IOP Concise Physics

Introduction to Beam Dynamics in High-Energy Electron Storage Rings

Andrzej Wolski

Chapter 1

Introduction

1.1 A brief history of electron storage rings and their uses

In the late nineteenth century, experiments with cathode ray tubes led to the discovery of the electron. The components of a cathode ray tube perform all the essential functions of the components in a modern accelerator and include: a particle source, some means of acceleration, components for controlling the trajectory of particles, and a way of detecting or observing the particles. By modern standards, however, the components in a cathode ray tube are extremely simple. A heated wire at a negative potential (the cathode) is the source of electrons, which are accelerated (in a vacuum) by a static electric field towards a fixed anode. The electrons may be steered using electric fields perpendicular to the accelerating field, or by a magnetic field generated using loops of wire carrying electric currents. The beam of electrons is observed either by the fluorescence of residual gas within the tube, or by the light produced when electrons strike a coating of fluoresent paint on the inner surface of the tube. In the next section of this chapter, we shall outline the corresponding systems in electron storage rings that are being constructed today. It is worth noting that nearly 30 years elapsed between the development of the first cathode ray tubes by William Crookes and others, around 1869, and the discovery of the electron by J J Thomson in 1897.

Accelerators were developed continually through the twentieth century [1], motivated by the desire to understand the physical properties of materials, and the laws of nature, at ever more fundamental levels. This required increases in beam energy and intensity, and improvements in the quality and stability of the beams being accelerated. One way to achieve a high energy is to use a sequence of electric fields. The strength of a static electric field is limited by the fact that any material will break down once the field reaches a certain value. However, by using oscillating fields, the ultimate energy is effectively limited only by the length of the accelerator.

doi:10.1088/978-1-6817-4989-1ch1

In a drift tube linac for example, particles pass through a sequence of conducting tubes. At any given moment, any two adjacent tubes are at potentials of opposite polarity: this means that the maximum potential can be kept below the breakdown limit. If the potentials of the tubes (and hence the fields between the tubes) are static, then particles are alternately accelerated and decelerated as they pass successive gaps, and there is no overall acceleration. However, if the potentials (and the fields) oscillate in synchronism with the particle motion, then it is possible to arrange for particles always to see accelerating fields, and for their energy then to increase over the entire length of the linac. Although linacs are still widely used today, they have a drawback in that the higher the energy required, the longer the accelerator has to be.

The invention of the cyclotron by Ernest O Lawrence in 1934 was a major step forward for achieving high energy: the key feature of a cyclotron is the use of a magnetic field to bend the particle trajectory into a spiral. This means that particles can be accelerated multiple times by an oscillating electric field, in the same way as in a linac; but the size of the accelerator is kept small by 'reusing' the same field. However, the energy that can be reached by particles in a cyclotron is limited by the size and strength of the magnets needed to maintain the spiral trajectory. Relativistic effects also lead to technical challenges in cyclotrons once particles reach energies at which they move at speeds close to the speed of light. Although cyclotrons are capable of accelerating particles to energies of several GeV, the magnets required are extremely large and can weigh tens or hundreds of tonnes. Nevertheless, cyclotrons are still widely used today for accelerating protons or ions (with mass larger than protons) for nuclear physics experiments, or for radiotherapy. However, electrons have a much lower mass than protons, so reach relativistic speeds at much lower energies, and this makes cyclotrons unsuitable for producing high-energy electron beams. As the energy of an electron increases, the radius of its trajectory in a magnetic field increases; however, its speed approaches a limit, the speed of light. Therefore, in a cyclotron, electrons start to take longer and longer to complete each turn of the spiral trajectory as their energy increases, until eventually their paths are no longer synchronised with the oscillating electric field needed to produce the acceleration.

Already in the first half of the twentieth century, it was realised that the limitations from relativity on electron beam energy could be overcome if particles could be kept on a circular, rather than spiral, trajectory as they were accelerated. This required the strengths of the magnetic fields controlling the trajectory to be increased in proportion to the momentum of the particles; conveniently, increasing the magnetic field created (by electromagnetic induction) an electric field that would accelerate the particles. With a suitable geometry, including some shaping of the poles of the magnet to control the field strength as a function of radius, a machine could be constructed that accelerated particles and at the same time confined them within a toroidal vacuum chamber. The resulting accelerator was known as a betatron. Although the concept dates back to 1922, the first successful machine was not demonstrated until 1940. Using betatrons, it became possible to accelerate electrons to tens of MeV; but the concept was soon superseded by the synchrotron. In a synchrotron, the magnetic fields used to guide the electrons are increased with beam energy, in the same way as in a betatron. However, rather than using electromagnetic induction to accelerate the particles, a synchrotron uses oscillating electric fields, in a similar way to a cyclotron. To compensate for changes in the speed of the electrons as they increase in energy, the oscillation frequency of the accelerating field is varied. In other words, the accelerating field oscillation is synchronised with the revolution frequency of the particles travelling around the ring: hence, the name 'synchrotron'. The first synchrotron was constructed (from a modified betatron) in the late 1940s, at the Woolwich Arsenal Research Laboratory in the UK. A second machine, purpose-built by the General Electric Company at Schenectady, New York, USA, followed soon after, and achieved an electron beam energy of 300 MeV.

After the first electron synchrotrons, a number of proton synchrotrons were developed with the aim of achieving beam energies of several GeV, motivated by the production of beams of high-energy particles for studies in nuclear physics. Synchrotrons were used to raise the energy of a beam injected at relatively low energy to the energy needed for an experiment, at which point the beam was extracted from the synchrotron and directed towards a target. However, it was soon realised that there were applications for high-energy beams (of protons or electrons) stored in synchrotrons for long periods. In principle, to operate the synchrotron as a storage ring all that was needed was to maintain the magnetic field strengths at constant values once the desired energy was reached. The use of synchrotron storage rings made it possible to perform colliding beam experiments, in which the trajectories of two beams of equal energy intersect at one or more points in the ring. This provides a significant advantage over fixed-target experiments, in which much of the energy of the incident particle is converted to kinetic energy of the collision products, and is therefore not available for generating new states of the interacting particles. When two particles collide with equal and opposite momenta, on the other hand, the total momentum is zero so all the energy is available for producing new states.

Colliders began to be developed in the 1960s. The first machine to be completed, AdA (Anello di Accumulazione, figure 1.1) was an electron–positron collider constructed in Frascati, Italy in 1961, by a team led by Bruno Touschek. The storage ring had a diameter of about 1.3 m and stored beams with energy 250 MeV. An electron–electron collider followed shortly after: VEP1 was operational by 1965 at the Institute of Nuclear Physics[1] in Novosibirsk, Russia, and achieved beam energies of 160 MeV in two storage rings each of 86 cm diameter. Numerous colliders have since been built, with each successive machine aiming for higher energy and/or luminosity. Machines have been constructed for colliding different combinations of lepton (electron or positron) and hadron (proton, antiproton, or ion) beams.

It has been known since the late nineteenth century that any charged particle will radiate energy in the form of electromagnetic waves when undergoing a change in

---

[1] INP is now named the Budker Institute of Nuclear Physics (BINP) in honour of its first director, Gersh Budker. Budker led the development of VEP1.

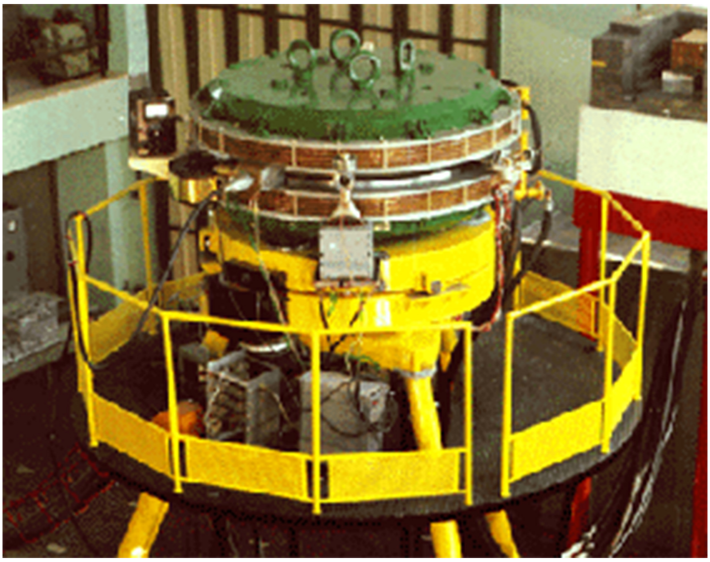

Figure 1.1. AdA (Anello di Accumulazione), the first electron–positron collider. AdA had a diameter of about 1.3 m, and stored beams with energy 250 MeV. It was constructed in Frascati, Italy, in 1961.

speed or direction. It was therefore expected, even before the first electron synchrotrons were constructed, that electrons in the magnetic fields in a synchrotron would emit radiation as a result of following a curved trajectory. The radiation from a beam of relativistic electrons is known as *synchrotron radiation* [2], and was observed for the first time in 1946, when light was seen emerging from the glass vacuum chamber of the General Electric Company synchrotron in Schenectady. Since synchrotron radiation limits the particle energy that can be achieved in an electron synchrotron, there were a number of studies to investigate its properties, and to confirm the theoretical predictions; the experience gained in handling synchrotron radiation for these studies opened the way for its use in a wide range of research. The value of synchrotron radiation lay in the intensity of the light that could be produced, compared to conventional sources, especially at short wavelengths. This enabled experiments to be performed that had previously been difficult or impossible, notably into the properties of materials (for example, auto-ionization in gases). Despite a growing community of researchers primarily interested in synchrotron radiation as a tool for scientific investigation, the use of synchrotron radiation was for many years ‘parasitic’ on machines built for high-energy physics; it was not until 1968 that Tantalus, a 240 MeV electron storage ring at the University of Wisconsin, became the first machine constructed specifically for the production of synchrotron radiation [3].

