Data acquisition at CERN: A future challenge

Marcel Zeiler

How did the universe look in the first moments after the Big Bang? Why does matter dominate over antimatter? What are the fundamental particles that make up the world as we see it today? To find answers to those and similar questions, the European Organization for Nuclear Research (CERN) is operating the world's largest particle accelerator, the Large Hadron Collider (LHC). The LHC accelerates protons to a velocity close to the speed of light and makes hundreds of them collide. The circumstances shortly after those collisions are representative to the universe's conditions only moments after the Big Bang. By analyzing thousands of such collisions, a steps can be made toward answering the previous questions.

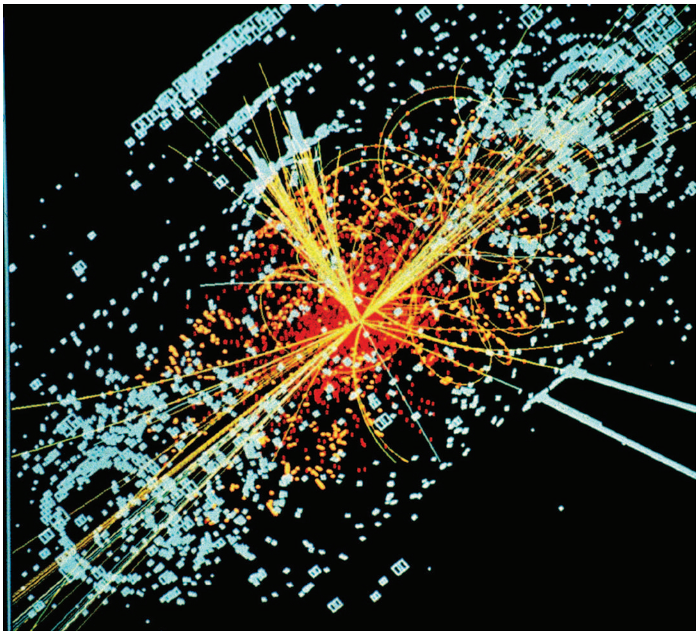

IMAGE COURTESY OF WIKIMEDIA COMMONS/CERN.

Before the collision events can be analyzed, the particle interactions must be detected. For this purpose, four huge particle detectors are installed alongside the LHC accelerator ring. The detectors comprise up to 10 million individual sensors arranged around the particle collision center. Each sensor can be regarded as a camera taking a picture of the particle collisions from a different angle. An entire collision event can later be reconstructed by combining all pictures taken at the same time. Since the particle collision rate is 40 MHz, pictures are taken every 25 ns.

This results in a tremendous amount of raw data produced by all the sensors in the four main detectors. In fact, raw measurement data in the order of 400 Tb/s has been created each second during LHC's Run 1 from 2009 to 2013. Operating the detectors for a full year, the raw data piles up to about 25 PB, which equals the storage capacity of more than 2.5 million DVDs. Before this data can be analyzed for interesting physics events, it needs to be sent from the detectors to data centers, where real-time data processing, filtering, and storage occurs.

As it is unaffordable to install and maintain a data acquisition (DAQ) system capable of storing and retrieving all raw measurement data on demand, multiple buffer-and trigger-stages are installed in the LHC's read-out chain. These stages are designed to select only new and rare particle collision events from the raw data and discard obviously well-known and frequent collision events that make up most of the measurement data. Special filtering criteria are established to decide whether or not a data set is to be kept. For example, if a set of collision data shows a specific combination of particle types and energy levels, it passes the filtering stage. Otherwise, the data is removed. In this way, the data stream that is actually stored and analyzed can be reduced to only a few gigabits per second. Using this solution makes data acquisition at CERN much more affordable and efficient.

With the LHC usually being upgraded every three to five years, its

Digital Object Identifier 10.1109/MPOT.2015.2441779
Date of publication: 6 May 2016

DAQ system needs to be modified too. The current upgrade schedule foresees the LHC being upgraded to the high-luminosity LHC (HL-LHC) by the middle of the next decade. Compared to the LHC that discovered the Higgs boson in 2012, the maximum particle collision energy will be doubled to 14TeV, and the average number of colliding particles will also increase considerably. Higher collision energies and more particles are expected to increase the produced raw data, and, therefore, also the required data processing rate, on average by a factor of 25. This sets unprecedented requirements on the DAQ systems needed by then.

