


The Compact Muon Solenoid Experiment

Detector Note

The content of this note is intended for CMS internal use and distribution only



31 October 2023

A Novel Timing Trigger with the CMS Hadron Calorimeter

Gillian Kopp, Chris Tully, Kiley Kennedy, Wonyong Chung, Svitlana Hoienco, Jeremy Mans, Michael Krohn,
Bryan Crossman, Joshua Hiltbrand, Andris Skuja, Long Wang

Abstract

A low-level hardware timing trigger using the Hadron Calorimeter (HCAL) in the Compact Muon Solenoid (CMS) experiment at CERN is designed and deployed for Run 3 of the Large Hadron Collider (LHC), with data taking beginning in 2022. The new Level-1 trigger relies on the Phase 1 upgrade of the HCAL, which provides greater segmentation granularity and an increased number of readout channels, and enables timing information to be used in the hardware trigger system. The timing trigger is designed to increase the CMS experiment's sensitivity to long-lived particles (LLPs) by introducing a trigger that processes data at 40 MHz to identify jets arriving at delayed times. In addition to identifying events with jets arriving with a time delay, the trigger also identifies events with unique depth signatures, both of which are characteristics of LLPs. This trigger is the first of its kind, in that timing and segmentation information is used at the hardware level to identify LLP candidate events. This paper reviews the algorithm, implementation, and performance of the trigger, and includes a discussion of the HCAL timing capabilities and online time alignment.

A Novel Timing Trigger with the CMS Hadron Calorimeter

Gillian Kopp^{*1}, Chris Tully¹, Kiley Kennedy¹, Wonyong Chung¹, Svitlana Hoienko¹,
Jeremy Mans², Michael Krohn², Bryan Crossman², Joshua Hiltbrand⁴, Andris Skuja³, and
Long Wang³

¹*Princeton University, Princeton NJ*

²*University of Minnesota, Minneapolis MN*

³*University of Maryland, College Park MD*

⁴*Baylor University, Waco TX*

October 31, 2023

*gillian.baron.kopp@cern.ch

Contents

1	Introduction	3
2	The CMS Detector and its Hadron Calorimeter	3
2.1	CMS Sub-Detectors	4
2.2	Hadron Calorimeter	4
2.2.1	Front-End Electronics	5
2.2.2	Timing Capabilities	5
2.2.3	Back-End Electronics	6
3	HCAL Timing and Alignment	6
3.1	Timing Alignment	7
3.2	QIE Phase Scan Method	8
3.3	QIE Scan in Collisions: TDC Based Time Alignment	8
3.4	Energy Impacts	10
4	Charge Weighted Time	11
4.1	CWT Calibration Procedure from QIE Scan	11
4.2	Timing in LED Runs	14
4.3	Validation of CWT Calibration using LED Data	18
5	Energy Scale Changes	19
5.1	Energy Response	19
5.2	Hit Occupancy Method	20
5.3	Validation of Method in HE	21
6	HCAL Timing in L1 Long-Lived Particle Trigger	23
6.1	uHTR Specifications	24
6.1.1	LLP flag – Depth bit	24
6.1.2	LLP flag – Timing bits	25
6.1.3	LLP flag – Energy requirements	25
6.2	L1 LLP Jet Trigger Algorithm	25
6.2.1	uGT Jet Bits	26
6.2.2	Trigger Menu Seeds	26
7	Monte Carlo Data Injection on Front-End Electronics	27
8	Performance	28
8.1	Trigger Performance in Simulation	28
8.2	Trigger Performance in Data	29
9	Summary	30

Abstract

A low-level hardware timing trigger using the Hadron Calorimeter (HCAL) in the Compact Muon Solenoid (CMS) experiment at CERN is designed and deployed for Run 3 of the Large Hadron Collider (LHC), with data taking beginning in 2022. The new Level-1 trigger relies on the Phase 1 upgrade of the HCAL, which provides greater segmentation granularity and an increased number of readout channels, and enables timing information to be used in the hardware trigger system. The timing trigger is designed to increase the CMS experiment's sensitivity to long-lived particles (LLPs) by introducing a trigger that processes data at 40 MHz to identify jets arriving at delayed times. In addition to identifying events with jets arriving with a time delay, the trigger also identifies events with unique depth signatures, both of which are characteristics of LLPs. This trigger is the first of its kind, in that timing and segmentation information is used at the hardware level to identify LLP candidate events. This paper reviews the algorithm, implementation, and performance of the trigger, and includes a discussion of the HCAL timing capabilities and online time alignment.

1 Introduction

Long-lived particles (LLPs) are present in many models beyond the Standard Model (SM) of particle physics and their unique experimental signatures provide a powerful way to search for new physics at CERN's Large Hadron Collider (LHC). LLPs may result from small couplings, suppressed phase space, and approximate symmetries. The signature of LLPs is often distinct from SM processes, potentially including displaced vertices, disappearing or a lack of tracks, and out-of-time decays with respect to collisions [1]. A key factor limiting LLP searches is that many of these events will escape the trigger system without dedicated triggers. To fully exploit the detector upgrades and the higher luminosity of Run 3 of the LHC, a dedicated LLP trigger utilizing the Phase 1 upgrade of the Compact Muon Solenoid (CMS) Hadron Calorimeter (HCAL) barrel is implemented [17].

The upgraded HCAL introduces depth segmentation and online timing measurements – both of which expand the phase space accessible in LLP searches. In particular, the HCAL behaves as an extended sampling calorimeter with depth segmentation and it is unique in being a timing detector with nearly 4π coverage. Both the depth segmentation and timing measurements are used in the trigger decision, as long-lived particles can create a jet with a time delay or with most of the energy deposited in higher calorimeter layers. LLPs decaying inside the calorimeter deposit significant energy in the higher depths and minimal energy in the lower depths. Similarly, LLPs decaying in or before the calorimeter arrive at a delayed time due to either the path length difference or a slow velocity. Jets containing multiple hits either with a delayed arrival time or with significant energy deposited in the later depth layers are selected by this trigger for further offline analysis. Thus, the new trigger targets the identification of long-lived particles (LLPs) with a new detector signature.

The calorimeter-based LLP trigger was deployed for Run 3 of the LHC, with data taking starting in 2022 from proton-proton collisions at a center of mass energy of 13.6 TeV. Run 3 will double the integrated luminosity of Run 2, which resulted in an integrated luminosity of 140 fb^{-1} , and thus presents a good opportunity for new physics searches. This paper covers the implementation of a hardware timing trigger using the HCAL of the CMS experiment, detailing the relevant hardware and firmware upgrades, trigger algorithm development, trigger commissioning, and performance in collisions.

2 The CMS Detector and its Hadron Calorimeter

The Compact Muon Solenoid experiment records proton-proton collisions at the LHC at CERN in Geneva, Switzerland. It is a general purpose experiment, designed to detect the Higgs boson and search for physics beyond the SM. In Run 3 of the LHC, protons collide at a center of mass energy of $\sqrt{s} = 13.6 \text{ TeV}$, and these collisions occur at a rate greater than 1 GHz. Bunches of protons are collided every 25 ns, and a two-tiered trigger system is used to select events for offline analysis. The Level-1 (L1) hardware trigger selects 100 kHz of events from the 40 MHz event rate, while the High Level Trigger further reduces the event rate to 1 kHz [4].

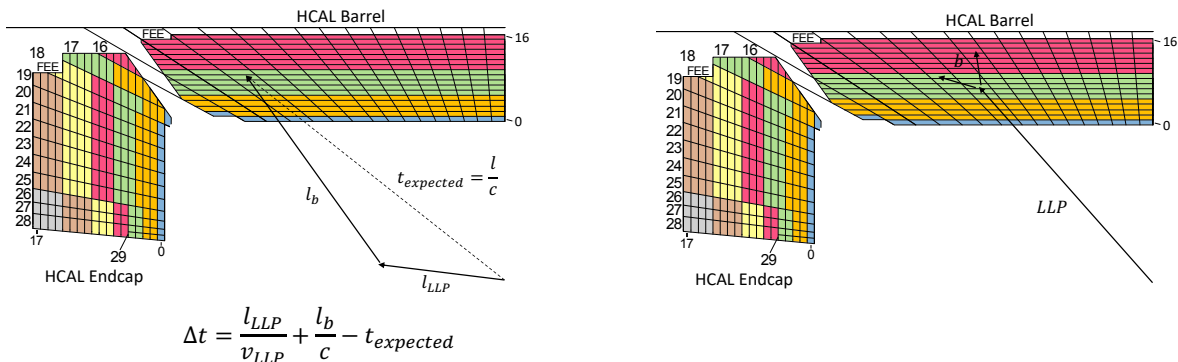
2.1 CMS Sub-Detectors

The CMS detector is formed of multiple sub-detector systems designed to detect specific particles. The innermost layer is the silicon tracker to identify the path of charged particles. Outside of the tracker is the electromagnetic calorimeter (ECAL) which measures energy from electrons and photons, and following the ECAL is the hadron calorimeter (HCAL) which measures the energy of composite particles. The superconducting solenoid sits outside of the HCAL and provides a magnetic field of 3.8 T. The outermost part of the detector is the muon chambers, as muons will pass undetected through the earlier calorimeters. A detailed description of the CMS detector is given in [2].

2.2 Hadron Calorimeter

The HCAL is comprised of four sections: barrel (HB), endcap (HE), forward (HF), and outer (HO). This work focuses on HB, which is a sampling calorimeter made of layers of plastic scintillator alternated with brass absorber. The light from the scintillation layers is transmitted through wavelength-shifting fibers to photodetectors for digitization.

During the Phase 1 upgrade of the CMS HCAL, the barrel was upgraded from hybrid photodiodes (HPDs) to silicon photomultipliers (SiPMs). SiPMs allow for higher photodetection efficiency and significantly larger gains than HPDs, and furthermore allow for an increase in the readout channels due to the improved signal-to-noise performance [6]. With more readout channels, the segmentation of the calorimeter is increased, giving precise time and energy measurements at each depth layer (four depths in HB, up to seven depths in HE, in Figure 1). The trigger development presented here takes advantage of this finer segmentation availability.



(a) Diagram of the time delay resulting from a LLP due to a combination of the path length difference and slow velocity.

(b) Diagram of LLP decaying within the HCAL volume, causing significant energy deposit to higher calorimeter layers.

Figure 1: Diagrams showing the Run 3 depth segmentation in HB (4 layers) and HE (up to 7 layers), and the unique signatures that LLP decays in or before the HCAL create. Figures adapted from [7].

Figure 1 illustrates a quarter barrel and endcap section of the HCAL, which has nearly 4π coverage. The physical towers depicted are labeled by $i\eta$, a detector coordinate that is related to η , the Lorentz invariant measure of pseudorapidity. Each $i\eta$ region corresponds to $\Delta\eta = 0.087$. The HCAL barrel (HB) extends from $|i\eta| = 1 - 16$ (about $0 < \eta < 1.4$), while the endcap (HE) extends from $|i\eta| = 16 - 29$ (about $1.4 < \eta < 2.5$). In $i\eta = 16$, there are three depth layers in the barrel and one in the endcap. The forward region (HF) extends from $|i\eta| = 29 - 41$, but is excluded from this trigger development. A single tower covers one $i\phi$ region, with a total of 72 $i\phi$ around the barrel, so each $i\phi$ is 5° in ϕ .

2.2.1 Front-End Electronics

In the Phase 1 upgrade to the HB and HE, the scintillation light is transmitted through wavelength-shifting fibers to SiPMs. The resulting SiPM signals are measured by charge integrator and encoder (QIE) ASICs in the front-end (on detector) electronics. The custom QIE11 chip performs an energy measurement with an analog-to-digital converter (ADC, via integration) and a rising edge timing measurement with a time-to-digital converter (TDC) [6]. The ADC code is reported in 8-bits while the TDC code is reported in 6-bits, providing 0.5 ns timing precision. Both measurements are performed in each 25 ns timeslice (TS) corresponding with the LHC collision rate of 40 MHz and are reported for each HCAL cell. An HCAL cell is defined as a single readout channel, which covers one $i\eta$, $i\phi$, and depth region. Per tower in the HB, there are 4 cells, one from each depth layer. Digitized signals from the QIE11 are then transmitted by the IGLOO2 FPGA to the back-end electronics (uHTRs) [3].

The QIE11 chip determines the ADC value by integrating the SiPM current output over the 25 ns TS. The charge is reported in 256 bins (8-bits) over a dynamic range of 400 pC. The QIE11 has four integrator ranges, each integrating a fraction of the current. The charge per bin is the smallest for lower input energies, as this approach allows for coverage of a large dynamic range while maintaining sensitivity to low energies.

To determine the TDC value, the current measured by SiPMs is compared directly to a flat current threshold. The threshold is an 8-bit programmable value, set at a current of 17.8 μA for Run 3 data taking. The time at which the current passes this threshold is the TDC value for this time slice and is reported in a 6-bit digital output, giving 0.5 ns bins within the 25 ns bunch crossing. This behavior is illustrated in Figure 2, where the time corresponding to when the pulse crosses the flat current threshold determines the TDC value reported.

