

STOCHASTIC COOLING AND THE ACCUMULATION OF ANTIPROTONS

Nobel lecture, 8 December, 1984

by

SIMON VAN DER MEER

CERN, CH- 1211 Geneva 23, Switzerland

1. A general outline of the $p\bar{p}$ project

The large project mentioned in the motivation of this year's Nobel award in physics includes in addition to the experiments proper described by C. Rubbia, the complex machinery for colliding high-energy protons and antiprotons (Fig. 1). Protons are accelerated to 26 GeV/c in the PS machine and are used to produce p 's in a copper target. An accumulator ring (AA) accepts a batch of these with momenta around 3.5 GeV/c every 2.4 s. After typically a day of accumulation, a large number of the accumulated \bar{p} 's ($\sim 10^{11}$) are extracted from the AA, reinjected into the PS, accelerated to 26 GeV/c and transferred to the large (2.2 km diameter) SPS ring. Just before, 26 GeV/c protons, also from the PS, have been injected in the opposite direction. Protons and antiprotons are then accelerated to high energy (270 or 310 GeV) and remain stored for many hours. They are bunched (in 3 bunches of about 4 ns duration each) so that collisions take place in six well-defined points around the SPS ring, in two of which experiments are located. The process is of a complexity that could only be mastered by the effort and devotion of several hundreds of people. Only a small part of it can be covered in this lecture, and I have chosen to speak about stochastic cooling, a method that is used to accumulate the antiprotons, and with which I have been closely associated.

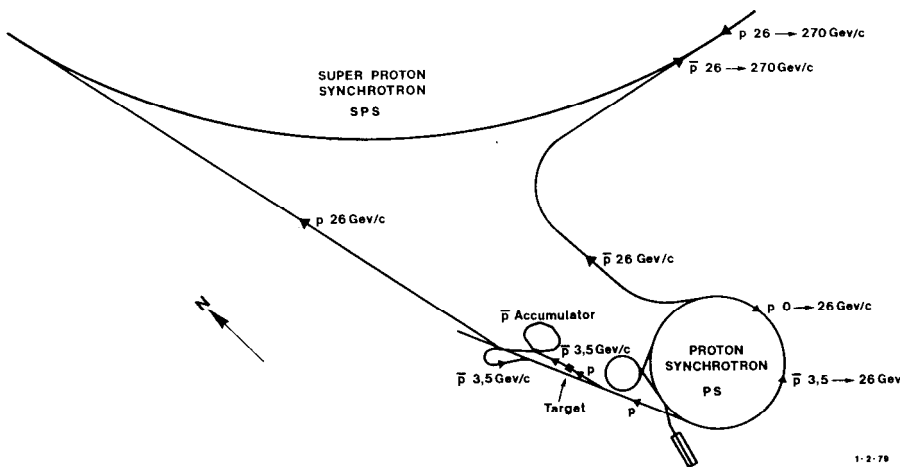


Fig. 1. Overall layout of the $p\bar{p}$ project.

2. Cooling, why and how?

A central notion in accelerator physics is phase space, well-known from other areas of physics. An accelerator or storage ring has an acceptance that is defined in terms of phase volume. The antiproton accumulator must catch many antiprotons coming from the target and therefore has a large acceptance; much larger than the SPS ring where the p's are finally stored. The phase volume must therefore be reduced and the particle density in phase space increased. On top of this, a large density increase is needed because of the requirement to accumulate many batches. In fact, the density in 6-dimensional phase space is boosted by a factor 10^9 in the AA machine.

This seems to violate Liouville's theorem that forbids any compression of phase volume by conservative forces such as the electromagnetic fields that are used by accelerator builders. In fact, all that can be done in treating particle beams is to distort the phase volume without changing the density anywhere.

Fortunately, there is a trick - and it consists of using the fact that particles are points in phase space with empty space in between. We may push each particle towards the centre of the distribution, squeezing the empty space outwards. The small-scale density is strictly conserved, but in a macroscopic sense the particle density increases. This process is called cooling because it reduces the movements of the particles with respect to each other.