In the decades following the construction of Tantalus, the use of synchrotron radiation for scientific research became more widely established [4]. Many new facilities were built, and advances in technology enabled significant improvements in the brightness, stability, and spectral range of the synchrotron radiation that could be provided for users. Conventionally, synchrotron light sources are classified in four ‘generations’. First-generation light sources consist of electron storage rings constructed primarily for high-energy physics applications (i.e. colliders), with the synchrotron radiation produced by the dipole magnets being used parasitically. Although the colliding beams are tightly focused at the interaction point to generate high luminosity, the beam size in the rest of the ring tends to be relatively large: this limits the brightness of the synchrotron radiation produced from the electron beams in the dipole magnets. Brightness is a measure of the radiation intensity per unit area of the source, and is an important figure of merit for many light source users. Second-generation light sources are those (such as Tantalus) constructed specifically for the production of synchrotron radiation, but with the radiation coming principally from the dipole magnets used to steer the beam around the storage ring. The beam size in a second-generation light source can be smaller than that in a collider, but the brightness of the synchrotron radiation is still low by modern standards.

Advances in the design of the magnetic lattice in electron storage rings led to the development of third-generation light sources. In particular, it was found from detailed studies of the effects of synchrotron radiation on the electron beam producing the radiation that lattice designs were possible, allowing an improvement in brightness by several orders of magnitude (see section 3.3). At the same time, magnets were developed specifically for generating synchrotron radiation with desirable properties for light source users. These *insertion devices* (undulators and wigglers [2, 5]) consist of sequences of short dipole magnets of alternating polarity; an example of an undulator, in the Advanced Photon Source at Argonne National Laboratory in the USA, is shown in figure 1.2. A key feature of third-generation synchrotron light sources is that the storage ring is designed to optimise the production of high-brightness, short-wavelength radiation from insertion devices (although the radiation from the dipole magnets is often used as well).

Particles passing through dipoles and insertion devices in storage rings generally produce radiation independently of other particles in the beam, so that the intensity of the radiation is proportional to the number of particles. However, if the size of a bunch of particles is small compared to the wavelength of the radiation being produced, then the particles can radiate coherently, i.e. acting as effectively a single particle. The intensity of the radiation in that case is proportional to the square of the number of particles. Since the number of particles is potentially very large (of order 10^9 or more), the intensity of coherent synchrotron radiation can be many orders of magnitude larger than for incoherent radiation. Since bunches in an electron storage ring are typically several millimetres in length, any coherent radiation from an entire bunch has a large wavelength and does not propagate efficiently through the vacuum chamber. However, it is possible in some circumstances for substructures to develop within a bunch, leading to coherent radiation at wavelengths below a millimetre. A *free electron laser* (FEL [6–9]) is an accelerator designed to produce coherent synchrotron radiation, either from an electron storage ring, or from a linear accelerator delivering the beam for a long undulator. FELs are sometimes known as fourth-generation light sources. Short-wavelength (extreme ultra-violet or x-ray) FELs are based on linacs rather than storage rings, and have a

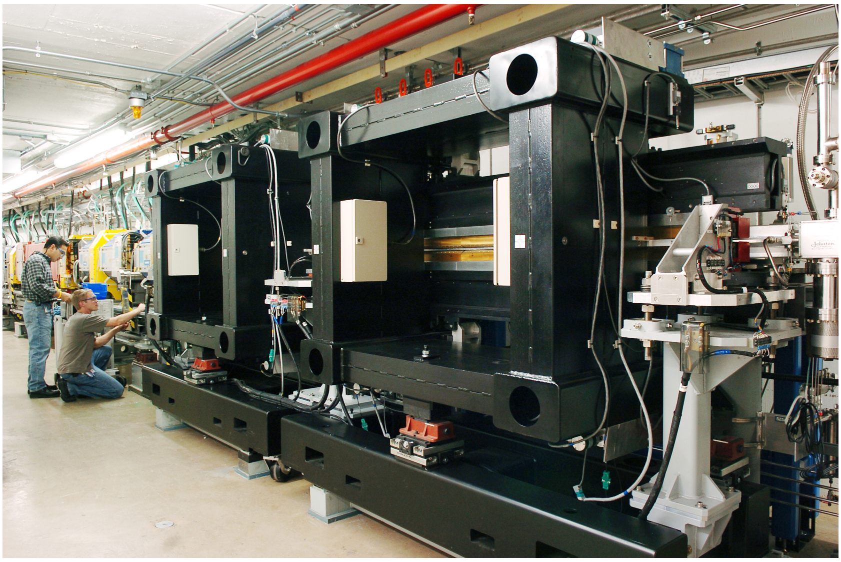

Figure 1.2. The canted undulator for sector 23 at the Advanced Photon Source. Sector 23 is operated by the General Medicine and Cancer Institutes Collaborative Access Team (GM/CA-CAT), a part of the Biosciences Division at Argonne National Laboratory. The GM/CA CAT has been established by the National Institutes of Health's National Institute of General Medical Sciences and National Cancer Institute to build and operate a national user facility for crystallographic structure determination of biological macromolecules by x-ray diffraction. (With permission from: Argonne National Laboratory.)

further advantage, in addition to the intensity of the radiation, in being capable of producing extremely short pulses of radiation, in some cases of order of a picosecond (10^{-12} s) or less. However, linac-based light sources are generally limited to serving only a few (perhaps three or four) users at any one time; third-generation light sources, however, provide much greater capacity. Depending on the size of the storage ring, a third-generation light source may serve several dozen user beamlines at once.

Developments in accelerator technology since the first-generation light sources have been matched by advances in the user beamlines: instruments designed for high-intensity, short wavelength radiation are required to steer and focus the beams of synchrotron radiation, to select particular wavelengths and polarisations, and to detect the radiation after interaction with a sample. The combination of accelerator and radiation beamline technology now available enables third-generation synchrotron light sources to support a highly diverse range of scientific research. The radiation can be used in various different ways [10–12]. For example, imaging can be performed in much the same way as in medical x-ray imaging, but the intensity of synchrotron radiation makes it possible to resolve much finer structures. The ability to select and combine images at different specific wavelengths also makes possible techniques that can provide images of structures on the scale of a few billionths of a metre. The ability to produce and select radiation at specific wavelengths from infrared to x-rays also enables spectroscopy techniques based on how different materials absorb or reflect light of different wavelengths. Detailed spectroscopic measurements can provide important information on the structure and properties of materials with applications in information technology, engineering, biology, and medicine. Finally, diffraction and scattering experiments can reveal the atomic structure of materials ranging from minerals and ceramics to polymers and proteins. Particularly in the life sciences, the ability to determine the structure of highly complex molecules has provided deep insights into numerous biological processes, often with important medical consequences (for example, in understanding how particular medicines work). The value of synchrotron light sources for fields as diverse as materials science, information technology, engineering, biology, and medicine is such that about fifty facilities are currently in operation around the world, with several new facilities either proposed, under construction or being commissioned. Electron storage rings seem certain to continue to play an important role in scientific research for many years to come.

1.2 General features and subsystems

The detailed design of a synchrotron storage ring will depend on the intended application, but the overall structure generally follows an established pattern [13]. The beam travels within a vacuum chamber in the form of a tube with an internal aperture of (typically) a few centimetres. The tube is bent in (roughly) the shape of a circle with a circumference that could be anything from a few hundred metres to several kilometres. Particles are guided and focused using magnets placed at intervals around the ring. Electrons travelling through magnetic fields lose energy by radiation (in the case of relativistic particles, this is termed *synchrotron radiation*); to replace the lost energy, radiofrequency (RF) cavities are used to provide an electric field that accelerates the particles on each revolution. The electric fields in the cavities oscillate at frequencies of (typically) a few hundred megahertz; the frequency of this oscillation must be synchronised with the revolution frequency of particles around the ring, so that a particle will arrive at a cavity at approximately the same phase of the electric field oscillation on every turn. In this way, the energy of each particle in the storage ring remains roughly constant. In electron storage rings, the particle energy is typically in the range from a few hundred MeV, to many GeV.

Understanding the principles of a synchrotron storage ring and understanding much of the beam dynamics requires some knowledge of the magnets and the RF cavities. However, storage rings rely for their operation on many other types of component; altogether, a storage ring will usually be constructed from several thousands, or tens of thousands, of separate components. Some familiarity with the different types of components needed for a storage ring is often helpful in understanding even some of the apparently more academic aspects of beam dynamics. This is especially true in the context of technical limits on beam parameters.

Components of all different types in a storage ring must, to some extent, work together; but for convenience, they are grouped into various subsystems. Thus, the magnets, perhaps with their power supplies, cables, and cooling systems, will form one subsystem; the RF cavities, with the RF power source, waveguides (for transporting the RF power to the cavities) and electronics for controlling the amplitude and frequency of the fields in the cavities, will form another subsystem. Further subsystems include the vacuum system, diagnostics, feedback systems, control system, injection system, and personnel protection system. In this section, we outline the principal components within each subsystem, and the roles that they play in a storage ring.

