The MAP Design & Simulation Program

10 August 2010

Abstract

This document (i) describes the design and simulation work requirements for a neutrino factory and a muon collider facility, (ii) compares these requirements with the present state of the art, (iii) identifies the main R&D issues, and (iv) describes the MAP R&D program we are proposing to address these issues, and its prospects for success.

Introduction

Design and simulations (D&S) play a major role in the proposed 7-year Muon Acceleration Program (MAP). Two of the chief goals of MAP are to produce a Design Feasibility Study for a muon collider and a Reference Design Report for a neutrino factory. These studies will include designs for all the important systems in these facilities and simulations of the expected machine performance. For MAP, the "design" activity discussed here includes system configurations, optics designs of machine lattices, and conceptual design work on required machine hardware, including magnets and RF cavities. "Simulations" include single-particle tracking through all the facility systems, studies of multi-particle effects, such as space charge and beam-beam interactions, and studies of critical processes that could possibly jeopardize the success of a muon-based facility, such as the mechanisms responsible for RF breakdown in magnetic fields.

In this document we will briefly discuss the current status of the muon facility design and simulation efforts and explain why we feel certain additional activities need to be undertaken as part of MAP. For each of the major (Level 2) D&S areas we will emphasize several of the most challenging proposed activities, which stand out either because they have never been part of any previous accelerator design, or because they may be difficult to achieve due to the parameter range in which we intend to use them.

There has been continuous work on the possibility of building a muon collider by MAP and its predecessor organizations since 1994 and studies on the possibility of building neutrino factories since 1999. Over that time, the machine designs have improved enormously, evolving from collections of simple concepts into integrated designs, many parts of which have been checked using realistic simulations. We feel that the design of neutrino factories is now quite advanced. Two detailed studies sponsored by Fermilab (Study 1) and by Brookhaven (Study 2) examined all aspects of the facility design [1,2]. This was followed by extensive simulation work as part of the APS Joint Study on the Future of Neutrino Physics, which led to a new neutrino factory design with improved performance and reduced cost (Study 2a) [3]. This was followed by participation in the International Scoping Study (ISS) [4], which compared the advantages and disadvantages of neutrino factory proposals from the US, Europe, and Japan. We are presently working with our collaborators on the International Design Study for a Neutrino Factory (IDS-NF) [5]. Important goals of this study are to improve the level of engineering detail in the design and to obtain a better cost estimate for the facility. This long history of successful studies gives us high confidence that the Reference Design Report work in MAP will be similarly successful.

The technical challenges for the muon collider design are much greater than those for the neutrino factory. There has been one previous study [6,7] that seriously looked at the design of a muon collider. However this study did not contain nearly the same level of detail as the subsequent neutrino factory studies. Fortunately, a large number of promising new ideas have been developed over the last decade, for example in the area of muon cooling. In addition, we have greatly improved the capabilities of our simulation tools. The development of the G4beamline code [8] by Muons Inc. was particularly noteworthy in this regard, due to its excellent graphical capabilities and its extensive treatment of particle interactions with matter based on the Geant4 code system. We believe that we have now identified most of the major issues for designing the collider and that we have possible solutions for addressing these issues. This gives us confidence that incorporating some combination of these new ideas into an integrated design will allow us to reach our goal for a high-luminosity collider. However, even if it turns out that we ultimately fail to reach all of our collider goals, the MAP work on the collider Design Feasibility Study should result in a clear answer to the question of whether muon colliders represent a viable path for future particle physics research.

At this time, a baseline configuration for the neutrino factory has already been chosen by the IDS-NF collaboration. For the muon collider there are still several possible technically-feasible choices for many of the machine subsystems. For this reason, a process of *down-selection* will be an important part of the work in MAP. We have established a series of milestones that will lead to a single collider machine configuration in FY13. A formal procedure for carrying out the process of making these choices is described in section 3.1.4 of the MAP proposal. In each case, a technical review will take place and advice will be solicited from outside reviewers. Among other things, the choice for each system will take into account the simulated performance, engineering feasibility and relative costs.

1. Proton Driver

Introduction and present status

A high-intensity proton driver is needed to produce the muons for a neutrino factory or muon collider. We assume that the proton driver will be based on the Project X driver that is currently under study at Fermilab. MAP design and simulation activities will concentrate on enhancements to the Project X baseline required by the muon program. The primary requirements for MAP include proton energy ~8 GeV, beam intensity ~4 MW, proton pulse width at the target ~2 ns, and the flexibility to adjust the pulse spacing and repetition rate.

It is obvious that the requirements on beam power, bunch length, and repetition rate, taken together, imply bunch intensities and peak bunch currents that will be difficult to achieve. Meeting those requirements, while also providing flexibility in the number and pattern of bunches delivered to the production target, will represent one of the important design activities for MAP. The current idea for controlling the temporal properties of the proton beam is to use an accumulator ring and a compressor ring after the H^- linac.

The Proton Driver Group has been closely following developments on Project X. Some preliminary work has been done with the Fermilab FESS group on possible site layouts for the facility. A preliminary study has also been completed on an upgrade path to make the beam from Project X meet the requirements for our muon facilities [9,10].

R&D issues and proposed R&D plan

i) Injection into the accumulator ring

The requirements for ~4 MW of proton beam power, together with space-charge considerations, lead to a design where the linac is followed by a small accumulator ring. This ring must be able to accept many turns of linac output at a rate of $\sim 4 \times 10^{15}$ protons/s and must produce the bunch structure needed by the muon facilities. An important issue that needs to be simulated is how to inject and capture this large number of protons in a small ring. Current ideas envision accumulating many turns of linac beam via charge-stripping of the H^- beam.

One issue that needs to be modeled is the feasibility of the foil stripping technique that is used. The incoming beam from the linac will be chopped to allow clean injection into pre-existing RF buckets and form the desired number of bunches. Painting will be necessary in transverse phase space and possibly also in longitudinal phase space. Very large transverse emittances must be prepared in order to control space-charge forces.

Another important issue will be preventing the foils from overheating. Several ideas have been proposed to accomplish this. The *resonant foil bypass* [11] would use AC dipoles to modulate the beam orbit at the foils at the same frequency as the incoming linac bunches. The foils could be separated longitudinally to provide more surface area for cooling. A third idea is to rapidly rotate a circular foil so that only an annular region in a corner of the aperture is exposed to the beam. The efficacy of the various alternatives will be assessed to identify the most promising approach.

ii) Producing a 2 ns proton bunch at the target

ICOOL simulations have shown that the performance of the phase rotation system in the front end is strongly affected by the bunch length of the pions produced at the target. This is, in turn, determined by the pulse width of the protons hitting the target. For good performance, MAP needs an rms proton pulse width ~2 ns. This is a challenging goal for an 8 GeV, high intensity beam. In addition, the neutrino factory and muon collider have different requirements for the muon pulse spacing and repetition rate. These quantities are determined by the proton pulse spacing and repetition rate at the target. It is highly desirable that the same proton driver complex be usable for either facility. The current idea for accomplishing this is to follow the accumulator ring with a proton bunch compressor ring and an external bunch combiner channel.

