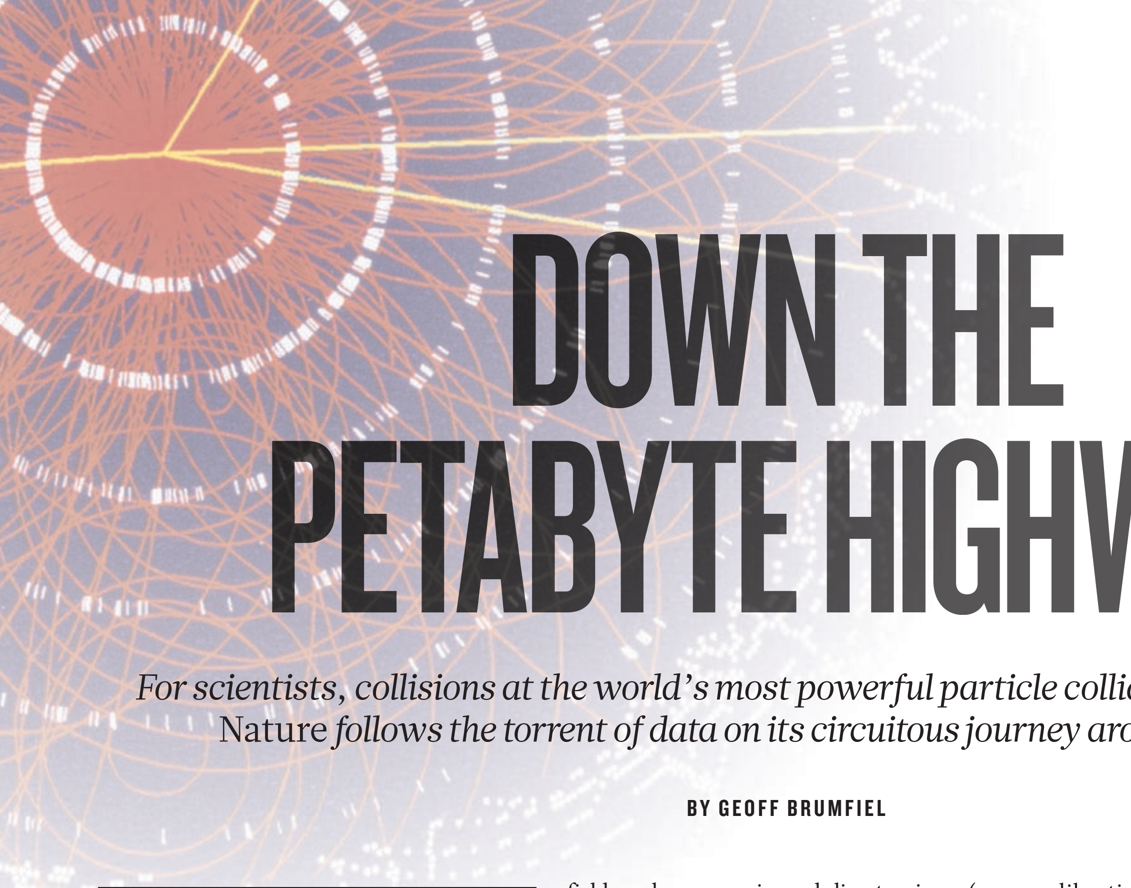

CERN

DOWN THE PETABYTE HIGHWAY

For scientists, collisions at the world's most powerful particle collider are just the start. Nature *follows the torrent of data on its circuitous journey around the world.*

BY GEOFF BRUMFIEL

ATLAS PARTICLE DETECTOR, SWITZERLAND, 30 MARCH 2010, 13:06 LOCAL TIME

Beneath gently rolling hills between the mountains of Switzerland and France, the world's greatest physics experiment starts its first real run. Two beams of high-energy protons meet head-on at almost the speed of light inside the Large Hadron Collider (LHC), a giant particle accelerator at CERN, Europe's high-energy physics lab. Nanoseconds after the protons crash together, their combined energy gives birth to heavier particles, which decay in an instant into a splatter of lighter debris.

