



CORNUCOPIA of familiar particles spraying out from each collision will include, just occasionally, something new and wonderful.

THE DISCOVERY MACHINE

A global collaboration of scientists is preparing to start up the greatest particle physics experiment in history **By Graham P. Collins**

You could think of it as the biggest, most powerful microscope in the history of science. The Large Hadron Collider (LHC), now being completed underneath a circle of countryside and villages a short drive from Geneva, will peer into the physics of the shortest distances (down to a *nano*-nanometer) and the highest energies ever probed. For a decade or more, particle physicists have been eagerly awaiting a chance to explore that domain, sometimes called the terascale because of the energy range involved: a trillion electron volts, or 1 TeV. Significant new physics is expected to occur at these energies, such as the elusive Higgs particle (believed to be responsible for imbuing other particles with mass) and the particle that constitutes the dark matter that makes up most of the material in the universe.

The mammoth machine, after a nine-year construction period, is scheduled (touch wood) to begin producing its beams of particles later this year. The commissioning process is planned to proceed from one beam to two beams to colliding beams; from lower energies to the terascale; from weaker test intensities to stronger ones suitable for producing data at useful rates but more difficult to control. Each step along the way will produce challenges to be overcome by the more than 5,000 scientists, engineers and students collaborating on the gargantuan effort. When I visited the project last fall to get a first-hand look at the preparations to probe the high-



KEY CONCEPTS

- The Large Hadron Collider (LHC), the biggest and most complicated particle physics experiment ever seen, is nearing completion and is scheduled to start operating this year.
- The LHC will accelerate bunches of protons to the highest energies ever generated by a machine, colliding them head-on 30 million times a second, with each collision spewing out thousands of particles at nearly the speed of light.
- Physicists expect the LHC to bring about a new era of particle physics in which major conundrums about the composition of matter and energy in the universe will be resolved.

—The Editors

SLIM FILMS

FAST FACTS

PROTON VELOCITY:
99.9999991% of light speed

PROTONS PER BUNCH:
up to 100 billion

NUMBER OF BUNCHES:
up to 2,808

BUNCH CROSSINGS
PER SECOND:
up to 31 million, at 4 locations

COLLISIONS PER BUNCH
CROSSING:
up to 20

DATA PER COLLISION:
about 1.5 megabytes

NUMBER OF HIGGS PARTICLES:
1 every 2.5 seconds (at full
beam luminosity and under
certain assumptions about
the Higgs)

CONTROL of the LHC accelerator
will be conducted from one sec-
tion (lower right) of the CERN
control center. The detectors
have their own control rooms.

energy frontier, I found that everyone I spoke to expressed quiet confidence about their ultimate success, despite the repeatedly delayed schedule. The particle physics community is eagerly awaiting the first results from the LHC. Frank Wilczek of the Massachusetts Institute of Technology echoes a common sentiment when he speaks of the prospects for the LHC to produce “a golden age of physics.”

A Machine of Superlatives

To break into the new territory that is the tera-scale, the LHC’s basic parameters outdo those of previous colliders in almost every respect. It starts by producing proton beams of far higher energies than ever before. Its nearly 7,000 magnets, chilled by liquid helium to less than two kelvins to make them superconducting, will steer and focus two beams of protons traveling within a millionth of a percent of the speed of light. Each proton will have about 7 TeV of energy—7,000 times as much energy as a proton at rest has embodied in its mass, courtesy of Einstein’s $E = mc^2$. That is about seven times the energy of the reigning record holder, the Tevatron collider at Fermi National Accelerator Laboratory in Batavia, Ill. Equally important, the machine is designed to produce beams with 40 times the intensity, or luminosity, of the Tevatron’s beams. When it is fully loaded and at maximum energy, all the circulating particles will carry energy roughly equal to the kinetic energy of about 900 cars traveling at 100 kilometers per hour, or enough to heat the water for nearly 2,000 liters of coffee.

