3.1 Protons on the Antiproton Target

3.1.1 Introduction

3.1.1.1 Justification

The goal of the Run IIb upgrades is to increase total antiproton yield to collisions. This should lead directly to the necessary luminosity required to reach the 15 fb^{-1} goal for Run II. Increasing the antiproton yield will involve increasing the proton flux on the antiproton target, improving the collection efficiency, and increasing the antiproton cooling rate. This project deals with increasing the proton flux on the antiproton target.

The machines responsible for providing protons to the antiproton target are the Linac, Booster, and Main Injector. The Linac provides about 50 mA of beam current to the Booster at 400 MeV. The Booster then accepts Linac beam for nine or ten revolutions to provide the Main Injector with about 5×10^{12} protons at 8 GeV. The Main Injector accelerates the beam to 120 GeV for the antiproton production target. The antiproton Debuncher ring and the Booster have the same circumference, so only a single Booster batch can be applied to the target before the Debuncher ring is full. The Main Injector has a circumference seven times the Booster's, so it can provide the target with beam and still have five batches of space for other uses (ie NuMI).

The most straight-forward method of increasing flux on target would be to increase the current in the Linac, or inject more rotations of Linac beam in the Booster. This would lead to higher current in the Main Injector. Unfortunately, space charge effects limit the amount of beam that can be injected into the Booster[7]. Injecting more than 500 mA into the Booster reduces its efficiency to the point where no more beam reaches the Main Injector. In order to increase beam intensity, we need to reduce the effect of space charge in the Booster, or we need to take advantage of the extra space in the Main Injector.

Slip stacking takes advantage of the extra longitudinal phase space in the Main Injector. It is a method of injecting two batches of Booster beam into the Main Injector and combining the two batches into one double charged batch before extracting to the antiproton target. Two batches of beam are injected consecutively into the Main Injector with slightly different momenta. The different momentum batches have slightly different velocities, and one batch eventually overtakes the other batch. When the two batches completely overlap, the RF voltage is increased to provide a bucket big enough to contain the entire momentum space of the two batches. There are two critical and conflicting parameters that determine the efficiency of slip stacking. First, the momentum separation between the batches must be large enough, compared to the bucket size, to minimize the interference between the two batches. Second, the momentum separation should be minimized at capture time in order to keep the longitudinal emittance of the final batch low.

Of course the momentum separation between the batches must not be larger than the momentum acceptance of the Main Injector. If the longitudinal emittance of the Booster beam is maintained at about 0.1 eV-s, the bucket height required to contain the beam in the absence of beam loading is about 9 MeV. The Main Injector has a measured momentum acceptance of about 1.4% (See Figure 3.1.1). This corresponds to about 124

MeV of acceptance or about fourteen Booster batches stacked end to end in energy space. This is more than enough aperture for successful slip stacking. Because of the large momentum aperture in the Main Injector, simplifying operational issues can take precedence over minimizing the momentum aperture required by the slip stacking process.

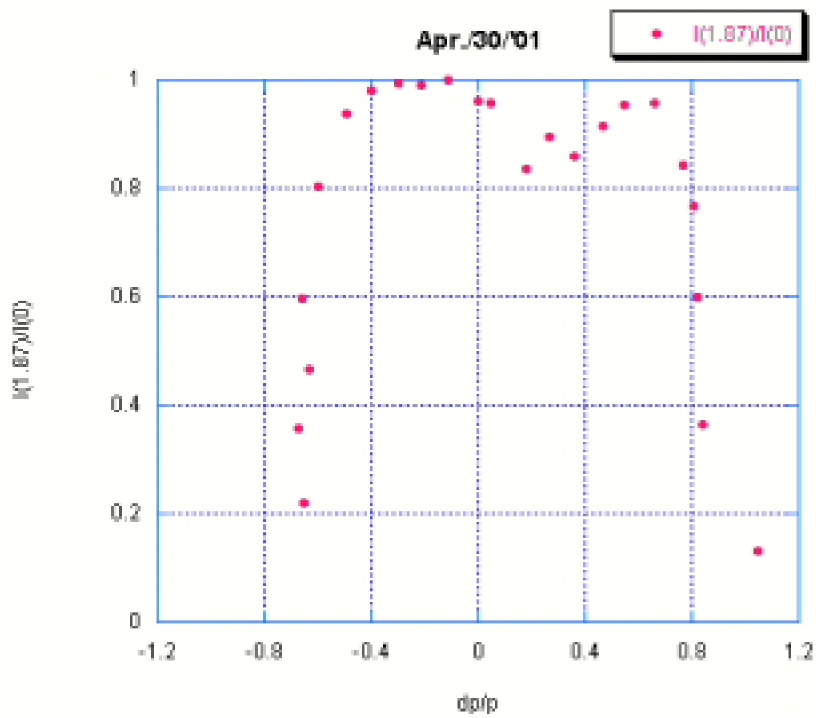

Figure 3.1.1 *Main Injector momentum aperture. This plot measures main injection efficiency as a function of momentum offset. The momentum offset is created with a three bump at the radial position feedback pickup*.

3.1.1.2 Slip Stacking Process

The slip stacking process in the Main Injector can be described in four phases: first batch injection, second batch injection, slipping, and capture. The first batch injection is essentially the same as any injection from Booster to Main Injector. If momentum aperture were a premium in the Main Injector, the first batch would be injected off momentum to the upper edge of the aperture. Because there is so much momentum aperture in the Main Injector, there is no need to inject the first batch off momentum. Therefore, the first batch is injected on the Main Injector's central orbit (Figure 3.1.2), just like all other Main Injector cycles. The only difference between first batch injection and standard injection is the capture voltage is about a factor of ten less for slip stacking. Since the Booster cannot match to such a low voltage, the Booster executes bunch rotation before extraction. Between the time that the first batch is injected and the second batch is injected, the first batch is decelerated off the central orbit sufficiently to make room for the second batch in longitudinal phase space (Figure 3.1.3).

After the first batch has been injected and decelerated, the RF voltage that is to be synchronous with the second batch is activated. (It is deactivated during first batch injection to minimize emittance growth caused by interference between the two RF voltages.) The second batch is also injected on the central orbit of the Main Injector, just behind the first batch (Figure 3.1.4). Since the first batch was decelerated off the central orbit, it now has a lower velocity than the second batch, and the particles from the second batch will overtake the first batch (Figure 3.1.5). The total RF voltage in the cavities and the total RF component of the beam current see 100% amplitude modulation as the batches slip past each other. Damage to the longitudinal emittance of the two batches is minimized because the frequency separation is significantly greater than the synchrotron frequency.

Once the two batches are completely aligned the capture process begins. The two RF voltage waveforms used for each batch jump in frequency to the same value. This value is the average frequency between the two slipping frequencies. Simultaneously, the total RF voltage is increased to create a bucket that encloses the entire momentum spread created by both batches. Presently, there is no good way to preserve the longitudinal emittance of the combined batch, and it is assumed that the beam will eventually filament to some emittance higher than twice the single batch emittance. After the beam is captured (Figure 3.1.6), it is accelerated to 120 GeV with the standard stacking ramp. It is important that the emittance is kept low enough to be accelerated through transition in the Main Injector. It also must be low enough to fit within the momentum aperture specifications of the accumulator.

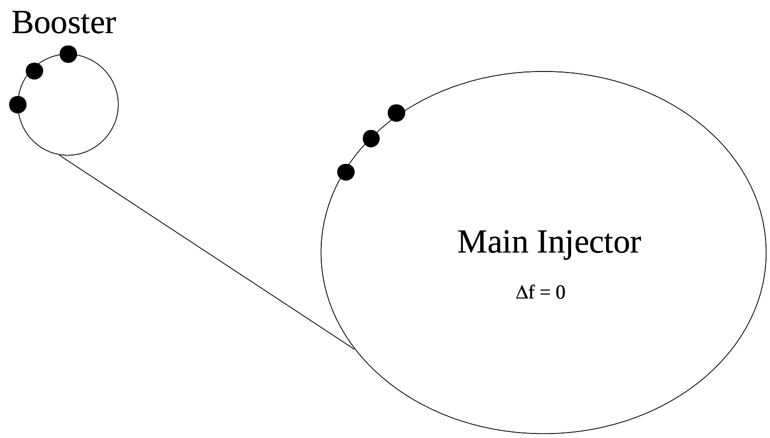

Figure 3.1.2 *First batch injection: First batch injected on Main Injector's central orbit.*

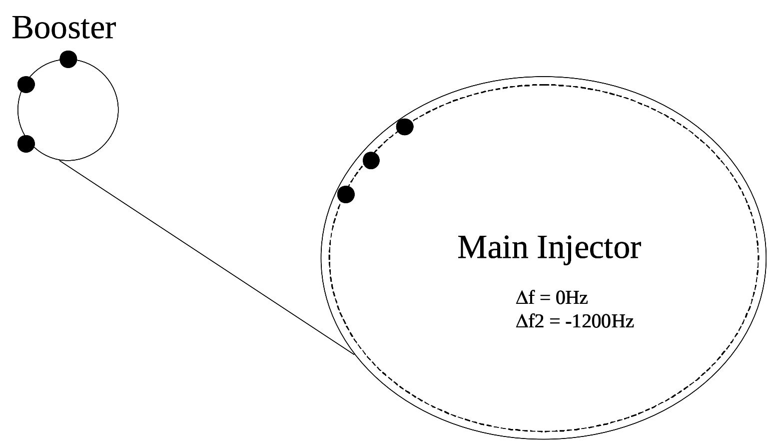

Figure 3.1.3 *First batch decelerated to make room for second batch.*

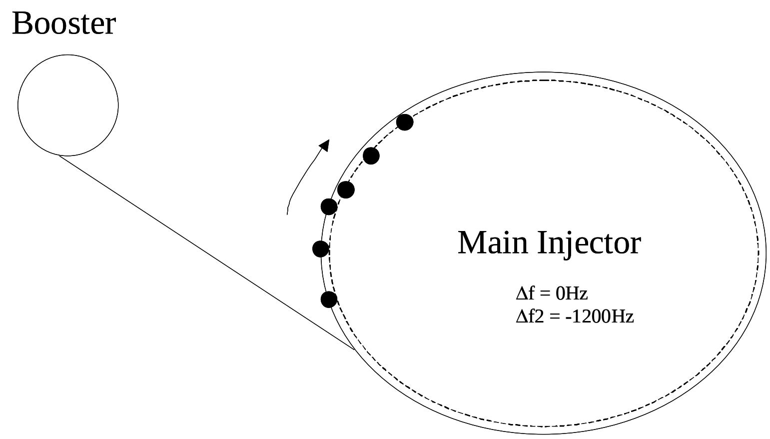

Figure 3.1.4 *Second batch injection on the Main Injector central orbit, just behind the first batch.*

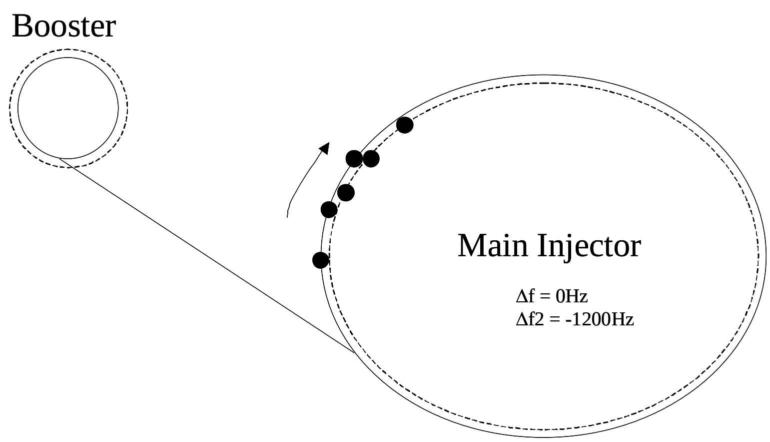

Figure 3.1.5 *Batches slipping. The second batch begins to overtake the first batch because of the differences in velocity.*

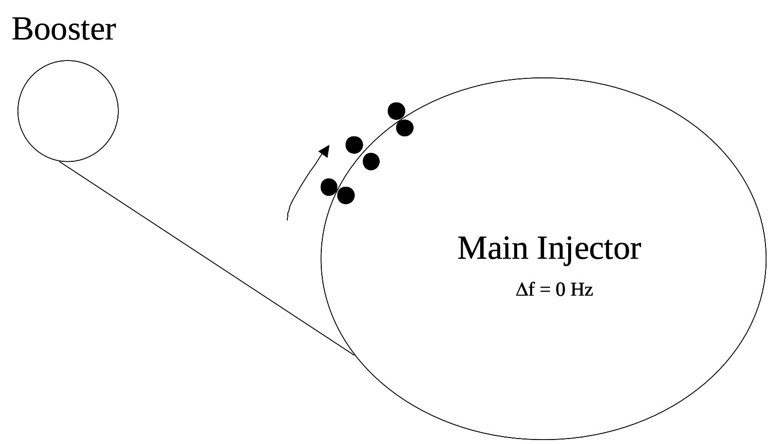

Figure 3.1.6 *Batch profile immediately after capture*

3.1.1.3 Slip Stacking Efficiency

Slip stacking will nearly double the overall flux on the antiproton target. There losses in efficiency due to the larger longitudinal emittance and extra cycle time is required to inject two batches and slip together. One Booster cycle occurs every 15Hz. Two batches can slip together in approximately one Booster batch time. This time becomes critical when NuMI operation is considered. For NuMI operation, after the two batches for stacking have been combined and the frequency offsets zeroed, five more batches will be injected. If the slipping time is even slightly greater than a Booster cycle, an entire Booster cycle will be wasted until NuMI injection can begin (the Booster magnet power supply is a resonant circuit and can't be held off). The following table illustrates the improvement over standard stacking given different slip stacking scenarios.

Because of the small buckets sizes and large beam currents involved in slip stacking, beam loading becomes the dominant barrier to a reliable slip stacking process. Beam will fall out of the buckets at the currents and voltages specified without proper beam loading compensation. These effects have been simulated, and some beam loading compensation systems have already been tested on the Main Injector. The operation systems, however, are not sufficient for slip stacking at high intensity. The most challenging aspect of the project is the design of a state of the art beam loading compensation system.

Slip stacking studies have already begun in the Main Injector. Many modifications to the RF system were required to facilitate the basic mechanics of slip stacking. The rest of the paper discusses the details of these modifications. It will also cover slip stacking simulation work, and it will cover the details of the beam loading issues and possible cures.

Operation Mode	Booster Cycles	Cycle Time	Relative Intensity Of Stacking Pulse	Stacking Improvement
Stacking	1	1.466s	1	1
Slip Stacking	3	1.6s	1	1.8
Slip Stacking w/ NuMI	8	1.93s	1	1.5

Table 3.1.1 *Relative Improvement in Proton Flux for Different Operating Scenarios*

3.1.2 Basic Hardware

Slip stacking is a predominantly longitudinal process, and the control for longitudinal processes lies in the RF system of the accelerator. The only hardware changes needed to facilitate slip stacking are in the RF control systems. The actual changes can be easily supported by the existing cavities and power supplies

3.1.2.1 LLRF

Many changes were made to the LLRF system when the Main Injector replaced the main ring. The system is now primarily digital. The voltage-controlled oscillator (VCO) was replaced with a numerically controlled oscillator (NCO) controlled by a DSP (Figure 3.1.7). Each NCO drives the IF ports of an I-Q modulator with its local oscillator input at 50 MHz. The NCOs have a 32 bit frequency register, making the frequency setting precise to about 2mHz. The DSP can update the NCO frequency value at about a 10 kHz rate, which is more than sufficient for the acceleration ramps in the Main Injector.

The new LLRF system also performs most of its feedback control digitally. It is equipped with a number of digitizers that look at error signals generated by the instrumentation. The instrumentation includes a phase detector for phase feedback and a beam position monitor for radial position feedback. The signals are sampled and stored in the DSP. The DSP then performs the appropriate filtering and gain control for the loops. The DSP will then update the frequency value of the NCO based on a combination of the frequency program and the feedback values.

Another task that the LLRF system must perform is the generation of the revolution marker. This marker remains synchronous with the beam revolution frequency, and it is used to define bucket enumeration. A marker that is synchronized with the first bunch in a batch of beam at injection will be synchronized with the same bunch for every injection. The marker is used to trigger kickers and instrumentation. Keeping the marker synchronized to the revolution frequency is straightforward. Since the RF output of the LLRF system must be synchronized with the beam, it can be used as a clock for a divide-by-N counter, where N is set to the harmonic number of the machine. To maintain consistent enumeration, either the marker must be reset by the same signal that triggers Booster extraction, or the marker itself must trigger Booster extraction. The Main Injector LLRF system uses both techniques. When the Main Injector is empty, the Booster resets the marker, and when there is beam already in the Main Injector, the Main Injector revolution marker triggers Booster extraction.

