

STOCHASTIC COOLING HARDWARE

D: 8503064774

C.S. Taylor
CERN, Geneva, Switzerland

INTRODUCTION

Particle physics research requires intense and dense beams to be produced by accelerators. Unfortunately, there is a tendency for beams to lose density in their passage through accelerators, and the 2-dimensional projections of the 6-dimensional hypervolume, representing the beam population in the longitudinal and transverse directions, suffer a growth in area. This growth is due to a number of effects such as linear, non-linear and resonant couplings, distortion and filamentation due to non-linear fields, and simple smearing out due to misalignments.

It has been an article of faith that Liouville's Theorem can be applied to the motion of particle beams. This theorem states that in a continuous medium, the density in the phase volume is invariant, and although it was an observable fact that the density in the 2-dimensional phase spaces did not remain constant, it was accepted that distortions of the hypervolume could result in such density changes in the projections while remaining itself invariant, much as the plan view of an incompressible bar will change if it is twisted and bent, or simply tilted - analogous to linear coupling. However, it did not seem likely that these processes could be reversed.

In the late 1960's and 1970's the challenge came from experimental particle physics to obtain higher densities in accelerators in order to have higher luminosities and event rates in colliding beam machines, and if possible to convert the vague cloud of antiprotons from a target (this represents a catastrophic loss of density) into a highly ordered beam useful for proton-antiproton collisions.

The first suggestion came from Novosibirsk, where Budker proposed using an ordered beam of electrons to knock a hadron beam into shape.


Simon van der Meer's answer to the physics challenge was to propose using chance variations to increase order.

This was surprising, as one tends to associate chance variations with randomness and disorder, and also because the method proposed relied on kicking the whole beam. How could this affect the density?


It was also extremely satisfying. It was as if, after watching helplessly for some years as the beam blew up through the accelerators, someone had told you how to reach into the phase space and re-pack the particles more tightly. Here was a way of reversing the disordering process, or of cooling with a reversed heat engine and reducing the entropy, and of course it led eventually to the W and Z⁰.

BASIC CONCEPTS OF STOCHASTIC COOLING


The first point to note is that a circulating beam of particles is both "grainy" and irregular. The "graininess" is analogous to the effect described by Schottky in the case of electrons travelling in a conductor; in the case of a mono-energetic circulating beam the passage of the particles past a perfect detector would produce a spectral line whose frequency would be the frequency of revolution of the particles and whose amplitude would be related to the current circulating. For a beam with a spread of energies the envelope of all the spectral lines would represent the distribution of the particle current with energy - the "Schottky" scan:



The irregularity of the particle population along its trajectory is random or "stochastic" and is exploited in stochastic cooling in the following way. As the beam moves past a given point the centre of charge will follow the fluctuations in charge density across the beam, and an electrode at this point, sensitive only to position, will produce impulse signals giving the instantaneous position of the beam centre:



Assuming that we could have a system of infinite bandwidth which could pass this sequence of delta functions, we could amplify these signals and apply them to a kicker in order to re-centre the measured sample:




What happens to the spread across the beam? (Note that as the distribution is re-centred, the squares of the extreme values are reduced faster than the squares of the smaller values).

Exercise for the Student

- 1) Ask for 5 random numbers from a normal distribution with $r = 0$ and $\sigma = 1$.
- 2) Find mean value \bar{r} ($\bar{r} \neq 0$ in general).
- 3) Subtract error in mean from each number to restore mean to zero.

- 4) Calculate new σ' .
- 5) Replace $\sigma = 1$ in the first step by $\sigma = \sigma'$ and continue.
- 6) Watch the progress of σ' . It sometimes is very irregular at the start (depending on your algorithm and choice of seed for the pseudo-random numbers), but eventually you are rewarded by the cooling of your distribution.



Note

The number of ways in which we can arrange the particles radially has been reduced, and therefore the entropy, ($\ln n$, where n is the number of ways), has been reduced. The order has increased.

Proposition


Insofar as the evolution of plants and animals progresses to a higher state of organization through chance mutations, Darwin's Theory of Natural Selection and stochastic cooling are similar in principle.

