

Chapter 7

Muon Storage Ring

7.1 Overall Design Motivation

The storage ring is the part of the Neutrino Factory that delivers the neutrino beam to the detector. As such, its effectiveness can be defined by the ratio of the number of neutrinos aimed in the direction of the detector to the total number of muons injected into the storage ring. Since we consider here sending neutrinos to a single detector, we use the storage ring geometry shown in Fig. 7.1. The straight sections are aligned parallel to a line from the storage ring to the detector. Thus, the muons decaying in one of the two straight sections (the “production” straight) contribute to the neutrinos headed toward the detector, while muons decaying in the rest of the ring do not. The effectiveness f_S of the storage ring is the ratio of the length of the production straight section to the circumference of the storage ring:

$$f_S = \frac{L_S}{C} = \frac{L_S}{2(L_S + L_A)}, \quad (7.1)$$

where L_S is the length of one straight section, and L_A is the length of 180° of arc.

In the case of a detector at the WIPP facility, the storage ring must be tilted 13.1° from horizontal to have a straight section pointing toward the detector. It follows that there must be a substantial elevation change from the top of the ring to the bottom. For this study, we have constrained ourselves to keep the bottom of the storage ring tunnel at least 10 feet above the water table. At the Brookhaven site, the water table averages 48 feet above sea level, with a seasonal variation of ± 5 feet. Thus, the bottom of the storage ring tunnel will be at 63 feet above sea level. The highest ground elevation on the Brookhaven site where the storage ring could be realistically placed is about 90 feet above

7.1. Overall Design Motivation

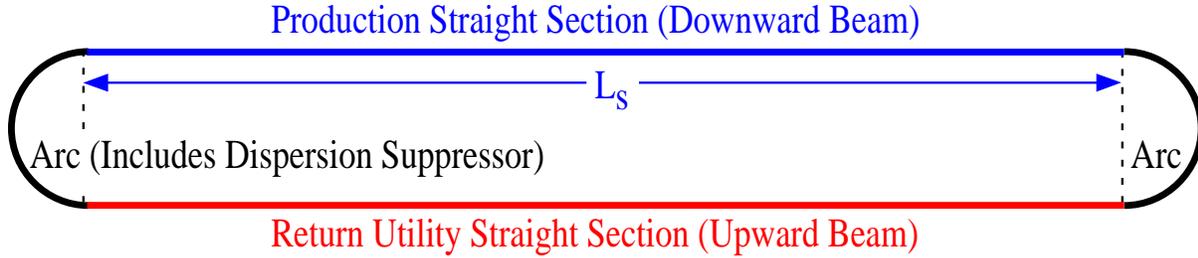


Figure 7.1: Storage ring diagram.

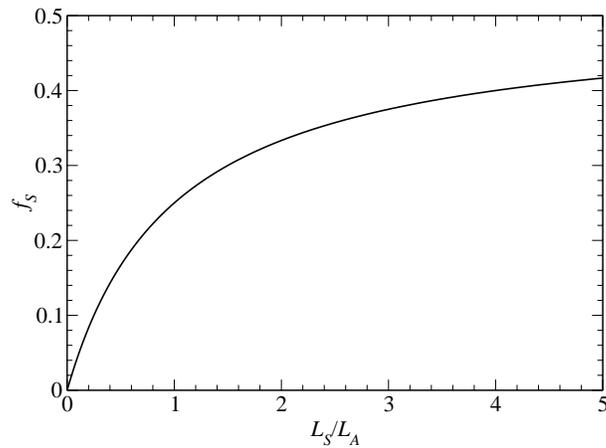


Figure 7.2: Fraction of decays in a straight (f_S) as a function of the ratio of the length of a straight to the length of an arc.

sea level. Thus, there is not enough room to keep the entire storage ring underground, and a hill is needed. For large L_S , the height of the hill increases linearly with L_S , and the volume of fill required for that hill increases as L_S^3 . Thus, there is an economic incentive to keep the ring circumference small.

One can see from Eq. (7.1) that f_S depends only on the ratio L_S/L_A . That dependence is plotted in Fig. 7.2. One can conclude two things from Fig. 7.2. First, for a given f_S , the shorter the arc is, the shorter the straight section can be. Second, beyond $L_S/L_A \approx 2$, it takes a very large change in L_S/L_A to increase f_S by even a small amount.

We summarize in Table 7.1 the values of some parameters of the storage ring.

Table 7.1: Muon storage ring parameters.

Energy (GeV)	20
Circumference (m)	358.18
Normalized transverse acceptance (mm-rad)	30
Energy acceptance (%)	2.2
Arc	
Length (m)	53.09
No. cells per arc	10
Cell length (m)	5.3
Phase advance (deg)	60
Dipole length (m)	1.89
Dipole field (T)	6.93
Skew quadrupole length (m)	0.76
Skew quadrupole gradient (T/m)	35
β_{\max} (m)	8.6
Production Straight	
Length (m)	126
β_{\max} (m)	200

7.1.1 Design Choices for Optimizing Arcs

It is clear from Fig. 7.2 that it is beneficial to minimize the arc length. Whereas Study-I focused on a storage ring with an energy of 50 GeV, here we consider a 20 GeV storage ring, at least partly for this reason. (there are several other reasons, related to machine cost and the insensitivity of CP violation physics to the beam energy, that contributed to this decision as well.)

One way to decrease the arc length of the storage ring is to use high-field bending magnets. Nb₃Sn superconductor can achieve very high fields, but due to its brittle nature, it is difficult to wind a $\cos\theta$ magnet using it. However, winding a pancake-type coil (Fig. 7.3) is not a problem with Nb₃Sn. For this reason, we have chosen to use Nb₃Sn pancake coil magnets in the arcs.

Another consideration for the arc magnet design is the decay of the muons. The superconducting magnet coils must be shielded from the decay electrons, which remain primarily in the horizontal plane of the beam. For a $\cos\theta$ magnet, this would be accomplished by putting a tungsten shield inside the magnet, increasing the required magnet aperture. The pancake coil configuration has the advantage that it can be designed with

7.1. Overall Design Motivation

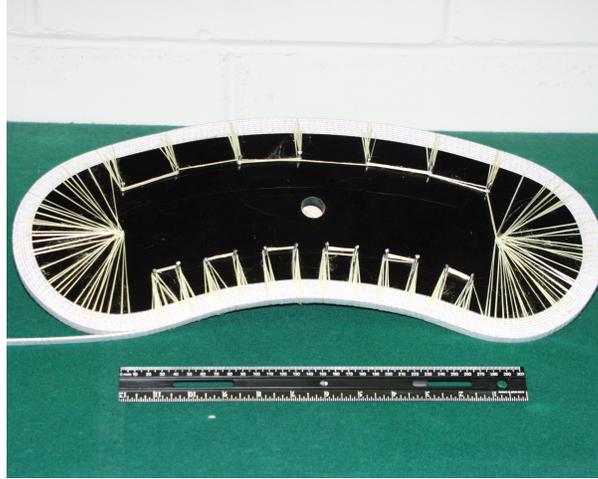


Figure 7.3: Photograph of a pancake coil. This also illustrates a method to provide sagitta and reverse bend using Kevlar strings.

no coil in the midplane, eliminating the necessity for coil shielding.

One of the primary obstacles to reducing the arc length is the required spacing between magnets. As demonstrated in Fig. 7.4, if one simply scales the magnets from the 50 GeV storage ring in Study-I by 40%, the arc is in fact longer than 40% of the 50 GeV arc length, since the inter-magnet spacing must remain roughly the same. Indeed, it might be necessary to increase the inter-magnet spacing due to larger beam sizes, and thus larger magnet aperture, at the lower energy. Some gain can be achieved by eliminating some of the gaps by using combined-function magnets. Ideally, it would be best to eliminate the inter-magnet gaps altogether. The pancake coil design achieves this by using coil configurations where one coil (of two) continues through each transition between magnets.

7.1.2 Choice of Straight Length

Once the length of the arc is minimized, the length of the straight determines the number of neutrinos decaying in the direction of the detector. For this study, there is an added consideration: a hill must be built to accommodate the storage ring. Thus, we chose a goal of $f_S = 0.35$. In the design a straight length of $L_S = 126$ m was chosen along with an arc length of $L_A = 53.09$ m. This requires a hill with a height of 43 m and having 6.4×10^5 m³ of fill.

