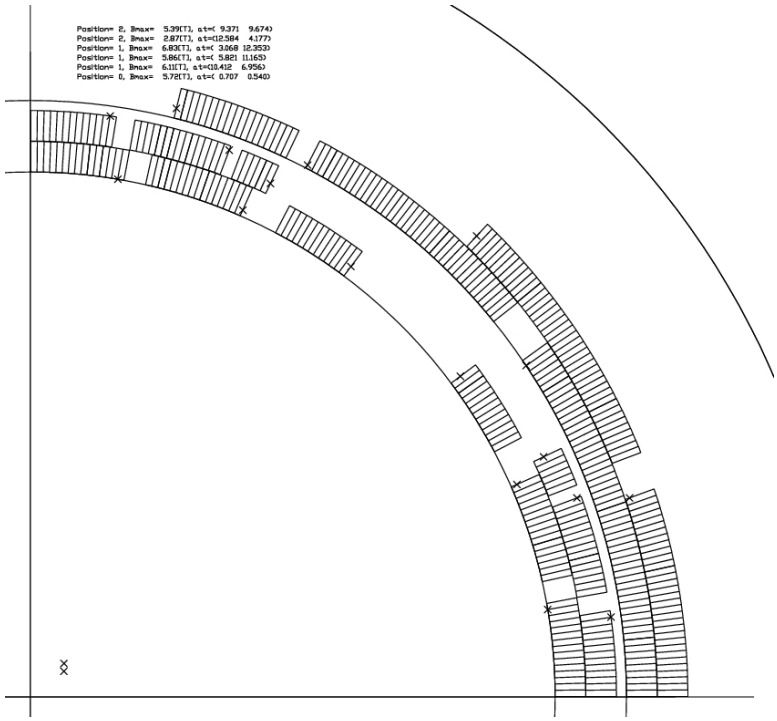

FIG. 37: First quadrant of the combined-function magnet cross section.

An alternative concept would be to use a single dipole-like design with laterally displaced poles (see Fig. 38) as discussed in Refs. [31, 32, 33].

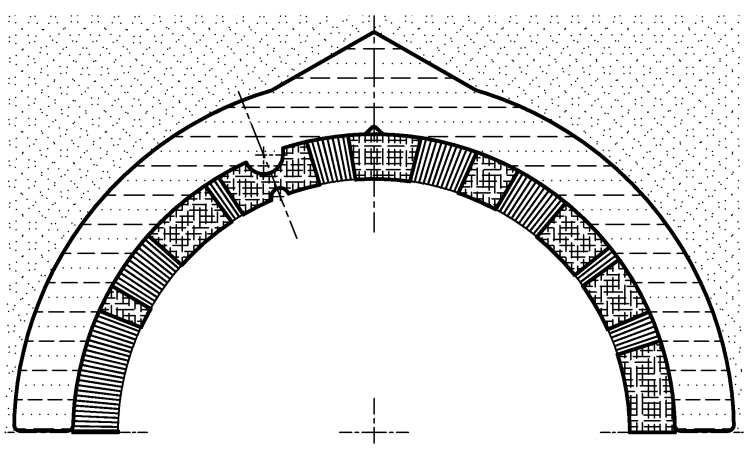

FIG. 38: Expanded view of the cross section of the superconducting combined function magnet used in the 50 GeV proton beamline for the J-PARC neutrino experiment [31].

V. MUON STORAGE RING AND PERFORMANCE

The storage ring in this study is assumed to be essentially identical to that in FS2. However, injection will be required in two opposite directions for the two differing signs. The injection lines must be designed such that when the train of one sign is traveling towards the detector, the train of the other sign is moving away from the detector. In this way the neutrinos of opposite kind arrive at well separated times, and the experiment can analyze their reactions separately. Another difference is that in this study, both straight sections must be designed with very high betas, so that the neutrino beams of both types are well collimated. Finally, it must be noted that the total energy deposited in the ring is doubled by the presence of equally intense muon beams, but now of two signs.

