Specifications for Third-Integer Resonant Extraction from the Cornell Synchrotron

Jeremy Perrin

August 23, 2014

1 Overview

Cornell's proposed dark photon search (EAGER) requires a high-intensity positron source to generate statistically significant results. One available option is direct extraction of these positrons from the Cornell Synchrotron, which provides a potential 10^8 positrons per cycle, at 60 cycles per second. This corresponds to a time-averaged current of about 1 nA. However, single-turn extraction of the entire beam would overwhelm the detector by generating many concurrent events. Therefore, it is necessary to extract the beam over many turns. This report discusses the potential benefits and drawbacks of third-order resonant extraction, which is one of several potential slow extraction methods.

When a particle beam has a tune near a third-integer, a localized sextupole moment can drive a resonance that shrinks the stable region of phase space in a controlled manner. This resonance can be understood intuitively by examining the horizontal transverse (or x-direction) phase space trajectory of a single particle with third-integer tune. Assume that the particle is initially located immediately before the sextupole element with an x betatron phase of zero. It first receives a strong positive kick from the sextupole, which has an x^2 field dependence. However, since the particle advances by a phase of $2\pi/3$ each turn, we can see that the next two kicks cancel each other out. After completing three full turns, the particle receives another strong positive kick from the sextupole.

Allowing the emittance and initial betatron phase of the particle to vary, we find that the resonance limits the stable phase space to a triangular region which shrinks as the sextupole strength is increased or the accelerator tune is moved toward the third-integer. (In principle, this resonance has no stopband – there is always some region of stable phase space unless the tune is exactly a third-integer.) Since a particle beam will have a range of phases and emittances, we can use this fact to selectively extract particles from the beam at a rate of our choosing. As the stable phase space shrinks, particles with large emittances are extracted first, followed by particles with smaller emittances. If the process is sufficiently adiabatic, particles with the threshold emittance become trapped at fixed points located at the corners of the stable triangle. A small decrease in the region of stable phase space then causes these particles to stream out along well-defined trajectories, allowing for controlled extraction of the beam.

While it is simple to show that this technique works in principle, there are numerous complications that must be addressed before resonant extraction can be implemented at any specific accelerator. In our case, the Cornell Synchrotron functions primarily as a pre-accelerator for CESR, a synchrotron light source and storage ring for beam dynamics research. As its primary function, the Synchrotron receives particles from a linac, accelerates them from around 120 MeV to 5 GeV, and injects them into CESR, all at 60 cycles per second. Therefore, all extraction magnets must operate as variations on this 60 Hz cycle, and particles must only be extracted near the peak of each cycle, when energies are sufficient to produce meaningful events.

In addition, the beam is adiabatically damped as it is accelerated. To accommodate the beam at low energies, the extraction septa must provide a significantly larger aperture than would be ideal for the extraction procedure. The result, as will be described later, is an irregular and large-emittance outgoing beam. It is possible to overcome this problem by introducing a bump in the closed orbit that moves the beam closer to the extraction septa. Introducing a closed bump also gives experimenters an additional degree of control, which could be useful if the lattice parameters are not exactly as the simulation predicts. These complications and others will be discussed in more detail in the following sections.

2 The Simulation

The results described in this report were found using a simulation built on the BMAD library, which allows for the detailed tracking of particles through a synchrotron lattice. Aside from its generalized tracking algorithms, specialized routines in the library calculate closed orbits and perturbations, including the chromaticity, dispersion, and beta functions of a given lattice. The resonant extraction simulation builds on the standard library functions by implementing time-variation in element strengths and beam emittances, which is necessary for a detailed examination of resonant extraction. Example code for the simulation can be found at `lepp.cornell.edu/~jdp279/darkphoton/programs/singletrack.f90` for single particle tracking and `lepp.cornell.edu/~jdp279/darkphoton/programs/resonantextraction.f90` for beam tracking.

2.1 Emittances, Apertures, and Beta Functions

Due to the lack of precise beam-measuring instrumentation in the Synchrotron, it was necessary for the simulation to make some assumptions about the machine's apertures and beam emittances. It is known that many of the synchrotron's bending magnets have a physical vertical aperture of 1.27 cm, and existing injection and extraction septa have a horizontal aperture of roughly 3 cm. Particles violating these apertures were therefore removed from the simulation. In addition, the bending magnets have a horizontal good-field aperture of roughly 2.5 cm, or slightly more. To take this into account, particles with a transverse position of greater than 3.5 cm were assumed to be lost, and care was taken to violate the 2.5 cm aperture as little as possible during extraction.

The simulation predicts beta functions at injection as shown in Figure 4. Together with these results, the apertures described above set limits on the maximum injection emittances

of the beam of 6×10^{-6} m·rad in the x-direction and 1×10^{-6} m·rad in the y direction. In addition, the dispersion of the lattice sets a limit in the energy variation at injection of roughly one part in 3000. The simulation assumes that the beam is injected with these maximum emittances. The beam is adiabatically damped as it is accelerated, so the emittances and energy variation at extraction are directly proportional, though not equal, to these values. The extraction procedure remains sound when the emittances at extraction are varied by factors of two, though significant modifications of certain extraction parameters are necessary when variations are increased to a factor of four. In particular, the transverse position of the second (magnetic) septum, the strength of the closed bump, and the specifics of the quadrupole sweep need to be adjusted somewhat to compensate for the change in emittance. For more details on these extraction parameters, see sections 6, 5, and 3, respectively.

The beta functions and betatron phase at the locations of the various extraction magnets are crucial to the success of the extraction procedure. In addition, the extraction procedure needs to be modified significantly if the fractional part of the synchrotron tune is not as the simulation predicts. As an input, the simulation has a lattice file detailing the positions and characteristics of the synchrotron bending magnets and quadrupoles. However, it is somewhat uncertain how closely the lattice file follows the actual accelerator, and how accurately the simulation can predict lattice parameters from a given input. It is known that the input file accurately gives the pattern of bending magnets (in particular, a repeating FOFODODO pattern with four short and two long drift insertions) and includes small adjustments in the locations of several magnets for fast extraction to CESR. Without these adjustments, the simulation predicts very regular lattice functions (see Figure 2), which are similar to those of a generic repeating FODO lattice (with the correct number of periods!). As expected, the existing adjustments cause some irregularities in the predicted beta functions (see Figure 3). These conclusions indicate that the simulation predicts the qualitative behavior of lattice functions with some accuracy. Regardless, it may not be a good idea to take these results on trust. If possible, it is highly recommended to measure the lattice functions and beam emittances prior to modification of the Synchrotron, or at the very least to confirm the detailed accuracy of the field strengths, position offsets, and apertures in the input lattice file. As a guide, the sensitivities of various elements of the resonant extraction procedure to changes in the lattice parameters will be discussed throughout this report.

The modified lattice file used for the simulation of resonant extraction can be found at `lepp.cornell.edu/~jdp279/darkphoton/files/layout/synch_layout_2.bmad`, and an unmodified lattice file, which likely represents the current state of the accelerator, can be found at `lepp.cornell.edu/~jdp279/darkphoton/files/layout/orig_synch_layout.bmad`.

3 The Quadrupoles

As described earlier, there are two parameters that control the size of the region of stable phase space – the synchrotron tune and the sextupole strength. Using the sextupole strength as the control parameter raises serious issues, since its strength must vary over at least an order of magnitude to sufficiently shrink the region of stable phase space. More precisely, the threshold sextupole strength for extraction of the highest-emittance particles depends strongly on the fixed detuning from resonance, but there is a tune-independent factor by

which the sextupole strength must increase to extract, say, 95 percent of the beam. The inductance of the sextupole magnet could make it very difficult to achieve the required agility and precision of control. It is therefore much simpler to fix the sextupole strength during extraction and vary the tune for fine control. The tune can be controlled by adjusting the strength of four main quadrupoles, located in L0 and L3. At the moment, these quadrupoles are used to keep the lattice parameters regular through the two long drift sections. By increasing their strength by roughly 14 percent (or around 1 T/m), we can adjust the fractional part of the tune from 0.614 to 0.667, pushing the beam onto the third-integer resonance.

We wish to extract the beam over as long a period as possible for several reasons. First, slower extraction reduces pile-up at the detector. The beam intensity in the synchrotron can then be increased, allowing for a higher integrated luminosity. Second, a slower extraction allows for finer control of the trajectory and time profile of the outgoing beam. As described earlier, the theory of the extraction procedure assumes that particles at the threshold emittance are trapped at the corners of the stable phase space, at which point they stream out in a controlled manner. Faster extraction makes this approximation less accurate, and the outgoing beam quality suffers as a result. In addition, the number of particles extracted per unit time becomes much more difficult to control when the strength of the resonance changes quickly. Finally, a faster extraction process requires rapid changes in the detuning from resonance, placing significant demands on the extraction magnets and their power supplies.

However, the energy of the circulating beam varies at 60 Hz. Therefore, when we extract over a longer period, the energy of the outgoing beam varies significantly over time. This isn't a serious issue for event reconstruction, since we know what the energy is at all times. Regardless, higher energy particles have a capacity to produce more interesting events. We have found that extraction over roughly 1000 turns best balances these effects, as the outgoing particle energy changes by at most 6 percent over this time.

To obtain an even extraction over 1000 turns, the region of stable phase space must shrink in a way that excludes a constant number of particles per unit time. Obviously, the details of the quadrupole sweep must then depend closely on the beam emittance and the distribution of particles in phase space. It will therefore be necessary to create a multi-parameter tunable sweep to account for the unknown beam parameters. The specific sweep described below serves as an example of what must be done in general, but is certainly not precisely what will be required.

The simulated sweep was pieced together from seven linear sections, all multiplied by the 60 Hz sinusoidal base sweep. The first and last section have unit value – that is, the quadrupole sweep is unchanged from its sinusoidal dependence in these sections. The second section raises the quadrupole strength to near the extraction threshold for the highest emittance particles in the beam, while the sixth lowers it back to its baseline. The remaining three sections have small, tunable slopes which cause the beam to spill out in a reasonably linear manner. Figure 5 shows the sample quadrupole sweep that was found to work in the simulation, both on its own and multiplied by the 60 Hz base sweep, and Figure 6 shows the resultant time profile of the extracted beam.

