With construction of the Large Hadron Collider at CERN currently under way, physicists face some exciting challenges in building what will be the world's biggest particle accelerator

Particle accelerators: to the LHC and beyond

Steve Myers

WHEN Ernest Walton was born 100 years ago this month on 6 October 1903, experimental particle physics – the field in which the Irishman later made his name – was still a simple subject. Laboratories were small and the revolution of modern physics had only just begun. Even in the early 1930s, when Walton and his Cambridge colleague John Cockcroft built their historic “atom-splitting” machine, particle accelerators were humble devices, constructed by one or two individuals working in isolation.

But as particle physics has advanced, so researchers have sought machines of ever higher energy. Today, modern particle accelerators are huge undertakings, costing billions of euros to design and build. They require massive interdisciplinary teams of physicists, engineers and IT specialists. The Large Hadron Collider (LHC) at CERN near Geneva – the largest accelerator currently being built – will cost a staggering SwFr 4.5bn (about €2.9bn). No fewer than 1300 giant superconducting dipole magnets, each with fields of more than 8 T, will be used to steer the protons around the machine, which will be housed inside the existing 27 km long tunnel of the now-defunct Large Electron Positron (LEP) collider.

My first encounter with accelerators took place in 1972 when I began working on the Intersecting Storage Rings (ISR) at CERN. I had come to Geneva straight after finishing my doctorate at the Queen's University of Belfast in Northern Ireland and had absolutely no experience with accelerators. For some reason, however, I was hired as one of the engineers in charge of operating the storage rings. I was not even a particle physicist, having trained as an electrical engineer.

Since then I have worked on two other, very different machines – LEP and the LHC. During that time, I have wit-

Massive attack – CERN's Large Hadron Collider, which will come online in 2007, will collide protons at energies of up to 14 000 GeV. The protons will be guided around the 27 km circumference ring by superconducting dipole magnets and focused by superconducting quadrupole magnets like the one shown here. The LHC will search for the Higgs boson and supersymmetry.

nessed some fascinating advances in the design and technology of particle accelerators. But with all three of these machines the guiding principles of accelerator design have stayed the same. The aim is to generate as many productive particle collisions (or “events”) as possible, at the highest possible energy, using the most appropriate particles for the job (see box on page 41).

My early years: the ISR

Early machines – including the Alternating Gradient Synchrotron at the Brookhaven National Laboratory in the US – tended to favour protons, but many accelerators since then have collided electrons or even heavy ions, such as gold. The ISR was a 26 GeV proton–proton storage-ring collider – the first of its type ever to be built. Although I was not around when it was designed, working on this machine was one of the most fulfilling experiences I have ever had. Some experts had initially feared that it would be impossible to keep the protons circulating for long enough to do experiments. But when I arrived in 1972 beams of about 8 A – roughly 40% of the design value – had already been stored and their lifetimes were reason-

ably good. The early worries had proved unfounded.

Over the next seven years, I became totally enthralled and absorbed by the physics of this wonderful machine. The ISR consisted of two intersecting rings each with a circumference of 954 m and covered an area about as big as two football pitches. It was not an accelerator as such, but stored high-energy protons that were fed into it from CERN's existing Proton Synchrotron (PS), which had been built in 1959. The ISR beams could be made to collide in eight interaction regions at energies of 11–26 GeV per beam – a range that was dictated by the energy of the PS.

X marks the spot – a crossing between two beam pipes of CERN's Intersecting Storage Rings, which was the world's first proton–proton collider. The machine consisted of two proton storage rings that crossed at eight separate places. It operated from 1971 to 1984 and accelerated the proton beams to a maximum total collision energy of 63 GeV.

After several years of operation and successive tweaks to the vacuum system, the ISR's beams could be made stable for very long periods of time. Indeed, towards the end of its life, the ISR's beams remained intact for a whole week. As for the beam current, it was enormous: 60 A was not uncommon. These large currents could only be stored in the rings by “stacking” the shots coming from the PS along side each other. As a result, the beams of particles were uniformly distributed around the circumference of the machine, rather than grouped in individual bunches.

The stability and long lifetime of the ISR made it the perfect place to search for unusual signals from the beam. Sensitive electrodes, for example, could pick up small “stochastic” electromagnetic signals that register the average conditions of the beam, such as its density. Using a high-frequency “kicker” that sends out rapidly varying electric and magnetic fields, it proved possible to control the stochastic signals – and hence the properties of the beam. This approach formed the basis of the “stochastic-cooling” technique, which can shrink the size of the beam in all three spatial dimensions and therefore increase the collision rate. (The beam is said to have been “cooled” because the particles occupy a smaller equilibrium volume, just as gas molecules would in a cooled balloon.)