Of course, we can only do this if we have information about the individual particle's position in phase space and if we can direct the pushing action against the individual particles. Without these two prerequisites, there would be no reason why particles rather than empty space would be pushed inwards. A stochastic cooling system therefore consists of a sensor (pick-up) that acquires electrical signals from the particles, and a so-called kicker that pushes the particles and that is excited by the amplified pick-up signals.

^P Such a system resembles Maxwell's demon, which is supposed to reduce the entropy of a gas by going through a very similar routine, violating the second law of thermodynamics in the process. It has been shown by Szilard¹ that the measurement performed by the demon implies an entropy increase that compensates any reduction of entropy in the gas. Moreover, in practical stochastic cooling systems, the kicker action is far from reversible; such systems are therefore even less devilish than the demon itself.

3. Qualitative description of betatron cooling

The cooling of a single particle circulating in a ring is particularly simple. Fig. 2 shows how it is done in the horizontal plane. (Horizontal, vertical and longitudinal cooling are usually decoupled.)

Under the influence of the focusing fields the particle executes betatron oscillations around its central orbit. At each passage of the particle a so-called differential pick-up provides a short pulse signal that is proportional to the distance of the particle from the central orbit. This is amplified and applied to the kicker, which will deflect the particle. If the distance between pick-up and kicker contains an odd number of quarter betatron wavelengths and if the gain is chosen correctly, any oscillation will be cancelled. The signal should arrive at




Fig. 2. Cooling of the horizontal betatron oscillation of a single particle




Fig. 3. Variation with system gain of the coherent cooling and incoherent heating effect

the kicker at the same time as the particle; because of delays in the cabling and amplifiers, the signal path must cut off a bend in the particle's trajectory.

In practice, there will not just be one particle, but a very large number (e.g. 10^6 or 10^{12}). It is clear that even with the fastest electronics their signals will overlap. Nevertheless, each particle's individual signal will still be there and take care of the cooling. However, we must now reduce the gain of the system because all the other particles whose signals overlap within one system response time will have a perturbing (heating) effect, as they will in general have a random phase with respect to each other. Fortunately, the perturbing effect is on average zero and it is only its second-order term that heats (i.e. increases the mean square of the amplitude). This is proportional to the square of the gain, whereas the cooling effect—each particle acting on itself—varies linearly with gain. As illustrated in Fig. 3, we may always choose the gain so that the cooling effect predominates.

4. Simplified analysis of transverse cooling

We shall now analyse the process sketched above in a somewhat approximative way, neglecting several effects that will be outlined later. The purpose is to get

some feeling about the possibilities without obscuring the picture by too much detail.

In the first place, we shall assume a system with constant gain over a bandwidth W and zero gain outside this band. A signal passed by such a system may be described completely in terms of $2W$ samples per unit time. If we have N particles in the ring and their revolution time is T , each sample will on average contain

$$N_s = N/2WT \quad (1)$$

particles. We may now consider the system from two viewpoints:

- a) we may look at each individual particle and combine the cooling by its own signal with the heating by the other particles,
- b) we may look at the samples as defined above and treat each sample as the single particle of Fig. 1; this is justified because the samples are just resolved by the system.

The two descriptions are equivalent and yield the same result. For the moment, we shall adopt b). Incidentally, the name "stochastic cooling" originated because from this viewpoint we treat a stochastic signal from random samples. However, viewpoint a) is more fundamental; cooling is not a stochastic process.

The pick-up detects the average position of each sample \bar{x} and the gain will be adjusted so that this is reduced to zero, so that for each particle x is changed into $x - \bar{x}$. Averaging over many random samples, we see that the mean square x^2 is changed into

$$\overline{(x - \bar{x})^2} = \overline{x^2} - \bar{x}^2.$$

Therefore, the decrement of x^2 per turn is $\bar{x}^2/x^2 = 1/N_s$, and the cooling rate (expressed as the inverse of cooling time) is $1/\tau = 1/N_s T$. In fact, we have to divide this by four. One factor 2 occurs because the betatron oscillation is not always maximum at the pick-up as shown in Fig. 2. Both at the pick-up and at the kicker we therefore lose by a factor equal to the sine of a random phase angle; the average of \sin^2 is $1/2$. Another factor 2 is needed because it is usual to define cooling rate in terms of amplitude rather than its square. So we have, using (1)

$$\frac{1}{\tau} = \frac{1}{4N_s T} = \frac{w}{2N}. \quad (2)$$

This result, although approximative, shows that stochastic cooling is not a practical technique for proton accelerators; for a typical accelerator $N \approx 10^{13}$, so that even with a bandwidth of several GHz the cooling would be much too slow compared to the repetition rate. In storage rings, however, the available time is longer and sometimes the intensity is lower, so that the technique may become useful.

