

Physics Performance and Fast Turnaround: the Challenge of Calibration and Alignment at the CMS Experiment during the LHC Run-II

Salvatore Di Guida^{*†}

Università degli Studi Guglielmo Marconi, INFN Sezione di Napoli, Italy

E-mail: salvatore.di.guida@cern.ch

Gianluca Cerminara

CERN, Geneva, Switzerland

Giovanni Franzoni

CERN, Geneva, Switzerland

Giacomo Govi

Fermi National Accelerator Laboratory, Batavia ILL, USA

Marco Musich

Université Catholique de Louvain, Belgium

Andreas Pfeiffer

CERN, Geneva, Switzerland

The CMS detector at the Large Hadron Collider (LHC) is a very complex apparatus with more than 70 million acquisition channels. To exploit its full physics potential, a very careful calibration of the various components, together with an optimal knowledge of their position in space, is essential. The CMS Collaboration has set up a powerful infrastructure to allow for the best knowledge of these conditions at any given moment. The quick turnaround of these workflows was proven crucial both for the algorithms performing the online event selection and for the ultimate resolution of the offline reconstruction of the physics objects. The contribution will report about the design and performance of these workflows during the operations of the 13TeV LHC RunII.

ICHEP2016: The 38th International Conference on High Energy Physics

3–10 August 2016

Chicago, IL (United States)

*Speaker.

†Also with CERN, EP Department.

1. Introduction

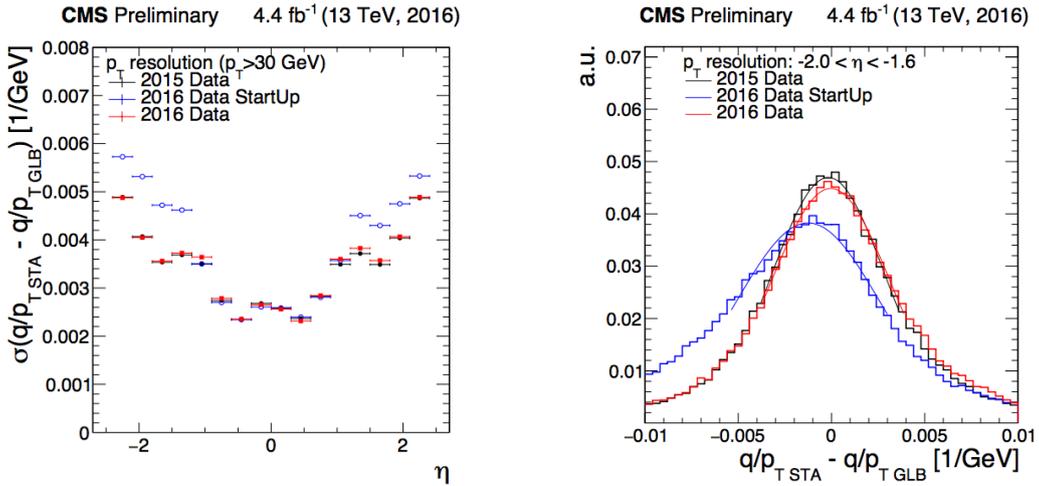
The Compact Muon Solenoid (CMS) is one of the two general purpose experiments collecting events from the collisions delivered by the Large Hadron Collider (LHC) at the European Organisation for Nuclear Research (CERN) in Geneva, Switzerland. The CMS primary goal is to explore physics at the TeV energy scale. This ranges from the search of the Higgs Boson to the searches for new physics beyond the Standard Model, to the precision measurements of already known particles and phenomena. The successful Run I data taking (2009-2013) at 7-8 TeV has led to the discovery of the Higgs Boson in 2012. The ongoing Run II data taking at 13 TeV is exploiting the sophistication of the detector for better digging into the complexity of the Standard Model, and for exploring the new energy domain for Supersymmetry and Exotica searches.

CMS has several cylindrical layers coaxial to the beam direction (the barrel) closed at both ends by disks which are orthogonal to the beam pipe (the endcaps). The central part of the detector is a 13 m long superconducting solenoid of 13 m diameter, providing a 3.8 T axial magnetic field. The magnet system is complemented by a steel flux-return yoke which has a cylindrical structure and is placed outside the solenoid.

