CERN 88-01
11 March 1988

ORGANISATION EUROPÉENNE POUR LA RECHERCHE NUCLÉAIRE
CERN EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

CAS CERN ACCELERATOR SCHOOL

Third John Adams Memorial Lecture

THE PROTON-ANTIPROTON COLLIDER
Lecture delivered at CERN on 25 November 1987

Lyndon Evans

GENEVA
1988

ABSTRACT

The subject of this lecture is the CERN Proton–Antiproton ($p\bar{p}$) Collider, in which John Adams was intimately involved at the design, development, and construction stages. Its history is traced from the original proposal in 1966, to the first $p\bar{p}$ collisions in the Super Proton Synchrotron (SPS) in 1981, and to the present time with drastically improved performance. This project led to the discovery of the intermediate vector boson in 1983 and produced one of the most exciting and productive physics periods in CERN's history.

RN — Service d'Information Scientifique-RD/757-3000-Mars 1988

CONTENTS

1. INTRODUCTION

The subject of this lecture is the CERN Proton–Antiproton ($p\bar{p}$) Collider and let me say straightaway that some of you might find that this talk is rather polarized towards the Super Proton Synchrotron (SPS). In fact, the $p\bar{p}$ project was a CERN-wide collaboration 'par excellence'. However, in addition to John Adams' close involvement with the SPS there is a natural tendency for me to remain within my own domain of competence.

Beam cooling is central to the concept of $p\bar{p}$ colliding beams and, therefore, I will start by giving a very brief and sketchy account of cooling techniques. This should be treated as no more than a definition of terms. Then, I will try to trace the history of the project from the early days. Finally, I will discuss some specific SPS problems and future plans.

When I was preparing this lecture it became even more clear to me that the Collider project was only made possible by the expertise built up inside CERN over the past 20 years or so. In fact, I could not think of one single major development, apart from the invention of the alternating gradient synchrotron itself, that was imported from outside. I will try to enlarge on this conjecture in this lecture.

At CERN we are frequently criticized for doing things too well. The term 'gold-plated engineering' was even used in a recent article in 'Nature'. When John Adams built the SPS, the possibility that it would ever operate as a storage ring was very remote indeed. Nevertheless, he built the machine in such a way that when the idea of colliding beams came along, we found ourselves in the unique position that it was possible, indeed relatively easy, to upgrade the machine to a storage ring. This was particularly true in the three important areas of magnetic field quality, vacuum, and reliability.

The idea of turning a synchrotron into a colliding-beam device by confronting hadrons with their antiparticles travelling in the opposite direction was clearly an attractive one and its feasibility had been proved in several electron–positron (e^+e^-) colliders already in existence. However, antiprotons are very much more difficult than positrons to produce in any useful quantity.

2. BEAM COOLING

In the first part of this lecture, I will try to outline the history of the Collider project from the early days. In doing so, I will frequently refer to beam 'cooling'. Cooling is fundamental to the production of antiproton beams of any useful intensity. This is simply because antiprotons just do not grow on trees! They must be created by using a high-energy accelerator to collide protons with a target and then accumulate them for many hours in a storage ring. They are created with a wide range of angles and energy spread which could never be accommodated within the 'acceptance' of a machine such as the SPS. Cooling allows the beam dimensions to be reduced or the phase-space density to be increased by many orders of magnitude.

There is a theorem by Liouville, for a long time the dictum of accelerator physics, which states that for a *continuous fluid* under the influence of *conservative forces* the density of phase space is invariant. Two methods of cooling attack separately the two constraints in italics above. Stochastic cooling relies on the fact that a particle beam, even though it contains very many particles, is not a continuous fluid. Electron cooling, which was the first method proposed, though not the first to be made to work, attacks the problem using non-conservative, i.e. collisional, forces.

The basic idea of stochastic betatron cooling is explained in Fig. 1. A beam pick-up detects random fluctuations in the centre of gravity of a sample of a beam circulating in a storage ring. The signal is amplified and sent to a kicker, where it arrives at the same time as the sample and corrects the error. Simon van der Meer [1] was the first to show that continuously correcting the *mean* of a sample slowly reduces the *sigma* or beam size—not intuitively obvious. In fact, the example as it is drawn, would not work at all. A fundamental requirement is that the statistical fluctuation is maintained between pick-up and kicker, and is 'smeared out' between kicker and pick-up again, to produce 'good mixing' so that the statistical distribution is continuously re-established. This means that the distance travelled from pick-up to kicker is generally much shorter than that from kicker to pick-up.

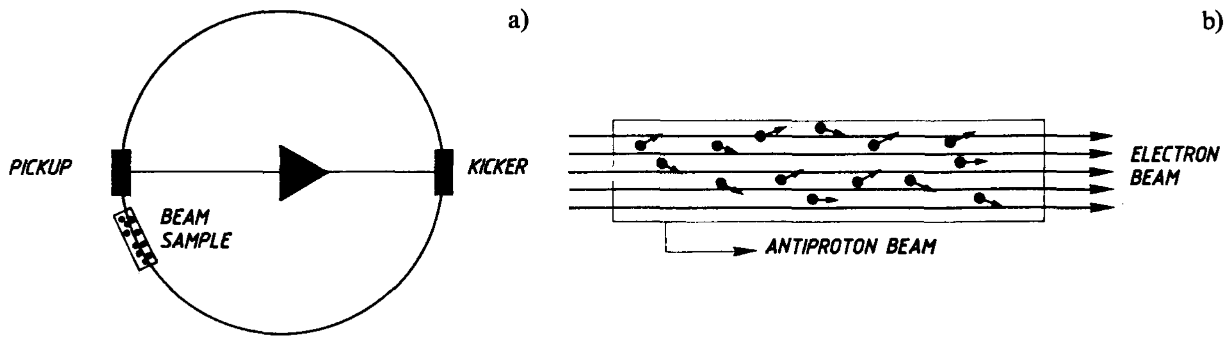

Fig. 1 (a) Stochastic cooling and (b) electron cooling

Electron cooling works by transferring energy from an antiproton beam circulating in a storage ring to a 'cold' electron beam travelling in synchronism with the antiproton beam over part of its path. The electron beam is continuously refreshed from an electron gun. Electron cooling works best at very low energy, whereas for the stochastic method the cooling rate is independent of energy.

A very important development in stochastic cooling was fast precooling of momentum spread by the filter method (Thorndahl cooling [2]). However, to explain this, a more detailed understanding of the nature of the signals desired from the beam is required and, although it proved to be the key element in the development of a colliding-beam facility based entirely on stochastic cooling, I will not discuss it here in detail.

