V. Stochastic cooling

A. Introduction/overview

Beam cooling is a technique whereby the transverse size and energy spread of a particle beam circulating in a storage ring is reduced without any accompanying beam loss. The goal is to compress the same number of particles into a beam of smaller size and energy spread, i.e. to increase the particle density. Phase space density can be used as a figure of merit for a particle beam, and cooling increases the density. On the surface, it would appear that stochastic cooling violates Liouville's Theorem, which states that phase space volume is conserved. However, Liouville's Theorem only applies to "conservative" systems and stochastic cooling, by definition, is not a conservative process. The cooling electronics act on the beam through a feedback loop to alter the beam's momentum or transverse oscillations.

Two types of beam cooling have been demonstrated and used at various laboratories, including Fermilab: electron cooling, which was pioneered by G. I. Budker and associates at Novosibirsk, and stochastic cooling, developed by Simon van der Meer of CERN. Electron cooling gets its name from the fact that an electron beam is used to cool the particles by removing energy. Stochastic cooling is so named because of the stochastic nature of the beam – i.e., particles move at random with respect to one another.

Theoretically, electron cooling works on the principle of a heat exchanger. Two beams travel a certain distance parallel to each other: a 'warm' beam of protons, antiprotons, or heavy ions with relatively large variation in transverse or longitudinal kinetic energy and a 'cold' beam of electrons having much less variation in kinetic energy. Both beams are tuned to travel at approximately the same velocity, and as the beams interact, the kinetic energy of the warmer beam is transferred to the electron beam. The electron beam can then be collected at the end of the cooling section, or recirculated. Note here that electron cooling is more effective longitudinally than transversely due to the limited transverse size of the electron beam.

Electron cooling was demonstrated at Fermilab in the early 1980's in a small storage ring known as the Cooling Ring which was located in a blue plywood racetrack-shaped building west of the Linac and Booster. It was on this machine, too, that stochastic cooling was first achieved at Fermilab.

During the design of the Fermilab Antiproton Source, electron cooling was not used because of the lack of proven high current relativistic electron sources. Since then, the technology has improved to the point that electron cooling is a viable alternative for future medium-energy storage rings. For that reason, electron cooling was developed for use in the Recycler Ring. Since the Antiproton Source only employs stochastic cooling at this time, the remainder of this chapter will concentrate on this technique for beam cooling. The stochastic cooling systems used in the Antiproton Source are either betatron or momentum. Betatron, βtron and transverse all refer to systems that reduce betatron oscillations in the horizontal and vertical transverse planes. Similarly, momentum, longitudinal, dp, and Δp are used interchangeably to describe systems that reduce the momentum spread of the beam.

B. Fundamentals

The terms beam temperature and beam cooling have been borrowed from the kinetic theory of gases. Imagine a beam of particles circulating in a storage ring. Particles will oscillate around the beam center in much the same way that particles of a hot gas bounce back and forth between the walls of a container. The larger the amplitude of these oscillations in a beam, the larger the beam size will be. The mean square velocity spread is used to define the beam temperature in analogy to the temperature of the gas. Beam cooling is desirable for applications such as:

- Providing a low emittance beam to a collider ring in order to maximize collision rate (luminosity).
- Accumulation of rare particles – cooling to make space available so that more beam can be stacked into the same storage ring (e.g. the Accumulator).
- Preservation of beam quality – cooling to compensate for various mechanisms leading to growth of beam size and/or loss of stored particles. Stochastic cooling was attempted (unsuccessfully) in the Tevatron for this reason and was known as "Bunched Beam Cooling".
- To provide a particle beam with an extremely small energy spread for precision experiments. The E760 and E835 experiments had successful runs in the 1990's, using the

Accumulator to collide antiprotons with hydrogen atoms from a gas jet.

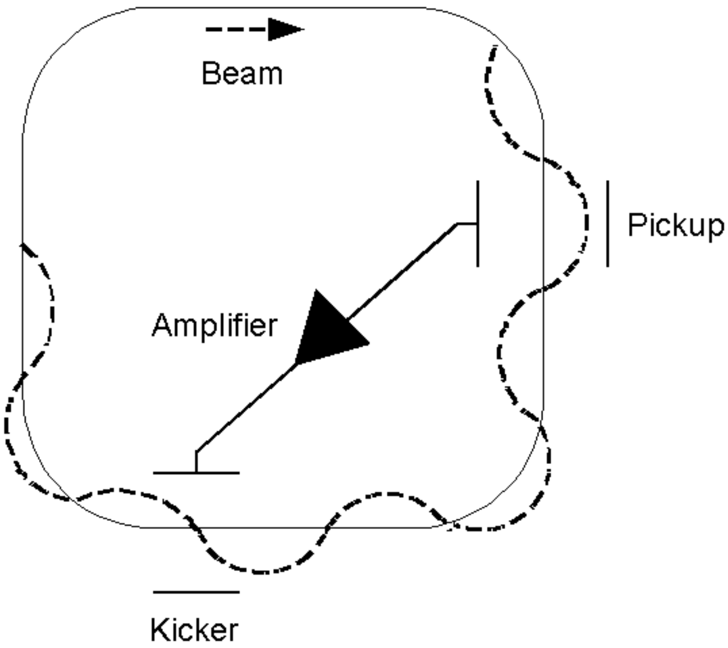

Figure 5.1 Single-particle model for a transverse stochastic cooling system

Consider a single particle circulating in a storage ring as shown in the single particle model depicted in figure 5.1. Assume that the particle has been injected with some error in position and angle with respect to the ideal orbit (the center of the beam pipe). As the focusing system tries to restore the resultant deviation, the particle oscillates around the ideal orbit. These betatron oscillations can be approximated by a purely sinusoidal oscillation. The cooling system is designed to damp the amplitude of this oscillation. A pick-up electrode senses the position of the particle on each revolution. The error signal is ideally a short bipolar pulse that has an amplitude that is proportional to the particle's deviation from the central orbit at the pick-up. The signal is amplified and applied to kickers which deflect the particle by an angle proportional to its error.

Specifically, consider a horizontal beam pick-up that consists of two plates (usually parallel) and is sensitive to either horizontal motion or equivalently a dipole oscillation. The pick-up is centered on the middle of the beam pipe, with one plate to the left of center and the other to the right. If the particle passes through the pick-up off-center, the plate which the particle passes closest to will have a greater current induced on it. If the signals are combined by measuring the difference between them in a so-called 'delta' or Δ mode, the output will be a measure of the relative particle position with respect to the center of the beam pipe. Generally, the output of several sets of electrodes is combined in phase to provide a signal of usable amplitude compared to the thermal noise floor. This signal is then amplified and applied with the most optimal averaged phase (timing) to the kickers. The kicker, like the pick-up, is an arrangement of plates on which a transverse electromagnetic field is created which can deflect the particle.

Since the pick-up detects a position error and the kicker provides a corrective angular kick, their distance apart is chosen to correspond to a quarter of a betatron oscillation (plus a multiple of π wavelengths if more distance is necessary). As shown in figure 5.2, a particle passing the pick-up at the crest of its oscillation will then cross the kicker with zero position error but with an angular deviation which is proportional to the displacement at the pick-up. Given a perfect kicker response and perfect betatron phasing, the trajectory of the particle would be corrected to that of the central orbit. A particle not crossing the pick-up at the crest of its oscillation would receive only a partial correction and require additional passages to eliminate the oscillation. Cooling systems, in fact, require many beam revolutions to cool the beam due to the large number of particles involved and the finite bandwidth of the hardware.

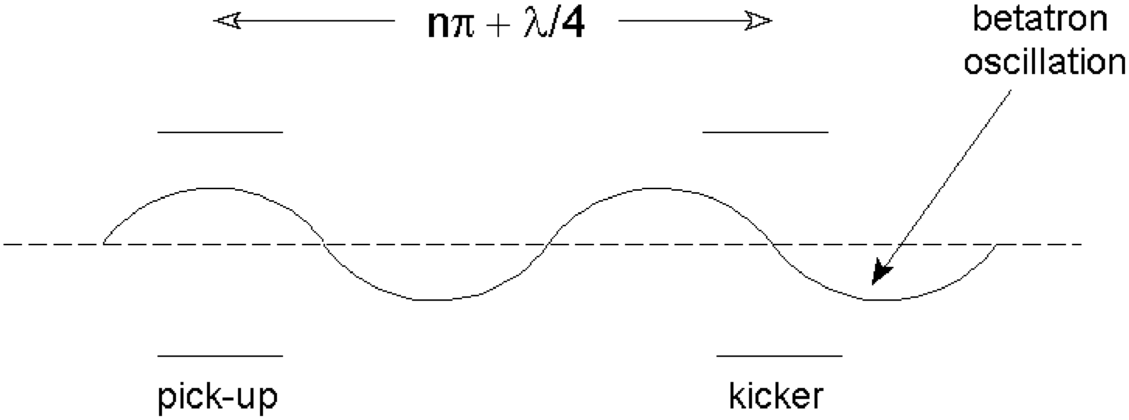

Figure 5.2 Optimum spacing between pick-up and kicker

There is another important aspect of stochastic cooling that this model can illustrate: the correction signal has to arrive at the kicker at the same time as the particle for optimum cooling. Since the signal is delayed in the cables and the amplifier, whereas the particle is moving at close to the speed of light, the path of the correction signal has to take a shortcut across the ring to reach the kicker at the correct time. For reasons explained below, applying the correction signal later than on the same revolution that it was created will lead to less efficient cooling or even heating.

