ONE

Before the Data

1.1 Why So Big?

The Large Hadron Collider (LHC) sits in a tunnel 27km long and about 100m underground. If you are familiar with London, it might help you to know that 27km is about as long as the Circle Line on the Underground, and the tunnel itself is similar in size to the Northern Line. If that doesn't help, then try this.

Imagine setting off from Meyrin, on the Swiss–French border near the airport, and driving towards the French countryside. The Jura Mountains are in front of you, Geneva Airport is behind. As you pass the border, you also pass the main site of the CERN laboratory on your left, and if you look to the right you will see a big wooden globe that looks like a sort of eco-nuclear reactor (it's not, it's an exhibition space, though it is eco-friendly, apparently), and you might catch a glimpse of the building housing the control room of the ATLAS experiment. You will know it if you see it, because it has a huge mural of the ATLAS detector itself on the wall.

Big though it is, the mural is painted to only one-third scale. ATLAS is very large, and is hidden underground, positioned at one of the interaction points of the LHC. These are the points where the two highest-energy particle beams in the world are brought into head-on collision. ATLAS is one of the two big general-purpose particle detectors designed to measure the results of these collisions.

Continue driving. You may imagine yourself in a nerdy little white van with a CERN logo on the side if this helps.

Pass through the village of St-Genis and continue into the Pays de Gex, in the foothills of the Jura Mountains. You are now surrounded by the LHC. If you are imagining yourself in winter, you might see the lifts of Crozet, the little Monts Jura ski resort, chugging away ahead of you. (Mont Blanc is behind you on the horizon, but keep your eyes on the road.) Keep driving, bear right towards Gex, maybe pass through the villages of Pregnin, Véraz and Brétigny. After about 25 minutes' driving through the French countryside – longer if you get stuck behind a tractor – you will get to the village of Cessy, near Gex. Here you will find the top of the shaft that leads down to CMS, the other big general-purpose detector on the LHC ring. ATLAS and CMS are independent rivals, designed differently by different collaborations of physicists, but with the same goal: to measure as well as possible the particles produced when protons collide in the LHC. They were designed to cross-check each other's observations, and to compete head-to-head for the quickest and best results.

All this time, on your journey from ATLAS to CMS, you have been inside the circumference of the world's biggest physics experiment. You entered it at the border when you passed ATLAS, and have now crossed its diameter.

The LHC is designed to collide subatomic particles at the highest energies ever achieved in a particle accelerator. We do this to study the fabric of the universe at the smallest distances possible, which for reasons to be described later also implies the highest energies possible. Given that the experiment is designed to look at very small things, it might be a surprise that it is so big. Building a long tunnel is very expensive, so why not make a smaller one?

In fact, it is the length of the tunnel that limits the energy of the colliding beams. If you accept the fact that to study small stuff you need high energies (please do, for now at least), you can understand why the LHC needs to be so big just from an understanding of fairly everyday physics.

Particles travel in a straight line at a constant speed, unless a force acts on them. This is one of Newton's laws of motion. In everyday life it isn't completely obvious (Newton was quite clever to work it out), but once you are aware of it, it is easy to see it in action.

The reason it is not completely obvious in everyday experience is that on Earth practically everything that moves has forces due to friction and air resistance acting on it, and everything experiences gravity. This is why if you set a ball rolling, it will eventually stop. Friction and air resistance act on it to slow it down. And if you throw a ball in the air, gravity will slow it down and eventually drag it back.

But in situations where friction or gravity can be ignored, things are clearer. Driving a fast car, or even a nerdy CERN van, you clearly have to apply a force, via the brakes, to slow it down. And more relevantly in the context of the LHC, if you want to change direction, to turn a corner at speed, this can only be done if there is sufficient friction between the tyres of the van and the road. Otherwise, you skid.

The driver and passengers experience a rapid turn of a corner as a sort of 'pseudo-force'. The van is turning, but your body wants to carry on in a straight line, so you feel as though you are being pressed against the sides of the van. It would be more true to our understanding of physics to think of the sides of the van as pushing against you, to force you to change direction, pushing you round the corner along with the vehicle.

The combination of speed and direction is called velocity. And if you combine the velocity and the mass of the object (the van, for example, or the passenger), you get the momentum. The bigger the mass, or the velocity, the bigger the momentum, and if you want to change the momentum of something, you need to apply a force to it.

I am being deliberately vague about how the velocity and mass combine to give momentum. At speeds much lower than the speed of light, it is good enough to just multiply – momentum is mass times velocity – and this is probably the right answer if you are taking a school course in physics. However, the exact expression is a little different, and the difference gets more and more important as speeds approach the speed of light. Then you need Einstein and relativity (of which more later), rather than Newtonian mechanics. But don't try this in a van.

Regardless of that, the larger the desired change in momentum, the bigger the force has to be. Hence the brakes on a lorry need to be able to exert more force than the brakes on a van, because even if the velocity is the same, the mass of the lorry is bigger so the change in momentum involved in making it stop is bigger.

This is the situation of the protons in the LHC tunnel. These are the highest-energy, and highest-momentum, subatomic particles ever accelerated in a laboratory. Even though the mass of a proton is tiny, their speed is tremendously high. They are really, really determined to travel in a straight line. So, to make the two beams of protons bend around the LHC and come into collision requires a huge force. The force is provided by the most powerful bending magnets we could build.

Given this maximum force, there is then a trade-off between how sharp the bend in the accelerator is and how high the proton momentum can be. Back to the van: this is exactly equivalent to the fact that there is a maximum speed at which you can take a given corner without skidding. If the corner is sharp, the speed has to be low, but for a gentle curve you can go faster. This, then, is why the LHC is so big. A big ring has more gentle curvature than a small one, and so the protons can get to a higher momentum without 'skidding'. Or, in their case, 'catastrophically escaping the LHC and vaporising expensive pieces of magnet or detector'. Something to be avoided.

The maximum bending power of magnets is thus the reason that proton accelerators need to be large if they are to get to high energies. For the other commonly collided particle, the electron, there is another reason that is worth looking at.

Before the LHC was installed, another machine occupied the 27km tunnel under the Swiss–French border. This was LEP – the Large Electron–Positron Collider. (Positrons are the anti-particle of the electron, carrying positive charge, in contrast to the electron's negative. LEP collided electrons and positrons together. Incidentally, people occasionally accuse particle physicists of hyping up their kit, but these are very descriptive, even dull, names.) LEP was turned off in the year 2000 because it had explored most of the physics within its reach and could not increase its energy further. The reason it could not go higher was, as with all the protons, also connected to the size of the tunnel, but in a different way.

This is to do with the fact that electrons have a mass about 1800 times smaller than the proton. Now, at the highest energies that doesn't make any significant difference to the force required to bend them round a corner. This is because, whether they are electrons or protons, they are moving very close to the speed of light, so you need the full special relativity expression for momentum, and the net result is that the mass they have when they are at rest is irrelevant for calculating the required force. So that wasn't the problem.

The problem was synchrotron radiation. This is the energy radiated by charged particles when they are accelerated. It is a universal phenomenon, roughly analogous to the wave a speedboat makes when it turns in the water. As a charged particle accelerates round a corner, photons fly off and carry away energy.

