# Tutorial 7: Particle accelerators
# With answers

Dr. M Flowerdew

January 27, 2016

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: 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 Introduction

The four most generally important accelerator parameters are

- which particle(s) it accelerates;
- the accelerator topology (i.e. linear or circular);
- the final energy of each accelerated particle (the *beam energy*);
- the instantaneous luminosity of collisions.

A more complete discussion would include items such as the duty cycle (which affects the integrated luminosity), details of the optics, losses due to impedance and other effects, beam-beam interactions and so on. Some key parameters of the LEP, Tevatron and LHC accelerators are given inTable 1 for reference. We will briefly discuss these four main parameters in turn.

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. Before this can happen, the relevant particles must be produced in their ionised state. For practical reasons (see Section 4), the particles are usually injected into the accelerator in discrete packets called *bunches*.

**Exercise 1:** Discuss the initial production of a) electrons, b) protons, c) antiparticles.

**Answer:** a) Electrons can be produced using an electron gun, essentially a hot anode where released electrons accelerate towards a high-voltage cathode. The cathode has an aperture so that the electrons can escape the device.

b) Protons are produced from molecular hydrogen. The hydrogen molecules are ionised and then accelerated by an electric field.

c) 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.

Acceleration of the particles up to their collision energy is usually a multi-stage process. 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.

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 LEP's normalised emittance $\gamma\epsilon$ is quoted for $E_{\text{beam}} = m_Z/2$. The symbols are described in the text of this section, except for $\epsilon$ which is the subject of Section 2.

| Quantity | Unit | LEP<br>$e^+e^-$<br>1989-2000 | Tevatron<br>$p\bar{p}$<br>1983-2011 | LHC<br>$pp$, $pA$, $AA$<br>2009-present |
|---|---|---|---|---|
| $E_{\text{beam}}$ | GeV | 80.5-104 | 900-1000 | 3500-6500 |
| Max. $N_{\text{bunch}}$ | | 12 | 103 | 2232 |
| Max. $N$ | $\times 10^{11}$ | 4 | 2.7 ($p$), 1.0 ($\bar{p}$) | 1.7 |
| $\gamma\epsilon$ | mm mrad | 18 ($y$)<br>1800 ($x$) | 63 ($p$)<br>47 ($\bar{p}$) | 2.5 |
| Peak $\mathcal{L}$ | $10^{33}$ cm$^{-2}$s$^{-1}$ | 0.1 | 0.52 | 7.7 |

The maximum beam energy that an accelerator can sustain depends upon several factors. One key factor is the maximum accelerating gradient ($\mathrm{d}E/\mathrm{d}z$) that can be achieved.[1] If the beams are required to bend then losses from synchrotron radiation must also be taken into account, as well as the magnetic field strength of the bending or *dipole* magnets. The magnetic