Particle beams, of course, are not composed of just a single particle. Rather, a beam is a distribution of particles around the circumference of the storage ring. Each particle oscillates with a unique amplitude and random initial phase and in this model the cooling system acts on a sample of particles within the beam rather than on a single particle. The number of particles in a sample, N_S, is given by:

$$N_s = \frac{N}{(2WT)}$$

where N is the number of particles in the beam, W is the bandwidth of the cooling system, and T is the beam's transit time around the ring. Using one of the Debuncher systems as an example with N = 1.8 X 10^8 particles, W = 1 GHz (Debuncher systems operate between 4 and 8 GHz, separated into 4 bands), and T = 1.695 μs, the number of particles $N_S \approx 53{,}000$ within each equally spaced sample. Making the bandwidth sufficiently large would, in principal, permit the single particle model above to be valid. However, designing the pick-ups and kickers to accomplish this is not practical.

The cooling process can be looked at as competition between two terms: (a) the coherent term which is generated by the single particle, and, (b) the incoherent term which results from disturbances to the single particle from its fellow sample members through the feedback loop. The coherent signal's contribution to the cooling process is linearly proportional to the system gain, while the incoherent heating term is proportional to the square of the system gain. If one plots these two terms as in figure 5.3, it is clear that there is some point at which the cooling term is

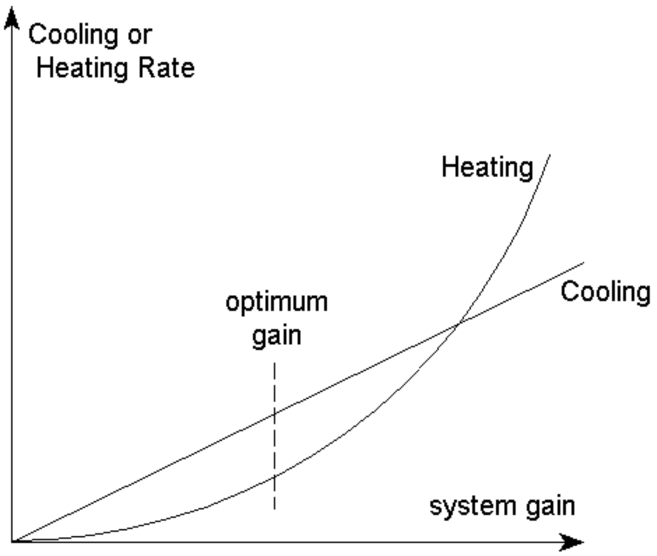

Figure 5.3 Heating and cooling terms as a function of system gain

maximized against the heating term. This is known as the optimum gain of the system. Note that this is usually different from the maximum gain of the system.

Mixing is a term used to represent how completely particles change position with respect to each other. Particles of different momenta "shear" away from each other due to path length differences as they traverse the ring. The stochastic cooling rate is maximized if an independent set of particles constitute each sample upon each revolution. This is sometimes referred to as "good" mixing. The term "stochastic cooling" is derived from the need for a random or stochastic sample of particles passing through the pick-up upon each revolution for cooling to work effectively. Partially random samples are produced because each particle is on a slightly different orbit due to the momentum spread of the beam. The lattice parameter known as the "slip factor," defined as $\eta = \frac{1}{\gamma_t^2} - \frac{1}{\gamma^2}$ where γ is the Lorentz factor ($\gamma = \frac{1}{\sqrt{1^2 + \frac{v^2}{c^2}}}$) and γ_τ is the Lorentz factor at transition, also contributes to the rate at which the particle samples are mixed from turn to turn. If the samples contain mostly the same particles on successive turns, then the cooling rate is decreased.

Although mixing of particles sampled at the pick-up is beneficial, no mixing is desired between the pick-up and the kicker. This is because the signal obtained at the pick-up should be applied at the kicker to the sample of beam creating the signal. Mixing between the pick-up and kicker is sometimes referred to as "bad" mixing. An ideal cooling system would have no mixing between the pick-up and kicker while having complete mixing between the kicker and the pick-up. In reality, the mixing factor present in an accelerator is somewhat less than ideal. The lattice of the storage ring and the momentum spread of the beam determine the mixing factor. It is for this reason that the spacing of pick-ups to kickers should be as small as reasonably achievable while maintaining adequate time for signal amplification and conditioning.

These factors can be written as an equation for the rate, $^1/_{\text{cooling time}}$ or $^1/_{\tau_{X^2}}$ (where τ is the cooling time constant), at which a beam is cooled:

$$\frac{1}{\tau_{x^2}} = \frac{2W}{N}\left[2g\left(1-\tilde{M}^{-2}\right) - g^2\left(M+U\right)\right]$$

where W is the bandwidth of the cooling system, N is the number of particles in the ring, g is the system "gain", or more accurately the number of particles multiplied by the electronic gain, $\tilde{M}$ is the ‘wanted’ mixing factor, M is the ‘unwanted’ mixing factor, and U accounts for random noise.

A list of selected references is included at the end of this chapter which forms the basis for this text and which can provide much more information to the reader on the theoretical aspects of stochastic cooling.

C. Betatron cooling

Betatron or transverse cooling is applied to a beam to reduce its transverse size, i.e. to reduce its horizontal or vertical emittance. The single particle model of cooling described above was that of a simple betatron cooling system. Betatron cooling systems use pick-ups in difference mode to generate the beam’s error signal. In the case of the Antiproton Source, both pick-ups and kickers are located in areas of low dispersion. This is so that any particles passing through the pick-ups off-center will have that position shift due only to transverse oscillations. In a high dispersion region, a particle's position could also be due to differences in momentum, and the resulting kicks could lead to unwanted momentum heating of the beam. The kickers apply a transverse field to the particles by applying the error signal to the kicker electrodes in "push-pull" fashion (one kicker plate has the same charge to push the beam, the opposing kicker plate has the opposite charge to pull the beam). Details of the specific transverse systems in the Antiproton Source are given below.

D. Momentum cooling

Momentum cooling systems reduce the longitudinal energy spread of a beam by accelerating or decelerating particles in the beam distribution towards a central momentum. In a momentum cooling system, the pick-up signals are combined in sum mode and similarly, the signal to the kicker electrodes is also applied in sum mode, providing longitudinal fields to accelerate or decelerate the passing particles.

Momentum cooling is used for several reasons in the Pbar source. Its function in the Debuncher is to further reduce the momentum spread of the beam (bunch rotation is the other mechanism used to reduce the momentum spread in the Debuncher). The stacktail momentum cooling system is used to cool the antiprotons deposited by ARF-1 at the edge of the stacktail by decelerating the antiprotons towards the core. The function of the core momentum systems is to maintain a small momentum spread on the particles in the core. This is desirable for two reasons, first to keep particles from being lost on the Accumulator momentum aperture and second to allow a denser bunch of antiprotons to be extracted during transfers. Accumulator momentum pick-ups are located in high dispersion areas and are positioned over the beam that is to be cooled (stacktail pick-ups over the stacktail, core pick-ups over the core). More details on each Δp system can be found in the following sections.

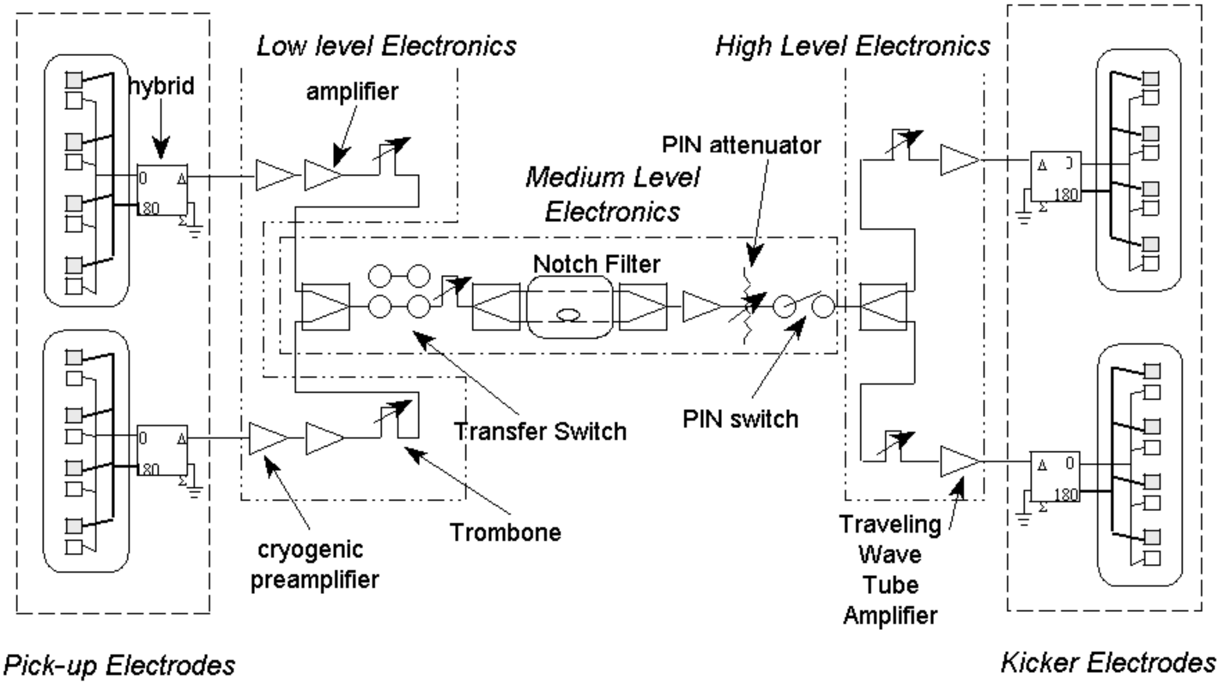

Figure 5.4 Basic stochastic cooling system layout

E. Specific systems

The stochastic cooling systems in the Debuncher and Accumulator are described below (use figure 5.4 as a reference). While each of the stochastic cooling systems perform different functions, they each have similar components, which will be subdivided into six basic parts for this discussion:

Beam pick-up electrodes or slotted waveguides: There are two different kinds of pick-ups used to sample the beam. The stacktail momentum and core momentum systems use beam pick-up electrodes. All Debuncher cooling as well as the Accumulator core transverse systems use slotted waveguide pick-ups. Both systems provide the same basic functionality, to provide a beam error signal to be processed by the cooling system.