The protons will travel in nearly 3,000

bunches, spaced all around the 27-kilometer circumference of the collider. Each bunch of up to 100 billion protons will be the size of a needle, just a few centimeters long and squeezed down to 16 microns in diameter (about the same as the thinnest of human hairs) at the collision points. At four locations around the ring, these needles will pass through one another, producing more than 600 million particle collisions every second. The collisions, or events, as physicists call them, actually will occur between particles that make up the protons—quarks and gluons. The most cataclysmic of the smashups will release about a seventh of the energy available in the parent protons, or about 2 TeV. (For the same reason, the Tevatron falls short of exploring terascale physics by about a factor of five, despite the 1-TeV energy of its protons and antiprotons.)

Four giant detectors—the largest would roughly half-fill the Notre Dame cathedral in Paris, and the heaviest contains more iron than the Eiffel Tower—will track and measure the thousands of particles spewed out by each collision occurring at their centers. Despite the detectors’ vast size, some elements of them must be positioned with a precision of 50 microns.

The nearly 100 million channels of data streaming from each of the two largest detectors would fill 100,000 CDs every second, enough to produce a stack to the moon in six months. So instead of attempting to record it all, the experiments will have what are called trigger and data-acquisition systems, which act like vast spam filters, immediately discarding almost all the information and sending the data from only the most promising-looking 100 events each second to the LHC’s central computing system at CERN, the European laboratory for particle physics and the collider’s home, for archiving and later analysis.

A “farm” of a few thousand computers at CERN will turn the filtered raw data into more compact data sets organized for physicists to comb through. Their analyses will take place on a so-called grid network comprising tens of thousands of PCs at institutes around the world, all connected to a hub of a dozen major centers on three continents that are in turn linked to CERN by dedicated optical cables.

Journey of a Thousand Steps

In the coming months, all eyes will be on the accelerator. The final connections between adjacent magnets in the ring were made in early




November, and as we go to press in mid-December one of the eight sectors has been cooled almost to the cryogenic temperature required for operation, and the cooling of a second has begun. One sector was cooled, powered up and then returned to room temperature earlier in 2007. After the operation of the sectors has been tested, first individually and then together as an integrated system, a beam of protons will be injected into one of the two beam pipes that carry them around the machine's 27 kilometers.

The series of smaller accelerators that supply the beam to the main LHC ring has already been checked out, bringing protons with an energy of 0.45 TeV "to the doorstep" of where they will be injected into the LHC. The first injection of the beam will be a critical step, and the LHC scientists will start with a low-intensity beam to reduce the risk of damaging LHC hardware. Only when they have carefully assessed how that "pilot" beam responds inside the LHC and have made fine corrections to the steering magnetic fields will they proceed to higher intensities. For the first running at the design energy of 7 TeV, only a single bunch of protons will circulate in each direction instead of the nearly 3,000 that constitute the ultimate goal.

As the full commissioning of the accelerator proceeds in this measured step-by-step fashion, problems are sure to arise. The big unknown is how long the engineers and scientists will take to overcome each challenge. If a sector has to be brought back to room temperature for repairs, it will add months.

The four experiments—ATLAS, ALICE, CMS and LHCb—also have a lengthy process of completion ahead of them, and they must be closed up before the beam commissioning begins. Some extremely fragile units are still being installed, such as the so-called vertex locator detector that was positioned in LHCb in mid-November. During my visit, as one who specialized in theoretical rather than experimental physics many years ago in graduate school, I was struck by the thick rivers of thousands of cables required to carry all the channels of data from the detectors—every cable individually labeled and needing to be painstakingly matched up to the correct socket and tested by present-day students.

Although colliding beams are still months in the future, some of the students and postdocs already have their hands on real data, courtesy of cosmic rays sleeting down through the Franco-Swiss rock and passing through their detectors sporadically. Seeing how the detectors respond



to these interlopers provides an important reality check that everything is working together correctly—from the voltage supplies to the detector elements themselves to the electronics of the readouts to the data-acquisition software that integrates the millions of individual signals into a coherent description of an "event."