1.2.1 Magnets

Magnetic fields are used to guide particles through the vacuum chamber, and to control the size (i.e. the cross-sectional area) of the beam as it moves around the ring. Dipole magnets produce a uniform vertical field that deflects the particles in the beam horizontally, so that the beam follows a defined path enclosed by the vacuum chamber. Quadrupole magnets (see, for example, figure 1.3) produce a magnetic field that varies linearly with distance from a path through the centre of the magnet: this type of field acts as a 'lens', providing the means to focus and control the size of the beam.

Other kinds of magnet are used for more refined control over the beam properties. For example, particles in the beam will have some (small) variation in energy, with the result that the focal length in a given quadrupole magnet will be different for different particles. This effect, termed *chromaticity* (see section 4.1), can be compensated by sextupole magnets, in which the field varies as the square of the distance from a line through the centre of the magnet.

Figure 1.3. The electron storage ring of the Advanced Light Source, Lawrence Berkeley National Laboratory. Dipole magnets (painted blue) steer the beam around the ring, while quadrupole magnets (red) provide focusing to control the beam size. Image courtesy of the Advanced Light Source, by SPat, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=25940137.

Most of the magnets in a storage ring will provide static fields; that is, the fields will be constant in time. The fields within these magnets must satisfy Maxwell's equations for magnetostatic fields[2]. Within the vacuum chamber, the current density (neglecting the beam itself) will be zero. This means that within the vacuum chamber, the magnetic field must have zero divergence and curl:

$$\nabla \cdot \vec{B} = 0, \quad \text{and} \quad \nabla \times \vec{B} = 0. \tag{1.1}$$

These equations are satisfied by fields with Cartesian components:

$$B_x = C_n r^n \sin(n\theta), \quad B_y = C_n r^n \cos(n\theta), \quad B_z = 0, \tag{1.2}$$

where C_n is a constant, r is the radial distance from the z-axis (along which the beam travels), and θ is the polar angle from the horizontal, transverse x-axis. Fields that can be expressed in the above form (1.2) are known as *multipole fields*. A field with $n = 0$ is a dipole field, $n = 1$ gives a quadrupole field, $n = 2$ a sextupole field, and so on. In terms of Cartesian co-ordinates, a general multipole field (with arbitrary n) is most easily expressed using complex notation:

$$B_y + iB_x = C_n r^n e^{in\theta} = C_n(x + iy)^n, \quad B_z = 0. \tag{1.3}$$

In a dipole, $B_x = 0$ and $B_y = C_0$. In a quadrupole, $B_x = C_1 y$ and $B_y = C_1 x$; and in a sextupole, $B_x = 2C_2 xy$ and $B_y = C_2\,(x^2 - y^2)$. In practice, field strengths in accelerator magnets can vary quite widely. The maximum achievable fields depend on the type of magnet, for example whether the conductors in an electromagnet are normal or superconducting, or on the type of material used in a magnet constructed from permanent magnetic material. Normal-conducting electromagnets are perhaps the most widely used type of magnet in electron storage rings; dipole magnets of this type typically achieve field strengths of order 1 T, and quadrupole magnets can achieve around 0.8 T at the aperture limit set by the pole tips.

Solutions to Maxwell's equations can be constructed by superposing multipole fields of different orders. The magnets in a storage ring will usually be constructed so that one particular multipole component is dominant; in the ideal case, the steering magnets (used to control the beam trajectory) will be 'pure' dipoles, and the focusing magnets (used to control the size of the beam) will be 'pure' quadrupoles. In practice, the field in any multipole magnet will contain components from all multipole orders, though one specific order will usually be very much larger than the others. In some cases, the steering magnets will be designed to control both the beam trajectory and the size of the beam: magnets in this case will contain both dipole and quadrupole components of significant strength.

It is also possible to rotate multipole fields about the z-axis. Fields given by the above expression (1.2) with real valued coefficients C_n are known as *normal* multipoles. A rotation of the field through an angle $\pi/2(n + 1)$ gives a *skew* multipole, which can be described by (1.2) with the relevant coefficient C_n having

---

[2] The electromagnetic theory needed to describe the fields in accelerator components can be found in many standard texts, for example [14, 15]. For a discussion more specialised for magnets in accelerators, see [16].

a pure imaginary value. A complex value for C_n (i.e. if C_n has real and imaginary parts) represents a superposition of a normal and skew multipole. A skew dipole will deflect a beam vertically rather than horizontally. A skew quadrupole will allow control over *coupling* in the beam (see section 2.6).

Magnets other than multipole magnets are used in storage rings for particular purposes. For example, strong solenoid fields are often used around the detector in a collider, although the role of the solenoid in this case is connected with the operation of the detector rather than the storage ring itself. *Insertion devices* consisting of sequences of (short) dipole magnets of alternating polarity are used in light sources to enhance the production of synchrotron radiation (see chapter 3).

Accelerator magnets can be electromagnets, or can be constructed from permanent magnetic materials. The advantage of electromagnets is that the field strength can be readily adjusted by controlling the flow of current through the coils of the magnet: although the field strengths in the magnets will be fixed during operation, commissioning and tuning a storage ring generally requires some adjustments to be made to the magnet strengths. The drawback of electromagnets is that high currents (several tens or hundreds of amperes) are usually needed to achieve the specified field strengths, so that providing the power, and cooling the magnets, can be an issue. Superconducting magnets are able to provide significantly higher field strengths than normal-conducting magnets; however, the additional cost and complexity associated with the cryogenics system needed to operate superconducting magnets means that normal-conducting magnets are usually preferred, where possible. Although magnets with adjustable field strength can be constructed using permanent magnetic materials, the mechanisms needed to provide the adjustment for such magnets have so far meant that (normal-conducting) electromagnets are usually the preferred option.

1.2.2 Radiofrequency cavities

Radiofrequency (RF) cavities [17] are used in electron storage rings to replace the energy that particles lose through synchrotron radiation. An RF cavity contains an oscillating electromagnetic field, with a dominant electric field component parallel to the trajectory of the beam as it passes through the cavity [18]. The energy gain of a particle of charge q as it passes through the cavity is $qV_0 \cos(\phi)$, where V_0 is the peak voltage across the cavity, and the particle arrives at a phase ϕ of the field oscillation. Although the energy lost in a single turn of the ring is usually only a small fraction of the energy of a particle, several cavities with peak voltages of the order of a megavolt or more are usually needed to maintain a beam in an electron storage ring operating with beam energy of a few GeV. The maximum voltage that can be achieved in a cavity is limited by the point at which electrons are stripped from the inner surface of the cavity, leading to field breakdown.

The oscillation frequency of the field in the RF cavities must be matched to the revolution frequency of particles in the ring: this is the basic principle behind operation of a synchrotron. A small change in the RF frequency will lead to a change in the beam energy, which will in turn change the revolution frequency: the dependence of the revolution frequency on the particle energy is characterised by the *phase slip factor* of the ring (see section 2.7) and allows stable operation of a storage ring even if the RF frequency is not set perfectly.

The phase of the RF field oscillation at which the voltage across the cavity exactly matches the energy lost by a particle to synchrotron radiation is known as the *synchronous phase*. If particles arrive at a phase slightly ahead of, or behind, the synchronous phase, then their motion around the ring can remain stable, through the mechanism of *phase stability* (see section 2.8). However, there are limits on the maximum distance from the synchronous phase for which stable motion can be maintained: particles arriving too far from the synchronous phase will be unable to adjust their energy while remaining within the storage ring, and will be lost from the beam. As a result, the beam in an electron storage ring will consist of bunches of particles, separated by gaps with length corresponding to the RF oscillation period. Typically, a bunch in a storage ring will be of order 10 ps (a few millimetres) in length, and the gaps between bunches will be of order 2 ns (about two-thirds of a metre), corresponding to an RF frequency of 500 MHz.

The shape of the field within each RF cavity must be carefully controlled to minimise any adverse impacts on beam behaviour. The oscillating electric field will induce a magnetic field in the cavity; in a simple cavity with a geometry that is approximately cylindrical, the electric field will be parallel to the axis of the cylinder, and the magnetic field lines will form circular loops centred on the axis. The magnetic field can deflect particles passing through the cavity, and since the strength of the magnetic field increases with distance from the axis, this can lead to focusing effects.

In addition to the 'fundamental' mode in which a longitudinal electric field accelerates particles passing through the cavity, the fields in an RF cavity can occur in *higher-order modes* in which the fields oscillate at higher frequencies, and form different patterns within the cavity. The fundamental mode is driven by electromagnetic fields generated by the RF power supply (such as a klystron, or a solid-state amplifier) and fed into the cavity through a waveguide and RF coupler. However, higher-order modes can be driven by the electromagnetic fields around the particles in the beam as they travel through the cavity. In some circumstances, the higher-order modes can reach amplitudes large enough that particle trajectories are deflected by an amount large enough for the beam to become unstable (see chapter 5). The cavity modes (the fundamental mode, as well as the higher-order modes) occur at discrete frequencies and with field patterns determined by the boundary conditions set by the shape of the cavity. If the frequencies of any of the higher-order modes coincide with frequencies present in the beam current spectrum, then resonances can occur in which the higher-order modes are driven to large amplitudes. An important step in the design of a storage ring is the optimisation of the design of the RF cavities, to avoid as far as possible any overlap between the cavity mode spectrum and the beam current spectrum. Nevertheless, the parameter regimes specified for modern electron storage rings can be extremely challenging, and feedback systems (as outlined below, in section 1.2.3) are often needed to maintain beam stability.

