PROCEEDINGS OF SCIENCE

Precision Timing with the CMS MIP Timing Detector for High-Luminosity LHC

John Dervan[a,*]

[a] *Department of Physics, Northeastern University, 100 Forsyth St., Boston, MA, United States*

E-mail: john.dervan@cern.ch

The Compact Muon Solenoid (CMS) detector at the CERN Large Hadron Collider (LHC) is undergoing an extensive Phase 2 upgrade program to prepare for the challenging conditions of the High-Luminosity LHC (HL-LHC). A new timing detector for CMS will measure minimum ionizing particles (MIPs) with a time resolution of about 30–40ps. The precise timing information from the MIP timing detector (MTD) will reduce the effects of the high levels of pileup expected at the HL-LHC, bringing new capabilities to CMS. The MTD will be composed of an endcap timing layer (ETL), instrumented with low-gain avalanche diodes and read out with the ETROC chip, and a barrel timing layer (BTL), based on LYSO:Ce crystals coupled to SiPMs and read out with the TOFHIR2 chip. An overview of the MTD design and its expanded physics capabilities, the latest progress towards prototyping and production, and the ultimate results demonstrating the achieved target timing resolution, are presented here.

The European Physical Society Conference on High Energy Physics (EPS-HEP2025)
7-11 July 2025
Marseille, France

PoS(EPS-HEP2025)519

*Speaker

1. Introduction

As of mid-2025, the Large Hadron Collider (LHC) has delivered approximately 400 fb^{-1} of proton–proton collisions to the CMS experiment—a mere 13% of the total dataset expected after the conclusion of the High-Luminosity LHC (HL-LHC) program. The HL-LHC aims to deliver more than 3 ab^{-1} of integrated luminosity and increase the instantaneous luminosity to $L_{\mathrm{inst}} \gtrsim 5 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, enabling precise measurements of Higgs boson properties, detailed studies of rare Standard Model processes, and significantly enhanced discovery potential for new physics.

Achieving this five-fold boost in the instantaneous luminosity requires substantial changes to the LHC accelerator complex, including the installation of new quadrupole magnets and crab cavities, as well as overhauls to collimation and an upgraded injector chain [2]. Alongside the accelerator upgrade, the experiments must contend with far harsher detector conditions: the average number of simultaneous interactions per bunch crossing (pileup) will rise from the 30–60 typical of Run 2 and Run 3 to around 140 in the nominal HL-LHC scenario and up to 200 in ultimate conditions.

At present, CMS manages pileup by associating charged tracks to reconstructed vertices; however, this becomes increasingly ineffective once the vertex line density at the interaction point approaches 1/mm at these pileup densities. This in turn degrades object reconstruction—particularly lepton isolation, jet energy resolution, and missing transverse energy—and reduces the overall performance of physics analyses. Radiation levels will also increase substantially in Phase II. The neutron-equivalent fluence in the innermost detector regions will reach an order of magnitude higher than in Phase I [1], necessitating new sensor technologies and readout systems capable of reliable operation over more than a decade of HL-LHC.

2. The CMS MIP timing detector

To mitigate these challenges, CMS will incorporate precision timing information for charged particles for the first time. The intrinsic spread of collision times within a bunch crossing is approximately 200 ps; by measuring per-track times with a precision of 30–40 ps in the barrel, individual interactions can be separated in time, restoring an effective pileup comparable to Phase I.

Delivering these capabilities is the minimum-ionizing particle (MIP) timing detector (MTD), an ultra-thin detector to be installed in front of the CMS calorimeters, providing nearly hermetic coverage for charged particles emitted at the interaction point. It will deliver precise time-of-flight measurements during Phase II operation. MTD will provide coverage up to a pseudorapidity of 3, with two distinct subdetectors: the Barrel Timing Layer (BTL), covering $|\eta| < 1.6$, and the Endcap Timing Layer (ETL). The design of each subdetector is shaped by different operational constraints, most notably the radiation environment, cost and power efficiency, and readout architecture. The MTD development and assembly schedule is also closely tied to the broader CMS upgrade timeline, with the two subsystems bookending the installation: the BTL must be completed ahead of the new outer tracker installation, while the ETL will be mounted on the nose of the High-Granularity Calorimeter (HGCAL) toward the end of the Phase II upgrade.

