

Inside the Large Hadron Collider

Inside the Large Hadron Collider

(by Robert Bee)



On 1st June, 2010, as part of Australian Astronomer Fred Watson's Stargazer II tour, we visited the most complex and expensive individual machine in the entire world – The Large Hadron Collider at CERN in Switzerland. I confess, as an avid reader of cosmology and the other end of the scale, particle physics, I was in a high state of anticipation, like a child having all his Christmases come at once. Or more prosaically, like a pig in mud.

I recognise that particle physics is not everyone's cup of tea. Okay, it's not most people's cup of tea. But even if you were of that mind, you could not go away, after a walking tour over the whole facility (excluding the 27 km of tunnel 100 meters below ground) with expert commentary from local physicists without having gained a sense of awe at the sheer complexity of the facility, the genius of engineering that went into designing and building it and the knowledge that, at that very moment, 100 metres beneath your feet, two beams of super high energy protons were rushing effectively at the speed of light to collide in titanic showers of exotic fundamental particles. Needless to say, I was 'impressed'.

When we arrived at the LHC, we were met by our personal guide for the day – Klaus Batzner. Klaus, it seems, is a 'legend' at CERN and the LHC in particular, having worked at CERN as a particle physicist for 30 years, then after retiring, as a volunteer guide for ten years. What Klaus didn't know about the LHC wasn't worth knowing. We were privileged to have him as our guide.



Klaus Batzner – our LHC guide

It would take a much longer article than this to convey, even at the simplest possible level, details of all the systems that combine to make this 'machine' work. There are so many of them, every one of which has to be operating at optimum levels. If any one sub-system falters, the whole LHC shuts down, potentially with catastrophic results. At such high energies and magnetic fields and low temperatures, there is no margin for error or failure. But I will attempt to give a summary overview. I hope you will be as impressed as I was, and still am.

Remember, it is this very machine which produced the recent (2012) much heralded breakthrough in knowledge of how our universe works – the discovery of the long illusive Higgs Boson, the 'giver of mass'. However, there is still much work for it to do.

What is the Large Hadron Collider?

Firstly, it is a large ‘collider of hadrons’, not a collider of ‘large hadrons’. The hadron is a class of atomic particle which are made up of quarks, the particle in question here being protons (made up of three quarks). The LHC is also designed to use lead nuclei which are much much heavier. Another description of the collider is ‘particle accelerator’. It accelerates atomic particles to incredible speeds so they can be made to collide at high energies. Large? Well, comprising a 27 km long tunnel (in a circle 8.5km diameter) running 100 metres beneath Switzerland and France near Geneva, the entire tunnel containing continuous super-conducting magnets, interspaced with four absolutely huge (by any standard) detectors, I think ‘large’ is a fair description.

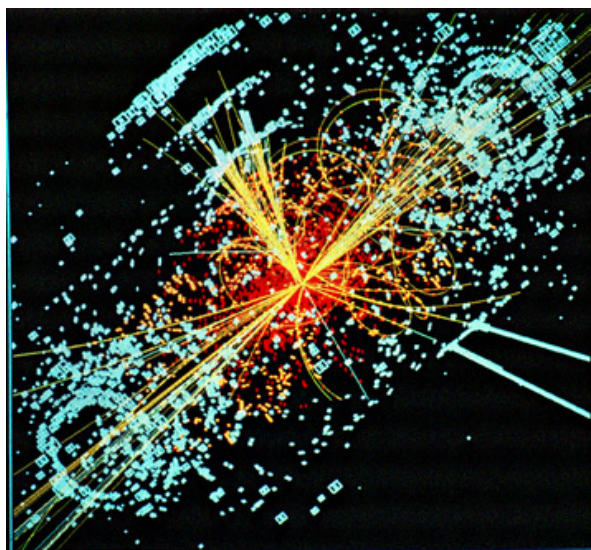


The path of the LHC tunnel between Lake Geneva and France’s Jura mountain range.

The LHC is a machine of superlatives. There is nothing ‘ordinary’ about it. In fact, it is a miracle it was ever built. A consortium of 20 nations agreed to spend around 8 billion euros in 1995 to build a fantastic machine for pure research purposes, to answer the biggest questions of science. Is there a Higgs boson? What is Dark Matter? (The question “what is dark energy” wasn’t even thought of at the time but can be addressed now.) Why were there more particles than anti-particles after the Big Bang? And many more. What was even more amazing, they went into it fully aware that there were, at that time, no known engineering and IT solutions to the problems they knew would arise. They would have to be solved on the way. Can you imagine that happening today? Yet they succeeded.

Large Hadron Collider - A Big Picture

In simple terms, two streams of ‘packets’ of protons travel near the speed of light in opposite directions in small 65mm diameter tubes around a 27 km-long circuit, held on course by superfluid helium cooled superconducting magnets. The proton packets are progressively accelerated to near light speed by sophisticated superconducting RF cavities at one point of the circuit. At four other circuit points there are huge particle detectors, each of a specific and unique design. Using magnets, the contra-rotating beams are brought to cross and collide at the centre of each detector, resulting in showers of fundamental, particles generated by the energy of the collisions. The data collected from all the particles from each of the collisions is collected in an unimaginably huge supercomputer system for future analysis. Here’s an example of a typical collision incident.

