

DRAFT

International Scoping Study of a Future Neutrino Factory and Superbeam Facility

Summary Report of Accelerator Working Group

Michael S. Zisman
for the
ISS Accelerator Working Group

March 31, 2007

1. Introduction

This document summarizes the findings of the Accelerator Working Group (AWG) of the International Scoping Study (ISS) of a Future Neutrino Factory and Superbeam Facility. The work of the group took place at three plenary meetings along with three workshops, and an oral summary report was presented at the NuFact06 workshop held at UC-Irvine in August, 2006. The goal was to reach consensus on a baseline design for a Neutrino Factory complex. One aspect of this endeavor was to examine critically the advantages and disadvantages of the various Neutrino Factory schemes that have been proposed in recent years. This comparison is discussed in Section 4.

The activities of the group were coordinated by an Accelerator Council, whose members are listed in Table 1-1. Initially, a series of issues and tasks was identified for each of the various subsystems that comprise a Neutrino Factory. Over the course of the one-year study, an attempt was made to address these. The status of this work is briefly summarized here. In addition, a list of required R&D activities was developed as a guide to future effort; its main items will be summarized in Section 7.

Table 1-1. Accelerator Council members.

R. Fernow	BNL
R. Garoby	CERN
Y. Mori	Kyoto University
R. Palmer	BNL
C. Prior	Oxford/ASTeC
M. Zisman, convener	LBNL

1.1 Issues addressed during the ISS

Proton Driver

- What is the optimum beam energy (which depends to some degree on the choice of target material)?
- What is the optimum repetition rate?
- What is the optimum bunch length?
- Is there a preferred hardware configuration (e.g., linac, synchrotron, FFAG ring,...)?

Target, Capture and Decay

- What is the optimum target material (high or low Z)?
- What are the target limitations on proton beam parameters at 4 MW (bunch intensity, bunch length, pulse duration, repetition rate)?
- How do superbeam and Neutrino Factory requirements compare?

DRAFT

Front End

- Compare existing schemes, both with and without cooling
- Consider effects of reduced operating specifications
- Examine trade-offs between cooling and downstream acceptance

Acceleration

- Compare alternative schemes (linac, RLA, FFAG) on equal footing
- Examine implications of increased acceptance
- Study dynamics and matching with errors

Decay Ring

- Design implications of final energy (20 vs. 50 GeV)
- Implications of keeping both muon signs
- Implications of two simultaneous baselines
- Optics requirements vs. emittance

1.2 Organization of report

In what follows, we discuss parameters and concepts for the proton driver (Section 2), the target (Section 3), the front end (Section 4, which also includes a performance comparison of the various designs that have been developed in recent years), the acceleration system (Section 5) and the decay ring (Section 6). Section 7 gives a description of the key elements of the R&D plan, many of which are already well under way. Section 8 provides a brief summary of what we have learned.

2. Proton Driver

2.1 Introduction

Many factors influence the specifications for the proton driver. Among these are the following:

- the required production of $\approx 10^{21}$ neutrinos per year
- muon yields as a function of the proton energy
- muon yields as a function of the target material
- heating and stress levels for the target material
- muon capture as a function of proton bunch length
- maximum acceptable duration of proton pulses on the target
- peak beam loading levels in the μ^\pm accelerators
- bunch train stacking in the μ^+ and μ^- decay rings

After considering all of these, the proton driver specifications for the ISS were set as indicated in Table 2-1. As can be seen, there are differing—and incompatible—

DRAFT

requirements for liquid-Hg and solid metallic targets. In this report, as discussed in Section 3, our baseline target choice is the liquid-Hg jet, so our efforts have focused mainly on those parameters in the designs described here. A solid target could be accommodated with some changes in design parameters.

Table 2-1. Proton driver requirements.

Parameter	Value
Average beam power (MW)	4
Pulse repetition frequency (Hz)	50
Proton energy (GeV)	10 ± 5
Proton rms bunch length (ns)	2 ± 1
No. of proton bunches	3 or 5
Sequential extraction delay (μ s)	≥ 17
Pulse duration, liquid-Hg target (μ s)	≤ 40
Pulse duration, solid target (μ s)	≤ 70

2.2 Proton Driver Options

A number of options were considered for a 4 MW, 50 Hz proton driver. These include:

- an H^- linac with a 50 Hz booster RCS and a 50 Hz NFFAG driver ring
- an H^- linac with pairs of 50 Hz booster and 25 Hz driver synchrotrons (RCS)
- an H^- linac with a chain of three non-isochronous FFAG rings in series
- an H^- linac with two slower cycling synchrotrons and two holding rings
- a full energy H^- linac with an accumulator and bunch compression ring(s)

Of these options, the most advanced design is for the first, and this is described here to give a sense of what a workable system must include. In the same spirit, the last option, using an energy at the low end of the desired range, is also briefly described in Section 2.5.

As can be seen in Table 2-1, we have chosen to specify proton driver performance requirements rather than a specific implementation. This is because there is not a unique solution to providing our requirements, and any solution that does so would be acceptable. The actual choice at a particular host site will undoubtedly be dictated by many factors, including cost, local expertise, and other possible uses of the proton driver complex. Thus, the examples here should be taken as indicative of possible approaches rather than as endorsements for a particular approach.

2.3 Proton and Muon Bunch Train Patterns

Proton bunch compression occurs in each 50 Hz cycle, with five bunches preferred.¹ To keep the pulse duration below 40 μs for the Hg-jet target, however, only three bunches can be used. Each proton bunch creates pions in the target, and these decay to give a single μ^\pm bunch, which is then transformed to a train of interleaved 80 μ^+ and 80 μ^- bunches in a bunch rotation scheme [2-1]. A schematic drawing of the proposed bunch patterns is given in Fig. 2.1.

For a uniform pattern in both decay rings, the three (397.5 ns) bunch trains are separated by 993.8 ns time gaps. If the rings are in a single tunnel, or if both beams are stored in a single ring, the μ^+ trains in one ring are interleaved in time with the μ^- trains in the other. In the example shown in Fig. 2.1, the time gaps between the μ^+ and μ^- bunch trains are 298 ns for the three trains. An RF system is needed in each ring to contain the 201.25 MHz bunches and preserve the gaps.

2.4 10 GeV, NFFAG Proton Driver Complex

Prior to the ISS, a 50 Hz, 10 GeV, 4 MW proton driver [2-2] was designed at RAL. During the study, the design was modified for a three (five), bunch compatibility between booster, driver and 20–50 GeV, muon decay rings.

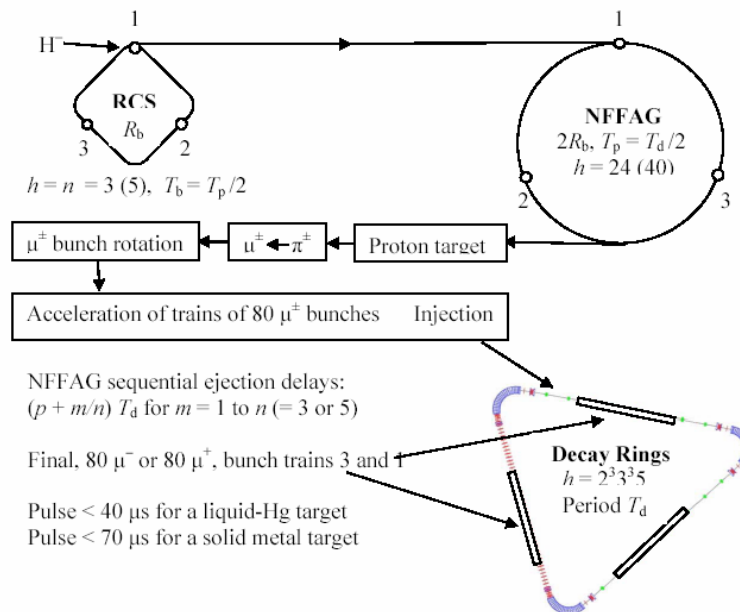


Fig. 2.1. Bunch patterns ($n = 3$) for the proton driver and decay rings. This example assumes triangular decay rings.