All Together Now

When everything *is* working together, including the beams colliding at the center of each detector, the task faced by the detectors and the data-processing systems will be Herculean. At the design luminosity, as many as 20 events will occur with each crossing of the needlelike bunches of protons. A mere 25 nanoseconds pass between one crossing and the next (some have larger gaps). Product particles sprayed out from the collisions of one crossing will still be moving through the outer layers of a detector when the next crossing is already taking place. Individual elements in each of the detector layers respond as a particle of the right kind passes through it. The millions of channels of data streaming away from the detector produce about a megabyte of data from each event: a petabyte, or a billion megabytes, of it every two seconds.

The trigger system that will reduce this flood of data to manageable proportions has multiple levels. The first level will receive and analyze data from only a subset of all the detector's components, from which it can pick out promising events based on isolated factors such as whether an energetic muon was spotted flying out at a large angle from the beam axis. This so-called level-one triggering will be conducted by hundreds of dedicated computer boards—the logic

MAGNET REPAIRS had to be carried out in 2007 after a design flaw came to light during a stress test.

CURIOSITIES

TILT!

The LHC's tunnel is tilted 1.4 percent from horizontal, to put as much of it as possible inside solid rock. It is about 50 meters deep on the Lake Geneva side and 175 meters deep on the other.

PHASES OF THE MOON

When the moon is full, at "high tide" the land near Geneva rises 25 centimeters, the LHC's circumference increases by 1 millimeter, and the beam energy changes by 0.02 percent. The experimenters must allow for this effect: they must know the beam energy to within 0.002 percent accuracy.


THE OCTAGON

The LHC tunnel is actually octagonal, with eight arcs connected by eight short, straight sections that harbor the four experiments and facilities related to controlling the beam.

[LHC AT A GLANCE]

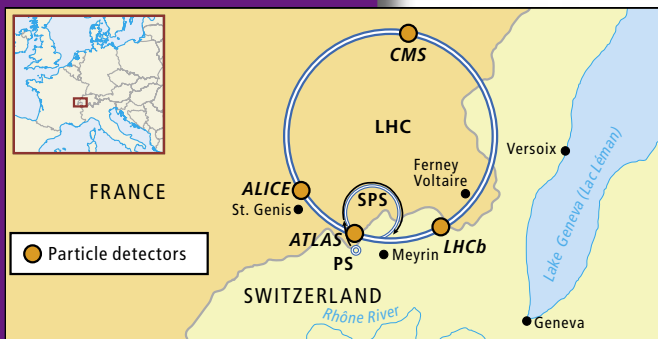
One Ring to Rule Them All

The Large Hadron Collider combines trusty old workhorses and trailblazing behemoths. Decades-old accelerators, including the Proton Synchrotron and the Super Proton Synchrotron, get the protons to 99.99975 percent of light speed. The LHC boosts the protons' energy by nearly 16 times and collides them 30 million times a second for up to 10 hours. Four major experiments generate more than 100 terabytes of collision data per second.



BRING YOUR PASSPORT

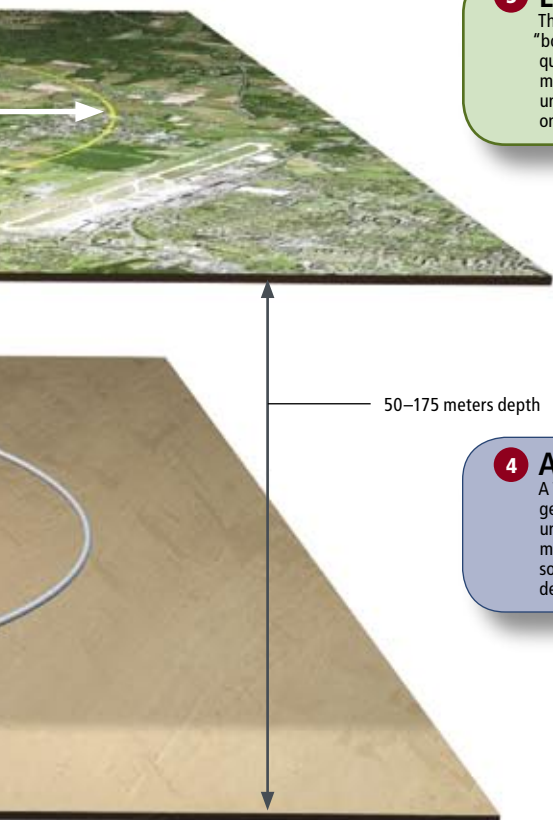
International in every respect, the LHC project involves scientists and funding from dozens of countries, including the 20 CERN member states in Europe, six observer states such as the U.S., Japan and Russia, and other countries such as Canada and China.