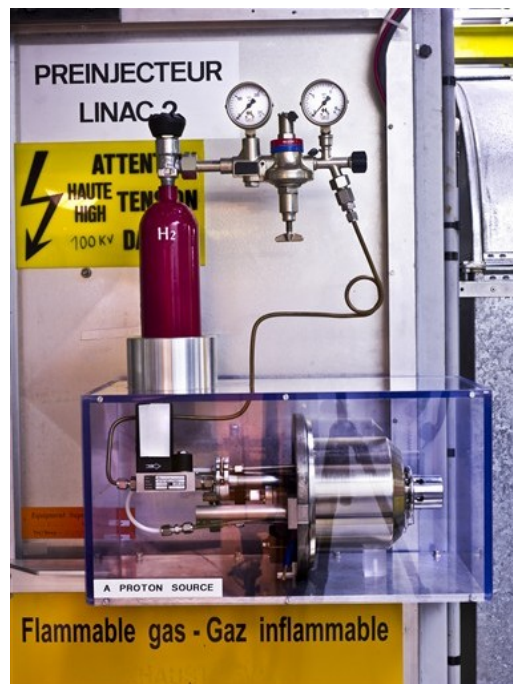


If that sounds complicated, keep in mind that there will be 30 million ‘crossings’ per second, each generating about 20 collisions. That’s around 600 million collisions per second. That’s a lot of data to analyse.

To help grasp the concept of what is going on and how complex the machine is, let’s break it down to some of its various components, or systems. We can’t cover all of them here but will look at the main ones. As you read the specifics of each system, I hope you will grasp the sense of awesomeness of the scientific and engineering achievement involved. This is no ordinary machine.

Proton Beams

There are two beams travelling in opposite directions in separate tubes about 65 mm diameter. Each beam consists of 2808 discrete packets of protons, each packet containing around 100 billion (10^{11}) protons. The packets trail each other by about 7.5 metres (or 25 microseconds). At near light speed, each packet will make 11,245 circuits each second. The process of starting the beams is, as you could imagine, quite complex. It starts from a very humble source of protons – a simple bottle of hydrogen gas.



Bottle of Hydrogen gas – source of the proton beams.

This bottle's pressure forces the hydrogen atoms out where an electric 'stripper' strips the electrons off the atoms and the protons are then channelled via various magnetic paths outside the main circuit, sped up and ultimately injected (by what they quaintly call 'kicker' magnets) into the main beam path. There, over a period of circuits taking 20 minutes, the position and stability of each packet being maintained at high precision by finely controlled feedback loops, the packets are sped up to near light speed. This gives each beam an awesome energy of 7 TeV, approximately 340 Megajoules. Put in lay terms, each beam will have as much energy as a family car travelling at 1,600 kph. At the same time, the energy stored in the magnets maintaining the beams (more about them later) is enough to melt 50 tonnes of copper. Impressive!

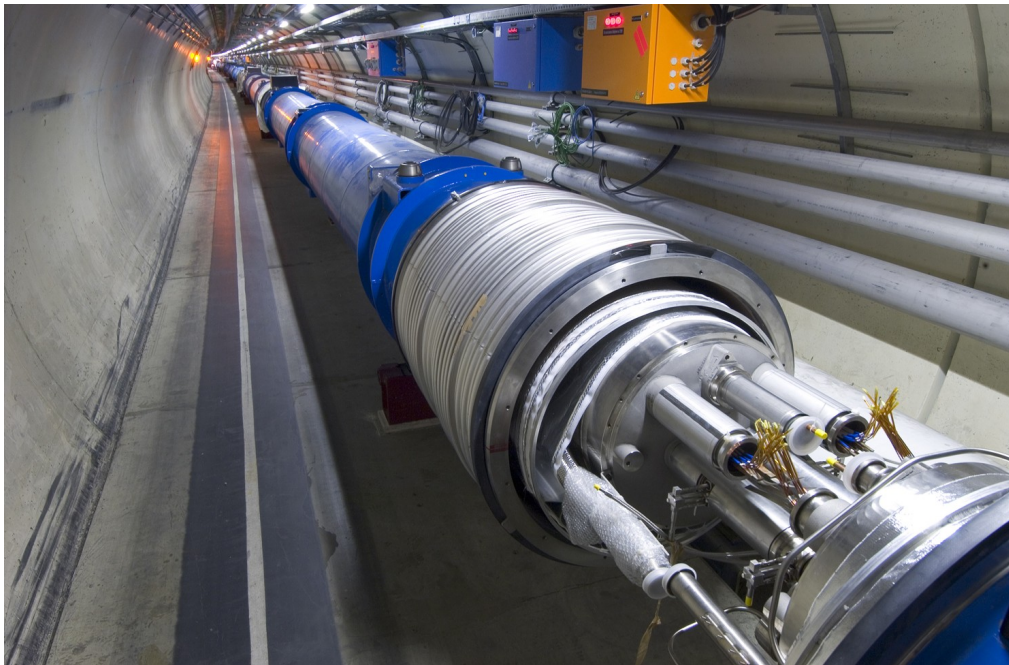
Once the beams are at this stage they can be kept in circulation (or 'stored') for up to 10 hours, during which time they have travelled a distance equivalent to Neptune and back. It is during this time when the experiments are conducted by colliding the beams at the various detectors around the beams' path.

With so many protons travelling at such high speed, obviously it is undesirable for them to encounter any stray particles in the beam tubes. This would cause proton scattering and reduce the lifetime of the beam. It can also have other more complex and undesirable consequences. For this reason, the tubes are evacuated to an extremely high level of vacuum, comparable to that in outer space. It is called Ultra High Vacuum (UHV) with values as low as 10^{-14} atmospheres, similar to the atmospheric pressure on the Moon.

Superconducting Magnets

Contrary to some expectations, most of the LHC does not accelerate the beams. Rather, via a huge number of magnets, it guides them from the exit of the accelerating point around the whole 27 km circuit back to the entry point. The focusing of the beam and gentle bending around the huge circle is achieved by vertical dipolar magnetic fields that act via the electromagnetic force.

An icon of the LHC is its amazing superconducting magnets.



