PROCEEDINGS OF SCIENCE

Trigger algorithms for alignment and calibration at CMS during LHC run 3

Philipp Nattland[a,*] **for the CMS collaboration**

[a] *RWTH Aachen University, 1. Physik. Insitut B, Aachen, Germany*

E-mail: philipp.nattland@rwth-aachen.de

During LHC Run 3, the CMS experiment faces challenging pileup and high event rate conditions. To efficiently select events of interest, CMS utilizes a two-tiered triggering system. This system consists of the hardware-based level-1 trigger (L1T) and the software-based high-level trigger (HLT) that runs on a computing farm. The L1T utilizes coarse calorimeter and muon detector information, while at the HLT, a streamlined event reconstruction using the complete detector information is performed. A critical asset of the experiment is the rapid derivation of alignment and calibration constants for both the HLT and offline reconstruction. To achieve this, a dedicated set of triggers is integrated into the data-taking process. This contribution outlines the CMS alignment and calibration workflows and the interplay between these dedicated triggers and the derivation of calibrations.

PoS(EPS-HEP2025)557

[*] Speaker

ISSN 1824-8039. Published by SISSA Medialab.

1. Introduction

At the CMS experiment [1, 2], collisions of bunches of protons and heavy ions occur at rates of up to 40 MHz. To efficiently select events of interest and save the corresponding detector readouts, a two-tiered trigger system is employed [3].
The first stage, referred to as the level-1 trigger (L1T) [4], is implemented in custom hardware and uses a reduced set of information from the calorimeters and muon system to perform the triggering decision. This enables triggering within the latency budget of $4\,\mu$s and reduces the rate to about 110 kHz.
In the second stage, known as the high-level trigger (HLT) [5, 6], the event rate is reduced further to $O(\mathrm{kHz})$. This software-based trigger operates on commercial CPUs and GPUs in a computing farm at the CMS experimental site. The HLT has access to the full detector readout, selecting events according to HLT paths. These paths are sets of algorithms that reconstruct physics objects, like tracks, electrons or jets, and apply selections on them. Based on the HLT paths that activated, events are grouped into non-exclusive data streams.
In order for the HLT to work efficiently and accurately, it is crucial to provide it with the most up-to-date calibration constants [7]. Low-latency calibrations also enable the analysis of all the data within a few hours from acquisition and reduce the need for further processing.
This contribution outlines the prompt calibration workflow of CMS, with a particular focus on the low-latency derivation of the tracker alignment.

2. Data streams

Conceptually, prompt calibration uses low-latency data streams (processed within 1-2 h) to calibrate the higher-latency physics data (available after $\geq$ 48 h). In CMS, this is implemented through the following dedicated data streams. [8]:

Low-latency streams:

- **Express:** Low-rate data selection for calibrations, with specialized offline reconstruction.
- **Calibration:** Specialized data stream with reduced event content to enable high rate for detector calibrations requiring large statistics.
- **Monitoring:** Used for data-quality monitoring.

Streams with delayed reconstruction:

- **Standard:** Main data stream used in physics analyses, also referred to as "prompt physics". Offline reconstruction 48 h after data collection.
- **Parking:** Streams for physics analyses, where offline reconstruction is performed only when resources are available. So far, in Run 3 (2022–2025), it was possible to reconstruct these streams at the same time as "standard".
- **Scouting:** Reduced event content, saving mostly objects from HLT reconstruction. Enables data collection at very high rate.

PoS(EPS-HEP2025)557

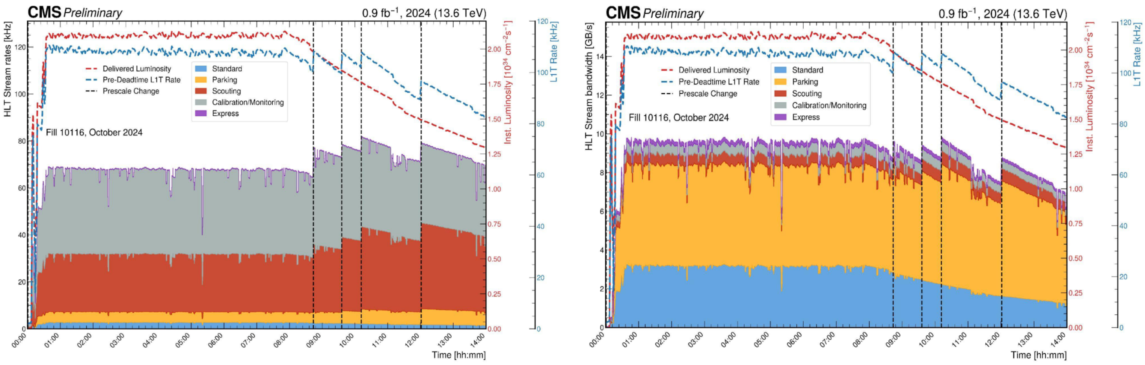

(a) **(b)**

Figure 1: HLT rate (left) and bandwidth (right), split by data streams, compared to the LHC instantaneous luminosity (red, dashed) and L1T rate (blue, dashed) as a function of time for one fill in 2024. Dashed horizontal lines correspond to changes in prescaling [8].

