Particle Accelerators, 1986, Vol. 18, pp. 167–182
0031-2460/86/1803–0167/$20.00/0

THE ISR AND ACCELERATOR PHYSICS†

KJELL JOHNSEN
CERN, 1211 Geneva 23, Switzerland

(Received May 3, 1985)

1. INTRODUCTION AND EARLY HISTORY

It was on 27 January 1971 that the very first hadron collisions took place in the Intersecting Storage Rings (ISR) at CERN, and it was on 23 December 1983 that the very last colliding-beam run was terminated, bringing to an end a very fruitful period of 12 years and 11 months of colliding-beam physics in this facility. The machine was working as well as ever, the physics program was as exciting as ever, but budgetary constraints imposed the decision to close this unique facility at the height of its performance.

The beginning of the history of colliding-beam devices can be traced to 1956, when a group at the Midwestern Universities Research Association (MURA) in the U.S. put forward the idea of particle stacking in circular accelerators. Of course, people who worked with particle accelerators had already speculated about the high center-of-mass energies attainable with colliding beams, but such ideas were unrealistic with the particle densities available in normal accelerator beams. The invention of particle stacking changed this situation significantly. It opened up the possibility of making two intense proton beams collide with a sufficiently high interaction rate to enable experimentation in an energy range otherwise unattainable by known techniques, except perhaps at enormous cost.

A group at CERN started investigating this possibility in 1957, first by studying a special two-way fixed-field alternating gradient (FFAG) accelerator and then in 1960 by turning to the idea of two intersecting storage rings that could be fed from the CERN 28-GeV Proton Synchrotron (PS). This change in concept for these initial studies was stimulated by the promising performance of the CERN PS at the very start of its operation in 1959.

The extensive study of the second possibility included the construction of an electron ring model (CESAR), which started operating in 1964 and which provided invaluable experience, in particular, in beam stacking and in the ultrahigh-vacuum technology so essential for colliding-beam devices. However, even earlier, in 1961, the Accelerator Research Division at CERN had gained sufficient confidence to present its first proposal for a 2×25-GeV storage ring

† An earlier version of this paper was presented at the final meeting of the Intersecting Storage Rings Committee, 27 January 1984, following the closing of the ISR for colliding-beam physics in December 1983.

system. This system was intended essentially for protons, but other particles were mentioned in the proposal. This led to a series of important actions. First, in 1962, France offered a site next to the original CERN site. The European Committee for Future Accelerators (ECFA) was then formed and in 1963 issued a strong recommendation in favor of a pair of 25-GeV proton storage rings, which it named the Intersecting Storage Rings (ISR). The recommendations also included a 300-GeV fixed-target accelerator.

In 1964 the Accelerator Research Division prepared a detailed design report for the ISR, which formed the basis for a formal proposal. In his last year as Director General, Professor Weisskopf, who had been one of the ISR's strongest proponents, saw the CERN Council adopt the ISR as a project by making the decision in principle, in June 1965, to construct this facility on a supplementary budget, then in December of the same year, by accepting the financial plan of the project and by voting construction funds from 1 January 1966.

The plans quoted a construction cost of 332 million Swiss francs (1965 value) and projected first operation of the facility by mid-1971. Both promises were fulfilled with some margin.

2. SOME FEATURES OF THE ISR CONSTRUCTION

Seen in retrospect, the ISR construction proceeded fairly uneventfully, although those involved experienced considerable anxiety at times. No detailed account of the construction will be given, but some features are worth recalling.

It was recognized that the ISR might encounter many unknown phenomena that might limit its future performance. For instance, unlike the case of electron rings, there would be no damping to keep beam sizes from growing because of small imperfections in the guiding field or perturbations arising in the beam. In short, it was considered a daring and bold project. Therefore, tight tolerances and flexibility for all components became important guidelines. This paid off handsomely, as tolerances and flexibility have been since stretched to their limit, and often far beyond, almost everywhere. A few examples can be mentioned.

The magnet system was designed with a set of unusually flexible pole-face windings. These windings became essential for the performance, in ways unforeseen before the start of operation. Many families of trimming magnets were incorporated, such as horizontal-field magnets, quadrupoles, skew quadrupoles, sextupoles, etc. They were all needed, and in some cases more had to be added.

Power supplies were designed to tight tolerances; non-the-less, some of the auxiliary supplies had to have their tolerances tightened by up to two orders of magnitude to satisfy the standards set by later operational experience.

During construction the system was changed by the adoption of the so-called suppressed bucket scheme, which made the stacking more efficient, at the same time simplifying the design of the injection system.

A vacuum system with an average pressure of 10^{-9} Torr was estimated to give adequate beam lifetime, and at the same time, such a vacuum was considered to be near the best achievable in such a large system. This was, therefore, taken as the design figure, and it was achieved with a good margin in hand. However, unexpected effects made even that inadequate, as I shall discuss below.

During construction a Beam/Equipment Interaction Committee was set up to monitor, at the design and prototype stage, all components that could be "seen" by the beam, the aim being to reduce as much as possible the coupling impedance between the beam and its surroundings. This was a difficult task, but an extremely important one. It is now clear how this foresight avoided much potential trouble.

Unusual care was taken to equip the machine with the best possible diagnostic equipment. Nevertheless, early operational experience led to the development of even more sophisticated techniques for beam observation.

These examples should illustrate that, although much care was taken during the construction of the machine, many of the components had subsequently to be developed and upgraded to the extent that the whole facility was a considerably more sophisticated device at the end of its life than when it was started up 13 years ago.