¹The use of the multi-bunch trains at 50 Hz is a change made during the study from the original single, 15-Hz train. The change was made to ease the production of the 2 ± 1 ns (rms) proton bunches, and to reduce the heavy beam loading in the μ^\pm accelerators.

DRAFT

The booster energy range is 0.2–3 GeV, and the proton driver energy range is 3–10 GeV. The driver is a new type of FFAG accelerator which uses a non-isochronous, non-linear, and non-scaling cell focusing structure (which we denote as “NFFAG”). Either three or five proton bunches may be used with the design.

A 90 MeV injector linac design at RAL [2-3] is extended to 200 MeV by adding a 110 MeV side coupled linac. The frequency is 324 MHz, and can be provided by commercially available klystrons. The pulse repetition frequency is 50 Hz, the peak current ~ 30 mA, the pulse duration ~ 400 μ s, and the duty cycle after chopping $\sim 70\%$.

A 50 Hz synchrotron is used for a booster. The injection scheme dominates the lattice design, which is based on that for a European Spallation Source (ESS) [2-4]. The straight sections are designed to be dispersion free. Ceramic vacuum chambers with contoured RF shields, based on designs used at the ISIS synchrotron [2-6], are required for the main and the correction magnets. There is a dedicated region for beam loss collimation in one ring superperiod. A momentum collimator protects ring components from longitudinal beam loss, and primary and secondary betatron collectors are used to localize the transverse beam losses in both planes.

A proton driver based on an NFFAG ring has several advantages:

- It can have a high duty cycle, and thus lower RF accelerating fields.
- Adiabatic compression is eased, as bunches may be held at the top energy of 10 GeV.
- It can utilize sturdy metallic vacuum chambers, in contrast with an RCS, which must have ceramic chambers with RF shields to limit the eddy currents.
- Single booster and driver rings and transfer lines can be used, saving cost.
- Low beam power loss during H^- injection and easier bunch compression compared with an option that uses a linac, an accumulator and a compressor ring [2-5].

The NFFAG cell layout is shown in the Figure 2.2. Three magnet types are used in the basic cell; the bd and BD units are non-linear, vertically focusing, parallel edged, combined function magnets, with bd and BD providing reverse and positive bending, respectively. The F magnet is a non-linear, horizontally focusing, positive bending, combined function magnet, whose edges are parallel to those of bd and BF. Beam loss collimation is a major design issue. The fractional loss in the collimators must be kept to about $1:10^3$, with that in the extraction region and elsewhere in the ring both less than $1:10^4$. Halo growth has thus to be limited. A layout of the complete proton driver system is shown in Fig. 2.3.

DRAFT




Figure 2.2: A single lattice cell of the 50 Hz, 4 MW, 10 GeV, NFFAG proton driver ring.

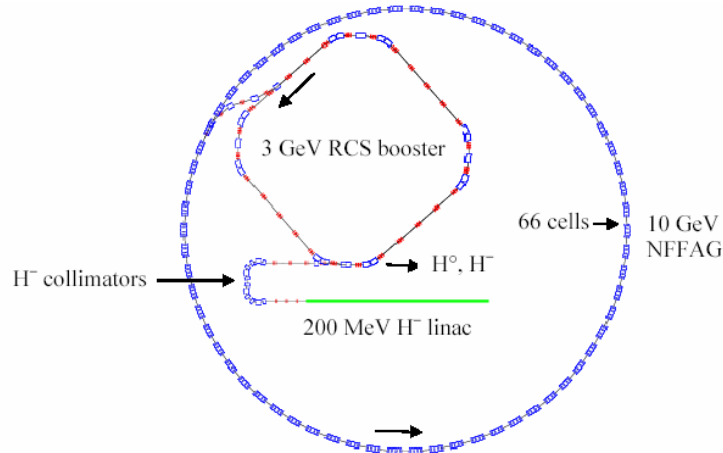


Figure 2.3: Schematic layout drawing of the linac, booster and NFFAG of the proton driver.

2.5 Linac Option

The low energy part of a proton accelerator complex always uses a linac. Above a certain kinetic energy, however, conventional setups make use of circular accelerators, which present a cost advantage because of their efficient use of RF systems. In the case of a multi-MW proton driver, that practice is worth reconsidering because of the need for fast acceleration (requiring fast cycling magnets, ceramic vacuum chambers and an expensive wide frequency range RF system providing a lot of voltage), which dramatically increases the cost of the circular accelerator. Today's linacs are capable of reliably providing tens of MW of beam power. For a Neutrino Factory however, fixed energy rings remain necessary to transform the long linac beam pulses into the required number of short bunches.

Linac-based proton drivers are being considered at FNAL [2-6] and at CERN. The CERN proposal [2-7] will be used as a typical example. The linac itself is a modified version of the 3.5 GeV Superconducting Proton Linac (SPL) whose Conceptual Design Report has recently been published [2-8]. Longer by 105 meters and equipped with 14 more 4 MW klystrons (for a total of 58), it accelerates protons up to 5 GeV in an overall length of only 534 m. Its structure is sketched in Fig. 2.4 and its main characteristics are shown in Table 2-2.

DRAFT

Table 2-2: Characteristics of the 5 GeV version of the SPL.

Ion species	H ⁻
Kinetic energy (GeV)	5
Beam power (MW)	4
Repetition rate (HZ)	50
Mean current during the pulse (mA)	40
Pulse duration (ms)	0.4
Bunch frequency (MHz)	352.2
Linac length (m)	534

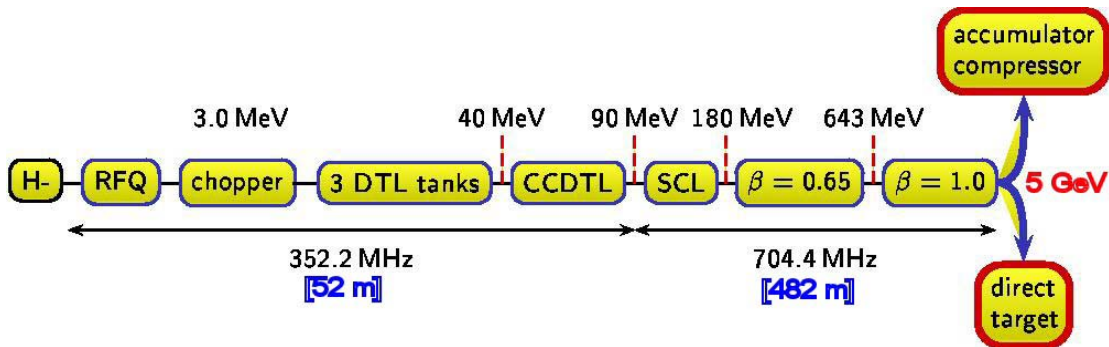


Figure 2.5-1: SPL block diagram.

Two fixed energy rings of approximately 300 m circumference are necessary to give the proton beam the required time structure for a Neutrino Factory. In the first one, the 400 μ s long linac beam pulse is accumulated using charge-exchange injection. Once accumulation is finished, bunches are transferred to a compressor ring where they are rotated in the longitudinal phase plane and ejected to the target when their length is minimum.

2.5.1 Superbeams

In the case of superbeams, the neutrino beam comes from the decay of pions and the muons themselves are not used. The infrastructure behind the target is thus limited to a focusing system and a decay tunnel. Therefore, every laboratory equipped with high energy proton accelerators has been the subject of superbeam proposals, and very different energies for the primary proton beam have been assumed, ranging from 3.5 to 400 GeV. The neutrino flux being directly proportional to beam power, the corresponding accelerators have to deliver a very high flux, sometimes well beyond what they were initially designed for. A representative list is given in Table 2-3.