To minimise the dissipation of RF power in the walls of the cavity, RF cavities must be made from materials with a good electrical conductivity. The material must

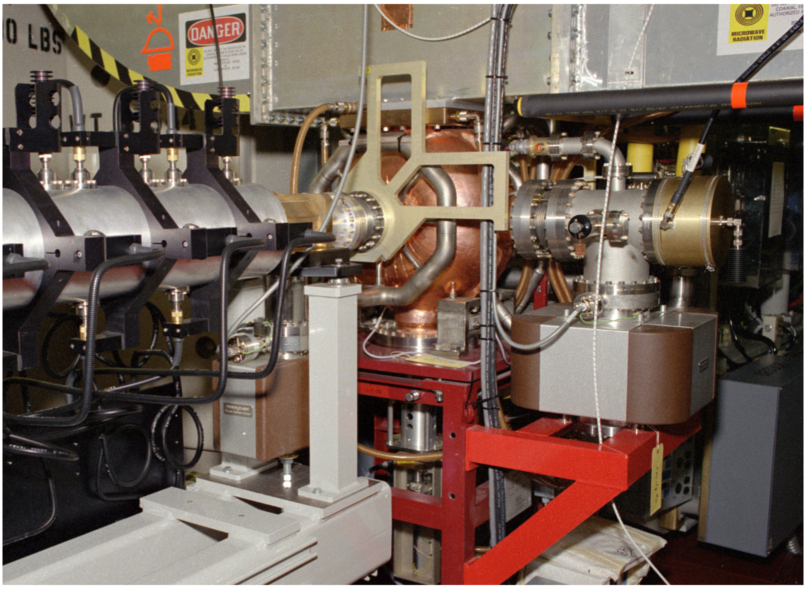

Figure 1.4. Radiofrequency cavity (copper chamber) in the storage ring of the Advanced Light Source at the Lawrence Berkeley National Laboratory. As it emits synchrotron light, the electron beam loses energy, which must be replaced if the beam is to continue circulating in the storage ring. The lost energy is put back into the beam by RF cavities, such as the copper structure in the centre of this photograph. Image courtesy of the Advanced Light Source, Lawrence Berkeley National Laboratory. Copyright 2010 The Regents of the University of California, through the Lawrence Berkeley National Laboratory.

also have appropriate mechanical and thermal properties to allow easy fabrication and stable operation. For cavities designed to operate at room temperature, copper is a common choice: see figure 1.4. However, at high field strengths, the amount of power dissipated in the walls of the cavity through induced electrical currents can be considerable, and it is usually necessary to provide cooling, by means of water flowing through pipes or channels fitted around the cavity. Some electron storage rings use superconducting RF cavities, which have the advantage of achieving relatively low power dissipation because of the extremely high conductivity of superconducting materials (although the DC resistivity of a superconductor is zero, oscillating currents are associated with some energy dissipation). The rate of decay of an oscillating electromagnetic field in a superconducting cavity can be slower by several orders of magnitude than the rate of decay in a comparable normal conducting cavity. The drawback with superconducting cavities is that it is necessary to operate below the critical temperature for the material from which the cavity is made. Although high-temperature superconducting materials do exist, their mechanical properties make them unsuitable for the fabrication of RF cavities. Most superconducting RF cavities are made from niobium[3], which has a critical temperature of 9.2 K. However, the operating temperature needs to be significantly

---

[3] Niobium has the highest critical temperature of any elemental superconductor.

lower than this (typically, around 4.5 K) so that the cavity remains superconducting in the presence of the strong magnetic fields that are inherent in the function of the cavity. The need for a cryogenic system, which brings additional cost and complexity to the RF system, often outweighs the benefits of superconducting technology: the choice between copper and niobium cavities is not a straightforward one.

1.2.3 Feedback systems

Storage rings usually contain many feedback systems with a wide variety of functions and operating over different timescales and parameter regimes. In general, the purpose of a feedback system is to maintain the stability of a measured parameter by making adjustments to components controlling that parameter. As an example, consider the beam trajectory (i.e. the orbit) in a storage ring: it is important that this is tightly controlled in a light source so that the beams of synchrotron radiation are directed accurately through the radiation beamlines to the experimental stations. In a collider, control over the orbit is needed to maintain good luminosity. The orbit is measured using a set of beam position monitors (BPMs) distributed around the ring, and can be adjusted using small dipole magnets that provide horizontal or vertical corrections to the beam trajectory. Over different timescales, small changes in orbit can occur from a number of causes, including (for example) electrical noise on the magnet power supplies, mechanical vibrations from pumps or other equipment, and temperature variations causing magnet supports to expand or contract. The first step to take to maintain orbit stability is to minimise the environmental effects affecting the orbit: magnet power supplies need to be of high quality (low noise), pumps should be mounted on supports that provide mechanical isolation from any vibration, magnet supports should be designed to minimise sensitivity to temperature variations, and the temperature in the tunnel housing the accelerator should be kept stable (in practice, often to a fraction of one degree Celsius). However, some residual beam motion is inevitable, so a feedback system is used to monitor the beam position around the ring, and to determine and apply the appropriate corrections to the strengths of the steering magnets to maintain a pre-defined orbit.

Calculating the appropriate correction to be applied by a feedback system from a set of measurements is not always a straightforward process. In the case of an orbit feedback system, for example, the way that the orbit responds to given changes in strength of the steering magnets depends on the strengths of all the other magnets in the storage ring. If the strengths of these magnets change slightly, then the orbit may respond in an unexpected way to changes in the steering magnets. If the feedback system is not carefully designed, then it may respond by attempting further corrections of steadily increasing amplitude, with the orbit becoming worse at each step until the beam is eventually lost from the storage ring. The same consequences can result if the signals from the BPMs are noisy, so that the correction applied by the feedback system is based on inaccurate data. Feedback systems are often needed on systems that are inherently unstable, and can be very sensitive to small environmental changes; an effective and reliable feedback system must be carefully designed, based on a good understanding of the system on which it will work. A number of standard algorithms have been developed, which can be applied in different situations; see, for example, [19].

Beyond achieving the required stability, a further challenge in many cases is the timescale on which a feedback system may need to operate. In the case of an orbit feedback system, orbit changes in response to ground motion or temperature variations may be on the timescale of minutes, hours or days: such timescales do not pose significant problems for modern feedback systems. However, electrical noise may cause oscillations in beam position with frequencies of the order of hundreds of hertz, or several kilohertz. Collecting data from a large number (maybe dozens) of BPMs, then calculating and applying a correction within a few milliseconds can be very challenging. In a storage ring, beam instabilities driven by wake fields (electromagnetic fields within the beam pipe generated by the beam itself) can occur on the timescale of a few microseconds: it is possible, using specialised fast feedback systems, to suppress such instabilities, allowing storage rings to operate in parameter regimes that would otherwise be inaccessible. Fast (bunch-by-bunch) feedback systems for suppressing beam instabilities are discussed further in section 5.6.

Other feedback systems in a storage ring may serve to maintain beam optics parameters (such as the betatron tunes), the beam energy and intensity; or they may operate within technical subsystems, for example to maintain the stability of the current from a magnet power supply, or to improve the stability of the frequency and amplitude of the fields in an RF cavity.

1.2.4 Vacuum systems

Even at very low pressure, gas in the beam pipe in an electron storage ring can have a number of effects that limit machine performance [20]. Particles in the beam may collide with gas molecules, leading to a loss of beam current. Gas molecules may become ionised by collisions, and the resulting positive ions can be ‘trapped’ in the negative electrical potential around a beam of electrons; interactions between the electrons and the ions can then cause the beam to become unstable. In positron storage rings, electrons from the ionisation of gas molecules can collide with the walls of the beam pipe, releasing additional electrons. Under some circumstances, the density of electrons (the *electron cloud*) can build up to the point where the beam becomes unstable. In colliders, scattering of particles in the beam from gas molecules in the vicinity of the interaction region can lead to backgrounds in the detector.

Gas pressures inside the vacuum chamber of a particle accelerator are commonly measured in torr (Torr): standard atmospheric pressure is 760 Torr (or 1.01×10^5 Pa in SI units). The pressure in the vacuum chamber of an electron storage ring is typically specified to be of order 1 nTorr, or roughly 10^{-12} times atmospheric pressure. At this pressure, the beam lifetime (the time taken for the beam current to fall by a given factor) is likely to be dominated not by gas scattering, but by other effects, in particular Touschek scattering (section 5.1). In practice, the pressure in a storage ring can vary considerably depending on the location around the ring and the operational conditions. When an electron beam is present, synchrotron radiation falling on the walls of the vacuum chamber can release gas molecules through photodesorption, leading to a significant increase in the local pressure. For this reason, the vacuum chamber is designed to minimise the amount of synchrotron radiation falling on the inner wall of the chamber. In a light source, the intense radiation beams from insertion devices are generally directed to radiation beamlines; but there can still be considerable amounts of radiation, for example from dipoles, remaining within the beam pipe. The fact that radiation can be reflected multiple times from the walls of the beam pipe can make it difficult to predict exactly where photodesorption will occur. A common technique to remove radiation from the main chamber is to shape the vacuum chamber so that it has a slot running along one side that opens into an 'antechamber'. The slot allows radiation to pass out of the main chamber, but since it has a lower conductance than the main chamber for the flow of gas, it is possible to maintain a lower pressure in the main chamber than in the antechamber.

