A ROOKIE'S GUIDE TO BOOSTER OPERATIONS

BOOSTER TECHNICAL NOTE
NO. 231

K. ZENO

September 29, 1998

ALTERNATING GRADIENT SYNCHROTRON DEPARTMENT
BROOKHAVEN NATIONAL LABORATORY
UPTON, NEW YORK 11973

A Rookie's Guide to Booster Operations

K. Zeno

September 29, 1998

Contents

1	Welcome to the Machine	1
1.1	Lost in Space	1
1.2	What Kind of Accelerator?	4
1.3	Physical System or Machine?	8
1.4	Why was the Booster Made?	9
1.4.1	Protons	10
1.4.2	Heavy Ions	11
2	Device Timing	12
2.1	The Need for Timing	12
2.2	Timelines	15
2.2.1	The Supercycle	15
2.2.2	The Booster Real Timeline	16
2.2.3	The Booster Gauss Timeline	16
3	Instrumentation	18
3.1	Current transformers	18
3.2	Multiwires	20
3.3	Beam Position Monitors and Pick-Up Electrodes	22
3.3.1	The Booster BPMs and PUEs	22
3.3.2	LTB BPMs	25
3.4	Beam Loss Monitors	27
3.5	Wall Current Monitor	31
4	Controlling and Monitoring through the Controls System	33

5	The Booster's Transverse Space	41
5.1	The Equilibrium Orbit and Betatron Oscillations	41
5.2	Stability and the Booster Lattice	46
5.3	The Transverse Equations of Motion	51
5.4	Minimizing the Effects of Resonances	53
5.4.1	Tune Control	53
5.4.2	Stopband Correction	56
5.5	Dispersion and Chromaticity	66
5.6	Chromaticity Control	71
5.7	Dipole Errors and Orbit Distortions	72
5.8	Orbit Correction	75
6	The <i>Rf</i>	80
6.1	The Rf System without Beam Control	81
6.2	The need for an Rf System with Beam Control	83
6.3	The Rf System with Beam Control	85
6.3.1	The Phase Loop	85
6.3.2	The Radial Loop	87
6.3.3	Using the Phase and Radial Loops	88
6.4	Controlling the Rf Voltage	91
7	Heavy Ion Operation	96
7.1	Injection	96
7.1.1	Delivery and Transport to the Booster	96
7.1.2	Multi-turn Injection	98
7.1.3	Coupled Multi-turn Injection	107
7.1.4	Is Coupling Used During Injection?	108
7.1.5	Tips for Setting Up Heavy Ion Injection	109
7.1.6	Interpreting <i>Structure</i> on the Injection Transformer . .	114
7.1.7	The Tandem Foils	117
7.1.8	Optimizing the Injection Efficiency	122
7.2	Rf Capture and Acceleration	125
7.2.1	Capturing the Stored Beam	125
7.2.2	Acceleration	127
7.2.3	Early Losses and the Booster Vacuum	136
7.3	Extraction	140
7.3.1	An Overview	140

7.3.2	Synchronizing the AGS and Booster Cycles	141
7.3.3	The F3 Kicker	145
7.3.4	The Kicker, the Bump, and the Septum	149
7.3.5	The Extraction Bump and its <i>Residuals</i>	150
7.3.6	Setting up Booster Extraction	158
8	Proton Operation I: <i>The Usual Suspects</i>	166
8.1	LTB	166
8.1.1	Overview	166
8.1.2	Steering and Matching	167
8.1.3	Tips for Improving Transmission	172
8.2	Injection	180
8.2.1	The Foil as a Proton <i>Source</i>	180
8.2.2	The Linac Pulse	182
8.2.3	Bunches, Turns, and Micropulses	186
8.2.4	The Chopper Program	187
8.2.5	Stripping the Incoming Beam	190
8.2.6	Low Intensity versus High Intensity	192
8.2.7	The E.O. and the Injection Trajectory	195
8.3	Capture	203
8.3.1	Measuring the Revolution Frequency	203
8.3.2	Longitudinal Matching at Injection	206
8.3.3	Distinguishing Capture Loss from Other Losses	209
9	Proton Operation II: <i>Loss Mechanisms</i>	211
9.1	Stopbands	213
9.1.1	Correcting Stopbands Systematically	213
9.1.2	Measuring the Tune	216
9.1.3	Correcting Stopbands at High Intensity	219
9.2	Keeping Away From Stopbands	221
9.2.1	The Emittance Dependence of the Transfer Efficiency	221
9.2.2	The Tunes Near Injection Momentum	226
9.2.3	The Tunes During Early Acceleration	231
9.2.4	The Second Harmonic Cavities	233
9.3	Reducing Losses Due to Increased Emittance	239
9.3.1	'Limiting Apertures' Early in the Cycle	239
9.3.2	Radial Steering and Orbit Correction	251

10 Proton Operation III: <i>Extraction</i>	256
10.1 Multiple Transfers, <i>Synchro</i> , and <i>Cogging</i>	259
10.2 The Synchro Loop	262
10.3 Cogging	265
10.4 ΔB^* and the Extraction Radius	266
10.4.1 Finding the Extraction Radius	266
10.4.2 Measuring ΔB^* with the Gauss Clock	269
10.4.3 ΔB^* and the Tune Functions	272
10.5 ΔB^* and the Extraction Momentum	273
10.6 A Survey of Extraction Losses	275
10.6.1 Diagnosing Problems with Power Supplies	276
10.6.2 Losses While the Synchro Loop is Closed	281
10.6.3 Coherent Instabilities	288
A A Transfer Matrix Description of the Extraction Bump	291
B Typical Operating Values as of the end of 1997	293
C Selected Bibliography	295

Chapter 1

Welcome to the Machine

1.1 Lost in Space

Understanding specific details about beam physics, pulsed power supplies, the layout of the accelerator complex, and a myriad of other things, is usually irrelevant to the operation of the Booster (or the AGS for that matter). What is relevant is the knowledge and experience gained from operating it, which is difficult, if not impossible to convey in a book. Reading this book is not a substitute for that experience.

It is as important to know what you don't need to know, as it is to know what you do need to know. For example, you don't need to know how many degrees a Booster main magnet bends the beam. Each one bends the beam 10 degrees, but that knowledge won't help you to operate the machine, or even understand it in a meaningful way. What you might need to know is that the magnets are on, and the field inside them is at the value it was when the machine was last running.

Most of the knowledge that is of value is of a relative sort. There is no absolute running state for the Booster. The Booster runs relative to how it has run. A mainly empirically based understanding of how things fit together when the machine is running well helps when you need to get back to that point in the *space* when the machine isn't running well. If your lost in the woods, how do you get back to civilization? You use a compass, and any other tools at your disposal, you look for trails and familiar landscapes, you look for clues. Similarly, to operate the machine, you try to find out, as

well as possible, where the machine *is* relative to where you want it to be, (using the available tools). You use any method available to you to move the machine in that direction. If you know the *landscape* it helps alot, but you don't have to recognize every tree. If you know tools that you have used successfully to get your *bearings*, you are likely to use them again. If you know methods to "*get the machine back*" that have worked in the past, you are likely to use them again. You don't have to understand why the method works, just like you don't have to know why civilization is at the other end of a particular trail.

Many of the problems that we encounter can be traced back to equipment failures, often power supplies. There are diagnostics in the control room which we use to check this equipment. In large part, these diagnostics are *routed* to the control room, through the *Controls system*. The vast majority of equipment is also *controlled* through this system. Problems associated with the *Controls system* can prevent us from fixing, or finding, a problem with the machine. Chapter 4 deals with the controls system *interface*.

I would like to think that the *right* directions *out of the woods* of *Booster space* are found by applying physical principles to the machine. And that *the way out* can be found, if one can understand where the machine *is*, in terms of *the physics*. Beyond the day to day equipment and *controls* problems, there is the issue of improving the machine's *performance*. I have found, at least for myself, that a successful approach to bringing the machine into a better state usually depends heavily on trial and error, and observations. It doesn't rely nearly as much on *my* physical understanding. When the machine runs *better* after I have changed some parameter (e.g., the injection field), I sometimes try to understand why it runs better, but rarely come close to finding an answer that satisfies me.

Most of getting around in *Booster space* isn't necessarily difficult to understand, but it is hard to describe. So, if I am in any way successful in my quest to convey the little I have learned about this *Booster Space*, I will have drawn a highly imperfect map of it, and provided some clues on how to move about in it. This map is a strange mixture of apparently unrelated topics that really are related to machine operation. The hard part is finding the relations between these things in the day to day operation of the machine. Abstractions about physics, information about controls and instrumentation, and purely empirical observations of how the machine *behaves*, are all part of it.

An interesting example of the confusion that surrounds the Booster (or AGS) is the riddle of *the switch*. People sometimes come into the Main Control Room (MCR), after a *maintenance day* expecting to find the machine up and running the way it was before the maintenance day. A *maintenance day* is a day when the machines are intentionally shut down to do work on systems that need work, but which cannot be worked on while the machines are running. We try to do all the necessary maintenance on different systems during this one period. It is typically scheduled to last about 12 hours. When we try to bring the machines on after a maintenance day, invariably there are problems, and the machines don't just come back on. The *switch* that turns the machines back on doesn't work. Why is that?

One obvious explanation is that the machines were not restored to their condition previous to the maintenance day. It is almost always the case that during the process of bringing the machine back on, we find systems that have not been restored correctly, or no longer work the same way they did before maintenance. But, even after we think we have found all these problems, the machine still is not exactly the same as it was. Maybe the intensity is a little lower, maybe the transfer from Booster to AGS is a little different, maybe injection is worse. It seems reasonable to suppose that this is because the state of the machine has just *drifted* a little from its pre-maintenance state. So making small adjustments in the right places *should* bring the machine very close to where it was before.

However, once the machine is back on, but not back to *where it was*, people start making small adjustments, called *tuning*, to restore the machine to its previous condition. But after the first few changes are made, if they are not the *right* changes, the state of the machine is different than it ever was, and the way back to the pre-maintenance machine may become hopelessly complicated. Yet, it is very difficult to make the right adjustments which do *bring the machine back*. In fact, it probably never happens. But adjustments need to be made, because the machine has *not recovered* from the maintenance day, and the physics program *has to be restored*. It seems like we hope for some inherent stability that will keep us close to, or even guide us to, a *good state* of the machine, but we sometimes find ourselves at the edge of, or down at the bottom of, a cliff.

Tuning involves making small steps, which you can retrace if things go wrong (if you can remember them). Steps so small, that the change in the landscape is barely, if at all, discernable when you take them. Yet the

accumulation of these steps brings you to a *better* place if you move in the *right* direction. You try a lot of different directions, and sometimes try to go back if they obviously fail. You don't want to lose *your place*. You don't want to get more lost than you already are, and lose what you already have.

1.2 What Kind of Accelerator?

The purpose of the Booster is to act as an injector for the AGS. It accelerates both protons and other ions. Proton acceleration is distinguished from the acceleration of other ions for several reasons. First, the experimental physics associated with protons, called *High Energy Physics* is different than that associated with other Ions, called *Heavy Ion Physics*. From the machine perspective, the process of injection of so called *Heavy Ions* (ions which are not protons), is distinctly different from that of protons. A different *preinjector*, or injector for the Booster, is used for each case. For Protons, a 200 MeV Linear accelerator (The Linac) serves as a preinjector; for Heavy Ions, the Tandem Van De Graaf (The Tandem) is the preinjector. An attribute of the circulating beam which determines to a large degree what problems and what type of machine setup is involved is the beam intensity. The beam intensity is simply the number of particles circulating in the machine at a point in time. For Heavy Ions, the 'intensity' is generally of the order 10^9 (10^{11} nucleons); for protons it is of the order 10^{13} . This is another reason why there is a distinction made between the two types of particles.

Instrumentation, accelerating power, radiation, and the 'machine physics' are different for either case. These differences are reflected in the operation of the machine. For example, the radiation loss monitor instrumentation, which is used to pinpoint the location of beam loss, is rarely used during Heavy Ion operation. This is because the signal strength from these loss monitors (which represents beam loss) is very low. In most cases, the signals are indistinguishable from noise. This is due to the properties of the Heavy Ions in the machine (that is, energy, intensity, and strength of their interaction with matter).

The processes involved in bringing the beam from the Linac (or Tandem) to the AGS can be grouped sequentially. Each step in the sequence must, to some degree at least, be complete before the next step can be taken.

The first problem encountered is getting the beam coming from the prein-

jector to the Booster. This involves sending it through a transfer line from the preinjector. For protons this line is the Linac to Booster (LTB) transfer line. For Heavy Ions, it is called the Tandem to Booster (TTB) transfer line. The function of a transfer line is two-fold. The first function, is the transport of the beam to the Booster. The second function is more subtle. The beam coming from the transfer line must have properties which allow the Booster to accept it. Some of these properties, like its momentum, and momentum distribution are determined by the preinjector itself. Others depend on the transfer line. In particular, the position of the beam in the beam pipe, its size, how its size is changing (its divergence), and other properties depend on the configuration of the transfer line. These properties must fall within some range which *matches* the requirements of the Booster. This *matching* is accomplished by the transfer line. For example, the beam must be steered through the transfer line so that it reaches the entrance to the Booster.

Once the beam has reached the Booster, it has to circulate, or *survive* for many revolutions around the machine. The process of *storing* the incoming beam inside the machine is called *injection*. Once it is *stored*, then it must be *captured* by the accelerating Rf. Once captured it must be *accelerated* to the energy required for *extraction*. Once near this energy, it must be extracted from the Booster in such a way that the AGS can accept it properly. Each step in the sequence has its own problems. Some problems span several or all of the steps. One of the major goals is keeping the beam survival as high as possible as each one of these steps is taken. This is problematic for several reasons which depend on the species of particle accelerated (Protons or Heavy Ions), and the step in the sequence.

The beam is kept in a roughly circular orbit by a magnetic field as it is accelerated. A uniform magnetic field causes a charged particle moving perpendicular to it, to travel in a circle. The radius of the circle is proportional to the ratio of the particle's momentum to charge state, and inversely proportional to the magnitude of the magnetic field. This type of motion is called cyclotron motion.¹ Similarly, in the Booster, dipole magnets are placed at regular intervals in the beam's path to keep the motion of the beam roughly circular.

The beam circulates in a roughly circular metal pipe that is called the

¹A derivation of this can be found in most textbooks on Electromagnetism. For example, see *Electrodynamics* by Griffiths.

beampipe. The beampipe is under vacuum, the pressure inside it is about 10^{-13} atmospheres. If the beampipe were not under vacuum, charge exchange and scattering between the particles in the beam and in the atmosphere would destroy the beam. The beampipe is about 202 meters in 'circumference'. This is, as nearly as possible, 1/4 the length of the AGS ring.

The beampipe (also called the vacuum chamber) passes in between the pole tips of magnets which bend and focus the beam. The beam also passes through Rf cavities in its trip around the ring. Inside the Rf cavities there are a number of *gaps* in which an oscillating electric field, parallel to the motion of the beam, is applied. This electric field does work on the beam as it passes through the Rf cavity. The oscillations in the electric field have a frequency close to an integer multiple of the beam's frequency of revolution around the *ring*. As a result, a particle in the beam experiences nearly the same electric field every time it passes through an Rf cavity. The repeated interaction with this electric field on successive revolutions acts constructively on the particle to increase (or decrease) its momentum.

Since the number of gaps, over which the electric field acts doesn't change, the electric field is equivalent to an Rf voltage wave of a certain amplitude, or *gap voltage*, which provides a kick to the beam on every revolution, or *turn*.

As the beam is accelerated it must be kept at a nearly constant radius to keep it in the beampipe. If the beam's momentum increases, its radius can be kept constant by increasing the applied magnetic field. In the Booster, as in any synchrotron, a magnetic dipole field is applied at right angles to the beam's motion to keep it moving in a circle. The magnitude of this field is increased as the beam is accelerated to keep the beam's radius constant. The frequency of the accelerating Rf voltage is changed to account for the change in beam revolution frequency associated with its acceleration.

Unavoidably, particles within the beam have slightly different momenta and positions. As a result, the beam as a whole will get larger as time passes if nothing is done. Quadrupole magnets are used to *focus* the beam in the plane perpendicular to its primary direction of motion around the ring. These magnets produce a magnetic field gradient the sign of which alternates at different regular locations in the Booster ring. This alternating gradient acts like a series of lenses to keep the beam from diverging indefinitely. ²

²The theory behind this so called *strong focussing* was first layed out in 'E.D. Courant

Both the quadrupole and dipole fields are produced by electromagnets fed by the same current. This current is changed as the momentum of the beam changes to keep the radius of curvature (dipole field), and focussing strength (quadrupole field) independent of momentum. The current is supplied by the Booster Main Magnet power supply.

The time dependence of the Rf voltage is used to keep the beam contained along its primary direction of motion. A specific particle's energy and position in the ring cause it to pass through an Rf cavity at a different time, or different *phase* on the wave. This phase dependence of the particle's transit through the Rf cavity is used to confine the particle within a range of possible energies and positions. This is called phase stability, and it is an integral part of the acceleration process.³

The Booster was designed to store the beam inside it for as long as it takes for the Rf to bring this beam to extraction energy, but essentially no longer. Other accelerators, like The Relativistic Heavy Ion Collider (RHIC) are designed, not only to bring the beam to some energy, but to *store* it for hours. These machines are typically *colliders*. In colliders, two opposing beams pass through each other, and collide. Since the collision rate is very small, the beams collide with each other over hours without the amount of beam in the machine going down appreciably. The AGS, sends the beam to targets. These targets make use of the beam they receive immediately; the beam is not *recycled*.

Essentially, the Booster and AGS accelerate beam from injection to extraction the same way every time they do it. Once the beam has been extracted from the Booster, its parameters are restored in preparation for the next *cycle*. The Booster cycle is generally on the order of 130-200 ms long. The length of the cycle may vary considerably depending on the type of particle being accelerated, its extraction energy, and other factors relating to the machine's configuration. This is particularly true for Heavy Ions. Most recently, the extraction energy for Protons was 1.96 GeV. For Au^{+32} , it was

and H.S. Snyder, "*The Theory of the Alternating Gradient Synchrotron*", Ann. Phys. 3,1 (1958).'

³An excellent and rigorous explanation of 'Phase Stability' in a Synchrotron is given in Chapter 2 of *An Introduction to the Physics of High Energy Accelerators*, by D.A. Edwards and M.J. Syphers, John Wiley and Sons, 1993. This book is also a great resource for anyone who would like a real understanding of other physics that goes on in synchrotrons such as the Booster.

95.2 MeV/nucleon. This *energy* is the kinetic energy, $E_k = (\gamma - 1)mc^2$. For Heavy Ions, the kinetic energy is given *per nucleon* (Gold has 197 nucleons).

1.3 Physical System or Machine?

Webster's dictionary defines a machine as, "*A system, usually of rigid bodies, constructed and connected to change, transmit, and direct applied forces in a predetermined way to accomplish a particular objective, as performance of useful work.*". The Booster, though it fits this definition quite well, also appears to fit a broader definition. The Booster is studied as a physical system by physicists, and its operation is far from being completely understood, as near as I can tell. The "*performance of useful work*", which I equate with the production of beam by it for the AGS, is *not* accomplished in a "*predetermined way*". The way it is produced is always in flux, and it is studied as if it were a physical system that is not completely understood. Somehow, this entity, the Booster, is a cross between a machine, and an apparatus for a physics experiment. You can treat it like a machine, or study it like a physical system, and both things are done. Sometimes these *ideas* about what the Booster is conflict. People look at the Booster (and AGS) in very different ways. I have heard the AGS described as a Power Amplifier by one technically minded person, and the Booster as a 30 million dollar attenuator by another (when during its early operation it initially *reduced* the output of the AGS).

You can treat it like a machine, and fix its parts if it breaks. You can find out what is wrong with it, if it doesn't work the way *its designed to work*. You can improve the parts to make it work better, and increase its output. Or you can think of it as a physical system, perform an experiment (or *study*), and hopefully gain some physical understanding. Then you can test that understanding by putting this physical system into a different, hopefully *better* state, and measuring its properties. If these properties show that the new state is more in line with the desired state of the system, then you *leave the changes in*. Now *the machine*, has changed its configuration, and its output has also changed, not as a result of changing a *part*, but by changing what it *is*. It is not uncommon, after such a change, to say that *it* is "*a different machine*".

The hybrid, and often conflicting nature, of people's ideas of what the

Booster is, are illustrated well by the term “*machine physics*”. One hears this term often in and around MCR. One thing is for certain, the reason *for* the Booster is to produce beam for the AGS, in order to meet the needs of the Physics programs.

Tuning is the approach used most often by operations to improve the machine’s performance. It involves trying to improve its performance using trial and error, while not necessarily constraining the things one tries, to what one understands. For the most part the changes to the *machine setup* take place below the level where the machine is well understood. Since, if you don’t know how things are supposed to fit together exactly, how can you say that it is wrong to change one thing or another by some small amount? If changing that one thing improves the performance of the machine, then why not leave it that way? For example, say I change the injection field in the Booster by a small amount, and the amount of beam injected increases dramatically. Say also, that I don’t know exactly what the field in Booster is, and I don’t know exactly where in the beampipe the beam should be for most efficient injection. Then why shouldn’t I leave the field at that value? At least as a starting point for further changes?

The Booster can be reconfigured in a vast number of ways. This flexibility allows one to study it as one would study the states of a physical system. It is like a computer that can be programmed to run an essentially infinite number of programs. The *Booster space* is related to the set of all possible configurations of the physical elements that comprise the Booster (and the beam that inhabits it). Most of these elements can be reconfigured through the *controls system* to make different *machines*, which are like different programs. Much of the following chapters will be concerned with the details of how these elements are configured through the *Controls System*.

1.4 Why was the Booster Made?

The Booster was made for two main reasons. For protons, the primary effect of the Booster is to increase the intensity of the beam delivered to the experiments. The intensity is one of the most important parameters for the experimenters. The Booster was also designed with Heavy Ions in mind. The Booster allows the acceleration of very Heavy Ions such as Gold. These *Heavy* Ions are used in experiments which study the collisions between their massive

nuclei, and the nuclei in targets often composed of similarly heavy atoms. The extremely high mass and energy densities that result are of interest in these experiments.

1.4.1 Protons

High Energy Physics (HEP) experiments at the AGS investigate rare decays of particles. The more beam they receive, the more likely they are to see the decay, or event, that they are looking for. The limitation on the experimental side, as far as how much beam they can use, has been the amount of data they can acquire from their detectors for analysis. As electronics, computers, and data acquisition have improved through the years, the amount of beam the experiments can use has also increased.

The intensity of the linac proton beam is such that the repulsive force that the beam has on itself from its *self-charge* can significantly counteract the effect of the magnetic focussing field. The result of this is a beam which is prone to many resonance conditions which can cause much, if not all, of the beam to be lost. The strength of this repulsive effect is proportional to the length of an accelerator ring, and decreases as the energy of the beam increases. At lower energies, the beam in a smaller accelerator ring is less susceptible to effects associated with the beam's self-charge. At higher energies these so-called *space charge* effects become less important. All but the ring circumference being equal, a smaller ring can accelerate more beam to an *intermediate* energy, than a large ring can. Once at that intermediate energy the beam can be transferred to a larger ring, where it can be accelerated further with less space charge related problems.⁴

Why not just accelerate it in a smaller ring to begin with? Well, first of all, the AGS has already been built. Secondly, higher energies are easier to accomodate in a larger ring because the magnetic field strengths required are less.

⁴Stephen D. Holmes, "*Low Energy Aspects of Circular Accelerators*", Fermilab- Conf-90/275, December 1990. This paper discusses the way space charge limits the intensity in a circular accelerator, and the use of a Booster to lessen the problem.

1.4.2 Heavy Ions

The acceleration of partially stripped ions such as Gold from the Tandem Van de Graaf energy requires a vacuum several orders of magnitude greater than that found in the AGS. The partially stripped ions have high cross sections for electron capture and stripping from the *residual gas* in the vacuum chamber. If an electron is lost, or gained, by an ion, it will be bent and accelerated differently, and will be lost from the machine.

It is thought that electron capture dominates at lower energies, and electron stripping becomes more important at higher energies.⁵ *Stripping foils* are used to remove electrons from the Gold atoms at Tandem. The Gold beam received from the Tandem is generally stripped of less than half its electrons before entering the Booster. When it enters the Booster the cross section for electron capture is relatively high, and an extremely good vacuum is required. After being accelerated by the Booster to AGS injection energy, it is extracted from the Booster, and stripped of almost all of the remainder of its electrons as it passes through the transfer line between the Booster and AGS. At this point, its energy is high enough that the electron capture cross section will not be too high even for the highly stripped Gold with the AGS vacuum. The electron stripping cross section decreases dramatically when its charge state is increased. So, the vacuum in the AGS is adequate for further acceleration of the high charge state Gold at these higher energies.

⁵H.M. Calvani and L.A. Ahrens, *Au³²⁺ Beam Intensity Losses in the AGS Booster Due to Charge Exchange Processes*, Booster Technical Note No. 228, October 1, 1996.

Chapter 2

Device Timing

2.1 The Need for Timing

The accelerators operate *in concert* to accelerate the beam from its creation at its source, through its extraction, and delivery to the experiments. The synchronization between accelerators, and devices within each accelerator, is accomplished by a distribution of timing triggers. Devices which require triggers synchronized to the acceleration cycle are connected to so-called *timelines*. These *timelines* are like timing busses inside a computer. Timing triggers are sent out along a *timeline* in the form of a digital code. Devices in the complex are equipped with timeline *decoders*. These decoders allow the devices to extract the triggers appropriate to their operation from the timelines.

For example, the Booster Main Magnet power supply¹ is programmed to start to follow a reference function when a time line trigger, corresponding to a specific hexadecimal number, occurs on a time line, and is 'decoded' by the power supply's time line decoder. This specific trigger, typically called a timeline *event*, is given a 3 letter acronym. In this case, the hexadecimal number is 000A, and its acronym is BT0. BT0 stands for Booster-Time-zero. The timeline *event* is given this name because it triggers the start of the Booster Main Magnet cycle. Since this timeline *event* (BT0) is distributed

¹A detailed description of the Booster Main Magnet Power Supply is given in: "AGS/AD/Op. Note No. 37, "Topics in Booster Main Magnet Operation", by Pepin Carolan, November 30, 1992

to all devices on this timeline, any other device may also be programmed to respond to it.

Another set of devices that uses BT0 as a trigger for its reference functions are the 'Tune Quad' power supplies. The tune quad power supplies feed current to coils wrapped around quadrupole magnets in the Booster. They are used to make relatively small adjustments in the focussing strength of these quadrupoles. Small adjustments like these are called 'trim' adjustments. For a given current in these coils, the focussing strength they have depends on the beam momentum.

For a specific particle type, there is only a small range of momenta that are found within the machine for a given Main Magnet (dipole) field. Why is this so? The beam must remain in the vacuum chamber (beampipe) to survive, and the chamber is only several inches wide in the horizontal plane. For a given Main Magnet field, a particle traces out a roughly circular *orbit* which has a *radius* determined by its momentum. If that radius falls outside the vacuum chamber, the particle will not survive. Although the Main Magnet field does not *do work* on the beam to increase or decrease its momentum, its momentum *is* constrained to be within a small range by the field's magnitude.

So, there is a relation between the magnitude of the Main Magnet field, and the current in these 'trim' coils, needed to produce the desired focussing strength. The required values of both of these are bound together by the value of the beam's momentum. The Main Magnet field, is a function of the current output of the Main Magnet power supply (magnetic field is roughly proportional to current). This current is very nearly a function of time from BT0 (current output follows its reference function).

How do you get the current in the coils to be related to the Main Magnet field? One way is to make the current in the coils a function of the time from BT0, and use the Main Magnet current (or field) reference function to find the field as a function of time, $B(t)$, in the cycle. This reference function exists as a data file accessible to any software in the computer network. From the desired focussing strength as a function of time, and the Main Magnet reference function, the tune quad current as a function of time can be calculated. There are other possible solutions to this problem. The above solution is roughly the way the designers of the Booster have chosen to solve it. A software application (OpticsControl) uses the Main Magnet function, and another function, which is essentially the desired focussing strength, to derive a current reference function for the tune quads that is a function of

time from BT0 (see figure 5.4).

This example illustrates why 'synchronization' between devices is important. In order for the beam to have the desired properties (momentum, size, intensity, etc), devices must act together, and evolve together in a particular way. The specifications for this synchronized evolution of device 'outputs' is determined by what is needed to achieve the desired beam characteristics, and the evolution of those characteristics. 'Jitter' that devices have with respect to each other, from cycle to cycle, may cause instability from cycle to cycle, and poor beam quality. An example follows.

For the beam to be injected into the Booster, and circulate, the magnetic field in the Booster must be within a small range of values. As mentioned earlier, for a given momentum, the beam radius is determined by the magnetic field. The beam momentum from Linac is fixed. It is about 640 MeV/c. This means the field has to be about 1500 gauss for the beam to fall on a radius approximately centered in the Booster vacuum chamber. For every +10 gauss (g) change in field, the radius changes by about -1 cm. If the *aperture* of the beampipe is, say 8 cm, then a change of 40 gauss from the 'ideal' field will force the beam's orbit outside of the beampipe, and the beam will not circulate.

For proton injection the Main Magnet field is typically increasing in and around injection time at about 30 g/ms. If the time the beam is injected relative to the 'ideal field' varies by more than $\frac{40}{30}$ ms, then the beam will not circulate in the machine, since the field will either be too high, or too low to keep it in the beampipe. The Booster Main Magnet Power Supply (BMMPs) follows a current reference that is a function of time from BT0. A timeline event is generated just before injection time. This timeline event is used to start a chain of events at Linac whose end result is the delivery of beam to the Booster. The Linac beam is delivered to the Booster about 2500 μ s after this timeline event.

Suppose that the injection timeline event occurs with a fixed time interval between it and BTO. Since the BMMPs follows a reference function that is a function of time from BTO, one might expect the field at the time that the beam enters the Booster to be reproducible from cycle to cycle, and set by the value of the current in the BMMPs reference function. However, there is some jitter between the actual field value at injection time, and its value derived from the reference function. There are several factors which might contribute to this jitter. The BMMPs is not perfect, and it does not follow

the reference function exactly. The temperature of the magnets effects the magnetic field dependence on current, and this temperature depends on a variety of factors (cooling water temperature, ambient temperature). The line voltage (from LILCO) also varies at some level, as do numerous other factors. Since the field value at injection is so critical, a different type of trigger is typically used. This trigger is generated by an instrument which measures the field within a Main Magnet dipole.

The instrument that measures the Main Magnet field is called the gauss clock.² A separate 'timeline' exists for triggers (or events) generated when the field value reaches a predetermined value. The *Gaussline* event that is used for injection is called BIJ.PEAKER.GT, or Peaker. It occurs when the field value is slightly lower than it is at injection. 2500 μ s after this event beam enters the Booster. This event occurs *close* to injection so that the jitter in the field from cycle to cycle at the time the beam from Linac enters the machine will be less.

2.2 Timelines

There are three timelines used to synchronize devices 'to the beam' in the Booster. They are the Supercycle, the Booster Gauss Timeline, and the Booster Real Timeline. A brief explanation of each follows.³

2.2.1 The Supercycle

The Supercycle is determined by a set of events from which all other events are delayed in one way or another. In this sense, the supercycle timing is the highest level of timing. Among other things, supercycle events define the duration and number of individual Booster cycles that occur within a supercycle. The 'resolution' of these events is $\frac{1}{60}$ second. The length of this supercycle, which generally includes one or more Tandem or Linac cycles, one or more Booster cycles, and one AGS cycle, is usually about 3 to 5 seconds. It is the mechanism which enables the machines to act in concert

²The Gauss clock is described in Booster Technical Note No. 175, "A Digital Voltage to Frequency Converter for the Booster Gauss Clock", By J. Geller, 7/25/90

³A detailed survey of Booster Timing is contained in AGS/AD/Op. Note 38, "Overview of Booster Timing", by K. Zeno, et al., February 17, 1993

to accelerate the beam from the source through extraction from the AGS. The software application through which the supercycle is setup, and *made active* is called "Superman". A Booster cycle is identified with a group of setpoints and lower level events (timing) through its 'PPM user'. PPM stands for pulse to pulse modulation, and refers to the fact that machine cycles within the same supercycle, for the same machine, can be run differently. For example, ppm user 1 may consist of a Booster cycle with a Main Magnet cycle which reaches 5400 gauss and is used as part of the acceleration cycle for the physics program. During the time the AGS is accelerating the beam for the physics program the Booster may be idle. However, that time within the supercycle could be used to study Booster phenomenon at injection energy (corresponding to about 1600 gauss). A ppm user, say user 2, may be defined on the supercycle with a magnet cycle that stays at injection energy. For this user, a group of setpoints and lower level timing can be sent to devices. This user can be set to occur during the time the Booster would have been idle, by programming it to occur at the appropriate time in the supercycle.

2.2.2 The Booster Real Timeline

These events are usually delayed from the supercycle event "BTO" (Booster T zero). These events have a resolution of one microsecond. For example, the current in a system of dipoles, used to correct the beam orbit, are set to follow a reference, or function (set through the BoosterOrbitControl application). This function may be set to begin at 12000 μ s after BTO. In this case the real time event "BRT.ORB_CORR.ST" is set to 12000 μ s. These events can be set through the application "Spreadsheet".

2.2.3 The Booster Gauss Timeline

These events occur when the Booster Main Magnet dipole field reaches a specified value. Often there are short real time delays from these events which are used to control processes which are highly dependent on the magnetic field. For example, the delay BIJ.FAST.TM is from the gauss 'time' event BIJ.PEAKER.GT. Both the delay and gauss 'time' event can be set through Spreadsheet. The delay is used to trigger a set of 4 dipoles which distort the dipole field as part of the injection process. These 4 dipoles are collectively known as the injection kickers, and make a bump in the orbit. The current in

these dipoles starts following a time varying reference (also set through the BoosterOrbitControl application) when the trigger BIJ.FAST.TM occurs. Since the Main Magnet dipole field, as a function of time from BT0 (its starting point), does not reproduce exactly from cycle to cycle, it is sometimes better to trigger devices (processes) according to the value of the field instead of when they should occur in time. In this example, BIJ.PEAKER.GT (or peaker) might be set to 15000 gauss clock counts, and BIJ.FAST.TM set to 2600 μ s. The gauss clock keeps track of the main dipole field, as a clock would the time. It 'ticks' in gauss clock counts instead of seconds. Each gauss clock count corresponds to a change in field of 0.1 gauss.

The field is measured by a Hall probe in a main dipole magnet outside the ring called the 'reference magnet'. It is in series with the Main Magnets in the Booster ring. Unlike time, field can go up and down, so the gauss clock has up and down counts. Every cycle the field is initialized with the value the hall probe measures. This initialization typically occurs just after BTO. The time is set by a real time event called BGT.JAM.VALUE.RT, which is usually set to 10 μ s. There is a spreadsheet 'device' called BGT.CALIBRATE which must be set to 'AUT', which stands for 'auto', for this to occur. BGT.CALIBRATE may also be set to 'MAN', which stands for 'manual'. In this case the value that the gauss clock is initialized to is contained in the reference value field for BGT.CALIBRATE on spreadsheet. The field gauss clock indicates the sum of the counts plus the initial value.

The Gauss timeline is initialized by the event BGT.SFIELD.BUFFER. Events that occur after the field set in BGT.SFIELD.BUFFER is reached are generated and sent out over the Gauss timeline. If BGT.SFIELD.BUFFER does not occur, say because the field is always higher than that set in BGT.SFIELD.BUFFER, then there will be no events that occur for that cycle. This is often the source of confusion.

The references that power supplies follow are not always constant (i.e.-Main Magnet, correction dipoles, injection kickers). These references consist of points. The Y axis is the amplitude, the X axis is time. The units of each may vary. To transform these mathematical reference functions into a time varying reference for power supplies, clocks are employed. For example, a 500 KHz clock is usually employed for the injection kickers. In this case, the time between points in the reference function is 2 μ s.

Chapter 3

Instrumentation

This section describes the instrumentation available and the tools used for monitoring it in MCR.

3.1 Current transformers

The beam current and intensity is measured using beam current transformers.¹ Current transformers convert the beam current into a voltage. The relation between the voltage from the transformer and the beam current is assumed to be linear. The calibration is determined from sending a pulse of known current (not the beam) through the transformer and measuring the voltage. Voltage signals from current transformers are available in MCR through the Xbar program. There are two current transformers in the Linac to Booster (LTB) transport line. One at the upstream end (XF011, 11 feet downstream of the first magnet in the line, DH1), and one at the downstream end (XF100). There are many in the Tandem to Booster (TTB) transport line (in sections 11, 16, 18, 25, 27, and 29). They are typically viewed as analog signals.

There are two current transformers in the Booster. The so-called *injection* transformer has the faster response time of the two. It is called the injection transformer because it was intended for use at injection. Injection is a fast process, and significant beam current structure can exist on the 10 μ s time

¹For more information see AGS/AD/Op. Note No. 35, "Booster Current Transformers and Booster Wall Monitors", by Jonathan Reich, November 13, 1992

scale. It treats the beam as a current winding in a voltage transformer. The other transformer, the *circulating* transformer, is not a transformer in the classical sense. It responds to a constant current with a voltage. Its response is not as fast, and so is used to look at processes that occur on a slower time scale. It is called a "DC current transformer" or DCCT. Both transformers are located in the same 'housing' at C6, which is near injection for both protons and Heavy Ions. Although the injection transformer was intended for use at injection, it is also often used throughout the acceleration cycle. There have been some problems with the response of the circulating transformer to high intensity protons. The time response of the injection transformer has been modified so that it can look at the entire acceleration cycle during these high intensity operations

The *calibrate pulse* is a pulse of known current that is sent through a current transformer to calibrate it. It is turned on/off through spreadsheet. For the Booster transformers, the calibrate pulse is turned on by setting two SLDs (simple logical devices) on spreadsheet. The SLD, BMD.INJXF.CAL, is set to CAL, and the SLD, BTG.BCBM.CAL.RT (a real timeline event), is set to ON. The value of the setpoint field for BTG.BCBM.CAL.RT determines when in the cycle it occurs. Its duration is on the order of milliseconds. It can be seen on the analog signals for the current transformers. It appears as a square pulse. The amplitude of the square pulse is measured to obtain a calibration. For example, if the calibrate pulse is 0.5 mA, and the height of the pulse is 5V, then the calibration is 0.1 mA/V. Both current transformers have different gain settings. The current of the calibrate pulse is typically different for different gains, as is the signal's response to it.

The Booster current transformers provide a measure of beam current. Often, one wants a measurement of the beam intensity, not the current. If the beam in the Booster is accelerating, then the current for a given beam intensity will increase with time. Therefore, the amplitude of the analog signal typically increases during the acceleration cycle. To get the number of particles in the beam, at a time in the cycle, the following formula is used:

$$N = \frac{6.24 \times 10^{18} \text{ ch/Coul}}{Q} I T_{rev}$$

N is the number of particles, I is the beam current (in Amperes), T_{rev} is the period of one revolution, and Q is the charge state. T_{rev} is the same as $1/f_{rev}$, where f_{rev} , the revolution frequency, is f_{Rf}/h . f_{Rf} is the Rf frequency,

and h is the Rf harmonic, corresponding to the number of bunches in the machine. The beam current, I , is obtained from measuring the voltage of the current transformer signal. The voltage is converted into current using the calibration obtained from the measurement of the calibrate pulse.

The current transformer signal can be divided by the Rf frequency to obtain a signal which is proportional to the beam intensity, not the current. This version of the current transformer signal is said to be normalized. The Booster current transformers are available either normalized or unnormalized through the mux system. The Booster beam intensity scalars at each MCR console are derived from a normalized signal.

3.2 Multiwires

Multiwires are planes of thin conducting wires (about 32) which are inserted into the beam pipe. They are oriented at right angles to the direction of the beam (half aligned vertically, the other half horizontally, evenly spaced in each plane). When inserted into the beam, particles within the beam hit the wires causing electrons to be *emitted* from the wires. Charge flows into each wire to restore the neutral charge of the wire. This charge is measured, and is proportional to how much beam hits the wire. This information yields information about the distribution of particles within the beam. In particular, the beam's center of mass, shape, and width. The shape of the beam is typically gaussian, though it is not always gaussian.

Beam *profiles* in the horizontal and vertical planes are obtained from the BeamLineInstrument program. Figure 3.1 shows a profile for Gold +32 beam in the TTb line. There are no multiwires in the Booster ring. They are generally only used in transport lines (though there is one at A20 in the AGS). Inserting a multiwire into an injection line does generally change the Booster performance. They are left out of the beam when not in use.

TTb multiwires are described by their section number and feet from the beginning of the section (sec.# /MW/ feet). There are 26 multiwires, they can be viewed through BeamLineInstrument. Usually, when working with the Booster only the last 4 multiwires (sections 27-29) are viewed in MCR. LTB has two multiwires, one at 35 feet (MW035) and 107 feet (MW107) from the beginning of LTB, which starts at the DH1 magnet.

Figure 3.1: Multewire profiles for Gold +32 in the TTB line

3.3 Beam Position Monitors and Pick-Up Electrodes

3.3.1 The Booster BPMs and PUEs

The center of charge (or mass) of the beam inside the Booster is measured by a pair of electrodes. At an azimuthal location a pair of electrodes is placed on opposite sides of the beampipe. They measure either the vertical or horizontal position relative to the center of the vacuum chamber (beampipe). Roughly speaking, there is a cylinder, shaped like the vacuum chamber, which the beam passes through. This cylinder composed of a conductor, is split into two pieces, or plates on a diagonal. The plates are insulated from each other. When the beam passes through them each plate has a voltage induced on it. If the beam passes through the center of the two plates, the voltage on each plate is the same. In the case of a vertical pair of PUEs (Pick-Up Electrodes), the diagonal cut extends accross the horizontal sides of the cylinder, forming a 'top' and bottom' plate. If the beam position is high, relative to center, then the top plate charges up to a higher voltage. By taking the difference in voltages, normalized to the beam intensity, a position measurement can be made. Mathematically the position, over a considerable range, is simply related to the potentials on the two plates,

$$y = k \left(\frac{V_{top} - V_{bottom}}{V_{top} + V_{bottom}} \right)$$

where y is the vertical position relative to the center of the plates, k is a calibration constant, and V_{top} and V_{bottom} are the voltages on the top and bottom plates respectively.

There are 2 groups of PUEs in the Booster ring. One group is for the horizontal plane, the other is for the vertical plane.² The horizontal PUEs are located at even numbered locations (i.e.-A2, A4, A6, A8, B2...), The vertical PUEs are at odd locations (A1, A3, ...). A few (about 3-4) are routed as analog signals to the control room. They come to the control room either as sum and difference signals for a pair of plates, or as signals from the

²*The Booster Design Manual*, December 5, 1988, pages 8-1 through 8-4, contains information on the original design and specifications of the PUEs.

individual plates. They are either hardwired or are available through Xbar (software routing). They can be viewed on a scope.

The BPMs (Beam Position Monitors) are derived from the PUEs. The signals from the PUEs are digitized. This digitization is done in order to measure the closed orbit in the Booster. The data from all the PUEs can be taken at one time in the Booster cycle, digitized, and turned into a controls system data object known as a CLD (Complex Logical Device). This CLD can be accessed by the BoosterOrbitDisplay program to give a representation of the closed orbit.

The digitization of the PUE signals must be synchronized to the time the individual bunches pass the PUEs. This synchronization is obtained from the Rf system. The BPM data is generally obtained by an integration of the PUE signal over one or more (as many as 255) bunches as they pass through the PUE. The gain of the BPM system can be changed through the program. The gain settings are $\times 0.1$ (signal attenuated to 1/10th), $\times 1$ (no attenuation or amplification), and $\times 10$ (signal amplified 10 times). For high intensity protons $\times 0.1$ is generally used; for heavy ions $\times 10$. The signal strength can also be changed by changing the number of bunches over which the integration is made. This is done through the program. The strength of the signals can also be viewed through the program.

There are several PUEs (associated with BPMs) that are used for specialized purposes. Two horizontal PUEs are used in the Rf system to measure the radius of the closed orbit (typically the PUEs at A2 and A8). The gain of these PUEs is set independently. There are also four more PUEs whose gains can be set independently. Two of these are used for measuring the betatron tune. All these gains are set through spreadsheet.

Figure 3.2 shows a typical closed orbit representation for proton running. The top graph is a representation of the horizontal closed orbit, the bottom the vertical. The ordinate is distance from the center of the beampipe in millimeters, the abscissa is the azimuthal location in the ring. The data points are shown, and straight lines connect them. In this figure it is easy to see that this particular horizontal closed orbit lies toward the inside of the beam enclosure. In both planes the closed orbit appears to oscillate around some value. An harmonic analysis of the data can be done within the program. The information on the left in the figure is specific to the orbit that is displayed, most of it is specified when the orbit is taken.

Figure 3.2: Closed Orbit representation using Booster BPM system and BoosterOrbitDisplay application

3.3.2 LTB BPMs

There are seven BPMs in the LTB line. These are referred to as Stripline BPMs. The principles behind their operation are more complicated than for those employed in the Booster. The horizontal and vertical BPMs are contained within the same "package". So, wherever a BPM is, both the horizontal and vertical positions are measured. The 'striplines' are plates that are longer than they are wide, aligned parallel to the direction of the beam, inside the beampipe. There are two pairs of striplines, one aligned with the horizontal plane, on either side of the beampipe, which is used to measure horizontal position. The other pair is aligned with the vertical plane, and is used to measure the vertical position.

The data from the BPMs is acquired either through the BeamLineInstrument program, or as analog signals through the multiplex (mux) system. The data viewed in the program is obtained from analyzing and digitizing data from part of the Linac pulse. The data viewed on the scope is the signal from the striplines before it is analyzed. It contains information about the time variation of position. However, the conversion of the signal voltage as a function of time into position is complicated.

Figure 3.3 is a sample LTB position measurement from the BeamLineInstrument program. Again, the data points are connected by straight line segments. The units of position are millimeters from the center of the beampipe. The S (or longitudinal) coordinate is distance from the start of the LTB line (from DH1) in feet. The top graph also records data from the LTB current transformers. The data from the BPMs in the top graph is obtained from the sums of the plates for the BPM, and could be used as a measure of intensity. However, this data is generally not used for intensity measurements. It can be used to see if the BPM is set up or working correctly. That is, the sum is an indicator of the signal strength. If it is too close to zero, the measurement will not be reliable; if it is too large, it may also not be reliable. A small signal may indicate that the measurement is being taken when there is little or no beam. If the signal is too large, then the signal might be saturating. The gain of the system may be controlled through the program, it is called the 'integrator gain'.³

³Information on the operation of the LTB BPMs, particularly with bunched beam, was obtained from a discussion with Dave Gassner

Figure 3.3: LTB BPMs as viewed through the BeamLineInstrument application.

There is some work required to insure that the BPMs are set up. There is a time window, or gate, during which the measurement is taken. This gate must occur while the beam is passing through the BPM. The time it occurs is a delay (in microseconds) from peaker, called the sample time. The width of this gate (Integrator hold) is also adjustable. A gate signal is available in MCR. It is put on a scope, together with a beam signal (either an LTB BPM or current transformer) and lined up so that it occurs during the beam pulse.

The BPMs operate on either chopped or unchopped beam. The setup is different depending on which type of beam is being measured. Chopped beam is produced by the Linac fast chopper. It is used to match the beam to the Rf in the Booster. The Linac "macropulse" is typically from 100 μ s to 500 μ s long. Within that 'macropulse' it may be chopped into smaller pulses, called turns. These pulses are generally less than 300 ns in width. Their width is described in terms of their width relative to the Booster Rf wave in degrees. It can range from $\sim 80^\circ$ to 360° . If the beam is being fast chopped then the 'mode' should be set to T/H, (track and hold). In this mode the BPM electronics are synchronized to the Rf (and chopper), and are able to look selectively at the parts of the macropulse where there is beam. This mode works only for turns whose width is between 180-330 degrees. If the fast chopper is off, then the mode 'TRK' must be used. Both these modes may be set through the program, or through spreadsheet. As can be seen in the figure, the statuses of the chopper, mode, sample time, and integrator hold are displayed in the program.

3.4 Beam Loss Monitors

The Booster Beam Loss Monitor system is used to isolate losses that occur in the Booster cycle to a place and time.⁴ Beam loss indicated on the Booster current transformers (by a negative slope) indicates that there is beam loss, but it does not indicate where that loss occurs. The Loss monitors are

⁴For technical details about the loss monitor system see: E. Beadle and G. Bennett, "Booster LRM System Hardware Specification", Booster Tech. Note No. 167, June 7, 1990. For a more operations oriented perspective see, OPM 6.12.11, "Booster Loss Monitor".

coaxial cables filled with Argon gas that are held at a potential. Beam loss causes radiation which ionizes the gas in the cables. A current flows due to this ionization. This current charges capacitors. Like the multiwires, these 'ionization chambers' operate in a voltage range where the current is proportional to the initial ionization by the beam. Hence, the charge (or voltage) on the capacitors reflects the amount of beam loss.

Most of these cables are located approximately 30 inches radially inward from the center of the beampipe, lying parallel to the direction of the beam, and slightly below beam elevation. There are 48 of these loss monitors in the ring. Each of these loss monitors is associated with one of the 48 subdivisions of the Booster ring. There are 8 loss monitors in each of the LTB and BTA beam lines, and 16 more that can be moved.

As with the BPMs, the voltage signal from the integrators can be viewed as either an analog signal through the mux system, or as digital data through the controls system. The configuration of the loss monitor system is rather complicated. There are two principal ways in which the data is typically viewed. The first is through the GPM (General Purpose Monitor) program. This program is used to create video displays, or software window displays of controls system data. In this case displays are created which look at points of particular interest in time (injection, extraction, etc.), and place (LTB, BTA, etc.). These displays are often used to monitor losses continuously.

In addition to GPM, the `BoosterLossMonitor` program is used to monitor losses. This program is also used to configure the loss monitor electronics. Figure 3.4 is a display of the `BoosterLossMonitor` program. The losses, displayed as 'counts', are in tabular form. They can also be displayed graphically. The loss monitors within the ring are named by their ring location. LTB loss monitors (top of the left column) are designated by their distance from DH1. BTA loss monitors, most in the bottom of the right column, are designated by their distance from the extraction point (F6). Some of the movable loss monitors have been assigned names which reflect their present locations. For example, C6_UPST is a movable loss monitor which is located at the upstream end of C6.

The losses are monitored over intervals of time within each cycle that are called 'windows'. At the end of each window the integrators in the loss monitor electronics are reset. The typical configuration consists of 5 windows. The first window is before injection. The second looks at injection. The third encompasses most of acceleration. The fourth window looks at

Figure 3.4: A tabular display of Booster losses using the BoosterLossMonitor application window

extraction, and the fifth occurs after extraction. These 'windows' are setup in this program. There can be as many as 16 windows. These windows can be changed easily. The window times are shown in the third line in the figure. 'Window 1' extends from Booster T zero (BT0) to BT0+10.2 ms, 'window 2' from BT0+10.2 ms to BT0+20 ms, and so forth. The GPM files that were made to look at a particular process in the cycle, say extraction, are set to look at losses in a particular window during which extraction typically occurs. However, there is no guarantee that the window timing has not changed since the file was made, or that the time extraction occurs has not changed since it was made.

The counts displayed for each loss monitor are proportional to the integrated current over the window. This integrated current, which is a voltage, is proportional to the amount of ionization. For a given momentum and particle type, the amount of ionization is proportional to the number of particles passing through the chamber. This leads one to develop a loss monitor calibration for a given momentum. For protons, at extraction momentum (2.3 GeV/c), this calibration is approximately 100,000 counts/ 10^{12} particles (for low gain). The response of the loss monitors may vary due to the specific character of a loss, but this variation is often neglected. The loss monitors are used primarily during proton operation. Their sensitivity is such that even in high gain it is difficult to see losses with Gold Ions, particularly at injection energy.

The loss monitors are also used to disable the Linac beam. There are two ways in which the loss monitors can disable the Linac beam. The first way uses the Fast Beam Inhibit (FBI) feature of the Linac. In this case, the Linac beam is inhibited when the integrated losses from a loss monitor reach a predetermined level which is set at Linac. This inhibit occurs while the Linac beam is being injected. It has the effect of terminating the Linac pulse prematurely. The LTB loss monitors typically cause this type of inhibit. However, the next Linac pulse occurs regardless of the losses on the previous cycle.

Some of the loss monitors can inhibit subsequent Linac pulses when their integrated loss reaches a set level. The loss monitors located in BTA, as well as eight of the 'movable' loss monitors, have this feature (called 'latching'). To enable the Linac beam, after a radiation inhibit, a manual reset must be done in MCR (Booster Rad Inhibit button). This feature is generally used as part of equipment protection. In particular, losses near the AGS A5

injection kicker, which may be damaged by excessive radiation, are monitored by movable Booster loss monitors, and can often exceed the inhibit level.

The analog signals show the time dependence of the integrated losses within a window. With these signals it is sometimes possible to distinguish between loss mechanisms which one would not be able to distinguish between using the digitized information. For example, a loss associated with the extraction kicker (F3) might be distinguished from a loss occurring due to the extraction bump. These losses occur within the same window and at the same place in the ring. But the kicker loss is faster and might be seen as an abrupt change in the loss *rate* on an analog signal.

3.5 Wall Current Monitor

A wall current monitor measures the current that flows lengthwise along the surface of the beampipe.⁵ Since the beampipe is a conductor, a charge is induced on it that is opposite in sign and the same magnitude as the beam charge. As the beam circulates this 'image' charge follows it as a current on the 'wall' of the beampipe. This image current is parallel to the beam's primary axis of motion, or longitudinal. The beampipe is interrupted by a resistor network. The image current passes through this resistor network. The resistor network is connected to 50 Ω coaxial cable. If the beam current is constant, the charge distribution inside the beampipe is constant, and so there is only an image charge on the beampipe. There is no image current there. Consequently, the 'wall monitor' does not 'see' the DC component of the beam current. The beampipe can still be treated as a 'perfect' conductor for much of the beam's frequency spectrum. The wall monitor's bandwidth, neglecting cable capacitance, is from about 1 MHz to about 1 GHz.⁶

There are two wall current monitors in the Booster. One is at D6, the other is at E3. The D6 wall current monitor is used for beam phase information in the Rf system, and a signal from it is not available in MCR. The voltage across the resistor network for E3 is available in MCR as a hardwired 50 Ω analog signal at MCR 4. This signal comes straight from the Booster ring. It is split, and also goes to the Rf console in bldg. 914. The signal

⁵see also: J. Reich, *Booster Current Transformers and Booster Wall Monitors*, Operations Note No. 35, Nov. 13, 1992. for more info. on the wall monitors.

⁶Booster Design Manual, Page 8-4

travels much further to get to MCR and the capacitance of the cable filters out much of the high frequency components of the signal. The waveform in bldg. 914 is more representative of the beam's distribution since it contains more of the signal's frequency components. The wall monitor can be used to measure the AC coupled beam intensity and shape. The amplitude of the voltage signal (from the baseline), $V_{bunch}(t)$, is proportional to the image current (using Ohm's law). That image current is proportional to the AC component of the beam current. The resistor network has an effective resistance of $R=3.125\Omega$. Therefore, the area under the curve for the image of a beam bunch is proportional to its charge. The bunch intensity (in charges) is,

$$N_{bunch} = \left[\int \frac{V_{bunch}(t)}{R} dt \right] (6.24 \times 10^{18} \text{ ch/Coul})$$

Due to the attenuation of the signal in MCR, from the cable capacitance, this measurement is not very reliable. However, a more reliable measurement can be made in Bldg. 914. This measurement can be used to verify the calibration of the current transformers. The longitudinal bunch shape, charge density, and other important beam characteristics are measured with this tool. This Booster wall monitor signal is also useful for setting up and checking the extraction kicker timing.

Chapter 4

Controlling and Monitoring through the Controls System

*Spreadsheet*¹ is an application for controlling and monitoring accelerator devices and timing, through the computer, as opposed to locally, at the device. It is mainly used to configure and monitor devices which have 'simple states'. Sending a constant 'reference' current to a power supply which feeds coils that are wrapped around a magnet is a typical use for Spreadsheet. By contrast, a power supply with a time varying current, $I(t)$, might be monitored via an analog signal from that supply, which could be viewed on a scope in MCR using the application Xbar. The reference function for $I(t)$ would *not* be sent through spreadsheet, but through an application that is designed for sending such relatively complex data types.

Spreadsheet is one of the primary tools used in MCR to *troubleshoot* equipment failures in the Booster (or anywhere in the AGS complex). Xbar, a signal routing application available MCR is another one of those tools.

I'll use a specific example to illustrate spreadsheet's usefulness. The first dipole magnet in the Booster to AGS transfer line (BTA), called DH1 (Dipole Horizontal One) is oriented to bend the beam horizontally. There are coils around it connected to a power supply. The magnetic field between the poles of this dipole magnet is proportional to the current that passes through these

¹The Program Overview help document, by Ted D'Ottavio, contains detailed information on the operation of Spreadsheet. It is the main resource for questions regarding it. It is available from the menubar in the application. Most other applications have similar program overviews.

coils. For a given particle momentum, charge state, and magnet, the angle through which the beam is deflected from passing through it is proportional to this field.