In HB (HE), a single QIE11 chip has 16 (12) channels, and each IGLOO2 encodes data from 8 (12) channels. One readout card contains 1 QIE11 and 2 (1) IGLOO2 FPGA, and a group of 4 readout cards comprises a Readout Module (RM), with one RM covering a total of 64 (48) channels. In HB, one RM covers one-half barrel $i\phi$ region. A Readout Box (RBX) is a group of 4 RMs, a calibration unit, and two ngCCMs. There are 18 RBXs per half barrel in both HB and HE. Each RBX also has a calibration module, used to send LED signals directly to the SiPMs for local run testing [3]. The physical layout of the boards is documented in [12], showing the 16-channel QIE11 and the IGLOO2 inputs, and similarly the 12-channel HE QIEs are documented [11].

The next generation clock and control module (ngCCM) is used to distribute the clocks and set the alignment of the QIE data and is implemented on a flash based FPGA. The QIE clocks are used to scan the time delay to test the trigger pathway and determine HCAL alignment.

2.2.2 Timing Capabilities

For the timing trigger, the TDC information from the sample of interest (SOI) is used and is reported as a 6-bit TDC address from the QIE11 ASIC. In this 6-bit address, the TDC is reported in 50 valid time bins within the 25 ns bunch crossing, along with multiple error or invalid codes. Specifically, TDC code 62 indicates that the pulse is already above the threshold and thus the pulse started earlier than the current bunch crossing. TDC code 63 indicates that the pulse never crosses the threshold for a valid timing measurement in the bunch crossing of interest, as there is not a high enough energy signal. Addresses 58-61 are error codes related to the phase-locked loop state.

In HB, the IGLOO2 firmware look up table (LUT) compresses the 6-bit TDC code to 2 bits. This reduction in information is required due to bandwidth constraints. The resulting 2-bit compressed TDC encodes four ranges: prompt, slightly delayed, very delayed, and invalid or error. The last value of the prompt range, denoted t_p , along with the boundary between the two delayed ranges, denoted t_d , defines the full 6:2 LUT. t_d is set 1 ns, or 2 TDC codes after t_p ($t_d = t_p + 2$). This gives the LUT:

- 00 prompt pulse (0-6 ns): $\text{TDC} \leq t_p = 12$
- 01 slightly delayed (6.5-7 ns): $t_p < \text{TDC} \leq t_p + 2 = 14$
- 10 very delayed (7.5+ ns): $t_p + 2 < \text{TDC} < 50$

- 11 invalid pulse: $50 \leq \text{TDC}$

The 6:2 LUT is applied per channel, and 8 different LUTs may be loaded per IGLOO2, allowing for a separate LUT to be used for each channel. The LUT has $2 \cdot 2^6$ bits to encode the whole mapping [10]. For the data presented here, a LUT that does not vary by position is used.

An illustration of the four timing ranges and the flat current threshold defining where the TDC is measured is shown in Figure 2. The pulse shown is an illustration, and demonstrates that SiPM pulses have relatively fast rise times, while the tail of the pulse extends past SOI.




Figure 2: Diagram of the 2-bit TDC code ranges as measured by the rising time of a pulse. The pulse drawn has TDC=00 in SOI.

2.2.3 Back-End Electronics

The front-end electronics are in the detector cavern and continuously send data via optical links to the back-end components, which are outside the shielding wall in the service cavern. The back-end electronics are the HCAL μ TCA Trigger and Readout modules (uHTRs), Advanced Mezzanine Card (ACM13), MCH, all contained in a μ TCA crate. Each uHTR receives data from a readout box (RBX) covering 20 degrees in ϕ (4 $i\phi$ regions), and a RBX contains or 4 Readout Modules (RMs). The uHTRs are split by HB-only, HB/HE overlap, and HE-only uHTRs (these three groups of uHTRs are defined by which detector region data is received from, HB, HE, or both) [14]. Multiple uHTRs are needed to take data from the entire RBX, due to the complex tower mapping to each uHTR. There are 18 μ TCA crates each with 12 uHTRs at P5, for a total of 216 uHTRs.

Each uHTR has two flash based FPGAs: the front reads out the front-end data and performs trigger processing, while the back transmits the data to the central DAQ [6]. The LLP trigger algorithm is implemented on the front uHTR FPGA, and described in detail in Section 6.1.

3 HCAL Timing and Alignment

The online time alignment of the HCAL is vital to ensure that prompt signals arrive within the same time window and delayed signals are properly identified. The QIE phase setting determines the prompt arrival time, ideally set at the start of the SOI. Additional adjustments to the TDC LUTs can further fine-tune the

prompt-delayed boundary per channel; however, the LUT settings serve only as correction factors and do not set the overall detector alignment.

3.1 Timing Alignment

The Run 2 method of HCAL time alignment is Charge Weighted Time (CWT), which is an estimator of the pulse time computed from the charge measured in each of the 8 timeslices. CWT was the best timing estimator available before the introduction of TDC with the Phase-1 Upgrade of the HB and HE. However, CWT assumes identical pulses across the detector and is very sensitive to fluctuations in the tail of the pulse – a longer pulse fall-off will increase the CWT value, while the pulse rising edge position is unchanged. CWT prioritizes placing the weighted “center” of the pulse in the same place, potentially causing the rising edge placement to be inconsistent. As a result, some pulses can have the rising edge placed relatively early or late. In contrast, TDC measures the rising edge of the pulse, avoids the drawbacks of CWT, and is a very precise time measurement available in Run 3. Thus, TDC is used for the improved HCAL alignment, and a detailed comparison to the CWT-based alignment is in Section 4.

Since the TDC must be compressed into 4 categories as detailed in Section 2.2.2, the time alignment cannot simply be done with offline timing, as the half ns timing granularity is not accessible offline. Instead, a useful approach is special “phase scan” runs that adjust the measured position of the pulse. By understanding the fraction of pulses in each timing range throughout the scan, an optimal online alignment is determined and deployed.

The QIE phases determine where the SOI-1 to SOI boundary falls and are adjusted such that the pulse arrival times are consistent across HB. Changing the QIE phase scans the HCAL clock and either pushes prompt jets into the delayed region or early region. By performing the QIE phase scan around well time-aligned phases (nominal phases), delayed jets occurring across the whole detector at a known time are artificially produced. This phase scan is an ideal (and one of the only ways) to measure the delayed jet turn-on with the real timing resolutions of jets in collisions. The phase scan also demonstrates the LLP timing trigger performance as a function of jet time delay.

With the phase scan data, a new time alignment method using the TDC codes is developed and demonstrated to align the detector to within 0.5 ns, a 10x improvement over the previous alignment. The TDC-based time alignment method relies on sensitivity to the rising edge of the pulse and is ideal for calorimeter time alignment. This method ensures that the rising edge of the pulse is fully captured, and minimizes sensitivity to fluctuations in the tail of the pulse due to developing hadronic showers – a sensitivity present in CWT.

Precise time alignment is vital for the calorimeter timing measurement. For optimal alignment, the prompt signal rising edge will be as close to the SOI-1 and SOI boundary as possible, while avoiding spillover into SOI-1. If the alignment places a prompt pulse significantly before the SOI, then the only valid time measurement will be in SOI-1, with no valid time measured in SOI. Similarly, if the alignment places a prompt pulse very late into SOI, it will be measured as a delayed pulse. Thus, precise and consistent phase alignment is vital for detecting late arriving pulses by clearly separating these from prompt pulses – which would not be possible if the alignment set prompt pulses late in the SOI.

In addition, since the alignment setting determines the boundaries of the sample of interest, the alignment impacts the event energy. The energy is measured by integrating the pulse, and for an accurate measurement, this integral must capture the bulk of the pulse where the prompt hadronic energy is contained. Having the rising edge of the pulse be late in SOI would cause a drop in the energy measurement since more of the energy will fall in SOI+1 and thus be excluded from the integral. Furthermore, the energy filtering algorithm used accounts for pileup by subtracting SOI-1 energy from SOI energy, and therefore it is important that the pulse not be pushed into SOI-1, otherwise a portion of the energy will be missed. Well-aligned and consistent phases improve both the timing and energy measurements in the calorimeter and are important for the timing trigger.

3.2 QIE Phase Scan Method

For the QIE scan, the IGLOO2 LUT is identical for all $i\eta$ and depth values in HB, setting $t_p = 12$, or 6 ns. This LUT creates three timing ranges, 0-6 ns, 6.5-7 ns, and 7.5-25 ns, corresponding to recorded TDC codes of 00, 01, and 10, with code 11 reserved for error codes.




Figure 3: Diagram of the shifts produced by changing the QIE delays. As the QIE phase is shifted by -4 ns, the pulse is moved earlier (ie, toward SOI-1).

The QIE scan is taken relative to nominal phases in 2 ns steps, from -4 ns to 10 ns (November 2022 scan, 13.6 TeV and 2400 colliding bunches). The prompt peak is optimized to be at the 0 ns phase (to within 2 ns), and further refinements for the final alignment phases are calculated from here. An illustration of the effect of shifting the pulse by -4 and +10 ns is in Figure 3. In May 2023, with 13.6 TeV collisions and 900 colliding bunches, a second phase scan was taken relative to the recommended 2023 HCAL phases. This scan covered 30 ns, in 1 ns steps, from -10 ns to +20 ns, and achieved detector alignment to within 0.5 ns. These new phases were used in collisions data taking starting in June 2023.

3.3 QIE Scan in Collisions: TDC Based Time Alignment

The Raw JetMET dataset is used for the phase scan analysis, as it provides events enriched in jets or missing transverse energy, which are crucial to this study. A minimum pulse energy of 4 GeV E_T in SOI is required for the pulse to be included in the analysis. For all channels with pulses satisfying this requirement, the distribution of TDC times is plotted, and QIE delays are adjusted such that arrival time distributions relative to the reference clock are the same across HB. Any data from a lumisection during the QIE phase scan where the HB QIE phase did not change (1 scan point), there were unexpected jumps in L1 rate (1-3 LS), or an emittance scan was in progress were removed as part of the data cleaning.

Figures 4 and 5 show typical TDC code distributions during the 2023 HCAL phase scan taken relative to the recommended 2023 phases. The analysis is performed for HB ($-16 \leq i\eta \leq 16$) where results are consistent across all $i\eta$, and shown here for $i\eta = -10$. The HCAL QIE11 chip reports a TDC value that is an online measure of the pulse arrival time. The TDC code fraction is the fraction of cells with each TDC code (4 codes) for cells over 4 GeV E_T . The four TDC codes are prompt (TDC ≤ 6 ns), slightly delayed ($6 < \text{TDC} \leq 7$ ns), very delayed (TDC > 7 ns), and no TDC (indicating that the TDC threshold was already crossed in a previous bunch crossing, TDC threshold is not crossed, or an error code). The prompt TDC distribution (blue) shows the prompt pulse arrival spread and is maximized at the reference phase (0 ns). With increasing phase offset, more delayed (green and orange) TDC codes are seen, as pulses arrive in the delayed region. As the next bunch crossing enters (QIE11 phase offset of 20 ns), the prompt fraction begins increasing. Ideal alignment places the prompt peak at a QIE11 phase offset of 0, such that the pulse rising edge is close to the start of the 25 ns sample of interest (SOI) window. This ideal alignment is achieved for

all channels in HB to within 0.5 ns, the TDC resolution. TDC codes for cells with a minimum of 4 GeV E_T are reported, which is sufficient to minimize the timeslew effects of the pulse rising edge. The phase scan spans 25 ns, though is not symmetric in 25 ns windows due to both statistical fluctuations during the scan and changes in bunch crossing containing the pulse arrival time.

The alignment procedure maximizes the percentage of arrival times in the TDC=00 (prompt) code. As seen in Figures 4 and 5, when the pulses are pushed early (below 0ns), the percentage of TDC=00 (prompt) rapidly drops while TDC=11 (error) rises, and there is a corresponding increase in TDC=10 (very delayed) in SOI-1. When the pulses are pushed late (at +2 ns and following steps), the percentage of TDC=00 (prompt) in SOI rapidly drops, while TDC=01 and TDC=10 increase, as expected.

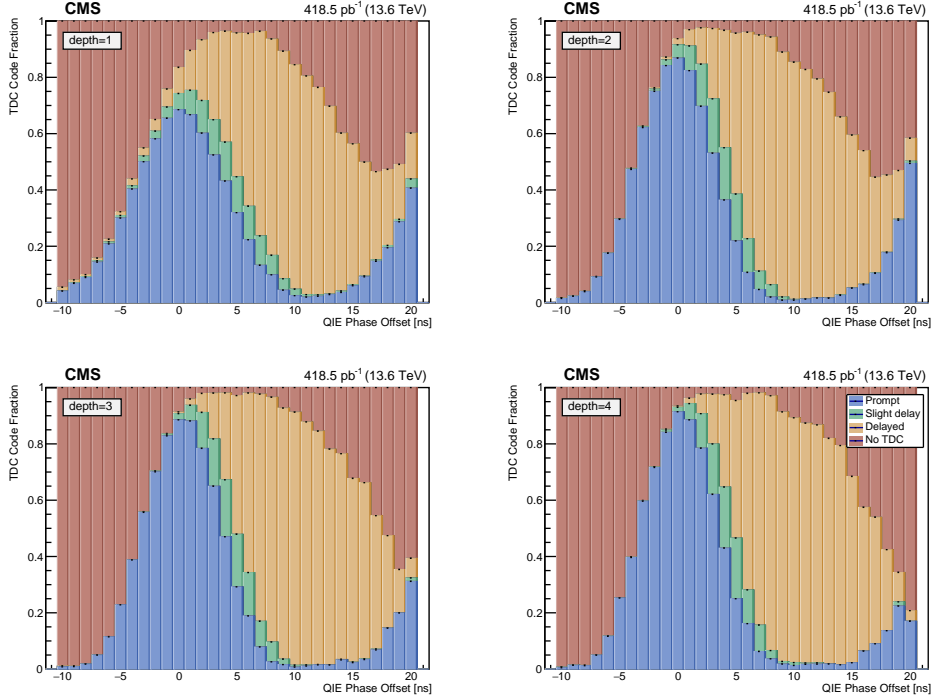


Figure 4: Distribution of TDC codes vs QIE scan, for each of 4 depths across HB, for $i\eta = -10$ in the SOI, during the May 2023 QIE phase scan. Published in [5].