The question then arises of the magnets' ability to handle the varied sweep profile. Here it seems we have two options – either input the varied signal into the existing quadrupoles, or construct new magnets with smaller inductances (say, air-core quadrupoles) to handle the

variations in the sweep while the existing quadrupoles operate on their standard 60 Hz sweep. These new quadrupoles would be placed directly adjacent to the existing quadrupoles.

While there are no easily locatable data on the inductance of the existing quadrupoles, a calculation based on estimates of the magnet parameters puts their inductance at around 2 mH. In contrast, separate air-core quadrupoles designed to make up the difference in field would have an inductance of roughly 0.3 mH, given a maximum current of 400 A. Enclosing the air-core quadrupoles in a box of ferrite effectively produces a magnetic short-circuit – when the maximum current and required magnetic field are held fixed, the ferrite box reduces the required number of coils by roughly a factor of two. Since the inductance of a magnet is proportional to the number of turns times the magnetic field per unit current, and both the field and current are held fixed, the addition of a ferrite box reduces the inductance by this same factor of two.

In addition, the presence of a ferromagnetic conductor in the existing magnets introduces nonlinear effects (in particular, eddy currents and hysteresis loss) that could limit their bandwidth irrespective of their inductance. It therefore seems that constructing new trim quadrupoles seems to be the better option in this situation. Note, however, that the higher currents in these magnets require fairly thick wires and/or cooling circuits to prevent overheating, and these may not be suitable for the geometry of the magnet. The frequency response of the existing magnets will need to be investigated in more detail before a decision can be made between these options.

The results of a Fourier-domain analysis of the trim quadrupoles' input signal are shown in Figure 7. In these figures, the trim quadrupoles are modeled as a resistance and inductance in series. The required current signal (Figure 5b) is Fourier-decomposed, and each mode is multiplied by the voltage-to-current ratio at that frequency, which is given by

$$V_n/I_n = L\left(R^2/L^2 + \omega_n^2\right)^{1/2}$$

The signal is then synthesized from the various Fourier modes. Note that a larger resistance-to-inductance ratio makes the voltages higher, but the signal more regular.

4 The Resonant Extraction Sextupole

A single sextupole must be placed in the lattice to generate the required third-integer resonance. As was previously stated, the field of the resonant extraction sextupole can follow the 60 Hz sweep during extraction, which should simplify the requirements for its power supply somewhat. However, the sextupole needs to be turned off during injection. At this time, the beam's large emittance causes the sextupole to induce significant irregularities in the beam size, despite the fact that the tune is far off resonance (see Figure 8). Assuming the injection emittances and apertures given earlier, this effect results in a loss of roughly half the beam. To prevent this, we must ramp the sextupole strength as shown in Figure 9. The ramp shown starts at turn 500 and ends at turn 1500 (with one full oscillation occurring over 6610 turns), but it may be extended to turn 2000 or otherwise smoothed out if the inductance of the sextupole prevents this level of agility.

Assuming the pole tips of the sextupole are 5cm from the center of the beam and the length of the magnet is 0.5 m, the optimal field at the pole tips during extraction was found

to be roughly 0.21 T. In other words, the sextupole should have a field strength of 84 T/m^2. For comparison, the sextupoles in CESR, which are used primarily to correct the lattice chromaticity, have a magnetic length of 0.23 meters and a field strength of about 1 T/m^2.

4.1 Location Sensitivity

The horizontal beta at the location of the sextupole modulates its apparent strength due to the strong radial dependence of its field. While it isn't necessary to place the sextupole directly at a location of maximum beta, this dependence should be kept in mind when choosing a location. More globally, the location of the resonant extraction sextupole determines the orientation of the stable phase space triangle and its outgoing trajectories everywhere in the lattice. Therefore, varying its position provides a good test of the sensitivity of the extraction procedure to more general uncertainties in betatron phase.

Due to space considerations, the location of the sextupole is much more flexible than the locations of the extraction septa, since the latter determines where the experiment itself will be held. Fixing the septa in their most convenient locations for extraction at L0, it was found that the resonant extraction sextupole could be placed exactly at the center of L4, a 6-meter empty section of the accelerator that once contained an RF cavity.

Varying the location of the sextupole, it was found that variations of up to 1.5 meters did not seriously influence the resonant extraction procedure. This adjustment corresponds to a betatron phase shift of 0.15 radians and a 17 percent change in the transverse beta at the sextupole. Variations at this scale could be corrected by fine-tuning the details of the quadrupole sweep. Beyond this point, though, other changes needed to be made – in particular, the changes in the transverse beta at the sextupole and the orientation of the outgoing trajectories meant that the extraction septa had to be moved transversely. (An adjustment of the relative transverse location of the septa compensates for changes in the orientation of the outgoing trajectories, while an adjustment of the absolute transverse location of both septa allows the extraction point to be moved inward in case of aperture violations elsewhere in the lattice.) With more significant changes, even these adjustments were insufficient. These results are summarized in Figure 10, which shows the how the extraction efficiency changes as the sextupole location is adjusted. Note that the overall phase relative to the sextupole isn't the only phase variable that matters in extraction – the relative phase of the extraction septa is also important, as well as their locations relative to the bump magnets, all of which will be described later.

4.2 Effects on Lattice Chromaticity

The unmodified lattice file presents a model of the accelerator with no sextupole moments whatsoever. As expected, the result is a negative chromaticity roughly equal to the tune of the accelerator, in both the x and y directions. The file also includes nominal values for the sextupole moments of the bending magnets, though they are 'commented out' so as to be ignored in actual calculations. Including them changes the x chromaticity from -12.9 to -11.8, and the y chromaticity from -12.7 to +3.3. The effect on the x chromaticity is small due to a strong cancellation from the focusing and defocusing magnets. Since the level of precision in the lattice file is unknown, these numbers may not be reliable, so they are best

used only as a comparison to the following values. Once the resonant extraction sextupole is turned on, the chromaticities change from -11.8 to -0.9 in the x direction and +3.3 to -4.6 in the y direction. Thus we see that the resonant extraction sextupole improves the x-chromaticity while worsening the y-chromaticity. Ignoring charge-dependent effects, these changes didn't seem to have much of an effect on extraction, at least when the variation in the particle energy at injection was assumed to be one part in 3000. The beam stability due to the chromaticity changes from the sextupole may or may not be an issue, though it's worth noting that the charge per bunch is fairly small. For example, if the Synchrotron delivers an outgoing current of 1 nA, and the circulating beam is divided over 150 bunches, the charge per bunch is 1×10^{-13} Coulombs. To examine these effects in a simulation, though, we would first need to better understand the Synchrotron's sextupole fields, and their variation between different bending magnets.

5 Bump Magnets

The introduction of a closed-orbit bump near the extraction septa is necessary to correct a rather serious issue caused by the 60 Hz energy oscillation of the Synchrotron. As described in the introduction, the beam is adiabatically damped as it accelerates from 120 MeV to 5 GeV. At injection, its emittances are therefore a factor of 40 or so larger than they are at extraction. During extraction, particles stream out along well-defined trajectories starting at the corners (fixed points) of the triangular separatrix, as shown in Figure 11. At a fixed x-coordinate, the step size along these separatrices (or the increase in a given particle's x-coordinate every three turns) depends strongly on the detuning from resonance and the resonance strength. To generate a sufficiently strong resonance (or sufficiently small detuning) to extract the beam after acceleration and adiabatic damping, the step size at the septa turns out to be fairly large, on the order of 1 cm. While this means that particles easily clear the septum without noticeable loss, the outgoing beam has a large emittance and is fairly irregular. The outgoing beam line then needs to have a larger width, requiring the outgoing focusing magnets to be stronger. Moreover, the septa need to have an aperture larger than the step size, and increasing the aperture of a given septum increases its required voltage or integrated magnetic field.

We want a way to manually control the trade-off between beam loss at the septum and the quality of the outgoing beam. As previously stated, we can do so by introducing a bump in the closed orbit at the extraction septa. This brings the beam closer to the septa at extraction, making the effective x-coordinate of the septa smaller, and therefore decreasing the step size at extraction. To create a closed bump with a betatron phase advance of roughly 2π, we place three dipole magnets near the septa. The first and last magnets open and close the bump, while the second magnet makes the bump slightly longer or shorter than 2π, as need be.

For extraction at L0, the first (electrostatic) and second (magnetic) septa are placed at T9 and T6, respectively. Descriptions of these septa and reasons for placing them at these locations will be explored in the following section. To create a closed bump with a maximum between these two septa, we place the first bump magnet at T12, the second 2.5 meters before the end of L0, and the third at T189. See Figure 1 for a visual representation

of this layout.

The strengths of these magnets scale roughly as the inverse square root of the beta functions at their location, since a larger beta makes a perturbation in momentum more significant. To create an appropriately-sized closed bump in the simulation, it was found that the first magnet should have an integrated kick of +0.00121 radians, the second +0.00057 radians, and the third -0.00178 radians. For an 0.3 meter magnetic length and a beam energy of 5 GeV, these kicks correspond to dipole field strengths of 0.067 T, 0.032 T, and 0.099 T, respectively.

The closed bump magnets must be turned off at injection to prevent aperture violations. Since they follow the 60 Hz sweep at extraction, the sweep pattern for these magnets can be made identical to that of the sextupole. See Figure 9 for details.

The description above assumes that the bump is created using separate dipole magnets placed between the existing dipoles, but conceivably, the bump could also be implemented by adding trim coils to the existing dipole magnets, allowing their strength to be modulated slightly. This method has a potential benefit – the existing dipoles would not need to be moved to accommodate the bump magnets. As we will see in the next section, though, this results in only a nominal reduction in perturbations to the lattice functions. The objection may be raised that adding trim coils to the existing dipoles modulates their quadrupole moments as well, but tests show that these effects may be compensated by moderate adjustments in extraction parameters. Therefore, both implementations of the closed bump could be successful.

