Recycler Rookie Book

Table of Contents

Acknowledgements: Part Karns compiled and wrote much of this rookie book. Other contributors included Cons Gattuso, Martin Hue, Dan Broemmelsiek, Maggie Stauffer, Chip Edstrom, and Steve Baginski.

Introduction

Geography

The Recycler ring is located in the southwest corner of the Fermilab site, adjacent to the Tevatron on one side and the Prairie Path bicycle trail on the opposite side. The Recycler Ring and The Main Injector share the same tunnel, the Recycler is located about 47 inches above the Main Injector, center line to center line.

Because the Recycler is located in the Main Injector tunnel, it follows the MI's naming convention. Because of this, Recycler numbers increase in the direction of proton travel rather than in the direction of antiprotons travel, even though antiprotons are its primary particle. There are six major "mileposts" at more or less equal intervals around the ring: MI-10, MI-20, MI-30, MI-40, MI-50, and MI-60. Protons are injected from Booster into the Main Injector at MI-10 and travel counterclockwise around the ring. At each of these mileposts there is a service building which houses equipment related to the accelerator components downstairs. Service buildings also provide access to the tunnel.

In between the major mileposts, intermediate numbers give locations; for example, between the MI-10 and the MI-20 service buildings the half-cell locations (quad location) start off with 100 and proceed up to 131 just below the MI-20 service building. At that location the numbering system switches over to a 200 series. This pattern repeats all the way around the ring.

Role of the Recycler

The purpose of the Recycler is to store antiprotons. The number of antiprotons available has always been an important limiting factor in producing the high luminosities desired for stores in the Tevatron. They are difficult, or at least time-consuming, to produce—and if enough are created they are difficult to store because large numbers of antiprotons in the Accumulator tend to develop instabilities and can be suddenly lost.

The Recycler ring is located directly above the Main Injector ring. When antiprotons are extracted from the Accumulator to the Recycler, they arrive in the same way they would during an HEP shot—through the AP-3, AP-1, P2 and P150 lines into Main Injector at MI-52, where they circulate for a few seconds (but are not accelerated) before being transferred into the Recycler at MI-22.

The Recycler then stores this beam for many hours. During this storage, the antiproton beam is "cooled." This cooling is done to reduce the longitudinal and transverse spread of the beam. More antiprotons can be injected and stored after they have been cooled, and the Recycler typically takes several injections from the Antiproton Source between Tevatron shots. In this way, the Recycler is able to provide higher intensity, lower emittance antiprotons bunches for the Collider Physics program.

Notes:

Magnets

The Recycler is much like Main Injector and Tevatron. They all use a repeating FODO lattice. They all use a 95 LCW system for cooling what needs to be cooled. The biggest difference in the Recycler is that its magnets are comprised mainly of permanent magnets. These magnets are made from magnetized strontium ferrite bricks. Interspersed between the bricks are nickel-steel “compensator” strips, included to balance out the temperature dependence of the strontium ferrite, and to balance the longitudinal magnet strength. The pole tips are made from low carbon steel, used for its lesser hysteresis effect, which causes magnetic memory effects in the pole tips. A small permanent change has been introduced into the magnetic field of the pole tips by accident before by running large currents down the beam pipe, but most all hysteresis effects seem to have been washed out during the assembly procedure as measured field values agree to within 0.01% of predicted values.

We have several types of permanent magnets in the Recycler, performing several different functions. There are gradient magnets that work as dipoles and quadrupoles. Specially made “Mirror Magnets” are used in the cramped space of the transfer lines in lieu of gradient magnets. There are quadrupole permanent magnets, and even the Lambertson magnets are permanent magnets. We will begin with the gradient magnets.

Permanent Magnets

Gradient Magnets (RGF, RGD)

In a FODO lattice, focusing and defocusing quadrupoles alternate, separated by dipoles. The Recycler uses combined function gradient magnets that work as combined dipoles and quadrupoles. There are four kinds of permanent gradient magnets used. They include a focusing and defocusing magnet for both regular arc cells and dispersion suppression cells.

Arc Gradient Magnets

The gradient magnets are quite similar to the combined function gradient magnets used throughout the Booster Ring. The pole faces are nearly identical, with minor differences because the Recycler permanent magnets do not need water-cooling. They are named RGF for focusing and RGD for defocusing. Figure 1, on the next page, shows the cross-sections and values of magnetic pole fields.

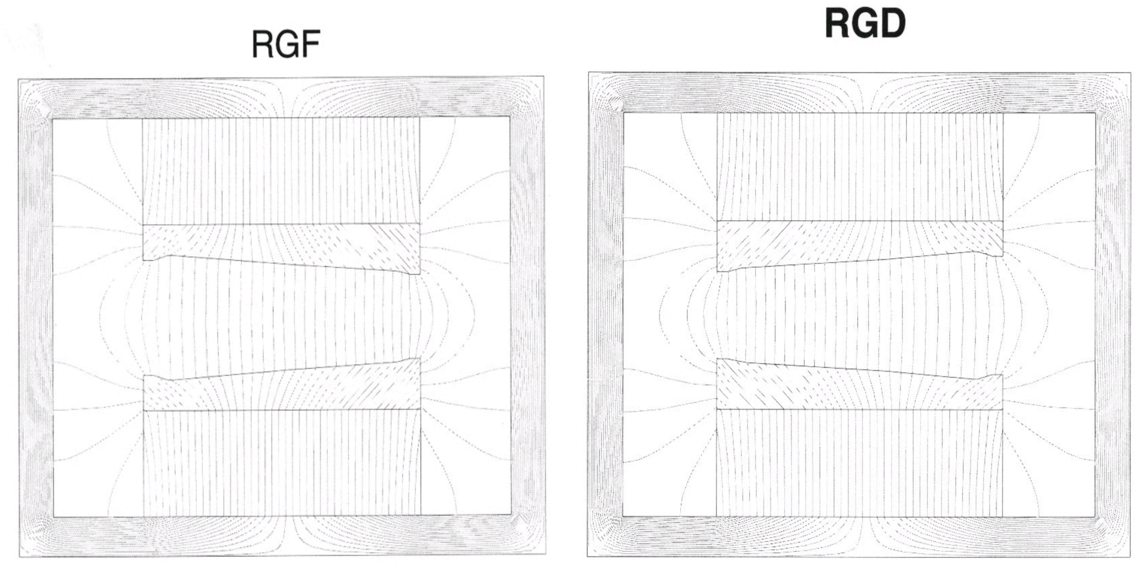

	RGF magnet		RGD magnet	
Height	9.0 in	22.9 cm	9.0 in	22.9 cm
Width	11.5 in	29.2 cm	11.5 in	29.2 cm
Physical Length	184 in	4.673 m	184 in	4.673 m
Magnetic Length	177 in	4.496 m	177 in	4.496 m
Dipole Field	1.375 kG		1.375 kG	
Quad Field	3.355 kG/m		-3.238 kG/m	
Sext Field	3.71 kg/m^2		-6.42 kG/m^2	

Figure 1: RGF and RGD magnet pole-face diagrams and field strengths.

Look closely at the field lines between the pole tips of the gradient magnets. Notice that they are not perfect dipole fields. Due to the arc of the magnetic field lines there are quadrupole fields introduced. We will leave it to the reader to use the right hand rule to prove the magnet configurations work with negatively charged antiprotons moving out of the page with a downward magnetic field. A similar discussion can be found in the Booster Rookie Book Chapter 6: Magnets.

Dispersion suppressor gradient magnets

Dispersion suppressor cells are needed on either end of the straight sections to reduce the dispersion function in the straight sections that do not use traditional lattice magnets. To this end, special gradient magnets have been made for these cells, named SGF for focusing, and SGD for defocusing. The dispersion suppressor magnets appear similar to the regular arc magnets. The main differences between them lie in the length of the magnets, and the strength of the quadrupole field. Figure 2 shows their pole-face diagrams and field strengths.

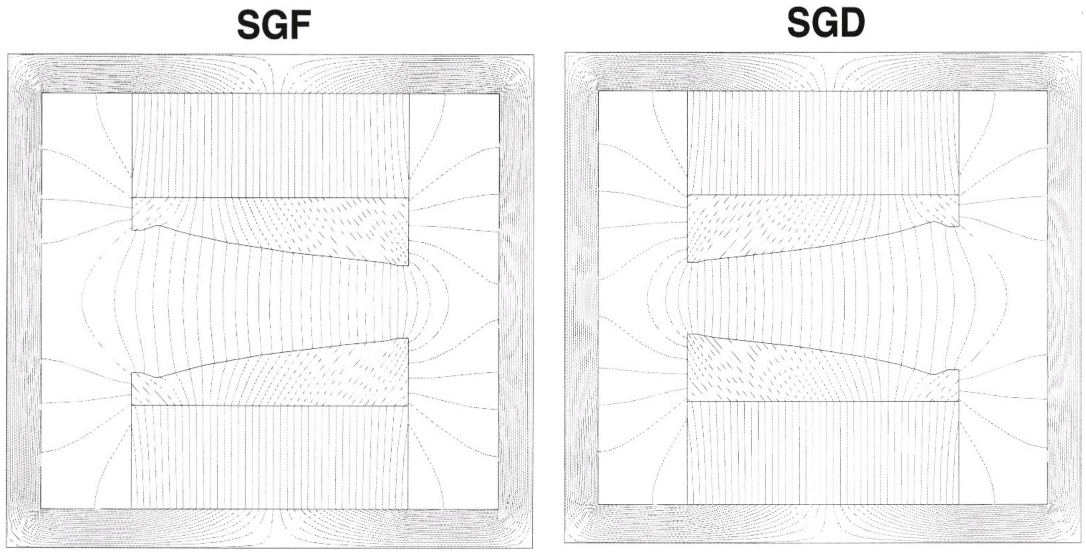

	SGF magnet		SGD magnet	
Height	9.5 in	24.1 cm	9.5 in	24.1 cm
Width	11.5 in	29.2 cm	11.5 in	29.2 cm
Physical Length	129.0 in	3.277 m	129.0 in	3.277 m
Magnetic Length	122.0 in	3.099 m	122.0 in	3.099 m
Dipole Field	1.330 kG		1.330 kG	
Quad Field	6.682 kG/m		-6.824 kG/m	
Sext Field	0.00 kG/m^2		0.00 kG/m^2	

Figure 2: Dispersion suppressor magnets pole face diagrams and field strengths. When compared to arc gradient magnets they have much stronger quadrupole fields despite their shorter length. This is achieved with steeper slopes on the magnetic poles.

Mirror Magnets

Mirror magnets were developed for use where gradient magnets are required, but there is not enough space on both sides of the beampipe to fit a conventional magnet (for example, where the transfer lines approach circulating beam pipe). Mirror magnets have all the magnetic material and the pole tip on one side of the beampipe, while a mirror plate occupies the other side of the beampipe. By putting a mirror magnet around both beam pipes, you now have fully functioning gradient magnets (at least as far as the beam is concerned) on both pipes. This reduces the separation to as little as 3.5” between the two beam pipes in the transfer lines, which allows the transfer lines to mimic the optics of the Main Injector and Recycler, making lattice matching much easier. This also allows designers to minimize the strength of the Lambertson and kicker magnets. Figure 3, on the next page, gives a cross section of these magnets.

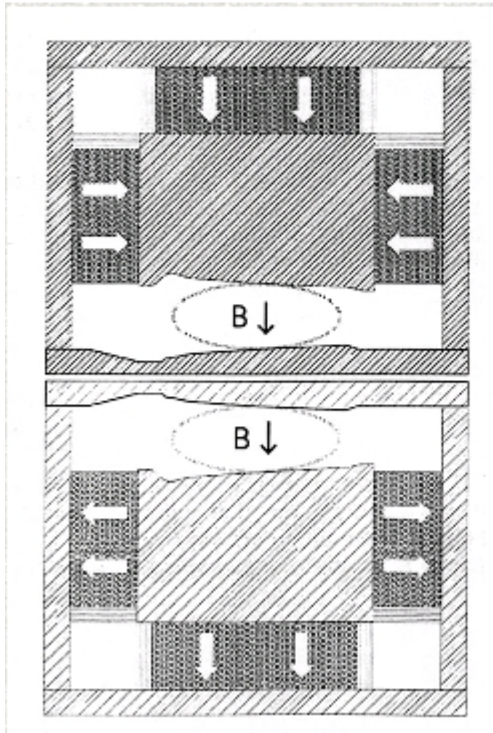

Figure 3: Mirror magnet cross-section.

There are three kinds of mirror magnets used in the RR transfer lines, with a grand total of 8 mirror magnets between them. Their layout will be covered in greater detail in the transfer line section, but we will mention their locations here as we discuss each type. Table 1 shows the differences between them.

	MGD magnet		MGS magnet		MDA magnet	
Height	10.0 in	25.4 cm	10.25 in	26.04 cm	10.00 in	25.4 cm
Width	12.5 in	31.8 cm	12.5 in	31.8 cm	12.50 in	31.8 cm
Physical Length	184.0 in	4.673 m	129.0 in	3.277 m	184.00 in	4.673 m
Magnetic Length	177.0 in	4.496 m	122.0 in	3.099 m	177.00 in	4.496 m
Dipole Field	1.375 kG		1.330 kG		1.375 kG	
Quad Field	3.238 kG/m		6.682 kG/m		0.00 kG/m	
Sext Field	-6.42 kG/m^2		0.00 kG/m^2		0.00 kG/m^2	

Table 1: Mirror magnet values. Field strengths vary by pole type for intended uses.

MGD magnets

There are 6 of these magnets used in areas where the beamlines meet circulating beam pipes. These magnets are magnetically equivalent to the RGD magnets. These are the only mirror magnets that are also used around circulating beampipe, so they require greater field quality across the horizontal aperture.

MGD

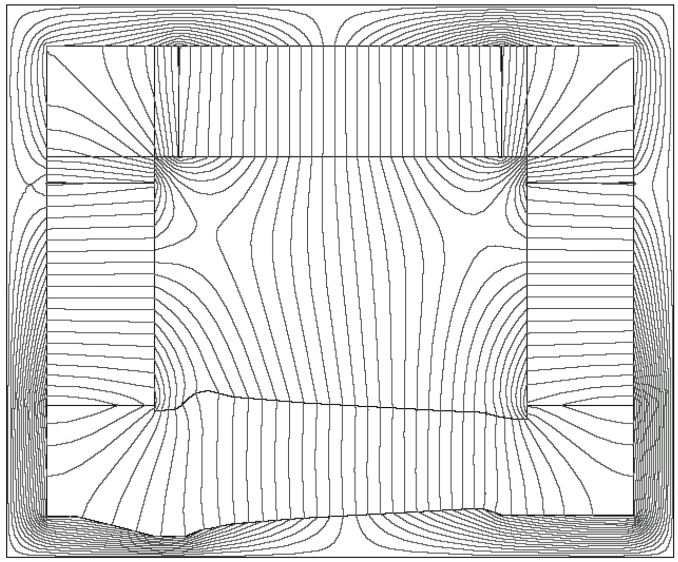

Figure 4: MGD magnet cross-section diagram.

MGS magnet

Only one of these magnets exist, and is used in the RR-32 beamline as opposed to making a large hole in a MI quad to allow for a regular gradient magnet. This is a magnetic clone of a SGF magnet, which requires the steeper slope on its pole tip face. The notch on the wide-gap side of the upper pole is required to terminate the field correctly.

MGS

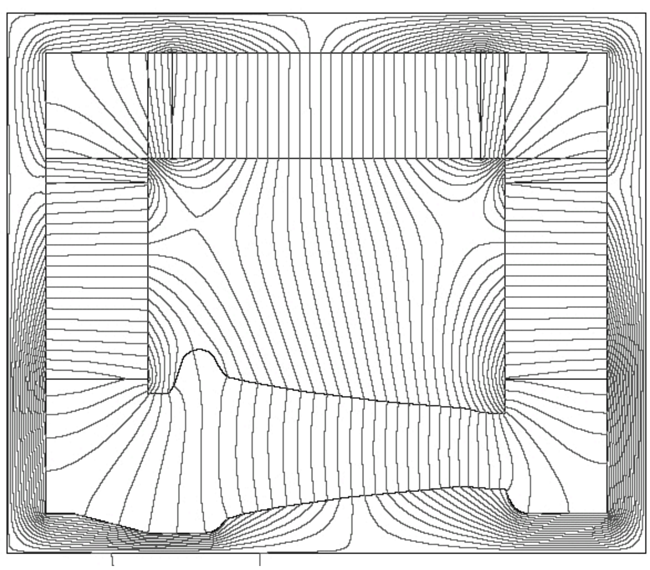

Figure 5: MGS magnet cross-section diagram.

MDA magnet

As with the MGS magnet, only one MDA magnet is required for the RR-40 abort line. This is a straight dipole field version of the MGD magnet that is used to merge the MI and RR abort lines just upstream of the dump.

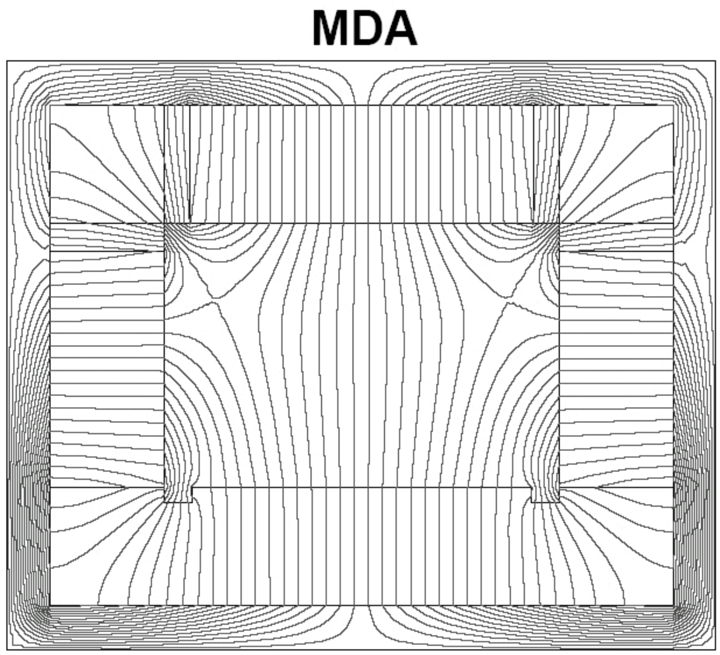

Figure 6: MDA magnet cross-section diagram.

Lambertson Magnets

Lambertsons are used as critical devices for many machines, but cannot be used as such in the Recycler because they are permanent magnets. Figure 7 shows their cross sections.

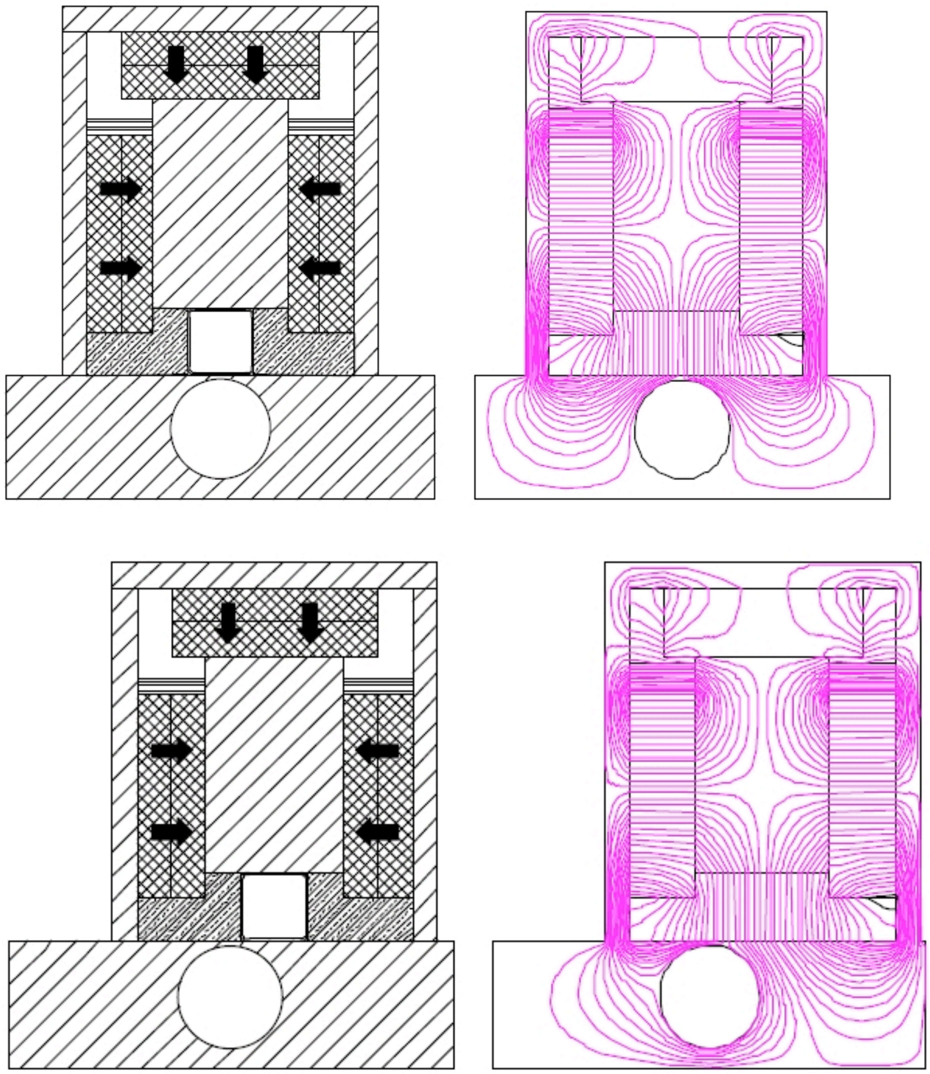

Front end

Back end

Figure 7: Lambertson magnet cross-sections. The field free region (bottom) angles through its length to reduce necessary field strength.

It is important to note that these Lambertsons, just like any other, do not have a true "field-free" region. There is bound to be some field leakage into this region, but we consider them "field-free" due to only small magnetic fields in this area.

An observant reader will notice the side ferrite bricks extend slightly past the bend region pole tip. This is done to help shape the dipole field. This also has the added benefit of allowing a flat pole tip, which is considerably less expensive to make.

Quadrupoles

While the gradient magnets used throughout the Recycler Ring do have a quadrupole field, in straight sections it is necessary to have a quadrupole field without the accompanying dipole field of the gradient magnets. Small quads were created for this purpose, the RQMF and RQMD quads.

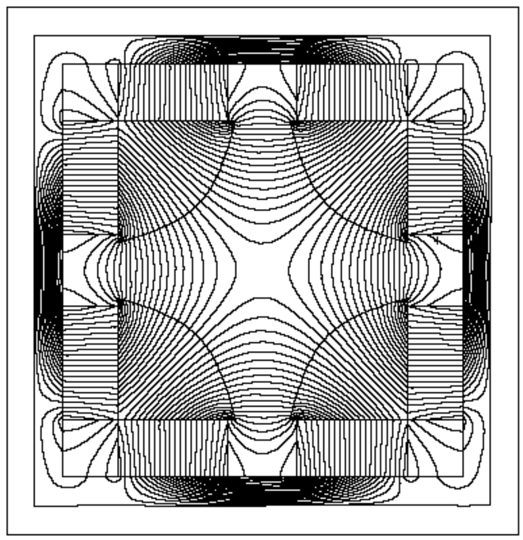

	RQMF magnet		RQMD magnet	
Height	8.25 in	20.95 cm	8.25 in	20.95 cm
Width	8.25 in	20.95 cm	8.25 in	20.95 cm
Physical Length	25.0 in	0.635 m	25.0 in	0.635 m
Magnetic Length	20.0 in	0.508 m	20.0 in	0.508 m
Dipole Field	0.00 kG		0.00 kG	
Quad Field	26.27 kG/m		-25.32 kG/m	
Sext Field	0.00kG/m^2		0.00kG/m^2	

Figure 8: RQMF/RQMD quad magnet cross-section and strengths.

Other quads were designed for special uses at different locations around the ring. Most have a stronger magnetic field than the RQMx quads to fulfill their specific purpose. First of these are the quads used in the abort line. These are named the RQAA, RQAB, & RQAC quads.

Another special purpose section is the phase trombone built at MI-60. We'll digress here for a few seconds to discuss a phase trombone. Typically, a synchrotron tune adjustment is done globally by varying the strength of the quadrupoles. As the Recycler cannot do this due to its permanent magnets, another method is required. Several designs were considered for this phase trombone, but the one settled on is to create an insert where you can very the strength of powered quads individually. This creates a section where you can match the beta phase, beta phase advance, and dispersion across the section, or vary them as you wish. By varying the beta phase through this section, you create a phase trombone. In the Recycler this was put in the 601-609 region. The RQSA, RQSB, RQSC, & RQSD magnets are used as phase trombone match quads. The RQSx quads are used on either end of this phase trombone to match the phase on both ends of the trombone to the connecting cells, allowing strict control of the betatron tune.

The Ecool insert in the RR-30 straight section requires another set of special purpose quads that change the lattice in the straight to optimize cooling. These quads RQRB-RQRH, and RQMH, stretch from half-cells Q301-Q309.

Permanent skew quads, named RQSS, are used in the MI-32 beamline to correct an unwanted skew quad field from the Lambertson. The comparable skew quads in the MI-22 line are powered, but they may in the future be replaced with permanent magnets once the field in this location is understood.

Sextupoles

Permanent sextupoles used around the extraction Lambertsons in the RR327–RR329 and RR213-RR215 sections. The normal gradient magnets surrounding the Lambertsons in these areas have been shifted off-center to create a permanent 3-bump. As these are gradient magnets, being off center creates a feeddown effect from the sextupole field. These permanent sextupoles are used to correct for this effect. These bear the names of RSXP, RSYP, and RSZP, and there is one of each in the 3-bump surrounding the extraction Lambertsons.

Notes:

Correction Elements

Correction elements are small, powered, air-cooled magnets located at each half-cell location in FODO lattice of the Recycler. Correction elements provide a means to compensate for imperfections in the main bending and focusing magnets, abnormality in the lattice configuration, and other destabilizing sources. In the Recycler, we use 3 different types of correction elements: dipole, quadrupole, and sextupole. The dipole can be oriented horizontally, vertically, or in a skewed configuration. The quadrupoles and sextupoles can be oriented as focusing, defocusing, or in a skewed configuration. The poles are arranged such that the angle between each pole is 360°/n, where n is the number of poles in the magnet. A skewed element is the same as a regular element except that all the elements are rotated 360°/2n. Below, is a description of each magnet, the effects its field has on beam, and the main function it serves.

Dipole

Dipoles have two pole faces orientated 180° from each other, as shown in Figure 9. Dipole fields bend the beam, thus changing its trajectory. Horizontal trims result in a horizontal bend, vertical trims result in a vertical bend, and skew dipole trims result in a bend in both planes. Skew correction dipoles are not used in the Recycler, as it is much easier correcting in one plane at a time. Combinations of correction dipole magnets can be used to steer beam through apertures or around obstacles, thereby reducing losses.

