
CMS Physics Analysis Summary

Contact: cms-pag-conveners-exotica@cern.ch

2025/07/07

Long-lived particle triggers at CMS: Strategy and performance in proton-proton collisions at $\sqrt{s} = 13.6$ TeV

The CMS Collaboration

Abstract

The CMS Run 3 (2022–2026) physics program expands the scope of the search for long-lived particles at the CERN LHC with the addition of dedicated triggers and the improvement of the existing triggers. These customized triggers are described in this note, along with their performance using several models of new physics and the proton-proton collision data collected by the CMS detector during 2022–2024 at a center-of-mass energy of 13.6 TeV.

Contents

1	Introduction	2
2	The CMS detector	3
3	The CMS trigger system	5
4	Event reconstruction	7
4.1	The particle-flow algorithm	7
4.2	Tracking and vertexing	7
4.3	Muons	9
4.4	Electrons and photons	10
4.5	Jets	11
4.6	Tau leptons	11
4.7	Hadrons, energy sums, and missing transverse momentum	12
5	Data and Monte Carlo simulation	13
6	Long-lived particle trigger algorithms and efficiencies	16
6.1	Tracker-based algorithms	19
6.1.1	p_T^{miss} -based disappearing-track triggers	19
6.1.2	Displaced-tau triggers	23
6.1.3	Displaced-jet triggers using the tracker	25
6.2	Calorimeter-based algorithms	28
6.2.1	Displaced-jet triggers using the HCAL	28
6.2.2	Delayed-jet triggers using ECAL timing	35
6.2.3	Delayed-diphoton triggers	38
6.2.4	Displaced photon+ H_T triggers	41
6.3	Muon spectrometer-based algorithms	42
6.3.1	Displaced single and dimuon triggers	42
6.3.2	Double displaced L3 muon triggers	50
6.3.3	Displaced L3 muon+photon triggers	51
6.3.4	Dimuon scouting triggers	52
6.3.5	Muon detector showers triggers with the CSCs	55
6.3.6	Muon detector showers triggers with the DTs	58
6.4	No-BPTX algorithms	59
6.4.1	Jet No-BPTX triggers	60
6.4.2	Muon No-BPTX triggers	60
7	Long-lived particle trigger acceptances	62
8	Summary	66
A	Glossary of acronyms	78

1 Introduction

While the standard model (SM) is a highly precise and predictive theory, it does not include observed phenomena such as gravity and dark matter. The SM also suffers from several theoretical shortcomings, such as the hierarchy and the baryogenesis problems. Thus, beyond-the-SM (BSM) phenomena must exist, although they have eluded us so far at the LHC [1–3].

Particles with macroscopically long lifetimes are an important possibility in the search for new phenomena and often appear in BSM theories, notably in models that describe the elementary particle nature of dark matter. Several physical mechanisms give rise to relatively long particle lifetimes in the SM, including approximate symmetries, small couplings, and hierarchies of scale [4]; such mechanisms are also present across a broad range of BSM theories [5, 6]. In particular, long-lived particles (LLPs) are specifically motivated in many models that address fundamental puzzles like the hierarchy problem, baryogenesis, dark matter, and the origin of neutrino masses [7]. In the last several years, LLPs have become a major focus in the search for new physics at the LHC.

LLPs produced inside the CMS detector at the LHC may decay at a measurable distance from the primary proton-proton (pp) interaction point or may completely pass through the detector before decaying, as opposed to promptly-decaying particles, which result in products consistent with an origin at the primary pp interaction. For example, neutral LLPs could travel a significant distance through the detector before decaying into displaced leptons, photons, or jets [8–10]. Heavy, slow-moving LLPs could decay to delayed particles a measurable amount of time after the primary pp interaction [11, 12]. A charged LLP could decay mid-flight into a neutral particle and a charged particle whose momentum is too small for its corresponding track to be reconstructed, producing a “disappearing” track [13], or into one detectable and one undetectable particle, producing a “kinked” track. Finally, monopoles or heavy stable charged particles can leave highly ionizing tracks in the detector [14].

Standard triggers, as well as offline object reconstruction and background estimation methods, are usually inadequate for LLP searches because they are designed for promptly decaying particles, and custom techniques are often needed to collect and reconstruct the data. The CMS trigger system is designed to collect data quickly, and to do so, the trigger algorithms generally assume promptly produced particles because displaced object reconstruction can often have a long processing time.

Many improvements and extensions for triggering on LLPs with the CMS experiment were made for and during Run 3 (2022–2026) of the LHC. In this note, we present the CMS Run 3 LLP trigger program and the performance of these custom triggers, using pp collision data recorded at $\sqrt{s} = 13.6$ TeV and collected during 2022–2024. We focus on new LLP triggers introduced for Run 3 data taking, but also describe triggers introduced earlier that have been maintained or improved since.

We organize the note as follows: in Section 2, we describe the CMS detector, and in Section 3, we briefly review the CMS trigger system in general. Section 4 describes the standard online and offline event reconstruction. In Section 5, we describe the data and simulation used. We describe the custom LLP trigger algorithms in Section 6 and show the efficiency of each trigger in data and simulation. In Section 7, we compare the performance of the different LLP triggers, showing their complementary acceptance in a benchmark LLP signal, in different fiducial regions of the detector. We summarize in Section 8.

2 The CMS detector

The CMS apparatus [15, 16] is a multipurpose, nearly hermetic detector, designed to trigger on [17–19] and identify electrons, muons, photons, and (charged and neutral) hadrons [20–22]. Its central feature is a superconducting solenoid of 6 m internal diameter, providing a magnetic field of 3.8 T. Within the solenoid volume are a silicon pixel and strip tracker, a lead tungstate crystal electromagnetic calorimeter (ECAL), and a brass and scintillator hadron calorimeter (HCAL), each composed of a barrel and two endcap sections. Forward calorimeters extend the pseudorapidity coverage provided by the barrel and endcap detectors. Muons are reconstructed using gas-ionization detectors embedded in the steel flux-return yoke outside the solenoid. More detailed descriptions of the CMS detector, together with a definition of the coordinate system used and the relevant kinematic variables, can be found in Refs. [15, 16]. A quadrant of the CMS detector is shown in Fig. 1.

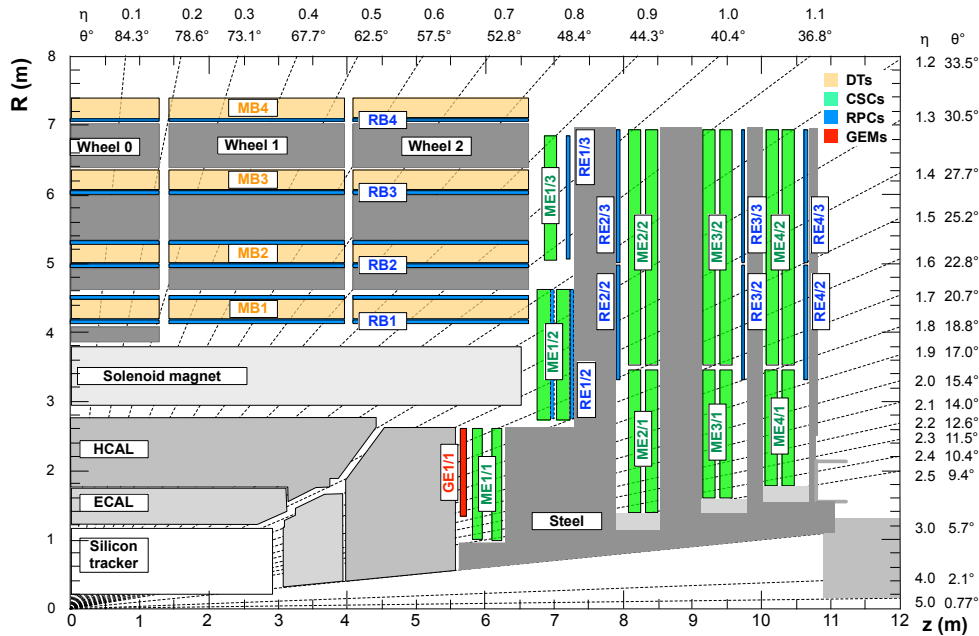


Figure 1: Schematic view in the R - z plane of a CMS detector quadrant at the start of Run 3, with the axis parallel to the beam (z) running horizontally and the radius (R) increasing upward. The interaction point is in the lower left corner. The locations of the various muon stations and the steel flux-return disks (dark areas) are shown, along with the silicon tracker, the electromagnetic calorimeter (ECAL), and the hadronic calorimeter (HCAL). The locations of the various muon stations are shown in color: drift tubes (DTs) with labels MB, cathode strip chambers (CSCs) with labels ME, resistive plate chambers (RPCs) with labels RB and RE, and gas electron multipliers (GEMs) with labels GE. The M denotes muon, B stands for barrel, and E for endcap. Figure taken from Ref. [16].

Between the beginning of the LHC operation in 2009 and the start of Run 3 data taking in 2022, the CMS detector has undergone several changes and upgrades, adapting the experiment to operating conditions at luminosities well beyond the original design [16].

The silicon tracker measures charged particles within the pseudorapidity range $|\eta| < 3.0$. During the LHC running period when the data used in this note were recorded, the silicon tracker consisted of 1856 silicon pixel and 15 148 silicon strip detector modules. The pixel detector was upgraded with the installation of a new detector in early 2017 [23]. In the new pixel detector, the number of barrel layers was increased from three to four, and the number of disks in each

endcap from two to three. The innermost barrel layer was completely replaced to ensure optimal performance until the end of Run 3. At the same time, the detector material budget was reduced, creating better tracking performance for pixel tracks with $|\eta| < 3.0$. For nonisolated particles of $1 < p_T < 10$ GeV and $|\eta| < 3.0$, the track resolutions are typically 1.5% in p_T and 20–75 μm in the transverse impact parameter d_0 [24].

The ECAL crystals are equipped with a digital clock synchronized to the LHC clock. The timing of an electromagnetic signature in the crystals can be precisely measured with these clocks [25]. The clocks on the ECAL crystals are synchronized so that the time at which an electromagnetic signature due to a particle produced at the proton collision vertex peaks in the crystal is at zero. Thus, the timing resolution of the crystals can be accurately measured from Z boson decays into dielectrons. The timing resolution of the ECAL for energies >10 GeV is <0.2 ns in the barrel.

In the ECAL, measures were taken to improve the monitoring and calibration of effects caused by radiation [16]. For example, before the start of Run 3, the trigger primitives' calibration algorithms were refined to identify and remove spurious "spike" signals created by direct energy deposits of particles in the barrel photodetectors.

In the region $|\eta| < 1.74$, the HCAL cells have widths of 0.087 in pseudorapidity and 0.087 in azimuth (ϕ). In the η - ϕ plane, and for $|\eta| < 1.48$, the HCAL cells map onto 5×5 arrays of ECAL crystals to form calorimeter towers projecting radially outwards from close to the nominal interaction point. For $|\eta| > 1.74$, the coverage of the towers increases progressively to a maximum of 0.174 in $\Delta\eta$ and $\Delta\phi$. Within each tower, the energy deposits in ECAL and HCAL cells are summed to define the calorimeter tower energies, which are subsequently used to provide the energies and directions of hadronic jets.

An upgrade of the HCAL was performed in stages that were installed between 2016 and 2019 [26]. In the HCAL barrel and endcaps, the readout bandwidth was increased to allow for more channels. As a result, the segmentation was increased to allow for both layer-dependent corrections for the observed radiation damage to the scintillating tiles [27] and to provide better rejection of energy deposits from PU interactions. The previous generation of photosensors was replaced by silicon photomultipliers, which measure the scintillator light output with a better signal-to-noise ratio. The readout electronics were upgraded to support the increased channel count, improve the precision, and add signal timing information. When combining information from the entire detector, the jet energy resolution amounts typically to 15–20% at 30 GeV, 10% at 100 GeV, and 5% at 1 TeV [28].

The CMS muon system consists of four types of gas-ionization detectors: drift tube chambers (DTs), cathode strip chambers (CSCs), resistive-plate chambers (RPCs), and triple-gas electron multiplier chambers (Triple-GEMs). The DT and CSC detectors are located in the regions of $|\eta| < 1.2$ and $0.9 < |\eta| < 2.4$, respectively, and are complemented by the RPCs in the range $|\eta| < 1.9$. The Triple-GEMs are located in the region of $1.55 < |\eta| < 2.18$. The chambers are arranged to maximize the coverage and to provide some overlap wherever possible. In both the barrel and endcap regions, the chambers are grouped into four "muon stations", separated by the steel absorber of the flux-return yoke. A detailed description of these detectors, including the gas composition and operating voltage, is reported in Ref. [21].

The GEM detector, consisting of four gas gaps separated by three GEM foils, was added in the endcaps in time for Run 3 data taking [16, 29]. The other subsystems, namely the DTs, CSCs, and RPCs, also underwent several upgrades. The top of CMS was covered with a neutron shield to reduce the background in the top external DT chambers. An outer ring of CSCs (ME4/2) was added after data taking ended in 2012 (at the end of Run 1), and in preparation

for the High-Luminosity LHC, the bulk of the CSC electronics upgrades that required chamber access were already performed during the long shutdown before Run 3 data taking. Outer rings of the RPC chambers in station four (RE4/2 and RE4/3) were also added. The endcap muon track finder of the L1 trigger was upgraded to utilize GEM-CSC joint track segments to optimize the final track reconstruction and resolution at the trigger level.

To cope with increasing instantaneous luminosities, the CMS data acquisition (DAQ) system underwent multiple upgrades [16]. A new optical readout link with a higher bandwidth of 10 Gb/s was developed. The bulk of the DAQ system downstream from the custom readout benefited from advances in technology to achieve a much more compact design, while doubling the event building bandwidth.

3 The CMS trigger system

The CMS Trigger selects events of interest using a two-tiered trigger system by reducing the input event rate of 40 MHz to several kHz. The first level (L1), composed of custom hardware processors, uses information from the calorimeters and muon detectors to select events at a rate of around 100 kHz within a fixed latency of approximately $4 \mu\text{s}$ [17]. The second level, known as the high-level trigger (HLT), consists of a farm of processors running a version of the full event reconstruction software optimized for fast processing and reduces the event rate to a few kHz before data storage [18].

After Run 1, the L1 trigger hardware was entirely upgraded and has operated successfully since 2016 [30]. For Run 3, although no major trigger hardware upgrade was performed, new capabilities have become available already through new algorithmic approaches, some of which are based on machine learning (ML) techniques. Software such as HLS4ML [31] facilitates the use of ML techniques in FPGAs. Developments for Run 3 within the L1 trigger primarily focus on expanding the physics reach of CMS through improved object measurement and calibration, as well as the introduction of dedicated triggers for LLPs and other exotic signatures. Some of these new triggers only became possible due to the enhanced capabilities of the global trigger logic and the increased trigger information provided by the calorimeters and muon systems.

The data processing of the HLT is structured around the concept of an HLT “path”, which is a set of algorithmic processing steps run in a predefined order that both reconstructs physics objects and makes selections on these objects based on the physics requirements. At the start of HLT path execution, each path requires events to pass specific L1 triggers (“L1 seeds”). The trigger objects generally used in L1 seeds are muons, $e\gamma$ objects (so-called since electrons and photons cannot be distinguished at the L1 trigger, as currently no tracking information is available there), jets, tau leptons, and event-level information like H_T and p_T^{miss} , which will be explained in Section 4. Then, each HLT path is implemented as a sequence of steps, generally of increasing complexity, reconstruction refinement, and physics sophistication. For example, the processing of intensive track reconstruction is usually performed only after some initial reconstruction and selection using data from the calorimeters and muon detectors. In case a path has requirements on two or more different kinds of physics objects, these different parts of the path are referred to as “legs” of the trigger. The reconstruction modules and selection filters of the HLT use the same software framework as used for offline reconstruction and analyses (CMSSW [32]).

The L1T and the HLT include primary triggers for analyses, as well as triggers for calibration, efficiency measurements, control region measurements, etc. that typically have looser requirements than the primary triggers. These latter triggers are often “prescaled”, meaning that they

select only a fraction of the events that satisfy their conditions to limit the storage rate. In contrast, an “unprescaled” trigger selects all events that satisfy the conditions of the trigger algorithm.

The HLT refines the purity and further reduces the rate of events previously selected by the L1T, targeting an average rate of several kHz for standard pp collision events for offline storage and prompt reconstruction. Additional storage beyond the nominal few kHz rate is available to “park” the data, whereby the offline reconstruction is postponed until computing resources are available to process the data. A higher rate of reduced-size events can be acquired using a technique known as “data scouting”, where only high-level physics objects, such as jets or leptons, reconstructed at the HLT are stored on disk. In data scouting, no raw data from the detector are stored for later offline analysis. In total, CMS employs three main approaches to data collection: the standard data, which uses the full event reconstruction performed offline immediately after data taking; the parking data, where events are saved for delayed offline reconstruction when computing resources permit; and the scouting data, in which no offline reconstruction is performed, and only the output of the HLT-level reconstruction is stored and used for analysis. The performance of parking and scouting is described further in Ref. [33].

Since 2016, the HLT has been operated using multithreaded event processing software, minimizing memory requirements by increasing the number of processes running in parallel [16]. Since the start of Run 3, the HLT has made use of graphical processing units (GPUs) in the trigger filter farm. Substantial improvements were achieved in the physics performance and speed of the software, as well as in the computing infrastructure. Some of the major changes are: support for multithreaded processes and utilization of GPUs; direct remote data access; and usage of high-performance computing centers. Algorithms implemented to run on both central processing units (CPUs) and GPUs are automatically directed to run on a GPU if one is available; otherwise, the CPU-based version of the algorithm is executed. The HLT can offload the track reconstruction based on the pixel detector and parts of the calorimeter reconstruction to GPUs. In particular, the Patatrack project [34] developed parallelized versions of pixel track and vertex reconstruction algorithms that can run on NVIDIA GPUs, while a collaboration between CMS and OpenLab ported the electromagnetic and hadronic calorimeter local reconstruction algorithms to GPUs [35, 36]. Based on these efforts, the overall event processing time has been reduced by about 40%.

Several improvements have been made in time for Run 3 data taking to better reconstruct physics objects at the HLT. For example, the tracking in the pixel and strip trackers, which is generally discussed in Section 4.2, was significantly revised. The HLT tracking is now typically performed using a single global iteration, as opposed to the three iterations that were used in Run 2 (2015–2018), and the pixel tracks are reconstructed by the Patatrack algorithm mentioned earlier in this section. This Patatrack algorithm offers improved performance over the HLT pixel tracking used in 2018 [34]. Furthermore, the identification of b jets at the HLT is essential to collect events containing such jets that would otherwise not pass the standard lepton, jet, or p_T^{miss} triggers at their nominal thresholds. Two new neural network taggers, DEEPJET [37] and PARTICLENET [38], were deployed in 2022, with improved performance over Run 2. In addition to tracks, the DEEPJET algorithm also uses information from neutral and charged particle-flow (PF) jet constituents, while the PARTICLENET algorithm provides a multinomial jet-flavor classification for categories of PF b, c, and light quarks, gluons, and hadronically-decaying tau leptons. The global PF algorithm [39], which aims to reconstruct all individual particles in an event by combining information provided by each CMS subdetector, is explained in Section 4.1.

In Run 3, data parking still targets B physics, as it did in Runs 1 and 2, but it also includes a

rich set of other physics data [33]. At the end of 2022, the parked data recorded events with a single displaced muon, events with two low p_T muons, and events with two low p_T , central electrons. In 2023, the parking strategy was extended by, for example, improving the purity of the dielectron triggers, such that now CMS also triggers on events with two b-tagged jets, Higgs events produced via vector boson fusion, and many LLP signatures.

A special version of the PF reconstruction algorithm using pixel tracks reconstructed with Patatrack was deployed for Run 3 HLT scouting [33]. As a result, the HLT scouting data were recorded at an increased rate of 30 (22) kHz in 2022 (2023) and with an event size of about 7 kB, compared to the full raw event size of about 1 MB. In addition to the PF objects that were already stored during Run 2, HLT scouting in Run 3 was expanded to include electrons and photons for the first time, as well as improved tracks with respect to those available in Run 2, which now are Patatrack tracker tracks built solely with pixel hits.

4 Event reconstruction

In this section, we briefly describe the standard event and object reconstruction algorithms, both for the offline reconstruction as well as the online reconstruction at the L1 and HLT. The details of the online reconstruction algorithms and their performance during Run 2 can be found in Refs. [17] and [19] for the L1T and the HLT, respectively. We explain here how the offline reconstruction algorithms are simplified and sped up for online use; in particular, online reconstruction algorithms usually assume that particles are promptly produced, allowing online objects to be reconstructed within the available CPU budget. We point out where the standard online reconstruction algorithms are inefficient for displaced objects, motivating the dedicated triggers described in Section 6.

4.1 The particle-flow algorithm

A PF algorithm [39] aims to reconstruct and identify each individual particle in an event, with an optimized combination of information from the various elements of the CMS detector. The energy of photons is obtained from the ECAL measurement. The energy of an electron is determined from a combination of the electron momentum at the primary interaction vertex as determined by the tracker, the energy of the corresponding ECAL cluster, and the energy sum of all bremsstrahlung photons spatially compatible with originating from the electron track. The full dedicated electron and isolated photon algorithms are time-consuming; therefore, a simplified version is run online, whose results are not fed into the online particle flow algorithm. The energy of a muon is obtained from the curvature of the corresponding track. The energy of a charged hadron is determined from a combination of the momentum measured in the tracker and the matching ECAL and HCAL energy deposits, corrected for the response function of the calorimeters to hadronic showers. Finally, the energy of neutral hadrons is obtained from the corresponding corrected ECAL and HCAL energies. The primary vertex (PV) is taken to be the vertex corresponding to the hardest scattering in the event, evaluated using tracking information alone, as described in Section 9.4.1 of Ref. [40].

4.2 Tracking and vertexing

Charged particle tracks in the tracker are reconstructed from hits in the pixel and strip tracker using a Kalman filtering technique [41]. The collection of reconstructed tracks is produced by multiple passes (iterations) of the track reconstruction sequence in a process called “iterative tracking” [22]. The basic idea of iterative tracking is that the initial iterations search for tracks that are easiest to find (e.g., of relatively large p_T , and produced near the interaction

region). After each iteration, the hits associated with the tracks are removed, thereby reducing the combinatorial complexity and simplifying subsequent iterations in a search for more difficult classes of tracks (e.g., low- p_T or greatly displaced tracks).

The offline tracking contains the following iterations. Iteration 0, the source of most reconstructed tracks, is designed for prompt tracks originating near the pp interaction point with $p_T > 0.8 \text{ GeV}$ that have three pixel hits. Iteration 1 is used to recover prompt tracks that have only two pixel hits. Iteration 2 is configured to find low- p_T prompt tracks. Iterations 3–5 are intended to find displaced tracks that originate outside the luminous region of the pp collisions and to recover tracks not found in the previous iterations. Figure 2 (left) shows the tracking efficiency for different iterations as a function of the track production vertex radius, for the standard offline Run 3 tracking with $t\bar{t}$ simulation.

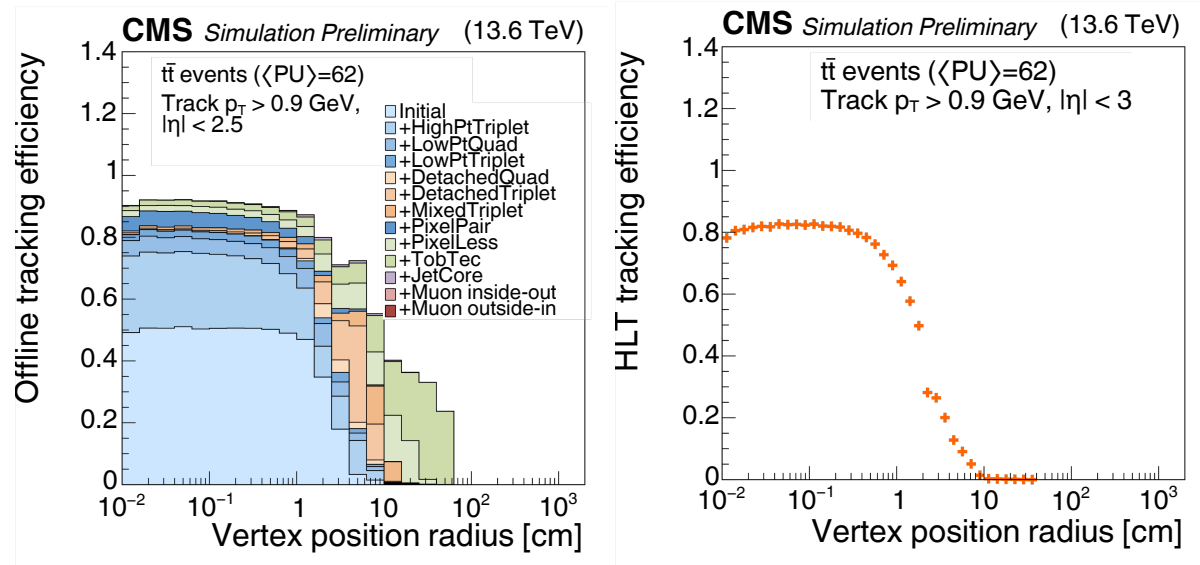


Figure 2: Standard tracking efficiency offline during Run 3 for different tracking iterations, as a function of simulated radial position of the track production vertex (left). Overall standard tracking efficiency at the HLT during Run 3, as a function of the simulated radial position of the track production vertex (right). In both figures, $t\bar{t}$ simulation in 2025 conditions and an average PU of 62 is used, and the tracks are required to have $p_T > 0.9 \text{ GeV}$ and $|\eta| < 2.5$.

Each iteration starts with a “seed” that defines the initial estimate of the trajectory parameters and their uncertainties using only two or three hits. Then, the Kalman filter extrapolates the seed trajectory along the expected flight path of the charged particle, searching for additional hits that can be assigned to the track candidate. The next step in each iteration is performed by the track-fitting module, which provides the best possible estimate of the parameters of each trajectory using a Kalman filter and smoother. Finally, the track selection sets quality flags and discards tracks that fail criteria specific to each iteration. The main differences among the iterations lie in the configuration of the seed generation and the final track selection.

The standard HLT tracking is similar to that used in the offline reconstruction, except that in Run 3, HLT tracking is now typically performed using a single global iteration targeting prompt tracks (formerly iterations 0–2 in Run 2), as discussed in Section 3. An additional tracking iteration, seeded by pixel hit pairs, was introduced in 2024 to mitigate issues in specific regions of the pixel detector, which had significantly reduced the efficiency of the standard iteration. Figure 2 (right) shows the HLT tracking efficiency as a function of the track production vertex

radius. The offline iterations that target displaced tracks (iterations 3–5) are not used in HLT tracking except in dedicated triggers such as the displaced-jets paths discussed in Section 6.1.3 and the displaced-tau paths discussed in Section 6.1.2.

Vertices are produced from clusters of tracks, and the same deterministic annealing algorithm [42] is used in the offline and HLT reconstruction [22]. The vertex position is fitted using an adaptive vertex fitter [43]. The HLT vertices (called “pixel vertices”) are produced from a subset of the pixel tracks described in Section 3, and the algorithm was modified in Run 3 to run on GPUs, if available.

4.3 Muons

Muon tracks in the offline analysis are first reconstructed independently in the tracker and in the muon system using Kalman filtering techniques [41]. In the muon system, “standalone muons” are reconstructed by combining the information from the DT, CSC, RPC, and Triple-GEM detectors. The signals from these four types of muon detectors are first combined into track segments, which are thereafter used as the input to a Kalman filtering algorithm. In the tracker, tracks are reconstructed iteratively, with each iteration seeded by a different combination of pixel or strip detector hits, to ensure all relevant track topologies are reconstructed [22]. Dedicated iterations seeded by standalone muons maximize the efficiency of muon track reconstruction. Tracker tracks and standalone muons with compatible momentum, direction, and position in the transverse plane are combined into “global muons”, and a combined fit to all associated tracker and muon system signals is performed to determine the parameters of the global track. The global muon trajectory is built “outside-in” in that it starts with the outer tracker layers and propagates inwards towards the interaction point. The p_T resolution of global muons is 1–3%, while for standalone muons, it is about an order of magnitude larger [21, 44]. The trajectories of muons with very low p_T are bent in the magnetic field to such a degree that they do not reach all stations of the muon system. To recover such muons, for which a standalone muon might not be reconstructed, tracker tracks are extrapolated to the muon segments and a loose geometrical matching to DT and CSC segments is performed. If at least one compatible segment is found, the track is classified as a “tracker muon”. The tracker muon trajectory is built from the center of the detector towards the muon system and is called “inside-out”. The offline muon reconstruction and performance are described in more detail in Refs. [21, 45].

In the trigger system, three types of muons are defined, referred to as L1, L2, and L3 muons [46]. An L1 muon is identified by the hardware-based L1 trigger system using only information from the muon detectors. The L2 and L3 muons, on the other hand, are defined within the HLT: an L2 muon is identified using information from the muon detectors alone, analogous to the offline standalone muon reconstruction algorithm, while an L3 muon additionally matches the muon detector track to a track from the tracker, analogous to the offline tracker and global muons, with both inside-out and outside-in trajectories. Displaced muons are found with both the inside-out and outside-in L3 muons. As is the case with their offline counterpart reconstruction algorithms, L3 muons have better p_T resolution than L2 muons. Standard muon HLT paths usually use the full L3 muon reconstruction, while muon HLT paths targeting displaced signatures often stop after the L2 muon reconstruction algorithm, including some paths described in Sections 6.3.3, 6.3.1, 6.3.2, 6.3.4, and 6.4.2. On the other hand, L3 muons can still have sensitivity to displaced muons as well, especially with the L3 muon reconstruction algorithms for displaced muons that are described in Sections 6.3.3, 6.3.1, and 6.3.2.

4.4 Electrons and photons

Electron and photon reconstruction is based on clusters of ECAL energy deposits [20]. In the offline reconstruction, these clusters are formed by grouping crystals with energies typically exceeding 80 MeV in the ECAL barrel and 300 MeV in the ECAL endcaps. The crystals must also be timed within 5 ns of the collision and corrected for the time of flight of a particle. A useful concept in the reconstruction of calorimeter objects is the transverse energy E_T , which is the energy of the object in question multiplied by $\sin(\theta)$, where θ is the polar angle with respect to the center of the CMS detector. Placing requirements on this quantity is an efficient way to reduce the number of objects reconstructed to a manageable level, and this is used when only calorimeter energy is available for a reconstruction element, without any associated primary vertex. A seed cluster is then defined as the one containing most of the energy deposited in any specific region with $E_T > 1$ GeV. Clusters of ECAL deposits within a certain geometric area around the seed cluster are combined into “superclusters” to include photon conversions and bremsstrahlung losses. A dedicated tracking algorithm, based on the Gaussian sum filter (GSF) [47], is used for electrons to estimate the track parameters. This GSF tracking step is seeded by both pixel detector seeds that are compatible with the supercluster position and the trajectory of an electron, as well as all tracker tracks with $p_T > 2$ GeV that are compatible with an electron trajectory. A dedicated algorithm [48] is used to find the tracker tracks that are likely to originate from photons converting into electron-positron pairs. Then, ECAL clusters, superclusters, GSF tracks, as well as conversion tracks and associated clusters, are all imported into the PF algorithm, which resolves the inputs into electron and photon objects.

