

Appendix A

Cost Estimates

A.1 Methodology and Facility Costs

A.1.0.1 Methodology

In this report we have described the components of a Neutrino Factory sited at BNL. The facility includes the following systems:

- Proton Driver (providing 1 MW of protons on target from an upgraded AGS)
- Target and Capture (a mercury-jet target in a 20-T superconducting solenoidal field to capture pions from the target)
- Decay and Phase Rotation (three induction linacs, with internal superconducting solenoidal focusing, to contain the muons from pion decays and provide nearly non-distorting phase rotation; a mini-cooling absorber section is included after the first induction linac)
- Bunching and Cooling (a solenoidal focusing channel with high-gradient rf cavities and liquid-hydrogen absorbers that bunches the 200 MeV/ c muons into 201.25-MHz rf buckets and cools their transverse emittance from 12 mm-rad to 2 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 recirculating linear accelerator to provide a 20 GeV muon beam)
- Storage Ring (a compact racetrack-shaped superconducting storage ring in which 35% of the stored 20 GeV muons decay toward a detector located 2900 km from the ring)

A.1. Methodology and Facility Costs

As part of the Study, we have specified each system in sufficient detail to obtain a “top-down” cost estimate for it. Clearly this estimate is not the complete and detailed cost estimate that would come from preparing a full Conceptual Design Report (CDR). Neither the definition of the various systems, nor the engineering effort available for this Study would permit this. On the other hand, there is considerable experience in designing and building accelerators with similar components, so we have a substantial knowledge base from which costs can be derived. The costs obtained for this Study were obtained mainly in that way.

Where available, we have used costs from existing components—scaled as needed to reflect essential changes in the key variables—to represent the expected costs to fabricate what we need. This applies to the Proton Driver, the superconducting and normal conducting magnets and their power supplies, the rf cavities, and conventional facilities and utilities. In some cases, we were able to take advantage of the experience with designing similar components in a different context. For example, the target facility we require is closely similar to that needed for the Spallation Neutron Source (SNS) project at ORNL, for which detailed CDR-level designs already exist and construction is under way. We made use of the expertise developed at SNS to estimate the facility costs for the Neutrino Factory target area. The superconducting target solenoid is not a standard device, yet even here there is a magnet of similar size and field strength, designed for the ITER project, that serves as a convenient scaling model. Given the limited time to arrive at cost estimates, we leaned heavily on experts in the various areas who could identify the key design parameters that influence costs and then scale accordingly from known costs. In the case of rf power sources, we made use of the multi-beam klystron (MBK) example developed at DESY for TESLA, along with expertise in developing other high-power tubes at U.S. Laboratories. As was done in Study-I, for devices such as the MBK, which are a significant extrapolation from existing hardware, allowance was made for a substantial development program, whose cost was amortized over the initial complement of devices needed for the Neutrino Factory.

A.1.0.2 Facility Costs

It should be noted that the design we have described in this report has erred on the side of feasibility rather than costs. Thus, we do not claim to present a fully cost-optimized design, nor one that has been reviewed from the standpoint of “value engineering.” In that sense, there is hope that a detailed design study will *reduce* the costs compared with what we estimate here. Only direct costs are included here, that is, the estimates do not contain allowances for EDIA, laboratory overhead burdens, or contingency. The breakdown by system is summarized in Table A.1; costs reported there are given in FY01

A.1. Methodology and Facility Costs

dollars. To facilitate comparison with the Feasibility Study-I estimate, we have converted the costs to FY00 dollars, shown in Table A.2, using the DOE-approved inflation factor of 2.5%. Following Study-I, we have put in Table A.2 an allowance of 10% for each of the systems to account for things we have not considered in detail at this stage.

It is interesting to compare our estimate with that of Study-I; in this study, we have **improved the performance by a factor of six** over that reached in Study-I, at a total cost (estimated in the same way for both designs) of about **3/4 of that in the original study**. This is an encouraging trend and, as noted, we have some hope that it will continue.

Table A.1: **Construction Cost Rollup per Components for Study-II Neutrino Factory.** All costs are in FY01 dollars.