Compared with a Neutrino Factory, the only requirements on the time structure of the proton beam come from the need to sufficiently reject background in the remote experiment. For that purpose, the duty factor of the proton beam has to be lower than 5×10^{-3} [2-9].

DRAFT

Table 2-3. Proposed superbeams.

	Proton beam energy (GeV)	Protons per pulse	Repetition period (s)	Beam power (MW)
CNGS+ [2-10]	400	$4.8-14 \times 10^{13}$	6	0.3-1.2
FNAL [2-11]	120	$9.5-15 \times 10^{13}$	1.5	1.1-2
JPARC [2-12]	50	33×10^{13}	3.64-1.6	0.6-1.5
BNL [2-13]	28	$9-25 \times 10^{13}$	0.4-0.1	1-4
FREJUS [2-14]	3.5	14.3×10^{13}	0.02	4

3. Target Issues

3.1 Beam Energy Choice

To determine the kinetic energy of the proton beam that is most efficient for the production of the soft pions, we process the produced pions through the entire front end of the Neutrino Factory front end using the Study 2a [3-1] configuration. As a figure of merit, we select surviving muons that are fully contained within the capture transverse acceptance (30 mm-rad) and the longitudinal acceptance (150 mm-rad) assumed for the subsequent accelerating section. The particle production model used was MARS V14 [3-2] and the propagation of the particles through the neutrino factory front end was done utilizing the ICOOL code [3-3]. The efficiency of the muon capture was computed by evaluating the number of collected muons at the end of the neutrino factory front end and normalizing the results to the power of the proton beam. Results utilizing a mercury-jet target are shown in Fig. 3.1. The target parameters such as radius, tilt angle, and longitudinal placement were previously optimized in Study 2a [3-1].




Fig. 3.1. Calculated production efficiency of positive and negative muons at the end of the Study 2a cooling channel, per proton and per GeV of proton beam energy, for a mercury-jet target. Although the curves are rather flat, an optimum energy, roughly 10 GeV, is discernible. Below about 5 GeV, the calculations show an abrupt fall-off in production. Above 10 GeV the fall-off is small but, from the muon production perspective, there is no benefit to increasing the beam energy beyond 10 GeV.

3.2 Choice of Target Material

We also investigated other candidate target types. Figure 3.2, shows an efficiency plot for a carbon target. Here, the optimal proton kinetic energy is centered around 5 GeV, somewhat lower than the case for mercury. As can be seen from comparing the two figures, the high-Z material shows the higher efficiency for soft-pion production, which will lead to the greatest number of captured muons. In evaluating the most efficient kinetic energy region for a mercury target, we find that 6–38 GeV protons give a sum of positive and negative pions within 10% of the maximum efficiency at 10 GeV.

3.3 Proton Beam Structure

3.3.1 Repetition Rate

For a given proton driver power, an increased repetition rate will lower the stress on the target (especially for a solid target) since the intensity per pulse is decreased. For the same pulse intensity and increased repetition rate, the proton driver power increases with a concomitant increase in stress on the target.² The primary downside of a higher repetition rate is the increased average power consumption of the RF systems. In Study II [3-4], the average power required for




Fig. 3-2. Calculated production efficiency for a carbon target. The yield per proton and per GeV is lower than for a mercury target and peaks at a lower energy.

²That is, the thermal load of each pulse on the target must be removed by the heat sink in a shorter time and the repetition rate will be limited by the ability to remove the dynamic stresses entirely between pulses.

DRAFT

these systems was 44 MW for a 15 Hz average repetition rate. This portion of the machine's power consumption will be proportional to the repetition rate.

Higher repetition rates reduce the amount of current per bunch train, which in turn reduces the beam loading in the RF cavities. Furthermore, some schemes for the storage ring require (superconducting) RF cavities to keep the beam bunched, and higher currents would require more RF power (and possibly more cavities) to compensate for beam loading there.

3.3.2 Pulse Length, Intensity and Structure

The pulse intensity, combined with the beam spot size, controls the quasi-static conditions of pressure and temperature generated in the target. Energy densities of up to 400 J/g, corresponding to $\sim 24 \times 10^{12}$ protons per pulse and $\sigma_r = 1$ mm, may be tolerated by some high performance solid materials. The pulse length controls the ensuing dynamic stresses and can play a significant role in determining whether a solid target survives the induced shock. Solid targets favor longer pulses because of the ability to relax during deposition. On the other hand, liquid-jet targets perform best at very short pulse lengths (a few ns), as the onset of jet destruction occurs much later. At the same intensity, a pulse having a uniform distribution over the same area as a Gaussian pulse (i.e., a 3σ spot) will reduce the stress and temperature in the target by approximately a factor of three.

3.3.3 Bunch Length

The proton bunch length has a strong influence on the usable muon intensity. The accepted muon density at the end of the cooling channel falls off with increasing proton driver bunch length on the target. This behavior can be partially understood by a simple theory that models the longitudinal dynamics of the muon beam through the RF components of the front end. Longer proton bunches produce initial longitudinal phase space areas that exceed the longitudinal acceptance of the front end.

4. Front End

The parts of the Neutrino Factory between the target and the beginning of the acceleration system are designated collectively as the *front end* [4-1]. There are two main requirements on the operation of the front end. First it has to collect the pions created in the target and form a beam from their daughter muons as efficiently as possible. Second it has to manipulate the transverse and longitudinal phase space of the muon beam so that it matches the accelerator acceptance as efficiently as possible. The Neutrino Factory front end described here is made up of the following subsystems:

- π/μ collection
- π decay region
- bunching
- phase rotation
- ionization cooling

DRAFT

The initial transverse phase space of the muon beam is determined mainly by the magnetic field strength in the channel, which provides the required focusing, and the radial aperture of the beam pipe. The longitudinal phase space can be modified by allowing the beam to travel a long distance in an empty magnetic lattice. This permits a correlation to develop between the temporal position and the energy of the particles in the bunch. Electric fields in RF cavities are then used to rotate the longitudinal phase space. This produces a longer particle bunch with a reduced energy spread. To ensure efficient acceleration, it is also necessary to bunch the beam to match the frequency of downstream RF cavities. Finally, it is also necessary to decrease the transverse emittance of the collected beam by means of an ionization cooling channel in order to optimize Neutrino Factory intensity.

4.1 Comparison of Front-end Systems

To understand the advantages and disadvantages of various cooling approaches, we compared the cooling channels proposed in the three published Neutrino Factory feasibility studies [4-1, 4-2, 4-3].

The single most significant difference in the various approaches compared is the choice of RF frequency:

- Japanese FFAG study [4-4], 5 MHz
- CERN linear channel studies [4-5–4-11], 88 MHz
- U.S. linear channel studies [4-12–4-14], 201 MHz

Another key choice involves the method of longitudinal capture, using either a single bunch in one RF bucket or a train of many bunches. This choice depends not only on the RF frequency but on the bunch structure of the proton driver. Because ionization cooling was included in most, but not all, designs, comparisons were made both with and without this feature. To permit valid performance comparisons, the ISS baseline decay ring configuration—a racetrack ring—was used for all cases.

