The Dissertation Committee for Nicholas John Evans
certifies that this is the approved version of the following dissertation:

Longitudinal Phase Space Tomography of Charged Particle Beams

Committee:

Sacha E. Kopp, Supervisor

Jack Ritchie

Michael Syphers

Roy Scwhitters

Chris Sneden

Longitudinal Phase Space Tomography of Charged Particle Beams

by

Nicholas John Evans, B.S.Phy.

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

August 2014

To my parents.

Acknowledgments

The last six years represent the hardest work I've ever done. I've learned a lot of physics, but the most important thing I've learned is that you can do much more with the support of people who believe in you than you can on your own.

Professionally this project would not have been possible without the help of Phil Adamson, who first propsed it. He has been an invaluable resource throughout, helping me navigate the technical ins and outs of Fermilab's accelerator systems and never hesitating to give a thorough answer. Duncan Scott was also instrumental in seeing this project through to the end. His labor, advice, and friendship helped immensely on a personal, and professional level. And I couldn't have done this project without working machines. I am indebted to the hard work of every physicist, operator, and technician who makes sure that protons make it safely to the Main Injector and beyond.

This degree has asked a lot of me intellectually of course, but personally as well. I never could have made it through the last six long, cold, tough years without the support of my friends and family, who often had no idea what I was working on but assured me I could do it anyway. I missed a lot of important things because I was working, but no one held it against me.

Thanks to: Meghan McAteer, Nathan Sellers, Benton Pahlka, and Re-

becca Bryzyski Pahlka for helping me endure Illinois and everything that entails; Casey, Corley, and Emma for always making me feel like I never left Texas; and Michael-Paul Hernández for consistently telling me that I could do this and sounding like he meant it.

And lastly, I express my gratitude for the four people most important to this endeavor. My mother, Nancy Coronado, who always made sure I knew I could do anything. My father, Carlie Evans, who instilled in me a respect for hard work and the people who do it, whatever their work is. I couldn't be a physicist without the aspiration, confidence, and respect I owe to them. Eryn Whitworth, whose belief in me as a person, and a scientist was unflagging even when my own was not. I hope I make her feel just as capable as she makes me feel. My advisor and mentor, Sacha Kopp, who believed in me from the very beginning and gave me a chance to become the scientist I am today. Without him I wouldn't have started graduate school, and I certainly wouldn't have finished. I've come to value his opinions and advice more than almost anyone else's. I hope he's prepared to keep offering both, becuase I'll keep asking.

And thanks to everyone I didn't mention by name, there are so many of you. Thinking back every person seems necessary, and I can't imagine how I would have succeeded any other way.

Longitudinal Phase Space Tomography of Charged Particle Beams

Publication No. _____

Nicholas John Evans, Ph.D.
The University of Texas at Austin, 2014

Supervisor: Sacha E. Kopp

Charged particle accelerators often have strict requirements on the beam energy, and timing to calibrate, or control background processes. Longitudinal Phase Space Tomography is a technique developed in 1987 to visualize the time, and energy coordinates of a beam. With non-invasive detectors, the beam can be visualized at any point during operation of a synchrotron. With the progress of computing power over the last 27 years, it is now possible to compute tomographic reconstructions in real time accelerator operations for many bunches around the accelerator ring. This thesis describes a real-time, multi-bunch tomography system developed and implemented in Fermilab's Main Injector and Recycler Rings, and a study of bunch growth when crossing transition. Implications of these studies for high intensity operation of the Fermilab accelerators are presented.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

For the past 100 years or so increasingly higher energy particles have allowed physicists to probe the universe at ever finer scales. The de Broglie wavelength of a particle is inversely proportional to its momentum, thus the higher the energy of a particle the shorter the wavelength and the smaller the features that can be resolved when using that particle as a probe.

Interactions involving higher energies also have more energy to create massive particles. But high energy particles must come from somewhere. Cosmic rays can have energies above 10^{20} eV, but these are incredibly rare, the flux of particles above 10^{20} eV is less than 1 $\text{km}^{-2}\text{century}^{-1}$ [11]. Because of the rarity, and unpredictability of cosmic rays, they are not an ideal source of high energy particles for studying the properties of fundamental interactions with any precision.

Particle accelerators are machines used to create more intense beams of particles under the control of experimenters, but at much lower energies than cosmic rays. Fig. 1.1a shows the history of particle accelerator beam energy on a log scale. As old technologies reach the limits and level off, new technologies have reliably stepped in to maintain the steady increase in beam energy. The

highest energy accelerator is currently the Large Hadron Collider at CERN, which can produces colliding proton beams of 4 TeV per proton each with plans to increase the energy to 7 TeV per proton. But the peak number of particles available for interactions, measured by a quantity called luminosity, is 10^{34} cm^{-2}s^{-1}, or on the order of 10^{53} km^{-2}century^{-1} in units comparable to the highest energy cosmic rays flux. Also notice in fig. 1.1a that even as one technology levels off, new machines are still being built with that technology. Energy is not the only consideration. The same greater intensity, and control that makes accelerators more desireable than cosmic rays in some cases, can also make more established technologies more desireable.

To observe a very rare process it is also necessary to observe many events. Fig. 1.1b shows the luminosity of various colliders versus the center of mass energy. While the relationship here appears mostly flat, the LHC, not yet built when this plot was made, shows an increase in both luminosity and center-of-mass energy to increase the physics reach through more interactions.

Increasingly, researchers are pushing for higher power machines with more current at more modest energies (still several GeV). High power beams allow researchers to obtain very precise measurements (e.g. muon g-2 measurements), observe very rare processes (e.g. flavor violating muon to electron decays), or to provide intense beams of secondary particles which are used as probes (e.g. spallation neutron sources).

The challenges to pushing the limits of energy and intensity are different, but both must often be addressed. To understand these limits requires a

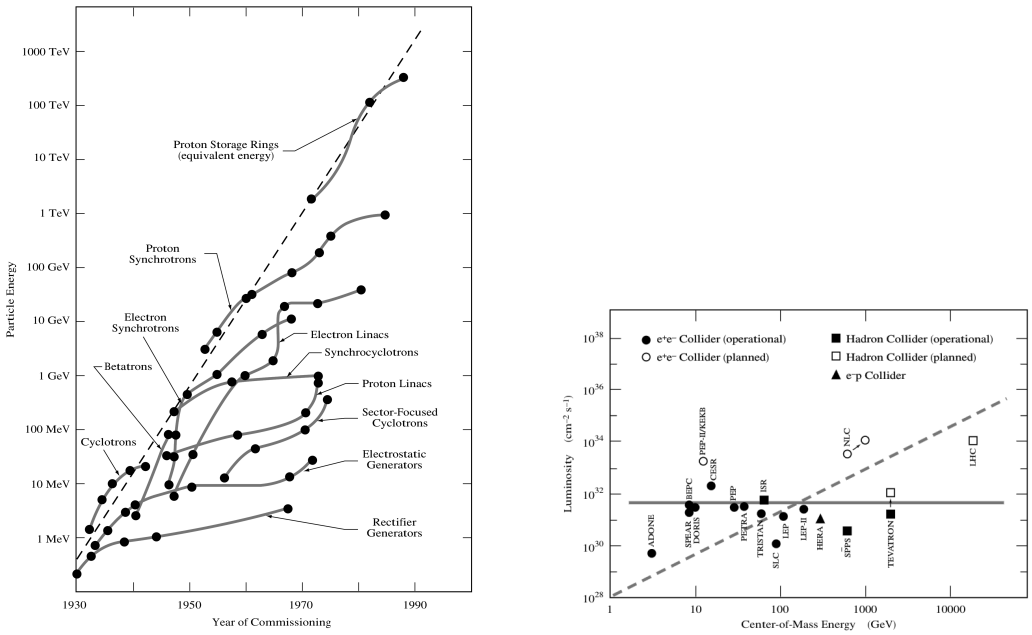

(a) Energy vs. Year (b) Luminosity vs. CM Energy

Figure 1.1: The evolution of particle accelerators. The Livingston plot in (a) shows the exponential increase in beam energy over time, including the technology used to achieve that energy. But even as an old technology reaches its limits in energy, new machines of the same kind are still built. (b) shows the luminosity, a measure of beam intensity, vs. center-of-mass energy in colliders. Luminosity doesn't increase at nearly the rate that energy does, but it is an important parameter for extending the physics reach of an accelerator. These plots were taken from [31]

basic understanding of the components of a particle accelerator.

1.1 Particle Accelerators

Particles accelerators utilize the Lorentz force to manipulate charged particles. Electric fields in the direction of motion impart energy to and accelerate particles, and magnetic fields steer particles onto desired orbits. The various ways of arranging these two functions leads to the vast array of modern particle accelerator types, from linear accelerators, to radio-frequency quadrupoles, and circular machines including: synchrotron, cyclotrons, microtrons, etc.

Many modern high energy accelerators, and all of the accelerators I will discuss in this document, are synchrotrons. Fig. 1.2 shows a simplified model of a synchrotron. A gap with an applied voltage, indicated as a dashed rectangle at the bottom of the image, provides an electric field in the direction indicated by the arrows on the three trajectories indicated by the black, red, and blue lines. The energy gained by a particle as it traverses this gap is equal to the charge of the particle, qe, times the voltage across the gap, V. Because this voltage is limited by the breakdown voltage of the materials used, the particle must traverse the gap many times to see an appreciable gain in energy. In this model only a single gap is used, but because of the magnetic fields perpendicular to the particle trajectories in the gray regions, the particles complete circular orbits and encounter the gap many times. One complete orbit, from cavity-to-cavity, is called a revolution.

For the field inside the gap to accelerate particles, it must oscillate in time as any closed loop integral in a static electric field results in zero increase in energy for a charged particle. The sinusoidal voltage will switch signs during half of the cycle, meaning particles that enter the gap at the wrong time will lose energy. This imposes a condition on the synchronization of the arrival time, or phase, of the particle at the gap providing the sinusoidal voltage. Hence the name synchrotron. Because the revolution period is normally in the radio-frequency range, the gap, which is physically realized as a resonant cavity, is called an RF cavity.

There will always be small errors from the ideal energy, and phase, with a distribution of values represented in a bunch of particles. How this manifests at the cavity is shown in fig. 1.3. Because the bending provided by the magnets in fig. 1.2 is dependent on momentum, an energy difference leads to a different revolution period. Different revolution periods mean that the phase of a particle with respect to the ideal particle changes over many turns. Properly phased, the RF cavity provides a restoring force to off momentum particles which leads to simple harmonic motion about the ideal energy and phase. This motion is called synchrotron oscillation, and the period of this motion is many machine revolutions.

For small errors relative to the synchronous phase and energy, the deviations in phase and energy are bounded and the motion is said to be stable. For larger errors the period of this motion increases, eventually stability is lost altogether and the motion is unbounded. The mathematical details of this

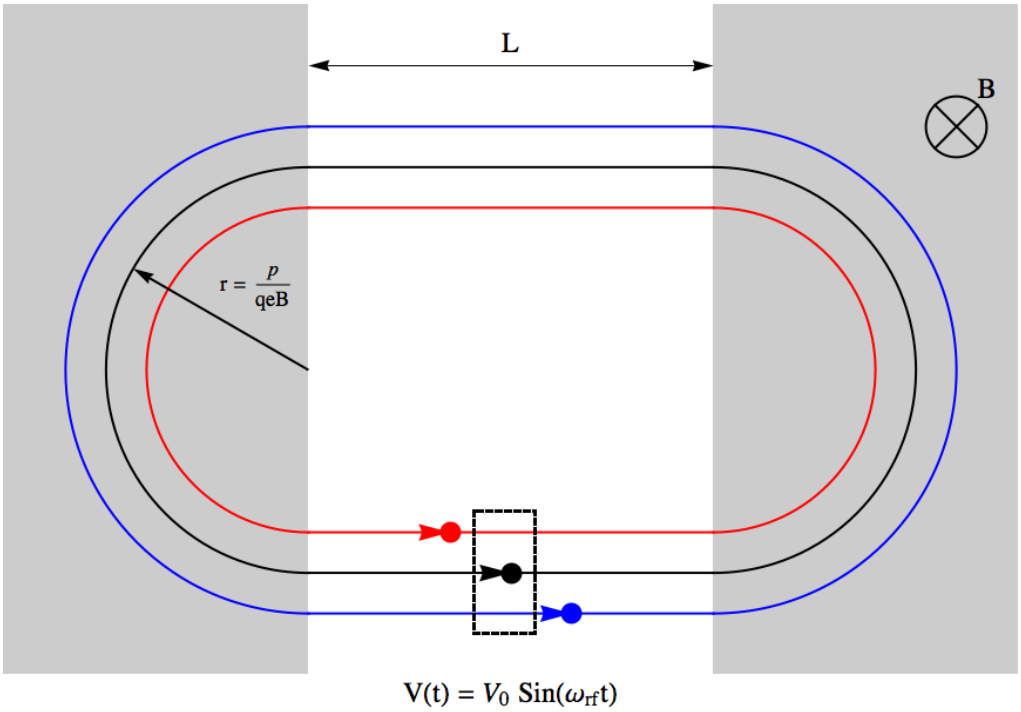

Figure 1.2: This image demonstrates the basic components of a synchrotron with respect to longitudinal dynamics. The grey blocks represent bending dipole magnets, with a magnetic field pointing into the page. Positive particles travel in the direction indicated by the arrows. The region between the dipoles is called a straight section. In the lower straight is a radio frequency cavity. The voltage that imparts energy to the three particles shown. The central trajectory is the design, or synchronous, particle. In a synchrotron the dipole field will increase as the synchronous particle gains energy so that the black path holds a fixed position. The blue particle has a lower momentum than the black particle, so the magnetic field produces a smaller radius. The red particle has a higher momentum, and a larger radius. The revolution period, τ of a particle is the time it takes to travel the complete path shown. A particle's path length, and velocity depend on its momentum relative to the sychronous particle.

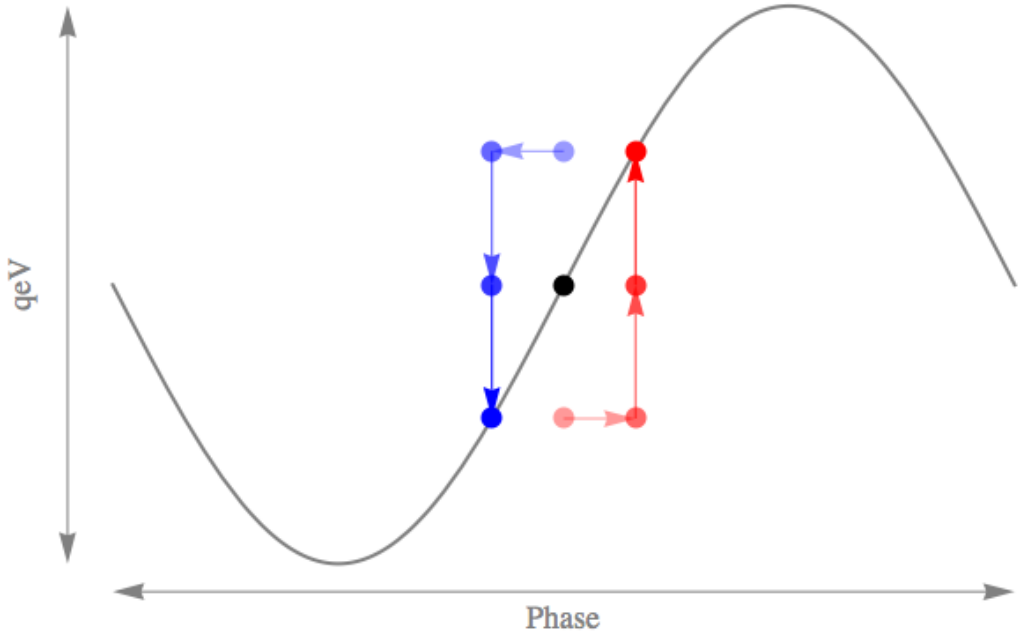

Figure 1.3: The vertical axis represents energy, and the horizontal axis is phase relative to the synchronous, or design particle, in black. Three particles from fig. 1.2 encounter the RF cavity on several revolutions. The cavity voltage scaled by the particle charge, qe, is indicated in gray. The particles first encounter the cavity in phase. Off-energy particles accrue phase on the next turn arriving out of phase relative to the synchronous particle, each now receives a voltage kick. Phased properly, this voltage decreases the energy of high energy particle(blue) that tends to arrive early, and increases the energy of low energy particle(red) that tends to arrive late. This motion about the synchronous point(black) is harmonic. This schematic exaggerates the step size for illustration.

process will be provided in ch. 2.

This stable motion, when viewed in the plane spanned by the time and energy coordinates which is known as longitudinal phase space, or just phase space, manifests as circular motion about the stable point, the coordinates of the synchronous particle.

Because the revolution frequency of the particle depends on its path length, the magnetic fields are implicated in the synchronicity condition when the design particle is accelerated, that is its energy changes from turn to turn. The bending radius of a charged particle in a magnetic field perpendicular to the velocity vector is given by eqn. 1.1.

$$r = \frac{p}{qB} \tag{1.1}$$

Where r is the bending radius, p the momentum of the charged particle, q its charge in units of the electron charge, and B the strength of the magnetic field in Tesla, and the radius points inward along the vector $q(\vec{v} \times \vec{B})$. This means that as momentum increases, either the field must increase, or the machine radius will increase. [1] It also means that particles near the synchronous energy

[1] At high enough energies there is an additional constraint imposed by radiative energy loss due to the radial acceleration produced by the bending magnets, called synchrotron radiation. The power radiated scales $\propto \gamma^4$, where γ is the Lorentz factor of the particle which is much higher for light particles of a given energy. Because of this scaling synchrotron radiation is more problematic in lepton accelerators (e^+, or e^-). When this energy loss over a single revolution exceeds the energy that can be supplied to the particle via the accelerating gap, the only solution is to make the machine larger, or to get rid of radial acceleration altogether with a linear accelerator.

will have slightly different path lengths.

The relationship between energy and revolution time exhibits a subtlety shown in fig. 1.4. The image demonstrates how velocity, and momentum, though related, must be considered separately when calculating the time it takes to traverse a length of the machine that includes bending magnets. At low energies particles with different momenta have different velocities. This difference in velocity dominates the time required to traverse a bending arc. At high energies the velocity of all particles approaches the speed of light. When this is the case the time to traverse a bending arc is dominated by the path length, which is proportional to the momentum alone.

The small differences in time to traverse a particular element are not important, only the total difference in revolution period of the particles is important, because this affects the phase of particles when they encounter the RF cavity.

There is a critical energy where the relationship between revolution period and energy relative to the synchronous particle changes, known as transition. At transition the frequency of all the particles is the same, and the oscillatory synchrotron motion facilitated by the RF cavity stops. This leads to problems with beam stability. The problems associated with transition will be discussed in detail in ch. 2.

At very high particle densities in the accelerator particles experience coulomb interactions that push the bunch apart. This is referred to as the

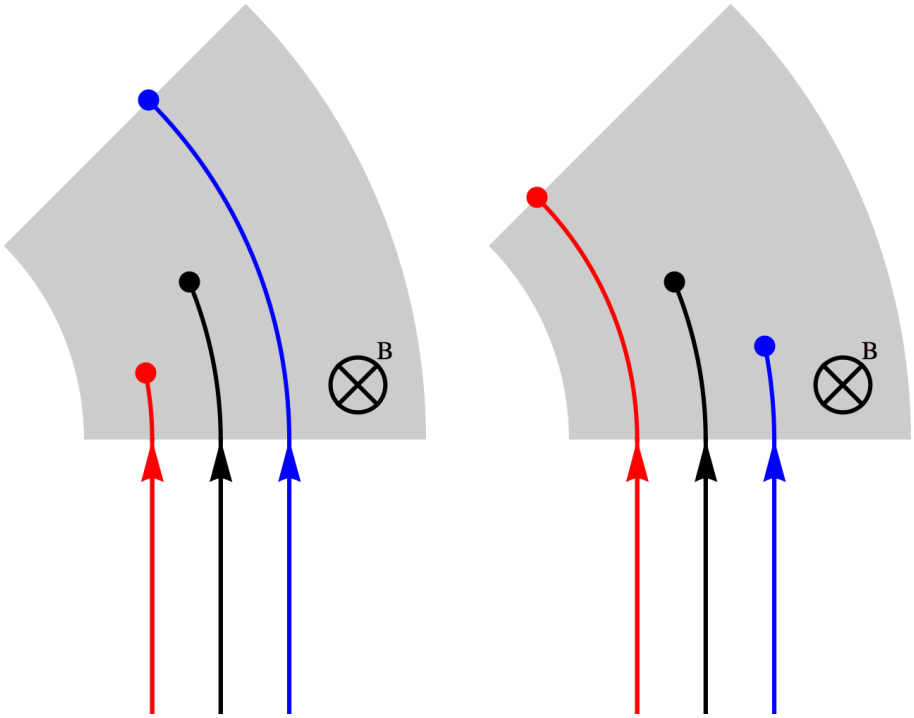

(a) $v \ll c$ (b) $v \approx c$

Figure 1.4: A magnetic field bends particles onto circular orbits. Panel (a) shows a snapshot of three positively charged particles with with velocities much less than the speed of light, and different momenta at the moment that the design particle has traversed half the bending magnet. The red, low energy particle has a smaller bending radius, but is also traveling more slowly so it lags behind the design particle. The opposite is true of the blue, high energy particle. Panel (b) shows the same scenario, but now the particles are very relativistic. Because the velocity of the particles is approximately the same, only the difference in path length contributes to the difference in transit time. Now the high energy particle takes longer to traverse the magnet. For a path parameterized according to azimuthal position around the ring, the high energy particle seems to move more slowly. The energy where this relationship changes is called the transition energy.

space-charge effect, or more generally a collective effect. Due to Lorentz contraction the electric field lines form a delta function in time pointing perpendicular to the direction of motion and the space-charge force does not effect the longitudinal distribution in the ultra-relativistic limit.

The term collective effects also includes the interactions of the coloumb fields of particles with conducting beam pipe, and other beamline elements. A variety of effects must be considered, the finite resistance, and the frequency dependent skin-depth of the wall, geometric discontinuities, resonant structures that may be excited by the passing beam. A high intensity bunch of protons may liberate electrons, which then accelerate towards a passing bunch, and continue accelerating toward the beam pipe when the bunch passes. With enough energy the electron can strike the pipe liberating more electrons. This process can lead to an electron cloud that interacts with the primary bunch and perturbing the motion. These kinds of interactions are numerous and represent an active field of study, as each new fix allows higher intensities, and more problems arise.

Collective effects are proportional to the intensity of the beam, or the beam current, and perturb the beam motion derived under the assumption of non-interacting particles. It is important to understand the collective effects in a high intensity machine to control the beam motion and maintain stable orbits.

1.2 The Importance of the Longitudinal Distribution

Stability means that excursions from the synchronous phase and energy are bounded. Collective effects are problematic because they perturb the potential provided RF cavity, eventually threatening stability. Particles which are unstable in phase simply lap the synchronous particle, much like a pendulum accruing phase by rotating about its axis. This is problematic for experiments that require particular beam structure. For instance, the planned mu2e experiment requires pulses of beam separated by long windows without beam during which background processes decay. Beam that strays into the clean window reduces the sensitivity of the experiment. Practically, this can be problematic because particles wind up in the portions of the ring normally left empty in order for extraction magnets to ramp up without spraying beam throughout the tunnel and causing activation of beamline elements, or radiation damage to sensitive components.

Particles that deviate in energy will also deviate in the transverse dimension when passing through magnets, as the two off-energy particles in fig. 1.2 indicate. If this deviation is large enough the particle will strike some physical restriction, an aperture restriction, either a magnet, or the beam pipe, or sometimes an intentional restriction called a collimator. Unstable motion in the energy coordinate also leads to beam loss on accelerator elements.

Unbounded deviations in energy are not acceptable in a machine designed to accelerate particles to a particular energy. Experiments often require small, known energy spreads to calibrate, or avoid backgrounds. Losses reduce

the amount of beam that reaches an experiment, and activate nearby equipment that can be safety hazards for technicians who have to repair machines, and damage equipment, or heating up superconducting elements that can lead to more down time, and further reducing beam to experiments.

1.3 Acceleration Through Multiple Rings

Fig. 1.5 shows the accelerator complex at Fermilab, and table 1.1 gives the size and energy range of each machine. Useful beam is extracted from many points in the complex depending on the particular operational mode.

Reaching the final beam energy in the multi-GeV range requires several machines each with a specific purpose. Moving beam into or out of a machine is called injection or extraction respectively, each operation has an associated injection, or extraction energy.

At low energies when the beam is not very relativistic collective effects are particularly problematic, and the velocity is changing rapidly. At higher energies, the velocity of the particle changes little, asymptotically approaching the speed of light. The momentum is changing however, which makes altering particle trajectories with a magnet more difficult, so higher energy machines are larger. At higher energies particles lost through collisions with the beam pipe, or other beam line elements deposit more energy. In addition, the saturation of magnet steel places a practical limitation, about a factor of 20, on the ratio of extraction energy to injection energy of a single machine [28].

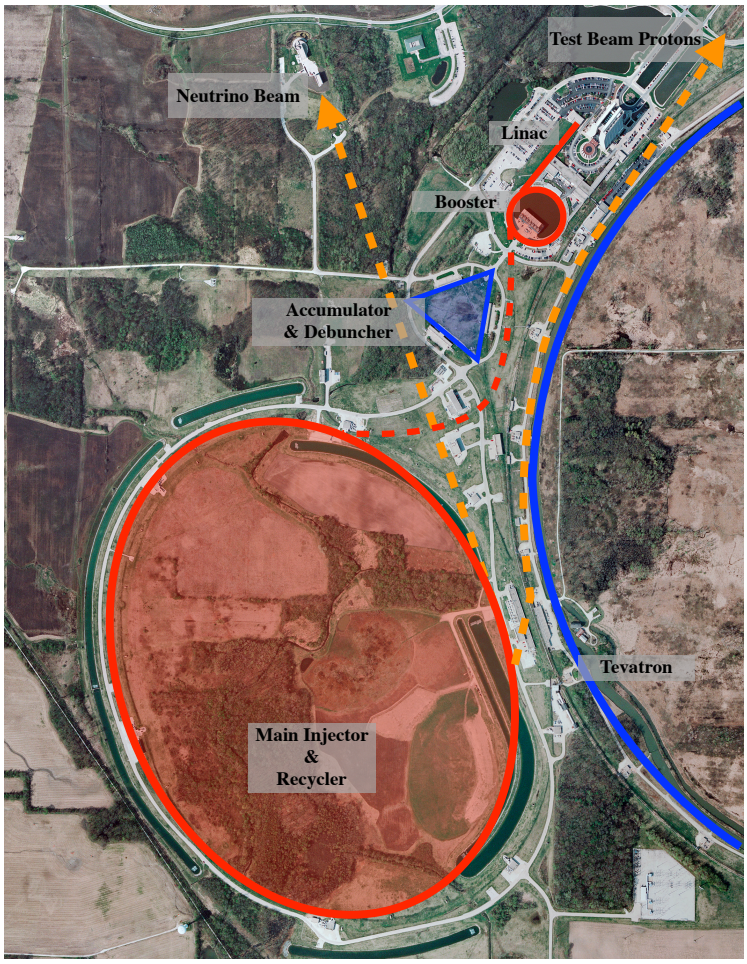

Figure 1.5: The machines in the Fermilab accelerator complex listed in tab 1.1 are shown. Machines currently used are outlined in red, machines that are not currently used are in blue. The red dashed lines indicate transfer lines between machines, and the orange dashed lines transfer lines to experiment halls. Beam starts at the source at the beginning of the Linac, proceeds to the Booster, then to the Recycler where it is slip-stacked to double intensity, and finally to the Main Injector which provides beam to neutrino experiments and the test beam area. The Accumulator and Debuncher rings are currently unused, but may be repurposed for 8GeV proton beam to supply muon experiments in the future.

Machine	Kinetic Energy	Circum[m]
RFQ	35-750keV	-
Linac	750kV-400MeV	-
Booster	400MeV-8GeV	474.2
Recycler	8GeV	3319.4
Main Injector	8-120GeV	3319.4
Accumulator	8GeV	474.0
Debuncher	8GeV	505.3

Table 1.1: Energy range and size of machines in the Fermilab complex.

Starting at the source, hydrogen gas is fed into a device called a magnetron that produces H^- ions. These ions are sent to the radio frequency quadrupole, which accelerates the beam of H^- ions from 35 keV to 750 keV kinetic energy. From the RFQ beam is passed to the linear accelerator(linac). The linac accelerates H^- ions from 750 keV to 400 MeV. The linac supplies protons to the neutron therapy center, a test facility, or the next accelerator, the Booster. This transfer from one machine to another is called injection. Upon injection from the linac to the Booster the H^- ions are passed through a thin foil, called a stripping foil, to remove both electrons, producing a beam of positively charged protons.

The remaining machines in the accelerator chain are synchrotrons like the simple model discussed above. The strength of the Booster's electromagnets ramp as part of a resonant circuit, meaning beam can only be accelerated at a rate equal to the natural frequency of the circuit(15 Hz). The duty factor of the Booster is limited to 7/15 cycles per second because of protons lost on beamline elements causing activation. After accelerating beam from 400

MeV to 8 GeV, the Booster can extract protons to an experiment hall, to the Recycler, or to the Main Injector.

The Recycler is a permanent magnet synchrotron. Because the magnet strength can't change it only operates at a beam energy of 8 GeV. The Recycler is mainly used to perform manipulations of the beam to increase intensity before sending it off to the Main Injector. In the future the Recycler may also provide beam to several muon experiments at Fermilab.

With a final energy of 120GeV/proton, the Main Injector is currently the highest energy synchrotron at Fermilab, it can accept beam either directly from the Booster, or Recycler at 8 GeV. The Main Injector magnets are not part of a resonant circuit, meaning the strength of the magnets can be programmed throughout the cycle. Because the magnet strength is variable, the Main Injector, like the Booster, can accelerate beam. Main Injector protons are either sent to the neutrino production beam line where they strike a target to produce a beam of secondary particles that decay into neutrinos used for several experiments, or extracted to other experimental areas. [2]

1.4 Longitudinal Diagnostics

One of the main tools used at Fermilab to non-destructively access information about the longitudinal distribution of particles in the beam is

[2]The CERN complex is similar but starts with a beam of protons, which are accelerated through a linac, the Proton Synchrotron Booster, the Proton Synchrotron, the Super Proton Synchrotron, and finally the Large Hadron Collider.

called a resistive wall current monitor(RWCM). A schematic of an RWCM is shown in fig. 1.6. An RWCM is an insert that replaces a small length of beam pipe with a non-conducting piece of ceramic. Resistors electrically connected to the beam pipe span the gap created by the ceramic. As a bunch of charged particles travels down the beam pipe the field generated by the bunch terminates at the conducting wall. The induced image charge follows the bunch creating an image current. When this current encounters the gap it is shunted through the resistors. By measuring the voltage across the resistors as a function of time, the total charge in the beam is deduced.

The resistive wall current monitor provides an impoverished view of the phase space. The theory of longitudinal phase space is developed with many references to the 2-D distribution in phase space, but the diagnostics usually available only provide a 1-D view of the phase space. A lot of information can be inferred from these 1-D projections if many projections are taken over time. Often this involves applying specific analysis techniques that require their own interpretation. A method known as longitudinal phase space tomography was first described by G. Jackson at Fermi National Accelerator Laboratory in 1987 [23], and further refined by S.Hancock et. al at CERN in 1998 [20].

Tomographic reconstruction is a non-destructive technique that allows the recreation of the 2-D distribution of particles energy, and phase relative to the synchronous particle at a particular machine revolution from measurements of the 1-D projections taken at over many subsequent revolutions. The reconstruction process converts a series of traces of the beam that contain use-

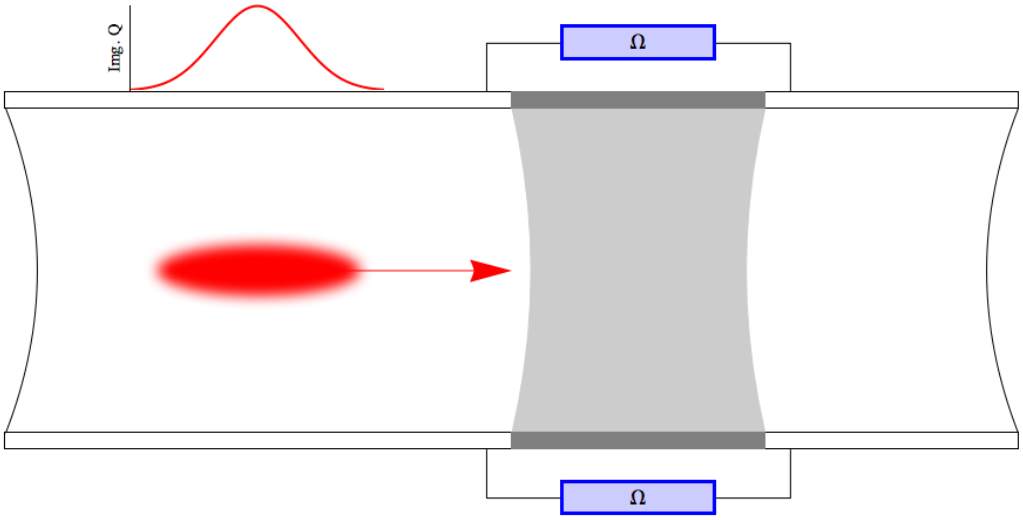

Figure 1.6: This image shows a cutaway of a beam pipe with a resistive wall current monitor installed. The particle bunch is shown in red propagating to the right. The gray band in the pipe wall represents the non-conducting ceramic insert. Two resistors of known resistance Ω are shown spanning the ceramic insert. The image current induced on the conducting pipe wall by the passing bunch is shown on top of the pipe. This current will be shunted through the resistors, a measurement of the voltage across the resistors as a function of time provides a measurement of the linear charge density of the bunch.

ful information into an intuitive view that anyone trained in accelerator physics will recognize immediately. The method is computationally intensive and, to date, the system at Fermilab is the first real-time, multi-bunch reconstruction system. Future increases in computing power will only make tomography more appealing as a non-destructive view into phase space.

Tomographic reconstruction provides measurements of the details of the energy coordinate, normally limited to a measurement of the centroid alone, as well as the area of phase space occupied by the particles in a bunch, an important quantity known as emittance, with fewer assumptions than other methods. Providing operators the ability to tune the beam phase, and energy, and identify errors in either coordinate when injecting from one machine to another quickly and intuitively has already proven useful during the commissioning of Fermilab's Main Injector, and Recycler in the Fall of 2013 after a year-long shutdown.

1.5 Organization of This Document

Chapter 2 describes longitudinal particle dynamics in a synchrotron. The stable, harmonic oscillation of particles about the design energy is derived, and the conditions required for stability are discussed. I will go into detail regarding the consequence of the phase slip indicated in fig. 1.4, a phenomenon known as transition which has implications for beam stability, and diffusion of the bunch of particles. Finally, I will mention some error conditions that arise when beam is transferred from one machine to another.

Chapter 3 deals with 2-D longitudinal phase space tomography, the technique that is the focus of this document. The algorithm used to reconstruct 2-D distributions of phase space from 1-D data is presented. I discuss the physical quantities that can be extracted from tomographic reconstructions, and some operational uses of the system implemented at Fermilab.

Chapter 4 is a demonstration of longitudinal tomography at Fermilab. In this chapter I describe the design of an experiment to measure the growth of particle distributions in phase space, described by a quantity known as emittance. Emittance is an important quantity measurements of which are either destructive, or require several approximations that limit the analysis that can be done. Chapter 5 presents the results of that experiment.

Chapter 6 will summarize the main points of the dissertation, and discuss the future applications of tomography. Useful extensions of the work presented here, and the challenges facing implementation will be discussed.

Chapter 2

Longitudinal Beam Dynamics

The motion of a particle beam in an accelerator can be divided into longitudinal, and transverse dynamics. Longitudinal beam dynamics describes particle motion in the direction of acceleration, and transverse dynamics describes motion along the two axes perpendicular to the direction of acceleration. Each of the three spatial axes can ordinarily be treated as uncoupled from the remaining two, and we will not consider any case otherwise.

The highest energy particle accelerators are synchrotrons, complex circular machines that use magnetic fields and radio frequency cavities synchronized to the revolution of charged particles to provide acceleration. The simplest model of a synchrotron, for the purpose of longitudinal dynamics, consists of dipole magnets, which provide the bending field necessary for a curved orbit, and a radio frequency cavity which imparts energy to the particles. This simple model is shown in fig. 1.2.

The black trajectory in fig. 1.2 is known as the design orbit, or design trajectory, followed by the design, or synchronous, particle. To first order, this path should be unchanged as the synchronous particle gains energy, the magnetic field ramping up so that the bending radius remains constant, and

the frequency of the RF cavity tuned to keep pace as the particle's velocity increases and the revolution period decreases. A red, and blue trajectory indicate particles with lower, and higher momenta respectively. Notice that the difference in momenta produces different bending radii. This orbit length difference combined with the different velocities of the off-momentum particles means the revolution times of off-momenta particles are different from that of the synchronous particle.

Fig. 1.3 shows how the high, and low momentum particles sometimes gain, and sometimes lose energy, depending on when they encounter the cavity, oscillating stably about the design values when the phase of the cavity is set properly relative to the synchronous particle. This energy oscillation also manifests as an oscillation in phase relative to the design particle, as a low energy particle takes more time to complete a revolution, and a high energy particle takes less time. For small deviations in energy this oscillation is simple harmonic motion, with a period much longer than a single revolution around the machine.

The phase-energy plane depicted in the simplified fig. 1.3 is used to describe longitudinal motion, and is called the longitudinal phase space. Fig. 2.1 shows the realistic trajectories of particles in this phase space. Notice that the axes are slightly different, where energy is actually given as ΔE, the deviation of the energy from the synchronous value. This is standard convention, and only done for convenience in calculation. [1]

[1]For those familiar with the phase space of a pendulum in classical mechanics, this image

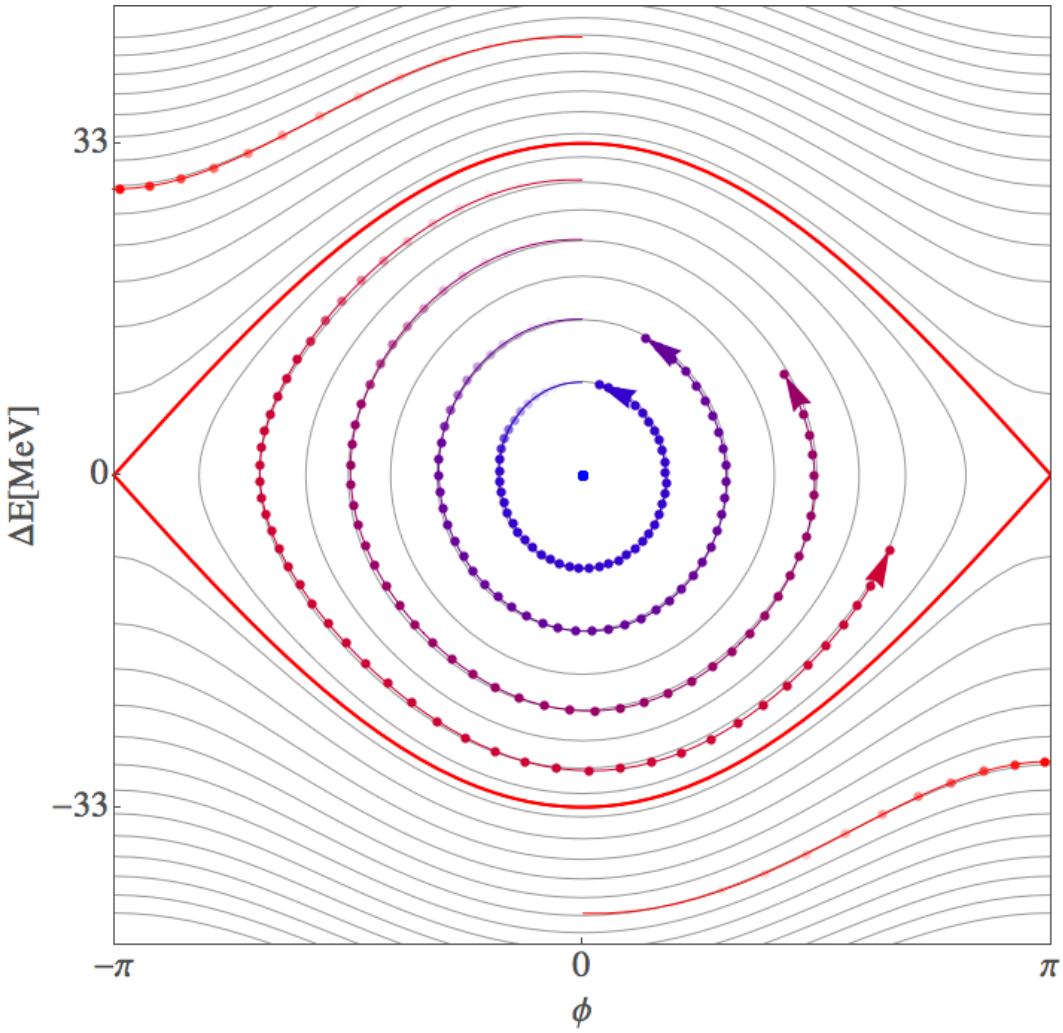

Figure 2.1: The central blue dot represents the energy and phase of the synchronous particle. The vertical axis is the energy deviation from the energy of the synchronous particle, ΔE. The horizontal axis represents the phase relative to the sine wave in the RF cavity depicted in fig. 1.2, ϕ. The time evolution of particles with energies and phases different from the synchronous value are shown as the paths surrounding the synchronous particle. Each point on a given curve represents the same particle on subsequent machine revolutions. A special trajectory, called the separatrix, that separates the regions of stable and unstable motion is indicated in red. The ϕ-ΔE plane is called the longitudinal phase space.

This chapter is dedicated to deriving the equations of motion governing the energy and phase of particles in a synchrotron, and exploring the features of such motion. One point in particular will occupy much of the discussion. A special moment in the cycle called transition. When the synchronous particle's energy is equal to the transition energy the increased velocity of a higher energy particle will be cancelled perfectly by the increased path length it experiences, similarly for slower particles. The revolution period becomes independent of momentum and stability is threatened. If the machine manages to make it past this energy, the relationship between relative momentum and revolution frequency will be reversed. Slower particles will seem to complete a revolution more quickly than faster particles. This change in dynamics can lead to unwanted diffusion of the beam.