Beam pick-up electrodes are quarter-wave loop (directional coupler) pick-ups that are contained within a tank assembly, which is kept under vacuum. The pick-up electrodes are striplines, with a terminating resistor on the adjacent grounded walls of the tank. Figure 5.5 illustrates the electric field lines generated by the passage of charged particles. More accurately, each antiproton generates a short pulse in the stripline as it traverses the gaps.

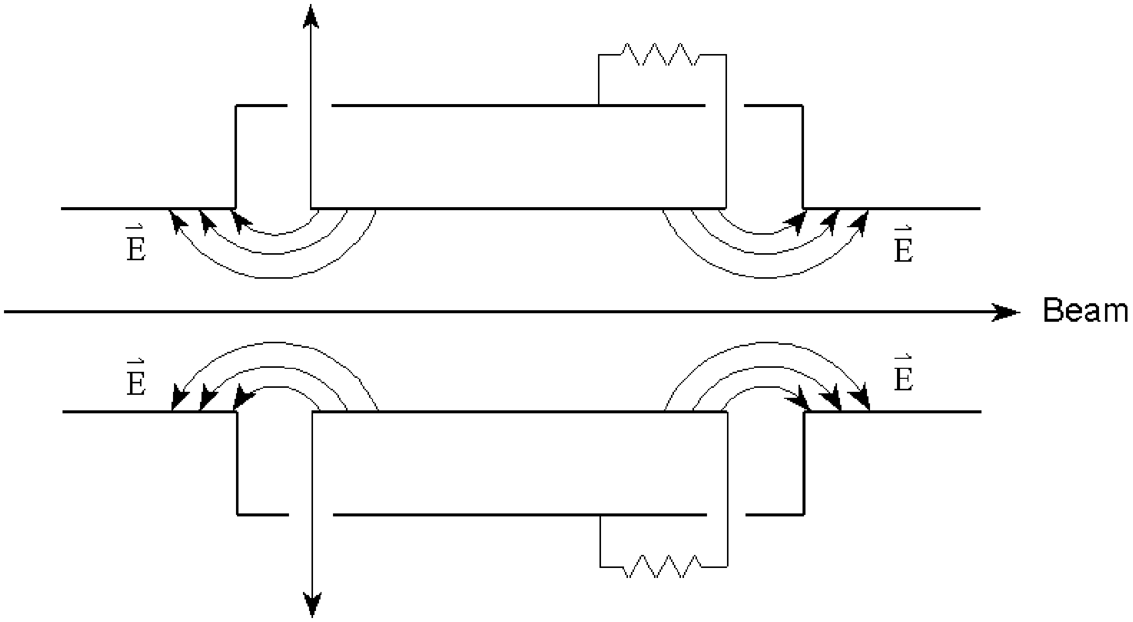

Figure 5.5 Stripline Pick-up

The pick-up plates form transmission lines that have a characteristic impedance. A series of pick-up electrodes are housed in a pick-up tank. Opposing electrodes (top and bottom or left and right, depending on the application) are combined in phase by combiner boards. Sum and difference signals are created by adding or subtracting signals between plates found on opposite sides of the vacuum chamber. Difference signals are used for betatron cooling; sum for momentum cooling. The sum and difference signals are created by passive devices known as hybrids.

The Slotted waveguide "slow wave" structure is shown in Figure 5.6, with the outline of a quarter at the bottom to provide a sense of scale. The rectangular beam pipe (blue box) is coupled to two rectangular waveguides (magenta boxes) by a series of slots. The transverse signal is derived from the difference of the two waveguides and the momentum signal is derived from the sum of the two waveguides. Beam traveling through the accelerator leaves a charged image current on the wall of the conductive beam pipe. The image current is interrupted by the slots and the electromagnetic waves are excited in the slots, which in turn excite traveling waveguide modes in the side waveguides and beam pipe.

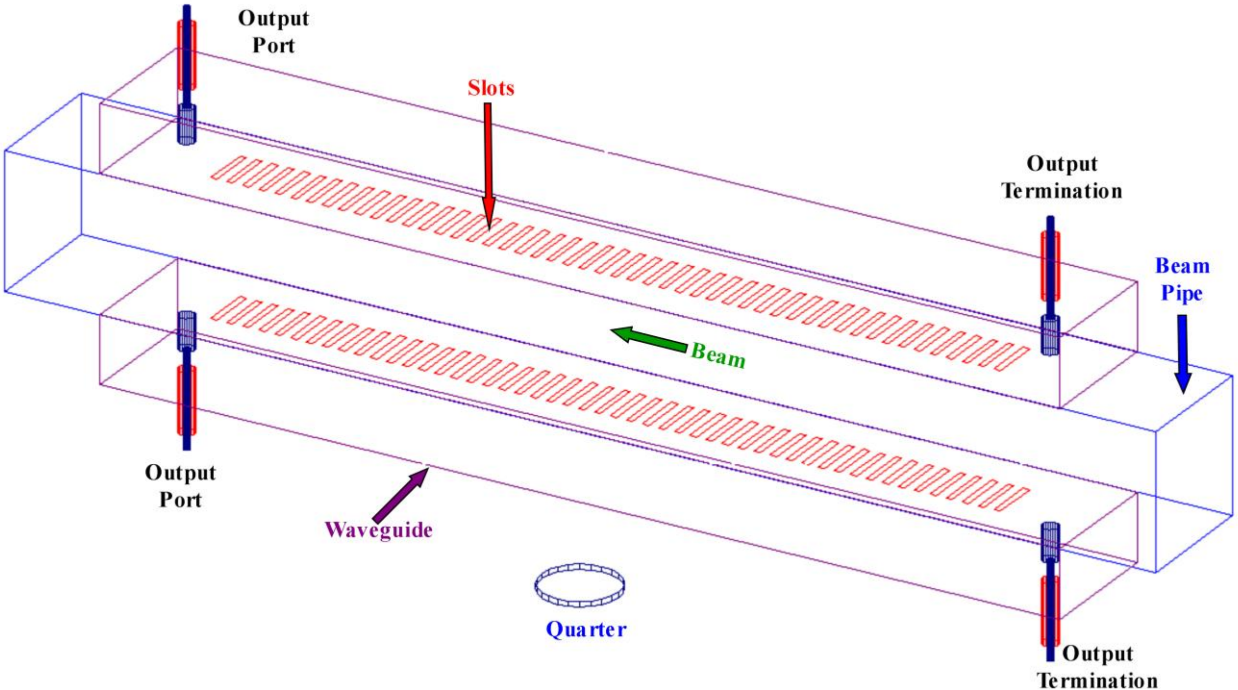

Figure 5.6 Slotted waveguide

The phase velocity is the rate at which the phase of the electromagnetic wave propagates through the waveguide. If there were no slots in the waveguide, the phase velocity would actually be faster than the speed of light. This does not violate relativity, since the speed at which energy is transported in the waveguide, called the group velocity, is not greater than the speed of light. This is similar to how water waves on a lake shore can appear to move much faster (phase velocity) than the actual movement of the water (group velocity). Back to our waveguides, the slots actually "slow down" the waveguide phase velocity modes by creating multiple reflections. The reduction in phase velocity is a function of the slot length and width and the spacing between the slots, and the coupling of the slots to the beam is proportional to the slot length. When the reduced phase velocity of the waveguide exactly matches the beam velocity, the coupling of the slots will add constructively. As a result, the output signal actually grows over the length of the slots like an "Electromagnetic whistle." The gain of the array is proportional to the number of waveguide slots and the length of the array; however, the bandwidth is inversely proportional to the length of the array.

The pick-up assemblies are cryogenically cooled with liquid helium to a temperature of about 4.5° Kelvin for the pick-up and 10° Kelvin for the amplifier. The signals are amplified with low noise cryogenic preamps, and narrow band filters reject signals outside of the desired frequency. Unlike the electrode pick-up systems that operate in the 2-4GHz range, the slotted waveguide system is designed to operate in the 4-8GHz range. This system provides a stronger response than the older pick-up electrode system, but the response is over a narrower bandwidth. As a result, slotted waveguide systems are divided into multiple bands. The Debuncher cooling systems have eight bands in each plane for the pick-ups (four bands each divided into upper and lower bands) This was done by making longer, narrower, band arrays that have higher sensitivity and also reduced by a factor of 2 the number of cold to warm feed through transitions in the pick-ups (for heat load considerations). The Debuncher kickers have four bands, as the fan-out system utilizes many TWTs to limit power dissipation at individual power feed-throughs. The Accumulator core transverse systems have three bands each. There is one pick-up array per sub-band per plane.

Low level electronics: the resultant sum and difference signal is amplified and added in phase with signals from other cooling tanks, if necessary, by means of mechanical delay lines known as trombones. The first stage of amplification is accomplished by GaAsFET preamplifiers, which in most cases are cryogenically cooled to reduce thermal noise. The Debuncher preamplifiers are cooled to liquid helium temperature, stacktail preamplifiers are cooled to liquid nitrogen temperature. Core systems do not require cryogenic cooling because there is a stronger signal from the beam. The 2-4

GHz core momentum system is the exception, the preamplifiers are cooled to liquid nitrogen temperature. Since the pick-up tank is located in A60 along with the stacktail pick-up tanks, there was little additional expense required to provide liquid nitrogen to preamplifiers. Ultimately, an amplified signal with a good signal to noise ratio is the input to the next level of the system.