The effect is actually much worse for particles with low mass. The amount of synchrotron radiation given off when a particle accelerates depends very strongly on the mass: it decreases like the mass to the fourth power. So, as the proton mass is 1800 times bigger, the energy lost on the bends for electrons is (1800 x 1800 x 1800 x 1800) or about 11 trillion times larger than it is for protons.

As the electrons and positrons squealed round the corners of LEP, photons were radiated this way, and with every revolution of the beam around the ring, more energy had to be pumped in to compensate. This is done by radio-frequency electromagnetic waves confined in big metal structures at intervals around the ring. Electric and magnetic fields oscillate in these structures precisely in time with the passing of the bunches of electrons, so that every time a bunch arrives it gets a kick from the field. This is true in all such machines. But at some point you reach a beam energy where so much is lost in synchrotron radiation that the electromagnetic waves in those structures cannot replace it. That's your maximum collision energy. LEP hit that wall.

This is where the size of the tunnel comes in again, of course. A 27km tunnel has a rather gentle curve. If it were smaller, the bends would be sharper, the acceleration would need to be bigger, so the energy lost through synchrotron radiation would be greater, and the maximum collision energy would be lower.

As an aside, this synchrotron radiation is very useful in other contexts. The Diamond Light Source at Harwell in Oxfordshire, for example, was built to produce it intentionally. The radiated beams of photons are used to study atoms, crystals, molecules, materials and surfaces. Many machines and laboratories originally built to study particle physics have been converted to become light sources once they have been superseded in the quest for higher energies. I have reason to be grateful for this personally, in fact. I did my doctoral work in Hamburg, at the DESY (Deutsches Elektronen-Synchrotron) laboratory. The particle physics of interest there at the time was the HERA electron–proton collider, where I worked in the ZEUS collaboration. But my then girlfriend was a crystallographer, using synchrotron light to work out the structure of proteins and other stuff. Because of the symbiotic relationship between particle-physics accelerators and synchrotron light sources, there is a branch of the European Molecular Biology Lab at DESY, and after a high-level discussion in the crowd at a St Pauli football match, Susanna managed to get her PhD supervisor to send her to Hamburg for most of her research. We've been married 20 years now, and it's all very fine and romantic. But synchrotron radiation is still a pain in the arse if you want a high-energy electron beam.

So, LEP was shut down in 2000 and dismantled, and installation of the LHC began. The LHC can get to higher energies because it collides protons with 11 trillion times less of a synchrotron radiation problem, but it requires the most powerful bending magnets you can make if you want to get to the highest possible momentum.

The formal approval for construction of ATLAS and CMS was given on 1 July 1997 by the then Director General of CERN, Chris Llewellyn Smith.[1]

LEP had been good, but the protons promised more.

Glossary: The Standard Model Particles and Forces

If you just want to crack on with the story and don't mind the odd unfamiliar word, you can skip these 'Glossary' bits. But without knowing something about the Standard Model, some of it might not make much sense.

The Standard Model of particle physics is our current best answer to the question 'What is stuff made of, if you break it down into its smallest components?'

Start with anything – a rock, the air, this book, your head – and pull it into its component parts (I recommend this remains a thought experiment). You will find fascinating layers of structure, micro- and nano-scale bits and pieces: fibres, cells, mitochondria.

You will eventually find molecules. With enough energy you can break them apart into component atoms. Atoms consist of a dense nucleus surrounded by electrons.

With a bit more energy, you can separate the electrons from the nucleus. With more energy still, the nucleus can be broken into protons and neutrons. With still more energy (and now you

---

[1] Incidentally, a man who previously, as head of physics at Oxford and afterwards as provost of UCL, seems to have had a period of following me around and being my boss.

do need a big collider!), you can see quarks inside those protons and neutrons.

We have never managed to see anything inside a quark, or break one into pieces.

If, at the 'atom-smashing' stage, we had ignored the nucleus and tried breaking up the electron, we'd have reached that point earlier. We have never managed to see anything inside an electron, or break one into pieces. This – the fact that we haven't managed to break one yet – is our working definition of what it means for a particle to be 'fundamental'.

And a key point is that wherever we had started, with whatever material, we would have ended up with electrons and quarks. In the Standard Model, they are the stuff that everything is made of, and they themselves are not made of other stuff.

You will come across a lot of particles in this book, but remember, there aren't many different kinds of fundamental ones when you get right down to it.

Electrons are an example of a class of particles called leptons. There are also muons and taus, which are just like electrons only heavier. The only other leptons are the three kinds of neutrino. Neutrinos do not interact much with other matter, but there are lots of them around. More than a trillion neutrinos pass through you from the Sun every second.

The other class of fundamental-matter particles consists of the quarks. There are six of them, too, just as there are six lepton types. They are called up, down, strange, charm, bottom and top, becoming more massive as you go (but peaking on whimsy in the middle).

Protons and neutrons are made of up and down quarks. Quarks are never found out on their own, they are always stuck together in bigger particles. These particles, the ones made of quarks, are generically called hadrons (hence the Large Hadron Collider, which mostly collides protons but occasionally collides atomic nuclei, which also have neutrons inside).

Those are all the matter particles we know of. They all have anti-particle partners, and they all interact with each other – attracting, repelling, scattering – via forces, which are carried by another kind of particle – vector bosons.

The electromagnetic force is carried by photons (quanta of light) and is experienced by all charged particles. That is, everything except the neutrinos.

The strong force is carried by gluons, and is only experienced by the quarks.

The weak force is carried by W and Z bosons, and all particles experience this.

To make the Standard Model work, and in particular to allow the fundamental particles to have mass, another unique and completely new object is also required – a Higgs boson. The hunt for this is, of course, the main topic of this story and I'll say much more about it later.

Gravity doesn't fit into the Standard Model. It is described by Einstein's theory of general relativity, but we do not know how to make a working quantum theory of that.

Those are the actors on the stage of the universe. There are lots of open questions in physics, but an astonishingly wide array of data – most of physics, chemistry and biology – from very large to very small distance scales, can be described astonishingly accurately by just these elements: quarks, leptons and the four forces between them, and the Higgs boson.

1.2 The 'No Lose' Theorem

I didn't begin working seriously on the LHC until around 2001. That was about nine years before we got our first high-energy collisions. Believe it or not, this makes me a bit of a Johnny-come-lately to the experiment. Options for a large hadron collider had already been considered in the design of the 27km tunnel for LEP, and were mentioned in the LEP design report in 1984, when I was just finishing secondary school and moving to sixth-form college. There would be many years of scientific, technical, financial and political discussions, followed by R & D, simulation and more politics, before the LHC gained approval in 1997.

Back in 1997, I had just moved to London from Hamburg and was still completely absorbed in work at HERA. It is a feature of big collaborations that you accumulate responsibilities along with a bank of experience and expert knowledge that can make it hard to disengage. It is difficult to climb a new learning curve on another experiment, with its confusing software and hardware and unfamiliar physics. Sometimes you need a bit of a shove to really start doing something else.

