A Cost-Effective Design for a Neutrino Factory

R.B. Palmer [a,1] C. Johnson [b] E. Keil [b]

[a] *Brookhaven National Laboratory, Upton NY, USA*

[b] *CERN, Geneva, Switzerland*

Abstract

The design of a neutrino factory based on a muon storage ring draws upon several tried and tested technologies, upon existing design work for other accelerator projects, e.g. neutron spallation sources, but it also depends on the development of technical solutions to certain specific requirements. These include the efficient capture of muons in a large volume of phase space, some reduction in overall phase space volume by ionization cooling, fast acceleration to the desired energy to avoid unacceptable decay losses and storage in a decay ring optimised for its purpose as a neutrino source. There is no obvious single combination of machines to achieve this aim. Here we present a scenario which relies to a large extent upon known technologies together with a relatively unambitious mix of new schemes. Some will be tried and tested during design and construction. Others during the early operational phase of the facility leading to a staged upgrade path – a well-proven strategy in the development of accelerator complexes.

1 Introduction and Design Philosophy

There is now a well-established interest in the generation of intense neutrino beams from the decay of muons circulating in a storage ring (1), (2). Conventional neutrino beams employ a proton beam on a target to generate pions, which are focused and allowed to decay into neutrinos and, incidentally, muons. The muons are dumped (stopped in shielding) and the neutrinos (muon-neutrinos) are directed to the detector. In a neutrino factory, pions are made the same way and allowed to decay, but it is the decay muons that are captured and used. The initial neutrinos are discarded, or used in a parasitic low energy neutrino experiment. The muons are accelerated and allowed to

[1] Supported by US Department of Energy contract DE-AC02-98CH10886

decay in a storage ring with long straight sections. It is the neutrinos from the decaying muons that are directed to a primary detector or detectors.

A muon-ring neutrino factory will be more expensive than a conventional source, but will offer a much higher neutrino fluence and a more complete range of neutrino experiments. Attention is now directed towards writing an optimal performance specification for such a machine and this will be an ongoing activity on a world stage for some time to come. This report represents one early and incomplete attempt to set a direction for this work. It is written for the specialist and assumes that the reader has prior knowledge of many aspects of the subject. All of the background material can be found in the status report of the Muon Collider Collaboration study (3). Muon current and polarization, neutrino energy and fluence, decay-ring geometry and siting, will be factors influencing the time scale and cost of the project. A staged approach in which all of these parameters will be upgraded, as the accelerator technology advances is natural and in keeping with all previous accelerator projects. We must devise a scenario that reaches the threshold of experimental interest soon after completion of the initial installation. This is our goal in describing this scenario despite our present sketchy knowledge as to where this threshold of experimental interest may lie.

Our design of the muon production system for a muon decay ring starts with the muon decay ring proper, in which the threshold-of-interest muon current circulates. We try to design the upstream components such that they do not cause bottlenecks in the process. Of course, we do not know at this stage that this can be actually be done. However, identifying such bottlenecks and finding a way of overcoming them is the name of the game. We outline the overall scheme in Section 2 and then in Section 3 we present our proposed solutions for each of the sub-systems working backwards from the decay ring. Finally in Appendix A we present a very preliminary set of parameters and performance estimates for this facility.

2 Overview

Fig. 1 shows the main components of the neutrino factory. A proton beam with momentum in the range 2 to 20 GeV/c and mean power initially in the region of 1 MW is targeted onto a liquid mercury jet within a 20 T solenoidal field - the configuration that has been optimised in the US muon collider study (3) for muons with longitudinal momentum in the range 100 to 300 MeV/c and transverse momenta from 0 to 225 MeV/c. The expected muon yield in this region is approximately 0.01 muons per proton per GeV (proton energy). The muon bunch length increases rapidly due to the relatively large velocity spread and the additional spread in the time of flight arising from the increase

in trajectory length with Larmor radius of the high p_T muons.

The initially very large momentum spread is reduced in two stages: a low-frequency RF phase rotation is followed by a short section of ionization cooling and then a long drift space to establish a correlation between momentum and longitudinal position, which is then used to perform a second phase rotation by means of an induction linac. This arrangement has the additional virtue of preserving muon polarization, which becomes correlated to laboratory momentum after the RF phase rotation, and correlated with time after the induction linac phase rotation. After this stage the very long bunches (80 m) are re-captured adiabatically at a frequency of 175 MHz.

To achieve ultimate performance in terms of maximizing the acceptance of muons into the downstream elements, an additional section of ionization cooling is being considered. The muons are then accelerated, initially to 600 MeV/c in an RF linac at 175 MHz, and then to 2 GeV/c in a 350 MHz super-conducting linac. Two recirculating linear accelerators (RLA's or Recirculators - as developed at the Thomas Jefferson Laboratory) accelerate the beam to the final desired energy whence it is fast-injected into a decay ring. This process is repeated at a repetition rate somewhere in the range 5 to 50 Hz depending on the proton driver.

