Recent progress in neutrino factory and muon collider research within the Muon Collaboration

Mohammad M. Alsharo'a,[1] Charles M. Ankenbrandt,[2] Muzaffer Atac,[2] Bruno R. Autin,[3] Valeri I. Balbekov,[2]
Vernon D. Barger,[4] Odette Benary,[5] J. Roger J. Bennett,[6] Michael S. Berger,[7] J. Scott Berg,[8] Martin Berz,[9]
Edgar L. Black,[1] Alain Blondel,[10] S. Alex Bogacz,[11] M. Bonesini,[12] Stephen B. Bracker,[13] Alan D. Bross,[2] Luca Bruno,[3]
Elizabeth J. Buckley-Geer,[2] Allen C. Caldwell,[14] Mario Campanelli,[10] Kevin W. Cassel,[1] M. Gabriela Catanesi,[12]
Swapan Chattopadhyay,[11] Weiren Chou,[2] David B. Cline,[15] Linda R. Coney,[14] Janet M. Conrad,[14] John N. Corlett,[16]
Lucien Cremaldi,[13] Mary Anne Cummings,[17] Christine Darve,[2] Fritz DeJongh,[2] Alexandr Drozhdin,[2] Paul Drumm,[6]
V. Daniel Elvira,[2] Deborah Errede,[18] Adrian Fabich,[3] William M. Fawley,[16] Richard C. Fernow,[8] Massimo Ferrario,[12]
David A. Finley,[2] Nathaniel J. Fisch,[19] Yasuo Fukui,[15] Miguel A. Furman,[16] Tony A. Gabriel,[20] Raphael Galea,[14]
Juan C. Gallardo,[8] Roland Garoby,[3] Alper A. Garren,[15] Stephen H. Geer,[2] Simone Gilardoni,[3]
Andreas J. Van Ginneken,[2] Ilya F. Ginzburg,[21] Romulus Godang,[13] Maury Goodman,[22] Michael R. Gosz,[1]
Michael A. Green,[16] Peter Gruber,[3] John F. Gunion,[23] Ramesh Gupta,[8] John R Haines,[20] Klaus Hanke,[3]
Gail G. Hanson,[24] Tao Han,[4] Michael Haney,[18] Don Hartill,[25] Robert E. Hartline,[26] Helmut D. Haseroth,[3]
Ahmed Hassanein,[22] Kara Hoffman,[27] Norbert Holtkamp,[20] E. Barbara Holzer,[3] Colin Johnson,[3] Rolland P. Johnson,[26]
Carol Johnstone,[2] Klaus Jungmann,[28] Stephen A. Kahn,[8] Daniel M. Kaplan,[1] Eberhard K. Keil,[2] Eun-San Kim,[29]
Kwang-Je Kim,[27] Bruce J. King,[30] Harold G. Kirk,[8] Yoshitaka Kuno,[31] Tony S. Ladran,[16] Wing W. Lau,[32]
John G. Learned,[33] Valeri Lebedev,[2] Paul Lebrun,[2] Kevin Lee,[15] Jacques A. Lettry,[3] Marco Laveder,[12] Derun Li,[16]
Alessandra Lombardi,[3] Changguo Lu,[34] Kyoko Makino,[18] Vladimir Malkin,[19] D. Marfatia,[35] Kirk T. McDonald,[34]
Mauro Mezzetto,[12] John R. Miller,[36] Frederick E. Mills,[2] I. Mocioiu,[37] Nikolai V. Mokhov,[2] Jocelyn Monroe,[14]
Alfred Moretti,[2] Yoshiharu Mori,[38] David V. Neuffer,[2] King-Yuen Ng,[2] James H. Norem,[22] Yasar Onel,[39]
Mark Oreglia,[27] Satoshi Ozaki,[8] Hasan Padamsee,[25] Sandip Pakvasa,[33] Robert B. Palmer,[8] Brett Parker,[8]
Zohreh Parsa,[8] Gregory Penn,[40] Yuriy Pischalnikov,[15] Milorad B. Popovic,[2] Zubao Qian,[2] Emilio Radicioni,[12]
Rajendran Raja,[2,*] Helge L. Ravn,[3] Claude B. Reed,[22] Louis L. Reginato,[16] Pavel Rehak,[8] Robert A. Rimmer,[11]
Thomas J. Roberts,[1] Thomas Roser,[8] Robert Rossmanith,[41] Roman V. Samulyak,[8] Ronald M. Scanlan,[16]
Stefan Schlenstedt,[42] Peter Schwandt,[7] Andrew M. Sessler,[16] Michael H. Shaevitz,[14] Robert Shrock,[37] Peter Sievers,[3]
Gregory I. Silvestrov,[43] Nick Simos,[8] Alexander N. Skrinsky,[43] Nickolas Solomey,[1] Philip T. Spampinato,[20]
Panagiotis Spentzouris,[2] R. Stefanski,[2] Peter Stoltz,[44] Iuliu Stumer,[8] Donald J. Summers,[13] Lee C. Teng,[22]
Peter A. Thieberger,[8] Maury Tigner,[25] Michael Todosow,[8] Alvin V. Tollestrup,[2] Yağmur Torun,[1] Dejan Trbojevic,[8]
Zafar U. Usubov,[2] Tatiana A. Vsevolozhskaya,[43] Yau Wah,[27] Chun-xi Wang,[22] Haipeng Wang,[11] Robert J. Weggel,[8]
K. Whisnant,[45] Erich H. Willen,[8] Edmund J. N. Wilson,[3] David R. Winn,[46] Jonathan S. Wurtele,[40,†] Vincent Wu,[47]
Takeichiro Yokoi,[38] Moohyun Yoon,[29] Richard York,[9] Simon Yu,[16] Al Zeller,[9]
Yongxiang Zhao,[8] and Michael S. Zisman[16]

[1]*Illinois Institute of Technology, Physics Division, Chicago, Illinois 60616, USA*
[2]*Fermi National Accelerator Laboratory, P.O. Box 500, Batavia, Illinois 60510, USA*
[3]*CERN, 1211 Geneva 23, Switzerland*
[4]*Department of Physics, University of Wisconsin, Madison, Wisconsin 53706, USA*
[5]*Tel Aviv University, Tel Aviv 69978, Israel*
[6]*Rutherford Appleton Laboratory, Chilton, Didcot, United Kingdom*
[7]*Indiana University, Physics Department, Bloomington, Indiana 47405, USA*
[8]*Brookhaven National Laboratory, Upton, New York 11973, USA*
[9]*Michigan State University, East Lansing, Michigan 48824, USA*
[10]*University of Geneva, DPNC, Quai Ansermet, CH1211 Geneve 4, Switzerland*
[11]*Jefferson Laboratory, 12000 Jefferson Avenue, Newport News, Virginia 23606, USA*
[12]*Istituto Nazionale di Fisica Nucleare, Italy*
[13]*University of Mississippi–Oxford, University, Mississippi 38677, USA*
[14]*Columbia University, Nevis Laboratory, Irvington, New York 10533, USA*
[15]*University of California–Los Angeles, Los Angeles, California 90095, USA*
[16]*Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, California 94720, USA*
[17]*Northern Illinois University, DeKalb, Illinois 60115, USA*
[18]*University of Illinois, at Urbana, Urbana-Champaign, Illinois 61801, USA*
[19]*Princeton University, Department of Astrophysical Sciences, Princeton, New Jersey 08544, USA*

[20]Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831, USA

[21]Institute of Mathematics, Prospect Academician Koptyug 4, 630090 Novosibirsk, Russia

[22]Argonne National Laboratory, Argonne, Illinois 60439, USA

[23]University of California, Davis, California 95616, USA

[24]University of California, Riverside, California 92521, USA

[25]Cornell University, Newman Laboratory for Nuclear Studies, Ithaca, New York 14853, USA

[26]Muons, Inc., Batavia, Illinois 60510, USA

[27]The University of Chicago, Chicago, Illinois 60637, USA

[28]KVI, Rijksuniversiteit, NL 9747 AA Groningen, Netherlands

[29]Pohang University of Science and Technology, Pohang, South Korea

[30]Northwestern University, Department of Physics and Astronomy, Evanston, Illinois 60208, USA

[31]Osaka University, Osaka 567, Japan

[32]Oxford University, Oxford, United Kingdom

[33]University of Hawaii, Department of Physics, Honolulu, Hawaii 96822, USA

[34]Princeton University, Joseph Henry Laboratories, Princeton, New Jersey 08544, USA

[35]Boston University, Boston, Massachusetts 02215, USA

[36]National High Magnetic Field Laboratory, Magnet Science & Technology, Florida 32310, USA

[37]Department of Physics and Astronomy, SUNY, Stony Brook, New York 11790, USA

[38]KEK High Energy Accelerator Research Organization, 1-1 Oho, Tsukuba 305, Japan

[39]University of Iowa, Iowa City, Iowa 52242, USA

[40]University of California, Berkeley, California 94720, USA

[41]Forschungszentrum Karlsruhe, Karlsruhe, Germany

[42]DESY-Zeuthen, Zeuthen, Germany

[43]Budker Institute of Nuclear Physics, 630090 Novosibirsk, Russia

[44]Tech-X Corporation, Boulder, Colorado 80301, USA

[45]Iowa State University, Ames, Iowa 50011, USA

[46]Fairfield University, Fairfield, Connecticut 06430, USA

[47]University of Cincinnati, Cincinnati, Ohio 45221, USA

We describe the status of our effort to realize a first neutrino factory and the progress made in understanding the problems associated with the collection and cooling of muons towards that end. We summarize the physics that can be done with neutrino factories as well as with intense cold beams of muons. The physics potential of muon colliders is reviewed, both as Higgs factories and compact high-energy lepton colliders. The status and time scale of our research and development effort is reviewed as well as the latest designs in cooling channels including the promise of ring coolers in achieving longitudinal and transverse cooling simultaneously. We detail the efforts being made to mount an international cooling experiment to demonstrate the ionization cooling of muons.

I. INTRODUCTION

Recent results from the SNO Collaboration [1] coupled with data from the SuperK Collaboration [2] have provided convincing evidence that neutrinos oscillate and that they very likely do so among the three known neutrino species. Experiments currently under way or planned in the near future will shed further light on the nature of these mixings among neutrino species and the magnitudes of the mass differences between them. Neutrino oscillations and the implied nonzero masses and mixings represent the first experimental evidence of

effects beyond the standard model (SM), and as such are worthy of vigorous scientific study.

This document indicates our progress along a path toward establishing an ongoing program of research in accelerator and experimental physics based on muon beams, and neutrino beams derived therefrom, that can proceed in an incremental fashion. At each step, new-physics vistas open, leading eventually to a neutrino factory and possibly a muon collider. This concept has aroused significant interest throughout the world scientific community. In the U.S., a formal collaboration of some 110 scientists, the neutrino factory and Muon Collider Collaboration, also known as the Muon Collaboration (MC) [3], has undertaken the study of designing a neutrino factory, along with research and development (R&D) activities in support of a muon

*Corresponding author.
  Electronic address: raja@fnal.gov
†Also at Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA.

collider design. The MC comprises three sponsoring national laboratories (BNL, FNAL, LBNL) along with groups from other U.S. national laboratories and universities and individual members from non-U.S. institutions.

One of the first steps toward a neutrino factory is a proton driver that can be used to provide intense beams of conventional neutrinos in addition to providing the intense source of low-energy muons (from pion decay) that must first be "cooled" before being accelerated and stored. Our vision is that while a proton driver is being constructed, R&D on collecting and cooling muons would continue. A source of intense cold muons could be immediately used for physics measurements, such as determining the electric and magnetic dipole moments of the muon to higher precision, muonium-antimuonium oscillations, muon spin rotation experiments, and rare muon decays. Once the capability of cooling and accelerating muons is fully developed, a storage ring for such muons would serve as the first neutrino factory. Its specific beam energy and its distance from the long-baseline experiment will be chosen using the knowledge of neutrino oscillation parameters gleaned from the present generation of solar and accelerator experiments (Homestake, Kamiokande, SuperKamiokande, SAGE, GALLEX, K2K, SNO), the next generation experiments (MiniBooNE, MINOS, CNGS, KamLAND, Borexino), and the high-intensity conventional beam experiments that would already have taken place.

A neutrino factory provides both ν_μ and $\overline{\nu}_e$ beams of equal intensity from a stored μ^- beam and their charge-conjugate beams for a stored μ^+ beam. Beams from a neutrino factory are intense compared with today's neutrino sources. In addition, they have smaller divergence than conventional neutrino beams of comparable energy. These properties permit the study of nonoscillation physics at near detectors, and the measurement of structure functions and associated parameters in nonoscillation physics, to unprecedented accuracy. Likewise, they permit long-baseline experiments that can determine oscillation parameters to unprecedented accuracy.

Depending on the value of the parameter $\sin^2 2\theta_{13}$ in the three-neutrino oscillation formalism, the oscillation $\nu_e \rightarrow \nu_\mu$ is expected to be measurable. By comparing the rates for this channel with its charge-conjugate channel $\overline{\nu}_e \rightarrow \overline{\nu}_\mu$, the sign of the leading mass difference in neutrinos, δm_{32}^2, can be determined by observing the passage through matter of the neutrinos in a long-baseline experiment. Such experiments can also shed light on the *CP*-violating phase, δ, in the lepton mixing matrix and enable the study of *CP* violation in the lepton sector. (It is known that *CP* violation in the quark sector is insufficient to explain the baryon asymmetry of the Universe; lepton sector *CP* violation possibly played a crucial role in creating this asymmetry during the initial phases of the big bang.)

While the neutrino factory is being constructed, R&D aimed at making the muon collider a reality would be performed. The muon collider will require muon beams that are more intensely cooled and have generally more challenging properties than those for a neutrino factory, so the latter forms a practical goal en route to the former. A muon collider, if realized, provides a tool to explore Higgs-like objects by direct *s*-channel fusion, much as the Large Electron-Positron (LEP) Collider at CERN explored the *Z*. It also provides a potential means to reach higher energies (3–4 TeV in the center of mass) using relatively compact collider rings.

A. History

The concept of a muon collider was first proposed by Budker [4] and by Skrinsky [5] in the 1960s and early 1970s. However, additional substance to the concept had to wait until the idea of ionization cooling was developed by Skrinsky and Parkhomchuk [6]. The ionization cooling approach was expanded by Neuffer [7] and then by Palmer *et al.* [8], whose work led to the formation of the neutrino factory and Muon Collider Collaboration [3] in 1995.[3]

The concept of a neutrino source based on a pion storage ring was originally considered by Koshkarev [12]. However, the intensity of the muons created within the ring from pion decay was too low to provide a useful neutrino source. The muon collider concept provided a way to produce a very intense muon source. The physics potential of neutrino beams produced by high-intensity muon storage rings was briefly investigated in 1994 by King [13] and in more detail by Geer in 1997 at a Fermilab workshop [14,15] where it became evident that the neutrino beams produced by muon storage rings needed for the muon collider were exciting in their own right. As a result, the MC realized that a neutrino factory could be an important first step toward a muon collider. With this in mind, the MC has shifted its primary emphasis toward the issues relevant to a neutrino factory. The neutrino factory concept quickly captured the imagination of the particle-physics community, driven in large part by the exciting atmospheric neutrino deficit results from the SuperKamiokande experiment. The utility of nonoscillation neutrino physics from neutrinos produced by muon storage rings has been studied in detail from 1997 onwards [16].

There is also considerable international activity on neutrino factories, with international conferences held at Lyon in 1999 [17], Monterey in 2000 [18], Tsukuba in 2001 [19], London in 2002 [20], and another planned in

---

[3] A good summary of the muon collider concept can be found in the status report of 1999 [9]; an earlier document [10], prepared for Snowmass-1996, is also useful reading. MC notes prepared by the Collaboration are available on the web [11].

New York in 2003 [21]. There are also efforts in Europe [22] and Japan [23] to study different approaches to realizing the neutrino factory. Recently a proposal has been submitted to perform an international muon ionization cooling experiment (MICE) to the Rutherford Appleton Laboratory [24].

B. Feasibility studies

Complementing the MC experimental and theoretical R&D program, which includes work on targetry, cooling, rf hardware (both normal conducting and superconducting), high-field solenoids, liquid-hydrogen absorber design, muon scattering experiments, theory, simulations, parameter studies, and emittance exchange [25], the Collaboration has participated in several paper studies of a complete neutrino factory design.

In the fall of 1999, Fermilab, with help from the MC, undertook a feasibility study ("study I") of an entry-level neutrino factory [26]. Study I showed that the evolution of the Fermilab accelerator complex into a neutrino factory was clearly possible. The performance reached in study I, characterized in terms of the number of 50-GeV muon decays aimed at a detector located 3000 km away from the muon storage ring, was $N = 2 \times 10^{19}$ decays per "Snowmass year" (10^7 s) per MW of protons on target.

Simultaneously, Fermilab launched a study of the physics that might be addressed by such a facility [27] and, more recently, initiated a study to compare the physics reach of a neutrino factory with that of conventional neutrino beams [28] powered by a high-intensity proton driver (referred to as "superbeams"). As will be described later in this paper, a steady and diverse physics program will result from following the evolutionary path from a superbeam to a full-fledged neutrino factory.

Subsequently, BNL organized a follow-on study ("study II") [29] on a high-performance neutrino factory, again in collaboration with the MC. Study II demonstrated that BNL was likewise a suitable site for a neutrino factory. Based on the improvements in study II, the number of 20-GeV muon decays aimed at a detector located 3000 km away from the muon storage ring was $N = 1.2 \times 10^{20}$ decays per Snowmass year per MW of protons on target. Thus, with an upgraded 4 MW proton driver, the muon-decay intensity would increase to 4.8×10^{20} decays per Snowmass year. (Research and development to develop a target capable of handling this beam power would be needed.) Though these numbers of neutrinos are potentially available for experiments, in the current storage ring design the angular divergence at both ends of the production straight section is higher than desirable for the physics program. In any case, we anticipate that storage ring designs are feasible that would allow 30%–40% of the muon decays to provide useful neutrinos.

Both studies I and II are site specific in that each has a few site-dependent aspects; otherwise, they are generic. In particular, study I assumed a new Fermilab booster to achieve its beam intensities and an underground storage ring. Study II assumed BNL site-specific proton driver specifications corresponding to an upgrade of the 24-GeV alternating-gradient synchrotron (AGS) complex and a BNL-specific layout of the storage ring, which is housed in an aboveground berm to avoid penetrating the local water table. The primary substantive difference between the two studies is that study II aimed at a lower muon energy (20 GeV), but higher intensity (for physics reach) than study I. Taking the two feasibility studies together, we conclude that a high-performance neutrino factory could easily be sited at either BNL or Fermilab. Figure 1 shows a comparison of the performance of the neutrino factory designs in studies I and II [27] with the physics requirements.

To put the above performance figures in context, it is important to note that a μ^+ storage ring with an average neutrino energy of 15 GeV and 2×10^{20} useful muon decays would yield (in the absence of oscillations) $\approx 30\,000$ charged-current events in the ν_e channel per kiloton year in a detector located 732 km away. In comparison, a 1.6 MW superbeam [28] from the Fermilab Main Injector with an average neutrino energy of 15 GeV would yield only $\approx 13\,000$ ν_μ charged-current events per kiloton year. In addition to having lower intensity than a neutrino factory beam, a superbeam would have


FIG. 1. (Color) Muon decays in a straight section per 10^7 s vs muon energy, with fluxes required for different physics searches assuming a 50 kT detector. Simulated performance of the two studies is indicated.

significant ν_e contamination, which will be the major background in $\nu_\mu \rightarrow \nu_e$ appearance searches. That is, it will be much easier to detect the oscillation $\nu_e \rightarrow \nu_\mu$ from a muon storage ring neutrino beam than to detect the oscillation $\nu_\mu \rightarrow \nu_e$ from a conventional neutrino beam, because the electron final state from the conventional beam has significant background contribution from π^0's produced in the events.

C. Neutrino factory description

The muons we use result from decays of pions produced when an intense proton beam bombards a high-power production target. The target and downstream transport channel are surrounded by superconducting solenoids to contain the pions and muons, which are produced with a larger spread of transverse and longitudinal momenta than can be conveniently transported through an acceleration system. To prepare a beam suitable for subsequent acceleration, we first perform a "phase rotation," during which the initial large energy spread and small time spread are interchanged using induction linacs. Next, to reduce the transverse momentum spread, the resulting long bunch, with an average momentum of about $250 \text{ MeV}/c$, is bunched into a 201.25-MHz bunch train and sent through an ionization cooling channel consisting of LH_2 energy absorbers interspersed with rf cavities to replenish the energy lost in the absorbers. The resulting beam is then accelerated to its final energy using a superconducting linac to make the beam relativistic, followed by one or more recirculating linear accelerators (RLAs). Finally, the muons are stored in a racetrack-shaped ring with one long straight section aimed at a detector located at a distance of roughly 3000 km. A schematic layout is shown in Fig. 2.

D. Detector

Specifications for the long-baseline neutrino factory detector are rather typical for an accelerator-based neutrino experiment. However, because of the need to maintain a high neutrino rate at these long distances (≈ 3000 km), the detectors considered here are 3–10 times more massive than those in current neutrino experiments.

Several detector options could be considered for the far detector such as the following:

A 50 kt steel-scintillator-proportional-drift-tube (PDT) detector.

A large water-Cherenkov detector, similar to Super-Kamiokande but with either a magnetized water volume or toroids separating smaller water tanks [30].

A massive liquid-argon magnetized detector [31].

For the near detector, a compact liquid-argon time projection chamber (TPC) (similar to the ICARUS detector [32]) could be used. An experiment with a relatively thin Pb target (1 L_{rad}), followed by a standard fixed-target spectrometer could also be considered.

E. Staging scenario

If desired by the particle-physics community, a fast-track plan leading directly to a neutrino factory could be executed. On the other hand, the neutrino factory offers the distinct advantage that it can be built in stages. This could satisfy both programmatic and cost constraints by allowing an ongoing physics program while reducing the annual construction funding needs. Depending on the results of our technical studies and the results of ongoing searches for the Higgs boson, it is hoped that the neutrino factory is really the penultimate stage, to be followed later by a muon collider (e.g., a Higgs factory). Such a collider offers the potential of bringing the energy frontier in particle physics within reach of a moderate-sized machine. Possible stages for the evolution of a muon-beam facility are described in Sec. III J.

F. R&D program

Successful construction of a muon storage ring to provide a copious source of neutrinos requires development of many novel approaches; construction of a high-luminosity muon collider requires even more. It was clear from the outset that the breadth of R&D issues to be dealt with would be beyond the resources available at any single national laboratory or university. For this reason, in 1995, interested members of the high-energy physics and accelerator physics communities formed the MC to coordinate the required R&D efforts nationally. The task of the MC is to define and carry out R&D needed to assess the technical feasibility of constructing initially a muon storage ring that will provide intense neutrino beams aimed at detectors located many thousands of kilometers from the accelerator site, and ultimately a $\mu^+\mu^-$ collider that will carry out fundamental experiments at the energy frontier in high-energy physics.

FIG. 2. (Color) Schematic of the neutrino factory study-II version.

The MC also serves to coordinate muon-related R&D activities of the NSF-sponsored University Consortium (UC) and the state-sponsored Illinois Consortium for Accelerator Research (ICAR), and is the focal point for defining the needs of muon-related R&D to the managements of the sponsoring national laboratories and to the funding agencies (both DOE and NSF). As already noted, though the MC was formed initially to carry out R&D that might lead eventually to the construction of a muon collider, more recently its focus has shifted mainly, but not exclusively, to a neutrino factory.

The MC maintains close contact with parallel R&D efforts under way in Europe (centered at CERN) and in Japan (centered at KEK). Through its international members, the MC also fosters coordination of the international muon-beam R&D effort. Two major initiatives, a targetry experiment (E951) in operation at BNL and a muon-cooling R&D program (MUCOOL), have been launched by the MC. In addition, the Collaboration, working in conjunction with the UC and ICAR in some areas, coordinates substantial efforts in accelerator physics and component R&D to define and assess parameters for feasible designs of muon-beam facilities.

G. Outline of report

In what follows, we give the motivation and a scenario for a staged approach to constructing a neutrino factory and eventually a muon collider. Section II discusses the physics opportunities, starting from conventional superbeams and going to cold muon beams, then a neutrino factory with its near and far detectors, and finally a muon collider. In Sec. III, we describe the components of a neutrino factory, based on the study-II design, and indicate a scientifically productive staged path for reaching it. Section IV covers our present concept of an entry-level Higgs factory muon collider. In support of the construction of a neutrino factory, an R&D program is already under way to address various technical issues. A description of the status and plans for this program is presented in Sec. V. Section VI describes current thinking about a cooling demonstration experiment that would be carried out as an international effort. Finally, in Sec. VII we provide a brief summary of our work.

II. PHYSICS MOTIVATION

In this Section we cover the physics potential of the neutrino factory accelerator complex, which includes superbeams of conventional neutrinos that are possible using the proton driver needed for the factory, and intense beams of cold muons that become available once the muon-cooling and collection systems for the factory are in place. Once the cold muons are accelerated and stored in the muon storage ring, we realize the full potential of the factory in both neutrino oscillation and

nonoscillation physics. Cooling muons will be a learning experience. We hope that the knowledge gained in constructing a neutrino factory can be used to cool muons sufficiently to produce the first muon collider operating as a Higgs factory. We examine the physics capabilities of such a collider, which if realized, will invariably lead to higher energy muon colliders with exciting physics opportunities.

A. Neutrino oscillation physics

Here we discuss [33] the current evidence for neutrino oscillations, and hence neutrino masses and lepton mixing, from solar and atmospheric data. A review is given of some theoretical background including models for neutrino masses and relevant formulas for neutrino oscillation transitions. We next mention the near-term and mid-term experiments in this area and comment on what they hope to measure. We then discuss the physics potential of a muon storage ring as a neutrino factory in the long term.