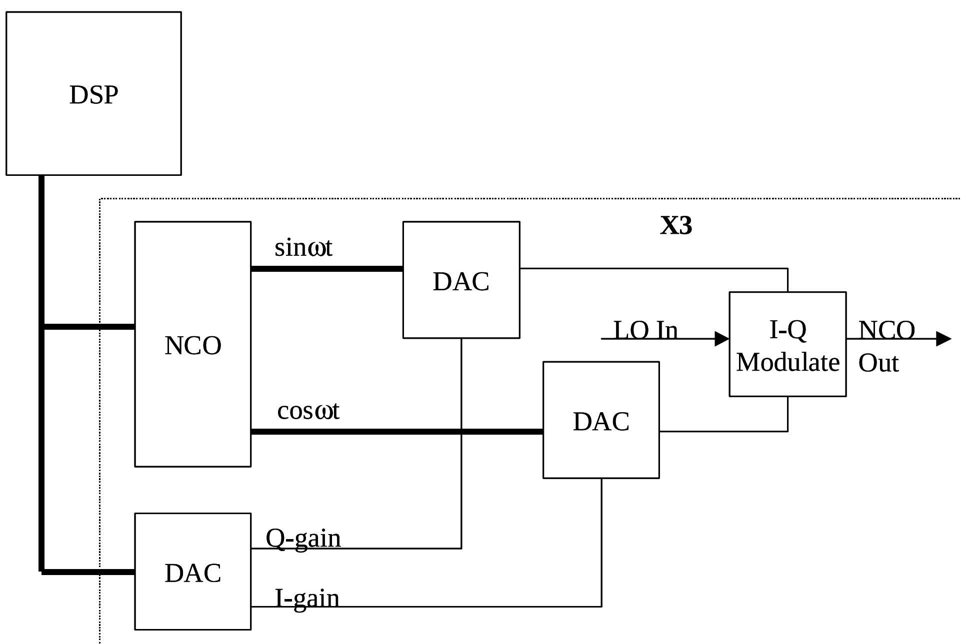

Figure 3.1.7 *Block diagram of NCO module. The DSP controls the NCO frequency setting, and the DAC that controls the phase shift prior to the I-Q modulator. The LO phase is matched on all three modulators, and the three NCOs are set to run the same frequency ramps. This makes the phase shift between outputs independent and precise.*

During some Main Injector cycles, the total RF voltage seen by the beam must be dropped to a very low value, sometimes zero. Coalescing is one application where the voltage must be varied. The RF voltage is dropped to a low level to rotate the 53 MHz buckets for minimum energy spread. Then, the 53 MHz voltage is turned off, and a 2.5 MHz RF system is enabled to rotate multiple 53 MHz bunches into one "superbunch". The 53 MHz system is then re-enabled to recapture the bunch into a single bucket. There is a lower limit to the anode bias voltage, and it is not low enough for coalescing. Paraphasing is one solution to the problem, and it was incorporated into the LLRF system.

Paraphasing is a process in which the RF cavities are divided into two groups, and each group is driven by a LLRF signal with a different synchronous phase. If the difference in synchronous phase between the two groups is 180° and if the total amplitude of the two groups is balanced, then the beam will see a zero net RF voltage. Paraphasing can develop arbitrary RF voltage levels without varying the anode bias if the difference between synchronous phase angles can be controlled arbitrarily (Figure 3.1.8).

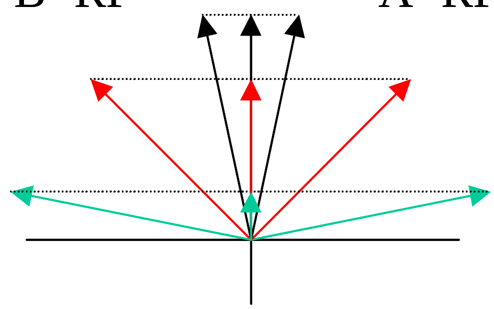

Figure 3.1.8 *Paraphase illustration. Arbitrary voltage levels, along the vertical axis, are generated by varying the phase difference between the 'A' and 'B' RF outputs.*

The LLRF system generates three independent RF outputs. Two of the outputs are used as the RF drive for the cavities, and each output drives a different group of cavities in the paraphase process (Figure 3.1.9). The third output is used as a beam synchronous signal for instrumentation and for generating the revolution marker. Independent NCO I-Q systems generate each output. Each NCO system has bipolar amplitude control on each I-Q leg, controlled by a DAC. This DAC is controlled by the DSP, and this gives the DSP the ability to place an arbitrary phase shift on any output.

For slip stacking, the Main Injector LLRF system must generate two RF signals with independent frequency controls. The distribution of the outputs to the cavities is already in place because of the paraphase system. Since each output has its own NCO system, one could program different frequency ramps to each output through the NCO. However, the NCOs have no means of communicating their relative phase offsets. The relative phase between the RF outputs becomes very important at the end of the slip stacking cycle. If the phase between the outputs is not adequately controlled, the RF seen by the beam could be paraphased at some arbitrary angle, causing large variations in capture voltage amplitude.

Instead of having independent frequency ramps for each NCO, it is better to use the independent paraphase control for each output. The paraphase system controls the phase shift on each output through an I-Q (In-phase-Quadrature phase) modulator. The phase shifter can generate a frequency offset on its output if the phase control is a ramped input (Figure 3.1.10). The DSP stores independent frequency offset tables for each output and converts the tables into phase ramps. It updates the I-Q DACs every time there is a phase step. This technique provides a significant advantage over changing the NCO frequency. The DSP maintains a very precise phase control between the two outputs, and it can dictate the exact paraphase angle at the end of the slip stacking cycle. Its greatest limitation is the DAC update rate, which limits the maximum offset frequency and the timing precision of the capture process.

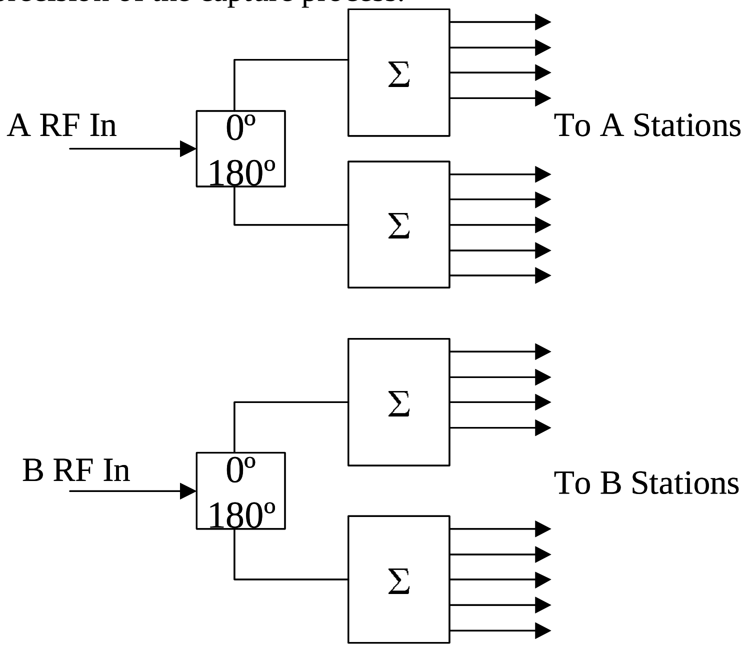

Figure 3.1.9 *Paraphase fanout. Two outputs of NCO module each drive half the RF stations. The hybrid is needed because the cavities are physically separated by half an RF wavelength.*

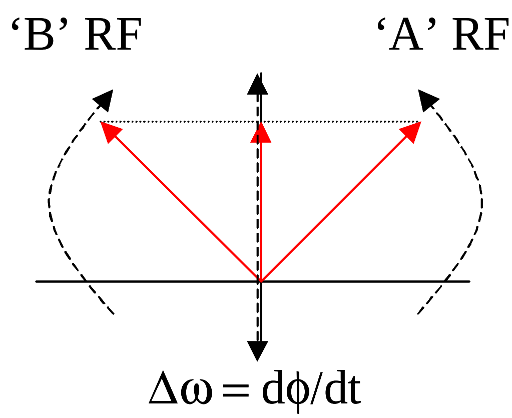

Figure 3.1.10 *Vector diagram showing paraphase control generating offset frequencies. The offset frequency is equal to the time derivative of the phase shift ramp.*

A preliminary slip stacking routine for the LLRF system has been tested on the Main Injector. First, both outputs are set to the nominal injection frequency before injection (Figure 3.1.11). This is used by the high level system to sample control loop errors for a smooth turn on at the end of the slip stacking cycle. The second RF output is given a positive offset frequency (1200 Hz) before injection to minimize destructive interference on the first batch of injected beam. The first batch of beam is injected synchronous with the first RF output. Both RF outputs are then decelerated until the second RF output has zero offset frequency, and the first output has a negative frequency offset (-1200 Hz). The second batch of beam is injected, and both outputs are accelerated until the frequency offsets on the two outputs are equal and opposite (+-600 Hz) (Figure 3.1.12). This is done so that the two batches are captured by a frequency with a net zero offset. Otherwise, the system must maintain the offset frequency after capture and into acceleration. Just before the batches are aligned, the frequencies are brought closer together in order to reduce the total bucked separation at capture time. Once the two batches are aligned, the frequency offsets are set to zero.

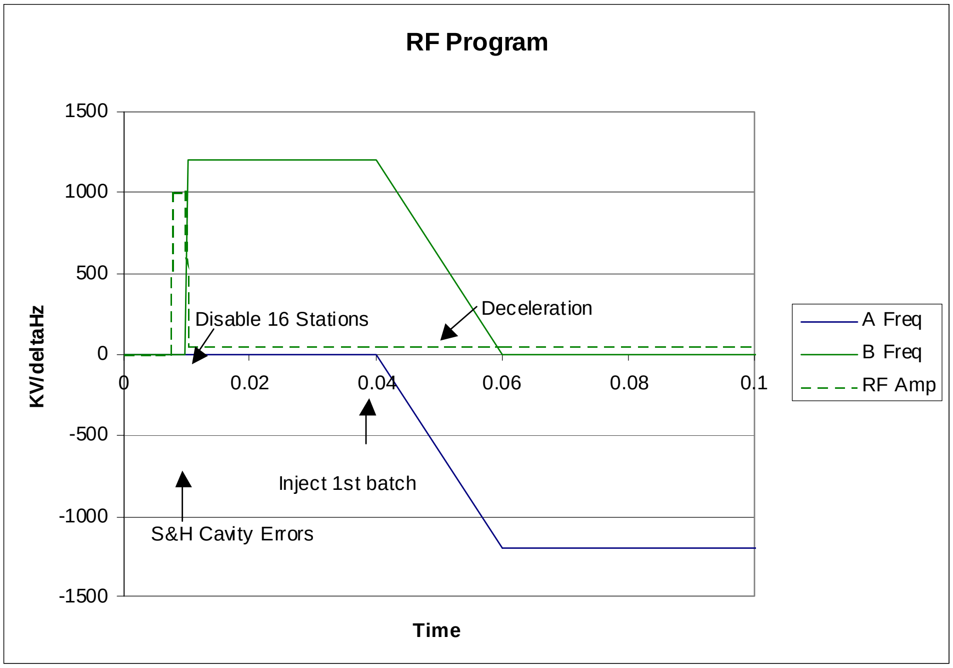

Figure 3.1.11 *Frequency offset program and RF drive level program for slip stacking. The error signals on the cavity feedback loops are sampled before the cavities are gated off. The first batch is injected and decelerated. The frequency difference between 'A' and 'B' outputs is kept constant.*

The next step in the process involves enabling the LLRF feedback loops and beginning the normal single batch acceleration process to 120 GeV. This has not been tried yet in preliminary tests. It appears to be possible to re-enable the loops after capture. The coalescing process already disables and enables feedback loops without noticeable effects on the beam quality.

Once the offset frequency signals are generated, it is relatively straight forward to generate the signals in the proper RF cavities. There are still some minor modifications to the high level system necessary to facilitate slip stacking in the Main Injector.

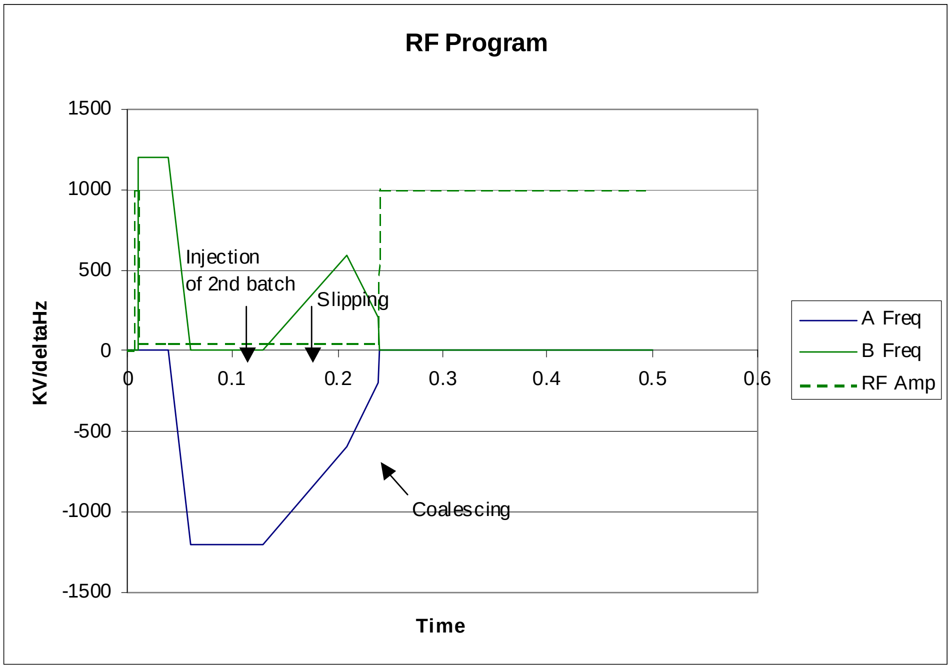

Figure 3.1.12 *Frequency offset program and RF drive level program for slip stacking. After the second batch is injected, the beam is accelerated so that the energy offsets are symmetric about the nominal orbit. Just before the batches are combined, the energy difference is reduced. The frequency offsets are set to zero, and the RF drive is increased simultaneously to hold both batches at capture.*

3.1.2.2 HLRF

The high level RF system is defined as all of the hardware used to get the signal from the LLRF distribution to the cavity gap. This includes the cavity itself and all RF power amplifiers downstream of the LLRF system (Figure 3.1.13)[8]. The cavity is a ferrite bias tuned resonant structure. The resonant frequency of the cavity varies from 52.8 MHz to 53.1 MHz from low to high ferrite bias current. The Q of the resonant cavity varies from 3500 to 5000, and the shunt impedance varies from 250 kOhms to 520 kOhms. The cavity has a 12:1 step up ratio from the RF power input to the gap. The primary power limit to the cavity comes from the ferrite tuners. The tuners cannot hold off a voltage greater than 9 kV without sparking.[9]

The cavity is driven by a Y567 tetrode power tube. This tube is capable of delivering about 75 kW of CW RF power. The plate of the power tube is biased by another Y567 tube in series, and this tube holds off the 30 kV DC bias from the power supply. The grid of the series tube controls the amount of bias on the plate of the power tube. A solid state power amplifier drives the cathode of the power tube with about 4 kW of power capability. The power tube also has a fixed screen voltage supply, and a programmable grid bias supply.[8]

There are four feedback loops around each cavity. One loop maintains the proper resonant frequency of the cavity through control of the ferrite bias current. The loop keeps the anode and cathode voltages in phase for maximum power efficiency and fundamental beam loading compensation. Another loop maintains the proper RF voltage on the cavity by measuring the DC screen current. If the screen current gets too large, the low level RF drive to the cavity is reduced. Likewise, if the screen current gets too low, the drive is increased so that the RF waveform fills the entire anode bias. This means that during normal operation, the anode bias dictates the RF voltage in the cavity. The signal from the cavity gap monitor is also fed back into the cavity drive, reducing the fundamental cavity impedance seen by the beam by a factor of 10. Finally, there is a fanout/fanback phase lock loop that compares the gap monitor voltage phase with the LLRF drive. This loop helps maintain a very precise phase relationship between different cavities for better paraphase operations.

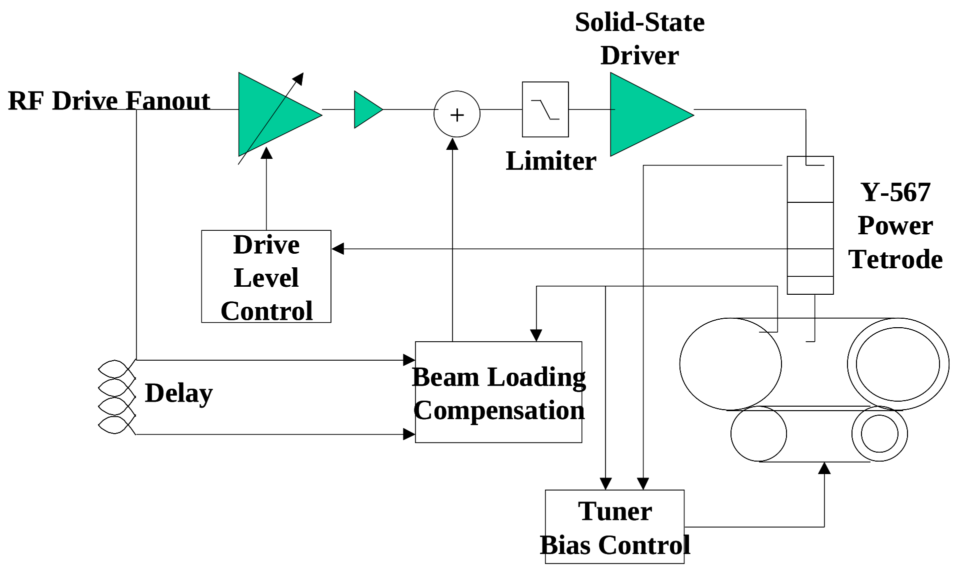

Figure 3.1.13 *Block diagram of Main Injector high level RF system. The drive level control receives its error signal from the DC screen current. Not pictured, the fanout-fanback phase loop.*

The modifications to the high level RF system necessary to facilitate slip stacking are a product of having to operate at a low voltage and still maintain beam loading compensation systems. Because of the way the power tube is biased, there is a limit to how low the plate can be biased. Therefore, the plate bias is kept high and only two stations are activated during the slip stacking process. Even though only two stations are active, the other sixteen stations must keep their RF feedback loops active. Instead of turning off the tubes completely, only the RF drive is disabled on the sixteen cavities. This keeps the tube biased and the feedback loop active during slip stacking. Special gates were added to the RF drive input to the amplifiers, so that single cavity drives could be activated without disabling the tube bias.

