5

Accelerators and beams

In a typical experiment, particles are prepared and allowed to interact and the interaction products are observed. The products may be the original particles, scattered in new directions, as in the experiments conducted by Rutherford and his colleagues, a process called elastic scattering. Alternatively, because energy may be converted into matter, the products may be new particles. At high energies this is the most likely outcome, with many particles being produced. Because these new particles are usually unstable and decay rapidly, experimenters study their decay products. In this chapter, I will consider how physicists produce particles to be used as projectiles.

Natural and man-made accelerators

For many years, the only particles available for experiments were those that occurred in nature, with energies that could not be controlled by the experimenter. In most cases the particle's energies were also not unique; only sources with a range of energies were available, such as the products of the radioactive decay of nuclei, as used in the experiments that led to the discovery of the nucleus. Here, the energies are in the MeV region, which is far too low for most modern particle physics experiments, although some experiments are conducted using the enormous flux of neutrinos emitted from nuclear reactors. These produce energy from the fission of radioactive nuclei, whose decay products are themselves subject to beta decay; the energies of the emitted neutrinos are typically a few MeV.

Much higher energies are available in the secondary particles produced by cosmic rays. Cosmic rays are particles, mainly protons, which impinge on Earth from space; their precise origin is unknown. Some have energies far greater than can be produced in laboratories. They interact with atoms in the atmosphere to produce particles that can be detected at the surface of Earth. Experiments using cosmic rays have yielded many remarkable discoveries, including that of the positron. They were also key to finding the first unstable short-lived particle that does not occur naturally, a boson with spin 0 called the pi meson or more usually just the 'pion' (denoted by the Greek letter pi, π). (*Meson* denotes a boson that experiences the nuclear strong interaction.) Pions exist in three states, one with positive electric charge (π^+), one with negative charge (π^-) and one without charge (π^0).

Although the study of cosmic rays and their products continues to be an interesting subject in its own right, nearly all experiments now use man-made beams produced by accelerators. An accelerator starts with a source of low-energy particles and boosts their energy by applying electric forces. Accelerators can only directly produce beams of stable, electrically-charged particles, such as protons, anti-protons or electrons but there are indirect methods for producing beams of unstable particles and even for producing some neutral particles, including neutrinos.

To begin the acceleration, low-energy particles are injected into an accelerator from a high-intensity source, such as a heated metal filament in the case of electrons. It may typically take several stages, each incrementally boosting the energy, to reach the maximum final energy. Accelerators have great advantages over natural sources of particles: the projectiles produced are of a single type and their energies can be controlled by the experimenter. The earliest accelerators were built in the 1930s but could only produce particles with energies of a few MeV. By the 1950s, energies of few GeV could be reached. The highest-energy machines currently available produce particles with energies of a few TeV and these can explore matter at the effective temperatures that existed very shortly after the Big Bang.

Linear and cyclic accelerators

Early examples of accelerators used static electric fields to accelerate particles. Such machines are limited to relatively low energies and so particle physicists turned to a different technique to produce beams for research: they boosted the initial energy of the particles by repeatedly applying a radio frequency (rf) electric field. A variety of machines use rf fields but the general principles that govern them are similar.

An example of a linear accelerator (also called a linac) using protons is the accelerated particles as shown in Figure 3.1. The protons from a low-energy source pass through a series of metal pipes, called drift tubes, that are located within another pipe, the vacuum pipe, which is maintained at a very high vacuum to minimise interactions between the protons and gas molecules.