For me, bizarrely, the shove was the birth of my first child. This was such an overriding priority that I managed to say no to a whole bunch of managerial and technical roles within the ZEUS collaboration. I wanted no responsibilities that would conflict with the terrifying challenges of looking after Susanna during her pregnancy and of being a dad after it.

As it turned out, the whole thing went very smoothly and was generally wonderful. So, as a bonus, I had lots of free time to think creatively about physics. One of the things I'd long been wanting to do was read enough and think enough to get my head around physics at the LHC, which was by then under construction at CERN. This holiday from HERA heat was the opportunity. With a couple of friends, Jeff and Brian, also HERA physicists,[2] I'd been thinking about what we might do – what the most exciting things to study would be. We were all very sceptical about new 'Beyond the Standard Model' physics and were keen to work on measurements of real things that would actually happen, rather than seeking evidence for speculative ideas to which we accorded little credibility. I think this may have been because we all came from a HERA background, where precision measurement was the main goal. Although to be honest, the main legacy of LEP was also precision measurement, so maybe it made no difference.

Anyhow. Not only did we not believe in such things as supersymmetry, or large extra dimensions, or Technicolor, all of them speculative extensions of the Standard Model designed to solve some of the problems with it. We didn't even believe in the Higgs boson – an integral part of the Standard Model, but one lacking experimental verification. So we asked ourselves, 'What is the most important and interesting thing to measure if there are no new particles?' A pessimist's approach, perhaps, but still fun.

The answer we came up with[3] was vector-boson scattering. This is a peculiar and rare scattering process that is expected to happen occasionally in very high-energy collisions, and it lay behind what was called the 'no lose' theorem at the LHC. It is

---

[2] Jeff Forshaw and Brian Cox, who amongst other things have also written physics books, though not about this yet.

[3] After some reading around – I'm not claiming we were the first to think of this!

very deeply connected to the reasons why the Higgs boson is so important. So it wasn't a bad choice for a first bit of LHC physics for us to look at, and it's worth spending a bit of time on now.

Vector bosons are force carriers. The photon, which is the quantum of light and carries the electromagnetic force, is a vector boson. Of more interest here, though, are the W, and to some extent the Z, bosons. These carry the weak force, and one of the oddest things about them is that, unlike the photon, they have mass.

In a proton–proton collision at the LHC, you have to picture two quarks, one from a proton in each beam, zooming towards each other. There is a small but non-zero chance that, as they do this, each will radiate a W boson. There's an even smaller, but still non-zero, chance that these W bosons will hit each other. That is vector-boson scattering – WW scattering in this case. It could happen with Zs or photons too. There are a bunch of different ways the bosons can bounce off each other, or fuse together and break up again. As is always the case in quantum mechanics, all the possibilities have to be taken into account and combined[4] – sometimes they add up, sometimes they subtract from each other. Put the whole thing together and you get the probability of the WW scattering occurring.

The 'no lose' theorem came from this calculation. Some of those scattering possibilities include a Higgs boson, and at the time there was no direct evidence for such a beast. However, if you do the sum and do not include a Higgs boson, then as you go to higher and higher energies, the probability of WW

---

[4] This includes taking into account time orderings other than the quarks-emit-Ws-that-then-collide one I gave here.

scattering grows and grows.[5] At some point you get nonsense answers involving probabilities bigger than one, or infinities. That is just a sign that your theory is broken – there won't be infinities in nature – but what it meant was that either a Higgs boson would be discovered at the LHC, or some other new physics would come into play and keep the calculation sane.

So in the pessimist's scenario of no Higgs boson, no black holes, whatever, measuring WW scattering might well be the only, or best, clue as to what was going on. It certainly had to involve either a Higgs boson or something else new – hence the 'no lose' theorem. By studying these scatterings we were certain we would discover some interesting physics.

Measuring WW scattering properly would be difficult, and we found lots of fun challenges. Manchester had done a deal with Apple and had a big new farm of Macs that had just started running Unix (OSX), making them useful to us (if still more expensive than the Linux boxes everyone else had). I have fond memories of sitting in a flat in Saddleworth, feeding the farm with lots of simulation jobs to test our ideas, then popping over the road to the pub for beer and dinner to argue about newer ideas. This was still before my son was born, but I had already divested myself of lots of other responsibilities. We submitted the paper in January 2002 and it was more or less ignored for six years, though it later became more fashionable and I'm very proud of it. One of the ideas we used turned out to be quite widely useful and would feature in the Higgs search itself.

---

[5] The possibilities involving the Higgs would contribute with a negative sign, and so they would stop this happening.

1.3 People Are Going to Be Interested in This . . .

While work was going on at CERN and around the world to construct the LHC and its detectors, it became increasingly obvious that quite a lot of people were going to be interested in the project, for all kinds of reasons. For the engineering and science, of course. But also because of the sheer scale, including the cost of the thing. The international collaboration and the sociology of several thousand physicists working together were intriguing to quite a few people, including academics in social sciences. Plus, of course, there were the two or three delusional publicity-seekers who thought, or claimed to think, that we were about to destroy Switzerland. Or the world. Or the entire universe.

The last bunch – the delusionists and conspiracy theorists – were bound to get lots of media coverage, 'because of balance'.[6] The only way to deal with that is to get real information out there. Also, since the European taxpayer had been investing something like the equivalent of a billion euros in CERN every year, we really owed it to people to explain what we'd done with the cash, and why.

Thoughts like these were passing through the minds of many people involved, including, I am sure, James Gillies, the head of communications at CERN, and many good science journalists. This is presumably why, in 2008, the doors of CERN were flung open to the world's media for 'Big Bang Day', when we switched on the machine.

---

[6] To quote David Shiffman: 'World's leading experts say there's a problem with false balance in environmental journalism, but Steve disagrees.'

Such thinking was certainly one reason why I agreed to be part of a series of short documentary films called *Colliding Particles*. These were a sort of 'fly on the wall' affair that started in the summer of 2008. Mike Paterson was the cameraman, producer, interviewer and director – everything, in fact, except animator and occasionally soundman. He won some support from the Science and Technology Facilities Council (STFC), the research council that funds particle physics in the UK, to make the films. They were aimed at schools, specifically at a then new part of the curriculum based on learning how science works. Apparently pupils would learn this by watching some physicists from behind Mike's camera.

Actually, it worked out very well. Amongst other things, Mike has some sort of genius for cutting and matching a long ramble by me to amazing pictures that make me appear coherent. There is roughly five minutes in the first film when I speak without ever finishing a sentence properly, while shots of the LHC and of the ATLAS detector being assembled fill the screen, imparting to the viewer a sense of the vision and wisdom I am pouring forth. Well, that's how it looked to me and my mum, anyway. The films also featured Adam Davison, who at the start was my PhD student, later a postdoc, and Gavin Salam, a theorist in Paris, and were (at least to start with) loosely based around a paper we'd written in 2007 with Gavin's student, Mathieu Rubin.

I mention all this now to show that, despite strong reservations from some particle physicists, we were taking public engagement quite seriously. The film was one of several initiatives at the time, a level of activity unprecedented in particle physics. Mike, Adam, Gavin and the films will pop up from time to time as we progress towards discovery.