3. HISTORICAL DEVELOPMENT OF THE $p\bar{p}$ COLLIDER CONCEPT

The first proposal for a $p\bar{p}$ colliding-beam device seems to have been made by Budker and Skrinskij at Orsay in 1966. This proposal was based on Budker's new idea of electron cooling.

Then in 1968, Simon van der Meer had the idea of stochastic betatron cooling but was so unsure of its usefulness that he did not even bother to publish it. Only in 1972, after Schottky signals demonstrating the discrete particle nature of the beam were observed and interpreted at the Intersecting Storage Rings (ISR), did he decide to publish his idea [1], although even then it seemed a little 'far-fetched' to him, as can be seen from his concluding remarks:

> 'This work was done in 1968. The idea seemed too far-fetched at the time to justify publication. However, the fluctuations upon which the system is based were experimentally observed recently. Although it may still be unlikely that useful damping could be achieved in practice, it seems useful now to present at least some quantitative estimation of the effect'.

Fortunately, it did not seem so far-fetched to Wolfgang Schnell at the ISR, who, taking the idea quite seriously, did a design study with the aim of trying it out at the ISR [3].

In October 1974 in fact, stochastic cooling was first observed at the ISR by Schnell, Thorndahl, and others [4].

Meanwhile, another CERN development gave important experimental support to the electroweak theory, the discovery of neutral currents, at Gargamelle in 1973, giving much more confidence that the mass of the intermediate bosons could be within the energy range of the Collider. This provoked Carlo Rubbia and collaborators to propose a colliding-beam experiment both at the SPS and at FNAL [5].

As a result, John Adams decided to bring together a small team to make an initial design study for an antiproton source. It is interesting to note that, even though stochastic cooling had already been demonstrated at that time, it was considered safest to base the design simply on electron cooling, as can be seen from the concluding remarks of the first design report.

'It would be most attractive to replace the entire electron cooling and stacking process by a stochastic cooling system. No deceleration would then be required, and a single ring, operated d.c., would be sufficient.

....

'However, the problem of stacking, say, 10^4 pulses and simultaneously cooling the injected pulses in the same ring, has not yet been solved. Although a solution is not quite excluded, it seems more reasonable at present to base the design on electron cooling, which requires much less extrapolation into unknown regions.'

One consequence of electron cooling is that it must be done at very low energy because of the very strong ($\gamma^{5/2}$) dependence of cooling time on energy. On the other hand, the best method of producing a sufficient antiproton flux at CERN is by colliding 26 GeV protons with a target, giving an optimum production energy of 3.5 GeV to the resulting antiprotons. Therefore, the first design study was based on a two-ring device; the first for bunch rotation followed by deceleration from 3.5 GeV to 100 MeV, and the second for electron cooling of the low-energy beam (Fig. 2).

It is remarkable how similar topologically this two-ring scheme is to the present Antiproton Accumulator/Antiproton Collector (AA/ACOL) complex, but at the time it was considered to be too complicated and too expensive to pursue. Stochastic cooling and stacking in a single ring at the $\bar{p}$ production energy of 3.5 GeV/c was clearly attractive but it was not known how to do this at the time. In fact, the design report was never published.

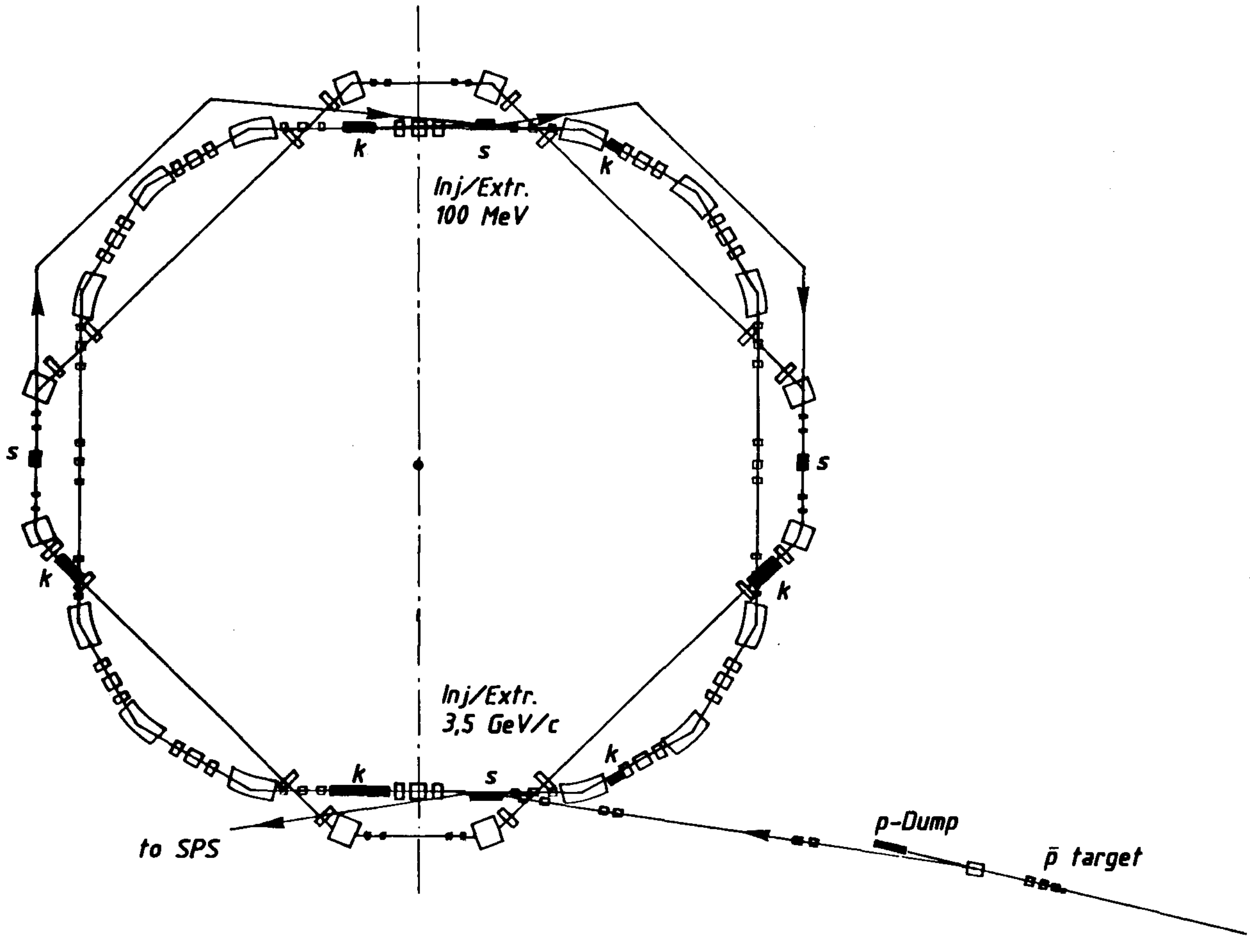

Fig. 2 The first proposal for an antiproton source consisted of a 3.5 GeV ring (shuttle ring) for bunch rotation and deceleration and a 100 MeV cooling ring using electron cooling

Another decision taken at that time was to convert the old muon storage ring into a 1.7 GeV/c storage ring in order to pursue the studies of stochastic cooling. This ring was named ICE (Initial Cooling Experiment).

