

The CMS Electromagnetic Calorimeter workflow

Chiara Rovelli^{1,*} on behalf of the CMS Collaboration

¹Istituto Nazionale di Fisica Nucleare, Sezione di Roma - Piazzale Aldo Moro 2, 00185, Roma (Italy)

Abstract. The CMS experiment at the LHC features an electromagnetic calorimeter (ECAL) made of lead tungstate scintillating crystals. The ECAL energy response is fundamental for both triggering purposes and offline analysis. Due to the challenging LHC radiation environment, the response of both crystals and photodetectors to particles evolves with time. Therefore continuous monitoring and correction of the ageing effects are crucial. Fast, reliable and efficient workflows are set up to have a first set of corrections computed within 48 hours from data-taking, making use of dedicated data streams and processing. Such corrections, stored in relational databases, are then accessed during the prompt offline reconstruction of the CMS data. Twice a week, the calibrations used in the trigger are also updated in the database and accessed during the data-taking. In this note, the design of the CMS ECAL data handling and processing is reviewed.

1 Introduction

The Compact Muon Solenoid (CMS) [1] is a general purpose detector operated at the CERN Large Hadron Collider (LHC). The CMS electromagnetic calorimeter (ECAL) [2] is made of lead tungstate scintillating crystals. Its design was mainly driven by the need of an excellent energy resolution, as dictated by the $H \rightarrow \gamma\gamma$ process which is characterized by a narrow peak over a large background. ECAL is the largest crystal calorimeter ever built for a high energy physics experiment. Although providing an optimal energy resolution, crystal calorimeters require constant monitoring to correct for radiation induced effects and periodic calibrations. The high ECAL complexity, together with the challenging LHC Run 2 operating conditions, required an elaborated workflow for the ECAL calibration and monitoring. Such workflow is documented in this note.

2 The CMS ECAL

ECAL is a compact, homogeneous calorimeter made of 75848 $PbWO_4$ scintillating crystals. The crystals are distributed in a barrel (EB), covering the pseudorapidity region $|\eta| < 1.48$, and two endcaps (EE) extending the coverage up to $|\eta| = 3$. A preshower detector (ES) based on lead absorbers equipped with silicon strip sensors is placed in front of the endcaps and covers the $1.65 < |\eta| < 2.6$ region. Silicon avalanche photodiodes and vacuum phototriodes, both with internal amplification factors, are used as photodetectors in the barrel and endcaps respectively. The front-end electronics uses 12-bit analogue-to-digital converters (ADC) to

*e-mail: chiara.rovelli@roma1.infn.it

sample the analogue signals from the detector at 40 MHz. In EB and EE, ten consecutive samples are stored for each received trigger, while in ES, three samples only are stored. The ECAL response varies under irradiation due to the formation of colour centres that reduce the crystals transparency. ECAL is equipped with a laser light injection system [3], that allows the monitoring of each crystal's transparency. Laser pulses are injected in each crystal by groups of a few hundreds crystals. A relative response measurement is obtained by normalising the crystal response to reference silicon PN photo-diodes. To ensure redundancy in the system and the possibility of probing the transparency with different wavelengths, two laser lights (blue and green) are used. Blue and orange LED light is also flashed in the endcap crystals.

3 ECAL calibration steps

Having accurate response corrections available with a short turnaround is crucial for ECAL. Online, a precise response allows an efficient event selection by the trigger system, while keeping the acquisition rate under control. Offline, frequently updated corrections guarantee a good dataset for analysis within a few hours from the data acquisition. The CMS strategy is based on a 48-hour delay between the data-taking and the reconstruction of the dataset used for physics analysis, called 'prompt reconstruction'. During the two-day time period, a few low-latency calibration workflows are run by the CMS sub-detectors. For what concerns ECAL, the transparency corrections and the electronics pedestal measurements are promptly computed and exploited for the first round of reconstruction. Other quantities, such as the alignment with respect to the CMS tracker or the energy scale, are updated once or more often during the year if needed due to detector changes. Finally, large data samples are needed to derive channel-to-channel intercalibration (IC) constants, which are therefore computed typically once at the end of each year data-taking. Either the prompt reconstructed data, or dedicated calibration streams, are used to monitor and calibrate the ECAL detector.

3.1 ECAL calibration streams

CMS uses a two-level trigger system to reduce the event rate [4, 5]. The first level (L1) uses custom electronics to analyse coarse data from the calorimeters and muon detectors only. The second level (High Level Trigger, HLT) runs on a computing farm and analyse the information from all sub-detectors in order to decide whether to record an event or not.