6 Septa

After particles leave the region of stable phase space, they travel on trajectories that spiral outward from the closed orbit. After roughly 30 turns, they are far enough away from the closed orbit that they encounter the first of two extraction septa. Depending on the step size of the particles on the outgoing trajectory, a number of these particles will hit the first septum, and the remainder will clear it. Therefore, making the first septum as thin as possible can greatly improve extraction efficiency. Magnetic septa with appreciable kicks (upward of 5 mrad) are limited to a thickness of roughly 1 mm due to the inverse relationship of wire area and resistance. Electrostatic septa, on the other hand, can be made as thin as 50 μm. Kicks from these septa, though, are limited to a few tenths of a milliradian, depending on the septum length and the electrostatic breakdown strength of the accelerator vacuum. Therefore, we use a two-septum system to extract particles from their outgoing trajectories – one electrostatic septum, and one magnetic septum. The strengths of these septa should follow the 60 Hz sweep of the Synchrotron dipoles, at least during extraction. In addition, the strength of magnetic septum may need to be modulated somewhat, as will be described in the following section.

6.1 Specifications

The simulated electrostatic septum is assumed to have a width of 100 μm, and is placed at T9. This location was chosen for two reasons. First, the magnetic septum should be placed

roughly $\pi/2$ radians in betatron phase downstream of the electrostatic septum. Doing so is important for the efficiency of the septa, because the $\pi/2$ phase advance projects the kick of the first septum fully onto the second septum. Based on space considerations for the target and particle detector, the magnetic septum must be placed at T6 or T7, which limits the possible locations for the electrostatic septum.

Second, both the electrostatic and magnetic septa must be placed at locations with large beta, so that the kick of the electrostatic septum may be amplified. It turns out that this consideration takes precedence over the first, as long as the phase advance is somewhat close to $\pi/2$ radians. Since T9 is close to a local beta maximum, this location was chosen for the electrostatic septum. T6, which is 0.9 radians downstream of T9, was chosen as the location of the magnetic septum. Placing this septum at T5 would be beneficial for extraction, since this location has a larger beta than T6 and the phase advance between the two septa would be closer to $\pi/2$. However, this change would probably not leave enough space for the experiment itself.

So that none of the beam hits it at injection, the electrostatic septum is placed 3.10 cm from the center of the beam pipe. Based on some trial and error and the ratio of beta functions at T9 and T6, the center of the magnetic septum is then placed 2.52 cm from the center of the beam pipe. As demonstrated in section 4.1, it is very helpful if the transverse position of one of the two septa is easily adjustable.

The magnetic septum was assumed to have a width of 1 mm and a kick of around 8 mrad. These numbers may be fairly challenging to realize, but they are certainly within the realm of possibility. For comparison, the two existing septa for fast extraction to CESR are significantly thicker, and have kicks of 5 and 8 mrad. These are preceded by a magnetic kicker with a strength of 1.2 mrad.

Clearing the 1 mm magnetic septum requires an electrostatic septum kick of 0.17 mrad. Assuming the electrostatic septum has a length of 0.5 m, its field would need to be 1.7 MV/m to generate such a kick. The required field of the first septum is roughly proportional to the width of the second.

Given the closed bump parameters described in the previous section, the step size at extraction is always less than 5 mm at both septa. Therefore, septum apertures of 5 mm are sufficient to prevent particles from hitting their outer walls.

6.2 Effects of Magnet-Moving

Since the two septa have lengths of around 0.5 meters, the adjacent bending magnets must be moved longitudinally outward to accommodate them. For the septa, the surrounding magnets are moved by 17 cm each. Assuming the closed bump is implemented using new magnets, these will also need to be accommodated via magnet-moving. For the first and third bump magnets, the surrounding magnets are moved by 9 cm each. The second bump magnet is located in L0, and therefore does not require any accommodation.

As with the magnet-moving for the existing extraction septa, these adjustments will cause perturbations to the lattice functions. The current lattice functions are given in Figure 3, while the lattice functions with all magnets moved are shown in Figure 4. Note that the majority of this change comes from the magnet-moving for the septa, and the implementation of the closed bump is relatively inconsequential.

6.3 Extraction

At this point, it may be helpful to provide a summary of the beam's behavior during extraction. As the triangular stable phase space of the beam shrinks, particles stream outward from its corners, their phase space trajectories rotating with the betatron phase. Since the tune is close to a third-integer, the outgoing particles will return to a given oscillation phase after three turns. However, their x-position has increased somewhat, by a value depending on the particle's position as it passes the extraction sextupole, and the strength of (and detuning from) the third-integer resonance. If everything is aligned correctly, the outgoing particles skim the inside of the magnetic septum three turns before extraction. Without the electrostatic septum, many would hit the magnetic septum on their following pass. On its extraction turn, a particle will clear the electrostatic septum, receiving a kick that places it just outside the magnetic septum. When everything is aligned correctly, the loss at the electrostatic septum is roughly 5 percent (for a 100 μm septum width) and the loss at the magnetic septum is similar. Turning up the kick from the first septum slightly (say, to 2 MV/m) and adjusting the transverse position of the second can eliminate the loss at the second septum altogether. In either case, the result is an extraction efficiency of between 80 and 90 percent, with between 5 and 10 percent of the beam remaining in the Synchrotron after extraction.

Single-turn snapshots of the horizontal phase space of the beam at the electrostatic and magnetic septa are shown in Figures 12 and 13, respectively. The trajectory of a single particle, tracked over the ten turns prior to extraction, is shown in Figure 14.

7 The Outgoing Beam

If the magnetic septum is placed at T6, the outgoing beam will have about 10 meters before the target. This space may be used to focus the outgoing beam in various ways, minimizing the experimental uncertainty that arises from its phase space profile.

7.1 Time-Varying Dipole Kicks

Without adjustment, we find that the average x-position and momentum of the outgoing beam change somewhat with time. We need a way to compensate for this drift, since it increases the time-integrated horizontal emittance of the outgoing beam and introduces a potential source of systematic uncertainty.

The drift can be caused by three separate effects. First, the 'closed bump' isn't always precisely closed. Since betatron phases and beta functions change with the main quadrupole sweep, the strengths of the bump magnets would need to change precisely with time to compensate. (This effect is reasonably minor, and doesn't influence any other aspects of extraction.) Second, the strength of the resonance changes with time, causing the step size at the septa to increase slightly. This increases the average x-position of the beam. Finally, because the outgoing trajectories begin at the corners of the stable phase space, which shrinks with time, their orientation in phase space changes over a cycle of extraction. Tests show that the third effect is the most prominent, though all three contribute somewhat. Note that the size of each effect can be directly related to the main quadrupole strength. In

fact, trial and error has shown that the magnitude of the required position and momentum compensation at any given time can be related *linearly* to the strength of the quadrupole sweep at that time, at least to a sufficient approximation.

Compensation for the drift can be done in a three-step process. First, the strength of the magnetic septum is modulated with time to compensate for the momentum drift. As usual, this modulation occurs as a product with the standard 60 Hz sweep. The magnitude of this compensation at any time is (roughly) linearly related to the magnitude of the quadrupole sweep (see Figure 5). While the quadrupole sweep represents a 14 percent modulation in the strength of the quadrupole magnets, the magnetic septum sweep represents a 5 percent modulation in its strength, assuming its base strength corresponds to an 8 mrad kick.

The position drift can be compensated for by introducing two dipole kicks at 4.5 m and 5.5 m after the point of extraction. The first kick can also act as a general bend magnet, steering the particle onto a trajectory toward the target. A time-dependent modulation on top of this strong bend, coupled with an equal and opposite modulation of the second dipole kick, can change the position of the beam without adjusting its momentum. Since the outgoing beam position drifts by several millimeters during a single extraction cycle, the two time-varying kicks should have strengths of several milliradians to compensate. Again, the required modulations, which are roughly proportional to the trim quadrupole sweep, occur as a product with the standard 60 Hz variation.

7.2 Quadrupole Focusing

Since there will be roughly 10 meters between the target and detector, we may minimize experimental uncertainties by making the standard deviation of the position distribution at the target (in mm) roughly 10 times the standard deviation of the momentum distribution (in mrad). This causes both types of uncertainty to contribute equally to the position at the detector.

To produce a beam with the required phase-space distribution, two quadrupoles are added after the dipole bending magnets. The first is located 6 m after the point of extraction and has a strength of 6.7 T/m, while the second is located 8.8 m after the point of extraction and has a strength of -7.7 T/m. These strengths assume a magnetic length of 0.5 meters and a 5 GeV beam. However, the quality of the focusing remains reasonable throughout extraction without modulating these magnets on a 60 Hz sweep. Therefore, they can be placed outside the beam pipe.

7.3 Results of Focusing

The phase-space profile of the outgoing beam 10 meters after the point of extraction without any focusing is shown in Figure 15. In contrast, the same beam profile with time-varying dipole kicks and quadrupole focusing is shown in Figure 16. The dipole kicks reduce the emittance of the (time-integrated) beam by roughly a factor of three, while the quadrupoles adjust the beam's phase space profile as needed.

8 Summary

As a summary, we will now list each of the elements described above, describe its specifications, and explain the aspects of the extraction procedure that can be modified by adjusting these specifications.

1. **The Resonant Extraction Sextupole.** Located at the center of L4, this magnet operates on a sweep that is zero at injection but sinusoidal at extraction. Its peak strength, assuming a magnetic length of 0.5 m, is around 84 T/m^2. Changing its location modulates its apparent strength (based on the local beta) and the orientation of the outgoing particle trajectories everywhere in the lattice (based on the local betatron phase). Increasing its strength can cause particles to violate apertures before extraction, but this effect can be somewhat compensated for by adjusting the quadrupole sweep and/or moving the extraction septa inward.

2. **The Trim Quadrupoles.** Four quadrupoles, located near the center of L0 and L3. Quadrupole magnets already exist at these locations, and the modulated signal may be fed into the existing magnets or new, adjacent quadrupoles. New trim quadrupoles require a field strength of around 1 T/m. Since the quadrupole sweep controls many aspects of extraction, these magnets must be fairly agile, and their sweep should be easily adjustable. See Figure 5 for the specific sweep that worked in simulation.