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

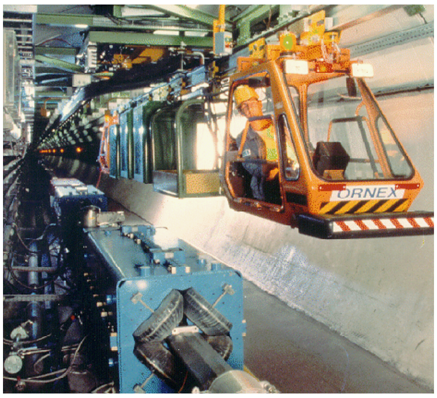


(a) LEP

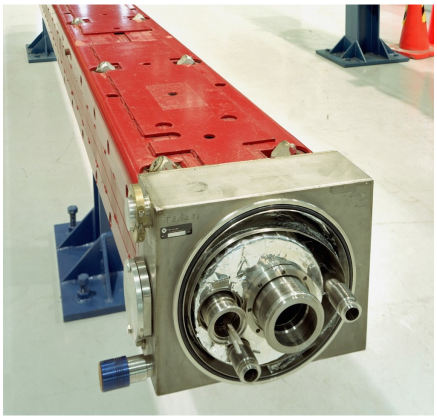

(b) Tevatron

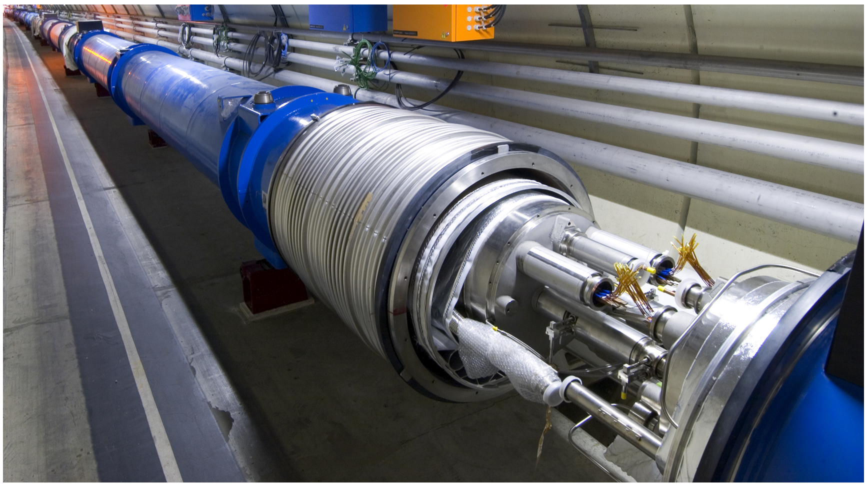

(c) LHC

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

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, which is produced by superconducting magnets cooled down to liquid helium temperatures.

**Exercise 2:** Briefly discuss the differences between linear the circular colliders. Consider their physical size, limitations on the interaction rate and any other factors you think are relevant. Consider also how

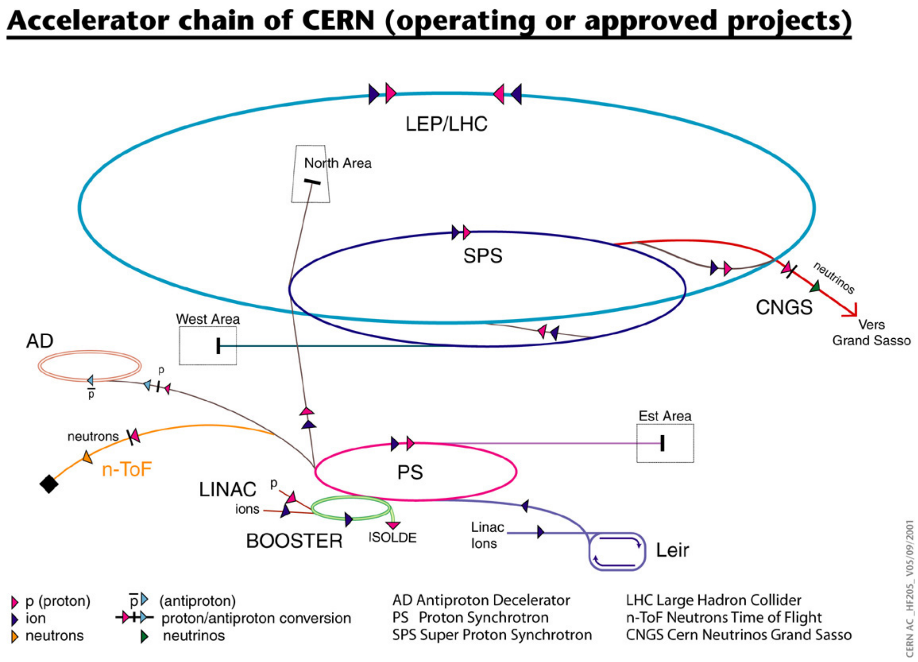


Figure 2: Accelerator complex at CERN, including preaccelerators for the LHC.

the key factors might differ for $e^+e^-$ and hadron colliders.