3. ACCELERATOR PHYSICS AT THE ISR

Because of the complexity of the ISR and the complexity of the phenomena observed on the machine (some of which were unexpected) beam studies and machine development were considered essential parts of the activity around the ISR during its entire lifetime. Figure 1 gives a typical example of how much time was devoted to these studies. It is not always easy to gain sympathy among users for such use of a large accelerator facility, but at the ISR this was never a serious problem, probably because some of the very early beam studies resulted in considerable performance improvements. The fact that over its lifetime about 15% of the "on-time" of the ISR was used for beam and component studies has

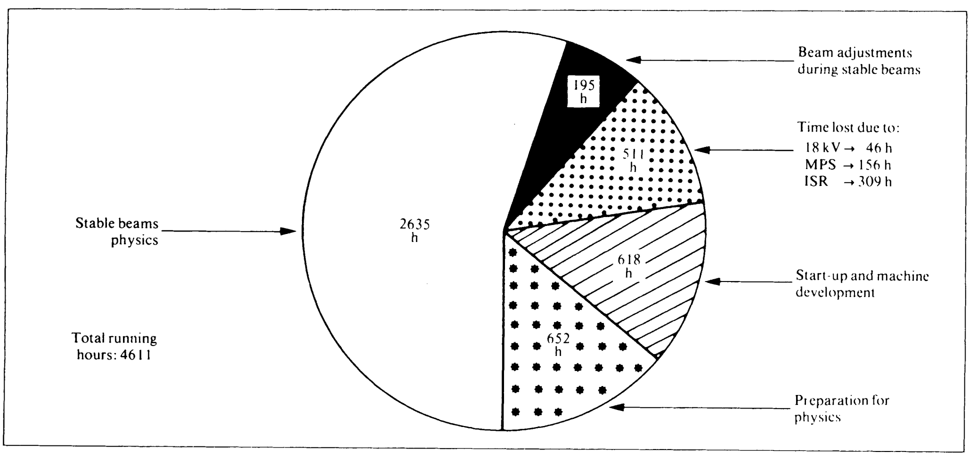

FIGURE 1 Distribution of 1981 ISR total operating time.

paid off in a rather spectacular improvement in performance over the years, as will be illustrated in what follows. Of course, these results were obtained not only because of the time devoted to such studies, but also because of the competence with which the studies were performed.

3.1. *"Mission-Oriented" Accelerator Physics*

Most of the machine-development studies carried out at the ISR had the specific aim of solving problems that arose in the effort to improve the performance of the facility. A few important examples of this kind of activity will be given in the following paragraphs, approximately in the sequence in which the problems first appeared.

3.1.1. The resistive-wall instability (*the "brick wall"*) During the very early operation of the ISR, we were only able to stack beam currents up to 2–3 A without difficulties. When we tried to stack to higher beam currents, instabilities arose that resulted in beam losses. Often a partial loss was enough to stabilize the beam, and stacking could continue. The left-hand side of Fig. 2 is a typical example of beam behavior under such conditions. The instability was associated with large coherent transverse oscillations, consistent with the theory of the resistive-wall instability. However, a sextupole field component had been built into the magnet profile to avoid this instability up to much higher beam currents than 3 A. At first, therefore, we were taken a little by surprise. In fact, what happened was that space-charge tune-shifts led to the violation of the stability

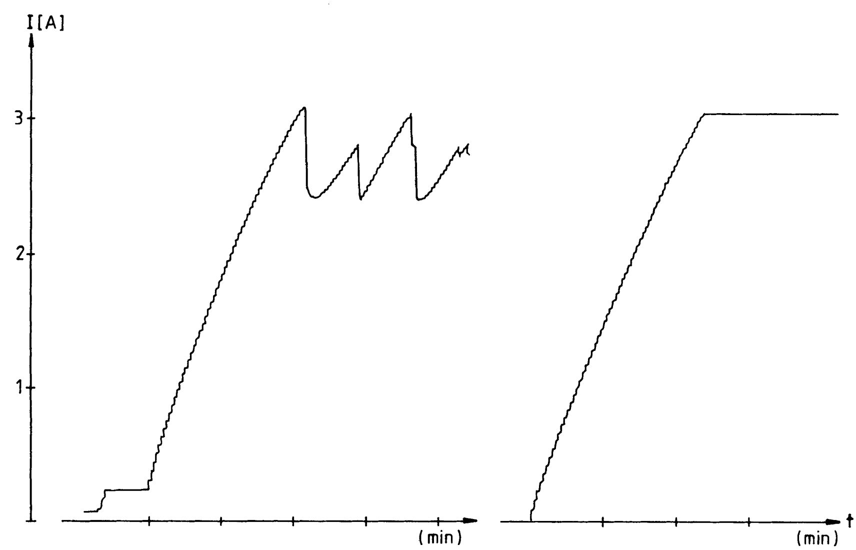

FIGURE 2 Beam current vs time plots, showing stacking operation hitting the "brick wall," at left, and an example of good stacking, at right.

requirement $dQ/dp > 0$ (Q being the betatron tune) in *parts* of the beam although $\Delta Q/\Delta p > 0$ was satisfied for the global beam. This also explained nicely why only small parts of the beam were often affected. The main cure was a very careful tailoring of the dependence of both Q's on momentum, the so-called working lines. Later such working lines required dynamic compensation of the space-charge effect as the beam current increased (Fig. 3), and as experience was gained, such compensation was carried out on-line by the control computer. Although this was simple in theory, it required extreme accuracy and flexibility of the components involved, as mentioned in Section 2.