3. **Electrostatic and Magnetic Septa.** The septa are located at T9 and T6 respectively. The electrostatic septum should be placed about 3.1 cm from the center of the beam pipe, and the magnetic septum about 2.5 cm from the center of the beam pipe. Both are roughly 0.5 meters in length. Their kicks are 0.17 mrad and 8 mrad, respectively, and they operate on the standard 60 Hz sweep during extraction, except for a 5 percent modulation in the magnetic septum that controls the momentum drift of the outgoing beam. It is important for the transverse positions of these septa to be easily adjustable, by at least a few millimeters, since their optimal locations depend on many unknown or variable extraction parameters.

4. **Closed Bump Magnets.** Three dipole kicks, located at T12, in L0, and at T189, are used to create a bump in the closed orbit at extraction. In simulation, the required integrated kicks of these magnets were +0.00121 radians, +0.00057 radians, and -0.00178 radians, respectively. These magnets must be turned off at injection, and therefore follow a similar sweep to that of the sextupole. Adjusting the overall strength of the closed bump changes the distance from the beam from the septa at extraction, controlling a trade-off between the emittance of the outgoing beam and the fraction of the beam lost at the electrostatic septum.

5. **Outgoing Beam Dipole Kicks.** One strong kick is used for general steering of the beam, and two (equal and opposite) weaker kicks correct for position drift of the outgoing beam. The strong kick may be generated by the same magnet as one of the weaker kicks. The weaker kicks require an agility similar to that of the trim quadrupoles, and have a magnitude of several milliradians. The magnets are located 4.5 m and 5.5 m after the magnetic septum.

6. **Two Outgoing Quadrupoles** to focus the beam onto the detector. These magnets do not need to follow a 60 Hz sweep. In the simulation, the first has a strength of 6.7 T/m and is located 6 m after the magnetic septum, while the second has a strength of -7.7 T/m and is located is located 8.8 m after the magnetic septum.

9 Figures

All of the images below are available at `lepp.cornell.edu/~jdp279/resonant`.

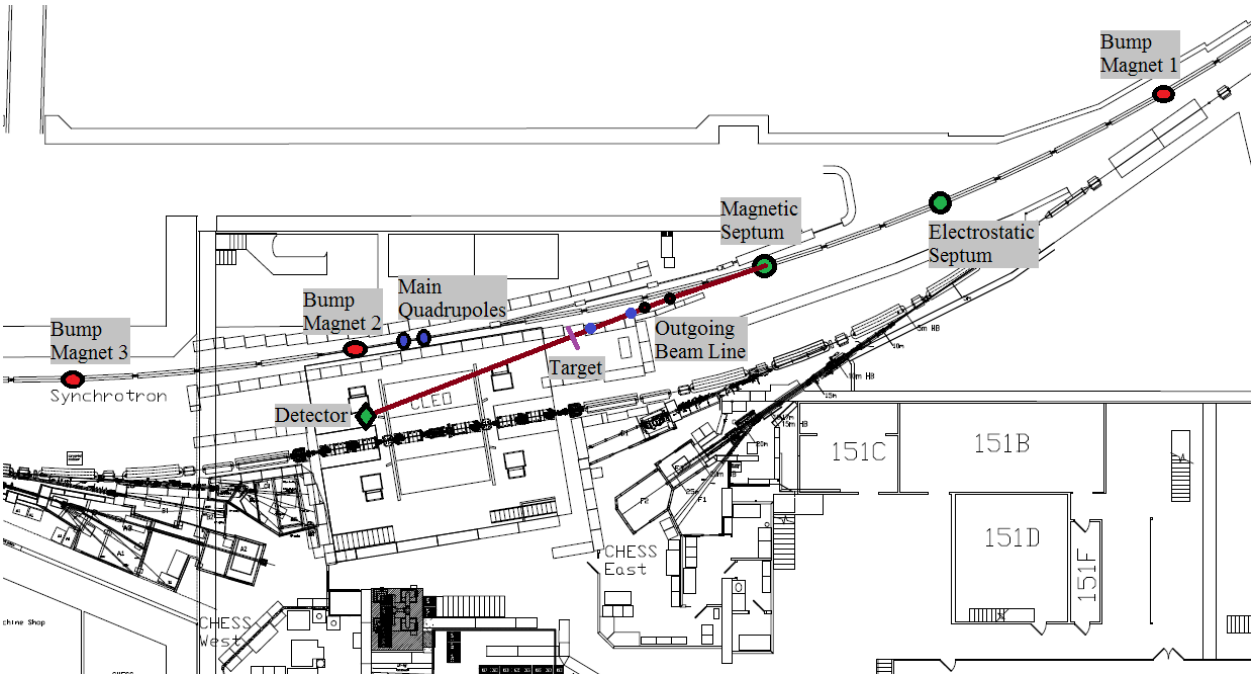

Figure 1: Layout of resonant extraction elements near L0. Not shown are the two main quadrupoles in L3, and the resonant extraction sextupole in L4.

.

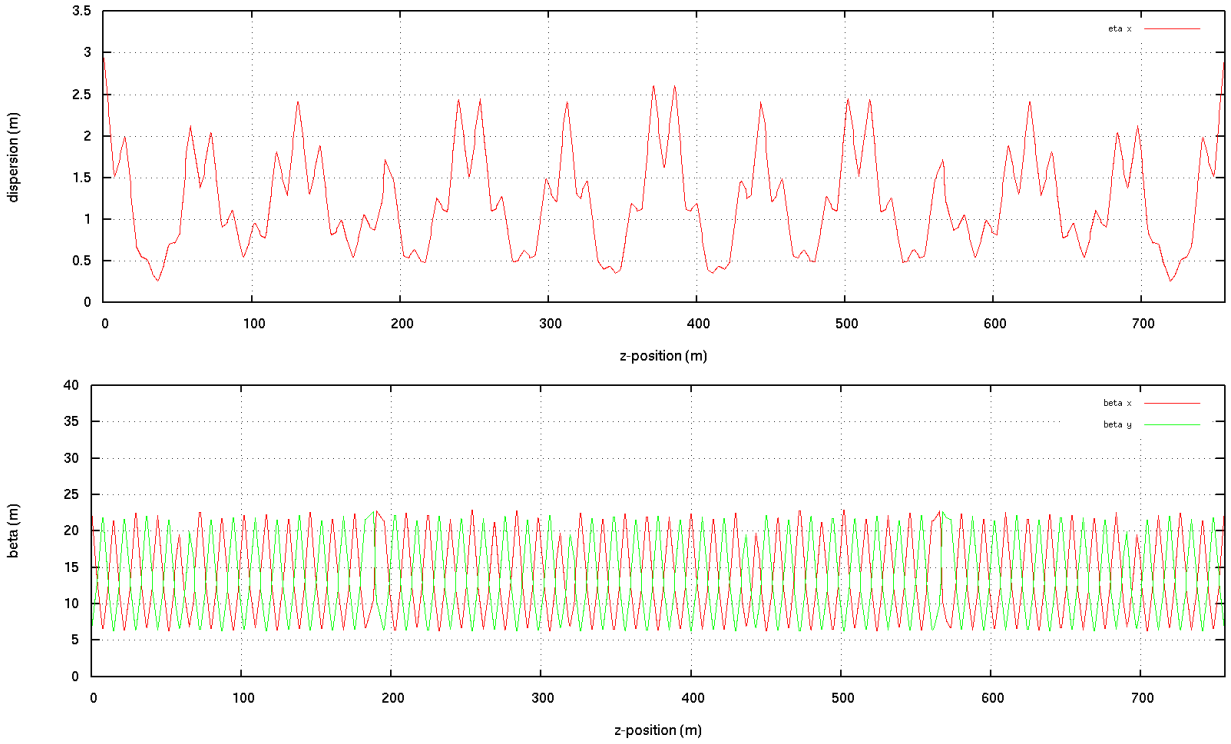

Figure 2: Dispersion and beta functions predicted by BMAD simulation for the original Synchrotron lattice. Existing adjustments in magnets near fast-extraction points are not included. In the second graph, the red line is the vertical beta and the green line is the horizontal beta.

.

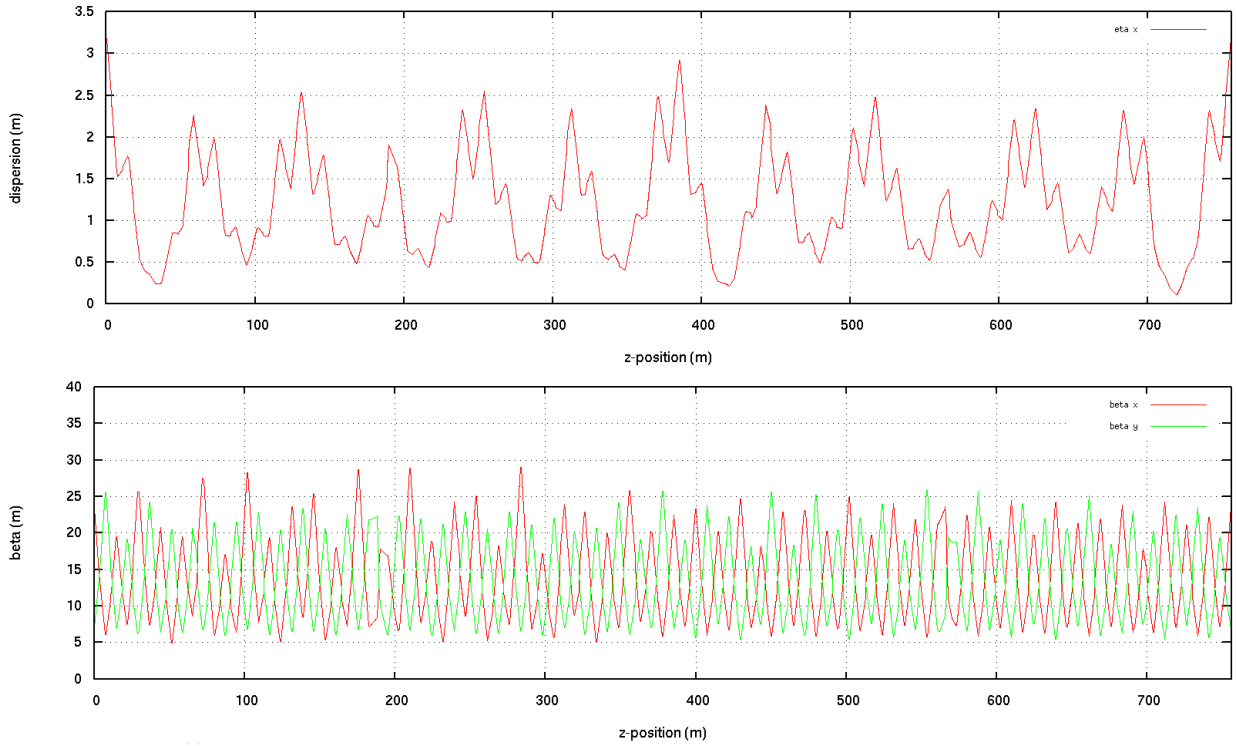

Figure 3: Dispersion and beta functions predicted by BMAD simulation for the current Synchrotron lattice. The differences from Figure 2 are due to magnet-moving to accommodate septa for fast extraction to CESR.