**Answer:** There are many possible answers to this question. One of the principle advantages of a circular collider is that the same particles can collide many times over several hours, whereas beams in a linear collider are used once and then dumped. For a circular collider of a fixed radius and energy, synchrotron radiation is much more important for $e^+e^-$ colliders than for hadrons, as the energy radiated depends inversely on the fourth power of the particle's mass. At LEP, this was already the limiting factor, which is why linear $e^+e^-$ colliders are being considered for the future. For protons, the maximum collision energy of a circular collider is typically determined by the dipole field.

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 $\text{cm}^{-2}\text{s}^{-1}$, and is defined so that $\sigma_{\text{tot}}\mathcal{L}$ is the interaction rate, if $\sigma_{\text{tot}}$ is the total interaction cross-section. The *integrated luminosity* $L = \int \mathcal{L}\,\mathrm{d}t$ 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

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

where $N$ is the number of particles per bunch, $N_{\text{bunch}}$ is the number of bunches in each beam, $f_{\text{rev}}$ 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 3:** Derive Equation (2).

**Answer:** We will solve this problem by calculating the total interaction rate, and deriving $\mathcal{L}$ from that. Consider the bunches travelling around the ring during a time period $1/f_{\text{rev}}$. In this time, each bunch of each beam will pass any fixed point once, and so there will be $N_{\text{bunch}}$ bunch-bunch collisions, or $N_{\text{bunch}} f_{\text{rev}}$ collisions per second. If we assume for simplicity that each beam is cylindrical, with volume $AL$, then each particle traverses a volume $L\sigma_{\text{tot}}$ of the other beam.[a] It therefore could interact with $NL\sigma_{\text{tot}}/AL = N\sigma_{\text{tot}}/A$ particles. Given that there are $N$ particles in each beam, we then obtain

$$\begin{aligned}\mathcal{L} &= \frac{1}{\sigma_{\text{tot}}} \cdot N_{\text{bunch}} f_{\text{rev}} \cdot \frac{N\sigma_{\text{tot}}}{A} \cdot N \\ &= \frac{N^2 N_{\text{bunch}} f_{\text{rev}}}{A},\end{aligned}$$

in agreement with Equation (2).

[a] This assumption is crude, and yet it captures the essential features we are after.

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 $N_{\text{bunch}} f_{\text{rev}}$, increasing the number of particles per bunch, or by decreasing the beam area $A = 4\pi\sigma_x\sigma_y$. 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 also difficult to increase much beyond $\mathcal{O}(10^{11})$, due to electromagnetic interactions of the bunches and also limitations on the total beam current. Small beam areas are achieved via focussing them close to the collision point. This is limited by considerations of the beam optics, which 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 $\epsilon_x$ is the volume of phase space occupied by some specified fraction of the beam, say 68%. The

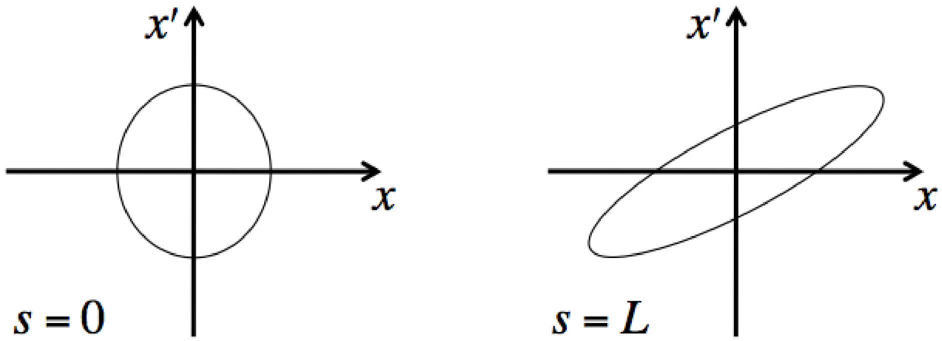


Figure 3: Sketch of a beam envelope in $x$–$x'$ space for a beam that freely propagates over a distance $L$. It is assumed that $x$ and $x'$ are uncorrelated at $s = 0$, for simplicity.