Magnets surrounding the beam tubes in the tunnel

These magnets come in many forms – dipoles, quadrupoles, sextupoles, octupoles, decupoles and dodecupoles (but the vast majority are simple dipoles, covering 85% of the 27 km circuit) – and encircle the beam tubes for the entire 27km, except for inside the four detectors. There are 1800 large magnets in all in the LHC, 1232 of which are 15 metre-long dipoles, weighing over 30 tonnes and costing 2.5 million euros each. That's over 3 billion euros in dipole magnets alone. There are a further 8,000 smaller magnets used for corrections, fine tuning and flexibility of beam manipulation.

To achieve the huge magnetic fields required to keep the tunnel at 'only' 27 km long, extra high electric currents are required and to achieve those currents in the necessary small dimensions of the magnets, superconducting materials were necessary. In the LHC niobium-titanium (Nb-Ti) is the superconducting material of choice. In an engineering marvel, huge numbers of 1mm diameter wires of 6 micro-m dia Nb-Ti embedded in copper are woven into filaments, then the filaments into cables to carry the magnets' currents. In the magnets at the LHC, a current density of 400A/mm² was achieved, providing a magnetic field of 8.3 tesla – 150,000 times stronger than Earth's magnetic field and one of, if not the, strongest magnetic fields created in the world. Normal 'warm' magnets could achieve only a maximum of 2 tesla.

Liquid Helium

To have superconducting magnets, super-cold temperatures are required. Cryogenic technology. This is achieved at the LHC using liquid helium, its normal temperature of 5.7°K being further lowered to a temperature of just 1.9°K (that's just 1.9° above absolute zero) where liquid helium enters a 'superfluid' state, pushing the Nb-Ti superconductors to new extremes. The vast majority of the hardware around the 27 km tunnel (at least all the bits inside the external vacuum filled insulating sleeves) is kept at this frighteningly low temperature. The LHC has a 'cold mass' of 37,000 tonnes cooled to 1.9°K. Nothing on this scale had ever been attempted before. The refrigeration plant to achieve this beggars belief, requiring an all-new technology.

During the entire period of operation of the machine, around 80 tonnes of superfluid helium has to be maintained. But before the 'cold mass' is lowered to 1.9°K, it is pre-cooled by 10,000 tonnes of liquid nitrogen. With all this, one can only imagine the engineering challenges in designing, building and maintaining operation of such a super-cold machine. It is reasonable to claim that the interior of the LHC is one of the coldest places in the universe.



End view of magnet, with beam tubes (centre) and helium pipes.

Beam RF Accelerator

The proton beams are accelerated at only one point in the 27 km circuit, but this happens 11,245 times each second for each of the 2808 proton packets. Over the start up period of 20 minutes, this leads to their final near light speed status of 99.9999991% of c.

This acceleration is achieved through an ingenious radio frequency (RF) technology, described modestly by CERN as a 'triumph of modern engineering'. Copper 'cavities', as shown in the photo below, using electric fields oscillating at 400Mhz, create stable volumes of space to confine the individual packets of 1011 protons. Though too complex to explain here, the principle is that the alternating current creates an alternating electric field which, when the timing of the packet's arrival is right (as it is carefully designed to be), the positive protons get a 'kick' in the forward direction and thus gain acceleration, incrementally approaching the speed of light.



RF Cavity

The performance of the RF cavity is greatly improved when cooled to superconducting temperatures. For purely engineering reasons, best performance was achieved by using eight RF cavities in line, resulting in a superconducting enclosed arrangement as shown below. It is through this arrangement (one for each beam tube) that the protons pass to be given their acceleration 'kicks'.



RF module with 8 RF cavities in series.

The eight accelerating cavities for each beam have another beneficial feature. They shorten the proton packets and make them more dense and once the maximum energy has been reached the cavities keep the packets nice and tightly packed. Clever!

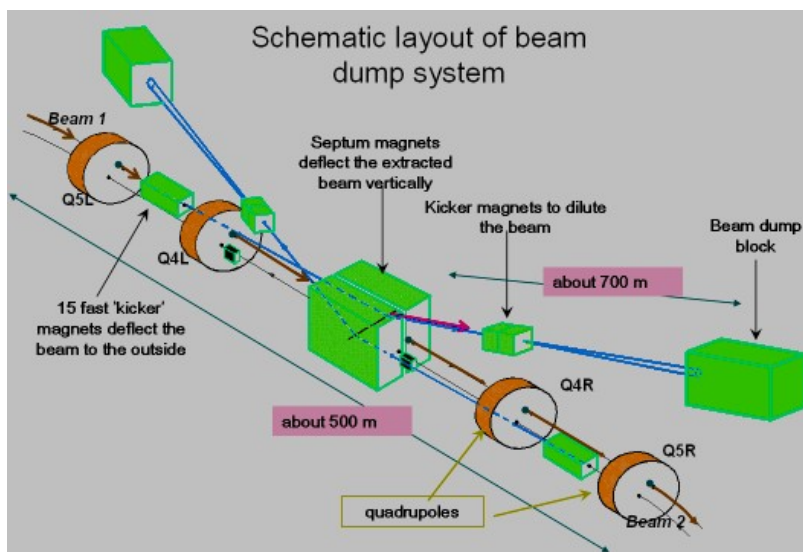
How Do I Stop This Thing? Beam Dump

We heard how the energy in the beams is that of a Toyota Camry at 1,600 kph, about 340 Megajoules. But when the experiments at the detectors are finished for the day – how do you turn then off? Where do all those protons and their energy go? I asked Klaus that question and, with a cheeky grin, he answered “Yes, we thought about that.”

There is a complete system for just that purpose. The Beam Dump. It is used in two circumstances. Firstly for the deliberate controlled ending of the beam after completion of a session. Secondly for the emergency dumping when something horrible goes wrong with the beam. The actual mechanism is both super-sophisticated and crude. That is, the control technology is like the sharpest laser scalpel while the conclusion is like a blunt axe.