5. Mixing and thermal noise

In deriving the cooling rate, we assumed that all samples have a random population, without correlation between successive turns. The main reason why the sample populations change is the spread in energy between the particles, which results in a revolution frequency spread. The particles overtake each other, and if the spread of revolution time is large compared to the sample duration, we speak of "good mixing"; in this case the derivation above is valid. In practice, it is rarely possible to achieve this ideal situation. In particular with strongly relativistic particles a large spread of revolution frequency can only be obtained by a large spread in orbit diameter; for a given aperture this reduces the momentum spread that is accepted by the machine.

We may see how bad mixing influences the cooling by replacing the correction \bar{x} in the derivation of the cooling rate by a smaller amount $g\bar{x}$. As a result we find in the same way

$$\frac{1}{\tau} = \frac{W}{2N} (2g - g^2). \quad (3)$$

Clearly, this is largest for $g = 1$.

It can be shown that the two terms correspond to the coherent, cooling effect (each particle cooled by its own signal) and the incoherent, heating effect from the other particles³. It is the second one that increases by bad mixing, because of the correlation between samples at successive turns. It may also increase if thermal noise is added to the signal (usually originating in the low-level amplifier attached to the pick-up). Thus, we may define a mixing factor M ($= 1$ for perfect mixing) and a thermal noise factor U (equal to the noise/ signal power ratio) and obtain

$$\frac{1}{\tau} = \frac{W}{2N} (2g - g^2(M + U)).$$

By optimizing g (now < 1) we find

$$\frac{1}{\tau} = \frac{W}{2N(M + U)} \quad (4)$$

6. Frequency domain analysis

This qualitative analysis may be made much more precise by considering the process from the frequency (instead of time) domain standpoint⁵.

Each particle produces in the pick-up (considered to be ideal) a delta-function signal at each passage. For a sum pick-up, where the signal is independent of the transverse position, the Fourier transformation into the frequency domain results in a contribution at each harmonic of the revolution frequency (Fig. 4) while for a difference pick-up the modulation by the betatron oscillation splits up each line into two components⁵. For a collection of many particles with slightly different revolution frequencies, these lines spread out into bands,

called Schottky bands because they represent the noise due to the finite number of charge carriers as described by Schottky⁶.

The width of these bands increases towards higher frequency. The total power is the same for each band. The power density is therefore lower for the wider bands at high frequency up to the point where they start to overlap; beyond this point the bands merge and their combined density is constant with frequency. This is illustrated in Fig. 5 for so-called longitudinal lines (from a sum pick-up).

The cooling process may now be seen as follows. Firstly, each particle will cool itself with its own (coherent) signal. This means that at the frequency of each of its Schottky lines the phase of the corresponding sine-wave signal must be correct at the kicker, so that the latter exerts its influence in the right direction. Secondly, the other particles produce an incoherent heating effect at