Exercise for the Student

Find a biologist and ask him if he thinks that evolution proceeds to higher states of order.

What have we assumed in this simple description ?

- 1) Infinite bandwidth.
- 2) The correcting kick arrives at the same moment as the sampled particles, i.e. the delay in the amplifier chain from the detector, ("pick-up" or "PU"), to the kicker is correct. Neither is there any circuit noise to blur the correcting kick.
- 3) All the particles in the sample arrive at the same moment, so that we correct the original sample, and not a modified sample produced by relative slip between particles of different revolution frequency - there is no mixing between pick-up and kicker.
- 4) We asked for a fresh set of numbers each time, so we have assumed complete mixing of populations by the time they get around again to the sampling point - "good" mixing between kicker and pick-up.




5) We have assumed complete correction of the errors.
How do real conditions modify this picture ?


To discuss this we should first see what the situation looks like in the frequency plane.

Bandwidth


With infinite bandwidth, the impulse from each particle will appear as a delta function recurring at $1/f_0$ seconds.




which will expand to an infinite series of integral harmonic frequencies.



Similarly, particles with slightly different revolution frequencies on either side will produce harmonics as shown.



The effect of finite bandwidth W is that we can now only observe time fluctuations in a spectral line of the order of $1/2W$ s, and we no longer look at a distribution along a transverse line as in the model, but have a finite sample length along the beam direction $\delta l = v/2W$, where v is the velocity of the particles (δl is about 15 cm for $W = 1$ GHz (10^9 Hz)).



We have lost resolution and expect this to slow down the cooling rate.

Recall (from other lectures):

$$1/T = W/2N \quad (2G - \text{incoherent term})$$

- where T is the cooling time,
- W the bandwidth,
- N the number of particles,
- G the overall gain or fractional correction.

Schottky bands as before, but they contain less precise information on the particle motion.

Note that if we use octave bands, i.e. from f to $2f$, then the bandwidth is f , and increasing the frequency will yield faster cooling for this reason alone (see below).


Question for the Student

Does finite pick-up length play a similar role in slowing down the cooling ?

Mixing

Pick-Up to Kicker

At first sight this will not hurt too much if the frequency slip lengthens the effective length of the sample ($v/2W$) by a small proportion (but see recent work by S. van der Meer¹ which is summarized in this School).




Kicker to Pick-Up

Complete mixing, or smearing out of the kick effects and renewal of the random population, producing again white noise at the pick-up, implies that there is no cumulative effect on a given particle from signals from other particles. Each particle in the sample length that can be resolved by the bandwidth will make a coherent contribution to the cooling, unaffected by the others.

This white noise condition is achieved by arranging the Schottky bands to fill the space between the harmonics, i.e. by the choice of momentum spread, frequency (the order of harmonic), and the machine dispersion, as was indicated on the sketch of momentum bands earlier.



At the other extreme of "no mixing", if we can imagine a machine in which a momentum spread produces no spread in revolution frequency, the kicker sample will return to the pick-up unchanged, apart from having had its mean corrected to zero, and the cooling process will stop.



For the in-between case of partial or "bad" mixing, a given particle will still make a contribution to the cooling, but now other particles close to its revolution frequency will turn up again and again on successive turns, and their signals can drive the resonant response of the circulating particle in an incoherent way, equivalent to a heating effect which combats the cooling. There is now an element of randomness in the cooling feedback akin to, and in addition to, the effect of amplifier noise.

Recall:

$$1/T = W/2N (2G - G^2 (M + u))$$

coherent term
heating term

mixing factor
M, = 1 for
complete mixing
amplifier noise/signal,
random

We note that if we raise the frequency of the cooling system, we improve the mixing through the widening of the Schottky bands and therefore improve the cooling rate, in addition to the benefit from the wider bandwidth mentioned above.

Incomplete Correction


In the presence of circuit noise and incomplete mixing between kicker and pick-up, the maximum cooling effect is obtained with a partial correction of the errors, that is the optimum gain is less than unity (differentiate above expression).