CMS consists of different subdetectors. Starting from the beam line, within the superconducting solenoid volume there are a silicon pixel and strip tracker to measure the momentum of the charged particles and to identify primary and secondary vertexes, a lead tungstate crystal electromagnetic calorimeter (ECAL) for measuring the energy and the position of photons and electrons, and a brass scintillator hadronic calorimeter (HCAL) to measure the energy of both neutral and charged hadronic particles. The CMS muon system is hosted within the flux-return yoke, and allows for the identification of muons and a precise measurement of their momentum. It relies on three kinds of gaseous detectors: drift tubes (DT), cathode strip chambers (CSC), and resistive plate chambers (RPC). A more detailed description of the CMS detector, including the experimental apparatus, the definition of the coordinate system, and the relevant kinematic variables, is provided in [1].

2. The Calibration Challenge

To achieve the CMS physics goals excellent performances in particle reconstruction and identification are necessary. A proper selection and reconstruction of collision events requires “secondary data”, the so-called condition data. They encapsulate the configuration required to bring the detector in running mode, describe the state of the experimental apparatus, and provide alignment and calibration constants. The workflows which yield calibration datasets are intrinsically iterative: they describe the status of millions of electronics channels (such as each crystal in ECAL) and modules (like in the silicon strips), so they depend on different procedures, which can be correlated with one another. Thus, the calibration workflows are able to produce data with improved precision when running on samples with larger statistics. This allows to improve the resolution of relevant physics observables computed from reconstructed data, as you can see in figure 1 for the muon transverse momentum. At the same time, while collecting data, physicists can understand better and better the evolution of the running detector, and develop newer and more efficient calibration algorithms, as well as define a finer-grained model representing detector conditions. This implies



(a) Width of the p_T resolution in bins of η using different muon track-based geometries. (b) The distribution of the p_T resolution is shown for the specific region $-2 < |\eta| < -1.6$

Figure 1: The muon p_T resolution is computed in a Z-enriched sample, by looking at the difference between the q/p_T for standalone and global muons. A selection based on the Muon p_T (greater than 30 GeV/c) and on the number of tracker hits and the χ^2 is performed. At low $|\eta|$ the newest geometry (2016 Data) performs similarly to the older ones, as DT have not been moved. At high $|\eta|$ region the new geometry brings a major benefit due to the alignment of the CSC disks.

that the software and computing infrastructure should be flexible enough in order to feed calibration workflows with dedicated event data, and also to quickly integrate and deploy fixes and upgraded implementations.

The possibility to produce calibrations with short turnaround has several advantages on detector operations, physics data production, and computing management. As far as data collection and selection is concerned, accurate calibrations allow for an efficient online event selection by the High Level Trigger (HLT) [2] while keeping the acquisition rate under control. A quick turnaround for the calibrations allows to process datasets within few hours, and deliver them for analysis within a few days from their acquisition. If the very first data reconstruction of physics objects consumes efficient calibrations, there can be less iterations of data re-processing: the computing resources of the experiment can therefore cope with an higher data rate.

3. The CMS Condition Database

Condition data change with time and require frequent time-based lookups. They are stored into relational databases and accessed during the execution of the reconstruction applications. The CMS Condition Database system [3][4] depicted in figure 2: it relies on one Oracle 11g Real Application Cluster and its Active Standby copy, housing two Database services:

- OMDS allows to handle the configuration and condition data produced online by the sub-detectors in relational tables;

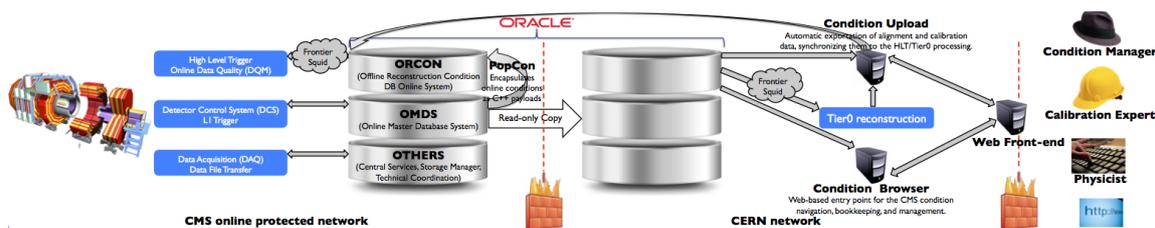


Figure 2: The CMS Condition Database model.