Consider this: The time for one revolution of a beam is less than 0.0001 seconds. If it is stopped dead in that time (as it has to be) the instantaneous power dissipated is 4×10^{12} watts, about 100 times the combined output of all Australia's power generators. The area density of this power is also enormous as the dimensions of the face of the beam is less than a square millimetre. This enormous power density would vaporise anything placed in its path. The damage potential of any straying of the beam from its designated path would be unthinkable. So how do they manage it? How to dispose of the beam from the main path in the time of only one beam circuit? It took CERN over twenty years to work out an acceptable system.

In very simple terms, there is at a specific point in the circuit a pair of 750m-long offshoot tunnels, much like escape ramps on a steep road. One tunnel for each beam.

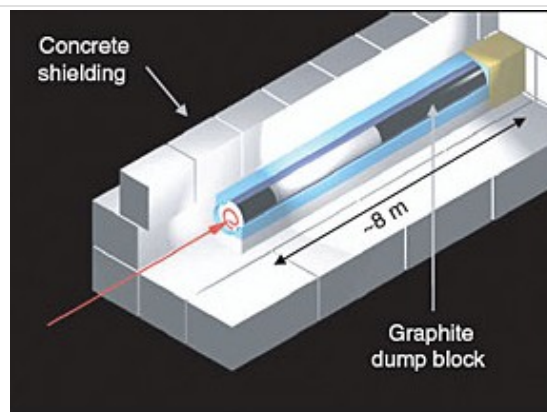


Schematic of Beam Dump System

At the point of departure from the true beam path to the slightly angled offshoot, there is a special magnet called the 'extraction kicker'. This kicker actually comprises 15 magnets. This is a magnet of truly special characteristics, enough to make an electrical engineer (such as myself) blanch. It has to move the beam horizontally by just 50mm, enough to move it just outside of the circulating beam aperture and into the high field aperture of the waiting Septum Magnet which remains permanently powered. The septum provides a strong vertical deflection, which after several hundred metres is enough to lift the beam above the superconducting magnets of the main LH ring and into the extraction line. While in this line, more kicker magnets 'dilute' the beam, effectively making an 'e' shape that dilutes the energy density on the target by a factor of 50. Eventually the beam arrives at an absorber block made of graphite, 700mm diameter and 7.7 metres long, shrunk fitted into a stainless steel jacket. This block is surrounded by 900 tonnes of shielding. That's the 'blunt axe'. Sounds simple but actually quite complicated, with many monitoring and fail-safe technologies woven into it.



Beam Dump tunnel at LHC



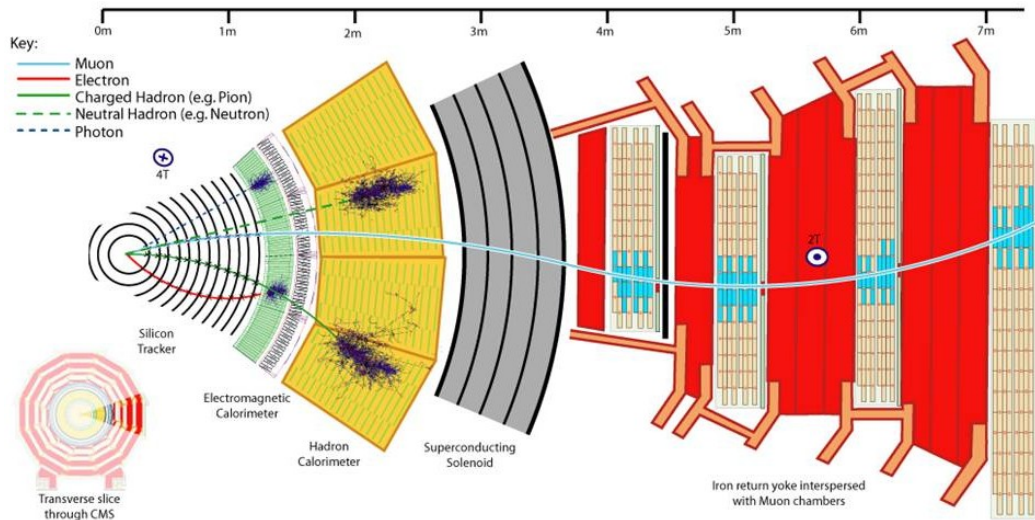
Beam Dump design

I said the kicker magnets were special. With a single turn coil, a current of 18,500 amps (under the pressure of 30,000 volts) is needed to reach the required magnetic field. This would be difficult enough for a normal magnet given time to build up its field over time, but for the LHC kickers, the full field of 0.34 tesla has to go from zero to full strength in less than 3 millionths of a second and stay at this value for at least one full turn of the accelerator to ensure all particles in the beam are extracted. This was an engineering triumph to achieve.

The Experiment Detectors

What is it all for? Ultimately, all of the forgoing infrastructure is for the sole purpose of arranging near light speed high energy packets of protons to collide on cue in a controlled area where super-detectors can 'detect' and record the myriad of particles that result so that knowledge of our universe can advance. This is done at the four detectors ATLAS, ALICE, CMS and LHCb.

These detectors, built in the path of and all around the crossing beams, are by any measure, gigantic. Space prevents a proper description of the nature of each but they can simply be described as containing multiple layers, each layer designed to perform a specific task and together they allow identification and precise measurement of the energies of all the particles produced in the proton-proton collisions. These layers are arranged like a cylindrical onion around the collision point, with further layers at each end of the cylinder to ensure all particles are captured. The diagram below shows the complexity of layers in the CMS Detector. The other detectors, though different in layout, follow similar principles.

