2. Introduction

A muon storage ring as a source of intense neutrino beams supersedes a standard neutrino source in many ways. Classical neutrino sources have long decay channels which are used to generate $\overline{\nu}_{\mu,e},\nu_{\mu,e}$ beams from pions coming from a target that is hit with an intense proton beam. In a muon storage ring the muons circulate after injection until they decay. A fraction of these muons will decay in the straight section, which will produce an intense, well collimated neutrino beam. If the muon beam divergence in the straight section is small compared to the decay angle, the opening angle of the neutrino beam is completely dominated by the decay kinematics. Given the energy of the muons this angle basically equals $1/\gamma_{\text{muon}}$. From the requirement to have the divergence of the muon beam in the straight section to be small compared to the divergence of the neutrino beam, an emittance goal for the muon source and the cooling channel can easily be defined.

A muon storage ring used to produce very intense and clean neutrino beams is most probably the first application of an intense muon source. After being generated from pion decay and cooled in an ionization cooling channel, the muon beam is accelerated and injected into a storage ring, where the muons decay while circulating. The neutrinos from the decay muons in the storage ring form a very intense and well collimated beam of electron and muon neutrinos (ν_e,ν_μ). The idea for such a neutrino source has been described many times (see section 2.1), but only recently with the progress being made on ionization cooling concepts within the Neutrino Factory and Muon Collider Collaboration, does an intense source seem feasible. With a new proton driver and a target that can withstand the power density and the intense radiation from the impinging proton beam, the source will produce enough muons to achieve $2\text{x}10^{20}$ muons decaying into neutrinos in one of the straight sections of the storage ring. In order to achieve this goal, very efficient and large aperture focusing solenoids and rf accelerating systems must be developed for the ionization cooling channel. The transverse emittance that has to be achieved in this channel, to be sufficient for a neutrino source on the other hand, has to be reduced by only a factor of approximately ten in both transverse dimensions. The longitudinal emittance coming from the source is almost of no importance, which makes longitudinal cooling unnecessary. Given the intensity goal of $2\text{x}10^{20}$ muons/year decaying in one straight section, an attempt has been made to investigate the technical feasibility of such a facility as a whole (see also [2]). This is reflected in the charge that was given to the study group and is summarized in the box to the right.

Charge for the purpose of this study

1. A design concept for a muon storage ring and associated support facilities that could, with reasonable assurance, meet performance goals required to support a compelling neutrino based research program.
2. Identification of the likely cost drivers within such a facility.
3. Identification of an R&D program that would be required to address key areas of technological uncertainty and cost/performance optimization within this design, and that would, upon successful completion, allow one to move with confidence into the conceptual design stage of such a facility.
4. Identification of any specific environmental, safety, and health issues that will require our attention.

Even though we have done our best to be complete and consistent in this report, the careful reader will note some inconsistencies throughout this report, particularly from chapter to chapter. The reason is that, early on, we had to specify many of the parameters so that we could make reasonably rapid progress on the engineering designs. However, optimization of the various subsystems continued vigorously during the six months of this study. The resulting inconsistencies are indications of ongoing progress in the design of this complicated and intertwined facility.

Later on in Table 2 the design parameters for the Neutrino Factory are presented in more detail. One of the main parameters, the neutrino flux that can be achieved as a measure of the performance, is worth discussing at this point. The goal was to achieve 2×10^{20} muon decays per year in one straight section of the storage ring. Given the numbers for muon survival from the different subsystems that will be described later in the report, the presented scenario will instead provide 6.0×10^{19} muon decays per year. This assumes perfect transmission between the different subsystems. Detailed error analyses still have to be done and are not included in the present performance.