DON FOLEY (illustration); CERN (aerial image); SIMON NORFOLK NB Pictures/Contact Press Images (LHC tunnel); STEFANO DAL POZZOLO Contrasto/Redux (CMS and ATLAS); CERN GENEVA (LHCb and ALICE); MAPPING SPECIALISTS (map)

1 LHC ACCELERATOR


Nearly 7,000 superconducting magnets steer the proton beams around the tunnel that was dug for the Large Electron Positron (LEP) collider in 1989 and focuses them to a hair's breadth.



50-175 meters depth

2 CMS

The Compact Muon Solenoid is one of two huge general purpose detectors that will lead the search for particles such as the Higgs and other new phenomena. It has five "barrel wheels" like the one shown here as well as end caps.




3 LHCb

This detector looks for "beauty" (or "bottom") flavored quarks and anti-quarks, to understand how the mysterious absence of antimatter in the universe came about. It monitors only one side of its collision point.




4 ATLAS

A Toroidal LHC Apparatus is the other general purpose detector, with a unique design based on toroidal magnets instead of the traditional solenoid. The "big wheels" (right) detect key particles called muons.



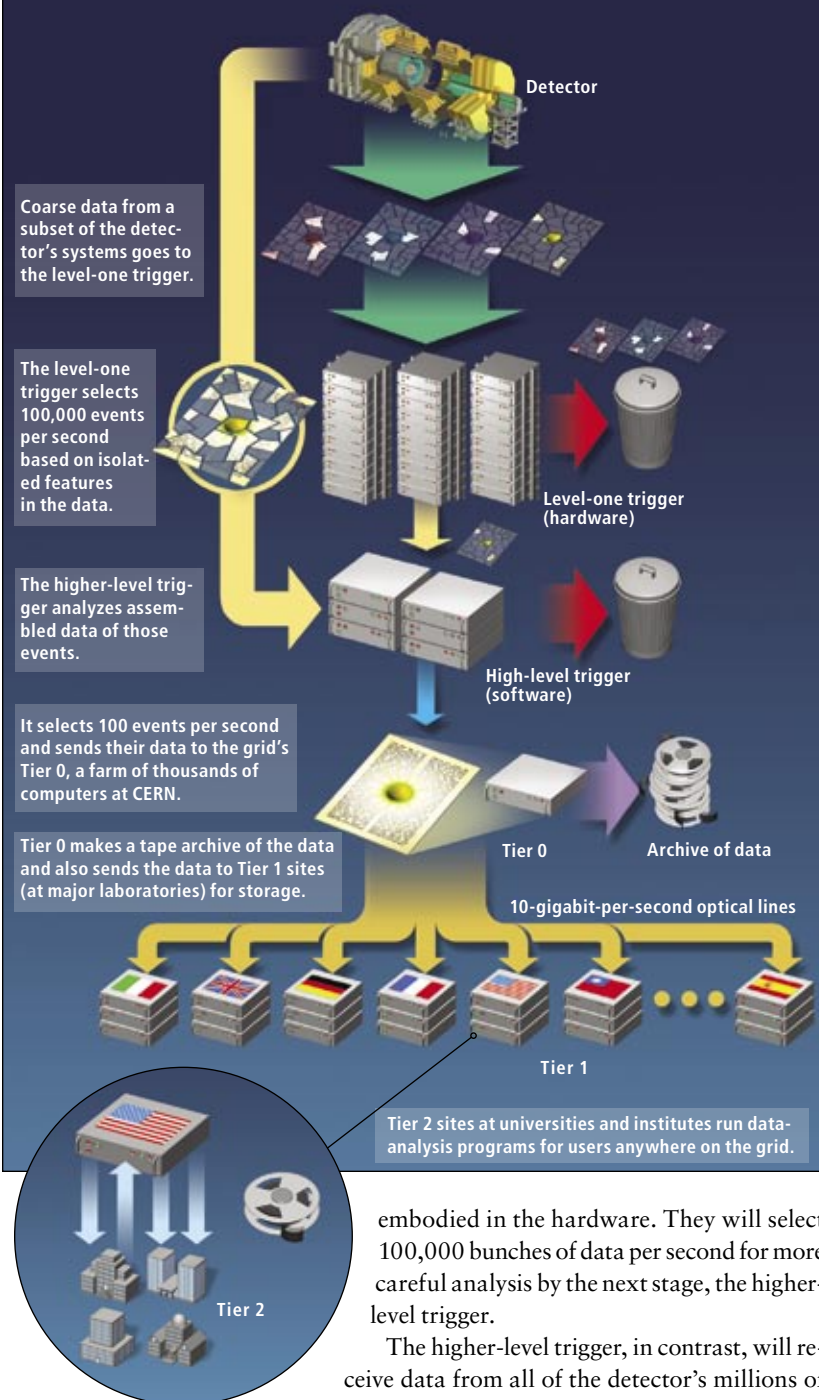
5 ALICE

A Large Ion Collider Experiment studies collisions of lead (Pb) ions producing primordial fireballs called the quark-gluon plasma. It also studies the proton-proton collisions as a reference point.



TOO MUCH INFORMATION

With up to 20 collisions occurring at 25-nanosecond intervals at the center of each detector, the LHC produces more data than can be recorded. So-called trigger systems select the tiny fraction of the data that has promising features and discard the rest. A global network of computers called a grid provides thousands of researchers around the world with access to the stored data and the processing power to analyze it.



embodied in the hardware. They will select 100,000 bunches of data per second for more careful analysis by the next stage, the higher-level trigger.

The higher-level trigger, in contrast, will receive data from all of the detector's millions of channels. Its software will run on a farm of computers, and with an average of 10 microseconds elapsing between each bunch approved by the level-one trigger, it will have enough time to "re-