The HLT rate and corresponding bandwidth of these streams are shown in Fig. 1 for one proton run in October 2024. The HLT rate is dominated by the Scouting and Calibration streams, as their reduced event content enables collection at very high rates (Fig. 1a). Consequently, these streams require comparatively little bandwidth, which is primarily allocated to the Standard and Parking streams (see Fig. 1b).

3. Prompt calibration loop

The low-latency calibrations are derived within the Prompt Calibration Loop (PCL) after each data acquisition run, which corresponds to a period of stable data-taking and may extend for several hours, depending on the duration of the proton or heavy-ion fill in the LHC and the operational status of the detector. As illustrated in Fig. 2a, the Standard streams are passed to a repacker and buffered on disk, while the Express and Calibration streams are passed through the PCL.

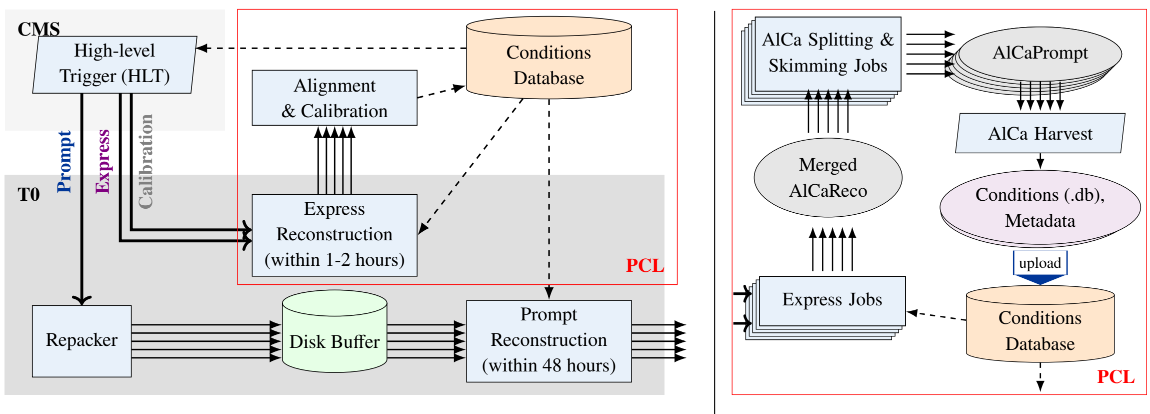

(a) **(b)**

Figure 2: Low-latency calibration workflow at CMS (left) and the Prompt Calibration Loop in detail (right). Computational steps are displayed as blue boxes, datasets as ellipses.

The first step in the PCL (see Fig. 2b) is the reconstruction of the data, parallelized across many Express jobs to stay within the 1–2 h time window. This reconstruction is specifically tailored to the Alignment and Calibration (AlCa) requirements, with the outputs saved and merged into *AlCaReco* files. Subsequently, all parallelizable computations of the AlCa algorithms are executed and stored in *AlCaPrompt* files. In the next so-called AlCa *harvest* step, all *AlCaPrompt* files are processed in a single job to derive the detector calibrations ("conditions") and save them to an *sqlite* file. Additionally, the *harvest* step outputs metadata, which is used to upload the conditions to the conditions database. These new conditions can then be used for the Prompt reconstruction of the same run and for the Express reconstruction of the following run.

4. Workflows in the PCL

The automated derivation of calibrations with the PCL has been crucial for physics performance during Run 3 of the LHC. There are many workflows implemented, including:

- beam spot calibration: The 3D position and width of the luminous region are fit, yielding up to one measurement per luminosity section, which lasts 23.31 s.
- strip tracker calibrations: The charge gains of the strip sensors are determined, problematic channels are identified, and the hit efficiency is monitored.
- pixel tracker calibrations: The bad components are determined and the Lorentz angle calibration performed.
- ECAL calibrations: The radiation damage is monitored and the ECAL pedestals are calibrated.
- alignment of the pixel tracker: The track-based derivation of the alignment constants of the inner tracking system is performed.

Additionally, the luminosity is measured through pixel cluster counting and the timing calibration of the Precision Proton Spectrometer is performed. Since a comprehensive description of all these calibrations is beyond the scope of this document, the track-based alignment is used as a representative workflow in the PCL to show the importance of low-latency calibrations for data-taking.