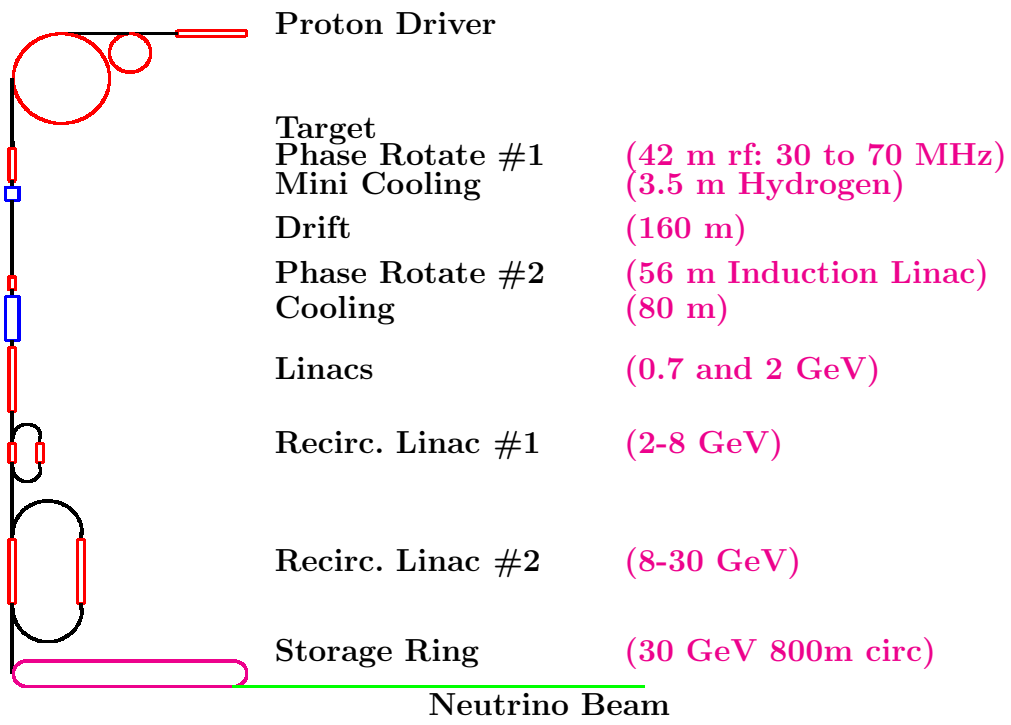

Fig. 1. Main components of the neutrino factory

3 Neutrino Factory Modules

The muons are accelerated in three stages, one or two straight linear accelerators, and two recirculating linear accelerators, μRLA1 and μRLA2, and stored in a muon storage ring μSR. Tab. 1 shows the parameters of the muon beam at their input ends. Our goal is to achieve *output* parameters of the various modules that are identical to the *input* parameter of the subsequent module. We believe that our parameters are feasible, but do not pretend that they are optimized.

Table 1
Beam Parameters at Input of Modules. The relative RMS momentum spread σ_e scales with energy E roughly like $1/E$, assuming that μRLA1 and μRLA2 are isochronous

	μSR	μRLA2	μRLA1	
Injection momentum injPc	30	8	2	GeV/c
Normalized RMS emittance $\epsilon_{\rm xn}$	1.667	1.671	1.830	mm
Relative RMS momentum spread σ_e	0.003	0.011	0.045	
Bunch train length L	<0.8	<0.8	<0.8	km
Beam current I	0.1			A

We define the normalized transverse emittance by $\epsilon_{\rm xn} = \beta\gamma\sigma_x^2/\beta_x$, where β and γ are the relativistic factor for the stored muons, while σ_x and β_x are RMS beam radius and β-function at the same place along the machine, respectively. Similarly, we define the normalized transverse acceptance by $A_{\rm xn} = \beta\gamma a_x^2/\beta_x$, where a_x is the vacuum chamber radius. We assume essentially round beams with similar emittance in the horizontal and vertical direction, and use the subscript x for any one of them. In parallel we use the unnormalized emittance ϵ_x and acceptance A_x that are a factor $\beta\gamma$ smaller than the normalized quantities. We typically assume that the ratio between vacuum chamber and RMS beam radius is sigN ≈ 3.

3.1 Muon Storage Ring

Muons are injected into the μSR with a full-aperture kicker system. The injected muon pulse must be shorter in time than a revolution period. There is no accumulation or stacking of muons in any direction of phase space. Accumulation would only be useful if the emittance and/or momentum band of the incoming muon beam were smaller than the acceptance of the decay ring. This may happen either because there is a bottleneck somewhere upstream that should be removed or the muon production, cooling and acceleration re-

sults in a smaller transverse emittance and momentum band that the accepted values, which we consider unlikely. In this case the acceptance of the muon decay ring could be smaller. The injection kicker pulse must decay faster than the difference between revolution period and muon pulse length. There is no requirement on the kicker rise time. The kicker parameters are reasonable (4).

Our design procedure for the μSR is implemented in a *Mathematica* notebook (4) that makes searches for plausible parameters easy, and explains the names of some variables. Tab. 2 shows the parameters that we assume for the first step, and those parameters that can easily be calculated from them. We assume that the μSR use has the shape of a racetrack. Adapting our procedure to the bow-tie or other configurations is straightforward.

Table 2
μSR Design Parameters. Assumed parameters are marked by an asterisk.

∗Collection momentum	30	GeV/c
∗Dipole field	7	T
∗Quadrupole field	4	T
Bending radius ρ	14.2956	m
∗Arc filling factor rhoR	0.6	
Average arc radius arcR	23.826	m
Total arc length totArcL	149.703	m
∗Total straight section length	650.297	m
Circumference cirC	800	m
Revolution frequency revolF	374.738	kHz

In the second step in our design procedure we assume that the arcs of the μSR use super-conducting magnets. We compute the quadrupole gradient from the ratio of the quadrupole field quadB and the vacuum chamber radius vacuumA that we don't know at this stage. We use the arc filling factor rhoR, the ratio of the sum of the dipole lengths and the arc length, as one of our adjustable parameters. We obtain the parameters of the FODO cells in the arcs shown in Tab. 3. The calculation uses the standard thin-element formulae, ignoring edge focusing (5). The free space in a half period freeS, i.e. the space not occupied by dipoles and quadrupoles, is controlled by the filling factor rhoR.

Simple two-cell dispersion suppressors with the same quadrupole strengths and weaker dipoles are used at the ends of the arcs. A complete arc consists of 7 FODO cells, surrounded by two dispersion suppressor cells on either side. The total bending in two dispersion suppressor cells is equal to that in a single FODO cell when edge focusing is neglected (5). Fig. 2 shows the schematic arrangement and the optical functions, computed with MAD (6)

Table 3
FODO Cell Parameters of μSR Arcs

Number of arc periods periodN	14	
Number of dispersion suppressor periods dispN	8	
Arc period length periodL	8.31685	m
Dipole length dipoleL	2.49505	m
$*$Arc cell phase advance/2π phaseAdv	0.25	
Arc tune tuneQ	5.5	
Maximum β-function $\beta_{\max}$	14.1977	m
Maximum dispersion $D_{\max}$	1.96477	m
Momentum compaction η	0.00924094	
Unnormalized emittance ϵ_x	5.87003	μm
RMS beam radius σ_x	10.8666	mm
$*$Number of standard deviations sigN	3	
Vacuum chamber radius vacuumA	32.5999	mm
Maximum half momentum spread Δ_e	0.0165922	
Focal length of quadrupoles focalL	2.94045	m
Quadrupole length quadL	0.277359	m
Free space in a half period freeS	1.38601	m

after matching with the *Mathematica* results as initial guesses, including edge focusing and finite-length elements.

The optical arrangement in the long straight sections consists of FODO cells without dipoles. The dispersion vanishes there. With the total straight section length shown in Tab. 2, about 36.7% of the muons decay into neutrinos flying towards the neutrino detector. The length of the lattice periods in the straight sections and their phase advance are chosen such that the ratio between the divergence of the muon beam and the opening angle of the neutrino beam, due to the muon decay kinematics, i.e. the normalized divergence $\gamma\sqrt{\epsilon_x\gamma_x}$, is about 0.1. In this case, the neutrino beam divergence increases by about 1%, if the two contributions are added quadratically. In a thin-lens FODO cell with length L and phase advance μ we have $\gamma_x = (1+\alpha_x^2)/\beta_x = 4\tan(\mu/2)/L$. Tab. 4 shows parameters of the long straight section lattice that achieve this goal. The matching between the arcs and the long straight sections and the adjustment of lattice parameters to the exact length of the long straight section remain to be done.

