CESR AT LOW ENERGY

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Robin Bjorkquist

May 2019

CESR AT LOW ENERGY

Robin Bjorkquist, Ph.D.

Cornell University 2019

Since its completion in 1979, the Cornell Electron Storage Ring (CESR) has been used to store electrons and positrons ranging in energy from 2 GeV to 6 GeV. Researchers at Cornell are now preparing to use CESR to demonstrate a new beam cooling technique called optical stochastic cooling (OSC). The OSC experiment will require 1 GeV electrons, well below the lowest energy previously stored in CESR. In this dissertation, I recount some of the challenges of low-energy operation and the results of initial attempts to store 1.5 and 1.0 GeV electrons in CESR.

BIOGRAPHICAL SKETCH

Robin Bjorkquist was born and raised in Portland, Oregon. She graduated from Reed College with a BA in physics in 2009, then took a staff position as the Associate Director of the Reed Research Reactor. She began her graduate studies at Cornell University in 2011.

For everyone who is drawn to physics

but has not yet found a home there.

ACKNOWLEDGEMENTS

I learned a great deal from each of my graduate research advisors, Lawrence Gibbons and David Rubin — not only about the doing-of-science, but also about the kind of human being, teacher and mentor I hope to become.

Although it is not the subject of this dissertation, I spent 80% of my time in graduate school working on the Muon $g-2$ experiment, a small particle physics experiment located at Fermilab. I am grateful to my Muon $g-2$ colleagues and friends for the companionship of those years, and most especially Team Fiber Harp. You're the best!

These have not been easy years for me, and I have benefitted tremendously from the wisdom, counsel, and sympathetic ears of many mentors and friends, including Jim Alexander, Janet Conrad, Renee Fatemi, Joe Grange, Fred Gray, David Griffiths, David Rubin, Darrell Schroeter, Erica Snider and Ron Zaraza. Each of you made me believe (at least once) that physics might have a place for me after all. Thank you!

I am grateful to everyone who gave me the opportunity to participate in science outreach during my time in graduate school, including Kaleigh Muller, Lora Hine and Erik Herman. My life is richer because of you.

I also gained some truly excellent friends at Cornell, especially Jennifer Chu, Pratiti Deb, Wee Hao Ng and Veronica Pillar. Much love to you all!

I have at various times been supported by fellowships from Cornell University, the National Science Foundation Graduate Research Fellowship Program (DGE-1650441) and the Fermilab URA Visiting Scholars Program. I've also been paid as a graduate research assistant through grants from the Department of Energy (DE-SC0008037) and the National Science Foundation (PHY-1734189).

TABLE OF CONTENTS

PREFACE

My dear readers: Hello. Welcome! It turns out that the Cornell University Graduate School provides very few hard-and-fast rules regarding the formatting, length, style, or content of a PhD dissertation — but they do recommend that it "*conform to the standards of leading academic journals in your field.*"

Yeah. About that. Let me tell you a secret (shhh!): I don't particularly enjoy reading academic physics journals. Some other activities I mostly-don't-enjoy: physics seminars, research group meetings, graduate-level physics classes, reading impenetrable textbooks, and (especially!) the interminable days of jargon-filled presentations that one must endure if one dutifully attends a science conference.

Maybe you're thinking, "*My god, Robin! Why on earth did you just spend eight years of your life in graduate school if you don't like any of those things?!?*" What can I say? It might have been a huge mistake.

But here are some things I *do* enjoy: I love interacting with students. I love being able to grasp the *how* and *why* of an experiment. I love talking about science with non-scientists (which I only get to do by virtue of *being a scientist*). I love asking and answering questions. I love seeing people's eyes light up when I tell them about my work. And, every now and then, I encounter a physics idea that fills me with delight.

I wanted very much to write a story that would contain some of that delight. Needless to say, this will not sound like a journal article — it will sound like me. I also hope my story will be accessible (at some level) to *everyone*, whether or not you have a background in physics. Have I succeeded? I don't know. If you start reading, but then find yourself becoming bored or confused, know this: I don't blame you! I myself fell in love with physics-in-person, not physics-on-the-page. Ask me sometime, and I'll tell you this story myself.

CHAPTER 1

FIRST LOOK AT A BEAM MEASUREMENT

I thought about starting with the introduction. My "Introduction to CESR" (now moved to Chapter 2) contains some of the things you'll need to know about the Cornell Electron Storage Ring (CESR) in order to understand what follows: its dimensions and layout, an account of the most important components, some details about the physics of particle beams. Delightful stuff, really.

However, I sometimes find it useful to jump into the middle of an experiment right from the start — to get a taste for the measurements that are involved, before then circling back to fill in all the details.

It's 10:00 pm on a Tuesday night. Half a dozen people are crowded into the CESR control room. After about 30 hours of effort (spread over several shifts), we've finally managed to inject & store a measurable number of 1.5 GeV electrons into CESR, but we remain puzzled by some of the features we see in our data. I'd like to start by showing you one of those measurements; it will be the anchor for the rest of the story.

1.1 Machine Studies

Six days a week, CESR operates as an x-ray generating machine. The storage ring runs through a circular tunnel, 40 feet below one of Cornell's sports fields. We fill the ring with bunches of particles: electrons going one way, positrons the other way. When the particles pass through special magnets called *undulators*, they produce beams of x-rays. This facility is called CHESS (the Cornell High Energy Synchrotron Source), and users come from all over the world to stick their samples into the x-ray beam lines for various kinds of imaging experiments.

On Tuesdays, however, we do *machine studies.* This can mean a lot of different things. Blocks of time are set aside for activities such as testing and calibrating instrumentation, developing new software tools, and investigating and correcting problems that have arisen in the operation of the accelerator. We also use CESR as a *test accelerator* to try out new accelerator-physics ideas and techniques. It is in this context that we are trying to *lower* the energy of the stored electrons to 1.5 GeV, and ultimately all the way down to 1.0 GeV.

Tu 4/24/18

- 00:00 – 07:00 – CHESS operations
- 07:00 – 08:00 – Characterization (SBP)
- 08:00 – 15:00 – Down / Access
- 15:00 – 16:00 – Reset Interlocks
- 16:00 – 17:00 – Test Fill (MJF)
- 17:00 – 18:00 – RF beamkiller diagnostics (MJF)
- 17:00 – 18:00 – Linac Section 7 tuning (GWC, JJR)
- 18:00 – 21:00 – OSC devel (SW, MPE, WFB, RAB)
- 21:00 – 00:00 – Phase measurement characterization (JSh)

Figure 1.1: Machine studies schedule for 4/24/2018.

Figure 1.1 shows the schedule for the Tuesday in question, 4/24/2018. Normal CHESS operations end at 7:00 am every Tuesday morning and resume 24 hours later. During the day shift on Tuesday (8:00 am to 4:00 pm), the accelerator is shut down, so that we can access the tunnel for maintenance. Between 4:00 pm and midnight, blocks of machine studies time are allocated to individuals or groups to work on various projects, depending on current needs. Finally, the time between midnight and 7:00 am is used to restore CHESS conditions and allow the machine to fully equilibrate before the x-ray users resume their experiments on Wednesday.

The measurement I'm going to show you was collected during the "OSC devel" time slot. OSC stands for *optical stochastic cooling*; that's the experiment that will require 1.0 GeV electrons.

1.2 Beam Position Monitors

When you're tuning the machine and trying to store a beam, it helps to be able to see what the particles are doing inside the beam pipe. Enter the *beam position monitors* (BPMs) — there are about 100 of these devices inside CESR, spaced approximately every 7.7 meters = 25 feet. Figure 1.2 shows a cross-sectional diagram of one BPM.

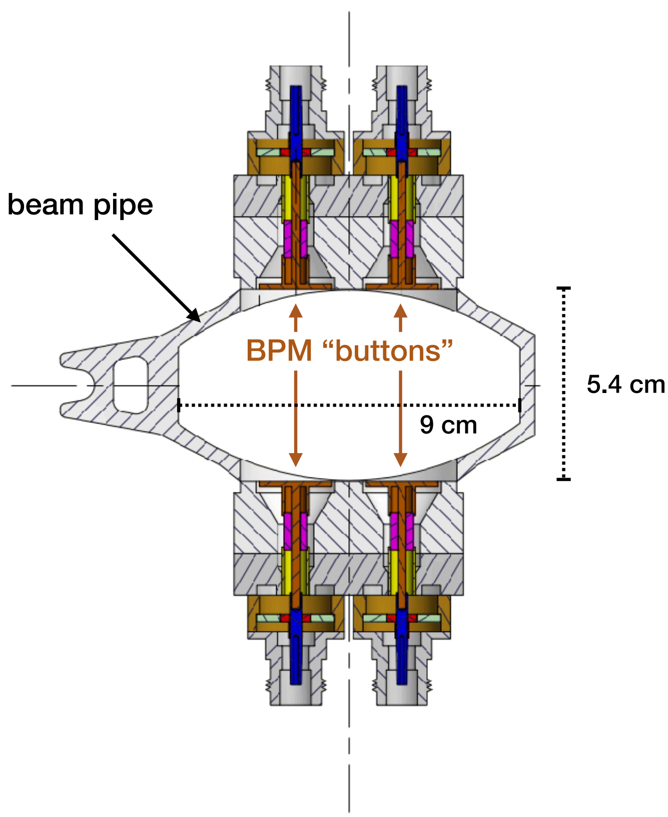

Figure 1.2: Beam position monitor.

The big opening at the center is the beam pipe; it has metal walls, and the inside is evacuated so the particle beam can travel freely without hitting too many air molecules (vacuum systems are essential for particle accelerators!). Take note of the dimensions: the beam pipe itself isn't terribly large — I could easily encircle it with my two hands, so long as I found a location with no other equipment or electronics getting in the way.

Now, imagine a little bunch of electrons whizzing by, straight into the page, traveling at very nearly the speed of light. After 2.5 microseconds, the same particles will be back again, having made a full lap around the storage ring. Ideally, the electrons will pass through at the exact center of the beam pipe — but things are rarely ideal. The job of the BPM is to tell us the horizontal and vertical position of the beam each time it passes this location inside the storage ring.

The BPM consists of 4 electrodes (we call them *buttons*) plus some readout electronics. When a charged particle bunch goes past, it has an electrical influence on each of the metal buttons, resulting in a voltage pulse that can be amplified and captured by the readout electronics. The size of the voltage pulse has to do with (a) the total charge of the particle bunch and (b) how close the particle bunch is to the button in question. Hopefully it sounds reasonable, then, that we can take the four button readings and compute the x (horizontal) and y (vertical) position of the particle bunch — because that's exactly what we do.

1.3 Some Beam Position Measurements

As promised, Figure 1.3 shows some BPM readings from the night of Tuesday 4/24/2018.

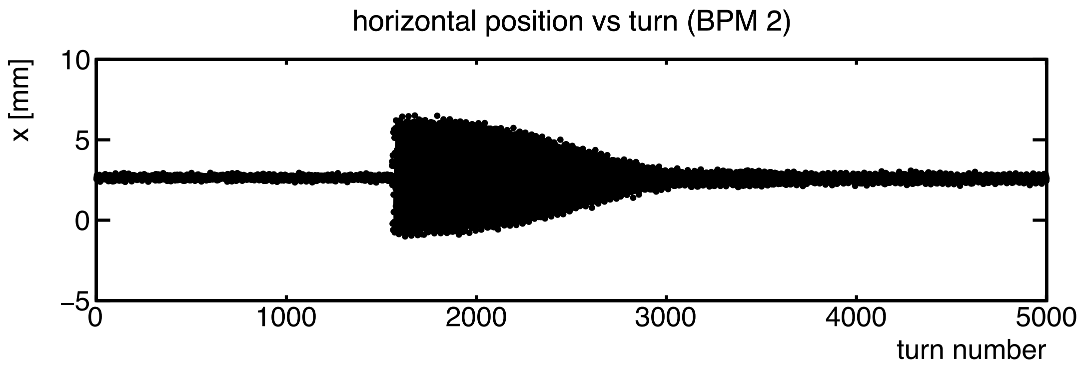

Figure 1.3: Sample BPM measurements, showing the horizontal position of the electron bunch at one location in CESR over the course of 5000 turns.

These measurements are from a single location in the storage ring. At that moment, we had approximately 10^9 electrons stored in CESR. I'm showing you 5000 measurements, as that bunch of electrons circled the ring 5000 times (which takes only a fraction of a second!). On this scale, the individual data points are all squished together; all you can really see is the *envelope* of the positions. That's enough, though, to immediately see that something dramatic happened around turn number 1500.

Yes it did. Something happened at that time because we *made it happen* — we used some pulsed magnets in the storage ring to give the beam a nudge and push it out of equilibrium.

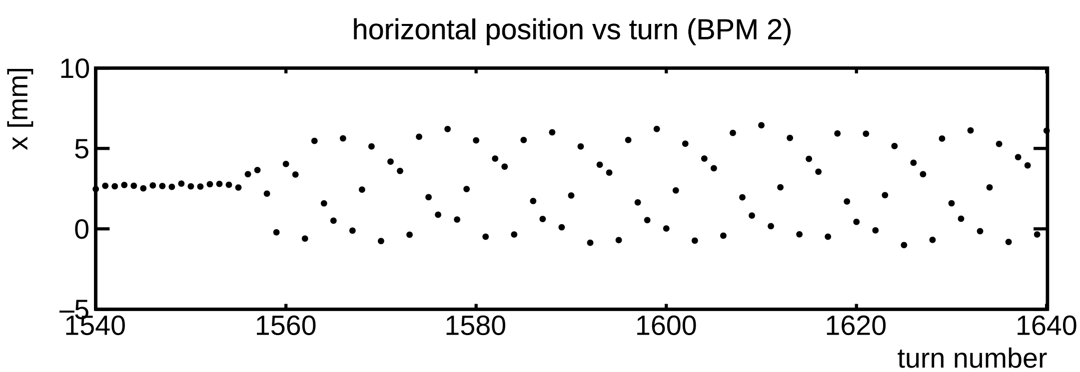

Figure 1.4: A closer view of the position measurements.

Let's take a closer look (Figure 1.4). Before turn number 1556, the beam is at equilibrium: every time it passes this BPM module, it is at the *same* horizontal position, about 2.5 mm away from the center (the small variations can be attributed to measurement uncertainty).

After we nudge the beam with our pulsed magnets, the position measurements are all over the place, as much as 4 mm to either side. Without any lines to guide the eye, it's hard to tell what's going on — but Figure 1.5 show the same plot again, now with the addition of a sinusoidal function.

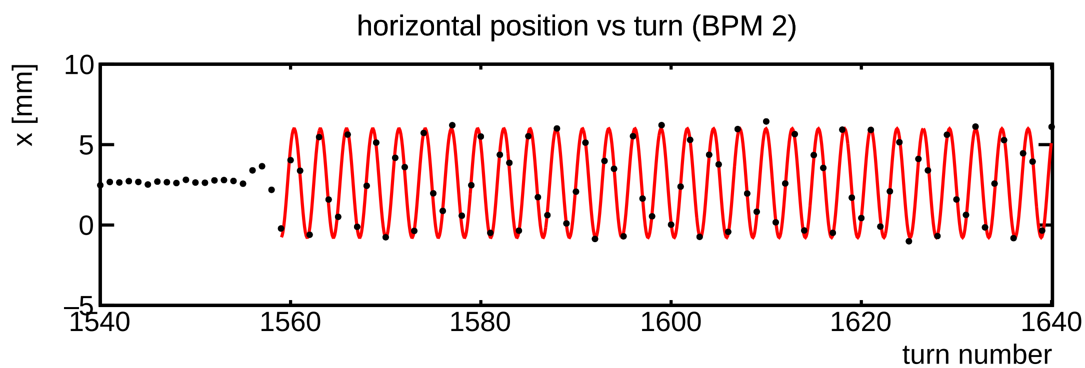

Figure 1.5: Sinusoidal fit to the measurements.

Now we can see that the beam is, in fact, *oscillating* around the equilibrium position, with a wavelength of about 2.8 turns. For a physicist, even a physicist who knows nothing about accelerators, this is not a strange result. After all, this is a *storage ring*, and any storage-thingy worth its salt *has* to be able to deal with small deviations from equilibrium — and if I've learned anything from my physics education, it's that a system close to a stable equilibrium will undergo sinusoidal oscillations.

The frequency of these oscillations — in this example, one cycle per 2.8 turns = 0.36 — is called the horizontal *tune* of the storage ring. If we had given the beam a vertical nudge (resulting in vertical oscillations of the beam position), we could have measured the vertical tune instead. The tunes are characteristics of the accelerator layout and all the magnet strengths; by measuring the tunes, we can verify that the accelerator is set up as intended (and if not, make the appropriate adjustments).

Our objective on 4/24/2018 was simply to check the tunes, but we noticed a puzzling feature in the data. If you look back at Figure 1.3, you can see that the initial oscillations die away over the course of some hundreds of turns. I will explain this phenomenon in Chapter 3. For now, I will just say that while we expected the oscillations to die away eventually, we thought they would last *much* longer than a thousand turns. This puzzle — and its resolution — provide some clues that might help to explain why we've had so much trouble lowering the energy of CESR to 1.0 GeV.