2.1 Physics Motivation for a Neutrino Source Based on a Muon Storage Ring

Recent measurements of atmospheric muon neutrino (ν_μ) fluxes from the Super—Kamiokande (SuperK) collaboration have shown an azimuth—dependent (→ baseline dependent) depletion that strongly suggests neutrino oscillations of the type $\nu_\mu \to \nu_x$. Since the atmospheric ν_e flux is not similarly depleted, ν_x cannot be ν_e and must therefore be either ν_τ, or ν_s (a sterile neutrino). These observations have inspired many theoretical papers, several neutrino oscillation experiment proposals, and much interest in the physics community. This interest is well motivated. Understanding the neutrino-mass hierarchy and the mixing matrix that drives flavor oscillations may provide clues that lead to a deeper understanding of physics at very high mass-scales and insights into the physics associated with the existence of more than one lepton flavor. Hence, there is a strong incentive to find a way of measuring the neutrino flavor mixing matrix, confirm the oscillation scheme (three—flavor mixing, four—flavor, n-flavor ?), and determine which mass eigenstate is the heaviest (and which is the lightest). This will require a further generation of accelerator based experiments beyond those currently proposed.

High energy neutrino beams are currently produced by creating a beam of charged pions that decay in a long channel pointing in the desired direction. This results in a beam of muon neutrinos ($\pi^+ \to \mu^+ + \nu_\mu$) or muon anti—neutrinos ($\pi^- \to \mu^- + \bar{\nu}_\mu$). In the future, we will need ν_e and $\bar{\nu}_e$ (as well as ν_μ and $\bar{\nu}_\mu$) beams to adequately unravel the mixing matrix. To illustrate this, consider neutrino oscillations within the framework of three-flavor mixing, and adopt the simplifying approximation that only the leading oscillations contribute (those driven by the largest Δm^2_{ij} defined as $\Delta m^2_{32} \equiv \Delta m^2_3 - \Delta m^2_2$, where m_i is the mass associated with mass eigenstate **i**.) The probability that a neutrino of energy E~(GeV) and flavor α oscillates into a neutrino of flavor β whilst traversing a distance L~(km) is given by:

$$P(\upsilon_e \to \upsilon_\mu) = \sin^2(\theta_{23}); \sin^2(2\theta_{13}); \sin^2(1.267\Delta m^2_{32}\, {}^{L}\!/_{E})$$

$$P(\upsilon_e \to \upsilon_\tau) = \cos^2(\theta_{23}); \sin^2(2\theta_{13}); \sin^2(1.267\Delta m^2_{32}\, {}^{L}\!/_{E})$$

$$P(\upsilon_\mu \to \upsilon_\tau) = \cos^4(\theta_{13}); \sin^2(2\theta_{23}); \sin^2(1.267\Delta m^2_{32}\, {}^{L}\!/_{E})$$

Each of the oscillation probabilities depends on Δm^2_{32} and two mixing angles θ_{ij}. To adequately determine all the θ_{ij} and sort out the various factors contributing to the $P(\nu_\alpha \to \nu_\beta)$ will require ν_e as well as ν_μ beams. In addition, there is a bonus in using ν_e beams since electron neutrinos can elastically forward scatter off electrons in matter by the charged current (CC) interaction. This introduces a term in the mixing matrix corresponding to $\nu_e \to \nu_e$ transitions that is not present for neutrinos of other flavors. Hence, if electron—neutrinos travel sufficiently far through the Earth, matter effects modify the oscillation probabilities. This modification depends on the sign of Δm^2_{32}, and provides a unique way of measuring which mass eigenstate is heaviest and which is lightest. We conclude that if we can find a way of producing ν_e beams *of sufficient intensity*, we are highly motivated to do so.

The obvious way to produce high energy ν_e beams is to exploit muon decays. Since muons live 100 times longer than pions, we need to avoid using a linear decay channel, which would be impractically long for high energy muons. The solution is to use a muon storage ring with long straight sections, one of which points in the desired direction. This yields a neutrino beam consisting of 50% ν_e and 50% $\bar{\nu}_\mu$ if μ^+ are stored, or 50% ν_μ and 50% $\bar{\nu}_e$ if μ^- are stored. Using a storage ring to produce secondary beams of $\mu^\pm$, $e^\pm$, $\bar{p}$ and ν was proposed by Koshkarev [3] in 1974. The idea is also ascribed to Wojcicki [4] as well as Collins [5]. The muons were to be produced by allowing high-energy pions to decay within the ring. The key questions that need to be addressed in order to produce a viable proposal for the production of secondary beams by this method are:

- How can enough particles be stored ?
- How can their phase-space be compressed to produce sufficiently intense beams for physics ?

