Tutorial 7: Particle accelerators

Dr. M Flowerdew

December 10, 2014

Now that we have covered the theoretical and historical background to the Standard Model, we will digress for a few weeks and discuss experimental matters. We will begin by examining particle accelerators, in particular what features determine their performance. To aid the discussion, three of the highest energy colliders of recent times will be used as examples: the LEP, Tevatron and LHC (see Figure 1). We will not be discussing accelerators for fixed target experiments, which have different requirements to collider experiments.

1 Key collider parameters

Four of the most important parameters of any accelerator are the type(s) of particle(s) it accelerates, the accelerator configuration (i.e. linear vs. circular), the final energy of each accelerated particle (the *beam energy*) and the luminosity that can be achieved. Example values for these quantites and others are shown in Table 1 for discussion throughout this tutorial.

To date, the only particles that have been successfully collided are electrons, positrons, protons and antiprotons, as well as nuclei and other ions (e.g. Pb or Ar). Only these particles are electrically charged and sufficiently stable to survive the early stages of collimation and acceleration. Positrons and antiprotons are not present in normal matter and must be made by colliding, say, protons into a dense target, followed by filtering and collimation of the debris. Ultimately, the choice of which particles to collide depends upon the goals of the experiments.

Exercise: Brainstorm some possible uses for the following combinations of colliding particles: e^-e^-, e^+e^-, $p\bar{p}$, pp, $e^{\pm}p$.

Exercise: What are some of the relative advantages and disadvantages of linear and circular colliders?

The maximum beam energy that an accelerator can sustain depends upon several factors. For linear accelerators, it is the length of the acceler-

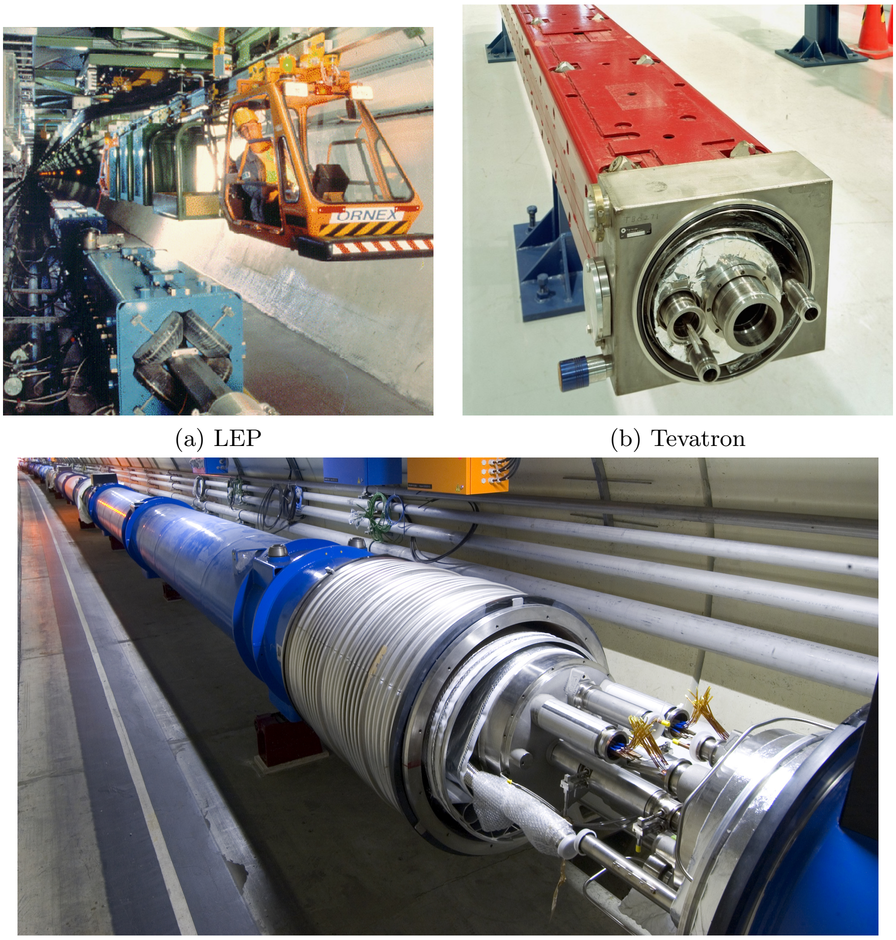

(a) LEP

(b) Tevatron

(c) LHC

Figure 1: Accelerator magnets from LEP, Tevatron and LHC.