CHAPTER 2

INTRODUCTION TO CESR

2.1 Layout

CESR is a *storage ring*, which means it is designed to hold a beam of high-energy particles. CESR can store electrons (traveling counter-clockwise), positrons (traveling clockwise), or, under the right circumstances, simultaneous counter-rotating beams of electrons and positrons.

By necessity, the storage ring must be paired with a particle accelerator, which produces the high-energy particles in the first place. The layout of the Cornell accelerator is shown in Figure 2.1.

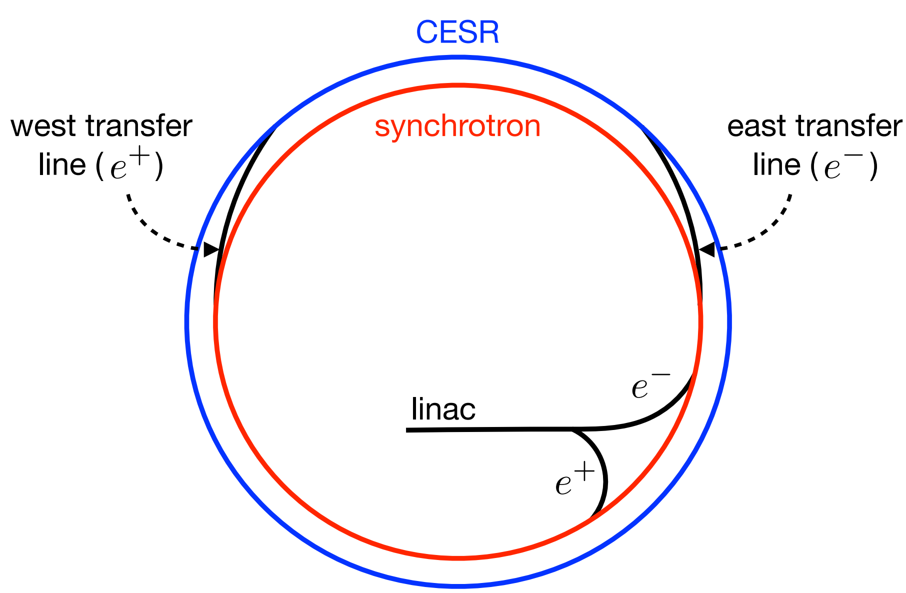

Figure 2.1: Layout of the Cornell accelerator.

The electron beam is created by an electron gun at the beginning of the linear accelerator ("linac"), which brings the particles up to an energy of 300 MeV. At that point, the electrons are transferred to the circular synchrotron, which ramps them up to their final energy over the course of some 3000 turns (about 7.5 ms). The synchrotron was designed as a 10 GeV machine, but can be adjusted to produce lower-energy particles as well. The synchrotron shares a tunnel with CESR — the synchrotron on the inner wall and CESR on the outer wall, as shown in Figure 2.2. After the particles complete their acceleration cycle in the synchrotron, they are transferred across the tunnel to CESR where they continue to circulate without gaining any more energy.

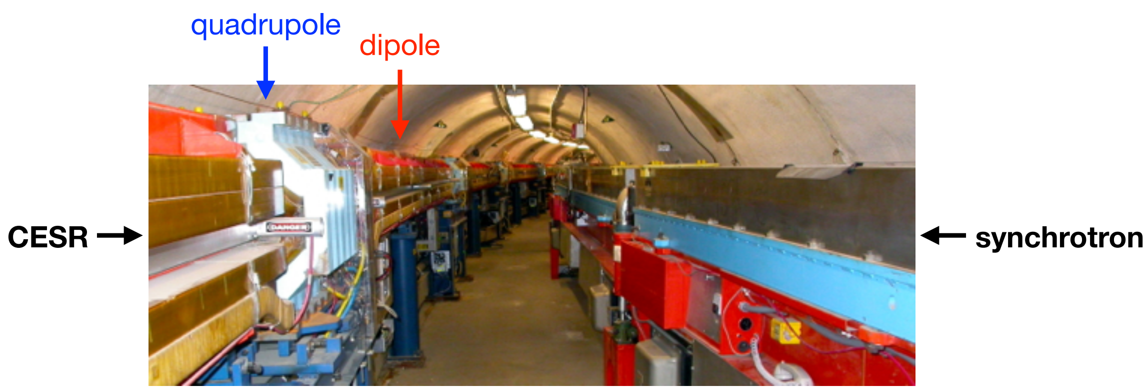

Figure 2.2: Photograph of the accelerator tunnel. CESR is on the left side of the tunnel and the synchrotron is on the right.

To create a positron beam, the process begins the same way, with the electron gun creating a beam of electrons. The electrons travel about halfway down the linac (gaining energy as they go), then smash into a tungsten target. The interaction of the electrons with the target produces some positrons, which can then be accelerated through the remaining half of the linac and injected clockwise into the synchrotron.

2.2 History

Cornell has been at the forefront of accelerator science almost from the time of the first particle accelerator. I quite enjoyed the historical account in Reference [1], so if you'd like to read a *much* more thorough version of this story, I warmly recommend that you look there.

The 10 GeV Cornell synchrotron was completed in 1967. This was actually the fourth Cornell synchrotron, following on the heels of earlier, smaller machines designed for 300 MeV (1949), 1.3 GeV (1952) and 2.2 GeV (1963). As you can see, the name of the game in those years was pushing the energy as high as possible, as quickly as possible, driven by a desire to access new energy regimes for particle physics experiments.

CESR was added to the Cornell accelerator complex in 1979. For the next three decades (1979 – 2008), CESR was home to CLEO, a large electron-positron collision experiment.

2.3 Magnets

Magnets are the true heart of CESR. If you learn nothing else about accelerators from me, I want you to know — really *know*, deep in your gut — what I mean when I say "magnet" in this context. So, let's talk about accelerator magnets!

I love magnets. I don't know about you, but when I hear the word *magnet*, I first think of the items stuck to my refrigerator. For instance, I have a set of magnetic physics words sent to me by the American Physical Society (Figure 2.3).

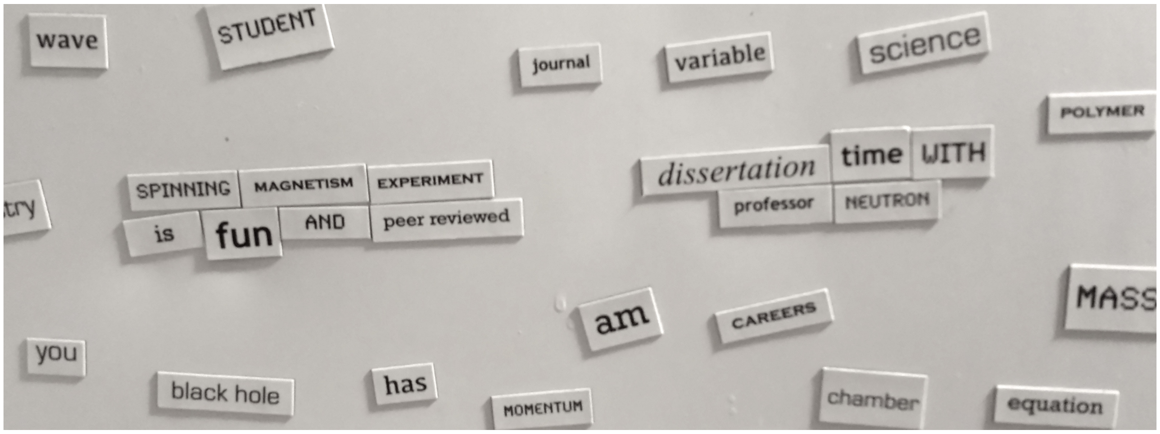

Figure 2.3: My refrigerator.

Even more fun than flat refrigerator magnets are the chunky magnets I played with as a child (and, okay, I admit it — also as an adult). I sincerely hope you have also done this. If not, I implore you: get yourself some magnets and play with them! Go ahead. I'll wait.

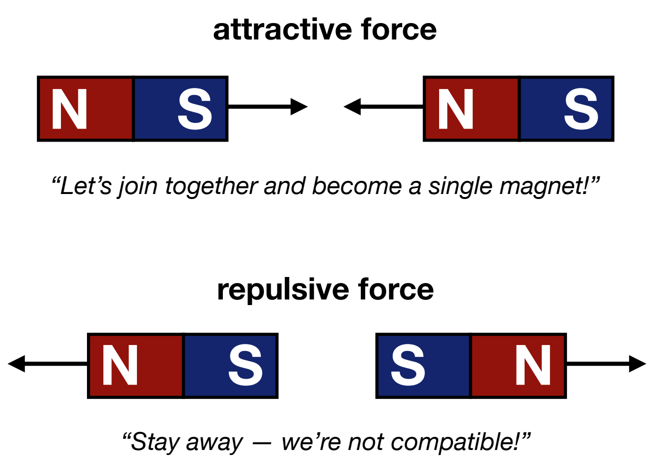

Figure 2.4: Opposite poles attract. Like poles repel.

Having played with magnets, you know from experience that the two opposite ends of a bar magnet are distinct and distinguishable. Depending on their relative orientation, two magnets might stick together or they might push apart. By convention, we call these two ends the north and south *poles* of the magnet. Opposite poles attract and like poles repel, as illustrated in Figure 2.4.

One reason I love playing with magnets is that you can *feel* the magnetic forces, even when the magnets themselves aren't touching. It becomes viscerally clear that the influence of a magnetic object extends outward through space, well beyond the edges of the object itself. Physicists describe this extended magnetic influence in terms of the *magnetic field*, which we can visualize by drawing magnetic field lines, like those shown in Figure 2.5. Each line emerges from the north pole of the magnet and loops around to return through the south pole.

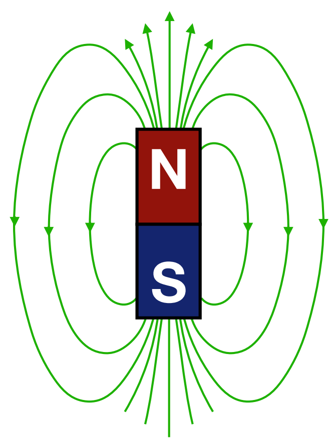

Figure 2.5: The magnetic field of a bar magnet.

The magnetic field is a *vector field*; at every location, the field has both a magnitude and a direction. You can get a sense for both of those aspects by looking a field diagram: the field direction is always tangent to the curved field lines (in the direction of the arrow), and the field is stronger where the lines are closer together.

Now, in an accelerator, the point isn't for the magnets to interact with each other, but instead for each magnet to have an influence on the motion of a charged particle passing nearby, as suggested in Figure 2.6. This is great, because the

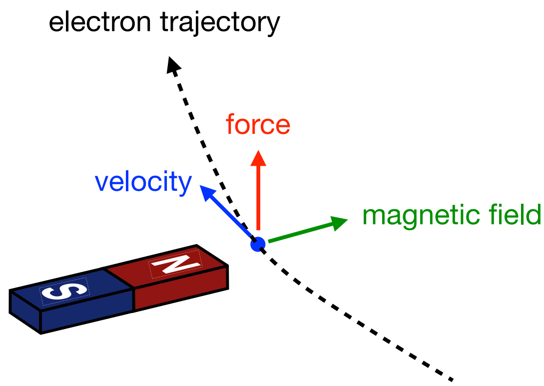

Figure 2.6: The magnetic force on a moving charged particle.

physics of a charged particle passing through a magnetic field is really very simple. Here it is: the magnetic force on the particle is

$$\vec{F} = q\vec{v} \times \vec{B}, \tag{2.1}$$

where q is the particle's electric charge, $\vec{v}$ is its velocity and $\vec{B}$ is the magnetic field at the location of the particle. The cross product in the equation mathematically encodes this fundamental truth about magnetic forces: the magnetic force is always perpendicular to the particle's direction of motion. That means the magnetic force can't change the energy of the particle, only its direction. In an accelerator, *electric fields* are needed to do the actual acceleration; magnetic fields are used for steering, focusing, and other corrections.

There's one more thing I need to tell you about magnets-in-general before delving into the specifics of some common accelerator magnets. A bar magnet is an example of a *permanent magnet*. It's a chunk of magnetized material that produces a magnetic field. However, you can also use electric currents to produce magnetic fields. A magnet that uses an electric current as the source of its magnetic field is called an *electromagnet*.

By winding a current-carrying wire into a coil, you can create a magnet with the same field shape as a bar magnet, as shown in Figure 2.7. You can also add an iron or steel core to the electromagnet, to help shape and strengthen the field.

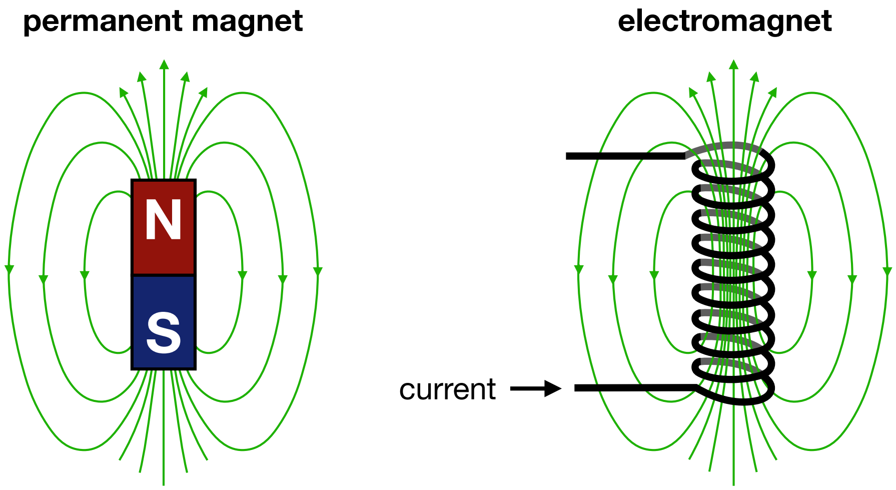

Figure 2.7: Permanent magnet and electromagnet.

Most of the magnets in CESR are electromagnets. The poles are made of steel with current loops wrapped around them. Electromagnets are generally more versatile and convenient than permanent magnets, because the strength of the magnet is adjustable — all you have to do is change the amount of current running through the coil. You can even swap the polarity of the magnet by reversing the direction of the current.

An accelerator magnet is a hefty device. Walking through the accelerator tunnel (Figure 2.2), one gets the sense that CESR is a bulky machine — but it's really the magnets that are bulky, thick layers of steel and aluminum surrounding a small beam pipe.

Next I want to tell you about two specific kinds of accelerator magnets: *dipole magnets* and *quadrupole magnets*.

2.3.1 Dipole Magnets

In the simplest interpretation of the word, a *dipole* is a just a magnet with two poles. The bar magnets we were talking about earlier? Those are dipoles. In an accelerator, however, "dipole" has a more specific meaning. First, the magnet is arranged so that the two poles face each other across a narrow gap, as shown in Figure 2.8.

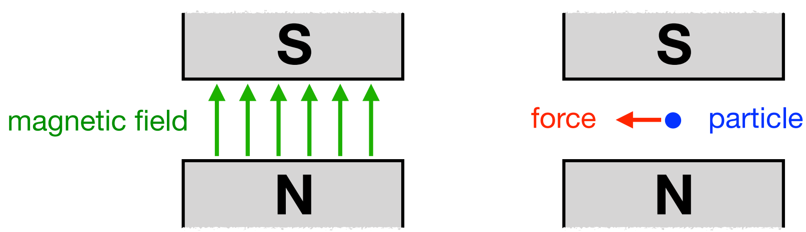

Figure 2.8: Dipole magnet.

Second, an accelerator dipole is designed in such a way that the magnetic field is uniform in middle of the gap, where the particles travel. This means (at least roughly speaking) that all the particles in the beam experience the *same* magnetic force as they pass through the dipole. Dipoles are steering magnets; their job is to bend the beam around a circular path.

CESR has 82 dipole magnets, collectively responsible for turning the beam through the 360° needed to make a full circuit of the ring. CESR has several types of dipole, each with its own specifications. The most common (what I'll call the "standard" CESR dipole) deflects the beam through an angle of 4.29 degrees.

Of course, magnet poles don't just float in the air as I've drawn them above. Zooming out a bit, Figure 2.9 shows a more complete view of a CESR dipole. In this cross section, you can see that the north and south pole are in fact connected (phew!). This is a "C-shaped" magnet, with the steel wrapping around on one

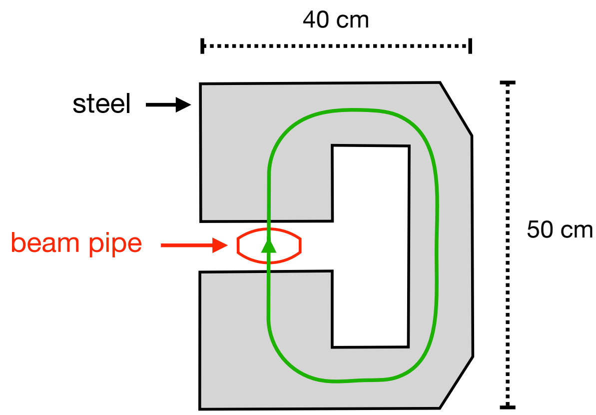

Figure 2.9: Dipole cross section.

side, creating a return path for the magnetic field lines; the other side is open, allowing access to the beam pipe. In this picture, electrons travel into the page (counterclockwise around CESR) and positrons travel out of the page (clockwise around CESR).