Medium level electronics: more amplification is applied and the signal is sent towards the kickers on a single coaxial cable known as a trunk line. Trombones are again used to ensure that the corrective signal arrives at the kickers at the appropriate time. Also included in the medium level electronics are variable PIN (P type, Intrinsic, N type semiconductor) attenuators which permit the gain of the system to be adjusted. Increasing the attenuation (expressed in units of dB's) will lower the power output of the system.

Another kind of component found in the medium level is two varieties of switches. Coaxial mechanical transfer switches break the continuity between the pick- up and kicker in order to make open loop transfer function measurements. The beam is a feedback element in this measurement. PIN diode switches are an additional means of opening and closing the circuit. PIN switches are used because they are solid state devices that do not have mechanical fatigue problems from frequent cycling. Most PIN switches have gating capability: the switch can be turned on (the circuit is closed), off (the circuit is open), or gated (the switch can be automatically turned on and off via timers). The core systems, for example, are gated during beam transfers so that the cooling is turned off when unstacking occurs and is turned back on after the transfer has been completed.

An important component of many of the system's medium level circuitry is notch filters. Notch filters act to remove undesired components of the signal from the pick-up before being applied to the kicker (in the case of the Accumulator stacktail and Debuncher betatron systems) or to shape the gain profile (as in the case of the Debuncher momentum system). Specific examples will be provided with the description of each cooling system below. Notch filters built for the cooling systems are of the correlator type, which use the constructive and destructive interference of the same signal transmitted over two transmission lines – like an interferometer. The basic components of the filters are a splitter, trombones, a delay element equivalent to a one turn delay and a hybrid. The splitter splits the medium level signal between two legs - a 'short' leg which is a straight ahead path for the incoming signal and a 'long' leg which consists of a Bulk Acoustic Wave (BAW) delay line, fiber optic link, or superconducting coaxial delay cable. The difference between the two paths is precisely equal to an integer number of revolution delays. Momentum cooling uses a one turn delay, Debuncher cooling systems utilize a one/two turn delay that is switched mid cycle. Debuncher transverse cooling utilizes two turn delays. Trombones are used to maintain the proper delay between the two legs and a 180-degree hybrid combines the two legs.

High level electronics: the signal from the medium level is fanned out to all of the kicker tanks and unraveled in time as appropriate by means of splitters and trombones. Prior to being applied to the kicker electrodes or slotted waveguides, the signals are further amplified at microwave frequencies through devices known as Traveling Wave Tubes or TWTs. Although part of the high level, the TWTs are treated separately here.

Traveling Wave Tube The TWT is a linear beam tube amplifier that provides 30-60 db of gain over octave bandwidths at microwave frequencies. Power levels of a few watts to thousands of watts are attainable. The TWTs used in the Antiproton Source for stochastic cooling operate over octave bandwidths of 2-4 GHz and 4-8 GHz. Each has a saturated power level of 200

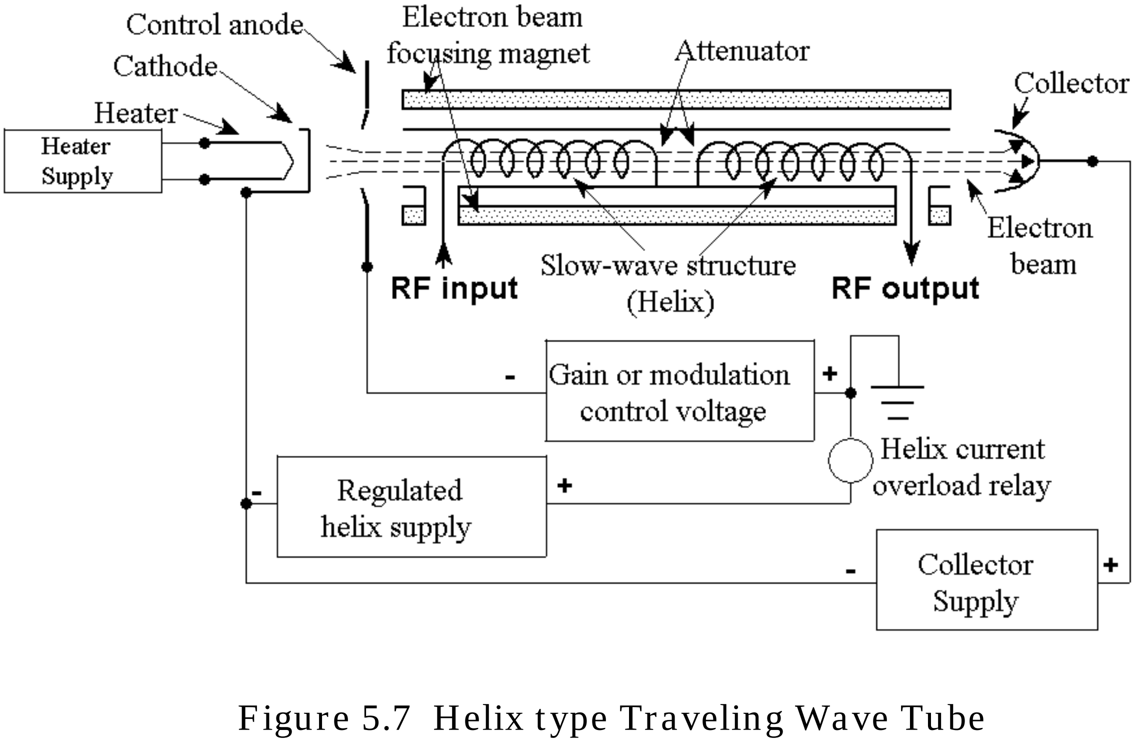

Figure 5.7 Helix type Traveling Wave Tube

watts and 40-50 db of gain, although they are normally run at 100 watts or less. Refer to figure 5.7, which diagrams a typical TWT, as you read the description that follows.

An electron beam is accelerated down the center of a helical 50 Ω transmission line, with the helix power supply providing the source of acceleration voltage. The kinetic energy of the electron beam is typically 3-10 keV and beam currents in the 200-500 mA range are produced from the TWT's used in the Antiproton Source. The microwave signal to be amplified is applied to the helical transmission line. Due to the relatively slow velocity of the electron beam, the helical transmission line acts as a "slow wave" structure forcing the propagating microwave signal to match the velocity of the electron beam. Adjustment of the helix supply is necessary to properly match the velocities and optimize tube performance. Propagating in "sync" causes a velocity modulation or bunching of the electron beam resulting in the electron beam imparting some of its energy to the latter part of the slow wave transmission line structure (i.e. gain).

The transmission line is not a resonant structure, hence a TWT can have a wide bandwidth of operation. An attenuating material is used to support the helical structure to provide isolation between the input and output (if the attenuation material is omitted, it is a BWO or Backward Wave Oscillator). The entire slow wave structure, electron source (cathode) and collector are housed in a sealed stainless steel vacuum envelope. The beam is confined within the helix with permanent magnet focusing. Some higher power TWTs use powered solenoid magnets, but those used in the Antiproton Source use rare earth magnets. The efficiency of TWTs is typically below 20% and those used for stochastic cooling in the Antiproton Source are about 10% efficient. The excess beam energy ends up in the collector. To improve efficiency, several stages of collector may be employed. While the stochastic cooling TWTs typically have one or two stages, some may have up to 4 collectors to improve efficiency. An anode may be added to the TWT to provide modulation or gain control. Only the 2-4 GHz TWTs at Fermilab are equipped with a modulation anode, but it is biased to the continuous mode.

The power supplies for a TWT must be very well regulated to produce a stable electron beam. The propagation time through a TWT is approximately 10-15 nanoseconds, while the stochastic cooling systems require timing precision to a few picoseconds. Voltage ripple of just a fraction of a percent is sufficient to cause enough propagation velocity variation in the electron beam to cause system timing problems.

Kicker electrodes or slotted waveguides: There are two different systems used to provide the corrective kick to the beam. The stacktail momentum and core momentum systems use kicker electrodes, while all of the Debuncher cooling as well as the Accumulator core transverse systems use slotted waveguides. Both systems provide the same basic functionality, to provide a