Taking a DIP

To meet these requirements, the *Intel–CERN European Doctorate Industrial Program* (ICE-DIP) was launched to address some of the emerging challenges. ICE-DIP is a public–private research and Ph.D. program funded by the European Commission. Both partners greatly benefit from this collaboration. CERN can investigate cutting-edge information and communication technologies for their applicability to the LHC, while Intel gets a unique chance to test and validate the functionality of their products in especially demanding environments. Each ICE-DIP researcher is also a Ph.D. student, with Maynooth University and Dublin City University in Ireland being the academic partners. The Ph.D. topics are defined such that they cover a wide range of aspects relevant to DAQ at CERN (see Figs. 1 and 2): optoelectronics, network and communication technology, and computer science. The general objective is to identify key-technologies for future DAQ systems. A particular emphasis lies on high-throughput and low-latency DAQ systems for high-energy physics experiments. The research outcomes, however, can also be transferred to other sciences where large amounts of data need to be analyzed.

A simplified schematic of the LHC's DAQ system from a particle detector to data storage and processing is depicted in Fig. 2. The very first step in the DAQ process is to transport particle collision data, produced in the detectors, to the low-level trigger system, where the first data filtering takes place. Fiber optical links provide the means of transportation for this purpose. Each optical link constitutes a transmitter, an optical fiber, and a receiver. In the LHC's fiber-optical installation, the optical transmitters sit very close to the center of the particle collisions. As a result, they are exposed to very high levels of radiation and strong magnetic fields. This makes it impossible to use commercial off-the-shelf optical transmitters. Consequently, dedicated research is needed to develop optical transmitters that can be operated in LHC's harsh environment.

The ICE-DIP's first work package covers research on low-cost, radiation-hard silicon photonics for data-intensive communication links. The emphasis lies on the applicability of silicon-photonic-based optical components for their utilization in radiation environments. In particular, the design and test of custom interferometric optical modulators that can withstand the radiation in the particle detectors is addressed. Silicon photonics technology has been chosen because high-performance optical components in silicon have

FIG1 The ICE-DIP network and its research topics.

Higher collision energies and more particles are expected to increase the produced raw data, and, therefore, also the required data processing rate, on average by a factor of 25.

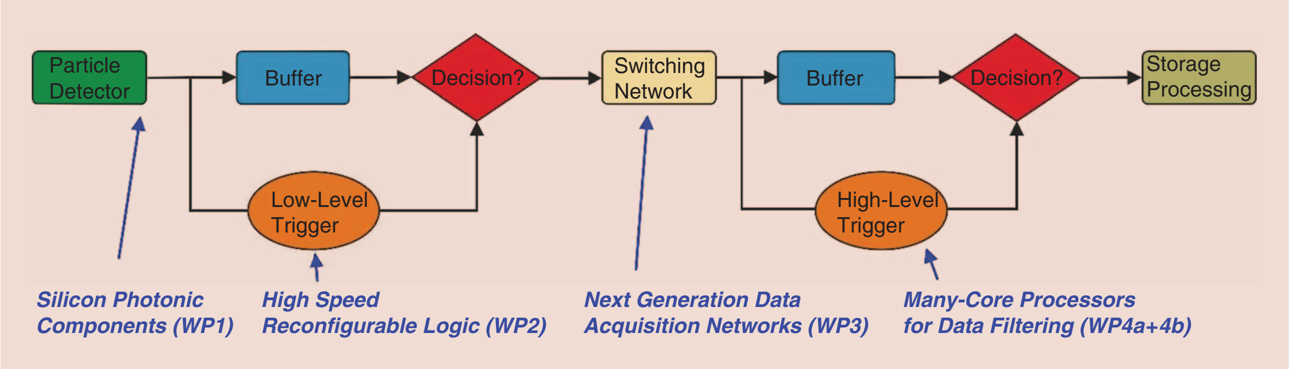

FIG2 A simplified schematic of the LHC's DAQ system, including the area of application of ICE-DIP's work packages (WPs). Before data is stored and processed, it undergoes several buffer-and trigger-stages, where it is decided if a data set is or is not interesting. If the decision is positive, the data set will be kept, otherwise it is discarded.

been demonstrated in recent years, inter alia by Intel. Furthermore, it promises low manufacturing costs due to the existing complementary-metal-oxide semiconductor fabrication infrastructure and the possibility for integration with silicon-based electronic drivers and particle detectors. An already proven radiation hardness for electronic components is an additional plus, making silicon a very promising candidate material for optical transmitters deployable inside the particle detectors at CERN.