Spreadsheet contains lists of devices, these lists are arranged in a 'tree' structure. The 'controls' for the supply that feeds current through the DH1 coils can be found by starting at the 'top level' in the tree structure. This level has several branches (Figure 4.1), each branch contains devices that are relevant to the branch within which they are contained. The branch 'Booster' contains 'sub-branches', one of these is called 'Extraction'. Within 'Extraction' is contained a 'sub-branch' called 'BTA', which stands for 'Booster to AGS transfer line'. This branch does *not* contain sub-branches. This branch contains 'devices', clicking on this 'branch' opens up a page with a list of these devices. These devices share one thing in common: whoever put these devices on this page thought that they were relevant to the operation of the BTA line.

The current output of the DH1 power supply can be set by changing the setpoint value (Setpt column) for the device ABI.DH1.SPRB. A current measurement, or 'readback' can be monitored in the column marked 'Measmnt'. The power supply can be shut off by setting ABI.DH1.STAT to 'OFF' in the 'C1' column. The state of the power supply (On or Off) can be monitored in the 'C2' column. If the Horizontal position on a multiwire in BTA downstream of this magnet was not centered, it could be centered by adjusting the current in this magnet through spreadsheet.

If there is a problem, that you think might be explained by a malfunctioning device in BTA, then this Spreadsheet page is a place to look for an indication of that problem. If you noticed that there was beam in the Booster, but none in the AGS, then you might suspect a problem in the transfer line between the two. If DH1, or any other *DC* magnet *tripped off* you could see that on spreadsheet. It would still have a command of 'ON', but the status of the power supply, in the C2 column, would read back 'OFF'. With this power supply off, the beam would not reach the AGS. Finding it off, you could try to turn it back on *through* spreadsheet. It could also still be 'ON', but not *at setpoint*. You could check its readback to determine this.

A *Controls system device* is distinguished from an actual physical device. For example, the power supply for DH1 is one physical device. There could be any number of controls devices *associated* with that physical device. In this specific case, there are two controls devices associated with the physical

Figure 4.1: Spreadsheet tree structure

device. One device turns the supply off and on (ABI.DH1.STAT), and the other device (ABI.DH1.SPRB) allows one to control its setpoint. These *Controls system devices* are called *Simple Logical Devices*, or SLDs for short.

SLDs can be power supply references, timing events, On/off statuses, vacuum gauge measurements, positions of stripping foils, drive positions for beam flags, etc.

Within the Booster branch of the spreadsheet tree, control and monitoring of practically all SLDs relevant to Booster operation is available. The branches and lists of devices are subject to change, and are not always as logical and comprehensive as they could be. *Some* devices have descriptions. These descriptions can be revealed by putting the cursor on the SLD name and pressing the right mouse button.

The most common branches used for operations are Injection, Extraction, ringTransverse, and ringLongitudinal. More examples of spreadsheet's use will be given throughout the book.

Roughly speaking, information flow, between devices and high level applications such as spreadsheet, is accomplished within a branching computer/electronic hardware network. The branches come together at *nodes*. These *nodes* are generally some type of electronic hardware or computer software. The *nodes* are places where information is gathered to send to, or receive from, places 'downstream' of it, or closer to the physical device. In general, the first 'node' in the path of connection between a high level application, such as Spreadsheet, and a physical device, is called a 'controls station'.

The controls system for the accelerator complex involves different sections (loosely associated with particular accelerators) wherein different controls hardware is used to construct the *nodes* within that section of the controls system. For example, Booster devices are controlled through stations that are Apollo (or Hewlett-Packard) computer nodes. In other parts of the complex, they can be VME computers (RHIC, FEB, SEB). VMEs are also referred to as Front End Computers (FECs) and are a more modern addition to the controls system. Intel controls stations are used in the older parts of the system, such as parts for the AGS, and TTB/Tandem controls. Each type of 'controls station' has its own peculiarities. The particular type of controls station used for a set of devices is mainly a function of when those devices were initially incorporated into the controls system. The AGS complex has been around nearly forty years, during that time the control system has had

to adapt to the changing technologies and demands. Its present state reflects this history.

Each controls station regulates the data flow between many devices 'downstream' of it, and the rest of the controls system upstream of it, including spreadsheet. It is typically located 'physically' in the same vicinity as the devices for which it has responsibility. Each station is associated with one or more device controllers, which are one step closer to the device. The station sends data to these device controllers, and polls them at regular intervals for information about the devices. In turn, the device controllers send to and receive information from the physical devices. They are the last step in the controls system. These controllers are linked to devices, through interfaces, often a programmable logic controller (PLC). The input to the interface (PLC) is the handover point between responsibilities. Sometimes, a communication problem lies here, not in the controls system per se. In a situation like this deciphering which group has the responsibility to fix the problem is the hardest part of the problem from the standpoint of operations. Generally, the best solution is to develop a line of communication between the controls group people, and the people who have responsibility for the device, so that together they can sort it out.

Spreadsheet is often where 'controls problems' first become apparent. If data from devices does not 'fill in', or 'update' on a spreadsheet page, then there is some interruption in the path between the device itself and the spreadsheet page. Such interruptions prevent one from controlling the effected devices. They also prevent one from monitoring these devices. If there is a problem in the machine, due to an equipment failure of some kind, it becomes nearly impossible to diagnose it in some circumstances. That is why it is a good idea to address controls problems as they arise. Say, for example, the machine was running fine, then all of a sudden, all the beam was lost at injection. You go to the Injection branch in Spreadsheet, and find that a large number of devices are not communicating properly. At that point, how can you diagnose the problem. It might be as simple as turning a power supply that has tripped off, back on. With communications working you could have restored the machine within a minute. But, without those communications it might take hours.

The condition of controls system elements (stations, controllers) can often be determined through the software applications in the MCR. There are certain symptoms that spreadsheet displays when the communications path

to the device fails. Sometimes, if communications problems exist, a 'C' indication will appear to the right of the device name. Normally, when the communications path is functioning properly a # character appears in this column, and changes color when the information is updated.

For a specific device, the communications path through the controls system elements can be found through spreadsheet. With the cursor on the device name, click the right mouse button. A window appears which contains two subwindows. The top window may contain a device description. The bottom window, called 'Device Details', contains control system information specific to the device. While most of this information is irrelevant from an operations perspective, there are a couple things that can be of interest. Near the top of all this information the controller and station names are written. In the case of ABI.DH1.SPRB the controller is CDC.BXT.914. The Apollo Station is called CST.914.2. The controller and/or station names typically contain information about their location. Here both elements indicate that they are in building 914. See Figure 4.2.

The application CST.914.2, called an Apollo station runs on an Apollo node. It can be checked through the Startup application. If it is not running, or not responding, then communications will be interrupted at this point. Often, if there is a problem with the station, it can be corrected here by stopping and restarting it. Sometimes, it is necessary to contact controls personnel if the communications cannot be reestablished. If the station does not indicate a problem, it may be the device controller. In this case the device controller is CDC.BXT.914. This controller is the final step in the communication with the device. Here the device is the reference for the DH1 magnet. An application called 'Configure' can be used to check this controller. It, like most other applications, is started through Startup.

Configure is an application used to check, and 'configure' the controls system hardware. The communications path can be checked through Configure by asking for reports from the different elements in the chain. Typically, reports will be received from elements upstream of the malfunctioning element. For example, if the controller is malfunctioning, then a report will be received from its station, but not from the controller itself, or a device downstream of it, when a request is made for reports from elements along a communications path.

How does one request such a reports? An SLD is selected, then from a pulldown menu named 'Reports', the item 'Test SLD communications path'

ABI.DH1.STAT	
Device Description (editable) <div> Status control for DH1 which is the first horizontal main steering magnet in BTA. DH1 is used for Proton/HI security interlock. Control status includes STAndby, OFF and ON. See SPRB device description. </div>	
Device Details (not editable) <div> Device: ABI.DH1.STAT Class: INJ_EXT.PSSTATS Controller: CDC.BXT.914 Apollo Station: CST.914_2 Tolerance: 0 tolerance offset: 0 Logical device number: 26 Controller block offset: 679 Amux input channel: 0 Max ppm users: 1 All user control: yes Legal commands: Command file name: INJ_EXT.PSSTATS.CF Number of used fields: 1 number of possible commands: 3 Command code: 6f command name: ON command_field: 1 associated com Command code: 4b command name: OFF command_field: 1 associated co Command code: 53 command name: STA command_field: 1 associated cor </div>	
Save	Revert
Close	Help

Figure 4.2: Spreadsheet description 'pop-up' window for ABI.DH1.SPRB

is also selected. This requests reports from the controls system elements within the communications path for this SLD. If no report is received from the controller, but a report is received from the upstream element (the station), then the controller is probably malfunctioning. The fix for this is typically a reset of the controller. This is accomplished manually, by going out to the controller, for DH1 in building 914, and pressing the reset button on it.

Chapter 5

The Booster's Transverse Space

In this section, I'll first describe, in general terms, the similarities and differences between a simple harmonic oscillator, and the strong focussing (or alternating gradient) aspect of the Booster. I'll need to define some terms, and concepts, that are commonly used to describe the motion in such a system. These descriptions will be made by way of analogies to a harmonic oscillator. They are only intended to give a '*Rookie*' the *flavor* of the theory involved. The real theory is described well in chapter 3 of *An Introduction to the Physics of High Energy Accelerators*, by D.A. Edwards and M.J. Syphers, John Wiley and Sons, 1993"

As the terminology and concepts are described, I'll try to show how they relate to the Booster. After that I'll describe the application of these concepts to the operation of the Booster. '*Tune control*', '*Stopband Correction*', and '*Orbit Correction*' are described. One of my intentions here is to show some of the motivation behind the design of the machine. Hopefully, this will help in gaining some initial perspective on the machine's operation.

5.1 The Equilibrium Orbit and Betatron Oscillations

The strong focussing aspect of the Booster pertains to the confinement of the beam in the plane *transverse* to its primary direction of motion. The Booster is designed to do work on the beam to increase its energy. This work is done to the beam by the electric fields in the Rf cavities. The force on

particles in an electric field is parallel to the electric field. Hence, a particle's momentum is increased in the direction parallel to the electric field. Ideally, there is no work done by the electric field (Rf) in the plane transverse to the electric field.

The increase in energy per revolution around the machine is small compared to the desired final energy of the particles in the beam. The beam needs to pass through the Rf cavities many thousands of times. Unavoidably, the particles have some motion in the transverse plane. This motion has to be contained within some reasonable range so that the beam will continue to pass through the Rf cavities.

Magnetic fields act at right angles to a charged particle's motion, and they do no work to change its energy. In the Booster, they are used to bound the motion in the transverse plane, so that the beam can continue to be *accelerated* by the Rf system.

For a specific charge state and momentum particle, known collectively as a particle's rigidity ($\frac{p}{q}$), there *may* exist a *closed orbit* inside the Booster. That is, if a particle finds itself on this orbit within the Booster, it will stay on this orbit, provided that the magnetic field and particle momentum remain unchanged. If a particle is displaced transversely from this orbit by some small amount, and the motion about it is bounded, this closed orbit would play an analogous role to the stable fixed point in a harmonic oscillator system. In a strong focussing system it is called an *equilibrium orbit*. Particles (of the same rigidity) which find themselves displaced, either in (transverse) momentum or position, from the equilibrium orbit (E.O.), will oscillate about this orbit as they circle the ring. This oscillation is caused by a restoring force which acts transversely to the beam's primary direction of motion. The restoring force is caused by the beam interacting with the field inside quadrupole magnets. For historical reasons, these oscillations are called *Betatron Oscillations*.

To be more precise, Betatron oscillations are in the plane perpendicular to the velocity vector of a particle that lies on the E.O. This plane can be formed by spanning the space formed by two orthogonal vectors that lie in it. One of those vectors can be chosen to lie in a plane parallel to the plane of the Booster. Where the plane of the Booster is *roughly* defined as perpendicular to the main dipole field, and passing through the vertical center of the beampipe. This vector, that lies parallel to the Booster plane, defines the horizontal direction. The vertical direction, is perpendicular to both this vector, and the velocity vector of a particle on the E.O. At a point

Figure 5.1: Local accelerator reference frame. The three orthogonal local right-handed coordinates are x and y (the transverse coordinates), and s (the longitudinal coordinate).

on the E.O., the vectors define a local right-handed coordinate system, with coordinates x (horizontal), y (vertical), and s (longitudinal). Figure 5.1 shows this graphically.¹

The betatron oscillations about the E.O. are generally resolved into horizontal and vertical components. In what follows I'll often just describe the motion in the horizontal dimension, but the same type of motion exists in the vertical dimension. In this ideal description, the motion in the two transverse dimensions is treated as if it were independent. This is often a reasonable approximation.

One thing that distinguishes a betatron oscillation from a harmonic os-

¹See also Edwards and Syphers, pgs. 66-67.

cillation is that the maximum possible amplitude of a betatron oscillation for an initial displacement in transverse momentum and position from the E.O. is a function of where the particle is in the Booster ring. Since it is a *ring*, I'll often describe things that pertain to the Booster in terms of their *azimuthal* location within the ring.

In a strong focussing system, the restoring force around the ring is *not* azimuthally uniform. If the restoring force in a harmonic oscillator began to vary periodically with time, and if the system remained stable, the maximum oscillation amplitude would become a function of time. In the Booster, the variable s takes the place of time. As the particle moves through the Booster, its s coordinate changes. The quadrupole field, which provides the restoring force, is a function of s , or more precisely, azimuthal location in the ring. Hence, the maximum oscillation amplitude of a particle is related to its azimuthal location.

The phase of the oscillation is determined by the particle's trajectory relative to the E.O. at the point where it enters the ring. Like a harmonic oscillator the amplitude of oscillations (at a specific azimuthal location) depend on the initial displacement and momentum of the particle from the E.O. in x - p_x phase space. Here, ' x ' is the distance from the E.O. in the horizontal plane, " p_x " is the component of the momentum in that plane (where $p_s \gg p_x$). p_x , for a given longitudinal momentum p_s , is (obviously) proportional to p_x/p_s . Additionally, p_x/p_s is proportional to dx/ds ,

$$\frac{p_x}{p_s} = \frac{\gamma m \beta_x c}{\gamma m \beta_s c} = \frac{\beta_x}{\beta_s} = \frac{dx}{ds}$$

dx/ds is called x' , and it is the angle (in the small angle approximation) that p_x makes with p_s . The space generally used in considering accelerators is x - x' space. For a given $p \simeq p_s$ this amounts to x - p_x phase space, and properties that normally hold in phase space hold there.

For a harmonic oscillator the area enclosed in phase space by the particle's path is a measure of the total energy contained in the oscillator. The same holds true for betatron oscillations provided the motion is stable. At a particular azimuthal location in the ring, a particle will fall on the outline of an ellipse in x - x' phase space every time it passes that location.

In either the x or y dimension, the phase space in which beam will circulate indefinitely can be represented by an ellipse. Although the orientation and proportions of the ellipse are different at different locations, the area

of the ellipses at different locations are the same. This phase space area is called the Booster acceptance. Somehow, the beam coming from either Linac or Tandem must be put into this acceptance. This process is the process of injection.

This property, that the phase space ellipse area occupied by the beam is independent of its azimuthal location, is exploited in order to make a definition of the beam's size that is independent of its azimuthal location. The ellipse area that contains N percent of the beam is called the N percent emittance.² In the Booster, or AGS, the most commonly used value for N is 95%.

As the beam's momentum is increased in the s -direction, none of this additional momentum is transferred to the x - y plane, since this plane is orthogonal to the s -direction. The average value of x' tends to get smaller as the beam accelerates since $x' = \frac{p_x}{p_s}$, p_x remains constant, and p_s is increasing. Due to the fact that x' is decreasing, the average value of x will decrease as well, since particles at smaller angles will make smaller excursions, and so the oscillations will get smaller in (x, x') space as the momentum increases. Consequently, the area in (x, x') space occupied by the beam shrinks. Additionally, because the acceleration process is very slow compared to the oscillations in phase space, the average values of x and x' both shrink 'together'. So, the shape of the ellipse in (x, x') space at a particular azimuthal location does not change as the beam's momentum increases.³

This so-called *adiabatic damping*, can be thought of as being due to the increase in the longitudinal component of the velocity (β_s), and the increase in γ as the beam accelerates.⁴ The emittance multiplied by $\beta\gamma$ is called the normalized emittance. For the ideal case, the normalized emittance is

²see Edwards and Syphers, pgs. 78-83, for a discussion of emittance. However, this discussion follows their discussion of the transverse equations of motion. It would be confusing if one was not familiar with those equations.

³The 'aspect ratio' does change in (x, p_x) space. A theorem from Thermodynamics, Liouville's theorem, states that the area in (x, p_x) space which encloses 'n' particles is conserved for a system for which energy can not come in or leave (See, F. Reif, *Fundamentals of Statistical and Thermal Physics*, Appendix 13, McGraw-Hill, Inc. 1965). Motion in the transverse dimensions meets this criterion. So, Liouville's theorem applies, and the phase space area is conserved in (x, p_x) coordinates. Since the average value of x decreases, the average value of p_x must increase to keep the area constant.

⁴An increase in γ can also be thought of as an increase in the relativistic mass (γm). As the beam accelerates, it becomes heavier, and therefore harder to deflect.

constant throughout the acceleration cycle.

The shape and orientation of the ellipse for a given azimuthal location is a property of the Booster magnet configuration. The area of the ellipse is determined by the beam's characteristics, including the momentum.

5.2 Stability and the Booster Lattice

The amplitude of an oscillation can be changed if there is a driving force which is correlated to the motion of the oscillator system. In a simple harmonic oscillator that driving force has a frequency. If the frequency is close to the natural frequency of the oscillator a resonance develops. In the accelerator, the correlation which causes a resonance is between the phase of the oscillation, and the azimuthal location. For example, say the driving force is produced when a particle passes through some magnetic field at an azimuthal location in the Booster. If the phase of the particle in phase space is correlated with its azimuthal location, then a resonance condition may develop. This resonance may cause the motion to become unstable.

If a particle returns to the same phase of its oscillation each time it passes one location it makes an integer number of oscillations in one revolution. If it only returns every second time at the same phase, then it makes a half integer number of oscillations per revolution, and so on. If n times the number of oscillations per revolution is an integer, then the possibility of a resonance condition exists. The integer ' n ' is correlated with the type of field error which produces a resonance.

A characteristic of the Booster (or any accelerator of this type) is that the number of revolutions per oscillation is independent of the starting azimuthal location. This number is called the Betatron tune, it is analogous to the natural frequency of a harmonic oscillator. Ideally, the betatron tune, or just 'tune', should be an irrational number to avoid potential resonances. Fields which drive resonances, called *field errors* always exist in a real machine, so potential resonances also always exist. In the simplest case, the horizontal and vertical oscillations are not coupled to each other. As a result, each plane has its own betatron tune.

One way to effect the response of the beam to field errors is to change the tunes of the particles in the beam. In this way the particles can be moved in 'tune space' to avoid the effects of field errors which are responsible for

resonances at specific tune values. For each plane the beam has a distribution of tunes. This distribution can be roughly characterized by a central tune value, and a tune spread, or width.

The restoring force and the tune are closely related. If the restoring force is larger, the wavelength of the betatron oscillations is shorter, and so the tune is higher. This is analogous to a harmonic oscillator where a larger 'spring constant', results in smaller amplitude oscillations, and a higher frequency. This restoring force is produced by a magnetic field which is at right angles to a particle's primary direction of motion. However, this is not the field responsible for keeping the beam circulating within the beampipe. For the Booster, this bending field is directed in the positive y direction. At a given time, and for a given azimuthal position, the bending field is constant in magnitude and direction. Hence the magnitude and direction of the force on a particle is independent of how far 'off center' the particle is. This field is produced by the main dipole magnets. There are 36 main dipoles in the Booster.

The field which produces the restoring force is a function of distance from the center of the beampipe. The field strength is proportional to the distance from the center of the beampipe. For example, B_y is the component of the field that causes a force in the horizontal direction. The field for horizontal oscillations is $B_y(s)=k(s)x$. $k(s)$ is independent of x and y position. The horizontal force on a particle is proportional to its distance from the beampipe center ($x=0$), and is directed towards (or away) from the beampipe center. This is analogous to the motion of a spring (Hooke's law). Here, $k(s)$ takes the place of the spring constant. This field provides a restoring force which keeps the beam oscillating around a central trajectory, the equilibrium orbit.

However, there are only certain field configurations that satisfy the following Maxwell's equations:

$$\nabla \cdot B = 0$$

$$\nabla \times B = 0$$

A configuration which satisfies these equations, and $B_y(s) = k(s)x$, is, ⁵

$$\overrightarrow{B(s)} = k(s)(y\hat{x} + x\hat{y})$$

⁵W.T. Weng and S.R. Mane, *Fundamental of Particle Beam Dynamics and Phase Space*, AGS/AD/91-2, Page 8.

Figure 5.2: Configuration of quadrupole magnet. Notice that a vertically focussing quad is obtained by rotating the horizontally focussing quad by 90 degrees, and vice versa. At $(x,y)=(0,0)$ the fields cancel each other.

This field is produced by a quadrupole magnet, which has four pole tips, see figure 5.2. If this field *focusses* in the x-direction ($k(s) > 0$), it *defocusses* in the y-direction, and vice versa. By alternating the sign of the field gradient $k(s)$ it turns out that it is possible to make stable oscillations in both planes. The beam sees a gradient which alternates periodically as a function of s . This is accomplished by rotating every other quadrupole by 90 degrees. There are 48 evenly spaced quadrupole magnets in the Booster ring.

The quadrupole magnet focusses like a lens. In the horizontal plane, consider a particle passing through a *focussing* quadrupole, moving parallel to the s direction ($\frac{dx}{ds} = 0$). Regardless of its x coordinate as it passes through the quadrupole, it will cross $x=0$ a distance ' l ' downstream of the quadrupole. In other words, the quadrupole has a focal length of ' l '.⁶

For much of the ring the gradient is zero. It is only at regularly spaced locations that it is not zero. Hence, unlike a harmonic oscillator the transverse gradient the beam sees is a function of time, or longitudinal position. Also,

⁶An illustration of how a quadrupole acts like a lens is given on pgs. 60-61 of Edwards and Syphers.

in the regions of the ring where the gradient is not zero, half of the time its sign leads to unstable oscillations. However, there exists a region of gradient strengths and quadrupole spacing for which there are stable oscillations in both planes. When the quadrupole spacing and strengths are within this region, the beam reaches its maximum horizontal size inside horizontally focussing quads, and minimum horizontal size in the vertical focussing quads. The vertical beam size has the opposite behavior, since the quads have the opposite effect in the vertical plane.

The arrangement of quadrupoles used in the Booster is called a FODO lattice. ⁷ FODO stands for Focussing-zero gradient - Defocussing - zero gradient. By convention, a focussing quad focuses in the horizontal and defocusses vertically. The Booster lattice is composed of 24 FODO sections put end to end. 75% of the zero gradient regions contain main dipole magnets. They are arranged in a periodic structure which repeats 6 times in one revolution around the ring. The six 'superperiods' which are identical with respect to the placement of main dipoles and quadrupoles are named A through F. Each superperiod has 8 sections. These sections are called half-cells because each contains half of a FODO cell. That is, either 'FO' or 'DO'. The half-cells are numbered 1-8. So, each half-cell in the ring has a name. For instance, B3, is the third half-cell in the B superperiod. Odd half-cells contain defocussing quads, even half-cells contain focussing quads. Main dipoles are missing from the 3rd and 6th half-cell of each superperiod. The half-cells without main dipoles are often called 'straight sections'. Figure 5.3 shows the layout of one superperiod graphically.

The quadrupole magnets are in series with the main dipoles. The field gradient (quadrupoles) and field strength (dipoles) are proportional to the current that passes through the magnets. As the momentum increases, the bending and focussing fields are increased proportionately to maintain the same bending and focussing.

⁷Lattice related information can be found in chapter 2 of the Booster Design manual

Figure 5.3: The Booster lattice. Above is an overview of the superperiods and where they relate to the transfer lines. Below is a superperiod with half-cells, where only the main magnets are shown.

5.3 The Transverse Equations of Motion

In this section the simple harmonic oscillator and the accelerator system are compared in a somewhat more mathematical way. This comparison is made so that concepts needed to adequately explain what the tools for operating the Booster are, why they exist, and how they are used, can be discussed.

Recall that, for either plane, motion in the accelerator, has similarities to a harmonic oscillator,

$$m \frac{d^2 x}{dt^2} = kx$$

which has the solution,

$$x(t) = A \cos(\phi(t) + \delta) = A \cos(\omega t + \delta)$$

What are some of the differences between simple harmonic motion and transverse motion in an accelerator?

- 1) The independent variable is s , not t .
- 2) The maximum amplitude of oscillations at a given $s=s_c$ depends on s . In other words, the oscillatory motion is *modulated* by an amplitude *function*.
- 3) The phase, called $\psi(s)$ cannot be reduced to " ωs " because $d\psi(s)/ds$ is not constant. It is not constant because the gradient in the ring is a function of s , so the rate of change of phase with distance is a function of s . The change in $\frac{\psi(s)}{2\pi}$, or the *phase advance* divided by 2π , that occurs in one revolution is called the betatron tune, denoted by the symbol " ν " or " Q ". The tune does *not* depend on where the particle's starting point is, because the total gradient a particle encounters in one revolution is independent of that starting point.

In its simplest form, the differential equation for the transverse motion in an accelerator is,

$$\frac{p}{q} \frac{d^2 x}{ds^2} = k(s)x$$

The ratio $\frac{p}{q}$ plays a role analogous to the mass, it is called the *rigidity*. Through the equation of cyclotron motion it can be expressed in terms of the dipole field (B) and the radius of curvature (ρ) of the main dipoles. The equation of cyclotron motion is $p = qB\rho$. So, the rigidity is also equal to $B\rho$.

The solution to the equation for accelerator motion is,

$$x(s) = A(s)\cos(\psi(s) + \delta)$$

The amplitude function, $A(s)$, can be decomposed into two parts. That is, $A(s)=AW(s)$, where A is constant and $W(s)$ is a function that accounts for the modulation due to the varying gradient of the lattice. For simple harmonic motion, a particle moves in an elliptical path in phase space. The area of the ellipse that its path encloses is a measure of the energy contained in its oscillations. This area is related to A , which is determined by the initial conditions (as is δ). For the accelerator, the constant term A is also determined by a particle's initial conditions. In an accelerator πA^2 is called the 'emittance', and is denoted by the symbol ϵ . Its value is related to the energy in the particle's oscillations.⁸

The motion in phase space is more complicated in an accelerator because $W(s)$ and $\psi(s)$ are not constants, and so $p_x(s)$ (or $x'(s)$) obtained from differentiating $x(s)$, is a much more complicated function than dx/dt is for a simple harmonic oscillator.

As mentioned earlier, at a particular azimuthal location, a particle's trajectory falls on an ellipse in (x, x') space. As the particle moves from one azimuthal location to another, its trajectory in (x, x') space is *mapped* from an ellipse associated with the first location onto an ellipse associated with the second location. After one full revolution, its trajectory in (x, x') space will again fall on the initial ellipse. The plot of its trajectory in (x, x') on many turns, and at two different azimuthal locations, is a plot of two ellipses. Each ellipse contains the particle's trajectories for one of the azimuthal locations.

The angular velocity $d\phi/dt$, or ω , has its analog in $d\psi(s)/ds$, whose inverse, instead of the oscillation period, is the 'local wavelength' of the betatron oscillations at the point s . It is called the Beta function, $\beta(s)$. It turns out that $\beta(s)^2 = W(s)$. So, the equation for betatron motion can be written as,

$$x(s) = \sqrt{\frac{\epsilon}{\pi}\beta(s)}\cos(\psi(s) + \delta) = \sqrt{\frac{\epsilon}{\pi}\beta(s)}\cos\left(\int \frac{ds}{\beta(s)} + \delta\right)$$

⁸This is only true for particles at one longitudinal momentum. Recall that the emittance shrinks as the momentum increases. This is just a consequence of how it is defined. The quantity $\beta\gamma\epsilon$, called the *normalized* emittance, is related to the oscillation energy independent of longitudinal momentum.

and $x'(s)$ is obtained from differentiating this.

Since the horizontal beam is largest at horizontally focussing quads, and smallest at vertically focussing quads, it follows that $\beta_x(s)$ is largest at horizontally focussing quads, and smallest at vertically focussing ones. The opposite is true for $\beta_y(s)$. Perhaps, the most important property of the beta function is that the *beam width* in either transverse dimension, is proportional to the square root of $\beta(s)$ at the azimuthal location associated with s .

In general, the average amplitude of betatron oscillations is inversely proportional to the tune. Or, the stronger the focussing (larger k), the more the beam oscillates (larger average $d\psi(s)/ds$), and the smaller the beam (smaller average $\beta(s)$).

For an harmonic oscillator, the location of a particle in phase space (x, p) can be parameterized as a function of t . This parametrization can be expressed in matrix form,

$$(x_2, p_2) = M_{t_2, t_1}(x_1, p_1),$$

where M_{t_2, t_1} is the *transfer* matrix. An analogous matrix formulation,

$$(x_2, x'_2) = M_{s_2, s_1}(x_1, x'_1),$$

describing particle motion in an accelerator, is an important tool. Particularly, in light of the periodic nature of the machine lattice, matrix algebra simplifies the description of particle motion.

The transfer matrix for an harmonic oscillator from t_1 to t_2 maps points on an ellipse to other points on the same ellipse. The transfer matrix for motion in an accelerator from s_1 to s_2 *maps* a point on an ellipse associated with s_1 to a point on an ellipse associated with s_2 . The transfer matrix for one revolution (from s_1 to s_1) *maps* a point on an ellipse to another point on the *same* ellipse.

5.4 Minimizing the Effects of Resonances

5.4.1 Tune Control

There are additional windings around the main quadrupole magnets which are not in series with the main magnet windings. Current is put through these

windings to make relatively fine adjustments to the gradients. The primary purpose of the extra windings is to make small changes to the betatron tunes. Recall that, increasing the gradient, which is like the spring constant in a harmonic oscillator, increases the tune. The tune is changed from the 'bare tune' (set by the main magnet current) in order to move the tune distribution of the beam away from resonances.

The OpticsControl program is used to make small adjustments to the tunes. There are two power supplies, referred to as the horizontal and vertical tune quad power supplies (located in bldg. 930A). They feed current through these extra 'trim' windings. The horizontal tune quad supply feeds current through the horizontally focussing quad trim windings. The vertical supply feeds the trim windings around the quads that are vertically focussing. Using the OpticsControl program the power supplies are configured to feed current into the windings to produce the desired tunes. The desired tunes as a function of real time (from Booster T zero) are set through the program for both the horizontal and vertical planes.

Figure 5.4 is a printout from the OpticsControl program.⁹ This is a typical configuration for proton running. The tunes are specified by assigning values to ν_x and ν_y for each time value specified in the bottom graph. The values of ν_x and ν_y are displayed in the middle graph. Straight lines connect the points in this graph. As many time values as necessary can be added.

The currents required in the two sets of quads are derived from these values. These currents are shown in the top graph. Straight lines connect the points here. However, the current between points is whatever is required to achieve the tunes requested in the middle graph. For example, if ν_x is set to be 4.6 at $t=30$ ms, and 4.7 at $t=40$ ms, a straight line is drawn between the two points. The requested value of ν_x at 35 ms is then 4.65. The currents required for $\nu_x=4.65$ at 35 ms will be sent to the supplies regardless of whether or not they fall on the straight line connecting the currents between those two points on the current graph.

The currents required are not only a function of the requested tunes, but also depend on the main magnet cycle. For instance, the current required to raise the tune 0.1 at injection field is much less than that required to raise it 0.1 at extraction field. The OpticsControl program looks at the Main Mag-

⁹Further information on the OpticsControl program is in 'AGS/AD/Op. Note No. 40, "Booster Tune Control", by Willem van Asselt, February 4, 1993.

Figure 5.4: The BoosterTuneControl 'Working Line' window in the Optic-sControl program.

nct function to derive these currents. The 'bare tunes' are also required to calculate the currents. Figure 5.5 shows the 'bare tunes' as a function of the main dipole field as viewed through the program. Ideally, the bare tunes as functions of dipole field would be flat. However, the quadrupoles start to saturate before the main dipoles. When magnets saturate the relation between current and field (or gradient) is no longer linear. The amount of current required to produce the same change in field or gradient becomes greater and greater. This is normally not a concern since proton and heavy ion rigidities are typically below the point where saturation becomes a problem.

5.4.2 Stopband Correction

The type of field error which is responsible for a resonance is determined by the fractional part of the tune. For example, a quadrupole field error causes a resonance when the change in the phase of oscillation per revolution at the azimuthal location of the field error is near 180 degrees. This can also be expressed in terms of the tune. If $2\nu_x$ is 'near' an integer, then quadrupole field errors may cause a resonance condition. Only the fractional part of $2\nu_x$ determines the proximity of $2\nu_x$ to an integer. Since the tune is independent of azimuthal location, the exact location of the field error is not important. Figure 5.6 illustrates why a quadrupole error causes a resonance when the fractional tune is one half.¹⁰

However, many field errors, are associated with magnetic fringe fields, and surveying errors, which may occur in patterns related to the lattice periodicity. Other errors may be randomly distributed around the ring. The collective effects of both these kinds of field errors can cause resonances. The effect of these resonances on the beam *does* depend on the locations of the field errors. In other words, a resonance condition can exist when there is a correlation between the phases of an oscillation at different azimuthal locations, and the azimuthal configuration of field errors. The field errors, as a function of azimuthal location, can be analyzed into their harmonic components. These field error harmonic components are associated with resonances. For example, if ν_x is near 4.5, then the $2\nu_x=9$ quadrupole resonance is associated with the ninth harmonic component of the quadrupole field er-

¹⁰R. Talman, from a course given on Accelerator Modelling at U. of Texas at Austin, in 1992.

Figure 5.5: The BoosterTuneControl *Bare Optics Functions* or *Bare Tunes* window in the OpticsControl program.

Figure 5.6: The effect of a quadrupole resonance. If a particle's tune is close to a half integer, the kicks that a particle receives as it passes through a quadrupole error field add constructively to increase the amplitude of the oscillations that particle makes about the equilibrium orbit.

rors. The range of tunes around 4.5 for which the beam is unstable is called a Stopband. If the amplitude of this harmonic component is large, then the effect of the resonance will be strong, and the stopband will be 'wide'.

The Booster typically operates with tunes, in both planes, in the range from 4.5 to 5. In practice, the field errors that have the largest effect on the beam are those associated with the lowest number of poles. In particular, the dipole (first order), quadrupole (second order), and sextupole (third order) resonances are most important. There are 'correction' magnets, of each 'order' arranged around the ring. These can be powered to 'null out' the harmonic components of field errors. There are two types of resonances. One is due to dipole field errors. These have the effect of distorting the equilibrium or closed orbit. As the tune approaches an integer ($n=1$), the orbit distortion, due to a dipole error grows. This is called a dipole resonance. It is distinguished from higher order resonances which are associated with 'stopbands'. Stopband resonances effect the transverse amplitude of oscillations, a dipole resonance effects the transverse position of the center of mass of the beam, or the E.O.

The lower order stopbands that are found located in the region of 'tune space' that the Booster was designed to operate in can be conveniently displayed in a 'tune space diagram'. Figure 5.7 shows such a diagram for this region of tune space. Whenever, $m\nu_x + n\nu_y = p$, where m, n , and p are integers, there is a resonance. The sum $|m| + |n|$ is the order of the multipole responsible for the resonance.

What makes a field an error field? A machine with ideal fields would be stable for all tune values. A real physical machine is unstable around certain tune values. For small imperfections, the turn to turn oscillations in the machine will differ only slightly from the ideal, even near a tune value associated with a resonance. However, near a resonance (within a stopband), the effect of the field imperfection on the oscillation 'adds' from turn to turn, until the oscillation amplitude can become so large that the beam is 'pushed' out of the acceptance. In other words, the condition for stability is barely met whenever $m\nu_x + n\nu_y = p$. If there is any deviation in the $|m| + |n|$ order field from ideal the stability condition may not be met. The strength of that error field relates to the width of the unstable region in tune space (the stopband width).

The application StopbandCorrect is used to 'null out' the error field components that cause stopbands. This application powers strings of correc-

Figure 5.7: Booster Tune Space Diagram showing half and third integer stopband resonance lines.

tor magnets so that they approximate harmonic components in order to correct the fields associated with a particular stopband. For example, the $2\nu_x=9$ stopband is excited by a field which has the azimuthal dependence, $C\sin(9\theta+\delta)=A\sin(9\theta)+B\cos(9\theta)$, where θ is the azimuthal angle starting at the beginning of the C1 half-cell. The amplitudes, A and B, of a 9th harmonic quadrupole field can be adjusted as a function of time.

Figure 5.8 is a printout of the half-integer stopband page. This is similar to the OpticsControl display. The bottom graph is used to insert time values. Instead of the tune, the values of the amplitudes for harmonic components are changed. In this case, the display shows the four components involved in the correction of half-integer stopbands ($2\nu_x=9$ and $2\nu_y=9$). For example, 'cos_9x' and 'sin_9x' are the amplitudes, A and B, of the $\cos 9\theta$ and $\sin 9\theta$ horizontally focussing quadrupole field components introduced by the correctors.

The half integer correctors consist of trim windings on the main quadrupole magnets. There are four strings of magnets used to correct these stopbands. The windings within each string are powered in series. However, within a string, the polarity of the windings vary. The power supplies which feed these strings are bipolar. The polarity is determined by the physical connection of the windings to the magnet in the ring. By reversing the sense of the connection the opposite field can be generated with the same current direction entering and leaving the magnet. Two strings of magnets are on the horizontal quads (QH1STR, QH2STR), and two are on the vertical quads (QV1STR, QV2STR).¹¹

Figure 5.9 shows the configuration of the four strings. A positive polarity (+1) corresponds to horizontally focussing regardless of what type of quad the windings are on. Negative polarity (-1) corresponds to horizontally defocussing, and '0' means there is no winding. Figure 5.10 shows how these strings can be powered to approximate a $\sin 9\theta$ horizontally focussing quadrupole field. For comparison, a $\sin 9\theta$ harmonic is also shown. This azimuthal field is generated by putting 1 'amp' into the QH1STR string, 1 'amp' into the QV1STR string, and -1 'amp' into the QV2STR string. Notice that for each string (figure 5.9) there are as many magnets in one polarity

¹¹Much of the information concerning the configuration of stopband correction strings comes from: Charles B. Whalen, Operations Note #36, *Stopband Corrections*, Nov. 30, 1992

Figure 5.8: Half-integer stopband page from StopbandControl program with typical high intensity Proton setup.

as the other. Hence, there is ideally no change in the tune when a string is powered.

The zero degree azimuth of the ring is taken to be the beginning of the C1 half-cell. A $\cos 9\theta$ quadrupole field can be approximated using another linear combination of string currents. By changing the relative amplitude of sine and cosine components the phase of the ninth harmonic is adjusted.

The strength of the error fields can be a function of time in the cycle. For example, a changing main magnet field can induce current along the surfaces of the beam pipe. These currents cause magnetic fields within the beam pipe. The strength of these fields is a function of dB/dt , which varies through the cycle. Hence, the amount of correction necessary to null them out will vary as well. Other field errors in the magnets themselves may be proportionately larger at lower (or higher) main magnet fields. The beam itself can even induce magnetic fields which vary as a function of time in the cycle.

The other stopbands that can be corrected through the StopbandCorrect program are: Linear Coupling, Third Integer Normal, and Third Integer Skew. The third integer resonances are associated with sextupole field errors. The linear coupling resonances are caused by quadrupole field errors.

A linear coupling resonance is produced by a quadrupole error which has components in both the horizontal and vertical planes. This is called a skew quadrupole field. These resonances can lead to either stable or unstable motion. There are two resonances of this type that are relevant. They are the $\nu_x - \nu_y = 0$, and the $\nu_x + \nu_y = 9$ resonances. Both resonances lead to energy transfer between the two phase space planes (horizontal and vertical). Recall that whenever $m\nu_x + n\nu_y = p$ there is a resonance. In the first case, called a difference resonance, $m=1$, and $n=-1$; in the second (a sum resonance), $m=n=1$. With a skew quadrupole field the motion in either plane is no longer independent of the motion in the other. If the tunes in either plane are such that the motion in the other plane is similar, a resonance involving the motion in both planes may develop. A difference resonance does not lead to unstable motion, only coupled motion. On the other hand, the sum resonance, $\nu_x + \nu_y = 9$, can lead to unstable motion.

This type of coupling between the motion in both planes is called linear because the equations of motion are first order in x and y . This is like the motion of coupled harmonic oscillators.¹²

¹²See chapter 5 of Edwards and Syphers.

Figure 5.9: Polarities of individual windings around quadrupoles used for half integer stopbands. Top Graph: QH1STR = dashed line, QH2STR = solid line. Bottom Graph: QV1STR = solid line, QV2STR = dashed line. Horizontal axis is the half-cell counting from A1.

Figure 5.10: Comparison of $\sin 9\theta$ (dashed line) and $QHSTR1+QVSTR1+QVSTR2$. Where there is one amp in each of these strings. This configuration generates something close to a $\sin 9\theta$ horizontal focussing field (beginning at C1). The tune remains almost unaffected since the net focussing in one revolution is almost unchanged.

The linear coupling correctors consist of skew quadrupoles placed next to some of the main quadrupoles. There are four strings of magnets used to correct for linear coupling. The windings within each string are powered in series. The sense of the wiring connections within each string (polarity) is the same for all these magnets. Each string is fed by its own bipolar power supply. The strings are called, QSSTR1, 2, 3, and 4.

The 'knobs' used to change the skew quad field through StopbandCorrect are called cos_0xy, and sin_0xy for the difference resonance; and cos_9xy, and sin_9xy for the sum resonance. For Heavy Ion injection the skew quads are also used to induce coupling between the two planes. The reason for this will be explained later.

The cos_0xy and sin_0xy knobs generate a field which is used to null out the zeroth harmonic component of the skew quadrupole field generated by the main magnets. They are used in correcting the difference resonance. Since $\sin 0\theta$ is identically zero, it is hard to understand the significance of this knob. In practice, a small change to this knob causes a large change in the quadrupole current.

The cos_9xy and sin_9xy knobs are used to eliminate the ninth harmonic component of the skew quadrupole field.

A sextupole field leads to oscillatory motion with a restoring force that is proportional to x^2 . This is different than motion from a quadrupole field, and its 'analog' the harmonic oscillator. The motion it generates is not linear since the equations of motion involve terms higher than first order in x and y . This fact causes a particle's tune to be a function of the amplitude of a particle's betatron oscillations. A sextupole resonance occurs when the 'tune' of a particle is close to a multiple of one third. Figure 5.13 illustrates why this is so. ¹³

5.5 Dispersion and Chromaticity

There is a range of momenta for which particles will circulate in the machine for a particular magnetic bending and focussing field. A particle which has a momentum, such that on average it passes through the center of the quadrupoles, is called an *on-momentum* particle. Ideally, an *on-momentum*

¹³R. Talman, from a course given on Accelerator Modelling at U. of Texas at Austin, in 1992.

Figure 5.11: Third integer resonance

particle, lying on the E.O., *sees* only the field from the dipoles since it passes through the center of quadrupoles where the field is zero.

The quadrupoles act to keep motion about this *equilibrium* orbit stable. When a particle is *off-momentum*, its *orbit* is larger or smaller. It then passes *off* center through the quads, and is bent by them as well as the dipoles. In general, there will still be a closed orbit for an *off-momentum* particle. However, the bending field is now a combination of the bending field of the dipole magnets, and the bending field experienced by a particle on the closed orbit as it passes off center through the quadrupoles.

Additionally, the quadrupoles still act to keep the motion around the *off-momentum* closed orbit stable. Say the position of an *off-momentum* particle in a quad on the E.O. is x_o , then the field that effects horizontal motion for particles of that momentum is $B_y = k(x - x_o) + kx_o$. The first term is the focussing field, the second is the bending field. The quadrupole strength, k , is the same as it is for an *on-momentum* particle. Since the focussing strength in a quadrupole is $\frac{2}{p}k$, a higher momentum particle is generally focussed less. Consequently, the tune of a particle, for a given main magnet field, is dependent on its momentum. Higher momentum particles generally have lower tunes.

For small momentum deviations, the factor that relates a particle's momentum deviation to its tune deviation from that of the *on-momentum* particle, is called the chromaticity. Expressed mathematically,

$$\frac{\Delta\nu}{\nu_o} = \xi \left(\frac{\Delta p}{p_o} \right)$$

Where ξ is the chromaticity, p_o is the central momentum, Δp is the deviation from the central momentum, $\Delta\nu$ is the deviation from the central tune, and ν_o is the tune of the on-momentum particle. In a machine, where no provisions have been made to alter the chromaticity, it is generally negative in both planes.¹⁴ A typical momentum spread for protons is about 1%. For typical chromaticity values of $\xi_x = -1.6$, and $\xi_y = -0.6$, and $\nu_x = \nu_y = 4.80$ this gives tune spreads of $\nu_x = 0.08$ and $\nu_y = 0.03$.

¹⁴see Edwards and Syphers, section 3.4.2 and 3.4.3 for more on chromaticity. Further information on Chromaticity is also given in: "AGS/AD/Op. Note No. 33, "Booster Chromaticity Measurement", by Pepin Carolan, September 1, 1992. This note gives the measured *natural* chromaticities in the Booster as $\xi_x = -1.574$ and $\xi_y = -0.626$.

The vertical E.O. is basically unaffected by small momentum deviations. The Horizontal E.O. has kinks in it at each of the quadrupoles. Since the placement of the quadrupoles is periodic, these kinks in the E.O. occur periodically. Within each superperiod the orbit distortion is almost identical for a given momentum deviation. This effect is called *Momentum Dispersion* and the resulting orbit is described by the *Dispersion function*.¹⁵

Figure 5.14 is a set of Booster orbits obtained on several Booster cycles with the BoosterOrbitDisplay program.¹⁶ They more or less *overlay* because the machine configuration, and the time in the cycle, are the same in all cases. The positions that the bpm's measure is that of the beam's center of charge (or mass). For a particle of a particular momentum, which undergoes betatron oscillations about its E.O., the average center of mass at a bpm describes the position of its E.O. Additionally, it turns out that for a particular central momentum, the amplitude of the orbit distortion at a particular azimuthal position is directly proportional to the momentum deviation.

So, the measurement from the bpm is the position of the E.O. for the average momentum of the beam. For example, if Δx is the deviation from the on-momentum E.O. for a momentum deviation Δp from a particular central momentum, then $\Delta x = C \Delta p$. Where C is the constant of proportionality. Say there are two particles, 1 and 2, with different momenta. Their E.O.s will have different positions at the location of a bpm, described by $\Delta x_1 = C \Delta p_1$ and $\Delta x_2 = C \Delta p_2$. The average position, Δx_{avg} , at that bpm is the position of the E.O. for a particle whose momentum is the average of the momenta of particles 1 and 2,

$$\Delta x_{avg} = \frac{\Delta x_1 + \Delta x_2}{2} = C \left(\frac{\Delta p_1 + \Delta p_2}{2} \right) = C \Delta p_{avg}$$

Consequently, the *orbit* displayed in the program represents the E.O. of a particle of average momentum. In this particular case, the average beam momentum was higher than the central momentum. Hence, the E.O. was towards the outside of the machine. The dots are the bpm measurements. They are connected by straight lines which do *not* represent the beam's position.

¹⁵see Edwards and Syphers, section 3.3 for more on Dispersion and the Dispersion function

¹⁶L. Ahrens, K. Brown, *Booster Protons II, Startup F.Y. '93*, pg. 52, 2 Mar 93

DATA: Current Orbit
 - FILE 2mar93+.7
 DATE: Tue Mar 2 22:05:34 1993
 PPM USER: Booster proteins 1
 GAIN: 10
 DIFFERING GAIN SETTINGS:
 None
 NUMBER OF GROUPS READ: 5
 NUMBER OF CYCLES: 5
 AVERAGING MODE:
 Overlay each group
 Display cycle 2
 NUMBER OF BUNCHES: 100
 BAD HORIZONTAL BPMS:
 NONE
 BAD VERTICAL BPMS:
 NONE

Figure 5.12: Orbits displayed in the BoosterOrbitDisplay program. They show the effects of Dispersion. The beam's average radius is towards the outside indicating a higher than central momentum. The orbit distortion repeats within each superperiod implying that it is largely a result of the beam momentum being higher than the central value.

Notice that within each superperiod the *orbit* is very similar. For example, at the 4 and 6 half cells the beam is further to the outside than at the 2 and 8 half cells. This is primarily due to the beam passing through the quads off center. It follows the Booster's Dispersion function.

5.6 Chromaticity Control

It turns out that the chromaticity can be controlled with the field produced by sextupole magnets. The field produced by a sextupole magnet is described by,

$$\vec{B} = \frac{1}{2}S \left[(2xy)\hat{x} + (x^2 - y^2)\hat{y} \right]$$

Where S is the sextupole strength. Consequently, the horizontal component of the gradient of the vertical field, $B_y = \frac{1}{2}S(x^2 - y^2)$, is $\frac{\partial B_y}{\partial x} = Sx$. This horizontal component of the gradient of B_y controls the horizontal focussing. Putting sextupole magnets in the ring introduces this gradient which is transverse position dependent. In a quadrupole that gradient is a constant k , in a sextupole it is Sx .

With sextupoles in the ring a radial shift in a particle's E.O. due to a change in momentum, will cause the focussing field that particle encounters to shift as well. The sextupoles can be adjusted to change the radial dependence of the focussing field. Consequently, the relationship between the deviation of a particle's momentum from that of an on-momentum particle and its tune can be adjusted. In other words, sextupoles can be used to adjust the chromaticity.

Similarly, a sextupole also has a horizontal field, $B_x = Sxy$. Its gradient controls vertical focussing, and it has the same radial dependence as the horizontal, $\frac{\partial B_x}{\partial y} = Sx$. As a result, the vertical tune dependence of the momentum can also be adjusted. It is easy to forget that although the vertical orbit does not really change with momentum (i.e.- there is no vertical *dispersion*), the vertical tune does. It is a function of the gradient in the vertical direction $\frac{\partial B_x}{\partial y}$. If this gradient is constant, the vertical tune will drop as the momentum is increased, just as the horizontal tune will. The sextupoles can adjust this dependence.

Sextupole magnets, which as the name implies have six poles, can have either of two polarities. In the Booster, the chromaticity sextupole mag-

nets are arranged in two strings that are powered separately. Sextupoles are placed next to each main quadrupole. One string connects all sextupoles next to horizontally focussing quads, one string connects all sextupoles next to vertically focussing quads. A linear combination of the sextupole currents in both strings allows the chromaticities in both planes to be adjusted independently.

The controls for the power supplies are SLDs on spreadsheet named BMM.CHROM.H_PS (at horizontally focussing quads) and BMM.CHROM.V_PS (at vertically focussing quads). The chromaticities are controlled through the OpticsControl program by adjusting the currents in each of these supplies. The Booster/TuneControl and Booster/ChromControl processes are contained within the OpticsControl program. They operate in much the same way. However, the supplies that power these sextupoles are not bipolar like the Tune quad supplies. This means that the range of possible chromaticities is somewhat limited. In general, it is not necessary to know the required currents in each string to get the requested chromaticity. This is taken care of by the program.¹⁷

5.7 Dipole Errors and Orbit Distortions

In an *'ideal'* machine the E.O. of an on-momentum particle, in both planes, is identical to its intended orbit. This intended orbit is often called a *reference orbit*.

A particle is bent when it passes through a dipole field. The deflection is proportional to the strength and length of the field. It is inversely proportional to the particle's rigidity. If the magnet is short, but the field is strong enough to deflect the particle, its transverse angle will change, but its position at the error will not change. It is simpler to consider dipole errors like this, where the only thing that changes when a particle passes through it is the particle's angle.

¹⁷Increasing the current in the sextupoles near the horizontally focussing quads increases ξ_x and decreases ξ_y . Increasing the current in the sextupoles near vertically focussing quads decreases ξ_x and increases ξ_y . The *chrom_H* magnets have a greater effect on the horizontal chromaticity than on the vertical. The *chrom_V* magnets have a greater effect on the vertical chromaticity. The *chrom_H* magnets have the opposite polarity of the *chrom_V* magnets.

Such an *unintended* dipole field, or *error*, causes the orbit to deviate from the intended orbit by *kicking* a particle on the E.O. as it passes through it. But, is there an E.O. which exists that takes this kick into account?

Consider the case of a lattice with only one dipole error. The path a particle takes around the machine just as it exits the error, until its next entrance is the same path it takes in an ideal machine with the same trajectory at the location where it exits the error. This is because the field that the particle sees is identical outside the error. Outside the error, the distorted E.O. is a path described by betatron oscillations about the reference orbit. The amplitude and phase of those oscillations is determined by the trajectory of a particle on the E.O. at the exit of the dipole error and the ideal machine lattice outside the error.

If the tune had an integer value, the orbit would close on itself at the location of the would-be error without the kick that the error provides. Conversely, any kick from the dipole error, when the tune is an integer, will prevent the orbit from closing on itself. Every time an integer tune particle passes through the error, the kick will add constructively to the kicks on previous turns. Consequently, the amplitude of the distortion will increase on every turn. There is no E.O. for a particle of integer tune. This is a dipole resonance.

From this perspective, it makes sense to suppose that the more '*in phase*' the oscillations of a particle on the E.O. are to the dipole error, the larger the orbit distortions will be for that error. As the tune approaches an integer, the orbit distortion increases indefinitely for a particular amplitude kick.¹⁸

There are many dipole errors scattered throughout a real machine like the Booster. They arise for various reasons. For example, a quadrupole that was not centered properly in the beam pipe will cause a dipole field at the beam pipe center. If its center were offset a distance δ from the beam pipe center, then the magnitude of that dipole field would be $k\delta$.

The details of the distribution of these errors around the ring determine the details of the orbit distortion. As mentioned earlier, the effect of resonances on the beam *does* depend on the location of field errors. Just like the higher order resonances, the degree to which the phase of betatron oscillations is correlated with the positions of the dipole errors determines how strongly those errors effect the beam.

¹⁸see Edwards and Syphers, section 3.4.1 for a mathematical description of this.

The dipole errors can be decomposed into their azimuthal harmonic components. A harmonic component of the total orbit distortion (multiplied by the square root of the β function) is due to the same harmonic component of the field error. Strictly speaking, it is *not* exactly the azimuthal angle dependence of these quantities that is involved. It is the phase advance from some arbitrary azimuthal location divided by the tune. Mathematically,

$$\phi = \frac{\psi(s - s_o)}{\nu}$$

where s_o is the arbitrary starting point. It turns out that, at least on the scale that is of concern here, the phase advance increases almost linearly around the ring. As a result, ϕ is a close approximation to the azimuthal angle from some starting location.

The formula that relates the dipole errors to the resulting orbit distortion is,

$$x(\phi) = \sqrt{\frac{1}{\beta(\phi)}} \sum_{k=0}^{\infty} \frac{\nu^2}{\nu^2 - k^2} [f_k \cos(k\phi) + g_k \sin(k\phi)]$$

Where $x(\phi)$ is the orbit distortion at ϕ , $\beta(\phi)$ is the β function there, and f_k and g_k are the amplitudes of the k th harmonic components of the dipole error field.¹⁹

When $x(s)\sqrt{\beta(s)}$ is viewed around the ring, it is a close approximation to a sine wave if it is generated by an azimuthal sine wave distribution of dipole errors. The phases of the sine waves (error and orbit) are essentially the same. The value on the sine wave divided by the square root of the β function at a particular azimuthal location gives the actual deviation of the orbit from the reference orbit. The orbits that one looks at in the orbit display program are a measurement of this orbit, plus the dispersion effects.

However, since the bpms are all located where the β function is essentially the same (at β_{max} points), the orbit one sees in the program will resemble a sine wave if it is due to a single harmonic. Figure 5.15 shows an orbit that is due to a fifth harmonic dipole error field.²⁰ The close relationship between

¹⁹Adapted from E.D. Courant and H.S. Snyder, "The Theory of the Alternating Gradient Synchrotron", Ann. Phys. 3,1 (1958).

²⁰The orbit distortion shown in figure 5.15 is $\eta(\phi) = 5\cos(5\phi) + 2\sin(5\phi)$, where $\eta(\phi) = x(\phi)\sqrt{\beta(\phi)}$.

Figure 5.13: A model of the vertical orbit produced by a 5ϕ harmonic error field. The solid line connects its values at the locations of the bpms. This is what one would expect to see with the orbit display program. The dashed line indicates a model of the actual orbit.

Orbit harmonics and *dipole harmonics* suggests a method to use to correct dipole errors.

5.8 Orbit Correction

The approach used to correct for dipole errors is similar to that used to correct for higher order resonances (stopbands). They both involve nulling out harmonic components of the error fields. There are 2 sets of dipole magnets distributed around the ring. One set is for the correction of errors that effect the horizontal orbit, the other set is for the vertical orbit. Both sets operate in a similar fashion. The horizontal correction dipoles are located within each even half-cell. Even half cells are where the horizontal β function reaches its maximum. The effect a dipole kick has on the orbit is proportional to the square root of the β function at the location of the error. The vertical correction dipoles are located in the odd half-cells, where the vertical β_{max}

occurs.