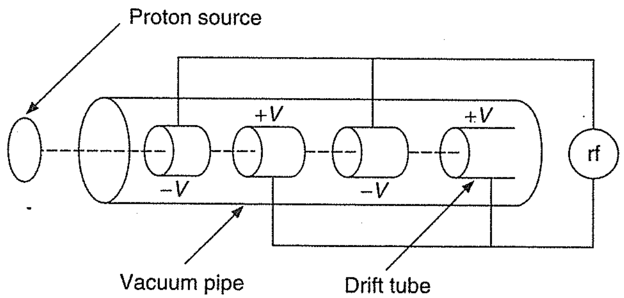

Figure 3.1 Acceleration in a linear proton accelerator

The drift tubes are connected successively to alternate terminals of an oscillator that produces an rf electric field. In Figure 3.1, V denotes the electric potential of the drift tubes and determines the force on the proton. With the signs of the potentials as shown, the protons will be accelerated towards the first drift tube. Because the tube is maintained at a constant potential, the protons will continue to drift through it at a constant velocity. If the rf oscillator can change the field direction (and so the signs of the potentials) as the protons pass through the first tube, the protons will again experience an accelerating force on their way between exiting the first tube and entering the second one, and so on. To be continuously accelerated, the particles must remain synchronised with the oscillating rf field as they pass through each successive tube. Since the speed of the protons increases as they move between the drift tubes, they will take a decreasing amount of time to travel the same fixed distance. If the oscillating field is of constant frequency, each subsequent drift tube must be longer, to ensure that the particles spend the same amount of time in each of the drift tubes. The beam must also remain focused, so that the particles do not hit the sides of the drift tubes during their passage through the accelerator. Proton linacs are often used in particle physics as injectors, that is, they produce proton beams of moderate energy that are injected into a more powerful machine, where they are accelerated to much higher energies.

For electrons, which rapidly approach the speed of light as they are accelerated, a variation of the method is used. The accelerator consists of a straight tube, in the form of a series of cylindrical metal cavities. Power is fed to the accelerator from devices called klystrons, which produce radiation in the form of pulses in the microwave region of the electromagnetic spectrum. These pulses are transported to the accelerator, where they generate an oscillating electric field pointing in the direction of the metal tube (similar to that produced by the rf oscillator in a proton linac) and a magnetic field circling the tube's interior. The metal cavities in the accelerator are of constant length and so the frequency of the microwaves is adjusted to ensure that the electrons arrive at each cavity at the optimal time to receive the maximum energy boost. Vitally, the magnetic field helps to keep the beam focused. The largest linac in the world is the Stanford Linear Collider (SLC) at the Stanford Linear Accelerator Laboratory (SLAC) in California, USA. It consists of eighty thousand copper cavities separated by copper discs, each with a small hole at the centre to direct the beam. The SLC has a maximum energy of 50 GeV. To achieve this energy it is very long – over three kilometres. Remarkably, it is built very close to the infamous San Andreas earthquake fault line!

In contrast to linear accelerators, cyclic accelerators use a circular, or near-circular, configuration. The earliest example is the cyclotron (see Figure 3.2), which consists of two d-shaped sections, contained in a high vacuum, across which an rf electric field is established. The dees are sandwiched horizontally between the poles of a magnet that produces a uniform magnetic field perpendicular to the dees. Charged particles are injected into the

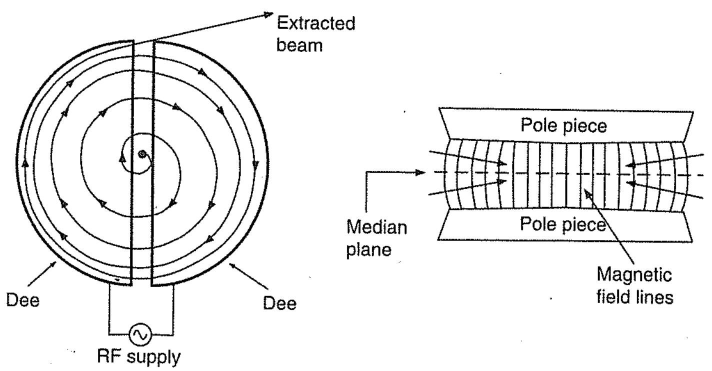

Figure 3.2 Schematic diagram of a cyclotron

machine near its centre, in a plane perpendicular to the magnetic field. Because the particles are moving perpendicular to the field, the force constrains them to follow outward spiral trajectories and the particles are accelerated each time they pass across the gap between the dees. At the maximum radius, which corresponds to the maximum energy, the beam is extracted. The shape of the magnetic field is curved at the edges of the magnet pole pieces. This produces forces, shown as arrows on the diagram, which move particles closer to the median plane, creating a stable beam. The first cyclotron was built by the American physicist Ernest Lawence in 1929 and was a mere thirteen centimetres in diameter.

The cyclotron works on the principle that particles always take the same amount of time to complete a circuit of the accelerator. However, as the speed of the particles increases, the effects of relativity become apparent and that assumption no longer holds; to maintain acceleration, more force is required with each circuit. Over time, the particles cease to be synchronised with the rf field and arrive late at the gap and so do not receive an energy boost. In practice, the cyclotron is only capable of accelerating particles to speeds that are just a few per cent of the speed of light.