We tested the cooling system by switching the kickers on and off repeatedly for long periods of several hours. Whenever the kickers were on, the event rate increased by a measurable amount. This indicated that the density of protons in the beam had risen during these periods and hence that the beam had successfully been cooled in the vertical plane. So a rudimentary stochastic-cooling system had been developed, even though the ISR was not the ideal place for it.

This “proof of principle” eventually proved crucial to CERN's decision in the late 1970s to convert its Super Proton Synchrotron (SPS) into a proton–antiproton collider. Sto-

chastic cooling was successfully applied to the antiprotons that were stored in the Anti-Proton Accumulation Ring and then accelerated through the PS into the SPS. The highly successful operation of the SPS with protons and antiprotons led to the discovery of the W and Z particles, for which Simon van der Meer and Carlo Rubbia shared the 1984 Nobel Prize for Physics (see “Carlo Rubbia and the discovery of the W and the Z” by Gary Taubes *Physics World* January 2003 pp23–28).

Next stop LEP

On my first visit to the ISR tunnel in 1972 I had been greatly impressed by the complexity and size of the machine. Imagine my reaction seven years later when I read the first design proposal to build a machine – the Large Electron Positron Collider (LEP) – that would be nearly 30 times larger. I was overawed and immediately decided to join the collider's then small design team.

LEP was a new adventure for CERN because all previous accelerators and storage rings built there had involved protons. Various designs of the LEP machine were produced until eventually in 1981 we received approval to build a version that would have a circumference of 27 km. But LEP was not only going to be big. It also set other technical challenges.

One critical difficulty centred on the unwanted nonlinear forces (or “kicks”) that are exerted by one beam whenever it collides with another travelling in the opposite direction. Caused by the strong nonlinear electromagnetic fields associated with the charge distribution of the beam, these forces can be characterized by a “beam–beam parameter” $\xi \sim I/\sigma$, where I is the intensity of the beam and σ is its cross-sectional area. Accelerator scientists strive for the brightest possible beams with the smallest possible cross-section – in other words they want to maximize ξ .

At low intensities, σ remains constant, which means that ξ increases linearly with intensity. At higher intensities, however, both the cross-section and the intensity increase at the same rate, which means that ξ eventually reaches a constant, maximum value. Since the event rate, or luminosity, of two counter-rotating beams of the same intensity is proportional to ξI , the plateau in ξ makes the luminosity and intensity rise hand in hand.

The maximum possible value of ξ – and hence luminosity – cannot be determined from first principles, but depends instead on how well the collider has been tuned. Dubbed the “beam–beam limit”, it dictates the ultimate performance of LEP and nearly all electron–positron colliders as well. In LEP's early days, the maximum value of ξ was about 0.015, whereas after many years of experience and optimization a record value of 0.083 was reached. This is also the highest value ever achieved in a collider.

Further increases in the beam intensity can, however, cause the charge distribution of the beams to become non-Gaussian, with long “tails”. This results in a catastrophic “blow up” of the beam cross-section, which swamps the detector with a massive background signal and forces data-taking to stop.

LEP came online in 1989 and was designed to be operated in stages. It first generated beam energies of about 50 GeV, which allowed precise measurements of Z bosons to be made. The energy was supplied by 128 room-temperature copper cavities located at four discrete points along the circumference of the beam pipe. Its second stage was reached in 1996

Accelerator basics

Different by design – (from left to right) the Relativistic Heavy Ion Collider at the Brookhaven National Laboratory in the US smashes gold particles head on; the Tevatron machine at Fermilab near Chicago collides antiprotons with protons that have been produced by smashing H^+ ions from this linear accelerator into carbon; the electron-positron linear collider at the Stanford Linear Accelerator Center in California; part of the planned TESLA electron-positron collider at the DESY lab in Hamburg.

The name of the game in experimental particle physics is to obtain as many productive collisions, or “events”, as possible in the hope that one of them produces something interesting. The number of events in a given time (or “event rate”) is highest if an accelerated beam is smashed into a stationary solid target. This approach was used at CERN’s Super Proton Synchrotron and at the “main ring” at the Fermi National Accelerator Laboratory near Chicago – and indeed for all machines with single-beam energies of up to about 500 GeV. But to obtain higher-energy beams, it makes sense to accelerate separate beams of particles in opposite directions and collide them with each other. Head-on collisions of this type have a much higher impact energy, which can therefore be transformed into heavier particles.