The raw measurement data is sent optically over thousands of parallel channels to the low-level trigger system, where it is predominantly acquired on FPGAs. At this first filtering stage, data preprocessing occurs to select potentially more interesting collision events for further processing and discard others. Apart from performing research on enhancing DAQ on FPGAs, the second ICE-DIP work package also focuses on the integration of high-speed, reconfigurable FPGAs and commodity CPUs. Tightly integrated FPGA/CPU modules will allow highly customizable and versatile data processing nodes. Additionally, FPGAs not just offer high computation performance but also show better performance/power consumption compared to CPUs and GPUs. Thus, they are of great interest to organizations like CERN and Intel when it comes to large amounts of data.

However, complex programming models for FPGAs and difficulties with CPU compatibility are principle issues that prevented this technology from widespread adoption in data center applications so far. Hence, the development of a prototype based on Intel's QuickPath Interconnect (QPI) interface for hardware integration of FPGAs with CPUs, as well as an Open Computing Language (OpenCL) for FPGAs for software integration, are key objectives of this work. This work will not only enable better exploitation of FPGAs for algorithm acceleration in real-time decision making (e.g., at the low-level trigger system in CERN's case), the outcomes can potentially also be used wherever a GPU is in use today.

Once the data has been preprocessed at the low-level trigger system, it is sent to the high-level trigger system for similar, but more sophisticated, data filtering. Special switching networks based on TCP/IP and Ethernet or InfiniBand technology are installed for this purpose. This is imperative because the many-to-one communication pattern, often used in DAQ networks, is susceptible to network overloading and collapsing when very high bursts of data arrive at the same time. This congestion must be suppressed as much as possible, as it entails the loss of data packets, and thereby, expensive retransmission would be required.

To mitigate this problem, work package three focuses on reliable high-bandwidth data acquisition networks based on commercial off-the-shelf components. Usually such networks are realized based on expensive core routers that have big buffers to account for very high data bursts. Within the ICE-DIP, it is investigated whether these core routers can be replaced by lower-priced commodity servers having a lot of DRAM memory and multiple network ports. Similar to core routers, these servers must still provide the same, or an even better, network performance, with loss-less data transmission rates of multiple terabits per second.

Filter, store, analyze

Now that the data has been routed through the switching network to the high-level trigger stage, it can be filtered further before it is eventually stored and analyzed thoroughly. This second filtering process currently relies mostly on a general purpose multicore processor to provide the necessary platform for the involved high-performance computing (HPC) tasks. The trend in HPC is moving toward heterogeneous solutions, for example, combining multicore and many-core processors.

Accordingly, the fourth work package investigates solutions to use many-core processors in high-throughput data filtering applications using commodity PCs. The work is subdivided into two sub-packages. Data access by many-core processors with respect to data transfer rate and access latency, as well as energy cost, is addressed in the first subpackage. This is important since the data to be filtered comes through the network instead of being directly loaded into the memory of the many-core processor platforms. As a consequence, many synchronizations issues that have not yet been solved occur.

Identification of an interface to optimize the available many-core processing power is pursued in the second subpackage. Software parallelization is regarded as the key enabler for this direction. To make use of parallelization, latency and bandwidth constraints must be identified and removed. Minimizing synchronization issues and maximizing parallelization will enable porting trigger tasks from multicore to many-core architectures, thereby optimizing the effectiveness of high-level trigger decisions.

The fraction of data that passes the last filtering stage is eventually sent to the LHC's worldwide computing grid for permanent storage. Scientists can access the data there and analyze it for rare and extraordinary collision events that might lead to new insights to fundamental questions about our universe.

To continue this ongoing research in particle physics after the LHC era, new advancements in high-end DAQ technology are mandatory. To support this evolution, novel key DAQ technologies are researched within the ICE-DIP project, whose outcomes will reveal new concepts in this field. Similarly, other disciplines beyond high-energy physics experiments that deal with real-time processing and analysis of large amounts of data will benefit from the conclusions as well. The project also highlights CERN's outstanding training opportunities for young scientists and engineers and demonstrates how they can contribute to future breakthroughs in the field of particle physics.

About the author

Marcel Zeiler (marcel.zeiler@cern.ch) earned his B.Sc. degree in physics from the Free University of Berlin in 2009 and his M.Sc. degree in physics from the Technical University of Berlin in 2012. Before joining CERN as a doctoral student in October 2013, he conducted research on directly modulated diode lasers for optical data communications at the Fraunhofer Heinrich-Hertz-Institute for Telecommunications in Berlin. He is currently investigating whether silicon photonic devices can be used as optical data transmitters in the harsh environments of high-energy physics experiments at the Large Hadron Collider.