System	Magnets (\$M)	RF power (\$M)	RF cav. (\$M)	Vac. (\$M)	PS (\$M)	Diagn. (\$M)	Cryo (\$M)	Util. (\$M)	Conv. Facil. (\$M)	Sum (\$M)
Proton Driver	5.5	7.0	66.1	9.8	26.6	2.2	28.5		21.9	167.6
Target Systems	30.3			0.8	3.5	8.0	18.8		30.2	91.6
Decay Channel	3.1			0.2	0.1	1.0	0.2			4.6
Induction Linacs	35.0		90.3	4.4	163.3	3.0	3.6		19.5	319.1
Bunching	48.8	6.5	3.2	2.7	2.1	5.0	0.3			68.6
Cooling Channel	127.6	105.6	17.7	4.3	4.8	28.0	9.5		19.5	317.0
Pre-accel. linac	46.3	68.4	44.1	7.5	3.0	6.0	13.6			188.9
RLA	129.0	89.2	63.4	16.4	5.6	4.0	28.9		19.0	355.5
Storage Ring	38.5			4.8	2.2	29.0	4.8		28.1	107.4
Site Utilities								126.9		126.9
Totals	464.1	276.7	284.8	50.9	211.2	86.2	108.2	126.9	138.2	1,747.2

A.2. Cost Reduction Options

Table A.2: Summary of Construction Cost Totals for Study-II Neutrino Factory. All costs are in FY01 dollars unless otherwise noted. ^a*Others* is %10 of each system to account for missing items, as was used in Study-I; ^b*Reconciliation* represents the Study-II costs given in FY00 dollars to permit direct comparison with Study-I costs. The inflation factor used is (1/1.025), per DOE official rates.

System	Sum (\$M)	Others^a (\$M)	Total (\$M)	Reconciliation^b (FY00 \$M)
Proton Driver	167.6	16.8	184.4	179.9
Target Systems	91.6	9.2	100.8	98.3
Decay Channel	4.6	0.5	5.1	5.0
Induction Linacs	319.1	31.9	351.0	342.4
Bunching	68.6	6.9	75.5	73.6
Cooling Channel	317.0	31.7	348.7	340.2
Pre-accel. linac	188.9	18.9	207.8	202.7
RLA	355.5	35.5	391.0	381.5
Storage Ring	107.4	10.7	118.1	115.2
Site Utilities	126.9	12.7	139.6	136.2
Totals	1,747.2	174.8	1,922.0	1,875.0

A.2 Cost Reduction Options

A.2.1 Introduction

For this study, an effort has been made to select specific and feasible technologies giving acceptable performance at reasonable cost. Nonetheless there are many alternative ideas that could be considered. In this chapter we discuss options that might lower the cost, improve performance, or be used as alternatives. In some cases, cost reductions may be possible with little sacrifice of performance; other choices would hurt performance, but by amounts that might be justifiable by the savings achieved. Some newer technologies might raise performance and lower costs simultaneously.

A.2. Cost Reduction Options

A.2.2 Capture Solenoid

A.2.2.1 Cost: Choice of Capture Field

Figure A.1 shows the efficiency for muon production *vs.* the axial peak field of the capture magnet. Maximum performance is achieved with the baseline value of 20 T, but the drop in efficiency is small for moderate reductions in this field. A drop from 20 T to 18 T would have an almost insignificant effect ($\approx 2\%$) and even a reduction to 15 T causes only $\approx 9\%$ reduction. The savings, even for a reduction to 18 T could be significant.




Figure A.1: Efficiency *vs.* capture field.

A.2.2.2 Cost/Performance: Use of Wrapped Insulation

Figure A.2 shows the field *vs.* power consumption for three different insert coil technologies (see Section 3.4.) The lowest curve is for the baseline design using MgO insulated hollow conductor giving 6 T with 12 MW. The dotted line above is for a wrapped ceramic insulation as being developed at CTD, Inc. [1]. With this conductor, for the same




Figure A.2: Efficiency of three types of inserts in 20 T magnet; lowest curve: mineral-insulated hollow conductor as developed for the JHF; middle curve: higher-performance hollow conductor under development; top curve: Bitter coil.

power consumption, the field from the hollow conductor would rise from 6 to 7.6 T, thus lowering the field needed from the superconducting coil from 14 T to 12.4 T, and offering significant savings. Alternatively, the gain in performance could be used to reduce the power consumption, or some combination of these two options could be considered.