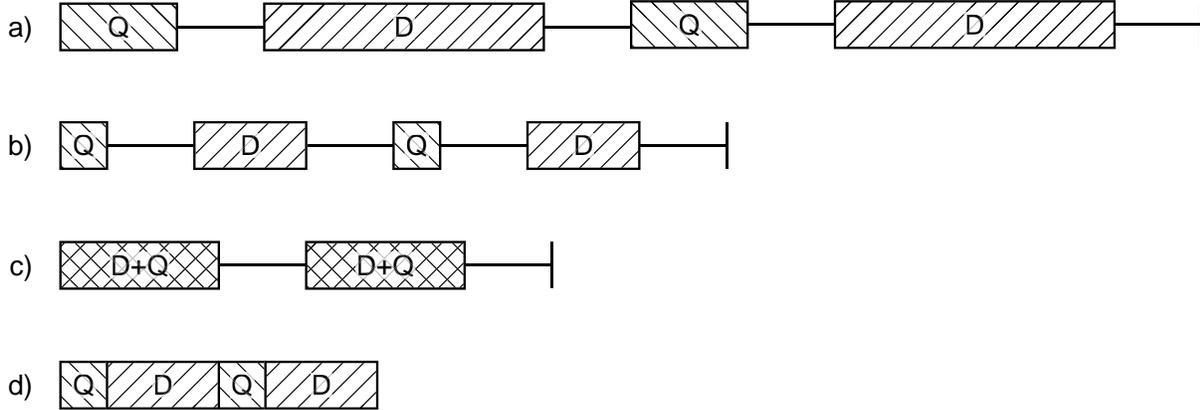


Figure 7.4: Demonstration of the effect of inter-magnet spacing on arc length. The top drawing (a) is the arc cell for a 50 GeV lattice from the Fermilab design study. Scaling that lattice to 20 GeV, but leaving the inter-magnet spacing fixed, does not reduce the cell length to 40% of the original length (b). In (c), we see that using combined-function instead of separated-function magnets can reduce the cell length substantially. In (d), we show what happens if the inter-magnet spacing is completely eliminated.

Table 7.2: Input beam parameters.

Energy (GeV)	20
Normalized transverse acceptance (mm-rad)	30
Ratio of full width to rms width	2.4
Relative momentum spread (full) (%)	± 1.9

7.1.3 Lattice Parameters

The input beam parameters are given in Table 7.2. There two primary issues related to these parameters that need to be dealt with. First, since the magnets operate at very high field, their aperture should be kept small to minimize peak fields and maintain field uniformity. Second, to give a reasonable uncertainty at the detector, the angular spread in the beam must be kept well below the angular spread in the neutrinos due to decays. Specifically, we require that $\sigma_\theta < 0.1/\gamma = 0.53$ mrad.

The first constraint requires that, in the arcs, the beta functions must be kept small. The second constraint requires that, in the production straight, the Twiss parameter gamma ($\gamma = (1 + \alpha^2)/\beta$) be small. Note that, in the production straight, not only does

7.2. The Lattice

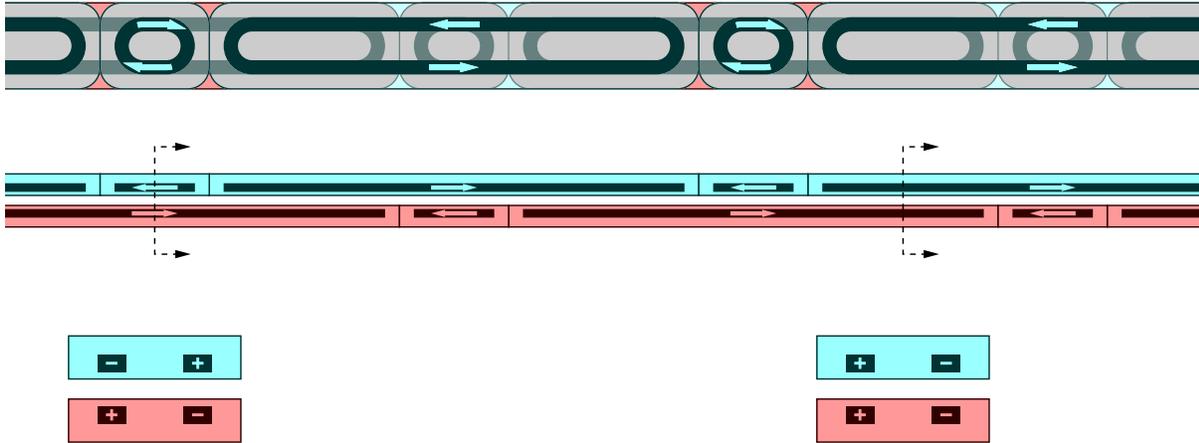


Figure 7.5: Arc magnet layout. Above is a top view of the arc (straightened out); in the middle is a side view. At the bottom are two cross sections of the magnet lattice, at the points indicated in the side view. Coil current directions are indicated, with + being out of the page.

β need to be large, but α must be small. The above requirement on σ_θ means that $\gamma < 0.011 \text{ m}^{-1}$ in the production straight. It will turn out that this constraint is not met for that lattice as is; later in this chapter we describe changes that must be made to achieve this. The reason for the difficulty is that, since β goes from being very small in the arcs to being very large in the production straight, there is necessarily an intermediate region which has moderate β and large α .

7.2 The Lattice

7.2.1 Effect of Magnet Choice on Lattice Design

Figure 7.5 shows a diagram of the arc cell layout. The arcs will consist of two sequences of racetrack coils, one placed above the other, but with the magnets overlapping. The direction of the current alternates from one magnet in the sequence to the next. There are two types of magnetic fields in this lattice, illustrated in the two cross sections in Fig. 7.5: dipole fields (the right cross section), and *skew* quadrupole fields (the left cross section).

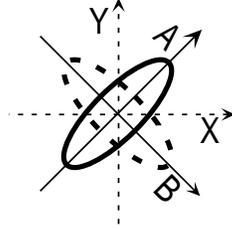


Figure 7.6: Eigenplanes for the pure skew decoupled lattice used for this storage ring. The eigenplanes are rotated 45° with respect to the horizontal.



Figure 7.7: Raising part of each coil to generate an upright quadrupole component in the bend.

Note that this lattice gives rise to skew quadrupole fields rather than upright quadrupole fields. A single coil on top covers one quadrupole and both bends in any given cell. That coil has no ends (and therefore no wasted space) at the ends of the skew quadrupole field in the section that it overlaps. The lower coils at that same skew-quadrupole region do in fact have ends. But, from an efficiency standpoint, those ends are not “wasted”: they merely create a transition from a bending region to a skew quadrupole region. In essence, the transition is really a combined-function section, rather than an empty section.

Thus, the lattice is completely skew, with the eigenplanes shown in Fig. 7.6. However, since the dipoles are focusing in the horizontal plane, which is not one of the eigenplanes, they introduce coupling between the two skew planes. To avoid this we create an upright quadrupole with

$$\frac{dB_y}{dx} = -\frac{eB_y^2}{2p}, \quad (7.2)$$

where p is the reference momentum. Then, the focusing is the same in the horizontal and vertical planes, and therefore cylindrically symmetric. Cylindrically symmetric focusing does not produce coupling between the skew planes. This upright quadrupole is created by raising or lowering one side of each coil pack within the bending region, as shown in Fig. 7.7. The amount that the coils need to be raised or lowered is actually very small, about 1 mm. Generating this amount of vertical shift in the coils will be straightforward.

7.2. The Lattice

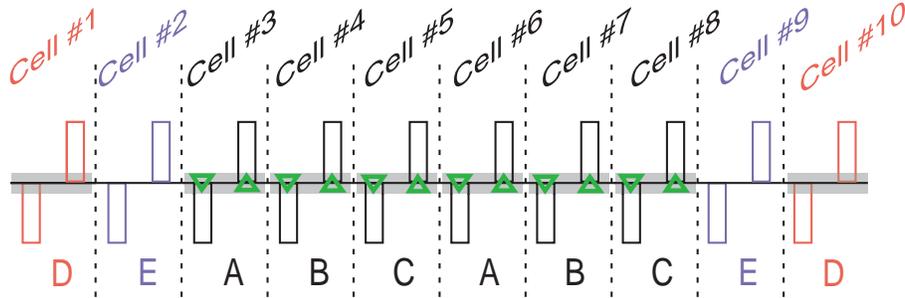


Figure 7.8: Layout of the arc. The “E” cells have no dipoles. The linear magnets in the central six cells are identical, and are the same as those in the “D” cells. The triangles indicate the placement of sextupoles.

7.2.2 Lattice Design

Figure 7.8 shows the layout of an arc. The phase advance per cell is chosen to be 60° . This gives reasonable values for the cell length, the beta functions, and the dispersion. It also gives a reduced swing in the dispersion functions, leading to a lower vertical dispersion.

To accomplish dispersion suppression, we employ the pattern of cells shown in Fig. 7.8. There are six central cells, surrounded by two cells having no bend, and on the ends there are two cells that are identical to the central cells. For a 60° phase advance per cell, this gives dispersion suppression and matches the dispersion in the central cells automatically. Arc cells with no dipoles will contain warm quadrupoles. This permits the use of collimation in those cells, and confines the decay shower to a region with warm magnets.

As discussed above, the production straight has large beta functions, so as to minimize the angular spread of the neutrinos decaying toward the detector. The other straight, where we inject, has beta functions roughly twice those in the arcs. That section will also be used to adjust the overall ring tunes.

Table 7.3 gives the parameters for the magnets used in the storage ring.