It should be noted that in the discussion so far we have been concerned with the spatial correction which we are applying to the beam for a given spatial error, and therefore the spatial sensitivity of the pick-up in say V/mm, and the efficacy of the kicker in say mm/V referred to the PU (or the equivalents in V/MeV, etc.) has to be included, along with the electronic gain, in order to arrive at the overall "gain".

We shall complete this discussion of basic concepts by considering some of the specific features of momentum and betatron cooling systems.

Momentum Cooling


Used to:



- pre-cool wide Δp from target.

- move particles into stack region for accumulation.

- cool stack to high density; portion of the stack core is periodically transferred to the collider. The cooled stack is in the regime of partial or "bad" mixing.




Error signals

Derived from:


- position sensitive electrodes as described ("Palmer method").

$p, f, r \rightarrow V$



Gain (slope of response) shaped to be zero at stack centre for stability and lower noise

- amplitude-to-frequency dependence imposed on sum signal by filter ("Thorndahl method").



Over a small frequency range, there is a linear change with δf , and a change of sign through zero (the filter method cannot be applied to overlapping Schottky bands).


Betatron Cooling

The first proposal for stochastic cooling concerned the damping of vertical betatron oscillations².

The principle is similar to that described above for cooling momentum spread, but it is applied to the incoherent betatron motion of particles in a machine.




The dipole moment "charge x displacement" expands to three frequencies. For one momentum we have:



where n is the harmonic number and q the fractional Q value of the machine.

We detect differences as before, but now there are two Schottky "sidebands" (and the revolution frequency band if we do not suppress it).




NB

Horizontal kicker does this \longleftrightarrow at $f_0(n \pm q)$

Momentum kicker does this \longleftrightarrow at nf_0

Correcting signals at sideband frequencies reduces the amplitudes at these frequencies, i.e. the betatron oscillation amplitudes.

Cooling momentum, and say, horizontal betatron motion does this:




Exercise for the Student

What do think will happen to the momentum cooling time if we increase the diameter of the ring? (Suggestion - invent several sets of constraints, including constant line charge or constant total charge in the ring, etc., and see what happens in these different cases). How would you approach the multi-parameter design problem?

HARDWARE

The problem: to transmit beam error signals (short impulses = strings of harmonics), highly amplified, with high fidelity, across the ring to meet the beam at the right moment as correcting kicks, with the overall gain, including the pick-up and kicker response as well as the electronic gain, corresponding to the optimum gain.



At present this is accomplished by analogue circuits.


General requirements:

- high gain - order of 120-150 db,
- additional delay added by components has to be kept within strict limits - order of 20-25 ns in the AA.
- low intermodulation distortion and independence of phase with amplitude.
- low circuit noise.

Pick-up Structures

Slot Couplers (L. Falin³)

These are a good solution where high sensitivity is not required, e.g. the stack core.




The coupler accumulates signal if the transmission along the coupling line is synchronous with the particle, and releases the sum of the amplitudes at the end as an output signal. The signal power varies as (frequency)².

Bandwidth is related to the degree of synchronism, which deteriorates if the slot size is increased, for example in order to increase the coupling (the coupling line is periodically loaded by the slots which slow down the signal propagation).

Slot couplers of the AA stack core system have slots of this form:


with vg line = 0.93 c
(AA \bar{p} 's have 0.96 c)
and bandwidth of 1-2 GHz.




Increasing the slot dimensions not only reduces the bandwidth, but also increases the loss of signal by re-radiation into the beam chamber. When the slot long dimension becomes equal to a half wavelength, it resonates, with heavy loading of the coupling line, and strong radiation into the beam chamber.


The sensitivity of a pick-up is expressed in various ways; for comparison of different kinds of structure it is convenient to use the coupling impedance, i.e. the ratio of the voltage output to the beam current.

For a pick-up, coupling slots can be arranged in an array of four slot boxes (as in the AA stack core system).



Summing the outputs of one pair of top and bottom boxes gives a spatial response  with horizontal position.


If then we difference left and right pairs we get this spatial response at the output of the difference network:



with the boxes spaced apart for low central gain for momentum cooling, and with the boxes close together to give high central gain for horizontal cooling.