- ORCON serves a subset of condition data needed by HLT, online and offline data quality monitoring (DQM), and offline reconstruction. Data are written and retrieved as C++ objects.

A REST based, hierarchical access layer (FRONTier [5]), by caching the results of the queries, allows an efficient access to the conditions data from the thousands of jobs processing the data on the LHC Computing Grid, as well as from the event selection jobs on the HLT computer farm. In addition to the Oracle based central storage, users can store their data in private files, which are also serving as interchange format. These files are handled with the SQLite software, using the same database structure as for the master database.

A *Payload* represents the set of parameters describing a condition data item to be consumed in event data processing. It is associated to an user-defined C++ entity class: the class instances encapsulate payload items. Payloads are transferred as binary arrays, with no assumptions on their layout: payloads objects are stored as BLOB data types. The time period in which a Payload provides a suitable description of the detector state is specified using an *Interval of Validity (IOV)*: time can be represented by a timestamp, a run number¹, or a combination of run and luminosity section number². A set of IOVs and the corresponding payloads builds up a consistent set of conditions, uniquely identified with a *Tag*. A *Global Tag* identifies a coherent set of tags providing a complete set of condition data to be consumed in a given workflow.

For security and availability reasons, the master database is sitting in the CMS online protected network. A central service, the *Condition Uploader*, provides a central access point for the upload of new calibration data into the database. It provides authentication and authorization for the access to the master database, synchronizes — when necessary — the IOV with HLT or Tier0 processing, and inserts the payload objects and the corresponding metadata into the production database.

4. Automated Calibrations

The requirements for calibration data production, in particular the low-latency use case, have a direct impact on the design of the online data collection, and the types of event streaming to be tailored:

- The event data streams for physics collision data have an average rate of $\approx 1\text{kHz}$;

¹A CMS run is a time period where the detector conditions and the data acquisition are uniform and coherent for recording and selecting physics events. A run can last between less than one and several hours depending on the length of the proton fill in the LHC and the detector status.

²A luminosity section is a time measurement within a CMS run: it corresponds to 23.3 s.

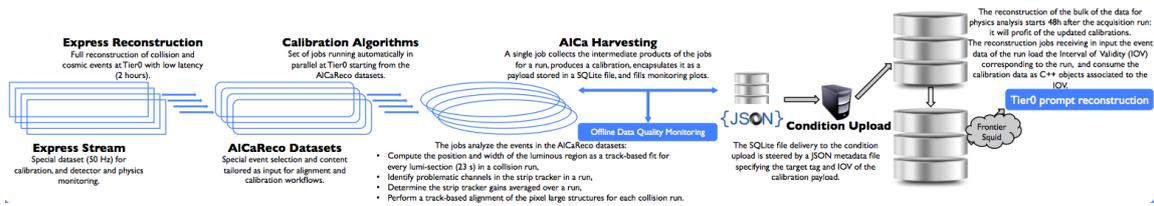


Figure 3: Schema of the automated calibrations running at the Tier0, based on Express Stream.

- A set of calibration streams, containing special kinds of events feeding calibration workflows;
- A limited fraction of physics events, constrained by the bandwidth (≈ 50 Hz), to assess physics performance with prompt feedback: this is the *Express Stream*;
- A dedicated stream of physics and calibration events for online DQM purposes (≈ 100 Hz), allowing to monitor the detector status while taking data.

The physics express and bulk streams produced by HLT are processed with different latency at the Tier0 farm at CERN:

- The *express processing* runs full physics reconstruction on the express stream as soon as available. The resulting physics objects for a run are accessible within few hours after data collection, so they allow to give prompt feedback about the detector status and physics performance, and to provide data for calibration workflows.
- The *bulk or prompt processing* performs the full reconstruction of the physics streams for physics analysis. It is delayed by 48 hours in order to allow the computation of some calibrations which can rapidly change.