ator itself and the electric field gradient ($\mathrm{d}E/\mathrm{d}z$) that are most important.[1] For circular accelerators, synchrotron radiation quickly becomes the dominant factor for electrons and positrons, as the accelerator has to replenish the energy lost on every turn. Circular (anti)proton accelerators are mainly limited by the magnetic field strength of the bending or *dipole* magnets. The magnetic field required to bend particles with an momentum of p around a circle of radius r is

$$B = \frac{3.336 \cdot p/\mathrm{GeV}}{r/\mathrm{m}}\ \mathrm{T}. \tag{1}$$

For the LHC, protons with $p = 7$ TeV are guided around a ring with a bending radius of $r = 2.8$ km. This requires a magnetic field of $B \approx 8.3$ T,

[1] Conventionally, the z direction is taken to coincide with the beam axis, locally at each point along the beam.

Table 1: Key parameter values for LEP, Tevatron and LHC. Values such as the peak instantaneous luminosity are quoted per experiment. The LHC parameters refer to pp operation, and the LEP's normalised emittance $\gamma\epsilon$ is quoted for $E_{\rm beam} = m_Z/2$. The symbols are described in the text of this section, except for ϵ which is the subject of Section 2.

Quantity	Unit	LEP e^+e^- 1989-2000	Tevatron $p\bar{p}$ 1983-2011	LHC pp, pA, AA 2009-present
$E_{\rm beam}$	GeV	80.5-104	900-1000	3500-4000
Max. $N_{\rm bunch}$		12	103	1380
Max. N	$\times 10^{11}$	4	2.7 (p), 1.0 ($\bar{p}$)	1.7
$\gamma\epsilon$	mm mrad	18 (y) 1800 (x)	63 (p) 47 ($\bar{p}$)	2.5
Peak $\mathcal{L}$	10^{33} cm^{-2}s^{-1}	0.1	0.52	7.7

which is produced by superconducting magnets cooled down to liquid helium temperatures. Acceleration to the final beam energy is rarely achieved in one step. For example, the accelerator complex at CERN is shown in Figure 2. This shows particle beams being produced for many experiments besides the LHC, but even for the LHC no fewer than five storage rings and two linear accelerators are required. For reasons of brevity, most of the discussion here will concern the final accelerator stage, e.g. the LHC ring itself in this case.

The *instantaneous luminosity* of a collider ($\mathcal{L}$) is a measure of how often particles have the opportunity to collide. It is usually measured in units of cm^{-2}s^{-1}, and is defined so that $\sigma_{\rm tot}\mathcal{L}$ is the interaction rate, if $\sigma_{\rm tot}$ is the total interaction cross-section. The *integrated luminosity* $L = \int \mathcal{L}\,{\rm d}t$ is then a measure of the total data collected by a collider experiment. If two identical beams with Gaussian profiles collide perfectly head-on, then the instantaneous luminosity can be written as

$$\mathcal{L} = \frac{N^2 N_{\rm bunch} f_{\rm rev}}{4\pi\sigma_x\sigma_y}, \tag{2}$$

where N is the number of particles per bunch,[2] $N_{\rm bunch}$ is the number of bunches in each beam, $f_{\rm rev}$ is the revolution frequency, and $\sigma_{x(y)}$ is the width of the beam in the $x(y)$ direction. If the beams are not perfectly aligned, $\mathcal{L}$ also depends exponentially on the square of the offset between the centres of the two beams.

Exercise: Derive Equation (2).

[2] Particles are typically accelerated in discrete bunches rather than continuous beams. One of the reasons for this is explored in Section 4.