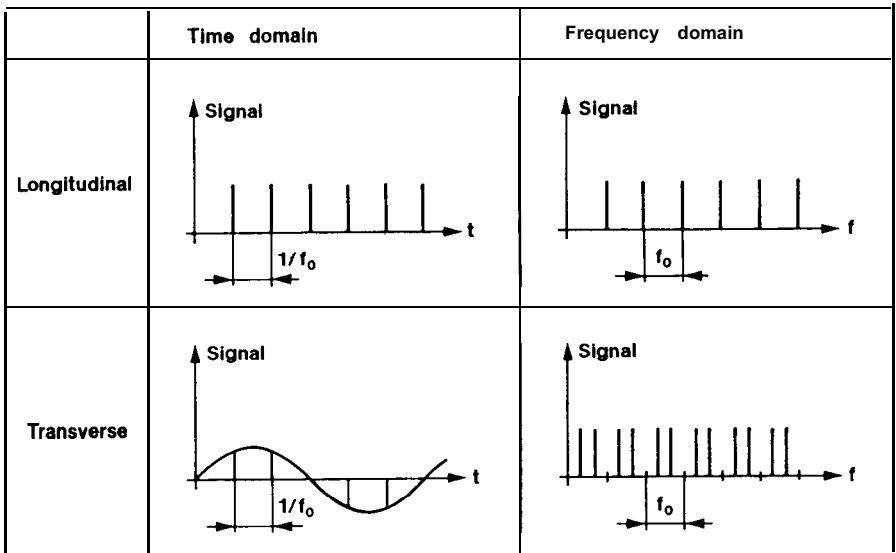


Fig. 4. Schottky signals in time domain and frequency domain



Fig. 5. Longitudinal Schottky bands originating from a large group of particles with slightly different revolution frequencies. At high frequencies the bands overlap.

each Schottky line proportional to the noise power density around that line⁷. Thus, only particles with frequencies very near to those of the perturbed particle will contribute. Any power density from thermal noise must of course be added to the Schottky power density.

For obtaining optimum cooling, the gain at each Schottky band should be adjusted so as to achieve an optimum balance between these two effects. If the bands are separated, the low-frequency ones have a higher density. This requires a lower gain and leads to less cooling for these bands. This is exactly the same effect that we called "bad mixing" in the time domain. At higher frequencies where the bands overlap we have good mixing and the gain should be independent of frequency.

Note that the picture given here (i.e. heating only caused by signals near the particle's Schottky frequencies) is completely different from the time-domain picture, where it seemed that particles in the same sample all contribute, independent of their exact revolution frequency. In fact, the latter is only true if the mixing is perfect and the samples are statistically independent. In the more general case, it turns out that both the optimum gain and the optimum cooling rate per line are inversely proportional to the density dN/df around that line, rather than to the total number of particles N . In the time domain treatment this was expressed by the mixing factor M , but the dependence of the parameters on frequency was lost.

There is yet another mixing effect that we have neglected so far. While moving from the pick-up to the kicker, each sample will already mix to a certain extent with its neighbours. This harmful effect may be described in the frequency domain as a phase lag increasing with frequency (particles with higher revolution frequency arrive too early at the kicker, so that their signal is too late). It appears quite difficult to correct this by means of filters at each Schottky band; on the other hand, in practical cases the effect is usually not very serious⁸.

7. Beam feedback

Another aspect that we have not yet considered is essential for the correct analysis of a cooling system. This is the feedback loop formed by the cooling chain together with the beam response (Fig. 6). Any signal on the kicker will

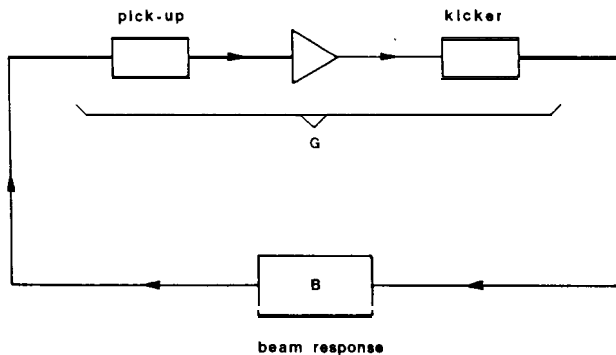


Fig. 6. Beam feedback effect. The loop is closed by the coherent beam response B




Fig. 7. Filter cooling

modulate the beam coherently (in position for a transverse kicker, in energy and density for a longitudinal one). The modulation is smoothed by mixing, but some of it will always remain at the pick-up, closing the feedback loop.

The beam response is a well-known effect from the theory of instabilities in accelerator rings. For cooling purposes, because the exciting and detecting points are separated in space^{5,9}, the treatment is slightly different. This is not the place to discuss the details; it may, however, be said that the response as a function of frequency can be calculated if the particle distribution versus revolution frequency is given, as well as some of the ring parameters.

It is found that for separated Schottky bands and with negligible thermal noise the optimum gain for cooling corresponds to an open-loop gain with an absolute value of unity and that the phase angle of the amplifier chain response must be opposite to the phase of the beam response*. As a result of this, it turns out that in the centre of the distribution the optimum loop gain becomes -1 for transverse cooling. The coherent feedback will then halve the amplitude of the Schottky signals as soon as the system is switched on. This is a convenient way of adjusting the gain; the correct phase may be checked by interrupting the loop somewhere and measuring its complete response with a network analyser¹⁰.