The above analysis was based on the muon capture rate per initial pion in the decay channel, and this final efficiency is given in Table 4-1. This approach avoids the complication of pion production uncertainties, and particularly their energy dependence. It is, however, useful to relate these efficiencies to a number of muons per year with one fixed assumption of the pion production per proton GeV. To do this, the following assumptions were made:

1. the number of captured pions per 24 GeV proton in Study 2a, 0.94, was used; this can be expressed in terms of pions per proton, per GeV of proton energy, as 39%
2. to correct for the MARS predicted improvement in performance at 10 GeV compared with the Study 2a choice of 24 GeV, the pion yield was increased by 10%

DRAFT

3. for the racetrack ring geometry, 38% of decays in the ring are assumed to take place in the production straight section of the decay ring
4. an average proton beam power of 4 MW was taken
5. a “Snowmass year” of 10^7 s was assumed

As can be seen in Table 4-1, the basic U.S. Study 2a scheme using 201 MHz is the only proposal that appears to meet the intensity requirement (106%) without modification. It does somewhat better than an 88 MHz scheme employing multi-bunch phase rotation (100%) on account of its smaller decay losses during acceleration due its greater acceleration gradients. This scheme has less transverse acceptance (42%) than the CERN 88 MHz case (50%) because it has less cooling.³

4.2 Cooling vs. Accelerator Acceptance

There is a trade-off that can be made between the amount of cooling that must be done and the acceptance of the downstream accelerators [4-15]. Clearly, if the accelerator acceptance were larger than the equivalent emittance of the collected muons after bunching and phase rotation, *no* cooling would be needed. This is an important concept that has significant cost implications for the Neutrino Factory design. An early study of this type is shown in Fig. 4.1. The line at 0.17 accepted muons per proton corresponds to the design goal for U.S. Study 2a.

The curves show the number of muons that are contained in various transverse phase space acceptances as a function of the length of the cooling channel. With the 30 mm-rad acceptance used in Study 2a, the length of the cooling channel must be 80 m. At the time this study was done, it appeared that increasing the accelerator acceptance to 45 mm-rad was possible, and would eliminate the need for any cooling at all. Subsequently it was discovered that it is very difficult to obtain transverse acceptances much larger than 30 mm-rad in the non-scaling FFAG accelerators, due to longitudinal phase-space distortions caused by the dependence of the time-of-flight on transverse amplitude. After these issues with the accelerator acceptance—and the ability of proposed solutions to mitigate the problem—are better understood, this question should be revisited.

³This was a deliberate choice to reduce cost. If more cooling, using a tapered channel and liquid hydrogen (as in FS2 and the CERN proposal) were used, somewhat higher performance could be achieved, though at significant incremental cost.

DRAFT

Table 4.1-1. Summary of efficiencies for different cases, including an estimate of useful muon decays per year assuming the same pion production estimate for 10 GeV protons. Parenthesized values are unpublished estimates or calculations made for this comparison study. The + signs indicates the inclusion of a system to separate and separately phase rotate each sign.

Freq MHz		Cool	A_{\perp} mm-rad	Phase rotation	η_{\perp} (%)	η_{\parallel} (%)	η_{accel} (%)	n_{\pm}	η_{all} (%)	$\frac{\mu}{year} \times 10^{21}$
5	Japan	No	30	No	(18)	(39)	50	1	3.5	0.11
5	Japan	No	30	Yes	(18)	(60)	50	1	5.4	(0.17)
5	Japan+	No	30	Yes	(18)	(60)	50	2	11	(0.34)
44-88	CERN	Yes	15	Yes	(50)	(15)	65	1	4.9	0.16
44-88	CERN	Yes	15	Neuffer	(50)	(48)	65	2	31	(1.0)
44-88	CERN	No	30	Neuffer	(20)	(48)	65	2	13	(0.41)
201	FS2	Yes	15	Multi	31	56	81	1	14	0.45
201	FS2+	Yes	15	Multi+	31	56	81	2	28	(0.9)
201	FS2	No	30	Multi	24	56	81	1	11	0.35
201	S2a	Yes	30	Neuffer	42	48	81	2	33	1.06
201	S2a	No	30	Neuffer	24	48	81	2	19	0.61




Fig. 4-1. Accepted number of muons per proton as a function of cooling channel length for various assumed downstream accelerator transverse acceptances.

4.3 Baseline Front-end Description

We describe here the baseline front-end design adopted for the ISS. One new feature of the adopted design is its ability to simultaneously accommodate muons of both signs. Provided the detector can handle both signs, this effectively doubles the number of useful muons per year. The schematic layout of the baseline front end is shown in Fig. 4.2.

DRAFT




Fig. 4.2. Schematic layout of ISS baseline front end.

The baseline proton driver beam has an energy of 10 GeV. The pion collection system begins with a 20 T solenoid with 7.5 cm beam radius surrounding the target. This is followed by a 12 m long channel where the solenoid strength falls adiabatically to 1.75 T and the channel radius increases to 25 cm. There then follows a 100 m long channel where the pions decay to muons and a correlation is built up between the time and energy of the muons. This correlation in longitudinal phase space is used by the 50 m long adiabatic buncher. Bunching is accomplished 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 in the 50 m long rotator section, with higher gradients and decreasing frequencies as we proceed down the beam line, is used to rotate the beam in longitudinal phase space in order to reduce its energy spread. The final rms energy spread in this design is 10.5%. An 80 m long 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.

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. 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 beta function eliminates the need to place the absorbers at the low-beta 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 beryllium will, in turn, have a thin coating of TiN to prevent multipactoring. The alternating 2.8 T solenoidal field is produced with one solenoid per half cell, located between the RF cavities. The channel produces a final value of $\epsilon_{\text{TN}} = 7.4$ mm rad, which 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 5.5 mm rad.

The cooling channel increases the number of accepted muons by about a factor of 1.6. Normalizing to the incident 10 GeV proton beam energy on the mercury

target, the figure of merit for the ISS front end is 0.0077 ± 0.0009 for the positive muons and 0.0089 ± 0.0010 for the negative muons. This efficiency is similar to the result from Study 2a [4-15] for 24 GeV proton interactions. In addition, 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 would deliver both (time tagged) neutrinos and antineutrinos to the detector. The beam at the end of the cooling section consists of a train of about 80 bunches with a varying population of muons in each one.

5. Acceleration System

The goal of the acceleration system is to increase the beam kinetic energy from 138 MeV (the average kinetic energy in the cooling section) to a final energy in the range of 20–50 GeV. The layout described here will accelerate to 25 GeV, with an option of doubling that final energy to 50 GeV.

The design of the acceleration system is based on consideration of both cost and performance, taking into account the penalty associated with decays and other losses. The chosen approach also aims to minimize transverse and longitudinal emittance growth during acceleration. To minimize the effects of muon decay, particles must be accelerated as rapidly as possible. This is made more difficult by the fact that the beam sizes, both transverse and longitudinal, are very large. For these studies, the transverse normalized acceptance is chosen to be 30 mm-rad, and the longitudinal normalized acceptance is 150 mm-rad. The transverse normalized acceptance is defined to be $a^2 p / \beta m c$, where a is the maximum half-aperture at a given location, β is the Courant-Snyder beta function at that point, p is the total momentum, m is the muon mass, and c is the speed of light. The longitudinal normalized acceptance is defined for an upright ellipse to be $\Delta t \Delta E / m c$, where Δt is the maximum half-width in time of the beam, and ΔE is the maximum half-width in energy of the beam.

The types of subsystems and their sequence are similar to what was proposed in [5-1]. However, significant changes in the details have occurred, in particular the energies for the transitions between the subsystems.

5.1 Overall Scenario

The acceleration system consists of several different types of subsystems. The choice of where to use which type of subsystem is governed by beam dynamics and cost considerations. At this point, a detailed cost optimization has not been performed and many of the beam dynamics issues are still being studied. Thus, the scenario chosen here is based on initial estimates of machine performance and on past experience with the cost behavior of these systems.

High average gradients are necessary to minimize the amount of muon decay. Superconducting cavities are used to keep the RF power required to achieve these high gradients modest. Since the RF systems (cavities, their cryostats, RF power systems, and the associated cryogenic systems) tend to be the most expensive

DRAFT

component of the acceleration systems, and since RF cavities generally operate most economically at or near their highest achievable gradient, minimizing cost implies minimizing the number of RF cavities in the acceleration system. To reduce the required number of RF cavities, we chose designs where the beam makes multiple passes through the RF cavities. The choice of which subsystem to use where is based primarily on the number of passes through the cavities it can accommodate.