Because there is no RF voltage in most of the cavities during slip stacking, many of the feedback loops around the stations will not regulate properly. These feedback loops have relatively large settling times, and if the errors are allowed to float to a rail, the beam will be adversely affected while the loop tries to settle again. In order to insure that the cavities regulate quickly after they're turned on, the errors from the feedback loops are sampled before the cavities are shut down. The errors are held during the slip stacking process and allowed to track again after the stations have been reenergized.

A preliminary RF voltage program has been tested on the Main Injector. First, all of the RF stations are active, and their feedback loop errors are sampled and held. Before the first batch is injected, all but one RF station is gated off. Just before the second batch is injected, the station with the RF synchronous with the second batch is turned on. All of the stations are activated precisely at capture time.

3.1.3 Low Intensity Slip Stacking

In order to better understand the basic mechanics of slip stacking, studies have been performed at intensities below the level where beam loading effects begin to dominate. These studies were performed in simulation and in the Main Injector itself.

3.1.3.1 Simulations

The simulations provide valuable information about idealized operating conditions for optimal slip stacking. These operating parameters would be difficult to determine empirically, because the operating conditions of the Main Injector change frequently during the course of studies. Without the aid of the simulation results, problems with slip stacking due to Main Injector setup errors would be much more difficult to diagnose.

"ESME is the program used to simulate longitudinal manipulations. The program ESME has been developed to model those aspects of beam behavior in a proton synchrotron that are governed by the radio frequency systems. It follows the evolution of a distribution in energy-azimuth coordinates turn-by-turn by iterating a map corresponding to the single-particle equations of motion. The map parameters may be updated each turn to reflect the action of the beam current on the individual particles through feedback loops, space charge, coupling impedance, etc. Over eighteen years of development it has been applied to a significant range of problems in rf capture, transition crossing dynamics, bunch coalescing, longitudinal single bunch and multi bunch stability. The standard output includes practically all of the information available from single particle dynamics like, for example, bucket area, bunch emittance, synchrotron frequency, slip factor, etc. Both input and output for collective motion calculations may generally be represented in frequency domain or time domain, as convenient in a particular application." [10]

There has been no need to modify the distribution version of ESME to do the slip stacking simulations; however, this does not mean that it has been unnecessary to write code. ESME is written with ten dummy entries for attaching application specific code. For the slip stacking application, three of these entry points are used. One subroutine, which has been around for a long time and more or less taken for granted by regular users, is used to make 166 copies of a 0.1 eVs bunch and place them at the desired bunch centers for the two batches (Figure 3.1.14). A second very short routine is used to test the location of marker particles located at the stable fixed points of the center bunches of each batch and to stop tracking when they are aligned in phase. These two routines are all that is required to find the emittance dilution and optimum frequency curves including performance degradation by beam loading. However only simplistic conclusions with respect to compensation techniques can be established without additional code.[11]

One of the first results derived from the simulation is the importance of energy separation between the batches while they are slipping (Figure 3.1.15) (Figure 3.1.16). Both bunch trains experience the full effect of both RF systems. At sufficient frequency separation, the extra frequency excitation averages to zero in a small fraction of a synchrotron oscillation period. At some energy separation, the phase motion of the two batches will be practically independent. For example, if the buckets overlap 50% in energy, the single particle motion is chaotic everywhere within the buckets. Tangential buckets define a lower limit for stable motion, but the simulations show rapid emittance growth. Emittance growth is not entirely absent at much larger separations, but simulations show an acceptably small amount of growth when the buckets are separated by four bucket-heights in energy. This corresponds to a space for a completely empty bucket between the upper and lower buckets.[11]

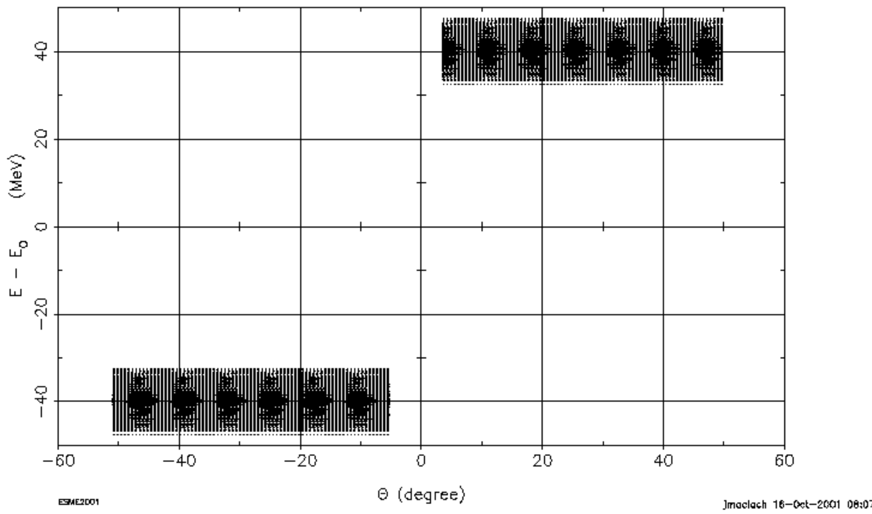

Figure 3.1.14 *Initial condition setup for slip stacking simulation. Two batches of 83 bunches each are injected with offset energies. Q represents the azimuthal position of the beam in the Main Injector.*

Parameter	Symbol	Value	Units
Mean reference orbit radius	R_O	528.30	M
Synchronous Energy	E_s	8938.28	MeV
Transition energy/$m_o c^2$	γ_T	18.6	
RF peak voltage, each system	V_o	150	kV
RF peak voltage at closest approach	V_o	85	kV
RF harmonic	h	588	
Shunt resistance of 20 RF cavities	R_{shunt}	2	MΩ
Loaded Q of cavities	Q	2000	
Synchrotron tune (150 kV)	ν_s	0.003	
Bucket height	H_B	12.9	MeV
Stationary bucket area	S_B		eV-s
Phase space area occupied by simulated particles	S_b		eV-s

Table 3.1.2 *Accelerator and beam parameters used in slip stacking simulation*[11]

The capture of two bunches into a single bucket will produce gross emittance dilution unless the bunches can be brought much closer together than four bucket-heights on center. It is possible to accelerate the two batches into each other just before capture. Although chaos does develop, it does take time. If the batches are accelerated quickly enough, there will be a net gain in overall captured emittance. Simulations show a better final emittance when the bunches are captured closer together in energy with an acceleration angle and slightly off angle with respect to RF phase alignment (Figure 3.1.17).

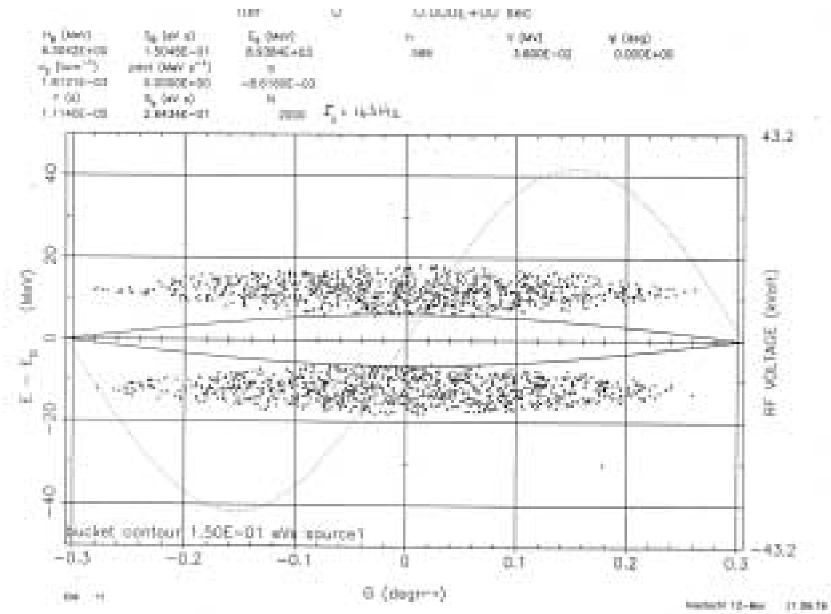

Figure 3.1.15 *This initial condition shows the two batches have already started to slip. Longitudinal emittance of each bunch is 0.1 eV-s.*

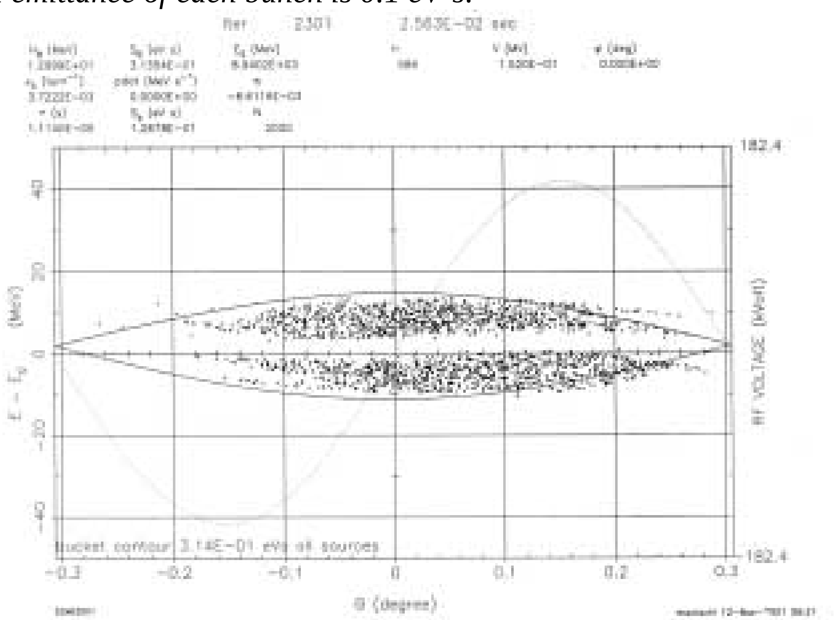

Figure 3.1.16 *This shows the beam profile at capture time. The longitudinal emittance of each bunch has not diluted significantly, and the final emittance is about three times the initial bunch emittance. These results do not include beam loading.*

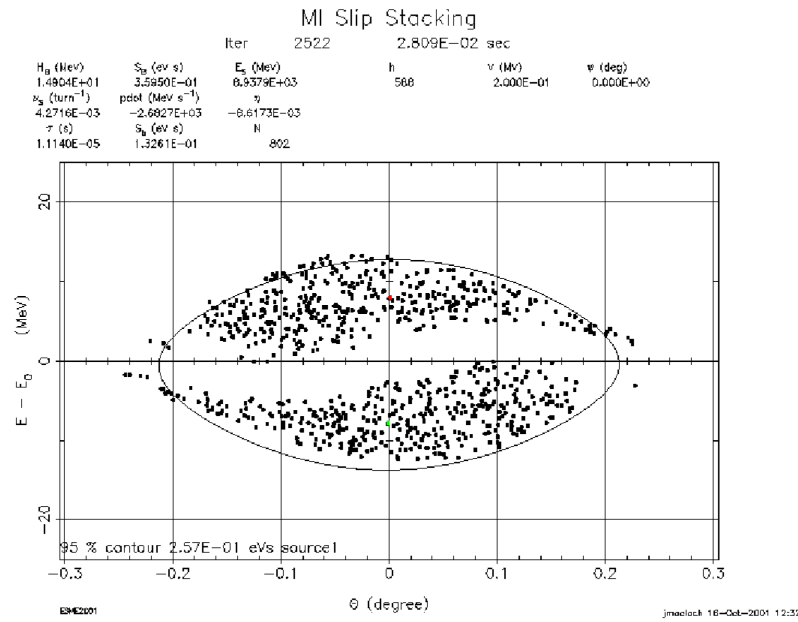

Figure 3.1.17 *Simulation of capture including reducing energy offset before capture. This improves final emittance by about 17%. Does not include effects of beam loading.*

3.1.3.2 Beam Studies

With many of the necessary hardware changes in place, it is possible to test most of the slip stacking process. In order to compare real beam emittance growth with simulations, it is important to eliminate injection mismatches. Because the slip stacking bucket height is so much smaller than the normal stacking bucket height, the Booster must rotate its bunches prior to extraction for proper bucket matching into the Main Injector (Figure 3.1.18, Figure 3.1.19).

Parameter	Symbol	Value	Units
RF voltage, each system	V_o	62.5	kV
Synchrotron frequency	f_s	219	Hz
Frequency separation	Δf	1200	Hz
Bucket area, each system	A_o	0.2	eV*s
Initial longitudinal emittance, 0.7e12 intensity	S_o	~0.1	eV*s

Table 3.1.3 *RF parameters for Main Injector slip stacking tests*

The Booster is also equipped with a paraphasing system for the purpose of eliminating the fundamental RF voltage during multi-turn injection from the Linac. The system is also utilized for bunch rotation just before extraction. The trigger time and paraphase angle are adjusted so that the beam performs a quarter synchrotron rotation prior to extraction. This technique works well at lower Booster intensities. At the present time, instabilities in the Booster are increasing the longitudinal emittance beyond the acceptance of a slip stacking bucket. These instabilities must be remedied in order to perform high intensity slip stacking (Figure 3.1.20).

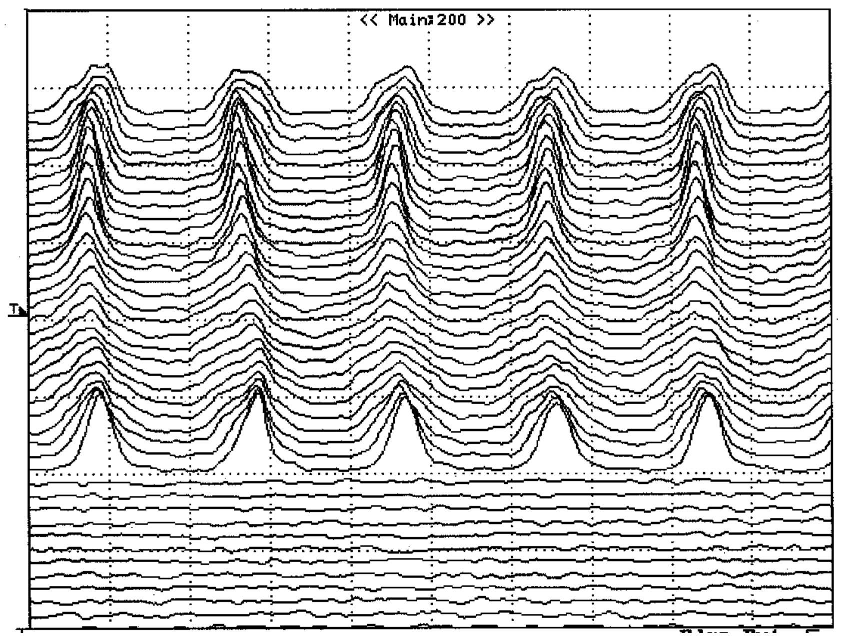

Figure 3.1.18 *Longitudinal pickup response of injection into the Main Injector slip stacking bucket without Booster bunch rotation as a function of time. Successive traces show the signal developing through the injection process (mountain range plot). There is a large quadrupole oscillation present. Intensity is 0.7e12 protons/batch.*

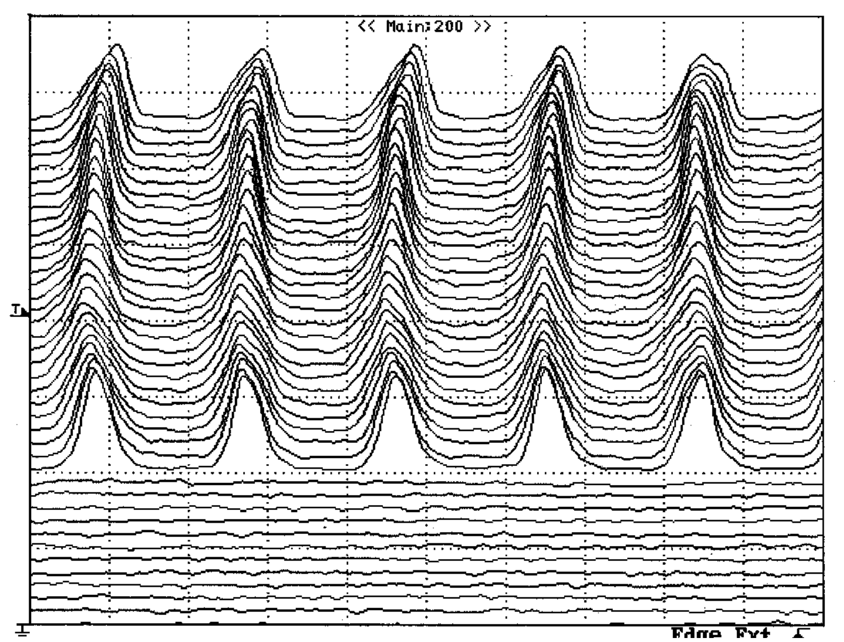

Figure 3.1.19 *Injection into Main Injector slip stacking bucket with bunch rotation in Booster tuned up. Intensity is 0.7e12 protons/batch.*

Once bunch rotation is tuned up in the Booster, the injection parameters for the first batch are tuned in the Main Injector. For this tuning, only the RF station that is synchronous with the first batch is enabled. There is a significant amount of RF leakage through all of the stations even when their drives are disabled, a mountain range plot shows the effect of the offset frequency leakage (Figure 3.1.21). To tune the injection parameters, all stations are set to the same frequency, making it easier to diagnose injection mismatch. By enabling a second batch synchronous RF station, one can test the bucket separation effects on beam emittance without the effect of beam loading (Figure 3.1.22) (Figure 3.1.23) (Figure 3.1.24).