8. Longitudinal cooling

So far, I have mainly discussed transverse cooling, i.e. reducing the betatron oscillations. Longitudinal cooling reduces the energy spread and increases the longitudinal density. This process, as it turns out, is most important for accumulating antiprotons.

One method of longitudinal cooling (sometimes called "Palmer cooling"¹¹) is very similar to the one of Fig. 2. Again, we use a differential pick-up, now placed at a point where the dispersion is high, so that the particle position depends strongly on its momentum. The kicker must now give longitudinal kicks.




Fig. 8a. Simple transmission-line filter




Fig. 8b. Amplitude and phase response vs. frequency.

A different method is to use a sum pick-up (Fig. 7) and to discriminate between particles of different energy by inserting a filter into the system ("Thorndahl method"¹²). This works because the Schottky frequencies of particles with different energy are different; the filter must cause a phase change of 180° in the middle of each band, so that particles from both sides will be pushed towards the centre. Such a filter may be made by using transmission lines whose properties vary periodically with frequency. The simple filter of Fig. 8a may serve as an example. The line, shorted at the far end, behaves as a short-circuit at all resonant frequencies, which may be made to coincide with the centres of the Schottky bands. Just above these frequencies the line behaves as an inductance, just below as a capacitance; thus, the phase jump of 180° is achieved (Fig. 8b). For relativistic particles, the length of the line must be equal to half the ring's circumference. More complicated filters, using several lines and/or active feedback circuits may sometimes be useful¹⁰.

The advantage of the filter method, especially for low-intensity beams, is that the attenuation at the central frequencies is now obtained after the preamplifier, instead of before it as with a difference pick-up. The signal-to-noise ratio is therefore much better. Also, at frequencies below about 500 MHz where ferrites may be used, sum pick-ups may be made much shorter than differential ones, so that more may fit into the same space. This again gives a better signal-to-noise ratio. Of course, for filter cooling to be practical, the Schottky bands must be separated (bad mixing).




Fig. 9. Loop-type and ferrite ring-type pick-ups (or kickers). Note that for loop-type kickers the beam direction should be inverted.

9. Pick-ups and kickers

Cooling systems often have an octave bandwidth, with the highest frequency equal to twice the lowest one. Pick-ups with a reasonably flat response may consist of coupling loops that are a quarter wavelength long in the middle of the band (Fig. 9a). At the far end, a matching resistor equal to the characteristic impedance prevents reflections (or, seen in the frequency domain, ensures a correct phase relationship between beam and signal). Two loops at either side of the beam may be connected in common or differential mode for use as a sum or differential pick-up. The same structure may function as pick-up or kicker. Sum pick-ups or kickers may also consist of a ferrite frame with one or more coupling loops around it (Fig. 9b).

At high frequencies (typically > 1 GHz), slot-type pick-ups or kickers¹³ become interesting (Fig. 10). The field from the particles couples to the transmission line behind the slots. If the latter are shorter than $\lambda/2$, the coupling is weak and the contributions from each slot may all be added together, provided the velocity along the line is equal to the particle velocity.

The signal-to-noise ratio at the pick-ups may be improved by using many of these elements and adding their output power in matched combiner circuits. A




Fig. 10. Slot-type pick-up or kicker. One end of the transmission line is terminated with its own characteristic impedance.




Fig. 11. Density distribution vs. revolution frequency in the Antiproton Accumulator. On the right, the stack; on the left, the newly injected batch, before and after precooling.

further improvement may be obtained by cryogenic cooling of the matching resistors and/or the preamplifiers.

Using many kickers reduces the total power required. The available power is sometimes a limitation to the cooling rate that may be obtained.




Fig. 12. Inside of a vacuum tank with precooling kickers at the left and space for the stack at the right. The ferrite frames of the kickers are open in the centre of the picture; they can be closed by the ferrite slabs mounted on the shutter that rotates around a pivot at the far right. Water tubes for cooling the ferrite may be seen.