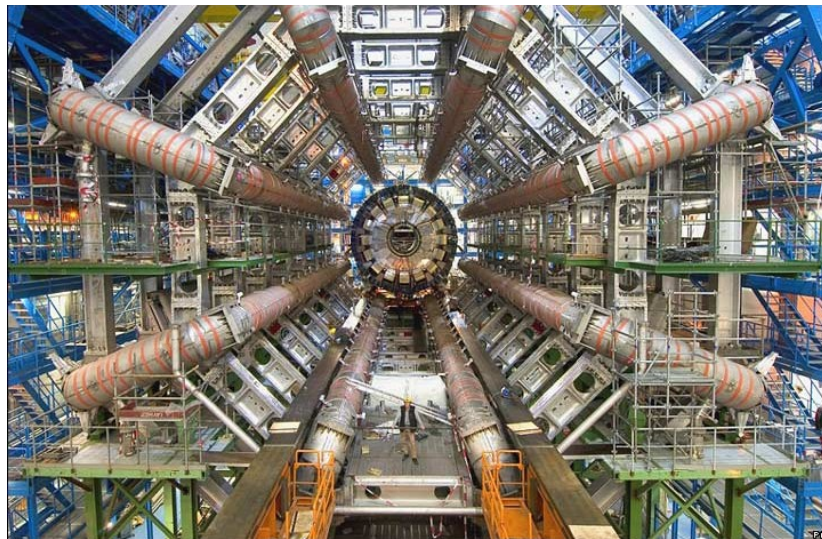


Slice of CMS Detector layers

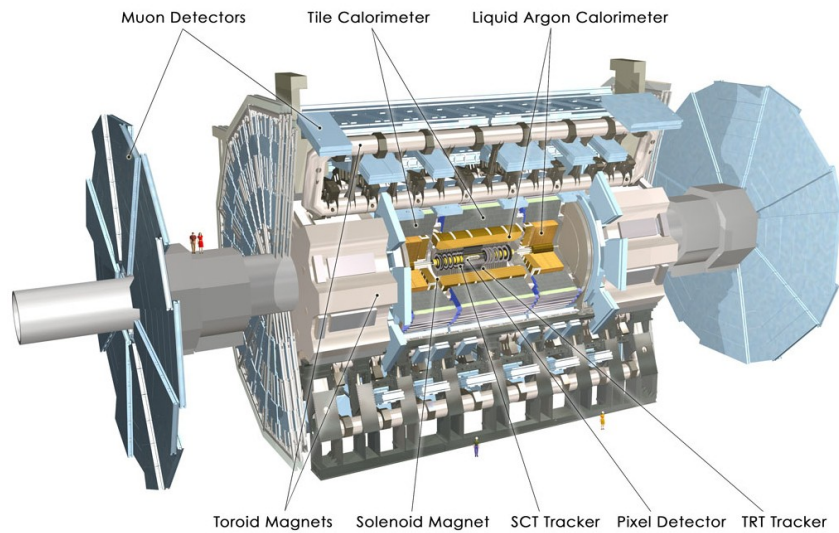
Here is a brief description of each detector:

ATLAS

One of the more excruciating acronyms, ATLAS stands for **A** Toroidal **L**HC **A**pparatu**S**. It is a scientific project in its own right, with around 2,500 scientists assigned to its work from 169 institutes and universities from around the world. The toroids in its name are shown in the following picture taken during its assembly 100 metres below ground level. Note the man in lower centre for scale.



As you can see from the diagram below, ATLAS is of immense proportions, worthy of its name. Note the two human figures for scale.



On the Stargazer II tour, we were fortunate to be able to observe the ATLAS control room, located comfortably at ground room. At the time, sadly, there was no beam running. Here you can see the physicists at their terminals and also the Stargazer group being 'lectured' by our guide Klaus.



ATLAS Control Room



Stargazer group and Klaus

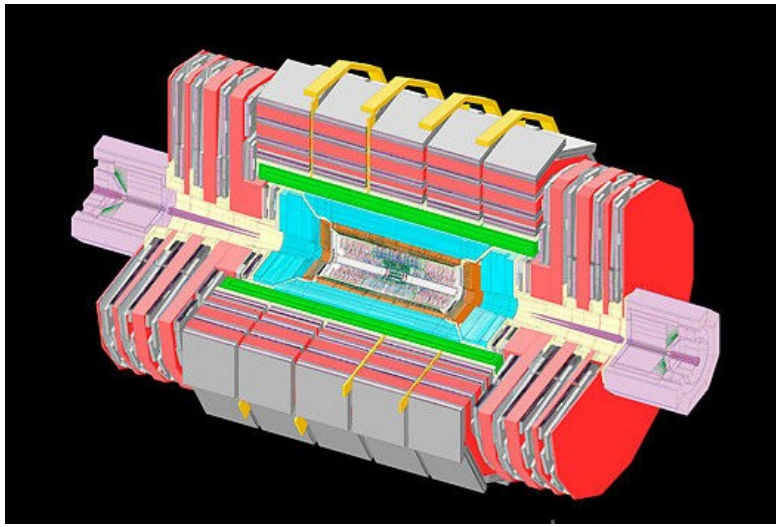
ATLAS was developed as a general-purpose detector to fully cover the rich physics potential of the LHC. This includes the search for the Higgs boson, dark matter, supersymmetric particles and extra dimensions. Here are a few of the key features of ATLAS:

- * It is 45 metres long and 25 metres high;
- * It weighs 7,000 tonnes;
- * With eight 25-metre-long superconducting coils and closed at each end by two end-cap toroids magnets looking like giant cog wheels, it is the largest toroid magnet system ever built;
- * It can store 1,600 Megajoules of magnetic energy, enough to lift the Eiffel Tower 16 metres off the ground;
- * There are over 100 million sensors to detect the collisions particles, able to locate a track to within 0.02 mm.