There is a Gaussian spread of arrival times of the pulse rising edge within each QIE phase setting, and this distribution of these arrival times is shown in Figure 4. In ideal alignment, the prompt window (TDC=00) is kept as narrow as possible. From the phase scan TDC distributions, the prompt fraction drops off rapidly at -2 and -4 ns, indicating that the arrival time distribution is close to the clock boundary. In addition, the prompt distribution does not plateau, indicating that the prompt window is not unnecessarily large.

The TDC-based time alignment study was also performed in groupings of $i\phi \bmod 4$ and split by RBX, to check for any dependence on the physical detector layout. No change was seen in the alignment results from these investigations. Furthermore, an analysis of the impact pileup on the TDC-based alignment showed minimal dependence. The most relevant pileup contribution for TDC data is the contribution from the previous bunch crossing. Therefore, the TDC code fraction study was performed with the dataset split by bunch position in the train, specifically comparing leading vs. non-leading bunches. The TDC=00 peak position differed by < 0.5 ns in leading vs. non-leading bunches, with Depth 1 exhibiting the largest effect (TDC=00 peak shifted on average 0.25 ns later in leading bunches). The TDC threshold is high enough to not fire from pileup in the previous bunch crossing, while other time approximation methods (CWT, discussed later), will be more susceptible to pileup. Finally, the impact of changing per cell energy threshold from 4 GeV to 8 GeV is studied. An increased energy threshold shows the expected timeslew of TDC,

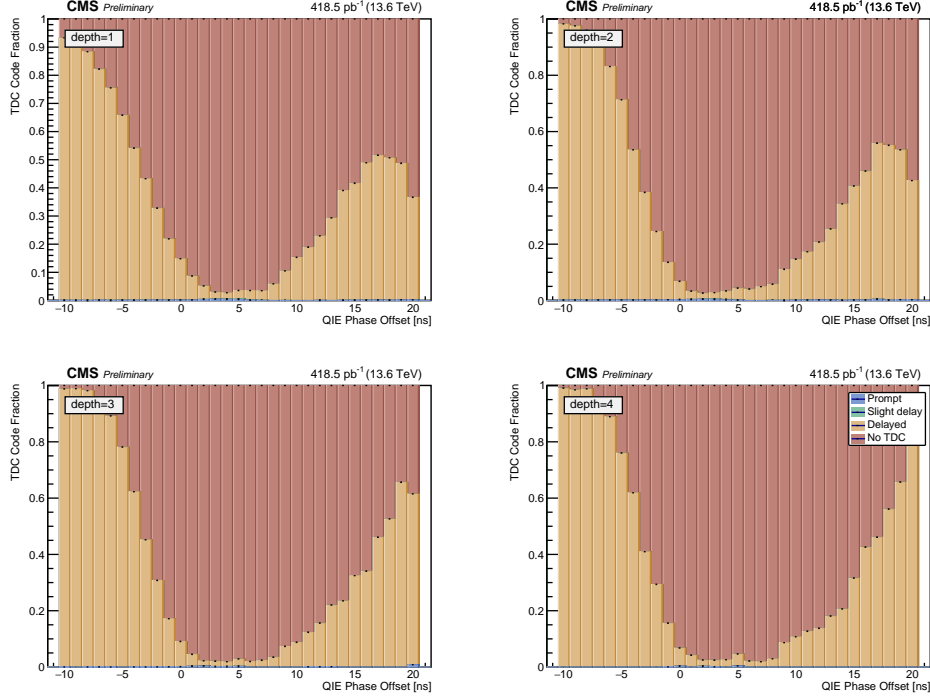


Figure 5: Distribution of TDC codes vs QIE scan, for each of 4 depths across HB, for $i\eta = -10$ in the SOI-1. The distribution of delayed codes in SOI-1 mirrors the distribution of no-TDC codes in SOI. If the TDC is set in SOI-1, no valid TDC is reported in SOI.

with the TDC=00 peak moving earlier by 3-5 ns. This leakage into SOI-1 is anticipated and leads to some inefficiency in very energetic delayed signals and a small prompt veto inefficiency. However, this is not a concern from the LLP trigger point of view and does not change the alignment conclusions as the effect is consistent across the detector.

3.4 Energy Impacts

Using the TDC code distribution is ideal for time alignment, as this can give alignment to within half a nanosecond. If pulses are vastly misaligned such that no valid TDC codes are seen within the SOI in a timing scan, an approximation can be done by fitting the pulse shape (based on raw ADC values in each timeslice) to a Landau and relying on the rising edge of the Landau distribution. This is done by defining a “turn on” point for the Landau($\mu = 0, \sigma = 1$) distribution (chosen to be the x-axis value corresponding to 66% of the maximum value in this case), and then scaling for all other pulse shape fits given the μ (most probable value) and σ (measure of the spread) fit parameters:

$$\text{TurnOn}(\mu, \sigma) = \mu + \sigma \cdot \text{TurnOn}(0, 1) \quad (1)$$


For a “well aligned” pulse (known from TDC information), the Landau fit turn-on was found to be at 3.015 (in TS units). All other channels were adjusted following:

$$\text{QIE Phase Adjustment (in ns)} = \frac{\Delta\text{Landau fit}}{0.04} \cdot \text{ns} \quad (2)$$

where $\Delta\text{Landau fit}$ is the difference between the ideal fit (3.015) and the current fit. This method of fitting is most sensitive to cases where there is significant spillover into SOI-1, however, only provides an approximate adjustment that is later refined with another timing scan and TDC code analysis.

The per-tower summed energy is also sensitive to the time alignment. In Run 3, HCAL uses the new PFA1’ algorithm, which subtracts a weighted fraction of the energy in SOI-1 from SOI. If the pulse arrival time

is in SOI-1, the summed energy is harmed both because the early energy is excluded from SOI and the SOI-1 energy is subtracted from the result. The early part of the pulse contains most of the electromagnetic energy. However, instead of just maximizing E_T in SOI, which is likely to put the rising edge in SOI-1, the pulse rising edge should be as close to the SOI start as possible, so the full early energy is captured. Figure 6 illustrates this with the fraction of TPs over 4 GeV decreasing as the phase scan progresses, while the average energy in SOI-1 is minimized at +2 ns, ensuring that the pulse rising edge is close to the start of SOI.



(a) Fraction of TPs over 4 GeV in SOI (TP TS2) vs. QIE phase offset. As the pulse moves later, the integrated charge drops. $i\eta = 4$ is shown and is representative of the HB.

(b) Average TP energy in SOI-1 (TP TS1) vs QIE phase offset, given that SOI $E_T > 4$ GeV. If the pulse is too early, see the pulse rising edge pushed into SOI-1. $i\eta = 4$ is shown and is representative of the HB.

Figure 6: Effects of pulse alignment shifts on the TP occupancy (fraction of towers over 4 GeV) and average TP energy in SOI-1. Both illustrate that the QIE phase offset of 0 produces ideal alignment, without causing the pulse to spillover into SOI-1.

4 Charge Weighted Time

Charge Weighted Time (CWT) is a method of estimating pulse timing from the energy distribution and has been used as a time-alignment variable in Run 2 and early Run 3. Charge weighted time is defined by:

$$\text{CWT} = \frac{\sum_i Q_i \cdot i \cdot 25}{\sum_i Q_i} \quad i = 0 \dots 7 \quad (3)$$

where i is the timeslice (each 25 ns), with $i = 3$ corresponding to the sample of interest (SOI).

However, CWT assumes that pulse shapes are the same across the whole detector, and is biased by late fluctuations in the hadronic shower. CWT prioritizes keeping the $\frac{SOI}{SOI+1}$ energy ratio consistent across the detector (since these are the two timeslices with the largest energy and thus largest impact on the CWT result), at the loss of sensitivity to the position of the rising edge. In contrast, for a given pulse shape, TDC peak and energy ratios are linear in time delay, and there is no assumption on pulse shape.

In addition, CWT uses charge measurements with a sampling frequency of 40 MHz, a granularity that is not small enough to pick up small changes in the pulse peak position. This leads to regions of CWT that have minimal sensitivity to pulse rising edge positioning, and is a drawback of CWT.

4.1 CWT Calibration Procedure from QIE Scan

TDC is a fundamentally accurate measure of pulse time that does not assume a uniform pulse shape, and the high precision of TDC is used to refine CWT. Given that the nominal (0 ns) reference point for the November and May QIE scans has ideal time alignment from TDC distributions, a set of corrections are

derived to allow for the use of CWT as an improved time estimator, after understanding the depth and QIE phase dependence. The plots in this section are from the November QIE phase scan.

CWT is computed for each depth and QIE delay, for pulses with over 5000 fC of total charge across the 8 timeslices. CWT has a depth dependence and is non-linear with respect to QIE delay. In particular, there is not a 1-1 correspondence between the change in QIE setting and CWT result, as when the QIE phase changes by 14 ns, CWT does not also change by 14 ns (the change in CWT is closer to 8 ns).




Figure 7: CWT from Equation 3, split by depth and QIE delay in HB, $1 \leq |i\eta| \leq 15$, and inclusive in $i\phi$.

In Figure 7, it is clear that Depth 1 behaves very differently than the higher depths. Depth 1 has lower CWT values, due to a cleaner pulse shape and fewer late time fluctuations. Due to the depth dependent non-linearity, choosing a set CWT value (89 ns in Run 2 and early Run 3) across all depths misaligns the pulse arrival times in HB. This is consistent with observations of time alignment performed with TDC methods. In this case, Depth 1 pulses are set to arrive late, while Depth 4 pulses arrive very early, leading to large position dependent alignment offsets within the detector.

The difference in late time fluctuations by depth is clear from the $\frac{SOI+1}{SOI}$ and $\frac{SOI-1}{SOI}$ plots, in Figure 8a and b. Depth 1 has less charge in SOI+1, due to faster falling tails. As CWT assumes the same pulse shape for each depth, this means that achieving the same CWT value for each depth requires pushing the Depth 1 pulse very late in SOI.

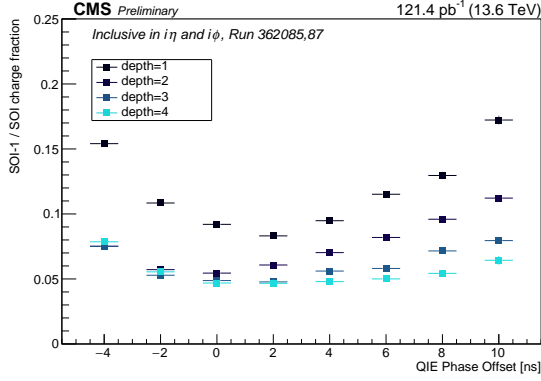
A re-weighting and calibration procedure is derived to correct for the pulse shape effects in CWT. First, the pulse shape used to calculate CWT is re-weighted. This is done by introducing a vector of weight factors (depth and timeslice dependent) into CWT, $w_{i,depth}$:

$$CWT_{depth} = \frac{\sum_i Q_i \cdot i \cdot 25 \cdot w_{i,depth}}{\sum_i Q_i \cdot w_{i,depth}} \quad i = 0 \dots 7 \quad (4)$$


and after re-weighting, a calibration (slope correction) is applied. The correction factors are determined from a linear fit to CWT_{depth} , giving the slope and intercept m_{depth} and b_{depth} . The final calibrated CWT is:

$$\text{Calibrated } CWT_{depth} = \frac{1}{m_{depth}} \cdot CWT_{depth} + b_{depth} \left(1 - \frac{1}{m_{depth}} \right) \quad (5)$$


where the CWT at QIE phase of 0 is unchanged, and now there is a one-to-one correspondence between changes in the QIE phase setting and the CWT result.



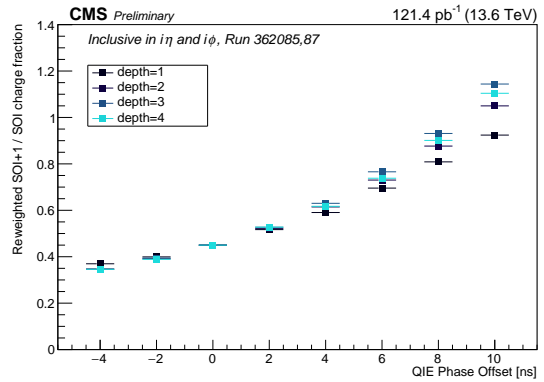
(a) $\frac{SOI-1}{SOI}$ from fC charge in each timeslice



(b) $\frac{SOI+1}{SOI}$ from fC charge in each timeslice



(c) $\frac{SOI-1}{SOI} \cdot w_{i=2,depth}$ from fC charge in each timeslice, which is the reweighted $\frac{SOI-1}{SOI}$ ratio



(d) $\frac{SOI+1}{SOI} \cdot w_{i=4,depth}$ from fC charge in each timeslice, which is the reweighted $\frac{SOI+1}{SOI}$ ratio

Figure 8: Plots of the original (a, b) and re-weighted (c, d) SOI±1 to SOI energy ratio, $1 \leq |i\eta| \leq 15$, and inclusive in $i\phi$.