At the collision point, 92 metres underground, the 7,000-tonne ATLAS detector sees everything. The debris particles pass first through the detector's inner tracker — a sophisticated layer of silicon electronics that records their paths. Beyond that lie systems that measure the energies of the particles. Some drag to a stop there, but heavy cousins of electrons called muons barrel along, flying metres from the collision point before being picked up by giant, mustard-coloured sensors.

Microprocessors convert the particles' paths and energies into electronic signals, and select a handful of promising collisions for a closer look. The data from the chosen collisions zip upstairs to a computer farm that discards the majority and creates a digital reconstruction of those that remain.

Even after rejecting 199,999 of every 200,000 collisions, the detector churns out 19 gigabytes of data in the first minute. In total, ATLAS and the three other main detectors at the LHC produced 13 petabytes (13×10^{15} bytes) of data in 2010, which would fill a stack of CDs around 14 kilometres high. That rate outstrips any other scientific effort going on today, even in data-rich fields such as genomics and climate science (see *Nature* **455,** 16–21; 2008). And the analyses are more complex too. Particle physicists must study millions of collisions at once to find the signals buried in them — information on dark matter, extra dimensions and new particles that could plug holes in current models of the Universe. Their primary quarry is the Higgs boson, a particle thought to have a central role in determining the mass of all other known particles.

The architects of the LHC decided in 2001 to deal with all that data by dividing and conquering. The results from the giant particle detectors get parcelled up and sent to a vast global network known as the Worldwide LHC Computing Grid, the most sophisticated data-taking and analysis system ever built. The network is as great a technological leap as the collider itself, and without it the project would quickly drown in its own data.

The Grid consists of some 200,000 processing cores and 150 petabytes of disk space, distributed across 34 countries through leased data lines (see 'March of the data'). By combining these resources, the Grid enables scientists to run vast analyses that would push the world's most powerful supercomputers to the edge.

CERN COMPUTING CENTRE, 30 MARCH 2010

Within minutes, the first collisions have made their way to a 1970s-era concrete building on the other side of CERN's campus. In a white, high-ceilinged room, racks containing 50,000 computing cores undertake a careful reconstruction of every selected collision. Details of each sub-detector's calibration, along with temperature readings and other environmental data from the cavern where ATLAS is housed, are used to piece each event back together. ATLAS scientists at CERN pull up reconstructions showing starbursts of narrow lines spreading from the collision points.

> **➲ NATURE.COM**
> For more on the Large Hadron Collider, visit: go.nature.com/lxtne9

In Grid terminology, the CERN computing centre is known as Tier 0. It undertakes an initial analysis of the data and stores one copy. The physics data from ATLAS on the first day of the March run total about 5.2 terabytes (5.2×10^{12} bytes), enough to fill around ten laptop computers, or five of the digital storage tapes kept on the floor below the rows of processors. The first day's harvest is modest compared with what will follow, but the ATLAS experiment has more than a thousand collaborators waiting for results. If all of them logged into CERN and attempted to pull the data from the first collisions back to their home institutions, the network would grind to a halt.

So instead, the Grid automatically spreads copies of the data geographically. Inside a small partitioned section of the computing centre, a wall of panels bristles with bright-orange fibre-optic cable. This is the heart of the system, and it routes data to sites across the globe at a blistering rate of 5 gigabytes per second.

OXFORDSHIRE, UK, 30 MARCH 2010

After CERN finishes the initial analysis, a dedicated fibre-optic link carries some of the data from the first round of collisions more than 800 kilometres to the Rutherford Appleton Laboratory, a sprawling research park nestled among muddy fields in rural Oxfordshire. Here, in a modern office building, a computing farm receives the data through a yellow cable only slightly thicker than a phone line. The lab is one of 11 Tier 1 centres spread around the world, where the data are further refined and split.