The compressor ring will be used to accept one or more bunches at a time from the accumulator. Then, a 90° bunch rotation in longitudinal phase space will be performed to shorten the bunches just prior to extraction. Of course, during this operation, the momentum spread will become large, of order 5%, so the ring must have a large momentum acceptance. Also, the space-charge tune shift will be large when the bunch length is short.

Space-charge issues in the compressor ring may preclude producing single 2 ns pulses at the target. In that case, there are ideas that need to be studied for using a beam *trombone and funnel channel* for the muon collider, as indicated schematically in Fig. 1.1. This allows different bunches to follow paths with different lengths and arrive simultaneously at the target from multiple directions.

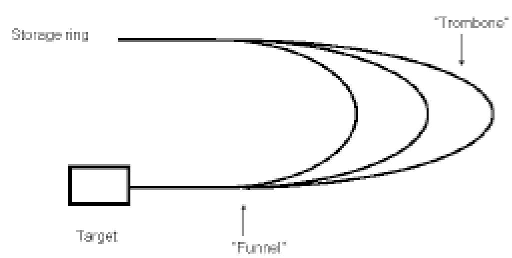

Figure 1.1 Schematic of the trombone and funnel concept.

iii) Increasing the power of the Project X beam

In order to obtain the required neutrino flux or collider luminosity, the MAP program may need proton intensities ~4 MW. Upgrades from the Project X baseline 8 GeV intensity of < 500 kW will be studied as part of the MAP activities. This would most likely be done by increasing the beam current, pulse duration, and repetition rate of the linac.

Down-selection

The designs of the accumulator and compressor rings for the proton driver are still at an early stage. Plans for MAP call for simulating these systems in greater detail and choosing a single baseline proton driver design in FY13 based on the down-selection process mentioned in the Introduction.

2. Front End

Introduction and present status

MAP defines the "front end" to encompass all the systems from the target to the ionization cooling channel used by the neutrino factory. Cooling used exclusively by the muon collider is considered separately in the following section. The front-end work is divided mainly between the target systems and the front-end channel. The MAP baseline target is a liquid-mercury jet inside a 20 T solenoid field as shown in Fig. 2.1. The target systems also include the tapered solenoid channel that extends ~15 m downstream from the target, the beam dump, shielding, and all the associated plumbing and infrastructure required by a high-power target hall [12]. Many of the target concepts have been validated by the results of the MERIT experiment [13], which studied the interaction of a proton beam with a mercury jet inside a magnetic field. The front-end channel is responsible for bunching the muon beam at a frequency of 201 MHz, reducing the energy spread of the beam to ~10%, and providing some initial cooling. The figure-of-merit for the system is the number of muons that can be captured inside the acceptance of the downstream acceleration system.

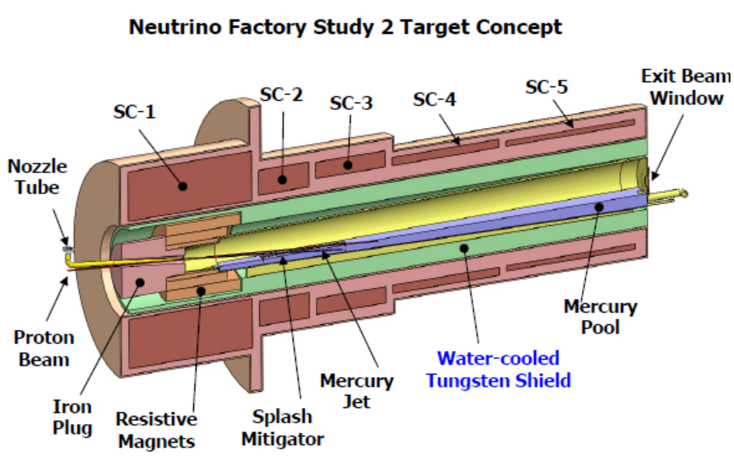

Figure 2.1 Target systems concept showing the mercury jet, superconducting solenoids, magnet shielding, and mercury pool.

The front-end and cooling systems represent some of the greatest departures from conventional accelerator designs. This follows from the difficulties of producing and capturing large numbers of muons and then transforming them into bunches suitable for injection into an accelerator. In the past decade, R&D in this area has focused on exploring a number of ideas for producing cooled muon bunches as efficiently as possible. The emphasis in MAP will be to explore the surviving ideas in more detail and to select the most promising schemes as the baselines for the neutrino factory and the muon collider. The baseline for the neutrino factory front end will be decided together with our partners in the IDS-NF collaboration. The design of the front end for the muon collider is coupled with the choice of the baseline cooling design. It would be desirable for facility-staging scenarios if the neutrino factory front end could also be used for the muon collider.

The only cooling method that can be effectively applied within the muon lifetime is ionization cooling [14]. This novel technique involves focusing the muon beam through material absorbers and then restoring the longitudinal momentum with RF cavities. A test of this idea is currently under way at the MICE experiment at Rutherford Appleton Laboratory.

Two grand schemes have been proposed for the design of the front end and the cooling systems for the muon collider. It is convenient to identify these schemes by the final transverse emittance they aim to produce, although they differ in many other parameters as well: the Low Emittance Muon Collider (LEMC), which aims to produce a normalized transverse emittance of 2 μm, and the High Emittance Muon Collider (HEMC) [15], which aims to produce a normalized transverse emittance of 25 μm. Some important properties for these scenarios are compared in Table 2.1.

Recent optimization work [16] has led to a front-end channel with a shorter version of the decay, bunching and phase rotation systems than used in the Study 2a design, as shown in Fig. 2.2. These preliminary designs show improved performance over the Study 2a design.

Table 2.1 Current 1.5 TeV Muon Collider Scenarios.

	LEMC	HEMC
Avg. Luminosity (10^{34} cm^{-2} s^{-1})	~2.1	~1.1
Avg. bending field (T)	10	8
Proton driver rep. rate (Hz)	65	15
β^* (cm)	0.5	1
Muons per bunch (10^{11})	1	20
Muon bunches in collider	10	1
Norm. Transv. Emittance (μm)	2.1	25
Norm. Long. Emittance (m)	0.35	0.07
Energy spread (%)	1	0.1
Muon survival (%)	31	7

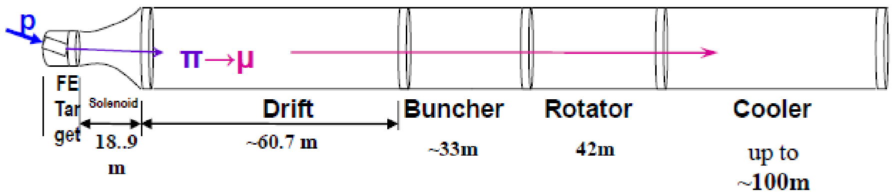

Figure 2.2 Layout of a possible front end channel.

Experimental measurements have shown that the maximum gradient in the RF cavities falls off with increasing strength of the magnetic field. There are two variants of the HEMC solution that address this problem. In the "hybrid" approach [17] the channel is filled with moderate-pressure hydrogen gas, just sufficient to allow operating the RF cavities at their design gradients, but most of the cooling takes place in conventional absorbers. In the other variant, the lattice is modified to use magnetic insulation. The LEMC scenario uses a quasi-isochronous helical cooling channel [18], which employs high-pressure hydrogen gas.

