Muon Collider Note 46
Version 9/3/99

Changes after 9/3/99 in bold face or listed here

Draft Parameters of a Neutrino Factory [1]

R.B. Palmer (for the collaboration)

Abstract Preliminary parameters of neutrino factories are given. Alternative parameter options are discussed. Where appropriate, the site specific parameters are given for both BNL and FNAL locations. Event rates are given for detectors at Gran Sasso, Soudan and SLAC.

Contents

[1] http://pubweb.bnl.gov/people/palmer/nu/params.ps

1 Introduction

This document is not intended as an actual parameter list for a neutrino factory. Too little simulation work has been done for that. The document is intended merely as a starting point from which we can work.

The scenario given is based on the example discussed in the scenarios working group at Lyon, which was, in turn based on the work of Palmer, Johnson and Keil (PJK); and this in turn owed much to the work by Geer et al on a lower energy neutrino beam, and to work done by the Muon Collider Collaboration. I list below some URL's where these sources can be found.

Muon Collider status report
http://prst-ab.aps.org/przv02i08tc.html
http://www.cap.bnl.gov/mumu/status_report.html

Muon Collider public web
http://www.cap.bnl.gov/mumu/

Geer et al
http://www.fnal.gov/projects/muon_collider/nu/nu.html
clic on TM 2073 for earlier 10 GeV ideas
clic on hep-ph/9906487 for physics and rates

PJK-ver1
http://nicewww.cern.ch/c̃dj/public/nu-fact/pjk-ver1/

A draft of a paper on polarization by Blondel:
http://alephwww.cern.ch/b̃dl/muon/nufacpol.ps

1.1 Scenario

The following figure illustrates the main components of the scenario described here.

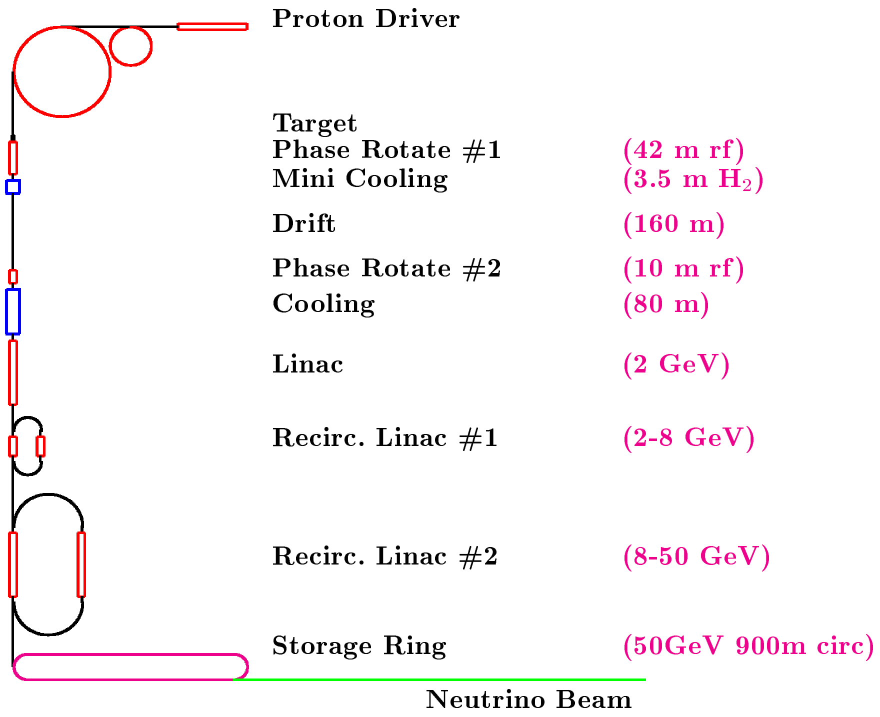

2 Components

2.1 Driver

The number of pions produced with an optimized system has been shown (see above figure from Mokhov) to be closely linear with the proton energy. Thus the number of pions, and muons into which they decay, is proportional to the proton power. This might suggest that the proton energy could be selected arbitrarily, but the situation is more complicated.

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 $c\ dt_{decay}$ that is characteristic to the decay process:

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

Where 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.

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

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

It also favors 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 favors 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.

Parameters for the BNL and FNAL cases, in different phases of their development, might be:

Energy	GEV	24	24	16	16
Power	MW	1	4	1	4
Repetition	Hz	2.5	5	15	15
p's/fill		10^{14}	$2\ 10^{14}$	$2.5\ 10^{13}$	10^{14}
bunches		6	6	4	4
circ.	m	807	807	474	474
spacing	m	135	135	118	118
sigma t	nsec	1	1	1	1

2.1.1 Options:

- Space bunches closer or further apart.

  Motive:

  Further spacing increases possible phase rotation, reduces dp/p, reduces needed gradient in capture and cooling rf.

  But further spacing increases the charge per bunch and thus the space charge tune shift. Further spacing also increases the length of phase rotation drift and thus decay loss.

Other drivers

A comparison with parameters using other possible drivers is:

		JHF Synch.	BNL/CERN Synch.	FNAL Synch.	CERN Linac
Proton E	GeV	50	24	16	4-2
Rep Rate	Hz	2.5	5	15	50
Proton Power	MW	4	4	4	4
Protons/pulse	10^{14}	2	2	1	1.3-2.6
Bunches		16	8	6	8
Protons/bunch	10^{13}	1.2	2.5	1.3	1.5-3
Relative Tune Shift		.25	1	.5	30-60*
Power: PR&Cool	MW	3	6	20	60
$n_\mu/(n_p\ E_p)$	GeV^{-1}	0.004	0.004	0.004	0.004
Duty Cycle	%	0.16	.3	1	3
$\rightarrow\ \mu^+\mu^-$		YES	YES	YES	NO**

Possible solutions:

*	External Bunch Compression to avoid Tune Shift in ring
**	Multiple Accumulators aimed at same Target