The value of the momentum compaction η in Tab. 3 takes into account the

Table 4
Long Straight Section Parameters

∗Length of cells periodS	100	m
∗Normalized divergence normD	0.1	
Phase advance in cells straightAdv	0.154702	
Maximum β-function maxstraightBeta	177.61	m
Minimum β-function minstraightBeta	64.5134	m
Maximum RMS beam radius straightSigma	32.2889	m
Aperture radius straightAx	96.8667	m
Focal length of quadrupoles focalS	53.5215	m
Length of quadrupoles straighQuadL	0.045278	m

effect of the long straight sections that contribute to the circumference, but not to orbit length changes for off-momentum particles. With our parameters, the bunch length increases by about 25 mm on a turn. An RF system with a synchrotron tune larger than the rms momentum spread is needed to preserve polarization (7). It may not be needed to simply keep the muon beam bunched and observable. It is an open question whether the chromaticity in such a machine has to be corrected. Matching the derivatives of the orbit functions α and β with respect to the momentum error between the arcs and the long straight sections, using a few sextupoles in the dispersion suppressors, may be good enough.

3.2 Recirculating Linear Accelerators

The higher the injection momentum into the recirculator, the easier it is to achieve the desired acceptance and momentum band. A smaller number of passes allows a larger momentum spread. A momentum gain of more than a factor 10 probably requires more than one recirculator. To emphasise this point, in the schematic layout in our Fig. 1 we propose two recirculators: the first with a gain per turn of 1.5 GeV/c taking the beam momentum from 2 GeV/c up to 8 GeV/c after four turns and a second with a gain of 5.5 GeV/c per turn giving a final momentum of 30 GeV/c after four turns. Here the overall momentum gain is a factor 15. The choice is a project planning issue.

Both μRLA1 and μRLA2 have some similarities with the electron recirculating linear accelerator ELFE at CERN (8) that is being studied now. The RF system uses super-conducting cavities very similar to those in LEP (9). Tab. 5 shows the main parameters of the straight sections containing the linear accel-

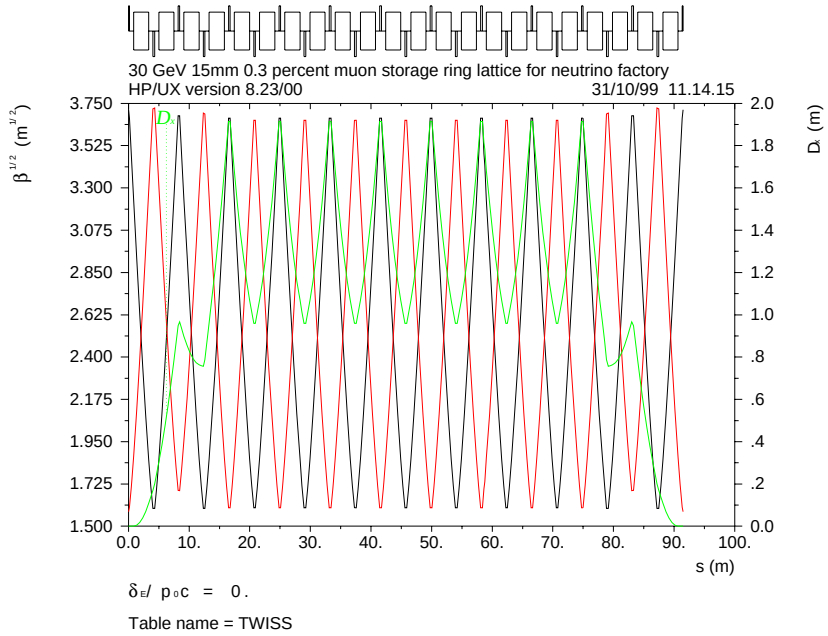

Fig. 2. Schematic layout and optical functions in an arc of the muon storage ring

Table 5
Parameters of the Linacs in μRLA1 and μRLA2. The optical parameters apply to the first pass.

	μRLA1	μRLA2	
∗Injection momentum injPc	2	8	GeV/c
∗Number of passes passN	4	4	
∗Ejection momentum ejPc	8	30	GeV/c
∗Frequency of RF system freqRF	352.209	352.209	MHz
Wavelength of RF system $\lambda_{\rm RF}$	0.851178	0.851178	m
∗Radius of beam ports cavityA	0.09925	0.09925	m
∗Number of RF modules moduleN	128	112	
Total accelerating voltage totalV	1.5	5.5	MV
∗Length of half cell/$\lambda_{\rm RF}$ halfCellL	4	17.5	
Total length of linac totLinL	435.803	1668.31	m
∗Phase advance in first pass linAdv	0.212071	0.212071	
Maximum β-function in linac lin$\beta_{\rm max}$	11.3383	49.6052	m
Normalized acceptance $A_{\rm xn}$	16.4681	15.0369	mm
Normalized RMS emittance $\epsilon_{\rm xn}$	1.82979	1.67076	mm

erator sections. We assume that the tapered sections at the ends of the LEP RF cavity modules are replaced by cylindrical ones. The aperture radius of new quadrupoles is at least equal to the beam port radius in the RF cavities. We choose a phase advance $\mu/2\pi = 0.212071$ which yields the smallest value of the maximum β-function at a given cell length. In μRLA2, we use the LEP RF cavity modules as they are, install one module in a half cell, arrive at the same half cell length as in ELFE (10), and achieve the specified acceptance. In μRLA1, we need a half cell length of 4 λ_{RF} in order to achieve the specified acceptance at 2 GeV/c. Therefore, we install single four-cell LEP RF cavities in each half cell. If there is not enough space for two warm-to-cold transitions in a half cell, we need cryo-modules containing both RF cavities and quadrupoles.

