

ENABLING NOVEL LLP SEARCHES THROUGH 5D
CALORIMETRY WITH THE COMPACT MUON SOLENOID
EXPERIMENT AT THE LARGE HADRON COLLIDER

GILLIAN BARON KOPP

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
PHYSICS
ADVISER: CHRISTOPHER TULLY

MAY 2024



© Copyright by Gillian Baron Kopp, 2024.

All Rights Reserved

Abstract

Long-lived particles are a compelling direction to search for physics beyond the Standard Model, and implementing dedicated triggers provides an excellent avenue to expand experimental coverage into this challenging parameter space. With the recent upgrades of the Hadron Calorimeter (HCAL) in the Compact Muon Solenoid (CMS) experiment at the Large Hadron Collider (LHC), the HCAL provides new timing capabilities for jets and hadronic tau decays with nearly 4π coverage, and sensitivity to highly displaced decays that occur within the calorimeter volume. The augmentation of the calorimeter information at the hardware trigger level enables quick identification of long-lived particle (LLP) decays using lower thresholds on energy-based event quantities. Both the depth segmentation and online timing capabilities are utilized in novel HCAL-based hardware-level triggers to identify displaced and delayed LLPs, either decaying inside the calorimeter volume (with sensitivity for decays up to almost 6 m from the collision point) or arriving at a delayed time. These triggers are deployed for Run-3 of the LHC, beginning data-taking in 2022. This two-pronged calorimeter trigger approach leverages the new capabilities of the CMS HCAL to expand the phase space accessible in ongoing LLP searches. In addition, online timing is utilized for detector alignment based on positioning the pulse rising edge, achieving an alignment accuracy of 0.5 ns, considerably higher than previous energy-weighting based approaches. The data collected with the new triggers provides a first look at the capabilities to capture softer events and expand the phase space accessible in LLP searches. A search is conducted for long-lived particles using the CMS experiment's 2022 and 2023 Run-3 datasets, comprising 67 fb^{-1} from proton-proton collisions at a center of mass energy of 13.6 TeV. LLPs have unique signatures distinct from Standard Model processes, and this analysis utilizes calorimeter and track-based variables such as shower radial distribution, neutral hadron energy, and trackless jet characteristics to identify LLP jets. The use of HCAL timing, energy, and positional (x, y, z) measurements in the trigger and analysis form the basis of 5D calorimetry. A background estimation is performed, and expected signal limits are set for an extended Higgs sector model.

Acknowledgements

Many wonderful individuals have made my graduate school experience immensely enjoyable and engaging over the last six years, and have encouraged me to challenge myself and do more than I imagined I could. Thank you to all who have brought opportunities, encouragement, and compassion, both academically and personally.

First, I would like to thank my advisor, Chris Tully, who has given me the freedom to explore and the trust to take on these projects. This has allowed me to take risks, push myself, work through setbacks, and spend the past years delving into the depths of this physics. I appreciate that I am constantly pushed to be a better researcher and physicist, and I'm always excited to dive back into work after our conversations. The intellectual freedom I have had, combined with Chris' care for me as a person first of all, has made my experience on CMS incredibly valuable.

The Princeton CMS group is a vibrant group of fellow particle physicists, and I have enjoyed the time I overlapped with everyone. Anywhere I go, between Princeton, CERN, and Fermilab, the group is always welcoming and makes each place feel like home. In particular, thanks to Kiley Kennedy for her excellent mentorship and guidance. To everyone, thanks for the physics discussions, debugging help, and entertaining chats: Marco Lucchini, Kelvin Mei, Wonyong Chung, Stephanie Kwan, Sam Higginbotham, Gage DeZoort, Bennett Greenberg, Nick Haubrich, Hichem Bouchamaoui, and Soohyun Yoon. My time at CERN was much enriched by the Princeton group there, thanks to Svitlana, Pallabi, Andrew, Adrian, Elliott, and Ashling for making B15 and B588 a welcoming place.

I am very thankful for the support from Princeton professors Isobel Ojalvo, Dan Marlow, and Jim Olsen. I also want to thank Chris, Isobel, Mariangela Lisanti, and Andrew Leifer for serving on my committee and Dan Marlow for being my dissertation second reader.

I greatly appreciated working with the HCAL operations group at CERN. In particular, my collaborations with Maria Toms, Pasha Parygin, Salavat Abdoulline, Seth Cooper, fellow DAQ and DOC on-call shifters, John Hakala, and Mariarosaria D'Alfonso taught me so much about running a detector as I learned the ins and outs of online HCAL operations and ran the phase scans. Collaboration with the UMN group, Jeremy Mans, Michael Krohn, Josh Hiltbrand, and Bryan Crossman, has been invaluable and enlightening, especially throughout the years of teststand and firmware work. Thanks to both Kelvin and Wonyong for their teststand expertise and dedication. From Caltech, Cristián Peña and Si Xie were some of the first people who introduced me to particle physics when

I joined Maria Spiropulu's group. It has been a joy to work with them on the LLP analysis at Fermilab after starting my particle physics journey there at an HGCAL testbeam in 2016.

Friends from each part of my life have all been a source of happiness, and I appreciate each of you. From Geneva, I have such fond memories of my time at Place de la Diversité with Stephanie Kwan and Bennett Greenberg. Thanks for being the best flatmates; you two truly made the experience of living abroad wonderful. Through our adventures moving abroad, CERN life and lunches, late-night chats, and incredible hiking trips, there was always laughter and support. Stephanie and I have been office-mates, apartment-mates (in three countries), and pandemic-mates, and truly, if we have gotten through graduate school and a years-long pandemic together, what can't we do? All my CERN and Geneva friends helped make a new country home through our time hiking, bouldering, game nights, and everything: thanks to Dan, Amy, David, Daniel, Christian, and Ricky.

I have many wonderful friends in Princeton who made returning there very enjoyable. Sara Sussman, I have fond memories of our time working together on Women in Physics and prelims, and I have appreciated reconnecting. Liz Helfenberger, you've been a wonderful friend since our first semester here, and it's fantastic to finish up graduate school as office-mates. I'm so proud to see you move on to your next position. Adri Dropulic, thanks for being there through all the highs and lows; your dedication, confidence, and unwavering truth to yourself are inspiring. Adri and Nora Linzer have been incredible housemates, with our lovely evening hangouts and catch-up chats. Sophie Dvali, I am so thankful for our friendship, our travels together, and how it feels like just yesterday, no matter how much time has passed since we've been together. Gio Arsenadze, thanks for sharing many excellent climbing and swimming evenings, for your incredible cooking, and for sharing Georgian cuisine. The hours spent talking into the evening after dinners with friends are always my favorite. Hichem Bouchamaoui, thanks for entertaining us with memes, and for lending me your office space to write this thesis.

My deepest thanks to William Ballinger for the constant support and unfailing faith in me, and for always believing in us. You bring such a wonderful sense of calm, and I cannot overstate how grounding that has been throughout the past years. Our travels and time spent together are invariably some of my best experiences.

And to everyone else, especially the Thursday movie night crew, Rockville climbing group, Totem bouldering crew, Fermilab lunch group, and friends from Geneva French class: thanks for bringing me away from work and reminding me to take a break to be in the broader world. All the love to

Core 6 (Gabby, Anusha, Carly, Anne, and Steph), my day one group of amazing and brilliant Caltech women. Each of you inspires me with your dedication and passion for your chosen path.

Finally, thanks to my family: my mom for her support and belief in me; my dad for inspiring a love of science from a young age; and my sister, Fionna, the strongest person I know.

To my grandparents.

Contents

Abstract	3
Acknowledgements	4
1 Motivation and Introduction	13
1.1 Standard Model of Particle Physics	14
1.1.1 Higgs Mechanism	17
1.1.2 Higgs Boson and the Yukawa Interaction in the Standard Model	18
1.1.3 Higgs Mass Quadratic Divergence	19
1.1.4 Hierarchy Problem	21
1.2 Physics Beyond the Standard Model	22
1.2.1 Limitations of the Standard Model	23
2 Long-Lived Particles	25
2.1 LLP Motivations	26
2.2 Experimental Signatures	27
2.2.1 Distinct Features of the Final State	28
2.2.2 LLP Landscape at the LHC	29
2.3 Exotic Higgs Decays to LLPs	31
2.3.1 $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$	31
2.3.2 Decays of an Exotic Higgs Boson	32
2.4 Improved LLP Search Opportunities	34
3 The Compact Muon Solenoid Detector	35
3.1 Large Hadron Collider	35
3.2 Compact Muon Solenoid Experiment	38

3.2.1	Detector Coordinate System	38
3.2.2	CMS Sub-Detectors	39
3.2.3	Tracker	40
3.2.4	Electromagnetic Calorimeter (ECAL)	41
3.2.5	Hadron Calorimeter (HCAL)	42
3.2.6	Superconducting Solenoid	47
3.2.7	Muon System	47
3.3	CMS Trigger System	48
3.3.1	Level-1 Trigger System	48
3.3.2	High Level Trigger	50
3.3.3	Run-3 Trigger Improvements	50
3.4	Offline Reconstruction	51
3.4.1	Jets	51
3.4.2	Particle Flow	53
3.5	Data Storage	55
4	Time Alignment of the Hadron Calorimeter	57
4.1	Time to Digital Converter (TDC)	57
4.1.1	Digitizing SiPM Signals with the QIE11	58
4.1.2	HB 6:2 TDC LUT	59
4.2	HCAL Time Alignment	59
4.2.1	QIE Phase Scans	60
4.2.2	Phase Scan Method	61
4.2.3	Phase Scan Analysis	62
4.2.4	2022 HB Time Alignment Issues	66
4.3	Time Alignment Validation	68
4.3.1	TDC Based Method of Time Alignment	69
4.3.2	Charge Weighted Time	70
4.3.3	Energy Response	79
4.4	Far-Reaching Impacts of HCAL Timing	84
4.5	Summary	85
5	Level-1 Delayed and Displaced Jet Triggers	87
5.1	Trigger Development	88

5.1.1	HCAL Trigger Pathway	89
5.1.2	TDC Considerations – Threshold and Amplitude Dependence	90
5.1.3	2018 HE TDC Data	91
5.2	Trigger Implementation	93
5.2.1	LLP flag: Timing trigger	93
5.2.2	LLP flag: Depth trigger	93
5.2.3	OR of Timing and Depth Towers	94
5.2.4	LLP Jet Flag	94
5.3	Firmware Specifications	95
5.3.1	uHTR Pathway	95
5.3.2	Level-1: Layer-1 and Layer-2 Pathway	97
5.3.3	uGT Pathway	97
5.4	Trigger Commissioning	98
5.5	Performance in Run-3 Collisions	99
5.6	Performance in Simulation	102
5.6.1	Efficiency and Rate Tables	104
6	HCAL-based High Level LLP Triggers	107
6.1	HLT Paths	108
6.1.1	Added Acceptance	109
6.1.2	HLT Rates	109
6.2	HLT Efficiencies	110
6.2.1	Improvements for 2024: CalRatio Trigger	115
6.3	HLT Performance	116
7	Datasets and Monte Carlo Samples	117
7.1	Run-3 Datasets	117
7.1.1	Skim for LLP Analysis	118
7.2	LLP Monte Carlo Samples	119
7.2.1	Masses and Lifetimes	119
7.2.2	Production of Samples	120
7.2.3	LLP Lifetime Re-weighting	120
8	Analysis Strategy	122

8.1	Object Selection	123
8.1.1	Particle Flow Jets	123
8.1.2	LLP-Jet Matching Efficiencies	123
8.2	Event Selection	126
8.2.1	Cut-Flow Table	126
8.2.2	W+Jets Selection	127
8.2.3	Minimal Signal Contamination	128
8.3	Sensitive Jet Object Variables	129
8.3.1	HCAL Rehit Based Variables	129
8.3.2	Track Based Variables	133
8.3.3	Discussion of Sensitive Jet Object Variables	134
8.3.4	Dependence on Jet Energy	139
8.4	Discriminator: Boosted Decision Tree	139
8.4.1	Training Variables	140
8.4.2	BDT Setup	142
8.4.3	BDT Performance	142
8.4.4	Performance of Triggered Events	145
8.4.5	BDT Cutflow Comparison	146
9	Backgrounds and Sensitivities	148
9.1	Background Estimations	148
9.1.1	Sources of Background	149
9.1.2	Fake Rate Method	150
9.1.3	Closure Tests	156
9.2	Systematic Uncertainties	157
9.2.1	Jet Energy Corrections	157
9.2.2	Custom Variable Contributions	157
9.3	Expected Sensitivities	158
9.4	Summary	159
10	Conclusions and Outlook	161
A	Run-3 Notable Runs and Significant Changes	163
B	HCAL QIE Phases and TDC-LUTs	166

B.1	TDC LUTs	166
B.2	Timing Adjustments	166
B.3	Compensating Response Corrections	167
C	Monte Carlo Simulation Generation	171
C.1	Monte Carlo for Trigger Development	171
C.2	Monte Carlo for LLP Analysis	172
C.2.1	Potential Future Extensions	173
D	Trigger Efficiencies and Integrated Luminosity Gains	174
E	Analysis Cutflow Tables	179
E.1	Event Selection	179
E.2	BDT Comparison	181
E.2.1	Additional ROC Curves	181
E.3	0, 1, and 2+ Tagged Jet Regions	182
E.3.1	BDT Discriminator Threshold 0.9999	182
E.3.2	Expected Sensitivities	182
E.3.3	BDT Discriminator Threshold 0.9995	185
F	MIP Timing Detector and Future LLP Searches	190
F.1	Barrel Timing Layer Testbeam	191
F.1.1	Response Uniformity	191
F.1.2	Efficiency	192
F.1.3	Cross-talk	193
F.2	Timing in Future Detectors	193

Chapter 1

Motivation and Introduction

Particle physics strives to answer questions about the fundamental constituents of the universe and how they interact to form the world as we know it. The idea that the universe is made of particles was proposed by Greek philosophers and has been iterated many times over as our understanding of *what* those most fundamental particles are improves. The advances made in the field over the last century have revolutionized our understanding, particularly with the quantum mechanical wave behavior of particles (1920s) and the adoption of the Standard Model of particle physics (1960s and 70s) – both considered by Thomas Kuhn as paradigm shifts in “The Structure of Scientific Revolutions” [64]. With the discovery of the Higgs Boson on July 4th, 2012, the full particle content of the Standard Model has been discovered. Since then, the model has continued to hold up remarkably well to experimental observations. The particles of the Standard Model form visible matter and carry the forces for three of the four fundamental interactions – mediating the strong, electromagnetic, and weak forces. In contrast, a particle mediator for gravity is unknown. The Standard Model describes all known fundamental particles and predicts their properties, is internally consistent, and is validated by many experimental observations – however, it is not a complete theory of fundamental physics.

As the Standard Model is incomplete and leaves numerous outstanding questions, we wonder if a more complete theory exists. Such a theory could be probed by observing deviations from the Standard Model predictions. Thus, the field is now at a fascinating juncture where the Standard Model has held up well through multiple experiments, yet there is much physics that it does not describe, and we are diving into this unknown. This thesis will provide an overview of how complex

experiments at the highest energy proton-proton collider in the world are used to investigate and make statements about models describing fundamental physics.

1.1 Standard Model of Particle Physics

The fundamental particles of the Standard Model (SM) are split into two groups: the twelve fermions (spin $\frac{1}{2}$) that comprise matter and the five bosons. The fermions are divided into quarks and leptons, each with three generations, with each generation more massive than the previous. The quarks interact via the strong force (as they carry color charge, described by QCD) and are found in bound states called hadrons [76]. Protons (uud) and neutrons (ddu) are examples of hadrons, and are baryons as they contain three quarks. Mesons, such as the pion ($\pi^+ = u\bar{d}$), contain two quarks. Quarks have color confinement, preventing them from existing outside of such bound states. The leptons include the electron, muon, tau, and their corresponding neutrinos. Leptons interact via the weak and electromagnetic interactions (only charged particles interact electromagnetically). There are five bosons: the four gauge bosons (spin 1) and the single scalar boson (spin 0, Higgs Boson).




Figure 1.1: Diagram of the particles in the Standard Model of particle physics. From [78].

Three of the known forces of nature are carried by the gauge bosons. These are the strong force,

carried by the gluon; the electromagnetic force, carried by the photon; and the weak force, carried by the W and Z bosons [76]. The “graviton” is expected to be the force-carrying particle for the gravitational force but has yet to be discovered and is not included in the Standard Model. Reconciling quantum mechanics and general relativity in one theory has been difficult – a shortfall of the Standard Model. Compared to the strong force, EM is weaker by a factor of 10^{-3} , weak force by a factor of 10^{-8} , and gravity by a factor of 10^{-37} (comparisons are quoted at the scale of a proton, 10^{-15} m, though the exact values of the forces will depend on the distance scale considered) [76].

There are 61 particles described by the Standard Model: 12 leptons (leptons, anti-leptons), 36 quarks (quarks, anti-quarks, each coming in three colors), twelve gauge bosons (eight gluons, photon, Z, W^+ , and W^-), plus the Higgs boson.

