Hydrogen balloon

You feel the shock wave as the fireball of a hydrogen explosion rises from a recently-lit balloon. This demonstration is a truly spectacular, if slightly gratuitous, reminder of where the protons accelerated in the LHC are found.

Apparatus

1 × hydrogen cylinder
1 × hydrogen regulator with flashback arrestor
3 × balloons (helium-grade latex deflate slower)
3 × balloon weights
string
1 × candle
1 × long stick
1 × lighter

Safety

This demonstration is both very hot and very loud. Be sure presenters wear eye and ear protection, that the audience is sat well back and that they cover their ears during the explosion. Also be aware of potential fire risks.

The demonstration

1. Before the show, fill the balloons with hydrogen and attach them with string to the weights.
2. Light the candle on the end of the stick.
3. Standing well back, use the candle to light the balloon.
4. Watch as a spectacular fireball ensues.

Two hydrogen balloons are exploded in tandem at the end of an *Accelerate!* show on tour at a science festival.

Vital statistics

reaction:

$$2H_2 + O_2 \rightarrow 2H_2O$$

adiabatic flame temperature:

2200°C

How it works

Hydrogen is used as a source of protons in the Large Hadron Collider. The simplest element, it is made up of just one proton and one electron, and that electron can be stripped off with electrostatic fields, leaving the proton to be accelerated in the collider.

A more visually exciting use of hydrogen is as a highly efficient fuel source: liquid hydrogen and liquid oxygen are used to create the huge controlled explosions which lift rockets off the ground and into space. Here, a balloon full of hydrogen and atmospheric oxygen are combined to make a slightly less controlled explosion.

We often explode two of these balloons to close the show as well as the one which is strictly necessary for explanation early on!

Van de Graaff generator

The Van de Graaff generator is a classroom classic with a surprising heritage in cutting-edge particle physics. As well as making your hair stand on end, these machines were used to accelerate particles through millions of volts.

Apparatus

1 × Van de Graaff generator
1 × electrically-insulating stool
Some confetti, or aluminium foil, or foil cake trays

Vital statistics

breakdown voltage of air:
30,000 V/cm

highest-voltage Van de Graaff:
25.5 MV

length of spark from a Van de Graaff LHC:
7 TV ÷ 30 kV/cm = **2300 km**

The demonstration

This demonstration involves high voltages, and so it should never be done by anyone with a pacemaker or other internal electrical device, or who thinks they might be pregnant.

The first part of this demo requires a volunteer from the audience. It works best on someone with long, light-coloured hair free from ties and styling products: light hair is often thinner, which means it will stand up more easily, and is also easier to see. Don't pick on an individual (they might be pregnant, or just shy!), but if you could encourage someone fitting that description that gives you the best chance of success.

1. Give your volunteer a round of applause and find out their name. Check that they aren't wearing any metal jewelry etc, and ask that they remove it if so. Put it safely to one side.

2. Get the volunteer to stand on an electrically-insulating stool. Place one hand on top of the generator dome, and get them to hold out their other hand flat. In this, place some confetti, pieces of aluminium foil, or cake tins. If you have a suitable light, dim the main lights and illuminate their head from behind to emphasise the forthcoming hairdo.

3. Check they're feeling OK, turn on the generator, and stand back. The only thing they need to do is not to take their hand from the dome and, should they do so, not try to replace it (or they'll get a shock!). Make sure the earthed globe is a long way from the main one to avoid shocking the volunteer too.

4. The objects in their hand will leap out and, by the time that's finished, their hair should be standing up pretty nicely. Get them to shake their head around a bit to encourage this.

5. Get the volunteer to take their hand off the dome, and jump down with both feet, and give them a round of applause!

Having shown the amusing effect of high voltage on a person, we can now

explore the limitations of these devices as particle accelerators. The problem is sparks, and we can use the sparks to work out the voltage to which we charged up our hapless volunteer!

1. If you haven't already, dim the lights.
2. Take the grounded sphere and place it near the dome of the Van de Graaff generator. When you get within a few centimetres, a spark should leap across with a crack. Do this a few times from different angles to show the audience.

How it works

The rubber belt inside the Van de Graaff generator runs between two rollers made of different materials, causing electrons to transfer from one roller to the rubber, and from the rubber onto the other roller, by the triboelectric effect. Brushes at the top and bottom provide a source and sink for these charges, and the top brush is electrically connected to the Van de Graaff's dome and so the charge will spread out across the dome.

This accumulated charge would like to distribute itself over as large a volume as possible, and so it will also spread out across anything you connect to the metal dome, including your volunteer. The reason it's important to stand them on something electrically-insulating is that the charge would like even more to spread out over the whole Earth, and connecting them to that will both massively reduce the effect, and also cause an electric shock as the current flows from the Van de Graaff to earth through its unfortunate human intermediary.

When insulated, the build-up of charge on the volunteer causes forces light objects to spread out as far as possible too, causing the confetti or foil to leap from their hand and then causing the individual hairs on their head to stand up. When they jump off the stool, the charge immediately flows to earth and their hair will immediately return to normal.

To work out the voltage of the Van de Graaff generator, and thus the voltage on our volunteer, we can use the length of the sparks in combination with the breakdown voltage of air—the voltage required to cause air to dissociate into ions and become conductive. This voltage is about 30,000 V/cm for dry air (hot, humid or lower-pressure air will tend to spark more easily). Sparks from the Van de Graaff are typically a few centimetres long, giving a voltage between 50,000 and 150,000 V.

A hair-raising experience with a Van de Graaff generator: illumination from behind helps emphasise the hairdo.

Their propensity to generate sparks is the fundamental limitation of Van de Graaff accelerators, or indeed any accelerator design based on a large, static voltage. Those used for research managed to get up to over 20 MV by clever use of insulating materials, right down to careful choice of the gas in which the generator sits to minimise the chance of sparks. A Van de Graaff can thus be used to accelerate particles up to reasonably high energies: moving an electron through 1 V gains it an energy of 1 eV, so energies of over 20 MeV are achievable by this method (and more if accelerating nuclei with greater than a single electron charge).

However, modern particle physics has gone some way beyond this: the Large Hadron Collider will ultimately use beams which have 7 TeV of energy each: equivalent to accelerating a proton through 7,000,000,000,000 V. If we divide that by the breakdown voltage of air, we can work out the length of a spark we might get from an LHC employing a single, giant Van de Graaff generator to accelerate its particles. We get 2,300 km: easily enough to stretch, for example, from Switzerland to anywhere in the UK.

Plasma ball

The graceful purple arcs of plasma dancing in a plasma ball are created by a large alternating voltage at its centre, and that alternating voltage creates an electromagnetic field with which we can light a fluorescent tube.

Apparatus

1 × plasma ball
1 (or more) × reinforced fluorescent tube(s)

The demonstration

1. Turn on the plasma globe.
2. Bring the fluorescent tube near the plasma ball. Before they touch, the tube should light!

How it works

At the centre of a plasma ball is a large alternating voltage, typically a few kilovolts oscillating at around 30 kHz. The low density of the gas in the globe (often neon) makes discharge significantly more favourable than it is in air at atmospheric pressure (the breakdown voltage of air which causes sparks from a Van de Graaff generator, for example, is 30,000 V/cm, whilst this can create arcs many centimetres long with just a few thousand volts). These fronds of plasma make their way from the centre of the globe to the edge, in a bid to reach earth. Creating an enhanced path to earth by touching the globe increases the strength of the discharge, which is why the arcs are attracted to your hand if you touch the globe.

