

17TH TOPICAL SEMINAR ON INNOVATIVE PARTICLE AND RADIATION DETECTORS
SIENA, ITALY
15–19 SEPTEMBER 2025

Track and vertex reconstruction with timing in CMS for the HL-LHC phase

P. Solanki ^a on behalf of the CMS collaboration

^aUniversity of Pisa and INFN Pisa,
Pisa, Italy

E-mail: prabhat.solanki@cern.ch

ABSTRACT: The upgrade of the CMS experiment for the High-Luminosity LHC (HL-LHC) will introduce unprecedented timing capabilities for charged particles through the MIP Timing Detector (MTD). As a key component of the Phase-II upgrade, the MTD will measure the arrival time of minimum ionizing particles with a resolution of $O(30)$ ps, covering the pseudorapidity region up to $|\eta| = 3$. This additional time information enables the reconstruction of interaction vertices in four dimensions (3 spatial + time), representing a major advancement in pileup mitigation and particle identification (PID) in the dense collision environment expected at the HL-LHC. This contribution presents recent developments in tracking and vertexing algorithms that exploit the MTD timing information, along with an assessment of the uncertainties in the time-of-flight estimation. The computational performance and the physics performance, in terms of vertex time resolution, reconstruction efficiency, pileup rejection and particle identification, will be discussed.

KEYWORDS: Particle identification methods; Particle tracking detectors; Timing detectors

Contents

1	Introduction	1
2	Vertex reconstruction with timing	2
2.1	Legacy 4D vertex reconstruction	2
2.2	3D+timing algorithm (“3Dt”)	2
2.3	Updated hybrid 4D algorithm	3
3	Performance and physics impact	3
4	Conclusion	5

1 Introduction

The high-luminosity LHC (HL-LHC) will deliver an integrated luminosity of an order of magnitude higher than that of the LHC, but this comes at the cost of a much higher pileup (PU) density. The expected average number of simultaneous proton-proton interactions per bunch crossing will increase from ~ 60 in Run 3 to 140–200 in the HL-LHC era [1]. To cope with this challenging environment, the CMS experiment is undergoing a complete Phase-2 upgrade, involving new silicon trackers, high-granularity calorimeters, and an upgraded muon system.

A key component of this upgrade is the new Minimum Ionizing Particle (MIP) Timing Detector (MTD) [1], designed to add precision time to the reconstruction of charged particle trajectories. The MTD consists of two sub-detectors: the Barrel Timing Layer (BTL) and the Endcap Timing Layer (ETL). The BTL is a cylindrical layer between the new outer tracker and the electromagnetic calorimeter barrel which utilizes LYSO:Ce crystals read out by Silicon Photomultipliers (SiPMs). With a total sensitive area of 38 m^2 and approximately 332,000 readout channels, the BTL covers the pseudorapidity region $|\eta| < 1.45$. The ETL consists of two disks located on the front face of the new High-Granularity Calorimeter end caps. It uses Low-Gain Avalanche Diodes (LGADs), has a sensitive area of 14 m^2 , and comprises approximately 8.5 million channels, extending the timing coverage to $1.6 < |\eta| < 3.0$. The MTD will provide a precise timestamp for charged particles with a resolution of 30–40 ps [1], relatively smaller than the typical RMS spread of collisions in time (~ 200 ps). This precise timing allows for the reconstruction of vertices in space-time, reducing the vertex merging rate to approximately 1% in space-time from 15% in space. Consequently, the effective pileup line density is reduced to $\sim 0.8 \text{ mm}^{-1}$ (equivalent to < 80 pileup collisions), restoring conditions where existing mitigation strategies are effective.

This contribution describes the algorithms developed for track and vertex reconstruction with timing in CMS for the HL-LHC. We first review the original “legacy” 4D vertexing approach, then introduce recent improvements that have led to a new, optimized 4D vertex reconstruction algorithm. The performance of this updated algorithm is compared to legacy 4D vertexing in terms of vertex time resolution, reconstruction efficiency, vertex separation, and pileup rejection.