$w_{i=2,depth}$ and $w_{i=4,depth}$ (SOI-1 and SOI+1) are determined from $\frac{SOI-1}{SOI}$ and $\frac{SOI+1}{SOI}$, while all other weights are 1, resulting in:

$$w_{i,depth} = [1, 1, w_{i=2,depth}, 1, w_{i=4,depth}, 1, 1, 1] \quad (6)$$

$$w_{i=2,depth} = \frac{\text{ideal}(i=2)}{\frac{SOI-1}{SOI}} \quad (7)$$

$$w_{i=4,depth} = \frac{\text{ideal}(i=4)}{\frac{SOI+1}{SOI}} \quad (8)$$

where $\text{ideal}(i=2)$ and $\text{ideal}(i=4)$ are the selected ideal charge ratios between $\frac{SOI-1}{SOI}$ and $\frac{SOI+1}{SOI}$, respectively, at QIE phase of 0. Deriving the weights based on $\frac{SOI\pm 1}{SOI}$ ratios normalizes the pulse shapes to be uniform in shape across depths, to align with the assumptions of CWT. Given these weights, the SOI-1 and SOI+1 plots are re-weighted, in Figure 8c and d. This is done with selected ideal ratio values of $\text{ideal}(i=2) = 0.05$ and $\text{ideal}(i=4) = 0.45$. The determined weight values are:

$w_{i,depth}$	Depth 1	Depth 2	Depth 3	Depth 4
i=2	0.543	0.918	1.03	1.07
i=4	1.25	1.04	1.02	0.991

(a) For $i\eta = -15$ to 15, inclusive.

$w_{i,depth}$	Depth 1	Depth 2	Depth 3	Depth 4
i=2	0.451	1.01	1.17	1.30
i=4	1.34	1.04	1.08	1.05

(b) For $i\eta = -8$ to 8, inclusive.

$w_{i,depth}$	Depth 1	Depth 2	Depth 3	Depth 4
i=2	0.382	0.811	0.827	1.02
i=4	1.36	0.957	1.01	0.986

(c) For $i\eta = -1$ only.

Table 1: Weights $w_{i=2,depth}$ and $w_{i=4,depth}$, determined with $\text{ideal}(i=2) = 0.05$ and $\text{ideal}(i=4) = 0.45$ at QIE offset 0ns, in collisions data. All other values in the weights vector are 1.

	Depth 1	Depth 2	Depth 3	Depth 4
Slope	0.42	0.49	0.56	0.57

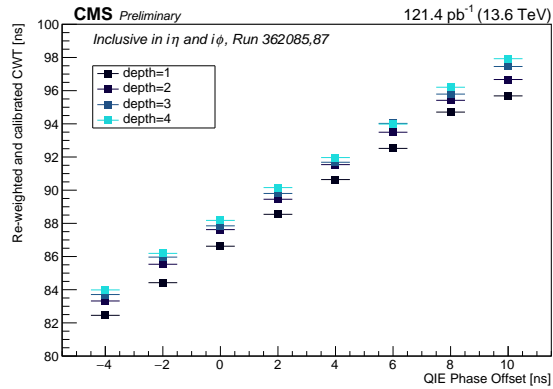
Table 2: Slopes for the collisions scan analysis.

Using the re-weighting and calibration, a final CWT is determined, plotted in Figure 9. This calibrated CWT is now linear in QIE delay, and much more consistent per depth. The re-weighting and linearization procedure for CWT improves its capabilities as a timing variable, as re-weighting accounts for pulse shape effects, and linearization calibrates CWT to give the expected response to a change in QIE phase.

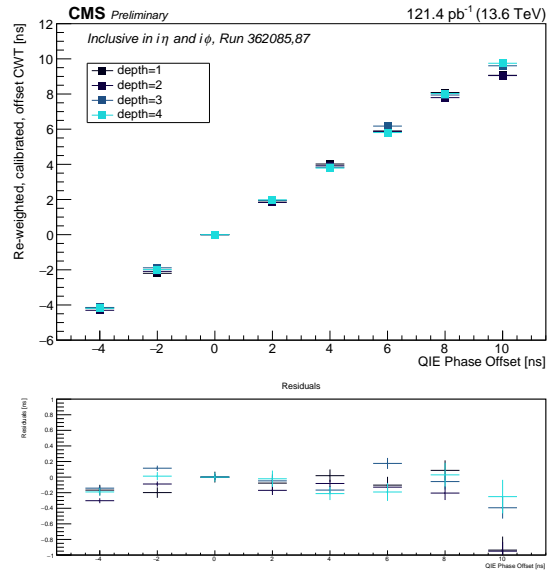
4.2 Timing in LED Runs

The behavior of the HCAL pulse shape and timing readouts during an LED scan was understood from a local run in November 2021. This is a similar scan as the timing scans in 13.6 TeV collisions used for HCAL alignment, but is from a HCAL data-taking run where an LED illuminates each SiPM. The Delay25¹, a programmable CMOS delay line that can be set to 0.5 ns intervals scans the timing of the LED pulse throughout the run. This scan was performed with a monochromatic LED pulse.

¹<https://proj-delay25.web.cern.ch>



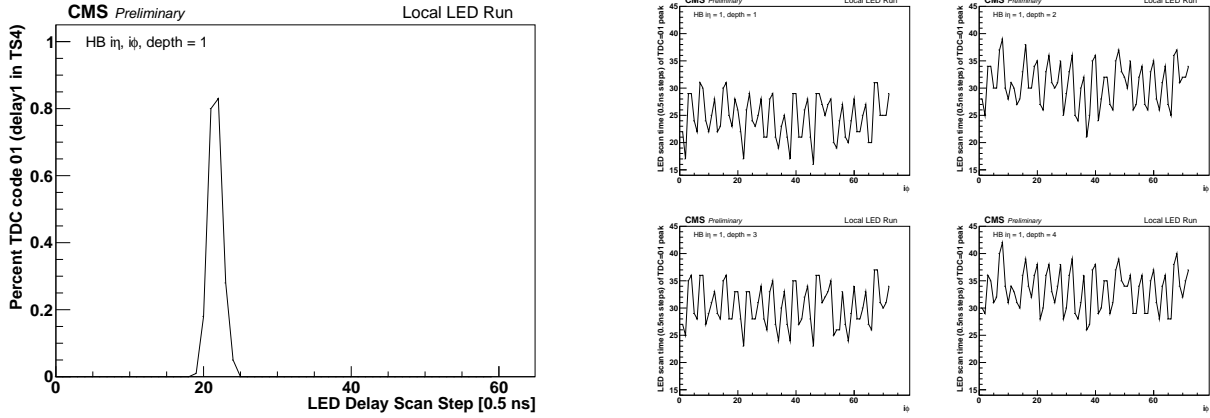
(a) Re-weighted and calibrated CWT.



(b) Re-weighted and calibrated CWT, with the depth dependent offset applied to shift the distributions to 0 ns at QIE phase offset of 0. The residual plot from $x = y$ line is also shown.

Figure 9: Re-weighted and calibrated CWT from Equation 5, split by depth and QIE delay in HB. The depth dependent offset is applied in b to shift the CWT distributions to intersect 0 at a QIE phase offset of 0 ns. The residuals show that the re-weighting and calibration procedure achieved agreement with QIE phase changes to within 0.5 ns over the range of the QIE scan.

The TDC = 01 peak is particularly helpful for aligning channels since it is a 1 ns wide bin. As seen in Figure 10a this peak is just 1-2 ns wide (resolution limited by the 1 ns step size of the LED scan) for a single channel.



(a) Percent of TDC = 01 in the local LED scan. The peak is clearly visible and spans ≈ 1 ns.

(b) LED delay scan setting (in 0.5 ns steps) where TDC = 01 peak occurs, vs. $i\phi$ for $i\eta = 1$ and separated by depth.

Figure 10: Plots of the TDC = 01 peak during an LED delay scan for individual channels.

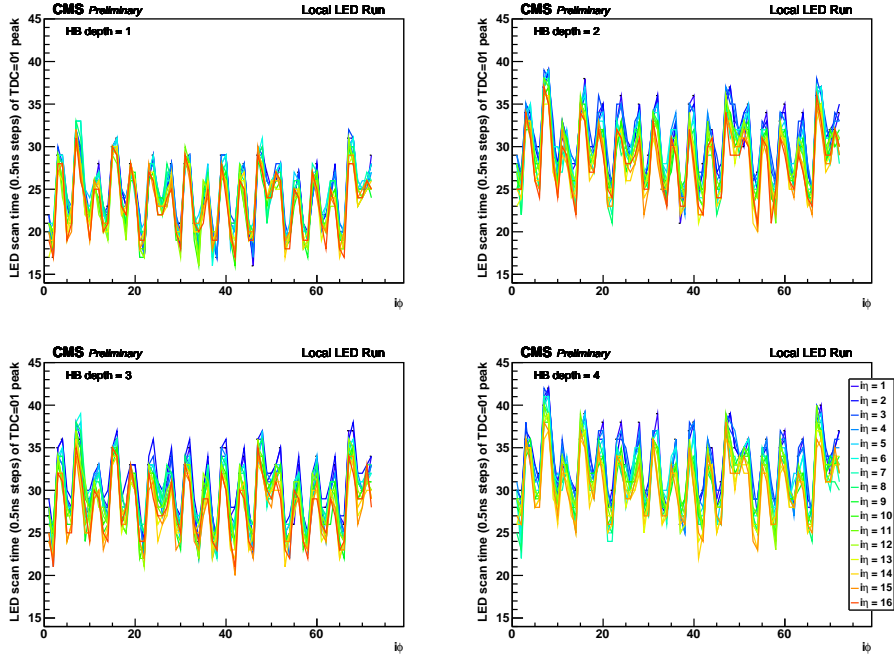


Figure 11: Overlay of scan time corresponding to TDC = 01 peak during an LED delay scan by $i\eta$. More central $i\eta$ positions are in blue, while larger ones are in orange. Plots are split by depth.

When split by $i\phi$, a pattern of 2 high - 2 low LED delay settings (at the time of TDC=01 peak) is observed, as seen in Figure 10b. Each RBX covers 4 $i\phi$, and contains 4 RMs, each covering 1 $i\phi$. There are 2 ngCCMs per RBX, each providing clock and control to 2 neighboring RMs, as seen in Figure 12. The cable length difference between the calibration unit and each half backplane causes the 5 ns delay seen in Figure 11, and this is an artifact of the LED scan - ie, this will not be seen in collisions as only the LED signal must travel

along this cable. The $i\phi$ differences are accounted for in the analysis of the LED data, by grouping together $i\phi$ ranges with the same cable length effects.

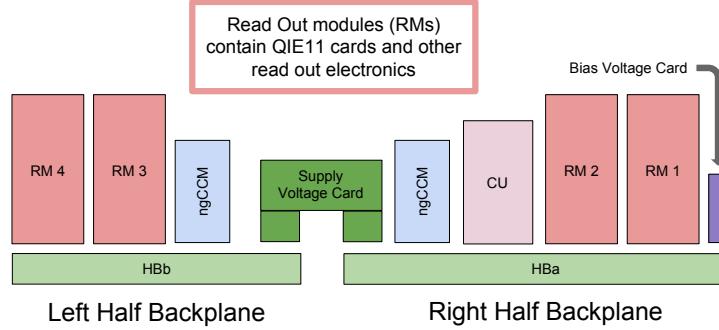
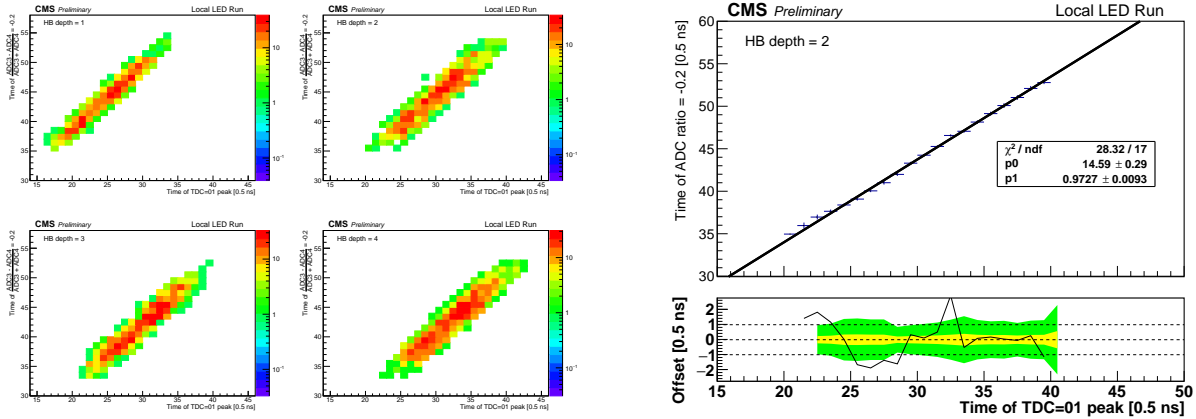


Figure 12: Diagram of the components in a single RBX: 4 RMs, 2 ngCCMs, a bias and supply voltage card, and a calibration unit (CU). Diagram from [16].