In the online reconstruction, the electron and photon candidates at the L1T are based on trigger towers. In the barrel, these trigger towers are defined by 5×5 ECAL crystals that are grouped with the HCAL tower directly behind them, while in the endcaps, the trigger towers are formed from groups of 5–25 crystals depending on their η - ϕ position [17]. The trigger tower with the largest E_T is clustered together with its adjacent E_T towers using a procedure that also trims the energy deposits to only include contiguous towers to match the electron or photon signature in the calorimeter. To form an L1 candidate, energy clusters must satisfy additional identification criteria and optionally isolation requirements.

The HLT electron and photon identification begins with a regional reconstruction of the energy deposited in the ECAL crystals around the L1 candidates [19]. Since Run 2, the signals in the ECAL crystals have been reconstructed by fitting the signal pulse with multiple template functions, to mitigate out-of-time “pileup” (PU), which are additional pp interactions within nearby bunch crossings. Superclusters are then built using the same reconstruction algorithm as used offline [49]. Not all the energy of the incident particle is captured by the superclustering algorithm because of energy threshold requirements on individual crystals or because the energy is deposited in an uninstrumented area of the detector. Therefore, in both the online and offline reconstruction, the supercluster energy is corrected to better reflect the energy of the incident particle. At the HLT, this correction is designed to be more tolerant of changing detector conditions and the presence of background, at the expense of ultimate achievable precision. The resulting superclusters can then have additional selections applied to them, such as further kinematic selections, selections based on shower shape, and selections based on isolation, depending on the HLT path. Electrons in particular can be further identified by looking for pixel hits that match the trajectory the electron took through the detector based on its supercluster position, in a process known as pixel matching. If further background rejection is required, a GSF track can be reconstructed using the trajectory of any matched pixel hits as a seed. In contrast to the offline reconstruction, GSF tracks are not seeded by generic tracker tracks at the HLT. Thus, any electron tracks at the HLT must have hits in the pixel detector and

therefore can not be significantly displaced. As a result, analyses that target displaced electron signatures often do not require a matched electron track at the HLT (Sections 6.2.4, 6.2.3, and 6.3.3). Since 2017, the pixel matching algorithm has required three pixel detector hits rather than two to maximize early background rejection, while a hit doublet is accepted only if the trajectory passes through a maximum of three active modules.

The electron momentum is estimated by combining the energy measurement in the ECAL with the momentum measurement in the tracker. The momentum resolution for electrons with $p_T \approx 45$ GeV from $Z \rightarrow ee$ decays ranges from 1.6 to 5%. It is generally better in the barrel region than in the endcaps, and also depends on the bremsstrahlung energy emitted by the electron as it traverses the material in front of the ECAL [20, 50].

4.5 Jets

Offline, hadronic jets are clustered from the energy deposits in the calorimeter towers using the infrared and collinear safe anti- k_T algorithm [51, 52] with a nominal distance parameter of 0.4. Jet momentum is determined as the vectorial sum of all particle momenta in the jet, and is found from simulation to be, on average, within 5 to 10% of the true momentum over the whole p_T spectrum and detector acceptance. The PU can contribute additional tracks and calorimetric energy deposits to the jet momentum. To mitigate this effect, charged particles identified to be originating from PU vertices are discarded, and an offset correction is applied to correct for remaining contributions. Jet energy corrections are derived from simulation to bring the measured response of jets to that of particle-level jets on average. In-situ measurements of the momentum balance in dijet, photon+jet, Z+jet, and multijet events are used to account for any residual differences in the jet energy scale between data and simulation [28]. The offline jet energy resolution amounts typically to 15–20% at 30 GeV, 10% at 100 GeV, and 5% at 1 TeV [28]. Additional selection criteria are applied to each jet to remove jets potentially dominated by anomalous contributions from various subdetector components or reconstruction failures.

Jets are also reconstructed at the HLT [19] using the anti- k_T clustering algorithm. The inputs for the HLT jet algorithm can be either calorimeter towers or reconstructed objects from the PF algorithm. Most HLT jet paths use the PF inputs (“PF jets”), whereas calorimeter jets (“calo jets”) are used as a first step to identify jet signatures and initiate the PF reconstruction. To account for detector and collision conditions, several corrections are applied to the estimated PF hadron energies, average PU energy, and jet energy scale.

Several HLT b-tagging algorithms designed for promptly produced b jets, such as the DEEPJET and PARTICLENET neural networks discussed in Section 3, can provide complementary efficiencies for displaced-jets signatures with $c\tau \lesssim 1$ mm. However, the HLT b-tagging algorithms take standard HLT tracks and secondary vertices as inputs, which are tuned for b-tagging and not efficient for LLPs with significant displacements or large masses. Furthermore, the DEEPJET and PARTICLENET neural networks are trained for the identification of the prompt b jets produced in SM processes. Therefore, their efficiency decreases rapidly for $c\tau \gtrsim 1$ mm, necessitating the development of dedicated triggers targeting displaced and delayed jets. Sections 6.1.3, 6.2.2, 6.2.1, and 6.4.1 describe triggers that target displaced and delayed jet signatures.

4.6 Tau leptons

Tau leptons decay hadronically (τ_h) to one or three charged hadrons most often, with zero, one, or two neutral pions accompanying them [4]. These final state particles may produce clusters in the calorimeters that are separated in ϕ because of the magnetic field. At the L1T, the τ_h reconstruction is based on an adaptation of the $e\gamma$ dynamic clustering, which is used

to reconstruct single clusters around local maxima or seeds that can subsequently be merged into a single τ_h candidate. Both offline and at the HLT, hadronic τ decays are reconstructed from jets using the hadrons-plus-strips (HPS) algorithm [53], which combines 1 or 3 tracks with energy deposits in the calorimeters to identify the tau lepton decay modes. Neutral pions are reconstructed as strips with dynamic size in η - ϕ from reconstructed electrons and photons, where the strip size varies as a function of the p_T of the electron or photon candidate. The HLT τ_h algorithm is composed of three steps. The first two steps, called L2 and L2.5, are the same as those used in the online τ_h reconstruction in Run 2, and they are described in detail in Ref. [53]. The L2 step uses the energy deposits in the calorimeter towers, while the L2.5 step complements energy deposits with the information from the pixel detector. The third step, L3, consists of the HPS algorithm, which receives the necessary inputs from the PF event reconstruction. The tracking reconstruction step within the PF algorithm is inefficient for displaced tau leptons because it assumes prompt tracks, and it is adapted in Run 3 for the displaced τ_h use case, as discussed in Section 6.1.2. In addition, the standard HPS algorithm requires that the τ_h constituents originate from the interaction point, making it inefficient for displaced τ_h objects. Overcoming this challenge at the HLT is also discussed in Section 6.1.2.

It is possible for light leptons and jets originating from the hadronization of quarks or gluons to be erroneously reconstructed as τ_h candidates. The deep neural network DEEPTAU [54] was developed to distinguish genuine τ_h from misidentified τ_h candidates. Information from all individual reconstructed particles near the τ_h axis is combined with properties of the τ_h candidate and the event. The rate of a jet to be misidentified as τ_h by DEEPTAU depends on the p_T and quark flavor of the jet. In simulated events from W boson production in association with jets, it has been estimated to be around 3% for a genuine τ_h identification efficiency of 60%, for jets with $20 < p_T < 100$ GeV. The misidentification rate for electrons (muons) is 2.60 (0.03)% for a genuine τ_h identification efficiency of 80 (>99)%. Section 6.1.2 describes triggers that target displaced τ_h signatures, and Section 6.2.2 describes triggers that make explicit use of L1 tau seeds.

4.7 Hadrons, energy sums, and missing transverse momentum

Charged hadrons are identified as energy clusters associated with charged particle tracks that are neither identified as electrons, nor as muons. Neutral hadrons are identified as HCAL energy clusters not linked to any charged hadron trajectory, or as a combined ECAL and HCAL energy excess with respect to the expected charged hadron energy deposit.

The total hadronic transverse momentum H_T is defined as the scalar p_T sum of all jets that meet certain selection criteria. While the details of the selection may vary among different offline analyses and different HLT paths, a common definition is to use all jets with $p_T > 30$ GeV and $|\eta| < 3.0$.

The missing transverse momentum vector \vec{p}_T^{miss} is computed as the negative vector sum of the transverse momenta of all the PF candidates in an event, and its magnitude is denoted as p_T^{miss} [55]. The \vec{p}_T^{miss} is modified to account for corrections to the energy scale of the reconstructed jets in the event. The PU per particle identification algorithm [56] is applied to reduce the PU dependence of the \vec{p}_T^{miss} observable. The \vec{p}_T^{miss} is computed from the PF candidates weighted by their probability to originate from the PV [55]. At the HLT, it is crucial to account for the instrumental effects of noise and beam-induced backgrounds to keep the rates of the p_T^{miss} triggers within reasonable limits. Thus, additional filtering algorithms are applied during the HLT p_T^{miss} reconstruction to achieve lower rates for p_T^{miss} triggers [19]. Calorimeter deposits consistent with noise signature or beam halo are removed from the energy sum computation

at the HLT. Online jet energy corrections can also be propagated to the calculation of the HLT p_T^{miss} , similar to that performed offline. An alternative p_T^{miss} trigger is based on a calculation that uses all the reconstructed PF objects except for muons, leading to the “ $p_T^{\text{miss}, \mu}$ ” HLT paths. Therefore, in this approach, events with high- p_T muons are also assigned large online p_T^{miss} , whereas for events with no reconstructed muons, the two calculations coincide. The main uses for this path are searches for new physics in final states with only jets and p_T^{miss} , which require the lowest p_T^{miss} thresholds possible. A dedicated trigger to select disappearing tracks was developed (Section 6.1.1) by requiring an isolated track in addition to $p_T^{\text{miss}, \mu}$, thereby allowing for the minimum threshold on $p_T^{\text{miss}, \mu}$ to be lowered.

5 Data and Monte Carlo simulation

The trigger performance studies use data collected with the CMS detector during 2022–2024. In addition, Monte Carlo (MC) simulation of both background and LLP BSM signal processes is used. The simulation of collision events is implemented through a fixed-order perturbative quantum chromodynamics (QCD) calculation of up to four noncollinear high- p_T partons, supplemented with a description of the underlying event, parton showering, and hadronization. The former step is usually performed by a matrix-element calculator and event generator, and the MADGRAPH5_aMC@NLO [57] and POWHEG [58] packages are used for almost all the studies presented in this note; the latter step is usually implemented by the PYTHIA 8 [59] generator. Combining the two steps is a delicate procedure; a matching procedure must be implemented to avoid double-counting of processes in the combination, with the exact recipe depending on the order of the perturbative calculation. The MLM merging [60] is used for leading-order (LO) calculations, while the FxFx [61] method is used for next-to-LO (NLO). A parton distribution function (PDF) must be used to map the simulated colliding protons to the initial state partons that are present in the matrix-element calculation, and the PYTHIA parameters must be adjusted to a set of values that better describe the observed dynamics of high-energy proton collisions, which is referred as a tune. The samples used here employ the NNPDF3.1 next-to-NLO PDFs [62] and the CP5 tune [63]. The detector response to simulated particles is modeled using the GEANT4 software [64]. Minimum-bias interactions are superimposed on each event to simulate the effect of additional PU interactions within the same or neighboring bunch crossings.

To demonstrate the performance of the triggers, a set of LLP signal processes described below is used.

Figure 3 shows Feynman diagrams for an anomaly-mediated supersymmetry (SUSY) breaking (AMSB) model [65, 66] in which a chargino $\tilde{\chi}_1^\pm$ is nearly mass-degenerate with a neutralino $\tilde{\chi}_1^0$, and the decay of the chargino is of the form $\tilde{\chi}_1^\pm \rightarrow \tilde{\chi}_1^0 + X$, where ‘X’ is an SM particle that can be a charged pion π^\pm , an electron, or a muon, depending on the type of neutralino. The $\tilde{\chi}_1^\pm$ is generated for masses between 100 and 1200 GeV and $c\tau$ values between 0.1 and 10 m.

Figure 4 shows the Feynman diagram for the $H \rightarrow XX$ process. This scenario makes use of a simplified model that is motivated by the twin Higgs scenario [67–69]. Generically, we define H to be a SM-like Higgs boson with a mass of 125 GeV (that undergoes an exotic decay) or a heavier BSM Higgs boson. The X particle can be a generic LLP. In most $H \rightarrow XX$ scenarios discussed in this note, however, the X is specifically a long-lived scalar particle and then denoted S. The masses of the H and X/S, as well as the lifetime and decay mode of the X/S, are specified in each scenario described below.




Figure 3: Feynman diagrams for the AMSB $\tilde{\chi}_1^\pm \rightarrow \tilde{\chi}_1^0 + X$ process.




Figure 4: Feynman diagram for the $H \rightarrow XX$ process.

Figure 5 shows the Feynman diagram for the $H \rightarrow Z_D Z_D$ process, which arises in a hidden Abelian Higgs model (HAHM) [70]. Here, H_D is the dark Higgs boson (generated with a mass of 400 GeV), which mixes with the 125 GeV SM Higgs boson via the mixing parameter κ , taken here to be 0.01, and gives mass to the long-lived dark photon Z_D . The Z_D is generated with masses between 10 and 60 GeV, and the Z - Z_D kinetic mixing parameter ϵ is varied between 10^{-7} and 2×10^{-9} . One Z_D is required to decay into two muons, but the other's decay is not restricted. In fact, the branching ratio of $Z_D \rightarrow \mu\mu$ depends on the mass, but for the masses considered in this note, it is around 10%. Thus, about 10% of the generated $H \rightarrow Z_D Z_D$ samples provide events with four muons.

Figure 6 shows the Feynman diagram for a gauge-mediated SUSY breaking (GMSB) model where a pair of $\tilde{\tau}$ sleptons is directly produced, followed by the decay of each $\tilde{\tau}$ to a τ lepton




Figure 5: Feynman diagram for the $H \rightarrow Z_D Z_D$ process.

and a neutralino $\tilde{\chi}_1^0$ or gravitino \tilde{G} . The $\tilde{\tau}$ is generated for masses between 100 and 500 GeV and $c\tau$ values of 0.01 to 1 m.




Figure 6: Feynman diagram for the direct $\tilde{\tau}$ pair production, followed by decay of each $\tilde{\tau}$ to a τ lepton and a neutralino $\tilde{\chi}_1^0$ (or a gravitino \tilde{G}).

Figure 7 shows an example Feynman diagram for another GMSB SUSY model, commonly referred to as the “Snowmass Points and Slopes 8” (SPS8) benchmark model [71], in which pair-produced squarks and gluinos undergo cascade decays and eventually produce a gravitino \tilde{G} , which is the lightest SUSY particle, and as such is stable and weakly interacting. In this benchmark model, the lightest neutralino $\tilde{\chi}_1^0$ is the next-to-lightest SUSY particle, and its nature largely determines the phenomenology of the cascade decay chains. The mass of the $\tilde{\chi}_1^0$ is linearly related to the effective scale of SUSY breaking Λ . Depending on the value of Λ , the coupling of the $\tilde{\chi}_1^0$ to the \tilde{G} could be weak and lead to long $\tilde{\chi}_1^0$ lifetimes. A benchmark Λ value of 100 TeV is generated for $c\tau$ values of 10 and 1000 cm. The dominant decay mode of $\tilde{\chi}_1^0$ is to \tilde{G} and a photon, resulting in a final state with one or two photons and p_T^{miss} . The dominant squark-pair and gluino-pair production modes also result in additional energetic jets.

Figure 8 shows a Feynman diagram for a compressed dark sector with a small mass splitting between a dark matter candidate χ^0 and a charged dark partner χ^\pm [72]. The small mass splitting induces the dark partners to be long-lived, decaying via an off-shell W boson to a dark matter candidate. At the LHC, this benchmark model would produce soft particles (e.g., leptons, as used here) and p_T^{miss} . A benchmark χ^\pm mass of 220 GeV and mass splitting of 20 GeV is




Figure 7: Feynman diagram for the GMSB SPS8 benchmark model, where pair-produced squarks and gluinos undergo cascade decays and eventually produce a gravitino \tilde{G} .

generated for $\chi^0 c\tau$ values of 3, 30, and 300 cm.




Figure 8: Feynman diagram for the singlet-triplet Higgs portal dark matter benchmark model, where a charged dark partner particle χ^\pm in a compressed dark sector decays via an off-shell W boson and produces a stable neutral dark matter candidate χ^0 .

6 Long-lived particle trigger algorithms and efficiencies

In this section, we describe the CMS LLP trigger program, namely, the trigger algorithms specifically designed to collect events that could contain LLPs. The rates are given throughout this note for an instantaneous luminosity of $2.1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ unless specified otherwise. Sometimes we indicate the “pure” rate of a trigger, which means the exclusive rate consumed by the trigger, but most often, we display the “total” rate of a trigger, which includes events that are also selected by other, similar triggers.

The total and pure rates of LLP triggers are shown in Table 1. The rates are calculated from a single run in 2024 data (run 386593), for an instantaneous luminosity of $2.1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, with a mean PU of 63.6. These instantaneous luminosity and mean PU values were typical of the maximum output of the LHC during a pp collision fill in 2024, although the peak instantaneous luminosity achieved in this year was $2.3 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. Here, the pure rates correspond to the unique events added from the logical OR of all the dedicated LLP HLT paths, on top of

the rate from the logical OR of the rest of the CMS HLT paths. The pure rate, therefore, quantifies the rate of completely new events saved by these triggers that would otherwise not be saved by CMS. The total rate of the LLP triggers accounts for about 10, 5, and 18% of the total rate collected by CMS for standard, parking, and scouting data taking, respectively. Most of the rate saved by LLP triggers is pure, constituting around 79, 78, and 90% of the total rate of LLP triggers for standard, parking, and scouting data taking, respectively, further demonstrating the extent to which these triggers extend the CMS phase space coverage. In addition to reporting the total and pure rates for standard, parking, and scouting data taking, Table 1 also reports those for the logical OR of standard and parking data taking, representing fully reconstructed events. The modest overlap between standard and parking LLP triggers, which is concentrated in the displaced and delayed jet paths, has been addressed for recent data taking.

Table 1: The total and pure HLT rate of all dedicated LLP triggers, as calculated from 2024 data for an instantaneous luminosity of $2.1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, with a mean PU of 63.6. Rates are reported for triggers in the standard, parking, and scouting data taking separately, as well as for combined standard and parking data to show the rate of all events that are fully reconstructed. The total rates include events that may or may not have been selected by other triggers, while the pure rates correspond to events that pass dedicated LLP HLT paths and do not pass non-LLP HLT paths. The pure rates are measured separately in standard and parked data. The combined rates are slightly less than the sum of the separate standard and parking rates because some events overlap. All rates shown have a statistical uncertainty of less than 1%.

Data	Total rate [Hz]	Pure rate [Hz]
Standard	393	311
Parking	234	182
Scouting	4200	3800
Full reconstruction: standard or parking	586	389

The various LLP trigger algorithms are outlined in Table 2. In this section, we describe the algorithms and any important updates or new triggers introduced for Run 3 data taking, show their rates and efficiencies, and mention which analyses use these triggers. The trigger efficiencies shown throughout this section are usually measured by requiring selection criteria similar to what is used in the offline analysis, and the efficiencies are typically shown with respect to offline object quantities, such as jet p_T or muon d_0 . (In contrast, Section 7 shows the acceptance of the different LLP triggers within a defined fiducial region as a function of LLP $c\tau$ or decay position, to demonstrate the complementarity of the different triggers.) We organize the algorithms' descriptions into four subsections: tracker-based algorithms (Section 6.1), calorimeter-based algorithms (Section 6.2), muon spectrometer-based algorithms (Section 6.3), and No-BPTX algorithms (Section 6.4).

At the start of Run 3, a few hundred Hz of additional trigger rate at the HLT was allocated for new LLP triggers. The HLT menu was significantly expanded to explore this type of new and unconventional physics signature. Additional information can be found in Refs. [73–77].

The HLT paths for hadronically-decaying displaced tau leptons are new for Run 3. These paths, which feature a dedicated reconstruction algorithm, will be described in Section 6.1.2.

There are now several flavors of displaced and delayed jet triggers available at the HLT in Run 3. As will be described in Sections 6.1.3, some of the displaced jet triggers were already available to a certain extent in Run 2, but have gone through major improvements over time, while the delayed jet triggers, described in Section 6.2.2 and 6.2.1, are completely new for Run 3.

Table 2: The LLP triggers and their total rates at the HLT in Run 3, calculated from 2024 data for an instantaneous luminosity of $2.1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, with a mean PU of 63.6. Triggers implemented for the first time in Run 3 are indicated by a dagger (\dagger). Rate values in parentheses correspond to the parked data rates; all others are standard data rates except for dimuon scouting. Nearly all rates shown have a statistical uncertainty of less than 1 Hz. “Disp.” is used as an abbreviation for “displaced” and “req.” is used as an abbreviation for “requirement”. The terms used in this table are explained in the corresponding subsection within Section 6.

Triggered signature	Trigger description	HLT rate [Hz]
Disappearing track	$p_T^{\text{miss}} > 105 \text{ GeV} + \geq 1 \text{ isolated track } (p_T > 50 \text{ GeV})$	4
Disp. tau	$\geq 2 \text{ disp. } \tau_h (p_T > 32 \text{ GeV}, d_0 > 0.005 \text{ cm})^\dagger$ $\geq 1 \text{ disp. } \tau_h (p_T > 24 \text{ GeV}) + \geq 1 \mu (p_T > 24 \text{ GeV})^\dagger$ $\geq 1 \text{ disp. } \tau_h (p_T > 34 \text{ GeV}) + \geq 1 e (p_T > 34 \text{ GeV})^\dagger$	36
Disp. jet	$\geq 2 \text{ jet } (p_T > 40 \text{ GeV}, \text{ inclusive tagging req.}) + H_T > 430 \text{ GeV}$ $\geq 2 \text{ jet } (p_T > 40 \text{ GeV}, \text{ disp. tagging req.})$ $+ H_T > 240 \text{ GeV} + \geq 1 \text{ L1 } \mu (p_T > 6 \text{ GeV})$	53 (163)
HCAL-based disp. and delayed jet	$\geq 2 \text{ jet } (p_T > 40 \text{ GeV}, \text{ displ. tagging req.}) + H_T > 170 \text{ GeV}^\dagger$ $\geq 2 \text{ jet } (p_T > 40 \text{ GeV}, \text{ inclusive. tagging req.}) + H_T > 200 \text{ GeV}^\dagger$ $\geq 1 \text{ jet } (p_T > 60 \text{ GeV}, \text{ neutral hadron energy fraction } > 0.7) + H_T > 200 \text{ GeV}^\dagger$	35
ECAL-based delayed jet	$\geq 1 \text{ inclusive and trackless jet}^\dagger$	37 (77)
Delayed diphoton	$\geq 2 \text{ ECAL superclusters (time } > 1 \text{ ns})^\dagger$	15
Disp. photon + H_T	$\geq 1 \gamma (p_T > 60 \text{ GeV}) + \text{PF } H_T > 350 \text{ GeV}$	12
Disp. single and dimuon	$\geq 2 \text{ L2 } \mu (p_T > 10 \text{ GeV}, d_0 > 1 \text{ cm})^\dagger$ $\geq 2 \text{ L3 } \mu (p_T > 16, 10 \text{ GeV}, d_0 > 0.01 \text{ cm})^\dagger$ $\geq 2 \text{ L2 } \mu (p_T > 23 \text{ GeV})$ $\geq 1 \text{ L2 } \mu (p_T > 50 \text{ GeV}, d_0 > 1 \text{ cm})^\dagger$ $\geq 1 \text{ L3 } \mu (p_T > 30 \text{ GeV}, d_0 > 0.01 \text{ cm})^\dagger$	165
Double disp. L3 muon	$\geq 2 \text{ L3 } \mu (p_T > 43 \text{ GeV})$	2
Disp. L3 muon+photon	$\geq 1 \text{ L3 } \mu (p_T > 43 \text{ GeV}) + \gamma (p_T > 43 \text{ GeV})$ $\geq 1 \text{ L3 } \mu (p_T > 38 \text{ GeV}, d_0 > 1 \text{ cm}) + \gamma (p_T > 38 \text{ GeV})$	5
Dimuon scouting	$\geq 2 \text{ scouting } \mu (p_T > 3 \text{ GeV})$	4200
MDS in CSCs	$\geq 1 \text{ CSC cluster } (\geq 200/500 \text{ hits in outer/inner rings})^\dagger$ $\geq 2 \text{ CSC clusters } (\geq 75 \text{ hits})^\dagger$	14
MDS in CSCs + X	$\geq 1 \text{ CSC cluster } (\geq 100 \text{ hits}) + \geq 1 e (p_T > 5 \text{ GeV})^\dagger$ $\geq 1 \text{ CSC cluster } (\geq 100 \text{ hits}) + \geq 1 \text{ L3 } \mu (p_T > 5 \text{ GeV})^\dagger$ $\geq 1 \text{ CSC cluster } (\geq 100 \text{ hits}) + \geq 1 \tau_h (p_T > 10 \text{ GeV})^\dagger$ $\geq 1 \text{ CSC cluster } (\geq 50 \text{ hits}) + \geq 1 \gamma (p_T > 20 \text{ GeV})^\dagger$	14
MDS in DTs	$\text{L1 } p_T^{\text{miss}} > 150 \text{ GeV} + \geq 1 \text{ DT cluster } (\geq 50 \text{ hits})^\dagger$ $\geq 1 \text{ L1 CSC cluster} + \geq 1 \text{ DT cluster } (\geq 50 \text{ hits})^\dagger$	9
Jet No-BPTX	$\geq 1 \text{ out-of-time jet } (E > 60 \text{ GeV})$	1
Muon No-BPTX	$\geq 1 \text{ out-of-time L2 } \mu (p_T > 40 \text{ GeV})$	7

The signal efficiency for the displaced-jet triggers was improved, targeting low-mass LLPs and making particular use of the new L1 triggers that use HCAL timing and depth information. There are also new HLT paths that exploit the timing of the electromagnetic calorimeter.

Triggers for displaced muons at the HLT have also been improved in Run 3. Displaced dimuon high-level triggers take L1 muons with low p_T thresholds and new displaced double muon triggers as input. These L1 muons feed into several types of displaced dimuon triggers at the HLT, designed to cover a wide range of displacements and improve the signal efficiency over that of Run 2. These displaced dimuon triggers are further described in Section 6.3.1.

In addition, as described in Section 6.3.4, the dimuon scouting triggers were improved in Run 3 to increase their performance for displaced muons.

Neutral LLPs with particularly long lifetimes could decay beyond the calorimeters, creating a high-multiplicity shower in the muon system. Such a muon detector shower (MDS) is expected to consist of hundreds of hits, but no reconstructed tracks or jets in the direction of the cluster of hits. New MDS triggers at L1 have been developed for Run 3 to collect these high-multiplicity events in the CSCs and provide input to triggers at the HLT. Several MDS HLT paths were developed, reconstructing clusters in both the CSCs (described in Section 6.3.5) and the DTs (described in Section 6.3.6).

6.1 Tracker-based algorithms

Here, LLP triggers that primarily employ tracking algorithms are described. These include p_T^{miss} -based triggers for disappearing tracks (Section 6.1.1), displaced-tau triggers (Section 6.1.2), and displaced-jet triggers (Section 6.1.3).

6.1.1 p_T^{miss} -based disappearing-track triggers

For the disappearing-tracks analysis [13], specialized trigger algorithms were developed in Run 2. The analysis searches for a disappearing-track signature, i.e., a track made by charged particles that live sufficiently long to decay into undetectable particles inside the volume of the tracker. This track can recoil off a jet produced from initial-state radiation (ISR) and is highly isolated, which results in a small amount of deposited energy in the calorimeter associated with the track, calculated in a cone of $\Delta R = 0.5$ around it (E_{calo}), and a large number of missing hits in the outermost region of the tracking detector. Tracks are said to have a missing hit if they are reconstructed as passing through a functional tracker layer, but no hit in that layer is associated with the track. A missing hit is described as “inner” if the missing layer is between the interaction point and the track’s innermost hit, “middle” if between the track’s innermost and outermost hits, and “outer” if it is beyond the track’s outermost hit. (Note that it is possible to have more than one hit per layer.) Offline, the disappearing track must have a large number of missing outer hits. Given that there is no tracking in the L1 trigger, a disappearing-track trigger must rely on only calorimeter p_T^{miss} (calo p_T^{miss}) at the L1T, and a combination of p_T^{miss} and an isolated track at the HLT. An event display of a simulated signal event can be seen in Fig. 9.

Below is a detailed description of the main HLT path used in the analysis, with the specifications used in Run 3. The disappearing-track triggers use the logical OR of several L1 seeds to select events with relatively high p_T^{miss} , that is, $p_T^{\text{miss}} > 70 \text{ GeV}$. The algorithm starts with the calo p_T^{miss} reconstruction and requires $p_T^{\text{miss}} > 105 \text{ GeV}$ to reject a large fraction of events and increase the algorithm’s speed. For events that pass this requirement, isolated tracks are reconstructed and required to have:

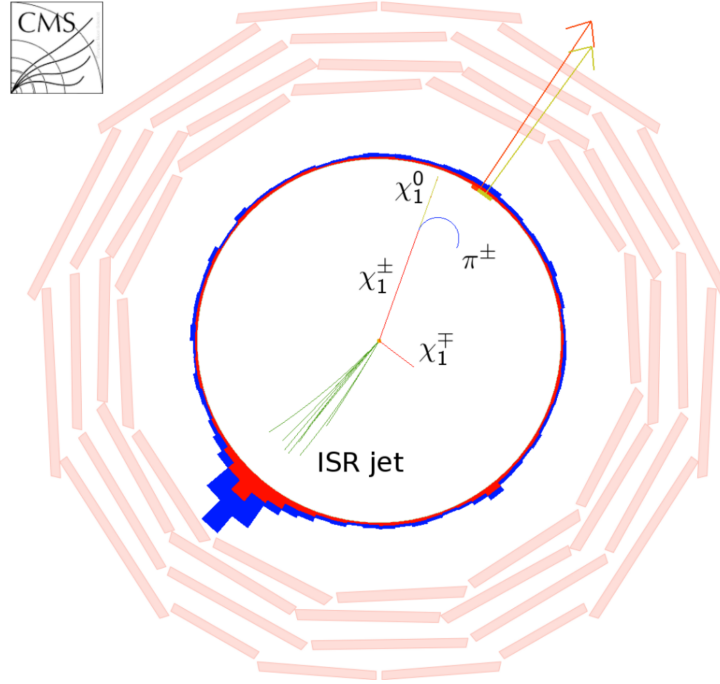


Figure 9: Event display of a simulated disappearing track. The $\tilde{\chi}_1^\pm$ (red line) is the LLP that decays in the tracker and recoils off an ISR jet (green lines), while $\tilde{\chi}_1^0$ (yellow line) is undetected by the detector, and π^\pm (blue curve) is not reconstructed because of its low p_T . The yellow and red arrows represent the calo and PF p_T^{miss} , respectively.