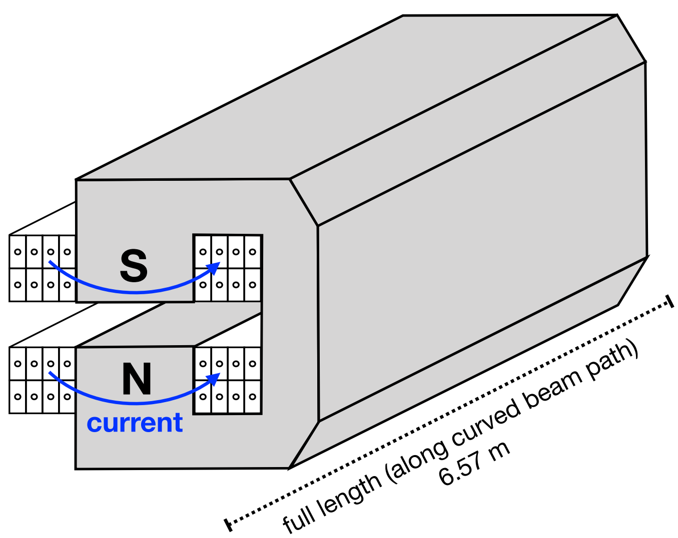

Figure 2.10: 3D view of a dipole.

Each of the standard CESR dipoles is 6.57 meters long (nearly 4 times as long as I am tall). That's how much magnet is required to turn the beam by just a few degrees — you could do it with a shorter dipole, but then would need a stronger magnetic field (everything is about tradeoffs!). Figure 2.10 is one more view of the standard CESR dipole, now showing the current loops as well. Each pole of the magnet is wrapped with eight loops of current-carrying aluminum. These conductors aren't so much "wire" as they are hefty aluminum bars. Cooling water flows through a channel in the center of each conductor.

There are accelerator physicists and engineers who specialize in magnet design. Even a dipole — the simplest of accelerator magnets — has a lot of fiddly details that need to be worked out. (What shape should the steel be to ensure field uniformity? How much cooling is required for the electrical coils? What should be done at the entrance and exit of the magnet so that it will have the desired properties?) When you consider that there are other types of magnets that are *much* more complicated (such as the undulators that wiggle the beam back and forth to produce x-rays for CHESS), you can start to see how a person might spend many years engrossed in these devices.

2.3.2 Quadrupole Magnets

As you might have guessed, a quadrupole has four poles, as shown in Figure 2.11

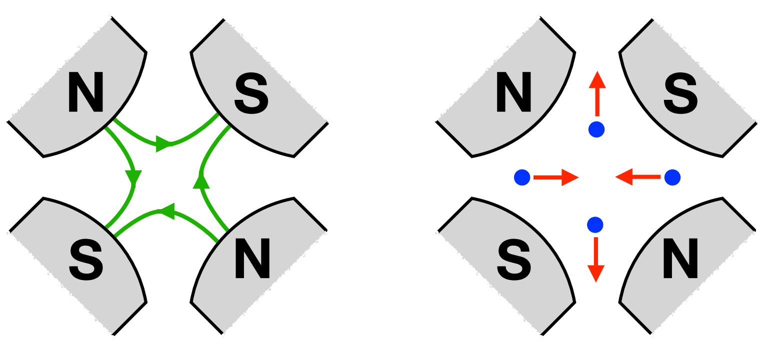

Figure 2.11: Quadrupole magnet.

Whereas a dipole has a uniform magnetic field, the field inside a quadrupole depends on location. If a particle passes through the very center of the magnet (at $x = 0$ and $y = 0$), it will experience no magnetic field at all. However, if the particle is off-center, it *will* encounter a field — and the farther it is from the center, the stronger that field will be.

The *direction* of the field also depends on location. For the configuration of north and south poles I've drawn above, the field orientation is such that a particle passing through with an x-offset will experience a magnetic force that pushes it back toward $x = 0$ (a *restoring force*), while a particle with a y-offset will experience a force that pushes it even farther away from $y = 0$.[1] In other words, this quadrupole is horizontally focusing but vertically defocusing. If you swap the north and south poles (which you can do by reversing the direction of the current through all the coils of the electromagnet), you'll get a magnet with the opposite effect: vertically focusing but horizontally defocusing.

In equations, the magnetic field inside a quadrupole is

$$B_x = k\left(\frac{P_0}{q}\right) y$$
$$B_y = k\left(\frac{P_0}{q}\right) x.$$

Note that the horizontal field depends linearly on y and the vertical field depends linearly on x. Because of the right-angle nature of magnetic forces, this actually means that x and y are *decoupled*: the horizontal force depends only on x and the vertical force depends only on y. The proportionality constant, $k(P_0/q)$, is split up into three pieces for convenience and clarity: P_0 and q are characteristics of the *particles* (P_0 is the reference momentum and q is the electric charge), while k (the

[1] Just as before, in CESR this could either be an electron going into the page or a positron coming out of the page — the result is the same for either particle species.

quadrupole strength) describes the size of the quadrupole's influence on the beam.

By convention, k is positive for a horizontally focusing quadrupole and negative for a horizontally defocusing quadrupole.[2]

It makes sense that P_0 and q appear in the field equations. If you increase the momentum of the particles, but meanwhile wish to keep the quadrupole strength k of the magnet unchanged, then you must strengthen the magnetic field — this is because a high-momentum beam is "stiffer" than a low-momentum beam; it takes a larger force to deflect higher-momentum particles by the same amount. Similarly, if you had some way to increase the electric charge of each particle, then you would need to reduce the magnetic field in order to maintain the same particle motion (because the magnetic force is proportional to q).

The quadrupole strength k has units of $1/\mathrm{m}^2$. Together with the length L of the magnet, it determines how much a particle is deflected when it passes through the quadrupole. I'll show you by working out the horizontal deflection: referring back to equation (2.1), we can evaluate the cross product and find that the horizontal magnetic force on the particle is

$$F_x = -qvB_y = -vkP_0x.$$

Here I assumed that the velocity v of the particle is entirely in the longitudinal z-direction. This is a reasonable approximation, because the transverse x- and y-components of the motion are always very small compared to the primary longitudinal motion.

If the quadrupole has length L, the time it takes a particle to pass through the

[2] I admit that I've been somewhat vague about conventions for signs and coordinate directions. That's because those things are always confusing as hell. If you ever find yourself in a position where you need to get it right, just remember: set it up so that positive k is horizontally focusing!

magnet is $\Delta t = L/v$. During this time, the particle's horizontal momentum will change by

$$\Delta P_x = F_x \Delta t = -LkP_0 x.$$

Assuming the particle is perfectly on-momentum ($P_{\text{total}} = P_0$), the horizontal angle of the particle's trajectory relative to the reference trajectory is $x' \approx P_x/P_0$, and that angle will change by

$$\Delta x' \approx \frac{\Delta P_x}{P_0} = -(Lk)x. \tag{2.2}$$

The deflection is proportional to x, but otherwise depends only on L and k, just as I promised.[3]

All this talk of focusing and deflection is reminiscent of another familiar object: the optical lens. Consider this: when parallel rays of light enter a simple converging lens, the lens focuses the light so that all the rays meet at a single point. The distance from the lens to this crossing point is called the *focal length* (f) of the lens.

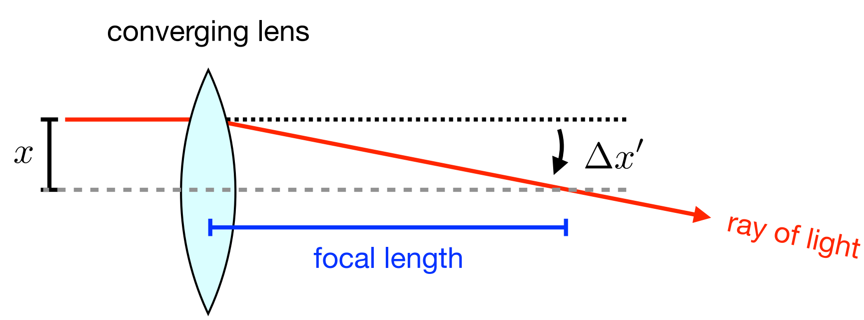

Figure 2.12: Converging lens.

[3]You might be thinking "Hey, wait a minute... if the beam is being deflected, then x must be changing! Which x are you supposed to use in equation (2.2)?" The answer: strictly speaking, it should be the average x along the length of the quadrupole — but if the quadrupole is short enough, it's probably okay to use the initial x. (This is akin to the "thin lens approximation" in optics.)

For this to work out, the deflection of each ray *must* be proportional to its initial displacement, as illustrated in Figure 2.12. Using our accelerator-language, that means the lens is designed to cause

$$\Delta x' = -\frac{x}{f}$$

and (drumroll please) *this is exactly the same as equation* (2.2) — provided that we identify the focal length of the quadrupole as

$$f = \frac{1}{kL}.$$

We see now that quadrupole magnets are truly the lenses of the accelerator. Their job is to focus the beam, thereby keeping the particles contained inside the beam pipe.

There is one major difference between optical lenses and quadrupole magnets. An optical lens is isotropic: it does the same thing in all transverse directions. The lens can be converging or diverging, but only one or the other. A quadrupole magnet, on the other hand, is *not* isotropic: if x is converging, then y is diverging, and vice versa. Luckily, if you alternate the two types of quadrupoles, the overall effect is to focus the beam in both transverse directions.

CESR has about 100 quadrupoles. There are a few different kinds (especially after the 2018 upgrade to the south arc), but the “standard” CESR quadrupole is 60 cm long, with the cross section shown in Figure 2.13. As you can see, the beam pipe fits snugly in between the four poles. The other open spaces in the diagram are actually filled by the conducting coils that wrap around each pole.

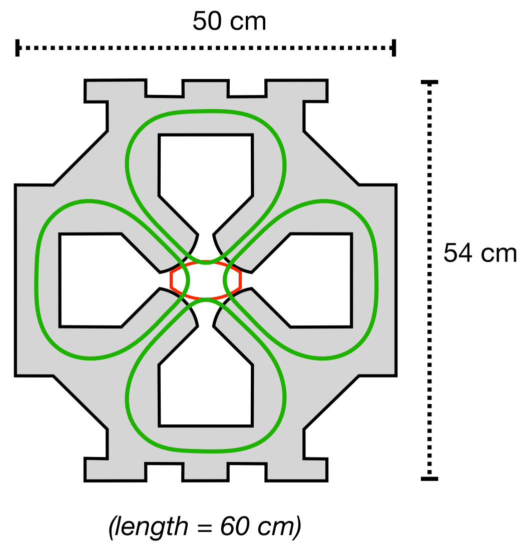

Figure 2.13: Quadrupole cross section.

2.4 Synchrotron Radiation

When a charged particle accelerates, it emits electromagnetic radiation. Figure 2.14 is a cartoon showing how that plays out in CESR. This radiation ends up being an important contributor to the characteristics of the stored beam.

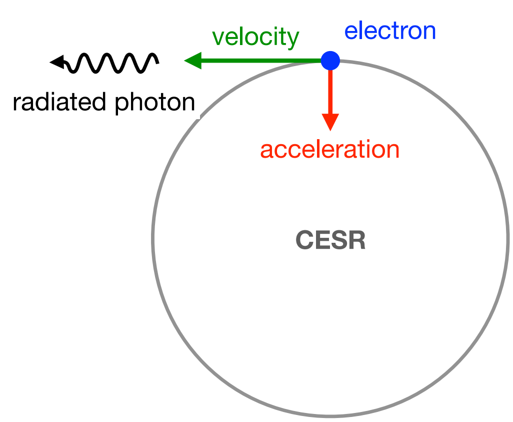

Figure 2.14: Synchrotron radiation.

Particles in CESR travel with (approximately) constant speed and energy. However, because the electrons continually change direction as they travel around the ring, they are accelerating. The acceleration vector points toward the center of the ring. In situations like this, where the velocity and acceleration are perpendicular to each other, the emitted radiation is called *synchrotron radiation*.

For an ultra-relativistic electron (moving at nearly the speed of light), the emitted radiation is concentrated in the forward direction. For my example of 1.5 GeV electrons, the relativistic gamma factor is $\gamma = 2900$, and the radiation is mostly contained within an opening angle of $1/\gamma = 0.0003$ radians $= 0.02$ degrees. That is a *very small* angle, so for most practical purposes, we can forget about it and imagine that the radiated photon is emitted precisely in the forward direction.

The total amount of energy is conserved, so when an electron radiates a photon (typically an x-ray), it loses a bit of its own energy. Hence, even though CESR is a storage ring and not an accelerator, it does need to include a mechanism to restore that lost energy. Otherwise, the storage ring would be useless — within not-very-much-time, the stored electrons would lose so much energy they would crash into the wall of the beam pipe (never a good outcome!).

2.5 Accelerating Cavities

The devices responsible for restoring the energy lost to synchrotron radiation are called *accelerating cavities* or *RF cavities*. CESR has four of these cavities, and each one is a complicated superconducting device (Figure 2.15).

CESR's accelerating cavities bear some similarity to the accelerating structures in the linac and synchrotron, but also important differences. In all cases, however, the end goal is to create a longitudinal electric field in the beam pipe, because it

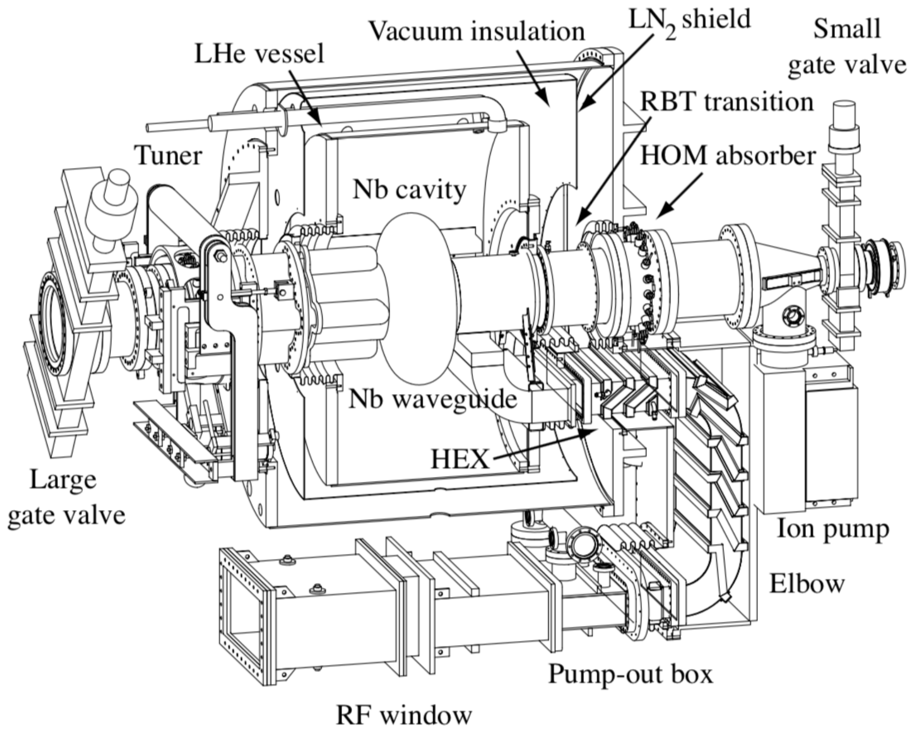

Figure 2.15: CESR accelerating cavity (from Reference [2]).

is the associated electric force that will give energy to the charged particles that make up the beam.

Static (unchanging) electric fields can't do the job, because one of the properties of static electric fields is that they are *conservative*. When a particle makes a full circuit of the ring (returning to the same place where it started), the total change in energy due to the influence of a conservative field is always zero — which is pretty useless if your goal is to add energy to the particles! To get around this limitation, accelerating cavities actually produce oscillating electric fields (a big part of why the machinery is so complicated), and the whole system is carefully timed so that the particles will pass through at a moment when the electric field is pointing in the appropriate direction for a forward force. This is also the reason why the accelerator has to have bunches of particles, instead of a continuous stream.

2.6 Energy Limits

Both the synchrotron and CESR can be adjusted for a range of particle energies. To change the energy, the most important requirement is that we adjust all the magnet strengths accordingly. Doubling the particle energy requires doubling all the magnetic fields; halving the particle energy requires halving all the magnetic fields.[4]

In general, the upper limit for the energy of a circular accelerator or storage ring will be set by one of two possible limiting factors: (1) the maximum strength of the bending magnets or (2) the ability of the RF cavities to restore the energy lost to synchrotron radiation.

Making an accelerator work properly at very *low* energy is a different problem entirely. I like to think about this issue in analogy to riding a bicycle. First, imagine riding a bike as fast as you can go — there is, of course, a limit: your legs (like the RF cavities) can only pedal so hard, and the drag from wind resistance (like synchrotron radiation) becomes stronger as your velocity (energy) increases. Now imagine riding a bike as slow as you can go. That is *also* a difficult task, but for very different reasons! A slow bicycle is wobbly and hard to control. It takes concentration and careful balance to keep from falling over.