2.1 Longitudinal Motion of Particles

Consider a charged particle of charge q, and energy E. Its momentum is $p/c = \beta E$, where $\beta = v/c$ is the velocity in units of the speed of light. In order to create a circular orbit a magnetic field of strength B is applied, as in fig. 1.2. When it encounters the magnetic field the particle travels in an arc with a radius, $r = p/qB$. If enough magnets are provided the particle completes a

should look familiar. The axes in the case of a pendulum would be angle, and angular velocity, or potential and kinetic energy. The trajectories then represent lines of constant energy. The limit of stable motion, indicated in red, would demarcate the energy above which the pendulum rotates about its axis in one direction. The sharp edges of this region at $\Delta E = 0$ would represent the unstable fixed point, where the pendulum is precariously balanced pointing upward.

full revolution and returns to its starting point. [2] The time it takes a particle to complete one revolution is called the revolution period, τ. In this simple model the full path length for a single revolution is $C = 2\pi r + 2L = 2(\pi p/qB + L)$, where L is the length of the straight section. Because of the bending magnets the path length is momentum dependent.

Now that the particle is traveling on a circular orbit, it will encounter the gap indicated at the bottom of the fig. 1.2 many times, once every τ seconds. Each time it passes the cavity is a chance to impart more energy to the particle. The gap is actually a resonant radio frequency(RF) cavity driven at the fundamental mode that has the main component of its electric field aligned parallel to the motion of the particles. The time dependent electric field is given by $E(t) = E\ \sin(\omega_{rf}t)$, where ω_{rf} is the resonant frequency of the cavity. This electric field provides a time-dependent accelerating kick to particles passing through the cavity. For the purposes of this analysis the RF cavity is modeled as a thin gap of length g, meaning E is static during the particle's transit of the cavity, and the voltage across the gap is $V(t) = E(t) \times g$. A particle of charge q that encounters the cavity at time t gains an energy of $qV(t)$ in units of electron-Volts, eV. (In reality there may be several cavities located around the circumference of the machine, but the analysis does not change significanly.)

[2] In this model there are two regions of magnetic field, called dipoles, that create the bending field for the circular orbit. The extent of the dipoles in this example is sufficient for a complete orbit. A real synchrotron may have many magnets each providing a much smaller portion of the arc.

It is helpful to give the arrival time of a particle in terms of phase relative to the RF oscillation, $\phi = \omega_{rf} t$. Any particle that arrives with a phase $\phi = 0, \pi, 2\pi...$ will receive no energy from the cavity. In order to provide acceleration, $\sin\phi$ must not be 0.

For acceleration the cavity must also provide coherent kicks to the particle from turn to turn. To do this the RF cavity must maintain a controlled phase relative to the arrival time of the particle every revolution period, τ. This relationship can be written:

$$h = \frac{\tau}{T_{rf}} = \frac{\omega_{rf}}{2\pi}\ \tau \tag{2.1}$$

Where h is a positive integer called the harmonic number and represents the number of RF periods, T_{rf}, per particle revolution, τ. This condition states that the particle's revolution period and the period of the cavity must be synchronized so that the particle sees the same phase from turn to turn.

This synchronicity condition defines a *synchronous particle*, also called the design particle. This is a conceptual particle that arrives at the cavity once every h oscillations of the RF cavity at the appropriate time, and with the appropriate energy. It is the black particle on the central orbit in fig. 1.2. All values describing the synchronous particle are given the subscript s. The phase ϕ_s, or velocity, β_s for instance. If the cavity imparts energy to the particle, its velocity will increase and τ will change, if I wish to maintain the relationship stated here, the RF frequency must also change. The magnetic

field will have to change if the particle is to maintain a constant orbit. The synchronous particle is an abstraction used to synchronize these changes to the systems of the accelerator.

Ultimately, the goal of acceleration is to produce more energetic particles. On each revolution the RF cavity can, in principle, provide an energy kick in the range $\pm qV$ depending on the arrival time of the synchronous particle. In order to keep proper accounting of the energy, we also need to know the arrival time of the particle, or it's phase since we're actually interested in the relative timing between the particle arrival and the field in the RF cavity. As a matter of convention, and convenience the synchronous particle always has the correct energy, E_s which we will use as the origin of our energy coordinate. The coordinate then is the deviation from this value, $\Delta E = E - E_s$. (For those keeping score, $(\phi, \Delta E)$ is not a canonically conjugate pair, $(\Delta t, \Delta E) = (\omega_{rf}(\phi - \phi_s), E)$ is however. I'll have more to say about this later.)

These two variables are all we need to describe the longitudinal motion of the beam. The $(\phi, \Delta E)$ plane, fig. 2.1, defines the *longitudinal phase space*, which I'll often call simply *phase space*. Phase accrues as the particles traverse their orbits, but the energy coordinate is only modified once per revolution, at the cavity. This means we can index the coordinates and keep track of discrete values at each turn, n.

In reality the accelerator handles more than one particle at a time, so we need to consider what happens to particles that deviate in phase or energy from the synchronous values. This is hinted at in fig. 1.2.

A particle with the synchronous energy (and charge) will experience the same bending radius as the synchronous particle, and have the same velocity, therefore it will traverse the ring in the same amount of time, and its phase is unaffected from turn to turn. If it enters the cavity with a different phase, however, it will receive a different kick than the synchronous particle and the energy will be different on the next revolution. Phase errors lead to energy errors.

A particle with the synchronous phase, but different energy will experience the same bending field but have different bending radius, and a different velocity, so in general it will complete a revolution in a different amount of time. As a result it will slip in phase over a turn and receive a different kick on subsequent turns. Energy errors lead to phase errors.

2.1.1 Longitudinal Stability

I've asserted that the relationship between phase and energy, stated here qualitatively, can lead to stable motion about the design energy, but I have yet to establish the conditions needed to obtain stability.

To say particle motion is stable means the particle with phase and energy initially near the synchronous particle should stay near the synchronous particle for all time. This is evident for the particles within the central region of fig. 2.1, bounded by thick red line called the *separatrix*. The two things that have not been quantified in the simple argument at the end of the last section are the amount of phase accrued over a single turn by a particle deviating

from the synchronous energy, and the energy difference accrued over a single turn by a particle deviating from the synchronous phase.

The difference in energy accrued on a single revolution for a particle with a phase different from the synchronous value is easy to calculate: $\Delta E = qV\ \sin\phi - qV\ \sin\phi_s$. For a particle with an energy error the phase accrued relative to the synchronous particle is a little more involved. It is given by $\delta\phi = (2\pi h) \times (\delta\tau/\tau_s)$. What we need to quantify is the fractional difference in revolution period between a particle and the synchronous particle, $\delta\tau/\tau_s$. The revolution time of the synchronous particle is just the path length, C, divided by the velocity.

$$\tau_s = v/C \tag{2.2}$$

Obviously the velocity depends on momentum, but because of the bending field so does the path length. To first order in momentum deviation, the fractional deviation in τ is calculated as follows:

$$\frac{\delta\tau}{\tau_s} = \frac{\delta C}{C_s} - \frac{\delta v}{v_s} = \left(\frac{1}{\gamma_t^2} - \frac{1}{\gamma_s^2}\right)\frac{\delta p}{p_s} \tag{2.3}$$

Where γ_s is the relativistic Lorentz factor, and γ_t is the proportionality between the path length deviation and momentum deviation, which must be calculated for the machine under consideration. There may actually be higher order dependences on $\delta p/p$, but generally the momentum deviations are so small that they can be ignored. An important point here is that depending on

the value of γ_t, and the particle energy, represented by γ_s, the relationship between revolution period and momentum can change. This means faster particles will have longer revolution periods when $\gamma_s > \gamma_t$, appearing to move more slowly than lower energy particles. And when $\gamma_s = \gamma_t$ the revolution period is independent of momentum.

Generalizing, the term $(1/\gamma_t^2 - 1/\gamma_s^2)$ is often written as the single parameter η, called the slip factor. Rearranging eqn. 2.3 gives a definition for η.

$$\eta = \frac{d\tau/\tau_s}{\delta p/p_s} \tag{2.4}$$

What η really tells us is how much of a time difference accrues for a particle given some deviation from the synchronous energy. Which is what we set out to calculate. There is one more matter of convention to discuss.

Because the momentum deviation $\delta p/p$ is typically small, it is standard to expand η in power series about $\delta p/p$ as in eqn. 2.5

$$\eta = \eta_0 + \eta_1 \delta p/p_s + \eta_2 (\delta p/p_s)^2 + ... + \eta_n \mathcal{O}((\delta p/p_s)^n) \tag{2.5}$$

When written this way, it is obvious that $(1/\gamma_t^2 - 1/\gamma_s^2)$ is just the lowest order term in the expansion, η_0. Unless explicitly stated, $\eta = \eta_0$. Keep in mind that η depends on γ_s, and changes throughout the cycle, changing sign as γ_s crosses γ_t. This is known as crossing transition.

I mentioned that γ_t is a relates the change in path length and change in momentum, this proportionality may itself have higher order terms. In general, fractional change in path length for a given momentum devaition is given by the momentum compaction in eqn. 2.6.

$$\frac{C}{C_s} = 1 + \alpha_0 \delta p/p_s (1 + \alpha_1 \delta p/p_s + \mathcal{O}((\delta p/p_s)^2) \tag{2.6}$$

The lowest order term is related to γ_t, with $\gamma_t = 1/\sqrt{\alpha_0}$. The higher order terms of η are often written in terms these values,

$$\begin{aligned} \eta_0 &= \alpha_0 - \frac{1}{\gamma_s^2} \\ \eta_1 &= \alpha_0 \alpha_1 - \eta_0 \alpha_0 + \frac{3}{2}\frac{\beta_s^2}{\gamma_s^2} \end{aligned} \tag{2.7}$$

These values are important for determining emittance growth near transition, where higher order terms of η become important.

With η defined I can quantify the difference in phase accrued on the $(n+1)^{th}$ turn by particles with energies deviating from the synchronous value. Together with the energy difference accrued for particles with phase errors the difference equations 2.8 describe the motion of particles in the longitudinal phase space from turn to turn.

$$\begin{aligned} \Delta E_{n+1} &= \Delta E_n + qV(\sin\ \phi_n - \sin\ \phi_s) \\ \phi_{n+1} &= \phi_n + 2\pi h \eta \frac{\Delta E_{n+1}}{\beta^2 E_s} \end{aligned} \tag{2.8}$$

Where I've replaced $\delta p/p_s$ with $1/\beta_s^2 \times \Delta E/E_s$. A separate $(\phi, \Delta E)$ pair describes the position of each particle. As simple as these equations seem, they are very powerful. The discrete particles indicated in fig. 2.1 were tracked using these equations. In order to understand the system analytically, though, we will develop a differential approximation. If we assume the changes in ϕ and ΔE are small over a single turn we can write the change in energy and phase per turn as differentials.

$$\begin{aligned} \frac{d\Delta E}{dn} &= qV(\sin\phi - \sin\phi_s) \\ \frac{d\phi}{dn} &= \omega_{rf}\tau\eta\frac{\Delta E}{\beta^2 E_s} \end{aligned} \tag{2.9}$$

Taking the derivative of the equation for ϕ, we can obtain a second order equation in ϕ alone:

$$\frac{d^2\phi}{dn^2} - \frac{\omega_{rf}\tau\eta qV}{\beta^2 E_s}(\sin\phi - \sin\phi_s) = 0 \tag{2.10}$$

Extending the single turn difference equation in this way allows us to analyze two important phenomena over many turns: harmonic oscillator behavior of particles relative to synchronous particle, and contours in phase space defining stable particles which can orbit the accelerator. The following discussion of these two effects will address the question of stability.

2.1.2 Harmonic Oscillator Approximation

Near the synchronous particle the RF cavity sine wave can be linearized, and a linear restoring force implies simple harmonic motion. In analogy with a simple harmonic oscillator, we should also expect an unstable fixed point. Because we control the phase of the restoring force in the RF cavity, we must make sure that the accelerator produces stable motion.

Considering small deviations around the synchronous phase allows us to write a linear form of the differential equation in terms of $\delta\phi = \phi_s - \phi$. Under this approximation $\sin\phi - \sin\phi_s \approx \delta\phi \ \cos\phi$ and eqn. 2.10 becomes:

$$\frac{d^2\delta\phi}{dn^2} + (2\pi\nu_s)^2\delta\phi = 0 \tag{2.11}$$

The equation of a simple harmonic oscillator with a natural frequency $2\pi\nu_s$ in units of [rad/machine revolution]. ν_s is called the synchrotron tune and gives the number of phase space oscillations per machine revolution, and is always less than one. The tune is just a frequency, but because the equations of motion are given with respect to machine revolutions, instead of time, the natural frequency is also in terms of machine revolutions. The expression for ν_s is:

$$\nu_s = \sqrt{-\frac{\omega_{rf}\tau_s\eta qV \ cos \ \phi_s}{\beta_s^2 E_s}} \tag{2.12}$$

The synchrotron tune can also be expressed as a synchrotron frequency, Ω_s, in units of Hz,

$$\Omega_s = 2\pi\frac{\nu_s}{\tau_s} = 2\pi\sqrt{-\frac{\omega_{rf}\eta q V\ cos\ \phi_s}{\beta_s^2 E_s}\tau_s} \tag{2.13}$$

It is well known that the frequency of a harmonic oscillator must be real in order to have stable motion. In the expression for ν_s only η and cos ϕ_s can change sign, all the remaining parameters must be positive. Thus, the condition $-\eta\ \cos\phi_s\ >\ 0$ must be maintained for stable motion.

Let γ_t be a given parameter, so we have no control over η at a particular energy. The sign of η is determined by the value of γ_s, or where we are in the acceleration cycle. ϕ_s is completely arbitrary though, and the values that lead to stable motion depend on η. Below transition $\eta < 0$ and $-\pi/2 < \phi_s < \pi/2$, and above transition $\eta < 0$ and $\pi/2 < \phi_s < 3\pi/2$ satisfy the stability requirement.

Additionally, in the tracking eqn. 2.8 the sign of sin ϕ_s determines whether the synchronous particle is gaining or losing energy. If the phase is 0 or π, depending on the sign of η, the synchronous particle receives no kick and the beam is said to be *stationary*.

A harmonic oscillator has constant total energy exchanging potential for kinetic as the system oscillates. In this case, it is action that is conserved. In the next section I will derive curves of constant action in the phase space, and use these to examine the limits of stability.

2.1.3 Particles Outside the Harmonic Oscillator Region

Particles only exhibit simple harmonic motion when the deviation from the synchronous particle is small. This is indicated in fig. 2.1 by the particle immediately orbiting the synchronous particle.The remaining particles that are within the separatrix will still exhibit oscillatory motion, but the frequency becomes dependent on action. The non-linearity in the focusing potential provided by the RF cavity creates this action dependence. Beyond the separatrix, motion is no longer stable. Because this nonlinearity is a feature of the system, not an error or perturbation, we need to consider it if we want to understand the full stability region.

We can start from eqn. 2.10 and produce an integral of the motion. i.e. there is a conserved term implied by the equations. This will be the phase space action I mentioned in the last section.

$$\int \left(\frac{\mathrm{d}^2\phi}{\mathrm{d}n^2} - \frac{\omega_{rf}\tau\eta q V}{\beta^2 E_s}(\sin\phi - \sin\phi_s) \right) \frac{\mathrm{d}\phi}{\mathrm{d}n}\mathrm{d}n = \int \mathrm{d}n \tag{2.14}$$

$$\frac{1}{2}\left(\frac{\mathrm{d}\phi}{\mathrm{d}n}\right)^2 + \frac{\omega_{rf}\tau\eta q V}{\beta^2 E_s}(\cos\phi - \phi\sin\phi_s) = \mathrm{C} \tag{2.15}$$

From eqn. 2.9 we can write this in terms of only ΔE and ϕ

$$\frac{1}{2}\Delta E^2 + \frac{\beta^2 E_s q V}{2\pi h\eta}(\cos\phi + \phi\sin\phi_s) = \mathrm{C} \tag{2.16}$$

This isn't strictly an action, though we could derive a Hamiltonian and convert into action-angle coordinates, I will refer to this constant as the action. Analogous to the pendulum, over time the particle swaps kinetic for potential energies, while maintaining a constant total value follows curves of constant action in the phase space. A special trajectory called the separatrix marks the boundary between stable and unstable motion. Fig. 2.2 shows contours of constant action for several values of ϕ_s.

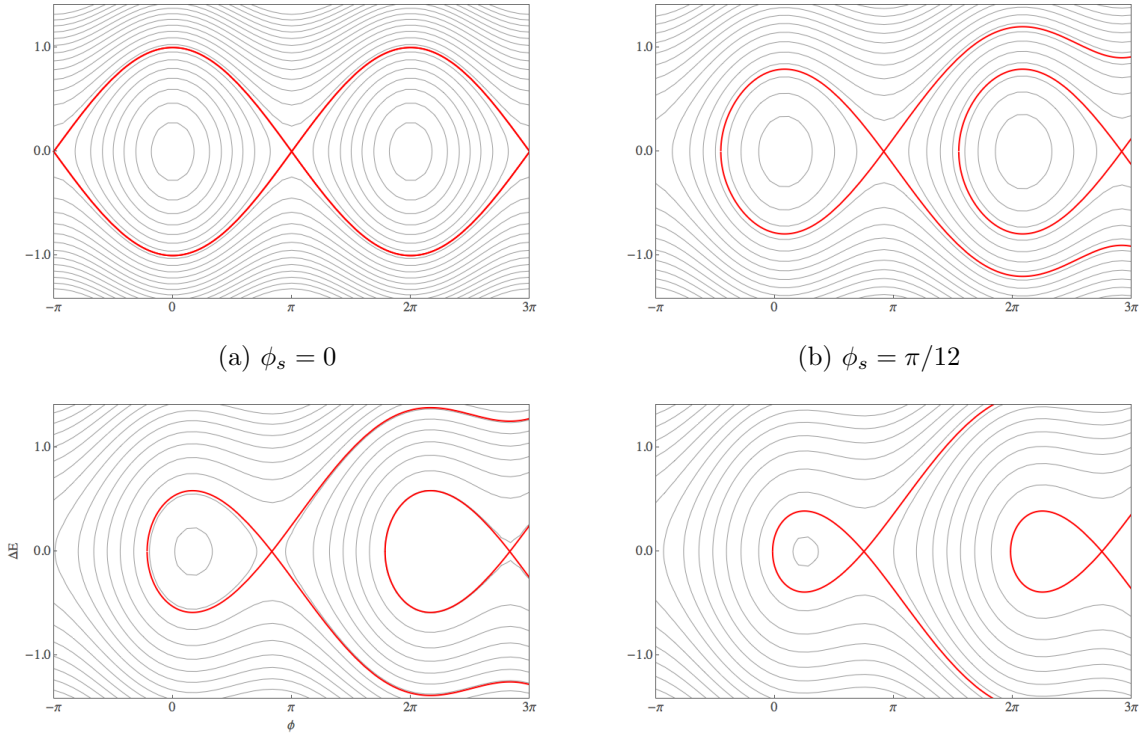

(a) $\phi_s = 0$ (b) $\phi_s = \pi/12$

(c) $\phi_s = \pi/6$ (d) $\phi_s = \pi/4$

Figure 2.2: Contours of constant action in the ΔE - ϕ longitudinal phase space as described by eqn. 2.16 for several values of ϕ_s. The regions enclosed by the red lines indicate buckets, the portion of phase space where motion is stable. A machine with harmonic number h will have h such stable buckets. The buckets shrink as ϕ_s deviates from the stationary bucket ($\phi_s = 0$). Particles outside of the bucket accrue phase indefinitely, but remain within a small range of energies when $\phi_s = 0$, and stream away from the synchronous energy otherwise.

In the case of a stationary bucket, $\phi_s = 0$, even the unstable region is only unstable in phase, returning to the same relative position every $2\pi h$ rad. This is analogous to a pendulum with enough energy to rotate fully around its axis. In the case of accelerating buckets, the unstable beam diverges in both energy and phase. In a real machine a large energy deviation will mean a large deviation in radial position, and eventually the particle will be lost on the beam pipe, or a magnet aperture.

The entire stable region represents an *RF bucket*. A bucket doesn't necessarily contain particles, it is a phase space region defined by the RF system, and machine parameters. When a bucket does contain particles the particles contained in the bucket are called a *bunch*.

The bucket height is an important number representing the maximum energy deviation from the synchronous value that can still undergo stable synchrotron motion. Because particles with energy deviations have different radii than the synchronous particle, and the beam pipe containing the particles is finite, an energy deviation that is too large can lead to beam being lost. Assuming that all of the beam is in the stable region, the bucket height must remain smaller than the aperture of the machine, which is the maximum energy deviation that can be physically accommodated.

Using eqn. 2.16 the phase space contour for the separatrix is determined by setting ϕ equal to one of the turning points, where $\Delta E = 0$. This gives the action on the separatrix. The maximum energy deviation, the bucket height, is obtained when $\phi = \phi_s$, minimizing sin ϕ_s. This is given in eqn. 2.17.

$$\Delta E_{MAX} = 2\sqrt{\frac{\beta^2 E_s q V}{2\pi\omega_{rf}^2 h|\eta|}}\left|\cos\ \phi_s - \frac{\pi - 2\phi_s}{2}\sin\phi_s\right| \tag{2.17}$$

Another important measure is the bucket area, A, the total area of phase space enclosed by the separatrix. When I discuss ensembles of particles I will discuss how to calculate the phase space area the beam occupies, which should be less than the area that the bucket occupies to ensure that all particles are executing stable motion in the longitudinal phase space. The bucket area can be found by integrating the area under the separatrix. The integral here can be solved exactly for the stationary bucket, $\phi_s = 0$ or π, or numerically otherwise.

$$A = 16\sqrt{\frac{\beta^2 E_s q s V}{2\pi\omega_0^2 h|\eta|}}\alpha(\phi_s) \tag{2.18}$$

Where the factor $\alpha(\phi_s)$ is the ratio of the area of the bucket with ϕ_s to the stationary bucket, eqn. 2.19.

$$\alpha(\phi_s) = \frac{1}{4\sqrt{2}}\int_{\phi_l}^{\pi-\phi_s}\sqrt{sgn(-\eta)[\cos\ (\pi-\phi_s) - \cos\phi + (\pi - \phi_s - \phi)\sin\phi_s]}\,\mathrm{d}\phi \tag{2.19}$$

and the limits of integration are the turning points, ϕ_l, and $\pi - \phi_s$. ϕ_l is the solution to the equation, $\cos\phi_l + \phi_l \sin\phi_s = -\cos\phi_s + (\pi - \phi_s)\sin\phi_s$, which usually has to be solved numerically.

Each panel of fig. 2.2 depicts the same rectangular phase space area, notice that as ϕ_s increases the stable area inclosed by the separatrix, given

by eqn. 2.18, decreases. Because of this the bucket area is important during the operation of a particle accelerator to ensure that all of the particles are contained within the stability region, especially when transferring particles between machines, or when accelerating a bunch of particles.

So far we've only discussed single particle dynamics, though I've alluded to the multi-particle bunches we will actually accelerate. The next section discusses the behavior of ensembles of particles within a bucket, and one measure used to characterize the ensemble.

2.1.4 Ensembles and Emittance

We don't accelerate particles one at a time, the ensemble of particles that occupy a particular bucket in the phase space is known as a bunch. The particles of a bunch are not necessarily confined to the linear region of the phase space very near the synchronous particle. As I showed in the last section, particles far from the synchronous particle in phase space, but still within the separatrix, have lower synchrotron frequencies. This change in frequency with increasing action causes the bunch to shear when many particles are present, but only if the bunch occupies a large enough spread in action.

A full treatment of the evolution of ensembles requires the Vlasov equation (which reduces to the collisionless Boltzmann equation when neglecting self forces) [9]. Tracking is much easier when only single particle dynamics are considered and for much of what we will cover, bunch self-forces, and beam-environment interactions, known as collective effects, are not necessary. But I

will discuss some collective effects later.

For a bunch of particles the emittance, ϵ, is often used as a measure of how diffuse the bunch is in phase space. Emittance is the average action of all the particles in a bunch, sometimes referred to as the phase space area occupied by the bunch. Fig. 2.3 shows how this second description can be misleading. The left panel shows a contrived bunch generated from a rectangular distribution. Because the actions spanned by this distribution in the given machine do not match the contours the distribution distorts over time filamenting to eventually fill the contour indicated in the right panel. The emittance is often written $\epsilon_{68\%}$ (RMS emittance), or $\epsilon_{95\%}$, which indicates the action such that 68%, or 95% of the particles in a bunch have an action less than that value. The emittance of a bunch should be less than the bucket area, eqn. 2.18, if all of the particles are to remain within the stable region defined by the separatrix.

The units of action, and therefore emittance, are energy $\times$ time, which seems to be at odds with the ϕ, and ΔE coordinates I've been using. Until now these coordinates have sufficed because they are intuitive, but technically they do not form a canonically conjugate pair. In order to properly measure emittance, ϕ should be replaced with $\delta t = \omega_{rf}(\phi - \phi_s)$ (Some authors may use $(\phi, W = \Delta E/\omega_{rf})$, $(\phi, \delta = \Delta p/p_0)$, which are equivalent).

Because this system is conservative, as long as changes to the system are adiabatic the action of any particle should be the same, and the emittance of a bunch should be preserved. Dissipative effects including noise in the RF

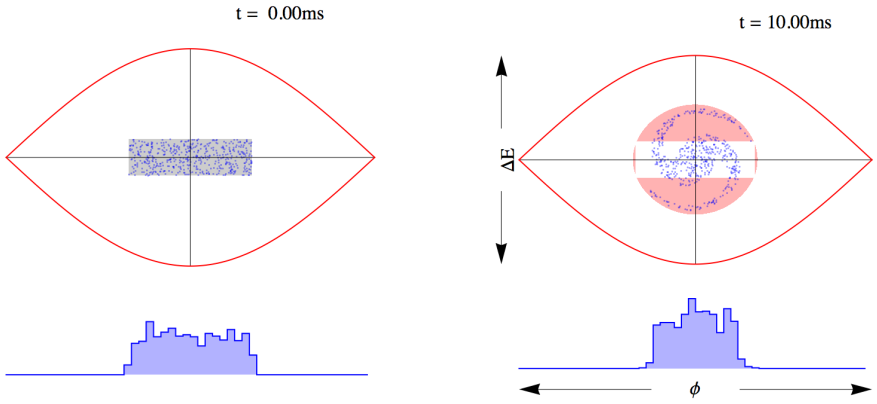

(a) (b)

Figure 2.3: Evolution of a particle distribution in phase space. Panel (a) shows a simulated distribution of particles in ϕ-ΔE phase space generated to lie within a rectangular region, and a histogram representing the projection of the distribution onto the ϕ axis. This projection is eqiuvalent to the signal that would be seen by an RWCM. Naively the area of the gray rectangle could be called the phase space area, or emittance, of the distribution. Panel (b) shows the same distribution tracked over many turns using eqn. 2.8. Because of the action dependent synchrotron frequency the bunch has filamented to fill the red area, which is larger than the initial rectangle. The emittance can be defined as the average action of the particles. The emittance of a bunch depends on the machine parameters defining the contours in the phase space.

system, beam scattering off of residual gas, or with other particles in the bunch can lead to emittance growth.

Non-adiabatic changes in machine parameters can also cause emittance growth. Transferring beam from one machine to another, and crossing transition are two examples of non-adiabatic changes that in machine parameters that I will discuss later in this chapter.

Because the stable area in phase space is limited, and beam losses occur for beam outside of the bucket, as discussed in the last section, the emittance growth throughout the acceleration cycle must be controlled.

2.2 Particle Stability During Acceleration

When $\phi_s \neq 0$ (or π above transition) the synchronous particle receives an energy kick each turn. All of the analysis thus far is unaffected by accelerating beam as long as the change from turn to turn is small, but the shape of the stability region changes. The bucket shape as a function of ϕ_s is shown in fig. 2.2. Notice in the figure that as ϕ_s increases the bucket area, width, and height all decrease. One unstable fixed point disappears and the buckets become asymmetric. Most dramatic of all is that the particles outside the stable region do not simple accrue phase indefinitely remaining at approximately the same energy. Now the particles outside the bucket stream away in the phase space. When far enough away in energy the particles will be lost as their radius decreases until they cannot be contained in the beam pipe.

When discussing the conditions for stability I mentioned that η, the phase slip factor, changes sign at a particular energy called the transition energy. Once we begin accelerating beam, it is necessary to understand something about transition, particularly if the machine in question accelerates through an energy range that includes transition. Not all machines cross transition, but transition will be important in later chapters.

2.2.1 Collective Effects

Until now I've ignored the fact that the particles in an accelerator are charged, except for their interaction with the accelerating cavity. In reality, particles interact with one another, and with the environment through their electromagnetic fields. By environment I mean all the elements needed to actually realize an accelerator, the resistive, conducting beam pipe that allows vacuum to be maintained, diagnostics, vacuum valves, and other discontinuities that interrupt the beam pipe. These interactions, and the dynamic consequences are referred to as collective effects.

The beam current is defined as the charge per unit time. For a 2-D phase space distribution is written as $\Psi(\phi, \Delta E)$, the linear charge density, $\lambda(\phi)$ is given by eqn. 2.20.

$$\lambda(\phi) = qeN \int \Psi(\phi, \Delta E)\, \mathrm{d}\Delta E \tag{2.20}$$

Where qe is the charge of a single particle, and N is the number of particles in the bunch. The beam current I is the charge per unit time. As the particle

bunch traverses the machine, the beam current interacts with the environment.

This interaction causes a modification of the energy if the perturbation to the dynamics described by eqn. 2.8 is significant the stability of the beam can be threatened, and the beam may be lost. The effect of this energy modification can be included by adding a term to the energy equation in eqn. 2.8 to get eqn. 2.21.

$$\Delta E_{n+1} = \Delta E_n + qV(\sin\phi_n - \sin\phi_s) + \delta E_Z(\phi_n) \tag{2.21}$$

Where δE_Z represents the energy modification for the particle under consideration. In general this contribution will change from turn to turn as the bunch evolves.

The energy loss due to some interaction is calculated by writing down the electric field on the bunch, and integrating over a single turn. This gives a voltage, which can be multiplied by the charge of the particle analogously to the voltage provided by the RF cavity. This voltage is then written as the product of the beam current and an impedance. Considering only longitudinal dynamics only fields directed azimuthally need to be considered, the presence of transverse field components also affects beam dynamics, but only transverse dynamics which I do not consider here.

The impedance can be reactive, or resistive. Resistive impedances cause net power dissipation, and reactive impedances do not, but they do distort the contours of constant action. The beam current is also often decomposed

into frequency components via a Fourier decomposition. The dependence of a reactive impedance on frequency characterizes the impedance as inductive or capacitive.

As an example, a common consideration is the space-charge field, the field of the bunch acting on itself.Because of boundary conditions on Maxwell's equations there is also a component that depends on the beam pipe. When the bunch is long compared to the beam pipe radius, the azimuthal component of the field due to the bunch, and the conducting wall only is given by, [3]

$$E_{sc} \approx \frac{qeg_0h^2}{4\pi\epsilon_0\gamma_s^2R_s^2}\frac{\partial\lambda}{\partial\phi} + E_w \tag{2.22}$$

Where R_s is the radius of the machine, and g_0 is a geometry factor related to the beam radius, a, and the pipe radius, b. For a cylindrical pipe $g_0 = 1 + 2\ln(b/a)$, and for a rectangular pipe of full width, w, and full height, h, $g_0 = 1 + 2\ln\left(\frac{2h}{\pi a}\tanh(\frac{\pi w}{2h})\right)$ [39], [29]. E_w is an additional contribution from the wall, which is often modeled as having a reactive coupling denoted by the impedance, L.

$$E_w \approx -\frac{qeh^2\omega_s^2L}{2\pi R_s}\frac{\partial\lambda}{\partial\phi} \tag{2.23}$$

Integrating the complete space-charge field around the circumference of the

[3]This discussion is only an overview to orient the reader to the idea of an impedance as it applies to a particle accelerator. For a more detailed discussion and derivations of these quantities see [9], and [30] or for specific applications near transition [39].

accelerator gives the voltage (which is the energy loss term for the space charge in eqn. 2.21), which can be written,

$$V(\phi) \approx \frac{2\pi qeh^2}{\tau_s}\frac{\partial\lambda}{\partial\phi}\left(\frac{g_0 Z_0}{2\beta_s^2\gamma_s^2} - \frac{2\pi L}{\tau_s}\right) \tag{2.24}$$

Where $Z_0 = 1/(\epsilon_0 c)$. Notice that the voltage depends on the derivative of the charge distribution, meaning a perfectly uniform longitudinal distribution has no longitudinal space-charge.

This voltage can also be expressed as the product of the beam current and the space-charge impedance per revolution harmonic, $Z_{||}/n$

$$|\mathrm{Z}_{||}/n| = i\left(\frac{g_0 Z_0}{2\beta_s^2\gamma_s^2} - \frac{2\pi L}{\tau_s}\right) \tag{2.25}$$

Because this particular effect is out of phase with the bunch, it is classified as reactive. This means the space-charge force does not dissipate energy. Resistive impedances will dissipate energy. The distortions in phase space, or energy dissipiation due to either type of impedance can lead to instability through perturbation to the dynamics described in the beginning of this chapter.

2.2.2 Transition Crossing

There are certain regimes where collective effects become particularly troublesome. This section describes transition crossing in more detail, and

specifically discusses how collective effects lead to bunch diffusion when crossing transition.

The first thing I did when writing down the equations of motion was to quantify the amount of phase slippage for off-momentum particles with the parameter η, eqns. 2.4, and 2.5. I pointed out that the slip-factor, η, is dependent on energy, and changes sign as the synchronous energy accelerates through the transition energy. Fig. 2.4 illustrates how particles slip relative to one another when the restoring force provided by the RF cavity is absent, both above and below transition, and the effect on particle motion when the proper restoring force is applied.

The change in sign of η causes the direction of the slippage shown in fig. 2.4 to change above transition. The reason for this is the difference in the way velocity and path length change with a change in momentum, see fig. 1.4. Below transition the velocity of a particle dominates its phase slippage, meaning higher momentum, faster particles traverse the ring more quickly and reach the RF cavity earlier. Above transition the velocity is effectively saturated, asymptotically approaching the speed of light, at which point the change in path length due momentum dependence of the bending radius dominates the difference in revolution time relative to the synchronous partilce, meaning higher momentum particles, with larger bending radii, traverse the machine more slowly, and arrive at the RF cavity later than the synchronous particle. (In both cases the opposite is true for lower momentum particles.)

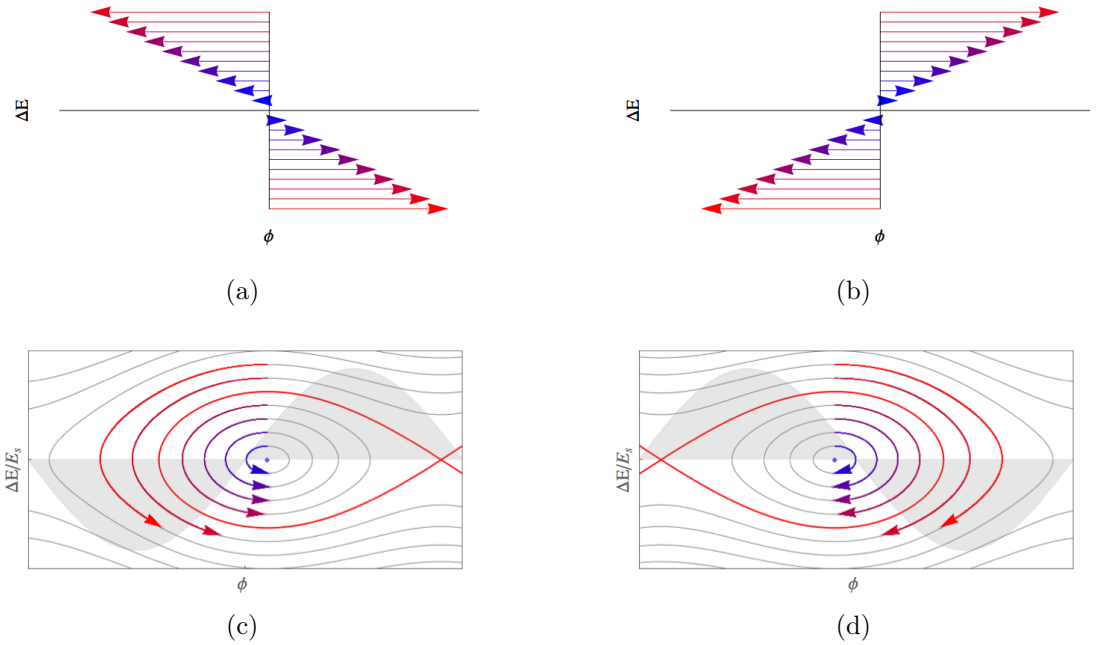

Figure 2.4: Panels (a) and (b) show particle trajectories in the longitudinal phase space in the absence of a restoring force. When the synchronous energy is below the transition energy the off-momentum particle's phase slip is dominated by the increase(or decrease) in velocity. Meaning higher momentum, faster, particles arrive earlier on subsequent turns. After transition, the increase(or decrease) in path length due to increasing bending radius dominates, causing higher momentum particles to arrive later on each subsequent turn. Because of this the potential that provides the restoring force must be shifted in phase by $\pi - 2\phi_s$ when the energy crosses the transition energy to maintain the same acceleration rate.

This change in relative slippage means that the slope of the restoring potential must change in order to maintain stability. This is evident from eqn. 2.12, and is also indicated in fig. 2.4. In order to change direction, the phase slip must pass through zero, which occurs at transition. When the synchronous energy equals the transition energy revolution frequency is independent of momentum. Because of this the oscillatory synchrotron motion demonstrated in the beginning of this chapter stops. This can also be seen in eqn. 2.12, since $\nu_s \propto \sqrt{\eta}$.

As a matter of convention when discussing transition, the moment of transition is usually taken as time $t = 0$. Times approaching transition are negative, and times exiting transition are positive. In addition, the acceleration rate, the change in the Lorentz factor with time, $\dot{\gamma}_s$ is taken to be a constant value. Near transition γ_s can be expressed as a function of time,

$$\gamma_s(t) = \gamma_t + \dot{\gamma}_s t \tag{2.26}$$

Which gives $\gamma_s(0) = \gamma_t$ as expected. The acceleration rate, $\dot{\gamma}_s$, is just the change in energy of the synchronous particle per revolution, divided by the time to complete one revolution, and the particle mass.

$$\dot{\gamma}_s = \frac{qeV \sin \phi_s}{m\tau_s} \tag{2.27}$$

Given this expression for $\gamma_s(t)$ near transition, it is also useful to have an expression for η_0 near transition. Using eqns. 2.4, and 2.26 and expanding around $t = 0$ gives,

$$\eta_0(t) = \frac{1}{\gamma_t^2} - \frac{1}{(\gamma_t + \dot{\gamma_s} t)^2} \approx \frac{2\dot{\gamma_s}}{\gamma_t^3} t \tag{2.28}$$

This section laid out the terms, and approximations used when discussing transition. Transition is taken as the $t = 0$. The acceleration rate is taken to be constant, this is almost always true at least in some neighborhood of transition. The slip factor is then determined by this rate, and is proportional to the time from transition. In addition to the assumptions made about the operating parameters of the synchrotron, a treatment of collective effects requires some assumptions about the bunch shape, since the energy difference is dependent on the bunch current.

2.2.3 A Bunch Model

The next few sections involve some calculations of emittance growth that require a bunch distribution to evaluate. Because the non-linear shearing is only applicable for large bunches, and difficult to evaluate, I will restrict these analytical considerations to the linear region near the synchronous particle, with motion described by eqn. 2.11. This can be derived from the Hamiltonian in eqn. 2.29, which is needed to specify a 2-D distribution in the phase space.

$$H = \frac{2\pi h \eta_0}{E_s \beta_s^2} \frac{\Delta E^2}{2} - qeV \cos\phi_s \delta \frac{\phi^2}{2} \tag{2.29}$$

When the distribution given by eqn. 2.30 is integrated according to eqn. 2.20

the result is the linear charge distribution in eqn. 2.31, which is parabolic in $\delta\phi$. The benefits of a parabolic distribution are: bunches are often gaussian, and appropriately normalized a parabolic approximates a gaussian distribution near the mean, and the space-charge voltage, proportional to $\partial\lambda/\partial\phi$, is linear.

$$\Psi(\phi, \Delta E) = \begin{cases} \frac{3}{4\pi H_{max}} \sqrt{1 - H(\phi, \Delta E)/H_{max}} & H < H_{max} \\ 0 & \text{otherwise} \end{cases} \tag{2.30}$$

Where H_{max} is the maximum action of the distribution given in terms of the bunch half width in ϕ, $\hat{\delta\phi}$, $H_{max} = -qeV \cos\phi_s \hat{\delta\phi}^2/2$.

$$\lambda_p(\phi) = \begin{cases} \frac{3qeN}{4\hat{\delta\phi}} \left(1 - \frac{\delta\phi^2}{\hat{\delta\phi}^2}\right) & |\phi| \leq \hat{\phi} \\ 0 & \text{otherwise} \end{cases} \tag{2.31}$$

Where q, e, N are the particle charge state, electron charge, and number of particles in the bunch respectively.

One of the questions I address in later measurements is whether the emittance growth functions derived under this approximation are applicable for more realistic beam distributions in Fermilab's Main Injector, which are often quite different from gaussian because of some RF manipulations detailed in ch. 4.