(a) Energy ratio and TDC peak time correlation.

(b) Linear fit on depth 2 of energy ratio and TDC peak time correlation, with residuals. Agreement between the two pulse time measurements is within 1.5 ns.

Figure 13: Correlation between LED scan times at which TDC = 01 peak was first reached, and when $\frac{ADC3-ADC4}{ADC3+ADC4} = -0.2$ was first reached. The 1-1 correlation shows that TDC timing is equivalent to pulse shape timing, and has a linear correlation with QIE phase scan change. This is for all $i\eta$, $i\phi$, and separated by depth.

Figure 13 demonstrates that the time alignment of the detector using TDC is equivalent to the previous method (ADC pulse shape). There is a pulse shape timing correlation between TDC (01 peak) and ADC ratio, showing that ADC ratios and TDC results both give precise and correlated timing information. The plot is the LED scan time of TDC=01 peak vs. time of $\frac{ADC3-ADC4}{ADC3+ADC4} = -0.2$. For a given pulse shape, TDC peak and ADC energy ratios are linear in time delay. This means that adjusting the HCAL clock (QIE phase) linearly adjusts the TDC peak location, which is the basis of the TDC-based time alignment. For HCAL time alignment, similar timing scans in 900 GeV and 13.6 TeV collisions are used.

A closure test of the CWT procedure is done with an LED run during the QIE phase scan, presented in the following section.

4.3 Validation of CWT Calibration using LED Data

After understanding the response of the detector during the LED timing scan described in the previous section, a QIE phase scan was taken during another local LED run to validate the CWT calibration and re-weighting procedures. This validation run used QIE phases set such that LED data “looks like” collisions timing – the pulse arrival time is close to the start of the SOI, and TDC=00 peak is at a phase offset of 0 ns. These phases are distinct from those that align pulses in collisions, due to LED signals having a slower rise time.

Importantly, LED signals have a different pulse shape than collisions pulses. As the LED is shined directly onto the SiPMs, the effects of hadron showering and dispersion contributing to the pulse shape differences by depth and $i\eta$ are removed. Since the LED provides a consistent pulse shape across the detector, all weights ($w_{i,depth}$) are expected to be ≈ 1 in the LED data. Thus, the LED data provides a closure test of the CWT procedure. In the LED phase scan, the cable length difference between the calibration unit and each half backplane must be accounted for, as this causes 5 ns delay between $i\phi$ groups.




Figure 14: Fraction of TDC codes throughout a QIE phase scan with local LED data, for $i\eta = -1$, and $i\phi = 1, 2, 5, 6, 9, 10 \dots 69, 70$. The consistent prompt TDC code distributions across depths demonstrate that these phases align the LED pulse arrival times well. QIE phase of 10 ns is excluded from the CWT analysis, as here the start of the prompt peak from the next BX is visible. This is due to the slow rise time of the LED.

The LED alignment worked particularly well for $-4 < i\eta \leq -1$ and $i\phi = 1, 2, 5, 6, 9, 10 \dots 69, 70$, and separately $1 \leq i\eta < 4$ and $i\phi = 3, 4, 7, 8, 11, 12 \dots 71, 72$, resulting in alignment to within 2 ns, shown in Figure 14. The standard CWT from Equation 3 is calculated for the LED pulses, and plotted by depth and QIE phase in Figure 15. The un-adjusted CWT is already linear in QIE phase and has no significant depth dependence. The adjustments from applying the derived CWT calibration and re-weighting procedure are minimal.

The depth-dependent weights are within 10% of 1 (excepting depth 1, for $i = 2$) and the slopes are within 5% of 1, listed in Tables 3 and 4. Additionally, after the re-weighting and calibration procedure, the LED data has an agreement with the QIE phase change to within 0.2 ns. This validates the CWT re-weighting and calibration procedure by demonstrating that when run on a uniform pulse shape, the weight and slope

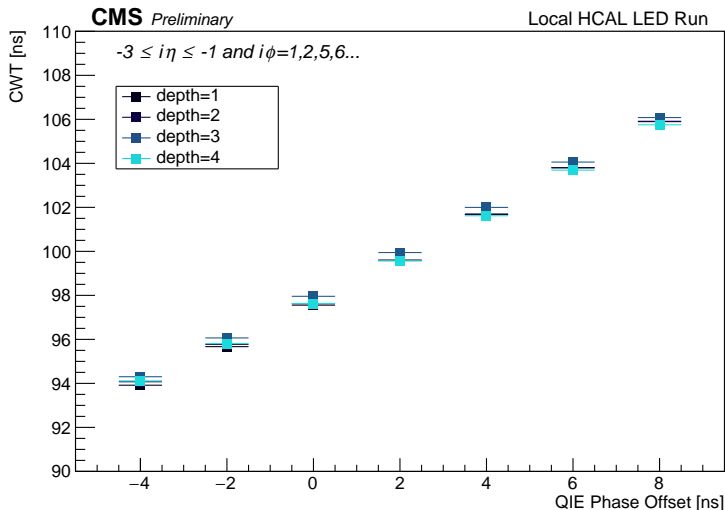


Figure 15: CWT from LED data, for $i\eta = -3$ to -1 , and $i\phi = 1, 2, 5, 6, 9, 10 \dots 69, 70$.

do not significantly differ from 1. Note that SOI-1 ($i = 2$) has very minimal charge in the LED run, and thus weights different from 1 do not significantly impact the results of the re-weighting procedure.

$w_{i,depth}$	Depth 1	Depth 2	Depth 3	Depth 4
$i=2$	1.11	1.00	0.92	1.00
$i=4$	0.95	0.98	0.95	1.03

Table 3: Weights $w_{i=2,depth}$ and $w_{i=4,depth}$, determined with $ideal(i = 2) = 0.005$ and $ideal(i = 4) = 1.0$ (at QIE phase offset 0 ns) from LED scan data, for $i\eta = -3$ to -1 , excluding $i\phi = 3, 4, 7, 8 \dots$. All other values in the weights vector are 1.

	Depth 1	Depth 2	Depth 3	Depth 4
Slope	1.05	1.00	1.01	0.96

Table 4: Slopes for the LED scan analysis for $i\eta = -3$ to -1 .

5 Energy Scale Changes

5.1 Energy Response

Altering the HCAL phase alignment affects the energy response, as more or less energy from the pulse is integrated. A pulse that arrives early (in SOI-1) will have a decreased integrated energy, as the rising edge will not be captured in the integral. Conversely, a pulse arriving very late will also be reduced, as a significant portion of the tail will fall in SOI+1, and therefore not contribute to the energy measurement. Ideally, the full prompt pulse will be captured, which the TDC-based alignment method is designed to do. To minimize the downstream effects on the L1 trigger and jet energies from an HCAL alignment adjustment, a corresponding set of response corrections is concurrently deployed along with the TDC-based time alignment. The response corrections are values the measured energy is scaled by and are designed to take into account detector effects, such that uniform energies are reported to the trigger system.

The change in average pulse energy vs. phase change is not a precise enough variable from which to estimate the response correction changes. This is expected, as the energy is from a falling exponential distribution, and its average will not change drastically with phase. Instead, for higher sensitivity, the hit occupancy

is used. This has a much stronger dependence on phase and allows for a direct estimation of response corrections.

5.2 Hit Occupancy Method

The change in the response corrections is estimated from a hit occupancy method. The hit occupancy is defined as

$$\text{Occupancy} = \frac{N_{hits}(\text{over } 4 \text{ GeV MAHI energy})}{N_{hits}(\text{over } 0 \text{ GeV MAHI energy})} \quad (9)$$

The hit occupancy will be maximized slightly before the optimal phases (at around -5 ns to -2 ns), as this is where the integrated pulse is largest as the full rising edge is included. The hit occupancy slightly decreases when the pulse arrives right at the start of the SOI, however, the prompt part of the signal is fully captured, which is best for energy reconstruction. Thus, this is a more accurate and less noisy alignment setting, though the hit occupancy is not maximized.

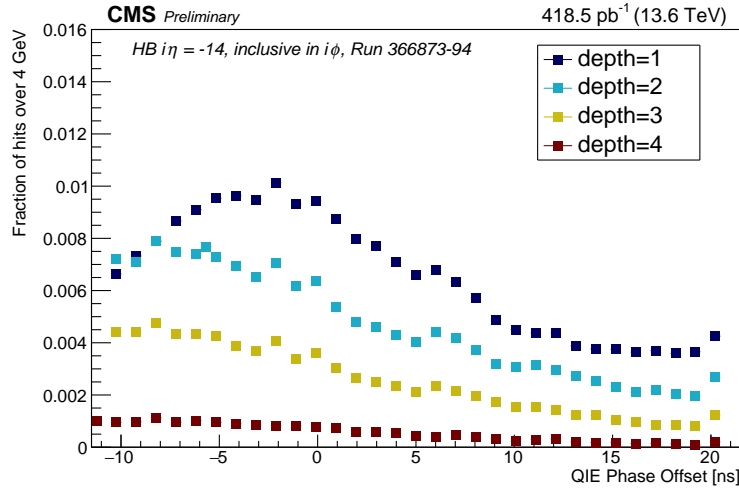


Figure 16: HB hit occupancy ≥ 4 GeV in MAHI energy for $i\eta = -14$ (representative of other regions).

Once the hit occupancy as a function of QIE phase offset is found, the factor needed to scale this by to achieve flat occupancy is found. These factors are the “energy scales“ to give uniform occupancy vs. phase. Importantly, for a given change in the QIE phase, these energy scales are how much the energy response must be scaled up or down by to give a consistent energy response.

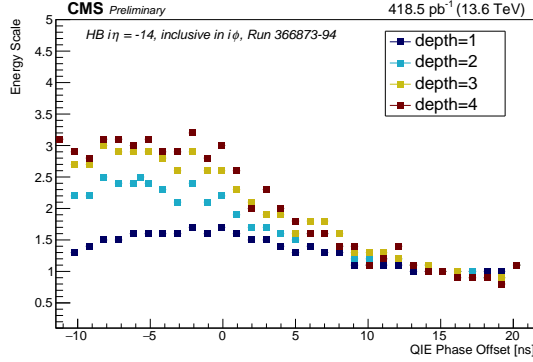
The scaled occupancy is computed as:

$$\text{Occupancy} = \frac{\int_4^\infty \text{MAHI energy}}{\int_0^\infty \text{MAHI energy}} \quad (10)$$

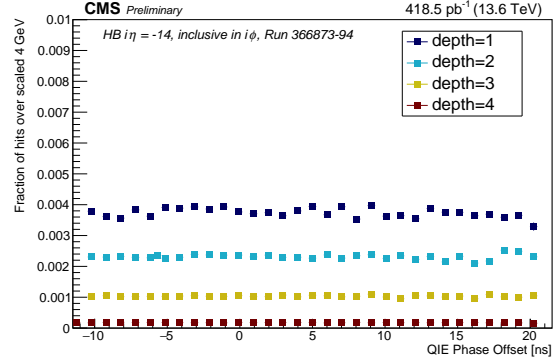
$$\text{Scaled Occupancy} = \frac{\int_{4 \cdot x}^\infty \text{MAHI energy}}{\int_0^\infty \text{MAHI energy}} \quad (11)$$

Scaling the energy down by x is equivalent to scaling the energy value of the integral lower bound up by x to count hits above a 4 GeV scaled energy. Thus, for each detector region, the “energy scale” values are scanned to determine what gives the same occupancy as the reference value.

$$\text{Energy}_{new, y \text{ ns}} = \frac{\text{Energy}_{initial}}{\text{scale}_{y \text{ ns}}} \quad (12)$$



(a) HB energy scale for MAHI energy to give flat hit occupancy, from Equation 12.



(b) HB scaled hit occupancy, $Energy_{new,y ns}$ from Equation 13.

Figure 17: HB energy scale and scaled occupancy for $i\eta = -14$, inclusive in $i\phi$.

The energy scale shows how each phase can be scaled to a reference phase and allows to predict changes in energy response when a given QIE phase change is made. For instance, moving to a new phase at y ns from an existing phase at x ns, while knowing the energy scale at all values of the phase scan means that the change in energy will be:

$$Energy_{old,x ns} = \frac{scale_{x ns}}{scale_{y ns}} \cdot Energy_{new,y ns} \quad (13)$$

The change in the response corrections when moving from QIE phases at x to y is then given by the factor $\frac{scale_{x ns}}{scale_{y ns}}$. The derived factors are expected to make the detector energy response more uniform, with the following trends seen in the barrel region:

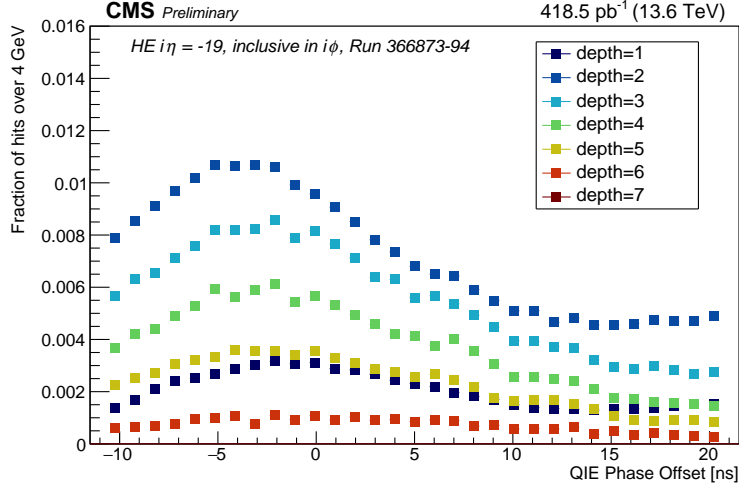
- Depth 1: Signal increases by 5-8%, as with the 2022 phases the pulse was arriving very late
- Depth 2 and 3: High $i\eta$ has a slight decrease in the signal, because of the very large late time component
- Depth 4: In the central barrel, there is a slight decrease in the signal as with the 2022 phases the pulse was placed right at the clock edge and now is moved a few ns after the clock edge. At high $i\eta$, the signal is increased, because with the 2022 phases the pulse was arriving very early (into SOI-1, so energy was missed in the integrated signal of SOI).