The magnets are powered by individual power supplies. The currents in these supplies can be configured to approximate azimuthal harmonic distributions. These *dipole harmonics* can be used to *null out* field errors. In general, the 5θ harmonic components (sine and cosine) are responsible for the largest orbit distortions since the tune is generally closest to the integer 5. In general, orbit distortions can cause beam loss because they can reduce the available aperture. However, this is not always the case. Distorting the orbit around some aperture can reduce beam loss. For example, moving the beam inward in areas where the beam is near some outside aperture can increase the acceptance. The innermost outside aperture is at the extraction point in the Booster. Moving the beam inward, away from this limiting aperture, might increase the acceptance.

The BoosterOrbitDisplay program can do a Fourier transform (FFT) of the bpm data as a function of azimuth around the ring. The most straightforward method of *correcting the orbit* uses the harmonic components of the orbit to determine which error field harmonic components need to be nulled out. In general, the $\cos 5\theta$ and $\sin 5\theta$ components of the orbit are the largest, so those are the most important dipole harmonics to null out. The distortion (only in the horizontal plane) due to dispersion is sometimes also significant. The dispersion distortion typically shows up as a sixth harmonic component in the orbit since there are six superperiods.

After an orbit is taken and analyzed, an appropriate harmonic component can be introduced into the dipole field through the correction dipoles. Another orbit can be taken afterwards to look at its effect. By continuing in this way, the orbit can be corrected.

The BoosterOrbitControl program is used to introduce dipole harmonics. For each harmonic number an amplitude for the cosine and sine components can be specified. This amplitude corresponds to the current, in amps, that goes to the correction dipoles. The ring is taken as beginning at the start of the C superperiod. Figure 5.16 shows the currents that would be sent to the vertical correction dipoles to generate a $5[Amps] \cdot \sin\theta$ field.

This graph is from the BoosterOrbitControl program. The currents can also be displayed in a table. Within the program it is necessary to specify a time in the cycle at which the adjustments to the orbit correctors will be made. This time is called the *control point*. Below the graph the time of the *Control point* is displayed.

CONTROL POINT(in ms) : 20

Figure 5.14: The currents sent to the dipole correctors for a $5Amps \cdot \sin\theta$ dipole harmonic. The x axis is the azimuthal position in the ring. The y axis is the current in the dipoles. This is from the BoosterOrbitControl program. The *Control Point* is the time in the cycle at which the currents in the dipoles have this value.

Each dipole corrector follows its own reference function. The currents in the orbit correctors are *incremented* from their starting values at the control point when an adjustment is made. The currents in the correctors need to *ramp up* to the value specified at the control point, and then *ramp down* to their original value over some period of time. The duration of these ramps are specified at the beginning of each *session* with the orbit correctors. The rate at which the current can rise is limited, so enough time must be given for the current to ramp up (and ramp down).

The supplies begin following their reference functions when a realtime event called BMC.CO_COR.ST.RT occurs. The program does not allow the ramps to start before this event occurs. There is also a readback SLD on spreadsheet that is a realtime event. It is BMC.CO_COR _SMPL.RT. This SLD is set to the control point by the program to get readbacks from the devices. The correctors have status and current readback SLDs on spreadsheet as well. There are also analog signals of the current in each corrector available through Xbar. It is often a good idea to check these signals when adjusting the orbit to confirm that the supplies are doing what you expect. The current range for each power supply is $\pm 25A$.

The dipole correctors can also be used to *make* orbit distortions, not just get rid of them. The dipole correctors were sent a $5\cos 5\theta$ and a $2\sin 5\theta$ harmonic. The orbit distortion one expects from this has the same harmonic components. The *expected* orbit is shown in figure 5.16. Figure 5.17 shows the orbit as it appeared in the OrbitDisplay program.²¹

Notice that the orbit measured by the bpms is remarkably similar to the *expected* orbit. The β function dependence of the orbit does not come into play because all the bpms are at places where the β function has a similar value.

As the tune is moved closer to an integer the amplitude of oscillations in the E.O. with a harmonic number equal to that integer will increase. However, regardless of the tune, the harmonic distortions of the orbit are nulled out by introducing *the same* dipole harmonics into the ring. Consequently, it is easier to null out a harmonic component of the field if the tune is close to the relevant integer because the orbit distortion is easier to see (and measure) in the orbit.

²¹Taken from AGS Studies Report No. 285, M. Tanaka, R. Thern, K. Zeno, *Measurement of Vertical Beam Size at the F6 Extraction Septum using Radiation Losses*, 1993

Figure 5.15: The Vertical Equilibrium orbit distortion due to the presence of a $5\cos 5\theta$ and a $2\sin 5\theta$ dipole harmonic. This orbit was taken just before extraction.

In practice, it is most important to correct the orbit when the beam is largest. The beam is largest near injection. Correcting the vertical orbit for High intensity proton injection is particularly difficult. This is because the vertical tune (in the OpticsControl program) is set very close to 5. Consequently, a relatively small fifth harmonic in the vertical dipole field results in a relatively large orbit distortion. The reason why the tune is set close to 5 has to do with space charge effects. This will be discussed in the chapter on Proton operation.

Chapter 6

The *Rf*

The longitudinal aspect of the Booster, doing work on the beam to increase its energy, is somewhat more removed from us than its transverse aspect, and the approach to it is markedly different. The system responsible for increasing the energy of the beam is commonly called the *Rf system*. However, it is really comprised of two interrelated, but different systems. Just as the elements that control the transverse behavior of the beam are driven by power supplies which follow a reference of some sort, so to are the Rf cavities. Within these cavities work is done on the beam to increase its energy by an electric field. But, whereas the transverse element power supply references are generally *static* functions of time or field that reside somewhere in a computer, or controls hardware, the waveform reference that the Rf *power supplies* follow is not.

A good introduction to Phase Stability and Synchrotron Oscillations is given in Chapter 2 of Edwards and Syphers, particularly pages 30-45. If you're not familiar with these concepts you should read this.

The waveform reference that the Rf *power supplies* follow comes from a servo system that uses as an input information from the beam.¹ The *power supplies* are Rf power amplifiers, which *feed* energy into the Rf cavities. The servo system is generally called the '*Rf Beam Control System*', or the "*Low Level Rf System*". It is distinctly different from the '*High Level Rf System*'. The *High Level system* feeds and regulates the flow of Rf energy into the

¹A description of this system can be found in "J. Michael Brennan, *Rf Beam Control for the AGS Booster*, BNL-52438, September 26, 1994". This report is the primary resource used for the information given here about the system.

cavities in the Booster ring through Rf power amplifiers. The *High Level System* follows a voltage amplitude reference which is a function of time, and a waveform reference from the Beam Control System. The High and Low level (Beam Control) systems are coupled to each other through the beam.

The response of the voltage in an Rf cavity to the voltage amplitude reference is not linear. It depends on the frequency of the Rf, as well as other factors. Therefore, a feedback system is required to achieve the requested voltage. This feedback system, not to be confused with the Beam Control System, is called the Automatic Gain Control (AGC) loop. It is part of the High Level Rf system.

6.1 The Rf System without Beam Control

The Beam Control system as a whole is designed to be a 'slave' to the main magnet field. Its operation makes essentially no assumptions about how the main magnet field is going to change. This is in sharp contrast to the way that transverse parameters are controlled.

The beam momentum can be determined from the magnetic field through the equation of cyclotron motion ($p = qB\rho$). From the momentum and the path length of the beam's orbit (roughly, the circumference), the revolution frequency can be found. This is an approximation since one never knows the magnetic field exactly, and the path length of the beam inside the vacuum chamber is also not exactly known. However, this serves as a first approximation for the beam frequency.

The magnetic field is obtained from the Gauss Clock. This is one of the two uses of the Gauss Clock. The other use, as discussed in chapter 2, is for control system timing. A revolution frequency table is calculated. It is like a column matrix which contains a list of frequencies. The index of each row in the column is the number of gauss clock *ticks* that corresponds to that frequency, the first row corresponding to zero field. A gauss clock tick is nominally 0.1 gauss. For example, the 20000th row would contain the revolution frequency for a beam that would be found circulating in the vacuum chamber when the field in the main magnet is 2000 ticks from zero field.

There are a few adjustments that can be made so that the frequency

table can better represent reality. There is a function (in the RfBeamControl application) which is used to calibrate the gauss clock *ticks* as a function of B. For example, when B is low, one gauss clock 'tick' might be slightly less than 0.1g, say 0.098g. When B is higher, it might be 0.103g. This function is used to derive the frequencies in the frequency table. As the gauss clock ticks occur, a 'pointer' is advanced to the row that corresponds to the next gauss clock tick. The increase in field per tick is controlled by this function. There is also a function that allows one to change the machine's radius as a function of field, $R(B)$. This may be useful since the optimum radius generally varies through the cycle for a variety of reasons.

The frequency table is calculated using these functions. The resulting frequencies are used as a *best guess* for the Rf frequencies as a function of main magnet field. These adjustments together with the resulting revolution frequency table are collectively known as the *frequency program*.

The pointer *register* is initialized to a field value contained in the SLD BRF.LL.B.STRT (*Bstart*) at the beginning of *every* cycle. It does not start advancing until a Gauss timeline trigger called BIJ.RF_TRAK_ST.GT (*Rf_Trak*) occurs. When this trigger occurs, the Rf frequency begins to *track* the Gauss clock, advancing (or receding) with every tick. The starting frequency corresponds to the field value in *Bstart*. Before *Rf_Trak* occurs, the pointer *does not* respond to Gauss Clock ticks.

Curiously enough, the value of *Bstart* is *not* the same as the value of *Rf_Trak*. It is important to realize that the two 'uses' of the gauss clock, as an advance for a pointer in the Rf system, and for control system timing, are distinct. The confusing thing about the triggering of the pointer for the Rf frequency table, is that it uses a trigger (*Rf_Trak*) that according to the control system is the value of the field in the Booster. But this number is not in general the same as the value for *Bstart*, which is a *field value* that the Rf system uses. *Bstart* is the field value in the pointer register, and it is independent of the Gauss timeline. The two uses of the gauss clock are essentially independent.

In an operational sense, *Rf_Trak* is the longitudinal *complement* to *peaker*. Adjusting *Peaker* changes the injection field, without changing the dependence of Rf frequency on field. Adjusting *Rf_Trak* changes the injection frequency, without changing the actual field at which injection occurs. However, changing *Peaker* does *indirectly* change the injection *frequency* as well as the injection field, whereas changing *Rf_Trak* only changes the injection

frequency. This makes moving around in *frequency, field space* confusing. Figure 6.1 shows an idealized manipulation one might make. We typically change Rf_Trak, but leave Bstart unchanged when making changes to the machine. Bstart is reserved for the Rf people.

6.2 The need for an Rf System with Beam Control

The main magnet field follows some function of time. In order for the Rf to accelerate the beam at exactly the correct rate throughout the cycle, many things would have to be known to a high degree of certainty. If dB/dt were not known exactly, then the correct kick to give to the beam per revolution could not be known. If the actual voltage amplitude in the Rf cavities were not known, then you would not know how much of a kick was being given. Similarly, the phase on the Rf voltage wave that the beam sees might be wrong for dB/dt .

The beam can be approximated as a *macroparticle*. Its center of mass, or charge, having some location in phase space. Unable to give exactly the correct kick on every turn, implies that this macroparticle is not the synchronous particle, because its energy is not correct to have it come back every turn at the correct phase. There is also no guarantee that the phase of the macroparticle will be correct. In fact, even if it started out with the right phase, since it has the wrong energy, it will undergo synchrotron oscillations and its phase will change with time.

The synchrotron oscillations that the macroparticle undergo, assuming it is able to stay in the machine, cause its emittance to increase. The beam as a whole oscillates around the stable fixed point where the synchronous particle would be. The approximation that the particle acts *coherently* as one macroparticle is only valid in the short run. For one thing, the synchrotron frequency is a function of the amplitude of the synchrotron oscillations. Particles with a larger amplitude take longer to make one revolution in phase space. Hence, the beam smears out to occupy an even larger area in phase space, and acts 'incoherently' after a time.

These *coherent* synchrotron oscillations are detrimental to the beam quality. It is important that the *macroparticle* keep as close to the *synchronous*

Figure 6.1: This illustrates how one adjusts the frequency and field at injection in the case where there is no feedback from the beam. Then the Rf frequency follows the frequency table. The injection revolution frequency depends on the radius and energy of the beam. The injection field value determines the average beam radius for a given energy. Typically the beam energy (also momentum or rigidity) is fixed by the preinjector (Tandem or Linac). Suppose one wanted to move from point (1) to point (3) in f_{Rf} and B . First, Peaker is changed to get the *right* field (pt.1 to pt.2), then Rf_Trak is changed to get the right frequency (pt.2 to pt.3). Ideally, the Rf frequency at injection might be $h * f_{Rf}$, where h is the Rf harmonic number. In practice, one rarely knows where in (f_{Rf}, B) one wants to go, but it is important to know how to move the beam around in this space.

particle as possible. The feedback system is designed to damp out the *coherent* synchrotron oscillations that the macroparticle might make. This is accomplished through the phase loop part of the Beam Control system.

6.3 The Rf System with Beam Control

6.3.1 The Phase Loop

The phase loop receives information about the phase of the beam from a wall current monitor. It compares the phase of the beam, to a *reference phase*, which is a *calculated* value for the synchronous phase. This *value* for the synchronous phase is determined from measurement of dB/dt , and the voltage on the cavities, as well the Booster's radius (R) and bending radius (ρ). dB/dt is obtained from the voltage induced on a single turn winding on the pole tip of the *reference magnet*. This is the same magnet that is used for the Gauss Clock. It is connected in series with the rest of the main magnets, but resides outside the ring in building 930A.

It turns out that a feedback loop which changes the Rf frequency in response to deviations in the phase of the *macroparticle* from this calculated value for the synchronous phase, can be configured to lead to well damped oscillations around this phase. This is what the phase loop within the Beam Control system does. It damps out coherent synchrotron oscillations that lead to emittance growth, by making small adjustments to the Rf frequency in response to a phase mismatch between the bunch and bucket. This is the most important function of the Beam Control System.

The difference between the phase reference and the measured beam phase gives the *phase error*. An analog signal of the phase error is available in MCR through the mux. It is called BXI.BEAM.PHASE.ERR. It is the oscillations visible on this signal that are damped out by the phase loop. Figure 6.2 shows the effect of turning on the phase loop on the phase error.²

The phase loop turns on well before the beam enters the machine on every cycle. However, it initially looks at the phase of a fixed oscillator, not the beam. When the beam enters the machine, it may, or may not have any phase dependence (it might be *dc*). Heavy Ions are generally injected in a *dc* beam pulse. The proton beam is chopped at the linac so that only beam

²Scope trace taken from J.M. Brennan, pg. 59

Figure 6.2: The effect of the phase loop on the beam phase error. Proton beam is injected into the Booster so that the phase of a *macroparticle*, is not at the synchronous phase of the Rf. It undergoes oscillations in energy and phase. The oscillations in phase between the reference phase and the beam phase are visible on the beam phase error signal (middle trace). When the step in the bottom signal occurs the phase loop starts to feedback on the beam phase (BRF.BEAM.CONT.ST occurs). From this point on the oscillations quickly damp out. This is the effect of the phase loop. The top trace is the wall current monitor. The sweep speed is 0.1 ms/div. The scope trace is taken from J.M. Brennan, page 59.

within a certain range in phase with respect to the Booster Rf phase enters the Booster (more about this later). That is why oscillations in the proton beam are immediately visible on the phase error signal in figure 6.2.

Roughly speaking, there are h macroparticles formed by the h batches of chopped beam from the linac. The harmonic number (h) for protons in the Booster is typically 2 or 3. Typically, once the chopped beam is in the machine, or the unchopped Heavy Ion beam has some bunch structure (induced from the Rf), the phase loop switches from feeding back on a fixed oscillator, to feeding back on the beam phase from the wall current monitor. The trigger that causes the switchover is the SLD, BRF.BEAM_CONT.ST (*Beam_Control*). It is a microsecond resolution delay from Peaker.

6.3.2 The Radial Loop

The primary feedback loop that operations is concerned with is the radial loop. A measurement of the average horizontal position of the beam is made by PUEs. This measurement is compared to a reference waveform that is configured through the '*radial steering function*' in the RfBeamControl application. The difference between these two waveforms is used to alter the phase reference for the phase loop. When it is altered the phase error becomes larger (it is generally very close to zero). The phase loop tries to damp out this deviation of the beam phase from the new phase reference by changing the Rf frequency. As the Rf frequency changes, the radius moves toward the requested radius, and the phase shift introduced by the requested radial shift decreases. When the phase error has returned to zero, the beam now has a different radius.

Operationally, this allows one to scan the beam radius to find an optimum setting (usually measured in terms of beam survival). Presumably, this optimum setting keeps the beam as far as possible from the inside and outside apertures, as well as any *dynamic* apertures that might be caused by resonances and the like. This allows us to make the most use of the available aperture inside the machine. This is probably most important early in the cycle, when the beam is largest. At the end of the cycle the radius becomes important for extraction. However, a technique is often used to match the extracted beam from the Booster to the AGS which involves turning off the radial loop. This process is called synchronization and will be discussed later in some detail.

The measurement of the beam radius is typically made by averaging the horizontal positions at 2 PUEs in the Booster. These PUEs are located about 1/2 of a betatron oscillation away from each other so that distortions in the orbit, which tend to have a wavelength on the order of a betatron oscillation, will cancel when the average of the beam positions at each location are taken.

The PUEs at A2 and A8 are used for this purpose. The radial measurement (relative to the surveyed center of the vacuum chamber and PUEs) is the average position at these two PUEs .

$$x_{avg} = \frac{1}{2}(x_{A2} + x_{A8})$$

Where x_{avg} is interpreted as the average horizontal (or radial) position of the beam with respect to the central radius. This is equivalent to the average *radius* of the Equilibrium orbit. We typically have an analog signal of the average of signals from these two PUEs in MCR (through Xbar, it is called BXI.RFE6.RADIAL_AVG). This signal is what the Radial loop compares to its reference function. Figure 6.3 shows a typical *Radial Steering function* from the RfBeamControl application.

We often have two other signals available, one from the A2 PUE and one from A8 (BXI.RFA6_RADA2 and BXI.RFE6_RADA8). As mentioned, the orbit distortion at A2 tends to be opposite in magnitude to the distortion at A8. The average of the two is a reasonable indicator of the radius. Figure 6.4 shows how the variations in position at A2 are similar to those at A8, but in the opposite direction.³ Here, the radial average is almost flat (not shown in the figure), but the positions at A2 and A8 vary considerably due to distortions in the orbit. These distortions are cancelled out by averaging to give a measurement of the beam radius.

6.3.3 Using the Phase and Radial Loops

The response time of the radial loop is very fast. Shifts in the radial steering function, which take 1-2 ms or more, are tracked well. However, fast radial shifts require the beam to accelerate to the new radius, and should generally be avoided. During a fast shift, the *bucket size* is reduced and some beam loss may occur. The radial loop does not work unless the phase loop is on

³John Donohue, HEP Store '96, section 3.6

Figure 6.3: The Booster radial steering function. This function is used as a radial reference by the *Radial Loop*. Even though the axis says *cm*, the units really are *not* centimeters, they are *Volts*. The calibration (Volts/cm) varies, depending on a number of factors. The function starts at Peaker and has microsecond resolution. The radial loop is turned on when the event BRF.RADIALLOOP.ST occurs. It is a delay from Peaker with microsecond resolution. Once the loop is on, the beam radius follows this function.

Figure 6.4: The A2 and A8 radial loop PUE signals. Orbit distortions during the cycle cause the positions at A2 and A8 to change. A2 and A8 are about $1/2$ a betatron oscillation apart, so the distortion in each is roughly opposite direction, and of the same amplitude. Averaging the two positions gives an estimate of the beam radius. The top trace is the current transformer, the second trace is A2, the third trace is A8. The sweep speed is 10ms/div. This is looking at the entire proton acceleration cycle. Taken from, "John Donohue, HEP Store '96, sec. 3.6"

beam control. The Radial loop is turned on when the microsecond delay from Peaker called BRF.RADIAL_LOOP.ST occurs.

The frequency of the beam before beam control is enabled is determined by the frequency program, and the value of Rf_Trak (and Bstart). Once Beam_Control occurs, the phase of the Rf will shift to keep the bunch phase the same as the synchronous phase, and the Rf will be h times the revolution frequency of the beam. If the beam is not bunched, then the phase information from the beam will be bad. The Rf frequency will make large excursions, and the beam will be lost from the machine. That is why beam control is not turned on as soon as Heavy Ion beam is injected. It must have *some* bunch structure before the phase loop can work. In this case, the Rf induces bunch structure, which the phase loop looks at once beam control enabled. With very low intensity beams, like early on in a Heavy Ion run when injection efficiency may be poor, there may not be enough signal from the wall current monitor to produce phase information that is reliable enough to accelerate the beam.

The gain of the radial loop PUEs was originally set through spreadsheet. However, it is now set by Rf personnel. Occasionally, one of the PUEs used for the radial loop will fail. If this is the case the beam may not accelerate, or accelerate poorly. The radial average signal will often show the beam trailing off in one direction of the other. If both PUE signals are available, they can be checked. As shown in figure 6.4, one PUE signal should look like the other reflected about the horizontal axis. This assumes that the reference is close to flat. It may also be the case that one of the signals is bad in MCR, and the PUE reference to the radial loop is OK. We have run the radial loop with just one PUE in the past. It is more difficult because orbit distortions effect the radius that the beam gets set to. At the very least it is necessary to adjust the reference function to run on one PUE. This is done in part to manually correct for orbit distortions that occur during the cycle. In the building 914 Rf control room there is a switch which allows one to run on the average of A2 and A8, or A2, or A8. This is generally set by Rf personnel.

6.4 Controlling the Rf Voltage

The net voltage wave that the beam sees on one revolution about the ring is the sum of the voltage waves in each cavity. There is a fixed phase difference

between the Rf cavities with respect to the beam. Normally, one might expect this phase difference to be zero so that the two voltage waves add constructively. Then the voltage wave the beam would see on one revolution would have twice the amplitude of the voltage in one cavity. Most of the time this is the case.

However, the Rf Voltage does not regulate as well at lower voltage amplitudes. The early part of the Heavy Ion cycle requires a very low Rf voltage. The reasons for this will be discussed in the section concerning Heavy Ion Rf Capture. One way to get around the lack of regulation is to give both cavities a higher amplitude reference, and change the relative phases of the two cavities. The so-called vector sum of the voltage in the two cavities, or the wave that the beam sees, will be reduced in amplitude, but the voltages on either cavity will be high enough so that they regulate well. This technique is called *Counterphasing*, and it is employed during the Heavy Ion capture process.

The voltage *amplitude* for each of the cavities is derived from the real time Scaler Volts per Turn function in RfBeamControl. Figure 6.5 shows this function. It is started from BRF.LL_VOL_CMD.ST, a real time delay from Booster Prepulse. This is typically set to 40001 μ s. The values on the y axis of the function are the values of the *scaler function*. They are simply the *scaler* sum of the two Rf cavities' voltage amplitude references. The scale may also not be *calibrated*.

The relative phase of the two Heavy Ion Rf cavities, A3 and B3, is set through the Counterphasing function. This function is associated with a special type of trigger called a *glitch bit*, or *Triggered Vector Advance*. This glitch bit is used to synchronize the counterphasing reference to injection.

In general, a glitch bit is a kind of trigger used in the Rf system which suspends the advance of the reference in time. This trigger is associated with a part of the reference function between two points that has a constant value, or level. Any segment of the function connected by two points is called a vector. When the start of the reference function begins, the reference given to the equipment advances at some rate given by whatever clock is controlling it. When this reference advances to the point where it falls on the constant level associated with the glitch bit it stops advancing, remaining at whatever

Figure 6.5: The ScalerVoltsperTurn function. This is the amplitude function sent to the Rf cavities. It has nothing to do with the relative phases of the cavities. This function is the Store archive for the 1997 Heavy Ion run.

level it is at. When the glitch bit trigger occurs that is associated with this vector, the reference starts advancing in time again. This time the advance begins at the next vector, or segment, in the reference function.

Figure 6.6 shows the counterphasing function, together with a glitch bit. What appears in the function is a representation of a trigger which is set through spreadsheet. It is an SLD. In this case, the SLD is called, BRF.PKR_TRG_C_PHS, it is a delay from peaker. This function starts from the real time delay from Booster Prepulse (BPP), called BRF.CNTR_PHS.ST, typically set to 40001 μ s. Both these SLDs can be found in the Aux_rf branch in Booster/ring_Longitudinal.

A value of 5 Volts typically means that the cavities are completely out of phase. In this situation there is no net voltage. At 0 Volts, the cavities are completely in phase. The analog signal BXI.BRF_VECTOR_SUM, available on Xbar, is the voltage wave the beam sees, the sum of the two counterphased waves.⁴

⁴More information on Rf timing is contained in pgs. 13-17C of AGS/AD/Op. Note 38, *Overview of Booster Timing*, by K. Zeno, et al., February 17, 1993

Figure 6.6: The Counterphasing function. For Heavy Ions this is used to change the relative phases of the two Heavy Ion Rf cavities, A3 and B3. The *glitch bit* appears as a small square above the constant level vector. When BRF.PKR_TRG_C_PHS occurs the function advances to the next vector. The glitch bit can be moved around within this program, but will only work if its time coordinate falls on a constant level line segment. Its vertical position has no significance. This function is the store archive for the Heavy Ion run.

Chapter 7

Heavy Ion Operation

7.1 Injection

7.1.1 Delivery and Transport to the Booster

The Heavy Ion beam comes from Tandem. It comes in a pulse, the width of this pulse varies from about 300 to 1000 μs . The current is typically on the order of 10 μA . For Gold +32 Ions (which are most commonly injected) this corresponds to about 10^9 ions. The Ions are delivered from Tandem through the TTB (Tandem to Booster) transport line. There are 29 sections to this line. Tandem delivers the beam to the beginning of the 27th section. We take it from there. There are multiwires (also known as Harps) which can be plunged into the beam, and from which profiles of the beam can be taken in both the horizontal and vertical planes.

The TTB line consists of dipole magnets which bend the beam, and quadrupole magnets which keep it focussed. Both types of magnets have DC current references which are set through spreadsheet. These values can not be different for different PPM users ('they are not PPM'). After section 27 these magnets really serve two purposes. The first is to keep the beam in the beampipe so that it reaches the Booster. The second, is to alter the properties of the beam as it enters the Booster, so that it can be accepted by the Booster configuration.

The dipoles in TTB alter the angle and position of the beam's center of mass in both the horizontal and vertical planes. In other words, they alter the *steering*. The quadrupoles alter the *optics*, or the shape and size of the

Figure 7.1: Multiwire profiles for Gold +32 in the TTB line

beam ellipse in phase space. The particular values of beam position, angle, and optical characteristics required are determined largely by the specific setup of devices in the Booster. Anything changing upstream of section 27 will typically alter these properties, and fine adjustments are typically necessary to keep the machine and beam *matched* at the injection point. Figure 7.1 shows typical multiwire profiles from the '97 Gold run. These are profiles from the last 4 multiwires in the TTB line, obtained with the BeamLineInstrument program.

Just before the beam enters into the Booster it passes through an electric field which is produced by an electrostatic device called *the inflector*. This essentially consists of two charged plates through which the beam passes. The electric field is horizontal, so it bends the beam horizontally (about

15°). The purpose of the inflector is to make the beam come in at as small a horizontal angle as possible to the outermost Booster orbit (more about the meaning of this later). An electrostatic device is used because it is generally easier, and less costly, to bend a relatively low velocity particle by a large amount with an electric, rather than magnetic, field. The inflector setpoint is also not PPM and is set through Spreadsheet (BIJ.SPTMC3.SPRB). The units are Kilovolts.

7.1.2 Multi-turn Injection

Multi-turn Injection is the name given to the type of injection employed for Heavy Ions. It refers to the fact that the injection process takes place over a period of time which is greater than the time it takes for the beam to make one revolution, or turn, around the machine. Therefore, beam must be injected into the machine while beam is already in the machine. The injection process takes this long because the number of particles that can be delivered from Tandem is a function of pulse width and current. The current has some peak value. Therefore, more beam can be delivered if the pulse width is increased. With multi-turn injection the percentage of beam 'stored' in the machine, of that which is delivered to the machine, called the injection efficiency, tends to decrease as the Tandem pulse width increases. However, the total amount of beam delivered tends to increase as the pulse width increases. There is some optimum where the maximum amount of beam can be 'stored' in the machine. For Au^{+32} the typical injection process takes place over about 40-50 turns. The revolution period is 15.1 μs , so 40-50 turns corresponds to a 600-760 μs long pulse. Typical values change from year to year, these were taken from the 1997 run.

The Tandem beam is not *born* inside the Booster acceptance. First, it must be transported to the physical location of the Booster. Then it must somehow be placed inside the Booster acceptance. The Tandem beam is put into the Booster acceptance by changing this acceptance during the injection process. This is accomplished by changing the shape of the E.O. in the injection region of the machine while the Tandem pulse is entering the machine. Dipole magnets in the Booster, capable of changing field rapidly, are employed for this purpose.

First, before any beam is injected, the E.O. is distorted (bumped outward) just *past* the opening in the beampipe, at C3. This is where the Tandem beam

Figure 7.2: A very rough idealization of the initial stage of injection. The initial Equilibrium orbit (thick line) is the E.O., if the inside edge of the inflector were not there. Notice that the TTB trajectory, or how the beam is coming into the machine, intercepts this line on a tangent. The distortion of the *E.O.*, which allows the trajectory to intercept the E.O. on a tangent, is *collapsing*. Due to the collapse of the distortion, or *injection bump*, the particle finds itself inside the Booster aperture, or acceptance, the next time it passes by the injection point.

enters the Booster, which is the inner edge of the inflector. At this point the Booster acceptance is zero, because there is no longer a closed orbit (E.O.) that lies within the machine. The TTB beam is steered so that it is as close to the inside edge of the inflector as possible. It then enters the machine near where the E.O. would be, if the inner aperture of the inflector did not exist. This is shown qualitatively in Figure 7.2.

However, if the orbit distortion is decreased in time, the position where the E.O. would be (if the inside aperture of the inflector wasn't there) moves into the Booster aperture. That is, at some point the E.O. will fall inside the machine, and the Booster acceptance will no longer be zero. If the rate

of decrease of the orbit distortion is such that the 'would be' E.O. is inside the machine before beam injected onto it has had time to come back to the injection point, then that beam will be stored in the machine.

This distortion, referred to as the injection bump, is intentionally decreased in size by reducing the current flow in the dipole magnets which produce the distortion. As it decreases, the acceptance increases. Ideally, particles are injected on the edge of the acceptance as it grows. The faster the acceptance increases (faster the bump decreases) the higher the percentage of injected beam is stored. The larger the acceptance of the Booster, without the distortion, the faster the acceptance can increase over a fixed time interval, and the higher the percentage of injected beam is stored.

Figure 7.3 shows a representation of the injection process in horizontal phase space. In this graph $x=0$ is the horizontal centerline of the Booster beam pipe. The line in the graph represents the position of the E.O. during the injection process. It moves over this range during the injection of the Tandem pulse, about $750 \mu s$ in this case. The four points on the line correspond to the points on the falling edge of the injection bump function. This is a software function used to configure the injection bump. The function, and the magnets that make the injection bump, are described below. For each point an ellipse is drawn (centered on that point) which represents the part of phase space where beam can be stored, or the *time dependent acceptance*. This phase space is drawn for the C3 inflector position in the Booster ring. At the right, the inner edge of the inflector is represented as a straight vertical line. The *time dependent acceptance* lies outside of the inner edge of the inflector because the injection bump is falling. The largest ellipse shows the acceptance after the injection bump has fallen completely. The acceptance is then limited by the position of the inflector. Ideally, the Tandem beam would be about the size of the initial acceptance ellipse. This ellipse lies completely in the inflector aperture when the first part of the Tandem pulse reaches the Booster. Note that the Tandem beam is very small relative to the Booster acceptance. It is because of this that the multiple turns of Tandem beam can be injected successfully.

For simplicity, this explanation neglects the *tune dependence* of the injection process. Consider a time slice of the injection process over which the increase in acceptance due to the falling bump is negligible. Consider the beam that is right at the entrance to the Booster at this *instant*. Follow this time slice of beam around one turn in the Booster. Its center of mass

Figure 7.3: *The Booster Acceptance during Injection.* The graph shows the *time dependent acceptance* ellipse as the Injection bump falls. The diagonal line shows the position of the Equilibrium Orbit during the process. The center of each ellipse corresponds to one of the 4 points on the line. The case where the E.O. is at $(x,x')=(0,0)$ occurs at the end of the injection process. The vertical line indicates the inflector.

will undergo betatron oscillations about the E.O. If the tune is close to an integer, then when it returns to the injection point it will be on the outside of the Booster acceptance. Some of it would scrape off if the bump fell more slowly. If the tune were close to a half integer, the first time around the time slice of beam would be on the inside of the acceptance. If the fall rate of the bump were decreased, this beam would not scrape off. The second time around it *will* be at the outside. At this point, the bump *must* have fallen the same amount as in the integer tune case to keep the time slice of beam in the Booster aperture. So, for the same injection efficiency, the fall rate of the bump could be slower if the tune were near a half integer. For the half integer tune case, the injection process could be lengthened, and more beam could be injected with the same efficiency. So, more total beam could be injected. The efficiency of Multi-turn injection is highly tune dependent.

The injection bump is produced by 4 ferrite dipole magnets. Their power supplies are located in bldg. 930UEB. They are located at C1, C3, C7, and D1. The bump is configured using the BoosterOrbitControl program. In the program this bump is called the Injection 4-order fast bump. This simply refers to the facts that it is used for injection, uses 4 dipoles, and can change quickly. One specifies the angle and position of the horizontal E.O. relative to the unperturbed E.O. at the injection point (C3) as a function of time. Figure 7.4 shows a typical injection bump from the '97 run as shown in the program. The two graphs on the top specify the position and angle. The amplitude and time of each point is specified through the program which connects the points. The current reference for the 4 dipoles is derived from this. The times at which each point occur are the same for x and x'. Those times are specified in the graph in the lower right. To make a bump function, the times are specified, then the amplitudes of both x and x' at those times may be adjusted. The number of points is arbitrary. The units of x and x' are shown on the respective axes. The time axes are in μs . However, the program assumes that the clock used is 1 MHz, actually it is 500 KHz. Consequently, 1 μs on the graph is actually 2 μs at the power supplies. The function begins, as mentioned previously, when the delay BIJ.FAST.TM occurs (this really *is* in microseconds). The graph in the lower left is not relevant right now.

The title of the program window also specifies the ppm user and says 'Other'. This bump can be configured differently for different ppm users.

Figure 7.4: Au^{+32} injection bump as shown in the BoosterOrbitControl program.

In this case the name of the user, which has the displayed configuration is Booster/AGS 1, or 'user 1' for short. 'Other' refers to the particle type. The program needs to know the particle type to calculate the dipole currents required for a particular $x(t)$ and $x'(t)$. What 'Other' means is specified elsewhere in the program. In this case the particle type is Au^{+32} . The currents derived from the x and x' functions for the 4 magnets are shown in figure 7.5. The "POS" above each magnet current refers to the polarity of each magnet. "POS" or positive polarity kicks the beam away from the center of the Booster, or in the positive x direction (i.e.- radially outward). These dipoles (also called injection kickers) can be set to negative polarity as well. Here, the ordinate is time. In this case the unit is $10 \mu\text{s}$, which is different then for the x and x' graphs. In other words, the duration of the actual injection bump (for this configuration) is about $1200 \mu\text{s}$.

Roughly speaking, as beam is injected into the phase space the edge of the acceptance ellipse is filled. The phase space well inside the acceptance ellipse is already filled. The size of the Tandem beam is of the order of the smallest ellipse shown in figure 7.3.

The currents in TTB and the Booster are measured using beam current transformers. Signals from the current transformers are available in MCR through the Xbar program. There are many current transformers in TTB. Comparing the current on different transformers, on the same Tandem pulse, can give information about how much beam is lost as it is transported through TTB. For example, if the peak current as measured at section 27 is $20 \mu\text{A}$ (2.0 V), but the peak current on the section 29 current transformer is $18 \mu\text{A}$ (1.8 V). Then 10% of the beam is being lost between the two transformers. This assumes that the responses of the transformers to beam (as seen on the scope in MCR) are identical.

There are two different current transformers in the Booster. The so-called injection current transformer, and the circulating transformer. The injection transformer responds faster to the beam current. It is generally used to look at the beam around injection time. Significant structure in the Booster current waveform exists on a short time scale ($< 100 \mu\text{s}$) during and just after the injection process. The circulating transformer is typically used later on in the cycle. Figure 7.6 shows a typical scope configuration for looking at the injection process. The sweep speed is 0.5 ms. The scope is triggered from Peaker. The top trace is the injection current transformer. The bottom trace is the TTB section 27 current transformer. The trace above it is the current

Figure 7.5: Currents in the four fast injection bump magnets as shown in BoosterOrbitControl during Au^{+32} running.

30-Dec-96
2:48:17

Figure 7.6: Heavy Ion injection as viewed on a scope. The *top* trace is the injection current transformer, the second from the top is the integral of the TTB current transformer signal (on the bottom), the third trace from the top is the current in an injection bump dipole. The sweep speed is 0.5 ms/div.

in one of the injection bump dipoles.

In this case, the Tandem pulse is about 780 μ s long. The average current is obtained by integrating the current over the entire pulse and dividing by the pulse width. Using the calibration 10 μ A/V an average current of 17.9 μ A is obtained. The other trace in this picture is the integrated current from the TTB current transformer. Notice that the current in the Booster (as seen on the injection transformer) starts to increase as the Tandem pulse arrives. The Booster current continues to increase until the Tandem pulse is over. This is the injection process, sometimes the increase in Booster current is referred to as *stacking*. The Tandem pulse, and the stacking of beam in the Booster, occur while the current in the injection dipole is decreasing. This corresponds to the decreasing injection bump. The incoming trajectory of the Tandem pulse remains fixed as the E.O. moves away from the inflector. Beam

fills the constantly increasing acceptance of the Booster. After the injection process is over some beam continues to circulate within the Booster.

7.1.3 Coupled Multi-turn Injection

Multi-turn injection fills the horizontal acceptance with the beam, but leaves the beam small vertically. Beam can only be successfully stored while the injection bump is falling. The rate at which the injection bump falls is closely related to the injection efficiency. If the bump falls too slowly, the acceptance will not increase fast enough, and the beam will find itself still outside the acceptance when it comes back around again.

Recall that Multi-turn injection is highly tune dependent, because the tune determines when a time slice of beam will come around again on the outside of the acceptance, or aperture. Imagine that the amplitude of the betatron oscillations could be reduced, or *damped*. Then, by the time the beam slice is by the outside aperture again, it is closer to the E.O., and further away from the outside aperture, than it would have been without the damping. This assumes that the beam slice being considered is not the slice injected onto the E.O., which is only an idealization that *may* roughly occur only right at the beginning of injection.

If this damping occurred, then the bump fall rate could be decreased further without decreasing efficiency. The amount of time beam could be injected successfully from Tandem could be increased, and the amount of beam stored could be increased.

Recall that the amplitude of a betatron oscillation is related to the energy contained in those oscillations. If the oscillations were damped, the energy would have to go somewhere. The vertical acceptance is not filled by the Multi-turn injection process. The beam can still be small vertically. If some of the energy in these oscillations could be transferred from the horizontal to the vertical plane, the horizontal oscillation amplitude would decrease. The vertical oscillation amplitude would increase. However, since the vertical acceptance is not filled, there should be *room* for some additional oscillation.

Skew quads couple the betatron motion in the horizontal and vertical planes. A particle with *only* horizontal displacement, passing through a skew quad, is *only* kicked vertically. In this way the horizontal displacement due to horizontal betatron oscillations can result in an increase in the amplitude of vertical oscillations.

Imagine that the horizontal and vertical tunes are the same, there is a skew quad in the Booster, and a particle is injected onto the vertical E.O., but off the horizontal E.O. Passing through the skew quad the first time, the vertical amplitude increases (since a horizontal displacement results in a vertical kick). Since the tunes are the same, there is a fixed relationship between the phases of oscillation in the two planes at the skew quad.

On subsequent turns, the horizontal displacement at the skew quad may, if the tunes are *right*, continue to go into increasing the amplitude of the vertical oscillations. As the vertical oscillations increase, the vertical displacement at the skew quad will increase. The vertical displacement will cause the beam to be kicked horizontally. If the particle's tunes, and the polarity of the skew quad are *correct*, this kick will act to reduce the amplitude of the horizontal oscillation. After a number of revolutions, the tables will turn, and the horizontal amplitude will start to increase again, as the vertical begins to decrease. The oscillation energy is essentially transferred back and forth between planes. This is an example of a $\nu_x - \nu_y = 0$ resonance.

With skew quads on in the Booster to couple the oscillations in both planes, the amplitude of the horizontal oscillations will decrease during the first few turns if things are set up right. This allows the bump to fall more slowly, and more beam to be injected successfully. By the time the horizontal amplitude starts to increase again, the bump has collapsed enough so that the time slice of beam is inside the aperture when it passes on the outside of the acceptance, through the injection area. The *standard* Heavy Ion injection setup uses the skew quads. Figure 7.6 is representative of a *standard* setup.

In the Booster, four strings of skew quadrupoles are used for the skew quadrupole field. They are controlled through the *Linear Coupling* part of the StopbandCorrect Application. Although they were initially included in the Booster design to *correct* for stopbands due to field errors, they are used here to *generate* a stopband in order to couple the motion in the two transverse planes. The closer the horizontal and vertical tunes are, the stronger the resonance condition will be. The horizontal tune is also constrained by the tune dependent nature of the injection process.

7.1.4 Is Coupling Used During Injection?

Now that you've read a highly qualitative oversimplification of Multi-turn coupled injection, its time to confront the real situation. There are *accelerator*

models which *predict* what the horizontal and vertical tunes should be, what the current and configuration of the skew quads should be, etcetera. When this configuration has been *put into the machine*, the injection efficiency has been relatively poor.¹ By *tuning* injection it can be improved. This is not unlike any other situation where a modeled setup is put into the machine. Everything needs tuning. The thing that is different here is that, by the time injection is relatively good, the configuration of the machine has strayed so far from the model's predictions that there is good reason to suspect that coupling is not being used, at least in the way intended, during injection.

Results of recent studies with Iron beam are consistent with this suspicion.² This study used 'chopped' beam so that the motion of the beam could be measured on PUEs on a turn by turn basis. These types of measurements will be described in more detail in Section 8.2.7. What they show is that the beam oscillates more or less as expected when the modeled values are introduced into the machine. And that when the motion is measured with the setup that results from 'tuning', the coupled motion required for 'Coupled Injection' is not there.

The fact is, like many things in the Booster, Heavy Ion injection remains largely a *black box*. However, these measurements have shed some light on it. Similar measurements have not yet been done with Gold beam, which we run more often and typically with somewhat better efficiency. It seems fair to say that the normal running state for injection is not the 'expected' one and that further investigation is required.

7.1.5 Tips for Setting Up Heavy Ion Injection

When trying to establish injection, one first tries to set up the profiles in TTB so that they resemble, as closely as possible, previous running conditions. The setpoints for devices in the Booster are set to their most recent *Store*³ values. The idea is to restore the beam to its previous running condition,

¹C. Gardner, *Fe/Au Startup Book II*, pgs. 99-103, 1997

²L. Ahrens, C. Gardner, AGS Studies Report No. 365, *A Study of the Effect of Linear Coupling on the Injection of Iron Ions into the Booster*, November 19, 1997

³The Store document (or *Store*) is operations attempt to document a machine setup. Once the machine is running relatively well, archives are made, analog signals are printed out, some measurements are also made. This is done for every running *mode*. For example, there are Store documents for HEP (protons) 97, HIP (Au) 96, etc.

by configuring the machine the same way as it was configured when it was running last.

Once this is done, one almost invariably finds that there is no sign of the beam inside the machine. What can one do now? Even though the beam profiles may look the same on the last multiwires as they do in the Store document, that doesn't necessarily mean that the trajectory (and optics) coming into the machine is the same as in that running condition. As for the trajectory, a variation in the inflector voltage of about 1 in a 1000 is significant. The inflector bends the beam just before it enters the machine, well past the last multiwire. Since it is past the last multiwire, any effect in the beam trajectory due to it is not observable on a multiwire. A slight change in the inflector voltage for a given setpoint, would be sufficient to change the angle of the beam's trajectory by a significant amount. Even if the trajectory were the same as it was, the machine may be set up differently, so that it may not accept a beam with the same trajectory.

When trying to set up injection one typically views the injection current transformer, a current transformer from the downstream end of TTB (say, section 29), and one of the injection bump power supply current signals. The scope is either triggered (through Xbar) from peaker, or T.T.BEAM_ON (Tbeam_on). Tbeam_on is a trigger that is typically generated 500 μ s after peaker occurs.⁴ The injection transformer should be in high gain (BLI.PROTON.TYPE set to PPR, BLI.XF.GN set to HGH). If you're lucky, a high gain PUE *may* be available just downstream of the injection point (say at C8 or D2). The PUE responds to changes in current. You might be able to see a signal on it as the leading and trailing edges of the beam pass it.

It is important to realize that the beam does not have to circulate, or *stack*, for it to be visible. Since the transformer is located just downstream of the injection point, it will see beam that is in the machine regardless of whether it is circulating. However, the amplitude of the signal will be very small, since it will only be *one or a few turn(s)* of beam. The first step, when trying to establish injection is to see any sign of beam, anywhere in the Booster. It is very rare to find beam that circulates for more than a few turns at this point.

The gain of the injection transformer is high enough that noise generated

⁴See *Overview of Booster Timing*, pgs. 18-20, for a detailed description of Tandem timing.

by the injection bumps is generally visible on the signal. The transformer is at C6, in the injection area. This noise can be subtracted out on the digital scopes in MCR. Other power supplies, such as the tune quads, have also been known to cause noise on the signal. It has occasionally been necessary, when setting up injection, to average the transformer signal over many Booster cycles to reduce certain types of noise, so that a signal from the beam will be visible.

Figure 7.7 shows the injection transformer averaged over many (20) Booster cycles, together with a TTB transformer.⁵ In this case, a time slice of beam is surviving for only a few turns. Notice that the injection transformer trace resembles the TTB transformer trace except that the current falls about 40 μ s *after* the current on the TTB trace falls. This may indicate that beam is surviving for several turns. The leading edge of the injection transformer also rises slower than the TTB transformer. This also indicates some accumulation of the beam. But the beam is not *stored*, because after the Tandem pulse is over, the injection transformer trace returns to zero. It may be that it takes 3 or 4 turns for the beam to return to the outside aperture (where the beam *scrapes*). If the tune is about 4.75, it would take 4 turns until the beam returns to the injection point after having made an integer number of oscillations since it was injected (4 turns*4.75= 19 complete oscillations). The revolution period is about 10 μ s, which implies about 4 revolutions before the beam scrapes off.

The last few dipoles in TTB (29DV1, 29TDH2, 29TDV3, 29TDH3) and the inflector are used to fine tune the injection trajectory. Ideally, the horizontal trajectory should be tangent to the E.O as the Tandem beam first enters the machine. This means that it should be up against the inside aperture (Booster side) of the TTB beampipe (which is the inside of the inflector at this point) as it joins with the Booster beampipe, and almost parallel to the Booster beampipe (as in Figure 7.2). It is difficult, if not impossible, to determine if this condition is met. When beam hits the inflector, a current between its two voltage plates is induced. A hardwired analog signal of the *inflector current* is often available in MCR. By adjusting the downstream horizontal magnets, and inflector you can force beam to hit the inflector, and current will be visible on the signal. If this is the inside aperture of the inflector, then the position of the beam (not necessarily the angle) may be

⁵From T. Roser, C. Gardner, and L. Ahrens, *Fe/Au Startup Book I*, pg. 18, 1997

Figure 7.7: Iron injection *without* stored beam. Beam survives for only a few turns. The top *square* pulse is the injection transformer, the bottom trace is a TTB current transformer. The scope appears to be triggered from Tbeam_on. It displays the average of twenty Booster cycles at injection time.

close to where you want it. If the beam is *moved back* just until no current is seen, then it might be where you want it.

The setup that exists towards the end of a run, when the *Store* archives are usually made, is typically *optimized*. With an optimized setup, the beam is right up against the inside wall of the inflector, the angle is close to tangent to the E.O.. The tunes are adjusted so that the bump can fall as slowly as possible, and the linear coupling is adjusted so that the bump can fall even slower than that. The bump falls slowly. A less than optimized setup would generally require a bump that falls faster to stack beam. The fact that beam doesn't *stack* when we *load the store values* means something is not optimized. If we don't know what it is, then one way to help the beam to stack is to increase the fall rate of the bump. This is normally what we resort to. In practice, you can do this after you have the beam where you think it should be at the injection point, making the assumption that the beam will not stack with the *optimized* bump.

After the fall rate has been increased, moving the last four magnets, and inflector, setpoints around, while looking at the injection transformer signal is occasionally sufficient to restore some form of stacking. One looks for increased amplitude (or any amplitude) on the transformer. If any signs of beam are evident there, one can then look for beam survival. That is, a beam current signal that remains after the Tandem pulse has stopped. Figure 7.8 (which will be discussed later) shows an example of surviving beam.

Since motion in the two planes is coupled, the vertical injection trajectory has a large effect on the injection efficiency. How the beam intercepts the vertical E.O., will effect the amplitude of horizontal oscillations, either for better or worse. It turns out that the beam usually wants to be well off center, lately about +10 mm on 29MW141 (as it is in figure 7.1). This may be a consequence of the interdependence of the horizontal and vertical trajectories. Alternately, it may imply that the vertical E.O. is not near the beampipe center at the injection point, since one *might* think that the optimum condition would have the beam injected near the E.O. However, the actual reason for this is unknown. Naively, one *might* expect the beam to be centered vertically on this multiwire.

7.1.6 Interpreting *Structure* on the Injection Transformer

Once stored beam is established, the injection efficiency is relatively low. Figure 7.8 represents a typical low efficiency injection setup.⁶ Notice that after the Tandem pulse ends, some beam remains in the Booster. Also, the Tandem pulse is injected during the falling edge of the injection bump. There are some other interesting features here. Once the Tandem beam stops, the Booster intensity begins to fall. As the beam is injected it *stacks*, but not all of it is stored. In fact, in this case most of it does *not* survive.

The slope of the Booster transformer is proportional to the rate at which the beam intensity at the transformer changes. The intensity at the transformer is the sum of the current that is presently being injected, plus the beam that makes it around 1 turn, and the beam that makes it around twice, etc. Or, put another way, it is the injected intensity, plus the *stored* intensity, plus the beam that circulates for 1 or several revolutions, but is not *stored*. When the injected beam goes to zero (the tandem beam pulse ends), then the transformer intensity is just the sum of the stored, and *not* stored beams. The slope on the transformer goes sharply negative after the Tandem pulse stops. This slope is from loss of the *unstored*, but circulating beam. It takes about 80 μs for this beam to stop circulating. This is Au^{+32} beam, which has a revolution period of about 15 μs . So, all the unstored beam is gone after about 5 revolutions.

You can see in figure 7.8 that the injection bump does not start to fall until about 50-100 μs *after* that Tandem pulse arrives. The beam that reaches the Booster does not get stored during this time because the acceptance is not increasing. Some beam may be visible on the transformer at this time, but it is not stored. The beam is stacking during the falling edge of the bump. Notice that during the second half of the stacking the slope decreases. This is probably due to the decreasing current in the Tandem pulse (it is *not* square).

The Booster intensity continues to decrease after the fast fall off right after the Tandem pulse ends, but at a much slower rate. This *slow* loss can usually be reduced, often at the loss of stacking efficiency, by reducing peaker. Reducing peaker causes the injection field to be lowered. The field

⁶C. Gardner, *Fe/Au Startup Book II*, pg. 15, FY 1997

Figure 7.8: Au^{+32} injection with *poor* injection efficiency. *Top trace*: Current in one of the injection dipoles. *Middle trace*: TTB section 29 current transformer. *Bottom trace*: Booster injection transformer. This scope picture displays some typical losses. Possible mechanisms are discussed in the text. An injection efficiency measurement (by C. Gardner) is also shown. The beam current at about 2 ms after the injection process is complete (209 μA) is compared to what it would be if all injected beam were *stored*. This is determined by multiplying the number of *turns* during which the beam is injected, by the average current in the beam pulse. The number of turns is found by dividing the pulse width (555 μs) by the revolution period (15.1 μs). This gives an efficiency of 24.7 percent.

is typically rising at about 1g/ms around injection time. If the beam orbit is close to the inside of the ring the beam will gradually scrape on an inside *limiting* aperture as the field slowly increases. This scraping, which removes beam from the outside of the beam's phase space ellipse may be, at least part of, the loss seen here. This scraping will occur only if the beam has not been *captured* by the Rf. Shortly after injection, the beam has *not* been captured by the Rf. Moving the injected beam outward (by reducing peaker) gives the Rf time to capture the beam before it would otherwise start to scrape. This may be why reducing peaker sometimes reduces this slope.

However, there is conflicting evidence to this hypothesis. Figure 7.9 shows the injection transformer with an optimized injection setup. It shows the transformer with and without the Rf on. The Rf is easily disabled by a button at each of the MCR consoles. This picture clearly shows what is called the *spiraling* beam. That is, the beam spirals inward and is lost on the inside limiting apertures of the machine as the field in the machine increases. It is evident from this picture that the beam begins *scraping* in earnest about 2 ms after injection in the no Rf case. There appears to be a distinct time when this begins to occur. From this time it takes about 1.5 ms for all the beam to be lost. So, apparently this earlier slow loss is not (in this case) attributable to scraping due to the beam spiraling inward.

It can also be noted from figure 7.9, that at least in this case, the Rf is not responsible for losses that occur *before* spiraling begins. The Rf is not responsible for the early slow loss shown here.

The skew quads follow a reference function which is typically set through the StopbandCorrect application. After the injection process they are not needed, but the fall rate of the current in the power supplies is limited. As a result they stay on after injection is over. With the configuration they have from tuning to optimize efficiencies their currents are correlated with a slow loss after injection.⁷ The reason for this is *not* known.

The slow loss just after injection can be lessened by raising the vertical tune as quickly as possible right at the end, and after injection. The vertical tune is generally raised from about 4.82 to 4.90 (as read in TuneControl) in about a millisecond. The tune quads can change their current much faster than the skew quads. Moving the horizontal and vertical tunes apart reduces the amount of coupling. Could the coupling be causing a loss somehow

⁷K. Zeno, *Heavy Ion Startup Book II*, pg. 81, 1996

Figure 7.9: Au^{+32} injection showing overlaid current transformer traces with and without the Rf on. The slow loss, shortly after injection, is still there.

(destructive stopband or some other mechanism)? Figure 7.10 shows the effect of shifting the vertical tune higher and *away* from the horizontal tune.⁸

Clearly, there are many things that are not understood about what is going on at and near Heavy Ion injection. Some have speculated that the slow losses are caused by some type of interaction with the residual gas in the beam pipe. These losses have a roughly exponential form which is what one would expect from a *vacuum related loss*. However, the conclusion of a detailed analysis of *Charge Exchange* loss mechanisms is that they are not sufficient to explain the size of the loss.⁹ This potential loss mechanism is described in more detail in section 7.2.3.

7.1.7 The Tandem Foils

A significant portion of what determines the efficiency of the injection process is *not* under our control. That is, the beam from Tandem. Various factors can

⁸K. Zeno, *Fe/Au Startup Book II*, pgs. 96-98, 1997.

⁹See H.M. Calvani and L.A. Ahrens.

Figure 7.10: The injection transformer with and without a rapid vertical tune shift just after injection. Top traces are the overlaid current transformer traces with and without the tune shift. Middle trace is current in one of the injection bump dipoles. Bottom trace is a TTB current transformer trace of the Au^{+32} beam.

influence how efficiently injection *can be* set up. Perhaps the most notable are the Tandem foils. The Tandem uses foils to strip the beam of electrons. They then select out the desired charge state, and send it down TTB to us. These foils have a finite lifetime, and the beam characteristics generally vary throughout their lifetime.

The lifetime of a foil can vary considerably, but is generally on the order of several hours during normal running conditions. The injection efficiency typically reaches a maximum within half an hour or so after a *foil change*. During normal running conditions half an hour corresponds to about 500 pulses. After that the efficiency remains steady for some time, maybe an hour or two. Then it trails off, usually over several hours. The trail off appears associated with a decreasing beam current, as well as a larger beam size.

The lifetime depends on a number of factors. One obvious factor is their thickness. Thicker foils usually last longer, but their thickness is generally inversely proportional to the injection efficiency. Tandem uses different thickness foils. The thinner foils are more difficult to manufacture, and are often *saved* for when unusually high intensity is needed. They have generally used either 2 or 3 $\mu\text{g}/\text{mm}^2$ foils. The use of the thinner foils is correlated with more efficient injection. There are also *good* and *bad* foils of the same thickness that generally come in batches.

The Tandem has these stripping foils on a wheel of some sort inside the Tandem machine. They generally go through the foils sequentially. Different batches are grouped together. It is not uncommon for the efficiency to change markedly when they switch from one batch to another.

Theoretically, the injection efficiency should improve with decreasing Tandem emittance. It seems reasonable to suppose that thinner foils have the potential to produce a lower Tandem emittance because the beam scatters less as it passes through the foil. This is thought to be why the use of thinner foils is correlated with more efficient injection.