This remedy was not enough to reach very high beam currents, so a transverse feedback system was developed and incorporated into the machine for stabilization of the lowest oscillation modes. This kind of development went on over many years and led to a gradual increase of the stability limit to around 60 A.

3.1.2. Pressure bumps After we started mastering the resistive-wall instability and beam currents of 4–5 A were obtained, another type of beam loss appeared. Again it resulted in a partial loss, as seen in Fig. 4. The phenomenon occurred on a much slower time scale than the "brick wall", and there were no associated beam oscillations. However, it was soon observed that dramatic local vacuum deteriorations preceded the beam loss. It was further observed that the pressure

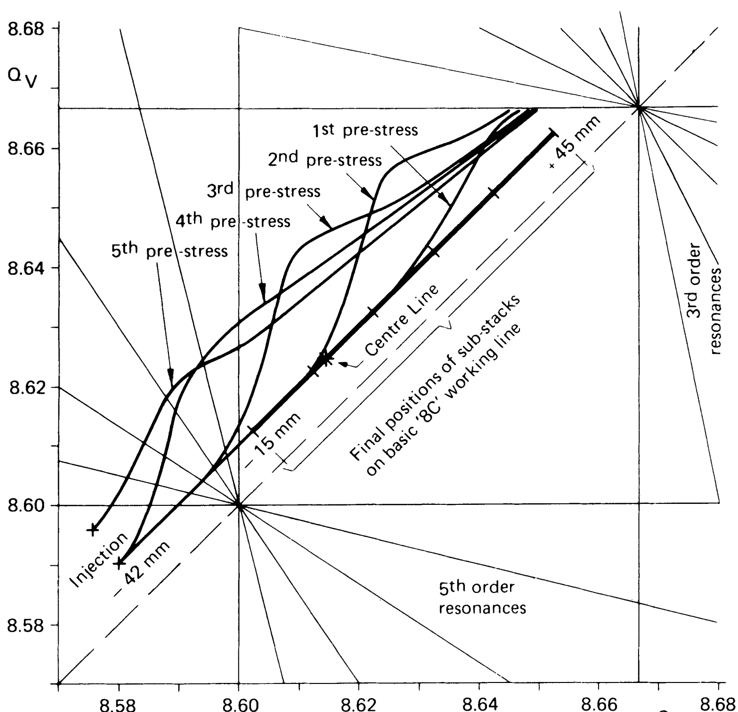

FIGURE 3 The family of prestressed working lines used at 22 GeV/c to stack 15 A in five steps of 3 A across the chamber from +45 to −15 mm (α_p average).

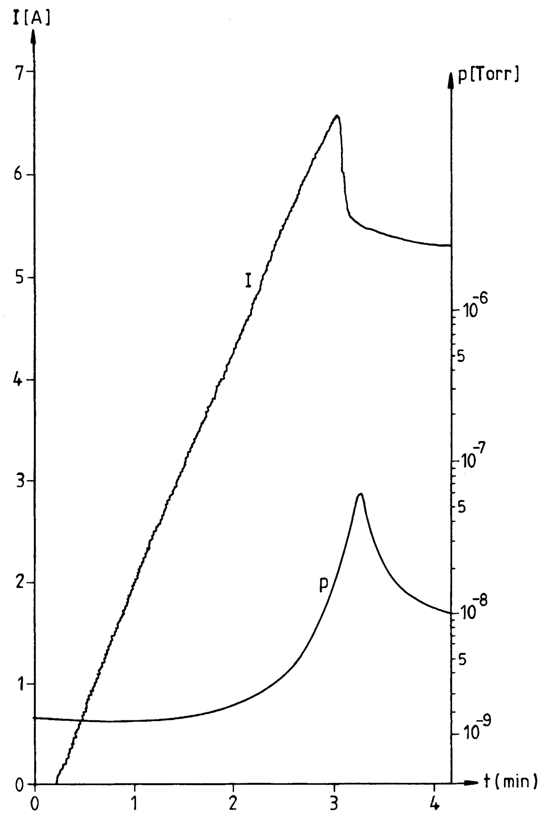

FIGURE 4 Example of a pressure increase caused by the stacked beam.

effect was not beam-loss dependent, but that it depended sharply on total beam current. Once these facts had been observed, the explanation became simple: The protons in the beam ionize part of the residual gas. These ions are driven into the walls of the vacuum chamber by the beam potential of 1 kV and penetrate considerably deeper into the wall material than a normal cleaning method does. This ion bombardment leads to desorption from the wall, resulting in more gas to be ionized by the protons. A runaway situation occurs above a certain critical beam current $I_c \propto S/\xi$, where S is the pumping capacity available and ξ is the desorption coefficient.

In the ISR, the cures were to attack these two parameters, S and ξ, by the addition of hundreds of titanium sublimation pumps, by an increase in the bake-out temperature of the vacuum components from 200°C to 300°C, and by glow-discharge cleaning of all the critical vacuum components. This was a gradual program. During each shutdown, the weakest point was attacked, which led to the optimum rate of improvement. The result was a steady advance over many years, leading to critical currents above 60 A and average pressures below 10^{-11} Torr. Figure 5 shows how the average vacuum improved over the years as a result of this program.