2.2 Target and Capture

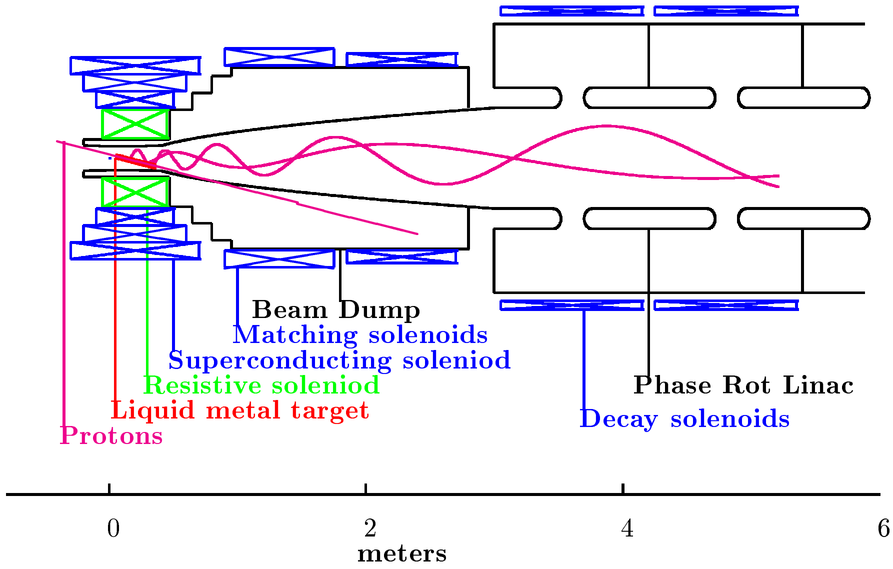

The target requirements are very similar to those for the muon collider, except that the instantaneous shock heating is somewhat less because of the larger number of proton bunches used. It is assumed here that the liquid mercury jet solution is used.

The capture solenoid is also assumed to be the same as in the muon collider status report.

material		Hg
radius	mm	7
length	m	0.45
angle	rad	0.15
solenoid field	T	20
solenoid radius	cm	7.5
maximum captured transvers mom.	MeV/c	225

2.2.1 Options:

- decrease jet radius

  Motive:

  reduce eddy current effects and total Hg mass flow

  But the efficiency will fall

- Use a moving band solid nickel target

  The somewhat lower density of the target will change the required optimization.

2.3 Phase Rotation #1

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 occured 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.

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.

The following table gives parameters obtained from a simulation using MCM with an initial proton energy of 24 GeV. One can assume that for an initial energy of 16 GeV, the ratio of muons to protons would be 2/3 of the values shown.

	len	rad	B	f	grad	n_μ/n_p	σ_{ct}	$<E>$	dE/E	emit
	m		T	MHz	MV/m		m	MeV	%	π 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	180	15	20

At the end of this first rotation, the simulated energy vs time distribution is shown below.

Color and symbols indicate Polarization:
red + P > .3
green o .3 > P > -.3
black − -.3 < P

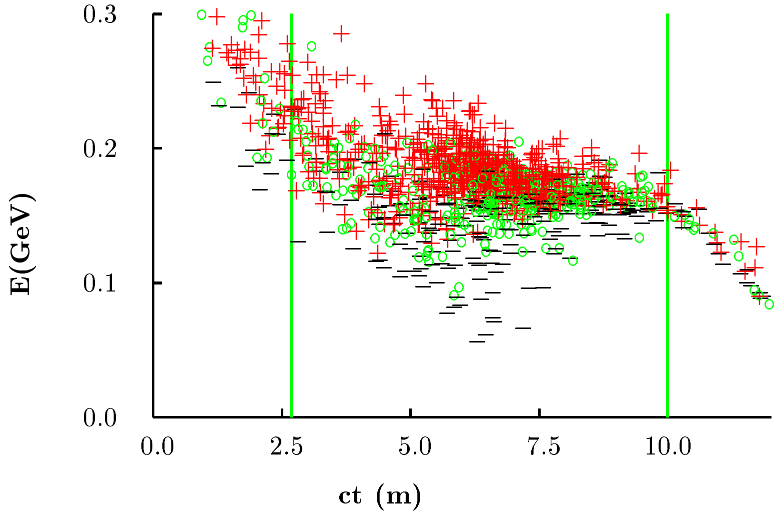

2.3.1 Options:

- This first stage of phase rotation can be eliminated, but the polarization separation will be lost.

2.4 Mini Cooling

Mini-cooling, which consists of a hydrogen absorber in a solenoidal field, serves two purposes. It reduces the muon energies so that the subsequent drift length for a second phase rotation could be kept short. It also lowers the transverse emittance by almost a factor of two.

In the current simulation, the mini cooling was 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.

	len	rad	B	f	grad	n_μ/n_p	σ_{ct}	$<E>$	dE/E	emit
	m		T	MHz	MV/m		m	MeV	%	π mm
mini-cooling	3.5	.3	1.25			.33	1.8	80	100	11

The following figure shows the energy vs time after the mini cooling.

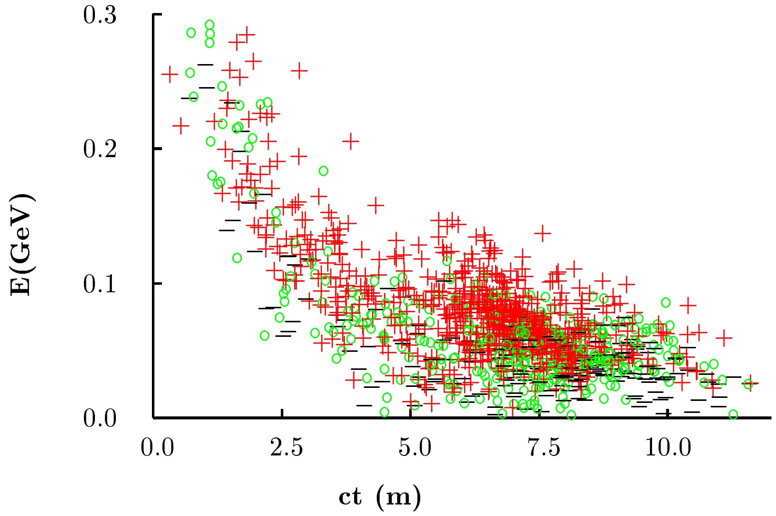

2.4.1 Options:

- Increase energy at the end of phase rotation #1 to increase mini cooling.
- Decrease gradients (for technical reasons) resulting in lower final frequencies or lower efficiency.
- Increase gradients (if possible) resulting in higher final frequencies or higher efficiency. Maybe better polarization.
- Less mini cooling, or a higher initial energy would reduce losses, but increase the needed length of drift in the second phase rotation.