So, what's "wobbly" about a low-energy accelerator? A few things come to mind:

[4] In the synchrotron, we actually have a choice: when we want to lower the output energy, we can either lower the magnetic fields as I just described *or* we can extract the particles early, before they reach the maximum energy set by the fields!

1. An electromagnet that is designed to operate at a particular current might have a poor field quality when you turn it down low. In other words, the difference between the *ideal* magnetic field and the *actual* magnetic field can become more significant — which could affect the particle dynamics in the ring.
2. Low-energy particles are more easily scattered than high-energy particles (off the thin beryllium window they need to pass through on their way into CESR, off residual gas molecules in the beam pipe, and especially off each other).
3. The amount of energy lost to synchrotron radiation depends sharply on the electron energy — a process that also affects the particle dynamics in the ring, as I will discuss in Chapter 3.

CHAPTER 3

INITIAL ATTEMPTS TO STORE LOW-ENERGY ELECTRONS

3.1 Spring 2018 Machine Studies

In spring 2018, we spent several machine studies shifts working to store low-energy electrons in CESR. The target for the OSC experiment is 1.0 GeV, but we approached that goal gradually — first storing a 2.1 GeV beam (which has been done before), then moving down to 1.5 GeV and finally 1.0 GeV.

We had some success at 1.5 GeV, but only up to a point; the injection efficiency was low, and we couldn't accumulate more than a small bunch. Working with 1.0 GeV electrons was even *harder* — we were barely able to accumulate a beam at all.

It turns out, though, that the measurement I showed you in Chapter 1 contains a clue that might help make this job easier; for the rest of this chapter, I'd like to explore that measurement in detail.

3.2 Back to Our Sample Measurement

Once again, Figure 3.1 is our beam measurement, recorded while we were attempting to store and accumulate 1.5 GeV electrons in CESR. Let me remind you of what this is. When we collected these data, there was a single bunch of electrons stored in CESR. The bunch contains about a billion electrons, built up bit by bit over many injection cycles. The physical size of the bunch depends on the settings of the accelerator, but might typically be something like 2 cm long, 0.5 mm wide

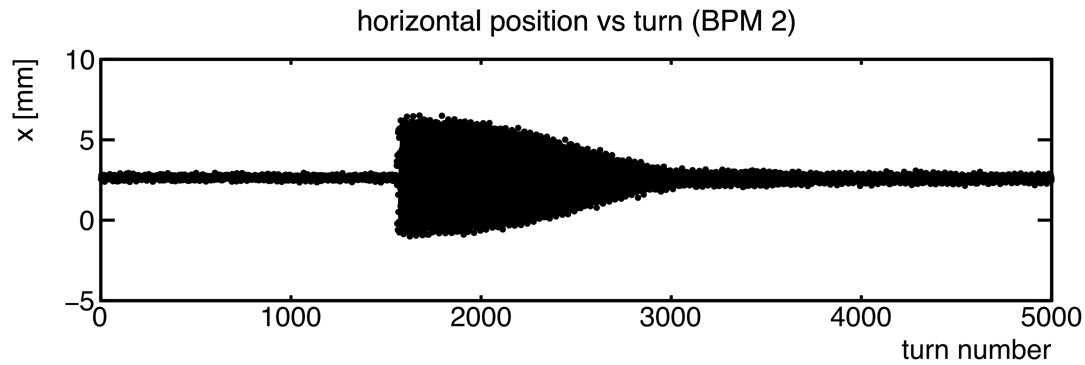

Figure 3.1: Sample BPM measurement.

and 0.05 mm tall. It's like a small, flat ribbon of charge — but a ribbon that happens to be zipping around an underground storage ring at nearly the speed of light.[1]

Each time the bunch passes a particular location in the ring, we use a Beam Position Monitor (BPM) to measure its horizontal and vertical position relative to the center of the beam pipe. In this plot, I am showing the horizontal positions (x) vs turn number as measured by one of the BPMs. At approximately turn number 1500, we gave the beam a little kick with a pulsed magnet, thereby knocking it out of equilibrium. This caused the bunch to start oscillating, swinging from side to side in the beam pipe. As you can see in the plot, those oscillations start out with an amplitude of about 4 mm, but they gradually die away over the course of the next several hundred turns.

[1]Depending on how fast you walk, it might take you about 10 minutes to walk around the half-mile-circumference accelerator tunnel; the electrons in CESR make 400,000 circuits every second.

3.3 Radiation Damping

As it happens, there is a well-understood mechanism that causes the transverse oscillations in the electrons' motion to die away. In general, the reduction of an oscillation amplitude is called *damping*, and in this case the physics process that is responsible is *radiation*.

As I explained in Chapter 2, the electrons emit synchrotron radiation as they circle the ring. CESR contains four accelerating cavities, whose job it is to restore the energy lost to radiation. The process of losing energy to radiation and then gaining it back in the accelerating cavities is what causes the transverse part of the electrons' motion to damp out. There are two parts to the reason why:

1. When an electron radiates a photon, it loses energy but continues on the same trajectory; the electron's momentum vector gets little bit shorter but doesn't change direction.
2. When the accelerating cavity restores the energy, it always adds momentum *in the longitudinal direction*.

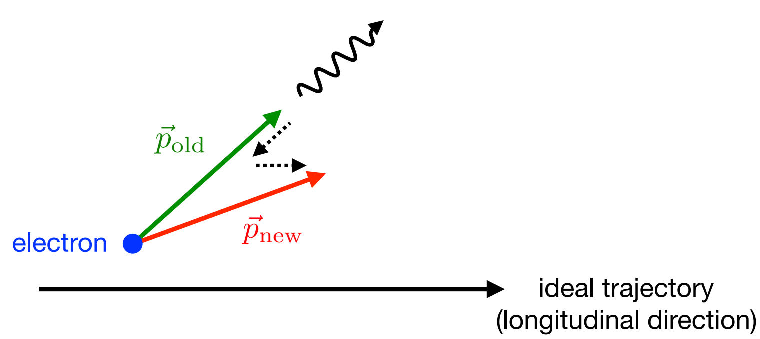

Figure 3.2: Radiation damping.

Thus, each time the electron goes around the ring, its momentum ($\vec{p}$) becomes a little bit less transverse and a little bit more longitudinal, gradually approaching the perfect design trajectory. This is illustrated in the cartoon in Figure 3.2. Watch out! I call this a "cartoon" in part because I did *not* draw it with realistic proportions (drawn realistically, you wouldn't be able to see anything in the diagram, so that would be pretty silly). I don't want you to get any wacky ideas, though, so I feel compelled to tell you a few things:

1. The transverse component of the momentum is always very small compared to the longitudinal component. In other words, the angle between the momentum $\vec{p}$ and the ideal trajectory is tiny right from the start. In my diagram, I happened to draw $\vec{p}_{\rm old}$ at an angle of about $40° = 0.7$ radians and $\vec{p}_{\rm new}$ at an angle of about $20° = 0.35$ radians. I did it like that so you would be able to see everything clearly — but those are ridiculous angles for the particles in an accelerator. A reasonable angle would be something on the order of a few milliradians or less.

2. The black dotted arrows in my diagram are meant to make you think of the changes in the electron's momentum during one turn around the ring: first the momentum lost to radiation and then the momentum gained from the accelerating cavity. In reality, those changes are *very* small compared to the overall momentum.

3. Finally, the energy loss and energy gain don't happen at the same time. The electron loses energy in little steps all the way around the ring (whenever it radiates a photon), while the energy gain happens in just one place, at the accelerating cavity. Furthermore, the momentum $\vec{p}$ is constantly changing as the electron oscillates back and forth on its trip around the ring. Sometimes

$\vec{p}$ points one way, sometimes the other way, sometimes straight ahead in the longitudinal direction; in order to calculate anything about the radiation damping, you have to remember to average across all those possibilities.

There's something else I need to tell you. Remember how I said (back in chapter 2) that we could forget about the opening angle of the radiation and pretend that the radiated photon is emitted in precisely the same direction as the electron's motion? Well, that's only true up to a point. The fact is that the electron *does* get a small sideways kick every time it emits a photon — sometimes to the left and sometimes to the right (there's a degree of randomness to this, just as with so much of subatomic physics). The effect of these sideways kicks is called *radiation excitation* and the result is that you can't ever have a perfectly damped beam. Instead, equilibrium occurs when the disruption to the orbit from the radiation excitation balances out any continued improvement from the radiation damping.

3.4 The Timescale of the Radiation Damping

So, can radiation damping explain what we saw in our measurement? The answer is *no*. Although radiation does cause the transverse oscillations to die away (just as we saw in our data), for 1.5 GeV electrons that process takes *much* longer than a few hundred turns in CESR.

Here's the thing: the radiated power — the energy lost per unit time because of synchrotron radiation — for an electron following a curved path inside an accelerator magnet is given by

$$P = \frac{1}{6\pi\epsilon_0}\frac{e^4}{m^4c^5}B^2E^2. \tag{3.1}$$

There are a bunch of constants in this equation (the permittivity of free space ϵ_0, the charge e and mass m of the electron, the speed of light c), but the main thing to notice is how the power depends on the squares of both the magnetic field B and the particle energy E.

Furthermore, B and E are related: if you want to force a relativistic particle with energy E to follow a circular path of radius ρ, you're going to need a magnetic field of strength

$$B = \frac{E}{ec\rho}. \tag{3.2}$$

Now, CESR is a circular machine with a radius of 122 meters. If we pretend for a moment that the electrons follow a perfectly circular path around CESR,[2] we can compute that the required magnetic field for our 1.5 GeV electrons would be

$$B = \frac{E}{ec\rho} = \frac{1.5 \text{ GeV}}{e(3.0 \times 10^8 \text{ m/s})(122 \text{ m})} = 0.04 \text{ Tesla}, \tag{3.3}$$

and each electron would radiate away its energy at a rate of

$$\begin{aligned} P &= \frac{1}{6\pi\epsilon_0}\frac{e^4}{m^4c^5}B^2E^2 = \left(\frac{ec^3}{6\pi\epsilon_0}\right)\frac{e^3B^2E^2}{(mc^2)^4} \\ &= \left(\frac{(1.6 \times 10^{-19} \text{ C})(3.0 \times 10^8 \text{ m/s})^3}{6\pi(8.85 \times 10^{-12} \text{ C/V}\cdot\text{m})}\right)\frac{e^3(0.04 \text{ T})^2(1.5 \text{ GeV})^2}{(0.511 \text{ MeV})^4} = 1.4 \text{ GeV/s}. \end{aligned} \tag{3.4}$$

That means the amount of energy lost in one turn around CESR (and subsequently restored by the accelerating cavities) would be

$$E_{\text{loss}} = (1.4 \text{ GeV/s})(2.5\ \mu\text{s}) = 3.5 \text{ keV}. \tag{3.5}$$

[2]In reality, the steering magnets cover only a fraction of the path, and are made stronger to compensate.

What does this have to do with the timescale of the radiation damping? Simply this: a reasonably good first estimate for the damping time (the time it takes to reduce the transverse amplitude by a factor of $e \approx 2.7$) is

$$\tau = \frac{E_{\text{particle}}}{E_{\text{loss}}} = \frac{1.5 \text{ GeV}}{3.5 \text{ keV}} = 430,000 \text{ turns.} \tag{3.6}$$

I've played fast and loose with a lot of details here, but the important thing is that the *order of magnitude* is correct — it does indeed take hundreds of thousands of turns for radiation damping to take effect for our 1.5 GeV electrons in CESR.

So, in our measurement, when we saw the transverse oscillations die out after only a few hundred turns . . . well, whatever caused that behavior, we know it *wasn't* radiation damping!

3.5 Simulating the Radiation Damping

Before I continue my story, I'd like to pause and say a few words about the fabulousness of computer simulations.

I began studying physics when I was junior in high school, almost sixteen years ago. In the intervening years, my understanding of what "physics" is (and how you do it) has changed a great deal. Here's one of the big-pictures lessons that I've learned: experimental particle physicists and accelerator physicists both invest a great deal of time and effort into computer simulations. Furthermore, in these fields, simulation tools are shared among scientists, often developed gradually over the course of years or even decades. The result is that modern simulation tools are very powerful and sophisticated, far better than anything a lone researcher could create if they started from scratch (though of course people still love to complain about the commonly-used tools!).

In the previous section, I made a rough estimate of the radiation damping time. What would it take for us to make this calculation more accurate? There are really two things:

1. We need to account for the subtleties of beam dynamics. It turns out that there's actually a factor of almost-two between the thing I calculated ($E_{\text{particle}}/E_{\text{loss}}$) and the real damping time. This is the kind of thing you can read about in accelerator physics textbooks (such as [3] and [4]).
2. We need to account for all the fiddly details of the real accelerator — in this case, the precise strengths and locations of all the magnetic fields in CESR.

There's surely an advantage to building a deep understanding of item (1), but item (2) is just tedious. Of course it's good to know the basic layout of your accelerator (and understand the importance of the various pieces), but there's no reason to sit around doing dozens of nearly-identical sub-calculations by hand. Computers are *really good* at repetitive, tedious calculations!

In fact, a simulation can help us with both items (1) and (2). Here at Cornell, we use a simulation tool called Bmad.[3] Using Bmad, I can easily load in a detailed description of CESR. Then I can give the simulation program a set of initial coordinates for a particle and ask it to compute the trajectory. The computer knows the electric and magnetic fields everywhere in the accelerator, so at each step it can find the electromagnetic force on the particle. The computer also knows where the beam pipe walls are, and will stop running if the particle crashes into something. If I specify that I want to include radiation in the simulation, the program will

[3]https://www.classe.cornell.edu/bmad/

use a random number generator to mock up the random nature of the synchrotron radiation. Where and when is a photon radiated? In what direction? All of that physics (in the form of equations describing probability distributions) has been programmed into the computer.

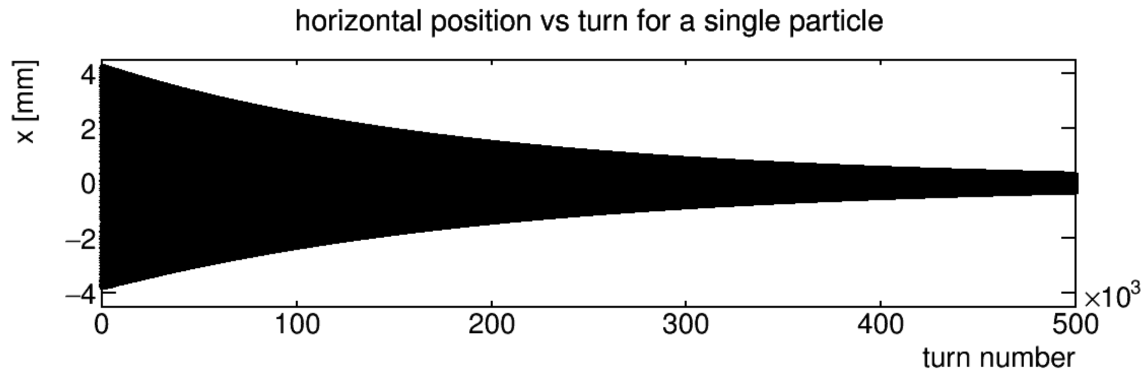

Figure 3.3: Simulated radiation damping.

Figure 3.3 shows the results of that simulation for a 1.5 GeV electron. The x axis covers a range of 500,000 turns — and you can see that this is indeed the appropriate timescale for the radiation damping!

A cautionary note: you should never blindly trust the results of a computer simulation. After all, the computer doesn't have magical knowledge about how the world works — it only ever does what you tell it to do (and sometimes what you told it to do isn't what you *thought* you told it to do). Often there are approximations built in to the program, and you (the user) need to make sure those approximations are valid for your situation. It's also entirely possible that the person who programmed in the physics made a mistake.

In this case, the simulation result reinforces my rough estimate. This agreement gives me more confidence that (1) I understand the physics and (2) I'm setting up the simulation correctly.

3.6 Decoherence

Now back to our mystery. Why *do* the oscillations in our data disappear so quickly? It turns out that it all comes down to these two truths:

1. Our BPM detectors measure the location of the bunch center, not the locations of any of the individual electrons.
2. The electron bunch is not a rigid object.

Earlier I described the electron beam as a tiny ribbon of charge, with a length, width and thickness. I do think that's a useful analogy, but (I confess!) it's also a bit misleading. Think of a normal, real-life ribbon: each tiny bit of the ribbon is bound to its neighbors in a fixed pattern. You can imagine a flexible ribbon, even a stretchy ribbon, but it's still a solid object. A particle bunch is *not like that*.

Before I say any more words, let me show you some illuminating simulation results.

One of the things our accelerator simulation tools can do for us is calculate the equilibrium properties of a particle bunch. (Remember the radiation excitation? That goes into this!). I asked the program to give me an equilibrium sample of 1000 particles for our 1.5 GeV accelerator setup. Each particle is created with an initial position, momentum, and energy, each randomly drawn from appropriate distributions. Figure 3.4 shows the distribution of horizontal positions for my 1000 virtual particles.

The center of this distribution is at $x = 0$ (the center of the beam pipe) and the standard deviation is $\sigma = 0.5$ mm (so the width of the bunch is ~ 1 mm, as you can see in Figure 3.4).