2 Vertex reconstruction with timing

2.1 Legacy 4D vertex reconstruction

The original 4D vertex reconstruction algorithm, the legacy 4D algorithm, extends the existing deterministic annealing (DA) vertex finder to include the time dimension [2]. The procedure is iterative and consists of three main steps: track clustering in space-time, an adaptive vertex fit, and the computation of the vertex time.

For each track with an associated MTD measurement, a time at the point of closest approach (PCA) to the beamline, t_{PCA} , is calculated by subtracting the particle’s time-of-flight (TOF) from the MTD time. Since the particle’s mass is unknown before vertexing, the pion mass is assumed as an initial hypothesis:

$$t_{\text{PCA}} = t_{\text{MTD}} - \text{TOF}(\pi). \quad (2.1)$$

To account for this mass assumption, the time uncertainty is inflated by adding in quadrature the difference in TOF between a proton and a pion:

$$\sigma^2(t_{\text{PCA}}) = \sigma^2(t_{\text{MTD}}) + \sigma^2(\text{TOF}) + [\text{TOF}(p) - \text{TOF}(\pi)]^2. \quad (2.2)$$

For low-momentum tracks, this inflated uncertainty is dominated by the TOF difference term, effectively down-weighting the MTD’s precision in the first clustering stage.

The algorithm proceeds in two iterations. First, tracks are clustered in (z, t) space using their t_{PCA} values and inflated uncertainties. After an adaptive vertex fit, a preliminary time for each vertex (t_{vtx}) is calculated via a simple uncertainty-weighted average of its constituent track times:

$$t_{\text{vtx}} = \frac{\sum_i t_i / \sigma^2(t_i)}{\sum_i 1 / \sigma^2(t_i)}. \quad (2.3)$$

This estimation is sensitive to the initial pion mass assumption, mainly if the vertex contains heavier particles like protons.

This preliminary vertex time enables an elementary particle identification using the Time-of-Flight Particle Identification (TOF-PID) method. Its principle is to test which mass hypothesis for a given track makes its space-time origin most consistent with its parent vertex. For each track associated with a vertex, a space-time compatibility χ_h^2 is calculated for each mass hypothesis $h \in \{\pi, K, p\}$, combining both longitudinal and temporal information:

$$\chi_h^2 = \frac{(z_{\text{PCA}} - z_{\text{vtx}})^2}{\sigma^2(z_{\text{PCA}})} + \frac{(t_{\text{PCA}}(h) - t_{\text{vtx}})^2}{\sigma^2(t_{\text{PCA}}(h))}. \quad (2.4)$$

The mass hypothesis yielding the minimum χ_h^2 value is assigned to the track. In the second iteration of the legacy algorithm, the track times and uncertainties are updated based on this assigned mass (after removing the inflation term), and the 4D clustering and vertex time calculation are repeated.

2.2 3D+timing algorithm (“3Dt”)

An alternative strategy, the 3D+timing (3Dt) algorithm, was developed to incorporate timing more efficiently. This approach decouples the spatial clustering from the time assignment. First, vertices are

reconstructed using the standard 3D DA algorithm, which uses only spatial track information. Then, a time is assigned to each 3D vertex *a posteriori* using the MTD measurements of its associated tracks.

The vertex time is determined not by a simple average, but by a more robust fit that minimizes a “free energy” cost function, F , treating the vertex time t_v as a free parameter [4]:

$$F = -T \sum_{i \in \text{tracks}} w_{0,i} \ln \left[Z_0 + \sum_{h \in \{\pi, K, p\}} \alpha_h \exp\left(-\frac{[t_i(h) - t_v]^2}{2T\sigma(t_{i,h})^2}\right) \right], \quad (2.5)$$

where the sum is over tracks i in the vertex, and h denotes the particle mass hypothesis. Here, $w_{0,i}$ is the track’s weight from the 3D spatial fit, $t_i(h)$ is its time under hypothesis h , T is the DA temperature, α_h are a priori probabilities for each particle type (0.7/0.2/0.1 for $\pi/K/p$), and Z_0 is an outlier rejection term.