We assume for the time being that in μRLA1 and μRLA2 the muons are accelerated at the crest of the RF wave. Since the arcs are assumed to be isochronous, phase oscillations are absent, and the RMS bunch lengths σ_s at the exit and at the entrance are the same. The curvature of the RF wave and the finite bunch length cause an increase of the relative RMS momentum spread σ_e in the bunches that is related to the wavelength of the RF system λ_{RF} by (8):

$$\sigma_e = \frac{1}{\sqrt{2}} \left(\frac{2\pi\sigma_s}{\lambda_{\mathrm{RF}}} \right)^2 \tag{1}$$

At the RF wavelength shown in Tab. 5, achieving $\sigma_e \approx 0.003$ requires an RMS bunch length $\sigma_s \approx 8.8$ mm at the entrance of μRLA1 and μRLA2. Two phenomena contribute to the relative momentum spread at the exit: (i) The momentum spread already present at the entrance and (ii) the momentum spread caused by bunch length and RF waveform (1). The former is adiabatically damped in the ratio of input and output energies, the latter is constant.

The passage of the bunch train through the RF cavities extracts energy from the energy W_s stored in the RF cavities, and causes a drop of the accelerating voltage along the train. The shunt impedance R, the cavity voltage U, and the RF power P are related by $R = U^2/P$. The quality factor Q, circular RF frequency ω, W_s and P are related by $Q = \omega W_s/P$. Taking the ratio of the two expressions and solving for W_s yields:

$$W_s = \frac{U^2}{\omega(R/Q)} \tag{2}$$

With typical LEP cavity parameters $U = 10\,\mathrm{MV}$, $\omega = 2\pi 352$ MHz, and $R/Q = 464\,\Omega$, we find $W_s \approx 97.4$ Joule. The energy W_e extracted by the a bunch train of average current $\bar{I} \approx 0.1$ A and duration $T \approx 2$ μs is $W_e = U\bar{I}T \approx 2$ Joule.

With these figures for W_s and W_e, we must expect a energy drop of about 1 % along the bunch train. The quality factor Q of the RF cavities is far too large to compensate the voltage drop during the train passage. In linear colliders, a similar problem is solved by running sections at slightly different frequencies (11). This method cannot be used in our super-conducting RF system, because the bandwidth is too small.

Since the muon velocity has $\beta < 1$, there is a phase slip along the linac that is largest for the first pass, and amounts to about 0.9235 r in μRLA1 when the RF cavities are installed for $\beta = 1$. By adjusting the phase of the muons at the entrance of the linac the effective acceleration remains about 90% of the nominal value for the first pass. For later passes, the phase slip is smaller and the entrance phase of the muons can be adjusted by fine tuning the lengths of the arcs.

The RF system in μRLA1 and μRLA2 need only accelerate pulses that are at most as long as the circumference of the decay ring. It may be pulsed and use room-temperature cavities if this is less expensive.

We assume that the beams from the first linac are fanned out vertically by a spreader, fed into four vertically stacked arcs with similar radii of curvature, combined into a single beam line by a combiner, and fed into the second linac. A similar arrangement of spreader, arcs and combiner is at the other end of the linac. The limit on the momentum spread and emittance in the beam arises from the requirement that the beams of successive passes must be fed into different beam apertures by the spreader (10). The relative momentum spread at the output momentum of the recirculating linear accelerator must be much smaller than the relative momentum difference between the last two passes, i.e. much smaller than $(1-24.5/30) \approx 0.1833$ and $(1-6.5/8) \approx 0.1875$, respectively, under our assumptions. This fact is one of the reasons why a smaller number of passes is more favourable.

Our design procedure for the arcs of μRLA1 and μRLA2 is also implemented in a *Mathematica* notebook. Tab. 6 shows the parameters that we assume, marked by an asterisk, and those parameters that can easily be calculated from them. The arcs resemble those of ELFE (10). The arcs consist of achromats in which the dispersion vanishes at either end. An achromat consists of two FODO cells, with equal and opposite quadrupole gradients, and dipoles only in the first and last half cell. The strengths of the quadrupoles are adjusted such that the dispersion vanishes at either end. The phase advance through an achromat, and the maximum values of the β-functions and dispersion follow. Since we are not concerned with the enlargement of beam emittance and momentum spread by synchrotron radiation, we can use super-conducting magnets in the arcs of μRLA1 and μRLA2. We compute the quadrupole gradient from the ratio of the quadrupole field quadB and the vacuum chamber radius vacuumA

Table 6
Arc Parameters for the last pass of μRLA1 and μRLA2. The calculation ignores the fact that the muon momentum on the last pass through the arcs is lower than the ejection momentum.

	μRLA1	μRLA2	
∗Dipole field dipoleB	7	7	T
∗Quadrupole field quadB	4	4	T
∗Arc filling factor rhoR	0.5	0.5	
Bending radius ρ	3.81216	14.2956	m
Average arc radius arcR	15.2486	57.1824	m
Total arc length totArcL	95.8101	359.288	m
Circumference cirC	531.613	2027.6	m
Revolution frequency revolF	563881	147855	Hz

that we don't know at this stage. We use the arc filling factor rhoR, the ratio of the sum of the dipole and quadrupole lengths and the length of a half cell, as one of our adjustable parameters. Tab. 7 shows a list of the arc parameters of μRLA1 and μRLA2 at the ejection momentum. Because of the adiabatic damping of betatron and synchrotron oscillations, the aperture of the arc at injection momentum is larger by a factor approximately equal to the square root of the ratio of ejection and injection momentum, i.e. about two.

Figs. 3 and 4 show the schematic layout and optical functions of the achromats in μRLA1 and μRLA2, computed with MAD (6) after matching with the *Mathematica* results as initial guesses, including edge focusing and finite-length elements. We have not checked the dynamic aperture by tracking. The critical part of this scheme is almost certainly the spreaders and recombiners because they must handle a large emittance and a large momentum band at the same time. We have not studied the obvious alternative, recirculators of dog-bone shape with one linac for acceleration in both directions, a horizontal spreader and combiner, and arcs of very different radii of curvature.

3.3 Linear accelerator

The long muon bunches prepared by the second phase rotation to have a reasonably flat momentum distribution and a momentum spread of $\approx 2\%$ rms are captured adiabatically at 175 MHz. This sub-harmonic is included because of the need to ensure adequate transverse acceptance up to the point where adiabatic damping has reduced the beam size sufficiently to pass cleanly into a second, lower aperture, 350 MHz linac. This will provide acceleration up

Table 7
Parameters of the achromats in the arcs of μRLA2 and μRLA1 at ejection momentum.