1. Evidence for neutrino oscillations

In a modern theoretical context, one generally expects nonzero neutrino masses and associated lepton mixing. Experimentally, there has been accumulating evidence for such masses and mixing. All solar neutrino experiments (Homestake, Kamiokande, SuperKamiokande, SAGE, GALLEX, and SNO) show a significant deficit in the neutrino fluxes coming from the Sun [34]. This deficit can be explained by oscillations of the ν_e's into other weak eigenstate(s), with Δm_{sol}^2 of the order 10^{-5} eV^2 for solutions involving the Mikheyev-Smirnov-Wolfenstein (MSW) resonant matter oscillations [35–38] or of the order of 10^{-10} eV^2 for vacuum oscillations (VO) [39]. Accounting for the data with VO requires almost maximal mixing. The MSW solutions include one for small mixing angle (SMA) and one for large mixing angle (LMA). Another piece of evidence for neutrino oscillations is the atmospheric neutrino anomaly, observed by Kamiokande [40], IMB [41], SuperKamiokande [42] with the highest statistics, and by Soudan [43] and MACRO [44]. These data can be fit by the inference of $\nu_\mu \rightarrow \nu_x$ oscillations with $\Delta m_{atm}^2 \sim 3 \times 10^{-3}$ eV^2 [42] and maximal mixing $\sin^2 2\theta_{atm} = 1$. The identification $\nu_x = \nu_\tau$ is preferred over $\nu_x = \nu_{sterile}$, and the identification $\nu_x = \nu_e$ is excluded by both the Superkamiokande data and the Chooz experiment [45]. In addition, the LSND experiment [46] has reported $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ and $\nu_\mu \rightarrow \nu_e$ oscillations with $\Delta m_{LSND}^2 \sim 0.1$–1 eV^2 and a range of possible mixing angles. This result is not confirmed, but also not completely ruled out, by a similar experiment, KARMEN [47]. The miniBOONE experiment at Fermilab is designed to resolve this issue, as discussed below. If one were to try to fit all of these experiments, then, since they

involve three quite different values of $\Delta m_{ij}^2 = m(\nu_i)^2 - m(\nu_j)^2$, which could not satisfy the identity for three neutrino species,

$$\Delta m_{32}^2 + \Delta m_{21}^2 + \Delta m_{13}^2 = 0, \qquad (1)$$

it would follow that one would have to introduce at least one further neutrino. Since it is known from the measurement of the Z width that there are only three leptonic weak doublets with associated light neutrinos, it follows that such further neutrino weak eigenstate(s) would have to be electroweak singlet(s) ("sterile" neutrinos).

Because the LSND experiment has not been confirmed by the KARMEN experiment, we choose here to use only the (confirmed) solar and atmospheric neutrino data in our analysis, and hence to work in the context of three active neutrino weak eigenstates.

2. Neutrino oscillation formalism

In this theoretical context, consistent with solar and atmospheric data, there are three electroweak-doublet neutrinos and the neutrino mixing matrix is described by

$$U = \begin{pmatrix} c_{12}c_{13} & c_{13}s_{12} & s_{13}e^{-i\delta} \\ -c_{23}s_{12} - s_{13}s_{23}c_{12}e^{i\delta} & c_{12}c_{23} - s_{12}s_{13}s_{23}e^{i\delta} & c_{13}s_{23} \\ s_{12}s_{23} - s_{13}c_{12}c_{23}e^{i\delta} & -s_{23}c_{12} - s_{12}c_{23}s_{13}e^{i\delta} & c_{13}c_{23} \end{pmatrix} K', \qquad (2)$$

where $c_{ij} = \cos\theta_{ij}$, $s_{ij} = \sin\theta_{ij}$, and $K' = \text{diag}(1, e^{i\phi_1}, e^{i\phi_2})$. The phases ϕ_1 and ϕ_2 do not affect neutrino oscillations. Thus, in this framework, the neutrino mixing relevant for neutrino oscillations depends on the four angles θ_{12}, θ_{13}, θ_{23}, and δ, and on two independent differences of squared masses, Δm_{atm}^2, which is $\Delta m_{32}^2 = m(\nu_3)^2 - m(\nu_2)^2$ in the favored fit, and Δm_{sol}^2, which may be taken to be $\Delta m_{21}^2 = m(\nu_2)^2 - m(\nu_1)^2$. Note that these Δm^2 quantities involve both magnitude and sign; although in a two-species neutrino oscillation in vacuum the sign does not enter, in the three-species oscillation, which includes both matter effects and CP violation, the signs of the Δm^2 quantities enter and can, in principle, be measured. For our later discussion it will be useful to record the formulas for the various neutrino oscillation transitions. In the absence of any matter effect, the probability that a (relativistic) weak neutrino eigenstate ν_a becomes ν_b after propagating a distance L is

$$P(\nu_a \rightarrow \nu_b) = \delta_{ab} - 4 \sum_{i>j=1}^{3} \text{Re}(K_{ab,ij}) \sin^2\left(\frac{\Delta m_{ij}^2 L}{4E}\right)$$
$$+ 4 \sum_{i>j=1}^{3} \text{Im}(K_{ab,ij}) \sin\left(\frac{\Delta m_{ij}^2 L}{4E}\right) \cos\left(\frac{\Delta m_{ij}^2 L}{4E}\right), \qquad (3)$$

where

$$K_{ab,ij} = U_{ai} U_{bi}^* U_{aj}^* U_{bj} \qquad (4)$$

and

$$\Delta m_{ij}^2 = m(\nu_i)^2 - m(\nu_j)^2. \qquad (5)$$

Recall that in vacuum, CPT invariance implies $P(\bar{\nu}_b \rightarrow \bar{\nu}_a) = P(\nu_a \rightarrow \nu_b)$ and hence, for $b = a$, $P(\bar{\nu}_a \rightarrow \bar{\nu}_a) = P(\nu_a \rightarrow \nu_a)$. For the CP-transformed reaction $\bar{\nu}_a \rightarrow \bar{\nu}_b$ and the T-reversed reaction $\nu_b \rightarrow \nu_a$, the transition probabilities are given by the right-hand side of (3) with the sign of the imaginary term reversed. (Below we shall assume CPT invariance, so that CP violation is

equivalent to T violation.) In most cases there is only one mass scale relevant for long-baseline neutrino oscillations, $\Delta m_{\text{atm}}^2 \sim \text{few} \times 10^{-3} \text{ eV}^2$, and one possible neutrino mass spectrum is the hierarchical one:

$$\Delta m_{21}^2 = \Delta m_{\text{sol}}^2 \ll \Delta m_{31}^2 \approx \Delta m_{32}^2 = \Delta m_{\text{atm}}^2. \qquad (6)$$

In this case, CP (T) violation effects may be negligibly small, so that in vacuum

$$P(\bar{\nu}_a \rightarrow \bar{\nu}_b) = P(\nu_a \rightarrow \nu_b) \qquad (7)$$

and

$$P(\nu_b \rightarrow \nu_a) = P(\nu_a \rightarrow \nu_b). \qquad (8)$$

In the absence of T violation, the second equality (8) would still hold in uniform matter, but even in the absence of CP violation, the first equality (7) would not hold. With the hierarchy (6), the expressions for the specific oscillation transitions are

$$P(\nu_\mu \rightarrow \nu_\tau) = 4|U_{33}|^2 |U_{23}|^2 \sin^2\left(\frac{\Delta m_{\text{atm}}^2 L}{4E}\right)$$
$$= \sin^2(2\theta_{23}) \cos^4(\theta_{13}) \sin^2\left(\frac{\Delta m_{\text{atm}}^2 L}{4E}\right), \qquad (9)$$

$$P(\nu_e \rightarrow \nu_\mu) = 4|U_{13}|^2 |U_{23}|^2 \sin^2\left(\frac{\Delta m_{\text{atm}}^2 L}{4E}\right)$$
$$= \sin^2(2\theta_{13}) \sin^2(\theta_{23}) \sin^2\left(\frac{\Delta m_{\text{atm}}^2 L}{4E}\right), \qquad (10)$$

$$P(\nu_e \rightarrow \nu_\tau) = 4|U_{33}|^2 |U_{13}|^2 \sin^2\left(\frac{\Delta m_{\text{atm}}^2 L}{4E}\right)$$
$$= \sin^2(2\theta_{13}) \cos^2(\theta_{23}) \sin^2\left(\frac{\Delta m_{\text{atm}}^2 L}{4E}\right). \qquad (11)$$

In neutrino oscillation searches using reactor antineutrinos, i.e., tests of $\bar{\nu}_e \rightarrow \bar{\nu}_e$, the two-species mixing hypothesis used to fit the data is

$$P(\nu_e \to \nu_e) = 1 - \sum_x P(\nu_e \to \nu_x)$$
$$= 1 - \sin^2(2\theta_{\text{reactor}})\sin^2\left(\frac{\Delta m^2_{\text{reactor}}L}{4E}\right), \quad (12)$$

where $\Delta m^2_{\text{reactor}}$ is the squared mass difference relevant for $\bar{\nu}_e \to \bar{\nu}_x$. In particular, in the upper range of values of Δm^2_{atm}, since the transitions $\bar{\nu}_e \to \bar{\nu}_\mu$ and $\bar{\nu}_e \to \bar{\nu}_\tau$ contribute to $\bar{\nu}_e$ disappearance, one has

$$P(\nu_e \to \nu_e) = 1 - \sin^2(2\theta_{13})\sin^2\left(\frac{\Delta m^2_{\text{atm}}L}{4E}\right), \quad (13)$$

i.e., $\theta_{\text{reactor}} = \theta_{13}$, and, for the value $|\Delta m^2_{32}| = 3 \times 10^{-3}$ eV^2 from SuperK, the CHOOZ experiment on $\bar{\nu}_e$ disappearance yields the upper limit [45]

$$\sin^2(2\theta_{13}) < 0.1, \quad (14)$$

which is also consistent with conclusions from the SuperK data analysis [42]. Further, the quantity "$\sin^2(2\theta_{\text{atm}})$" often used to fit the data on atmospheric neutrinos with a simplified two-species mixing hypothesis, is, in the three-generation case,

$$\sin^2(2\theta_{\text{atm}}) \equiv \sin^2(2\theta_{23})\cos^4(\theta_{13}). \quad (15)$$

The SuperK experiment finds that the best fit to their data is $\nu_\mu \to \nu_\tau$ oscillations with maximal mixing, and hence $\sin^2(2\theta_{23}) = 1$ and $|\theta_{13}| \ll 1$. The various solutions of the solar neutrino problem involve quite different values of Δm^2_{21} and $\sin^2(2\theta_{12})$: (i) LMA solution: $\Delta m^2_{21} \simeq$ few $\times 10^{-5}$ eV^2 and $\sin^2(2\theta_{12}) \simeq 0.8$; (ii) SMA solution: $\Delta m^2_{21} \sim 10^{-5}$ eV^2 and $\sin^2(2\theta_{12}) \sim 10^{-2}$; (iii) LOW: $\Delta m^2_{21} \sim 10^{-7}$ eV^2, $\sin^2(2\theta_{12}) \sim 1$; and (iv) "just so": $\Delta m^2_{21} \sim 10^{-10}$ eV^2, $\sin^2(2\theta_{12}) \sim 1$. The SuperK experiment favors the LMA solutions [34]; for other global fits, see, e.g., Ref. [34]. We have reviewed the three-neutrino oscillation phenomenology that is consistent with solar and atmospheric neutrino oscillations. In what follows, we will examine the neutrino experiments planned for the immediate future that will address some of the relevant physics. We will then review the physics potential of the neutrino factory.

3. Relevant near- and mid-term experiments

There are currently intense efforts to confirm and extend the evidence for neutrino oscillations in all of the various sectors—solar, atmospheric, and accelerator. Some of these experiments are running; in addition to SuperKamiokande and Soudan-2, these include the Sudbury Neutrino Observatory, SNO, and the K2K long-baseline experiment between KEK and Kamioka. Others are in development and testing phases, such as miniBOONE, MINOS, the CERN–Gran Sasso program, KamLAND, Borexino, and MONOLITH [48]. Among the long-baseline neutrino oscillation experiments, the

approximate distances are $L \simeq 250$ km for K2K, 730 km for both MINOS (from Fermilab to Soudan) and the proposed CERN–Gran Sasso experiments. K2K is a ν_μ disappearance experiment with a conventional neutrino beam having a mean energy of about 1.4 GeV, going from KEK 250 km to the SuperK detector. It has a near detector for beam calibration. It has obtained results consistent with the SuperK experiment and has reported that its data disagree by 2σ with the no-oscillation hypothesis [49]. MINOS is another conventional neutrino beam experiment that takes a beam from Fermilab 730 km to a detector in the Soudan mine in Minnesota. It again uses a near detector for beam flux measurements and has opted for a low-energy configuration, with the flux peaking at about 3 GeV. This experiment is scheduled to start taking data in 2005 and, after some years of running, to obtain higher statistics than the K2K experiment and to achieve a sensitivity down to the level $|\Delta m^2_{32}| \sim 10^{-3}$ eV^2. The CERN–Gran Sasso program will also come on in 2005. It will use a higher energy neutrino beam, $E_\nu \sim 17$ GeV, from CERN to the Gran Sasso deep underground laboratory in Italy. This program will emphasize detection of the τ's produced by the ν_τ's that result from the inferred neutrino oscillation transition $\nu_\mu \to \nu_\tau$. The OPERA experiment will do this using emulsions [50], while the ICARUS proposal uses a liquid-argon chamber [51]. For the joint capabilities of MINOS, ICARUS, and OPERA experiments, see Ref. [52]. Plans for the Japan Hadron Facility (JHF), also called the High Intensity Proton Accelerator, include the use of a 0.77 MW proton driver to produce a high-intensity conventional neutrino beam with a path length of 300 km to the SuperK detector [53]. Moreover, at Fermilab, the miniBOONE experiment is scheduled to start data taking in the near future and to confirm or refute the LSND claim after a few years of running. There are several neutrino experiments relevant to the solar neutrino anomaly. The SNO experiment is currently running and has recently reported their first results that confirm solar neutrino oscillations [1]. These involve measurement of the solar neutrino flux and energy distribution using the charged-current reaction on heavy water, $\nu_e + d \to e + p + p$. They are expected to report on the neutral current reaction $\nu_e + d \to \nu_e + n + p$ shortly. The neutral current rate is unchanged in the presence of oscillations that involve standard model neutrinos, since the neutral current channel is equally sensitive to all three neutrino species. If however, sterile neutrinos are involved, one expects to see a depletion in the neutral current channel also. However, the uncertain normalization of the ^{8}B flux makes it difficult to constrain a possible sterile neutrino component in the oscillations [54]. The KamLAND experiment [55] in Japan started taking data in January 2002. This is a reactor antineutrino experiment using baselines of 100–250 km. It will search for $\bar{\nu}_e$ disappearance and is sensitive to the solar neutrino oscillation

scale. KamLAND can provide precise measurements of the LMA solar parameters [56] and recently the first results from KamLAND have confirmed the LMA solution [57]. A global analysis of the KamLAND and solar neutrino data has further restricted the solar δm^2 range and the best fit value currently is 7×10^{-5} eV^2 [58,59]. On a similar time scale, the Borexino experiment in Gran Sasso is scheduled to turn on and measure the ^{7}Be neutrinos from the Sun. These experiments should help us determine which of the various solutions to the solar neutrino problem is preferred, and hence the corresponding values of Δm_{21}^2 and $\sin^2(2\theta_{12})$. This, then, is the program of relevant experiments during the period 2000–2010. By the end of this period, we may expect that much will be learned about neutrino masses and mixing. However, there will remain several quantities that will not be well measured and which can be measured by a neutrino factory.

4. Oscillation experiments at a neutrino factory

Although a neutrino factory based on a muon storage ring will turn on several years after this near-term period in which K2K, MINOS, and the CERN–Gran Sasso experiments will run, it has a valuable role to play, given the very high-intensity neutrino beams of fixed flavor-pure content, including, uniquely, ν_e and $\bar{\nu}_e$ beams in addition to ν_μ and $\bar{\nu}_\mu$ beams. A conventional positive charge selected neutrino beam is primarily ν_μ with some admixture of ν_e's and other flavors from K decays [O(1%) of the total charged-current rate], and the fluxes of these neutrinos can only be fully understood after measuring the charged particle spectra from the target with high accuracy. In contrast, the potential of the neutrino beams from a muon storage ring is that the neutrino beams would be of extremely high purity: μ^- beams would yield 50% ν_μ and 50% $\bar{\nu}_e$, and μ^+ beams, the charge-conjugate neutrino beams. Furthermore, these could be produced with high intensities and low divergence that make it possible to go to longer baselines.

In what follows, we shall take the design values from study II of 10^{20} μ decays per "Snowmass year" (10^7 s) as

being typical. The types of neutrino oscillations that can be searched for with the neutrino factory based on the muon storage ring are listed in Table I for the case of μ^- which decays to $\nu_\mu e^- \bar{\nu}_e$: It is clear from the processes listed that since the beam contains both neutrinos and antineutrinos, the only way to determine the flavor of the parent neutrino is to determine the identity of the final state charged lepton and measure its charge. A capability unique to the neutrino factory will be the measurement of the oscillation $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$, giving a wrong sign μ^+. Of greater difficulty would be the measurement of the transition $\bar{\nu}_e \rightarrow \bar{\nu}_\tau$, giving a τ^+ which will decay part of the time to μ^+. These physics goals mean that a detector must have excellent capability to identify muons and measure their charges. Especially in a steel-scintillator detector, the oscillation $\nu_\mu \rightarrow \nu_e$ would be difficult to observe, since it would be difficult to distinguish an electron shower from a hadron shower. From the above formulas for oscillations, one can see that, given the knowledge of $|\Delta m_{32}^2|$ and $\sin^2(2\theta_{23})$ that will be available by the time a neutrino factory is built, the measurement of the $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$ transition yields the value of θ_{13}. To get a rough idea of how the sensitivity of an oscillation experiment would scale with energy and baseline length, recall that the event rate in the absence of oscillations is simply the neutrino flux times the cross section. First of all, neutrino cross sections in the region above about 10 GeV (and slightly higher for τ production) grow linearly with the neutrino energy. Secondly, the beam divergence is a function of the initial muon storage ring energy; this divergence yields a flux, as a function of θ_d, the angle of deviation from the forward direction, that goes like $1/\theta_d^2 \sim E^2$. Combining this with the linear E dependence of the neutrino cross section and the overall $1/L^2$ dependence of the flux far from the production region, one finds that the event rate goes like

$$\frac{dN}{dt} \sim \frac{E^3}{L^2}. \tag{16}$$

We base our discussion on the event rates given in the Fermilab neutrino factory study [27]. For a stored muon

TABLE I. Neutrino oscillation modes that can be studied with conventional neutrino beams or with beams from a neutrino factory, with ratings as to the degree of difficulty in each case; $* =$ well or easily measured, $\sqrt{} =$ measured poorly or with difficulty, — $=$ not measured.

| Measurement	Type	Conventional beam	Neutrino factory
$\nu_\mu \rightarrow \nu_\mu$, $\nu_\mu \rightarrow \mu^-$	Survival	$\sqrt{}$	$*$
$\nu_\mu \rightarrow \nu_e$, $\nu_e \rightarrow e^-$	Appearance	$\sqrt{}$	$\sqrt{}$
$\nu_\mu \rightarrow \nu_\tau$, $\nu_\tau \rightarrow \tau^-$, $\tau^- \rightarrow (e^-, \mu^-), \ldots$	Appearance	$\sqrt{}$	$\sqrt{}$
$\bar{\nu}_e \rightarrow \bar{\nu}_e$, $\bar{\nu}_e \rightarrow e^+$	Survival	—	$*$
$\bar{\nu}_e \rightarrow \bar{\nu}_\mu$, $\bar{\nu}_\mu \rightarrow \mu^+$	Appearance	—	$*$
$\bar{\nu}_e \rightarrow \bar{\nu}_\tau$, $\bar{\nu}_\tau \rightarrow \tau^+$, $\tau^+ \rightarrow (e^+, \mu^+), \ldots$	Appearance	—	$\sqrt{}$

FIG. 3. Fit to muon neutrino survival distribution for $E_\mu = 30$ GeV and $L = 2800$ km for ten pairs of $\sin^2 2\theta$, δm^2 values. For each fit, the 1σ, 2σ, and 3σ contours are shown. The generated points are indicated by the dark rectangles and the fitted values by stars. The SuperK 68%, 90%, and 99% confidence levels are superimposed. Each point is labeled by the predicted number of signal events for that point.

energy of 20 GeV, and a distance of $L = 2900$ to the WIPP Carlsbad site in New Mexico, these event rates amount to several thousand events per kton of detector per year, i.e., they are satisfactory for the physics program. This is also true for the other path lengths under consideration, namely, $L = 2500$ km from BNL to Homestake and $L = 1700$ km to Soudan. A usual racetrack design would allow only a single path length L, but a bow-tie design could allow two different path lengths (e.g., [60]). We anticipate that at a time when the neutrino factory turns on, $|\Delta m_{32}^2|$ and $\sin^2(2\theta_{23})$ would be known at perhaps the 10% level (while recognizing that future projections such as this are obviously uncertain). The neutrino factory will significantly improve precision in these parameters, as can be

seen from Fig. 3 which shows the error ellipses possible for a 30 GeV muon storage ring. In addition, the neutrino factory can contribute to the measurement of (i) θ_{13}, as discussed above; (ii) a measurement of the sign of Δm_{32}^2 using matter effects; and (iii) possibly a measurement of *CP* violation in the leptonic sector, if $\sin^2(2\theta_{13})$, $\sin^2(2\theta_{21})$, and Δm_{21}^2 are sufficiently large. To measure the sign of Δm_{32}^2, one uses the fact that matter effects reverse sign when one switches from neutrinos to antineutrinos and carries out this switch in the charges of the stored $\mu^\pm$. We elaborate on this next.

5. Matter effects

With the advent of the muon storage ring, the distances at which one can place detectors are large enough so that for the first time matter effects can be exploited in accelerator-based oscillation experiments. Simply put, matter effects are the matter-induced oscillations that neutrinos undergo along their flight path through the Earth from the source to the detector. Given the typical density of the Earth, matter effects are important for the neutrino energy range $E \sim \mathcal{O}(10)$ GeV and $\Delta m_{32}^2 \sim 10^{-3}$ eV^2, values relevant for the long-baseline experiments. Matter effects in neutrino propagation were first pointed out by Wolfenstein [35] and Barger, Whisnant, Pakvasa, and Phillips [36]. (See the papers [61–76] for details of the matter effects and their relevance to neutrino factories.) In brief, assuming a normal hierarchy, the transition probabilities for propagation through matter of constant density are [74,77]

$$P(\nu_e \to \nu_\mu) = x^2 f^2$$
$$+ 2xyfg(\cos\delta\cos\Delta + \sin\delta\sin\Delta) + y^2 g^2, \tag{17}$$

$$P(\nu_e \to \nu_\tau) = \cot^2\theta_{23} x^2 f^2$$
$$- 2xyfg(\cos\delta\cos\Delta + \sin\delta\sin\Delta)$$
$$+ \tan^2\theta_{23} y^2 g^2, \tag{18}$$

$$P(\nu_\mu \to \nu_\tau) = \sin^2 2\theta_{23}\sin^2\Delta + \alpha\sin 2\theta_{23}\sin 2\Delta\left(\frac{\hat{A}}{1-\hat{A}}\sin\theta_{13}\sin 2\theta_{12}\cos 2\theta_{23}\sin\Delta - \Delta\cos^2\theta_{12}\sin 2\theta_{23}\right), \tag{19}$$

where

$$\Delta \equiv |\delta m_{31}^2|L/4E_\nu$$
$$= 1.27|\delta m_{31}^2/\text{eV}^2|(L/\text{km})/(E_\nu/\text{GeV}), \tag{20}$$

$$\hat{A} \equiv |A/\delta m_{31}^2|, \tag{21}$$

$$\alpha \equiv |\delta m_{21}^2/\delta m_{31}^2|, \tag{22}$$

$$x \equiv \sin\theta_{23}\sin 2\theta_{13}, \tag{23}$$

$$y \equiv \alpha\cos\theta_{23}\sin 2\theta_{12}, \tag{24}$$

$$f \equiv \sin[(1 \mp \hat{A})\Delta]/(1 \mp \hat{A}), \tag{25}$$

$$g \equiv \sin(\hat{A}\Delta)/\hat{A}. \tag{26}$$

The amplitude A for $\nu_e e$ forward scattering in matter is

given by

$$A = 2\sqrt{2}G_F N_e E_\nu$$
$$= 1.52 \times 10^{-4} \text{ eV}^2 Y_e \rho(\text{g/cm}^3)E(\text{GeV}). \quad (27)$$

Here Y_e is the electron fraction and $\rho(x)$ is the matter density. For neutrino trajectories that pass through the Earth's crust, the average density is typically of the order of 3 gm/cm^3 and $Y_e \simeq 0.5$. For neutrinos with $\delta m_{31}^2 > 0$ or antineutrinos with $\delta m_{31}^2 < 0$, $\hat{A} = 1$ corresponds to a matter resonance. Thus, for a neutrino factory operating with positive stored muons (producing a ν_e beam) one expects an enhanced production of opposite sign (μ^-) charged-current events as a result of the oscillation $\nu_e \rightarrow \nu_\mu$ if δm_{32}^2 is positive and vice versa for stored negative beams. Figure 4 [73] shows the wrong-sign muon appearance spectra as a function of δm_{32}^2 for both μ^+ and μ^- beams for both signs of δm_{32}^2 at a baseline of 2800 km. The resonance enhancement in wrong-sign muon production is clearly seen in Figs. 4(b) and 4(c). By comparing these (using first a stored μ^+ beam and then a stored μ^- beam) one can thus determine the sign of Δm_{32}^2 as well as the value of $\sin^2(2\theta_{13})$. Figure 5 [73] shows the difference in negative log-likelihood between a correct and wrong-sign mass hypothesis expressed as a number of equivalent Gaussian standard deviations versus baseline length for muon storage ring energies of 20, 30, 40, and 50 GeV. The values of the oscillation parameters are for the LMA scenario with $\sin^2 2\theta_{13} = 0.04$. Figure 5(a) is for 10^{20}

decays for each sign of stored energy and a 50 kt detector and positive δm_{32}^2; 5(b) is for negative δm_{32}^2 for various values of stored muon energy. Figures 5(c) and 5(d) show the corresponding curves for 10^{19} decays and a 50 kt detector. An entry-level machine would permit one to perform a 5σ differentiation of the sign of δm_{32}^2 at a baseline length of ~ 2800 km. For the study-II design, in accordance with the previous Fermilab study [27], one estimates that it is possible to determine the sign of δm_{32}^2 even if $\sin^2(2\theta_{13})$ is as small as $\sim 10^{-3}$.

6. CP violation

CP violation is measured by the (rephasing-invariant) product

$$J = \text{Im}(U_{ai}U_{bi}^*U_{aj}^*U_{bj})$$
$$= \frac{1}{8}\sin(2\theta_{12})\sin(2\theta_{13})\cos(\theta_{13})\sin(2\theta_{23})\sin\delta. \quad (28)$$

Leptonic *CP* violation also requires that each of the leptons in each charge sector be nondegenerate with any other leptons in this sector; this is, of course, true of the charged lepton sector and, for the neutrinos, this requires $\Delta m_{ij}^2 \neq 0$ for each pair ij. In the quark sector, J is known to be small: $J_{\text{CKM}} \sim \mathcal{O}(10^{-5})$. A promising asymmetry to measure is $P(\nu_e \rightarrow \nu_\mu) - P(\bar{\nu}_e - \bar{\nu}_\mu)$. As an illustration, in the absence of matter effects,

FIG. 4. (Color) The wrong-sign muon appearance rates for a 20 GeV muon storage ring at a baseline of 2800 km with 10^{20} decays and a 50 kt detector for (a) μ^+ stored and negative δm_{32}^2, (b) μ^- stored and negative δm_{32}^2, (c) μ^+ stored and positive δm_{32}^2, (d) μ^- stored and positive δm_{32}^2. The values of $|\delta m_{32}^2|$ range from 0.0005 to 0.0050 eV^2 in steps of 0.0005 eV^2. Matter enhancements are evident in (b) and (c).