Cyclotrons played an important role in the early days of particle physics; Fermi discovered the first hadron resonance using one. Because of their limits, they have since been replaced in particle physics research by cyclic accelerators, *synchrotrons*. Synchrotrons operate similarly to linear accelerators but the acceleration takes place in a circular or near-circular orbit rather than in a straight line. The particles travel in an evacuated tube, the beam pipe, and are kept on their paths by the fields created by an array of magnets that bend the beam, an assembly referred to as the beam line. As the beam repeatedly traverses one or more cavities placed in the circle, where energy is given, the particles are accelerated.

Charged particles travelling in a circular orbit continuously emit electromagnetic radiation, called synchrotron radiation in this context. For a given energy, the losses due to synchrotron radiation increase very rapidly as the mass of the particle decreases and so are very severe for electrons, which are very light. To compensate for the losses, the accelerators input large amounts of rf power, which limits the energies that electron synchrotrons can reach. No similar limit affects electron linacs.

The momentum of an orbiting particle in a synchrotron is proportional to the product of the strength of the magnetic field times the radius of the orbit. Because the momentum increases during acceleration, the magnetic field must be steadily increased if the radius of the orbit is to remain constant, in contrast to the cyclotron, which uses a constant magnetic field. The maximum magnetic field that can be produced over an adequate region is limited and so very-low-temperature technology ('superconducting' or 'cryogenic' technology) is used to reduce power consumption. To achieve very high energies and limit synchrotron radiation, the radius of the accelerator ring must be very large. The Tevatron accelerator at the Fermi National Laboratory in Chicago, which accelerates protons and anti-protons to energies of 1 TeV, has a radius of one kilometre.

In the course of its acceleration, a beam may traverse the ring many millions of times before reaching its maximum energy. Consequently, physicists strive hard to keep the particles in stable orbits; if they meander, they will not achieve optimal acceleration and, if they strike the walls of the vacuum tube, may even be lost from the beam. The particles are accelerated in bunches, each being synchronised with the rf field; oscillations in their trajectories are controlled by another series of magnets, placed at intervals around the beam line, which act like optical lenses. Each magnet focuses the beam in one direction and defocuses it in a direction perpendicular to that, so alternate magnets have their field directions reversed to maintain a stable orbit.

Fixed-target machines and colliders

There are two types of linacs and synchrotrons – 'fixed-target' and 'colliding beam'. In fixed-target machines, particles are accelerated to the highest operating energy and then the beam is extracted from the machine and directed on to a stationary target, usually a solid or a liquid. In this respect, fixed-target machines operate like the earlier cyclotrons. Much higher energies have been achieved for protons than for electrons, because of the large synchrotron radiation losses inherent in electron machines. The intensity of the beam produces large numbers of interactions that can either be studied in their own right or used to produce secondary beams. From $E=mc^2$, it follows that the heavier the particles that are produced, the higher the energy that is required, which reveals a disadvantage of fixed-target experiments when very high energies are needed. When a fast-moving projectile hits a stationary target, some of the energy of the projectile is transferred to the target. This 'energy of motion' of the target is now no longer available to produce new particles and a 'law of diminishing returns' soon sets in, where an increasing fraction of the initial energy is wasted.

This disadvantage may be overcome by using a moving target, to produce a head-on collision between two beams of particles. The SLC and the Tevatron are both colliding beam machines, more commonly known as *colliders*. The two beams often consist of particles and their anti-particles, for example electrons and positrons, circulating in opposite directions in the same machine. The beams are allowed to intersect at a fixed number of points around the accelerator circumference, where experiments are sited. The two beams usually have the same energy and so their combined energy is entirely available for particle production. The products of the collision are scattered rather uniformly without significant momentum in a preferred direction, which has consequences for the way in which particles are detected.

Colliding beam experiments are not without their drawbacks. The colliding particles have to be stable and charged (because they have to be capable of being accelerated by electric fields), which limits the interactions that can be studied. In addition, the probability of a collision in the region where the two beams intersect is far smaller than that achieved in fixed-target experiments, because the particle densities in the beams are very low compared to those in a solid or a liquid target.