For a given Booster acceptance, the bump must fall faster to store a larger emittance beam. This can be seen from figure 7.3. In order to have efficient injection the acceptance ellipse must enclose the Tandem beam as it enters the machine. If the incoming beam is larger, the acceptance ellipse must be placed further outside the machine to place the Tandem beam inside it. As the beam is injected, the bump must then fall faster to store the beam. The details of the most efficient configuration for the Booster to accept the

Tandem beam depend strongly on the emittance of the Tandem beam.

In practice, the size of this acceptance ellipse, and the bump's fall rate, which are both determined by the Booster setup, remain unchanged, and the Tandem emittance varies. This variation can be due to a number of factors, some of which have already been discussed. Some possible mechanisms which could effect the injection efficiency are: 1) The use of different thickness foils, 2) The age of the foil, or how much beam has passed through the present foil, 3) The amount of beam that is put through the foil per pulse, and 4) The size and duration of the beam that passes through the foil. Undoubtedly, there are other factors as well.

For reasons which are unclear, the efficiency of the beam that makes it into the early part of the cycle, say 5 ms after injection, is inversely related to the intensity of the Tandem beam. However, it is not clear that the proportion of beam initially stored to that injected necessarily has this relationship. This intensity dependent effect is largely related to a slow loss that occurs after the injection process is complete.

Figure 7.11 shows beam sizes on the 29MW090 harp, versus injection efficiency over a period of 8 days during the Au^{+32} run. These measurements were extracted from the *Morning Numbers*.¹⁰ The *Morning Numbers* are measurements that are taken every day to document the evolution of the machine during a *physics run*. The injection efficiency is determined by doing a measurement identical to the one made in figure 7.8. On the sixth day the type of foil being used is changed. Clearly the beam size at the multiwire increases when this change is made. During the period before the change in foil type, the beam size (emittance) remains fairly constant. It also remains fairly constant after the change but is at a different level. All this assumes that the *optics* that transport the beam to this multiwire from Tandem, as well as other properties of the beam haven't changed.

The efficiency increase during the first few days is most likely due to changes made to Booster injection through tuning. The drop in efficiency *after* the change may be due to an increase in Tandem emittance. The relatively high efficiency measured on the sixth day is a mystery.

¹⁰K. Zeno, *Fe/Au Startup Book II*, pg. 105, FY 1997

Figure 7.11: Beam size on 29MW090 and injection efficiency versus time. The type of foil was changed on day 5 (the sixth day). The change in foil *type* is correlated with an increase in beam size in both planes at the multiwire. The lines on the graph only connect the data points.

7.1.8 Optimizing the Injection Efficiency

The optimization of the injection efficiency is one of the most tuning intensive areas involved in operating the Booster (or the AGS). It is a long, and sometimes tedious process. It normally takes over a month of running to achieve something resembling optimized injection. The setup for Gold running has reached a relatively optimized state during several runs. The Iron runs that have taken place have never reached a comparable condition. This is primarily because the Iron runs have not been long enough. They also involve a lot of intermittent running, which makes tuning more difficult.

The highest efficiencies for Au^{+32} have been in the range of 60-70 percent with a 600-700 μs pulse, rather high current (20-30 μA), and a *good* Tandem foil. This corresponds to something in the neighborhood of 2.5×10^8 ions at 2 ms after injection stored in the machine. Figure 7.12 is a scope picture of a highly optimized injection setup at *high* Tandem current.¹¹ About 20-25 percent of the loss in efficiency is due to losses that occur after the injection process. That is, after the beam has been *stored* in the machine. As discussed earlier, the tune (figure 7.10) and skew quads, appear to have an effect on this *slow* loss.

In figure 7.12, the Tandem pulse is quite square, and the slope of the injection transformer is relatively constant during the injection process. This indicates that the injection efficiency does not vary by much through the injection process.

The type of tuning that I find most effective involves changing injection related parameters by small amounts in a loop of some sort. This *loop* is iterated many times. Sometimes it brings about an improvement, many times it does not. The loop hinges around the injection bump. Unfortunately, there are no independent parameters. A given injection trajectory will give the best efficiency for a given bump. A change in the trajectory requires a change in the bump to give the optimum efficiency. That optimum condition may be better or worse than the initial optimum condition. Here we are just dealing with a few variables, actually there are many more: Both tunes, both trajectories (vertical and horizontal), the skew quads, the injection field (peaker), the orbit correctors, the TTB optics, the Tandem beam characteristics, and others.

How does one navigate through this maze? I think that keeping in mind

¹¹K. Zeno, *Fe/Au Startup Book III*, pg. 32, FY 1997

23-Jan-97
19:01:18

Figure 7.12: Relatively high efficiency injection at high Tandem current using a *good* Tandem foil. Top trace is the injection transformer, middle is the integral of the Tandem transformer, third trace is an injection kicker current signal, and the fourth trace is a TTB current transformer.

the relations between different parameters as one changes them is helpful. It seems to me that the most important element is the injection bump. The changing acceptance of the Booster is what allows beam to be stored in the machine. The acceptance and its rate of change are determined primarily by the injection bump. The other end of this is the trajectory. It also has to be reasonable for beam to be stacked. With coupling, the trajectory in the vertical becomes a more integral part of the injection process. Without coupling, the horizontal tune effects how fast the beam comes around again on the outside of the inflector. With coupling, the separation between the tunes effects how fast it comes around. The skew quads also effect this. Peaker effects how large the bump can be, and how much available aperture there is. And so on...

Perhaps a concrete, yet hypothetical (and oversimplified) example is in order. Consider figure 7.8 which shows an inefficient injection setup. What are the major things that could be wrong with it? First one considers the pulse width. The pulse is too wide for the bump. But we want that pulse width, so that we can inject more beam. How can we make the Booster accept that pulse width? Well, the bump isn't wide enough, so widen it. So you widen it, but in the process you necessarily decrease its fall rate. The beam stacks more evenly, but less is stacked. So you increase its amplitude. Now the beam gets stacked better, but the beginning of the pulse doesn't stack anymore. Why doesn't it stack? Maybe the orbit is too far outside the machine at the beginning of the pulse. If you raise Peaker, then the orbit will move to the inside. You do that, now the beginning stacks, but the end is worse. By moving Peaker maybe you moved the orbit so far inward that you reduced the available aperture of the machine. By adjusting the inflector you find you can improve stacking at the end. You did this because you remember that by adjusting the inflector you have been able to improve the stacking towards the end of injection. This is all very confusing, and is purely hypothetical. The point is you shouldn't be afraid to try things, and the things you try might be more successful if they are based, however loosely, on something.

7.2 Rf Capture and Acceleration

The beam from Tandem is delivered to the Booster as an almost constant current (dc) pulse, typically anywhere from 200 to 1000 μs long. During injection, the main magnet field is ramping up at a slow rate, about 1g/ms. It ramps slowly for two main reasons. First, a large dB/dt would reduce the available acceptance for the injection process. Second, Rf manipulations must be done before the beam is able to be accelerated efficiently because the beam is *not bunched*.

The Rf manipulations necessary to efficiently accelerate Heavy Ions are quite different from those used to accelerate protons because proton injection is synchronized to the Rf. From an operations perspective, the principle Rf problems one faces in the early stages of the Heavy Ion cycle are: 1) Capturing the stored beam, 2) Keeping the longitudinal emittance as small as possible, and 3) Accelerating the beam as quickly as possible, without seriously compromising 1) or 2). The main magnet cycle (and therefore the acceleration cycle) are set up with these things in mind.

7.2.1 Capturing the Stored Beam

Just as beam comes to be stored during the injection process, beam becomes able to be accelerated through the *Capture* process. In a sense, *Capture* is the longitudinal analog of injection. While the beam is being injected, the Rf is on at a low voltage, and the beam is being bunched. The capture process *should* not normally interfere with injection. In figure 7.9, the current transformer traces for the cases where the Rf is on and off overlap during injection. For this particular setup, there does not appear to be a difference in injection efficiency between the two cases. The capture process did not noticeably effect the injection efficiency.

One typically wants to *capture* all the stored beam. A dc beam has to be captured at a low $\frac{dB}{dt}$ in order to capture all of it. This is because the Rf is only able to capture beam that finds itself on all the phases of the Rf wave when $\frac{dB}{dt}$ is zero.¹² Since the beam contains no phase information, Beam Control is not turned on right as beam is injected. The Rf acts on the beam once injected, and begins to *bunch* it.

¹²If you're not familiar with longitudinal 'dynamics' you should read Chapter 2 of Edwards and Syphers.

Once the beam has been bunched to some degree, and contains some phase information, Beam Control is typically turned on. During the '97 run, this delay from Peaker, was often set to about 900 μ s. In this case, Beam Control turns on towards the end of the injection process. There is sometimes evidence of bunched beam on a PUE this early in the cycle.¹³ The time when Beam Control can be effectively turned on depends on how soon phase information is available. Higher beam intensity would provide a larger signal, or a higher Rf voltage amplitude would bunch the beam faster. However, the Rf voltage is kept low during the initial stage of capture to keep the longitudinal emittance small. This will be discussed in more detail below.

Since the Rf voltage is low, it is necessary to set the frequency of the Rf very accurately. Due to the low voltage, the Rf buckets will be relatively small, allowing only a small range of particle energies (and hence frequencies) to fall inside the bucket. The energy of the Tandem beam is fixed. One adjusts the Rf frequency to shift the energy (or frequency) of the Tandem beam with respect to the Rf. It is important to line them up initially so that the beam falls inside the bucket. It is also important because beam control is not on, and cannot compensate for an error in the injection frequency.

Typically, when first setting up acceleration, the revolution frequency of the beam is measured. Then the Rf frequency at injection is set to hf_{rev} by the Rf group. During running conditions, this frequency is adjusted with Rf_Trak. The revolution frequency measurement is tricky since the beam is not bunched.

In the past, a kicker, such as the F3 extraction kicker is set to fire just after injection.¹⁴ The F3 kicker amplitude is set to kick the beam hard enough that it does not get around the machine. This is not difficult since this is an extraction kicker, and the beam is at injection energy. The width of the kicker pulse is small compared to the revolution period of the beam since the beam moves so slowly at injection. A *hole* is kicked out of the circulating beam. The beam is no longer dc, and is now visible on a PUE signal. To find the revolution frequency, one counts a number of revolutions of the beam on the PUE signal, and divides this by the amount of time it takes to go that many revolutions. This can be done rather easily on a digital scope, though

¹³C. Gardner, *Fe/Au Startup Book I*, pg. 54, FY 1997

¹⁴K. Zeno, *Fe/Au Startup Book I*, pg. 110, FY 1997

it gets tedious counting many turns. For Au^{+32} , counting 100 turns, f_{rev} was measured to be 66.3559 KHz for the fiscal year (FY) '97 run. f_{Rf} was set to be $530.85 = hf_{rev}$, where $h = 8$.

7.2.2 Acceleration

A large Rf voltage at injection causes particles that are located at phases on the Rf wave to trace out relatively large amplitude oscillations in longitudinal phase space. Take for example a particle that finds itself far from the synchronous phase (which is about 0° on the Rf sine wave at injection). Every time it passes through an Rf cavity it will get a bigger kick if the Rf voltage amplitude is higher. It will do oscillations about the synchronous phase, but the oscillations in energy will be greater for a given phase if the Rf voltage is higher. This will give the beam a larger longitudinal emittance, and this emittance cannot be reduced.

The optimum voltage from this perspective is one that is high enough to form a bucket that is big enough to fit the initial energy spread of the beam, but not higher. The energy spread of the Tandem beam requires only a small voltage, providing the frequency of the Rf is matched to the beam's frequency.

Raising the voltage *rapidly* after the initial stage of capture also causes the emittance of the beam to increase. This is because the system is thrown out of equilibrium by the changes. It is like what happens when the voltage is set too high for the energy spread of the newly stored beam at the initial stage of capture. The emittance of the beam will increase. But if the voltage is increased slowly with respect to the motion of the particles in phase space, the system will remain in equilibrium, and the emittance will not increase. This is the ideal, it is called *adiabatic capture*. Before the field starts to rise too fast, it is important that the transition between an almost non-accelerating, or *stationary* condition, and the accelerating condition, can be done without taking the system out of an equilibrium condition. The voltage is slowly increased during this period, and the beam becomes tighter in phase. The Rf needs some time to accomplish this.

Figure 7.12 is a scope picture illustrating the capture process.¹⁵ The bunched beam is visible on the C6 sum PUE signal. Notice that the signal

¹⁵C. Gardner, *Fe/Au Startup Book I*, pg. 54, FY 1997

Figure 7.13: Rf capture of Au^{+32} . The top trace is the C7 injection kicker. The second from the top is the normalized injection transformer. The third signal is the sum signal from the C6 PUE. The bottom trace is the Rf Vector Sum. The C6 Sum signal shows bunch shape oscillations like those that might occur if the Rf voltage were too high for the energy spread of the incoming beam.

from the bunched beam is small during the injection process even though the current transformer indicates that there is beam circulating in the machine. The PUE signal does not see the unbunched component of the beam. The bottom trace is a high frequency Rf signal which is a superposition of the voltage waves in the two Rf cavities (located at A3 and B3), called the Rf Vector Sum.

The C6 PUE signal shows evidence of *bunch shape oscillations*. The bunch in phase space does not have the same shape as the bucket. As the bunch undergoes synchrotron oscillations it alternately narrows and widens in real space. When the bunch is narrow its peak current is higher. The C6 PUE sum sees the ac, or bunched, component of the beam current. These oscillations are symptomatic of some type of *mismatch* between the beam and Rf. For example, the Rf voltage may be too high, or the frequency of the Rf may not be hf_{rev} . This type of oscillation is a mechanism whereby

the emittance grows. It can cause beam loss.

As discussed in section 6.4, the voltage the beam sees is determined by both the ScalerVoltsperTurn and Counterphasing functions. As an example, I'll look at the configuration for the Fiscal Year (FY) '97 run. This configuration changes from day to day, and year to year, but looking at one case still illustrates some points.

In this case, the glitch bit in the counterphasing function is set to about 500 μ s after peaker, just as the beam starts to be injected. The cavities start out almost completely counterphased, with a reference of about 4.7 V, where 5.0 V is completely counterphased (see figure 6.6). From this point the cavities are brought into phase over a period of a few milliseconds. The ScalerVoltsperTurn function is low, about 10 kV on the function (see figure 6.5). However, the Heavy Ion cavities (A3 and B3) supply only 17 kV each. The ScalerVoltsperTurn function is not calibrated correctly for Heavy Ions. 90 kV on the function is really about 34 kV on the cavities, or the maximum voltage. So, 10 kV on the function is more like 3 or 4 kV. Due to the counterphasing, the voltage the beam sees is much less than this.

Figure 7.14 shows B and $\frac{dB}{dt}$ throughout the Au^{+32} cycle.¹⁶ This picture was taken just after acceleration had been established. Therefore, it is not optimized, and the losses are higher than normal. It shows that $\frac{dB}{dt}$ is low during and just after injection, then about 4-5 ms after injection it starts to rise to its maximum value. The current transformer shows a loss 2-3 ms after injection. This is probably the loss of the uncaptured beam. As the field rises, beam that is not captured, falling outside the buckets, is not accelerated. It spirals inward as the field increases, finally being lost on the inside apertures. Compare this to figure 7.9 where all the beam is lost in this way when the Rf is off. The early acceleration part of the Booster Main Magnet function is shown in figure 7.15.

The loss about 30 ms into the cycle is about where the minimum bucket size in the cycle occurred. If the longitudinal emittance is large, a loss like this might occur because the bucket is not large enough for the beam. Figure 7.16 shows a similar picture about 5 days later in the run.¹⁷ A loss at about the same time is apparent. In this case, the gap volts (ScalerVoltsperTurn) were not set at their maximum value. Raising them higher got rid of this

¹⁶C. Gardner, *Heavy Ion Startup Book I*, pg. 28, 1997

¹⁷G. Marr, *Heavy Ion Startup Book I*, pg. 71, 1997

Figure 7.14: The Au^{+32} cycle showing $\frac{dB}{dt}$ and B analog signals together with the current transformer. Beam is injected at the time indicated by the arrow at the bottom of the graph. 2-3 ms later there is a capture loss visible on the transformer. After about 30 ms there is another loss, possibly due to insufficient bucket size (or equivalently an emittance that is too large). The sweep speed is 10 ms/div. The B and $\frac{dB}{dt}$ analog signals are available through Xbar.

Figure 7.15: The Main Magnet cycle during the FY '97 Au^{+32} run. The top trace shows the magnetic field function during the early part of the cycle. Below it is the $\frac{dB}{dt}$ function for the entire magnet cycle. Compare these to figure 7.14

Figure 7.16: Beam loss due to insufficient bucket size during the acceleration cycle. The top signal is the normalized circulating current transformer (BXI.CIRC.XFMR_NORM). The bottom trace is the A3 cavity gap voltage signal (BXI.RFA3.V_GAP_DC). It shows the *amplitude* of the A3 gap volts. The picture shows the entire acceleration cycle (10ms/div). The beam loss in question occurs about 30 ms after injection.

loss. This is consistent with a bucket size/emittance problem since raising the gap volts increases the bucket size.

There are essentially two ways to adjust the *vector* voltage (the voltage the beam sees) during the first few milliseconds of the cycle. You can use the counterphasing function, or the ScalerVoltsperTurn function. The counterphasing function is more closely synchronized to injection because of the glitch bit that is a delay from peaker. For that reason, I find it easier to adjust the rate of change of the vector voltage during the first few milliseconds using the counterphasing function. After the first few milliseconds the counterphasing function is typically set to zero for the remainder of the cycle. This corresponds to zero counterphasing. After the first few milliseconds, the vector voltage needed is high enough that the counterphasing is no longer needed. Then only the ScalerVoltsperTurn function is used to adjust the voltage.

It is also difficult from the `ScalerVoltsperTurn` function to know when, with respect to the beam, you are changing the voltage. However, you can see that on an individual cavity voltage signal. Usually, when you are adjusting the Scaler voltage it is helpful to have a cavity signal up on the scope (BXI.RFB3.V_GAP_DC or BXI.RFA3.V_GAP_DC), together with a current transformer signal. The circulating transformer is better for looking over the entire cycle though it has been noisy during the early part. The injection transformer is typically used to look at the early part.

Figure 7.17 shows the effect of adjusting the initial counterphasing level. Here the level is changed from less to more counterphasing with positive results.¹⁸ In fact, the loss it improves is apparently a component of this generic *slow* loss that happens after injection. So, now there are three ways to effect this loss, the skew quads, the tune shift, and the vector voltage. The total loss rate during early acceleration may be caused by several unrelated mechanisms. It is interesting to consider the connection between figure 7.13 and 7.17 since a high vector voltage may cause bunch shape oscillations.

Figure 7.18 shows the circulating current transformer throughout the cycle with two slightly different magnet cycles.¹⁹ The cycle with the lower $\frac{dB}{dt}$ has *less* loss early in the cycle. However, since extraction must remain unaffected, the $\frac{dB}{dt}$ is higher late in the cycle (to reach the same extraction field), and as a result, the Rf voltage is not sufficient to accelerate more beam past the bucket size minimum. This later loss is *not* at issue here. It could be removed by lowering the $\frac{dB}{dt}$ later in the cycle, and allowing extraction to occur later.

The key point about figure 7.18 is the reduction in the *capture loss* for the low $\frac{dB}{dt}$ case. That is, the loss due to beam that is never captured inside the Rf buckets. One possibility is that the lower $\frac{dB}{dt}$ at injection allows for a wider bucket in phase which allows more of the *dc* beam to initially fit in the bucket. The magnet cycle has been made to rise as fast as it does because of the hypothesis that certain loss mechanisms are stronger at lower momentum, particularly vacuum related loss mechanisms (to be discussed below). Therefore, the sooner the beam reaches a higher momentum, the smaller the total loss will be due to these mechanisms. However, this observation supports lowering the initial acceleration rate.

¹⁸K. Zeno, *Heavy Ion Startup Book II*, pg. 97, 1997

¹⁹K. Zeno, *Heavy Ion Startup Book II*, pg. 112-113, 1997

Figure 7.17: The effect of the initial counterphasing level on losses after Au^{+32} injection. The overlayed traces are the injection transformer with the initial counterphasing level set at 4.0V and 4.6V. The trace that appears to have less loss *after* injection has the level set at 4.6V. The bottom traces are an injection kicker's current, and a TTB current transformer. Figure 6.6 shows a typical function for Heavy Ions.

Figure 7.18: The effect of changing $\frac{dB}{dt}$ during early acceleration on the early slow loss. The two overlaid circulating transformer traces show this effect (top traces). Lowering $\frac{dB}{dt}$ relative to normal running reduced the *capture loss*. The third trace is the Booster main magnet current, the fourth trace is the BX1.RFB3.V_GAP_DC (B3 gap volts).

One might suppose that raising the initial vector voltage could prevent this capture loss at the normal $\frac{dB}{dt}$. However, raising the vector voltage (lowering counterphasing) causes losses as well (see figure 7.17). The mechanism whereby higher voltage causes loss during this part of the cycle is unclear. Higher voltages cause bunch shape oscillations (figure 7.13), which cause larger longitudinal emittance. This, by itself, does not adequately explain loss early in the cycle where the available voltage is ample. One possible explanation is that there is some type of *momentum* aperture in the machine. A larger energy (or momentum) spread, leads to an increase in the transverse beam size due to dispersion effects. Raising the voltage causes the momentum spread to increase, which increases the beam size. Raising it too much causes beam loss, perhaps like what is seen in figure 7.17.

7.2.3 Early Losses and the Booster Vacuum

The slow loss also appears sensitive to the Tandem pulse width. Figure 7.19 shows the first few milliseconds of the cycle with a 300 μ s and a 600 μ s wide Tandem pulse.²⁰ When *normalized* to the stored intensity, the slow loss is less in the 300 μ s case. Over the 3.5 ms after injection, the *normalized* beam loss is approximately,

$$\frac{\text{amount of beam lost}}{\text{amount of beam initially stored}} = \frac{1.5V}{6.9V} = 0.22$$

in the wide pulse case. This compares to a loss that is not measureable on the scale shown in the short pulse case, and therefore less than $\frac{0.2V}{1.9V} = 0.11$. The pulse width was varied by adjusting the width of the source pulse at Tandem (TTM.SOURCE_WIDTH on spreadsheet). Varying this pulse width should not be done without consulting with Tandem operations first. Due to the nature of the injection process, one would expect the transverse beam size to be smaller with a shorter pulse. Also, less total loss generally occurs with a shorter pulse.

There has been speculation that one of the loss mechanisms may be related to the gas that resides in the vacuum chamber, called *residual* gas.

²⁰C. Gardner, *Heavy Ion Startup Book III*, pg. 34-35, 1997. A Studies report was also written which analyzed this, and other related data. see L. Ahrens, AGS Studies Report No. 353, "Gold Injection into the Booster: Beam Survival as the Length of the Tandem Current Pulse is Varied", Feb 24, 1997.

Figure 7.19: The effect of the Tandem pulse width on the loss rate *after* injection. The top trace is a current transformer, the middle (noisy) pulse is the vertical tune quad current, and the bottom two traces are an injection kicker current, and a TTB current transformer. The top picture is the 600 μ s case, below it is the 300 μ s case. Notice that the vertical tune quad current varies with respect to injection time. This is because the time after BT0 that injection occurs varies. The vertical tune quads follow a real time reference, and injection is triggered by a gauss event (Peaker).

Beam loss occurs when an ion changes its charge state due to interaction with the residual gas. Afterwards, the ions's rigidity is wrong for the main magnet field. Its orbit will fall outside the vacuum chamber, and no longer circulate.

For a constant concentration and make-up of this gas, and constant velocity beam, one would expect the loss rate, due to this *background gas*, to be proportional to the *instantaneous* beam intensity. Expressed mathematically,

$$\frac{dN}{dt} = -\alpha N(t), \text{ or } N(t) = N(0)e^{-\alpha t}$$

where N is the number of particles in the machine at a time t from injection, and α is the loss rate per ion. This equation leads to exponential decay if the above conditions remain valid.

Over the first few milliseconds of the cycle, $\frac{dB}{dt}$ is low, so the average particle velocity doesn't change much. Additionally, if the make-up and concentration of gases remains unchanged, and other loss mechanisms are not considered, one might expect the loss due to this vacuum related effect to follow the above equation. If N_w is the Booster intensity with the wide pulse, N_n the narrow pulse intensity, and t_1 is a *small* time after injection ($t=0$), then

$$\frac{N_w(t_1) - N_w(0)}{t_1} \frac{1}{N_w(0)} \approx \frac{dN_w(0)}{dt} \frac{1}{N_w(0)} = -\alpha$$

and,

$$\frac{N_n(t_1) - N_n(0)}{t_1} \frac{1}{N_n(0)} \approx \frac{dN_n(0)}{dt} \frac{1}{N_n(0)} = -\alpha$$

so,

$$\frac{N_w(t_1) - N_w(0)}{t_1} \frac{1}{N_w(0)} \approx \frac{N_n(t_1) - N_n(0)}{t_1} \frac{1}{N_n(0)}.$$

Therefore, the measurement made above for the two pulse widths should give a similar result in both cases if the loss were due mainly to this vacuum effect. As discussed earlier, there are other mechanisms which contribute to this beam loss. Since the wide and narrow pulse *normalized* loss rates are not the same, mechanisms besides this vacuum effect are responsible for a significant amount of the loss.

Besides the loss mechanisms already discussed, another vacuum related mechanism has been proposed. Since this *normalized* loss rate seems to increase as the injected intensity is increased, intensity related vacuum effects have been investigated. One speculation is that the vacuum is deteriorated by beam loss. When beam hits the surface of the vacuum chamber it causes particles to be ejected into the vacuum chamber, which further deteriorate the vacuum. However, there are no vacuum measurements that have been made which support this hypothesis.²¹

Since the beam size with the shorter pulse width tends to be smaller, this simple pulse width test might lead one to suspect that the *transverse* emittance of the stored beam is somehow related to the (normalized) slow loss rate. It is important to realize that the Booster *transverse* emittance is determined by the Booster injection setup, *not* the *transverse* emittance of the Tandem beam.

The losses during early acceleration are sensitive to the following:

- 1) Transverse
 - The skew quads, or linear coupling.
 - A tune shift, particularly a vertical tune shift (figure 7.10).
- 2) Longitudinal
 - The counterphasing, or vector voltage (figures 7.13 and 7.17). This may effect the losses associated with the dispersion related component of the transverse beam size.
 - $\frac{dB}{dt}$. Lowering $\frac{dB}{dt}$ may improve capture by increasing the bucket width (figure 7.18).
- 3) Vacuum related losses

Many of these mechanisms may be interrelated. In particular, the tune shift may reduce the coupling which is associated with the skew quads. The beam size may be affected by these mechanisms, which is also related to the vector voltage through the dispersion. Studies have been performed which support the idea that a component of this slow loss is momentum dependent. These momentum dependent losses are believed to be vacuum related.²²

As can be seen in figure 7.18, with the normal magnet cycle, most of the loss occurs before about 10-20 ms after injection. Aside from the possibility

²¹H.M. Calvani and L.A. Ahrens, *Au³²⁺ Beam Intensity Losses in the AGS Booster Due to Charge Exchange Processes*, Booster Technical Note No. 228, October 1, 1996.

²²H.M. Calvani and L.A. Ahrens, page 1

of a loss near the bucket minimum (see figures 7.14 and 7.16), accelerating the beam from this point to extraction is not usually a problem.

7.3 Extraction

7.3.1 An Overview

There are essentially two different types of extraction used within the AGS complex. They are called Fast and Slow extraction. The names refer to how long it takes for the beam to be extracted. Fast extraction takes place on a time scale the order of a microsecond, whereas the time scale for slow extraction is on the order of a second. Fast extraction, the kind used in the Booster for transfer to the AGS, is further subdivided. Fast extraction can either be synchronized to the AGS Rf system, or not synchronized. The kind typically used for HIP is *not* synchronized to the AGS Rf system. Non-synchronized fast extraction is relatively straightforward. This type of extraction is what will be described here. I'll try to describe *synchronized extraction* later, in the Proton section.

In fast extraction, the beam is extracted in *bunches*. In the Booster, all of the beam that is extracted, is extracted during one revolution of the beam around the machine. There are typically 8 bunches in the Booster when running Au^{+32} . A dipole magnet is used to *kick* the bunches out of the machine. In order not to disrupt the beam *before* it is extracted, the field in this magnet changes very quickly. This magnet, called a kicker magnet, is located just *upstream* of where extraction occurs.

A reasonably complete description of extraction from the Booster must involve certain aspects of the AGS configuration. The beam coming from the Booster has to be accepted by the AGS. In non-synchronized extraction, there is no Rf voltage in the AGS when the beam from the Booster enters it. The beam comes into the AGS in bunches. These bunches circulate around the AGS. The bunches gradually spread out due to the momentum spread of the particles within each bunch. After about a millisecond, the beam is more or less uniformly distributed about the AGS ring. This process is called *debunching*. After the beam has debunched in the AGS, the Rf voltage is turned on, and the beam is captured and accelerated.

The AGS has to be configured to accept a beam, whose momentum

matches that of the beam extracted from the Booster. The timing of the AGS injection equipment must also be synchronized to Booster extraction, and Booster extraction equipment. Hence, many of the triggers used to initiate Booster extraction are used for AGS injection equipment as well.

Extraction is only one part of the process of transferring the beam from the Booster to the AGS.²³ Each part of the transfer process depends on the other parts. The parts should not be considered in total isolation from each other. Additionally, what goes on in the Booster up until extraction effects how well the beam is extracted, and transfer in general. This is mainly because the efficiency of the extraction (and transfer) depends on the characteristics of the beam. These characteristics are determined by what precedes extraction. Probably the most important characteristics of the beam are the transverse and longitudinal emittances. If the beam is larger, it will be harder to extract.

There is a stripping foil used in BTA to change the Au^{+32} beam into Au^{+77} beam. Other charge states are made as well but about 65% of the Au^{+32} comes out of the foil as Au^{+77} . The Au^{+77} beam has scattered, and deposited energy in the foil. It now has a larger momentum spread (corresponding to a larger longitudinal emittance) due to the energy deposition, and a larger divergence (corresponding to a larger transverse emittance) due to its scattering off the particles in the foil. In particular, the momentum spread may cause some difficulties during capture in the AGS. The rigidity of the Au^{+77} beam is $\frac{32}{77}$ that of the Au^{+32} beam. The magnet settings on the Au^{+77} side of the foil are scaled by this fraction to accomodate this lower rigidity.

7.3.2 Synchronizing the AGS and Booster Cycles

The Booster and AGS cycles need to be synchronized, so that the beam is injected into the AGS at the appropriate time in the AGS cycle. There needs to be a fixed, periodic relationship between the Booster and AGS cycles for this to occur on a regular basis. This fixed, periodic relationship in time is set by the supercycle. The supercycle repeats itself on the order of every few seconds. Within each supercycle, specific supercycle events (or triggers) are

²³The BTA line and other transport issues are covered in AGS/AD/Op. Note No. 43, "Transport from the Booster to the AGS", by Paul Sampson, September 6, 1994.

generated at the same times relative to the beginning of each supercycle.

Some of these events are responsible for initiating Booster and AGS cycles. Consequently, the Booster and AGS cycles have a fixed, periodic relationship to each other. The AGS T zero supercycle event, whose mnemonic is AT0, is used to trigger the reference function for the AGS Main Magnet. As described previously, Booster T zero (BT0), triggers the Booster Main Magnet reference function. Consequently, there is a fixed, periodic relationship between the field in the Booster and that in the AGS. Typically, the AGS Main Magnet function, and the placement of AT0 in the supercycle, are configured so that Booster extraction occurs at the appropriate time(s) in the AGS cycle.

Figure 7.21 is the Superman application main window. It shows an *archive* of a Supercycle that was active during an Au^{+32} run. The top graphic shows the relative timing of the Booster and AGS magnet cycles. There are 4 Booster cycles in this 4 second (240 jiffy) supercycle. A *timeline* representation of these supercycle events, together with other supercycle events, is shown below the magnet cycles.

This timeline representation is worth looking at in some detail. Each PPM user, within a supercycle is initiated by a *user reset*. AU1 (AGS User 1 Reset), BU2 (Booster User 2 Reset), and BU3 (Booster User 3 Reset) appear in this particular supercycle. After each user reset, a Prepulse event occurs. As shown in the timeline representation, BPP (Booster Prepulse) and APP (AGS Prepulse) appear in this supercycle. As can be seen in the timeline, a prepulse event also occurs before every T zero event. Prepulse events are used to trigger equipment which needs more time to get ready for the upcoming cycle than a T zero event would provide. When a T zero occurs the respective magnet function, and other equipment start to follow their references. Some equipment needs to be prepared for this. For example, the High level Rf systems, and equipment that needs to charge up (such as extraction kickers) use prepulse events.

After each prepulse event is a *field reset* event. A field reset event is used in the initialization, or *calibration* of the gauss clock for the next cycle. On a cycle with this event the Gauss clock is *jammed* with a value for the field when the real timeline event BGT.JAM_VALUE occurs (see section 2.2.3 for more information). BFR (Booster Field Reset) and AFR (AGS Field Reset) occur in this supercycle.

After each field reset, a T zero event occurs. There are four BT0 events

occurring on this supercycle, and one AT0. Therefore, there are 4 Booster cycles, and 1 AGS cycle in this supercycle. However, only three Booster magnet cycles are visible in the top graphic. These three magnet cycles are User 2 magnet cycles. This is apparent because they start at the same time on the timeline as the BT0 events between the BU2 and BU3 events.

There is a user 3 Booster cycle, with its associated BT0, which follows the BU3. However, it does not show up on the top graphic. This user 3 cycle was used as a *study cycle*. In general, it is not necessary to associate a particular magnet function with this, or any user for a supercycle to be *well defined*. The supercycle and its archive consist only of timing events. The magnet function is changed through the Booster (or AGS) MainMagnet application. The Superman application reads the live magnet function through the controls system and simply displays it. These supercycle events generally follow certain *rules*, mostly about their relative timing, that are defined within the program. Any deviations from the rules show up in the first two *timelines* called '*Codes Added*' and '*Codes Deleted*'. Apparently, due to these *deviations* the Superman program does not display the Booster user 3 cycle.²⁴

At one point the live Main Magnet function for user 3 ramped up to a relatively low field value, say 1000 gauss, and stayed there for about 2 seconds. If it had appeared on this representation, it would have started at the BT0 following the Booster user 3 reset (BU3). The extraction events would typically be disabled on a studies cycle. Notice that user 3 does not interfere with the HIP parts of the supercycle. The Booster would not have been used during this part of the supercycle for anything.

The TT0 supercycle events, visible on the Tandem timeline representation, prepare the Tandem for the beam request that occurs at Tbeam_on within the Booster cycle immediately following it. The TT0 events determine the Booster cycles that will have beam. The beam is typically extracted from the Booster near the peak of its magnet cycles. Notice that the peaks of the last two user 2 Booster magnet cycles, occur while the AGS Main Magnet is at its lowest field value. This level is called the *dwell* field. The dwell field typically corresponds to the field value required for the beam coming from the Booster. It is set in the AGSMagnet application. The time that AT0 occurs in the supercycle can be shifted in jiffy sized increments across

²⁴As per a discussion with N. Williams.

the supercycle. It is preferable to have the AGS Main Magnet field ramp up as soon as the beam in the AGS has been captured.

The first Booster user 2 cycle has no associated TT0, and therefore no beam on it. This is a typical setup. The first cycle is often called a *dummy* cycle. Some equipment, particularly the Main Magnet, responds differently to the same reference functions on the first cycle within a regular group of cycles. In order to make all the cycles with beam as much alike as possible, each group of Booster cycles starts out with a *dummy cycle* that is used to initialize equipment for that group of cycles.

The second Booster User 2 cycle does not have a TT0 associated with it either. This cycle could be used to get higher intensity in the AGS. Inserting a TT0, associated with this cycle, into the supercycle, causes beam to appear on this cycle. Since it is the same user, the setpoints for the devices on this cycle will be essentially the same as for cycle 3. The beam from both cycles will be transferred to the AGS, resulting in more beam in the AGS.

Events such as BT1, BT3, and BT4 are called *flavored* T zeros. They are used primarily to configure the Rf system properly for Booster to AGS transfer. Their use will be described below in more detail. The AGE and BGE events are AGS and Booster *group ends*, respectively. They are primarily used by the controls system to collect data, like readbacks, from the previous cycles.

7.3.3 The F3 Kicker

When the beam is ready to be extracted, the kicker's field should be that required to kick the beam out of the machine. The 8 bunches have been circulating around the machine up until this point. Ideally, they should *not* have been affected by the field in the kicker up until this time. The bunches pass by the *kicker magnet* at the Rf frequency. This is typically about 5 MHz at extraction energy. So, one bunch passes through it every 200 ns. When the beam is ready to be extracted, the field in the kicker magnet rises from zero to the required value for extraction. It is able to do this in the time between the passing of one bunch and the arrival of the next!

The kicker is located at F3. The opening in the beampipe that the extracted beam enters is at F6. Figure 7.21 shows extraction with the current

Figure 7.21: Au^{+32} beam kicked by the F3 extraction kicker. Top trace is the current in the extraction kicker. This signal is hardwired at MCR4. Below it is another hardwired signal, the C4 PUE sum signal. It is the sum of the two PUE plates at C4. This signal gives an (ac) intensity measurement. See text for more information.

waveform for the F3 kicker, and the Booster beam.²⁵ The PUE at C4 is used to look at the bunched Booster beam. The 8 bunches pass by this PUE repeatedly, showing up as *spikes* on the C4 trace. Once the beam is kicked a *hole* appears in the beam seen at C4, indicating that some of the bunches are no longer in the machine. However, some beam still remains circulating because the kicker pulse is not long enough.

The kicker *risetime* is about $0.2 \mu s$ according to the kicker trace. The kicker is timed here so that the first bunch to be kicked, which sees the rising edge of the kicker pulse, is only partially kicked. The kicker field only remains at a value high enough to kick the beam *out of the Booster aperture* for about $1.2 \mu s$. Since it takes $\frac{8}{5 MHz} = 1.6 \mu s$ for a bunch to make one revolution, all the bunches do not get kicked out of the aperture. It kicks about 6 of the 8 out of the machine. At least part of the bunch that was disturbed by the rising edge of the kicker pulse stays circulating. By the time the bunch preceding this partially kicked bunch makes another revolution, the kicker

²⁵K. Zeno, *Heavy Ion Startup Book I*, pg. 78, 1995-96

current has fallen back close to zero, and it receives very little kick. This bunch also continues to circulate. All this can be seen rather clearly in figure 7.21.²⁶ Figure 7.22 illustrates it in a diagram. The fact that the PUE and the extraction aperture are at different positions, C4 and F6 respectively, and cable delays in the signal transmission back to MCR, accounts for the delay seen on the scope between the kick and the disappearance of the beam.

Under the best circumstances about 6 of the 8 bunches can be seen on an AGS PUE or wall current monitor. Presently, the F3 kicker power supply is in the process of being replaced. In the future, its pulse width should be long enough to kick out all of the Au^{+32} beam.

For extraction to be reproducible from cycle to cycle, the time that the kicker *fires* must be fixed with respect to the time that a bunch passes through the kicker. Since the Rf and beam *phases* are essentially locked through the Rf beam control system, a trigger synchronized to the beam is available through the beam control system. The Rf system uses a trigger from the control system called BRF.COGGING_ARM to generate the kicker trigger. When BRF.COGGING_ARM occurs, it *arms* part of the low level Rf system called the *cogging box*. When the phase of the Rf matches some predetermined value, the cogging box generates a trigger which is sent to the kicker. This trigger is also sent to the A5 kicker which is in the AGS. The A5 kicker kicks the beam onto the AGS orbit, just as the F3 kicker kicks the beam off the Booster orbit. Injection into the AGS is very much like extraction from the Booster reversed in time.

Once the F3 receives its trigger, there is an additional delay before the kicker fires. The trigger from the Rf system reaches the F3 kicker electronics at the same time relative to the bunches every cycle. It synchronizes the bunches to the kicker. A delay after this trigger is employed to adjust the relative timing of the kicker pulse and bunch. This delay is set through the SLD BXT.KDHF3.IO, it has nanosecond resolution. Another similar delay is used for the A5 kicker. It is called ABI.KDHA5.IO.

The A5 kicker needs to know when the beam is going to pass through it. It receives the same trigger from the cogging box that the F3 does. To within a negligibly small range, the beam has the same momentum at extraction

²⁶In this case, where the kicker pulse is not wide enough to kick all the bunches, the optimum timing may have the kicker current rising *while* a bunch is passing it (as in Figure 7.21)

Figure 7.22: The bunches kicked by the F3 kicker. The bunches in the Booster are illustrated at four times relative to the kicker pulse. (1) is at the start of the pulse, all the bunches are still circulating. (2) is the situation after half a revolution around the machine. The picture is simplified by treating any beam passing through the kicker while the kicker is on as immediately extracted from the machine. (3) shows the situation after $\frac{3}{4}$ of a revolution. That is, just before the kicker pulse ends. (4) shows the situation after another $\frac{1}{2}$ revolution. The two bunches in figure (3) pass through the kicker *after* the pulse is over. Compare this to figure 7.21.

every cycle (to be discussed below). So, the transit time of the beam from when it is kicked at F3, to when it passes through the A5, is essentially fixed. As a result, the time interval between when the Rf trigger is generated, and when the beam passes through the A5 is fixed. This assumes the F3 fires the same way every time and its delay is not changed. Therefore, by adjusting the A5 kicker delay, the A5 can be set to fire just before the beam passes through it.

7.3.4 The Kicker, the Bump, and the Septum

In general, the kicker is not, by itself, sufficient to kick the beam efficiently out of the machine. There are two main reasons for this. First, because of the very strong time constraints placed on it, rising to its peak current between bunches, it was not designed to have a field that is strong enough to accomplish this. The closer to the outside of the beampipe the E.O. is *at the extraction point*, the *less* the beam has to be kicked to extract it. And consequently, the less current the kicker has to have. By distorting the Equilibrium orbit around the extraction point (F6), the circulating beam is brought closer to the opening in the beampipe. This orbit distortion is called the *Extraction Bump*.

After the beam is extracted from the Booster, it passes through the *F6 Extraction Septum magnet*. The F6 septum magnet, which is in the F6 straight section, has a vertical division in it. On the *beam left* side of the division, there is a magnetic field which bends the beam away from the Booster (to the left) and into BTA. On the right side of this division, or Septum, there is no field. Beam on the right side continues to circulate in the machine. The septum has a thickness of 10 mm or less.²⁷ The magnetic field bends the beam on the left side of the septum away from the Booster by a large amount.²⁸ Both the position and angle of the beam as it enters the F6 septum magnet are important for efficient transmission through it, and the upstream elements of the BTA line.

The extraction bump is not only important for moving the E.O. up against the inside edge of the extraction septum (or the outside of the Booster beampipe). It also allows the position and angle of the beam as it enters the

²⁷*The Booster Design Manual*, December 5, 1988, pages 5-3

²⁸142 milliradians, from *The Booster Design Manual*, pages 5-3

left side of the extraction septum (the extraction trajectory) to be adjusted independently. The extraction bump is composed of four dipoles. With four dipoles, the angle and position of the E.O. at one point can be adjusted independently. In this case, that point is designed to be the s coordinate of the start of the F6 extraction septum.

The kick from the F3 kicker changes the extraction trajectory. However, for a given change in position at F6, the change in the angle is determined by the details of the Booster lattice from F3 to F6, not by the kicker. In other words, one cannot adjust the position and angle of the extraction trajectory independently with just F3. The extraction bump is used to adjust the two independently. The intent of the F3 is to change the position of the beam at F6 by changing its angle (relative to the E.O.) at F3. When the beam is kicked at F3 it starts to do Betatron oscillations about the E.O. It behaves as if it were a particle initially displaced in angle from the (distorted) E.O. As it passes the extraction point, the phase and amplitude of the oscillation at that point may place it *outside* the Booster acceptance, and hopefully, on the left side of the F6 septum, and not on the septum itself.

Typically, the F3 kicker runs near its maximum value. In general, the smaller the horizontal beam size (or emittance) is at extraction, the easier it is for the kicker to kick all the beam from one side of the septum to the other. If the beam is small, the extraction bump can move the circulating beam closer to the septum before it starts to scrape there. Consequently, the kicker does not have to kick the beam as hard to get it to the other side. Additionally, since the beam is smaller, it does not have to be kicked as far over on the other side in order to miss hitting the septum.

7.3.5 The Extraction Bump and its *Residuals*

The extraction bump is created by putting current through additional windings on four of the main magnet dipoles near the extraction point. The *main windings*, through which a current flows that produces the main bending field, are powered by the Main Magnet power supply. These additional windings (incorrectly called *Backleg windings*) are wound around the main magnets, but are not connected, in series or in parallel, to the *main* windings.

The dipole field in each magnet is approximately proportional to the sum of the currents in its windings. Consequently, current that passes through the backleg windings (or '*backlegs*' for short) is subtracted or added to the current

that passes through the main windings to alter the field. Any deviation from the main field can be thought of as an additional dipole with a polarity either the same or opposite to that of the Main Magnet field. The main dipole magnets at F2, F4, F7, and A1 have backleg windings powered by separate supplies. The F2 and F7 backleg windings add to the main bending field, kicking the beam inward. The F4 and A1 windings have the opposite polarity. The setpoints and On/Off status for them are controlled through SLDs found in the Booster/Extraction/Ring branch in Spreadsheet.

To complicate matters, the current in the backleg windings for a given setpoint is a function of the dB/dt of the Main Magnet. The changing Main Magnet field induces a voltage on these windings.²⁹ This induced voltage causes a deviation in the current from the setpoint. The details of the response of the current to $\frac{dB}{dt}$ have to do with the details of the power supplies which feed current to the windings.

As with the injection bump, an application *code* was made to adjust the extraction bump. It can be found in the BoosterOrbitControl program (under 'process /orbit bumps /extraction /TDH slow bump /4-bump). However, the application code does not compensate for the effect of the induced voltage on the power supply currents. For example, if the application code configures the bump to have $x=31$ mm and $x'=-2$ mrad at the extraction point, it will send the appropriate setpoints for this. However, it assumes that the readbacks will follow the setpoints. But the readbacks do not follow the setpoints, so the bump that results will not be the one requested.

There is a significant advantage to using an application code over adjusting the four magnets independently (as SLDs on spreadsheet). The advantage is that an application code can configure the currents in the magnets to make an orbit *bump*, not just an orbit *distortion*. Ideally, an orbit bump is an orbit distortion *localized* around some point at which one intends to change the orbit's position and/or angle.

Figure 7.23 shows a *model* of a typical Heavy Ion Extraction bump. The initial deflection of a particle on the E.O. occurs as it enters F2. Given some initial (x, x') of the E.O. relative to the center of the beampipe at F2, the position of the E.O. at the F6 septum could be controlled by adjusting the amplitude of the kick at F2. However, the orbit's angle, x' , at F6 is dependent on the orbit's position, x , at F6. By adjusting the kick at F4 as well, both the

²⁹The induced voltage is described by Faraday's law, $\nabla \times E = -\frac{1}{c} \frac{\partial B}{\partial t}$

Figure 7.23: A model of the Booster extraction *4-bump* in the extraction region. On the top of the graph the locations of various elements that determine the shape of the bump, and the F6 septum, are shown. The beam exits the machine at the F6 septum. Before the F3 kicker fires, the E.O. of an *on-momentum* particle follows the solid curve. The solid curve is the orbit distortion caused by the four backleg windings (at F2 (146.8 A), F4 (424.5 A), F7 (219.5 A), and A1 (595.2 A)). (x, x') at F6 is (16.6 mm, -1.0 mrad). When the F3 kicker is energized, or *ON*, the beam follows the *extraction trajectory* and exits at the F6 septum.

position and angle at F6 can be determined independent of each other. The A1 backleg winding (BLW) produces the last kick the beam receives before it leaves the extraction region. In order for no orbit distortion to occur *outside* the extraction region the beam has to have the same (x, x') as it would for an undistorted orbit as it exits the A1 magnet. Since the position (x) at A1 is not affected by a kick at A1, it must be the same at A1 as it would be if there were no distortion. The kick from F7 is used to make x *the same* at A1. Now, the angle (x') at A1 will need to be adjusted to make it *the same* as for the undistorted case. The A1 kick accomplishes this.³⁰

Also shown in figure 7.23 are the locations of the main quadrupoles in the region. As the beam on the distorted E.O. passes through these *off center* it is also kicked by them. The effects of these kicks change the shape of the orbit bump. Horizontally focussing quads (in even numbered half-cells) kick the beam back towards the beampipe center. The vertically focussing quads (in odd numbered half-cells) kick the beam away from the center. The angle with respect to an idealized closed orbit passing through the center of the beampipe (and the quads), x' , is also plotted. Notice that changes to x' as it passes through the quads are proportional to how far *off center* the beam is at that s coordinate.

If the action of the 4 bump magnets causes the (x, x') of the E.O. coming out of the A1 magnet to be different than it would be for the undistorted case, then the orbit around the machine will be altered. Figure 7.24 is a model of an extraction bump which does not have the same (x, x') at the exit to A1 as the orbit which has not been distorted by the extraction bump magnets. For simplicity, this undistorted orbit is taken to be $(x, x') = (0, 0)$ throughout the ring, ie.- a flat line. Of course, in a real machine this is never the case. The orbit displayed here can also be thought of as a so-called *difference orbit*. A difference orbit looks at the difference in the orbit between two different states. Here, the states are the bump 'on' and the bump 'off'. The difference orbit shows the effect of the bump on the orbit. Real difference orbits can be taken with the BoosterOrbitDisplay program. If the bump were 'perfect', this difference orbit would be flat and run through the center of the beampipe at all points except between F2 and A1. The fact that it does not indicates that it is not 'perfect'.

In fact, the amplitude of the orbit distortions outside the extraction re-

³⁰See Appendix A for a more analytical description of this *4-bump*.

Figure 7.24: A model of an imperfect Booster extraction bump and the orbit distortions it creates around the ring. Also shown are the extraction trajectory, and locations of the bump magnets and F6 septum.

Figure 7.25: A model of a Booster extraction bump with small *residuals*. The solid curve is the orbit distortion caused by the four backleg windings (at F2 (182.7 A), F4 (-588.2 A), F7 (103.6 A), and A1 (-677.0 A)). (x, x') at F6 is (23.0 mm, -1.5 mrad). Also shown are the extraction trajectory, and locations of the bump magnets and F6 septum.

gion, called '*bump residuals*', are a significant fraction of the bump's amplitude. These *residuals* effect the bump itself because they alter (x, x') at the entrance to F2. Figure 7.25 shows an extraction bump that is closer to ideal, at least in so far as the *bump residuals* are concerned. Figure 7.26 is an actual extraction orbit obtained through the OrbitDisplay program. Though it is for protons, it is very similar to an orbit one might expect to see for Heavy Ions.

At first glance it might appear that the orbit distortions outside the extraction region are to be avoided. They have the potential of causing beam loss outside the extraction region because the beam makes excursions far from the beampipe center. However, they are not necessarily detrimental. The Booster was designed so that the outside aperture which limits the Booster's acceptance is the entrance to the F6 septum. Roughly speaking, this is the innermost outside aperture of the machine. Beam that scrapes off on the outside of the beampipe, will tend to scrape off there, not at other places.

Figure 7.26: An orbit from the BoosterOrbitDisplay program taken just before extraction, while the extraction bump is on. It is a proton orbit, but is similar to what one would expect for Heavy Ions. The bpm measurements are the dots. The dots are connected by straight lines.

However, large residuals may cause the beam to scrape off on the inside apertures as well. The Beam Dump is the outermost inside aperture. It will be described in section 9.3.1. In particular, large residuals may cause beam loss there.

In part, due to the complications associated with the $\frac{dB}{dt}$ dependent response of the backleg winding power supplies the extraction bumps have always been tuned using the SLDs on spreadsheet. This type of tuning generally results in bumps with large residuals. However, the transverse emittance of the beam at extraction energy, particularly for Heavy Ions, is relatively small compared to the machine's acceptance. Or, in english, the beam at this energy is small relative to the beampipe. So, the beam has room to move around inside it without scraping.

There are analog signals available on Xbar for each of the extraction bump power supplies. The initial timing event for extraction, which plays a role analogous to the one that *peaker* plays for injection, is called BRF.XTRCN-STRT.GT ('*extraction start*' for short). It is a gauss event, that occurs almost precisely 5 ms before extraction. The trigger, described earlier, that is responsible for arming the Rf cogging box used in kicker timing is a microsecond resolution delay from this event. For extraction that is *not* synchronized to the AGS Rf, this is set to 5000 μs . The kicker trigger occurs within microseconds after *cogging arm* occurs. Consequently, *Extraction start* determines the field at which extraction will occur, as well as the time in the cycle. The *discharge triggers* for the extraction bump and F6 septum power supplies are microsecond resolution delays from this event as well. BXT.B.BLW_DISCHARGE and BXT.F6_DISCHARGE are their respective SLD names (found in Booster/Extraction/timers on spreadsheet). All these supplies use capacitor banks which charge up during the cycle, and discharge when these triggers occur. They share a common discharge trigger called BXT.F6&BLW_STCHARGE found on the same page.³¹

The extraction bump p.s. current pulses have the shape of a half sine wave and last about 5 ms. Figure 7.27 shows the current pulses for each of the four power supplies during a Heavy Ion run. The spikes on each of the current pulses, near but just after their peaks, are due to noise generated by

³¹A detailed description of extraction timing and more information about extraction is given in "AGS/AD/Op. Note No. 42, *Booster Extraction*, by Dave Gassner, January 25, 1994.

Figure 7.27: The extraction bump power supply current outputs.

the discharge of the F3 kicker. They indicate when extraction occurs. Ideally, extraction should occur at the peak of the current waveforms. Particularly, if it occurs *after* the peak, the bump will reach its maximum value *before* extraction. This may cause unnecessary beam loss. The current pulses could be moved earlier by moving the discharge trigger earlier.

7.3.6 Setting up Booster Extraction

Typically, the first step in setting up extraction is the restoration of the last run's Store archive. There are usually some changes to the machine setup relative to the last run. These often include changes to the Main Magnet function. Consequently, the magnet cycle may have a different length, or the extraction momentum may be different. Additionally, the response of the radial loop may be different, so the beam may be at a different *radius* at extraction. It is also difficult to verify that the current in the F6 septum is the same. Changes to this current on the level of 1 part in 1000 have an effect on the beam that is measurable. If the magnet cycle (or function) is

different, many functions may have to be adjusted. Among them are the functions for the tune and dipole correction magnets.

Of the four multiwires in BTA, the first is particularly useful in setting up and optimizing Booster extraction. When displayed through the Beam-LineInstrument program, profiles from these multiwires, like those in TTb, are analyzed, and measurements of the beam's position and width in both the horizontal and vertical planes are shown alongside with the profile. In addition to this information, the area under the profile curves is calculated and displayed. This area is called the '*Intensity Sum*', because it is actually the sum of the voltages on all the wires. For a given multiwire, in a particular gain setting (High or Low, High is used for Heavy Ions) and momentum, this area measurement is a very useful way to gauge the *relative* intensity. This measurement does not (presently) give an 'absolute' number for the intensity at the multiwire. The program allows one to acquire multiwire data on a continuous, supercycle to supercycle, basis. As changes are made upstream of a multiwire, the *Sum* (horizontal and/or vertical) can be watched for changes, which are proportional to the beam intensity. Figure 7.28 is a BeamLineInstrument display for MW006 during a Fe^{+10} run.

MW006 sees the beam just as it exits the F6 septum. There are no BTA elements (dipoles or quads) upstream of it. On this multiwire the beam is typically wider vertically than horizontally.³²

After loading the Store archive and confirming that supplies are working, its a good idea to check that beam is being extracted from the machine. The quick way is to look on MW006 for signs of beam. From the Booster side, a scope setup like that in figure 7.21 can be used to look for the disappearance of the bunches when the kicker fires. More than likely, the fine delay on the kicker timing will need to be adjusted. However, its more worthwhile to concentrate on this when setting up the timing for the A5 kicker while working on *AGS* injection.

As far as kicker timing goes, its important to check that the kicker's discharge time is *locked* to the beam's phase. No jitter should be visible between the kicker pulse and the bunches. If there is, the problem is usually associated with the *cogging box* timing. The most common cause of this

³²This is because the value of the Beta function here in the horizontal plane is small, and it is large in the vertical. Typically, the horizontal *emittance* is somewhat larger than the vertical. It was alluded to in section 5.3 that $x \propto \sqrt{\beta}$.

Figure 7.28: Multiwire MW006 in the BTA line during a Fe^{+10} run.

problem is the lack of *flavored* T0s at the right times in the supercycle. These are used in the cogging box electronics. BT4 is required to *reset* the cogging electronics. It must occur on the supercycle. The other *flavored* T0s (BT1, BT2, BT3) occur on each of the cycles with beam and extraction *enabled*. For example, in figure 7.20, BT1 occurs on the only cycle with extraction enabled. BT4 occurs after this.

As with much of machine setup, extraction setup sort of works backwards. If beam is not visible on MW006 after the initial setup, then you check to see what might be wrong. First, you might confirm that the F6 septum is on and near setpoint. If it is, since you don't know precisely enough what the current in the septum is (or should be), you might scan its current and look for signs of beam on MW006. If this doesn't work, you can check a current transformer, the wall current monitor, or a PUE, to see if the beam stops circulating at extraction time. If it does *not* disappear from the machine, you might check the kicker current pulse, verifying that it is occurring at the right time in the cycle. That is, within microseconds of *extraction start* + 5 ms. This can be done by triggering a scope at this time, and finding the kicker current pulse. If it is firing, but not at the right time in the cycle, *cogging arm* (BRF.COGGING_ARM) may not be set to the correct time (5000 μ s for *unsynchronized* transfer).