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 $s$ 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 be reversed by focussing.

**Exercise 4:** Consider the magnetic field produced by a quadrupole magnet (for $x, y \ll$ the magnet separation)

$$B = \begin{pmatrix} -gy \\ -gx \\ 0 \end{pmatrix}. \tag{3}$$

Sketch the field lines and indicate schematically where coils would be placed in a real system. What are the forces exerted on a particle

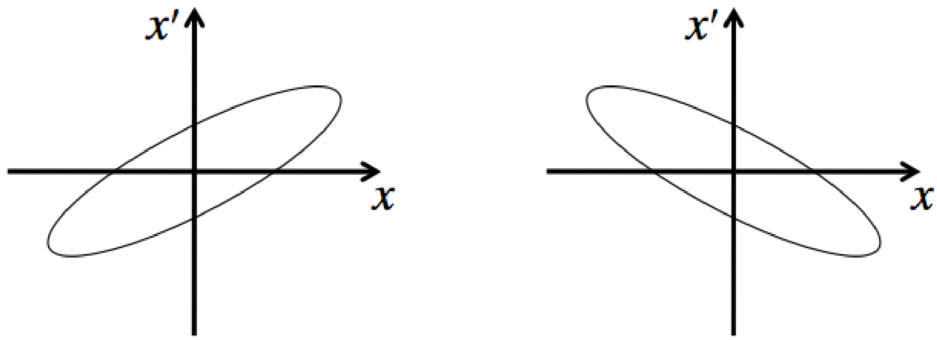


Figure 4: Sketch of a beam envelope in $x$–$x'$ space before (left) and after (right) focussing in the $x$ direction. It is assumed that the focussing is perfectly tuned to the beam, i.e. that $x' \to -x'$ for all particles in the beam.

traversing the quadrupole in the $x$ and $y$ directions? You may assume that the particle is highly relativistic, i.e. $v \approx c$. Convince yourself that, depending on the magnet's length in $z$, the beam can be focussed in the $x$ direction. What effect does this field have on $y'$?

**Answer:** This is a field directed along the radius vector for $y = -x$, and against the radius vector for $y = x$. There is a saddle point at the origin. Coils could be placed around the $y = \pm x$ axes to produce a field approximating Equation (3) for small displacements.

The Lorentz force on a particle with coordinates $(x, y, z)$ and velocity $v = (0, 0, c)$ is

$$F = ev \times B = \begin{pmatrix} 0 \\ 0 \\ c \end{pmatrix} \times \begin{pmatrix} -gy \\ -gx \\ 0 \end{pmatrix} = \begin{pmatrix} -egvx \\ egvy \\ 0 \end{pmatrix}.$$

The force in the $x$ direction is proportional to $-x$, i.e. it is 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 $x$. The force in the $y$ direction is proportional to $y$ and therefore has a defocussing effect, increasing the divergence of a beam with a profile similar to Figure 3 (right).

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.

As found in Exercise 4, focussing in $x$ is detrimental to the beam quality in $y$. 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

$$\frac{1}{f} = \frac{1}{f_1} + \frac{1}{f_2} - \frac{d}{f_1 f_2}. \tag{4}$$

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)). If the $x$-width of the ellipse in Figure 3 (left) is to become substantially narrower, then the size in $x'$ must increase by Liouville's theorem. 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. It is also clearly desirable to keep the overall emittance as small as possible, which can be achieved via 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 $\gamma\epsilon_{x,y}$.[2]

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 (right)) 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.

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

[2]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.

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 $\gamma^4/r$, 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).

## 4 RF cavity acceleration

To accelerate particles to energies of hundreds of GeV, alternating electric and magnetic fields must be used. 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.

**Exercise 5:** Why are static electric fields unsuitable for accelerating particles to GeV-scale energies?

**Answer:** From the definition of an electron-Volt, acceleration to 1 GeV would require a static electric field of 1 GV (or 0.5 GV if positive and negative polarities are used). Even putting aside the obvious safety issues, it would be impractical to build such a machine. For example, the breakdown field strength (dielectric strength) of air is 3 MV/m, and so the high-voltage portion would have to be located several hundreds of metres away from the ground supply.