Because of the need to produce higher energies in modern particle physics research, almost all the new machines built today are colliders. To date, the largest collider is the Large Hadron Collider (LHC), which collides protons with protons. It has been built deep underground, at depths from 50 to 175 m, in a tunnel at the European Centre for Nuclear Research (CERN), at a cost of approximately £2.8 billion. The LHC is massive, with a circumference of 27 km and straddles the French-Swiss border, near Geneva. Each beam has been designed to have a maximum energy of 7 TeV. Building the tunnel housing the LHC and filling it with equipment was a major international civil engineering project. It became operational in late March 2010, although initially at half the full beam design energy while commissioning tests take place.

Figure 3.3 shows a section of the beam line at the LHC. One of the cylinders is schematically shown open to reveal the two beams of protons, the bending magnets that keep them in circular orbits and the liquid helium system for the magnets. In total, over 1600 superconducting magnets are installed at the LHC, both to bend the beam and to focus it, with most weighing over twenty-seven tonnes. Approximately ninety-six tonnes of liquid helium are needed to keep the magnets at their operating temperature of 1.9 K, making the LHC the largest liquid helium cryogenic facility in the world.

A schematic diagram of the LHC and some of the other accelerators at the CERN site is shown in Figure 3.4. This

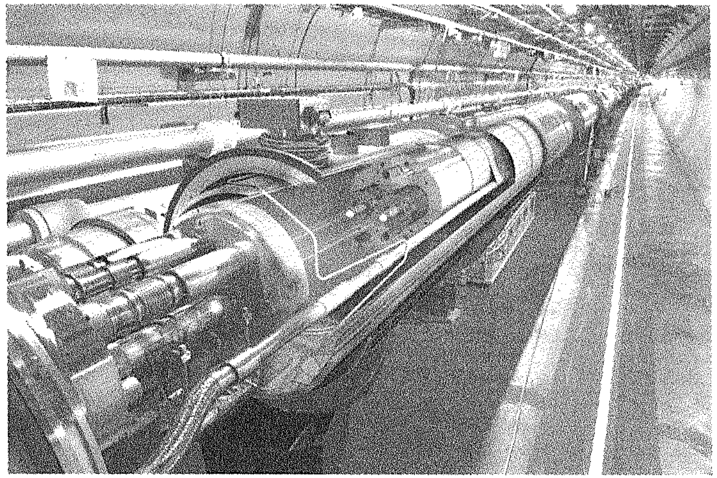

Figure 3.3 LHC tunnel and beam line (photo CERN-AC-091118801 by M. Brice; reproduced by permission of CERN)

illustrates the multi-stage process that particle physicists use to obtain high energies. The acceleration begins with a linac, whose beam is boosted in energy in the Proton Synchrotron Booster (PSB) and passed to the Proton Synchrotron (PS), a machine that is also used in lower-energy experiments. The beam is then transported to the Super Proton Synchrotron (SPS), where its energy is increased still further before being injected into the LHC. There are four beam intersection points in the LHC and experiments (denoted by the names ALICE, CMS, LHC-b and ATLAS) are located at each of these points. I will return to some of these experiments and their findings in later chapters. The neutrino beam, marked at the bottom of Figure 3.4, is directed to the Gran Sasso laboratory in Italy, 730 km away. The fact that the beam can travel such large distances is a good example of the very low probability of neutrino interactions mentioned in Chapter 1.