The next real breakthrough came when Lars Thorndahl invented the filter method [2] for fast momentum cooling which allowed the momentum spread of low-intensity pulses to be reduced to a manageable value within about a second, between two PS pulses. This method was successfully tested (in ICE) in 1977. This, together with substantial developments in the theory of stochastic cooling, was the key to the implementation of an antiproton source based completely on this process.

In the same year, experiments started at the SPS to convert the machine from a pulsed synchrotron into a storage ring. In parallel, a second design report [6] was quickly put together and things moved so fast that by the summer of 1978 the $p\bar{p}$ project was approved by the CERN Council.

The construction of the AA started towards the end of that year and by June 1980 protons were already circulating in it!

One substantial change after the design report was published was the choice of injection energy into the SPS. Experiments had shown that injection of dense single bunches at low energy produced tremendous space-charge problems. In addition, transition energy had to be crossed around 23 GeV. At low intensity this creates no great problems but for colliding-beam operation the single-bunch intensity is more than a factor of 20 higher, and it was shown that fast longitudinal (negative mass) and transverse (head-tail) instabilities limited the achievable bunch intensity to about a factor of 5 lower than the required intensity. At the time there were two conflicting schools of thought: to try to combat the stability problems in the SPS or to transfer the problem to the PS. Consequently, after a study of the various alternatives, John Adams decided to use the PS (which is much better suited for acceleration of low-energy, high-intensity beams) to accelerate the antiprotons from 3.5 to 26 GeV/c.

A small aside at this point. One suggestion was to locate the AA in the South Hall of the PS. One objection to this was that the hall was too small, which was correct but not an insurmountable problem. The other objection was voiced, not very loudly, by a small group of people working on a small storage ring [the Low Energy Antiproton Ring (LEAR)], to be located in that hall and which turned out to be one of the other great successes of the $p\bar{p}$ project.

In 1980 the SPS was shut down for 11 months in order to complete the substantial civil-engineering modifications required for colliding-beam operation (UA1, UA2, and TT70) and for the required modifications to the machine, vacuum improvements, low-β insertions, etc.

Then in July 1981, the first $p\bar{p}$ collisions were observed in the SPS. Towards the end of that year, the first technical run took place in which 0.2 nb^{-1} of integrated luminosity was obtained for the whole period, with something like 700 h scheduled and 140 h realized (we now get that in about 10 min of colliding-beam operation).

In autumn 1982 the first real physics runs took place where 28 nb^{-1} of luminosity was produced and the first W candidates were found. In 1983, the integrated luminosity was increased to 153 nb^{-1} and the Z^0 particle was seen. Since then the performance of the complex has been steadily improved.

4. THE SPS AS A COLLIDER

4.1 Operation

This is how the whole thing works. After stacking antiprotons for many hours, the PS and SPS prepare for a fill. Firstly, three proton bunches are sent to the SPS at 2.4 s intervals and stored at 26 GeV. These three bunches are placed azimuthally symmetrically around the SPS circumference. They have to be specially 'shaped' in the PS by a complex RF procedure until they fit into a single 5 ns long SPS 'bucket'.

Then a single antiproton bunch is unstacked from the AA and injected into the PS at 3.5 GeV, accelerated to 26 GeV, and injected into the SPS on the correct azimuthal position to within a fraction of a nanosecond for collision in the detectors. This is followed at 2.4 s intervals by two more antiproton bunches. The machine then accelerates both the proton and antiproton beams to storage energy (315 GeV), squeezes them at the low-β insertions (30 quadrupoles) in a second or so, and keeps the two beams in storage for up to 24 h whilst the next batch of antiprotons is being prepared in the AA.

4.2 Performance

Table 1 shows the evolution of the Collider performance over the years 1982 to 1985. The 1986 run is not included because it was cut short due to a major problem with the UA1 detector.

It can be seen that most of the parameters of the design report, apart from the antiproton bunch intensity, have either been achieved or exceeded. The beam energy was increased in 1984 from 270 to 315 GeV by installing extra water pumps in the tunnel. Both the low-β insertions and the proton-bunch intensity have exceeded their design values by a substantial amount. The number of bunches per beam is limited to three by the scarcity of antiprotons, although six-bunch operation has been achieved in machine studies.

Overall, the best achieved peak luminosity is about a factor of 3 lower than the design report. In the end, it is the *integrated* luminosity that counts, and this depends on the luminosity lifetime, which is now in excess of 24 h.

Table 1
Evolution of Collider performance 1982 to 1985
compared with the design report (DR)

	DR	1982	1983	1984	1985
Beam energy (GeV)	270	273	273	315	315
β_H	2.0	1.5	1.3	1.0	1.0
β_V	1.0	0.75	0.65	0.5	0.5
N_p (10^{11})	1.0	1.0	1.5	1.6	1.6
$N_{\bar{p}}$ (10^{11})	1.0	0.1	0.1	0.2	0.2
Number of bunches	6	3	3	3	3
Integrated luminosity (cm^{-2} s^{-1})					
average per coast		0.5	2.1	5.3	8.2
average per day		0.4	1.8	5.1	5.8
per year		28	153	395	655
Luminosity (10^{29} cm^{-2} s^{-1})					
peak	10.0	0.5	1.7	3.5	3.9
average per coast		0.1	0.5	1.0	1.3
Hours scheduled		1750	2064	2136	2688
Hours realized		748	889	1065	1358
Number of coasts		56	72	77	80
Average coast duration (h)		13	12	15	17
Percentage of coasts terminated by faults		41	40	32	18

5. IMPACT OF EARLIER CERN DEVELOPMENTS

I should now like to talk about some of the earlier developments at CERN that contributed to the success of the Collider.

The most obvious of these is the impact of the ISR. First of all, the main Collider experiments themselves were designed by a nucleus of physicists from the ISR. It was there that they learned to do physics in the difficult conditions of colliding hadron beams and to build the right kind of detectors.

In the areas of accelerator physics and technology the impact of the ISR has also been tremendous. For example, the vacuum systems of both the AA and the SPS were built by people who obtained experience of ultra-high vacuum on an industrial scale in the extremely demanding environment of the ISR. The ISR also taught us a great deal about the behaviour of beams in hadron storage rings, in particular beam stability, bunched beam lifetime due to RF noise, effect of neutralization, and many other phenomena.