Anyway, here we were, after many years of R & D and eight years of construction, at the switch-on.

It was 10 September 2008. In CERN, the control room was packed with journalists and Brian Cox. BBC Radio 4 made a day out of it – 'Big Bang Day'. I had come back from CERN and was in Westminster, in a big hall with the minister (John Denham, Secretary of State for Innovation, Universities and Skills), many other of the great and occasionally good, and more journalists. This was all very exciting and totally new territory for us in terms of engagement with the media and politicians, but still the most terrifying and exciting thing was the fact that after so many years of preparation we were finally about to switch on our experiment.

The switch-on was pretty closely choreographed, with Lyn Evans, the LHC project leader, the ringmaster in the LHC control room. The beams were to be sent into the LHC octant by octant. That is to say, initially the beams would go an eighth of the way around the 27km tunnel and hit a beam blocker. Then a quarter, three-eighths, and so on until hopefully they got all the way around and came back again, registering two dots (one on the way in, one on the way back) on a scintillation counter that was the centre of attention for thousands of physicists and a significant chunk of the world's media.

One nice thing was that when the beam hit a blocker in front of one of the detectors, a spray of particles would be produced that would register in the detector – the first beam activity we'd seen in these highly complex and sensitive pieces of kit. Less nice from my point of view was that because Lyn chose to send the clockwise beam around first, ATLAS was the last in line, and therefore the last to see particles. Still, the next step after us was a full circuit. Lyn counted down: 3 . . . 2 . . . 1.

A moment of nerves when nothing appeared . . . then bingo! Two dots. The most exciting two dots I've seen before or since. For the first time, a beam had successfully completed a circuit of the LHC.

As the day progressed, beams were sent round in both directions and stored successfully. Sheer exhaustion at Westminster led us to the pub next door, though the accelerator teams at CERN were still hard at work. I will never forget drinking a beer at lunchtime and seeing progress updates of my own physics experiment on the BBC news ticker at the bottom of the pub TV screen. It was hard not to be triumphalist about it all. The headlines the day after, declaring that we hadn't destroyed the world, were fun if a bit premature (we hadn't yet collided the beams, after all), but we had got ourselves a working experiment and we had managed to share our excitement with the people who were paying partners in the enterprise. I was looking forward to welcoming the new PhD students to UCL with the promise of imminent data.

Nine days later, it had all gone catastrophically wrong.

1.4 Breakdown

As I have already described, the limiting factor on the maximum energy the protons can get to in the LHC is the centripetal force – the force needed to bend them round the corners so that they stay in the ring. This force is provided by huge magnets. Imagine whirling a brick round your head on a thin piece of string. If you whirl it too fast, the string will break. The protons are the brick, our magnets are the string. We really do not want them to break.

In the days after the start-up, many tests were performed on the LHC. In particular, when the beams were circulated on 10 September, the magnets had not been running at full strength.

The magnets are electromagnets, meaning that the magnetic field is generated by circulating electric currents. As well as the impressive engineering and industrial know-how that goes into making them (there are 1232 dipole magnets to bend the beam, each of them is 15m long and has a mass of 35 tonnes), there is a lot of physics here.

The facts that electric currents create magnetic fields and that magnetic fields bend electric currents were observed and measured by Michael Faraday in the 19th century, and built into the theory of electromagnetism by James Clerk Maxwell. Maxwell's equations are one of the highlights of a physics degree, and are arguably the first expression of a mathematical unification of two apparently different forces (electrostatics and magnetism), setting a trend that physics has followed ever since.

To produce the force required to bend the proton beams at the LHC, very high currents are needed. At full power the magnets need to carry a current of nearly 12,000 amps (A). This is about fifty thousand times bigger than the current drawn by a typical incandescent household light bulb.

When an electric current flows through a normal material (such as the light-bulb filament), the electrons carrying the current collide with the vibrating atoms of the material. This makes the electrons lose energy and the atoms vibrate more vigorously, heating up the material. This is electrical resistance, and it is a big problem if you want a current of 12,000 A. Essentially, any normal material will vaporise.

The discovery of superconductivity changed this. Superconducting materials offer zero resistance to the flow of an electric current. This is a rather amazing quantum mechanical effect, understood in the Bardeen, Cooper and Schrieffer theory. At low temperatures, the electrons form pairs and start to behave like bosons.[7] This allows them all to sink into the same quantum state – a condensate – and overlap with each other. At that point, it takes quite a lot of energy to have any effect on a pair, because that would mean changing the whole condensate, which is in a coherent quantum state. In general, the collisions with the material don't have enough energy to do this – at least not when the material is very cold, when the atoms in the material are hardly vibrating at all. So the pairs of electrons flow on unimpeded, losing no energy and experiencing no resistance.

The magnets in the LHC are superconducting. They are cooled to a temperature of 1.9 kelvins (-271.3°C), using pressurised liquid helium.[8]

On 19 September, the LHC team were testing the magnets up to full electric current – the current at which they could bend beams around the LHC at full energy. The LHC operates as eight independent sections, which can be powered, warmed or cooled separately. They had commissioned seven of these octants to full current and were on the eighth and last, almost ready for first collisions.

At this point, all information from the monitors and sensors in that sector suddenly ceased.

---

[7] See Glossary: Bosons and Fermions (pp.38–40).

[8] This is, as is frequently pointed out, colder than outer space. The cosmic microwave background has a temperature of 2.7 kelvins. However, it is not, as is sometimes said, the coldest place in the universe. Apparently the Boomerang Nebula is at 1 kelvin. I have no idea why.

Watching the dashboard remotely in London, I just saw a note that 'first collisions' had been delayed for at least a few days. But the reality was much worse. There had been a catastrophic explosion. Several of the huge magnets had been ripped out of their concrete moorings. We were not going to get our first proton–proton collisions for more than a year.

I had to explain to various journalists, including a live phone call to breakfast TV, what had happened. To tell the truth, other than that it was very bad, we didn't completely know at that stage. By far the worst experience was having to stand up in front of the PhD students and tell them they wouldn't be getting any collision data for a while longer. I phrased it as 'two steps forward, one step back', but the step backwards felt enormous.

The full story gradually emerged days and weeks later. There had been a fault in one of the connectors between two magnets. This was due to a flaw in the welding. The connector developed a small electrical resistance. On its own that would have been a serious but not catastrophic problem. Part of a superconducting system suddenly developing a resistance is an occupational hazard of the technology; it leads to what is called a 'quench', and the superconducting magnets have elaborate protection against quenches, meaning that the enormous energy in the electric currents is safely dissipated before it heats up and damages the magnet.

Unfortunately, the quench protection didn't extend to the connectors. The current was not safely dissipated, and the connector was vaporised.