For vertical cooling we can take the differences between the top and bottom slot rows, and sum the left and right differences.

In the AA stack core system there is one tank for the betatron cooling, and one for momentum, as shown.



Loops

These are tuned devices which are particularly useful where sensitivity is required. Loops have been used so far up to 500 MHz, sometimes with the incorporation of ferrites to increase the coupling, but development work towards the application of loops in the GHz region has been actively pursued recently.



The action of the wall loops can be thought of as the inducing of a signal as the particle passes the discontinuity at A, to which is added a signal as it passes B, the latter signal arriving back at A in phase with the first (two quarter wavelengths and a change of sign).


These loops are frequency sensitive, with an amplitude response varying as $\sin 2\pi L/\lambda$.

High Frequency Problems with Pick-Ups

- Short wavelengths are less forgiving of mechanical irregularities, which are a larger fraction of a wavelength than at lower frequencies. Reflections from imperfections can add their effects at certain frequencies and perturb the amplitude and phase transmission through the structure.
- A beam chamber wide enough to take the beam becomes a good waveguide at high frequencies, with the number of possible waveguide modes increasing with frequency (waveguide propagation starts at a wavelength equal to twice the largest dimension of the chamber); the beam can excite a choice of field patterns corresponding to these modes, resulting in irregular spatial responses with particle position. In the wide beam chambers which we are usually concerned with, the "TE" waveguide modes with a vertical "E" vector are the problem.

It is desirable to have but one mode of propagation, the TEM mode, as in a parallel plate transmission line, for which spatial responses are smooth and fairly simple to calculate.

This problem of the suppression of waveguide modes is a fairly common one in accelerators. The solution adopted in the AA stack core system is to make the pick-up beam chamber approach as closely as possible a parallel plate transmission line, by fitting the side walls with the impedance of free space, $120\pi \Omega$.




$120\pi \Omega$ ferrites backed by spaced conductors

NOTE: Even with the suppression of waveguide modes within the pick-up chamber, we can still be vulnerable to vertical E fields launched into the chamber from elsewhere, for example from waveguide modes at the revolution frequency in the adjoining chambers, excited by the interaction of beam and chamber wall discontinuities, or from other sources.

Kickers

Kickers for stochastic cooling are essentially pick-up structures working in reverse.

For example, the AA stack core kickers consist of 4 slot boxes arranged in a square. It is convenient to kick in a region where the beam is small in transverse dimensions, partly because the small kicker aperture economises in kicker power, but also because ideally one would like to have no kicker power escaping from the kicker and propagating through the beam chamber around to the pick-up position; one obvious way to prevent this is to make the kicker aperture sufficiently small so that there can be no waveguide mode propagation within the kicker structure itself, another is to damp with attenuators any propagation through the beam chamber from K to PU.



In the square array of the AA, one excites top and bottom slot boxes in the difference connection for vertical kicking, and the left and right for the horizontal (kick constant with frequency in difference mode).

For momentum kicking we exploit the feature of this slot coupler that "sum" connection produces a longitudinal kick (kick proportional to frequency in sum mode).


The loop kickers in the AA pre-cooling system also use ferrites to increase the coupling.

Typical power levels in kickers can be seen from the AA figures:

- max. 2 kW for the pre-cooling (fast momentum)
- approx. 500 W for the stack tail
- less than 10 W for each stack core system (Δ , V and H)

Summing and Differencing Networks

These can be obtained commercially as 4-port hybrid networks. A hybrid is a useful little passive device which can add and subtract signals, and divide by two as a splitter. It is not for nothing that its waveguide forerunner was known as a "Magic Tee".



In addition to the 3 dB loss from an input to an output port, there is a transmission loss in each path, typically around 0.3 dB for the 1-2 GHz band for example. The differencing behaviour comes from an added phase shift in the circuit which results in 180° from one input to the difference port, and 0° from the other. The circuit is usually of printed stripline.

As a practical point, it is not worth worrying too much over the sign of the cooling feedback since one needs only flip a hybrid over to reverse the sign.

Low Noise Amplifiers

Commercial solid-state amplifiers give reasonable noise figures, e.g. less than 2 dB over the 1-2 GHz octave.