The calculated beam fluxes used in [3], [4] and [5] were too low to motivate the construction of a secondary beam storage ring. A viable solution to the key question (how to make sufficiently intense beams) was implemented at the beginning of the 1980's for antiproton production, leading directly to the CERN proton—antiproton collider and the discovery of the weak Intermediate Vector Bosons. The solution to the intensity question involved using lithium lenses to collect as many negative particles as possible, and stochastic cooling to reduce the phase-space of the $\bar{p}$ beam before acceleration. In 1980 it was suggested [6] that the negative particle collection ring (the Debuncher) at the proposed

Fermilab antiproton source could be used to provide a neutrino beam downstream of one of its long straight sections. The Debuncher collects negative pions (as well as antiprotons) which decay to produce a flux of captured negative muons. The muon flux in the Debuncher was subsequently measured and found to be modest. The short baseline neutrino oscillation experiment proposal (P860 [7]) that was developed following these ideas was not approved ... the problem of intensity had not been solved!

In order to make progress we need a method of cooling muon beams and a way of producing more muons. Stochastic cooling cannot be used since the cooling time is much longer than the muon lifetime. Ionization cooling was proposed as a possible solution (see [8]). A way of collecting more pions (that subsequently decay into muons) using a very high-field solenoid was proposed by Djilkibaev and Lobashev [9] in 1989. Thus by the end of the 1980's the conceptual ingredients required for very intense muon sources were in place, but the technical details had not been developed. Fortunately in the 1990's the desire to exploit an intense muon source to produce muon beams for a high energy muon collider motivated the formation of a collaboration (back then, the *Muon Collider Collaboration*, subsequently renamed as *The Neutrino Factory and Muon Collider Collaboration*). This has resulted in a more complete technical understanding of the design of an intense muon source [10]. In 1997 it was proposed by Geer [11] to use a muon collider type muon source to produce an intense beam of muons at low energy, rapidly accelerator the muons to high energies, and inject them into a dedicated muon storage ring with long straight sections, to produce a very intense neutrino source. It was shown that this "neutrino factory" was sufficiently intense to produce thousands of events per year in a reasonably sized detector on the other side of the earth! The intensity problem had been solved! In addition, it was shown that the ring could be tilted at large angles to provide beams for very long (trans-earth) neutrino oscillation experiments, and that muon polarization could in principle be exploited to turn on/off the initial ν_e flux [11]. This proposal came at a time of increasing interest in neutrino oscillation experiments due to the SuperK results, and also at a time when the particle physics community was/is considering possible facilities needed at its laboratories in the future [12]. Thus, the neutrino factory concept quickly caught the imagination of the physics community. At the very beginning a specific scenario was picked to investigate the technical feasibility of many of the components necessary for such a source. Given the knowledge at that time about the intensity that would be required to have a compelling physics program, the energy and the baseline length was specified as well. These parameters, the basis for the accelerator facility feasibility study, are presented in Table 1. The physics study [1] on the other hand, investigated a much broader spectrum of experiments, which most probably would lead to a reconsidered set of parameters that would be used later on.

Energy of the Storage Ring should be 50 GeV

Number of neutrinos/straight section is $2\text{x}10^{20}$ per year

No polarization

Capability to switch between μ^+ and μ^+

Baseline for facility Fermilab to SLAC/LBNL

Table 1: Set of parameters chosen for the feasibility study following a very early assessment of the goals for the physics study.