Again, on its own this would have been a big problem, leading to months of delay, but not a catastrophic one. However, the huge current now had nowhere to go and sparked across the gap left by the vaporised connector. The spark punctured the liquid-helium containment vessel and suddenly tonnes of pressurised liquid helium became a gas. Very quickly. This was the explosion, powerful enough to tear some of those 35-tonne magnets from their concrete moorings. Several of them were destroyed or damaged, and the precision instrumentation and delicate cryogenics were turned into a mess of twisted metal. It all meant that we had a very long wait, and a lot of work to do, before the machine would deliver physics results.

1.5. While We Were Waiting . . .

The aftermath of this disaster was quite instructive.

As with any big project, there are people who object to the amount of resources spent. There are people who resent the huge public and media interest that is denied to lots of other good science. Also, apparently, particle physicists are sometimes perceived as arrogant, though I can't think why. For whatever reason, while there was much genuine sympathy for us disappointed LHC physicists, there was also *Schadenfreude*.

And remember, a substantial section of the particle-physics community thought the whole media- and public-engagement exercise around the start-up a big mistake. Many colleagues saw it as at best a hostage to fortune, and at worst unscientific media hype. The failure of the machine nine days after such a high-profile public event must have seemed a massive opportunity to say 'I told you so'. I was mainly just miserable about the delay, and did indeed have moments when I felt we'd made massive idiots of ourselves in a very public fashion and should have kept quiet until we had the results.

In October 2008, very soon after the catastrophe, we had the 'inauguration' of the LHC. This was a very odd event. It was a formal celebration, planned before the disaster. It was held in the huge magnet-testing hall, soon to be reopened to test (re-test) the repaired and refurbished magnets, which would be needed to repair the wounded machine below our feet. Despite being very interested in CERN and the LHC, Lord Drayson, the UK science minister, didn't show up. To be honest, I don't blame him. Frankly, it was a depressing affair.

But these feelings gradually changed.

The embarrassment became almost a source of pride. The failures in the magnet connector should not have happened, of course. But listening to accelerator physicists and engineers diagnose and discuss the systems involved just emphasised the amazing complexity of the LHC and the amount of new technology that had been integrated, on an industrial scale, into this huge machine. Not only were we on the edge of physics, we were on the edge of engineering, too. Also, no one had been hurt. In fact the whole LEP and LHC civil-engineering project, which was on the scale of the Eurotunnel construction, was carried out with remarkably few casualties.

Research carries risks. As the project leader Lyn Evans said, many times in many interviews, the LHC was its own prototype. Nothing like it had been done before. A notable feature of the wide-eyed physicists staring into multiple cameras on 10 September was that we were more nervous and excited about the possibility our experiment might, or might not, work than we were about our live-TV debut. The nearest equivalent is the visible nervousness on the faces of space scientists when a precisely engineered satellite they have built is making its way into orbit on top of a plume of flame.[9]

This was really doing science in public, and the reality is that science is not a seamless progression of triumphant advances. Two steps forward, one step back, indeed.

From the point of view of the media, this twist just prolonged a good story, and they were remarkably reasonable in the way they treated us, despite (or including, really) a certain amount of mocking on the occasional game show.

As for the misery . . . well, that was harder to shake, but there were some important things to be getting on with. We'd locked ourselves into an obsolete version of the Linux operating system, fearing to upgrade with data imminent. We could fix that – and a lot of other bugs frozen into our software because we didn't have time to fix them properly. One of the most significant things we did was to get rid of an obsolete jet finder. But I realise that's going to take some explaining.

Glossary: Quarks, Gluons and Jets

Jets are what quarks and gluons make when they try to escape.

Every proton (in fact, every hadron) is made up of quarks stuck together by gluons. As already mentioned, these gluons are the force carriers of the strong force, just as photons carry the electromagnetic force, and the W and Z bosons carry the weak force. The role played by the electric charge in electromagnetism

---

[9] In a display cabinet at UCL's Mullard Space Science Laboratory, they have the twisted remnants of delicate electronics from the first version of the Cluster mission, fished out of the Kourou swamps in French Guiana after the Ariane launch failed in 1996. Chastening stuff. But Cluster flew again, as Cluster II, and did the science it was built for.

is played by a quantity called ‘color’ in the case of the strong force. This color has nothing to do with the colour we perceive with our eyes, and I’ll use the US spelling to distinguish between them, since the physicist who introduced it, Oscar W. Greenberg, is an American. However, there is an analogy in there.

For electric charges, there is only one way to get a neutral charge – you have exactly as many anti-charges (negative charges) as positive charges. So an atom is neutral because the number of protons in the nucleus (carrying one positive charge each) is exactly equal to the number of electrons in the cloud around the nucleus (each carrying a negative charge). The result of adding up all the positive and negative charges is zero – they cancel each other out, and the atom is neutral.

For color, the same thing can happen. If we (arbitrarily) name one of the colors ‘red’, then there is also an ‘anti-red’ color (which you might want to call cyan since it’s the complementary colour, though personally I think that pushes the analogy between the quanta of charge for the strong force and visible colours a bit too far). Color-neutral objects called mesons can be made by combining a color with its anti-color.[10] But one difference between electromagnetism (more correctly known in its quantum version as quantum electrodynamics, QED) and the strong force (known as quantum chromodynamics, QCD) is that there is an additional possible way of getting a color-neutral object.

There are three possible colors, often (still arbitrarily) called red,

---

[10] The name meson comes from the Greek *meso*, meaning ‘medium’, because its mass is lighter than that of protons and neutrons but heavier than an electron’s. A meson contains a quark and an antiquark, and it is color-neutral because these will have, for example, red and anti-red color respectively, which cancel each other out. Or, if you prefer, red and cyan, making white.

green and blue. If you have one of each, you also get a color-neutral object. This is where the analogy with mixing three primary colours to get white comes in. Protons and neutrons are made this way. They contain three quarks, one of each color, and so also end up being colorless. Particles made of three quarks like this are called baryons (from the Greek, meaning ‘heavy’) and protons and neutrons are the most common examples. Mesons and baryons are both subclasses of hadrons, of course, since anything made of quarks is a hadron.

An odd feature of all this is that quarks are never observed on their own, a long way away from other quarks. They are always confined inside color-neutral hadrons of one kind or another, due to another peculiarity of the strong force.

Most of the fundamental forces get weaker with distance – the attractive force between a positive and a negative electric charge, for example, is weaker the further away from each other the charges are (it falls off as the inverse of the separation squared – $1/r^2$). But the strong nuclear force is different. The force between two quarks actually gets stronger as you pull them apart. It is as though they are attached to each other by an elastic band or a piece of string. As they move apart, the string becomes taut, and a large amount of energy is stored in the string tension.

Inside the LHC, when two quarks inside protons bounce off each other, they head away from each other at practically the speed of light and with an enormous amount of energy. Initially, the ‘string’ is slack and they feel very little force. This phenomenon is called ‘asymptotic freedom’, and the 2004 Nobel Prize in Physics was awarded jointly to David J. Gross, H. David Politzer and Frank Wilczek ‘for the discovery of asymptotic freedom in the theory of the strong interaction’. It means that when quarks are inside the proton, you can for some purposes and to some approximation treat them as though they are free, as though they aren't bound together at all.