It was ATLAS, with another detector independently, that identified the potential Higgs boson announced on 12th July 2012. Possibly one tick on the board.

CMS - The Compact Muon Solenoid

This is another of the general-purpose experiments built to explore the physics at the TeV scale. It also involves over 2,500 scientists and engineers from over 180 institutes and 38 countries. A huge project! As its name suggests, the CMS was essentially designed for the accurate detection of muons, particles that are like heavy electrons using its unique solenoid magnet. And it may be 'compact' but it is still the heaviest of the four LHC detectors, weighing in at 12,500 tonnes.

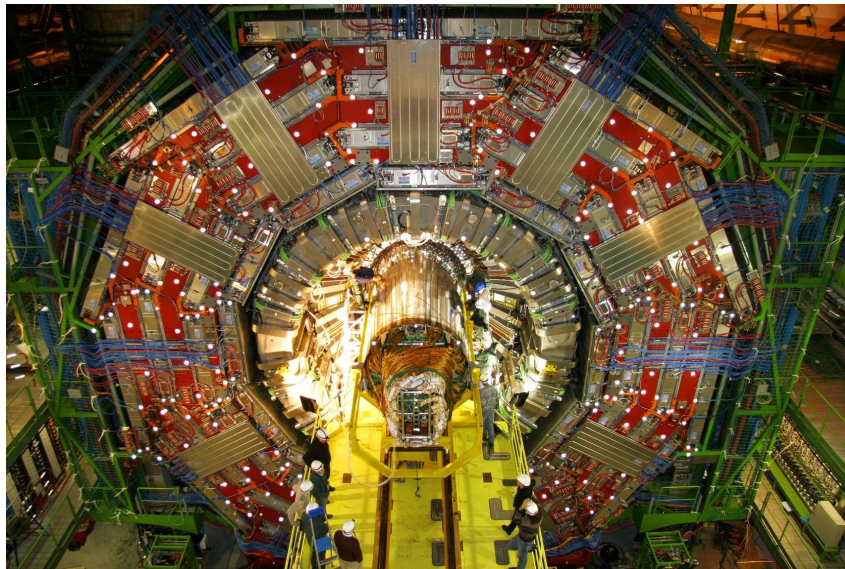


Cutaway of the CMS detector

CMS has the world's largest solenoid magnet, 13 metres long and 6 metres in diameter. Overall, the CMS detector is 21 metres long and 15 metres in diameter.

Its research targets using its ability to detect muons include the Higgs Boson – a 'discovery' it recently shared with the ATLAS detector – as well as dark matter, anti-matter and other exotic cosmological conundrums.

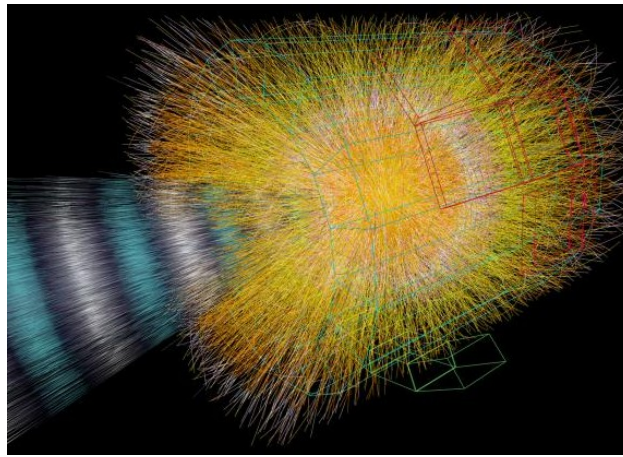
While hidden well beneath the ground level, the CMS's technical complexity is suggested by the brilliant colouring of its many thousands of special components. It simply exudes sophistication.



A section of CMS during installation

ALICE - A Large Ion Collider Experiment

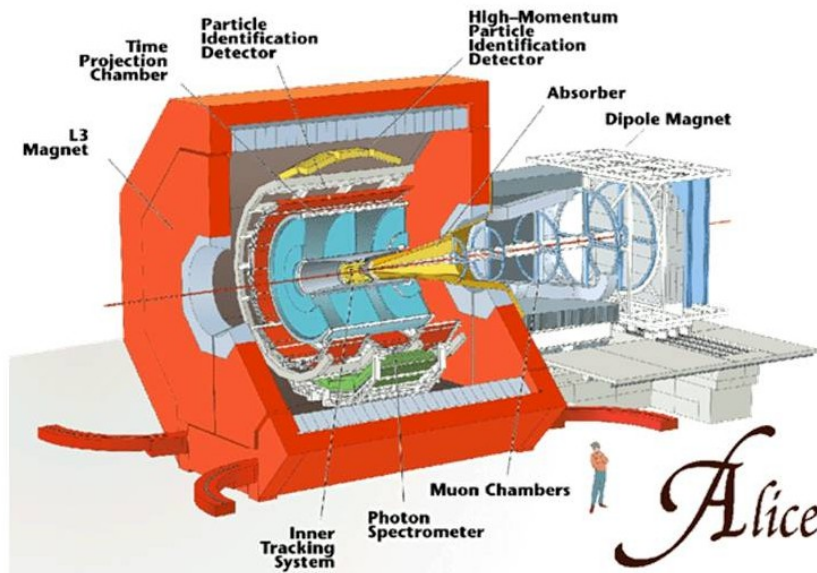
ALICE is designed specifically to look for the quarks and gluons present in the quark-gluon soup immediately after the Big Bang. Not literally, of course, but by reproducing the conditions of that early era by observing the collision of lead ions (hence the Large Ions in the title). The combined energy of two colliding 7TeV beams of lead ions is expected to produce temperatures up to 100,000 times those at the heart of our Sun, just as after the Big Bang. These temperature should liberate quarks and gluons and ALICE is designed to capture and identify them.