PoS(EPS-HEP2025)519

2.1 Physics motivation and performance impact

A precision timing layer located outside the CMS tracking volume is expected to deliver significant improvements not only in pileup mitigation and object reconstruction, but also in searches for physics beyond the Standard Model. By introducing time-of-flight information to the event record, the Lorentz factor β of charged particles can be directly measured (Eq. 1)

$$\frac{1}{\beta} = \frac{c(t_{\mathrm{MTD}} - t_{\mathrm{vtx}})}{L}, \tag{1}$$

where t_{MTD} and t_{vtx} denote the arrival time of the particle at MTD and the reconstructed vertex time, respectively, and L is the corresponding flight path. Measurements of β are especially valuable in heavy ion and flavor physics, where plotting $1/\beta$ as a function of particle momentum allows separation of protons, kaons, and pions. MTD will also reduce the rate of incorrect track-vertex associations by roughly a factor of two, and will improve the performance of physics objects such as missing transverse energy ($E_{\mathrm{T}}^{\mathrm{miss}}$), b-tagging, and lepton isolation [3].

From a physics perspective, time-of-flight measurements provide new handles in searches for long-lived particles (LLPs) and heavy stable charge particles (HSCPs). For LLPs, kinematic closure techniques can be used to reconstruct particle masses directly from vertex timing information; for HSCPs, deviations in $1/\beta$ distributions, arising from delayed MTD hits, offer a distinctive means of discrimination from Standard Model backgrounds. These capabilities enable a new level of directness in such searches.

3. Barrel timing layer

Precision timing measurements in the radial region of CMS between the outer tracker and the barrel calorimeter are provided by BTL, which aims to maintain 30–40 ps timing resolution, increasing to 60 ps by end of life, while sustaining a fluence up to 2×10^{14} over the course of CMS Phase II operations.

3.1 Design and sensor technology

The fundamental detection unit of BTL consists of a roughly $3 \times 4 \times 56$ mm^3 bar of cerium-doped lutetium-yttrium oxyorthosilicate (LYSO:Ce) scintillating crystal affixed with a pair of silicon photomultipliers (SiPMs). By using pairs of SiPMs, the calculated MIP time of arrival is independent of the point of impact along the bar. LYSO:Ce is chosen for its exceptional radiation hardness (<10% light output loss by end-of-life conditions), short rise (<100ps) and decay (~40ns) times, and high light yield (~40kγ/MeV). Meanwhile, the compactness, robustness, and insensitivity of SiPMs to magnetic fields made them ideal candidates for BTL photosensors. Readout of the BTL signals is performed by the dedicated Time-Of-Flight HiRate (TOFHIR) ASIC [4].

The effective operation of BTL hinges on maintaining consistent thermal conditions during runtime. Many SiPM behaviors which contribute to time resolution are highly temperature dependent; a combination of evaporative cooling to −35°C in the ambient infrastructure with dual-phase CO_2 circulated throughout each BTL tray, together with thermoelectric coolers (TECs) maintaining −45°C in contact with the SiPMs, are used to ensure a stable operating temperature.

