

Exploiting Apache Spark platform for CMS computing analytics

M Meoni¹, V Kuznetsov², L Menichetti³, J Rumševičius⁴, T Boccali⁵
and D Bonacorsi⁶

¹ University of Pisa & INFN Pisa, Italy

² Cornell University, NY, U.S.A.

³ CERN, Switzerland

⁴ Vilnius University, Lithuania

⁵ INFN Pisa, Italy

⁶ University of Bologna and INFN, Italy

E-mail: marco.meoni@cern.ch, vkuznet@gmail.com, luca.menichetti@cern.ch,
justinas.rumsevicius@gmail.com, tommaso.boccali@pi.infn.it,
daniele.bonacorsi@unibo.it

Abstract. The CERN IT provides a set of Hadoop clusters featuring more than 5 PBytes of raw storage with different open-source, user-level tools available for analytical purposes. The CMS experiment started collecting a large set of computing meta-data, e.g. dataset, file access logs, since 2015. These records represent a valuable, yet scarcely investigated, set of information that needs to be cleaned, categorized and analyzed. CMS can use this information to discover useful patterns and enhance the overall efficiency of the distributed data, improve CPU and site utilization as well as tasks completion time. Here we present evaluation of Apache Spark platform for CMS needs. We discuss two main use-cases CMS analytics and ML studies where efficient process billions of records stored on HDFS plays an important role. We demonstrate that both Scala and Python (PySpark) APIs can be successfully used to execute extremely I/O intensive queries and provide valuable data insight from collected meta-data.

keywords: CERN, CMS, HEP, Big Data, Hadoop, Spark, Analytics, Machine Learning.

1. Introduction

The need for agile data analytics platforms is increasing in High Energy Physics (HEP). The current generation of HEP experiments has globally surpassed the Exa-Byte of storage with deployed computing resources exceeding one million computing cores. On top of this, computing infrastructures are highly distributed, and data transfer between sites is a major activity in the computing operations. Systems of such complexity require a careful monitoring in order to be maintained in a healthy state. At this scale processing TBytes of data on a daily basis become a real challenge. Big Data derived approaches have started to be introduced in order to get a better understanding of the computing operations and utilization of computing resources. Among others, the Elasticsearch, Hadoop Map-Reduce, and Spark are attractive solutions to handle large data sets. CERN itself is moving most of its computing dashboards from a standard RDBMS oriented approach (current dashboards use Oracle as a back-end) towards Hadoop based solutions [1]. The expected outcome of this transition is to efficiently process large data sets and aggregate information across distributed data providers. Such studies are becoming more and

more important due to the need of reducing computing resources and the associated cost with respect to general increase of physics data collected by the LHC experiments. Funding agencies and review committees are, in fact, asking for increasingly optimized computing operations which can only benefit from a deeper understanding of computing resource utilization.

Since 2015, the Compact Muon Solenoid (CMS) experiment [2] has stored large sets of computing logs (“meta-data”) about user accounting, jobs life-cycle, resources utilization, sites performance, software versions, and file access, on Hadoop file systems. The Apache Spark open-source cluster-computing framework has been evaluated as a valuable candidate to handle large amount of this meta-data stored on Hadoop file system (HDFS). Spark provides a scalable and high-performance data processing solution based on APIs offered in a variety of languages.

The benefit of using Spark for CMS experiment is two-fold. On one hand, it reduces the processing time on large datasets (DS) and manages effectively heterogeneous meta-data. On the other, it allows to perform Machine Learning (ML) studies and/or computing analytics over large DS from distributed data sources. Starting from unstructured raw data, the data preparation and analysis steps are completely defined in Hadoop infrastructure. This opens up a possibility to offload production RDBMS workload onto a Hadoop+Spark environment which is more suitable for large data aggregations and distributed joins.

In this paper we address two main use cases:

- CMS data popularity which provides various metrics about CMS user activities based on data available at CERN HDFS system. In particular, we discuss the aggregation pipeline from multiple data-streams and present useful information about user access patterns, such as number of accesses, CPU metrics versus data-tiers, sites, in visualization dashboards;
- Intelligent data placement based on modeling (via ML) and predicting popular CMS datasets to be used in analysis jobs. This activity includes fast queries for reprocessing monitoring time intervals, file-to-DS aggregation and correlation, as well as the evaluation of predictive models.