.

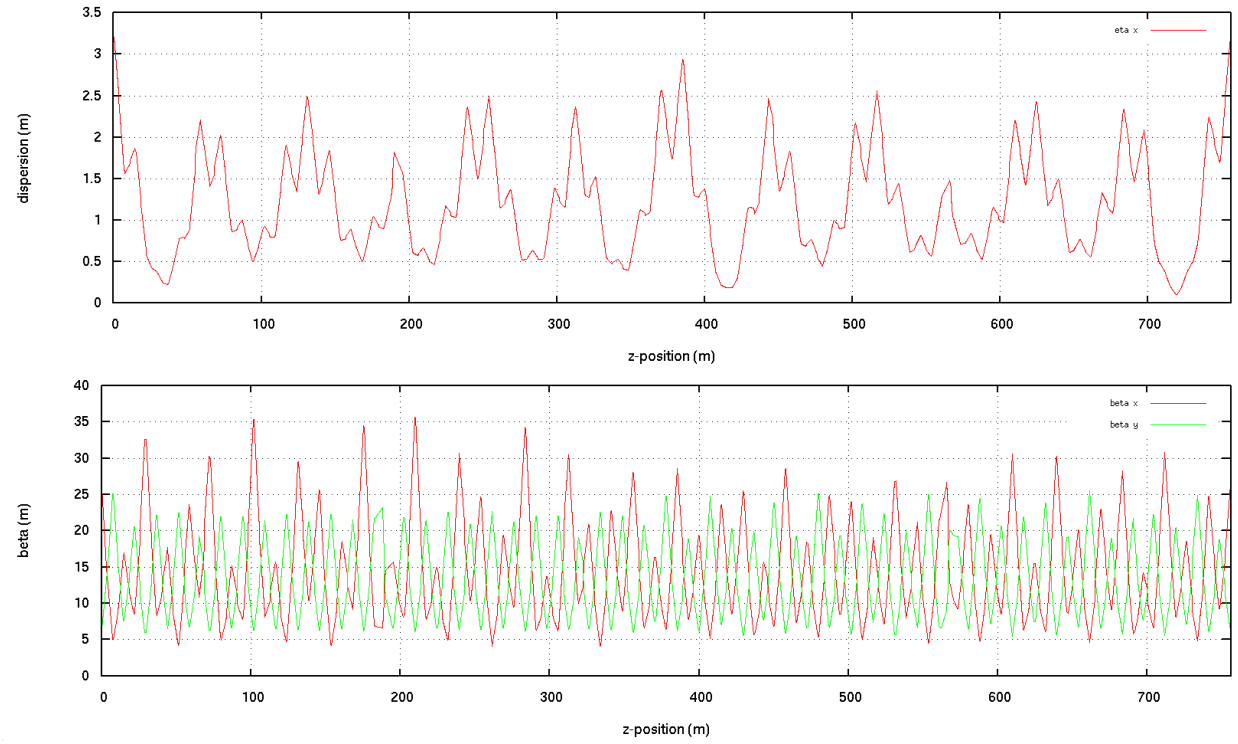

Figure 4: Dispersion and beta functions when dipoles are moved to make room for resonant extraction septa and bump magnets. Note that none of the extraction elements are 'on,' meaning that these changes would influence the Synchrotron in both normal operation and slow extraction.

.

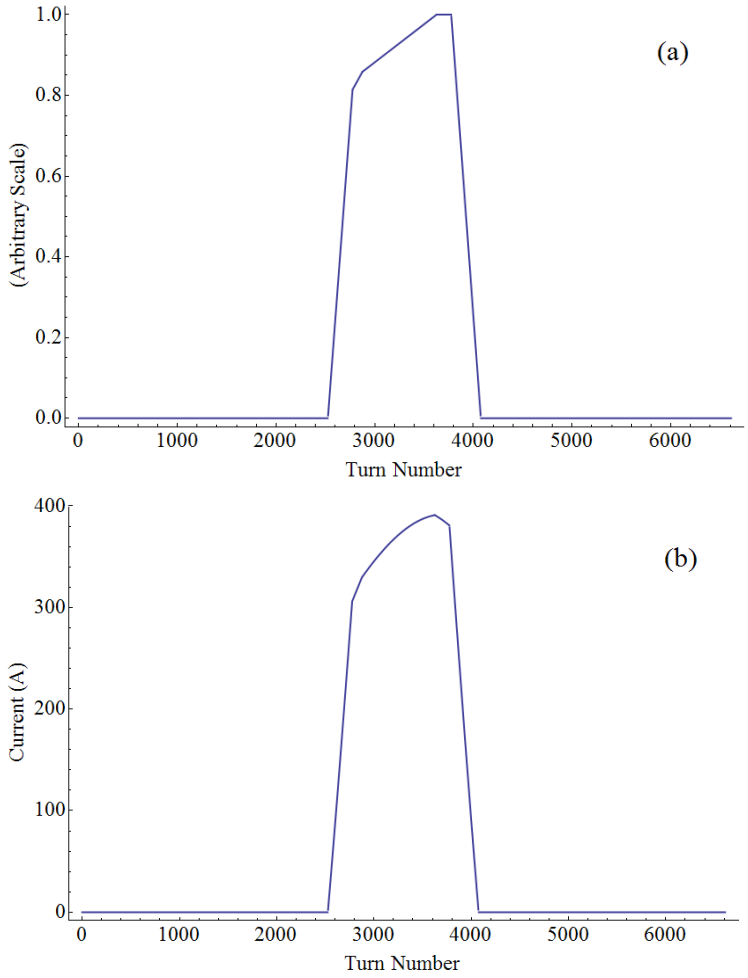

Figure 5: Trim quadrupole sweep used in simulation to push the Synchrotron tune onto the third-integer resonance. Figure 5a shows an energy-normalized sweep, while Figure 5b shows the same sweep multiplied by the Synchrotron's 60 Hz modulation. The latter is shown in units of current, assuming a 400 A maximum current through the trim quadrupoles.

.

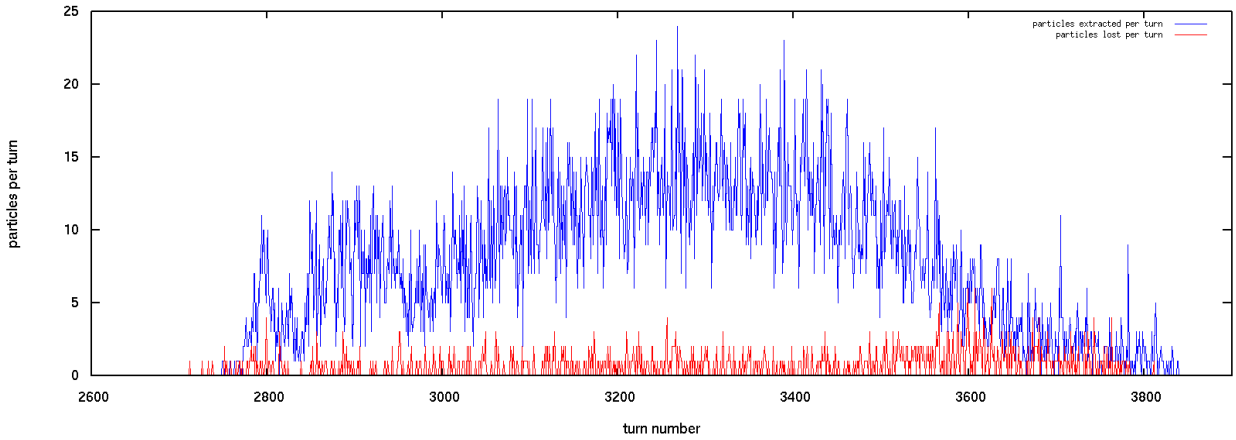

Figure 6: Time profile of the outgoing particle beam for the quadrupole sweep shown above. The blue line represents particles extracted per turn, and the red line represents particles lost per turn. The increase in particle loss near the end of extraction was due to a slight misalignment of the magnetic septum. The simulation for this figure tracked 10^4 particles, so the statistical variations shown here should be significantly greater than those for the actual beam, which should contain about 10^8 particles. However, note that much of the variation (including some local structure) does not appear to be statistical. The quadrupole sweep could potentially be modified to remove this structure.

.

