

"What we want to work on is the big items, things like true deep copy or a seamless way to spread parallel regions across multiple devices, or load balancing across the GPU *and* the CPU and how do you manage that. Those are big items; those are what users really want," said Wolfe.

He added, "Last summer I was visiting CSCS in Lugano, Switzerland, and each node of the cluster they host for the weather forecasting service MateoSwiss has four K80s, so [eight GPUs per node](#). Well how do you manage that? Is it easy or is there a way to make it even easier, that's a big challenge, and we're ready to take it on."

Community engagement and education

Via partnerships with Oak Ridge National Laboratory and other member orgs, OpenACC continues to offer hackathons around the world. More info is at [http://www.openacc.org/hackathons](#). Oak Ridge and OpenACC will also be conducting a series of free two and half day workshops starting next year. These are designed to introduce developers to the framework and are a new addition to OpenACC's [training and education program](#).

SC16 activities include:

[OpenACC Birds of a Feather, Wed. Nov. 16th 5:30–7:00PM in room 155-C](#). Discussion will include such topics as "Should OpenACC and OpenMP ever merge".

Free "[Parallel Programming with OpenACC](#)" books will be signed by author Rob Farber Monday, November 14 from 7:00 to 9:00 pm in the OpenACC booth #634.

[Bringing About HPC Open-Standards World Peace](#), Nov. 16th, 10:30 am 255-BC

Members will be available for questions in the OpenACC booth #634.

Visit [http://www.openacc.org/sc16](#) for a list of all OpenACC member activities.

The post [OpenACC Expands Community, Reveals Roadmap Details](#) appeared first on [HPCwire](#).

Categories: [RSS Feed Reader](#)

[Intel Collaborates with CERN to Support Upgraded LHC Experiments](#)

[HPC Wire](#) - Fri, 11/04/2016 - 14:49

Much of the media attention given to the particle accelerator experiments that happen at the European Organization for Nuclear Research, known as CERN, is focused on the Large Hadron Collider (LHC). It's no surprise, given the LHC is the world's largest, most complex machine, unravelling some of the toughest scientific problems by accelerating particles (protons or heavy ions) and making them collide in a gigantic 27-kilometer ring. But the work that happens immediately after particles collide in the LHC is not only critical to science, it's also quite interesting and important from a computing and data processing perspective. After all, the creation of particles or results in the LHC is only significant if scientists can quickly isolate them from millions of inconsequential signals for further study. That means ongoing advancements in trigger and data acquisition systems are essential to fully reaping the rich potential of the LHC. And, as you can imagine, the networking and computing challenges are extreme in nearly every dimension.

Historically, CERN's trigger and data acquisition systems have relied heavily on custom technology. For the next upgrade cycle scheduled to start around the end of 2018, however, CERN engineers and scientists were aware that the scalability limitations and costs of their custom solutions could start to limit progress. This led to a new collaborative project with Intel, through [CERN openlab](#), focused on exploring the feasibility of complementing one of the LHC's detectors with off-the-shelf data acquisition, data movement, and data filtering technology from Intel, including an FPGA platform. If the proof of concept is successful, it could have a significant impact on the future design and efficiency of trigger systems, including the remaining LHC detectors and other scientific instruments.

The Large Hadron Collider tunnel is located 100 meters underground on the Franco-Swiss border, near Geneva. Source: CERN

Preparing for enormous data growth

Simply phrased, the LHC fires high-energy particle beams at each other in a 27-kilometer ring. The detectors on the LHC include tracking devices that plot the trajectory of particles following collisions, as well as calorimeters that measure their energy, which helps to narrow down their identity. Niko Neufeld, a deputy project leader at CERN who works on the LHCb experiment, [likened the computing challenges](#) for the LHC experiments to solving millions of small puzzles involving up to a billion proton-proton collisions every second to retain the most interesting ones for deeper analysis.

Given the task at hand, the near-detector "online" computing challenges at CERN have always been extreme. And when CERN upgrades the LHC and detectors from late 2018 to early 2021 as part of its regular upgrade cycle, the data rates running through the various systems will jump significantly. Neufeld provides perspective on the jump: "Network scaling needs have really grown. For example, the largest networks we currently run at CERN have total bandwidth of around 800 GB per second. Following the upgrade work, our networks will need to carry between 40 and 50 TB of data per second. If you compare that to a Google data center, it may not sound impressive, but for a scientific instrument it's a huge step in terms of bandwidth." Neufeld said that the computing challenges have also grown quite complex. "We cannot simply scale up computing by a factor of 100 because we have whatever Moore's law gives us... We may be able to grow our computing farms by a factor of 1 or 2, but not much more. The rest has to come from more clever processing models," he added.