The final front-end system is the cooling for the neutrino factory. The baseline used in the International Scoping Study was the Study 2a transverse cooling channel. The dependence of the accepted number of muons for the IDS-NF baseline channel is shown as a function of the RF gradient in the cooling channel in Fig. 2.3. The channel acceptance falls off by 22% if the gradient is reduced from the design value of 15 MV/m to 10 MV/m. Variations of this channel that are being simulated and that address the RF problem also include filling the cavities with moderate-pressure hydrogen gas or using a magnetically-insulated lattice. One problem with these transverse cooling channels is the loss of muons that fall out of the RF bucket. This could be reduced by using a 6D cooling channel instead, such as the helical FOFO snake channel, discussed in the following section.

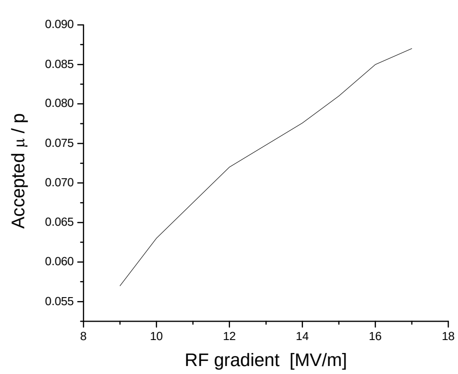

Figure 2.3 Accepted muons per incident proton as a function of the gradient in the cooling channel for the IDS-NF baseline front end channel. (Note the suppressed zero on both axes.)

R&D issues and proposed R&D plan

i) Simulation of the mercury jet

An important simulation task for the front end is understanding the interaction of the proton beam with the target. We have chosen a liquid-mercury jet as the target in order to accommodate the very high proton beam power. The MERIT experiment at CERN represents the best demonstration of the behavior of a mercury jet in the simultaneous presence of a proton beam and a magnetic field. It is important that simulation efforts continue trying to model the results of that experiment. Since the target is a conducting liquid, it is necessary to understand distortions in the shape of the jet as it traverses the magnetic field and the possibility of cavitation voids forming inside the jet. Simulations on the shape of the jet nozzle could result in an improved design that produces a more laminar jet. Simulations also need to be made of the collection of the spent mercury from the jet and stopping the non-interacted proton beam in the mercury pool.

ii) Shielding the magnets near the target

An enormous flux of neutrons and charged particles is produced in the mercury target. This particle flux presents a challenging problem for shielding the high field superconducting magnets closest to the target. The deposited energy results in heat loading in the 4 K cryogenic system, could result in quenching the Nb_3Sn superconductor, and likely results in lifetime issues in materials from radiation damage. Detailed studies need to be made, optimizing the type and location of the shielding around these magnets. This may require redesigning the coil configuration that creates the tapered field downstream from the target. We also plan to explore whether new developments in superconductor materials might allow an improved design of the 20 T solenoid that surrounds the target.

iii) Target facility design

A realistic target facility involves a number of auxiliary subsystems that need to be designed. These include a closed mercury handling loop that takes the used mercury and pumps it back into the jet nozzle, cryostats for the magnets, beam windows, and a beam dump. We also need a design for a robotic system for remote handling of repairs or replacements of target area components.

iv) Pion production characteristics

The neutrino fluxes at a neutrino factory and the luminosity at a muon collider are both directly related to the pion production rate at the target. It is thus very important to have reliable simulations of the pion production distributions and to optimize the capture of the resulting pion flux. Recent studies with MARS15 have resulted in a significant increase in the number of accepted muons. It is important to compare these results with other production codes and to validate the predictions of these codes with measured results from the HARP and MIPP experiments. Studies on pion production and collection schemes have a direct influence on the required proton driver power. Simulations with particle production and tracking codes need to clarify the optimal values for the beam energy, jet dimensions, and intersection angles. Further improvements in collection may be possible by modifying the solenoid field taper downstream from the target.

v) Design of bunching and phase rotation systems

The bunching and phase rotation systems in the front end are responsible for forming the collected muons into 201 MHz bunches suitable for subsequent cooling and acceleration, and for reducing the energy spread in the muon beam. Our current designs need to be simulated in greater detail, including any required matching sections.

vi) Simulation of RF breakdown

An important issue in the design of the front end and the cooling channels is the behavior of normal-conducting RF cavities in strong magnetic fields. Initial measurements have shown that the cavity gradient falls off rapidly with increasing field. The necessity of looking at many alternative designs for the front end and cooling channels comes from the fact that we do not yet have definitive experimental results, and therefore must simulate the expected behavior of a number of proposed solutions to the problem. These proposed solutions fall into three categories:

1) If the problem can be fixed using cavities made from beryllium or utilizing atomic layer deposition, then we can continue to use vacuum-filled pillbox RF cavities with thin beryllium windows covering the irises.
2) It is possible that the problem can be fixed by using hydrogen gas in the RF cavity. High-gradient channels filled with high-pressure gas is the solution assumed in the LEMC scenario. The hybrid approach uses a lower pressure to run the cavities at a more moderate gradient. For these scenarios it is important to understand the effects of high intensity beams on gas-filled cavities.
3) It may be possible to solve the problem by using open-iris vacuum-filled RF cavities in specially designed magnetic lattices. This is the basis for studies of magnetically insulated and bucked-coil lattices. Simulations of breakdown mechanisms [19,20] should help us interpret the results of the MTA RF experiments as they become available.

Down-selection

The design of the front end system still has a number of options. The baseline configuration for the target system will be decided following a technical review in FY10. The choice of the front end channel will be dependent on the results of the experimental program on RF cavities now under way at the MTA. A single baseline design for the front end channel will be chosen in FY11 following a technical review.

3. Cooling

Introduction and present status

Cooling the 6-dimensional emittance of the muon beam by a factor of ~10^6 is among the most challenging requirements for a muon collider. Ionization cooling has never been used in an accelerator design before, so there are no examples on which to base a design. After many years of investigation, we have identified several promising approaches to addressing the cooling problem. Simulating these remaining approaches in much more detail and selecting a single baseline system for the collider are major goals for MAP.

There are two general approaches for accomplishing the required cooling. These scenarios involve complicated arrangements of 6D cooling channels, final transverse cooling channels, charge separation and recombination systems, bunch train merging systems and many matching sections. In addition, several specialized cooling techniques may be useful for replacing parts of the main cooling scenarios.

The goal of the HEMC scenario [21] is to achieve a normalized transverse emittance of 25 μm. It has the advantage that preliminary simulation work has shown that this may indeed be possible. It has the disadvantages that it requires 2×10^{12} muons in each bunch, which may cause problems with RF cavity loading, and that optimal performance requires the development of 40–50 T solenoids for the final cooling. The current version of the scenario is outlined in Fig. 3.1. This figure shows the various stages of cooling on a plot of normalized transverse versus longitudinal emittance. The beam from the front end starts on the upper right and the requirement for the collider is on the upper left.