construct" each event. In other words, it will project tracks back to common points of origin and thereby form a coherent set of data—energies, momenta, trajectories, and so on—for the particles produced by each event.

The higher-level trigger passes about 100 events per second to the hub of the LHC's global network of computing resources—the LHC Computing Grid. A grid system combines the processing power of a network of computing centers and makes it available to users who may log in to the grid from their home institutes [see "The Grid: Computing without Bounds," by Ian Foster; *SCIENTIFIC AMERICAN*, April 2003].

The LHC's grid is organized into tiers. Tier 0 is at CERN itself and consists in large part of thousands of commercially bought computer processors, both PC-style boxes and, more recently, "blade" systems similar in dimensions to a pizza box but in stylish black, stacked in row after row of shelves [see illustration on opposite page]. Computers are still being purchased and added to the system. Much like a home user, the people in charge look for the ever moving sweet spot of most bang for the buck, avoiding the newest and most powerful models in favor of more economical options.

The data passed to Tier 0 by the four LHC experiments' data-acquisition systems will be archived on magnetic tape. That may sound old-fashioned and low-tech in this age of DVD-RAM disks and flash drives, but François Grey of the CERN Computing Center says it turns out to be the most cost-effective and secure approach.

Tier 0 will distribute the data to the 12 Tier 1 centers, which are located at CERN itself and at 11 other major institutes around the world, including Fermilab and Brookhaven National Laboratory in the U.S., as well as centers in Europe, Asia and Canada. Thus, the unprocessed data will exist in two copies, one at CERN and one divided up around the world. Each of the Tier 1 centers will also host a complete set of the data in a compact form structured for physicists to carry out many of their analyses.

The full LHC Computing Grid also has Tier 2 centers, which are smaller computing centers at universities and research institutes. Computers at these centers will supply distributed processing power to the entire grid for the data analyses.