Mathematically, the Standard Model of particle physics is described by the gauge symmetry group:

$$SU(3)_C \times SU(2)_L \times U(1)_Y \quad (1.1)$$

Fermion fields (leptons and quarks) transform under $SU(2)_L$, with L indicating left chirality. Right chiral fermions transform as singlets, and left chiral fermions transform as doublets under $SU(2)_L$ [41]. The three generations of fermions are indicated by $i = 1, 2, 3$, where the lepton generations are the electron, muon, and tau, and the quark generations are the up / down, charm / strange, and top / bottom. Equation 8.3 in [41] describes the transformation of the fermion fields under $SU(2)$:

$$\ell_{1L} = \begin{pmatrix} \nu_e \\ e^- \end{pmatrix}_L \quad \ell_{2L} = \begin{pmatrix} \nu_\mu \\ \mu^- \end{pmatrix}_L \quad \ell_{3L} = \begin{pmatrix} \nu_\tau \\ \tau^- \end{pmatrix}_L \quad (1.2)$$

$$e_{1R} = e_R^- \quad e_{2R} = \mu_R^- \quad e_{3R} = \tau_R^- \quad (1.3)$$

Neutrinos (ν) are only left-handed. Therefore, there is no equivalent ℓ_R representation for them, while there is for the electrons, muons, and taus. The quarks also have three generations and are organized into doublets (left-handed chirality) and singlets of up or down type (right-handed

chirality) [41]:

$$q_{1L} = \begin{pmatrix} u \\ d \end{pmatrix}_L \quad q_{2L} = \begin{pmatrix} c \\ s \end{pmatrix}_L \quad q_{3L} = \begin{pmatrix} t \\ b \end{pmatrix}_L \quad (1.4)$$

$$u_{1R} = u_R \quad u_{2R} = c_R \quad u_{3R} = t_R \quad (1.5)$$

$$d_{1R} = d_R \quad d_{2R} = s_R \quad d_{3R} = b_R \quad (1.6)$$

Quarks are termed “up-type” (up, charm, top) and “down-type” (down, strange, bottom). There are anti-fermions for all the leptons and quarks mentioned above. In addition, quark fields transform under the $SU(3)_C$ gauge couplings (C is color, for quantum chromodynamics). Quarks and gluons have color charge, and their transformation under $SU(3)_C$ is quantum chromodynamics, which is invariant under color rotation. The quarks are color triplets that are transformed by a phase transformation, while the leptons are color singlets and are unchanged (Equation 8.4 in [41]):

$$q_{L,R}(x) \rightarrow e^{-ig_s \alpha_s^a(x) \lambda^a / 2} q_{L,R}(x) \quad \ell_{L,R}(x) \rightarrow \ell_{L,R} \quad (1.7)$$

g_s and α_s^a are gauge couplings of $SU(3)_C$ and λ^a are the Gell-Mann matrices, the generators of $SU(3)$. Gluons, g^a , also transform under the adjoint representation of $SU(3)_C$.

The generators of $SU(2)_L$ are taken to be three $W_{1,2,3}^\mu$ gauge bosons, and the generator of $U(1)_Y$ as the gauge boson B^μ , where Y is hypercharge [41]. The gauge bosons comprise the $SU(2)_L \times U(1)_Y$ symmetry of the electroweak interaction, and with exact electroweak symmetry, all are massless. However, we observe massive $W^\pm$ and Z gauge bosons, which are accounted for by the spontaneous breaking of the electroweak symmetry by the Higgs mechanism. The $SU(2)_L$ Higgs field doublet breaks the symmetry $SU(2)_L \times U(1)_Y$ to $U(1)_{\text{EM}}$, while $SU(3)_C$ remains unbroken [77].

The Higgs field transforms under $SU(2)_L$. It is unique in having a “Mexican hat-shaped” potential, where the origin does not have the minimum potential but rather a non-zero vacuum expectation value. The fact that the origin does not have the lowest potential, but the lowest energy state is instead a radius of points around the origin, is considered spontaneous symmetry breaking (movement from Point A to B in Figure 1.2). This non-zero value breaks the $SU(2)_L \times U(1)_Y$ symmetry and allows other particles to acquire mass through the Higgs mechanism.




Figure 1.2: The Higgs potential shape, with a non-zero value at the origin. Figure from [55].

1.1.1 Higgs Mechanism

This spontaneous symmetry breaking from $SU(2)_L \times U(1)_Y$ to $U(1)_{\text{EM}}$ is due to a minimization of the Higgs potential, from the Lagrangian

$$\mathcal{L} = \frac{1}{2} \partial_\mu \phi \partial^\mu \phi - \left(\frac{1}{2} \mu^2 \phi^2 + \frac{1}{4} \lambda \phi^4 \right) \quad (1.8)$$

with $\mu^2 < 0$ and $\lambda > 0$. The scalar Higgs field $\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}$ (a doublet) has a non-zero vacuum expectation value, from the minimization of its potential term in the Lagrangian giving $\phi = \pm \sqrt{\frac{-\mu^2}{\lambda}} = v$, and thus:

$$\langle \phi \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix} \quad (1.9)$$

The vacuum expectation value (VEV), or $v = \sqrt{1/\sqrt{2}G_F} \approx 246$ GeV, contributes to the mass of the W and Z bosons through EW symmetry breaking; however, the photon remains massless. The boson masses are (Equation 8.6 in [41]):

$$M_W = \frac{1}{2} g_2 v \quad M_Z = \frac{1}{2} \sqrt{g_Y^2 + g_2^2} v \quad (1.10)$$

where g_Y and g_2 are the gauge coupling constants from $U(1)_Y$ and $SU(2)_L$, respectively. Now the fields for the physical gauge bosons ($W^\pm$, Z, and photon) are found by mixing the three $SU(2)$ gauge

bosons $W_{1,2,3}$ and the $U(1)$ gauge boson B (Equation 8.7 in [41]):

$$W^{\mu,\pm} = \frac{1}{\sqrt{2}(W_1^\mu \mp iW_2^\mu)} \quad Z^\mu = -\sin\theta_W B^\mu + \cos\theta_W W_3^\mu \quad A^\mu = \cos\theta_W B^\mu + \sin\theta_W W_3^\mu \quad (1.11)$$

$\sin\theta_W$ is the weak mixing angle, $W^\pm$ are the charged W bosons, and Z is the neutral boson. A is the massless photon – the only boson that has not acquired mass through the Higgs mechanism. Note that A only depends on W_3 and B . The charge of a boson particle depends on its interaction with the photon A (carrying the EM force), and because W_3 and $W_{1,2}$ do not commute, the W bosons are charged while the Z is not. The W and Z bosons have absorbed three degrees of freedom from the initial doublet, while the fourth gives the Higgs boson [77].

The remaining neutral scalar after symmetry breaking is the Higgs boson. The Higgs field contributes to the fermionic terms in the Lagrangian in the form of ϕ and ϕ^C (Higgs doublet field and the corresponding charge conjugated doublet) in the Yukawa interaction [41]. The charged lepton and quark masses are directly proportional to v and a mass matrix for each generation (Equation 8.11 in [41]):

$$\text{charged lepton mass} = m_{e;ij} = \frac{1}{\sqrt{2}} f_{ij}^{e*} v \quad (1.12)$$

$$\text{quark mass} = m_{d,u;ij} = \frac{1}{\sqrt{2}} f_{ij}^{d,u*} v \quad (1.13)$$

The fermion mass terms in the SM Lagrangian are not invariant under electroweak symmetry ($SU(2)_L \times U(1)_Y$), so they have a bare mass of zero, and their physical mass is acquired through the spontaneous symmetry breaking [49]. Thus, the SM fermion and gauge boson masses all depend directly on the Higgs vacuum expectation value.

1.1.2 Higgs Boson and the Yukawa Interaction in the Standard Model

In the SM, the Higgs boson is a $SU(2)_L$ doublet with $\langle\phi\rangle = 1/\sqrt{2} \begin{pmatrix} 0 \\ v \end{pmatrix}$ and a hypercharge value of $Y_\phi = 1$. The SM Lagrangian coupling of the Higgs to fermions involves both the field and its charge conjugate, ϕ, ϕ^C – but they have different Yukawa interaction couplings for down and up type right-handed chiral fermions. The field ϕ couples with the down type right-handed fermions and massive right-handed leptons (e, μ, τ), while the conjugate couples with the up type right-handed fermions. The SM Lagrangian terms containing fermions and the Higgs are (Equations 8.9 and 8.10

in [41]):

$$\mathcal{L}_1 = -f_{ij}^{e*} \bar{\ell}_{iL} \phi e_{jR} - f_{ij}^{d*} \bar{q}_{iL} \phi d_{jR} + h.c. \quad (1.14)$$

$$\mathcal{L}_2 = -f_{ij}^{u*} \bar{q}_{iL} \phi^C u_{jR} + h.c. \quad (1.15)$$

where:

- f_{ij} indicates the coupling strength between the fermions and Higgs field, and the superscript indicates which fermion this coupling strength is for
- $\bar{\ell}_L$ indicates a left-handed anti-lepton (electron, muon, tau, and three associated neutrinos)
- $\bar{q}_L$ indicates a left-handed anti-quark (up, down, charm, strange, top, bottom)
- ϕ and ϕ^C are the Higgs field and the charge conjugated Higgs field, with $\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}$ and $\langle \phi \rangle = 1/\sqrt{2} \begin{pmatrix} 0 \\ v \end{pmatrix}$, and v is the vacuum expectation value
- e_R and d_R are down type right-handed chiral fermions (electrons, muons, taus and down type quarks: down, strange, bottom)
- u_R are up type right chiral fermions (up type quarks: up, charm, top)
- indices i, j are generation indices for the three generations of fermions
- $h.c.$ indicates the Hermitian conjugate of the previous terms

Note that the Yukawa interaction couples two different generations together – anti-quarks are coupled with quarks from other generations ($f_{ij}^{d*} \bar{q}_{iL} \phi d_{jR}$, $f_{ij}^{u*} \bar{q}_{iL} \phi^C u_{jR}$ as indicated by the i, j indexes), while anti-leptons are coupled with leptons from different generations ($-f_{ij}^{e*} \bar{\ell}_{iL} \phi e_{jR}$). Additionally, this Lagrangian restricts left-handed particles to only interact with right-handed particles. Since neutrinos are only left-handed (ν_R is not known to exist), there is no leptonic term in $\mathcal{L}_2$ above, while $\mathcal{L}_1$ has a lepton and quark term.

1.1.3 Higgs Mass Quadratic Divergence

Due to the Higgs coupling to fermions, a fermionic loop contributes a correction to the Higgs boson two-point function (propagation amplitude), which gives the observed Higgs mass. Following Chapter 1 in “Theory and Phenomenology of Sparticles” [41]: The Lagrangian for the Higgs decaying

to a fermion and anti-fermion is given by the Yukawa coupling to the Higgs field (Equation 1.3 in [41]):

$$\mathcal{L} = -\frac{\lambda_f}{\sqrt{2}} h \bar{f} f - \frac{\lambda_f v}{\sqrt{2}} \bar{f} f \quad (1.16)$$

From the second term of the Lagrangian $\lambda_f v/\sqrt{2} \cdot \bar{f} f = m_f \bar{f} f$, it can be seen that the fermion mass is $m_f = \lambda_f v/\sqrt{2}$, and from the first term, that the vertex factor of $h \rightarrow f \bar{f}$ is $-i\lambda_f/\sqrt{2}$. Thus, an allowed Feynman diagram is a fermionic loop, which is the one-loop correction to the Higgs two-point function.




Figure 1.3: The one loop correction to the Higgs propagator with a fermion, anti-fermion loop. Diagram from [40].

The above Feynman diagram has two vertex factors of $-i\lambda_f/\sqrt{2}$, and two fermionic propagators each contribute factors of $i/(\not{k} - m_f)$. In calculating the propagation amplitude of the Higgs with the single fermionic loop correction, there is a dependence on the coupling constant λ_f^2 and the fermion mass m_f , though there is no dependence on the Higgs mass. The result is quadratically divergent in the momentum space integral (Equation 1.4 in [41]):

$$\Pi_{h \rightarrow f \bar{f} \rightarrow h}(0) = (-1) \int \frac{d^4 k}{(2\pi)^4} \text{Tr} \left(-i \frac{\lambda_f}{\sqrt{2}} \right) \frac{i}{\not{k} - m_f} \left(-i \frac{\lambda_f}{\sqrt{2}} \right) \frac{i}{\not{k} - m_f} \quad (1.17)$$

$$= -2\lambda_f^2 \int \frac{d^4 k}{(2\pi)^4} \left[\frac{1}{k^2 - m_f^2} + \frac{2m_f^2}{(k^2 - m_f^2)^2} \right] \quad (1.18)$$

An integral of the form $\int d^d x \cdot (2\pi)^{-d} (x^2 + a)^{-n}$ is divergent if $n \leq d/2$. In this case, $d = 4$ and $n = 2 = d/2$, and the above integral is quadratically divergent. Typically, dimensional regularization and renormalization are used to solve these integrals and cancel divergences, and to do so a cutoff parameter is introduced to restrict the integral only to be logarithmically divergent.

In the case of the above integral, the cutoff parameter Λ is chosen to be the reduced Planck mass scale of $m_{\text{Pl, reduced}} = m_{\text{Pl}}/\sqrt{8\pi} = 2.4 \cdot 10^{18}$ GeV, as this is the largest scale known in physics [41]. The one-loop correction to the two-point function for the Higgs will give corrections to the bare Higgs mass, $m_{H,0}^2$, resulting in the physical Higgs mass m_H^2 . Physical mass accounts for virtual loop

corrections due to couplings allowed by the Lagrangian, while bare mass is a fundamental parameter without these corrections. Setting the cutoff bound to be m_{Pl} causes this single loop correction to be 10^{30} times larger than the bare Higgs mass, which is restricted to be less than 1 TeV due to the known $W^+W^- \rightarrow W^+W^-$ process which occurs through a Higgs mediator [41]. This is an extremely large correction for a quadratic divergence – causing concern about the Standard Model.

Further concerns are caused by the fact that this one-loop correction to the Higgs mass is dependent on the square of the cutoff energy [42]:

$$m_H^2 = m_{H,0}^2 + \alpha \Lambda^2 \quad (1.19)$$

This means that the Higgs mass is directly sensitive to physics at much higher energy scales – and either requires that these cutoff energy scales are relatively low (implying expected discovery of physics beyond the SM at energies lower than a TeV) or SM parameters are extremely fine-tuned [42]. To ensure the Higgs mass remains at 125 GeV (as needed from observation), “fine-tuning” on the order of $1/10^{26}$ would be required [41]. If this fine-tuning is not done to such a precise level, the Higgs mass would be far too large to match observations. These are the fine-tuning and Higgs mass divergence problems, two of the issues discussed in Section 1.2.1.

1.1.4 Hierarchy Problem

The hierarchy problem is the fact that the electroweak interaction scale is incredibly small compared with the Planck mass scale, at 250 GeV and $1.2 \cdot 10^{19}$ GeV, respectively. The Planck mass is $m_{\text{Pl}} = \sqrt{\hbar c/G} = 1.2 \cdot 10^{19}$ GeV (10^{15} higher than the energy scales accessible with the LHC!), and is the scale at which the gravitational force strength is comparable with the other forces, so quantum mechanics is expected to be impacted by gravity (quantum gravity). Thus, the scales of the fundamental forces are vastly different, as gravity is weak (contributing to the Planck mass being so large) and the electromagnetic force is strong. The Standard Model precisely and accurately explains many observations of fundamental particles; however, it seems strange that the underlying parameters (the scales of the fundamental forces) differ by such large factors.

This is also seen in the above description of the Higgs mass dependence on the Planck scale. Since the Higgs receives large mass corrections from every particle that couples to the Higgs field, the Higgs potential is very sensitive to any physics beyond the Standard Model [66]. There are terms in the Higgs mass correction that depend on the mass of the heaviest particle coupling to the Higgs. If this

scalar mass is very high, then it would be expected that the Higgs mass should be much higher than the observed 125 GeV, or that there are very exact cancellations (extreme fine tuning) to counteract this term [66]. Within the Standard Model, there is no explanation for how scalars remain relatively low mass (compared to the highest physical mass scales) without a specific symmetry protecting them [6].

Furthermore, the gauge boson, charged lepton, and quark masses depend on the Higgs mechanism, so changing the Higgs VEV would affect many other observables. This is a highly relevant issue, as the dependence of the Higgs mass on the highest energy cutoff scale does not explain why these fundamental particles remain at relatively low masses.

The question arises: why is the Higgs mass close to the electroweak energy scale, yet has extremely large loop corrections (even at leading order), with no physical explanation for this fine-tuning? Many models of physics beyond the Standard Model start to address this by introducing symmetries to prevent or cancel the loop corrections or lowering the Planck scale [4].

1.2 Physics Beyond the Standard Model

The Standard Model (SM) of particle physics includes all the known particles, and its predictions have been experimentally supported up to the TeV scale at high-energy colliders. However, many questions remain unanswered, as dark matter, gravity, and dark energy are not explained by the SM, yet are supported through experimental observations. Extensive observational evidence supports the existence of dark matter from galaxy rotation, galaxy formation, and gravitational lensing. The large velocities of stars around the center of galaxies cannot be explained only by the observed visible matter, and this discrepancy between visible matter and the observed effects of non-visible matter implies physics beyond the Standard Model. In addition, neutrino oscillations suggest that neutrinos are not massless, though they are massless particles in the Standard Model.

Furthermore, the Standard Model must be extended to explain physics at higher energies. The SM does not extend to the Planck scale, the highest energy scale known in physics, of $2.4 \cdot 10^{18}$ GeV. This scale is where quantum gravitational effects must be accounted for [66]. There are 16 orders of magnitude between the EW scale and the Planck scale, and it is possible that currently undiscovered physics exists at these higher energy scales. To explain these discrepancies, along with the hierarchy problem and divergence of the Higgs mass (discussed in Sections 1.1.3 and 1.1.4), many extensions to the SM have been proposed.