Vital statistics

oscillating voltage:
2–5 kV at around 30 kHz

gas inside the globe:
usually neon

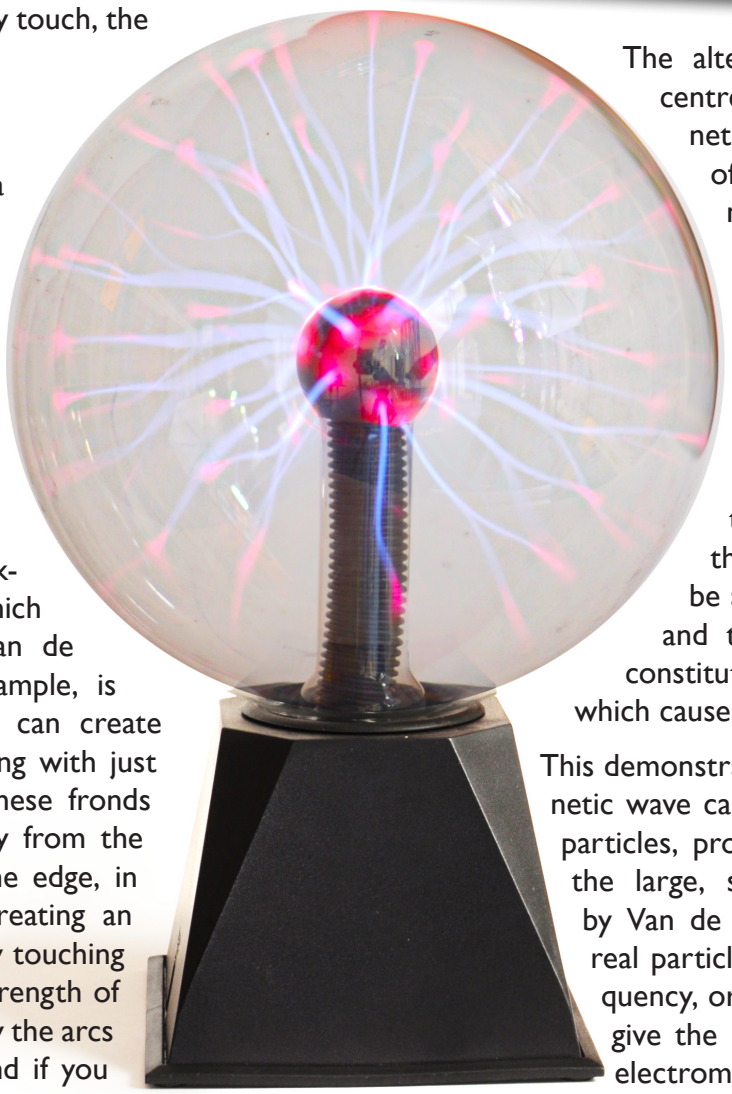

The alternating voltage at the centre creates electromagnetic waves, and the arcs of plasma act as antennae, meaning that the extent of the electromagnetic field surrounding the ball is significantly larger than the bounds of the glass globe. Bringing the fluorescent tube near to the plasma ball allows the electrons inside to be accelerated by this field, and those moving electrons constitute an electric current, which causes the bulb to light up.

This demonstrates that an electromagnetic wave can be used to accelerate particles, providing an alternative to the large, static voltages supplied by Van de Graaff generators. In a real particle accelerator, radio-frequency, or RF, cavities are used to give the particles a kick with an electromagnetic standing wave.

Beach ball accelerator

How can we use a wave to accelerate particles? The whole audience is involved in this massive, chaotic, hands-on visualisation of what that rather cryptic idea means, as we accelerate beach balls using a Mexican wave.

Apparatus

n × beach balls
$n/2$ × electric pumps

The demonstration

The first demonstration shows how a wave can be used to accelerate particles:

1. Place approximately as many beach-balls as there are rows in the audience at one side of the auditorium.
2. Encourage the audience to do a 'Mexican wave', transporting the beach-ball along its crest and accelerating it from stationary to nearly the speed of light.

What it shows

In an accelerator, the charged particles are accelerated by an electromagnetic standing wave, resonating within a radio-frequency, or RF, cavity. That's obviously something of a mouthful, and indeed a mind-full, so the concept is got across in stages, the first of which being how it's possible to use a wave to accelerate a particle. The beach-ball, or proton, surfs along the 'Mexican wave' of humans, gathering speed. This is in much the same way as a surfer surfs on a water wave, gaining speed—and therefore energy—as they do so. Similarly, we can use an electromagnetic wave to accelerate particles surfing along on it.

Another important feature of RF cavities which this demo could be used to illustrate to a more advanced audience is the importance of synchronisation: if an audience member stands up before the beach-ball arrives, or after it's passed, there won't be any additional acceleration. By analogy, it's like making sure you push a roundabout when one of the bars is passing your hand, or you won't provide any energy to increase its rotational speed. Similarly, the rapidly-alternating fields inside an RF cavity must be synced precisely with the particles zooming through them lest they fail to accelerate or, worse, actively decelerate the beam!

Vital statistics

scale model protons:
a 1 m beach ball is around 10^{15} times larger than a proton