2.5 Phase Rotation #2

The object is to minimize the momentum spread at the expense of lengthening the bunch up to a distance approaching the initial proton bunch spacing (approx 135 m in this case). The very long resulting bunch is then re bunched at a higher frequency (175 MHz), yielding a train of about 30 individual muon bunches for every initial proton one.

In addition, this rotation results in the polarization being correlated with time, ie bunch number, instead of energy. The polarization is now preserved.

The rotation is performed by a drift (160 m) followed by energy correction followed by bunching.

In this case, an induction linac is used (80 MeV acceleration) in which the pulse shape is tailored to correct the time-energy correlation generated by the drift. The induction linac must supply a rapid train of accelerations, spaced by the proton bunch spacing, and equal in number to the number of proton bunches (see above under driver).

The following figures show the energy vs time before and after the energy correction.

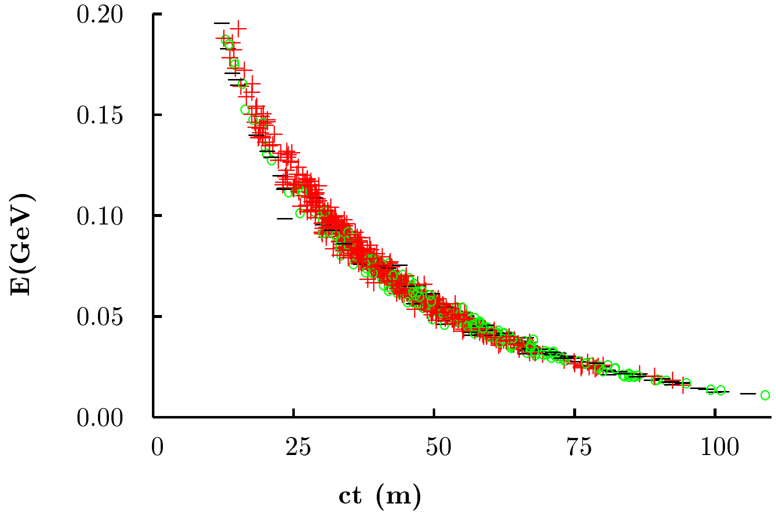

After Acc to 100 MeV

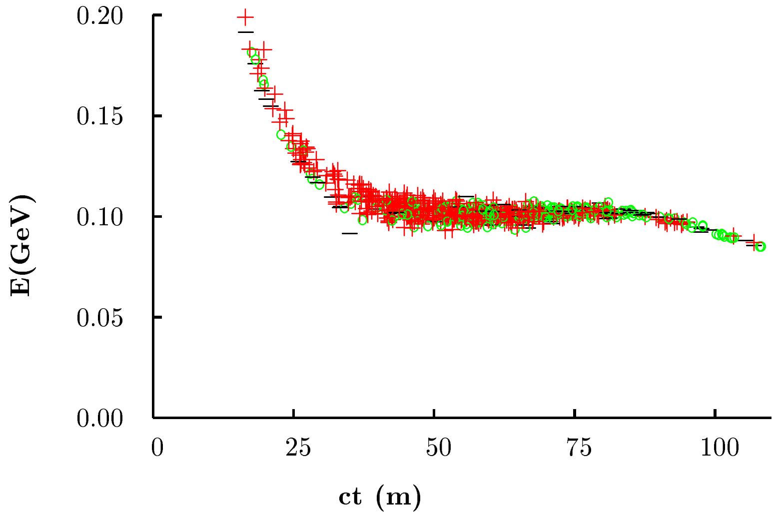

The simulated final polarizations vs bunch position are:

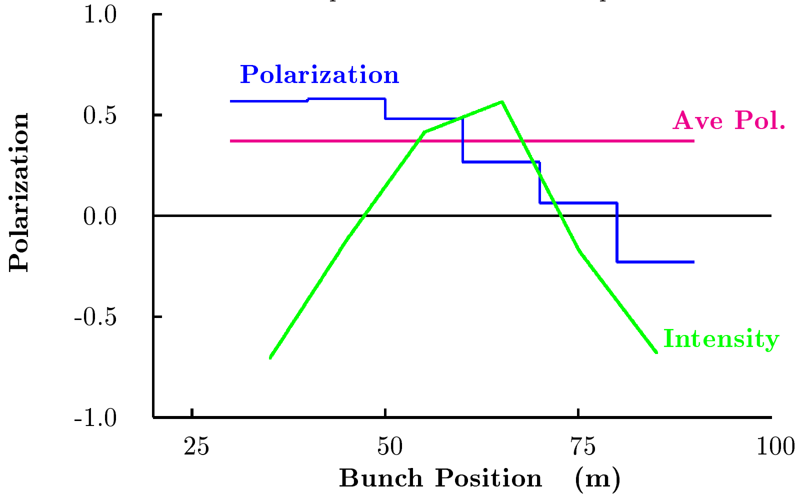

The sensitivity of the maximum polarization to the initial proton bunch length is:

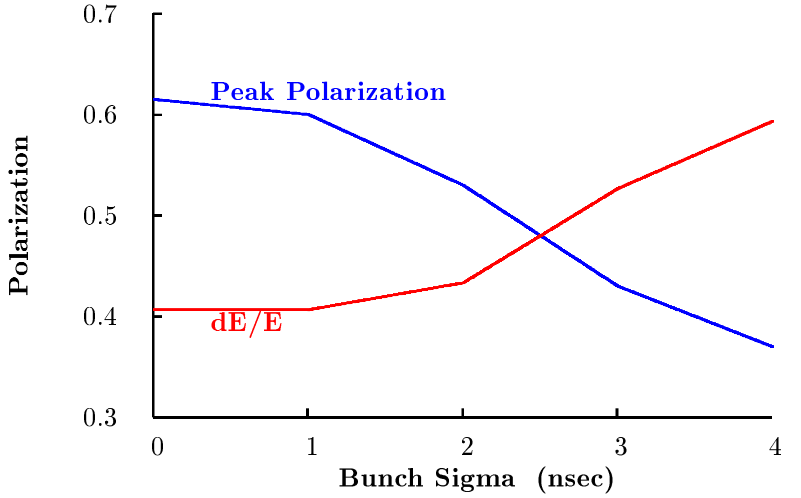

	len	rad	B	f	grad	n_μ/n_p	σ_{ct}	$< E >$	dE/E	emit
	m		T	MHz	MV/m		m	MeV	%	π mm
drift	160	.37	.6			.28				
Ind Lin	56	.25	1.25		1	.28	20		5	10
Selected						.2	13	100	3	9

Average Polarization = .37 (peak value of .6 for 20%)
dp/p ≈ 2 %

2.5.1 Options:

Increase drift to do more phase rotation, reduce the bunched dp/p, reduce required rf gradients in bunching and cooling.