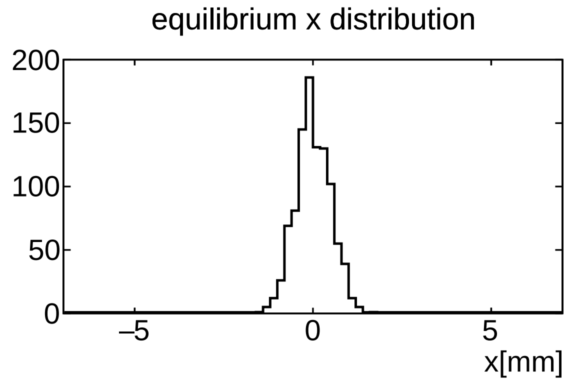

Figure 3.4: Simulated particle bunch distribution.

Next I gave my simulated particle bunch an initial horizontal offset of 4 mm (to mock up the situation in our measurement), and tracked the motion of each particle for 1000 turns around CESR. Figure 3.5 shows what happens to the average position of the electrons, which is what our detectors measure.

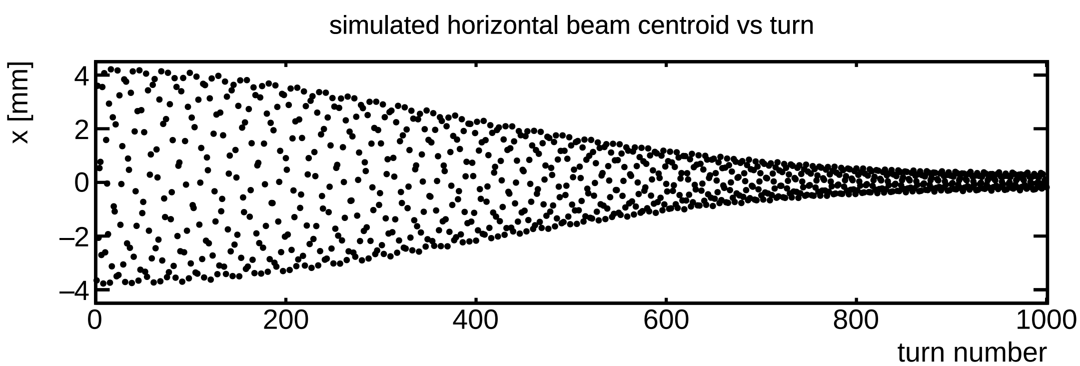

Figure 3.5: Simulated bunch center vs time.

Look at that! The initial amplitude dies away on a timescale of hundreds of turns, very much like what we saw in the data. The simulation doesn't match the measurement *exactly*, but the basic character is there. This is very encouraging news, because it suggests that the physics responsible for our observations is already part of the simulation, which in turn means that we can use the simulation to figure out what's going on.

The wonderful thing about a simulated world is that you (the programmer /physicist) have godlike powers. In real life we have no way to measure the trajectories of individual electrons, but in the simulation we know *everything*. (It's not always obvious what to look at, but in principle the information is there!)

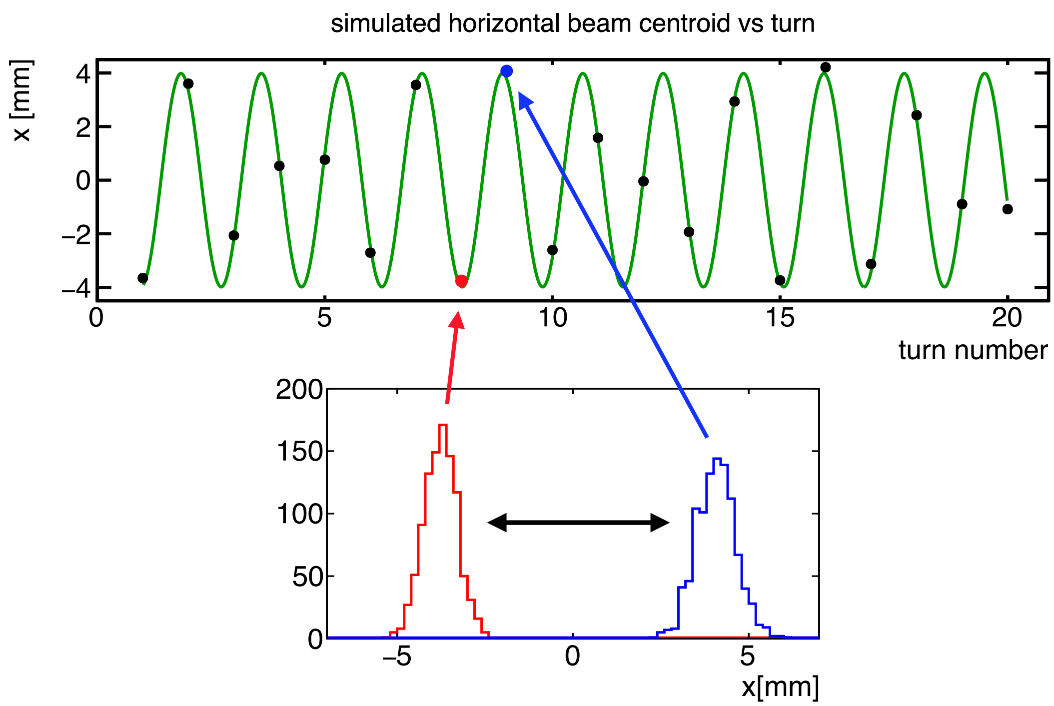

Figure 3.6: Simulated horizontal beam centroid vs turn.

Figure 3.6 is a zoom-in of the previous plot, showing the first 20 turns around CESR. I've also added a best-fit sine function to guide the eye. You can see that the bunch center is indeed oscillating with an amplitude of 4 mm, just as I intended when I gave the bunch an initial 4-mm offset. Below that, I've plotted histograms showing the horizontal positions of all the particles at two key moments in time (which is information we don't have in real life).

I've shown the x distribution at the two extremes of the oscillation: when the bunch is all the way to the left and when the bunch all the way to the right. In each case, the distribution looks pretty much the same as the equilibrium distribution, and you can imagine this distribution sliding back and forth along the x-axis as the bunch oscillates. It's all very nice and tidy.

Okay, but what about later? As shown in Figure 3.7, after 400 turns, the oscillation amplitude of the bunch center has shrunk to only 2 mm. But look what happened to the distributions!

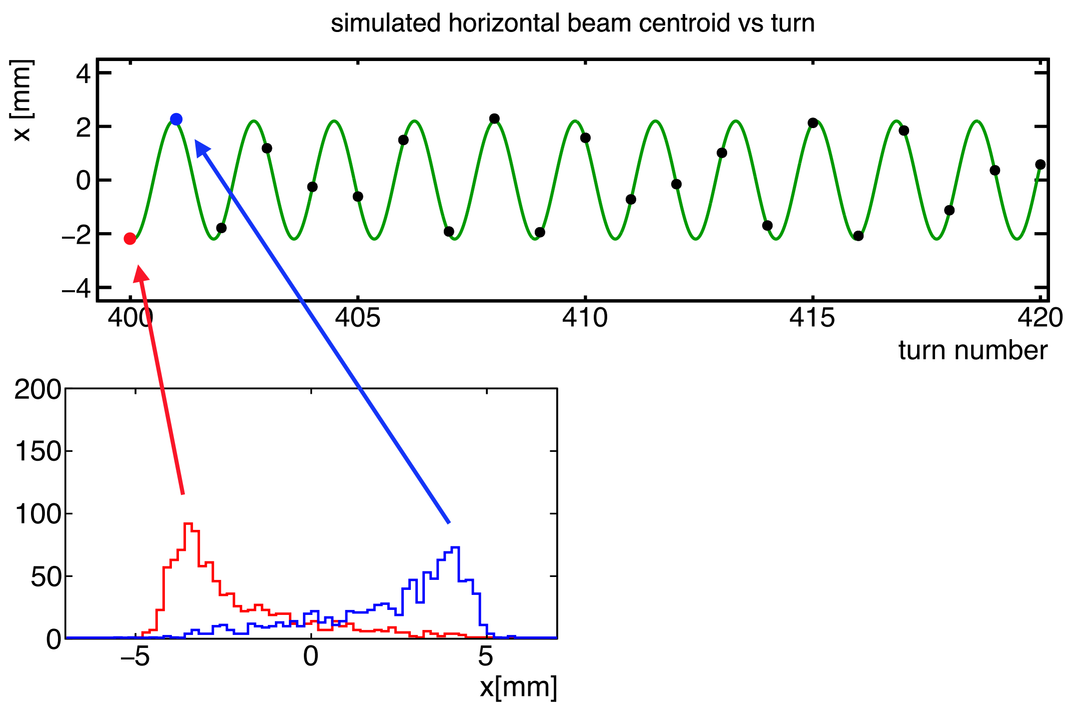

Figure 3.7: Simulated horizontal beam centroid vs turn after 400 turns.

If you only gave me a moment to glance at these distributions, and asked me what the average positions are, I would probably say −4 mm (red) and +4 mm (blue), because that's where the peaks are. However, if you look more closely, you see that the distributions are no longer symmetric — each has developed a long, one-sided tail. When the bunch is all the way to the left, most of the particles are clustered around $x = -4$ mm (just like before), but there are a significant number of stragglers, stretching all the way across to the other side. As a result, the mathematical averages are only −2 mm and +2 mm.

What is this madness? It has a name, and that name is *decoherence*. If a collection of things oscillates together in lock-step, all perfectly in phase with each other, we say that their motion is *coherent*. Before I did these simulations, when I thought about the electron bunch, I naively pictured coherent motion — probably because I was inadvertently thinking of the bunch as a rigid object. Now, though, the real picture is emerging: After 400 turns, each particle is still oscillating back and forth across the full range of motion (−4 mm to +4 mm), but they've started to get out of sync with each other.

And in Figure 3.8, you can see that after 800 turns, the decoherence is nearly complete — there's hardly any difference between the extremal distributions.

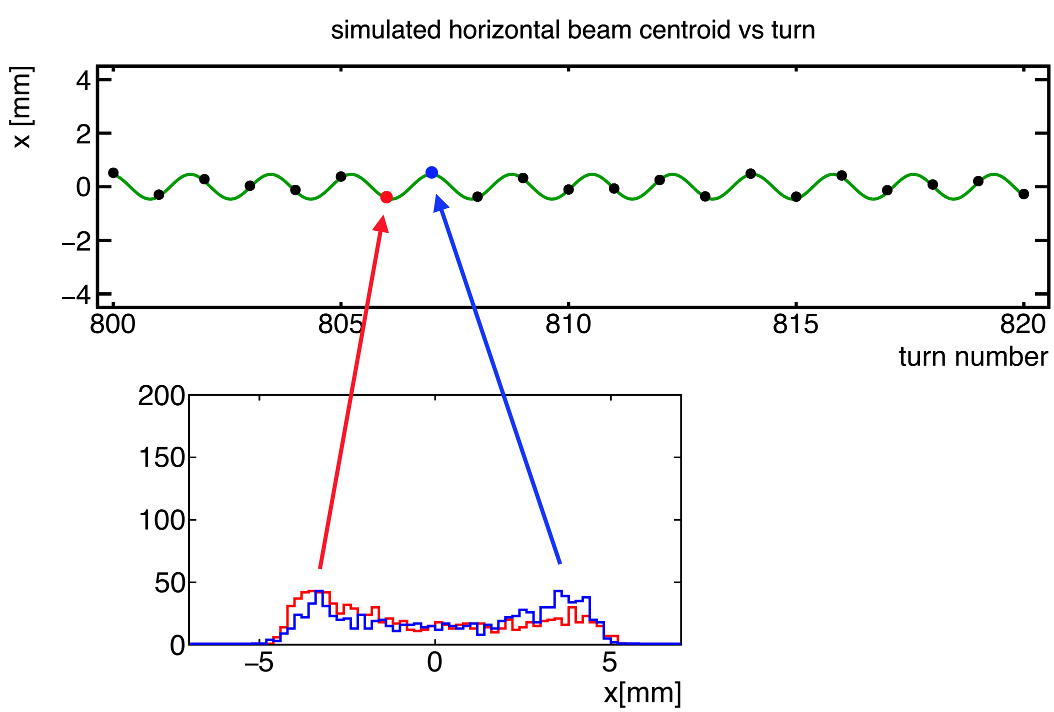

Figure 3.8: Simulated horizontal beam centroid vs turn after 800 turns.

Why did this happen? It must be the case that the various particles in the bunch each have a slightly different oscillation frequency.

3.7 Damping and Decoherence in Phase Space

I want to show you one more way of looking at the decoherence. Accelerator physicists are constantly plotting beam distributions in something called *phase space*. In a moment, you'll see how powerful and lovely this can be. First I need to define another coordinate (x'), which is the phase space partner to the horizontal position x.

An accelerator is designed with a particular *reference trajectory* in mind — this is the path that an ideal particle (perfectly on-axis, no transverse oscillations) would follow. The transverse coordinates of an actual particle are measured with respect to this reference trajectory, as shown in Figure 3.9.

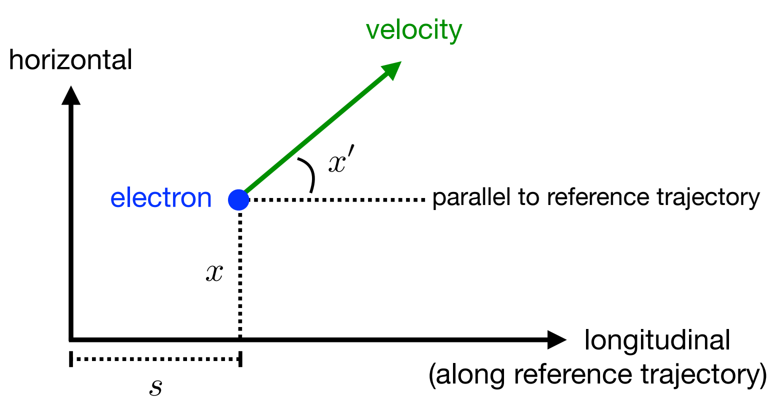

Figure 3.9: Definition of the accelerator coordinates s, x and x'.

The coordinate s tells you where you are in the accelerator — it measures the longitudinal position of a particle along the reference trajectory. The coordinate

x is the horizontal position and x' is the rate of change of x with respect to s,

$$x' \equiv \frac{dx}{ds}. \tag{3.7}$$

Because it's so small, x' is approximately equal to the horizontal angle of the particle's motion with respect to the reference trajectory (which is how I drew it on my diagram); I like this geometrical way of thinking about it because I find that it's easy to picture and hold in my mind.[4]

If you're only given the coordinates of a particle at one location, you need both x and x' in order to judge how close that particle is to the reference trajectory. Let's say a particle is following the path shown in the cartoon in Figure 3.10, oscillating around the reference trajectory.

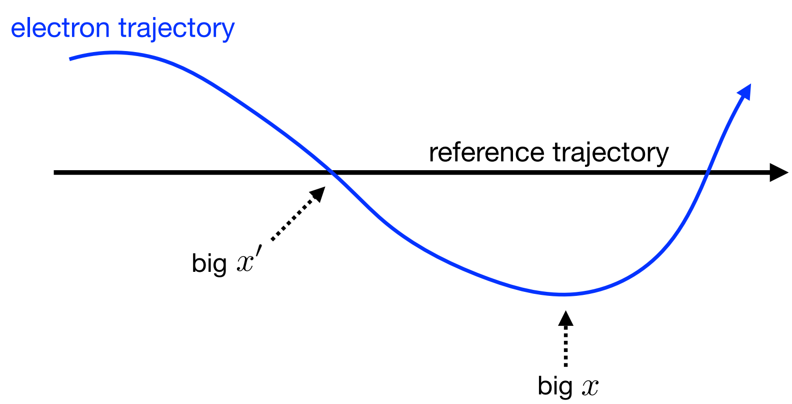

Figure 3.10: Cartoon trajectory showing the interplay of x and x'.

There are places along the path where x happens to be zero — but the large angle x' at those locations reveals the non-ideal nature of the particle's motion. It's this complementary nature of x and x' that makes them good phase space partners.

[4]If you're a physicist or are otherwise familiar with phase space, you're probably expecting the coordinate x' to have something to do with the horizontal momentum. It does. In another small-angle approximation, $x' \approx p_x/p_{\text{total}}$.

Okay! now we're ready to make some phase space plots (Figure 3.11). On the left is the bunch at equilibrium; this is once again the 1000 particles from my simulation, now with a dot drawn in phase space to mark the coordinates of each particle. When I give the bunch an initial x-offset of 4 mm, that's the same as sliding the whole distribution over by 4 mm in the diagram, and then letting it go.