Two dedicated trigger streams have been developed to have frequent and high granularity corrections of the ECAL conditions. They run at the online computing farm on the events accepted by the L1 trigger. Only a reduced event content is stored, with a typical size of a few kBytes per event while the RAW data size is around 1 MB/event. This allows the selection of events at a high trigger rate even if the relevant triggers would be heavily prescaled in the standard CMS Physics trigger menu.

The first stream aims at selecting photon pairs produced in neutral pion or η decays ($\pi^0 \rightarrow \gamma\gamma$, $\eta \rightarrow \gamma\gamma$) in QCD events or in events with electromagnetic objects accepted by the L1 trigger. To comply with the demanding environment of the online farm, the selection of π^0 and η candidates is based on ECAL variables only, from localized regions around the L1 candidate. The ten digitized samples (ECAL digis) from all the crystals in regions where a π^0 or η candidate is reconstructed are stored, while the remaining part of the event is dropped to limit the required trigger bandwidth. With this approach, a trigger rate up to 7 kHz for π^0 's and η 's in different detector regions could be sustained during the full LHC Run 2. Data selected by this stream are used to monitor the ECAL response stability in time and to compute IC constants.

The second stream is based on a zero bias trigger, i.e. a beam bunch crossing time trigger. To limit the bandwidth, only the digis corresponding to single crystal energy deposits not consistent with the expected noise are recorded, while the remaining part of the event is dropped. In this way it was always possible to maintain the event size below a few kBytes only during the LHC Run 2, and to sustain a trigger rate of 3 kHz or more for the stream. Data selected by this stream are used for the detector calibration and for the tuning of the local reconstruction.

3.2 ECAL prompt calibration

During the 48-hour delay between the data-taking and the prompt reconstruction, several CMS workflows run automatically on computing farms once per run or more often. They produce conditions which are automatically uploaded to the Condition Database and used for the prompt reconstruction of the data of the same run, after a few hours. Checks are also run and performance plots produced in an automatic way.

ECAL laser calibration triggers are issued at a 100 Hz rate using the empty part of the LHC orbit. The laser monitoring system allows the observation of each crystal's transparency every 40 minutes. Transparency corrections are computed on a dedicated computer farm. The processing starts at the end of each run and takes a few hours for a typical 12 hour fill. Monitoring plots are produced and checked by the ECAL operations team. Fig. 1 shows the time evolution of the ECAL energy scale measured from the invariant mass distribution of $\pi^0 \rightarrow \gamma\gamma$ decays in the barrel. The π^0 candidates are selected via the high-rate stream discussed in Sec 3.1, which allows a fine monitoring of the corrections approximately every 5 minutes of data-taking. The monitoring system tracks response changes in the π^0 mass with very good precision.




Figure 1. Time evolution of the ECAL energy scale measured from the invariant mass distribution of $\pi^0 \rightarrow \gamma\gamma$ decays in the barrel. The energy scale is measured by fitting the invariant mass distribution of approximately 500k photon pairs in the mass range of the π^0 meson. Each point corresponds to approximately 5 minutes of data-taking. Data with (green dots) and without (red dots) the applications of the laser monitoring corrections are shown. The right-hand panel shows the projected relative energy scales.

During the LHC Run 2, the amplitude of signals collected in the ECAL crystals has been reconstructed by a template fitting technique [6]. The algorithm requires a precise knowledge of the average pedestal value of the digitised pulse shape and a monitoring of its time evolution and fluctuations. The pedestals are computed using events of the calibration sequence, with an automatic workflow which runs at the CMS Tier-0 [7]. One measurement per channel is computed for each run and used in the prompt reconstruction. Monitoring plots are also produced and uploaded to the Data Quality Monitoring system.

An accurate ECAL energy response is also fundamental for triggering purposes, both at L1 and at HLT, to stabilize the response of the detector and the trigger rate. For usage at trigger level, transparency corrections and pedestals are not computed with the same workflow as offline. During the 2018 data-taking, they were manually computed and updated twice/week in the database.

3.3 Offline calibration

The computation of the ECAL energy scale and crystal-to-crystal intercalibration constants requires a large amount of data and therefore these quantities are not updated within 48 hours for the prompt reconstruction. The energy scale is monitored using $Z \rightarrow ee$ decays, and updated offline as often as needed - usually two or three times per year of data-taking. The full statistics of one year of data-taking instead is used to compute new IC constants. These calibrations are therefore used for the data reprocessing which typically takes place after the end of the data-taking.