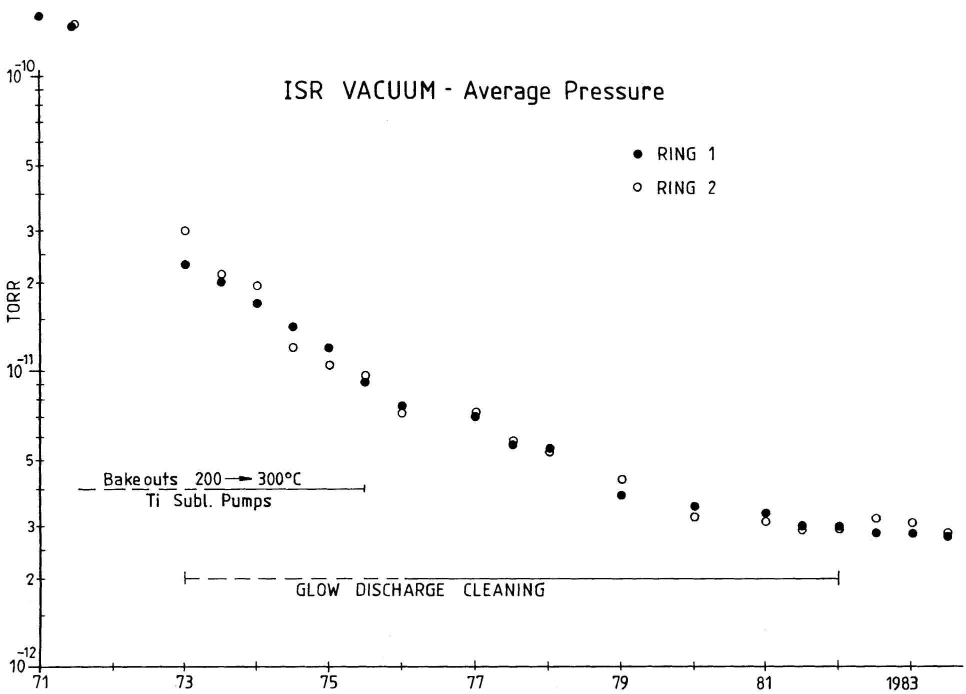

FIGURE 5 The average pressure of the ISR vacuum for the years 1971–83.

3.1.3. Development of vacuum chambers for experimental areas Experimental areas had their special requirements for the vacuum system. For background reasons, it was necessary to produce a pressure well below the average in the rest of the machine, and in the last years of operation, 10^{-13} Torr was achieved. These chambers had to have special shapes, and at the same time, the wall thickness had to be kept to a minimum in order not to interrupt the secondary particles. This sometimes conflicted with sound mechanical engineering, and a few collapses of vacuum chamber walls occurred when safety factors were reduced too far. Development over the years nevertheless resulted in very satisfactory designs, giving reliable operation in spite of working very close to the limits.

3.1.4. Beam clearing However good the vacuum is, there will always be some residual gas that will be subject to ionization due to the bombardment of the protons in the beam. The electrons created by this ionization tend to gather in the potential well of the beam, thus leading to the neutralization of the beam. This was considered to be unacceptable, because it would shift the betatron tune of the machine, and for this reason, clearing electrodes working at several kilovolts were provided.

In the machine design, it was estimated that, if a clearing of about 90–95% could be achieved, no problems would arise. Early operational experience demonstrated this to be wrong. The clearing had to be improved to very close to 100%, and it was particularly important to clear local pockets of electrons that had a tendency to be trapped in unforeseen places. In fact, the ultimate problem turned out not to be the shift in betatron tune but rather a sophisticated coupling between the proton beam and the electron cloud in such pockets, leading to beam oscillations ("e–p lines" at the frequencies of electron oscillations in the beam's potential well) and consequential beam losses. Considerable detective work was needed in the early years of operation to find all the areas of inadequate clearing. This effort removed all these problems.

It is worth noting that this is a problem specific to coasting beams. Bunched beams clear themselves adequately.

3.1.5. Background for physics experiments Apart from perhaps the first few weeks, the ISR beam lifetime has always been adequate for efficient operation. However, it was soon discovered that background conditions were bothersome or unacceptable at loss rates far below those acceptable from the lifetime point of view. This led to the development of a series of remedies to prevent beam losses. The main attack was to keep the stored beam as resonance-free as possible. It was, for instance, necessary to completely avoid the fifth-order resonance (Fig. 3), and it was preferable to stay away from all orders up to the eighth order as well. This became possible after the transverse feedback system was developed, opening up the clean area with a betatron tune between 8.9 and 9.

In addition, very good beam collimation was necessary, and it became customary to scrape the halo off the beam at intervals during a physics run. Figure 6 illustrates a typical loss rate of the order of one part per million per minute and shows how scraping operations (the large peaks in the figure) help maintain this to the end of a 56-hour run.

3.1.6. Longitudinal instabilities Since longitudinal instabilities were among the better-known phenomena, exceptional care was taken during construction to avoid such difficulties (see Section 2). This effort paid off handsomely, and no real difficulties arose at low frequencies. However, in the microwave range, instabilities were observed when the injected beam was being debunched on the stacking orbit. This led to a reduction in stacking efficiency. A special cavity working on the third harmonic of the main radio frequency was constructed to enhance the nonlinearities of the stacking "bucket." This provided Landau damping, which suppressed the instability.