The only problem with particle-particle collisions is that the density of a beam of particles is much less than that of a stationary target, which leads to a low event rate. To increase the number of events, all machines that have been built since the early 1970s have been designed to allow bunches of particles to interact over and over again. This is done by sending bunches in opposite directions around a circular collider so that they crash into each other once every lap.

Most particle accelerators either collide electrons and positrons or protons and antiprotons, but if all things were equal, the electron would be the particle of choice. It has no internal structure, which means that it produces “clean” collisions that are easy to study.

However, the electron is about 2000 times lighter than a proton, which means that it loses about a factor of 10^{13} more energy per turn through synchrotron radiation than a proton. (The energy lost per turn is proportional to $(E/E_0)^4$, where E_0 is the rest mass of a particle of energy E .) The problem with protons is that they are made up of quarks and therefore produce complex collisions.

The Large Electron Positron (LEP) collider at CERN, for example, produced synchrotron radiation losses of about 3 GeV per turn when operating at energies of 100 GeV. If protons had been used, the equivalent figure would have been just 0.003 eV. The energy losses were replenished by giving the particles an electromagnetic “kick” on each turn using “resonant cavities”. These devices generate radio-frequency electromagnetic fields that are synchronized with the passage of the bunches.

The energy losses due to synchrotron radiation can also have an adverse effect on machines – like the forthcoming Large Hadron Collider – that use superconducting magnets to bend the particles around the ring. Any radiation losses have to be absorbed by a powerful cryogenic system to ensure that the magnet stays cold enough to superconduct. Due to the large synchrotron-radiation losses produced by high-energy electron circular accelerators, LEP was surely the largest and last accelerator of its kind. For colliding electrons and positrons at energies much higher than LEP the only practical solution is a linear collider.

when LEP was generating energies of 80–86 GeV per beam – high enough to allow pairs of W^+ and W^- bosons to be produced. In the final years of LEP – when the emphasis of the machine was placed on finding “new” physics – the beam energy was pushed significantly above the design value to a maximum of 104.5 GeV.

These high-energy runs, which culminated in the autumn of 2000, gave some tentative signs of the Higgs particle (see “CERN extends search for Higgs” *Physics World* October 2000 p5). Despite the possible evidence, however, LEP was finally shut at the end of 2000 to make way for construction of the LHC. Verification of these results will need higher-energy machines like the proton-antiproton Tevatron collider at the Fermi National Accelerator Laboratory near Chicago, or – more likely – the LHC itself.

Planning for the LHC

Following the closure of LEP in 2000, the remaining 150 or so staff who were working on this machine turned their full attention to the LHC project. When it comes online in 2007, the LHC will collide protons (and later heavy ions as well) with a total energy of 14 000 GeV. These huge collision energies, which are almost 100 times greater than those achieved at LEP, will usher in a new era in high-energy physics. We can expect physicists to detect supersymmetry and the elusive Higgs particle, which is believed to provide all other particles with mass.

The LHC beam will consist of a “train” of about 3000 bunches of protons, each of which will be about 8 cm long. Each train will be followed by a gap lasting 3 μ s, in which no beam is present. The 8.3 T superconducting dipole magnets

Tour de force – the ALEPH experiment at CERN's Large Electron Positron (LEP) collider, which ran from 1989 to 2000. ALEPH initially measured the properties of the Z boson that were produced at collision energies of 91 GeV. It was later tuned to run at energies of up to 200 GeV to study the production of pairs of W⁺ and W⁻ bosons. All three particles are carriers of the weak force. The ALEPH experiments is shown here being dismantled.

that will steer the protons around the ring will produce by far the highest fields ever used for a high-energy collider. The fields will be created by circulating currents totalling 13 000 A around superconducting coils that have been chilled with liquid helium to just 1.9 K. The cutting-edge technology required to make these magnets has now been mastered by European industry, with the equipment being built by firms such as Alstom (France), Ansaldo (Italy) and Noelle (Germany). The dipoles are currently arriving in large numbers at CERN, where they are being subjected to rigorous testing for mechanical, electrical and magnetic integrity.

In a machine with such an enormous amount of stored electrical energy – and enough beam energy to melt almost 500 kg of copper – safety issues and protection systems are crucial to the reliability and even survival of the equipment. Using superconducting magnets to steer intense beams of protons with a large destructive power could be dangerous. The beams could actually destroy large sections of the machine unless suitable protection systems are in place.