2.2 Basis for the Accelerator Facility Layout

If the μ-beam divergence in the straight section is small compared to the decay angle, the opening angle of the neutrino beam is dominated by the decay kinematics. Given the energy of the muons this angle equals $\approx 1/\gamma_{\text{muon}}$, where γ_{muon} is the muon energy in units of $m_{\text{muon}}c^2$. Thus, the opening angle of the muon beam in the storage ring, in order to not contribute significantly to the divergence of the neutrino beam, has to be small compared to $1/\gamma_{\text{muon}}$. A reasonably large β-function in the decay straight allows therefore a comfortably large emittance which in our case has to be approximately a factor of 10 smaller (in both dimensions) than the emittance coming from the target. This defines the goal for transverse emittance cooling. The total flux which can be achieved defines the performance of a neutrino source. (This is quite different from a muon collider where the luminosity is proportional to the square of the number of particles per bunch.) Therefore the longitudinal emittance is of no direct interest and the flux is simply proportional to the number of particles per pulse - independent of the longitudinal distribution. The longitudinal emittance has to be small enough to manipulate and finally cool the beam, which is the only requirement that has to be met.

In the simple version of a racetrack shaped storage ring with two long straight sections considered in this study, more than one third of the muons will decay in each straight section. Given the large number of different and technically demanding sub-systems required for the entire facility, the charge for the feasibility study was focused on basic questions one would have to answer for such an facility (see charge in the box on the first page). Given the large variety of possibilities for short (~500 km), long (~3000 km) and very long baseline (>8000 km) experiments which all influence the technical layout one way or the other, a choice had to be made to focus on one design. A specific set of accelerator parameters was chosen based on the physics goals known at that time. Other boundary conditions were:

1) Given the experience with the simulations being done for the Muon Collider and earlier studies of a neutrino source [13], a reasonable assumption had to be made for the number of muons one could obtain per incident proton on target. This number, which includes all the decay losses and the beam loss during cooling and acceleration, is a very critical number because it defines the requirements on the proton driver and target. The goal for this study is to achieve 0.1 muon injected into the ring per incident proton on target.
2) We required at least one third of the muons circulating in the ring to decay in each straight section, given the available space and assuming a racetrack shape for the storage ring.
3) Because this is a pulsed accelerator, the average current that has to be accelerated to achieve 2×10^{20} neutrinos/year, critically depends on the accumulated operating time per year. More operating time reduces the investment cost in the high power rf systems which are expected to dominate the cost. An optimistic assumption led to 2×10^{7} sec/year assumed for the purpose of this study.
4) A storage ring tunnel with an acceptable slope would be the only possible design that could reasonably be investigated over the time scale that was available. This choice is particularly relevant since cryogenics and civil engineering would be based on experience with other more standard type installations.
5) Abandoning polarization for this study had two advantages. A very low frequency, high gradient rf system that was proposed directly after the target [13] would not be necessary, because the correlation in longitudinal phase space for the forward and backward polarized pions does not have to be preserved. For the same reason the proton bunch length in the proton accelerator could go up to 3 nsec instead of 1 nsec, which is a significant relief.

The final list of parameters for this study is shown in Table 1. This table, together with assumptions 1, 2 and 3 above, led to most of the specifications that were necessary to start the design work on the accelerator complex. These specifications are summarized in Table 2.

1. Given the ongoing study at Fermilab for a fast cycling proton synchrotron (15 Hz) with 16 GeV extraction energy, the number of protons per pulse required on target is at least 2×10^{13}. This as approximately 1 MW beam power on target.
2. The transverse emittance of the muon beam after the cooling channel has to be small enough, in order to have the beam divergence in the straight section to be less then 1/10 of the decay angle, which is $1/\gamma = 2$ mrad. Given an invariant emittance of $\gamma\cdot\varepsilon=3.2\pi\cdot\mathrm{mm}\cdot\mathrm{rad}$ the β-function would be ~400 m. This seemed reasonable.
3. Following the assumption of having ten protons per one muon injected in the storage ring, 2×10^{12} muons per pulse are required after the cooling channel and have to be accelerated.
4. No polarization.
5. The Neutrino beam is directed from Fermilab to SLAC/LBNL with a distance of ~3000 km. This sets the slope of the storage ring with respect to the earth surface at 22% or 13 deg. Gentle enough to think of conventional installation methods.