Lattice functions for the 20 GeV muon storage ring using compact skew combined function arc cells are shown in Figs. 7.9 and 7.10. Here the beta functions, (β_A, β_B) , are given for the 45° -rotated betatron eigenplanes (A,B) shown in Fig. 7.6, but the eigenplane dispersion functions (η_A, η_B) are projected to dispersion in the normal horizontal-vertical coordinate system according to the relationships, $\eta_x = \frac{\eta_A + \eta_B}{\sqrt{2}}$ and $\eta_y = \frac{\eta_A - \eta_B}{\sqrt{2}}$.

By design, the dispersions in the A and B eigenplanes are nearly equal, so the effective vertical dispersion is much smaller than the horizontal dispersion. With this skew lattice,

Table 7.3: Magnet parameters. Only four of the utility straight quadrupoles require the strength shown; others are as low as 6.7 T/m. The arc dipole and SC quadrupole parameters are really idealized parameters for a lattice made from standard magnets; the actual fields are provided by overlapping coils as described in the text.

	Number	Length (m)	Field (T)/ Gradient (T/m)	R_{pole} (mm)	B_{pole} (T)
Arc dipole	32	1.89	6.93		
Arc SC quadrupole	32	0.76	35		
Arc NC quadrupole	16	0.65	27.2	47	1.28
Production straight quadrupole	2	1.85	5.0	175	0.87
	2	5	3.6	200	0.72
	2	2.25	6.4	125	0.8
	2	0.76	23.2	50	1.16
Utility straight quadrupole	28	1.9	11.6	70	0.81

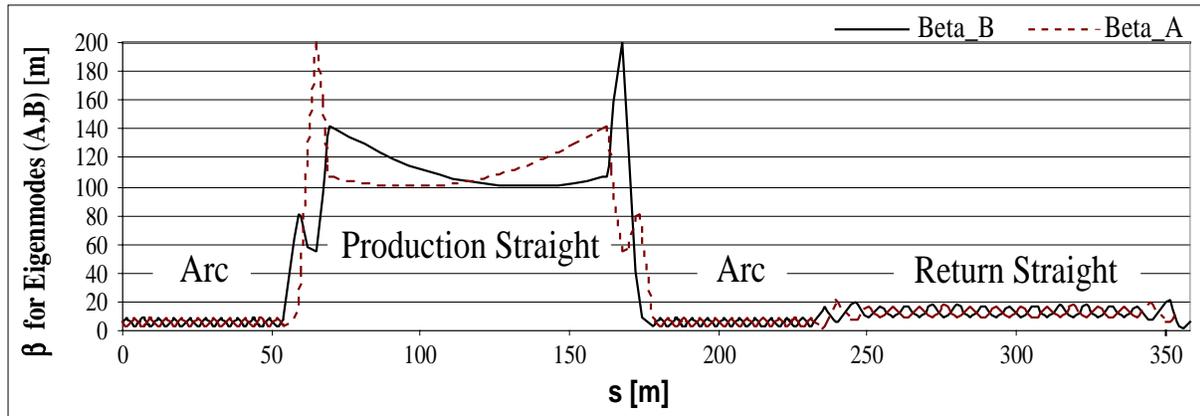


Figure 7.9: Beta functions for the storage ring, in the 45°-rotated eigenplanes.

the horizontal dispersion is nearly constant across the arc, whereas the vertical dispersion oscillates with small amplitude about zero. The arc cells without bending match the dispersion to zero for both eigenplanes in the straight sections.

The lattice shown in Fig. 7.9 has a ratio between the lengths of arcs and straight sections such that the geometric decay ratio, f_s , is just over 35%. The central 93 m of the production straight has a Twiss gamma function of 0.01 m^{-1} , which meets the requirements for the angular divergence described in Sec. 7.1.3. Within that central

7.2. The Lattice

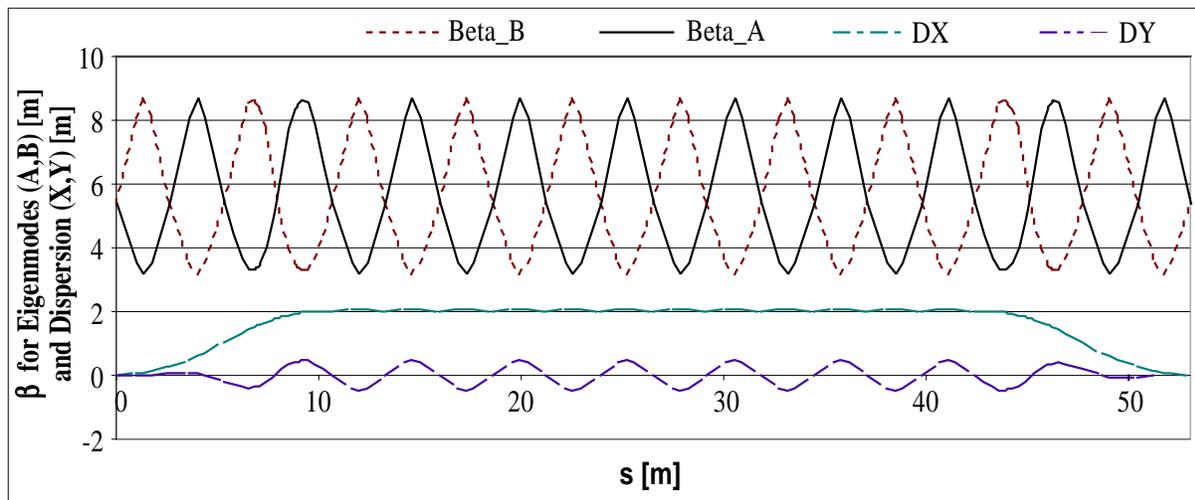


Figure 7.10: Beta functions and the dispersion function for an arc of the storage ring. The beta functions are in the 45° -rotated eigenplanes, while the dispersion functions are projected into the standard horizontal and vertical planes.

section, there are no quadrupoles; this appears to be the most straightforward way to minimize α , and therefore γ , in that section. Beta functions in the return straight are intermediate in magnitude between the values in the arcs and the production straight in order to facilitate injection.

7.2.3 Chromatic Correction Sextupoles

For chromatic correction, a skew sextupole has been placed at each skew quadrupole in the central section of the arcs (see Fig. 7.8). A skew sextupole, like the dipole and skew quadrupole magnets, requires no coil in the midplane of the magnet. These sextupoles must be very strong, since they correct not only the chromaticity generated in the arcs, but also the chromaticity generated in the straights. Since the straights are significantly longer than the arcs, the sextupoles require high strengths.

The sextupoles are divided into A, B, and C families, as indicated in Fig. 7.8. All sextupoles in a given family have identical strengths (this is true separately for the sextupoles at focusing and defocusing quadrupoles). Due to the 60° phase advance per cell, there is an automatic first-order cancellation between the nonlinear terms from these sextupoles, as shown in Fig. 7.11. The families can be chosen to give a second-order cancellation of

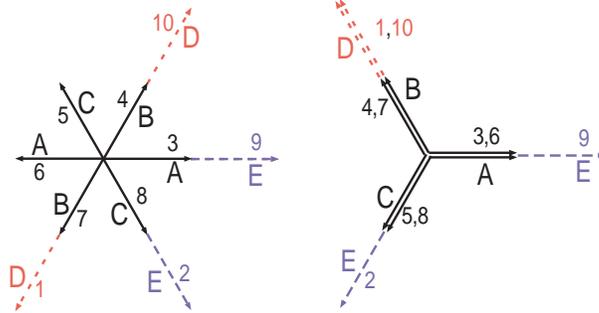


Figure 7.11: Diagram showing cancellation of the sextupole nonlinearities due to the choice of families and the 60° phase advance per cell. The left diagram shows the first-order cancellation, the right diagram shows that, given the right relationship between the A, B, and C family strengths, there will be a second-order cancellation as well. Note that putting sextupoles in the D or E regions would not give a similar cancellation.

the nonlinearities as well, as illustrated in the second part of Fig. 7.11.

The required sextupole strength to achieve zero chromaticity is $S = B'' = 78 \text{ T/m}^2$. This sextupole strength could be reduced by putting sextupoles in the remaining arc cells, but there is no simple cancellation scheme for the nonlinearities in that case, as illustrated in Fig. 7.11.

With one degree of freedom in our three families needed to correct the chromaticity, and another needed to correct the second-order nonlinearities, there is still one remaining degree of freedom in each plane. We use this to minimize the nonlinear momentum dispersion in the production straight. Writing the transverse momentum of the closed orbit in the center of the production straight as a function of $\delta = \delta p/p$ as $p_{x0}(\delta) = D_{p1}\delta + D_{p2}\delta^2 + \dots$, the sextupoles could be used to eliminate D_{p2} (note that D_{p1} is already eliminated in our dispersion suppression scheme).