2.2.4 The Non-Adiabtic Regime

To maintain the phase space area of a distribution parameters defining phase space contours should be changed adiabatically. That is the fractional

change in the synchrotron frequency, Ω_s in eqn. 2.13, caused by the changing parameter should be small relative to the synchrotron frequency or,

$$\frac{1}{\Omega_s^2}\frac{\mathrm{d}\Omega_s}{\mathrm{d}t} \ll 1 \tag{2.32}$$

For a constant acceleration near transition, it is reasonable to assume that only η is changing, in which case the adiabatic condition becomes,

$$\frac{1}{\Omega_s\eta_0}\frac{\mathrm{d}\eta_0}{\mathrm{d}t} \ll 1 \tag{2.33}$$

For any rate of change of η, while η is near 0 on either side of transition there will be some time during which this condition is not satisfied, and motion is non-adiabatic. This is called the non-adiabatic time. [4] Using eqns. 2.33 and 2.28 the non-adiabatic time can be calculated,

$$T_{na} = \left(\frac{\pi E_s \beta_s \gamma_t^3 \tau_s}{eV|cos\ \phi_s|\dot{\gamma_s} 2\pi h}\right)^{1/3} \tag{2.34}$$

Because η is going to 0, the motion in ϕ slows down eventually stopping altogether, but the motion in ΔE continues. This leads to the bunch becoming narrower in ϕ, and taller in ΔE. The height of the bunch depends on how long the beam remains in the region where motion in ϕ is negligible, characterized by the non-adiabatic time. The width in ϕ, and height in ΔE, of a parabolic distribution dimension at the moment of transition crossing are given by:

[4] The non-adiabatic time is sometimes called the characteristic time for non-adiabtic motion and is denoted by T_c

$$\hat{\sigma}_\phi(\epsilon) = 0.52\sqrt{\frac{2\pi h\epsilon}{T_{na}qeV|\cos\ \phi_s|}}$$
$$\hat{\sigma}_\delta(\epsilon) = 0.71\frac{1}{\tau_s E_s \beta_s^2}\sqrt{2\pi h T_{na}qeV|cos\ \phi_s|\epsilon} \tag{2.35}$$

Notice that the bunch widths depend on T_{na}, and the emittance, ϵ. A full derivation of these quantities involves a detailed treatment of the envelope function for bunches near transition, which I do not cover here. A full treatment is given in [39].

If there are no higher order terms in η, this process is reversible and there will be no emittance growth. But non-linear terms of η change the effectiveness of the focusing potential near transition and lead to emittance growth.

Because each higher order of η in eqn. 2.5 depends on another factor of $\delta p/p_s$, which is much less than one, each subsequent term is important on a shorter timescale, and is thus less important unless the acceleration is halted near transition. I will only consider the first non-linear term here.

Fig. 2.5 shows the phase space when only the first non-linear term of the slip-factor, η_1 acts on particles. Notice how the trajectories form a parabola around the synchronous particle because the non-linear term adds an extra factor of $\delta p/p_s$. Because η_1 provides the same slip on either side of the synchronous energy, only the momenta above or below the synchronous value will experience a restoring force during the time this non-linear term dominates. And the opposite side will be stable after transition. This leads to

asymmetric tails that increase the bunch area. This is called a single particle effect, because it happens regardless of the number of particles in the bunch. Whether the momenta above, or below receive a restoring force above or below transition depends on the sign of η_1.

Because of the spread in the bunch the second order term, η_1, affects the high momentum tail of the bunch when the product of the momentum and the non-linear term are greater than the first order term, given by the non-linear condition in eqn. 2.36

$$|\eta_1 \delta p/p_s| \geq |\eta_0| \tag{2.36}$$

Using the values of of η_0, η_1 given in eqn. 2.7, the maximum spread in momentum given by eqn. 2.35, and the condition in eqn. 2.36 a non-linear time can be computed, during which the second order term dominates the motion.

$$T_{nl} = \frac{|(\alpha_1 + 3/2\beta_s^2)|\hat{\sigma}_\delta \gamma_t}{2\dot{\gamma_s}} \tag{2.37}$$

Where α_1 is the first non-linear term of the momentum compaction factor discussed earlier in this chapter, and introduced in eqn. 2.6. This non-linear time characterizes the time during which unfocused particle trajectories resemble fig. 2.5, rather than those in the upper plots of fig. 2.4.

The selective defocusing due to the momentum spread of the bunch leads to a chromatic mismatch when the first order term of η once again dominates. Defocusing occurs for both the high and low energy tails, but on

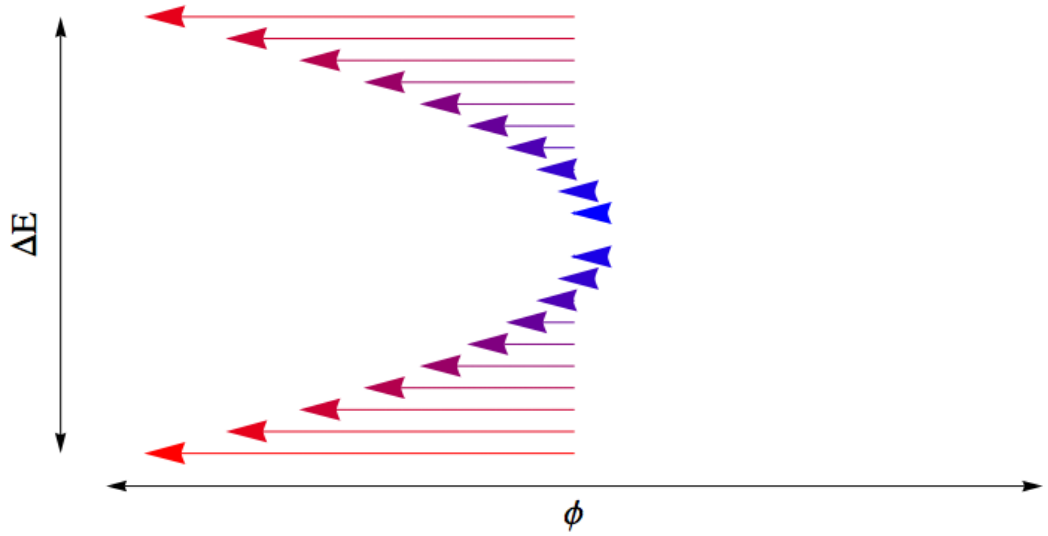

Figure 2.5: Particle trajectories in the longitudinal phase space in the absence of a restoring force when only the first non-linear term is present, and positive. Particles above and below the synchronous energy slip in the same direction and the restoring potential provided by the RF cavity is only effective on the high momentum particles. When the phase of the RF changes at transition, the first non-linear term does not change sign so the RF cavity then focuses only the low momentum particles. This selective defocusing creates tails in momentum that increase the bunch area in proportion to the time it takes to accelerate half the bunch across transition.

opposite sides of transition. Where each momentum tail is defocused depends on the sign of η_1. Fig. 2.5 depicts a positive η_1, which would defocus the high momentum particles before transition, and the low momentum particles after transition. The particles that are defocused will lie on different action contours when the first order term once again becomes dominant, this leads to emittance growth when crossing transition. Because these tails are a consequence of the momentum spread of the bunch through T_{nl}, this effect is called a chromatic mismatch.

The amount of emittance growth is limited by the ratio of the time during which motion is non-adiabatic, and the time during which the non-linear term dominates. There are two cases for the emittance growth given in eqn. 2.38 [39].

$$\frac{\Delta\epsilon}{\epsilon} \approx \begin{cases} 0.76\frac{T_{nl}}{T_c} & T_{nl} \ll T_c \\ e^{\frac{4}{3}\left(\frac{T_{nl}}{T_c}\right)^{3/2}} - 1 & T_{nl} \gg T_c \end{cases} \tag{2.38}$$

Where $\Delta\epsilon/\epsilon$ is the fractional emittance growth, the difference in final and initial, over the initial value. [5] When the non-linear effects are present outside of the non-adiabatic regime, the growth is exponential relative to the ratio of characteristic times. Usually, though $T_{nl} > T_c$, and the growth is only linear ($T_{nl} > T_c$ also indicates roughly that the beam is below the microwave

[5] What constitutes the 'initial' and 'final' emittance though? Because the growth here is caused by the non-adiabatic, and non-linear motions of the particles, the initial, and final emittances should be measured at times far away from transition compared to the longer of the non-adiabatic, or non-linear times.

instability threshold). This growth is known as a single particle effect, because the only consideration is the dynamics of single particles.

Because of the bunch narrowing I described, the current of the bunch increases, reaching a maximum value at transition. The minimum width of a parabolic bunch is given in eqn. 2.35. This increase in current means the bunches couple to the machine through the impedance discussed in sec. 2.2.1.

2.2.5 Collective Effects at Transition

In sec. 2.2.1 I discussed effects caused by the fields of the charged particles in the bunch. The interaction of these fields was characterized by the longitudinal impedance, $Z_{||}$. Because of non-adiabatic motion the bunch narrows approaching transition, and the peak current increases. This means that for a given impedance, collective effects become worse at transition.

In the previous section, I discussed how the first non-linear term of η created an asymmetry in the particle motion that counteracted the focusing potential, and how this perturbation does not change sign with the phase shift induced upon crossing transition. A similar factor leads to emittance growth when collective effects are considered.

Fig. 2.6 shows an example of how the space charge voltage modifies the focusing potential provided by the RF cavity. The voltage induced by the beam current coupling to the impedance adds a perturbation the focusing potential provided by the RF cavity. Panel (a) shows the linear charge density of a gaussian bunch, and its derivative, which is proportional to the space

charge impedance given in eqn. 2.25. The beam current times the impedance gives the perturbing voltage, which is linear near the synchronous value for this distribution.

Panels (c), and (d) show how the voltage induced by the space charge impedance perturbs the nominal potential provided by the RF cavity, below and above transition respectively. A phase space contour for the nominal potential , and the perturbed potential which is slightly distorted are indicated. In the linear region this perturbation mimics a change in the voltage of the RF cavity. Outside of the region where the perturbation, and the nominal potential are linear the effect will be more complicated but sill represents a distortion of the phase space contours. Panel (b) isolates the contours indicated in panels (c) and (d) to emphasize the difference before and after transition. Because the perturbation doesn't change when the nominal potential does, a bunch initially filling the contour before transition will be mismatched relative to the same contour after transition. This leads to emittance growth.

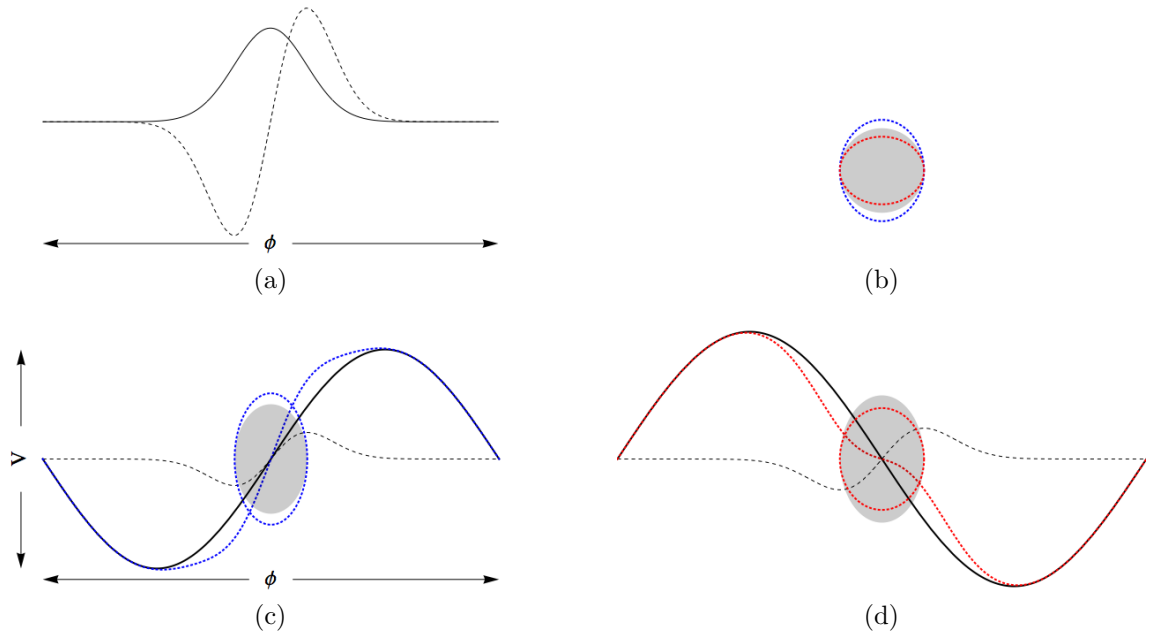

Figure 2.6: Growth across transition due to the space-charge impedance. (a) A gaussian linear charge distribution, $\lambda(\phi)$, (solid) and its derivative (dashed), which is proportional to the space-charge impedance, eqn. 2.25. (c), and (d) show the focusing potential provided by the RF cavity with the appropriate phase shift (solid), below and above transition, the space-charge voltage (dashed), and the potential resulting from the sum of the two, (dotted, in color). The phase space contour of the unperturbed potential is shown as a gray disk, and a perturbed contour as a dotted circle. (b) Shows how the perturbed potentials do not match before and after transition. A bunch which initially fills the blue area will spread out in the phase space causing emittance growth.

This example is just used to illustrate how an impedance, which is not controllable, perturbs the bunch shape differently on either side of transition, causing emittance growth. There are actually several mechanisms that lead emittance growth in this way, each characterized by an impedance.

There are also impedances associated with strong resonances, typically from higher order modes in cavities or resonant structures in the beamline. I do not consider those modes here. Impedances which are not due to sharp resonances are called broadband, and act over a range of frequencies.

Because the voltage induced on the bunch is the beam current times the impedance, to give values for the emittance growth, it is necessary to assume a beam model and provide current measurements. In the case of a parabolic bunch, the peak current is given as,

$$\hat{I} = \frac{3\pi h N_b q e}{2\tau_s \hat{\sigma}_\phi(\epsilon)} \tag{2.39}$$

Where N_b is the number of particles in the bunch, this number may have some correlation to the emittance, ϵ. $\hat{\sigma}_\phi$ is the minimum bunch width, given in eqn. 2.35, which depends on the initial emittance.

The fractional emittance growth due to a broadband reactive impedance is,

$$\frac{\Delta\epsilon}{\epsilon} \approx \frac{h\hat{I}(\epsilon)|Z_{||}/n|}{3V|\cos\ \phi_s|\hat{\sigma}_\phi(\epsilon)^2} \tag{2.40}$$

And the fractional emittance growth due to a broadband resistive impedance is,

$$\frac{\Delta\epsilon}{\epsilon} \approx \frac{\hat{I}(\epsilon)R}{\sqrt{6}V|\cos\ \phi_s|\hat{\sigma}_\phi(\epsilon)} \tag{2.41}$$

Where R is the value of the resistive impedance.

Both of these expressions depend on ϵ through the peak current, if there is a correlation between bunch intensity and emittance, and through the bunch width.

2.3 Simulation of Injection Errors

I stated in ch. 1 that any high energy particle accelerator complex will employ a series of machines to reach the final desired energy. Transferring beam from one machine to another can be thought of as a non-adiabatic change in parameters since a single turn is the smallest time interval typically considered. Thus the transfer of beam from one machine to another can lead to emittance growth, and losses.

I discuss two types of errors here, offset errors, and matching errors. Sometimes called dipole and quadrupole errors.

2.3.1 Dipole errors

Fig. 2.7 illustrates the mechanism leading to dipole errors. Transferring beam from one machine to another can be thought of as changing all the

machine parameters in a single turn. If the only parameters to change are the synchronous energy, or phase, the distribution of particles is suddenly offset from the bucket center. This offset in the phase space is called a dipole error, since an offset charge has a dipole moment. If the error were large compared to the bunch size some particles may fall into the non-overlapping region between the red and gray separatrices, these particles would be lost.

Fig. 2.8 shows the time evolution of the dipole error created in fig. 2.7. Because the central moment of the beam is now at a larger action, the emittance has increased. Because the particles nearer the separatrix have lower rotation frequencies in the phase space the bunch also begins to distort over time. In the final panel of fig. 2.8, after much time as elapsed, the bunch is much longer than it initially was, and has a larger energy spread. If special care was taken to ensure either a short bunch, or a small energy spread the dipole error has undone the effort.

However, a dipole error is not necessarily problematic if the available stable phase space is ample compared to the bunch size, and the precise structure of the beam is not important. But for some machines or experiments phase space may be at a premium, or the structure may be important. For example, a light source, or Free Electron Laser might have strict requirements on the longitudinal structure of the beam in order to provide optimal light yield. Or beam from multiple bunches may be combined into a single bucket, which I will discuss in a later chapter. In this case the phase space area occupied by the bunch needs to be controlled to ensure no beam is lost.

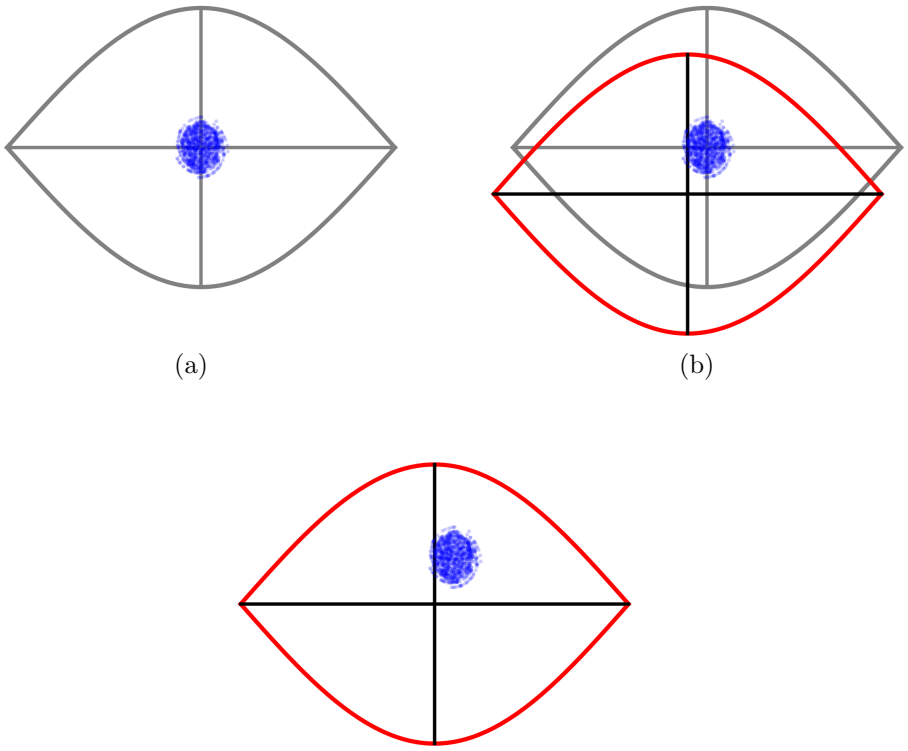

Figure 2.7: Beam transfer between machines exhibiting a dipole error. (a) shows the phase space coordinates of a bunch of particles centered in the region of stability. (b) shows the same bunch of particles upon transfer to a different machine. Because this transfer occurs in one turn, the distribution does not evolve significanly between machines. In this case the RF systems are out of phase, and the synchronous energy is not the same in both machines. The contours showing the stable area for each machine overlap, but not completely. The distribution is not centered relative on the synchronous particle in the new machine. (c) removes the separatrix from the first machine to make clear how a dipole error results from the beam transfer when the synchronous energy and phase are not matched between machines. Fig. 2.8 shows the time evolution this dipole error.

A dipole error can sometimes be intentional. Because beam subject to a dipole error diffuses and fills up more of the phase space, increasing the bunch length, the particle density is lower. When coulomb interactions of particles with one another, or the surrounding beam pipe become an issue this decreased charge density can help prevent certain instabilities.

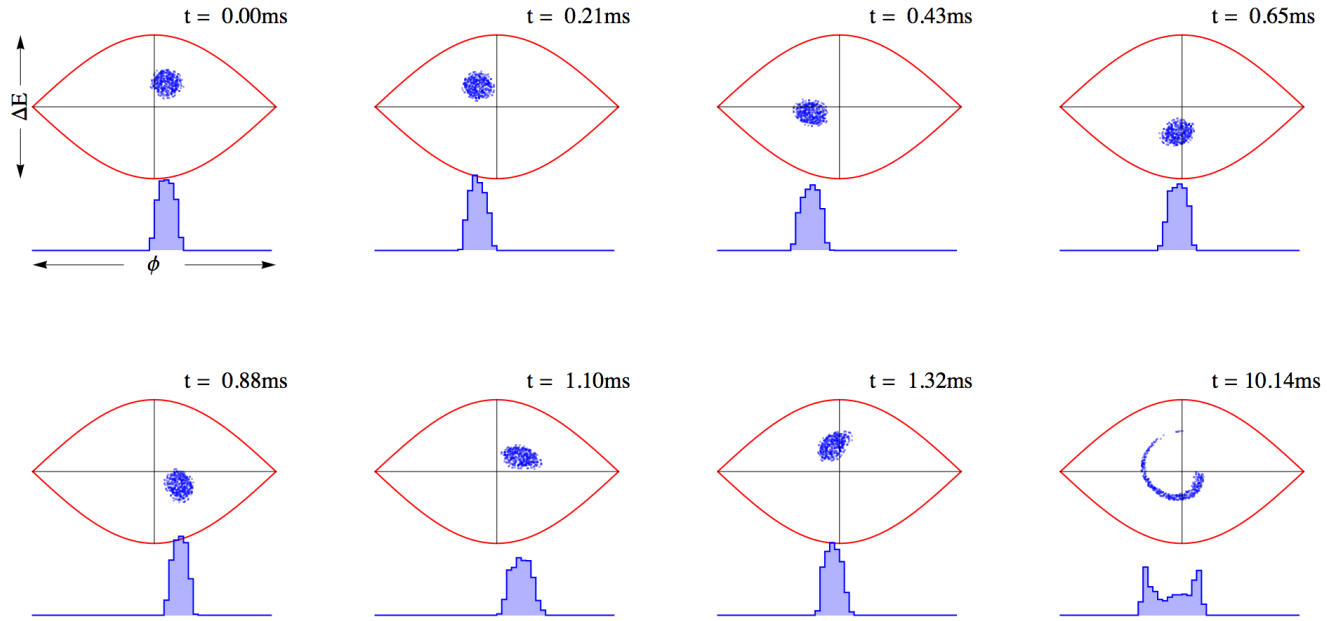

Figure 2.8: These frames show the time evolution of the ϕ-ΔE phase space coordinates of the distribution of particles from fig. 2.7 using typical Main Injector parameters for tracking. The histogram included at the bottom of each frame shows the projection of the distribution onto the ϕ axis. This histogram is equivalent to the signal that would be seen by an RWCM. Time $t = 0$ represents the moment of injection, notice the error in ϕ, and ΔE. Because particles of different action have different frequencies in the phase space the distribution does not move coherently about the center, but shears with increasing time. The final frame is much later, illustrating how the bunch filaments due to the dipole injection error.

2.3.2 Quadrupole errors

Fig. 2.9 illustrates another error that occurs during beam transfer, a bucket shape mismatch, also called a quadrupole error. A quadrupole error occurs when the contours of the first machine do not have the same aspect ratio as the those in the second machine. In fig. 2.9 the gray contours illustrate the bucket of the initial machine. The separatrix is shown, but so is the contour indicating the maximum action represented in the distribution. When the synchronous energy, and phase are aligned between the machines, but one of the parameters defining the bucket height in eqn. 2.17 has an error, the voltage for instance, the aspect ratio of the buckets will be different. When the new values are established the bunch will not fill the new contour. This is called a bunch mismatch, or a quadrupole error.

Fig. 2.10 shows the time evolution of the quadrupole error depicted in fig. 2.9. The distribution appears to tumble in the phase space as the particles at higher actions take longer to complete a phase space revolution than the particles near the center of the bucket. A RWCM will only have access to the projections of this motion on the time axis, also shown in fig. 2.10. In projection the bunch seems to breathe, the width of the projection expanding and contracting as the bunch tumbles. The final frame in fig. 2.10 shows the distribution after many rotations in the phase space, the beam has filamented to fill a much larger area than it initially did.

Notice in fig. 2.10 that even in this extreme case the bunch rotates almost rigidly, for the first synchrotron period or so, only filamenting over

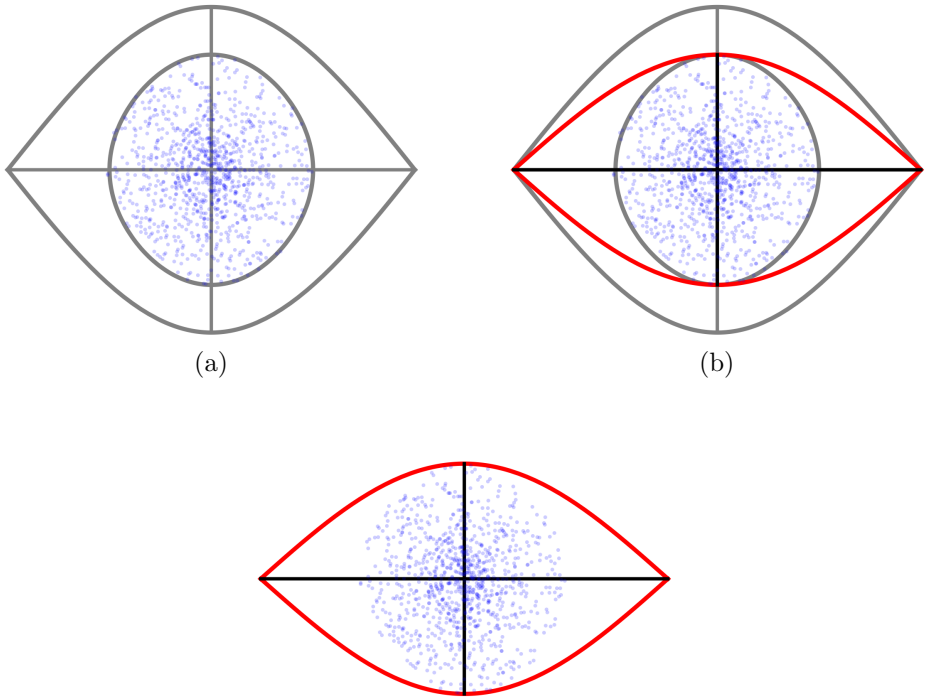

Figure 2.9: Beam transfer between machines exhibiting a quadrupole error. (a) shows the phase space coordinates of a bunch of particles centered in the region of stability. The phase space contour containing the particles is also indicated. (b) shows the same bunch of particles upon transfer to a different machine. In this case the bunch shapes are mismatched, the voltage of the first machine is double the voltage in the second. The distribution is centered, the synchronous phase and energy are the same, but beam does not conform to the contours of the second bucket. (c) removes the separatrix from the first machine to make more clear how a quadrupole error results from the beam transfer when the bucket shapes are mismatched. Fig. 2.10 shows the time evolution of this quadrupole error.

many turns. This can be used to match beam in some cases. If the first machine has a high voltage, as in this example, the voltage can be turned down a quarter of a synchrotron period before the beam transfer. This will induce a rotation in the bunch, compare the first and third times in fig. 2.10. If the beam is transferred after a quarter of a synchrotron period, the aspect ratio has changed. This may help matching, and reduce the overall emittance growth due to filamentation.

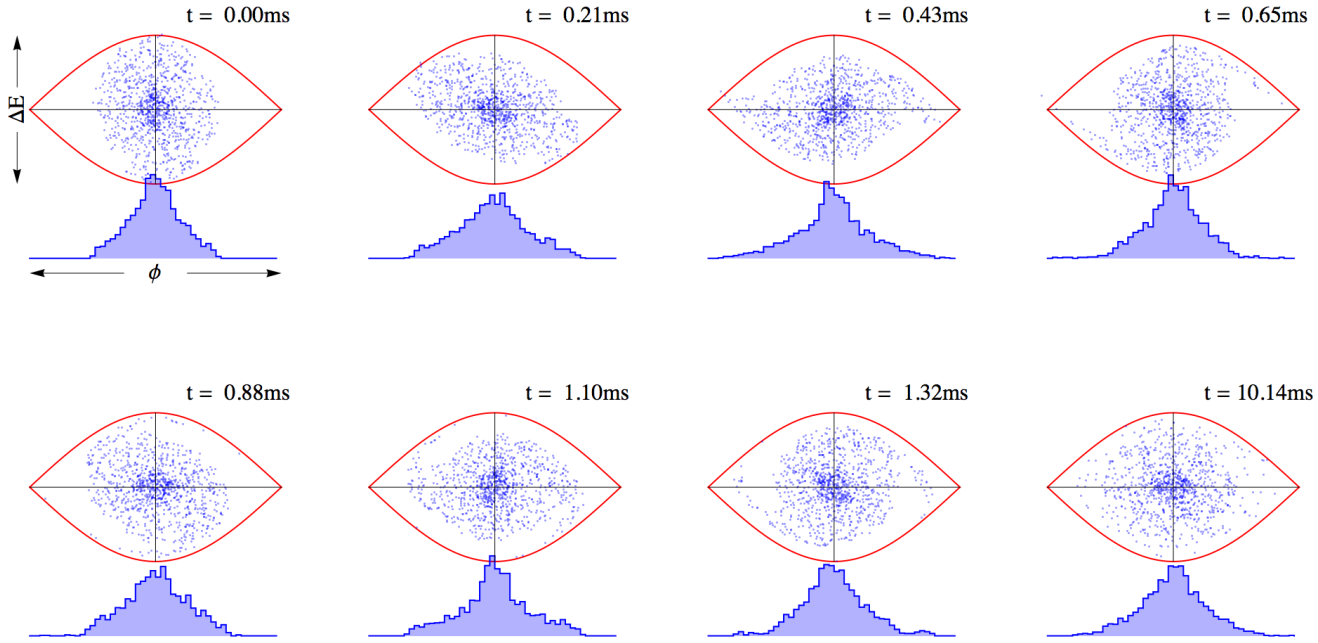

Figure 2.10: The time evolution of the phase space coordinates of the distribution of particles from fig. 2.9 using typical parameters from the Main Injector for tracking. The histogram included at the bottom of each frame shows the projection of the distribution onto the ϕ axis. This histogram is equivalent to the signal that would be seen by an RWCM. Time $t = 0$ represents the moment of injection, there is no error in ϕ, and ΔE, but the distribution shape is mismatched relative to the contours. Because particles of different action have different frequencies in the phase space the distribution shears with increasing time. The final frame is much later, illustrating how the bunch filaments due to the quadrupole injection error.

2.4 Summary of Longitudinal Beam Dynamics

This chapter gave a brief introduction to longitudinal beam dynamics. Only the details relevant to the rest of this document were presented, but there is a wealth of existing knowledge, and active research regarding longitudinal dynamics, transition crossing, and instabilities. To the interested reader I recommend the following texts, [12], [27], and [9].

In order to discuss particle acceleration, I described a simple model of a synchrotron. Bending magnets provide the fields needed to circulate energetic particles of a given momentum at a given radius. A radio frequency cavity provides a time varying field that is synchronized with the revolution of a particle known as the synchronous particle.

When the RF cavity is properly phased, there is a spread in time, and energy for which particles will remain close to the synchronous particle. This shape stable region in phase space, called a bucket, was derived, and it's behavior under normal operating conditions described. Because of this stability it is possible to accelerate ensembles of particles. Under ideal conditions the phase space area occupied by an ensemble in units of $[eV \cdot sec]$, called the emittance, remains constant. In the case of errors, or dissipative effects the emittance increases. Increases in emittance can lead to beam loss, and distortion of the energy, and time profiles of the bunch.

The special point known as transition occurs when the slip-factor changes sign. This means that relative to the synchronous particle, higher energy par-

ticles actually traverse the ring more slowly, because the contribution to revolution frequency from increased path length becomes more important than the smaller increase in velocity. Transition is a dangerous time when many of the assumptions used to derive the longitudinal equations of motion breakdown. The change in parameters becomes non-adiabatic, non-linear terms begin to contribute significantly to the phase space motion, a characteristic bunch narrowing increases the beam's coupling with the environment. All of these considerations can cause emittance growth and must be evaluated when designing a machine that will cross transition.

In order to reach increasingly higher energies it is usually necessary to transfer beam between machines. This transfer can be thought of as a non-adiabatic change in parameters. In order to prevent emittance growth it is necessary to match the incoming beam and the machine. Dipole errors arise from the bunch being injected with a coherent error in either energy or phase. A shape mismatch, also called a quadrupole error, occurs when the parameters defining the phase space contours do not match, such that the incoming bunch is distorted in aspect relative to the contours of the new machine. Both of these errors lead to emittance growth through filamentation.

A detailed knowledge of the distribution of particles in the longitudinal phase space is necessary for the efficient operation of a machine, and the transfer of particles between machines. In order to control beam losses, and the longitudinal properties of the beam we need intuitive, precise diagnostics. Tools that allow the operators of a machine to quickly diagnose problems

and extract quantities of interest will allow physicists and engineers to push to ever higher beam power in the face of increasingly stringent operational tolerances. The next chapter will introduce the theory of longitudinal phase space tomography, which allows the longitudinal phase space to be visualized as presented in this chapter.

Chapter 3

Tomography

Tomography is the process of building an image of an object from slices, or projections, of the object. Such an image is called a tomogram. Using many angularly separated projections along one dimension tomography enables reconstruction of an image in two dimensions. The tomogram itself is simply a recasting of existing data into a more meaningful coordinate system.

The most widely known examples of tomography are in medicine, where 1-D information obtained non-invasively can be used to reconstruct 2-D images of the inside of a patient. In an x-ray CT scanner, shown in fig. 3.1, x-rays are emitted on one side of the body and the transmitted x-rays recorded on a detector on the opposing side. The attenuation of the x-rays gives information about the total density along the rays from emitter to detector plane. The opposing detector/emitter system rotates around about a point located in the body providing many 1-D attenuation curves of one 2-D slice at various angles to reconstruct that slice in 2-D, shown in fig. 3.1. Many 2-D images can be obtained at small displacements along the rotation axis and then stacked to build up a 3-D image, also shown.

There are many technologies used to achieve the same end, Positron

Emission Tomography(PET), for example, uses detectors placed around the patient to reconstruct the origin of pairs of gamma rays produced by e^+-e^- annihilations from radioactive isotopes injected into a patient. Selective absorption of the particular radioactive element used allows the reconstruction of the location of certain cell types.

The earliest mathematical statement of the problem was in 1917 by Radon [32], who described an integral transform between a function in x,y to a function of distance along an axis of projection x' and the projection angle θ. To reconstruct a distribution requires two pieces of information: projections at many angles, and knowledge of the transformation that governs the evolution of the distribution between angular acquisitions, which may be a spatial rotation or some more complicated dynamics. As long as these two conditions are met, the technology used to realize the reconstruction is largely irrelevant, although details in implementation might change.

Fig. 3.2 depicts a schematic of a tomographic reconstruction setup. An object is located at the center of a ring of detectors. (Alternatively the image could represent one detector which can be swiveled around the object, keeping track of the angle used for each acquisition.) Each detector only records a 1-D trace of the object along a local horizotnal axis. This information may be the attenuation of x-rays in a medical scan, or some other measure of the total density along a line through an object. The information from all of the detectors is combined in a process known as backprojection. The resulting guess is then compared to the data and iteratively refined until projections of

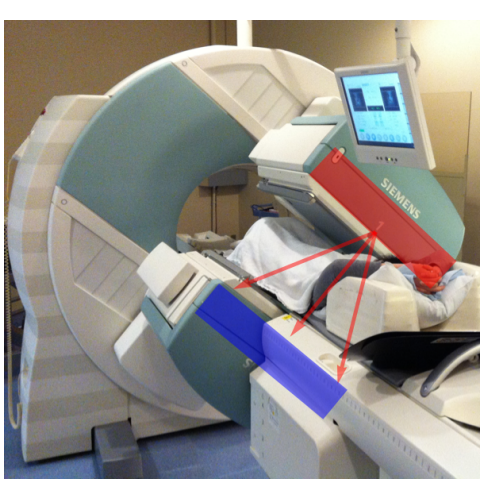

(a) Scanner

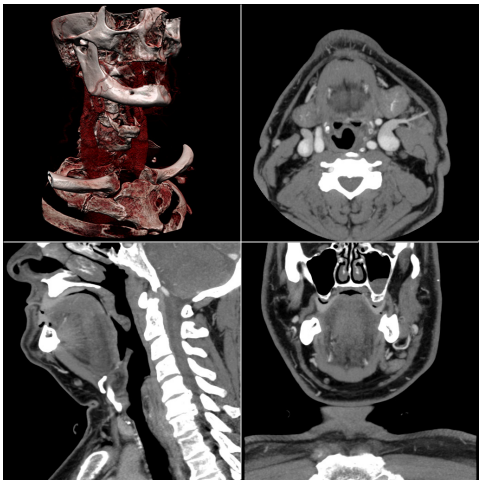

(b) Reconstruction

Figure 3.1: An x-ray CT scanner used to see inside a patient, and the resulting images. X-rays are emitted from one of the flat planes surrounding the patient in (a) and travel through the body. The transmitted rays are detected on the opposite side of the body. This is repeated at many angles relative to the patient. Tomographic reconstruction uses these 1-D transmission profiles to reconstruct 2-D images of the body like those in (b). A 3-D image of the body can even be constructed by stacking many 2-D images.

the guess agree with the original data to a predefined level of precision.

In the previous chapter I showed that an ensemble of particles in an accelerator undergo oscillatory motion about a fixed point, tumbling within the 2-D phase space. Using standard accelerator diagnostics projections of the ensemble can be obtained over many machine revolutions, sampling the distribution at various angles as it tumbles. This provides the raw input for a tomographic reconstruction algorithm, with the 2-D ϕ-ΔE phase space distribution playing the role of the patient. As a technique for particle accelerators the method was first described, and implemented at Fermi National Accelerator Labortory by G.Jackson [24], [23]. The method was further developed at CERN in the 1990's by Hancock et al. [20], [22]. It has been used by Angoletta, Findlay, and Hancock in LEIR to make an independent measurement of the accelerating cavity voltage [18], a parameter normally difficult to accurately obtain.

Modern computing power is such that data acquisition, reconstruction, and analysis can be done in the control room in real-time for very large sets of data. This opens up diagnostic possibilities not previously available. Particularly the ability to view many bunches at once.

In this chapter I will describe the theory of longitudinal phase space tomography, and how it is used to extract physics quantities of interest from data, and identify errors in machine operation. In chs. 4 and 5 I will show how one can use a multi-bunch, real-time tomographic reconstruction system to make meaningful measurements of machine parameters.

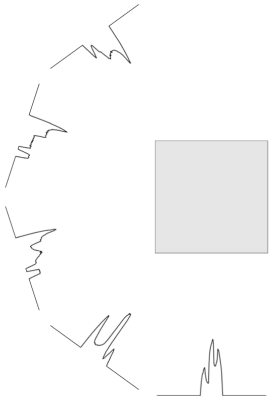

(a) Take Data

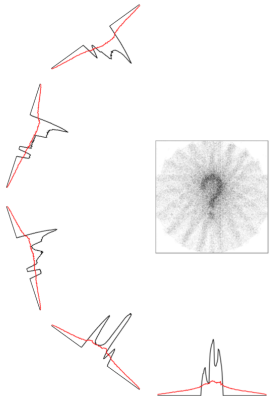

(b) Initial Guess

(c) Final Guess

Figure 3.2: Panel (a) shows an initially unknown object in a gray box. Projections of the objected are measured at various angles. Panel (b) shows the naive back projections of the data onto the area being scanned, constructive interference provides an initial guess at the object. The detectors in panel (b) show the original data in black, and the data resulting from the guess in red. The error between these two is then back projected to correct the guess. Panel (c) shows the result of many iterations. The data can be obtained by a single detector which is allowed to rotate between acquisitions, or a detector at each location. This image based on an image by S.Hancock [17].

3.1 Theory of 2-D Phase Space Tomography

Traces like those shown in fig. 3.3 represent a standard view of longitudinal phase space available in the control room. As the beam traverses the accelerator it induces image currents on the conducting beam pipe. A device called a resistive wall current monitor is used to measure these image currents [13], [38]. [1] The measured image currents represent a projection of the ϕ-ΔE phase space onto the ϕ axis. Each horizontal trace in fig. 3.3 is the measured beam current of a single bunch on a single turn. The axis going into the page, labelled in ms, gives the time of the revolution that a particular trace was acquired. Because of the oscillatory motion described in ch. 2 the subsequent projections represent the 2-D phase space distribution at different angles, this can be seen in the evolution of the traces with time.

The measurements taken with the resistive wall current monitor are non-destructive, and over many turns acquire the angular views of the phase space, similar to those presented in fig. 3.2, needed to reconstruct the 2-D distribution of particles in ϕ-ΔE phase space. In the simple example depicted in fig. 3.2 the angular position of each detector is known. Since the object depicted in the figure exists in real space, the transformation needed to combine the angular information is simply a rotation around the center of the object. In an accelerator this rotation is provided by the longitudinal dynamics de-

[1] A small length of the beam pipe is replaced with a non-conducting ceramaic gap to interrupt the image current. The gap is spanned by resistors of known value. The image currents are shunted through these resistors and measured.

rived in ch. 2 that cause the distribution to tumble over many turns. This means that in addition to the data acquired a model of the machine must be used to reconstruct the 2-D distribution.

Fig. 3.4 is the tomographic reconstruction made from all of the data in fig. 3.3, but the image only shows the distribution at the moment of the first, red trace. In this case there are two distinct tadpole shaped distributions that will circle one another, clockwise, following the equations of motion derived in ch. 2. The image shows that these distinct beamlets occupying the stable area of the phase space were initially separated in energy. This was the result of a scheme to increase beam intensity employed at Fermilab called slip-stacking that I will discuss in detail in ch. 4.

There are many techniques for solving tomographic reconstruction problems and recreating images such as 3.4 from data such as 3.3. I employ an algebraic reconstruction technique (ART), which is an iterative algorithm, meaning a solution is built up from an initially crude guess through successive refinement. ART was originally developed for use in tomographic reconstruction by Gordon, et. al. in 1970 [16], but is similar to a method for solving linear systems of equations developed by Kaczmarz in the 1937 [25].

A single iteration can be summed up in a few simple steps that provide a convincing argument for the method before we add any mathematical justification. Fig. 3.5 details a single iteration of the ART algorithm in the case of longitudinal phase space.