Losses are summarized in table X. We define η to be the probability that a muon makes it successfully into the storage ring. The number of decays N_μ, of each sign, injected into the storage ring in a 10^7 second year is given by:

$$N_\mu = 10^7 \ f \ N_p \ \mu/p \ \eta \approx 10^7 \ \times \ 15 \ \times \ (17 \times 10^{12}) \ \times \ 0.17 \ \times \ 0.67 \approx 2.9 \times 10^{20}.$$

We define η_{straight} to be the length of the straight section pointing to the detector divided by the circumference of the storage ring. The number of decays N_ν, of both signs, in the storage ring, decaying towards the detector, in a 10^7 second year is given by:

$$N_\nu \ = 2 \ N_\mu \ \eta_{\text{straight}} \approx 2 \times 2.9 \times 10^{20} \times 0.35 \approx 2.0 \times 10^{20}.$$

This is a factor of two greater than that reported in FS2. Note that if the proton driver power could be raised to 5 MW (4 MW has been discussed in a further upgrade of the BNL AGS), then the number of neutrinos per year would match the high performance goal (10^{21}) suggested at the first NuFact Workshop in Lyon, France [34].

VI. REQUIRED R&D

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.

TABLE X: Particle losses from cooling to storage ring

	loss %
Match into linac	15
Linac	6.0
RLA	6.1
FFAG 1	5.1
FFAG 2	6.5
Total losses	33

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 in substantial magnetic fields. Observations of breakdown in 805 MHz cavities have shown [35] serious reductions of attained rf gradients when the cavity is operated in a field. It is not clear how to scale these observations to the 201.25 MHz case, so experimental tests are urgently needed. If the required gradients cannot be achieved in the specified magnetic fields, then significant redesign will be needed.

The acceleration section demands high gradients from superconducting rf (SCRF) cavities at this frequency; our requirements are somewhat beyond the performance reached to date for cavities in this frequency range [36].

Development and testing of efficient high-power rf sources at a frequency near 200 MHz is also needed.

The ability of the target to withstand a proton beam power at 1 MW and above 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. COST ESTIMATE: ASSUMPTIONS AND ALGORITHM

For this study (ST2B) substantial effort has been directed at simplifying the design and thus hopefully reducing the costs of three major components of a neutrino factory: phase rotation, cooling, and the higher energy part of the muon acceleration. A preliminary comparison with FS2, which contained detailed cost estimates with significant engineering input, shows that a great deal of progress has been achieved. Starting from the FS2 work breakdown schedule, we derived element costs per unit length, integral rf voltage, or net acceleration. For all but the final FFAG acceleration, these costs were then applied to the ST2B parameters after scaling for magnetic fields, radii, stored energy, rf gradient, etc. For the FFAG costs a new cost algorithm had to be developed [27]. Further details on the costing algorithms and their application to this new design can be found in a recent reference [37].

There are a number of reasons why we believe this new design should be significantly less expensive than the previous one described in FS2. For phase rotation we have replaced 260 m of expensive induction linacs with 54 m of more conventional rf cavities. The 108 m long cooling channel in FS2 was replaced with an 80 m channel. The new cooling channel uses a simplified magnetic lattice with reduced peak solenoid fields; it also replaces liquid hydrogen absorbers with solid LiH absorbers. For acceleration to 20 GeV we have made use of recent advances in the development of non-scaling FFAG accelerators to replace most of the recirculating linear accelerator used in FS2. There are of course some changes in the new design, such as the adopted increase in transverse acceptance in the accelerators, which will increase the costs over FS2. However, our examination shows that the design changes should lead to an overall reduction in costs.

A summary of the preliminary estimates for the percentage cost reductions for the ST2B neutrino factory design is presented in Table XI. It is most likely that a proton driver will first be built in conjunction with a neutrino super-beam experiment, so we begin the neutrino factory systems with the target and capture section. Excluding the proton driver the new design should cost $\approx 35\%$ less than the FS2 design.