Simulated collision of Lead Ions in ALICE

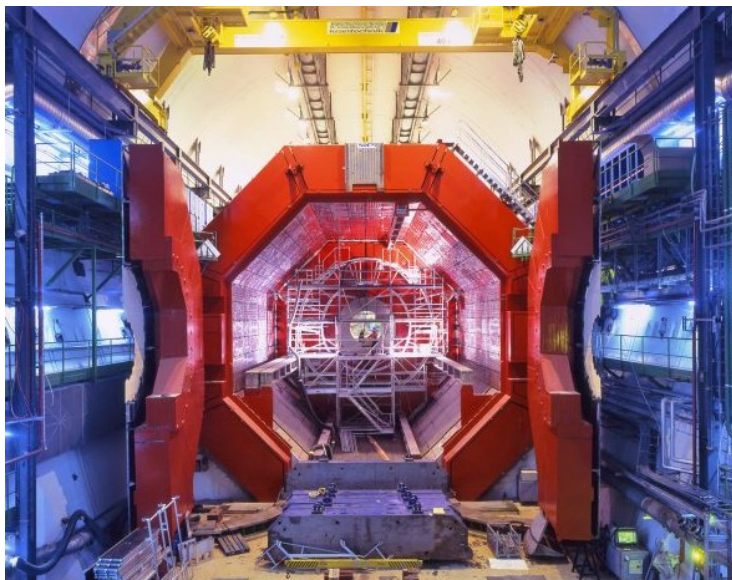
Imagine this. In ALICE's heart, after the lead ion collision, the created quark-gluon plasma survives less than 10-18 seconds. In cooling, the quarks give birth to some 20,000 new particles – for each of the 8,000 collisions per second. Each of the eighteen sub-detectors is designed to determine one property or the identity of more one or more sorts of particles. Gathering this information helps physicists accumulate evidence about the creation of the plasma.

Referred to as 'one of the small detectors', ALICE stands 16 metres tall, 16 metres wide and 26 metres long. It weighs in at about 10,000 tonnes.



A cutaway of the ALICE Detector

ALICE's internal detection systems are radically different to those of ATLAS and CMS, as befits its different purpose. Too complex to describe here, it contains such detectors as a Silicon Pixel Detector (SPD), a very novel technology; a Silicon Drift Detector (SDD); a Time Projection Chamber (TPC) to allow 3-dimensional tracking; and a Time of Flight System (TOF) which is designed to distinguish between pions, kaons and protons with a time resolution of a staggering 50 picoseconds (50×10^{-12} s).

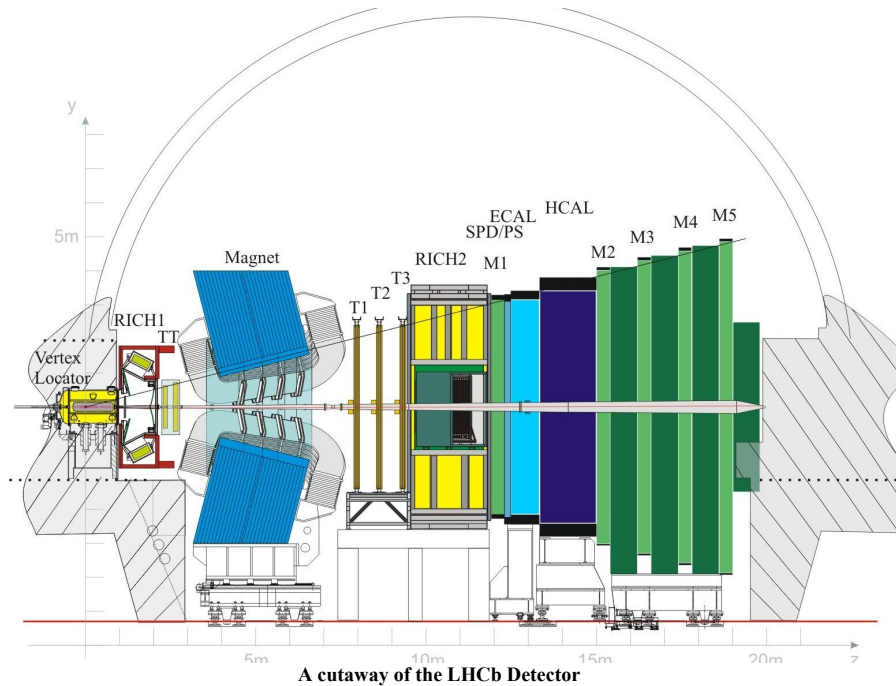


Straight out of a technological wonderland, ALICE sounds appropriately named.

LHCb

The LHCb Detector's prime aim is to answer the question: why was there more matter than anti-matter created after the Big Bang? Thankfully for us there was or we wouldn't be here. But the fact that there was makes physicists suspect that there may be subtle differences between matter and anti-matter. The LHCb experiment hopes to shed light on this question, looking at any differences in their decay rate, for example. The LHCb will study the decay rates of particles with the beauty anti-quark. Since beauty quarks and anti-quarks are heavy and didn't survive in today's universe, by reproducing the high energy, high temperature conditions of the Big bang it is anticipated they will observe these 'beautiful' particles and study their differences. That is where the 'b' comes from in its name: for the 'beauty' quark.

The LHCb collaboration involves 690 scientists from 48 institutions and 15 countries.