3.1.7. Beam-beam problems A beam-beam limit in the normal meaning of the word was never observed under normal operation of the ISR, in spite of extensive studies. The beam-beam effect was probably too weak for that. (A little more on this later.) However, from 1978 a beam-beam instability did occur occasionally at beam currents above about 30 A when the beams were steered head-on to each other. The most critical separation seemed to be one-and-a-half to two beam widths, but the probability of losing the beam by such a manipulation was only

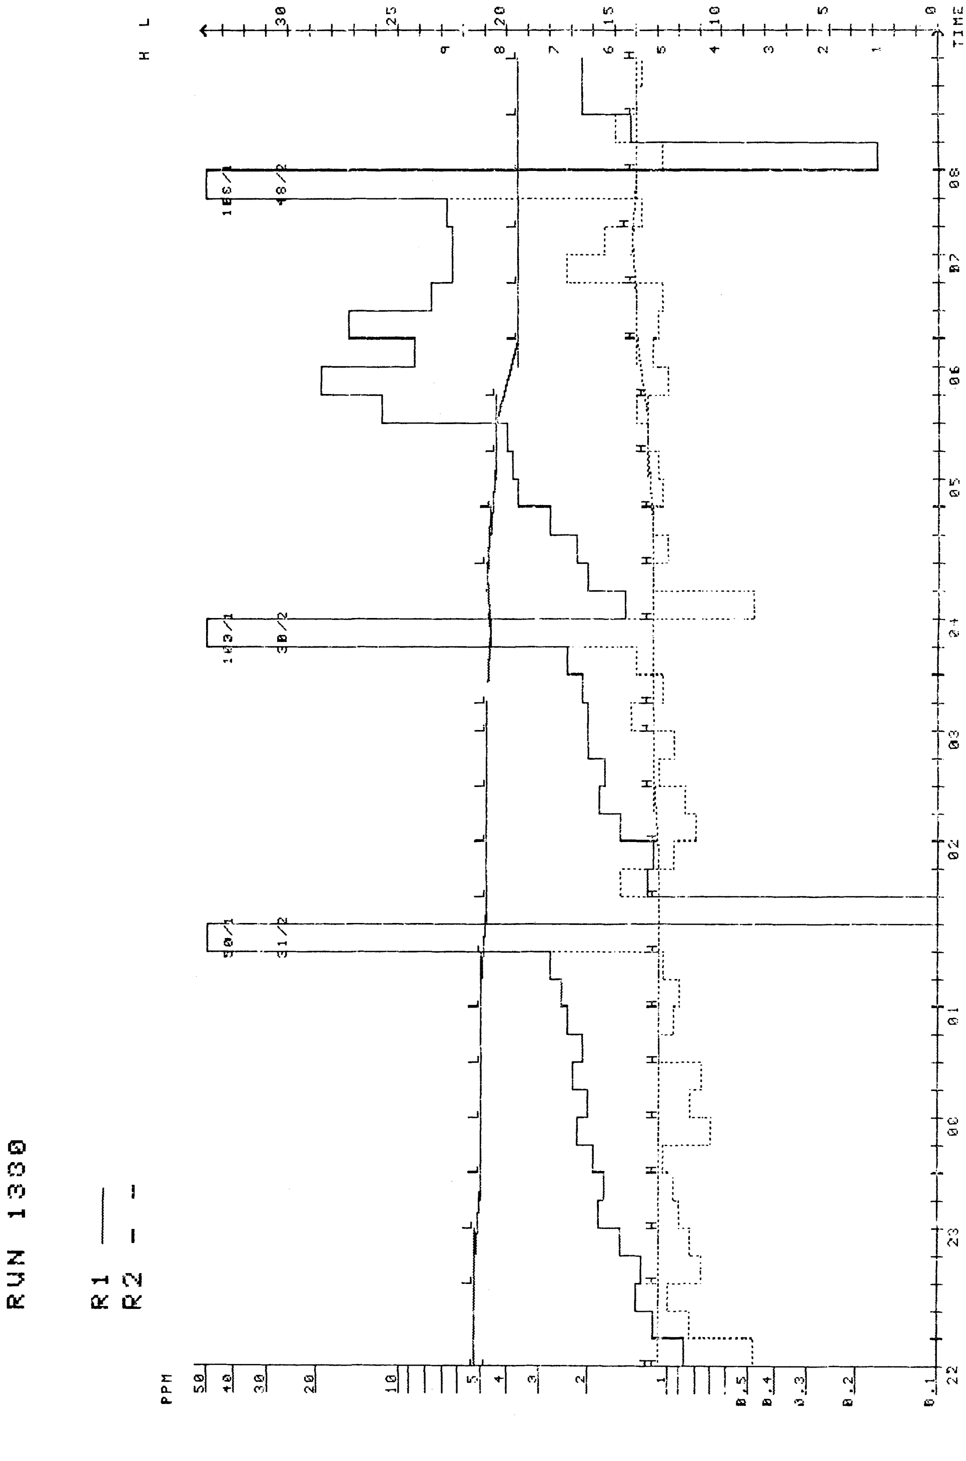

FIGURE 6 Example of an ISR run summary.

about 20%. This erratic behavior made systematic studies of the phenomenon rather difficult.

There have been quite a few theoretical speculations on the cause. A recent theory on coherent response of a coupled-beam system indicates that the critical frequency is outside the range of the transverse feedback system. This should have been tested during the last hour of ISR operation on 23 December 1983, but the beams had to be dumped prematurely. A quote can be given from the last report on the subject: "The ISR ghost disappeared with the last beam."

3.1.8. Phase-displacement acceleration The magnet system of the ISR had a sufficient safety factor that it could accept particle energies up to 31.4 GeV. However, a high-current stack has an energy spread of up to 3%, which is much more than the rather low-power system of the ISR could rebunch and accelerate from the injection and stacking energy of 26 GeV. During the early operation of the ISR, small currents were rebunched and accelerated to 31 GeV in the normal way, but the results were unsatisfactory for physics owing to the low luminosity that this gave.