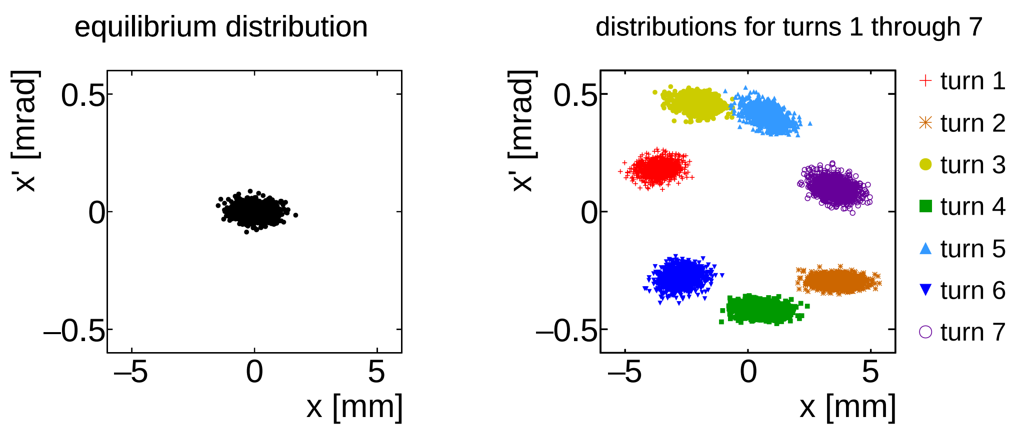

Figure 3.11: Phase space plots.

The figure on the right shows the resulting distribution on the next seven turns (always measured at the moment when each particle passes the starting location). You can see that the bunch remains pretty much the same size and shape, but appears at a different location in phase space each time it whizzes past.[5] After radiation damping does its work the bunch will be back at equilibrium, but that will take a long time (hundreds of thousands of turns).

Long before that happens, though, we get distributions like those in Figure 3.12. This is what decoherence looks like in phase space! These are the same distributions that I plotted in the previous section — two examples near turn 400 (on the left) and two examples near turn 800 (on the right).

[5]This is equivalent to being at a different phase of the sine wave I plotted earlier — hence the name *phase space*.

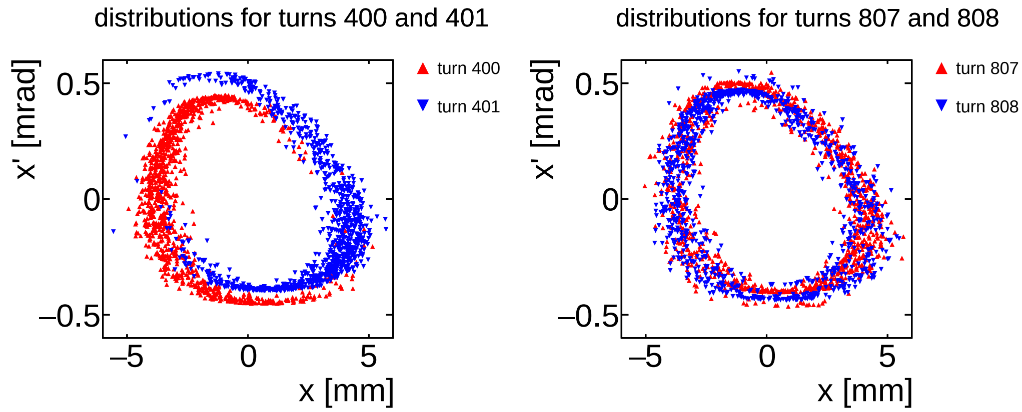

Figure 3.12: Phase space plots after more time has elapsed.

Look again at the plot on the left, the distributions for turns 400 and 401. Here in phase space, we can see something that we couldn't see before. Particles with larger amplitude (on the edge of the bunch that is farthest away from the phase space origin) are smearing counter-clockwise, while particles with smaller amplitude are smearing clockwise. This is a clue! It suggests that the source of the decoherence is an amplitude-dependent effect in the accelerator.

3.8 So What?

This story began with a strange and unexpected feature in our data. At the time, it struck me as a curiosity, but mostly just an inconvenience. After all, we were trying to get the particle bunch to oscillate so that we could feed the data through a Fast Fourier Transform (FFT) algorithm and extract the oscillation frequency. With the oscillations dying out so fast, we didn't have much data to work with. (*"Well that's annoying,"* I thought.)

But is it *important*? The answer might well be yes. The decoherence that we observed is a sign that there is an unusually large range of oscillation frequencies present in the particle bunch. This large tune spread might be interfering with our ability to accumulate particles in CESR. If we can mitigate that effect, it might improve the injection efficiency and/or the beam lifetime, thereby making it significantly easier for us to store, monitor, and tune up our 1 GeV electron beam. I will pick up this story again in Chapter 5, in which I investigate the possible causes of the tune spread.

First, I'd like to describe in some detail one of the successes of our Spring 2018 work with 1.5 GeV electrons — the orbit correction.

CHAPTER 4

CORRECTING THE ORBIT

4.1 Introduction

The theme for this chapter is imperfection. You see, the accelerator you build in real life is never quite the same — never quite as perfect — as the accelerator in your mind. So, if you're in the accelerator business, you'd better be prepared to deal with those imperfections. I want to tell you a story about one type of imperfection, the consequences of that imperfection, and the way we mitigate those consequences.

Let's get warmed up by looking once again at this measurement (Figure 4.1). When I showed you these data before, I drew your attention to the way the beam oscillates after we give it a magnetic kick. In particular, we focused on the time it takes for those oscillations to die away.

That's all very interesting, but I skipped over an important feature of this plot. I swept it under the rug, hoping you wouldn't notice and get suspicious!

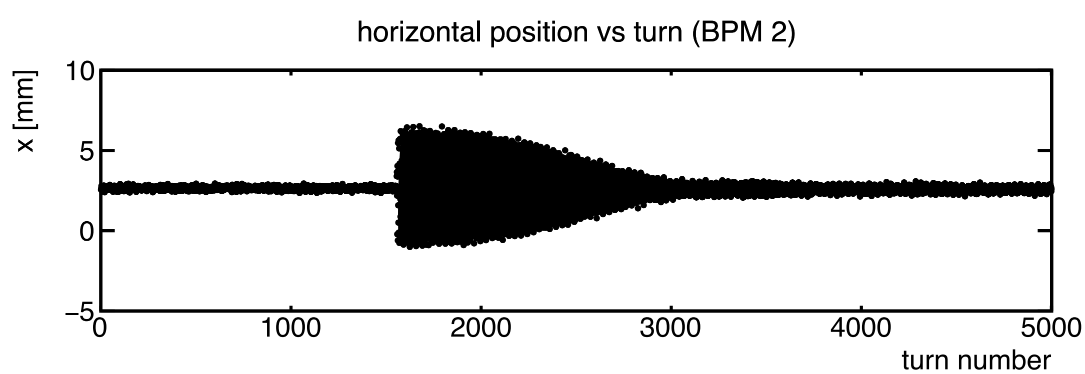

Figure 4.1: Our favorite BPM data, one more time.

Take a look at the beginning of the plot, *before* the kick happens. During this time, the beam is at equilibrium (no oscillations), which means it follows the same path every time it goes around the ring. At the location of BPM number 2, we see that the bunch is always approximately 2.5 mm to the right of the beam pipe's center.

I hope the following question just popped into your mind:
2.5 mm? How come the equilibrium position isn't zero?

Hang onto that question.

4.2 Measuring The Orbit

Here's the plan for this section:

1. First, some terminology: I will tell you what I mean by the "orbit."
2. We will delve into the details of an orbit measurement — not because you need to understand those details in order to read what follows (you can skip this part if you like), but because measurement lies at the very heart of experimental science, and (according to me) is always worth a close look.
3. I will show you the orbit we measured on 4/5/2018, the day we first stored 1.5 GeV electrons in CESR. (Spoiler: it doesn't match the design orbit!)

The *orbit* — also known as the *equilibrium orbit* or the *closed orbit* — is the path taken by an ideal reference particle as it travels around the storage ring. It's called the *closed orbit* because it closes on itself after just one turn around the ring. Furthermore, the closed orbit is unique: there is only one such trajectory.

Real particles, of course, are never quite perfect. Even when the bunch as a whole is at equilibrium — so that the center of the bunch travels along the closed orbit — each individual particle still oscillates around that orbit, albeit with an oscillation amplitude that may be very small.

If you hear an accelerator physicist talking about "designing a lattice," they mean choosing the location and specifications of the sequence of elements (mostly magnets) that make up the machine. Once you have a lattice, you can compute the associated closed orbit.

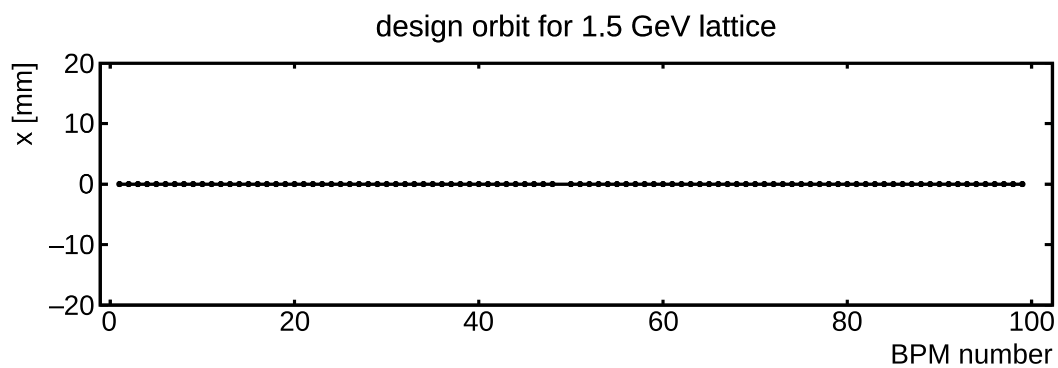

Figure 4.2: Design orbit for 1.5 GeV lattice.

For our 1.5 GeV CESR lattice, the ideal closed orbit is shown in Figure 4.2. That is, the design orbit goes right down the center of the beam pipe, all the way around the ring. This is almost a tautology: *of course* the accelerator is designed so that the perfect reference particle goes right down the middle! I mean, really: the entire coordinate system had to be defined relative to *some* reference trajectory, so why wouldn't that reference be the same as the closed orbit?

It turns out, though, that the design orbit isn't always so boring. For example, Figure 4.3 shows the orbit that was used for CHESS runs up until summer 2018. You'll notice that there are actually *two* orbits: one for positrons, which travel clockwise in CESR (left to right in the plot) and one for electrons, which travel

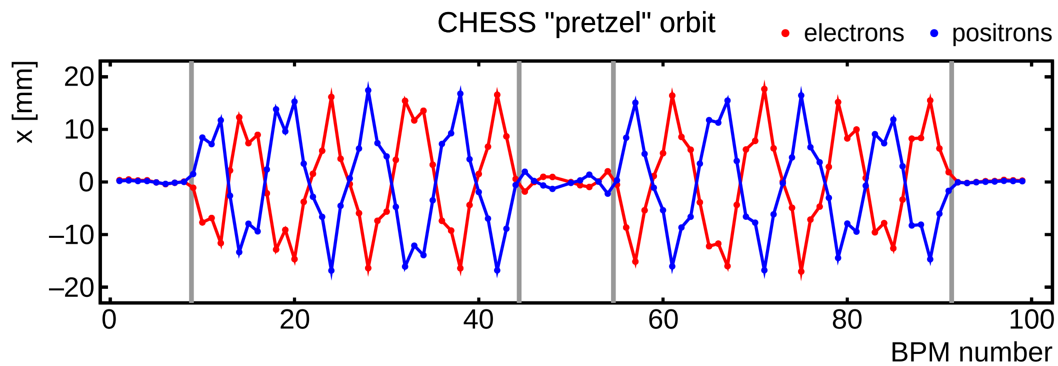

Figure 4.3: The CHESS "pretzel orbit."

counterclockwise (right to left). In the old CHESS setup, some of the x-ray beam lines were fed by electrons and others were fed by positrons. Thus, in order for CHESS to run at full efficiency, we needed to have electrons and positrons in CESR at the same time.

This so-called "pretzel orbit" allows the two counter-rotating beams to exist in the machine simultaneously without ever colliding. The electron and positron orbits *do* cross at several locations, including an extended overlap on the south side of the ring (from BPM 92 wrapping back around to BPM 8), but the particle bunches can be spaced in such a way that they always miss each other.

From 1979 until 2008, CESR was home to the CLEO detector, which sat at location 0 in the ring. A similar pretzel orbit was used in those days, with one important difference: the electron and positron bunches were timed so that they *would* collide, right in the middle of CLEO. Experimenters studied those collisions in order to learn more about the underlying particle physics.

The pretzel orbit is created by a set of four *electrostatic separators*. I marked the locations of these devices in grey on the orbit plot. The electric field inside each separator deflects the electrons and positrons in opposite directions, thereby establishing different orbits for the two particle species.

As of this year, however, the days of the pretzel orbit are over. In summer 2018, CHESS and CESR underwent a big upgrade. All the CHESS x-ray lines now emerge in the *same* direction from CESR, which means that only *one* particle species is needed for CHESS running,[1] and we can use a simple straight-down-the-middle orbit.

The OSC experiment — the impetus for us to push down the energy of CESR — will also use only one particle species (electrons), and can have a simple orbit.

Now that we know what the orbit *should* be (zero everywhere!), let's take a closer look at the first 1500 turns from our sample data set, before the kick, when the beam is still at equilibrium (Figure 4.4).

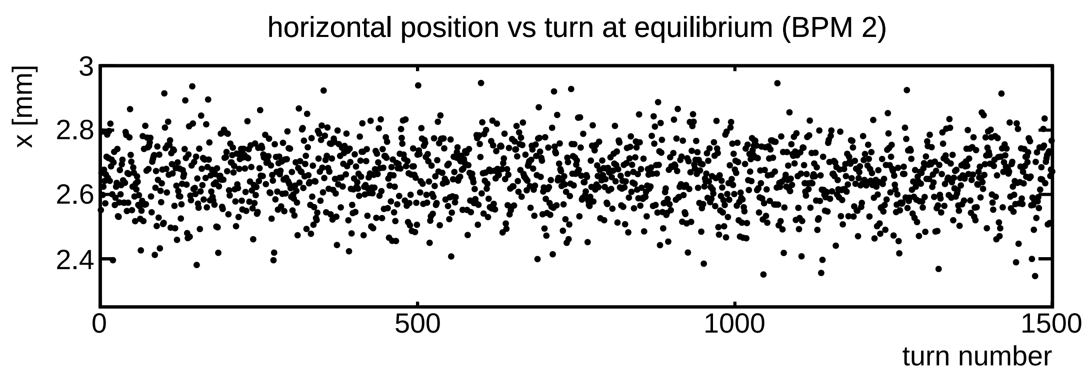

Figure 4.4: Horizontal position measurements of the beam at equilibrium.

With this zoomed-in view, we can see that the measurements are not, strictly speaking, identical: they range from about 2.4 mm to 2.9 mm. There's no discernible pattern, and each chunk of time looks the same as all the others — which means that it doesn't actually make much sense to display these data as a time series. We can learn more by reformatting the same information into a histogram,

[1] For now, CHESS is using positrons, but I've heard talk that they may someday swap the polarity of the synchrotron and storage ring, allowing the use of clockwise electrons for CHESS.

showing the number of times each possible x value appears in the data (Figure 4.5).

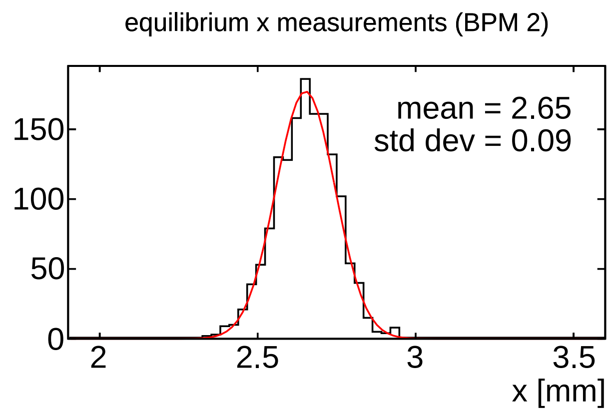

Figure 4.5: Equilibrium x measurements.

Aha! Now we see that our equilibrium position measurements have a classic bell curve shape, also known as a *normal distribution* or *gaussian distribution*. I've drawn the histogram in black and the best-fit gaussian function in red. Having shown that our equilibrium position measurements are indeed gaussian, we can reduce all 1500 measurements to just two numbers: the mean and the standard deviation of the distribution.

The mean ($\langle x \rangle = 2.65$ mm) is the center of the distribution — the number I initially eyeballed as 2.5 mm on the zoomed-out plot. Now that I've actually calculated the mean from the data, I can amend my previous statement: the equilibrium horizontal position of the beam at this location is 2.65 mm. (You should still be asking, *"How come it's not zero?"* but now our mystery is more precise!)

The standard deviation ($\sigma = 0.09$ mm) describes the width of the distribution. For a gaussian distribution, 68% of values fall within one standard deviation of the mean (in this case, between 2.56 and 2.74 mm), and 95% fall within two standard deviations (2.47 to 2.83 mm).

I suspect this is all rather hum-drum for any scientists in the audience, because it comes down to one of the universal and inescapable truths of experimental science.

universal and inescapable truth: Every measurement has uncertainty.

corollary: The uncertainty is usually gaussian.

That means gaussian distributions pop up all the time in science, especially when we measure the same thing over and over again.

But what's really going on here? There are two main possibilities:

1. Some random effect is causing the bunch to genuinely be in a different place each time it passes this BPM (sometimes at 2.5 mm, sometimes at 2.8 mm, etc.).
2. The beam is truly at 2.65 mm on every turn, but the individual measurements are smeared out due to electronic noise in the BPM button signals.