The paper is organized as follows. Sect. 2 outlines the related works, with focus on efforts applied to Big Data analytics in HEP. Then we provide description of system setup in Sect. 3 where we discuss technical parts of the data flow and underlying technologies. Then, we address the CMS data popularity use case in Sect. 4. Finally, we demonstrate how Hadoop and Spark platforms can be used to analyze large data sets in order to predict user behavior in Sect. 5 followed by discussion of using these predictions in sites caching strategies.

2. Related Work

After the introduction of Spark platform [3] for the general public in 2010 and following adaptation of SQL [4] and ML [5] within Spark, many businesses realized its power within the domain of data analytics. Small and large companies successfully applied this framework to different use cases which required processing large datasets. Up until now, the HEP uses GRID infrastructure to process individual jobs at distributed sites. In order to leverage Map-Reduce paradigm and distributed nature of Hadoop+Spark platform many aspects should be modified. Among the others, changes to middle-ware in order to read underlying ROOT data-format, adaptation of data workflow scheduling and submission tools as well as overall design of data processing [6]. Such effort is undergoing within recently funded projects such as DIANA-HEP [7]. The analytics use cases, e.g. extracting data insight from available meta-data (logs), is very well studied and adapted in a business world. Recently, HEP experiments realized that they can explore their meta-data to gain additional insight about user activities and apply this knowledge to improve their computing models [1]. In particular, exploiting Machine Learning techniques on computing meta-data is a hot topic in the HEP community at large [8]. The data popularity problem was studied by all major HEP experiments, e.g. in CMS [9, 10], LHCb [11], and

ATLAS [12]. It was found that ML approach can be useful addition to custom data placement solutions adopted by experiments. In this work, we extend previous CMS attempts [9, 10] to use CMS meta-data in the context of dataset popularity and extend discussed techniques to Spark platform.

3. System Setup and Implementation

The European Organization for Nuclear Research (CERN) has deployed two large Hadoop clusters featuring nearly 5 PBytes of raw storage. Currently, they consist of around 50 nodes equipped with 64 GBytes of RAM and 32 cores/node on a mix of Intel/AMD CPUs and CentOS7/SLC6 operating systems.

The Spark cluster is tuned and optimized to reach maximum scalability and parallelism. Every job is calibrated in favor of local processing within the same executor and reducing the data exchange between cluster nodes. The Spark paradigm consists of reading data once and processing them as much as possible in memory, minimizing data shuffling, network communication between executors, and disk I/O operations. In addition, together with the cluster resource manager (YARN in our case), Spark jobs can profit from dynamic allocation of their executors, automatically scaling in and out their number when necessary. This allows users to run the same job against folders of the same dataset without adapting different parameters according to their individual sizes.

The CERN Hadoop service provides scientists with a variety of tools and Big Data frameworks, such as Spark, Pig, Hive, which are available in different clusters. These are distinct but not isolated environments where data stored in one HDFS namespace can be accessed by jobs running indiscriminately in whatever cluster. Obviously jobs reading the local namespace will be faster, and if necessary data can be moved from one cluster to another using Map-Reduce jobs.

CMS data is stored in daily folders in CSV, JSON, or Avro format. Spark 2.0 framework, via the SparkSQL context, is equipped to read these formats. These components take as input an HDFS path and may require additional parameters depending on the kind of the format, e.g. a header of CSV format, a data schema or quotation symbols, and so on. All of them return a DataFrame used in forthcoming computations.

Since Spark is written mainly in Scala and Java, we originally used the former language for work discussed in Sect. 5. Later we found that the Scala programming language is not (yet) well adopted in HEP community. Therefore, we made significant effort to adopt our code base to PySpark framework, as explained in Sect. 3.2.

In next Sections 3.1 and 3.2, we provide details of our experience with Scala and PySpark APIs. We mostly outline pros and cons of both languages based on our practice with them.

3.1. Scala usage

Our original attempt to process large DS on HDFS was done in the Scala programming language and it reflects results described in Sect 5. Here we would like to highlight our observation of using Scala for this task. Fig. 1 shows the processing time required for DS popularity monitoring queries. The first plot shows the amount of time, in seconds, that was needed by a number of monitoring queries. The names on the x-axis are borrowed from the corresponding materialized views (MV) implemented in the current Oracle-based monitoring infrastructure. The second plot shows the reprocessing time to reproduce the results of an Oracle DB view. We found that regardless of the Oracle view or selected time period the average processing time remains in a range of few minutes¹. This makes Spark and Scala a very effective combination for reprocessing

¹ In order to exploit built-in distribution and persistence, we relied on `repartition()` and `persist()` API to split data uniformly across the computing nodes and cache them during the aggregation steps.

popularity queries on Hadoop. In fact, the same results take continuous running of incremental MV updates in Oracle back-end.