1.2.1 Limitations of the Standard Model

The previous sections detailed the Standard Model and illuminated some of its inconsistencies or shortfalls. Here, a few key issues are further discussed, along with the approaches of models proposing physics beyond the SM.

- *Gravity*: The Standard Model describes the particle origins of three of the four fundamental forces, without a particle explanation for gravity. The “graviton” is the hypothesized gravitational force carrier but requires modifications to the Standard Model. Making the SM consistent with general relativity generally requires the energy scales of the cosmological constant and spontaneous EW symmetry breaking to be similar.
- *Hierarchy problem*: The gravitational force and weak force differ by 10^{24} , without a clear underlying reason. This is related to the fine-tuning problem, as the calculated corrections to the Higgs boson mass are significantly larger than the bare mass.
- *Fine-tuning*: The renormalization (the process of applying corrections to the fundamental value) of the Higgs mass by the loop diagrams pushes the mass to very high values. If precise cancellations occur, the Higgs mass returns to the much lower observed value of 125 GeV. These cancellations are termed “fine-tuning”, as there is no other known reason why those parameters have the needed values.
- *Why are masses and coupling constants the values we see?* The Standard Model has 25 free parameters. These parameters are chosen to have values matching observation, but without a theoretical motivation [76]. However, there are hints at underlying symmetries, such as from the fact that the masses within each of the three fermion generations are similar.
- *Dark matter*: Observational cosmological evidence supports the existence of dark matter, a form of matter that interacts gravitationally but not electromagnetically (therefore is considered “dark”). Galaxy rotation curves, where the velocity is higher at higher radii than expected from the observed matter, and gravitational lensing both indicate a significant presence of non-visible matter and are strong evidence of dark matter. No particle explanation for dark matter has yet been observed.
- *Dark energy*: Measurements of the cosmic microwave background support the Λ CDM model, which has only 5% of the energy density of the universe as baryonic matter, 27% as dark matter, and the remaining 68% as dark energy, which contributes to the accelerating expansion of the

universe. Though dark energy does not currently directly impact the Standard Model, the constraints from cosmological observations on the composition of the universe do. As most dark matter is “cold” (non-relativistic), this constrains the neutrino masses to be under 1 eV [76].

- *Imbalance of matter and antimatter*: After the Big Bang, it is expected that equal amounts of matter and antimatter would be produced; however, we observe more matter than antimatter. An imbalance would require an interaction with baryon and lepton number violation, CP symmetry violation, and which occurs when the universe is not at thermal equilibrium (identified by Sakharov) [49]. CP violation is seen in the SM in the neutral kaon decay.
- *Neutrino mass*: The neutrino mass is much smaller than the mass of other fermions, meaning they would need a very small Yukawa coupling to the Higgs. Current experiments focus on neutrino-less double beta decay, which would only be observed if the neutrinos were Majorana particles [76].

Models of physics beyond the Standard Model are motivated by addressing these issues while remaining within the constraints set by the experimental limits on known Standard Model processes. In particular, this dissertation focuses on long-lived particles, a type of particle ubiquitous both within the SM and in many well-motivated theories beyond the SM, as detailed in Chapter 2. However, identifying long-lived particle events is difficult due to the large backgrounds and assumptions of the detector reconstruction, so implementing dedicated low-level techniques is critical. Without accurate identification at the first levels of data collection, these events are not saved for offline analysis, which is a limitation of the past approaches. To address this, unique properties of long-lived particle events are identified and used to design and implement specific data collection pathways, as detailed in Chapters 5 and 6. A first-of-its-kind analysis, relying both on the new data collection capabilities and upgraded detectors, is underway on Run-3 data, with details presented in Chapters 8 and 9.

Chapter 2

Long-Lived Particles

Physics beyond the Standard Model has not been observed at colliders, and new physics is expected to either be at energy scales above those currently probed at the LHC or in a region that is accessible but has been overlooked in past searches. Many models predict the existence of long-lived particles (LLPs) with kinematically accessible states resulting from the decay of these high-mass particles; therefore, LLPs are a natural direction to search for new physics at high energy scales. These models include split supersymmetry (SUSY), exotic Higgs decays, dark matter models (R-parity conserving models generally), and models with heavy stable charged particles. Long-lived particles are a promising direction for new physics searches, particularly as they have not been probed extensively in previous experiments. Many early searches for BSM physics at the LHC focused on missing energy (MET) or short lifetimes (tracker-based searches). Thus, they would be blind to particles decaying beyond the tracker but within the detector, a unique signature with no MET. Informed by the types of models the LHC has already searched for, it is now essential to explore LLPs. They present a unique signature that could be overlooked by the prompt and MET-focused searches while being physically well-motivated. In some sense, the models excluded by the early searches make searching in these new regions more critical, as the lack of earlier observation incrementally increases the probability of finding new physics in these unexplored regions. This chapter outlines the motivations for LLPs and detection possibilities at the LHC.

2.1 LLP Motivations

The Standard Model has particles across a large range of lifetimes, including some particles with “long-lifetimes”, as can be seen in the many orders of magnitude spanned by SM particle lifetimes in Figure 2.1 [72]. LLPs are also common in BSM physics theories and are possible for the same reasons as SM particles are long-lived: small couplings, scale suppression from a heavy mediator, or suppressed decay phase space (often due to small mass splitting). New, unobserved as of now, weak-scale particles may be long-lived due to these reasons, and this dissertation focuses on improvements to LLP triggering and analysis in the CMS experiment.




Figure 2.1: Examples of the range of lifetimes present in Standard Model particles, with $c\tau$ ranging from stable (electron), very long (proton), intermediate (muon), and prompt (Z-boson). Leptons are in red, mesons in green, baryons in blue, bosons in orange, and quarks in purple. Figure from [2].

The lifetime τ of a particle is the inverse of the decay width, Γ , which is given by [76]:

$$d\Gamma \propto \frac{1}{m} |\mathcal{M}_{decay}|^2 d\Pi \quad (2.1)$$

A small decay width Γ gives rise to sizable lifetimes for particles. To have a small decay width, either the available final state phase space ($d\Pi$) or the decay matrix ($\mathcal{M}$) must be small. The matrix element can be small due to small couplings or suppression by a heavy intermediary. The Standard Model contains examples of long-lived particles due to both mechanisms.

- *Scale suppression (with a heavy mediator)*: The muon, charged pion, and kaon all decay via an off-shell particle (the W boson mediator, as a virtual particle). The large mass dif-

ference between the W mass and these particles contributes to a small matrix element, as $\Gamma_{\pi^+} \propto \frac{g_W^2}{8\pi} \left(\frac{m_\pi}{M_W}\right)^4 m_\pi$ [71]. The charged pion decays to $\pi^+ \rightarrow \mu^+ \nu_\mu$ with a lifetime of $2.6 \cdot 10^{-8}$ s. Both the kaon and muon decays are also mediated by the weak interaction, producing relatively long lifetimes. The long-lived kaon K_L primarily decays either into three pions ($3\pi^0$ or $\pi^+ \pi^- \pi^0$), $\pi^\pm e^\mp \nu_e$, or $\pi^\pm \mu^\mp \nu_\mu$, while the charged kaon K^+ decays primarily into two or three pions or $\mu^+ \nu_\mu$. The charged kaons and K_L lifetime is around 10^{-8} s. The muon decays into $\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu$, with a lifetime of $2.2 \cdot 10^{-6}$ s [79].

- *Suppressed decay phase space:* The neutron has a lifetime of about 900 s, decaying to a proton, e^- , and $\bar{\nu}$. The decay phase space for the neutron is suppressed as the neutron is very close to the proton mass, contributing to the long neutron lifetime. The neutron mass is 939.6 MeV, while the proton mass is 938.3 MeV, a difference of only 0.14%, significantly constraining the decay phase space. The decay width is given by $\Gamma_n \propto \frac{g_W^2}{8\pi} \left(\frac{m_n - m_p}{M_W}\right)^4 (m_n - m_p)$, and the 1.3 MeV differences in the proton and neutron mass contribute to the small decay width [71].

In many models, the small coupling between the LLP and the decay products contributes to long lifetimes. BSM models often assume that the exotic particles couple weakly to SM particles (as these processes are expected to be rare), contributing to the long lifetimes.

Soon into Run 1 and Run 2 of the LHC, prompt, R-parity conserving SUSY was significantly constrained. However, LLPs are a “loophole” to this since standard tracking algorithms do not identify them. Searching in regions with low SM backgrounds is a powerful direction to search for LLPs. This leaves a few primary LLP directions to search: split SUSY, gauge mediation, hidden valley, and R-parity violation [3]. This search focuses on twin Higgs models, where the Higgs weakly couples to a LLP, which then decays back to standard model particles.

2.2 Experimental Signatures

Long-lived particles are an increasingly compelling direction for new physics searches due to the many well-motivated BSM models including LLPs, and the lack of evidence of BSM physics in prompt-based searches. The experimental coverage of LLPs is limited due to the assumptions governing many searches – primarily that particles originate from the collision point. This means we could overlook many signatures in which a LLP decays away from the collision point, producing a displaced jet signature.

LLP searches generally focus on a “signature” based approach, focusing on a particular signature

that may appear in multiple models. These signatures can be classified by decay location (tracker, calorimeters, muon systems) and final products (hadronic jets, photons, muon system showers). As LLPs have unique signatures distinct from Standard Model processes, searching for LLPs is an excellent way to probe physics beyond the Standard Model. This thesis focuses on a hidden sector model where an exotic Higgs boson decays to two LLPs, each decaying to two b quarks.

2.2.1 Distinct Features of the Final State

The long-lived particle experimental signature heavily depends on the mass and $c\tau$ of the LLP, along with its detectable decay products. Figure 2.2 shows a range of possible LLP decay signatures, with BSM particles shown in red.




Figure 2.2: A schematic of various LLP signatures in a general-purpose particle detector, such as CMS, with the detector cross-section shown. This work focuses on displaced jets and trackless, low-electromagnetic fraction jets. Figure from [55].

In particular, the tracker is very valuable in searches for LLPs with small $c\tau$, as displaced or non-pointing tracks distinguish the signal region. LLPs could show up with “disappearing tracks” where a charged LLP decays into an exotic particle, or “kinked tracks” where the charged LLP decays into one detectable and one exotic particle. Several innovative parameters, such as impact parameters, have been designed for these searches. Impact parameters measure how close the track approaches the primary vertex, as d_{xy} and d_z , and their significance is d_{xy}/σ_{xy} and d_z/σ_z , with σ_{xy} and σ_z

as the associated uncertainties. The impact parameter significance has been instrumental in track-based LLP searches. Track multiplicity is also used, with the caveat that displaced tracks lead to track reconstruction inefficiencies in simulation.

Beyond the tracker, calorimeters become central in LLP searches with a higher $c\tau$. Timing information can identify jets that are delayed relative to the primary proton-proton collision, resulting from either a slow-moving LLP (low β) or increased travel time due to the path length of the decay. The CMS electromagnetic calorimeter has a time resolution of about 100 ps, while the hadron calorimeter has a resolution on the order of 1 ns. Furthermore, the energy split between each calorimeter (ECAL and HCAL in CMS, for example) or between depth layers of one calorimeter can provide hints about LLPs decaying inside the calorimeter. The shower shape obtained from the calorimeter readouts is another valuable variable to identify a jet resulting from a LLP decay [59]. The shower may have an elliptical shape or a unique incident angle if it results from a LLP decay.




Figure 2.3: A schematic illustrating how the decay position of an LLP (tracker vs. calorimeter) impacts the observed jet characteristics. It is these different jet qualities that are used at the trigger and analysis levels to distinguish LLP candidate events from backgrounds, in Chapters 5 and 8, respectively.

2.2.2 LLP Landscape at the LHC

As shown in Figure 2.2, a wide range of LLP signatures is possible. The experiments at the LHC have done well in Run-2, particularly with tracker-based searches for precision at low $c\tau$ values, setting the best limits to date for mean proper decay lengths of 3-300 mm [20]. The muon system

has been utilized as a sampling calorimeter for LLP searches, adding sensitivity to a broad range of decays that create an isolated and localized cluster of hits in the muon system, setting limits on LLPs with mean proper decay lengths above 4 m [27]. More calorimeter-based searches have recently contributed to the intermediate lifetime region, utilizing the precision timing and energy measurements available from the calorimeters to identify trackless and delayed jets [25]. Figure 2.3 illustrates that LLPs decaying in the calorimeter have vastly different jet characteristics, and segmented calorimeters are a powerful handle in LLP identification. For instance, jets from LLPs decaying in the calorimeter are expected to be trackless, have clustered calorimeter hits, and have potential time delays from the decay of high-mass particles. The difference in the distribution of hits from a prompt and LLP jet is shown in Figure 2.4.




Figure 2.4: Prompt jets tend to have a widespread distributions of hits in the calorimeter system, while jets resulting from LLPs decaying in the calorimeters are trackless and have highly clustered energy deposits.

Sensitivity to the unique decay properties of LLPs is greatly improved by the increased granularity of the calorimeter segmentation available in Run-3, both in the longitudinal and radial directions, and utilizing this information at the data collection, reconstruction, and analysis levels expands the coverage of potential new physics. Chapter 5 details the use of the increased calorimeter information at the trigger level, while Chapters 8 and 9 utilize the 5D calorimetry capabilities (energy, timing, and positional information) for improved LLP reconstruction and in a Run-3 analysis. The model used in both trigger development and in the LLP search is reviewed in Section 2.3.

2.3 Exotic Higgs Decays to LLPs

The LHC was designed to measure the properties of the Higgs boson and, therefore, is a Higgs factory with high production rates. With the current experimental constraints on the Higgs, there are still significant possibilities of Higgs field couplings to BSM physics (often at high energy scales) that may be explored further. Given the Standard Model Higgs production rates, the CMS Run-2 combination upper limit for the branching ratio of Higgs to undetected is $\text{BR}(H \rightarrow \text{BSM}) < 16\%$ at 95% confidence level [23]. The model investigated in this thesis is an exotic decay of the Higgs boson (or decay of an exotic, high-mass Higgs-like boson) to two long-lived particles, each decaying to two b-quarks: $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$.

2.3.1 $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$

Higgs portal theories (also called hidden valley, hidden sector, or dark sector models) are a broad class of models where BSM particles couple to the SM-like Higgs boson, and are well-motivated as they are renormalizable. Exotic decays of the Higgs boson are a signature of many theories, due to the small width and thus the sensitivity of the Higgs phenomenology to new physics couplings, since large BSM branching fractions are possible [37]. As the Higgs decay width is narrow at 4.1 GeV in the SM, even a small coupling to an exotic sector could have a significant branching ratio, as $BR = \Gamma_{\text{specific}}/\Gamma_{\text{total}}$. This is a neutral naturalness model, meaning that the Higgs boson couples the SM particles to the hidden sector, where particles carry no color charge. Neutral naturalness models address the hierarchy problem by protecting the weak scale of the Higgs mass with discrete symmetries that avoid the Higgs mass divergence due to couplings to high energy scales [35].

A simple extension to the SM that provides coupling between the Higgs doublet and a hidden sector with long-lived particles is created by adding a real singlet scalar (S) that interacts with the SM Higgs through a renormalizable Lagrangian (Equation 29 in [65]):

$$\mathcal{L}_{\text{Higgs portal extension}} = -\frac{\epsilon}{2}S^2|H|^2 + \frac{\mu_s}{4!}S^4 + \mu_H^2|H|^2 - \lambda_H|H|^4 \quad (2.2)$$

If the Higgs doublet H and the singlet scalar S both have non-zero vacuum expectation values, then their corresponding physical scalar particles mix with the mixing angle $\sin \theta$, which is proportional to their vacuum expectation values and inversely proportional to the mass difference between the particles [65]. The coupling to SM fermions is also proportional to $\sin \theta$ in the Lagrangian, and therefore the process $H \rightarrow SS \rightarrow f\bar{f}f\bar{f}$ is possible, as shown in Figure 2.5. At small $\sin \theta$, the new

particle S is long-lived, contributed to through both a small coupling and suppressed decay phase space. The decay width of S to fermions ($S \rightarrow f\bar{f}$) is proportional to $\sin^2 \theta$, and the decay phase space is suppressed if S is not allowed to decay into other hidden sector particles [65]. This SM extension with a Higgs coupling to a long-lived particle is accessible at the LHC if the mass of S is less than half the Higgs mass, as then two long-lived particles can be pair-produced through an exotic Higgs decay.




Figure 2.5: Feynman diagrams of an exotic Higgs decay to two long-lived particles, producing multiple jets. In $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ production, the Higgs may be produced through the dominant gluon fusion or the non-dominant VBF or VH production modes. Any of these production modes may be targeted in trigger development.

2.3.2 Decays of an Exotic Higgs Boson

Along with the exotic decays of the Higgs boson presented above, similar models postulate exotic Higgs that decay via the same pathway (two long-lived particles, each decaying to two fermions). In this case, a similar singlet-scalar SM extension gives an exotic Higgs boson with the same production mechanisms as the Standard Model Higgs [2]. The exotic Higgs mass is a free parameter, and with the flexibility in the choice of m_H , m_S , and $c\tau$, a wide parameter space can be scanned.