- $E_{\text{calo}} < 100 \text{ GeV}$;
- isolation (I) such that the scalar sum of the p_T of all other tracks within $\Delta R < 0.3$ of the candidate track must be $< 5\%$ of the candidate track's p_T ($I < 0.05$);
- number of missing inner hits = 0;
- number of missing middle hits = 0;
- number of hits > 5 ;
- $|\eta| < 5.5$;
- $p_T > 50 \text{ GeV}$.

The total rate of this HLT path is 4 Hz.

The efficiency of the HLT path is evaluated with data containing at least one muon and $W \rightarrow \ell\nu$ simulation, to select events that contain significant p_T^{miss} (coming from the neutrino) and an isolated track (coming from the muon). The efficiency is also calculated with the AMSB signal benchmark model, the Feynman diagram for which is shown in Fig. 3, as described in Section 5, for three different values of LLP proper lifetimes. Given the small mass difference between the chargino and the neutralino ($\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) \sim \mathcal{O}(m_{\pi^\pm})$), the outgoing SM particle does not have sufficient energy for it to be reconstructed as a track in the detector, and the chargino proper lifetime is large enough for it to travel through the inner layers of the tracker subdetector. In addition, the neutralino does not interact with the CMS detector material. This signature is the target of the disappearing-tracks analysis: a track of a charged particle that vanishes inside the volume of the tracker.

To calculate the efficiency, charginos with a mass of 900 GeV, which is the smallest chargino

mass not yet excluded by the disappearing tracks search, and three distinct proper lifetimes $c\tau$ are used: 0.1, 1, and 10 m, where τ is the chargino proper lifetime and c is the speed of light. The efficiency is calculated as the number of events passing a given HLT path, divided by the total number of events. Table 3 shows the offline requirements applied when calculating the efficiency in data and simulation. To study the trigger efficiency, requirements are applied to data and $W \rightarrow \ell\nu$ simulation to select a control sample of muons. On the other hand, requirements are applied to the signal simulation to obtain a sample of candidate disappearing tracks. This second set of requirements mimics the offline analysis requirement for the events to have an isolated, high p_T track. In the table, the tight muon identification [21] refers to global muons that have a given number of hits in the muon chamber included in the global-muon track fit, pixel tracker, tracker layers, and matched muon station segments, more specifically >1 , >1 , >10 and >2 , respectively. They are also required to be consistent with the PV with $d_0 < 0.2$ cm and to have $\chi^2/\text{ndof} < 10$ with respect to the global muon fit. The track high purity identification [78] is based on the selection of tracks with a deep neural network (DNN). The DNN algorithm makes use of several kinematic and geometric parameters of the track and its hits, and its identification value is assigned based on the score of the network.

Table 3: Event requirements for the disappearing-track trigger efficiency calculation, for data and $W \rightarrow \ell\nu$ simulation (left) and signal simulation (right). The selection criteria are applied sequentially over the given objects, and an event is selected if at least one object per event passes all requirements.

Data and $W \rightarrow \ell\nu$ simulation		Signal simulation	
Object	Selection	Object	Selection
≥ 1 isolated muon with $p_T > 24$ GeV at the HLT		≥ 1 tracks	$ \eta < 2.5$
≥ 1 muons	$p_T > 55$ GeV		passing high purity ID
	$ \eta < 2.1$		$d_0 < 0.02$ cm
	passing tight muon ID		$d_z < 0.5$ cm
	$I < 0.15$		hits ≥ 4
	missing inner hits =0		missing inner hits =0
	missing middle hits =0		missing middle hits =0
	match HLT muon ($\Delta R < 0.1$)		$I < 0.01$

The left plot of Fig. 10 shows the overall efficiency of the main disappearing-track HLT path in $\tilde{\chi}_1^\pm \rightarrow \tilde{\chi}_1^0 + X$ simulated events, for $c\tau = 10, 100, \text{ and } 1000$ cm. The efficiency is shown without corrections to improve the agreement between the simulation and the data. The events in this figure are required to have $p_T^{\text{miss}, \mu} > 120$ GeV to replicate the disappearing-tracks search selection. The efficiency is plotted as a function of the number of layers with valid measurements of the track that pass the offline selection. This is done to reflect the signal region bins used in the search, showcasing the different efficiencies for tracks of different sizes. The figure shows that the chargino proper lifetime has a large impact on the efficiency. In particular, for the 10 cm $c\tau$ line, the efficiency is smaller in the “4 layers” than in the “5 layers” bin, even though most of these tracks would disappear before the fifth layer of the tracker detector. This smaller efficiency comes from requiring a minimum of 5 hits in the online tracks. The offline tracks, which are used to estimate the efficiencies, are reconstructed with more robust algorithms, as is only possible offline. This step can cause some of the offline tracks to have 4 layers with measurements, even when the HLT track has 5 hits.

The right plot of Fig. 10 shows the overall efficiency of the main disappearing-track HLT path in 2022 and 2023 data and $W \rightarrow \ell\nu$ background simulation. The overall efficiency is defined as the number of events that pass the offline selection from Table 3 and the disappearing-track




Figure 10: L1T+HLT efficiency of the disappearing-track trigger: Efficiency as a function of the number of tracker layers with valid measurements of the track that pass the offline requirements, in $\tilde{\chi}_1^\pm \rightarrow \tilde{\chi}_1^0 + X$ simulated events in 2022 conditions, where $m_{\tilde{\chi}_1^\pm} = 900$ GeV and $\tilde{\chi}_1^0$ is nearly mass-degenerate with $\tilde{\chi}_1^\pm$ (left). The efficiency is shown for LLPs with $c\tau = 10, 100,$ and 1000 cm in black, blue, and red, respectively. Comparison of efficiencies calculated with 2022 data (black), 2023 data (blue), and $W \rightarrow \ell\nu$ simulation (red), as a function of offline reconstructed PF $p_T^{\text{miss}, \mu}$ (right). The efficiency follows the turn-on shape but does not reach 100% because of the isolated track leg of the algorithm.

HLT path (including the L1T, HLT p_T^{miss} , and HLT isolated-track requirements), divided by the number of events that pass the offline selection from Table 3. Because the disappearing-track triggers use calo p_T^{miss} , the efficiency is measured with respect to the offline $p_T^{\text{miss}, \mu}$.

The overall efficiency is the convolution of the per-leg efficiencies, which are shown in Fig. 11 left and right. The L1T + HLT p_T^{miss} -leg efficiency (Fig. 11 left) is defined as the number of events that pass the offline selection from Table 3 and the L1T and HLT p_T^{miss} requirements, divided by the number of events that pass the offline selection from Table 3. The isolated track leg efficiency (Fig. 11 right) is defined as the number of events that pass the offline selection from Table 3 and the disappearing-track HLT path (including the L1T, HLT p_T^{miss} , and HLT isolated-track requirements), divided by the number of events that pass the offline selection from Table 3 and the L1T and HLT p_T^{miss} requirements.

The overall efficiency curve (Fig. 10 right) does not reach 100% because of the selections made in the isolated track leg of the HLT path (as can be seen in Fig. 11 right), and because the HLT tracking efficiency does not reach 100% as well [79].

The efficiency in 2022 and 2023 data is very similar, but there is a large data-to-simulation discrepancy in Fig. 10 right and Fig. 11 right. This difference between data and background simulation comes from slightly higher efficiencies in simulation in the HLT selections made on the track isolation and the number of missing inner and middle hits. When measuring these properties offline, 9.2% more MC-simulated events than data events pass these requirements. This behavior is expected since there are no corrections applied to the MC simulation.

An important note is that the disappearing-tracks signature may also be triggered by standard




Figure 11: L1T+HLT efficiency of each leg of the disappearing-track trigger in 2022 data (black), 2023 data (blue), and $W \rightarrow \ell\nu$ simulation (red). Efficiency of the L1T + HLT p_T^{miss} leg as a function of offline reconstructed PF $p_T^{\text{miss}, \mu}$ (left). Efficiency of the full HLT path, taking into account only events that already passed through the p_T^{miss} leg, as a function of the selected muon p_T (right).

$p_T^{\text{miss}, \mu}$ triggers since those also target events containing a large amount of calo p_T^{miss} . However, the disappearing-track triggers increase the efficiency to trigger on the disappearing tracks, in part because they lower the threshold on $p_T^{\text{miss}, \text{no mu}}$ from 110 GeV in the standard $p_T^{\text{miss}, \text{no mu}}$ triggers to 105 GeV. Using standard $p_T^{\text{miss}, \mu}$ triggers in logical OR with the disappearing-track triggers to select simulated signal events provides access to additional phase space for a small increase in rate, with respect to using standard $p_T^{\text{miss}, \mu}$ triggers alone.

6.1.2 Displaced-tau triggers

Displaced tau leptons represent a well-motivated signature in many BSM theories, in particular, in some GMSB models [80–82] that include the process described in Section 5 and shown in Fig. 6. About one-third of the time, the tau lepton will decay into a displaced muon or displaced electron, and two neutrinos. Dedicated HLT reconstruction techniques are available for displaced muons, as described in Sections 6.3.1 and 6.3.2, while for displaced electrons, the options available at the HLT are to select superclusters as described in Section 6.2.3 or photons as proxies for displaced electrons, as described in Section 6.3.3. Almost all the remaining decay final states of tau leptons contain hadrons and a neutrino. The reconstruction of displaced τ_h particles is particularly challenging, both at the offline and trigger levels. In particular, there was no online τ_h reconstruction sequence in CMS addressing this topology in Run 2, and the previous search [83] had to rely on a combination of p_T^{miss} -based triggers and prompt di- τ_h triggers to select events with at least one of the final state tau leptons decaying hadronically. Figure 12 shows an event display of a simulated GMSB benchmark signal event, with two τ_h s in the final state.

A dedicated set of HLT paths has been designed and included in the menu for the Run 3 data taking since 2022, targeting the direct reconstruction of displaced τ_h using a combination of a pixelless tracking iteration on top of the standard tracking and a modified τ_h reconstruction

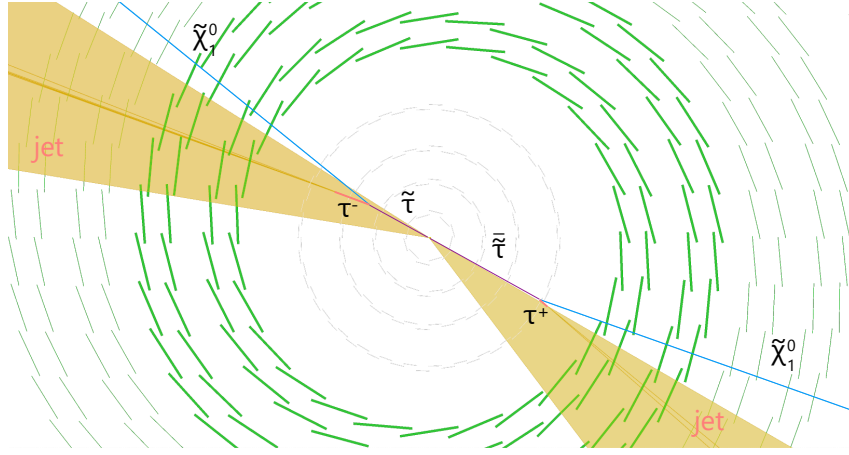


Figure 12: Event display of simulated $\tilde{\tau}$ pair production in a GMSB benchmark model, followed by the decay of each $\tilde{\tau}$ to a τ lepton and a neutralino $\tilde{\chi}_1^0$, with the τ leptons decaying hadronically. The magenta lines indicate the $\tilde{\tau}$ particles, the light blue lines indicate the neutralinos, and the dark orange lines indicate the τ leptons. The shaded dark yellow cones show the two reconstructed jets, and the orange lines inside the jets are the hadrons from the τ decay.

algorithm.

Studies performed on the standard L1 τ_h reconstruction algorithm have shown that it was not necessary to develop a dedicated displaced τ_h L1 seed. Therefore, the HLT displaced reconstruction is seeded by standard L1 τ_h objects.

The dedicated displaced reconstruction is a modification of the standard HPS algorithm, which is described in Section 4.6 and is used in the third step of the online reconstruction of promptly decaying τ_h s. To enable the HPS algorithm to reconstruct displaced τ_h s, the standard tracking algorithm has been extended to include a dedicated displaced iteration, which starts from combinations of hits in the strip tracker (pixel-less seeds) and builds tracks without assuming their provenance from the interaction point. To meet the CPU constraints, this iteration is only used in regions around isolated L2 tau leptons. Modifications to the HPS algorithm itself have also been introduced. The standard algorithm includes requirements on the τ_h constituents (charged and neutral hadrons) to originate from the interaction point, effectively preventing the reconstruction of displaced τ_h objects. These requirements have therefore been loosened. While the nominal algorithm performs poorly for τ_h decaying significantly beyond the τ proper lifetime scale, the dedicated one provides the ability to target τ_h s with decay lengths as long as 20 cm.

A suite of trigger paths using the modified HPS algorithm has been developed, targeting final states with two τ_h or with one τ_h and one displaced lepton. In the di- τ_h path, each τ_h is required to have $d_0 > 0.005$ cm and $p_T > 32$ GeV. The cross-flavor paths, which have been available online since 2024, require $p_T(\tau_h) > 24$ (34) GeV and $p_T(\ell) > 24$ (34) GeV for the muon (electron) flavored one, respectively. The total rate of these HLT paths is 36 Hz.

Figure 13 shows the trigger efficiency in simulated $pp \rightarrow \tilde{\tau}\tilde{\tau}(\tilde{\tau} \rightarrow \tau\tilde{\chi}_1^0)$ events, emulating 2022 data-taking conditions, with each τ decaying hadronically and the $\tilde{\tau}$ having a simulated $c\tau$ of 10 cm. The displaced di- τ_h efficiency and the combination of this path's efficiency with those of the previously available Run 2 p_T^{miss} -based and prompt di- τ_h path efficiencies are shown with filled markers, while the efficiencies of the Run 2 trigger options are shown with open markers. The improvement introduced with the novel trigger path can be seen in the lower panels and

is noticeable at low to medium d_0 values and low p_T^{miss} values, where the previously available triggers could not be efficient.

Figure 14 shows the rate of the trigger path in 2022 and 2023, as a function of PU. The plots indicate a linear rate response with respect to PU, although the rate increased slightly in 2023 because of the data-taking conditions.

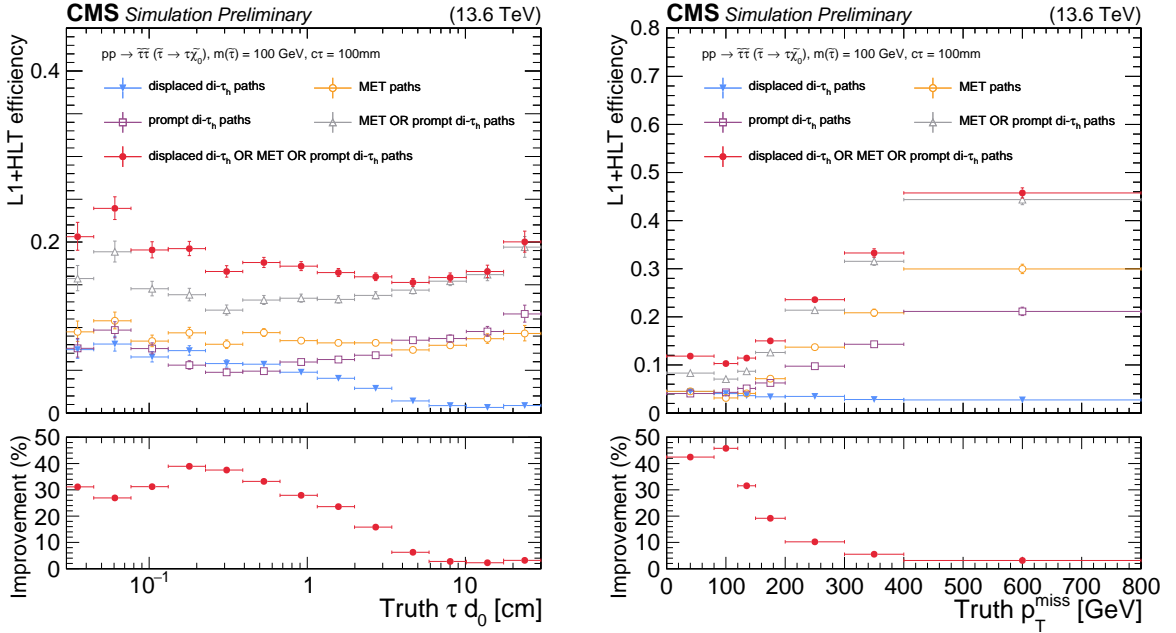


Figure 13: L1T+HLT efficiency of the displaced τ_h trigger, for simulated $pp \rightarrow \tilde{\tau}\tilde{\tau} (\tilde{\tau} \rightarrow \tau\tilde{\chi}_1^0)$ events, where each τ decays hadronically and the $\tilde{\tau}$ has a simulated $c\tau$ of 10 cm. The efficiency is shown for the displaced di- τ_h trigger path (blue filled triangles), the previously available p_T^{miss} -based paths (orange open circles), the previously available prompt di- τ_h paths (purple open squares), the combination of the p_T^{miss} -based and prompt di- τ_h paths (gray open triangles), and the combination of the p_T^{miss} -based, prompt di- τ_h , and displaced di- τ_h paths (red filled circles), using 2022 data-taking conditions. The efficiency is evaluated with respect to truth-level quantities. Efficiency of the highest- p_T τ of the event as a function of the d_0 (left). Efficiency as a function of p_T^{miss} (right). A selection on the visible component of the truth-level tau lepton $p_T > 30$ GeV and its pseudorapidity $|\eta| < 2.1$ is applied. The lower panels show the ratio (gain in %) of the trigger efficiency given by the combination of the displaced di- τ_h trigger path with the p_T^{miss} -based and prompt di- τ_h paths, divided by that of the combination of the previously available p_T^{miss} -based and prompt di- τ_h paths.

6.1.3 Displaced-jet triggers using the tracker

A suite of displaced-jet triggers was already available in Run 2 [84], utilizing the tracking information at the HLT system. Displaced jets are identified, or “tagged”, in the HLT as jets with a small number of prompt tracks or with the presence of displaced tracks. With these triggers, the Run 2 [84] and Run 3 [85] displaced-jets searches set stringent limits on a large number of BSM models with hadronically decaying LLPs and with LLP masses as low as $\mathcal{O}(100$ GeV), including split SUSY [86–92], GMSB [80–82], R -parity violating (RPV) SUSY [93–99], covering a variety of final-state topologies. Such a displaced-jets signature is well-motivated in many BSM scenarios, especially the ones involving Higgs-portal hidden sectors. Figure 15 shows a display of an $H \rightarrow SS \rightarrow qq\bar{q}\bar{q}$ event with $m_S = 55$ GeV, where one long-lived S decays inside

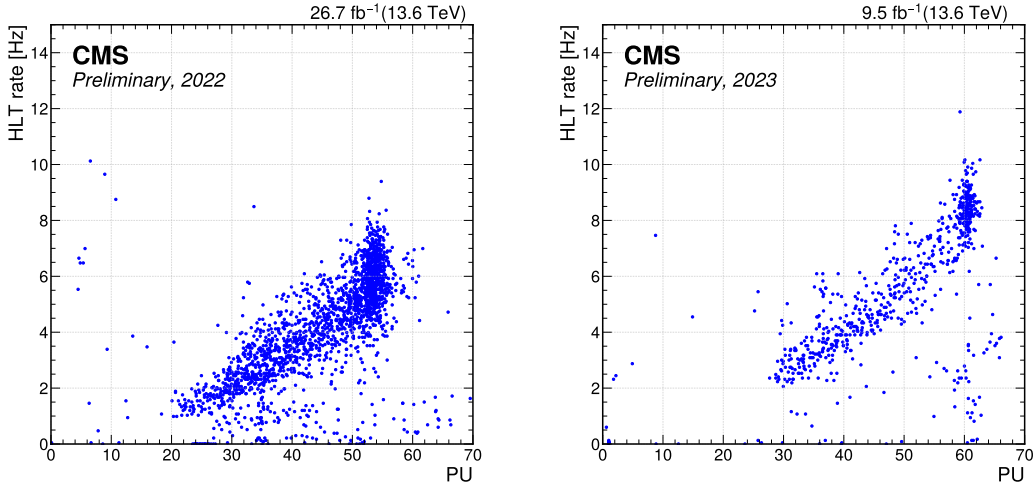


Figure 14: Total rate of the displaced τ_h trigger in 2022 (left) and 2023 (right) data, as a function of PU.

the tracker, producing a pair of displaced jets accompanied by displaced tracks and a displaced vertex. The Feynman diagram for this process is shown in Fig. 4.

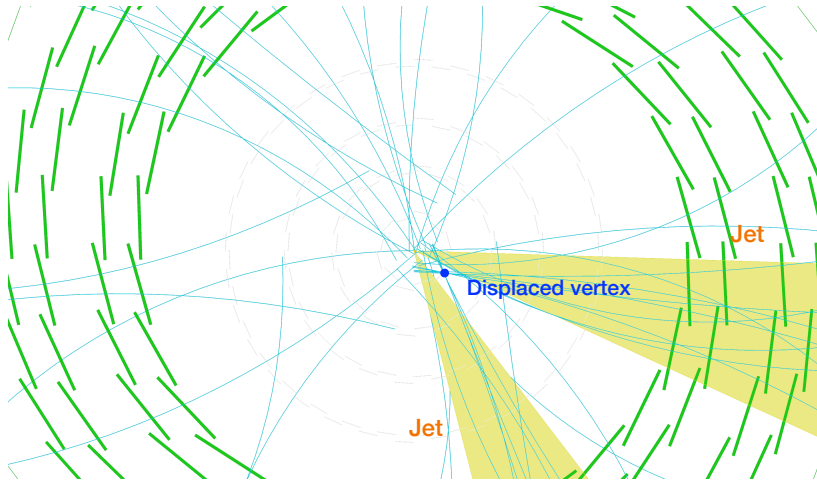


Figure 15: Event display of a pair of displaced jets arising from an LLP decay, producing displaced vertices and tracks. The simulated process is $H \rightarrow SS \rightarrow q\bar{q}q\bar{q}$. The blue curves indicate the reconstructed tracks. The yellow cones indicate two reconstructed jets, with a number of associated displaced tracks. The blue point indicates the displaced vertex induced by the LLP decay.

In Run 3, the displaced-jets triggers have gone through major improvements to increase the trigger efficiencies for low-mass LLPs with masses smaller than 100 GeV, such as those expected in the exotic Higgs boson decays. The main improvements are in the tracking-based displaced-jets tagging requirements. Displaced jets are tagged at the HLT using two types of requirements. The first requires calo jets with at most one associated prompt track, referred to as the “inclusive” tagging requirement. Prompt tracks are tracks that have $p_T > 1$ GeV and d_0 measured with respect to the leading PV smaller than 0.5 mm. Tracks are associated with the calo jets by requiring the angular distance ΔR between a given track and a given jet to be smaller than 0.4. The second requirement, referred to as the “displaced” tagging requirement, starts with the inclusive tagging requirement and additionally requires that if there is exactly

one associated prompt track, there should also be at least one associated displaced track with $p_T > 1 \text{ GeV}$ and $d_0 > 0.3 \text{ mm}$.

Two sets of displaced-jets triggers are implemented in standard data taking: a main set, which provides most of the data and signal efficiency, and an auxiliary set, which was added to recover some additional signal efficiency. The main displaced-jet trigger is seeded by an H_T L1 trigger and then requires $H_T > 430 \text{ GeV}$ at the HLT, where H_T is the scalar sum of p_T for all calo jets satisfying $p_T > 40 \text{ GeV}$ and $|\eta| < 2.5$. The trigger also requires the presence of at least two calo jets with $p_T > 40 \text{ GeV}$ and $|\eta| < 2.0$ that also pass the inclusive tagging requirement. Since 2023, the HLT H_T threshold of this trigger has been lowered to 390 GeV , with correspondingly lower L1 H_T thresholds, and this set of triggers has been moved to the parked data.

The auxiliary set of displaced-jet triggers in standard data taking is seeded by an L1 trigger that requires $H_T > 240 \text{ GeV}$ and the presence of a muon with $p_T > 6 \text{ GeV}$, to further improve the efficiencies for LLPs with heavy flavor decays. The trigger further requires the presence of at least two calo jets with $p_T > 40 \text{ GeV}$ and $|\eta| < 2.0$ that also pass the displaced tagging requirement.

As described in Section 4, the CMS experiment employs an iterative process for the track reconstruction, and the displaced tracking iterations are usually not used at the HLT because they are CPU-intensive. However, a displaced tracking iteration is used in the displaced-jet triggers. To meet the stringent HLT CPU timing constraints, only the seeds near jets are considered for track reconstruction in the displaced-jet triggers. Furthermore, the next tracking iteration proceeds only if a given event has two tagged displaced jets based on the tracks from earlier iterations.

The total HLT rate of these displaced-jet paths is $53 (163) \text{ Hz}$ in standard (parked) data taking. The efficiency of the main displaced-jets triggers is shown in Figs. 16, 17, and 18.

Figure 16 shows the efficiency of the $H_T > 430 \text{ GeV}$ (the 2022 minimum threshold) and $H_T > 390 \text{ GeV}$ (the 2023 and later minimum threshold) HLT requirements measured in data. The measurements are performed using events collected with an isolated single-muon trigger that requires the presence of a muon with $p_T > 27 \text{ GeV}$. The efficiency of the HLT $H_T > 390 (430) \text{ GeV}$ requirement is larger than $\approx 98\%$ when the offline H_T is larger than $400 (440) \text{ GeV}$. The left plot of Fig. 17 shows the efficiency of an offline calo jet to pass the online p_T requirement in displaced-jets triggers as a function of the offline jet p_T . The right plot of Fig. 17 shows the efficiency of an offline calo jet in simulated signal events to pass the online inclusive tagging requirement as a function of the number of offline prompt tracks.

Figures 16 and 17 split the 2023 data before and after an update of the HCAL energy response corrections (and a simultaneous minor update to the HCAL gains). This conditions update was primarily due to a change in the ϕ -dependent timing alignment of the HCAL barrel, which increased the energy response. The change in energy response can impact the trigger turn-on curve, as observed for the tracking-based displaced jet triggers.

Figure 18 shows the fraction of displaced b quarks, d quarks, and τ leptons that pass the displaced-jets tagging requirement at the HLT as a function of the truth-level L_{xy} , indicating that the displaced-jets triggers have good performance for a variety of final states across the tracker volume. The overall trigger efficiencies are $\approx 0.4\text{--}1.0\%$ for $H \rightarrow SS \rightarrow q\bar{q}q\bar{q}$ signatures with $c\tau < 1 \text{ m}$. Figure 19 shows the gain in trigger efficiency for the displaced-jet triggers in Run 3 compared to Run 2. The Run 3 trigger efficiencies are larger than those of Run 2 by a factor of 4 to 17 for $10 < m_S < 60 \text{ GeV}$ and $1 < c\tau < 1000 \text{ mm}$. Compared to the standard trigger that requires $H_T > 1050 \text{ GeV}$ at the HLT, the displaced-jets triggers have more than a

factor of 20 higher efficiency for the $H \rightarrow SS \rightarrow q\bar{q}q\bar{q}$ signatures with $c\tau \gtrsim 10$ mm.

The data collected with the displaced-jets triggers in 2022 has been analyzed in a Run 3 displaced-jets search [85]. This search has achieved up to a factor of 10 increase in sensitivity compared to previous searches [84, 100] for $H \rightarrow SS \rightarrow q\bar{q}q\bar{q}$ signatures, despite analyzing a data set of a much smaller integrated luminosity. The full Run 3 data set collected with the displaced-jets triggers is expected to further improve the CMS sensitivity to this signature significantly.

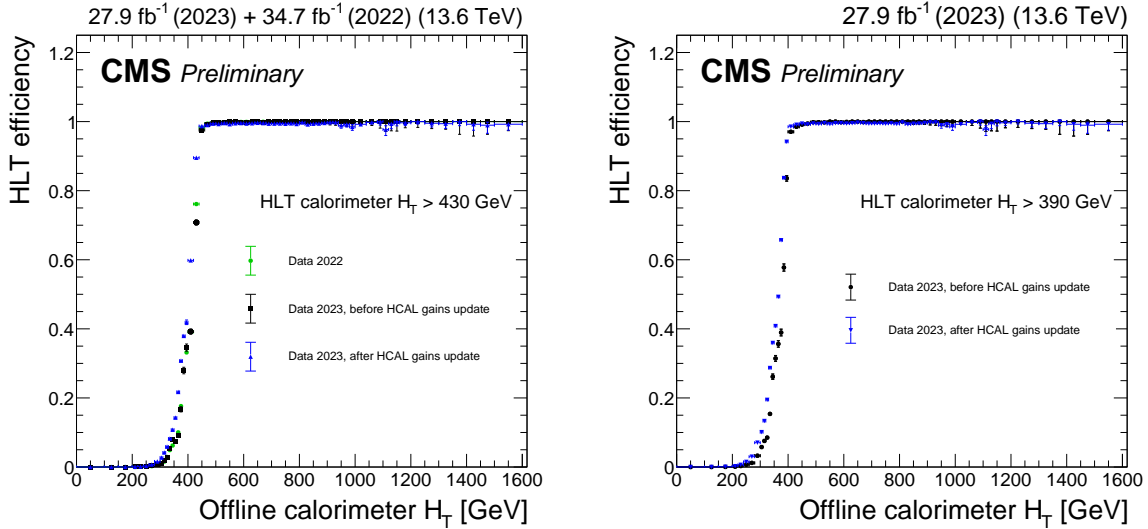


Figure 16: HLT efficiency for a given event passing the main displaced-jets trigger to satisfy HLT calorimeter $H_T > 430$ GeV (left) and $H_T > 390$ GeV (right) as a function of the offline calorimeter H_T . For this trigger, the minimum calorimeter H_T threshold was 430 (390) GeV in 2022 (2023 and later). The measurements are performed in data collected in 2022 (green circles), in 2023 before an update of the HCAL gains and energy response corrections (black squares), and in 2023 after the update (blue triangles).

6.2 Calorimeter-based algorithms

We now describe the LLP trigger algorithms that feature calorimeter information. In this section, we cover displaced-jet triggers that use the HCAL (Section 6.2.1), delayed-jet triggers using ECAL timing (Section 6.2.2), delayed-diphoton triggers (Section 6.2.3), and displaced photon+ H_T triggers (Section 6.2.4).