A.2.3 Bitter Magnet

The upper dashed line in Fig. A.2 is for a Bitter magnet.

The Bitter magnet design is the invention of Prof. Francis Bitter of MIT, who in the late 1930's first used such magnets to generate 10 T in a 5 cm bore. The design has the potential to be a very efficient insert for the pion capture magnet. The windings of a Bitter magnet are sets of thin annular plates, each like a big washer, slit along a radius, as in a lock washer. In each plate a voltage difference between the two edges of the slit forces the current to flow circumferentially, the long way around from one edge of the slit to the other, before entering the next plate. Tie rods or components of the

A.2. Cost Reduction Options

magnet housing keep the plates in good registration and provide the axial clamping for good electrical contact over the sectors in which current transfers from one plate to the next.

The Bitter design has many virtues. It possesses great inherent strength and permits the use of a wide range of conductors, such as heavily cold-worked copper, with excellent combinations of strength and electrical conductivity. Therefore the conductor can resist the huge tensile hoop stresses that arise in generating intense fields. The fraction of conductor in a Bitter magnet typically is much higher than in a magnet built from hollow conductors. One reason is that only a thin film between adjacent plates suffices to confine the current to its desired path, because the potential difference between adjacent plates is only a few volts. Another reason is that cooling passages may be very small, because they are so short. This is true especially if one cools the magnet radially, by means of shallow grooves etched into one face of each plate (or each pair of plates; one can mate each etched plate with an unetched one). The cooling passage length in such a magnet is its “build” (outer radius minus inner radius). If, instead, one chooses to cool the magnet axially, through holes punched in each plate and insulator, the cooling passage length will be the magnet length. For the pion capture insert coil, axial passages are several times longer than radial ones—but still short, by an order of magnitude, relative to those in a magnet employing hollow conductors (See, Section 3.4). The favorable cooling geometry enables Bitter magnets to operate at very high power densities. Another virtue of the Bitter magnet design is the ease with which desired field profiles can be achieved by employing turns of the appropriate thickness in each of many axial zones. One need only change the thickness of plates comprising a turn, or change the number of plates making up each turn.

Figure A.3 plots the relative costs of various systems, each with the peak field of its associated superconducting magnet. The set labeled “unshielded” employs a Bitter magnet whose bore accommodates just the pion capture beam tube and radial clearance for an annulus to bring water to the radial cooling passages. The annulus is tapered from the upstream to the downstream end, in order to maintain a water velocity in the annulus of about 10 m/s. The annular height is at most 1.8 cm. For the set labeled “shielded” the bore accommodates 10 cm of shielding with water-cooled tungsten carbide, just as for the magnet with hollow conductors. Each magnet has an outer diameter of 80 cm if shielded, 40 cm if not—values close to the optimum.

To consider a Bitter magnet for the insert to the pion capture system will require an R&D program that validates three issues:

- verification that radiation will not immediately induce arcing so severe that a substantial fraction of current flows through the arc instead of through the copper




Figure A.3: Relative cost of pion capture magnet, as function of the power consumed by its Bitter magnet, with and without 10 cm of shielding with water-cooled tungsten carbide. Decreasing the power of the Bitter magnet by a factor of four from the 8-10 MW maximum plotted here entails a $\approx 50\%$ increase in system cost; the needed field contribution from the superconducting magnet rises from ≈ 12 T to ≈ 15 T.

windings

- development of an insulator—undoubtedly a ceramic—that will withstand not only the intense radiation emanating from the target but also the environment of a Bitter magnet. (Even without radiation this environment involves high clamping pressure, high temperature and high water velocities.)
- verification that conductors will not deteriorate too much in strength and ductility when irradiated for at least a few months.

If so, one can save many megawatts of power consumption and/or many millions of dollars of capital cost in superconducting magnets—a tantalizing prospect for economy for the Neutrino Factory.