7.2.4 Coil End Effects

There are two end effects with which we will be primarily concerned. The first is a solenoid generated by the ends of the coils. The reason for the solenoid field is seen in Fig. 7.12. Where one coil ends, there is a transverse current, which leads to a longitudinal field. (In standard magnet designs, the symmetry of the return coils eliminates this solenoid field. Since our design lacks this symmetry, there is a net solenoid field.)

7.2. The Lattice

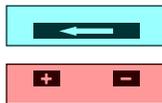


Figure 7.12: Arc magnet cross section at the skew quadrupole end of the upper coil pack.

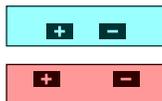


Figure 7.13: Cross-section as one approaches the end of the top coil, showing the coils on top moving closer to each other. This has a nonzero sextupole field.

It is important to note that this solenoidal field will affect the linear lattice; it is not just a nonlinear effect. There are two effects: the first is a rotation, the other is a focusing effect. The focusing can just be taken into account in the linear lattice, but the rotation is more difficult since it can couple the eigenplanes together. The coupling can be removed from the lattice, if necessary, by rotating the skew quadrupoles by the amount that each end rotates. One quadrupole sign is rotated in one direction; the other is rotated in the opposite direction. The amount of rotation can be calculated precisely knowing only the current in the coil and the height of the coil above the axis. The coils need to be displaced by only about 1 mm to put this rotation in, which is comparable to the coil displacement we already must achieve to put the upright quadrupole field into the bending sections.

The second end effect we consider is a sextupole contribution. This is not the sextupole-order contribution caused by the longitudinal derivative of the dipole field. Rather, it comes from the fact that the coil symmetry changes when the coil starts to cross at the end. Consider the dipole cross section in Fig. 7.5. That cross section will have a nonzero sextupole field unless the coils are placed in precisely the correct places. (For a single wire coil, that would be along a line through the center, 30° from horizontal.) The main body of the magnet is designed with the coils placed to eliminate that sextupole contribution. However, when the coils are moved out of the position that zeros the sextupole contribution so that they can cross over, as depicted in Fig. 7.13, there will be a net sextupole field at that point. This sextupole field can be hundreds of times as large as the sextupole component in the body of a realistic magnet.

Uncorrected, this sextupole field would decrease the dynamic aperture of the ring.

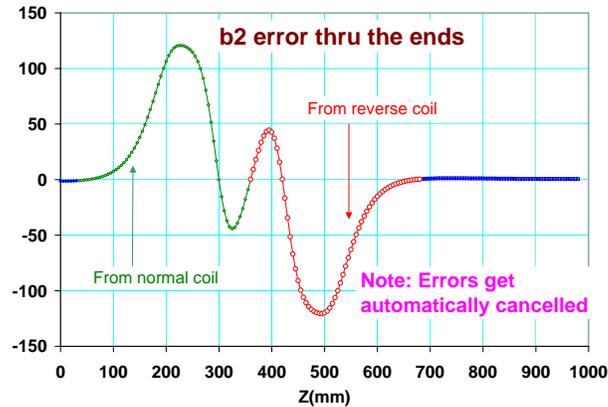


Figure 7.14: Longitudinal variation in the sextupole field at the transition from the dipole to the skew quadrupole region.

However, immediately adjacent to the end that generates this sextupole is another coil end that generates an opposite sextupole. As shown in Fig 7.14, the integrated sextupole is zero over a very short distance (a few tens of cm), and should not be a problem for beam dynamics.

7.2.5 Correction and Tuning

In principle, since the short coils in the skew quadrupole regions can be powered separately from the longer coils, it is possible to perform both dipole and quadrupole corrections by varying their current. To create a “pure” correction requires more than one magnet working in concert, since one coil does not produce a pure dipole or a pure skew quadrupole at a given point.

Alternatively, an additional coil could be added opposite the short coil to allow pure dipole or pure skew quadrupole corrections at a given point. The coil could be either superconducting or normal conducting, since it need not generate large fields.

To adjust the overall tune of the ring, the quadrupoles in the utility straight will be used.

7.2.6 Tracking

Tracking studies were performed on a single arc cell using COSY INFINITY [1]. A more detailed analysis is given in [2]. The arc cell used is slightly different from the one described earlier in this chapter. A diagram of the cell used in the tracking is shown in

7.2. The Lattice

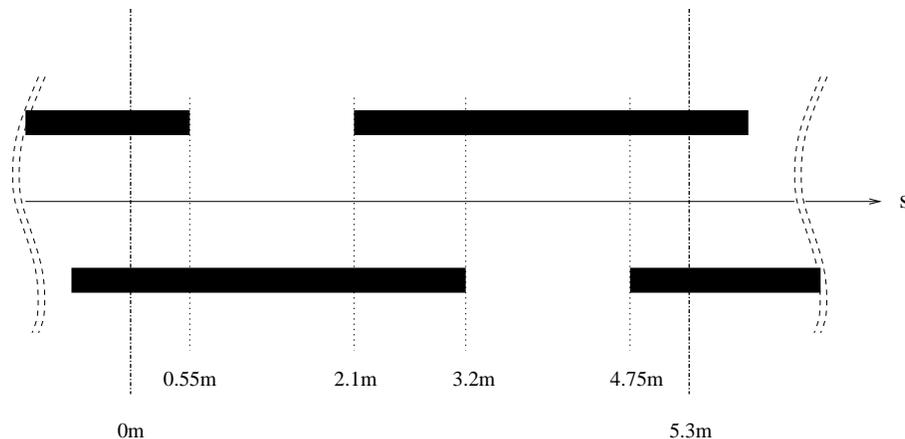


Figure 7.15: Arc coil configuration used for tracking.

Fig. 7.15. Note that the short coils with a reversed current direction are missing. Instead of having a pure skew quadrupole section, there is a combined-function section which is part dipole and part skew sextupole. The phase advance in this cell is 60° , just as in the actual lattice. We expect the results from tracking this cell to be similar to what we find from tracking the actual lattice, since the only real difference between the two is a slight redistribution of the linear components. Table 7.4 lists the magnet parameters used in the COSY-INFINITY tracking. In the combined-function areas, there are regions of overlapping bend, upright quadrupole, skew quadrupole, and skew sextupole (these are the regions with only a single coil, labeled as (b)/(b') and (e)/(e') in Table 7.4). To model this properly in COSY INFINITY, these sections are divided into 10 subsections of length $\Delta L = 0.155$ m, each of which consists of a bend/upright-quadrupole section of length ΔL , a negative drift of length $-\Delta L$, and a skew-quadrupole/skew-sextupole section of length ΔL . The skew sextupole strength is chosen to be sufficient to precisely cancel the chromaticity. COSY INFINITY also uses the full kinematic Hamiltonian (the full square root, as opposed to simply $p_x^2/2 + p_y^2/2$) in its tracking.

Figure 7.16 shows tracking results for the arc cell as described above. The dynamic aperture for this idealized lattice is much larger than needed. Figure 7.17 demonstrates, that for amplitudes within the magnet aperture, the lattice is extremely linear and well decoupled.

The next step is to add end fields to these magnets. An Enge-function falloff model [3] was used to model these end fields. Fig. 7.18 shows the results. Note that there is now significant coupling between the planes. Linear coupling is, in principle, correctable as described earlier, but there is a nonlinear coupling that will inevitably be there, as is

Table 7.4: Optical description and parameters for the arc cell. $k_1 = (\partial B_y / \partial x) / (B\rho)$, $k_2 = (\partial^2 B_y / \partial x^2) / (B\rho)$, $(B\rho)$ for a 20 GeV muon is 67.064332 Tm, B_D is the dipole field strength, and B_Q and B_S are the quadrupole and sextupole field strength at the aperture, $r = 6.5$ cm, divided by 7.02296 T.

Section	Starting position (m)	Length (m)	Tilt angle ($^\circ$)	Deflection (B_D)	k_1 (m^{-2}) (B_Q)	k_2 (m^{-3}) (B_S)
(a)	0.00	0.55	45	57.6 mrad (7.02296 T)	-0.00548 (-0.02389 T)	0
(b)	0.55	1.55	45	81.8 mrad (3.51148 T)	-0.00137 (-0.00597 T)	0
(b')			0	0 mrad	-0.30269 (-1.31950 T)	-0.01932 (-0.002737 T)
(c)	2.10	0.55	45	57.6 mrad (7.02296 T)	-0.00548 (-0.02389 T)	0
(d)	2.65	0.55	45	57.6 mrad (7.02296 T)	-0.00548 (-0.02389 T)	0
(e)	3.20	1.55	45	81.8 mrad (3.51148 T)	-0.00137 (-0.00597 T)	0
(e')			0	0 mrad	0.30269 (1.31950 T)	0.01317 (0.001866 T)
(f)	4.75	0.55	45	57.6 mrad (7.02296 T)	-0.00548 (-0.02389 T)	0

the case in an upright lattice. The beam still is well within the dynamic aperture.