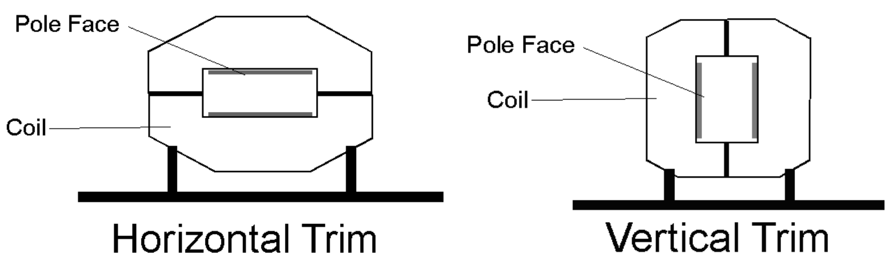

Figure 9: The horizontal trim and vertical trim have the same construction but different orientation.

Combinations of dipoles used for steering beam, as shown in Figure 10, are called “bumps.” Generally, a single correction dipole magnet is not changed by itself because it would disturb the entire Recycler orbit. Bumps are set-up in such a way that they make only local changes. Bumps are named for the number of like-correction dipole elements that are used. For example, a vertical 3-bump consists of three consecutive vertical correction dipoles working together to introduce a local vertical offset to the orbit. Bumps can have many elements in them but 3-bumps and 4-bumps are the most useful, introducing the most local bump to the orbit, and they are therefore the bumps most often used. If there is an aperture restriction at RR-103 and a 3-bump is performed around R:V103, the vertical correction element at RR-103, the orbit should only change between RR-101 and RR-105. Correction dipole magnets are used for bumping in both the horizontal and vertical planes and are highly effective at 8 GeV.

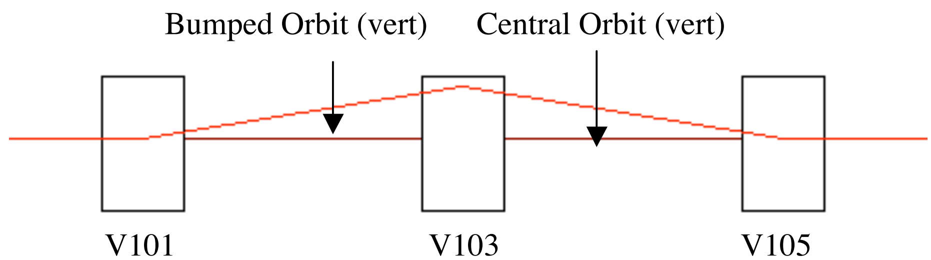

Figure 10: This is a typical vertical 3-bump around R:V103. Note that the orbit is changed very little outside of the local orbit bump.

Quadrupole

Quadrupoles, as shown in Figure 11, have four pole faces oriented 90° from each other. Quadrupole fields focus or defocus beam that does not pass through the center of the field, this also corresponds to the center of the beam pipe if the magnet is aligned properly. If a quadrupole field is set up to focus beam in the horizontal plane it will also defocus beam in the vertical plane due to the arrangement of the poles as illustrated in Figure 12. The focusing and defocusing by the gradient magnets, as shown in Figure 13, along with the permanent quadrupoles (not to be confused with the correction element quadrupoles) determines the tune of the machine in both the horizontal and vertical planes. The tune is simply the number of oscillations the beam makes about the closed orbit.

It is impossible to make a mono-energetic beam, or a beam consisting of particles with the same energy. This translates to a proportional momentum spread that affects the degree to which magnetic fields will deflect the beam's constituent particles. In a quadrupole, higher momentum particles will be focused less than lower momentum particles. This results in a lower tune for the higher momentum particle and a higher tune for the lower momentum particle. The dependence of the machine's tune spread due to the spread in momentum is called chromaticity.

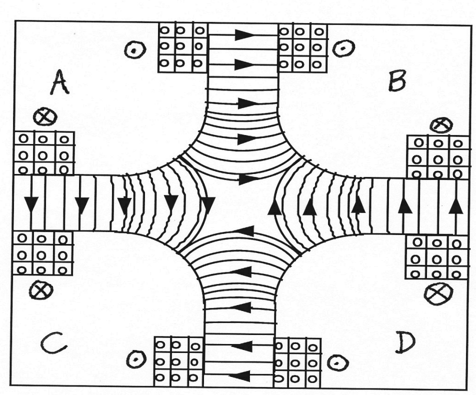

Figure 11: The four (usually iron) magnetic pole faces are labeled A, B, C, & D. The arrowed lines indicate the magnetic field. The direction of current flow (into or out of the page) is labeled near the square bus work. Notice the holes in the bus for cooling LCW. A given quadrupole can only focus the beam in one plane—while defocusing the beam in the other.

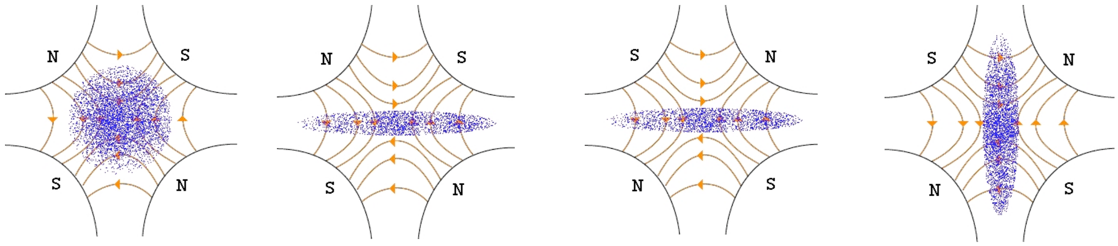

Figure 12: Given a round charge distribution acted upon by a quadrupole field, the charges will focus in one plane and defocus in the other. When the poles are reversed the plane that had been focusing will defocus and the plane that had been defocusing will focus. By convention, a quadrupole is named for its focusing property in the horizontal plane.

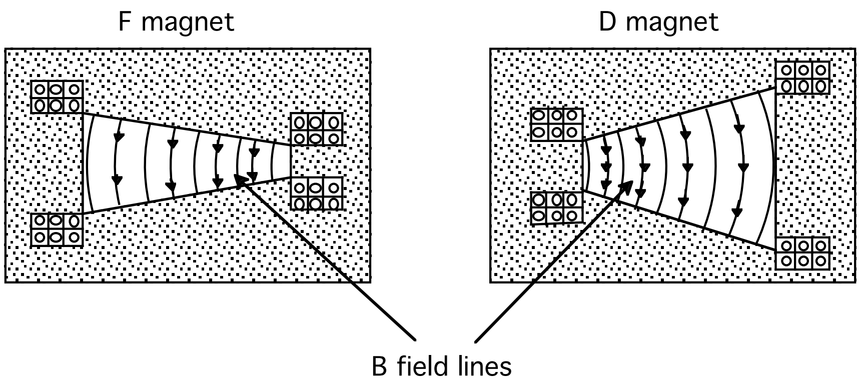

Figure 13: Typical gradient magnets in the Booster (they are very similar to Recycler gradient magnets). Unlike most quadrupoles, they also have a dipole component that bends the beam while focusing or defocusing. This dual-functionality comes from the unique shape of the magnet. In this figure, the "F magnet" is the focusing gradient magnet and the "D magnet" is the defocusing gradient magnet.

Quadrupole correction magnets can also be rotated by 45° so the pole faces are parallel to the horizontal and vertical planes. This is referred to as a skew correction quadrupole, as shown in Figure 14. The skew quadrupole correction elements are used to control the coupled resonance that occurs because the horizontal and vertical tunes are not completely independent of one another. The primary reason for the two planes being coupled is that the ARC gradient magnets have higher order fields present. Misalignment in these gradient magnets contributes to the coupling of the planes, but to a lesser degree.

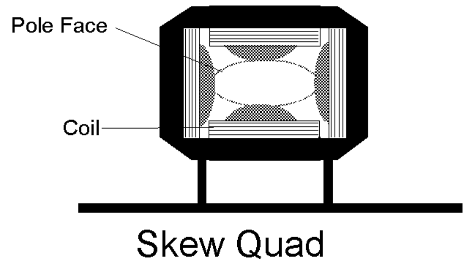

Figure 14: This is a typical skew quadrupole. Note that the pole faces are now at the top, bottom, and sides rather than at the corners. They have been offset by 45° from a regular quadrupole.

Sextupole

Sextupoles have six pole faces orientated 60° from each other, as shown in Figure 15. Sextupoles are used to adjust the chromaticity introduced by the quadrupoles. When a particle enters a sextupole field, the tune correction corresponds to the particles distance from the central orbit. The further the particle is from the central orbit, the larger is its resulting change in tune. Sextupole magnets use the dispersion caused by dipole fields to control the chromaticity arising from the quadrupole magnets. Both horizontal sextupoles and vertical sextupoles are used to control the tune spread in both planes. Keeping the tune spread as small as possible is important because of closely spaced resonance lines. If the tune spread were too large, many resonance lines could be crossed resulting in loss of beam.

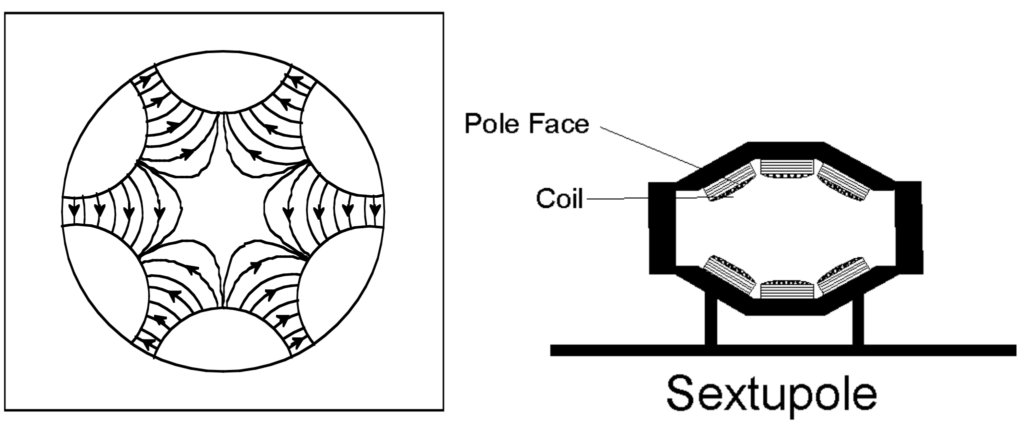

Figure 15: On the left is a field line diagram of a sextupole. Each of the poles has the opposite charge of its neighbors. The field is strongest near the poles with no field in the very center. On the right is the design of sextupoles used in the recycler.

Each of the correction elements has its corresponding optical analogy. Dipoles, which correct trajectory, are analogous to mirrors. Quadrupoles, which focus or defocus the beam, are analogous to lenses. Sextupoles, which adjust chromaticity, act as non-linear optical elements with variable focal length that increases or decreases with distance from the central axis through the element.

Correction Element Ramps

Why do we ramp the correction elements? In the Booster, Main Injector, and Tevatron we also have ramped correction elements, but their function is to counter the effect of acceleration and maintain a relatively constant orbit. The Recycler, running at a constant 8 GeV, does not have to compensate for its own acceleration. Instead, the recycler correction elements ramp to counter the effect of the Main Injector's ramp on beam in the Recycler. Without compensation, beam in the recycler heats up and lifetime suffers dramatically.

Even though all of the above correction elements have the ability to be ramped, only the dipole correction elements are ramped, since higher order effects due to the Main Injector are minimal. The ramps are generated by a C453 card located in the electronics room of a service building. These ramps are based on the measured Main Injector current as read from MDAT. The two correction element bulk supplies and the correction element regulator chassis are also located in the electronics room near the C453 card.

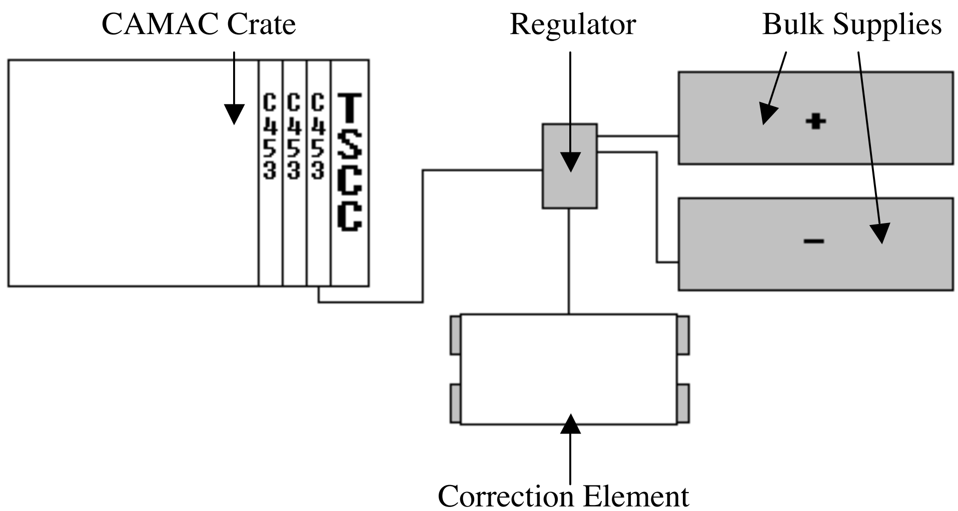

Figure 16: The correction element ramp components. The C453 supply the ramp to the correction element regulator that, in turn, provides the necessary current from the bulk supplies to the correction element.

The bulk supplies provide positive and negative voltage to the regulator, as shown in Figure 16. The regulator then regulates the required output according to the waveform from the C453 card. The regulator chassis contains status and control cards for both the correction elements and the bulk supplies.

Orbit Compensation

Because the MI and Recycler share the same tunnel, there are stray magnetic fields from the MI bus that interfere with the Recycler beam orbit. A fair amount of magnetic shielding absorbs some of these stray fields, but not all of them.

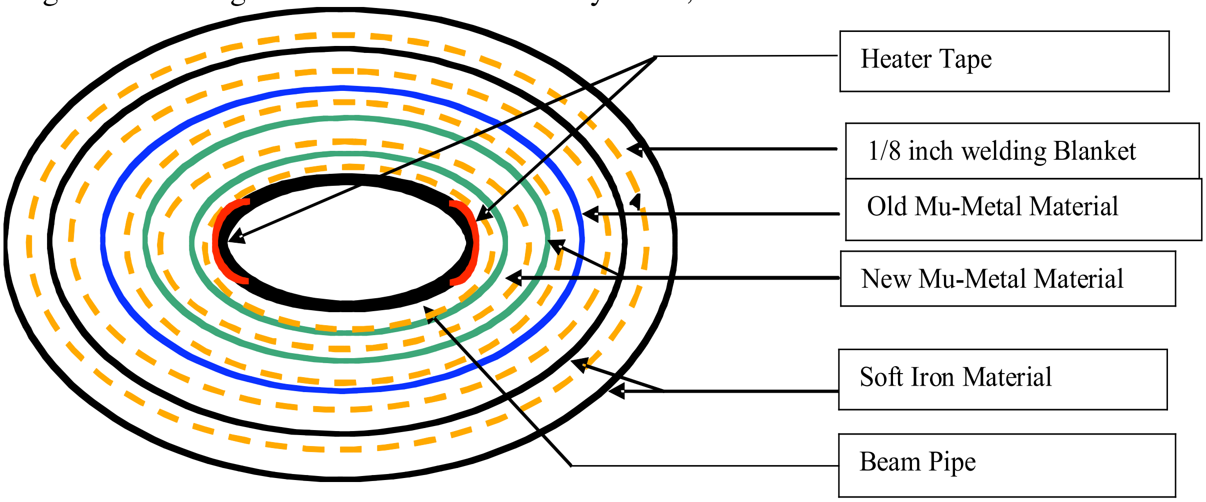

Figure 17: Beam pipe cross-section diagram.

It is necessary for the Recycler to have some method or methods of compensation for the Main Injector to keep the orbit stable. The Recycler employs four different methods to overcome the influence of the Main Injector on its beam. Each works to counteract a different problem.

Quad Compensation Loop (QCL)

It is important to remember that the Main Injector has ramping quadrupoles, not just dipoles. In fact there are two sets of ramping quadrupoles in the MI. The focusing and defocusing quads each have their own bus in the tunnel. They are completely separate busses, with current flowing in one direction for the focusing bus, and the other direction for the defocusing bus. Now if these two busses were both the same distance from the Recycler, then the magnetic fields from the two busses would cancel each other out. Generally this is the case, but what if the 2 busses carry different currents through them? This is indeed the case with the MI quad's busses.

The MI horizontal tune is 26.425, and the vertical tune is 25.415. This difference in the integer number of the tunes means that the current through the defocusing buss is smaller than the focusing bus. Different currents mean different magnet field strengths, so they won't cancel for the Recycler.

This is fixed by running a loop of cable around the ring, parallel to the quad busses. A PS named R:QCLP housed at MI-60 drives current through the cable, compensating for the current difference in the quad busses of the MI. This supply references a 465 card with tables for each MI ramp event that listen to the MI quad difference current (MI MDAT 62) and ramps the current in the cable with respect to this input.

Ramp Compensation

With the permanent magnets and the powered correction elements, a base orbit can be defined. The correction elements are not set to zero; instead they have base values that are determined by the $2x event G tables in their 453 cards. These reference the MI measured current to determine their value during each MI ramp cycle. In this way the base orbit is compensated if the MI bus is energized or not.

In addition to the base orbit, an orbit correction must be applied to compensate for the MI dipole field while ramping. This is done using the $2x H tables, which again reference the MI measured current to decide the amount to ramp the dipole correction elements. The quads, skew quads, and sextupoles are not ramped.

Path Length Correction

Now that the transverse effects of the MI have been corrected by the dipoles, one must turn attention to the path length distortions that this causes. As dipoles ramp, they cause the beam to change its path length around the Recycler ever so slightly. Add up each of these corrections, and the beam is making a longer journey around the ring. This "breathing" in turn causes the beam to take longer to make one complete circuit around the ring. If this keeps happening the beam quickly becomes out of phase with the RF, and the beam will begin to slosh around in the bucket.

This is corrected with a 5-bump in the 520 region that changes with each ramp cycle to better match phase. This bump utilizes the $2s F(t) tables for these correctors: R:H518, R:H520, R:H522, R:H524, & R:H526. Figure 18 below shows the $23 ramp for R:H518. This will show you each of the ramps associated with the base orbit, the ramp compensation, and the path length correction.

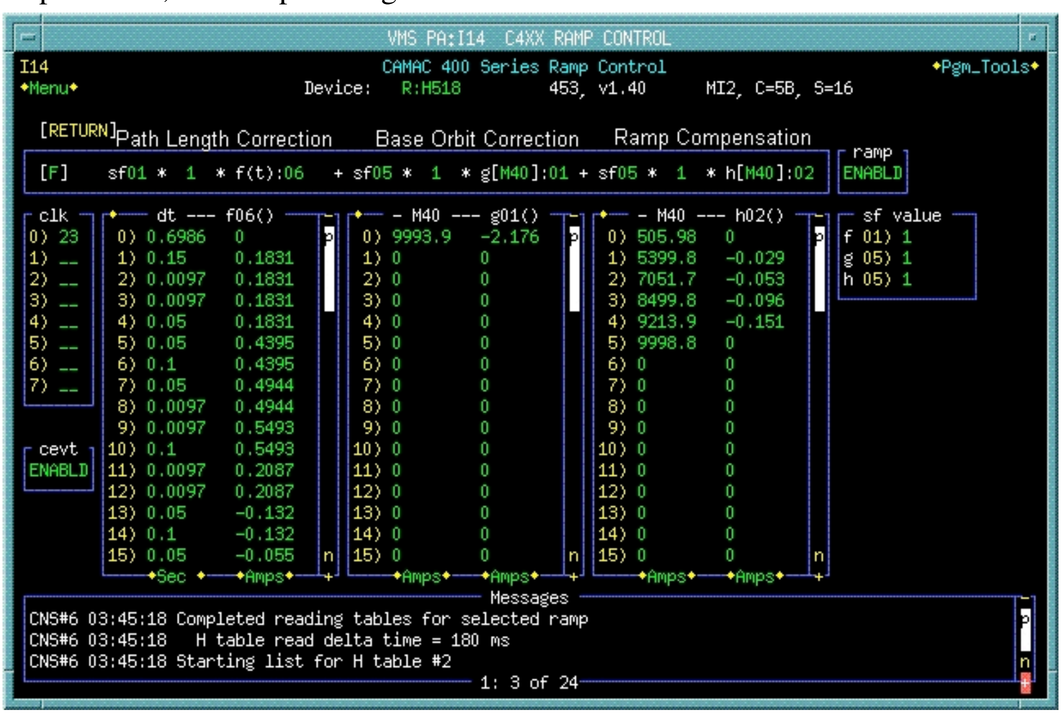

Figure 18: $23 ramp table for R:H518 showing Path Length, Base Orbit, and Ramp Compensation correction ramps.

Injection Counterwave

During extraction from the Recycler, a 40mm oscillation is needed from the kicker to reach the field region of the Lambertson. The problem with this is that the Recycler circulating beam pipe aperture is not that big. To find a way around this, a 20mm counterwave is introduced via a horizontal 3-bump between the kicker and the Lambertson. Thus the total oscillation is within the beam pipe aperture until the moment of extraction. This bump is also needed when injecting beam from the MI to move beam close enough to the central orbit for the kicker to finish the job. This counterwave is defined by the $Ex F(t) tables for each injection bump:

RR22 antiproton injection: R:H130, R:H210, & R:H212
RR32 proton injection: R:H330, R:H340, & R:H400

As a result of the number of oscillations beam goes through in the arc cell magnets, and thus their sextupole feeddown effects, a separate time bump is included to reduce the tune shift caused by the RR22 injection bump as it ramps. This 3-bump uses the following correctors: R:H410, R:H412, & R:H414. There is also a comparable counterwave performed in the Main Injector.

Critical devices

As mentioned before, the injection Lambertsons in the Recycler are permanent, and cannot be used as critical devices. Instead two pairs of vertical dipoles in the beamlines are used. The devices are named R:VDPA1, and R:VDPA2. This stands for Vertical (Vernier) Dipole, Potted, Type A.

Each PS powers elements in both the MI-22 and MI-32 transfer lines, that way it gives the requisite two inputs to the safety system that will prevent beam from reaching the Recycler. The magnets appear in pairs, so that there are a total of eight of these magnets, four for each major bend in each line.

Transfer Lines

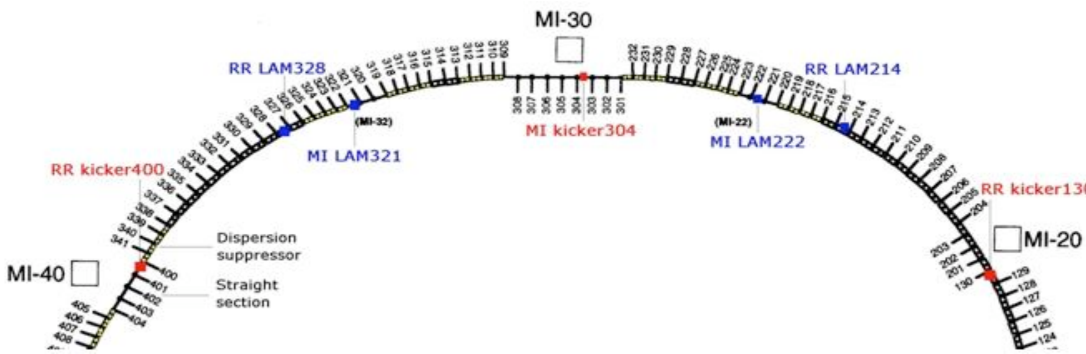

Now that we have learned about the magnets that make up the Recycler Ring, we will move on to how they are used in the transfer lines. Even in the transfer lines, permanent magnets are used predominantly. There are a few exceptions, such as the critical devices that are powered, but again most elements are of the permanent magnet variety.

There are three beam lines that transfer beam into and out of the Recycler. Two of them originate in the MI-30 straight section, heading in opposite directions to meet the Recycler just outside of straight sections RR-22, and RR-32. The other is the Recycler abort line that sends circulating protons to the MI abort at MI-40. As a cost-cutting measure, the same kickers are used in several different transfer schemes. This allowed the Recycler to use only three kickers, instead of the six you would expect for a machine with three different transfer lines.

MI-22 Antiproton MI to RR transfer Line

The MI-22 transfer line is the line used to inject antiprotons into the Recycler from the Main Injector. It can also be used to send protons from the Recycler to the MI, but as it is primarily a antiproton transfer line, we will concentrate on that mode. Powered elements in this line are given numbers in the 700's, like the AP-2 line in the Pbar Source. Figure 19 on the next page shows a drawing of the beamline with all the elements shown.

The antiproton beam transfer begins with a horizontal kick to the outside from R:KPS3A at the 304 location. This kicker, while residing in the MI, is given an R: device name as its only function is to kick beam into or from the Recycler. From there beam travels to the field region of the RR Lambertson at 222. This bends the beam upwards with a 23mr angle. As the MI-22 area is a short straight section, permanent quadrupoles are used instead of gradient magnets.

Following this is the pair of VDPA magnets powered by R:VDPA1. These bend the beam trajectory down, parallel to the Recycler. After this bend one of the CE skew quads used to correct skew quad fields from the Lambertson.

The end of the parallel portion of the beamline comes as we enter the second pair of VDPAs, these ones controlled by R:VDPA2. These bend the beam upward again toward the Recycler.

Now as we near the Recycler we meet two mirror magnets, both of the MGD variety. This leaves just the kicker-kicker combination to complete the line. The Lambertson bends beam down the needed 23mr to bring it level with the Recycler beampipe. Beam is still radially to the inside of the Recycler orbit, so we need a kicker to put the beam on the circulating orbit. R:KPS2A gives it a 1.17mr kick to the outside. This is the beamline as antiprotons see it going into the Recycler. For protons out of Recycler, start at the end with KPS2A and read backwards.