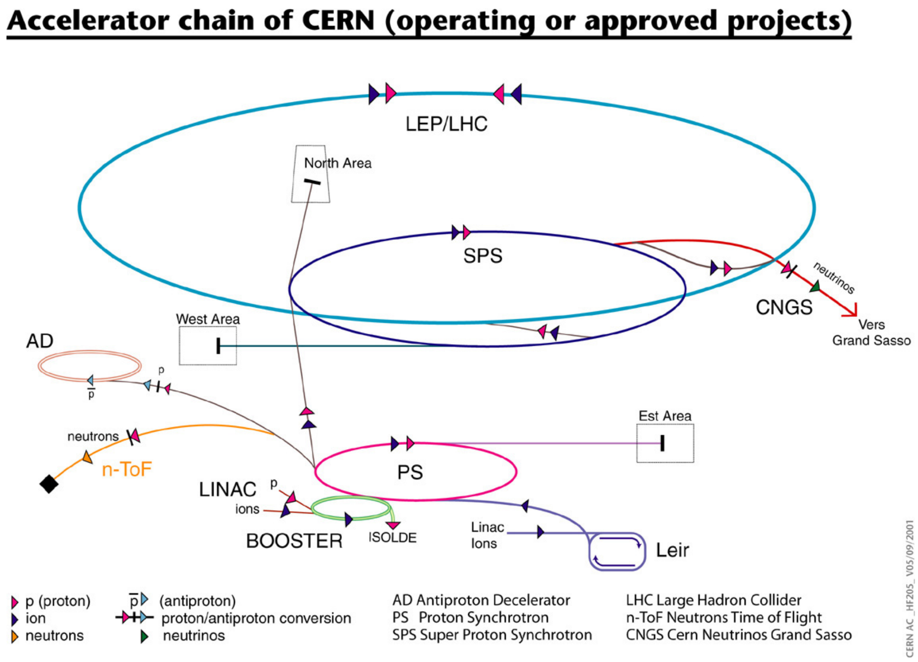

Figure 2: Accelerator complex at CERN, including preaccelerators for the LHC.

Increasing the instantaneous luminosity allows processes with smaller cross-sections to be explored. Higher luminosities can be achieved by increasing the rate of bunch-bunch collisions $N_{\text{bunch}} f_{\text{rev}}$. Having a large number of bunches in a single beam is technically complex, and a relatively recent development (see Table 1). The number of particles per bunch is difficult to increase much beyond $\mathcal{O}(10^{11})$, due to electromagnetic interactions of the bunches ($10^{11}e = 16$ nC) and also the total beam current is limited by the power required to accelerate and bend the beam. The final factor in Equation (2) is the beam area $A = 4\pi\sigma_x\sigma_y$, which should be small to achieve a high instantaneous luminosity. For this reason, beams are usually heavily focussed close to the collision point, however this is practically limited by considerations of the beam optics. This is what we consider next.

2 Emittance and focussing

The emittance of a beam describes the range of deviations from the ideal path that the particles take, and is a key parameter in determining the final focussing of a beam at an interaction point. If we consider just one direction perpendicular to the beam, say x, then the emittance ϵ_x is the volume of phase space occupied by some specified fraction of the beam, say 68%. The phase space is parameterised by x and $x' = \mathrm{d}x/\mathrm{d}s \approx \frac{1}{c}\mathrm{d}x/\mathrm{d}t$, where s

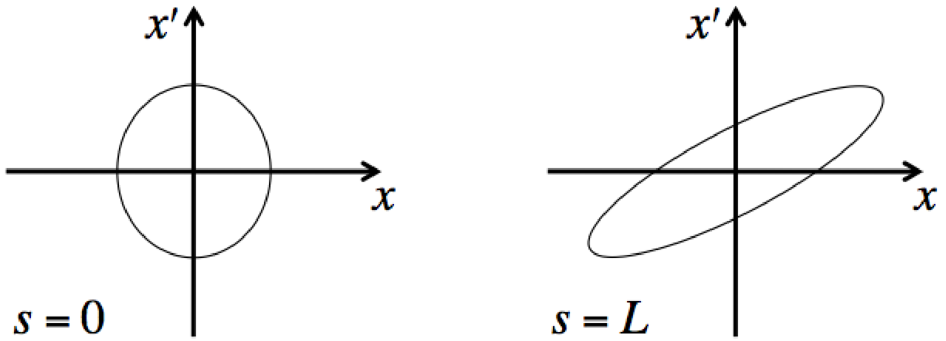

Figure 3: Sketch of a beam envelope in $x - x'$ space for a beam that freely propagates over a distance L. It is assumed that x and x' are uncorrelated at $s = 0$, for simplicity.

parameterises the distance along the ideal beam line. Due to Liouville's theorem, this volume is conserved as it propagates through the accelerator.