6.2.1 Displaced-jet triggers using the HCAL

The recent Phase I upgrade to the CMS HCAL [26] introduced both depth segmentation and precision cell timing via a time-to-digital converter (TDC) to the HCAL barrel (HB). Although the primary motivation for these enhancements is to maintain high efficiencies for particle identification in increasingly high-PU environments, both aspects of this upgrade also enable the identification of displaced- and delayed-jet signatures. Making full use of the upgraded HCAL requires new L1 trigger algorithms, where the new HCAL depth and timing information are available. Unlike many other triggers presented in this note, this required implementing new hardware-level LLP jet trigger algorithms that seed displaced and delayed jet HLT paths. This section focuses on a novel L1 trigger targeting LLPs that decay hadronically before or within

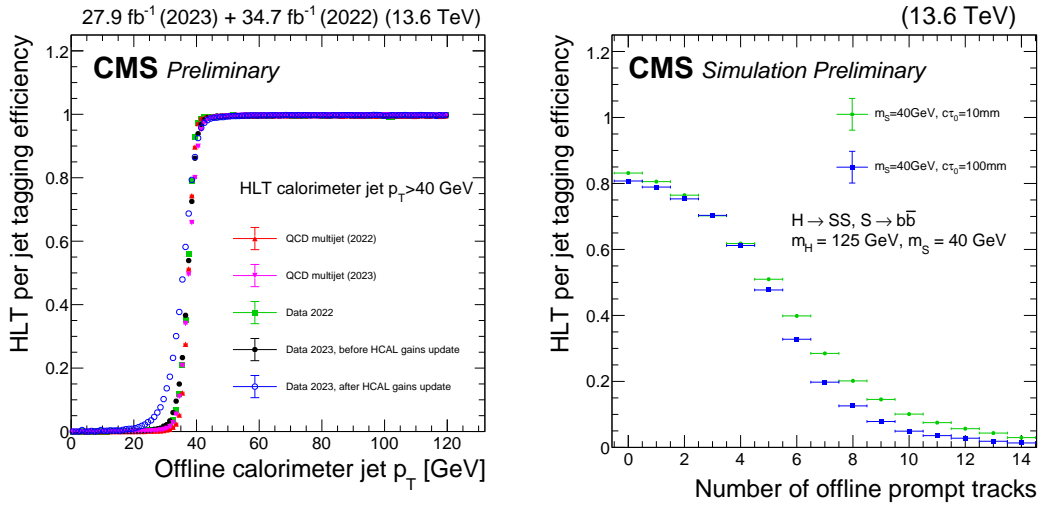


Figure 17: HLT efficiency of the main displaced-jets trigger: Efficiency of an offline calorimeter jet to pass the online p_T requirement in displaced-jets triggers (left), which require $p_T > 40$ GeV, in data collected in 2022 (green squares), in 2023 before an update of the HCAL gains and energy response corrections (black filled circles), and in 2023 after the update (blue open circles). The efficiencies measured with QCD multijet simulation are also shown, with 2022 conditions (red triangles) and 2023 conditions (purple triangles). These measurements are performed using events collected with a prescaled trigger that requires $H_T > 425$ GeV at the HLT. An offline $H_T > 450$ GeV selection is also applied to ensure the prescaled trigger reaches its plateau. The efficiency is $>96\%$ when the offline jet p_T is >40 GeV. The efficiency has a broader turn-on in the later 2023 data because of the update of the HCAL conditions. Efficiency of an offline calorimeter jet to have at most one HLT prompt track in 2022 conditions, as a function of the number of offline prompt tracks, in simulated $H \rightarrow SS, S \rightarrow b\bar{b}$ signal events where $m_H = 125$ GeV and $m_S = 40$ GeV (right). Two proper decay lengths of the S particle are shown: $c\tau = 10$ mm (green circles) and $c\tau = 100$ mm (blue squares). For jets in signal events, when the number of offline prompt tracks is <4 , the tagging efficiency is larger than 70%.

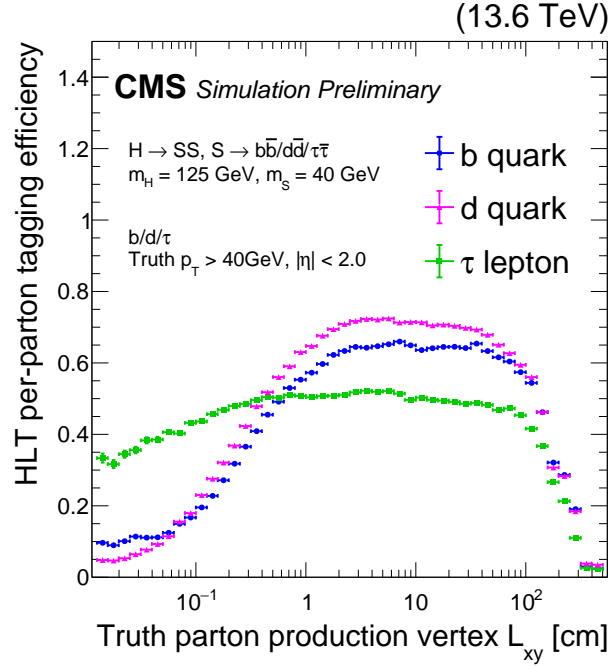


Figure 18: HLT efficiency of the main displaced-jets trigger in 2022 conditions, for $H \rightarrow SS$ signal events where $m_H = 125 \text{ GeV}$ and $m_S = 40 \text{ GeV}$. The per-parton (quark or lepton) HLT displaced-jet tagging efficiency as a function of the truth-level L_{xy} of the parton is shown for displaced b quarks (blue circles), d quarks (purple triangles), and τ leptons (green squares) with $p_T > 40 \text{ GeV}$ and $|\eta| < 2.0$.




Figure 19: The ratio between the Run 3 displaced-jets trigger efficiency and the Run 2 displaced-jets trigger efficiency as a function of LLP $c\tau$, in simulated $H \rightarrow SS, S \rightarrow b\bar{b}$ signal events where $m_H = 125 \text{ GeV}$ and $m_S = 15$ (blue triangles), 40 (green squares), or 55 (red circles) GeV. The Run-3 displaced-jets trigger efficiencies are measured in 2022 conditions.

the HB volume. The targeted signature is illustrated in Fig. 20. The signal model used to evaluate the trigger performance is an exotic decay of the Higgs boson (or decay of a high-mass, exotic Higgs boson) that produces two scalar LLPs (S), which each decay into two b quarks ($H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$), as shown in Fig. 4. The scalar S can be long-lived, with a $c\tau$ on the order of meters, permitting decays before and within the HCAL volume. The mass of the scalar is constrained to be $m_S \leq m_H/2$.

The L1 HCAL-based LLP triggers significantly increase the LLP acceptance with respect to standard triggers because of their lower jet p_T thresholds. In particular, the LLP acceptance is increased by up to a factor of 4 for a signal with $m_H = 125 \text{ GeV}$, $m_S = 50 \text{ GeV}$, and $c\tau = 3 \text{ m}$ [73]. These triggers will be used in an upcoming Run 3 search for LLPs, which relies on the upgraded HB information at both trigger and offline analysis levels. This search aims to increase the coverage of LLP models decaying hadronically in the calorimeter region, while also adding sensitivity to lower jet and event energies.

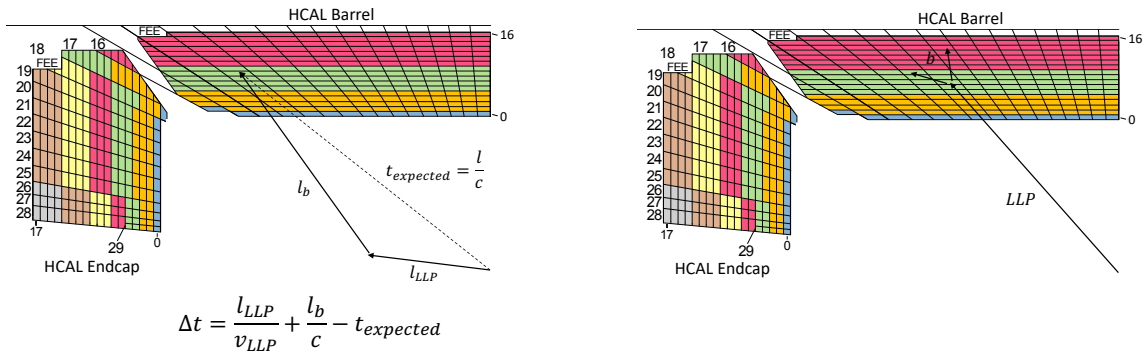


Figure 20: Diagram showing the lateral view of the CMS HCAL barrel and endcap in Run 3. The distinct depth segmentation is indicated by the different colors, where the HCAL barrel consists of 4 depth layers and the HCAL endcaps contain up to 7 depth layers. A delayed jet signature (left), where the time delay results from a combination of the low LLP velocity due to its relatively high mass and the path length difference with respect to a promptly produced particle, is shown. A displaced jet signature (right) resulting from LLPs decaying within the HCAL volume, producing significant energy deposits deeper in the HCAL with minimal energy deposits in shallower calorimeter layers, is also depicted.

The TDC provides high-precision timing information for each HCAL channel and is recorded as a 6-bit value used in the trigger pathway and offline analysis. Within each 25 ns bunch crossing, the TDC encodes fifty valid time bins in half-ns steps, along with error and invalid codes. Given the bandwidth constraints in sending timing information off-detector in the HB, each 6-bit TDC address is compressed into a 2-bit value corresponding to the following four timing values: prompt, slightly delayed, very delayed, and invalid or error. These categories are defined as:

- Prompt: $t_{\text{pulse}} \leq 6 \text{ ns}$
- Slightly delayed: $6 < t_{\text{pulse}} \leq 7 \text{ ns}$
- Very delayed: $7 < t_{\text{pulse}} < 25 \text{ ns}$
- Invalid pulse: invalid or error code

In Run 3, the barrel region of HCAL is segmented into four depths, each with time and energy readouts, as illustrated in Fig. 20. Each HB depth is made of interleaved layers of scintillator and brass absorber, or stainless steel for the first and last absorbers [26]. To accommodate the

different behavior and pulse shapes observed in the innermost depth (layer 0), two updates are made for 2024 data taking [101]: first, the prompt range for depth 1 is set to $t_{\text{pulse}} \leq 9$ ns; second, the slightly delayed range is extended to encompass a 3 ns wide window, with $9 < t_{\text{pulse}} \leq 12$ ns for depth 1 and $6 < t_{\text{pulse}} \leq 9$ ns for all other depths.

The complete information from each of the four HB depth segmentations cannot be sent to the trigger because of bandwidth constraints. Therefore, quantities are calculated per-trigger tower (a group of cells across four depths at the same η and ϕ coordinates) and encoded in a per tower LLP flag that is used to compute the per jet LLP flag. An HCAL tower will pass the tower LLP flag if it satisfies either of the following:

- Delayed tower: Includes at least one delayed cell with energy >4 GeV and no prompt cells with energy >4 GeV
- Displaced tower: Includes a cell in depths 3 or 4 with energy >5 GeV and both cells in depths 1 and 2 with energy <1 GeV.

Any L1 jet containing at least two LLP-flagged towers within the 9×9 HCAL towers included is flagged as an L1 LLP jet. Additional kinematic selections are then applied, including $H_T > 200$ GeV and jet $p_T > 60$ GeV for the lowest unprescaled L1 single LLP jet trigger and two jets with $p_T > 40$ GeV (with no additional H_T requirement) for the L1 double LLP jet trigger. The stringent requirement of identifying multiple delayed or displaced towers at L1 permits a comparably low H_T minimum requirement, increasing sensitivity to a range of LLP models [80–82, 86–99].

Timing scans in the HCAL were performed in 2022, 2023, and 2024 [101]. Timing scans are performed at the start of data taking each year and demonstrate the stability of HCAL timing, with no major timing adjustments needed between 2023 and 2024. These scans produce artificially delayed jets across the whole detector at a known time in collision data and are crucial to understanding the detector and trigger performance. From the timing scan data, the delayed tower and jet turn-on are measured, with the timing resolutions of jets in collisions. Figures 21 and 22 show the efficiency of the per tower LLP flag and the per jet LLP flag, demonstrating the high degree of sensitivity of these flags to jet delay. In particular, jets with an arrival time ≥ 6 ns are identified as delayed with a high efficiency.

Figure 23 shows the L1T efficiency of the HCAL-based LLP jet trigger as a function of offline H_T and calorimeter jet p_T , for jets matched to LLPs decaying in HCAL depths 3 and 4, that is, where $214.2 \leq R < 295$ cm and $|\eta| \leq 1.26$. The efficiency plateau is reached at 300 GeV in event H_T and 125 GeV in calorimeter jet p_T . Notably, this enables these dedicated LLP jet triggers to increase the efficiency for relatively low-energy events that are inaccessible with purely energy-based triggers. The HCAL-based LLP jet triggers increase the acceptance at L1 to signals with $m_H = 125$ GeV, $m_S = 50$ GeV, and $c\tau_S = 3$ m by up to a factor of 4 times that of standard L1 H_T triggers. These paths provide added efficiency for low p_T and H_T LLP events that are more difficult to trigger on with the displaced-jet HLT paths seeded by L1 H_T triggers that are discussed in Section 6.1.3.

Without a dedicated L1T HCAL-based LLP trigger, standard L1T jet or H_T triggers could be used to seed the displaced-jet paths at the HLT. At the L1T in Run 3, the lowest minimum threshold for an unprescaled single jet trigger is at 180 GeV, and this trigger reaches 80% efficiency at an offline jet p_T of 200 GeV. As of 2023, the lowest minimum threshold for an unprescaled L1T H_T trigger is at 280 GeV, and this trigger reaches 80% efficiency at an offline H_T of 320 GeV. Before the start of 2023 data taking, the lowest minimum threshold for an unprescaled L1T H_T trigger was at 360 GeV, and this trigger reaches 80% efficiency at an offline

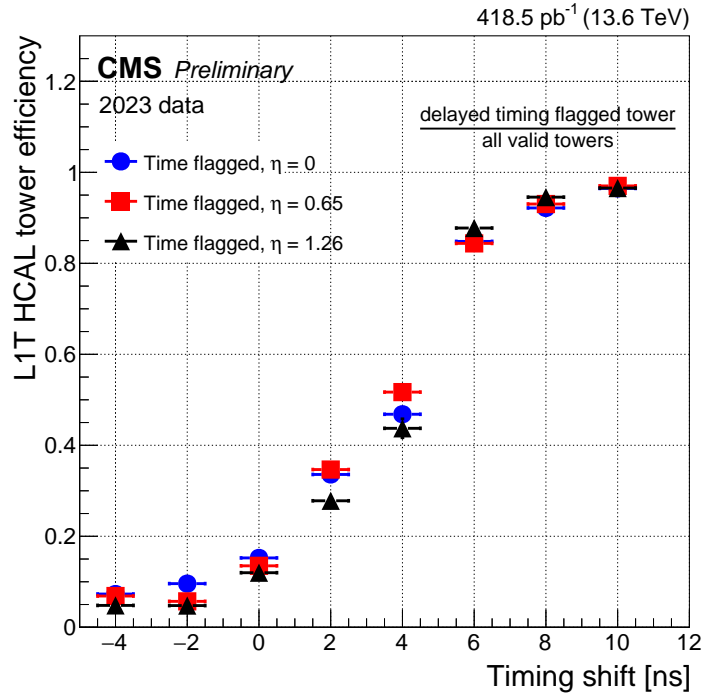


Figure 21: L1T HCAL tower efficiency of the timing-flagged towers in 2023 data: HCAL delayed timing tower efficiency during an HCAL timing phase scan during 2023, with efficiencies split by trigger towers centered at $\eta \approx 0$ (blue circles), 0.65 (red squares), 1.26 (black triangles) and with width $\Delta\eta = 0.087$. The sharp turn-on between timing delays of 0-6 ns is expected with the prompt timing range set at $t_{\text{pulse}} \leq 6$ ns, demonstrating the timing trigger performance. The timing-flagged towers must have at least one delayed cell, no prompt cells, and energy over 4 GeV. The efficiency is calculated relative to towers with any valid timing code, meaning the tower contains at least one cell with energy ≥ 4 GeV and a TDC code of prompt, slightly delayed, or very delayed. Multiple flagged towers are required for the HCAL-based displaced and delayed jet L1T to be set, and this shows the turn-on at a per-tower level relative to incoming pulse timing.

H_T of 400 GeV. A L1T H_T seed of 360 GeV has been the lowest-threshold seed used for the L1T H_T -seeded displaced dijet HLT paths in Run 2 and Run 3, and thus is referenced for comparison.

Figure 24 demonstrates the efficiency of the HCAL-based LLP jet trigger, as a function of LLP decay radius $R = \sqrt{\Delta x^2 + \Delta y^2}$. The LLP is required to decay within $|\eta| \leq 1.26$ and be matched to an offline jet with $p_T > 100$ GeV. The L1 trigger has an efficiency of more than 90% for LLP decays within HCAL depths 3 and 4 ($214.2 \leq R < 295$ cm), the region targeted by the depth trigger to identify displaced LLP jets.

In general, the high-mass sample reaches a higher efficiency at L1, as heavier LLPs are more likely to produce jets that pass the stringent calorimeter energy deposits (multiple cells with over 5 GeV energy deposited in an individual depth) for the depth-based trigger. Low-mass samples are typically more difficult to trigger, especially as their H_T spectra are lower, and therefore, energy-based selections are more difficult to satisfy. The high-mass sample also reaches higher L1T efficiencies in jet p_T and event H_T because of the decay kinematics. In this sample, the LLP is more boosted, leading to more collimated decay products. This is beneficial as the trigger requires multiple nearby depth-flagged trigger towers. The LLP $c\tau$ primarily impacts the signal acceptance rather than the efficiency.




Figure 22: L1T efficiency of the LLP jet trigger in 2023 data: The HCAL LLP-flagged L1T trigger delayed jet fraction versus jet E_T during the 2023 HCAL phase scan demonstrates that the delayed jet fraction approaches unity as the timing shift, with units in ns, is increased. The figure shows results inclusive in pseudorapidity for the HCAL barrel, corresponding to $|\eta| \leq 1.35$. The fraction of LLP-flagged L1 jets is compared to all L1 jets from a dataset of events enriched with jets or p_T^{miss} . No explicit selection criterion is applied on the jet E_T , though the implicit requirement for a jet to have at least two cells with $E_T \geq 4$ GeV shapes the resulting jet turn-on curve.

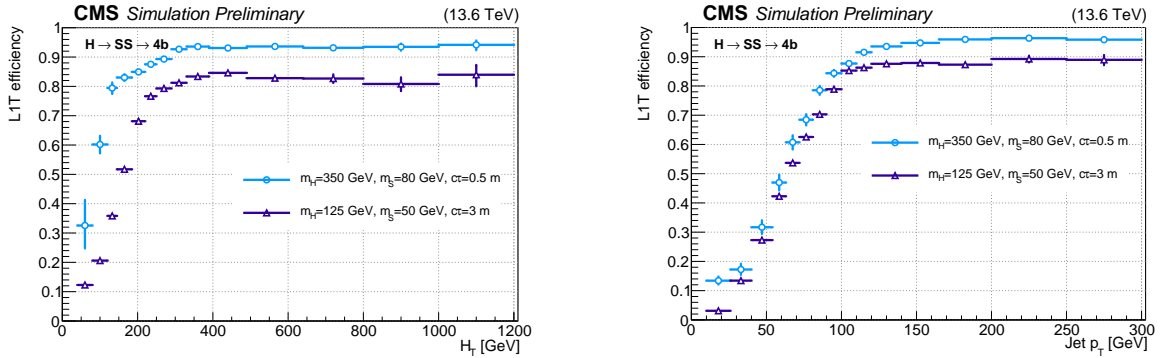


Figure 23: L1T efficiency of the HCAL-based LLP jet triggers, as a function of event H_T (left) and jet p_T (right), for $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ events with $m_H = 350$ GeV, $m_S = 80$ GeV, and $c\tau_S = 0.5$ m (light blue circles) and $m_H = 125$ GeV, $m_S = 50$ GeV, and $c\tau_S = 3$ m (purple triangles), in 2023 conditions. The trigger efficiency is evaluated for LLPs decaying in HB depths 3 or 4, corresponding to $214.2 \leq R < 295$ cm and $|\eta| \leq 1.26$. These LLPs are also required to match to an offline jet in HB.

A suite of displaced- and delayed-jet HLT paths is seeded by the HCAL-based LLP L1 jet triggers. The total rate of these HLT paths is about 35 Hz. Since the start of Run 3, a number of these paths follow the HLT approach detailed in Sections 6.1.3 and 6.2.2. Owing to the seeding




Figure 24: L1T efficiency of the HCAL-based LLP jet triggers as a function of LLP decay radius R for $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ events with $m_H = 350$ GeV, $m_S = 80$ GeV, and $c\tau_S = 0.5$ m (light blue circles) and $m_H = 125$ GeV, $m_S = 50$ GeV, and $c\tau_S = 3$ m (purple triangles), in 2023 conditions. The trigger efficiency is evaluated for LLPs within $|\eta| \leq 1.26$ where either the LLP or its decay products are matched to an offline jet in HB with $p_T > 100$ GeV.

with selective L1 triggers, the HLT H_T requirement is reduced by 230 GeV, enabling the lowest threshold trigger with $H_T > 170$ GeV and calo jet $p_T > 40$ GeV. The lowest HLT H_T threshold of 170 GeV benefits from the unprescaled L1 double-jets seed, which has no H_T requirement at L1.

In 2024, an additional HLT path was added to target highly displaced jets arising from LLPs decaying hadronically inside the calorimeters, complementing the depth-based L1 seed. This HLT path, referred to as the “CalRatio” trigger, selects for jets that have a large neutral hadron energy fraction (NHF). The composition of these objects is not necessarily enriched in neutral hadrons; instead, this selection is associated with a large fraction of energy deposited in the HCAL with minimal ECAL and tracker activity—a defining feature of displaced jets produced inside the calorimeters. This trigger selects events with $H_T > 200$ GeV containing at least one calo jet with $p_T > 60$ GeV, $|\eta| < 1.5$, and hadronic energy fraction > 0.7 ; to further suppress background and keep rates reasonable, the trigger additionally requires at least one PF jet with the same kinematic requirements and $NHF > 0.7$. The HLT efficiency of the path, together with characteristic distributions of this variable in data, signal, and background, is shown in Fig. 25. The efficiency plateaus at around 90% for $NHF > 0.8$; the reason it does not reach 100% is largely because the additional offline kinematic selections in the trigger (e.g., H_T and jet p_T), which are equivalent to their online quantities, are below their respective efficiency plateaus.

6.2.2 Delayed-jet triggers using ECAL timing

Dedicated triggers targeting delayed jet signatures were introduced at the HLT during Run 3 to improve the performance for nonprompt jets [12]. These triggers are highly efficient for LLPs that decay into jets within the tracker or the ECAL volume, as shown in Fig. 26. Many extensions of the SM predict the existence of LLPs that decay into jet final states, including GMSB models [82] and hidden valley models [102]. The delayed jet triggers utilize the timing capabilities of the ECAL [103] to search for such signatures.

These HLT paths are seeded by L1 triggers that require either $H_T > 350$ GeV or a τ_h candidate with $p_T > 120$ GeV. At the HLT, requirements on the jet timing are placed on single- or double-inclusive and trackless calo jets. An inclusive jet is defined as a calo jet that is matched to a

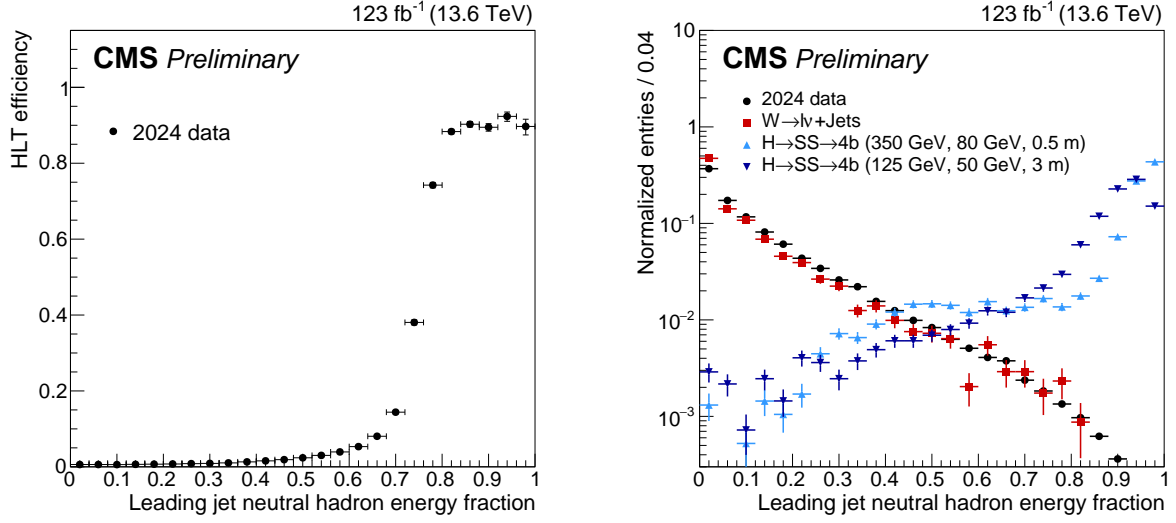


Figure 25: HLT efficiency of the CalRatio trigger as a function of the leading PF jet NHF in 2024 data, measured with respect to a logical OR of the HCAL-based LLP L1 jet triggers (left). Distribution of the leading PF jet NHF (right) in 2024 data (black circles), $W \rightarrow \ell\nu$ background simulation in 2024 conditions (red squares), and $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ signal simulation in 2023 conditions (blue and purple triangles). Events are required to have $H_T > 200$ GeV and the leading jet is required to have $p_T > 60$ GeV and $|\eta| \leq 1.5$, which are equivalent to the respective HLT jet object selections. The signal distributions additionally require the leading jet to be matched to an LLP decaying anywhere inside the barrel calorimeter volume ($129 \leq R < 295$ cm). The clear separation between the displaced signal and the prompt background in the right plot motivates the development of the CalRatio trigger.

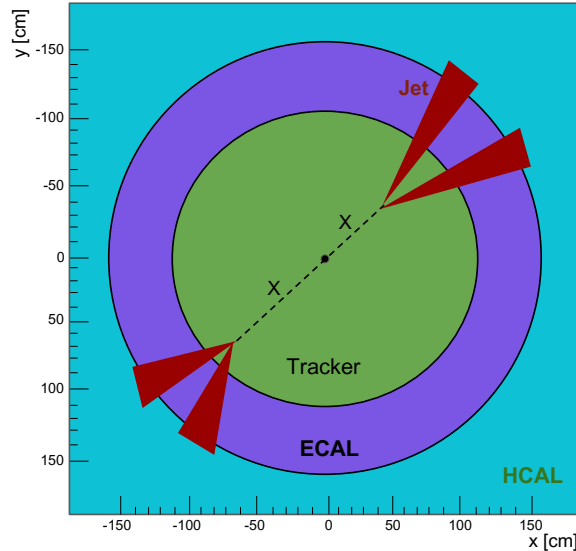


Figure 26: Diagram of a typical delayed-jet signal event.

reconstructed track from the tracker, and a trackless jet is defined as a calo jet that does not have a matched prompt track. The triggers at the HLT include a variety of timing thresholds to balance the trigger rates against the signal efficiency. For example, the timing thresholds for the single jet triggers at the HLT are ≥ 2 ns, and the timing resolution is around 0.4 ns. The delayed jet triggers are designed to be model-independent and performant for lower H_T final

states than those of the Run 2 analysis, which targeted final states with $p_T^{\text{miss}} > 300 \text{ GeV}$ [12]. Figure 27 shows the improvement in the trigger efficiency in Run 3 with the dedicated delayed-jet triggers, compared to that of Run 2, where a standard trigger that requires $H_T > 1050 \text{ GeV}$ was available. The improvement is particularly noticeable at low H_T , for a signal model where very heavy BSM Higgs bosons decay into generic LLPs (X) and subsequently into 4 b jets, $H \rightarrow XX \rightarrow b\bar{b}b\bar{b}$, the Feynman diagram for which is shown in Fig. 4. Some trigger inefficiencies are introduced because H_T is calculated offline using jets with $p_T > 40 \text{ GeV}$ and differs from the H_T calculated at the HLT.

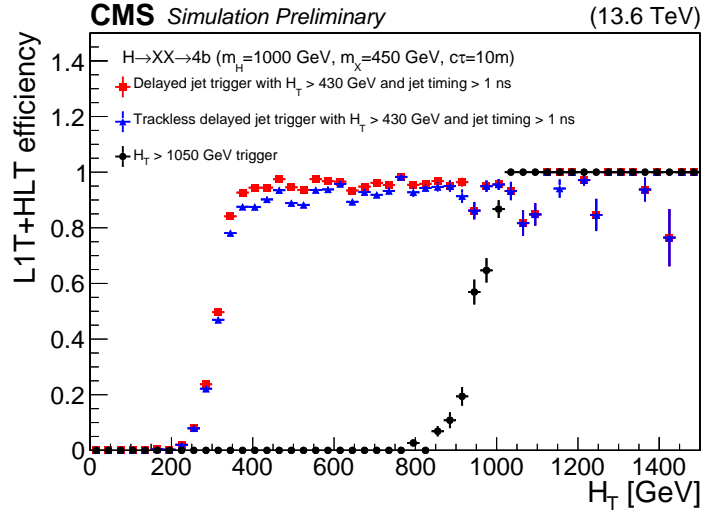


Figure 27: L1T+HLT efficiency of the inclusive (trackless) delayed-jet triggers introduced in Run 3 in red squares (blue triangles), in 2022 conditions, and the H_T trigger (black circles), which was all that was available in Run 2, for a $H \rightarrow XX \rightarrow b\bar{b}b\bar{b}$ signal with $m_H = 1000 \text{ GeV}$, $m_X = 450 \text{ GeV}$, and $c\tau = 10 \text{ m}$. The addition of these delayed jet triggers results in a significant improvement in the efficiency of the signal for $430 < H_T < 1050 \text{ GeV}$.