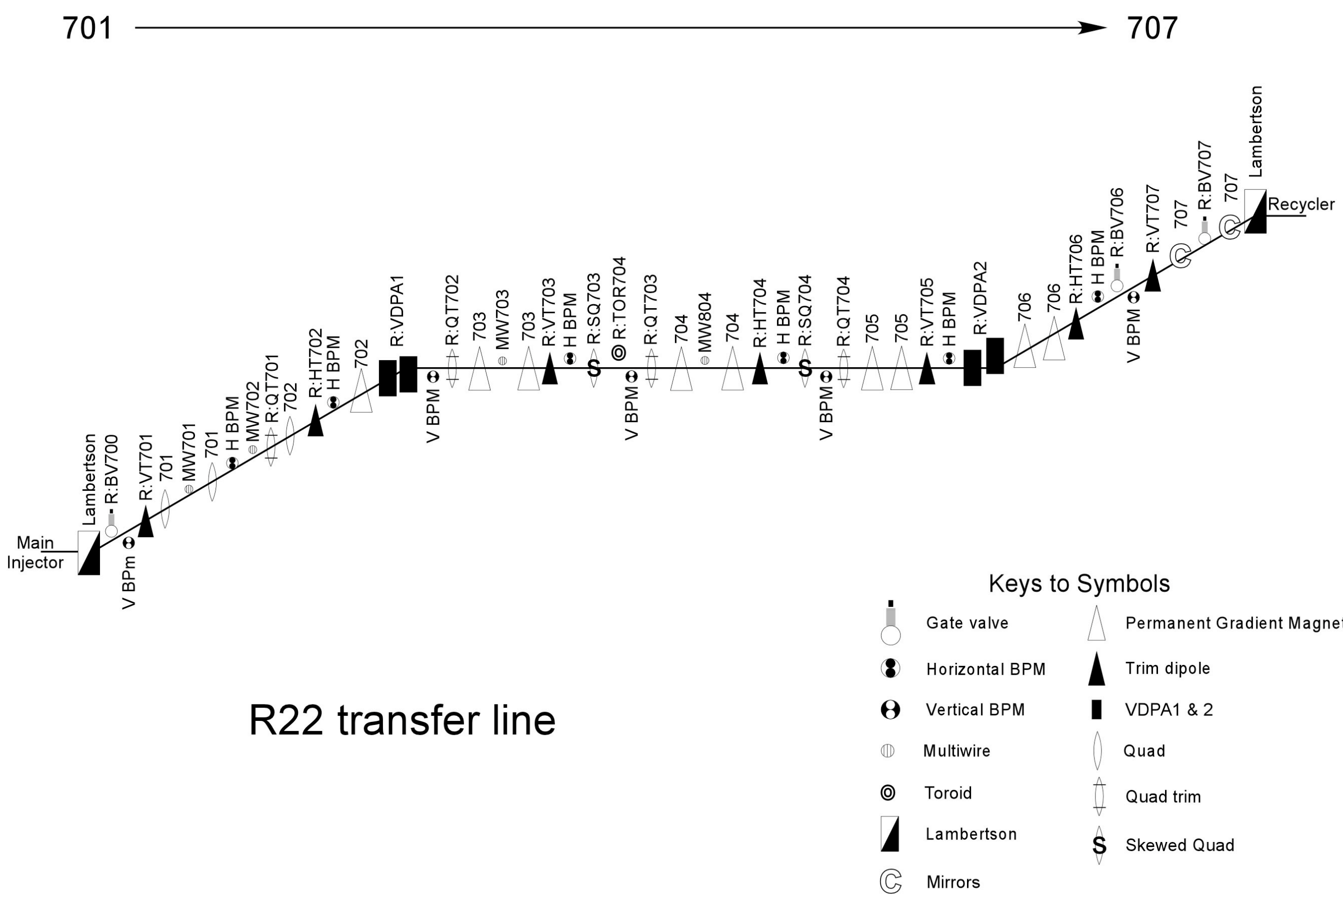

Figure 19: MI-22 Transfer line.

MI-32 Antiproton RR to MI transfer line

As with the MI-22 line, we will follow the MI-32 line with the antiproton trajectory, starting in the Recycler. This is easier to do as the beamline is numbered to follow the antiproton direction instead of the proton. The naming convention for the powered elements in the line puts these devices in the 800s. Refer to Figure 20 for the layout of the line.

Pbar extraction begins with R:KPS4A at the 400 location giving a 0.92mr kick to the outside. This kick, plus the bump from the counterwave, puts beam in the Lambertson field region at 328 to bend the beam down with a 23mr angle. Immediately following the Lambertson is a permanent trim skew quad, an RQSS magnet that corrects the skew field from the Lambertson. The beam progresses through two MGD mirror magnets, and the first VDPA dipole pair in this line, powered by R:VDPA2 that bends the beam parallel to the Recycler.

After a short section running parallel to the Recycler a downward kick is given from the two VDPA dipoles powered by R:VDPA1.

A 32mr upward bend comes from the Lambertson at 321. From here it is a quick jaunt to the R:KPS3A kicker at 304 to nudge the beam onto the MI orbit with a 1.16mr kick.

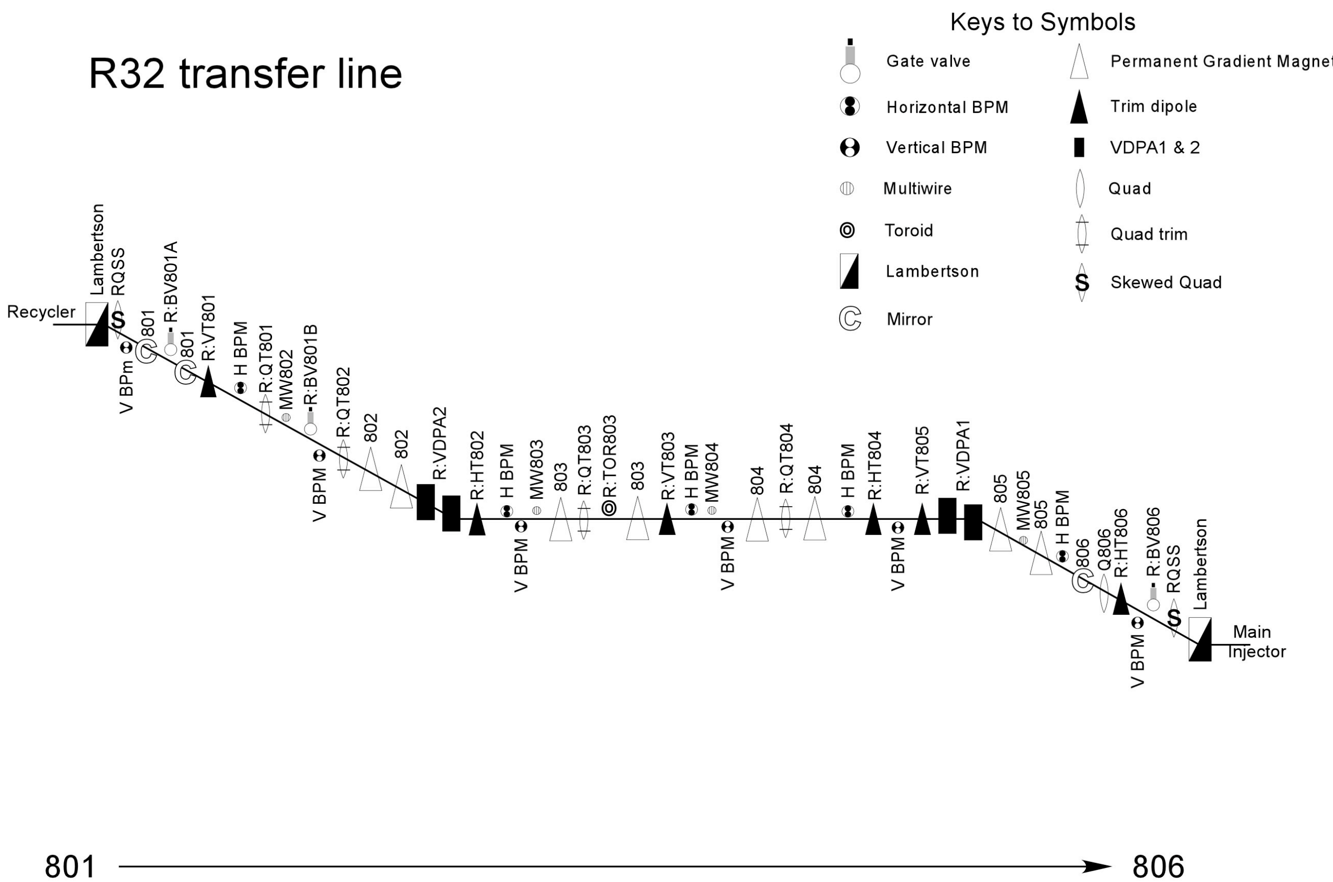

Figure 20: MI-32 Transfer Line

RR40 Recycler Proton Abort Line

This beamline has just one purpose. When the antiprotons circulating in the Recycler need to be aborted, they are just scraped away with collimators, or are slammed into beam valves when they closed for access. With circulating protons, there exists a possibility for greater intensities that require a proton abort line. The Recycler partially shares its abort line with the Main Injector as a way to cut costs. Powered elements in the line are numbered in the 000's. Figure 21 shows the RR-40 abort line layout.

The Recycler abort line starts with a kick to the outside from R:KPS4A in the 400 area. After the kick it is a short distance to the MDA mirror magnet that serves to merge the two abort lines into one.

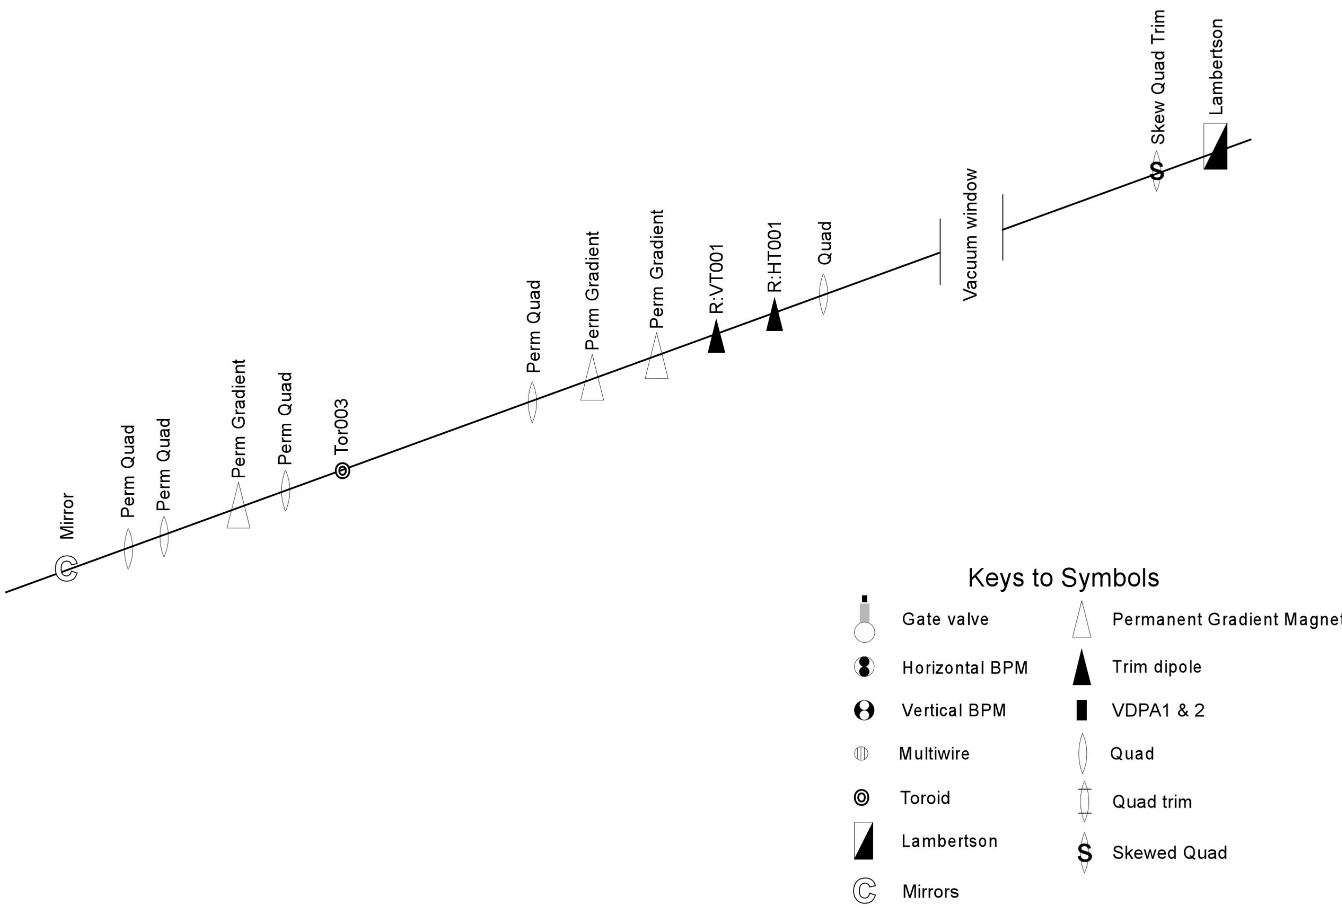

Figure 21: RR-40 Abort Line

Recycler Injection Scenarios

With the various components of the Recycler injection/extraction lines being used for multiple uses, it is worthwhile to familiarize yourself with the roles these components take, and the timing that allows this all to happen smoothly. We will discuss the VDPAs and their associated events first. Figure 22 shows how the VDPA power supplies power the magnets in the tunnel.

As we mentioned before, the VDPAs are vertical dipoles in the MI-22, and MI-32 transfer lines that are used as critical devices to prevent transfer between the MI and Recycler. Two power supplies deliver current to magnets in each line simultaneously as the magnets are wired in series. So, when the antiproton beam is being injected into the Recycler through the MI-22 line, the VDPA magnets in the MI-32 line are also being ramped.

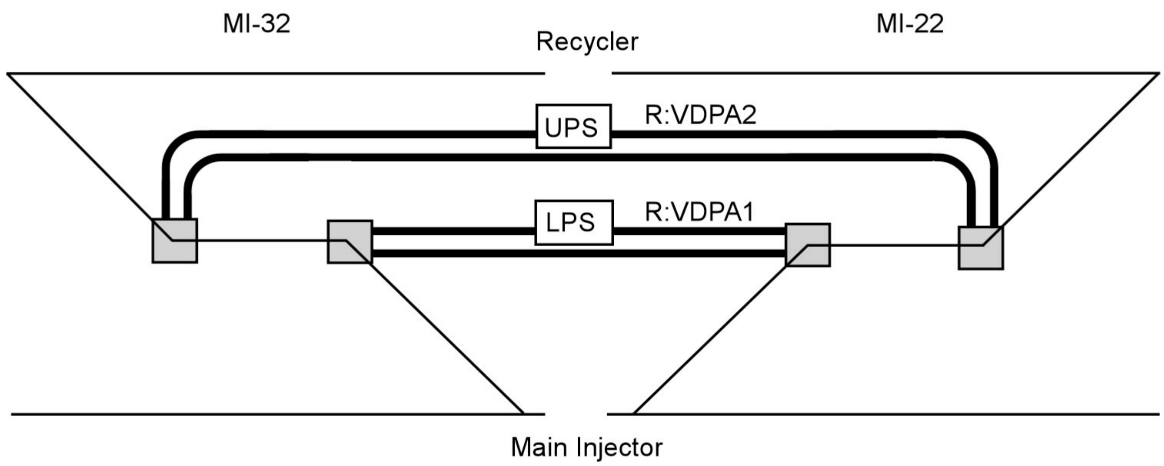

Figure 22: VDPA PS and Magnet Layout.

You may have noticed that the MI-22 line is half a cell longer than the MI-32 line. As a result of this, the VDPAs have less overall length to accomplish their bends. This requires that the magnets run at different currents in each line. This is done by using a 055 card and setting up different ramps for each injection/extraction event. Table 2 lists the associated events and device database names for each transfer.

TCLK	Transfer	Transfer Line	Database Names
$E0	Antiprotons from MI to RR	MI-22	R:VDUE0 & R:VDLE0
$E2	Protons from MI to RR	MI-32	R:VDUE2 & R:VDLE2
$E3	Protons from RR to MI	MI-22	R:VDUE3 & R:VDLE3
$E4	Antiprotons from RR to MI	MI-32	R:VDUE4 & R:VDLE4
$ED	Prepare for Injection	Either	R:VDUED & R:VDLED
$EE	Prepare for Extraction	Either	R:VDUEE & R:VDLEE

Table 2: Timing events for various Recycler injections/extractions.

As a result of the different uses for each line, different amplitudes are needed from the kickers. The RR 20 kicker (R:KPS2A) has two purposes: Pbar injection from MI into RR, and proton extraction from RR into MI. Due to the different electric fields from each particle (+ for proton, - for antiproton), different kicker strengths are needed. Each of these kicks requires different amplitudes and requires different triggers, as shown in Table 3.

TCLK	RRBS	Reflected TCLK	Description	Database Names
$E0	$A0	$F0	Antiprotons from MI to RR	R:KPS2A0
$E3	$A3	$F3	Protons from RR to MI	R:KPS2A3

Table 3: RR-20 kicker timing information.

The RR 30 kicker (R:KPS3A) has even more to do than the RR 20 kicker. This kicker handles all beam transfers between the MI and RR. Of course, each of these kicks have different amplitudes and triggers listed in Table 4.

TCLK	RRBS	Reflected TCLK	Description	Database Names
$E0	$A0	$F0	Antiprotons from MI to RR	R:KPS3A0
$E2	$A2	$F2	Protons from MI to RR	R:KPS3A2
$E3	$A3	$F3	Protons from RR to MI	R:KPS3A3
$E4	$A7	$F7	Antiprotons from RR to MI	R:KPS3A7

Table 4: RR-30 kicker timing information.

The last of the three kickers at RR-40 (R:KPS4A) has three different purposes, again with different amplitudes and triggers listed in Table 5.

TCLK	RRBS	Reflected TCLK	Description	Database Names
$E2	$A2	$F2	Protons from MI to RR	R:KPS4A2
$E4	$A7	$F7	Antiprotons from RR to MI	R:KPS4A7
$E7	$D0	$F5	Protons to RR abort	R:KPS4D0

Table 5: RR-40 kicker timing information.

The different amplitudes can be found on ACNET page R66, shown in Figure 23. Each kicker and its referenced events are laid out on pages for easy reference.

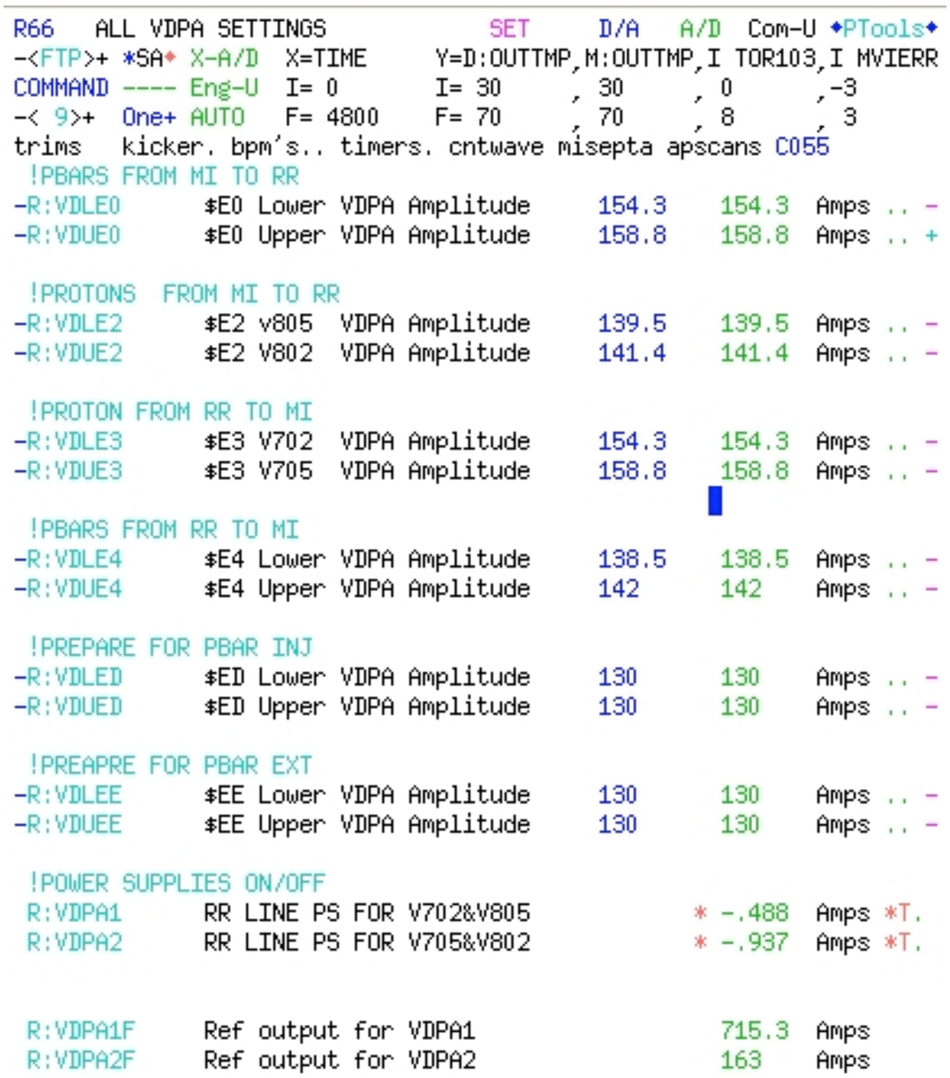

```
R66   ALL VDPA SETTINGS            SET     D/A    A/D   Com-U *PTools*
-<FTP>+ *SA* X-A/D  X=TIME     Y=D:OUTTMP,M:OUTTMP,I TOR103,I MVIERR
COMMAND ---- Eng-U  I= 0       I= 30     , 30     , 0      ,-3
-< 9>+  One+ AUTO   F= 4800    F= 70     , 70     , 8      , 3
trims   kicker. bpm's.. timers. cntwave misepta apscans C055
 !PBARS FROM MI TO RR
-R:VDLE0       $E0 Lower VDPA Amplitude     154.3    154.3  Amps .. -
-R:VDUE0       $E0 Upper VDPA Amplitude     158.8    158.8  Amps .. +

 !PROTONS  FROM MI TO RR
-R:VDLE2       $E2 v805  VDPA Amplitude     139.5    139.5  Amps .. -
-R:VDUE2       $E2 V802  VDPA Amplitude     141.4    141.4  Amps .. -

 !PROTON FROM RR TO MI
-R:VDLE3       $E3 V702  VDPA Amplitude     154.3    154.3  Amps .. -
-R:VDUE3       $E3 V705  VDPA Amplitude     158.8    158.8  Amps .. -

 !PBARS FROM RR TO MI
-R:VDLE4       $E4 Lower VDPA Amplitude     138.5    138.5  Amps .. -
-R:VDUE4       $E4 Upper VDPA Amplitude     142      142    Amps .. -

 !PREPARE FOR PBAR INJ
-R:VDLED       $ED Lower VDPA Amplitude     130      130    Amps .. -
-R:VDUED       $ED Upper VDPA Amplitude     130      130    Amps .. -

 !PREAPRE FOR PBAR EXT
-R:VDLEE       $EE Lower VDPA Amplitude     130      130    Amps .. -
-R:VDUEE       $EE Upper VDPA Amplitude     130      130    Amps .. -

 !POWER SUPPLIES ON/OFF
 R:VDPA1       RR LINE PS FOR V702&V805           * -.488  Amps *T.
 R:VDPA2       RR LINE PS FOR V705&V802           * -.937  Amps *T.


 R:VDPA1F      Ref output for VDPA1                 715.3  Amps
 R:VDPA2F      Ref output for VDPA2                 163    Amps
```

Figure 23: Picture showing different amplitudes of VDPA field. The VDPAs have different fields because of the different lengths of the lines.

Recycler Beam Permit System & the RR-40 Dump

The Recycler has a dedicated beam permit/abort loop. The Recycler Beam Permit (BP) Loop is a serial loop of CAMAC 200, or C200, modules. A separate C201 module at MI-40 sends out a 5 MHz signal. If each C200 module has no faults, this signal is passed along and returned to the C201. If any of the eight inputs to a C200 module is low, the 5 MHz signal is not passed along and the loop collapses. Inputs to the RR BP can be made at Abort Fan In-Panels and C200 modules located in the Main Injector Service Building electronics rooms. Only a limited number of devices will pull down the permit loop. A list of the abort devices for each Recycler house can be found in the table 6 on the next page. These inputs, along with their trip and mask status can be found on the Recycler Beam Permit Status page, R67.

House	Input	Description
MI8	-	None
RR10	2	Vacuum System (BV101)
	3	Vacuum System (BV103)
RR20	2	Vacuum System (BV201)
	3	Vacuum System (BV213)
RR30	1	Safety System (CBV 302&308)
	2	R:VDPA1
	3	R:VDPA2
	4	R:BV700 (R22 beam line)
	5	R:BV806 (R32 beam line)
	6	Beam valve 305
	7	Beam valve 307
RR40	2	Beam valve 401
RR50	-	None
RR60S	-	None
RR60N	2	Beam valve 610

Table 6: Recycler Beam Status

Unlike the Tevatron, where dropping the beam permit will drop the abort link and fire the kickers, the Recycler has no abort link and therefore does not fire the kickers when the beam permit drops out. Furthermore, there are three kickers used for beam transfers to and from the Recycler. Only one of these, R:KPS3A, is connected to the Recycler Beam Permit Loop. The beam permit system provides status to Beam Switch Sum Box (BSSB) and Beam Sync Clocks, RRBS and MIBS for the Recycler and Main Injector respectively. If the beam permit is down, the BSSB inhibits the pulse shift and the MIBS and RRBS systems inhibit the proton beam sync transfer events. In this way, the RR BP system functions like the PBAR beam permit system. This is only effective for protons. Antiprotons from the Accumulator do not need a pulse shift to be extracted and they may be injected into the Recycler without a beam permit present.

So how do we send protons to the dump? We use R:KPS4A, the kicker at RR-40, to transfer the protons down an abort line to the dump. The following is the sequence of events that is followed to extract to the dump:

- The Recycler Sequencer (R48) issues a TCLK $E7
- This triggers a RRBS $D0, subsequent to the next RRBS $AA
- A C279 triggers the kickers to fire upon receipt of the $D0

The RRBS $AA, like the $AA marker for any of the other machines, marks the starting point of beam in the Recycler and is synchronized with the RR RF system. The $C0 can be called the zero azimuth marker and is defined as the starting point (bucket zero) from which all the other bucket locations (up to 587) are defined. The kickers can be disabled through the C279 card.

Recycler RF

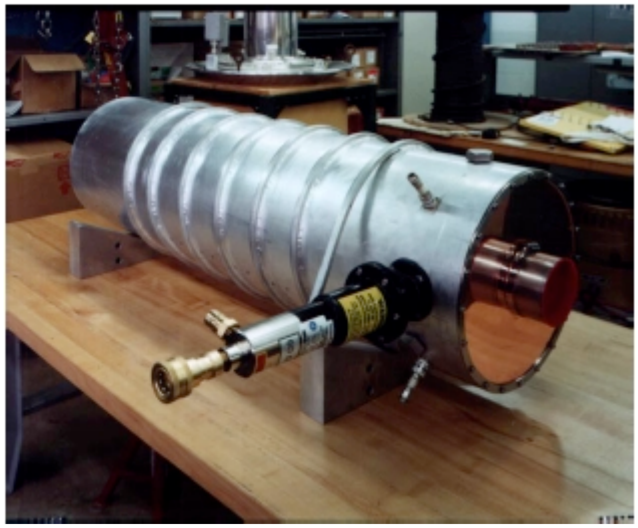

The Recycler uses RF to maintain control over its beam, like all the other machines at Fermilab. Unlike other machines however, it uses a barrier bucket system mainly to encapsulate the beam, not to accelerate it.

