The Compact Muon Solenoid Experiment
Conference Report

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland

27 April 2020 (v4, 07 August 2020)

New challenges for distributed computing at the CMS experiment

Natascha Krammer for the CMS Collaboration

Abstract

The Large Hadron Collider (LHC) experiments soon step into the next period of run-3 data-taking with an increased data rate and high pileup requiring an excellent working computing infrastructure. In the future High-Luminosity LHC (HL-LHC) data-taking period, the compute, storage and network facilities have to be further extended by large factors and flexible and sophisticated computing models are essential. New techniques of modern state-of-the-art methods in physics analysis and data science, Deep Learning and Big Data tools, are crucial to handle high-dimensional and more complex problems. Beside flexible cloud computing technologies the usage of High Performance Computing (HPC) at the LHC experiments are explored. In this presentation, I will discuss the LHC run-3 and future HL-LHC runs computing technologies and the utilisation of modern physics analysis and data science methods for the increasing and complex demands of large scale scientific computing.

Presented at *INSTR20 Instrumentation for Colliding Beam Physics*

Prepared for submission to JINST

INST20 - INSTRUMENTATION FOR COLLIDING BEAM PHYSICS

24 – 28 FEB. 2020

BUDKER INSTITUTE OF NUCLEAR PHYSICS AND NOVOSIBIRSK STATE UNIVERSITY, NOVOSIBIRSK
(RUSSIAN FEDERATION)

New challenges for distributed computing at the CMS experiment

Natascha Krammer on behalf of the CMS collaboration

Institute of High Energy Physics, Austrian Academy of Science, Vienna, Austria

E-mail: natascha.krammer@oeaw.ac.at

Abstract: The Large Hadron Collider (LHC) experiments soon step into the next period of run-3 data taking with an increased data rate and high pileup requiring an excellent working computing infrastructure. In the future High-Luminosity LHC (HL-LHC) data-taking period, the compute, storage and network facilities have to be further extended by large factors and flexible and sophisticated computing models are essential. New techniques of modern state-of-the-art methods in physics analysis and data science, Deep Learning and Big Data tools, are crucial to handle high-dimensional and more complex problems. Beside flexible cloud computing technologies the usage of High Performance Computing (HPC) at the LHC experiments are explored. In this presentation, I will discuss the LHC run-3 and future HL-LHC runs computing technologies and the utilization of modern physics analysis and data science methods for the increasing and complex demands of large-scale scientific computing.

Keywords: Computing, Data processing methods, Analysis and statistical methods

Contents

1	Introduction	1
2	Helix Nebula Science Cloud	2
3	HEPCloud at Fermilab	3
4	Usage of Opportunistic CPU resources at the CMS experiment	3
5	Usage of High Performance Computing resources at the CMS experiment	4
6	Big Data solution - Analytics and Monitoring Infrastructure at CERN (Aka MONIT)	4
7	Machine Learning in High Energy Physics - Particle Classification with Deep Learning using Neural Networks	4
8	Conclusions	6

1 Introduction

The next LHC runs with increasing data rate and high pileup interactions will imply major computing challenges. At the LHC run-3 the computing experts have to deal with a peak luminosity (rate of collisions) of $2.0 \times 10^{34} \text{cm}^{-2} \text{s}^{-1}$. During the next step, the HL-LHC upgrade in 2024, its peak luminosity will increase by a factor of five beyond the original design value and the integrated luminosity (total collisions created) by a factor ten. This results in an expected peak luminosity up to $5 - 7.5 \times 10^{34} \text{cm}^{-2} \text{s}^{-1}$ and after the estimated time of operation of 10 years, a total integrated luminosity seen by the CMS detector [1] of up to $3.000 - 4.000 \text{fb}^{-1}$. Assuming a flat budget the resulting computing resource requirements for run-3 will go well beyond any capacity increase expected from improvements in technology, with 15% as optimistic estimate. Such a development is shown in Fig. 1. However, many indications suggest that 10% is much more realistic [2]. A constant budget growth until run-4 is essential for HL-LHC with an expected increase of 60 times more computing power and 12 times more storage compared to run-2.