Decrease drift to allow more, but smaller, bunches in driver (see above on driver bunch spacing)

2.6 Bunching (not yet simulated)

As noted above, this can be done either before, or after, the energy correction. It must be noted that quoting the resultant momentum spread by its rms width is somewhat inappropriate: the bunch after rotation will be far from Gaussian. The important parameter is the efficiency of capture of the resulting bunch by the following cooling stage. This is yet to be determined.

The frequency used here is chosen to be a multiple of 350 MHZ, the frequency of the LEP sc cavities that are assumed to be used in the later acceleration.

frequency	175 MHz
λ	1.7 m
bunches	≈ 40
σ_z	9 cm
rms dp/p	8 %
emit	9 π mm

Note: these numbers are only estimates.

2.6.1 Options

- Lower frequencies would allow the use of higher harmonics to improve bunching efficiency.

  If more mini-cooling could be used, the beam radii would be less and higher harmonics might be usable even with 175 MHz.

- bunching followed by energy correction with rf

  This "bunched beam phase rotation" requires that the frequencies of the bunching and energy correction cavities need to track the changing bunch spacing that will occur prior to the completion of the energy correction. Such a system would save the greater length and expense of the induction linac, but has yet to be simulated.

- induction linac energy correction followed by a second mini-cooling and then bunching with acceleration.

Neuffer obtains more efficient bunching if the initial energy is lower and the bunching is done together with acceleration. This suggests that a second minicooling with about 1 m of hydrogen could be used to lower the muon energies to about 25 MeV followed by the bunching and acceleration back to 100 MeV. This is being studied by Fukui.

The addition of the second mini-cooling will further reduce the required conventional cooling to follow.

2.7 Cooling

A simple comparison of the total produced six dimensional emittance and the total acceptance of a plausible storage ring indicates that cooling should not be needed. But without cooling, the muon accelerator would have to have a transverse rms acceptance of 18 π mm rad (full acceptance $\approx$ 0.18 π m rad). This we have shown is possible with large aperture solenoid focusing and low frequency rf, but would be expensive.

A more reasonable acceleration scheme has been studied by Eberhard Keil, that has a transverse an rms transverse acceptance of 1.5 π mm radians. The following cooling is designed to match into this scheme.

The current simulation, using a fixed lattice, cools to below 3 π mm radians, but does not achieve the required 1.5 π mm radians, because coulomb scattering at the end. Other lattices, with stronger fields easily reach the lower requirement but do not accept the full initial emittance. We will thus require two different lattices with suitable matching between them, or a lattice in which the field and resulting beta functions can be smoothly varied along the cooling channel.

Coil Blocks

Cell length = 2.5 m

start	dl	radius_1	r_2-r_1	I/A
m	m	m	m	A/mm^2
0.280	0.400	0.700	0.150	88.48
1.100	0.300	0.700	0.150	-38.21
1.820	0.400	0.700	0.150	88.48