It could also be some combination of (1) and (2). In general, the spread in a set of repeated measurements can be attributed to many simultaneous factors. If you want to truly understand your measurements (and a physicist usually does!), it behooves you to think about all the possibilities. In this case, I can tell you that (2) substantially outweighs (1): the variation that we see in our x values is largely caused by noisy signals.[2]

[2]For this example measurement, the beam current in CESR was only 82 microamps. With a higher current, the signal-to-noise ratio is higher and the corresponding position uncertainty is smaller.

Identifying all possible sources of uncertainty (and studying them in detail) is especially important if you're trying to do a precision measurement. Lucky for us, we're *not* trying to do a precision measurement right now: we know that the design orbit has $x = 0$ at the location of BPM 2 (indeed, at *all* the BPM locations), whereas our measurement shows $x = 2.65$ mm. Whether it was actually 2.6 mm or 2.7 mm instead of 2.65 mm doesn't much matter, because those are small differences compared to the deviation from the design value.

Nevertheless, I want to share one more factoid about gaussian distributions. When you calculate the mean by averaging N values, the uncertainty on the mean is given by

$$u = \frac{\sigma}{\sqrt{N}}.$$

The more measurements you include in the average, the better you can pinpoint the center of the distribution. For our example measurement, I included 1500 measurements in the average, so the uncertainty on the mean is

$$u = \frac{0.09 \text{ mm}}{\sqrt{1500}} = 0.002 \text{ mm}.$$

This means I'm justified in my choice to report the mean $\langle x \rangle = 2.65$ mm out to two decimal places, and, if that's all the precision we need, there's no need to include any additional measurements in the average.

To measure the full orbit, all we need to do is repeat the same averaging procedure for all 100 BPMs. Figure 4.6 shows the orbit we measured when we first stored 1.5 GeV electrons in CESR.[3] You can see that the orbit is pretty wild, with horizontal deviations from the design orbit as large as 1 cm. The beam pipe walls

[3] If you have sharp eyes, you might notice that BPM 2 (the first data point) has $x \approx 4$ mm in this plot, not the 2.65 mm from my sample measurement. Indeed, that is true! We measured this initial orbit on 4/5/2018, whereas the data I've been using as an example were collected on 4/24/2018, after we had already taken steps to improve the orbit. What are those steps? Keep reading, and I'll tell you!

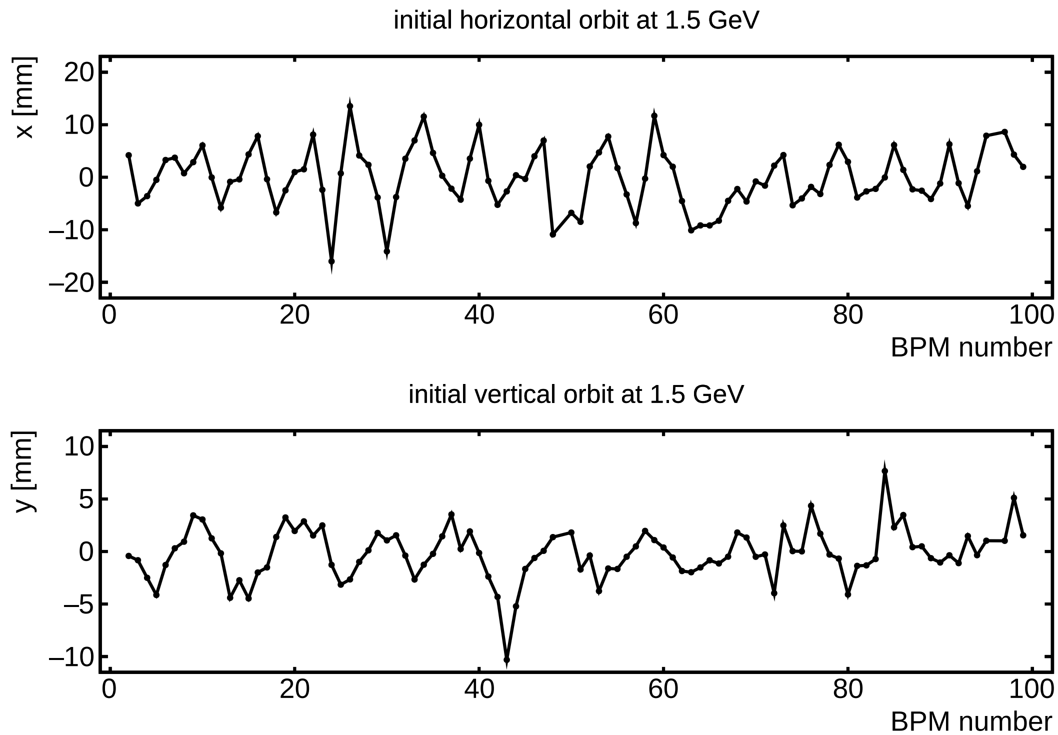

Figure 4.6: Initial orbits at 1.5 GeV.

are at $x = \pm 4.5$ cm, so that's a pretty big deviation (though still more contained than the pretzel orbit!).

We're left with the question *why.* Why does the measured orbit differ from the design orbit? Why does it look so wacky?

The problem is that the storage ring magnets aren't perfectly aligned.

4.3 Effects of Quadrupole Misalignment

Earlier I made this claim: the measured orbit differs from the design orbit because the storage ring magnets aren't perfectly aligned.

Now, if a dipole magnet isn't in quite the right place, that doesn't really matter, because the dipole field is uniform. A quadrupole field, on the other hand, is position-dependent — if the quadrupole isn't precisely centered on the design trajectory, there are consequences. For example, let's suppose a quadrupole is shifted a little bit to one side. That means the entire magnetic field will be shifted, as shown in Figure 4.7.

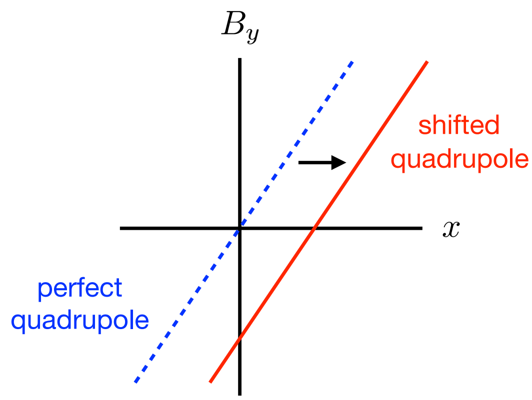

Figure 4.7: The magnetic field of a misaligned quadrupole

But we can also think about this a different way. As shown in Figure 4.8, the magnetic field of the misplaced quadrupole is the same as what you would get from a combination of two perfectly-placed magnets: a quadrupole and a dipole! For this reason, the field error caused by a misplaced quadrupole is called a *dipole error*. It's just like having an accidental extra dipole in the ring.

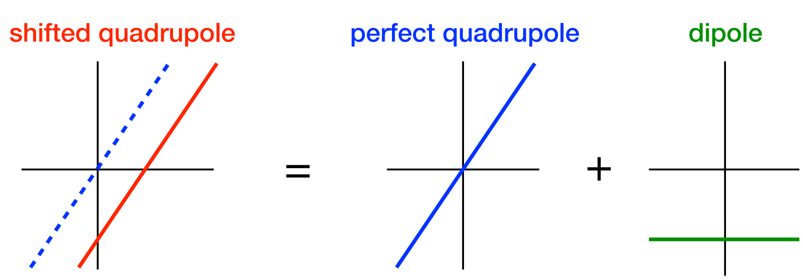

Figure 4.8: The field of a misaligned quadrupole is the sum of a perfect quadrupole field and a dipole field.

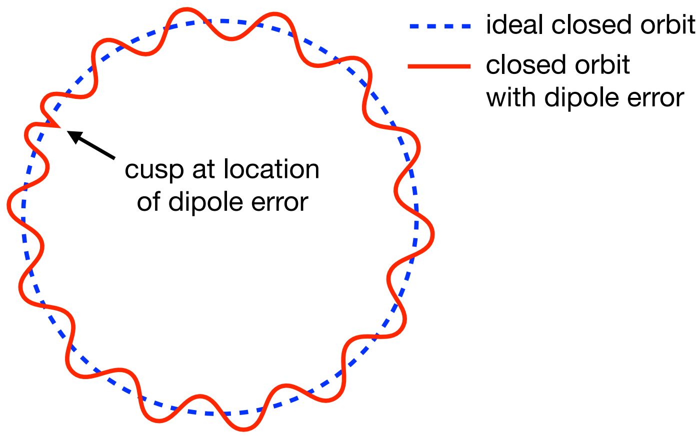

Figure 4.9: Cartoon version of the effect of a dipole error on the closed orbit.

Adding an extra dipole to the ring *definitely* changes the closed orbit. How could it not? Every time a particle goes past that location, it gets an extra kick — a kick which must now be incorporated into the orbit. As a result, the new orbit will have a cusp at the location of the dipole error. Furthermore, the closed orbit is truly a global quantity — you can't change one part of the orbit without changing the whole thing. In this case, pinching the orbit in one spot sends ripples through the entire closed solution. The result will look something like Figure 4.9.

Figure 4.10 shows a more realistic picture of the same scenario. I used our simulation tools to add a 1-mm horizontal offset to one of the CESR quadrupoles, then re-calculated the closed orbit.[4] I've marked the location of the misplaced quadrupole in red. This *looks* messier than my cartoon version, but look — it has the same features: (1) a cusp at the location of the dipole error and (2) oscillations around the perfect design orbit.

[4]Maybe you want to know how the simulation finds the closed orbit. Here's how it works: the computer chooses a set of initial coordinates for a particle, then tracks that particle once around the ring. If the final coordinates match the initial coordinates, then that's it — we've found the closed orbit! Otherwise, the computer changes the initial coordinates a little bit and tries again, until it eventually finds the closed solution. This kind of multi-dimensional optimization problem comes up a lot in accelerator physics (sometimes with a *lot* of variables), so it's good that we have computers at our disposal.

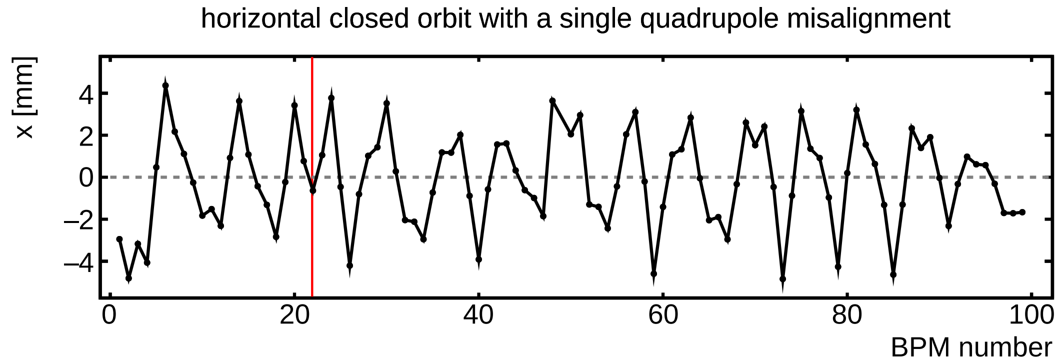

Figure 4.10: Simulated closed orbit with a single quadrupole misalignment.

With just *one* misplaced quadrupole in the ring, it's pretty easy to spot the cusp and identify the culprit. In real life, however, there will be dipole errors (some smaller and some larger) associated with *all* the quadrupoles. If you look back at the measurements I showed you earlier (Figure 4.6), you can clearly see the oscillatory nature of the orbit, but it's hard to tell which quadrupole(s) are responsible.

4.4 Correcting the Orbit

We want the orbit to be as close as possible to the design orbit. I hope I've convinced you that the discrepancy is caused by quadrupole misalignments, but that still leaves the question: *How do we fix it?*

Conceptually, this problem has two parts:

1. Locating the dipole errors.
2. Un-doing the dipole errors.

Let's start with (2). You might be thinking, "Easy — go down in the tunnel and do a better job lining up the magnets!" But that's not the answer. Aligning massive chunks of metal is just *hard* — and no matter how careful you are, there will always be some amount of error. Besides, we really need a more flexible solution, something that can be implemented and changed quickly.

The answer, it turns out, is to add *even more* dipoles to the machine. We don't try to fix the underlying quadrupole misalignment (unless it's *really* egregious); instead, we treat the symptoms: when a dipole error kicks the beam one way, we use one of our extra dipoles to kick it back the other way.

These extra dipoles are called *steering magnets*. They give us the handles we need to control the orbit. CESR has about 50 horizontal and 50 vertical steering magnets. These aren't actually stand-alone magnets. Instead, they consist of extra coils wrapped around the poles of existing magnets. Most of the horizontal steerings are built on top of existing dipole magnets and most of the vertical steerings are built on top of sextupole magnets. All the main dipoles in the ring are wired in series and controlled together, but each steering coil has its own power supply; this allows us to independently strengthen or weaken the magnetic field in each dipole (or sometimes each pair of dipoles) in order to compensate for nearby dipole errors.

Let me show you how an orbit correction plays out for a simple example. I already showed you my simulated orbit for a lattice that is perfect everywhere except for a single misaligned quadrupole. Figure 4.11 shows that orbit again (in black), together with the corrected orbit (in blue) that can be achieved by carefully adjusting the strength of just one steering magnet. As you can see, the improvement is dramatic!

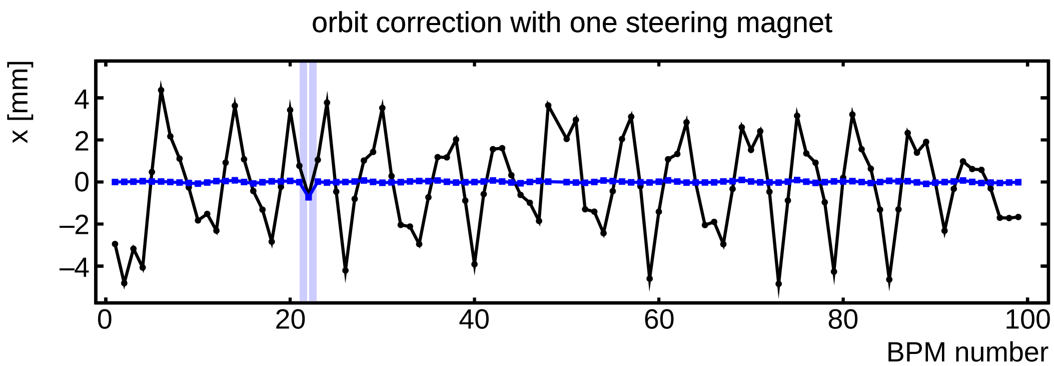

Figure 4.11: Simulated orbit correction using a single steering magnet.

In this case, the closest horizontal steering is built on top of a pair of dipoles, one on either side of the misaligned quadrupole. I marked the location of these two dipoles with shaded bars on the plot. The cusp at the location of the misaligned quadrupole is still there (because the dipole error is still there!), but the steering correction does a lot to flatten the orbit around the rest of the ring.

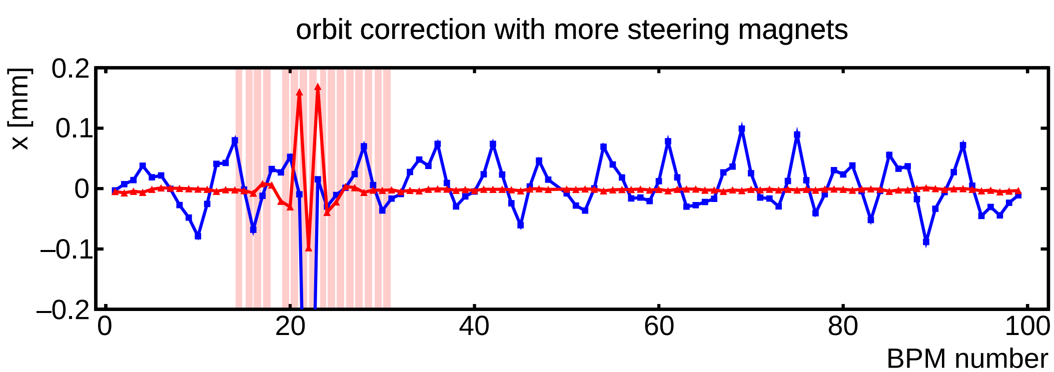

Figure 4.12: Simulated orbit correction using nine steering magnets.

Things get even better if we use more than one steering magnet. Figure 4.12 shows the best possible correction that we can achieve using the 9 red-marked steering magnets (2 single dipoles and 7 pairs of dipoles). Note the change in scale — the blue orbit is once again the 1-steering correction, and the red orbit is the 9-steering correction. As before, there is an (unavoidable) disturbance close to the location of the dipole error, but the orbit everywhere else is extremely flat.

I'll tell you a secret: I didn't figure out the best settings for those 9 magnets by myself. I didn't even figure out the 1-steering correction by myself. That would have been extremely tedious. No, I used a computerized optimization routine to tell me the best settings for the steering magnets.

CESR orbit corrections are typically done via computer optimization, but the procedure is finicky enough that human eyes and brains are still required (and it can take many hours to get right, especially for a new lattice!). Now, earlier I said that an orbit correction requires two steps: locating the dipole errors, then un-doing those errors. When I said that, I was thinking like a human — a human who likes to know the *why* of things. In truth, the computer doesn't need to know which quadrupoles are misaligned in order to correct the orbit — it simply does an optimization to find the set of steering values that will produce the best possible orbit.