The dominant production mode of the Higgs is gluon fusion, but the Higgs decaying to LLPs can also be produced through vector boson fusion (VBF) or Higgs-strahlung (VH) production, opening up the channels for triggering [2]. In the models considered in this work, the Higgs is produced through gluon-gluon fusion, the LLP is produced from an on-shell decay of the SM-like Higgs boson, and the LLP mass is constrained by $m_S \leq m_H/2$ (on-shell decay). The mass and lifetime values chosen for the analysis and trigger studies are listed in Chapters 5 and 7.

The exotic Higgs boson decays into two scalar long-lived particles, each decaying to two b quarks ($H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$), and has the unique signature of displaced vertices. The coupling of the LLPs to SM particles is inherited from the mixing with the Higgs, and thus the Yukawa weighted branching

ratio prefers decays to third generation fermions ($b\bar{b}$ or $\tau\bar{\tau}$) [37]. The SM branching ratios are $H \rightarrow b\bar{b}$ (58.1%) and $H \rightarrow \tau\tau$ (6.3%). Current limits set for $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ exclude branching fractions larger than 10% for $c\tau$ between 1 mm and 1 m, while less stringent limits have been set for lifetimes over 1 m [20]. The Run-2 branching fraction exclusion plots are shown in Figures 2.6 and 2.7.




Figure 2.6: Expected and observed 95% CL upper limits on the branching fraction of $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$, assuming the SM Higgs gluon-gluon fusion production cross section. Plot from [20], with data from Run-2 of the LHC.




Figure 2.7: Observed 95% CL exclusion limits on the branching fraction of $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$, assuming the SM Higgs gluon-gluon fusion production cross section. Limits are set from three different hadronic long-lived particle analyses. Plot from [43].

2.4 Improved LLP Search Opportunities

Long-lived particles are a promising direction for new physics searches. They are supported by well-motivated theories, and have not been excluded by previous collider searches. Improved experimental techniques can expand LLP sensitivity, as experimental coverage of LLPs is restricted at the trigger level, contributing to a key limitation in LLP searches. Currently, the trigger system primarily focuses on the decay of particles from the main interaction vertex. This means that we are blind to the new physics signature of many long-lived particles (LLPs) that decay to observable particles within the detector, yet are displaced from the main interaction point. In addition, the majority of past LHC searches relied on the assumption that particles outside the Standard Model would escape and thus not interact within the detector – leading to missing energy in the observed event. The delayed signature of LLPs provides an avenue to search for new physics independent of missing energy.

CMS has done well in LLP searches, but many unexplored regions remain. The $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ model is particularly challenging as the lack of dedicated triggers in the past has led to high H_T cuts or MET-based searches in Run-2, limiting the performance at lower Higgs masses.

By developing new triggers sensitive to both the time delay from the displaced LLP vertex and the delayed impact from the massive particles' finite velocity, we can capture otherwise inaccessible and unrecorded physics, and increase sensitivity to a wide range of physics models. The recent upgrade of the Hadron Calorimeter in the CMS experiment provides an additional handle to increase the LLP coverage at the trigger level and is the focus of this work. The upgrade and trigger implementation are discussed in detail in Chapters 3 and 5, and Chapters 8 and 9 detail the LLP analysis approach and results.

Detectors are generally designed for prompt physics searches, making LLP searches challenging, though detector-focused approaches also provide interesting opportunities for LLP event selection and search improvement. These opportunities encompass the trigger, event/object reconstruction, and data analysis approaches. Custom techniques are used to improve the data collection (trigger level) and analysis (new variables) to expand the LLP phase space covered.

Chapter 3

The Compact Muon Solenoid Detector

The Large Hadron Collider (LHC) is the highest energy particle collider in the world, built near Geneva, Switzerland, under the border between Switzerland and France. LHC proton-proton collision began in 2010 at a center of mass energy of 7 TeV. Since then, there have been three run periods with increasing beam energy. Currently, Run-3 is ongoing at a center of mass collision energy of 13.6 TeV. The beam energies and luminosities are designed to probe physics at the TeV scale. The accelerator complex consists of a 27 km circumference tunnel of superconducting magnets and beampipes, where particles are accelerated to close to the speed of light. There are four collision points around the ring where various detectors are located. The four main detectors are ATLAS (A Toroidal LHC ApparatuS), CMS (Compact Muon Solenoid), ALICE (A Large Ion Collider Experiment), and LHCb (LHC-beauty). CMS and ATLAS are general-purpose particle detectors, while ALICE focuses on heavy ion collisions and LHCb focuses on b-physics. The collider is described in Section 3.1, and the CMS experiment is described in Section 3.2, followed by specifics of the trigger system and offline reconstruction approach in Sections 3.3 and 3.4.

3.1 Large Hadron Collider

The European Organization for Nuclear Physics (CERN, in French *Conseil européen pour la Recherche nucléaire*) operates the LHC. The accelerator complex consists of a linear accelerator (LINAC 4,

the first step in proton acceleration), the proton synchrotron booster (PSB), and two initial accelerators, the proton synchrotron (PS) and super proton synchrotron (SPS). LINAC4 accelerates hydrogen ions, and the electrons are removed in PSB, so only protons circulate after this. After this, protons have reached 450 GeV and are injected into the main LHC tunnel, where beams are contained in two beam pipes and circulate in opposite directions. Here, the beams are accelerated to their full collision energy of 6.8 TeV, achieved in 20 minutes. Once beams are injected into the main collider, the “fill” lasts many hours, often 8 – 10 hours. These two beams are directed to collide at four interaction points around the ring, where the CMS, ATLAS, ALICE, and LHCb experiments are located.




Figure 3.1: Diagram of the CERN accelerator complex, with the LINAC4 (purple), proton synchrotron booster (PSB, light purple), proton synchrotron (PS, red), super proton synchrotron (SPS, blue), and finally the LHC (dark blue) illustrated. CMS is at the top. ATLAS, ALICE, and LHCb are placed much closer together in reality at “Points 1, 2, and 8”, respectively, while CMS is at “Point 5” (each point is $\frac{1}{8}$ of the ring). Figure from [11].

Both the collision energy and the luminosity indicate the capabilities of the collider. More energy

is available at higher collision energies to produce new particles; at a higher event rate, there is a larger chance to observe rare events. The number of events per second is dependent on both the luminosity (a quality of the accelerator, reported in $\text{cm}^{-2} \text{s}^{-1}$) and the event cross-section:

$$\frac{dR}{dt} = \mathcal{L} \cdot \sigma_p \quad (3.1)$$

with the cross-section as σ_p and the number of events per second as dR/dt [53]. Thus, the total number of events produced is $N = \mathcal{L}_{int}\sigma_p$. The integrated luminosity, $\mathcal{L}_{int}$, is the integral of the instantaneous luminosity ($\mathcal{L}$, maximum number of interactions per second) over time, given by $\mathcal{L}_{int} = \int_0^T \mathcal{L}(t)dt$.

The crossing angle and beam profile are important contributors to the total luminosity that can be achieved. For a Gaussian beam profile, with two beams colliding head-on, the luminosity is:

$$\mathcal{L} = \frac{N_1 N_2 n_b f}{4\pi\sigma_x\sigma_y} \quad (3.2)$$

where the number of particles per bunch are N_1 and N_2 (in the LHC, $N_1 \simeq N_2$), the revolution frequency is f , and the total number of bunches is n_b . σ_x and σ_y are the Gaussian profile of the beam [53]. A narrower beam, a more densely packed beam, and a greater number of colliding bunches will all produce a higher luminosity. The design luminosity of the LHC is $\mathcal{L} = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, and in the 2022 and 2023 portions of the Run-3 dataset, 67 fb^{-1} of integrated luminosity were collected. Femtobarns, with $\text{fb} = 10^{-39} \text{ cm}^2$, report the integrated luminosity collected in a run period.

Several practical effects lower the achieved luminosity, particularly the beam crossing angle, collision offset, and non-Gaussian beam profiles. A carefully calibrated crossing angle in the LHC is used to avoid unwanted collisions, as the proton bunches are very closely spaced. A precisely engineered set of dipole magnets to adjust the particle beam trajectory and quadrupole magnets to adjust the beam shape are placed along the LHC beampipe. Each of these magnets is superconducting and needs to be cooled with liquid helium to reach the required temperatures. The LHC crossing angle is $300 \mu\text{rad}$, and this non-zero crossing angle avoids unwanted collisions [53].

In the LHC, proton bunches are spaced every 25 ns, with about $1.15 \cdot 10^{11}$ protons per bunch [15]. Recording and analyzing the large number of proton collisions produced (about 1 billion per second, from an event rate of 40 MHz) is a significant challenge for the detector design of each experiment. During the bunch crossings, multiple proton-proton collisions occur, most of which are considered

pileup. The collision event of interest needs to be distinguished from the overlapping collisions, which becomes more difficult at higher luminosities.

3.2 Compact Muon Solenoid Experiment

The Compact Muon Solenoid (CMS) experiment records proton-proton collisions at the LHC at CERN in Geneva, Switzerland. The CMS experiment is about 21 m long, 15 m tall, and weighs about 14,000 tons. It is a general-purpose experiment designed to detect the Higgs boson and search for physics beyond the Standard Model.

3.2.1 Detector Coordinate System

The coordinate system CMS uses sets the x -axis to point to the center of the LHC ring, the y -axis upwards, and the z -axis along the beam.




Figure 3.2: Diagram of the coordinate system used by the CMS experiment, showing the x , y , and z directions relative to the LHC. Figure from [55].

The coordinates η and ϕ are defined using this coordinate system. ϕ is the radial coordinate (azimuthal angle) in the x - y plane, while η is related to the polar angle θ . η and ϕ are defined as:

$$\eta = -\ln \left(\tan \left(\frac{\theta}{2} \right) \right) = \tanh^{-1} \left(\frac{p_L}{|\mathbf{p}|} \right) \quad \phi = \tan^{-1} \frac{y}{x} \quad (3.3)$$

ϕ is the angle in the x , y plane (z is along the beam axis). η is the pseudorapidity and defines the position relative to the beam direction. p_L is the longitudinal momentum and θ is a polar angle, with $\theta = 0$ along the beam axis, corresponding to $\eta = \infty$, while $\theta = \pi/2$ corresponds to $\eta = 0$.




Figure 3.3: Pseudorapidity and θ correspondence, from [55].

Transverse energy is defined as $\vec{E}_T = E \cdot \frac{\vec{p}_T}{|\vec{p}|}$ and $|p| = p_T \cosh \eta$. Therefore,

$$E = E_T \cosh(\eta) \quad (3.4)$$

Due to the hyperbolic cosine scaling, the same E_T corresponds to a much larger energy deposit at higher η values.

Pseudorapidity is useful because the production of particles as a function of rapidity is constant, or the chance of producing a particle within a $\Delta\eta$ does not depend on the absolute η value, just the range. Differences in pseudorapidity are Lorentz invariant quantities, so this is a valuable quantity because it quantifies the distance between particles independent of the rest frame. Technically, this is true for rapidity, and pseudorapidity is rapidity in the massless limit.

3.2.2 CMS Sub-Detectors

The CMS detector comprises multiple sub-detector systems designed to detect specific particles. The innermost layer is the silicon tracker, which is used to identify the path of charged particles. Outside 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 the HCAL and provides a magnetic field of 3.8 T. The outermost part of the detector is the muon chambers, as muons will not be stopped in the calorimeters. A detailed description of the CMS detector is given in [16], and an illustration in Figure 3.4. Details of each sub-detector are covered in the following sections.

The radial extent of each subdetector is:

- Tracker: 2.9 – 125 cm

- Electromagnetic calorimeter: 129 – 177 cm
- Hadron calorimeter: 177 – 295 cm
- Superconducting solenoid: 295 - 400 cm
- Muon system: 400 - 720 cm




Figure 3.4: Illustration of each component of the CMS detector. From innermost to outermost, there are the silicon trackers, ECAL, HCAL, superconducting solenoid, and finally, the muon chambers. Note that the calorimeters have barrel and endcap sections. Figure from [70].

3.2.3 Tracker

The tracker is designed to measure the position and momentum of charged particles very close to the interaction point using ionization energy deposits. When charged particles traverse the silicon tracker material, their EM interactions produce a hit, which the digitizers process. Tracks are formed by linking these hits to describe the trajectories charged particles take, and since the tracker is inside the solenoid magnet, precise measurements of p_T are also achieved. The momentum is measured from the curvature of the particle's path in the magnetic field, experiencing a force $Bvq = mv^2/r$, and thus $p = Bqr$. With p in GeV, B in Tesla, and r in m, this becomes $p = 0.3Br$. Tracking is a key aspect of particle identification and is heavily used in Particle Flow (PF), covered in Section 3.4.

The tracker has multiple regions, each with a different detector design to handle the unique constraints. Three barrel pixel detector layers are used close to the beamline and interaction point (radius < 10 cm) to measure the secondary vertices and track impact parameters. The pixels are made of silicon modules with sensors $100 \times 150 \mu\text{m}^2$, accommodating the very high particle flux of around 10^7 s^{-1} and maintaining an occupancy of around 0.01% per pixel [15].

Further out, above a radius of 20 cm, ten layers of silicon microstrip detectors are used. Silicon detectors are advantageous as they have a high granularity and, thus, the precision needed for the high track density. In the intermediate region ($20 < r < 55$ cm), silicon detectors have a minimum cell size of $10 \text{ cm} \times 80 \mu\text{m}$, while in the outer region ($55 < r < 110$ cm) the cell size is increased to $25 \text{ cm} \times 180 \mu\text{m}$ due to the lower particle flux [15]. These choices maintain the occupancy below $\approx 3\%$.

3.2.4 Electromagnetic Calorimeter (ECAL)

The ECAL measures the energy of electromagnetic particles with photodiodes. When electrons or photons reach the ECAL, they produce scintillation in the 75,000 lead tungstate (PbWO_4) crystals, which allows for energy measurements with a high resolution, as nearly all of the electron and photon energy is deposited in the crystals. Lead tungstate crystals are advantageous in the ECAL as they are relatively radiation hard and emit most of their scintillation light within the 25 ns bunch crossing time [15]. The scintillation light is measured by avalanche photodiodes (APDs), which have a high gain to compensate for the low light yield of the crystals. In the endcap, vacuum phototriodes (VPTs) are used. The ECAL signal is digitized with a 12-bit ADC, and the noise has been measured to be around 40 MeV per channel in the supermodules. Above an electron energy of 50 GeV, the ECAL energy resolution reaches $\sigma_E/E = 0.5\%$ [15]. The high ECAL energy resolution is advantageous for the $H \rightarrow \gamma\gamma$ search. The ECAL time resolution is under 100 ps.

The radiation length of a material is a measure of the scattering of a particle as it traverses the material, specified as the mean length into the material at which the energy is reduced by $1/e$. ECAL provides about 30 radiation lengths, while the HCAL barrel adds about 60-100 radiation lengths (from the brass absorber, which has a radiation length of 1.49 cm) [15].

The ECAL channels are grouped into “supermodules” in the barrel region, with each supermodule covering $0 < \eta < 1.479$ and 20° in ϕ , for a total of 36 supermodules in EB. The endcap comprises two semicircle “Dees” of supercrystals [15]. The barrel ECAL crystals each cover 1° , or

$\Delta\eta = \Delta\phi = 0.0174$, and are arranged in an $\eta - \phi$ grid. In contrast, the HCAL crystals behind the ECAL instead cover 5° , or $\Delta\eta = \Delta\phi = 0.087$ in the barrel region.

3.2.5 Hadron Calorimeter (HCAL)

The HCAL comprises alternating layers of brass absorber and plastic scintillator and is designed to measure the energy of hadrons (composite particles made of quarks and gluons). An incident particle undergoes a nuclear interaction in the absorber, and the resulting shower of secondary particles is detected in the scintillation layers. The scintillation light is extracted on wavelength-shifting fibers and subsequently observed on SiPMs (silicon photomultipliers formed of Geiger mode avalanche photodiodes). The HCAL readout electronics process the information from each readout channel and perform charge integration and digitization (ADC). The front-end electronics also have a time measurement with 0.5 ns granularity, based on time to digital conversion (TDC), providing a resolution of about 1 ns. Thus, energy and time measurements are available from each HCAL cell.

HCAL is an “under-compensating” calorimeter because there is a lower response at low energy hits. This is due to the high percentage of nuclear interactions in the shower that do not cause scintillation. The HCAL response corrections discussed in Chapter 4 compensate for this effect to provide accurate measurements for jet reconstruction and particle flow. In addition, the HCAL response depends on how much of the jet energy is deposited in ECAL, which is accounted for in the recalibration.

The HCAL comprises four sections: barrel (HB), endcap (HE), forward (HF), and outer (HO). HF is placed in the far forward region, providing coverage from $2.85 < |\eta| < 5.2$, and starting at 11.15 m from the interaction point. HE covers $1.3 \leq |\eta| \leq 3$. HO has a similar $|\eta|$ coverage as HB, $|\eta| < 1.3$ and is placed outside the solenoid. HO is made of five rings, with 12 sectors in ϕ , and has two layers of scintillators (SiPMs) in the central region and one layer at higher η s. The central scintillators use the iron magnet return yoke as absorber material.

HB and HE are sampling calorimeters made of layers of plastic scintillator alternated with brass absorber (using brass from WWII shell casings). There are 17 layers in the barrel section of HCAL, each approximately 6 cm. The materials are primarily brass absorber (C26000, or cartridge brass, 70% copper and 30% zinc), plastic tile scintillator (Kuraray SCSN81 plate, except for layer 0, which uses Bicron BC408), and stainless steel support structures (inner and outer layers, for structural

support). The materials in each layer are:

- Layer 0: 9 mm scintillator, 61mm stainless steel.