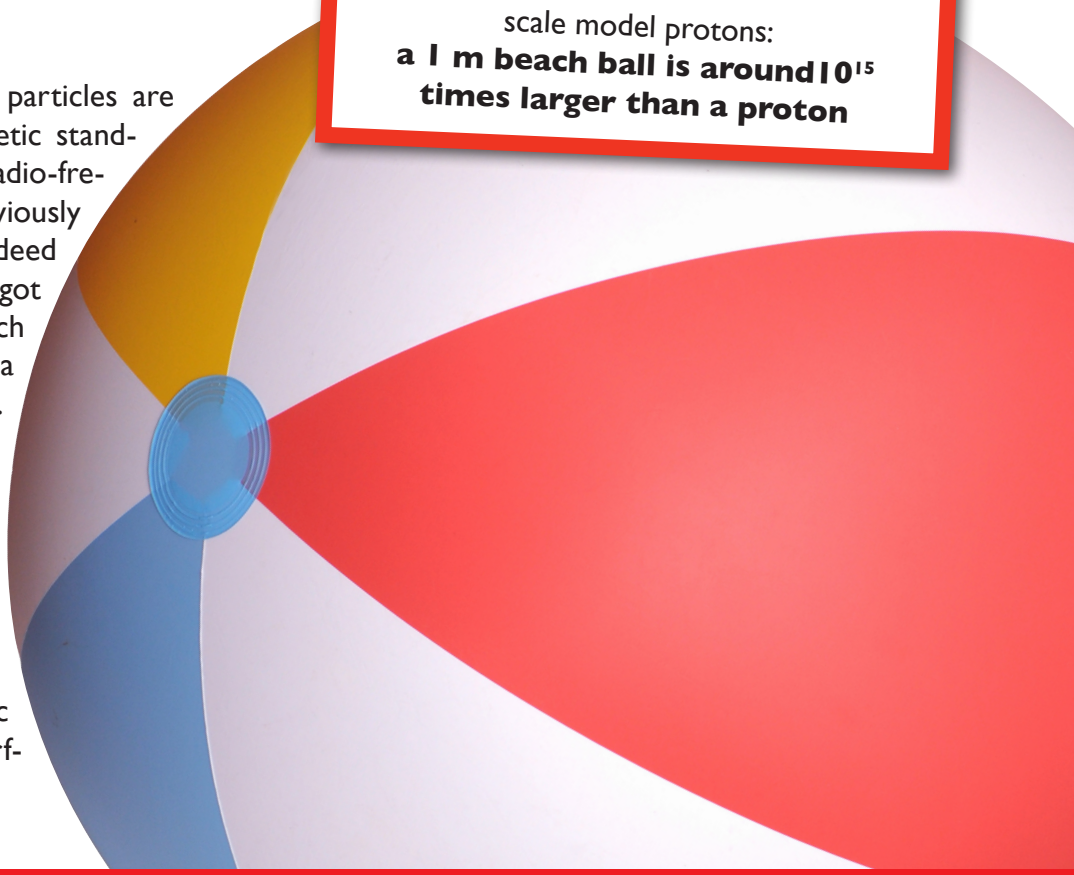

Cathode ray tube

The cathode ray tube was a scientific curiosity discovered in the late 19th century, and a mainstay of display technology in the late 20th. We now know that the mysterious 'cathode rays' are in fact electrons—and we can use magnets to bend their path.

Apparatus

This experiment obviously requires a cathode ray tube filled with gas which glows when electrons hit it. The ideal CRT is enclosed by Helmholtz coils to allow a varying magnetic field to be applied. In the absence of Helmholtz coils, a strong neodymium magnet should suffice to bend the electron beam.

In addition to a cathode ray tube, you'll probably need a sensitive camera to show your audience the results of this experiment. The beams of electrons are too dim for anything except a very small audience to see directly, and are something of a challenge for video equipment too! A camera with a night mode, or manual control over gain (or ISO) and shutter speed will probably be necessary.

If you've not got a cathode ray tube, an old CRT TV or computer monitor and a strong magnet will provide a more qualitative version of this demo.

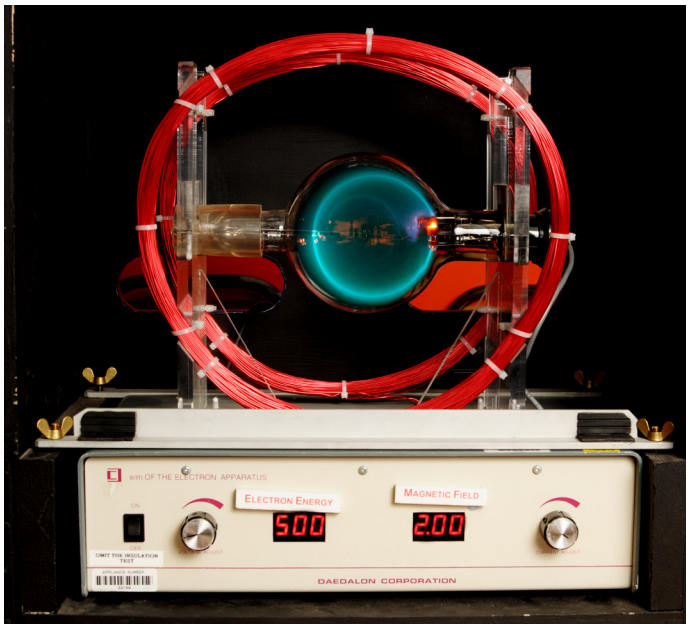

A cathode-ray tube with controls for both the electron energy and the magnetic field applied. The blue-green glowing trail in the gas in the globe shows the path of the electrons.

The demonstrations

Cathode ray tube

1. Dim the lights and turn on the camera if you're using one.
2. Turn up the energy of the electron beam until the gas inside the globe is clearly glowing.
3. If your CRT doesn't have Helmholtz coils, simply wave the neodymium magnet near the CRT to show the beam bending. You may need to do this quite slowly if the camera is set to a low frame rate to increase its low light sensitivity.
4. If your CRT does have Helmholtz coils, turn up the current in them until the beam bends.
5. Having curved the path of the beam, turn up the energy further and show that the curvature decreases with increasing electron energy.
6. Apply a higher magnetic field to demonstrate that the curvature can again be increased by increasing the magnetic field strength.

Vital statistics

speed of an electron accelerated through 1 V:
600 km/s

strength of LHC bending magnets:
8.36 T

CRT TV/monitor + magnet

1. Get an image on the television or computer screen. If it's a computer screen simply plugging it into a laptop should work. For a TV, many

camcorders and digital stills cameras will have an S-video, component or composite connection; older camcorders may have these directly, but newer camcorders or digital cameras may have a bespoke cable which plugs into a mini-USB or similar jack on the camera and feeds out to multiple types of connector for insertion into the TV. A relatively still, bright image or video makes the effect we're about to observe easier to distinguish.

2. Put the strong magnet near to the TV screen. The image will warp, and sweeping trails of colour will appear.

3. If the distortion and colours remain after taking the magnet away from the TV, turning it off and on again should force the TV to 'degauss' which will fix the problem—this is signified by the distinctive clunk which often accompanies a CRT turning on. Sometimes, often after repeated cycling, the TV will fail to degauss. In this case, turn it off, leave it for a short period, and turn it on again.

How it works

The key here is that magnetic fields will bend the path of a moving charged particle, and we can make use of this effect to control a beam. Crucially for the Accelerate! recipe, you need a larger magnetic field to bend a faster-moving particle.

In the cathode ray tube, electrons are ejected from the cathode and accelerated through a voltage, gaining some 600 km/s for every volt they are accelerated through. Some of these fast-moving electrons crash into the gas inside the tube, causing it to glow, which allows us to see the path of the beam. Helmholtz coils can then be used to apply a quantifiable magnetic field by passing a known current through them.

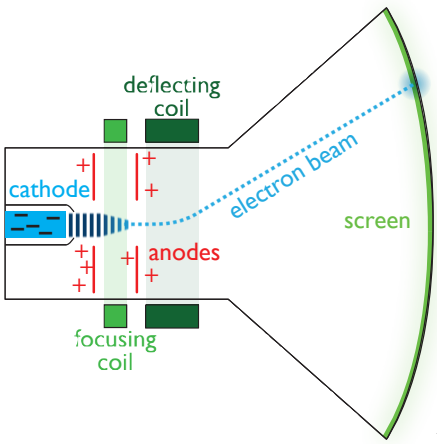

A schematic cross-section of a cathode-ray tube television or computer monitor. The cathode creates the electrons, and anodes accelerate them, and the focusing and deflecting coils use magnetic fields to control the beam to the desired point on the screen, where electron impact causes coloured phosphors to glow and an image to form.

A magnetic field will cause a force to act on the electrons which is perpendicular to both their direction of travel and the magnetic field. This causes a charged particle in a magnetic field to follow a circular path. The faster the motion of the particle, the larger the circle traced out for a given field or, conversely, the larger the field needed for a given radius of curvature of the beam. Making this quantitative point is impossible without control over both particle energy and magnetic field, so this will need to be stated if your demo doesn't have both of these.