However, it is in the development of stochastic cooling, both experimentally and theoretically, that the ISR accelerator physicists obviously made their greatest contribution to the Collider.

I have already mentioned the fact that the ICE project gave the vital input needed to base the whole scenario entirely on stochastic cooling, especially the experimental verification of Thorndahl cooling. By the way, at CERN electron cooled beams were obtained for the first time in ICE in 1979.

The PS plays a key role in the Collider scheme. In particular, some of the manipulations required to produce beams for antiproton production and for the Collider are at the very limits of stability. They are only made possible by the knowledge and understanding built up by many years of development of that machine. No other machine in the world to my knowledge manipulates particle beams with such dexterity.

Finally, I have cited the production and focusing of secondary beams. The antiproton beam for the AA is obtained from a target which produces both a large angular and a large momentum spread. The orthodox way to focus the beam from the target is to use a conventional 'quadrupole doublet'. However, conventional methods prove to be totally inadequate for the required transverse and longitudinal acceptance and depth of focus. Instead a 'horn' is used whose working principle is shown in Fig. 3. A pulsed current of about 160 kA passes through the horn material. Particles inside the inner conductor of the horn see no magnetic field, whereas those between the two conductors experience a poloidal field identical to that of a line current on the horn axis, which deflects the particle. This field falls off as $1/r$, r being the radial distance from the axis. So if the time the particle spends between the two conductors is proportional to r^2, meaning that the inner surface is a parabola, the net integral of the field and distance $B d\ell$ the deflecting angle is proportional to the radius, i.e. the horn acts as a linear lens. This lens has a better angular and momentum acceptance and depth of focus than any conventional quadrupole system and it is at least a factor of 2 more efficient. I mention this specially because it was invented at CERN in 1962 by a clever young man named Simon van der Meer [7].

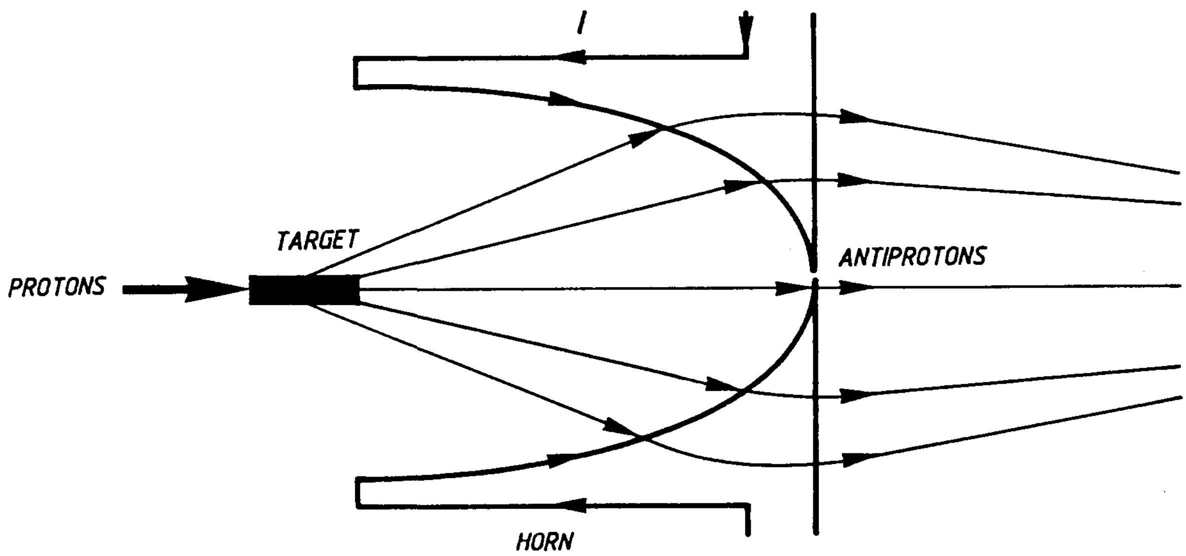

Fig. 3 Focusing horn for antiproton production

6. PROPHETS OF DOOM

Throughout the history of the development of particle accelerators there has always seemed to be a breed of pessimists who can invariably prove that a machine will never work. Way back in the design stage of the PS it was predicted that the non-linearities in the guide field would render particle orbits unstable. Indeed, John Adams and Mervyn Hine pioneered the use of digital computers at Harwell in the mid-1950's to investigate these effects by numerical simulation, a technique very much in use to this day. Before the ISR came on line there were predictions that the machine would not work for much the same reasons, fortunately unfounded. Even around 1977 a well respected group proved conclusively that it was impossible to cool simultaneously in all three phase planes. As it happens, ICE was performing just that impossible task at the same time.

However, it was in the SPS Collider that the prophets got closest to the mark. The problem has yet again to do with non-linear fields, but this time they are generated by the colliding beams and are much stronger than those generated by any reasonably well designed magnets. But it is not only the non-linearity that gives difficulties. Particles are subjected to *non-linear periodic* kicks as they encounter the opposing bunches every revolution and these give rise to non-linear resonances in the beams.

Non-linear resonances have been the plague of particle accelerators ever since their conception, but generally resonances of order higher than 5 are harmless. This means that the forbidden tune values are limited to fairly narrow bands and there is plenty of space between for the beam to sit without being affected.

The beam-beam interaction produces non-linear resonances of much higher order and severely limits the performance.

All existing e^+e^- colliders are limited in their performance by this beam-beam interaction. The question was how to extrapolate to the $p\bar{p}$ Collider, the first hadron bunched-beam collider, very much like an e^+e^- machine but without the advantage of synchrotron radiation damping. For once, the ISR could not tell us the answer since it worked in a completely different regime.

Figure 4 shows the effect of the beam-beam interaction in the SPS. On the left, a chart recorder output of the intensity of a strong proton bunch and a weak antiproton bunch is shown on a sensitive scale. On the right is a small part of the tune diagram in the vicinity of the operating point of the SPS. Good lifetime is only achieved when the beam is free of 10th order resonances. Notice that the proton bunch is unaffected by the tune changes because it does not experience any substantial beam-beam effect from the much weaker antiproton bunch.