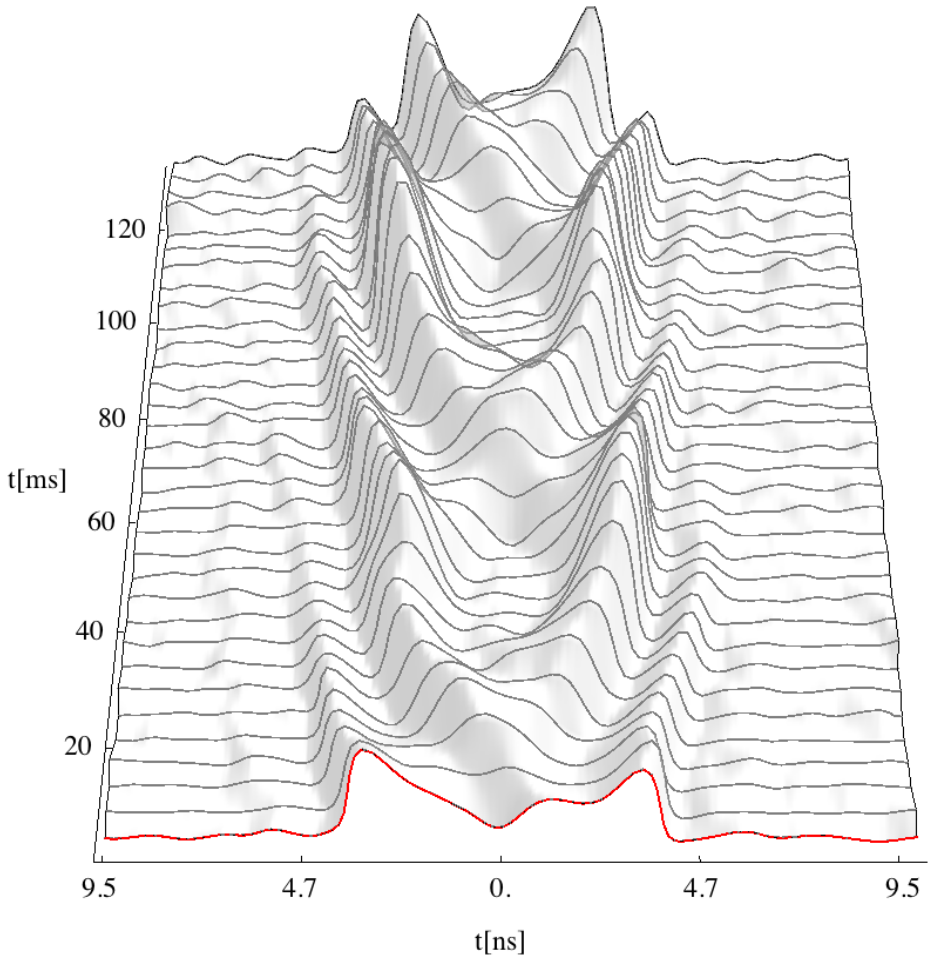

Figure 3.3: Each 1-D horizontal trace represents a measurement of current induced by a passing bunch of charged particle on the beam pipe at fixed location in the accelerator as a function of time, or phase. This is eqiuvalent to the projections shown beneath the phase space distributions in figs. 2.8, and 2.10. The synchronous phase is placed at the origin of the horizontal axis. The time axis going into the page, marked in ms, represents the time of the revolution on which the trace was taken. This is equivalent to the individual frames in figs. 2.8, and 2.10. This set of traces is more difficult to interpret than the 2-D distributions shown in those figures.

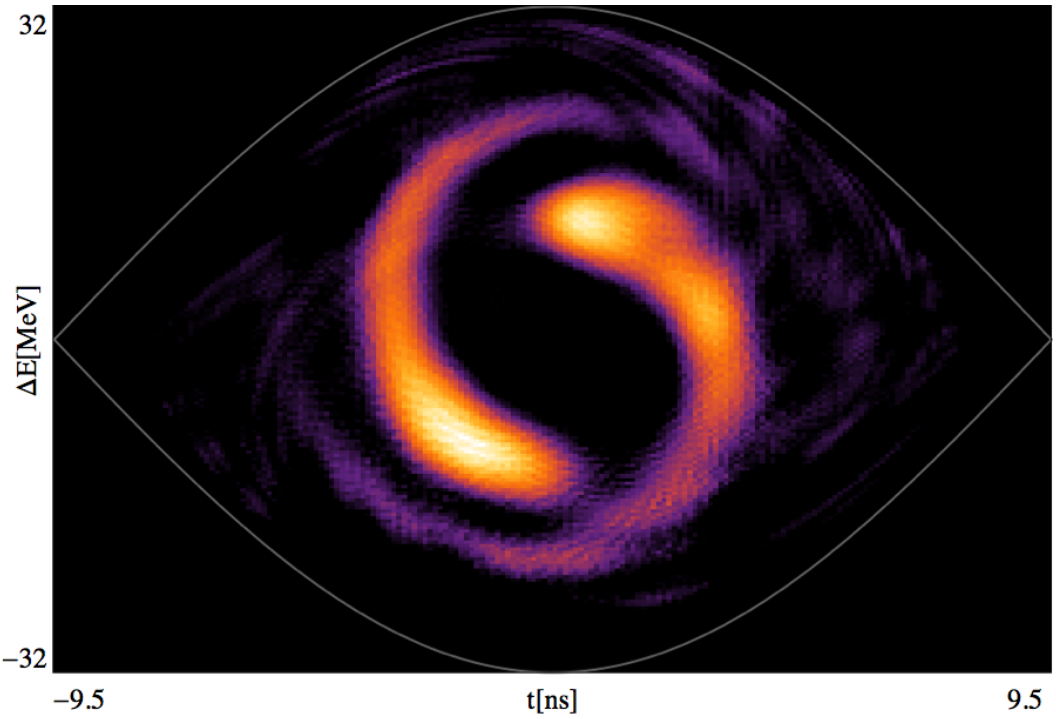

Figure 3.4: The 2-D time-ΔE, (time $= \omega_{rf}(\phi - \phi_s)$ phase space is depicted, showing a reconstruction of the distribution that produced the traces in panel fig. 3.3. The distribution is shown at the moment the first, red trace was taken. To produce this reconstruction from the traces also requires a model of the machine dynamics, provided by eqn. 2.8. The separatrix bounding the area of stable motion in phase space is indicated with a gray line. A reconstruction like this provides a view of the phase space that is easier to interpret than the raw traces for someone familiar with the standard treatment of longitudinal dynamics presented in ch. 2.

The actual distribution is located in the upper left of the image. You can see the question mark tumbling in the 2-D phase space, this is the distribution I wish to reconstruct from the projections immediately below the 2-D phase space. These projections represent the data that would be read out from the resistive wall current monitor, like the traces shown in fig. 3.3.

The iterative reconstruction process starts with a crude guess at the 2-D distribution. The guess may contain prior knowledge, but this isn't required. A uniform distribution, gaussian, or an empty phase space are all valid initial guesses. In fig. 3.5 the phase space has been divided into pixels each representing a small area of the phase space contained within the stable region, indicated by the separatrices that should by now be familiar from ch. 2. The only restriction placed on the guess at this point is that it must lie within the stable region, and must represent the same integrated charge as the data.

The initial guess is indicated in the top three left images. The guess is just the total charge divided evenly amongst the bins, some statistical noise is obvious in this image. The three frames in the top row represent the guess distribution tracked using eqn. 2.8. Immediately below the frames representing the guess at various angles are the projections of the guess at those angles. The projections are roughly parabolic because of the initial uniform guess, and the shape of the separatrix.

The next row shows the errors between the data frames and guess projections. Where the guess overestimates the data the errors are negative, and where the guess underestimates the data the errors are positive. The

errors are then back-projected, meaning each bin, or sample, in the 1-D traces is distributed evenly over all the pixels that are able to contribute to that bin. Notice in the image that the backprojection only distributes the bin over pixels lying within the separatrix.

The backprojected errors must then be reoriented to the same angle so that the reoriented errors can then be added to the initial guess to refine it. The same tracking equations used to propagate the guess can be used to reorient the backprojections, but now the equations must be used to track backward in time. Exactly how this is achieved will be discussed later in the chapter.

Once the guess has been corrected with the reoriented errors an iteration is complete. The refined guess can be taken as the answer, or used as the initial guess for the next iteration to create a better guess. Typically some measure of convergence is used to determine when iteration should halt. I use the 'discrepancy' discussed in [15] and used by Hancock et al.[2] The discrepancy is defined as the root mean square of the errors, it gives a measure of the agreement between the 1-D data, and the data that would be produced by the 2-D guess.

[2]The best measure of convergence is a complicated mathematical issue taken up in inverse problem theory, and is beyond the scope of this document. See [37].

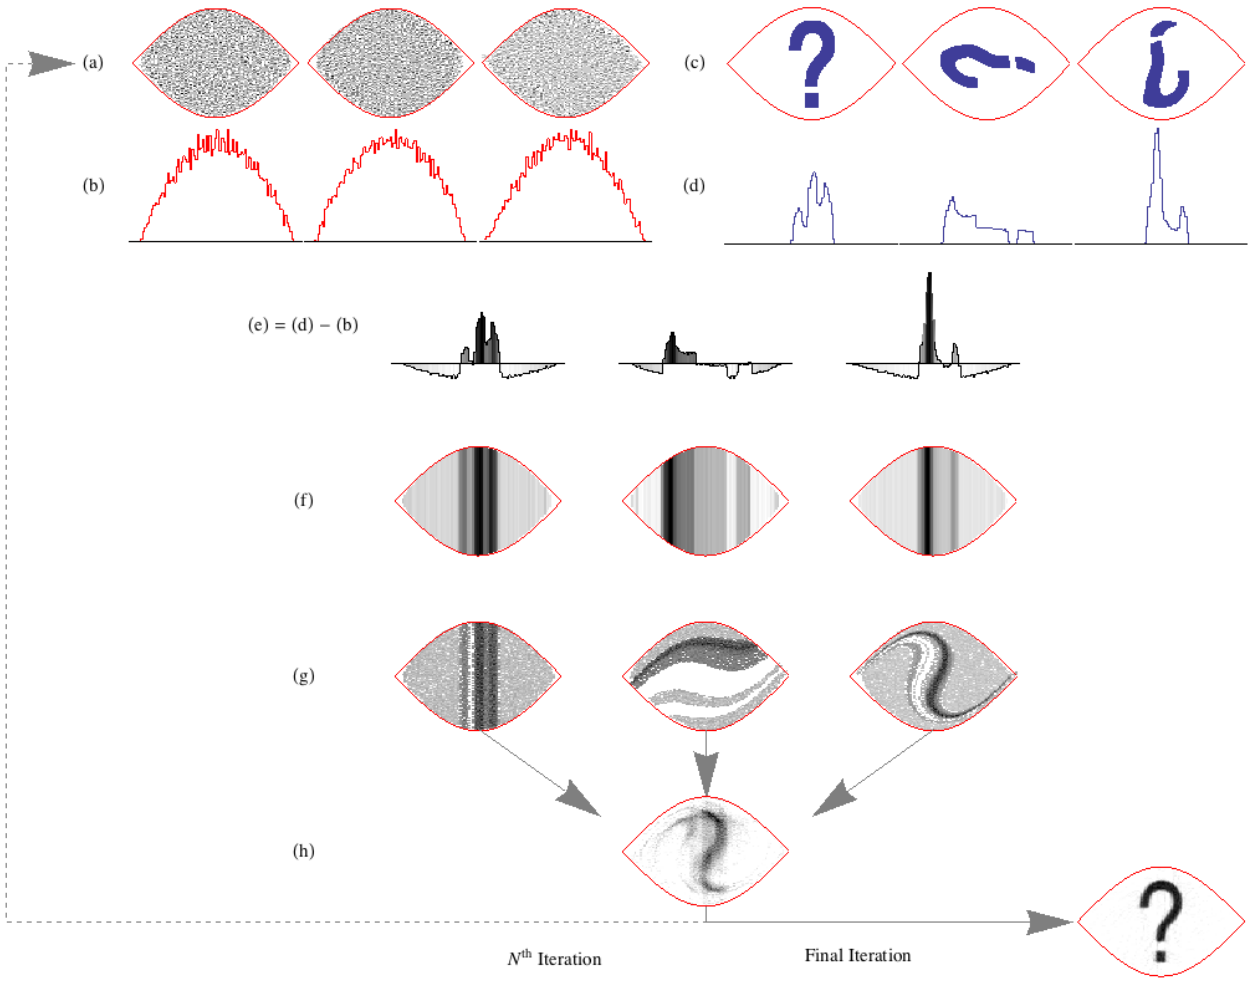

Figure 3.5

Figure 3.5: This image shows how an iterative algorithm is used to reconstruct a longitudinal phase space distribution at one moment in time. Row (c) is the actual distribution in ϕ - ΔE phase space, initially unknown, with the stable area enclosed in red. Each subsequent column in (c) shows the distribution tracked over a quarter of a phase space rotation. Row (d) is the projection of each frame onto the ϕ axis, this is the input data. Row (a) represents a guess at the object in row (c), also tracked. Initially this guess is very bad, here it is a uniform distribution. Row (b) is the projection of row (a) onto the ϕ axis. Row (e) is the error projection, or the difference between the data (d) and the projections of the initial guess (b). Because the projections represent a loss of information, we assume each projection was the result of a distribution with the same profile in ϕ, but evenly distributed over ΔE, this is called a backprojection. Row (f) represents the backprojection of the errors onto the phase space. In order to combine this information the backprojections must be reoriented, or tracked backward in time. (In the case of a rigid body, this would simply be a rotation by $-\theta$.) Row (g) is the sum of the elements in row (f), this is the first correction to the guess, and the end of the first iteration. From here row (g) is either fed back into row (a) and another iteration performed, or row (g) is taken as the answer. After many iterations the solution should converge on something similar to the original distribution.

The only remaining question is how the rotational operation, and it's inverse are carried out. Because we can't write the coordinate transformation of the phase space analytically, we rely on discretized approximations known as maps. There will be a separate map from the initial time $t = 0$ to each turn $t = n\tau_s$, where $n\tau_s$ are multiples of the revolution time of the synchronous particle. The process used to build part of one map is outlined in fig. 3.6.

The phase space in fig. 3.6 should look familiar, but it has been discretized into pixels. From each pixel, i_0 (where the subscript denotes turn number) some number of test particles are launched, shown in blue in the image. These particles are tracked using the difference equations from ch. 2. Because of the action dependent frequencies, the test particles, indicated in red, spread out. By tallying the number of particles launched from pixel i_0 that now occupy each pixel j_n, and normalizing by the number launched from i_0 one creates a map of probabilities that pixel i_0 contributes to pixel j_n. This is shown as the background color of the final pixels in the second frame of the figure. The third frame just removes the test particles, indicating pixel i_0 in blue. This allows the propagation of a weight in pixel i_0 to the n^{th} machine revolution. This takes care of the rotation, but to compare to data a guess must also be projected onto the ϕ axis, which is all the detector has access to. This is indicated by the histogram in the last frame of the figure. This operation must be carried out for every pixel in the image, and every turn in order to have a useable set of maps.

Similar maps must also be made in reverse from turn n to turn 0 by re-

versing the phase space rotation. These maps are used for the backprojection, and reorientation of the data, and error traces from fig. 3.5.

Once all the maps have been calculated, they can be used with simple matrix multiplication to transform a distribution between frames. The discretization of the phase space should not cover any more than a single bucket, since outside of this area the motion is no longer rotational, and does not provide useful information. It is possible attempt to reconstruct the beam outside of the separatrix, and the algorithm may still converge, but it is effectively meaningless. It is also possible to reconstruct more than one bucket at a time, but the maps become large, and repetitive since the dynamics within each bucket are identical. It is better to split the data into bucket sized frames and reconstruct each bucket separately, this also allows multiple buckets to be reconstructed concurrently.

The presentation so far has been heavily geometrical, which is useful for gaining intuition, but once the maps have been obtained tomographic reconstruction reduces to an inverse, linear algebra problem.

The forward problem is to determine the data that would be generated by a particular distribution given some equations of motion. We seek the inverse, that is given data and equations of motion, what distribution generated the data? We wish to solve eqn. 3.1 for the model, m.

$$d = G\ m \tag{3.1}$$

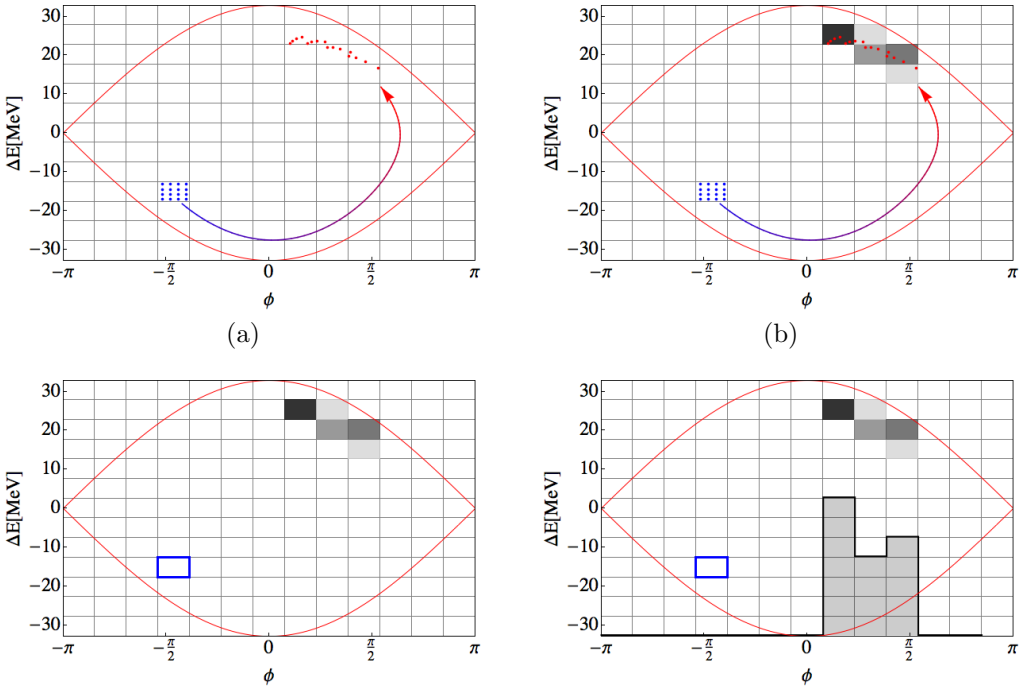

Figure 3.6: These images show the 2-D ϕ-ΔE phase space discretized into pixels. Panel (a) shows test particles, in blue, launched from a pixel and tracked over several turns. The final positions of the test particles are indicated in red. In panel (b) the red particles are binned, supplying weights for each pixel. In panel (c) the counting operation is complete. The rotation of the blue pixel over n turns has been discretized, and can be used to transform a discrete image of the phase space. In panel (d) the weights have been projected onto the ϕ axis. The inverse of the map must also be calculated in order to backproject the errors, as in fig. 3.5.This allows the guess in fig. 3.5, to be rotated in the phase space, and then projected so it can be compared to the data.

Where d is the data, m is the 2-D distribution and G represents the transformation that takes the distribution from $t = 0$ to the data. G comprises both the time evolution of the distribution and the projection operation. All of the physics of the system is contained in G.

The data will be discretized by the sampling procedure, and because of the map making, so will m, though this isn't strictly required by the theory of inverse problems. But the discretization of m, and d, do not need to share any particular relation as long as the maps, G have the appropriate dimensions. The answer can be of arbitrary dimension. [3]

The ART method discussed qualitatively earlier can be written as:

$$m_i^{q+1} = max\left[0, m_i^q + G_{i,j}^{-1}\left(d_j - G_{j,i}m_i^q\right)\right] \tag{3.2}$$

Where m is still the 2-D distribution, d, the 1-D data, and G, G^{-1} represent the forward and backward mapping/projection operations. The subscripts i,j indicate pixel, and data sample respectively, and the superscript q iteration number. The max operator ensures that no negative values creep into the distribution. There are many variants on this basic formula that allow negative values in the distribution during intermediate steps, or only update

[3] As with measuring convergence, the optimal discretization is a complicated problem without an established answer. Most sources suggest a heuristic approach to determining the discretization for a particular problem. e.g. Tarantola [37]. We undertook studies to determine a good discretization and found that the number of pixels in each direction should be at least the number of samples in a frame n_f, and that depending on the application a multiple of n_f can be used. The effectiveness of a particular discretization is tied to the number of test particles launched per pixel.

a pixel with an error value if the error is larger than some bound, etc. These variations are detailed in a tutorial on ART by R. Gordon [15]. After each iteration the distribution should be renormalized, or else the projections will quickly become larger than the data and the error term will grow without bound.

Aside from providing a roadmap for implementing a tomographic reconstruction algorithm, casting this problem in terms of inverse problems lets us say a few things about the solutions we obtain. From inverse problem theory we know that the problem is likely ill-posed, violating the unique solution condition. This is the same as saying the inverse problem is underdetermined. Because the system we are studying is physical, we assume there is a solution, though certain un-modeled effects can lead to data that is not consistent with any physical distribution. In practice, ART algorithms are extremely tolerant of this sort of error and will still converge, though they may begin to diverge if reconstruction is not halted. This is one reason to use a convergence criterion instead of a fixed number of iterations, the other being speed if convergence is reached quickly.

Additionally, in creating the maps we must have measurements for all of the parameters that appear in eqn. 2.8, the amplitude of the cavity voltage, the slip factor(η), harmonic number, and the synchronous energy, and phase. Treating tomography as an inverse problem, these parameters can be probed as well as the 2-D distribution [37]. The problem is then: given a set of data, what 2-D particle distribution *and machine parameters* are the most likely

solution?

Of all the numbers in the tracking equations, RF Voltage is the most difficult number to pin down precisely, so let us take this as an example. We can reconstruct the same data many times, scanning the parameter space of voltage near the expected value with each new reconstruction. The value of voltage that minimizes the discrepancy, meaning the data are reproduced with the smallest error, can be interpreted as a measurement of the voltage [18].

Any realization of a tomographic reconstruction algorithm will have to make many design choices. How many particles should be launched from each pixel? What is the optimal number of frames to use? How should the phase space be discretized? I've alluded to some of these questions throughout this section, for a discussion of some of these questions see appendix A.

3.2 Deriving Physics Quantities Using Tomography

The previous chapter introduced several errors, and quantities of interest to the evolution of a particle beam in longitudinal phase space. This section will show how a tomogram is used to extract the quantities of interest, and identify errors in machine operation. I will start with emittance and then move onto injection errors, beam halo, energy spread, and collective effects. I've also mentioned several times throughout this chapter that tomograms allow the experimenter to apply intuition to machine commissioning problems for which a dedicated analysis has not been developed, a short section will be devoted to an example of this.

3.2.1 Emittance

Emittance is an important measure of the beam density in phase space, defined in ch. 2. Standard methods of measuring the beam emittance are destructive, or involve measuring the time spread of the beam, and then making assumptions about the beam equilibrium to say something about the energy spread, which is normally difficult to measure [1]. With tomography, it is not necessary to make any assumptions about the beam distribution, since an arbitrary distribution can be reconstructed.

Once a 2-D beam distribution has been obtained calculating emittance is simple. Fig. 3.7 shows a view of a reconstruction with curves of constant action overlaid to demonstrate how emittance is measured. At each pixel the action is calculated, and binned according to the annuli defined by the constant action contours. Below the reconstruction is a plot of the distribution of particles as a function of action. The bins in the lower plot can be summed until the appropriate fraction of the beam has been reached, and the emittance read off of the horizontal axis.

3.2.2 Injection Errors

Using tomography it is easy to resolve injection errors like those depicted in figs. 2.8, and 2.10. Those were contrived examples for the sake of illustration, fig. 3.8 shows the reconstruction of an actual phase error the Recycler ring, which receives beam from the Booster. In this case, the first row is data acquired at injection and subsequent rows represent approximately 100ms

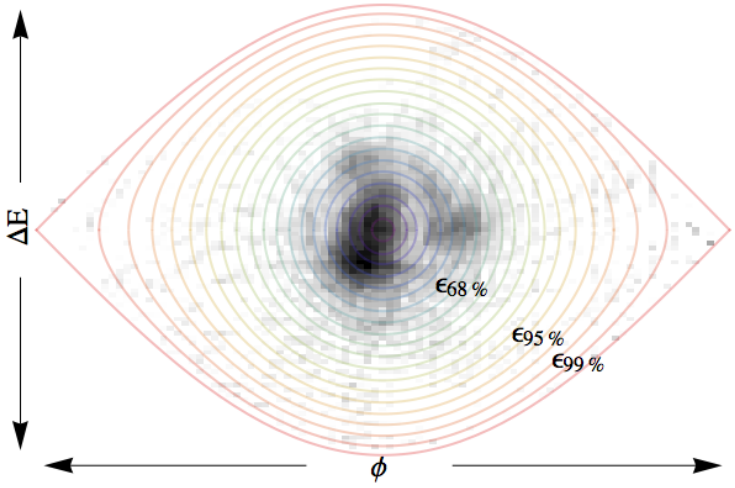

(a) Reconstruction

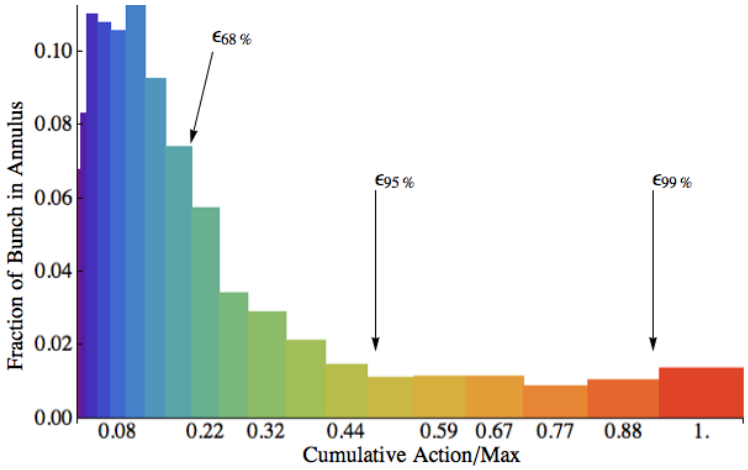

(b) Distribution

Figure 3.7: Plot (a) shows the ϕ-ΔE phase space reconstruction of a bunch in grayscale, with contours of constant action indicated in color. The colors of the contours correspond to the bars in plot (b), which represents the distribution function of the beam as a function of the action binned according to the annuli in the upper image. The horizontal axis is normalized to the maximum action in the image, multiplying by the bucket area gives the emittance in $eV \cdot sec$. The emittance is the area enclosing x%. Commonly used values of x are: 68%(rms), 95%, or 99%. These contours are indicated in both panels.

steps. The emittance growth from filamentation that occurs over many turns is obvious. The input data, reconstruction projections, and reconstruction of the phase space are all shown.

At injection the beam centroid is offset by about a sixth of a bucket width, $\approx$ 3ns. This dipole error is intentional. As I mentioned in ch. 2, collective effects depend on the charge density. In this case, collective effects were causing uncontrolled emittance blow-up, and unacceptable beam losses during Recycler commissioning. A small phase error was introduced to allow the beam to filament, increasing the spread of the beam in the ϕ coordinate, and reducing the space charge effects. Ultimately this created a lower final emittance, with a more controlled beam shape.

Fig. 3.9 shows the input data, reconstruction projections, and reconstructions of beam under the same operating conditions as depicted in fig. 3.8. In this image however, another source of error is contributing a dipole error in energy. Notice that the reconstruction in column (a) shows the bunch centroid low in the energy coordinate by about a third of the bucket height. This error was caused by a temperature fluctuation in the power supplies that maintain the extraction energy of the Booster beam. This additional error causes more filamentation than desired. Tomography allows this sort of error to be easily discovered and corrected.

Tomography is also able to resolve injection errors even after filamentation has occurred. Fig. 3.10 shows just such an error. In this example tomography was being used online in the Main Injector during commissioning.

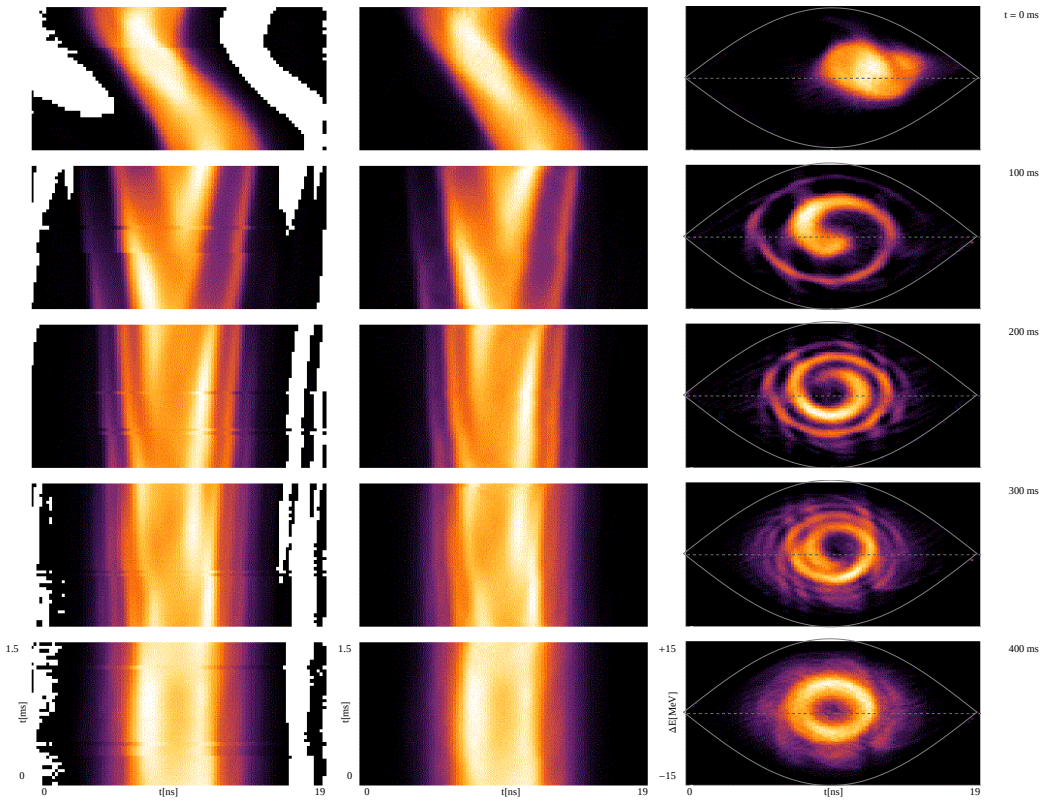

Figure 3.8: The three columns of this image represent the input data, the profiles created by the tomography algorithm, and the reconstructed 2-D phase space distribution of the sum of all bunches in a Booster batch in the Recycler. Starting from the top of the image each row of frames represents a step of about 100 ms. Each frame is from a different Booster batch under the same operating conditions. Notice the offset in phase in the first frame. This intentional phase error causes filamentation meant to mitigate collective effects, which were leading to uncontrolled emittance blow-up in the Recycler.

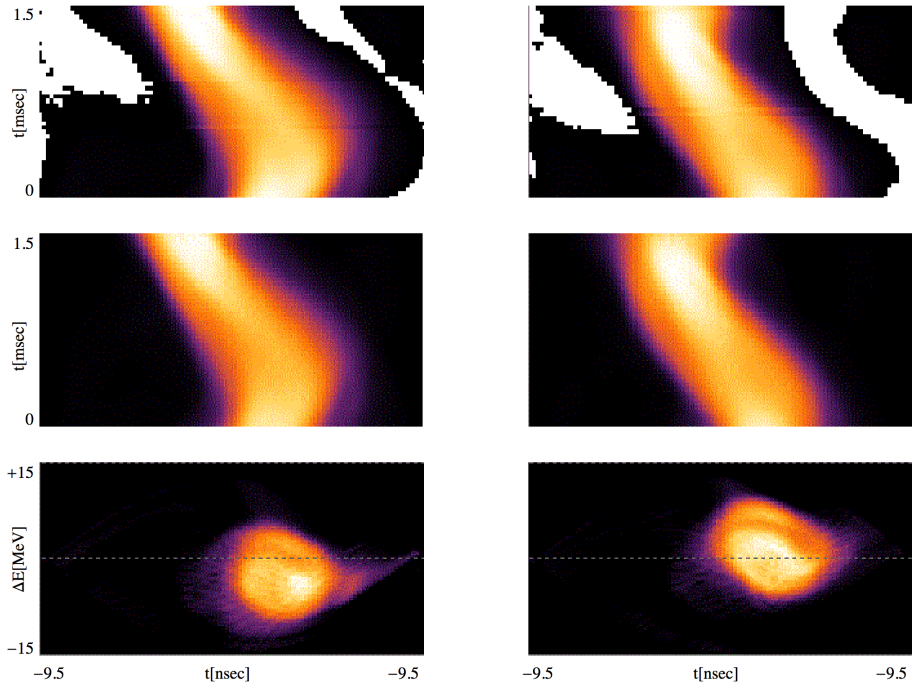

(a) Phase Offset Plus Energy Error (b) Energy Error Corrected

Figure 3.9: Column (a) shows the input signal, reconstructed profiles, and 2-D phase space reconstruction of the same intentional phase offset indicated in fig. 3.8 with an addition energy error caused by a temperature fluctuation in Booster magnet power supplies. The error in this case causes the injected bunch to have a lower energy than the design energy of the Recycler. Reconstructions make tuning easy, as it is obvious what the problem is at a glance. This allows the machine to run reliably in a configuration like the one shown here, where specific problems are alleviated with precise tuning. Column (b) shows the same data, but with the error corrected.

The beam was passing from the Booster, to the Recycler, and then to the Main Injector, a common mode of operation. The figure shows the resistive wall current monitor data in the upper right. The phase space reconstruction is in the lower left. Above and to the right of the reconstruction are the projections of the reconstruction onto the time, and energy axes.

Even though the bunch appears to be centered in the Main Injector according to the data from the resistive wall current monitor, and the projections, the tomographic reconstruction shows a spiral shape indicative of filamentation. Eventually this was tracked to a Booster setting that was not changed when operation switched from Booster-Main Injector to, Booster-Recycler-Main Injector. The beam was injected into the Recycler with an error, where it filamented before being transferred to the Main Injector. The problem in this case did not destroy the beam, but the emittance increase seen in the Main Injector could become problematic with increasing intensity, and would be much more difficult to track down without the 2-D reconstruction.

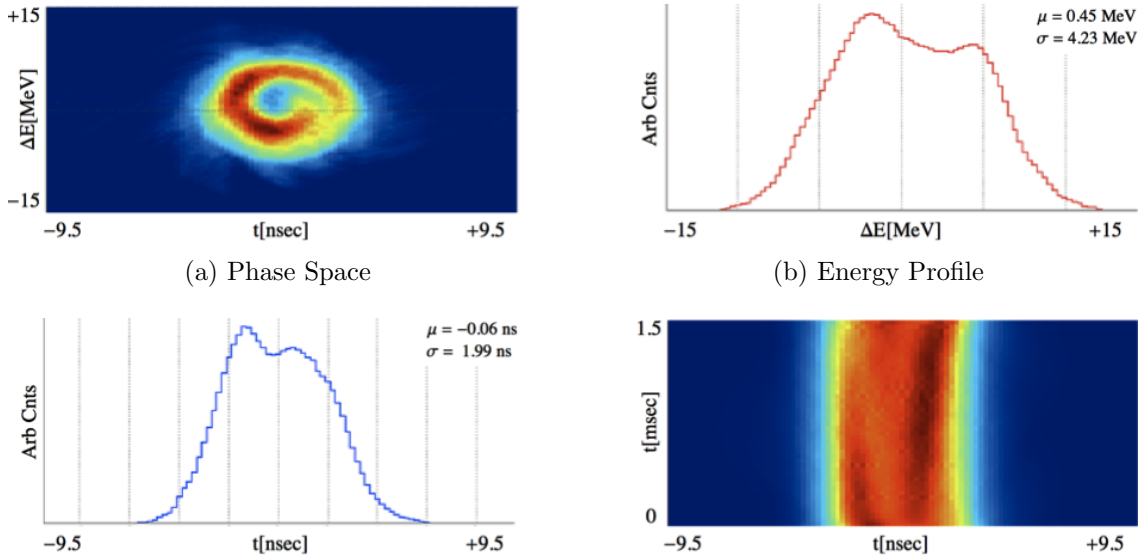

(a) Phase Space

(b) Energy Profile

(c) Time Profile

(d) Data

Figure 3.10: The presence of a dipole error in the Recycler is seen at the moment of injection to the Main Injector. (a) shows the 2-D phase space reconstruction of the data in (d). (b), and (c) to the show the projections of (a) onto the energy and time axes respectively. The hollow shape in (a) is a signature of the dipole error in the Recycler that has filemented before being transferred to the Main Injector.

Typically when beam is transferred from the Booster to the Main Injector or Recycler, the buckets will not be matched, as in panel (a) of fig. 3.11. This quadrupole error, like the simulated error depicted in fig. 2.10, is corrected via bunch rotation. Bunch rotation is the term used to describe a phase space manipulation that corrects the quadrupole mismatch of the upright distribution from panel (a) of fig. 3.11 by lowering the RF cavity voltage in the Booster for a short period in order extract beam that more closely matches the contours of the Main Injector. There are alternate schemes to acheive a similar result [41], [2]. Tomography could be useful in examining the effect of these schemes on beam quality.

3.2.3 Energy Spread

The plots shown in fig. 3.12 demonstrate how tomography provides energy spread measurements. The image depicts several plots of beam that has been slip-stacked, which is a particular method of increasing beam intensity used in Fermilab's Recycler. Slip-stacking involves 'stacking' beam in the energy coordinate of longitudinal phase space out of phase. Because of the slip factor, η, defined in eqn. 2.4, and the energy difference, the two bunches eventually slip into phase. At this point they are recaptured. The two bunches with an energy spread can be seen in the reconstruction in the lower left of fig. 3.12 separated in energy. The resistive wall current monitor data is shown in the upper right. The projections of the reconstruction onto the ϕ and ΔE axes are shown histograms above and to the right of the reconstruction.

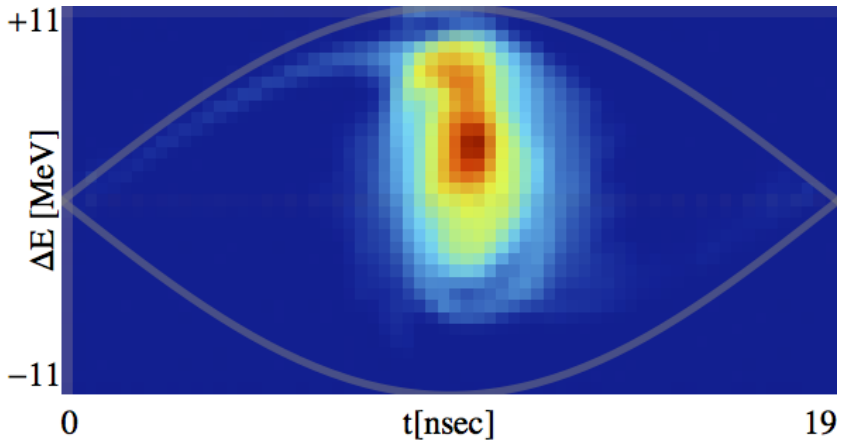

(a) Quadrupole Mismatch

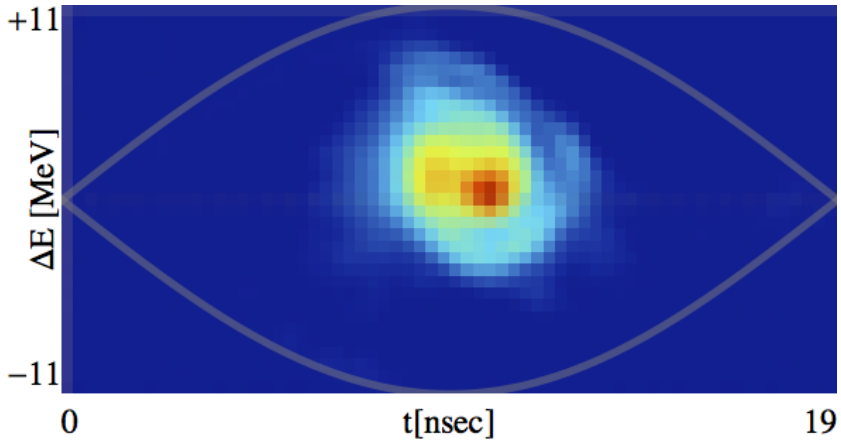

(b) Mismatch Corrected

Figure 3.11: Panel (a) show the 2-D reconstruction of ϕ-ΔE phase space for a bunch subject to a quadrupole mismatch upon transfer from the FNAL Booster into the Main Injector, and panel (b) the correction. The signature of quadrupole mismatch in an image like this is a distorted aspect ratio relative to the shape of RF bucket contours. After many revolutions, the mismatched beam in panel (a) would filament to fill nearly the entire stable bucket. Because the reconstructions depict bunches at injection, the filamentation has not yet occurred.

Energy spread is typically a difficult measurement to make, similar assumptions must be made as in the measurement of emittance. That is one must assume that the bunch is in equilibrium, or gaussian in order to calculate an energy spread from a time spread. A large energy spread can lead to losses if the spread exceeds the momentum aperture, it may signal a problem somewhere in the accelerator chain. Energy spread can also be problematic during slip stacking. If the tails of either bunch exceed the limits of the larger recapture bucket, losses can occur during acceleration. In fig. 3.12, not only is it easy to extract the energy separation of the two slip-stacked bunches, but also the energy spread of each bunch separately.

3.2.4 Beam Halo

Beam halo is diffuse beam that accumulates, for whatever reason, around the primary beam core. Fig. 3.12 also exhibits some beam halo. Notice in the reconstruction, that both the primary distributions of charge have a small accumulation of charge to their left that appears to separated from the main bunches. This can be seen as a shoulder in the upper left plot of the time projection. The exact shape of this halo is difficult to determine from projections alone, but is obvious in the reconstruction.

In this particular case, the halo is likely due to the Booster bunch rotation described in the last section, but this is not certain. Halo can be due to RF noise, intra-beam scattering, or power supply ripple among other things. Halo can cause small, but noticeable losses if it escapes the region of stable motion

in the phase space. In some experiments, the time structure of the beam can be relevant, halo contributes to the tails in both time and energy, oscillating between the two due to synchrotron motion. If the detector can resolve halo, profiles are enough to the prove existence of halo, but only tomography will provide detailed distribution of halo, which may be operationally relevant, or may provide insight into the cause. The existence of artifacts in reconstructions, typically near the edges of a distribution, means that one should be careful when assessing halo. Multiple reconstructions can be done of the same beam pulse changing reconstruction parameters, or using frames slightly advanced in time. Generally artifacts will not persist the same way actual beam will. This is a well known technique [14].