FIG. 5. The statistical significance (number of standard deviations) with which the sign of δm_{32}^2 can be determined versus baseline length for various muon storage ring energies. The results are shown for a 50 kt detector, and (a) 10^{20} μ^+ and μ^- decays and positive values of δm_{32}^2; (b) 10^{20} μ^+ and μ^- decays and negative values of δm_{32}^2; (c) 10^{19} μ^+ and μ^- decays and positive values of δm_{32}^2; (d) 10^{19} μ^+ and μ^- decays and negative values of δm_{32}^2.

$$P(\nu_e \to \nu_\mu) - P(\bar{\nu}_e \to \bar{\nu}_\mu) = 4J(\sin 2\phi_{32} + \sin 2\phi_{21}$$
$$+ \sin 2\phi_{13})$$
$$= -16J \sin \phi_{32} \sin \phi_{13} \sin \phi_{21}, \tag{29}$$

where

$$\phi_{ij} = \frac{\Delta m_{ij}^2 L}{4E}. \tag{30}$$

In order for the *CP* violation in Eq. (29) to be large enough to measure, it is necessary that θ_{12}, θ_{13}, and $\Delta m_{\text{sol}}^2 = \Delta m_{21}^2$ not be too small. From atmospheric neutrino data, we have $\theta_{23} \simeq \pi/4$ and $\theta_{13} \ll 1$. If LMA describes solar neutrino data, then $\sin^2(2\theta_{12}) \simeq 0.8$, so $J \simeq 0.1 \sin(2\theta_{13}) \sin \delta$. For example, if $\sin^2(2\theta_{13}) = 0.04$, then J could be $\gg J_{\text{CKM}}$. Furthermore, for parts of the LMA phase space where $\Delta m_{\text{sol}}^2 \sim 4 \times 10^{-5}$ eV^2, the *CP*-violating effects might be observable. In the absence of matter, one would measure the asymmetry

$$\frac{P(\nu_e \to \nu_\mu) - P(\bar{\nu}_e \to \bar{\nu}_\mu)}{P(\nu_e \to \nu_\mu) + P(\bar{\nu}_e \to \bar{\nu}_\mu)} = \frac{\sin(2\theta_{12}) \cot(\theta_{23}) \sin \delta \sin \phi_{21}}{\sin \theta_{13}}. \tag{31}$$

However, in order to optimize this ratio, because of the smallness of Δm_{21}^2 even for the LMA, one must go to large path lengths L, and here matter effects are important. These make leptonic *CP* violation challenging to measure, because, even in the absence of any intrinsic *CP* violation, these matter effects render the rates for $\nu_e \to \nu_\mu$ and $\bar{\nu}_e \to \bar{\nu}_\mu$ unequal since the matter interaction is opposite in sign for ν and $\bar{\nu}$. One must therefore subtract out the matter effects in order to try to isolate the intrinsic *CP* violation. Alternatively, one might think of comparing $\nu_e \to \nu_\mu$ with the time-reversed reaction $\nu_\mu \to \nu_e$. Although this would be equivalent if *CPT* is valid, as we assume, and although uniform matter effects are the same here, the detector response is quite different and, in particular, it is quite difficult to identify $e^\pm$. Results from SNO and KamLAND testing the LMA [56] will help further planning. The neutrino factory provides an ideal set of controls to measure *CP*-violation effects since we can fill the storage ring with either μ^+ or μ^- particles and measure the ratio of the number of events $\bar{\nu}_e \to \bar{\nu}_\mu / \nu_e \to \nu_\mu$. Figure 6 shows this ratio for a neutrino factory with 10^{21} decays and a 50 kt detector as a function of the baseline length. The ratio depends on the sign of δm_{32}^2. The shaded band around either curve shows the variation of this ratio as a function of the *CP*-violating phase δ. The number of decays needed to produce the error bars shown is directly proportional to $\sin^2 \theta_{13}$, which for the present example is set to 0.004. Depending on the magnitude of J, one may be driven to build a neutrino factory just to understand *CP* violation in the lepton


FIG. 6. (Color) Predicted ratios of wrong-sign muon event rates when positive and negative muons are stored in a 20 GeV neutrino factory, shown as a function of baseline. A muon measurement threshold of 4 GeV is assumed. The lower and upper bands correspond, respectively, to negative and positive δm_{32}^2. The widths of the bands show how the predictions vary as the *CP*-violating phase δ is varied from $-\pi/2$ to $\pi/2$, with the thick lines showing the predictions for $\delta = 0$. The statistical error bars correspond to a high-performance neutrino factory yielding a data sample of 10^{21} decays with a 50 kt detector. The curves are based on calculations presented in [72].

sector, which could have a significant role in explaining the baryon asymmetry of the Universe [78].

B. Physics potential of superbeams

It is possible to extend the reach of the current conventional neutrino experiments by enhancing the capabilities of the proton sources that drive them. These enhanced neutrino beams have been termed superbeams and they form an intermediate step on the way to a neutrino factory. Their capabilities have been explored in recent papers [28,79,80]. These articles consider the capabilities of enhanced proton drivers at (i) the proposed 0.77 MW 50 GeV proton synchrotron at the JHF [53], (ii) a 4 MW upgraded version of the JHF, (iii) a new ~1 MW 16 GeV proton driver [81] that would replace the existing 8 GeV Booster at Fermilab, or (iv) a fourfold intensity upgrade of the 120 GeV Fermilab Main Injector (MI) beam (to 1.6 MW) that would become possible once the upgraded (16 GeV) booster was operational. Note that the 4 MW 50 GeV JHF and the 16 GeV upgraded Fermilab Booster are both suitable proton drivers for a neutrino factory. The conclusions of both reports are that superbeams will

extend the reaches in the oscillation parameters of the current neutrino experiments but "the sensitivity at a neutrino factory to *CP* violation and the neutrino mass hierarchy extends to values of the amplitude parameter $\sin^2 2\theta_{13}$ that are 1 to 2 orders of magnitude lower than at a superbeam" [79,80].

To illustrate these points, we choose one of the most favorable superbeam scenarios studied: a 1.6 MW NuMI-like high-energy beam with $L = 2900$ km, detector parameters corresponding to the liquid-argon scenario in [79,80], and oscillation parameters $|\delta m_{32}^2| = 3.5 \times 10^{-3}$ eV^2 and $\delta m_{21}^2 = 1 \times 10^{-4}$ eV^2. The calculated three-sigma error ellipses in the $[N(e^+), N(e^-)]$ plane are shown in Fig. 7 for both signs of δm_{32}^2, with the curves corresponding to various *CP* phases δ (as labeled). The magnitude of the $\nu_\mu \to \nu_e$ oscillation amplitude parameter $\sin^2 2\theta_{13}$ varies along each curve, as indicated. The two groups of curves, which correspond to the two signs of δm_{32}^2, are separated by more than 3σ provided $\sin^2 2\theta_{13} \gtrsim 0.01$. Hence the mass hierarchy can be determined provided the $\nu_\mu \to \nu_e$ oscillation amplitude is not more than an order of magnitude below the currently excluded region. Unfortunately, within each group of curves, the *CP*-conserving predictions are separated

from the maximal *CP*-violating predictions by at most 3σ. Hence, it will be difficult to conclusively establish *CP* violation in this scenario. Note for comparison that a very long-baseline experiment at a neutrino factory would be able to observe $\nu_e \to \nu_\mu$ oscillations and determine the sign of δm_{32}^2 for values of $\sin^2 2\theta_{13}$ as small as $\mathcal{O}(0.0001)$. This is illustrated in Fig. 8. A neutrino factory thus outperforms a conventional superbeam in its ability to determine the sign of δm_{32}^2. Comparing Figs. 7 and 8 one sees that the value of $\sin^2 2\theta_{13}$, which has yet to be measured, will determine the parameters of the first neutrino factory.

Finally, we compare the superbeam $\nu_\mu \to \nu_e$ reach with the corresponding neutrino factory $\nu_e \to \nu_\mu$ reach in Fig. 9, which shows the 3σ sensitivity contours in the $(\delta m_{21}^2, \sin^2 2\theta_{13})$ plane. The superbeam $\sin^2 2\theta_{13}$ reach of a few $\times 10^{-3}$ is almost independent of the subleading scale δm_{21}^2. However, since the neutrino factory probes oscillation amplitudes $O(10^{-4})$ the subleading effects cannot be ignored, and $\nu_e \to \nu_\mu$ events would be observed at a neutrino factory over a significant range of δm_{21}^2 even if $\sin^2 2\theta_{13} = 0$.

C. Nonoscillation physics at a neutrino factory

The study of the utility of intense neutrino beams from a muon storage ring in determining the parameters governing nonoscillation physics was begun in 1997 [14]. More complete studies can be found in [27] and recently

FIG. 7. Three-sigma error ellipses in the $[N(e^+), N(e^-)]$ plane, shown for $\nu_\mu \to \nu_e$ and $\bar{\nu}_\mu \to \bar{\nu}_e$ oscillations in a NuMI-like high-energy neutrino beam driven by a 1.6 MW proton driver. The calculation assumes a liquid-argon detector with the parameters listed in [28], a baseline of 2900 km, and 3 years of running with neutrinos, 6 years running with antineutrinos. Curves are shown for different *CP* phases δ (as labeled), and for both signs of δm_{32}^2 with $|\delta m_{32}^2| = 0.0035$ eV^2, and the sub-leading scale $\delta m_{21}^2 = 10^{-4}$ eV^2. Note that $\sin^2 2\theta_{13}$ varies along the curves from 0.001 to 0.1, as indicated [79].

FIG. 8. Three-sigma error ellipses in the $[N(\mu+), N(\mu-)]$ plane, shown for a 20 GeV neutrino factory delivering 3.6×10^{21} useful muon decays and 1.8×10^{21} antimuon decays, with a 50 kt detector at $L = 7300$ km, $\delta m_{21}^2 = 10^{-4}$ eV^2, and $\delta = 0$. Curves are shown for both signs of δm_{32}^2; $\sin^2 2\theta_{13}$ varies along the curves from 0.0001 to 0.01, as indicated [79].

FIG. 9.   Summary of the 3σ level sensitivities for the observation of $\nu_\mu \to \nu_e$ at various MW-scale superbeams (as indicated) with liquid-argon "A" and water-Cerenkov "W" detector parameters, and the observation of $\nu_e \to \nu_\mu$ in a 50 kt detector at 20, 30, 40, and 50 GeV neutrino factories delivering 2×10^{20} muon decays in the beam-forming straight section. The limiting 3σ contours are shown in the $(\delta m_{21}^2, \sin^2 2\theta_{13})$ plane. All curves correspond to 3 yr of running. The grey shaded area is already excluded by current experiments.

a European group has brought out an extensive study on this topic [82]. A neutrino factory can measure individual parton distributions within the proton for all light quarks and antiquarks. It could improve valence distributions by an order of magnitude in the kinematical range $x \gtrsim 0.1$ in the unpolarized case. The individual components of the sea ($\bar{u}$, $\bar{d}$, s, and $\bar{s}$), as well as the gluon, would be measured with relative accuracies in the range of 1%–10%, for $0.1 \lesssim x \lesssim 0.6$. A full exploitation of the neutrino factory potential for polarized measurements of the shapes of individual partonic densities requires an *a priori* knowledge of the polarized gluon density. The forthcoming set of polarized deep inelastic scattering experiments at CERN, DESY, and RHIC may provide this information. The situation is also very bright for measurements of C-even distributions. Here, the first moments of singlet, triplet, and octet axial charges can be measured with accuracies that are up to 1 order of magnitude better than the current uncertainties. In particular, the improvement in the determination of the singlet axial charge would allow a definitive confirmation or refutation of the anomaly scenario compared to the "instanton" or "skyrmion" scenarios, at least if the theoretical uncertainty originating from the small-x extrapolation can be

kept under control. The measurement of the octet axial charge with a few percent uncertainty will allow a determination of the strange contribution to the proton spin better than 10%, and allow stringent tests of models of SU(3) violation when compared to the direct determination from hyperon decays. A measurement of $\alpha_S(M_Z)$ and $\sin^2 \theta_W$ will involve different systematics from current measurements and will therefore provide an important consistency check of current data, although the accuracy of these values is not expected to be improved. The weak mixing angle can be measured in both the hadronic and leptonic modes with a precision of approximately 2×10^{-4}, dominated by the statistics and the luminosity measurement. This determination would be sensitive to different classes of new-physics contributions. Neutrino interactions are a very good source of clean, sign-tagged charm particles. A neutrino factory can measure charm production with raw event rates up to 100×10^6 charm events per year with $\simeq 2 \times 10^6$ double-tagged events. (Note that charm production becomes significant for storage ring energies above 20 GeV.) Such large samples are suitable for precise extractions of branching ratios and decay constants, the study of spin-transfer phenomena, and the study of nuclear effects in deep inelastic scattering. The ability to run with both hydrogen and heavier targets will provide rich data sets useful for quantitative studies of nuclear models. The study of Λ polarization both in the target and in the fragmentation regions will help clarify the intriguing problem of spin transfer.

Although the neutrino beam energies are well below any reasonable threshold for new physics, the large statistics makes it possible to search for physics beyond the standard model. The high-intensity neutrino beam allows a search for the production and decay of neutral heavy leptons with mixing angle sensitivity 2 orders of magnitude better than present limits in the 30–80 MeV range. The exchange of new gauge bosons decoupled from the first generation of quarks and leptons can be seen via enhancements of the inclusive charm production rate, with a sensitivity well beyond the present limits. A novel neutrino magnetic moment search technique that uses oscillating magnetic fields at the neutrino beam source could discover large neutrino magnetic moments predicted by some theories. Rare lepton-flavor-violating decays of muons in the ring could be tagged in the deep inelastic scattering final states through the detection of wrong-sign electrons and muons, or of prompt taus.

D. Physics that can be done with intense cold muon beams

Experimental studies of muons at low and medium energies have had a long and distinguished history, starting with the first search for muon decay to electron plus gamma ray [83], and including along the way the 1957 discovery of the nonconservation of parity, in which the g

TABLE II.   Experiments which could beneficially take advantage of the intense future stopped muon source. The numbers were worked out for scenarios at a future stopped muon source (SMS) of a neutrino factory at CERN [87]. They are based on a muon flux of 10^{21} particles per annum in which beam will be available for 10^7 s. Typical beam requirements are given in Table III.

Type of experiment	Physics issues	Possible experiments	Previously established accuracy	Present activities (proposed accuracy)	Projected for SMS @ CERN
"Classical" rare and forbidden decays	Lepton number violation; searches for new physics: SUSY, *L-R* symmetry, *R*-parity violation, …	$\mu^- N \rightarrow e^- N$ $\mu \rightarrow e\gamma$ $\mu \rightarrow eee$ $\mu^+ e^- \rightarrow \mu^- e^+$	6.1×10^{-13} 1.2×10^{-11} 1.0×10^{-12} 8.1×10^{-11}	PSI, proposed BNL (5×10^{-17}) Proposed PSI (1×10^{-14}) Completed 1985 PSI Completed 1999 PSI	$<10^{-18}$ $<10^{-15}$ $<10^{-16}$ $<10^{-13}$
Muon decays	G_F; searches for new physics; Michel parameters	τ_μ non $(V - A)$	18×10^{-6} Typically, few 10^{-3}	PSI ($2\times$), RAL (1×10^{-6}) PSI, TRIUMF (1×10^{-3})	$<10^{-7}$ $<10^{-4}$
Muon moments	Standard model tests; new physics; *CPT* tests, *T*, respectively; *CP* violation in 2nd lepton generation	$g_\mu - 2$ edm_μ	1.3×10^{-6} $3.4 \times 10^{-19} e$ cm	BNL (3.5×10^{-7}) Proposed BNL ($10^{-24} e$ cm)	$<10^{-7}$ $<5 \times 10^{-26} e$ cm
Muonium spectroscopy	Fundamental constants, μ_μ, m_μ, α; weak interactions; muon charge	M_{HFS} M_{1s2s}	12×10^{-9} 1×10^{-9}	Completed 1999 LAMPF Completed 2000 RAL	5×10^{-9} $<10^{-11}$
Muonic atoms	Nuclear charge radii; weak interactions	μ^- atoms	Depends	PSI, possible CERN ($\langle r_p \rangle$ to 10^{-3})	New nuclear structure
Condensed matter	Surfaces, catalysis, bio sciences, …	Surface μSR	n/a	PSI, RAL (n/a)	High rate

value and magnetic moment of the muon were first measured [84]. The years since then have brought great progress: limits on the standard-model-forbidden decay $\mu \rightarrow e\gamma$ have dropped by 9 orders of magnitude, and the muon anomalous magnetic moment $a_\mu = (g_\mu - 2)/2$ has yielded one of the more precise tests (≈ 1 ppm) of physical theory [85]. The front end of a neutrino factory has the potential to provide $\sim 10^{21}$ muons/yr, 5 orders of magnitude beyond the most intense beam currently available.[4]

Such a facility could enable a rich variety of precision measurements. In the area of low-energy muon physics a majority of experiments with a high physics potential is limited at present by statistics. The list of conceivable projects includes the following (see Table II): precise determinations of the properties characterizing the muon, which are the mass m_μ, magnetic moment μ_μ, magnetic anomaly a_μ, charge q_μ, and lifetime τ_μ; measurements of the muon-decay parameters (Michel parameters); *CPT* tests from a comparison of μ^- and μ^+ properties; measurements of fundamental constants of general importance (e.g., the electromagnetic fine structure constant α or the weak interaction Fermi constant G_F; sensitive searches for physics beyond the standard model either through measuring differences of muon parameters from standard model predictions or in dedicated searches for rare and forbidden processes, such as $\mu \rightarrow e\gamma$, $\mu \rightarrow eee$, $\mu^- N \rightarrow e^- N$ conversion, and muonium-antimuonium (M − $\overline{\text{M}}$) conversion or searches for a permanent electric dipole moment d_μ of the particle;

searches for *P* and *T* violation in muonic atoms; precise determinations of nuclear properties in muonic (radioactive) atoms; applications in condensed matter, thin films, and at surfaces; applications in life sciences; and muon catalyzed fusion (μCF).

A detailed evaluation of the possibilities has recently been made by a CERN study group, which assumed a facility with a 4 MW proton driver [87].

In the search for "forbidden" decays, Marciano [88] has suggested that muon lepton-flavor violation (LFV) (especially coherent muon-to-electron conversion in the field of a nucleus) is the "best bet" for discovering signatures of new physics using low-energy muons. The MECO experiment [89] proposed at BNL offers, through a novel detector concept, very high sensitivity and some 4 orders of magnitude improvement over the current limits from PSI [90]. At a future high muon flux facility, such as the neutrino factory, this could be improved further by 1–2 orders of magnitude.

The search for $\mu \rightarrow e\gamma$ is also of great interest. The MEGA experiment recently set an upper limit $B(\mu^+ \rightarrow e^+ \gamma) < 1.2 \times 10^{-11}$ [91]. Ways to extend sensitivity to the 10^{-14} level have not only been discussed [92] but also have led to an active proposal at PSI [93]. The experiment aims for an improvement of 3 orders of magnitude over MEGA which had systematics limitations. The μ-to-e-conversion approach has the additional virtue of sensitivity to new physics that does not couple to the photon.

An observation of a nonzero electric dipole moment (EDM) of the muon, d_μ, could prove equally exciting; This has generated a Letter of Intent [94] to observe d_μ,

---

[4]The beam at PSI, Villigen, providing a maximum rate of [86].

TABLE III.   Beam requirements for new muon experiments. We show the needed muonic charge q_μ and the minimum of the total muon number $\int I_\mu dt$ above which significant progress can be expected in the physical interpretation of the experiments. Measurements which require pulsed beams are sensitive to the muon suppression I_0/I_m between pulses of length δT and separation ΔT. Most experiments require energies up to 4 MeV corresponding to 29 MeV/c momentum. Thin targets, respectively, storage ring acceptances, demand rather small momentum spreads $\Delta p_\mu / p_\mu$ [87].

Experiment	q_μ	$\int I_\mu dt$	I_0/I_μ	δT (ns)	ΔT (ns)	E_μ (MeV)	$\Delta p_\mu / p_\mu$ (%)
$\mu^- N \to e^- N$	$-$	10^{19}	$<10^{-9}$	≤ 100	≥ 1000	<20	$1 \cdots 5$
$\mu \to e\gamma$	$+$	10^{16}	n/a	Continuous	Continuous	$1 \cdots 4$	$1 \cdots 5$
$\mu \to eee$	$+$	10^{15}	n/a	Continuous	Continuous	$1 \cdots 4$	$1 \cdots 5$
$\mu^+ e^- \to \mu^- e^+$	$+$	10^{16}	$<10^{-4}$	<1000 s	≥ 20000	$1 \cdots 4$	$1 \cdots 2$
τ_μ	$+$	10^{13}	$<10^{-4}$	<100	≥ 20000	4	$1 \cdots 10$
non$(V - A)$	$\pm$	10^{13}	n/a	Continuous	Continuous	4	$1 \cdots 5$
$g_\mu - 2$	$\pm$	10^{15}	$<10^{-7}$	≤ 50	$\geq 10^6$	3100	10^{-4}
d_μ	$\pm$	10^{16}	$<10^{-6}$	≤ 50	$\geq 10^6$	≤ 1000	$\leq 10^{-5}$
M_{HFS}	$+$	10^{15}	$<10^{-4}$	≤ 1000	≥ 20000	4	$1 \cdots 3$
M_{1s2s}	$+$	10^{14}	$<10^{-3}$	≤ 500	$\geq 10^6$	$1 \cdots 4$	$1 \cdots 2$
μ^- atoms	$-$	10^{14}	$<10^{-3}$	≤ 500	≥ 20000	$1 \cdots 4$	$1 \cdots 5$
Condensed matter (including bio sciences)	$\pm$	10^{14}	$<10^{-3}$	<50	≥ 20000	$1 \cdots 4$	$1 \cdots 5$

which proposes to use the large electric fields associated with relativistic particles in a magnetic storage ring. As CP violation enters in the quark sector starting with the second generation, the muon is a particularly valuable probe in this regard, despite the already low limits for electrons. Moreover, there exist some models in which the electric dipole moment scales stronger than linearly [95].

It is worth noting that for searches of rare muon decays and for d_μ that the standard model predictions are orders of magnitude below the presently established limits. Any observation which can be shown to be not an artifact of the experimental method or due to background would therefore be a direct sign of new physics.

There are three experiments going on currently to improve the muon lifetime τ_μ [96]. Note that the Fermi coupling constant G_F is derived from a measurement of τ_μ. The efforts are therefore worthwhile whenever experimental conditions allow substantial improvement. One should however be aware that a comparison with theory in this channel is presently dominated by theoretical uncertainties.

In the case of precision measurements (τ_μ, a_μ, etc.), new-physics effects appear as small corrections arising from the virtual exchange of new massive particles in loop diagrams. In contrast, LFV and EDMs are forbidden in the standard model, thus their observation at any level would constitute evidence for new physics.

The current status and prospects for advances in these areas are shown in Table II, which lists present efforts in the field and possible improvements at a neutrino factory or muon collider facility. The beam parameters necessary for the expected improvements are listed in Table III.

It is worth recalling that LFV as a manifestation of neutrino mixing is suppressed as $(\delta m^2)^2 / m_W^4$ and is thus entirely negligible. However, a variety of new-physics scenarios predict observable effects. Table IV lists some examples of limits on new physics that would be implied by nonobservation of μ-to-e conversion ($\mu^- N \to e^- N$) at the 10^{-16} level [88].

The muon magnetic anomaly (muon g-2 value [97]) has been measured recently at the Brookhaven National Laboratory (BNL) with 0.7 ppm accuracy [85]. At present, no definite statement can be made whether this result agrees or disagrees with standard theory, which is sensitive to electroweak corrections. The theory has recently come under severe scrutiny and in particular an error was found in the calculation of hadronic light by light scattering [98]. The theoretical calculations are being improved upon, and with more data, there is a good chance that this might eventually lead to evidence far beyond the standard model effects [99]. The final goal of the experimental precision is 0.35 ppm for the current set of experiments. This value could be improved by an order of magnitude at a neutrino factory, provided cold muons of energy 3.1 GeV are made available. This could further spur more accurate theoretical calculations that improve upon contributions from hadronic vacuum

TABLE IV.   Some examples of new physics probed by the nonobservation of $\mu \to e$ conversion at the 10^{-16} level (from [88]).

New physics	Limit
Heavy neutrino mixing	$\lvert V_{\mu N}^* V_{eN} \rvert^2 < 10^{-12}$
Induced $Z\mu e$ coupling	$g_{Z_{\mu e}} < 10^{-8}$
Induced $H\mu e$ coupling	$g_{H_{\mu e}} < 4 \times 10^{-8}$
Compositeness	$\Lambda_c > 3000$ TeV

polarization and hadronic light by light scattering [100]. In addition, the muon g-2 experiments at CERN have provided the best test of *CPT* invariance at a level of 2×10^{-22} which is more than 3 orders of magnitude better than the mostly quoted result $K^0 - \overline{K}^0$ mass difference [101]. A $g - 2$ measurement at the neutrino factory front end that uses muons of both charges would lead to further improvement in these *CPT* limits.

Precision studies of atomic electrons have provided notable tests of QED (e.g., the Lamb shift in hydrogen) and could in principle be used to search for new physics were it not for nuclear corrections. Studies of muonium ($\mu^+ e^-$) are free of such corrections since it is a purely leptonic system. Muonic atoms can also yield new information complementary to that obtained from electronic atoms. A number of possibilities have been enumerated by Kawall *et al.* [102], Jungmann [103], and Molzon [104].