Both the Antiprotons Accumulator and the Recycler are Antiprotons storage machines; however, they differ in the way they store beam. The Accumulator stores the beam in a DC "ribbon" to the radial inside of injected beam. When the time comes to extract, it uses a 2.5 MHz system that slowly captures some of this beam and moves it out to the extraction position.

The Recycler stores beam longitudinally between barrier buckets. When it is ready to extract, the beam is separated into 9 longitudinal slices that are slid ahead of the rest of the beam one at a time. Each slice, or parcel, then has a 2.5 MHz sinusoidal signal applied before being extracted to the MI.

In this chapter we will discuss how this is accomplished using the Recycler RF system, and how the system itself has been designed.

Low-Level RF System

The Recycler houses its LLRF in a VXI crate in the MI-60 control room. This crate houses all the electronics necessary to generate the signals sent to the HLRF system.

Refer to Figure 24 for a diagram of the LLRF system. The signals begin with a 53 MHz digital oscillator that provides a synchronized signal to three different Direct Digital Synthesizers (DDS). These DDS's all run from the same clock, so there is no need for phase lock loops or feedback systems to maintain frequency synchronization.

Low-Level Block Diagram

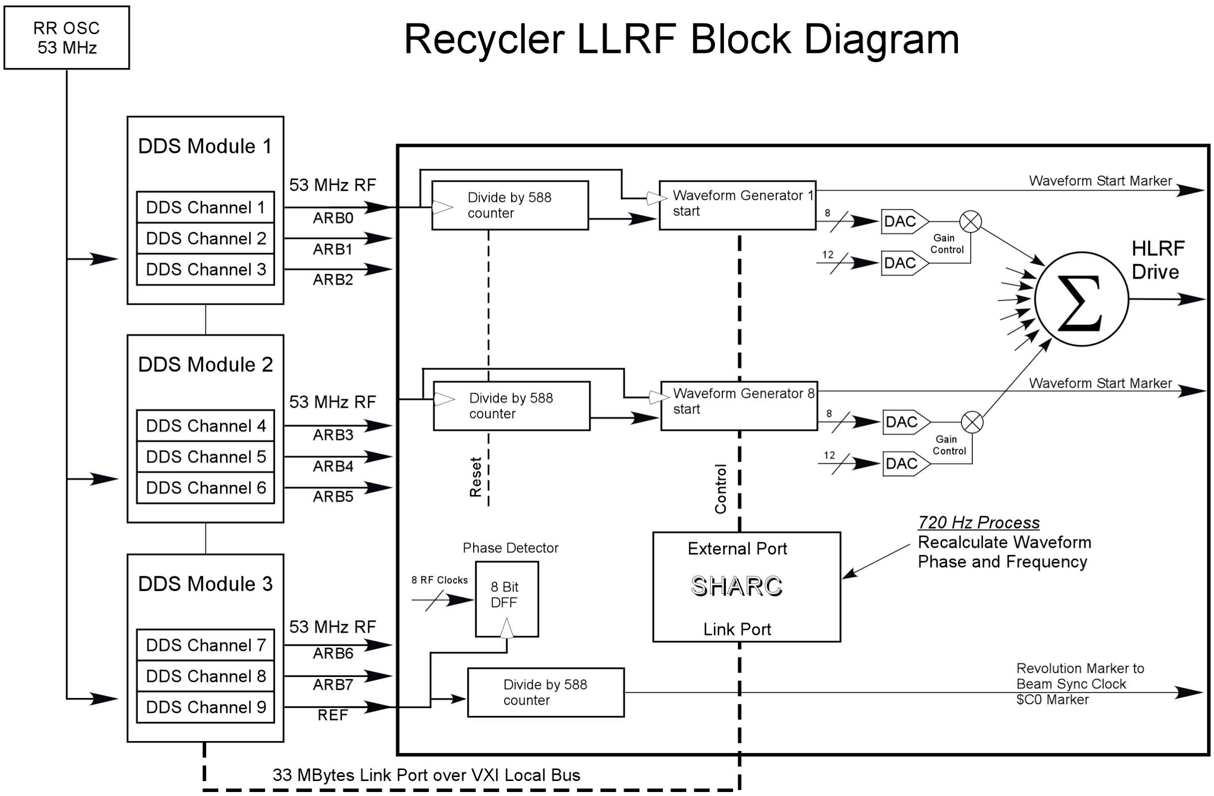

Figure 24: RR LLRF schematic.

Each of these Direct Digital Synthesizers has three outputs. From these outputs we have the origins of our 8 arbitrary waveforms (ARBs). Table 7 lists the ARBs and their common uses. Each can and will be used in different ways to support different functions, such as momentum mining and extracting.

ARB0	Injection barrier bucket (down)
ARB1	Injection barrier bucket (up)
ARB2	2.5 MHz
ARB3	Transition barrier bucket (down)
ARB4	Transition barrier bucket (up)
ARB5	AA marker
ARB6	Cold beam barrier bucket (down)
ARB7	Cold beam barrier bucket (up)

Table 7: RF waveforms and typical usage

The best-known waveforms are ARB6, and ARB7, which are used to hold the cold beam. ARB3 was used as the linear bucket waveform before Ecool came online. You might wonder why there are only 8 waveforms listed, but we have a total of 9 waveforms available from the DDS's. The last one is used as the $C0 marker. This marker tells the machine where the 0^{th} bucket in the machine is.

All the waveforms are generated by sending the DDS channel sine waves into comparators that generate pulses at 52.8 MHz. There is a divide by 588 counter that generates a pulse each turn, which then triggers the arbitrary waveform generators.

The arbitrary waveform generators are told which waveform to apply to each channel by software. The waveforms are then summed and sent to the HLRF system via cables.

Now would be a good time to cover some common uses of waveforms. We know what ARB6 and ARB7 are typically used for, as well as ARB0 and ARB1. However there are several processes that require the use of more ARB waveforms.

Injection into the machine requires the use of those four waveforms, plus ARB2 that is a 2.5 MHz signal. Technically all the waveforms are still being summed into the one HLRF drive signal, but the ones not needed for the chosen manipulation have the gain of that waveform set to zero. Figure 25 shows the RF waveforms as used for the injection process.

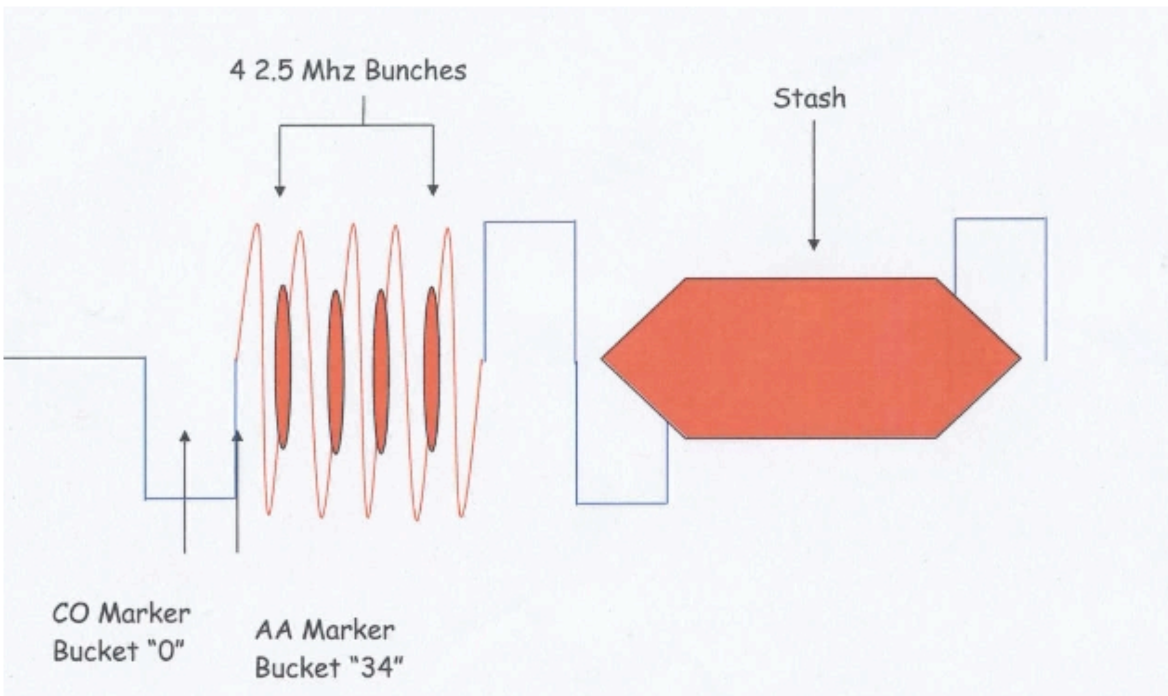

Figure 25: Injection RF scheme for RR.

Note the use of ARB0 and ARB1 and within them ARB2 that is holding the beam in 2.5 MHz buckets. After the beam is injected, the gain of ARB2 is lowered to zero, then the positions of ARB0 and ARB1 are moved to merge the new hot beam into the cold beam.

For the extraction process, a similar overlapping of ARBs is done to accomplish the goal of extracting 9, low-momentum spread, high-intensity bunches. Once ARB6 and ARB7 are in the final position, beam is mined into those nine bunches.

Manipulations such as momentum mining require the use of all the waveforms. Momentum mining is the process by which a lower potential bucket is swept across the cold beam bucket capturing the high momentum particles, and removing them from the colder beam. While this reduces intensity a small amount, the gains in having a smaller momentum spread are worth it. Tevatron luminosity is dependant on beam quality nearly as much as beam intensity and these lower momentum spread bunches enable the RR to send bunches to the Tevatron with smaller longitudinal emittance. Figure 26 is a drawing of the mining process.

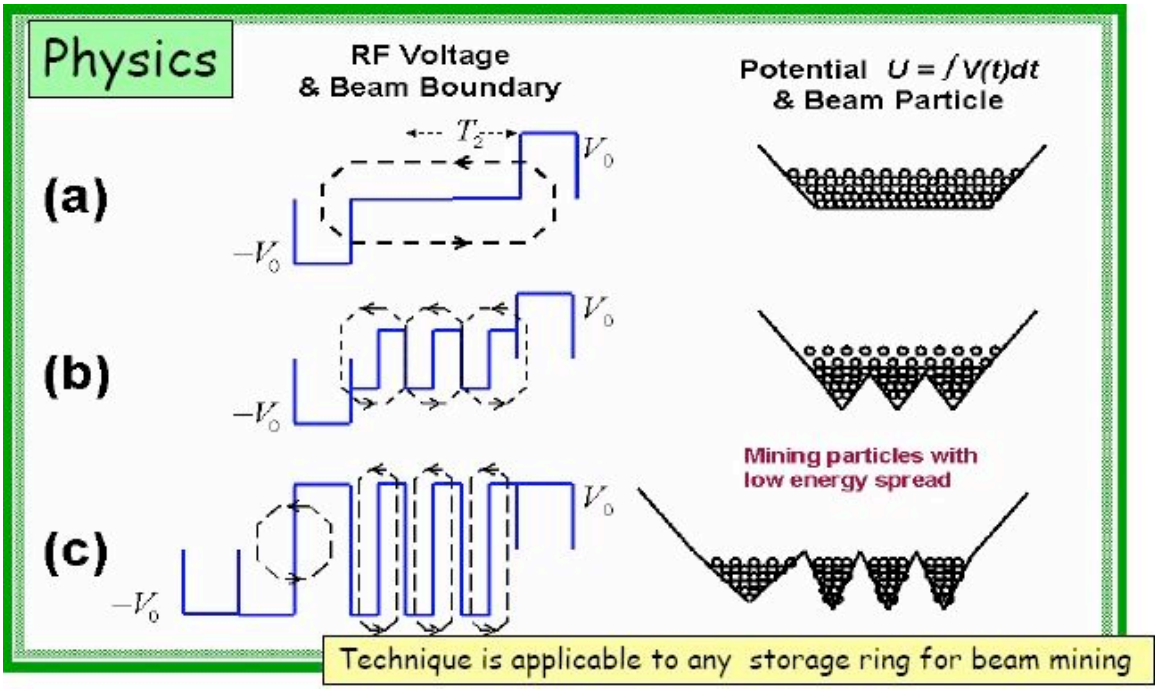

Figure 26: Momentum mining display. Beam starts in 1 bucket (a). Mining is done while the bunches are being formed (b). Then another bucket is formed to allow higher momentum particles to spill out of the bunches to be extracted (c).

Once these high momentum particles are swept away, 9 equally spaced bunches (called parcels) are formed. One at a time these 9 parcels are moved to the extraction region, spread out a little bit, and a 2.5MHz waveform is applied to the beam, giving 4 bunches per each parcel for a total of 36 bunches of antiprotons that are sent to the Tevatron.

LLRF Software

The user interface with the RR LLRF FE is done via an ACNET application page R6 (RR LLRF CONTROL). From here the waveforms can be edited and sent to the ARB generator. There are many waveforms that can be sent. Each of the functions for beam manipulation is given a State, or Cycle as it is referred to in the program. Within each state, there are typically several different files. The active state is the one currently in loaded for that cycle. There are also studies and operational states that can be used as well. Operational files are typically used for a defined purpose and should not be edited without a thorough understanding of what its effects will be. Studies files can be edited to suit a certain request or study.

Once read each file can be modified and then sent to the LLRF FE. This updates the current operational file and has it ready for the next time its particular RF state is issued and activated. This can be done from a parameter page, though is typically handled through the RR sequencer. The state (cycle) must be activated and then a $EC will actually tell the FE to issue the active state. This is what is considered episodic mode for the Recycler RF.

Episodic mode occurs when G:TLGRR≠0, as was described earlier. Any value other than 0 will cause the LLRF FE to not listen to beam transfer events coming from the timeline, and instead will be triggered by the $EC event. In this state V:TLGRR (RR RF state device that reflects the current state of RR RF) will reflect the value of G:TLGRR. These state values correspond to the RF state loaded by R6. This is helpful because many RF manipulations need to be done that do not involve transfers. In this state the sequencer alerts the LLRF of impending transfers and triggers the necessary RF states.

When G:TLGRR=0, the Recycler RF is in periodic mode. In periodic mode, the LLRF gets its commands from the Time Line Generator (TLG). In this mode V:TLGRR will not necessarily reflect G:TLGRR as V:TLGRR will be set by the timeline while G:TLGRR will remain zero. This happens most frequently during transfers into and out of the machine. In this mode, RF manipulations are triggered by beam transfer events ($E0, $E2, $E3, $E4).

High-Level RF

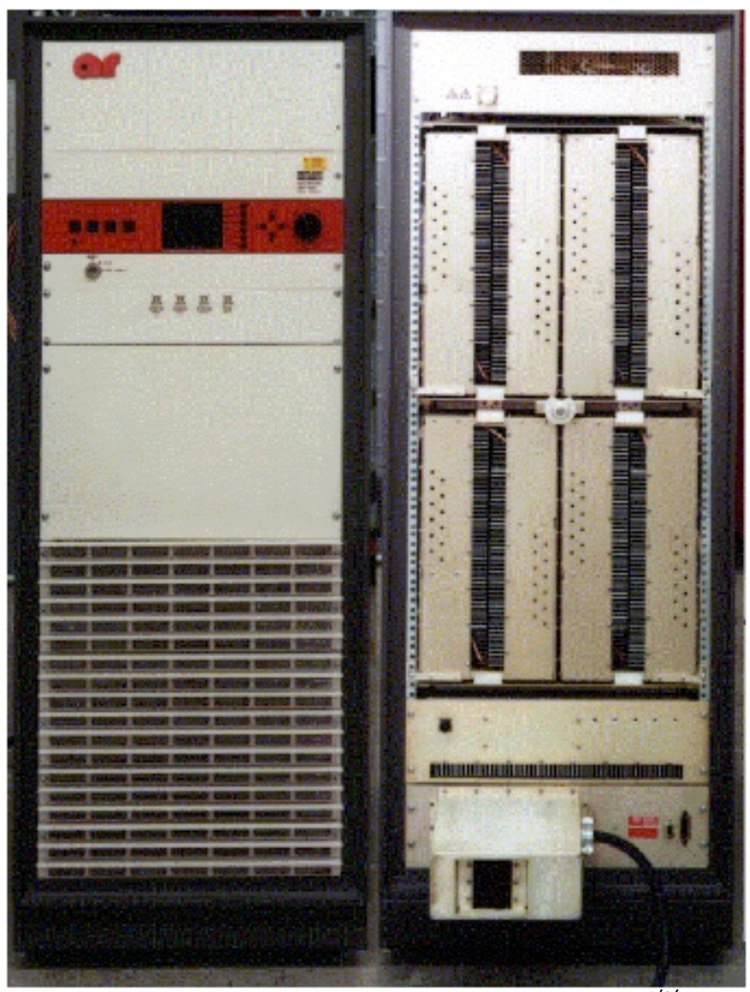

Now that the LLRF has summed its 8 inputs into one signal, it needs to be amplified and applied to the beam. The signal is amplified with 4 broadband solid-state amplifiers. These supplies are capable of 3500W output in a range of frequencies from 10 kHz to 100 MHz. The Recycler doesn't use the RF to accelerate the beam to higher energies, so it is not locked to a synchronous frequency.

These amplifiers are found in the north side of the MI-60 service building. They were originally air cooled, but have been switched to water-cooling which taps off the MI RF 95 water system. Figure 4 shows a picture of the front and back of the broadband amplifiers. In the back you can see 4 identical RF modules. You can also see from the picture that they were still air cooled when the picture was taken.

Figure 27: RR RF High-Level Amplifiers

Cavities

Once amplified, the signal is sent to the 4 cavities via a 7/8th inch coaxial line. Figure 28 is a cross-section of the RF cavities. It has a water cooled outer shell made of aluminum, with an aluminum inner conductor, and a 4 inch stainless steel beampipe that has a 1 inch ceramic gap connected electrically with a beryllium-copper finger stock. A thin copper end plate completes the electrical connection between the beampipe and the outer cavity shell. The beampipe is thermally isolated from the cavity. It is wrapped with heating tapes to allow a bake-out of the vacuum system.

The yellow objects are Mn-Zn ferrite cores (11.5" OD x 6" ID x 1" thick), with the three smaller ones near the accelerating gap being Ni-Zn ferrite cores (10" OD x 6" ID x 1" thick). These ferrite blocks are air cooled, spaced ½ inch apart and supported by Kapton spacer blocks.

Connected directly across the gap is a 10 kW RF load modified with a 60 ohm resistor. This RF load is water-cooled. The resistor in parallel with the stack of ferrites creates an impedance of 50 ohms over the 100 kHz to 20 MHz frequency range. Both the input drive and the 60-ohm resistor are connected to the inner conductor at the gap with 1" wide x 4" long flat copper strips. The voltage drop across the gap is monitored by two resistive voltage dividers.

The water-cooling from the cavities in the tunnel is fed directly from the MI 95 Cavity LCW system in the tunnel.

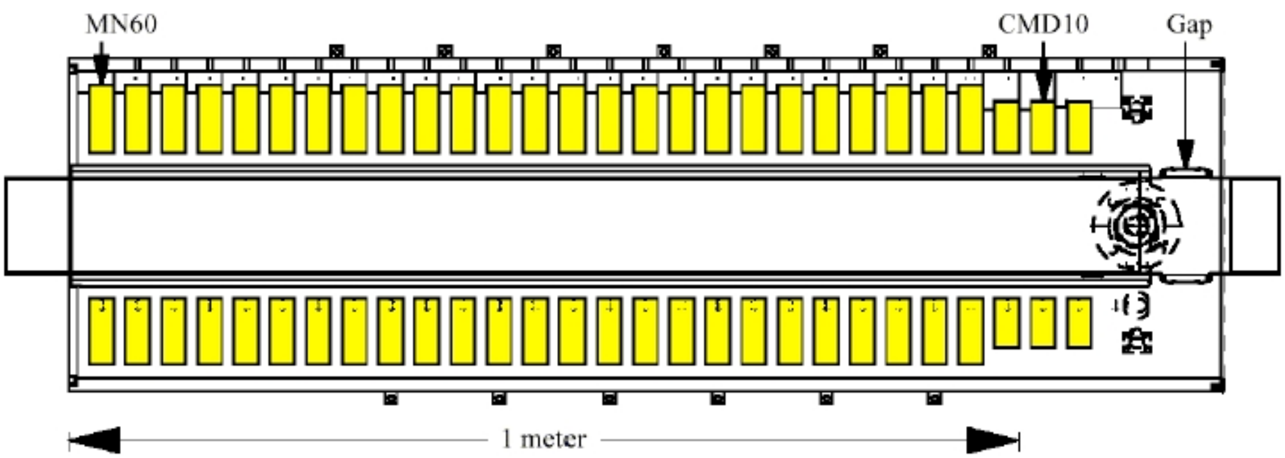

Figure 28: Cross section of RR Wideband RF cavity

Cooling

Stochastic Cooling

When it comes to storage rings, lifetime of the particles being stored is prized above all else. Superior vacuum is important, as you don't want your important antiprotons colliding with random stray molecules. Good vacuum will only get you so far as you need a way to keep the particles you want to keep from interfering with each other. There are several different ways to do this. The Recycler uses electron cooling and stochastic cooling.

Electron cooling is a method of cooling the beam by Coulomb interaction of the antiprotons and electrons. "Hot" antiprotons will exchange momentum with "cold" electrons, thus cooling the antiprotons. This will be explained in greater detail later in the book in the Ecool section, so for now we will focus on the stochastic cooling method.

Stochastic cooling involves measuring a random sample of beam and applying a small kick to correct for errors in its position or momentum. The kick must be applied to the same sample of beam that was measured; otherwise the beam will be heated instead of cooled. In order for this to work, placement of the kicker relative to the pickup for each system must be precise. As pickups measure position errors and kickers give beam an angular kick, their placement for the transverse systems must be at least ¼ betatron oscillation apart, or 90 degrees of betatron phase advance. See the Pbar Rookie Book for a detailed explanation of the stochastic cooling process.

The Recycler phase advance per cell is roughly 85 (79) degrees horizontally (vertically), which adds up over the 35 or so half cells between the pickups and the kickers. With the pickups placed between 211 and 213, the kickers placed between 102 and 103, we get a net phase advance near 90 degrees, which is close enough to allow for optimum cooling. For those keeping track, those numbers for phase advance relate to roughly 4.25 (3.75) betatron oscillations.

There are 4 stochastic cooling systems used in the Recycler. The three transverse systems operate in the 1-2 Ghz (horizontal) and 2-4 GHz (both horizontal and vertical) range. The longitudinal system operates in the 0.5 GHz-1 GHz range. The frequency ranges for the longitudinal system was chosen to accommodate the +/- 20 MeV/c momentum spread that is maintained in the Recycler. The transverse bandwidths are also chosen in part due to this +/- 20 MeV/c momentum spread. The intrabeam scattering that drives the longitudinal spread will have little effect on the transverse emittances, and the transverse emittances are likely to be driven by other mechanisms. Nonetheless, the 2-4 GHz bandwidth is enough to keep transverse emittances below 10 pi-mmmr.

Figure 29 shows the relative locations of the components of the stochastic cooling systems. Each system begins with a Schottky pickup in the 212 region and the signal is then sent into the MI-21 peanut. In the tunnel the pickups are oriented such that the beam encounters the horizontal pickups (longitudinal pickups before transverse) and then the vertical pickups (again longitudinal before transverse). Tank 1 for both planes is used for the longitudinal system, and the other tanks are used for the appropriate transverse system. There are many amplifiers, hybrids and other various pieces of equipment that can be tuned by system experts, but for the sake of simplicity, we will just consider them part of the amplification of the signal, and leave it at that. If you would like to know the "nuts and bolts" of the system, please refer to the particular system drawings.