The type of setup technique that I'm describing can be summarized as follows:

- 1) Confirm that the machine is as close to its previous running state as reasonably possible. This is done by confirming that the *Store* archive has actually been loaded to all the relevant devices, and that those devices are on and at appropriate setpoints. It is generally not worthwhile to *go crazy* and check every single device meticulously. When you start looking at the beam, its behavior will allow you to diagnose the problem. The beam is the primary *diagnostic tool*.

- 2) Start looking at places where you expect (or hope) the beam will be. If it is not there, diagnose the problem using the tools available, and your understanding of the processes involved. Start with the most basic and fundamental aspects. Is the beam reaching extraction? Is the multiwire working (maybe the beam is there)? Is it being kicked out? Are the bumps, the kicker, and/or the septum on and at setpoint?

- 3) As a problem is narrowed down, you can start looking for more obscure causes.

Tips for Optimizing Booster Extraction

Once beam is visible on MW006, the rudimentary aspects of extraction are established, and a valuable tool for setting up extraction properly (MW006) is now usable. With the beam visible on MW006, it can now be centered horizontally there by changing the F6 septum current.

Its often a good idea to scan the extraction radius at this point and try to optimize the sum on MW006. The radius is changed by adjusting the radial steering function during this time. The radial steering function starts at peaker. To *adjust* the radius, put 3 points on the function, as in figure 6.3. In this case, the first point (the one that occurs at about 40 ms from peaker in figure 6.3) should be put about 10 ms before extraction. The second point might occur, say 5 ms, after that. And the third, well after extraction time. Set the level on the last two points to be the same. In addition to the current transformer trace, a radial signal BXI.RFE6.RADIALAVG, called the '*radial average*' signal could also be watched.

Changing the radius, not only changes the average position of the E.O. in the beampipe, it also changes the shape of the orbit, and the extraction momentum. As the E.O. is moved away from the central radius, it passes off center through the quads, and other lattice elements. This causes an orbit distortion which has the same periodicity as the number of superperiods in the ring (6), and which repeats within each superperiod. The bump magnets act to distort this *off momentum* E.O. orbit. Consequently, the orbit at F6, relative to the beampipe center, is not just moved in or out by some amount due to the radius change, its angle relative to the beampipe center also changes. Additionally, changing the momentum of the extracted beam will have consequences for the transport of the beam through BTA, but its effects are particularly noticable in the AGS, especially during the Rf capture process. Sometimes this change in momentum, due to a radius change, can be compensated for by a change in the extraction field. The extraction field can be changed by changing *extraction start*.³³

It is also important to check for beam loss in the Booster *before* extraction. This beam loss can occur if the bump is too large, has large residuals, or has the wrong shape. Notice that the *peak* at F6 in figure 7.25 is more pronounced than it is in figure 7.24. Also, in figure 7.24, the amplitude of the distortion

³³The ability to compensate for a momentum change becomes much more complicated when extraction *is* synchronized to the AGS Rf.

at F7 relative to F6 is considerably greater than it is in figure 7.25. In figure 7.24 the bump shape is less optimum because of the relatively large excursion at F7. The horizontal beam size is smaller at F7 than at F6 since F7 is a horizontal ' β minimum' and F6 is a horizontal ' β maximum'. The outside edge of the beampipe at F6 is also closer to the beampipe center than it is at F7. Both these factors tend to reduce the amount of beam loss that may occur near the F7 peak, but it is doubtful that they make that beam loss impossible. A poorly shaped bump, may cause the beam to hit an aperture within the extraction region, say at F7, which is not the *intended* limiting aperture (the F6 septum).

One can check for beam loss associated with the bump on a current transformer signal just before extraction time. Its a good idea to look at an extraction bump p.s. current signal at the same time to check for a correlation between its amplitude and any beam loss. Beam loss associated with the bump will be visible as a drop in the current transformer amplitude as the bump rises. It may also be observable on a loss monitor analog signal (such as the F6 loss monitor). This beam loss may be easier to distinguish from losses immediately at extraction time (when the kicker fires) if the kicker is turned off.

If there is beam loss, the bump and/or radius probably needs further adjustment. Particularly with Heavy Ions, there is some flexibility in the extraction setup due to the fact that the configurations of the bump, radius, and kicker are *not* independent of each other. In fact, an optimum setting for one is very dependent on the settings of the others. The flip side of this is that there can be too many variables, and no clearly correct values for any of the parameters when they are viewed in isolation from each other. For example, the circulating beam does not *necessarily* need to be right up against the septum. Therefore, it might be possible to decrease the bump amplitude and/or move the radius to the inside to reduce the beam loss, while maintaining or improving extraction efficiency by increasing the kicker current. As the radius is moved to the inside and/or the bump is reduced, the loss will disappear. The sum on MW006 may go *down*. The F3 kicker can then be raised to recover this beam, and possibly find more.

Alternately, there may be no beam loss before extraction due to the bump and/or radius settings, but the kicker may not be strong enough to kick all the beam out, or may kick it out at a bad angle. It might be possible to find a better setpoint for F3 by scanning its setpoint, and looking at MW006 sum

for an increase in the extracted beam. The position on MW006 can then be centered with F6. As F6 is adjusted, the MW006 Sum will probably change. If its maximum occurs far from the multiwire center (more than 5-6 mm), this may indicate that the beam is not coming out of the machine well, and the bump and/or radius may require additional adjustment.

Even after these steps are '*complete*' it may not be possible to get close to a *good* position and angle at the extraction point without further adjustment of the bump. In general, the bump can be adjusted by using the bump SLDs. Unfortunately, it's difficult to know how the changes effect the bump residuals, the bump's shape, or the position and angle at F6.³⁴ There is a measurement readback facility in BoosterOrbitControl which allows one to get a position/angle measurement at F6 from the p.s. readbacks. This facility also calculates the residuals, decomposed into cosine and sine parts. If orbit data is available, residuals can also be measured directly from the orbit using the undistorted orbit as a reference, as in Figure 7.24. By calculating the (x, x') and the residuals of the bump from the SLD measurement readbacks, as these changes are made, it may be possible to track the effect of some of the changes on the bump and orbit. It may also be a good idea to watch the current transformer for any signs of beam loss *before* extraction.

The tune is also an important parameter. For a given set of currents in the bump magnets, changing the tune will alter the bump and the residuals.

Without an absolute intensity measurement downstream of F6 it is in general impossible to know how efficient extraction is. The AGS intensity is used to get an absolute measurement of the *transfer* efficiency. The difference between the early intensity in the AGS and the late intensity in the Booster may be due to losses that do not occur at extraction. Losses during the transport of beam from the Booster to AGS, and AGS injection, will both cause the early AGS intensity to be lower than the late Booster intensity.

Complicating matters further is the fact that improving Booster extraction efficiency will typically increase transport and AGS injection losses. A change to the Booster extraction setup generally changes the trajectory of the beam exiting the Booster. Consequently, BTA steering and AGS injection will need to be changed to compensate. So, it is not a good idea to

³⁴The increment feature of the extraction bump application code may not be as susceptible to the error introduced by the disagreement between setpoint and readback for small changes.

look at the early intensity in the AGS when just trying to improve Booster extraction. This number is primarily valuable as an indicator of the overall transfer efficiency between the two machines. The MW006 sum is the most direct way to gauge the effect of the changes one makes on the extraction efficiency.

Chapter 8

Proton Operation I: *The Usual Suspects*

8.1 LTB

8.1.1 Overview

The Linac delivers a 200 MeV H^- beam to the Booster through the LTB line. After the beam leaves tank 9 of the linac it passes through a straight section about 60 feet long. At this point the beampipe splits. One beamline continues on a straight path. The other line bends towards the left. This is the LTB beamline. The DH1 magnet is located at the fork and bends the beam by about 7.5° when it is on. Otherwise, the beam continues moving on a straight path.

The straight beamline that begins at tank 9 and continues past the beginning of LTB is called '*HEBT*'. *HEBT* is an acronym which stands for *High Energy Beam Transport*. Before the Booster was built, beam would pass down this line for injection into the AGS. The section downstream of DH1 is still used occasionally for measuring the characteristics of the Linac beam.

The magnet settings in the part of HEBT upstream of DH1 are typically set by Linac personnel. However, there are two horizontal and two vertical dipoles in this part of the beamline that we do control. MCR personnel control all the magnets in LTB.

The LTB line, from DH1 to the injection point in the Booster, is about

122 feet long. It is shown in figure 8.1.¹ As noted in Chapter 3, there are multiwires located at 35 and 107 feet from DH1, and a number of bpms in the beamline.

As can be seen in figure 8.1, the LTB line contains a very large bend. The beam is bent using the magnets DH2, DH3, DH4, and DH5. These magnets are powered together in series. The DC reference for the current is set through an SLD on spreadsheet (BLI.DH2-5.SPRB). All magnets in the LTB line are *DC*. They can all be found on the LTB page under Booster/Injection in Spreadsheet. The magnets DH2, 3, 4, and 5 are all the same type. Hence, they all bend the beam by the same amount. *Steering* magnets are denoted by their distance from the start of LTB (in ft.). *Steering* magnets are used for fine adjustments. There are only steering magnets for the vertical plane. The LTB line also contains 13 quadrupole magnets.

8.1.2 Steering and Matching

The description of particle motion in a beamline is deceptively similar to its description in a periodic lattice. This often leads to confusion because there are important differences. In particular, there is no closed equilibrium orbit in a beamline. Locations in the beamline are assigned values for the β function, as well as other lattice-like parameters. However, these parameters are *not* intrinsic properties of the beamline. They *are* intrinsic properties of a lattice. There is a special trajectory about which one can describe particle oscillations in a beamline. This trajectory is the path of a particle at the transverse center of mass of the beam.

The path the beam's center of mass takes through the beamline is "*well-defined*" assuming some simplifying assumptions are made.² The beam's center of mass trajectory is described relative to some reference trajectory. This is typically a trajectory along the beampipe center. The beampipe center is surveyed to coincide with the center of the quadrupole magnets.

The trajectory of the beam's center of mass after passing through a dipole *magnet* is the same as that of a single particle with the same initial trajectory because all the particles in the beam are bent by the same amount.

¹Taken from notes from the lecture, *An Overview of the Booster*, given by Ken Reece, 7 Aug 1991.

²In particular, all the particles being considered have the same momentum.

Figure 8.1: The LTB line.

When the beam passes through a quadrupole *magnet* off-center it is bent by a dipole and quadrupole field. Let the beampipe (and quadrupole magnet) center be $(x, x') = (0, 0)$, and B_y be the field made by a quadrupole magnet in the horizontal plane. The vertical field a particle experiences passing through the magnet is $B_y = k(x - x_o) + kx_o$. x_o is the position of the beam's center of mass. All the particles are kicked by a field of strength kx_o . This is called *quad steering*. The beam is *focussed* by a gradient k about the *trajectory* of the center of mass. In this case, the quadrupole *magnet* acts as a combination of a dipole (with strength kx_o) and a quad (with strength k) centered at the beam's center of mass.

Similarly, when the beam as a whole passes through a quadrupole it passes through a field that is proportional to the distance from the beampipe center ($B_y = kx_o$). Hence, the beam as a whole undergoes betatron oscillations about the beampipe center. For example, a positive kick by a dipole will result in movement in the negative direction at points where the phase advance from the kick is between π and 2π , and no movement at all at a point where the phase advance is exactly π . The total phase advance from the end of Linac to the Booster is about 4π . The phase advance from the start of LTB to the Booster is about 2π . It is somewhat less in the vertical than in the horizontal plane.

The beam enters the beamline with some distribution in phase space. This distribution can be enclosed by an ellipse of some particular shape, orientation, and area. The area of that ellipse is the emittance. In general, particle motion through quadrupoles *maps* points on an ellipse in (x, x') at one longitudinal location onto an ellipse at another longitudinal location. So, an ellipse in phase space at the beamline's starting point containing say 95% of the beam will be mapped into an ellipse with the same area, but different shape and orientation at a downstream location.

Unlike Heavy Ion injection, the Linac beam is injected into the full Booster acceptance. This acceptance does not change during injection, and it is described as the area inside some ellipse in phase space. The shape and orientation of this ellipse are determined by the Booster lattice, not the beam. If the emittance of the proton beam in the Booster is to be kept at a minimum, it is important that the orientation and shape of the *beam ellipse* at the injection point is as close as possible to the orientation and shape of the Booster's acceptance ellipse.

The function of LTB then is two-fold. First, it is to transport the beam

to the Booster from the Linac with a minimum of losses. Secondly, it is to alter the linac beam's phase space distribution (beam ellipse) with the quads so that it conforms to the shape and orientation of the ellipse that defines the Booster's acceptance. These conditions are met by reshaping and reorienting the linac beam ellipse to match the Booster's acceptance ellipse without compromising the beam transmission through LTB.

How much beam from the Linac actually gets stored in the Booster is dependent on the interplay between these two factors. For example, it might be that getting the beam ellipse to have the proper shape at the Booster injection point requires a beam so large at a point upstream that beam loss occurs there. It is often difficult to determine which part of the puzzle is responsible for a drop in injection efficiency. With protons (or H^- in LTB) loss monitors *are* helpful.

In practice, it is difficult to determine very accurately to what degree the beam is matched to the acceptance. Data from multiwire profiles taken with different quad strengths, together with a model of the beamline, is used to find values for the emittance and *optical parameters* of the beam.³ The beam ellipse is matched to the Booster acceptance by configuring the quadrupoles in the beamline based on this information. The BeamLineEmit program is used for this purpose.⁴

Figure 8.2 shows a *projection* of the beam *envelope* (or width) from the end of linac to the Booster. This is from the BeamLineEmit program. The quad settings from the 1997 run, together with values for the linac beam characteristics at the start of the line (found with this program in 1993) were used to arrive at this projection. The 95% unnormalized emittance of the linac beam was measured using this program to be $7.8\pi \pm 1.8\pi$ mm-mrad in the horizontal, and $16.6\pi \pm 4.4\pi$ mm-mrad in the vertical.⁵

³These are parameters which define the beam ellipse. They are called Twiss parameters. They are α , β , and γ . The general equation of an ellipse is $\frac{\epsilon}{\pi} = \gamma x^2 + 2\alpha x x' + \beta x'^2$. ϵ is the area of the ellipse. In a machine lattice β is identified with the value of the β function, $\alpha = \frac{1}{2} \frac{d\beta}{ds}$, and $\gamma = \frac{1+\alpha^2}{\beta}$. These parameters are determined by the machine lattice. In a beamline they are determined by the beamline and the initial ellipse that the beam occupies.

⁴The BeamLineEmit program and its application to the BTA line is described in: Paul Sampson, AGS/AD/Op. Note No. 43, "Transport from the Booster to the AGS", Appendix I, September 6, 1994.

⁵This corresponds to a horizontal 1 σ emittance of $1.3\pi \pm 0.3\pi$ mm-mrad, and a 1 σ vertical emittance of $2.8\pi \pm 0.7\pi$ mm-mrad. For the record, the following data fitting

The momentum spread of the Linac beam also increases the beam's size in the horizontal plane as it moves through LTB. For example, a particle with a higher momentum than the central momentum will not be bent as easily as the central momentum particle. Its '*center of mass*' trajectory will be different. This is called Dispersion, and it is like the effect that occurs in the machine. This effect is not taken into account in the beam sizes shown in figure 8.3. The dispersion also makes the particle trajectory at the injection point momentum dependent. As described in section 5.5, the closed orbit a particle follows is also momentum dependent. The relationship between these two dispersion related effects at injection can be important.

8.1.3 Tips for Improving Transmission

At the startup of a run, the settings from the last run are usually loaded. The beam profiles on MW035 and MW107 are related to the quadrupole configuration (or *optics*) of the line, as well as the shape of the linac beam in phase space. Figure 8.3 shows typical beam profiles for these multiwires.

MW035 is located just before the large bend, and MW107 towards the end of the straight section that connects the end of the bend and the Booster (see figure 8.1). MW107 generally has some problems. In particular, many of the wires often do not work. Profiles from MW107 of the quality shown in figure 8.3 are rare.

If the profiles in either multiwire appear very wide in one plane it is probably because a quad is off, or at the wrong setting. The beam is typically centered within a few millimeters on either multiwire, and in both planes.

Under ideal conditions, the beam is centered passing through the quads. This situation is ideal because the closer the beam is to the center of the beampipe, the smaller the betatron oscillations about the beampipe center will be. If these oscillations are small, the beam is less likely to scrape somewhere in the line. It is relatively straightforward to determine if the beam is nearly centered in a quad. If the beam is not centered, then changing the quad current, will change the amount of quad steering. The position on a multiwire (or bpm) downstream of the quad will change. By adjusting

and modeling techniques were used: Beamline is "LTB93.bmlm_dir", Measurement segment is "LTB.mmnt_segment_dir", Measurement strategy is "Q10MW107.mmnt_strategy_dir", and archived measurement tactics are "matched_3_12_93.mmnt_tactic_dir".

steering elements upstream of the quad, and repeating this exercise, the beam can be centered in the quad. The beam can be centered most easily by starting with quads at the upstream end of the line, and working downstream. However, working on centering the beam in the quads is destructive to the beam in the machine. Consequently, the normal running condition is far from this.

Typically, the bpms, multiwires, loss monitors, and current transformers are all used to optimize the transmission in the LTB line. As a first pass, starting from the *Store* values, the beam can be centered in the horizontal plane on MW035 with DH1, then centered on MW107 with DH2-5. In the vertical plane, the pitching magnets in HEBT (NP043 and/or NP053) can be used to center it on MW035, and DV083 on MW107.

Figure 8.4 illustrates the effect of a dipole kick on the beam's trajectory using a model of the beamline in the BeamLineEmit program.⁶ In this case, a kick in both planes is given at the start of HEBT. The beam undergoes betatron oscillations about the beampipe center because of the steering effect of the quadrupoles. Notice for example, that a positive kick at the start of HEBT results in movement in the negative direction on MW035 in both planes. On MW107 this kick moves the beam in the positive direction in the horizontal, and the negative direction in the vertical.

In order to center the beam, and also make its trajectory parallel to the reference trajectory at one point (at s^*), two upstream dipoles are needed. It's preferable if one of those magnets is close to that point ($\psi(s^* - s_1) \approx 0$ where s_1 is its location), and the other magnet has a phase advance which turns its kick into a pure position change at the point ($\psi(s^* - s_2) \approx \frac{\pi}{2}$, where s_2 is its location). In figure 8.4, the horizontal kick at the start of HEBT results in little position change at DH3 ($\psi(33m - 0m) \approx n\pi$), and a large change near QH1 ($\psi(24m - 0m) \approx \frac{n\pi}{2}$). DH3 and QH1 are about 10 m (or 30 ft) apart and the phase advance between them is about $\frac{\pi}{2}$. Consequently, to move the beam at a particular point, a rule of thumb is to use a dipole about 30 feet upstream of it. DH2-5, since it is composed of several magnets, each at different phase advances from some upstream point, cannot be thought of in this relatively simple way.

⁶The BeamLineEmit program calls another program called MAD. MAD is used to model the beamline given as inputs the initial beam parameters, and the configuration of magnets in the beamline.

In practice, the *right* magnets to use are typically found by trial and error. That is, look at the position on the bpm and change an upstream dipole. If it has a large effect on the position there relative to the position at other bpms, then that dipole probably effects the position at the bpm more than the angle. Alternately, a dipole that has little effect on the position at the bpm in question, but a significant effect on it at other bpms, probably changes the angle more than the position at the bpm.

There are seven bpms, more or less evenly spaced in each plane in LTB (see section 3.1.3 and figure 3.3). Their locations in the beamline are shown at the bottom of figure 8.4. With this number of measurement locations, it is easy to get an idea of the effect a dipole change has on the overall trajectory. In particular, one can see the effect of the quadrupoles on the trajectory after a change to an upstream dipole that is reminiscent of the effect seen in figure 8.4 where the hypothetical kick occurs at the start of HEBT.

The final goal here is *not* to get the beam centered in the beamline. One of the goals *is* to maximize the efficiency of the transmission through the line. Centering the beam is correlated with transmission efficiency. For this reason, the diagnostics (bpms and multiwires) are used to help get the beam close to centered. However, even if the bpms and multiwires worked “*perfectly*”, and the beamline was surveyed “*perfectly*”, that does not mean the beam is centered at places other than the bpms and multiwires. In practice, it is impossible to center the beam in the beampipe at all points. But *centering* the beam in the beampipe is not the goal, it is only the first step in an approach for optimizing the transmission. If the beam happens to be centered when the transmission is optimized that's OK, but it is also largely irrelevant.

There are current transformers at 11 and 100 feet in LTB (called XF011 and XF100).⁷ There are analog signals from these current transformers available through Xbar. They are designed to have a similar response to beam (a similar calibration). If there is perfect transmission from 10 feet (before the big bend) to 100 feet (well after the big bend) then the signals from these two transformers should overlap. If there is less than perfect transmission, then the XF100 signal will be smaller. Once the beam is well centered on the multiwires and bpms, steering changes can be made while looking at these two signals. If the transmission improves, then the new steering, though the

⁷see section 3.1.1

beam may no longer appear to be centered, is usually preferable.

Figure 8.5 shows analog signals for XF011 and XF100 overlayed on a scope.⁸ Notice that the time dependence of the current measured on both transformers is similar, but the amplitude of one of the traces is slightly less than the other. The lower amplitude trace is XF100, and the difference in amplitudes is taken to be due to beam loss between the two transformers. A similar measurement can be made using the injection current transformer in the Booster, and an LTB transformer. It is located at C6, just downstream of the injection point. However, this measurement is more involved because it requires that the beam does not circulate in the Booster. It has the advantage of including all of LTB past 11 feet in the measurement.

Another tool which can be used to address “*the bottomline*” is the loss monitor system (see section 3.1.2). Certainly, the less loss in LTB for a given intensity linac beam, the better the transmission. In general, one expects higher losses to coincide with poorer transmission. The current transformers and the loss monitors can both be used to *gauge* the transmission. Having two different ways of measuring something gives one more confidence in both measurements if they are consistent (or less if they are not). Although the current transformer measurement may be more reliable because it is a more direct way of measuring the quantity of concern (transmission), the loss monitors allow one to better isolate the location of a loss.

Figure 3.4 is a display of loss monitor data from the BoosterLossMonitor program. There are eight loss monitors in LTB. They are designated by their distance from DH1 in feet. There are also analog signals available for each of these loss monitors.

An approach often used when working with loss monitors involves moving the loss *downstream*. At some point the LTB loss would then become a Booster injection loss. If losses are high near DH1 (ltb008). One would try to adjust steering to reduce this loss. In many cases, such a loss will be reduced, but losses further downstream will probably increase (say at ltb038). Then one concentrates on those losses, moving them further downstream, and so on...

It is also helpful to determine whether the loss is due to a vertical or horizontal aperture. If the loss is insensitive to changes in upstream horizontal dipoles, but is sensitive to changes in vertical dipoles, then it is probably

⁸Taken from D. Gassner and J. Reich, Booster Protons 3, pg. 73, FY 1994

Figure 8.5: Scope setup for an LTB transmission efficiency measurement. Analog signals from XF011 and XF100 are overlayed. The ratio of their amplitudes is the efficiency of the transport from 11 feet to 100 feet in LTB. This can be measured on the scope by measuring the voltage level in both transformers at the same time in the pulse. In this picture the level of XF011 (the higher amplitude trace) is being measured. It is 4.51V. Similarly, the level of XF100 can be measured. The transmission efficiency in this case is about 90-95%.

due to a vertical aperture problem. Suspecting that it is a vertical aperture problem, one can concentrate on improving the vertical trajectory.

In practice, there is a significant loss (as seen on ltb008) just downstream of DH1 which is difficult to remove. This loss, if it is associated with losses in the transport of beam from linac to the Booster, should have to do with how the beam goes through DH1, and its size near the loss monitor. According to figure 8.2, the beam size is rather large in the vertical plane upstream of DH1.⁹ The HEBT steerers could be used to try to reduce this loss, at least moving it downstream. But this approach is typically more or less unsuccessful. If the beam is too large steering alone will not fix the problem. Quadrupoles need to be changed. In this case, adjusting the quadrupoles in HEBT might effect the loss.

A rough calibration of the loss monitors at injection energy was done in 1997. The result is that about 1200 loss monitor counts is equivalent to a loss of 1 Tp of beam ('Tp' stands for Teraparticle= $1 \cdot 10^{12}$ particles).¹⁰ The percentage of the linac beam lost in LTB during this calibration was in the 5% range. This is probably a typical value.

If the beam steering is optimized, but beam loss still occurs, it may be because the optics of the beamline need to be adjusted. For example, the vertical beam size around DH5 as shown in figure 8.2 may be larger than the beampipe. Adjusting the quadrupoles upstream of it can make the vertical beam size smaller here. However, such an adjustment may cause the horizontal beam size to increase to such an extent that it will cause loss somewhere else.

So much is effected when a quadrupole is changed that it is difficult to know what the beam is really sensitive to. For example, a change in a quadrupole's current, may also cause the trajectory to change due to quad steering. If the injection efficiency improves, it may be because the steering somewhere has changed, not because the transmission is better because the beam is smaller at DH5. This seems to me to be a primary reason why it can be important to center the beam going through the quadrupoles. The more centered the beam is in the quadrupoles, the more independent steering and optics changes become, and the easier it is to know what effect quad changes

⁹Figure 8.2 should not be relied on too heavily. It is only a model of the beamline which may or may not reflect reality on any given day.

¹⁰K. Zeno, HEP Setup Book II, FY 1997, pgs. 6-11

have.

A change in a quadrupole also effects the matching to the Booster. Making a quadrupole change may improve the transmission. But, improving the transmission isn't much good if the beam can no longer be injected into the machine due to some injection mismatch. In any event, because quadrupoles change so much about the setup its not a good idea to change them without closely observing their effect on injection.

8.2 Injection

8.2.1 The Foil as a Proton *Source*

Proton injection is very different from Heavy ion injection. The beam enters the Booster as one particle, H^- , and is stored in the machine as another, H^+ . The H^- beam is stripped of its electrons by passing through a *foil* located in the C5 half-cell in the Booster. This foil is located *inside* the Booster's proton acceptance, and acts as a *source* for protons. Protons are created in the Booster at this source, and hence do not have to be put into the Booster acceptance using phase space gymnastics such as those used for Heavy Ions. The linac accelerates H^- so that injection can occur this way.

Recall that the difficulty with Heavy Ion injection was that a particle outside the Booster acceptance could not be put into the acceptance without changing that acceptance. In the case of protons, the *particle* is changed, not the acceptance. H^- beam entering the Booster passes through a hole in the outside of the C5 bending magnet. It is bent by the field inside the magnet in the opposite sense to that of a proton until it is roughly tangent to the E.O. of the protons that are created at the foil. The foil is located at this point, which is just downstream of the C5 main magnet. The beam passes through it, and finds itself on, or close to the E.O., as a proton.

Everytime beam passes through the foil it scatters and its momentum spread increases. If the beam circulating in the Booster were to pass through the foil on every turn its transverse and longitudinal emittances would both become very large, causing beam loss, and other problems. Consequently, the foil is placed away from the beampipe center, where the circulating beam normally resides, and towards the inside. During injection the orbit is distorted so that it passes through the foil, in order to inject the beam created

at the foil near to its E.O. The betatron oscillations about the orbit are thereby smaller than they would be if the E.O. ran through the beampipe center at the foil. Such large oscillations would cause beam loss and a larger beam. After the beam pulse from Linac has ended, the orbit is put in to its undistorted state in a controlled way, and the beam circulates near the middle of the beampipe, away from the foil.

The magnets used to change the acceptance during Heavy Ion injection are used to move the orbit into the foil during proton injection and keep it there until injection is over. Their function in each case is totally different, though they are the same (the C1, C3, C7, and D1 ferrite magnets) and they are controlled by the same software (BoosterOrbitControl). In each case, the bump they produce is called the *injection bump*. For Heavy Ions, the injection bump moves the E.O. to the outside at the injection point (the C3 inflector) and needs to fall during injection in order for beam to be stored. With protons, the bump moves the E.O. to the inside at the injection point (the C5 foil) and needs to keep the orbit passing through (or near) the foil during injection in order for beam to be stored optimally.

It turns out that during *high intensity* proton running the charge density of the beam can be high enough to significantly change the betatron tunes of the particles within the beam. This leads to problems with resonances that can cause beam loss. For that reason, the beam is intentionally made larger so that it will be less dense, and thereby less susceptible to these so-called *space charge* effects.

One way the beam density is reduced is by injecting the beam while the injection bump is collapsing. The collapsing injection bump causes the E.O. to move outward at the foil's location. The protons continue coming out of the foil at the same (x, x') relative to the beampipe center. As the bump collapses, the E.O. gets further away from the incoming beam at the foil, and the betatron oscillations about the E.O increase. The resulting stored beam is spread out over a larger area in the transverse plane, and is consequently less dense. The behavior of the bump during this time has some superficial similarity to the Heavy Ion situation, where beam is injected on the bump's falling edge. However, the reason in this case for injecting the beam while the bump is collapsing is totally different.

8.2.2 The Linac Pulse

The energy of the Linac beam is about 200 MeV ($p \approx 645 MeV/c$, $\beta = \frac{v}{c} = 0.56$). It is delivered in a pulse which can be up to about $400 \mu s$ long. The linac current is about 30-40 mA. This is about one thousand times higher than the Tandem pulse current. The momentum spread of the linac beam, $\left\langle \frac{\Delta p}{p} \right\rangle_{fullwidth}$, is on the order of 10^{-2} .¹¹ The momentum spread of the Tandem beam is about one tenth of this.¹²

The size of the linac beam as it hits the foil is determined by its emittance, and the orientation of its phase space ellipse at the foil. The momentum spread also increases the horizontal beam size at the foil by a finite amount. Taking these elements into account and using the same model of LTB as in figures 8.2 and 8.4, a very rough estimate of the size of the linac beam at the foil can be made. This model gives a horizontal half-width of the linac beam of 12 mm, and a vertical half-width of 9 mm.¹³

The pulse that comes out of the Linac source is chopped into an approximately square pulse (as in figure 8.5). The *Slow Chopper* does this. It is located in *LEBT*, which is the line the beam from the linac source passes through before it enters linac.¹⁴ Initially, the slow chopper is in a state which does not allow beam to pass through it without being deflected. When the event LPI.CHOP_ON_DLY.RT occurs, the chopper allows beam to pass through it, and into linac, unhindered. This event is a microsecond resolution delay from Peaker, typically set to about $2500 \mu s$.¹⁵ The width of the pulse allowed to pass through the slow chopper is set by the SLD, LPI.CHOP_WIDTH. Its value is in microseconds.¹⁶

¹¹Keith Zeno, AGS/AD/Op. Note No. 39, “Strategies for Measuring and Estsblishing Values for Proton Injection Parameters in the AGS Booster”, pg 11, December 22, 1992.

¹²L. Ahrens, HIP setup Fe/Au Booster/AGS, Book I, pg. 22, FY '98. This measurement was done with Fe^{+10} .

¹³The formula used to calculate this is, $\langle \Delta z \rangle^2 = \beta \frac{\epsilon}{\pi} + D \left\langle \frac{\Delta p}{p} \right\rangle_{halfwidth}^2$. In the horizontal plane, $\beta = 13.0m$, the dispersion $D = 1.15m$, and $\epsilon = 7.8\pi mm mrad$. In the vertical plane, $\beta = 4.6m$, $D = 0m$, and $\epsilon = 16.6\pi mm mrad$. $\langle \Delta z \rangle$ is the half width in either plane.

¹⁴LEBT stands for “Low Energy Beam Transport”.

¹⁵This SLD is not a Booster Real timeline event as the suffix, ‘.RT’, might imply. It is a *Linac* event, as the prefix ‘LPI’ implies. ‘LPI’ stands for, *Linac PreInjector*.

¹⁶Detailed information on injection timing for protons is given in “AGS/AD/Op. Note No. 38, “Overview of Booster Timing”, pgs. 13-17C, by K. Zeno et. al., February 17,

The reason why the slow chopper is called '*slow*', is because there is another chopper in LEBT, downstream of it, which chops the *macropulse* that comes out of the slow chopper into many *micropulses*. These *micropulses* are contained within the envelope of the slow chopper macropulse. The *fast chopper*, which makes the micropulses, is synchronized to the drive signal from the low level Booster Rf system. This signal also drives the Rf in the cavities. The fast chopper chops out parts of the macropulse at the Booster Rf frequency, leaving holes and micropulses. Each micropulse takes the same amount of time to pass through the linac, and the time each micropulse enters the linac is synchronized to the Booster Rf. Consequently, the micropulses that emerge from the linac, and enter the Booster, are synchronized to the Booster Rf.

Figure 8.6 is a scope picture of the fast chopped beam as it appears in LTB.¹⁷ The microstructure is visible here on an analog signal derived from the sum of the voltage on all 4 plates of the bpm at 19 feet in LTB (bpm019). This signal is available on Xbar. Compare this picture to figure 8.5 where the beam is *not* fast chopped. This linac pulse is also much wider than the one in figure 8.5. Here the pulse width is about 320 μs , whereas in figure 8.5 it is about 40 μs .

Measuring the current (or intensity) of the fast chopped beam in LTB, or Linac, is difficult. Transformers respond differently to it, and their response is not generally known. The width of the micropulses vary as well. This makes calibration an even more complicated issue. In addition, efficiency measurements in LTB, like those associated with figure 8.5, are typically made with unchopped beam.¹⁸ The signals are cleaner. The assumption made in generalizing this measurement to fast chopped beam, is that the chopping process does not have a significant effect on transmission.

However, the chopper does have the potential of increasing the beam emittance. The beam is chopped out when it passes through a transverse voltage in the fast chopper. When the beam is meant to pass through the chopper undeflected, the voltage is brought to zero. Something like this certainly has the potential of effecting the emittance of the micropulse, particularly at its edges where this voltage must be changing. The multiwires in LTB do not

1993.

¹⁷Taken from K. Zeno, Booster Protons 4, pg. 90, FY 1994.

¹⁸Beam that is not fast chopped, but is slow chopped, is typically called unchopped.

19-Jun-94
1:15:34

Fast Chopped Beam in LTB

Figure 8.6: Fast chopped beam in LTB. The top trace is BXI.LTB.-BPM019.4SUM. The bottom signal is the current in the D1 injection bump magnet (BXI.KDHD1). The linac pulse width is 320 μs .

show anything obvious. The chopper could also potentially change the steering, though this effect has not been observed. The fast chopper is turned on and off with the SLD, LIN.CHOPPER_ON_OFF.

The current coming out of Linac is measured with an analog signal from the tank 9 current transformer, which is available in Xbar. Its calibration is 10 mA/V with unchopped beam, unterminated. Its response is too slow to see the fast chopped beam. The *Booster Input* scaler at each console in MCR is typically calibrated for the fast chopped beam. The current transformer that is used for this digital readout is subject to change, and always a bit of a mystery to me. It can be the current transformer from tank 9, or one in HEBT or LTB.

The strength of detrimental space charge forces decreases as the momentum is increased. The more time the beam spends at lower momentum, the more time the stronger space charge forces have to act on the beam, and the greater their overall effect. For this reason, it is advantageous to accelerate the beam at a high rate away from injection energy as soon as the beam is injected. The rate of momentum increase is proportional to the $\frac{dB}{dt}$ of the Main Magnet. The Main Magnet cannot change its $\frac{dB}{dt}$ very quickly. Consequently, the fastest way to get the beam to a higher momentum is to inject the beam while the Main Magnet field is rising rapidly.

For a given Rf voltage, the width of the stable phase region of the Rf wave shrinks as the acceleration rate increases.¹⁹ This is related to the fact that the *average* kick that a particle receives must be enough to keep its energy increasing at the required rate. As a result, a particle's average phase on the Rf wave gets closer to the Rf wave's peak amplitude, and the spread in phase about which the motion is stable is necessarily reduced.

The micropulses enter the Booster with a fixed phase relationship to the Rf. That is, the time that they pass through an Rf cavity is correlated with the phase of the Rf. Consequently, the micropulse can be described in terms of the phase of the Rf wave. Its center passes through an Rf cavity when the phase of the low level Rf is at some value, ϕ_o . Also, it enters the Rf cavity when the Rf phase is ϕ_1 , and exits it when it is ϕ_2 . The width of the micropulse in phase is $\Delta\phi = \phi_2 - \phi_1$. The phase of the micropulse, ϕ_o , can be shifted relative to the phase of the Rf. It is shifted by adjusting the phase

¹⁹Acceleration rate is somewhat of a misnomer. More precisely, the rate at which the beam's *energy* increases is the relevant parameter, not its velocity.

of the fast chopper using the SLD, LPI.CHOPPER_PHASE. For example, this “*Chopper Phase*” can be adjusted so that the micropulse phase coincides with the center of the stable phase region, which is the synchronous phase.

The fast chopper’s primary purpose is to chop out the beam that lies outside the stable region. This reduces the radiation due to beam loss by disposing of the beam that cannot be accelerated in the Booster while its energy is still very low (in LEBT). The fast chopper is useful for other reasons which will be described later.

$\Delta\phi$ has to be smaller than the width of the stable phase region. The width of the micropulses, $\Delta\phi$ is changed using a software application called *Chopper*. A parameter in the program called *Bunch Width* is used to adjust it. It is called the bunch width because the micropulse width is closely related to the width of the bunches that circulate in the machine.

8.2.3 Bunches, Turns, and Micropulses

In the simplest case, the fast chopper chops the macropulse into a series of micropulses at the Rf frequency. The Rf frequency is an integer multiple of the beam’s revolution frequency. For every 2π advance in phase on the Rf wave a micropulse is generated. That micropulse has the phase ϕ_o within that Rf cycle of phase extent 2π . During each revolution period there are h micropulses injected into the Booster (h is the Rf harmonic number). Each of these is injected at the same phase on the Rf wave, ϕ_o , but is associated with a different individual Rf wave cycle.

Assume for the sake of simplicity that there is only one Rf cavity. The first micropulse injected passes through this Rf cavity during the ‘*first*’ Rf cycle at ϕ_o . The next micropulse passes through it during the ‘*second*’ Rf cycle at ϕ_o . After the first micropulse has made one revolution, or *turn*, it again passes through the cavity on the $h + 1$ Rf cycle at ϕ_o . The $h + 1$ microparticle passes through this cavity on the $h + 1$ cycle at ϕ_o as well. Therefore, micropulses 1 and $h + 1$ are in the same place, but have made a different number of turns around the Booster. Subsequent micropulses continue to be injected, and overlay.

The micropulses are classified by what revolution, or turn, the beam initially injected is making when they are injected. Micropulses 1 through h are called the first turn, $h + 1$ through $2h$ are the second turn, and so on. The beam within each bucket is called a bunch, and the bunches are numbered

as well. The first bunch is the bunch that contains the micropulses: $1, h + 1, 2h + 1, \dots, nh + 1$. The second bunch, assuming there is more than one, contains the micropulses: $2, h + 2, 2h + 2, \dots, nh + 2$.²⁰ During the past few years, h has been two. When the Booster first started running, h was three. There are indications that it will change again. Since I am most familiar with the $h = 2$ setup, this is what I'll be describing unless otherwise noted.

The revolution period at injection is $1.2 \mu s$. A typical high intensity slow chopper pulse is about $300 \mu s$ long. Consequently, the typical number of turns injected, n , is about $\frac{300 \mu s}{1.2 \mu s} = 250$. Actually, the fast chopper does not start to work until $10 \mu s$ after the slow chopper starts letting beam through at LPI.CHOP_ON_DLY.RT. Before this time the beam that passes through the fast chopper is deflected and not injected into Linac. As a result, the length of the Linac macropulse is $10 \mu s$ shorter than the slow chopper width when the fast chopper is on.

8.2.4 The Chopper Program

The actual situation is considerably more complicated. The micropulses will only stack up directly on top of each other if they are injected at the synchronous phase, and their energy (or revolution frequency) is exactly right to keep them at that phase on subsequent revolutions. Additionally, the Booster field is increasing rapidly during injection (at $30 g/ms$) and the Linac energy remains constant. The Rf voltage is typically at its maximum value ($\approx 100 kV/turn$) during this time.

In an ideal simple case, the Rf frequency tracks the B field, and *Beam-Control* is not enabled. In this situation, the Rf frequency increases during the injection process as the field increases. The first turns of beam injected are accelerated throughout injection. The last few turns of beam injected enter the Booster still at the Linac energy, which will be less than the *present* energy of the beam from the first few turns that has been accelerating. The result is an increase in the energy spread of the Booster beam over that of the Linac beam that has been injected.

The chopper program sends instructions to the fast chopper to chop the micropulses so that they fit into an Rf bucket. Since the Rf frequency is

²⁰I invented the term *micropulse* in order to describe turns and bunches. You won't hear people using this word in the control room.

increasing, the central energy of the Rf bucket is changing relative to the energy of the beam being injected. Parameters such as the rate of Rf frequency change are input into this program. They are used to calculate the instructions to be sent to the fast chopper. The Chopper program main window is shown in figure 8.7.

The amount of information in this window is rather overwhelming. Fortunately, there are only a handful of items here that are typically of any interest to operations. One of those items is called the bunch-turn matrix. The bunch-turn matrix is displayed here as a matrix of ones and zeros in the middle of the window. Each element in this matrix refers to a specific micropulse. A micropulse is identified by its row and column number. The row number is the same as the bunch number. The column number is the same as the turn number. There are actually only two rows with six hundred columns shown here. For example, row 2, column 110 is actually the tenth column from the left in the fourth *line*. This corresponds to the micropulse injected into bunch 2 on the 110th turn.

The bunch-turn matrix determines whether a micropulse will pass through the fast chopper and into the linac. If the micropulse is associated with a '1', then the fast chopper 'creates' it. If it is associated with a '0' then no micropulse is created. In this case, the matrix is full of zeros after the 250th turn. This means that beam is not injected after the 250th turn. Also, there are zeros every tenth turn. Therefore, there is no beam injected into the machine on these turns. This pattern of regularly occurring zeros within a field of ones is called a *Sieve*. The percentage of turns with beam is set in the *beam options* pull down menu under *Sieve Options*. In this case, the percentage is 90%.

The number of turns injected is set in the *Inj turns* field on the right of the page under *New Physics Parameters*. In this case, it is set to 250. This corresponds to how far out 'ones' extend in the bunch-turn matrix. The fast chopper can be used to inject beam into only one of the bunches by sending zeros to all the turns for one of the rows. This is accomplished by selecting '1' under '*Number of Bunches*' on the right.

The '*bunch width*' parameter sets the width of the micropulses. The micropulse width can vary throughout the injection process so that the resulting micropulses conform to the contours of an idealized Rf bucket. This bucket is displayed in the phase space diagram at the bottom. The x axis is phase, the y axis is deviation from the synchronous energy. The solid region inside

the bucket is the bunch that results from the input parameters. The width of this bunch in phase is set with the *bunch width* parameter. This phase space diagram is highly idealized.

The slow chopper width can be set through this program as well as with the SLD. Obviously, the fast chopper can only produce micropulses if beam is passing through it. If the micropulse from the slow chopper ends before the requested number of turns is produced, then fewer turns than the requested amount will be produced. Consequently, to produce n turns, the slow chopper width must be at least $(1.2\mu s)n + 10\mu s$.

8.2.5 Stripping the Incoming Beam

The stripping foils are attached to a wheel that is inside the Booster vacuum enclosure, and driven by a motor. There are six places around this wheel where a foil can be held. An SLD, BXO.CARBON_FOIL, controls which of the six possible locations is inserted into the beampipe. All the foils are made of the same material. They differ by how far inside from the beampipe center their outside edge is when inserted. The Description pop-up window for BXO.CARBON_FOIL contains a list of foils and their edge positions relative to beampipe center. Typically, a foil whose outside edge is 1 inch to the inside from the beampipe center is used (foil 3 or 4). There is no foil attached to the wheel at location '1'. BXO.CARBON_FOIL is set to '1' during Heavy Ion operation.

A measurement of the horizontal position at C6, is used to get an idea of the incoming beam's position at the foil. The position at C6 is measured using analog signals from the C6 PUE (if they are available). This PUE is located about 0.6m downstream of the foil. For all the particles in a beam with a 15 mm (0.6 in) half-width to pass through the foil, and for a foil with an edge 1 inch from the beampipe center, the incoming beam's position at the foil would have to be at least $1.0in + 0.6in = 1.6in$ (41 mm) to the inside. Since there are no magnetic elements between the PUE and the foil, the beam's positions relative to the beampipe center, at the two locations are related by the angle of the incoming beam. That is, $x_{C6} = x_{foil} + l \cdot x'_{foil}$, where l is the distance between the foil and the C6 PUE. Typically, since the PUE is 'close' to the foil, and x_{foil} has a relatively large negative value, $|x'_{foil} \cdot l| \ll |x_{foil}|$, and $x_{foil} \approx x_{C6}$.

Figure 8.8: An injection trajectory measurement using the C6 PUE A and B plates and a half-turn of beam. The position in millimeters is found by dividing the difference between the two voltages by their sum and multiplying by 100.

To measure the position at C6, the chopper program is loaded with a bunch-turn matrix that requests beam on only one turn, and in only 1 bunch. This configuration is called a 'half-turn'. Hardwired PUE signals may be available either as signals from each of the two plates (called the A and B plates), or as sum and difference signals from the plates ($V_A - V_B$, and $V_A + V_B$).²¹ In figure 8.8, the C6 PUE A and B plate signals for a half-turn are shown.²² The scope may be triggered using a microsecond resolution delay from Peaker. Here, the C6 position is measured to be -37 mm.

When the C6 PUE is working, and analog signals are available from it, this measurement is taken every day during normal running as part of the *Morning Numbers*. A Typical value for the C6 position, during high intensity

²¹A is the outer plate, B is the inner plate.

²²Zeno et. al., Booster Protons 2, pg. 16, 1994

running, is $-50\text{mm} \pm 10\text{mm}$. The C6 position can be adjusted using DH115, the last magnet in LTB. Increasing DH115 moves the beam to the inside. It also makes its angle at the foil more negative. The *optimum* position will vary depending upon how other parameters are set, and, in particular, the injected beam intensity.

The *bottomline* is to get as much of the incoming beam to pass through the foil as possible, but not move the incoming beam so far to the inside that you run the risk of hitting the inside of the beampipe with it. The foil is wide enough vertically, that getting the beam to hit the foil in the vertical plane is not generally regarded as a significant problem. However, the aperture of the beampipe is much smaller vertically than horizontally, and the vertical trajectory is critical.

The C6 PUE plates are typically hit by the electrons that are stripped off the injected beam during the injection process. Because of this, the foil is removed when this position measurement is taken. When first looking at injection (with the foil in), the first sign of beam in the Booster is often the voltage induced by these stripped electrons hitting the plates of this PUE.

8.2.6 Low Intensity versus High Intensity

The details of the injection process are markedly different for '*high*' and '*low*' intensity. Even though high intensity is the rule, the description of the low intensity setup is a good starting point. The high intensity setup grows out of the low intensity setup, and relies even more than usual on trial and error.

The injected beam intensity can be adjusted by using the Sieve option in the Chopper program. To a good approximation, one is only changing the '*effective*' linac current by adjusting '*the Sieve*'. Consequently, the Booster setup should not have to be changed to maintain the same efficiencies if the injected intensity is increased this way, unless the optimal setup is '*intensity dependent*'. One invariably finds that, as the injected intensity is increased using this method, the Booster setup with the highest efficiency begins to deviate markedly from the initial *low intensity* setup.

The deviation is generally attributed to the effects of the beam's charge on itself, or '*space charge effects*'. In order to understand where the machine *is* in the high intensity case, it helps to have a reference setup from which deviations can be measured. The low intensity setup, besides being important in itself, can serve as a sort of reference. Additionally, the machine is usually

started up with low intensity. By the way, high intensity effects are not obvious with less than about 5 Tp injected.

As with the Heavy Ion description, I'll first concentrate on the transverse aspects of injection.

High Intensity Archives in a Low Intensity Setup

In the absence of space charge effects, the intent is to keep the emittance of the beam as small as possible. Accordingly, the injection bump is setup so that the E.O. coincides with the trajectory of the incoming beam at the foil. A smaller beam is more likely to spiral, and be injected into the acceptance efficiently. If the undistorted E.O. runs through the beampipe center, and the beam is injected so that its outer edge is just to the inside of the foil's outer edge (one inch from center), then the injection bump should bump the orbit about 40 mm to the inside to accomplish this.

However, because the beam is injected directly into the Booster's acceptance, the injection bump is not nearly as critical as it is in the Heavy Ion case. In fact, it is possible, though not advisable, to successfully store a little beam without it. Consequently, even though low intensity is being considered, the high intensity Store archive for the injection bump is typically more than adequate as a starting point.

There are, however, other high intensity Store archives that are decidedly *not* adequate as a starting point. In particular, the TuneControl functions invariably need to be modified. At high intensity, the beam effectively defocusses itself, and the betatron tunes of a particle for the same quadrupole magnet field are lower. Additionally, how much the tunes are depressed from the low intensity tunes depends on the details of the particle's motion through the other particles. For a beam with a gaussian distribution, particles which undergo relatively small betatron oscillations have a relatively large tune depression.²³ The range of particle tunes within the beam due to space charge effects can be 0.5 or more.

With such a large range of tunes within the beam, it is impossible to keep all the particles from passing through stopbands. Naturally, it is important to keep as much beam as possible away from the strongest stopbands in order to keep losses at a minimum. The lowest order stopbands (the half

²³See Edwards and Syphers, pgs.172-178 for more on Space Charge effects.

integer, $\nu = 4\frac{1}{2}$, and the third integer, $\nu = 4\frac{2}{3}$) are typically regarded as being the stopbands which have the potential to produce the most beam loss. Empirically, it has been found that it is necessary to move the *requested* tunes higher in and around injection, with respect to their low intensity values, to minimize losses (or, equivalently, increase injection efficiency). In particular, it has been found that the vertical requested tune needs to be moved much more than the horizontal from its low intensity value. In fact, the *requested* vertical tune is typically at, or slightly above 5.0 during high intensity running.

As the tune is raised, and approaches an integer, the amplitude of orbit distortions increases (see Section 5.7). The actual tunes of particles in a high intensity beam are less for the same 'external' focussing field than in the low intensity case. So, even though the requested tunes for the high intensity case are near, or at an integer, the actual tunes are not, and the orbit is not as bad as it would be if there were no space charge effects. Conversely, when the high intensity tune functions are loaded, but the beam is low intensity, the orbit distortions are worse than in the high intensity case because the *requested* tunes are close to the *actual* tunes. The orbit can be so bad, that the beam will not circulate.

Not surprisingly, when setting up, or running, the machine at low intensity the tunes should be reduced from their high intensity values. This is most important early in the cycle, and particularly at injection, when space charge effects are the strongest. Figure 5.4 shows a typical high intensity tune function. In this cycle, injection occurs around 10 ms from BT0. At 10 ms, $\nu_x \approx 4.85$, and $\nu_y \approx 5.01$. For low intensity running, ν_x is normally about 4.75, and ν_y to about 4.80 at injection, and kept at or below these values throughout the cycle.

Since the orbit improves as the tunes move away from the integer, it is easier to get the beam to spiral if the tunes are lower. Alternately, if the tunes have to be lowered more than usual to get the beam to spiral, then there may be a problem with the orbit correctors. If this is the case, it is a good idea to check that the archive for the orbit correctors is a good one, and that they are on and functioning properly. The orbit at high intensity, particularly in the vertical, is typically tuned extensively so that the tune can be raised as high as possible, and is usually a well corrected orbit. As a result, the high intensity Store archive for the orbit correctors is usually a good choice for use at low intensity.

8.2.7 The E.O. and the Injection Trajectory

The Effect of the Rising Magnetic Field on the E.O.

The bottom trace in figure 8.6 is the current running through the D1 injection bump magnet, or *injection kicker*. Notice that the beam is injected while this current is decreasing. This is in apparent contradiction to what was said in section 8.2. In that section, it was stated that the injection bump is used to keep the orbit passing through the foil during injection. More specifically, in the case of low intensity, its purpose is to keep the E.O. tangent to the incoming beam so that it is injected onto the E.O., and thus kept as small as possible.

So, why is the injection bump decreasing in amplitude during injection? This injection bump '*ramp*' was first used to compensate for the increasing Main Magnet field during injection. Due to the increasing field, the E.O. of the **linac energy** proton beam is moving to the inside during injection. To keep the E.O. at the same position at the foil, the amplitude of the injection bump has to decrease.²⁴

Adjusting the bump is only one way to change the location of the E.O. at the foil. Its location at the foil can also be '*fine tuned*' by changing the injection field, which for a given Main Magnet function, is set by Peaker. However, adjusting Peaker moves the beam radially everywhere, changes the beam's revolution frequency, and moves the time that injection occurs with respect to real time functions such as the Tune, Orbit Correctors, Chromaticity, and Main Magnet. The advantage of changing the bump is that it doesn't change everything else. The advantage of Peaker is that it's quick.

If you change Peaker, and efficiency gets better, you generally don't know why it gets better. If you change the bump and it gets better, at least you have a clue. This is a general rule: '*Knobs*' which change many parameters, like Peaker, can fix a lot of things, or screw up a lot of things. They are the most powerful, and often allow the kind of quick fix that is so *critical* during normal operations. However, when you use them, they give you few if any clues, as to what the problem was, how it was fixed, or how the setup has changed.

What the optimum value for Peaker is depends on how everything else

²⁴The trajectory of the incoming beam, even though it passes through part of the C5 main magnet, whose field is increasing, changes relatively little during this time.

is set up. Within some range, it's not clear that one value is necessarily better than another. One value may be better for a particular setup of other parameters, but not as good as another value with a different setup. Yet, since the radius of the E.O. varies during the injection process, one might think that its average value, or value at the middle of the injection process, would be constrained by the beampipe's width, to be near the beampipe center.

Let's take a look at how strong this constraint might be. The injected beam begins to enter the Booster a few microseconds after the slow chopper starts letting beam pass through LEBT. This happens at LPI.CHOP_ON_DLY.RT, which is normally set to about $2500\mu s$ after Peaker. Peaker also determines what the radius of the initial E.O. is around the machine. If injection is $300\mu s$ long, the field will change by, $30g/ms \cdot 0.3ms = 9g$ during injection. A change of 1 gauss in the injection field, changes the radius of the E.O. by about 1 mm, so this amounts to a change of about 1 cm. The available horizontal aperture is about 4 inches, or 10 cm. So, the constraint on the E.O. due to the beampipe's width may not be very important.

Recall that the Rf is on at its maximum value. Once a turn of beam is injected, the Rf accelerates it one way or the other. The effect of the Rf may counteract, or contribute to any overall radial movement. In other words, it is not clear what kind of value for Peaker *'makes sense'*.

At the start of a run, after the Store archive is loaded, and the tunes are lowered to accomodate lower intensity, a scope is setup as in figure 8.6, but with the addition of a Booster PUE and/or a current transformer signal. Then Peaker is scanned up and down while looking for circulating beam on the PUE or transformer. If the beam circulates, the injection trajectory is tuned to improve the survival, and Peaker is adjusted further. This is done a few times, and Peaker is left at the value which provides the best injection efficiency.

Coherent Betatron Oscillations

There is a way to measure how far the incoming trajectory is from its E.O. A half-turn of beam can be placed at any time within the macropulse from the slow chopper. Say, the normal linac pulse is $300\mu s$ long, resembling figure 8.6. Now a half-turn is placed in the middle of this pulse, this is now the only injected beam. There are usually signals from the C7, C8, and D2 PUEs

(recall that the C6 PUE is not usable when the foil is in). The half-turn can be watched on each revolution as it passes through a particular PUE. When the half-turn is first injected, the particles within it act together, and it undergoes betatron oscillations about its E.O. This is not unlike the betatron oscillations of the center of mass particle moving through a beamline such as LTB.

However, in this case, the beam comes around to the same measuring device repeatedly. Its position after each revolution will normally be different. Its (x, x') , or equivalently (y, y') , on each turn will fall on an ellipse which is specific to the lattice location of the PUE. The x , or y , position is what the PUE measures.

After several revolutions, the particles within the beam tend to act less as a whole. This is due primarily to the beam's momentum spread. Unless the chromaticity is exactly zero, particles will have slightly different tunes. Roughly speaking, particles will then fall on the ellipse at slightly different (x, x') with respect to each other after each revolution. As the number of revolutions increases, the spread in (x, x') will increase, and the *coherent* motion of the half-turn visible on the PUE signal, will gradually disappear. The oscillations of the half-turn about the E.O. will be unobservable since the particles will cover the entire ellipse on each turn. Figure 8.9 shows these coherent betatron oscillations on the C7 PUE, which looks at the vertical position, or (y, y') .²⁵

The bottom trace, the C7 difference signal, shows the oscillations of the injected beam about the E.O. The top trace is the sum signal from the PUE. The coherent oscillations do not '*damp out*' significantly over these forty or so revolutions. This is probably because the vertical chromaticity is close to zero. Normally, the chromaticities are set very close to zero at injection for high intensity running. Presumably, this allows greater control over the beam's position in tune space. It reduces the tune spread slightly, by minimizing the component of the tune spread due to the momentum spread. However, the reason it is set there is purely empirical. That is, all other things being equal, the injection efficiency seems better that way, for whatever reason.