3.2.5 Collective Effects

Often the particles are treated as non-interacting, point-like particles that simply follow the constant action contours derived in ch. 2. The charge is only relevant as it couples particles to the electric field provided by the RF cavity. Collective effects are all other interactions that a particle bunch experiences as a result of the charge of the particles. This includes self-forces, the interaction of bunch particles with one another, and interactions with the beam pipe, and beam line elements. These interactions can cause instabilities in the beam that blow-up emittance, or in extreme cases cause complete loss of the beam. Fig. 3.13 shows the evolution of a train of bunches subject to collective effects in the Main Injector. Each row of the image represents a train

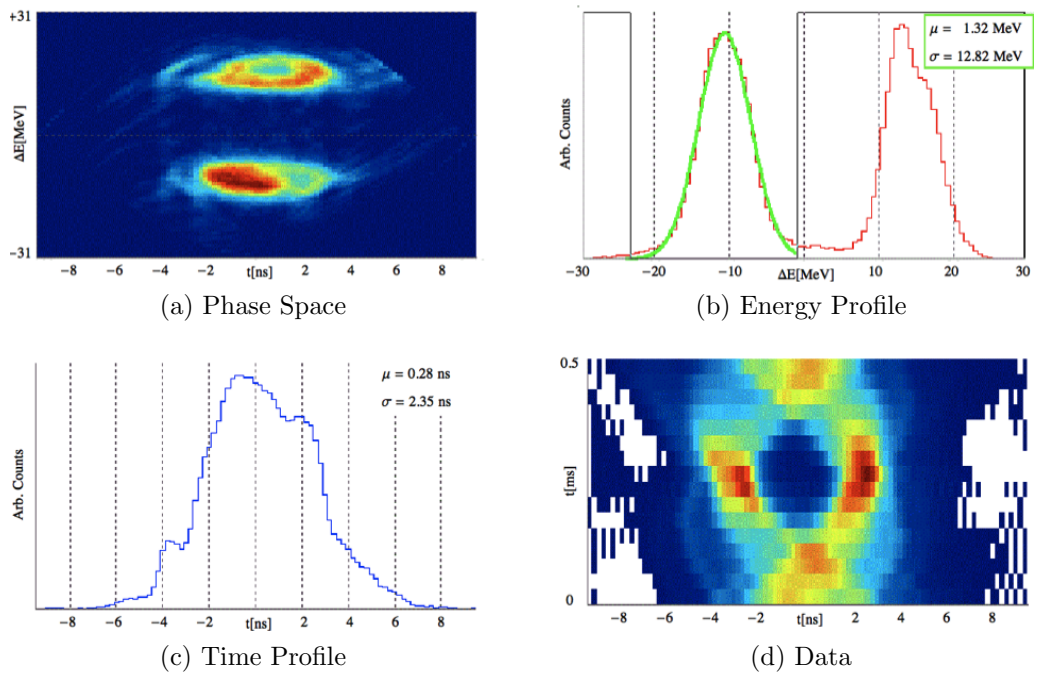

(a) Phase Space

(b) Energy Profile

(c) Time Profile

(d) Data

Figure 3.12: Reconstruction of slip-stacked beam demonstrating energy spread measurements, and beam halo. (a) shows the 2-D phase space reconstruction of the data in (d), the detailed structure of the slip-stacked beam can be seen. (b) to the right shows the projection onto the energy axis, which provides a measure of the energy spread of each beam separately, and the energy separation between beams. (c) shows the time profile of the beam at the moment of the reconstruction.

of ten bunches in the Main Injector the rows are taken at 89ms intervals. From the very first interval it is obvious that the bunches are not simply tumbling in the phase space, as predicted by the tracking equations from ch. 2 as there is some evolution of the bunch shape. As time advances the emittance blows up noticeably.

The exact cause of this instability is not known, but it was observed while the feedback system that is meant to control such longitudinal instabilities was active. The phase space structure of the beam may provide future experimenters some hints as to what is going on.

Collective effects are not included in the tracking equations we used to build the maps, so the questions naturally arise: can bunches subject to effects not included in the tracking be reconstructed? how do those effects manifest in reconstructions?

There are two cases to consider regarding whether a bunch subject to collective effects can be reconstructed. If the collective effect is due to some particular set of conditions that can be modified, or that disappear after some point in the cycle, it may be possible to see the effect of the interaction on the beam after the interaction is no longer important.

If the reconstruction is done while the interaction is important, there may still be some hope of obtaining reconstructions. If the growth rate of the instability present is shorter than the one half synchrotron period needed to reconstruct the longitudinal phase space, it is likely that the effect cannot be

reconstructed with any accuracy. In order determine if the image produced by the tomography algorithm is reliable the following procedure can be used. First, more frames of data should be taken than needed, at least a full synchrotron period. The data can be reconstructed at a frame near one half a synchrotron period. This reconstruction can be done with either the frames preceding the frame to be reconstructed, or the frames after the frame to be reconstructed. If there are dynamics present changing on a timescale faster than half a synchrotron period, there should be a significant difference between the two reconstructions of the same frame.

One type of collective effect, called coupled bunch modes, are due to wakefields, the electric fields excited by passing bunches that then decay away but may influence trailing bunches. The fields can excite modes in the machine resonant with the bunch spacing, and the synchrotron motion. This leads to the excitation of azimuthal modes which impose an n-fold symmetry on the beam distribution, where n is the order of the instability. Because of this symmetry, the projections of the beam will show a structure that repeats n times per synchrotron period. This repetitive structure looks like a symmetry to the reconstruction algorithm. Because the effect is not included, it is likely that measurements taken from tomograms influenced by collective effects will have certain errors, though qualitatively the instability is obvious. These can either be investigated offline in more detail including space charge in the tracking, or with more application specific tools. There are many fantastic references discussing collective effects in more detail, see [9], and [33].

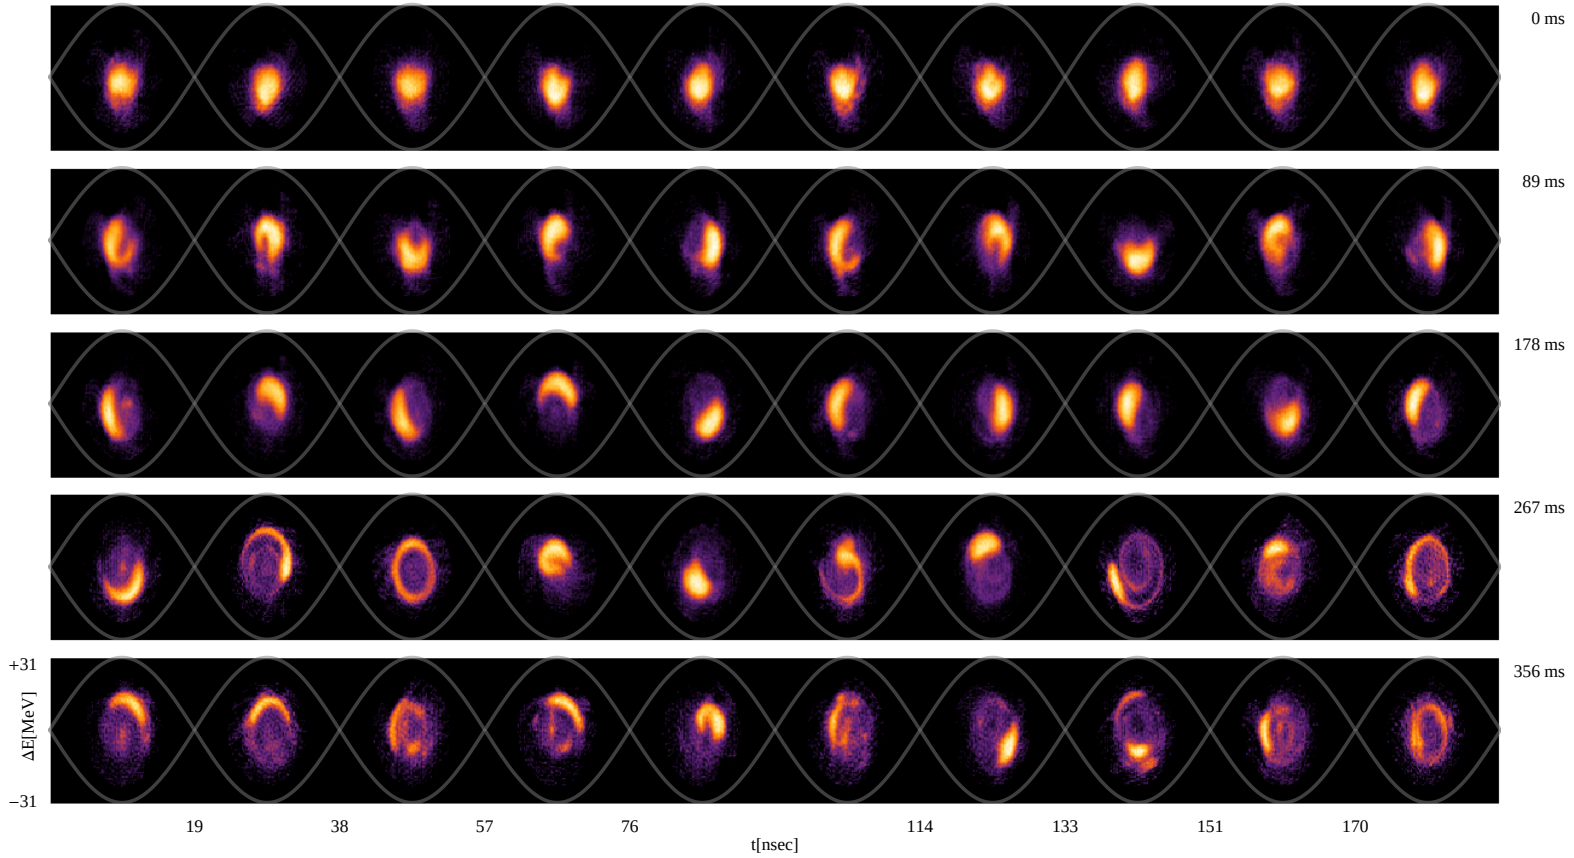
0 ms
89 ms
178 ms
267 ms
356 ms
+31
ΔE[MeV]
−31
19
38
57
76
114
133
151
170
t[nsec]

Figure 3.13

Figure 3.13: Each row shows the ϕ-ΔE phase space reconstruction for the same 10 bunches in the Main Injector at a different time, indicated on the right. The axes are labelled on the bottom row since all images have the same range. As time progresses the bunches distort, and emittance grows because of an instability in the Main Injector that is controlled with a damper system. Without the damper the beam would be lost completely. The limits of stability for each bucket are drawn in gray. The dynamics driving this instability are not included in the tracking equations, but the phase space can still be reconsrtucted if the instability grows on a timescale longer than $\approx$ half a synchrotron period, the time required to make a full rotation in the phase space.

Other times the collective effect may be present, but not have such a clearly defined signature. Fig. 3.14 shows five buckets of beam in the Main Injector. These buckets were created when beam that was bunched by a 53MHz RF system (rougly 20ns bucket width), was coalesced with a 2.5MHz RF system (roughly 400 ns bucket width). The first reconstructed bunch shows the expected behavior. The bunch is diffuse and evenly distributed through the portion of the bucket that it occupies. Some striations left over from the original beam structure can be seen, but these appear to be perturbations to the main distribution. The bunches trailing the first bunch become increasingly distorted, showing large spiral arms. The spiral shape evident in the trailing bunches indicates that a wakefield from the leading bunches is probably to blame since the effect gets worse along the bunch train. Additionally, beam is only supposed to occupy four buckets here. The instability has caused beam to leak into the fifth bucket trailing the bunch train. If the fifth bucket were assumed empty this beam could be kicked out of the machine by extraction kickers and represent serious losses.

Tomography allows experimenters to visualize effects like this without first developing formal metrics to quantify things like the spiraling shown. It indicates that there is a problem, and as in the case presented in fig. **??**, allows experimenters to quickly test solutions and decide on effectiveness without lengthy data analysis normally required of longitudinal problems.

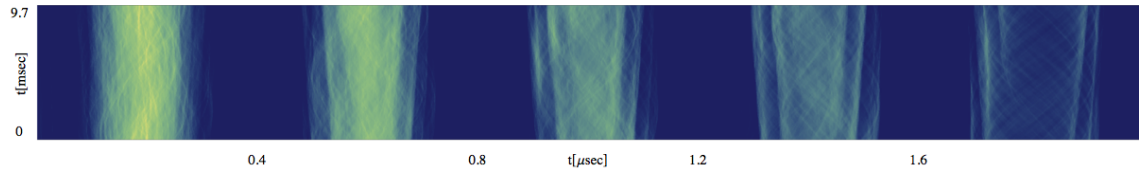

(a) Data

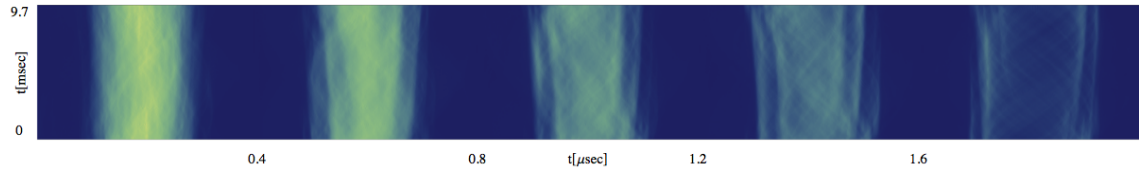

(b) Reconstruction Projections

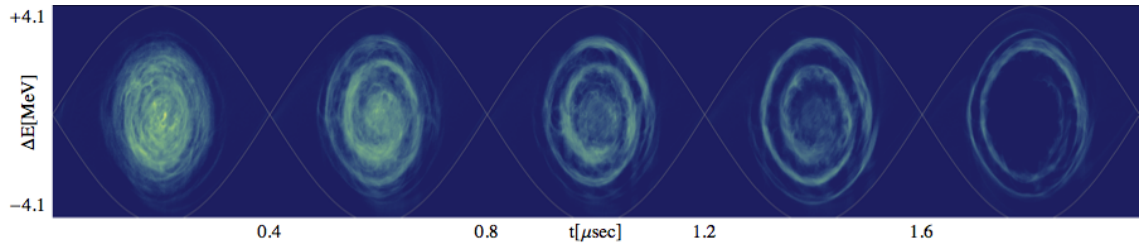

(c) 2-D Reconstruction

Figure 3.14

Figure 3.14: Row (a) shows traces from the wall current monitor of five bunches in the Main Injector. The horizontal axis is position around the ring, and the vertical axis represents discrete machine revolutions. Row (b) shows the same information but as the projections of the distribution found through tomographic reconstruction. Notice that some of the detail is lost. Row (c) represents the longitudinal phase space reconstructions of row (a). This data was taken during a study of beam bunching for the planned g-2 experiment. Only 4 of the buckets should be occupied, beam loading, or RF phase error created beam loss leading to particles in the 5th bucket. This represents decreased efficiency of the accelerator complex as some beam intended for the experiment will not be utilized. This could also be problematic with respect to losses if kicker timing assumes an empty fifth bucket. Some wakefield effects are obvious in the spiraling pattern in the last 4 bunches. Notice how evenly distributed the beam occupying the first bunch is. Each bunch should contain roughly the same number of particles and a similar distribution. An error such as this would be difficult to detect without tomography.

3.3 Some Comments on Implementation

The tomography algorithm discussed here was implemented in a program called Tomography and Related Diagnostics in Synchrotrons, (TARDIS). All of the images shown were produced using TARDIS, either in offline mode, or online during machine operation.

TARDIS allows for real-time, multi-bunch reconstruction of longitudinal phase space using the ART algorithm; several variants are implemented. TARDIS handles all pre-processing, plotting, and post-processing analysis. It allows data to be saved for further analysis offline. The user interface is designed to be simple to use and interpret for anyone with a basic knowledge of longitudinal beam dynamics.

In addition to the simple, almost single button, operation mode that defines most of the day-to-day use of TARDIS at Fermilab, experts can interface with almost every aspect of the program to undertake novel studies. All configurations can be saved so the expert can set the system for a particular use case and recall all the relevant parameters in the future in order to provide easy access to useful measurements.

3.4 Summary

In this chapter I've shown how a standard 1-D diagnostic, and knowledge of longitudinal dynamics can be combined to recreate a 2-D image of a charge distribution via a tomographic reconstruction algorithm.

I also outlined some of the physics quantities that can be measured using the resultant image, and some of the common machine errors that can be easily spotted.

With only the 1-D information a dedicated analysis would have to be done searching for each error condition. The 2-D analysis allows the experimenter to focus analysis on only the features that seem worthwhile, and in some cases diagnose problems that would slip through a standard analysis. Certain assumptions can be relaxed with respect to similar measurements by other techniques, so the experimenter needs less facility with this particular method to interpret those measurements.

In addition to measures of standard quantities like emittance, or injection mismatch, bunch structure can be investigated in more detail. Some beam manipulations are sensitive to the detailed momentum distribution, e.g. slow-extraction to switch yard at Fermilab. Consistent features can be seen and their causes diagnosed significantly speeding up troubleshooting, or more in depth studies.

In the following chapters I will introduce the Fermi National Accelerator Laboratory complex, specific operational concerns, and an experiment to measure transition growth carried out with the newly commissioned tomography system.

Chapter 4

Design of an Experiment to Measure Emittance Growth Across Transition

Proposed experiments at the intensity frontier, such as those at Fermilab, will require more protons on target with more stringent timing and energy spread requirements in order to precisely measure rare phenomena. In addition to the beam quality specifications of particular experiments, beam must meet certain requirements to operate the accelerators safely. At higher intensities beam losses must be lower to reduce component damage, environmental irradiation, and maintain a safe environment for workers. These may not seem like scientific concerns, but they all affect beam operations, and down time, which ultimately impact the experiments carried out with particle beams.

Problems with the longitudinal dynamics of the beam can lead to losses through instabilities, injection errors, and other unforseen causes of emittance growth. The ability to quickly isolate and solve problems is crucial. Problems arising from intensity dependent effects cannot be induced below a certain threshold and so present a particular challenge in terms of finding solutions during operation.

In this chapter I will present the design of an experiment to probe

longitudinal emittance growth across transition. Because the measurements were performed at Fermilab, I will discuss the Fermilab accelerator complex, and a particular method for doubling beam intensity used at Fermilab called slip-stacking. This study shows that in the simple case of single batch beam the analytical growth predicitons given in ch. 2 are consistent with the observed growth. A fit to the the growth determines the longitudinal broadband impedance, and first non-linear momentum compaction. The numbers found are consistent with the growth observed for slip-stacked beam, but there is additional growth not accounted for by this model.

I will discuss the sources of error, and the precision attainable when measuring emittance growth with tomography. Ch. 5 will present the results of this experiment, and implications for future operation.

4.1 Fermilab Accelerator Complex

The machines in the Fermilab complex are detailed in table 1.1, the physical layout of the site is shown in fig. 1.5. A brief overview of the whole facility was given in ch. 1. Here I will focus on the Main Injector, and Recycler.

The Main Injector and Recycler are housed in the same tunnel and each have a circumference 7 times the Booster circumference, with the same 53MHz fundamental RF frequency, and thus the same bucket spacing of $\approx$ 19ns. [1]

[1]There is actually a slight difference in the operating radii of the Recycler and Main Injector, which means the Booster energy needs to be adjusted to inject into one machine or the other.

Both machines can accommodate 6 Booster batches in series. Space must be left to allow injection/extraction kickers time to ramp up and kick the beam.

The Recycler is a permanent magnet machine, which means the bending field cannot be changed so it is only used for beam storage and RF manipulations at 8GeV, not acceleration. In the future the Recycler RF will be modified to help prepare beam for the planned muon experiments which will use a 2.5MHz RF system in conjunction with the Accumulator, and Debuncher rings to deliver beam to precision muon experiments.

The Booster represents an intensity limitation at Fermilab. Because space-charge effects are worse at lower energies, only so much beam can be injected into the Booster at 400 MeV before the beam becomes unstable. To increase the intensity of beam delivered to neutrino experiments the Recycler performs a process called slip-stacking [8], [4]. Figs. 4.1 shows the entire process from above a synchrotron, in real-space, and fig. 4.2 shows the same process in longitudinal phase space, but only depicts a few bunches.

First, 6 Booster injected in series around the Recycler ring. These can be diverted from the central orbit by slightly decreasing the beam energy. At the same bending strength the radius will also decrease. Then another 6 batches are injected on the design orbit. A separate RF system is used for each set of 6 batches, operating with two slightly different frequencies. Because of the slightly different energies, the two collections of batches will accrue phase difference. Once the two beams are in phase a much larger RF voltage can be applied, with a synchronous particle on the design orbit that lies between

the two slip-stacked batches. This RF system now captures two bunches per bucket. The phase space reconstruction of a bunch shortly after slip-stacking is shown in fig. 3.4. The energy separated bunches are obvious above and below the synchronous point at $(\phi = 0, \Delta E = 0)$. Slip-stacked beam is still at 8 GeV and is extracted to the Main Injector for acceleration.

The Main Injector uses electromagnets, and accelerates 8 GeV beam to 120 GeV for use in neutrino experiments, or at the fixed target test area, known as switch yard. The Main Injector has a transition gamma of $\gamma_t = 21.6$, which means it must pass through transition to accelerate from 8 GeV injection, to 120 GeV extraction energy. Crossing transition necessarily involves emittance growth when dealing with intense beams. Though transition is well understood in general, it has not been studied in detail experimentally in the Main Injector.

The Accumulator and Debuncher share a tunnel, and are both currently unused. In the future they may be used to prepare beam for precision muon experiments. Historically these machines were used to accumulate and prepare anti-protons for use in the Tevatron using a process similar to slip-stacking.

Currently the tomography system TARDIS is only available in the Main Injector and Recycler, which means we can study beam at any point during the Recycler or Main Injector cycles, from injection to extraction. The Booster can also be probed by examining beam in either the Recycler or Main Injector immediately upon injection. Longitudinal phase space tomography can be used in all of the synchrotrons at Fermilab, but the Accumulator and Debuncher are not currently in operation, and the Booster lacks the triggering hardware

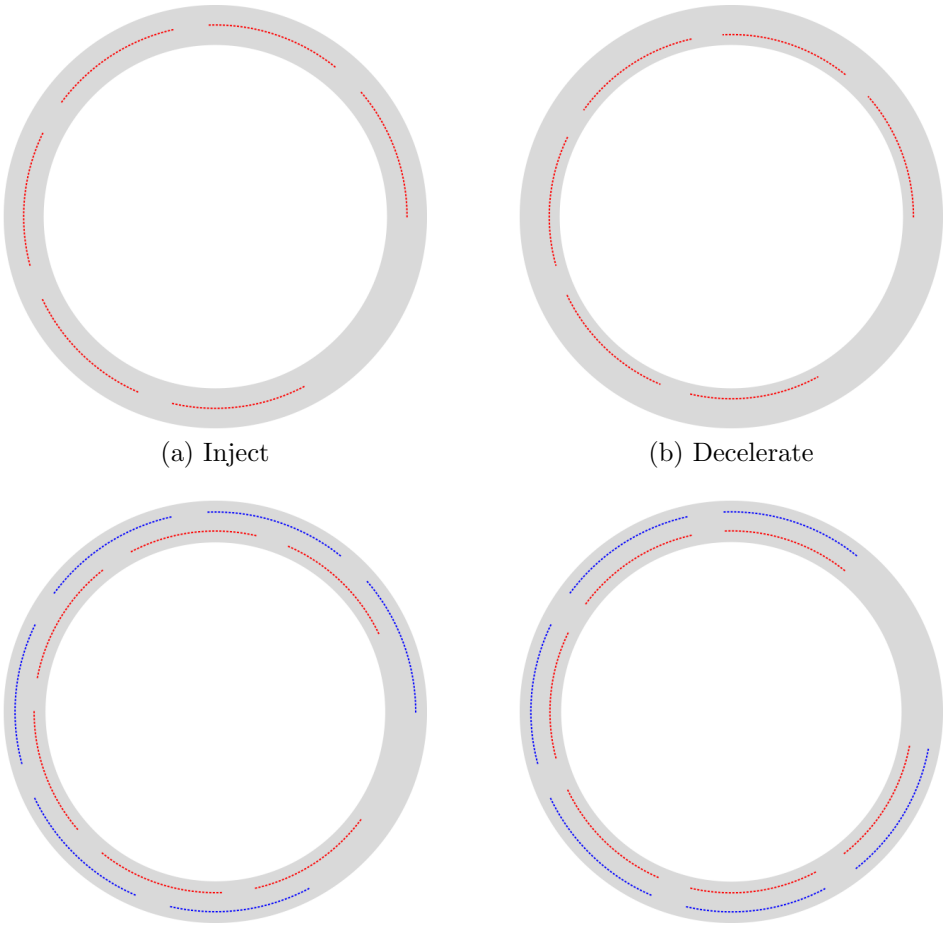

(a) Inject (b) Decelerate

(c) Slip (d) Recapture

Figure 4.1: A schematic view of slip-stacking in the Main Injector or Recycler to double beam intensity. (a) Six batches of Booster beam, each 84 bunches, are injected into the machine at the design energy. (b) The first six batches are accelerated to an energy slightly lower than the design energy, decreasing the orbit radius. (c) Six more batches are injected, but with an energy slightly higher than the design energy, thus a larger orbit radius. The two fills have slightly different revolution frequencies. (d) The two fills circulate long enough for the batches to slip in phase until they are azimuthally aligned. A 1MV RF is then turned on to capture two bunches per bucket. Fig. 3.4 shows the same process in longitudinal phase space.

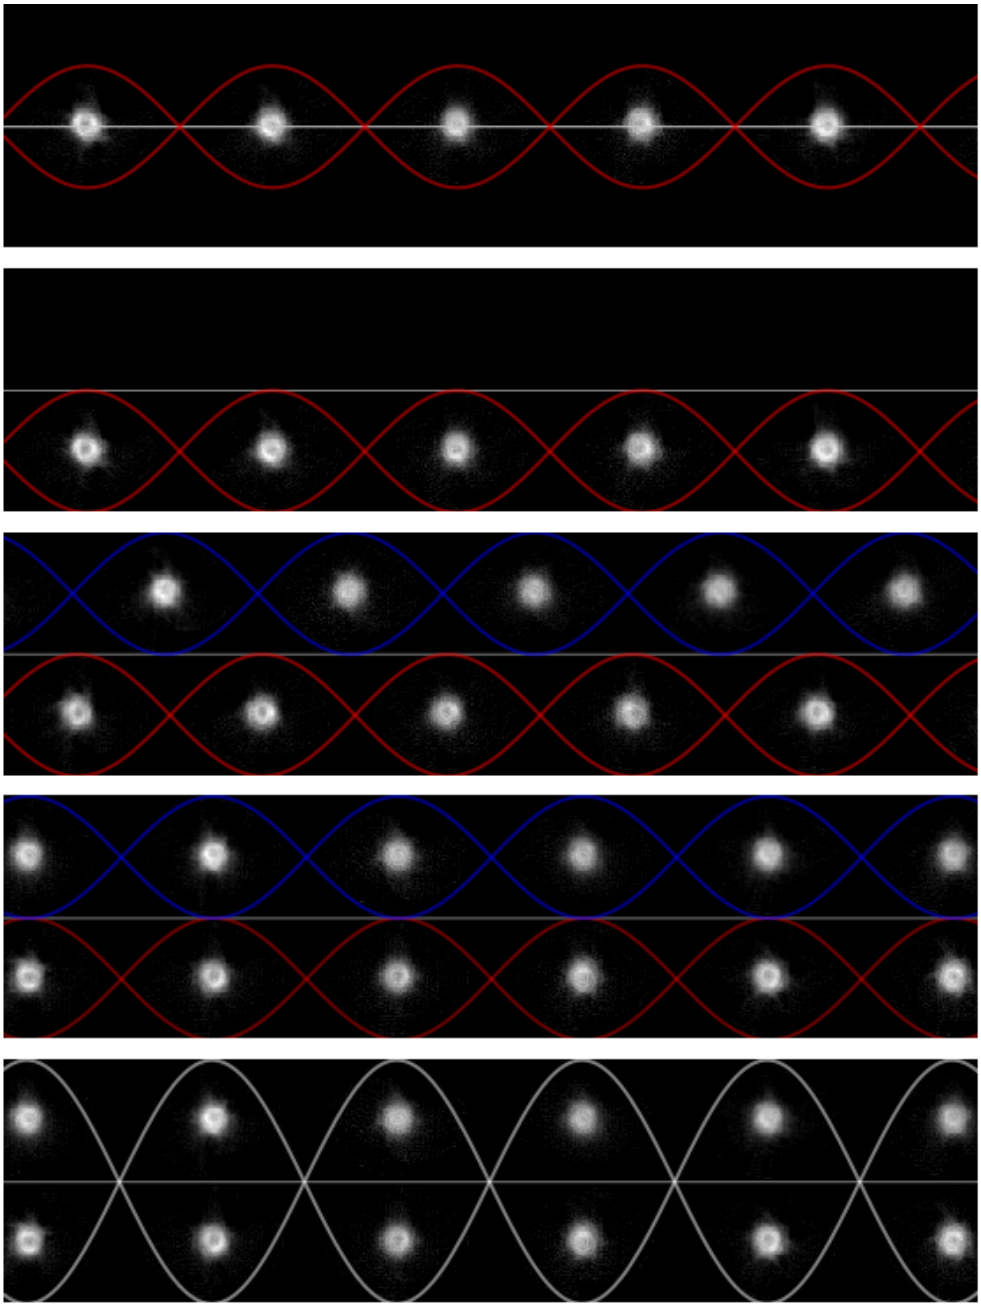

Figure 4.2

Figure 4.2: Longitudinal phase space of two Booster batches during slip-stacking, corresponding to the panels shown in 4.1. The top image shows the first injection on the central orbit. Five batches are injected in series on this orbit, and then decelerated to a smaller orbit, in the second panel. The third panel shows the second set of injections at an energy higher than the central orbit. The two sets of injections are initially out of phase, but they slip relative to one another. When the batches are aligned a stronger RF voltage is applied, such that two bunches are confined to a single bucket.

necessary to commission TARDIS.

4.2 Emittance Growth at Transition Crossing

Currently losses are acceptable for operation in the Main Injector, but future upgrades will require a better understanding of loss mechanisms in order to keep running at higher intensities. Losses at lower energies, before acceleration, are always preferred since the same number of protons lost result in lower deposited energy and lower activation of the tunnel and beamline elements, affecting both machine reliablity, and safety. For this reason emittance growth induced during the acceleration cycle needs to be well understood.

Ch. 2 discussed longitudinal emittance growth crossing transition, defined as: $(\epsilon_A - \epsilon_B)/\epsilon_B$. Where ϵ_B, and ϵ_A are the emittance before, and after transition. Though emittance growth is inevitable for high intensity beam crossing transition, growth does not necessarily mean beam loss. An understanding of the driving contribution to emittance growth is critical in determining what steps will minimize emittance growth. For instance, growth driven by an instability may require more powerful hardware for the longitudinal damper system, growth driven by non-linear terms of the slip-factor may require some modification of the magnets used to focus the beam transversely that determine the slip-factor. Any growth would be ameliorated by a higher rate of acceleration crossing transition, but this may not be the most cost-effective, or simplest solution. It could be that some minor operational considerations previously unimportant could improve performance.

The study presented in ch. 5 uses the recently installed tomography system to look at emittance growth crossing transition as a function of input emittance. The expressions for emittance growth as a function of emittance written explicitly in eqns. 2.38, 2.40, and 2.41 are combined in the single eqn. 4.1 used to fit the observed emittance growth. The three fit parameters unique to each growth mechanism are: $\alpha_1, |\mathrm{Z}_{||}/n|$, and R. Each is explained in detail in ch. 2.

$$\frac{\Delta\epsilon}{\epsilon}(\epsilon) = \frac{\Delta\epsilon}{\epsilon}(\epsilon;\alpha_1)_{SingleParticle} + \frac{\Delta\epsilon}{\epsilon}(\epsilon;|\mathrm{Z}_{||}/n|)_{Reactive} + \frac{\Delta\epsilon}{\epsilon}(\epsilon;R)_{Resistive} \quad (4.1)$$

By measuring the emittance on either side of transition, I am able to calculate the emittnace growth without reconstructing the beam very near to transition, where tomography becomes increasingly difficult due to the lack of phase space rotation, rapidly changing machine parameters, and increasingly distruptive collective effects. Under normal circumstances emittance should not grow once transition has been crossed, as long as the measurements are taken far enough from the time where non-linear, non-adiabatic dynamics dominate. A meausrement of emittance growth should indicate the dominant source of emittance growth when crossing transition.

Tomographic reconstruction is particularly well-suited to this problem because it is not necessary to wait for bunches to equilibrate, or filament, before making a measurement. This allows the measurement to be made very close to transition in an attempt to minimize the effect of any other emittance

growth that may be present during the acceleration cycle due to noise in the RF phase or voltage, or magnet power supplies, etc.

Shown in fig. 4.3 before, and after transition crossing, single batch beam typically satifies all the assumptions used to derive the growth equations in sec. 2.2.2. The bunches are mostly confined to the linear region, and approximately gaussian. Because of this the single batch beam is used to measure the fit parameters in eqn. 4.1.

Slip-stacked beam fills more of the bucket than single batch beam, and is initially concentrated in two clumps. Fig. 3.4 shows slip-stacked beam near the moment of recapture. The two distinct clumps are dispersed somewhat during acceleration but do not filament completely by the time the beam crosses transition, resembling fig. 4.4 before and after transition. The growth of these bunches across transition is compared to the prediction based on the values measured in the slip-stacked case.

There is no machine knob for emittance, though there are techniques to increase emittance in a controlled way, such as the intentional filamentation shown in 3.8. In this study the natural variation of emittance in the Main Injector provides the only variation in input emittance. There is a batch-to-batch variation of the mean emittance, as well as variation within a batch. The exact cause of this inter-batch variation is not teased out here, but will be discussed in ch. 5. But if the emittance variation is a consequence of dynamics in the Main Injector, the same dynamics may affect emittance growth at transition crossing.

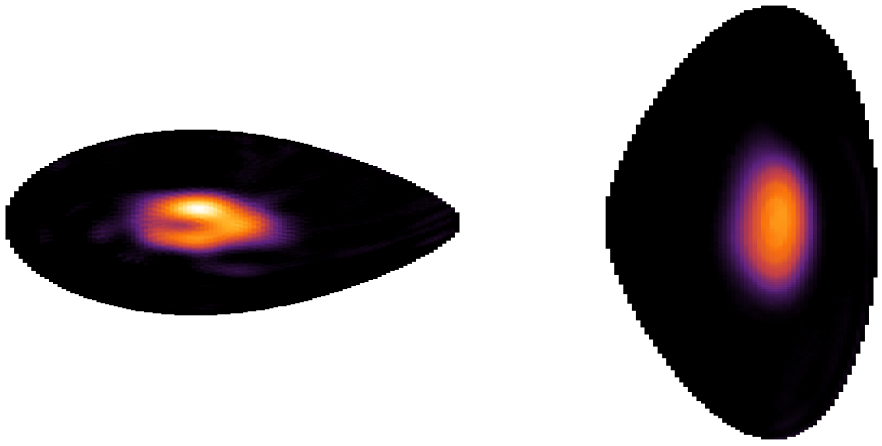

(a) Below Transition (b) Above Transition

Figure 4.3: A representative bunch from the single batch configuration reconstructed above and below transition. The bunches are properly scaled in ϕ and ΔE, and a common color scale is used to indicate charge density across images. The ϵ_{rms} growth across transition ($\Delta\epsilon/\epsilon$) is $\approx 32\%$ for this bunch. Notice that the charge distribution is relatively smooth, confined to the central, linear portion of the bucket, and nearly gaussian compared to the bunch shown in 4.4.

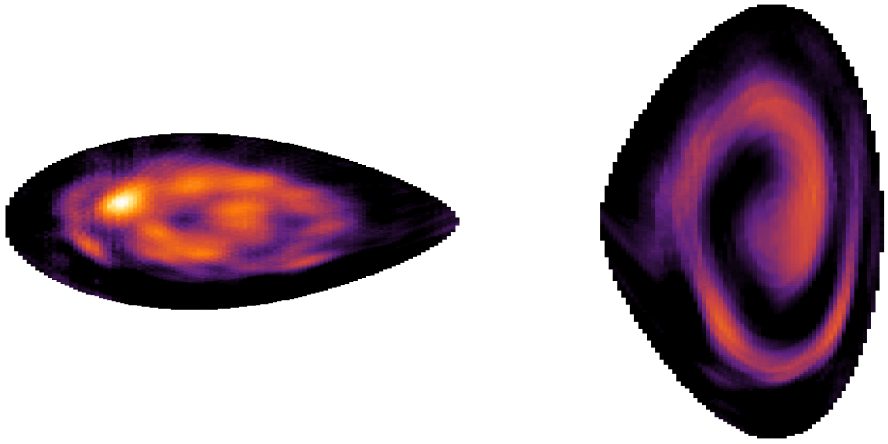

(a) Below Transition (b) Above Transition

Figure 4.4: A representative bunch from the slip-stacked configuration is reconstructed above and below transition. The bunches properly scaled relative to one another, and a common color scale is used to indicate charge density across images. The ϵ_{rms} growth across transtiion($\Delta\epsilon/\epsilon$) is $\approx$ 75% for this bunch. Notice how much more of the bucket the distribution occupies below transtion, and how irregular the distribution is. This bunch is neither gaussian, nor confined to the central, linear portion of the bucket.

As a note, voltage is a knob we have control over, even if it is one of the main sources error, and a good case could be made to study transition growth under various voltage configurations crossing transition. Expressing the equations for emittance growth, 2.38, 2.40, and 2.41 in terms of voltage for a constant emittance would allow some differentiation in effects, and offer additional insight. This has been done in Fermilab's Main Ring in the past [26], but I did not explicitly probe the voltage in this study.

4.3 Predicted Growth and Precision

Table 4.1 gives the parameters necessary to estimate emittance growth in the Main Injector crossing transition indicating single batch, or slip-stacked values as appropriate. The values at the top of the table are input parameters, and the values at the bottom are calculated. Precise values of the frequency, for instance, depend on the exact time in the acceleration cycle since this value varies to maintain synchronicity with the particle bunches. The ramp files used to control the machine are available, and were used for reconstruction purposes, this process will be discussed in ch. 5.

For both single batch, and slip-stacked cases $T_{na} > T_{nl}$, defined in eqns. 2.34, 2.37, which means the single particle contribution to the growth given by eqn. 2.38 is simply, $\approx 0.76\ T_{nl}/T_{na}$, and the beam is below the microwave instability threshold.

The values $|\mathrm{Z}_{||}/n|$, R, and α_1, the broadband reactive impedance, broadband resistive impedance, and first non-linear term of the momentum

Parameter	Value Single(Slip)
h	588
f_s[MHz]	≈53
N/Bunch	5.9E9(4.8E10)
γ_t	21.6
$\|Z_{\|\|}/n\|$ [Ω]	≈ 13
α_1	0.996
V[MV]	≈ 3.3(3.6)
ϕ_s[rad]	≈ 0.83(0.74)
T_{na}[msec]	1.88(1.76)
T_{nl}[msec]	0.37(0.72)

Table 4.1: Values used to predict emittance growth.

compaction (related to the slip-factor), used to estimate growth are design values and estimates taken from [30], [7], and [40]. These values will be estimated from the fit obtained to the growth measurements. [2]

Histograms showing the input emittance for the combined data in the single batch, and slip-stacked cases are shown in figs. 4.5, and 4.6. The variance of a gaussian fit for each case is ≈ 0.009 eV · sec, and ≈ 0.02 eV · sec for single batch and slip-stacked beam respectively. [3]

Given the relationship to input emittance from eqns. 2.40, and 2.38,

[2]Measuring the values, $|Z_{||}/n|$, R, and α_1 independently, and calibrating the growth model would allow for online monitoring of the sources of emittance growth. And using emittance growth across transition may be a viable way to tune the various RF parameters for minimal growth. Once calibrated, changes in the fit parameters would signal something amiss. Using this analysis during operation may be beneficial for machine tuning.

[3]The single batch beam seems to be bi-modal. Becuase this data was taken parasitically during normal beam operation, this probably represents some normal tuning of the Booster, which is supplying beam to the Main Injector. These plots are only used to estimate the range of growth that should be observed, so the bi-modal distirbution is not problematic.

Single Batch	Mechanism	Mean Growth	Range
	Reactive	18%	±5%
	Single Particle	14%	±5%
	Resistive	≈ 0	≈ 0
Slip-Stacked			
	Reactive	12%	±5%
	Single Particle	20%	±3%
	Resistive	≈ 0	≈ 0

Table 4.2: Growth predictions

table 4.2 gives the growth prediction for the mean input emittance in the single batch, and slip-stacked cases, as well as the range of growth given the spread measured in the histograms 4.5, and 4.6.

Fig. 4.7 shows the growth predicted by eqns. 2.38, 2.40, and 2.41 using the values estimated in tab. 4.2, and curves produced a ±50% change in each of the fit parameters. The black dashed lines represents the reactive term, to which the growth is most sensitive, the red lines the single particle term, and the blue lines the resistive term, barely distinguishable from the nominal value. This insensitivity to the resitive growth manifests as a very large error on the measurement of the parameter, and is due to the suppression of this effect by a factor of h, and one over the bunch width relative to the reactive growth.