An acceleration method called phase-displacement acceleration had been invented by the MURA group and had formed a part of their original paper on stacking. The method consists of repeatedly moving "empty buckets" through the beam from above to below. This shift downwards in energy of empty phase space leads to a corresponding shift upwards in energy of the phase space occupied by the beam. No large voltages (buckets) are needed. Small buckets merely require more passages. In 1973 it was decided to try this method at the ISR. It required very low noise operation of the RF system and very fine variation of bucket size and magnet power supplies. The working line also required very fine adjustments during the acceleration to avoid instabilities. This has been developed over the years into a very sophisticated procedure where all the operations are controlled on-line by the computer, with only occasional operator intervention. The result has been that 2×31.4-GeV beams have been available for physics with essentially no loss of luminosity or other beam qualities.

3.1.9. Low-β insertions In the original design of the ISR some care was taken to make the vertical β in the intersection somewhat smaller than in the rest of the machine. However, a low-β insertion, in the proper sense of the word, was designed and installed only in 1974, after construction was finished. This low-β section consisted largely of borrowed quadrupoles from the PS, DESY, and the Rutherford Laboratory. It increased the luminosity by a factor of 2.3 in this particular intersection, a very welcome improvement.

At about the same time, however, work began on superconducting quadrupoles for a more powerful low-β insertion. Such an insertion was installed in 1981 and gave a 6.5-fold improvement in luminosity, which brought the operational luminosity to above 10^{32} cm^{-2}s^{-1}, a record that has not been beaten by any other colliding-beam machine. Figure 7 illustrates how the effective beam height was

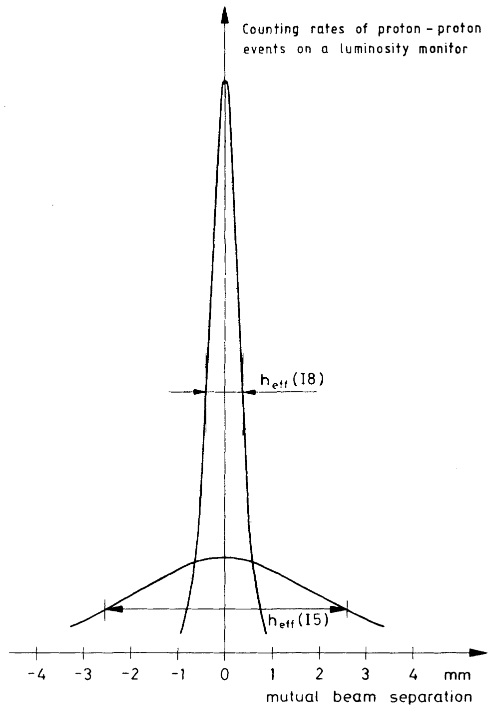

FIGURE 7 Counting rates of proton-proton events on luminosity monitors, as a function of the beam separation, showing the effective height of the beam (h_{eff}) at a normal interaction point (I5) and in the low-β insertion at I8.

reduced in this insertion by a large factor, compared with the normal intersections. (Such curves result from the normal luminosity-measuring method to be described in the next section.)

Moreover, this improved insertion yielded other important benefits as well. It gave us experience in constructing superconducting magnets of accelerator quality (precision, reliability, etc.), and it gave us operational experience under the very stringent conditions imposed by circulating stored beams.

3.1.10. Diagnostics As already mentioned in Section 2, many diagnostic methods were used in the ISR, some of which were developed far beyond their planned capability, largely because of the operational demands that arose after the start-up of the machine. Some of these developments were among the most exciting experiences at the ISR, and it is with regret that I describe only a few typical examples.

Over the years, luminosity measurements became more significant than envisaged during the planning and construction period, as these measurements entered so directly into the accuracy of some of the most important experimental results from the ISR. The method of measurement consisted of precision steering for beam separation at the interaction points, giving very accurate values for the effective height of the beams (see Fig. 7). This, together with highly accurate measurements of the circulating currents, gave correspondingly accurate values for the luminosity. Considerable development and careful checking was needed to reduce both systematic and random errors to desired values. Operationally, the procedure was used to calibrate the physics monitors at the beginning of a run, after which the monitors recorded subsequent variations in the beams during the run.

It was foreseen neither in the planning of the ISR nor during its construction to use Schottky noise as an element in beam diagnostics. It came, in fact, as a surprise to observe this noise on the ISR beam in 1972. However, as soon as it was observed, its potential became apparent, both as a tool for diagnostics and for the practical development of stochastic cooling (see later). The fact that a circulating beam consists of a finite number of protons gives rise to statistical fluctuations in the beam current and in the beam's center of gravity. These fluctuations are very small and had not been previously observed in accelerator beams. However, the electronics available in the early seventies had improved, and integration methods could be used with beams that had lifetimes as long as those at the ISR. It thus became possible to observe these fluctuations and to make use of them, both to monitor the distribution of particles in longitudinal momentum and to measure the extremes of the tune values in the stack without any interference with the coasting beam. It also became instrumental in detecting the growth of betatron amplitude at particular orbits in a stack, which helps to discern the presence and strength of various nonlinear resonances. This became an operational tool from about 1974.

However, continued development led to new and better applications. An example of this development is the use of the beam with its transfer function (the meaning of transfer function being that used in servo-mechanism theory) as an element in a feedback loop to study the stability of the system. The method consists of giving the beam a small and harmless perturbation. The result is observed by Schottky scans and fed through a fast Fourier transform to give a stability diagram of the Nyquist type. This is a very powerful way of taking preventive action against instabilities, and even on-line compensation methods have been used.

The Schottky scan method has possibly become the most powerful of all beam observation methods, not only for the ISR but also for the other collider projects in operation, under construction, or in the planning stage.