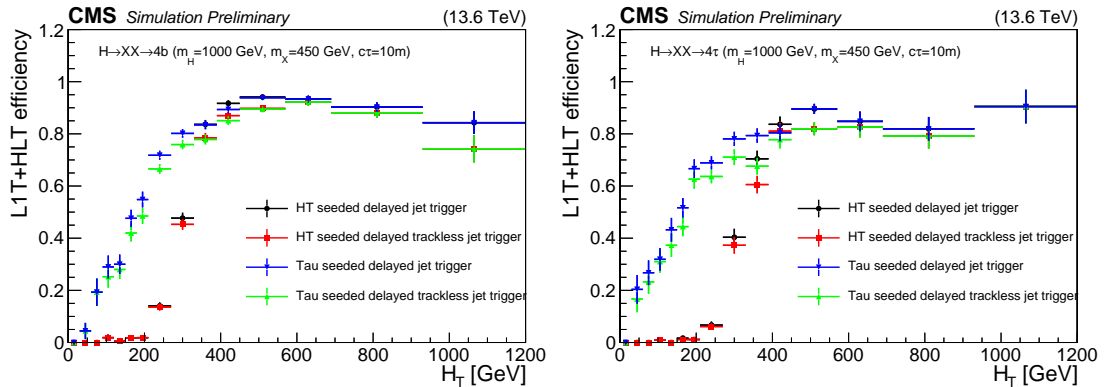


Figure 28: L1T+HLT efficiency of the delayed-jets triggers for signal models with $H \rightarrow XX \rightarrow b\bar{b}b\bar{b}$ (left) and $H \rightarrow XX \rightarrow 4\tau$ (right), with $m_H = 1000 \text{ GeV}$, $m_X = 450 \text{ GeV}$ and $c\tau = 10 \text{ m}$, in 2022 conditions. The improvement from the tau triggers (blue and green triangles) can be seen in the $H_T < 430 \text{ GeV}$ region compared to the H_T -seeded triggers (black circles and red squares). These plots include events with jets with $p_T > 40 \text{ GeV}$, number of ECAL cells > 5 , barrel region with $|\eta| < 1.48$ and jet timing $> 2 \text{ ns}$.

The inclusion of the tau candidate seed at the L1T improves the analysis sensitivity to low- H_T signal models, targeting LLPs that decay into tau leptons in addition to jets. Figure 28 demon-

strates the improvement in signal efficiency because of the L1 tau seed addition, particularly at the low- H_T final states for $H \rightarrow XX \rightarrow b\bar{b}b\bar{b}$ and $H \rightarrow XX \rightarrow 4\tau$ decays.

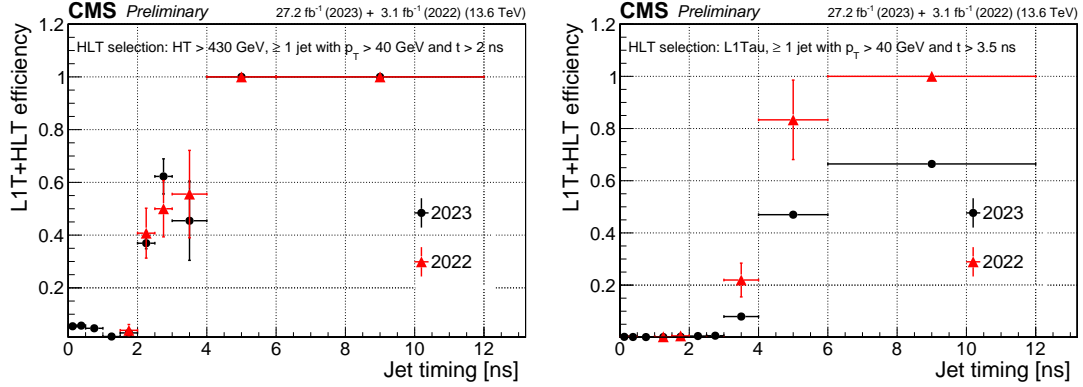


Figure 29: L1T+HLT efficiency of the delayed-jets triggers as a function of jet timing for 2022 and 2023 data-taking periods for the H_T -seeded trigger (left) and the τ -seeded trigger (right). A clear turn-on feature can be seen around the threshold values. The plots include events that pass the $p_T^{\text{miss}} > 200$ GeV trigger and have at least one barrel jet with $p_T > 50$ GeV, number of ECAL cells > 8 , and ECAL energy > 25 GeV. The H_T is calculated using the scalar sum of jets with offline $p_T > 40$ GeV, and this is different from the H_T calculation used at the HLT level, which can cause trigger inefficiencies. The maximum triggerable jet time is 12.5 ns.

The total trigger rate of the ECAL-based delayed-jet triggers at the HLT is 37 Hz in standard data taking. Figure 29 shows the trigger performance in Run 3 data collected in 2022 and 2023 as a function of jet timing. A clear turn-on behavior can be seen at the trigger timing threshold.

To further increase the analysis sensitivity, triggers with lower timing thresholds are introduced in the parked data with an upper bound on timing to ensure minimal overlap with the triggers in standard data taking [33]. These triggers take advantage of the higher rate allocation in the parked data and have a lower timing threshold, thus providing higher signal acceptance. The parking triggers increase the signal acceptance by more than 30% compared to their counterparts in the standard data taking. The total trigger rate for these triggers in the parked data is 77 Hz.

6.2.3 Delayed-diphoton triggers

Electrons and photons produced from the decay of an LLP whose mean path length is of the order of the CMS tracker dimensions will produce a distinct signature in the ECAL. The peak signal time of such an electromagnetic signature will be delayed compared to an electron or photon created from the proton collision vertex at the center of the detector. This is conceptually demonstrated in Fig. 30. The time delay at the seed crystal of the ECAL supercluster, called e/γ seed time, is chosen as a representative value of the time delay of the delayed ECAL signature. For a displaced vertex in the tracker inner pixel region, the expected delay in the peak ECAL electromagnetic signature is of the order of nanoseconds, which is significantly higher than the ECAL timing resolution. Figure 31 shows the expected ECAL time delay for different LLP $c\tau$ values and 2023 data with only online ECAL calibrations. In data, the two largest p_T electrons that pass the tight electron identification [20] and have invariant mass between 86 and 98 GeV are considered for the distribution. In the MC simulation, an ECAL supercluster reconstructed with the online reconstruction algorithm is angularly matched to a truth electron. The LLP decay appears in the ECAL $\mathcal{O}(ns)$ later than the prompt $Z \rightarrow ee$ data. (The peak of the $Z \rightarrow ee$ data distribution is at less than 0 ns because only online calibrations are used in this plot.) The

ECAL timing measurement is influenced by the radiation damage to the ECAL crystals, which reduces the efficiency to trigger on delayed electrons and photons.




Figure 30: Diagram of a prompt (orange dotted line) and displaced (red solid line) electron, which comes from the decay of a chargino, reaching the ECAL in the x - z plane. The ECAL crystal clocks are shifted so that a prompt electron arrives at the same time stamp for any crystal. An electron produced at a displaced vertex thus arrives delayed compared to a prompt particle. This is diagrammatically shown in the clocks above the ECAL crystals in blue.

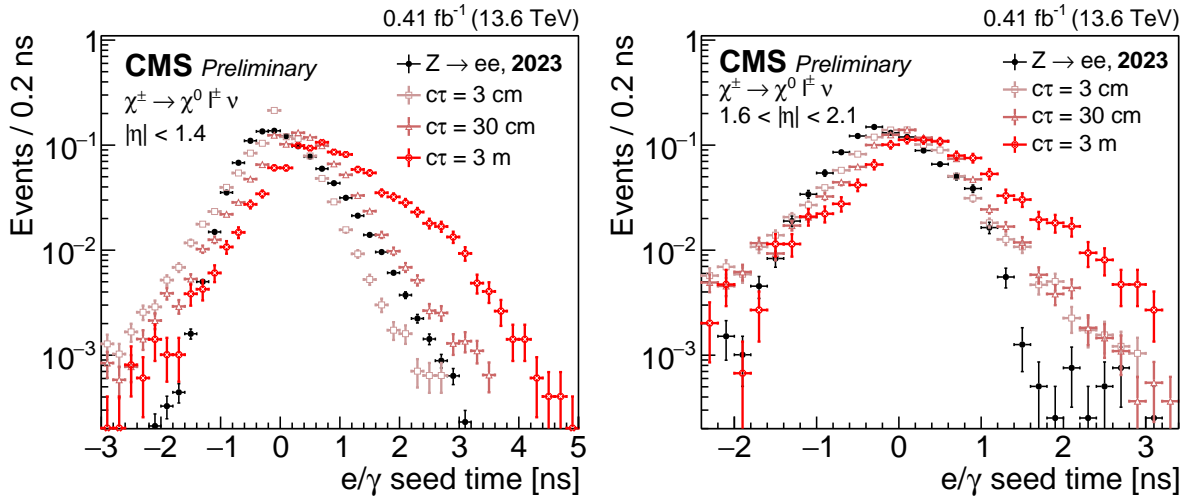


Figure 31: ECAL time delay of the e/γ L1 seeds in the barrel (left) and endcap (right). The distributions are shown for 2023 $Z \rightarrow ee$ data (black points) and $\chi^0 c\tau$ values of 3 (light red circles), 30 (darker red circles), and 300 cm (bright red circles), assuming the singlet-triplet Higgs dark portal model ($\chi^\pm \rightarrow \chi^0 \ell \nu$, where the χ^\pm has a mass of 220 GeV and the χ^0 has a mass of 200 GeV). The distributions are normalized to unity.

A new trigger was introduced in Run 3 to improve the CMS sensitivity to delayed electrons and photons arising in the singlet-triplet Higgs portal dark matter model discussed in Ref. [72]. This trigger uses single and double e/γ seeds at the L1T, and it selects both delayed electrons and photons with the same HLT path by making requirements only on ECAL superclusters. No electron track requirement is made to create specific trigger paths for electrons because the typical electron track reconstruction algorithm is not performant for displaced electrons. Although this trigger is highly efficient for both delayed electrons and photons, we call it “the delayed-diphoton trigger” in what follows for simplicity.

The delayed diphoton trigger requires at least two ECAL superclusters with $|\eta| < 2.1$ and a minimum energy of 22 (12) GeV for the leading (subleading) supercluster. This asymmetric minimum energy threshold is reduced to a symmetric 10 GeV threshold in both superclusters if the event contains PF $H_T > 300$ GeV. The time delay for the crystal with the largest energy is required to be larger than 1 ns for both the ECAL superclusters.

The total rate of these HLT paths is about 15 Hz. Figure 32 (left) shows the variation of the delayed-diphoton trigger rate with integrated luminosity and PU in 2024. The trigger rate decreases nonlinearly during a single fill because the increasing opacity in the ECAL crystals causes the timing distribution to shift towards negative values. This opacity from radiation damage recovers by the start of the next fill, with very little permanent damage between individual fills. Over longer periods of collisions, this radiation damage could substantially reduce the signal sensitivity. However, periodic updates to the online calibrations mitigate this damage. The calibrations correct for a shift in timing towards negative values and maintain the trigger's efficiency for signal processes with a delay of 1 ± 0.2 ns. The widening of the timing resolution throughout the year from the radiation damage results in an increased trigger rate. In 2024 data taking, the rate of this trigger with respect to PU is linear, as shown in Fig. 32 (right).

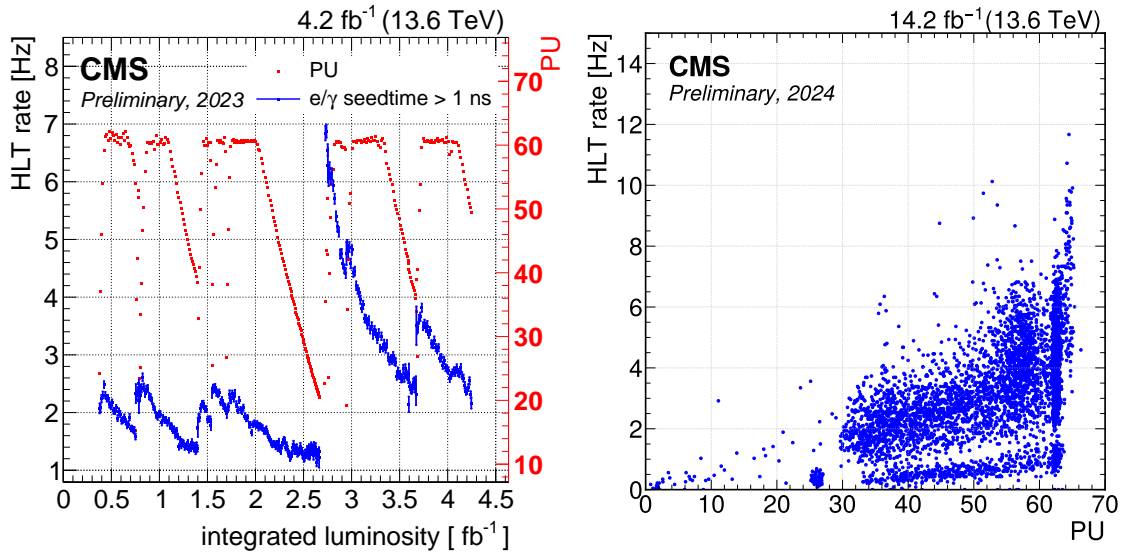


Figure 32: HLT rate (blue points) of the delayed-diphoton trigger in the first data collected in 2023, corresponding to an integrated luminosity of 4.2 fb^{-1} , compared with the PU during the same data-taking period (red points), as a function of integrated luminosity (left). The rate decreases nonlinearly, and then changes to linear asymptotically through a single collision fill. It recovers by the start of the next fill with $< 1\%$ reduction in rate between the fills. The rate generally increased throughout the year because of periodic online calibrations to mitigate the loss in trigger efficiency due to radiation damage of the ECAL crystals. The delayed-diphoton trigger rate shows a linear dependency on PU in 2024 data, at an instantaneous luminosity of approximately $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ (right).

Data collected in the later part of 2023 and corresponding to an integrated luminosity of 10.82 fb^{-1} are used to measure the trigger efficiency of the delayed diphoton HLT path, using the tag and probe method with $Z \rightarrow ee$ decays. The coarse online calibration results in a large number of Z boson decays with a prompt signature delayed by more than 1 ns. The tag electron was required to have $p_T > 20 \text{ GeV}$, $|\eta| < 2.1$, cluster seed time > 1 ns, and to have

passed the tight identification [20]. Probe electrons are only considered if the tag and probe electron pair are oppositely charged and in the invariant mass window $84 < m_{ee} < 96 \text{ GeV}$, consistent with the mass of the Z boson. The probe electron is required to have $p_T > 4 \text{ GeV}$ and $|\eta| < 2.5$. The trigger efficiency is defined as the ratio of the probe electrons passing the HLT path to all the probe electrons passing the selection criteria. Figure 33 shows the trigger efficiency as a function of the subleading probe electron seed time, which sharply rises at a seed time of 1 ns and plateaus thereafter. Figure 34 shows the efficiency as a function of the subleading probe electron p_T and η , for events with a seed time $> 1 \text{ ns}$.




Figure 33: L1T+HLT efficiency of the delayed-diphoton trigger as a function of the subleading probe electron (e_2) supercluster seed time, measured with data collected in 2023. At the HLT, the subleading e/γ supercluster (e/γ_2) is required to have $p_T > 12 \text{ GeV}$, $|\eta| < 2.1$, and a seed time $> 1 \text{ ns}$. The trigger is fully efficient above 1 ns.

6.2.4 Displaced photon+ H_T triggers

A trigger designed to collect events containing a displaced photon in addition to hadronic activity was originally deployed online in 2017, during Run 2. This displaced photon+ H_T trigger was designed for the search for delayed photons [11], which targets displaced and delayed photons coming from the decays of long-lived neutralinos in a GMSB model [71]. The delayed photon signature is shown in Fig. 35. At the HLT, this trigger requires a photon with $p_T > 60 \text{ GeV}$; a set of loose, calorimeter-based identification requirements; loose isolation; and a loose set of requirements on the calorimeter shower shape, such that the shower is compatible with a displaced and delayed photon. In particular, the shower would be more elliptical in the η - ϕ plane for displaced photons than for prompt photons, and so the length of the major (minor) axis of the shower S_{major} (S_{minor}) is required to be less than 0.4 (1.5). In addition, in Run 3, the trigger requires PF $H_T > 350 \text{ GeV}$, where H_T is the scalar sum of p_T for all jets satisfying $p_T > 30 \text{ GeV}$ and $|\eta| < 2.5$, to keep the rate around 12 Hz at the HLT. The L1 seeds of the HLT path require at least one e/γ object, tau lepton, or jet.

The efficiency of the displaced photon+ H_T trigger is shown in Fig. 36. Events in 2017 data (Run 2) and a GMSB signal benchmark with $\Lambda = 100 \text{ TeV}$ and $c\tau = 10$ or 1000 cm in 2017 conditions are selected that pass a trigger requiring isolated muons with $p_T > 27 \text{ GeV}$ and at least one offline photon with $\eta < 1.4442$ and passing the displaced photon identification criteria as described in Ref. [11]. Only the leading photon satisfying these requirements is

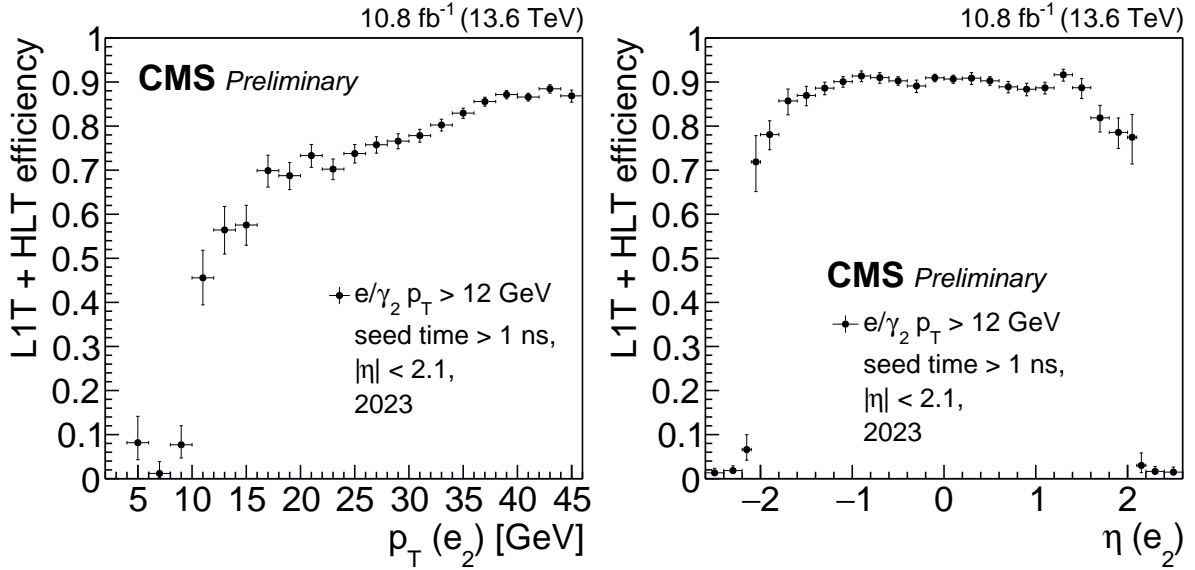


Figure 34: L1T+HLT efficiency of the delayed-diphoton trigger as a function of subleading probe electron (e_2) p_T (left) and η (right), measured with data collected in 2023. At the HLT, the subleading e/γ supercluster (e/γ_2) is required to have $p_T > 12$ GeV, $|\eta| < 2.1$, and a seed time > 1 ns. The efficiency rises sharply for $p_T > 12$ GeV and plateaus for $p_T > 35$ GeV. The slow rise in between is from additional L1 H_T requirements. The trigger is efficient in the region $|\eta| < 2.1$.

considered in the trigger efficiency measurements. Figure 36 shows the efficiency of the trigger as a function of offline photon p_T (left) and the event H_T (right). These plots indicate that the trigger is fully efficient for photon $p_T > 70$ GeV and $H_T > 400$ GeV. Note that in 2017, the HLT path required $H_T > 350$ GeV for jets satisfying $p_T > 15$ GeV and $|\eta| < 3$. The HLT requirements were subsequently changed to $H_T > 300$ GeV for jets satisfying $p_T > 30$ GeV and $|\eta| < 2.5$. The efficiency of the standard photon identification and the displaced photon identification in the displaced photon+ H_T trigger is over 95%. For a typical high-PU run in 2017, the displaced photon identification requirements reduced the rate of the trigger by 20%.

In 2022, the rate of this trigger increased more than initially expected as a result of the larger PU. The HLT rate as a function of PU was also nonlinear, and thus difficult to maintain. The trigger was improved for 2023 data taking by increasing the minimum p_T requirement on the jets included in the PF H_T calculation at the HLT. After raising the minimum jet p_T threshold from 15 to 30 GeV, the total HLT rate was reduced back to 12 Hz (about 4 Hz pure), and the behavior with respect to PU was again linear. Figure 37 shows the rate of this trigger in 2022 and 2023 data as a function of PU.

6.3 Muon spectrometer-based algorithms

This section describes trigger algorithms designed to select LLP decays in the muon system. These include displaced-dimuon triggers (Section 6.3.1), double displaced L3 muon triggers (Section 6.3.2), displaced muon+photon triggers (Section 6.3.3), dimuon scouting triggers (Section 6.3.4), and MDS triggers with the CSCs (Section 6.3.5) and with the DTs (Section 6.3.6).

6.3.1 Displaced single and dimuon triggers

An LLP produced in a hard interaction of the colliding protons can traverse a considerable distance in the detector before decaying into muons. When muons are produced well within




Figure 35: Diagrams of a promptly-produced photon in the x - y plane showering in the ECAL (left), a delayed photon produced from a long-lived neutralino $\tilde{\chi}^0$ in the x - y plane and showering later in the ECAL (middle), and an elliptical shower in the η - ϕ plane produced from a delayed photon (right). The lengths of the major (S_{major}) and minor (S_{minor}) axes of the shower and the reconstructed hits (RecHits) that compose the shower are labeled.




Figure 36: L1T+HLT efficiency of the displaced photon+ H_T trigger as a function of photon p_T (left) and event H_T (right), for 2017 data (black points) and a GMSB signal with $\Lambda = 100$ TeV and $c\tau = 10$ cm (blue triangles) and $c\tau = 1000$ cm (red squares). Left: The efficiency to obtain a photon at the HLT passing the trigger requirements is shown. This includes the L1T component. Right: The efficiency to pass the H_T trigger requirement is shown.

the silicon tracker, they can be reconstructed by both the tracker and the muon system. However, if the muons are produced with larger displacements, they can only be reconstructed by




Figure 37: Total rate of the displaced photon+ H_T HLT path in 2022 data (left), at an instantaneous luminosity of approximately $1.8 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, and 2023 data (right), at an instantaneous luminosity of approximately $2.0 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, as a function of PU. The rate vs PU behavior was nonlinear in 2022 and fixed in time for 2023 data taking.

the muon system. A pair of reconstructed muon tracks is fitted to a common secondary vertex forming a “displaced dimuon”. All reconstructed dimuon events are classified into two mutually exclusive categories: (1) where both muons are global or tracker muons (TMS-TMS category), or (2) where both muons are standalone muons (STA-STA category). (The TMS-STA category is not considered because it provides only a small fraction of events for the considered signals.) Owing to the excellent resolution of the tracker, the position and momentum resolution of global and tracker muons are far superior to that of their standalone muon counterparts. This enhanced position resolution leads to an improvement in the resolution of the TMS-TMS dimuon vertex of a factor of 100 to 1000 for simulated HAHM events [70], depending on the lifetime of the LLP and the masses of the exotic Higgs boson and the LLP. Figure 38 shows an event display that illustrates the displaced dimuon signature.

Since Run 2, a set of dedicated triggers has existed that relies on muons reconstructed only in the muon system using algorithms that are similar to those used in the standalone offline muon reconstruction. These were deployed in 2016 and 2018 and included a beamspot constraint in the muon track fits at the L1T, but not at the HLT. This beamspot requirement caused a large reduction in displaced muon reconstruction efficiency for large LLP lifetimes. Each HLT muon was required to be within the region $|\eta| < 2.0$ and to have $p_T > 28$ (23) GeV for 2016 (2018) data taking. The CMS Collaboration performed a Run 2 data analysis that used these triggers in a search for LLPs decaying to a pair of muons [104]. During Run 3, these triggers are also used, maintaining the p_T threshold of 23 GeV from 2018.

To enhance the signal efficiency, a novel set of triggers, both at the L1T and HLT, was introduced at the beginning of Run 3. These improvements made an early search [105] possible, and its results are competitive with those of Run 2, using only 2022 Run 3 data with less than half of the Run 2 integrated luminosity.

At the L1T, two additional sets of triggers were introduced. One set comprises double-muon triggers with either no or a very low (4.5 GeV) minimum muon- p_T threshold. These triggers require that the L1 muon candidates have opposite signs, segments in at least three different




Figure 38: A Run 3 event containing a candidate LLP decay into a pair of muons away from the interaction point, reconstructed in the CMS detector. The red lines correspond to the two muons, which are detected only in the muon system. The muon tracks are used to calculate a dimuon vertex, indicated by the white circle, where the LLP is hypothesized to have decayed.

muon stations, angular separation $\Delta R_{\mu\mu} = \sqrt{(\Delta\eta_{\mu\mu})^2 + (\Delta\phi_{\mu\mu})^2}$ that does not exceed a maximum threshold that varies between 1.2 and 1.4, and $|\eta| < 1.5$. The other set takes advantage of a new track-finding procedure in the barrel section of the L1 muon trigger in the barrel muon track finder (BMTF) algorithm. This novel technique enables the reconstruction of L1 muon candidates and the determination of their p_T without using a beamspot constraint [16]. These triggers require $p_T > 15$ and 7 GeV for the leading and subleading L1 muon, respectively. The minimum p_T thresholds are lowered to 6 and 4 GeV when $d_0 > 25$ cm for both L1 muons.

In the overlap and endcap sections of the L1 muon trigger, the determination of muon p_T without a beamspot constraint was implemented for 2024 data taking, extending the coverage of L1 displaced muon triggers to the full η range of CMS.

In the overlap section, to allow for the beamspot unconstrained p_T measurement, the overlap muon track finder (OMTF) algorithm first extrapolates the direction of a reference hit selected from hits of two innermost muon stations to other muon stations. Then, the differences between ϕ values of hits in these other stations and extrapolated ϕ values from the reference hit are used by the OMTF algorithm [106] to reconstruct the muons and determine both beamspot-constrained and -unconstrained p_T .

In the endcap section, the endcap muon track finder (EMTF) algorithm uses a neural network (NN) to determine beamspot unconstrained muon p_T and d_0 using 29 features extracted from the muon track. In parallel to the NN, the original EMTF algorithm, based on a boosted decision tree [107], is used to determine the beamspot-constrained p_T .

Figure 39 shows the efficiency of the BMTF, OMTF, and EMTF algorithms for both beamspot-constrained and -unconstrained p_T , using a simulated sample of LLPs that decay into a pair of muons.

At the L1 Global Trigger, an additional displaced dimuon trigger algorithm was introduced in 2024. This trigger uses the outputs of BMTF, OMTF, and EMTF and requires two high-quality muons with beamspot unconstrained $p_T > 6$ GeV. It covers the range $|\eta| < 2$, thus extending the η range of the displaced dimuon triggers used previously.

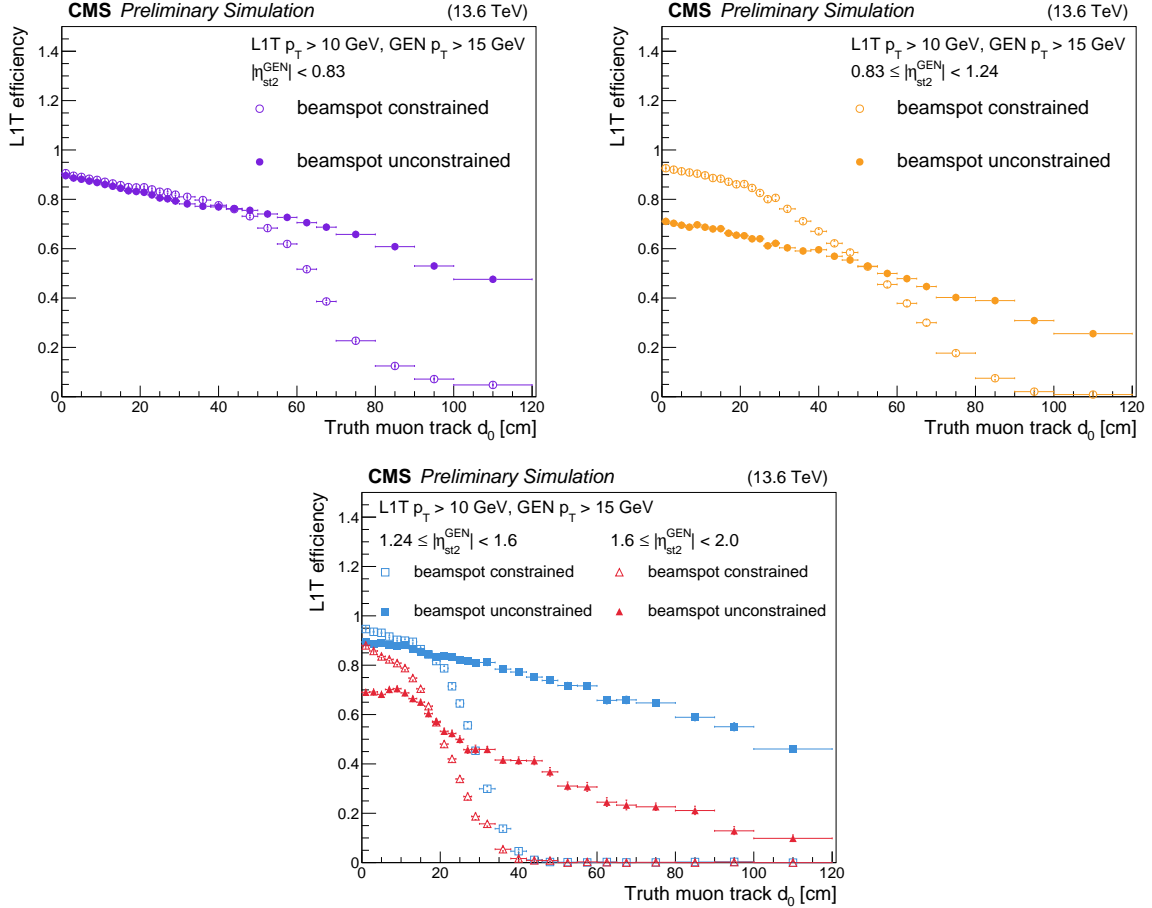


Figure 39: BMTF (upper left), OMTF (upper right), and EMTF (lower) L1T efficiencies for beamspot-constrained and -unconstrained p_T assignment algorithms for L1T $p_T > 10$ GeV with respect to truth-level muon track d_0 , obtained using a sample which produces LLPs that decay to dimuons. The L1T algorithms and data-taking conditions correspond to 2024. A selection on the truth-level muon track $p_T > 15$ GeV is applied to show the performance at the efficiency plateau. The truth-level muon tracks are extrapolated to the second muon station to determine the η_{st2}^{GEN} values that are used in the plot. The solid markers show the new vertex-unconstrained algorithm performance, while the hollow markers show the default beamspot-constrained algorithm performance. In the EMTF plot, the different colors show different $|\eta|$ regions: $1.24 < |\eta_{st2}^{GEN}| < 1.6$ (blue), $1.6 < |\eta_{st2}^{GEN}| < 2.0$ (red).