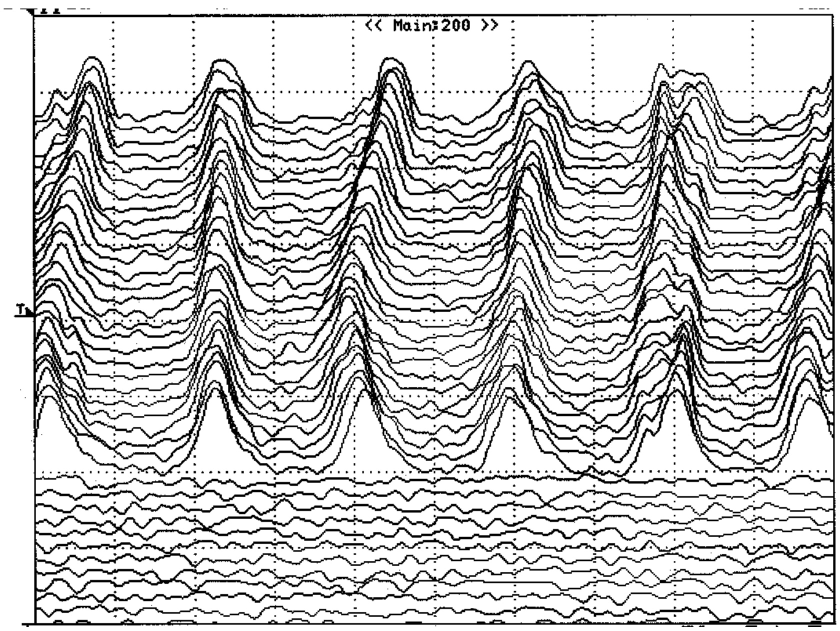

Figure 3.1.20 *Injection into Main Injector slip stacking bucket with bunch rotation in Booster tuned up. Intensity is 1.5e12 protons/batch. Problems with bunch energy errors are already evident.*

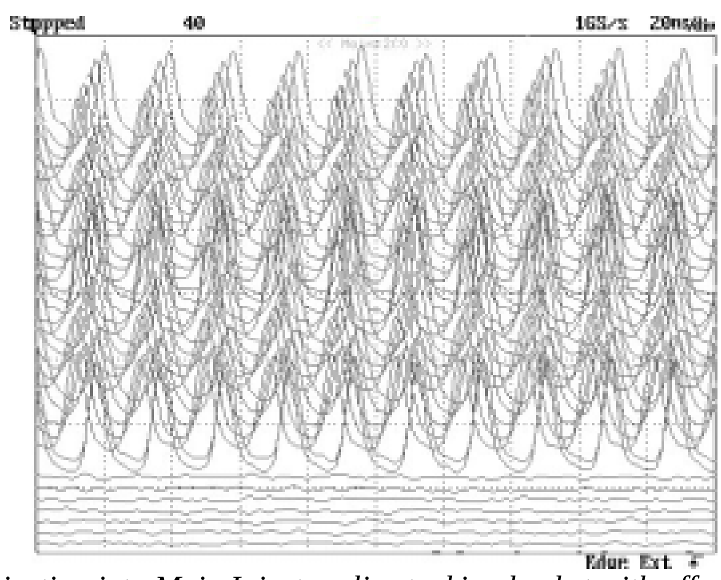

Figure 3.1.21 *Injection into Main Injector slip stacking bucket with offset frequency output from LLRF but offset stations gated off. The effect of RF leakage is apparent.*

Tuning the second batch is a bit more complicated than tuning the first batch. The bucket offset and phase offset are dictated by the integral of the frequency offset curve for the second batch. Both offset frequency ramps are programmed to follow the second batch ramp, and only the second batch RF station is enabled. The second batch is injected at the time in the cycle when it would normally be injected, and the injection phase and bucket offset are tuned for minimal injection mismatch and batch separation respectively.

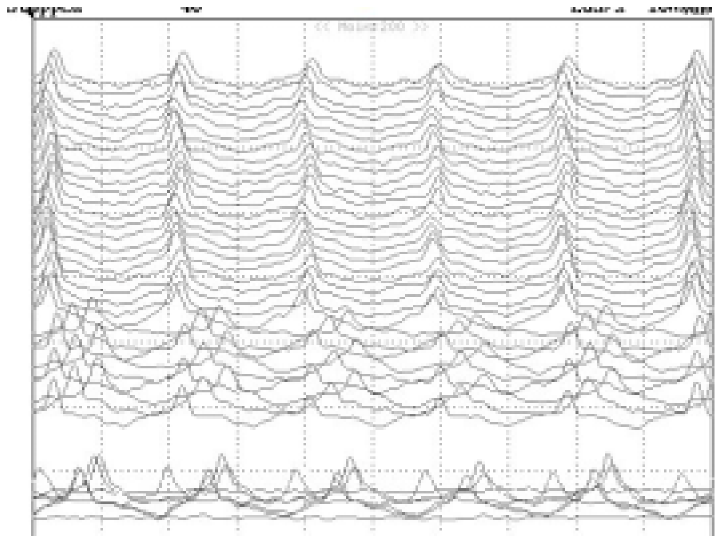

Figure 3.1.22 *Emittance preservation of first batch with 600Hz frequency separation. This separation is not enough to keep the buckets from overlapping.*

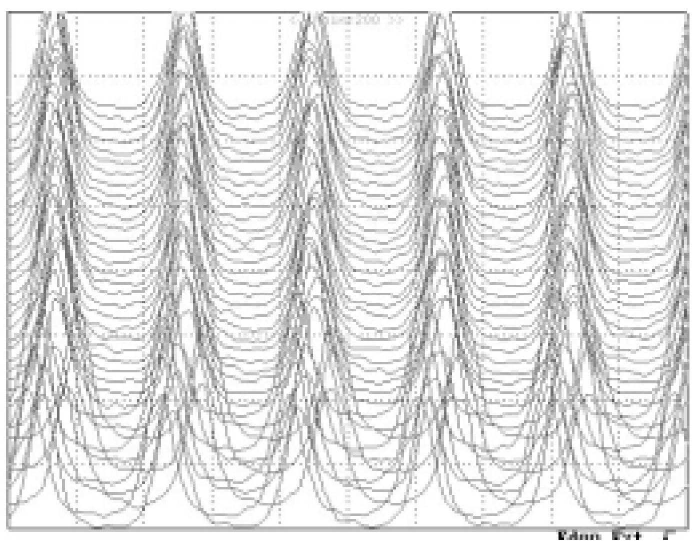

Figure 3.1.23 *Emittance preservation of first batch with 1200Hz frequency separation. The energy separation between slipping buckets is slightly greater than four times the RF bucket height.*

After the second batch injection is tuned, it is straight-forward to test the entire low energy process. In the first attempt with the Main Injector, the frequencies were kept at a constant separation during the slipping process. Some time slightly before capture, the frequencies were ramped to a smaller separation to reduce the total emittance occupied by the two batches prior to capture. At capture time, the frequency offsets were set to zero, and all of the RF cavities were enabled (Figure 3.1.25) (Figure 3.1.*26*) (Figure 3.1.*27*) (Figure 3.1.28) (Figure 3.1.29).

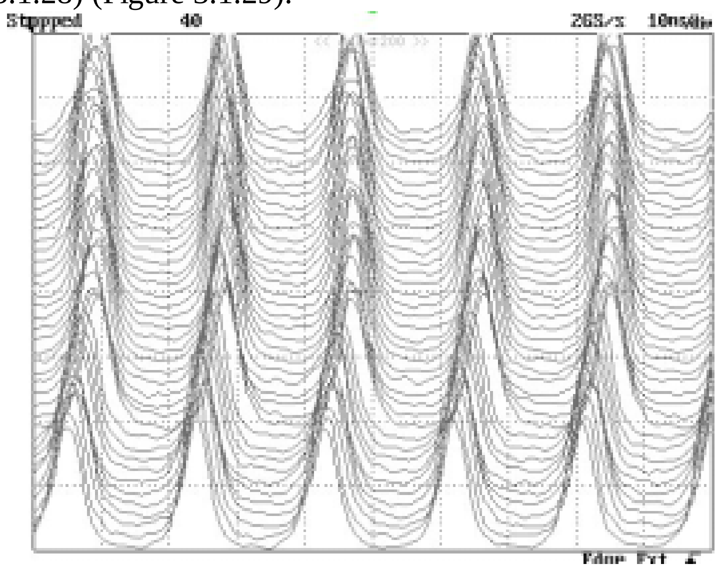

Figure 3.1.24 *Emittance preservation of first batch with 1800Hz frequency separation. This separation is well above having a whole bucket separation in energy between the slipping batches, although not much improvement over 1200Hz separation.*

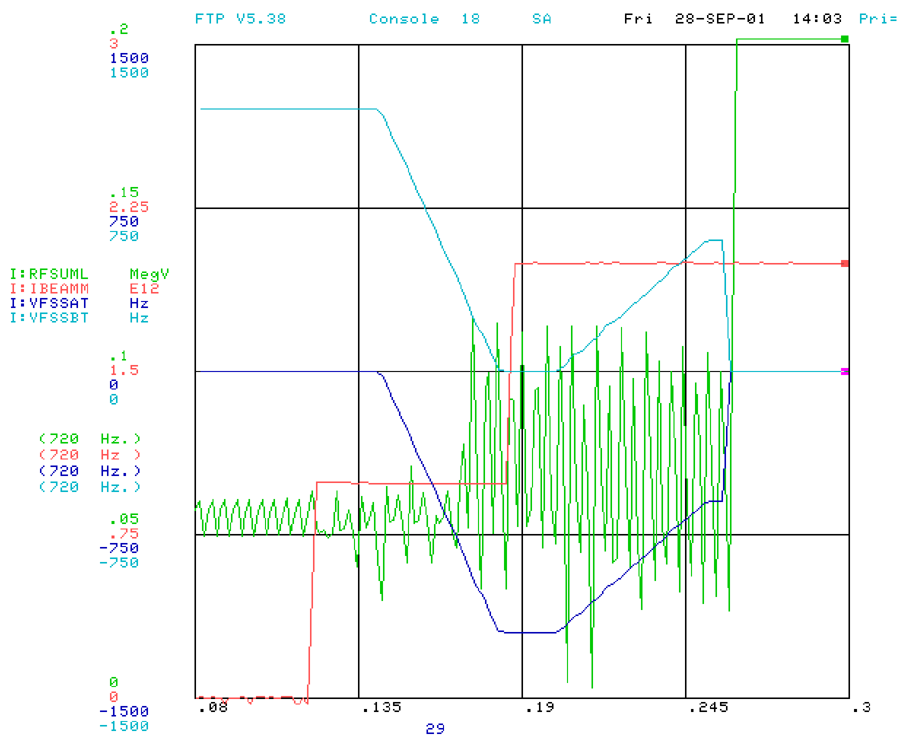

Figure 3.1.25 *Plot of Main Injector slip stacking cycle. I:RFSUML is the total RF voltage seen by the beam. Notice the large modulation when the 'B' RF is activated. I:IBEAMM is the total beam intensity in the Main Injector. I:VFSSAT and I:VFSSBT are the offset frequencies of the 'A' RF and 'B' RF respectively.*

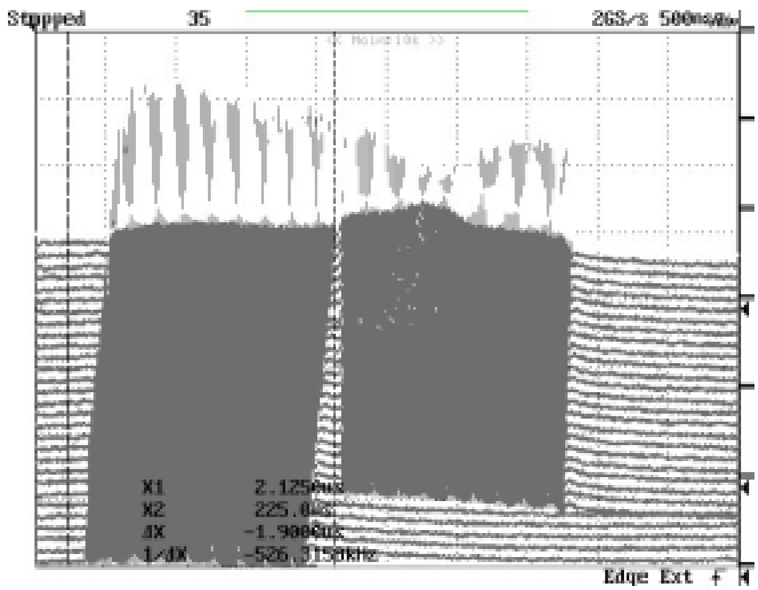

Figure 3.1.*26 Mountain range plot showing Main Injector slip stacking just after the second batch is injected*

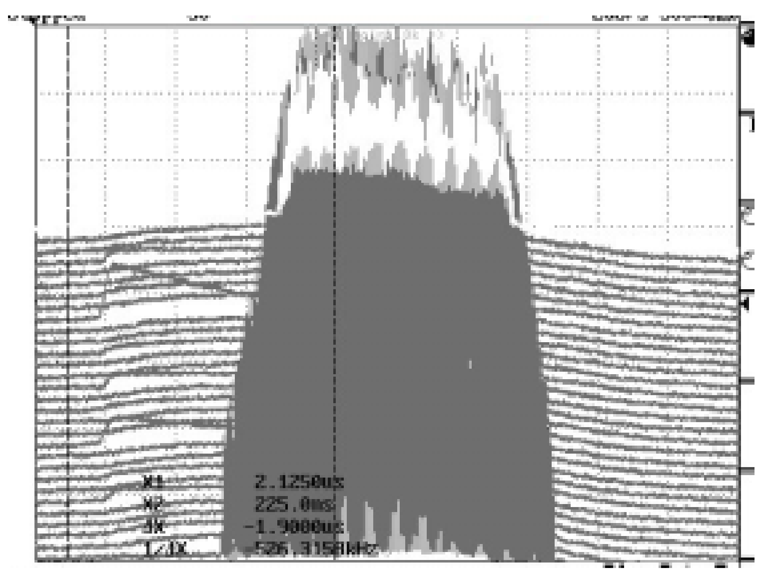

Figure 3.1.*27 Mountain range plot showing Main Injector slip stacking when the two batches are captured into a single batch*

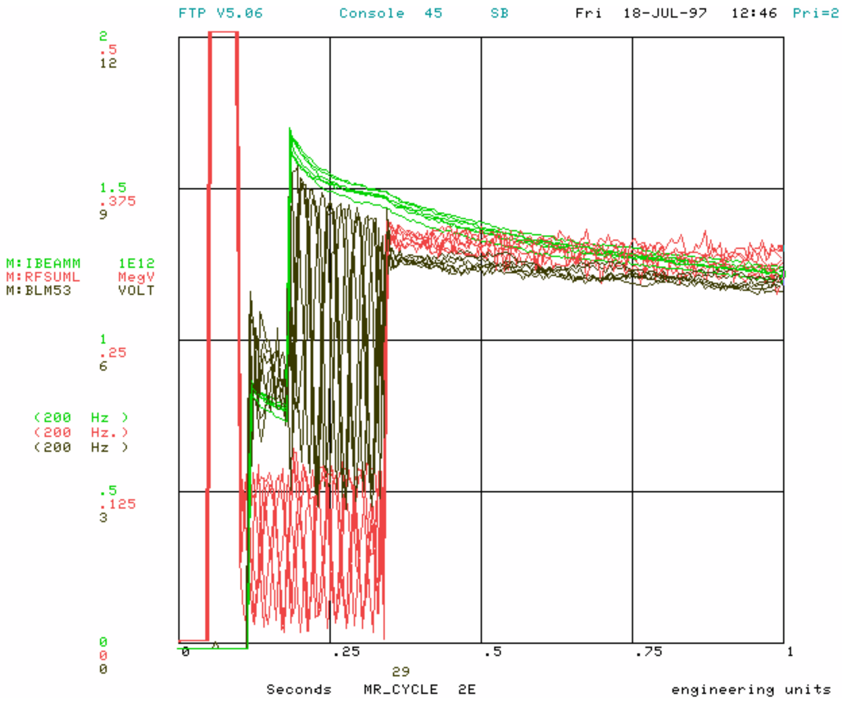

Figure 3.1.28 *Plot of main ring slip stacking cycle. M:RFSUML is the total RF voltage seen by the beam. M:IBEAMM is the total beam intensity in the main ring. M:BLM53 represents the 53 MHz component of beam current. This gives a measure of the amount of beam still left in the bucket.*