However, that apparent freedom ends rather quickly when you try to remove a quark from a proton – for instance by hitting it with another quark from another proton going the opposite way around the LHC. Though they fly away from each other initially (and even radiate more gluons and quarks as they accelerate), the string gets taut almost immediately, and the quarks and gluons know they aren't really free.

What happens next is intriguing, though. There is so much energy stored in the tension of the string between two quarks, as the force pulling them back together doesn't fall off with distance, that it becomes energetically possible, and indeed favourable, to make a new quark and an antiquark. The cost of doing this in terms of energy (E) is the mass of the quark plus the mass of the antiquark, multiplied by the speed of light squared ($E = mc^2$, but you probably knew that). But the benefit is that you can have much shorter strings and so much less potential energy stored in the string tension.

You can think of the quarks as being the ends of the string. They fly away from each other until at some point the string snaps and two new ends (new quarks) are produced.

Eventually, we see a spray of hadrons. You might think that this is a bit useless if we really want to see what is going on with the fundamental particles – the quarks and the gluons and so on. But all is not lost. Because the initial quarks get kicked so hard, the sprays of hadrons are shaped into narrow jets. All the splitting and production of new quarks and gluons shuffles energy around, but the amounts of energy shuffled that way are much smaller than the initial kick the quarks get from the collision. So in the end, the direction of the jet reflects pretty closely the initial direction of the quark.

Of course, ‘pretty closely’ is not a very scientific term. We need to quantify that and be as precise as possible about it. Jet algorithms are the tools that let us do that. They give a recipe of how to combine the observed hadrons produced in a collision to get objects (jets) with energy and momenta that can be compared to a theoretical prediction. You can imagine many ways of doing that, but some ways are definitely better than others.

One issue in designing (or choosing) a good jet algorithm is the fact that the theory really doesn’t know how to predict what happens at low energies. These low energies correspond to (relatively) long distances, where the strings are snapping, hadrons are being formed and lots of low-energy gluons get thrown around. Since you can’t predict how many of these low-energy gluons might be produced – and since anyway that is not something we can measure, or want to measure, and it will fluctuate a lot – having a jet finder that is insensitive to the number of low-energy gluons seems like a good idea. In fact it is essential. The jargon for this insensitivity is ‘infrared safety’. One of the things we did while the LHC was under repair was switch to an infrared-safe jet algorithm.

1.6 Names, Inertia and the Media

In 2008, the main jet algorithm used in ATLAS (and CMS) analysis code was not infrared safe. After the breakdown, we managed to change this to a newer, better, infrared-safe algorithm. This would make a huge difference to the quality of the physics we could do later.

You might wonder why, if it made such a big difference, we didn't change earlier. It's an interesting question, and the answer tells you something about doing science in very large collaborations (as well as something about physics).

Back when I was still in Hamburg working on HERA, LEP was still running at CERN and the Tevatron in Chicago was finding the top quark, proto-collaborations had formed to design and propose possible detectors for the coming LHC. ATLAS was formed out of two of these – EAGLE and ASCOT. This kind of thing often happens. It is very important not to use up your best collaboration name on the first proposal, since you will almost certainly have to merge with some other proposal and therefore have to pick a new name at some point. I can only presume CMS made this mistake. Probably their proto-collaborations were called cool things like TITAN or JOR-EL[11] but they had to merge so often they ran out of ideas and ended up as CMS.

Anyway, ATLAS is a good name, so well done, Peter Jenni et al.

Marginally more important than choosing the name is showing that the detector will be able to do the physics required. Will it have a good enough resolution? Will there be enough bandwidth to read out the data? Will there be enough detectors in the right places to measure everything you want to see? And will you be able to afford it? To answer these questions convincingly, you have to write a technical design report. To do this and persuade people that you should actually be allowed to build your ideas, you need lots of results from test beams, where you fire particles into prototypes to show you understand them. But you also need huge amounts of software, some of it to simulate what the physics and the detector might

---

[11] They weren't.

look like, and some of it to reconstruct from the simulated data (or from test-beam data) what the measurement might eventually look like.

This means that by the time first data seem imminent, there are people who have been working on the experiment for ten or more years, and they have grown used to using the tools that were available ten years earlier. Attempts to change this will be met with huge inertia, even when there are much better tools available and when physics has moved on a lot due to data from elsewhere during those ten years. When the priority is to get everything ready for first data, the reluctance to change or reinvent anything is understandable. This is the state we were in on ATLAS in 2007.

The jet algorithm is just such a tool. The understanding of jets and the strong interaction, QCD, improved greatly in the 1990s and 2000s, due to a lot of work by theorists and a lot of data, mainly from HERA and LEP. The problems with infrared safety were understood and a new generation of jet finders was proposed. Unfortunately, ATLAS (and also, to a large extent, the Tevatron experiments) had already started using the old jet algorithms. And the new algorithms had problems: some of them were too slow, and lots of them made jets with irregular shapes, which made understanding the experimental resolutions and efficiencies harder. The HERA, LEP and Tevatron experiments did make measurements with the new algorithms (in fact HERA and LEP eventually moved over to them completely), but there were doubts about whether they could really be used at the LHC, and these doubts, combined with both inertia and the pressure of time, meant we did not change in time for 2008.

Of course, everyone knew that once we'd taken some data with the old algorithms it would only get more difficult to change to the new ones. So once we had an unlooked-for extra year before data arrived, it was a golden opportunity to switch, once and for all, to the new, better technology. Crucial, too, was the fact that the problems with speed and irregular-shaped jets seemed to have been solved.[12] Coordinated by our jet conveners, dozens of postdocs and students began checking whether the new algorithm really worked, not just theoretically but in the software ATLAS would use to select and analyse the data. Some of those people were working at UCL, and I edited the enormous internal note that, taken as a whole, finally persuaded everyone to make the switch. I think this was time very well spent. It certainly cheered me up a lot.

For me, another galvanising event, not directly related to the LHC, was a meeting on 18 May of the Skeptics in the Pub at Penderel's Oak in Holborn. This was the first Skeptics in the Pub meeting I'd been to. It's a bit of a random thing to throw into a more-or-less chronological run-through of the LHC story, but it does connect, trust me.

Simon Singh, a well-known science writer with a PhD in particle physics – he worked on LEP at Cambridge with some of my present ATLAS colleagues – was being sued by the British Chiropractic Association (BCA) for writing that it was happily promoting what he felt were 'bogus treatments' without taking account of the lack of reliable evidence showing the treatments worked. Since Singh clearly believed the treatments were bogus and since the BCA was obviously promoting them, you might have thought the argument would hinge on how happy the

---

[12] For those who want to know the specific science, the best solution turned out to be the so-called anti-k-T algorithm by Matteo Cacciari, Gavin Salam and Gregory Soyez (http://arxiv.org/abs/0802.1189).

association was. Unfortunately, British libel law was a complete ass and it looked as though Singh, in order to defend an honestly made comment, was going to have to prove that the BCA was knowingly lying; that meant that so long as they could show that they didn't believe the treatments were bogus they would be home and dry. Finally, Singh persuaded the court of appeal that he didn't have to prove the BCA knew the treatments to be bogus, just that *he* had good reasons for contending that they were. The BCA withdrew its legal action at that point.