Next, multipole components in the body of the magnets are added. We started with multipole components from a 2D model of the region with both coils, which was computed using Poisson [4, 5] and Opera 2D. Independently, we performed a 3D field simulation using a bar magnet model, where the field can be expressed analytically [6, 7]. Two models were used: one where the bar magnets were infinite in extent, and a second where there were four magnets laid out as in Fig. 7.15. The bar magnet models were constructed to match as closely as possible the 2D model described above. The multipole fields in the real system were computed by starting with the values from the original 2D model, and scaling them by the ratio of the values in the second bar magnet model to those in the

7.2. The Lattice

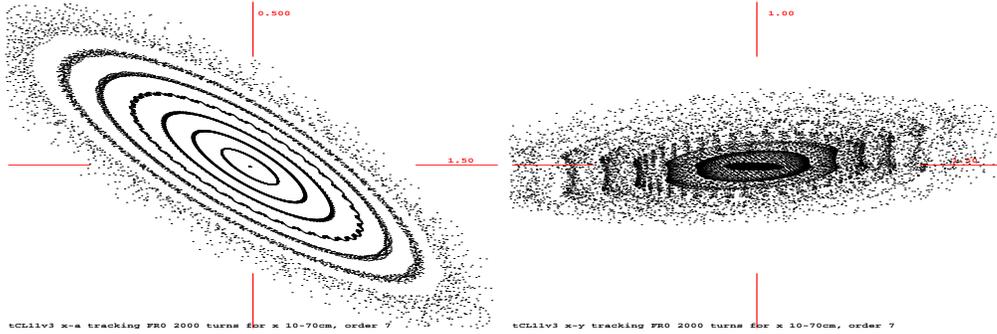


Figure 7.16: Tracking particles in the arc cell starting at $A = 10, \dots, 70$ cm, showing A - p_A motion (left) and A - B motion (right). The scales of the pictures cover ± 150 cm in the A plane, ± 100 cm in the B plane, and ± 0.5 in the p_A plane. Recall that A and B are the diagonal eigenplanes.

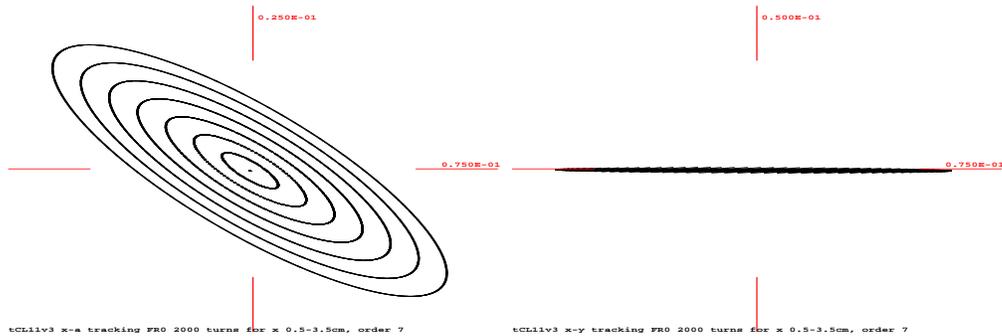


Figure 7.17: Tracking particles in the arc cell, starting at $x = 0.5, \dots, 3.5$ cm, showing A - p_A motion (left) and A - B motion (right). The scales of the pictures cover ± 7.5 cm in the A plane, ± 5 cm in the B plane, and ± 0.025 in the p_A plane.

first bar magnet model. The values of these computed multipole components are given in Table 7.5.

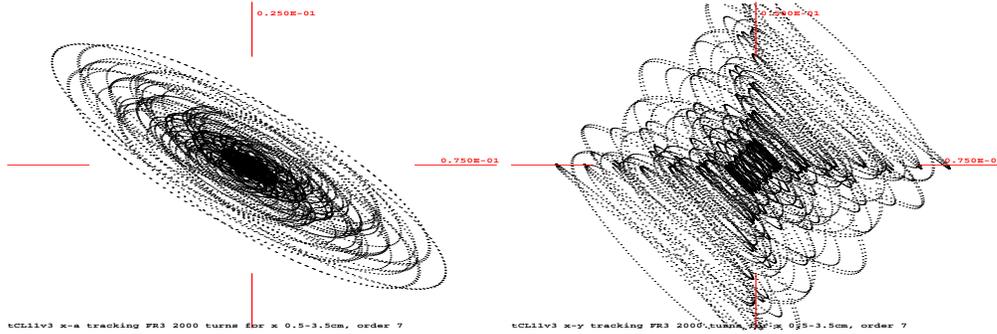


Figure 7.18: Tracking with end fields. Initial conditions and scales as in Fig. 7.17.

Table 7.5: Multipole components in the arc cell. The value given in the table is the maximum of the magnetic field in Tesla for that multipole component at a radius of 6.5 cm.

		Double Coil Region Normal	Single Coil Region (b)	
			Normal	Skew
2	Sextupole	-0.721874127471	-0.360937063736	0
3	Octupole	0	0	0.100208577080
4	Decapole	-0.325677875510	-0.162838937755	0
5	Dodecapole	0	0	0.105845309541
6	14-pole	-0.048154527567	-0.024077263783	0
7	16-pole	0	0	-0.111799108203

Results of tracking with these multipole components are shown in Fig. 7.19. The dynamic aperture is unacceptably small. However, the multipole components shown in Table 7.5 are extremely large. As described later in Sec. 7.3.1, we believe we can make those multipole components much smaller. It was found that if the sextupole, skew octupole, and decapole components were reduced to 10% of the values in Table 7.5, the dynamic aperture was acceptable, even with fringe fields (see, Fig. 7.19).

The bar magnet model gives a solenoidal field as shown in Fig. 7.21. In that model, the solenoidal field doesn't become small anywhere. The solenoidal field was modeled using several short solenoids with a Gaussian-shaped profile. There are three effects of this

7.2. The Lattice

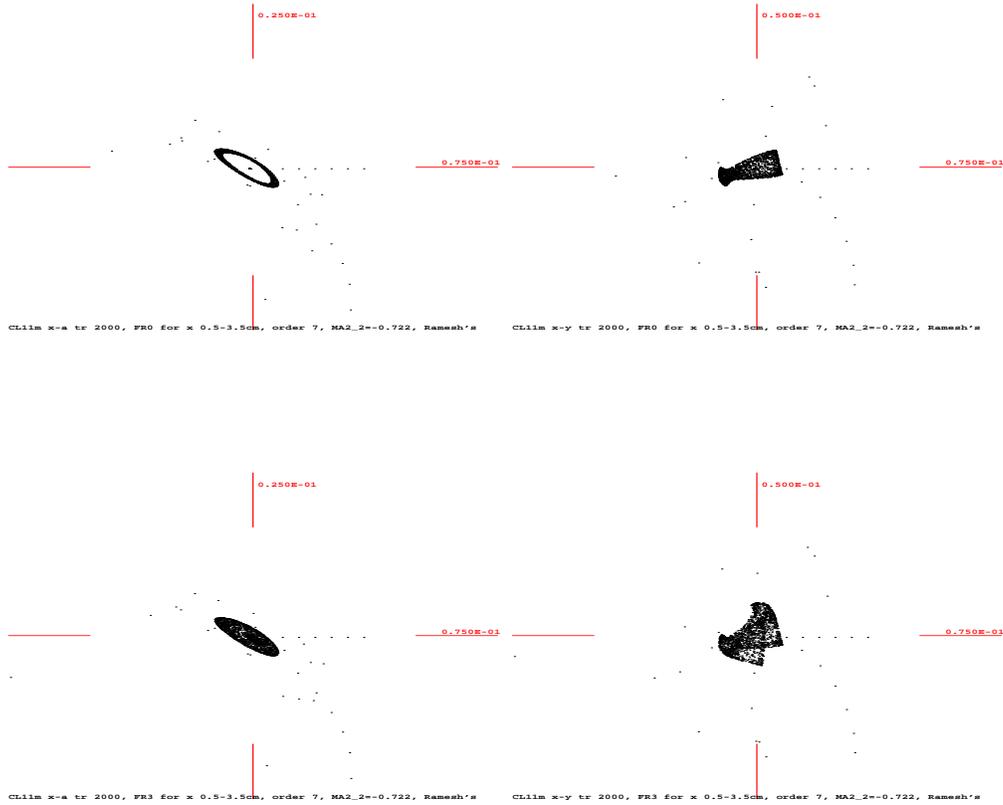


Figure 7.19: Tracking the arc cell with the additional multipole components in Table 7.5. Without (top two) and with (bottom two) fringe fields.

solenoidal field: linear coupling, linear focusing, and additional nonlinearities (from the longitudinal derivative of the field). Figure 7.22 shows the results of tracking with these solenoidal fields added. The solenoidal fields have a significant impact on the dynamic aperture of the machine. Further study is needed to examine to what extent this effect can be mitigated.