Figure 5.1 shows a diagram of the entire acceleration system. The following subsections will explain the different types of subsystems, why they were chosen, and the reasons behind the energy transition points.

5.1.1 Recirculating Linear Accelerators

Recirculating linear accelerators (RLAs) are machines that take one or more linacs and connect them by a series of arcs. After each pass through the linac, the beam enters a different arc, which will transport it to the next linac or the next pass through the same linac. The switchyard, where the beam from the linac is transported into each individual arc, uses fixed-field magnets. Because of the nonzero energy spread in the beam, the nonzero transverse beam size, the space required for magnet coils, and other considerations, the number of separate arcs the beam can be directed into is typically limited to 4 or 5. This, in turn, limits the number of passes through the cavities that an RLA can achieve.

One way to increase the efficiency of an RLA is to change its geometry. Figure 5.2 shows two different layouts: a racetrack layout and a dogbone layout. The racetrack layout is, in principle, more straightforward to design and build: the arcs bend in only one direction, and there is no need to introduce vertical bending to avoid beamline crossings that occur when one tries to minimize arc length. The dogbone geometry, on the other hand, is more efficient. In particular, since the energy separation at the switchyard effectively limits the number of passes one can make through the linac, the dogbone layout allows one to make twice as many passes through the linac than the racetrack layout, and is thus preferred.




Fig. 5.1. Layout of the acceleration system.

DRAFT




Fig. 5-2. RLA geometries: dogbone layout (above), racetrack layout (below).

The baseline design has two dogbone RLAs. We chose the maximum RLA energy to be 12.6 GeV. Using two RLAs allows a lower injection energy in the first one than if only one RLA is utilized. Furthermore, it potentially increases the amount of synchrotron oscillation in the RLAs, reducing the effects of the time-of-flight dependence on transverse amplitude in the RLA linacs (see below) and differential beam loading down the bunch train.

5.1.2 Pre-Acceleration Linac

We begin acceleration with a linac. This avoids problems found in a recirculating accelerator with large beam sizes and large relative energy spreads (in particular, the variation of the velocity with energy, which means that if the phase relationship between cavities in the RLA linac is correct for the final linac pass, it is incorrect for the initial pass). The linac is used to accelerate to a point where such effects can be handled in the RLA.

Because the time-of-flight variation with transverse amplitude causes particles with large transverse amplitudes to fall behind in RF phase, we choose the final energy of the pre-accelerator linac to be 0.9 GeV. While a detailed parameter optimization has not been completed, this energy is expected to be the lowest energy that will keep the phase slip tolerable due to the velocity difference between the first and last linac passes.

5.1.3 Fixed Field Alternating Gradient Accelerators (FFAGs)

To avoid the limitations of the RLA switchyard, it is possible to utilize a single arc for all beam energies. This is what is known as a fixed-field alternating-gradient (FFAG) accelerator. All FFAGs consist of a sequence of simple, identical cells with RF cavities in most of them. The design of the cell determines the type of the FFAG and the method by which beams must be accelerated. For this study, both scaling and non-scaling FFAGs were considered.

Scaling FFAGs

Scaling FFAGs are the original type of FFAG that was first described and built in the 1950s [5-2, 5-3, 5-4]. A design study for a Neutrino Factory based solely on

DRAFT

FFAGs for muon acceleration was completed in 2001 [4-4]. We have chosen not to use scaling designs in the ISS baseline design for two reasons.

1. Scaling FFAGs generally require relatively low frequency RF cavities in order to accelerate muons, of the order of 15 MHz, due to the relatively large time-of-flight variation with energy in these machines. This would require the earlier capture systems to use the same low RF frequency, which significantly decreases the capture efficiency of the machine.
2. Scaling FFAGs require relatively large aperture high-field superconducting magnets, which are expensive.

At lower energies, it may be possible overcome these difficulties of using scaling FFAGs.

Linear Non-Scaling FFAGs

Linear non-scaling FFAGs [5-5,5-6] attempt to address the two main difficulties of scaling FFAGs (large aperture and large time-of-flight variation with energy) by addressing the underlying reason for these problems—in a linear non-scaling FFAG, most of the bending is placed in the defocusing magnets. As a result, for an equivalent energy range, magnet apertures in a non-scaling FFAG can be reduced compared with a scaling device. Furthermore, at least for high energies, the ring can be made isochronous at a single energy within the energy range of the machine. This is shown in Fig. 5.3, which gives the time of flight dependence on energy in an example linear non-scaling FFAG. The relatively small time-of-flight variation with energy in these machines allows the use of relatively high frequency RF, such as the 201 MHz RF that is used in the bunching, phase rotation, and cooling channels. This permits reasonably high accelerating gradients.

Linear non-scaling FFAGs become more efficient at higher energies [5-7], as it is possible to make more passes through the cavities. As a result, our preference is to use FFAGs only at the higher energies, where they become more efficient than RLAs.

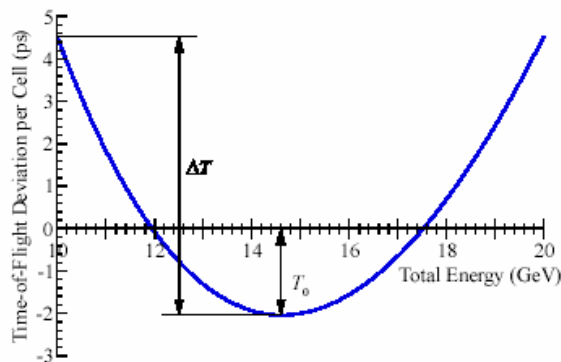


Fig. 5-3. Time of flight as a function of energy in a linear, non-scaling FFAG cell.

DRAFT

Based on our studies, it appears that a factor of two energy gain is roughly the optimal acceleration range for an FFAG stage; aperture requirements and time-of-flight range increase very rapidly beyond that. The primary difficulty with linear non-scaling FFAGs is that the time of flight depends on transverse amplitude [5-8]. Thus, large amplitude particles arrive at a later RF phase than low amplitude particles. This becomes problematic with multiple stages, since large amplitude particles should arrive earlier, not later, than low amplitude particles for optimal transmission in the next stage.

Despite this shortcoming, linear non-scaling FFAGs seem to be the best option for accelerating to the highest energies. Improvements that help to address the time-of-flight dependence on transverse amplitude are being examined and will be included in future designs. Preliminary simulations [5-9] suggest that two FFAG stages should result in a tolerable level of longitudinal emittance dilution. With this in mind, we choose (see Fig. 5.1) to use a non-scaling FFAG to reach 25 GeV, with a second stage being used to reach 50 GeV if required.

6. Decay Ring

Conceptual designs have been obtained for racetrack, triangular and bow-tie shaped, μ^+ and μ^- decay rings. A 20 (upgradeable to 50) GeV energy has been considered [6-1], with neutrino detectors at distances of 7500 and ~ 3500 km. For these baseline distances, racetrack designs need the ring planes tilted downwards by $\sim 36^\circ$ and $\sim 18^\circ$, respectively. Triangle and bow-tie designs need side-by-side rings in a (near) vertical plane, with detectors in (nearly) opposite directions from the rings, in gnomonic projection. If suitable detector sites are available, the triangle or bow-tie rings are favored, as their $\sim 40\%$ greater production efficiencies make them the more cost effective. If suitable sites are not available, the use of racetrack rings in separate tunnels provides the better solution. Both designs are compatible with the Neutrino Factory's pattern of three or five bunch trains, as described in the Section 2.3.

Recently, it has been recognized that a bow-tie shape of decay ring has several advantages if neutrino detectors are at 7500 and ~ 3500 km distances, as specified in the study. Unfortunately, a bow-tie shape preserves the muon polarization, and interferes with the accuracy of the related beam instrumentation [6-2]. A possible scheme to overcome this drawback is being considered.