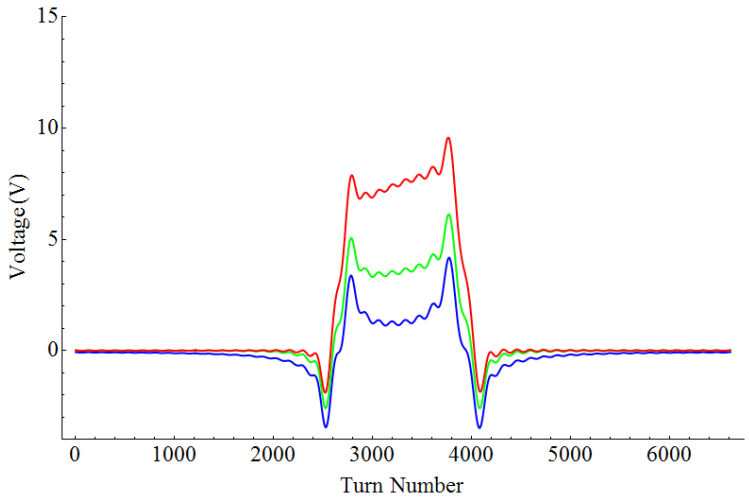

Figure 7: Possible voltage sweeps for an air-core trim quadrupole, assuming the current sweep shown in Figure 5b. The inductance is assumed to be 0.3 mH. The resistance of the quadrupole is negligible for the blue curve, 9 mOhm for the green curve, and 20 mOhm for the red curve. Only the first 50 Fourier modes (up to 3000 Hz) are taken into account.

.

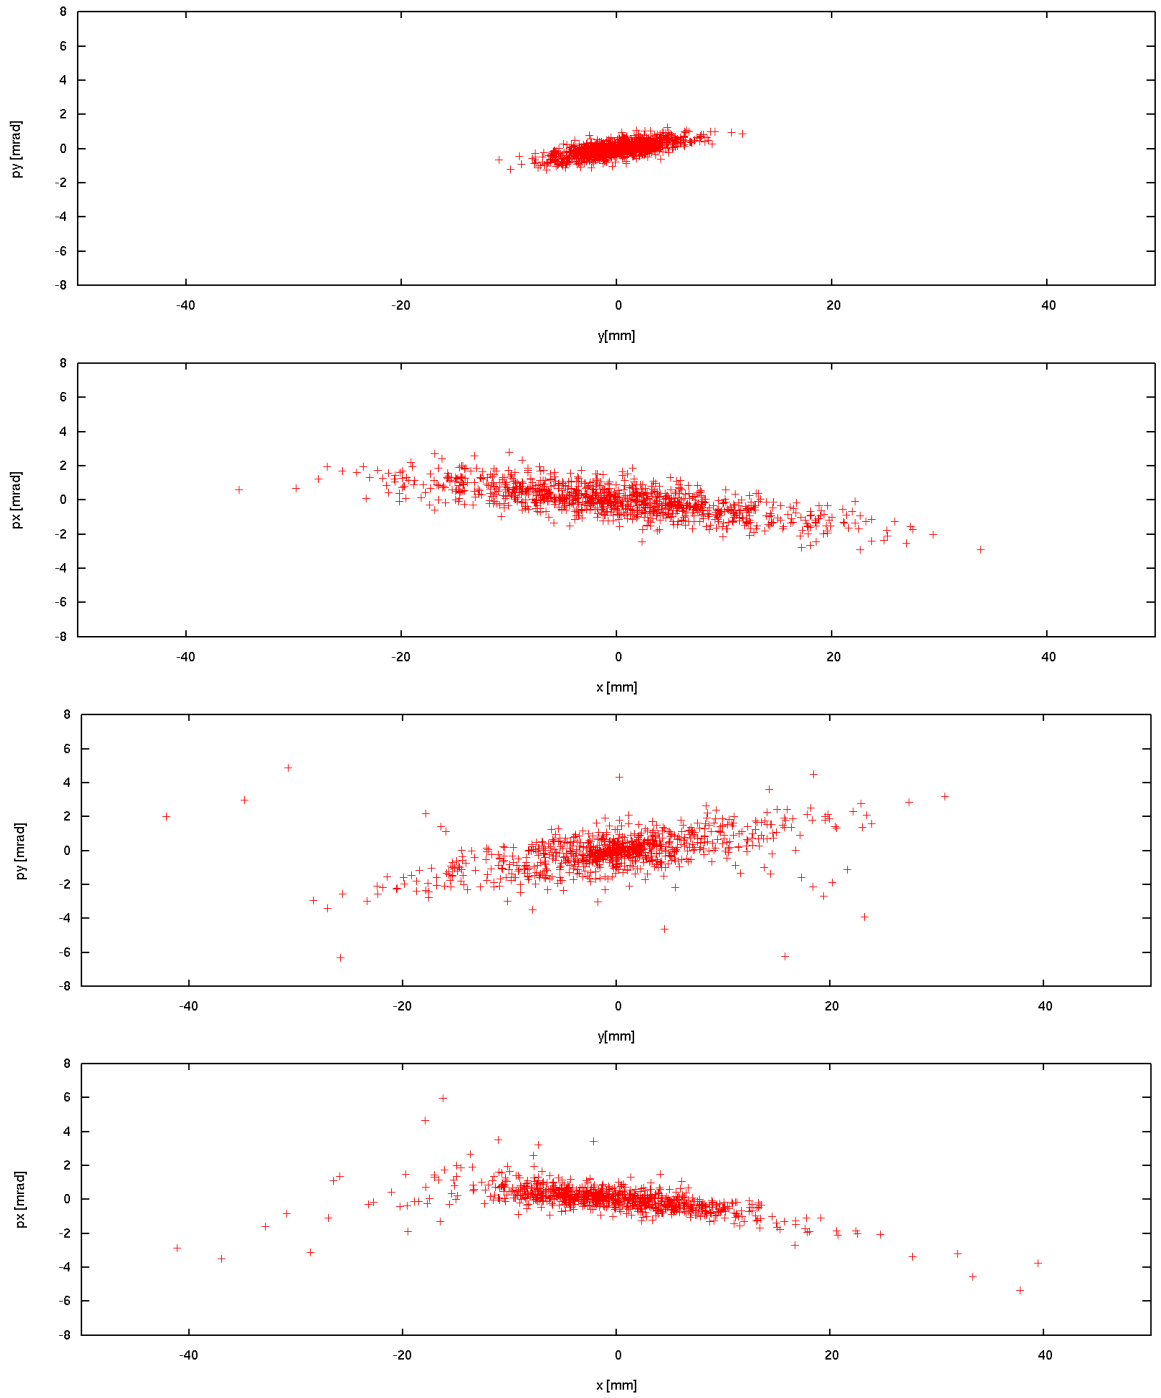

Figure 8: Demonstration of beam loss at injection due to the resonant extraction sextupole. For these images, all apertures in the Synchrotron were turned off so that particles were not removed from the beam at large transverse displacement. The first two figures are snapshots of the y and x phase space of the beam, 10 turns after injection, with the sextupole turned off. The third and fourth figures represent the same situation with the sextupole turned on. Despite the large detuning from the third-integer at injection, the sextupole causes a loss of about half the beam, mainly from vertical aperture violations.

.

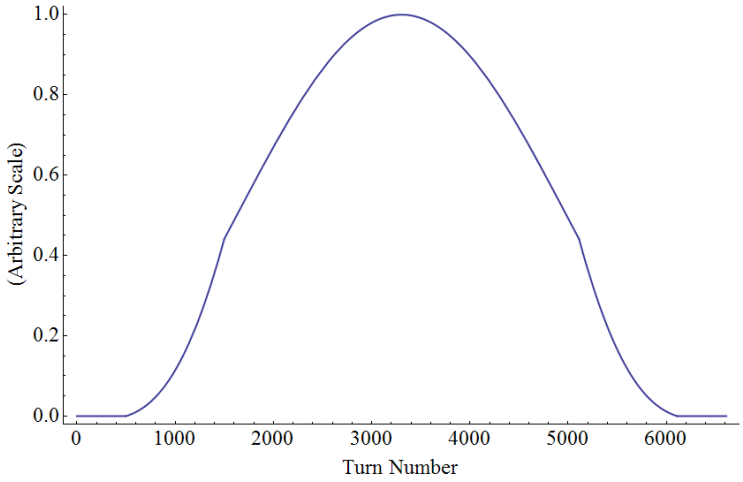

Figure 9: Possible sweep pattern for the resonant extraction sextupole. The magnet is turned off at injection, and follows the other Synchrotron magnets between turn 1500 and turn 5110.

.

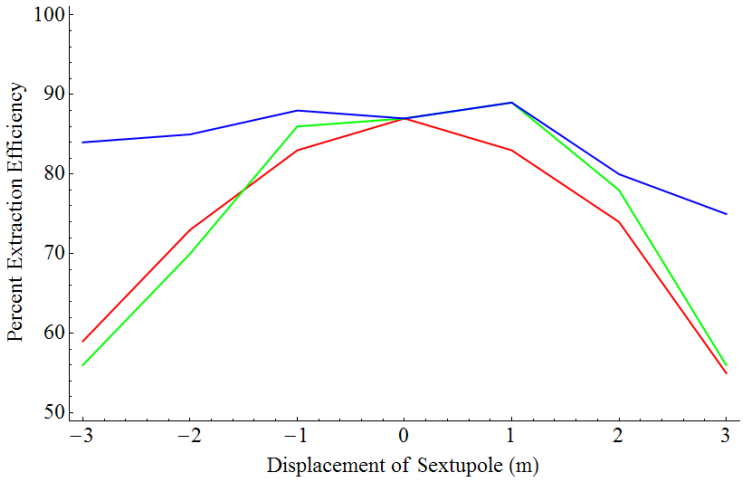

Figure 10: Variation of extraction efficiency with changes in the location of the resonant extraction sextupole. Zero displacement corresponds to the center of L4. The red curve shows the extraction efficiency without adjustment of any additional extraction parameters. The green curve includes adjustment of the quadrupole sweep, and the blue curve includes adjustment of both the quadrupole sweep and transverse septum locations. The quadrupole sweep was tuned primarily to correct the time profile (not shown) but also had some positive influence on the extraction efficiency for small adjustments in the sextupole location. The data shown here should be viewed primarily as qualitative, since there was a statistical fluctuation in extraction efficiency between runs on the order of several percent. In addition, all adjustments of the quadrupole sweep and septum locations were done by hand, so the given data may not represent true optimum values.

.