Hydrogen Absorber

length	36 cm	
Al Window:	50 micron,	20 cm rad
H_2-rf gap:	2 cm	

rf

5 RF pillbox cavities

lengths 42 cm

freq 175 MHz

grad 15 MV/m
windows 250 micron Be, 30 cm rad

Note: The window thicknesses are arbitrary

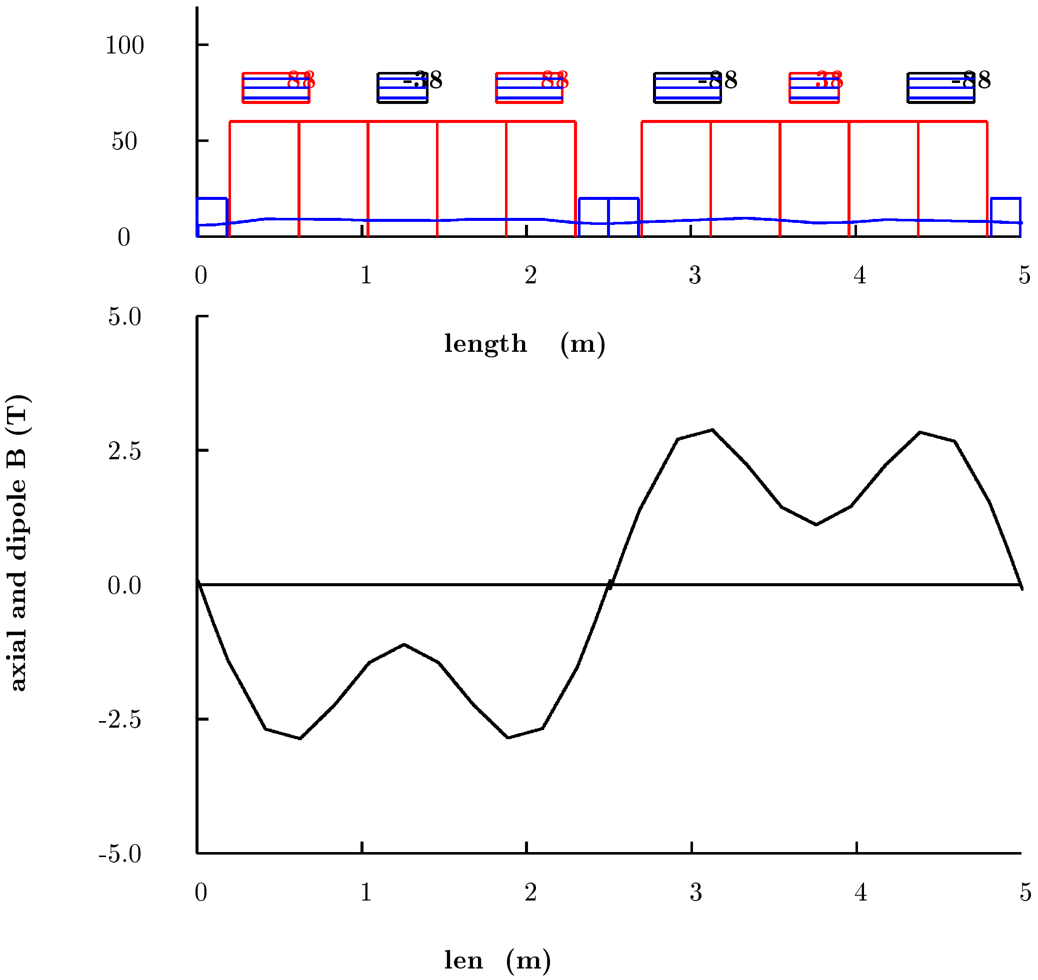

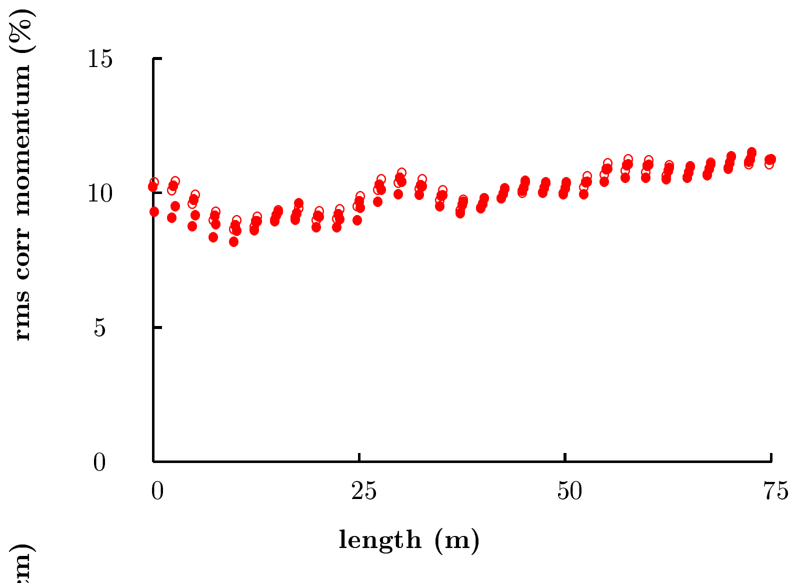
rms corr momentum (%)
15
10
5
0
0
25
50
75
length (m)

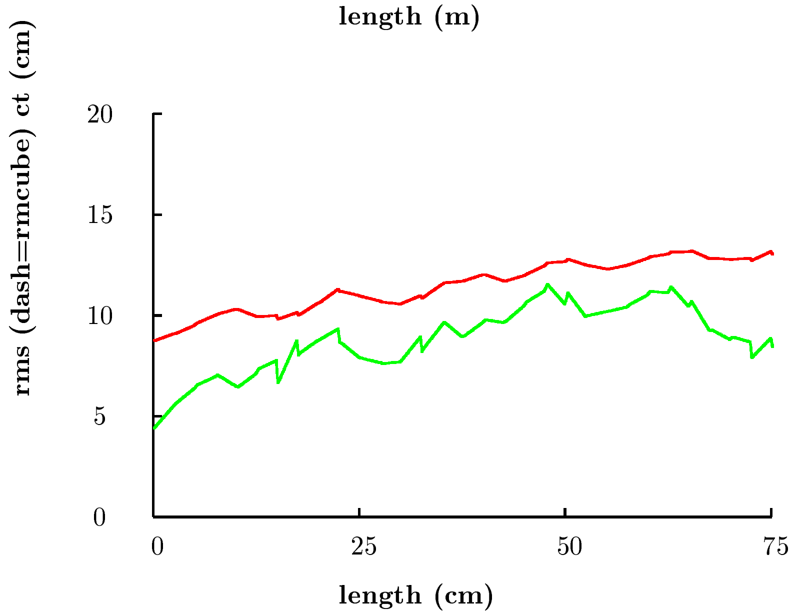
rms (dash=rmcube) ct (cm)
20
15
10
5
0
0
25
50
75
length (cm)

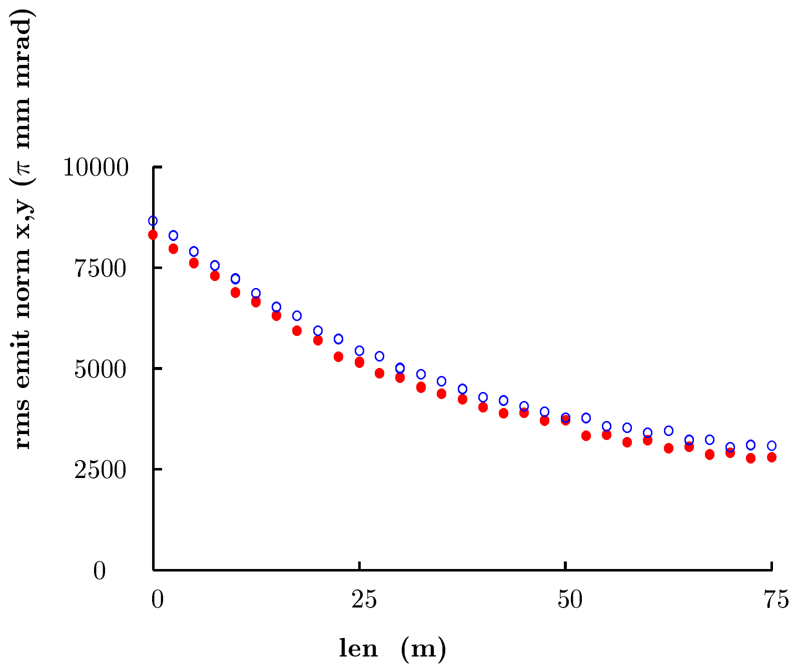
rms emit norm x,y (π mm mrad)
10000
7500
5000
2500
0
0
25
50
75
len (m)