By making measurements on the muonium system, for instance, one can produce precise measurements of the fundamental constants and also do sensitive searches for new physics. The muonium ground state hyperfine structure has been measured to 12 ppb [105] and currently furnishes the most sensitive test of the relativistic two-body bound state in QED [103]. The precision could be further improved significantly with increased statistics. The theoretical error is 120 ppb. The uncertainty arising from the muon mass is 5 times larger than that from calculations. If one assumes the theory to be correct, the muon-electron mass ratio can be extracted to 27 ppb. A precise value for the electromagnetic fine structure constant α can be extracted. Its good agreement with the number extracted from the electron magnetic anomaly must be viewed as the best test of the internal consistency of QED, as one case involves bound state QED and the other that of free particles. The Zeeman effect of the muonium hyperfine structure allows the best direct measurement of the muon magnetic moment, respectively, its mass, to 120 ppb, improved by higher precision measurements in muonium and muon spin resonance. These are also areas in which the neutrino factory front end could contribute. Laser spectroscopy of the muonium $1s$-$2s$ transition [106] has resulted in a precise value of the muon mass as well as the testing of the muon-electron charge ratio to about 2×10^{-9}. This is by far the best test of charge equality in the first two generations.

The search for muonium-antimuonium conversion had been proposed by Pontecorvo three years before the system was first produced by Hughes *et al.* [107]. Several new-physics models allow violation of the lepton family number by two units. The current limit is $R_g \equiv G_C/G_F < 0.0030$ [108], where G_C is the new-physics coupling constant, and G_F the Fermi coupling constant. This sets a lower limit of 2.6 TeV/c^2 (90% C.L.) on the mass of a grand-unified dileptonic gauge boson and also strongly disfavors models with heavy lepton seeded radiative mass generation [108]. The search for muonium-antimuonium

conversion has by far the strongest gain in sensitivity of all rare muon-decay experiments [103].

The high-intensity proton machine needed for the neutrino factory can also find use as a new generation isotope facility which would have much higher rates compared to the present ISOLDE facility at CERN. Nucleids not yet studied could be produced at quantities which allow precision investigations of their properties [87]. The measurements of muonic spectra can yield most precise values for the charge radii of nuclei as well as other ground state properties such as moments and even $B(E2)$ transition strengths for even-even nuclei. An improved understanding of nuclear structure can be expected which may be of significance for interpreting various neutrino experiments, rare decays involving nuclei, and nuclear capture. An urgent need exists for accurate charge and neutron radii of francium and radium isotopes which are of interest for atomic parity violation research and *EDM* searches in atoms and nuclei.

Muonic x-ray experiments generally promise higher accuracy for most of these quantities compared to electron scattering, particularly because the precision of electron scattering data depends on the location of the minimum of the cross section where rates are naturally low. In principle, for chains of isotopes charge radii can be inferred from isotope shift measurements with laser spectroscopy. However, this gives only relative information. For absolute values, calibration is necessary and has been obtained in the past for stable nuclei from muonic spectra. In general, two not too distant nuclei are needed for a good calibration.

The envisaged experimental approaches include (i) the technique pioneered by Strasser and Nagamine [109], which involves cold films for keeping radioactive atoms and as a host material in which muon transfer takes place; (ii) merging beams of radioactive ions and of muons; and (iii) trapping of exotic isotopes in a Penning trap which is combined with a cyclotron trap. Large formation rates can be expected from a setup containing a Penning trap [110], the magnetic field of which also serves as a cyclotron muon trap [111]. For muon energies in the range of electron binding energies the muon capture cross sections grow to atomic values, efficient atom production results at the rate of approximately 50 Hz. It should be noted that antiprotonic atoms could be produced similarly [112] and promise measurements of neutron distributions in nuclei.

E. Physics potential of a low-energy muon collider operating as a Higgs factory

Muon colliders [113,114] have a number of unique features that make them attractive candidates for future accelerators [9]. The most important and fundamental of these derive from the large mass of the muon in comparison to that of the electron. The synchrotron radiation loss in a circular accelerator goes as the inverse fourth power

of the mass and is 2×10^9 times less for a muon than for an electron. Direct s-channel coupling to the Higgs boson goes as the mass squared and is 40 000 greater for the muon than for the electron. This leads to (i) the possibility of extremely narrow beam energy spreads, especially at beam energies below 100 GeV; (ii) the possibility of accelerators with very high energy; (iii) the possibility of employing storage rings at high energy; (iv) the possibility of using decays of accelerated muons to provide a high-luminosity source of neutrinos as discussed in Sec. IIA4; (v) increased potential for probing physics in which couplings increase with mass (as does the SM $h_{\rm SM} f\bar{f}$ coupling).

The relatively large mass of the muon compared to the mass of the electron means that the coupling of Higgs bosons to $\mu^+\mu^-$ is very much larger than to e^+e^-, implying much larger s-channel Higgs production rates at a muon collider as compared to an electron collider. For Higgs bosons with a very small (MeV-scale) width, such as a light SM Higgs boson, production rates in the s channel are further enhanced by the muon collider's ability to achieve beam energy spreads comparable to the tiny Higgs width. In addition, there is little beam-strahlung, and the beam energy can be tuned to one part in a million through continuous spin rotation measurements [115]. Because of these important qualitative differences between the two types of machines, only muon colliders can be advocated as potential s-channel Higgs factories capable of determining the mass and decay width of a Higgs boson to very high precision [116,117]. High rates of Higgs production at e^+e^- colliders rely on substantial VV Higgs coupling for the Z + Higgs (Higgs-strahlung) or $WW \rightarrow$ Higgs (WW fusion) reactions. In contrast, a $\mu^+\mu^-$ collider can provide a factory for producing a Higgs boson with little or no VV coupling so long as it has SM-like (or enhanced) $\mu^+\mu^-$ couplings.

Of course, there is a trade-off between small beam energy spread, $\delta E/E = R$, and luminosity. Current estimates for yearly integrated luminosities (using $\mathcal{L} = 1 \times 10^{32}$ cm^{-2} s^{-1} as implying $L = 1$ fb^{-1}/yr) are $L_{\rm year} \gtrsim$ 0.1, 0.22, and 1 fb^{-1} at $\sqrt{s} \sim 100$ GeV for beam energy resolutions of $R = 0.003\%$, 0.01%, and 0.1%, respectively; $L_{\rm year} \sim 2$, 6, and 10 fb^{-1} at $\sqrt{s} \sim 200$, 350, and 400 GeV, respectively, for $R \sim 0.1\%$. Despite this, studies show that for small Higgs width the s-channel production rate (and statistical significance over background) is maximized by choosing R to be such that $\sigma_{\sqrt{s}} \lesssim \Gamma_h^{\rm tot}$. In particular, in the SM context for $m_{h_{\rm SM}} \sim 110$ GeV this corresponds to $R \sim 0.003\%$.

If the $m_h \sim 115$ GeV LEP signal is real, or if the interpretation of the precision electroweak data as an indication of a light Higgs boson (with substantial VV coupling) is valid, then both e^+e^- and $\mu^+\mu^-$ colliders will be valuable. In this scenario the Higgs boson would have been discovered at a previous higher energy collider (even possibly a muon collider running at high energy),

and then the Higgs factory would be built with a center-of-mass energy precisely tuned to the Higgs boson mass. The most likely scenario is that the Higgs boson is discovered at the Large Hadron Collider (LHC) via gluon fusion ($gg \rightarrow H$) or perhaps earlier at the Tevatron via associated production ($q\bar{q} \rightarrow WH, t\bar{t}H$), and its mass is determined to an accuracy of about 100 MeV. If a linear collider has also observed the Higgs via the Higgs-strahlung process ($e^+e^- \rightarrow ZH$), one might know the Higgs boson mass to be better than 50 MeV with an integrated luminosity of 500 fb^{-1}. The muon collider would be optimized to run at $\sqrt{s} \approx m_H$, and this center-of-mass energy would be varied over a narrow range so as to scan over the Higgs resonance (see Fig. 10).

1. Higgs production

The production of a Higgs boson (generically denoted h) in the s channel with interesting rates is a unique feature of a muon collider [116,117]. The resonance cross section is

$$\sigma_h(\sqrt{s}) = \frac{4\pi\Gamma(h \rightarrow \mu\bar{\mu})\Gamma(h \rightarrow X)}{(s - m_h^2)^2 + m_h^2(\Gamma_{\rm tot}^h)^2}. \qquad (32)$$

In practice, however, there is a Gaussian spread ($\sigma_{\sqrt{s}}$) to the center-of-mass energy and one must compute the effective s-channel Higgs cross section after convolution assuming some given central value of $\sqrt{s}$:

$$\bar{\sigma}_h(\sqrt{s}) = \frac{1}{\sqrt{2\pi}\sigma_{\sqrt{s}}} \int \sigma_h(\sqrt{\hat{s}}) \exp\left[\frac{-(\sqrt{\hat{s}} - \sqrt{s})^2}{2\sigma_{\sqrt{s}}^2}\right] d\sqrt{\hat{s}}, \qquad (33)$$

$$\simeq^{\sqrt{s}=m_h} \frac{4\pi}{m_h^2} \frac{{\rm BF}(h \rightarrow \mu\bar{\mu}){\rm BF}(h \rightarrow X)}{[1 + \frac{8}{\pi}(\frac{\sigma_{\sqrt{s}}}{\Gamma_h^{\rm tot}})^2]^{1/2}}. \qquad (34)$$

$m_{h_{\rm SM}} = 110$ GeV, ϵL=0.00125fb^{-1} per bin

FIG. 10. Number of events and statistical errors in the $b\bar{b}$ final state as a function of $\sqrt{s}$ in the vicinity of $m_{h_{\rm SM}} = 110$ GeV, assuming $R = 0.003\%$, and $\epsilon L = 0.001\,25$ fb^{-1} at each data point.

It is convenient to express $\sigma_{\sqrt{s}}$ in terms of the rms Gaussian spread of the energy of an individual beam, R:

$$\sigma_{\sqrt{s}} = (2 \text{ MeV})\left(\frac{R}{0.003\%}\right)\left(\frac{\sqrt{s}}{100 \text{ GeV}}\right). \qquad (35)$$

From Eq. (32), it is apparent that a resolution $\sigma_{\sqrt{s}} \lesssim \Gamma_h^{\text{tot}}$ is needed to be sensitive to the Higgs width. Further, Eq. (34) implies that $\overline{\sigma}_h \propto 1/\sigma_{\sqrt{s}}$ for $\sigma_{\sqrt{s}} > \Gamma_h^{\text{tot}}$ *and that* Γ_h^{tot} is not so large that $\text{BF}(h \to \mu\bar{\mu})$ is extremely suppressed. The width of a light SM-like Higgs is very small (e.g., a few MeV for $m_{h_{\text{SM}}} \sim 110$ GeV), implying the need for R values as small as $\sim 0.003\%$ for studying a light SM-like h. Figure 10 illustrates the result for the SM Higgs boson of an initial centering scan over $\sqrt{s}$ values in the vicinity of $m_{h_{\text{SM}}} = 110$ GeV. This figure dramatizes (i) that the beam energy spread must be very small because of the very small $\Gamma_{h_{\text{SM}}}^{\text{tot}}$ (when $m_{h_{\text{SM}}}$ is small enough that the $WW^\star$ decay mode is highly suppressed); (ii) that we require the very accurate *in situ* determination of the beam energy to one part in a million through the spin precession of the muon noted earlier in order to perform the scan and then center on $\sqrt{s} = m_{h_{\text{SM}}}$ with a high degree of stability. If the h has SM-like couplings to WW, its width will grow rapidly for $m_h > 2m_W$ and its s-channel production cross section will be severely suppressed by the resulting decrease of $\text{BF}(h \to \mu\mu)$. More generally, any h with SM-like or larger $h\mu\mu$ coupling will retain a large s-channel production rate when $m_h > 2m_W$ only if the hWW coupling becomes strongly suppressed relative to the $h_{\text{SM}}WW$ coupling.

The general theoretical prediction within supersymmetric models is that the lightest supersymmetric Higgs boson h^0 will be very similar to the h_{SM} when the other Higgs bosons are heavy. This "decoupling limit" is very likely to arise if the masses of the supersymmetric particles are large (since the Higgs masses and the superparticle masses are typically similar in size for most boundary condition choices). Thus, h^0 rates will be very similar to h_{SM} rates. In contrast, the heavier Higgs bosons in a typical supersymmetric model decouple from VV at large mass and remain reasonably narrow. As a result, their s-channel production rates remain large.

For a SM-like h, at $\sqrt{s} = m_h \approx 115$ GeV and $R = 0.003\%$, the $b\bar{b}$ rates are

$$\text{signal} \approx 10^4 \text{ events} \times L(\text{fb}^{-1}), \qquad (36)$$

$$\text{background} \approx 10^4 \text{ events} \times L(\text{fb}^{-1}). \qquad (37)$$

2. What the muon collider adds to LHC and linear collider (LC) data

An assessment of the need for a Higgs factory requires that one detail the unique capabilities of a muon collider versus the other possible future accelerators as well as comparing the abilities of all the machines to measure the same Higgs properties. Muon colliders, and a Higgs factory, in particular, would become operational only after the LHC physics program is well developed and, quite possibly, after a linear collider program is mature as well. So one important question is the following: If a SM-like Higgs boson and, possibly, important physics beyond the standard model have been discovered at the LHC and perhaps studied at a linear collider, what new information could a Higgs factory provide? The s-channel production process allows one to determine the mass, total width, and the cross sections $\overline{\sigma}_h(\mu^+\mu^- \to h \to X)$ for several final states X to very high precision. The Higgs mass, total width, and the cross sections can be used to constrain the parameters of the Higgs sector. For example, in the minimal supersymmetric standard model (MSSM) their precise values will constrain the Higgs sector parameters m_{A^0} and $\tan\beta$ (where $\tan\beta$ is the ratio of the two vacuum expectation values of the two Higgs doublets of the MSSM). The main question is whether these constraints will be a valuable addition to LHC and linear collider (LC) constraints. The expectations for the luminosity available at linear colliders has risen steadily. The most recent studies assume an integrated luminosity of some 500 fb^{-1} corresponding to 1–2 yr of running at a few $\times$ 100 fb^{-1} per year. This luminosity results in the production of greater than 10^4 Higgs bosons per year through the Bjorken Higgs-strahlung process, $e^+e^- \to Zh$, provided the Higgs boson is kinematically accessible. This is comparable or even better than can be achieved with the current machine parameters for a muon collider operating at the Higgs resonance; in fact, recent studies have described high-luminosity linear colliders as "Higgs factories," though for the purposes of this report, we will reserve this term for muon colliders operating at the s-channel Higgs resonance. A linear collider with such high luminosity can certainly perform quite accurate measurements of certain Higgs parameters, such as the Higgs mass, couplings to gauge bosons, and couplings to heavy quarks [118]. Precise measurements of the couplings of the Higgs boson to the standard model particles are an important test of the mass generation mechanism. In the standard model with one Higgs doublet, this coupling is proportional to the particle mass. In the more general case there can be mixing angles present in the couplings. Precision measurements of the couplings can distinguish the standard model Higgs boson from that of a more general model and can constrain the parameters of a more general Higgs sector.

The accuracies possible at different colliders for measuring m_h and Γ_h^{tot} of a SM-like h with $m_h \sim 110$ GeV are given in Table V. Once the mass is determined to about 1 MeV at the LHC and/or LC, the muon collider would employ a three-point fine scan [116] near the resonance peak. Since all the couplings of the standard model are

TABLE V. Achievable relative uncertainties for a SM-like $m_h = 110$ GeV for measuring the Higgs boson mass and total width for the LHC, LC (500 fb^{-1}), and the muon collider (0.2 fb^{-1}).

	LHC	LC	$\mu^+\mu^-$
m_h	9×10^{-4}	3×10^{-4}	$1–3 \times 10^{-6}$
Γ_h^{tot}	> 0.3	0.17	0.2

known, $\Gamma_{h_{\text{SM}}}^{\text{tot}}$ is known. Therefore a precise determination of Γ_h^{tot} is an important test of the standard model, and any deviation would be evidence for a nonstandard Higgs sector. For a SM Higgs boson with a mass sufficiently below the $WW^\star$ threshold, the Higgs total width is very small (of the order of several MeV), and the only process where it can be measured *directly* is in the s channel at a muon collider. Indirect determinations at the LC can have higher accuracy once m_h is large enough that the $WW^\star$ mode rates can be accurately measured, requiring $m_h > 120$ GeV. This is because at the LC the total width must be determined indirectly by measuring a partial width and a branching fraction, and then computing the total width,

$$\Gamma_{\text{tot}} = \frac{\Gamma(h \to X)}{BR(h \to X)}, \qquad (38)$$

for some final state X. For a Higgs boson so light that the $WW^\star$ decay mode is not useful, the total width measurement would probably require use of the $\gamma\gamma$ decays [119]. This would require information from a photon collider as well as the LC and a small error is not possible. The muon collider can measure the total width of the Higgs boson directly, a very valuable input for precision tests of the Higgs sector.

To summarize, if a Higgs is discovered at the LHC or possibly earlier at the Fermilab Tevatron, attention will turn to determining whether this Higgs has the properties expected of the standard model Higgs. If the Higgs is discovered at the LHC, it is quite possible that super-symmetric states will be discovered concurrently. The next goal for a linear collider or a muon collider will be to better measure the Higgs boson properties to determine if everything is consistent within a supersymmetric framework or consistent with the standard model. A Higgs factory of even modest luminosity can provide uniquely powerful constraints on the parameter space of the supersymmetric model via its very precise measurement of the light Higgs mass, the highly accurate determination of the total rate for $\mu^+\mu^- \to h^0 \to b\bar{b}$ (which has almost zero theoretical systematic uncertainty due to its insensitivity to the unknown m_b value), and the moderately accurate determination of the h^0's total width. In addition, by combining muon collider data with LC data, a completely model-independent and very precise deter-

mination of the $h^0\mu^+\mu^-$ coupling is possible. This will provide another strong discriminator between the SM and the MSSM. Further, the $h^0\mu^+\mu^-$ coupling can be compared to the muon collider and LC determinations of the $h^0\tau^+\tau^-$ coupling for a precision test of the expected universality of the fermion mass generation mechanism.

F. Physics potential of a high-energy muon collider

Once one learns to cool muons, it becomes possible to build muon colliders with energies of ≈ 3 TeV in the center of mass that fit on an existing laboratory site [9,120]. At intermediate energies, it becomes possible to measure the W mass [121,122] and the top quark mass [121,123] with high accuracy, by scanning the thresholds of these particles and making use of the excellent energy resolution of the beams. We further consider here the ability of a higher energy muon collider to scan higher-lying Higgs-like objects such as the H^0 and the A^0 in the MSSM that may be degenerate with each other.

Heavy Higgs bosons

As discussed in the previous section, precision measurements of the light Higgs boson properties might make it possible to not only distinguish a supersymmetric boson from a standard model one, but also pinpoint a range of allowed masses for the heavier Higgs bosons. This becomes more difficult in the decoupling limit where the differences between a supersymmetric and standard model Higgs are smaller. Nevertheless with sufficiently precise measurements of the Higgs branching fractions, it is possible that the heavy Higgs boson masses can be inferred. A muon collider light Higgs factory might be essential in this process. In the context of the MSSM, m_{A^0} can probably be restricted to within 50 GeV or better if $m_{A^0} < 500$ GeV. This includes the $250–500$ GeV range of heavy Higgs boson masses for which discovery is not possible via H^0A^0 pair production at a $\sqrt{s} = 500$ GeV LC. Further, the A^0 and H^0 cannot be detected in this mass range at either the LHC or LC in $b\bar{b}H^0$, $b\bar{b}A^0$ production for a wedge of moderate $\tan\beta$ values. (For large enough values of $\tan\beta$ the heavy Higgs bosons are expected to be observable in $b\bar{b}A^0$, $b\bar{b}H^0$ production at the LHC via their $\tau^+\tau^-$ decays and also at the LC.) A muon collider can fill some, perhaps all of this moderate $\tan\beta$ wedge. If $\tan\beta$ is large, the $\mu^+\mu^-H^0$ and $\mu^+\mu^-A^0$ couplings (proportional to $\tan\beta$ times a SM-like value) are enhanced, thereby leading to enhanced production rates in $\mu^+\mu^-$ collisions. The most efficient procedure is to operate the muon collider at maximum energy and produce the H^0 and A^0 (often as overlapping resonances) via the radiative return mechanism. By looking for a peak in the $b\bar{b}$ final state, the H^0 and A^0 can be discovered and, once discovered, the machine $\sqrt{s}$ can be set to m_{A^0} or m_{H^0}

FIG. 11.   Separation of A and H signals for $\tan\beta = 5$ and 10. From Ref. [116].

and factorylike precision studies pursued. Note that the A^0 and H^0 are typically broad enough that $R = 0.1\%$ would be adequate to maximize their s-channel production rates. In particular, $\Gamma \sim 30$ MeV if the $t\bar{t}$ decay channel is not open, and $\Gamma \sim 3$ GeV if it is. Since $R = 0.1\%$ is sufficient, much higher luminosity ($L \sim 2$–10 fb^{-1}/yr) would be possible as compared to that for $R = 0.01\%$–0.003% required for studying the h^0.

In short, for these moderate $\tan\beta - m_{A^0} \gtrsim 250$ GeV scenarios that are particularly difficult for both the LHC and the LC, the muon collider would be the only place that these extra Higgs bosons can be discovered and their properties measured very precisely.

In the MSSM, the heavy Higgs bosons are largely degenerate, especially in the decoupling limit where they are heavy. Large values of $\tan\beta$ heighten this degeneracy. A muon collider with sufficient energy resolution might be the only possible means for separating out these states. Examples showing the H and A resonances for $\tan\beta = 5$ and 10 are shown in Fig. 11. For the larger value of $\tan\beta$ the resonances are clearly overlapping. For the better energy resolution of $R = 0.01\%$, the two distinct resonance peaks are still visible, but become smeared out for $R = 0.06\%$.

Once muon colliders of these intermediate energies can be built, higher energies such as 3–4 TeV in the center of mass become feasible. Muon colliders with these energies will be complementary to hadron colliders of the Superconducting Supercollider (SSC) class and above. The background radiation from neutrinos from the muon decay becomes a problem at ≈ 3 TeV in the CoM [124]. Ideas for ameliorating this problem have been discussed and include optical stochastic cooling to reduce the number of muons needed for a given luminosity, elimination of straight sections via wigglers or undulators, or special sites for the collider such that the neutrinos break ground in uninhabited areas.

III. NEUTRINO FACTORY

In this Section we describe the various components of a neutrino factory, based on the most recent feasibility study (study II) [29] that was carried out jointly by BNL and the MC. We also describe the stages that could be constructed incrementally to provide a productive physics program that evolves eventually into a full-fledged neutrino factory. Details of the design described here are based on the specific scenario of sending a neutrino beam from Brookhaven to a detector in Carlsbad, New Mexico. More generally, however, the design exemplifies a neutrino factory for which our two feasibility studies demonstrated technical feasibility (provided the challenging component specifications are met), established a cost baseline, and established the expected range of physics performance. As noted earlier, this design typifies a neutrino factory that could fit comfortably on the site of an existing laboratory, such as BNL or FNAL.

A list of the main ingredients of a neutrino factory is given as follows:

Proton driver: Provides 1–4 MW of protons on target from an upgraded AGS; a new booster at Fermilab would perform equivalently.

Target and capture: A high-power target immersed in a 20 T superconducting solenoidal field to capture pions produced in proton-nucleus interactions.

Decay and phase rotation: Three induction linacs, with internal superconducting solenoidal focusing to contain the muons from pion decays, that provide nearly nondistorting phase rotation; a "minicooling" absorber section is included after the first induction linac to reduce the beam emittance and lower the beam energy to match the downstream cooling-channel acceptance.

Bunching and cooling: A solenoidal focusing channel, with high-gradient rf cavities and liquid-hydrogen absorbers, that bunches the 250 MeV/c muons into 201.25 MHz rf buckets and cools their transverse normalized rms emittance from 12 to 2.7 mm rad.

Acceleration: A superconducting linac with solenoidal focusing to raise the muon beam energy to 2.48 GeV, followed by a four-pass superconducting RLA to provide a 20 GeV muon beam; a second RLA could optionally be added to reach 50 GeV, if the physics requires this.

Storage ring: A compact racetrack-shaped superconducting storage ring in which $\approx 35\%$ of the stored muons decay toward a detector located about 3000 km from the ring.

A. Proton driver

The proton driver considered in study II is an upgrade of the BNL AGS and uses most of the existing components and facilities; parameters are listed in Table VI. To serve as the proton driver for a neutrino factory, the existing booster is replaced by a 1.2 GeV superconducting

TABLE VI. Proton driver parameters for BNL and FNAL designs.

	BNL	FNAL
Total beam power (MW)	1	1.2
Beam energy (GeV)	24	16
Average beam current (μA)	42	72
Cycle time (ms)	400	67
Number of protons per fill	1×10^{14}	3×10^{13}
Average circulating current (A)	6	2
No. of bunches per fill	6	18
No. of protons per bunch	1.7×10^{13}	1.7×10^{12}
Time between extracted bunches (ms)	20	0.13
Bunch length at extraction, rms (ns)	3	1

proton linac. The modified layout is shown in Fig. 12. The AGS repetition rate is increased from 0.5 to 2.5 Hz by adding power supplies to permit ramping the ring more quickly. No new technology is required for this—the existing supplies are replicated and the magnets are split into six sectors rather than the two used presently. The total proton charge (10^{14} ppp in six bunches) is only 40% higher than the current performance of the AGS. However, due to the required short bunches, there is a large increase in peak current and concomitant need for an improved vacuum chamber; this is included in the upgrade. The six bunches are extracted separately, spaced by 20 ms, so that the target, induction linacs, and rf systems that follow need only deal with single bunches at an instantaneous repetition rate of 50 Hz (average rate of 15 Hz). The average proton beam power is 1 MW. A possible future upgrade to 2×10^{14} ppp and 5 Hz could give an average beam power of 4 MW. At this higher intensity, a superconducting bunch compressor ring would be needed to maintain the rms bunch length at 3 ns.

If the facility were built at Fermilab, the proton driver would be a newly constructed 16 GeV rapid cycling booster synchrotron [125]. The planned facility layout is shown in Fig. 13. The initial beam power would be 1.2 MW, and a future upgrade to 4 MW would be possible.

The Fermilab design parameters are included in Table VI. A less ambitious and more cost effective 8 GeV proton driver option has also been considered for Fermilab [125]; this too might be the basis for a proton driver design.

B. Target and capture

A mercury jet target is chosen to give a high yield of pions per MW of incident proton power. The 1 cm diameter jet is continuous and is tilted with respect to the beam axis. The target layout is shown in Fig. 14. We assume that the thermal shock from the interacting proton bunch fully disperses the mercury, so the jet must have a velocity of 20–30 m/s to be replaced before the next bunch. Calculations of pion yields that reflect the detailed magnetic geometry of the target area have been performed with the MARS code [126]. To avoid mechanical fatigue problems, a mercury pool serves as the beam dump. This pool is part of the overall target—its mercury is circulated through the mercury jet nozzle after passing through a heat exchanger.