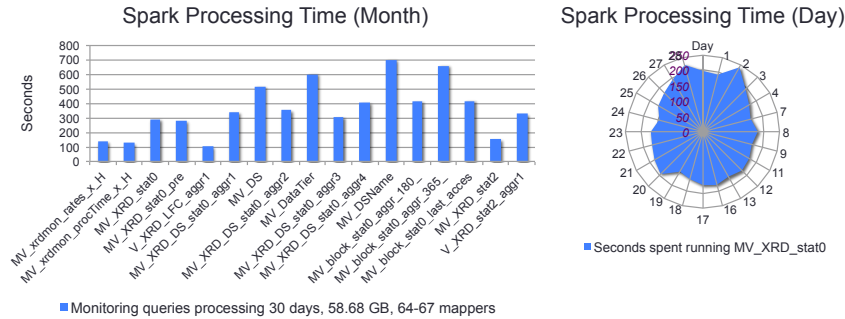


Figure 1. Fast reprocessing of DS popularity monitoring query on Apache Spark. The left plot shows monthly processing time for individual Oracle DB views implemented in Spark. The plot on a right shows processing time of one Oracle DB view for daily records.

3.2. PySpark usage

The transition from Scala to PySpark was trivial. But we found a few particularly interesting issues every programmer should know about when developing code using Python (PySpark) APIs.

Careful crafting of data is required to fit it into worker’s node memory. For instance, most of the run-time errors we experienced were related to `java.lang.OutOfMemoryError`, Java heap, and GC errors issues. A typical example would be the temptation to use `df.collect()` PySpark function to collect intermediate results. Instead, the code should perform series of operations over the data and either write them out to HDFS or run aggregation steps. Further, when the data are skewed not all workers obtain the right amount of workload. For instance, grouping or shuffling operations of non equally distributed data can cause the slowness of entire job within executor. To take advantage of parallelism the code should not iterate over the containers, instead, optimized PySpark APIs should be used, e.g. `df.foreach` in favor of `for item in df.collect()` pattern or `df.treeReduce(lambda x, y : x + y)` instead of local non cluster-distributed functions `np.sum(df.collect())`.

Users are advised to rely on PySpark functions to avoid slowness of built-in Python functions over DataFrame operations, e.g. `pyspark.sql.functions.split` should be used instead of `split` Python built-in function.

To avoid these types of pitfalls we developed our own framework called CMSSpark [13] which provided all necessary tools for code submission to Spark platform, set up necessary Java libraries for data processing (e.g. support for CSV or Avro data formats) as well as set up proper run parameters optimized for CERN Spark cluster.

This framework significantly helped new developers to easily adopt their business logic into framework and quickly produce desired results. For instance, any user with access to the CERN Hadoop cluster can quickly launch distributed (Py)Spark job to process CMS data, focusing on what operation to run rather than struggling on how to properly submit YARN-based batch job.

4. CMS Data Popularity

The common approach of data analytics is based on RDBMS solutions. At CERN, the Oracle DB is a de-facto standard for many bookkeeping and analytics services. Even though it works extremely well we see that the amount of information HEP experiments want to process is growing exponentially each year. We already found certain limitations on data processing pipelines within Oracle DB. For instance, the CMS most populated DBs, Data Bookkeeping System (DBS) and Data Location System (PhEDEx) [14], already contain a few hundred GBytes of data. Merging data only from those two DBs becomes an issue. Even though they reside in one physical DB, the cross-joins among them caused large latencies on complex queries due to large data volume. Moreover, growing number of collected meta-data about user jobs which resides outside of Oracle DB may provide an additional insight into data utilization by our infrastructure. Spark platform provides a valuable solution to overcome these obstacles.

The CMS, among other HEP experiments, has joined a new trend of storing various meta-data to HDFS starting in 2015. This effort was driven by individual users and groups willing to investigate a new technology to analyze unstructured data, such as logs from user based jobs. By 2017 we already collected billions of records from various data providers. Here we present a successful effort showing utilization of Hadoop+Spark platform for analytics purposes and modeling of DS popularity based on available meta-data. We start our discussion with CMS analytics use case where significant portions of these data were used.