In the case of the CRT TV, the paths of the electrons are distorted by the magnet being brought near the screen. The picture on the screen is dependent on the electrons precisely hitting phosphors on the back of the screen, which emit different colours of light when impacted. The electrons are thus forced to land in the wrong place, causing the distortion of the image and the psychedelic colours.

A CRT television first without and then with a magnet brought near to it. The patterns of colour are related to the magnetic field lines of the large circular magnet brought near to the TV: its poles are on its faces.

Electrical resistance

Electrical resistance can be considered rather dryly as the R in $V = IR$. However, we can show its effects more spectacularly as a shower of sparks when wire wool is placed over the terminals of a battery.

Apparatus

5 × battery
some wire wool
1 × heatproof mat

Safety

Safety specs should be worn for this demo because it creates very hot sparks. It is largely smokeless, but do be aware of any smoke alarms, heat detectors etc.

Vital statistics

melting point of iron:
1535°C

current carried in LHC magnets:
11,850 A

The demonstration

1. Place a clump of wire wool across the terminals of a suitable battery.
2. Small, bright sparks will start crawling along the length of some of the wires. When this starts happening, take the wire wool away from the battery and wave it through the air to fan the flames. If you just leave it on the battery, the results may be unspectacular!
3. Take care that the sparks don't land on anything which may catch fire or melt, including flooring—use of a heatproof mat is advised.

How it works

This demonstration shows the effect of electrical resistance. When an electric current flows through an ordinary conductor, the electrons don't flow smoothly like water through a pipe. They bump into the material they're flowing through, which means that they lose energy to the material, which therefore heats up. This phenomenon is observed in a rather more controlled fashion in electric kettles, hobs and heaters, when resistive heating warms your water, your dinner or the air in a room.

This poses a fundamental limit on how much electric current can be passed through a conventional conductor before it heats up to a dangerous degree and melts or catches fire. This effect is used in fuses to cut off the flow of electricity if a certain threshold current is exceeded. However, when designing a high-field electromagnet for a particle accelerator, this is a design limitation, and motivates the use of superconducting magnets.

There are two ways to work out just how hot this wire wool gets. The colour a hot material glows can often give a clue as to its temperature. Objects start giving off a significant amount of visible light when they exceed 500°C, becoming red hot. The wire wool glows white in this experiment, indicating temperatures in excess of 1500°C, further confirmed by the fact that the metal will sometimes melt, cooling to form small globules. Iron melts at 1535°C, so this stuff really is very hot!

Liquid nitrogen

Liquid nitrogen, boiling furiously despite being hundreds of degrees below freezing, is thoroughly captivating. You can make a whole show with this stuff, so *Accelerate!* concentrates on the safety and engineering challenges of using cryogenic liquids.

Apparatus

A few litres of liquid nitrogen
1 × bunch of flowers
1 × small bowl
Warm water
Washing up liquid
1 × glass bottle
1 × balloon
1 × length of rubber tubing

Vital statistics

boiling point of nitrogen:
−196°C (77 K)

expansion from liquid to gas:
600 to 700 times

The demonstrations

The first two of these are done by audience volunteers, so get a couple down and kit them out with safety visors and cryogenic-safe non-porous gloves.

Smashing flowers

1. Give the first volunteer a bunch of flowers and encourage them to dunk them into a bucket of liquid nitrogen.
2. When the boiling stops, get them to pull the flowers out and smash them (hard!) onto a nearby surface.

Nitrogen bubbles

1. Give the second volunteer a polystyrene cup filled with nitrogen, and get them to throw it all into a bowl containing warm, soapy water.
2. Watch as a cloud of steam clears to leave a huge mass of bubbles.

Filling a balloon

1. Place a small volume of liquid nitrogen in the bottom of a glass bottle.
2. Stretch a balloon over the neck of the bottle. The balloon will fill up with nitrogen gas, probably enough to explode it, or at least make it fly off comically around the room.

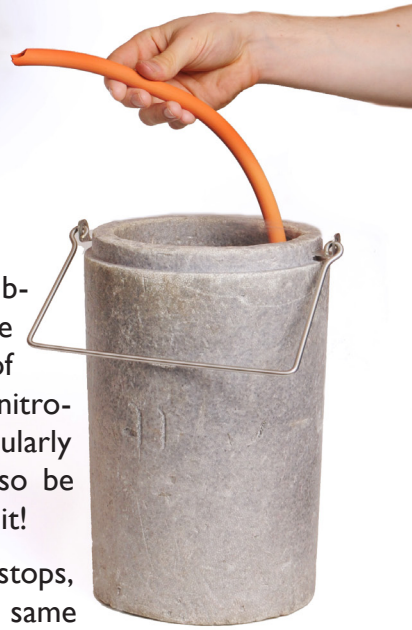

Smashing rubber

1. Take a length of rubber tube and place one end in a bucket of liquid nitrogen. Liquid nitrogen will spray spectacularly out of the other end, so be careful where you point it!
2. When the spraying stops, the tube will be at the same temperature as the nitrogen. Remove it, and tap it on a nearby surface to demonstrate that its properties have changed somewhat…
3. Smash it with a hammer!

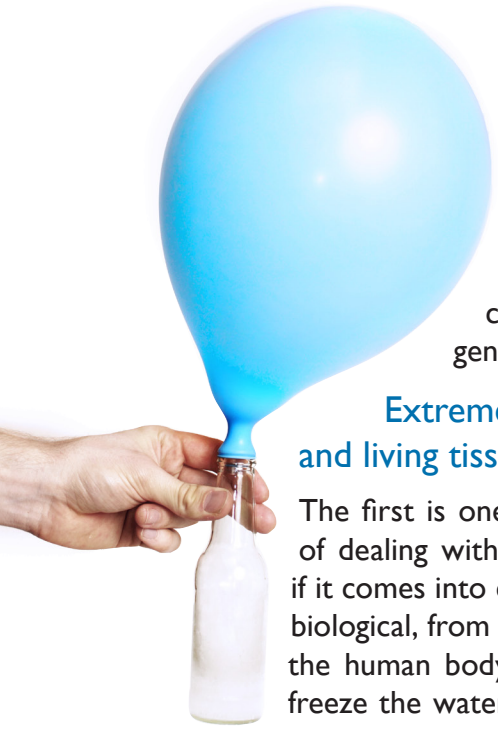

How it works

These four demonstrations are centred on three basic concepts relating to liquid nitrogen, and cryogenic liquids in general.

Extreme cold and living tissue

The first is one of the safety issues of dealing with the extremely cold: if it comes into contact with anything biological, from a bunch of flowers to the human body, it can very rapidly freeze the water inside and between the cells, thus killing them and also making the tissue very brittle. The flowers can be smashed: since you would likely have the restraint not to do the same to your finger, the more likely consequences are severe frostbite. That's why we have to be careful with these cryogens, wearing safety gear when handling them but, more importantly, trying to make sure they don't come into contact with people at all.

Liquid–gas expansion

The second is the massive expansion as you change a liquid into a gas. Liquid nitrogen has a density a little lower than water, but when it vapourises that density immediately drops several-hundredfold. Then, as it warms to room temperature, it expands by a further factor of three or so (its temperature changes from 77 K to 300 K, and we know that ideal gases' volumes are proportional to their temperatures). The end result is an expansion of 600 to 700 times as liquid nitrogen is warmed. That's why a huge, foaming tower of bubbles can be created from just a cupful of nitrogen and, slightly more quantitatively, why a whole balloon can be more than filled with a tiny amount of nitrogen in the bottom of a bottle.