Note that this layer has a larger scintillator directly behind ECAL to capture low energy showering particles from support material between HCAL and ECAL. In addition, there is no absorber before the scintillator (no dead material), but a significant shower can develop from the ≈ 1 interaction length provided by ECAL. HCAL depth 1 is unique because it also contains a neutral density filter.

- Layer 1-8: 3.7 mm scintillator, 50.5 mm brass
- Layer 9-14: 3.7 mm scintillator, 56.5 mm brass
- Layer 15+16: 3.7 mm scintillator, 75 mm stainless steel, 9mm scintillator

The layers are combined in four groups, forming the four depths of the upgraded HB. Depth 1 is comprised only of layer 0, depth 2 of layers 1-4, depth 3 of layers 5-9, and depth 4 of layers 10-16 as shown in Figure 3.5. Knowing the composition of each layer allows the number of interaction lengths in each depth to be calculated. An interaction length (λ_I) is the mean distance a hadron travels before being absorbed in the material due to a nuclear interaction. In HB, the average hadronic interaction length is 16.42 cm. Each barrel depth can be measured in terms of how many interaction lengths it contains, detailed in Table 3.1 and calculated from the specifications in [16].

At 90 degrees ($i\eta = 1, -1$), the HB absorber thickness is $5.82 \lambda_I$ [30]. The interaction lengths increase with η as $1/\sin\theta$ and reach $10.6 \lambda_I$ at $\eta = 1.3$ (end of HB) due to the larger volume of material traversed at the high η regions. The ECAL in front of the HCAL also contributes about $1.1 \lambda_I$ of material [16].

	Depth 1	Depth 2	Depth 3	Depth 4
λ_I (brass)	-	1.23	1.57	1.72
λ_I (steel)	0.58	-	-	0.72
λ_I (total)	0.58	1.23	1.57	2.44

Table 3.1: 14 layers of brass (743 mm) gives 4.53 interaction lengths. Stainless steel (136 mm) would need to provide 1.3 interaction lengths to achieve 5.82 total interaction lengths [30]. Given this assumption, this table estimates the interaction lengths per depth.

Run-3 HCAL Upgrade and Segmentation

During the Phase 1 upgrade of the CMS HCAL, the barrel was upgraded from hybrid photodiodes (HPDs) to silicon photomultipliers (SiPMs) from Hamamatsu Photonics. SiPMs have higher




Figure 3.5: Diagram of the radial distances to each depth layer in HCAL, computed from the values in [29].

photodetection efficiency and significantly larger gains than HPDs, resulting in an improved signal-to-noise performance. The improved performance of SiPMs over HPDs allows for an increase in the number of readout channels, allowing for enhanced depth segmentation [34]. The new front-end electronics support the increased granularity and record precise time and energy measurements at each depth layer (4 depths in the barrel, up to 7 depths in the endcap, in Figure 3.6). The trigger development and implementation presented here exploit this finer segmentation availability.

In Run-3 (2022-2025), there are four depth layers available in the HCAL barrel region (HB), and up to seven depth layers available in the HCAL endcap (HE) region, as shown in Figure 3.6. The depth segmentation of the barrel region is a new capability not available in Run-2. The information available from the HCAL at the L1 trigger includes tower transverse energy (E_T , reported in ADC counts) and TDC measurements at each cell. The total energy of a trigger tower is found by summing the energy of each cell in the tower, and the noise per trigger tower is ≈ 200 MeV. Run-3 uses a new pulse shape filter algorithm, PFA1', for HCAL energy reconstruction. This is designed to reduce the out-of-time (OOT) pileup contributions by subtracting a baseline from the time slice of interest and is based on a single time slice instead of the two time slice energy reconstruction previously used. PFA1' also demonstrates better agreement between the trigger primitive energy values and offline energy reconstruction [22].

HCAL Detector Geometry

Figure 3.6 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 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 (also $\Delta\phi = 0.087$, the same radial distance as each cell $\Delta\eta$).




Figure 3.6: In Run-3, there are four depth layers in HB and up to seven in HE (present since 2018). Figure from [38].




Figure 3.7: Before Run-3, there are one or two depth layers in HB and up to three in HE. Figure from [38].

Trigger primitives are created following the trigger tower geometry, defined by $i\eta$ and $i\phi$, the detector coordinates. In the HCal barrel, there are 72 $i\phi$ sections; however, the high $i\eta$ regions of the endcap

have 36 $i\phi$ sections. Thus, in HE, the energy per tower is split and assigned to two towers to keep the η , ϕ segmentation of the HCAL uniform for the L1 trigger [13].

Front-End Electronics: QIE11s and IGLOO2

In HB (HE), a single QIE11 chip has 16 (12) channels, and each IGLOO2 encodes data from 8 (12) channels. The QIE chip is a custom ASIC, the charge (Q) integrator and encoder, that measures the pulse energy and arrival time. 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) SiPM 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 next-generation clock and control modules (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 [17]. The physical layout of the boards is documented in [48], showing the 16-channel QIE11 and the IGLOO2 inputs, and similarly, the 12-channel HE QIEs are documented [47]. The data encoded by the IGLOO2 is transmitted to the back-end electronics, where it is further processed and eventually input to the trigger system.

The ngCCM is used to distribute the clocks and set the alignment of the QIE data and is implemented on a flash-based FPGA (field programmable gate array). The QIE clocks are used to scan the time delay to test the triggering pathway and determine HCAL alignment.

The front-end electronics are in the detector cavern and contain the RBX with RMs and ngCCMs. The front-end electronics continuously send data via optical links to the back-end components, which are outside the service cavern's shielding wall. The Lmap gives a 1-1 correspondence between RBX, RM, QIE card, QIE channel (front-end coordinates) and $i\eta$, $i\phi$, depth. In contrast, the Emap gives this correspondence between back-end coordinates (crate, slot, fiber, fiber channel) and detector coordinates.

Back-End Electronics: uHTR

The back-end electronics are the uHTRs (HCAL uTCA Trigger and Readout modules), Advanced Mezzanine Card (ACM13), and MCH, all contained in a uTCA crate. Each uHTR receives data from a full RBX (readout box, covering 20° in ϕ) but split by HB-only, HB/HE overlap, and HE-only uHTRs [61]. Multiple uHTRs are needed to take data from the entire RBX due to the complex tower mapping to each uHTR. There are 18 uTCA crates, each with 12 uHTRs at P5, for a total of 216 uHTRs (4x as many RBXs).

The uHTR has two flash-based FPGAs: the front reads out the front-end data and performs trigger processing, and the back transmits the data to the central DAQ [34]. The LLP trigger algorithm is implemented on the front uHTR FPGA and is described in detail in Section 5.3.1. The trigger payload specifications allow each data packet in HB to carry data from eight trigger towers, using 8 bits for energy and 6 extended bits per tower, detailed in [61] and [31].

3.2.6 Superconducting Solenoid

A high magnetic field produced by a superconducting solenoid with a 5.9 m inner diameter and length of 13 m is used in CMS to bend charged particles to determine their momentum. The magnet weighs 12,000 tons, making up a large portion of the total weight of CMS. The produced field is 4 T, and the high field allows for good momentum resolution [15]. For example, the momentum resolution is 10% for 1 TeV muons. To achieve this, the CMS solenoid is made of an aluminium-stabilized conductor with four winding layers [15]. The return yokes are contained in the iron of the muon system, causing charged particles to reverse their direction of curvature once reaching the muon system.

3.2.7 Muon System

The muon systems sit outside the rest of the detector and comprise barrel drift tubes (DTs) and endcap cathode strip chambers (CSCs) along with resistive plate chambers (RPCs). The muon system covers particles with $|\eta| \leq 2.4$. The DTs are four layers in the muon barrel region ($|\eta| < 1.2$), interleaved between the magnet steel return yoke layers [15]. Six layers of barrel RPCs are alternated with the DTs and return yoke. The return yoke for the magnet system uses 1.5 m of iron, split into four sections, and contributes a magnetic field in the muon system of 2 T. The CSCs cover the endcap $1.6 < \eta < 2.4$ region with four layers, while the CSCs and RPCs have three layers in the region $0.8 < \eta < 1.6$. Both in the barrel and endcap regions, the layers of RPCs, CSCs, or DTs are interleaved with the return yoke.

The CSCs, DTs, and RPCs are all forms of gaseous ionization chambers. The CSCs are made of gas regions and a gap region with cathode strips and anode wires set at a 90° angle to each other. When a particle traverses the system, the gas is ionized, and the resulting charged particles travel over the gap region, causing an electron avalanche that is detected [15].

3.3 CMS Trigger System

Buried within this enormous quantity of data collected at CMS from the LHC collisions are events potentially containing indications of new physics, and one of the primary challenges is saving the results of these events without requiring an exorbitant amount of memory. The current solution to this challenge is the CMS trigger system, which strives to identify the most physically significant fraction of the events for offline analysis. CMS utilizes a two-level trigger system, with a hardware trigger (Level-1, or L1) and a software trigger (High Level Trigger, or HLT). After the Level-1 Accept signal, the event is reconstructed based on the full readout of the detector. Given this information, the HLT decision is made, reducing the trigger output to 1 kHz [21]. Thus, 1/40,000 events are saved for offline analysis – and it is crucial to ensure the trigger is selecting events that further our offline physics analysis and avoid blind spots to new physics.

3.3.1 Level-1 Trigger System

The CMS trigger system uses a hardware trigger to select 100 kHz of events from the 40 MHz event rate to save for offline analysis. The L1 trigger relies on detector readouts from the calorimeter and muon systems and then uses preset algorithms, or “seeds”, to check events against the criteria set in the trigger “menu” [73]. If an event satisfies at least one L1 seed, then the Level-1 Accept signal is sent, initiating the complete detector readout that is sent to the HLT. HLT further refines the event selection to 1 kHz in 100 ms (compared to the 4 μ s latency for the L1 trigger), using additional information from tracking and reconstructed vertices.

The Level-1 trigger system combines information from the muon system (cathode strip chambers, resistive plate chambers, and drift tubes), as well as the ECAL and HCAL. Trigger primitives (local energy deposits) from the ECAL and HCAL are sent to the Layer-1 calorimeter trigger. After calibration, merging, and sorting the calorimeter energy deposits, the trigger towers are used at the Layer-2 calorimeter trigger to reconstruct physics objects, such as jets. Figure 3.8 shows the flow of information in the L1 trigger system.

In particular, jets (comprised of $\approx 75\%$ hadrons and $\approx 25\%$ photons) are one of the most essential objects at Level-1 and a focus of the trigger development presented in this thesis. Jets are reconstructed from a 9x9 region of HCAL and ECAL trigger towers, and a “chunky donut” algorithm performs pileup subtraction. The aim is to perform dynamic pileup subtraction on an event by event level, by estimating the surrounding energy density. The chunky donut algorithm sums the




Figure 3.8: Diagram of the information available to the L1 trigger. The L1 trigger takes input from the calorimeters (HCAL and ECAL trigger primitives) and the muon system, specifically the cathode strip chambers (CSC), resistive plate chambers (RPC), and drift tubes (DT) trigger primitives. These are then processed via the concentration preprocessing and fan-out (CPPF) and demultiplexing card (DeMux) before being sent to the global trigger. [73]

calorimeter tower energy in the four 3×9 towers surrounding the 9×9 L1 jet, and subtracts the sum of the lowest three of these from the total jet energy [9]. Alternative ϕ ring pileup subtraction algorithms have also been proposed to further improve energy measurements.

Energy-based quantities are also computed at Calo L2, specifically jet transverse energy (E_T), along with overall transverse energy (H_T , Equation 3.5) and missing transverse energy (MET, Equation 3.6). All are calculated in the transverse plane (xy plane), and MET is the momentum imbalance in an event, while E_T and H_T are the total transverse momentum of the jet or event, respectively.

Specifically, MET and H_T are:

$$H_T = \sum E_T(\text{all jets with } |\eta| < 2.5, p_T > 30) \quad (3.5)$$

$$MET = \sum \vec{E}_T(\text{all trigger towers, all } \eta) \quad (3.6)$$

H_T is the scalar sum of the energy of all jets in the event, giving a measure of the total event

energy. MET is “missing transverse energy” to quantify any energy that has escaped the detector (neutrinos, or as-yet undiscovered particles). At the online trigger level, jets are reconstructed based solely on calorimeter clusters (individual particles are not differentiated) without including tracker or muon information, and thus L1 H_T and MET are also independent of those two. However, full jet reconstruction and tagging (identifying τ , b, etc. jets) rely on the tracker.

Finally, events are formatted in a demultiplexer for the global trigger (μ GT), where final L1 processing is performed [73]. The six highest p_T jets in each half η are sent to the Global Trigger. At the μ GT, there are 12 jets, 12 e/γ , eight tau, and eight muon candidates. The μ GT runs the L1 menu, which can contain up to 512 algorithms, and makes a final “L1 Accept” decision determining whether an event is kept. The L1 menu has a rate of 100 kHz, which is then cut down to 1 kHz at the HLT, where algorithms are implemented in software.

3.3.2 High Level Trigger

The second level of the CMS trigger system is the High Level Trigger. This further reduces the events selected by the L1 trigger from 100 kHz to 1 kHz, utilizing offline reconstruction methods that take about 300 ms per event (as compared with the 4 μ s latency for the L1 decision) [12]. This is accomplished with a dedicated computing farm with significant CPU resources. The reconstruction performed at the HLT is a simplified version of the complete offline event reconstruction, and the output of the HLT is primarily for physics analysis, with a small portion dedicated to data quality monitoring (DQM). As opposed to the L1 trigger, which only uses inputs from the calorimeters and muon systems, the HLT combines information from the tracker, calorimeters, and muon system, along with linking information from each of those via particle flow (Section 3.4). This is beneficial for refining further and selecting specific event topologies.

The HLT rate is assigned to each object type or broad analysis category, such as Jet-MET or b-tagging, and Exotic or Higgs physics. The Exotica group has a total HLT rate of ≈ 750 Hz, shared rate of ≈ 250 Hz, and a pure rate of ≈ 60 Hz [12]. The total rate is all events triggered by that group, the shared rate is split equally among all groups that trigger the event, and the pure rate is for events *only* triggered by that group.

3.3.3 Run-3 Trigger Improvements

In Run-3, CMS focuses on diversification of the physics program, with new triggers, looser triggers, and increased use of the scouting and parking trigger opportunities [69]. We expect exciting physics

in Run-3; expanding our data collection and analysis strategies is a way to explore the potential new physics best. Run-3 has a slight jump in energy (small impact on physics) and aims to record 300 fb^{-1} by 2026 (likely to be closer to 250 fb^{-1} due to 2022 and 2023 operational issues, with the power cut and the LHC magnet quench accident). Thus, investments in new trigger and analysis strategies are critical.

The Phase 1 upgrade of the Level-1 trigger system was significant, including a complete replacement of hardware and software using uTCA technology (allowing for a high bandwidth backplane), FPGA processing cards, and high-speed optical links (allowing for multiple Gb/s) [75]. Along with detector upgrades, this has enabled a wide range of new triggers in Run-3, including improvements to B physics triggers, Higgs physics parking triggers, new seeds for $W \rightarrow 3\pi$, new LLP triggers (muon system-based and HCAL-based), and the addition of scouting triggers with lower energy thresholds.

Standard events have 2.5 kHz at 1.5 MB per event and are processed promptly. To allow for more events and data to be saved, “scouting” and “parking” are both options. Scouting is saving specific trigger streams with a reduced event size but at a higher rate, and parking is reconstructing a trigger stream opportunistically instead of promptly to allow for lower threshold triggers that would otherwise be too high a rate. Parking triggers have a rate of about 3 kHz at 1.5 MB per event, and the data is stored on tape and reconstructed when resources are free. Scouting saves events with a size of about 15 kB (100x smaller than standard events) and at a higher rate of 30 kHz. This is enabled by the CPU and GPU resource investment at HLT and is investigated at L1 for HL-LHC.

3.4 Offline Reconstruction

Once the readouts from each detector have been digitized, interpreted, and locally reconstructed, a global reconstruction begins. The reconstruction approach is based on particle flow, which combines information from each subdetector to identify individual particles in the event.

3.4.1 Jets

Jets are used to group reconstructed hadronic and electromagnetic showers within the CMS detector, resulting from collimated showers of quarks and gluons that hadronize. The standard jets in CMS are clusters of hits and particles within $\Delta R \leq 0.4$, which matches the barrel 9×9 trigger tower region

used for L1 jet formation [73]. “Large-R” jets with $\Delta R \leq 0.8$ are often used for complex objects with more spread. Jet reconstruction combines information from the tracker (identifying charged particles) and calorimeter clusters from the ECAL and HCAL, and a single jet may combine multiple calorimeter clusters [74].

This analysis uses AK4 CHS jets, which use the anti- k_T clustering algorithm to associate particle flow candidates within $\Delta R \leq 0.4$ of the jet (giving the acronym AK4). The anti- k_T algorithm considers the distance between particles and the jet d_{ij} , and each object (jet or particle) and the beam d_{iB} . The distance definitions for two objects (particles or jets) are [10]:

$$d_{ij} = \min \left(\frac{1}{p_{T,i}^2}, \frac{1}{p_{T,j}^2} \right) \cdot \frac{\Delta_{ij}^2}{R^2} \quad (3.7)$$

$$d_{iB} = \frac{1}{p_{T,i}^2} \quad (3.8)$$

$p_{T,i}$, η_i , and ϕ_i are the particle’s transverse momentum, rapidity, and azimuthal angle, respectively, and $\Delta_{ij}^2 = (\eta_i - \eta_j)^2 + (\phi_i - \phi_j)^2$. R is a radial parameter, and the radius of interest is set at 0.4 for AK4 jets. If the smallest distance is d_{ij} , then the particle is grouped with the jet, and if the smallest distance is d_{iB} , then the object is considered a jet and removed from the remaining list of objects to group [10]. The factor $\min \left(1/p_{T,i}^2, 1/p_{T,j}^2 \right)$ means that for particles equally spaced from the jet object, particles at a higher momentum will be selected first in the clustering, biasing the jet center toward higher momentum particles. This process is repeated iteratively until all objects are clustered. Finally, charged hadron subtraction (CHS) is performed on AK4 jets to remove charged hadrons associated with tracks that do not lead to the primary vertex.