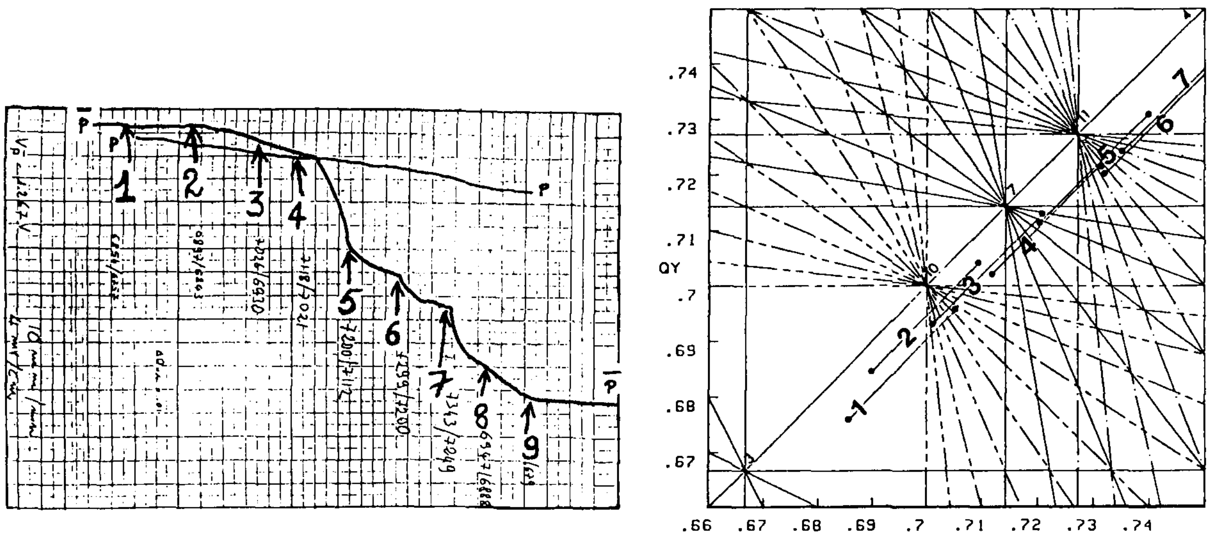

a) b)

Fig. 4 The beam-beam effect in the SPS. (a) A chart recorder output of the intensity decay (zero suppressed) of a strong proton bunch colliding with a weak antiproton bunch as the working point (b) is varied.

The net result of all this is that the total tune spread must be kept as low as possible in order to keep the working point away from these very high order resonances. In particular, for six bunches per beam, which is the design objective, the tune spread is double because the beams collide at 12 points, 10 of which cannot be used for experiments. The tune spread spans points 1 to 4 in the diagram and the lifetime is correspondingly bad.

In order to make the lifetime acceptable with 6 bunches per beam, we have to separate the beams at all the unwanted crossing points except one by using two pairs of electrostatic separators to produce a global orbit distortion all around the ring except in the region between the two experiments (Fig. 5). The distortion is in opposite directions for protons and antiprotons so the beams miss one another in nine of the unwanted crossings. The improvement in lifetime is spectacular (Fig. 6). Therefore the Collider would not have worked under the conditions of the design report. For once, the pessimists were right.

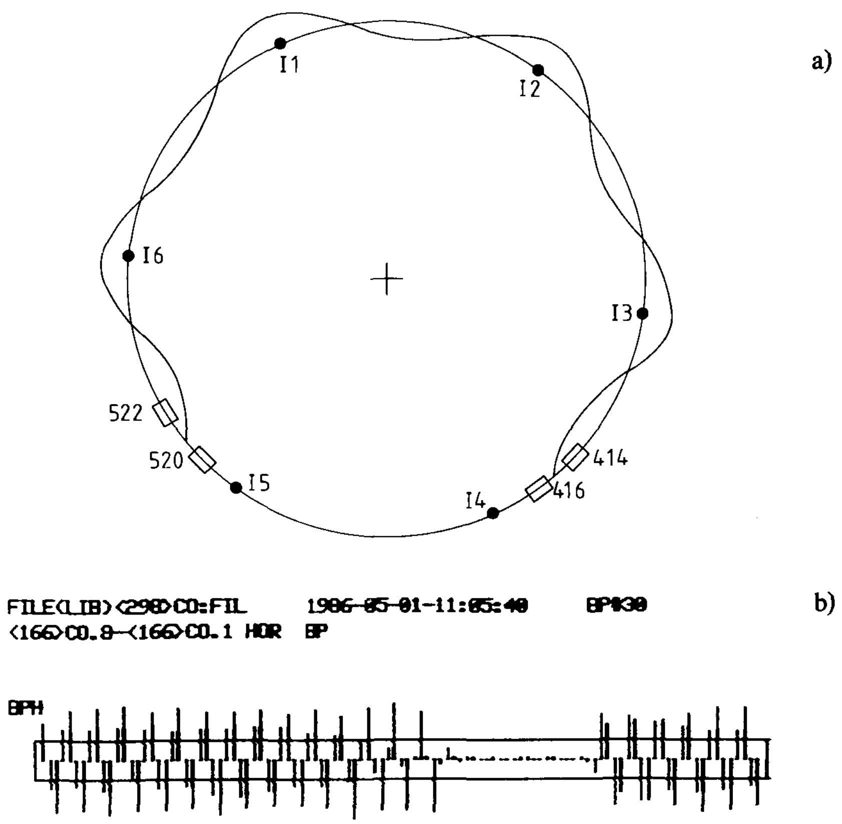

Fig. 5 Beam separation using electrostatic deflectors. (a) Schematic of SPS separation scheme with electrostatic deflectors. (b) Measured closed orbit with separation.

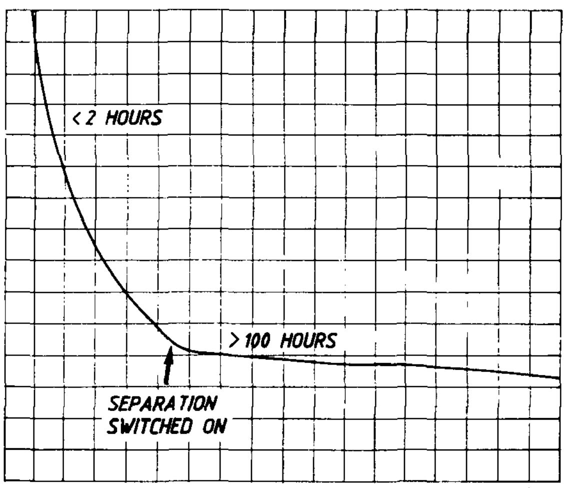

Fig. 6 Chart recorder output of the intensity decay of a single antiproton bunch in the presence of 6 dense proton bunches (zero suppressed) as separation is switched on.

7. IMPACT ON THE WORLD SCENE

It would be something of an understatement to say that the $p\bar{p}$ Collider has not gone unnoticed. In the first place, its initial success was partially responsible for the cancellation of the ISABELLE project at Brookhaven in 1983, though many question the wisdom of this decision.

The Collider has shown that clean physics can be achieved with colliding hadron beams—initially doubted by part of the physics community. The immediate reaction on the other side of the Atlantic was the Superconducting Super Collider (SSC) proposal, a giant hadron collider of 20 + 20 TeV with a circumference of around 80 km, three times that of LEP.