The growing computing needs make it necessary for the CMS experiment to investigate the use of resources beyond the traditional grid-provided systems. During the LHC runs the computing needs of experiments are not constant over time. The main computing activities, data (re)processing, simulation data generation, reconstruction and analysis activities have an irregular time schedule with peaks dictated by the deployment of software releases, analysis results needed for important conferences and the timeline of data taking. Dynamical resources provided by commercial cloud services are necessary to meet the growing

computing and storage needs. Another resource pool, the HPC systems, with computers with a high level of performance compared to a general-purpose computer are deploying Exa-scale machines in the near future (one exaFLOP is a billion billion calculations per second). Physics results require to deal with and analyze a growing amount of stored data in the next years. State-of-the-art technologies are used in High Energy Physics (HEP) like Big Data tools to handle an enormous amount of data and Deep Learning methods to support sophisticated software development. Deep Learning is a class of Machine Learning algorithms inspired by the structure and function of the brain called Artificial Neural Networks.

In the next chapter I will introduce cloud centers in Europe and in the United States as additional resources for HEP computing and tests performed by the CMS experiment to use opportunistic and HPC computing resources. Afterwards I will describe available Big Data solutions and Machine Learning and Deep Learning use cases in HEP.

Figure 1. Run-3 (starting 2021) computing and storage (disk and tape) capacity needs. The broken green line assumes technology gain of 15% with flat budget [2].

2 Helix Nebula Science Cloud

For CERN and their partner institutes an initiative, the European hybrid cloud platform, called the Helix Nebula Science Cloud was established to support high-performance and data-intensive scientific computing. End-users from many research communities like HEP, astronomy, life sciences can use the Science Cloud resources. The project is sponsored by 10 of the European leading public research organizations and co-funded by the European

Commission (Horizon 2020). The procurers of the initiative (CERN, CNRS, DESY, EMBL-EBI, ESRF, IFAE, INFN, KIT, SURFSara and STFC) provided funds and manpower to deploy use-cases with applications, data and in-house IT resources [3].

3 HEPCloud at Fermilab

The HEPCloud [4] at the Fermilab computing center provides for the scientific computing community resources beyond local worker nodes and grid-computing centers by expanding into HPC and cloud computing. Based on the policy for a particular experiment, workflow requirements, costs and efficiency of accessing the various resources the HEPCloud routes jobs to the best local or remote computing resources available, as illustrated in the scheme in Fig. 2.

Figure 2. HEPCloud resources at Fermilab [4].

4 Usage of Opportunistic CPU resources at the CMS experiment

The Batch on EOS Extra Resources (BEER) project was developed at CERN to profit from periods of low average CPU utilization of several hundred EOS file storage servers and to utilize these resources for batch processing jobs [5]. Tests have shown little interference between batch and storage services on a node. A model for production was deployed on part of the CERN EOS production service. The CMS experiment has successfully performed a test run in 2020 to profit of the BEER resources with 1.7k additional cores provided to run jobs.

5 Usage of High Performance Computing resources at the CMS experiment

In the mid- to long term future runs (run-3 and HL-LHC) the LHC experiments aim towards increasing the usage of HPC resources. The growing funding in HPC infrastructures initiate the deployment of Exa-scale machines in the near future. Countries and funding agencies push the HEP communities to make use of these resources and it is in the interest of the HEP experiments to have access to the best technologies available. HPC future contributions are regarded as integral part of the Worldwide LHC Computing Grid (WLCG) strategy towards HL-LHC. CMS has investigated the usage of HPCs resources with multithreaded jobs submitted through the Fermilab HEPCloud. The running jobs on 55k cores have been using Cori, the newest supercomputer at NERSC (National Energy Research Scientific Computing Center) located at Lawrence Berkeley National Laboratory.

6 Big Data solution - Analytics and Monitoring Infrastructure at CERN (Aka MONIT)

A new unified monitoring architecture (MONIT) for the CERN Data Centers and for the WLCG Infrastructure is based on established open source technologies to collect, stream, store and access monitoring data. The tools used are popular technologies such as Collectd, Kafka, Spark, Elasticsearch, InfluxDB and Kibana and Grafana for visualization. An overview of the available technologies is shown in Fig. 3. The monitoring infrastructure, fully based on CERN cloud resources, covers the whole workflow of the monitoring, from collecting and validating metrics and logs to making them available for dashboards and reports. In case of anomalous conditions alarms are raised and the relevant teams are automatically notified. Data sources from different subsystems are used to build complex workflows and predictive analytics (data popularity, smart caching, transfer latency) and for performance studies. The Grid Data Centers use these WLCG monitoring tools to achieve the required levels of stability, scalability and performance [6].