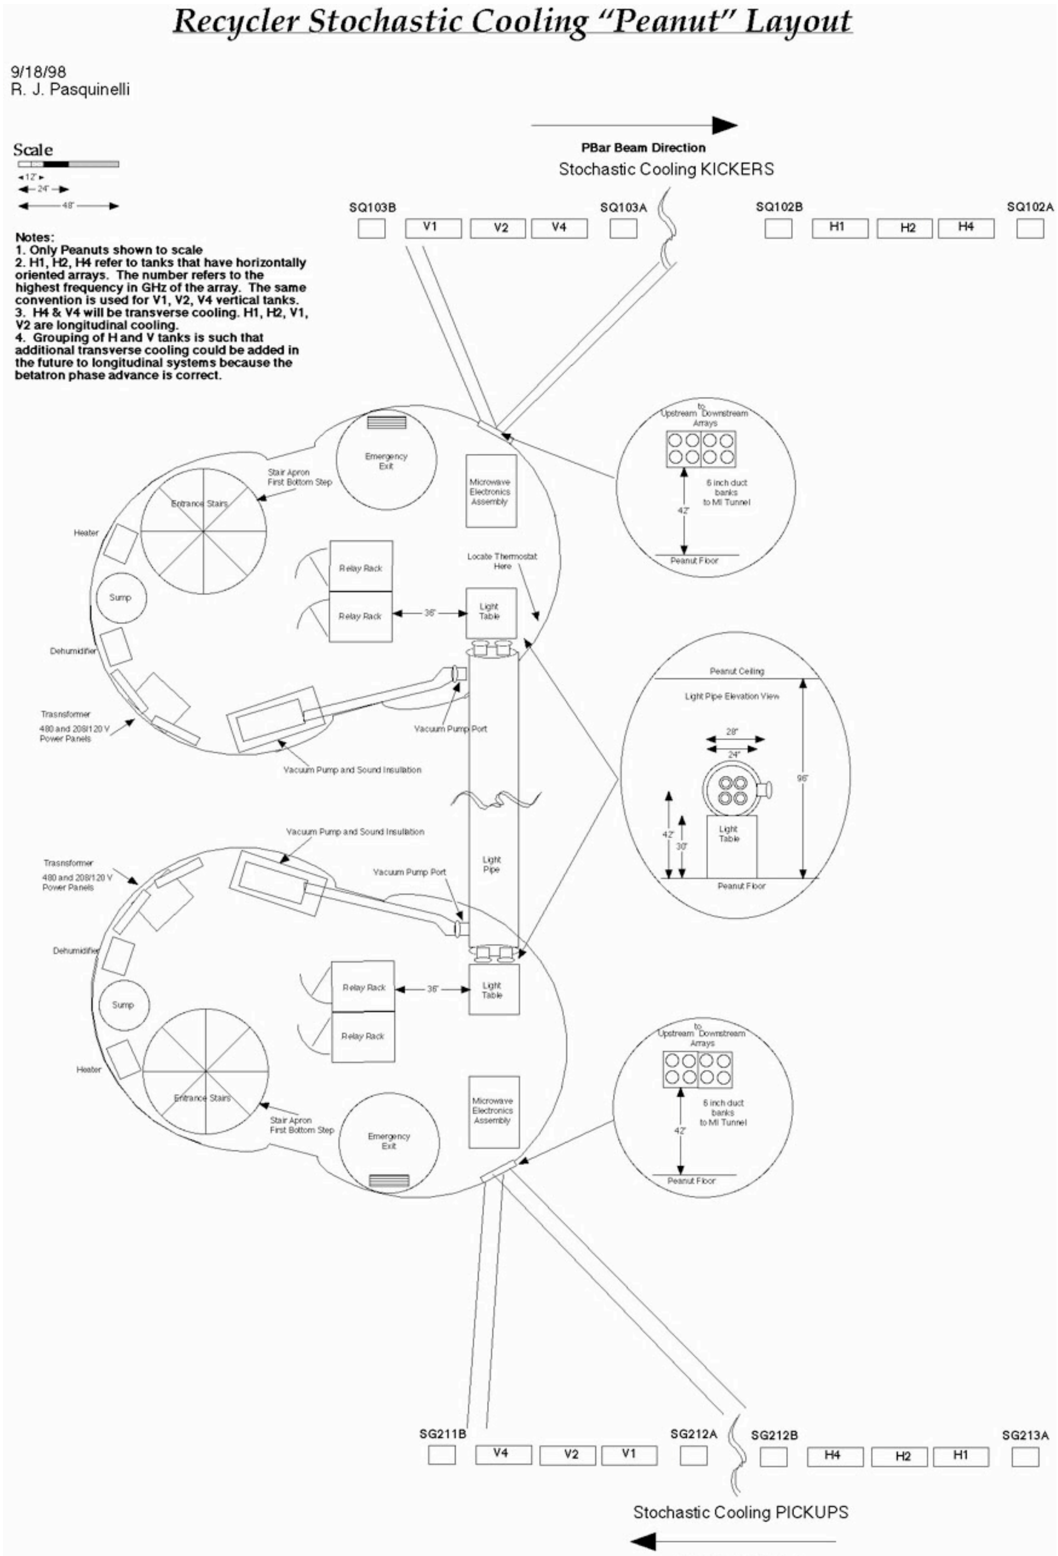

Figure 29: Stochastic cooling peanut layout.

0.5 – 1 GHz Longitudinal System

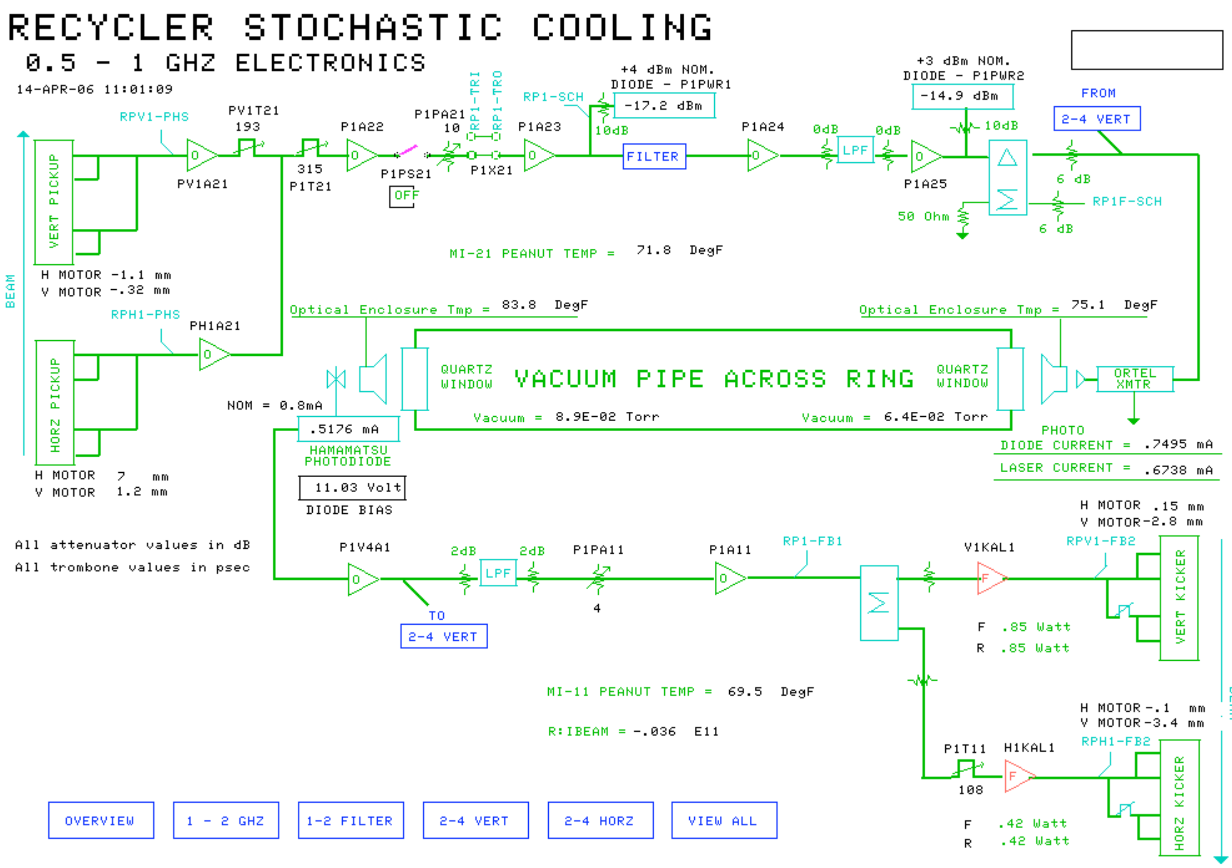

Figure 30: 0.5 – 1 GHz Cooling System.

This system has two pickups, one horizontal and one vertical. Each pickup has a plate on either side of the beam. For the longitudinal cooling these two plate's signals are summed. The signals are then combined to give a better signal-to-noise ratio. This could be done with one pickup, either horizontal or vertical, but having both assures us of a better signal.

This signal is measured shortly after the combiner to check for signal quality. There is a spectrum analyzer operators can use to measure the "front end" cooling, and an amplifier used to nominalize the S/N of the system without allowing for distortion of the signal. The proper signal is selected with a switch-tree that allows one to choose from the four available system signals.

Before the longitudinal signal is sent across the laser link, it is passed through an optical notch filter. As the central peak in a Schottky signal is the on-momentum particles, it would not be wise, or efficient to try to cool these. This is where the notch filter is used. It uses a 3 GHz signal split onto two paths, one that is kept 1 whole turn (of beam around the Recycler) longer than the other with a trombone. This is then mated with another signal that is used to keep the phase difference between those two signals such that it gives zero voltage. This notch is applied after the signal to noise is set, and before the laser link transmits the signal to MI-11.

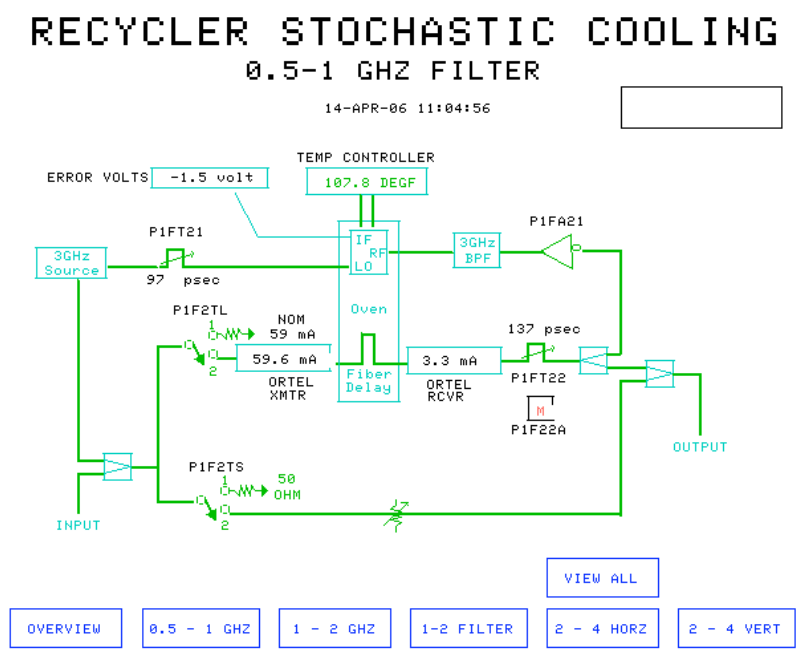

Figure 31: Optical notch filter diagram.

Another signal is added to the 0.5-1 GHz longitudinal signal before transmitted via the laser link. This signal is the 2-4 GHz vertical system. These signals are simply summed together. As the frequency bands do not overlap, it is easy to separate them after transmission, and in combining them; the cost of one laser link is deducted.

Now the signal is ready to be sent to the kicker side of the cooling system. This is done with an infrared laser (1440 nm) through an evacuated chamber to the MI-11 peanut. The chamber is evacuated to keep thermal effects and obvious interfering particles out of the way of the laser. The pump for this is controlled by a PLC from the MI-11 peanut that keeps the vacuum near 0.1 Torr.

In the MI-11 peanut, the signals are received with a photodiode that requires alignment from time to time by operators. This is done by flipping a target into the laser beam and centering the beam on the target.

After the beam is centered, the signal is split in two, one for this system, and one for the 2-4 GHz vertical system. Now the signal is sent through a low pass filter to remove the 2-4 GHz signal and we are left with our original signal. Once we have only the 0.5-1 GHZ signal again we need to check the signal suppression of this system. We use another signal analyzer to compare the beam signal with and without the cooling system on. Again, a switch-tree is used to select the appropriate system. When the system-off signal is measured, the final amplifier before the switch-tree is turned off. The antiproton system opens pin switches to remove power from its cooling systems. Here the amplifier is turned off. Either way, no signal is being applied. Look at Figure 3 to see a typical signal suppression signal. One can see the longitudinal system has a dip in the middle of the signal. This is the effect of the optical notch filter described above.

Once the correct amplification of the signal is set, it is then split and sent to the kickers in the tunnel in the 102-103 section. Like the pickups, there is a horizontal and a vertical kicker. All this happens before the beam travels from 211-103.

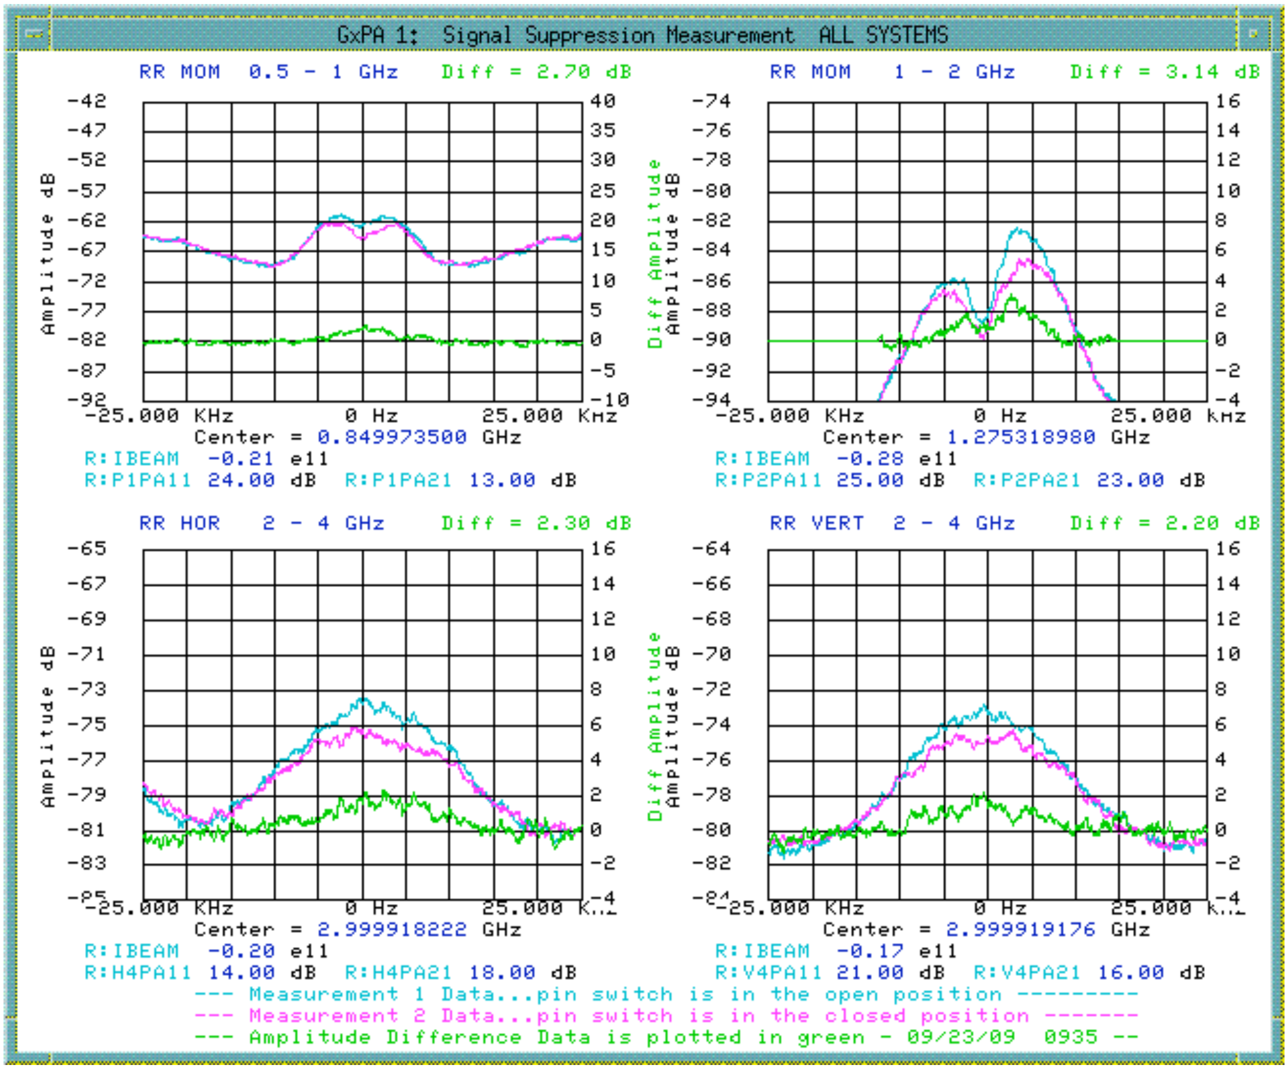

Figure 32: Typical signal suppression image generated by ACNET application R16.

1 – 2 GHz Longitudinal System

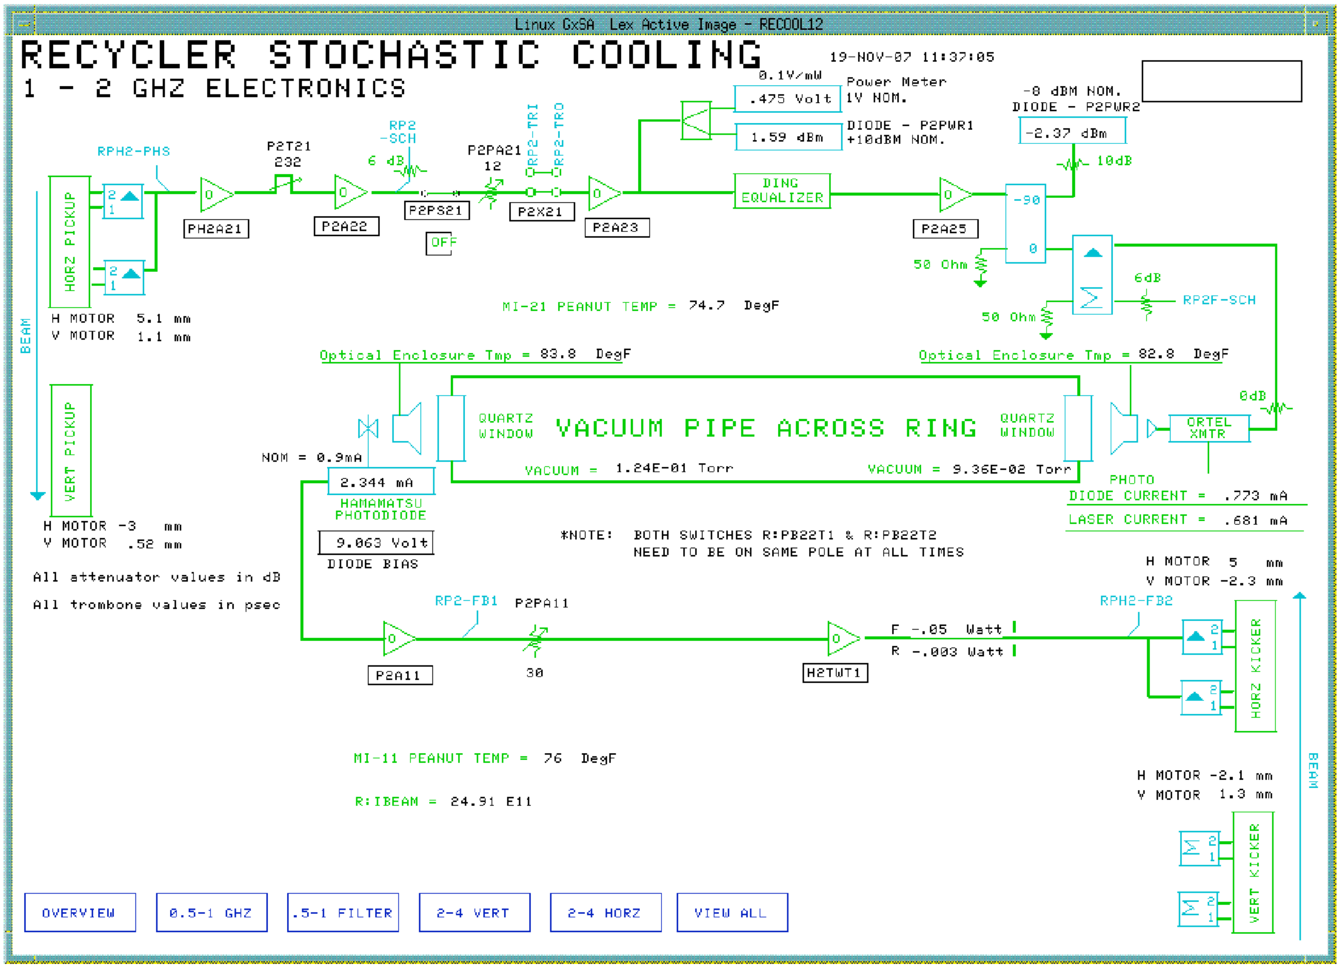

Figure 33: 1 – 2 GHz Horizontal System

This system is set up quite similar to the 0.5-1 GHz system. Its pickups are just downstream of the 0.5-1 GHz pickups. From here it follows a similar path through the “front-end” setting. As the range of this system is close to that of the transverse systems it is not summed with one of them before sent, and therefore has its own laser link. The signal then travels to the MI-11 peanut where it is centered on the pickup and amplified to provide the correct amplitude via signal suppression, then sent to the kickers. The kickers are placed just downstream of those for the 0.5-1 GHz system.

2 -4 GHz Horizontal System

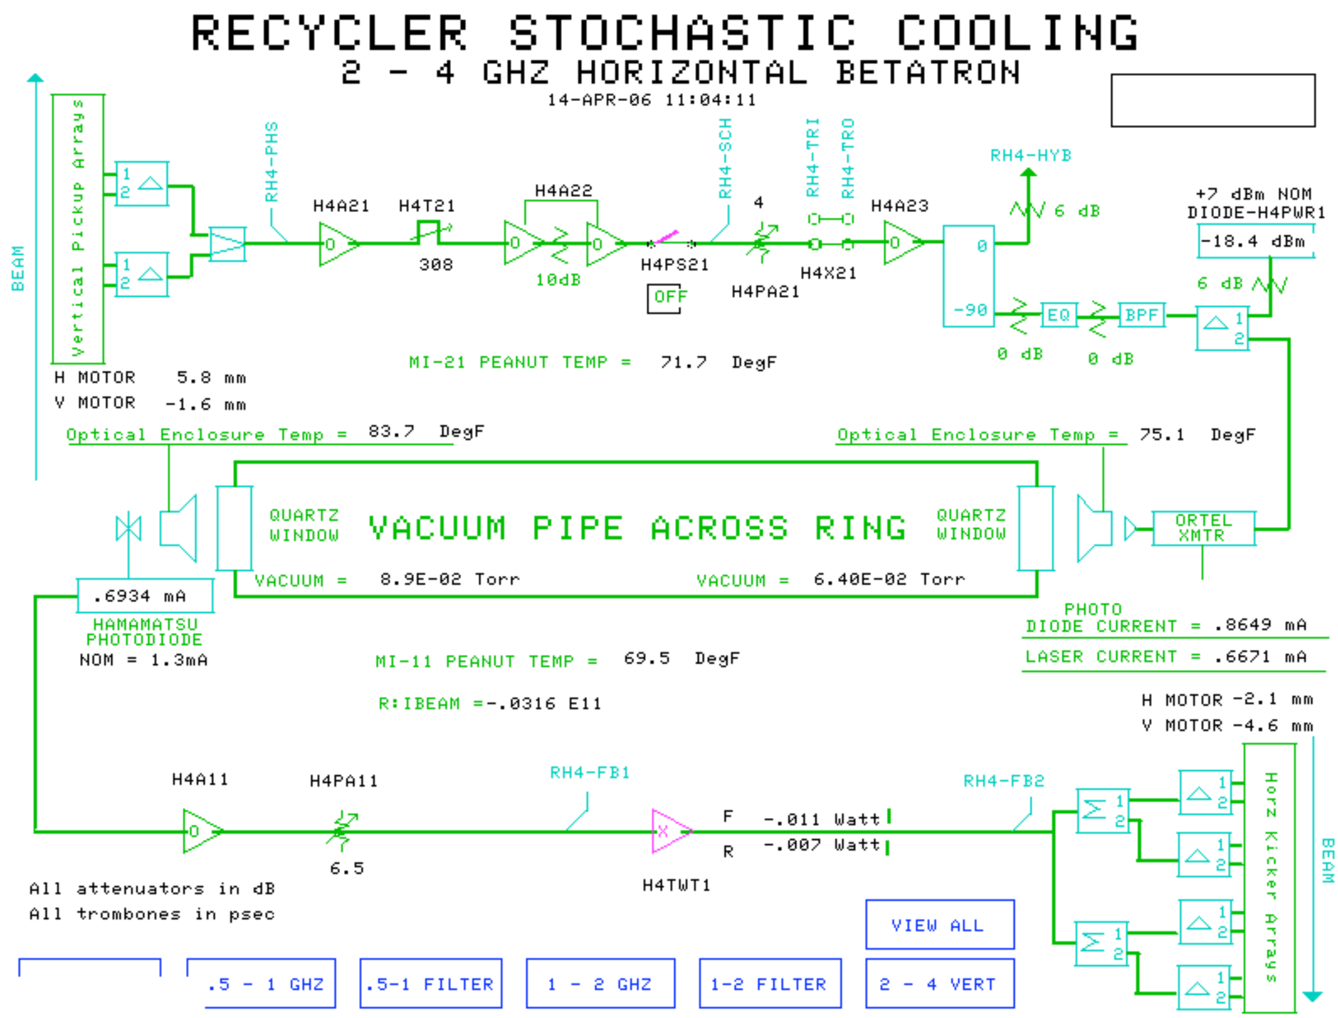

Figure 34: 2 –4 GHZ Horizontal System.

The transverse systems are a little simpler than the longitudinal system. There is just the one Schottky tank with 2 pickups, one on either side of the beam. The signals from the two plates are subtracted to obtain a horizontal offset from the two signals in this case. This difference is what is then sent through the amplification chain where the “front-end” is set and then through the laser link to MI-11. The photodiodes are aligned and send the signal on to set the signal suppression then down into the tunnel to the horizontal kicker downstream of the longitudinal system’s horizontal kickers.