At the point of production, the position and direction of the particles are essentially uncorrelated, meaning that the volume defining the emittance is an ellipse (as in Figure 3 (left)). Even though individual particles have complicated paths through the accelerator, we can understand the propagation of the beam as a whole by considering just this envelope. This is because particles in the beam that start inside the envelope cannot cross it – two particles with the same position and velocity will experience the same force, and their future paths must be identical.

We begin by considering the free propagation of the particles within this initial ellipse. The regions of the ellipse with positive x' will migrate to higher values of x over time, and regions with negative x' will migrate to negative values of x. This is illustrated in Figure 3 (right). This is an intuitive result: the beam spreads out in x over time due to the initial spread in x velocities.

This effect can of course be reversed by focussing. If a magnetic field is arranged such that its magnitude in the y direction varies as $-gx$, where g is a constant, then the force experienced in the x direction by a particle with velocity $v \approx (0, 0, c)$ will be $-gcx$, i.e. a restoring force. This field acts like a lens with a focal length of p/egl, where p is the beam energy, and l is the length of the focussing magnet, reducing x' for parts of the ellipse with positive x, and vice versa for negative x. Figure 4 shows an example where the focussing exactly compensates for the increased beam width, effectively reversing the sign of x' with respect to the beam before focussing. After propagating for another distance L, the original shape of the beam from Figure 3 (left) can be recovered. Thus, with repeated focussing, the overall beam size in x can be maintained.

The field required to focus in the x direction can be achieved with a

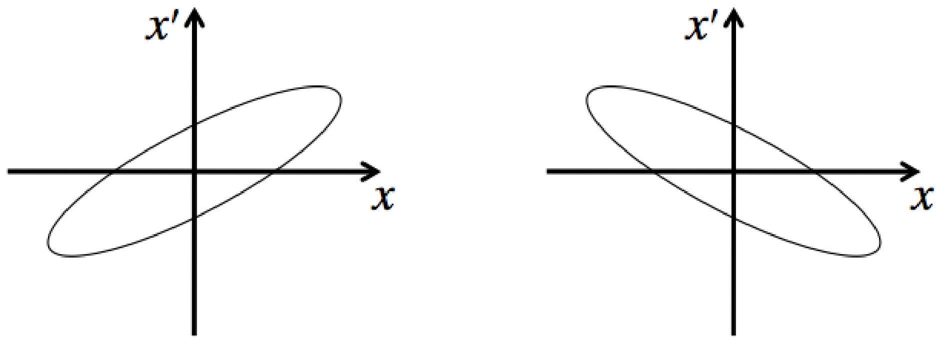

Figure 4: Sketch of a beam envelope in $x - x'$ space before (left) and after (right) focussing in the x direction. It is assumed that the focussing is perfectly tuned to the beam, i.e. that $x' \to -x'$ for all particles in the beam.

quadrupole magnet (for an example, see Figure 1 (a)). However, the full quadrupole field in the $x - y$ plane has components $B = (-gy, -gx, 0)$, giving a total force on the particle of $F = v \times B = (-gvx, gvy, 0)$. This will defocus the beam in the y direction, increasing its divergence. Fortunately, it is possible to achieve focussing in both directions by alternating quadrupole magnets that focus in x and y. To see why this works, recall that the focal length f of two lenses with focal lengths f_1 and f_2 and separated by a distance d is given by

$$\frac{1}{f} = \frac{1}{f_1} + \frac{1}{f_2} - \frac{d}{f_1 f_2}. \tag{3}$$

If f_1 and f_2 have opposite sign, then f can always be made positive (and therefore focussing) as long as d is sufficiently large. In particular, if $f_2 = -f_1$, then f is always positive.