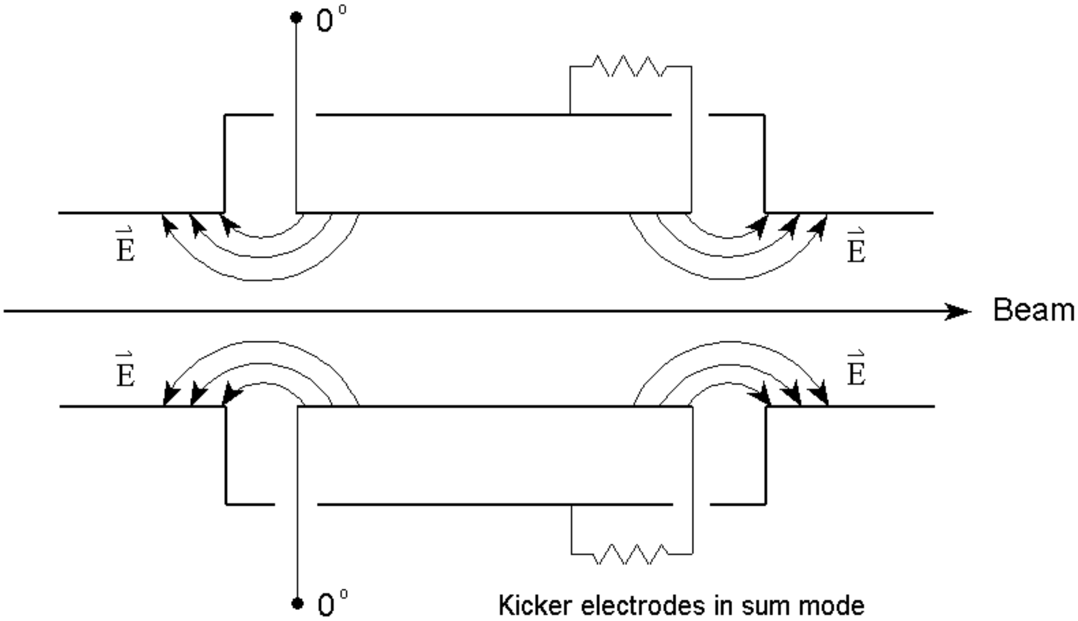

Kicker electrodes in sum mode

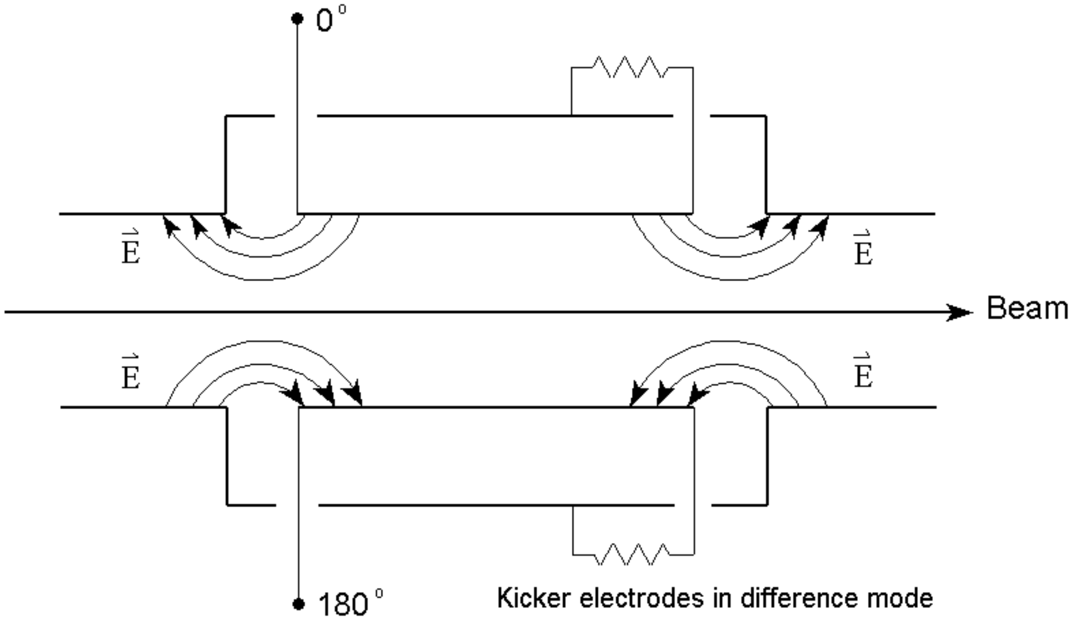

Kicker electrodes in difference mode

Figure 5.8 Kicker electrodes in sum and difference mode correction to the error signal measured by the pick-up electrodes or slotted waveguides.

Kicker electrodes are physically identical to their pick-up counterparts. Each loop is terminated with a resistor and is rated to handle up to 10 Watts of microwave power. The stacktail and core kicker tanks in straight section 30 are outfitted with a design of kicker array referred to as a planar loop, which are made on a printed circuit board to simplify fabrication.

The kicker arrays and terminating resistors are cooled with water provided by a closed-loop chilled (55° F) system. Make-up water to the system comes from Pbar 95 LCW, but there are no de-ionizing cartridges used to preserve the low conductivity. The cooling water is usually referred to as "clean" water, and has excess heat removed by a small refrigeration unit that is located on the skid. Chilled water was originally used for cooling the tanks, but proved to be too dirty and caused clogged flow turbines and reduced cooling efficiency.

Although kicker electrodes for transverse and longitudinal cooling systems are physically the same, there is a difference in how the correction signals are applied to them. Simplified diagrams of kickers in both sum (longitudinal) and difference (transverse) modes are illustrated in Figure 5.8. As with pick-up electrodes, excitation of the beam takes place at the gaps between the pick-up and grounded wall. Note that in sum mode the signals applied to the kicker electrodes are in phase with each other. When in sum mode, the electric fields are oriented so that a longitudinal kick is applied to the beam. In difference mode the signals are 180° out of phase with respect to each other and the electric fields result in a transverse kick to off-center particles.

As with their pick-up counterpart, the kicker waveguides are based on the principle of slowing the phase velocity of the waveguide modes in the beam-pipe and input/output waveguides to match the velocity of the beam. A transverse correction can be made to the beam by applying a correction with the phase between opposing input/output waveguides at 180 degrees. Likewise, a longitudinal correction can be made by applying the correction with the phase between opposing input/output waveguides at 0 degrees. Slotted waveguide kickers are physically similar to their slotted waveguide pick-up counterparts. In the Debuncher's case, the pick-ups are twice as long as the kickers. There are only four kicker bands per plane, compared to the eight pick-up bands. Each TWT puts out about 150W. One important design consideration is stopping the microwave energy from the waveguides from propagating around the ring. To achieve this, LCW cooled microwave absorbers were added in the tanks at each end in order to absorb stray microwave power. The LCW cooling is provided by the same cooling water system that is used for the kicker electrodes.

System	Debuncher Horizontal	Debuncher Vertical	Debuncher Momentum
Slotted Wavegide Pick-up location	D10	D10	D10
Slotted Waveguide Kicker Location	D30	D30	D30
# of bands	8 pick-up & 4 kicker	8 pick-up & 4 kicker	8 pick-up & 4 kicker
Overall Bandwidth	4-8 GHz	4-8 GHz	4-8 GHz
Bandwidth (Band 1)	4.0-4.95 GHz	4.0-4.95 GHz	4.0-4.95 GHz
Bandwidth (Band 2)	4.85-5.82 GHz	4.85-5.82 GHz	4.85-5.82 GHz
Bandwidth (Band 3)	5.8-6.9 GHz	5.8-6.9 GHz	6.0-7.1 GHz
Bandwidth (Band 4)	6.65-8.1 GHz	6.65-8.1 GHz	7.2-8.3 GHz
# of TWTs per band	4	4	8
# of TWTs (Total)	16	16	32
TWT operating power (each)	150 watts peak	150 watts peak	150 watts peak
TWT trip level	175 watts	175 watts	175 watts

Table 5.1 Debuncher cooling systems

1. Debuncher Betatron

The Debuncher Betatron systems reduce the transverse emittances of beam in the Debuncher so that the pbars will transfer efficiently into the Accumulator. Each system presently reduces the emittance from about 30 to 3 pi-mm-mrad in 2.2 seconds. The bandwidth of these cooling systems is 4-8 GHz but is comprised of 4 discrete cooling systems of approximately 1 GHz bandwidth each. Because the Pbar intensity in the Debuncher averages only 1.5–2.0 x10^8 particles, the electrodes and preamplifiers are cooled with liquid helium. This serves to reduce the thermal noise, which would contribute to the less efficient cooling. Unwanted signals are also removed by the use of correlator Bulk Acoustic Wave (BAW) two turn delay notch filters. The BAW filters notch out unwanted thermal noise at harmonics and half harmonics (between the betatron sidebands) of the revolution frequency. Then, leaving only the signals from the betatron sidebands, the signals are amplified by the TWTs and applied to the kickers. By increasing the signal to noise ratio, less

TWT power is produced as noise that would heat the beam, while leaving more power to cool the beam. Debuncher cooling signals from upstream and downstream pick-up tanks exit different 10 sector stub rooms, but arrive at the same stub room in 30 sector, as shown in Figure 5.9.

There are a total of 8 kicker tanks in straight section 30, each is used for both transverse and momentum cooling. One kicker tank is used for each band in each plane, with the horizontal and vertical tanks of each band also utilized for momentum cooling. Due to the length of the pick-up and kicker arrays and the need to keep the proper phase advance between the pick-ups and kickers, the tanks are separated by 180° of betatron phase advance and combined with a 180° hybrid. There are 16 TWTs in each plane operating at 150W peak output each.