2 – 4 GHz Vertical System

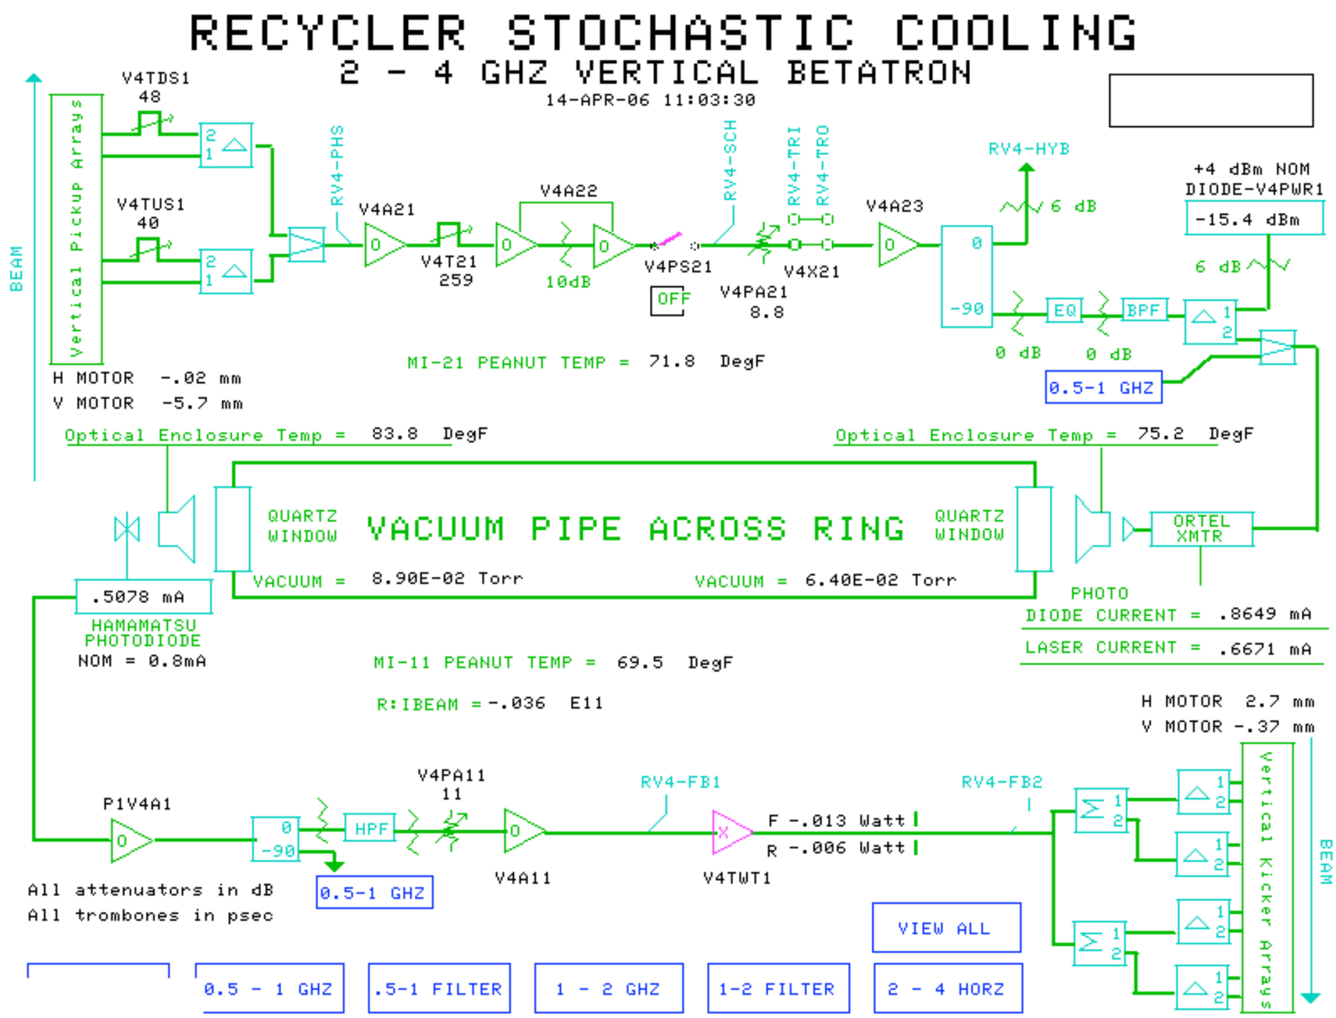

Figure 35: 2 – 4 GHz Vertical System

The vertical system's pickup tank is the last of the stochastic cooling pickups at 211 that beam encounters. Again the signals form the two plates (top and bottom) are subtracted and sent through the "front-end" setting. This signal is combined with the 0.5-1 GHz longitudinal signal and sent over the laser link. Once the photodiode picks it up, the signal is split (the other part going to the 0.5-1 GHz system), filtered to remove the 0.5-1 GHz signal, and then amplified after signal suppression is checked. From there it heads into the tunnel and to the kicker right before the 103 section.

Electron Cooling

Mary Sutherland has written articles on this subject that can be found at the following Recycler website: http://www-ecool.fnal.gov/ecool_rookie_book.html

Diagnostics

With the magnets to guide the beam and the RF in place to control the beam, all that is needed is a way to monitor the beam. The Recycler has several methods to allow us to see where the beam is, what it is doing, and how it is behaving. The first of these are the BPMs, which will tell us where the beam is in the aperture. Flying wires tell us how big the beam is transversely. Schottky detectors are used for stochastic cooling pickups, and resistive wall monitors give us a picture of the beam as a whole.

BPMs

The hardware consists of a pre-amplifier connected to a split-plate BPM, an analog differential receiver-filter module, and an 8-channel 80MHz digital down converter VME board. The system produces position and intensity with a dynamic range of 30 dB and a resolution of +/- 10 microns. The position measurements are made on 2.5 MHz bunched beam and barrier buckets of the un-bunched beam. The digital receiver system operates in one of six different signal processing modes that include 2.5MHz average, 2.5MHz bunch-by-bunch, 2.5MHz narrow band, un-bunched average, un-bunched head/tail, and 89kHz narrow band. Receiver data is acquired on any of up to sixteen clock events related to Recycler beam transfers and other machine activities. Data from the digital receiver board are transferred to the front-end CPU for position and intensity computation on an on-demand basis through the VME bus. Data buffers are maintained for each of the acquisition events and support flash, closed orbit and turn-by-turn measurements. A calibration system provides evaluation of the BPM signal path and application programs.

System Overview

1. Beam position split-plate capacitive pickups in the beamline vacuum.
2. Signal preamplifiers in the beam enclosure near the beamline.
3. Analog signal receiver (transition modules) and digital signal processing electronics in the service buildings.
4. Timing signal generators.
5. VME based “front-end” data acquisition and control computers and support hardware.
6. Integrated position and intensity calibration system.

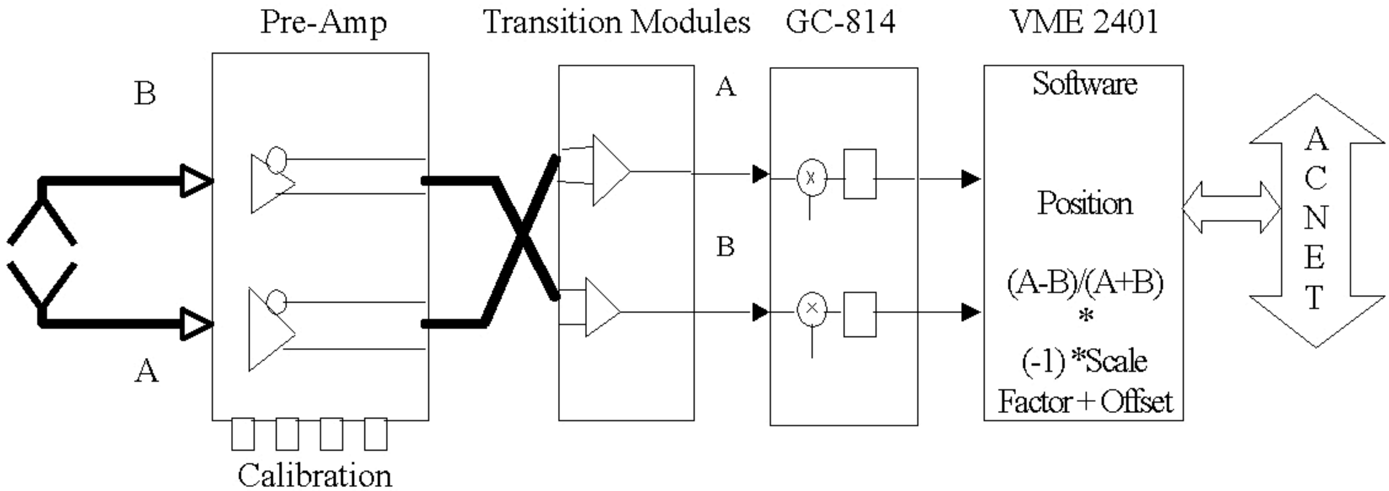

Figure 36: Recycler BPM system block diagram.

There are 211 BPM locations (104 horizontal, 107 vertical) served by this system in the Recycler Ring and 26 locations in the associated beam lines.

Beam position is determined from the relative amplitude from opposing electrodes in a BPM pick-up. In this design, the individual A and B signals follow parallel, but independent, signal processing paths from the pick-up through digitization and hardware digital processing. Signal differencing and normalization to obtain beam position information is done in software in the VME CPU. The accuracy of the BPM system is limited by how well the A and B channel pairs are matched, calibrated, and controlled in gain, offset, and frequency response. The BPM system is required to report the A+B signal sum as well as beam position.

The BPM signals are a function of pick-up design, beam intensity, the temporal structure of the beam current, and the beam position. BPM system design is thus influenced by each of these parameters. The Recycler functional specification identifies two beam structures (2.5 MHz, and un-bunched) to be measured, the range of beam intensities for each, and the required accuracy for each. This sets the required scale for determination of signal magnitude dynamic range, bandwidth, and A/B channel matching. This BPM system uses existing tunnel-to-service building twisted pair cabling with runs ranging from 150 ft to 1300 ft. Signal attenuation over this range of cable lengths varies by about a factor of six, impacting the overall system dynamic range requirements

Signal digitization and hardware digital processing is performed in EchoTek GC814 VME digital down-converter/digital receiver (DDC) modules. The EchoTek module requires different set-ups for processing signals from the different beam structures and for different measurement modes, e.g. single bunch, head vs. tail, etc.

The Timing Signal Generator interfaces to the Recycler Beam Sync Clock and provides programmable triggers to the EchoTek modules. This is an Industry Pack (IP) style module that resides on a carrier board on each VME CPU module.

The VME front-end computer receives digital A and B signal data from the EchoTek modules, performs final beam position and intensity (sum) signal processing, and controls the numerous aspects of timing and the EchoTek module. It also provides calibration system set-up, and interfaces to ACNET.

Beams in The Recycler

The BPM system is required to measure proton beams and antiproton beams. The Recycler has up to three partitions of beam at a time-cold “de-bunched” beam, warm “de-bunched” beam, and injected or to-be-extracted beam with 2.5 MHz structure.

The 2.5 MHz structure is made up of four bunches in successive 2.5 MHz RF buckets (396 ns spacing). Bunch lengths range from a Gaussian sigma of 25 ns to 50 ns. The intensity range for bunched beam is 2.0e10 to 30e10 particles in four bunches. The required absolute position relative to the center of the BPM is +/- 1.0 mm. The relative difference between two measurements on subsequent pulses with stable beam is +/- 0.40 mm.

A partition of de-bunched beam is contained between RF barriers created by a wideband RF system. The width of a typical barrier is 906 ns, it can be as small as 679 ns and as large as 1132 ns. The de-bunched beam has no intentional modulation structure over the length of the bunch. The dynamic range is 20e10 to 400e10 particles for de-bunched beam.

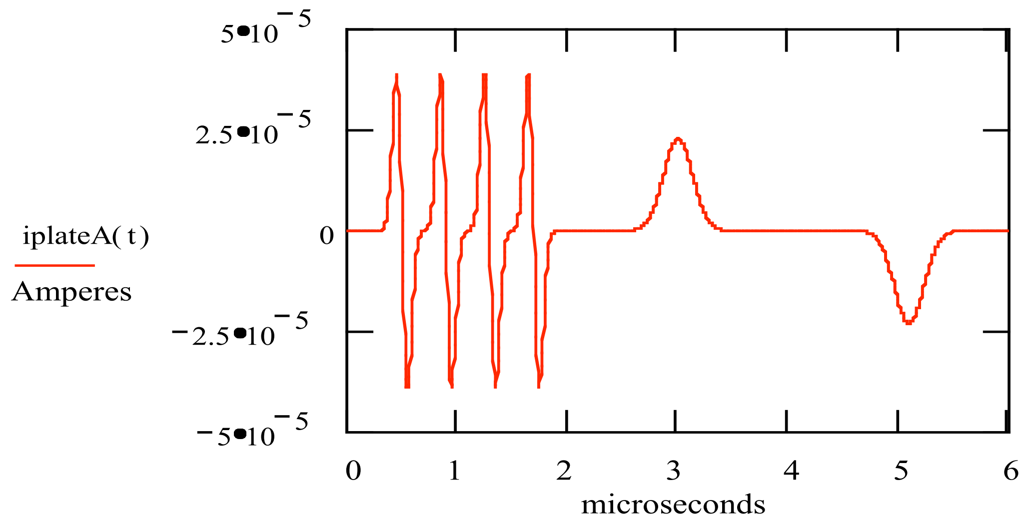

Electrode current for 2E10 total in four 2.5 MHz ó = 50 ns bunches and 20E10 in 2.094 ì sec bunch with 566 ns edges

Figure 37: RF Barriers

Hardware Signal Processing Modes

The Recycler BPM system allows selection of one of six signal-processing modes in the EchoTek digital receiver boards. The modes are:

1. 2.5 MHz Ensemble.

   This mode results in measurement of the magnitude of the 2.5 MHz signal integrated over an interval corresponding to about 1.6 microseconds (four 2.5 MHz periods), the full extent of the 4-bunch beam structure.

2. MHz Bunch-by-Bunch.
   This mode results in measurement of the magnitude of the 2.5 MHz signal integrated over an interval corresponding to about 0.4 microseconds (one 2.5 MHz period). The measurement duration is about 1.6 microseconds (four 2.5 MHz periods) allowing measurement of each of the four 2.5 MHz bunches.
3. 2.5 MHz Narrow Band.
   This mode results in measurement of the magnitude of the 2.5 MHz signal integrated over an interval corresponding to about 700 microseconds (about 64 TBD Recycler turns).
   This mode may be used to measure the closed orbit of all beam in the Recycler Ring that generates 2.5 MHz signal frequency components.
4. Un-bunched Ensemble.
   This mode results in measurement of the magnitude of the signal from beam confined in a barrier bucket otherwise lacking intentional modulation. The signal is effectively integrated over an interval corresponding to the full barrier bucket length. The actual EchoTek setup is determined algorithmically based on the barrier bucket length as determined from MDAT, the accelerator data distribution system. This mode is used to measure the "average" position of the barrier bucket structure beam.
5. Un-bunched Head or Tail.
   This mode results in measurement of the magnitude of the signal from beam in the leading or trailing edges of a barrier bucket. The signal is effectively integrated over an interval corresponding to the typical rising or falling time of the barrier bucket beam structure. This mode measures the position of the head and/or the tail of the barrier bucket structure beam.
6. 89 kHz Narrow Band.
   This mode results in measurement of the magnitude of the 89 kHz signal integrated over an interval corresponding to about 700 microseconds (about 64 Recycler turns). This mode is used to measure the closed orbit of all beam in the Recycler Ring that generates 89 kHz signal frequency components.

Preamp Design

Each preamp is a two-channel circuit required to:

a. Receive via coaxial cables and into well-controlled impedances two beam signals (A and B) from the respective BPM electrodes.
b. Provide suitable gain and impedance matching to drive the beam signals on twisted-pair cables to the service buildings.
c. Provide suitable connections and necessary DC isolation from the electronics to apply ion-clearing voltages up to 500 volts to the BPM electrodes.
d. Provide connections, control, and signal conditioning for BPM calibration signals for application to the normal beam signal paths at the preamp inputs.

e. Sustain without damage the peak input signal levels presented by specified beams (e.g. 53 MHz) outside the normal functional range for beam position measurement.

The preamp is designed to present a resistive load impedance of 1.5 kΩ to beam signals from the pick-up electrodes. This resistance coupled with the parallel electrode and coaxial cable capacitance provides a “high-pass” transfer function from beam current to preamp input voltage with a corner frequency of about 1.5 MHz. The preamplifier is designed for a mid-band voltage gain of about 3.5 with a bandwidth of approximately 23 MHz. The final amplifier stage provides a differential output to drive a 100Ω twisted-pair cable in a balanced fashion. Series resistors provide a back termination for the cable. The preamp is capable of driving a differential voltage up to 3.7 V peak-to-peak on the cable downstream of the back-terminating resistors.

RG58 cables with SMA connectors connect the preamp to the BPM electrodes. The calibration signal waveform is connected to and daisy chained through the preamp on RG58 cables with BNC connectors. A calibration control signal is connected to and daisy chained through the preamp on RG58 cables with BNC connectors. Ion clearing voltages connect to and daisy chain through the preamp on RG58 cables with SHV connectors. A cable with four twisted-pair wires connects the preamp in the tunnel and the Transition Module crate in a service building to carry the preamp A and B output signals from the tunnel and to provide DC power to the preamp.

Transition Module

The transition module input is configured using a differential receiver op-amp (AD8130) converting the differential input to a single ended signal. The main function of the transition module is to compensate for cable attenuation, provide an anti-aliasing filter, and normalize the input voltage to use the full dynamic range of the A/D converter.

Signal Digitization and Processing on EchoTek Board (DDC)

EchoTek GC814 cards have eight independent channels each structured as a set of A/D converter, digital down-converter and FIFO memory elements. Each channel digitizes and filters one of the A and B beam signals from the transition module at 80 MHz asynchronously.

Each digital down converter (DDC) device has four channels in parallel of a NCO, CIC filter and a two user-programmable LPF. Each channel can produce an “I and Q” pair as a 16-bit word and stored in the FIFO as one 32-bit long word. The data are transferred to the front-end processor where position and intensity are computed by processing the “I and Q” pair.

Our system uses two of the four DDC device channels, one channel generates the “I” term, and the second channel the “Q” term. This reduces decimation in the CIC filter, increases computation speed, and allows for the design of a fairly wide bandwidth filter. It also allowed us to meet the narrower bandwidth requirements where high decimation rates were needed in the case of turn averaging.

The DDC programmable filters were setup in asymmetric mode. The first filter stage (CFIR) was coded with an 11 tap Gaussian filter with the purpose of reducing data jitter introduced by asynchronous sampling. The second filter (PFIR) was setup as a boxcar averaging filter whose length was determined by the clock rate at the input of the 64 tap filter, the number of taps used, and the period of the data being sampled.

Data transfers between the EchoTek board and the front-end processor (MVME 2401) are done under DMA control (DMA32 or DMA64) or direct bus transfers depending on the type of measurement performed. Data transfer rate measurement provide a figure-of-merit determining the number of data points produced at the output of different types of filters as well as set the maximum number of transfers achievable in the most stringent case where data is obtained every turn (turn-by-turn) as opposed to single shot data acquisition mode.

Ion-Clearing Electrodes

Recycler BPM pickups serve a dual purpose. Not only do they report back position, but each is also used as ion clearing electrodes. It was decided to cover this here instead of in the vacuum section as they use the BPM electrodes, and are used for sensitive vacuum readbacks.

The Ion Clearing Field (ICF) system connects high voltage to the Recycler BPM electrodes in order to collect ions drifting in the beam pipe. One ICF power supply chassis exists in each of six Main Injector service buildings, and provides HV to all the BPMS in that respective sector. There is a circuit in the BPM preamps that isolates this HV signal from the BPM position measurements. The HV is distributed to the BPM electrodes in a daisy-chain fashion (both upstream and downstream) with separate paths for the A and B electrode cables. There are roughly 35 BPMs in each string upstream and downstream. This results in four cables being routed upstairs to a relay rack and the PS chassis that is positioned near the BPM electronics.

This PS chassis contains 2 programmable power supplies, a 4-channel current measuring circuit, and analog voltage and current monitoring circuit, and a logic board that provides status, and local/remote control. Each of the two HV power supplies drives two current measuring circuits. This provides 4 HV and associated current references. These are then cabled and run down to the BPM electrodes.

These HV lines are not interlocked to the safety system, as they have been deemed similar to the HV lines used for ion pumps.

The current measuring circuit mentioned above is used to measure the voltage drop across a large resistor in series with the load. The circuit is capable of measuring current in the pA range (yes, that's picoamps) While that is slightly offset by constraints from leakage current in electrodes and cables and MI ramp effects, it is indeed quite effective at measuring current draw in the nanoamp range that they currently provide. With proper scaling, this can be converted to a usable vacuum readback for quality of beam tube vacuum.

Flying Wires

Flying wires are similar to multiwires in that a wire is used to obtain a picture of the beam. However they are more akin to SEMs as they use loss monitors downstream of the wire to obtain the “picture” of the beam. What the flying wire does that the other two systems do not is take this “picture” while moving.

As the wire flies through the beam it creates collisions with the beam streaming past. These collisions create secondary particles that are detected by loss monitors. The electronics measure the losses and correlate them with the position of the wire. This gives a good profile of the beam from which one can determine transverse emittances. This also has a bad effect on the beam as one might imagine, it is after all creating losses. Wire flies are typically done on smaller stashes to minimize the impact.

The horizontal flying wire is located just above MI Q620, and the vertical wire is just above MI Q619.

Follow along with Figure 40 as it all begins with an ACNET console sending commands to a PC at MI-62. This PC talks with a VME crate that houses all the electronics for the flying wire system readbacks. Along with the commands from the PC, the crate also receives RRBS, MDAT and TCLK signals. This allows the crate to delay flies until a certain clock event has occurred.

The PC software tells the motor controller to start the wire moving. Figure 38 shows a cross section of the wire and beam interaction region.

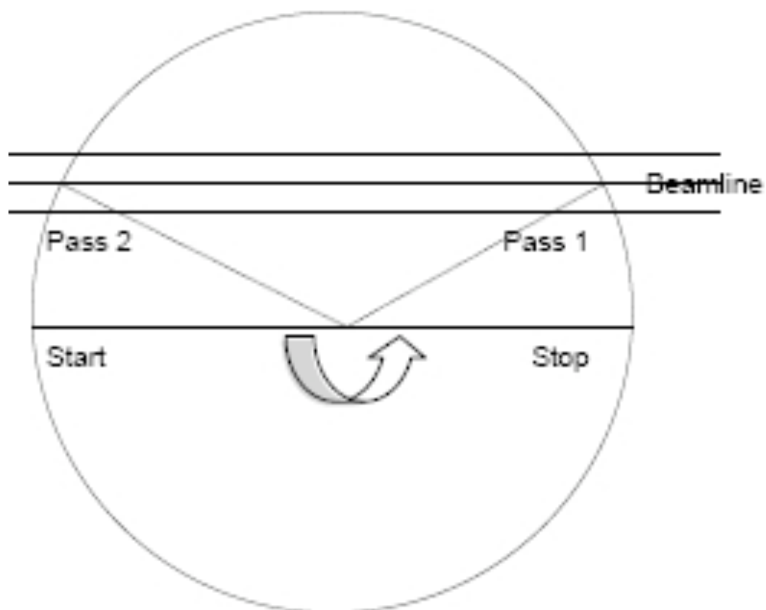

Figure 38: Flying wire interaction area.

Each fly of the wire involves 1.5 turns. The first ½ turn is to allow time for the wire to accelerate to its measurement speed. The 2nd half turn is where the measurement takes place as the wire pass through the beam twice. The last half turn is to allow the wire time to decelerate and stop and prepare for the next measurement. Each successive fly will travel in opposite directions. This allows for less time between measurements. Figure 39 shows a schematic of the wire, and its apparatus.

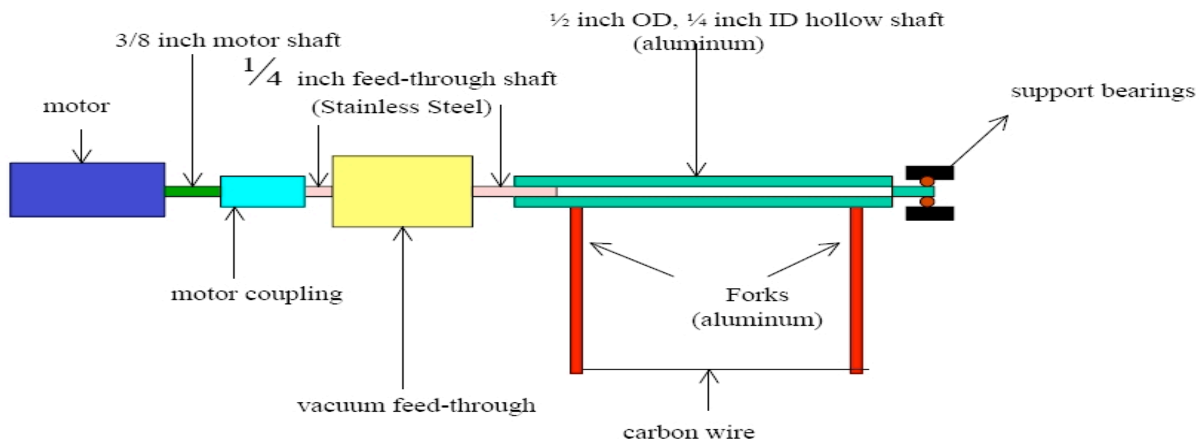

Figure 39: Flying wire schematic.

The motor controller reads back its position to the VME crate using a Resolver mounted on the motor shaft. This Resolver is capable of counting tics every 51mm, or 1.31 arc minutes.

Now that the wire has flown through the beam, loss monitors must catch the secondaries that come from this interaction. Scintillators readback through photomultiplier tubes that amplify the signal and readback to a digital/analog converter in the VME crate at MI-62. The crate uses this signal as well as the resolver signal and feeds this information to the PC software that fits it to a Gaussian curve and extracts emittances.

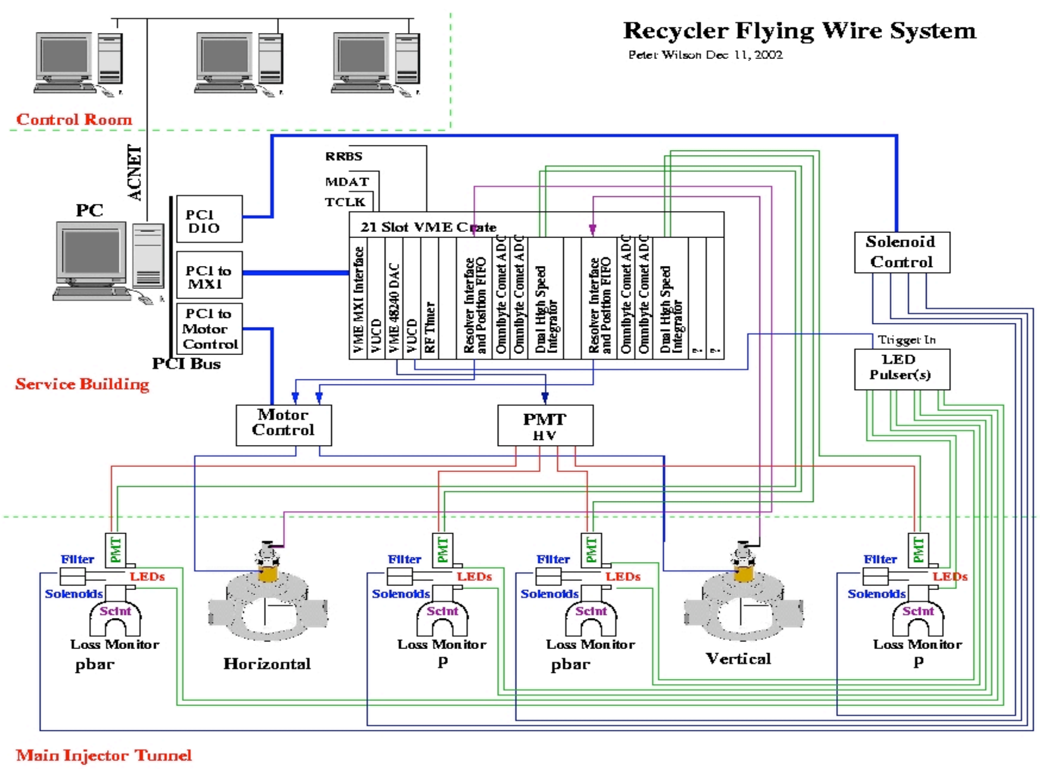

Figure 40: Schematic of flying wire system.

Schottky Detectors

Moving charged particles generates electrical current. The finite electrical charge of these particles gives rise to small statistical variations, or noise, in the electrical current. This Schottky effect was observed first by a German physicist named Walter Schottky in 1914. The information contained in this Schottky noise is quite useful as a non-destructive means of acquiring longitudinal and transverse properties of particle beams, and for stochastic cooling measurements.

