



[Home \(/\)](#) / [News \(/news/\)](#) / [CERN Prepares for New Computing Challenges with Large Hadron Collider](#)

CERN Prepares for New Computing Challenges with Large Hadron Collider

Michael Feldman (/project/top500-news-team/) | May 9, 2018 07:31 CEST

@ E-mail (mailto:?)

subject=CERN%20Prepares%20for%20New%20Computing%20Challenges%20with%20Large%20Hadron%20Collider&body=http%3A%2F%2Ftop500.org%2F

🐦 Tweet (https://twitter.com/share?)

url=http%3A%2F%2Ftop500.org%2FtFBMg&text=CERN%20Prepares%20for%20New%20Computing%20Challenges%20with%20Large%20Hadron%20Collider

f Like (https://facebook.com/sharer/sharer.php?u=http%3A%2F%2Ftop500.org%2FtFBMg)

G +1 (https://plus.google.com/share?url=http%3A%2F%2Ftop500.org%2FtFBMg)

in Share (https://www.linkedin.com/shareArticle?mini=true&url=http%3A%2F%2Ftop500.org%2FtFBMg)

Thanks to the discovery of the Higgs boson in 2012, CERN's Large Hadron Collider (LHC) has probably become the most widely recognized science project on the planet. Now almost 10 years old, the 27-kilometer ring of superconducting magnets is the world's largest and most capable particle accelerator. As such, it enables physicists to push the envelope of particle physics research.

Less well-known is the computing infrastructure that supports this effort – that of the The Worldwide LHC Computing Grid (WLCG), a network of more than 170 computing centers spread across 42 countries. Because LHC experiments can involve processing petabytes of data at a time, the computational, networking, and storage challenges for the project are immense. And when the next-generation High-Luminosity LHC (HL-LHC) is launched in 2026, these challenges will become even more formidable.

To find out more about what this entails, we asked Dr. Maria Girone, CERN's openlab CTO, to describe the high performance computing technology that undergirds the LHC work and talk about what kinds of hardware and software are being considered to support the future HL-LHC machine. Below is a lightly edited transcript of our conversation.

TOP500 News: Can you outline a typical computing workflow for an LHC application – for example, the workflow that resulted in the discovery of the Higgs boson particle?

Maria Girone: Workflows in high-energy physics typically involve a range of both data-intensive and compute-intensive activities. The collision data from the cathedral-sized detectors on the Large Hadron Collider needs to be filtered to select a few thousand interesting collisions from as many as one billion that may take place each second. The search for new phenomena is like looking for needles in enormous haystacks.

Once interesting collision events have been selected the processing-intensive period begins. The particles from each collision in the detectors are carefully tracked, the physics objects are identified, and the energy of all the elements are measured with extreme precision.

At the same time, simulation takes place on the Worldwide LHC Computing Grid, the largest collection of computing resources ever assembled for a single scientific endeavor. The WLCG produces a massive sample of billions of simulated beam crossings, trying to predict the response of the detector and compare it to known physics processes and potential new physics signals. In this analysis phase, the data from the detector is examined against predictions based on known background-only signals. When the data diverges statistically significantly from the background-only signals, we declare a discovery.

TOP500 News: Do any of the experiments in the LHC project currently use deep learning or some other form of AI?

Girone: Neural networks and machine-learning techniques have been used in high-energy physics for many years. Optimization techniques, such as boosted decision trees, have been widely used in analysis. The field is now looking to expand the use of deep-learning and AI techniques based on the progress made by industry in these areas.

There is potential for applications throughout the data-selection and processing chain, which could increase the efficiency and performance of the physics searches. Other areas we are exploring include object identification based on 3D image-recognition techniques, improved simulation using adversarial networks, better monitoring via anomaly-detection techniques, and optimized resource use through machine-learning algorithms.

TOP500 News: How are the computing challenges for the Large Hadron Collider different from typical HPC simulations?

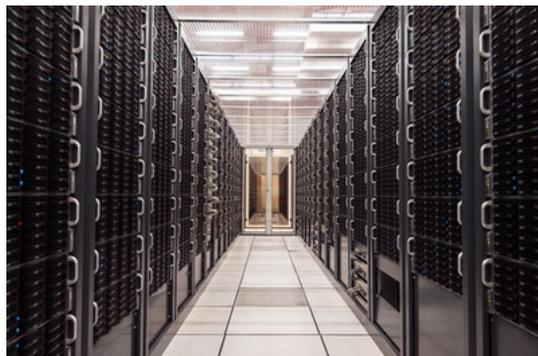


Girone: The computing challenges of the LHC differ from typical HPC applications in the structure of the problem, the time-scale of the program and the number of contributors to the code. Whether processing a data event or producing simulations, each collision event can be treated independently. This means that the application lends itself to simple parallelization across many nodes.