7 Machine Learning in High Energy Physics - Particle Classification with Deep Learning using Neural Networks

Machine learning has been applied in particle physics to solve several problems, beginning with use-cases in high-level physics analysis in the 1990s and 2000s, followed by a burst of applications in particle and event identification and reconstruction in the 2010s. Future research and development areas for machine learning in particle physics have been investigated [7]. Machine Learning methods are used for various applications, e.g. signal versus background classification with Neural Networks and Boosted Decisions Trees, object reconstruction, tagging (e.g. deep jet tagging) and for fast simulations (e.g. in calorimeter showers and time projection chamber clusters).

High Energy Physics with their particle detectors and sensors have to handle data-intensive streams. Data pipelines that move data from detectors to the data centers operate at high throughput. With particle collision events happening every 25 ns the amount of

Figure 3. Big Data solutions for computing monitoring and analytics (Aka MONIT) at the LHC experiments [6].

data flowing through the online systems of the LHC experiments is currently of the order of 1 PB/s. The collected data are filtered and data not interesting for physics analysis are thrown away to reduce the rate to the available bandwidth and storage capacity. The fast data processing is performed by a trigger system, which is organized in multiple levels. The rate transferred to the storage for later analysis is currently around 1-10 GB/s.

Improvements in the accuracy of the event filtering system of an experiment reduces wasted computing, storage and network resources, which are strongly needed for data analysis. A possible solution can be provided by topology classification with Deep Learning using Neural Networks. A classifier is used to improve the purity of data samples selected at trigger level, by distinguishing three different event topologies of interest, for example ("W + jet", "QCD", " $t\bar{t}$ "). The workflow is shown in Fig. 4. The classifier in the online filtering system of the experiment aims to improve accuracy and in particular reduce the number of false positives compared to current systems, where complex selection rules are implemented. The data processing pipeline profits from a more efficient usage of computing resources and provides improvements in data analysis. The pipeline can use tools from Big Data systems, notably Apache Spark, BigDL and Analytics Zoo [8] [9].

Figure 4. Particle Classification with Deep Learning using Neural Networks [9].

8 Conclusions

For the runs at HL-LHC the computing power and storage requirements will be increased to a much higher level, which can't be provided by technology improvements only. New computing models and more efficient software have to be developed and additionally cloud computing and HPC resources have to be used and to profit of the best technologies available. Tests on HPC and cloud facilities have proven the capability to execute efficiently both data-intensive and CPU-intensive workflows. The needs and demands of the scientific community are time-dependent with peaks and cloud resources can be added to the existing computing facility. The European Helix Nebula Science Cloud as well as the HEPCloud at Fermilab have been installed to provide the needed additional computing infrastructure.

Common Big-Data solutions based on open-source are developed to handle the vast amount of data in monitoring and operating tasks. Modern tools and methods in Machine Learning and Deep Learning are implemented to improve the program performance and efficiency and are used in many HEP tasks like analysis optimization, filtering systems and particles classification.

References

- [1] S. Chatrchyan et al., The CMS experiment at the CERN LHC. JINST 3 (2008) S08004.
- [2] I. Bird, WLCG Status Report(2019).
- [3] Helix Nebula Science Cloud <https://www.helix-nebula.eu>.
- [4] Fermilab HEPCloud <https://hepcloud.fnal.gov>.
- [5] Smith, David and Di Girolamo, Alessandro and Glushkov, Ivan et al., Sharing server nodes for storage and compute, CERN-IT-Note-2019-001, [https://cds.cern.ch/record/2653012\(2018\)](https://cds.cern.ch/record/2653012(2018)).
- [6] A. Aimar, A.A. Corman, P. Andrade et al., MONIT: Monitoring the CERN Data Centres and the WLCG Infrastructure, EPJ Web Conf. Volume 214, <https://doi.org/10.1051/epjconf/201921408031>, [http://monit.web.cern.ch/monit\(2019\)](http://monit.web.cern.ch/monit(2019)).
- [7] K. Albertsson, P. Altoe, D. Anderson et al., Machine Learning in High Energy Physics Community White Paper., arXiv:1807.02876v3 [physics.comp-ph](2018).
- [8] T.Q. Nguyen, D. Weitekamp, D. Anderson et al., Topology Classification with Deep Learning to Improve Real-Time Event Selection at the LHC., Comput Softw Big Sci 3, 12., arXiv:1807.00083v3 [hep-ex], [https://doi.org/10.1007/s41781-019-0028-1\(2019\)](https://doi.org/10.1007/s41781-019-0028-1(2019)).
- [9] Particle Classification with Deep Learning using Neural Networks Blog <https://db-blog.web.cern.ch/blog/luca-canali/machine-learning-pipelines-high-energy-physics-using-apache-spark-bigdl>.