We can identify three kinds of calibration workflows:

- *Online*: a set of applications produces conditions to be consumed by HLT and express processing within few minutes. These describe the detector configuration and state, such as modules included in data taking, hardware trigger settings, and magnetic field.
- *Low-latency*: the express reconstruction yields pre-defined datasets which are used to compute a well identified subset of calibrations consumed by bulk processing at Tier0;
- *Offline*: a fine-grained set of calibrations of the various parts of the detector are produced from the the full reconstruction both of physics streams in bulk processing, and of calibration streams, possibly using dedicated calibration data tiers, and keeping into account the interdependencies between them; the new datasets are later consumed in data reprocessing campaigns ensuring ultimate accuracy.

In the following, we will concentrate on the low latency calibration workflows, in particular on their automation. The automated workflows are executed for each CMS run. Figure 3 shows a schematic view of the infrastructure set up to compute the calibration data. The raw data in the express stream for a run are processed by parallel jobs using the CMSSW framework [6] on the Tier0 computing

farm [7]. These jobs perform the reconstruction of the physics objects for analysis purposes; at the same time, they split the output event collections into smaller datasets, labeled as `AlCaReco`. The `AlCaReco` datasets are tailored for the calibration algorithms and optimized both in terms of event selection and event content. Next, each calibration algorithm is run in parallel automatically at the Tier0, with one `AlCaReco` dataset as input. The job splitting feeds each application with events in the same luminosity section: this allows to compute calibration constants for a run with the granularity of at least one luminosity section. The jobs produce as output intermediate products specific to the given quantity that needs to be calibrated. In the final step, called *AlCa Harvesting*, a single job collects the intermediate products per each calibration workflow, produces one or more calibration constants, and encapsulates them as condition payload objects, which are stored, together with the corresponding IOV, into a `SQLite` file. The `SQLite` file is then sent to the Condition upload service together with a metadata file in `JSON` format specifying the target tag and IOV of the calibration payloads. Monitoring histograms in `ROOT` format are also produced by modules running the DQM component [8], and uploaded to the web-based Graphical User Interface for inspection of the performance of the calibration algorithms by the detector experts.

These are the calibration workflows currently running at the CMS Tier0:

- Identification of problematic channels in the silicon strip tracker in a run, for event-by-event optimization of the pattern recognition reconstructing the trajectories of charged particles;
- Determination of the charge gains in the silicon strip tracker sensors averaged over a run, for correcting radiation induced effects;
- Computation of the position and width of the luminous region as a track-based fit for every luminosity section in a collision run: the procedure and some results of the 2016 proton-proton collision data are described in figure 4;
- Track-based alignment of the large structures of the silicon pixel detector in each collision run: the details and the results are shown in figure 5.

5. Concluding Remarks

An efficient reconstruction of the particles and a correct measurement of the physics observables rely on the alignment and the calibrations of each sub-detector, as well as between each other. The calibration workflows must provide efficient calibrations to ensure the full resolution and the best physics performance of the CMS detector. The CMS collaboration deployed a robust and flexible infrastructure delivering accurate calibrations with a fast turnaround so that they can be consumed by the bulk processing at Tier0, thus minimizing the need for offline reprocessing of the data. With the huge amount of data collected up to now in the second LHC run, calibration algorithms became more and more complex. This implies that the statistics accumulated over a single run may not be sufficient to reach the required precision of the calibration constants. After an intense development cycle in 2016, in view of the 2017 proton-proton collision data taking, the infrastructure will be able to compute calibration data by harvesting over multiple runs with similar configuration. This topic is covered in detail in [10].