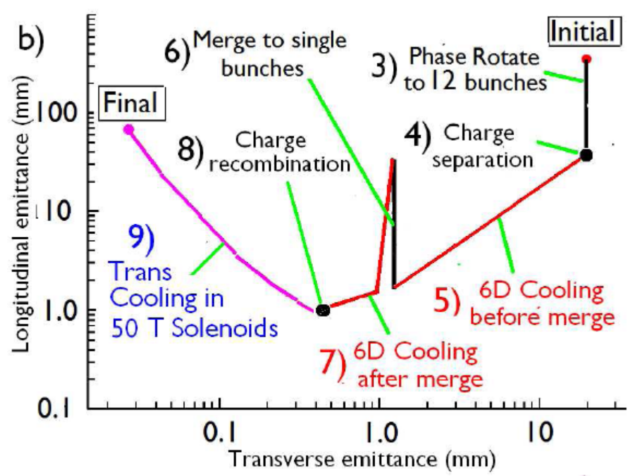

Figure 3.1 Stages of cooling in the HEMC scenario.

The HEMC scenario could begin with a Helical FOFO Snake channel [22], which can simultaneously transport both muon charges. The partially cooled beam is then separated by charge and sent to tapered Guggenheim [23] cooling channels. The bunch train from the front end is merged to a single bunch midway through the Guggenheim channel in this scenario. Preliminary simulations have been made of cooling in Guggenheim channels, charge separation and bunch merging. The Guggenheim channel is followed by a charge recombination system. The beam at this point should have a longitudinal emittance smaller than is required by the collider. This allows a final transverse cooling channel to reduce the transverse emittance while allowing the longitudinal emittance to grow. Preliminary simulations have indicated that a 50 T channel should meet the requirements for the HEMC scenario [24]. Studies are under way to learn the consequences for performance if fields lower than 50 T prove to be the practical limit. At the end of the cooling channel a collimation system is used to remove the tails of the muon beam before starting acceleration.

The goal of the LEMC scenario [25] is to achieve a normalized transverse emittance of 2 μm. This would allow using fewer muons per bunch, which would reduce RF loading issues in the accelerator and muon decay background issues in the detector. The initial 6D cooling is done by sending charge-separated muon beams into tapered Helical Cooling Channels (HCC) [26]. Preliminary simulations have shown good 6D cooling in the HCC channel. The LEMC scenario proposes reaching its final emittance goal by following the HCC with Epicyclic Parametric Ionization Cooling (EPIC) [27] and Reverse Emittance Exchange (REMEX) channels. Cooling down to 2 μm emittances has not been demonstrated yet with realistic simulations. The charges of the cooled beams must be recombined before acceleration. The LEMC scenario proposes partially merging the bunch train from the front end in a high energy ring.

There are several other cooling techniques that might substitute for part of these scenarios if they could reach the required emittances and if there is some advantage in doing so. The most important among these is cooling using lithium lenses.

The operation of RF cavities in strong magnetic fields is also an issue for the cooling channels. Apart from the HCC and the hybrid cooling approach, the other cooling channels operate with vacuum RF cavities. Possible solutions to the breakdown problem discussed in the previous section also need to be investigated in this context. Preliminary simulations of one of the Guggenheim stages have shown that magnetically insulated lattices give similar performance to the standard channel.

R&D issues and proposed R&D plan

i) Cooling the 6D emittance of the muon beam by a factor ~10^6

One of the primary requirements for the muon collider is to rapidly cool both the transverse and longitudinal emittances of the captured muon beam. It is the task of the 6D cooling channels to accept the relatively large emittance bunch trains from the front end and to produce low emittance bunches that can be accepted by the final cooling channels. The beam at the end of the 6D channel should have a longitudinal emittance smaller than is required by the collider. As we have seen, there are two major scenarios for how to do the 6D cooling, along with a number of variants, chiefly concerned with possible solutions to the RF cavity problem. Presently we are investigating three 6D channels that could play some role in the collider design, as shown in Fig. 3.2. A goal for MAP is to make detailed simulations and enough engineering studies of these channels that an informed choice can be made of which channel(s) to use for the collider study.

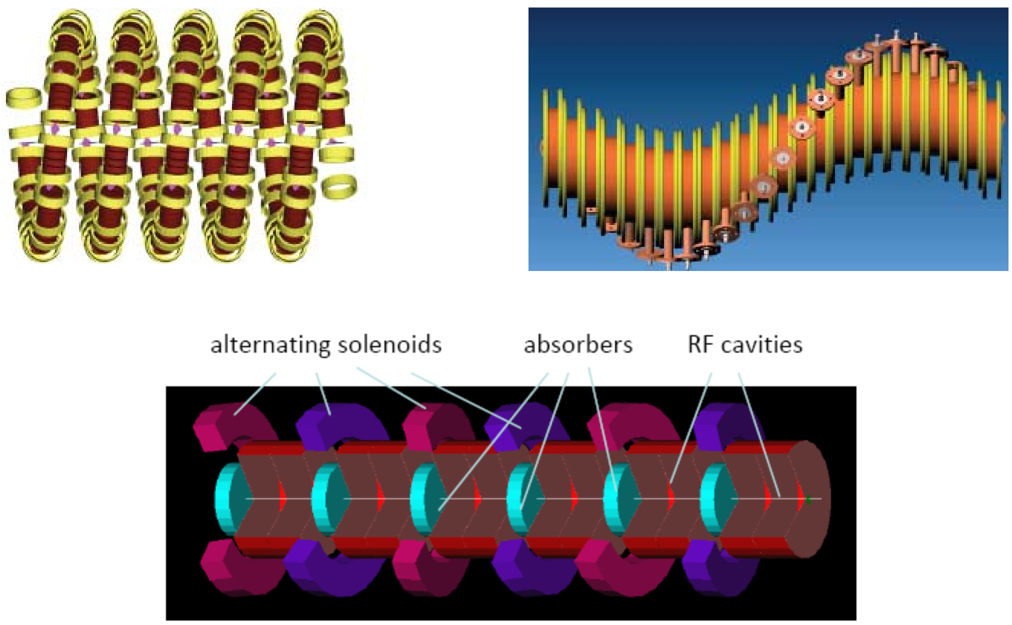

Figure 3.2 6D cooling channels: Guggenheim (upper left), HCC (upper right), Helical FOFO snake (bottom).

The current version of the HEMC scenario begins with a Helical FOFO Snake channel. This 6D cooling channel has the advantage that it can simultaneously transport both signs of muon charges. However, preliminary simulations indicate that this channel cannot efficiently reach the required emittance for the end of the 6D channel, so we expect it to be used only for the initial cooling (e.g., to make charge separation easier and more efficient).

The Guggenheim cooling channel has the advantage that it uses a large-pitch helix, so that inserting RF cavities in the channel is straightforward. The channel uses periodic solenoid focusing, which allows absorber material to be placed only at low beta function locations and enables the channel to achieve low final emittances. It has a lower initial cooling rate than the HCC channel, so that it involves a longer channel with more muon decays, and it is susceptible to the problem of operating vacuum RF cavities in strong magnetic fields. A Guggenheim channel only works for one sign of the muon charge, so the beam must go through a charge separation system before entering this channel.