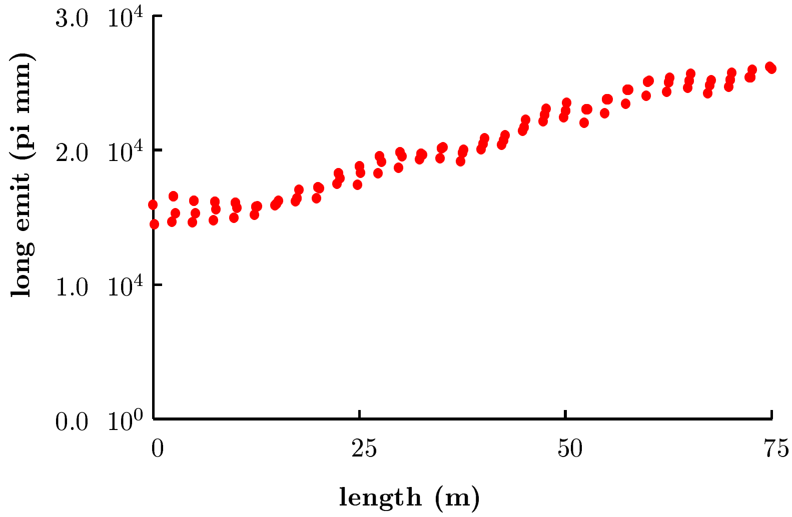
long emit (pi mm)
3.0 10^4
2.0 10^4
1.0 10^4
0.0 10^0
0
25
50
75
length (m)

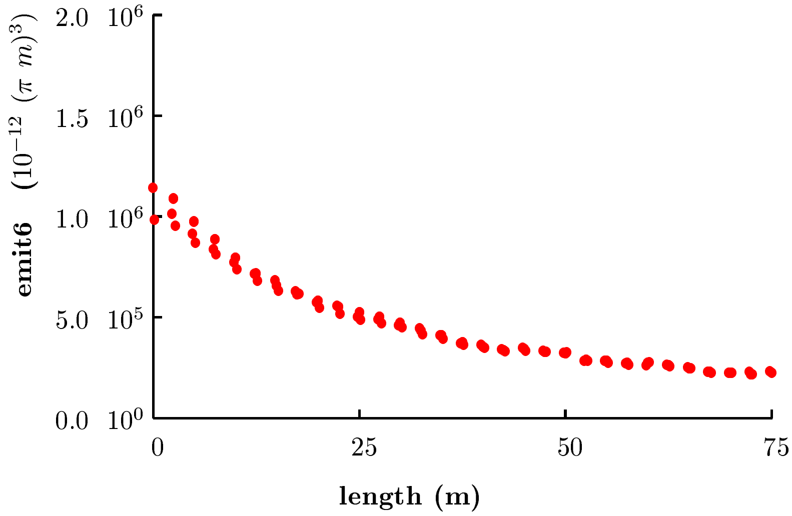

2.0 10^6
1.5 10^6
1.0 10^6
5.0 10^5
0.0 10^0
emit6 $(10^{-12}\ (\pi\ m)^3)$
0
25
50
75
length (m)

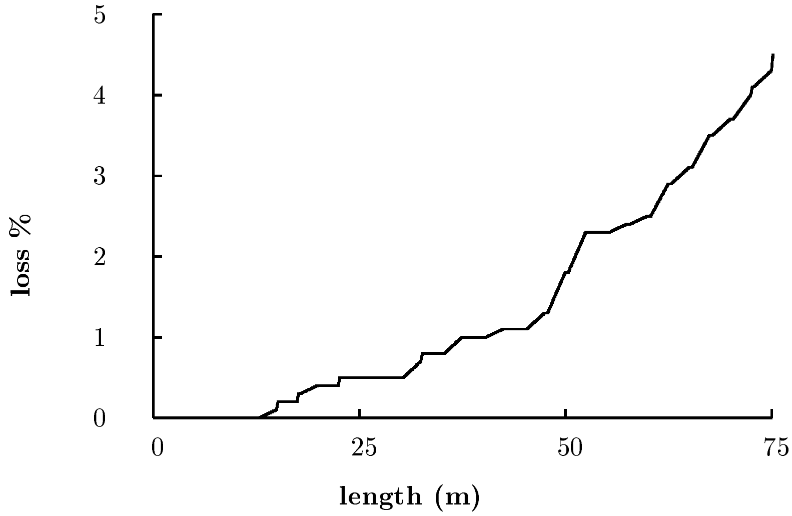

5
4
3
2
1
0
loss %
0
25
50
75
length (m)

momentum	190 MeV/c
frequency	175 MHz
λ	1.7 m
Maximum B	3 T
Axial field gradient	15 MV/m
Cell length	2.5 m
Liquid hydrogen absorber length	.36m
number of cells	30
total length	80 m
σ_z	9 $\rightarrow$ 12 cm
rms dp/p	8$\rightarrow$11 %
emit$_\perp$	9 $\rightarrow$1.5 π mm
emit$_\parallel$	15 $\rightarrow$28 π mm

2.7.1 Options:

- Less cooling to final emittance 3 π mm rms, can be done (and has been simulated) with a single kind of cell (no field grading or matching into a second lattice to get to the lower emittance).
- Cooling at higher energy to reduce grad requirement ? But length will increase rapidly.
- Cooling at lower frequency to reduce grad requirement
- Cooling from lower initial dp/p (by using more phase rotation) to reduce grad
- Use a second mini-cooling and drastically reduce or eliminate this conventional cooling

2.8 Amplitude-energy correlation

If the phase rotation drift and buncher have a lattice with the same amplitude-z velocity properties as the cooling lattice, then the correlation is automatically generated. Remember: the drift sorts particles by their z velocity, not their energy. After the drift, their longitudinal position is a function of that z velocity which is the required correlated combination of energy and amplitude. The bunching, done in the same lattice (or one with the same properties) is also a bunching by z velocity, not energy, so the correlation is preserved. And so into the cooling. Note that a simple solenoid will NOT do for the drift or bunching, since the z velocity is a function not only of amplitude, but also of angular momentum. A solenoid of one sign gives a higher z velocity for one angular momentum sign than the other. Both drift and bunching must done with alternating fields of some kind. I would suggest that for out studies we actually use the same super fofo lattice as in the first cooling stage. Then all should be well. Whether we will need to do something clever to match into the start of the accelerator is another question.