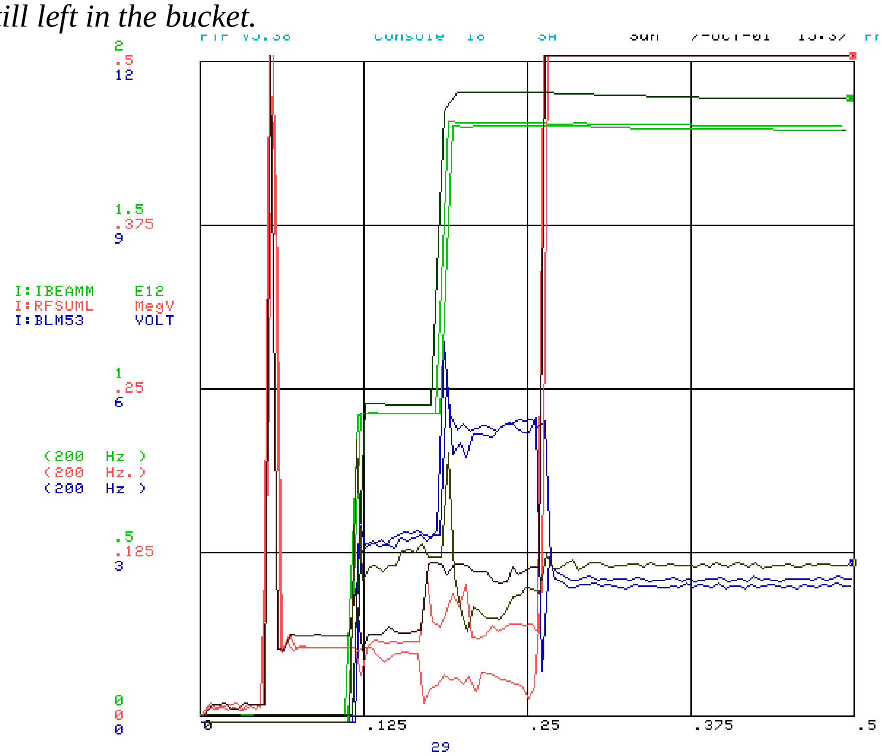

Figure 3.1.29 *Plot of Main Injector slip stacking cycle.*

3.1.4 Beam Loading

3.1.4.1 Beam Loading Theory

The greatest limitation to practical slip staking intensities is beam loading in the RF cavities. There are many aspects of beam loading that must be addressed in order to satisfactorily accelerate high intensity beam with low voltages. First, the RF cavities must provide enough average power to contain the beam in the bucket. This not only includes maintaining the proper bucket height, it also includes maintaining enough extra voltage for focusing in the presence of beam loading. In the absence of any feedback on the beam or cavity, this aspect is referred to as a Robinson criterion.[12] In general, this criterion states that the total power delivered to the cavity by the beam current must be about equal to or less than the power delivered to the cavity by the power amplifiers.

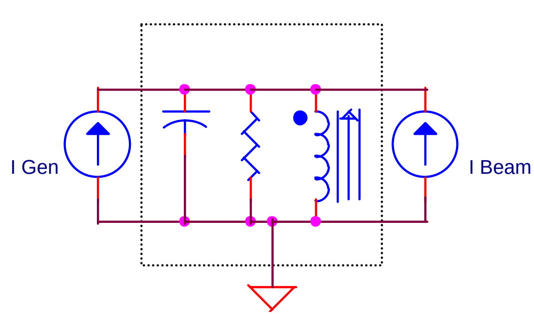

Figure 3.1.30 *Cavity block diagram. Total current (It) is IGen + IBeam.*

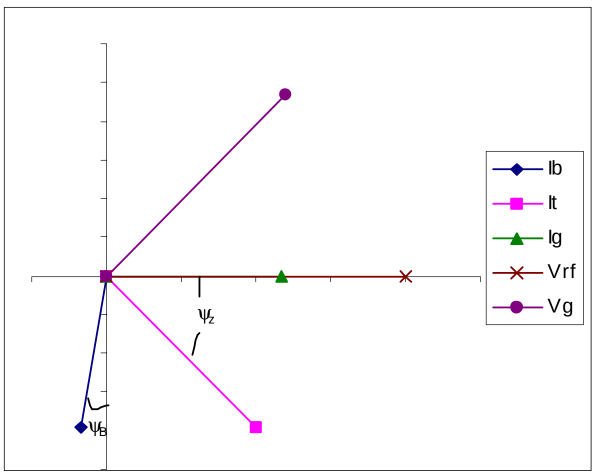

Figure 3.1.31 Vector sum of IGen and IBeam with cavity detuned such that IGen and Vrf are in phase. ψ_z is the detuning angle

To understand this, one can view the cavity as a resonant circuit that is driven by two current sources (Figure 3.1.30).[12] The two current sources represent the drive from the power amplifier (Ig) and the beam itself (Ib). These sources are out of phase by 90° for a stationary bucket, and they are out of phase by 90° plus the synchronous phase angle for an accelerating bucket. For optimum power efficiency, the amplifiers should see a purely resistive load. Thus, the resonant frequency of the cavity is detuned until the total voltage is in phase with the current generated by the power amplifier (*Figure 3.1.31*).

A Robinson instability occurs when the phase of the generator voltage (Vg), which is defined as the voltage in the cavity without beam, is 180° out of phase with the beam current (Figure 3.1.32). At this point, beam with a small synchronous phase error will not receive the necessary focusing power needed to bring it back into the bucket. The exact definition of the threshold for this instability is given in Eq. (3.1.1), where ψ_L is the load angle between Vrf and Ig. The condition for $\sin(2\psi_L)>0$ is discussed later in the introduction.

$$0 < \sin(2\psi_Z) < 2Y\cos(\psi_B)$$
$$Y = \frac{I_B}{\mathrm{Re}(I_T)} \qquad (3.1.1)$$

The threshold defined above assumes that no other feedback loops act on the beam and cavity voltage. Usually, feedback loops will control the load angle, the synchronous phase, and the cavity voltage amplitude. Instabilities develop in these systems at large beam intensities because the different loops begin to couple.[13] For example, a change in synchronous phase for an intense beam will also cause a change in the load angle, and a change in the voltage amplitude. These problems have been analyzed, and the conditions for stability are more complex. However, in most cases, the loops still remain stable as long as the power generated in the cavity by the beam does not exceed the power generated by the power amplifier.[13]

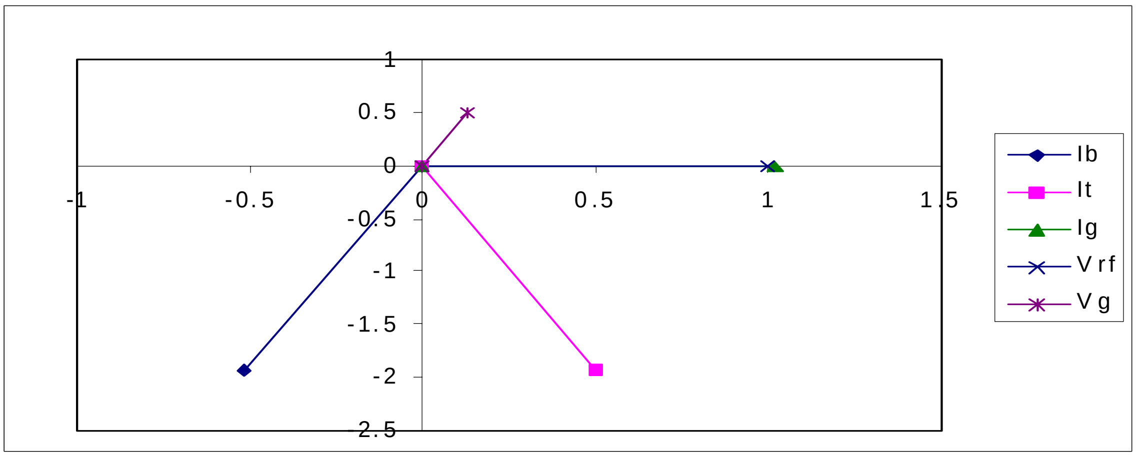

Figure 3.1.32 *Vector diagram of a cavity at the limit of the Robinson instability. Vg and Ib are at opposite phases.*

Of course situations arise in which the amount of beam in an accelerator is far beyond the stability limits explained above. It may not be possible for the power amplifiers to provide enough voltage to maintain the stability, and it may not be desirable to have such a large bucket size.

Another aspect of beam loading is referred to as transient beam loading. As a bunch of beam first passes through the RF cavity, it loads the cavity down by an amount proportional to the intensity of the bunch. The next bunch through sees a smaller RF voltage than the first bunch, and the last bunch of the batch sees the smallest voltage of all. Such a distribution in amplitudes could lead to large differences in synchronous phase and cause emittance blowup. For slip stacking, the transient voltage in the cavity varies by as much as 40 kV.[14] This is 40% of the total voltage on a batch and will certainly cause adverse effects, as the simulations will show. Also, during slip stacking, as the two batches of beam are slipping past each other, the fundamental component of the beam current is 100% AM modulated with the difference in frequencies. This implies that there are components of transient beam loading within the fundamental resonance of the cavity.

3.1.4.2 **Beam Loading Simulation**

Simulations reveal the adverse affects of beam loading in the slip stacking process. The simulations use the cavity parameters such as Q, shunt impedance, and resonant frequency to generate the appropriate wake fields generated by the beam. (Figure 3.1.33) (Figure 3.1.34) (Figure 3.1.35) (Figure 3.1.36)

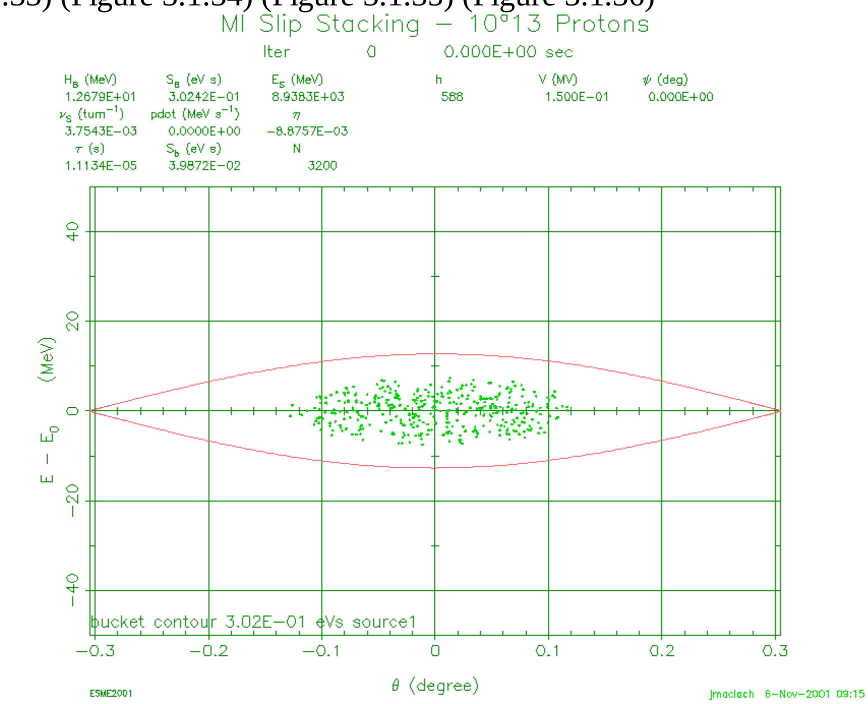

Figure 3.1.33 *Initial condition setup for a single bunch in a batch for slip stacking including beam loading. Initial longitudinal emittance is 0.1 eV-s.*

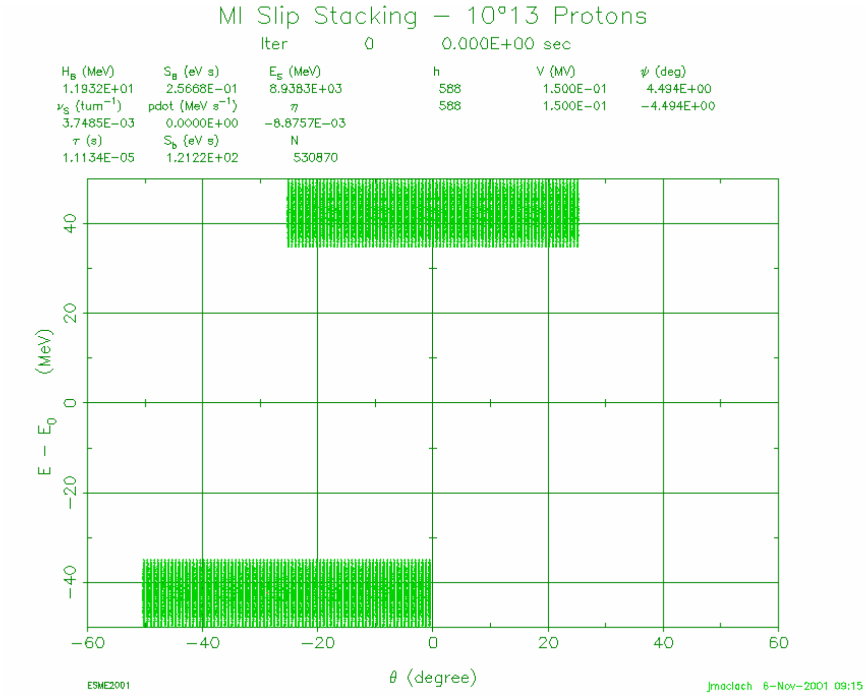

Figure 3.1.34 *Initial condition setup for both batches of beam for slip stacking including beam loading. Intensity is 5e12 protons/batch. θ represents azimuthal position around the Main Injector.*

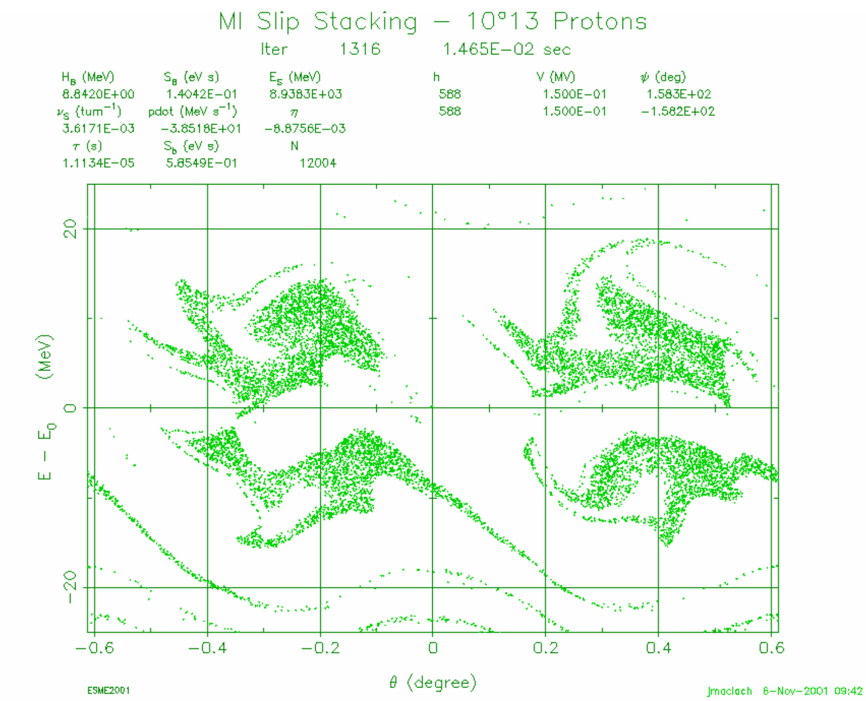

Figure 3.1.35 *Final profile for captured beam for two bunches in a batch for slip stacking including beam loading with no compensation.*

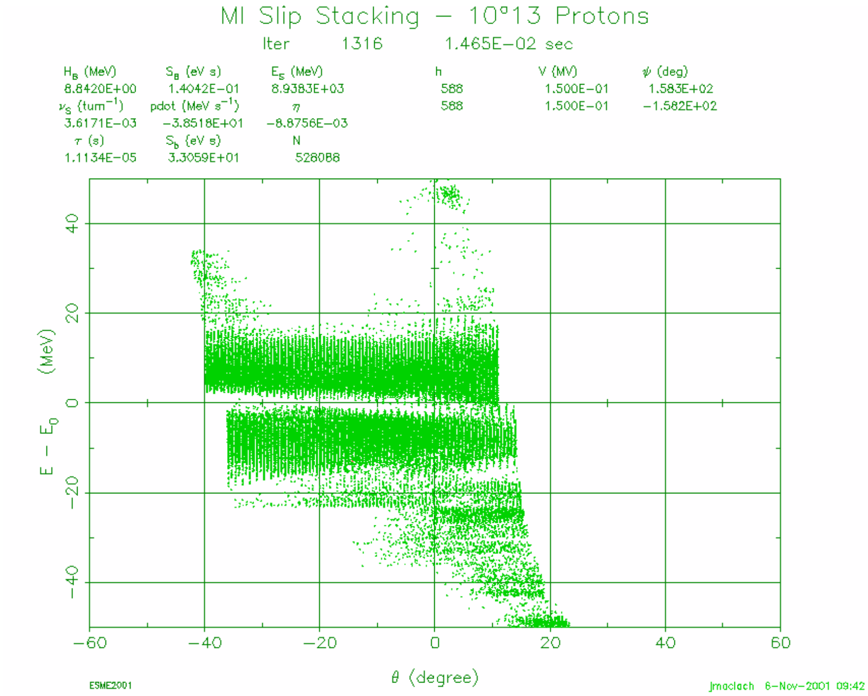

Figure 3.1.36 *Final profile for captured beam for both batches for slip stacking including beam loading with no compensation.*

3.1.5 Beam Loading Compensation

In order to overcome the limitations on slip stacking imposed by beam loading, some method of reducing the effect beam loading must be applied to the cavities. One possibility for reducing the effect of the beam in the cavity is by decreasing the cavity impedance. This will also require more power from the RF power amplifier, negating the benefit of the low acceleration power required. Another possibility is to virtually reduce the cavity impedance through feedback. By applying the beam generated voltage to the input of the power amplifiers, the cavity impedance seen by the beam is reduced without requiring more power from the amplifiers. However, implementation of a feedback system requires careful design to ensure there is enough stability margin in the loop itself.