One of the biggest concerns centres on any unwanted or unforeseen sources of heat, which could arise from local faults or movements in the superconducting coils or even though the impact of a small portion of the high-energy beam on the magnet. The heat could locally warm the superconductor and turn it into a “normal” conductor, massively increasing its resistance in the process. Fortunately, CERN's engineers have developed a system to protect the magnets from such an event, which is known as a “quench”. The system first protects the magnet in which the quench occurred and then safely extracts the energy from all of the other magnets that are connected to it in the same sector of the ring.

The “quench-protection system” initially detects any increase in voltage across the coil of the magnet, which indicates that it has gone from a superconducting to a normal state. Over the next few milliseconds, the system sends current into “heating blankets” that warm the whole magnet coil and distribute the heat throughout its entire volume. This reduces the local heat density caused by the fault and stops

the coil from melting at the point where the quench occurred. After about a second, the coil of the quenched magnet is short-circuited by cold diodes to stop it receiving any of the energy stored in the other magnets of the sector. Eventually – after about 2 minutes – the stored energy in all the magnets of the sector is extracted into resistors.

Without this protection, any quench would almost certainly destroy the magnet. But with the protection in place, a quench simply makes the beam go down, causing at worst a few hours' disruption to users. Clearly, this protection system must be ultra-reliable to prevent the magnets from being destroyed: even tiny deposits of the proton beam on the superconducting coil can induce a quench. Larger quantities of lost beam can severely damage the magnets and any other material that gets in its path.

To protect the sensitive elements of the LHC, beam losses are deliberately localized at the “collimation system”. Consisting of pairs of composite-carbon “jaws” that are positioned very close to the beam, the system ensures that any particles that deviate from it are lost here and not at sensitive elements such as the magnets.

Given the extremely high destructive power of the LHC beam, the only way of getting rid of its total energy is to use “beam dumps”. There will be two such dumps on the LHC – one per beam. They will be housed in dedicated caverns at the end of 600 m long tunnels that run off tangentially to the main collider tunnel. Each beam dump will extract the beam from the machine, dilute its energy density and transfer it into a 7 m long steel-coated carbon cylinder that is some 70 cm in diameter.

To send the beam into the dump, a pulsed electromagnetic field will first be used to deflect the beam horizontally from its normal path. The field will be switched on at the moment when the 3 μ s gap between the beam's 3000 bunches occurs. This will ensure that no particles experience the field as it builds up to its maximum value but that all 3000 bunches to be dumped experience the full power of the deflecting field, which is held constant as the bunches pass by.

The field, which will be produced by 15 fast pulsed magnets, will deflect the beam horizontally by about 0.27 mrad for a beam with an energy of 7000 GeV. This will be enough to direct the beam into the magnetic gap of special magnets that then deflect it vertically by 2.4 mrad. The combined horizontal and vertical deflections “extract” the beam into the dump tunnel. Special elements in the tunnel increase the cross-section of the beam, and thereby reduced its density before it hits the face of the dump block.

Another design challenge with the LHC is to ensure that the fields of the superconducting magnets vary as little as possible as a function of the beam's transverse displacement, x . Unfortunately, there is always some nonlinearity in the fields. For example, the fields from “residual quadrupoles” increase linearly with x , while those from residual sextupoles increase with x^2 . These nonlinear fields, which depend critically on the how the superconducting coils are arranged, can cause the particles in the beam to become unstable in the horizontal and vertical directions, and eventually crash into the wall of the collider. In fact, they would actually be lost at the site of the collimation system.

Therefore, when the magnets are built, great care must be taken to minimize the nonlinear fields by careful winding of the coils and by ensuring that they are precisely located.

The magnets are also equipped with “correction elements” that are independently powered and can compensate for any local imperfections in the field. Furthermore, each magnet is subjected to a careful measurement procedure to quantify the nonlinear fields and identify any magnets that lie outside the required tolerance.

Another potential problem with superconducting cables is that they subject the beams to nonlinear fields that vary with time, as well as position. These fields, which depend on the recent magnetic history of the magnet, can be corrected using online measurements from the beam itself. The fields can also be evaluated by making magnetic measurements in real time on a “reference” magnet that is situated outside the ring but is otherwise identical to – and powered in series with – the ring magnets.