Table 2: Specifications for the accelerator complex of the neutrino source.

It was also recognized very early, and it is worth noting here, that because of the high energy (50 GeV) and high average current (6×10^{20} muons per year in 2×10^{7} seconds) the average muon beam power would be 240 kW during operation. One of the highest pulsed power lepton beams in the world. This would clearly be a cost driver and led very soon to a focus on lower energy. A unique and very interesting feature of neutrino sources was identified as a result of this discussion, namely the possibility to balance cost between higher energy and higher intensity (basically a more expensive accelerator) and a larger and more advanced detector. This will certainly be used for cost optimization in later studies.

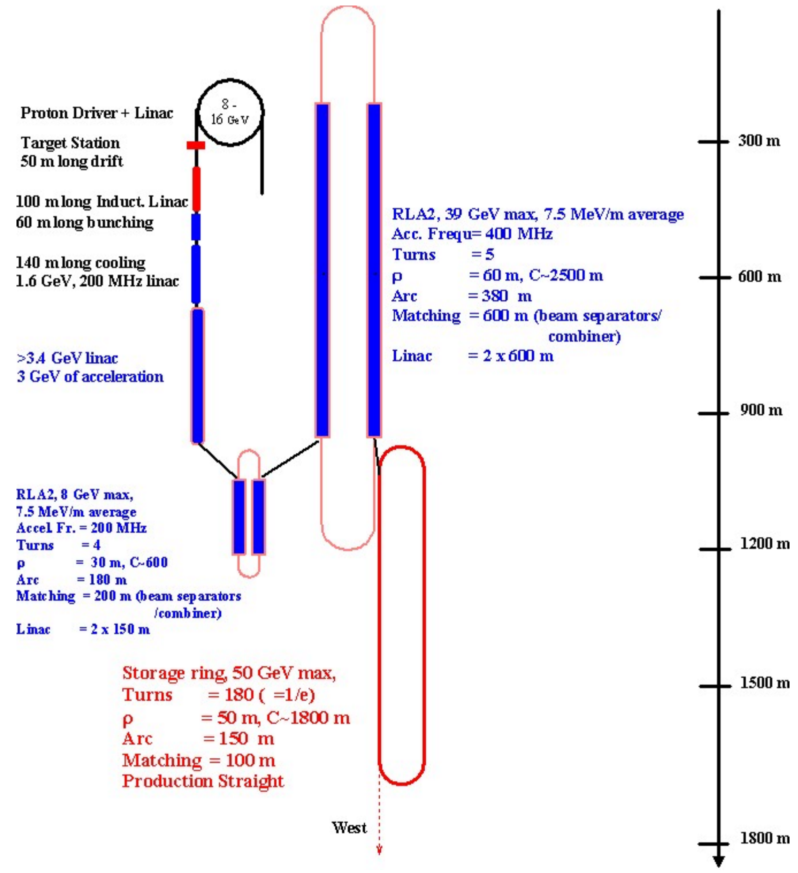

Figure 1: Footprint (to scale) of the whole facility.

For the storage ring an acceptance which is twice as large as the expected emittance from the cooling channel was chosen, because of the uncertainties in the cooling channel design. The aperture was designed to accommodate $\pm 3\sigma$ of 3.2 π mm-rad rms normalized emittance. This allows for a total emittance growth of approximately a factor of 2 in the accelerating systems, once the muon beam has been cooled down to the goal value of 1.6 π mm-rad or it would allow for a somewhat larger emittance coming out of the cooling channel.