The engineering implications of this are that any heat leak in a piece of cryogenic apparatus can be catastrophic: the gas produced from a large volume of liquid cryogen can be enormous, and in an enclosed space without appropriate safety valves, could cause an explosion!

Material properties at low temperature

The final point is that materials' properties can be radically altered by cooling to nitrogen temperatures: in the case of the flower, that was simply water freezing, but rubber remains solid whilst changing from flexible to rather glassy. This means we have to take care when designing these machines that the materials we use will work as expected at low temperatures. For example, using rubber seals on a machine containing liquid nitrogen would obviously be a bad idea because the rubber becomes very brittle and might shatter, and could no longer contain the liquid.

Cryogenics in the LHC

In the LHC, the extreme cold is required to keep the superconducting magnets working. In fact, the superconductor used wouldn't function at liquid nitrogen temperatures, and instead liquid helium is used.

Helium boils at just 4.2 K, or around −269°C. Further, the pressure above the liquid is lowered by pumping and this reduces its temperature even more, to a rather chilly 1.9 K.

The 'cosmic microwave background', an electromagnetic echo from the Big Bang which permeates the Universe, is at 2.7 K, which gives rise to the claim that the inside of the LHC is 'colder than outer space'.

Safety

These demonstrations also illustrate why safety is of paramount importance when dealing with nitrogen in a science show. Always be sure to wear appropriate clothing and use safety equipment when handling it. The greatest risk from liquid nitrogen is probably asphyxiation, resulting from a large volume of nitrogen gas displacing oxygen in the room. Be sure to conduct a risk assessment in conjunction with your local safety office before using or transporting liquid nitrogen.

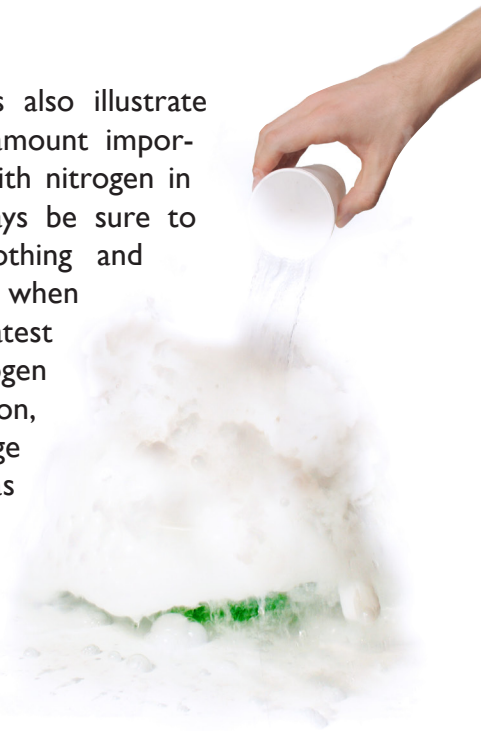

Superconductivity

The sight of a chunk of superconductor floating serenely over some super-strong magnets is a genuine wow moment in the show. These amazing materials are used to create the enormous magnetic fields used to control the beam in the LHC.

Apparatus

5 × neodymium-iron-boron magnets (their poles are on their faces: arrange them as shown)
1 × piece of high-temperature superconductor
1 × polystyrene cup, 2/3 full of liquid nitrogen
1 × plastic tweezers

The demonstration

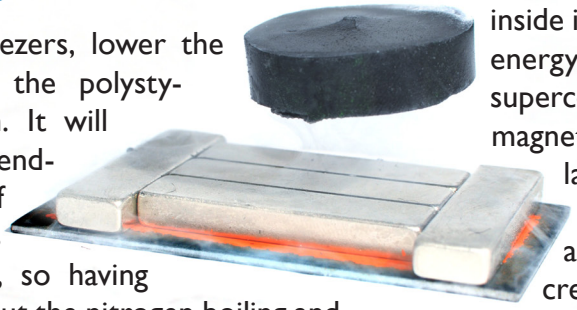

1. Using the plastic tweezers, lower the superconductor into the polystyrene cup of nitrogen. It will bubble fiercely. Depending on your piece of superconductor, this may take some time, so having some patter ready about the nitrogen boiling and sucking the heat out of the superconductor is advisable.
2. When the boiling stops, pull the superconductor out of the nitrogen again using the tweezers, and lower it vertically onto the five magnets. When you release the tweezers, it will just continue to float there!
3. When the superconductor warms above its critical temperature, it will fall onto the magnets. Place it back into the nitrogen with the tweezers. Being so cold, it should cool down rather more quickly this time. Repeat!

How it works

Superconductors are materials which display very unusual properties when cooled down below their critical temperature: chief amongst them, as implied by their name, they suddenly lose all electrical resistance.

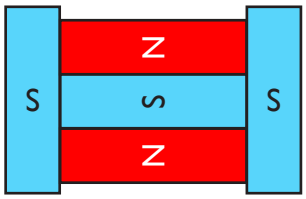

The arrangement of the magnets to allow stable superconducting levitation. Note that absolute polarity is not important!

Related to this vanishing resistance is the effect on display here: the Meissner effect, where a superconductor excludes all magnetic flux from inside its volume. This is because it costs no energy to induce a supercurrent inside the superconductor and, when brought near a magnet, that current will be set up by Lenz's law to oppose the change causing it. Consequently, the superconductor acts something like a magnetic mirror, creating an image of the same magnetic pole as it's brought near, and repelling it. This allows it to defy gravity, and levitate.

The reason we need multiple magnets is that when levitated over a single magnet, the superconductor is in an unstable equilibrium: the more complex configuration of fields over the arrangement of magnets suggested means that there's a stable position over them, and the superconductor will even spring back if displaced slightly.

This effect isn't exactly what we use in the LHC, though it does implicitly demonstrate the zero electrical resistance which is used. It's also not this type of superconductor which is used in the LHC. This demo relies on so-called 'high-temperature' superconductors which can superconduct at liquid nitrogen temperatures. These superconductors are brittle, layered materials which don't lend itself to extrusion into long wires so niobium–titanium, a superconducting alloy, is used instead, and cooled by liquid helium to a temperature of just 1.9 K (−271°C).