It is the deviation from the *baseline* of the PUE difference signal that gives a measure of the deviation of the beam's position from the beampipe center.

²⁵From N.W. Williams and H.M. Calvani-Garcia, HEP Setup Book I, pg. 85, FY '96.

Figure 8.9: Coherent Betatron Oscillations at injection using a half-turn on the C7 vertical PUE. The top trace is the C7 Sum, the bottom is the C7 difference signal.

In this case, the baseline undulates considerably over the 40 revolutions.

Notice that it takes about 23 revolutions for the beam to return to the same y value. For example, when the beam's position is $y = 0$ there is no deviation of the difference signal from the baseline. This occurs about $10\mu s$ (or 2 boxes) after the half-turn is injected. On subsequent revolutions, the beam position at the PUE becomes more negative, reaching a maximum at about 20-25 μs after injection. Then it moves back towards the beampipe center, arriving there about 35 μs after injection.

If the tune is exactly an integer, the phase advance in one revolution will bring the beam back to the same (y, y') on the ellipse at a particular location, say C7. Therefore, the position seen there, y , will remain the same on each revolution. If the tune is near, but not at an integer, (y, y') on each revolution will be slightly different than on the previous, and y will change slightly each time. These changes will be cumulative as the difference in the phase from that of the integer tune increases on each turn. Mathematically, let $\Delta\psi$ be the difference in the beam's phase on the ellipse from that associated with the nearest integer tune after one revolution. Then $n\Delta\psi$ will be the difference in phase from a particle with that integer tune in n revolutions. If $y = 0$ on the k th revolution, then y will again be near zero when $(n - k)\Delta\psi$ approaches π .

In the case shown in figure 8.9, $(n - k)\Delta\phi \approx \pi$ when $(n - k) = 23$. So, it takes about 46 revolutions to make one full rotation about the ellipse at C7. So, $46 \cdot \Delta\psi \approx 2\pi$, and $\frac{\Delta\psi}{2\pi} \approx \frac{1}{46} = 0.02$. Since, it is assumed that ν_y is near 5, ν_y is either 5.02, or 4.98. This is near the requested high intensity injection tune. Since only a half-turn is injected, the beam's tune is not effected by space charge. Consequently, the actual tune is very close to the requested tune, which in this case is about 5.

PIP

A program called PIP²⁶ is used to analyze the coherent oscillations on PUE signals. Position measurements for each revolution are obtained from PUE sum and difference signals. PUE data is obtained from the digital scopes in MCR through a Labview program and General Purpose Interface Busses (GPIB). The PIP program uses these connections to get the digitized PUE

²⁶PIP stands for Proton Injection Parameters.

Figure 8.10: A fit of the position on each revolution (or turn) to a curve at C7. The y-axis is the position. The x-axis is the turn, which is proportional to the amount of rotation about the phase space ellipse located at the PUE.

signals from a scope.

The positions on each revolution are fit to a curve. Figure 8.10 shows the fit for the data in figure 8.9. Each revolution (or turn) the phase on the ellipse shifts by $\Delta\psi$. The turn number, is represented on the x axis. The positions on each turn are represented on the y-axis. Notice that in this case, the beam makes almost one complete oscillation. Also, the oscillations are nearly centered about $y = 0$. So, the position of the E.O. at C7 is near the beampipe center.

Given the lattice parameters at the PUE which describe the shape of the ellipse, the area of the ellipse that the center of mass traces out can be found

using the fitted data. The so-called *Courant-Snyder Invariant*, or CSI, is this area divided by π . If the beam were one particle, the particle's emittance would be equal to this area. PIP calculates this number. It is called an invariant because its value is the same at every azimuthal location in the ring. It is the emittance of the center of mass particle. It gives a measure of how far off the E.O. the beam is injected. If the beam were injected directly onto the E.O. this quantity would be zero.

The (y, y') relative to the E.O. on the first revolution can also be found. This (y, y') can be propagated back to the foil using the transport matrix between the PUE and the foil to arrive at the trajectory of the injected beam relative to the E.O. at the foil. However, the parameter of primary interest here is the CSI. By using the last few steering magnets in LTB this can be adjusted (DV95 and DV083). Alternately, it can be adjusted by changing the E.O. In the vertical plane, this can be done with the vertical orbit correctors. There is a vertical injection bump that is composed of 4 of these dipoles in the injection area. It is controllable through the BoosterOrbitControl program.

The value of the vertical CSI is critical at high intensity. This isn't surprising since it is related to the charge density in the vertical dimension. Recall that in the low intensity setup, one tries to inject the beam tangent to the horizontal E.O., which is moving during injection. The vertical E.O. is not moving.

Adjusting the Injection Bump's $\frac{dB}{dt}$ Ramp

Figure 8.11 shows a typical proton injection bump's (x, x') and time dependence at the foil. The third and fourth points on the angle and position functions can be used to adjust the ramp of the injection bump to match the E.O. while the B field increases. The ratio of the position to angle is usually around six. The range over which this ratio can be changed is fairly limited without changing the polarity of some of the injection bump magnets. The trajectory of the incoming beam can be adjusted to try to match this ratio using magnets in the downstream end of LTB (DH115 and DH088).

Without the injection bump, as the field increases, the E.O. moves inward at the foil. Its angle relative to the beam pipe center also changes there. The position and angle of the unbumped E.O. at the foil change in about a six to one ratio as the E.O. moves to the inside. To null out the effect of the increasing field on the E.O. at the foil, the position and angle parts of the

Figure 8.11: A typical proton injection bump as viewed through the BoostOrbitControl program.

injection bump should collapse in the same ratio. This is why the bump is setup as it is.

The SLD BIJ.FAST.TM, or *Fast.tm*, which sets the time the bump begins from Peaker, can be used to move the bump with respect to the beam. If the beam comes in on the ramp, moving the ramp, by adjusting *Fast.tm* is an easy way to adjust the bump's x and x' together. In general, *Fast.tm* is a very useful tuning *knob*.

The chopper can be used to inject half-turns at the beginning, end, and middle of injection. PIP can be used to measure the CSI at these times, and the points on the injection bump can be adjusted to null out the CSI's time dependence. However, although the ramp was initially introduced to compensate for the increasing B field, at high intensity it plays a more important role.

8.3 Capture

8.3.1 Measuring the Revolution Frequency

Just as it's important to inject the beam near the E.O. in transverse space, it's also important to inject the beam near the synchronous phase and energy, the longitudinal analog of the E.O. To match the synchronous energy, the Rf frequency is set to match the desired harmonic of the revolution frequency of the linac energy beam.

Once spiraling beam has been established, the revolution frequency, f_{rev} , is measured. It can be measured using a method similar to that discussed in 7.2.1 for Heavy Ions. In that case, while the Rf is off, a hole is kicked in the beam so that it becomes visible on a PUE signal. Then a number of turns are counted, and the time it takes to make those turns is measured. By contrast, with protons, since the beam is routinely chopped, it is immediately visible on the PUEs, or the E6 Wall Current Monitor (see section 3.1.5).

The Wall current monitor is a *resistive* pick-up, unlike PUEs which are capacitive. Consequently, the 'Wall monitor' can better resolve the time structure of the bunch. As the bunch circulates, its width increases due to its momentum spread. It is the revolution frequency of a particle in the center of the bunch's momentum distribution that we seek to measure. The more turns are counted, the more resolution the measurement has. But, the

longer the beam has been in the machine, the wider the bunch becomes, and the harder it is to find its center. This is because the particles in the bunch have slightly different momenta. So, it's important to make the injected bunch as narrow as possible. The Chopper program is used to configure the fast chopper to send one narrow micropulse, which can be seen on the wall monitor. This bunch typically has a width of 90° .

The wall monitor signal can also be fed into a *Mountain Range Display*.²⁷ Figure 8.12 is a Mountain Range Display from the instructions in the Morning Numbers logbook. To create a '*Mountain Range*', a guess is made for the value of the revolution frequency. Then, a scope is triggered repeatedly near some designated harmonic of that frequency. The scope is setup to take multiple traces that occur after it receives a start trigger. The number of traces to be taken is also specified. The Mountain Range Display program displays the traces that result consecutively, one above another.

If the triggering frequency is nearly an integer multiple of f_{rev} , the bunch on each trace will appear almost directly above the bunch on the previous trace. As the bunch is followed from one trace to the next it will move further to the right or left, depending on whether the triggering frequency is just greater, or just less, than the harmonic of the revolution frequency. The triggering frequency is adjusted so that the bunches in each trace line up vertically. When this happens, the triggering frequency, which is known, will be a harmonic of f_{rev} .

This method for measuring f_{rev} is generally regarded as more accurate and less prone to error than the counting method. It is used in the Morning Numbers measurement of f_{rev} , which is made to check the linac energy.

It's important to realize that the mountain range would look entirely different if the Rf was on. Of course, the behavior of the beam would be entirely different. But, the important point is that a revolution frequency measurement like the one described would not be possible. If the bunch was in the Rf bucket, and *Beam_Control* was off, the bunch's frequency would oscillate about the frequency corresponding to the synchronous energy, *not*

²⁷A description of the generic Labview Mountain Range Display program is given in: B. Tamminga and W. Van Asselt, AGS/AD/Op. Note No. 44, *Mountain Range Display*, December 12, 1995. The program used for the Booster revolution frequency measurement is the *turn by turn* Mountain Range Display, which is similar, but not identical to this. Instructions on how to take a revolution frequency measurement with it are contained in the Morning Numbers Logbook.

the linac energy. If *beam_control* was on, the Rf frequency would typically shift to damp out these synchrotron oscillations, but not before the Rf did work on the beam to change its energy from its initial value.

Once f_{rev} is measured, the Rf group usually adjusts the Rf frequency so that it is hf_{rev} at injection. Afterwards, the Rf is 'turned on', and they adjust things to accelerate the beam to extraction.

8.3.2 Longitudinal Matching at Injection

Once the Rf is set to the measured value of hf_{rev} , and the Rf is turned on, the bunch can be viewed on the mountain range. If *Beam_Control* is off, the bunches will probably still oscillate about the synchronous phase. With the Rf frequency set to hf_{rev} , the energy is presumably very close to the synchronous energy. Any oscillations are then mostly due to a mismatch between the Rf phase and the bunch's phase at injection. As mentioned in section 8.2.2, the bunch's phase relative to the Rf can be adjusted with the *Chopper_Phase* SLD.

Another way to see the synchrotron oscillations that a bunch undergoes is with the *Radial_Average* and *Beam_Phase_Error* analog signals (see section 6.3). Oscillations on the radial signal represent changes in energy from the synchronous energy, oscillations on the phase error signal represent changes in phase from the synchronous phase. Figure 8.12 shows both these signals with about 100-150 turns of injected beam.²⁸

Before the beam is injected these signals are flat. During injection, the signals are hard to decipher. Injection ends well before $300\mu s$ from the start of the traces. After it is over, the oscillations in both signals are presumably a result of a 'mismatch' between the beam and the Rf. Notice that the oscillation in one signal is about 90° 'out of phase' with the oscillation in the other. In longitudinal phase space, $(\Delta E, \Delta\phi)$, the beam's center of mass moves around the stable fixed point defined by the synchronous energy and phase. When it is far from the synchronous phase, it is close to the synchronous energy, and vice versa.

In principle, if the energy is matched, the oscillations can be nulled out by adjusting the chopper phase. However, since this beam is injected over a finite time, and the Rf frequency is changing during that time, the real

²⁸From K. Reece, *Booster Setup Book VII*, pg. 83, 1993

Figure 8.13: The radial average signal (top) and the beam phase error signal at injection, with a narrow injection pulse and beam control off. The radial average signal is BXI.RFE6.RADIAL_AVG, and the phase error signal is BXI.BEAM_PHASE_ERR

picture is far more complicated. It's instructive to move the chopper phase around, and watch its effect on these signals, and also on the losses very early in the cycle. One type of early loss is a capture loss. If the chopper phase is moved far enough from the synchronous phase, part of a bunch will not be injected into the Rf bucket, and will spiral inward, and be lost. This type of loss is easy to see on a current transformer.

Recall that *Rf_track* is used to change the Rf frequency at injection. Naturally, its value will have an effect on the matching. However, the synchrotron oscillations will typically damp out quickly if *Beam_Control* is on (as seen in figure 6.2). If *Beam_Control* is on, the Rf will move towards the beam's frequency, and only the initial frequency will be solely determined by *Rf_track*. *Beam_Control* is typically set to turn on very early in the injection process, sometimes right as the beam begins to come in ($2500\mu s$ after Peaker). It is set there because efficiencies are typically better. But, if the facility exists for the Rf to damp out synchrotron oscillations, and match itself to the beam, why shouldn't it be used as soon as reliable information about the beam is available? Recall that the beam's phase information comes from the D6 wall current monitor. Unlike Heavy Ions, the beam is bunched coming in, and phase information from the wall monitor is immediately available.

The same is true for the *Chopper_Phase*. If *Beam_Control* is on, the Rf will move to match the beam's phase. However, even with *Beam_Control* coming on very early, both the *Chopper_Phase* and *Rf_Track* have 'optimum' values. Presumably, this is because the response time of the Rf system to the beam information is finite.

Capture loss and high longitudinal emittance go together. Typically, the more a bunch is *mismatched* to the bucket, the larger the longitudinal emittance and capture loss will be. However, paradoxically, when running high intensity, beam survival early in the cycle may sometimes be improved by mismatching the bunch and bucket. Presumably, this is for the same reason that injecting the beam off the E.O. can improve efficiencies. That is, because it reduces the charge density. In this case, it is the longitudinal 'line' charge density that is being reduced. This longitudinal charge density not only effects the beam's longitudinal behavior, it also effects its properties in the transverse plane. If the longitudinal density is higher, a transverse cross section of the beam will have a higher charge density, and the transverse motion will be different. In this way, longitudinal and transverse motion become inextricably woven together at high intensity.

Other, less destructive, methods are used to modify the bunch shape besides an intentional mismatch. In particular, the *second harmonic Rf cavities* have been used to modify the bunch shape. There will be more about these later.

8.3.3 Distinguishing Capture Loss from Other Losses

There is a trick that can be used to look at the capture loss. The wall current monitor only sees the AC, or bunched component, of the beam. If the bunch shape does not change appreciably over some interval, the peak current will be proportional to the total current of the bunched beam. A current transformer measures the total current, the sum of the bunched and unbunched parts of the beam. Since the magnetic field is increasing rapidly, the beam that is not captured is lost very soon after injection. This uncaptured beam is not bunched, and all of the beam remaining after the uncaptured beam is lost is bunched.

In figure 8.13, the injection transformer and the D6 wall current monitor traces are shown at, and just after injection.²⁹ The baselines of both traces have been set to be the same. Their gains have been adjusted so that the wall current monitor peak voltage is the same as the current transformer voltage a couple milliseconds after injection. The amplitude difference between the two traces before this time presumably reflects the difference between the bunched and total (bunched + unbunched) beam currents, or intensities.

Notice that the total beam intensity is initially higher, and then decreases, until it matches the bunched intensity. This is the loss of the unbunched component of the beam. This unbunched beam is apparently not captured or accelerating, and the loss is due to this beam spiraling inward and hitting the beampipe because of the increasing Main Magnet field.

²⁹Taken from J.M. Brennan, *Booster Book IX*, pg. 86-89, 1993

Figure 8.14: The Injection transformer and the D6 wall current monitor during early acceleration. There are 300 turns injected, and the bunch width is 120° . The peak total intensity is 16 Tp, the peak bunched intensity is 13 Tp. The bunched intensity was measured using the method described in section 3.1.5 using the wall current monitor. Another trace is also visible in the picture. This trace is the injection transformer with the Rf off, and shows the spiraling beam. The latter part of it is obscured by the wall current monitor signal.

Chapter 9

Proton Operation II: *Loss Mechanisms*

There are a number of factors that effect beam survival. The interplay between these factors becomes very complicated when running high intensity. I emphasize that the Booster intensity is *not* limited as much by the finite Linac current, as by 'problems' associated with high intensity effects inside the Booster. Since the name of the game is 'High Intensity', solutions to these 'problems' are sought. These '*solutions*' are primarily found using 'trial and error' methods, and often move the machine far away from any well understood setup.

At high intensity, the emphasis is on correlating observations with parameters, without necessarily knowing the reasons behind such correlations. The ways in which parameters effect beam survival, and relate to the configuration of other parameters, can still be described with some allusions to physical mechanisms, but these descriptions are often highly qualitative and speculative. Yet, the descriptions are helpful because they provide a framework that can, on occasion, provide some direction and insight. However, relying too heavily on these descriptions can be a mistake, because they are often crude, and tend to be incomplete or even wrong.

It's instructive to give a real example of how the machine can '*behave*' in a way opposite to the way one might expect. It is a common observation that raising the requested vertical tune above its low intensity value can dramatically reduce early losses. Even though space charge effects become less important later in the cycle, theoretically, they are still there. If raising

the requested tune early in the cycle helps because it moves the beam away from stopbands, then one would not think lowering it *below* the low intensity tune would improve the machine's performance. However, the fact is that the amount of beam that can be transferred to the AGS often increases when the requested tunes are lowered *below* the low intensity tunes *after* the early part of the cycle. A typical configuration for the tune functions during high intensity is shown in figure 5.4. Notice that the requested tunes between 30 and 50 ms, the middle of the cycle, continue to drop, to levels below those that are typical for low intensity.

It is believed that most of the transverse losses that occur at high intensity are in one way or another due to stopbands. There are two primary ways that losses due to stopbands are reduced in the Booster: 1) Correcting the field errors that cause stopbands, and 2) Keeping the beam away from stopbands in tune space. To a great degree, the correction of stopbands can in principle be done in a systematic manner that does not involve much 'trial and error'. By contrast, beam is kept away from stopbands largely by using 'trial and error' methods.

The beam is kept away from stopbands in several different ways. First, as discussed earlier, the beam is accelerated quickly away from injection momentum. Since the space charge related tune spread is greater at lower momentum, it is harder to avoid stopbands when the momentum is low. Raising the beam's momentum, reduces the size of the beam in tune space, which keeps the beam away from stopbands.

Secondly, the charge density is reduced in both the transverse and longitudinal dimensions to reduce the space charge tune spread. In the transverse plane, the charge density is reduced by injecting the beam away from the E.O. In the longitudinal plane, the second harmonic cavities are used.

Lastly, the beam is moved in tune space away from the offending stopbands by adjusting the Tunes. The Orbit is affected by these tune changes (see section 5.7), and the Orbit Correction system must be used extensively to improve it. Adjusting the Chromaticity can also have some small, but significant effect.

There is still more that has to be done to limit losses. Some losses are not a direct result of stopbands, but result from the increased transverse emittance associated with them. The emittances, both longitudinal and transverse, also tend to increase at high intensity just because it's harder to inject a lot of beam into a small area, especially when the beam current is finite,

and $\frac{dB}{dt}$ is high. Consequently, maximizing the acceptance using the orbit correctors, and keeping the beam in the center of it with radial steering, are very important. This is particularly true early in the cycle, when the beam is the largest.

In what follows, I'll try to address these pieces of the high intensity puzzle. The way the machine actually runs is a compromise between all these factors. For example, the tune spread will decrease if the beam is injected further from the E.O., but the emittance of the beam will increase, and so the injection efficiency will be reduced. In an optimized high intensity machine, increasing the amount of beam initially stored by injecting the beam closer to the E.O. will usually increase the stopband losses, because the charge density will be higher, and actually decrease the amount of beam making it to extraction energy! This is inferred from the observation that adjusting the vertical steering using DV095 can increase the Booster's early intensity while simultaneously decreasing its late intensity.

9.1 Stopbands

9.1.1 Correcting Stopbands Systematically

At low intensity the beam has a tune spread that is only due to its momentum spread. This tune spread may be reduced even further by setting the chromaticities to zero. The low intensity beam occupies a rather small area in tune space. Consequently, the beam can be moved away from strong stopbands by modifying the tune functions, and stopbands are not a significant problem. Alternately, the beam as a whole can be moved into a resonance condition by adjusting the tune functions. Typically, a tune function is modified so that the beam passes through a stopband at a particular time in the cycle. A loss will normally occur as the beam moves through the stopband. This loss is visible on the current transformer, and its magnitude is a measure of the stopband's strength.

The loss on the current transformer is measured, and expressed as a fraction of the beam's intensity just before the loss. The StopbandCorrect program is used to adjust the relevant harmonic components of the field responsible for the stopband (see section 5.4.2). For example, say the stopband that I am trying to correct is $2\nu_y = 9$. The amplitudes of the $\cos 9\theta$ and $\sin 9\theta$

quadrupole field components determine the strength of this stopband. The amplitude of each of these components, during the time in the cycle that the loss occurs, is scanned separately. For each component, the loss is recorded at several different amplitudes. One will typically find that the amount of loss passes through some minimum value as the component's amplitude is scanned. The smaller the loss, the less the strength of the field error responsible for the loss. The amplitude that minimizes this loss is the value used 'to correct' that harmonic component of the field responsible for the stopband. This procedure can be repeated for each stopband that has correction magnets.

Figure 9.1 shows the loss on the current transformer before and after a correction is applied for the $2\nu_y = 9$ quadrupole stopband.¹ The top trace is the current in the vertical tune quads. The vertical tune is lowered through $\nu_y = 4.5$ and then raised back up through it again. The tune closely mimics the current shown in the quad signal. In the corrected case, a loss occurs each time the tune passes through the half-integer stopband. In the uncorrected case, little or no beam survives through the first passage.

To complicate matters, the strength of field errors is dependent on what the Main Magnet is doing. Since, field errors are produced primarily by the main magnets, one expects that the current required to correct them will be highly dependent on the Main Magnet function. However, unlike functions for other parameters that are highly dependent on the Main Magnet, such as the Tune and Chromaticity functions, until recently StopbandCorrect functions were in no way derived from the Main Magnet function. Consequently, this dependence had to be introduced manually into the functions. The procedure described above only works for a particular Main Magnet function at a particular time in the cycle. Correcting all the stopbands throughout a magnet cycle using this method is an extremely time consuming task.

The strength of stopbands is particularly dependent on the main magnet's $\frac{dB}{dt}$. A changing magnetic field induces an electric field, which causes currents to flow along the vacuum chamber. These *Eddy* currents, produce magnetic fields which cause stopbands. The most recent version of StopbandCorrect allows one to make correction functions that are functions of B , $\frac{dB}{dt}$, and $\frac{d^2B}{dt^2}$.

Figure 5.8 is a typical set of correction functions. These are for the half-integer stopbands, $2\nu_x = 9$ and $2\nu_y = 9$. These functions have been arrived

¹ Adapted from, Y. Shoji, Booster Book VIII, pg. 67, 1993

7-May-93
23:59:25

$$2Q_Y = 9$$

Figure 9.1: The $2\nu_y = 9$ stopband, before and after correction using the half-integer stopband correctors. The top trace is the current in the vertical tune quads. The bottom traces are the current transformer, before and after the correction has been found using the method described in the text. Incidentally, 'Q' is another symbol commonly used for the tune.

at primarily by using this systematic approach. However, while running high intensity the stopband correction functions continue to be tuned. More about this later.

Fortunately, stopbands are only an unavoidable loss mechanism during the 'early' part of the cycle. Presumably, this is because the space charge tune spread becomes small enough later in the cycle that the major stopbands can be avoided by adjusting the tunes. However, even though stopbands may not cause significant beam loss later in the cycle, they may still have an important effect on the beam's emittance.

9.1.2 Measuring the Tune

When trying to correct stopbands, the question arises: "How do I know which stopband is causing a loss?" If I know which stopband is causing a loss I will know which magnets to use to correct it. To know which stopband is responsible, I need to know what the tunes are. One can 'believe' the tune functions, or better still one can measure the tunes. In section 8.2.7, the tune was measured by counting how many revolutions it takes a 'half-turn' of beam to trace out one turn about its phase space ellipse at C7. Recall that this information only yielded the difference between the tune and the closest integer. In the example, ν_y was measured to be ± 0.02 away from an integer. Since, the model of the Booster lattice expects the tune to be about 5, the tune was taken to be either 4.98 or 5.02.

So, the tune can be measured at injection. But, coherent betatron oscillations typically disappear shortly after injection. However, a magnet with fast rise and fall times, like the F3 kicker, can be fast enough to kick a bunch on only one revolution. If the kick is not too strong, the bunch will not be kicked out of the machine, but will only be displaced from the E.O. It will then undergo coherent betatron oscillations, which will be visible on a PUE.

In figure 8.9, the 'envelope' which encloses the coherent oscillations resembles a sine wave of the form, $x = A \sin(2\pi \Delta\nu f_{rev} t)$, where $\Delta\nu = \frac{\Delta\phi}{2\pi}$ is the difference between the tune and the integer closest to the tune. In the figure, $\Delta\nu = 0.02$ and $f_{rev} = 842.2 \text{ KHz}$.² The envelope's 'oscillation period', $\frac{1}{\Delta\nu f_{rev}}$, is about $60 \mu s$, which is consistent with the figure. This frequency, $\Delta\nu f_{rev}$, can be thought of as a 'beat frequency' between the betatron os-

²This is f_{rev} at injection. It is the value measured in figure 8.12

cillation **frequency** and nf_{rev} , where n is the closest integer to the tune. The closer these frequencies are, the slower this 'beat frequency' is. A Fast Fourier transform, or 'FFT', of the PUE signal is performed to find this beat frequency, and thereby determine the fractional part of the tune.

There is a Tune Meter kicker for each plane. Their power supplies are located in building 930A. Real timeline events control their charging and discharging. These events are located on spreadsheet together with the on/off and setpoint controls for the power supplies. The BoosterTuneMeter program can also be used to set the kick level (in kV), and the kick (or discharge) time (in microseconds from BT0). The BoosterTuneMeter program uses PUEs whose gains can be changed independently of the other PUEs (or bpms).

Figure 9.2 shows a tune measurement as displayed in the BoosterTuneMeter program window.³ Both tune meters are setup to discharge at 66.000 ms and at 10 kV. The '*Kick Plane*' sets which kickers 'fire'. The '*Digitize Plane*' sets the plane from which data is displayed. The number of measurements to average over is set in '*Num Averages*'. The top traces are the PUE signals. Notice that oscillations are visible on each of the PUE signals. These oscillations begin when the kickers discharge. The horizontal oscillations damp out faster than the vertical. This may indicate that, in this case, the vertical chromaticity is closer to zero than the horizontal. The bottom traces are FFTs of the PUE data. Below the bottom traces are values for the tunes. These '*values*' are just the peak amplitudes in each FFT. One can zoom in near an FFT peak and select a different value if desired. These FFTs are particularly clean and unambiguous, others may require closer inspection.

There are analog signals available on Xbar for the discharge current in each of the tune kickers. They can be found under 'Booster/Diagnostics'. The tune meter measurement is complicated in that it requires several distinct elements to work together. The controls have to send the request from the program, the kicker has to kick, the PUE has to work, and the program has to get the PUE data back to analyze it. With a kicker and PUE signal on a scope one can at least confirm that the kick is occurring at the right time. One sees the discharge current and the oscillations the kick induces on the PUE signal.

The tune meter measurements yield the tune of a bunch as it oscillates about the E.O. This is called the '*coherent*' tune because it is the tune of all

³J.M. Reich, *Booster Protons* 4, pg. 118, 1994

Figure 9.2: The Booster TuneMeter program window.

the particles moving together. When the intensity is low, this coherent tune is close to the tunes of the particles within the beam. At high intensity, the coherent tune and the tunes of individual particles can be quite different.

To a great degree, the space charge tune shift that particles *within* the beam experience occurs because of fields generated directly from charges and currents *inside* the beam.⁴ On the other hand, there are fields that are external to the beam that result from the interaction of the beam's charge and current with its environment. These fields are the space charge related fields that primarily effect the coherent tune. For example, the image charge on the inside surface of the vacuum chamber, that results from the beam's charge, produces an electric field that can decrease the focussing field in both transverse planes.

In short, because of space charge, the tunes of individual particles can vary considerably from one place to another within the beam, and these tunes result in part from fields that do not effect the coherent tune. This distribution of tunes *inside* the beam is called the '*incoherent*' tune distribution, and unlike the coherent tune, it is very difficult, if not impossible, to directly measure it.

The '*tune spread*' refers to the maximum range of particle tunes within the beam's tune distribution. The average tune of the particles is sometimes called the *incoherent tune*. The '*average tune shift*' is the difference between the tune without space charge effects, about the same as the requested tune, and the average tune.

Figure 9.3 is an estimate of the incoherent tunes throughout a high intensity Booster cycle. The maximum and minimum tunes of particles within the beam are plotted for both planes. Compare these tunes to the high intensity tune functions shown in figure 5.4.

9.1.3 Correcting Stopbands at High Intensity

Even after many of the stopbands have been corrected in a systematic manner, one finds that changes to the correction functions can occasionally improve the machine's performance. Figure 9.4 shows the effect of a $3\nu_y = 14$

⁴The beam's charge generates an electric field, and its current generates a magnetic field. The strength of these fields depends on the position relative to the bunch's transverse and longitudinal center.

Figure 9.3: Estimated Maximum and minimum incoherent tunes. The maximum and minimum incoherent tune shifts were added to the requested tunes throughout the cycle. The tune shifts were determined using a model for the tune shifts, and parameters from the high intensity setup. The top trace is with the minimum vertical shift, the second from the top is with the minimum horizontal shift, the third trace is with the maximum vertical shift, and the bottom trace is with the maximum horizontal shift. I emphasize that this is only a model, there is little doubt that the actual tune shifts are significantly different from these. Taken from, K. Zeno, Booster Tech. Note No. 227, "Coherent Tune Shifts in the Booster", September 6, 1995.

(skew sextupole) stopband correction on losses at high intensity.⁵ This correction was arrived at by scanning the correctors and looking at their effect on losses.

9.2 Keeping Away From Stopbands

Minimizing the effects of stopbands is an extremely time consuming task. Their effect on the amount of beam that *can* be transferred to the AGS is the most important consideration. This amount does not only depend on the intensity at extraction energy, it also depends on the beam's emittance there. Just because the beam is in the Booster, doesn't mean it can be extracted, transported, and successfully injected into the AGS. All other things being equal, the smaller the beam is, the more efficient the transfer will be. Both the intensity and emittance may be affected by stopbands. Naturally, the amount that can be transferred also depends on how extraction, BTA, and the AGS are set up.

9.2.1 The Emittance Dependence of the Transfer Efficiency

During proton running, there are several Booster cycles in quick succession within each supercycle. The beam from each cycle is extracted, transported, and injected into the AGS while the AGS main magnet remains at a constant field appropriate for storage of the Booster beam. This part of the AGS cycle is called the '*injection porch*'. The Booster cycle is typically about 150 ms long. In between each cycle, the beam from previous cycles continues to circulate in the AGS on the injection porch. In the past, there have been 4 transfers during each AGS cycle, this configuration requires some of the beam to stay on the porch from extraction on the first Booster cycle to extraction on the last Booster cycle, or $3 \cdot 150ms = 450ms$. During this time losses occur in the AGS.

The transfer efficiency is defined as the ratio of "*the intensity in the AGS at the end of the injection porch*" to "*the Booster intensity just prior to extraction summed over all the Booster cycles*". The transfer efficiency

⁵T. Roser, Booster Book IX, pg. 107, 1993.

8-May-93
22:36:32

Figure 9.4: The effect of the Skew Sextupole $3\nu_y = 14$ stopband correction on high intensity beam. The two traces are the injection transformer with and without the correction. The late intensity in the corrected case is about 10 T_p .

depends on the amount of loss incurred during extraction, BTA transport, and injection into the AGS on each transfer, as well as the beam survival on the injection porch.

As mentioned in earlier chapters, there are digital readouts, called '*scalers*', located at each console in MCR. Two of these scalers display the numbers used to determine the transfer efficiency in 'real time' during every super-cycle. They are labeled '*AGS CBM*'⁶ and '*Booster Late*'. The intensity early in the Booster cycle, and the intensity of the beam coming from Linac, are also displayed on scalers. These scalers are labeled '*Booster Early*', and '*Booster Input*', respectively. The times in the Booster and AGS cycles that the measurements are taken are set in SLDs.⁷

The Booster and AGS measurements come from sampling the Booster and AGS current transformers at the times set in these SLDs. For each of the two Booster scalers, the data from the sampling is summed over all the Booster cycles to yield the total intensity at that time in the Booster cycle, in each AGS cycle. For example, let the *Booster Late* sample time be set to 60000 gauss clock counts, and the measured intensities on the four cycles at that time be 14.37, 15.22, 17.34, and 13.12 Tp. Then the *Booster Late* scaler will read 6005, corresponding to 60.05 Tp, the sum of these numbers. The *Booster Input* scaler is also summed for each AGS cycle.

The '*scalers*' are a very convenient tool used during tuning to gauge the effects of the changes one makes. Since the scalers are such a convenient and simple tool, they can also be relied upon too heavily. For example, one might assume that the Booster is running better because the '*Booster late*' scaler is reading higher. However, the emittance of the Booster beam is also an important parameter when it comes to gauging how well the Booster is running. But it doesn't have a scaler associated with it, so it has a tendency to be overlooked. Nevertheless, the numbers on these scalers are still regarded by many as 'the last word' when it comes to machine performance.

⁶CBM stands for Circulating Beam Monitor.

⁷The SLDs can be found under *AGS/Timing/Instruments/cbm*. The SLD for *Booster Input* is BGN.INTEGRATE.P.SP, for *Booster Early* it is BGN.EARLY.CBM.SP, for *Booster Late* it is BGN.LATE.CBM.SP, and for *AGS CBM* it is AGN.CBM.SP. The Booster SLDs are gauss time line events, and the AGS SLD is an AGS real timeline event. Also of interest is the SLD AGN.INJ.CBM.SP, which sets the time for the *AGS first transfer* scaler. This time is set to measure the AGS transformer right after the first Booster transfer. It is an AGS real timeline event.

Changes in the widths of the beam's profiles on the multiwires in BTA can reflect changes in the Booster beam's emittance. These widths can be changed without effecting the Booster's late intensity by modifying the vertical tune during the acceleration cycle. The amount of beam that reaches the end of the injection porch in the AGS is strongly correlated with these changes. It follows from this that the transfer efficiency is dependent on the emittance of the beam in the Booster, and that this emittance is not necessarily correlated with the Booster's intensity.

In figure 9.5, the transfer efficiency is plotted against the horizontal beam width on MW060. The only parameter adjusted during these measurements was the vertical tune between 29 and 57 ms from BT0. This period is after the early losses and before the extraction related processes begin. The Booster late scaler read about 46 Tp, and varied by less than 1 Tp over the course of the measurements. The requested vertical tune was scanned between 4.60 and 4.76. The width tended to increase as the tune was lowered.⁸

Presumably, the emittance increases because of the effects of stopbands. In this case, it may be the $2\nu_y = 9$ stopband that is responsible. Perhaps surprisingly, the vertical width on MW060 did *not* change significantly. There are two points here: 1) The emittance can be increased without effecting the Booster's intensity, and 2) This emittance change *can* effect the transfer efficiency. The second point illustrates how some losses that occur during transfer can have their 'roots' in what happens well before it.

One might try to conclude from this experiment that the Booster's acceptance is larger than the acceptances encountered during transfer. However, this conclusion is not entirely valid. The statement has to be qualified since some of the reduction in transfer efficiency may be due to less than optimized extraction and AGS injection, as well as steering problems in BTA.

One might also be tempted to conclude from the observation that the horizontal width on MW060 increases, that the horizontal acceptance upstream of MW060 is *not* responsible for the reduction in transfer efficiency. Since, if it were responsible, how could the wider beam make it down to the multiwire?

Suppose that the emittance were such that the beam just fit through an aperture. If the emittance were then increased, particles with x amplitudes larger than the aperture would be lost. But, particles with smaller x ampli-

⁸K. Zeno, IIEP Startup Book II, pgs. 58-66, FY '95

Figure 9.5: Transfer Efficiency versus Horizontal beam width (FWHM) on MW060 in BTA. This plot shows the effect of emittance on transfer efficiency.

tudes and large x' amplitudes would not be lost. In general, an increase in emittance corresponds to an increase in the average x and x' amplitudes of particles in the beam distribution. As a consequence, particles coming out of the limiting aperture after such an increase in emittance have, on average, larger x' amplitudes. Such particles, moving at larger angles with respect to $x = 0$, will naturally develop larger x amplitudes as they move down the beamline. Therefore, the width at MW060 will generally increase if there are limiting apertures *upstream* of it.

9.2.2 The Tunes Near Injection Momentum

In the previous section, the emittance was increased by moving the beam in tune space. This caused losses during transfer. Losses very early in the Booster cycle can also be highly dependent on the requested tunes.

The Vertical Tune at and near Injection

The first step in raising the Booster's intensity is to increase the amount of injected beam. This is done either by increasing the number of turns, the *bunch width*, or both of these parameters.⁹ One invariably finds that the Booster's efficiency, defined here as *Booster Late* over *Booster Input*, will start to drop at some point as the amount of beam coming from the Linac is increased.

At this point, one may sometimes find that raising the requested vertical tune near and at injection, away from its low intensity value can improve the Booster's efficiency. However, if the vertical tune is raised beyond a certain point, *Booster Late* will again begin to drop.

Since distortions in the vertical orbit typically increase in amplitude as the vertical tune is raised closer to an integer, it's natural to suspect that the drop in *Booster Late* as the requested tune is raised may be related to such distortions. One might expect this deterioration of the orbit to show up in orbit data acquired from the bpm system. Adjusting the vertical orbit's fifth harmonic at injection with the orbit correctors will often bring *Booster Late* back to what it was when the tune was slightly lower, or even raise *Booster Late* higher. Adjusting the vertical steering at the downstream end

⁹It is often also necessary to increase the slow chopper width, LPI.CHOP.WIDTH.

of LTB, with DV095 and/or DV083, can also restore or even increase the late intensity.

It's important to keep in mind what conditions are altered when these types of changes are made. For one thing, the orbit correctors, the requested tune, and the LTB steering, all effect the match between the incoming beam and the E.O. For example, since the vertical orbit changes when the tune is changed, the E.O. at the foil will change, and so the trajectory of the incoming beam relative to the E.O. will be different as well. Similarly, when the orbit correctors are changed, the vertical orbit changes, and the match is affected. Therefore, it's not surprising that the vertical steering has a different optimum position when the tune and orbit corrections are changed.

However, *Booster Late* often *increases* when these kinds of changes are made. One might expect it to be possible to restore the intensity by correcting the orbit and by adjusting the matching, but how could it be *increased*? The obvious explanation is that the orbit and matching were not optimized initially, and tuning them both caused the efficiency to improve. It is probable that this is often a large part of the reason for the increase.

Is it *necessary* to raise the requested vertical tune to reduce losses to the extent that they are reduced when the machine is optimized? There are several potentially important aspects of the Booster setup that change when the requested tune, at and near injection, is changed. In addition, there are many parameters that have to be adjusted to optimize the efficiency for a different requested tune. So, it's hard to confirm 'empirically' that the tune change is *necessary* to reduce losses to this extent.

In practice, when injecting high intensity, the Booster efficiency has clearly been its best when the requested vertical tune is high. In theory, I think it makes sense that the requested tunes might have to be adjusted as the intensity is increased because of the expected tune shift and spread due to space charge, and the existence of stopbands. Hence, it seems *reasonable* that raising the requested vertical tune near injection 'by itself' reduces losses, and that it does so because it moves the beam away from stopbands.

From the purely practical standpoint of 'Tuning For Intensity', the requested vertical tune, at injection and during the first 10 ms or so of acceleration, seems to play a role as important in some ways as the one the injection bump plays for Heavy Ions.

Figure 9.6 is a tune space diagram. The shaded box indicates where the beam is in tune space at injection using the model for the incoherent tunes

shown in figure 9.3. This model is for an optimized high intensity setup. Since this is only a model, the figure should only be taken qualitatively. Incidentally, the density of particles expected near the center of the shaded box is greater than the expected density near its edges.

From the figure you can see that the requested vertical tune at injection is set at about 5.00. This is typically the case when the machine's intensity late in the cycle is at its highest. PIP measurements, like the one in figure 8.9, confirm that the actual tune in the absence of space charge effects is in this vicinity. If it is raised only slightly higher than its optimum value, say by 0.01, the injected intensity often falls dramatically.¹⁰ If it is lowered slightly from this value, the intensity may stay about the same, or drop slightly. If it is lowered a little further, say by 0.02, then the intensity will generally begin to drop, but *not* dramatically.

The requested tune and the injection mismatch are highly dependent on each other since the tune shift and spread are affected by the mismatch. Once the machine is optimized for high intensity the most common tuning adjustments made *to keep* it optimized are small changes to the vertical tune at and near injection, and to DV095. DV095, which controls the vertical injection matching, typically needs the most 'attention'.

The Horizontal Tune at and near Injection

In figure 9.6, one can see that the requested horizontal tune at injection is not nearly as high as the vertical. The modeled tunes in the horizontal are also significantly lower. There are several possible explanations for this. First, since the tunes are shifted to avoid stopbands, what the beam tunes are when

¹⁰An hypothesis developed to explain this dramatic intensity drop is as follows. The tune of the beam that first enters the Booster is initially not affected by space charge forces, since there is no other beam in the machine. Consequently, its tune is at or near the integer, and it is *not* stable. In the case where the tune is raised too high, the beam will not survive long because it is very close to the dipole resonance. If the initial tune is slightly lower it will survive longer because it is not as close to this resonance. Beam continues to be injected, and this newly injected beam lowers the tune enough so that the beam initially injected becomes stable, and is stored. The space charge tune shift is *used* to move the vertical tune away from the integer, but it can only work if the beam initially stays in the machine for more than some minimum amount of time. The requested vertical tune has a threshold above which beam does not survive long enough to allow the tunes to be lowered by space charge so that the beam will store.

Model of Beam Distribution in Tune Space at injection

Figure 9.6: A model of the particle tunes in tune space at injection during high intensity running. The shaded box in the upper right indicates the beam's tune distribution taken from figure 9.3. The 'X' represents the position of the requested tunes.

the Booster efficiency is its highest can potentially tell us something about the locations of the stopbands that cause the most losses. In looking at the figure, one might be led to speculate that the requested (and actual) vertical tune needs to be so high because the $3\nu_y = 14$ stopband is very strong. And similarly, that the horizontal tune has no reason to be as high because there is no stopband of comparable strength in a position that would require the tune to be raised to avoid it. Figure 9.4 shows that the potential for a very strong $3\nu_y = 14$ stopband is there.¹¹

Alternately, one might speculate that the effect of space charge is less for some reason in the horizontal plane than it is in the vertical. Figure 9.6 does indicate that the average tune shift expected from this model in the horizontal plane is somewhat less than that expected in the vertical. However, the requested horizontal tune is still normally raised as intensity is increased. The optimum value at low intensity is about 4.77 to 4.80, and at high intensity it is about 4.83 to 4.87.

Another possible reason why the horizontal tune is not raised as high as the vertical is because of the injection bump. Recall that if a bump is not 'perfect', the orbit will be distorted *outside* the region where the intent is to distort it (i.e.-outside the bump) with so-called '*bump residuals*' (see section 7.3.5). As the horizontal tune is moved closer to the integer, these orbit distortions will probably increase. The orbit bump produced by the currents in the injection bump magnets also changes as the tune changes. So, when the tune is changed the orbit distortions increase, the bump residuals change, and the bump itself changes. This will tend to make the horizontal orbit harder to correct as the tune is raised.

Because of the injection bump, mismatching the E.O. with the injected beam in the horizontal plane is quite different than in the vertical. Aside from adjusting the LTB steering to improve the Booster's efficiency, principally with DH115 and DH088, the beam is often injected, at least in part, on the rapidly falling edge of the injection bump.

Notice in figure 8.6 that the beam is injected on the $\frac{dB}{dt}$ ramp. During typical high intensity running the bump's start time, *Fast.tm*, is often adjusted so that the last part of the linac pulse is injected about $50 - 100\mu s$

¹¹Recall that when a stopband is '*corrected*', the losses it produces are minimized as *much as they can be* with the 'correction magnets'. That does not mean that there will no longer be losses produced by the stopband.

after the end of the $\frac{dB}{dt}$ ramp. In figure 8.6 this would correspond to moving the bump earlier by 1 to 2 divisions, while injecting the linac beam at the same time. This has the effect of smearing the last part of the beam out in phase space and thereby reduces its density. The slope of the ' $\frac{dB}{dt}$ ramp' can also be adjusted to *smear out* the injected beam in phase space. However, just moving *fast.tm* appears to have a similar net result, and is much quicker.

A similar capability does not exist in the vertical plane because there is no way to change the vertical orbit this quickly. The fact that the Booster typically runs with the injected beam partially on the rapidly falling edge of the injection bump suggests that this so-called '*phase space painting*' has some positive effect on the charge density distribution that mismatching doesn't have.

Figure 9.7 illustrates the effect of adjusting both the requested tunes and orbit correctors near injection, as well as the LTB steering. Notice that not only is the injected intensity increased, but the slow loss, associated with the slope of the current transformer traces has also been reduced.

9.2.3 The Tunes During Early Acceleration

In figure 5.3, the tune functions '*optimized*' for high intensity are shown. In the figure, the requested vertical tune drops rapidly over the first 20 or so milliseconds of acceleration, and drops the most rapidly just after injection. If this behavior is interpreted as governed by the effects of space charge, it's reasonable to speculate that these effects are becoming less important as the beam accelerates. This seems reasonable because the requested tune is moved back towards its 'low intensity value' as the beam's momentum increases. In the absence of space charge effects, this 'low intensity tune' is presumably where it's easiest to minimize losses. So, the smaller the space charge effects, the closer to this tune the value of the *optimized* requested tune becomes. This speculation is also consistent with the prediction that the average tune shift and spread will decrease as the beam's momentum increases.¹²

¹²Roughly speaking, the space charge defocussing force is dependent on the momentum largely because it is the result of the action of two opposing forces. The magnetic force induced directly by the beam's current is directed towards the center of the beam's distribution, so it resembles a focussing force. Since the magnetic force is proportional to a charged particle's velocity, this focussing force increases as the momentum increases. On the other hand, the electric field induced by the beam's charge is away from the center of

9-Apr-01
5:47:01

Calculating XFMR before and after tuning Booster

Figure 9.7: The effect of adjusting the tunes and orbit correctors near injection, and the LTB steering, on early losses. The higher trace is the current transformer after tuning, the lower before tuning.

Figure 9.3 shows graphically a model's prediction of how the tune spread and average tune shift decrease as the beam is accelerated. For instance, notice that the vertical tune spread, the distance between the maximum and minimum incoherent vertical tunes, decreases with time from about -0.3 at injection to -0.1 at 30 ms. The maximum vertical tune is essentially the same as the requested tune since it corresponds to the location in the beam where the charge density is close to zero. So, the *average tune shift* decreases as the beam is accelerated.

Since stopbands are suspected as the mechanism which prevents the tune from being lowered without increasing losses, it's natural to look for a 'stop-band line' in the tune space diagram that the *optimized* tune distribution seems to avoid. In figure 9.3, the modeled minimum vertical tunes stay around 4.65 during this period. This is in the vicinity of the $3\nu_y = 14$ stop-band line, which occurs at $\nu_y = 4\frac{2}{3}$.

It's tempting to infer from this that the requested vertical tune is set high primarily to avoid this stopband. But, it's likely that other stopbands also enter significantly into the picture, even after they have been 'corrected'. Nevertheless, this description highlights the key elements involved, and the kind of 'squeeze play' that seems to exist between them.

The behavior of the horizontal tunes in figures 5.3 and 9.3 is less dramatic, but qualitatively similar to the vertical case. This leads one to suspect that a similar mechanism is at work.

9.2.4 The Second Harmonic Cavities

The second harmonic cavities were not part of the original Booster design. There are two sets of Rf cavities in the Booster. One set was designed for use with Heavy Ions, the other set was designed for use with protons. The Heavy Ion Rf cavities, located at A3 and B3, were designed to operate in a lower frequency range than the proton Rf cavities, at A6 and E6, since Heavy Ions come into the Booster at a much lower velocity. The proton Rf cavities

the distribution, so it is defocussing. It is not directly dependent on the beam's velocity. The magnetic component of the force is smaller than the electric component, so the net force is defocussing. But as the beam's velocity increases the net force decreases because of the increasing contribution of the magnetic force. This picture is further complicated by the fact that the beam shrinks as its momentum increases, as well as a number of other factors. See Edwards and Syphers. pgs. 174-180 for more on this.

operate over a frequency range from 2.3 MHz to about 5.0 MHz. The Heavy Ion cavities originally operated over a frequency range from about 0.6 MHz to 3.0 MHz.¹³ The Heavy Ion cavities were later modified to operate at twice the proton Rf frequency for use as '*second harmonic*' cavities. They can also still operate at the much lower frequencies required for Heavy Ions.

The peak amplitude of the Rf voltage wave in a Heavy Ion cavity is about 20 kV. However, further modification may allow this voltage to rise higher. Presently, the Heavy Ion cavities can deliver a *vector* voltage wave with a peak amplitude of about 40-45 kV. The peak amplitude of the Rf voltage waves in the Proton cavities are about 45-50 kV. So, they can deliver a *vector* voltage wave with a peak amplitude of about 100 kV.

When A3 and B3 are configured as second harmonic cavities, the phase of their vector voltage wave can be adjusted with respect to the phase of the A6 and E6 vector voltage wave. From MCR, their relative phase is adjusted through the counterphasing function. Recall that for Heavy Ion operation, the counterphasing function is used to adjust the phase relationship *between* A3 and B3 to adjust their vector voltage. Here it is used to adjust the phase of their vector voltage with respect to the A6 and E6 vector voltage.

The motion in longitudinal space can either be stable or unstable. Whether a particle's motion is stable depends, in part, on its location on the Rf wave relative to the synchronous phase, ϕ_s .¹⁴ In the Booster, ϕ_s is on the rising side of the Rf wave. Particles associated with a higher voltage on the wave than the voltage at ϕ_s will get more of a kick than a particle at ϕ_s , and their revolution *period* will decrease. If the particle's energy is not too far from the synchronous energy, the particle will move back in phase towards ϕ_s on subsequent turns. If a particle received less of a kick at a higher phase than at ϕ_s , then it would tend to move away from ϕ_s , and the motion would not be stable. The dependence of the voltage on the phase of the Rf wave is key. Within some range in particle energy and for phases greater than ϕ_s , if moving higher in phase on the wave results in a voltage higher than that at ϕ_s , then the motion will be stable. The higher that voltage is with respect to the synchronous phase, the faster particles will tend to move back towards it, and the more the bunches will 'tighten up'.

¹³"J. Michael Brennan, *Rf Beam Control for the AGS Booster*, BNL-52438, pg. 14, September 26, 1994".

¹⁴It also depends on its energy relative to the synchronous energy.

A steep $\frac{dV}{dt}$ in the vicinity of ϕ_s will increase the charge density because it will cause the bunches to 'tighten up' more. The second harmonic cavities are used to reduce $\frac{dV}{dt}$, particularly in the vicinity of ϕ_s . The center of the bunch is typically at the synchronous phase. Ideally, the phase of the second harmonic voltage with respect to the accelerating voltage is adjusted with the counterphasing function so that it reduces $\frac{dV}{dt}$ around ϕ_s . This is accomplished by placing the 'zero-crossing' of the 'falling side' of the second harmonic wave at the synchronous phase on the accelerating wave. In this way, the second harmonic cavities do not do any net work on the bunch to accelerate it, its energy is spent pushing the bunch apart.

Typically, both the accelerating and the second harmonic Rf are on throughout the cycle. The actual voltage waveform that determines the particle motion is the sum of both the second harmonic and the accelerating Rf waveforms. The resulting voltage waveform has a lower $\frac{dV}{dt}$ near a bunch's center than without the second harmonic, since the $\frac{dV}{dt}$ of the second harmonic wave is negative there. Consequently, the net 'restoring force' near the center of the bunch is reduced.

The overall effect of both sets of cavities is to distribute the beam more uniformly within the Rf bucket. Consequently, the peak charge density is less for the same bunch intensity. Just as important is the fact that the second harmonic cavities change the shape and size of the Rf bucket. Compared to the bucket at injection depicted in figure 8.7, which only takes into account the accelerating cavities, the actual bucket is thought to be more *square* and have a larger area.¹⁵ For a given number of turns, more beam can be injected into a bucket if the bucket is bigger. Its true that more beam can be injected into the original bucket by increasing the number of turns, but in that case the charge density will tend to increase.

F_Dot

At least in the past, the chopper program has not taken the effect of the second harmonic cavities into account. As a result, the chopper parameters have been used to modify the injected beam in a less than straightforward way. In particular, the *F_Dot* parameter in the chopper program has been adjusted to modify the shape of the bunch sent by the fast chopper.

¹⁵J.M. Brennan, personal communication

For a given setpoint of the *bunch width* parameter, the chopper program sends instructions to the fast chopper for the width of each micropulse within a bunch. Each micropulse is chopped to a certain width determined by these instructions, so that the bunch that results has the same shape as the contours of the bucket and has the specified *bunch width*. Changing *F_Dot* changes the shape of the bucket that the chopper is instructed to fill. It does *not* change the actual bucket, which is determined by the Rf voltages and the like. It just instructs the chopper to chop the beam differently.

F_Dot is $\frac{df}{dt}$, the rate of frequency change, or the acceleration rate. At high *F_Dot* the micropulses are chopped more narrowly at the beginning and the end of the macropulse than in the center. This results in a bunch that looks rounded in phase space. Notice how the bunch in figure 8.7 has rounded left and right sides.

The response time of the linac tank 9 current transformer is not fast enough to resolve the micropulses. For example, if the fast chopper is off, the current measured would be about 35 mA (3.5 V). If it is on, and chopping out half the beam, the current measured would be about 17mA. However, the response time *is* fast enough to see some structure within the linac pulse. For example, if the fast chopper chopped out more beam early in the macropulse than towards the end of it, the current measured during the early part of the pulse would be less than that measured at the end. The current in the Linac pulse as measured on the transformer would increase with time.

When *F_Dot* is high, the envelope of the injected beam as viewed on this current transformer will have a rounded appearance, since the micropulses are chopped more narrowly at the beginning and the end. When *F_Dot* is lowered, the width of the micropulses at the beginning and the end of the macropulse come closer to the width of the micropulses in the center. Consequently, the envelope of the injected beam on the tank 9 transformer looks more like a square pulse. More to the point, the amount of Linac beam that is injected for a given *bunch width* increases as *F_Dot* is decreased.

In theory, if the value of *F_Dot* is lowered below the actual $\frac{df}{dt}$, and the second harmonic cavities are off, this 'extra' beam, that comes in the form of wider micropulses at the beginning and end of the macropulse, will not be captured because it will not fit in the bucket. If the second harmonic cavities are then turned on, the bucket shape will be more *square* and some of this beam *will* be captured. Consequently, one way to tell that the second harmonic cavities are having a positive effect, is to lower *F_Dot*. If lowering

it increases *Booster Late*, then it probably is doing its job.

This last statement is contingent upon having a reasonably optimized high intensity setup in the first place. In such an optimized setup, when the second harmonic cavities are off, *F_Dot* will typically be set to about 30 *KHz/ms*. Scanning the *bunch width* one will find a value of it which maximizes *Booster Late*. Then *F_Dot* can be scanned to maximize *Booster Late* again. Probably, one will find that its initial value was close to optimum. If the second harmonic cavities are then turned on, and are setup correctly, one might expect that lowering *F_Dot* will then increase *Booster Late*. Under some circumstances, the optimum value of *F_Dot* with the second harmonic cavities on has been in the range of 0 to 10 *KHz/ms*. Typically, the optimum *bunch width* will also increase somewhat when they are on and setup well, from say 220° to 260°. These numbers all refer to the $h = 2$ Rf setup.

The Counterphasing Function

Figure 9.8 shows B , $\frac{dB}{dt}$, and the injection transformer across the cycle.¹⁶ Notice that beam is injected as $\frac{dB}{dt}$ is increasing. When the beam is injected, $\frac{dB}{dt}$ is about 30 *g/ms*, it then proceeds to rise to a maximum value of 90 *g/ms*, and stays there for much of the acceleration cycle. The vector voltage of A6 and E6 is typically at or near its maximum value until late in the cycle. As a consequence of the constant Rf voltage and the fact that the beam requires more of a kick per turn at higher $\frac{dB}{dt}$, the synchronous phase increases as $\frac{dB}{dt}$ increases. Since the 'zero-crossing' of the second harmonic waveform is ideally located at the center of the bunch, and the center of the bunch is at the synchronous phase, one might suspect that the phase of the second harmonic voltage wave relative to the accelerating Rf voltage wave would have to change as $\frac{dB}{dt}$ changes.

The level of the counterphasing function sets the phase difference between the 'accelerating' voltage wave and the 'second harmonic' voltage wave. If the bunch position on the accelerating voltage wave shifts in phase, the counterphasing level would need to shift to keep the second harmonic phase at the same place relative to the bunch.