5. Track-based alignment of the pixel detector

The pixel tracking system is a crucial component of the CMS detector, facilitating, among others, the precise reconstruction of the primary interaction vertex (PV). To be able to fully exploit this, the position and orientation of all pixel modules must be precisely known. This is achieved at CMS by using particle tracks as inputs to a procedure called track-based alignment [9].

5.1 Lorentz drift

The charge carriers (electrons and holes) that are created by ionizing particles passing through the silicon pixel sensor, are affected by the strong magnetic field and therefore deflected in their

PoS(EPS-HEP2025)557

path to the electrode. This deflection is referred to as Lorentz drift. The drift angle (Lorentz angle) depends not only on the strength of the electric and magnetic fields, but also on the mobility of the charge carriers, which is affected by radiation damage. The Lorentz angle can be corrected for in the tracker alignment, if modules with inward- and outward-facing electric fields are aligned independently.

5.2 Alignment workflow in the PCL

Since effects such as magnet cycles and temperature variations affect the physical alignment over time and the Lorentz angle is affected by radiation damage, the alignment has to be regularly updated. To this end, multiple alignment workflows are implemented within the PCL:

- Low-granularity (LG) alignment: The six large structures of the pixel tracker — comprising two pixel half-barrels (BPIX) and four forward half-cylinders (FPIX) — are aligned. Each structure has six degrees of freedom (three rotational and three translational), resulting in a total of 36 alignment parameters for the fit. This workflow uses only data collected with minimal triggering requirements (MinBias).

- High-granularity (HG) alignment: The modules in BPIX and FPIX are mounted on ladders and panels, respectively. Instead of the large structures, in the HG alignment, these ladders and panels are aligned, resulting in about 3600 free parameters, which are fit using only MinBias data.

- High-granularity-combined (HGC) alignment: The same parameters are fit as for the HG alignment, but in addition to MinBias data, tracks from $Z \rightarrow \mu\mu$ decays are used.

The HG alignment was introduced in Run 3 to help mitigate radiation damage effects that manifest in modules with inward- and outward-facing electric fields with opposite sign, such as the Lorentz angle described earlier. A figure of merit to compare the performance of these algorithms is the distribution of median residuals (DMR). For this, each track is refitted without a hit of interest, and the distance between the measured hit and the prediction by the fit is calculated. The median of these so-called residuals is taken for each module, and the difference of the mean value between inward- and outward-facing modules can be derived.

This is shown for the modules of BPIX layer 1 in Figure 3 as a function of recorded luminosity during 2022. At the start of the year, up to the technical stop at about 11 fb^{-1}, the LG alignment was active during data taking. The effects of radiation damage are clearly visible as a trend in the data points. After the technical stop, the HG alignment was deployed in the PCL. The residual difference could be significantly improved, such that the HG PCL alignment could also directly be used in the end-of-year re-reconstruction. The increase in alignment precision comes with the drawback of a significantly larger parameter space that cannot be fully constrained by the MinBias tracks. This results in a bias in PV-related observables, as the PV reconstruction relies on the pixel tracker. To evaluate this, the PVs are redetermined without one track of interest and the distance of closest approach in z between the recalculated PV and the track, d_z, is shown as a function of the track pseudorapidity η in Fig. 4a. The mean dz shows a clear dependence on the track η, a bias introduced by the HG PCL alignment. To resolve this, the HGC alignment was introduced in the

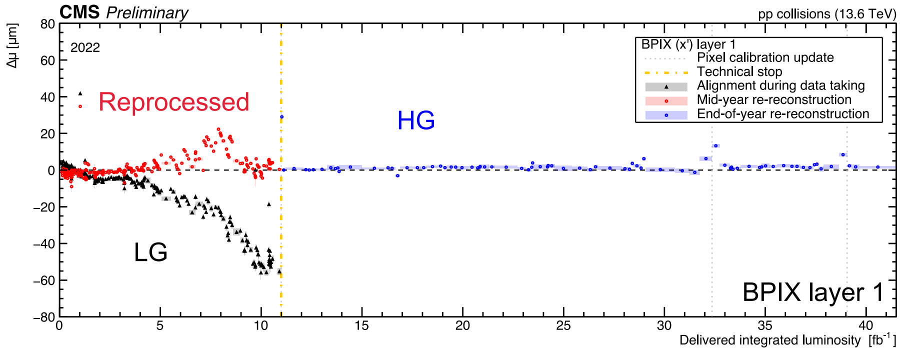

Figure 3: Difference in mean value of DMR $\Delta\mu$ between radially inward- and outward-facing modules as a function of delivered integrated luminosity during 2022. The red symbols correspond to the reprocessed alignment, the black symbols to the LG alignment, the blue symbols to the HG alignment, which was good enough to be used in the end-of-year re-reconstruction without need for any reprocessing. [10]