The HCC channel is filled with high-pressure hydrogen gas. The beta function is constant and muon interactions with material occur continuously. The advantage of this technique is that it allows the RF cavities to run at a higher gradient in a magnetic field and gives a fast initial rate of emittance reduction, which has been confirmed in simulations. There are several potential difficulties with this approach. We do not know yet if the interactions of the intense muon beam with the gas will drain the stored energy from the RF cavities. There are engineering issues with getting the RF cavities inside the helical channel, particularly in the final stages. An HCC channel only works for one sign of the muon charge, so the beam must go through a charge separation system before entering this channel.

ii) Cooling the final normalized transverse emittance to 2–25 μm

Another difficult requirement for the muon collider is to reduce the transverse emittance of the muon beam to the level required by the collider ring. The two collider scenarios have adopted very different solutions for achieving this final cooling, as shown in Fig. 3.3. The HEMC scenario uses a 40–50 T cooling channel, while the LEMC scenario uses the EPIC/REMEX channels. The other alternative possibility, still under investigation, is to use high-field lithium lenses. A goal for MAP will be to make detailed simulations and enough engineering studies of these channels so that an informed choice can be made of which channel to choose for the collider study.

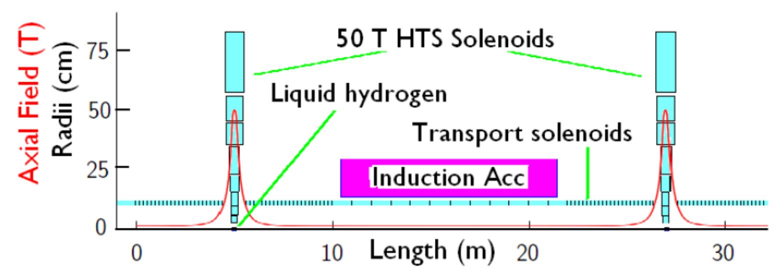

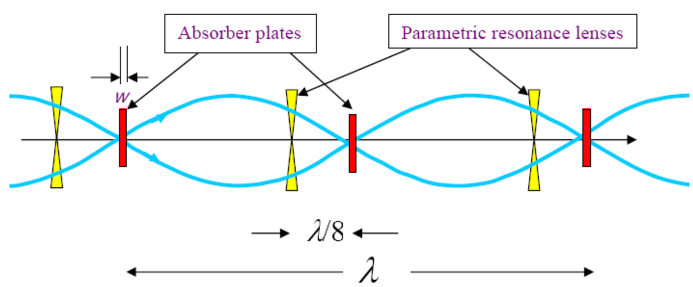

Figure 3.3 Final transverse emittance cooling channels: an example stage for the 50 T channel (upper), PIC channel (lower).

The current scenario for the final cooling channel uses 19 stages, starting with a 35 T field and ending at 50 T. The energy of the muon beam is reduced as it proceeds down the channel. The bunch length increases in this channel, so reacceleration in the final stages may require using induction linacs. The main advantage of a 50 T channel is that it has been shown to reach the required emittance values in preliminary simulations. The main disadvantage is that meeting the optimal luminosity requirements for the collider requires the development of 50 T solenoid magnets. The EPIC/REMEX channels would have the advantage of reaching a smaller final transverse emittance if they can be shown to work in realistic simulations. Simulations have shown that lithium lenses can reach low emittances, but not yet those required by the collider. This last option also has the disadvantage that it would require extensive hardware R&D that is currently not being supported.

iii) Design of a charge separation system

The front end produces a train of interleaved positive and negative muon bunches. The beam must be separated by charge before it enters the HCC or Guggenheim cooling channels. Preliminary simulations have indicated that this could be done using large-aperture bent solenoids. These simulations need to be continued in more detail.

iv) Design of a bunch recombination system

We need to combine the muon bunches in the train from the front end into a single bunch of each charge for the collider ring. It is desirable to do the bunch recombination after the beam has undergone some initial cooling. Preliminary simulations have indicated that this could be done using a combination of longitudinal and transverse stacking. Low energy bunch merging systems have been developed using planar wigglers or helical channels. Preliminary studies have also been done of a high energy bunch merging ring. These simulations need to be continued in more detail and a single approach chosen for the collider study.

v) Simulation code development

Simulating the behavior of the muon beam in a cooling channel has required the development of new simulation codes, such as ICOOL and G4beamline, with a mixture of features from accelerator and particle interaction physics. We have found that having two codes available for simulating cooling channel performance has been extremely useful for detecting mistakes and for estimating uncertainties. As an example, a recent comparison of the transverse emittance in the front end channel is shown in Fig. 3.4. The agreement between the codes for complicated channels is typically ~5-10%. These codes need to be continually updated as more of the channel is modeled and as ideas for cooling methods evolve. Eventually the complete baseline cooling channel needs to be simulated in a high statistics end-to-end simulation.

Down-selection

The design of the cooling system still has a number of options. The choice of the 6D and final cooling channels will be dependent on the results of the experimental program on RF cavities now under way at the MTA. This, together with a technical review of the performance and engineering feasibility of the various options, will allow a single baseline design for the cooling channel to be chosen in FY12.

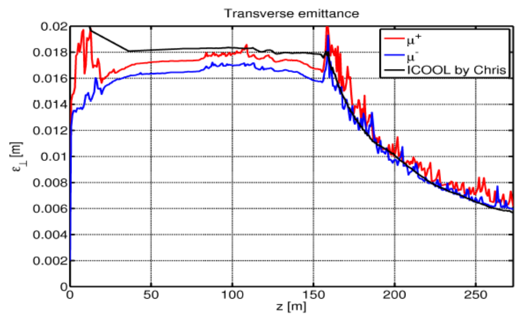

Figure 3.4 Transverse emittanceas a function of distance down the front-end channel. The black curve is for μ^+ from ICOOL. The red curve is for μ^+ from G4beamline.

4. Muon Acceleration

Introduction and present status

Following cooling, the muon bunches need to be accelerated to the final energy required by the facility, 25 GeV for the neutrino factory and 750 GeV for the muon collider. Since the acceleration represents a large fraction of the cost the facilities, it is important that this be done as efficiently as possible, particularly for the collider. Because of the muon lifetime, the acceleration needs to be done quickly. The cooled muon bunches for the neutrino factory still have large normalized rms emittance (~7 mm) by normal accelerator standards, so these acceleration systems must have large acceptance.

For the neutrino factory, the proposed solution adopted by the IDS-NF study [28] is to use a straight linac up to 0.9 GeV, a pair of Recirculating Linear Accelerators (RLAs) that accelerate the beam to 3.6 and 12.6 GeV respectively, and a Fixed Field Alternating Gradient (FFAG) accelerator up to 25 GeV, as shown in Fig. 4.1. Preliminary lattice designs for the linacs and RLA arcs have been made for these systems and some tracking has been done. Simulations have been done for some matching sections and for the chromaticity compensation. The initial straight linac was chosen because of the need for large aperture and tight focusing in both planes. The RLAs were selected because of their acceptance and the economy of reusing the RF cavities in the straight section. A "dog-bone" was chosen over a racetrack configuration to allow better orbit separation at the linac ends and to allow the beams to traverse the linac in both directions. A non-scaling FFAG [29] was selected for the final acceleration because early costing studies indicated it was less expensive than a third RLA.