Fig. 13. Precooling $6 \times 10^6 \bar{p}$'s in 2 seconds. Longitudinal Schottky band at the 170th harmonic (314 MHz) before and after cooling.

10. Accumulation of antiprotons; stochastic stacking

It is now possible to explain how the antiproton accumulator works. It should, however, be made clear first that stochastic cooling is not the only method available for this purpose. In fact, already in 1966, Budker¹⁴ proposed a pp collider scheme where the cooling was to be done by his so-called electron cooling method. A cold electron beam superimposed on the \bar{p} beam cools it by electromagnetic interaction (scattering). We originally also planned to use this idea; it turns out, however, that it needs particles with low energy to work well with large-emittance beams. An additional ring to decelerate the antiprotons would then have been needed. The simpler stochastic method, using a single ring at fixed field was preferred.

In Fig. 11 we see how the particle density depends on revolution frequency (or energy, or position of the central orbit; the horizontal axis could represent any of these). On the right, the so-called stack, i.e. the particles that have already been accumulated. On the left, the low-density beam that is injected every 2.4 seconds. The latter is separated in position from the stack in those regions of the circumference where the dispersion of the lattice is large. In such a place the injection kicker can therefore inject these particles without kicking the stack. Also, the pick-ups and kickers used for the first cooling operation (longitudinal precooling) are placed here so that they do not see the stack. They consist, in fact, of ferrite frames surrounding the injected beam (Fig. 12). The pick-ups are therefore sum pick-ups (200 in total, each 25 mm long in

beam direction) and the Thorndahl type of cooling, with a filter, is used¹⁵. Figure 13 shows how the distribution is reduced in width by an order of magnitude within 2 seconds. The number of antiprotons involved is about 6×10^9 , the band used is 150-500 MHz.

After this precooling, one leg of the ferrite frames is moved downwards by a fast actuator mechanism¹⁶ so that the precooled beam can be bunched by RF and decelerated towards the low-frequency tail of the stack (Fig. 11). The whole process, including the upward movement of the "shutter" to restore the pick-ups and kickers, takes 400 ms. The RF is then slowly reduced¹⁷ so that the particles are debunched and deposited in the stack tail.

They must be removed from this place within the next 2.4 seconds because Liouville's theorem prevents the RF system from depositing the next batch at the same place without simultaneously removing what was there before. A further longitudinal cooling system, using the 250-500 MHz band, therefore pushes these particles towards higher revolution frequencies, up against the density gradient¹⁸.

This so-called stack tail system should have a gain that depends on energy (or revolution frequency). In fact, the density gradient increases strongly towards the stack core (note the logarithmic scale), and the gain for optimum cooling should vary inversely with this. We achieve this by using as pick-ups small quarter-wave coupling loops, positioned underneath and above the tail region, in such a place that they are sensitive to the extreme tail, but much less to the far-away dense core. This results in a bad signal-to-noise ratio for the region nearer to the core. Therefore, two sets of pick-ups are used, each at a different radial position and each with its own preamplifier and gain adjustment. With this set-up we obtain fast cooling at the stack edge where the particles are deposited, and slow cooling at the dense core, where we can afford it because the particles remain there for hours.

A problem is that the tail systems must be quite powerful to remove the particles fast enough. As a result, their kickers will also disturb the slowly-cooled stack core (the Schottky signals do not overlap with the core frequencies, but the thermal noise does). The problem exists because the kickers must be at a point where the dispersion is zero to prevent them from exciting horizontal betatron oscillations. They therefore kick all particles (tail or core) equally.

A solution is found by using transmission-line filters as described above to suppress the core frequencies in the tail cooling systems. These filters also rotate the phase near the core region in an undesirable way; this does not matter, however, because the cooling of the core is done by a third system of larger bandwidth (1-2 GHz).

While the particles move towards the core, they are also cooled horizontally and vertically, first by tail cooling systems, then by 1-2 GHz core systems. The layout of the various cooling circuits is shown in Fig. 14. In the general view of Fig. 15, some of the transmission lines transporting the signals for the pick-ups to the kickers may be seen.