At the HLT, two new triggers were introduced. The new triggers are characterized by minimum p_T thresholds that depend on the muon d_0 . The first algorithm, labeled “Run 3 (2022, L2)”, attempts to reconstruct the muon candidates at the L3 stage (like global muons), taking advantage of their superior tracker resolution for determining muon d_0 , instead of stopping the online reconstruction at the L2 stage (muon system alone), as was done in the Run 2 trigger. If either of the two L2 muon candidates is reconstructed as a L3 muon with $d_0 < 1$ cm, the event is discarded, since such an L2 muon candidate is likely to originate from the background processes. The resulting trigger operates with minimum L2 muon p_T thresholds at 10 GeV and

significantly improves the signal efficiency while contributing around 10 Hz of additional HLT rate.

The second new HLT algorithm, labeled “Run 3 (2022, L3)”, relies entirely on the online L3 muon reconstruction. A moderate minimum threshold on the impact parameter of each muon, $d_0 > 0.01$ cm, allows the minimum p_T thresholds to be fairly low: 16 (10) GeV for the leading (subleading) L3 muon. The resulting trigger greatly improves the signal efficiency at smaller $c\tau$ compared to the Run 3 (2022, L2) trigger, while adding only another 10 Hz of HLT rate.

An equivalent trigger to the Run 3 (2022, L3), labeled “Run 3 (2022, L3 dTks)”, is also included in the trigger menu. This new trigger uses a newly developed L3 muon reconstruction algorithm that closely resembles the algorithm used offline. The main improvements are the inclusion of the displaced tracker muons and the optimization of the seeding stage of the muon trajectory building, which is made more efficient for displaced muons by removing the d_0 requirement. The new trigger, currently in the commissioning phase, has the same signal efficiency but with about 15% less rate than the Run 3 (2022, L3) trigger.

The combined L1T+HLT efficiency of the various displaced-dimuon triggers and their logical OR as a function of $c\tau$ is shown in Fig. 40 for simulated HAHM signal events, where the dark Higgs boson (H_D) mixes with the SM Higgs boson (H) decays to a pair of long-lived dark photons (Z_D), each with a mass $m_{Z_D} = 20$ GeV, as shown in Fig. 5. The addition of new triggers at the L1T and HLT improves the efficiency for Z_D with $m_{Z_D} > 10$ GeV and $c\tau \gtrsim 0.1$ cm by a factor of 2 to 4, depending on $c\tau$ and mass.




Figure 40: L1T+HLT efficiencies of the various displaced-dimuon triggers and their logical OR as a function of $c\tau$ for the HAHM signal events with $m_H = 125$ GeV and $m_{Z_D} = 20$ GeV, in 2022 conditions. The efficiency is defined as the fraction of simulated events that satisfy the detector acceptance and the requirements of the following sets of triggers: the Run 2 (2018) triggers (dashed black); the Run 3 (2022, L3) triggers (blue); the Run 3 (2022, L2) triggers (red); and the logical OR of all these triggers (Run 3 (2022), solid black). The lower panel shows the ratio of the overall Run 3 (2022) efficiency to the Run 2 (2018) efficiency. Figure adapted from Ref. [105].

An efficiency measurement is performed using $J/\psi \rightarrow \mu\mu$ events recorded using jets and p_T^{miss} triggers with 2022 and 2023 data. This data set corresponds to an integrated luminosity of 64.3 fb^{-1} , which is a little more than is used in other sections of this note focused on 2022 and 2023 data, since in this case only the subdetectors needed for muon reconstruction are

required to provide good-quality data. The efficiencies are defined as the subset of these events that also pass the Run 3 (2022, L3) path. The measurement is compared with the efficiency in simulation, defined as the fraction of events that pass the same trigger, obtained by combining various simulated samples of B hadrons that decay into $J/\psi \rightarrow \mu\mu$. Figure 41 shows trigger efficiencies as a function of the minimum p_T of the two muons in the event $\min(p_T)$ (upper left), the maximum p_T of the two muons $\max(p_T)$ (upper right), and the minimum d_0 of the two muons $\min(d_0)$ (lower). For each measurement, the other two variables are required to be at the trigger efficiency plateau. The efficiency plateau, where the efficiencies are close to 65% in data and 70% in simulation, is reached for $\min(p_T) > 12$ GeV, $\max(p_T) > 18$ GeV, and $\min(d_0) > 0.015$ cm. The main loss of efficiency at the plateau is caused by double-muon L1 triggers when selecting two muons with small spatial separation. This inefficiency is particularly prominent in $J/\psi \rightarrow \mu\mu$ events, as they typically have $0.05 < \Delta R(\mu\mu) < 0.15$. There is some remaining efficiency of around 15% below the trigger thresholds. This is because of differences between the online and offline L3 muon reconstructions, which can provide different estimations of the p_T and d_0 parameters.

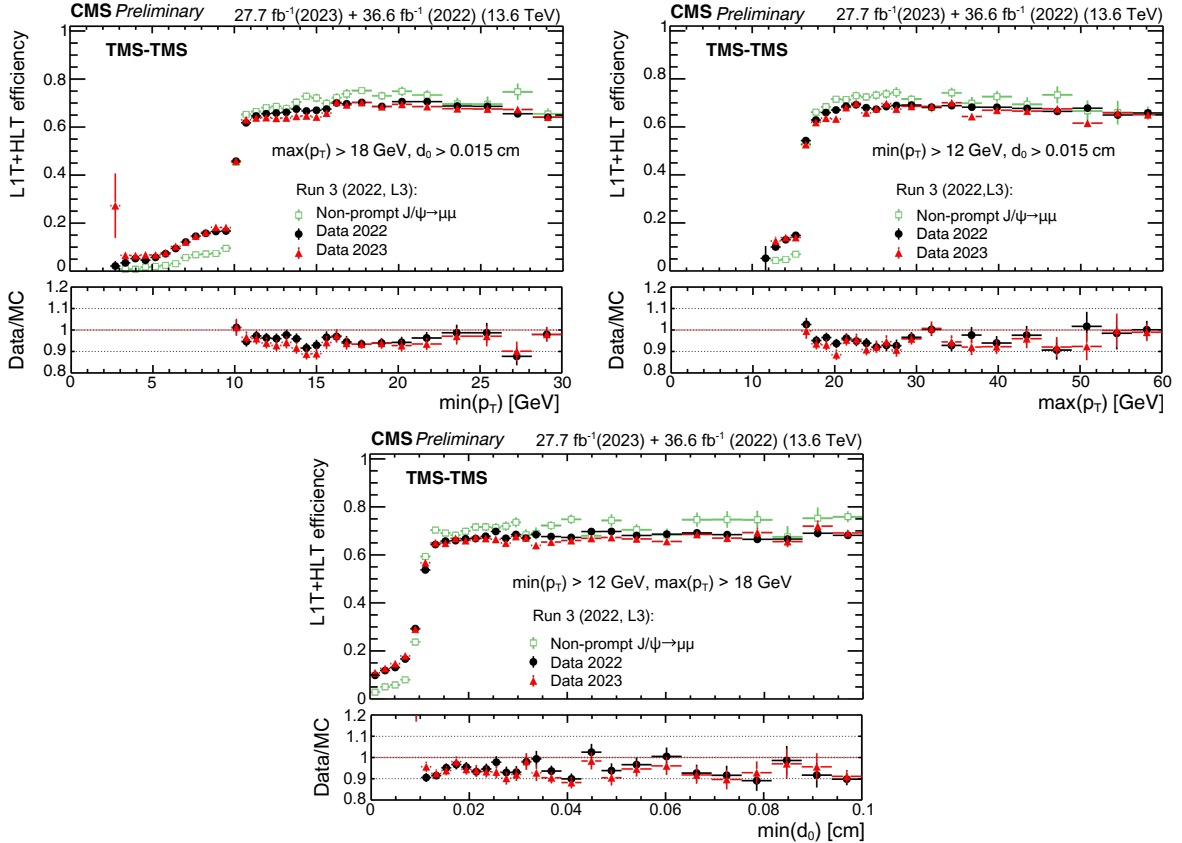


Figure 41: L1T+HLT efficiency of the Run 3 (2022, L3) triggers in 2022 data (black), 2023 data (red), and simulation (green) as a function of $\min(p_T)$ (upper left), $\max(p_T)$ (upper right), and $\min(d_0)$ (lower) of the two muons forming TMS-TMS dimuons in events enriched in $J/\psi \rightarrow \mu\mu$ events. Efficiency in data is the fraction of $J/\psi \rightarrow \mu\mu$ events recorded by the triggers based on the information from jets and p_T^{miss} that also satisfy the requirements of the Run 3 (2022, L3) triggers. It is compared to the efficiency of the Run 3 (2022, L3) triggers in a combination of simulated samples of $J/\psi \rightarrow \mu\mu$ events produced in various B hadron decays. The lower panels show the ratio of the data to simulated events.

We also use $J/\psi \rightarrow \mu\mu$ TMS-TMS dimuon candidates with an invariant mass of $3.0 < m_{\mu\mu} <$

3.2 GeV recorded with the Run 2 (2018) triggers to evaluate the efficiency of the Run 3 (2022, L3) and Run 3 (2022, L3 dTks) triggers. This is shown in Fig. 42 (upper left) for the Run 3 (2022, L3) trigger and split in the 2022 and 2023 years, with an efficiency higher than 90% for dimuons with $\min(d_0) > 0.012$ cm, and in Fig. 42 (upper right) for the combined 2022 and 2023 years and both triggers simultaneously. Figure 42 (lower) shows the invariant mass distribution for TMS-TMS dimuons in events recorded by the Run 2 (2018) triggers, and the subset of events also selected by the Run 3 (2022, L3) or Run 3 (2022, L3 dTks) triggers. This last trigger provides larger signal efficiency and background rejection than the Run 3 (2022, L3) trigger.

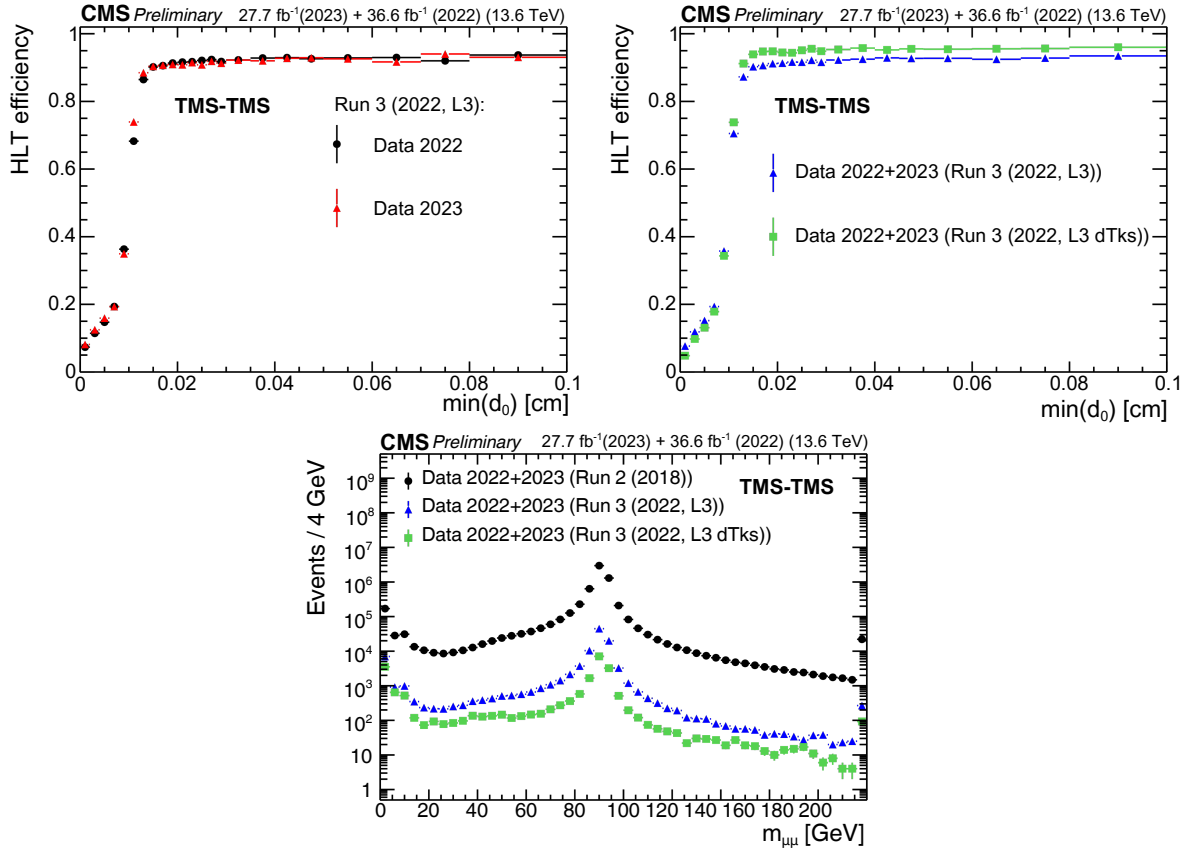


Figure 42: HLT efficiency, defined as the fraction of events recorded by the Run 2 (2018) triggers that also satisfied the requirements of the Run 3 (2022, L3) triggers, as a function of offline-reconstructed $\min(d_0)$ of the two muons forming TMS-TMS dimuons in events enriched in $J/\psi \rightarrow \mu\mu$ (upper left). The black circles represent efficiencies during the 2022 data-taking period, and the red triangles represent the 2023 period. For dimuons with offline $\min(d_0) > 0.012$ cm, the combined efficiency of the L3 muon reconstruction and the online $\min(d_0)$ requirement is larger than 90% in all data-taking periods. The HLT efficiency of the Run 3 (2022, L3) triggers, shown with blue triangles, and the Run 3 (2022, L3 dTks) triggers, shown with green squares, in $J/\psi \rightarrow \mu\mu$ events in the full 2022 and 2023 data set (upper right). Invariant mass distribution for TMS-TMS dimuons in events recorded by the Run 2 (2018) triggers in the combined 2022 and 2023 data set (black circles), and in the subset of events also selected by the Run 3 (2022, L3) trigger (blue triangles) and Run 3 (2022, L3 dTks) trigger (green squares), illustrating the background rejection of the L3 triggers (lower).

We use STA-STA dimuons that have a three-dimensional (3D) opening angle compatible with a cosmic ray muon traversing CMS recorded with the Run 2 (2018) trigger to evaluate the efficiency of the Run 3 (2022, L2) triggers. The efficiency as a function of $\min(d_0)$ is shown in

Fig. 43 (left). It is greater than 95% for dimuons with $\min(d_0)$ up to 350 cm in both 2022 and 2023 years. Figure 43 (right) shows the invariant mass distribution for TMS-TMS dimuons in events recorded by the Run 2 (2018) triggers, and the subset of events also selected by the Run 3 (2022, L2).

In 2024, a few displaced single-muon triggers were added. At the L1T, a single-muon BMTF trigger requires a high-purity muon with beamspot-unconstrained $p_T > 15$ GeV. This trigger seeds two paths at the HLT: one path that requires at least one L2 muon with $p_T > 50$ GeV and rejects events with at least one L3 muon with $d_0 < 1$ cm, and another path that requires at least one L3 muon with $p_T > 30$ GeV and $d_0 > 0.01$ cm.

The total rates of the displaced single and dimuon triggers described here are 67 and 102 Hz, respectively.

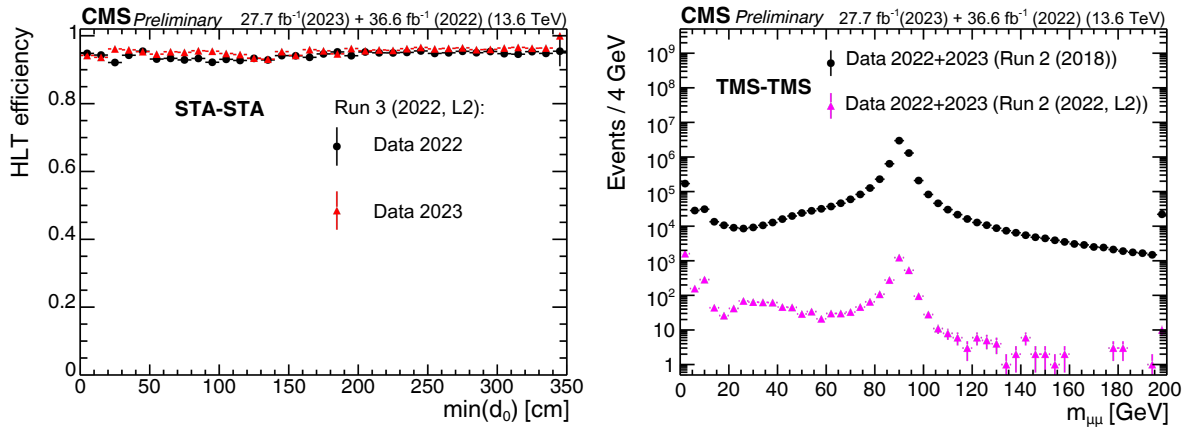


Figure 43: HLT efficiency, defined as the fraction of events recorded by the Run 2 (2018) triggers that also satisfied the requirements of the Run 3 (2022, L2) triggers, as a function of offline-reconstructed $\min(d_0)$ of the two muons forming STA-STA dimuons in events enriched in cosmic ray muons (left). The black circles represent efficiencies during the 2022 data-taking period, and the red triangles represent the 2023 period. For displaced muons, the efficiency of the online $\min(d_0)$ requirement is larger than 95% in all data-taking periods. The invariant mass distribution for TMS-TMS dimuons in events recorded by the Run 2 (2018) triggers in the combined 2022 and 2023 data set (black circles), and in the subset of events also selected by the Run 3 (2022, L2) triggers (pink triangles), illustrating the background rejection of the Run 3 (2022, L2) triggers (right).

6.3.2 Double displaced L3 muon triggers

Another displaced-dimuon trigger has also taken data since Run 2. This trigger is designed for the displaced leptons analysis, described in Ref. [108], where the two leptons are not required to come from a common displaced vertex. This signature is shown in Fig. 44 and can arise for example in an RPV SUSY model with long-lived top squarks that decay into displaced quarks and leptons [109], a GMSB model with long-lived sleptons that decay into a gravitino and a lepton [110], and a model with BSM Higgs bosons that decay into two long-lived scalars, which each decay into two displaced leptons [111]. At the HLT, this trigger requires at least two L3 muons, each with $p_T > 43$ GeV, whose track fit is not updated at the PV, which provides the trigger with higher acceptance of displaced muons. This HLT path is seeded by L1 dimuon triggers. The total rate of these HLT paths is about 2 Hz.

Figure 45 shows the efficiency of this trigger in HAHM signal events, as a function of the $\min(p_T)$ of the two global or tracker muons in the event. The HAHM signal is used here,




Figure 44: Diagram of a simulated signal event in the inclusive displaced-leptons search, from a transverse view of the interaction point, in the analysis presented in Ref. [108]. The black arrows indicate the lepton d_0 vectors. Figure taken from Ref. [112].

rather than one of the signals mentioned above that are targeted by this trigger, in order to more readily compare the efficiency to that of other displaced muon triggers. The turn-on is shown in muon p_T , illustrating that the double displaced L3 muon trigger is fully efficient for offline muons with $\min(p_T) > 48$ GeV. Furthermore, this trigger has 100% efficiency for all displacements. The HLT path does not exhibit a turn-on at small $\min(d_0)$ values because no displacement requirement is made in this trigger.




Figure 45: L1T+HLT efficiency of the double displaced L3 muon trigger as a function of $\min(p_T)$ of the two global or tracker muons in the event. The efficiency is plotted for HAHM signal events in 2022 conditions with $m_{Z_D} = 50$ GeV and $\epsilon = 4 \times 10^{-9}$ (black triangles), $m_{Z_D} = 60$ GeV and $\epsilon = 2 \times 10^{-9}$ (red triangles), and $m_H = 125$ GeV in both cases. The events are required to have at least two good global or tracker muons with $p_T > 23$ GeV.

6.3.3 Displaced L3 muon+photon triggers

Several cross-object triggers, which select at least one photon and at least one displaced muon at the HLT, have been used in CMS since Run 2. These triggers are designed for an analysis, described in Ref. [108], that searches for two LLPs that are produced in pp collisions and decay in the detector, one of which to a displaced muon and other decay products, the other of which to a displaced electron and other decay products. This signature was shown above in Fig. 44 and can arise in the same theoretical scenarios discussed in Section 6.3.2. The displaced electron and muon would not necessarily come from the same displaced vertex. In this way, the HLT photon is a proxy for a displaced electron. To be sensitive to larger displacements and limit bias towards lower displacements, the displaced electron track is not reconstructed at the HLT, but rather only in the offline analysis.

At the HLT, the first muon-photon cross-object trigger requires an L3 muon with $p_T > 43$ GeV whose track fit is not updated at the PV to minimize the constraints to the interaction point and maximize the efficiency for displaced muons. This trigger also requires a photon at the HLT with $p_T > 43$ GeV and a set of loose, calorimeter-based identification selections. In addition, there is a second muon-photon cross-trigger available at the HLT, which makes similar requirements except it lowers the minimum p_T requirements on both the HLT photon and the muon to 38 GeV, and to compensate for the rate increase, the muon is required to have $d_0 > 1$ cm. In this way, the efficiency to trigger on lower p_T signals is increased, while paying a modest price in displacement. Both muon-photon cross-object HLT paths are seeded by muon-e/ γ L1 triggers.

The total rate of these HLT paths is about 5 Hz. The performance of the displaced L3 muon leg of this trigger can be seen from the performance of double-displaced L3 muon triggers, as described in Section 6.3.2 and shown in Fig. 45. The performance of the photon leg of the trigger, which is a standard photon reconstructed at the HLT, can be seen in Ref. [19].

6.3.4 Dimuon scouting triggers

Dimuon scouting triggers, which are fully documented in Ref. [33], have been available since 2017. Their purpose is to substantially lower the muon minimum p_T thresholds compared to standard triggers. In particular, the dimuon scouting triggers require both muons to have $p_T > 3$ GeV at the HLT, as opposed to standard dimuon triggers, which historically required at best minimum p_T thresholds of 17 and 8 GeV for the leading and subleading p_T muons, respectively. Since Run 3, an inclusive dimuon trigger strategy has been implemented to collect events in the low-mass region (below the Y peak) using the parking data, with p_T thresholds as low as those in scouting. However, dedicated scouting muon reconstruction still makes scouting the optimal strategy for displaced muon signatures compared to parking. This is because the dimuon scouting triggers use HLT reconstruction algorithms that lack any association between the muons and the PV. This enables scouting searches for resonances that have nonzero displacement from the PV, such as the search described in Ref. [113]. This search targets a displaced dimuon signature similar to the one described in Section 6.3.1, but it targets signals with lower dimuon masses, such as a scenario with an SM-like Higgs boson that decays to four leptons via two intermediate dark photons or through the Higgs portal via a dark Higgs boson [70]. Figure 38, which is shown above in Section 6.3.1, illustrates an example signal event that is also applicable for the search that uses dimuon scouting triggers.

In Run 2, the dimuon scouting HLT paths required that muon tracks deposit energy in at least two layers of the pixel tracker detector. This requirement limited the maximum muon d_0 to about 11 cm. At the beginning of Run 3, this pixel hit requirement was removed from the scouting muon reconstruction algorithm, which increased the efficiency to trigger on displaced

muons, as they may be sufficiently displaced to not create hits in the pixel detector. This change occurred at the same time as the L1T BMTF improvement mentioned in Section 6.3.1, which also improved the displaced-muon trigger performance.

Figure 46 shows the distribution of the dimuon vertex transverse displacement L_{xy} , for events that contain at least one pair of opposite-sign muons associated with a selected secondary vertex. Events are collected with dimuon displacements up to about 100 cm, corresponding to the end of the sensitive region of the tracker. At the positions of the pixel layers, with radii of 29, 68, 109, and 160 mm, photons undergoing conversion processes in the material lead to peaks in the L_{xy} distribution. These peaks are less pronounced in the Run 3 distribution because of the removal of the pixel-hit requirement, which leads to higher efficiency, but also lower purity, if no additional analysis-specific quality criteria are required, as is the case here.




Figure 46: Comparison between the L_{xy} distribution for dimuon scouting events, for Run 2 (orange) and Run 3 (blue) events in data that contain dimuon pairs with a common displaced vertex and a minimal selection on the vertex quality. The dashed vertical lines, placed at radii of 29, 68, 109, and 160 mm, correspond to the positions of the pixel layers where photons undergo conversion processes in the material, causing the observed peaks in the L_{xy} distribution. Figure taken from Ref. [33].

A combination of L1T algorithms requiring one of two options is employed. These options are to require either (a), at least two muons with minimum p_T thresholds of 15 and 7 GeV for the highest- p_T and lowest- p_T muon, respectively, or (b), two opposite-sign muons with p_T thresholds as low as 0 GeV and an additional selection on the invariant mass or angular separation (for better handling the trigger rate). The proportional rate of these triggers at the L1T amounts to about 10 kHz at the highest luminosity for an average PU of 60. Events selected by at least one of the L1 seeds are then reconstructed at HLT with a minimum requirement of 3 GeV, reducing the HLT rate of the dimuon scouting trigger to about 4.2 kHz.

The efficiency of the dimuon scouting trigger is shown in 2024 data in Fig. 47 and in simulated HAHM signal events for two illustrative masses of $m_{Z_D} = 2.5$ and 14 GeV in Figs. 48 and 49. Figure 47 shows that the dimuon scouting trigger efficiency is driven by the L1T, which requires high-quality muons. Figures 48 and 49 demonstrate that the dimuon scouting triggers are highly efficient for HAHM events, especially for muons coming from decays of dark photons with sufficiently large m_{Z_D} and with a less boosted decay (ensuring good angular muon separation), and where also a significant fraction of the muons are reconstructed even for the

longest considered lifetimes. Good efficiency is maintained for L_{xy} values up to about 40 cm.




Figure 47: L1T+HLT efficiency of the dimuon scouting trigger as a function of the subleading muon p_T , for 2024 data. The efficiency of the L1T dimuon seeds (pink squares) and the HLT dimuon scouting trigger with the vertex-unconstrained reconstruction algorithm (blue triangles) is shown. At least two scouting muons with $\chi^2/\text{dof} < 3$ and $\Delta R > 0.1$ are required.

Figure 50 shows the p_T resolution of scouting muons with respect to offline muons, as a function of the scouting muon p_T , for 2024 data events. The resolution is shown for two different scouting muon reconstruction algorithms: one that is constrained to the primary vertex, and one that is unconstrained. The vertex-constrained algorithm is appropriate for prompt muons, while the vertex-unconstrained algorithm provides higher efficiency for displaced muons. The vertex-unconstrained algorithm provides p_T resolutions ranging from 0.5 to 1% for $p_T < 50$ GeV and increasing to 2% at 100 GeV, for both barrel and endcap muons. The vertex-constrained algorithm provides very good p_T resolutions for both the barrel and endcaps, that is, less than 0.5% for all p_T values in the barrel and between 0.5 and 1% in the endcaps.

6.3.5 Muon detector showers triggers with the CSCs

During pp collisions, LLPs could be produced at the CMS interaction point and, if they have a sufficiently large lifetime, travel to the muon system before decaying, creating an MDS signature. The defining feature of LLP decays within the CMS muon system is the multiplicity of hits created by particle showers. This starkly contrasts what a typical muon would deposit in the detector, which are several hits in each detection layer that form a trajectory pointing to the interaction point. Some Run 2 analyses have used this MDS signature offline, as described in Refs. [114, 115], relying on other objects for the trigger strategy. To improve the CMS sensitivity to this striking signature, several MDS triggers have been created in time for Run 3 data taking.

The overall strategy at the L1T is to first identify the LLP showers by counting the number of hits at individual CSC chambers, and to then aggregate and condense the presence of showers from the chambers through the rest of the backend electronics. Hits per bunch crossing




Figure 48: L1T+HLT efficiency of the dimuon scouting trigger as a function of the generator-level L_{xy} , for HAHM signal events in 2024 conditions. The efficiency is shown for $m_{Z_D} = 14$ GeV and $c\tau = 100$ mm (pink squares) and $m_{Z_D} = 2.5$ GeV and $c\tau = 100$ mm (blue triangles). The muons are required to have $p_T > 15$ GeV and $|\eta| < 2.4$ at the generator level.

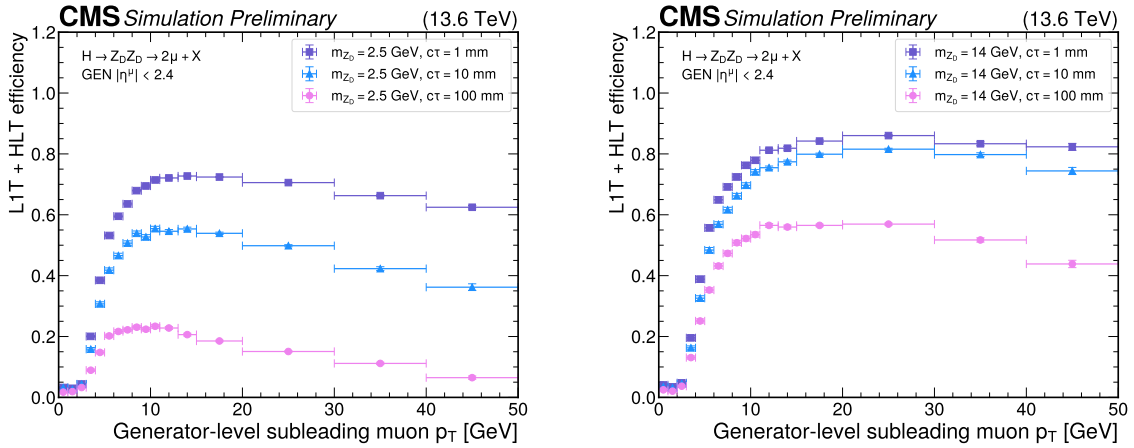


Figure 49: L1T+HLT efficiency of the dimuon scouting trigger as a function of the generator-level subleading muon p_T , for HAHM signal events in 2024 conditions. The efficiency is shown for Z_D masses of 2.5 (left) and 14 GeV (right), and $c\tau$ values of 1 (purple squares), 10 (blue triangles), and 100 mm (pink circles). The muons are required to have $|\eta| < 2.4$ at the generator level.

are counted in each CSC chamber on the CSC local trigger, which has both anode-wire and cathode-strip readout. Optimized thresholds are set on each sector individually to evaluate whether the found number of hits is above or below such thresholds. The larger available resources of the optical trigger motherboard allow the trigger to detect the coincidence of anode and cathode counts above thresholds for chambers in ring 1, while an anode-only logic is




Figure 50: The p_T resolution of scouting muons with respect to offline muons, as a function of the scouting muon p_T , for 2024 data events. The dimuon ΔR is required to be greater than 0.2, and the scouting muon p_T is required to be greater than 3 GeV. The resolution is shown for muons in the barrel (blue filled points) and the endcaps (purple filled triangles) that are reconstructed with the unconstrained vertex reconstruction algorithm, as well as for muons in the barrel (red filled squares) and the endcaps (orange unfilled squares) that are reconstructed with the constrained vertex reconstruction algorithm. The figure is made using a special monitoring data set that collects events triggered by a mixture of HLT paths (both scouting and standard triggers) with a very high prescale, in which all information about the muon objects stored from the offline and scouting reconstruction.

set for chambers in rings 2 and 3. This information is encoded in several bits called “loose”, “nominal”, and “tight”, with different thresholds for each, and then transferred to the EMTF sector processors. Likewise, the EMTF transmits information about the existence or absence of any chamber passing the counting criteria per sector to the global muon trigger. Finally, the global muon trigger transfers the information to the global trigger, and the L1 physics seeds (One-Nominal, One-Tight, and Two-Loose) are filled within the L1 menu. In the case of the Two-Loose strategy, which targets pair-produced LLPs, the global muon trigger evaluates if there are at least two different sectors with any chamber above loose thresholds, which vary depending on the chamber. A display of one CMS event triggered in 2022 by this logic is shown in Fig. 51.