The Compact Muon Solenoid (CMS) is one of two large general-purpose detectors on the LHC. The image above captures a candidate proton-proton event as a part of the CMS search for the Higgs boson. Source: CERN

According to Neufeld, the detector teams face a host of challenges in preparing for the data-rate jump. Neufeld said that the customized trigger (hardware and software) on the front-end of the existing detector system had a long and expensive refresh cycle. “The engineering resources for ASICs [application-specific integrated circuits] and the FPGAs [Field Programmable Gate Arrays] in high energy physics are limited compared to industry, and the tight integration with the detectors makes upgrades outside of our major maintenance periods impractical,” Neufeld explained. “We thought that moving to a more software-centered approach using off-the-shelf technology could greatly reduce these limitations and expand the developer base that is available in our community. Physicists are usually knowledgeable in some programming language, however, HDL [hardware description language] is a different challenge with a very long and steep learning curve.”

Olof Bärring, a deputy group leader of computing facilities at CERN, added that cost and energy considerations were also an important part of the equation. The lab needs to continue addressing greater computing, data moving and storage challenges with a more or less flat budget and within the datacenter’s existing energy envelope.

Exploring options for three critical challenges

In 2014, the former CTO of CERN openlab, Sverre Jarp, invited Intel’s Karl Solchenbach, director Exascale Labs Europe, and Steve Pawlowski, former vice president of advanced computing solutions at Intel, to discuss CERN’s technical challenges with near-detector online computing to see if any Intel technology might be useful in addressing them. During the meeting, Neufeld presented three main challenges, and the participants worked together to map technologies to the challenges.

Challenge one: real-time or near-real-time data-processing with very short (order of 10 microseconds) latencies

The Intel team considered an Intel Xeon processor and Altera Stratix V FPGA integration along with another spatial architecture, and ended up recommending the Intel Xeon processor/FPGA configuration. Neufeld says that the reasoning was that the Intel Xeon processor/FPGA configuration should allow LHC experiments to potentially replace part, and in some cases all, of the custom electronics used in the first step of online data filtering. “That meant we would be able to use off-the-shelf hardware programmed using ‘high-level’ general purpose languages instead of HDLs, which was an important step for us,” he explained.

Challenge two: very high throughput local area networks

For data transport, Intel recommended investigating the potential of Intel Omni-Path Architecture (Intel OPA) as an alternative to deep-buffer Ethernet switching. Neufeld said the main driver for considering Intel OPA over traditional approaches was cost. “It is certainly technically possible to build the network in a more traditional way, but it has become prohibitively expensive, given our budget,” he explained.

Challenge three: the need of massive data-processing for data-reduction

For the filtering of detector data in software, the new Intel Xeon Phi processor was an attractive potential solution, given that the filtering process itself is quite parallel and individual collisions in the LHC are statistically independent.

Not your average POC

Once the CERN and Intel teams agreed that the identified off-the-shelf solutions for each challenge had potential, the hard work of proving the viability of each solution in the next-generation data acquisition environment needed to begin in earnest. In 2015, CERN and Intel decided to expand their existing collaboration through CERN openlab and signed an agreement for a joint three-year project called the High Throughput Computing Collaboration (HTCC). As of writing, the HTCC project has reached the halfway point and the core team, which includes seven CERN scientists and one Intel engineer on site, has made progress in each key challenge area.

The main CERN data center. The 110,000 processor cores and 10,000 servers hosted in its three rooms run 24/7. Source: CERN

Networking milestones

In the networking area, Neufeld said a big part of the challenge is simply getting access to a system big enough to test software. “To properly prepare our software, we need access to complete supercomputers. Fortunately, Intel can provide access to clusters that are up to the job,” he explained. Neufeld said that the team recently had an important success while running its software on an Intel OPA cluster with more than 500 nodes. “We have already been able to achieve a full-duplex, high-throughput transit of 70 TB per second of data flying through the cluster, so that’s already half of what we need by the time of the upgrades.”

FPGA-related milestones

When CERN updates the LHC and the experiments from late 2018 to early 2021, its detectors will support trigger-free readouts. The LHC generates up to around 1 billion collisions per second in the experiments, and the goal is to read them all. Moving forward, a flexible software-based trigger system running in a large (up to 4000 nodes) computing farm will select the interesting collisions. In the meantime, CERN is investigating which technology options are the best fit for accelerating its algorithms. For its initial FPGA proof of concept, the HTCC team deployed a two-socket Intel Xeon processor/FPGA machine, which included the following hardware connected by the Intel QuickPath Interconnect:

- Intel Xeon CPU E5-2680 v2
- Altera Stratix V GX A7 FPGA with 234,720 adaptive logic modules (ALMs)

The team is particularly interested in the potential compute and power efficiency gains that are possible with using OpenCL in a combined CPU and FPGA system. With respect to the FPGA platform testing, Neufeld said that the team dealt with host of technical challenges because it wanted to do a meaningful comparison of OpenCL and to get a sense of the costs using a high-level framework. "Just for illustrative purposes, it took us two weeks to set up a new kernel using OpenCL compared to three months to complete an equivalent Verilog implementation, and we had a very skilled engineer on that job," explained Neufeld. "In the end, I think it was a good investment because we needed to prove to the electronic engineers that the new technology actually provided a less painful way to get the results we need."