Beach ball collider

How do the different types of particle collision differ? We learn the difference between particle–particle collisions of the sort used in the LHC and fixed-target collisions, by throwing more beach balls around…

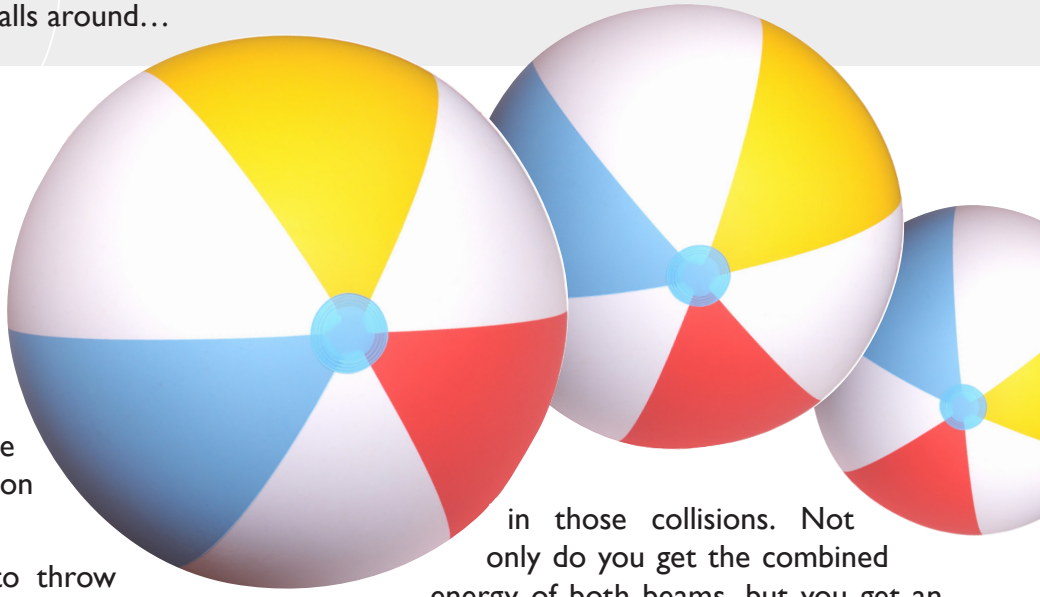

Apparatus

n × beach balls
$n/2$ × electric pumps

The demonstrations

First, we try particle–particle, or centre-of-mass collisions:

1. Place half the beach balls on one side of the audience and half on the other.
2. Get the audience members to throw the balls at a point somewhere in the middle of the room, and see how many (or few!) collide.

Second, we try fixed-target collisions:

1. Pick a fixed target. It's most fun if this is a person, so a presenter, or a member of the audience if they're up for it!
2. Distribute the beach balls around the room and get audience members to throw them at the fixed target.

What it shows

Of the two kinds of particle collision, smashing particles into a fixed target is clearly simpler. In fact, if you think colliding beach balls thrown from opposite sides of the room is hard, try colliding two beams of protons, each the width of a human hair, moving at nearly the speed of light, which have travelled in opposite directions around the 27 km ring of the LHC! That's exactly what happens at the collision points around the ring of the Large Hadron Collider.

The reason we bother with all this precision is because it gives us much more energy to play with in those collisions. Not only do you get the combined energy of both beams, but you get an additional bonus from conservation of momentum. The reason that particle–particle collisions are also known as centre-of-mass collisions is because the lab frame is also the centre-of-mass frame, and it's possible for the products of a collision to be net stationary afterwards. (It's actually a little less likely in the LHC because protons are composite particles and each quark carries a slightly different share of the momentum…but that's another story.) In a fixed-target collision, the products need to be zooming off in the direction of the original beam's momentum, rather like when you smash the cue-ball into the top of the triangle during a break in snooker, and that means that a lot of energy can't be used to make new particles and find new physics, but is 'wasted' in the form of the products' kinetic energy.

Alternatives

If you're on tour and want a more compact solution, or simply want something which takes a bit less time to inflate, these particle collision demos work just as well with smaller balls.

Cloud chamber

A cloud chamber makes the invisible visible, allowing us to see delicate, wispy proof that there are tiny particles whose story starts in outer space shooting through all of us, every minute of every day.

Apparatus

1 × cloud chamber
Propanol (aka isopropyl alcohol or IPA)
Dry ice (solid carbon dioxide)

The demonstration

The cloud chamber is prepared and placed down. Moments later, wispy streaks of cloud appear, seemingly spontaneously, inside the chamber. These tiny clouds show the path of charged particles through the chamber—and, since there are no obvious sources of charged particles around, they're evidently natural and omnipresent…

How it works

The base of the cloud chamber is filled with dry ice, and an absorbent material near the top thoroughly soaked in propanol. Propanol is quite volatile, and so forms a vapour at the top of the chamber. As the vapour falls, it cools rapidly due to the dry ice and the air becomes 'supersaturated': the propanol really wants to condense, but there is nothing for it to condense onto. Charged particles passing through the chamber cause the propanol molecules to gain an electric polarisation, and be drawn towards those particles, and one-another. This provides the impetus for them to condense into tiny liquid droplets in the chamber which show up as white streaks of cloud along the path of the particles.

Vital statistics

muon mass:
207× electron mass

muon flux at sea level:
10,000 muons/m²/minute

This demo is often done with a radioactive source, with alpha and beta particles causing propanol to condense, but it actually works in the absence of a source too, because of cosmic ray muons passing through the apparatus. The muons are produced high in the atmosphere by protons (the 'cosmic rays') smashing into the nuclei of gases. These produce a variety of daughter particles, but the only ones typically long-lived enough to make it to the Earth's surface are muons.

Muons are heavy electrons, and decay into an electon and a neutrino with a mean lifetime of 2.2 μs. This actually provides an interesting test of special relativity: muons are typically produced around 15 km up in the atmosphere, a distance which takes around 50 μs to traverse at the speed of light—over 20 muon lifetimes. Thus we'd expect barely any to make it! However, since they *are* travelling quite near the speed of light, time in their frame of reference is significantly dilated as seen by an observer on Earth, meaning that a significant fraction can, in fact, make it to the surface.

ATLAS & CMS

ATLAS and CMS are the two 'general-purpose' detectors at the LHC. They're looking for any new particles or unknown physics which the LHC's record-breakingly high energies might allow us to observe for the first time.

The Higgs boson

Probably the most famous goal of ATLAS and CMS is to spot the Higgs boson—a particle predicted independently and almost simultaneously by three groups of physicists, including Peter Higgs, which allows all other particles to have mass. It's a bizarre problem in modern physics that we can explain the properties of subatomic particles to incredible precision, but we're unable to explain why they have mass—in other words, what makes it hard to accelerate or decelerate them.

The Higgs boson is the only particle in the 'standard model' of particle physics which hasn't been observed. There are two types of particle in the model: fermions, which are 'stuff' (eg electrons, and the quarks which make up protons and neutrons); and bosons, which transmit forces (eg photons, which are particles of light, and transmit electromagnetic forces). The Higgs boson is one of the latter and, if it exists, would be responsible for giving particles mass. The theory says that the Universe is filled with a sticky soup of Higgs particles, and those particles which interact most strongly with the Higgs particles are bogged down by them. This is what we think of as mass—a tiny, light electron barely sees the Higgs particles, whilst a proton (which is 2,000 times heavier) is wading through a dense sea of them.

If we find it, it would be the last piece in the standard model's mathematical house of cards and, further new physics notwithstanding, our understanding of subatomic particles would be nearly complete. If we don't, it's perhaps more exciting for particle physicists, because suddenly a cornucopia of new potential theories is unleashed, ready for the LHC to test.