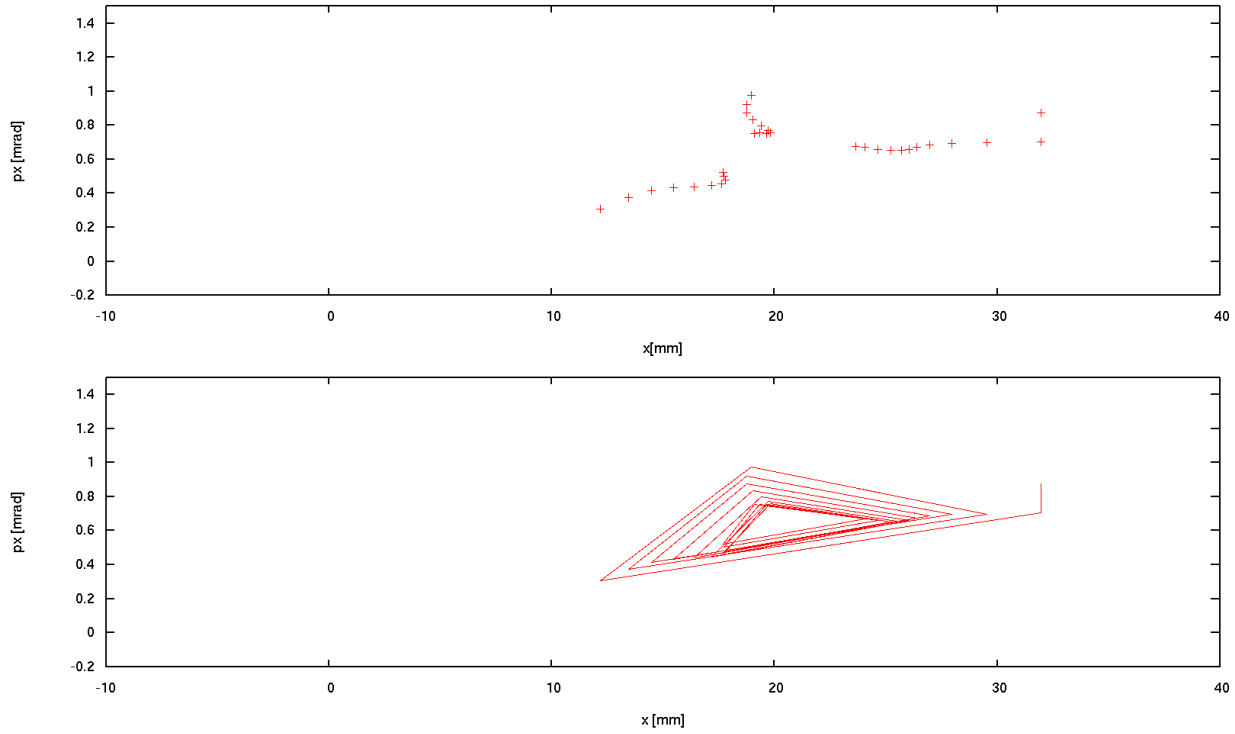

Figure 11: Turn-by-turn path of a single particle for the 30 turns prior to extraction. The particle's horizontal transverse location is taken at the electrostatic septum. The two images show the same data, but the bottom image connects points taken on subsequent turns. On the extraction turn, points immediately before and immediately after the electrostatic septum are shown to illustrate the magnitude of the septum kick.

.

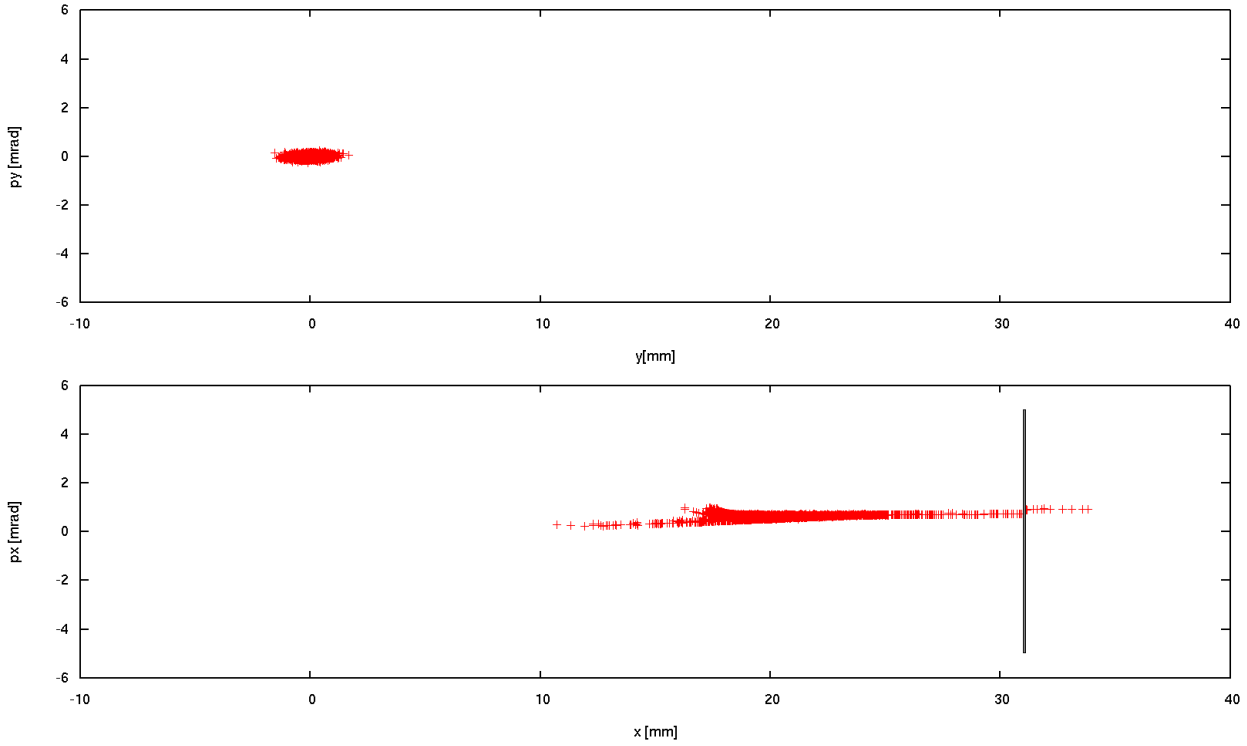

Figure 12: Transverse phase space profile of the beam during extraction, taken at the electrostatic septum. The first image shows the vertical phase space, while the second shows the horizontal phase space. The gray bar in the second image shows the width and location of the electrostatic septum – all particles outside it receive a kick of 0.17 mrad, which is visible here as a small vertical displacement. Note that no particles hit the electrostatic septum on this turn, although several were fairly close.

.

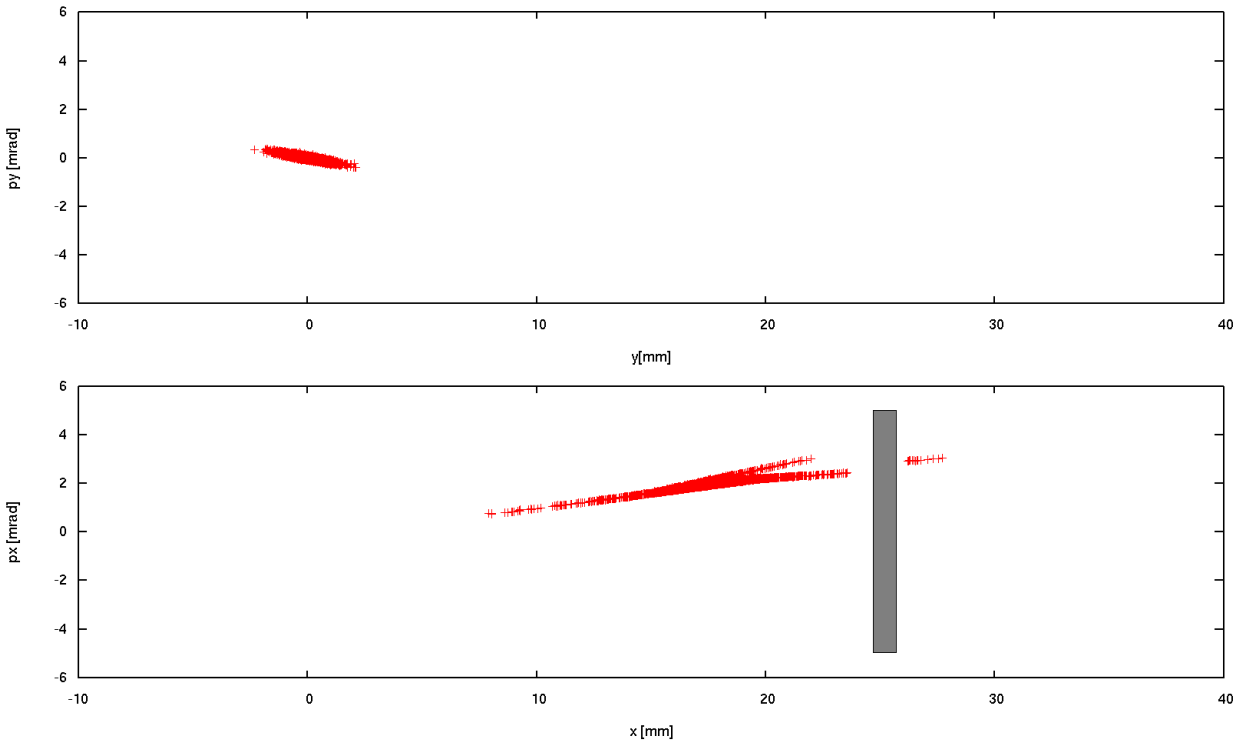

Figure 13: Transverse phase space profile of the beam during extraction, taken at the magnetic septum. The first image shows the vertical phase space, while the second shows the horizontal phase space. The gray bar in the second image shows the width and location of the magnetic septum – all particles outside it are removed from the circulating beam and transferred to the outgoing beam line. Note that this figure depicts the same turn in the same simulated extraction cycle as Figure 12, so the particles which cleared the electrostatic septum in that image are the same as those outside the magnetic septum in this image.

.