3.1.5.1 Beam Loading Compensation Theory

One of the most common methods of beam loading compensation is direct RF feedback. In this method, a signal derived from the actual cavity voltage is fed back into the cavity drive (Figure 3.1.37). This reduces the effective impedance seen by the beam and the fanout drive as show in Equation (3.1.2). For stability purposes, the beam reacts dynamically as if it was under the influence of a much higher voltage. This new, effective voltage is proportionally greater than the actual drive voltage by the open loop gain of the feedback loop (Figure 3.1.38) (Figure 3.1.39). If the feedback loop has an ideal response with no delay, the effective shunt impedance and Q of the cavity are reduced by the open loop gain.

$$Z^* = \frac{Z}{1+ZGS} \qquad (3.1.2)$$

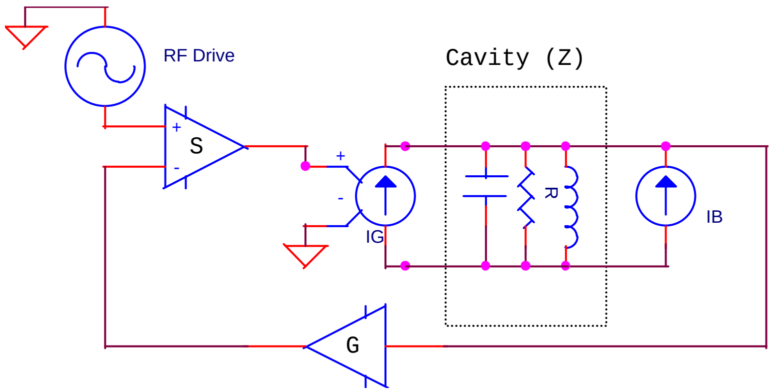

Figure 3.1.37 *Block diagram of direct RF feedback. The feedback gain is G, and the transconductance is S.*

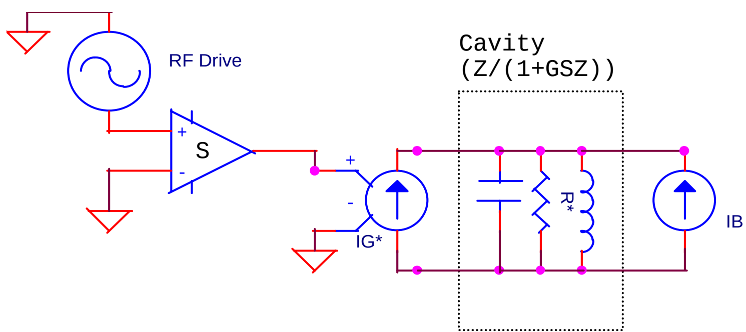

Figure 3.1.38 *Block diagram of equivalent circuit as seen by the beam and the RF drive.*[14]

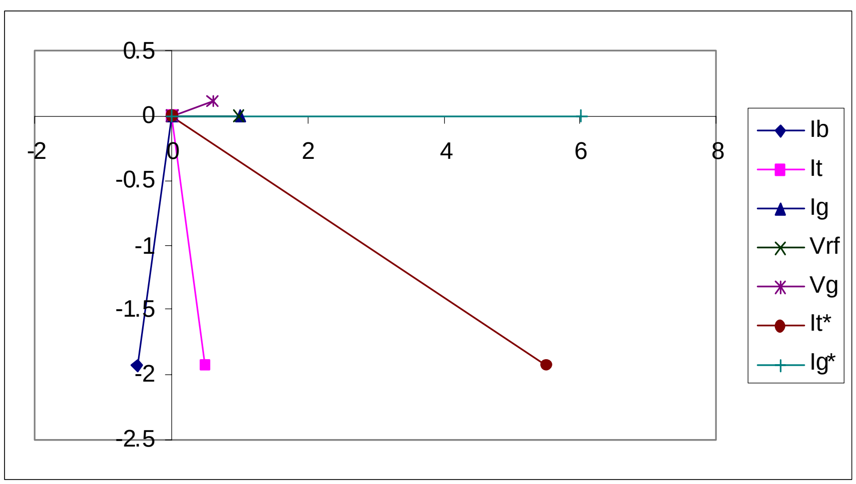

Figure 3.1.39 *Vector diagram of cavity with feedback compensation. The loop gain (GS) is 20dB. Ig* and It* are the virtual currents seen in the equivalent circuit.*

Of course it is impossible to physically construct a zero delay system, especially when the summing point and output are separated by 100 feet of cable. All physical feedback systems have stability limits on the amount of open loop gain. The Nyquist stability criterion states that a system will be stable as long as the complex mapping of its open loop response does not encircle the (1,0) point. For most of the feedback systems in the Main Injector, this criterion is over specified. The open loop response of the system is not allowed to have the real part of its response exceed one. This limits the amount of gain that the direct RF feedback system can achieve without driving the cavity unstable as shown in Equation (3.1.3).

$$G\max = \frac{\pi Q}{2\omega_r \tau}$$
$$R\min = \frac{R}{G\max} \qquad (3.1.3)$$

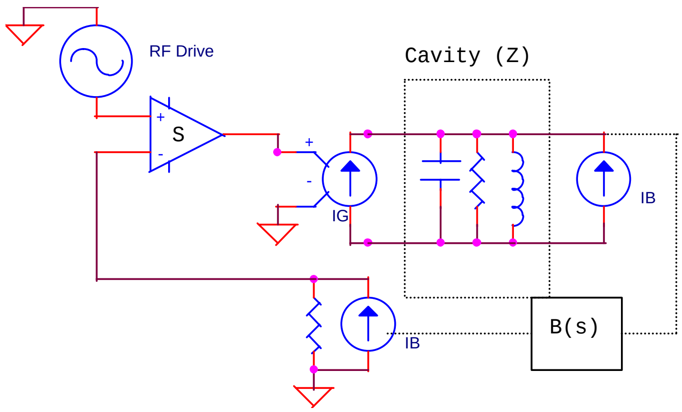

Figure 3.1.40 *Feedforward block diagram where S is the transconductance of the amplifier and B(s) is the beam response*

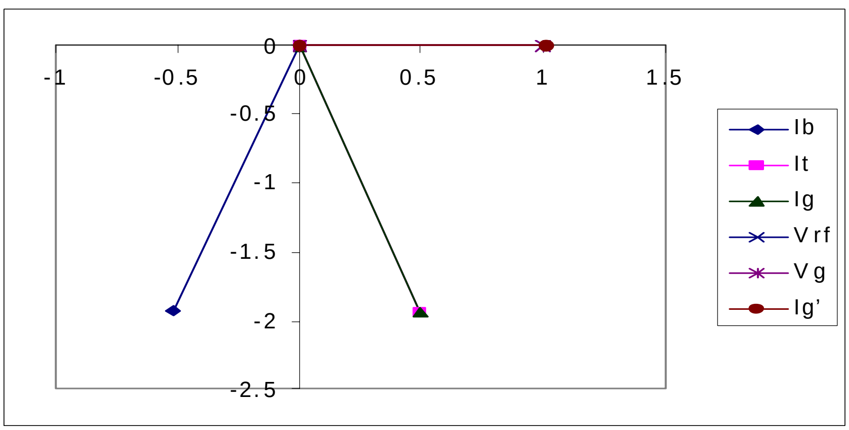

Figure 3.1.41 *Vector diagram of a cavity with ideal feedforward compensation. The generator current and total current are always the same. This puts the generator voltage always in phase with the RF voltage and no Robinson instability can occur.*

Another method of beam loading compensation is called feedforward compensation. In this method, a low impedance beam current detector creates the signal that is applied to the RF drive. The amplification of the signal from the detector to the cavity is set so that the power amplifier in the cavity produces a current that exactly opposes the beam current (Figure 3.1.40). If the match is perfect, the cavity voltage becomes completely decoupled from the beam current (Figure 3.1.41).[15] Because matching is so critical in this system, it is essential that the signal corresponding to a particular bunch from the detector be applied to the same bunch as it comes around in the cavity. In a highly relativistic system, it is impossible to outrun the beam with signal propagation. Therefore, the signal must be fed back on the bunch after approximately one revolution.

In the feedforward system, the amount of compensation achieved is not limited by Nyquist stability limits like the direct RF feedback system. The limits of the feedforward system are determined by the ability to match the beam current with the power amplifier. Many factors can affect this matching, the worst being changes in the gain of the power amplifier chain and the non-linearity of the power tube transconductance. This greatly limits the robustness of the feedforward system and makes it more difficult to engineer than the direct RF feedback system.[16] Also, there is still a weak feedback loop involving the beam response to changes in the cavity voltage. The beam could be driven unstable longitudinally without proper design of the feedforward response.

3.1.5.2 Beam Loading Compensation Simulation

Some simulations have been performed which describe the effect of beam loading compensation on slip stacking. These simulations show the effect of an ideal direct RF feedback system. They show how the slip stacking would respond to a cavity with reduced shunt impedance and Q, which implies a zero delay system (Figure 3.1.42) (Figure 3.1.*43*) (Figure 3.1.44) (Figure 3.1.45). Some work still needs to be done in order to simulate more physical systems.

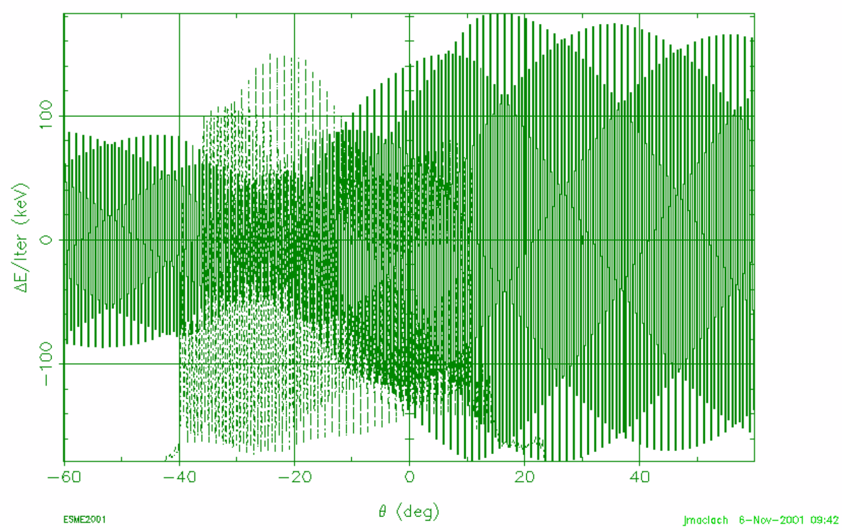

Figure 3.1.42 *Inverse FFT of beam current and cavity voltage in slip stacking simulation with no beam loading compensation. Intensity is 10^{13} protons. Time in cycle is close to capture time.*

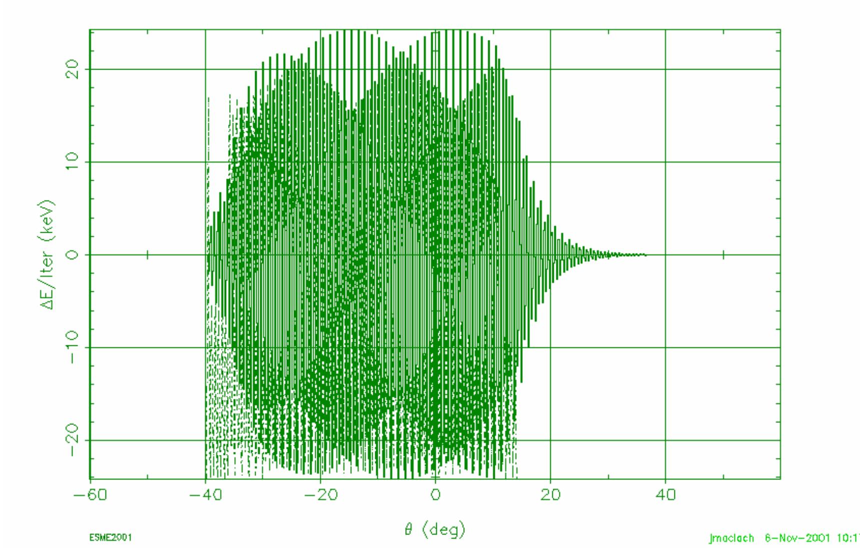

Figure 3.1.*43 Inverse FFT of beam current and cavity voltage in slip stacking simulation with ideal direct RF feedback with 40dB of loop gain. Intensity is 10^{13} protons. Time in cycle is close to capture time.*

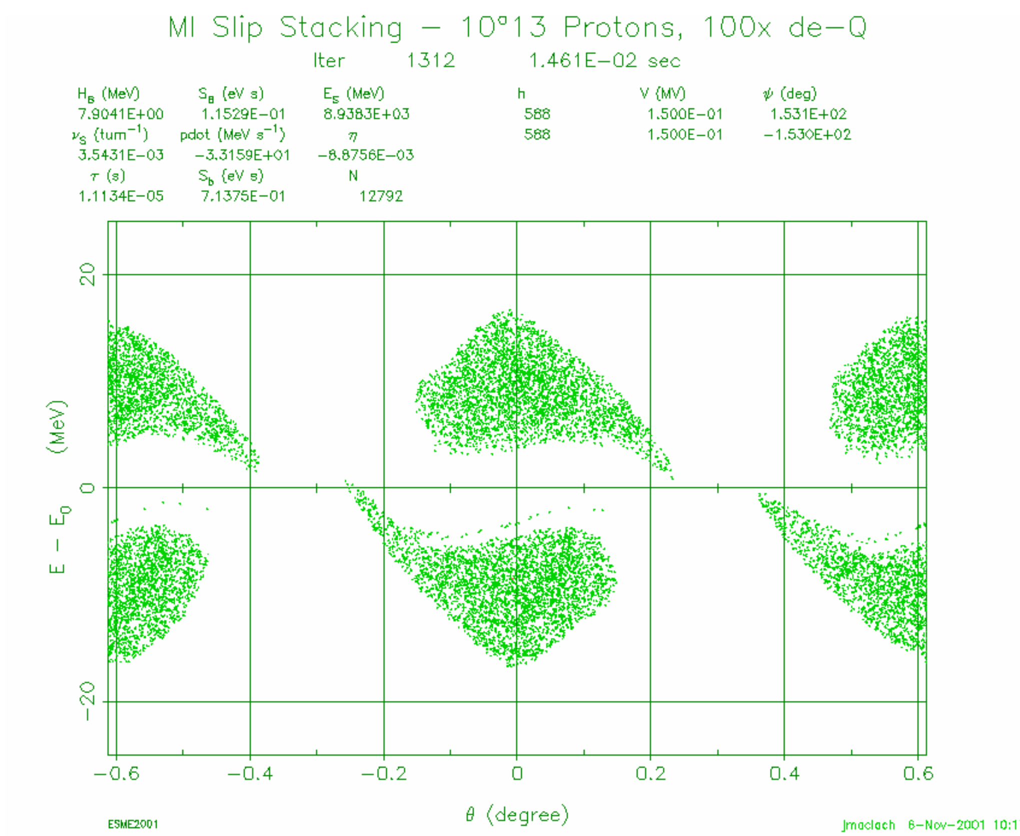

Figure 3.1.44 *Final profile for captured beam for two bunches in a batch for slip stacking including an ideal direct RF feedback system with 40dB of loop gain. Intensity is 10^{13} protons*