With the LHC program running over multiple decades, there is need for software to be continuously improved. Occasionally, big components are reworked entirely, and there is also a lot of legacy code and services to support.

For many applications used in high-energy physics application, it is the case that several hundred people may well have contributed to the code base over many years. Traditional HPC simulations are often developed by much smaller groups of contributors, with more specific expertise in this area. Our use of code developed by very large numbers of contributors makes it challenging to reach the level of optimization often achieved with other HPC codes.

TOP500 News: Given that the computing grid is spread around the world, what types of challenges are encountered with regard to sharing the large datasets



associated with LHC work?

Girone: Data management has been a consistent area of development in LHC computing. We move petabytes per day and all the data needs to be monitored for consistency. We have become leaders in moving data using global networks. In the last few years, we have augmented the traditional techniques for moving and replicating data to provide real-time remote access to data files across the globe. Our global data-access model has helped to optimize the use of processing and storage resources, as well as making it possible to use commercial cloud and HPC resources in an opportunistic manner.

TOP500 News: Can you describe what the High-Luminosity LHC will be able to do that cannot currently be done with the present-day LHC?

With the HL-LHC, about five to ten times more beam crossings will take place compared to today. Each of these crossings will result in about five times as many individual proton-proton collisions. This increase will help us to search for rarer signals and more precisely measure rare phenomena.

TOP500 News: As you look ahead to the High Luminosity LHC, what emerging technologies in hardware and software do you think are the most promising?

Girone: Looking forward to the HL-LHC, there are many interesting new technologies. The continued improvements in networking technologies will help us to continue distributing data efficiently. The progress made with various types of accelerators is being explored too. We have programs with GPUs and FPGAs that have the potential to dramatically improve the performance of the computing systems.

For software, better optimization and code modernization also hold great promise. Finally, new techniques like advanced data analytics and deep learning have the potential to change how analysis and reconstruction are performed, thus enabling us to process more data more efficiently.

Dr. Girone will present greater depth on the subject of computing challenges at CERN during the opening keynote of the 2018 ISC High Performance Conference (<http://ct.isc-events.com/click.php?id=55>), which will take place on June 24-28 in Frankfurt, Germany. Her keynote address will take place on Monday, June 25.

Images: Maria Girone; CERN IT center. © CERN

[← previous \(/news/cavium-releases-thunderx2-arm-processor/\)](#)

[next → \(/news/swedish-university-to-break-new-ground-with-scandinavias-most-powerful-supercomputer/\)](#)

Current rating: **4.6** ● 1 ● 2 ● 3 ● 4 ● 5

Comments

There are currently no comments

New Comment

Name

required

Email

required (not published)

Website

optional

Comment

Comment

(/lists/2017/11)

(/project/call_for_proposals/)

Tweets by @top500supercomp

TOP500
@top500supercomp

@cray_inc has teamed up with @DigiCatapult to help UK-based startups building artificial intelligence products get access to #HPC resources in Cray's Accel #AI Lab. ow.ly/XmqN30kmJsn



Jun 10, 2018

TOP500 Retweeted

Berkeley Lab CS
@LBNLcs

RT @kamcmahon: Will be attending this one @ISChpc #ISC18 on the discussion of Moore's Law and Dennard scaling - what is changing in processors.

[Embed](#)

[View on Twitter](#)

[Like](#) 4.6K people like this. Be the first of your friends.

Recent Posts

[Summit Up and Running at Oak Ridge, Claims First Exascale Application \(/news/summit-up-and-running-at-oak-ridge-claims-first-exascale-application/\)](/news/summit-up-and-running-at-oak-ridge-claims-first-exascale-application/)

[University of Sydney Adds AI to Its Repertoire with New Supercomputer \(/news/university-of-sydney-adds-ai-to-its-repertoire-with-new-supercomputer/\)](/news/university-of-sydney-adds-ai-to-its-repertoire-with-new-supercomputer/)

[AMD Previews Future GPUs and CPUs for the Datacenter \(/news/amd-previews-future-gpus-and-cpus-for-the-datacenter/\)](/news/amd-previews-future-gpus-and-cpus-for-the-datacenter/)

[As Moore's Law Winds Down, Chipmakers Consider the Path Forward \(/news/as-moores-law-winds-down-chipmakers-consider-the-path-forward/\)](/news/as-moores-law-winds-down-chipmakers-consider-the-path-forward/)