At CERN, a study group has been working on a more modest proposal for a hadron collider in the LEP tunnel, the Large Hadron Collider (LHC). Since the tunnel already exists, it is imperative that the bending magnets work at the highest possible field. Both the SSC and LHC would be pp colliders in the first instance, since the expense of producing an antiproton source and associated equipment for the very high luminosity operation which is imperative at these energies, offsets the cost of a second ring. Nevertheless, $p\bar{p}$ collisions would certainly be another option in the LHC.

8. CONCLUDING REMARKS: THE FUTURE OF THE $p\bar{p}$ COLLIDER

We no longer have the monopoly of high-energy colliding beams at CERN. The FERMILAB Tevatron has recently come on line and although its present performance is still below that of the SPS, its potential for improvement is great. It operates at 1.8 TeV in the centre of mass and although its luminosity will be comparable with that of the present SPS Collider, the cross-section for rare processes is considerably higher.

In order to remain competitive, the luminosity of the SPS Collider must be increased by about an order of magnitude. In order to do this a second ring, the Antiproton Collector (ACOL) [8] has been built and is being commissioned at the present time. Superficially, the ACOL/AA complex looks quite similar to the shuttle/cooling ring of the proposal made in 1976. Its function is to increase substantially the number of particles that can be captured from the production target, both in angular and energy spread.

Coupled with ACOL, SPS improvements have been made to allow us to digest the increased antiproton flux. These include the electrostatic separators already mentioned, a new RF system, and improved chromaticity correction.

These improvements, together with major upgrades of the main detectors should keep the SPS Collider competitive into the next decade.

REFERENCES

[1] S. van der Meer, Stochastic damping of betatron oscillations, Internal Report CERN/ISR/PO/72-31 (1972).

[2] G.C. Carron and L. Thorndahl, Stochastic cooling of momentum spread with filter techniques, Internal Report CERN/ISR-RF/78-12 and ISR-RF/Note LT/PS (1977).

[3] W. Schnell, About the feasibility of stochastic damping in the ISR, Internal Report CERN/ISR-RF/72-46 (1972).

[4] P. Bramham, G. Carron, H.G. Hereward, K. Hübner, W. Schnell and L. Thorndahl, Stochastic cooling of a stored proton beam, Nucl. Instrum. Methods **125**, 201 (1975).

[5] C. Rubbia, P. McIntyre and D. Cline, Producing massive neutral intermediate vector bosons with existing accelerators, Proc. Int. Neutrino Conf. Aachen, 1976 (Vieweg Verlag, Braunschweig, 1977), p. 683.

[6] Design study team, Design study of a proton-antiproton colliding beam facility, Internal Report CERN/PS/AA 78-3 (1978).

[7] S. van der Meer, CERN 62-16 (1962).

[8] E.J.N. Wilson (ed.), Design study of an Antiproton Collector for the Antiproton Accumulator (ACOL), CERN 83-10 (1983).

1966	Budker and Skrinskij propose $p\bar{p}$ colliding beams based on electron cooling.
1968	van der Meer. Idea of stochastic cooling.
1972	Schottky signals observed and interpreted at the ISR. van der Meer publishes paper on stochastic betatron cooling of transverse emittance.
1972	Schnell feasibility study of stochastic cooling experiment at the ISR.
1975	First experimental demonstration of stochastic cooling. Rubbia et al. propose $p\bar{p}$ colliding-beam experiment at the SPS.
1976	First design report based entirely on electron cooling. ICE construction started.
1977	Thorndahl invents filter method of fast momentum cooling. Thorndahl cooling tested on ICE. Refinement of theoretical understanding. Hereward and Sacherer. SPS storage experiments started.
1978	Second design report based entirely on stochastic cooling. July: Authorization of $p\bar{p}$ project.
1980	Start of 11 months shutdown for SPS modifications. June: Protons circulating in AA.
1981	July: First proton–antiproton collisions in SPS. November: First technical run (0.2 nb^{-1}).
1982	First real physics run (28 nb^{-1}). W found.
1983	January: W announcement. April–July: Collider run. Z^0 found.

Plate 1 Project History

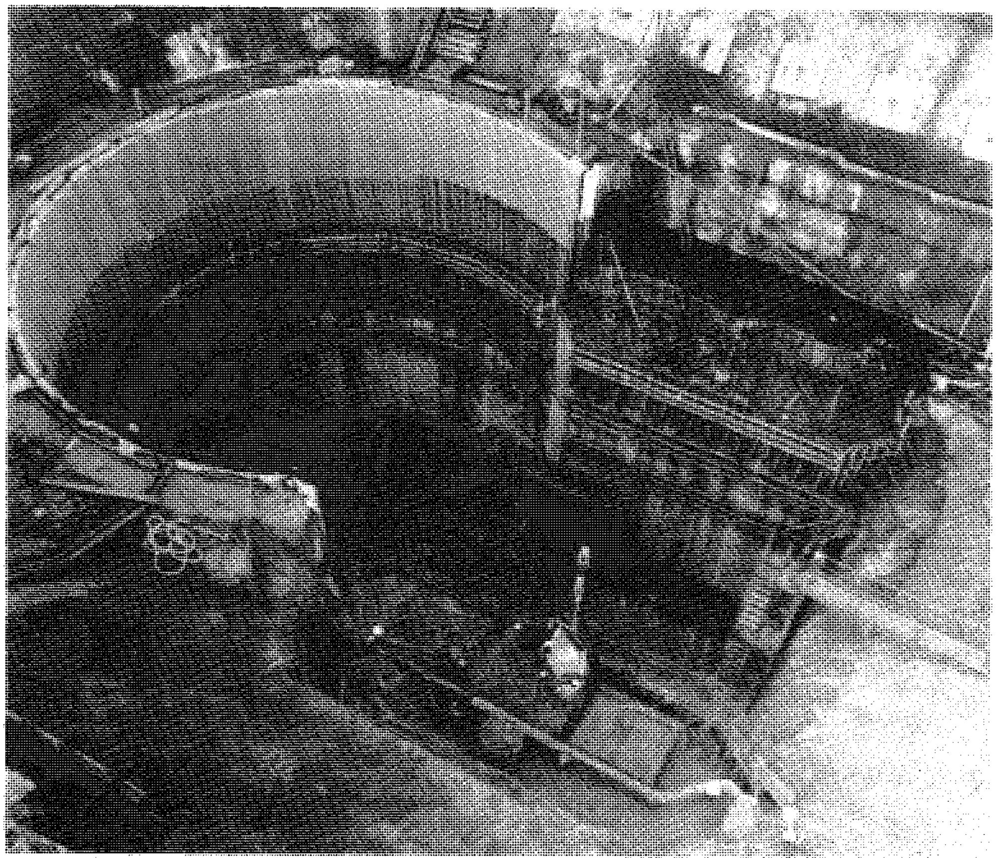

Plate 2 Part of the open excavation for the UA1 detector hall. The photograph shows the 'garage' area where the detector was assembled. The SPS tunnel covered with earth crosses the bottom right-hand corner. During the 1980 shutdown the tunnel wall was opened and the second half of the detector hall was excavated.