The second harmonic cavities are most critical early in the cycle and at injection. They are used there to capture a larger beam and reduce space

¹⁶C. Gardner, HEP Setup Book II, pg. 47, 1997

8-May-97
9:39:01

REMOTE ENABLE

3
10 ms
0.50 V

2
10 ms
1.00 V

4
10 ms
0.50 V

GO TO
LOCAL

~19 TP
INJ-NORM

10 ms

← 8 ms

1 1 V DC
2 1 V DC
3 .5 V DC
4 .5 V DC

Ext10 DC 3.65 V 50Ω

500 kS/s

□ STOPPED

Figure 9.8: B and $\frac{dB}{dt}$ across the acceleration cycle, together with the normalized injection transformer. The $\frac{dB}{dt}$ signal being used is BXI.STA1.PS.V, which is really a M.M. power supply voltage. On this scale it is very close to the actual $\frac{dB}{dt}$. Similarly, the B signal is actually the main magnet current, BXI.MAIN.MAG.CUR. The intensity late in the cycle is about 19 Tp.

charge effects by reducing the beam's peak longitudinal charge density. The counterphasing reference signal is presently available in MCR. The glitch bit, BRF.PKR_TRG_C_PHS, used in the Heavy Ion cycle, can also be used. In some cases, it appears that the second harmonic cavities increase the Booster's late intensity on each cycle by several T_p , which results in an overall increase of 10-15 T_p on *Booster Late*.

When the counterphasing function is tuned to maximize the *amount* of beam that can be stored in the AGS successfully, the function that often results consists of two levels, and a shift between them. One level is before injection and one is after. The value shifts from one level to the next in about 1 millisecond right at injection. The function does not seem strongly dependent on the increasing $\frac{dB}{dt}$, since it does not typically ramp during the time that it is changing. The change in the synchronous phase from injection to the high $\frac{dB}{dt}$ is about 10 degrees. In the scheme of things, this may be too small of an effect to worry about. I don't have a good explanation for why 'the shift' is there. Figure 9.9 shows such a counterphasing function.

The Rf setup tends to change from year to year. There are indications that it will change again for the next run, and that A3 and B3 may not be used as second harmonic cavities.

9.3 Reducing Losses Due to Increased Emittance

9.3.1 'Limiting Apertures' Early in the Cycle

Since the emittance typically shrinks as the momentum increases due to adiabatic damping (section 5.1), so-called 'limiting' apertures are more important earlier than later in the acceleration cycle. The losses that occur due to the large emittance associated with high intensity happen because the beam scrapes somewhere in the machine. Naturally then, it's important to become familiar with the locations in the ring with '*apertures*' that have the greatest potential for causing beam loss, and to find out how to get around them. There are two such 'limiting' apertures that are significant enough to warrant special attention. Probably, the more important of the two is called the '*Beam Dump*'.

Figure 9.9: A typical counterphasing function used for the second harmonic cavities. Note that the *glitch bit* is positioned on the first line segment. This is typically set at about 1000 to 2500 μs after peaker, so that the shift occurs around injection time.

The Beam Dump and the Dump Bump (a.k.a. 'the Bump Dump')

The Booster was designed under the premise that it is better if all beam loss is localized, than spread out around the ring. Wherever beam loss occurs material becomes 'activated'. If all the activation is in one place, it will be easier to 'manage'. In the Booster that place is aptly named the *Beam Dump*. The Beam Dump was designed to be the location where beam that is 'lost' or 'not useful' will be 'dumped'. The Beam Dump is located in the D6 '*straight section*'. The Beam Dump itself is a large mass specifically designed to 'absorb' the dumped beam.

There are many examples of beam that is 'lost' or 'not useful'.¹⁷ Sometimes the beam is purposely not extracted from the Booster. For example, when performing a 'study', the Rf may be set to turn off before extraction time, but after the phenomenon being studied has occurred.¹⁸ The beam in the machine after the relevant time in the cycle has passed is not 'lost', but it is no longer of any 'use'. With the Rf off, it will spiral inward, and presumably dump into the Beam Dump. More importantly, during normal running, beam that is not successfully stored or captured should be dumped here, as well as beam that is lost due to stopbands, or just about anything else.

Even under the best circumstances, this beam loss does not all get dumped at D6. The Beam Dump was designed to be the Booster's outermost inside aperture. Whether it actually is, depends on the Booster's configuration, specifically the horizontal orbit. As long as it is the outermost inside aperture, much of this beam will be dumped there.

The BoosterLossMonitor program is used to see where in the ring the beam that is lost is 'dumped'. For example, one can look in the window during which injection occurs if one wants to see where the beam loss associated with injection is 'dumped'. Figure 3.4 shows the losses around injection time, which in this case is window 2. In the figure, there are peaks in the amount of beam lost just downstream of the foil (C5-C8), and around the

¹⁷No sarcasm intended.

¹⁸The Rf 'gap volts' can be 'turned off' at a particular time in the cycle using an SLD and a pushbutton at each of the MCR consoles labeled, 'BOOSTER ETO1'. When this SLD, BRFEARLY-OFF1.RT, is 'on', the Rf will shut off at the time in the cycle specified in this SLD when any of these pushbuttons is 'enabled'. This '*early turn off*' can be 'disabled' by pressing one of these buttons while it is 'enabled'.

dump (D5-D7). This is a typical loss pattern for injection, and shows that a lot of the beam loss does indeed occur at the dump. In general, the location where a peak in the amount of loss occurs will be near a 'limiting aperture'. At injection, the Beam Dump, and the injection area are ordinarily the 'limiting apertures' associated with the most beam loss. The loss monitor analog signals are used to look at the time dependence of losses more closely than can be done with the loss monitor program.

It is ironic that 'the dump', which exists to reduce 'the effects' of radiation, may actually *cause* beam that would otherwise survive to be lost since it is designed to be a limiting aperture. An orbit bump exists that can be used to bump the circulating beam away from the Beam Dump to prevent this kind of beam loss. When the orbit is bumped away from the dump, losses elsewhere will typically increase. The amplitude of the dump bump that minimizes the total amount of beam lost depends on how the rest of the machine is configured.

The dump bump is usually 'on' during high intensity running, and turning it off normally reduces the intensity significantly in the Booster. Turning it off will normally reduce *AGS CBM* as well, but not by as much as *Booster Late*. It is likely that the AGS intensity does not drop as much as the Booster's when the bump is turned off, because the Booster's acceptance in this case tends to decrease. Since the emittance is often limited by this acceptance, the emittance will also tend to be less, and the transfer efficiency will often improve. Turning the bump 'on' again will tend to increase the emittance at extraction, and thereby trade beam loss near injection energy for beam loss at extraction energy. This is not a good thing because losses at higher energy produce more radiation.

Yet, since the machine is typically optimized with the Dump Bump 'on', it doesn't necessarily follow that the Dump Bump is solely responsible for the increase in the acceptance (and intensity) observed when it is turned back on. The Dump Bump can certainly be used to *effect* the acceptance, but that does not necessarily mean it can *increase* it beyond what would be possible without it. It exists because of the recognition that the dump, being the outermost inside aperture, has the *potential* to cause beam loss. However, the fact that the Booster, optimized for high intensity, does run with it 'on', *suggests* that it allows one to increase the acceptance and reduce the total amount of loss, beyond what would be possible otherwise, or at least make it easier to do so.

Figure 9.10 illustrates the dramatic effect the dump bump *can* have on the injection efficiency.¹⁹ It shows the spiraling beam at injection with and without the bump. One of the reasons for the difference between the two current transformer traces is probably the increase in the Booster's acceptance with the bump on. However, as mentioned, the effect of the bump on the acceptance is highly dependent on how the rest of the machine is configured. For example, the value of *Peaker*, which effects the injection field, can be set so that the radius of the injected beam is towards the inside. Then, the beam will tend to scrape on the inside and a larger bump will be required to move the beam away from the dump. As a result, the injection efficiency will become very *dependent* on the dump bump. By lowering *Peaker* and adjusting many other parameters, it may be possible to configure the machine optimally with a much smaller Dump Bump, and so that it does not have a dramatic effect on the injection efficiency.

The difference in the two transformer traces is consistent with what one expects from a change in the Dump Bump amplitude since, one expects, that changing its amplitude will change the distance between the circulating beam and the closest inside aperture. Beam that is injected during the early part of the linac pulse will initially have a relatively large radius because the field is lower at that time. It may not scrape more in the case where the bump is off, because the dump is an *inside* aperture, and the beam is as far to the *outside* as it is going to be at this time. This agrees with the observation that the rate of rise in both of the traces, corresponding to the injection efficiency, is initially the same. However, the slope of the lower trace starts to drop, and eventually becomes negative before the slope in the higher trace does so. Presumably, this is because, as the beam spirals inward, it starts to scrape at the dump sooner when the bump is off. It does this because, when the bump is off, the inside aperture is closer to the spiraling beam. An analog signal from a loss monitor near the dump, say at D6 or D7, could be used to see the time dependence of losses at the dump to check this hypothesis.

Just as the optimum value for the dump bump at injection is affected by the injection radius, which is set by *Peaker*, its optimum value during acceleration is also affected by the beam's radius. If the radius early in the

¹⁹From ,H.M. Calvani-Garcia, HEP Setup Book I, pg. 49, 1996. In this case the Rf is off. If the Rf were on, the effect of the dump bump on injection losses could change considerably.

Figure 9.10: Spiraling beam with and without the dump bump. The two top overlaid traces show the beam at injection with the Rf off. The higher trace shows the normal running condition with the dump bump. The lower of the two traces is the transformer without the bump. The other traces are the radial average signal in both cases.

cycle is towards the outside, then the optimum Dump Bump amplitude, *for this radius*, will tend to be smaller. From the standpoint of localizing beam loss, the lower the bump's value the better.

However, if the emittance is nearly the same as the acceptance, which presumably, it often is early in the cycle, the 'optimum' radius, will be constrained to be near the center of the acceptance. This 'optimum' radius corresponds to the radial steering function that *'threads'* the Booster's acceptance with a minimum of beam loss. At this point, if the acceptance is made larger by increasing the bump's amplitude, the center of the acceptance will shift inward, since the extra room that has been created will be on the inside. Accordingly, the 'optimum radius' will shift further to the inside as well.

As you can see, when tuning the dump bump early in the cycle to reduce losses, one of the key parameters to adjust along with it is the radial steering function. The Dump Bump's amplitude and the radius are often adjusted in an attempt to minimize losses as viewed on a current transformer signal. It seems possible that if losses are reduced in this way it is because they are associated with the Beam Dump. However, the `BoosterLossMonitor` program could also be used to confirm that the losses one is trying to reduce are actually occurring at the Beam Dump. Certainly, the radius may not have been optimized for the initial acceptance, so changing it could reduce losses, and this might have nothing to do with losses at the dump.

For a given radius, the bump has an amplitude at which losses are minimized, finding the combination of bump amplitude and radius that minimizes losses the most is a major part of the problem. Further complications arise from the fact that the horizontal orbit, as determined by the horizontal orbit correctors and the tune, will also effect where the beam scrapes as it spirals inward. These parameters are also 'fair game' when trying to optimize the dump bump, but do not ordinarily seem as important as the radial steering.

Optimizing the bump and radius may be less complicated if the horizontal orbit is 'well-corrected', and correcting it will also tend to increase the acceptance. To start to correct it early in the cycle, one might work at a particular time *after* injection, say 5 to 10 ms after it, where losses are still large. One sets the ramp times up in the `BoosterOrbitControl` program so that the ramp starts *after* injection so as not to disturb injection (see section 5.8). If 'reasonable' orbits are obtainable through the `BoosterOrbitDisplay` program, the orbit correctors can be used to minimize the amount of orbit

distortion. Typically, the largest harmonic component of the orbit is the fifth, which can be minimized using a fifth harmonic dipole correction.²⁰

Figure 9.11 illustrates the effect of the dump bump on the orbit.²¹ This is a '*difference orbit*' that was taken a few milliseconds after injection. The orbit without the bump has been subtracted from the orbit with the bump to show the effect of it on the closed orbit. There is no bpm at D6. It's likely that the position at D6 is further to the outside than at D4 and D8 since the orbit in the vicinity of D6 is expected to look roughly like a half-sine wave with D4 and D8 on opposite sides of the peak.²²

The dump bump is produced by current in backleg windings at C7, D2, E1, and E4. In figure 9.11, you can see the effect of the kicks at these locations on the orbit. Unlike the extraction bump backleg windings, the current in each of these is not independently controllable. This is because there is only one power supply for all of the four windings. The current reference for the power supply is set through a function in `BoosterOrbitControl`. The power supply is bipolar, and the current is able to follow a reference function that changes sign during the cycle. Ordinarily, the current in the supply early in the cycle is about 20 A.

The dependence of the amount of total beam loss on the dump bump's amplitude decreases as the beam size decreases. Therefore, after about 20 ms of acceleration, its amplitude can be varied over a considerable range without an appreciable effect on the loss rate. Consequently, its reference can be set at zero during the middle and late parts of the cycle, and it ordinarily is set at zero here.

There is an analog signal for the power supply current called `BXI.DUMP_BUMP.I` available on Xbar. There is an SLD available for the on/off state (`BGN.DUMP_BMP.STAT`) which also has a current readback. The time the current is read is set with the real timeline event, `BGN.DMP_BMP.RDTIME`. The function also has a start time set with the real timeline event, `BGN.DMP_BMP.STTIME`. This time is rarely if ever changed. These SLDs, and a few

²⁰Note that the dump bump distorts the orbit, but this distortion is 'intended'. Therefore, it's often advisable to remove the data from the bpts inside the bump (at C8, D2, D4, D8, and E2) from the orbit, so that they do not effect the results of the FFT harmonic analysis that are being used to 'correct' the orbit. This data can be removed using the 'edit bpm status' item in the diagnostics menu of `BoosterOrbitDisplay`.

²¹C. Gardner, *HEP Setup Book I*, pg. 37, 1997.

²²It's not exactly a half sine wave because the orbit is modulated by the β function.

DATA: Current Orbit
 - FILE bump_off
 DATE: Fri Mar 21 17:21:41 1997
 PPM USER: Booster/AGS --- 2
 GAIN: 1
 DIFFERING GAIN SETTINGS:
 LLRF A2 A:0.1 B:0.1
 LLRF A8 A:0.1 B:0.1
 NUMBER OF GROUPS READ: 3
 NUMBER OF CYCLES: 5
 AVERAGING MODE:
 Average over all groups
 Display cycle 2
 NUMBER OF BUNCHES: 50
 BAD HORIZONTAL BPMS:
 A2
 A8
 C6
 BAD VERTICAL BPMS:
 C7

GROUPS ARE AVERAGED - CYCLE NO. 2:

• 20.002 ms

Figure 9.11: A 'difference orbit' showing the difference between the orbit with and without the bump. This particular bump moves the position at D6 to the outside. This is with 10 Amps in the bump's power supply. The positions are in millimeters.

others, are found under 'Booster/ring-Longitudinal/Dump'.²³

The dump bump can also be used with Heavy Ions. The amount of radiation that Heavy Ion beam loss causes is orders of magnitude less than the amount that Proton beam loss causes. Consequently, from a radiation standpoint, the Beam Dump is only important for Proton running. However, it can still be a limiting aperture for Heavy Ion running, and the dump bump can be used to move around it. In practice, the dump bump doesn't seem as important in reducing beam loss during Heavy Ion running, though it may have some small positive effect.

The dump bump is another reason why keeping the horizontal orbit 'corrected' while the horizontal tune is raised is more difficult than keeping the vertical orbit 'corrected' while the vertical tune is raised. There is another reason as well, it is called the *Slow Injection Bump*.

The Slow Injection Bump

The injection foil is inside the beampipe, and its outer edge is normally located about 1 inch from the beampipe's center. Although the circulating beam routinely passes through it while the injection bump is on, it also has the potential of passing through it after it has collapsed. Every time beam passes through the foil, it is scattered. This scattering will increase the angles of particles in the horizontal and vertical planes, and consequently tend to increase the emittance in both transverse planes. Particles will also lose varying amounts of energy as they pass through the foil. This will cause an increase in the beam's momentum spread. This increased momentum spread will tend to increase the longitudinal emittance, as well as the component of the transverse beam size associated with Dispersion. It is for these types of reasons that the foil is not put in the center of the beampipe.

The foil is not an 'impenetrable' limiting aperture like the Beam Dump, but it probably is something to be avoided if possible. The Booster beampipe is 6 inches wide. The usable beampipe is 4 inches wide at the foil, since the foil's edge is 1 inch from the beampipe center. The foil position is not fixed, it can be changed by selecting a different foil on 'the wheel'. Foils whose outer edges are further to the outside, are easier to hit with the incoming beam, and require a smaller fast injection bump, but it is harder for the circulating

²³Even though the dump bump is a transverse parameter.

beam to avoid them. Alternately, foils whose outer edges are further to the inside are easier for the incoming beam to miss, and require a larger bump, but it is easier for the circulating beam to avoid them.

What is the horizontal size of the circulating beam at the foil at injection energy? One approach to estimating this is to assume that the *horizontal* emittance is limited by the vertical acceptance.²⁴ Why would one think this is true? Well, the beampipe is circular outside the main magnet dipoles, but inside them, it is narrower vertically than horizontally. Consequently, the acceptance is smaller in the vertical plane than in the horizontal plane. During high intensity, measurements of the emittance in both planes using profiles on MW006 have found that they appear to be, at least sometimes, about the same.²⁵ The fact that they appear to both be about the same leads one to speculate that there is some kind of coupling between the motion in both planes. If there is some mechanism, like coupling, which keeps the emittance in both planes about the same, and the vertical acceptance is smaller than the horizontal acceptance, then one might expect both the vertical and the horizontal emittance to be limited by the vertical acceptance.

From the equation of motion,

$$y(s) = \sqrt{\frac{\epsilon}{\pi}} \beta(s) \cos(\psi(s) + \delta)$$

one can see that the maximum excursion a particle, of emittance ϵ , makes as it traverses the ring, $y(s)_{max}$, is $\sqrt{\frac{\epsilon}{\pi}} \beta(s)$. In the dipoles, the vertical half-width of the beampipe is 35 mm, if the vertical half-width here limits the size of the beam vertically, $y(s)_{max}$ should be the same as this width. The maximum β function in either plane is about 14 m, and it sometimes has that value in the vertical when the beampipe is narrow vertically. Therefore, using this formula, one gets that the maximum ϵ is about 90π mm mrad. Assuming that the maximum emittance in the horizontal plane is the same, this means the beam has a half-width of about 31 mm, or 1.2 inches, at the foil.²⁶

²⁴L. Ahrens suggested this approach.

²⁵See K. Zeno, Booster Tech. Note No. 227, "Coherent Tune Shifts in the Booster", pg. 13-14, September 6, 1995

²⁶ β_x at the foil is 10.5 m, using the same formula, $x_{max} = \sqrt{\frac{\epsilon}{\pi}} \beta(s)$, one gets $x_{max} = 31\text{mm}$.

This estimate is crude, but it can be used to gauge how important this concern about beam passing through the foil might be. If the circulating beam were centered in the beampipe at the foil, any particle whose maximum excursion, x_{max} is greater than 1 inch will pass through the foil. For a half-width of 31 mm in a gaussian beam distribution, about 20% of the beam would continue to pass through the foil after injection. To make matters worse, the beam that continues passing through the foil is composed of particles with the largest oscillation amplitudes, and these amplitudes will continue to increase due to scattering in the foil.

If the vertical acceptance limits the horizontal emittance, why be concerned about the Beam Dump as a limiting aperture? Well, it turns out that the dump is only a horizontal limiting aperture if the beam is roughly centered, or towards the inside, at the dump. If it is centered there, then the horizontal emittance (at injection) is limited to about 75π mm mrad. But, the aperture at the Beam Dump is big enough to let a 160π mm mrad beam through it. So, the beam dump is a concern because the beam has to be moved to the outside there, but it is not a concern because the beampipe is small there. Similarly, if the beam were centered in the space between the foil and the outside of the beampipe, the emittance could be as large as 230π mm mrad before the beam started to hit the foil. It is the vertical acceptance that is typically blamed for being the 'hard' limit on the emittance in both planes.

So, you can see that the potential exists here for beam loss and an increase in emittance due to the foil. It's possible to move the beam radius to the outside to avoid the foil, but this might bring the beam closer to outside apertures elsewhere in the machine. Incidentally, the innermost outside aperture is the F6 septum. The orbit correctors could also be used to move the orbit to the outside in the vicinity of the foil. But there is already a bump that exists for this purpose. This bump is called the *Slow Injection Bump*, and it is produced by current in backleg windings at C4, C8, and D1.

Unlike the dump bump, the currents in the slow injection bump windings can be varied independently. Also, unlike the dump bump, the references for the slow injection bump windings are DC setpoints in SLDs, not functions. They are also not bipolar. In fact, they trip off if you send them a positive setpoint. They have start and stop times that are controlled through SLDs.²⁷

²⁷The on/off SLD for the C4 power supply is BMM.TDHC4.STAT. Similar SLDs ex-

Analog signals are available through Xbar for the currents in each of these windings. This bump is considered to be most important early in the cycle, since the beam is most likely largest then.

There is a high level code through which one can adjust the bump's amplitude in a 'sensible' way. It can be found in the `BoosterOrbitControl` program. However, in many cases, the currents in each SLD are adjusted independently, using as a basis for the changes, improvements in the Booster's late intensity. The `BoosterOrbitControl` program also calculates the bump's amplitude and 'residuals' from the currents in the windings. Since the SLDs are often adjusted independently, these 'residuals' are typically large.

Although this bump doesn't really have a standard setup, it's not uncommon for it to move the orbit about 10 mm to the outside at the foil at injection energy. The bump is usually turned off in the middle of the cycle. Because of the large residuals that it typically has, a loss *can* occur when the bump shuts off. This is because the orbit can be quite different without it.

9.3.2 Radial Steering and Orbit Correction

In general, radial steering is used to center the E.O. in the middle of the horizontal acceptance. The orbit corrections are used to maximize both the horizontal and vertical acceptances. Both the vertical and horizontal orbit correctors are sensitive, particularly early in the cycle.

The `BoosterOrbitDisplay` program can be used in conjunction with radial steering to adjust the radius of the beam's orbit. It can also be used together with the orbit correctors to minimize distortions in the E.O, except where they exist to move the beam away from limiting apertures. It seems natural to think that the radius should be roughly centered, and the orbit distortions, not associated with moving away from limiting apertures, should be minimized. What 'minimized' means is not so clear because of the limiting apertures, and for similar reasons, on some level, the need for the beam to be centered is not clear either. Yet, it's probably a good idea to roughly center the beam in the beampipe, and to null out obvious orbit distortions, such as a large fifth harmonic component.

ist for the C8 and D1 power supplies. There are also current readback SLDs for each winding, they have '.SPRB' suffixes. The start and stop times are microsecond resolution delays from BT0 called `BIJ.SLO.TM.ON` and `BIJ.SLO.TM.OFF`. All these SLDs are in `Booster/Injection/Ring`.

The BoosterOrbitControl program also enables one to put bumps into the orbit around every ' β_{max} ' in the ring. Since the beam is largest where the β function is largest, these locations are particularly susceptible to beam loss. In the horizontal plane, β is a maximum in the even numbered half-cells, since this is where the horizontally focussing quads are located. Similarly, β is a maximum in the vertical plane at the quads in odd numbered half-cells.

These bumps are produced by the corrector nearest the β_{max} , and the two other correctors that are closest. For example, to make a bump at F6, the dipole correctors at F4, F6 and F8 are used. This is called a '3-bump', and in this case the bump would be in the horizontal plane. To make a '3-bump', a time in the cycle is specified, as well as ramps, just as in the case of the harmonic corrections.

For example, while keeping an eye on the losses on the current transformer, the orbit could be roughly centered using radial steering, and its distortions somewhat minimized using the harmonic dipole corrections. Then one might see that losses are high at D7 using the loss monitor program. One could then 'construct' a 3-bump at D7 with the dipole correctors at D5, D7, and E1, and scan the bump's amplitude, while looking at losses at D7.²⁸ One can often minimize losses at a particular location with this method. However, one sometimes finds that these losses are just shifted to another location in the ring, say E7 for example.

Figure 9.12 shows an orbit during normal running, about 1 ms after injection.²⁹ This orbit was acquired after the vertical orbit was 'flattened out' using 3-bumps. There are a few things that are noteworthy. First, the dump bump is clearly visible. Secondly, the beam is *not* centered horizontally, it is towards the outside. Thirdly, the orbit appears to be very well corrected in the vertical plane.

It is most important to correct the orbit at, and just after injection. In the horizontal plane this is difficult because the orbit, even outside the injection region, changes when the injection bump shuts off. Ideally, the way to get around this is by insuring that the injection bump does not have significant residuals. This can be done by looking at a difference orbit between the orbit when the bump is on, and the orbit just after it has collapsed. The smaller

²⁸In figure 3.4, there is a peak in the losses at A6, one might use a 3-bump centered there to reduce this loss.

²⁹K. Zeno, Booster Protons 4, pg.91-93, 1994.

Figure 9.12: A typical 'corrected' orbit just after injection.

the difference orbit, the smaller the residuals. If there are residuals, one can try to correct them using a facility in the BoosterOrbitControl program, but this may also change injection. Additionally, when one works with the correctors very close to injection, one has little choice but to adjust the orbit at injection as well.

Though one might think that the radius should be centered, as in LTB, the goal is *not* to center the beam, it is to minimize the losses. There are a variety of possible explanations why losses might be minimized when the beam is towards the outside. For example, its optimum value is *expected* to be dependent on the amplitude of the dump bump and/or the slow injection bump. One reason why the vertical orbit is flatter than the horizontal is because it has intentionally been flattened out. The same attention has not been given to the horizontal plane. However, in the horizontal, one *expects* that the optimum orbit may have distortions around the dump and possibly at the foil.

Unlike the vertical, the corrected orbit in the *horizontal* plane is not *expected* to be flat. It may seem that the orbit distortions in the horizontal have 'no business' being there. But remember that this orbit results from a compromise between a lot of effects. The radius may 'want to be' to the outside because of the dump bump's amplitude. Because the radius is to the outside, the effect of dispersion will distort the orbit.³⁰ These orbit distortions move the beam with respect to apertures. But it may be that this particular set of orbit distortions is consistent with keeping the beam away from the most significant 'limiting' apertures, such as the beam dump, the foil, and the F6 septum.

In particular, the orbit at the dump and foil is to the outside, and it's not clear where the orbit at F6 is. So, nothing in the orbit data contradicts this line of reasoning. Just like in LTB, the beam doesn't 'care' that it's not centered. What effects the beam are the things that it hits, and sometimes you can get a hint as to where those things are by looking at where the orbit is when the losses, on the current transformer or the loss monitors, are at their lowest.

There are also vertical and horizontal 4-bumps for the injection area con-

³⁰Compare this orbit with the orbit in figure 5.14, where the orbit distortions are thought to be primarily due to dispersion. Orbit distortions due to dispersion can be nulled out by adjusting the sixth harmonic dipole correction, even though the distortions are not due to dipole 'errors' per se.

structed from the orbit correctors. These are called 'injection correction bumps'. These 4-bumps allow the position and angle of the E.O. at injection to be adjusted. This is particularly useful in the vertical plane, since the fast injection bump does essentially the same thing as this in the horizontal plane. These bumps effect the matching of the incoming beam to the E.O. without having to change the incoming beam's trajectory.

Orbits taken at different times in the cycle, with the orbit correctors off, tend to overlay for the most part. This is thought to be because the dipole errors that exist at one main magnet field value, to first order, remain the same at other field values, except that their strength increases at the same rate as the main magnet field increases. Consequently, if the orbit is corrected at one momentum (or main magnet field), it can be corrected at other momenta by scaling the correction currents. It is typically corrected just after injection.

This scaling can be done automatically by selecting 'extrapolate' under the 'Correction Options' menu. The 'scaling' can be confirmed by looking at the current in the correction magnets. Its current should look like the current in the main field magnets. Additionally, if the orbit is corrected to move away from some aperture at one momentum, the amount of current required at a higher momentum to move the same amount away from the aperture will also scale.

Chapter 10

Proton Operation III: *Extraction*

Proton injection is markedly different from Heavy Ion injection, but such is not the case with extraction. The same devices are used to extract the beam for both Protons and Heavy Ions, and they are used in the same way. However, there is an additional 'constraint' placed on the way that Protons are extracted which has not typically been in place during Heavy Ion extraction.

This constraint is related to the way the Booster and AGS Rf systems are synchronized at Booster extraction. It modifies the way that extraction is 'optimized', and the space in which the Booster operates at and near extraction. This 'constraint' boils down to the requirement that there be a fixed relationship between the phase and frequency of the Booster and AGS Rf at Booster extraction. The process through which this condition is achieved is referred to as '*Synchro*'.

The need for *Synchro* stems from the fact that, for protons, we usually transfer beam from more than one Booster cycle into the AGS before 'accelerating' in the AGS. It is possible to have 'multiple transfers' of Booster beam within one AGS cycle without it, but not without much higher injection losses in the AGS.

Figure 10.1 shows what multiple transfers within one AGS cycle look like on the Booster and AGS current transformers.¹ In this case, there are four Booster cycles per AGS cycle. With each transfer of a '*batch*' of Booster

¹B. Tamminga, HEP Setup Book II, pg. 53, 1996.

beam, the intensity in the AGS jumps up. On the AGS current transformer, it looks as if the beam from one transfer is 'stacked' on top of the beam from the previous transfers.

From a radiation standpoint, the higher amount of beam loss associated with multiple transfers in the absence of *Synchro* is not a concern with Heavy Ions, but it is a concern with high intensity protons. *Synchro* is *needed* to reduce these losses during high intensity proton operation. It has also been true that 'multiple transfers' in one AGS cycle, as well as the increased transfer efficiency afforded by *Synchro*, have both been required to increase the AGS's intensity to meet the expectations of the High Energy Physics (Proton) experiments, whereas this has not been the case for Heavy Ion Physics experiments.

Multiple *batches* of *Heavy Ion* beam have also been transferred during one AGS cycle using *Synchro*. In fact, because of the 'improvements' in efficiency and beam quality that result, it is likely that it will be used during the acceleration of Heavy Ions for RHIC.

The details of the *Synchro* process are essentially the same whether Protons or Heavy Ions are considered. In addition, it has always been used with high intensity protons, but only rarely has it been used with Heavy Ions. So, it seems natural to describe it within the context of Proton operation.

There are other aspects of extraction that pertain only to Protons, but are not necessarily related to *Synchro*. For example, loss monitors are a much more useful diagnostic tool with protons. Also, there are phenomena, perhaps associated with high intensity, that are only a concern during proton operation. For example, the sensitivity of the transfer efficiency to the tunes during the middle and late acceleration cycle.

Because of the need for efficient multiple transfers, and therefore the need for *Synchro*, the AGS has been inextricably linked to the Booster. As a consequence, it is harder to find the line between Booster extraction and AGS injection. For example, adjusting the AGS Rf frequency during the injection porch changes the Booster Rf frequency at extraction. Since the AGS Rf frequency does not effect the Booster's extraction *field*, this change in the Booster's extraction frequency results in a change to the beam's momentum and radius at extraction. Before considering these relationships much further, it seems like a good idea to explain in some detail what *Synchro* is, and how it relates to the transfer efficiency.

Figure 10.1: The AGS current transformer, AXI.F15_IXFMR_NORM, and the normalized injection transformer in the Booster during the AGS injection porch (and early acceleration). In this picture there are four Booster cycles during the injection porch. During each Booster cycle, the beam is transferred, it disappears from the Booster transformer, and reappears on the AGS transformer. This display is from the labview program, 'Machine-Efficiencies'. This program can be used to measure efficiencies during the transfer process using the current transformers in the Booster and AGS. The x-axis is time in milliseconds from AT0, the y-axis is intensity in Tp.

10.1 Multiple Transfers, *Synchro*, and *Cogging*

It was mentioned in section 7.3.3 (on Heavy Ion extraction) that AGS injection is like Booster extraction except that it is reversed in time. Just as the beam is kicked *off* the E.O. to extract it from the Booster, it is also kicked *onto* the E.O. in the AGS to inject it. The beam that is already stored in the AGS passes through the injection kicker's aperture on every revolution. When the injection kicker fires, any stored beam that is passing through it, regardless of whether it is already stored, or just being injected, will be kicked. Herein lies the problem, to inject beam *into* the AGS, you have to kick beam *out* that is already there. Unless of course, there is no stored beam in the kicker's aperture when it fires.

The Booster was designed so that its 'central orbit', would be as close as possible to one quarter the length of the 'central orbit' in the AGS. For example, say the beam was extracted from the Booster at a momentum that placed its orbit at the center of the Booster's beampipe (i.e.- at the 'central orbit'). It would pass through BTA, and be injected into the AGS, at the same momentum. Say also, that the main magnet field in the AGS was adjusted so that the orbit of this injected beam passed through the center of the AGS beampipe. Then, if these conditions were met, and the Booster was designed perfectly, the revolution frequency in the AGS would be exactly one quarter of the revolution frequency in the Booster.

This beam, once injected into the AGS, would initially cover at most one quarter of the AGS's 'circumference'. If nothing were done to contain this beam, it would spread out, covering more and more of the AGS on every revolution, since it would be composed of particles moving at different velocities and with different path lengths. Eventually, it would be spread out evenly, in azimuth, about the AGS ring. By the time the next batch of Booster beam was ready for transfer, about 150 ms later, the beam previously transferred would be spread out evenly about the AGS. As the beam being transferred was kicked into the machine, the stored beam that happened to be passing through the kicker's aperture at the same time would be kicked. Kicking the stored beam generally results in beam loss and increased emittance.

However, consider a beam that is 'contained' within some fraction of the AGS. 'Contained' in the sense that at any instant there are some azimuthal

locations without beam. Then, every revolution there will be a time *window* during which no stored beam passes through the injection kicker. During this *window* the kicker could fire without disturbing the stored beam. Consequently, the Booster beam could be injected during this *window*. This assumes that the kicker current is able to rise to the required value, stay on at that value while the beam coming from the Booster passes through it, and fall back to nearly zero before the stored beam enters it again.

One way to keep the beam 'contained' during the time between transfers is to transfer it into Rf buckets in the AGS. If the beam is injected into Rf buckets, it will remain bunched, and not spread out. This is typically what is done.

Recall that for Booster injection, the fast chopper chops the beam into micropulses that are injected directly into the Booster's Rf buckets. Just before extraction, the beam is already in 'bunches'. These 'bunches' play a role analogous to that of an injected turn of micropulses. In the case of Booster injection, the time that a turn comes into the Booster is locked to the phase of the Booster Rf, so that the micropulses within that turn will be injected into the Booster buckets. In the case of AGS injection, the time that a *batch* of beam from the Booster comes into the AGS is locked to the phase of the AGS Rf, so that the bunches within that *batch* will be injected into the AGS buckets.

For Booster injection, the SLD *Chopper_Phase* is used to adjust the relative phase of the chopper and the Booster Rf, and hence the micropulse and the Booster Rf bucket. For AGS injection, the SLD, BRF.BSTR_AGS-PHASE, or *Booster_AGS_Phase*, is used to adjust the relative phase between the Booster and AGS Rf at extraction. In the Booster, the phase at which a bunch falls on the Booster Rf wave is locked to the synchronous phase through the phase loop. So, adjusting the relative phase of the Booster and AGS Rf waves adjusts where that bunch will fall on the AGS Rf wave when it is injected into the AGS. This, in turn, relates the Booster bunch to the AGS Rf bucket. An additional requirement is also needed to fix the phase of the Booster bunches to the AGS Rf wave, and it is that the kicker fire at the same phase of the Booster Rf wave on every cycle. This condition was also needed for 'unsynchronized' extraction. Recall that the *Cogging Box* in the low level rf system is used to generate kicker triggers that are 'locked' to the beam (or Rf) phase (see section 7.3.3).

In the past, the batch transferred to the AGS has contained more than

one bunch. Hence, it has been necessary that the spacing between bunches be the same as the spacing between Rf buckets in the AGS. In the case of $h=2$ in the Booster, it has been necessary that there be two Rf buckets in the AGS occupying one quarter of the AGS circumference. The Rf buckets occupy one quarter of the AGS circumference because that is the same as the length of the Booster circumference, which contains the two evenly spaced bunches. To transfer those bunches to the AGS and into AGS Rf buckets, the spacing between the AGS Rf buckets needs to be nearly the spacing between the Booster Rf buckets.

Since two buckets occupy one quarter of the AGS circumference, there will be eight buckets total in the AGS, two for each quarter. Consequently, the AGS has run at $h=8$. Since there are 8 AGS buckets, and two bunches per Booster batch, there will be four transfers required to fill all eight AGS buckets. I hasten to add that this is only one approach to containing the stored beam so that injected beam can pass through the kicker when no stored beam is passing through it. It has been used in the past, but it is likely that a different approach will be used in the future. It may be that the Booster will run with $h=1$, and the AGS with $h=6$. If the Booster only has one bunch, then the spacing between bunches in a batch is no longer a constraint. The process by which the batches from each Booster cycle are transferred into the empty buckets is called *Cogging*.

If the beam is to be injected into AGS Rf buckets, the AGS Rf frequency must be very close to a harmonic of the beam's revolution frequency in the AGS. Since the Booster is one quarter the circumference of the AGS, the revolution frequency in the Booster will be four times what it is in the AGS. But since $h_{AGS} = 4h_{Booster}$, the Rf frequencies in both machines, at transfer, need to be the same. Whatever the details of the setup, there needs to be a fixed relationship between the Booster and AGS Rf frequencies at transfer because there is a fixed relationship between the Booster and AGS circumferences, and because the beam that is extracted has the same momentum as the beam that is injected into the AGS.²

²This last condition, which may seem obvious, is not met in the case of Heavy Ions, because the Heavy Ion beam loses energy as it passes through the stripping foil in BTA. The situation there is slightly more complicated.

10.2 The Synchro Loop

As I've attempted to show, the phase of the Booster's Rf needs to be synchronized to the phase of the AGS Rf at extraction. During much of the cycle, the phase of the Booster's Rf is *not* synchronized to the AGS Rf phase. How then does it become synchronized? It becomes synchronized through the use of a feedback loop in the Booster low level Rf system. This feedback loop is called the Synchro loop.

The details of how the Synchro loop works are quite complicated.³ From the perspective of Operations, I think that at least some kind of qualitative understanding of what the Synchro loop does is helpful. The low level Rf system runs on the radial loop during much of the cycle. When the switchover to the Synchro loop begins, the beam control system is switched off the radial loop. The time that the switch occurs is set in BRF.SYNCH_LOOP.ST. This is a microsecond resolution delay from the gauss timeline event BRF.SYNCH_LOOP.STRT.GT.

BRF.SYNCH_LOOP.STRT.GT, or '*Synch Start*', is nominally set to the same field value as '*Extraction Start*'. Typically, BRF.SYNCH_LOOP.ST is set to be several hundred microseconds. Since '*Extraction Start*' occurs 5 ms before extraction, the Synchro loop comes on slightly less than 5 ms before extraction. Once it is on, the beam radius is no longer controlled by the radial steering function.

The Synchro loop is designed to regulate the phase of the beam with respect to the phase of some reference oscillator signal. The Synchro loop tries to make the phase difference zero by changing the Rf frequency. Without the action of the Synchro loop, if the frequency of the beam signal and the reference signal differ, the beam signal will constantly slide in phase with respect to the reference signal. So, with the Synchro loop 'closed', if the frequency of the beam signal were far from the frequency of the reference signal, the Booster Rf frequency would change in an attempt to make the phase difference zero, and this would drive the beam out of the machine.

The AGS frequency must remain essentially constant, because the beam on the injection porch is affected by it. On the other hand, the beam in

³The detailed description of how Synchronization and Cogging is accomplished is found in pages 74-101 of "J. Michael Brennan, *Rf Beam Control for the AGS Booster*, BNL-52438, September 26, 1994." Much of what follows results from my mangled interpretation of this material.

the Booster is accelerating up until extraction. Hence, the Booster's Rf frequency is changing, and is different than the AGS Rf frequency. How can the Synchro loop work to 'lock' the beam's phase to the phase of the AGS Rf if the frequency of the Booster Rf (or the beam) is different from the AGS Rf frequency and is changing? A signal derived from the Booster beam signal is made which has a similar frequency to the AGS Rf signal, and the phase difference between them is used as an input to the Synchro loop.

It is possible to make this signal, with Booster beam phase information and a similar frequency to the AGS, because one can anticipate what the difference between the Booster beam frequency and the AGS Rf frequency will be. With the Synchro loop off, and the radial loop on, the beam is accelerated to extraction. Since the Booster magnet cycle is quite reproducible from cycle to cycle, the difference between the Booster Rf (or the beam's) frequency and the AGS Rf frequency is a reproducible function of time within each cycle.

Using a Frequency and Time Interval Analyzer in the Booster Rf control room (in building 914), like the one at MCR3, the frequency during the last 10 ms of acceleration is recorded every $2\ \mu\text{s}$ and stored in RAM. This frequency is subtracted from a number called the 'target frequency', which represents the AGS Rf frequency, to obtain $\Delta f(t) = f_{\text{target}} - f_{\text{beam}}(t)$.

The frequency given by $f_{\text{mr}}(t) = f_{\text{beam}}(t) + \Delta f(t)$, called the moving reference frame frequency, is very nearly constant and very close to the AGS Rf frequency. The Rf signal derived from the beam and a synthesized Rf signal derived from $\Delta f(t)$ are added vectorially to get an Rf signal whose frequency is very close to the AGS Rf frequency and that contains the beam's phase information.

The phase difference between this 'moving reference frame' signal and the AGS Rf signal is measured, and used as the input to the Synchro loop. The reference input for the Synchro loop is zero. Once the Synchro loop is closed, the difference between the phase of the beam expressed in this moving reference frame and the AGS Rf phase is kept locked to the reference input (which is zero). This is true except for a constant phase shift between them which can be adjusted in the *Booster_AGS_Phase* SLD.

Once the Synchro loop is closed, there is a fixed relationship between the beam's phase and the AGS Rf phase determined by the value of *Booster_AGS_Phase*. At the moment of extraction, $\Delta f(t)$ is nominally zero, and the AGS and Booster Rf frequencies are very nearly the same. Hence, the phase and

frequency of the Booster Rf are 'locked' to the phase and frequency of the AGS when transfer occurs.

Turning on the Synchro Loop

Before running the machine with Synchro, extraction needs to be set up using the radial loop. Heavy Ion extraction is set up in this way. Typically, one tries to adjust the extraction radius so that it matches the radius when transfer was optimized. This radius is typically obtained from an orbit that was taken when the machine was optimized (say from Store or the morning numbers). In the past, the radius has often been slightly towards the outside, say +5 mm, at extraction. So, the extraction radius is adjusted using radial steering so that an orbit indicates that the extraction radius is about +5 mm.⁴ Normally transfer, including spiraling in the AGS, is initially set up while Synchro is off.

Since there will be more beam loss in the AGS if multiple transfers are used, the 'pre-Synchro' set up is usually not done with multiple transfers. Beam is requested from Linac on only one of the Booster cycles using the Superman program. Additionally, it is done at relatively low intensity.

It's also a good idea to make the extraction radius flat during the time that the Synchro loop will be on (i.e.- the last 5 ms of acceleration). This is done by making the radial steering function flat during this time. Once it seems like the extraction radius is correct, the extraction frequency can be measured. This frequency can be used as the target frequency for the 'frequency' table that will contain $\Delta f(t)$. Then a frequency table is made by the Rf people.

Afterwards, Synchro is enabled by turning on the *Synch_Start* and *Synch-loop_start* SLDs. This is typically done by the Rf personnel. At this point, extraction is occurring with 'Synchro'. The trigger *Cogging_Arm* also has to be changed from 5000 μ s to 100 μ s when Synchro is on.

⁴Figure 7.26 is a fairly typical proton extraction orbit. Though, the radius here is not towards the outside.

10.3 Cogging

Recall that for Heavy Ions, the '*Cogging Box*' in the low level Rf electronics was responsible for sending out kicker triggers so that the kickers would fire at the right time (or equivalently phase) with respect to the bunches in the Booster. Since the bunches in the Booster are synchronized to the AGS Rf buckets, if they are kicked out at the same phase on every transfer, they will enter the AGS at the same phase with respect to the AGS Rf buckets on every transfer. Consequently, *Booster_AGS_Phase* can be used to match up the Booster bunches with the AGS buckets.

This works fine for a single transfer, but when multiple transfers are involved it is not sufficient because the beam may be injected into buckets that are already full. Somehow, whatever generates the kicker trigger has to keep track of which buckets are full so that it can fire, not only when the phase is correct, but also when the bucket spinning around the AGS is in the appropriate spot.

The *Cogging Box* is not only responsible for locking the time that the kicker fires to the position of the bunch in the Booster, it is also responsible for sending the trigger so that the bunch is injected into the correct AGS bucket. The metaphor of 'cogging' is used for this kind of transfer. For $h=2$ in the Booster, and $h=8$ in the AGS, one can picture a two gears, each one representing a machine. The Booster's 'gear' has two gear teeth or 'cogs', the AGS gear has eight cogs. These gears are locked together, and each cog represents a bucket. These cogs can be numbered.

Before the first transfer, AGS cogs 1 and 2 are locked to the Booster cogs. When the first batch is transferred, the two of the AGS buckets, corresponding to cogs 1 and 2, are filled and the AGS gear advances by two cogs. Now cogs 3 and 4 are locked to the Booster cogs. When the second batch is transferred, the buckets corresponding to cogs 3 and 4 are filled, and the AGS gear advances to the next two cogs. Now cogs 5 and 6 are locked to the Booster cogs. The next batch of Booster beam is transferred into the buckets corresponding to cogs 5 and 6, and the AGS gear advances so that the last two cogs, 7 and 8, are locked to the Booster cogs. The last batch is transferred into the buckets corresponding to these two cogs. The *Cogging Box* keeps track of which buckets have been filled, and generates the kicker triggers at the appropriate times.

The kicker triggers are generated at the correct phase, and for the appro-

priate bucket during the revolution of the AGS beam right after the trigger *Cogging_Arm* occurs.⁵ This is also the case for Heavy Ions. As with Heavy Ions, certain supercycle events have to be in place (BT1 and BT4) and two other SLDs have to be set properly. The SLD, BRF.COGGING.RESET, resets the Cogging Box for the next supercycle. It is a delay (in microseconds) from the supercycle event BT4. The other SLD, BRF.AGS_RF_READY, puts the Cogging Box in the initial state, waiting for the first transfer to occur. It is delayed from the supercycle event BT1.

10.4 ΔB^* and the Extraction Radius

10.4.1 Finding the Extraction Radius

When Synchro is 'on' there is no straightforward way to adjust the radius at extraction. Consequently, one has to adjust the parameters involved in extraction differently. For example, since the radius is no longer easily controllable, one may tend to rely more heavily on the extraction bump to optimize extraction. It's true that extraction is initially set up as in the Heavy Ion case with Synchro 'off', but when this is done, it is at low intensity, and is typically done only as a 'first pass'. Once Synchro is turned on, it stays on, essentially for the entire run. Extraction, as well as transfer, are optimized while Synchro is 'on'.⁶

With the radial loop closed, variations in the radius at extraction do not *in principle* occur because the radial loop keeps the radius 'fixed' to the reference provided by the radial steering function. Yet when the Synchro loop is closed, the machine is no longer on the radial loop, and so the radius is no longer 'fixed' to any reference function. What is 'fixed' is the extraction frequency, which is determined by the 'frequency table' (i.e.- $\Delta f(t)$) because of the action of the Synchro loop.⁷ It turns out that the extraction radius is *uniquely* determined by the extraction frequency and field.

The radius is a critical parameter in the extraction setup, and variations in it can cause problems. Sometimes, when extraction or transfer has

⁵For unsynchronized extraction, this trigger is delayed from *Extraction_Start* and is set to 5000 μs , but for synchronized extraction it is not delayed from this trigger.

⁶When Synchro is first turned on, it should not effect the extraction efficiency.

⁷This assumes that the AGS Rf frequency or the 'target frequency' have not changed.

deteriorated, it may be because the extraction radius has changed for some reason. Therefore, it's important to know what the extraction radius is when the machine is running well, so you can check to see if it has changed. A 'nominal' extraction radius can be measured by taking an orbit at extraction when extraction and transfer look 'nominal'. The crudest way to gauge how well extraction is set up is to look at the transfer efficiency (i.e.- *AGS CBM/Booster Late*). Other, more 'analytical' methods will be described later.

Figure 7.26 shows a typical extraction orbit optimized while Synchro is 'on'. As mentioned, it is often the case that the radius is set to be 5-10 mm to the outside of beampipe center, even though this is not the case here. The 'radius' that I'm speaking of is actually the $\cos 0\theta$ harmonic component of the orbit obtained using the 'quadratic fit' option in the BoosterOrbitDisplay program. Due to the conventions that are used, the average deviation of the beam *at the bpms* from the beampipe center is actually one-half of the amplitude of the $\cos 0\theta$ component. It is the amplitude of the $\cos 0\theta$ component that is typically measured to be $+5 - 10$ mm.⁸ This extraction radius measurement is taken every day as part of the *Morning Numbers*.

It's not really important which number you use for the extraction radius. What's important is that you know what the number is when extraction is OK. Then you will know what to compare the radius to when a problem arises. Another way to measure the value of the radius is with the radial average signal. It's not uncommon to be unable to obtain a reasonable orbit using the bpms, but one typically has a radial average signal available. Figure 10.2 shows the radial average signal during the latter part of the cycle.⁹ The voltage level of the radial average signal could be measured at extraction to gauge the extraction radius.

This picture also shows the switchover from radial loop to the Synchro loop. The third trace from the top is the *Synchro Phase Error* signal. It shows the deviation of the input phase (the difference between the phases of the moving reference frame and AGS waves) from the phase reference. Before the Synchro loop is closed it oscillates rapidly. Once the Synchro loop is

⁸More precisely, the beam *radius* is the average position of the beam around the ring with respect to the beampipe center, which is not the same as the average position of the beam at the bpms. This is because the average value of the dispersion at the bpms is not the same as the average value of the dispersion.

⁹Greg Marr, Brian Bukala, *HEP Setup Booster-AGS*, pg. 17, FY 1997

Figure 10.2: Extraction with Synchro. The top trace is the AGS current transformer (AXI.F15_IXFMR_NORM), the second trace is the radial average signal (BXI.RFE6.RADIAL_AVG), the third trace is the Synchro error signal (BXI.SYNCHRO_ERROR), and the bottom trace is the Booster circulating current transformer (BXI.CIRC_XFMR_NORM). The arrow indicates where the scope trigger occurs. The trigger is *Extraction_Start* on cycle 4. The sweep speed is 2ms/div.

closed, the oscillations in the *Synchro Phase Error* signal stop. At this point the beam's phase is essentially 'locked' to the AGS Rf phase. Notice that the radial average changes slightly when the Synchro loop is closed. There can be oscillations on the Synchro error signal after *Synch_Start*. These are associated with making the switchover to the Synchro loop stable. They are typically of a much slower frequency than the oscillations visible *before Synch_Start*.

10.4.2 Measuring ΔB^* with the Gauss Clock

In section 10.2 it was mentioned that stable extraction (and transfer) is based on the premise that the magnet cycle is reproducible during the time in the cycle that Synchro is 'on'. This is simply a result of the fact that the radius depends on the field. When Synchro is 'on', the Rf system does not 'correct' for variations in the magnet cycle. Extraction is very sensitive to the value of the extraction radius. The extraction efficiency will vary if the field varies, as well as the way the beam comes out of the machine, and the extraction momentum. When the radial loop is on, the extraction efficiency and trajectory are not nearly as sensitive to variations in the extraction field since the radius is held constant.

There are different ways in which the extraction field varies. For example, the extraction field could be a function of the Booster cycle number within each supercycle. In that case, the extraction field on cycle 2 might be higher than the extraction field on cycle 5. Alternately, the extraction field may have changed by a similar amount on all the cycles within a supercycle from what it was the other day.

The primary tool for measuring variations in the extraction field is the gauss clock. In section 2.2.3 the Gauss Clock was described. The Gauss Clock up and down counts are routed to MCR. These Gauss Clock signals can be attached to scalers like the ones used for the intensity measurements. The scalers provide a 'digital readout' of the number of Gauss Clock counts over some time interval. The Booster Gauss Clock counts have been routed to a scaler at MCR 4. The interval over which the number of counts is measured is determined by the start and stop triggers that are set through Xbar.

As described in section 7.3.5, extraction occurs 5000 μs after *Extrac-*

tion_Start.¹⁰ Since *Extraction_Start* is a 'field' trigger, it will occur when the field reaches the value specified in the SLD. The value of *Extraction_Start* plus the amount of gauss clock counts that occur from *Extraction_Start* to extraction, ΔB^* , is a measurement of the field at extraction.

Extraction_Start is generally known, what one needs to find out is ΔB^* . The scalers are used for this purpose. The start trigger for the scalers is *Extraction_Start* (called B.XTRCN.START.GU on Xbar). The stop trigger for the scalers is typically *Extraction_Start* + 5000 μs .¹¹ The number of up counts minus down counts during this interval is ΔB^* . The number of down counts is typically negligible.

Xbar allows one to specify during which cycle(s) and user(s) the triggers will occur. Since the value of *Extraction_Start* is the same on all cycles for the same user, any variation in ΔB^* accross cycles corresponds to a change by the same amount in the extraction field. Just as with Heavy Ions, there is normally a *dummy* magnet cycle before the magnet cycles with beam within a supercycle (see section 7.3.2). If one measures ΔB^* on this cycle, one will typically find that it is significantly less than the values on subsequent cycles within the same supercycle. This is one of the main reasons why a dummy cycle is used. If there was beam on that cycle, its radius (and momentum) would be significantly different than on the othercycles. When multiple transfers are used, it is clearly important that the Booster cycles within each AGS cycle be as similar as possible.

ΔB^* for each of the cycles with beam is measured as part of the morning numbers each day. The variation from cycle to cycle within a supercycle will in general be associated with a variation in radius on each cycle. If the radius varies from cycle to cycle within each supercycle, losses and extraction efficiency will vary for each cycle as well. The extraction momentum will also be different on each cycle.

How sensitive is the radius to variations in the extraction field? There is a general equation which relates small variations in the field to the variations in the radius at a fixed frequency. For a fixed frequency, the change in radius, ΔR , due to a change in field, ΔB , is described by the following *differential*

¹⁰Actually extraction occurs a microsecond or so after that due to cogging.

¹¹The trigger B.COGGING_ARM.RF could also be used as the stop trigger. It is essentially the trigger generated by the Rf system that is sent to the kickers to make them fire. Hence it is locked in phase to the beam and the kickers, and occurs *after* the *Cogging_Arm* event occurs.

relation,¹²

$$\frac{\Delta B}{B_o} = (\gamma^2 - \gamma_{tr}^2) \frac{dR}{R_o}$$

where B_o is the initial field, R_o is the Booster's radius, γ is the relativistic gamma, and γ_{tr} is read as 'gamma transition', and has a value of about 4.9 for the Booster. If $\gamma < \gamma_{tr}$ then a positive change in the field (i.e. $\Delta B > 0$) causes the radius to move inward (i.e. $\Delta R < 0$). If $\gamma > \gamma_{tr}$ then the opposite is true. In the Booster, the beam has never been accelerated to an energy where γ even comes close to γ_{tr} , so γ is always less than γ_{tr} .¹³

At extraction energy in the Booster, $B \approx 6400$ g and $\gamma \approx 3.1$. The central radius, R , of the Booster is 32.114 m. Plugging these numbers into the equation above, one finds that an increase of three gauss in the Booster's extraction field corresponds to a radial change of about -1 mm. Typically, ΔB varies by less than about 1 gauss (10 counts) over the cycles with beam, corresponding to about a millimeter of movement.

The unnormalized horizontal emittance that contains 95% of the beam at extraction is about $35 \pi \text{ mm mrad}$.¹⁴ Since the β function has an average value of about 7.5 m, the half-width of the beam, using $x = \sqrt{\frac{\epsilon}{\pi} \beta}$, is on the order of 25 mm, which is large compared to 1 mm. So, even if the beam was already scraping on some aperture, the *percentage* increase in beam loss if it moved by a millimeter would not be large. However, it might be measureable with the loss monitors. Variations in the amount of beam loss associated with extraction as a function of the cycle number within a supercycle are often observed.

Changes in the extraction field don't always come from where you might expect. The Booster Main Magnet Power Supply is quite reproducible, yet the extraction field *can* change significantly. Strangely enough, one of the most common reasons for a change in the extraction field is a shift in the requested tunes after *Extraction_Start*.

¹²In this equation, as well as the other *differential relations* to follow, ΔB represents the change in the extraction field. *not* the amount the field changes in the 5 ms leading up to extraction called ΔB^* .

¹³The kinetic energy, $E_k = (\gamma - 1)mc^2$, where $\gamma = \gamma_{tr}$ is called *Transition energy*, and its value is related to the details of the machine lattice. In the AGS, the beam is accelerated through *transition energy*.

¹⁴This was estimated using the profile width on MW006 in BTA. See: K. Zeno, Booster Tech. Note No. 227, "*Coherent Tune Shifts in the Booster*", pg. 14, September 6, 1995.