The efficiency of the trigger to select LLPs is evaluated with a sample of $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ events, the Feynman diagram for which is shown in Fig. 4. Signal efficiencies are shown in Tables 4 and 5 for simulated events with one and two LLPs decaying in the CSC detector acceptance, respectively. In the latter case, it can be observed that the presence of the Two-Loose L1 seed yields up to a 10% increase in the signal efficiency.

At the HLT, hits from all the CSC chambers or stations are available and can be reconstructed as cluster objects. This allows the HLT to capture the properties of signal showers across multiple CSC chambers, thus rejecting more backgrounds. The CSC 3D spatial points (hits) are clustered using the Cambridge-Aachen (CA) algorithm [116, 117] in the η - ϕ space with a distance parameter of 0.4. A cluster object is defined to have at least 50 hits, a threshold approximately

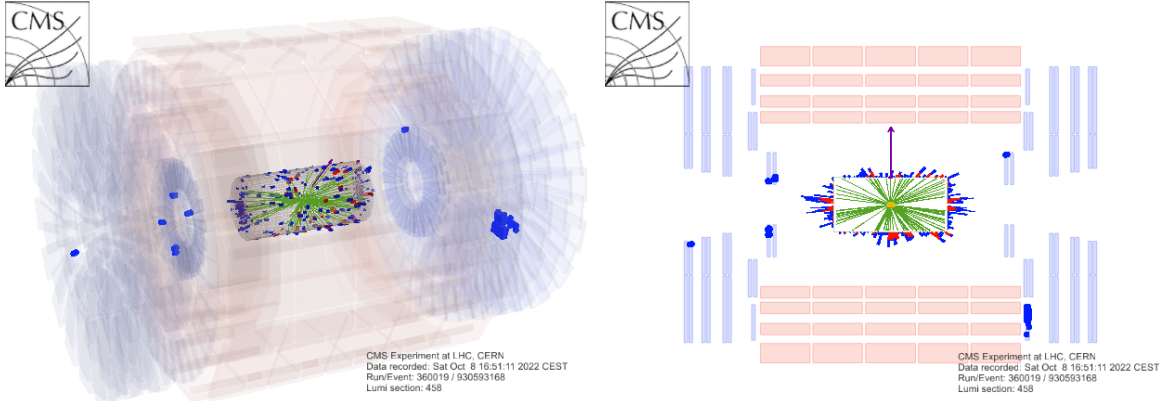


Figure 51: Event display of a collision triggered by the CSC MDS trigger. The CSC reconstructed hits are represented by blue dots in the muon end-cap region. This event features a CSC cluster of 210 hits in the ME1/3 ring. The event was recorded on October 8th, 2022.

Table 4: The L1T CSC MDS trigger efficiency ϵ in $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ events, for different choices of m_H , m_S , and $c\tau_S$, for simulated events with at least one LLP decaying in the CSC detector acceptance. The ϵ_{Acc} refers to the percentage of events within the CSC detector acceptance. The $\epsilon_{\text{One-Nominal}}$ ($\epsilon_{\text{One-Tight}}$) requires the number of hits in at least one CSC chamber to pass the nominal (tight) thresholds.

m_H [GeV]	m_S [GeV]	$c\tau_S$ [mm]	ϵ_{Acc}	$\epsilon_{\text{One-Nominal}}$	$\epsilon_{\text{One-Tight}}$
125	12	900	16.2%	43.1%	39.2%
125	25	1500	15.9%	43.7%	40.0%
250	60	10000	10.3%	55.5%	52.1%
250	120	10000	15.4%	29.9%	27.7%
350	160	10000	17.9%	39.1%	36.7%

two times higher than a typical muon would create in the muon system.

The additional cluster properties available at HLT allow us to design various trigger paths to target different LLP production modes, as shown in Table 2. The single MDS path, which was introduced in time for 2022 data taking, is the most generic path, whereas the double MDS path (introduced in 2023) targets pair-produced LLP models more efficiently because of its looser cluster selections. For LLPs produced together with a lepton or photon, additional cross-trigger paths were introduced in time for 2024 data taking, where the trigger thresholds for both the cluster and the lepton/photon selections are lowered.

The L1T efficiency as a function of the reconstructed CA cluster for one of the CSC rings and L1T rates for One-Nominal and Two-Loose as a function of PU are shown in Fig. 52. (The efficiencies for other CSC rings are shown in Ref. [76].) To evaluate the efficiency of individual CSC rings, CSC clusters created by muon bremsstrahlung-induced, electromagnetic showers are used as a proxy for LLP signal clusters in data. The efficiency is defined as the fraction of events that pass the CSC MDS One-Nominal L1 seed, given that an uncorrelated trigger selects the data and that a muon is matched to a CSC cluster. Because the L1 thresholds of different CSC rings are different, the efficiency is measured separately for each CSC ring. Thus, the centroid position of the cluster is restricted to be within a single CSC ring (e.g., ME2/2), and at least 90% of the hits in the clusters are contained in the same CSC ring.

Figure 52 also shows that the L1T rate of both the One-Nominal and Two-Loose seeds are

Table 5: The L1T CSC MDS trigger efficiency ϵ in $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ events, for different choices of m_H , m_S , and $c\tau_S$ for simulated events with two LLPs decaying in the CSC detector acceptance. The ϵ_{Acc} refers to the percentage of events within the CSC detector acceptance. The $\epsilon_{\text{One-Nominal}}$ ($\epsilon_{\text{Two-Loose}}$) requires having the number of hits in at least one (two) CSC chamber(s) passing the nominal (loose) thresholds.

m_H [GeV]	m_S [GeV]	$c\tau_S$ [mm]	N_{acc}	$\epsilon_{\text{One-Nominal}}$	$\epsilon_{\text{One-Nominal OR Two-Loose}}$
125	12	900	0.7%	69.6%	75.0%
125	25	1500	0.7%	75.1%	79.7%
250	60	10000	0.3%	74.0%	78.2%
250	120	10000	1.0%	44.4%	49.6%
350	160	10000	1.1%	63.4%	69.6%

proportional to the average PU. The cluster selections at the HLT further reduce the L1T rates by a factor of 170, such that the single and double MDS paths have a combined total rate of 14 Hz. The CSC MDS cross triggers provide an additional 14 Hz total rate at the HLT at the same instantaneous luminosity.

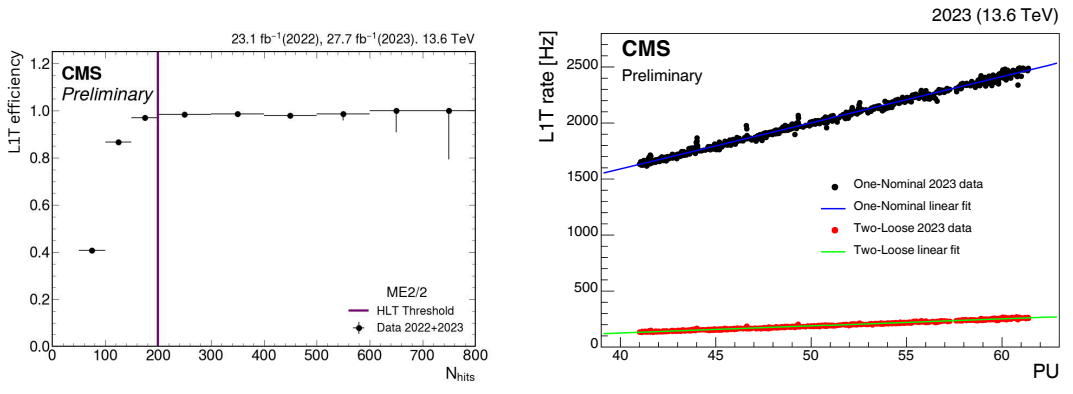


Figure 52: L1T efficiency of the One-Nominal CSC MDS trigger as a function of CSC hit cluster size, measured with muon bremsstrahlung-induced electromagnetic showers using data taken in 2022 and 2023 (left). L1T rate as a function of PU in a CMS run for the One-Nominal and Two-Loose seed for a 2023 data-taking run (right). The rate dependence on PU is extracted by using a linear fit.

6.3.6 Muon detector showers triggers with the DTs

Similarly to the MDS signature in the CSCs described in Section 6.3.5, an LLP decaying in the DTs in the muon system can produce showers from the resulting decay products. These showers can be identified using clusters of hits in the DTs, using this subsystem as a sampling calorimeter. The DTs provide complementary coverage to the CSC MDS triggers described in Section 6.3.5. In Run 2, CMS performed a search for LLPs using the DT cluster showers and provided competitive sensitivity limits to various models, including Higgs boson to LLP decays and first LHC limits for dark sector coupling to the Higgs boson [115].

In Run 3, dedicated triggers have been introduced that target a cluster of hits in the DTs opposite the jet coming from initial or final state radiation. Figure 53 provides a visual representation of this signature. This trigger enables a more model-independent search for LLPs decaying in the DTs. The trigger uses p_T^{miss} with a minimum threshold of 150 GeV in the L1 seed, as implementing a dedicated MDS trigger at the L1T in DTs was impossible because of the lack of available bandwidth. At the HLT, DT clusters are reconstructed similarly to the CSCs, and a

minimum size of 50 hits is required. Note that the hits from the first DT station, MB1, are excluded in the clusters, as it has the least amount of material upstream towards the interaction point and thus higher backgrounds. Furthermore, another HLT path was implemented that requires the One-Nominal MDS CSC cluster at the L1T and at least one DT cluster at the HLT. This path recovers some efficiency for the case that multiple clusters are created in the transition region between the barrel and the endcaps. The total trigger rate at the HLT of the DT MDS paths is 9 Hz.




Figure 53: Diagram of the MDS signature in the DTs. The LLP decays in the DTs, opposite a jet from initial state radiation.

Figure 54 shows good trigger efficiencies on simulated signal events with respect to offline reconstruction quantities. The cluster used at the HLT is identical to the offline one, which leads to close to 100% efficiency when there are at least 50 hits in the cluster. Less than 50 hits are not allowed by the trigger algorithm, and so there are no measurement points in Fig. 54 right below 50 in the cluster size.

6.4 No-BPTX algorithms

The No-BPTX trigger algorithms are unique in that they are active only when no proton bunches are colliding in the detector. They are designed to trigger on LLPs that have especially long lifetimes, are slow-moving, and could come to a stop in the detector (if the material is dense enough) and then decay sometime later. These triggers collect data when the beam pickup timing device (BPTX) does not detect the proton beam, and hence they are called the “No-BPTX” triggers.

Each LHC beam consists of proton bunches arranged into an irregular pattern of “trains” [118]. Within a train, the proton bunches are nominally spaced 25 ns apart, with larger spacings between trains to account for the needs of the injection process. In an LHC orbit, there are 3564 bunch slots (BXs), which are 25 ns long. Each BX could be filled with proton bunches, which usually occupy the first 2.5 ns of the BX, or could be empty. The trains may be spaced such




Figure 54: HLT efficiency of the DT MDS triggers as a function of p_T^{miss} (left) and cluster size (right), for $H \rightarrow SS \rightarrow 4b$ events with $m_H = 125$ GeV, $m_S = 40$ GeV, and $c\tau_S = 1$ m, in 2023 conditions. Events are required to have at least one cluster with more than 50 hits (left) and $p_T^{\text{miss}} > 250$ GeV (right).

that multiple empty BXs could be between filled BXs. The maximum occupancy of the LHC in Run 3 was about 2400 colliding bunches, depending on the year. At the end of each LHC orbit, there is a $3.15 \mu\text{s}$ long “abort gap” which is kept free of proton bunches to avoid losses during the rise time of the LHC beam dump kickers.

To search for LLP decays during these empty BXs, the main No-BPTX triggers select events at least two BXs away from any proton bunches. Thus, these triggers are live only during these specific time windows. This distance of two BXs is chosen so that we maximize the search time window while suppressing most of the events from secondary pp interactions and from “beam halo”, which are mostly muons traveling outside the LHC beam that are produced by LHC beam-collimator scattering. As a result, these triggers collect most of their data during the LHC abort gap. The No-BPTX algorithm is present in dedicated L1 seeds with jet and muon variants. Higher p_T thresholds for the jets and muons are employed at the HLT, to suppress backgrounds and lower the rate to a few Hz. We will now discuss the jet and muon No-BPTX triggers in more detail below, in Sections 6.4.1 and 6.4.2, respectively.

6.4.1 Jet No-BPTX triggers

The jet No-BPTX triggers have been used since Run 1 to trigger on hadronic decays of LLPs that become stopped in the dense inactive material in the CMS calorimeters [119, 120]. Such stopped LLPs arise in BSM scenarios with long-lived gluinos [87, 89] or top squarks [121–123], for example. This signature is illustrated in Fig. 55.

The jet No-BPTX triggers are designed to trigger on particularly out-of-time calo jets at the HLT with energy greater than 60 GeV and $|\eta| < 3$ that are at least two BXs away from pp collisions. A requirement is made on the total energy of the jets instead of the transverse energy because the stopped particles will decay isotropically, and not necessarily predominantly in the transverse plane. These triggers primarily select events with beam halo and noise in the calorimeters; thus, their rates are largely independent of PU. The total rate of these HLT paths




Figure 55: Diagram of a stopped particle event, which can be selected by the jet No-BPTX triggers. The dotted black arrow indicates the LLP, such as a gluino, that travels through the detector before coming to a stop in the dense material of the HCAL. After some time, this stopped particle decays hadronically, producing a significant energy deposit in the HCAL.

is less than 2 Hz at the maximum number of colliding proton bunches.

6.4.2 Muon No-BPTX triggers

The muon No-BPTX triggers, like the jet No-BPTX triggers, select events when there are no colliding bunches of protons in the detector, i.e., empty BXs. The muon versions of these triggers select out-of-time L2 muons and have taken data since Run 1. The main trigger selects events at least two BXs away from the pp collision time with at least one muon reconstructed in the muon system with $p_T > 40$ GeV. This trigger primarily selects muons from atmospheric cosmic rays that have produced showers of particles in the CMS detector or the rock above it. Thus, the rate of this trigger depends on the number of filled bunch crossings but is largely constant as a function of PU.

The rate of the main muon No-BPTX HLT path as a function of the number of colliding bunches is shown in Fig. 56. The rate decreases linearly as the number of proton bunches increases and there are fewer empty BXs when the No-BPTX trigger selects events. The slight rate increase from Run 2 to Run 3 is mostly likely a result of the increase in rate from the backgrounds in the EMTF, and the increase in rate of the LHC backgrounds overall from the increase in center-of-mass energy. The total rate of these HLT paths is about 7 Hz at the maximum number of colliding proton bunches.

The muon No-BPTX triggers are used to trigger on the muonic decays of stopped LLPs [120], that arise, for example, in models with long-lived gluinos [87, 89] or multiply charged particles [124]. This signature of a stopped particle decaying into two back-to-back muons is illustrated in Fig. 57 and was probed in the latest stopped particle analysis [120]. It should be noted that the muon No-BPTX triggers are also performant for more signatures than this, as the triggers simply require at least one out-of-time muon.

7 Long-lived particle trigger acceptances

In this section, we show the trigger acceptance, as a function of LLP $c\tau$ or decay position, for the different types of LLP triggers described in Section 6, for a given LLP signal benchmark model described in Section 5. In this way, the diversity and complementarity among the different trigger strategies are demonstrated. The trigger acceptances for models that predict neutral LLPs are shown first, for hadronic and leptonic signatures, followed by the acceptances for models that predict charged LLPs.




Figure 56: Rate of the main muon No-BPTX HLT path as a function of the number of colliding bunches, for different years in Run 2 and Run 3.

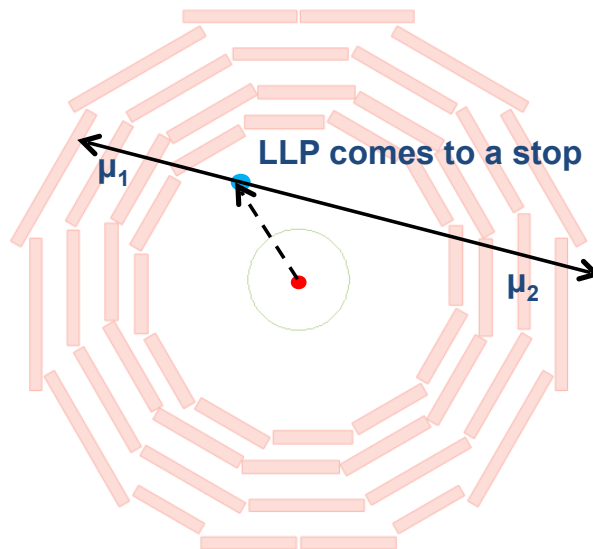


Figure 57: Diagram of a stopped particle event, which can be selected by the muon No-BPTX triggers. The dotted black arrow indicates the LLP that travels through the detector before coming to a stop in the iron yoke in the muon system. After some time, this stopped particle decays to two back-to-back muons.

The trigger performance of different LLP triggers is shown here for the benchmark $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ process, the Feynman diagram for which was shown above in Fig. 4.

Figure 58 shows the trigger acceptance for the CSC and DT MDS triggers as a function of LLP $c\tau$. Figure 58 left illustrates the improvement of the dedicated CSC MDS triggers in Run 3 with respect to the Run 2 trigger strategy, which relies on triggering on p_T^{miss} . Figure 58 right illustrates the improved performance of the DT MDS triggers in Run 3 over that of the Run 2 trigger strategy, namely, triggering on p_T^{miss} at the L1 and HLT. Since the Run 3 MDS triggers

in the DTs still rely on p_T^{miss} at the L1T because of rate constraints, the improvement is less significant than that of the CSC trigger paths.

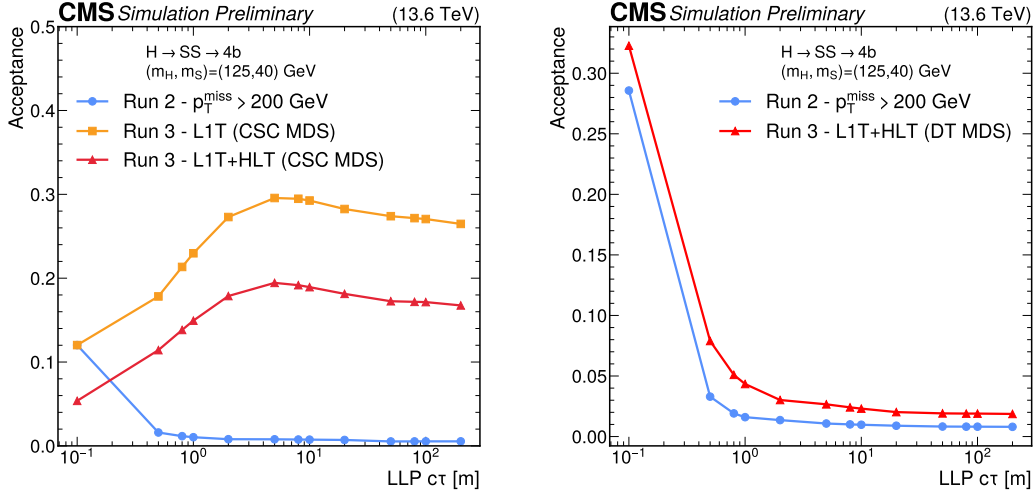


Figure 58: Comparison of the acceptance in Run 2 and Run 3 for the CSC (left) and DT (right) MDS triggers at the L1T and HLT as a function of the LLP lifetime, for $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ events with $m_H = 125$ GeV and $m_S = 40$ GeV, in 2023 conditions. The acceptance is defined as the fraction of events that pass the specified selection, given an LLP decay in the fiducial region of the CSCs (left) or DTs (right). The left plot compares the acceptance of the Run 2 strategy of triggering on p_T^{miss} (blue circles), which was required to be >200 GeV offline, with that of the Run 3 strategy of triggering on the MDS signature in the CSCs, where the L1T (L1T+HLT) acceptance is shown with orange squares (red triangles). The right plot compares the acceptance of Run 2 strategy of triggering on p_T^{miss} (blue circles) with the Run 3 strategy of triggering on the MDS signature in the DTs, where the L1T+HLT acceptance is shown with red triangles.

Figure 59 shows the L1T and HLT acceptances for the CSC MDS trigger as a function of the LLP decay position along the z -axis. The L1T (HLT) acceptance is approximately 40 (10)% for showers in the ME1/1 region because high thresholds are needed in this CSC station to suppress hadronic punchthrough. However, the acceptance increases to around 70 (60)% in the other stations. The shaded regions represent steel within the CSC system, and the structure observed in the acceptance distribution partly reflects variations in steel thickness, which influences shower development. The corresponding acceptance distribution for the DT MDS trigger is shown in Fig. 60. In this case, there is no dedicated L1T seed, so acceptance is shown both for all events and for those that pass the L1T p_T^{miss} trigger. Hadronic punchthrough is of reduced concern for the DT MDS trigger, owing to the lower particle flux in the barrel region and the presence of the solenoid, which results in higher HLT acceptance in the first station compared to the CSC.

Figures 61 and 62 show the trigger acceptance for the CSC and DT MDS triggers, respectively, as a function of the LLP decay radius and z position. These figures illustrate that the MDS trigger program is comprehensive and allows CMS to trigger on the MDS signature throughout the entire muon system.

Figure 63 shows the trigger acceptance at the HLT for various hadronic LLP triggers in $H \rightarrow XX \rightarrow b\bar{b}b\bar{b}$ events, as a function of the LLP decay radius. The figures correspond to different choices of m_H and m_X . The signal acceptance used in this figure is defined as the ratio of




Figure 59: L1T (blue circles) and L1T+HLT (orange squares) acceptance for the CSC MDS trigger as a function of LLP decay positions in the z -direction, for $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ events with $m_H = 350 \text{ GeV}$, $m_S = 80 \text{ GeV}$, and $c\tau_S = 1 \text{ m}$, in 2023 conditions.




Figure 60: HLT (blue circles) and L1T+HLT (orange squares) acceptance for the DT MDS trigger as a function of LLP decay positions in the radial direction, for $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ events with $m_H = 350 \text{ GeV}$, $m_S = 80 \text{ GeV}$, and $c\tau_S = 1 \text{ m}$, in 2023 conditions.




Figure 61: L1T (left) and L1T+HLT (right) acceptance for the CSC MDS trigger as a function of LLP decay positions, for $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ events with $m_H = 350 \text{ GeV}$, $m_S = 80 \text{ GeV}$, and $c\tau_S = 1 \text{ m}$, in 2023 conditions.




Figure 62: HLT acceptance for the DT MDS trigger as a function of LLP decay position, for $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ events with $m_H = 350$ GeV, $m_S = 80$ GeV, and $c\tau_S = 1$ m, in 2023 conditions. The L1T acceptance that is based on p_T^{miss} is not included.

the number of LLPs passing the respective HLT requirements, given that the LLPs decay in the corresponding fiducial region of the detector shown in Fig. 64. For example, the signal acceptance for MDS triggers using CSCs is computed with respect to the decay of LLPs within the fiducial region of the CSCs. The fiducial region used in the tracker-based displaced jet triggers is defined as $R < 300$ cm, $|z| < 560$ cm, and $|\eta| < 2.0$. Here, the fiducial region extends from the tracker into the calorimeters because some of the tracking-based displaced jets algorithms rely on calo jets with a prompt track veto, enabling sensitivity to displaced LLP decays beyond the tracker volume itself. For the ECAL (HCAL)-based delayed jet triggers, the fiducial region is defined as $R < 152$ (300) cm and $|\eta| < 1.48$ (1.26), which corresponds to the barrel region of ECAL and the barrel region of HCAL with full segmentation information. For the MDS triggers using the DT and CSCs, the fiducial region is defined as $200 < R < 800$ cm and $|z| < 661$ cm for the DTs, and $R < 695$ cm, $400 < |z| < 1100$ cm, and $0.9 < |\eta| < 2.4$ for the CSCs, which are extended in front of the muon system to include the LLPs that initiate the hadronic showers in the solenoid and HCAL.

Figure 63 demonstrates the complementarity of the LLP triggers, as the displaced-jet triggers using the tracker have high acceptance for the smallest decay radii, the delayed-jet triggers using ECAL timing provide higher acceptance up until the start of the ECAL, the displaced-jet triggers using the HCAL provide acceptance through the HCAL, and the MDS triggers with the DTs and CSCs provide acceptance inside the muon system.

Figure 65 shows the complementarity in coverage, in the HLT muon p_T - d_0 plane, of the various dimuon triggers that target displaced signatures. The displaced dimuon triggers (Section 6.3.1, covering L2 muon paths in blue and L3 muon paths in red), the double displaced L3 muon triggers (Section 6.3.2, in red), and dimuon scouting triggers (Section 6.3.4, in green) are shown. The coverage of these triggers somewhat overlaps in this plane, but each trigger provides access to additional phase space.

The acceptance of the displaced τ_h trigger in $pp \rightarrow \tilde{\tau}\tilde{\tau} (\tilde{\tau} \rightarrow \tau\tilde{\chi}_1^0)$ events where each τ decays hadronically, is shown in Fig. 66.

8 Summary

The CMS Run 3 (2022–2026) long-lived particle (LLP) trigger program, which includes many dedicated triggers designed to select different LLP signatures, has been presented. These ded-

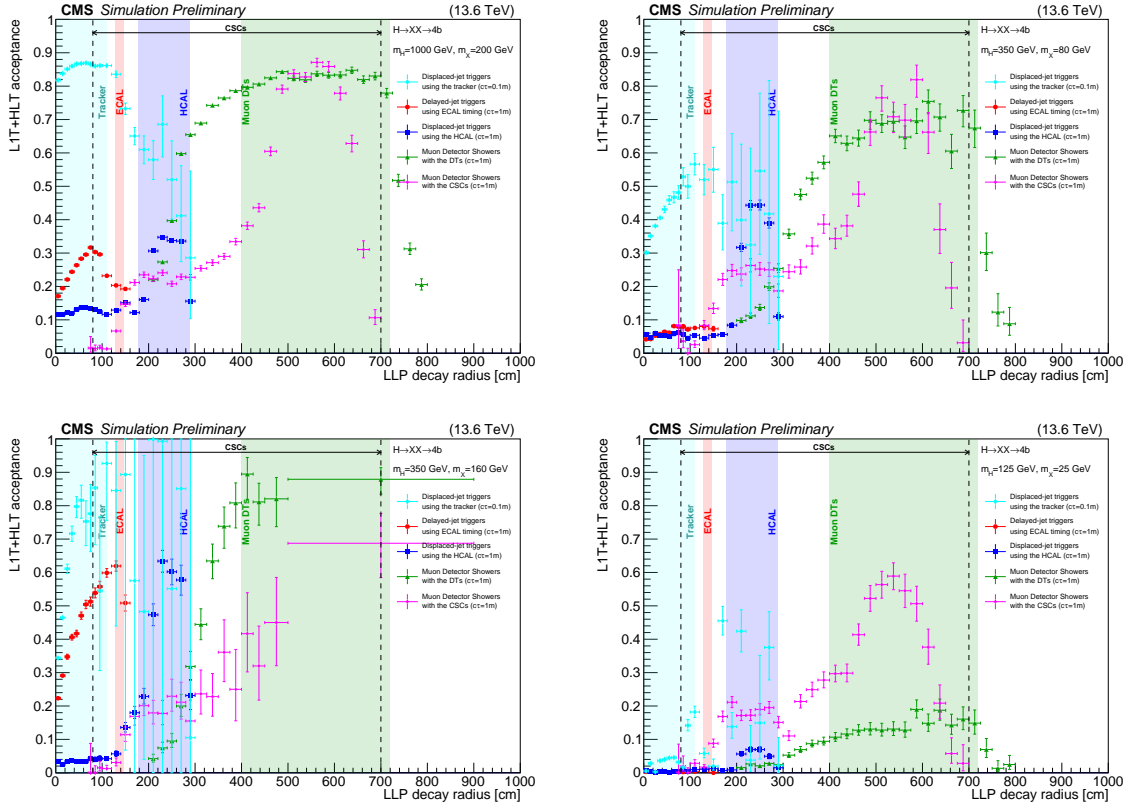


Figure 63: L1T+HLT acceptance for various LLP triggers using different subdetectors, as a function of LLP decay radius, for $H \rightarrow XX \rightarrow b\bar{b}b\bar{b}$ events in 2023 conditions. The plots correspond to different signal points with $m_H = 1000 \text{ GeV}$, $m_X = 200 \text{ GeV}$ on the upper left, $m_H = 350 \text{ GeV}$, $m_X = 80 \text{ GeV}$ on the upper right, $m_H = 350 \text{ GeV}$, $m_X = 160 \text{ GeV}$ on the lower left, and $m_H = 125 \text{ GeV}$ and $m_X = 25 \text{ GeV}$ on the lower right. In each of these plots, the $c\tau$ is 0.1 m for the displaced-jet triggers using the tracker and 1 m for the other triggers. The acceptance is shown for the displaced-jet triggers using the tracker (cyan points), for the delayed-jet triggers using ECAL timing (red circles), for the displaced-jet triggers using the HCAL (blue squares), for the MDS triggers with the DTs (green triangles), and for the MDS triggers with the CSCs (pink points). The boundaries of the tracker, ECAL, HCAL, DTs, and CSCs are also shown in the figures.




Figure 64: The fiducial regions defined for the calculation of each trigger acceptance in Fig. 63. The text contains the exact definitions of each region.




Figure 65: Diagram of the overlapping coverage, in the HLT muon p_T - d_0 plane, of the dimuon triggers that target displaced signatures. The L2 muons are shown in blue, the L3 muons are shown in red, and the scouting muons are shown in green. The displaced dimuon triggers, the double-displaced L3 muon triggers, and the dimuon scouting triggers are shown. The p_T of the highest- p_T muon (p_T^1) is shown on the y -axis, and the minimum p_T thresholds on each triggered muon ($p_T^{1,2}$) are indicated within the colored parts of the diagram. The value of 100 (1000) cm given on the x -axis for the end of the tracker (muon system) is approximate.