4.1. Data Processing

We used Spark jobs to acquire various information about user activities such as total number of accesses to dataset and processing time of analysis jobs at specific sites. We identified four data-streams (in the form of logs) representing this information: AAA [15] logs of remote XrootD accesses; EOS [16] logs of file accesses on CERN EOS system; CMSSW [17] logs capturing information about CMS software framework, and CRAB [18] logs representing user analysis jobs. All of these data are accessible on HDFS.

Each data-stream contains various metrics associated with a job in question, e.g. name of the file used during a job, user Distinguished Name (DN), processing time of the job, and similar. In order to get insight into data present in these data streams they should be properly cleaned, sorted, filtered, and joined with data from other sources such as DBS and PhEDEx DBs which provide information about dataset, block and site names. The latter two DBs are quite large, their current size is at the level of a few hundred GBytes each. In order to answer simple questions about users accessing a particular dataset on a certain site(s), it was required to join various attributes from multiple streams along with DB tables from DBS and PhEDEx Oracle DBs. We found that processing time to join and aggregate data across DBS/PhEDEx DBs and data-streams is impossible to achieve in a reasonable amount of time using a vertical approaches based on RDBMS solutions². At the same time a simple dump of DBS and PhEDEx DBs into HDFS and performing a regular Spark job to join desired data-streams with them allows quickly (on a scale of minutes) to get insight into data we were looking for. For instance, a daily snapshot consists of about 130M log records which were joined at a file level with DBS and PhEDEx DBs, grouped and aggregated to roughly a few thousand records in under an hour of Spark processing time. We measured that monthly stats can be obtained in about 12 hours. This allowed us to build a series of dashboards to look at data utilization in the CMS experiment.

² A typical job to aggregate information from Oracle DBs and a single data-stream required many hours to complete using materialized view across all data.

4.2. Data Analytics and Monitoring

There are different approaches to data monitoring. On one hand, someone can delegate pushing information from individual GRID sites and applications into the central monitoring system. On the other, such information can be obtained from an individual logs stored in a common place (e.g. central HDFS) via external jobs. The former approach requires proper instrumentation of monitoring metrics and coordination of data injection from all participating data providers. The latter relies on significant processing power to collect and aggregate information in a reasonable time.

The ability to efficiently process large data sets on HDFS via Spark jobs opens up a possibility to build various dashboards based on user access metrics extracted from various data-streams. For that purpose we used CERN MONIT system [1] based on ElasticSearch and Kibana tools. The data processed by Spark jobs were injected as JSON documents into CERN MONIT system via Apache ActiveMQ Stomp [19] protocol. We achieved a reduction factor of 5000 coming from meta-data records (logs) on HDFS to aggregated records pushed into CERN MONIT system. Fig. 2 shows one part of the CMS popularity dashboard we built.

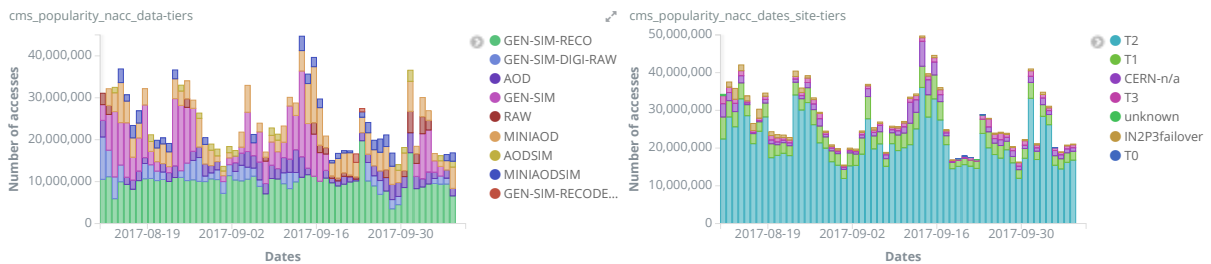


Figure 2. CMS popularity dashboard based on processing four CMS data streams: AAA, EOS, CMSSW and CRAB. The displayed part of the dashboard represents the number of accesses vs data-tiers and site-tiers as time series plots.