I went along to the meeting partly out of a sense of outrage that the law could be used to effectively silence a science-based critique, and partly just to have a beer with some friends who were going. Holborn is very near the UCL campus in Bloomsbury. I got a lot more out of this than beer. I met a whole bunch of intelligent, diligent and well-informed journalists working in newspapers and broadcasting, and I have to say this came as a bit of a shock. I met several excellent writers who wrote blogs and other online stuff. And I became aware of a set of people who knew and cared more about science and rationality than I, in my arrogance and ignorance, had expected to find outside the world of scientific academia I knew. This was as heartening, and as exciting, as the libel situation was depressing and dangerous. This wasn't just about Singh. Dr Peter Wilmshurst was being sued for analysing the health outcomes of a heart implant, and others have written better than I can about the chilling effect bad libel laws can have on scientific discussion when exploited by rich people or big companies. When Singh announced he was going to fight on, he appeared brave but unlikely to succeed. However, there was to be a campaign not just to help him, but to get the law changed, and there was at least some hope that it would be successful.

Some of the writers became good friends, and the campaign led to a measure of success in reforming the law – the new libel laws, which included some specific protections for peer-reviewed science, came into effect at the start of 2014. And Singh saw off the BCA. All of that was good, but looking back on this period, there was an additional and important benefit for me.

The campaign helped to break down some of the barriers that exist between the world of science, the world of academia (which overlaps with but is very different from the world of science) and other worlds, including those of the media, comedy and politics. This was a new experience for me, and proved rather important in influencing how I talked to people about what we were doing at CERN. Brian Cox, who was a friend and collaborator before he stormed those barricades with huge success, made it easier for us all (and indeed he was at the Penderel's Oak meeting), but there was more to it than that. While I only participated in the libel-reform campaign in a very small way by writing letters and attending meetings, select committees and the like, I still made connections and gained an understanding of those other worlds in the process.

This removed a lot of my fear and suspicion of the media and politics, and also showed how lobbying could be done effectively. Along with a 'Science Blogging Talkfest' organised in 2010 by Alice Bell (then at Imperial) and Beck Smith of the Biochemistry Society, the libel-reform process did a huge amount to make me more comfortable with discussing science in public and via the media.

This permeability between different worlds seems essential to me, and while it requires some face-to-face interaction (and for me, at least, alcohol helps with this), online tools such as blogs and social media also have a very positive impact. It was after Alice and Beck's talkfest that Alok Jha, who was one of the panellists, asked me to join a new *Guardian* science-blogging initiative.

Exchanging ideas on Twitter in particular is something that has helped me a lot. Having a blog and a Twitter account that are read by a fairly large number of my colleagues in science and by contacts in the media turned out to be invaluable for my confidence when talking to bigger audiences elsewhere. Months later, I was sitting in a taxi on the way to be interviewed by John Humphrys on BBC Radio 4's *Today* programme about why the Higgs search was a waste of time. Knowing that if I said something stupid, or was misrepresented, I could at least tweet or blog about what I had really meant to say was a great comfort. As a scientist, it is a daunting thing to pop up briefly in the media and be given a few minutes, more or less at the mercy of professionals, to describe your work and why it is worthwhile. To have the means to reach people directly makes a big difference.

It turned out for me that in the particular example of the *Today* programme, John Humphrys was great and I didn't need to apologise for anything. But I still love that the tool invented by Tim Berners-Lee in the office downstairs from mine at CERN is supporting, amongst many other things, improved engagement between the science of CERN and the rest of the world.

Glossary: Bosons and Fermions

The word boson causes no end of trouble when people report on the search for the Higgs boson. It regularly gets spelled or pronounced 'bosun', and once, while being interviewed on TV and reading the autocue over the presenter's shoulder, I saw it spelled as 'bosom'. Like a true professional, Krishnan Guru-Murthy said 'boson' without blinking.

Boson is the name for a generic class of particles. The Higgs boson is one, but so are many other particles. In the Standard Model, all the particles that carry forces – gluons, the W and the Z, photons, plus the graviton, if there is one – are bosons.

Quarks, electrons and neutrinos, on the other hand, are fermions.

The difference between them is just spin. But in this context, spin is a quantum of angular momentum. It is a bit like the particle is spinning, but that is really just an analogy, since point-like fundamental particles could not spin, and anyway fermions have a spin such that in a classical analogy they would have to go round twice to get back to where they started. Quantum mechanics is full of semi-misleading analogies like this.

Regardless, spin is important. Bosons have, by definition, integer spin. The Higgs has zero, the gluon, photon, W and Z all have one, and the postulated graviton has two units of spin. Quarks, electrons and neutrinos are fermions and all have a half unit of spin. This causes a huge difference in their behaviour.

The best way we have of understanding fundamental particles is quantum field theory. In quantum field theory, a 'state' is a configuration describing all the particles in a system (say, a hydrogen atom). The maths is such that if you swap the places of two identical fermions, with identical energies (say, two electrons), then you introduce a negative sign in the state. If you swap two bosons, there is no negative sign.

Since swapping two identical particles of the same energy makes no physical difference to the overall state, you have to add up the two different cases (swapped and unswapped) when calculating the actual probability of a physical state occurring.

Adding the plus and the minus in the fermion case gives zero, but in the boson case they really do add up. This means any state containing two identical fermions of the same energy has zero probability of occurring. Whereas a state with two identical bosons of the same energy has an enhanced probability.

This fairly simple bit of maths is responsible for the periodic table and the behaviour of all the elements. Chemical elements consist of an atomic nucleus surrounded by electrons. Because electrons are fermions, not all the electrons can be sucked into the lowest energy level around the nucleus. If they were, the probability of that ‘state’ happening would be zero, by the argument above. So as more electrons are added around a nucleus, they have to sit in higher and higher energy levels – less and less tightly bound to the nucleus. The behaviour of a chemical element – how it reacts with other elements and binds to form molecules, and where it sits in the periodic table – is driven by how tightly bound its outermost electrons are.

When bosons clump together they do some fascinating stuff too. The condensate that is responsible for the superconductivity in the LHC magnets, for instance. But it’s hard to beat being responsible for the whole of chemistry, and therefore biology. And the rest.

Some theories extend the Standard Model by relating force-carrying bosons to matter-particle fermions. They do this by introducing a new symmetry between them. This symmetry is so mathematically compelling that it is called ‘super’ – supersymmetry – but that’s another story.

1.7 Boost One

Something else that we had time for, while the LHC was being fixed, was a lot more study with simulated data.

To be honest, by now a lot of us were pretty sick of this. I'd left the analysis of real data when I left the ZEUS experiment in 2005, and had been doing little but simulation studies since. There's a joy in a well-written program, and a lot you can learn from modelling things in computer code. But without actual experimental results it all begins to feel a bit self-referential.

We did now have better software tools, some of which had been improved using information gleaned from the few 'beam splash' events we had recorded before the LHC broke down, and also by studying cosmic-ray data taken with the full ATLAS detector. Cosmic rays are particles from space that bombard the Earth constantly.[13]

And, very specifically, Adam and I could do some more work on our fancy boosted Higgs analysis. I mentioned this in section 1.3; now may be a good time to expand on it.