The vertex time t_v is found by iteratively minimizing this cost function. The procedure follows the principles of deterministic annealing: the minimization begins at a high “temperature” T , which acts as a variance parameter governing the resolution scale of the clustering, allowing the fit to broadly consider contributions from all mass hypotheses. The temperature is then gradually lowered, causing the fit to “cool” and converge robustly to the optimal value of t_v , maximizing the overall likelihood. By considering all mass hypotheses simultaneously, this method avoids the biases intrinsic in the legacy algorithm’s initial pion assumption and provides a more precise vertex time estimation.

2.3 Updated hybrid 4D algorithm

The updated 4D algorithm is a hybrid approach that combines the strengths of the previous methods. It replaces the computationally intensive first iteration of the legacy algorithm with the fast 3Dt method. The workflow is as follows:

1. Standard 3D vertexing uses DA to find spatial vertex candidates.
2. The 3Dt time-fitting algorithm (eq. (2.5)) is applied to assign an initial time to each 3D vertex.
3. Using these initial vertex times, a mass hypothesis is assigned to each track through TOF-PID.
4. A single, final 4D clustering in (z, t) is performed using the refined track times (with non-inflated uncertainties). This step allows for separating spatially overlapping vertices based on their time.
5. The final vertex times are recalculated using the robust DA time fit as in the 3Dt algorithm.

This approach provides the full physics performance of a two-pass 4D algorithm but reduces the vertex reconstruction CPU time by approximately 30% [3]. As shown below, it also improves key physics metrics compared to the legacy algorithm.

3 Performance and physics impact

The performance of the vertexing algorithms was evaluated using simulated $t\bar{t}$ events at $\sqrt{s} = 14$ TeV with an average PU of 200. Reconstructed vertices are classified as *real* if matched to a simulated primary interaction or *fake* otherwise, based on a shared-track matching procedure [3].

Figure 1 shows the vertex time resolution and pull distribution for the signal vertex. The updated 4D algorithm achieves a time resolution of 7.1 ps, a significant improvement over the 8.7 ps resolution of the legacy 4D algorithm. The updated method also removes a negative bias in the legacy algorithm, which was likely caused by the initial pion mass assumption for all tracks. The pull distribution for the updated algorithm is closer to a unit Gaussian, indicating more reliable uncertainty estimates.

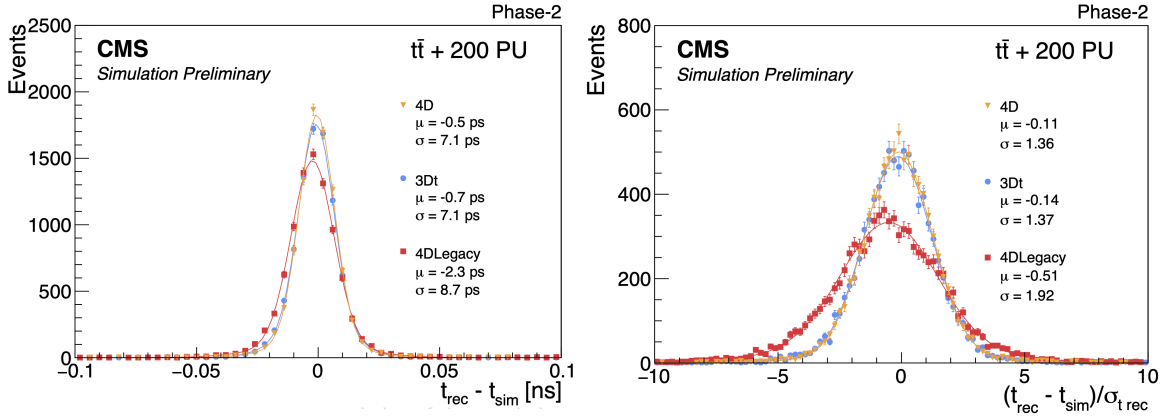


Figure 1. Vertex time resolution (left) and pull (right) for signal vertices for updated 4D (orange triangles), 3Dt (blue circles), and legacy 4D (red squares) approaches. Reproduced with permission from [3].