Pions emerging from the target are captured and focused down the decay channel by a solenoidal field that is 20 T at the target center, and tapers down, over 18 m, to a periodic (0.5 m) superconducting solenoid channel ($B_z = 1.25$ T) that continues through the phase rotation to the start of bunching. Note that the longitudinal directions of the fields in this channel do not change sign from cell to cell as they do in the cooling channel. The 20 T solenoid, with a resistive magnet insert and superconducting outer coil, is similar in character to the higher field (up to 45 T), but smaller bore, magnets existing at several laboratories [128]. The magnet insert is made with a hollow copper conductor having ceramic insulation to withstand radiation. MARS [126] simulations of radiation levels show that, with the shielding provided, both the copper and superconducting magnets could have a lifetime greater than 15 yr at 1 MW.

In study I, the target was a solid carbon rod. At high beam energies, this implementation has a lower pion yield than the mercury jet, and is expected to be more limited


FIG. 12. (Color) AGS proton driver layout.

FIG. 13. (Color) FNAL proton driver layout from Ref. [125].

in its ability to handle the proton beam power, but should simplify the target handling issues that must be dealt with. At lower beam energies, say, 6 GeV, the yield difference between C and Hg essentially disappears, so a carbon target would be a competitive option with a lower energy driver. Present indications [129] are that a carbon-carbon composite target can be tailored to tolerate even a 4 MW proton beam power—a very encouraging result.

FIG. 14. (Color) Target, capture solenoids, and mercury containment.

Other alternative approaches, including a rotating Inconel band target and a granular Ta target are also under consideration, as discussed in study II [29]. Clearly there are several target options that could be used for the initial facility.

C. Phase rotation

The function of the phase rotation section in a neutrino factory is to reduce the energy spread of the collected muon beam to a manageable level, allowing reasonable throughput in the subsequent system components. The following description refers specifically to the properties of the U.S. feasibility study II for a neutrino factory. The initial pions are produced in the mercury target with a very wide range of momenta. The momentum spectrum peaks around 250 MeV/c, but there is a tail of high-energy pions that extends well beyond 1 GeV. The pions are spread in time over about 3 ns, given by the pulse duration of the proton driver. After the 18 m long tapered collection solenoid and an 18 m long drift section, where the beam is focused by 1.25 T solenoids, most of the low-energy pions have decayed into muons. At this point the muon energy spectrum also extends over an approximately 1 GeV range and the time spectrum extends over approximately 50 ns. However, there is a strong correlation between the muon energy and time that can be used for phase rotation.

TABLE VII. Properties of the induction linacs used in feasibility study II.

Induction Linac		1	2	3
Length	m	100	80	100
Peak gradient	MV/m	1.5	−1.5	1.0
Pulse FWHM	ns	250	100	380
Pulse start offset	ns	55	0	55

In the phase rotation process an electric field is applied at appropriate times to decelerate the leading high-energy muons and to accelerate the trailing low-energy ones. Since the bunch train required by a neutrino factory can be very long, it is possible to minimize the energy spread using induction linacs. The induction linac consists of a simple nonresonant structure, where the drive voltage is applied to an axially symmetric gap that encloses a toroidal ferromagnetic material. The change in flux in the magnetic core induces an axial electric field that provides particle acceleration. The induction linac is typically a low gradient structure that can provide acceleration fields of varying shapes and time durations from tens of nanoseconds to several microseconds. Some properties of the induction linacs are given in Table VII.

Three induction linacs are used in a system that reduces distortion in the phase-rotated bunch and permits all induction units to operate with unipolar pulses. The induction units are similar to those being built for the DARHT project [130]. The 1.25 T beam transport sole-noids are placed inside the induction cores in order to avoid saturating the core material, as shown in Fig. 15.

Between the first and second induction linacs two liquid-hydrogen absorbers (each 1.7 m long and 30 cm radius) are used to (1) provide some initial cooling of the transverse emittance of the muon beam and (2) lower the average momentum of the beam to better match the downstream cooling-channel acceptance. This process is referred to as minicooling. The direction of the solenoid magnetic field is reversed between the two absorbers. The presence of material in the beam path destroys the conservation of canonical angular momentum that occurs when a particle enters and leaves a solenoid in vacuum. The buildup of this angular momentum would eventually lead to emittance growth. However, this growth can be minimized by periodically reversing the direction of the field.

The beam at the end of the phase rotation section has an average momentum of about 250 MeV/c and an rms fractional energy spread of $\approx 4.4\%$. Figure 16 shows the evolution of the beam distribution in the phase rotation section.

D. Buncher

The long beam pulse (400 ns) after the phase rotation is then bunched at 201.25 MHz prior to cooling and acceleration at that frequency. The bunching is done in a lattice identical to that at the start of the cooling channel and is preceded by a matching section from the 1.25 T solenoids


FIG. 15. (Color) Cross section of the induction cell and transport solenoids.

FIG. 16.   Evolution of the beam distribution in the phase rotation section. The graphs show the distribution before the phase rotation, after the first induction linac (top row, left to right), after minicooling, and after the second and third induction linacs (bottom row).

into this lattice. The bunching has three stages, each consisting of rf (with increasing acceleration) followed by drifts (with decreasing length). In the first two rf sections, second-harmonic 402.5 MHz rf is used together with the 201.25 MHz primary frequency to improve the capture efficiency. The 402.5 MHz cavities are designed to fit into the bore of the focusing solenoids, in the location corresponding to that of the liquid-hydrogen absorber in the downstream cooling channel. Their aperture radius for the 402.5 MHz cavities is 20 cm at the IRIS, while that of the 201.25 MHz cavities is 25 cm. The gradients on axis in the cavities are 6.4 MV/m for the 402.5 MHz cavities, and range from 6 to 8 MV/m for the 201.25 MHz cavities. The resulting bunches fill the 201.25 MHz stationary rf bucket. Figure 17 shows the evolution of the longitudinal distribution in the buncher.

E. Cooling

The transverse emittance of the muon beam after phase rotation and bunching must be reduced in order to fit into the downstream accelerators and storage ring. Ionization cooling is currently the only feasible option for cooling the beam within the muon lifetime. In ionization cooling the transverse and longitudinal momenta are lowered in the absorbers, but only the longitudinal momentum is restored by the rf. The following description refers specifically to the properties of the U.S. feasibility study II for a neutrino factory. Transverse emittance cooling is achieved using cooling cells that (1) lower the beam energy by 7–12 MeV in liquid-hydrogen absorbers, (2) use 201 MHz rf cavities to restore the lost energy, and (3)

use 3–5 T solenoids to strongly focus the beam at the absorbers. At the end of the cooling channel the rms normalized transverse emittance is reduced to about 2.5 mm rad.

Each cell of the lattice contains three solenoids. The direction of the solenoidal field reverses in alternate cells in order to prevent the buildup of canonical angular momentum, as mentioned above in the discussion of minicooling. In analogy with the FODO lattice this focusing arrangement is referred to as a (S)FOFO (super-FOFO) lattice. Multiple Coulomb scattering together with the focusing strength determines the asymptotic limit on the transverse emittance that the cooling channel can reach. The focusing strength in the channel is tapered so that the angular spread of the beam at the absorber locations remains large compared to the characteristic spread from scattering. This is achieved by keeping the focusing strength inversely proportional to the emittance, i.e., increasing it as the emittance is reduced. The solenoidal field profile was chosen to maximize the momentum acceptance ($\pm 22\%$) through the channel. To maintain the tapering of the focusing it was eventually necessary to reduce the cell length from 2.75 m in the initial portion of the channel to 1.65 m in the final portion. A layout of the shorter cooling cells is shown in Fig. 18.

Figure 19 shows a simulation of cooling in this channel. The transverse emittance decreases steadily along the length of the channel. This type of channel only cools transversely, so the longitudinal emittance increases until the rf bucket is full and then remains fairly constant as particles are lost from the bucket. A useful figure of merit

FIG. 17. (Color) Evolution of beam in buncher. Plots are at the beginning of the buncher (top left), and at the ends of the three bunching stages (top right, bottom left, and bottom right, in that order).

for cooling at a neutrino factory is the increase in the number of muons that fit within the acceptance of the downstream accelerators. This is shown in Fig. 20. At each axial position the number of muons is shown that fall within two acceptances appropriate to a downstream accelerator. Both acceptances require the muon longitudinal phase space be less than 150 mm. The density of particles within a normalized transverse acceptance, for example, steadily increases by a factor of about 3 over the channel length, clearly showing the results of cooling. The saturation of the yield determined the chosen channel length of 108 m.

F. Acceleration

The layout of the acceleration system is shown in Fig. 21, and its parameters are listed in Table VIII. The acceleration system consists of a preaccelerator linac followed by a four-pass recirculating linac. The recirculating linac allows a reduction in the amount of rf required for acceleration by passing the beam through linacs multiple times. The linacs are connected by arcs, and a separate arc is used for each pass. At low energies, however, the large emittance of the beam would require a much shorter cell length and larger aperture than is desirable and needed at higher energies. This, combined with

FIG. 18. (Color) Two cells of the 1.65 m cooling lattice.

FIG. 19. The transverse (filled circles, in mm) and longitudinal (open circles, in cm) emittances, as a function of the distance down the cooling channel.

FIG. 20. Muons per incident proton in the cooling channel that would fall within a normalized transverse acceptance of 15 mm (open circles) or 9.75 mm (filled circles).

difficulties in injecting the large emittance and energy spread beam into the recirculating accelerator, and the loss of efficiency due to the phase slip at low energies lead to the necessity for a linac that precedes the recirculating linac.

A 20 m SFOFO matching section, using normal conducting rf systems, matches the beam optics to the requirements of a 2.87 GeV superconducting rf linac with solenoidal focusing. The linac is in three parts. The first part has a single 2-cell rf cavity unit per period. The second part, as a longer period becomes possible, has two 2-cell cavity units per period. The last section, with still longer period, accommodates four 2-cell rf cavity units per period. See Tables IX and X for details of the rf cryostructures and cavities. Figure 22 shows the three cryomodule types that make up the linac.

This linac is followed by a single four-pass RLA that raises the energy from 2.5 to 20 GeV. The RLA uses the same layout of four 2-cell superconducting rf cavity structures as the long cryomodules in the linac, but utilizes quadrupole triplet focusing, as indicated in Fig. 23.

The arcs have an average radius of 62 m and are all in the same horizontal plane. They also utilize triplet focusing. There are around 120 arc cells, with 2.15 m dipoles, and triplet quadrupoles which are very similar to those in the linacs. The required full quadrupole apertures vary from 20 to 12 cm, and the dipole gaps vary from 14 to 9 cm. All magnet pole tip fields are under 2 T, except in the switchyard where they are as high as 2.3 T in some cases (and the magnet apertures rise to 21 cm).

The 4.5 K-equivalent cryogenic load for the entire acceleration system is 27.9 kW. In study I, where the final beam energy was chosen to be 50 GeV, a second RLA is needed. This second RLA is similar to the RLA just described, but considerably larger.

G. Storage ring

After acceleration in the RLA, the muons are injected into the upward-going straight section of a racetrack-shaped storage ring with a circumference of 358 m. Parameters of the ring are summarized in Table XI. High-field superconducting arc magnets are used to minimize the arc length. Minimizing the arc length for a given length of straight maximizes the fraction of the circumference contained in the straight section, thereby maximizing the fraction of neutrinos (around 35% in our case) decaying toward the detector.


FIG. 21. (Color) Accelerating system layout.

TABLE VIII.   Main parameters of the muon accelerator.

Injection momentum (MeV/c)/kinetic energy (MeV)	210/129.4
Final energy (GeV)	20
Initial normalized transverse acceptance (mm rad)	15
rms normalized transverse emittance (mm rad)	2.4
Initial longitudinal acceptance, $\Delta p L_b / m_\mu$ (mm)	170
momentum spread, $\Delta p/p$	± 0.21
bunch length, L_b (mm)	± 407
rms energy spread	0.084
rms bunch length (mm)	163
Number of bunches per pulse	67
Number of particles per bunch/per pulse	$4.4 \times 10^{10}/3 \times 10^{12}$
Bunch frequency/accelerating frequency (MHz)	201.25/201.25
Average beam power (kW)	150

TABLE IX.   Parameters for three types of linac cryomodules.

Cavities per period	1	2	4
Period length (m)	5	8	13
Number of periods	11	16	19
Cavity type	A	A	B
Solenoid full aperture (cm)	46	46	36
Solenoid length (m)	1	1	1.5
Maximum solenoid field (T)	2.1	2.1	4.2

TABLE X.   Parameters for superconducting cavities.

	A	B
Frequency (MHz)	201.25	201.25
Cells per cavity	2	2
Aperture diameter (cm)	46	30
Energy gain (MV)	22.5	25.5
rf pulse length (ms)	3	3
Input power (kW)	980	1016
Peak surface field (MV/m)	23.1	24.3
Q_0	6×10^9	6×10^9

Furthermore, the beta functions in the downward-going straight (which is pointed toward the detector) are made large to reduce the angular divergence of the beam. This ensures that the angular divergence of the beam is dominated by the calculable relativistic angular divergence of the decay neutrinos. The goal of this is not only to make the angular divergence of the neutrino beam as small as possible and therefore maximize the flux, but it also reduces the experimental uncertainty associated with an uncertainty in the flux that would come from an uncertainty in the angular divergence of the muon beam.

All muons are allowed to decay; the maximum heat load from their decay electrons is 42 kW (126 W/m). This load is too high to be dissipated in the superconducting coils. For study II, a magnet design (see Fig. 24) has

been chosen that allows the majority of these electrons to exit between separate upper and lower cryostats, and be dissipated in a dump at room temperature. To maintain the vertical cryostat separation in focusing elements, skew quadrupoles are employed in place of standard quadrupoles. The result is a skew FODO lattice, giving diagonal


FIG. 22. (Color) Layouts of short (left), intermediate (middle), and long (right) cryomodules. Blue lines are the SC walls of the cavities. Solenoid coils are indicated in red.

FIG. 23. (Color) Layout of an RLA linac period.

oscillations, as opposed to the horizontal and vertical oscillations of the usual upright FODO lattice. In order to maximize the average bending field, Nb_3Sn pancake coils are employed. One coil of the bending magnet is extended and used as one-half of the previous (or following) skew quadrupole to minimize unused space. Figure 25 shows a layout of the ring as it would be located at BNL. (The existing 110 m high BNL smokestack is shown for scale.) For site-specific reasons, the ring is kept above the local water table and is placed on a roughly

TABLE XI.   Muon storage ring parameters.

Energy (GeV)	20
Circumference (m)	358.18
Normalized transverse acceptance (mm rad)	30
Energy acceptance (%)	2.2
Momentum compaction	0.028
Arc	
Length (m)	53.09
Number of cells per arc	10
Cell length (m)	5.3
Phase advance (deg)	60
Dipole length (m)	1.89
Dipole field (T)	6.93
Skew quadrupole length (m)	0.76
Skew quadrupole gradient (T/m)	35
β_{max} (m)	8.6
Production straight	
Length (m)	126
β_{max} (m)	200

30 m high berm. This requirement places a premium on a compact storage ring.

The beam is allowed to debunch in the storage ring. In one muon lifetime (0.42 ms), a bunch with the full energy spread ($\pm 2.2\%$) will have its total length increase by 0.51 μs (the storage ring is 1.19 μs long, and the bunch train starts out 0.33 μs long). If one wishes to avoid the increase in the bunch train length, one could install rf cavities, but the voltage required to avoid energy spread increase for the momentum compaction in this ring is prohibitive: a better solution would be a ring redesigned to have very low momentum compaction.

For study I, a conventional superconducting ring was utilized to store the 50 GeV muon beam. The heat load from muon decay products in this scenario is managed by having a liner inside the magnet bore to absorb the decay products. This approach is likewise available for BNL, provided some care is taken to keep the ring compact; acceptable lattice solutions have been found for this option as well.

H. Overall machine summary

Figure 26 shows the muons per incident proton in the "front end" (everything before the acceleration) of the study-II neutrino factory. Table XII gives the values at


FIG. 24. (Color) Three-dimensional view of the storage ring magnets.

FIG. 25. (Color) Top view and cross section through 20 GeV ring and berm. The existing 110 m tower, drawn to scale, gives a sense of the height of the ring on the BNL landscape.

the ends of several sections and the losses in those sections. These significant losses are a necessary cost of making a low-emittance beam that can be accelerated and injected into a storage ring.

An overall layout of the neutrino factory on the BNL site is shown in Fig. 27. Figure 28 shows the equivalent picture for a facility on the Fermilab site. In this latter case, the layout includes the additional RLA and longer storage ring needed to reach 50 GeV. Clearly the footprint of a neutrino factory is reasonably small, and such a machine would fit easily on the site of either BNL or Fermilab.

FIG. 26.   Muons per incident proton in the study-II neutrino factory front end.

TABLE XII. Muon survival and losses in several sections of the study-II neutrino factory. μ/p is the number of muons per proton at the end of that section, and "Loss" is the loss in that section.

Section	μ/p	Loss
Phase rotation	0.60	...
Buncher	0.47	22%
Cooling	0.22	53%
Accelerator aperture	0.16	26%
Preaccelerator linac	0.15	10%
Recirculating linac	0.13	10%

I. Detector

The neutrino factory, plus its long-baseline detector, will have a physics program that is a logical continuation of current and near-future neutrino oscillation experiments in the U.S., Japan, and Europe. Moreover, detector facilities located in experimental areas near the neutrino source will have access to integrated neutrino intensities 10^4–10^5 times larger than previously available (10^{20} neutrinos per year compared with 10^{15}–10^{16}).

The detector site taken for study II is the Waste Isolation Pilot Plant (WIPP) in Carlsbad, New Mexico. The WIPP site is approximately 2900 km from BNL. Space is potentially available for a large underground physics facility at depths of 740–1100 m, and discussions are under way between the DOE and the UNO project [30] on the possible development of such a facility.

1. Far detector

Specifications for the long-baseline neutrino factory detector are rather typical for an accelerator-based neutrino experiment. However, the need to maintain a high neutrino rate at these long distances requires detectors 3–10 times more massive than those in current neutrino experiments. Clearly, the rate of detected neutrinos depends on two factors—the source intensity and the detector size. In the final design of a neutrino factory, these two factors would be optimized together.

Two options are considered for the WIPP site: a 50 kt steel-scintillator PDT detector or a water-Cherenkov detector. The PDT detector would resemble MINOS. Figure 29 shows a 50 kt detector with dimensions 8 m × 8 m × 150 m. A detector of this size would record up to 4×10^4 ν_μ events per year.

A large water-Cherenkov detector would be similar to SuperKamiokande, but with either a magnetized water volume or toroids separating smaller water tanks. The detector could be the UNO detector [30], currently proposed to study both proton decay and cosmic neutrinos. UNO would be a 650 kt water-Cherenkov detector segmented into a minimum of three tanks (see Fig. 30). It would have an active fiducial mass of 440 kt and would record up to 3×10^5 ν_μ events per year from the neutrino factory beam.

Another possibility for a neutrino factory detector is a massive liquid-argon magnetized detector [31] that would also attempt to detect proton decay, as well as solar and supernova neutrinos.


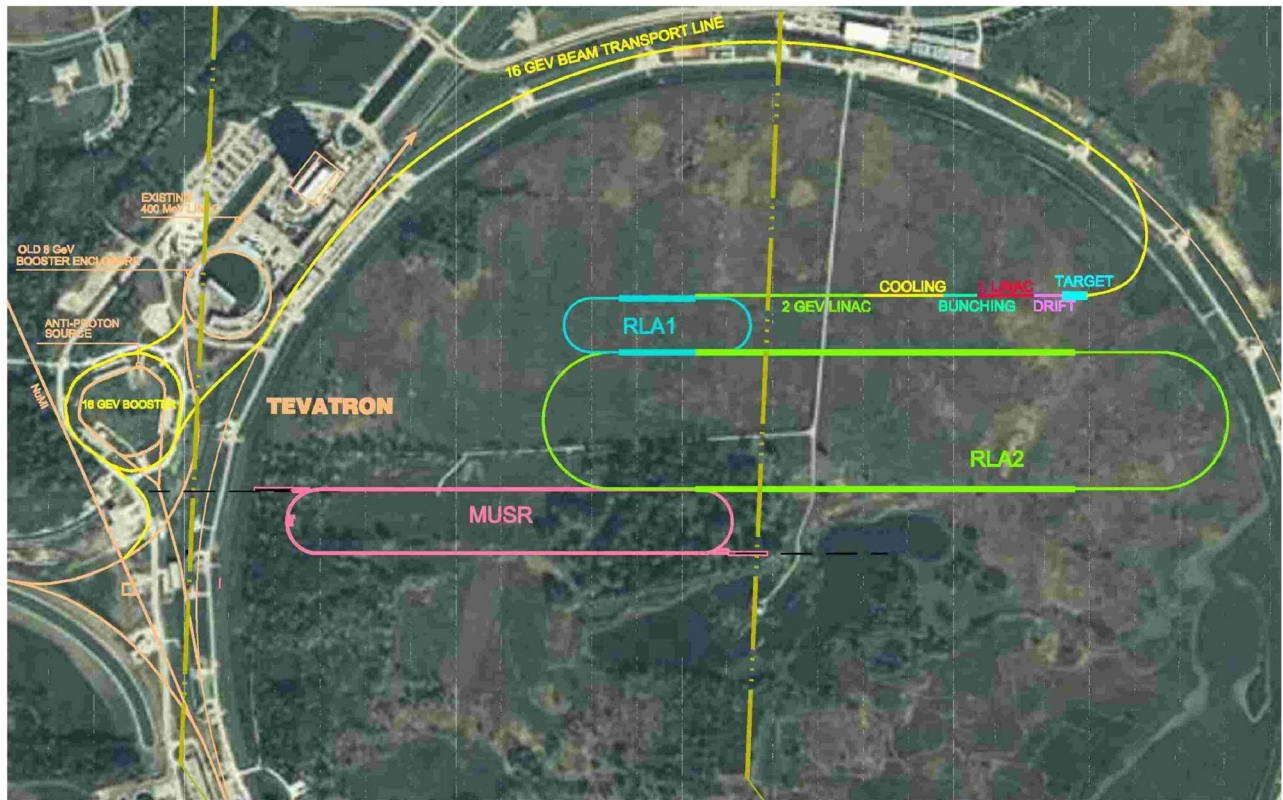FIG. 27. (Color) Schematic of a 20 GeV neutrino factory at BNL.

FIG. 28. (Color) Schematic of a 50 GeV neutrino factory at Fermilab.

2. Near detector

As noted, detector facilities located on site at the neutrino factory would have access to unprecedented intensities of pure neutrino beams. This would enable standard neutrino physics studies, such as $\sin^2\theta_W$, structure functions, neutrino cross sections, nuclear shadowing, and pQCD to be performed with much higher

precision than previously obtainable. In addition to its primary physics program, the near detector can also provide a precise flux calibration for the far detector,

FIG. 29. (Color) A possible 50 kt steel-scintillator-PDT detector at WIPP.

FIG. 30. (Color) Block schematic of the UNO detector, including initial design parameters.

though this may not be critical given the ability to monitor the storage ring beam intensity independently.

A compact liquid-argon TPC (similar to the ICARUS detector [32]), cylindrically shaped with a radius of 0.5 m and a length of 1 m, would have an active mass of 10^3 kg and a neutrino event rate $\mathcal{O}(10$ Hz). The TPC could be combined with a downstream magnetic spectrometer for muon and hadron momentum measurements. At these neutrino intensities, it is even possible to have an experiment with a relatively thin Pb target (1 L_{rad}), followed by a standard fixed-target spectrometer containing tracking chambers, time of flight, and calorimetry, with an event rate of $\mathcal{O}(1$ Hz).

J. Staging Options

It seems quite possible—perhaps even likely—that the neutrino factory would be built in stages, both for programmatic and for cost reasons. Here we outline a possible staging concept that provides good physics opportunities at each stage. The staging scenario we consider is not unique, nor is it necessarily optimized. Depending on the results of our technical studies and the results of continued searches for the Higgs boson, it is hoped that the neutrino factory is really the penultimate stage, to be followed later by a muon collider (Higgs factory). We assume this possibility in the staging discussion that follows. Because the physics program would be different at different stages, it is impractical at this time to consider detector details.

1. Stage 1

In the first stage, we envision a proton driver and a target facility to create superbeams. The driver could begin with a 1 MW beam level (stage 1) or could be designed from the outset to reach 4 MW (stage 1a). (Since the cost differential between 1 and 4 MW is not expected to be large, we do not consider any intermediate options here.) It is assumed, as was the case for both studies I and II, that the target facility is built from the outset to accommodate a 4 MW beam. Based on the study-II results, a 1 MW beam would provide about 1.2×10^{14} μ/s (1.2×10^{21} μ/yr) and a 4 MW beam about 5×10^{14} μ/s (5×10^{21} μ/yr) into a solenoid channel.

In addition to the neutrino program, this stage will also benefit π, K, and $\overline{p}$ programs, as discussed in [131,132].

2. Stage 2

In stage 2, we envision a muon beam that has been phase rotated (to give a reasonably low momentum spread) and transversely cooled. In the nomenclature of study II, this stage takes us to the end of the cooling channel. Thus, we have access to a muon beam with a central momentum of about 200 MeV/c, a transverse (normalized) emittance of 2.7 mm rad, and an rms energy

spread of about 4.5%. The intensity of the beam would be about 4×10^{13} μ/s (4×10^{20} μ/yr) at 1 MW, or 1.7×10^{14} μ/s (1.7×10^{21} μ/yr) at 4 MW. If more intensity were needed, and if less cooling could be tolerated, the length of the cooling channel could be reduced. As an example, stopping at the end of lattice 1 instead of the end of lattice 2 in the study-II cooling channel would result in an increase of transverse emittance by roughly a factor of 2. This is an appropriate stage to mount an experiment to search for a nonzero muon electric dipole moment.

3. Stage 3

In stage 3, we envision using the preacceleration linac to raise the beam energy to roughly 2.5 GeV. At this juncture, it may be appropriate to consider a small storage ring, comparable to the $g - 2$ ring at BNL, to be used for the next round of muon $g - 2$ experiments.

4. Stage 4

At stage 4, we envision having a complete neutrino factory operating with a 20 GeV storage ring. This stage includes the RLA and the storage ring. If it were necessary to provide a 50 GeV muon beam as stage 4a, an additional RLA and a larger storage ring would be needed.