Fast cooling of injected beams requires high gain, and output power is dominated by amplifier noise from the first stage.

Therefore, the next generation of cooling systems will require cryogenic cooling of amplifiers and input terminations, with noise temperatures around 20-40 K, in order to limit the noise power at the output.


The problem with low-noise GaAsFET devices is the the matching over a wide band. Development until recently has been for narrow band radio-astronomy.

For reference:
 NF dB = 10 log F

$$F = \frac{(\text{signal/noise})_{IN}}{(\text{signal/noise})_{OUT}}$$
 Noise temperature (K)
 $T_n = (F - 1)290 \text{ K}$
 Noise power out
 $P = F \cdot G \cdot k \cdot T_0 \cdot W$
 G power gain
 T_0 input temp. (°K)
 k Boltzmann const.

Power Amplifiers

These can be solid-state in the 100-500 MHz range. Above 1 GHz the Travelling Wave Tube (or TWT) Amplifier has until now been the obvious choice.



The TWT amplifies by interaction between the slow wave on a helix and an electron stream.

The TWT amplifiers are available in octave bands, 1-2 GHz, 2-4 GHz, etc., as well as in narrower bands developed for satellite communication. Over octave bands, powers can be obtained up to 3 kW CW.

The wide-band TWT's were mainly required for electronic counter-measures (jamming, etc.), where agility in working frequency demands large bandwidths.

Tight control of the amplitude and phase response of the TWT's has apparently not been a major requirement for military applications, and a ripple of several dB in amplitude and tens of degrees of phase is not uncommon.

In the GHz frequency range, solid-state replacements for TWT's are not yet above about 10 W CW for octave bands. Narrower band devices in combination might however become interesting in the future.

NOTE: A problem common to many high-frequency amplifiers is that the extrapolation of phase advance with frequency back to zero frequency does not always produce an intercept of 0 or 180 degrees, but sometimes an awkward value such as 70 or 130 degrees, which also varies between individuals of the same type. This makes selection necessary.

Filters

Required in the "Thorndahl" method of momentum cooling, and have also been used for the correction of sideband phases.

A transmission line of length L with total reflection at the far end will resonate with frequencies of the harmonic series. For example, if a line is short-circuited and lossless, the line has an input impedance:

$$Z_i = j \tan 2\pi L / \lambda ,$$

which has zeros repeating at


$$L = \lambda/2, 2\lambda/2, 3\lambda/2 \dots$$

i.e. at frequencies:

$$f = c/2L, 2c/2L, 3c/2L \dots$$


These zeros can be lined up with the harmonics of the revolution frequency, giving zero response at the harmonic, with the change of phase desired and an approximately linear response through the harmonic frequency.

The sharpness of the phase and amplitude response is affected by the losses in the line and by the coupling to the transmission line carrying the signal, in other words we are concerned with the loaded Q of the system.




The phase behaviour around the revolution frequency can be reversed by using an open-circuited line in place of the short circuit.

More flexible control of filter performance can be obtained with a variant known as a "transmission filter" or sometimes "correlation filter", which effectively sets up interference patterns, not between incident and reflected waves as in the "stub" filter described above, but between two transmission paths of different lengths, that is, an interferometer.



Filters can be conveniently set up by measuring the notches produced in the amplifier noise.

A recent development by C. Metzger uses three sets of fringes, and enables one to correct the errors in horizontal sideband phase (due to PU to K distance being not equal to odd quarter betatron wavelengths) without perturbing the particles moving into the stack.



Some of the problems with filters are:

- irregularity in the line or connectors produces errors in notch frequencies;
- dispersion, or wave propagation varying with frequency, causes a systematic variation in notch spacing

(The filter lines used to date have been either copper coaxial lines with dielectric spacers or foam insulated cables. Work is in progress in the Fermilab on small diameter (=1.6 mm) superconducting cable for filters; good notch precision has already been demonstrated on prototypes).

- filter lines at room temperature need to be well insulated, if not temperature controlled, to prevent length changes and drifting of the notch frequencies. Notch frequency precision of the order of 10^{-5} is sought.