A.2. Cost Reduction Options

A.2.4 Phase Rotation

A.2.4.1 Cost: Combining Induction Linacs 2 and 3

In the baseline design there are 3 induction linacs. The first linac must be separate from the other two in order to achieve non-distorting phase rotation, but the second and third linacs are separate only in order that they each be unipolar. A single second linac with a bipolar pulse approximately equal to the sum of the two opposite polarity pulses would perform equally well. This appears possible and would be somewhat less expensive.

A.2.4.2 Cost: Fewer Induction Linacs

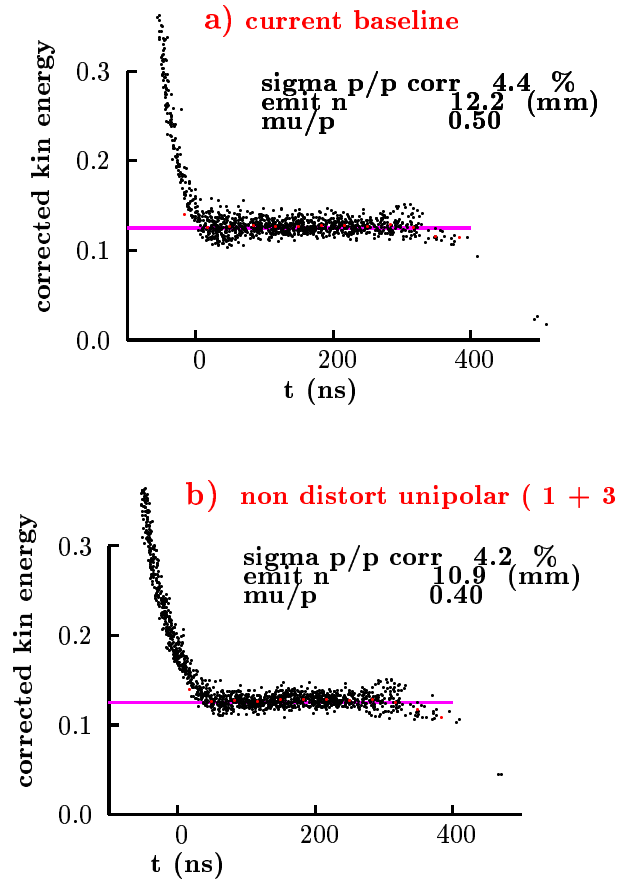


Figure A.4: Final energy *vs* time for different phase rotation systems: a) baseline; b) with IL2 removed.

A.2. Cost Reduction Options





Figure A.5: Final energy *vs* time for different phase rotation systems: c) without IL1; d) without IL1 and IL2.

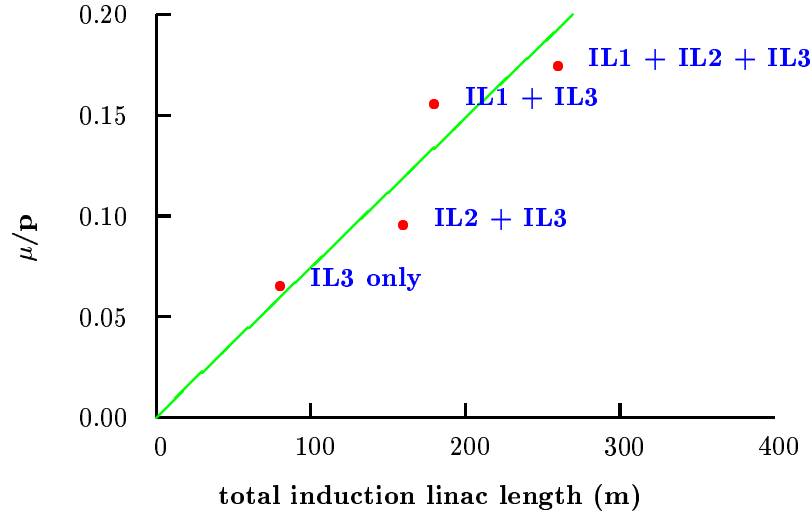


Figure A.6: Efficiency *vs.* length of induction linacs. The μ/p ratio that characterizes the performance of the front end is measured at the end of the cooling channel.