5. Stage 5

In stage 5, we could envision an entry-level muon collider operating as a Higgs factory. Because the initial muon beam must be prepared as a single bunch of each charge, an additional ring for the proton driver to coalesce proton bunches into a single pulse is anticipated. The cooling will have to be significantly augmented. First, a much lower transverse emittance is needed, and second, it will be necessary to provide longitudinal cooling (emittance exchange) to maintain a reasonable transmission of the muons. The additional cooling will permit going to smaller solenoids and higher frequency rf systems (402.5 or perhaps 805 MHz), which should provide more cost-effective cooling. Next, we will need considerably more acceleration, though with smaller energy acceptance and aperture requirements than at present. Lastly, we will need a very low β^* lattice for the collider ring, along with mitigation of the potentially copious background levels near the interaction point. In this case the detector is, in effect, part of the collider ring, and its design must be an integral part of the overall ring design.

IV. MUON COLLIDERS

The primary advantage of using muons in a lepton collider arises from the fact that the muon is ≈ 200 times heavier than the electron. It is thus possible

FIG. 31. (Color) Comparative sizes of various proposed high-energy colliders compared with the FNAL and BNL sites. The energies in parentheses give for lepton colliders their CoM energies and for hadron colliders the approximate range of CoM energies attainable for hard parton-parton collisions.

to accelerate muons using circular accelerators that are compact and fit on existing accelerator sites. See Figure 31 for a comparison of relative sizes of muon colliders ranging from 500 GeV to 3 TeV center of mass energies with respect to the LHC, SSC, and NLC (Next Linear Collider). Once the problem of cooling a muon beam to sufficiently small emittances is solved and the beams can be accelerated, higher energies are much more easily obtained in a muon collider than in the linear electron-positron collider. Because the muon is unstable, it is necessary to cool and accelerate the beam before a substantial number has decayed. The number of turns in a muon lifetime is independent of the muon momentum for a given magnetic field, since both the circumference and the muon lifetime in the laboratory frame scale with muon momentum. With typical bending magnetic fields (≈ 5 T) available with today's technology, the muons last ≈ 1000 turns before half of them have decayed in the collider ring.

Muon decay also gives rise to large numbers of electrons that can affect the cryogenics of the magnets in the machine as well as pose serious background problems for detectors in the collision region. The 1999 status report [9] contains an excellent summary of the problems (and possible solutions) one faces on the way to a muon collider.

Figure 32 shows a schematic of such a muon collider, along with a depiction of the possible physics that can be addressed with each stage of the facility. Some of the ideas needed to obtain longitudinal cooling necessary for the muon collider are discussed in Sec. IV B and some of the parameters of the accelerator system for higher energy colliders are discussed in Sec. IV C below.


FIG. 32. (Color) Schematic of a muon collider.

A. Higgs factory requirements

The emittance of the muon beam needs to be reduced by a factor of $\approx 10^6$ from production [9] to the point of collision for there to be significant luminosity for experiments. Table XIII lists the transverse and longitudinal emittances at the end of the decay channel, study-II [29] cooling channel, and those needed for a 0.1 TeV center of mass energy muon collider, also known as a Higgs factory. It can be seen that one needs to cool by a factor of ≈ 20 in the transverse dimension and ≈ 3 in the longitudinal dimension from the study-II emittances to achieve the emittances necessary for a Higgs factory. This can be achieved by ionization cooling similar to the scheme described in Sec. III. The transverse emittance is reduced during ionization cooling, since only the longitudinal energy loss is replaced by rf acceleration. However, due to straggling, the longitudinal energy spread of the beam increases, even if the average longitudinal energy of the beam is kept constant. The longitudinal emittance thus grows in a linear cooling channel. In order to cool longitudinally, one needs to create dispersion in the system and have wedge absorbers at the point of maximum dispersion so that the faster particles go through the thicker parts of the wedge. This results in a reduction in longitudinal emittance accompanied by an

TABLE XIII. Transverse and longitudinal emittances at the end of the decay channel, study-II cooling channel, and the Higgs factory cooling channel.

Emittance at end of	Transverse emittance (π mm)	Longitudinal emittance (π mm)
Decay channel	17	150
Study-II cooler	2.6	30
Higgs factory cooler	0.14	9

FIG. 33. (Color) Plan of a 0.1-TeV-CoM muon collider, also known as a Higgs factory.

increase in transverse emittance and is thus called emittance exchange.

The status report [9] outlines the details of the acceleration and collider ring for the 0.1 TeV Higgs factory, shown schematically in Fig. 33. Table XIV gives a summary of the parameters of various muon colliders including three different modes of running the Higgs collider that have varying beam momentum spreads. Additional information about the muon collider can be found in [133,134].

B. Longitudinal cooling

At the time of writing of the status report [9] there was no satisfactory solution for the emittance exchange problem and this remained a major stumbling block towards realizing a muon collider. However, ring coolers have been found to hold significant promise in cooling in 6D phase space. Another advantage of ring coolers is that one can circulate the muons many turns, thereby reusing the cooling-channel elements. Several meetings on emittance exchange were held [135] and a successful workshop [136] was held in 2001, where we explored in some depth several kinds of ring coolers. These options differ primarily in the type of focusing used to contain the beam. We describe the current status of our understanding of three types of ring coolers here.

TABLE XIV. Baseline parameters for high- and low-energy muon colliders. Higgs/year assumes a cross section $\sigma = 5 \times 10^4$ fb; a Higgs width $\Gamma = 2.7$ MeV; 1 yr $= 10^7$ s.

|				0.1	
CoM energy (TeV)	3	0.4		0.1	
p energy (GeV)	16	16		16	
p's/bunch	2.5×10^{13}	2.5×10^{13}		5×10^{13}	
Bunches/fill	4	4		2	
Repetition rate (Hz)	15	15		15	
p power (MW)	4	4		4	
μ/bunch	2×10^{12}	2×10^{12}		4×10^{12}	
μ power (MW)	28	4		1	
Wall power (MW)	204	120		81	
Collider circumference (m)	6000	1000		350	
Average bending field (T)	5.2	4.7		3	
rms $\Delta p/p\%$	0.16	0.14	0.12	0.01	0.003
6D $\epsilon_{6,N}$ $(\pi \text{m})^3$	1.7×10^{-10}	1.7×10^{-10}	1.7×10^{-10}	1.7×10^{-10}	1.7×10^{-10}
rms ϵ_n (π mm mrad)	50	50	85	195	290
β^* (cm)	0.3	2.6	4.1	9.4	14.1
σ_z (cm)	0.3	2.6	4.1	9.4	14.1
σ_r spot (μm)	3.2	26	86	196	294
σ_θ IP (mrad)	1.1	1.0	2.1	2.1	2.1
Tune shift	0.044	0.044	0.051	0.022	0.015
n_{turns} (effective)	785	700	450	450	450
Luminosity cm^{-2} s^{-1}	7×10^{34}	10^{33}	1.2×10^{32}	2.2×10^{31}	10^{31}
Higgs/year			1.9×10^3	4×10^3	3.9×10^3

FIG. 34. (Color) Layout and parameters of the solenoid based ring cooler. Circumference: 36.963 m; nominal energy at short straight section: 250 MeV; bending field: 1.453 T; normal field gradient: 0.5; maximum solenoid field: 5.155 T; rf frequency: 205.69 MHz; accelerating gradient: 15 MeV/m; main absorber length: 128 cm; LiH wedge absorber: 14 cm; gradient of energy loss: 0.75 MeV/cm.

1. Solenoidal ring coolers

The basic design of the solenoidal ring cooler [137] is presented in Figure 34. Eight focusing dipole magnets with an index $n = -\frac{1}{2}$ are used for bending and focusing of the beam. Each of these dipoles bends the beam through 45° with a central orbit bending radius of 52 cm. We have done calculations to show that such dipoles are buildable. Figure 35 shows a configuration of such a dipole and the resulting magnetic field components calculated using a 3D field calculation program. There are four long solenoids containing rf cavities and liquid-hydrogen absorbers for transverse cooling. A magnetic field of 2.06 T at the end regions of the solenoids provides the same transverse focusing as the bending magnets. The magnetic field adiabatically increases to 5.15 T towards the center of the solenoid in order to produce a small β function (25–30 cm) at the absorbers. The short solenoids are designed to create an appropriate dispersion function that is zero at the long solenoids, which house the 200 MHz rf cavities. Their field is ±2.06 T at the edges and ±2.75 T centrally. A symmetric field flip is required in the short solenoids to prevent the buildup of canonical angular momentum. This field flip causes the dispersion in the long solenoids housing the rf cavities to be zero while permitting a nonzero dispersion at the lithium hydride wedge absorbers at the centers of the short solenoids which then produce longitudinal cooling via emittance exchange.
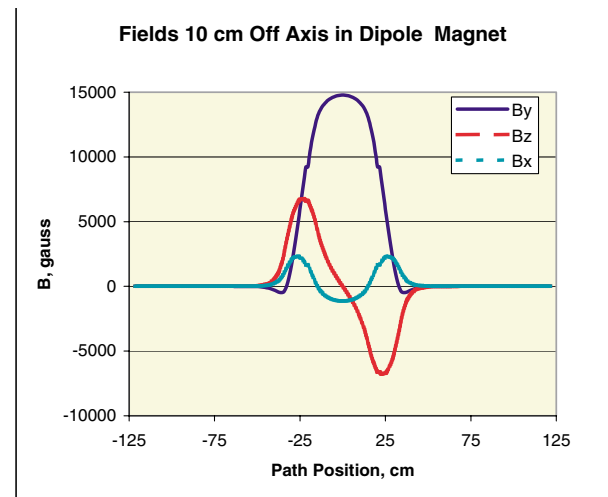

FIG. 35. (Color) The figure shows a computer model of a 52 cm radius dipole with index $n = -\frac{1}{2}$ and the calculated field components radially 10 cm off axis.

xEvolution of the beam emittance and transmission is shown in Fig. 36 as a function of the number of turns in the ring. In 15 turns, the transverse emittance decreases from 1.2 to 0.21 cm yielding a cooling factor of 5.7, the longitudinal emittance decreases from 1.5 to 0.63 cm (cooling factor 2.4), and the 6D emittance decreases from 2.2 to 0.028 cm^3, with an overall cooling factor $\approx$ 79. The transmission is 0.71 without decay and 0.48 with decay. We define a merit factor for cooling that is the total transmission including decay times the 6D cooling factor. The merit factor for this ring is then 38. This implies that transverse emittance at the ring cooler is about the same

FIG. 36. (Color) Evolution of the beam emittance/transmission at the ring cooler. After 15 turns, the transverse emittances are reduced from 1.2 to 0.21 cm and the longitudinal emittance is reduced from 1.5 to 0.63 cm. The 6D emittance is therefore reduced from 2.2 to 0.028 cm^3, giving a 6D cooling factor of 79. The transmission without decay is 71%, and the transmission with decay is 48%. The resulting merit factor is 38.

as at a linear SFOFO cooling channel employed in study II [29], whereas the longitudinal emittance is noticeably less.

This cooler provides mainly transverse cooling and can be used as a part of a neutrino factory or a muon collider. A cooler specially designed for strong longitudinal cooling (''bunch compressor'') can also be created using a similar scheme. Such a compressor would be a part of a muon collider to shorten muon bunches from 6–8 m (minimal length after π-μ decay and phase rotation, see Ref. [9]) to 0.6–0.8 m acceptable for further cooling by a 200 MHz channel.

Two options for the bunch compressor are considered in Ref. [138]. The first one is a two-step cooler where each step is very similar to the ring cooler shown in Fig. 34. The main difference is that the primary goal in the first cooler is the longitudinal bunching of the beam. This leads to a uniform magnetic field in the long solenoids and lower frequency/voltage of the accelerating rf system (15.6 MHz/4 MeV/m at the first stage vs 62.5 MHz/8 MeV/m at the second one). Another option is a 15 MHz octagonal cooler composed of the same cells as in Fig. 34, but with half the bending magnet angle. Decrease of longitudinal emittance from 43 to 2.5–3 cm, as required for the muon collider, is obtained in both cases.

We are proceeding with a realistic simulation of this system using GEANT and ICOOL that employs realistic magnetic fields [139] produced by field calculation programs.

After the two stage cooler, we still need a factor of $\approx$ 30 in transverse cooling, but we are within a factor of 4 in longitudinal cooling relative to the Higgs factory goals. Lithium lens cooling, which with its strong focusing will cool transversely further while degrading longitudinally due to straggling, is a possibility and is being investigated.

2. RFOFO ring coolers

The cooling lattice for the neutrino factory (see Sec. III) employs a configuration of fields known as an SFOFO lattice (super-FOFO) where the axial magnetic field profile changes polarity in alternate cells of the lattice. For the ring cooler design under consideration here, we employ an RFOFO lattice (regular-FOFO) where the axial field profile changes polarity in the middle of a cell. As a result all cells in an RFOFO lattice are identical.

The ring cooler design employs a single cell for both transverse cooling and emittance exchange. It uses solenoids for focusing, giving large angular and momentum acceptances. The cell includes dispersion, acceleration, and energy loss in a single thick hydrogen wedge. Figure 37 shows the layout of the cooling ring drawn to scale. The RFOFO lattice was chosen because, unlike in the SFOFO case used in study II, all cells are strictly identical, and the presence of an integer betatron resonance within the momentum acceptance is eliminated.

The basic 33 m circumference ring is made up of 12 identical 2.75-m long cells. In the figure, this symmetry is broken for injection and extraction, but the magnetic fields in this insertion are nearly identical to those in the rest of the ring. Figure 38 shows a detailed view of three cells of the lattice, in plan (a) and side (b) views.

The longitudinal field on axis has an approximately sinusoidal dependence on position. The beam axis is displaced laterally with respect to the coil centers [as shown in Fig. 38(a)] to minimize horizontal fields that cause vertical beam deviations. The lattice transmits


FIG. 37. (Color) Layout of an RFOFO cooling ring.

FIG. 38. (Color) Three cells of the RFOFO lattice; (a) plan, (b) side.

FIG. 39. (Color) Transmission, normalized transverse emittance, normalized longitudinal emittance, normalized 6-dimensional emittance, and the merit factor, as a function of distance.

particles in the momentum band from 150 to 250 MeV/c. The minimum value of the beta function at the central momentum is 40 cm.

Bending round the ring and the required dispersion are provided by applying an approximately 0.125 T vertical bending field generated by alternately tilting the solenoids [as shown in Fig, 38(b)]. There is no attempt to set the field index n (where $B \propto r^n$) to the 0.5 value, so the focusing is not identical in x and y.

It is found that the acceptance is reduced as the bending field is increased. We thus use a wedge with maximum possible angle (giving zero thickness on one side) and the lowest bending field consistent with adequate emittance exchange. The dispersion at the absorber is approximately 8 cm in a direction 30° from the y axis. The dispersion at the center of the rf is of the opposite sign, and also mostly in the y direction.

The liquid-hydrogen wedge has a central thickness of 28.6 cm and a total wedge angle of 100° and is rotated 30° from the vertical. The rf cavities are at a frequency of 201.25 MHz and have an accelerating gradient of 12 MV/m.

The ICOOL [140] simulation (with results shown in shown Fig. 39) used fields calculated from the specified coils and thus neglects no end field effects. But in this simulation, no absorber, or rf, windows, are included, nor did it include the injection/extraction insertion. The rf was represented by the fields in perfect pillbox cavities. The input tracks were taken from a study-II [29] simulation, using distributions from just upstream of the transverse cooling system. The use of study-II simulated distributions is intended to allow a more realistic estimate of the ring's performance. No attempt was made to match this beam to the ring dispersion or slight differences in the transverse beta functions.

Figure 39 shows the simulated transmission, transverse emittance, longitudinal emittance, 6D emittance, and a merit factor M vs length in the ring. M is given by

$$M = \frac{\epsilon_6(\text{initial})}{\epsilon_6(\text{final})} \times \text{transmission}.$$

Initially, the x emittance falls more rapidly than the y, because it is the y emittance that is exchanged with the longitudinal. But the Larmor rotations soon mix the x and y emittances bringing them to a common value.

After a distance of 400 m (≈ 12 turns), the 6D emittance has fallen by a factor of 238, with a transmission of 48% (66% without decay). The merit factor is 136. The same factor for the study-II cooling lattice, also without windows, is 13. Studies with realistic windows and the injection/extraction insertion added, show lower merit factors, but always far better than the study-II example.

The design of the injection/ejection channels and kickers will be challenging, and this ring could not be used, as is, to replace the study-II cooling channel because the bunch train in that case is too long to fit in the ring. Both problems would be removed in a helical cooling channel. The merit factor for such a channel could be even better than that of the ring because it would be possible to "taper" the optics, as a function of distance down the channel, and thus lower the final equilibrium emittance.

3. Quadrupole ring coolers

Alternative ring designs have been explored which are based on storage rings which consist of conventional quadrupole and dipole magnetic elements instead of solenoids [141]. The strategy has been to utilize the SYNCH storage ring design code [142] to develop linear lattice solutions and then transfer the lattice parameters into the ray-tracing tracking code ICOOL [140] in which absorbers and energy recovery with rf cavities can be included for full simulation. An example of such a ring is shown in

FIG. 40. (Color) Schematic diagram of half of a 22.5° bending cell. A wedge absorber is located in the middle of the cell.

Fig. 40 where the elements of a half-cell for a 22.5° bending cell are depicted schematically. The correspondence between the beam envelope beta and dispersion functions resulting from a simulation with the ICOOL tracking code and the SYNCH calculated values are shown in Fig. 41. The full 16-cell ring is shown in Fig. 42.

In general, we find that the performance of the rings, as measured in emittance reduction along with par-

FIG. 42. (Color) Top view of a sixteen cell muon-cooling ring.

ticle transmission and decay losses, improves when more compact lattice designs are considered. In Fig. 43, the variation of the normalized emittances as a function of ring turns is shown for an 8-cell ring. A reduction of normalized emittance is observed for all three dimensions. This particular ring has a total circumference of 30.9 m. Each half-cell contains one 22.5° combined function dipole proceeded and followed by a single horizontally focusing quadrupole. The average muon beam momentum is 250 MeV/c and liquid-hydrogen absorbers with wedge opening angles of 40° are used. For each cell, the central beam orbit traverses 24 cm of absorber. The energy loss in the wedge absorbers is compensated with 201 MHz rf cavities with peak on-axis gradients of 16 MV/m.


FIG. 41. (Color) The β_x, β_y, and D (dispersion) in a 22.5° bending cell. SYNCH calculations (solid curves) and beam parameters from an ICOOL simulation (marked points) are compared.


FIG. 43. (Color) The evolution of x, y, z normalized emittances in 32 full turns.

FIG. 44. (Color) The transmission and the figure of merit factor as a function of the arc length in 32 full turns.

The muon transmission efficiency and total merit factor (muon survival rate times the ratio of initial to final 6D emittance) is shown in Fig. 44 as a function of ring turns. The merit factor reaches 18, while the muon transmission efficiency including decay losses after 18 full turns is 40%.

Rings in which the focusing function is handled exclusively by the dipole elements have also been explored. In this case, the natural focusing power of the dipole is utilized for horizontal focusing while the entrance and exit dipole edge angles are adjusted to provide the required vertical focusing. Several examples of such lattices have been examined. Because these lattices are more compact than lattices which also include quadrupoles, the performance of these rings tends to be better. As an example, performances with merit factors on the order of 100 have been observed with rings based on 4-cell lattice with eight 45° dipoles. For one case, the entrance and exit angle dipole faces are 7.4° and 21.8° relative to the normal of the beam trajectory. The ring for this example has a circumference of only 9.8 m and the design of injection/ejection cells may prove to be very challenging.

4. Injection into ring coolers

The most serious technical problem facing the ring cooler approach is the injection system which may require a very powerful kicker magnet [143]. The energy stored in the injection kicker goes as the square of the emittance of the beam and inversely as the circumference of the ring. A promising injection scheme that does not use kicker magnets, but instead uses absorbers to degrade the beam energy and rf phase manipulations has been proposed [144] and is being studied.

C. Higher energy muon colliders

Once the required cooling has been achieved to make the first muon collider feasible, acceleration to higher energies becomes possible. Colliders with 4 TeV center of mass energy have been studied [9] and Table XV lists the parameters for such a collider. The muons are accelerated initially by a linear accelerator followed by a series of RLA's followed by rapid cycling synchrotrons (RCS's). The radiation from the neutrinos from the muon decay begins to become a problem at CoM energies of 3 TeV [124].

There have been preliminary attempts to study colliders of even higher energy, starting at 10 TeV all the way up to 100 TeV in the center of mass and we include the references to these studies [145] for the sake of completeness.

D. Muon collider detectors

Figure 45 shows a straw-man muon collider detector for a Higgs factory simulated in GEANT. The background from muon-decay sources has been extensively studied [9]. At the Higgs factory, the main sources of background are from photons generated by the showering of muon-decay electrons. At the higher energy colliders, Bethe-Heitler muons produced in electron showers become a problem. Work was done to optimize the shielding by using specially shaped tungsten cones [9] that reduce the backgrounds resulting from electromagnetic showers from entering the detector. The occupancy levels due to

TABLE XV. Parameters of acceleration for a 4 TeV muon collider.

|	Linac	RLA1	RLA2	RCS1	RCS2
E (GeV)	$0.1 \to 1.5$	$1.5 \to 10$	$10 \to 70$	$70 \to 250$	$250 \to 2000$
f_{rf} (MHz)	$30 \to 100$	200	400	800	1300
N_{turns}	1	9	11	33	45
V_{rf} (GV/turn)	1.5	1.0	6	6.5	42
C_{turn} (km)	0.3	0.16	1.1	2.0	11.5
Beam transit time (ms)	0.0013	0.005	0.04	0.22	1.73
$\sigma_{z,beam}$ (cm)	$50 \to 8$	$4 \to 1.7$	$1.7 \to 0.5$	$0.5 \to 0.25$	$0.25 \to 0.12$
$\sigma_{E,beam}$ (GeV)	$0.005 \to 0.033$	$0.067 \to 0.16$	$0.16 \to 0.58$	$0.58 \to 1.14$	$1.14 \to 2.3$
Loss (%)	5	7	6	7	10

FIG. 45. (Color) Cut view of a straw-man detector in GEANT for the Higgs factory with a Higgs $\rightarrow b\bar{b}$ event superimposed. No backgrounds shown. The tungsten cones on either side of the interaction region mask out a 20° area.

background photons and neutrons in detectors were shown to be similar to those predicted for the LHC experiments. It still needs to be established whether pattern recognition is possible in the presence of these backgrounds, especially the Bethe-Heitler muons, which are a unique source of background to muon collider detectors.

V. R&D STATUS

A successful construction of a muon storage ring to provide a copious source of neutrinos requires several novel approaches to be developed and demonstrated; a high-luminosity muon collider requires an even greater extension of the present state of accelerator design. Thus, reaching the full facility performance in either case requires an extensive R&D program.

The required R&D program has been identified for each of the major neutrino factory systems. In particular, some critical components must be prototyped and their performance verified. For example, the cooling channel assumes a normal conducting rf (NCRF) cavity gradient of 17 MV/m at 201.25 MHz, and the acceleration section demands similar performance from superconducting rf (SCRF) cavities at this frequency. In both cases, the requirements are beyond the performance reached to date for cavities in this frequency range. The ability of the target to withstand a proton beam power of up to 4 MW must be confirmed, and, if it remains the technology of choice, the ability of an induction linac unit to

coexist with its internal superconducting (SC) solenoid must be verified. 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.

Below we give an overview of the MC R&D program goals and list the specific questions we expect to address. We also summarize briefly the R&D accomplishments to date.

A. R&D program overview

A neutrino factory comprises the following major systems: proton driver, target and (pion) capture section, (pion-to-muon) decay and phase rotation section, bunching and matching section, cooling section, acceleration section, and storage ring. These same categories apply to a muon collider, with the exception that the storage ring is replaced by a collider ring having a low-beta interaction point and a local detector. Parameters and requirements for the various systems are generally more severe in the case of the muon collider, so a neutrino factory can properly be viewed as a scientifically productive first step toward the eventual goal of a collider.

The R&D program we envision is designed to answer the key questions needed to embark upon a zeroth-order design report (ZDR). The ZDR will examine the complete set of systems needed for a neutrino factory and show how the parts can be integrated into a coherent

whole. Although a fully engineered design with a detailed cost estimate is beyond the scope of a ZDR, enough detailed work must been accomplished to ensure that the critical items are technically feasible and that the proposed facility could be successfully constructed and operated at its design specifications. By the end of the full R&D program, it is expected that a formal conceptual design report (CDR) for a neutrino factory could be written. The CDR would document a complete and fully engineered design for the facility, including a detailed bottom-up cost estimate for all components. This document would form the basis for a full technical, cost, and schedule review of the construction proposal. Construction could then commence after obtaining approval.

The R&D issues for each of the major systems must be addressed by a mix of theoretical calculation, simulation modeling, and experimental studies, as appropriate. A list of the key physics and technology issues for each major system is given.

1. Proton driver

Production of intense, short proton bunches, e.g., with space-charge compensation and/or high-gradient, low frequency rf systems.

2. Target and capture section

Optimization of target material (low Z or high Z) and form (solid, moving band, liquid-metal jet).

Design and performance of a high-field solenoid (≈ 20 T) in a very high radiation environment.

3. Decay and phase rotation section

Development of high-gradient induction linac modules having an internal superconducting solenoid channel.

Examination of alternative approaches, e.g., based upon combined rf phase rotation and bunching systems or fixed-field, alternating gradient (FFAG) rings.

4. Bunching and matching section

Design of efficient and cost-effective bunching system.

5. Cooling section

Development and testing of high-gradient NCRF cavities at a frequency near 200 MHz.

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

Development and testing of liquid-hydrogen absorbers for muon cooling.

Development and testing of an alternative gaseous-absorber cooling-channel design incorporating pressurized, high-gradient rf cavities.

Development and testing of candidate diagnostics to measure emittance and optimize cooling-channel performance.

Design of beam line and test setup (e.g., detectors) needed for demonstration of transverse emittance cooling.

Development of a full 6D analytical theory to guide the design of the cooling section.

6. Acceleration section

Optimization of acceleration techniques to increase the energy of a muon beam (with a large momentum spread) from a few GeV to a few tens of GeV (e.g., recirculating linacs, rapid cycling synchrotrons, FFAG rings) for a neutrino factory, or even higher energy for a muon collider.

Development of high-gradient SCRF cavities at frequencies near 200 MHz, along with efficient power sources (about 10 MW peak) to drive them.

Design and testing of components (rf cavities, magnets, diagnostics) that will operate in the muon-decay radiation environment.

7. Storage ring

Design of large-aperture, well-shielded superconducting magnets that will operate in the muon-decay radiation environment.

8. Collider

Cooling of 6D emittance (x, p_x, y, p_y, t, and E) by up to a factor of 10^5–10^6.

Design of a collider ring with very low β^* (a few mm) at the interaction point having sufficient dynamic aperture to maintain luminosity for about 500 turns.