2. Debuncher Momentum

Antiprotons that circulate in the Debuncher have their momentum spread further reduced after bunch rotation and adiabatic debunching by means of momentum cooling systems. Momentum cooling was an upgrade that was installed in 1989, and was later updated to a multiple band 4-8GHz slotted waveguide system in 2000. The current system uses the same pick-up and kicker slotted waveguides as those in the Debuncher betatron systems. Instead of using the signals from the pick-ups in the difference mode, however, the signals are summed. Similarly, the signal applied by the kickers to the beam is in the sum mode. The frequency range of this system is 4-8 GHz, and like the betatron cooling is

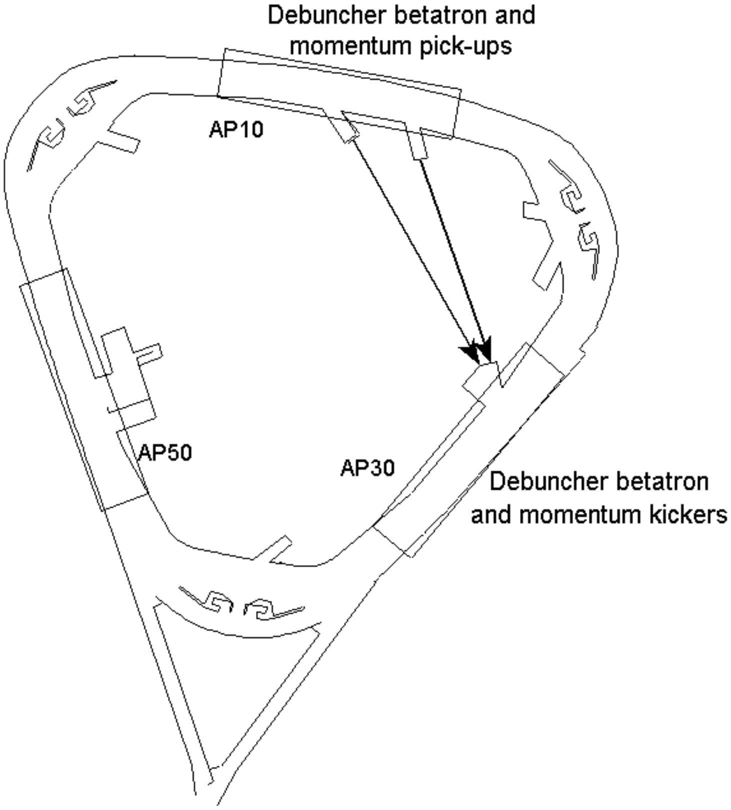

Figure 5.9 Location of Debuncher stochastic cooling components

comprised of 4 discrete cooling systems of approximately 1 GHz bandwidth each. This system currently reduces the Debuncher $\Delta p/p$ (momentum spread) from ~ 0.30% to < 0.14% in 2.2 seconds.

All of the Debuncher transverse pick-up and kicker slotted waveguides are used for the momentum system – the kickers are driven with both momentum and transverse signals. 32 TWTs are dedicated to momentum cooling, again mounted on the kicker tanks, and run at 150 watts peak per TWT watts. This system also has a notch filter that provides the gain shaping necessary for momentum cooling. The filter utilizes a fiber optic delay that switches between a one turn and two turn delay. At the beginning of the stacking cycle, the momentum spread of the beam is wide and requires

System	Stacktail Δp	Core 4-8 Horizontal	Core 4-8 Vertical	Core 2-4 Δp	Core 4-8 Δp
Pick-up Type	Pick-up Electrode	Slotted Waveguide	Slotted Waveguide	Pick-up Electrode	Pick-up Electrode
Pick-up Location	A60	A10	A10	A60	A20
Kicker Type	Kicker Electrode	Slotted Waveguide	Slotted Waveguide	Kicker Electrode	Kicker Electrode
Kicker Location	A30	A30	A30	A30	A50
# of pick-up sets	256 at +13.7 MeV (2 tanks with 128 each) 48 at –6.4 MeV 16 at –22.9 MeV	One slotted waveguide for each band	One slotted waveguide for each band	16 at core orbit 16 at central orbit	32
Number of bands	1	3	3	1	1
Bandwidth	2-4 GHz	**4-8 GHz** Band 1: 4.35-5.65GHz Band 2: 5.35-6.65GHz Band 3: 6.35-7.65GHz	**4-8 GHz** Band 1: 4.35-5.65GHz Band 2: 5.35-6.65GHz Band 3: 6.35-7.65GHz	2-4 GHz	4-8 GHz
# of amplifiers	32 sum 4 delta TWTs	3 five Watt Solid State amps	3 five Watt Solid State amps	1 TWT	2 TWTs
# of kicker pairs	256 with 64 delta kicker pairs (half vertically half horizontally oriented)	One slotted waveguide for each band	One slotted waveguide for each band	32	64
Typical operating power	1,000 watts	150W peak	150 W peak	40 watts	0-10 watts

Table 5.2 Accumulator cooling systems

a single turn notch filter. At mid-cycle, the two turn delay is switched in optically to create a steeper gain profile, further reducing the momentum spread.

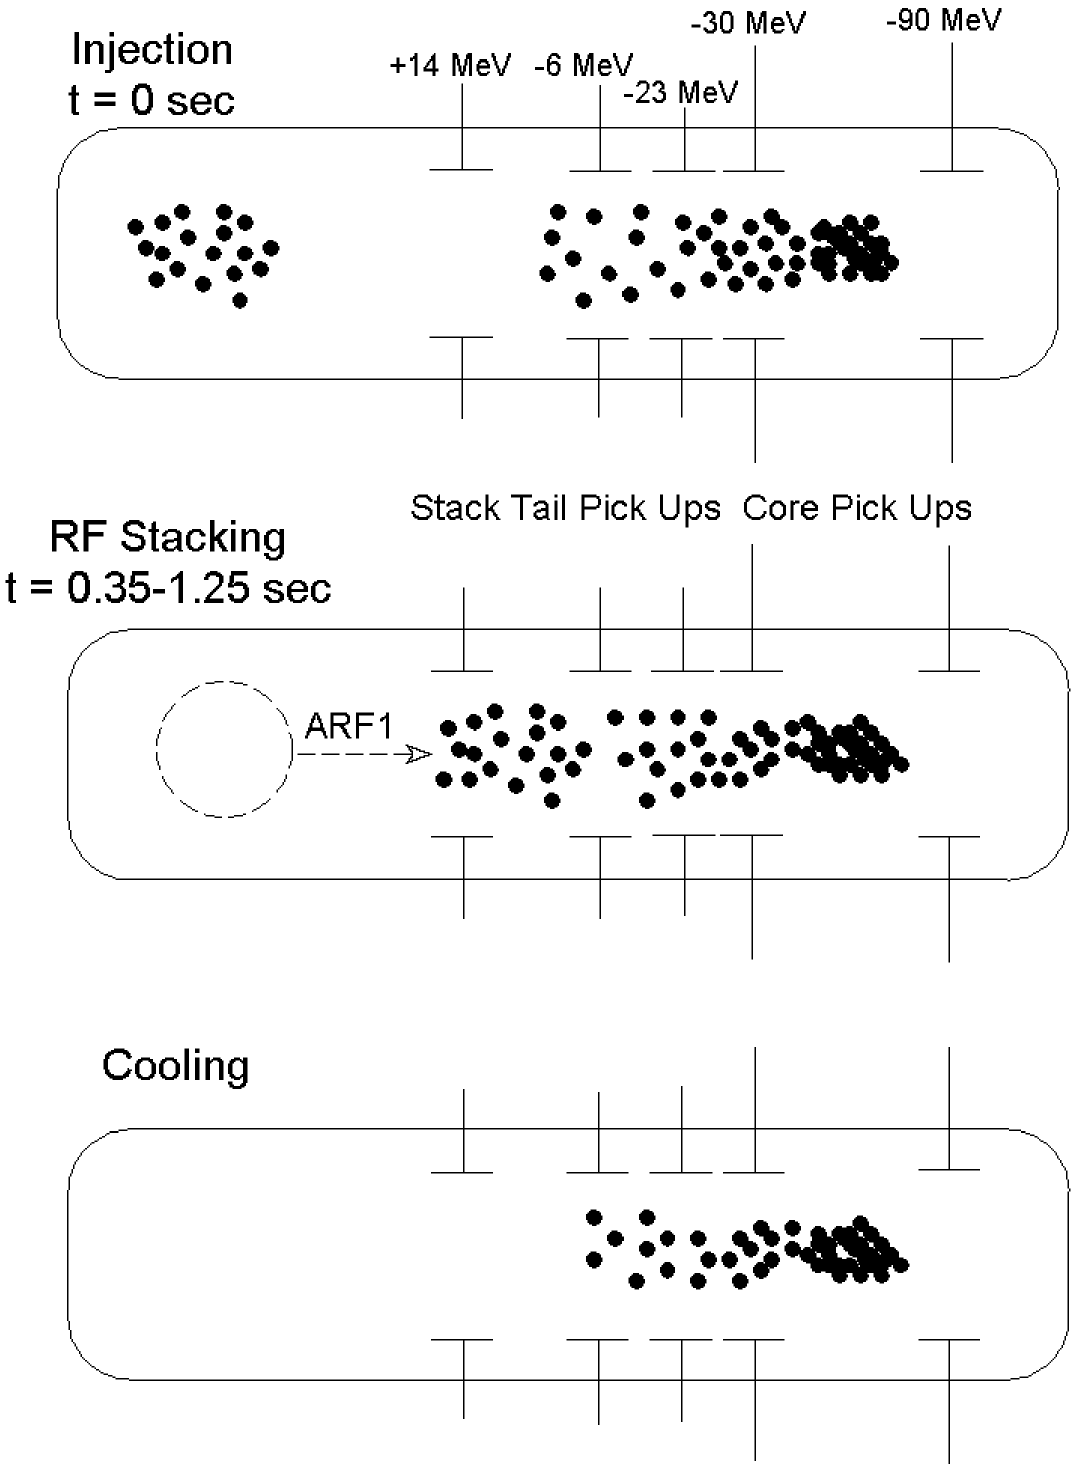

Figure 5.10 Stacktail and core momentum pick-up

3. Accumulator Stacktail Momentum

After antiprotons have been injected into the Accumulator, the particles must be decelerated by roughly 150 MeV to reach the core. The first 60 MeV of deceleration is handled by ARF-1 while the final 90 MeV is accomplished by the 2-4 GHz stacktail momentum system. Because an RF bucket displaces beam that it passes through, it is not possible to use an RF system to decelerate beam the full 150 MeV to the core.