The *Colliding Particles* films I mentioned, made by Mike Paterson and featuring Gavin, Adam, me and a cast of dozens, were popular with teachers. They show a lot about how particle physics is done. As intended, they seemed to be doing a decent job of hitting the ever-changing National Curriculum on 'how science works'. But even though they are partly based around our scientific paper about one way we might find the Higgs boson,[14] they do not contain much actual physics. There was a

---

[13] See 2.2 Minimum Bias and 4.2 Science Board for more about them.

[14] Referred to in the films as the 'Eurostar paper', as a sort of code for an initially underground London/Paris collaboration.

review of the films in *Physics World*, the magazine of the Institute of Physics, which gently pointed this out.

Fair comment, which I felt ought to be addressed. So I wrote my first ever blog post in an attempt to explain the physics behind the paper. Here is a rewritten version of that, which concentrates on explaining the new ideas in the paper rather than giving a summary of why the Higgs is interesting, which will come later . . .

What we knew at the point the paper was written was that if the Higgs boson existed, and if its mass was what seemed to be the most likely value (around 120 GeV[15]), then lots of Higgs bosons would be produced at the LHC. The difficult trick would be to pick them out from all the other things that would be going on.

The Standard-Model Higgs boson is a wave, or, more correctly, an excitation, in a quantum field[16] that fills the universe. Interacting with this field is what gives the other particles mass. Particles get different amounts of mass depending on the way they ‘couple’, or stick, to the field. This meant that the Higgs boson itself would couple to the mass of every particle, and would therefore be very likely to decay into the heaviest particle it could. These decays would happen super-quickly, so all we would see would be the things into which it had decayed. We would have to work out from them whether or not there was, briefly, a Higgs boson produced by the LHC.

For a 120 GeV Higgs, the heaviest thing it could decay into

---

[15] GeV stands for gigaelectronvolts, the particle physicist’s preferred unit of both mass and energy. I will say more about what they mean in section 2.1.

[16] See Glossary: Fields, Quantum and Otherwise (pp. 70–3).

would be a pair of bottom (b) quarks. Quarks come in pairs. Down and up quarks are all you really need to make protons and neutrons, but for some not fully understood reason we also have strange and charm quarks (similar to the down and up respectively, but more massive) and bottom and top. The strange quark was so named because the first evidence for it was particles with strange properties (basically living longer than expected and decaying oddly) seen in cosmic-ray interactions. 'Charm' seems a very whimsical name, but I guess the quark was quite charming when it turned up because it solved some tricky issues with the weak interaction. When the last two quarks were discovered, the names beauty and truth were proposed, I guess extrapolating from 'charm'. This shows you how unreliable extrapolation can be . . . Anyway, everyone calls them top and bottom now. Top only showed up at Fermilab in 1995. This was a great relief as it finally laid to rest the theoretical models in which the bottom quark had no partner, usually giving rise to tediously suggestive names – bare-bottom models, topless models, you get the idea . . .

I digress. The Higgs, if it were to have a mass of around 120–130 GeV, would not have enough energy to decay to top quarks, or to W and Z bosons, and so would mostly decay to bottom.

Bottom quarks also then decay (after travelling a few 100 microns or so) and each one gives a spray, or jet, of hadrons. We could see some of these particles in our detectors, and would be able to reconstruct the fact that two bottom quarks decayed, and that they came from a Higgs decay.

So: if a Higgs boson were to be produced, it would decay to two b quarks, and they would give two jets of hadrons.

The problem is, lots and lots of b quarks and jets of hadrons are produced at the LHC, and most of them have nothing to do with a Higgs. Before our paper, it looked like this background noise would completely swamp the signal and we would have to rely on other, rarer, Higgs decays to find it. Not only did this make finding the Higgs harder, but even if you found the Higgs some other way, seeing the decays to b quarks is actually pretty important in proving that whatever you might have found really was a Standard Model Higgs boson.

The idea we had in the paper was to look at those collision events where a Higgs is not just made, but is made and given a lot of kinetic energy – i.e. it is moving very fast, at speeds that are a substantial fraction of the speed of light. This happens in about 5 per cent of the Higgs production events we were simulating (assuming the LHC design beam energy of 14 TeV), so we would miss a lot of Higgs bosons by only looking at the fast ones. In fact, with the LHC starting at lower energies we would waste even more of them. But the advantage of this approach would be that we would lose still more of the backgrounds, the non-Higgs events, because the background jets usually have a lower energy.

Something happens when you look at these fast-moving Higgs decays. The faster the Higgs moves, the smaller the opening angle between the two b quarks it produces when it decays. In fact, very often the jets from the two quarks merge into a single jet.

This is a problem if, as had been done in all the studies until then, you look for two b-quark jets as your telltale signature that a Higgs boson was there. In our paper we turned this problem into an advantage. By looking at the internal structure of this single jet, at its substructure, we could see evidence for the two b quarks and the Higgs decay, get rid of even more background, and measure the mass of the Higgs boson well enough to make it stand out over the remaining background. What had looked like a hopeless case was recovered as a promising way of finding the Higgs at the LHC.

The idea of looking at jet substructure to find the decays of fast-moving particles had been put forward earlier by Mike Seymour,[17] and again in the paper[18] I mentioned in 1.2, which I wrote with Brian and Jeff. But Gavin had a better way of doing it. And together with Adam and Mathieu we were the first to apply these ideas to the Higgs boson.

The unexpected delay of a year meant that – with Erkcan Ozcan, a UCL postdoc, and a group from Freiburg – Adam and I could try it out on a fully simulated ATLAS detector. We found it should still work pretty well even with a more realistic estimate of the experimental errors and a more complete study of the backgrounds.

A sign that these ideas were picking up fans was that out of the blue, Gavin and I received an invitation to go to the SLAC National Accelerator Laboratory to talk about them at a meeting called 'Giving New Physics a Boost'. Or just 'Boost'. This was July 2009, and there were by now quite a few papers coming out developing new ways of finding the decays of boosted particles and highlighting applications at the LHC. If you have enough energy in your collider, boosted heavy particles turn up quite often. So, even though London to SLAC is a long way to travel for a two-day meeting, it was worth the trip.

SLAC is in Menlo Park, California, about 40km from San Francisco. I'm not good at long-haul flights and the two days

---

[17] See http://inspirehep.net/record/359650?ln=en.

[18] See http://arxiv.org/abs/hep-ph/0201098.

passed in a bit of a daze, to be honest, though from the discussions it was clear that exciting physics was happening, to the extent that a longer, follow-up meeting was already planned for 2010 in Oxford. I remember sitting on the seafront in San Francisco eating clam chowder from a bread bowl, waiting for the flight home and wondering what had just happened. Whatever it was, it was good.

Anyway, by the autumn of 2009 the LHC was repaired, at least to the extent that CERN was confident it could get the beams up to half the design energy – still a factor of 3.5 higher than any previous experiment. Slightly chastened but wiser, we were ready to try again for real data.