SYSTEM MEASUREMENTS

Beam Feedback


Modulation of the beam by the kicker sends a wave to the pick-up, which modifies the pick-up signals^{1,4,5}.

The sign is such that:

- Δp Schottky signals should be decreased in amplitude at the distribution edges when the cooling loop is closed.

In the AA stack core the most noticeable effect is that the peak is higher with the loop closed.

- Transverse Schottky signals are decreased in amplitude with the loop closed.



This effect of transverse signal suppression is an extremely useful tool for checking and optimizing the delay and gain of the betatron cooling systems.

If no other instrumentation but a spectrum analyser is available, one can, with patience, get close to the best settings by re-touching delay and gain while opening and closing the loop, starting from a completely misadjusted system.


Moving for example the delay, one proceeds from "no effect" to "one sideband cooling and the other heating" at a mid-band harmonic, then "both cooling", and then one can check at the bottom and top of the band and retouch the delay until there is acceptable cooling over the whole band, then trim the gain.

Amplifier Gain

Optimization can be speeded up by setting gain and delay from measurements of amplitude and phase over the band, using as a reference previous results obtained from good cooling settings.


This is useful for checking the electronics when malfunctioning is suspected, or for trimming the system when a component is replaced by one whose performance is slightly different.

(In the AA stack core, amplitude and phase are measured from the pick-up output through the transmission line and amplifiers to the kicker output).



The measurement employs a microwave network analyser, which is a form of high frequency bridge.

A signal is split into two equal "test" and "reference" paths. The unknown, in this case the system from PU to K, is included in the "test" path, and the phase and amplitude of the "test" and "reference" transmission paths are compared and displayed on a screen.




The phase convention is that a test path which is electrically longer than the reference appears as a phase trace sloping down towards the high frequencies, and vice versa.

Beam Transfer Function

This measurement is similar in principle to the "amplifier gain", but includes in addition the beam.


The feedback loop is opened to permit excitation of the beam by a signal generator feeding the kicker, and the result seen by the pick-up is carried back to the point at which the loop is opened.



The network analyser measures the amplitude and phase transmission from A to B, and care is taken in the switching and balancing of cables to make A and B exactly in the same plane.

This is the most informative measurement which has been devised for cooling system analysis, since the "Beam Transfer Function" or "BTF" carries information about the whole system and its stability over the band.


Examining a given harmonic we find the amplitude and phase analogue traces to be typically as follows:



One can read off the system gain from the amplitude scale, and the phase at the nf_0 and $(n \pm q_0)f_0$ points, which should be 0° for momentum and $+180^\circ$ for the betatron sidebands.


Repeating this at other harmonics over the band reveals the awful truth about the amplitude and phase performance of the system. It also indicates whether the delay is in error, and by how much.

The BTF display is used mainly in the AA stack core system for this full-bandwidth diagnostic work rather than for the analysis of individual harmonics.



Taking this example of horizontal betatron sidebands, the group of a dozen or so sidebands on the left show a slight trend downwards in phase towards the higher frequencies (upper trace), which could then be corrected by shortening the variable delay.

We might then measure as carefully as possible from the expanded display on the right the phase and peak amplitudes at the $(n \pm q)f_0$ frequencies, at intervals across the band, and plot the results as shown (the curve demonstrates the effect of the Metzger filter on the horizontal sideband phases):



SYSTEM IMPERFECTIONS


- One of the time consuming problems encountered with the AA 1-2 GHz system was the coupling between the output at kicker power level, and the input at PU level. Coupling via the beam chamber as mentioned earlier was not in fact a problem, whereas external leakage through connectors, bad joints, and inadequate shielding of TWT amplifiers, and coupling in and out of the vacuum vessels via pumps, viewing ports, etc. all combined to make the output-input coupling a puzzling phenomenon. A sensitive indication of this type of problem is the BTF measurement showing phase and amplitude information in the absence of beam (this became known as the hardware transfer function or HTF).

A feature of the coupling was its rapid variation with frequency, which could suggest that one is dealing with a transmitting array of leakage sources at the kicker end and a receiving array at the pick-up end, with lobe intersections changing rapidly with frequency.