	μRLA1	μRLA2	
Number of achromats achroN	8	16	
Number of FODO cells cellN	16	32	
Arc cell length cellL	5.98813	11.2277	m
$*$Dipoles per half cell dipoleN	1	1	
Dipole length dipoleL	1.49703	2.80694	m
$*$Achromat phase advance/2π phaseAdv	0.570352	0.570352	
Arc tune tuneQ	4.56282	9.12563	
Maximum β-function in arcs $\beta_{\max}$	10.9294	20.4926	m
Maximum dispersion in arcs $D_{\max}$	2.68166	2.51406	m
Unnormalized emittance at ejection ϵ_x	24.1644	5.8843	μm
RMS beam radius at ejection σ_x at $\beta_{\max}$	22.4948	13.3217	mm
Vacuum chamber radius vacuumA	67.4845	39.9652	mm
Maximum half momentum spread Δ_e	0.0251652	0.0158967	
Focal length of quadrupoles focalL	1.91736	3.59506	m
Quadrupole length quadL	0.234806	0.27811	m
Free space in a half period freeS	1.26223	2.52882	m

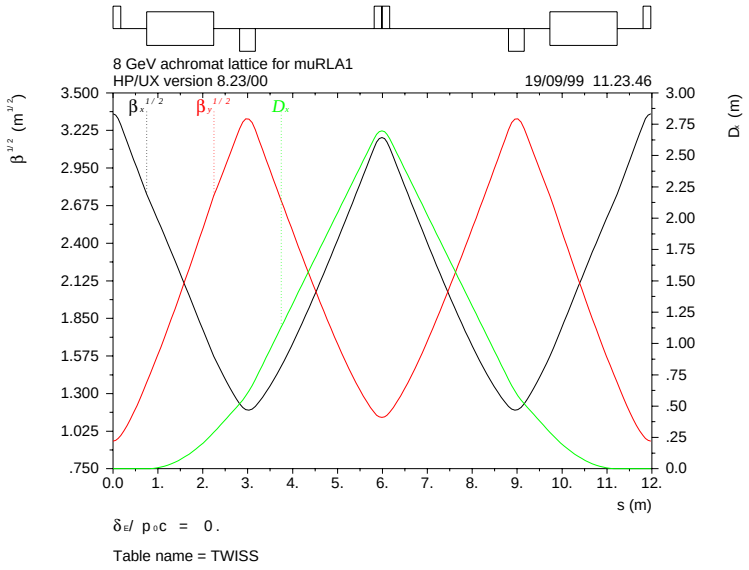

Fig. 3. Schematic layout and optical functions of the achromat in the μRLA1 arcs

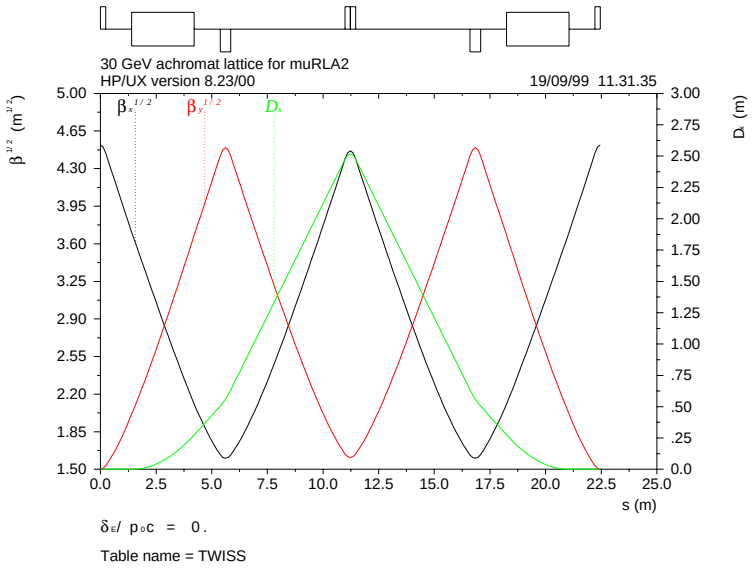

Fig. 4. Schematic layout and optical functions of the achromat in the μRLA2 arcs

to the injection energy of the first recirculator. Large beam apertures, and thus low RF frequency, are particularly advantageous for the linac front end, particularly if it turns out that we wish to maintain solenoid focusing over this section.

In our scenario, the acceptance of the first linac up to 0.6 GeV/c must not be smaller than the acceptance of the downstream modules, i.e. 15 mm. Hence, we propose 175 MHz cavities. For further acceleration up to the recirculator injection momentum of 2 GeV/c, we propose 350 MHz super-conducting cavities and a FODO lattice. Again, a pulsed RF system, that most likely has a higher accelerating gradient, a shorter length, and hence less muon losses due to decay, is probably the most attractive.

3.4 Muon cooling

A transverse ionization cooling system may be installed in front of the muon accelerator. Its purpose is to reduce the transverse muon beam emittance, if this is found to be desirable and affordable. It is useful if and only if the emittance of the muon beam at this stage, determined by targeting, pion production, and muon decay kinematics, is larger than the emittance that can be handled by the components installed downstream. Such a cooling system can be installed after the facility starts up as a part of the improvement programme. Preliminary estimates show that, if the full advantage is to be taken of the muon production and capture section designed by the US Muon Collider Collaboration, MCC, some transverse cooling will be needed.

On the other hand if, to reduce costs, a target station with lower emittance and hence lower muon yield is chosen some or all of the transverse cooling may be omitted. Momentum cooling (phase-space exchange using wedges) can also be considered as an upgrade path if it appears cost effective – this is related to the recirculator spreaders and recombiners and may have to be anticipated in the design as a future upgrade in order to overcome a bottleneck in the recirculators (see Section 3.2).

3.5 Muon Capture and Phase Rotation

The proton production target will be inside a 20 T hybrid solenoid the pions leaving via a solenoid matching section into a 1.25 T solenoidal decay line which incorporates the first stage of phase rotation to reduce the longitudinal momentum spread. This intricate system is the heart of the facility and much work has already been done on its optimisation for maximum muon yield into a muon collider by the MCC study. For the neutrino factory the beam requirements are more relaxed: several proton bunches, longer than for the collider, can be targeted per repetition cycle. Muon emittance must satisfy the constraints of the downstream systems, in particular the muon accelerator and the decay ring transverse acceptances and, as we already mentioned, there is an upper limit on the momentum spread that can be accepted in the recirculator.