Understanding this noise will give a better understanding of how the Schottky detector works, and how it can be used.
Imagine a beam composed of N particles randomly distributed around the ring, each with charge q, orbiting a circular accelerator with a revolution period of Tn. If you stop time at a predetermined point that a particle passes, the current created by these moving particles would be this sum of delta functions

$$I(t) = q\sum_{n=1}^{N} \sum_{k=-\infty}^{\infty} \delta(t - t_n - kT_n)$$

To envision this in the frequency spectrum one would do a Fourier transform of this current to yield another infinite train of delta functions

$$\tilde{I}(\Omega) = q\sum_{n=1}^{N} \omega_n \sum_{k=-\infty}^{\infty} \delta(\Omega - k\omega_n) e^{-i\Omega t_n}$$

This shows that the current's frequency spectrum has peaks at each harmonic of the revolution frequency wn = 2P/Tn. The power spectrum of I(t) can be determined as well, but first we will look at the distribution of the frequencies wn.
Assume the revolution frequencies vary around a mean value of w0. This variation is proportional to the deviation of the particle momenta p from their mean of po as such:

$$\frac{\delta\omega}{\omega_0} = \eta \frac{\delta p}{p_0}$$

where the frequency dispersion η is

$$\eta = \gamma^{-2} - \alpha_p$$

g is the relativistic Lorentz factor and ap is the momentum compaction factor of the ring lattice. Once h is known, it is possible to use the beam current frequency spectrum to compute the momentum width of the beam. From this you can derive the harmonic number kmax above which the Schottky harmonics start to overlap

$$k_{max} > \frac{1}{2}\left|\eta \frac{\delta p_{max}}{p_0}\right|^{-1}$$

If we are to try to get a longitudinal signal of this beam, lets assume that the current signal is independent of position, i.e. that S(x,y) ≈ S0. We can assume this as the sensitivity is constant over the range of betatron amplitudes at the detector. In this case the signal spectrum is directly proportional to the current spectrum. We will let our momentum distribution Ψ(δp/p0) be normalized to N. In the case where the Schottky bands do not overlap at frequencies of Ω(δp)=kω0(1+ηδp/p0), ZL is the line impedance of the detector, then the power spectrum is then:

$$P(\Omega) = \frac{(qf_0)^2}{4Z_L}\left|\tilde{S}_0(\Omega)\right|^2 \frac{\Psi(\delta p / p_0)}{|\eta k|}$$

This power spectrum has a few important properties:

- The total power of each harmonic is proportional to N. Once properly calibrated the total power is a good measure for beam current. This can be handy at lower intensities where current transformers may be subject to noise.
- It is proportional to the square of the charge. Heavy ions such as Gold are easier to see than protons.
- The width of the signal is proportional to the harmonic number k, and the power density is inversely proportional to the absolute value of both k and h.

Figure 41 shows a hypothetical longitudinal Schottky power spectrum. The spectrum is pretty broad to exaggerate the concepts covered above.

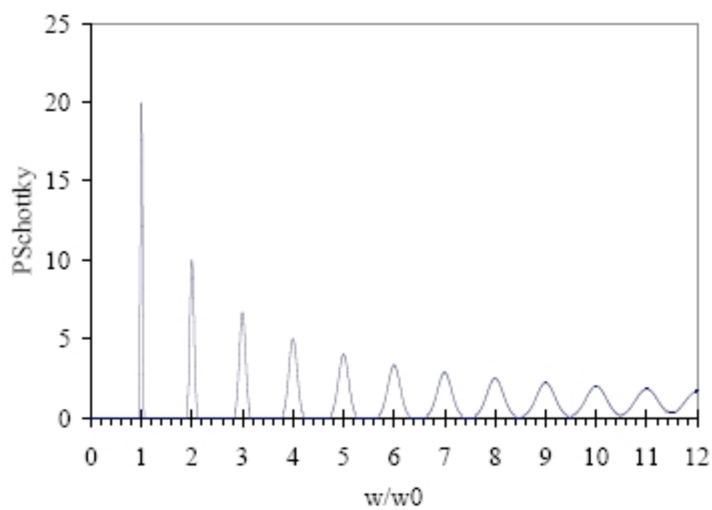

Figure 41: Hypothetical Schottky spectrum.

With one detector we can easily get a longitudinal spectrum, but what if we would like to get a picture of how the beam behaves in the transverse dimensions? Suppose we take two Schottky detectors and place them on opposite sides of the beam and take a difference signal, in much the same way that a BPM works. This will give us a way to look at the transverse power spectrum.

To do this, imagine a beam performing horizontal betatron oscillations around an electrical center x=0 of the detector.

$$x_{kn} = A_n \sin(2\pi Q_x k) + \mu_n$$

μ_η is the random betatron phase, and A_η is the betatron amplitude related to a single particle n with emittance $\varepsilon_{x,\eta}$ and horizontal beta function β_x at the detector where

$$A_n = \sqrt{\varepsilon_{x,n}\beta_x}$$

In the rhythm of the oscillations about the electric center, the detector output changes signs. This modulation produces sidebands around the revolution harmonics

$$\omega_{k,\pm} = (k \pm Q_x)\omega$$

Ignoring chromatic effects at high enough harmonics, the power spectrum assuming right and left bands are not overlapping is:

$$P(\omega_{k,\pm}) = \langle \varepsilon_x \rangle \beta_x \frac{(qf_0)^2}{16 Z_L} |S'(\Omega)|^2 \frac{\Psi(\delta p / p_0)}{|\eta k|}$$

where $<\varepsilon_x>$ is the mean emittance at the given momentum deviation. If you replace all the x subscripts with y subscripts then you have the vertical Schottky power. This means that from the transverse Schottky power spectra we can determine:

- The mean emittance, even if it is a function of the off-momentum particles, assuming the power spectral density is gauged well enough.
- The betatron tune
- The linear approximation of the chromaticity at low harmonics.

Remember this derivation was for a coasting DC beam. Bunching the beam will cause some extra signals due to the oscillations inherent to the bunch motion in the RF bucket. Nonetheless this should give you an idea as to how the Schottky noise comes about. There was a fair amount of math there, and several steps were skipped, so I will try to summarize with a more qualitative explanation.

As a charged particle passing through space creates an electrical current. Many charged particles will create a large current. In a synchrotron, these particles are not all traveling at exactly the same velocity, on the same orbit, or cross our detector at the same position. Due to these small variations in the orbits of these particles, there will be small variations in the current our detector sees from time to time. These small variations in the current observed are called Schottky noise. With a sensitive enough detector we can measure this noise. Each revolution this beam passes the Schottky detector it creates a band of frequencies where the detector "sees" the beam. This band of frequencies has sidebands that are detectable if the detector is sensitive over a broad enough span. Investigating these sidebands that show up can give us information as to the properties (both longitudinal and transverse) of the beam without ever interfering or disturbing the beam at all.

Now that we understand the process by which this diagnostic signal is created, we will discuss the physical detectors used in the Recycler.

There are several Schottky detectors in the Recycler. There are 6 in use for the stochastic cooling systems, 3 in each plane. There is a 30 MHz detector at MI-30 that is currently used for tune measurements. Above MI Q602 is a 79 MHz longitudinal Schottky detector. Figure 42 is of one of the 1.75 GHz detectors in the MI-62 section that is used to compute longitudinal and transverse emittances, among other things. The frequencies of the detectors were chosen so that we could get the desired information without serious signal interference, such as the coherent signal produced by bunched oscillating beam.

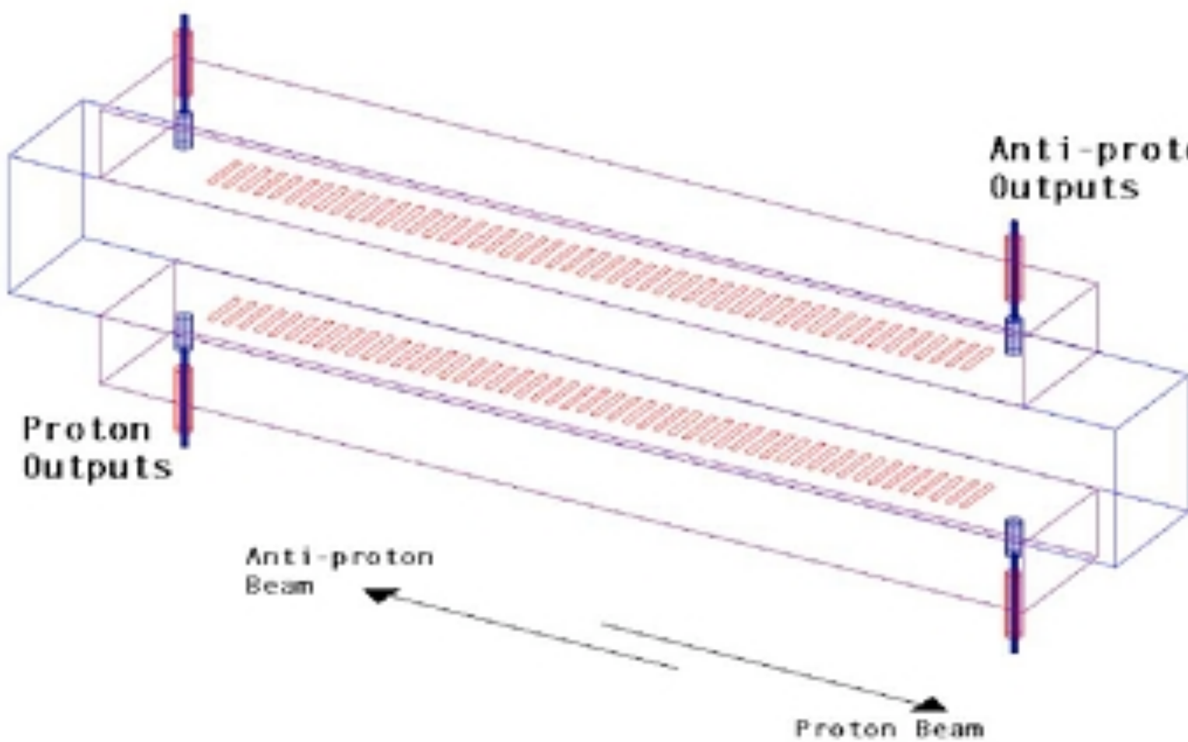

Figure 42: Diagram of Schottky pickup tank.

Resistive Wall Monitor

A resistive wall current monitor, or RWM, measures the image charge the flows along the beam pipe following the beam. This image charge has an equal magnitude, but opposite sign of the beam. In order to measure this image current, the beam pipe is cut and a resistive gap is inserted. Ferrite cores are used to force the current through the resistive gap, instead of other possible conducting paths.

Along with the image current, there are other currents that can flow along the beam pipe. These are negated by placing the resistive gap and ferrite cores inside a metal can to shunt these extraneous currents around the gap, rather than through it.

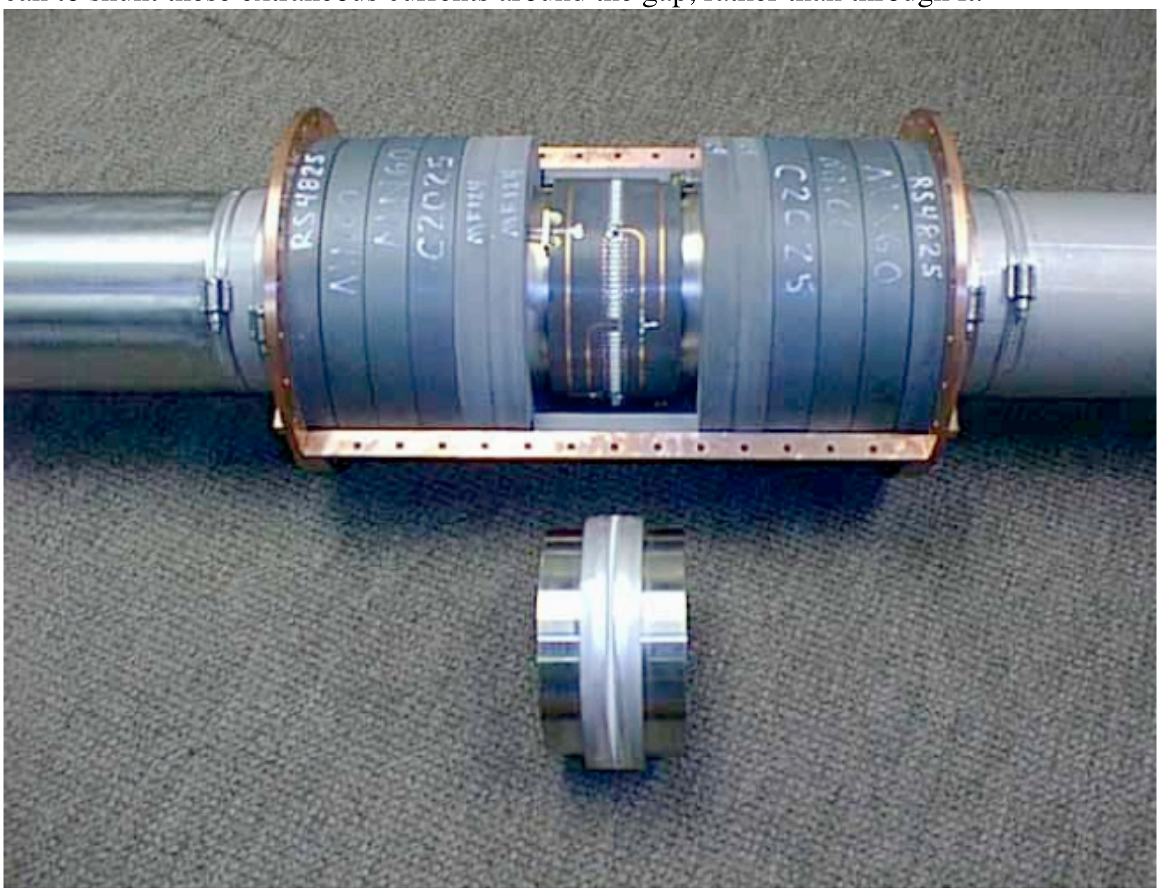

Figure 43: Picture of Resistive Wall Monitor.

The inductance of the ferrite cores and the resistance of the gap form a high pass filter with a corner frequency of R/2pL, which is usually a few kilohertz. Anything above this frequency the cores induce a current through the resistive gap that just cancels the beam current, thus minimizing the net current through the cores center.

The impedance of the gap is chosen to be much less than that of the ferrite cores. Several types of cores and microwave absorbers are used to maximize impedance and minimize resonances within the chosen bandwidth. Any time beam passes a discontinuity in the beam pipe electromagnetic energy is launched into the beam pipe. This energy can travel both against and with the beam, but is generally slower than the beam. A RWM cannot distinguish between this energy and the beam current, so microwave absorbers are needed on each side of the resistive gap.

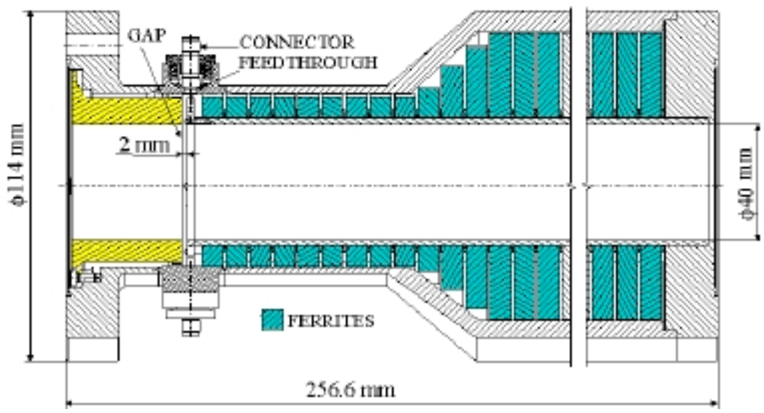

Figure 44: General RWM cross-section picture.

The beam current traveling across this resistive gap gives a voltage drop across the resistor. Measuring this voltage drop as the beam passes allows for a very precise measurement of the beam. A RWM allows one to look at the beam in a bunch-by-bunch basis. The Jell-O display's bottom trace of beam is created using a RWM.

The Recycler actually has two RWMs, both in MI-60. The 1-ohm RWM is used for the Jell-O display as seen in Figure 45, and a separate 25-ohm RWM is used for studies and RF system diagnostics.

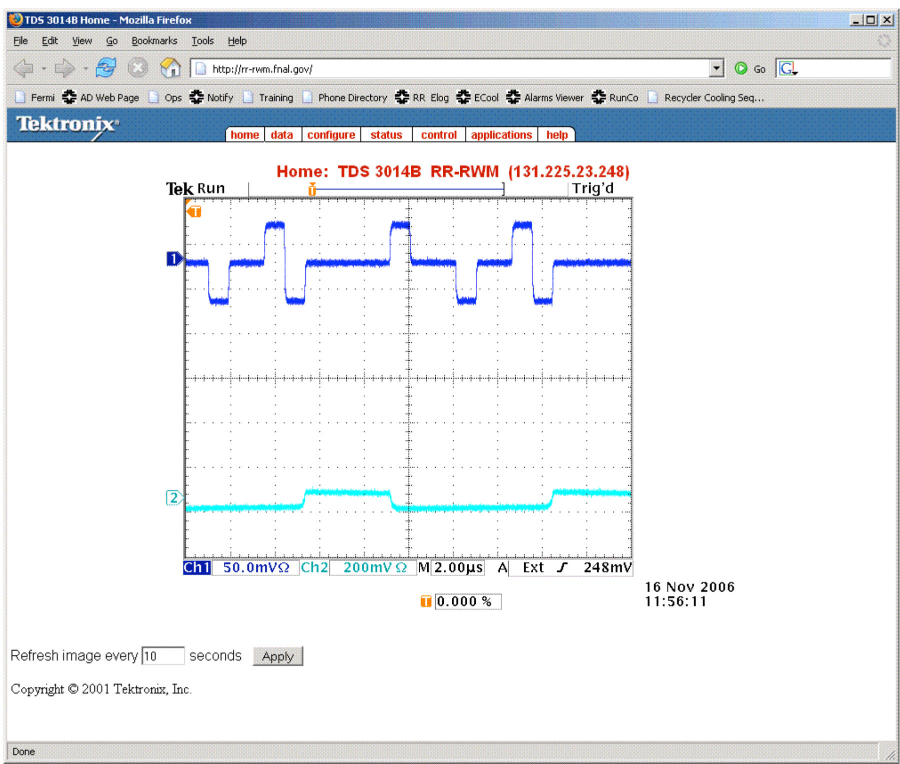

Figure 45: RWM output as viewed on scope. The top trace is the RF system output and the bottom trace is the RWM output, commonly called the Jell-O Display.

Utilities

Vacuum

The beam pipe in the Recycler, as with any accelerator, must be kept under vacuum. Any molecules in the pipe act as obstacles to circulating beam. As individual protons or antiprotons strike the errant air molecules in the beam pipe they are deflected. If the deflection is big enough, they will be absorbed by the beam pipe or otherwise leave the machine. If the deflection is smaller, the deflected particle may stay in the machine but its orbit will be different due to the transfer of energy from the deflection. In general, better vacuum will result in fewer of these deflections, which means that more of the beam will remain in the machine for a longer amount of time.

The beam in Recycler must circulate for hours, days, or weeks. This requires extremely good vacuum. The Recycler maintains a vacuum on the order of 2 X 10-10 torr. To gain a feeling of exactly how good of vacuum this is, we can compare it to more conventional units of measurement. Atmospheric pressure, for example, is traditionally measured in inches or millimeters of mercury as measured by a barometer. One torr is equivalent to one mmHg (one millimeter of mercury). For comparison, the pressure at sea level, one atmosphere (1 atm), is 760 mmHg or 760 torr. Another unit of pressure is the micron. One torr is equivalent to 1,000 microns.

Vacuum Pumps

Different types of vacuum pumps are effective at different pressures. When starting from atmospheric pressure, the beam pipe in the Recycler is pumped down in three stages, using roughing and turbo pumps, ion pumps and Titanium Sublimation pumps (TSP's) in that order. Of these, only the ion pumps and TSP's are permanently affixed to the ring, and are responsible for maintaining the ultra-high vacuum in the pipe. The roughing pumps and turbo pumps, combined on portable carts as a single unit, are moved to locations where the vacuum is especially poor, such as during initial pumpdown.

Roughing pumps and Turbo Pumps

Roughing pumps use a mechanical piston to remove air from the beam tube. It is the only type of pump effective at atmospheric pressure. It is capable of pumping down to about 10 torr.

Turbo pumps are basically multi-tiered "fans" which drive the air molecules out of the pipe, as shown in Figure 46. They are turned on when the roughing pump has done what it can to improve vacuum. The blades spin at a rate of several tens of thousands of RPM; the ends of the blades actually move faster than the molecules they are trying to hit, forcing them out of the beam pipe. The turbo will pump the beam pipe down to approximately 10-5 torr.

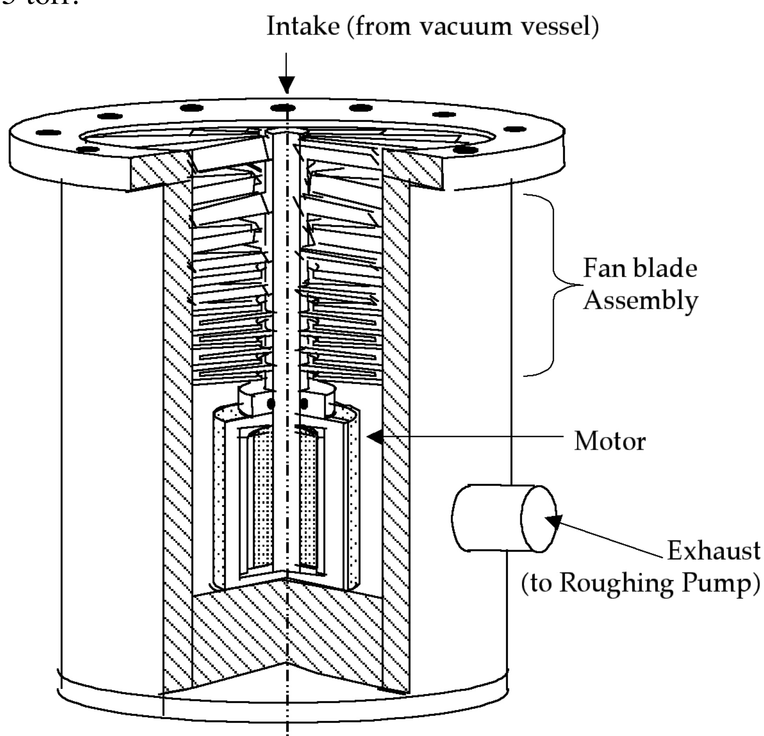

Figure: 46 – A turbo pump assembly

Pumping ions

Any pumping scheme that expels air outside the system through mechanical processes, such as the roughing and turbo pumps, will have some problems with "back streaming," or the tendency for air to slip back in from the outside through the pumping apparatus. Ion pumps avoid this problem because they do not expel gasses out side of the system. Instead, gas molecules are literally buried in the pump itself.

Inside an ion pump, 5200V is placed across a stainless steel or titanium anode and one or two titanium cathodes depending on the design. Free electrons in the electric field will be accelerated and strike the air molecules, freeing even more electrons. The resulting positive ions will also pick up energy from the electric field and be pulled toward the cathode. The more chemically reactive molecules, such as oxygen (and to a lesser extent, nitrogen) will combine chemically with the titanium. Other molecules, such as the chemically inert argon, may still have enough energy to bury themselves in the metal of the cathode.

The high voltage for the ion pumps comes from individual power supplies in the service buildings. Red cables (a safety standard for high voltage) connect the power supplies to the pumps. These supplies are always on, even during accesses, to maintain a good vacuum in the beam pipe. It is important to remember not to lick the leads when accessing the enclosure, nor to use wet dirty scissors when severing the cable.

The cathode and anode assembly is encapsulated in permanent magnets, as shown in Figure 47, which bathe the components in a magnetic field of about one kilogauss. The magnetic field causes all of the charged particles to spiral, increasing their path length and their chances for ionizing ambient air molecules.

The pumping capacity of an ion pump, measured in units of liters/second (L/s), is proportional to the number of ions removed from the beam pipe. The vast majority of ion pumps in the Recycler ring are 30 L/s. There are a few locations where more powerful pumps are needed, such as the regions connected to the beam transfer lines, and in the vicinity of the RF cavities, lambertsons, cooling tanks and special diagnostic equipment. Additional ion pumps have been added here to handle the high gas load.

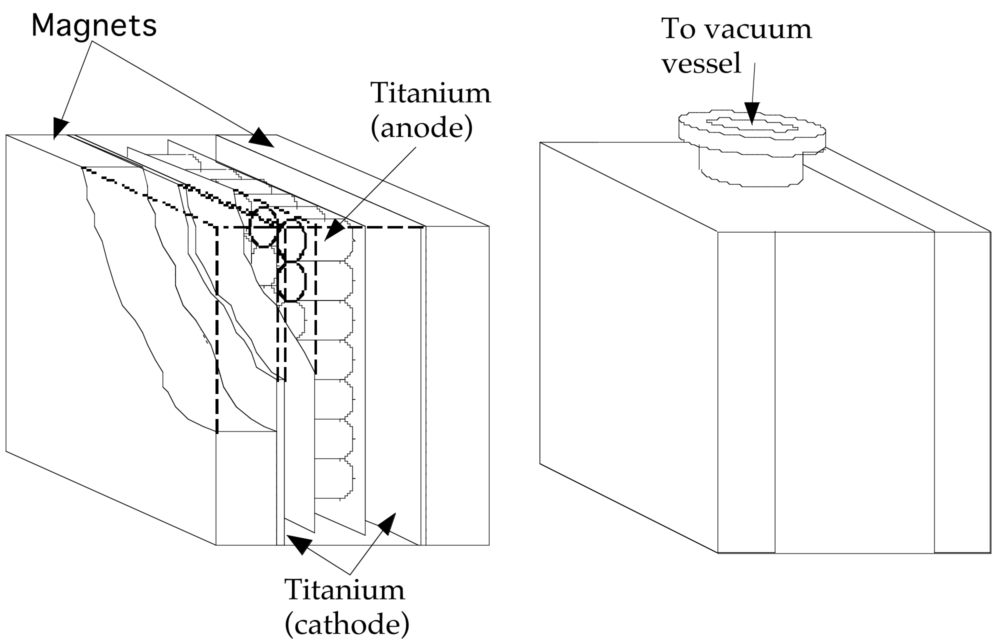

Figure 47: Cut-away and exterior views of an ion pump. The pump shown here has a titanium anode and two titanium cathodes.