2.9 Acceleration

Acceleration from 186 MeV/c ($\approx$ 100 MeV) to 50 GeV is achieved by a linac followed by two recirculating linear accelerators (RLA's).

The following parameters are taken directly from preliminary work by Eberhard Keil for the Lyon Nufac99 workshop. They will be updated as better numbers are available.

The assumption made is that the larger second (and possibly also the first) recirculating accelerator uses LEP superconduction cavities, or cavities with the same parameters and dimensions. The use of these cavities sets constraints on the minimum energy that allows the required emittance to be transported in them. If the full (approximately 10 m long) cryostats, containing four cavities, are used as is, then this minimum energy is approximately 7.5 GeV. This then is taken as the approximate injection energy into the second RLA. If the cavities are used are rehoused individually in new cryostats, then the minimum energy is approximately 1.8 GeV. This is uses as the approximate injection energy the first RLA.

		Linac 1	Linac 2	RLA #1	RLA #2
E	GeV	.1-.6	.6-2	2 -7.5	7.5-50
freq	MHz	175	350	350	350
Grad	MV/m	15	10	10	10
Δ E/turn	GeV	.5	1.3	1.1	3.2
n		1	1	4	6

2.9.1 Options

- If a system could be designed, presumably using special rf cavities, that could accept larger transverse emittance (eg 3 π mm radians) then the needed cooling could be reduced.

2.10 Storage Ring

2.10.1 Geometries

The geometry of the storage ring is site specific, being a function of bot the ring and detector locations. The following map shows directions and direct distances from rings at BNL or FNAL to Gran Sasso, Soudan and SLAC.

For physics reasons (to separate MSW from vacuum oscillations), two differing ring to detector distances are required. If the two detectors lie in approximately opposite directions from the ring then it seems reasonable to design the ring with long sides that point to the two detectors, adding, if needed, a third straight to close the ring. Two geometries are of particular interest:

- a triangular geometry lying in a tilted plane. This minimizes the amount of bending required and maximizes the total straight for a given circumference. But, as Blondel has noted, the lengths of the straights pointing at the two distant detectors is NOT maximized.

- a bow tie, or figure of eight geometry, also lying in a plane. This geometry uses more total bending, but does maximize the important straights. It also has the interesting feature of not precessing the spins. A variant of the bow tie looks much the same but does not lie in a plane, so that there is a significant separation of the beams where they cross. In this case there is a slow precession of spin.
- The bow tie can be made assymetric so as to maximize the length of the upward straight.

The following figure shows the two geometries for the BNL to Gran Sasso and SLAC case.

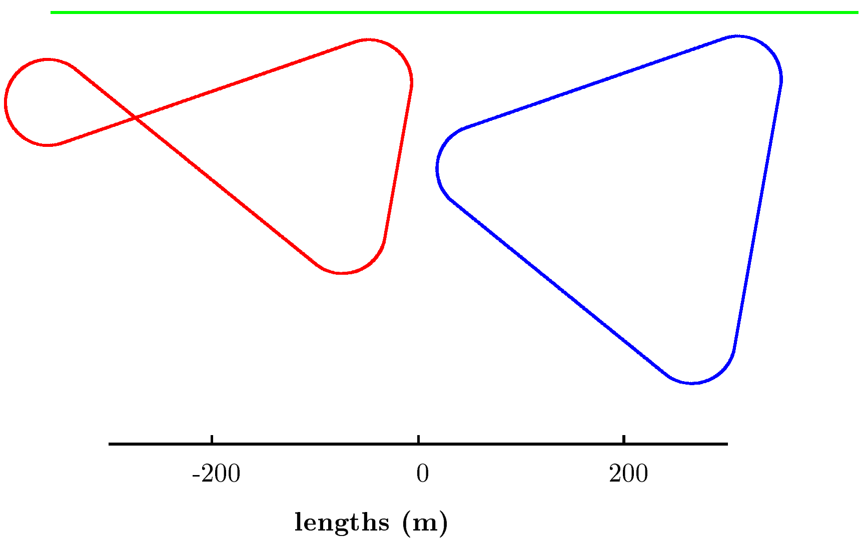

General:

location		BNL	BNL	FNAL	FNAL
Detector 1		Gran Sasso	Gran Sasso	Gran Sasso	Gran Sasso
Detector 2		SLAC	Soudan	SLAC	Soudan
Distance 1	km	6500	6500	7100	7100
Distance 2	km	4100	2000	2800	730
Dip 1	deg	30.5	30.5	34	34
Dip 2	deg	19	8.5	13	3
horizontal ang	deg	40	55	30	97
space angle	deg	63	65	55	98
tilt angle	deg	54	40	60	34

The "space angle" is 180 degrees minus the angle in space between the vectors pointing to the two detectors.

The "tilt angle" is the angle of the plane containing the two vectors (and the ring) with respect to the horizontal, and is the steepest angle that the tunnel has.

For triangular ring:

location		BNL	BNL	FNAL	FNAL
Detector 1		Gran Sasso	Gran Sasso	Gran Sasso	Gran Sasso
Detector 2		SLAC	Soudan	SLAC	Soudan
ave. bending $< B >$	T	4	4	4	4
circumference	m	1000	1000	1000	1000
total arc length	m	260	260	260	260
minor straight	m	254	260	234	317
major straights	m	243	240	253	210
fraction of circ.		.243	.240	.253	.210

For bow tie ring:

location		BNL	BNL	FNAL	FNAL
Detector 1		Gran Sasso	Gran Sasso	Gran Sasso	Gran Sasso
Detector 2		SLAC	Soudan	SLAC	Soudan
total bend angle	$\times$ 360 deg.	1.35	1.36	1.31	1.53
ave. bending $< B >$	T	4	4	4	4
circumference	m	1000	1000	1000	1000
total arc length	m	352	353	339	400
minor symmetric straights	m	64	68	54	97
minor one side straight	m	128	136	108	194
major straights	m	260	254	276	201
fraction of circ.		.260	.254	.267	.201

2.10.2 Beam Shield

The heating from muon decay in the storage ring is significant. Mokhov (see Status Report) has calculated the attenuation of decays from 50 GeV muons from 2 cm of tungsten shielding to be 0.06. If such shielding is introduced around the beam pipe in the arcs then the power dissipated in the superconducting magnet coils is estimated to be:

type		triangle		bow tie	
proton power	MW	1	4	1	4
muon power	kW	200	800	200	800
pipe heating	kW	70	280	70	280
arc heating	kW	18	73	20	80
magnet heating in arc	kW	1.1	4.3	1.2	4.8