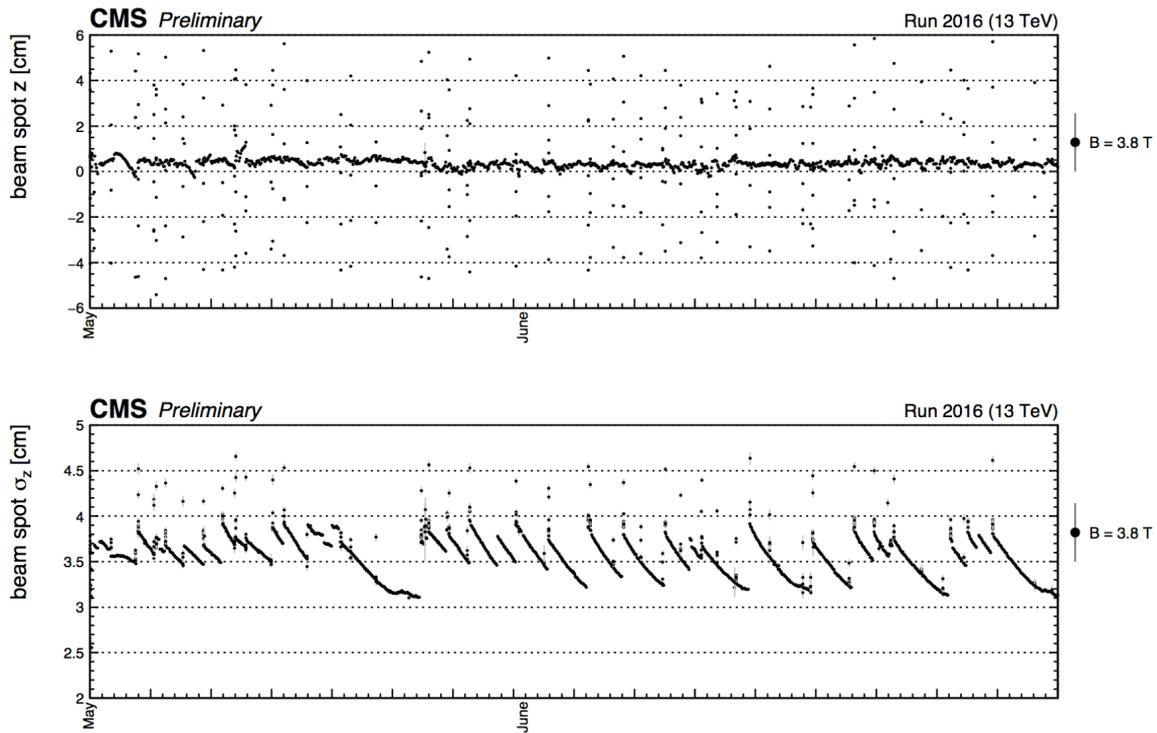


Figure 4: z position and z -longitudinal dimension of the luminous region in CMS interaction point, with respect to the CMS reference frame centred on the mechanical support of the Tracker. The luminous region quantities are determined through a fit to the three-dimensional distribution of the primary vertexes and through a fit to the track impact parameter d_0 vs track azimuthal angle ϕ distribution. These fits are performed for each luminosity section. Subsequently, the results are averaged over contiguous luminosity sections (IOV). An IOV is closed and a new one is opened if the boundary of 60 luminosity sections is reached or if any quantity drifts significantly. The resulting values for the z position and z -longitudinal dimension of luminous region for each IOV are shown in figure as a function of time using the single luminosity section as unit. The data considered in this plot have been collected by CMS in 2016 during proton collisions at 13 TeV.

References

- [1] CMS Collaboration, “The CMS experiment at the CERN LHC”, *JINST* **3**(2008) S08004, doi:10.1088/1748-0221/3/08/S08004.
- [2] CMS Collaboration, “CMS High Level Trigger”, CERN/LHCC **2007-021** (2007).
- [3] Di Guida S. et al., “The CMS Condition Database system”, *J. Phys.: Conf. Ser.* **664** 042024 (2015), doi:10.1088/1742-6596/664/4/042024
- [4] Di Guida S. et al., “Multi-threaded Object Streaming”, *J. Phys.: Conf. Ser.* **664** 042044 (2015), doi:10.1088/1742-6596/664/4/042024
- [5] Dykstra D., “Scaling HEP to Web size with RESTful protocols: The frontier example”, *J. Phys.: Conf. Series* **331** 042008 (2011).
- [6] Jones, C.D. and Sexton-Kennedy, E., Stitches Together: Transitioning CMS to a Hierarchical Threaded Framework, *J. Phys.: Conf. Ser.* **513** 022034 (2014), doi:10.1088/1742-6596/513/2/022034