Rocky Road

With all the novel technologies being prepared to come online, it is not surprising that the LHC has experienced some hiccups—and some more


serious setbacks—along the way. Last March a magnet of the kind used to focus the proton beams just ahead of a collision point (called a quadrupole magnet) suffered a “serious failure” during a test of its ability to stand up against the kind of significant forces that could occur if, for instance, the magnet’s coils lost their superconductivity during operation of the beam (a mishap called quenching). Part of the supports of the magnet had collapsed under the pressure of the test, producing a loud bang like an explosion and releasing helium gas. (Incidentally, when workers or visiting journalists go into the tunnel, they carry small emergency breathing apparatuses as a safety precaution.)

These magnets come in groups of three, to squeeze the beam first from side to side, then in the vertical direction, and finally again side to side, a sequence that brings the beam to a sharp focus. The LHC uses 24 of them, one triplet on each side of the four interaction points. At first the LHC scientists did not know if all 24 would need to be removed from the machine and brought aboveground for modification, a time-consuming procedure that could have added weeks to the schedule. The problem was a design flaw: the magnet designers (researchers at Fermilab) had failed to take account of all the kinds of forces the magnets had to withstand. CERN and Fermilab researchers worked feverishly, identifying the problem and coming up with a strategy to fix the undamaged magnets in the accelerator tunnel. (The triplet damaged in the test was moved aboveground for its repairs.)

In June, CERN director general Robert Aymar announced that because of the magnet failure, along with an accumulation of minor problems, he had to postpone the scheduled start-up of the accelerator from November 2007 to spring of this year. The beam energy is to be ramped up faster to try to stay on schedule for “doing physics” by July.

Although some workers on the detectors hinted to me that they were happy to have more time, the seemingly ever receding start-up date is a concern because the longer the LHC takes to begin producing sizable quantities of data, the more opportunity the Tevatron has—it is still running—to scoop it. The Tevatron could find evidence of the Higgs boson or something equally exciting if nature has played a cruel trick and given it just enough mass for it to show up only now in Fermilab’s growing mountain of data.

Holdups also can cause personal woes through the price individual students and scien-



tists pay as they delay stages of their careers waiting for data.

Another potentially serious problem came to light in September, when engineers discovered that sliding copper fingers inside the beam pipes known as plug-in modules had crumpled after a sector of the accelerator had been cooled to the cryogenic temperatures required for operation and then warmed back to room temperature.

At first the extent of the problem was unknown. The full sector where the cooling test had been conducted has 366 plug-in modules, and opening up every one for inspection and possibly repair would have been terrible. Instead the team addressing the issue devised a scheme to insert a ball slightly smaller than a Ping-Pong ball into the beam pipe—just small enough to fit and be blown along the pipe with compressed air and large enough to be stopped at a deformed module. The sphere contained a radio transmitting at 40 megahertz—the same frequency at which bunches of protons will travel along the pipe when the accelerator is running at full capacity—enabling the tracking of its progress by beam sensors that are installed every 50 meters. To everyone’s relief, this procedure revealed that only six of the sector’s modules had malfunctioned, a manageable number to open up and repair.

When the last of the connections between accelerating magnets was made in November, completing the circle and clearing the way to start cooling down all the sectors, project leader Lyn Evans commented, “For a machine of this complexity, things are going remarkably smoothly, and we’re all looking forward to doing physics with the LHC next summer.” ■

THOUSANDS of processors at CERN are linked together to provide the computing power needed to manage the data as they flow from the experiments.

➔ MORE TO EXPLORE

The Large Hadron Collider. Chris Llewellyn Smith in *Scientific American*, Vol. 283, No. 1; pages 70–77; July 2000.

Discovering the Quantum Universe. Available at www.interactions.org/quantumuniverse/qu2006

CERN’s pages for the public are at <http://public.web.cern.ch/public>

Links to the LHC experiments are at http://lhcb.web.cern.ch/lhc/LHC_Experiments.htm

US LHC Blogs. Monica Dunford, Pamela Klabbers, Steve Nahn and Peter Steinberg. Available at www.uslhcb.org/blogs

Track the daily status of the accelerator commissioning at <http://lhcb.web.cern.ch/lhc>