2.10.3 Lattice

Allowing 20 % emittance growth in the accelerator in each of three directions; and requiring an acceptance of 3 sigma in each of the 6 dimensions:

energy	GeV	50
transverse rms emittance	π mm radians	1.8
transverse acceptance	π mm radians	16
longitudinal rms emittance	π mm dp/p	33
longitudinal acceptance	π mm dp/p	300

If the bunch spacing is 1.7 m (corresponding to the 175 MHz bunching used here) then a reasonable maximum rms bunch length in the collider would be 6 cm. Thus the minimum momentum spread in this case would be

$$\sigma_p \text{ (minimum)} \ = \ 0.1\%.$$

It is desirable that the beam divergence be small compared with the decay angles ($1/\gamma$). The ability, for instance, of the polarization to control the anti electron neutrino flux is dependent on this condition.

The rms beam divergence in the straight sections should probably be below 1/10 th of $1/\gamma$. This is in order a) to maximize the dependence of the ν_μ to ν_e ration on the polarization; and b) to assure that the flux observed is not significantly effected by the exact magnitude of this divergence. To achieve this low divergence, the required beta function in the major straights is:

$$\beta_{\text{major straights}} \ \geq \ 75 \qquad \text{m}$$

For the up going straight, aimed at a near detector on the surface, there is probably not such a stringent requirement on the beam divergence. If the divergence is required to be below 1/3 th of $1/\gamma$ then

$$\beta_{\text{upgoing straight}} \ \geq \ 8 \qquad \text{m}$$

2.11 Muon Budget

p driver energy (GeV)		24	16
	Factor	μ/p	μ/p
Pions after Match (< 1 GeV, forward)		0.66	0.44
After Phase Rotation #1 (selected)	0.45	0.3	.2
After Phase Rotation #2 (selected)	0.7	0.21	.14
After RF Capture	0.7	0.15	.1
After Cooling	0.9	0.13	.09
After Acceleration	0.7	0.092	.061
$n_\mu/(n_p\ E_p)$ (GeV^{-1})		.0038	.0038

2.12 Events

$$N_{\mathrm{events}} \propto N_\mu\ E_\mu^3\ L^{-2}\ m$$

For proton power= 1.5 MW, and the above muon capture efficiencies, we would,

in a 2.5 10^{20} second year, obtain 4×10^{20} muons per year decaying in the storage ring. If we take the fraction of the ring pointing to a given detector to be 0.25 (approximately as in the bow tie geometries) then the number of decays pointing to the given detector will be approximately 10^{20}.

The following table (compiled by Steve Geer) gives charged current neutrino interaction rates (per kt-year) as a function of baseline length L for an E_μ = 50 GeV muon storage ring in which there are 1×10^{20} unpolarized muon decays per year within a neutrino beam–forming straight section. The rates are listed for

a) $\nu_e \to \nu_\mu$ oscillations with $\Delta m^2 = 1 \times 10^{-4}$ eV^2/c^4 & $\sin^2 2\theta = 1$
b) $\nu_e \to \nu_\tau$ oscillations with $\Delta m^2 = 3.5 \times 10^{-3}$ eV^2/c^4 & $\sin^2 2\theta = 0.1$
c) $\nu_\mu \to \nu_\tau$ oscillations with $\Delta m^2 = 3.5 \times 10^{-3}$ eV^2/c^4 & $\sin^2 2\theta = 1$

Also listed are the rates for the un oscillated neutrino interactions, the corresponding statistical significance of the disappearance signal (numbers in parentheses), and the rates for the antineutrino interactions.

			BNL-GS	BNL-SLAC	BNL-Soudan	FNAL-GS	FNAL-SLAC	FNAL-Soudan
					L	(km)		
sign	case	mode	≈6500	≈4100	≈2000	7332	≈2800	732
μ^+	(a)	$\nu_e \to \nu_\mu$	5×10^{-2}	0.86	1.5	3×10^{-5}	1.3	1.6
		$\nu_e \to \nu_e$	1500	3800	16000	1200	8200	1.2×10^5
		sig/bkg	.00003	.0002	.0001	$2\ 10^{-8}$	.00015	.00001
		$\overline{\nu}_\mu \to \overline{\nu}_\mu$	890	2200	9400	700	4800	7.0×10^4
μ^+	(b)	$\nu_e \to \nu_\tau$	31	60	70	20	67	73
		$\nu_e \to \nu_e$	1400	3700	1.6×10^4	1100	8000	1.2×10^5
		sig/bkg	.022	.016	.004	.018	.008	.0006
			(2.4σ)	(2.7σ)	(1.5σ)	(1.9σ)	(2.0σ)	(0.6σ)
		$\overline{\nu}_\mu \to \overline{\nu}_\mu$	890	2200	9400	700	4800	7.0×10^4
μ^-	(c)	$\nu_\mu \to \nu_\tau$	450	570	650	410	620	680
		$\nu_\mu \to \nu_\mu$	760	3100	1.7×10^4	490	8000	1.4×10^5
		sig/bkg	.6	.18	.04	.8	.08	.005
			(35σ)	(23σ)	(12σ)	(40σ)	(16σ)	(4.6σ)
		$\overline{\nu}_e \to \overline{\nu}_e$	770	1900	8100	600	4100	6.1×10^4

For comparison, the approximate numbers of events in the proposed CERN - Gran Sasso experiment (NGS), and Minos experiments, are given below. It is seen that the numbers of events with the 1.5 MW neutrino factory, in a detector at the same 730 km, is approximately 100 times that in the NGS, or about 40 times the highest energy Minos example. If the proton driver were upgraded to 4 MW, the factors are approximately 300 and 100.

The figutes in the following table need checking.

	νfactory	CERN-NGS	FNAL Minos		
$<E_\nu>$ (GeV)	40	26	3	6	12
km	730	730		730	
$\nu_\mu \to \nu_\tau \to \tau$	680	≈ 7	$\approx$0	$\approx$30	$\approx$40
$\overline{\nu}_\mu \to \overline{\nu}_\mu$	140k	1.5k	.46k	1.4k	3.2k