Now consider strong focussing close to an interaction point, necessary to increase the instantaneous luminosity (Equation (2)). The emittance ellipse of the right-hand part of Figure 4 would stretch up and down to more extreme values of x'. Due to Liouville's theorem, the width of the ellipse would shrink, and after a short propagation time the beam size in x would be much smaller than it started. This is where the interaction point should be located. However, high values of x' correspond to particles travelling at large angles with respect to the ideal beam line, and so the beam will quickly diverge again, requiring more focussing to avoid losses from collisions with the beam pipe wall. Therefore, the final focussing magnets should be as close to the interaction point as possible, to increase the maximum tolerable beam divergence. In addition, it is desirable to have the overall emittance as low as possible, which allows for a smaller beam size for a given maximal divergence. This is achieved through beam cooling.

3 Beam cooling

The restrictions of Liouville's theorem only apply to a closed system that does not exchange energy with its surroundings. If this assumption is broken, then it is possible to alter the emittance of the beam. Even the act of accelerating the beam (see Section 4) reduces the transverse emittance as defined, because p_z increases while p_x and p_y remain the same. Therefore, $x' = p_x/p_z$ is reduced, and similarly for y'. However, cooling is usually understood to mean a reduction in the normalised emittance $\gamma\epsilon_{x,y}$.[3]

One of the simplest ways to reduce the (normalised) emittance of a beam is to wait until it is spatially extended (as in Figure 3) and insert a beam stop restricting its width. Due to the correlation between x and x' in that case, this will also reduce a substantial fraction of the particles that contribute most to the beam divergence. This technique is most useful in the early stages of beam production, where the energy per particle is relatively small.

The primary method for cooling high-energy beams is in a damping ring. For a circular collider, the collider ring itself can act as a damping ring, while for a linear collider the damping ring would be a separate component away from the main accelerator. As the particles circulate around the damping ring, they lose energy to synchrotron radiation. This reduces all components of the particles' momenta, while acceleration to maintain a constant beam energy only increases p_z. Thus, over time, the beam divergence decreases.

The energy lost per particle per revolution due to synchrotron radiation scales as γ^4/r, where r is the radius of curvature, assuming $\beta \approx 1$. Thus, electron and positron beams can be cooled very effectively with very low beam energies. At high energies the discrete nature of the photon emission process adds noise to x' and puts a limit on the lowest achievable emittance that depends on the beam energy. This is the main reason for the relatively poor value of $\gamma\epsilon_x$ quoted for LEP in Table 1. Synchrotron radiation also ultimately limits the beam energy that a circular e^+e^- collider can sustain.

The radiative damping time for (anti)protons is much longer than for $e^\pm$ for the same accelerator parameters. For this reason *stochastic cooling* is often used to accelerate the cooling process. This uses readings taken of the beam in one part of the ring to correct the beam profile in another part of the ring, which is possible because the straight-line distance between the two points is shorter than the path taken by the beam. It is best if corrections can be applied to parts of a bunch, rather than the whole bunch, and so typically the beam is stretched in z before the corrections are applied. Over time, the average deviations from the ideal beam line can be reduced, thus cooling the beam. The invention of this procedure led to the discovery of the W and Z bosons and the awarding of the 1984 Nobel Prize to Simon van der Meer (together with Carlo Rubbia).

[3] Actually $\beta\gamma\epsilon_{x,y}$, but we are assuming that $\beta \approx 1$.

Figure 5: Example RF cavity proposed for a future linear e^+e^- collider.

4 RF cavity acceleration

To accelerate particles to energies of tens to thousands of GeV, alternating electric and magnetic fields must be used (**Question:** why not static electric fields?). This is normally achieved using radio-frequency (RF) cavities like the one illustrated in Figure 5. Each cell of the RF cavity oscillates in anti-phase to its neighbours, at a characteristic frequency determined by the geometry of the cavity. If the particle bunches are timed correctly, they will pass through two cavities in the oscillation period, and thus be accelerated by every cell.