Currently, the number of ion pumps in the Recycler is approximately one per half-cell. They are named for their three-digit location number in the tunnel. For example, the ion pump at Recycler tunnel location 629 is called IP629. Some half-cells have more than one ion pump, so they require further identification. In the ARC Regions, the ion pumps are given numerical designations. At the Q504 location, for example, we have two ion pumps 4051 and 4052. In the straight sections, the ion pumps are given alphabetical designations. At Q100, for example, the ion pumps are designated IP100A through IP100G.

Under normal circumstances, the ion pumps are sufficient to maintain proper vacuum in the Recycler ring. If, however, there is an air leak, the pump will respond by drawing more current as it impounds the air molecules. If vacuum becomes worse than approximately 10-5 torr, or if the pump is shorted, the titanium cathode will overheat and begin to outgas its store of trapped molecules. At that point the ion pump is unable to improve vacuum, and it trips off. If the leak is bad enough, several pumps will trip off sequentially, in a process called cascading, as each pump tries to take over the work of those who have tripped before them. Sometimes it is possible to turn the pumps back on after they have had a few minutes to cool, but if not, an access must be made to fix the leak, and the turbo pumps are brought in to recover the vacuum.

Titanium Sublimation Pumps

While found elsewhere at Fermilab, TSPs are most prevalent in the Recycler with TSPs at each half–cell location. These pumps are necessary to maintain the level of vacuum required for the Recycler. The TSP cans are adjacent to the beam pipe rather than being the beam pipe itself.

Sublimation pumps are a form of “getter” pump, operating on the principle that chemically stable compounds can be formed between gas molecules (H2, O2, N2, CO, CO2) and the getter (titanium). In this context, a “getter” is the chemically reactive material that gas molecules combine with to form stable compounds. Getters cannot pump chemically un-reactive noble gases, such as argon and helium.

Unlike ion pumps, which are powered all of the time, the sublimation pumps are powered infrequently. The sublimation pumps are "fired" over a 5-minute period to sublimate approximately 10-20 monolayers of titanium onto the pump's interior surface. As gas molecules make contact with the getter film, stable compounds are formed and the vacuum pressure improves since there are fewer gas molecules in the beam pipe volume. Because of the chemical nature of the TSP, there is no direct way of measuring a TSP’s pumping rate. To determine if the TSP’s are still pumping, one needs to look at the ion pump read back levels to determine if the TSP’s in a particular region need to re-fire.

During normal operation, sublimations are spaced months apart. Each Recycler sublimation pump contains two filaments to extend the lifetime of the pump, with only one filament at a time is sublimated. Because TSP’s have no effect on inert gases, ion pumps are still an integral factor in keeping the Recycler vacuum at its best achievable level.

Sector Valves

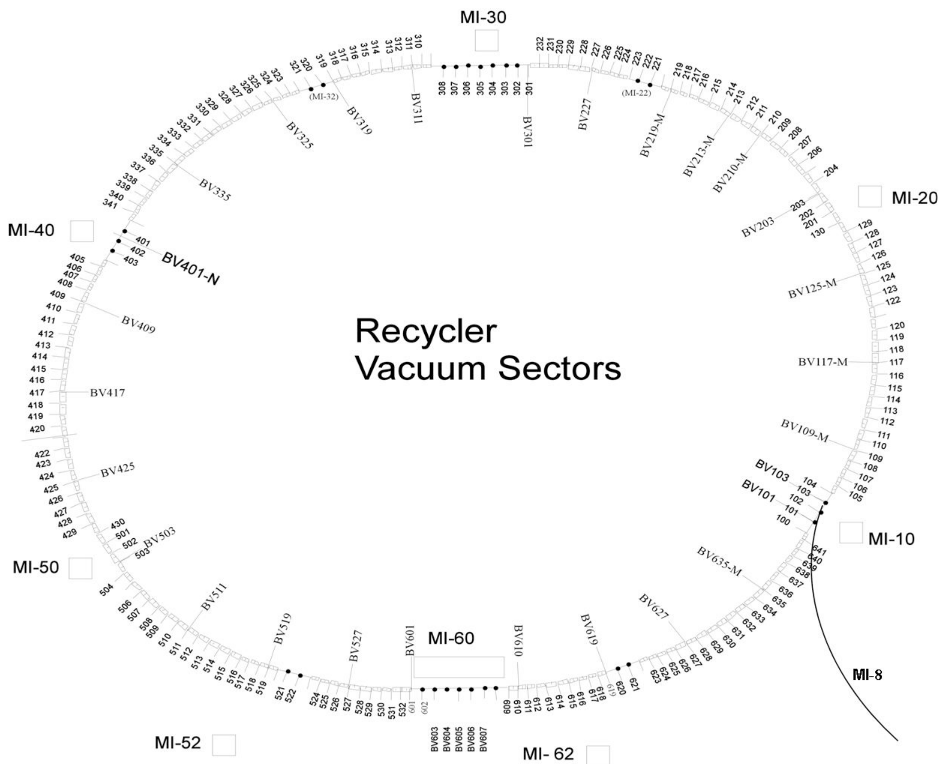

Figure 48: The Recycler Ring broken into its vacuum sectors.

The Recycler ring is divided into 30 vacuum sectors, as shown in Figure 48. Sector valves can be closed to create a barrier to bad vacuum, thereby isolating vacuum problems to relatively short sections of the ring. Coincidentally, the valves are also effective barriers to the beam. For this reason, sector valves are also commonly referred to as beam valves. The status of the valves can be read through the controls system. The beam valves are designated “BV” followed by the location number.

The length of the sectors varies between 7 and 9 half-cells. The beam valves are closer together where vacuum instabilities are most likely, or where there are connections to external beamlines. Sector valves bound most of the straight sections. The Cooling Sectors have several beam valves internal to the sector, called intermediate valves, for isolating short groups of cooling cavities. There are also two beam valves, BV303 and BV307, which are used by the safety system to control circulating beam in the Recycler ring. These coasting beam valves will close if the electrical permit is present and the radiation permit is not if the electrical permit is not present, it is impossible to circulate beam, and there is no need to close the valves.

The sector valves are held open by air pressure. Air compressors at MI-20, MI-40, and MI-60S maintain pressure of 80 to 100 psig in copper tubes. The air is then usually cooled and dehumidified by a dryer unit. The unit can be manually valved out if necessary. The tubes enter the tunnel through penetrations and run parallel to the cable trays on the ceiling and can be tapped from the tubes at intervals for use by the valves. Although each valve is provided with a ballast tank, the valves will eventually close in the event of a compressor failure. This is why the pneumatic valves are found closed following a power outage.

Fast-Acting Valves

This section describes the Recycler vacuum protection system that would be triggered in the event of a failure of the vacuum tubes inside the Electron Cooling System's pelletron. The pelletron is pressurized with sulfur hexafluoride (SF6) gas at 75 psi. A catastrophic vacuum tube failure would cause a pressure wave of SF6 gas to travel through the Electron Cooling vacuum system toward the Recycler at velocities greater than the speed of sound, reaching the Recycler in 10 ms and contaminating the entire Recycler vacuum system in 5.2 seconds. In the event of such a failure, the fast-acting valve system described here will protect the Recycler vacuum system against SF6 contamination as well as high internal pressures that could rupture vacuum components.

This system is located in the MI-31 building and in the 30 Sector of the Recycler Ring. Figure 49 shows the layout of the components. Two fast-acting valves are located on the MI-31 side and two on the Recycler side. The function of the first set of valves (in Electron Cooling's Supply and Transfer Lines) is to minimize the amount of SF6 gas that will reach the Recycler. The second set of fast valves, located on either side of the cooling section in the Recycler, will stop the flow of gas from reaching the rest of the Recycler vacuum system.

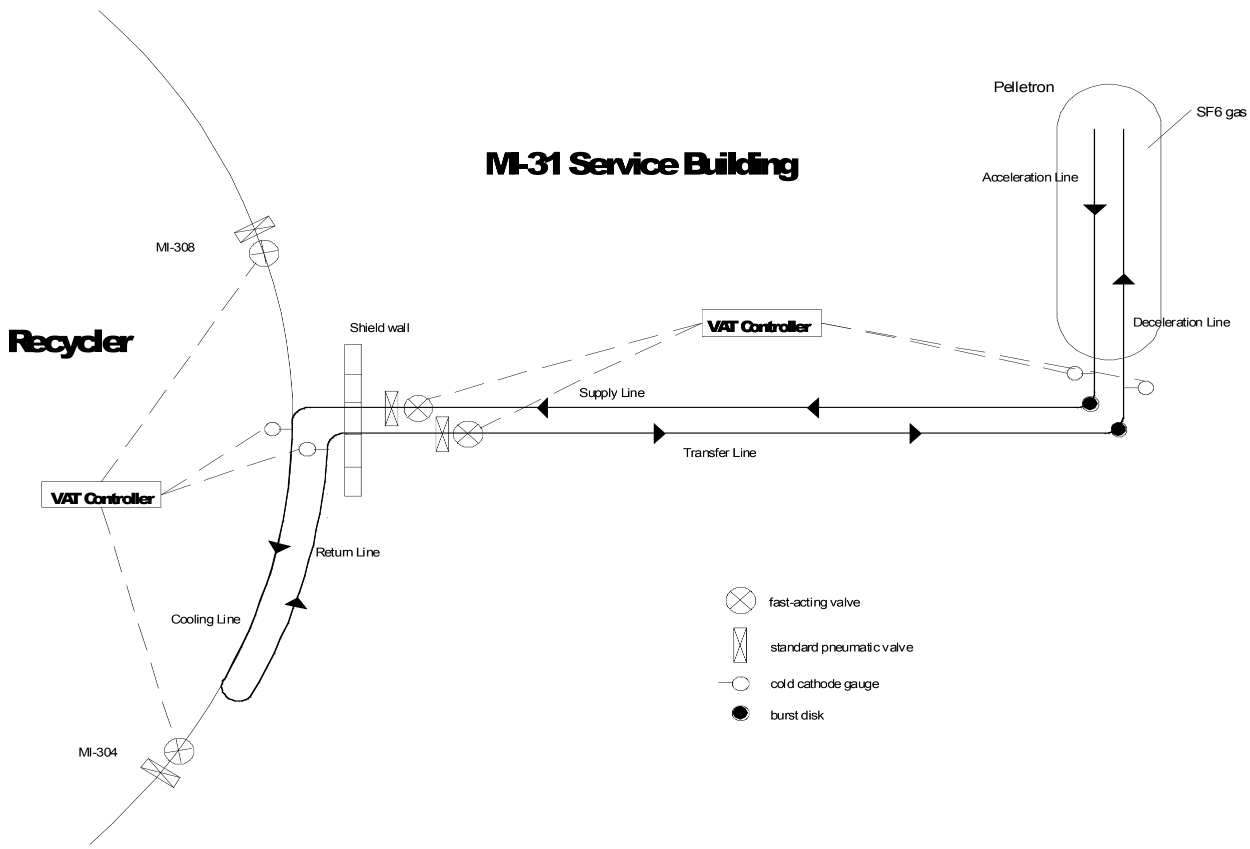

Figure 49: Schematic drawing of the Recycler/Electron Cooling SF6 vacuum protection system.

The valves on the MI-31 side are controlled by the cold cathode gauges and controllers on that side. The Recycler fast valves are controlled by the gauges and controllers on the Recycler side. A trip on the MI-31 side will not affect the Recycler valves. Specifically, the Supply and Transfer Line fast valves (BVS04F and BVT03F) will close if either VAT gauge under the pelletron (VATA06, VATD06) reads a pressure greater than 5E-4 Torr. Likewise, the Recycler fast valves (BV304F, BV308F) will both close if either VAT gauge on the 90˚ bend magnets in the Recycler (VATS06, VATT01) reads greater than 5E-4 Torr.

The distance from the first set of sensors (under the pelletron) to the Transfer/ Supply Line valves is approximately 11 m. A pressure wave of SF6 gas could travel this distance in 4.5 ms. In the Recycler proper, the distances between the sensors and fast valves are 34.4 m and 16.33 m for the RR308 and RR304 locations, respectively. The gas would reach the RR308 fast valve 112 ms after passing the sensors. It would reach the RR304 valve 53 ms after being seen by the sensors.
The four fast-acting valves are not completely leak tight. Thus, each one is backed by a standard pneumatic valve, which is activated through the CIA crate vacuum interlocks. They take a few seconds to close. These valves will close if three or more of the four selected ion pumps on either side of the valve trip off (this usually occurs at approximately 10E-6 Torr). In the tunnel, these valves are named BV304S and BV308S. In MI-31, they are BVS04S and BVT03S. The ion pump permits will close the fast and slow valves.

Two vacuum burst disks are located on the 90˚ bend magnets under the pelletron. In the event of vacuum system overpressure, the disks will rupture, relieving the pressure in the system. This will occur between 15 and 25 atm. Relief of this system is in accordance with the current ODH status of the MI-31 high-bay (ODH Class 0) and silo (ODH Class 1).

Vacuum Instrumentation

There are three methods for measuring vacuum in the Recycler. The first is by means of a Pirani (Pe-RAH-nee) gauge. Piranis work by heating a wire with a constant current source and measuring its resistance. The resistance of the wire increases with temperature. Air molecules near the wire carry away the heat. The better the vacuum, the fewer the air molecules and the hotter the wire, so the resistance of the wire can be interpreted by the applications software as pressure. Pirani gauges read back measurements from normal atmospheric pressure down to 10-3 torr and are most accurate at the lower end of that range.

One alias for a Pirani gauge is “convectron” gauge, because heat is carried away by the convection of the air. There is one Pirani gauges in every sector, including the “subsectors” in the cooling sectors. They are found near the bottom of the “tee” or “cross” assemblies attached to the TSP cans, and can be recognized by their blue caps.

The second method of measuring vacuum is through the ion pumps. The arc of an ion pump creates a current that can be read by the power supply. Since the current is proportional to the number of molecules ionized, it can be interpreted by applications software as pressure. In this way, the vacuum readback on an ion pump is derived from the arc current. The readback is valid for the range in which the ion pumps are active, from about 10-4 torr down to at least 10-9 torr.

The third method of measuring vacuum in the Recycler is using the ion gauges. The ion gauge works by heating a filament grid, which releases electrons within the gauge. These electrons ionize the ambient air molecules, which are attracted to a collector. Meanwhile, the electrons are then recaptured on the grid. The current that is measured from the collector is directly proportional to the vacuum pressure. It should also be noted that these gauges are gas specie-sensitive.

Vacuum Controls

On the vacuum applications page, currently R55, a sector is usually named for the service building at which its equipment resides. If there is more than one vacuum sector associated with a service building they are numbered consecutively. At MI-10, for example, there are three Recycler vacuum sectors named RR10-1, RR10-2, and RR10-3.

The subpages on R55, such as RR10 and RR20, separate the vacuum system into "houses" for easier viewing. A house is defined by the service building at which the individual controls or power supply that controls the vacuum device can be found. The "Global" page has readbacks for the average ion pump pressure readings for each house, the open/close status of the beam valves, and their permit status. This page also allows the user to see status of and control the ion pumps, and to see ion gauge readings for the entire ring.

With the exception of RR10 and some of the very short sectors, each sector has three sets of ion pumps that can be used to interlock the valves. Each set includes four pumps, at least two of which need to be on in order to grant a permit.

At each service building, the vacuum system is controlled by the CIA (Controls Interface Adapter) crate. There are several kinds of cards in the CIA crate. These cards control the on/off status of vacuum components as well as returned read backs for these devices. As a general rule, there is one ion pump card for every 6 ion pumps. The first four channels on each card can be used to establish the beam valve permits. There are usually two ion pumps monitored by each card that are not represented on the front panel.

In some of the sectors, three ion pump cards, controlling 12 ion pumps, are tied into the beam valve permit. The three groups of pumps are spaced apart to more accurately sample the length of the sector. If three pumps within one of the selected groups fail, the permit is withdrawn.

The sector valve card is what actually decides whether or not a beam valve should close based on the status of the permit. It directly controls only the upstream beam valve for the sector. When the permit is dropped, the appropriate sector valve cards close their associated valve. To re-open a valve between two sectors, a permit must be obtained from both the upstream and downstream sectors. For example, in order to open BV101, permits must be granted from both RR10-2 and RR30-2. There are only 8 pneumatic sector valves placed around the Recycler are interlocked to selected readbacks from the ion pumps. Given that, most of the Recycler sector valves can only be closed manually.

The sector valve card also displays the permit status and the open/closed status of each of the two valves. There is also a "request" bit which indicates that a valve has been asked to open but has not yet complied. The same information is available on R55 on the "Sector Valve" sub page.

As an example of how the ion pumps interact through the CIA crate to close beam valves, consider the scenario in Figure 50. In the figure, IP641, IP100, IP1021, IP1022, IP1023 and IP1024 have all tripped off. With the majority of the pumps off, the ion card representing those pumps interpret the information as a vacuum problem, which could spread to other sectors. It sends a "NO" bit to the sector card for RR10-3, which orders BV101and BV103 to close.

The Recycler CIA crate interfaces to the outside world using an Arcnet loop, not through CAMAC as with some of the other machines. There is a 186-bearing card in the rightmost slot which functions as the interface. All of the CIA crates in the Recycler talk to a dedicated VME front end, called MIVAC, at MI60S, rack 123. The Arcnet loop uses the CATV cables—the same system that carries TV images—to relay data back and forth from the CIA crates to MIVAC. The VME, in turn, is linked to the general controls system through Ethernet. MIVAC uses a MOOC software platform to communicate with the CIA crates, so many of the controls related error messages will read back as MOOC errors.

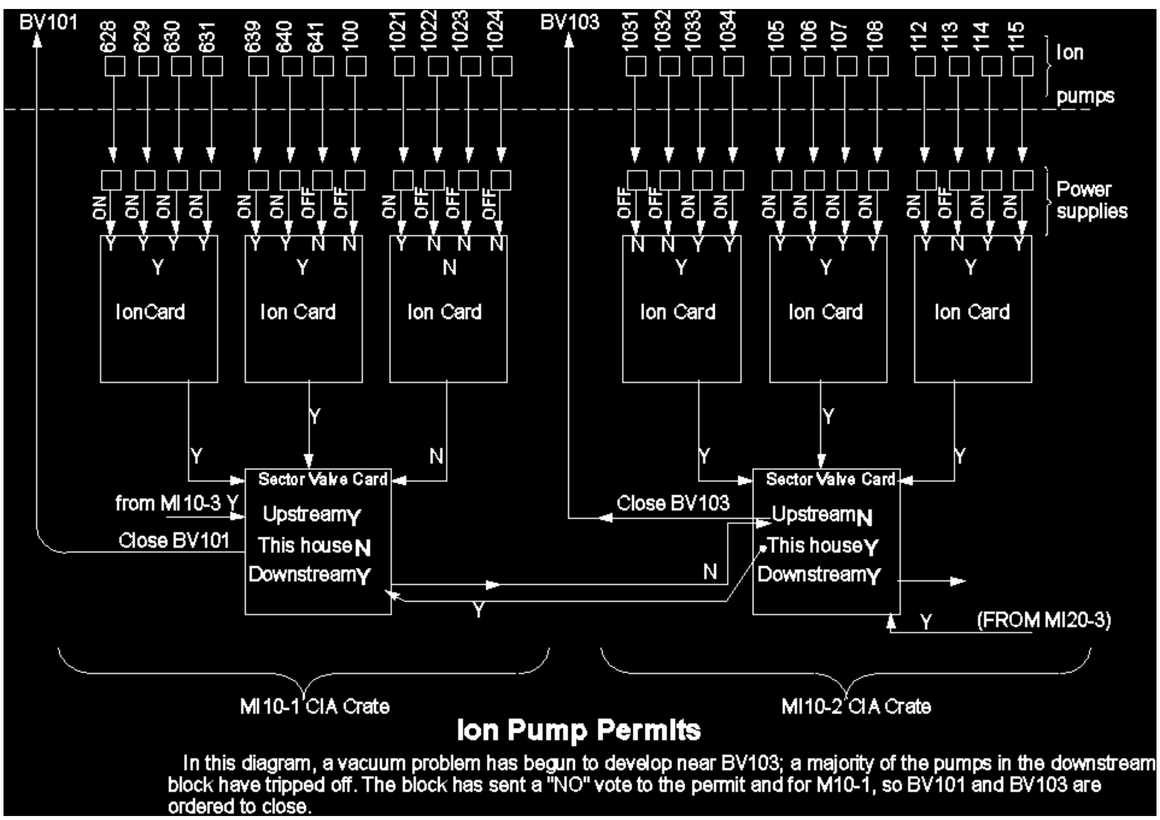

Figure 50: A hypothetical vacuum scenario to demonstrate the permit logic

MIVAC relays information on the vacuum for use by the rest of the controls system. It also analyzes the ion pumps tripped off and generates alarms. Alarms are generated when any three consecutive ion pumps are off, or if 20% of the total number of pumps in a sector trip. The alarms are independent of the valve permits; they only provide a warning that vacuum problems are developing.

Water cooling

The Recycler is mostly comprised of permanent magnets that need no water-cooling. There are several components that do need water-cooling.
The VDPA magnets in the transfer lines are water-cooled. They are tapped off the main 95 LCW header in the tunnel. The VDPA PS's upstairs are water-cooled as well, likewise tapped off the main header. The powered skew quadrupoles in the MI-22 transfer line are water cooled as well.

The stochastic cooling kicker tanks are water cooled, as are the TWT's in the amplification chain.
The RF cavities downstairs in the MI-60 region are water-cooled. The PS's upstairs in the North end of the MI-60 service building will be water-cooled in the near future as well.

Ecool at MI-31 has its own closed loop system that makes up water from the main header at MI-30. This make-up must be done manually.

Power

Much like the water-cooling section, the power section of the utilities is pretty short. That is one nice advantage to a mostly permanent magnet ring.

All of the correction elements at each house for the Recycler are powered with house power. One thing to remember is that when Feeders 86, 87, and 89 are switched off for MI tunnel access, these correction elements are still powered. Keep that in mind next time you are climbing on top of the MI when in the tunnel for a controlled access. This is allowable because the connections in the tunnel are covered, and these magnets carry less than the 25A allowable under controlled access restrictions. There are exceptions to this however. The skew quad trims are tied into the safety system as their leads are not completely covered.

The VDPA PS's at MI-30 are powered through the DSPHP-MI30-1A-7 safety switch at MI-30. This is the switch taken during the MI switch-off procedure, and ensures these dipoles are not powered during an access.

References

Bringing together all the information contained in this book would not have been possible without the many contributions from Cons Gattuso. I would also like to thank Martin Hu, and Dan Broemmelsiek for helping to explain the many things contained in this book that's would have been unattainable without their help. I would also like to thank operators: Maggie Stauffer, Chip Edstrom and Steve Baginski for their input on the book as well. Here is a list of all the documents used, hopefully in roughly the order they were used:

1. G. Jackson, "The Fermilab Recycler Ring Technical Design Report", FERMILAB-TM-1991.
2. G.W. Foster, "Recycler Lambertson Magnetic Design", MI note 256, March 1999.
3. G.W. Foster, "Magnetic Design of Recycler Ring v.18 Gradient Dipoles", MI note 227, January 1998.
4. G.W. Foster, "Recycler Quadrupole Magnetic Designs", MI note 249, September 1998.
5. G.W. Foster, "Mirror Magnet Designs for the Recycler Ring", MI note 212, August 1998.
6. D. Johnson, "Revised Recycler Pbar Transfer Lines", MI note 161, December 1997.
7. Marchionni, T.G. Anderson, G.W. Foster, C. Gattuso, M. Hu, D. Johnson, C.J. Johnson, C.S. Mishra, A.R. Oleck, H. Piekarz, M. Syphers, M.J. Yang, "Performance of Antiproton Injection and Extraction Transfer Lines of the Recycler Ring at Fermilab", Proceedings of the 2003 PAC.
8. Bhat, "Applications of Barrier-Bucket RF Systems at Fermilab", Beams Document 2209-v1, March 2006.
9. G. Jackson, "Operational Experience with the Permanent Magnet Recycler Ring at Fermilab", Proceedings of EPAC 2000.
10. C. Bhat, "A First Look at the Longitudinal Emittance in the Fermilab Recycler Ring", MI note 287, August 2002.
11. C. Bhat, B. Chase, B. Foster, H. Kang, J. Maclachlan, K. Seiya, P. Varghese, D. Wildman, "Update on Fast Recycler Stacking", 2006.
12. C. Bhat, "Pbar Momentum Mining in the Recycler", Beams Document 933-v2, January, 2004.
13. J. Dey, D. Wildman, "Wideband RF System for the Fermilab Recycler Ring", Proceedings of the 1999 PAC.
14. G. Jackson, "Specifications for the TE01 Signal Transmission of the Recycler Stochastic Cooling Systems", MI note 233, May 1998.
15. R. Pasquinelli, "Recycler Stochastic Cooling 'Peanut' Layout", Recycler Department Webpage, September 1998.
16. J. Marriner, "Recycler Instrumentation Status", DOE Review, July 2003.
17. P. Odier, "A New Wideband Wall Current Monitor", CERN.

18. F. Nolden, “Instrumentation and Diagnostics Using Schottky Signals”, Proceedings of DIPAC 2001.
19. B. Fellenz, J. Crisp, “An Improved Resistive Wall Monitor”.
20. F. Caspers, R. Jones, “Specifications for the LHC TW Pickup”, BNL Tune Feedback Workshop, March 2005.
21. T. Linnecar, W. Scandale, “Continuous Time Measurements Using the Schottky Detector”, IEEE Transactions on Nuclear Science, Vol NS-30, No. 4, August 1983.
22. D.A. Goldberg, W. Barry, G.R. Lambertson, F Voelker, “A High Frequency Schottky Detector for Use in the Tevatron”, Proceedings of PAC 1987.
23. P. Wilson, “Recycler Flying Wire System Status”, December 2003.
24. D. Sun, “Design and RF Measurement Results of Schottky Pickup for Recycler and Tevatron”, RF Technote 49, January 2003.
25. R. Pasquinelli, “Microwave Schottky Pickups at Fermilab”, Proceedings of Meeting on Measuring Tunes in TeV, January 2004.
26. R. Webber, J. Crisp, P. Prieto, D. Voy, C. Briegel, C. McClure, R. West, S. Pordes, M. Mengel, “Fermilab Recycler Ring BPM Upgrade Based on Digital Receiver Technology”, Beams Document 1328-v1, September 2004.
27. S. Bhashyam, N. Chester, D. Hicks, A. Marchionni, J. Misek, V. Razmyslovitch, P. Wilson, J. Zagel, “On the Mechanics of Recycler Flying Wire Assembly”, Technical Division Note TD-02-044, November 2002.
28. J. Fitzgerald, “Recycler Ion Clearing Field Power Supplies”, Beams Document 414-v1, January 2003.
29. L. Nobrega, “Recycler/Electron Cooling SF6 Vacuum Protection System”, Fermilab Accelerator Division Mechanical Support Department Engineering Specification No. 1302-ES-296422, September 2005.