Pileup impacts both the tracking and calorimeter clusters: higher pileup means more overlapping tracks and more vertices in the tracker, along with larger energy in calorimeters, either increasing the energy of the cluster of interest or creating a new cluster. During Run-3, CMS experiences about 54 pileup events per collision. Offline, pileup is mitigated through multiple strategies, one of which is CHS. This approach uses the tracks matched to a calorimeter cluster and removes particles with tracks not originating from the primary vertex to improve jet performance and energy measurements. Another approach is pileup per particle identification (PUPPI) that computes the likelihood that each particle in an event originated from pileup.

Jet energy scale and resolution corrections are applied to account for detector effects, such as detector regions with inactive material (at endcap and barrel transition region, for example). Ideally,

reconstructed jets match the truth jets well, and the ratio between them is a Gaussian centered at 1 [74]. The jet energy scale (JES) is applied to shift the Gaussian center to 1, while jet energy resolution (JER) adjustments are applied to achieve a narrow Gaussian width (about 5% above 1 TeV).

Of the energy in a jet, about 2/3 is from charged particles, while 1/3 is from neutral particles. As the tracker performs best at low p_T and the calorimeter performs best at high p_T , a combination of multiple detectors is optimal for reconstruction. The technique many experiments use to link track and calorimeter information is “Particle Flow”, or PF [74].

3.4.2 Particle Flow

Particle flow (PF) is the process of combining the information from all detector layers (tracks, calorimeter clusters, muon hits) to reconstruct the particle content of the event. PF was initially used at LEP and has been successful in the hadron collider environment primarily thanks to the high granularity of the detectors, which provide enough spatial information for particle identification. Particle flow takes in tracks, calorimeter, and muon system deposits to reconstruct a particle. The high magnetic field in CMS benefits the PF algorithm as it allows for the separation of charged and neutral particles. At the calorimeter level, PF aims to measure the direction and energy of particles, distinguish charged and neutral hadrons, identify electrons, and refine energy measurements for particles with low-quality tracks [19].

Particle flow uses multiple stages to identify different types of particles. At each stage, after a track or cluster has been used in PF identification, it is removed from the remaining elements used in the rest of the event identification. Leptons (electrons and muons) are used in the background selection for W+jets events, and an overview of the reconstruction of these particles is given next, followed by the reconstruction of hadrons.

Muons

First, muon candidates are identified based on a cluster of hits in the muon system that are matched to a charged track. Isolated muons are identified by associated tracks and calorimeter deposits within $\Delta R < 0.3$ of the muon system deposit, but not containing over 10% of the muon p_T [19]. The rejection of candidates with significant calorimeter energy prevents the misidentification of hadrons as muons. However, a stricter tight selection is applied if muons are in a jet (non-isolated muons) to reduce the high p_T hadrons that punch through to the muon system and create spurious hits.

Electrons and Photons

Next, electrons are identified from the tracker and calorimeter information and have similar energy patterns as photons due to electrons emitting bremsstrahlung photons or photons converting to e^+e^- pairs. Electrons are identified based on ECAL energy deposits matched to a charged track, with no corresponding HCAL deposit, and with a veto if the cluster is linked to three or more tracks. Isolated photons (which may either be converted to e^+e^- pairs or result from electrons via bremsstrahlung emission) are also identified, based on ECAL energy over 10 GeV and no associated charged track. In both electron and photon identification, calorimeter clusters with mostly ECAL energy are required, as the HCAL energy from cells with $\Delta R < 0.15$ must be under 10% of the ECAL supercluster energy [19].

Electron candidates are further identified using a machine learning algorithm, specifically a BDT, that uses variables relating to the associated track, ECAL cluster, and calorimeter energy distribution to identify isolated and non-isolated electrons. Photons have a looser PF selection, while tighter selections are typically applied at the analysis level.

Hadrons

At this stage, the remaining tracks are associated with ECAL and HCAL energy deposits and identified as charged hadrons, neutral hadrons, or non-isolated photons ($\pi^0 \rightarrow \gamma\gamma$). Neutral pions decay to two photons, and to reconstruct this accurately, ECAL clusters over 400 MeV and determined to be photons are paired. If the pair energy is over 1.5 GeV this is considered a π^0 , and is well validated with the photon pair invariant mass peak at 135 MeV consistent with a neutral pion [19]. Charged hadrons have a calorimeter cluster linked to a track but no muon signal. Thus, if the track energy approximately matches the calorimeter energy, this is considered a charged hadron, and the particle momentum and energy are measured from the track momentum. The high track resolution improves the energy and momentum measurements, motivating its use for particle energy measurements when possible. High track energies are masked, and though occasionally (10% of the time) these are matched to charged hadrons, their energy is generally better measured by the calorimeter than by the tracker. Charged hadrons can be mis-reconstructed as muons due to punch-through, i.e., when the hadronic shower extends to the muon system. To reduce this, muon clusters and their associated tracks are cross-checked against ECAL and HCAL deposits.

As photons and neutral hadrons have an ECAL or HCAL energy deposit but no associated track, if the track energy is less than the calorimeter energy (or there are no associated tracks), the

ECAL energy is reconstructed as photons. The remaining HCAL energy is reconstructed as neutral hadrons. A neutral hadron will deposit only 3% of the jet energy in ECAL, while photons carry 25% of the jet energy – motivating the assumption that ECAL energy is photons and the remainder can be assumed as neutral hadrons. After these steps (and subtracting the cluster energies already accounted for), any remaining excess ECAL and HCAL energy is considered as photons and neutral hadrons, respectively [19].

For the case where a combined ECAL and HCAL shower is not linked to a track (such as outside tracker acceptance of $|\eta| < 2.5$), the assumption that the ECAL energy is solely photons is no longer valid. Instead, this energy is grouped with the HCAL energy as either a charged or neutral hadron, and the remaining ECAL clusters without a HCAL cluster are photons. PF does not distinguish particles within the charged or neutral hadron categories.

Particle reconstruction is performed by particle flow. Considering how pions are reconstructed gives a helpful understanding of the various types of energy deposits:

- $\pi^0 \rightarrow \gamma\gamma$ (99% BR) is seen as ECAL energy and reconstructed as photons (not neutral hadrons).
- $\pi^\pm, K^\pm$ generally undergo nuclear interactions before decaying. The primary decays are to $\mu^\pm \nu_\mu$ with BRs of 99.99% and 64%, respectively. These are reconstructed as charged hadrons.

Hadrons (π, K , protons) will hadronically interact with the material, and the shower of absorbed secondary particles provides the signal. In a jet, one-third of the energy is generally from neutral particles (reconstructed as neutral hadron energy and photon energy), while two-thirds is from charged particles (reconstructed as charged hadron energy, electron, and muon energy). As a general approximation, a 40 GeV jet will have a particle multiplicity of 20 neutral hadrons, 8 charged hadrons, and 20 photons (or EM particles) [18].

3.5 Data Storage

The collisions at the LHC produce about 1 billion collisions per second, which translates to 1 petabyte of collision data per second. After the trigger systems reduce this by a factor of 40,000, the CERN Data Center processes the remaining data for reconstruction and archives it on tape. On average, 1 petabyte of data per day is processed and then distributed to various computing centers that provide CERN users access to the data via a distributed computing system. HTCondor and CRAB

(CMS Remove Analysis Builder) are two tools that allow CERN users to utilize grid computing resources for efficient data analysis by providing access to significant computational power over an extended time period (high-throughput computing, or HTC).

Chapter 4

Time Alignment of the Hadron Calorimeter

The Phase-1 upgrade of the CMS hadron calorimeter involved the upgrade of both the photodetectors and front-end ASICs and was completed in October 2019. Upgrading the photodetectors from hybrid photodiodes to silicon photomultipliers provides better signal-to-noise performance. These changes and an increase in the number of readout channels allow for increased calorimeter longitudinal segmentation. Instead of the previous 1 or 2 (2 or 3) depth-layer readout segments in the HCAL barrel (endcap) detector, there are now 4 (up to 7), as shown in the differences between Figures 3.6 and 3.7. The increased longitudinal segmentation allows for integrated charge (as ADC) and time (as TDC) measurements at each depth. These precise online pulse timing measurements are used for detector alignment, covered in this chapter, and for a novel LLP trigger, as detailed in Chapters 5 and 6. The results in these chapters are also presented in CERN CMS DN-2023/022 [63].

4.1 Time to Digital Converter (TDC)

A time-to-digital converter on the custom QIE11 ASIC chip provides high-precision timing information for each HCAL channel. This is recorded as a 6-bit TDC value for use in the trigger pathway and offline analysis.

4.1.1 Digitizing SiPM Signals with the QIE11

In the Phase 1 upgrade to the HB and HE, the scintillation light is transmitted through wavelength-shifting fibers to SiPMs. The SiPM signals are measured by charge integrator and encoder (QIE) ASICs in the front end (on detector) electronics, and the upgraded QIE11 chip performs both analog to digital conversion (ADC, via integration) and rising edge time-to-digital converter (TDC). The digitized signals from the QIE11 are transmitted by the IGLOO2 FPGA to the back-end electronics (uHTRs) [17]. This is performed in each 25 ns timeslice (TS) corresponding with the LHC collision rate of 40 MHz and spacing between bunch crossings. An ADC code (8-bits) and a TDC code (6-bits, with 0.5 ns precision) are recorded for each timeslice.

The QIE11 functionality is described in detail in [17], and a summary is given here. The SiPM current output is integrated over the 25 ns TS to determine the ADC value. This is reported in 256 charge 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.

The SiPM output (a current vs. time pulse) is compared directly to a flat current threshold to determine the TDC value. The threshold is an 8-bit programmable value set at $18.7 \mu\text{A}$ for Run-3 data taking. The time at which the current passes this threshold is the TDC value and is reported in a 6-bit digital output, giving 0.5 ns bins within the 25 ns bunch crossing.




Figure 4.1: Diagram of how a pulse shape is digitized on the QIE chip, resulting in a 6-bit TDC code to encode 50 time bins (with 25 ns of range) and invalid or error codes. Figure from [54].

4.1.2 HB 6:2 TDC LUT

The TDC information from the sample of interest (SOI, corresponding to $TS = 3$ out of 8 total timeslices) is used and is reported as a 6-bit TDC address from the QIE11 ASIC. Within the 25 ns bunch crossing, this encodes fifty valid time bins (0-49, in half ns steps) along with error and invalid codes. Specifically, TDC code 62 indicates that the pulse started high, and code 63 indicates that the pulse began low and remained low. Addresses 58-61 are error codes relating to the phase-locked loop state.

In HB, the IGLOO2 firmware LUT compresses the 6-bit TDC address into 2-bits, required due to bandwidth constraints [28]. This 6:2 LUT is defined from per- $i\eta$ and depth “prompt” and “delayed” TDC ranges. 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: $TDC \leq t_p$
- 01 slightly delayed: $t_p < TDC \leq t_p + 2$
- 10 very delayed: $t_p + 2 < TDC < 50$
- 11 invalid pulse: $50 \leq TDC$

The 6:2 LUT is applied per channel, and eight different LUTs may be loaded per IGLOO2, allowing for a separate LUT for each channel ($i\eta$, $i\phi$, depth). In Run-3, the LUTs were not $i\phi$ dependent. The LUT has $2 \cdot 2^6$ bits to encode the whole mapping, as documented in [46].

4.2 HCAL Time Alignment

Time alignment of the HCAL is vital to ensure that prompt signals arrive within the same time window across the detector and that delayed signals are correctly identified. There are two main components to ensuring the detector is aligned well: setting the QIE phases so that prompt signal arrival time is at the start of the timeslice and adjusting the TDC LUTs to place the prompt-delayed boundary properly.




Figure 4.2: 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.

4.2.1 QIE Phase Scans

Precise detector alignment is vital for the calorimeter's timing and energy measurements. The QIE phases set where the boundary between timeslices falls relative to the pulse arrival time. With the phase scan data, a new time alignment method using the TDC codes is developed and demonstrated to be very accurate. The alignment setting defines the boundaries of the sample of interest, from which the primary event energy and time are reported. TDC is based on the rising edge of the pulse, while the energy is based on the pulse integral, and thus, the TDC-based time alignment method is sensitive to the pulse's rising edge. Relying on the TDC information for calorimeter time alignment is ideal, as this 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.

During the phase scan, the fraction of each TDC code is evaluated at each scan point, and an optimal alignment is determined. The phases are adjusted so that the pulse arrival times are consistent across HB. For optimal alignment, the rising edge will be as close to the start of SOI as possible, while avoiding spillover into SOI-1. This ensures that a valid time measurement is recorded and that the energy measurement captures the bulk of the pulse. The energy filtering algorithm accounts for pileup by subtracting SOI-1 from SOI energy; therefore, it is essential that the pulse is not pushed into SOI-1; otherwise, some energy will be missed. Having the rising edge of the pulse placed too late in SOI would also cause a drop in the energy measurement since more of the energy will fall in

SOI+1.

The phase scan data has three main goals. First, it will be used to time align the HCAL based on TDC measurements and for the corresponding energy response corrections. Second, it is a valuable way to demonstrate and study the L1 LLP trigger performance. Finally, this data will enable the first depth-dependent pulse shape measurements in both HB and HE.

4.2.2 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 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. An illustration of the four timing ranges and the flat current threshold defining where the TDC is set is shown in Figure 4.2.




Figure 4.3: Diagram of the shifts produced by changing the QIE delays. As the QIE phase is shifted by -4 ns (i.e., increased by 8, as QIE phase values are in half ns steps), the pulse is seen earlier (i.e., toward SOI-1). Correspondingly, shifting the QIE phase by +10 ns (i.e., decreasing the QIE phase value by 20), the pulse is seen later in SOI.

The first QIE phase scan of Run-3 was taken in August 2022, relative to the 2022 operating phases, in 2 ns steps from -2 to 8 ns. This was during 13.6 TeV collisions with 2400 colliding bunches, and the data from this scan demonstrated timing misalignment present in 2022 (Section 4.2.4).

The November 2022 QIE phase scan was taken relative to recommended phases (proposed after the August scan) in 2 ns steps, from -4 ns, -2 ns, 0 ns, 2 ns, 4 ns, 6 ns, 8 ns, and 10 ns. This was with 13.6 TeV collisions at 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 this scan, resulting in the recommended 2023 phases. An illustration of the effect of shifting the pulse by -4

and +10 ns is in Figure 4.3.

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 phases were used in collisions data taking starting in June 2023.

4.2.3 Phase Scan Analysis

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 for a single HCAL cell 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 lumisections), or an emittance scan was in progress were removed as part of the data cleaning.

Figures 4.4 and 4.5 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 $i\eta$, shown here for $i\eta = -10$. The HCAL QIE11 chip reports a TDC value, 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. The prompt fraction increases as the next bunch crossing enters (QIE11 phase offset of 20 ns). 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, sufficient to minimize the timeslew effects of the pulse rising edge. The phase scan spans 25 ns, though it 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.




Figure 4.4: 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 [26].

The alignment procedure maximizes the percentage of arrival times in the TDC=00 (prompt) code. As seen in Figures 4.4 and 4.5, when the pulses are pushed early (below 0 ns), 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 in SOI rapidly drops, while TDC=01 and TDC=10 increase, as expected.

There is a Gaussian spread of arrival times of the pulse rising edge within each QIE phase setting, and the distribution of these arrival times is shown in Figure 4.4. The prompt window (TDC=00) is kept as narrow as possible for ideal alignment. 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 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




Figure 4.5: 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 that of no-TDC codes in SOI. If the TDC is set in SOI-1, no valid TDC is reported in SOI.

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 most significant effect (TDC=00 peak shifted on average 0.25 ns later in leading bunches). The TDC threshold is high enough not to 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, with the TDC=00 peak moving earlier by 3 – 5 ns. This leakage into SOI-1 is anticipated to lead to inefficiency in very energetic delayed signals and a slight prompt veto inefficiency. However, this is not a concern from the LLP trigger point of view, and the alignment conclusions are not changed as the effect is consistent across the detector.

Comments on Depth 1 and Changes for 2024

Depth 1 exhibits different behavior than the other depths, with a slightly lower prompt fraction (65-70%, compared with 90-95% in all other depths) and a wider Gaussian. In addition, the delayed signal plateau is reached about 3 ns later in depth 1 than in the other depths. To adjust for these differences, for 2024 data taking, it was decided to 1) delay depth 1 by 3 ns in the QIE phase (to allow the prompt peak to reach the high-efficiency plateau), 2) increase the prompt TDC LUT range to 0-9 ns in depth 1 only (3 ns shift needed to remain in phase with the QIE 3 ns change). All depths will have the slightly delayed range increased from 1 ns to 3 ns, as the 1 ns window is too narrow to provide detailed pulse placement information.