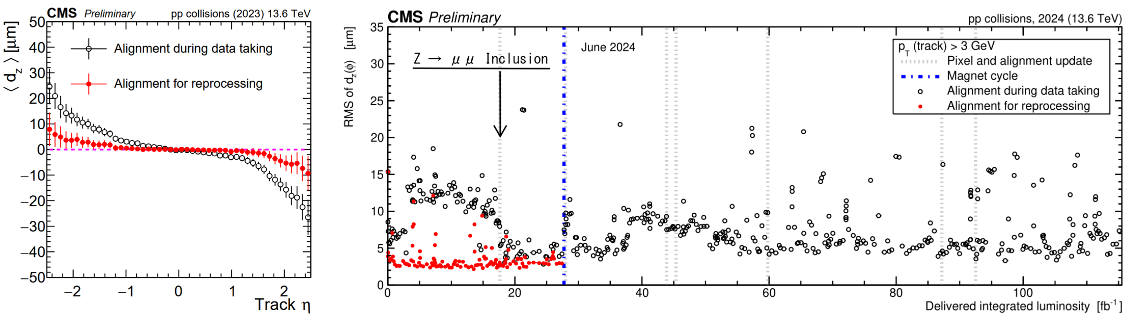

(a) **(b)**

Figure 4: Mean $d_z(\eta)$ as a function of track η for the Alignment during data taking and the Alignment for used for reprocessing on the left and the RMS of $d_z(\phi)$ as a function of delivered integrated luminosity on the right. After the introduction of the HGC alignment at around 18 fb^{-1}, the RMS is reduced significantly [11].

PCL, which uses mass constraints on the $Z \to \mu\mu$ tracks to constrain additional degrees of freedom. The root mean square (RMS) of $d_z(\phi)$, with the azimiuthal angle ϕ, is shown as a function of the luminosity in Fig. 4b. The HGC PCL alignment was introduced at around 18 fb^{-1}, where a clear drop in the RMS of $d_z(\phi)$ is visible for the alignment during data taking, showing that through inclusion of the HGC alignment in the PCL the bias can be reduced effectively.

6. Summary

Low-latency calibrations enable efficient HLT triggering and are essential for precise prompt reconstruction for physics analyses. One example of this is the automated alignment implemented in the Prompt Calibration Loop (PCL) of CMS. The implementation of the high-granularity and high-granularity-combined alignment algorithms in the PCL helped significantly reduce the effects of radiation on the Lorentz angle, while preventing bias in the PV reconstruction. This allows for use of the prompt alignment in the end-of-year re-reconstruction without the need to recalculate.

PoS(EPS-HEP2025)557

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, Development of the CMS detector for the CERN LHC Run 3, JINST 19 (2024) P05064, doi:10.1088/1748-0221/19/05/P05064

[3] CMS Collaboration, The CMS trigger system, JINST 12 (2017) P01020, doi:10.1088/1748-0221/12/01/P01020

[4] CMS Collaboration, Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13\,\mathrm{TeV}$, JINST 15 (2020) P10017, doi:10.1088/1748-0221/15/10/P10017

[5] CMS Collaboration, Performance of the CMS high-level trigger during LHC Run 2, JINST 19 (2024) P11021, doi:10.1088/1748-0221/19/11/P11021

[6] S. Varghese, CMS High Level Trigger Performance for Run 3, PoS 449 (EPS-HEP2023) 517, doi:10.22323/1.449.0517

[7] G. Cerminara *et al.*, Automated workflows for critical time-dependent calibrations at the CMS experiment, J.Phys.Conf.Ser. 664 (2015) 7 072009, doi:10.1088/1742-6596/664/7/072009

[8] CMS Collaboration, Output rates and bandwidth of the High-Level Trigger in the 2023 and 2024 proton-proton runs, CMS-DP-2025-015 (2025), cds.cern.ch/record/2931957

[9] CMS Collaboration, Strategies and performance of the CMS silicon tracker alignment during LHC Run 2, : Nucl. I. Methods A 1037 (2022) 166795, doi.org/10.1016/j.nima.2022.166795

[10] CMS Collaboration, Tracker alignment performance in 2022 (addendum), CMS-DP-2022-070 (2022), cds.cern.ch/record/2845618/

[11] CMS Collaboration, CMS Tracker Alignment Performance Results for 2024 Data Taking and Early Reprocessing, CMS-DP-2025-029 (2025), cds.cern.ch/record/2936312

PoS(EPS-HEP2025)557