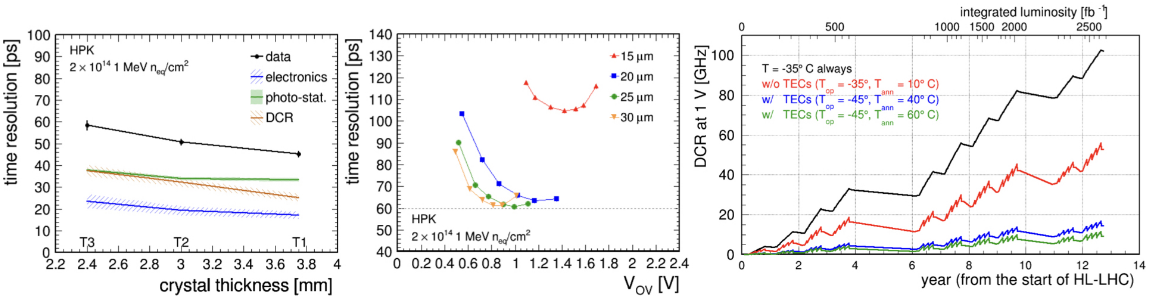

Figure 1: Performance validation of the barrel timing layer. Left: time resolution as a function of LYSO face thickness; center: time resolution as a function of overvoltage for different cell sizes in HPK modules [5]. Right: projected DCR versus integrated luminosity for different operation and annealing scenarios [6].

3.2 Performance validation

The central performance metric of the barrel timing layer is the time resolution, which must be kept as low as possible not only at the start of Phase II operations but also through the end of life of the CMS detector. This is accomplished through a combination of the material and technology selection discussed above and a careful tuning of the operating parameters of the detector through the evolution of CMS conditions. The contributions to the time resolution are summarized in Eq. 2, which represents the quadrature sum of the statistical term, electronics term, dark-count rate (DCR) term, digitization term, and clock distribution term.

$$\sigma_t^{\mathrm{BTL}} = \sigma_t^{\mathrm{stat}} \oplus \sigma_t^{\mathrm{elec.}} \oplus \sigma_t^{\mathrm{DCR}} \oplus \sigma_t^{\mathrm{digi.}} \oplus \sigma_t^{\mathrm{clock}} \tag{2}$$

A series of test beam campaigns have been conducted in recent years aimed at optimizing the statistical, electronics, and DCR contributions (Figure 1). Since the statistical term scales as $1/\sqrt{N_{\mathrm{pe}}}$, where N_{pe} is the number of photoelectrons, reducing the statistical contribution was accomplished by optimizing the LYSO:Ce thickness and packaging to maximize the light output, with a final chosen thickness of $3.1 \times 3.75\mathrm{mm}^2$ The electronics term scales as $\sigma_t^{\mathrm{elec}} \propto \sigma_{\mathrm{noise}}/N_{\mathrm{pe}}$; reducing this term involved optimizing the cell size against the SiPM power consumption, while maximizing the photon detection efficiency and overvoltage; test beam results demonstrated an optimal working point of 25μm at 1V overvoltage. Since the thermal noise (DCR) term scales as $\sqrt{\mathrm{DCR}}/N_{\mathrm{pe}}$, reduction of this contribution relies on DCR mitigation. DCR is induced by radiation damage to the SiPM silicon matrix, and mitigating this consists of reversing the TEC polarity during downtime to anneal the matrix; additionally, the TOFHIR ASIC performs a dedicated DCR cancellation algorithm. As a result of these optimizations, irradiated sensors have been shown to match the target time resolution performance set in the MTD technical design report [3, 7].

3.3 Construction status

As of July 2025, construction of BTL is well underway, with the on-detector cooling infrastructure fully assembled at the CMS Tracker Integration Facility (TIF) completed by February 2025. Upon assembly, cooling trays are shipped from CERN to the four BTL assembly centers located at the California Institute of Technology, the University of Virginia, INFN Milano-Bicocca, and Peking University. Here, the trays are populated with completed detector modules (Figure 2) and

Figure 2: Sensor side of first BTL tray constructed at CERN. Image: S. T. Rodriguez / CERN.

outfitted with the full complement of readout electronics before being shipped back to CERN for commissioning and installation. As of July 2025, the production of LYSO:Ce and SiPMs is fully complete, with around 45% of the full ten thousand sensor modules assembled. Integration of the fully-assembled trays at TIF is expected to commence in the second half of 2025, targeting the BTL commissioning and installation deadline of November 2026.