Study of muon beam dynamics at a large longitudinal space-charge parameter and at high beam-beam tune shift.

9. Detector

A study of cost trade-offs between increasing the detector mass compared with increasing the beam intensity.

Simulation studies to define acceptable approaches for a muon collider detector operating in a high-background environment

Most of these issues are being actively pursued as part of the ongoing MC R&D program. In a few areas, notably the proton driver and detector, the MC does not currently engage in R&D activities, though independent efforts are under way. Longer-term activities, related primarily to the muon collider, are supported and encouraged as resources permit.

B. Recent R&D accomplishments

1. Targetry

A primary effort of the Targetry Experiment E951 at BNL has been to carry out initial beam tests of both a solid carbon target and a mercury target. Both of these tests have been made with a beam intensity of about 4×10^{12} ppp, with encouraging results.

In the case of the solid carbon target, it was found that a carbon-carbon composite having nearly zero coefficient of thermal expansion is largely immune to beam-induced pressure waves. Sublimation losses are a concern. However, a carbon target in a helium atmosphere is expected to have negligible sublimation loss. If radiation damage is the limiting effect for a carbon target, the predicted lifetime would be about 12 weeks when bombarded with a 1 MW proton beam. Carbon targets would therefore seem viable for beam powers up to 1 MW, or perhaps a little higher.

For a mercury jet target, tests with about 2×10^{12} ppp showed that the jet is not dispersed until long after the beam pulse has passed through the target. Measurements of the velocity of droplets emanating from the jet as it is hit with the proton beam pulse from the AGS (≈ 10 m/s for 25 J/g energy deposition) compare favorably with simulation estimates. High-speed photographs indicate that the beam disruption at the present intensity does not propagate back upstream toward the jet nozzle. If this remains true at the higher intensity of 1.6×10^{13} ppp, it will ease mechanical design issues for the nozzle. The tests so far have used a jet with an initial mercury velocity of 2 m/sec. At a neutrino factory a 20 m/sec jet is envisioned. A prototype is under development.

2. MUCOOL

The primary effort in the muon ionization cooling (MUCOOL) R&D has been to complete the Lab G high-power rf test area at Fermilab and begin high-power tests of 805-MHz rf cavities. A test solenoid for the facility, capable of operating either in solenoid mode (its two independent coils powered in the same polarity) or gradient mode (with the two coils opposed), was commissioned up to its design field of 5 T.

An 805 MHz open-cell cavity has been tested in Lab G to look at gradient limitations, magnetic field effects, and compatibility of the rf cavities with other systems. We have measured the dark currents over a range covering 14 orders of magnitude and accumulated data on the momentum spectrum, angular distribution, pulse shape, dependence on conditioning, and dependence on magnetic fields [146]. The dark currents seem to be described by the Fowler-Nordheim field emission process, which results from very small emitter sources (submicron sizes) at very high local electric fields (5–8 GV/m). This im-plies that the emitter fields are enhanced by large factors, $\beta_{FN} = \sim 500$, over the accelerating field. (At these electric fields the electrostatic stress becomes comparable to the strength of hardened copper.) We have shown how both normal conditioning and nitrogen processing can reduce dark currents. Our data from the 805 MHz cavity have been compared with other data from NLC cavities, superconducting TESLA cavities, and 200 MHz proton linacs, showing that all cavities seem to be affected by the same processes.

We have also looked at damage produced on irises and windows, primarily when the system is run with the solenoid magnet on. A number of effects are seen: copper splatters on the inside of the thin Ti window, burn marks on the outside of the window due to electron beamlets, and some craters, evidently produced by breakdown on the irises. The electron beamlets burned through the windows twice. We have measured the parameters of the beamlets produced from individual emitters when the magnetic field is on, and we have seen ring beams, presumably produced by $E \times B$ drifts during the period when the electrons are being accelerated. The radius of the beamlets is found to be proportional to E/B^2.

We are proceeding with an experimental program designed to minimize the dark currents using surface treatment of the copper cavity.

A second cavity, a single-cell pillbox having foils to close the beam iris, has been tuned to final frequency, shipped to Fermilab, and testing has begun at Lab G. This cavity will permit an assessment of the behavior of the foils under rf heating and give indications about multipactor effects. It will also be used to study the dark current effects discussed above. An advantage of the pillbox cavity is that its windows can be replaced with ones made from (or coated with) various materials and cleaned or polished by various techniques.

Development of a prototype liquid-hydrogen absorber is in progress. Several large diameter, thin (125 μm) aluminum windows have been successfully fabricated by machining from solid disks. These have been pressure tested with water and found to break at a pressure consistent with design calculations. A new area, the MUCOOL test area (MTA), is being developed at FNAL for testing the absorbers. The MTA, located at the end of the proton linac, will be designed to eventually permit beam tests of components and detectors with 400 MeV protons. It will also have access to 201-MHz high-power rf amplifiers for testing of future full-sized 201-MHz cavities.

Initial plans for a cooling demonstration are well under way. This topic is covered separately in more detail in Sec. VI.

A parallel cooling-channel development effort based on the use of gaseous hydrogen or helium energy absorber has begun. Muons Inc. [147] has received a DOE Small Business Technology Transfer grant with the Illinois

Institute of Technology to develop cold, pressurized high-gradient rf cavities for use in muon ionization cooling. These cavities will be filled with dense gas, which suppresses high voltage breakdown by virtue of the Paschen effect and also serves as the energy absorber. A program of development for this alternative approach to ionization cooling is foreseen that starts with Lab G tests, evolves to an MTA measurement program, and leads to the construction of a cooling-channel section suitable for tests in MICE.

3. Feasibility study II

The MC has participated in a second feasibility study for a neutrino factory, cosponsored by BNL. The results of the study were encouraging (see Sec. III), indicating that a neutrino intensity of 1×10^{20} per Snowmass year per MW can be sent to a detector located 3000 km from the muon storage ring. It was also clearly demonstrated that a neutrino factory could be sited at either FNAL or BNL. Hardware R&D needed for such a facility was identified and is a major part of the program outlined here.

4. Beam simulations and theory

In addition to work on study II, our present effort has focused on longitudinal dynamics [148]. We are developing theoretical tools for understanding the longitudinal aspects of cooling, with the goal of developing approaches to 6D cooling, i.e., "emittance exchange." This is a crucial aspect for the eventual development of a muon collider, and would also benefit a neutrino factory. Improved designs for the phase rotation, bunching, and acceleration systems are also being explored, with an emphasis on preparing the way for a future design study in which the performance obtained by the study-II design is maintained, but with a reduction in cost.

5. Other component development

At present, the main effort in this area is aimed at development of a high-gradient 201-MHz SCRF cavity. A test area of suitable dimensions has been constructed at Cornell. In addition, a prototype cavity has been fabricated for the Cornell group by our CERN colleagues. Mechanical engineering studies of microphonics and Lorentz detuning issues are being carried out.

6. Collider R&D

Studies of possible hardware configurations to perform emittance exchange, such as the compact ring proposed by Balbekov *et al.* [137], are now getting under way. A ring cooler has the potential to cool in 6D phase space, provided the beam can be injected into and extracted from it. A series of workshops have been held on the topics of emittance exchange and ring coolers that have helped further our understanding of both.

VI. INTERNATIONAL MUON IONIZATION COOLING EXPERIMENT

A. Motivation

Ionization cooling of minimum-ionizing muons is an important ingredient in the performance of a neutrino factory. However, it has not been demonstrated experimentally. We seek to carry out an experimental demonstration of cooling in a muon beam. Towards this goal, we have developed (in collaboration with a number of physicists from Europe and Japan interested in neutrino factories) a conceptual design for an international muon ionization cooling experiment. A proposal for MICE has recently been submitted to the Rutherford Appleton Laboratory in England [24].

The aim of the proposed cooling experimental demonstration is as follows:

To show that we can design, engineer, and build a section of cooling channel capable of giving the desired performance for a neutrino factory.

To place it in a beam and measure its performance, i.e., experimentally validate our ability to simulate precisely the passage of muons confined within a periodic lattice as they pass through energy absorbers and rf cavities. The experience gained from this experimental demonstration will provide important input to the final design of a real cooling channel. The successful operation of a section of a muon-cooling channel has been identified (most recently by the U.S. Muon Technical Advisory Committee [149]) as a key step in demonstrating the feasibility of a neutrino factory or muon collider.

B. Principle of the experiment

Fundamentally, in a muon-cooling experiment one needs to measure, before and after the cooling channel, the phase space distribution of a muon beam in six dimensions [150]. Such a measurement must include the incoming and outgoing beam intensities and must avoid biases due to the decay of muons into electrons within the channel and due to possible contamination of the incoming beam by nonmuons [151]. Two techniques have been considered: (i) the multiparticle method, in which emittance and number of particles in any given volume of phase space are determined from the global properties of a bunch; and (ii) the single-particle method, in which the properties of each particle are measured and a "virtual bunch" formed off-line. The full determination of the covariance matrix in six dimensions is a delicate task in a multiparticle experiment, and the desired diagnostics would have to be developed specifically for this purpose; moreover, a high-intensity muon beam bunched at an appropriate frequency would need to be designed and built. For these reasons, the single-particle method is preferred. The single-particle approach, typical of particle-physics experiments, is one for which experimental

FIG. 46. (Color) Conceptual layout of MICE upstream spectrometer: following an initial TOF measurement, muons are tracked using detector planes located within a solenoidal magnetic field. Although in principle three x, y measurements as shown suffice to determine the parameters of each muon's helical trajectory, in practice additional measurement redundancy will be employed; for example, a fourth measurement plane can be used to eliminate very-low-momentum muons that would execute multiple cycles of helical motion. A similar spectrometer (but with the time-of-flight measurement at the end) will be used downstream of the cooling apparatus.

methods already exist and suitable beams are already available.

In the particle-by-particle approach, the properties of each particle are measured in magnetic spectrometers before and after the cooling channel (Fig. 46). Each spectrometer measures, at given z positions, the coordinates x, y of every incident particle, as well as the time. Momentum and angles are reconstructed by using more than one plane of measurement. For the experimental errors not to affect the measurement of the emittance by a significant factor, the rms resolution of the measurements must be smaller than typically 1/10th of the rms equilibrium beam size in each of the six dimensions [152].

C. Conceptual design

Figure 47 shows the layout under consideration for MICE, which is based on two cells of the feasibility study-II "lattice 1" cooling channel. The incoming muon beam encounters first a beam preparation section, where the appropriate input emittance is generated by a pair of high-Z (lead) absorbers. In addition, a precise time measurement is performed and the incident particles are identified as muons. There follows a first measurement section, in which the momenta, positions, and angles of the incoming particles are measured by means of tracking devices located within a uniform-field solenoid. Then comes the cooling section itself, with hydrogen absorbers and 201 MHz rf cavities, the lattice optics being provided by a series of superconducting coils; the pairs of coils surrounding each absorber have opposite magnetic fields ("bucking" solenoids), providing tight focusing. The momenta, positions, and angles of the outgoing particles are measured within a second solenoid, equipped with a tracking system identical to the first one. Finally, another time-of-flight (TOF) measurement is performed together with particle identification to eliminate those muons that have decayed within the apparatus.

D. Performance

Simulations of MICE have been carried out for a configuration including four tracking stations per

FIG. 47. (Color) Schematic layout of the MICE apparatus.

FIG. 48. (Color) A possible MICE tracking-detector configuration.

spectrometer, each station consisting of three crossed planes of 500-μmmicron-thick square-cross-section scintillating fibers (Fig. 48), immersed in a 5 T solenoidal field. TOF is assumed to be measured to 70 ps rms. As shown in Fig. 49, measurement resolution and multiple scattering of the muons in the detector material introduce a correctable bias in the measured emittance ratio of only 1%. (For this study the effect of the cooling apparatus was "turned off" so as to isolate the effect of the spectrometers.)

Figure 50 illustrates the muon-cooling performance of the proposed MICE cooling apparatus. The normalized transverse emittance of the incoming muon beam is reduced by about 8%. The longitudinal emittance increases by about the same amount, thus the net cooling
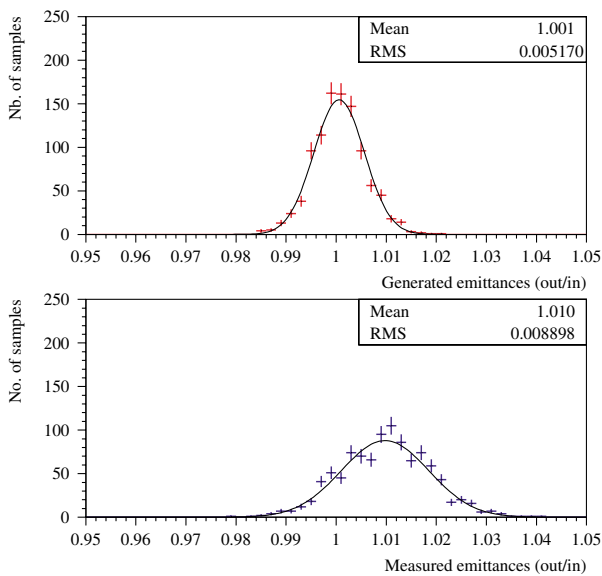

FIG. 49. (Color) Distribution of ratios of output to input 6-dimensional emittance for 1000 simulated experiments, each with 1000 accepted muons. The top figure shows the distribution of this ratio for the emittances as generated by simulation; the bottom figure, as "measured" in the simulated experiments. The curves are Gaussian fits to the points.
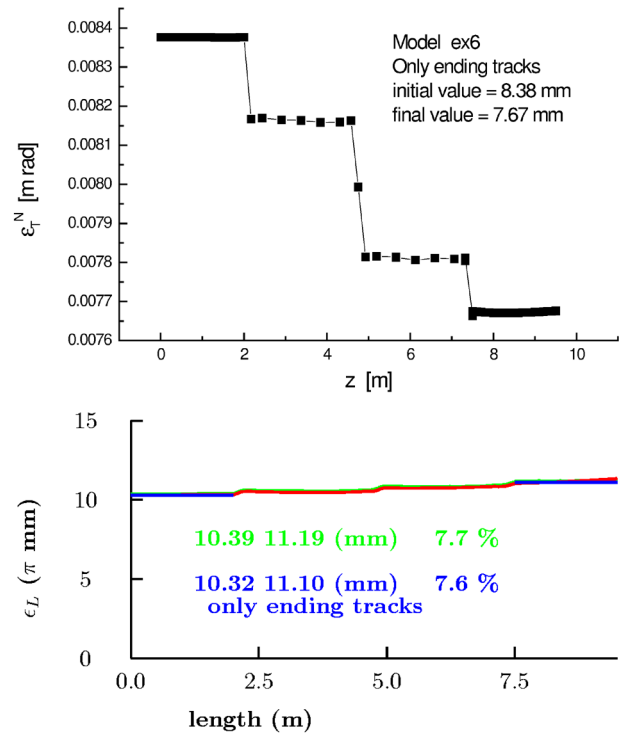

FIG. 50. (Color) Results from ICOOL simulation of MICE: normalized transverse (top) and longitudinal (bottom) emittances vs distance.

in six dimensions is also about 8%. These are large enough effects to be straightforwardly measured by the proposed spectrometers.

The CERN Neutrino Factory Working Group has studied a variant of the proposed MICE cooling apparatus, in which 88-MHz rf cavities are employed in place of the 201-MHz devices (the 88- and 201-MHz designs have similar cooling performance) [153]. Figure 51 (from the CERN study) elucidates further experimental issues. As shown in Fig. 51(a), for input emittance above the equilibrium emittance of the channel (here about 3500 mm mrad), the beam is cooled, while for input
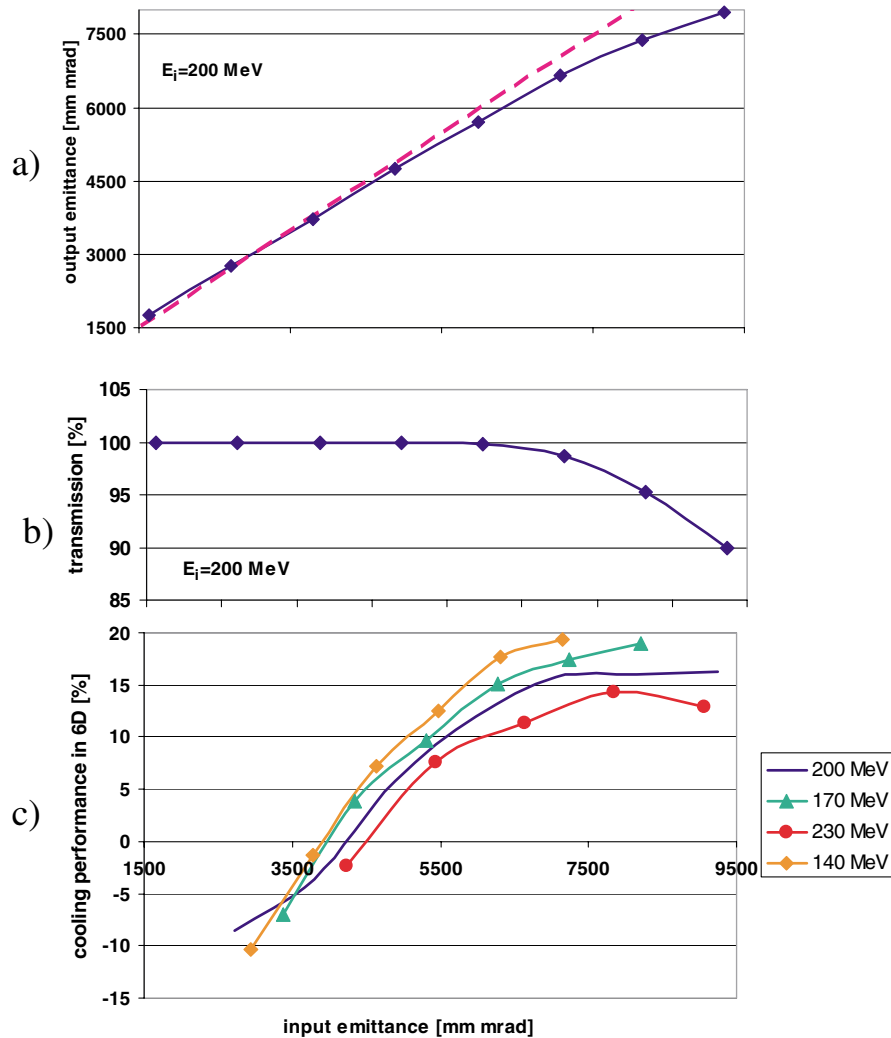
FIG. 51. (Color) Simulation results for 88-MHz variant of MICE apparatus: (a) output emittance vs input emittance, with 45° line (dashed line) superimposed; (b) beam transmission vs input emittance; (c) cooling performance (see text) vs input emittance for various beam kinetic energies (top to bottom: 140, 170, 200, and 230 MeV).

emittance below equilibrium it is heated (and, of course, for an input beam at the equilibrium emittance, the output emittance equals the input emittance). Figure 51(b) illustrates the acceptance cutoff of the cooling-channel lattice; for input emittance above 6000 mm mrad, the transmission probability falls below 100% due to scraping of the beam. Figure 51(c) shows the effect of varying the beam momentum: cooling performance improves as the momentum is lowered,[5] as quantified here in terms of the fractional increase in the number of muons within the phase-space volume accepted by a hypothetical acceleration section downstream of the cooling channel.

The goal of MICE includes verification of these effects in detail in order to show that the performance of the cooling apparatus is well understood. Subsequent running could include tests of additional transverse cooling cells, alternative designs, or emittance exchange cells.

One critical aspect of this experiment is operation in the presence of backgrounds due to dark currents from the rf cavities. While it is possible to operate the experiment using comparatively low rf gradients, it would be highly desirable to produce cavities which would yield less dark current at higher gradients. This would permit more efficient use of the rf cavities and power supplies. We are trying to develop cavities with low dark currents.

VII. CONCLUSIONS

In summary, the Muon Collaboration is developing the knowledge and ability to create, manipulate, and accelerate muon beams. Our R&D program will position the

---

[5]Despite the increased cooling efficiency at low muon momentum, simulations of an entire muon production section and cooling channel suggest that momenta near the ionization minimum represent the global optimum for neutrino factory performance.

high energy physics community such that, when it requires a neutrino factory or a muon collider, we shall be in a position to provide it. A staged plan for the deployment of a neutrino factory has been developed that provides an active neutrino and muon physics program at each stage. The requisite R&D program is diversified over laboratories and universities and has international participation.

The very fortuitous situation of having intermediate steps along this path, that offer a powerful and exciting physics program in their own right, presents an ideal scientific opportunity, and it is hoped that the particle-physics community will be able to take advantage of it.

———

[1] SNO Collaboration, Q. R. Ahmad *et al.*, nucl-ex/0204008; Q. R. Ahmad *et al.*, nucl-ex/0204009; Q. R. Ahmad *et al.*, Phys. Rev. Lett. **87**, 071301 (2001).

[2] Super-Kamiokande Collaboration, Y. Fukuda *et al.*, Phys. Lett. B **433**, 9 (1998); Phys. Lett. B **436**, 33 (1998); Phys. Rev. Lett. **81**, 1562 (1998); Phys. Rev. Lett. **82**, 2644 (1999).

[3] The MC Collaboration website is http://www.cap.bnl.gov/mumu/

[4] G. I. Budker, in *Proceedings of the 7th International Conference on High Energy Accelerators, Yerevan, 1969*, edited by A. I. Alikhanian (Publishing House of the Academy of Sciences of Armenia SSR, Yerevan, USSR, 1970), p. 33; extract in *Proceedings of the Physics Potential and Development of $\mu^+\mu^-$ Colliders: Second Workshop*, edited by D. Cline, AIP Conf. Proc. No. 352 (AIP, New York, 1996), p. 4.

[5] A. N. Skrinsky, in Proceedings of the International Seminar on Prospects of High-Energy Physics, Morges, 1971 (unpublished); extract in *Proceedings of the Physics Potential and Development of $\mu^+\mu^-$ Colliders: Second Workshop* (Ref. [4]), p. 6.

[6] A. N. Skrinsky and V. V. Parkhomchuk, Sov. J. Part. Nucl. **12**, 223 (1981).

[7] D. Neuffer, Part. Accel. **14**, 75 (1983).

[8] R. B. Palmer, D. Neuffer, and J. Gallardo, in *A Practical High-Energy High-Luminosity $\mu^+\mu^-$ Collider*, Proceedings of the Advanced Accelerator Concepts: 6th Annual Conference, edited by P. Schoessow, AIP Conf. Proc. No. 335 (AIP, New York, 1995), p. 635; D. Neuffer and R. B. Palmer, in *Progress Toward a High-Energy, High-Luminosity $\mu^+\mu^-$ Collider*, Proceedings of the Future of Accelerator Physics: The Tamura Symposium, edited by T. Tajima, AIP Conf. Proc. No. 356 (AIP, New York, 1996), p. 344.

[9] Muon Collider Collaboration, Charles M. Ankenbrandt *et al.*, Phys. Rev. ST Accel. Beams **2**, 081001 (1999).

[10] BNL Report No. BNL-52503, 1996, Fermilab Report No. Conf-96/092, 1996, LBNL Report No. LBNL-38946, 1996.

[11] Mucool Notes, http://www-mucool.fnal.gov/notes/notes.html

[12] D. Koshkarev, CERN Report No. CERN/ ISR-DI/74-62, 1974.

[13] B. J. King, physics/9907026 (unpublished).

[14] S. Geer, in *Workshop on Physics at the First Muon Collider and at the Front End of a Muon Collider*, edited by S. Geer and R. Raja, AIP Conf. Proc. No. 435 (AIP, New York, 1998), p. 384.

[15] S. Geer, Phys. Rev. D **57**, 6989 (1998).

[16] D. A. Harris and Kevin McFarland, in [14], hep-ex/9804009; hep-ex/9804010; B. J. King, hep-ex/9907035; I. Bigi *et al.*, hep-ph/0106177.

[17] NuFact99, Lyon, France, http://lyoinfo.in2p3.fr/nufact99/. *Proceedings of the ICFA/ECFA Workshop on Neutrino Factories based on Muon Storage Rings, Lyon, France, 1999*, edited by W. Barletta [Nucl. Instrum. Methods Phys. Res., Sect. A **451**, 1–384 (2000)].

[18] NuFact00, Monterey, California, http://www.lbl.gov/Conferences/nufact00/. *Proceedings of the International Workshop NuFact'00, Muon Storage Ring for a Neutrino Factory, Monterey, CA, 2000*, edited by W. Barletta [Nucl. Instrum. Methods Phys. Res., Sect. A **472**, 323–652 (2001)].

[19] *Proceedings of the Third International Workshop on Neutrino Factories based on Muon Storage Rings, Tsukuba, Ibaraki, Japan, 2001*, edited by Shinji Machida and Koji Yoshimura [Nucl. Instrum. Methods Phys. Res., Sect. A, 503, 1–424 (2003)].

[20] *Proceedings of NuFact02, Imperial College, London*, http://www.hep.ph.ic.ac.uk/NuFact02/

[21] *Proceedings of NuFact03, Columbia University, New York*, http://www.cap.bnl.gov/nufact03

[22] Neutrino factory and muon storage rings at CERN, http://muonstoragerings.web.cern.ch/muonstoragerings/

[23] The Japanese neutrino factory, http://www-prism.kek.jp/nufactj/

[24] The MICE proposal can be found at http://hep04.phys.iit.edu/cooldemo/notes/notes.html

[25] MUCOOL home page, http://www.fnal.gov/projects/muon_collider/cool/cool.html; http://hepunx.rl.ac.uk/~edgecock/muons/muscat.html; emittance exchange home page, http://www.fnal.gov/projects/muon_collider/eexchzange/; targetry home page, http://www.hep.princeton.edu/mumu/target/.

[26] N. Holtkamp and D. Finley, Fermilab Report No. Fermilab-Pub-00/108-E, 2000, http://www.fnal.gov/projects/muon_collider/nu-factory/nu-factory.html

[27] C. Albright *et al.*, Fermilab Report No. FN692, 2000, hep-ex/0008064; http://www.fnal.gov/projects/muon_collider/nu/study/study.html

[28] V. Barger, R. Bernstein, A. Bueno, M. Campanelli, D. Casper, F. DeJongh, S. Geer, M. Goodman, D. A. Harris, K. S. McFarland, N. Mokhov, J. Morfin, J. Nelson, F. Peitropaolo, R. Raja, J. Rico, A. Rubbia, H. Schellman, R. Shrock, P. Spentzouris, R. Stefanski, L. Wai, and K. Whisnant, Fermilab Report No. FERMILAB-FN-703, hep-ph/0103052, 2001.

[29] S. Ozaki, R. Palmer, M. S. Zisman, and J. Gallardo, BNL Report No. BNL-52623, 2001; http://www.cap.bnl.gov/mumu/studyii/FS2-report.html.