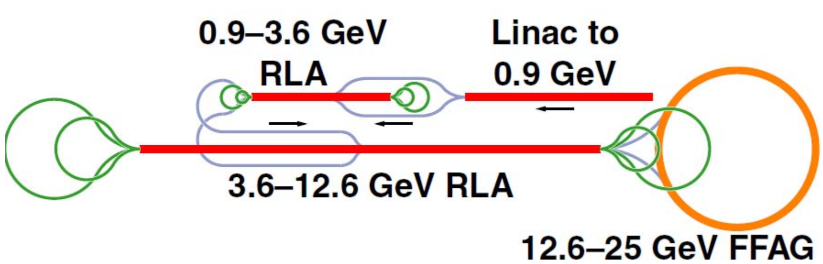

Figure 4.1 Layout of acceleration system for the IDS-NF neutrino factory.

There are several options for the accelerator system for the collider. One possibility is to use the 25 GeV neutrino factory accelerator as the first stage of acceleration for the collider. The acceleration to higher energy could use an additional set of RLAs [30] or could use a series of Rapid Cycling Synchrotrons (RCS) [31]. An RLA is potentially a straightforward option for accelerating to high energies. Its primary disadvantage is the practical limitation on the number of passes the beam can make through the linac due to the complexity of the switchyard for the return arcs. The RCS uses a combination of fixed and varying field dipoles whose polarities oppose each other at injection and act in unison at high energy. Preliminary design studies have been made for both of these options. A number of alternative scenarios have also been considered. One is to use an RLA with fast-ramping magnets, allowing for a greater number of passes. Using FFAGs is another possibility. This choice is potentially advantageous, since FFAGs generally become more efficient at higher energy. Preliminary studies have also been made of using an FFAG arc lattice in an RLA in order to minimize the number of return arcs.

R&D issues and proposed R&D plan

Complete acceleration systems need to be designed as part of MAP and the expected beam behavior needs to be simulated in detail. For each of the potential acceleration systems, design work is needed on refining the lattice optics. Single-particle tracking simulations need to be done for the whole accelerator chain, which will probably require some additional code development. Dynamic aperture and resonance crossing need to be further investigated. Conceptual designs need to be made of required magnets and RF cavities. Auxiliary systems need to be designed for injection, extraction, and orbit corrections. Matching sections and transfer lines need to be designed between the stages.

i) Design of a low energy accelerator

For a staged approach it must be determined if the 25 GeV neutrino factory accelerator design is suitable for the first stage of acceleration for the collider. If we decide not to make use of FFAGs, an alternative system using three RLAs would need to be investigated. Beam spreader and recombiner systems need to be designed for all the RLAs. Simulation studies will be used to indicate which alternative has the best performance and has the potential to be most cost effective.

The acceleration system for the neutrino factory makes use of 201 MHz superconducting cavities. We need to quantify the impact of fringe fields on cavity operation, since it has a big impact on component spacing, and hence acceleration system costs.

ii) Design of a high energy accelerator

Accelerating the muon beams to high energy is likely to be a cost driver for the muon collider. Part of the MAP studies will be to complete the designs for the RLA and RCS acceleration systems in sufficient detail so that a single baseline system can be selected for the collider study. Scaling up the RLA design to 750 GeV should be straightforward in principle, although a lot of detailed work will be necessary. A comparison must be made of the advantages of using a racetrack versus a dog-bone configuration for this accelerator.

The advantage of an RCS for acceleration is that it allows a larger number of passes through the RF cavities, reducing both the capital and operating costs of the machine. For this reason, the RCS is our choice for an initial design configuration. The current design for the lattice cell uses a combination of fixed superconducting 8 T dipoles together with ±1.8 T normal-conducting iron dipoles. As a consequence, the particle orbit moves across the aperture during the acceleration cycle. There are a number of challenges that must be addressed before deciding if this is a practical solution. The rapid variation of the magnetic fields and the short ramping time requires magnets with very thin laminations in order to manage eddy currents. The influence of the ramped fields on the 8 T superconducting magnets and the possibility of quenching needs to be examined. The changing beam orbits must be kept synchronized with the RF cavities. Impedance issues with the fast-ramping iron magnets and the choice of vacuum chamber material need to be examined.

iii) Effects of 2×10^{12} muons in a bunch

Because the intensity of a coalesced bunch for the collider will be quite high, collective effects constitute a potential operational limitation. There are several such effects to consider, and these must be simulated to assess their impact on performance. Although the muon beam spends only a short time in the accelerator complex, its individual bunches have a substantial charge, and impedance-driven collective effects may well be important. For acceleration, the major contribution to the impedance will be the RF cavities. For the collider application, the charge in a single bunch is large enough to extract a substantial fraction (~8% per bunch passage) of the stored energy from one of these cavities. As this is a nonstandard operating regime, we must study its beam dynamics implications in detail. We will study the effect of short-range wakes, probably the most important effect, as well as long-range wakes. We will also consider the effects of having both signs of muons in the various acceleration stages simultaneously, as most acceleration scenarios envisage this. The bunches will thus collide parasitically many times during acceleration. The large bunch charge means that the crossings could substantially perturb the beam, so the importance of this must be quantified.

Down-selection

Two major choices must be made for the acceleration systems. First we must determine if the neutrino factory accelerator design is satisfactory for the collider. Second we must choose whether to accelerate to high energy with an RLA or an RCS. These choices will be made following technical reviews of the simulated performance, engineering feasibility, and relative costs. Plans for MAP call for choosing a single baseline neutrino factory accelerator design in FY11 and a single baseline muon collider accelerator design in FY13.

5. Collider Ring

Introduction and present status

The final part of the muon collider facility is the collider ring, where the muon beams collide at low-beta interaction points. The proper design of this ring is essential for achieving luminosities ~10^{34} cm^{-2} s^{-1}. The design of the interaction region (IR) is strongly coupled with the design of the detector. All collider ring systems need to be designed as part of MAP and the expected beam behavior needs to be simulated in detail.

Preliminary work has been done on several ring designs. Most of these have assumed the high normalized transverse emittance (~25 μm-rad) from the HEMC scenario. They differ in the location of the closest dipole to the interaction point (IP) and in the arrangement of the sextupole families for chromaticity correction. A ring design for the LEMC scenario based on a low normalized transverse emittance of ~2 μm-rad is also under investigation. The collider design should be flexible enough to accept a wide range of muon beam emittances since there is uncertainty in the amount of emittance reduction that can be obtained from the cooling channel.

We currently have a promising ring design [32] with β* = 1 cm, large momentum acceptance and good dynamic aperture. The beta function and the dispersion in the IR are shown in Fig. 5.1.

The local chromaticity correction in this design utilizes three sets of sextupoles in each IR. The IR produces large positive contributions to the momentum compaction factor that must be compensated by a negative contribution from the arcs. The arc cells use a flexible momentum compaction design. Preliminary simulations have shown that this optics design is relatively insensitive to the beam-beam effect. A summary of the current collider ring parameters are given in Table 5.1. Reducing the β* to 0.5 cm would require a transverse emittance of 12.5 μm.