The layout of the LHCb detector is radically different from that of ATLAS and CMS. While they were of quasi-cylindrical shapes, optimised for particle detection perpendicular to the beam axis, the geometry of the LHCb resembles that of a reclining pyramid, with its apex located at the collision point. The LHCb's 4,500 tonne detector has been designed to efficiently detect B mesons produced by b and anti-b quarks and to study the products of their decays. Instead of having the 'onion skin' structure of the other LHC particle detectors, the LHCb experiment stretches out 20 metres along the beam pipe, with its detectors spatially organised to efficiently measure the particles of interest.



The LHCb installation completion celebration

Data Analysis

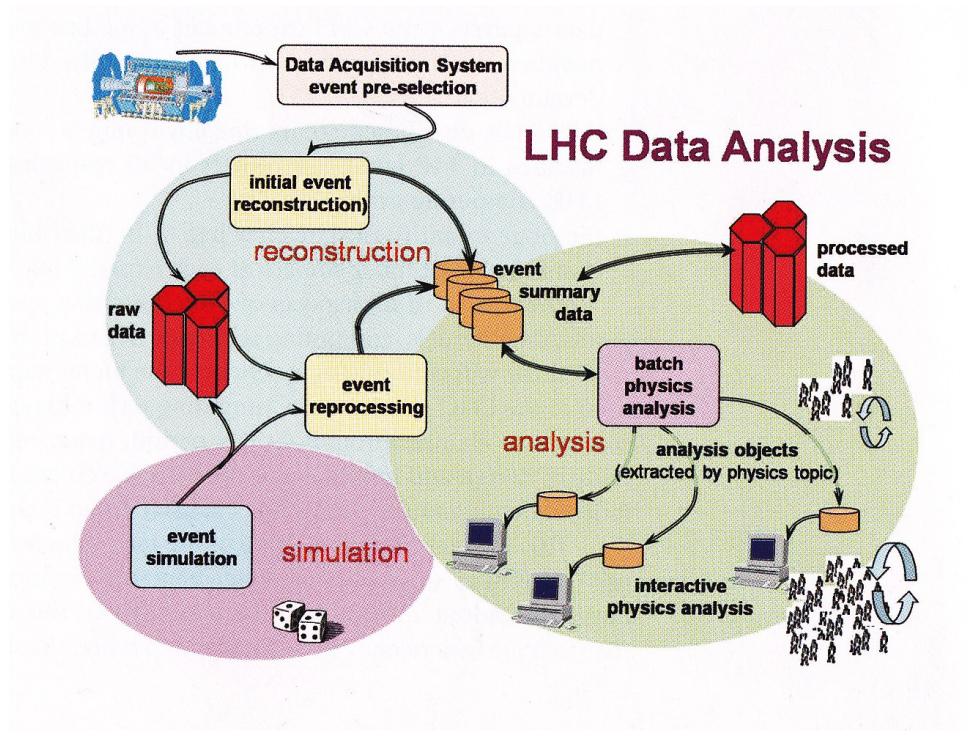
When operating, the LHC produces an enormous amount of data. If it was all recorded, it would equal a rate of data of 1 petabyte per second. 1 petabyte (PB) = 10^{15} bytes approx = 200,000 DVDs full of data. Per second. Overwhelming!

Thankfully, by using special electronics embedded in the detectors themselves and dedicated on-line computers, this is selectively reduced by six orders of magnitude to a much more 'manageable' rate of 1 gigabyte per second. This reduced data volume is recorded in magnetic storage for later analysis.

But even with this huge filtering and reduction process, there is still a formidable amount of data, approximately **30 PB per year**. To put that into perspective: A standard DVC disc is 1mm thick. **30PB equates to a stack of DVDs 6 kilometres high**. That data has to be managed and distributed to the many thousands of scientists around the world who are analysing the physics from the LHC.

As you may imagine, this is a task of gargantuan proportions. Understandably, automatic computer-based algorithms are an essential part of managing such enormous amounts of data which are piling up on the recording system every 25 nanoseconds as each packet of protons collide and produce their shower of particles. These very sophisticated algorithms, purpose designed by the scientists in each detector experiment, are applied successively to the raw data generated and collected for each experiment. This allows data in the form of observable events and physical quantities (such as charge, speed, momentum, energy etc) that can be compared to the theoretical predictions made by computer simulations. Effectively, the algorithms, applied at super-high speeds, sort out the 'chaff from the grain' before storage of the 'interesting' data.

This is all highly complex and a detailed reading of technical descriptions of the data systems, hardware and software involved tends to boggle the mind of laymen such as myself. A simplified 'high level' view of the data flow is shown in the diagram below.



To give an, albeit inadequate, glimpse of the logistics involved and the increased complexity of the LHC computing system from previous high energy accelerators and collider projects, consider these factors:

* For all four detectors – CMS, ATLAS, ALICE and LHCb – there are over 6,000 scientists and engineers involved, an unprecedented number. A significant fraction of these are involved just in algorithm and program development.

* The computer environment to manage all this data analysis is very widely distributed, literally worldwide. There are approximately 100,000 processors installed in 140 computer centres in 35 countries, all integrated into the LHC computing grid, called 'The Grid'.

A key to this grid infrastructure is that the design and construction for it started over ten years ago and its job will continue for many more years. Yet, computer technology moves on at a prodigious rate and all the components of this complex system had to be designed to allow for and accommodate new computing technologies, even technologies not conceived of at the start. Flexibility for seamless transition is the key.

Conclusion

Our tour of the LHC sadly came to an end after a visit to the CERN cafeteria where we shared lunch with what seemed like all 3,000 of the scientists and engineers working at the LHC. It was as chaotic as the inside of one of their detectors. On driving away, we compared personal impressions. One thing was predominant in our minds: The Large Hadron Collider's claim to be the most complex single 'machine' ever built was no exaggeration.

* * *

References:

"The Large Hadron Collider: A Marvel of Technology" edited by Lyndon Evans.

"Destination Universe" Published by CERN