The specified pressure within the vacuum chamber of an electron storage ring is usually only achieved after a long period (maybe some weeks or months) of pumping and conditioning. Different types of pumps are needed to cover the range of pressures and species in the residual gas. For example, turbo-molecular pumps can operate from atmospheric pressure to below 0.01 nTorr, and can pump any kind of gas molecule. However, the pumping is localised so that the required pressure may only be achieved in the near vicinity of the pump; this limitation can be overcome by installing large numbers of pumps, but their high cost can make this prohibitive. Titanium sublimation pumps can achieve high pumping speed at relatively low cost, and can be implemented to provide 'distributed' pumping, i.e. operating to pump long sections of the beam pipe. However, titanium sublimation pumps cannot operate above pressures of around 10^{-4} Torr, and will not pump molecules of inert gases such as argon and methane. Similarly, different types of instrumentation (pressure gauges) are needed to ensure accurate readings of pressure in different regimes and environments.

Materials used in the construction of an accelerator vacuum chamber must have certain properties. For example, the surface of any material will absorb gas molecules when exposed to gas at high (atmospheric) pressure, then release the gas molecules when the pressure is reduced (in a vacuum vessel) in a process known as *outgassing*. The outgassing rates from materials used in an accelerator beam pipe should be low enough to allow the specified pressure to be maintained without excessive numbers of pumps. The chamber walls should have good electrical conductivity to minimise the electromagnetic fields generated by particles in the beam as they travel through the chamber (in particular, to minimise resistive-wall wake fields, see section 5.3). The material should have good mechanical properties to allow the different parts of the chamber to be fabricated easily. In positron storage rings, it may be necessary to choose a material with a low probability of releasing electrons, either from photoemission or from the incidence of primary electrons or ions, to avoid the build-up of an electron cloud (section 6.2.3). Finally, the chamber should be able to withstand heating by a couple of hundred degrees, which may be required to condition the chamber or to activate certain types of coating (in particular, non-evaporable getter, or NEG coatings) used to provide additional pumping. Aluminium is a common choice of material for the vacuum chambers of electron storage rings. Some of the properties of aluminium are not ideal; for example, it can lose mechanical strength when heated to the temperatures need to activate NEG coatings. Nevertheless, aluminium can meet many of the requirements for the vacuum chamber in an electron storage ring, at a reasonable cost.

The geometry of the vacuum chamber is also important for a number of different reasons. First, the aperture should be large enough to accommodate the beam, allowing for particles performing transverse oscillations of large amplitude, or with significant energy errors. Steering errors and processes occurring during injection can also lead to large deviations of the beam trajectory from the ideal orbit. From point of view of the vacuum, a large aperture helps to achieve a good gas conductance through the chamber, which can improve the pumping rate achieved with a fixed number of pumps. Second, the chamber should be designed to allow as much synchrotron radiation as possible to leave the main chamber without striking the walls; the impact of synchrotron radiation on the walls of a vacuum chamber can lead to photodesorption, or to local heating. The operation of some beam diagnostic devices can be affected by synchrotron radiation, and may need to be shielded in some way. In a light source, extraction ports are necessary to allow the radiation to exit the ring towards experimental areas, but ports may also be necessary to extract the synchrotron radiation from colliders. Third, the chamber needs to allow some flexibility to allow for changes in position or length in different sections, in response to the motion of components attached to the chamber or to changes in temperature. Some components fixed to the chamber (for example, beam position monitors) also need to be isolated as far as possible from vibrations or other mechanical motion that may be transmitted along the chamber. Flexibility and a certain amount of mechanical isolation can be provided by ‘bellows’ located at appropriate points around the ring. Finally, the geometry of the inside of the chamber needs to be as smooth as possible, avoiding gaps or sudden changes between different apertures or cross-sections. This is because abrupt transitions can trap electromagnetic fields generated by particles in the beam, leading either to localised heating as the energy of the fields is dissipated in the walls of the chamber, or to wake fields that act back on the beam potentially causing beam instabilities.

The requirements on the chamber geometry are often in conflict; for example, it can be difficult to provide ports for vacuum pumps or for extracting synchrotron radiation while maintaining the ‘smoothness’ necessary to minimise wake fields. Bellows also tend to be associated with cavities that can lead to strong wake fields. The solution is often to provide some ‘RF shielding’ using strips of metal arranged to present a smooth, unbroken screen to electromagnetic fields oscillating at microwave (and lower) frequencies, while allowing the relevant components (vacuum pumps, bellows, etc) to perform their appropriate function. Optimising the design of a vacuum chamber can be a complex, iterative task involving modelling the vacuum, mechanical, and electromagnetic properties, and often requiring some compromises to be made to achieve a practical solution. For further information on a range of aspects of vacuum systems for accelerators, see [21].

1.2.5 Diagnostics

Beam diagnostics [22] are critical for the commissioning and effective operation of a storage ring. The behaviour of the beam in a storage ring is sensitive to a large number of environmental factors, and can be affected in particular by field and alignment errors on the magnets, frequency and amplitude errors in the RF cavities, and timing errors on the injection system. A variety of collective effects (some of which are discussed further in chapter 5) can impact the beam parameters and behaviour in ways that can be difficult to predict, and that depend on the beam intensity, the magnetic lattice and the design of the vacuum chamber. No matter how carefully a storage ring is designed and assembled, injecting and storing a beam and achieving the specified beam quality and stability depends on detailed, accurate, and complete information on the properties and behaviour of the beam. Any storage ring must include a wide range of diagnostics devices to provide information on the charge in individual bunches, the position of the beam at different points around the storage ring, and the transverse beam size and bunch length. The design of the diagnostics system needs to take into account the most critical points at which particular measurements need to be made, the parameter range over which the diagnostics should provide accurate information, the rate at which measurements need to be taken, and the impact that the diagnostics devices themselves will have on the beam. Installing additional diagnostics after construction and commissioning can be difficult, so it is important to include the diagnostics system as an integral part of the design process for a storage ring.

Many diagnostic devices work on the principle of directly detecting the electromagnetic fields around a bunch of charged particles; but the detection can be done in many different ways. For example, to measure the total amount of charge in a bunch, a metal block can be placed in the path of the bunch. When the charge hits the block, the electric potential of the block changes by $V = Q/C$ where Q is the charge on the bunch and C is the capacitance of the block. Unfortunately, this technique destroys the bunch in the process of making the measurement, so is usually not suitable for a storage ring. It can, however, be a useful method in a single-pass accelerator, such as a linac, or in the transfer lines used for injecting a beam into a storage ring. In practice, the block needs to be carefully shaped so that as many electrons as possible are trapped within the block, rather than simply being scattered around it: this generally leads to a *Faraday cup* design, in which a bunch passes through an aperture into a cavity within the block. Furthermore, the voltage signal can be very small and last for a very short time, so the electronics used to detect the signal need to be very sensitive, with low noise, and high bandwidth (i.e. able to process signals that change rapidly with time). Within a storage ring, charge measurements are more routinely made using either a *DC current transformer* (DCCT) to measure the voltage induced in a coil of wire by the magnetic field around the beam, or a *wall-current monitor* placed across an insulating gap in the vacuum chamber to detect the electric currents induced in the chamber by the beam. These devices are non-destructive, in the sense that they allow the beam to continue to circulate during and following the measurement.

After measurement of the bunch charge, the trajectory of the beam is usually the next critical measurement needed during commissioning and tuning of a storage ring. In general, the trajectory is measured using a series of *beam position monitors* (BPMs), each of which consists of a set of electrodes that protrude into the vacuum chamber but are insulated from it. A simple BPM pickup may consist of one electrode protruding into the top of the chamber, and another protruding from the bottom. As a bunch passes between the electrodes, it will induce a voltage on each electrode, with the size of the voltage depending on the proximity of the beam to the electrode: the closer the beam to the electrode, the larger the voltage. Thus, the difference between the voltage signals is related to the vertical position of the beam in the chamber. The sum of the voltages induced on the electrodes can give an indication of the total charge in the bunch, although specialised devices (such as a DCCT or a wall-current monitor) can usually provide more accurate and reliable charge measurements. Of course, the voltages induced on the electrodes also depend on the horizontal position of the bunch within the vacuum chamber; a common arrangement in a storage ring is therefore to use four electrodes to measure the horizontal and vertical position of the beam at the same time. The shape of the electrodes must be optimised to maximise the sensitivity of the BPM, and to minimise wake fields (see section 5.3): in electron storage rings ‘button’ electrodes are commonly used, in which the electrode terminates in a flat plate that is flush with the surrounding surface of the vacuum chamber. Electrodes consisting of ‘strip lines’ can also be used; the optimum geometry depends on the parameter regime. Very precise beam position measurements (with resolution less than 1 μm) can be made from the oscillating fields induced when a bunch passes through a cavity. However, the wake fields resulting from cavity BPMs usually preclude their use in storage rings.

Interpreting the signals from a beam position monitor is not completely straightforward. Typically, the BPMs in a modern electron storage ring are required to provide measurements of the position of the beam centroid with a resolution of order 10 μm (which may be comparable to the transverse size of the beam). The variation of the voltage on each electrode with beam position will be nonlinear, and will depend on the shape of the surrounding vacuum chamber. The voltage will also depend on the bunch charge. The signal returned by the electronics used to process the voltages on the electrodes will be subject to some level of noise, and will also depend on temperature. Determining the beam position from the signals returned by a BPM requires careful calibration, and will be subject to random and systematic errors. Random errors (from noise in the electronics) can be reduced by making repeated measurements, although it may be difficult to distinguish between real changes in beam position and random fluctuations in the BPM signals. Systematic errors (for example, from effects associated with the local geometry of the vacuum chamber, or temperature variations within the electronics) can be more difficult to characterise. Readings from a BPM are often best regarded as providing relative rather than absolute measurements of beam position.