When the stack contains a sufficient number of antiprotons (typically 2×10^{11}), a fraction of these ($\sim 30\%$) is transferred to the PS and from there to




Fig. 14. Plan of the AA ring with its 7 cooling systems. L = longitudinal, V = vertical, H = horizontal.

the SPS machine. This is done by bunching a part of the stack, of a width that may be adjusted by properly choosing the RF bucket area¹⁹. These are accelerated until they are on the same orbit where normally particles are injected. They can then be extracted without disturbing the remaining stack. This process is repeated (at present three times); each time one RF bucket of the SPS is filled. The remaining p's form the beginning of the next stack.

11. Design of longitudinal cooling systems; Fokker-Planck equation

The main difference between transverse and longitudinal cooling systems is that the latter will change the longitudinal distribution on which the incoherent (heating) term depends, as well as effects such as the beam feedback. This complicates the theory; still, everything can be calculated if all parameters are given.

It is convenient to define the flux Φ , i.e. the number of particles passing a certain energy (or frequency) value per unit time. It may be shown⁵ that

$$\Phi = F\Psi - D\delta\Psi/\delta f_0, \quad (5)$$

where Ψ is the density dN/df_0 , while F and D are slowly varying constants, depending on various system parameters as well as on the particle distribution. The first term represents the coherent cooling, the second one the incoherent (diffusion) effect that has the effect of pushing the particles down the gradient under the influence of perturbing noise.

By using the continuity equation

$$\delta\Psi/\delta t + \delta\Phi/\delta f_0 = 0,$$

expressing that no particles are lost, we find the Fokker-Planck-type equation

$$\frac{\delta\Psi}{\delta t} = -\frac{\delta}{\delta f_0} (F\Psi) + \frac{\delta}{\delta f_0} (D\frac{\delta\Psi}{\delta f_0}) \quad (6)$$

that allows us to compute the evolution of the density versus revolution frequency f_0 and time given the initial distribution. The particles deposited at the edge are introduced as a given flux at that point.

The constants F and D depend on many system parameters (pick-up and kicker characteristics, amplifier gain, filter response, beam distribution, etc.). Their value is found through summing the contributions of all Schottky bands. Analytic solutions of (6) do not exist in practice and a complicated numerical treatment is indicated.

Such calculations resulted in the design of the antiproton stacking system. At the time this was done, tests in a small experimental ring (ICE) had confirmed the cooling in all planes at time scales of the order of 10 seconds. However, it was not possible to check the stacking system (increasing the density by four orders of magnitude) in any way, and it may be argued that we took a certain risk by starting the project without being able to verify this aspect. Fortunately,



Fig. 15. View of the Antiproton Accumulator before it was covered by concrete slabs. The silvered material around the vacuum tanks is insulation, needed because everything may be heated to 300° to obtain ultra-high vacuum. The transmission lines crossing the ring and carrying the cooling signals may be seen.

everything behaved according to theory and although the number of p 's injected is smaller than was hoped for by a factor 3.5, the cooling works largely as expected.

12. Other applications of stochastic cooling; future developments

At present, stochastic cooling is used at CERN in the \bar{p} accumulator and in the low energy ring (LEAR) where the p 's may be stored after deceleration in the PS. Before the intersecting storage rings (ISR) were closed down last year, they also used the antiprotons and contained cooling equipment.

In the SPS where the high-energy collisions take place, cooling would be attractive because it would improve the beam lifetime and might decrease its cross-section. However, a difficulty is formed by the fact that the beam is bunched in this machine; the bunches are narrow (3×4 ns). In fact, owing to the bunching each Schottky band is split up into narrow, dense satellite bands and the signals from different bands are correlated". Nevertheless, a scheme is being considered that might improve the lifetime to a certain extent".

In the United States, a \bar{p} accumulator complex similar to the CERN one and also using stochastic cooling is being constructed²². This machine is expected to have a stacking rate an order of magnitude higher than the CERN one because it uses a higher primary energy to produce the antiprotons and higher frequencies to cool them. In the meantime, we are building a second ring at CERN,

surrounding the present accumulator (Fig. 16), with a similar performance. It will have stronger focusing, so increasing both transverse acceptances by at least a factor 2, and the longitudinal one by a factor 4. The increased focusing strengths will diminish the mixing; consequently, higher frequencies (up to 4 GHz) will be used for cooling. The present AA will be used to contain the stack and its cooling systems will also be upgraded.