Becuase this study is done parasitically, one reasonable possibility to consider is a correlation between emittance and bunch charge. This is not problematic, but we need to understand what effect such a correlation would have on the observed growth. The single particle growth, eqn. 2.38 depends

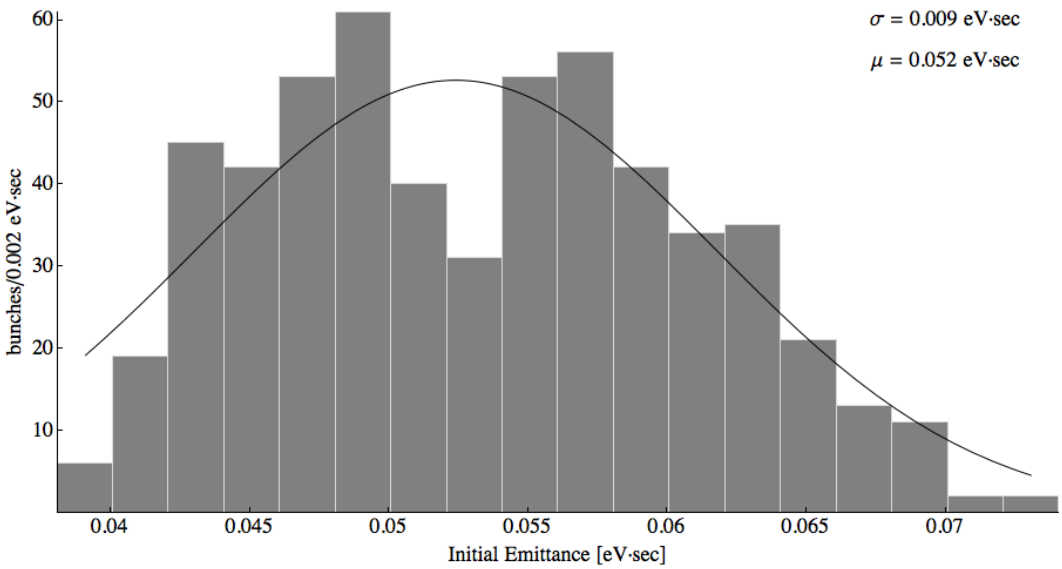

Figure 4.5: Initial $\epsilon_{68\%}$ of 693 bunches used to probe emittance growth across transition. The spread represents the natural emittance variation seen in the Main Injector under similar operating conditions, including the inter-batch variation shown in fig. 4.3. The emittance spread seen here correlates with bunch position in a batch and is likely caused in the early part of the Main Inejctor cycle, but may also come from the Booster. This variation allows us to probe the dependence of emittance growth on input emittance.

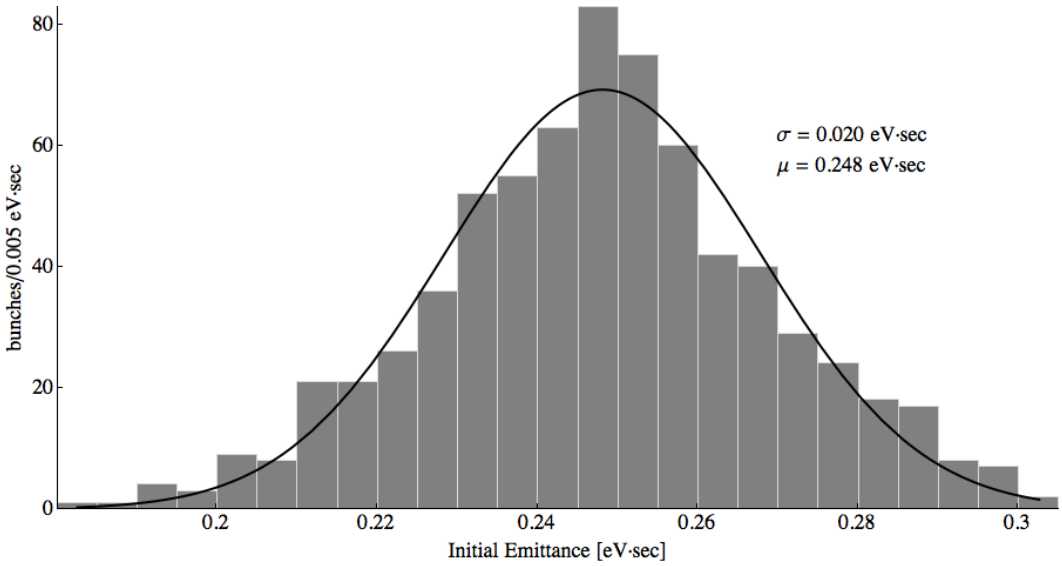

Figure 4.6: Initial $\epsilon_{68\%}$ of 728 bunches used to probe emittance growth across transition with slip-stacked beam. The spread represents the natural emittance variation seen in the Main Injector under similar operating conditions, including the inter-batch variation shown in fig. 4.4. The emittance spread seen here correlates with bunch position in a batch and is likely caused in the early part of the Main Inejctor cycle, but may also come from the Booster. This variation allows us to probe the dependence of emittance growth on input emittance.

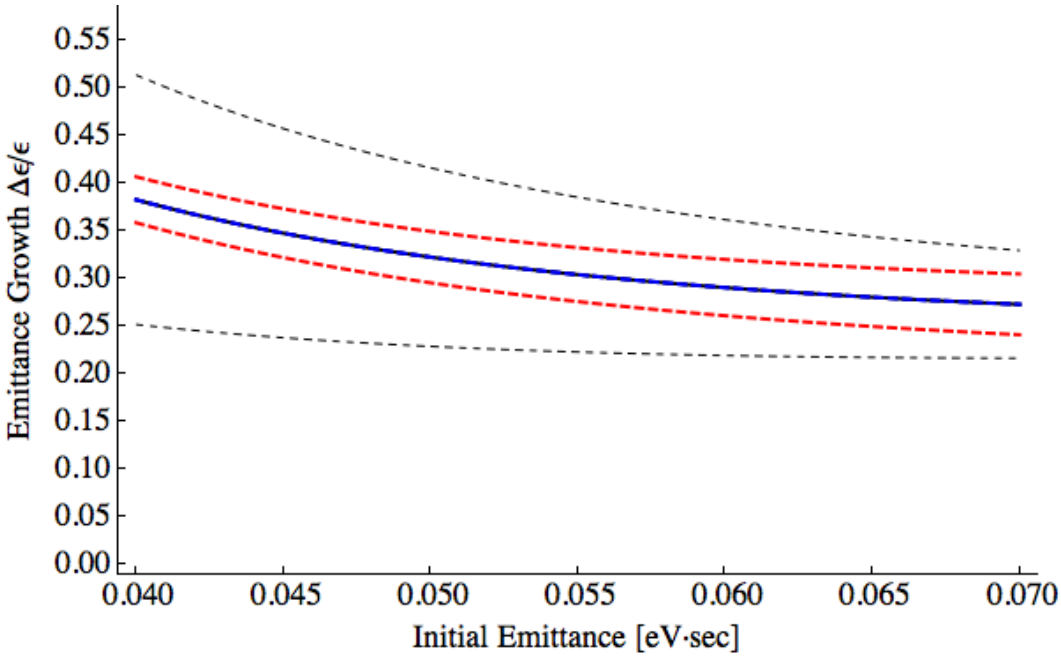

Figure 4.7: Predictions for the single batch emittance growth. The solid black curve indicates the nominal prediction. Growth is also indicated for each parameter independently varied by ± 50% of the value given in table 4.1 as dashed lines. The first non-linar momentum compaction (α_1, red), and reactive impedance ($|Z_{||}/n|$, black) terms dominate. The resistive(R, blue) curves are hardly distinguishable from the nominal growth.

only on emittance so would be unaffected. But the growth due to broadband impedance, eqns. 2.40, and 2.41 is proportional to bunch charge (through the beam current). Fig. 4.8 shows how a linear charge correlation would effect the slope of the emittance growth as a function of emittance using the fit parameters estimated in tab. 4.1. The blue, and red lines show the growth for bunches with 6% more, and less charge than the nominal value for single batch beam in the Main Injector. The line that connects the two is the observed growth for a postive linear correlation of charge with emittance over the full range of input emittance values in fig. 4.5.

Because the growth is a combination of single particle, and impedance effects we can't correct the growth measurements without knowing the value of the single particle fit parameter, the first non-linear momentum compaction, but we can correct the fitting functions by giving the beam current in eqns. 2.40, and 2.41 as a function of emittance.

4.4 Sources of Error

A schematic of the RWCM used in the Reycler and Main Injector is shown in fig. 1.6. As defined in eqn. 2.20, the beam current measured by the RWCM is the projection of the 2-D phase space distribution onto the ϕ, or time, axis. Fig. 3.3 shows many beam current measurements taken with the RWCM over successive turns. These 1-D measurements of the $\phi - \Delta E$ phase space distirbution are then used to recreate the full 2-D distribution shown in fig 3.4.

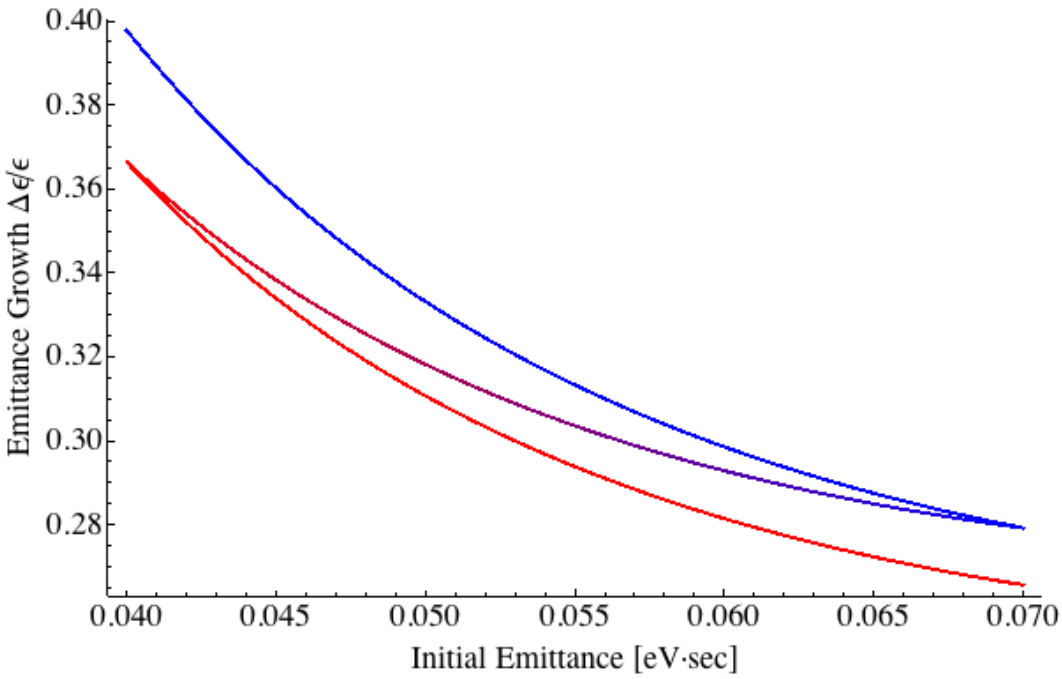

Figure 4.8: The effect of a linear correlation between charge and emittance on observed emittance growth. The blue, and red lines represent emittance growth for bunch charges ±6% from the nominal value for single batch beam taken from tab 4.1. The blue line has increased growth because of a larger bunch charge. The line connecting the two includes a linear correlation between charge and emittance. The decrease in the slope becuase of this correlation mimics an increase in the single particle effect. The opposite correlation would resemble an increased impedance.

The design considerations for this RWCM are described in detail in [13]. The RWCM shows a flat frequency response from 3kHz − 4GHz. This is on par with the 4GHz analog bandwidth of the scope used for acquisition. This means we should not expect to resolve features below ≈ 250ps.

The frequency response of the Main Injector type RWCM has been well characterized in [5]. The frequency response is then converted to a time domain response function and this response is deconvolved from the data. A single turn of the data shown in fig. 3.3 before(dashed line), and after(solid line) the frequency response correction is shown in fig. 4.9. The traces in fig. 3.3 have been processed in this way.

The RF frequency in the Main Injector and Recycler rings is about 53MHz, giving a bucket spacing of roughly 19 nsec, or 76 samples at 4GHz. The bucket height changes throughout the cycle, for this study the bucket height below transition is about 57MeV, and above transition is about 132MeV. If we assume the same discretization in both coordinates, a reasonable assumption, we have energy resolution of ≈ 0.75 MeV, and ≈ 1.75 MeV. This equates to a pixel area of 0.0002 eV · sec, and 0.0005 eV · sec which we can interpret as the resolution of the detector in phase space area. (The question of resolution is a difficult one that was mentioned briefly in ch. 3 with resepect to image discretization. The resolution here is also not a property of the detector, since it depends on the bucket height. In fact, it is probably variable over the bucket, and also connected to the number, and spacing of frames in the data.) Although there may be several pixels at the same action, we can estimate the

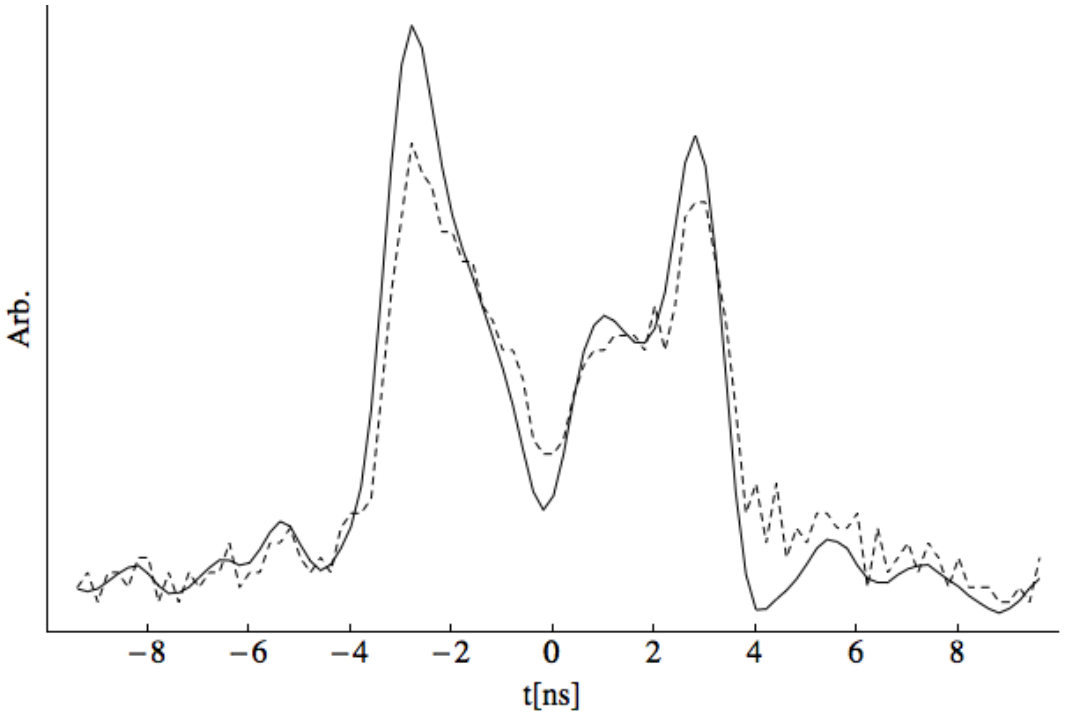

Figure 4.9: The normalized wall current data from a single bunch, on a single turn. The dashed line indicates the raw trace, and the solid line the trace corrected for the frequency response of the RWCM.

emittance measurement has a resolution of a similar order of magnitude.

For bunches of initial emittance ≈ 0.05 eV $\cdot$ sec, and ≈ 0.25 eV $\cdot$ sec for single batch, and slip-stacked beam respectively this means we should be able to resolve emittance on the order of a percent in either case. This should be taken as the maximum precision of the detector hardware. Other sources of error will degrade the precision.

Random errors due to electronics noise, or jitter in the acquisition trigger, are difficult to evaluate becuase of the nature of the reconstruction algorithm. In place of trying to provide an analytic expression of the error in the final image, it is more useful to employ a technique that uses real data to estimate the errors on a particular physical quanitiy, here the emittance.

The procedure is illustrated in fig. 4.10. Twice as many frames as necessary are used, or one full synchrotron period. The frame exactly halfway through the set, indicated by a white line, is reconstructed with only the first half of the data, and then independently with only the second half.

Fig. 4.11 shows histograms of the error between the emittance take from two such reconstructions for all bunches in the single batch data set, both above and below transition. The emittance from the reconstruction done using the first half of the dataset is taken to be the 'real' value, and the error is always the percent difference of the emittance measured from the reconstruction done with the second half of the dataset for consistency. This percentage was calculated for each bunch in the data set, and then binned to

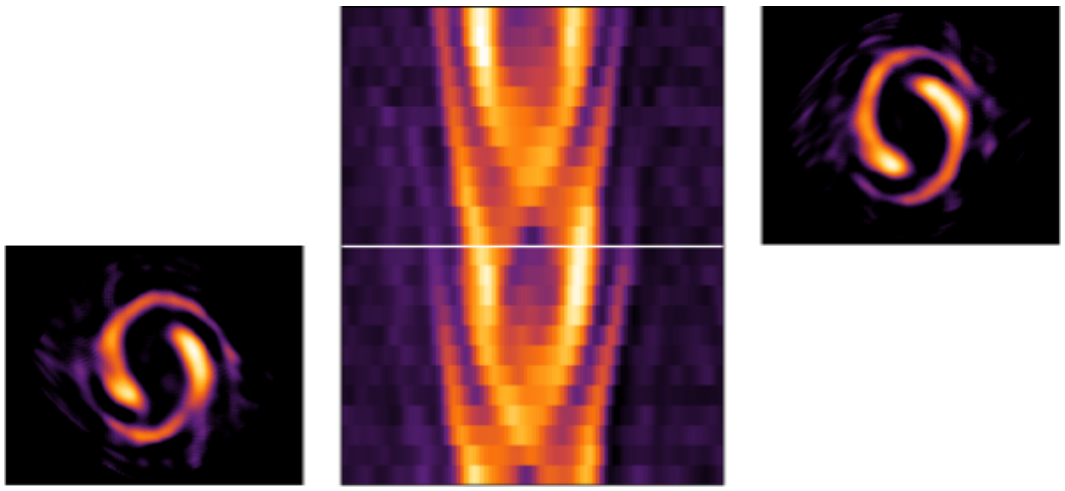

Figure 4.10: An illustration of the procedure used to determine the error on emittance measurements due to random errors in the data. A full synchrotron period of data is used. The phase space distribution is reconstructed twice at the same turn, indicated by a white line. Once using the first half synchrotron period, and again using the second half. The emittance is measured for each reconstruction and compared.

find the variance of the errors.

To reconstruct the 2-D phase space distribution, we must know all the parameters in eqn. 2.8: E_s, β, ω_{rf}, τ, ϕ_s, and V_{rf}. The energy, RF frequency, and η are very well known, to within a percent. Because τ is related to ω_{rf} through h, it is also known. This only leaves the RF voltage and phase. These are more difficult to pin down. The RF voltage is notoriously difficult, and tomographic reconstruction has actually been used to make a measurement of the voltage in the LEIR ring at CERN [18]. The RF cavity is subject to beam loading, several cavities are usually used, so an effective voltage must be calcualted, often the read-out is not calibrated, etc. These are not typically issues when running, because the programmed value is an approximate voltage, and a feedback system varies ϕ_s to maintain the proper momentum for the precisely known bending field. As a result the value of the RF voltage is among the largest sources of error.

The uncertainty in voltage is $\approx 10\%$. A similar method as used for the random errors can be used to calculate an uncertainty in the emittance due to the uncertainty in the voltage. By reconstructing with the nominal voltage, and a range of voltages within the uncertainty, the spread in emittance values is used to give an estimate on the emittance. For the measurments presented in this study, this adds approximately 2% error on the emittance measurements. Because ϕ_s and the voltage are related through the acceleration rate, which is known well, I take this 2% error to cover errors in either voltage or ϕ_s. From table 4.2, we should be able to resolve the expected growth, which should have

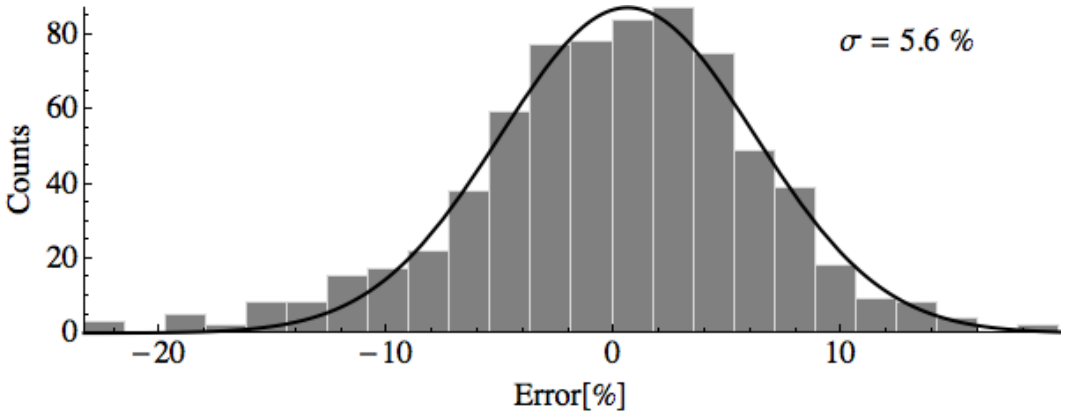

(a) Below Transition

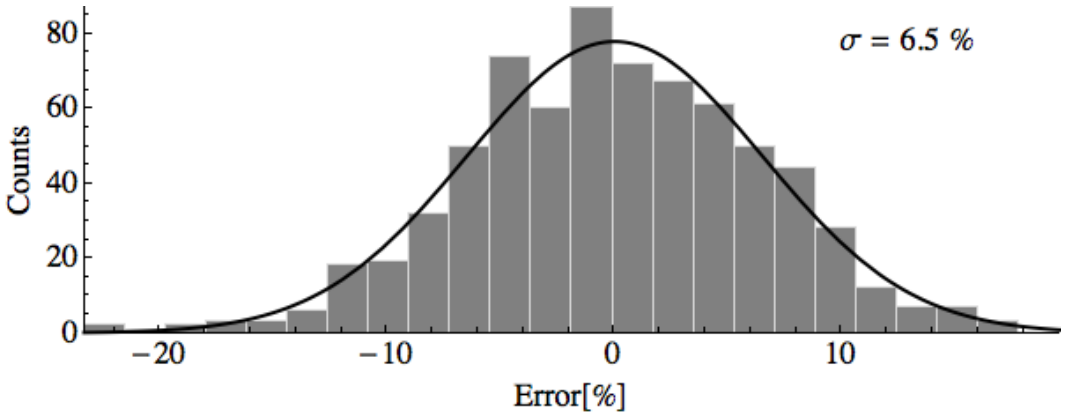

(b) Above Transition

Figure 4.11: The errors for single batch beam below and above transition obtained using the method described in fig. 4.10. The percent error is the difference between the emittance measured from the two reconstructions divided by measurement from the reconstruction using the first half of the data. The results for the slip-stacked beam were similar.

a range of approximately $15 - 20\%$ over the input emittances.

Adding the two contributions in quadrature, the total error from the uncertainty in voltage, and the random contributions is 7%, rounded to the nearest percent. Becuase these procedures are so approximate I don't feel that a more precise measure of the error is justified.

4.5 Summary

Transition is an unavoidable source of emittance growth. Even though losses during and after transition in the Main Injector are not problematic now, they may become problematic in the future at higher intensities. It is important to understand the leading cause of emittance growth when crossing transition to plan for future increases in intensity. Usually, instability is the main concern, but at very high intensities, either peak intensity, or high repetition rate, small losses can affect machine performance. Becuase of this it is not enough to establish thresholds for instability, but to understand precisely what affect transition crossing has on the beam.

This chapter presented the design of a study that was carried in Fermilab's Main Injector to understand emittance growth induced by transition crossing using the newly comissioned tomography system. The precision of the emittance measurements was evaluated and determined to be good enough that a meaningful study of the dependence of emittance growth on emittance can be probed parasitically, using only the natural cycle-to-cycle, and inter-batch variation in emittance. This measurement will provide critical machine

parameters necessary for evaluating the transition crossing. The next chapter presents the results of this study.

Chapter 5

Measurement of Emittance Growth Across Transition

This chapter presents a study done in Fermilab's Main Injector to determine the dominant emittance growth mechanism across transition as discussed in ch. 2, and to compare the predications made by the analytical formulas given to measurements of single batch, and slip-stacked beam.

The newly commissioned tomography system TARDIS is used to reconstruct bunches above and below transition during the same injection, and provide emittance measurements. The measurement is done bunch-by-bunch allowing the relationship between input emittance and emittance growth to be probed using the natural emittance variation present during normal operation. Over many pulses the emittance spread shows a variance of $\approx 10 - 20\%$ under the same operating conditions, some of this results from batch-to-batch variation, and some from inter-batch variation.

I will present the results of emittance growth measurements for two beam configurations in the Main Injector: single batch, and slip-stacked. Single batch means that each Main Injector bucket contains the beam from a single Booster bucket. This situation is close to ideal, since each bunch is a

single, nearly gaussian distribution near the center of the bucket similar to what was used to derive emittance growth predictions in ch. 2. Under normal Booster operating intensities I show emittance growth across transition is dominated by space charge growth, and that the single batch configuration is consistent with the theory presented in ch. 2. The slip-stacked configuration shows a deviation from the expected growth, but this is not unreasonable. Possible causes of this discreapncy, and methods for determining the validity of each cause are discussed at the end of this chapter.

5.1 Measurement of Emiitance Growth

In order to measure growth on a bunch-by-bunch basis, each dataset should contain a nearly symmetric window in time around transition. The emittance growth derived in sec. 2.2.2 occurs within a small window around transition. Becuase emittance measurements taken via tomographic reconstructions don't require that beam filament, the measurements can be made anywhere sufficiently far from transition that no more emittance growth is occuring. I use ±60msec, $\approx \pm 30 \times T_c$. This is sufficiently far from transition that the non-adiabatic effects, and the bunch narrowing can effectively be ignored. Fig. 5.1 shows the peak current for a single bunch as a function of revolution. The peak near the center represents transition, where the bunch has narrowed. Out near $\pm 60 msec$ the peak current is relatively flat, indicating that the beam width is changing slowly, if at all, relative to the sharp peak within $\approx \pm$10msec of transition.

Moving too far from transition risks conflating other effects that may lead to emittance growth. For instance, slip-stacked beam resembles two fish swimming tail-to-tail when initially recaptured, fig. 3.4, but has diffused such that the initial bunches are hardly recognizable by the time the measurement begins, fig 4.4. Without a measurement of the emittance at injection, I cannot determine if this is simple filamentation, or if there is emittance growth. If there is emittance growth, whatever is causing it may affect the results of the transition growth measurement.

Tomographic reconstruction of data requires several machine parameters, RF frequency and voltage, slip-factor, harmonic number, and synchronous phase and energy, to create the maps necessary for the backprojections. The programmed ramp values are available to obtain relevant parameters, but must be synchronized with the data. Limits on the memory of the digitizer mean the data acquisition starts in media res, with respect to the ramp. We need some reference point to orient to the ramp timing. Dynamics at transition provide such a reference point.

As discussed in detail in ch. 2 a bunch narrows approaching transition as η approaches 0, widening on either side of transition. This moment of peak beam current provides a fixed point we can use to orient the data relative to the ramp. The same figure, 5.1, I used to determine how far from transition constitutes 'far enough' is used to orient the data to the specified ramp file. Time $t = 0$ is transition, a convention established in ch. 2.

Once an approximate frame has been found for transition, the pro-

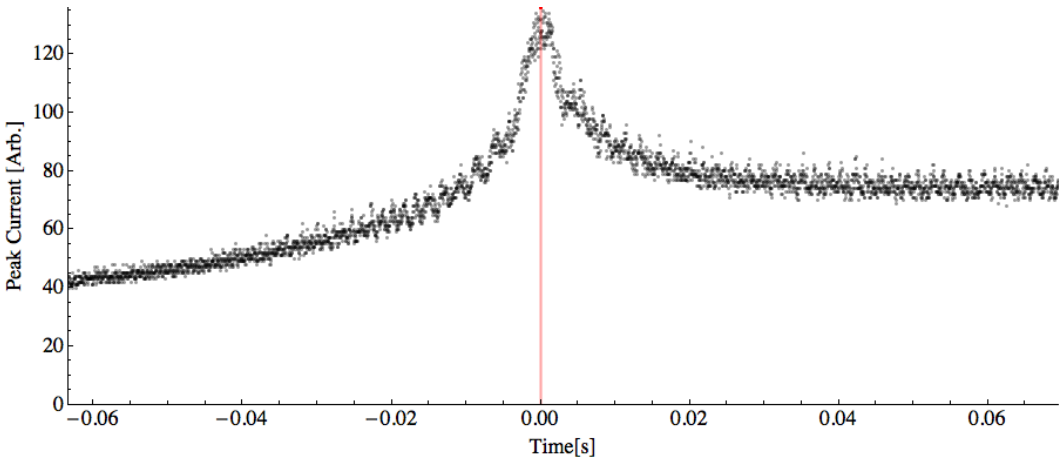

Figure 5.1: Peak current of a bunch is plotted as a function of time from transition. As mentioned in Ch. 2, the bunch width is minimal at transition($t = 0$), increasing on either side. This bunch narrowing means the peak current is maximal at transition. This high current means that coupling to machine impedances can become problematic. This point is used to orient the data relative to the programmed ramp. Before transition some oscillation can be seen in the peak current, this is likely due to the breathing motion of a quadrupole mode. The integrated charge remains the same, but the peak current oscillates as the bunch expands and contracts in the time coordinate.

grammed ramp values can be interpolated to determine the necessary parameters. Bunches are spaced 2π radians apart in phase, or $1/f_{rf}$ seconds, and the sample rate is known. These two pieces of information are used to further refine the timing.

Starting from some position near the first bunch, it is necessary to define a new bunch every $1/f_{rf}$ seconds. If the frequency is correct, the bunch centroids should be consistently spaced, even if they are not precisely centered. Any slope in the bunch centroid position across a batch indicates a frequency error which should be corrected. It is important to note here that centroid doesn't mean the centroid of a single frame, but the centroid of the synchrotron motion over many turns.

With the reconstruction parameters established for the data at $\approx 30 \times T_{nl}$ each bunch can be reconstructed. Figs. 4.3, and 4.4 show typical reconstructed bunches for single batch, and slip-stacked beam respectively. The bucket size, shape and orientation change across transition. The change in size comes from ramping the voltage, done to minimize losses, and the change in η as transition is crossed. The bucket appears flipped about the vertical axis because of the necessary phase shift induced at transition crossing, detailed in ch. 2. In addition to the required phase shift, the phase above transition is not exactly $\pi - 2\phi_{sBelow}$, but is determined by the voltage in order to maintain the programmed acceleration rate.

Once the bunches have been reconstructed, extracting emittance measurements is simple. Ch. 3 details the process for obtaining an emittance mea-

surement from a tomographic reconstruction. Figure 5.2 shows the normalized distritbution of particles as a function of action in the bunches presented in fig. 4.3, and 4.4 both above and below transition. The growth in emittance is easier to see when the data is presented this way, but recasting the data in this way is not necessary to do the measurement.

Fig. 5.3 shows the input emittance values for both single batch, and slip-stacked beam at the time of the 'below transition' measurement, about 60 msec below transition. The single batch beam shows relatively little variation in emittance throughout the batch. Most of the emittance variation comes from batch-to-batch variation. In contrast, the slip-stacked beam shows a definite correlation of emittance with bunch position in the batch. The exact cause of the emittance variation in a batch is not known, but is likely due to wakefields, since the trailing bunches show larger emittance.

Further study could pin down the cause of this variation. First one would need to establish whether the variation comes from the Booster or the Main Injector. Acquiring several frames of data at multiple points in the Main Injector cycle and comparing the emittance would indicate if there is emittance growth taking place between injection and transition. If the same trend in emittance is present at injection, however, the variation would be due to a machine earlier in the acceleration cycle, likely the Booster.

If the variation is being induced by the Main Injector study cycles could be used to further illuminate the cause of this emittance variation. Specially prepared batches that contain a large gap between the head and tail of

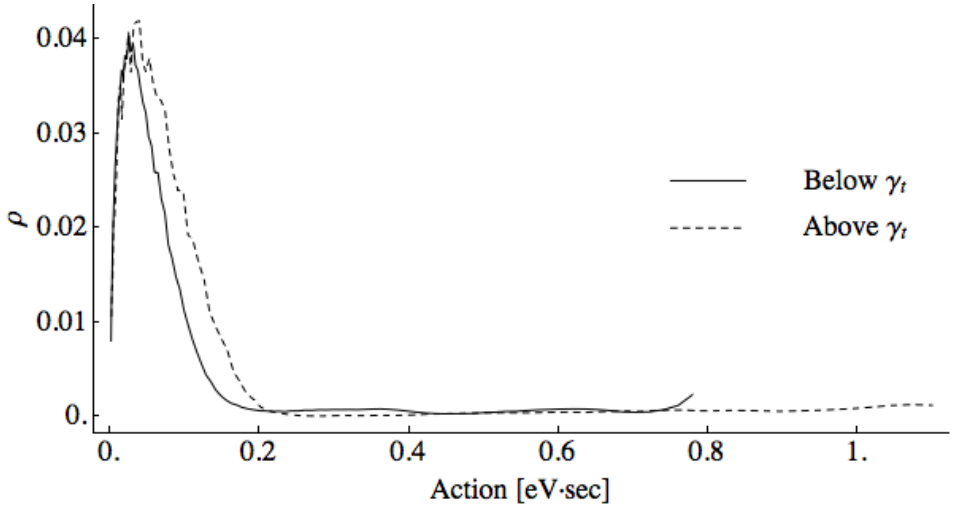

(a) Single Batch

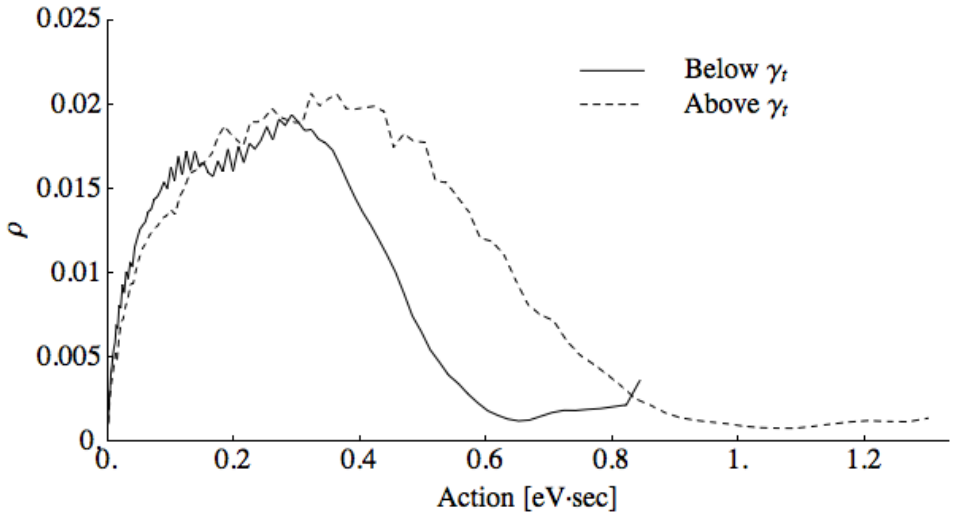

(b) Slip-Stacked

Figure 5.2: Normalized distributions of particles in action for the bunches in figs. 4.3, and 4.4. The solid (dotted) lines indicate below (above) transition. The curves below transition do not extend as far because the bucket area is smaller below transition. Growth is much more pronounced in the slip-stacked case, the distribution widening significantly. In both cases the curve below transition seems to increase right near the end, this is a consequence of beam becoming 'stuck' to the separatrix during reconstruction. This reconstruction artifact is cut when calculating emittance. See appendix A.

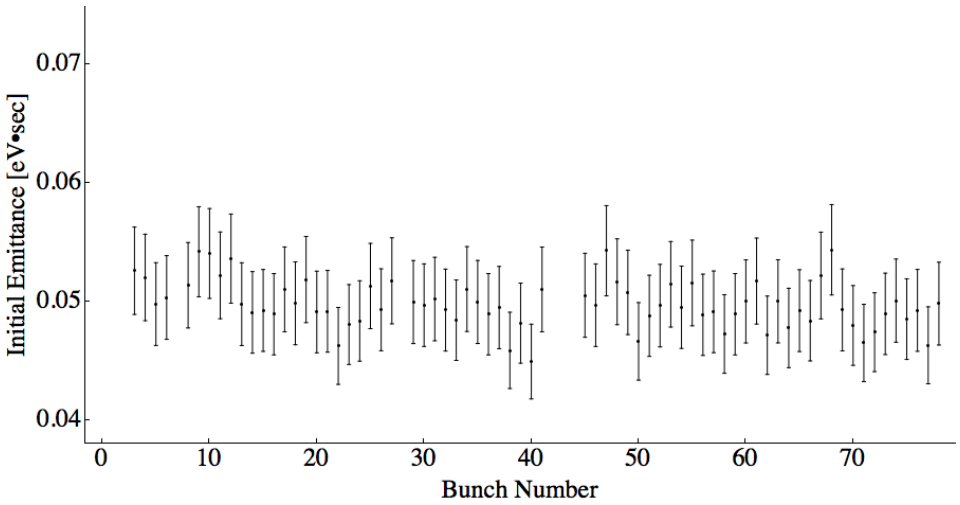

(a) Single Batch

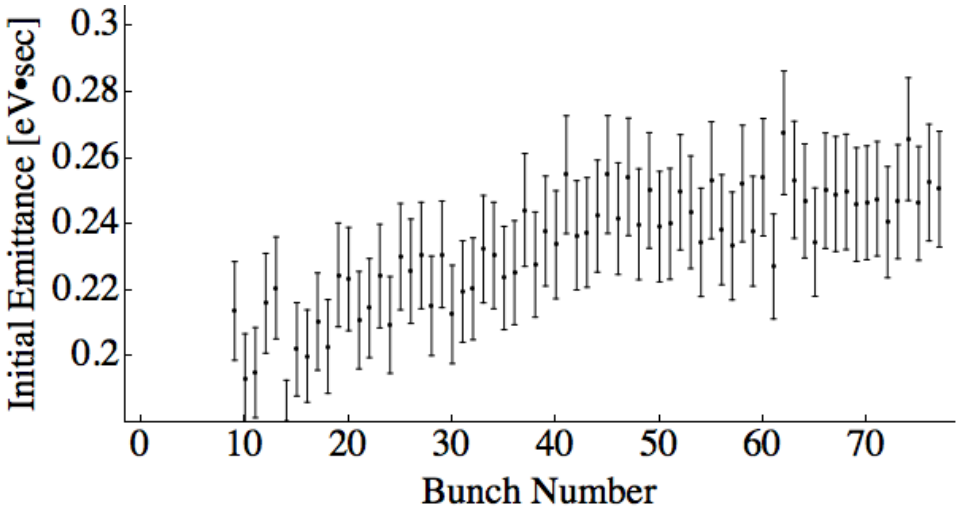

(b) Slip-Stacked

Figure 5.3: Emittance before transition plotted as a function of bunch position along the batch. A single example is shown for the single batch, and slip-stacked cases. Bunch 1 is the leading edge of the batch. Some of the leading bunches have been cut because they do not contain enough particles. The batchwide variation is random in the single batch case and small compared, but the slip-stacked beam shows a correlation between emittance and position, likely the result of wakefields.

the batch could indicate the range of the wakefields generating the variation. Fig. 3.13 showed some evidence that there are instabilities, mostly controlled by the longitudinal dampers, that increase emittance in the trailing bunches. This needs to be studied further to make any definitive statements, though.

Fig. 4.8 demonstrates the effect of a charge correlation on emittance growth, the measured correlation for the single batch, and slip-stacked beam is shown in fig. 5.4. This data is binned in the same way as the growth data in terms of initial emittance. The charge in each bin is scaled to the mean charge, and the result of a linear fit to the data is used as a correction to the beam current when fitting the growth, as described in the ch. 4.

To obtain this correlation, the total beam current is measured with a toroid in the Main Injector and reported in the control room as the number of protons in the machine. This value varies by about 1% from pulse to pulse, and is used to calibrate the RWCM data. I assume the charge in each batch is equal. But the charge in each bunch is calculated using the mean integrated number of digitizer counts over all frames used to reconstruct a particular bunch. A cut is made at 75% of the mean charge, bunches below this threshold are discarded. This cut mainly serves to get rid of the leading and trailing bunches that are affected by kicker magnets.

5.2 Results

Figs. 5.5, and 5.6 are the growth measurements for the single batch, and slip-stacked beam respectively. The values were binned along the abcissa

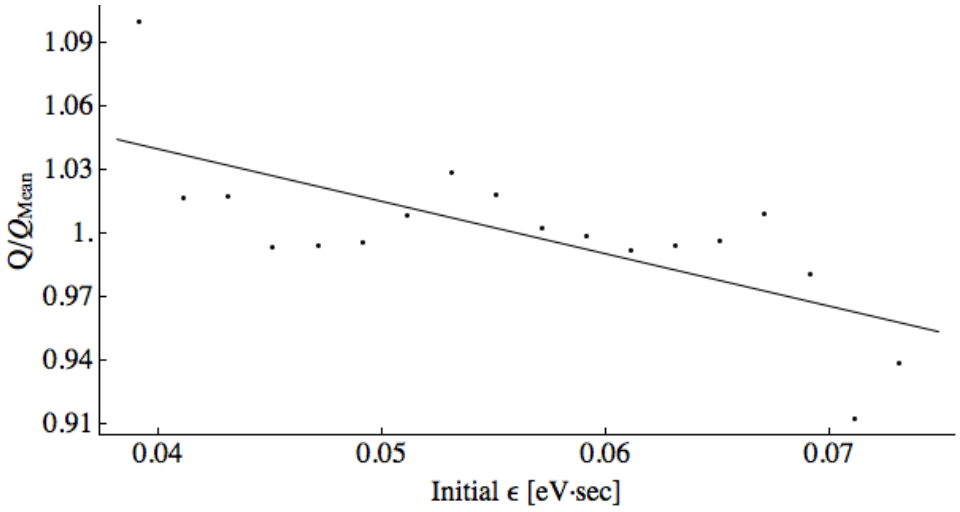

(a) Single Batch

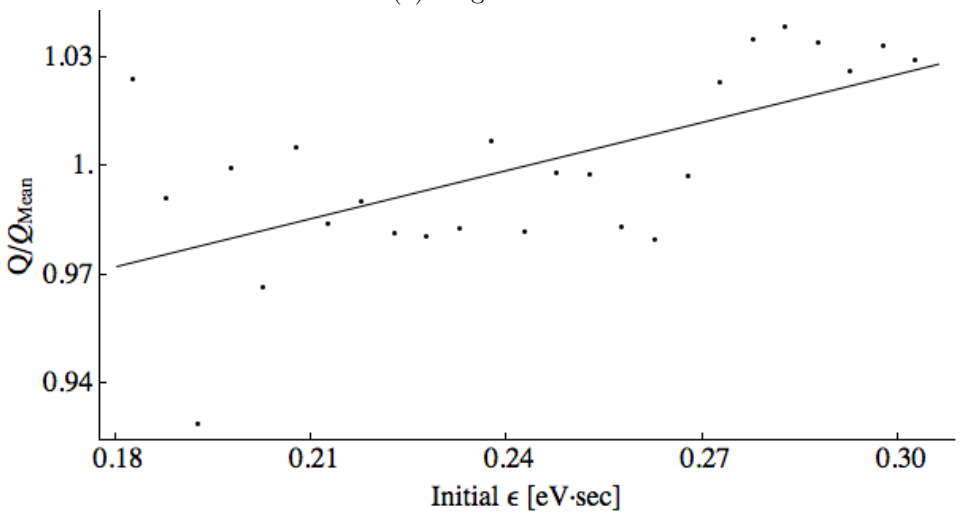

(b) Slip-Stacked

Figure 5.4: Bunch charge plotted as a function of initial emittance. A cut has been made for bunches with less than 75% of the maximum charge have been cut, typically the leading and trailing bunches. Charge is binned along emittance axis using the same binning that is used in the growth measurement. A linear fit to the correlation is shown. This fit is used to scale the peak current for initial intensity to correct for varying bunch charge on the slope of the growth. The correlation is not consistent for the two beam configurations.

to reduce statistical uncertainty, and fit using eqn. 4.1.

