

THE BIG REBOOT

As the Large Hadron Collider prepares to come back to life after a two-year hiatus, physicists are gearing up to go beyond the standard model of particle physics.

Mike Lamont grabs the last croissant from a table and eats it as he walks through the control centre at CERN, the European laboratory for particle physics just outside Geneva, Switzerland. It is mid-morning, and the vast blue room is full of physicists staring into computer screens. Lamont, the operations manager of CERN's beams department, explains that they are running tests to ensure that an unexpected computer outage would not affect the network of electronics, vacuum pipes and superconducting magnets that comprise CERN's Large Hadron Collider (LHC), the most powerful particle accelerator in the world.

BY MATTHEW CHALMERS

This is one of numerous checks that are helping Lamont and his colleagues to sleep better at night. They are nearing completion of a major refurbishment that has been under way since March 2013. They have already started cooling down the accelerator's 27-kilometre ring of superconducting magnets in preparation for a restart next year. But when the LHC comes back to life, circulating its twin beams of protons in opposite directions around the ring, Lamont and his colleagues will be pushing to get as close as possible to its design energy of 7 trillion electron volts (TeV) per beam — nearly twice what the LHC has managed so far. Each beam will pack as much energy as a speeding freight train.

Lamont knows all too well what can happen if things go wrong. He was here in September 2008, when the team last attempted to ramp up the US\$5-billion collider to such energies — and ended up triggering an electrical fault that knocked it out of commission for more than a year and cost tens of millions of dollars to repair.

"We've learned a lot about the machine since then," says Lamont. The researchers managed to patch it up and get it working again by the end of 2009, although they ran it at only half its design energy to avoid another shutdown. Still, that was enough for collisions between the beams to produce conclusive evidence for the long-sought Higgs boson — the last unconfirmed prediction of the 40-year-old standard model of particle physics, which describes the behaviour of every particle and force known except gravity.

But for all the acclaim that greeted the announcement of the Higgs particle in July 2012 — and the 2013 Nobel Prize awarded to the theorists who first conjectured its existence — there is much more that the LHC physicists hope to learn from the machine's upcoming run. Is the newly discovered Higgs the only particle of its kind, as the standard model predicts, or is it just the lightest member of a whole family? If there are more Higgs particles, some of them might appear at higher collision energies. Or perhaps the high energies will produce other new, exotic particles that have no place in the standard model.

Theorists have been predicting such particles for decades. Supersymmetry, an extension of the standard model first proposed in the early

1970s, holds that each particle has a heavier counterpart or 'sparticle', and that the two differ in predictable ways.

One or more such sparticles might turn out to be the constituents of dark matter — an invisible haze that is massive enough to control the motion of galaxies, but that is unaccounted for in the standard model. Finding these sparticles, assuming that they are not too heavy to be produced at LHC energies, will thus be a prime goal for the refurbished machine. It might even turn up still more exotic results, such as evidence for spatial dimensions beyond the familiar three. But first, Lamont and his team have to get the LHC running at full capacity.

CLICKS AND HISSES

After a short drive from the control centre, Lamont dons a helmet, steel-capped boots and emergency breathing equipment, then steps into an elevator for a journey 100 metres underground. The elevator doors open onto a service corridor. From there, a short walk leads to the LHC tunnel, where a string of cylindrical, bright-blue magnets curves gently into the distance.

Even after 25 years at CERN, Lamont says that he is still in awe of the power and complexity of the machine. It seems light years away from the cerebral calm of the control room. Down here, the LHC hums, clicks and hisses, and its tunnel smells of metal, dust and warm circuitry. The 15-metre-long, 35-tonne magnets are held off the concrete floor by heavy-duty jacks, and are packed with intricate wiring and plumbing that encases the airtight beam pipes running through their centres. To avert another short circuit, the LHC has been fitted with sensors and thousands of kilometres of cable to detect the slightest sign of a surge in voltage. Crucially, 10,000 of the superconducting connectors that link the magnets have been reinforced or replaced — a task that took more than 250 people more than a year to complete.

Since June, the team has been cooling the magnets down towards their operating temperature of 1.9 kelvin — at which the current-carrying cables that generate the magnetic fields become superconductive. To keep this process manageable, the LHC ring is divided into eight sections that can be refrigerated independently. Once the magnets have been chilled, which takes two months per section, the team will carry out electrical tests to ensure that they can operate at high energy. Lamont already knows that things will not go smoothly. There is a batch of magnets that performed perfectly in tests above ground, but, for some reason, 'quench' or lose their superconductivity when they reach magnetic fields equivalent to beam energies of just 6.5 TeV. This is not a disaster — fixing the magnets is just a matter of cycling each one through several quenches until it settles down and works properly. But it does take time, he says. "And there are hundreds of the buggers!"

Eventually, however, proton beams will once again be threaded



ILLUSTRATION BY ANDY POTTS

through the LHC — a milestone now scheduled for March 2015. And then, after another few weeks of testing, the physicists will start carefully steering the beams into collisions and checking that it is safe for the detectors to start data collection.