4. Endcap timing layer

The ETL will extend the precision timing capabilities of CMS to the forward region up to $|\eta| < 3$, where the expected neutron-equivalent fluence reaches $1 \times 10^{15}\ \mathrm{n_{eq}/cm^2}$ over the HL-LHC lifetime. Operating in this environment requires radiation-tolerant sensors and readout capable of maintaining per-hit timing resolutions better than 50 ps throughout Phase II. ETL is implemented as a pair of disks on each endcap, ensuring that most forward tracks receive two precise timing measurements and enabling robust pileup mitigation and time-of-flight–based reconstruction.

4.1 Design and sensor technology

The central sensor component of ETL is an array of low-gain avalanche diodes (LGADs), which are ultra-fast solid state detectors with reverse-biased *p-n* junctions. The LGAD design for ETL is a 16×16 array with $1.3 \times 1.3\mathrm{mm}^2$ active area per LGAD. These utilize a 50-micron thick active layer, targeting better than 8 fC at end-of-life conditions; this involves a trade-off between signal size and primary ionization time jitter. The chosen gain of between 10 and 30 is selected for producing an improved signal-to-noise ratio. The design specifications of ETL are suitable for the anticipated $1 \times 10^{15}\mathrm{n_{eq}/cm^2}$ fluence in the innermost (highest η) region.

Radiation hardness in such a harsh environment is a central consideration. Work was done with multiple vendors to optimize sensor uniformity, fill factor, and production yield per wafer, ultimately achieving better than 70%. The target sensor performance after irradiation was achieved by increasing the bias voltage. Sparking damage was observed for sensors irradiated with a Sr-90 β source and operated above 11.5V/µm; however, target performance for ETL is achievable below this working point (Figure 3) [8]. After optimization, the ETL sensors exhibit excellent uniformity and good efficiency, achieving the target time resolution.

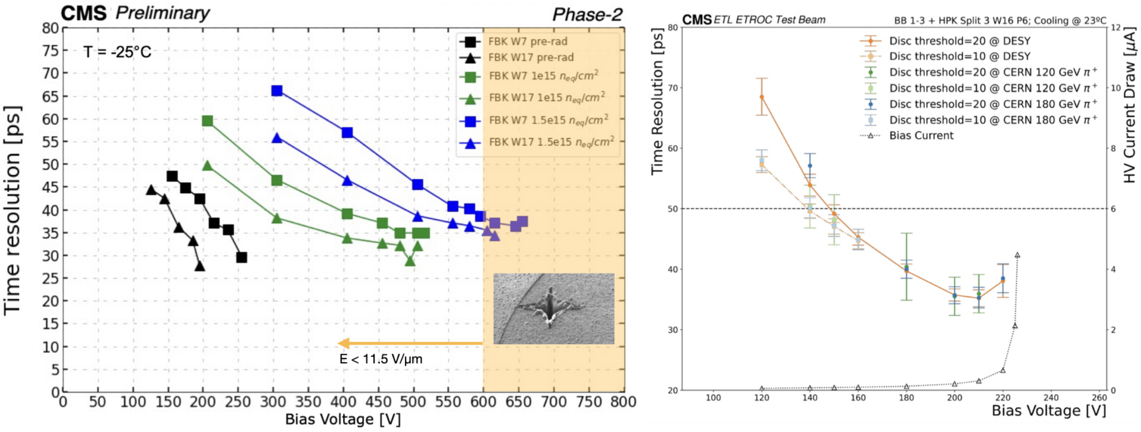

Figure 3: Performance validation plots for the endcap timing layer. Left: operational bias voltage range outside of sparking regime for various FBK sensors (modified from [8]). Right: the full demonstrated 70-volt operational range within which ETL exhibits target timing resolution [9].