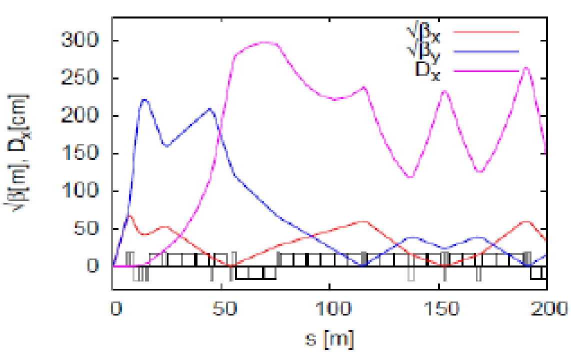

Figure 5.1 Layout and optics functions for the collider IR.

Table 5.1 Collider ring parameters (HEMC scenario).

Parameter	Value
Beam energy (TeV)	0.75
Number of IPs	2
Circumference (km)	2.73
β^* (cm)	0.5 – 2
Momentum compaction, α_p (10^{-5})	-1.3
Normalized transverse emittance (µm)	25
Momentum spread (%)	0.1
Bunch length, σ_s (cm)	1
Muons per bunch (10^{12})	2
Beam-beam parameter per IP	0.09
Dynamic aperture (σ)	5.7
Static momentum acceptance (%)	±1.2
Average luminosity per IP (10^{34} cm^{-2} s^{-1})	1.1

R&D issues and proposed R&D plan

i) Basic ring design and optimization

We need to carry out additional basic lattice design studies of the interaction region, taking into account the choice of beam energy and constraints due to quadrupole gradients and practical magnet apertures. We need to continue examining various chromatic correction schemes, such as special correction sections versus local correction within the IR. We need to study the trade-offs of using FODO cells versus achromats for the arcs. We will also examine the performance trade-offs of having one versus two IRs. Complete ring lattices need to be designed, including special matching sections, injection, collimation, and beam abort. Work on collider lattices must go hand-in-hand with the hardware studies of magnets and superconducting RF, and with the detector studies.

We need to design, analyze, and simulate the RF system for the collider ring. We will optimize the design of the accelerating structure, including a higher-order-mode (HOM) analysis. We will then perform wakefield and impedance simulations to evaluate the requirements for HOM damping and/or feedback systems.

We will develop detailed scenarios for closed-orbit correction and explore other tuning algorithms suitable for these short-lived beams. We will examine the suitability for muon beams of the injection, beam abort, and collimation systems. We will also assess the efficacy of various beam instrumentation devices in the harsh collider environment.

ii) Beam dynamics studies

Simulation work needs to be done on the behavior of the muon beams in the ring. The chromaticity needs to be studied in higher order and the design of the correction schemes needs to be optimized. Tracking studies will be made, taking into account the effects of magnet imperfections (fringe fields, nonlinear components, and alignment errors) on the dynamic aperture. We need to examine the need for compensating these imperfections with dedicated correctors. We also need to understand the effect of synchrotron oscillations on the dynamic aperture. We need to calculate the impedance budget and do a stability analysis of the coherent motion of the muon beams. Beam-beam interactions need to be simulated. Ideas for the suppression of beam-beam effects, such as using a plasma, need to be evaluated.

iii) Achieving $\beta^* \leq 1$ cm at the IP

High luminosity in the muon collider requires a small beta function at the crossing point. Achieving the required beta function, ~1 cm at the IP, will be challenging. This requires a very large beta function in some of the IR quadrupoles and makes the design susceptible to alignment, jitter, and other errors. The small beta function requires that the ring have large momentum acceptance. We need to determine the limits that arise from the chromaticity of the final focusing quadrupoles and from aperture restrictions on the magnets in the IR. On the other hand, the fact that the muons only survive for ~1000 turns gives hope that the beam dynamics may be rather forgiving in many respects. A dynamic aperture of only 3–4σ may suffice to store the beam for this number of turns. Thus, within reason, magnet imperfections, chromatic effects, and beam-beam effects should have relatively little impact on performance.

iv) Problems with electrons from muon decays

One big difference between a muon collider and previous colliders is that the muons are not stable particles. All the muons in the collider ring will eventually decay and produce a large background of high energy electrons. In the arcs these electrons represent an additional heat load for the cryogenics and they can cause radiation damage in insulation and other materials. These problems are also present in the IR but, in addition, the electrons represent a major background in the detector and clearly a dominant issue for the detector design.

Down-selection

Plans for MAP call for simulating the collider ring systems in greater detail and choosing a single baseline ring design in FY12 based on a technical review and the down-selection process discussed in the Introduction.

6. Machine-Detector Interface

Introduction and present status

The design of the collider ring interaction region and the design of the detectors for the muon collider are intimately related. The size and location of magnets near the IP and backgrounds from muon decays constrain the design of detector components. Likewise, limitations on acceptable background rates in detector elements place demands on the IR magnet design. Figure 6.1 shows a MARS15 simulation of particle tracks in the IR. The ability to measure possible physics processes depends critically on the detector characteristics. In addition, work on the IR and the detector must go hand-in-hand with hardware design work on the magnets. These coupled design requirements point to the necessity for coordination and iterative design work on the detector and the IR, which we address by the creation of a dedicated MAP group responsible for the design and simulation of the machine-detector interface.

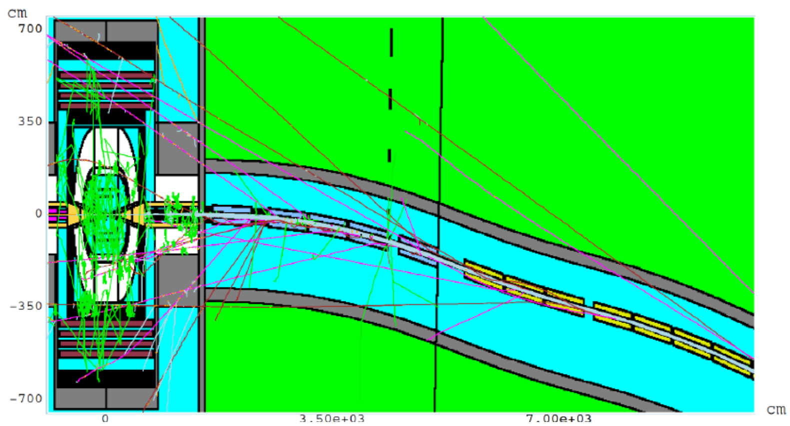

Figure 6.1 Particle tracks in the muon collider IR.

Preliminary designs have been made [33] of an interface region between the collider ring IR and a prototype detector. The radiation fluxes have been greatly reduced using a sophisticated tungsten absorber cone around the beam pipe, liners in the magnet apertures, tungsten masks between the IR quadrupoles, and using open-midplane dipoles.