Figure 66: L1T+HLT acceptance of the displaced τ_h trigger, for simulated $pp \rightarrow \tilde{\tau}\tilde{\tau} (\tilde{\tau} \rightarrow \tau\tilde{\chi}_1^0)$ events, where each τ decays hadronically and the $\tilde{\tau}$ has a simulated $c\tau$ of 10 cm. The acceptance is shown for the displaced di- τ_h trigger path using 2022 data-taking conditions and is plotted with respect to the generator-level $\tilde{\tau}$ decay radius. Selections on the visible component of the generator-level tau lepton p_T ($p_T(\tau) > 30$ GeV), its pseudorapidity ($|\eta(\tau)| < 2.1$), and its decay radius ($R < 115$ cm) are applied.

icated LLP triggers extend the standard CMS trigger program and provide crucial access to unconventional LLP event signatures. In time for 2022 data taking, some LLP triggers, such as the tracking-based displaced-jet, displaced-muon, and dimuon-scouting triggers, have been improved. Other LLP triggers, such as the displaced-tau trigger, delayed-jet trigger using the electromagnetic calorimeter timing, displaced-jet trigger using the hadronic calorimeter, and the muon detector shower trigger, were newly added in 2022. The performance of these triggers has been shown with several new physics models and with 2022–2024 proton-proton collision data collected at $\sqrt{s} = 13.6$ TeV.

These triggers greatly improve the CMS sensitivity to long-lived, beyond-the-standard model particles, enhancing the experiment's ability to search for new physics. At the time of writing, two CMS LLP analyses based on early Run 3 data collected with dedicated LLP triggers have been published. These analyses are a search for displaced jets [85] and a search for displaced dimuons [105]. In the next years, many more LLP analyses using these data and the trigger presented here are expected.

References

- [1] A. Pomarol, "Beyond the standard model", in *2010 European School of High Energy Physics*, p. 115. 2012. [arXiv:1202.1391](#).
- [2] E. Ponton, "TASI 2011: Four lectures on TeV scale extra dimensions", in *Theoretical Advanced Study Institute in Elementary Particle Physics: The Dark Secrets of the Terascale*, p. 283. 2013. [arXiv:1207.3827](#). [doi:10.1142/9789814390163_0007](#).
- [3] D. Saikumar, "Exploring the frontiers: Challenges and theories beyond the standard model", 2024. [arXiv:2404.03666](#).
- [4] Particle Data Group Collaboration, "Review of particle physics", *Phys. Rev. D* **110** (2024) 030001, [doi:10.1103/PhysRevD.110.030001](#).
- [5] J. Alimena et al., "Searching for long-lived particles beyond the standard model at the Large Hadron Collider", *J. Phys. G* **47** (2020) 090501, [doi:10.1088/1361-6471/ab4574](#), [arXiv:1903.04497](#).
- [6] L. Lee, C. Ohm, A. Soffer, and T.-T. Yu, "Collider searches for long-lived particles beyond the standard model", *Prog. Part. Nucl. Phys.* **106** (2019) 210, [doi:10.1016/j.pnpnp.2019.02.006](#), [arXiv:1810.12602](#). [Erratum: [doi:10.1016/j.pnpnp.2021.103912](#)].
- [7] D. Curtin et al., "Long-lived particles at the energy frontier: The MATHUSLA physics case", *Rept. Prog. Phys.* **82** (2019) 116201, [doi:10.1088/1361-6633/ab28d6](#), [arXiv:1806.07396](#).
- [8] ATLAS Collaboration, "Search for displaced vertices of oppositely charged leptons from decays of long-lived particles in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector", *Phys. Lett. B* **801** (2020) 135114, [doi:10.1016/j.physletb.2019.135114](#), [arXiv:1907.10037](#).
- [9] ATLAS Collaboration, "Search for heavy neutral leptons in decays of W bosons produced in 13 TeV pp collisions using prompt and displaced signatures with the ATLAS detector", *JHEP* **10** (2019) 265, [doi:10.1007/JHEP10\(2019\)265](#), [arXiv:1905.09787](#).

- [10] LHCb Collaboration, “Search for $a' \rightarrow \mu^+\mu^-$ decays”, *Phys. Rev. Lett.* **124** (2020) 041801, doi:10.1103/PhysRevLett.124.041801, arXiv:1910.06926.
- [11] CMS Collaboration, “Search for long-lived particles using delayed photons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **100** (2019) 112003, doi:10.1103/PhysRevD.100.112003, arXiv:1909.06166.
- [12] CMS Collaboration, “Search for long-lived particles using nonprompt jets and missing transverse momentum with proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Lett. B* **797** (2019) 134876, doi:10.1016/j.physletb.2019.134876, arXiv:1906.06441.
- [13] CMS Collaboration, “Search for disappearing tracks in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Lett. B* **806** (2020) 135502, doi:10.1016/j.physletb.2020.135502, arXiv:2004.05153.
- [14] ATLAS Collaboration, “Search for magnetic monopoles and stable high-electric-charge objects in 13 TeV proton-proton collisions with the ATLAS detector”, *Phys. Rev. Lett.* **124** (2020) 031802, doi:10.1103/PhysRevLett.124.031802, arXiv:1905.10130.
- [15] CMS Collaboration, “The CMS experiment at the CERN LHC”, *JINST* **3** (2008) S08004, doi:10.1088/1748-0221/3/08/S08004.
- [16] 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, arXiv:2309.05466.
- [17] CMS Collaboration, “Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **15** (2020) P10017, doi:10.1088/1748-0221/15/10/P10017, arXiv:2006.10165.
- [18] CMS Collaboration, “The CMS trigger system”, *JINST* **12** (2017) P01020, doi:10.1088/1748-0221/12/01/P01020, arXiv:1609.02366.
- [19] 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, arXiv:2410.17038.
- [20] CMS Collaboration, “Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC”, *JINST* **16** (2021) P05014, doi:10.1088/1748-0221/16/05/P05014, arXiv:2012.06888.
- [21] CMS Collaboration, “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **13** (2018) P06015, doi:10.1088/1748-0221/13/06/P06015, arXiv:1804.04528.
- [22] CMS Collaboration, “Description and performance of track and primary-vertex reconstruction with the CMS tracker”, *JINST* **9** (2014) P10009, doi:10.1088/1748-0221/9/10/P10009, arXiv:1405.6569.
- [23] Tracker Group of the CMS Collaboration, “The CMS phase-1 pixel detector upgrade”, *JINST* **16** (2021) P02027, doi:10.1088/1748-0221/16/02/P02027, arXiv:2012.14304.

-
- [24] CMS Collaboration, “Track impact parameter resolution for the full pseudo rapidity coverage in the 2017 dataset with the CMS phase-1 pixel detector”, CMS Detector Performance Summary CMS-DP-2020-049, 2020.
- [25] CMS Collaboration, “Time reconstruction and performance of the CMS electromagnetic calorimeter”, *JINST* **5** (2010) T03011, doi:10.1088/1748-0221/5/03/T03011, arXiv:0911.4044.
- [26] CMS Collaboration, “CMS technical design report for the Phase 1 upgrade of the hadron calorimeter”, CMS Technical Proposal CERN-LHCC-2012-015, CMS-TDR-010, 2012. doi:10.2172/1151651.
- [27] CMS Collaboration, “Measurements with silicon photomultipliers of dose-rate effects in the radiation damage of plastic scintillator tiles in the CMS hadron endcap calorimeter”, *JINST* **15** (2020) P06009, doi:10.1088/1748-0221/15/06/P06009, arXiv:2001.06553.
- [28] CMS Collaboration, “Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV”, *JINST* **12** (2017) P02014, doi:10.1088/1748-0221/12/02/P02014, arXiv:1607.03663.
- [29] CMS Collaboration, “The Phase 2 upgrade of the CMS muon detectors”, CMS Technical Proposal CERN-LHCC-2017-012, CMS-TDR-016, 2017.
- [30] CMS Collaboration, “CMS technical design report for the Level-1 trigger upgrade”, CMS Technical Proposal CERN-LHCC-2013-011, CMS-TDR-012, 2013.
- [31] J. Duarte et al., “Fast inference of deep neural networks in FPGAs for particle physics”, *JINST* **13** (2018) P07027, doi:10.1088/1748-0221/13/07/P07027, arXiv:1804.06913.
- [32] C. D. Jones et al., “Using the CMS threaded framework in a production environment”, in *Proc. 21st International Conference on Computing in High Energy and Nuclear Physics (CHEP 2015): Okinawa, Japan, April 13–17, 2015*. 2015. [J. Phys. Conf. Ser. 664 (2015) 072026]. doi:10.1088/1742-6596/664/7/072026.
- [33] CMS Collaboration, “Enriching the physics program of the CMS experiment via data scouting and data parking”, *Phys. Rept.* **1115** (2025) 678, doi:10.1016/j.physrep.2024.09.006, arXiv:2403.16134.
- [34] A. Bocci et al., “Heterogeneous reconstruction of tracks and primary vertices with the CMS pixel tracker”, *Front. Big Data* **3** (2020) 601728, doi:10.3389/fdata.2020.601728, arXiv:2008.13461.
- [35] CMS Collaboration, “Developing GPU-compliant algorithms for CMS ECAL local reconstruction during LHC Run 3 and Phase 2”, *J. Phys. Conf. Ser.* **2438** (2023), no. 1, 012027, doi:10.1088/1742-6596/2438/1/012027.
- [36] CMS Collaboration, “Heterogeneous computing for the local reconstruction algorithms of the CMS calorimeters”, *J. Phys. Conf. Ser.* **1525** (2020), no. 1, 012040, doi:10.1088/1742-6596/1525/1/012040.
- [37] E. Bols et al., “Jet flavour classification using DeepJet”, *JINST* **15** (2020) P12012, doi:10.1088/1748-0221/15/12/P12012, arXiv:2008.10519.

- [38] H. Qu and L. Gouskos, “ParticleNet: Jet tagging via particle clouds”, *Phys. Rev. D* **101** (2020) 056019, doi:10.1103/PhysRevD.101.056019, arXiv:1902.08570.
- [39] CMS Collaboration, “Particle-flow reconstruction and global event description with the CMS detector”, *JINST* **12** (2017) P10003, doi:10.1088/1748-0221/12/10/P10003, arXiv:1706.04965.
- [40] CMS Collaboration, “Technical proposal for the Phase-II upgrade of the Compact Muon Solenoid”, CMS Technical Proposal CERN-LHCC-2015-010, CMS-TDR-15-02, 2015.
- [41] R. Frühwirth, “Application of Kalman filtering to track and vertex fitting”, *Nucl. Instrum. Meth. A* **262** (1987) 444, doi:10.1016/0168-9002(87)90887-4.
- [42] K. Rose, “Deterministic annealing for clustering, compression, classification, regression, and related optimization problems”, *IEEE Proc.* **86** (1998) 2210, doi:10.1109/5.726788.
- [43] R. Frühwirth, W. Waltenberger, and P. Vanlaer, “Adaptive vertex fitting”, *J. Phys. G* **34** (2007) N343, doi:10.1088/0954-3899/34/12/N01.
- [44] CMS Collaboration, “Performance of CMS muon reconstruction in pp collision events at $\sqrt{s} = 7$ TeV”, *JINST* **7** (2012) P10002, doi:10.1088/1748-0221/7/10/P10002, arXiv:1206.4071.
- [45] CMS Collaboration, “Identification of low-momentum muons in the CMS detector using multivariate techniques in proton-proton collisions at $\sqrt{s} = 13.6$ TeV”, *JINST* **20** (2025) P04021, doi:10.1088/1748-0221/20/04/P04021, arXiv:2412.17590.
- [46] CMS Collaboration, “Performance of the CMS muon trigger system in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **16** (2021) P07001, doi:10.1088/1748-0221/16/07/P07001, arXiv:2102.04790.
- [47] W. Adam, R. Frühwirth, A. Strandlie, and T. Todorov, “Reconstruction of electrons with the Gaussian-sum filter in the CMS tracker at LHC”, *J. Phys. G* **31** (2005) N9, doi:10.1088/0954-3899/31/9/N01, arXiv:physics/0306087.
- [48] CMS Collaboration, “Performance of photon reconstruction and identification with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV”, *JINST* **10** (2015) P08010, doi:10.1088/1748-0221/10/08/P08010, arXiv:1502.02702.
- [49] CMS Collaboration, “Performance of electron reconstruction and selection with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV”, *JINST* **10** (2015) P06005, doi:10.1088/1748-0221/10/06/P06005, arXiv:1502.02701.
- [50] CMS Collaboration, “ECAL 2016 refined calibration and Run 2 summary plots”, CMS Detector Performance Summary CMS-DP-2020-021, 2020.
- [51] M. Cacciari, G. P. Salam, and G. Soyez, “The anti- k_T jet clustering algorithm”, *JHEP* **04** (2008) 063, doi:10.1088/1126-6708/2008/04/063, arXiv:0802.1189.
- [52] M. Cacciari, G. P. Salam, and G. Soyez, “FastJet user manual”, *Eur. Phys. J. C* **72** (2012) 1896, doi:10.1140/epjc/s10052-012-1896-2, arXiv:1111.6097.
- [53] CMS Collaboration, “Performance of reconstruction and identification of τ leptons decaying to hadrons and ν_τ in pp collisions at $\sqrt{s} = 13$ TeV”, *JINST* **13** (2018) P10005, doi:10.1088/1748-0221/13/10/P10005, arXiv:1809.02816.

-
- [54] CMS Collaboration, “Identification of hadronic tau lepton decays using a deep neural network”, *JINST* **17** (2022) P07023, doi:10.1088/1748-0221/17/07/P07023, arXiv:2201.08458.
- [55] CMS Collaboration, “Performance of missing transverse momentum reconstruction in proton-proton collisions at $\sqrt{s} = 13$ TeV using the CMS detector”, *JINST* **14** (2019) P07004, doi:10.1088/1748-0221/14/07/P07004, arXiv:1903.06078.
- [56] D. Bertolini, P. Harris, M. Low, and N. Tran, “Pileup per particle identification”, *JHEP* **10** (2014) 059, doi:10.1007/JHEP10(2014)059, arXiv:1407.6013.
- [57] J. Alwall et al., “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”, *JHEP* **07** (2014) 079, doi:10.1007/JHEP07(2014)079, arXiv:1405.0301.
- [58] S. Frixione, P. Nason, and C. Oleari, “Matching NLO QCD computations with parton shower simulations: The POWHEG method”, *JHEP* **11** (2007) 070, doi:10.1088/1126-6708/2007/11/070, arXiv:0709.2092.
- [59] T. Sjöstrand et al., “An introduction to PYTHIA 8.2”, *Comput. Phys. Commun.* **191** (2015) 159, doi:10.1016/j.cpc.2015.01.024, arXiv:1410.3012.
- [60] J. Alwall et al., “Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions”, *Eur. Phys. J. C* **53** (2008) 473, doi:10.1140/epjc/s10052-007-0490-5, arXiv:0706.2569.
- [61] R. Frederix and S. Frixione, “Merging meets matching in MC@NLO”, *JHEP* **12** (2012) 061, doi:10.1007/JHEP12(2012)061, arXiv:1209.6215.
- [62] NNPDF Collaboration, “Parton distributions from high-precision collider data”, *Eur. Phys. J. C* **77** (2017) 663, doi:10.1140/epjc/s10052-017-5199-5, arXiv:1706.00428.
- [63] CMS Collaboration, “Extraction and validation of a new set of CMS PYTHIA8 tunes from underlying-event measurements”, *Eur. Phys. J. C* **80** (2020) 4, doi:10.1140/epjc/s10052-019-7499-4, arXiv:1903.12179.
- [64] GEANT4 Collaboration, “GEANT4—a simulation toolkit”, *Nucl. Instrum. Meth. A* **506** (2003) 250, doi:10.1016/S0168-9002(03)01368-8.
- [65] G. F. Giudice, M. A. Luty, H. Murayama, and R. Rattazzi, “Gaugino mass without singlets”, *JHEP* **12** (1998) 027, doi:10.1088/1126-6708/1998/12/027, arXiv:hep-ph/9810442.
- [66] L. Randall and R. Sundrum, “Out of this world supersymmetry breaking”, *Nucl. Phys. B* **557** (1999) 79, doi:10.1016/S0550-3213(99)00359-4, arXiv:hep-th/9810155.
- [67] Z. Chacko, H.-S. Goh, and R. Harnik, “The twin Higgs: Natural electroweak breaking from mirror symmetry”, *Phys. Rev. Lett.* **96** (2006) 231802, doi:10.1103/PhysRevLett.96.231802, arXiv:hep-ph/0506256.
- [68] D. Curtin and C. B. Verhaaren, “Discovering uncolored naturalness in exotic Higgs decays”, *JHEP* **12** (2015) 072, doi:10.1007/JHEP12(2015)072, arXiv:1506.06141.

- [69] H.-C. Cheng, S. Jung, E. Salvioni, and Y. Tsai, “Exotic quarks in twin Higgs models”, *JHEP* **03** (2016) 074, doi:10.1007/JHEP03(2016)074, arXiv:1512.02647.
- [70] D. Curtin, R. Essig, S. Gori, and J. Shelton, “Illuminating dark photons with high-energy colliders”, *JHEP* **02** (2015) 157, doi:10.1007/JHEP02(2015)157, arXiv:1412.0018.
- [71] B. C. Allanach et al., “The Snowmass points and slopes: Benchmarks for SUSY searches”, *Eur. Phys. J. C* **25** (2002) 113, doi:10.1007/s10052-002-0949-3, arXiv:hep-ph/0202233.
- [72] F. Blekman et al., “Soft displaced leptons at the LHC”, *JHEP* **11** (2020) 112, doi:10.1007/JHEP11(2020)112, arXiv:2007.03708.
- [73] CMS Collaboration, “Performance of long lived particle triggers in Run 3”, CMS Detector Performance Summary CMS-DP-2023-043, 2023.
- [74] CMS Collaboration, “Level-1 trigger algorithm for long-lived particle jets in Run 3”, CMS Detector Performance Summary CMS-DP-2024-058, 2024.
- [75] CMS Collaboration, “CSC high multiplicity trigger in Run 3”, CMS Detector Performance Summary CMS-DP-2022-062, CERN-CMS-DP-2022-062, 2022.
- [76] CMS Collaboration, “High multiplicity trigger for long-lived particles in CMS detector in 2022 and 2023”, CMS Detector Performance Summary CMS-DP-2024-099, CERN-CMS-DP-2024-099, 2024.
- [77] CMS Collaboration, “Displaced photons HLT trigger on 2017 data at 13 TeV”, CMS Detector Performance Summary CMS-DP-2020-015; CERN-CMS-DP-2020-015, 2020.
- [78] CMS Collaboration, A. Giannini, “A DNN for CMS track classification and selection”, in *26th International Conference on Computing in High Energy & Nuclear Physics*. 2023. arXiv:2311.05157.
- [79] CMS Collaboration, “Performance of track reconstruction at the CMS high-level trigger in 2023 data”, CMS Detector Performance Summary CMS-DP-2024-013, CERN-CMS-DP-2024-013, 2024.
- [80] M. Buican, P. Meade, N. Seiberg, and D. Shih, “Exploring general gauge mediation”, *JHEP* **03** (2009) 016, doi:10.1088/1126-6708/2009/03/016, arXiv:0812.3668.
- [81] P. Meade, N. Seiberg, and D. Shih, “General gauge mediation”, *Prog. Theor. Phys. Suppl.* **177** (2009) 143, doi:10.1143/PTPS.177.143, arXiv:0801.3278.
- [82] G. F. Giudice and R. Rattazzi, “Theories with gauge mediated supersymmetry breaking”, *Phys. Rept.* **322** (1999) 419, doi:10.1016/S0370-1573(99)00042-3, arXiv:hep-ph/9801271.
- [83] CMS Collaboration, “Search for direct pair production of supersymmetric partners of τ leptons in the final state with two hadronically decaying τ leptons and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **108** (2023) 012011, doi:10.1103/PhysRevD.108.012011, arXiv:2207.02254.
- [84] CMS Collaboration, “Search for long-lived particles using displaced jets in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **104** (2021) 012015, doi:10.1103/PhysRevD.104.012015, arXiv:2012.01581.

-
- [85] CMS Collaboration, “Search for light long-lived particles decaying to displaced jets in proton-proton collisions at $\sqrt{s} = 13.6$ TeV”, *Rep. Prog. Phys.* **88** (2025) 037801, doi:10.1088/1361-6633/adaa13, arXiv:2409.10806.
- [86] N. Arkani-Hamed and S. Dimopoulos, “Supersymmetric unification without low energy supersymmetry and signatures for fine-tuning at the LHC”, *JHEP* **06** (2005) 073, doi:10.1088/1126-6708/2005/06/073, arXiv:hep-th/0405159.
- [87] G. F. Giudice and A. Romanino, “Split supersymmetry”, *Nucl. Phys. B* **699** (2004) 65, doi:10.1016/j.nuclphysb.2004.08.001, arXiv:hep-ph/0406088. [Erratum: doi:10.1016/j.nuclphysb.2004.11.048].
- [88] J. L. Hewett, B. Lillie, M. Masip, and T. G. Rizzo, “Signatures of long-lived gluinos in split supersymmetry”, *JHEP* **09** (2004) 070, doi:10.1088/1126-6708/2004/09/070, arXiv:hep-ph/0408248.
- [89] N. Arkani-Hamed, S. Dimopoulos, G. F. Giudice, and A. Romanino, “Aspects of split supersymmetry”, *Nucl. Phys. B* **709** (2005) 3, doi:10.1016/j.nuclphysb.2004.12.026, arXiv:hep-ph/0409232.
- [90] P. Gambino, G. F. Giudice, and P. Slavich, “Gluino decays in split supersymmetry”, *Nucl. Phys. B* **726** (2005) 35, doi:10.1016/j.nuclphysb.2005.08.011, arXiv:hep-ph/0506214.
- [91] A. Arvanitaki, N. Craig, S. Dimopoulos, and G. Villadoro, “Mini-split”, *JHEP* **02** (2013) 126, doi:10.1007/JHEP02(2013)126, arXiv:1210.0555.
- [92] N. Arkani-Hamed et al., “Simply unnatural supersymmetry”, 2012, arXiv:1212.6971.
- [93] P. Fayet, “Supergauge invariant extension of the Higgs mechanism and a model for the electron and its neutrino”, *Nucl. Phys. B* **90** (1975) 104, doi:10.1016/0550-3213(75)90636-7.
- [94] G. R. Farrar and P. Fayet, “Phenomenology of the production, decay, and detection of new hadronic states associated with supersymmetry”, *Phys. Lett. B* **76** (1978) 575, doi:10.1016/0370-2693(78)90858-4.
- [95] S. Weinberg, “Supersymmetry at ordinary energies: 1. Masses and conservation laws”, *Phys. Rev. D* **26** (1982) 287, doi:10.1103/PhysRevD.26.287.
- [96] L. J. Hall and M. Suzuki, “Explicit R -parity breaking in supersymmetric models”, *Nucl. Phys. B* **231** (1984) 419, doi:10.1016/0550-3213(84)90513-3.
- [97] R. Barbier et al., “ R -parity violating supersymmetry”, *Phys. Rept.* **420** (2005) 1, doi:10.1016/j.physrep.2005.08.006, arXiv:hep-ph/0406039.
- [98] C. Csaki, E. Kuflik, and T. Volansky, “Dynamical R -parity violation”, *Phys. Rev. Lett.* **112** (2014) 131801, doi:10.1103/PhysRevLett.112.131801, arXiv:1309.5957.
- [99] C. Csaki, E. Kuflik, O. Slone, and T. Volansky, “Models of dynamical R -parity violation”, *JHEP* **06** (2015) 045, doi:10.1007/JHEP06(2015)045, arXiv:1502.03096.
- [100] ATLAS Collaboration, “Search for light long-lived particles in pp collisions at $\sqrt{s} = 13$ TeV using displaced vertices in the ATLAS inner detector”, *Phys. Rev. Lett.* **133** (2024) 161803, doi:10.1103/PhysRevLett.133.161803, arXiv:2403.15332.

- [101] CMS Collaboration, “A novel timing trigger with the CMS hadron calorimeter”, CMS Detector Performance Summary CERN-CMS-DN-2023-022, 2023.
- [102] M. J. Strassler and K. M. Zurek, “Echoes of a hidden valley at hadron colliders”, *Phys. Lett. B* **651** (2007) 374, doi:10.1016/j.physletb.2007.06.055, arXiv:hep-ph/0604261.
- [103] CMS Collaboration, “The CMS electromagnetic calorimeter project: Technical Design Report”, CMS Technical Design Report CERN-LHCC-97-033, CMS-TDR-4, 1997.
- [104] CMS Collaboration, “Search for long-lived particles decaying to a pair of muons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **05** (2023) 228, doi:10.1007/JHEP05(2023)228, arXiv:2205.08582.
- [105] CMS Collaboration, “Search for long-lived particles decaying to final states with a pair of muons in proton-proton collisions at $\sqrt{s} = 13.6$ TeV”, *JHEP* **05** (2024) 047, doi:10.1007/JHEP05(2024)047, arXiv:2402.14491.
- [106] CMS Collaboration, “The algorithm of the CMS Level-1 overlap muon track finder trigger”, *Nucl. Instrum. Meth. A* **936** (2019) 368, doi:10.1016/j.nima.2018.10.173.
- [107] CMS Collaboration, “Boosted decision trees in the Level-1 muon endcap trigger at CMS”, *J. Phys. Conf. Ser.* **1085** (2018) 042042, doi:10.1088/1742-6596/1085/4/042042.
- [108] CMS Collaboration, “Search for long-lived particles decaying to leptons with large impact parameter in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Eur. Phys. J. C* **82** (2022) 153, doi:10.1140/epjc/s10052-022-10027-3, arXiv:2110.04809.
- [109] P. W. Graham, D. E. Kaplan, S. Rajendran, and P. Saraswat, “Displaced supersymmetry”, *JHEP* **07** (2012) 149, doi:10.1007/JHEP07(2012)149, arXiv:1204.6038.
- [110] J. A. Evans and J. Shelton, “Long-lived staus and displaced leptons at the LHC”, *JHEP* **04** (2016) 056, doi:10.1007/JHEP04(2016)056, arXiv:1601.01326.
- [111] M. J. Strassler and K. M. Zurek, “Discovering the Higgs through highly-displaced vertices”, *Phys. Lett. B* **661** (2008) 263, doi:10.1016/j.physletb.2008.02.008, arXiv:hep-ph/0605193.
- [112] CMS Collaboration, “Dark sector searches with the CMS experiment”, *Phys. Rept.* **1115** (2025) 448, doi:10.1016/j.physrep.2024.09.013, arXiv:2405.13778.
- [113] CMS Collaboration, “Search for long-lived particles decaying into muon pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV collected with a dedicated high-rate data stream”, *JHEP* **04** (2022) 062, doi:10.1007/JHEP04(2022)062, arXiv:2112.13769.
- [114] CMS Collaboration, “Search for long-lived particles decaying in the CMS endcap muon detectors in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. Lett.* **127** (2021) 261804, doi:10.1103/PhysRevLett.127.261804, arXiv:2107.04838.
- [115] CMS Collaboration, “Search for long-lived particles decaying in the CMS muon detectors in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **110** (2024) 032007, doi:10.1103/PhysRevD.110.032007, arXiv:2402.01898.

- [116] Y. L. Dokshitzer, G. D. Leder, S. Moretti, and B. R. Webber, “Better jet clustering algorithms”, *JHEP* **08** (1997) 001, doi:10.1088/1126-6708/1997/08/001, arXiv:hep-ph/9707323.
- [117] M. Wobisch and T. Wengler, “Hadronization corrections to jet cross-sections in deep inelastic scattering”, in *Workshop on Monte Carlo Generators for HERA Physics (Plenary Starting Meeting)*, p. 270. 1998. arXiv:hep-ph/9907280.
- [118] R. Bailey and P. Collier, “Standard filling schemes for various LHC operation modes”, Technical Report LHC-PROJECT-NOTE-323, 2003.
- [119] CMS Collaboration, “Search for decays of stopped long-lived particles produced in proton-proton collisions at $\sqrt{s} = 8$ TeV”, *Eur. Phys. J. C* **75** (2015) 151, doi:10.1140/epjc/s10052-015-3367-z, arXiv:1501.05603.
- [120] CMS Collaboration, “Search for decays of stopped exotic long-lived particles produced in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **05** (2018) 127, doi:10.1007/JHEP05(2018)127, arXiv:1801.00359.
- [121] D. Abercrombie et al., “Dark matter benchmark models for early LHC Run 2 searches: Report of the ATLAS/CMS dark matter forum”, *Phys. Dark Univ.* **27** (2020) 100371, doi:10.1016/j.dark.2019.100371, arXiv:1507.00966.
- [122] L. M. Carpenter, R. Colburn, J. Goodman, and T. Linden, “Indirect detection constraints on s and t channel simplified models of dark matter”, *Phys. Rev. D* **94** (2016) 055027, doi:10.1103/PhysRevD.94.055027, arXiv:1606.04138.
- [123] L. Covi and F. Dradi, “Long-lived stop at the LHC with or without R-parity”, *JCAP* **10** (2014) 039, doi:10.1088/1475-7516/2014/10/039, arXiv:1403.4923.
- [124] P. Langacker and G. Steigman, “Requiem for a fractionally charged, massive particle”, *Phys. Rev. D* **84** (2011) 065040, doi:10.1103/PhysRevD.84.065040, arXiv:1107.3131.

A Glossary of acronyms

AMSB	Anomaly-mediated supersymmetry breaking
BMTF	Barrel muon track finder
BPTX	Beam pickup timing device
BSM	Beyond the standard model
BX	Bunch slot
CA	Cambridge–Aachen
CMS	Compact Muon Solenoid
CSC	Cathode strip chamber
CPU	Central processing unit
DAQ	Data acquisition
DNN	Deep neural network
DT	Drift tube
ECAL	Electromagnetic calorimeter
EMTF	Endcap muon track finder
GEM	Gas electron multiplier
GMSB	Gauge-mediated supersymmetry breaking
GPU	Graphical processing unit
HAHM	Hidden Abelian Higgs model
HCAL	Hadronic calorimeter
HB	Hadronic calorimeter barrel
HLT	High level trigger
HPS	Hadron plus strip
ISR	Initial state radiation
L1	First level
LHC	Large Hadron Collider
LLP	Long-lived particle
LO	Leading order
MC	Monte Carlo
MDS	Muon detector shower
ML	Machine learning
NLO	Next-to-leading order
NN	Neural network
OMTF	Overlap muon track finder
PDF	Parton distribution function
PF	Particle flow
PU	Pileup
PV	Primary vertex
QCD	Quantum chromodynamics
RPC	Resistive-plate chamber
RPV	R -parity violating
SM	Standard model
SUSY	Supersymmetry
TDC	Time-to-digital converter
3D	Three-dimensional