Fig. 16. The new ACOL ring (under construction) around the AA. This ring will increase the stacking rate by an order of magnitude. The stack will still be kept in the AA ring.

ACKNOWLEDGEMENTS

The development of the stochastic cooling theory owes much to H. G. Hereward, D. Möhl, F. Sacherer, and L. Thorndahl. The latter also made important contributions to the construction of most cooling systems at CERN and it is doubtful if the accumulator would have been feasible without his invention of the filter method. It is also a pleasure to acknowledge the invaluable contributions of G. Carron (hardware), L. Faltin (slot pick-ups) and C. Taylor (1-2 GHz systems).

During the construction of the Antiproton Accumulator, R. Billinge was joint project leader with myself and it is mostly because of his contributions to the design and his able management that the machine was ready in a record time (2 years) and worked so well.

REFERENCES

1. L. Szilard, Gber die Entropieverminderung in &cm thermodynamischen System bei Eingriff-
fin intelligenter Wesen, Zcitsch. f. Physik 53 (1929) 840.
2. S. van der Meer, Stochastic damping of betatron oscillations in the ISR, CERN/ISR-PO/ 72-
31 (1972).
3. D. Möhl, Stochastic cooling for beginners, Proceedings of the CERN Accelerator School on
Antiprotons for Colliding Beam Facilities, CERN 84- 15 (1984).
4. F. Sacherer, Stochastic cooling theory, CERN/ISR-TH/78- 11 (1978).
5. D. Möhl, G. Petrucci, L. Thorndahl, S. van der Merr, Physics and technique of stochastic
cooling, Phys. Reports 58 (1980) 73.
6. W. Schottky, Ann. Physik 57 (1918) 541.
7. H. G. Hereward, The elementary theory of Landau damping, CERN 65-20 (1965).
8. S. van der Meer, Optimum gain and phase for stochastic cooling systems, CERN/PS-AA/83-
48 (1983).
9. S. van der Meer. A different formulation of the longitudinal and transvrse beam response,
CERN/PS-AA/80-4 (1980).
10. S. van der Meer, Stochastic cooling in the Antiproton Accumulator, IEEE Trans. Nucl. Sci.
NS28 (1981) 1994.
11. R. B. Palmer, BNL, Private communication (1975).
12. G. Carron, L. Thorndahl, Stochastic cooling of momentum spread by filter techniques,
CERN/ISR-RF/78-12 (1978).
13. L. Faltin, Slot-type pick-up and kicker for stochastic beam cooling, Nucl. Instrum. Methods
148 (1978) 449.
14. G. I. Budker, Proc. Int. Symp. on Electron and Positron Storage Rings, Saclay, 1966, p. 11 -I-
1; Atom. Energ. 22 (1967) 346.
G. I. Budker et al., Experimental studies of electron cooling, Part. Acc. 7 (1976) 197.
15. S. van der Meer, Precooking in the Antiproton Accumulator, CERN/PS-AA/78-26 (1978).
16. D. C. Fiander, S. Milner, P. Pearce, A. Poncet, The Antiproton Accumulator shutters: design,
technology and performance, CERN/PS/84-23 (1984).
17. R. Johnson, S. van der Meer, F. Pedersen, G. Shering, Computer control of RF manipulations
in the CERN Antiproton Accumulator, IEEE Trans. Nucl. Sci. NS-30 (1983) 2290.
18. S. van der Meer, Stochastic stacking in the Antiproton Accumulator, CERN/PS-AA/78-22
(1978).
19. R. Johnson, S. van der Meer, F. Pedersen, Measuring and manipulating an accumulated stack
of antiprotons in the CERN Antiproton Accumulator, IEEE Trans. Nucl. Sci. NS-30 (1983)
2123.
20. H. Herr, D. Möhl, Bunched bram stochastic cooling, CERN/EP/Note 79-34 (1979).
21. D. Boussard, S. Chattopadhyay, G. Dôme, T. Linnecar, Feasibility study of stochastic cooling
of bunches in the SPS, CERN/SPS/84-4 (1984).
22. Design Report Tevatron I project, Fermi National Accelerator Laboratory, Batavia, 111. (1983).