3.2. Accelerator Studies Initiated by General Interest

Quite a few accelerator studies were started in the spirit of answering general questions in accelerator physics and technology, rather than because of the pressure of performance requirements. Quite often these studies also resulted in practical applications in the ISR. Only a few of these studies will be listed below.

3.2.1. Cold bore In an accelerator with superconducting magnets, it is natural also to keep the vacuum chamber at liquid-helium temperature and thus to rely upon the pumping action of the cold chamber walls to achieve the required vacuum. This gives a high pumping capacity. However, for a storage ring, we have seen in Section 3.1.2 that the desorption coefficient is an equally important element in avoiding beam-induced pressure bumps. Since a cold chamber wall pumps by collecting a certain layer of frozen gas on the surface, there was some fear that the surface might develop a large desorption coefficient that could spoil the good vacuum performance. In order to investigate this experimentally, a section of cold vacuum chamber was installed and tested in the ISR. The tests demonstrated critical currents above the achievable ISR currents, and the conclusion could be drawn that a cold vacuum chamber is perfectly acceptable from a pressure-bump point of view.

3.2.2. Lifetime of bunched beams There was, at one time, some uncertainty about the lifetime of bunched beams compared with coasting beams, and careful studies were undertaken in the ISR and later in the PS and the SPS. The most conclusive results came from the SPS, but the ISR contributed to the conclusions that the main source of difference was the noise in the system and that as this was eliminated the lifetime of bunched beams approached that of coasting beams.

3.2.3. Beam-beam effects When two beams cross, one beam behaves as a lens acting on the particles in the other beam. The linear part of this lens only shifts the tune of the beam, and this can be easily corrected. However, this kind of lens has very strong nonlinear components, which are harmful to the beams. (In spite of the fact that the linear part is harmless, the linear tune shift is normally used as a measure of the strength of the perturbation of the nonlinear part.) This is the main limitation on luminosity observed on e^+e^- colliders. In the ISR this effect was too small to be observed under normal conditions. However, it constituted an uncertainty in the design of other hadron colliders, and therefore considerable studies were undertaken to try to find the limits that the beam-beam effect imposes.

A early experiment simulated the nonlinear beam-beam effect with a nonlinear lens. Another approach was to increase the β value in the crossing region and to decrease drastically the energy of the "probing" beam to enhance the effect. The effect on bunched beams was also investigated. These experiments showed that the loss-rate enhancement due to beam-beam effects starts appearing at tune shifts of the order of 0.02 in a single intersection for coasting beams and at 0.003 per intersection for bunched beams. For bunched beams the SPS $p\bar{p}$ Collider has later given more conclusive results.

3.2.4. Stochastic cooling Stochastic cooling was invented in 1968, but at that time it was considered unrealistic for practical applications. This changed completely in 1972, when technology had advanced sufficiently that Schottky noise was observed and used for beam observation. This development made it natural to try an experimental verification of stochastic cooling on an ISR beam.

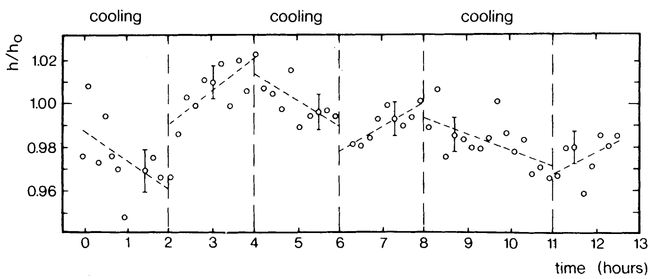

FIGURE 8 Observation of stochastic cooling in the ISR through measurements of the effective beam height (h/h_0) as a function of time, decreasing when cooling is applied and increasing when not applied. The cooling equipment, installed in only one ring, detects and corrects statistical fluctuations of average beam position. Luminosity is inversely proportional to the effective beam height.

Cooling equipment was built and installed in the ISR, and after some initial difficulties, the stochastic cooling effect was clearly demonstrated, as illustrated in Fig. 8.

As is well known, the most spectacular use of this cooling technique so far has been in the Antiproton Accumulator (AA) at CERN. However, other important applications were also made directly in the ISR. Special cooling systems were built for the circulating antiproton beam in the ISR and for proton beams up to 10 A (the highest intensity ever cooled). This led first to an initial significant increase in luminosity during $p\bar{p}$ physics runs. Second, it made it possible to keep the antiproton beams circulating for an incredibly long time (the record was 345 hours) without a decrease in luminosity. Third, the antiproton loss rate was unmeasurably small, which led, for instance, to a lower estimate of antiproton lifetime of more than 1000 hours (in the rest frame of the particle).

Another application of stochastic cooling in the ISR is for the hydrogen-jet experiment. A circulating antiproton beam in the range 3.5 to 5.72 GeV/c is cooled in the transverse and longitudinal planes. In the latter, the relative momentum spread is cooled down to $\pm 3.5 \times 10^{-4}$. This "fixed-target" experiment will be the very last experiment on the ISR, to take place during the spring of 1984.

3.3. *Resulting ISR Performance*

In the above summary of accelerator studies on the ISR it has only been possible to list some of the highlights. A very much wider spectrum of studies is behind the impressive improvements of the ISR performance from the start in 1971 to the final colliding-beam run in 1983. The performance can be illustrated by a list of figures and a few examples.

The highest centre-of-mass energy has been 62 GeV (equivalent to a fixed target accelerator of 2 TeV). This has been an operational energy used for a large fraction of the physics runs and with close to maximum luminosity.