4.2 Readout and performance validation

Development of the Endcap Timing Read-Out Chip (ETROC) has been a hallmark of the ETL research and development process. Operation of ETL requires a chip with low noise and power consumption, excellent radiation hardness and uniformity, and sensitivity to LGAD signals even at end-of-life conditions. The detector utilizes a two double-face disc design so that 85% of tracks are measured with two hits, giving a per-hit time resolution (from the combination of LGAD and ETROC contributions) of approximately 50ps and an overall time resolution of 35ps.

Development of ETROC was performed in stages. The ETROC0, developed in 2018, consisted of a single pixel channel, a preamplifier, and a discriminator; this was scaled in 2019 to the ETROC1, comprised of a 4×4 clock tree and time-to-digital converter. This demonstrated the first full-chain readout for ETROC. The final production design, developed in 2022, was the 16 × 16 ETROC2, attached to the LGAD array via bump-bonding. The combined time resolution of the ETROC2 and LGAD was measured in an electron test beam at DESY and pion test beam at CERN; results demonstrated a full 70V bias range where the sensor produces the target time resolution (Figure 3).

5. Conclusion and outlook

The CMS MIP timing detector will play a central role in maintaining robust event reconstruction and physics performance during the high-pileup conditions expected in HL-LHC. By providing per-track time measurements with 30–40 ps precision, MTD restores the effective pileup to Phase I levels and provides new object reconstruction and search capabilities for new physics.

Extensive research and development and test beam campaigns have validated the performance of the BTL LYSO:Ce-SiPM modules and the ETL LGAD+ETROC assemblies under irradiation levels consistent with the expected Phase II environment. Construction of BTL is in progress with cooling trays being exchanged with the four assembly centers ahead of the November 2026 installation deadline. The final 16×16 design of the ETL readout chip, meanwhile, has demonstrated the target timing resolution and shows excellent uniformity.

References

[1] CMS Collaboration, "RadSim Results for CMS Phase-2 v6.0.0.2 Radiation Simulation," CMS Internal Report DP-2023/087 (2023).

[2] G. Apollinari *et al.*, *High-Luminosity Large Hadron Collider (HL-LHC) Preliminary Design Report*, CERN-2015-005 (2015).

[3] CMS Collaboration, "A MIP Timing Detector for the CMS Phase-2 Upgrade," CERN-LHCC-2019-003, CMS-TDR-020 (2019).

[4] E. Albuquerque *et al.*, "TOFHIR2: The readout ASIC of the CMS Barrel MIP Timing Detector," CERN-Document-Server record 2753168 (2024), https://cds.cern.ch/record/2753168/files/2404.01208.pdf.

[5] F. Addesa *et al.*, "Optimization of LYSO crystals and SiPM parameters for the CMS MIP timing detector," *JINST* **19** (2024) P12020.

[6] A. Bornheim *et al.*, "Integration of thermo-electric coolers into the CMS MTD SiPM arrays for operation under high neutron fluence," *JINST* **18** (2023) P08020.

[7] CMS Collaboration, "Barrel Timing Layer Performance Plots," CMS Detector Performance Summary CMS-DP-2023-093 (2023). https://cds.cern.ch/record/2884701.

[8] CMS Collaboration, "Test Beam Characterization of 16×16 LGADs for MTD ETL," CMS Detector Performance Summary CMS-DP-2024-035 (2024).

[9] M. Safdari *et al.*, "Beam test results for the Endcap Timing ReadOut Chip (ETROC2) for CMS ETL upgrade," in *Proc. 2024 IEEE NSS/MIC/RTSD*, Tampa, FL, USA, 2024. doi: 10.1109/NSS/MIC/RTSD57108.2024.10656441. Slides: https://lss.fnal.gov/archive/2024/slides/fermilab-slides-24-0293-cms.pdf.