The beam optics is highly non-linear. The phase space of the pion and muon flux from the target must be computed by particle tracking and compared to the acceptances of the downstream modules. We do not expect the acceleration and beam transport systems to accept all of the muons created in the decay channel and so it is at this level that the muon yield is determined. In this section we present the results of yield estimates based on the MCM muon capture and phase rotation simulation program (12). This program contains numerous approximations (it uses, for instance, Gaussian approximations of scattering and straggling), but probably gives a reasonable estimate of expected performance. This work indicates that a final yield of approximately 0.004 muons per proton GeV ($\approx$ 0.1 muons per 24 GeV proton) can be obtained as follows. Muons in the decay line are initially phase rotated along a 42 m section of RF cavities with frequencies in the range 30 to 70 MHz. Fig. 5 shows a scatter plot of energy vs. time after the first phase rotation. This is followed by a mini-ionization-cooling section comprising a 3.5 m length of a hydrogen absorber in a solenoidal field, serving two purposes. It reduces the muon energies so that the subsequent drift length for a second phase rotation can be kept short. It also lowers the transverse emittance by almost a factor of two. An early, high gradient, phase rotation is required if polarization is to be selected without particle loss. Forward decays, having one polarization, yield higher energy muons than backward decays which have the other. If full

phase rotation occurred before decay then polarization and final energy are fully correlated, but significant correlation is obtained even with partial rotation before decay. The essential requirement is that significant energy changes occur before the decay. Phase rotation after decay cannot distinguish energy changes due to decay kinematics from the energy spread of the initial pions and there is no way to separate the different polarizations.

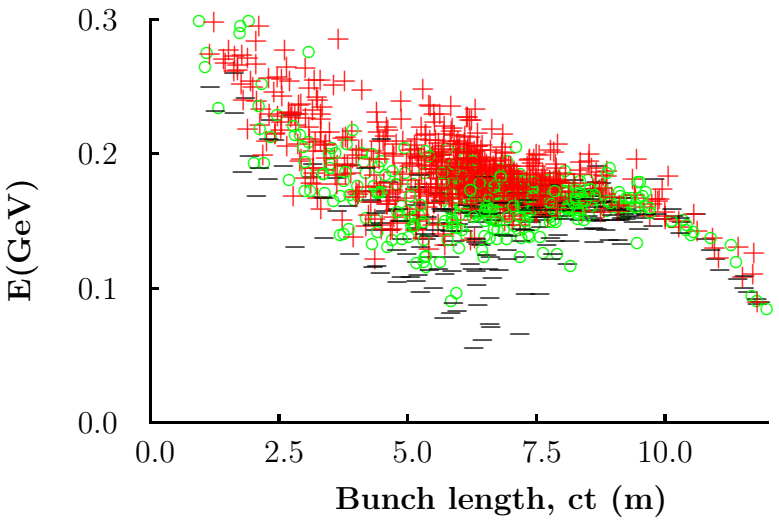

Fig. 5. This scatter plot shows a simulation of the energy vs time after the first phase rotation. The symbols indicate polarization: Red + for $P > 0.3$, green o for $-0.3 < P < 0.3$, black $-$ for $P < -0.3$.

At the end of this first stage, the bunch length has increased by a factor of 6 and the energy spread has decreased by the same amount. In the simulation illustrated in Fig. 6, the mini cooling is done in a single hydrogen absorber placed in a fixed magnetic field of 1.25 T. Such cooling introduces canonical angular momentum and it will probably be desirable to do the mini-cooling in two stages with a field reversal between them. Tab. 8 shows the parameters of the mini cooling section.

Table 8
Parameters of the mini cooling section

len	rad	B	n_μ/n_p	σ_{ct}	$< E >$	dE/E	emit
m		T		m	MeV	%	mm
3.5	.3	1.25	.33	1.8	80	100	11

Then a drift length of 160 m is used to establish the momentum correlation needed for a second stage of phase rotation. At this point the bunch length is many meters long (80 to 100 m), with a strong correlation between position and momentum. A bipolar 50 MV induction linac (an alternative RF capture and phase rotation scheme is also being considered) is used to remove this correlation, yielding an rms momentum spread of approximately 2% and

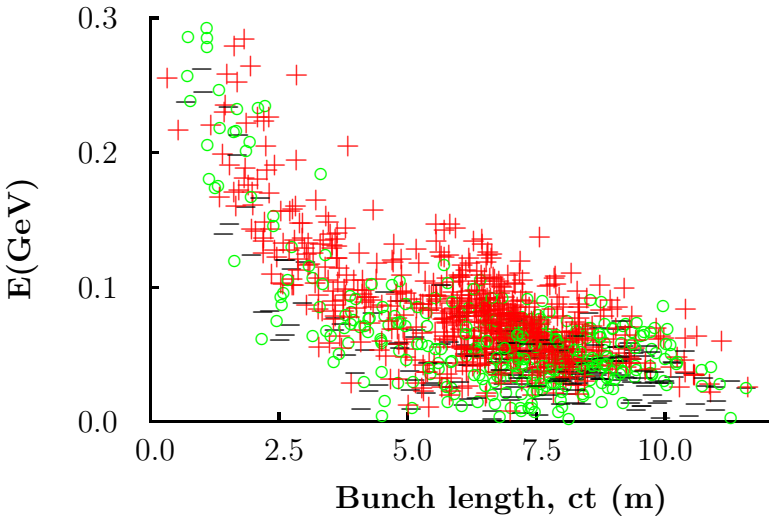

Fig. 6. This scatter plot shows the energy vs time after the mini-cooling.

transverse emittance 10 mm. There are approximately 0.008 (0.2 muons per 24 GeV proton) selected muons per GeV proton at this point. Fig. 7 shows the scatter plot of the energy vs. time distribution before and after the second phase rotation.

These muons are selected for capture into a 175 MHz linac with 40 bunches per original proton driver bunch. This process has not yet been included in the simulation and so we infer a capture efficiency in the region of 70% from the work of Neuffer et al. (13) (14). After capture the muon yield is estimated to be 0.006 muons per GeV proton, the rms momentum spread $\sigma_e \approx 6\%$, the rms bunch length $\sigma_z \approx 0.17$ m, and the transverse emittance still 9 mm, requiring an acceptance of the order of 80 mm. With only 15 mm transverse acceptance in the downstream modules we have the expected bottleneck. This can only be overcome by adding a stage of transverse ionization cooling, estimated to require 80 m of alternating liquid hydrogen and RF cells comprising 0.8 m length hydrogen and 175 MHz RF at 15 MV/m. Simulations using ICOOL (15) predict a reduction in the transverse emittance by a factor of three at the expense of an increase in the momentum spread from 6% to 12%. Further work on this system is expected to achieve the remaining factor of 2 without further increase in the momentum spread. With $\sigma_e \approx 12\%$ and $\sigma_z \approx 0.17$ m, the rms bunch area is is about 0.0068 eVs. This is larger than the assumed bunch area in the μSR, i.e. 0.0026 eVs, by a factor of three. It does not seem impossible to redesign the μRLAs and μSR for a larger longitudinal acceptance. The performance gain versus cost of this crucial stage can be reliably estimated after some development work, partly in collaboration with the manufacturers of high-power RF sources in the frequency range from 30 to 800 MHz. This industrial collaboration has already begun as a part of the MCC.