The cure, after experiments with microwave absorption materials, circuit paint, Al foil, etc., was fortuitous, and resulted from the covering of the ring with concrete shielding blocks, which acted as effective absorbers and scatterers of microwave radiation.

- Another related problem is that of coupling between cooling systems working in different frequency bands, for example by high power leakage saturating a low power amplifier outside its frequency range, leading to non-linear products within its working band.

In the AA, the AA stack core betatron PU originally sat in a region of high radial extension ($\alpha_p = 10$ m), rather close to the precooling kickers. Enough energy in the 150-500 MHz range from these kickers reached the 1-2 GHz PU (in the form of TE transmission with a vertical E field) to saturate the pre-amplifiers in their low frequency roll-off. This produced harmonic distortion peaks in the 1-2 GHz band.



This was partially cured by damping with ferrites and blocking the leakage paths in and out of sublimation pumps, and completely eliminated by installing a new betatron pick-up away from the offending kickers.

- The last problem we shall discuss is the suppression of the sum signal in a differencing system - the common mode rejection.

If a beam is accurately centred in position and the mechanical and electrical balance is good, then one should see very little signal at the revolution harmonic frequency.



This is of interest if the power amplifiers have not much reserve in hand, since additional power loading by the unwanted common mode signal could drive them into saturation.

The contribution of the hybrid units to the common mode signals is usually quite low .. of the order of 30 dB down - and cables and lines can be cut and measured quite accurately (for example electrical length measurements reproducible to tenths of a mm in 200 m).

The main problem is in assuring the mechanical and electrical symmetry of the pick-up structures. Vacuum feedthroughs for example have to be very well matched over the frequency band if their contribution to transmission unbalance is to be insignificant. As was mentioned earlier in connection with waveguide mode propagation in pick-ups, the beam itself can also feed revolution harmonic signals into a difference pick-up from excited discontinuities inside or outside the pick-up chamber and these have to be damped in some way.

CONCLUSIONS

In the course of this paper the reader may have remarked that a stochastic cooling system does not require very high powers or voltages; even with projected systems calling for microwave powers of the order of 10 kW in total, this will be split up into several 50 Ω circuits and the working voltages will never exceed a few hundred volts. Therefore voltage breakdown is not usually a problem, and apart from the final amplifier one is dealing with low level electronics, and reliability is good. Furthermore, the moderate power level means that there is no microwave radiation hazard - because of the very high electronic gain from the pick-up to the kicker, leakage from connectors has to be kept very low, in fact well below the levels of mW's per cm² typical of health hazard limits.

Finally, the general impression one gets from working with stochastic cooling is that although it is delicate in conception it is quite robust in practice. As mentioned earlier one can arrive at an approximation to optimum settings by watching the "open loop" and "closed loop" Schottky signals as parameters are varied. Similarly, if something drifts or is adjusted away from its correct value in error, it is likely that the beam will continue to be cooled, although it may take longer and may not reach a low value of momentum spread or betatron amplitude.

The real problems are in the system design stage, that is in the optimization of a large number of inter-related parameters, and in the details of RF and microwave design, where components have to be matched for high performance over a wide frequency band in order to assure that the correct amplitude and phase is applied to each harmonic within the band.

ACKNOWLEDGEMENTS

I am grateful to many members of the AA Group for reading drafts and offering useful suggestions.

REFERENCES

1. S. van der Meer, Optimum Gain and Phase for Stochastic Cooling Systems, CERN/PS-AA/83-48, 1983.
2. S. van der Meer, Stochastic Damping of Betatron Oscillations in the ISR, CERN/ISR-PO/72-31, 1972.
3. L. Falin, Slot-Type Pick-Up and Kicker for Stochastic Beam Cooling, Nucl. Inst. Meth. 148 (1978), pp. 449-455.
4. D. Möhl, G. Petrucci, L. Thorndahl, S. van der Meer, Physics and Technique of Stochastic Cooling, Phys. Rep. 58(2) 73-119(1980), Part. I.
5. D. Sacherer, Stochastic Cooling Theory, CERN/ISR-TH/78-11, 1978.