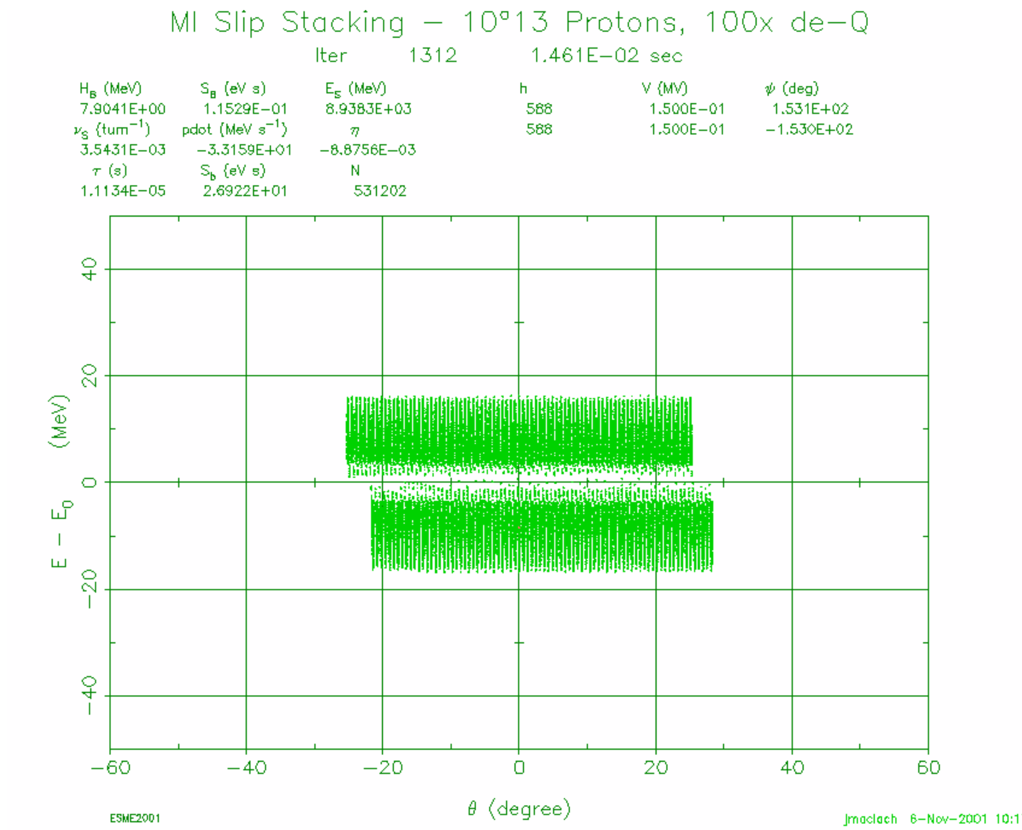

Figure 3.1.45 *Resulting profile for captured beam for both batches for slip stacking including beam loading with an ideal direct RF feedback system with 40dB of loop gain. Intensity is* 10^{13} *protons*

3.1.5.3 **Direct RF Feedback**

The Main Injector is equipped with a direct RF feedback system (Figure 3.1.46).[8] Each cavity has an independent feedback system. The system consists of a module that converts the signal from the cavity gap monitor to baseband (Figure 3.1.47). The signal is low-pass filtered, up-converted, and combined with the fundamental amplifier drive signal. It is important that the phase of the open loop response remain 180° at the fundamental frequency for maximum stability margin. Because the system works through the entire Main Injector ramp, the feedback system must maintain the proper phase intercept as the fundamental frequency changes. There is enough fixed delay in the loop to cause a 60° error in the phase intercept over the frequency ramp without compensation. The system maintains the proper phase by using different delays for the up-convert and down-convert RF references in the feedback module. The up-convert reference delay is matched to the LLRF fanout delay to the cavity, and the down-convert delay is matched to the cavity gap signal from the tunnel. With the proper delays on the references, the feedback module will adjust its delay to maintain the proper phase intercept for the system.

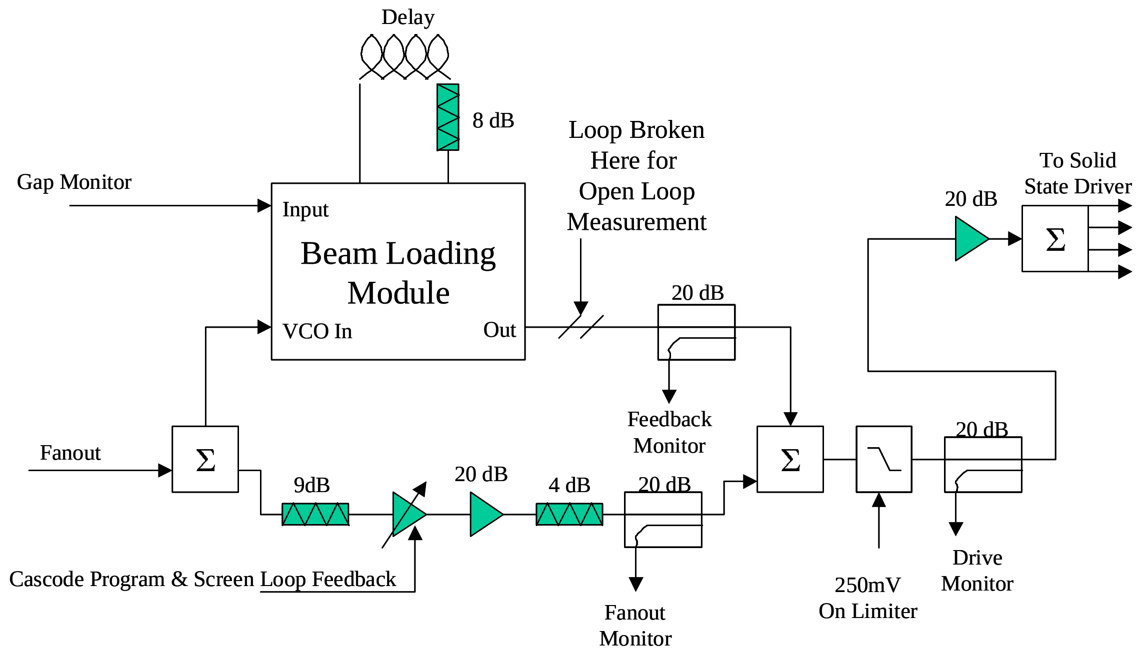

Figure 3.1.46 *Direct RF feedback system in the Main Injecto*r

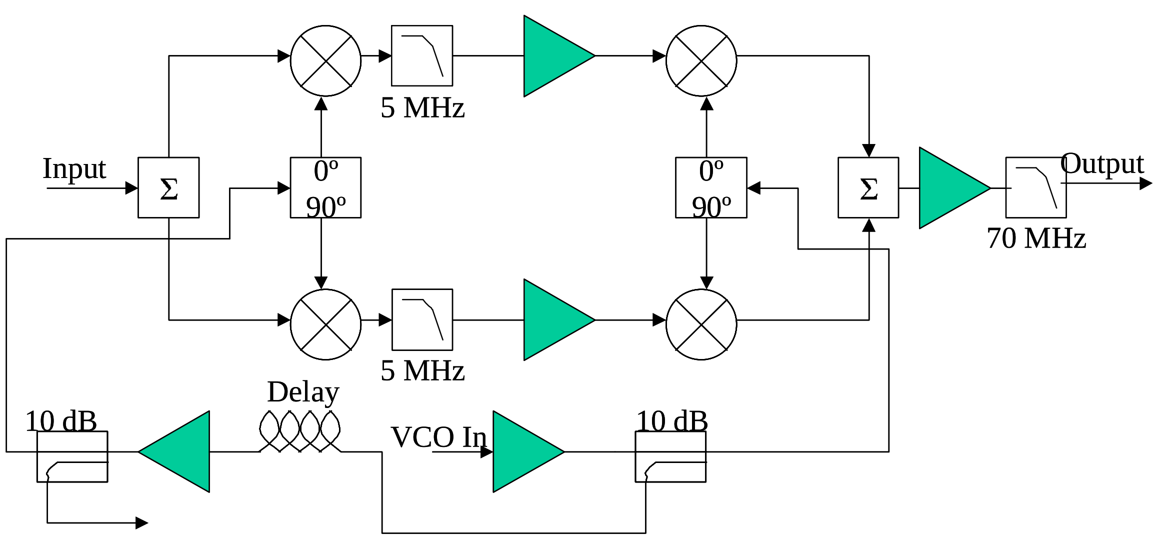

Figure 3.1.47 *Block diagram of the beam loading module. The superheterodyne structure is designed to track the phase response with the changing VCO frequency. The downconvert reference is synchronized with the cavity gap signal, and the upconvert reference in synchronized with the fanout.*

Maintaining proper phase intercept improves the stability margin, but there is still a stability limit on the allowable open loop gain on the system. The current Main Injector system will only allow an open loop gain of about 26 dB with a reasonable gain margin (Figure 3.1.48) (Figure 3.1.49). The feedback module has a fixed gain profile, in frequency, over many revolution harmonics, so the cavity response dictates the open loop bandwidth. Because of the high Q of the cavity, the open loop gain of the system rolls off quickly. Thus, the system performs insignificant beam loading compensation at any revolution harmonics other than the fundamental.

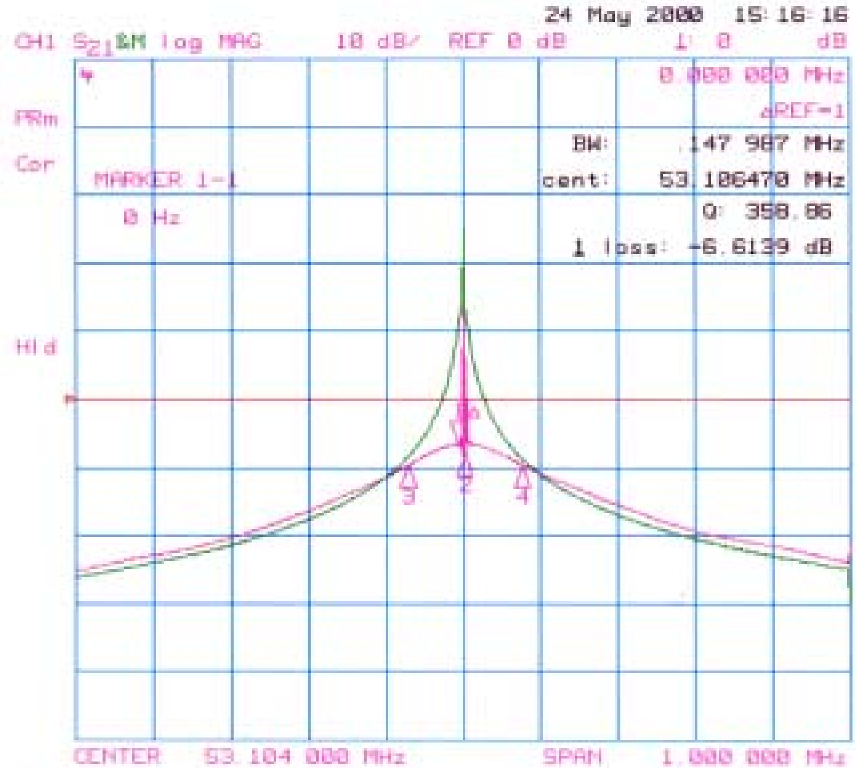

Figure 3.1.48 *Cavity response with 20dB of direct RF feedback beam loading compensation (pink) and without (green).*

In order to provide beam loading compensation at the revolution harmonics, the feedback module must be modified. First, to ensure the proper open loop phase intercept for multiple revolution harmonics, the system must have a delay equal to some multiple of the revolution period. Second, the bandwidth of the filter should not be dictated by the cavity, since this is not optimal for stability. The bandwidth of the system could be reduced to the point of just containing the frequency difference between the two batches in a slip stacking cycle. Of course the filter would necessarily have the same shape around the fundamental frequency as well as multiple revolution lines. This implies some kind of digital filter sampling at the fundamental frequency with taps at multiples of the revolution frequency.[17]

A digital filter for beam loading compensation is already being designed. The current design uses a DSP with a highly parallel architecture, clocked at a multiple of the fundamental frequency (Figure 3.1.50). The down-converted signal from the cavity gap is digitized and stored in FIFO memory blocks. Data from the memory blocks are burst into the DSP, and the DSP performs the filtering calculations. Output data from the DSP is burst into another set of FIFO memory blocks which drive a DAC. The FIFO memory blocks maintain the system delay at one revolution period. The output of the DAC is up-converted and combined with the cavity fanout drive. Calculations done for an IIR filter in the DSP show that open loop gains on the order of 40dB are achievable (Figure 3.1.51) (Figure 3.1.52). To maximize the gain margin for the revolution harmonics, the signals for the revolution harmonics will follow a different path than the fundamental, so that they can receive a 90° phase shift. This is to compensate for the cavity response, which is reactive at the revolution harmonics.

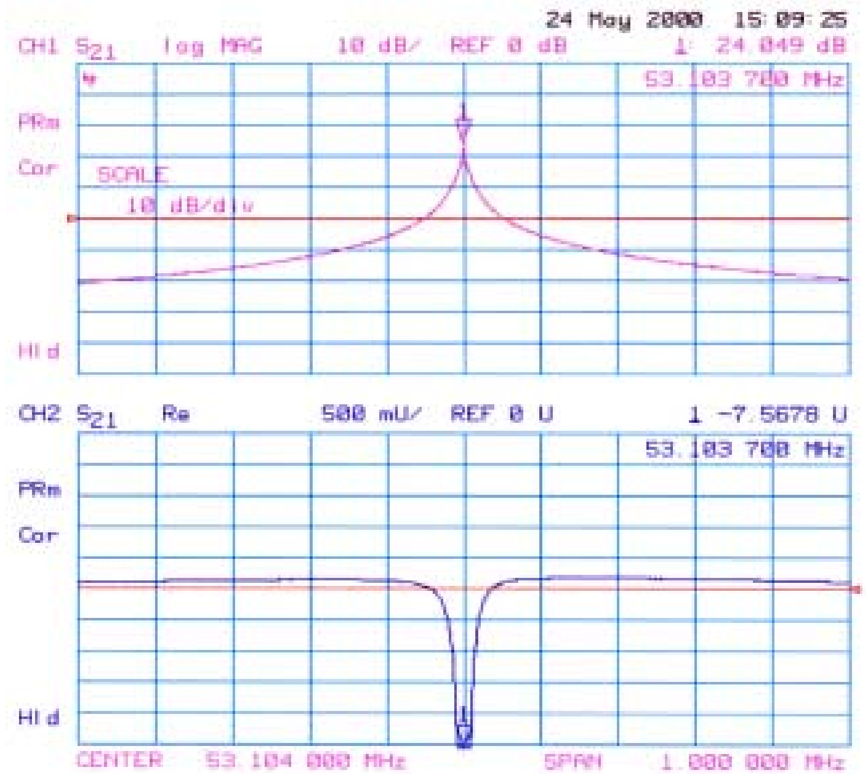

Figure 3.1.49 *Cavity response with 20dB of direct RF feedback beam loading compensation. Notice that the real part of the response is well below +1.*

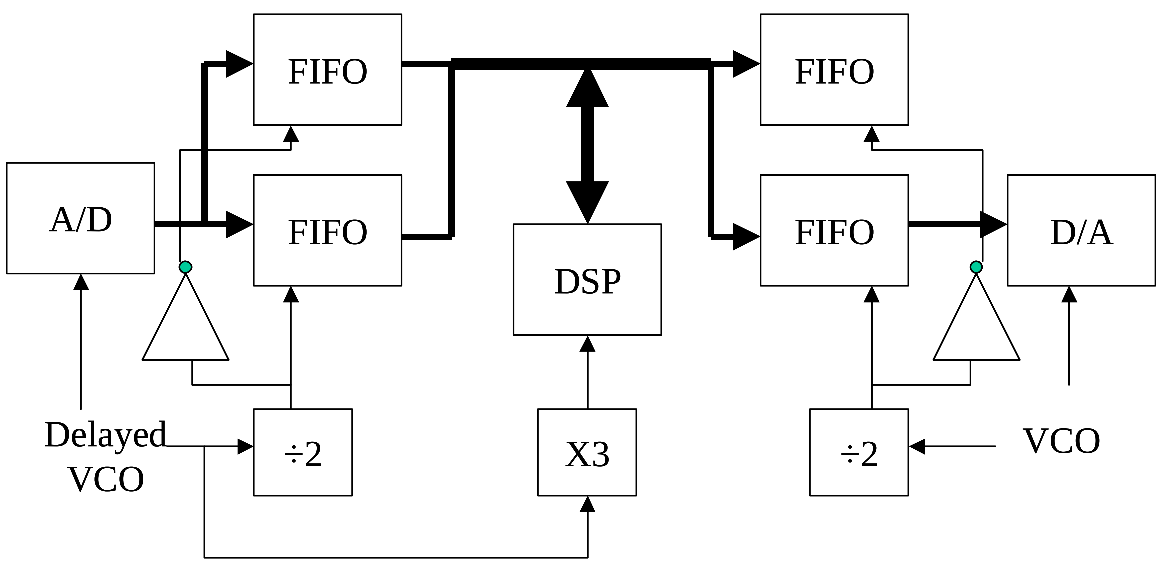

Figure 3.1.50 *Block diagram of digital direct RF feedback system. The DSP calculates the filter response and bursts data to and from the FIFOs. The digitizer and DAC are clocked with references at differing delays to track the changing revolution period of the Main Injector.*