Table 7.6 indicates the strong effect of the solenoidal fields on the tune. In particular, the correction scheme described in Sec. 7.2.3 for canceling geometric nonlinearities will not work as well when the phase advance is not exactly 60° . The tune split comes from coupling from the solenoids, which can be corrected as described in Sec. 7.2.4. Thus, this

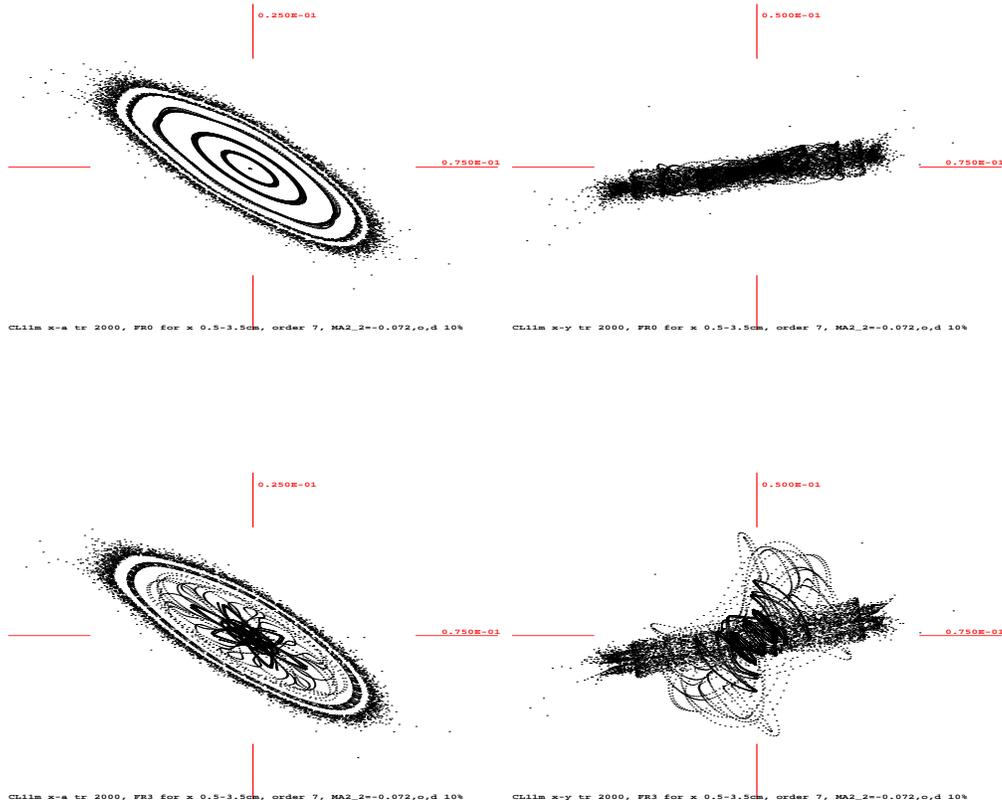


Figure 7.20: Tracking with the additional multipole components, with the normal sextupole, skew octupole and normal decapole strengths reduced to 10% of their initial values. Without/with (top/bottom) fringe fields.

decoupling, along with an adjustment of magnet strengths to compensate for the focusing effect from the solenoids, may restore the dynamic aperture.

7.3 Magnets

Figures 7.23 and 7.24 show diagrams of the arc magnets to be used in the storage ring. Each coil pack with its associated hardware will be built in the lab, and the individual cryostats will then be connected together appropriately in the tunnel.

7.3. Magnets

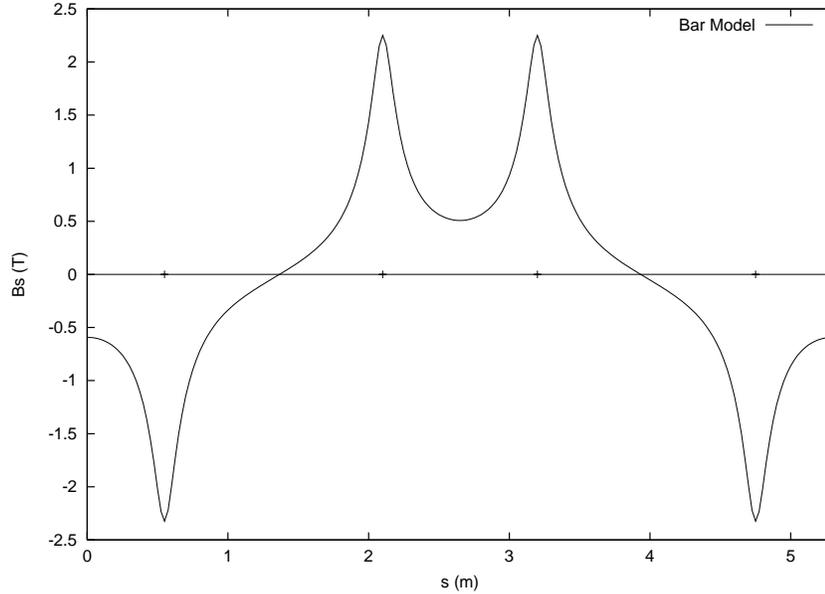


Figure 7.21: Field profile of solenoidal field component from the bar magnet model.

Table 7.6: Linear tunes in each study case for the two orthogonal planes.

	No fringe field effects	With fringe field effects
Initial approximation	0.166667 0.166667	N/A
Thick lens model	0.168422 0.168422	0.168040 0.166919
With solenoids	0.162584 0.174157	0.162190 0.172703

The vertical aperture of the dipole is 80 mm. The actual vertical separation between the upper and lower coils is 130 mm, to accommodate space taken by the support structure, heat shield, and cryostat. The aperture in the horizontal plane is much larger than in the vertical plane, since the space under the coils is also available due to the open-midplane design (see Figs. 7.23 and 7.24). The actual horizontal coil aperture, 240 mm, is dictated by the minimum bend radius at the end of 120 mm. Of course the beam cannot take advantage of the region under the coil, since the field quality is poor there. It is important to know the field in this region, however, since it is needed for tracking of

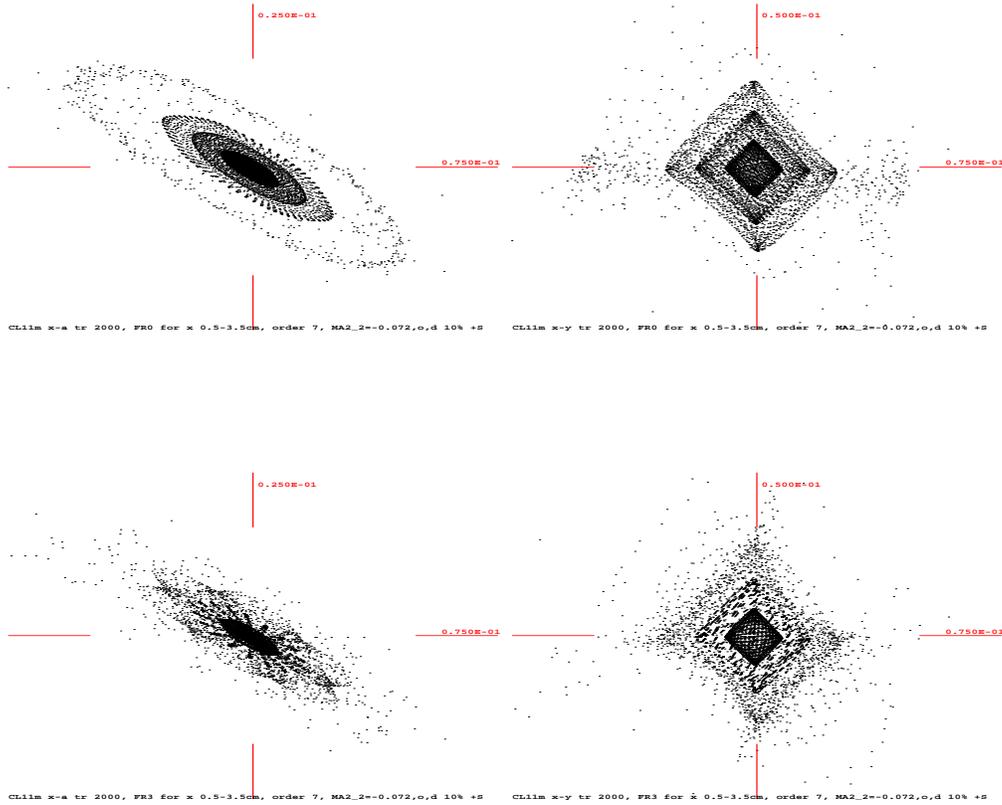


Figure 7.22: Adding the solenoid components to tracking with the additional multipole components, with the sextupole, octupole and decapole strengths 10% of the initial. Without/with (top/bottom) fringe fields.

decay electrons. The region of good field quality can be determined from the harmonics given in Section 7.3.1, or by a field profile within the beam tube.