The precision of the crystal intercalibration directly affects the detector energy resolution. IC constants are computed from LHC data using independent methods and then combined. The algorithms exploit the azimuthal symmetry of the energy flow in zero bias events, the invariant mass of photon pairs from $\pi^0 \rightarrow \gamma\gamma$ decays, the energy over track momentum (E/p) ratio of isolated electrons from $W \rightarrow ev$ and $Z \rightarrow ee$ decays, and the invariant mass of electron pairs from $Z \rightarrow ee$ decays. For the first two techniques, data collected with the streams detailed in Sec. 3.1 are used. The E/p and $Z \rightarrow ee$ methods exploit a dedicated data format, storing the ECAL hits and other information only for events with at least one good reconstructed electron. Based on this data format, the reconstruction of the ECAL-related part of the event is possible, with the flexibility of using different calibration and correction constants. The reprocessing of a one-year dataset using this framework takes about 3-4 days on the CERN batch queue system. Given the complementarity of the different intercalibration techniques, an improved performance can be obtained by combining them accounting for the different precision as a function of the pseudorapidity. Details on the intercalibration algorithms and performance in Run 2 can be found in [8].

The CMS experiment performed a detector re-calibration for the full Run 2 dataset in 2019, to have a stable energy scale throughout the entire dataset and the best energy resolution. This calibration set is currently being used for a full Run 2 data reprocessing. For what concerns ECAL, the recalibration took around 2 months per year of data-taking. All the inputs to energy and time reconstruction were rederived, including those already computed for the prompt reconstruction. Figure 2 shows the energy resolution measured with electrons on data collected in 2017, for the prompt reconstruction and for the refined calibration. A large improvement is observed, in particular in the high pseudorapidity region. A relative energy resolution between 2% and 4% is achieved at all pseudorapidities.

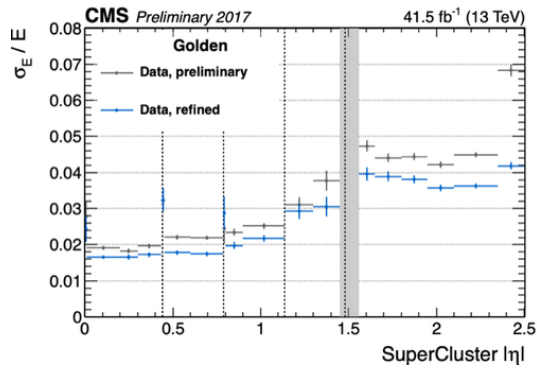


Figure 2. Relative electron (ECAL) energy resolution unfolded in bins of pseudorapidity for the ECAL barrel and the ECAL endcaps. Electrons from $Z \rightarrow ee$ decays are used. The resolution measured on 2017 data is shown for very low bremsstrahlung electrons ('golden'). The relative resolution is extracted from an unbinned likelihood fit to $Z \rightarrow ee$ events, using a Breit-Wigner function convoluted with Gaussian as the signal model.

4 Conclusions

Crystal calorimeters provide excellent energy resolution, but require continuous monitoring and corrections for ageing effects, as well as periodic channel-to-channel calibrations. Fast and reliable workflows have been developed to deliver a first set of corrections for the CMS ECAL within 48 hours of data-taking, making use of dedicated data streams and processings. During LHC Run 2, at the end of each year of running all the inputs to ECAL energy reconstruction and calibration were rederived, to achieve optimal energy resolution and performance.

References

- [1] CMS Collaboration, The CMS experiment at the CERN LHC, JINST 3 (2008) SO8004
- [2] CMS Collaboration, The electromagnetic calorimeter project, Technical Design Report, CERN/LHCC 97-033, CMS TDR 4 (1997)
- [3] M. Anfreville et al., Laser monitoring system for the CMS lead tungstate crystal calorimeter, Nucl. Instrum. Meth. A 594 (2008) 292
- [4] The CMS collaboration, CMS TriDAS project: Technical Design Report, Volume 1: The Trigger Systems, CERN/LHCC 2000-038 (2000)
- [5] The CMS collaboration, CMS high level trigger, CERN/LHCC 2007-021 (2007)
- [6] E. Di Marco, CMS electromagnetic calorimeter calibration and timing performance during LHC Run I and future prospects, <https://cds.cern.ch/record/1975982>
- [7] D. Hufnagel et al, The architecture and operation of the CMS Tier-0, J. Phys.: Conf. Ser.331 032017
- [8] F. Cavallari and C. Rovelli, Calibration and Performance of the CMS Electromagnetic Calorimeter in LHC Run 2, *these proceedings*.