The HB energy response corrections that correspond to the 2023 TDC-based QIE phase alignment are listed in Table 5.

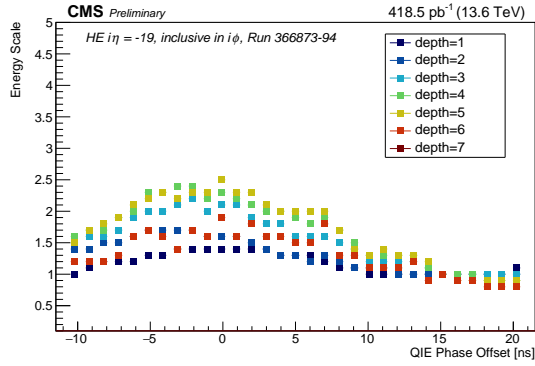
5.3 Validation of Method in HE

The hit occupancy method to estimate changes in the energy response corrections is new, and before deployment in HB this method was validated in HE. HE had a QIE phase adjustment implemented between 2022 Era D and Era E, with Era D QIE phases at an offset of about 4-5 ns relative to the May QIE phase scan. Since the May 2023 QIE phase scan included both HB and HE, this data and the same method described in the earlier section are used for validation.

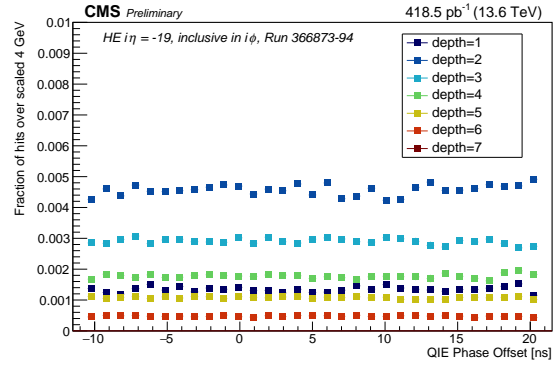
The predicted changes generally reduced the response corrections by 10-20%, which is expected as the new HE alignment increases the energy response by placing more of the prompt signal within the integration window. This is consistent with the HE response corrections deployed (and obtained by an independent method). The two methods are within 10% (20%) of each other for 70% (80%) of the depth and $i\eta$ regions. The best agreement is in depths 2-5, as higher depths are likely affected by HE containment corrections as well. This is deemed sufficiently in agreement to proceed with the HB response corrections derived from the hit occupancy methods.



(a) HE hit occupancy ≥ 4 GeV in MAHI energy.



(b) HE energy scale for MAHI energy to give flat hit occupancy, from Equation 12.



(c) HE scaled hit occupancy, $\text{Energy}_{new,y_{ns}}$ from Equation 13.

Figure 18: HE hit occupancy, energy scale, and scaled occupancy for $i\eta = -19$ and inclusive in $i\phi$ (representative of other regions). $i\eta = -19$ has 6 depths.

$ i\eta $	Depth 1	Depth 2	Depth 3	Depth 4
-16	1.00	1.13	1.20	-
-15	0.94	1.16	1.08	1.07
-14	0.94	1.09	1.15	0.97
-13	0.95	0.88	0.91	0.85
-12	0.84	0.96	1.00	0.94
-11	0.85	0.89	0.87	1.03
-10	0.95	1.04	1.00	1.03
-9	0.89	1.00	0.91	1.03
-8	0.89	1.00	0.91	1.00
-7	0.94	1.00	1.03	0.94
-6	1.00	0.96	1.00	1.06
-5	0.94	0.92	0.97	1.03
-4	1.00	0.92	1.00	0.97
-3	0.94	0.92	1.00	1.06
-2	1.00	0.93	1.00	1.06
-1	1.00	0.88	0.96	1.00
1	1.00	0.89	1.00	1.00
2	1.07	0.96	1.00	0.94
3	0.94	0.93	0.97	1.00
4	0.94	0.96	1.00	1.07
5	1.00	0.92	1.03	1.06
6	0.94	1.00	0.94	1.06
7	0.94	0.96	1.00	0.95
8	0.94	1.00	1.03	1.06
9	0.94	1.00	0.94	0.92
10	0.94	1.00	0.97	0.94
11	0.96	0.93	1.06	0.97
12	0.94	1.08	1.03	1.00
13	0.89	1.04	1.10	0.97
14	0.94	1.12	1.13	0.93
15	0.94	1.05	1.04	1.13
16	1.00	1.13	1.19	-

Table 5: HB energy response adjustments with the change from the October 2022 phases to the June 2023 phases. This is derived from the hit occupancy method.

Further validation comes from the standard response correction analysis using the E_γ dataset in 2023C and 2023D data for IsoTrack calibration. This analysis, performed by collaborators in the HCAL DPG group, validates that the predicted HB compensating response corrections deployed with the new time alignment agree with the IsoTrack response correction method to within 5%.

6 HCAL Timing in L1 Long-Lived Particle Trigger

The novel L1 delayed jet trigger relies upon the new HCAL capabilities and re-programmable firmware to significantly improve the LLP physics reach. The trigger identifies delayed hadronic jets using TDC and ADC information as input. The full information from each HCAL depth segmentation layer cannot be sent to the trigger due to bandwidth constraints. Thus, per-tower timing and depth dependent quantities are computed and packed into bits sent to the L1 trigger. Ultimately, the trigger utilizes a timing feature bit with 1 bit per trigger tower (TT) transmitting information from HCAL through L1.

The new LLP trigger is sensitive to both time delays resulting from the path length difference traveled by a LLP and its decay product and from the slow velocity of the LLP. In addition, the trigger maintains

sensitivity to LLPs decaying within the detector volume, creating large energy deposits in higher layers of the calorimeter (Figure 1).

The Phase-1 HCAL upgrades provide handles to trigger on displaced objects, as energy deposits and timing information are recorded at each cell. This delayed jet trigger is based on the delayed timing and energy deposits of pairs of energetic hits that are indicative of a LLP decay in or before the HCAL. Per trigger tower (a group of multiple depths at the same $i\phi$ and $i\eta$) a LLP flag is defined, either by:

- including an energetic delayed cell (over timing value, ie. TDC=01 or 10, and energy > 4 GeV) and not including any energetic prompt cells (in the prompt range, ie. TDC=00, and over 4 GeV)
- including a large energy deposit in the higher depths (> 5 GeV) and very little energy in the lower depths (< 1 GeV)

The trigger forms L1 delayed jet objects, which are required to have at least 2 LLP-flagged trigger towers within the 9x9 jet region and pass jet energy requirements. To control rates at the L1 trigger, event energy thresholds are also required.

6.1 uHTR Specifications

The calorimeter trigger towers (TTs) in the HCAL are comprised of physical calorimeter towers grouped together in depth. The HCAL trigger primitives (TPs) are formed by combining the information from a single calorimeter TT, such as summing the energies deposited in each depth. In addition to the standard TP generation in HBHE, six “extended bits” or “feature bits” of information are generated and transmitted to the Level-1 trigger. For the delayed jet trigger, these feature bits encode information about (i) the longitudinal shower profile data, and (ii) the shower time data constructed from the TDC information available in each constituent channel of the trigger primitive. Thus, feature bits mark delayed arrival times and distinctive energy deposits that are characteristics of signals from long-lived exotic particle decays.

The HCAL μ TCA Trigger and Readout modules (uHTR) firmware sets the six feature bits per trigger tower and these are sent to L1 for the delayed jet trigger. The full details are given in the *Specifications for HCAL uHTR Firmware* [14]. Currently, the last two bits (bit[4] and bit[5]) are reserved. The first four bits carry information relating to the arrival time and depth profile of the shower, and will be used to inform the delayed jet trigger:

- bit[0] depth flag
- bit[1] prompt timing flag
- bit[2] delayed timing flag
- bit[3] very delayed timing flag
- bit[4] reserved
- bit[5] reserved

Following the uHTR, to stay within bandwidth constraints in the Layer-1 and Layer-2 trigger pathways, one bit per trigger tower will be used to send the delayed or displaced LLP signal information to Level 1 Global Trigger, where a delayed/displaced jet flag is set.

6.1.1 LLP flag – Depth bit

To set the depth-based fine-grain bit, there must be significant energy deposited in the higher depths of the calorimeter, while little energy is deposited in lower depths. If energies in both depths 1 and 2 are under 1.0 GeV E_T (16 ADC), and either energy in depth 3 or 4 is at least 5.0 GeV E_T (80 ADC), then the depth-based fine grain bit is set for HB. This is stored in `bit[0]` of the fine-grain bits.

- bit[0] depth flag (< 1 GeV in each of depths 1 and 2, and ≥ 5 GeV E_T in at least one of depth 3+)

In the HE implementation, which was designed but not implemented for 2022 and 2023 data taking, energies in both depths 1 and 2 must be under 1.0 GeV (16 ADC), and the energy in any of depths 3 - 7 is at least 5.0 GeV (80 ADC) for the depth bit to be set.

6.1.2 LLP flag – Timing bits

Feature bits are set on a per-trigger tower basis. To set the three timing-based fine grain bits, cells must both be above the energy threshold (≥ 4.0 GeV E_T , or linearized ADC ≥ 64 , using a single bit in the energy LUT) and have a valid TDC value (00, 01, or 10). Three fine-grain bits are set to reflect the presence of an energetic cell within the tower with a given timing range.

Per tower, bit[1], bit[2], and bit[3] of the fine grain bits are set based on:

- bit[1] is set if there is a prompt cell (TDC=00) within the tower above the energy threshold
- bit[2] is set if there is a delay 1 cell (TDC=01) within the tower above the energy threshold
- bit[3] is set if there is a delay 2 cell (TDC=10) within the tower above the energy threshold

If the tower has at least one delayed cell and no prompt cells, the trigger tower is flagged as a LLP (feature bit = 1). Towers with prompt cells are veto-ed regardless of the presence of delayed cells in the tower. The timing flag bit logic is `!bit[1] & (bit[2] | bit[3])`.

6.1.3 LLP flag – Energy requirements

The firmware implementation and configuration preserves the option of having different energy requirements for HB and HE, as the LUTs used to implement the energy requirements can be adjusted without firmware reprogramming. The uHTR Group 0 LUT is used to impose energy requirements for the set bits. Full details are in Section 3.5 of the uHTR specifications. 4 bits of the energy LUT are required - bits 12, 13, 14, and 15. Bits 14 and 15 simply set energy thresholds for the delayed and prompt cells, while bits 12 and 13 are depth dependent, and are used in the depth profile determination.

These bits are set based on the linearized ADC energies. These are the 0-255 raw values after they have been pedestal subtracted and gain corrected, and have a direct correlation to transverse energies in GeV.

6.2 L1 LLP Jet Trigger Algorithm

L1 Layer-1 performs the reduction from 6 HCAL feature bits to 1 LLP flagged bit per tower, and forwards this single bit to Layer-2. In order to send LLP flagged trigger towers to Layer-2, the logical OR of depth and timing flagged towers is taken, such that a single bit is set when the tower is either tagged due to depth or timing (or both). The 6:1 bit logic is

`Depth OR Timing = Depth OR ((Not prompt) AND (Delay1 OR Delay2))`

Following the bit specifications listed above, this is:

`bit[0] | (!bit[1] & (bit[2] | bit[3]))`

The L1 delayed jet trigger is based on a pair of energetic hits that are indicative of a LLP decay before or within the hadron calorimeter. Thus the delayed jet objects are required to have at least 2 LLP flagged trigger towers within the 9x9 jet region, and have a jet energy over $E_T = 40$ GeV (50, 60 GeV). To control rates at the Level 1 trigger, event energy over 120 (160, 200) GeV is required. An illustration of the data pathway from the HCAL through L1 is in Figure 19.

Current triggers based on H_T or multi-jet triggers impose high hadronic energy thresholds in order to contain the rate [4]. Many LLP events present a difficulty for triggering, as the hadronic jets are relatively low energy and the H_T is low [4]. By designing a calorimeter trigger based on time delays and lowering the H_T threshold, we can increase sensitivity to LLP models with long lifetimes. The delayed time signature will allow for background rejection as well, as the LLP events are expected to have a much more significant time delay as compared to the rest of the bunch crossing and background events.