The operating field of the dipole is 6.93 T, and the quench field is over 8 T. This gives an operating field margin of over 15%. As mentioned earlier, the maximum field on the conductor at quench excludes the possibility of using NbTi at 4.2 K operating temperature. The coils, therefore, are made of Nb₃Sn superconductor. Large bend radius in the ends allows the use of the “react and wind” technique in a pancake (racetrack) coil geometry. These pancake coils must have a large sagitta, due to the small size of the

7.3. Magnets

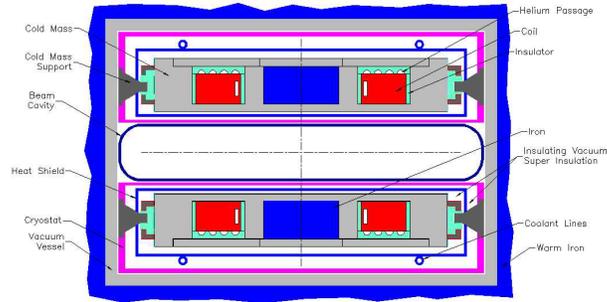


Figure 7.23: An engineering design of the magnet cross section.

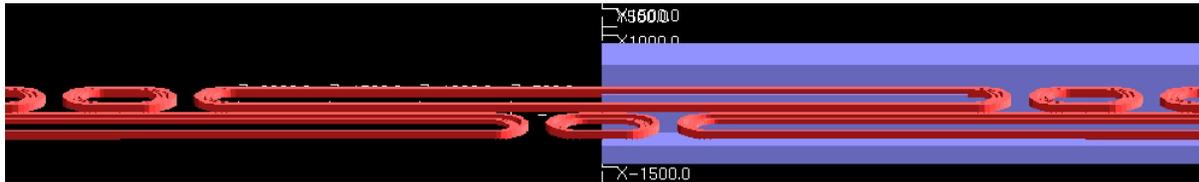


Figure 7.24: Three-dimensional view of the storage ring magnets. Iron (blue) is shown for only half the length.

ring. The reverse curvature in the coil is provided by Kevlar strings as shown in Fig. 7.3.

In this design, most of the energy from the decay particles is deposited in the warm-iron yoke. A warm-iron design allows the heat generated by decay particles to be removed efficiently. This means that a tungsten inner liner to protect the superconducting coils is not needed.

An engineering design of the magnet is shown in Fig. 7.23. The coils are located in a cryostat that is placed inside the warm iron yoke. Upper and lower coils will have separate cryostats that are placed in a common vacuum vessel. A major design consideration is to minimize the cryostat volume, in particular the vertical separation between the upper and lower coils.

The Lorentz forces are large, due to the high field and large aperture of this magnet. Horizontal forces are contained in a self-supporting collar structure. The vertical forces are transmitted to the cryostat with a coldmass support post, which in turn is connected to the iron and contains the force. The posts are designed to minimize the heat leak, which in the present design is comparable to the heat load from the decay products.

Table 7.7: Estimated field errors in the pure dipole magnet section at a 20 mm reference radius. $\langle b_n \rangle$ and $\langle a_n \rangle$ are the expected means of the normal and skew terms. $d(b_n)$ and $d(a_n)$ are systematic uncertainties arising from design and manufacturing errors, and $\sigma(b_n)$ and $\sigma(a_n)$ are the random uncertainties in those values. Note that $n = 2$ corresponds to the sextupole term.

n	$\langle b_n \rangle$	$d(b_n)$	$\sigma(b_n)$	$\langle a_n \rangle$	$d(a_n)$	$\sigma(a_n)$
1	0	0.2	0.2	0	1	2
2	-1	1	2	0	0.1	0.5
3	0	0.1	0.1	0	0.3	1
4	-1	1	1	0	0.05	0.2
5	0	0.03	0.03	0	0.1	0.5
6	-0.3	0.2	0.1	0	0.03	0.1
7	0	0.03	0.01	0	0.03	0.1
8	-0.1	0.1	0.02	0	0.03	0.1
9	0	0.03	0.01	0	0.03	0.1
10	-0.03	0.02	0.02	0	0.03	0.1

7.3.1 Magnetic Design

The magnet cross section is based on two double-layers of pancake coils (total four layers) with a curvature that approximately follows the beam trajectory. The total height of the coil pack is 40 mm (made with four stacks of 10-mm-wide cable), and the width is 60 mm. The cross section has spacers (equivalent to the wedges in a cosine theta design) to optimize the field quality and to reduce the peak field in the conductor. The peak field in the conductor is higher when the coil and magnet are made compact, in particular, when there is a large gap between the upper and lower coils at the midplane of the magnet. In the present design, the peak field is about 50% of the central field.

The expected skew and normal harmonics are given in Tables 7.7 and 7.8. These are not final values, as the magnetic design is not yet optimized. Table 7.7 gives the field errors in the dipole section, where the upper and lower coils have the same polarity, and Table 7.8 gives the field errors in the skew quadrupole section, where the upper and lower coils have opposite polarity. Field harmonics are expressed in terms of the normal and

7.4. Beam Flux to Detector

Table 7.8: Field errors in the skew quadrupole magnet section at a 20 mm reference radius. See Table 7.7 for definitions.

n	$\langle b_n \rangle$	$d(b_n)$	$\sigma(b_n)$	$\langle a_n \rangle$	$d(a_n)$	$\sigma(a_n)$
1	0	0.2	0.2	0	1	2
2	-0.5	0.5	1	0	1	0.5
3	0	0.1	0.1	2	2	1
4	-0.5	0.5	0.5	0	0.05	0.2
5	0	0.03	0.03	1	1	2
6	0	0.2	0.1	0	0.03	0.1
7	0	0.03	0.01	0.5	0.5	0.3
8	0	0.1	0.05	0	0.03	0.1
9	0	0.03	0.01	0.1	0.03	0.1
10	0	0.02	0.01	0	0.03	0.1

skew harmonic coefficients, b_n and a_n , defined by

$$B_y + iB_x = 10^{-4} B_0 \sum_{n=0}^{\infty} (b_n + ia_n) \left(\frac{x + iy}{R} \right)^n, \quad (7.3)$$

where x and y are the horizontal and vertical coordinates, R is the reference radius, and B_0 is the field strength at $x = R$ and $y = 0$ in the dipole straight section. Field harmonics are normalized to this same B_0 in both sections. The reference radius R is 20 mm.

7.4 Beam Flux to Detector

There are really two quantities that we must deliver to the detector: a sufficient flux of neutrinos, and a sufficiently small uncertainty in that flux. It turns out that the latter may be the more difficult challenge. To understand this, consider Fig. 7.25. In that figure, we show event rate as a function of beam divergence (for a round beam), and fit a curve to that. This is based on the fact that, for an elliptical Gaussian beam divergence, with angular divergences of σ_p and σ_q in the two directions, and a Gaussian divergence of the neutrino decays with RMS divergence σ , the central flux is expected to be proportional to

$$\frac{1}{\sqrt{(1 + \sigma_p^2/\sigma^2)(1 + \sigma_q^2/\sigma^2)}}. \quad (7.4)$$

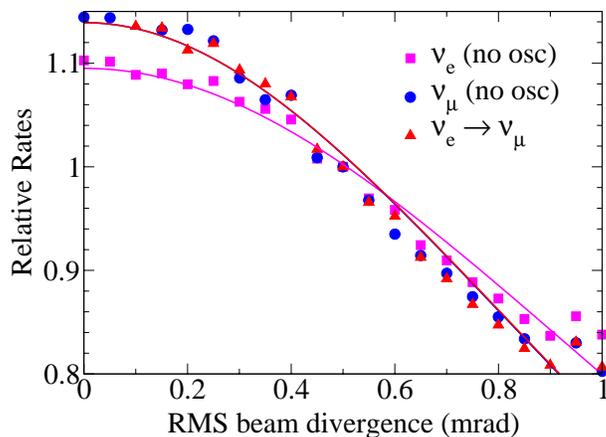


Figure 7.25: Relative event rate at the detector for a 30 GeV muon beam as a function of the rms beam angular divergence, with data taken from [8]. Curves are fit to this data with the functional form $A/(1+\sigma_\theta^2/\sigma^2)$, where σ_θ is the beam divergence, and σ and A are fit parameters. This form is the central flux of two overlapping Gaussians, one with divergence σ_0 (the beam) and the other with divergence σ (the decay divergence). The fit σ_0 values are 1.64 mrad (ν_e) and 1.41 mrad (ν_μ and $\nu_e \rightarrow \nu_\mu$).

While the basis for this model is not particularly accurate, it nonetheless gives a reasonably accurate representation of the simulated data when you fit to σ . Based on the fit to the data in Fig. 7.25, we take $\sigma = 0.42/\gamma$.