There is a faint smell of burning in the tunnel. Lamont explains that a vacuum pipe is being heated to drive off stray molecules. He walks past the magnets to a point where the bare beam line plunges through a massive copper and steel wall. On the other side lies ATLAS, one of the LHC's four main particle detectors (see 'The refurbished ring'). Soon, bunches of high-energy protons will be fired past this point into ATLAS, where they will slam into equally energetic protons circulating in the other direction and send collision debris streaming outwards through the detector. "You think: Jesus, they let us steer a beam through here?" says Lamont. "I still can't believe I get paid to play around with this thing."

SYSTEM UPDATES

Some 8.5 kilometres away from ATLAS, on the opposite side of the LHC ring, Tiziano Camporesi stares up at the 12,500-tonne Compact Muon Solenoid (CMS) — and marvels at the audacity of the physicists who designed it 30 years ago. "They must have been crazy," he says. Many people declared the machine — a vast cylinder with concentric layers of silicon particle sensors, superconducting magnets and massive iron 'yokes' to contain the magnetic field — too intricate to ever work. But it did, says Camporesi, and "far better than we ever expected". CMS and ATLAS were the detectors that identified the Higgs boson in 2012.

"YOU THINK: JESUS, THEY LET US STEER A BEAM THROUGH HERE? I STILL CAN'T BELIEVE I GET TO PLAY AROUND WITH THIS THING."

Camporesi, who was elected spokesman of the 3,800-strong CMS collaboration earlier this year, is now coordinating activities for next year's high-energy run. His team, like those on all of the LHC's main experiments (which also include the more specialized ALICE and LHCb detectors located elsewhere around the ring) have been making some much needed repairs and upgrades during the downtime. They have had one piece of good news: in the central region of the detector, where the beams intersect and newly created particles

explode outwards from the collision point, the sensitive silicon trackers have so far survived without radiation damage. But CMS physicists have replaced several faulty photomultiplier tubes that were giving false results, making it seem as though an exotic new particle had been produced when none had.

Camporesi is particularly proud of the four disc-shaped chambers added to each end of the CMS to improve its ability to sense particles called muons. This upgrade, in turn, will beef up the detector's 'trigger' — a combination of electronics and software that will monitor the particles streaming through the detector from the collisions and look for patterns that could signify an event worthy of further study. Physicists have been using such triggers for decades, says Camporesi. But the LHC's next run will not only boost the energy of each beam, it will also increase the number of protons they carry. The result will be some 1–2 billion collisions per second in the CMS. Particles from one collision will still be making their way out of the detector while as many as 50 new collisions occur behind them. From all those events, the trigger

THE REFURBISHED RING

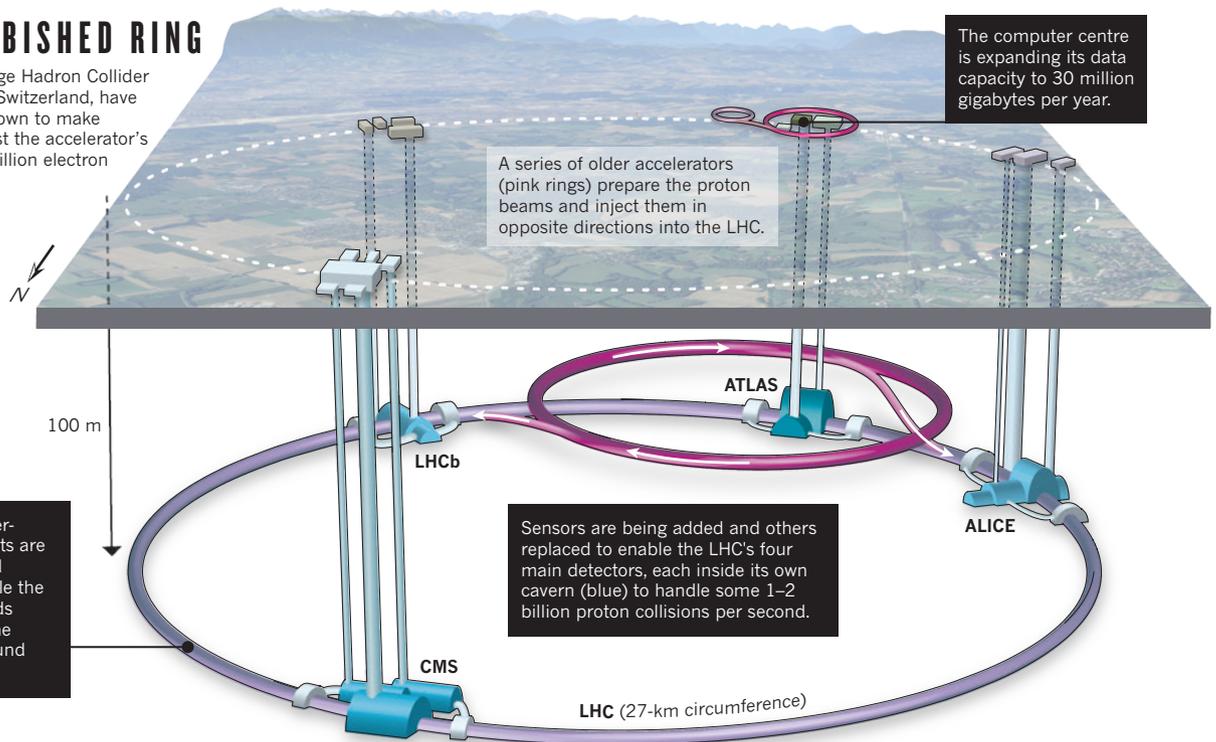
Physicists at the Large Hadron Collider (LHC) near Geneva, Switzerland, have used a 2-year shutdown to make repairs that will boost the accelerator's energy to nearly 7 trillion electron volts per beam.