In both figures, the obvious downward slope indicates that collective effects are in fact dominating emittance growth across transition. The values obtained from the fit on the single batch beam are reported in tab. 5.1.

To within the indicated errors the fit parameter for the dominant contribution, $|Z_{||}/n| = 18.0 \pm 10.5\ \Omega$ agrees with the estimate presented in tab. 4.1.

The value of the first non-linear momentum compaction, $\alpha_1 = -1.6 \pm 0.8$, is not near the estimate in tab. 5.1. This is not entirely surprising. To my knowledge the first non-linear momentum compaction has not been directly measured in the Main Injector. The value of this parameter is determined, in part, by the sextupole fields that are used to set the chromaticity of the beam in the transverse dimension, a topic beyond the scope of this document. Since its design the chromaticity in the Main Injector has been increased to combat transverse instabilities [34]. There is at least one study indicating that large chromaticities would yield a value close to -1.5 for a lattice similar to the Main Inejctor [36]. (According to eqn. 2.38 $\alpha_1 = -1.5$ gives minimal growth for beams with $\beta \approx 1$.)

The resistive fit parameter is not very informative, but gives a value $R = 49.2 \pm 60400\ \Omega$. The large error means the growth is relatively insensitive to the broadband resistance. This was discussed in ch. 4. The resistive component is highly suppressed compared to the reactive component because of the large harmonic number of the machine and, an additional factor of the bunch width.

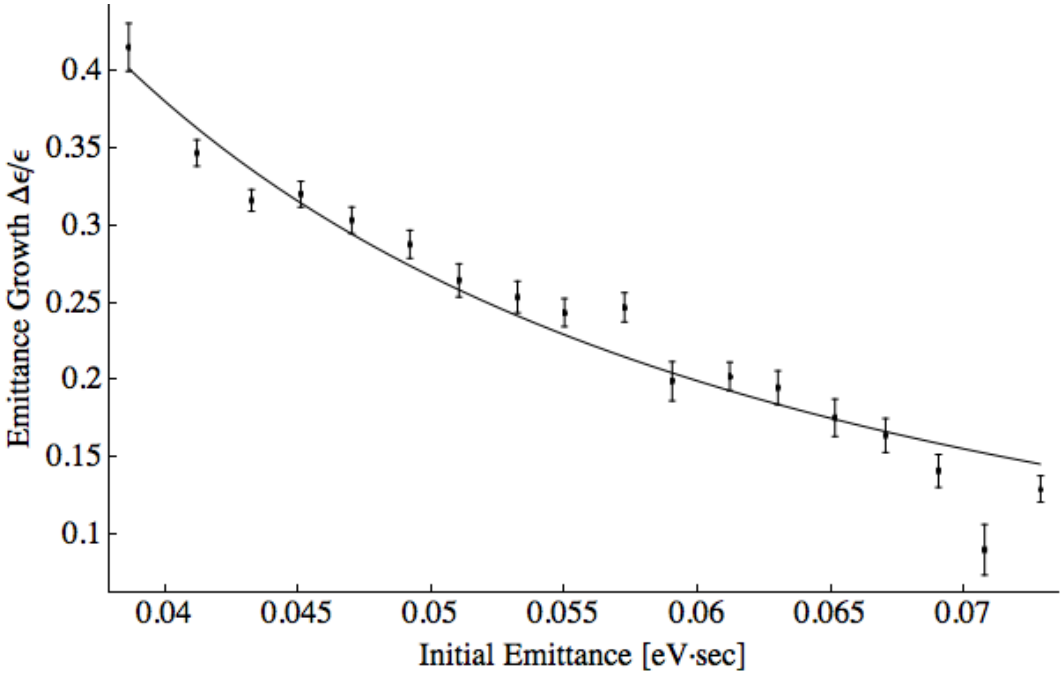

Figure 5.5: Fractional growth of ϵ_{rms} across transition for single batch beam plotted as a function of input emittance. 693 bunches were used and the growth averaged over all bunches within 0.001 $eV \cdot sec$ bins. A fit to eqn. 4.1 is shown as a solid line, with best fit values given in the accompanying text. The reactive impedance dominates the growth.

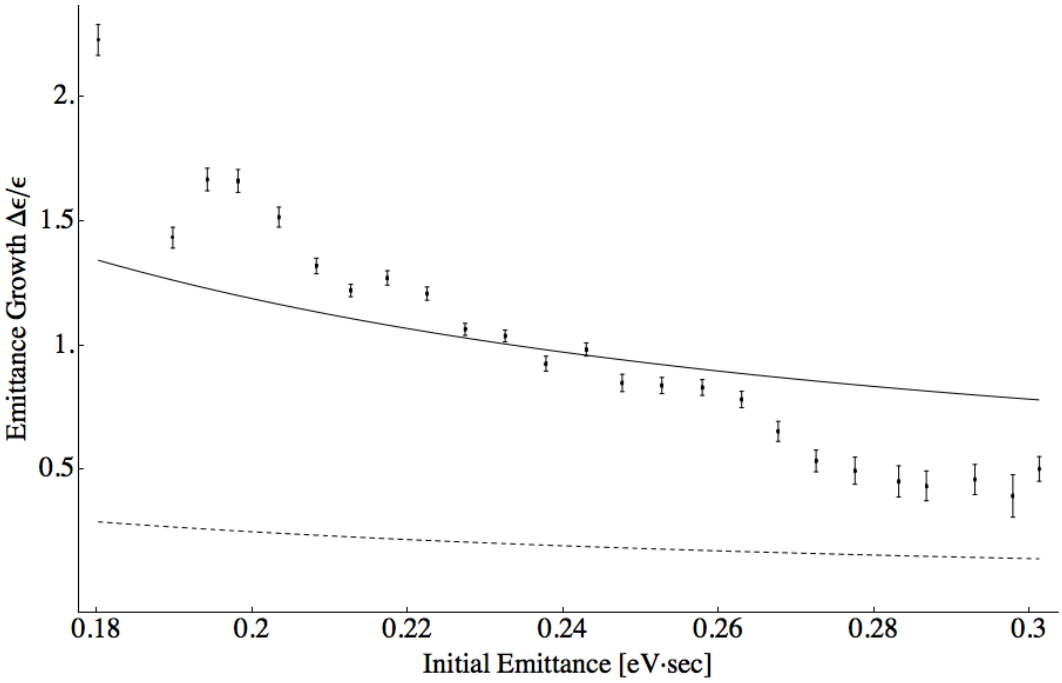

Figure 5.6: Fractional growth of ϵ_{rms} across transition for single batch beam plotted as a function of input emittance. 728 bunches were used and the growth averaged over all bunches within 0.005 eV $\cdot$ sec bins. The solid line is a fit to eqn. 4.1 and is not very good. The dotted line shows the prediction for the growth using values obtained from the single batch growth fit in fig. 5.5. The growth deviates significantly from the predictions using the values obtained for single batch beam. This is discussed at length in the text, in short I attribute this discrepancy to the violation of several assumptions used in deriving the growth contributions.

In fig. 5.6 the growth expected from the slip-stacked configuration using the paremeters obatined in the single batch fit is shown as a dashed line, and a fit obtained independently on the slip-stacked data is shown as a solid line. Obviously, neither of these lines is a good description of data. The independent fit is included because the growth expected from the values obtained using the single batch beam don't describe the data well. It is highly unlikely, however, that the machine changed significantly between the acquisition of the two datasets, though some variation is possible if components were fixed (or broken), or magnets settings were significantly different. (The space charge impedance has a weak dependence on the transverse beam size, a, that scales as $Log(1/a)$. This was completely ignored, but even allowing the beam size to change within reasonable limits would produce only a modest variation in the impedance [30].)

More likely is the violation of the assumptions that the beam can be modeled as parabolic, and that it is confined by the linear region of the RF cavity voltage used to derive the analytical model of growth presented in ch. 2.

5.3 Summary of Findings

Using the values of $|Z_{||}/n|$, α_1, and R as fit parameters the growth due to mismatches from the non-linear momentum compaction, and broadband impedance in ch. 2 do a good job of describing the observed emittance growth. The value of $|Z_{||}/n|$ determined by the fit agrees with the value calculated in [30]. The fit is insensitive to the resistive growth, so it is reasonable to

	Mechanism	Obs.$\Delta\epsilon/\epsilon$	Parameter	Best Fit Val.
Single	Reactive	0.13-0.37	$\lvert Z_{\lvert\rvert}/n\rvert$	$18.0 \pm 10.5\ \Omega$
	Resistive	$< 10^{-3}$	R	$49.2 \pm 60400\ \Omega$
	Non-Linear	$< 10^{-3}$	α_1	-1.6 ± 0.8
	Total	0.13-1.37		
Slip-Stacked	Reactive	0.81-1.66	$\lvert Z_{\lvert\rvert}/n\rvert$	$73.6 \pm 83.8\ \Omega$
	Resistive	$< 10^{-3}$	R	$45.1 \pm 252000\ \Omega$
	Non-Linear	$< 10^{-3}$	α_1	-2.8 ± 2.6
	Total	0.81-1.66		

Table 5.1: A summary of emittance growth measurements and best-fit parameters.

consider this value undetermined with respect to this measurement. The value of α_1 does not agree with the only published value I could obtain, but there is reason to believe that the number measured is reasonable.

In the case of the slip-stacked beam the same fitting procedure yields a value of $|Z_{||}/n|$ that is much higher than expected, and appears to be a poor match for any combination of the fit parameters considered. Using the value for $|Z_{||}/n|$ obtained in the single batch case, which I expect to satisfy all the assumptions used in deriving the analytical expressions, predicted growth is below the observed growth.

There are two interpretation: Either the fit is correct, and something was changed in the machine between the two measurements that introduced some other source of emittance growth, or the more likely case that the slip-stacked beam is exhibiting excess growth due to violation of one of the assumptions necessary for the analytic derivations presented in ch. 2.

Without further study the cause of the deviation between the expected

and observed growth can't be pinned down, but there are some hints. The following list outlines several possibilities for the divergence, the liklihood that each is responsible for some or all of the observation, and a brief explanation of the steps needed to evaluate each factor.

1. Violation of linear assumption: The derivations of emittance growth presented assume a beam contained in the linear portion of the bucket. This is the case for single batch, but not slip-stacked beam. Beam outside the linear region will necessarily lead to some emittance growth because the buckets before and after cannot be matched perfectly, given the opposite facing tails on either side of transition. This will also likely complicate the analytical functions for growth derived. This is a very likely candidate for at least some of the discrepancy seen in the slip-stacked case. The input distributions obtained here can be used as the input to simulations of the Main Injector to evaluate the actual emittance growth on realistic distributions. I have performed some preliminary simulations which suggest that beam in the non-linear regions of the bucket can contribute emittance growth of the order needed to explain this divergence with only linear dynamics, but I have not investigated the relationship between growth and initial emittance systematically to date.

2. Violation of parabolic assumption: The derivations also assume that the bunch is parabolic. This is a resonable approximation when the bunch is

gaussian, and simplifies calculation. This assumption holds for the single batch case, but not for the slip-stacked beam. Slip-stacked distributions tend to contain non-uniform clumps of beam. This seems a less likely contributor than the non-negligible occupation of the non-linear region of the bucket. Again simulations could help clear this up, spefically using the distributions obtained here to evaluate the effect of particular bunch distribution on emittance growth. This seems less likely, because the effect witnessed has such a coherent behavior, and bunch shape is generally quite chaotic. Unless some features of the bunch shape can be identified that correlate with input emittance (or bunch position, since earlier bunches tend to have lower emittance), this probably contributes to the variance at a particular point, rather than a coherent alteration of the functional relationship.

3. Interaction of growth mechanisms: Each growth mechanism here distorts the bunch in a different way. The reactive impedance squishes, or stretches the bunch in the phase space. The non-linear momentum compaction shears the bunch in one direction in time (the direction depends on the sign), creating tails. The resistive component creates a depression in the energy of the bunch proportional to the current. I have assumed that each of these emittance growth phenomena can be considered independently, and their contributions simply added together. In the case that the growth from two of the three is small compared to the remaining effect, this should be valid, which the single batch case suggests is true.

However, the interaction of these three distortions is difficult to write analytically, and may cause some of the observed discrepancy. This should be investigated with a series of studies, and simulations. Because the single batch case clearly has the reactive impednence as the dominant effect, and the measurement done in this study suggests that the single particle growth is minimal, this seems unlikely.

4. Change of machine impedance: The data for the single batch, and slip-stacked beam were taken at different times. It is possible that the impedance of the machine increased between the two measurements. This seems unlikely, as the shape of the emittance growth is not consistent with any of the chromatic effects mentioned here for any realistic values of the fit parameters, and does not resemble the microwave instability. But a large impedance cannot be completely ruled out as a contributing factor.

5. Mistuned transition crossing: There may be a phase, or energy, error present after crossing transition. This is certainly not true in the case of single batch beam, the signature would be a dipole oscillation that is not seen. In the slip-stacked case, there is obviously some filatmentation that could be due to a transition tuning error, see fig. 4.4. This can be modeled as an injection error. Such an error leads to fractional emittance growth that has a functional form consistent with the observed discrepancy, but would require an error larger than a bucket in either phase or

energy to account for the all of the observed growth. The reconstructions of the after-transition bunches in the slip-stacked case show that whatever error is present decreases with position in the bunch train.

6. Presence of high-Q resonator: All of the impedances considered here are broadband, it is possible that there is a high-Q resonator contributing to emittnace growth. Typically high-Q resonators are RF cavities that haven't been correctly shorted, or cavity-like structures. Whatever the offending element, it would need a mode near some harmonic of the beam frequency. This seems unlikely as a high-Q resonator would affect the beam prior to transition, and continue after transition, as well as affecting the single batch beam. Unless, again something was introduced into the Main Injector between measurements.

7. Instability near transition: There are slow-growing instabilities that can form [10], [39] when the beam is left near transition (head-tail, etc.). The acceleration rate through transition is fast, and the microwave instability is usually the first candidate for instability near transition. Because the beam is below the microwave instability threshold, and the characteristics of microwave growth don't seem to be present, I've assumed these other instabilities are not problematic.

A dedicated set of studies should start simple using single batch, low intensity beam so that collective effects are negligible. For the same intensity, beam should be intentionally filamented with a phase error to increase

emittance. The range over which this will still satisfy the assumptions used to derive emittance growth will be fairly small, but this would allow the experimenter to probe the initial emittance in a more controlled way. If the value measured here for the first non-linear momentum compaction factor is correct, the emittance growth should approach zero as the bunch charge is decerased.

If the phase error is increased beyond where the linear assumption is satisfied, eventually the beam will filament with a hollow center in phase space, as in fig. 3.8. When this ring is concentrated at the same action as the two beam distributions captured in slip-stacking, this would represent a clean, simulated slip-stacked distribution for the purpose of study.

Or low intensity beam could be slip-stacked. If the same emittance growth behavior as observed in this study is present, then the observed growth could be attributed to the violation of the linear assumption. If this is not the case, beam intensity can then be used as an explicit parameter in a more controlled way to probe the onset of the large divergence between theory and measurement in the slip-stacked case. In order to assure crossing errors are not to blame, the phase shift should be tuned for minimum dipole oscillation after transition. Quadrupole oscillation may still be present due to mismatches.

Simulations could answer three questions: how do the various growth mechanism affect one another when they are of the same order, what effect does the bunch distribution have on the growth when it is not parabolic, and how does the growth scale with input emittance when the beam begins to populate the non-linear region at an appreciable level. The tomograms obtained for the

study presented in this paper can be used as input in order to help answer the latter two questions. Realistic input distributions can be propagated through simulations and compared to the output distributions, an example of measured input and output distributions are shown in fig. 4.4. In general, the ability to simulate realistic distributions obtained through tomographic reconstructions is one appeal of doing tomographic reconstruction.

An independent measurement of the coupling impedances in the Main Injector would be helpful regardless of the outcome of this study. In addition to providing a check on the meaurements here, and suggesting possible solutions to the problem of steep depedence of growth on input emittance, an independent calibration could be used as a baseline allowing this method to be used as diagnostic to monitor the emittance growth over time. Deviations from historical norms could be used as an indication that transition is tuned incorrectly, or that there is a problem with the impedance of the machine.

Unless specially designed to avoid transition, sometimes at considerable effort, some synchrotron in an accelerator chain is likely to cross transition, and emittance growth will undoubtedly occur. And, because special RF beam prepration techniques like slip-stacking will continued to be used to overcome the intensity limitations of low energy machines necessary to reach final beam energies required by experimenters, it is important to understand the effect that non-gaussian distributions have on the evolution of particle distributions. Simulations are helpful, but can only go so far in this regard without some grounding in experiment. Beams are always becoming more intense, and tol-

erances on losses more strict, because of this it is critical that we understand how the beam evolves, and have ways to extract realistic distributions obtained experimentally to refine simulation, and control the beam quickly and intuitively in the control room. This study showed how tomography, and similar diagnostic techniques that use machine modeling, and measurement can be used to understand particle accelerators.

Chapter 6

Summary And Outlook

This chapter will briefly review each section of this document, highlighting the most important points before discussing the implications of this work, and the outlook for the future of particle accelerators.

6.1 Longitudinal Dynamics

Ch. 2 introduced longitudinal beam dynamics. The accelerator modeled in fig. 1.2 was used as a minimal example to discuss the synchronous particle, composed of the radio-frequency cavity used to accelerate particles, and magnets to steer particles, and the synchronous particle to which all the accelerator systems are synchronized, and which is the basis of most analysis. I discussed the effect of small errors in either particle energy or phase relative to the synchronous particle.

Particles exhibiting small errors in energy take more or less time to traverse the ring, because of the devaition in velocity, and path length due to bending magnets. This timing error leads to a phase slip for off-energy particles relative to the arrival time of the sycnhronous particle. When a sinusoidal focusing potential is applied with a radio frequency cavity synchronous with

the revolution frequency of the design particle, the restoring force exerted reduces the energy error of particles as they slip away in phase. Over many turns this leads to harmonic motion about the design values of phase, and energy in the longitudinal phase space. This motion was analyzed, and the conditions for stability established.

A small natural spread in energy, and timing with respect to the synchronous values is expected. These errors are described by a quantity called the emittance, often referred to as the phase space area the bunch occupies, but more accurately described as the average action of the particles in the distribution. Errors in transferring beam from one machine to another, or noise on magnets, or RF power supplies can cause the emittance of a bunch of particles to grow leading to operational problems. Particles that exceed the stable phase space area are lost either during acceleration, crossing transition, or upon extraction from the machine.

The electric fields produced by particles in a bunch can perturb longitudinal dynamics through interactions of the particles with one another, and with the environment. These intensity dependent, collective effects modify the focusing potential of the RF cavity distorting the trajectories in phase space, and depend on the actual linear charge distribution. If the perturbation is large enough beam motion can become unstable and beam is again in danger of being lost.

As particle accelerators push the limits of beam power into the regime of many megawatts the tolerance for losses must become increasingly strict.

Losses near the sensitive superconducting equipment increasingly necessary for high power accelerators can lead to quenches and affect machine reliability. Nuclear activation of the equipment in beam enclosures must be kept to a minimum both because of environmental concerns, and worker safety. The downtime necessary to fix damaged components, or to allow safe access may seem like mundane annoyances, but every second an accelerator is offline is a second that it is not carrying out its mission for science or industry.

Losses are monitored and the accelerator is shut down when thresholds are exceeded, but this is a coarse control. In addition to direct measures, diagnostics used to preciesly tune beam to the optimal conditions, and monitor operations are crucial in maintaining reliability, and saftey as beam power increases. Some diagnostics provide very precise, and accurate direct, destructive measures that require no reconstruction. These serve a necessary purpose in bringing a machine up to its operational state, or in the case of linear accelerators where beam is dumped after every cycle. However, reliable, continuous operation of high power accelerators benefit from real-time, non-destrutive, easy to interpret diagnostics that allow operators to quickly identify aberrant machine behavior.

6.2 Tomographic Reconstruction

In analogy to medical tomography, tomography in particle accelerators allows for the reconstruction of two-dimensional phase space relying only on non-destructive measurements of beam projections. The method has histor-

ically been used with several algorithms, filtered backprojection(FIR), maximum entropy(MENT), and algebraic reconstruction techniques(ART) to name a few. Because of the unique non-linear nature of the problem in particle accelerators, algebraic reconstruction techniques tend to work best. They are fast, robust, and conceptually simple. ART has been applied to longitudinal [23], [20], and transverse [19] dimensions. The work presented here was mostly concerned with providing an online, real-time, multi-bunch system for large synchrotrons like Fermilab's Main Injector, and Recycler Rings.

The method has been expanded to include the space charge contribution [21] but other energy loss mechanisms have not been included to date. These could be added to the tracking equations to allow application to rings such as electron based light sources where radiative energy losses are part of the design, and normal operation of the machine. The inclusion of space-charge forces requires that the reconstruction maps be calculated separately for each bunch, as the energy loss is dependent on the precise structure of the beam current. This is problematic for real-time systems in large synchrotrons at present.

For more intense beams space charge considerations become increasingly crucial to acheiving very good reconstructions. In order to properly include all of the energy loss effects, the impedance of the machine in question must be known very well. Very intense beams are becoming more common as new uses are devised for particle beams, and precision experiments are pursued.

The tomography algorithm is separable into a map building, and a recosntruction portion. Leveraging the extensive work put into tracking programs used at the major accelerator complexes to create the maps, where all of the beam physics knowledge is needed, so that the tomographic reconstruction program is not reinventing the wheel is a practical step that would go a long way toward making tomography available in more machines. There will certainly be tradeoffs when the program is not built from the ground up to take all reasonable approximations into account for a particular machine, but this may be offset by the effort of much more talented computer scientists working to optimize performance, and the very detailed machine models that are typically available.

In the future, full 6-D phase space reconstruction should be possible for the longitudinal, and both trasnverse phase spaces. Such a full 6-D reconstruction could be used to propagate the distribution measured at one point in a ring to obtain the 3-D real-space bunch distribution at collider interaction points for calculation of the overlap integral of colliding bunches. This method may be able to improve luminosity calibration, which represents one of the largest uncertainties in collider experiments. In the LHC for instance, the luminosity measurement has an error of 11% [3].

A technical extension of tomography that has not been implemented yet, but will likely be necessary for full 6-D phase space reconstructions is the inclusion of coupling. Typically the longitudinal and transverse coordinates can be treated independently, though not always. In the same way, each of

the transverse coordinates can be treated independently. This separation of coordinates is more likely to be violated, as simple angular alignment errors in magents can couple the motion of particles in two transverse dimensions. To my knowledge, beam exhibiting coupled motion in the transverse plane has not been reconstructed to date.

Currently the only way to determine errors in reconstruction is empirically. A bunch must be reconstructed many times while varying parameters to map out the distortion of the bunch as a function of the parameter under investigation, or to relate errors in parameters to errors in particular physics quantities. This is possible offline, but online, or in the case of space charge this quickly becomes overwhelming. More work to understand how various factors affect the error in reconstructions, or the development of simpler algorithms to determine errors would help tomography reach full potential for scientific, rather than just operational, applications. Without work of this kind, it will be difficult for the technique to be used by anyone other than experts in the diagnostic itself for serious studies.

Tomography is an analysis technique applicable regardless of the detector technology used to obtain projections. The RWCM detector technology used at Fermilab is very good, it is non-destructive, and has a large bandwidth (3kHz-4GHz) [13], but it was not developed with tomography in mind. The 20ns buckets in the Main Injector and Recycler are near the limits of what can be reconstructed with the existing RWCM's. In fact, reconstructing in the Main Injector during a particular 120GeV mode called resontant extraction

that uses very short bunches is almost impossible. Because the features of the very short bunch cannot be resolved nearly all reconstructions occupy on a few pixels.

The current limit on precision of reconstructions, and consequently physics quantities, is the detector bandwidth even though it is large. As tomography becomes a more viable real-time diagnostic, dedicated detectors should be designed to address the particular challenges of tomography: high bandwidth, non-destructive, accurate reproduction of high frequency features, turn-by-turn, and bunch-by-bunch capabilities. To date Fermilab's tomography system has not been calibrated using a destructive, direct method of phase space visualization, but this could greatly improve the quality of the non-destructive tomography, increasing faith in the precision of measurements.

It should be possible to extend the reconstruction technique used here to very stable linear machines. [1] The mathematics of reconstruction only requires that there be some known transformation relating projection data to some initial 2-D distribution. Varying the phase with which a single pulse encounters an RF cavity, or a series of RF cavites, over many pulses may provide the necessary transformation. If the beam distribution is stable from pulse to pulse it may be possible to reconstruct the 2-D distribution. Much work needs to be done to determine the limits of such a technique, but the algorithm would be essentially the same ART technique used here. In this case however,

[1] In fact some work is ongoing at the FACET facility at SLAC National Laboratory to do just this.

it may be more useful to use a destructive, direct technique on a single bunch, as the required transformation may be so disruptive to normal operations that it is effectively destructive. Still, the normal destructive methods for imaging very short bunches can be prohibitively expensive, such as the RF deflecting cavity, XTCAV [6], used at SLAC to reconstruct the longitudinal phase space of LCLS beam after the FEL, which requires a full 50 MW power supply system and costly X-band cavities near the beam dump, far away from the acclerating stations.

As I showed in this document with measurements near transition, tomography can be used to probe dynamics that are not modelled by looking at regions where the perturbing effects are negiligible to infer the properties of the machine. The number of situations where this sort of technique can be employed is certainly limited, but there may be some niche applications for which it is helpful.

In a similar way, suppose some machine suffers from space-charge effects, or some unknown instability, but is an intermediate part of an accelerator chain. Beam could be injected into this machine and allowed to evolve for some short time, t. After t seconds have elapsed the beam is extracted to another machine that is not experiencing such problems and the distribution is reconstructed. By varying the amount of time the beam is allowed to evolve in the problem machine, the evolution of beam in that machine could be imaged directly.

Many of the operational benefits of having a real-time, multi-bunch

system already realized at Fermilab were presented in this document. These benefits support the argument that longitudinal tomographic reconstruction should become a part of nearly every circular machine. The necessary diagnostics equipment is relatively modest, and with increases in computing power, a dedicated personal computer can provide close to real-time images of nearly every bunch even in large synchrotrons on a time-scale appropriate for machine tuning.

6.3 The Future of Accelerators

Particle accelerators provide intense, high energy beams of particles. These machines serve the scientific purposes of giving experimenters access to regimes of energy and statistics not available from natural sources to probe fundamental interactions, or create intense beams of radiation. They serve practical purposes by allowing engineers to precisely control beams of particles used for everything from welding, to cancer treatment. The job of the accelerator scientist is always to push the operational boundaries, increasing the possibilities for science, and industry. This can take many forms. Sometimes we push the boundaries of single particle production, where precise knowledge of the energy and timing of one particle is important for instrument calibration, but more often we want more particles, confined to smaller regions, at higher energies, more reliably. The first two goals are often acheived in tandem, creating very intense beams to produce very bright, coherent beams of radiation, or very precise measurements of rare phenomena. The highest en-

ergies are currently only being chased at a single machine, the Large Hadron Collider(LHC), where scientists are uncovering the most fundamental laws of interaction of the natural world.

To acheive very intense beams at more modest energies than the 4 TeV of the LHC, but still in the many GeV range, requires accelerator scientists use many methods to overcome the limitations inherent in the activity of confining charged particles to small volumes. At low energies the electric force between particles tends to push them apart, as they traverse the metallic enclosures they also interact with the environment. The more particles there are the more strongly these collective effects interact with the beam. These interactions are usually less problematic at high energies as well. There are many schemes to increase the amount of beam in a high energy machine, while reducing the amount of beam in a low-energy machine, for instance the slip-stacking done at Fermilab.

There is a direct increase in beam intensity that comes from stacking mulitple beams, but in the case of Fermilab, slip-stacking the Reycler while Main Injector accelerates is another way to boost average beam power. By running these operations in parallel, the complex becomes more efficient, allowing more protons to be sent to experiments. It is improvements like this that will continue to encourage physicists to explore alternative beam delivery schemes with existing accelerators.

Each experiment has specific requirements, and sometimes it isn't the total intensity that needs to be increased, but the peak intensity. In this case,

two RF systems with different harmonic numbers may be used to capture many small buckets into one large bucket created by the lower frequency RF. This allows beam initially distributed in series with respect to time, to be rotated in the phase space to occupy a small time spread, with a larger energy spread. This 'coalescing' is actually the problem tomography was devised to monitor at Fermilab in 1987 [23].

These schemes, sometimes called RF gymnastics, will only multiply as experimenters demand more beam, and accelerator scientists are faced with large existing complexes that still have plenty of operational life. Although particle accelerator simulation software is a very developed technology it still needs to be checked against experiment, and beams still need to be monitored for all the myriad ways that things can go wrong. The non-gaussian, and unmatched distributions that necessarily arise from techniques like slip-stacking can't be treated analytically in terms of simple distributions, and simple measures of peak current, or emittance. Because of this it is necessary to have diagnostics that give a clear indication of how beam is evolving, and that provide realistic particle distributions to check the results of simulations, and provide distributions for the input to simulations.

Tomography offers the opportunity to view the actual 2-D distributions generated by these schemes, and to study the way they evolve in machines. This document presented one study well-suited to tomographic reconstruction where emittance measurements need to be taken before a beam filaments. This can be important when a beam spends little time in a particular machine, and

the researcher cares about the actual phase space distribution of beam.

There is no doubt that going forward acclerators will play a key role in helping scientists uncover the laws of nature. As the desire for more beam pushes up against the abilities of existing machines, accelerator physicists will devise new machines, and new ways to use old machines. Tomography, and other techniques that blend machine modeling, high-powered computers, and live data will play a critical role in increasing the potential to do cutting edge science by providing unprecedented views into the operation of these great machines.

Appendices

Appendix A

Errors and Reconstruction Artifacts

In tomographic reconstructions spurious features are called artifacts. There are many sources of artifacts, each with their own distinguishing features. The following is a brief list of artifacts, their causes, and ways to remove them if available. This section is intended as a brief guide to the qualitative signatures of various errors, not a comprehensive account of the quantitative effects of such errors. Typically figuring out the quantitative errors introduced must be done on a case by case basis, as they can vary with many parameters and may affect different measurements done using the reconstruction differently.

There are several ways to identify artifacts. One is to reconstruct an object several times, using the output of one run with a particular feature removed and observing whether the algorithm converges to an image containing that feature under subsequent trials [14]. This is similar to starting a fitting algorithm at a different point to identify whether an extremal is local or global. Another option is to simulate an object, and the data that object would generate and then distort the data using a model of the effect under consideration to see how the object is distorted. This appendix uses the second method for

all illustrations. For all examples in this appendix the profiles shown in fig. A.1 are used as the baseline matched distribution. This parabolic distirbution is generated according to the bunch model used in sec. 2.2.3.

The perfectly matched distribution shown should represent a bunch that is smoothly distributed out to a given radius, notice that there are already some artifacts present in the reconstructed profiles before we've begun to add error conditions. These are intrinsic to the reconstruction process and manifest in the reconstruction as non-smooth areas, in particular a cross-hatching and a soft edge to the distribution are obvious in this reconstruction. These can be eliminated by increasing the discretization, or the number of test particles launched per cell, but typically these actions will increase the reconstruction time and may lead to their own artifacts, so I use this example as a baseline.

A.1 Angular Range

In a rigidly rotating system, a patient in a medical setting for instance, the data should represent projections taken over π radians. [1] Fig. A.2 shows several reconstructions of a simple matched distribution with the same angular resolution, but different angular spans, from $0.2\pi - 2\pi$ in terms of synchtrotron rotations. The range plotted for each of the projections is the same, with the gray region indicating frames not included for the reconstructions using less

[1] The projection onto 0 and π radians will provide the same information, but if the angular steps are such that subsequent rotations provide angular interpolation, more than π radians may be desireable.

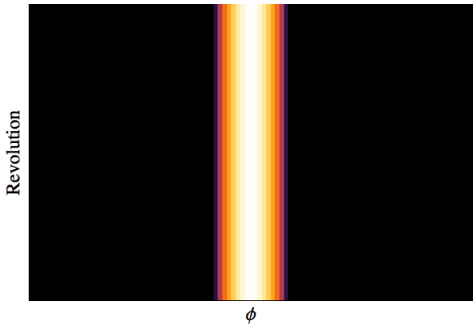

(a) Input

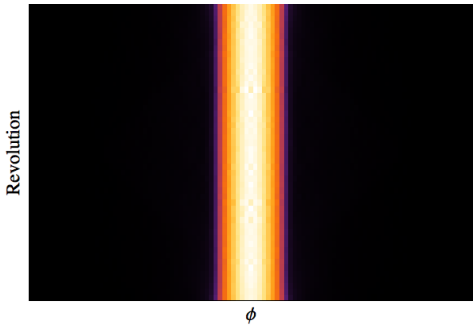

(b) Reconstructed Profiles

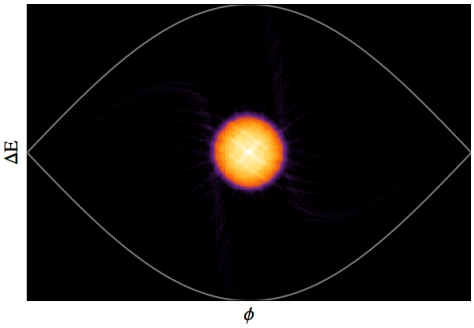

(c) Reconstruction

Figure A.1: The model parabolic distribution used to illustrate some common artifacts. (a) are the input profiles. (b) the baseline reconstructed profiles, and (c) the reconstructed 2-D distribution. Even before adding any error conditions some artifacts intrinsic to the reconstruction algorithm are present in the reconstructed profiles, and 2-D disribution. Note the spurious cross-hatching, and soft edge on the distribution in (c).

than a full synchrotron period.

Notice that beyond a half a synchrotron period the reconstruction does not change. In the reconstructions spanning less than π radians there are large streaks evident in the reconstruction. If this sort of streaking shows up the number of turns used for reconstruction should be increased.

Because the synchrotron frequency drops to 0 at the separatrix, the angular range requirement will never be completely satisfied throughout the bucket, though the change is relatively slow until just near the separatrix so in practice there is no problem. Additionally, the beam is often confined to the central region of the bucket and this effect isn't noticeable. If the beam is very near the separatrix, and this is either uncorrectable, or desired, more turns should be used for the reconstruction.

A.2 Angular Resolution

If the frames of data span the full π radians or more but the angular separation between frames is too large, the reconstruction may show artifacts like those in fig. A.3. This can be problematic even if the acquisition is turn-by-turn if the synchrotron period is very short. The solution to this problem is to simply add more data. It is highly unlikely that the synchrotron period will be precisely an integer number of turns, so the additional frames should provide angular interpolation beyond π radians. In this case the number of frames should only be increased until the artifact disappears. A more rigorous condition, time permitting, is to add frames until the discrepancy is minimized.

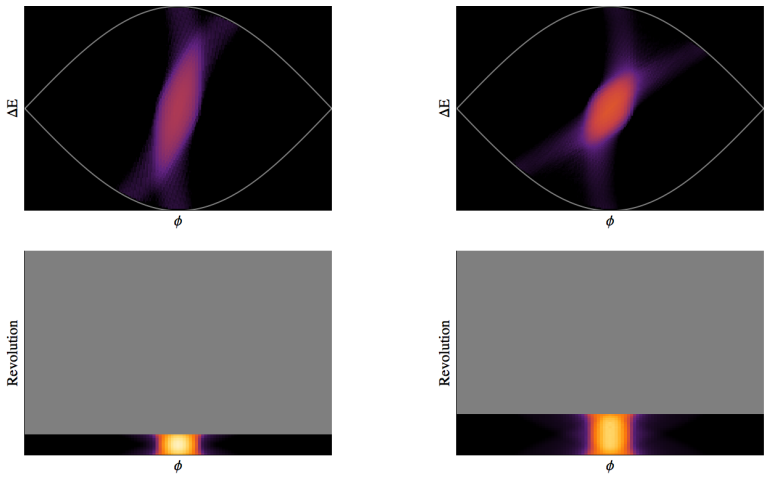

(a) 0.1 T_s

(b) 0.2 T_s

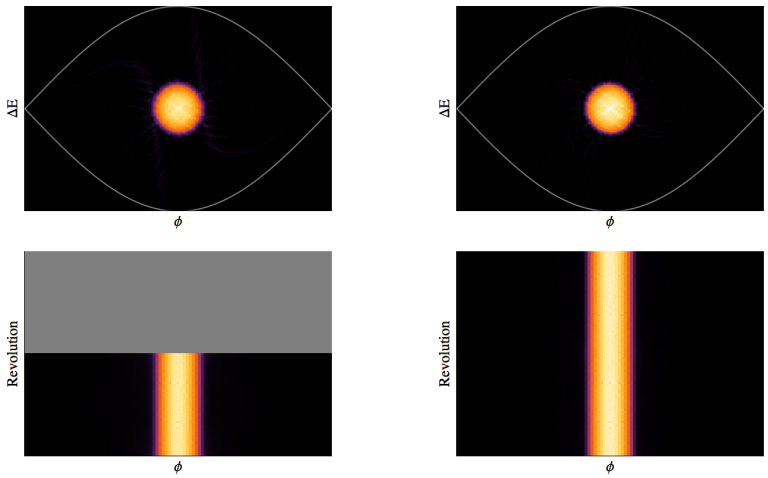

(c) 0.5 T_s

(d) 1.0 T_s

Figure A.2: The effects of insufficient angular range on reconstructions is shown. Each reconstruction uses some fraction of a single synchrotron period, T_s, taken at the same rate, one frame per revolution. Notice that beyond half a synchtrotron period, the additional information does not improve the reconstruction.

As soon as adding more frames increase the discrepancy, no more frames should be added.

A.3 Random Noise

Regardless of the source of random noise in the data, the effect is generally the same. Noise in the data frames contributes noise to the reconstructed image. Fig. A.4 shows reconstructions with random noise generated to lie within a full width of $\pm 10\%$ of the peak signal. Becuase this artifact depends on the data being distorted, the input data is shown, the data from A.1 is used as the baseline. The reconstructed profiles, and 2-D phase space are also shown. Notice how the noise effects the reconstruction most near the edges of the separatrix. The separatrix is suppressed in the reconstructed image so that the high intensity pixles near the edge can be seen.

Because an iterative algorithm weights the pixels of the image by the amount of constructive interferrence, random noise, which should be uncorrelated from frame to frame, should be suppressed relative to the actual signal by a factor on the order of the number of frames of data. Essentially, each frame is an independent measurement of the distribution. But beacuase the signal represents a line integral along the energy coordinate of the phase space which means the far edges of the image suffer more from noise because each line integral is taken along a smaller number of pixles in the ΔE dimension becuase of the tapering of the separatrix along the edges.

A cut can be made as a function of action near the separatrix. The

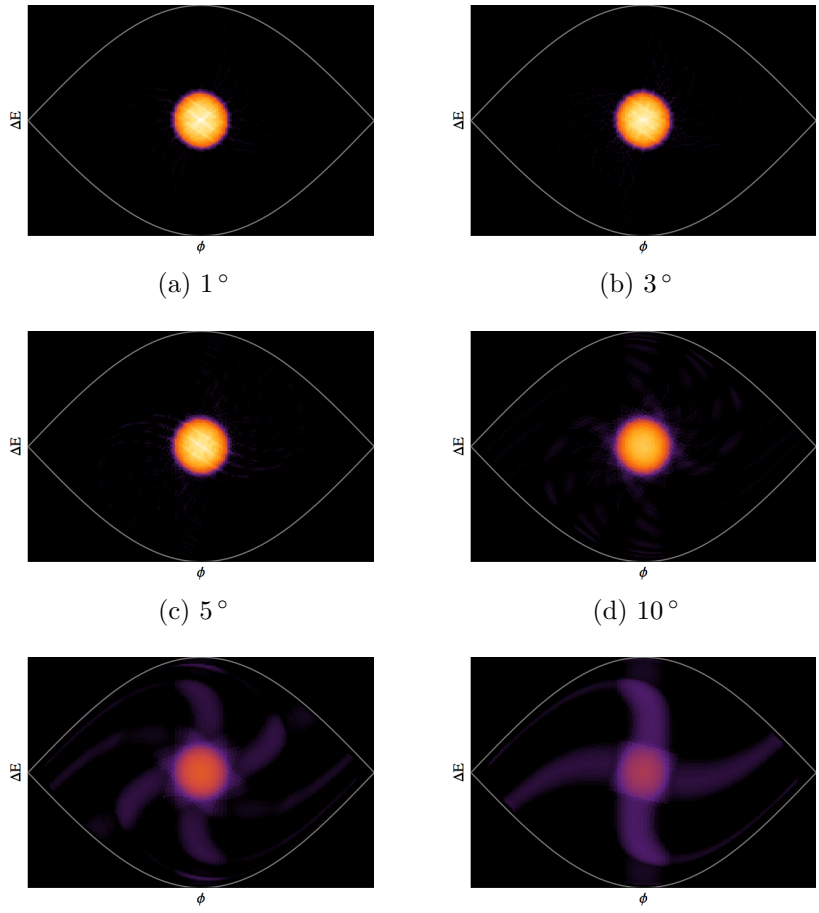

(a) 1 °

(b) 3 °

(c) 5 °

(d) 10 °

(e) 20 °

(f) 30 °

Figure A.3: The effects of insufficient angular resolution on 2-D $\phi - \Delta E$ reconstructions. Each reconstruction uses the same data spanning half a synchrotron period, but fewer frames of the data representing acquisitions at fewer, more largely spaced angles. The angular spacing of acquisitions is indicated beneath the reconstruction, relative to a synchrotron period. For instance, one acquisition at $t = 0$, and $t = \mathrm{T}_s/2$ would give an angle of 90 °.

reconstructed image with this cut is also shown in fig. A.4. The gray areas surrounding the main distribution show what has been cut. The remaining phase space shows some evidence of noise, but the main core of the distribution is still apparent. This noise can affect emittance measurements in particular.