From this curve, one can understand why a larger divergence beam contributes more to the uncertainty: the uncertainty comes from the uncertainty in the angular spread times the slope of that curve. When one has a larger divergence, the curve has a larger slope. Another way to view this is that if the angular spread in the beam is much smaller than the angular divergence from the decays, the uncertainty in that angular spread has little effect on the total angular spread in the beam. Using this model, we can plot the flux and the uncertainty in that flux relative to a beam with zero angular spread. This is shown in Fig. 7.26. The resulting relative flux is 0.759 ± 0.014 . Thus, despite the fact that the flux uncertainty from the middle of the straight section is almost exactly 1% (add the two planes in quadrature), the ends of the production straight contribute disproportionately to the uncertainty, despite their relatively small contributions to the total flux. Also, note the reduced flux: instead of having 35% of the muons decaying in a very low divergence beam toward the detector, we have only 27% of the muons decaying

7.4. Beam Flux to Detector

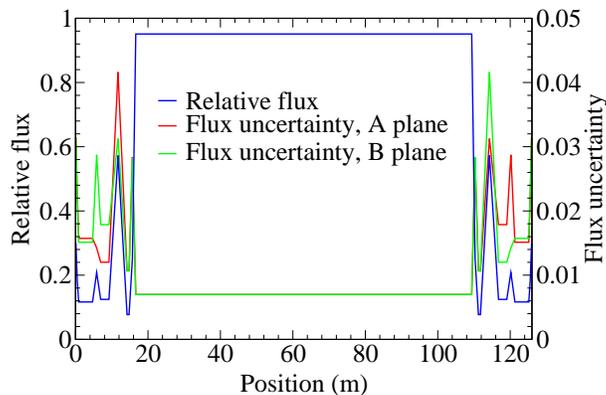


Figure 7.26: Flux at the detector, and the uncertainty in that flux, as a function of position in the production straight. A relative uncertainty in the angular divergence of 15% in each plane is assumed and, within a given plane, those uncertainties are assumed to be completely correlated. The correlation assumption is correct for the contribution to the uncertainty from the emittance uncertainty, but only partially correct for the contribution from the uncertainty in the lattice functions (the uncertainties must be propagated properly).

from the low divergence portion of the straight section.

It is clearly important to eliminate this uncertainty from the ends of the production straight. These ends are a matching section from the relatively low beta functions in the arcs to the high beta function in the straight. Thus, they will necessarily have a large angular divergence. One workaround is to point the matching sections in a different direction from that of the detector by bending slightly less than 180° in the arcs, and inserting a bend just before the “good” section of the production straight. One drawback to this is that the production straight becomes shortened. However, since the ends contribute relatively little to the total flux, not much is lost: whereas originally the relative flux was 0.759 ± 0.014 , the relative flux would become 0.700 ± 0.007 after the workaround.

The required bending angle has been estimated to be 29 mrad. This calculation assumed that the ends produced a flat beam with an angular spread in the wide dimension equal to the largest angular spread anywhere in the end, and required that this beam contribute less than 0.1% to the total flux.

Table 7.9: Required precision for muon beam measurements [8].

momentum	$\delta p_\mu \leq 0.3\%$
momentum spread	$\Delta\sigma_p/\sigma_p < 0.17$
divergence	$\sigma_\theta \leq 0.1/\gamma$
	$\delta\sigma_\theta/\sigma_\theta \leq 0.2$
polarization	$\delta P_\mu < 0.01$
direction	$\delta\theta_\mu < 0.6 \sigma_\theta$

7.5 Instrumentation

The storage ring presents some new beam instrumentation problems. There are of course the usual measurements of emittance, divergence, closed orbit, injection, extraction, beam loss and beam energy. However, since we are interested in the properties of the neutrino beam produced by the muon beam, it is helpful to have additional information. The precision with which the parameters of the muon beam must be known to achieve sufficient precision in the neutrino beam flux have been determined by Geer [8], and these are listed in Table 7.9. As indicated in this table, for instance, the polarization of the muon beam should be measured, and precision measurements of the beam direction in the straight section would be useful. The majority of instrumentation for the muon beam in the storage ring should utilize proven technology. The primary difficulty is that precision measurements are complicated by the presence of decay electrons in the beam, since electrons can cause showers.

The fraction of primary decay electrons in the beam is $L/\gamma\tau c$, where L is a path length in the storage ring, and $\gamma\tau c \sim 125$ km is the decay length at 20 GeV. In bending magnets, these electrons would be swept out within a few meters of their creation, so the electron contribution would be negligible. Quadrupoles produce less efficient sweeping. The fraction of muons that will decay in the 126 m straight sections, however, is ~ 0.001 , and the electron/muon ratio at the downstream end of the straight will be $\sim 0.001F_s$, where F_s is a factor that depends on the probability of electrons showering and being swept from the vacuum pipe. These backgrounds are relatively low, but they may not be insignificant. Since estimates of the electron background are difficult, it may be desirable to have precision measurements external to the ring for determining the neutrino beam direction, profile and divergence.

We anticipate that the 6D “pencil” beams used to tune up the accelerator will also be useful in tuning up and operating the storage ring.

Semertzidis and Morse [9] have looked at using the $g - 2$ frequency of the muons to

7.5. Instrumentation

determine the muon beam energy. They consider measurement using the very substantial signal from synchrotron radiation from decay electrons. Synchrotron radiation is non-linear in energy, which amplifies the oscillations. The beam momentum spread is also measurable using $g - 2$, since the beam will be dephased by the spread in γ .

The beam size and divergence in the quadrupole-free decay straight section of length, $2L$, can be measured by comparing beam profile measurements at the ends, σ_e , and middle, σ_m , of the straight section. We assume that $\alpha = 0$ in the center of the straight section. The beam divergence is $\sigma_\theta = \sqrt{\sigma_e^2 - \sigma_m^2}/L$, and the beam emittance is $\epsilon_\perp = \sigma_m \sigma_\theta$. Since there will be no focusing in this part of the straight section, it should be possible to perform a Monte Carlo calculation and subtract the contributions from any decay electrons. The beam size can be measured using visible transition radiation from foils inserted in the beam, or a variety of other fast detectors.

In order to separate the contributions to the neutrino flux from decays in the upstream and downstream matching sections, bending magnets have been introduced. Using a near detector located a few hundred meters from the straight section, the precise profile of the contributions from the three sections to the downstream detectors can be evaluated with statistics 10^5 times larger than will be available in the far detector. We assume a dense, fine-grained detector consisting of tungsten or other heavy plates interspersed with hodoscopes or liquid-argon calorimeters. This could be located in a 52 m deep shaft, 200 m downstream of the decay straight, where the three “beams” would have a Gaussian radius of about 1 m. Rates could be high, on the order of 25 events/fill for a 1-m-thick detector (~ 0.5 kHz). With this rate, it may be possible to measure and subtract the background contribution of neutrinos produced in the upstream and downstream matching sections. Such a system needs further analysis, and is not considered part of the baseline design.

The polarization of the muons in the storage ring can be measured by looking at the momentum distribution of the decay electrons moving in the beam direction. Roughly 8×10^4 decays/m per turn will generate about 100 W/m of signal. These electrons are produced close to the beam as part of the fan of electrons that is swept inward by the bending magnets. The muon beam polarization is measured from the electron decay spectrum. A shower calorimeter, which can absorb the shower from forward-going electrons close to the beam, can be instrumented to look at the power deposition rate at the μs time scale. Detection with a calorimeter should be relatively linear with electron energy. The precision with which the Monte Carlo calculations could be done would be crucial, and this work is under way.

Bibliography

- [1] M. Berz. COSY INFINITY Version 8 reference manual. Technical Report MSUCL-1088, National Superconducting Cyclotron Laboratory, Michigan State University, East Lansing, MI 48824, 1997. see also <http://cosy.nscl.msu.edu>.
- [2] K. Makino, C.J. Johnstone, M. Berz and B. Erdelyi. Nonlinear effects in the 20-GeV Feasibility II muon storage ring. Technical Report MSUCL-1197, National Superconducting Cyclotron Laboratory, Michigan State University, East Lansing, MI 48824, 2001.
- [3] M. Berz, B. Erdelyi, and K. Makino. Fringe field effects in small rings of large acceptance. *Physical Review ST-AB*, 3:124001, 2000.
- [4] K. Halbach, “A Program for Inversion of System Analysis and Its Application to the Design of Magnets,” Lawrence Livermore National Laboratory report UCRL-17436 (1967); CONF-670705-14
- [5] J. H. Billen and L. M. Young, “POISSON/SUPERFISH on PC Compatibles,” Proceedings of the 1993 Particle Accelerator Conference, Vol. 2 of 5, 790-792 (1993).
- [6] M. Berz. *Modern Map Methods in Particle Beam Physics*. Academic Press, San Diego, 1999.
- [7] M. M. Gordon and T. Taivassalo. The z^4 orbit code and the focusing bar fields used in beam extraction calculations for superconducting cyclotrons. *Nuclear Instruments and Methods*, 247:423, 1986.
- [8] C. Crisan and S. Geer, *How well do we need to know the beam properties at a neutrino factory?* MUC Note 87 and FERMILAB-TM-2101, February 13, 2000.
- [9] Y. Semertzdis and W. Morse, Unpublished Note, BNL (2000).