All of the stacktail pick-ups are located in the A60 high dispersion region and are subdivided into three separate arrays called the +13.7 MeV (leg 1), –6.4 MeV (leg 2) and –22.9 MeV (leg 3 or compensation leg) pick-ups. Figure 5.10 shows the relative positions of the stacktail and core momentum pick-ups. The pick-up names identify the part of the stacktail for which the particular pick-up array is most sensitive to, relative to the central orbit of the Accumulator. The stacktail extends from about +30 MeV where ARF1 deposits beam to the edge of the core at about -30 MeV. In the high dispersion region, where the pick-ups are located, a difference in energy results in a primarily horizontal position shift (there is very little vertical dispersion in the Accumulator). A notable difference in the three arrays is in the number of pick-up elements each one contains. The +13.7 MeV pick-ups are made up of 256 individual pick-up electrode pairs divided evenly between

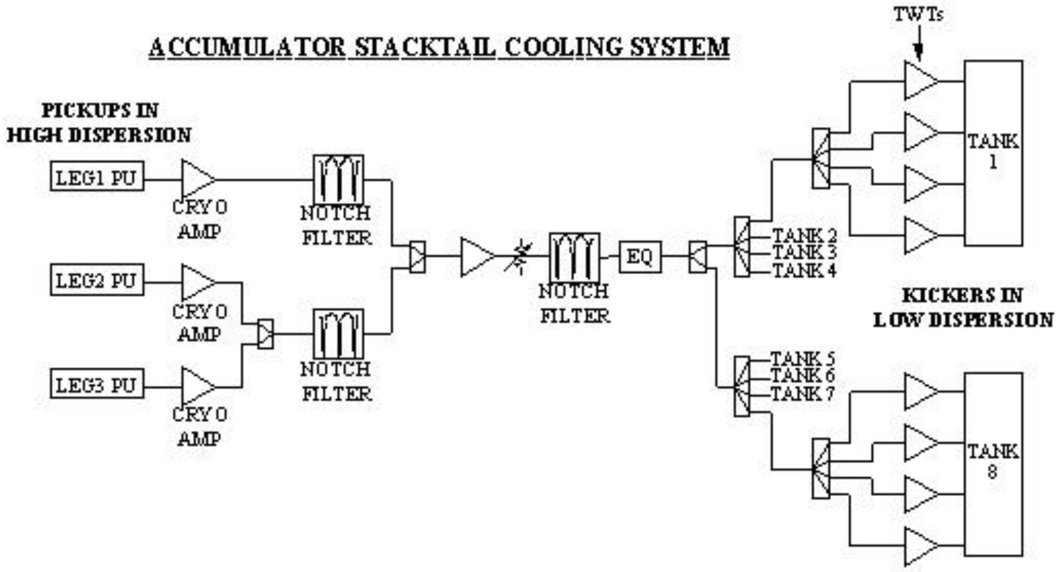

Figure 5.11 Diagram of the stacktail system

two different tanks. The –6.4 MeV pick-ups, made up of 48 electrodes, and the –22.9 MeV pick-ups having only 16 electrodes, are located inside another tank. Figure 5.11 provides a simplified diagram of the stacktail system and Figure 5.12 shows the signal path from the pick-up tanks in the A60 straight to the kickers in the A30 straight section.

To understand why there are so many pick-up electrodes at +13.7 MeV and so few at –22.9 MeV, consider how beam is distributed in the stacktail. At the deposition orbit, the point where ARF-1 drops off the beam, there is a relatively small amount of beam for the +13.7 MeV electrodes to detect. For the stacktail system to work effectively, a certain amount of the beam signal must be detected above the background noise. Thermal noise from the pick-ups is reduced by cooling parts of the pick-up assemblies to liquid nitrogen temperature. To achieve an adequate amount of beam signal above the noise floor from the +13.7 MeV array, it is necessary to have a large number of pick-ups. The –22.9 MeV pick-ups, on the other hand, are located much closer to the core where there is considerably more beam. Sixteen electrodes are adequate to produce a reasonable signal to noise ratio.

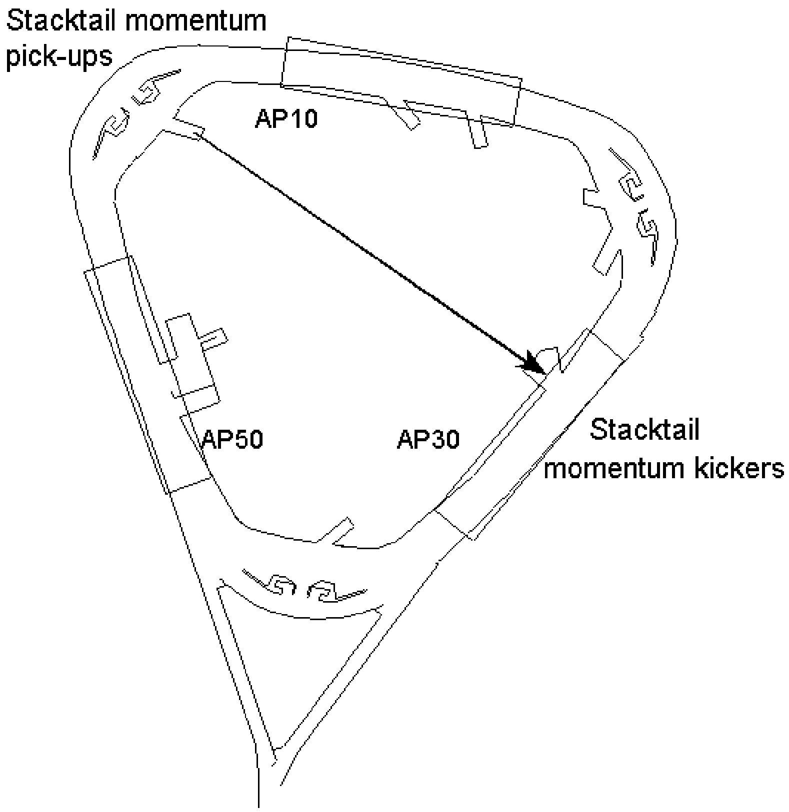

Figure 5.12 Location of stacktail momentum cooling system components

The signals coming from the pick-up arrays are modified by the stacktail electronics to provide the phase and gain characteristics necessary to effectively momentum cool the beam. This must be accomplished while minimizing effects on beam in the core. The system gain changes nearly exponentially across the stacktail, and is highest where ARF-1 drops beam off and lowest at the edge of the core. Because of this, the high energy beam arriving at the edge of the stacktail moves very rapidly away from the deposition orbit. It is important for the stacktail system to have this feature since any beam remaining near the deposition orbit will be RF displaced into the field region of the injection kicker when ARF-1 pulses on the next stacking cycle. Low energy beam on the core side of the stacktail moves very slowly and tends to "pile up" against the core, giving the stacktail its characteristic shape, which is illustrated in figure 5.13.

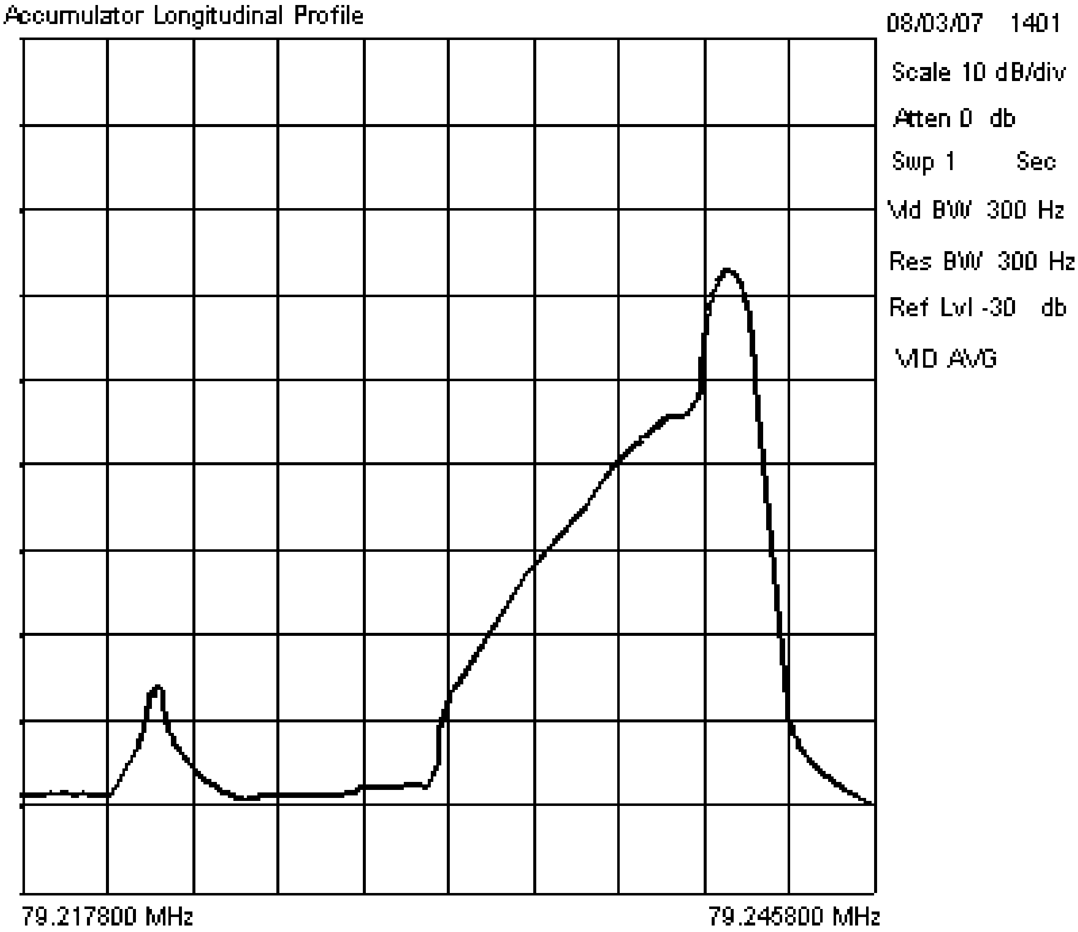

Figure 5.13 Accumulator longitudinal spectrum analyzer display

Transverse kicks induced by the stacktail momentum system, mostly due to imperfect hybrids and kicker misalignment, lead to betatron heating of the beam in the stacktail and core. This can be partially overcome in the stacktail system by applying a small part of the signal from the kicker electrodes in the difference or delta mode (recall that momentum pick-ups and kickers are normally in the sum mode). The first and last kicker tanks in the A30 straight section are stacktail tanks used as "delta kickers". These tanks were selected because they are nearly 90° out of betatron phase with each other. Half of all of the stacktail momentum kicker electrodes are oriented horizontally and the other half are oriented vertically. The delta kickers apply the difference signal to the kicker electrodes, resulting in a transverse kick to the beam. The delay and attenuation values for the delta kickers are calculated using network analyzer beam measurements, then fine-tuned empirically. Delta kickers in the Stacktail have not been used effectively since the upgrade to 2-4 GHz at the beginning of Run II.