The highest stacked current ever seen in a single beam was 57 A. Normal operational currents during high-luminosity runs have been 30 to 40 A. The highest luminosity at the start of a physics run was $1.4 \times 10^{32}\,\text{cm}^{-2}\,\text{s}^{-1}$, achieved in December 1982. Figure 9 illustrates the development of the maximum luminosity over the years of operation of the ISR. Loss rates during physics runs have typically been kept to one part per million per minute, which rendered very good background conditions. In fact, most of the background that the experiments struggled with came from the unwanted parts of the pp collisions.

Very long uninterrupted physics runs could be provided—of 50 to 60 hours if desired. This was important since all the manipulations from switch-on of the machine to good beam conditions for physics typically took 10 hours, and long stable operation was also high—86% of scheduled physics time during the last year of operation.

The ISR has also been used for antiprotons, α particles, and deuterons. In 1977 deuterons and protons were stacked in the same ring, and in a rather impressive display, the proton stack was decelerated through the deuteron stack, illustrating how particles with different revolution frequencies can be selectively treated in the same ring.

In the last three years of ISR operation, several p$\bar{\text{p}}$ runs over the energy range 15–31 GeV took place. A record antiproton current of 13.82 mA and a record luminosity of $4.5 \times 10^{28}\,\text{cm}^{-2}\,\text{s}^{-1}$ (in the superconducting low-β insertion) was

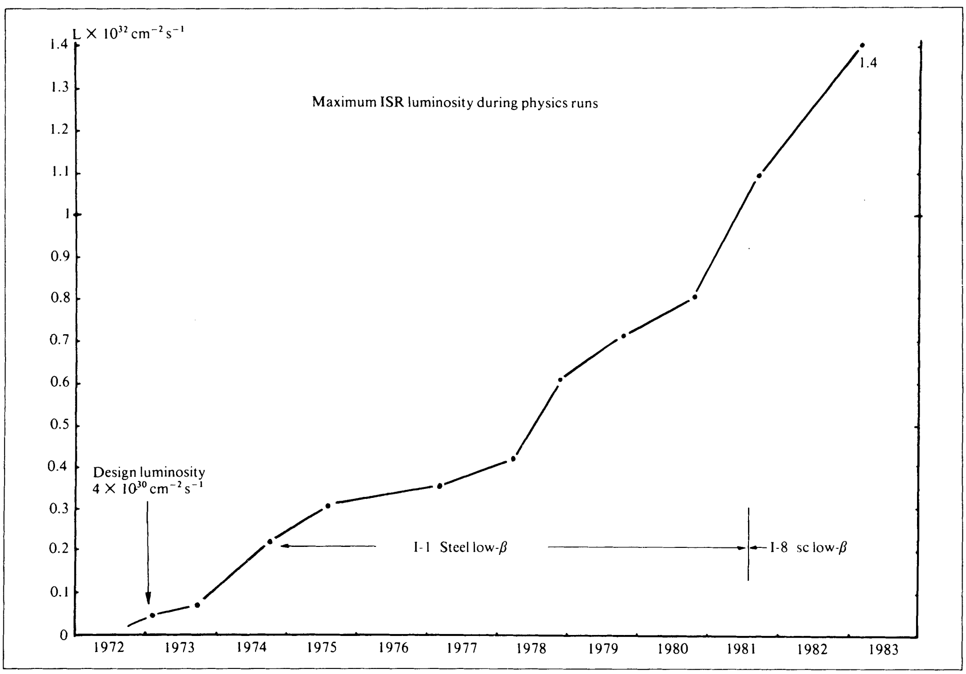

FIGURE 9 ISR luminosity during physics runs: September 1971 — First ISR experiment to be completed, R101; maximum luminosity = $1.3 \times 10^{29}\,\text{cm}^{-2}\,\text{s}^{-1}$. December 1982 — Highest luminosity achieved for physics (R807) = $1.4 \times 10^{32}\,\text{cm}^{-2}\,\text{s}^{-1}$.

achieved at the very end of 1983. The quality of the ISR is illustrated by the exceedingly low background, which made it possible to distinguish events in the highly asymmetric situation of antiproton beams of the order of milliamperes colliding with proton beams of up to 20 A.

4. CONCLUDING REMARKS

The ISR performance has improved over the years of operation far beyond the most optimistic hopes of its planners and builders, and the fine use of this facility made by the physics community is very much appreciated.

In addition, the ISR has been the finest instrument one can imagine for research in accelerator physics. The accelerator studies performed have led to technological inventions and developments in such areas as vacuum, diagnostics, stochastic cooling, controls, etc.

From the general performance of the ISR and the related accelerator development has emerged a general confidence in our ability to predict the performance of other hadron colliders, most of them still being planned or built. In fact, a considerable change in attitude has taken place. Before the time of the ISR, only fixed-target facilities were on people's minds for hadron physics. Nowadays, one talks almost exclusively about colliders. This has been a fantastic transformation for those of us who remember the reluctance of many physicists to accept the idea of colliders as a useful physics tool at the time the idea of the ISR was launched.

In this review of the development of the ISR, I have avoided quoting names. This may be unfair to those who have made outstanding contributions and I apologize to them. However, it would probably be more unfair to quote some names and leave out others. Hundreds and hundreds of devoted people have been involved in the ISR development from 1960 till today and in the exploitation of the facility. The achievements I have described (and many more) are the results of a remarkable cooperation among all these people.

The greatest strength of the accelerator community at CERN has been its ability to pull together, and the finest example of that cohesion of effort has been the ISR program, from its beginning to its end.