Measurements of the transverse beam size can be obtained by placing a screen in the path of the beam. The screen may be made from a material that fluoresces when hit by a bunch of charged particles, such as yttrium aluminium garnet (YAG); or, it may be a simple metal foil from which *transition radiation* may be observed from a charged bunch crossing the boundary from the vacuum to the interior of the metal. In either case, the light produced can be observed with a camera. A screen using transition radiation has the advantage that the radiation pulse length is equal to the length of the electron bunch producing the radiation (typically of order several picoseconds), whereas the light from a YAG screen can persist for a considerable time (maybe hundreds of nanoseconds) after the bunch has passed. Thus, a screen using transition radiation can, with a suitable fast camera, provide information on the size of an individual bunch, whereas a YAG screen may, depending on the bunch spacing, provide a signal representing an integration over many bunches. On the other hand, a YAG screen is usually more sensitive than a screen using transition radiation, and may therefore be more appropriate at relatively low beam currents and beam energies. Both techniques are, of course, destructive, and therefore not suitable for a storage ring in regular operation. Instead, measurements of beam size may be made from observation of the synchrotron radiation produced from charged particles moving through magnetic fields. Observing the radiation is straightforward, but again, obtaining the desired information (in this case, of the size of the beam producing the radiation) requires a detailed understanding of the physical principles and practical implementation of the various components of the system, and careful analysis of the signals obtained.

Measurements of bunch length are also important in optimising the performance of an electron storage ring, and can (like measurements of the transverse beam size) be performed using synchrotron radiation. A common instrument used for bunch length measurements is a *streak camera*. A streak camera consists of a plate that emits electrons when struck by a pulse of radiation. The electrons are accelerated (by a static electric field) towards a detector, but between the emitter and the detector pass through a transverse electric field that is oscillating at high frequency. The variation of the transverse field means that the electrons receive a transverse deflection with a magnitude that depends on the time at which they pass through the field. Early electrons then arrive at the detector at a different transverse position from late electrons. The transverse spread of the electrons on the detector provides a measurement of the length of the electron bunch in the storage ring. The importance of synchrotron radiation for electron beam diagnostics is such that electron storage rings often have synchrotron radiation beamlines dedicated to diagnostics for the beam in the storage ring.

A wide range of diagnostics, beyond those mentioned here, have been developed to provide more complete and detailed information on beam properties in storage rings (and other types of accelerator). New diagnostic techniques are still being developed, either to widen the scope of information available, or to provide better precision or accuracy in various parameter regimes than can be achieved using existing techniques, or to minimise the effects of the measurement on the beam. In some cases, valuable information can be obtained by using a range of diagnostics while manipulating the beam in various ways. For example, the position of a beam in a quadrupole magnet can be found by changing the strength of the magnet while observing the beam orbit using the BPMs; this method, sometimes called *beam-based alignment*, can be used to steer the beam through the centres of the quadrupole magnets, which is often an important step in tuning an electron storage ring for optimum performance. Any tuning procedure requires not only suitable diagnostics, but also a good understanding of those diagnostics: this will involve not just an appreciation of the physical principles on which the measurement is based, but also an awareness of the appropriate parameter regimes, the limitations of the diagnostics in terms of accuracy and the sources and sizes of random and systematic errors, the ways in which any measurements from the diagnostics may be verified, and the potential impact on the beam. Often, the best approach is to regard the beam and the diagnostics as a single physical system, and to view the signals from the diagnostics in that context rather than taking them at face value.

1.2.6 Control systems

The purpose of the control system [23] in an accelerator is to allow the operators (through a convenient interface) to set the parameters for the various components and subsystems, to read signals from beam diagnostics and other components, and to display and record these signals as required. A modern storage ring can include thousands of separate components, and the inevitable complexity of a system on such a scale means that careful consideration needs to be given to the design, implementation, and maintenance of the control system. Typically, the architecture of a control system consists of three layers. In the bottom layer are computers that communicate directly with a small number of components (such as magnet power supplies, beam position monitors, vacuum pumps etc). The top layer consists of computers that display information for the operators, and allow the operators to send commands to the various components and subsystems. The middle layer consists of a set of powerful data servers that channel information between the top and bottom layers, and record data from different sources. The required data transfer rates may be very large, for example in the case of high-resolution camera images collected at high frame rates; the data files that need to be recorded may similarly be of a considerable size. The computer hardware and software need to be carefully selected to match the performance requirements.

In practice, a modular approach is often the most practical, with each of the different layers structured to reflect different sets of components or subsystems in the accelerator. Different software systems may be used in different parts of the control system; there are then potential issues, for example regarding compatibility of data formats. However, tools are generally available that allow different platforms to communicate with each other efficiently.

1.2.7 Injection systems

The beam in an electron storage ring follows a closed orbit, in which each bunch of particles returns to the same position after each turn of the ring. An injection system must control the trajectory of a bunch arriving at the storage ring, so that the new bunch follows the closed orbit once it is in the ring [24]. The simplest system conceptually is 'single-turn' injection, in which a dipole magnet is located at the point where the trajectory of an incoming bunch crosses the closed orbit. Under normal circumstances, this dipole magnet is turned off, so that it does not deflect bunches that are already on the closed orbit. However, during injection the dipole is turned on, so that an incoming bunch, initially travelling at some angle to the closed orbit, is deflected so that it follows the closed orbit on leaving the dipole. It is then necessary to switch off the dipole before another bunch already on the closed orbit arrives, otherwise this bunch will be deflected out of the storage ring. This scheme requires a dipole magnet that can be turned on and off very quickly: a magnet of this type is often called a *kicker magnet*. Even with a fast kicker magnet, single-turn injection usually means leaving a gap in the train of bunches long enough for the field in the kicker to be turned on and off. A gap in the bunch train in an electron storage ring may be needed in any case, to avoid ion trapping (see section 5.2). However, another problem with single-turn injection is that the source must be capable of delivering in a single pulse a train of bunches each with the full charge specified for each bunch in the storage ring. This can place severe demands on the particle source.

An alternative injection scheme involves filling the storage ring over several turns, merging additional charge with the charge already in the storage ring. This can be accomplished using a *septum magnet*. A septum magnet consists of a dipole magnet with two apertures separated by a narrow plate, or 'blade'. By passing an electric current through the blade, it is possible to arrange for the dipole field to appear only in one of the apertures. The aperture on the other side of the blade, in which there is zero field, is the aperture through which the beam passes when on the closed orbit. The injection process follows a number of steps. First, a set of dipoles is used to move the closed orbit so that it comes very close to the septum blade. Then, a bunch from the source is directed through a transfer line towards the storage ring, where it arrives at the septum on the side of the blade where there is a strong dipole field. The field deflects the incoming bunch so that it is travelling parallel to the closed orbit, but at some displacement from it. Since the injected bunch is not on the closed orbit, it will oscillate around the closed orbit as it travels around the ring (see section 2.2). The magnets in the storage ring are set so that once the bunch arrives back at the septum after completing a turn of the storage ring, it is at a phase of its orbit oscillation such that it passes through the septum through the field-free aperture. The final step in the process is then to turn off the dipole magnets that were used to distort the closed orbit, so that (usually after several turns) the beam follows its original closed orbit. Meanwhile, the oscillations of the newly injected bunch around the closed orbit will steadily decrease in amplitude as a result of processes such as synchrotron radiation damping (see section 3.2.2). Thus, after several turns, the charge in the ring has increased and the beam is once again on the original closed orbit. This process can be repeated as often as necessary to reach the full beam current needed in the storage ring. Although dipoles are still needed that can be turned on and off quickly, the timescale is now on the order of several revolution periods, rather than a fraction of a revolution period. The demands on the particle source are also greatly relaxed, compared to single-turn injection. A drawback with multi-turn injection is that there is a risk of significant beam losses on the septum, or other components: the septum blade in particular can easily be damaged if the heat load from particles lost from the injected or stored beams is too high. The injection system needs to be carefully designed and the entire storage ring needs to be finely tuned to minimise beam losses during injection.

There are a number of variations on the scheme outlined above for multi-turn injection, which may have advantages in some cases, usually by reducing injection losses. For example, it is possible to take advantage of the variation in particle trajectory with the energy of the particle (an effect known as *dispersion*—see section 2.5) by injecting bunches with some difference in energy: a 'dispersion bump' can then be used in addition to (or even instead of) an 'orbit bump' in the region of the septum.

Particles in the beam in an electron storage ring can collide with residual gas particles in the vacuum chamber or with other particles in the beam (see section 5.1): these collisions or scattering events can result in the loss of particles from the beam, leading to a decay in the beam current. If no additional particles are injected into the ring to replace the particles lost by scattering, the beam current can fall by half over the course of a few hours; the exact timescale can vary quite widely, depending on the beam parameters and the pressure of the residual gas in the vacuum chamber. Traditionally, electron storage rings have been operated in such a way as to refill the ring a couple of times each day. This reduces the impact on users from variations in beam orbit and beam size associated with the injection process; however, it has the disadvantage that there are large variations in the brightness of the radiation beams (in a synchrotron light source) or the luminosity (in a collider) over time. Also, the heat load on different components in the accelerator can change substantially in response to changes in current, making it difficult to maintain temperature stability. To address these problems, light sources have developed ways of operating with 'top-up' injection, in which small amounts of charge are injected at short, regular intervals of perhaps no more than a few minutes. If the effects on the beam from the injection process can be kept small enough, then top-up injection provides better overall beam stability for users than the traditional mode of refilling at intervals of several hours.