Vital statistics: ATLAS

A Toroidal **L**HC **A**pparatu**S**

diameter:	25 m
length:	46 m
mass:	7,000 tonnes

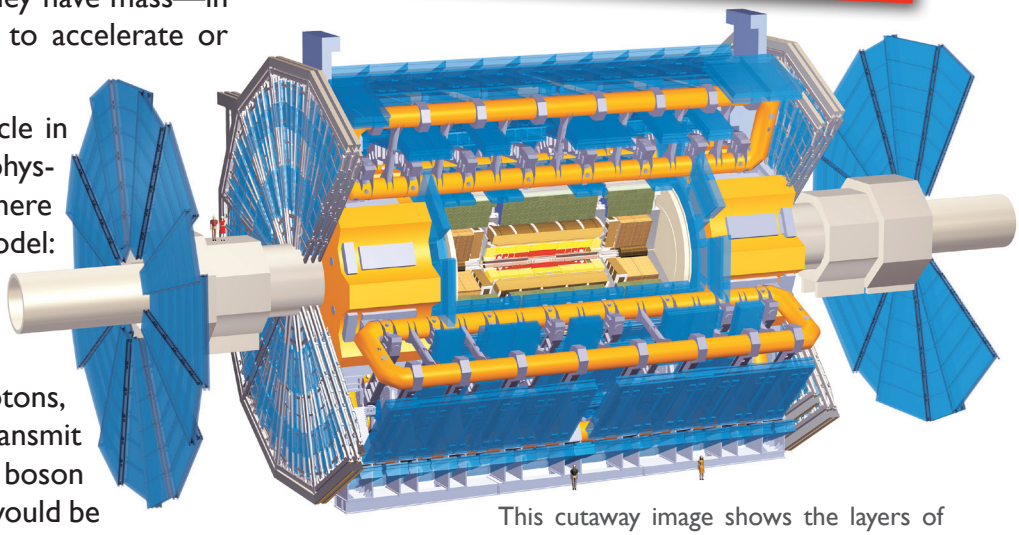

This cutaway image shows the layers of the ATLAS detector. Protons collide in the centre, and each layer is designed to detect different particle properties. Image © CERN.

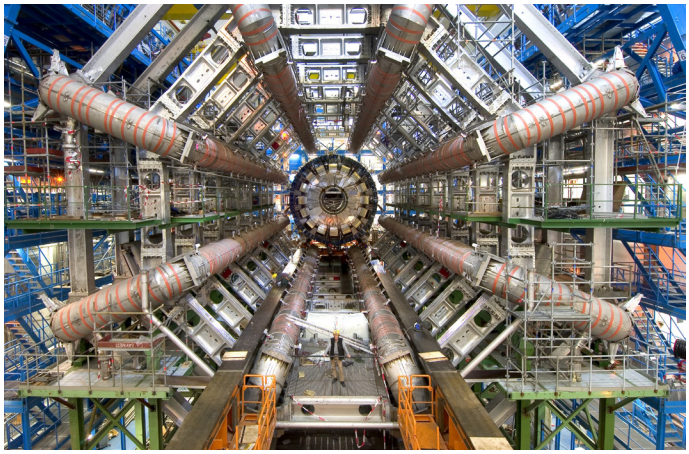

This iconic image shows the ATLAS detector cavern before the detector electronics were inserted. The person in a hard hat gives a sense of scale! Image © CERN.

Vital statistics: CMS

Compact **M**uon **S**olenoid

diameter:	15 m
length:	21.5 m
mass:	12,500 tonnes

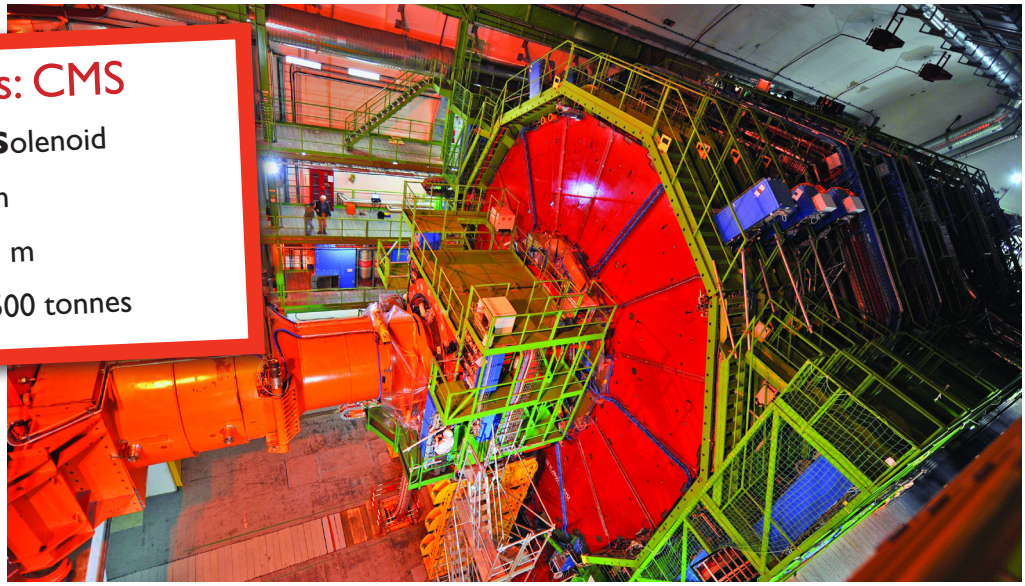

The completed CMS detector in its cavern. Protons travel through the orange pipe to the centre of the detector where the particles collide. Image © CERN.

Dark matter

Astronomers have known for some time that there's something missing in our understanding of the Universe. We can predict with extremely high precision how objects in the cosmos should orbit one-another: in the Solar System, for example, this works extremely well, and we can predict the motion of planets and moons with incredible accuracy. However, when we try to apply these laws to whole galaxies, something goes wrong: stars near the edges of the galaxies are rotating around galactic cores much more quickly than we'd expect given the matter we can see with telescopes. This effect scales up to clusters of galaxies too. We can measure their mass using 'graviational lensing', where gravity bends the light from further-away objects. The curvature of the light's path is related to the amount of mass bending it, and this allows us to deduce that there's much more there than just the visible mass in the clusters.

What this mass is composed of we just don't know, but it's possible that it's made up of heavy, weakly-interacting subatomic particles. The hope is that some of these dark matter particles will be produced in the LHC.

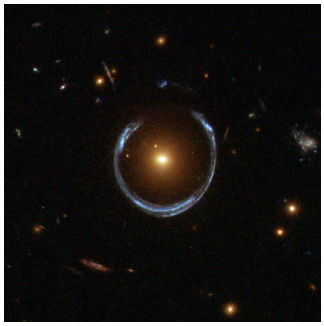

The near-complete circle in this image is the light from a distant galaxy bent around the nearby galaxy at its centre. Images like this allow us to estimate the mass of the galaxy bending the light, and suggest the existence of dark matter. Image courtesy of ESA/Hubble and NASA.

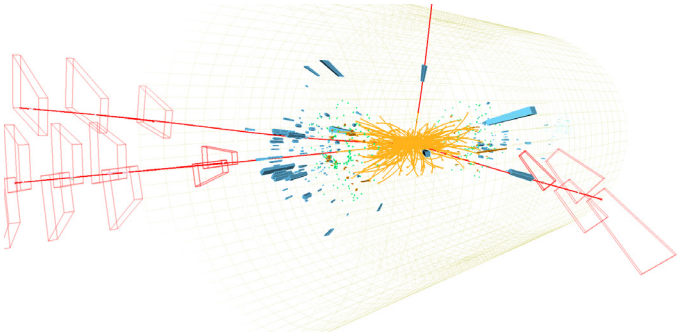

A candidate Higgs event recorded at CMS, showing one of its characteristic signatures, decay into four muons (the red lines). Image © CERN.