Further cost savings could be achieved if one or more of the linacs were eliminated and the remaining linacs re-optimized. This has been studied assuming a fixed geometrical layout, so that the option of upgrading to the original baseline design is retained. Figures A.4 and A.5 show 3 such cases, together with the baseline design. Figure A.6 shows the muon production efficiency (at the end of the cooling channel) for the four cases, plotted against the sum of the lengths of the remaining linacs. The losses in efficiency are large if the first linac is eliminated, but less severe (11%) if only the second linac is removed. Removing IL2 would provide a cost saving of about 4%, so its presence is favorable from a cost-benefit standpoint.

A.2.5 Cooling

A.2.5.1 Cost: Less Cooling

Figure A.7 shows the muon production into the defined accelerator acceptance as a function of length. Table A.3 shows the values for three cooling lengths. It is seen that a

A.2. Cost Reduction Options

reduction in cooling length from 108 to 88 m, which would offer significant savings, reduces the performance by only 3.4%. Looked at in terms of marginal costs, however, we note that the baseline scenario still appears cost effective.

Table A.3: Efficiency for three cooling lengths.

Cooling length (m)	μ/p	Loss (%)	Savings (%)
108	0.174	0	0
88	0.168	3.4	2.5
68	0.150	13.8	6.7

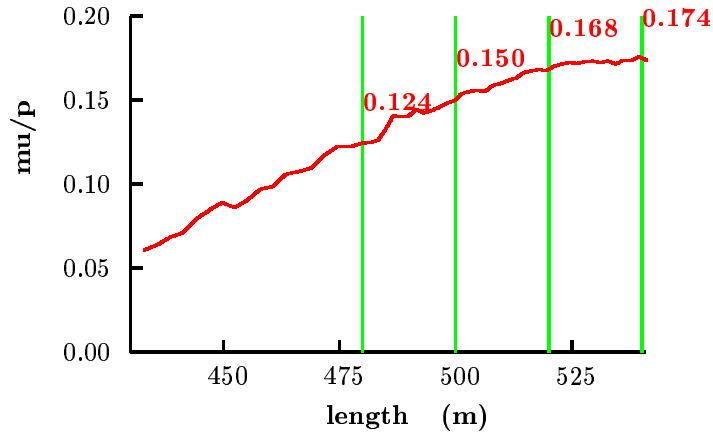


Figure A.7: Efficiency *vs.* length of cooling.

A.2.5.2 Cost: Fixed Field Alternating Gradient

Fixed Field Alternating Gradient (FFAG) acceleration offers the possibility of savings. There would be no multiple arcs, and no switchyards: the lattice would have a large enough momentum acceptance to circulate the muons from initial to final energy. The

number of turns could now be raised, limited only by muon decay considerations, thus lowering the needed rf acceleration per turn.

Lattices have been designed with momentum acceptances of more than a factor of 2-3. Injection and extraction would be performed using kickers. Designs being studied at KEK [2] employ low frequency, low accelerating gradient rf and accept relatively large decay loss. Work in the US [3] has concentrated mainly on higher gradient superconducting rf with fewer turns and less loss. The main problem in this approach is assuring that the rf phase is set correctly at each pass. The ideal solution is a ring that is exactly isochronous, but the best current designs are less than ideal and require phase control of the rf corresponding to frequency variations of the order of 10^{-4} . This would be easy for conventional rf, but is difficult in a superconducting cavity. The use of ferrites weakly coupled to such cavities is being studied.

A.2.6 Summary

Although we believe that the current Study-II baseline represents a feasible and reasonably costed high performance design, there are many possibilities for cost reduction that could be considered for an initial implementation.

A.2. Cost Reduction Options

Bibliography

- [1] J. Rice, DoE SBIR Phase I Final Report, *Ceramic Insulation for Heavy Ion Fusion and Other High Radiation Magnets*, DoE Grant No. DE-FG03-00ER82979 (2000).
- [2] Y. Mori, *KEK FFAG Program*, presented at NUFACT01
http://www-prism.kek.jp/nufact01/May25/WG3/25wg3_mori.pdf.
- [3] C. Johnstone, *US FFAG Program*, unpublished.