6.1 Ring Features

The use of a single racetrack ring in each of two separately oriented tunnels is proposed as the baseline design, as this helps in finding suitable detector sites. An alternative is the use of two separate isosceles triangle (or bow-tie) shaped rings in a common, larger tunnel. The μ^+ and the μ^- bunch trains are injected into separate rings, though each racetrack ring has a μ^+ and μ^- counter-rotating beam

DRAFT

option. Stored muons decay to neutrinos, which pass from the ring straight sections to the detectors. Each racetrack ring is aligned to its own detector, and the racetracks have either one neutrino production region, or two of a slightly lower individual efficiency.

Each triangle ring has two downward sloping production regions, and the μ^+ trains in one ring are interleaved in time with the μ^- trains in the other. The two detectors accept neutrinos from both rings. Baseline distances needed are 7500 km for one detector, and 2500–3500 km for the other. For these distances, the smallest triangle apex angle (and the best production efficiency) is about 50° , when detectors are in opposite directions (in a gnomonic projection) from vertically aligned rings. Some incline of the rings to an exact vertical plane is expected for most pairs of suitable detector sites.

6.2 Ring Specifications

Detector locations for accelerator sites have to be defined before final parameters may be set. Specifications depend on input beam parameters along with detector and bunch pattern requirements. A compatible set of rings (with the booster orbit half the length of the driver injection orbit, and the proton rotation period at 10 GeV in the driver half that for the muons in the decay rings) are a booster of circumference 400.792 m, an NFFAG with injection (ejection), orbit lengths 801.584 (801.447) m, and a 20 GeV decay ring of circumference 1608.802 m.

An important parameter for the production straights is the ratio of the muon rms divergence angles to the rms opening angles of the decay neutrinos. The ratio has to be ≤ 0.14 at both 20 and 40–50 GeV, for the normalized, rms transverse input μ^\pm beam emittances of 4.8 mm-rad. The decay ring apertures are set a factor 50% larger than the beam envelopes, to allow the use of muon beam loss collectors, which are needed due to the megawatt muon beam powers involved.

6.3 Lattice Designs

Both the racetrack and triangle rings use bend magnets at the ends of the production straights to separate off neutrinos that arise from the muons of large divergence angles. These magnets complicate the lattice designs by creating dispersion in the matching sections to the main arcs. Six-parameter matching is needed, with the dispersion kept small in the regions of large betatron amplitudes. The lattices for the racetrack and triangle rings have different designs for the arcs, the production straights and the matching sections, though the designs are interchangeable.

6.3.1 Racetrack Ring

A layout of a 1608.802 m circumference, racetrack ring is shown in Figure 6.1. There are two arcs, each with 15 FODO cells of superconducting dipole and quadrupole magnets. If a single neutrino production straight is used, it has a length of 600.2 m, a production efficiency of 37.3%, and the other long straight

DRAFT

section has the collimators, tune control and rf systems. If a counter-rotating, μ^\pm beam option is used, the lattice is modified for two, shorter production straights of a slightly reduced efficiency.

Separate tunnels are used for the two racetrack rings as they are aligned to different detectors. Ring planes are sloped downwards at an angle of $\sim \arcsin(L/2R)$, where L is the distance to the detector and R is the equatorial radius. For the distances proposed of 7500 and 3500 km, the tilt angles are $\sim 36^\circ$ and $\sim 16^\circ$, and the maximum depths of the tunnels are ~ 435 m and ~ 185 m, respectively. The width of the tunnels is not as great as that of the common tunnel used for the triangle rings. The μ^+ and μ^- beam lines from the 20(50) GeV, μ^\pm accelerating ring have branches passing to each racetrack, for a total of four beam line tunnels. Extra services and service buildings have to be provided for the additional ring and beam line tunnels.

The racetrack decay rings are based on a design from an FNAL study [4-13] of year 2000, though parameters and some ring elements have changed. The β values in the production straight are reduced to ~ 153.0 m, while the transverse acceptances of the ring are increased to 67.5 mm-rad. The dispersion introduced by the dipoles at the ends of the production straight increases throughout the arc matching sections until a six-parameter match to the arcs is obtained.

For the proposed racetracks, the ratio of the muon rms divergence angles to the rms opening angles of the decay neutrinos is ~ 0.11 at 20 GeV, assuming a normalized rms muon emittance of 4.8 mm-rad. If the upgrade lattice is unchanged, the ratio scales with $\sqrt{\gamma}$ at higher energy, becoming ~ 0.17 at 50 GeV.




Fig. 6.1. Schematic layout for a racetrack shaped, muon decay ring.




Fig. 6.2. Betatron and dispersion functions in a racetrack ring.

6.3.2 Triangle Ring

A layout drawing of an isosceles triangle ring, with a 1608.802 m circumference and a 52.8° , apex angle, is shown in Figure 6.3. Two 398.45 m long downward sloping production straights give a neutrino production ratio efficiency of $2 \times 24.8\%$ (whereas a 60° angle gives $2 \times 23.9\%$). A maximum efficiency results when the apex angle is minimum, with the detectors in opposite directions (in gnomonic projection) from two vertically aligned rings in the same tunnel. When the detector sites are not opposite, it is necessary to tilt the plane of the rings about a production straight axis, and increase the apex angle until the straights and detectors are again re-aligned.

The arc cells have a FODO design, and the maximum values for betatron and dispersion functions are $\beta_v = 12.67$ m, $\beta_h = 12.67$ m and $D_h = 1.438$ m. At the ends of the arcs, the dispersion function is $D_h = 0.845$ m.

One noteworthy feature of this design is that focusing in each production straight is provided by eight 4.04 T, superconducting solenoids, arranged symmetrically. A figure-of-merit for the production straight focusing is given by the inverse of the maximum lattice $\beta\gamma$ function, which is ~ 1 for solenoids, but is $(1 - \sin \mu)/2$ for thin lens FODO cells, where μ is the half-cell phase advance. For equal muon divergence angles, the maximum β value for solenoid lenses is thus about half the value found in FODO focusing cells. This reduces the beam size in the production straights and the adjacent matching sections, and thus improves the ring dynamic aperture. The maximum beam diameter in the production straight is 265 mm at 20 GeV for 30 mm-rad normalized transverse emittances and a β of 94.3 m at the solenoid focusing waists. Aperture diameters are 50% larger, at ~ 398 mm. The ratio of the muon beam rms divergence angles, to the rms opening angles of the decay neutrinos, is 0.1, assuming a normalized rms emittance of 4.8 mm rad. Figure 6.4 shows the lattice functions over the entire ring.

Lattice modifications are planned when the rings are upgraded from 20 to 40 or 50 GeV. Some matching components would be repositioned and the solenoid focusing would be weaker (though the fields are increased). The two rings would require a realignment as a result of the modifications, and the fields at the central orbits of the superconducting arc magnets must be increased to ~ 5.6 T at 50 GeV.

DRAFT

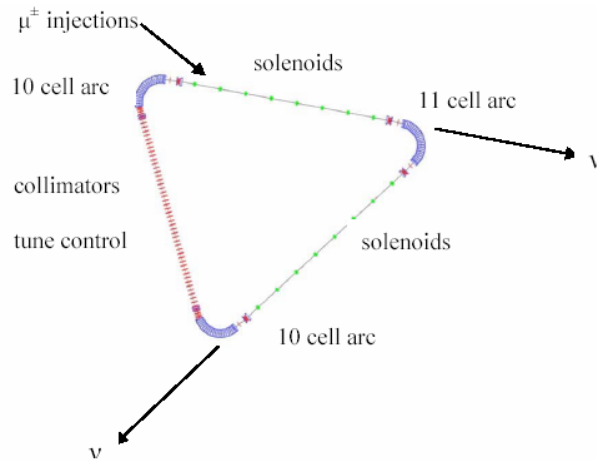


Fig. 6-3. Schematic layout for a 52.8° apex angle, isosceles triangle, muon decay ring.