New particles, unknown physics

It's easy to forget when we have some very definite ideas of what we might find at the LHC that we might see some things which we've never thought of before. The LHC is the highest energy particle collider which has ever been built, and might well allow us to create particles which are completely new to science. ATLAS and CMS would be the likely places that these kinds of results would be unearthed.

Sensitive measurements of the energies of particles created in a smash allow us to work out the mass of a new particle created thanks to $E = mc^2$. Then, more subtle measurements of what kind of daughter particles follow and their speed and direction of flight will hopefully allow scientists to work out exactly what they've made and work out what, if any, effect this has on existing theories—or entirely new ones!

LHCb

LHCb is looking specifically for certain types of particle decays involving B-mesons: particles containing a beauty quark. These reactions will allow us to very precisely differentiate matter and antimatter—but why does that, er, matter?

Where has all the antimatter gone?

All particles have an antimatter twin with the opposite electrical charge. For example, the negatively-charged electron's antiparticle is called a positron, for obvious reasons. As the Universe cooled from being a searing ball of pure energy in the moments after the Big Bang, it's thought that an equal amount of matter and antimatter should have been created.

Matter and antimatter are not happy bedfellows: when one comes into contact with the other, they annihilate in a flash of gamma radiation. Consequently, it's pretty obvious that there isn't a lot of antimatter here on Earth, otherwise life would be rather more hazardous than it is, with things forever disappearing in huge flashes of energy!

In fact, it seems that the observable Universe is pretty much entirely matter (there are no huge regions of annihilating matter and antimatter between the stars) which begs the question: where has all the antimatter gone?

LHCb is looking for tiny differences in the way that matter and antimatter react, in an attempt to understand this mystery. Perhaps a little bit less antimatter was created in the first place, or perhaps it behaved a bit differently in subsequent reactions meaning that, when all the annihilating was over just after the birth of the Universe, a small excess of matter was left over which, ultimately, went on to form us.

It's looking at these tiny differences primarily using particles called B-mesons, which contain a beauty quark. Also known, less romantically, as a bottom quark, these are the second heaviest kind of quark known, a thousand times heavier than the up and down quarks which make up protons and neutrons.

Vital statistics: LHCb

LHC beauty

length:	20 m
mass:	4,500 tonnes
angle covered:	17° × 14°

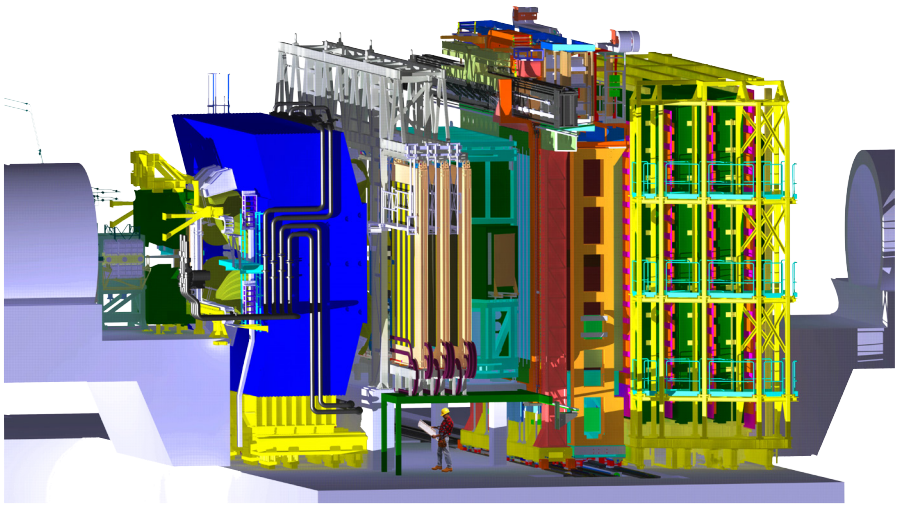

The layout of the LHCb detector, with a picture of a construction worker for scale. LHCb looks quite different to the other experiments in the LHC, in that it's not cylindrical. That's because the B-mesons which it's looking for tend not to fly out in all directions, but instead fly at small angles to the direction of the incoming protons. Image © CERN.

ALICE

The ALICE detector is optimised for examining collisions between lead ions rather than protons, to probe the quark–gluon plasma which, at a staggering two trillion degrees, filled the Universe moments after the Big Bang.

Vital statistics: ALICE

A Large **I**on **C**ollider **E**xperiment

diameter:	16 m
length:	26 m
mass:	10,000 tonnes

Part of the ALICE detector known as the time projection chamber. Image and ALICE logo © CERN.

What did the Universe look like just after the Big Bang?

The ALICE experiment is designed rather differently to the other detectors at the LHC. Rather than being optimised to examine collisions between individual protons, ALICE is used when the LHC is filled with lead ions, which contain 82 protons and about 125 neutrons each.

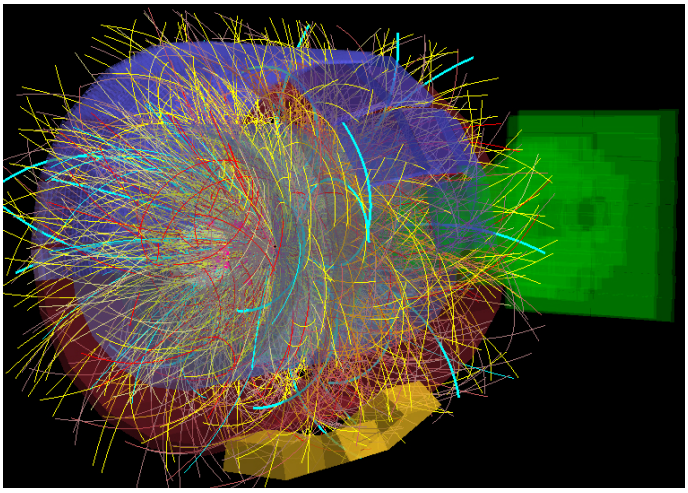

The chaotic aftermath of a lead-ion collision recorded in ALICE. Studying the tracks of these particles gives us an insight into what conditions were like in the quark–gluon plasma in the very early Universe. Image © CERN.

When these (relatively!) enormous ions smash together, they create a state of matter called quark–gluon plasma. Quarks are normally trapped inside protons and neutrons, stuck there very tightly by particles called gluons. In this plasma, quarks and gluons are all buzzing around completely free. Unsticking protons and neutrons requires an incredible amount of energy, meaning that the mixture must be incredibly hot: around two trillion degrees, or about 100,000 times hotter than the centre of the Sun! Temperatures like this only existed in nature a couple of millionths of a second after the Big Bang.

ALICE hopes to understand why quarks are never seen alone in the modern Universe, only ever assembled into particles, like protons and neutrons. The mechanism for this so-called 'confinement' is currently not fully understood. ALICE is also trying to work out why protons and neutrons weigh so much: if you add up the masses of the three quarks which make up the proton, you only get about 1% of the total mass. A deeper understanding of how quarks and gluons interact could explain where 99% of the mass of everything around us comes from.