A initial cost-saving strategy could be to add a second mini-cooling stage in lieu of the full ionization cooling section, with the latter planned as a subsequent performance upgrade.

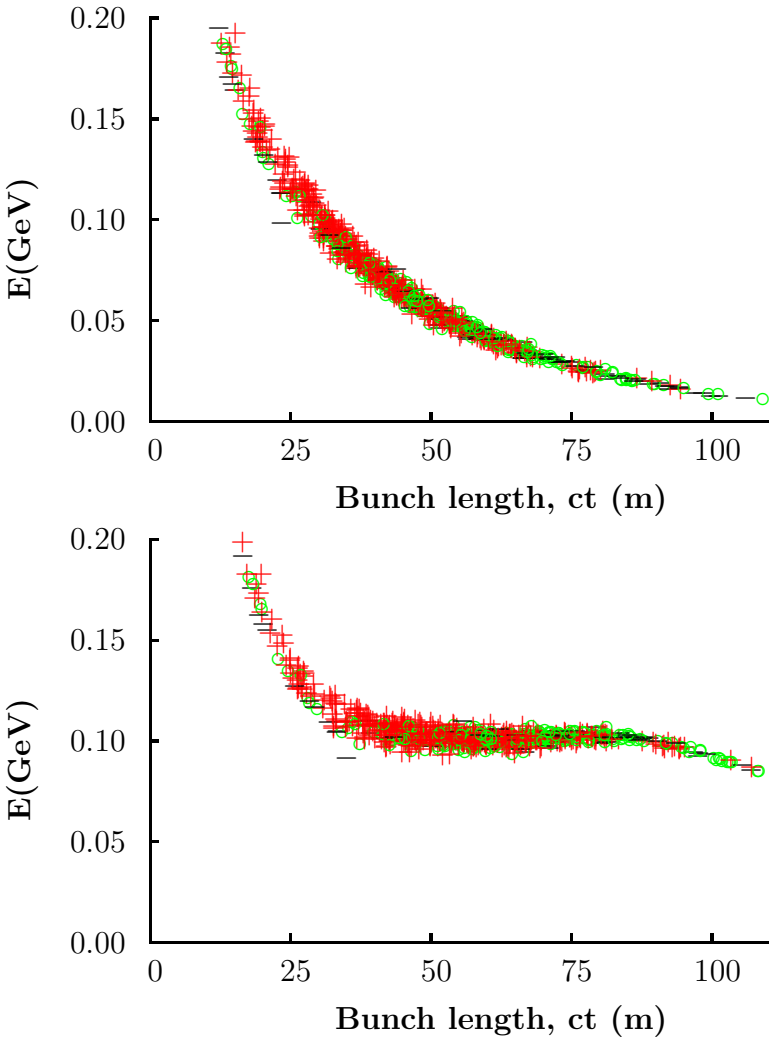

Fig. 7. These scatter plots show the energy vs time before (above) and after (below) the energy correction that effects the second phase rotation

3.6 Target Station

The target station is the only system in the facility that must be constructed to the final design specification since the extremely high radioactivity levels will preclude any easily staged upgrade. The target station comprises a liquid mercury jet – a relatively simple pumped mercury line producing a continuous jet for proton repetition rates greater than about 25 Hz. Otherwise, below 25 Hz the jet could be pulsed, although the added complication of pulsing and synchronising the jet may well be considered unnecessary. Around the target is a tungsten radiation shield that absorbs a large part of the incident beam power. This must be cooled, probably by a primary circuit of mercury so as to maintain a high effective density of the shield. In the forward direction, where the majority of the high-energy secondaries are stopped (90% of the beam power) one option is to install a mercury beam dump that would be engineered

in a way similar to the spallation or transmutation targets of other multi-megawatt proton machines. The mercury would be circulated in the same way as the target jet for cooling via a heat exchanger. It may be advantageous for this dump to be in the form of an open trough of mercury in the lower part of the tungsten shield with the primary proton beam directed downwards towards the trough. This avoids the need for a dump window, but necessitates a window downstream from the target section to exclude mercury vapour from the first phase rotation RF cavities. We also consider the possibility of reprocessing the mercury by distillation to remove non-volatile spallation products. This would avoid the buildup of very high levels of induced activity and limit radiation damage to the dump – both major problems at the intended beam power levels. The reprocessing could be done in batch mode or continuously in a bypass line.

A full engineering study of the target station is needed to complement the target dynamic and magneto-hydrodynamic studies that have begun as a part of the Muon Collider Collaboration (3). A reasonable assumption would be an initial beam power level in the region of 1 MW rising to 5 MW (possibly 20 MW) with proton driver upgrades.

3.7 Proton Driver

The muon flux into a given acceptance and momentum band has been shown to be proportional to the beam power on the target (16). Our design for the maximum muon pulse length ensures that the ratio between instantaneous and average proton beam power is as low as possible. This scheme favours a fast cycling synchrotron with as high a repetition rate as can reasonably be achieved. The bunch separation must be more than twice the bunch length after the second phase rotation (in our example 135 m) to allow time for the induction linac pulse to rise and fall. This also ensures that adjacent bunches do not overlap in the second phase rotation section.

The total six dimensional emittance of the produced muons depends on, among other things, the pion bunch length, and thus on the proton bunch length σ_p if that length is longer than a length $c\ dt_{\rm decay}$ that is characteristic of the decay process:

$$dt_{\rm decay} = \frac{(m_\pi - m_\mu)}{m_\pi}\ \frac{1}{\gamma_\pi^2}\ \tau_\pi \qquad (3)$$

Here τ_π is the pion life time and $m_\pi\gamma_\pi$ is the pion energy. For a maximum E_π of 300 MeV gives $dt \approx 1$ nsec. This, if the proton energy is low, can imply a

large tune shift in the proton ring prior to extraction.

$$\Delta\nu \propto \frac{n_p C}{\sigma_t \epsilon_n \gamma_p^2} \propto \frac{n_p}{< B > \sigma_t \epsilon_n \gamma_p} \tag{4}$$

Here C is the circumference, $< B >$ is the average bending field, and ϵ_n is the transverse emittance of the protons. The above dependency favours a higher proton energy.

It also favours a high repetition rate with relatively fewer protons per bunch, but once again the situation is complicated. The total six dimensional emittance of the produced pions depends also on the number of proton bunches employed to fill the storage ring. This favours a small number of large proton bunches in the driver, and thus a larger tune shift.