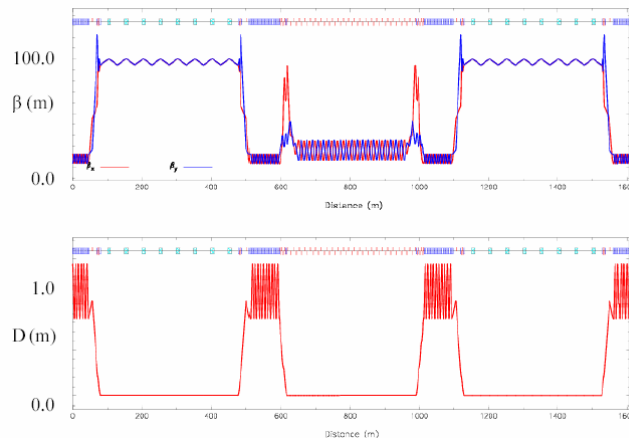


Fig. 6.4. Lattice betatron and dispersion functions for the 1608.802 m, triangle decay rings.

7. R&D Needs

R&D activities in support of a Neutrino Factory have been ongoing for many years. In recent years, the effort has become more and more a coordinated international effort. In what follows we will call out some of the main R&D issues that need to be studied in preparation

7.1 Proton Driver

In Section 2, we described a nonlinear non-scaling FFAG ring that could serve as a proton driver ring. As this concept is new and untried, the fabrication and testing of a low-energy electron model is called for. Continued development and testing of tracking codes adequate for this parameter regime are also important. For example, space-charge issues will be significant. Collimator development to protect key machine components is needed. Both primary beam loss and beam halo need to be considered.

DRAFT

In the case of the RF systems, development efforts aimed at improved designs for low-frequency high-gradient cavities must continue. Because the beam power is high, beam loading is a matter for concern and must be studied computationally and ultimately experimentally.

For the linac-based designs, details of the ancillary rings (accumulator and compressor) need to be specified. To permit hands-on maintenance, beam losses must be kept to a minimum. This will involve careful studies of vacuum issues, instabilities, and beam halo formation. Both J-PARC and LHC are developing the tools for this, and participation in such activities will be a help for proton driver development.

7.2 Target

Work on the liquid-Hg jet target are well along in the context of the MERIT experiment. This work needs to be completed and analyzed with high priority. Determination of acceptable single pulse and pulse train durations must be made. As a possible follow-on, it will be worth exploring other high-Z targets that are not liquid at room temperature but have a low melting point, e.g., a Pb-Bi eutectic. Much of this can be done off-line with a modified MERIT apparatus. An assessment of the need for beam tests should be part of this program.

Solid targets remain a possibility for a Neutrino Factory, at least at the 1 MW proton driver level. Continued tests to determine the power-handling capability of solid materials should continue in order to identify the practical limits of this technology. This will involve shock tests and irradiation studies to understand the changes in materials properties in the Neutrino Factory target environment. Determining acceptable single-bunch and bunch-train spacing parameters is necessary for solid targets also. Development of one or more practical implementation options for solid targets is needed. A beam test of such a system is also highly desirable.

At present, our information on pion production rates and their dependence on the proton beam parameters (particularly bunch length and energy) comes solely from model calculations. Completing and analyzing the measurements of production rates will be critical in deciding on the optimum parameters of the proton driver.

7.3 Front End

Foremost in this area is the demonstration of ionization cooling that will take place in the MICE experiment. This will take several more years to complete. Developing the various components needed remains a high priority. In particular, high-gradient RF cavities that operate in a strong solenoidal field are needed. Both vacuum cavities having irises closed with beryllium disks and H₂-gas-filled cavities have been proposed and both need further study. For the vacuum cavity,

DRAFT

the primary issue is the observed degradation of gradient in a strong magnetic field. For the gas-filled cavity, the main issue is whether the gas maintains its desirable insulating properties when subjected to an intense beam of ionizing radiation. This test requires an intense beam, but does not require muons. Absorber thermal tests with LiH sandwiched in beryllium must be carried out. As a follow-on to MICE, building and testing a section of “Guggenheim” cooling channel will provide options for providing 6D cooling, thus improving the compatibility between Neutrino Factory and Muon collider designs.

Experimental studies of muon multiple scattering are being analyzed and should be included in simulations of the cooling process. It is not expected that the cooling performance will be markedly changed by such details, but this needs to be confirmed.

In terms of simulations, we must optimize the machine by balancing the cooling channel performance against the acceptance of the acceleration system. We also need to evaluate the robustness of our technical solutions by means of error studies

7.4 Acceleration

The primary acceleration system R&D activity will be to participate in the EMMA experiment to test an electron model of a non-scaling FFAG. This type of accelerator is presently untested, and we need to know whether our performance simulations are correct. There are many beam dynamics issues that have arisen during the course of the ISS, most notably the dependence of time-of-flight on transverse amplitude, that need to be fully understood. At present, it appears that no more than two FFAG systems can be cascaded. We need to develop and test mitigation techniques to improve the situation.

Because the acceleration system layout tends to be tightly spaced (both RLAs and FFAGs), we need to demonstrate that we can operate superconducting RF cavities in close proximity to high-field magnets, and that we can achieve the requisite gradients at 201 MHz. Initial work on this at Cornell was encouraging, but much remains to be done.

Another area that needs exploration is the use of high-frequency cavities in a scaling FFAG. If there were intractable issues that arose with non-scaling FFAG designs, the so-called harmonic-number-jump acceleration scheme might be a viable fallback. This needs first to be studied in detail with simulations, but could lead to a hardware test if the calculation results look encouraging.

Because we require the largest practical acceptance for the acceleration system, the normal paraxial approximation does not hold. New tracking tools are being developed for this purpose, and these need to be checked carefully with other codes and with experiments whenever possible.

DRAFT

7.5 Decay Ring

The decay ring requires several novel superconducting magnets. They must be combined-function devices and must accommodate the substantial heat load from decay electrons from the muon beam. Designs for these magnets are needed, along with the corresponding cost estimates. Large aperture injection kickers capable of operating at 50 Hz must be developed.

Tracking studies with errors need to be continued, using specialized codes like Zgoubi [6-4] that can handle this parameter regime. Polarization studies are needed to see whether a bow-tie ring is a possible configuration for the decay rings.

8. Summary

In this document we have briefly summarized the findings of the ISS Accelerator Working Group. A lengthier report is in preparation that gives much more of the detail of our work.

We have developed parameters for the proton driver, determined an optimum target implementation, defined a front-end scenario, and proposed a viable acceleration scheme. Several decay ring geometries have been considered and compared. The present baseline assumes a pair of racetrack rings. The alternative triangle geometry, which has somewhat higher efficiency, would be preferred if suitable detector sites are available.