10.4.3 ΔB^* and the Tune Functions

If the requested tunes are shifted from one value to another at a time in the cycle after *Extraction_Start* and before extraction, one often finds that ΔB^* , as measured by the gauss clock, will change. The conventional explanation for this observation has to do with induced voltage on the main magnet windings. The tune quad windings are *trim* windings around the main quadrupoles (see section 5.4.1). When voltage is applied to these windings it induces a $\frac{dB}{dt}$ in the quadrupoles. This $\frac{dB}{dt}$ induces a voltage on the main windings on the magnet, which are in series with windings around the other main magnets. Consequently, the $\frac{dB}{dt}$ of the 'main magnet' changes.¹⁵

In figure 5.4, there is a shift in the tune functions at around 50 ms from BT0. Introducing a shift of the magnitude shown in this figure after *Extraction_Start*, but before extraction, in a function that was initially flat there, can change the extraction field by several gauss (corresponding to several tens of gauss clock counts). This would cause the extraction radius to shift by a few millimeters.¹⁶

On the other hand, the effect of shifting the tune before *Extraction_Start* on the extraction field is typically negligible. Since *Extraction_Start* is a field trigger, it occurs at the *time* the field reaches the value specified in the SLD regardless of what time that is. A change in $B(t)$ prior to *Extraction_Start*, such as that associated with a tune shift before it, will tend to change the time that it occurs, but not the field at which it occurs. Consequently, if the tunes are changed *before* *Extraction_Start* occurs, *Extraction_Start* will still occur at the correct field. The effect on ΔB^* will generally be small since the tune is not shifted during the period that ΔB^* is measured, and so extraction will occur at nearly the same field.

As mentioned in Chapter 9, the transfer efficiency is a function of the tunes during the middle and late parts of the acceleration cycle. It is also true that the transfer efficiency is a function of the tunes at extraction, but

¹⁵The induced voltage is described by Faraday's law, $\nabla \times E = -\frac{1}{c} \frac{\partial B}{\partial t}$.

¹⁶An effect reminiscent of this one can occur when the tunes are changed *before* injection. The average $\frac{dB}{dt}$ can change over the interval between BT0 and *Peaker*, causing the time that the field reaches the value specified in *Peaker* to change. Since many parameters follow functions of time, most notably the tune functions themselves, their values at injection can change because the *time* that injection occurs changes. This hypothesis is motivated by the observation that changing the tunes *before* injection time, but not at injection *time*, can have a large effect on the injection efficiency.

for different reasons. Changing the horizontal tune at extraction changes the extraction bump, the bump's residuals, the E.O., and the trajectory coming out of the machine. Similarly, in the vertical plane, changing the tune, changes the vertical E.O., and the vertical trajectory coming out of the machine.

Recall that the transfer efficiency can also be altered by adjusting the tunes *before* *Extraction_Start* (see section 9.2.1).¹⁷ However, this dependence is most likely due to the effects of stopbands and other high intensity phenomena. For example, this dependence could, in part, be associated with a phenomenon related to high intensity called a *Coherent Instability*.

Since the reasons for the tunes being what they are at extraction and during late acceleration are different, there is little reason to think that their optimum values should be the same. Hence, one would expect in an optimized machine, some type of shift in the requested tunes just before extraction. The shifts visible in both planes in the high intensity tune functions shown in figure 5.4 are such shifts. In both planes, the requested tunes during middle and late acceleration are generally lower than they are right at extraction. Hence, it simplifies matters if the 'acceleration' tunes are shifted to their extraction values *before* *Extraction_Start* occurs.

A change in ΔB^* can also be observed when the extraction bump is changed. In section 7.3.5 it was mentioned that the $\frac{dB}{dt}$ of the main magnet effects the current in the backleg windings associated with the extraction bump. Similarly, the voltage across the backleg windings causes a $\frac{dB}{dt}$ that induces a voltage in the main magnet windings and changes its $\frac{dB}{dt}$. Therefore, changing the backleg winding currents changes ΔB during the interval from *Extraction_Start* to extraction, just as changing the tunes then does. The difference in this case is that the extraction bump is only on during this time. In practice, it seems that this effect is less important than the effect of tune shifts on extraction.

10.5 ΔB^* and the Extraction Momentum

When Synchro is off, variations in the extraction field still effect the extraction momentum, but the relationship between the two is different. When

¹⁷When this data was taken the tunes after *Extraction_Start* were intentionally kept the same because of the effect a tune shift during this time can have on ΔB .

Synchro is off, the radial loop is on, and extraction occurs at a fixed radius. For a fixed radius, the change in momentum, Δp , due to a change in field, ΔB , is described by the relatively simple relation,

$$\frac{\Delta p}{p_o} = \frac{\Delta B}{B_o}$$

where B_o is the initial field, and p_o is the initial momentum. This *differential relation* comes from taking differentials in the equation which defines *rigidity*, $p = qB\rho$, and dividing the result by the equation itself. A change of a part in one thousand in field results in the same fractional change in momentum. At extraction, $B \approx 6400g$ and $p \approx 2.7 \text{ GeV}/c$. So, an increase of 1 gauss in the extraction field results in an increase in the momentum of 0.42 MeV/c.

How does this compare with the case where the frequency is fixed, not the radius? For a fixed frequency, the change in momentum, Δp , due to a change in field, ΔB , is described by,

$$\frac{\Delta p}{p_o} = \frac{\gamma^2}{\gamma^2 - \gamma_{tr}^2} \frac{\Delta B}{B_o}$$

An increase of 1 gauss in the extraction field results in a change in the momentum of -0.28 MeV/c. The momentum change in either case is of the same order of magnitude, but is opposite in sign. A higher field at fixed frequency decreases the momentum.

As found in the previous section, when Synchro is on, and the field is increased, the beam moves to the inside. So, if the frequency is fixed, and the field is increased slightly, the beam moves to the inside and its momentum is reduced. Table 10.1 summarizes these relationships.

Extraction Field	Synchro Loop	Radial Loop
$\Delta B > 0$	$\Delta p, \Delta R < 0$	$\Delta p > 0$ $\Delta R = 0$
$\Delta B < 0$	$\Delta p, \Delta R > 0$	$\Delta p < 0$ $\Delta R = 0$

Table 10.1: The effect of variations in the extraction field on the extraction momentum and radius. The effect depends on whether the Synchro or radial loop is used at extraction.

Small changes to the radius are probably more important than small changes to the momentum as far as extraction efficiency is concerned. However, the changes in momentum are very important as far as injection into the AGS is concerned. They could also effect the transport through BTA significantly. The *transfer* efficiency is effected by such changes. It's important to note that changes to the extraction setup, such as a change in the extraction radius or field, not only effect extraction, they effect transfer, due largely to the change in momentum that results.

10.6 A Survey of Extraction Losses

Particularly with high intensity, the transfer efficiency is prone to vary considerably from cycle to cycle, or supercycle to supercycle, in an apparently unpredictable way. One of the major tasks we are faced with is isolating the cause of these variations so that it can be fixed. In section 9.2.1, the transfer efficiency was defined, and the components that go into its definition were briefly mentioned. Clearly, variations in the way beam is extracted, or how much beam is extracted may cause variations in the transfer efficiency. Often what happens *downstream* of the Booster can be used to diagnose what is happening in the Booster. For example, a position variation on multiwires in BTA often results from variations in the way the beam comes out of the machine.

Finding the causes for this 'unstable transfer', as it is commonly called, can be a good way to gain an understanding of how the machine runs. For example, the correlation between the position fluctuations at MW060 with a variation in the amplitude of an extraction bump backleg winding analog signal, tells one something about the importance of that particular backleg winding.

The causes of transfer problems fit into a few major categories: 1) Power supplies that are for some reason not doing what they have been instructed to do, 2) Problems associated with the process of Synchronization, and 3) Variations in the *quality* of the beam that is extracted.

10.6.1 Diagnosing Problems with Power Supplies

A problem with a power supply is usually easier to find when the supply fails completely, than when it is intermittent. Sometimes the beam will not transfer for a cycle or two, then transfer will return to normal, only to fail again at 3 o'clock in the morning. This intermittent kind of problem can also occur for supplies not associated with extraction, but for some reason extraction supplies seem particularly susceptible to this type of failure. The symptoms associated with power supply failures depend, of course, on the power supply.

The F3 kicker

In the past, there have been quite a few problems with the F3 kicker. However, the power supply has been replaced, so there may be fewer problems in the future. Often the problems are not associated with the power supply itself, but with the high voltage cables and the electromagnetic noise generated when the kicker fires.

For protons, the kicker generally runs at about 35 kV, its highest possible voltage, and so it runs at the edge of its capacity. A complete failure will result in no beam being extracted from the machine. As viewed on a current transformer, the beam will accelerate to extraction time, but not get kicked out. Therefore, a scope triggered by the *Cogging_Arm* trigger will continue to show bunches on the wall current monitor after the trigger has occurred. Additionally, if one triggers the scope with *Cogging_Arm* one may find that the kicker current pulse is nowhere to be found.

In general, this will only be true on the first cycle with beam within an AGS cycle. On subsequent cycles, the beam will be lost *before* extraction time, but after the Synchro loop comes on. Before beam enters the AGS, the Rf frequency is set to a predetermined value. The extraction frequency of the first transfer is determined by this frequency. Several milliseconds after the first transfer, *Beam_Control* in the AGS is typically set to turn on. At this point the phase loop comes on.

In the AGS, a *frequency loop* comes on at or just after *Beam_Control* is turned on. This *frequency loop* is like the radial loop, except that the beam follows a *frequency steering* function, not a *radial steering* function.¹⁸

¹⁸Just as the radial loop will not work until *Beam_Control* is on, the frequency loop

Once *Beam_Control* is 'on' in the AGS, the AGS Rf tries to match the phase of the beam. If there is no beam, the AGS Rf frequency will go 'haywire' as the phase loop tries to match the phase of a beam that does not exist. Therefore, once *Beam_Control* is 'on', there has to be beam in the AGS, for the AGS Rf to have the frequency set by the frequency steering function. If it does not have that frequency during the time that the Synchro loop is on, the AGS Rf frequency will be very different from the moving reference frame frequency. The Booster Rf frequency will change to try to match the AGS and moving reference frame phases which are now *sliding* with respect to each other because of the frequency difference (see section 10.2). As a result, the beam's radius will change, usually resulting in the beam being lost on the beampipe. So, if beam is not injected on the first transfer, beam will not survive to extraction on subsequent transfers because the AGS Rf frequency will be wrong. That is why, if the F3 kicker doesn't fire on the first transfer, beam will be accelerated through extraction on the first transfer, but not survive to extraction on subsequent transfers.

This behavior, associated with the action of the Synchro loop, can prevent the root cause of the extraction problem from being discovered. Often Synchro itself will be blamed because it appears to be malfunctioning. The reason why the beam is lost in the Booster is because the beam is missing in the AGS. One can set *Beam_Control* in the AGS to turn on *after* the last transfer to prevent this type of behavior from clouding the issue.¹⁹

If the F3 kicker trips off, no beam will be injected into the AGS, and an alarm for it may appear on the ADT alarm screen. Secondly, if the Booster intensity is high a 'chipmunk' radiation interlock may also occur. This radiation interlock will close the LTB beamstops, and prevent beam from being injected into the Booster. Typically, the chipmunk above the F6 septum (NMO58) will interlock, or at least alarm, due to the increased radiation levels there. These radiation levels may be particularly high if the kicker stays on, but the current is not at setpoint. Since, in this case, the beam will be kicked into the metal part of the septum, and not the aperture.

The loss monitors in the extraction region will also indicate much higher

will not work either. Radial and frequency control are mutually exclusive. The SLD 'ARF.FREQ_LOOP.ST' is used to turn on the frequency loop. It is located under "AGS/Rf/llrf/Feedback_loops".

¹⁹The SLD is located in "AGS/Rf/llrf/Feedback_loops" and is called ARF.BEAM_CONT.ST.

losses if the kicker is not at setpoint, or fires at the wrong time, and the horizontal position on multiwires in BTA will shift. Fortunately, the hardwired current signal for the F3 kicker has fine enough resolution to easily detect significant variations in the kicker current. Figure 7.21 shows the kicker current signal. Variations in its amplitude may be correlated with the phenomena just described. Additionally, one can look at the wall current monitor, together with this signal, and check for variations in the time that the kicker fires with respect to the beam. As mentioned in section 7.3.6, this type of variation is associated with the cogging process.

There is a voltage signal available on Xbar that is often helpful in diagnosing problems with the kicker. A capacitor bank discharges when the kicker fires and charge flows through the kicker magnet. Figure 10.3 shows the cap bank voltage signal together with the normalized injection current transformer.²⁰ This figure shows how the cap bank charges up when it is working normally. Notice that the charging takes much of the Booster cycle.

A number of different problems can be detected with this signal. When the kicker discharges its voltage falls to zero. This normally occurs at extraction. However, it can sometimes discharge before extraction, and this can be seen easily on this signal. It may also be missing a discharge trigger, in which case it will not discharge at all. In that case, the signal would not fall to zero at extraction. It may also not charge at all, this behavior could indicate a short somewhere. Often this short is associated with the high voltage cables, and may be inside the ring enclosure.²¹

The F3 kicker setpoint, as well as the F6 septum setpoint, are not controlled by SLDs. The setpoints of each Booster cycle within a Booster user can be controlled separately. Recall that SLDs allow one to control the setpoint on a particular user, but the setpoint on each cycle for that user must be the same. CLDs, or Complex Logical Devices, are used to control the setpoints for different cycles within a user.²² CLD setpoints are controlled using a *CLD pop-up window* that can be brought up by clicking the middle mouse button on the setpoint field associated with the CLD on spreadsheet. For example, the F3 kicker setpoint CLD is called 'BXT.KDHF3.SP_CLD', and

²⁰Greg Marr, FY '97 HEP Store, section 4.1.3

²¹The time that the F3 kicker cap bank starts to charge up is set in the SLD 'BXT.F3&A5_ST_CHRGE', a microsecond resolution delay from 'BXT.914.PRE_DELAY', which is a microsecond delay from Booster prepulse (BPP).

²²The A5 kicker and L20 septum in the AGS have CLD setpoints as well.

12-Jun-97
6:03:41

Figure 10.3: The top trace is the F3 kicker charging as shown on the 'cap bank' voltage signal (BXI.F3.CAP_BANK_V). This signal is found under 'Booster/Power_Supplies:Extraction'. The normalized injection current transformer is also shown. The voltage starts to rise about 15 ms after injection, and the cap banks *discharges* at extraction.

it is found in the 'Booster/Extraction/Ring' branch on spreadsheet. Since cycle 1 is normally a *dummy* cycle, the setpoint on this cycle is often set to zero for the F3 kicker.

The Extraction Bump

When the extraction bump comes on, losses can occur due to the orbit distortion that results. These losses do not occur when the beam is extracted, they occur before extraction. Figure 7.27 shows the current waveforms for the extraction bump backleg windings. Extraction normally occurs at their peak. They are on for about two milliseconds before extraction. Losses during this period are often associated with the extraction bump. These losses can be distinguished from the losses that occur at extraction by looking at their time dependence with the loss monitor analog signals. They can also be seen on the current transformer.

Within a loss monitor *time window*, the integrated loss is shown on a loss monitor analog signal. The losses normally occur in the extraction region. For example the loss monitors at F6 can be used to look at this loss. Losses at the dump (D6) might increase as well. The slope of the signal is proportional to the loss rate. Losses associated with the extraction bump coming on occur *before* the kicker fires. They can also be seen on the current transformer during this period, but the loss monitor signals also tell you where they occur.

These losses are not normally associated with a power supply problem, although they would typically increase if one of the power supplies malfunctioned. Typically, the backleg winding power supplies fail in one of two ways. Sometimes they just trip off, and so their current falls to near zero. It may not be exactly zero because of the induced voltage from the main magnet $\frac{dB}{dt}$. If it does trip off the status SLD may *not* indicate that the supply is off. However, the current readback will typically read very low, and the amplitude of the analog signal will be very small as well.

Sometimes the power supply does not trip off, but the current varies from cycle to cycle. A significant variation in the current is relatively easy to see on the analog signal for the power supply. The position in BTA will normally also vary, as well as the extraction losses, and transfer efficiency.

The F6 Septum

The F6 septum runs at thousands of amps, and a variation of 10 amps can be significant. If there is a variation in the horizontal position on BTA multiwires, but the extraction radius, ΔB^* , the AGS Rf frequency, the backleg winding currents, and the F3 kicker current are stable, it's a good bet that there is a problem with the F6 septum.

The analog signal for the F6 septum current, BXI.SPTMF6.IRDBK, has too much noise on it to measure the kind of variation that could cause unstable extraction. The current readbacks on spreadsheet are also not accurate enough. Therefore, although one may suspect that the F6 septum is responsible for unstable extraction by a process of elimination, it is difficult to prove. The digital scopes in MCR have a feature called *extended resolution* which can filter out a lot of the noise. It may be possible to correlate the amplitude of the filtered signal with the horizontal position say at MW060.

10.6.2 Losses While the Synchro Loop is Closed

As mentioned, problems with the beam in the AGS can cause beam loss in the Booster that appears to be associated with Synchronization. There are other problems that can occur while the Synchro loop is on, which may not be due to Synchro malfunctioning.

The Gap Volts at and Near Extraction

The lower the Rf voltage at extraction, the smaller the momentum spread of the extracted beam. A smaller momentum spread tends to improve transfer for a number of reasons. For example, the momentum spread contributes to the horizontal beam size. Reducing it will therefore tend to reduce losses associated with the horizontal beam size. There are other reasons why a lower momentum spread might reduce transfer losses. The point is that the gap volts are normally lowered, from their value during the rest of the cycle, to improve transfer. The gap volts are often kept just above the point where beam loss will occur. Consequently, if some parameter then varies by a small amount, beam loss may start to occur, even though it did not occur when the voltage was lowered initially. The amount of beam lost may start to vary greatly from cycle to cycle in an unpredictable and apparently chaotic way.

The extent to which the Rf voltage can be lowered near extraction has to do with the $\frac{dB}{dt}$ at and just before extraction. The $\frac{dB}{dt}$ that the magnet cycle has been configured to have tends to change considerably from year to year. The lower it is, the less Rf voltage is required for acceleration, and so the more the voltage can be lowered.

However, the Rf cavity voltage sometimes has problems regulating at low $\frac{dB}{dt}$, and lowering it may cause the beam to become unstable because of these regulation problems. The point at which the beam becomes unstable also tends to be intensity dependent. The Rf system must compensate for the voltage induced on the cavities by the beam, and this becomes more difficult as the intensity is increased. Sometimes, when the stability of the transfer is highly intensity dependent, it is because the gap volts are too low. Figure 10.4 shows this type of behavior. The beam's stability was improved in this case by raising the gap volts.²³ In the figure, the lower trace is the A6 gap volt signal. Oscillations that start to occur on it a millisecond or so before beam loss occurs (top trace). Oscillations also occur on the radial signal. The beam loss tends to occur near the peaks in these oscillations, possibly in part because the beam has been moved into an aperture at these times.

Adjusting the Radius at the Switchover

Another type of loss that occurs while synchro is on is not associated with oscillations on the gap volts. Adjusting the radius with the radial steering function just before the synchro loop comes on will often reduce this loss and make the beam more stable. In figure 10.2, one can see that the radius moves when 'control' switches to the synchro loop. The idea is to reduce, or at least change, the shift that occurs when the beam switches from the radial to the synchro loop by changing the initial radius.

The reason why the beam becomes unstable in the first place is unclear to me. It may be because $B(t)$ has changed during this part of the cycle since the frequency table was made. For example, this might happen if a tune shift was introduced around the time of the switchover. Alternately, the radial steering function may have been changed, making the radius 'less favorable' for a stable switchover. Changes to the frequency in the AGS could also effect the switchover, since they would change the radius the beam moves

²³K. Zeno, *HEP Startup Book IV*, pgs. 36-37, FY '95

6-Jun-95
1:00:35

Figure 10.4: A loss that occurred while the Synchro loop was on. The amount of beam lost on each cycle varied considerably. The stability of the beam survival and transfer were improved by raising the Rf voltage during this part of the cycle. The top trace is the Booster current transformer, the middle trace is the radial average signal, and the bottom trace is the A6 gap volt signal (BXI.RFA6.GAP_V_DC).

to once on the synchro loop. In short, the beam can become unstable at the switchover, and 'empirically' it seems that one of the relevant parameters is the radial shift that occurs there.

Extraction_Start

Moving the values of *Extraction_Start*, as well as *Synch_Start*, can also effect the transfer stability and efficiency. If just 'synchro' were considered, one might think that both these timers should be kept at the values they had when the frequency table was made. However, adjusting *Extraction_Start* will change the extraction field, and possibly the extraction radius. It changes the extraction field because it moves the *time* that extraction occurs on the B(t) function. Consequently, the extraction momentum can be adjusted by changing *Extraction_Start*. Such a change can make transfer more efficient, but it can also move the machine into a state that will make the transfer unstable. Therefore, if extraction and transfer are unstable, it might be because one or both of these SLDs have been moved to a bad place as far as the synchronization process is concerned.

The Frequency in the AGS

The Rf frequency in the AGS before *Beam_Control* and the frequency loop come on is determined by an oscillator, and this frequency does not normally vary significantly once it is set. This oscillator is called '*the crystal*'. However, once the frequency loop is turned on, normally about 5-25 ms after the first transfer, the AGS Rf frequency is determined by a voltage reference signal. This reference signal is derived from the frequency steering function found in the *AGSRFBeamControl* program. Unlike the frequency on *the crystal*, without changing the reference function this voltage, and therefore the AGS Rf frequency, has been known to change. In other words, to keep the same voltage reference for the frequency loop, so that the frequency along the AGS injection porch doesn't change, it is sometimes necessary to change the frequency steering function.

Clearly, the AGS Rf frequency before the frequency loop comes on can be different from what it is once the loop is on. Additionally, since it is then controlled by a time dependent reference,²⁴ it can be different at each of the

²⁴That is, the frequency steering function doesn't have to be flat across the injection

times at which transfers occur. This may cause the extraction frequency, radius, and momentum to vary systematically from cycle to cycle within a supercycle. If ΔB^* does not vary significantly and in a systematic way over the transfers, but the extraction radius does, it may be because of a frequency variation across the injection porch. The extraction losses on each of the four transfers will also vary systematically.

Figure 10.5 is a graphical display of the losses within one supercycle.²⁵ The window during which extraction occurs is displayed for each of the four transfers during that supercycle. The losses in the F superperiod are associated with extraction. The 'MSK' and 'GLI' loss monitors are BTA loss monitors. Losses there are associated largely with transport. The losses are highest on the last cycle (cycle 5), and are smallest on the first cycle (cycle 2). This type of variation in the loss pattern, if it is reproducible from supercycle to supercycle, may be caused by a variation in the AGS frequency across the injection porch. However, there are other mechanisms which could cause it. For example, it might be caused by a systematic variation in ΔB^* or a power supply across the cycles within the supercycle.

The frequency variation across the injection porch can be measured with the *Frequency and Time Interval Analyzer* at MCR 3. Figure 10.6 shows the frequency across the injection porch as measured with this instrument.²⁶ An AGS Rf signal is used as input. The frequency to the left of the initial oscillation is the frequency before *Beam_Control* in the AGS is turned on. In this case, that frequency is about 200 Hz less than the frequency while the frequency loop is on. Therefore, since the AGS and Booster Rf frequencies are the same, the Booster extraction frequency will be about 200 Hz lower on the first transfer than on subsequent transfers.

What effect does this have on the extraction radius? A change of 200 Hz in the AGS Rf frequency corresponds to a change of 200 Hz in the Booster extraction Rf frequency. This is a change of 100 Hz in the beam's revolution frequency in the Booster. The *differential relation* that relates a change in radius to a change in revolution frequency at fixed field is,

$$\frac{\Delta R}{R_o} = \frac{\gamma_{tr}^2 - \gamma^2}{\gamma^2} \frac{\Delta f}{f_o}$$

porch

²⁵ *HEP Setup Book II*, pg 22, FY 1996

²⁶ L. Ahrens, *HEP Setup, Booster/AGS*, pg. 104, FY 1997

HP 5371A Frequency and Time Interval Analyzer

TVar: Frequency A 16 Jun 1997 13:52:03
 ● Mkr y: 2.8090667 MHz

Figure 10.6: The AGS Rf frequency across the injection porch as measured with the *Frequency and Time Interval Analyzer* at MCR 3 using an AGS Rf signal. The vertical axis is frequency, 200 Hz/div. The horizontal axis is time, 50 ms/div.

For $h=2$ in the Booster and $h=8$ in the AGS, the Booster's revolution frequency at extraction is the AGS Rf frequency divided by two, $f_o = \frac{2.809MHz}{2} = 1.4045MHz$. So, if $\Delta f = 100Hz$, the radius will shift by about 3 mm.

The frequency shift in the AGS can be corrected by adjusting the frequency steering function to match the frequency before *beam_control* comes on. However, this 'correction' will change the AGS setup as well.

10.6.3 Coherent Instabilities

As illustrated in section 9.2.1, the transfer efficiency and beam size can be affected by changing the tunes during acceleration. One of the reasons for this dependence may be the effect of stopbands. Any mechanism that increases the emittance will tend to have an effect on extraction losses. One possible mechanism for emittance 'blow-up' that has not been considered until now is called a *Coherent Instability*. At high intensity, the beam can interact with the fields outside it to produce coherent motion. This motion can be either transverse or longitudinal. In the transverse plane, it can resemble the coherent betatron oscillations described in section 8.2.7, or be somewhat more complicated. Unlike the coherent oscillations that were described in that section, these arise during the acceleration cycle, and are not directly associated with the injection mismatch. Like stopbands it is a kind of resonance phenomenon that may lead to beam loss and emittance 'blow-up'.

Particularly in the vertical plane, there is evidence for *coherent instabilities*. Figure 10.7 shows beam loss that appears to be associated with a vertical coherent instability.²⁷ The amplitude of the oscillations, as shown in the vertical PUE difference signal, increases reaching a maximum value around when beam loss occurs on the current transformer. Presumably, these oscillations about the E.O. cause the beam to scrape on vertical apertures and result in beam loss. In this case, the loss was removed by adjusting the vertical chromaticity so that it was further away from zero during this part of the cycle. This also reduced the amplitude of the oscillations visible on the PUE signal.

Instabilities like this one may occur without causing beam loss, because they may not become strong enough to cause the beam to scrape. However,

²⁷T. Roser, *Booster VII*, pg. 118-199, FY 1993

20-Apr-93
19:27:01

VERTICAL INSTABILITY

Figure 10.7: Evidence for a vertical coherent instability. The top trace is the E7 vertical PUE difference signal. The bottom trace is the current transformer.

the emittance may still increase because of them. Their strength tends to vary unpredictably from cycle to cycle. Hence, if the transfer efficiency varies a lot from supercycle to supercycle, it may be due to such an instability. Since they can cause beam loss, they may also cause *Booster Late* to vary considerably from supercycle to supercycle.

For the most part, these instabilities seem to occur in the vertical plane. The beam will in general become more susceptible to them as the chromaticity is moved towards zero and the tune is moved towards the integer. Hence, coherent instabilities are a reason to *lower* the vertical tune during acceleration. Coupling between the motion in the two transverse planes can also make the beam more stable. Moving the tunes closer to each other increases the amount of coupling. The optimized requested tunes during high intensity are shown in figure 5.4. Their values in both planes during the middle and late acceleration cycle, when the space charge tune shifts are not as large as they are earlier in the cycle, are close to each other and relatively far from the integer.

Appendix A

A Transfer Matrix Description of the Extraction Bump

At the end of section 5.3 the idea of a *transfer matrix*, which relates (x, x') at s_1 to (x, x') at s_2 was mentioned (s_1 and s_2 are longitudinal positions). This *transfer matrix* is 2×2 , and is denoted by $M_{s_1 \rightarrow s_2}$. A *4-bump* can be described using this idea. Consider a particle upstream of the first kick and on the undistorted E.O. Its trajectory relative to the undistorted E.O. at the second kick (F4) is related to its trajectory as it exits the first kick (at F2) by, $(x, x')_{F4} = M_{F2 \rightarrow F4}(x, x')_{F2}^*$. '*' denotes $(x, x')_{F2}$ after the kick. But, $x_{F2} = 0$ since the kick has just occurred, and $x'_{F2} = \delta_{F2}$ (where δ_{F2} is the angular deflection from the kick at F2). Similarly,

$$\begin{pmatrix} x \\ x' \end{pmatrix}_{F6} = M_{F4 \rightarrow F6} \begin{pmatrix} x \\ x' \end{pmatrix}_{F4}^*$$

the trajectory just *before* the kick at F4 is,

$$\begin{pmatrix} x \\ x' \end{pmatrix}_{F4} = M_{F2 \rightarrow F4} \begin{pmatrix} 0 \\ \delta_{F2} \end{pmatrix}$$

so,

$$\begin{aligned} \begin{pmatrix} x \\ x' \end{pmatrix}_{F6} &= M_{F4 \rightarrow F6} \left[M_{F2 \rightarrow F4} \begin{pmatrix} 0 \\ \delta_{F2} \end{pmatrix} + \begin{pmatrix} 0 \\ \delta_{F4} \end{pmatrix} \right] \\ \begin{pmatrix} x \\ x' \end{pmatrix}_{F6} &= M_{F2 \rightarrow F6} \begin{pmatrix} 0 \\ \delta_{F2} \end{pmatrix} + M_{F4 \rightarrow F6} \begin{pmatrix} 0 \\ \delta_{F4} \end{pmatrix} \end{aligned}$$

This last matrix equation can be expressed as two non-matrix equations,

$$x_{F6} = (m_{12})_{F2 \rightarrow F6} \delta_{F2} + (m_{12})_{F4 \rightarrow F6} \delta_{F4}$$

$$x'_{F6} = (m_{22})_{F2 \rightarrow F6} \delta_{F2} + (m_{22})_{F4 \rightarrow F6} \delta_{F4}$$

The subscripts 12 and 22 refer to the matrix elements. These are two equations, with two unknowns (δ_{F2} and δ_{F4}). For a given (x, x') at F6, the required kicks at F2 and F4 can be obtained from these two equations.

Similarly, starting at A1 and working backwards towards the desired (x, x') at F6 will give the kicks at F7 and A1 needed to put the E.O. back to its undistorted state downstream of A1. The transfer matrices used are from A1 to F7, and from F7 to F6. They are the inverse of the matrices from F7 to A1, and F6 to F7, respectively.

Appendix B

Typical Operating Values as of the end of 1997

<i>Parameter</i>	Protons	<i>Au⁺³²</i>	<i>Fe⁺¹⁰</i>
Injection Tunes			
<i>Horizontal</i>			
OpticsControl	4.83	4.73	4.70
Measured coherent	4.82	4.73	4.68
<i>Vertical</i>			
OpticsControl	5.00	4.82	4.84
Measured Coherent	4.90	4.82	4.85
Injection Chromaticity <i>(OpticsControl)</i>			
Horizontal	-0.4	-1.5	0
Vertical	0	-0.4	0
Skew Quads @ injection			
QS1	N/A	12.7 A	6.5
QS2	N/A	5.0 A	6.5
QS3	N/A	-5.0 A	6.5
QS4	N/A	2.5	6.5
Injection Time (from BT0)	17 ms		15.1 ms
<i>f_{rev}</i> at Injection	842.2 KHz	66.3559 KHz	103.88 KHz
Rf Harmonic	2	8	8

<i>Parameter</i>	Protons	Au^{+32}	Fe^{+10}
Ext. Time (from BT0)	85 ms	?	?
Cycle Length	150 ms	200 ms	200 ms
f_{rev} at Extraction	1.405 MHz	625 KHz	625 KHz
Extraction Energy	1.96 GeV	95 MeV/n	95 MeV/n
Extraction Momentum	2.7 GeV/c	430 MeV/c/n	428.8 MeV/c/n
Injection Energy	200 MeV	0.92 MeV/n	2.27 MeV/n
Injection Momentum	640 MeV/c	41.5 MeV/c/n	64.9 MeV/c/n
Injection Field	1550 g	615 g	868 g
Injection $\frac{dB}{dt}$	33 g/ms	1-2 g/ms	1-2 g/ms
$\langle \frac{\Delta p}{p} \rangle_{fullwidth}$ at injection	7×10^{-3}		7×10^{-4}
Injection foil position	1" from center (3 or 4)	out (1)	out (1)
Inflector Readback	0	48.0 kV	21.8 kV
Typical Late Intensity	$2 \cdot 10^{13}$	$1 \cdot 10^9$ ions	$3 \cdot 10^9$ ions

Parameter	Value
Radius (R_o)	32.114 m
Bending radius (ρ)	13.87 m
γ_{tr}	4.9
Horizontal β_{max}	14.0 m
Vertical β_{max}	13.2 m
Horizontal β_{min}	3.8 m
Vertical β_{min}	3.9 m
Average Dispersion	1.6 m
Maximum Dispersion	3.0 m
Minimum Dispersion	0.5 m
Vertical Acceptance	90π mm mrad

Appendix C

Selected Bibliography

Concepts

1. D.A. Edwards and M.J. Syphers, *An Introduction to the Physics of High Energy Accelerators*, John Wiley and Sons, 1993.
2. Stephen D. Holmes, “*Low Energy Aspects of Circular Accelerators*”, Fermilab-Conf-90/275, December 1990.
3. W.T. Weng and S.R. Mane, *Fundamental of Particle Beam Dynamics and Phase Space*, AGS/AD/91-2.

Booster Design

1. *The Booster Design Manual*, December 5, 1988.

Operations and Other Related Topics

1. L. Ahrens, C. Gardner, AGS Studies Report No. 365, *A Study of the Effect of Linear Coupling on the Injection of Iron Ions into the Booster*, November 19, 1997.
2. L. Ahrens, AGS Studies Report No. 353, “*Gold Injection into the Booster: Beam Survival as the Length of the Tandem Current Pulse is Varied*”, Feb 24, 1997.
3. Willem van Asselt, AGS/AD/Op. Note No. 40, “*Booster Tune Control*”, February 4, 1993.
4. J. Michael Brennan, *Rf Beam Control for the AGS Booster*, BNL-52438,

September 26, 1994.

5. H.M. Calvani and L.A. Ahrens, *Au³²⁺ Beam Intensity Losses in the AGS Booster Due to Charge Exchange Processes*, Booster Technical Note No. 228, October 1, 1996.
6. Pepin Carolan, AGS/AD/Op. Note No. 33, *“Booster Chromaticity Measurement”*, September 1, 1992.
7. Pepin Carolan, AGS/AD/Op. Note No. 37, *Topics in Booster Main Magnet Operation*, November 30, 1992.
8. Dave Gassner, AGS/AD/Op. Note No. 42, *Booster Extraction*, January 25, 1994.
9. Paul Sampson, AGS/AD/Op. Note No. 43, *“Transport from the Booster to the AGS”*, September 6, 1994.
10. Jonathan Reich, AGS/AD/Op. Note No. 35, *Booster Current Transformers and Booster Wall Monitors*, November 13, 1992.
11. B. Tamminga and W. Van Asselt, AGS/AD/Op. Note No. 44, *Mountain Range Display*, December 12, 1995.
12. Charles B. Whalen, Operations Note #36, *Stopband Corrections*, Nov. 30, 1992.
13. K. Zeno, et al., AGS/AD/Op. Note 38, *Overview of Booster Timing*, February 17, 1993.
14. K. Zeno, AGS/AD/Op. Note No. 39, *“Strategies for Measuring and Establishing Values for Proton Injection Parameters in the AGS Booster”*, December 22, 1992.
15. K. Zeno, Booster Tech. Note No. 227, *“Coherent Tune Shifts in the Booster”*, September 6, 1995.

Index

- ΔB^* , 270
- $\Delta f(t)$, see 'frequency table'
- ρ , see 'Bending radius'
- Acceptance
 - Defined 45
 - Time dependent 100
- Acceleration
 - Heavy Ions 127-136
 - Protons (losses during) 211-255
- Adiabatic capture, see 'capture'
- Adiabatic damping 45
- AFR (AGS field reset) 142
- AGE (AGS Group End) 145
- AGS CBM 223, 242, 267
- Alternating Gradient 48
- Analog Signals
 - Beam Phase Error
 - BXI.BEAM_PHASE_ERR 86, 207
 - Booster Current transformers
 - Injection transformer 106, 112, 115, 117, 118, 123, 128, 134, 137, 210, 222, 238, 244, 258, 279
 - BXI.CIRC_XFMR_NORM 132, 135, 232, 268
 - AXI.F15_IXFMR_NORM 258, 268
 - Dump Bump
 - BXI.DUMP_BUMP_I 246
 - Extraction bump backlegs 158
 - Heavy Ions 106, 115, 118, 123, 128, 134, 137
 - Protons 184
 - F3 kicker
 - Current 146
 - BXI.F3_CAP_BANK_V 279
 - F6 Septum Current
 - BXI.SPTMF6_L_RDBK 281
 - Inflector Current 111
 - Loss Monitors 163, 242, 280
 - LTB BPMs
 - BXI.LTB.BPM019_4SUM 184
 - LTB current transformers 178
 - Main magnet current
 - BXI.MAIN_MAG_CUR 130, 135, 238
 - Main magnet voltage
 - BXI.STA1_PS_V 130, 238
 - Orbit Correctors 78
 - PUEs
 - C4 146
 - C6 128, 191
 - C7 198
 - E7 289
 - Radial signals
 - BXI.RFA6_RADA2 90
 - BXI.RFE6_RADA8 90
 - BXI.RFE6_RADIAL_AVG 207, 268, 283
 - Rf voltage
 - AGS Rf 287
 - BXI.RFA3.V_GAP_DC 132
 - BXI.RFB3.V_GAP_DC 135
 - BXI.RFA6.GAP_V_DC 283
 - BXI.BRF_VECTOR_SUM 128
 - Synchro Phase Error
 - BXI.SYNCHRO_ERROR 268
 - TTB current transformer 106, 112, 115, 118, 123, 134, 137
 - Tune quad current 137, 215
 - Tune kicker current 217

- Wall Current Monitor 86, 205, 210
- Aperture, see 'limiting apertures'
- AT0 (AGS T-zero) 141-144, 258
- AU1 (AGS User 1) 142
- Au^{32+} , see 'Gold'
- Backleg windings 150
- Bare tunes 56-57
- Batch 256-257
- BeamLineEmit 174-175
- BeamLineInstrument Application 20-21
- Beam_Control* 87
- Beam dump 241-248
- Beam Position Monitors (bpms)
 - Booster 22-24
 - LTB 25-27
- Bending Radius, see 'radius of curvature'
- Beta function 52
- Betatron Tunes, see 'Tunes'
- Betatron oscillations 42
- BFR (Booster Field Reset) 142
- BLW, see 'Backleg Windings'
- BMMPs, see 'Main Magnet power supply'
- Booster Early* 223
- Booster Input* 223
- Booster Late* 223-228, 239, 242, 267, 290
- Booster Lattice 46-50
- BoosterLossMonitor 27-31, 177, 286
- BoosterMainMagnet 131
- BoosterOrbitControl
 - Dipole Correctors
 - Harmonic Corrections 75-79
 - 3- bumps 252-255
 - Dump Bump 246
 - Extraction Bump 151
 - Correction 4-bumps 254-255
 - Fast Injection Bump 102-105
 - Slow Injection Bump 251
- BoosterOrbitDisplay 23-24, 69-70, 78-79, 153, 156, 245-247, 251, 267
- BPMs
 - defined 22-23
 - LTB 25-27
 - Booster 23-24
- Booster Rad Inhibit button 30
- BoosterRfBeamControl
 - Counterphasing function 91-95
 - Frequency Program 81-83
 - ScalerVoltsperTurn function 92-93
 - Radial Steering function 87-89
- BoosterTuneMeter 217-218
- BGE (Booster Group End) 145
- BPP (Booster Prepulse) 92, 94, 142
- Bstart* 82, 91
- BTA (Booster to AGS) line 33, 141, 158-165, 221-226, 261
- BT0 (Booster T-zero) 16
- BT1 145, 161, 266
- BT4 145, 161, 266
- BU2 (Booster User 2) 144
- Bucket
 - and phase loop 85
 - and Heavy Ions 126, 127-129
 - and Chopper 186-190
 - and matching 207-208
 - and second harmonic cavities 235-237
- Bucket size 88, 129-130, 133, 135, 139
- Bump residuals
 - defined 153-155
- Bunch
 - and current transformers 20
 - and BPMs 233
 - and Wall monitor 31-32
 - and Chopper 185-190
- Bunch shape oscillations 128, 136
- Bunch-turn matrix 188-189
- Bunch width 188
- C3 inflector 98-106, 111-113, 124, 181
- Calibrate* 17
- Calibrate pulse 19-20
- Capture
 - Adiabatic 126
 - Heavy Ions 125-129, 133-136
 - Protons 203-210
- Capture loss 133-136, 208-210

- Chipmunk 277
- Chopper, see 'Fast' or 'Slow'
- Chopper application 187-190
- Chopped beam 183
 - and LTB BPMs 27
- Chopper_Phase* 186, 206-208
- Chopper_Width* 182, 187, 226
- Chromaticity
 - defined 66-68
 - and betatron oscillations 197
 - and tune spread 212, 213
 - and main magnet 214
 - and tune measurement 217
 - and coherent instabilities 288
- Chromaticity Control 71-72
- CLD (Complex Logical Device)
 - and BPMs 23
 - and extraction equipment 278-280
- CLD pop-up, see 'Spreadsheet'
- Closed Orbit, see Equilibrium Orbit
- Coherent Betatron Oscillations 196-201, 288
- Coherent Synchrotron Oscillations 85
- Coherent Instability 273, 288-290
- Coherent tune 217-219
- Cogging 259-261, 265-266, 278
- Cogging_Arm* 147, 161, 270
- Cogging Box 147, 157, 159
- Communications path 38-40
- Configure application 38
- Controller, see 'Device Controller'
- Control Point 76-77
- Controls Station 36
- Counterphasing
 - defined 92
 - and Heavy Ion injection 129-134
 - and second harmonic 237-239
- Coupled Multi-turn Injection 107-109
- Coupling and beam size 249
- Courant-Snyder Invariant 201-203
- Crystal, the 284
- CSI, see 'Courant-Snyder Invariant'
- Current Transformers,
 - Calibrating 19
- Circulating 104, 132-133, 135, 268
 - defined 18-20
 - gain 19, 110
 - Injection 18-19, 106, 110-117, 122-124, 209-210, 237-238, 258
 - LTB 18, 176-178
 - Measuring intensity with, 19
 - Normalized 20
 - 'TTB' 104-106
- Cyclotron motion 5, 51, 81
- $\frac{dB}{dt}$,
 - at bucket size minimum 129-135
 - and eddy currents 214
 - at extraction 281-282
 - at injection 125, 185, 195-196
 - and injection bump ramp 201-203
- Device Controllers 37-40
- Device,
 - Controls System 34-36
 - Physical 34-37
- Device details 38
- DH115 192, 201, 230
- Difference orbit 153, 246, 252, 254
- Differential relations 270-271, 274
- Dipoles
 - Backleg windings 150
 - Correctors 75-79
 - Main 7, 13-15, 49-50
- Dipole Errors,
 - described 72-75
 - harmonic correction of 75-79
- Dispersion
 - effect on beam size 136, 139, 248
 - and BPMs 267
 - effect on orbit 69-71, 76
 - in LTB 172, 182
 - and the radius 254
- Dummy cycle 145, 270, 280
- Dump Bump 241-248, 252-254
- DV095 213, 226-228
- Dwell field 144
- Early Turn off 241
- Eddy currents 214

- Emittance
 - and BTA foil 141
 - and high intensity 193, 212-213,
 - and injection foil 180-181, 248-251
 - and Losses 239-255
 - and transfer 221-226, 259, 271. 288-290
 - Longitudinal 83, 85, 125-135, 208,
 - of Tandem beam 119-121
 - Transverse 45, 52, 139, 150, 157, 159, 169-170, 182, 183, 201, 216
 - Normalized 45-46, 52
- Equilibrium Orbit
 - and bump residuals 150-155
 - defined 42
 - and dipole errors 72-75
 - and dispersion 69-70
 - and H. I. injection 98-102
 - and proton injection 180-181, 195-203
- Extraction
 - Radius 158, 162-163
 - and the supercycle 141-145
 - Synchronized 256-290
 - and Transfer 164-165, 275
 - Types 140
 - Unsynchronized 145-165
- Extraction Bump,
 - with Heavy Ions 150-158
 - with Synchro 266
 - effect of $\frac{dB}{dt}$ 151, 273
 - and residuals 153-157
- Extraction field
 - and momentum 162, 273-275
 - and radius 270-271
 - and ΔB^* 270
 - and Tune functions 272-273
- Extraction Losses 158-161, 275-290
- Extraction_Start*
 - defined 157
 - and Extraction field 162
 - and transfer stability 284
 - and tune functions 272-273
- F3 Extraction Kicker
 - and Fast extraction 140, 145-150
 - and BPP 142
 - and cogging 260, 265-266
 - problems 159-160, 276-280
- F6 Extraction Septum
 - described 149
 - and position at MW006 161, 163-164
 - power supply triggers 157
 - problems 281
- Fast Chopper
 - and Chopper program 187-190
 - defined 183
 - and emittance 183, 184
 - and LTB BPMs 27
 - phase 185
 - and radiation 186
- Fast Injection Bump
 - and BoosterOrbitControl 102-105, 202
 - magnets 102, 181
 - and *Fast.tm* 102, 203
 - Heavy Ions 102-106, 122-123
 - Protons 181, 201-203
- Fast.Tm* 17, 102, 203, 230
- Fe^{10+} , see 'Iron'
- FBI, see 'loss monitors, FBI'
- FFT 76, 217, 246
- Field errors 46, 56, 63, 73-79, 108, 212, 214
- Field reset 142
- Flavored T-zeros 145, 161, 266
- Fdot 235-237
- Foils
 - BTA 141, 261
 - Proton injection 180-181, 182, 190-192, 193, 195, 197, 201, 227, 241, 248-251
 - Tandem 117-121, 121-123
- Foil change 119
- FODO 48-50
- Fourier Transform, See 'FFT'
- Frequency
 - Revolution 6, 19, 81-82, 91, 126-127, 186, 187, 190, 195, 203-208, 259,

- 261, 285-288
- Rf 19, 82-90, 126-128, 145, 183, 186-188, 203-208, 233-234, 257, 261-264, 266, 276, 281, 284-288
 - Moving reference frame 263, 267, 277
 - Target 263, 264, 266
- Frequency Loop (in AGS) 276-277, 284-285
- Frequency Program 82, 91
- Frequency Table 81, 82,
 - and Synchro 264, 266, 282, 284
- Frequency Steering (in AGS) 277, 284-288
- Functions (reference) defined 12-15, 17
 - Chrom Control 72
 - Counterphasing 92, 94-95, 129, 237-240
 - Dipole Correctors 76-79
 - Dump bump 246
 - Fast Injection Bump
 - Heavy Ions 102-105
 - Protons 201-203
 - Frequency Steering (in AGS) 284
 - Main Magnet
 - B(t) 12-13, 130-131, 143-144
 - DB/dt 130-131
 - Radial Steering 87-91
 - ScalerVoltsperTurn 92-93, 129
 - Stopband Correction 61-62, 214-215
 - skew quads 116
 - Tune Control 12-15, 54-55
- Gauss Clock
 - Defined 17
 - and Gauss Time line 15-17
 - and Rf 81-82, 85
 - Scalers 269-270
 - and Supercycle 142
- Glitch bit 92-95, 129, 132, 239
- Gold 8, 10-11, 96-97, 104, 109, 117-121, 136-139
 - see also 'Heavy Ions'
- GPM (General Purpose Monitor) 28-30
- Group end 145
- Hall probe 17
- Harmonic number '*h*' 19, 87, 186, 261
- Harmonic corrections, see 'Dipole Errors'
- Harps (see multiwires)
- Heavy Ions
 - Operation 96-158
 - and vacuum 11, 136-139
 - and loss monitors 30
- HEBT 166, 174-179
- Incoherent tune 219
- Induced voltage,
 - on extraction backlegs 151, 280
 - from tune quads 272-273
- Inflector, see 'C3 inflector'
- Inflector current 111
- Injection Bump, see 'Fast injection Bump'
- Injection kickers, see also 'Fast Injection Bump'
 - in Booster 16-17, 103-106, 123, 128, 134, 137, 195
 - in AGS (A5) 30, 259-260
- Injection porch (AGS) 221-224, 257-258, 263, 284-288
- Instrumentation 18-32
- Intensity 4
 - and beam loss
 - Heavy Ions 138-139
 - Protons 211-213
- Intensity Scalers, see 'Scalers'
- Iron 109, 112, 122
- Jam_Value* 17, 142
- Jiffy 142-144
- Kinetic Energy 8, 271
- LEBT 182, 183, 186, 196
- Limiting aperture 76, 116, 163, 226, 239-255
- Linac 4-5, 15, 27, 30, 84, 166, 172, 175, 182-190
- Longitudinal, defined 31
- Longitudinal charge density 208, 212, 233-239

- Longitudinal coordinate
 - defined 42-43
 - and LTB 25
- Longitudinal matching 125-132, 206-209
- Longitudinal Emittance, see 'Emittance, Longitudinal'
- Longitudinal charge density 208, 212, 233-239
- LTB (Linac to Booster) line 166-180
- Loss Monitors
 - Application, see 'BoosterLossMonitor'
 - Analog signals 28, 163, 177
 - Calibration 30
 - Latching 30, 285-286
 - FBI (Fast Beam Inhibit) 30
 - and Proton extraction 280
 - and Proton injection 177
- Low Level Rf, see 'Rf Beam Control'
- Macroparticle 83
- Macropulse 183
- Main Magnet Power Supply (BMMPS) 12-13
- Main Magnet field 12-17, 42, 49-50, 129-131
- Main Magnet dB/dt 83, 85, 125-136, 185, 195-196, 201, 212
- Matching 96-97, 128, 201-203, 206-209, 228, 230
- Micropulse 186-187
- Mismatch 128, 167-172, 201-203, 208, 228, 230
- Momentum 5-7
 - and loss monitors 30
 - and space charge 185
 - and vacuum effects 136-139
- Momentum Dispersion, see 'Dispersion'
- Momentum Spread
 - and chromaticity 66-71
 - and Heavy Ions 136, 141
 - and protons 172, 180, 182, 197, 203-205, 248, 281
- Morning Numbers 120, 191, 205, 264, 270
- Mountain Range 203-206
- Moving reference frame 263, 267, 277
- Multi-turn Injection 98-109
- Multiple Transfers 256-261
- Multiwires 20-21
 - in BTA 160-165, 224-226
 - in LTB 172-174
 - in TTB 20-21, 97
- Nodes 36
- OpticsControl
 - ChromControl 71-72
 - TuneControl 53-57
- Orbit distortions 72-75
- Orbit harmonics 72-79
- Orbit Correction 75-79
- Orbit correctors, See 'dipoles'
 - (*Peaker*) 12-17, 82-84, 115-116, 195-196, 243, 272-273
- Phase advance 51, 174-175, 291-292
- Phase Loop 85-87
- Phase space
 - Transverse 44-46
 - Longitudinal 188-190
- Phase space painting 231
- PIP 199-201
- PLC (Programmable Logical Controller) 37
- PPM 16
- PUE (Pick-up electrode) 22-24
 - see also 'analog signals, PUEs'
- Quadratic fit 267
- Quadrupole 46-50
- Quadrupole, skew 63, 66
 - see also 'Coupled Multi-turn Injection'
- Radius of Curvature (ρ), 7
- Radial Loop 87-91
- Radial Steering 87-91
- Radiation
 - Chipmunk 277

- and the Fast Chopper 186
 - and the Beam Dump 241
 - see also, 'Loss Monitors'
- Reference magnet 17, 85
- Reference orbit 72
- Reference phase 85
- Reports 38-40
- Residuals,
 - see 'extraction bump, residuals'
 - and injection bump 230, 252-254
- Residual gas 136-139
- Resonance 46
- Rf 80-95, 233-239
- Rf Beam Control 83-87
- Rf Cavities 91-95, 233-234
- Rf Vector Voltage 91-95, 132
- Rf_Track* 82-84
- Rigidity 51
- Scalers 20, 223, 269
- Scraping 116
- Second Harmonic Cavities 233-239
- Sieve 188
- Signals, see Analog Signals
- Sextupoles
 - Chromaticity Correction 71-72
 - Stopband Correction 63, 66, 219-222
- Skew Sextupoles 219-222
- SLD (Simple Logical Device)
 - Defined 36
 - In Spreadsheet 33-36, 39
- SLDs
 - ABI.DH1.SPRB 34
 - ABI.DH1.STAT 34
 - AGN.INJ.CBM.SP 223
 - AGN.CBM.SP 223
 - ABI.KDHA5.IO 147
 - ARF.BEAM.CONT.ST
 - Beam_Control* 276, 277, 284, 285, 288
 - ARF.FREQ.LOOP.ST 276
 - BIJ.FAST.TM
 - Fast_Tm* 17, 102, 203, 230, 231
 - BIJ.PEAKER.GT see *Peaker*
 - BIJ.RF_TRAK.ST.GT
 - Rf_Track* 82, 84, 91, 126, 208
 - BIJ.SLO.TM.OFF 251
 - BIJ.SLO.TM.ON 252
 - BIJ.SPTMC3.SPRB 98
 - BGN.DMP.BMP.RDTIME 246
 - BGN.DMP.BMP.STTIME 246
 - BGN.DUMP.BMP.STAT 246
 - BGN.EARLY.CBM.SP 223
 - BGN.INTEGRATE.P.SP 223
 - BGN.LATE.CBM.SP 223
 - BGT.CALIBRATE
 - Calibrate* 17
 - BGT.JAM.VALUE
 - Jam_Value* 17, 142
 - BGT.SFIELD.BUFFER 17
 - BL1.DH2-5.SPRB 167, 174
 - BL1.PROTON.TYPE 110
 - BL1.XF_GN 110
 - BMC.CO_COR_SMPL.RT 78
 - BMC.CO_COR.ST.RT 78
 - BMD.INJXF.CAL 19
 - BMM.CHROM.H.PS 72
 - BMM.CHROM.V.PS 72
 - BMM.TDHC4.STAT 250
 - BRF.AGS_RF_READY 266
 - BRF.BEAM.CONT.ST 86, 87, 91, 204, 206, 208
 - BRF.BSTR.AGS.PHASE 260
 - BRF.CNTR.PHS.ST 94
 - BRF.COGGING_ARM
 - Cogging_Arm* 147, 264, 266, 270, 276
 - BRF.COGGING.RESET 266
 - BRF.EARLY_OFF1.RT 241
 - BRF.LL.B.STRT
 - Bstart* 82, 84
 - BRF.LL.VOL.CMD.ST 92
 - BRF.PKR_TRG_C.PHS 94, 95, 239
 - BRF.RADIAL_LOOP.ST 91
 - BRF.XTRCN.STRT.GT
 - Extraction_Start* 157, 161, 162, 262, 266, 268, 270-273, 284
 - BTG.BCBM.CAL.RT 19
 - BXO.CARBON_FOIL 190

- BXT.B.BLW_DISCHARGE 157
- BXT.F3&A5.ST_CHARGE 278
- BXT.F6_DISCHARGE 157
- BXT.F6&BLW_STCHARGE 157
- BXT.KDHF3.IO 147
- BXT.KDHF3.SP_CLD 278
- LIN.CHOPPER_ON_OFF 185
- LPI.CHOPPER_PHASE
 - Chopper_Phase* 186, 187, 196
- LPI.CHOP_ON_DLY.RT 182
- LPI.CHOP_WIDTH
 - Chopper_Width* 182, 187, 226
- TTM.TBEAM_ON
 - Tbeam_on* 110
- TTM.SOURCE_WIDTH 136
- Slow Chopper 182
- Slow Injection Bump 248-251
- Space charge 10, 219-220, 228-229
- Spiraling beam 116-117, 243-244
- Spreadsheet 33-40
- Stacking,
 - during Heavy Ion injection 106
 - with multiple transfer in AGS 256-258
- Station, see 'Controls Station'
- Steering magnets 167
- Stopbands 56-66
- Stopband Correction application 59-66
- Stopband Corrections 56-66
 - Half Integer 61-65
 - Linear coupling 63, 66
 - Third integer 63
- Store 109, 158
- Straight section 49-50
- Stripping foils, see 'foils'
- Strong Focussing 7, 41-50
- Supercycle events 15-17, 141-145
- Superman 141-145
- Switchover 87
- Synchronization (*Synchro*) 256-258, 262-264
- Synchronous phase 85, 127-128, 186, 187, 203, 206-208, 233-239
- Synchro Loop 262-264
- Synchrotron 6, 80
- Synchrotron oscillations 80, 83-87, 128, 204-209
- Tandem current 96, 122-123, 136
- Tandem foil, see 'Foils, Tandem'
- Tandem foil change, see 'foil change'
- Tandem pulse width 98, 136-139
- Tandem Van De Graaf
- Tbeam_on 110
- Timelines
 - Gauss 16-17
 - Real 16
 - Supercycle 15-16
- Timeline Decoder 12
- Time dependent acceptance 100-102
- Trajectory 44, 47, 52, 99, 106, 149, 167-172
- Transfer efficiency 221-226
- Transverse 41, 41-79
- Triggered Vector Advance, see 'glitch bit'
- Trim windings 54
- TTB (Tandem to Booster) line 5, 96-98
- TT0 (Tandem T-zero) 144
- Tunes 46, 51-53, 53-57
- Tune Kicker 216-219
- Tune Shift,
 - near Heavy Ion injection 137, 139
 - before proton extraction 272-273
- Tune space diagram 59-60, 227-229
- Tuning 3-4, 9, 122-124, 227
- Turns, number of
 - Heavy Ions 98
 - Protons 186-187
- User reset 142
- Vacuum 5-6
- Vacuum Chamber 6
- Vacuum Related Loss 11, 136-139
- Wall Current Monitor,
 - and Capture 203-206
 - Intensity Measurement 31-32

Revolution frequency measurement
203-206

Working line 55

Xbar 33

Xbar signals, see 'Analog Signals'

XF011, see 'current transformers, LTB'

XF100, see 'current transformers, LTB'