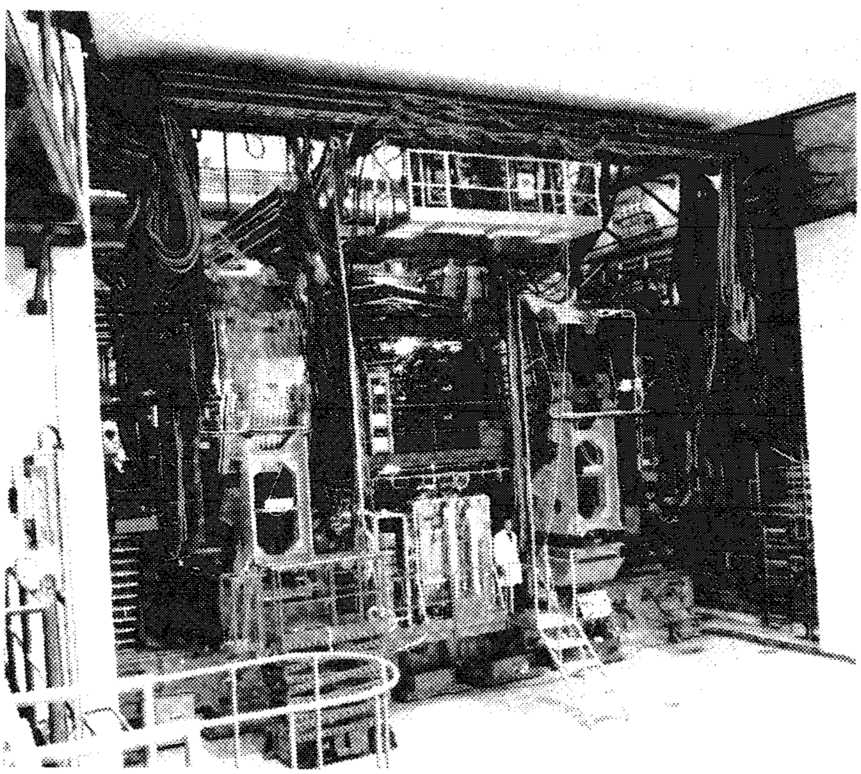

Plate 3 The UA2 detector in position and ready for data taking.

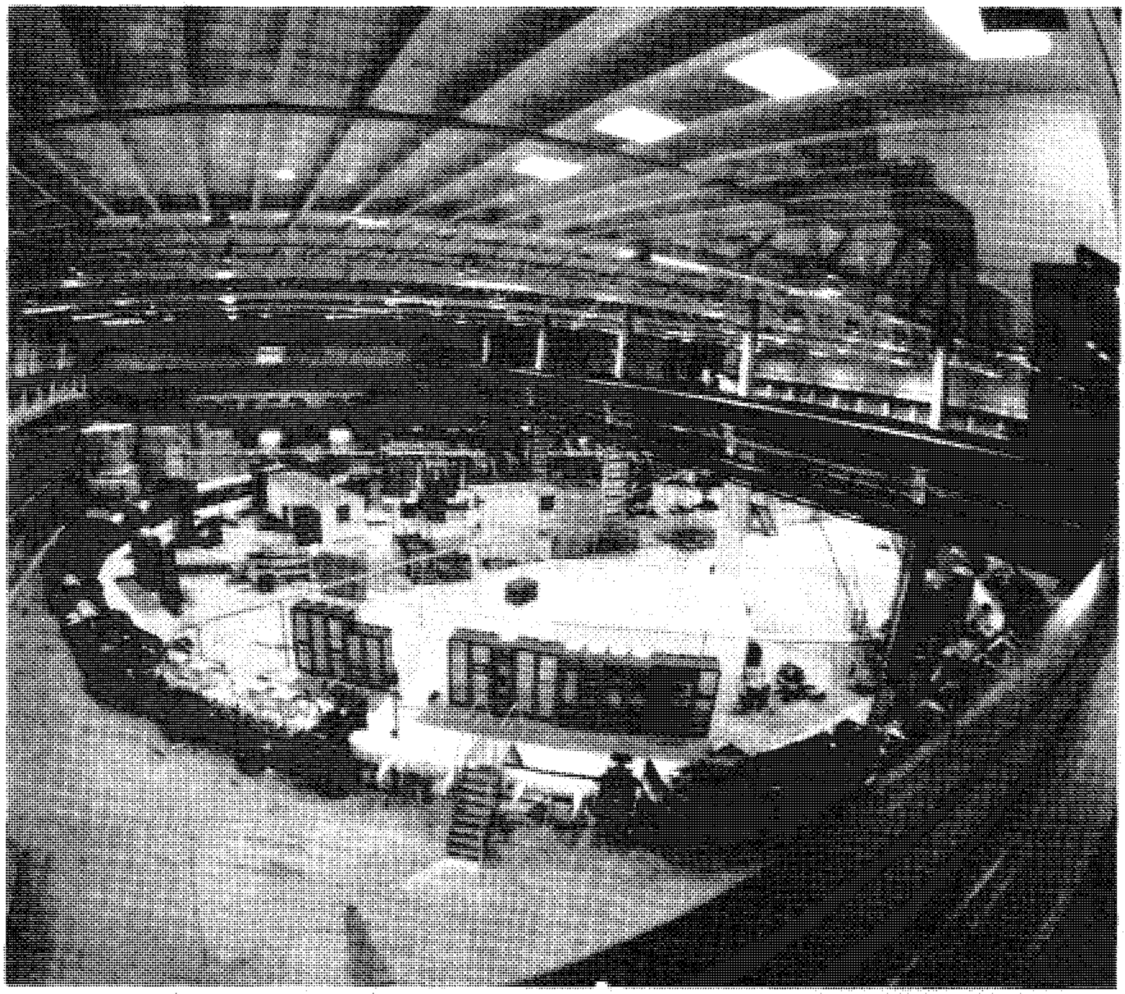

Plate 4 A view of the AA ring. The coaxial lines linking stochastic cooling pick-ups and kickers can clearly be seen.