The footprint of the total facility is comparatively small and fits easily under several existing laboratory sites. The same is certainly true for any detector that could be considered. A generic sketch of the accelerator facility which is made to scale is shown in Figure 1. This figure shows the logical relationships between the various subsystems. The largest subsystems are the accelerating linac in the cooling channel, the superconducting linac after the cooling and the recirculating accelerators (RLA1 and RLA2). The total area required in order to provide a 50 GeV muon beam to a storage ring is approximately 1.0 × 2.0 km. A more elaborate, site specific picture is shown in Chapter 13. There, using minimal deviations from this logical layout, we have integrated the facility onto the Fermilab site. The proton driver is placed near the Main Injector, and the rest of the facility easily fits inside the Tevatron.

The basic argument leading to the generic layout is that bending between the different subsystems should be minimized. This will minimize muon loss because of the large transverse emittance that will have to be transported. The same number of passes through each linac of the RLAs is another criterion that was applied to make the beam loading equal on both sides of each RLA, which leads to identical rf system requirements for both sides. Coming out of the last RLA, the muon beam would be gently bent downward into the storage ring tunnel and injected into the straight section pointing to the long baseline experiment. Another remarkable result of this layout, given the earlier boundary conditions, is that the direction the proton beam hits the target defines the natural direction of the neutrino beam going to the experiment. Therefore once the location of the detector is fixed, the layout is constrained, or one of the boundary conditions have to be given up, which will most probably increase cost or decrease performance.

REFERENCES

[1] Neutrino factory physics study coordinators: S. Geer and H. Schellman. See http://www.fnal.gov/projects/muon_collider/nu/study/study.html

[2] Neutrino factory technical study coordinators: N.Holtkamp and D. Finley. See http://www.fnal.gov/projects/muon_collider/nu-factory/nu-factory.html

[3] D.G.Koshkarev, Preprint ITEP-33, 1974 and “Proposal for a decay ring to produce intense secondary particle beams st the SPS”, CERN/ISR-DI/74-62.

[4] S. Wojcicki, unpublished (1974).

[5] T. Collins, unpublished (1975).

[6] D. Cline and D. Neuffer, AIP Conf. Proc. 68, 846 (1980).

[7] W. Lee (spokesperson) et al., FNAL Proposal P860, “A Search for neutrino oscillations using the Fermilab Debuncher”, 1992.

[8] A.A. Kolomensky, Sov. Atomic Energy Vol. 19, 1511 (1965).

[9] R.M. Djilkibaev and V.M. Lobashev, Sov. J. Nucl. Phys. 49(2)}, 384 (1989).

[10] "$\mu^+\mu^-$ Collider: A Feasibility Study", BNL-52503, Fermilab-Conf. 96/092, LBNL-38946, and C. Ankenbrandt et al. (Muon Collider Collaboration), Phys. Rev. ST Accel. Beams 2, 081001 (1999).

[11] S. Geer, “Neutrino beams from muon storage rings: characteristics and physics potential”, FERMILAB-PUB-97-389, 1997; Presented in the Workshop on Physics at the First Muon Collider and Front-End of a Muon Collider (Editors: S. Geer and R. Raja), November, 1997, and published in S. Geer, Phys.Rev.D57:6989-6997,1998.

[12] J. Ellis, E. Keil, G. Rolandi, “Options for future colliders at CERN”, CERN—SL/98-004 (1998)., B. Autin, A. Blondel, J. Ellis (editors), “Prospective study of muon storage rings at CERN”, CERN 99-02, April (1999).,D. Finley, S. Geer, J. Sims, “Muon Colliders: A vision for the future of Fermilab”, FERMILAB-TM-2072, Jun 1999., S.-C. Ahn et al., “Muon Colliders: A Scenario for the evolution of the Fermilab accelerator complex”, FERMILAB-FN-677, Aug 1999.

[13] "The Cost Effective Design for a Neutrino Factory", R. Palmer, C .Johnson, E. Keil, CERN-SL-99-070-AP, Nov 1999. 22pp, CERN Library, MUCOOL note Nr. 67.