DRAFT

References

- [2-1] D. Neuffer, in the U S Neutrino Factory Design Study Report 2A (2004).
- [2-2] G. H. Rees, FFAG Studies for Neutrino Factory Accelerators, Nuclear Physics B, 155 (2006), p 301.
- [2-3] C. Plostinar, RAL, private communication (2006).
- [2-4] H. Lengeler (ed), The European Spallation Source Study, Vol 3, ESS-96-53-M (1996), p 3-15.
- [2-5] C. R. Prior and G. H. Rees, Synchrotron-based Proton Drivers for a Neutrino Factory, Proc. of the 7th European Particle Accelerator Conference, Vienna, Austria, 26-30th June, 2000, p 963; A. Blondel *et al.* (editors), ECFA/CERN Studies of a European Neutrino Factory Complex, CERN-2004-002 and ECFA/04/230, 13 April, 2004.
- [2-6] Fermilab and Muons proton driver, D. Neuffer, <http://www.cap.bnl.gov/mumu/project/ISS/060424/talks/DNeuffer1.ppt>
- [2-7] SPL-based 5 GeV proton driver, R. Garoby, http://nufact06.physics.uci.edu/Workshop/Slides/RGaroby_SPL3_Pdriver.ppt, presented at NuFact06, August 24 – 30, 2006, Irvine (CA).
- [2-8] Conceptual Design of the SPL II - a high power superconducting H- linac at CERN, F. Gerigk (editor), CERN-2006-006.
- [2-9] Time structure of a superbeam, M. Mezzetto, http://nufact06.physics.uci.edu/Workshop/Slides/MMezzetto_Presentation.pdf
- [2-10] Neutrino oscillation physics at an upgraded CNGS with large next generation liquid Argon TPC detectors, A. Meregaglia and A. Rubbia, <http://www.arxiv.org/abs/hep-ph/0609106v1>
- [2-11] Physics at the Fermilab proton driver, M.G. Albrow *et al.*, <http://www.arxiv.org/pdf/hep-ex/0509019>
- [2-12] Proton driver: the evolution of JPARC, S. Machida, Proc. of NuFact05, Nuclear Physics B, Proceedings Supplements, Vol.155, p.58 and http://www.lnf.infn.it/conference/nufact05/talks/Plenary/Machida_Plenary.ppt
- [2-13] Proton driver: prospects in the US, G. Appollinari, Proceedings of NuFact05, Nuclear Physics B, Proceedings Supplements, Vol.155, p.61 and http://www.lnf.infn.it/conference/nufact05/talks/Plenary/Apollinari_Plenary.ppt
- [2-14] The Theta_13 and Delta_CP sensitivities of the SPS-Frejus project, J.E. Campagne, A. Caze, <http://doc.cern.ch/archive/electronic/cern/others/neufact/nufact-note-142.pdf>
- [3-1] J.S. Berg *et al.*, “Cost-effective design for a neutrino factory”, Phys. Rev. Spec. Top.-Acc. Beams 9 (2006) 011001.
- [3-2] N.V. Mokhov *et al.*, Fermilab-Conf-98/379; LANL Report LA-UR-98-5716(1998); <http://www-ap.fnal.gov/MARS>.

DRAFT

- [3-3] R.C. Fernow, "Recent developments on the muon-facility design-code ICOOL", PAC'05, Knoxville, TN, 2005, p. 2651, <http://www.jacow.org>.
- [3-4] M. Alsharo'a et al., "Recent progress in neutrino factory and muon collider research within the Muon Collaboration", Phys. Rev. Spec. Top.-Acc. Beams 6 (2003) 081001
- [4-1] R.C. Fernow, Muon front end for the neutrino factory, Nuc. Phys. B (Proc. Suppl.) 155:74-78, 2006.
- [4-2] J.S. Berg et al., Cost effective design for a neutrino factory, Phys. Rev. Special Topics- Accelerators and Beams 9,011001 (2006).
- [4-3] ECFA/CERN studies of a European neutrino factory complex, CERN-2004-002.
- [4-4] NuFactJ Working Group, A feasibility study of a neutrino factory in Japan, v. 1.0, May 2001; <http://www-prism.kek.jp/nufactj/nufactj.pdf>.
- [4-5] A. Lombardi, "[A 40-80 MHz system for phase rotation and cooling](#)", April 2000, <http://slap.web.cern.ch/slap/NuFact/NuFact/nf20.pdf>.
- [4-6] A. Lombardi, "[A 40-80 MHz System for Phase Rotation and Cooling](#)", <http://slap.web.cern.ch/slap/NuFact/NuFact/nf34.pdf>.
- [4-7] S. Gilardoni, A. E. Ball, A. Blondel, N. Vassilopoulos, "[Updated Results of the Horn Study For the NuFact](#)", Sept 2000, <http://slap.web.cern.ch/slap/NuFact/NuFact/nf42.pdf>
- [4-8] R. Garoby, "[Current Activities for a Neutrino Factory at CERN](#)" <http://slap.web.cern.ch/slap/NuFact/NuFact/nf74.pdf>
- [4-9] R. Garoby and F. Gerigk, "[Cavity Design for the CERN Muon Cooling Channel](#)", Sept, 2001, <http://slap.web.cern.ch/slap/NuFact/NuFact/nf87.pdf>
- [4-10] H. Hanke, "Muon Phase Rotation and Cooling: Simulation Work at CERN"; presented at Nufact02, London, July 1-6 2002.
- [4-11] G Franchetti et al, Proc. Nufact02, London, July 1-6 2002
- [4-12] N. Holtkamp and D. Finley. eds., "A Feasibility Study of a Neutrino Source Based on a Muon Storage Ring," Fermilab-Pub-00/108-E, 2000.
- [4-13] R.B. Palmer, Cooling optimization, http://mice.iit.edu/nfmcc06/nfmcc06_palmer_tocoolornot.pdf.
- [4-14] S. Ozaki, R. Palmer, M. Zisman, and J. Gallardo, eds., "Feasibility Study-II of a Muon-Based Neutrino Source," BNL-52623 (2001); see http://www.cap.bnl.gov/mumu/studyii/final_draft/The-Report.pdf.
- [4-15] S. Geer and M. Zisman, eds., "Neutrino Factory and Beta Beam Experiments and Development, <http://www.cap.bnl.gov/mumu/study2a/>.
- [5-1] J. S. Berg *et al.*, Phys. Rev. ST Accel. Beams **9**, 011001 (2006).
- [5-2] K. R. Symon *et al.*, Phys. Rev. **103**, 1837 (1956).
- [5-3] F. T. Cole *et al.*, Rev. Sci. Instrum. **28**, 403 (1957).
- [5-4] D. W. Kerst *et al.*, Rev. Sci. Instrum. **31**, 1076 (1960).
- [5-5] C. Johnstone, W. Wan, and A. Garren, in *Proceedings of the 1999 Particle Accelerator Conference, New York, 1999*, A. Luccio and W. MacKay, eds., IEEE, Piscataway, NJ (1999), p. 3068.

DRAFT

- [5-6] F. E. Mills and C. Johnstone, in *Proceedings of the 4th International Conference on Physics Potential & Development of $\mu^+ \mu^-$ Colliders*, San Francisco, CA (UCLA, Los Angeles, CA, 1999), pp. 693-698.
- [5-7] J. Scott Berg, in *The International Workshop on FFAG Accelerators, October 13–16, 2004, KEK, Tsukuba, Japan*, S. Machida, Y. Mori, and T. Yokoi, eds. (2005), p. 1, available from http://hadron.kek.jp/FFAG/FFAG04_HP/. Also MUC-CONF-ACCELERATION-309.
- [5-8] S. Machida, in *The International Workshop on FFAG Accelerators, December 5–9, 2005, KURRI, Osaka, Japan*, Y. Mori, M. Aiba, and K. Okabe, eds. (2006), p. 65. http://hadron.kek.jp/FFAG/FFAG05_HP/
- [5-9] S. Machida, “FFAGs as Muon Accelerators for a Neutrino Factory,” to appear in the proceedings of the 10th Biennial European Particle Accelerator Conference, EPAC’06 (2006).
- [6-1] G. H. Rees, C. Johnstone and F. Meot, “20–50 GeV Muon Decay Rings for a Neutrino Factory,” Proc. of 10th European Particle Accelerator Conference, Edinburgh, U.K., 26-30th June, 2006.
- [6-2] A. Blondel *et al.* (editors), ECFA / CERN Studies of a European Neutrino Factory Complex, CERN-2004-002 and ECFA/04/230, 13 April, 2004.
- [6-4] F. Meot and G. Rees, Report DAPNIA-06-04, CEA, Saclay, 2004.