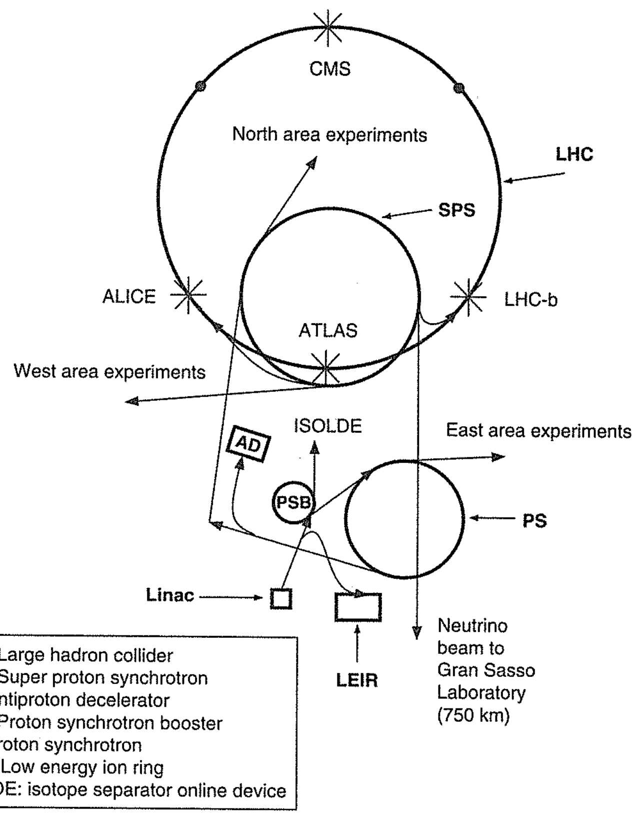

Figure 3.4 Schematic diagram of the accelerators at the CERN site

A particular type of collider has become of increasing interest in recent years. This is the so-called 'particle factory', in which the energies of two colliding beams are tuned so that the total available energy peaks at a particular value corresponding to the mass of a specific particle that physicists wish to study. At present there are two '*B*-factories', one in Japan and one in the United States, that

accelerate beams of electrons and positrons, whose total energy is tuned to produce copious numbers of the *B*-meson, a particle that may reveal details of the weak interaction. The *B*-meson is a very short-lived boson, with a lifetime of just 10^{-12} seconds. Even if it were moving at the speed of light, it would only travel a distance of about 3×10^{-4} m before decaying, too short a distance to be measured with even the most refined techniques currently available. Therefore, *B*-factories adopt a novel approach: the beams of electrons and positrons have unequal energies and hence different speeds. As a consequence, the 'debris' (which includes many *B*-mesons) that is produced at the point where the two beams annihilate is moving in the direction of the faster beam. When quoting particle lifetimes, it is always understood that the particle is at rest. However, if the particle is moving, there is an effect in special relativity called 'time dilation' which means that the lifetime as observed in the laboratory is greater than the 'at rest' value. For this reason, the *B*-mesons travel sufficiently far in the laboratory to be detected (although the distance is still only measured in millimetres). Essentially, the *B*-factory sacrifices some energy to produce new particles in favour of the ability to actually detect them.

Particle beams

Accelerators can only directly produce beams of stable charged particles, which in colliders often include anti-particles. Ideally, in a collider, beams of anti-particles should have comparable particle densities to those of the particles with which they collide and also have a single energy. But what is the source of anti-particles? For example, anti-protons could initially be produced in collisions of protons with a target but typically only one anti-proton will be produced for every million protons, so the resulting anti-proton beam would have a very low density and also have a wide spread of momentum, unless a way can be found of 'compressing' the beam.

An ingenious technique for doing this was first suggested by the Dutch applied physicist, Simon van der Meer, in 1968. It uses the electrical signals produced by small groups of particles within a larger bunch to drive an electromagnetic device that applies a correction (a steering pulse or kick) to reduce the momentum spread within the bunch. These individual corrections are applied continuously and over an extended time, so the average tendency of individual particles to move away from the other particles in the beam is reduced and the beam is 'compressed'. This is referred to as 'cooling' the beam. It is appropriate to think of this technique as cooling because the particles in the bunch can be characterised by an internal effective temperature, as we saw when discussing photons. If the average momentum of the bunch were to be subtracted from the momentum of each of its constituent particles, the particles would appear to be moving randomly; the more vigorous the motion, the 'hotter' the bunch, much like the molecules in a gas. Cooling times range from a second to several minutes, depending on the degree of cooling required. Van der Meer applied his technique to anti-protons produced at the SPS at CERN and was able to create an anti-proton beam. This beam was used in a proton–anti-proton collider experiment that discovered the $W^{\pm}$ and Z^{0} bosons. Van der Meer shared the 1984 Nobel Prize in Physics for his work on cooling, a rare example of the prize being awarded to an accelerator physicist/engineer.

Physicists are also interested in the interactions of unstable particles, such as pions. Beams of unstable particles can be formed, provided their constituents live long enough to travel appreciable distances in the laboratory. One way of doing this is to direct an extracted primary beam from a synchrotron or linac on to a heavy target. In the resulting interactions with the target nuclei, many new particles are produced. Some of these particles are electrically

charged and their trajectories can be guided by applying magnetic fields.

If the new beam is composed of unstable particles, it can be used to produce further beams from their decay products. For example, beams of muons may be made from the decays of pions. There are even ingenious methods to prepare beams of neutrinos from the weak decays of particles such as pions and muons.