Even though various dashboard plots were useful to get a broader view on daily activities of CMS users, we were also able to address specific questions raised by various users. For example, how much data resides on all T2 sites with detailed breakdown of number of events and replica sizes of common used data-tiers.

While looking at the current status of resource utilization, which is a very important component of our daily operations, we oversee that it can be further enhanced by predicting user behavior and adjusting our resources accordingly.

5. Modeling CMS Popularity

Usage of Machine Learning in HEP was mostly used in physics analysis such as discrimination of signal versus background, clustering algorithms, object reconstructions, or similar tasks. Recently, ML ideas started penetrating into other domains of HEP such as helping physicists better understand computing resources. For instance, ML modeling has been discussed in the scope of dataset placement optimization, reduction of transfer latencies as well as in network-aware applications for identifying anomalies in network traffic, predicting congestions and WAN path optimization. Here we present results build on top of previous studies for predicting dataset popularity in context of better data placement strategies. We extend these results by demonstrating how to predict dataset popularity using larger datasets and use them to build strategies for better data placement at participating sites. In fact, DS placement at most WLCG experiments is based on DS popularity which should be considered to make an optimal choice of

replication to maximize data availability for processing and analysis needs. As shown in [9, 20], CMS dataset popularity can be successfully predicted using ML approach.

In this work, the generation of ML dataframes involved several processing steps on Spark and it was achieved through a Scala reduction pipeline shown in Fig. 3. It highlights the number of samples and the computing time at each step of the reduction chain. In the shown example, we start from two billion AAA file access logs, merge them with PhEDEX DB to extract dataset and site information and end-up with half a million records suited for ML studies.

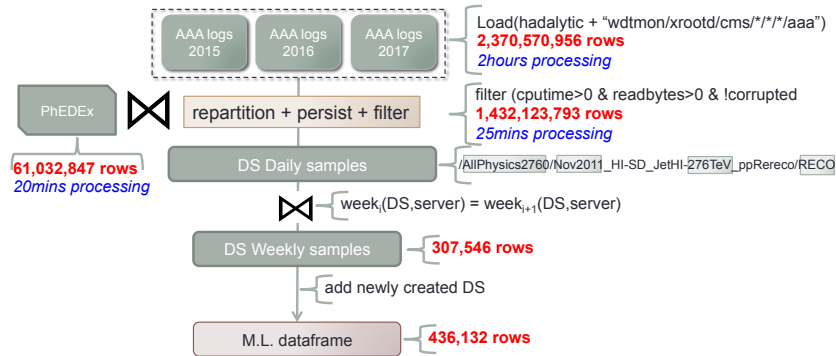


Figure 3. An example of reduction pipeline of Spark components from file AAA access logs level to a dataframe suitable for Machine Learning input.

The input dataframe for popularity predictions consisted of a set of attributes and a popularity label. This label is computed on the basis of the usage metric designed as popularity feature. When usage metric in some $time_period_{i+n}$ is higher than a given threshold, the samples corresponding to $time_period_i$ is marked as popular. Our choice for timespan n was chosen to be one week. This was justified by empirical observation. When the temporal window in the training set was increased to 3 weeks the number of positive samples decreases remarkably and prediction accuracy decreased. In fact, Fig. 4 shows that the number of DS accesses over 3 weeks time frame is about 50% lower than across 2 weeks, which explains that most DS have a time-limited access pattern.




Figure 4. Dataset access patterns during 2 and 3 week period of time.

After applying a cut on number accesses to a dataset, we converted used dataset into classification problem which later was trained on Spark via various classifiers available in Spark MLlib. We achieve the same level of accuracy as reported in [9].

The outcomes from the classifiers are used to implement a data caching policy based on popularity predictions. The datasets are stored at various CMS Grid sites to minimize the network traffic and spread the load for analysis jobs. The current site placement policy is based on assumptions on the expected needs for datasets, mostly using the data-tier. The use of storage, at least for small-to-medium sites, is based on Least Recently Used (LRU) caching systems in order to select which elements should be removed when a site reaches its capacity. Instead, we design a different caching policy that relies on knowledge of DS popularity in the

next week and keeps DS in the cache if it is still popular. If a given DS is selected for eviction but is marked as popular by the prediction model, then it is maintained inside the cache.