[30] UNO Official Home Page, http://superk.physics.sunysb.edu/uno/

[31] D. B. Cline, F. Sergiampietri, J. G. Learned, and K. McDonald, astro-ph/0105442; K. T. McDonald,

hep-ex/0204037; D. Ayres *et al.*, physics/9911009. An early proposal for a liquid-argon detector, ARGONAUT, may be found in G. Harigel, H. Kautzky, P. McIntyre, and A. Van Ginneken, Fermilab Report No. 601D, 1978.

[32] F. Arneodo *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **455**, 376 (2000).

[33] Some of the text quoted here appears in the study-II document [29]. We would like to thank Robert Shrock for helpful comments.

[34] Fits and references to the Homestake, Kamiokande, GALLEX, SAGE, Super Kamiokande, and SNO data include N. Hata and P. Langacker, Phys. Rev. D **56**, 6107 (1997); J. Bahcall, P. Krastev, and A. Smirnov, Phys. Rev. D **58**, 096016 (1998); J. Bahcall and P. Krastev, Phys. Lett. B **436**, 243 (1998); J. Bahcall, P. Krastev, and A. Smirnov, Phys. Rev. D **60**, 093001 (1999); hep-ph/0103179; M. Gonzalez-Garcia, C. Peña-Garay, and J.W. F. Valle, Phys. Rev. D **63**, 013007 (2001); M. Gonzalez-Garcia *et al.*, Phys. Rev. D **63**, 033005 (2001); V. Barger, D. Marfatia, K. Whisnant, and B. P. Wood, Phys. Rev. D **64**, 073009 (2001); P. Krastev and A. Smirnov, Phys. Rev. D **65**, 073022 (2002); J. N. Bahcall, M. C. Gonzalez-Garcia, and C. Peña-Garay, hep-ph/0111150; M. Maltoni, T. Schwetz, and J.W. F. Valle, Phys. Rev. D **65**, 093004 (2002). Recent discussions of flux calculations are found in J. Bahcall, Phys. Rep. **333**, 47 (2000), hep-ph/0009044, Proceedings of Neutrino-2000, and http://www.sns. ias.edu/~jnb/; V. Barger, D. Marfatia, and K. Whisnant, Phys. Rev. Lett. **88**, 011302 (2002); Phys. Lett. B **509**, 19 (2001); V. Barger, D. Marfatia, K. Whisnant, and B. P. Wood, hep-ph/0204253. Super Kamiokande data are reported and analyzed in SuperKamiokande Collaboration, Y. Fukuda *et al.*, Phys. Rev. Lett. **82**, 1810 (1999); S. Fukuda *et al.*, Phys. Rev. Lett. **86**, 5651 (2001). For recent reviews, see, e.g., Y. Suzuki, Nucl. Phys. B, Proc. Suppl. **91**, 29 (2001); Y. Takeuchi, in *Proceedings of the International Conference on High Energy Physics, Osaka*, http://ichep2000.hep.sci. osaka-u.ac.jp; Proceedings of the Fifth Topical Workshop at the Gran Sasso National Laboratory: Solar Neutrinos, Assergi, Italy, 2001 (to be published).

[35] L. Wolfenstein, Phys. Rev. D **17**, 2369 (1978).

[36] V. Barger, K. Whisnant, S. Pakvasa, and R. J. N. Phillips, Phys. Rev. D **22**, 2718 (1980).

[37] V. Barger *et al.*, Phys. Rev. D **22**, 1636 (1980).

[38] S. P. Mikheyev and A. Smirnov, Yad. Fiz. **42**, 1441 (1985) [Sov. J. Nucl. Phys. **42**, 913 (1986)]; Nuovo Cimento Soc. Ital. Fis., C **9**, 17 (1986); S. P. Rosen and J. Gelb, Phys. Rev. D **34**, 969 (1986); S. Parke, Phys. Rev. Lett. **57**, 1275 (1986); W. Haxton, Phys. Rev. Lett. **57**, 1271 (1986); T. K. Kuo and J. Pantaleone, Rev. Mod. Phys. **61**, 937 (1989).

[39] V. Barger, R. J. N. Phillips, and K. Whisnant, Phys. Rev. D **24**, 538 (1981); S. L. Glashow and L. M. Krauss, Phys. Lett. B **190**, 199 (1987).

[40] Kamiokande Collaboration, K. S. Hirata *et al.*, Phys. Lett. B **205**, 416 (1988); **280**, 146 (1992); Y. Fukuda *et al.*, Phys. Lett. B **335**, 237 (1994); S. Hatakeyama *et al.*, Phys. Rev. Lett. **81**, 2016 (1998).

[41] IMB Collaboration, D. Casper *et al.*, Phys. Rev. Lett. **66**, 2561 (1991); R. Becker-Szendy *et al.*, Phys. Rev. D **46**, 3720 (1992); Phys. Rev. Lett. **69**, 1010 (1992).

[42] Y. Fukuda *et al.*, Phys. Lett. B **433**, 9 (1998); Phys. Rev. Lett. **81**, 1562 (1998); **82**, 2644 (1999); Phys. Lett. B **467**, 185 (1999); H. Sobel, Proceedings of Neutrino-2000, Nucl. Phys. B, Proc. Suppl. **91**, 127 (2001); T. Toshito, in ICHEP-2000. Recent discussions of flux calculations are in T. Gaisser, Nucl. Phys. B, Proc. Suppl. **87**, 145 (2000); P. Lipari, Astropart. Phys. **14**, 153 (2000); G. Battistoni, hep-ph/0012268; G. Fiorentini, V. Naumov, and F. Villante, Phys. Lett. B **510**, 173 (2001).

[43] W. Allison *et al.*, Phys. Lett. B **449**, 137 (1999); A. Mann, Proceedings of Neutrino-2000, hep-ex/0007031.

[44] M. Ambrosio *et al.*, Phys. Lett. B **478**, 5 (2000); B. Barish, Proceedings of Neutrino-2000, Nucl. Phys. B, Proc. Suppl. **91**, 141 (2001).

[45] M. Apollonio *et al.*, Phys. Lett. B **420**, 397 (1998); **466**, 415 (1999).

[46] LSND Collaboration, C. Athanassopoulous *et al.*, Phys. Rev. Lett. **77**, 3082 (1996); **81**, 1774 (1998); K. Eitel, Proceedings of Neutrino-2000.

[47] KARMEN Collaboration, K. Eitel, in Proceedings of Neutrino-2000 [Nucl. Phys. B, Proc. Suppl. **91**, 191 (2000)].

[48] References and websites for these experiments and future projects can be found, e.g., at http://www.hep.anl.gov/ndk/hypertext/nu_industry.html

[49] M. Sakuda and K. Nishikawa, in *Proceedings of the 30th International Conference on High Energy Physics, Osaka, 2000*, edited by C. S. Lima and T. Yamanaka (World Scientific, Singapore, 2000); S. H. Ahn *et al.*, Phys. Lett. B **511**, 178 (2001).

[50] OPERA Collaboration, CERN Report No. CERN-SPSC-97-24, 1997, hep-ex/9812015.

[51] ICARUS/ICANOE Collaboration, F. Cavanna *et al.*, INFN Report No. LNGS-P21-99-ADD-1,2, 1999; A. Rubbia, hep-ex/0001052.

[52] V. Barger, A. M. Gago, D. Marfatia, W. J. Teves, B. P. Wood, and R. Zukanovich Funchal, Phys. Rev. D **65**, 053016 (2002).

[53] JHF Neutrino Working Group, Y. Itoh *et al.*, Letter of Intent: A Long Baseline Neutrino Oscillation Experiment using the JHF 50 GeV Proton Synchrotron and the Super-Kamiokande Detector, 2000.

[54] V. D. Barger, D. Marfatia, and K. Whisnant, Phys. Rev. Lett. **88**, 011302 (2002).

[55] P. Alivisatos *et al.*, Stanford Report No. STAN-FORD-HEP-98-03, 1998; http://www.awa.tohoku.ac.jp/KamLAND

[56] V. D. Barger, D. Marfatia, and B. P. Wood, Phys. Lett. B **498**, 53 (2001).

[57] KamLAND Collaboration, K. Eguchi *et al.*, Phys. Rev. Lett. **90**, 021802 (2003).

[58] V. Barger and D. Marfatia, Phys. Lett. B **555**, 144 (2003).

[59] G. L. Fogli, E. Lisi, A. Marrone, D. Montanino, A. Palazzo, and A. M. Rotunno, Phys. Rev. D **67**, 073002 (2003).

[60] Z. Parsa, in *Proceedings of NNN99*, AIP Conf. Proc. No. 533 (AIP, New York, 1999), p. 181.

[61] A. De Rujula, M. B. Gavela, and P. Hernandez, Nucl. Phys. **B547**, 21 (1999).

[62] V. Barger, K. Whisnant, and R. J. N. Phillips, Phys. Rev. Lett. **45**, 2084 (1980).

[63] L. Sulak, in *Proceedings of the First Workshop on Grand Unification: New England Center, University of New Hampshire, 1980*, edited by Paul H. Frampton, Sheldon L. Glashow, and Asim Yildiz (Math Science Press, Brookline, MA, 1980), p. 163; D. Ayres, T. Gaisser, A. K. Mann, and R. Shrock, in *Proceedings of the 1982 DPF Summer Study on Elementary Particles and Future Facilities, Snowmass, CO*, edited by R. Donaldson, R. Gustafson, and F. Paige (Fermilab, Batavia, IL 1982), p. 590; D. Ayres, B. Cortez, T. Gaisser, A. K. Mann, R. Shrock, and L. Sulak, Phys. Rev. D **29**, 902 (1984).

[64] P. Krastev and S. Petcov, Phys. Lett. B **205**, 84 (1988).

[65] A. J. Baltz and J. Weneser, Phys. Rev. D **37**, 3364 (1988).

[66] S. Petcov, Phys. Lett. B **434**, 321 (1998); M. Chizhov, M. Maris, and S. Petcov, hep-ph/9810501; M. Chizhov and S. Petcov, Phys. Rev. D **63**, 073003 (2001); Phys. Rev. Lett. **83**, 1096 (1999).

[67] E. Akhmedov, A. Dighe, P. Lipari, and A. Smirnov, Nucl. Phys. **B542**, 3 (1999); E. Akhmedov, Nucl. Phys. **B538**, 25 (1999); hep-ph/0001264.

[68] P. Krastev, Nuovo Cimento Soc. Ital. Fis. **103A**, 361 (1990); R. H. Bernstein and S. J. Parke, Phys. Rev. D **44**, 2069 (1991).

[69] V. Barger, S. Geer, and K. Whisnant, Phys. Rev. D **61**, 053004 (2000).

[70] I. Mocioiu and R. Shrock, Phys. Rev. D **62**, 053017 (2000); in *Proceedings of NNN99* (Ref. [60]), p. 74; J. High Energy Phys. 0111 (2001) 050.

[71] V. Barger, S. Geer, R. Raja, and K. Whisnant, Phys. Rev. D **62**, 013004 (2000).

[72] V. Barger, S. Geer, R. Raja, and K. Whisnant, Phys. Rev. D **62**, 073002 (2000).

[73] V. Barger, S. Geer, R. Raja, and K. Whisnant, Phys. Lett. B **485**, 379 (2000); Phys. Rev. D **63**, 033002 (2001).

[74] A. Cervera, A. Donini, M. B. Gavela, J. Gomez Cadenas, P. Hernandez, O. Mena, and S. Rigolin, Nucl. Phys. **B579**, 17 (2000); **B593**, 731(E) (2001).

[75] M. Freund and T. Ohlsson, Mod. Phys. Lett. A **15**, 867 (2000); T. Ohlsson and H. Snellman, J. Math. Phys. (N.Y.) **41**, 2768 (2000); Phys. Rev. D **60**, 093007 (1999); Phys. Lett. B **474**, 153 (2000); M. Freund, M. Lindner, S. Petcov, and A. Romanino, Nucl. Phys. **B578**, 27 (2000); M. Freund, P. Huber, and M. Lindner, Nucl. Phys. **B585**, 105 (2000); M. Freund, M. Lindner, and S. Petcov, Nucl. Instrum. Methods Phys. Res., Sect. A **451**, 18 (2000); K. Dick, M. Freund, P. Huber, and M. Lindner, Nucl. Phys. **B588**, 101 (2000); **B598**, 543 (2001).

[76] S. M. Bilenky, C. Giunti, and W. Grimus, Phys. Rev. D **58**, 033001 (1998); K. Dick, M. Freund, M. Lindner, and A. Romanino, Nucl. Phys. **B562**, 29 (1999); M. Tanimoto, Phys. Lett. B **462**, 115 (1999); A. Donini, M. B. Gavela, and P. Hernandez, Nucl. Phys. **B574**, 23 (2000); M. Koike and J. Sato, Phys. Rev. D **61**, 073012 (2000); **62**, 079903(E) (2000); **62**, 073006 (2000); M. Koike, T. Ota, and J. Sato, Phys. Rev. D **65**,

[77] M. Freund, Phys. Rev. D **64**, 053003 (2001); V. Barger, D. Marfatia, and K. Whisnant, Phys. Rev. D **65**, 073023 (2002).

[78] T. Yanagida, in *Proceedings of NuFACT'01, Tsukuba, Japan, 2001* (Ref. [19]).

[79] V. Barger, S. Geer, R. Raja, and K. Whisnant, Phys. Rev. D **63**, 113011 (2001).

[80] V. Barger, D. Marfatia, and K. Whisnant, hep-ph/0108090.

[81] See http://www-bd.fnal.gov/pdriver

[82] M. L. Mangano *et al.*, CERN Report No. CERN-TH/2001-131, 2001, hep-ph/0105155.

[83] E. P. Hincks and B. Pontecorvo, Phys. Rev. **73**, 257 (1948).

[84] R. L. Garwin, L. M. Lederman, and M. Weinrich, Phys. Rev. **105**, 1415 (1957).

[85] Y. Semertzidis, in *Proceedings of ICHEP02, Amsterdam, 2002*, http://www.ichep02.nl/index-new.html; for an earlier result, see H. N. Brown *et al.*, Phys. Rev. Lett. **86**, 2227 (2001), and references therein.

[86] R. Edgecock, in *Proceedings of the Workshop on Instrumentation for Muon Cooling Studies, 2000* (Illinois Institute of Technology, Chicago, 2000), http://www.iit.edu/~bcps/hep/Nov2000_mucool.html

[87] J. Aysto *et al.*, CERN Report No. CERN-TH/2001-231, 2001, hep-ph/0109217.

[88] W. Marciano, in *Proceedings of the Fermilab Workshop on the First Muon Collider and at the Front End of the Muon Collider, Batavia, IL, 1997* (Ref. [14], p. 58.

[89] James L. Popp *et al.*, hep-ex/0101017.

[90] W. Honnecker *et al.*, Phys. Rev. Lett. **76**, 200 (1996).

[91] M. L. Brooks *et al.*, Phys. Rev. Lett. **83**, 1521 (1999).

[92] M. Cooper, in *Proceedings of the Fermilab Workshop on the First Muon Collider and at the Front End of the Muon Collider, Batavia, IL, 1997* (Ref. [14], p. 443.

[93] T. Mori *et al.*, PSI Report No. R-99-05, 1999.

[94] R. M. Carey *et al.*, "AGS Letter of Intent–Search for a Permanent Muon Electric Dipole Moment," 2000; see also Y. Semertzidis *et al.*, J. Mod. Phys. A **16**, 287 (2001).

[95] J. Ellis *et al.*, hep-ph/0111324; see also K. S. Babu *et al.*, hep-ph/0006329.

[96] J. Kirkby *et al.*, PSI Report No. R-99-06, 1999; D. W. Hertzog *et al.*, PSI Report No. R-99-07, 1999; K. Nagamine *et al.*, RIKEN-RAL muon lifetime experiment, 1998.

[97] F. J. M. Farley and E. Picasso, in *Quantum Electrodynamics*, edited by T. Kinoshita (World Scientific, Singapore, 1990), p. 479; V. W. Hughes and T. Kinoshita, in *Muon Physics I*, edited by V. W. Hughes and C. S. Wu (Academic Press, New York, 1977), pp. 12–199.

[98] M. Knecht and A. Nyffeler, Phys. Rev. D **65**, 073034 (2002); M. Knecht *et al.*, Phys. Rev. Lett. **88**, 071802 (2002); M. Hayakawa and T. Kinoshita, hep-ph /0112102; J. Bijnens *et al.*, Nucl. Phys. **B626**, 410 (2002); I. Blokland, Phys. Rev. Lett. **88**, 071803 (2002).

[99] A. Czarnecki and W. Marciano, hep-ph/0102122.

[100] W. Marciano and B. L. Roberts, hep-ph/0105056.

[101] R. Bluhm, V. A. Kostelecky, and C. D. Lane, Phys. Rev. Lett. **84**, 1098 (2000).

[102] D. Kawall *et al.*, in *Proceedings of the Fermilab Workshop on the First Muon Collider and at the Front End of the Muon Collider, Batavia, IL, 1997* (Ref. [14]), p. 486.

[103] K. Jungmann, in *The Hydrogen Atom*, edited by S. Karshenboin *et al.* (Springer, Heidelberg, 2001), p. 81.

[104] W. Molzon, in *Proceedings of the Fermilab Workshop on the First Muon Collider and at the Front End of the Muon Collider, Batavia, IL, 1997* (Ref. [14]), p. 152.

[105] W. Liu *et al.*, Phys. Rev. Lett. **82**, 711 (1999).

[106] V. Meyer *et al.*, Phys. Rev. Lett. **84**, 1136 (2000).

[107] V. W. Hughes *et al.*, Phys. Rev. Lett. **5**, 63 (1960).

[108] L. Willmann *et al.*, Phys. Rev. Lett. **82**, 49 (1999); R. Abela *et al.*, Phys. Rev. Lett. **77**, 1950 (1996).

[109] P. Strasser and K. Nagamine, Hyperfine Interact. (to be published).

[110] J. Dilling *et al.*, Hyperfine Interact. **127**, 491 (2000).

[111] G. L. Borchert *et al.*, Acta Phys. Pol. B **29**, 131 (1998).

[112] CERN ISOLDE Group, R. Hayano, in Proceedings of the Workshop on Radioactive Antiprotonic and Muonic Atoms, Switzerland, 2001 (unpublished).

[113] The material on the Higgs factory physics is excerpted from V. Barger, M. S. Berger, J. F. Gunion, and T. Han, hep-ph/0110340.

[114] D. B. Cline and G. Hanson, in *Proceedings of Snowmass 2001*, http://www.slac.stanford.edu/econf/C010630/papers/M103.PDF

[115] R. Raja and A. Tollestrup, Phys. Rev. D **58**, 013005 (1998).

[116] V. Barger, M. S. Berger, J. F. Gunion, and T. Han, Phys. Rep. **286**, 1 (1997).

[117] V. Barger, M. S. Berger, J. F. Gunion, and T. Han, Phys. Rev. Lett. **75**, 1462 (1995); V. Barger, T. Han, and C. G. Zhou, Phys. Lett. B **480**, 140 (2000).

[118] M. Battaglia and K. Desch, hep-ph/0101165.

[119] J. F. Gunion, L. Poggioli, R. Van Kooten, C. Kao, and P. Rowson, hep-ph/9703330.

[120] S. Ahn *et al.*, Fermilab Report No. Fermilab-FN-677, 1999.

[121] V. Barger, M. S. Berger, J. F. Gunion, and T. Han, Phys. Rev. D **56**, 1714 (1997).

[122] M. S. Berger, in *Proceedings of the Fermilab Workshop on the First Muon Collider and at the Front End of the Muon Collider, Batavia, IL, 1997* (Ref. [14]), p. 663.

[123] M. S. Berger, in *Proceedings of the Fermilab Workshop on the First Muon Collider and at the Front End of the Muon Collider, Batavia, IL, 1997* (Ref. [14]), p. 797.

[124] B. J. King, in *Proceedings of the 1999 Particle Accelerator Conference, New York* (IEEE, Piscataway, NJ, 1999), pp. 318–320; physics/9908017; N. V. Mokhov and A. Van Ginneken, pp. 3074–3076; Fermilab Report No. Fermilab-Conf-99/067, 1999; N. V. Mokhov and A. Van Ginneken, in *Proceedings of the International Conference on Radiation Shielding (ICRS-9), Tsukuba, 1999* [J. Nucl. Sci. Technol., Suppl. 1, 172–179 (2000)].; Fermilab Report No. Fermilab-Conf-00/065, 2000; B. J.

King, in *Proceedings of the HEMC99, Montauk, NY, 1999*, edited by B. J. King, AIP Conf. Proc. No. 530 (AIP, New York, 1999), pp. 165–180.

[125] W. Chou, A. Ankenbrandt, and E. Malamud, Fermilab Report No. FERMILAB-TM-2136, 2000.

[126] N. Mokhov, Fermilab Report No. Fermilab-Conf-01/134, 2001, see [127], p. 745; http://www-mucool.fnal.gov/mcnotes/public/ps/muc0194/muc0194.ps.gz; N. Mokhov, http://www-ap.fnal.gov/MARS/; nucl-th/9812038.

[127] *Proceedings of the 2001 Particle Accelerator Conference, Chicago* (IEEE, New York, 2001), http://accelconf.web.cern.ch/AccelConf/p01/

[128] R. J. Jayakumar *et al.*, IEEE Trans. Appl. Supercond. **10**, 1 (2000).

[129] J. R. Haines and C. C. Tsai, ORNL Report No. ORNL/TM-2002/27, 2002.

[130] M. J. Burns *et al.*, in *Proceedings of the 1999 Particle Accelerator Conference, New York* (Ref. [124]), p. 617.

[131] Fermilab Proton Driver Physics Study Website is http://projects.fnal.gov/protondriver/

[132] *Proceedings of the Workshop on a Low-Energy $\overline{p}$ Storage Ring ($\overline{p}2000$)*, edited by D. M. Kaplan, H. A. Rubin, and K. Seth (Illinois Institute of Technology, Chicago, IL, 2001), http://www.iit.edu/~bcps/hep/pbar2000.html

[133] G. Hanson, in *Proceedings of NuFact01, Tsukuba, Japan* (Ref. [19]).

[134] Higgs Factory Report, edited by D. Cline and G. Hanson, submitted to Snowmass, 2001, http://www.physics.ucla.edu/snowmass/

[135] *Proceedings of the Emittance Exchange Workshop, Brookhaven National Laboratory, 2000*, http://www.cap.bnl.gov/mumu/exchange/; *Proceedings of the Workshop on Instrumentation on Muon Cooling Studies, IIT, 2000* (Ref. [86]).

[136] *Proceedings of the Workshop on Ring Coolers and Emittance Exchange, Lawrence Berkeley Laboratory, 2001*, http://www.fnal.gov/projects/muon_collider/eexchange/workshop/ee_agenda.htm

[137] V. Balbekov, S. Geer, N. Mokhov, R. Raja, and Z. Usubov, in *Proceedings of the 2001 Particle Accelerator Conference, Chicago* (Ref. [127]), p. 3867.

[138] V. Balbekov, "Solenoid Based Ring Coolers," 2002, http://www.capp.iit.edu/~capp/workshops/mumice02/mumice02_schedule.html

[139] S. Kahn *et al.*, in *Proceedings of the 2001 Particle Accelerator Conference, Chicago* (Ref. [127]), p. 3239, http://accelconf.web.cern.ch/AccelConf/p01/PAPERS/RPPH010.PDF

[140] R. Fernow, in *Proceedings of the 1999 Particle Accelerator Conference, New York* (Ref. [124]), p. 3020, http://pubweb.bnl.gov/people/fernow/

[141] The working group on the quadrupole ring cooler consists of Y. Fukui, D. Cline, A. Garren, P. He (UCLA), A. Bogacz (Jefferson Lab), D. Errede (UIUC), H. Kirk (BNL), and F. Mills (Fermilab).

[142] A. A. Garren, A. S. Kenney, E. D. Courant, A. D. Russel, and M. J. Syphers, Report No. SSCL-MAN-0030, 1993.

[143] R. Palmer, "Muon Cooling Rings and Kickers," 2002, http://pubweb.bnl.gov/people/palmer/exch/kicker2.pdf

[144] D. Neuffer, Mucool Report No. 211, http://www-mucool.fnal.gov/mcnotes/public/pdf/muc0211/muc0211.pdf, 2001.

[145] *Colliders and Collider Physics at the Highest Energies: Muon Colliders at 10 TeV to 100 TeV: HEMC'99 Workshop*, edited by Bruce J. King, AIP Conf. Proc. No. 530 (AIP, New York, 1999); *Muon Colliders at the High Energy Frontier, Proceedings of the 6-Month Feasibility Study*, edited by B. J. King and A. Caldwell (Rinton Press, Princeton, NJ, 2001).

[146] J. Norem *et al.*, Mucool Note No. 226, 2001, http://www-mucool.fnal.gov / mcnotes / public / pdf / muc0226/muc0226.pdf; J. Norem, V. Wu, A. Moretti, M. Popovic, Z. Qian, L. Ducas, Y. Torun, and N. Solomey, Phys. Rev. ST Accel. Beams **6**, 072001 (2003).

[147] Muons, Inc. is a private corporation headed by Rolland Johnson that has received the DoE STTR grant to carry out this research, http://www-mucool.fnal.gov/mcnotes/public/ps/muc0247/muc0247.ps.gz

[148] Chun-xi Wang and Kwang-Je Kim, Phys. Rev. Lett. **88**, 184801 (2002); Kwang-Je Kim and Chun-xi Wang, Phys. Rev. Lett. **85**, 760 (2000); G. Penn and J. S. Wurtele, Phys. Rev. Lett. **85**, 764 (2000); V. M. Malkin and N. J. Fisch, Phys. Rev. Lett. **85**, 5575 (2000).

[149] H. Edwards *et al.*, Muon Technical Advisory Committee, "MUTAC Committee Report on the Muon Collaboration Review, 2001."

[150] M. Aleska *et al.*, CERN Report No. CERN-Nufact-Note-108, http://molat.home.cern.ch/molat/neutrino/nf108.pdf, 2002.

[151] L. Cremaldi and D. Summers, Mucool Note No. 221, 2001, http://www-mucool.fnal.gov/mcnotes/public/pdf/muc0221/muc0221.pdf; D. Bartlett *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **260**, 55 (1987).

[152] A. Blondel, contribution to the CERN Neutrino Factory Working Group, http://nfwg.home.cern.ch/nfwg/nufactwg/38/blondel.PDF

[153] K. Hanke, in *Proceedings of the Workshop on an International Muon Ionization Cooling Experiment, CERN, 2001*, http://muonstoragerings.cern.ch/October01WS/oct01ws.html