These differences in depth 1 are likely due to its placement directly behind the ECAL and lack of absorber material. Depth 1 is a very thin layer, made of only one scintillator layer: 9 mm of Bicron as compared to four to six 3.7 mm thick layers of SCSN81 in all other depths. This scintillator is placed directly behind the ECAL support structures without a dedicated HCAL absorber layer (the ECAL support structure acts somewhat like an absorber layer, though). Furthermore, depth 1 has a neutral density filter to attenuate the signal by a factor of 3 (note there is no NDF in HE depth 1, as it was removed in the 2017 year-end technical stop).

Consider showers of pions in ECAL: $\pi^0 \rightarrow \gamma\gamma$ with a 99% branching ratio. If a high-energy shower occurs in the deeper part of ECAL, HCAL depth 1 (layer 0) will detect some of the energy of this high EM fraction shower. By HB depth 2, 61 mm of stainless steel and multiple layers of brass are traversed, so the EM shower will be significantly reduced. Thus, the rising edge in depth 1 could be broadened by the contributions of EM shower spillover. Altogether, these choices mean that ECAL leakages may be present, and depth 1 suffers from the largest variations since it is most exposed and sensitive to shower fluctuations, contributing to the broader spread of arrival times. The +3 ns delay and increased prompt window in depth 1 are expected to make this depth behave more similarly to depths 2 through 4 and reduce the remaining pre-firing.

In addition, a pre-firing suppression algorithm is added to prevent the trigger firing in BX-1 when the signal falls in BX0. This filtering algorithm zeros timeslices except for the highest one (expected to fall in SOI or BX0) for very high-energy pulses. To do so, the reported trigger primitive energy is zeroed if the energy in the following timeslice is higher and above a threshold. This threshold is expected to be 5 GeV for 2024.

4.2.4 2022 HB Time Alignment Issues

The multiple phase scans taken were vital to understanding time misalignment in the HCAL barrel present in 2022. In particular, from the August 2022 QIE phase scan taken relative to the 2022 operating phases, it was found that the pulse arrival time in depth 1 was very late (prompt peak occurring at about -5 ns in Figure 4.6, so the pulses are arriving late in SOI), while the pulse arrival time in depths 3 and 4, particularly at high $i\eta$, was up to 10 ns early (prompt peak occurring at +10 ns in Figure 4.6, so the pulses are arriving in the previous bunch crossing). The 2022 and early 2023 alignment differences across depths are clear in Figure 4.6, as the position of the TDC=00 peak changes significantly with depth. With the adjusted alignment (Figure 4.4), the TDC=00 peak is consistent across the detector. This significant position-dependent misalignment causes errors in a wide range of quantities, such as energy response, pileup corrections, and discrepancies in the time domain responses (delayed signals would appear prompt in some regions, while prompt signals can appear delayed in other regions – and this cannot be well accounted for in the detector simulation). For an accurate energy response, the prompt part of the HCAL signal must be integrated, as the late part is disproportionately from the low hadronic response. The pileup correction algorithm subtracts energy from SOI-1 and harms the energy response if the signal starts in SOI-1. Finally, the wide distribution of signal arrival times can increase jet energy resolution systematics due to the mismatch with simulation. Clearly, the wide-ranging effects of HCAL time alignment mean that it is vital to set the alignment as accurately as possible.

After identifying the misalignment present in the 2022 phases from the August scan, a new set of adjusted phases is derived for the November 2022 scan. These phases are slightly refined (using a Gaussian fit to the TDC=00 peak and centering this at 0 ns) for the May 2023 scan, where alignment to within 0.5 ns is achieved. The phases derived for the May scan were deployed online in June 2023 to fix the HB alignment issues for the remainder of 2023 Run-3 data taking. The final adjustment to the QIE phases from 2022 to 2023 is detailed in Table B.6. Figures 4.7 and 4.8 illustrate the position-dependent timing misalignment of HCAL before the TDC-based phases are implemented (situation during 2022 and early 2023). The 2023 alignment is the peak at 0 ns offset in Figure 4.7, demonstrating the consistency in alignment achieved with the TDC-based method.

HB TDC LUTs in Run-3

The IGLOO2 HB TDC-LUT sets t_p , t_d from 4.1.2 for $i\eta$ 1-15, depths 1-4, and $i\eta$ 16, depths 1-3. The IGLOO2 LUT is re-programmable, and multiple LUTs have been used throughout Run-3. All LUTs




Figure 4.6: Distribution of TDC codes vs. QIE scan, for each of 4 depths across HB, for $i\eta = -10$ in the SOI, during the August 2022 QIE phase scan. The differing positions of the TDC=00 peak by depth show that the pulse arrival time alignment is not consistent across HB.




Figure 4.7: Position dependent timing misalignment, where pulses arrived very late in depth 1 and very early in depth 3 and 4 at high $i\eta$ values.




Figure 4.8: Alignment of recommended 2023 phases (adopted June 12, 2023) as compared with October 2022 phases. There is a significant spread in TDC-based pulse arrival times with operating phases. New phases resolve this, achieving time alignment to < 1 ns.

in Appendix B list the last TDC value for the prompt (00) range, or t_p from Section 4.1.2.

First, a v1 TDC LUT based on simulation is used. This sets t_p at the TDC value where 90% of the QDC background falls below, giving the LUT in Table B.1. Given that the misaligned phases were still in use for 2022 and early 2023 data taking, the TDC LUT was adjusted to compensate as best as possible. In October 2022, the LUT in Table B.2 was implemented to adjust the delayed ranges with the TDC distributions seen in data for the remainder of 2022 and the start of 2023 data-taking. The v2 LUT significantly widens or reduces the prompt range to account for the position-dependent alignment differences. The effective available delayed range with the v2 TDC LUT is shown in Table B.5. Deployed along with the HCAL time re-alignment in 2023, the v3 TDC LUT values set the prompt range evenly across the detector, encompassing 0-6 ns (Table B.3). The v4 TDC LUT (Table B.4) will be used for 2024 datataking, and will compensate for the wider shower spread observed in depth 1.

4.3 Time Alignment Validation

As 2023 is the first time TDC-based HCAL time alignment was used, the approach was validated by comparing it to alternative time alignment methods and understanding any discrepancies. In addition, the time alignment procedure was performed on data from a LED phase scan, which

provides a consistent pulse shape across the detector to eliminate the effects of shower propagation. Finally, a set of energy response corrections are derived and validated for deployment along with the TDC-based time alignment.

4.3.1 TDC Based Method of Time Alignment

Basing the HCAL time alignment on the TDC code distribution is ideal, achieving alignment to within half a nanosecond. The general approach is to align the phases such that the prompt peak is at a QIE phase offset of 0 ns, adjusting from Figure 4.6 to Figure 4.4, as was done between 2022 and 2023 of Run-3. However, suppose pulses are vastly misaligned such that no valid TDC codes are seen within the SOI in a timing scan. This was the case in high depths, high $i\eta$, during the 2022 August phase scan. In that case, an approximation can be made 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:

$$\text{TurnOn}(\mu, \sigma) = \mu + \sigma \cdot \text{TurnOn}(0, 1) \quad (4.1)$$

For a pulse known to be “well aligned” from TDC information, the Landau fit turn-on was found to be $\text{TurnOn}(\mu, \sigma) = 3.015$. All other channels were adjusted following:

$$\text{QIE Phase Adjustment (in ns)} = \frac{\Delta\text{Landau fit}}{0.04} \cdot \text{ns} \quad (4.2)$$

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, it only provides an approximate adjustment that is later refined with another phase scan and TDC code analysis. This method was used to determine adjustments needed to the phases from the August 2022 phase scan for the November 2022 scan, when the prompt peak was not clearly visible in the TDC code fraction distributions.

4.3.2 Charge Weighted Time

Charge weighted time (CWT) estimates pulse timing from the energy distribution and has been used as a time-alignment variable. CWT is defined by:

$$\text{CWT} = \frac{\sum_i Q_i \cdot i \cdot 25}{\sum_i Q_i} \quad i = 0 \dots 7 \quad (4.3)$$

where i is the timeslice (8 total timeslices, 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 maintaining a consistent SOI/(SOI+1) energy ratio 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. However, with differing pulse shapes across HB, CWT leads to inconsistent alignment. In contrast, the prompt TDC peak is linear in time delay, and there is no assumption on the pulse shape in a TDC-based alignment approach. The impact of differing pulse shapes is demonstrated by comparing the CWT values throughout phase scans in collision data and LED data.

CWT Calibration Procedure from QIE Scan

TDC is a fundamentally accurate pulse time measure 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 data.

CWT is computed for each depth and QIE delay, for pulses containing over 5000 fC of total charge across the 8 timeslices. CWT has a depth dependence and is non-linear with respect to clock delay. In particular, there is not a 1-1 correspondence between the change in clock setting and CWT result, as when the QIE phase changes by 4 ns, CWT does not also change by 4 ns. In Figure 4.9, it is clear that depth 1 behaves very differently than later 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. In this case, depth 1 pulses are set to arrive late, while depth 4 pulses arrive very

early, leading to large alignment offsets within the detector. This is consistent with observations of time alignment performed with CWT methods.




Figure 4.9: CWT from Equation 4.3, split by depth and QIE delay in HB, $1 \leq |i\eta| \leq 15$, and inclusive in $i\phi$.

The difference in late time fluctuations by depth is clear from the (SOI+1)/SOI and (SOI-1)/SOI plots, in Figure 4.10a and b. Depth 1 has less charge in SOI+1, due to more quickly falling tails. As CWT assumes the same pulse shape for each depth, achieving the same CWT value for each depth requires pushing the depth 1 pulse very late in SOI to increase the energy observed in SOI+1.


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, $w_{i,depth}$:

$$\text{CWT}_{\text{depth}} = \frac{\sum_i Q_i \cdot i \cdot 25 \cdot w_{i,depth}}{\sum_i Q_i \cdot w_{i,depth}} \quad i = 0 \dots 7, \text{ depth} = 1 \dots 4 \quad (4.4)$$


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

$$\text{Calibrated CWT}_{\text{depth}} = \frac{1}{m_{\text{depth}}} \cdot \text{CWT}_{\text{depth}} + b_{\text{depth}} \left(1 - \frac{1}{m_{\text{depth}}} \right) \quad (4.5)$$


where the CWT at QIE phase offset 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 4.10: Plots of the original (a, b) and re-weighted (c, d) SOI $\pm$ 1 to SOI energy ratio, $1 \leq |\eta| \leq 15$, and inclusive in $i\phi$.

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

$$w_{i,depth} = [1, 1, w_{i=2,depth}, 1, w_{i=4,depth}, 1, 1, 1] \quad (4.6)$$

$$w_{i=2,depth} = \frac{\text{ideal}(i=2)}{\frac{SOI-1}{SOI}} \quad (4.7)$$

$$w_{i=4,depth} = \frac{\text{ideal}(i=4)}{\frac{SOI+1}{SOI}} \quad (4.8)$$

where $\text{ideal}(i=2)$ and $\text{ideal}(i=4)$ are the selected ideal charge ratios between (SOI-1)/SOI and (SOI+1)/SOI, respectively, at QIE phase of 0. Deriving the weights based on (SOI $\pm$ 1)/SOI ratios normalizes the pulse shapes to be uniform in shape across depths, aligning with the CWT assumptions. Given these weights, the SOI-1 and SOI+1 plots are re-weighted, in Figure 4.10 c and d. This is done with selected ideal ratio values of $\text{ideal}(i=2) = 0.05$ and $\text{ideal}(i=4) = 0.45$. Table 4.1 lists

the determined weight values.

$w_{i,depth}$	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 4.1: 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 0 ns, 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 4.2: Slopes for the collisions scan analysis.

The slope m_{depth} in Table 4.2 is very close to 0.5 for all depths. The depth dependence is relatively small after the re-weighting, however, the fact that the slope differs significantly from unity further indicates that CWT is non-linear in QIE phase offset. Along with the dependence of CWT on pulse shape, these two factors make un-weighted CWT an imprecise time estimator.

Using the re-weighting and calibration, a final CWT is determined, plotted in Figure 4.11. 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 the QIE phase.

Timing in LED Runs

The behavior of the HCAL during a LED scan was understood from a scan taken in a local run in November 2021, used to verify the new TDC LUTs. This is a similar scan as the timing scans in 13.7 TeV collisions used for HCAL alignment, but is run during a local HCAL data-taking run where a LED illuminates the SiPMs, and the timing of the LED pulse was changed throughout the




Figure 4.11: Re-weighted and calibrated CWT from Equation 4.5, split by depth and QIE delay in HB. A depth-dependent offset has been applied in b to shift the distributions to 0 ns. The residuals show that the re-weighting and calibration procedure achieved agreement with the QIE phase changes to within 0.5 ns over the range of the QIE scan.

run. The Delay25¹, a programmable CMOS delay line that can be set to 0.5 ns intervals, is used for the LED scan. This scan was performed with a monochromatic LED pulse.




Figure 4.12: Plots of the TDC = 01 peak during a LED delay scan for individual channels.

The TDC = 01 peak is particularly helpful for understanding channel alignment since it is a 1 ns wide bin. As seen in Figure 4.12a this peak is just 1-2 ns wide (resolution limited by the 1 ns step size

¹<https://proj-delay25.web.cern.ch>




Figure 4.13: 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.

of the LED scan) for a single channel. When all $i\phi$ are combined, the TDC=01 peak is significantly widened.





Figure 4.14: 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 [67].

When split by $i\phi$, a pattern of 2 high - 2 low LED delay settings (in nanoseconds) is observed, as seen in Figures 4.12b and 4.13. 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 4.14. The cable length difference between the calibration unit and each half backplane causes the 5 ns delay seen in Figure 4.13, and this is an artifact of the LED scan – i.e., 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 LED data analysis, by grouping together $i\phi$ ranges with the same cable length effects.



(a) Energy ratio and TDC peak time correlation.

(b) Linear fit on depth 2 to energy ratio and TDC peak time correlation, with residuals.

Figure 4.15: Correlation between LED scan times at which TDC = 01 peak was first reached, and when $\frac{\text{ADC3}-\text{ADC4}}{\text{ADC3}+\text{ADC4}} = -0.2$ was first reached. This is for all $i\eta$, $i\phi$, and separated by depth.

Figure 4.15 demonstrates that the time alignment of the detector using TDC is equivalent to the previous method (ADC pulse shape). A pulse shape timing correlation between TDC (01 peak) and ADC ratio shows that ADC ratios and TDC LUT results both give precise and correlated timing information. This plots the LED scan time of TDC=01 peak vs. time of $\frac{\text{ADC3}-\text{ADC4}}{\text{ADC3}+\text{ADC4}} = -0.2$. 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.

Validation of CWT Calibration with 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 – set so that 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, but a consistent shape

across HB. As the LED is shined directly onto the SiPMs, the effects of hadron showering and dispersion causing 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 4.16: Fraction of TDC codes through 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 demonstrate that these phases align the LED pulse arrival times well. QIE phase of 10 ns is excluded from the CWT analysis, as the start of the prompt peak from the next BX is visible here. 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 4.16. The standard CWT from Equation 4.3 is calculated for the LED pulses, and plotted by depth and QIE phase in Figure 4.17. 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 listed in Tables 4.3 and 4.4 are within 10% of 1 (excepting




Figure 4.17: CWT from LED data, for $i\eta = -3$ to -1 , and $i\phi = 1, 2, 5, 6, 9, 10 \dots 69, 70$.

depth 1, for $i = 2$), and the slopes are within 5% of 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. Additionally, after the re-weighting and calibration procedure, the LED data agrees with the QIE phase change to within 0.2 ns. This validates the CWT re-weighting and calibration procedure by demonstrating that the weight and slope do not significantly differ from 1 when run on a uniform pulse shape.

$w_{i,depth}$	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

(a) Weights using ideal($i = 2$) = 0.005 and ideal($i = 4$) = 1.0. For $i\eta = -3$ to -1 , excluding $i\phi = 3, 4, 7, 8 \dots$

$w_{i,depth}$	Depth 1	Depth 2	Depth 3	Depth 4
$i=2$	1.17	1.08	0.95	1.00
$i=4$	0.95	0.98	0.98	1.04

(b) Weights using ideal($i = 2$) = 0.005 and ideal($i = 4$) = 0.9. For $i\eta = 1$ to 3 , excluding $i\phi = 1, 2, 5, 6 \dots$

Table 4.3: Weights $w_{i=2,depth}$ and $w_{i=4,depth}$, determined with ideal($i = 2$) = 0.005 and varying ideal($i = 4$) (at QIE phase offset 0 ns) from LED scan data. The analysis is split by $i\phi$ groupings. 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 4.4: Slopes for the LED scan analysis for $i\eta = -3$ to -1 .


4.3.3 Energy Response

Altering the HCAL time 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 have a reduced energy measurement, 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. A corresponding set of energy response corrections is concurrently deployed to minimize the downstream effects on the L1 trigger and jet energies from an HCAL alignment adjustment.


The change in average pulse energy vs. phase change is not a precise enough variable from which to estimate the response corrections. 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.

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 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, capturing the full early energy.

For example, Figure 4.18 shows the fraction of TPs over 4 GeV, which is maximized below a phase offset of 0 ns when the pulse starts in SOI-1 but still deposits significant energy in SOI. The average energy in SOI-1 is minimized close to a phase offset of 0 ns. These figures illustrate that time alignment is more subtle than simply maximizing energy recorded, as this loses sensitivity to the pulse rising edge. Additionally, a small change in phase setting has a large impact in the energy split between neighboring BXs.