1.2.8 Personnel protection

Accelerators present a number of hazards, including components or substances at very high or very low temperature, high voltages, toxic chemicals, confined spaces (where there may in some cases be a risk of oxygen depletion from the accidental release of liquified gases), laser beams, and radiation. Systems need to be in place to minimise the risks from all hazards. The first step should be to eliminate each hazard if at all possible, for example by designing systems to avoid the possibility of components reaching temperatures high enough to present a risk of burns. Inevitably, however, there will be instances where hazards cannot be avoided, and precautions must then be taken to reduce the risks. Possible measures may include: enclosing or shielding the hazard; limiting access to only those personnel with a need to work in the area affected, and who have been properly trained and authorised; ensuring that clear warning notices are displayed; and providing monitoring systems and alarms The use of personal protective equipment should usually be the last line of defence.

Although all hazards need to be properly considered and addressed, systems associated with protection from radiation are often the most visible in an accelerator facility [25]. Radiation may be generated from any high-energy particles that are incident on matter; the type of radiation and its properties will depend on the type of particle involved, the particle energy, and the materials struck by the particles. Synchrotron radiation from relativistic particles undergoing acceleration (see chapter 3) is also a hazard. In electron storage rings, the highest levels of radiation are produced while the accelerator is operating. The ring will be enclosed either in an underground tunnel, or within a 'surface tunnel' constructed from large concrete blocks. Access to the tunnel will be through a limited number of entrances, usually with a labyrinth so there is no direct line of sight to the accelerator from outside the tunnel. Each entrance will have an interlocked gate or door, so that the accelerator is automatically and immediately switched off if anyone tries to enter the tunnel during operation. Before the accelerator is turned on, the tunnel will be searched according to a well-defined procedure to ensure that there is no-one in the tunnel before the entrance interlocks are activated. In the case of an electron storage ring, after the accelerator is turned off the radiation from any materials that have been activated will normally decay quite rapidly; nevertheless, measures will be taken to ensure that radiation in the area around the accelerator is below specified limits before allowing general access.

Personnel protection systems need to be completely reliable. For that reason, systems (including alarms and interlocks) are designed with significant redundancy in different components, and so that any failure will lead to a safe condition; for example, an interlock on the entrance to the accelerator tunnel should actively report when the entrance gate is closed, so that if the interlock fails the system interprets the lack of a signal as meaning that the gate is open. Although the personnel protection system may communicate with the accelerator control system, the need for high reliability often means that critical safety systems are controlled directly by dedicated hardware, and are not under computer (software) control.

1.3 Some examples of electron storage rings

Although electron storage rings in third-generation synchrotron light sources have a number of common features, they can also vary quite widely in their structure and parameters. It is difficult to identify any one ring as 'typical'; however, the parameters for some representative examples are given in table 1.1. Even in the cases given, the parameters can be changed depending on the needs of the users. For example, some light source users need radiation in short pulses (tens of picoseconds) with relatively long gaps (hundreds of nanoseconds) in between: a time structure of this kind can be achieved by injecting charge into the ring in just one or two bunches. The synchrotron radiation produced by the beam in this case will have low average

Table 1.1. Parameters of some storage rings for third-generation synchrotron light sources. It is important to note that most light sources have significant flexibility in their operating parameters, so as to be able to vary the characteristics of the synchrotron radiation that they produce for users. Parameters marked with an asterisk (*) in the table can be varied, in some cases over a very wide range: the values given should be taken as being indicative.

	ALS	SPring-8	AS	DLS	MAX IV
Start of operation	1993	1999	2007	2007	2017
Circumference (m)	197	1436	216	562	528
Lattice type	TBA	DBA	DBA	DBA	7BA
Number of arc cells	12	44	14	24	20
Beam energy* (GeV)	1.9	8.0	3.0	3.0	3.0
Beam current* (mA)	500	100	200	300	500
Natural emittance* (nm)	2.0	2.4	7.1	3.2	0.3
RF frequency (MHz)	500	509	500	500	100
RF voltage* (MV)	1.1	16	3.2	2.5	1.5
Bunch length* (ps)	21	13	22	20	33
Beam lifetime* (h)	8	200	16	10	10

intensity, because the charge that can be contained in a single bunch is limited by collective effects. Other users may be less sensitive to the time structure of the radiation, but need a high average intensity: for these users, as much charge as possible will be injected into the ring, usually in several hundred bunches separated by gaps of a few nanoseconds. In addition to changes in filling pattern, there can be variations in the beam optics that can affect the beam size, which is determined (in part) by the emittance of the beam[4]. Storage rings used in light sources also often have upgrade programmes intended to extend the life of the facility by taking advantage of developments in technology that allow improved performance. Individual upgrade projects may range in scale from the installation of a small number of new components, to the replacement of essentially the entire storage ring.

The Advanced Light Source (ALS [26], see figure 1.3) in Berkeley, California, USA is one of the earliest third-generation light sources, i.e. one of the first light sources to be developed with a low-emittance lattice for the production of high-brightness synchrotron radiation. The low emittance is achieved by the use of a triple-bend achromat (TBA) structure for the lattice (see section 3.3): each arc cell includes three dipole magnets. Although other light sources have been built using TBA arc cells, the double-bend achromat (DBA, with just two dipoles in each arc cell) has been a more popular choice. Recently, however, there has been interest in lattice styles using a larger number of dipoles in each arc cell, to achieve very low emittances: an example is MAX IV in Lund, Sweden [27, 28].

---

[4] Emittance is discussed in more detail in section 2.3; but briefly, emittance is the product of the beam size and divergence (the rate of change of the beam size). Emittance is a more useful measure of beam quality than beam size, since the emittance of a beam is constant as the beam moves around a storage ring, whereas the beam size varies depending on the local focusing provided by quadrupole magnets.

SPring-8 (Super Photon Ring—8 GeV [29]) in Hyogo Prefecture, Japan, is one of five light sources around the world operating with electron beam energies above 5 GeV. Using a higher energy for the electron beam makes it easier to generate high-brightness synchrotron radiation at shorter wavelengths (i.e. in the hard x-ray region of the electromagnetic spectrum), which is needed for some applications. There are also some benefits of higher energy in terms of the accelerator physics: for example, the beam becomes less sensitive to certain collective effects, and tends to have a longer lifetime. The drawback of a higher beam energy is that to achieve a low emittance, which is necessary for generating high-brightness synchrotron radiation, a large circumference, accommodating a large number of arc cells, is needed (see section 3.3). A number of the subsystems, in particular the RF system, also need to be larger in rings operating at higher energy. The result is that both the construction cost and the running costs of the facility increase with increasing energy. Many third-generation light sources are designed to operate with a beam energy around 3 GeV, which provides a good compromise between the capability to produce high-brightness, short-wavelength synchrotron radiation and the cost of the facility.

The Australian Synchrotron (AS [30]) in Melbourne, Australia, and the Diamond Light Source (DLS [31], figure 1.5) in Didcot, UK, are examples of synchrotron light sources constructed in the first decade of the twenty-first century. These, and many other facilities constructed around the same period, are based on double-bend achromat lattices, with 3 GeV beam energy. There can still be significant differences between such machines, however. DLS, for example, has more than twice the

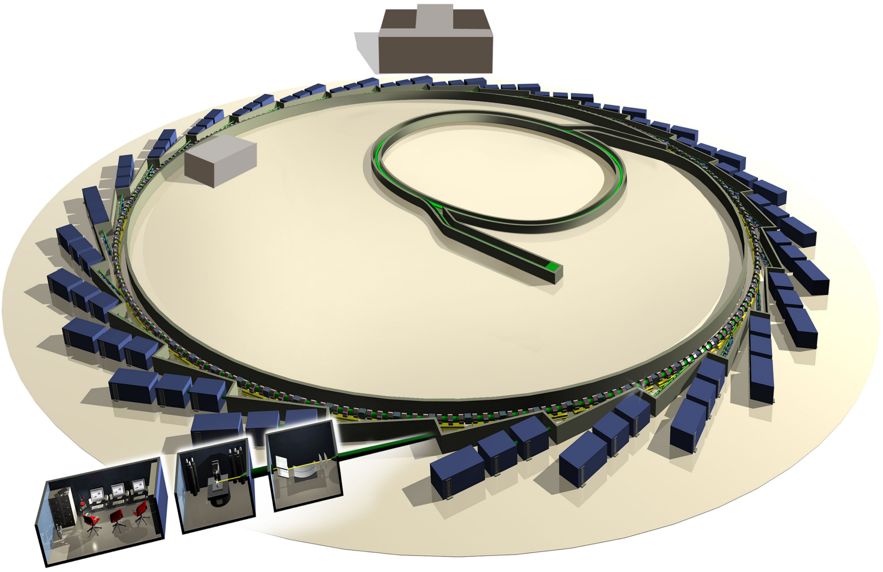

Figure 1.5. Schematic of the DLS. The electron source and booster synchrotron (for raising the particle energy before injection into the storage ring) are located within the main storage ring. Synchrotron radiation beamlines, constructed tangential to the main storage ring, carry synchrotron radiation from dipoles and insertion devices to the experimental areas. Image courtesy of Diamond Light Source.

Table 1.2. Parameters of some storage rings for electron–positron colliders. As with light sources, storage rings for colliders usually have significant flexibility in their operating parameters. Parameters marked with an asterisk (*) in the table can be varied, in some cases over a very wide range: the values given should be taken as being indicative.