Fast forward to the future

High-energy physicists have for several years been considering what to build after the LHC. One option is a “neutrino factory”, which would consist of muons that have been accelerated and held in a storage ring. The muons would then decay to create an intense neutrino beam that would be directed towards a detector thousands of kilometres away (see “Neutrinos aim for the big time” *Physics World* August 2002 pp5–6). Other options include the Very Large Hadron Collider – a circular machine some 10 times bigger in circumference than the LHC – and an electron–positron linear collider.

The consensus is that the next machine should be an electron–positron linear collider with collision energies of 0.4–1.0 TeV or more. The machine would consist of two linear accelerators that accelerate electrons and positrons in opposite directions and collide them with a very small beam cross-section. Just as LEP provided precision measurements of the Z and W particles that had been discovered with the earlier SPS, so the electron–positron linear collider could provide precise measurements of the Higgs boson and supersymmetry that might be discovered with the LHC.

An electron–positron linear collider would have one key advantage over the LHC: electrons and positrons have no internal structure, which leads to very “clean” collisions. However, the particles must travel in a straight line – rather than a circle – to prevent them from losing lots of energy by emitting synchrotron radiation. The most crucial components in a linear collider are therefore the high-frequency cavities that must accelerate the beams from rest to high energy in a single trip down the linear accelerator. The final energy is simply the accelerating gradient – measured in megavolts per metre – multiplied by the total length of the collider.

Despite some friction over whether it should be built in Europe, the US or Japan (see “Emotions run high in race for collider” *Physics World* August 2001 p10), there is international agreement that the electron–positron collider will be built as a worldwide collaboration. Three main projects are currently

Testing times – one challenge with the LHC is to ensure that the fields of the superconducting magnets vary as little as possible as a function of the beam’s transverse displacement. Any “nonlinear” fields, which depend critically on how the superconducting coils are arranged, can cause the protons in the beam to become unstable and crash into the wall of the collider. The coils must therefore be positioned as precisely as possible.

being considered.

- The TESLA project in Germany will use a superconducting radio-frequency acceleration system with a gradient of 25 MV m^{-1} at a frequency of 1.3 GHz to create beams with an energy of 500 GeV. A possible extension of this first-phase collider to 800 GeV is possible by replacing the radio-frequency system with an upgraded 35 MV m^{-1} system.

- A joint US/Japan project at 6–11 GHz and gradients of $30\text{--}50 \text{ MV m}^{-1}$ using copper cavities at room temperature. The present design energy of this collider is 500 GeV, but it could be extended to 1000 GeV by adding cavities.

- A CERN study (CLIC) for a machine that would have accelerating gradients of 150 MV m^{-1} (generated by an ingenious two-beam scheme) and operate at 30 GHz with an energy range of 500–3000 GeV depending on the total length.

The technology that is associated with TESLA is currently the most advanced of these three projects. The CERN design has the most potential, but still needs further R&D work. The TESLA and the US/Japan projects are expected to have a price tag of about SwFr 5bn (about €3bn) – slightly more than the

LHC. Although a precise estimate of the cost of CLIC is not yet available, it will probably be about the same as the other two projects but produce much higher energy. The final choice is expected to be made by a specially elected international committee next year. To ensure that the collider comes online as planned in 2015, construction will have to begin as soon as possible after a site is chosen in 2007.

The next 10–15 years will be exciting for the particle-physics community. The LHC will take us into a new energy era with accelerator technologies pushed to the absolute limit. We can anticipate wonderful new physics results as the machine performance is pushed towards its design value – and hopefully beyond. Looking further ahead, the design and construction of the next linear collider will present further enormous challenges to the ingenuity and inventiveness of the world’s accelerator builders. History has shown, however, that they have met – and gone beyond – similar challenges in the past.

Further reading

- P J Bryant and K Johnsen 1993 *The Principles of Circular Accelerators and Storage Rings* (Cambridge University Press)
- A W Chao 1993 *Physics of Collective Beam Instabilities in High Energy Accelerators* (Wiley-Interscience, New York)
- A W Chao and M Tigner 1998 *Handbook of Accelerator Physics and Engineering* (World Scientific Publishing, Singapore)
- S Y Lee 1999 *Accelerator Physics* (World Scientific Publishing, Singapore)
- S Tazzari and M Ferrario 2003 Trends in high-energy particle accelerators *Rep. Prog. Phys.* **66** 1045–1094
- H Wiedemann 1993 *Particle Accelerator Physics* (Springer, New York)

Steve Myers is in the Accelerators and Beams Division, CERN, CH-1211-Genève-23, Switzerland, e-mail steve.myers@cern.ch