However a high driver repetition rate with smaller numbers of protons per fill would not increase the emittance per fill and would still reduce the tune shift. The difficulty with this approach is that the higher repetition rate increases the wall power required for the pulsed rf needed for acceleration and cooling.

Note that a proton synchrotron accelerating 10^{14} protons to 20 GeV at a repetition rate of 3 Hz delivers a proton beam power of approximately 1 MW. Thus an upgraded CPS or AGS could plausibly achieve our design goal.

4 Conclusions

The design of neutrino factories is a very new and rapidly developing activity. We stated in the introduction that our present work is incomplete and the reader will certainly have appreciated that by now.

Our goal was assembling a neutrino factory from modules that "fit together", such that a module accepts a beam with parameters as delivered by the previous module. We have achieved this goal for two rather long strings of modules, the first string ranging from targetting, through capture, phase rotations, cooling, etc., and the second string ranging from the recirculating linear muon accelerators to the muon storage ring proper (I don't mention the linear accelerators here). We missed our optimistic goal at the boundary between the first and the second string of modules: The beam emittances at the output end of the first string are larger than those accepted at the input of the second string by factors of order three in all three degrees of freedom. To remove this bottleneck, we could either reduce the emittances by further cooling, or increase the acceptances of the downstream string, as already mentioned.

We have not carried out extensive parameter searches for the modules. Instead, we concentrated on one design, basing our choice on experience, intuition, prejudices, and simulation. We do not claim that our choices are anywhere near to the optimum for either the individual modules, varying their internal parameters, but leaving the beam parameters at either end unchanged, or for the whole neutrino factory by also redefining the module boundaries. Nevertheless, our work should be a firm basis for such quantitative optimizations in the future.

A very desirable improvement on the present proposal would be, as mentioned in the text, the replacement of the induction linac by an RF system for the second phase rotation. This "bunched beam phase rotation" requires that the frequencies of the bunching and energy correction cavities need to track the changing bunch spacing (arising from the velocity spread) that will occur prior to the completion of the energy correction. Such a system would avoid the greater length and expense of the induction linac, but it has yet to be studied in detail.

References

[1] S. Geer, *Neutrino beams from muon storage rings: Characteristics and physics potential*, Phys. Rev. **D 57** (1998) 6989, and erratum Phys. Rev. **D 59** (1998) 039903.

[2] B. Autin; A. Blondel; J. Ellis, eds. *Prospective Study of Muon Storage Rings at CERN*, CERN-99-02, (1999).

[3] Muon Collider Collaboration, *Status of Muon Collider Research and Development and Future Plans*, Phys.Rev. ST Accel.Beams **2**, 081001 (1999), also BNL-65-623, Fermilab-PUB-98/179, LBNL-41935 (1999).

[4] E. Keil, *Muon Storage Ring Design with Simple* Mathematica *Packages*, CERN SL/99-053 AP (1999).

[5] E. Keil, *Single-Particle Dynamics – Linear Machine Lattices*, CERN 77-13 (1977) 22.

[6] H. Grote and F.C. Iselin, *The MAD Program Version 8.16, User's Reference Manual*, CERN-SL-90-13 (AP) (1990)

[7] A. Blondel, *Muon polarisation in the neutrino factory*, these proceedings.

[8] G. Geschonke and E. Keil, *A Recirculating Electron Accelerator (ELFE) using the LEP Super-Conducting RF Cavities*, CERN SL/98-060 (RF) (1998).

[9] LEP 2 Design Team, *LEP Design Report, Vol. III: LEP 2*, CERN-AC/96-01 (LEP 2) (1996).

[10] H. Burkhardt (ed.), *ELFE at CERN*, to be published

[11] F. Hinode ed., *ATF Accelerator Test Facility: Design and Study Report*, KEK Internal 95-4 (1995).

[12] R.B. Palmer and J.C. Gallardo, *MCM Fortran program to simulate the front end and cooling section of a muon collider*, unpublished.

[13] S. Geer, C. Johnstone, D. Neuffer, *Muon Storage Ring Neutrino Source: The path to a muon collider?*, Fermilab TM-2073 (1999).

[14] S. Geer, C. Johnstone, D. Neuffer, *Design Concepts for a Muon Storage Ring Neutrino Source*, FERMILAB-Pub-99/121 (1999)

[15] R.C. Fernow, *ICOOL: A Simulation Code for Ionization Cooling of Muon Beams*, Proc. 1999 Particle Accelerator Conference (New York, 1999) 3020.

[16] N. Mokhov, *Pion Yield from a Mercury Target versus Proton Energy*, These proceedings, (1999).

A Parameter List for Neutrino Factory

Phase Rotation

	len m	rad	B T	f MHz	grad MV/m	n/n	σ_{ct} m	dE/E %	emit mm
target	.45	.075	20			.66	.3	100	-
match	3	.2	2						
rf	3	.3	1.25	60	8				
rf	33	.3	1.25	30	5				
rf	7.5	.3	1.25	45	7	.35	1.8	15	20
cool 1	3.5	.3	1.25			.33	1.8	15	11
drift	170	.37	.6			.28			
Ind Lin	56	.25	1.25		1	.28	20	5	10
Selected						.2	13	3	9

Polarization = .48 (peak value of .65)

Momentum spread = 2 %

Bunching (not simulated)

freq	175 MHz
λ	1.7 m
bunches	40
σ_z	17 cm
rms dp/p	6 %
emit	9 mm

Cooling

freq	175 MHz
λ	1.7 m
B	2 T
Grad	15 MV/m
Cell	5 m
H2	.8m
cells	16
length	80 m
σ_z	20 cm
rms dp/p	6→12 %
emit	9 →3 mm

Acceleration

min E for full sc cells 7.5 GeV min E for 1/4 sc cells 1.8 GeV

		Lin 1	Lin 2	Recirc 1	Recirc 2
p	GeV/c	.1-.7	.7-2.0	2-8.2	8.2-30
freq	MHz	175	350	350	350
Grad	MV/m	15	10	10	10
Δ p	GeV/c	.6	1.3	1.5	5.5
n		1	1	4	4

Muon Budget

	Factor	μ/24 GeV proton
Muons after Match (below 1 GeV)		0.66
Muons after Phase Rotation #1 (selected)	0.45	0.3
Muons after Phase Rotation #2 (selected)	0.7	0.21
Muons after RF Capture	0.7	0.15
Muons after Cooling	0.9	0.13
Muons after Acceleration	0.7	0.092

Polarization vs. position along the bunch train