The computer centre is expanding its data capacity to 30 million gigabytes per year.

A series of older accelerators (pink rings) prepare the proton beams and inject them in opposite directions into the LHC.

Thousands of superconducting magnets are being repaired and reinforced to handle the high magnetic fields needed to guide the proton beams around the LHC's ring.

Sensors are being added and others replaced to enable the LHC's four main detectors, each inside its own cavern (blue) to handle some 1–2 billion proton collisions per second.



will have to decide which to store for further analysis; the goal is to bring the final recorded event rate down to several hundred per second. “It’s occupying a lot of our time right now,” says Camporesi.

BIG DATA CHALLENGE

Once the refurbished LHC is running again, the raw electronic signals from CMS and the other detectors will flow back to the main CERN campus through optical fibres that link directly into the laboratory’s computing centre — a stuffy, windowless room in which row upon row of racks hold some 100,000 processors, and cooling fans work noisily to control the heat.

The processors will analyse the incoming data with algorithms that determine the identity, energy and direction of each particle emerging from each collision. The results will then be stored on magnetic tape — an old-fashioned medium that has the advantage of being cheaper and more durable than digital storage.

But just storing the information at CERN would not satisfy researchers’ near-insatiable appetite for it. Today’s particle physicists spend most of their time writing thousands of lines of computer code designed to search millions of collisions for unusual signals. To get data to these researchers, CERN set up the Worldwide Computing Grid, in which the computer centre streams copies of the data to 13 ‘tier-1’ computer centres worldwide. These centres, in turn, are connected to more than 150 smaller computer clusters called tier-2 nodes, most of which are at universities.

Fortunately for the end users, they do not have to know about any of this. A physicist just has to submit a program to the grid and specify which collision events are to be examined. Grid software will then automatically shunt the job to a centre that has enough processing power and disk space free to run it, then return the results (see *Nature* 469, 282–283; 2011). On this particular day at the CERN computer centre, a live screen shows that 10,500 programs are running at this facility alone — and it represents only 6% of the grid’s resources. Were it not for the grid, says Jeremy Coles, a physicist at the University of Cambridge, UK, who is the grid’s UK coordinator, his colleagues would probably still be searching for the Higgs boson.

The challenge for the future, says Coles, is dealing with the sky-high event rates to come. During the LHC’s first run, despite radical pruning by the detectors’ triggers, the data still piled up at the rate of 15 petabytes

(15,000 trillion bytes) a year — more than in all the videos uploaded to YouTube annually. When the LHC starts up again next year, the doubling of the collision rate will bump that up to roughly 30 petabytes per year — an average of about 1 gigabyte per second.

Coles is confident that the grid will be able cope with that increase — not least because of technological advances that have enabled much closer integration between computer centres. “Networking has come on very quickly in the past 10 years,” he says — “more than we thought.” Just last year, for example, CERN expanded the capability of its data centre, which is at the limits of available space and power, by linking it up with a facility in Budapest through two fibres that transfer data at 100 gigabits per second. From an operations point of view, says Coles, it is just like having the machines in the room next door.

But the data onslaught will not stop there. Planned upgrades to the LHC will send its data output soaring to 110 petabytes a year by the early 2020s, and, eventually, to as much as 400 petabytes a year. “We currently have no way to deal with this,” says Coles. Making matters worse, computer chip speeds are stagnating. The best commercial chips now available often contain two, four or eight processors to boost their power. Future chips are likely to have even more. But the code that runs the LHC was written to be run on one processor at a time, says Coles. Finding ways to run data analysis on many processors in parallel would mean rewriting 15 million lines of code written by thousands of people over many years.

Still, when CERN physicists needed a better way to share information in the late 1980s, the result was the World Wide Web. And when they needed a better way to access computer resources back in the 1990s, they invented the world’s largest computing grid. The LHC scientists seem confident that they will invent a way around this problem as well.

Lamont seems to feel similarly confident when he talks about the ‘next big collider’ in particle physics. Although CERN just celebrated its sixtieth anniversary, and the LHC still has another 20 years of proton smashing to go, the laboratory is exploring the feasibility of a collider 80–100 kilometres around that would drill even deeper into the structure of matter (see *Nature* 503, 177; 2013). Lamont says that he would be lucky to see such a machine built during his lifetime, but he points out that the LHC, which came online for the first time in 2008, was first sketched out in 1984. “We’ve got to start thinking about the next machine now,” he says. ■

Matthew Chalmers is a freelance writer in Bristol, UK.