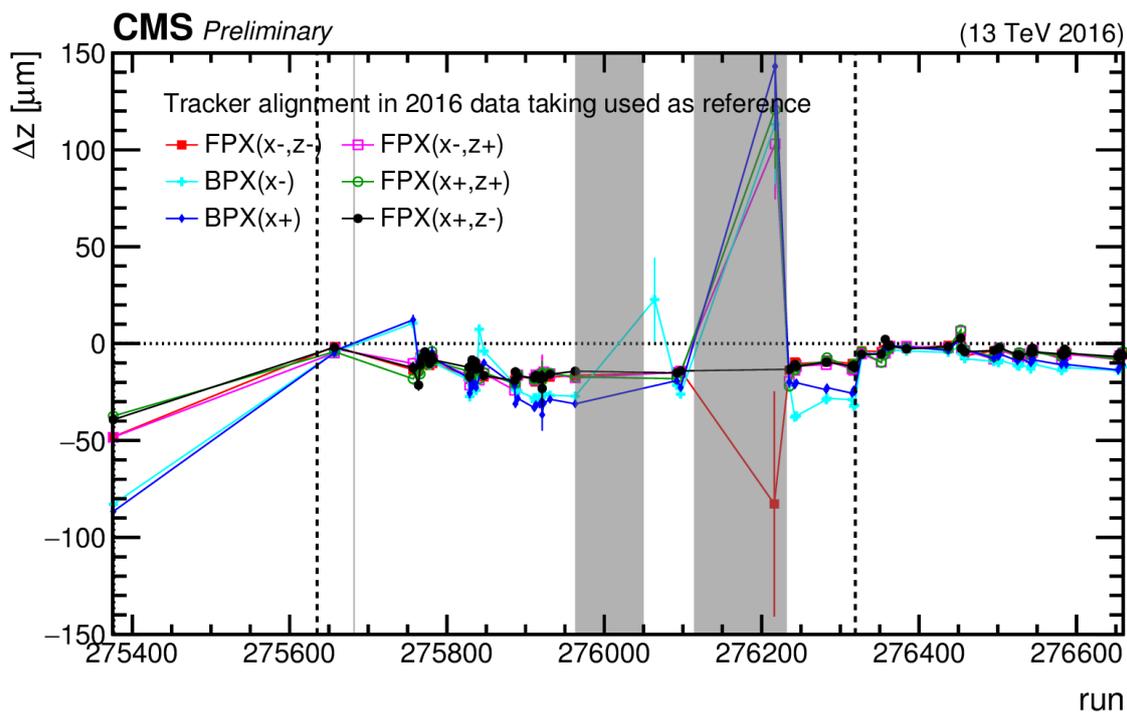


Figure 5: CMS is employing an automatic procedure to continuously monitor movements of the high-level structures of the pixel tracker [9], which occur for example due to temperature changes or changes of the magnetic field. For each run with more than 20,000 events, an alignment of the high-level structures is performed automatically, measuring the movements relative to the geometry used in data processing. When appropriate, the geometry is updated with the results of this alignment. The picture shows the observed movements in the global z direction of the high-level structures of the pixel detector monitored as a function of run number. The run range covers the time from June 21 to July 12 2016, and the recorded data corresponds to an integrated luminosity of 7 fb^{-1} . Error bars represent the statistical uncertainties of the measurement. Grey bands represent runs during which CMS magnet was not at 3.8 T. Vertical dashed lines illustrate updates of the pixel high-level structure reference geometry, after which the mis-alignment is cured. Typical movements during magnet-off periods are smaller than $50 \mu\text{m}$ in x and y , and smaller than $150 \mu\text{m}$ in z . These movements are mostly recovered after a magnet cycle, but this is not expected.

[7] Jones, C.D. et al., Using the CMS Threaded Framework In A Production Environment, *J. Phys.: Conf. Ser.* **664** 072026 (2015), doi:10.1088/1742-6596/664/7/072026

[8] Rovere, M., The Data Quality Monitoring Software for the CMS experiment at the LHC, *J. Phys.: Conf. Ser.* **664** 072039 (2015), doi:10.1088/1742-6596/664/7/072039

[9] CMS Collaboration, Alignment of the CMS tracker with LHC and cosmic ray data, *JINST* **9**(2014) P06009, doi:10.1088/1748-0221/9/06/P06009

[10] Oramus, P., Continuous and fast calibration of the CMS experiment: design of the automated workflows and operational experience, *CHEP 2016 Proceedings*, to appear.