The inclusion of timing also impacts vertex finding efficiency and fake rate. As shown in figure 2, the updated 4D algorithm reconstructs more real vertices than the legacy algorithm across the full PU range, indicating a higher efficiency. While the fake rate of the updated 4D algorithm is slightly higher than that of the legacy algorithm, it remains lower than that of the 3Dt algorithm.

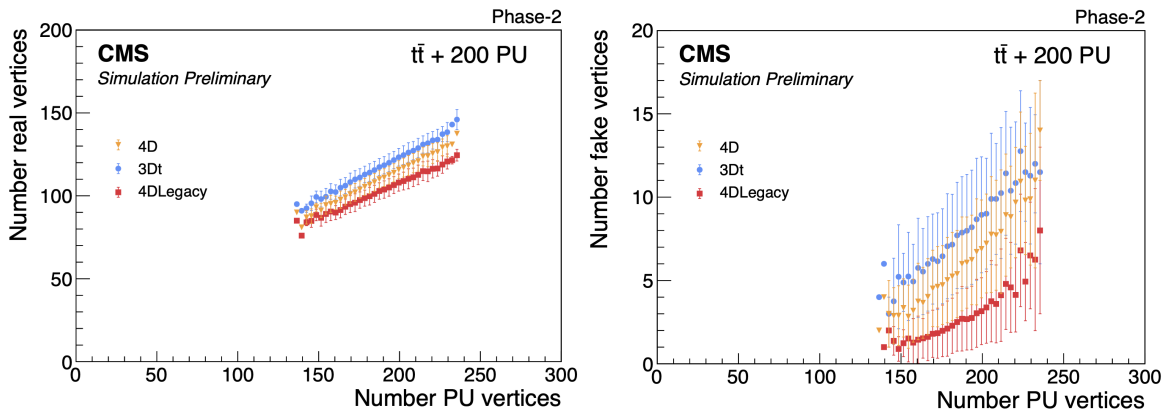


Figure 2. Number of reconstructed real vertices (left) and fake vertices (right) as a function of the number of true PU interactions for updated 4D (orange triangles), 3Dt (blue circles), and legacy 4D (red squares) algorithms. Reproduced with permission from [3].

A critical advantage of 4D reconstruction is its ability to resolve distinct collisions nearly coincident in space. Figure 3 shows the distribution of the longitudinal separation, $|\Delta z|$, for pairs of real reconstructed vertices. Standard 3D algorithms, and by extension the 3Dt method, cannot resolve vertices with minimal spatial separation and typically merge them; this is visible as the sharp drop in the 3Dt distribution for $|\Delta z| \lesssim 0.3$ mm. In contrast, the updated 4D algorithm, which uses timing in its

clustering step, reconstructs a significantly larger number of real vertex pairs at tiny separations, even approaching $|\Delta z| = 0$. This is a significant improvement over the legacy 4D algorithm and provides a direct demonstration of timing’s ability to separate vertices that are spatially inseparable.