[UK AI Startups Will Get Access to Cray Supercomputers \(/news/uk-ai-startups-will-get-access-to-cray-supercomputers/\)](/news/uk-ai-startups-will-get-access-to-cray-supercomputers/)

Archive

2018

June (/news/archive/2018/6/) (7)

May (/news/archive/2018/5/) (25)

April (/news/archive/2018/4/) (19)

March (/news/archive/2018/3/) (18)

February (/news/archive/2018/2/) (16)

January (/news/archive/2018/1/) (23)

2017

December (/news/archive/2017/12/) (15)

[November \(/news/archive/2017/11/\)](/news/archive/2017/11/) (24)
[October \(/news/archive/2017/10/\)](/news/archive/2017/10/) (24)
[September \(/news/archive/2017/9/\)](/news/archive/2017/9/) (19)
[August \(/news/archive/2017/8/\)](/news/archive/2017/8/) (23)
[July \(/news/archive/2017/7/\)](/news/archive/2017/7/) (21)
[June \(/news/archive/2017/6/\)](/news/archive/2017/6/) (31)
[May \(/news/archive/2017/5/\)](/news/archive/2017/5/) (24)
[April \(/news/archive/2017/4/\)](/news/archive/2017/4/) (21)
[March \(/news/archive/2017/3/\)](/news/archive/2017/3/) (20)
[February \(/news/archive/2017/2/\)](/news/archive/2017/2/) (21)
[January \(/news/archive/2017/1/\)](/news/archive/2017/1/) (17)

2016

[December \(/news/archive/2016/12/\)](/news/archive/2016/12/) (15)
[November \(/news/archive/2016/11/\)](/news/archive/2016/11/) (26)
[October \(/news/archive/2016/10/\)](/news/archive/2016/10/) (23)
[September \(/news/archive/2016/9/\)](/news/archive/2016/9/) (22)
[August \(/news/archive/2016/8/\)](/news/archive/2016/8/) (25)
[July \(/news/archive/2016/7/\)](/news/archive/2016/7/) (23)
[June \(/news/archive/2016/6/\)](/news/archive/2016/6/) (27)
[May \(/news/archive/2016/5/\)](/news/archive/2016/5/) (23)
[April \(/news/archive/2016/4/\)](/news/archive/2016/4/) (7)
[March \(/news/archive/2016/3/\)](/news/archive/2016/3/) (4)
[February \(/news/archive/2016/2/\)](/news/archive/2016/2/) (7)
[January \(/news/archive/2016/1/\)](/news/archive/2016/1/) (4)

Categories

[Blog \(/news/category/blog/\)](/news/category/blog/) (18)
[Feature Article \(/news/category/feature-article/\)](/news/category/feature-article/) (177)
[In Depth \(/news/category/in-depth/\)](/news/category/in-depth/) (167)
[Intersect360 Research Contribution \(/news/category/intersect360-research-contribution/\)](/news/category/intersect360-research-contribution/) (1)
[ISC Events \(/news/category/isc-events/\)](/news/category/isc-events/) (1)
[press-release \(/news/category/press-release/\)](/news/category/press-release/) (10)
[sticky-green \(/news/category/sticky-green/\)](/news/category/sticky-green/) (13)
[This Week in HPC \(/news/category/this-week-in-hpc/\)](/news/category/this-week-in-hpc/) (183)
[Whispers \(/news/category/whispers/\)](/news/category/whispers/) (4)

Tags

[whispers \(/news/tag/whispers/\)](/news/tag/whispers/) (3) [green500 \(/news/tag/green500/\)](/news/tag/green500/) (20) [hpe \(/news/tag/hpe/\)](/news/tag/hpe/) (4)

Feeds

[RSS \(/news/feeds/rss/\)](/news/feeds/rss/) / [Atom \(/news/feeds/atom/\)](/news/feeds/atom/)

[IMPRINT \(/IMPRINT\)](/imprint/) [CONTACT \(/CONTACT\)](/contact/) [LOG IN \(/ACCOUNTS/LOGIN/?NEXT=/NEWS/CERN-PREPARES-FOR-NEW-COMPUTING-CHALLENGES-WITH-LARGE-HADRON-COLLIDER/\)](/accounts/login/?next=/news/cern-prepares-for-new-computing-challenges-with-large-hadron-collider/) [SIGN UP \(/ACCOUNTS/SIGNUP/?NEXT=/NEWS/CERN-PREPARES-FOR-NEW-COMPUTING-CHALLENGES-WITH-LARGE-HADRON-COLLIDER/\)](/accounts/signup/?next=/news/cern-prepares-for-new-computing-challenges-with-large-hadron-collider/)

Copyright 1993-2018 TOP500.org (c)