The computer program that contains our orbit correction routine is called CESRV ("Virtual CESR"). CESRV includes an entire suite of routines to help with CESR operations. Here's an outline of how an orbit correction works in practice:

1. We measure the orbit (CESRV has a command for this, though we discovered it doesn't always work well for low beam currents).
2. CESRV does an optimization to find the steering magnet strengths needed to make the *simulated* orbit (in a perfect machine with no misalignments!) match our measured orbit. We then conclude that if we make the *opposite* changes to the real steering magnets, the orbit will improve.
3. CESRV is connected to the real steering magnet controls, so we can issue a command to load the new settings into the machine.

This procedure is complicated by the fact that making *any* change to the machine can have unintended consequences. Sometimes step (3) causes us to lose the beam, thereby requiring additional time (and sometimes additional tuning of the injection process) to re-accumulate particles in CESR.

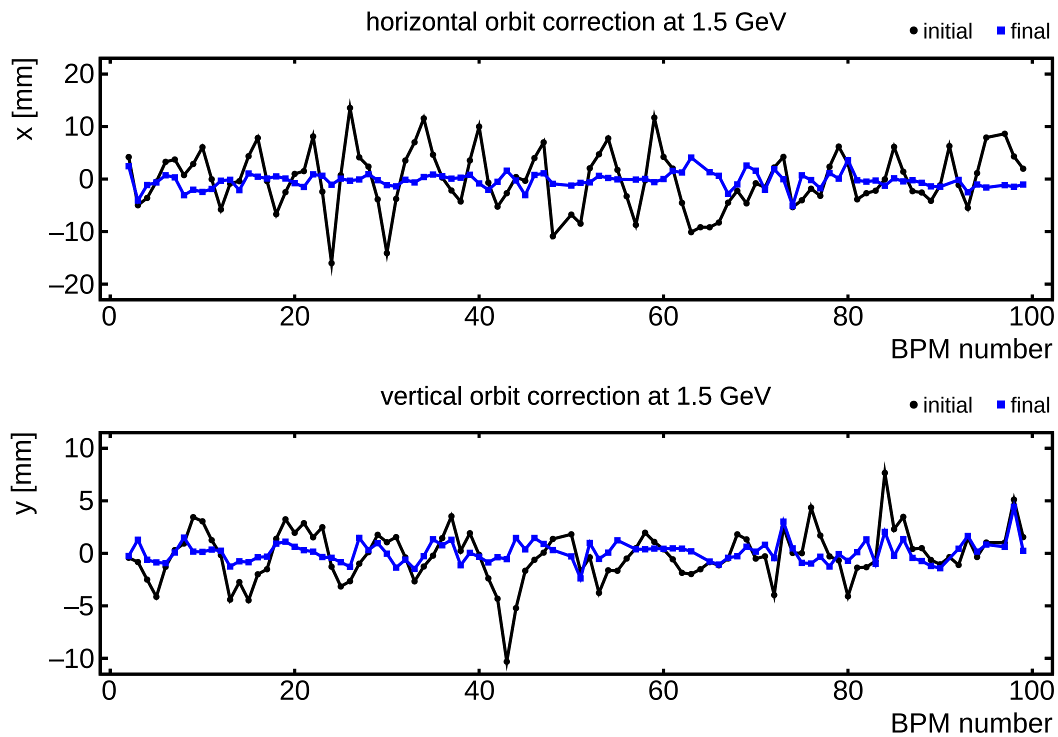

Figure 4.13: Successful orbit correction for 1.5 GeV electrons in CESR.

In spring 2018, we spent about 35 hours (in 6 separate shifts spread over a few weeks) working to store 1.5 GeV electrons in CESR. It took a while to store any particles in CESR at all — but once we managed to accumulate a measurable beam, one item on our agenda was to do an orbit correction. Figure 4.13 shows the results of that effort. In black is the initial orbit, collected on 4/5/2018. The final orbit is in blue, measured on 4/24/2018 after a couple successful iterations of our orbit correction procedure (as well as the usual assortment of problems, mysteries, and unsuccessful attempts).

CHAPTER 5

NONLINEARITY AND THE AMPLITUDE-DEPENDENT TUNE

As we saw in Chapter 3, the decoherence we observed in our stored 1.5 GeV electron beam revealed what must be a large tune spread within the bunch. Furthermore, simulations revealed that it appeared to be an amplitude-dependent effect, with large-amplitude particles oscillating at a different frequency than small-amplitude particles. An amplitude-dependent tune is a clear sign of nonlinearity in the machine.

Dipole and quadrupole magnets both have linear fields; if (ideal) dipoles and quadrupoles were the only elements present in CESR, we would expect the resulting beam dynamics to be linear as well — the motion of a large-amplitude particle could simply be scaled down, and it would match exactly the motion of a small-amplitude particle.

However, there *are* elements in CESR which have nonlinear magnetic fields. Some questions arise:

1. Which element(s) are responsible for the large tune spread we observed at 1.5 GeV?
2. Is this nonlinearity interfering with our ability to inject, accumulate and store electrons in CESR?
3. Is there anything we can do about it?

This is another case where our simulation tools prove useful: we can easily simulate the motion of particles with various amplitudes, and we also have the ability to test scenarios where the various nonlinear elements of CESR are adjusted or even removed entirely.

Figures 5.1 and 5.2 show the amplitude dependence of the horizontal and vertical tunes in our 1.5 GeV lattice. In each case, I tracked particles up to the largest amplitude that could be stored without hitting an aperture.

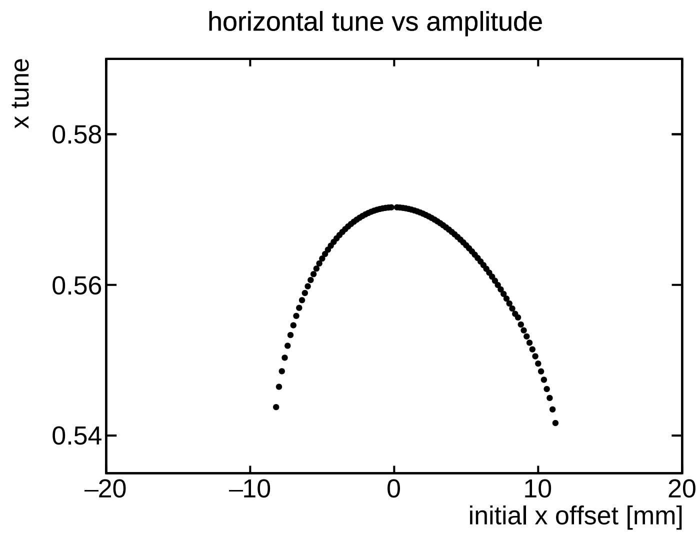

Figure 5.1: Simulated horizontal tune vs particle amplitude.

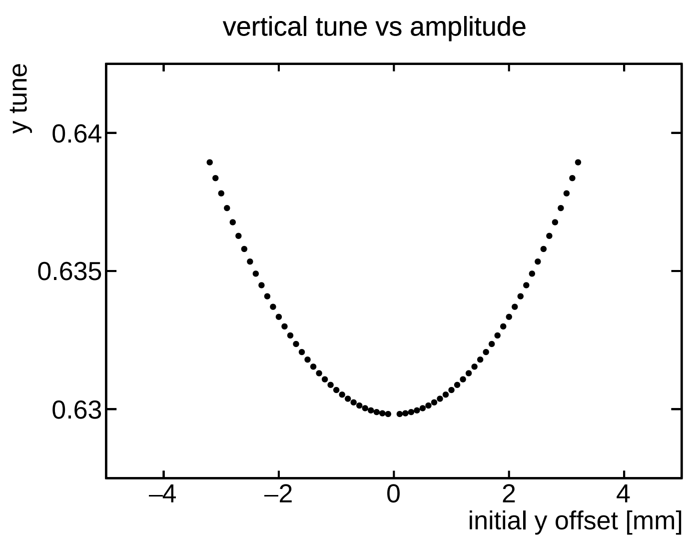

Figure 5.2: Simulated vertical tune vs particle amplitude.

Two of the sources of nonlinearity in CESR are the sextupole magnets and the undulator magnets. My simulations show that the sextupoles are important for the horizontal amplitude-dependent tune (Figure 5.3), but have very little influence on the vertical. Conversely, the nonlinear component of the undulator field is important in the vertical (Figure 5.4), but not at all in the horizontal.

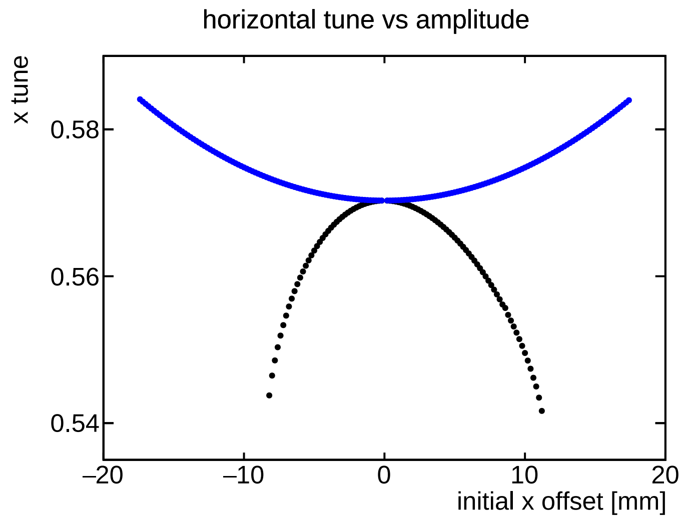

Figure 5.3: Simulated horizontal tune vs particle amplitude for two cases: sextupoles on (black) and sextupoles off (blue).

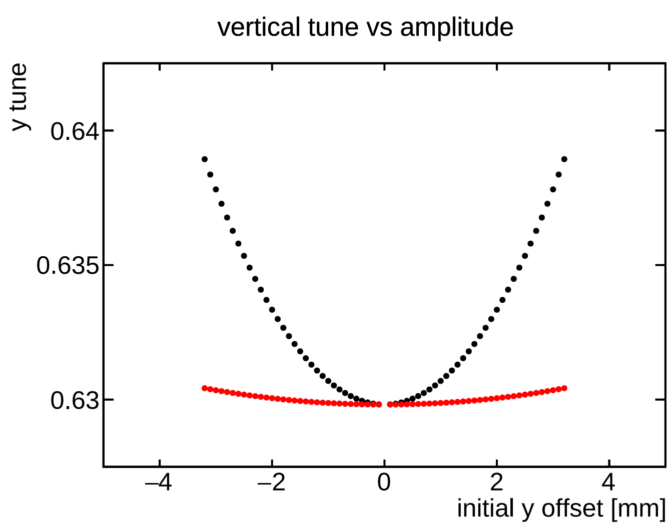

Figure 5.4: Simulated vertical tune vs particle amplitude for two cases: full undulator fields (black) and undulator fields with nonlinear component removed (red).

Now, while it's possible to change the sextupole strengths (or even turn them off entirely), we're pretty much stuck with the undulators as they are. Luckily, the situation in the vertical is much less problematic than situation in the horizontal.

The overall tune can be adjusted by strengthening or weakening all the quadrupole magnets together — and it turns out that the particle dynamics sometimes sharply depend on this base tune.

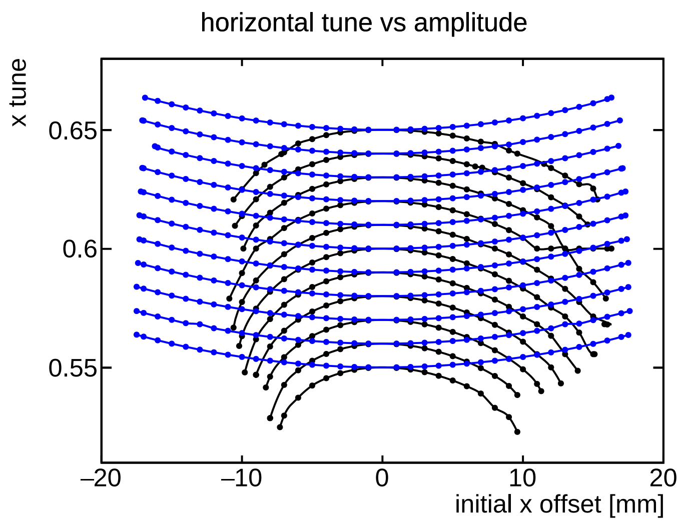

Figure 5.5: Horizontal tune scan.

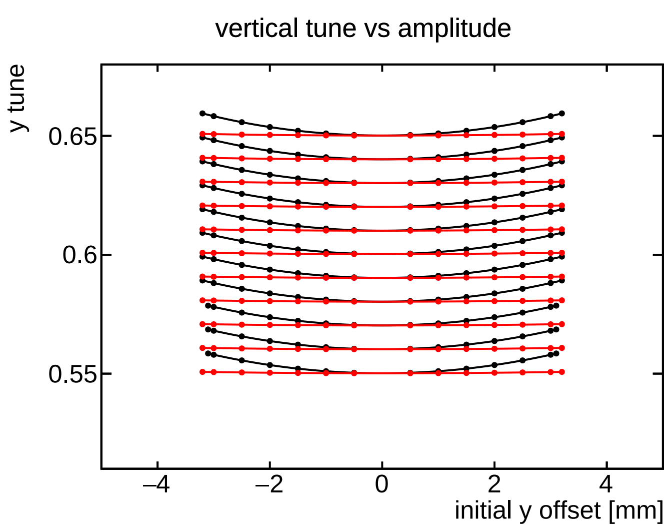

Figure 5.6: Vertical tune scan.

Figures 5.5 and 5.6 show the same kind of simulation results, but for a range of base tunes. You can see that for the vertical motion, the results don't much depend on the tune. Furthermore, the maximum-achievable amplitude changes very little when the nonlinear fields are removed. This is because the vertical motion is constrained by a very small *physical* aperture at the undulator; the vertical amplitude can't ever be large enough for the nonlinearity to have a big effect.

The horizontal case is different. With the sextupoles on, we see that the maximum-achievable amplitude does depend on the tune — and in all cases it is smaller than what is achievable when the sextupoles are turned off. The sextupoles *do* (of course) serve a function in CESR, so turning them off is not without consequence, but they become most essential at high beam current. For the OSC experiment, we require only a small beam current, so it may be advantageous to operate without sextupole fields.

CHAPTER 6

CONCLUSION

In summer and fall 2018, the Cornell accelerator complex was shut down for a major upgrade to the CHESS x-ray facilities. This involved replacing the entire south arc of CESR (approximately 1/6 of the ring). Consequently, our efforts to run CESR at low energy have been on hold since May 2018. Now, with the CHESS upgrade nearing completion, it will soon be time to resume our low-energy investigations.

We can't quite pick up where we left off, because CESR has *changed* in the meantime. If anything, the task will be more difficult now, because of the addition of several new small-aperture undulators to the CESR lattice. Nevertheless, we can carry forward some lessons that we learned in spring 2018.

There's a bit of a catch-22 involved in the early stages of establishing a new beam configuration: we need to accumulate a measurable number of particles in CESR in order to make the necessary corrections to the storage ring optics (including the orbit correction I described in Chapter 4, as well as corrections to the tunes, the phase advance between each pair of detectors, and the coupling between the horizontal and vertical particle motion) — but at the same time it can be *very* difficult to inject, store, and accumulate a beam before those corrections are made. You can't correct what you can't see, and you can't see what you can't inject!

This is a situation where *any* insight that might help us store and measure that first beam has the potential to be *extremely* useful. As I described in Chapter 3, last April we observed an unexpectedly quick decoherence in the horizontal motion of our stored 1.5 GeV electrons. Using our simulation tools, I identified the source of that decoherence: the nonlinear magnetic fields inside the sextupole magnets were causing a large amplitude-dependence in the horizontal tune.

Furthermore, my simulations showed that the *dynamic aperture* — the maximum possible oscillation amplitude for stored particles, taking account of all nonlinear effects — depends strongly on the horizontal tune. Now, during initial tuning (before measurements are possible) we're never quite sure what the tunes in the machine actually *are*, which means that a strong tune-dependence will only make the job more complicated. My simulations show that when the sextupole magnets are turned off completely, not only is the dynamic aperture larger, but it's also much less sensitive to the base tune. This suggests that turning off the sextupoles — at least initially — would streamline the tuning process.

BIBLIOGRAPHY

[1] M. Tigner and D. G. Cassel, “The Legacy of Cornell Accelerators,” Annu. Rev. Nucl. Part. Sci. **65**, 1 (2015).

[2] S. Belomestnykh, P. Barnes, E. Chojnacki, R. Ehrlich, R. Geng, D. Hartill, R. Kaplan, J. Knobloch, H. Padamsee, S. Peck, et al., in *Proceedings of the 1999 Workshop on RF Superconductivity* (1999), pp. 24–30.

[3] D. A. Edwards and M. J. Syphers, *An Introduction to the Physics of High Energy Accelerators* (WILEY-VCH Verlag GmbH & Co. KGaA, 2004).

[4] M. Sands, “The Physics of Electron Storage Rings: An Introduction,” SLAC Report **121** (1970).