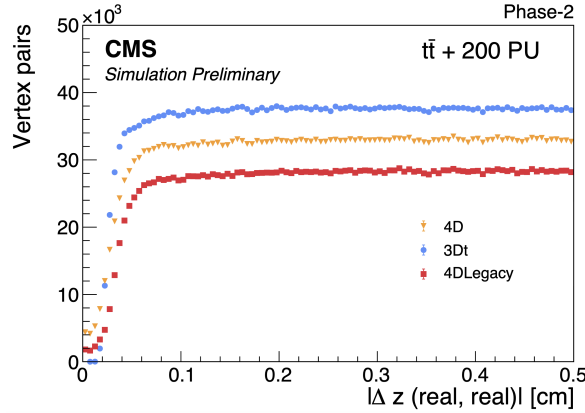


Figure 3. Distribution of the longitudinal separation, $|\Delta z|$, between pairs of real reconstructed vertices for updated 4D (orange triangles), 3Dt (blue circles) and legacy 4D (red squares) algorithms. Reproduced with permission from [3].

Finally, improved vertex reconstruction leads to cleaner physics objects. Figure 4 shows the PU contamination in the leading primary vertex. The relative contribution of PU to jet-based quantities is estimated by clustering jets without the PU tracks, recomputing the observables, and subtracting them from those computed using all charged tracks. Compared to the legacy algorithm, the updated 4D algorithm reduces the fraction of PU tracks associated with the leading vertex. This leads to a corresponding reduction in the contribution of PUs to higher-level objects.

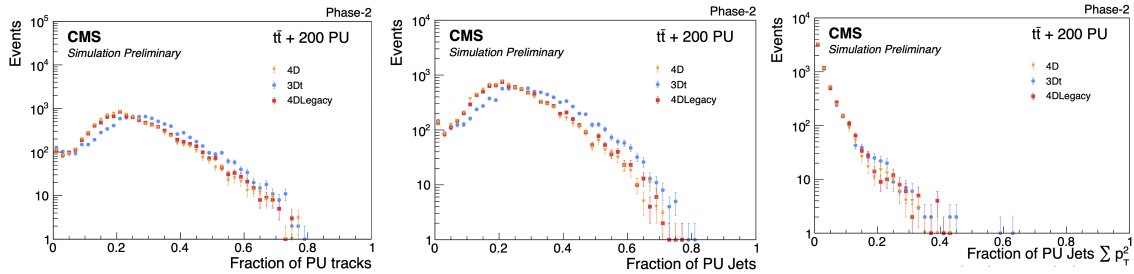


Figure 4. Relative contribution of PU to the leading vertex’s track multiplicity (left), jet multiplicity (middle), and associated jet $\sum p_T^2$ (right) for updated 4D (orange triangles), 3Dt (blue circles), and legacy 4D (red squares) algorithms. Reproduced with permission from [3].

4 Conclusion

In conclusion, the updated 4D vertex reconstruction algorithm successfully leverages the MTD’s precision timing to enhance vertexing performance in the challenging HL-LHC environment. It improves vertex time resolution from 8.7 ps to 7.1 ps and demonstrates the ability to resolve spatially overlapping vertices, recovering separate interactions even at $|\Delta z| \approx 0$. Furthermore, the updated algorithm reduces the contribution of pileup tracks to the primary vertex, while also reducing computational time by approximately 30% compared to the legacy approach.

Acknowledgments

The work presented in this paper has been funded by the European Union - Next Generation EU, Mission 4 Component 2, CUP I53D23001520006.

References

- [1] CMS collaboration, *A MIP Timing Detector for the CMS Phase-2 Upgrade*, [CERN-LHCC-2019-003](#), [CMS-TDR-020](#) (2019).
- [2] CMS collaboration, *Description and Performance of Track and Primary-Vertex Reconstruction with the CMS Tracker*, *2014 JINST* **9** P10009 [[arXiv:1405.6569](#)].
- [3] CMS collaboration, *Update of the vertex reconstruction using track time from MTD*, [CMS-DP-2024-085](#) (2024).
- [4] K. Rose, *Deterministic annealing for clustering, compression, classification, regression, and related optimization problems*, *IEEE Proc.* **86** (1998) 2210.