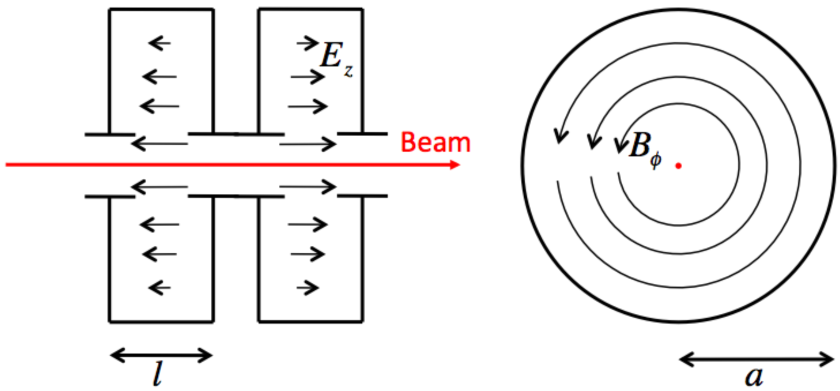


Figure 6: Schematic of a simple "pill box" RF accelerator cavity. Projections parallel and perpendicular to the beam are shown, together with arrows indicating the directions of the electric and magnetic field inside the cavity. With the electric field in this configuration, the particle bunch should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

Realistic cavities, such as the one in Figure 5, are highly optimised to produce the best field properties for acceleration. A simpler variation called a pill-box cavity is 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 $\phi$, as shown.

**Exercise 6:** What boundary condition(s) must the electric field in Figure 6 satisfy? Derive or look up the lowest-frequency solution to Maxwell's equations in this case. Show that the resonant frequency is given by $f = 2.405c/2\pi a$, where $a$ is the cavity's radius. Find the necessary distance between the centres of adjacent cavities if they are to both accelerate the particles in-phase.

**Answer:** The boundary conditions at the cavity walls mean that the electric field must vanish at $r = a$, where $r$ 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 lowest-order Bessel function, $J_0(\rho)$, vanishes when $\rho = 2.405$. In this case, $\rho = \omega r/c = 2\pi f r/c$, and substitution gives the answer we require. The required distance between adjacent cavities is $d = c/(2f) = \pi a/2.405$,[a] where the factor of 2 arises because we want the particle to arrive at the next cavity after half an oscillation.

[a]Assuming $v = c$.

With careful timing, it is possible for the RF cavities to focus the bunches longitudinally, i.e. in the $z$ direction. This is achieved if the bunches arrive just before the maximum of the oscillation in the electric field, rather than at the maximum, as illustrated in 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 successive stable equilibria. Particles arriving later than $U$ will eventually end up in the next bunch, after a period of deceleration.

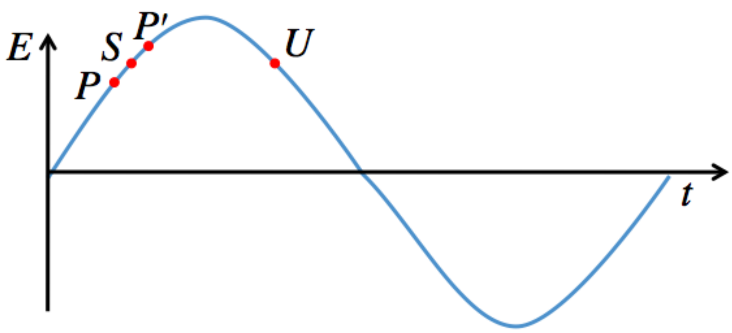


Figure 7: Sketch of the electric field at the centre of an RF cavity, as a function of time. The points indicate different times at which particles to be accelerated may pass through the cavity. $S$ and $U$ show stable and unstable equilibrium positions, respectively, while particles passing through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