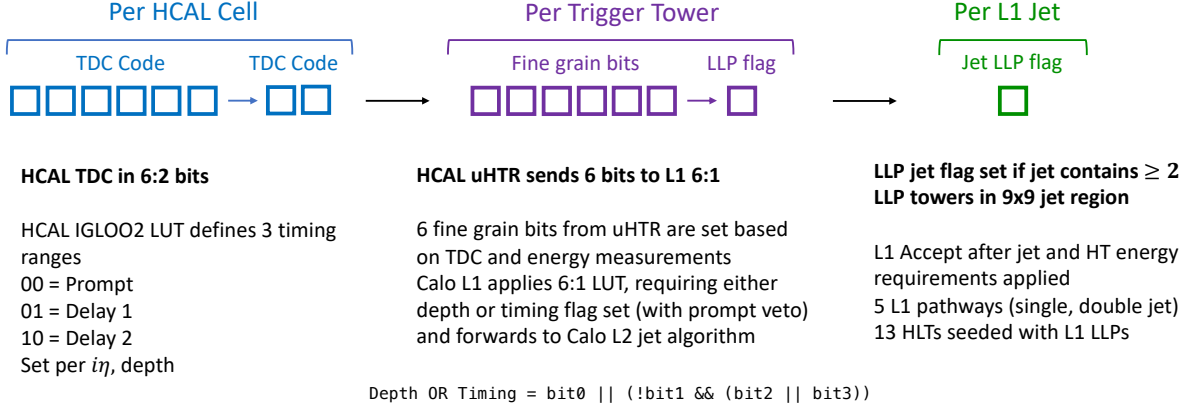


Figure 19: Dataflow from the HCAL, through the HCAL uHTR, L1 Layer-1 LUT, and Layer-2 jet algorithm.

6.2.1 uGT Jet Bits

At Layer-2, the number of flagged trigger towers in a L1 jet region is summed (using the 9x9 towers around the center of a L1 jet). When there are at least two flagged trigger towers within the 9x9 jet region, `jet.HwQual` is set to flag the jet. A full emulator in CMSSW was written and validated with firmware.

At the global trigger (uGT), there are five reserved and quality bits for a jet object (bits 27-31). Bit 27 (a quality bit) is used to flag a jet as delayed based on HCAL timing and depth profiles that are indicative of a LLP decay. This bit is referred to as “DISP”. When this bit is set to 1, the jet has been tagged as a LLP jet.

L1 jet word bit(s)	Parameter stored
31..30	spare
29..28	quality flags
27	DISP
26..19	φ
18..11	η
10..0	E_t

Table 6: Data structure of a jet object, modified from “Scales for inputs to uGT”, Table 11 [13], and “CaloL2 input/output specifications”, Page 4 [8].

An optional cut on this bit is used in combination with the requirement that jet $E_T \geq 40$ GeV in order for the jet to be flagged as delayed at L1. In addition, the L1 event HTT is required to be at least 120 GeV, with at least one delayed jet found in order to set the L1A signal for the delayed/displaced jet trigger.

6.2.2 Trigger Menu Seeds

In the trigger menu, a new LLP delayed/displaced jet seed is defined using the displaced jet bit along with jet E_T requirements. The DISP type cut on a jet is done with a 1-bit LUT, resulting in 2 states, [0] and [1]. In TME, the displaced jet is `JET-DISP_LLP = [1]` and a non-displaced jet is `JET-DISP_noLLP = [0]`. Jet E_T requirements from 40-70 GeV are used in the L1 trigger in various pathways. A LLP single jet seed to trigger on a jet as low as 40 GeV is:

$$\text{SingleLLPJet40} = \text{JET40}[\text{JET-DISP_LLP}]$$

If η restrictions are added, this would become `SingleLLPJet40 = JET40[JET-ETA_2p52, JET-DISP_LLP]`. In the menu, the AND of the LLP jet seed with low threshold HTT seeds is taken. For example, this seed is used in conditions such as:

L1_HTT120_SingleLLPJet40 = JET40[JET-DISP_LLP] AND L1_HTT120

This seed indicates that a single jet as low as 40 GeV will trigger as long as the displaced bit is set and the event H_T is over 120 GeV. A selection of seeds with tighter jet and event energy requirements along with a dijet LLP seed are added to the trigger menu.

7 Monte Carlo Data Injection on Front-End Electronics

It is possible to inject user-specified data into the RAM buffers of the uHTR FPGAs to effectively run Monte Carlo directly on detector front-end electronics. The data is injected into the input stage of the uHTR front-end FPGAs and replaces the data packets received from the QIEs (i.e. the bits for ADC energy, TDC timing, and TTC codes) per bunch crossing, per detector cell. The injected data is run through the full trigger and DAQ paths of the FPGA and trigger primitive outputs are propagated to L1 as real data (Figure 20). These mock trigger primitives are useful for a variety of firmware and trigger diagnostics such as algorithm tests and packet alignment checks. Localized hits can be manufactured in detector coordinates to represent jets, with pulse shapes and timings to mimic long-lived particle decays. These methods were used at P5 to test L1 and HLT trigger pathways for the LLP trigger.

If extended to a full simulation workflow, digitized calorimeter hit collections from detector simulations can be mapped to uHTR detector coordinates and repackaged into bunch crossings according to the uHTR packet specifications (Figure 21), allowing for the direct representation of MC events as genuine hardware events in the HCAL. In operation, care must be taken to limit MC bunch crossing rates to comply with L1 rate limits.

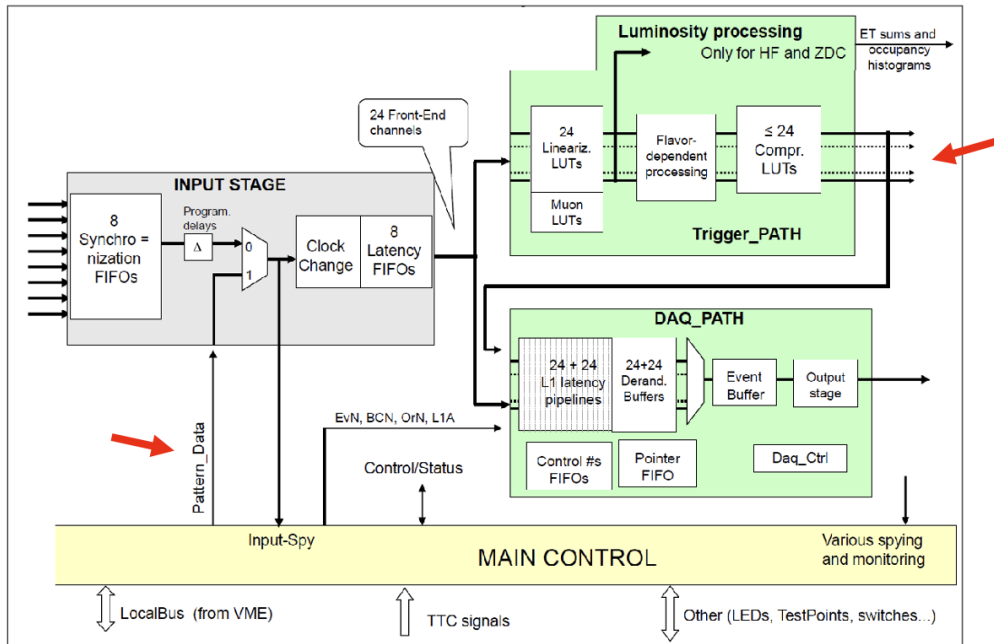


Figure 20: Schematic diagram of the FPGA showing input stage, trigger and DAQ pathways, and outputs [9]. Mock data in RAM is injected (left red arrow) and runs through both pathways as real data. Trigger primitive output is indicated by the red right arrow. Note: the schematic shown is from the previous generation HTR module; no same diagram is available for the current uHTR. The HTR schematic shown is adequate as a conceptual representation.

8-Channel Format

Byte	7	6	5	4	3	2	1	0
0	K28.5 Comma Character							
1	Reserved (0)				CapId		CE	BC0
2	QIE ADC 0							
3	QIE ADC 1							
4	QIE ADC 2							
5	QIE ADC 3							
6	QIE ADC 4							
7	QIE ADC 5							
8	QIE ADC 6							
9	QIE ADC 7							
10	LE TDC 3		LE TDC 2		LE TDC 1		LE TDC 0	
11	LE TDC 7		LE TDC 6		LE TDC 5		LE TDC 4	

- BC0 : bunch counter alignment marker, sent once per orbit at an agreed phase
- CE : indicates a mismatch in capid between the four channels
- QIE ADC : Encoded charge-integrated amplitude
- LE TDC : leading-edge TDC measurement from inside the QIE, remapped into a limited set of bins by the Igloo2 FPGA. The specific remapping may be indicated by the reserved bits

Figure 21: Packet specification for the QIE11 PPOD input links to the uHTR (8-bit channel) [15]. Each packet represents a single bunch crossing, for a single detector cell.

8 Performance

8.1 Trigger Performance in Simulation

Plots for the timing trigger efficiency in simulation vs. LLP displacement and jet energy are shown in Figure 22. LLP jet efficiencies increase with LLP displacement, until the end of the HCAL volume (up to 6 m). Jet efficiencies significantly increase after 40 GeV, motivating this as the choice of lowest threshold Level 1 trigger seed. There is little dependence of the efficiencies on the HTT 120 selection.

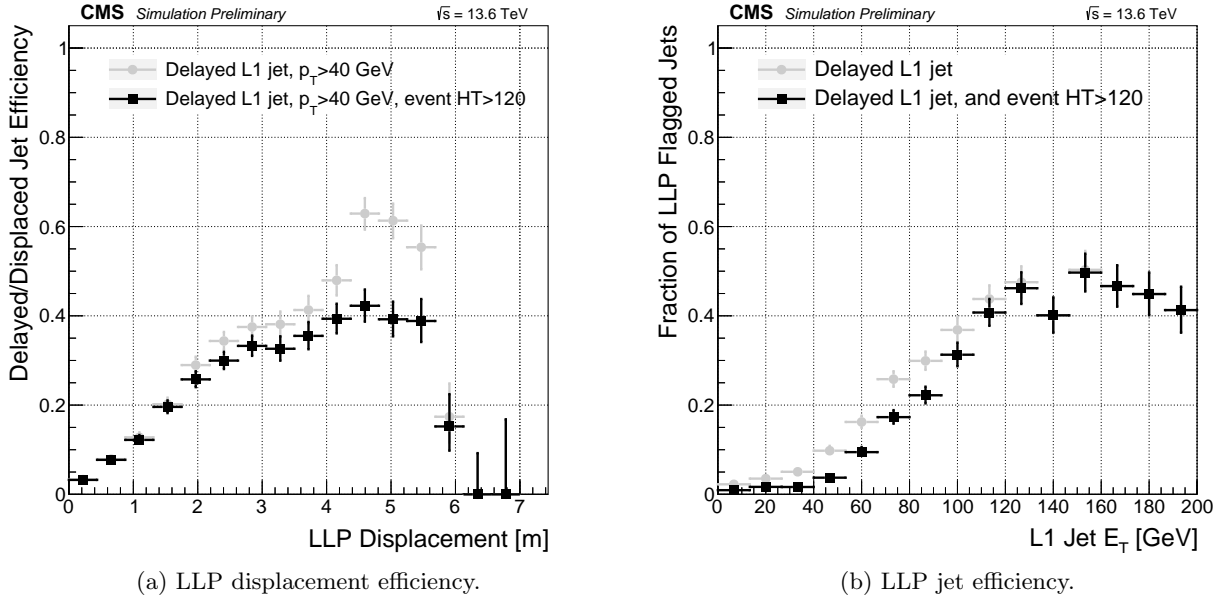
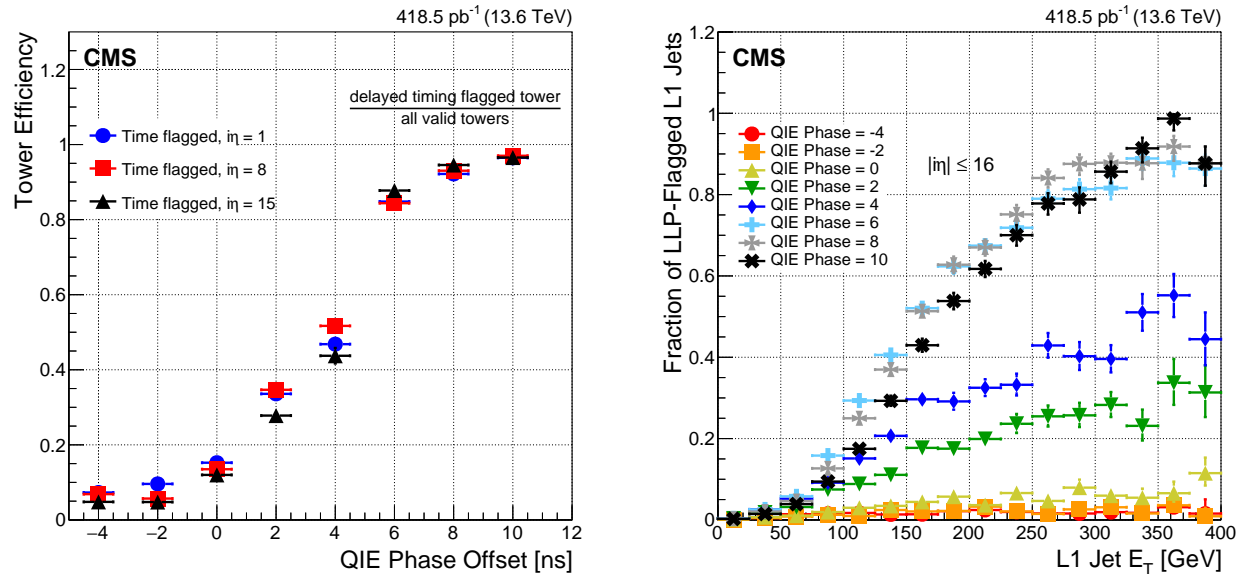


Figure 22: Efficiency plots with LLP flagged L1 jets. Four Monte Carlo samples are combined in these plots ($m_H = 125, c\tau = 3\text{m}$, $m_H = 250, c\tau = 10\text{m}$, $m_H = 350, c\tau = 10\text{m}$, $m_H = 1000, c\tau = 10\text{m}$) to show trigger performance across a representative range of LLP displacements and jet efficiencies.