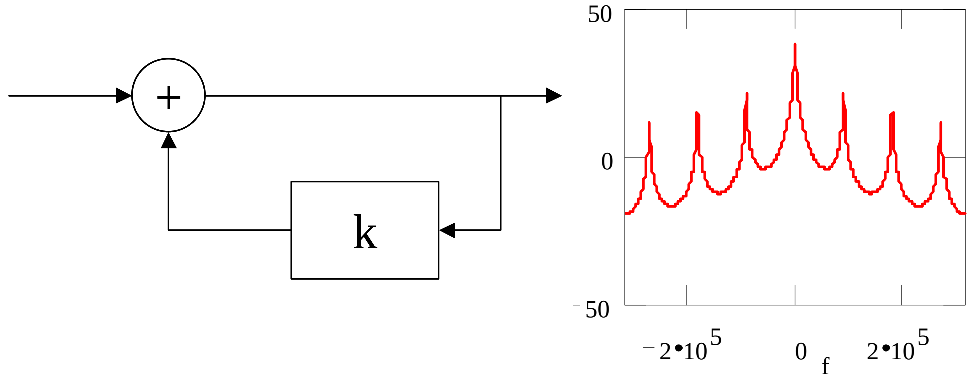

Figure 3.1.51 *Program for IIR comb filter in the DSP and simulated open loop response of system through cavity when* $k = 1 - 2^{4}$

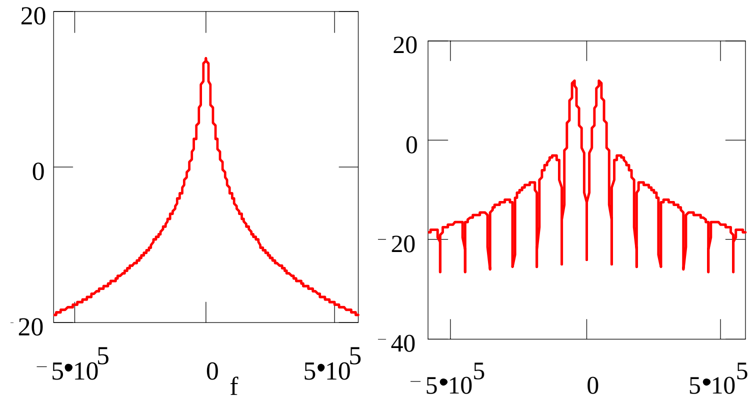

Figure 3.1.52 *Sample of cavity response without beamloading compensation and sample of cavity response with digital direct RF feedback.*

There are three operating conditions that could stress the RF power limitations of the high level system with beam loading compensation. First, fast changes in the RF fanout drive will cause the feedback system to overdrive the power amplifier. A RF fundamental feedback system forces more current into the cavity when there are transients in the drive. This is how the feedback system makes the impedance and Q of the cavity appear lower to the fanout. Unfortunately, transients that occur at moderately high RF voltage levels will attempt to drive the power amplifiers well beyond there power limits (i.e. transition crossing). To compensate for this, a limiter is placed downstream of the feedback summing junction. Although the limiter does prohibit effective beam loading compensation at transition crossing, the RF voltage is very high at this point in the cycle, minimizing beam loading effects.

The second RF power limitation of the fundamental feedback system is the drive power required at the summing junction. The gain from the fanout drive to the cavity is reduced by the feedback gain. If the loop gain is high enough, high power RF amplifiers would be necessary to drive the low level summing junction of the cavity. This is very costly when multiplied by 18 stations, and it limits the loop gain in the current Main Injector direct RF feedback system. Currently, the loop gain in the feedback system is not time variable, and the amount of power required to drive the summing junction at peak RF power output pushes the limits of the medium power drive amplifiers upstream of the summing junction. The 20dB limit on the loop gain could be improved with higher power RF amplifiers or variable gain feedback loops that can increase the loop gain when less (or zero) cavity voltage is required.

The third and most prohibitive RF power limitation observed on the direct RF feedback system is the current drive required for effective compensation. When the fundamental frequency is the only component compensated by the feedback system (and the cavity is allowed to tune itself for beam loading), the power amplifiers are not required to supply more power, except for dynamic changes to the fundamental component of the beam or the fanout drive. Compensation of frequency components other than the fundamental (or even the fundamental if the cavities are not tuned to compensate for beam loading), such as revolution harmonics, will require that the power amplifier have enough available power to match the current in the other frequency components. There is very little power dissipated in the cavity at these frequency components, because the impedance of the cavity is usually very low at these frequencies, and the feedback causes the impedance to look even smaller. The critical factor is the available current in the power amplifier circuit.

As mentioned earlier, the Main Injector RF power amplifiers utilize a cathode follower circuit on the power tube. This means that the solid state drivers must provide all of the current required by the power tube. The two specifications that determine the power amplifier's ability to provide effective compensation are the power tube's maximum current capacity and the solid state drivers maximum power output. The current capacity of the tube should be enough to handle the slip stacking load. Although the peak current of the beam during slip stacking exceeds the tube's maximum average current rating, this rating is a thermal rating and can be scaled with duty factor. The solid state driver also has enough power to handle the slip stacking current.[9]

3.1.5.4 Feedforward

Some tests of the power amplifier chain's ability to handle the current requirements of beam loading compensation have already begun. To test the amplifier's response to transient beam loading, a feedforward system is used instead of a digital feedback system. Constructing a compensation system that is stable and operates on revolution harmonics is easier with feedforward. The feedforward system currently being tested in the Main Injector uses a wall current monitor for its beam current source. The signal from the wall current monitor is down-converted, filtered, and delayed digitally (Figure 3.1.53). The output of the digital delay drives a special cavity fanout system. Instead of having a system of equal length cables, this fanout system is designed to have delays different by the beam transit time between cavities. At each of the cavities, the signal is upconverted and combined with the drive (Figure 3.1.54).

The disadvantage of the feedforward system is the beam signal and power amplifier current must match very closely. There is no inherent correction mechanism like there is in a feedback system. Therefore, the system can only operate in very well defined conditions. It cannot track energy changes or changes in RF amplitude or operating conditions. Also, it is extremely important that the signal path be completely linear, otherwise the feedforward signal will be too distorted to cancel out the beam signal when the two meet in the cavity. The power tube is a major source of nonlinearity in the signal path. It is so bad that, without compensation, only very specific beam profiles will work with the feedforward system. Regular stacking cycles are the only cycles where significant feedforward compensation is observed in the cavities (Figure 3.1.55).

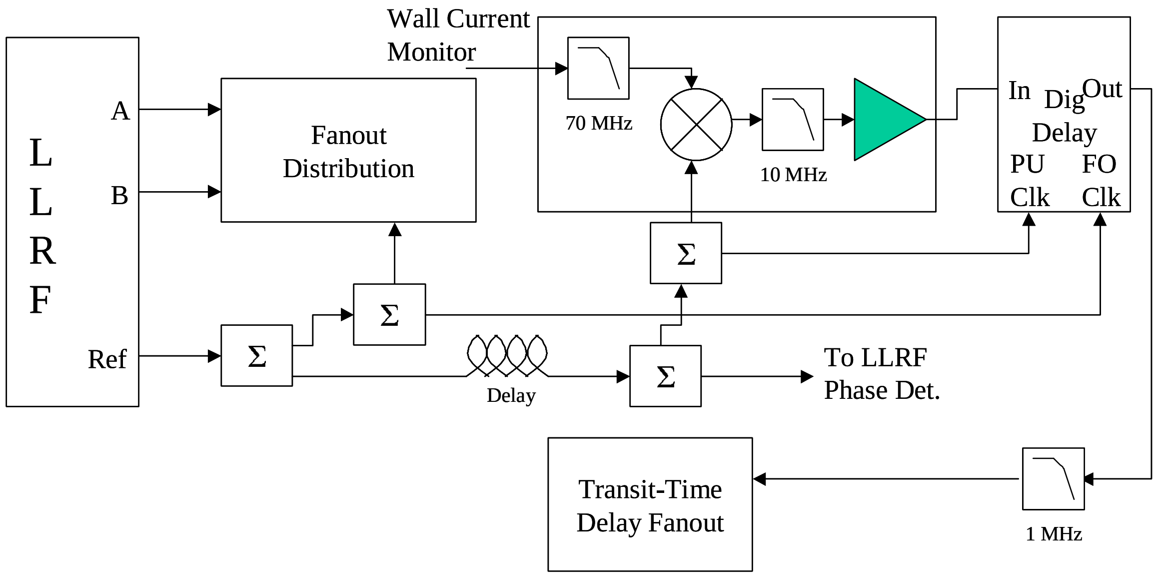

Figure 3.1.53 *Feedforward beam signal processing. Signal from the wall current monitor is downconveted, filtered, and delayed digitally. The output from the delay enters the transit-time delay fanout.*

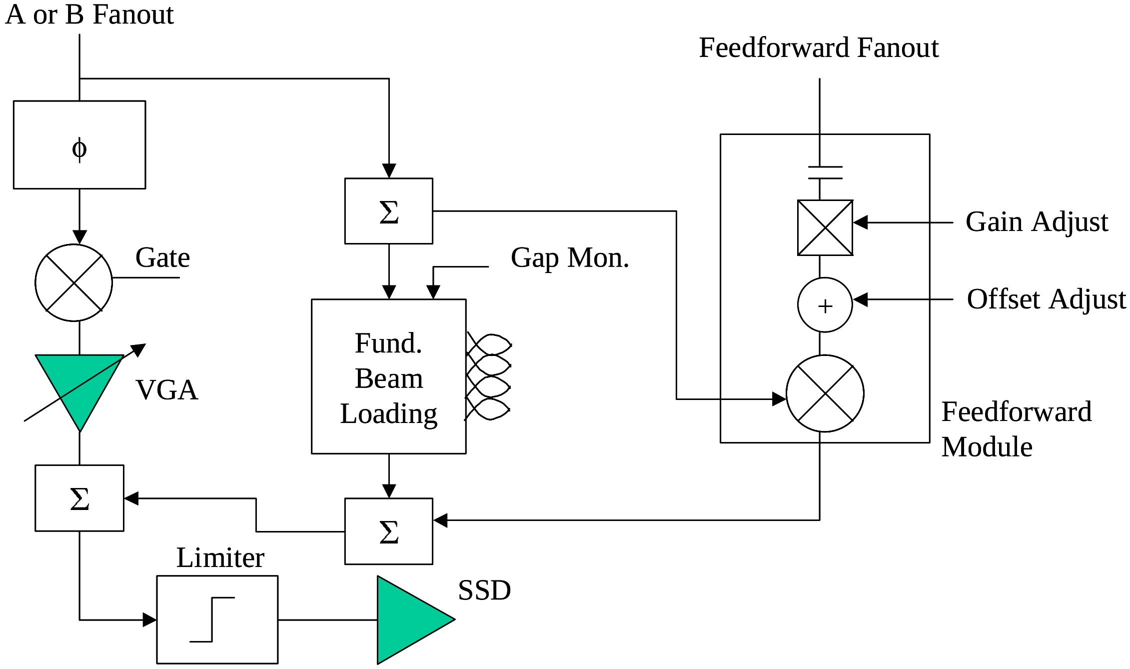

Figure 3.1.54 *The signal from the feedforward processing is upconverted and combined with the fanout and direct RF feedback signals.*

It would be desirable to have a feedforward system working in conjunction with a feedback system for slip stacking. Both systems would help each other and may ease some of the specifications on each system for reaching the beam loading compensation goals. Beam loading compensation is most important during the slipping process when the RF buckets are very small. This process has the advantage that it occurs at a single energy, and the cavities are fixed at a certain voltage level. The challenge is to deal with the multiple fundamental components of beam.

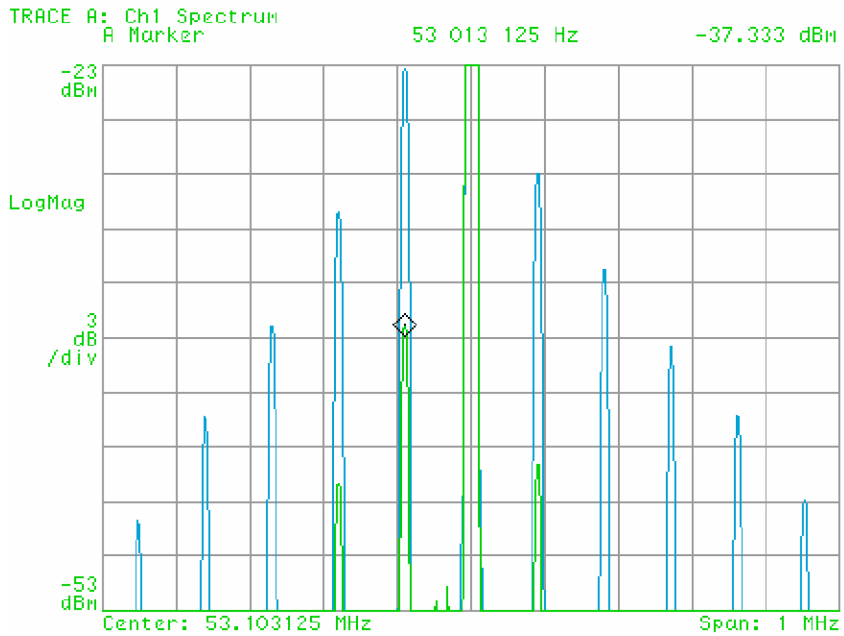

Figure 3.1.55 *Results of feedforward compensation on a single cavity during a normal stacking cycle at high field. The blue trace is the spectrum of the cavity gap without compensation, and the green trace is with compensation.*

3.1.6 Other Slip Stacking Issues

There are some other miscellaneous issues that need to be addressed to make slip stacking a reliable means of increasing antiproton flux. Once the two batches are captured into a single batch, the main ring cycle must continue, as it does for a standard stacking cycle, and accelerate beam to 120 GeV. The difference between this and a standard stacking cycle is that the beam now has twice the intensity and three times the emittance.

Before beam can be accelerated, the main ring LLRF system must reestablish lock to the beam, and the high level feedback loops must be re-enabled. As was mentioned previously, the high level loops have been modified to allow for smooth turn on after capture. The LLRF system should not have a problem reestablishing lock to the beam. It already successfully reestablishes lock to beam after a coalescing cycle where synchronization is transferred from the 53 MHz system to the 2.5 MHz system and back again.

Another issue in the acceleration process is transition crossing. The Main Injector was designed to handle as much as 0.5 eV-s of longitudinal emittance at transition without significant emittance dilution.[18] This should be tested as part of the slip stacking operational commissioning. If there is significant emittance dilution at 0.3 eV-s, then some kind of gamma-t jump may need to be implemented.

Finally, coupled bunch instabilities may develop at the new, high intensities. For slip stacking only, there should not be a serious problem. The design goal for slip stacking is 1e13 protons per pulse, and the Main Injector has already achieved its commissioning goal of 2e13 protons for fixed target with only simple coupled bunch feedback loops that are still available for use. When NuMI begins its operation, however, the intensity goal of the Main Injector will get to 3e13 protons per pulse. This may lead to faster growing instabilities.

The digital beam loading compensation system can be modified easily to become a synchrotron motion damper as well. By changing the dsp program, the digital filter can have a response at the revolution frequency and the synchrotron sidebands. It may even be possible to automatically track the frequencies of the synchrotron sidebands. The digital compensation system can also be modified to act as a transverse damper. It will not be the case of just changing the dsp software as it is for the longitudinal case, but once the dsp, its clocks, and memory interfaces all work together, adding extra channels for processing stripline pickups is relatively straight forward.

3.1.7 Conclusions

There are four distinct paths of effort required to successfully commission slip stacking. First, low intensity operational studies one the Main Injector must continue. These studies include tuning the capture process, finding the optimal frequency separation, accelerating to 120 GeV, and documenting beam emittance for different beam loading compensation conditions. Second, the simulations must prove their reliability against the beam studies. The simulations must model beam loading accurately, and they must have a reliable model for physical beam loading compensation, not just zero delay feedback. Third, studies on beam loading compensation systems must continue. A simple feedforward system must be commissioned in the Main Injector to insure that the high level system has the available power to compensate beam loading during slip stacking. Finally, work should continue on a digital beam loading compensation system. Although there is not enough data to prove that we need a digital system yet, it is highly likely that we will, given the beam loading parameters slip stacking necessitates. Since a complicated digital feedback system requires a significant amount of lead time, work should continue in parallel with the other studies.

The goal of the project is to have high intensity slip stacking commissioned by the middle of FY03. Achieving this goal will require the following budget and full time equivalents:

Budget:

Digital feedback system x 18 cavities + spares	\$190,000
High level RF system upgrades	\$70,000

This budget assumes that no major upgrades of the RF power amplifiers or cavities are necessary. As it stands, no calculations have revealed any short comings in the high level system to handle beam loading compensation with slip stacking.

FY02 Full time equivalents:

2 Engineers – Design and commissioning of beam loading compensation hardware, modification of high level systems as necessary, modification of low level RF system as necessary.

2 Physicists – Study slip stacking operation, commission slip stacking, simulate slip stacking in the presence of beam loading and beam loading compensation.

2 Technicians – Construction of beam loading compensation hardware, modifications of high level systems as necessary.

FY03 Full time equivalents:
1 Engineers, 1 Physicists 1 Technicians

	Total	M&S	Labor	Phys.	Eng.	Draft	Tech	CP
FY02	770	170	600	2	2	0	2	0
FY03	390	90	300	1	1	0	1	0
FY04	0	0	0	0	0	0	0	0
FY05	0	0	0	0	0	0	0	0
Project	1160	260	900	3	3	0	3	0

Table 3.1.4 *Funding profile for slip stacking project.*