TABLE XI: Preliminary cost comparison between FS2 and ST2B.

System	Reduced Cost %
Target, capture, 18 m drift	98
Bunching and Phase Rotation	39
Cooling	60
Acceleration	77
Ring	100
Total without driver & controls	65

VIII. 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 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.

A substantial 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 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, which now captures both muon signs simultaneously. The non-scaling FFAG acceleration concept has been further developed and used for accelerating the muons up to the 20 GeV design energy. The performance of the new system equals that of the earlier FS2, for each of two neutrino states (ν and $\bar{\nu}$) that are generated essentially simultaneously. The performance is thus effectively twice that of FS2. At the same time, the facility is simpler than that in FS2 and of the order of 35% less costly.

R&D is also continuing to confirm needed component performance and establish the phys-

ical concepts used. Continued optimization is ongoing, and is expected to further improve performance and reduce the cost.

Acknowledgments

We would like to thank C. Johnstone, S. Geer, and Y. Fukui for helpful discussions. This research was supported by the U.S. Department of Energy under Contract Nos. DE-AC02-98CH10886, DE-AC02-76CH03000, and 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), also, D. Neuffer, *High-Frequency Buncher and Phase Rotation for the Muon Source*, MUC-NOTE-181 (2000). 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 and A. Van Ginneken, Proceedings of the 2001 Particle Accelerator Conference (2001), http://accelconf.web.cern.ch/Accel/Conf/p01/PAPERS/TPPH162.pdf.

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

[15] 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.

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

[17] 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/.

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

[19] A. Chao and M. Tigner, eds., *Handbook of Accelerator Physics and Engineering* (World Scientific, 1999).

[20] R. Fernow, *Heating in the study 2a absorber window*, MUC-NOTE-317 (2005).

[21] K. R. Symon, D. W. Kerst, L. W. Jones, L. J. Laslett, and K. M. Terwilliger, Phys. Rev. **103**, 1837 (1956).

[22] C. Johnstone, W. Wan, and A. Garren, in *Proceedings of the 1999 Particle Accelerator Conference*, edited by A. Luccio and W. MacKay (IEEE, Piscataway, NJ, 1999), p. 3068.

[23] F. E. Mills and C. Johnstone, in the transparency book for the 4th International Conference on Physics Potential & Developent of μ^+ μ^- colliders, San Francisco, CA (1999), UCLA, Los Angeles, CA, pp. 693–698.

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

[25] 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.

[26] J. S. Berg, *Recent results from optimization studies of linear non-scaling FFAGs for muon acceleration*, MUC-CONF-ACCELERATION-309, Proceedings of FFAG 04, KEK, Tsukuba,

Japan, 13–16 October (2004), `http://www-mucool.fnal.gov/notes/noteSelMin.html`.

[27] R.B. Palmer, J.S. Berg, *A Model for Determining Dipole, Quadrupole, and Combined Function Magnet Costs*, Proceedings of EPAC (2004).

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

[29] S. Koscielniak and C. Johnstone, in [40], pp. 1831–1833.

[30] 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).

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

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

[33] 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).

[34] N. Autin, ed., Nucl. Instrum. & Meth. **A451** (2000), Proceedings of the ICFA/ECFA Workshop NUFACT'99: Neutrino Factories based on Muon Storage Rings, Lyon France.

[35] J. Norem, *et al.*, Phys. Rev. ST Accel. Beams **6**, 072001 (2003), also MUC-NOTE-226 (2001).

[36] R. L. Geng et al., in [40], pp. 1309–1311.

[37] R. Palmer and M. Zisman, *Estimation of Study IIB costs based on Study II*, MUC-NOTE-322 (2005).

[38] 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).

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

[40] J. Chew, P. Lucas, and S. Webber, eds., *Proceedings of the 2003 Particle Accelerator Conference* (IEEE, Piscataway, NJ, 2003).