Realistic cavities, such as the one in Figure 5, are highly optimised to produce the best field properties for acceleration. We will now consider a simpler variation called a pill-box cavity, shown in Figure 6. In this case, the resonant volume is a simple cylinder. There are solutions of Maxwell's equations for this geometry where the electric field points purely along z, while the magnetic field is circular in ϕ, as shown. The boundary conditions at the cavity walls mean that the electric field must vanish at $\rho = a$, where ρ is the radius variable in cylindrical coordinates. In this case, both electric and magnetic fields are described by Bessel functions, and for the lowest frequency mode the maximum electric field strength is found along the axis of the cylinder. The resonant frequency is fixed by the cavity's radius, and is $f = 2.405c/2\pi a$. This in turn determines the ideal distance between adjacent cavities, $d = \pi a/2.405$,[4] such that successive cavities accelerate the particles constructively.

Counter-intuitively, RF cavities are usually designed so that the bunches arrive just before the maximum of the oscillation in the electric field, rather than at the maximum. This is to allow longitudinal focussing, which can be understood using Figure 7. Point S is defined as the ideal time of arrival for a bunch to be accelerated, where it will just reach the next cavity at the same point in its oscillation. A particle travelling slightly faster than the average will arrive early, perhaps at point P. In this case, it experiences a lower electric field, and will be accelerated less in this cycle. Conversely, a late particle (at P') will be accelerated more than average. These lead to small oscillations around S as particles gain and lose energy in different RF cavities. For very late particles, point U marks the divide between two

[4] Assuming $v = c$.

successive stable equilibria. Particles arriving later than U will eventually end up in the next bunch, after a period of deceleration.

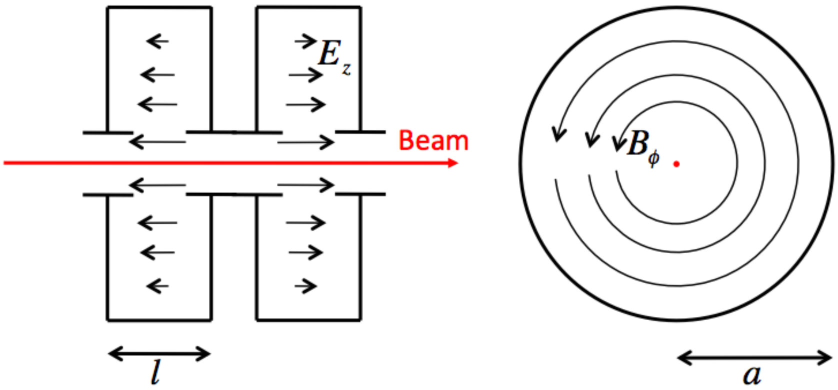

Figure 6: Schematic of a simple "pill box" RF accelerator cavity. Projections parallel and perpendicular to the beam are shown, together with arrows indicating the directions of the electric and magnetic field inside the cavity. With the electric field in this configuration, the particle bunch should be in the right-hand cavity.

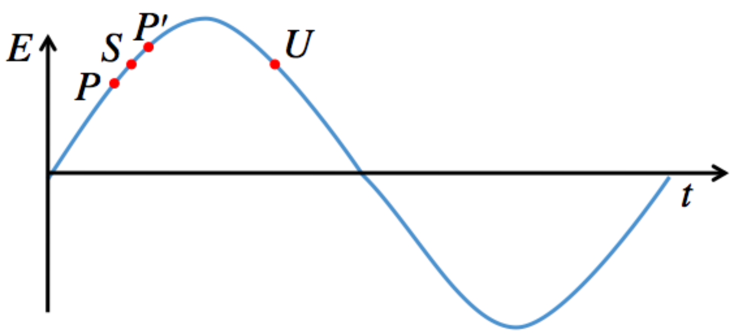

Figure 7: Sketch of the electric field at the centre of an RF cavity, as a function of time. The points indicate different times at which particles to be accelerated may pass through the cavity. S and U show stable and unstable equilibrium positions, respectively, while particles passing through at P and P' are both pushed towards S.