Particle physics is a bit like investigating a mid-air collision. Nobody is there to witness it; instead, the debris is painstakingly collected and reassembled to give investigators hints as to what happened. In this case, physicists divide up the different kinds of particles for study. One group looks at muons, for example, while another focuses on high-energy γ-rays. The computers at the lab help by creating dozens of copies of the data, focusing on various aspects of the collision. They are given names like data10_7TeV.00152166.physics_MinBias.merge.DESD_PHOJET.* — which contains data on photons and narrow jets of particles.

CHICAGO, ILLINOIS, 15 MAY 2010

A team of US researchers sends a request for data out on the Grid, and information on several subsets of the collisions from 30 March travels from Oxfordshire via New York to a post-war University of Chicago building just two blocks from the site of the Manhattan Project's first nuclear reactor.

Rob Gardner, the physicist in charge of the computing facility, says, "What we've assembled here is a data centre just about as cheaply as we can put silicon on the floor." It looks like a smaller version of the computing centres in Geneva and Oxfordshire, but with one importance difference: researchers can bring coffee into the Chicago site. "It's not a clean environment," says Gardner.

His cluster of computers is one of the Grid's 140 Tier 2 sites. Unlike Tier 1s, which undertake serious reconstructions of the data, Tier 2 centres mainly provide storage and computing resources and can be accessed by users all over the world.

In an office above the cluster, postdoc Antonio Boveia sits at a metal desk with his laptop. His machine is at the far end of the Grid from CERN, with lines of code scrolling against the black screen. To conduct an analysis — such as one on the decay of the Higgs boson into heavy particles known as W-bosons — he types in commands in the common programming language C++. For just one of Boveia's analyses, he must study tens of millions of collisions. Even if his laptop's hard drive were 4,000 times its current size and could accommodate the data, his processor would still take a few years to complete the work. "It would be impossible," he says.

The Grid makes it possible by splitting the task. When Boveia enters his request, the Grid pulls data from sites such as the one in Oxfordshire, then parcels the analysis into thousands of separate pieces and spreads it across the network. The pieces might be processed at CERN, or at a facility in Italy, or, more likely, in many places at once. In a matter of days, Boveia receives an e-mail alert telling him that the analysis is complete.

The operation does not always work so smoothly. The Tier 1 and Tier 2 centres are managed locally, which means that they each have their own protocols — and problems. In the summer of 2009, as simulated data was flowing through the Grid in advance of the first real collisions, fluff from local cottonwood trees clogged the Chicago centre's air-conditioning unit and forced a shutdown. The same year, road workers severed one of CERN's fibre-optic links in Switzerland, and a fire brought down the Tier 1 centre in Taipei, Taiwan, for months. When things go wrong, alerts are dispatched by e-mail or, occasionally, by phone to an assortment of emergency contacts around the globe.

The system relies on goodwill, says Jamie Shiers, a group leader in CERN's computing department. "We have no line management over these people whatsoever," he says. But somehow, the global cooperative produces results.

CERN, 24 DECEMBER 2010, 11:54

The ATLAS team posts an initial analysis from the Chicago group onto the pre-print server arXiv.org (ATLAS Collaboration. Preprint at http://arxiv.org/abs/1012.5382; 2010). The report — on W-bosons produced through mechanisms other than the decay of Higgs bosons — includes collisions from the first day's run, along with many others. Measurements of the W-bosons produced show good agreement with existing theories.

The physics data set from 30 March now makes up just 0.02% of the total data collected by the ATLAS detector. Most physicists on the collaboration are using that initial set without even realizing, as they acquire sections for analysis and combine them with other data sets. The first hints of a Higgs boson may already be stored on a computer disk in Mumbai, Melbourne or one of the many other sites to which LHC data are distributed. But even if it is there, the Higgs will stay hidden until many more petabytes have flowed through the Grid. ■

Geoff Brumfiel *is a senior reporter for* Nature *based in London.*

MARCH OF THE DATA

The Worldwide LHC Computing Grid harnesses 200,000 computing cores in 34 countries. The central node of the Grid, called Tier 0, is housed at CERN in Geneva, Switzerland; there are 11 Tier 1 sites and 140 Tier 2 sites. This story follows data on a trip from CERN to Chicago, Illinois.