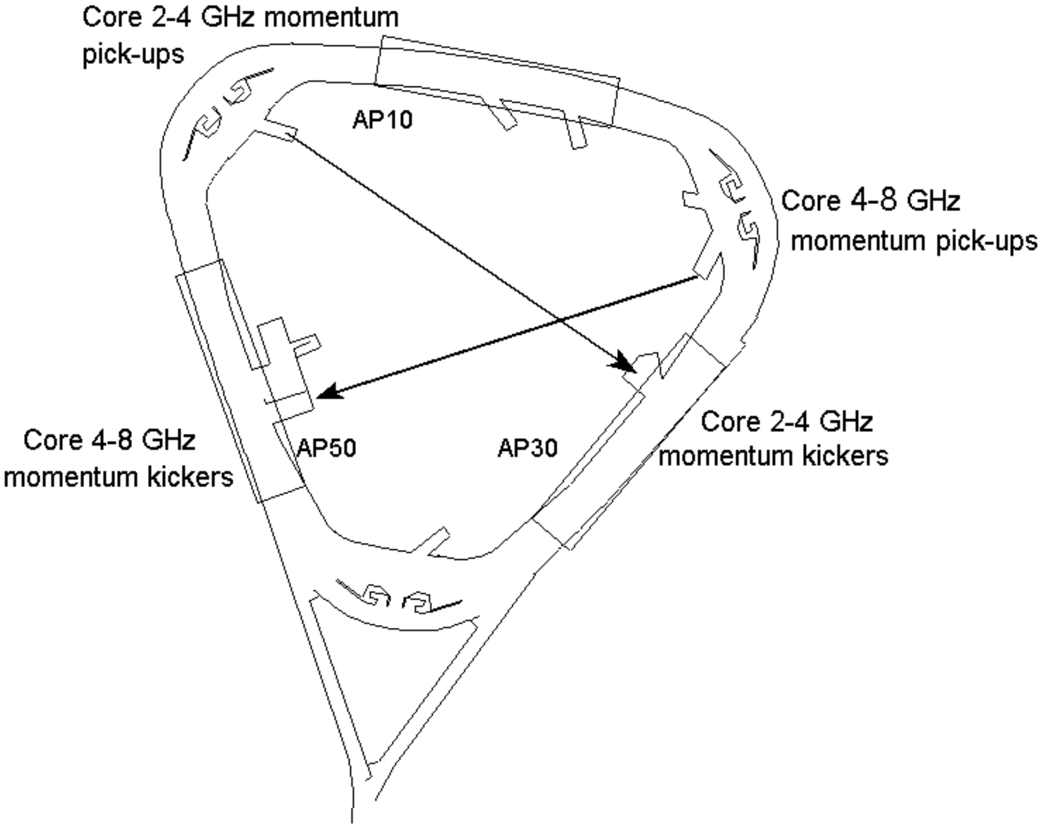

Figure 5.14 Location of core momentum cooling systems

4. Core Momentum

The Core momentum cooling systems keep the antiproton core contained by decelerating high energy particles and accelerating low energy particles. There are two core momentum systems currently in use. The original 2-4 GHz system has its pick-up tank in the A60 high dispersion straight section and kickers in the A30 area. The 4-8 GHz system, added in 1989, includes a pick-up tank in the A20 section and a kicker tank in A50 (see figure 5.14).

The 2-4 GHz and 4-8 GHz systems are used together to provide momentum cooling for the core. The 4-8 GHz system is able to cool the core to a smaller momentum spread with decreased cooling time (because of its larger bandwidth), while the 2-4 GHz system has a greater frequency “reach”. The 4-8 GHz pick-up arrays are moveable, so that system can cool beam away from the core for beam studies. The two core momentum systems are stand-alone systems, so they need to be kept aligned to optimize performance. Normally the pick-ups of the 4-8 GHz system are positioned so that both systems are cooling to the same revolution frequency.

5. Core Betatron

There are three horizontal and three vertical transverse cooling systems spanning the 4-8 GHz band. Each of the three systems in the transverse planes operates over only part of the octave band, as summarized earlier in Table 2. These systems are used to control the transverse emittances of particles in the core. Pick-up tanks are located in the A10 low dispersion straight section, an area where any sensed position error will be due to transverse oscillations rather than energy. The kickers are in the A30 straight section. Both pick-up and kicker tanks use slotted waveguides, like the Debuncher cooling systems.

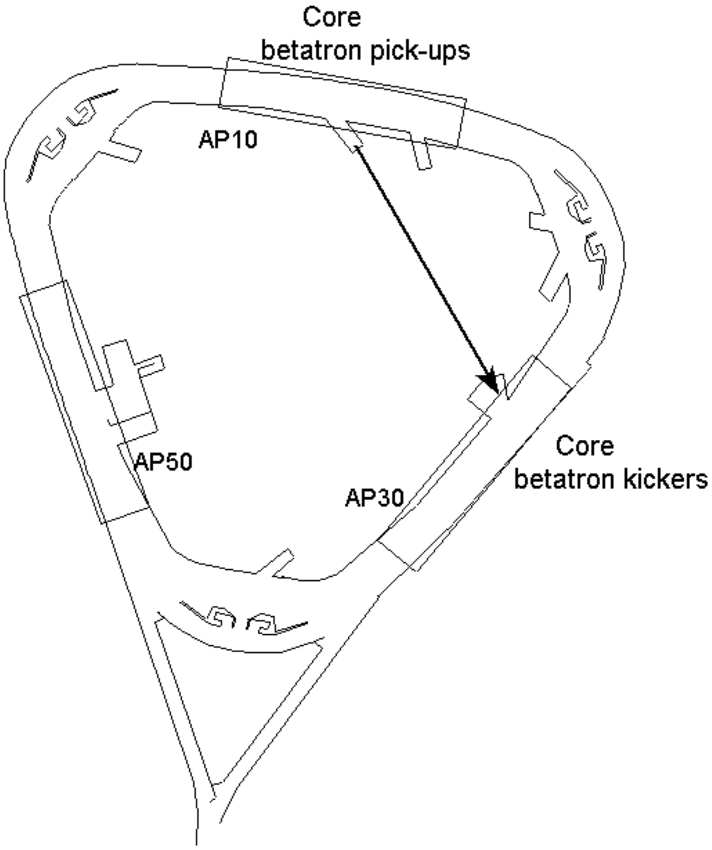

Figure 5.15 Location of core betatron cooling systems

References

1) S. van der Meer, An Introduction to Stochastic Cooling, AIP Proceedings 153, SLAC Summer School, 1985; Fermilab Summer School, 1984; pp. 1628-1649.

2) A. V. Tollestrup and G. Dugan, Elementary Stochastic Cooling, unpublished
3) John Marriner, David McGinnis, An Introduction to Stochastic Cooling, AIP Conference Proceeding #249 VI, 693-761 (1990)
4) D. Möhl, G. Petrucci, L. Thorndahl, S. van der Meer, Physics and technique of stochastic cooling, Phys. Reports 58 (1980) 73-119.
5) Robert L. Shrader, Electronic Communication - fourth edition, McGraw-Hill Book Company.
6) D. Möhl, Stochastic Cooling for Beginners, PS/LEA/Note 84-12.
7) Design Report Tevatron I Project, September 1984, Fermilab.
8) Conversations with Ralph Pasquinelli.
9) D. Möhl, CERN Internal Report, CERN 87-03 V2:453-533 (1987)
10) Bisognano, Leeman AIP Conference Proceeding #87 Fermilab Summer School on Accelerator Physics 1981
11) D.A. Goldberg, G.R. Lambertson, Dynamic Devices; A Primer on Pickups and Kickers, AIP Conference Proceeding #249 VI, 537-600 (1992)
12) Conversations with Dave McGinnis
13) Conversations with Paul Derwent
14) Dave McGinnis, The 4-8 GHz Stochastic Cooling Upgrade for the Fermilab Debuncher, Pbar Note 625, http://www-bdnew.fnal.gov/pbar/documents/pbarnotes/pdf_files/PB625.PDF (November, 1999).
15) Dave McGinnis, Slotted Waveguide Slow-Wave Stochastic Cooling Arrays, Pbar Note 626, http://www-bdnew.fnal.gov/pbar/documents/pbarnotes/pdf_files/PB626.PDF (November, 1999).
16) Paul Derwent, Debuncher Stochastic Cooling Upgrade, Powerpoint presentation (March 22, 2000).
17) John Marriner, Debuncher Stochastic Cooling Upgrade for Run II and Beyond, Pbar Note 573, http://www-bdnew.fnal.gov/pbar/documents/pbarnotes/pdf_files/PB573.PDF (April 1, 1998).
18) Conversations with Dave Vander Meulen