A.4 Timing Offsets

Timing offsets due to triggering problems can lead to errors in reconstruction. There are two components to consider, a predictable jitter that comes from the interaction of the timebase on the scope and the timebase of the accelerator timing system, and a random jitter.

The non-random component occurs becuase a machine revolution is not a perfect multiple of the sample spacing on the scope. Over subsequent turns, the distribution can be seen to shift slightly from turn to turn, but not by more than a single sample. This error is simulated in fig. A.5. Notice in the input data how the error resest after a given number of turns. Because the magnitude of this error is typically small, the effect on recosntructions is also small, as can be seen in A.5. The bunch exhibits some irregularity around the edges of the distribution.

To control this error the sampling rate of the scope can be set higher until error is no longer visible. If setting the sampling rate higher increases the time to reconstruct, the number of pixels in the reconstruction can be reduced relative to the number of samples per bunch. This is especially helpful if the sampling rate is set to a value higher than the analog bandwidth of the scope,

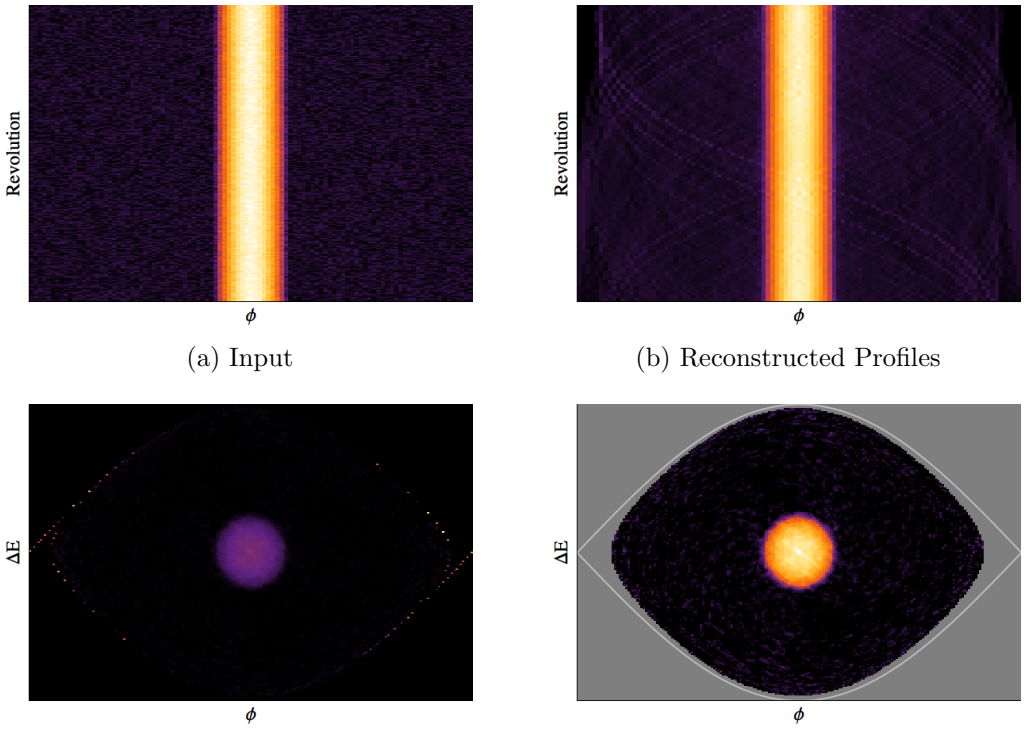

(a) Input

(b) Reconstructed Profiles

(c) Image

(d) Image with cut

Figure A.4: The effect of random noise in data on a tomographic reconstruction. (a) shows the data from fig. A.1 with added noise. (b) shows the reconstructed profiles. (c) is the initial reconstruced image with high intensity pixels near the separatrix. In (d) a cut has been applied removing pixles above an action just inside the separatrix, cut pixles are shown in gray.

or detector. See sec. A.7 for a discussion of pixel number and reconstruction.

A random jitter, much worse than what is typically observed, is shown in fig. A.6. Noise was modeled here as a randomly generated timing offset for each frame in the range of ±200ps and ±400ps full width. Actual jitter tends to be about 200ps, and only occurs occasionally. Without diagnosing the source of the random jitter in the electronics this component cannot be removed. A reconstruction of a simulated, perfectly matched bunch is shown in fig. A.6, with and without a single frame offset by the 200ps. Even with these very severe jitter models, the effect is merely to smear the jitter in the projections, and blur the edges of the distribution in the 2-D reconstruction. Some halo is also visible around the central bunch distribution. This will not affect most qualitative assesments, but will introduce a small error in emittance measurements.

A.5 DC Offset

A DC offset in the data contributes a serious error in the reconstruction. Part of the preprocessing is to remove any DC offset, but fig. A.7 shows the effect if one is present. The input data from fig. A.1 with a DC offset equal to 10% of the peak signal added. The reconstructed profiles, and image are shown as well. Notice how the same beam near the edges of the bucket appears as in the case of random noise.

A perfectly uniform projection tends to reconstruct as a halo concentrated at the separatrix, in practice this normally manifests as intense beam,

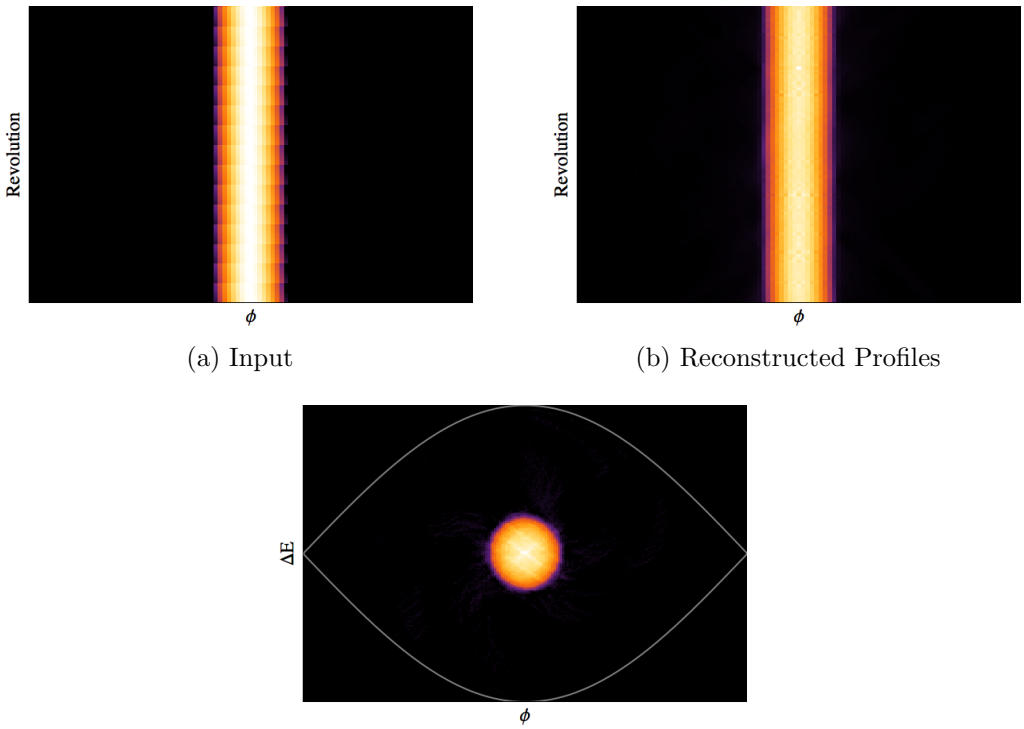

(a) Input

(b) Reconstructed Profiles

(c) Reconstruction

Figure A.5: The effect of mismatch between the scope timebase and the revolution period of the beam. The beam advances in phase by a fraction of a bin width on each turn becuase the timebase of the scope and the accelerator trigger are not perfectly matched. This is seen when the number of samples per bucket is not large enough. The resulting error is not large, but it does smear the bunch.

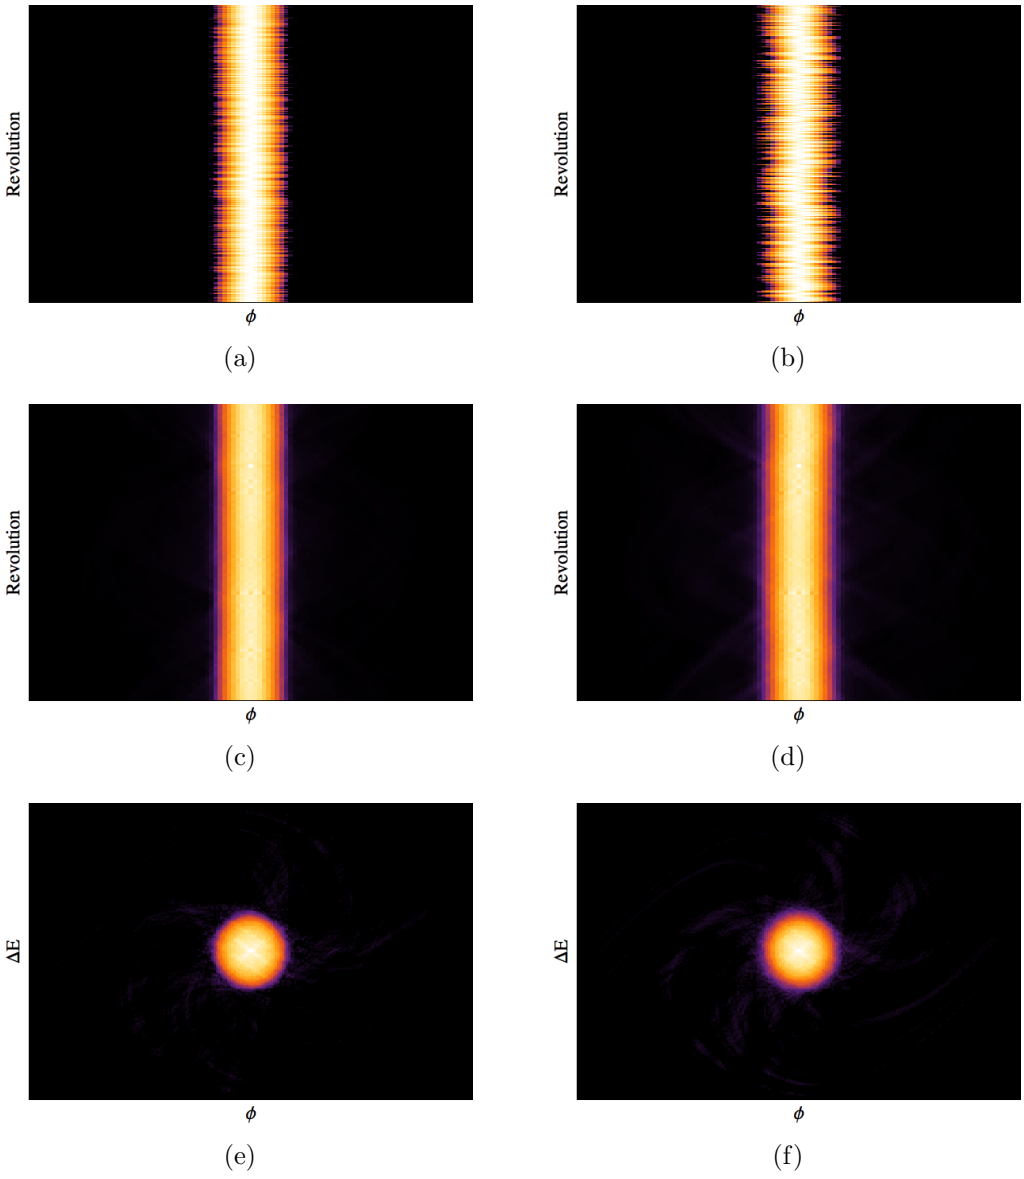

Figure A.6: The effect of random trigger jitter in traces on a tomographic reconstruction of those traces. Two cases of a severe jitter model are shown. The left column shows the input data, reconstructed projections, and 2-D phase space reconstruction for data generated by applying a ±200ps ($\approx 1\%$ of bunch width) on every frame. The right column shows the same information for a ±400ps jitter. The effect of the jitter is to blur the distribution and introduce a small amount of halo around the primary distribution.

usually much higher than any actual signal, around the edges of the bucket. A DC offset has a very similar effect to random noise. This is handled with a cut for pixels above a certain action, in the same was as the random noise case. Any DC offset present in the data should be removed during preprocessing though, so this should not be necessary.

A.6 Frequency Dependent Bunch Meausurement Errors

The bandwidth of the RWCM used a Fermilab is about $3\text{kHz} - 4\text{GHz}$, in addition to this there is a frequency dependent attenuation in the cables carrying the RWCM monitor from the tunnel to the digitizer [13], [35]. Each of these errors affects the shape of the traces used for reconstruction, but in grossly the same way.

The cable attuation can be reduced by using cables with lower attenuation per unit length over a large bandwidth, or reducing the length of cable between the detector and digitizer. The frequency response of the detector can also be measured and corrected for. [2] The impulse response of the RWCMs used at Fermilab in [5]. Fig. 4.9 shows a trace corrected for the detector response in a real case. Fig. A.8 shows the convolution of the detector response with the matched data used throughout this appendix. The shoulder introduced by the high frequency roll off is apparent. In this example, however, the

[2]This correction process is not perfect. At high enough frequencies the attenuation means that the signal will be completely lost, obviously a deconvolution can't restore this lost information.

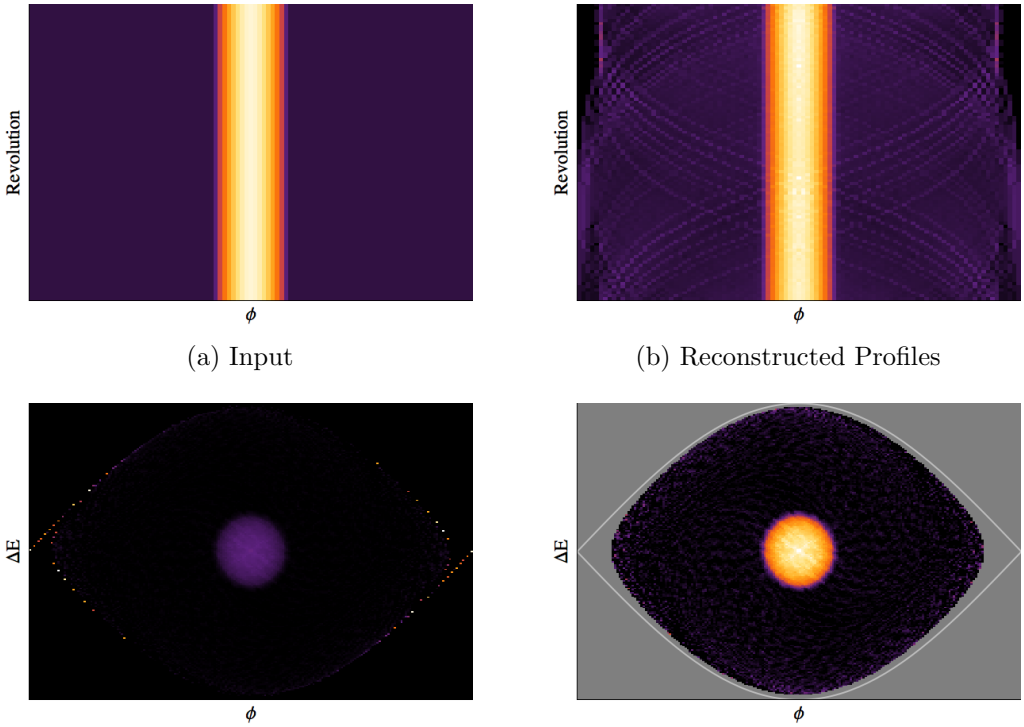

(a) Input

(b) Reconstructed Profiles

(c) Reconstructed Image

(d) Image with Cut

Figure A.7: The effect of data with a DC offset on a tomographic reconstruction. (a), and (b) show the data, with a DC offset equal to 10% of the peak signal, and a reconstruction respectively. Notice how the spurious points in (c) gather near the separatrix, which is suppressed to show the intense pixles around the edges. If this occurs, a cut can be made as a function of action to suppress the points near the edge, as in (d). The gray area in (d) indicates where the cut applies. With the cut on (d) the main distribution can be seen, but what looks like random noise is also present in the image.

shoulder is larger than usual because of the high frequency content needed to realize the sharp cut-off of the parabolic distribution at the full-width.

Fig. A.9 shows the input data used to evaluate the effect of detector response on reconstruction, the data is the trace from A.8 repeatd over many turns. The figure also shows the reconstructed profiles, and image. Qualitatively the effect on the reconstruction is quite small. This is largely due to the asymmetry in this effect. Becuase the shoulder is much more pronounced on the trailing edge of the bunch, it is not consistent with the model of longitudinal dynamics. The algorithm attempts to place some of this assymetric beam near the separatrix, where synchrotron motion is slowest. The effect we've seen many times of intense beam on the separatrix shows up again. Again a cut can be made for actions larger than some value near, but inside the sepatatrix. Figure (d) shows the reconstruction after this cut has been made, the gray region indicates phase space subject to the cut. After the cut there is some blurring of the distribution around the edges, but the reconstruction is still quite good. If the detector response is known, a correction should be attempted.

A.7 Discretization Artifacts

The first step in building the maps used to reconstruct a 2-D image of the phase space is to chop the phase space into pixels, or discretize it. Because the maps serve to transform the 1-D data collected at many times into the 2-D reconstruction, and the maps are constructed as a part of the reconstruction

process, the number of pixels in the 2-D image is arbitarary. (The dimensions of the reconstruciton in physical units will typically be limited to a single bucket width and height.) In ch. 3 I mentioned the difficulty in determining the optimal discretization of the phase space.

The sample rate used during acquisition sets a lower limit on the number of pixels in the time dimension equal to the number of samples per bucket, a number calculated during preprocessing as (sample rate/f_{RF}). Using fewer bins would represent a loss of precision. [3] Additionally, the discretization of the energy dimension is truly arbitrary, though it is sensible to use the same lower limit as for the time dimension. But what about upper limits?

Adding more bins means more test particles will be launched and maps will be larger, both increase the processing time. In addition to the increase in processing time artifacts can arise if the discretization is too fine. These two considerations mean we would like to set an upper limit as well as a lower limit, though this limit is less obvious. As a rule of thumb, similar to the rule for estimating the final digit of precision with any measurement device, the discretization can be set to twice the lower limit described here. (In TARDIS this is acheived by setting the number of pixels to -2. Where negative numbers, -n, indicate that $n\times$(samples per bucket), as determined by preprocessing, will

[3]This is actually a slight oversimplification. Sometimes the sample rate is set higher than is justified by the precision of the detector to combat a particular triggering problem encountered at Fermilab discussed in sec.A.4. In this case, the lower limit of the discretization may be resonably set lower than the number of samples per bunch. Other artifacts can arise if the number of image pixels in ϕ is not rational with the number of samples per bunch.

be used.)

A.8 Centering Error

Another possible error that can cause artifacts is the misidentification of the center of the synchrotron motion, which can be thought of as a small timing error. Fig. A.10 shows reconstructions, both the profiles and the phase space distribution, for several errors. (A negative error means the algorithm thinks the center is farther to the left of the image than it is.) The top row of phase space images share a common color range, so camparisons can be made, as do the profiles on the lower row. Notice how the profiles become wider, and the bunch shape distorts as the magnitude of the error increases. The effect of these small timing offsets are not large in terms of most operational concerns, but they can affect measurements made from the reconstructions, notably emittance measurements. They may also add a symmetry to the bunch that resembles a coupled bunch mode. Notice how the tails that develop around the central bunch can be used to determine which direction the offset is in. This is not always possible when the distribution is very clumpy, but it can offer a hint for manual corrections.

Typically finding the center of the synchtron motion is handled by the preprocessing in TARDIS, but the algorithm occasionally becomes confused, and one may wish to add an additional timing offset. This can be done in the configuration menu under 'Tomography'.

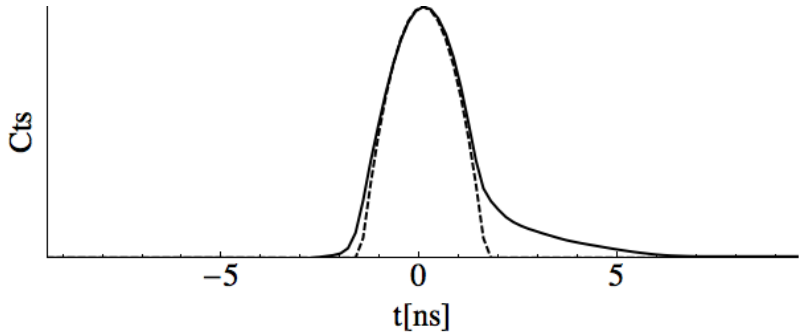

Figure A.8: A single frame of the baseline matched distribution time projection is convolved with the measured impulse response of the RWCM detector, as presented in [5]. The high frequency roll of of the detector above about 4GHz creates a shoulder seen here. This shoulder is more prounced than normal because of the high frequency content in the parabolic distribution, which ends abruptly at the full width. Fig. A.9 shows the effect fo this kind of error on reconstruction.

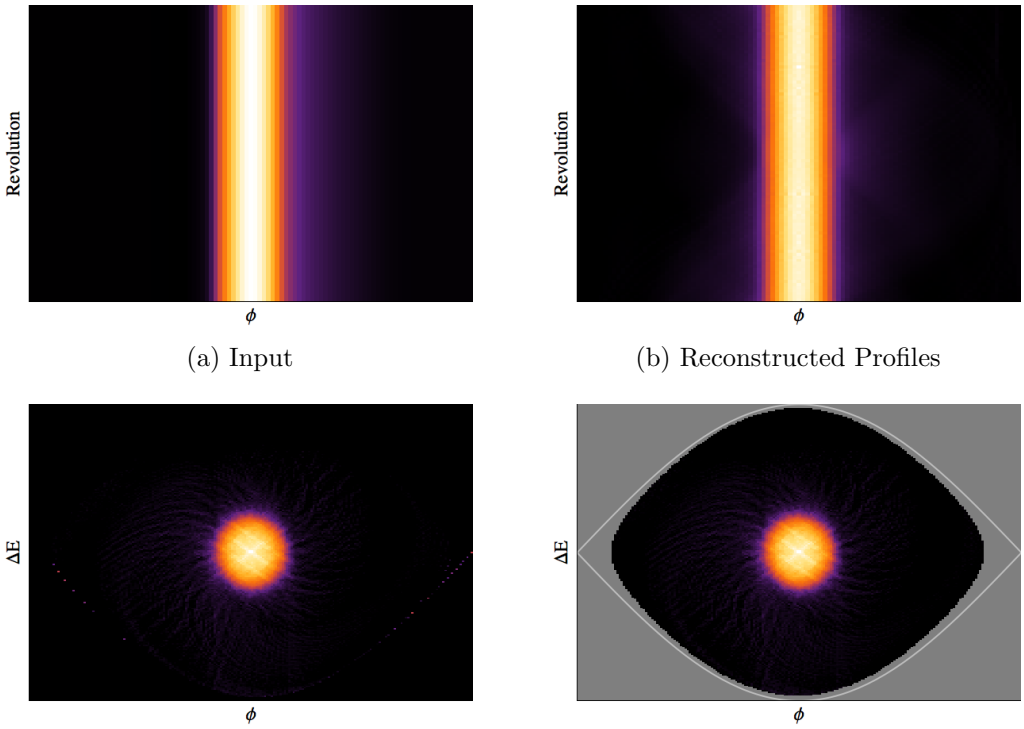

(a) Input

(b) Reconstructed Profiles

(c) Reconstructed Image

(d) Image with Cut

Figure A.9: The effect of detector impulse response on a tomographic reconstruction of those traces. (a) shows several frames of the trace from A.8 which provides the input data. (b) shows the projections reconstructed from (a). (c) and (d) show the 2-D reconstruction, and the 2-D reconstruction with a cut for action near the separatrix.

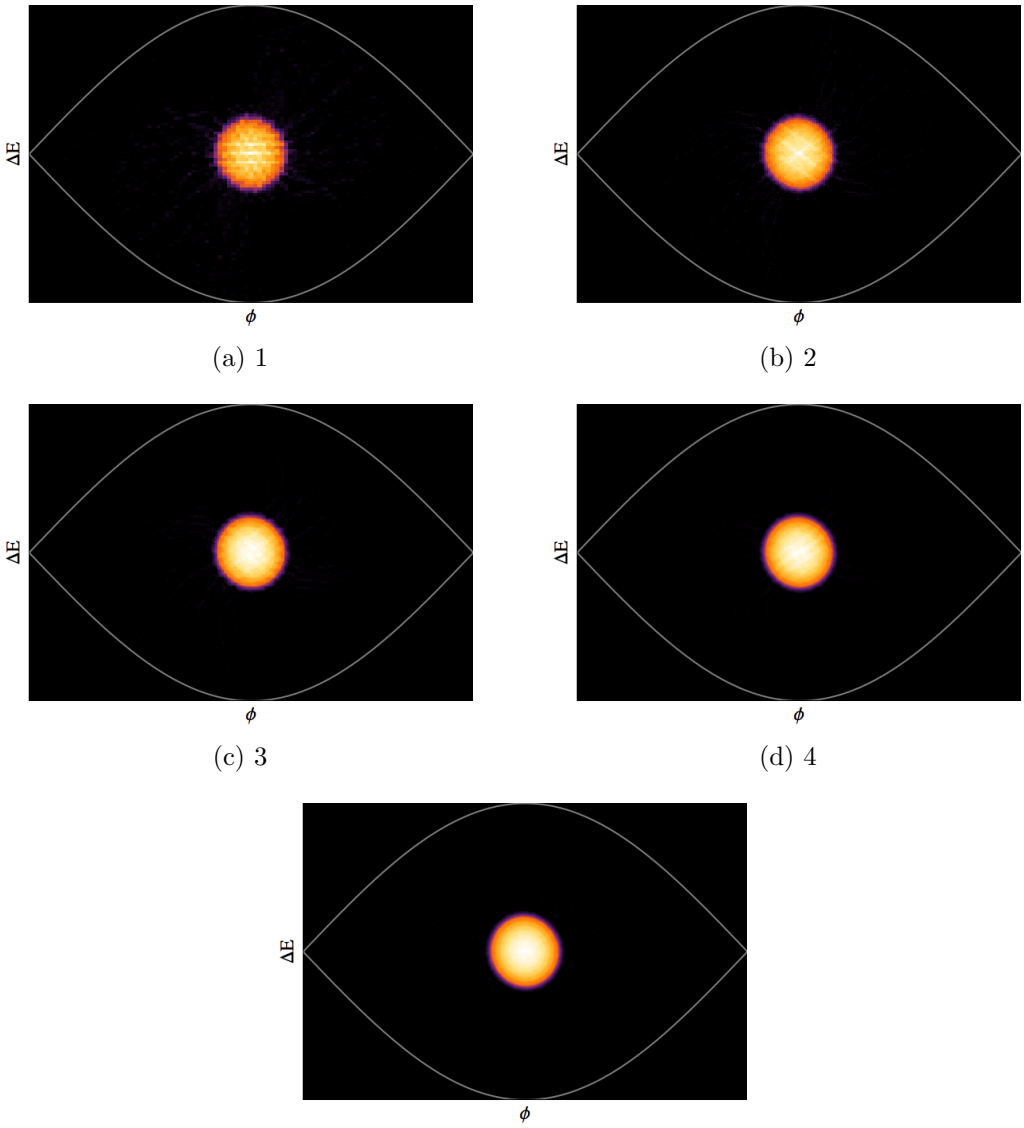

(a) 1 (b) 2

(c) 3 (d) 4

(e) 5

Figure A.10: The effect of discretization on reconstructions of the same data. For data with 100 samples per bunch, the number below each image is the multiplier used to decide the number of pixels in the ϕ, and ΔE dimensions. The total number of pixels in the image increases with this multiplier squared. The image becomes much clearer, but takes longer to calculate. The cross-hatching in the coarser images are artifacts, as are the rings seen in image (e). The rings are much more apparent in data that is not matched.

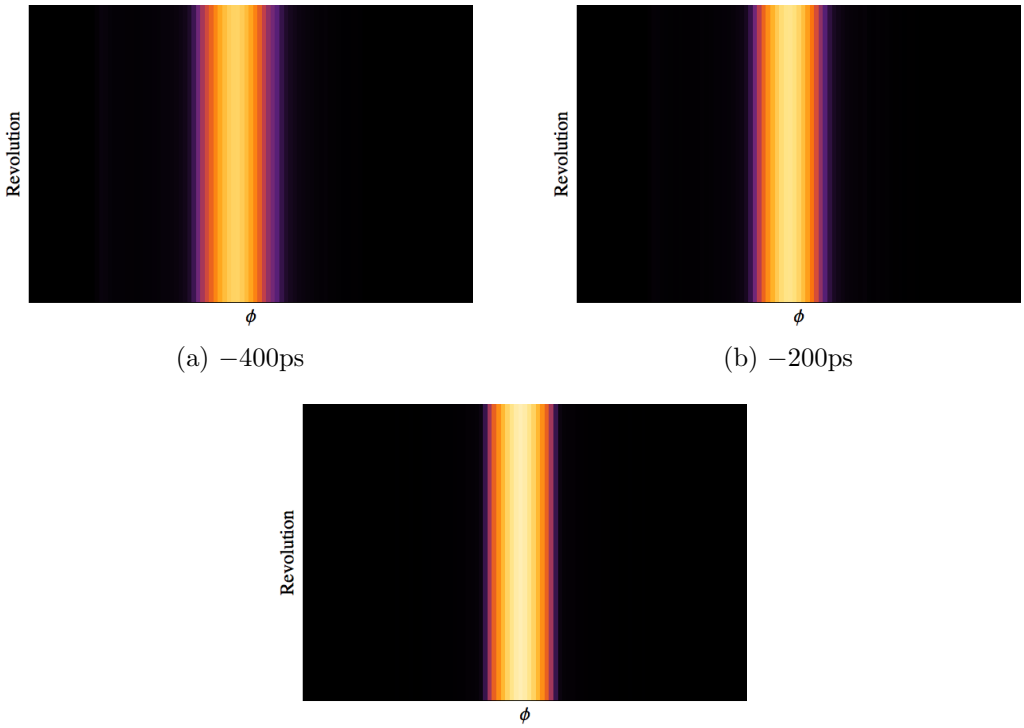

(a) −400ps (b) −200ps

(c) 0ps

(d) +200ps (e) +400ps

Figure A.11: The effect of a misidentified center on reconstruction. Each panel indicates an error introduced in locating the center of up to ±400ps (≈ 2% of a bucket), this is equivalent to a timing offset. Each recosntruction shown on the top shares a common range to allow comparison between reconstructions. Notice how the projections show an asymmetry that tends to smear in the direction of the actual center.

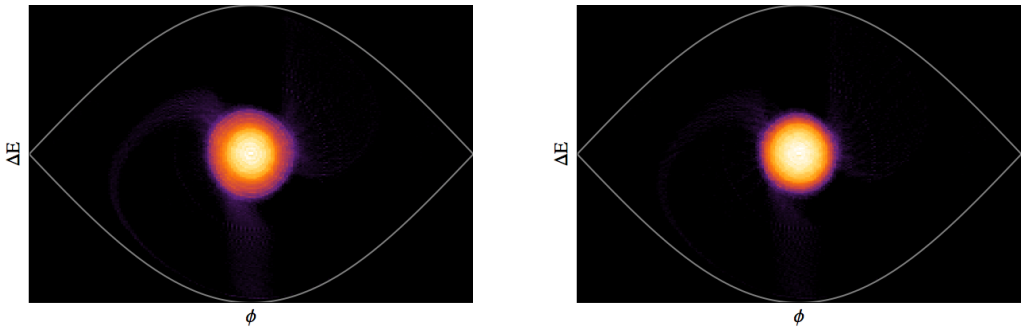

(a) −400ps (b) −200ps

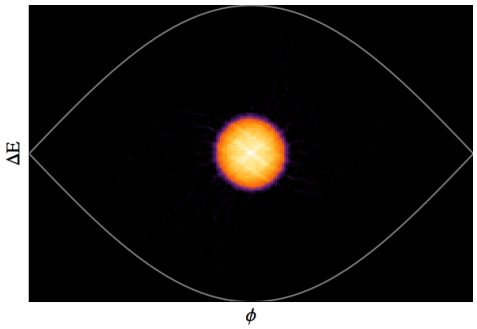

(c) 0ps

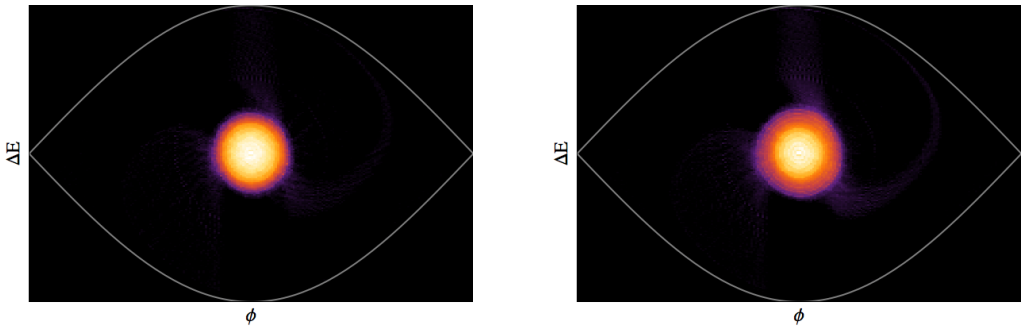

(d) +200ps (e) +400ps

Figure A.12: The effect of a misidentified center on reconstruction. Each panel indicates an error introduced in locating the center of up to ±400ps, this is equivalent to a timing offset. Each recosntruction shown on the top shares a common range to allow comparison between reconstructions. Notice how the bunch shape distorts and the intensity near the center drops. The effect is not large when looking by eye, but can have an effect on emittance measurements, increasing emittance as the bunch is spread out to larger action.

Bibliography

[1] *Signal Processing For Longitudinal Parameters Of The Tevatron Beam*, 2005.

[2] *Momentum Spread Reduction at Beam Extraction from the Fermilab Booster at Slipstacking Injection to the Main Injector*, 2007.

[3] Measurement of CMS Luminosity. 2010.

[4] C. Ankenbrandt. Slip Stacking: A New Method of Momentum Stacking. 1981.

[5] M. Backfish. MiniBooNE Resistive Wall Current Monitor. Technical Report FERMILAB-TM-2556-AD, Fermi National Accelerator Laboratory, 2013.

[6] C. Behrens, F. J. Decker, Y. Ding, V. A. Dolgashev, J. Frisch, Z. Huang, P. Krejcik, H. Loos, A. Lutman, T. J. Maxwell, J. Turner, J. Wang, M. H. Wang, J. Welch, and J. Wu. Few-femtosecond time-resolved measurements of x-ray free-electron lasers. *Nat Commun*, 5, 04 2014.

[7] C. Bhat. Main Injector Intensity Limitations at Transition Crossing, 1997.

[8] D. Boussard and Y. Mizumachi. PRODUCTION OF BEAMS WITH HIGH LINE DENSITY BY AZIMUTHAL COMBINATION OF BUNCHES IN A SYNCHROTRON. (TALK). *IEEE Trans.Nucl.Sci.*, 26:3623–3625, 1979.

[9] Alex Chao. *Physics of Collective Beam Instabilities in High Energy Accelerators*. Wiley, 1993.

[10] Alexander Wu Chao and Maury Tigner. *Handbook of Accelerator Physics and Engineering*. World Scientific, Singapore, 1999.

[11] The Pierre Auger Collaboration. Correlation of the highest-energy cosmic rays with nearby extragalactic objects. *Science*, 318(5852):938–943, 2007.

[12] Donald Edwards and Michael Syphers. *Physics of High Energy Accelerators*. Wiley, 1993.

[13] Brian Fellenz, , Brian Fellenz, and Jim Crisp. An improved resistive wall monitor.

[14] R Gordon. Artifacts in reconstructions made from a few projections. pages 275–285, 1973.

[15] Richard Gordon. A tutorial on art (algebraic reconstruction techniques). *IEEE Trans. Nucl. Sci. NS-21*, (3):78–93, 1974.

[16] Richard Gordon, Robert Bender, and Gabor T Herman. Algebraic reconstruction techniques (art) for three-dimensional electron microscopy and x-ray photography. *Journal of theoretical Biology*, 29(3):471–481, 1970.

[17] S. Hancock. Iterative tomography. https://tomograp.web.cern.ch/tomograp/overview/iterative-tomography.html. Accessed: July 23, 2014.

[18] S Hancock, ME Angoletta, and A Findlay. LEIR RF Voltage Calibration using Phase Space Tomography. Nov 2010.

[19] S Hancock, A Jansson, and M Lindroos. Tomographic reconstruction of transverse phase space from turn-by-turn profile data. (CERN-PS-99-029-OP):6 p, Apr 1999.

[20] S Hancock, P Knaus, and M Lindroos. Tomographic Measurements of Longitudinal Phase Space Density; 1998 ed. (CERN-PS-98-030-RF):4 p, Jul 1998.

[21] S. Hancock, M. Lindroos, and S. Koscielniak. Longitudinal phase space tomography with space charge. *Phys. Rev. ST Accel. Beams*, 3:124202, Dec 2000.

[22] S Hancock, M Lindroos, E McIntosh, and M Metcalf. Tomographic Measurements of Longitudinal Phase Space Density; 1999 ed. Technical Report CERN-PS-99-002-OP. CERN-IT-99-001, CERN, Geneva, Jan 1999.

[23] G Jackson. A phase space tomography (PST) monitor for adjusting bunch rotation during coalescing. Technical Report FERMILAB-FN-469. FN-469, FERMILAB, Batavia, IL, 1987.

[24] G. Jackson. Design, implementation, and results from a longitudinal phase space tomography (pst) monitor in the fermilab main ring. In *Particle Accelerator Conference, 1993., Proceedings of the 1993*, pages 2418–2419 vol.3, May 1993.

[25] S. Kaczmarz. Angenäherte auflösung von systemen linearer gleichungen. *Bulletin International de l'Académie Polonaise des Sciences et des Lettres*, 35:355–357, 1937.

[26] Ioanis Kourbannis, Keith Meisner, and King-Yuen Ng. Experimental study of the main ring transition crossing. Technical report, Fermi National Accelerator Laboratory, 1990.

[27] S. Y. Lee. *Accelerator Physics*. World Scientific, 2nd edition, 2004.

[28] P. Limon. The optimum energy of the really large hadron collider sited at Fermilab. *eConf*, C960625:ACC048, 1996.

[29] King-Yuen Ng. Exact Solutions for the Longitudinal and Transverse Impedances of an Off Centered Beam in a Rectangular Beam Pipe. *Part.Accel.*, 16:63, 1984.

[30] King-Yuen Ng. Some estimation concerning crossing transition of the main injector. Technical report, Fermi National Accelerator Laboratory, 1990.

[31] Wolfgang K.H. Panofsky. Evolution of particle accelerators and colliders. *SLAC Beam Line*, 27N1:36–44, 1999.

[32] Johann Radon. On the determination of functions from their integral values along certain manifolds. *Medical Imaging, IEEE Transactions on*, 5(4):170–176, Dec 1986.

[33] Frank James Sacherer. Methods for computing bunched-beam instabilities. Technical report, CM-P00063598, 1972.

[34] D.J. Scott, D. Capista, B. Chase, J. Dye, I. Kourbanis, et al. Coalescing at 8 GeV in the Fermilab Main Injector. *Conf.Proc.*, C1205201:3779–3781, 2012.

[35] Duncan J Scott. Distortion in resistive wall current monitor signal transmission lines. Technical Report Beams-doc-4436-v3, Fermi National Accelerator Laboratory, August 2013.

[36] J. P. Shan, S.G. Peggs, and S.A Borgacz. Analytical evaluation of the second order momentum compaction factor and comparison with mad results. In *Particle Accelerator Conference, 1993., Proceedings of the 1993*, pages 74–76 vol.1, May 1993.

[37] Albert Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, 2005.

[38] R Thurman-Keup, C Bhat, W Blokland, J Crisp, N Eddy, B Fellenz, R Flora, A Hahn, S Hansen, T Kiper, A Para, S Pordes, and A V Tollestrup. Longitudinal bunch monitoring at the fermilab tevatron and

main injector synchrotrons. *Journal of Instrumentation*, 6(10):T10004, 2011.

[39] Jie Wei. *Longitudinal dymaics of the nonadiabatic regime on alternating-gradient synchrotrons*. PhD thesis, State University of New York at Stony Brook, 1989.

[40] Jie Wei. Transition crossing in the main injector. Technical report, Brookhaven National Laboratory, 1990.

[41] Xi Yang, Charles M. Ankenbrandt, Rene Padilla, William A. Pellico, Joseph E. Dey, and Kiyomi Koba. Reducing the momentum spread of 8-GeV proton beam via the bunch rotation in Booster. 2005.