Following a series of test cases, such as sorting and calculating the Mandlebrot fractal, to understand the potential of the Xeon/FPGA system, the HTCC team developed an FPGA fine-tuned for the rigors of RICH (ring-imaging Cherenkov) reconstruction. Only then did it begin doing LHC-specific workload analysis.

Cherenkov angle reconstruction is used for particle identification in the detectors. Based on initial tests the Xeon/FPGA machine shows promise for processing greater numbers of photons after the LHC upgrade. Source: CERN

In coming months, the HTCC team will also be testing a newer system that is built using a combined Intel Xeon CPU and FPGA in a single package. It will include the new high-performance Arria 10 FPGA from Altera as well as a faster interconnect of the CPU and FPGA.

Intel Xeon Phi processor milestones

Neufeld said that testing of the Intel Xeon Phi processor platform will be a major focus of the HTCC team for the next year or so. He noted that like everyone else, the team needs to figure out how to adapt well to the new architecture and different level of parallelism. To achieve this the HTCC team has been working with Intel engineers on benchmarking and understanding the different algorithms' implementations using Intel analysis tools, such as Intel VTune Amplifier XE and Intel Advisor XE, as well as different performance models, such as the roofline model.

"In addition to the inclusion of bootable sockets, in-package memory and high main-memory bandwidth, what is particularly interesting with Intel Xeon Phi processors is the integrated fabric and its potential to quickly distribute workloads to where they fit best. We will test both data movement aspects on the Intel Xeon Phi processor as well as the distribution of the algorithms between the Intel Xeon and Intel Xeon Phi processor using the fabric as an interconnect," Neufeld added.

Pushing the boundaries of precision

When asked what the progress made on the data acquisition systems for the LHC might mean for wider applications, Neufeld said the value is all about greater precision for complex experiments. "To some extent, it's just statistics. Either you really increase the amount of data and the precision by a significant factor, or you stop doing it," Neufeld explained. "This work should lead to an important jump in precision. For example, the LHCb experiment's online collection and analysis system currently selects just 1 million of the 40 million bunches of protons that cross in the accelerator every second, with the others being discarded based on less-precise hardware-calculated signatures. After the LHCb upgrade, the number of collisions is set to grow yet further, and we will look at all of them in the software, in order to take the best physics out of there. And since there are actually only a couple of milliseconds to do that for each collision, it's really quite a leap forward." Direct discussions with the Intel development team will continue to be invaluable to getting things right on projects as the team races to meet its deadline for the start of the upgrade work.

The post [Intel Collaborates with CERN to Support Upgraded LHC Experiments](#) appeared first on [HPCwire](#).

Categories: [RSS Feed Reader](#)

IDC's Conway Sets Stage for SC16 Precision Medicine Panel

[HPC Wire](#) - Fri, 11/04/2016 - 09:56

Q: The fact that Precision Medicine is the opening panel at SC strongly suggests the growing importance of HPC in making PM and basic life science research possible. Recognizing SC is primarily a technology conference, could you frame the goals of this panel?

Precision medicine, also called personalized medicine, promises to transform medical practice and healthcare spending by enabling personalized diagnoses and treatment plans that are custom-tuned for each patient's physiology, symptoms, medical history, DNA and even lifestyle. What constitutes a good outcome for a broken hand may be different for an office worker and a concert violinist. HPC is already playing a key role in early precision medicine initiatives around the world, by speeding up genome sequencing and by making it possible to quickly sift through millions of archived patient records to identify treatments that have had the best success rates for patients closely resembling the patient under investigation. Biology is fast becoming a digital science and healthcare analytics is one of the fastest-growing new market segments for HPC. Precision medicine is happening at the intersection of biology, medical practice, healthcare economics, and data science. The expert panel at SC'16 will explore this emerging domain from these varied perspectives, with special emphasis on the major role HPC has already started to play.

Warren Kibbe, NCI

This is a pretty august group:

- Mitchell Cohen, Director of Surgery, Denver Health Medical Center; Professor, University of Colorado School of Medicine.
- Martha Head, Senior Director, The Noldor; Acting Head, Insights from Data at GlaxoSmithKline Pharmaceuticals
- Warren Kibbe, Director, Center for Biomedical Informatics and Information Technology (CBIIT); Chief Information Officer; Acting Deputy Director; National Cancer Institute (NCI)