R&D issues and proposed R&D plan

i) Simulation of radiation levels

Simulations of particle fluxes are needed to help determine lifetime limitations on magnets and other components in the IR. Radiation levels must be computed in the vicinity of the IR and the detector. Heat loads must be determined for cryogenic components throughout the ring.

ii) Design of the IR absorber cones

The flux of electrons from muon decays near the beam axis at the IP is so intense that it must be shielded from the collider detectors. A cone-shaped collimator surrounding the beam pipe that extends into the detector must be designed to protect the tracking systems in the detector from unacceptable backgrounds. The cone angle must be large enough to reduce the backgrounds in the rest of the detector to acceptable levels, yet small enough that it does not seriously impact the ability of the detector to measure the desired physics processes. Making accurate predictions of the backgrounds requires detailed, high statistics simulations of the flux of particles from muon decays. The detailed shape of the opening needs to be optimized.

iii) Control of muon beam halo

The muon beams will develop a halo that needs to be suppressed or else it will represent another background for the detector. The halo muons are too energetic to be directly collimated, so special deflection systems need to be designed to remove these particles from the circulating beams.

iv) Design of auxiliary IR systems

A number of auxiliary systems need to be specially designed for the IR, including the beam pipe, vacuum system, cryogenics, and conventional services.

v) Control of neutrino-induced radiation

The design of the IR also influences the potential issue of off-site radiation caused by neutrino interactions. These neutrinos originate in the muon decays and are most intense in a narrow plane parallel to the collider ring. We do not believe that this is a serious problem at 1.5 TeV, but this needs to be checked by simulating radiation levels for the site boundary.

Down-selection

The baseline design configuration for the machine-detector interface will be chosen following a technical review. This will occur after the specification of the baseline collider ring design in FY12 and the specification of the baseline detector design.

7. Milestones and deliverables

Milestones and deliverables for the Design and Simulation work are given in Table 7.1.

Table 7.1 Design & Simulation Milestones

Date	Milestone	Deliverable
FY10	specify *target* initial configuration	MAP Rev, Des Report
FY11	specify *front end* initial configuration specify *NF μ acceleration* initial configuration	MAP Rev, Des Report MAP Rev, Des Report
FY12	specify *collider ring* initial configuration specify *cooling* initial configuration	Ext Rev, Des Report MAP Rev, Des Report
FY13	specify *proton driver* initial configuration specify *MC μ acceleration* initial configuration	Ext Rev, Des Report MAP Rev, Des Report
FY14	finish D&S for Interim MC DFS report finish D&S for **Final IDS-NF RDR report**	Formal Report Formal Report
FY15	provide specifications & parts count for MC costing	Design Report
FY16	provide description of remaining MC R&D items finish D&S for **Final MC DFS report**	Design Report Formal Report

Down-selections in the early years of this plan will result in a single, complete specification for the muon collider by FY13.

References

[1] A feasibility study of a neutrino source based on a muon storage ring (Study 1), Fermilab report PUB-00/108-E, April 2000.
[2] Feasibility study-II of a muon-based neutrino source, BNL report 52623, June 2001.
[3] J.S. Berg *et al.*, Cost effective design for a neutrino factory, Phys. Rev. ST-AB 9:011001, 2006.
[4] J.S. Berg *et al.*, Accelerator design concept for future neutrino facilities, International Scoping Study of a future Neutrino Factory and Superbeam facility, RAL report TR-2007-23, February 2008.
[5] J.S. Berg, International Design Study of a Neutrino Factory, Proc. IPAC10, p. 3597.

[6] $\mu^+\mu^-$ Collider: a feasibility study, Snowmass 1996, BNL-52503, Fermilab CONF-96/092, LBNL-38946.
[7] C. Ankenbrandt *et al.*, Status of muon collider research and development and future plans, Phys. Rev. ST-AB 2:081001 (1999)
[8] T. Roberts *et al.*, Particle tracking in matter dominated beam lines, Proc. IPAC10, p. 1871.
[9] Project X as a proton driver for a neutrino factory and/or a muon collider, Muons Inc.-Fermilab report, 26 January 2010.
[10] G. Flanagan *et al.*, Using Project X as a proton driver for muon colliders and neutrino factories, Proc. IPAC10, p.4452.

[11] M. Popovic *et al.*, CW SRF h- linac as a proton driver for muon colliders and neutrino factories, presented at Workshop on Applications of High Intensity Proton Accelerators, Fermilab, October 2009.
[12] H. Kirk *et al.*, A 4 MW target station for a muon collider or a neutrino factory, Proc. IPAC10, p. 3590.
[13] K. McDonald *et al.*, The MERIT high-power target experiment at the CERN PS, Proc. IPAC10, p. 3527.
[14] D. Neuffer, Principles and applications of muon cooling, Part. Accel. 14:75, 1983.
[15] R. B. Palmer *et al.*, A complete scheme for a muon collider, Proc. COOL 2007 Workshop, p.77.
[16] D Neuffer *et al.*, Muon capture in the front end of the IDS neutrino factory, Proc. IPAC10, p. 3500.
[17] J. Gallardo and M. Zisman, A possible hybrid cooling channel for the neutrino factory, Proc. IPAC10, p. 3515.
[18] C. Yoshikawa *et al.*, Quasi-isochronous muon collection channels, Proc. IPAC10, p. 3512.
[19] Z. Insepov *et al*., Modeling RF breakdown arcs I: arc evolution, arXiv:1003.1736v2 [physics.acc-ph] 11 June 2010; Z. Insepov *et al*., Modeling RF breakdown arcs II: plasma/materials interactions, arXiv:1006.2370v1 [physics.acc-ph] 11 June 2010.
[20] D. Stratakis *et al.*, Effects of external magnetic fields on the operation of high-gradient accelerating structures, Nucl. Instr. Meth. A 620 (2010) p. 147.
[21] R. Palmer *et al.*, Scheme for ionization cooling for a muon collider, in Proc. NuFact08 Workshop, Proc. Sci. (NuFact08) 019, p. 1-8.
[22] Y. Alexahin, Circularly inclined solenoid channel for 6D ionization cooling of muons, in Proc. PAC09.
[23] R. Palmer, Towards end-to-end 6D cooling simulations, https://mctf.fnal.gov/meetings/2010/04_28/palmer_1004-v3.pdf/view
[24] R. Palmer, Cooling in 50 T solenoids, http://www.cap.bnl.gov/mumu/conf/collider-091201/talks/RPalmer1-091201.pdf
[25] R. Johnson and Y. Derbenev, Low emittance muon colliders, Proc. PAC07, p. 706.
[26] K. Yonehara *et al.*, A helical cooling channel system for muon colliders, Proc. IPAC10, p. 870.
[27] V. Morozov *et al.*, Epicyclic twin-helix magnetic structure for parametric-resonance ionization cooling, Proc. IPAC10, p. 166.
[28] A. Bogacz *et al.*, Status of linac and RLAs, IDS-NF Meeting, Fermilab, 9 April 2010; https://www.ids-nf.org/wiki/FNAL-2010-04-08/Agenda/Accelerator/Files?action=AttachFile&do=get&target=S3-Bogacz-1-v1.ppt
[29] J. Pasternak *et al.*, Injection/extraction system of the muon FFAG for the neutrino factory, Proc. IPAC10, p. 3476.
[30] S.A. Bogacz *et al.*, Recirculating linear accelerators for future muon facilities, Proc. IPAC10, p. 3602.
[31] D. Summers *et al.*, Muon acceleration to 750 GeV in the Tevatron tunnel for a 1.5 TeV $\mu^+\mu^-$ collider, Proc. PAC07, p. 3178.
[32] Y. Alexahin *et al.*, Conceptual design of the muon collider ring lattice, Proc. IPAC10, p. 1563.
[33] Y. Alexahin *et al.*, Muon collider interaction region design, Proc. IPAC10, p. 1566.