Preliminary results are detailed in [20], with the main difference w.r.t LRU observed at lower caching size threshold. We always assume a cold start at each site which generates a number of compulsory cache misses corresponding to the first reference to each distinct DS. Our caching strategy outperforms LRU by up to 20% when the cache size is limited. This makes it very effective in production sites with limited storage or bandwidth. It paves the way for further comparison with others caching strategies and for research aiming at enabling smart policies for dataset replica management based on demonstrated ML approach.

6. Conclusions and Outlook

Spark has proven to be a very solid framework for data analytics and ML. It is extremely effective for crunching large datasets available on HDFS and extracting aggregation metrics from underlying meta-data. It can be integrated into HEP workflow via PySpark APIs or stand-alone application written in Scala or Python. We discussed possible pitfalls programmers should be considered using either of the programming languages.

We demonstrated that Spark framework represents a valuable alternative to RDBMS solutions for analytics needs. It allows to (re-)process and aggregate various metrics from multiple data-streams efficiently in distributed nature. We were able to access logs from different data streams including AAA, EOS, CMSSW and CRAB and join them with the HDFS dumps of two largest CMS databases, DBS and PhEDEx. As a result the data can be easily aggregated into set of dashboards based on CERN MONIT system and be used for various monitoring tasks within collaboration.

We also implemented a scalable pipeline of Scala components to accomplish mission-critical data reduction and ML processing from billion of records available at CERN HDFS system and build ML models that can be used towards data-driven approach in CMS computing infrastructure. For example, dataset popularity predictions can play a significant role in intelligent data placement at GRID sites both for newly created datasets and holding dataset samples frequently accessible by CMS physicists. A clever caching strategies based on ML predictions can lead to cost-effective data placement across GRID sites and better utilization of available computing resources.

In conclusion, CMS considers the Spark platform as a potentially crucial component in the ecosystem that would enable LHC experiments to attack Big Data problems in the long run.

References

- [1] Aimar A et al 2017 *Unified Monitoring Architecture for IT and Grid Services*, *J. Phys.: Conf. Series* **898** 092033
- [2] Chatrchyan S et al 2008 *The CMS Experiment at the CERN LHC*, *JINST* **3** S08004T
- [3] Matei Z et al 2010 *Spark: Cluster Computing with Working Sets*, *HotCloud* 10–10
- [4] Armburst M et al 2015 *Spark SQL: Relational Data Processing in Spark*
- [5] Meng X et al 2016 *MLlib: Machine Learning in Apache Spark*, *J. Machine Learning Research* **17** 1-7
- [6] Gutsche O et al 2016 *Big Data in HEP: A comprehensive use case study*, *Preprint* arXiv:1703.04171
- [7] DIANA HEP project, part of NSF's Software Infrastructure for Sustained Innovation (SI2) program, <http://diana-hep.org/pages/activities.html>
- [8] HEP Software Foundation white papers: <http://hepsoftwarefoundation.org/activities/cwp.html>
- [9] Kuznetsov V, Li T, Giommi L, Bonacorsi D and Wildish T 2016 *Predicting dataset popularity for the CMS experiment*, *J. Phys.: Conf. Series* **762** 012048
- [10] Meoni M, Boccali T, Magini N, Menichetti L and Giordano D 2017 *XRootD popularity on Hadoop clusters*, *J. Phys.: Conf. Series* **898** 7 072027
- [11] Hushchyn M, Charpentier P and Ustyuzhanin A 2015 *Disk Storage Management for LHCb Based on Data Popularity Estimator*, *J. Phys.: Conf. Series* **664** 042026
- [12] Beermann T et al 2014 *Popularity Prediction Tool for ATLAS Distributed Data Management*, *J. Phys.: Conf. Series* **513** 042004

- [13] <https://github.com/vkuznet/CMSSpark>
- [14] Giffels M et al *The CMS Data Management System*, *J. Phys.: Conf. Series* **513** 042052
- [15] Xrootd project page: <http://www.xrootd.org/>
- [16] Mascetti L et al 2015 *Disk storage at CERN*, *J. Phys.: Conf. Series* **664** 042035
- [17] CMS software framework, <https://github.com/cms-sw/cmssw>
- [18] Spiga D et al 2008 *CRAB: The CMS distributed analysis tool development and design* Nuclear Physics B - Proceedings Supplements 177-178 **1** 267-268
- [19] Apache ActiveMQ Stomp, <http://activemq.apache.org/stomp.html>
- [20] Meoni M, Perego R and Tonello N 2017 *Popularity-based caching of CMS datasets*, ParCo