	CESR	LEP	DAΦNE	SuperKEKB	
				LER (e^+)	HER(e^-)
Start of operation	1979	1989	1999	2018	2018
Circumference (m)	768	26 659	98	3016	3016
Beam energy* (GeV)	5	104	0.51	4	7
Beam current* (mA)	365	5.2	5000	3600	2600
Natural emittance* (nm)	300	21	1000	3.2	2.4
RF frequency (MHz)	500	352	368	509	509
RF voltage* (MV)	5	3500	0.254	8.4	6.7
Bunch length* (mm)	20	10	30	6	5
Beam lifetime* (h)	5	5	2	0.2	0.2

circumference of AS. This not only allows the ring to accommodate a larger number of arc cells, enabling a lower emittance (and higher brightness for the synchrotron radiation beams), but also provides for a much larger number of radiation beamlines, thus accommodating a larger number of user experiments.

Storage rings in colliders are in many aspects more varied than the storage rings used in light sources. Table 1.2 shows some of the parameters for a number of electron–positron colliders. The scientific goals for a collider tend to be more specific than for light sources, which serve a more diverse user community. Broadly speaking, colliders aim to discover new phenomena either by achieving higher centre-of-mass (collision) energies than previous colliders, or by achieving higher luminosities. The collision energy determines the processes that can occur during a collision between two particles; for a process to occur, the collision energy will normally have to exceed some threshold. The luminosity determines the rate at which the various processes occur in a collider: the higher the luminosity, the more events of a given type will be observed, allowing more precise measurements of the parameters associated with that event (for example, the mass of a particle produced in a collision).

Large Electron–Positron Collider (LEP), at CERN, Geneva, Switzerland [32] operated from 1989 to 2000, reaching the highest collision energy ever achieved in an e^+e^- collider. The final beam energy of 104 GeV was only possible because of the large circumference of the machine. The energy lost from the beam through synchrotron radiation has to be replaced by the RF cavities; the larger the circumference, the lower the synchrotron radiation power, and the lower the cost of the RF system. Hadron colliders can achieve much higher collision energies because the synchrotron radiation losses depend on the mass of the radiating particle: the heavier the particle, the smaller the radiation power. The drawback with hadron colliders is that the collisions are far more complex than in lepton colliders, because of the internal structure of hadrons, and the ability of hadrons to interact via the strong nuclear force. It is unlikely that any future electron–positron colliders based on storage rings will exceed the centre-of-mass energy achieved in LEP. There have been extensive studies of linear colliders that avoid the problem of synchrotron radiation by accelerating and colliding electrons and positrons in a machine in which the beam essentially follows a straight path; however, there are significant technical difficulties that would need to be overcome in such a machine.

Other storage ring colliders, including the Cornell Electron Storage Ring (CESR, at Cornell University, Ithaca, New York, USA [33, 34]), DAΦNE (the Double-Annulus Phi factory for Nice Experiments, at LNF, Frascati, Italy [35]) and SuperKEKB (Super KEK B-factory, at KEK in Tsukuba, Ibaraki Prefecture, Japan [36]) aim for maximum luminosity, rather than the highest possible energy. Each of these machines is optimised for a particular process, or a small number of processes, to be studied in detail. For example, SuperKEKB is intended to study the physics of B mesons: the centre-of-mass beam energy is set at the peak of the cross-section (i.e. the probability) for production of ϒ(4S) mesons, which decay to B mesons. Of particular interest is the asymmetry between the B meson and its antiparticle. The asymmetry in the electron and positron beam energies in SuperKEKB (and in the earlier colliders, PEP-II and KEKB) allows precise measurements of the lifetimes of the particle–antiparticle pairs of B mesons produced in the collisions.

The very wide variation between parameters that can be seen from the examples in table 1.2 results from the fact that the different colliders are optimised for different physics goals. Each different parameter regime tends to involve different challenges for the design, construction, and operation of the storage ring.

References

[1] Sessler A and Wilson E 2007 *Engines of Discovery: A Century of Particle Accelerators* (Singapore: World Scientific)

[2] Rubensson J-E 2016 *Synchrotron Radiation: An Everyday Application of Special Relativity* (San Rafael, California, USA: Morgan & Claypool Publishers) (Bristol, UK: IOP Concise Physics, IOP Publishing)

[3] Rowe E M and Mills F E 1973 Tantalus I: a dedicated storage ring synchrotron radiation source *Part. Accel.* **4** 211–27

[4] Winick H and Doniach S (ed) 1980 *Synchrotron Radiation Research* (New York, NY, USA: Plenum)

[5] Clarke J A 2004 *The Science and Technology of Undulators and Wigglers* (Oxford, UK: Oxford University Press)

[6] Freund H P and Antonsen T M Jr 1996 *Principles of Free-electron Lasers* II edn (London, UK: Chapman and Hall)

[7] Saldin E L, Schneidmiller E A and Yurkov M V 2008 *The Physics of Free Electron Lasers* (Berlin, Germany: Springer)

[8] Szarmes E B 2014 *Classical Theory of Free-electron Lasers* (San Rafael, California, USA: Morgan & Claypool Publishers) (Bristol, UK: IOP Concise Physics, IOP Publishing)

[9] Kim K-J, Huang Z and Lindberg R 2017 *Synchrotron Radiation and Free-electron Lasers* (Cambridge, UK: Cambridge University Press)

[10] Willmott P 2011 *An Introduction to Synchrotron Radiation: Techniques and Applications* (Chichester, UK: Wiley)

[11] Mobillo S, Boscherini F and Meneghini C (ed) 2015 *Synchrotron Radiation: Basics, Methods and Applications* (Heidelberg, Germany: Springer)

[12] Margaritondo G 2002 *Elements of Synchrotron Light for Biology, Chemistry and Medical Research* (Oxford, UK: Oxford University Press)

[13] Winick H (ed) 1994 *Synchrotron Radiation Sources: A Primer* (Singapore: World Scientific)

[14] Griffiths D J 2017 *Introduction to Electrodynamics* IV edn (Cambridge, UK: Cambridge University Press)

[15] Jackson J D 1998 *Classical Electrodynamics* III edn (New York: Wiley)

[16] Wolski A 2014 October Maxwell's equations for magnets *Technical Report* arXiv:1103.0713v2

[17] Vretenar M 2011 Radio frequency for particle accelerators—evolution and anatomy of a technology *Proc. CERN Accelerator School 2010: RF for accelerators (Ebeltoft, Denmark, 8–17 June 2010)* ed B Roger CERN–2011–007 (Geneva, Switzerland: CERN), pp 1–14

[18] Wolski A 2011 Theory of electromagnetic fields *Proc. CERN Accelerator School 2010: RF for accelerators (Ebeltoft, Denmark, 8–17 June 2010)* ed B Roger CERN–2011–007 (Geneva, Switzerland: CERN), pp 15–66

[19] Taylor C J, Young P and Chotai A 2013 *True Digital Control: Statistical Modelling and Non-minimal State Space Design* (New York: Wiley)

[20] Mistry N B and Li Y 2013 Requirements for vacuum systems *Handbook of Accelerator Physics and Engineering* II edn ed A Wu Chao, K Hubert Mess, M Tigner and F Zimmermann (Singapore: World Scientific), pp 431–32

[21] Brandt D (ed) 2007 June *Proc. CERN Accelerator School Course on Vacuum in Accelerators (Platja d'Aro, Spain, 16-24 May 2006). Technical Report CERN-2007-003* (Geneva, Switzerland: CERN)

[22] Brandt D (ed) 2009 August *Proc. CERN Accelerator School Course on Beam diagnostics (Dourdan, France, 28 May-6 June 2008). Technical Report CERN-2009-005* (Geneva, Switzerland: CERN)

[23] Rehlich K 2013 Controls and timing *Handbook of Accelerator Physics and Engineering* II edn ed A Wu Chao, K Hubert Mess, A Wu Chao, M Tigner and F Zimmermann (Singapore: World Scientific), 760–3

[24] Rees G H 2013 Ring injection and extraction *Handbook of Accelerator Physics and Engineering* II edn ed A Wu Chao, K Hubert Mess, M Tigner and F Zimmermann (Singapore: World Scientific)

[25] Roesler S and Silari M 2013 Radiation protection principles *Handbook of Accelerator Physics and Engineering* II edn ed A Wu Chao, K Hubert Mess, M Tigner and F Zimmermann (Singapore: World Scientific), pp 767–9

[26] Advanced Light Source https://als.lbl.gov [Online; accessed 29 October 2017]

[27] MAX IV https://www.maxiv.lu.se [Online; accessed 21 March 2018]

[28] Leemann S C, Sjöström M and Andersson Å 2017 May First optics and beam dynamics studies on the MAX IV 3 GeV storage ring *Proc. Eighth International Particle Accelerator Conference (Copenhagen, Denmark)* pp 2756–59

[29] SPring–8 http://www.spring8.or.jp/en/ [Online; accessed 29 October 2017]

[30] Australian Synchrotron http://www.synchrotron.org.au [Online; accessed 20 March 2018]

[31] Diamond Light Source http://www.diamond.ac.uk [Online; accessed 20 March 2018].

[32] Aßmann R, Lamont M and Myers S 2002 A brief history of the LEP collider *Nucl. Phys. B (Proc. Suppl.)* **109B** 17–31

[33] Berkelman K 2004 *A Personal History of CESR and CLEO* (Singapore: World Scientific)

[34] CESR https://www.classe.cornell.edu/Research/CESR/WebHome.html [Online; accessed 21 March 2018]

[35] DAΦNE http://www.lnf.infn.it/acceleratori/ [Online; accessed 21 March 2018]

[36] SuperKEKB http://www-superkekb.kek.jp/documents.html [Online; accessed 20 March 2018]