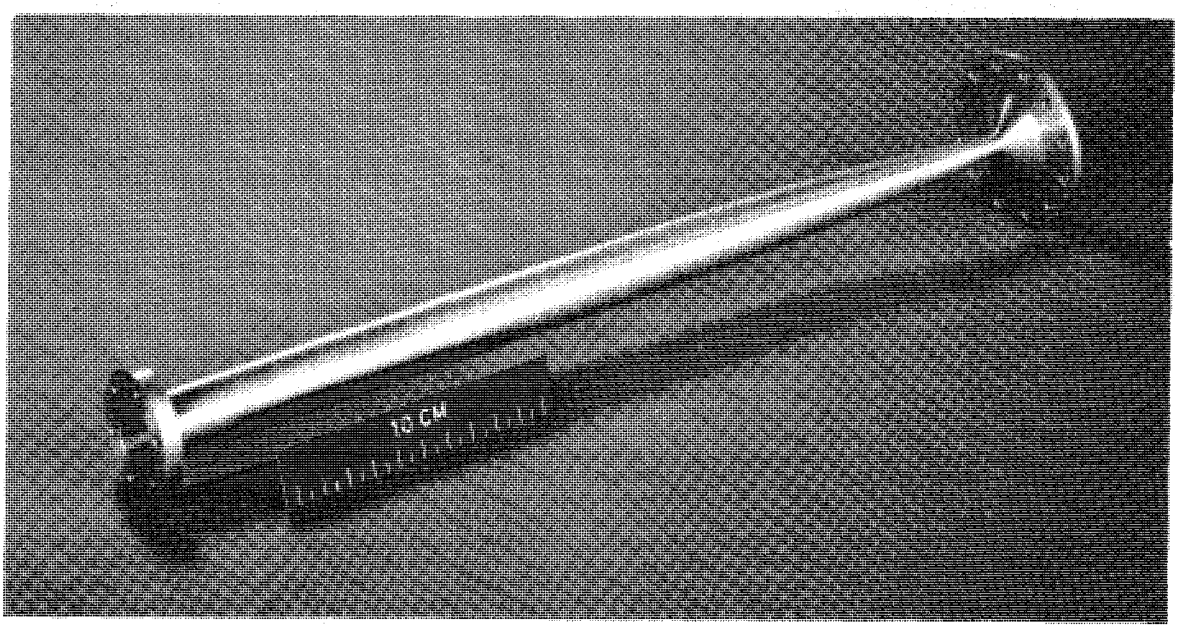

Plate 5 The inner conductor of the coaxial horn for the focusing of antiprotons from the target.

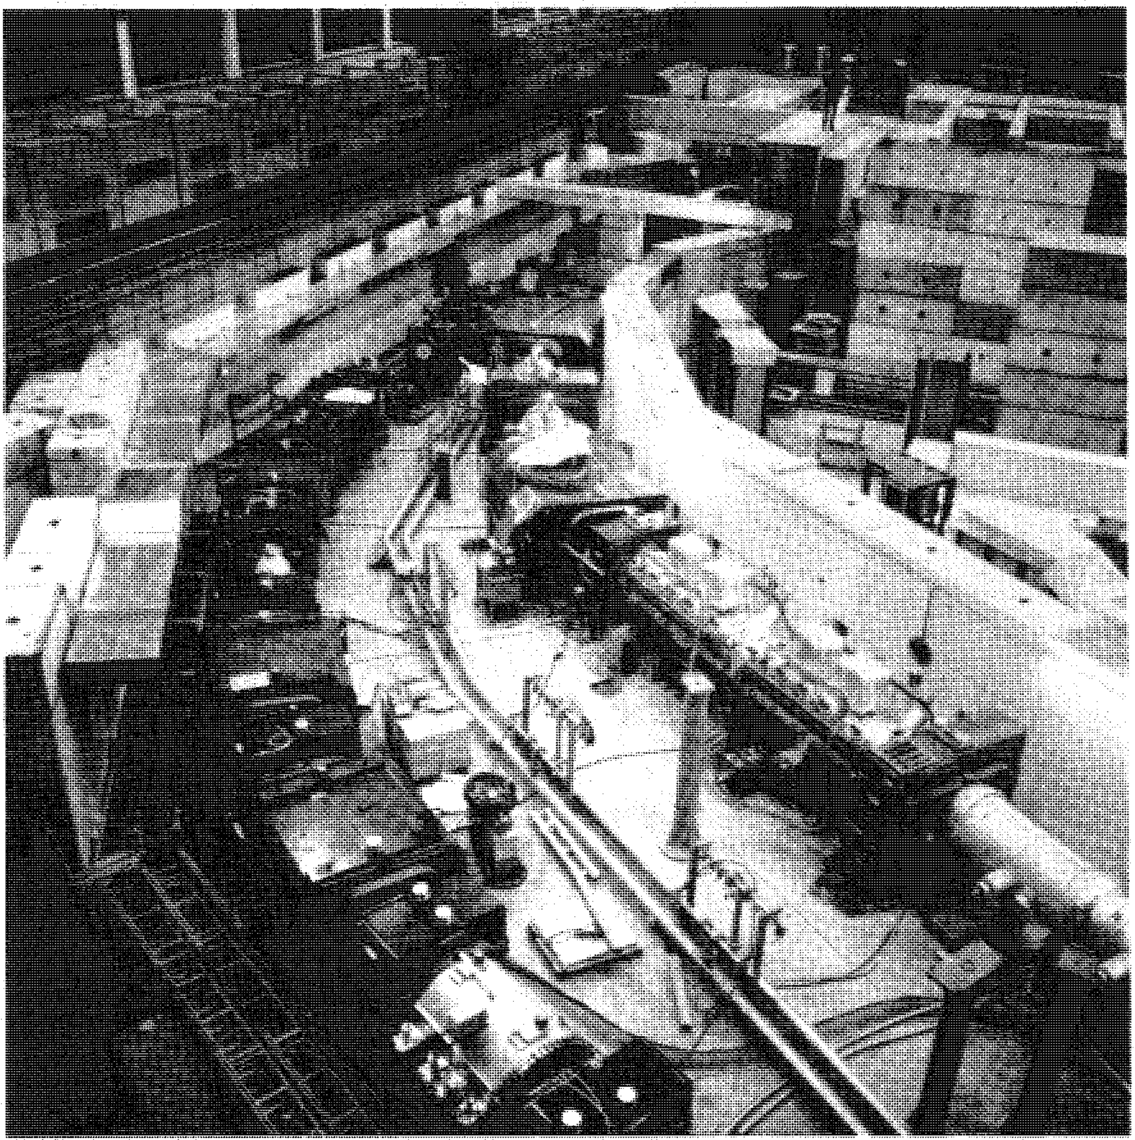

Plate 6 Part of the ACOL ring on the outside together with the AA on the inside. The purpose of ACOL is to increase the $\bar{p}$ collection rate by an order of magnitude, acting as a precooling ring for the AA.