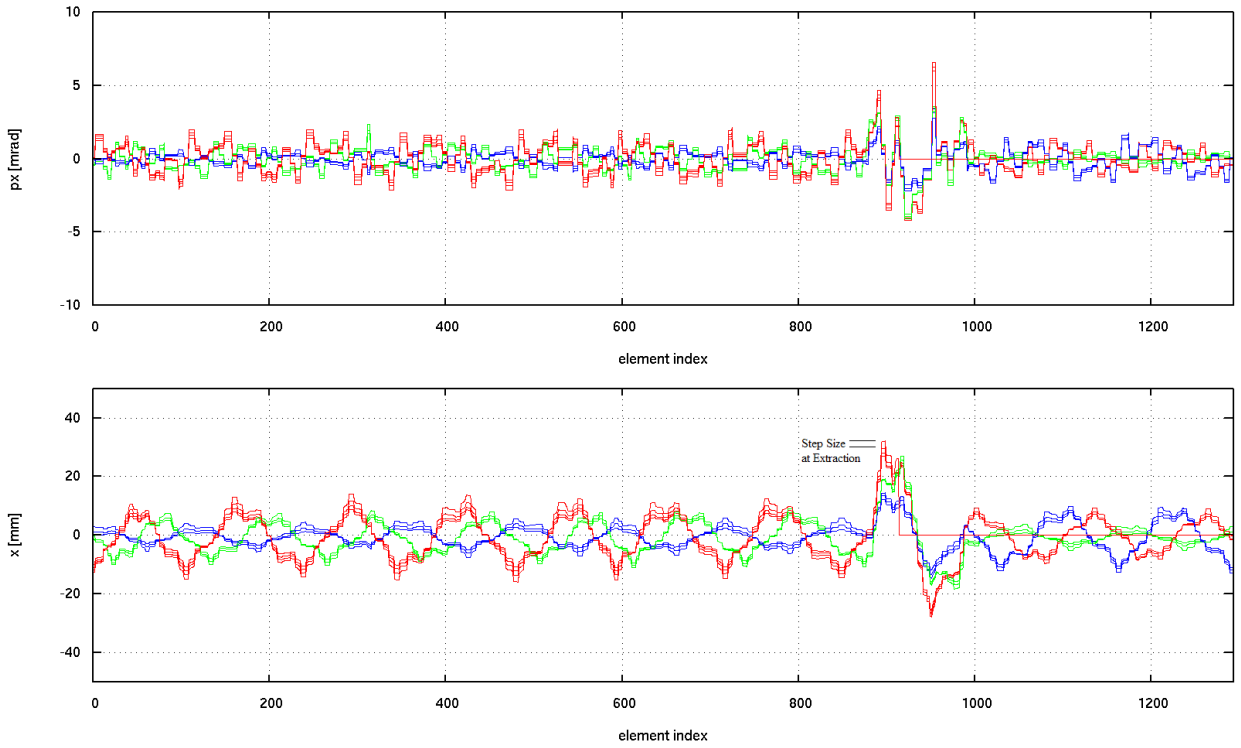

Figure 14: Horizontal transverse trajectory of a single particle, tracked over the 10 turns prior to extraction. Since the tune is very close to a third integer, the particle's path on a given turn tends to be similar to its path three turns prior. Therefore, turns of a different number modulo 3 are shown in different colors above. For example, turns 1, 4, 7, and 10 are shown in red. Note that the particle does not exactly retrace its steps every three turns. Instead, its horizontal displacement is slightly larger than that which it had three turns prior. This effect is the essence of the extraction procedure, and can be thought of as particles 'streaming out' along well defined trajectories starting at the corners of stable phase space. The step size is defined as the change in horizontal displacement every three turns, and the step size at extraction near the electrostatic septum is labeled above.

.

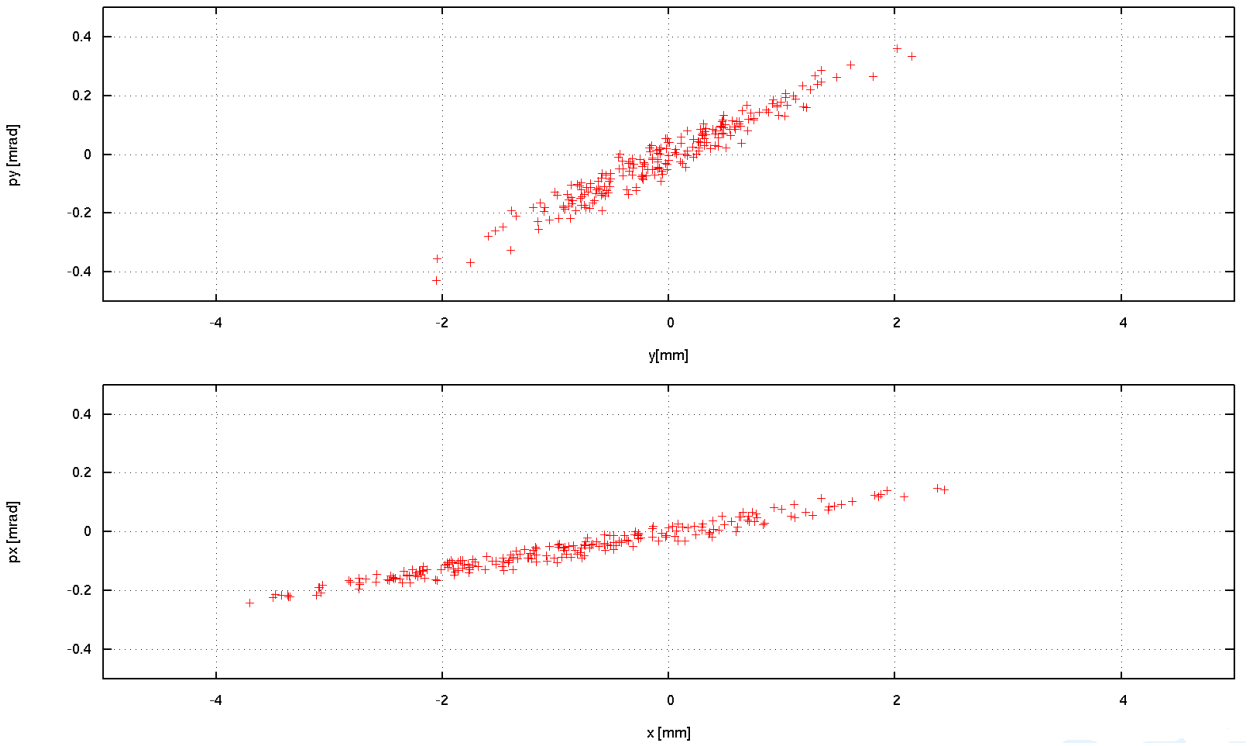

Figure 15: Time-integrated outgoing beam, taken 10 meters from the extraction point, without any time-varying dipole kicks or quadrupole focusing. On any specific turn, the horizontal emittance of the outgoing beam is about three times smaller than that shown above. However, due to the position and momentum drift of the outgoing beam, the time-integrated emittance is relatively large.

.

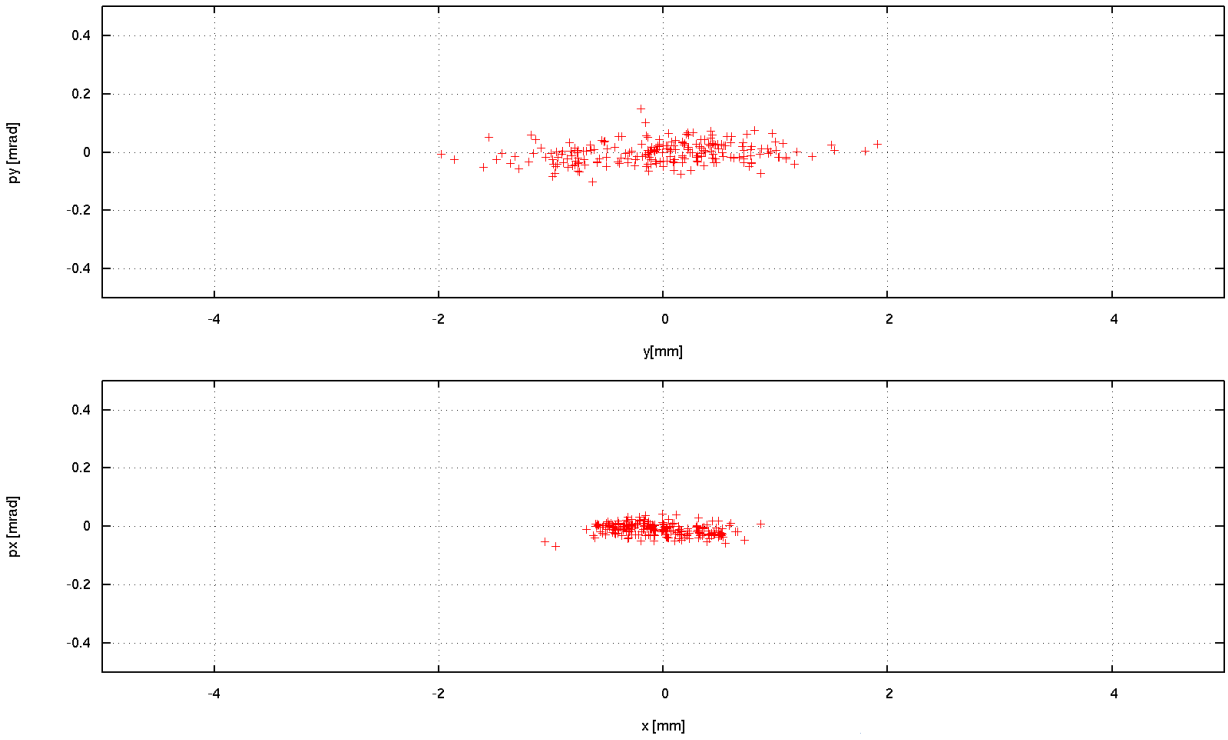

Figure 16: Time-integrated outgoing beam, taken 10 meters from the extraction point, with both time-varying dipole kicks and quadrupole focusing in the outgoing beam line. Note that the apparent horizontal emittance decreases by roughly a factor of three compared to that of Figure 15.

.