8.2 Trigger Performance in Data

The QIE phase scan is one of the best ways of demonstrating the delayed jet performance in collisions data and validating the trigger pathway. Scanning the QIE phases adjusts the HCAL clock and pushes prompt jets into the delayed region, which is vital for LLP trigger validation. As the scan is taken relative to nominal phases, it artificially produces delayed jets, which occur across the whole detector at a known time. This phase scan is an ideal (and one of the only ways) to measure the delayed jet turn-on with the real timing resolutions of jets in collisions. The efficiency of the per tower delayed flag and the per jet delayed flag in the QIE scan show the turn-on curves with respect to phase delay. These are in Figure 23 a and b.



(a) Timing flagged tower efficiency. HCAL delayed timing tower efficiency during the 2023 HCAL phase scan, with efficiencies split by $i\eta = 1, 8, 15$ ($\eta \approx 0.04, 0.65, 1.26$). The delayed timing towers require at least one delayed cell and no prompt cells in the four depths contained in the tower, and the efficiency is calculated relative to towers with any valid timing code. Here, towers have a valid timing code if they contain one or multiple cells with energies over 4 GeV and with TDC code prompt, slightly delayed, or very delayed. A prompt cell has TDC value ≤ 6 ns, and a delayed cell has TDC value > 6 ns. The timing turn-on curve shows a sharp turn-on between QIE11 phase delays of 0-6 ns, as expected with the prompt timing range set at 6 ns by the HCAL TDC look-up-table (LUT). At low values of the phase offset, there is a small fraction of towers flagged as delayed, due to the pulse arrival time spread. With prompt collisions pulses, the tails extend into the delayed region, causing a few towers to be flagged as delayed.

(b) LLP jet efficiency. The HCAL LLP-flagged Level-1 Trigger (L1T) delayed jet fraction vs. jet E_T during the 2023 HCAL phase scan demonstrates that the delayed jet fraction reaches 1 as the phase delay is increased. A L1 jet is LLP-flagged when there are at least two flagged HCAL LLP towers (depth or timing-based) contained in the jet. The fraction of LLP-flagged L1 jets is compared to all L1 jets from a dataset of events enriched with jets or missing transverse energy (Raw JetMETdataset). The delayed jet fraction decreases at the largest delays (10 ns and above), as at these large delay settings the total hadronic jet energy is reduced since a significant amount of the jet energy has been pushed into the subsequent bunch crossing. No direct cut is made with respect to Jet E_T : the implicit requirement for a jet to have at least two cells with $E_T > 4$ GeV sculpts this distribution.

Figure 23: Efficiency plots vs. QIE delay in 13.6 TeV collisions, both published in [5].

Figure 24 shows the rates of a selection of L1 triggers during the QIE phase scan. This illustrates the strong impact the timing setting has on the LLP trigger rates. As the phase offset is increased to 6 ns, the LLP rates are maximized. The rates are greatly suppressed at a phase offset of 20 ns (equivalent to -5 ns), and in particular, the double jet trigger shows suppression of over two orders of magnitude.




Figure 24: L1 trigger rates for selected triggers during the May 2023 phase scan. The total trigger rates of the five L1 LLP triggers during the 2023 HCAL phase scan demonstrate the strong dependence of the LLP trigger rate on HCAL timing. Relative QIE11 phase values during the scan are shown, including a reference point of the 0 ns in early 2022 data-taking, while the 0 ns point of the phase scan is the recommended 2023 HCAL time alignment. L1 LLP trigger rates are maximized at HCAL timing delay of +6 ns, as jets have been pushed into the delayed region. At higher HCAL timing delays, the LLP trigger rates are decreased, as a significant amount of the jet energy is pushed into the subsequent bunch crossing. At optimal time alignment, the trigger rates for the five L1 LLP triggers implemented have a prescaled rate below 1 kHz. The trigger rates are shown vs. lumisection, which is a sub-section of a run (≈ 23.3 seconds) during which time the instantaneous luminosity is unchanging. Published in [5].

Five Level 1 trigger seeds are used in 2022 and 2023. These are designed as delayed jet triggers (Section 6.2), with increasing event energy (HT) thresholds and jet energy thresholds. A double jet trigger with low energy requirements is also used. These L1s are used to seed HLT delayed and displaced jet triggers, each of which is listed in Tables 7 and 8.

L1 Trigger	Pre-scale factor	Rates
		Hz, after prescale
L1_HTT120_SingleLLPJet40	100	25.6
L1_HTT160_SingleLLPJet50	50	27.2
L1_HTT200_SingleLLPJet60	1	769
L1_HTT240_SingleLLPJet70	1	463
L1_DoubleLLPJet40	1	34.4

Table 7: L1 trigger paths in Run 3, rates referenced from Run 368822 after the 2023 HCAL time alignment. The updated time alignment decreased LLP rates by a factor of 3, as expected. The total rate is under 1 kHz.

9 Summary

The development of a novel timing trigger using hadron calorimeter online timing information is presented, along with an alignment method relying upon TDC. The calorimeter depth and timing trigger is designed to increase sensitivity to LLPs by extending the LLP acceptance to particles decaying both before and within the calorimeter volume. This LLP timing trigger has been implemented in hardware and software in the CMS Experiment for Run 3 of the LHC and targets both the delayed time of arrival of LLPs along with

HLT Trigger	Rate at $2 \cdot 10^{34}$ (Hz)
HLT_HT200_L1SingleLLPJet_DisplacedDijet40_Inclusive1PtrkShortSig5_v3	19.31 (disabled)
HLT_HT240_L1SingleLLPJet_DisplacedDijet40_Inclusive1PtrkShortSig5_v3	14.08
HLT_HT280_L1SingleLLPJet_DisplacedDijet40_Inclusive1PtrkShortSig5_v3	9.52
HLT_HT170_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack_v3	4.34
HLT_HT200_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack_v3	2.97
HLT_HT200_L1SingleLLPJet_DisplacedDijet60_DisplacedTrack_v3	1.25
HLT_HT270_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack_v3	1.44
HLT_HT320_L1SingleLLPJet_DisplacedDijet60_Inclusive_v3	3.31
HLT_HT420_L1SingleLLPJet_DisplacedDijet60_Inclusive_v3	1.23
HLT_HT200_L1SingleLLPJet_DelayedJet40_SingleDelay1nsTrackless_v3	4.63
HLT_HT200_L1SingleLLPJet_DelayedJet40_SingleDelay2nsInclusive_v3	0.14
HLT_HT200_L1SingleLLPJet_DelayedJet40_DoubleDelay0p5nsTrackless_v3	2.11
HLT_HT200_L1SingleLLPJet_DelayedJet40_DoubleDelay1nsInclusive_v3	[low stats]

Table 8: HLT trigger paths in Run 3, rates referenced from Run 368822 (PU 54, n. bunches 2400). The rate reduction factor for HCAL-based HLT paths compared to rates pre-HCAL re-alignment is about 3, as expected. Standard jet triggers are minimally affected by the time re-alignment, with rate reduction factors of 1.02 for low p_T jet triggers and up to 1.08 for high p_T jet triggers.

unique energy deposits from LLPs.

The online TDC from the hadron calorimeter allows for time alignment to be achieved within 0.5 ns. The new time alignment method using data from a phase scan is detailed. An additional alignment method based on re-weighted and calibrated pulse shape distributions is presented and this method is calibrated with the higher precision TDC. The re-weighted pulse shape timing method provides additional avenues for calorimeter timing, such as corrections when a phase scan cannot be done, or for use in software triggering and pulse reconstruction offline. Both methods are validated with LED data, which provides a consistent (non-position dependent) pulse shape.

This dedicated LLP trigger creates a novel capability for Run 3 of the CMS experiment that is central to understanding physics beyond the Standard Model and aims to contribute to a comprehensive search for new physics with few underlying assumptions.

References

- [1] Juliette Alimena et al., *Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider*, Journal of Physics G: Nuclear and Particle Physics **47** (2020), no. 9, 090501.
- [2] CMS Collaboration, *The CMS experiment at the CERN LHC*, Journal of Instrumentation **3** (2008), no. 08, S08004–S08004.
- [3] ———, *CMS Technical Design Report for the Phase 1 Upgrade of the Hadron Calorimeter*, Tech. Report CERN-LHCC-2012-015. CMS-TDR-010, CERN, Geneva, September 2012, Final version.
- [4] ———, *The Phase-2 Upgrade of the CMS Level-1 Trigger*, Tech. Report CERN-LHCC-2020-004. CMS-TDR-021, CERN, Geneva, April 2020, Final version.
- [5] ———, *Performance of long lived particle triggers in run 3 (cms dp-2023/043)*, Tech. report, CERN, cms-trigger-coordinator@cern.ch, July 2023, Available at <https://twiki.cern.ch/twiki/bin/view/CMSPublic/Run3LLPHLT>.
- [6] Seth I. Cooper, *Phase I Upgrade of the CMS Hadron Calorimeter*, Nuclear and Particle Physics Proceedings **273–275** (2016), 1002–1007, 37th International Conference on High Energy Physics (ICHEP).
- [7] Jay Richard Dittmann, *HCAL Tower Depth Segmentation Figures*, CMS DocDB, June 2017, Available at <https://cms-docdb.cern.ch/cgi-bin/DocDB/ShowDocument?docid=13342>.
- [8] G.Iles, A.W.Rose, and A.Tapper, *Layer 2 Input and Output Specification*, 26 April 2014, Available at https://twiki.cern.ch/twiki/pub/CMS/CMSUKL1TriggerUpgrade/L2_Input__Output_Specification_v1.docx.
- [9] Tullio Grassi, *HCAL Trigger and Readout (HTR) Reference Guide*, CMS DocDB, 2011, Available at <https://cms-docdb.cern.ch/cgi-bin/DocDB/ShowDocument?docid=4769>.
- [10] ———, *HB TDC: Proposed upgrade of the RM Igloo2 FPGA firmware (F/W)*, HCAL Operations/HB I&C, November 2019, Available at https://indico.cern.ch/event/864899/contributions/3644121/attachments/1946309/3229581/HB_RM_igloo2_18Nov2019.pdf.
- [11] Tullio Grassi and Theresa Shaw, *CMS HE FE QIE card - Upgrade Production Version*, CMS DocDB, 2018, Available at https://cms-docdb.cern.ch/cgi-bin/DocDB/RetrieveFile?docid=12782&filename=CMS_HE_FE_Production_Schematic_18MAR2016.pdf&version=2, <https://cms-docdb.cern.ch/cgi-bin/DocDB/ShowDocument?docid=12782>.
- [12] ———, *CMS HB 16-Channel Production QIE Card*, CMS DocDB, 2020, Available at <https://cms-docdb.cern.ch/cgi-bin/DocDB/ShowDocument?docid=13544>, https://cms-docdb.cern.ch/cgi-bin/DocDB/RetrieveFile?docid=13544&filename=CMS_HB_Production_FE_Module_schematic.pdf&version=1.
- [13] H.Bergauer, J.Ero, M.Jeitler, J.Wittmann, C.-E.Wulz, C.Foudas, K.Bunkowski, M.Konecki, L.Uvarov, C.Battilana, D.Rabady, H.Sakulin, J.Brooke, G.Iles, A.Tapper, D.Acosta, I.Furic, A.Madorsky, M.Matveev, P.Padley, P.Klabbers, and W.H.Smith, *Scales for inputs to μGT (φ , η , p_t/E_t , and others)*, 23 February 2021, Available at https://globaltrigger.web.cern.ch/globaltrigger/assets/upgrade/ugt/scales_inputs_2_ugt_2021Feb23.pdf.
- [14] Gillian Kopp, Chris Tully, Owen Long, and Georgia Karapostoli, *Specifications for HCAL uHTR Firmware*, CMS Public DocDB, August 2021, Available at https://cms-docdb.cern.ch/cgi-bin/PublicDocDB/RetrieveFile?docid=12306&filename=LLPbits_uhtr_spec.pdf&version=22.
- [15] Jeremiah Mans, *HCAL μTCA Trigger and Readout Module (uHTR)*, CMS Public DocDB, August 2021, Available at <https://cms-docdb.cern.ch/cgi-bin/PublicDocDB/RetrieveFile?docid=12306>.
- [16] Grace E. Cummings on behalf of the CMS Collaboration, *Overview of CMS HCAL Barrel Phase 1 Upgrade*, LHC Students Poster Session, Winter LHCC, 2019, Available at https://indico.cern.ch/event/797774/contributions/3321893/attachments/1802113/2939720/GECummings_HBUpgrade_LHCC2019.pdf.

- [17] A Tapper and Darin Acosta, *CMS Technical Design Report for the Level-1 Trigger Upgrade*, Tech. report, CERN, 2013, Additional contacts: Jeffrey Spalding, Fermilab, Jeffrey.Spalding@cern.ch Didier Contardo, Universite Claude Bernard-Lyon I, didier.claude.contardo@cern.ch.