(a) Fraction of TPs over 4 GeV in SOI (TP TS2) vs. QIE phase offset. As the pulse moves later, the integrated charge drops.



(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.

Figure 4.18: 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 ns produces ideal alignment without causing the pulse to spill over into SOI-1. $i\eta = 4$ is shown and is representative of the HB.

Hit Occupancy Method

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

$$\text{Occupancy} = \frac{N_{hits}(\text{over 4 GeV MAHI energy})}{N_{hits}(\text{over 0 GeV MAHI energy})} \quad (4.9)$$

MAHI stands for Minimization At HCAL Iteratively and is the local reconstruction algorithm for HCAL pulses. The inputs are three consecutive time slices (SOI-1, SOI, and SOI+1) that are fit to a pulse shape template, and a system of linear equations is solved with iterative minimization of the resulting χ^2 value. The uncertainties on the pulse shape are used to account for noise. The result of MAHI estimates the energy contained in a pulse and is used as the reconstructed per-cell energy at HCAL.

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.

Once the hit occupancy as a function of QIE phase offset is found, the phase-dependent factor needed to scale this to achieve flat occupancy is found. These factors are the “energy scales” for uniform




Figure 4.19: HB hit occupancy ≥ 4 GeV in MAHI energy for $i\eta = -14$ (representative of other regions).

occupancy vs. phase. 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:

$$\text{Occupancy} = \frac{\int_4^\infty \text{MAHI energy}}{\int_0^\infty \text{MAHI energy}} \quad (4.10)$$


$$\text{Scaled Occupancy} = \frac{\int_{4 \cdot x}^\infty \text{MAHI energy}}{\int_0^\infty \text{MAHI energy}} \quad (4.11)$$

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, the “energy scale” values are scanned for each detector region to determine what gives the same occupancy as the reference value.

$$\text{Energy}_{\text{new}, y \text{ ns}} = \frac{\text{Energy}_{\text{initial}}}{\text{scale}_{y \text{ ns}}} \quad (4.12)$$

The energy scale shows how each phase can be scaled to a reference phase and allows us 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:

$$\text{Energy}_{\text{old}, x \text{ ns}} = \frac{\text{scale}_{x \text{ ns}}}{\text{scale}_{y \text{ ns}}} \cdot \text{Energy}_{\text{new}, y \text{ ns}} \quad (4.13)$$



(a) HB energy scale for MAHI energy to give flat hit occupancy, from Equation 4.12.

(b) HB scaled hit occupancy, $\text{Energy}_{\text{new},y\text{ ns}}$ from Equation 4.13.

Figure 4.20: HB energy scale and scaled occupancy for $i\eta = -14$, inclusive in $i\phi$.

The change in the response corrections when moving from QIE phases at x to y is given by the factor $\text{scale}_{x\text{ ns}}/\text{scale}_{y\text{ ns}}$. 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: The signal slightly decreases at high $i\eta$ because of the large late time component.
- Depth 4: There is a slight decrease in the signal in the central barrel 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 arrived very early (into SOI-1, so energy was missed in the integrated signal of SOI).

The HB energy response corrections corresponding to the 2023 TDC-based QIE phase alignment are listed in Table B.7.


Validation of Method in HE

The hit occupancy method to estimate changes in the energy response corrections is new, and before deployment this method was validated in HE. HE had a QIE phase adjustment done between 2022 Era D and Era E, with Era D placing the QIE phases at an offset of about 4-5 ns on the 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 response corrections are generally reduced by 10-20%, which is expected as the new HE alignment increases the energy response by placing more of the prompt signal within the inte-



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



(b) HE energy scale for MAHI energy to give flat hit occupancy, from Equation 4.12.



(c) HE scaled hit occupancy, $\text{Energy}_{\text{new}, y \text{ ns}}$ from Equation 4.13.

Figure 4.21: 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.

gration 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.

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%.

4.4 Far-Reaching Impacts of HCAL Timing

The misalignment of pulses in HCAL in 2022 and early 2023 contributes to pre-firing at L1, which occurs when the L1 trigger conditions are satisfied *before* the BX of interest. Misaligned HCAL has pulses arriving very early (up to 10 ns) at high depth, high η , which places significant energy in BX-1 (enough to pass the L1 trigger conditions) and contributes to pre-firing. The large reduction of L1 pre-firing after the time alignment in Figure 4.22 is a success of the new HCAL TDC-based time alignment.




Figure 4.22: Fraction of pre-firing in Run-2 and Run-3, split by runs before and after HCAL timing alignment. The pre-firing is significantly reduced after the time re-alignment (run 368822) [68].




Figure 4.23: Fraction of pre-firing in Run-3, split by runs before and after HCAL timing alignment. The pre-firing is significantly reduced after the time re-alignment (run 368822) [39].

Most pre-firing is at $\eta \approx 1.2$, corresponding to $i\eta = 12 - 16$, where the HCAL alignment change is the most significant. In particular, two isolated HCAL deposits with $p_T > 30$ GeV (a dijet) can fire the tau trigger, causing pre-firing if this occurs in BX-1. This impact is particularly large at high mass dijet events, which often have $|\eta|$ near the edge of the barrel. This high η region is most misaligned, further increasing the likelihood of the tau trigger pre-firing.




Figure 4.24: Distribution of pre-firing by $i\eta$ in Run-3, before the HCAL alignment. The high $i\eta$ regions ($i\eta = 15$ is $\eta = 1.26$) contribute most to pre-firing and have the largest alignment correction. [68].

There are additional impacts on the jet energy. With the misaligned HCAL, 10% of the energy leaks into SOI-1 (for TeV jets). This causes a 5% drop in the energy scale of high energy (TeV scale) jets, due to the deposit of ≈ 50 GeV of energy in BX-1. This effect can be corrected for in reconstruction, while the pre-firing can be modeled, but not corrected for in data [44].

In addition, the HCAL energy resolution is improved by the TDC-based time alignment methods. After the alignment, the energy resolution of the barrel is improved by 10%, from 20% to 18% (and to 16.6% with corrections) [56].

4.5 Summary

A calorimeter alignment method relying upon TDC has been developed and validated in collisions and LED scans to evaluate the impact of varying pulse shapes. The new alignment deployed in

June 2023 used the TDC-based phases and decreases in pre-firing have been seen. An additional alignment method based on re-weighted pulse shape distributions is refined and calibrated with the higher precision TDC to correct for the pulse shape dependence of CWT. The pulse shape timing method provides additional avenues for calorimeter timing, such as corrections when a clock scan cannot be done, or for use in software triggering and pulse reconstruction.

Chapter 5

Level-1 Delayed and Displaced Jet Triggers

A delayed jet hardware and software trigger using the upgraded Phase-1 HCAL is developed and 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 aims to double the integrated luminosity of Run-2, which resulted in an integrated luminosity of 150 fb^{-1} . The trigger targets the identification of LLPs, utilizing the new capabilities of the CMS hadron calorimeter. In particular, the HCAL behaves as an extended sampling calorimeter with depth segmentation and is also unique in being a timing detector with nearly 4π coverage. Both the depth segmentation and timing measurements are used in the trigger decision. Jets containing multiple hits either with a delayed arrival time or with significant energy deposited in the higher depth layers are selected by this trigger for further offline analysis. This LLP timing trigger has been implemented in hardware and software for Run-3 and targets both the delayed arrival time of LLPs and unique energy deposits from LLPs.

The HCAL Phase-1 upgrade allows for the deployment of a new delayed and displaced jet trigger in Run-3, designed to capture decays from long-lived particles with lifetimes from 0.5 - 6 m, such that the LLP decays are before or within the HCAL volume. These new triggers expand the ability to trigger on hadronic LLP events at lower thresholds (H_T , jet E_T , etc.) and provide sensitivity to LLP decays from the tracker through calorimeter volume.

The developed Level-1 (L1) delayed jet trigger relies upon the new detector capabilities and re-

programmable firmware to significantly improve the LLP physics reach. The trigger identifies delayed hadronic jets using TDC (time to digital conversion) and ADC (analog to digital conversion, energy digitization) information as input. The complete 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 transmitting information from HCAL through L1. 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 allows for a comprehensive search for new physics with few underlying assumptions.


Trigger development involves three main stages: LLP signature identification, trigger algorithm development, and trigger implementation in firmware. Following firmware and software development, the L1 LLP seed is used at HLT (Chapter 6), and trigger performance in collisions is evaluated. This chapter covers the implementation of a hardware timing trigger, detailing the relevant hardware and firmware upgrades, trigger algorithm development, trigger commissioning, and performance in collisions.

5.1 Trigger Development

Detecting LLPs presents a unique challenge, as the standard algorithms to reconstruct events assume particles originate from the main interaction vertex, and this assumption is not valid for LLP decays. LLPs decay away from the primary proton-proton collision interaction point and, therefore, have unique:

- *Time stamps*: delayed relative to particles originating from the primary vertex due to the path length difference traveled by the LLP and decay products compared with the shortest straight line path length (illustrated in Figure 5.1a)
- *Energy deposits*: often large energy deposits in higher layers of the calorimeter, with little energy deposited in the lower layers (illustrated in Figure 5.1b)

The new LLP trigger is sensitive to both time delays resulting from the path length difference traveled by an LLP and its decay product and from the slow velocity of the LLP (Figure 5.1a). In addition, the trigger maintains sensitivity to LLPs decaying within the detector volume, creating large energy deposits in higher layers of the calorimeter (Figure 5.1b).



(a) Diagram of the time delay resulting from an 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 5.1: Diagrams showing the Run-3 depth segmentation in HB (4 layers) and HE (up to 7 layers), and the unique signatures created by LLP decays in or before the HCAL. Figures adapted from [38].

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 indicative of an 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, i.e., TDC=01 or 10, and energy > 4 GeV) and not including any energetic prompt cells (in the prompt range, i.e., 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 must have at least 2 LLP-flagged trigger towers within the 9x9 jet region and pass jet energy requirements. Event energy thresholds are also required to control rates at the L1 trigger.

5.1.1 HCAL Trigger Pathway

The calorimeter trigger towers 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 tower, such as summing the energies deposited in each depth. In addition to the standard TP generation in HB and HE, 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 characteristic of signals from long-lived exotic particle decays.

The HCAL uTCA Trigger and Readout modules (uHTR) firmware sets the six feature bits per trigger tower, which are sent to L1 for the delayed jet trigger. Specifications for the uHTR firmware are outlined in Section 5.3.1. Following the uHTR, to stay within bandwidth constraints in the Layer-1 and Layer-2 trigger pathways, one bit per trigger tower will send the delayed or displaced LLP signal information to Level-1 Global Trigger, where a delayed/displaced jet flag is set.

5.1.2 TDC Considerations – Threshold and Amplitude Dependence

Configurable look-up tables determine which time windows within the bunch crossing of interest are represented by the available TDC codes. These time window boundaries have a granularity of 0.5 ns, detailed in Section 4.1.2. TDC values are set once the current pulse passes a fixed current threshold. To ensure this threshold is set at a reasonable value, the rate of TDC=62 (error code where the pulse starts high) is evaluated at the nominal threshold ($18.7 \mu A$), 2x, 3x, and 4x nominal threshold. Ideally, the TDC threshold will be set as low as possible while maintaining an acceptable error code rate. In HB MC, the TDC=62 rate is $\leq 3\%$ in all depths for all TDC thresholds tested. Since HB has much lower rates than HE, and since 2018 HE data can be used for comparison, the TDC = 62 rate in HE is focused on in determining the Run-3 TDC threshold.

In 2018, HE was upgraded to use SiPMs and the QIE11 with TDC timing capabilities at a TDC threshold of $18.7 \mu A$. The error code rate in 2018 data is compared to the Run-3 simulation. Single muon and isolated bunch data are considered – isolated bunch for no PU conditions and single muon for standard PU conditions. Single muon data¹ is relatively unbiased for HCAL measurements, and the error code rate is $< 3\%$ for depths 2 through 7 (and significantly lower in isolated bunch data²). Based on this acceptably low error code rate in data and the fact that if there were significant problems with the 2018 TDC threshold, this would have been caught, we proposed to leave the TDC threshold at its default value of 18.7 for both HB and HE Run-3 data taking. However, this is a significantly lower error code rate than seen in the Run-3 simulation. In simulation, a threshold of

¹/SingleMuon/Run2018A-v1/RAW

²/IsolatedBunch/Run2018A-v1/RAW

$4 \cdot 18.7 \mu A$ is similar to what is seen in data, and therefore, thresholds from $18.7 \mu A$ to $4 \cdot 18.7 \mu A$ are used in the following studies.

An injected signal in MC is used to understand the amplitude dependence of TDC. This allows various energies to be input for an otherwise identical signal. The measured TDC value vs. the ADC values of the signals is plotted in Figure 5.2. At high energy (above $ADC = 64$, or 4 GeV), most TDC values are around $2\text{--}3 \text{ ns}$, and the amplitude dependence is low. Results are consistent across HB and HE, among depth layers, and with varying TDC threshold values, indicating that at the operating TDC threshold, amplitude dependence is a minor effect for signals above 4 GeV .




Figure 5.2: Amplitude dependence (TDC vs. ADC) for an injected signal at various TDC thresholds. The blue line indicates the ADC value where most large TDC / low amplitude hits fall below. The energy threshold for the timing trigger is 64 ADC , well above these low amplitude hits.

Applying an energy requirement of $ADC = 64$ means the ADC threshold is $\approx 3 - 4x$ larger than the TDC threshold (even up to a TDC threshold in the simulation of 56.1 or $74.8 \mu A$). This indicates a low dependence of timing values on pulse amplitude for the signals used in the trigger.

5.1.3 2018 HE TDC Data

Isolated bunch data from 2018A and 2018D is used to study the timing distributions. A few observations were noted:

- At low $i\eta$, all TDC values shifted earlier as the run progressed (comparing 2018A and 2018D)
- Higher depths had TDC values shifted earlier as the run progressed
- A peak at $\text{TDC} = 3$ is present in 2018 data and Run-3 MC and is both more prominent in data and at low $i\eta$
- In 2018 data, TDC shifts earlier as energy increases. This timeshift may be due to QIE11 gain changing
- TDC value below which 90% of prompt hits arrive differs by 1-2 ns in Run-3 MC and 2018 data (earlier in data)

This data emphasizes the importance of adjusting clock boundaries and TDC threshold to measure an accurate time. Along with the Run-3 simulation, 2018 data informs settings for Run-3. After the time alignment adjustments in Chapter 4, the pulse is contained well in SOI.




Figure 5.3: SOI-1 and SOI TDC distribution from isolated bunch 2018A data in HE, for pulses over 13k fC (left) and 7k fC (right). `/IsolatedBunch/Run2018A-v1/RAW`.

Additionally, since prompt signals have a significant tail of TDC values throughout the bunch crossing, the rate at which this prompt tail falls off compared to the LLP tail impacts our discovery potential with the dedicated LLP timing trigger. A few irreducible things contribute to the prompt tail, notably: the QIE11 phase delays cannot put the pulse precisely at the front edge of the prompt distribution because there is pulse rise time and amplitude dependence that range over 5 ns or so. Thus, the tail of delayed components from the LLP sample is vital since relying on this and prompt vetos allows for identifying LLPs.

5.2 Trigger Implementation

The timing resolution of the HCAL is relatively low at 0.5 ns, so we mainly target LLPs with a relatively large $c\tau$ (on the order of 0.5 – 10 m). With this timing resolution, a $c\tau \approx 10$ cm would be challenging to detect, but later decays are evident due to the resulting more significant time delay. Using timing information in the trigger allows for sensitivity to exotic decays of low to intermediate Higgs masses, a mass region not currently covered by pre-existing triggers. Longitudinal depth segmentation allows for sensitivity to exotic decays within the calorimeter volume, extending the trigger sensitivity to additional LLP decay kinematics.

5.2.1 LLP flag: Timing trigger

Feature bits are set on a per-trigger tower basis. To set the three timing-based feature 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 feature 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 feature 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 an LLP (LLP flag = 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])`.

5.2.2 LLP flag: Depth trigger

Significant energy must be deposited in the higher depths of the calorimeter to set the depth-based feature bit, 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 feature bit is set for HB and the trigger tower is flagged as an LLP. The depth bit is stored in `bit[0]` of the feature bits.

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.

5.2.3 OR of Timing and Depth Towers

L1 Layer-1 reduces from 6 HCAL feature bits to 1 LLP flagged bit per tower and forwards this single bit to Layer-2. 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

$$\text{Depth OR Timing} = \text{Depth OR } ((\text{Not prompt}) \text{ AND } (\text{Delay1 OR Delay2}))$$

Following the bit specifications listed above, this is:

$$\text{bit}[0] \mid (\mid \text{bit}[1] \ \& \ (\text{bit}[2] \mid \text{bit}[3]))$$

5.2.4 LLP Jet Flag

The delayed jet trigger is based on the delayed timing of pairs of energetic hits indicative of a LLP decay before or within the hadron calorimeter. A delayed jet object must have at least 2 LLP-flagged trigger towers within the 9x9 jet region, with jet energy over $E_T \geq 40$ GeV (50, 60, 70 GeV). To control rates at the Level-1 trigger, event energy over $H_T \geq 120$ GeV (160, 200, 240 GeV) is required. An illustration of the data pathway from the HCAL through L1 is in Figure 5.4.




Figure 5.4: Dataflow from the HCAL, through the HCAL uHTR, L1 Layer-1 LUT, and Layer-2 jet algorithm.

Current triggers based on H_T or multi-jet triggers impose high hadronic energy thresholds to contain the rate [21]. Many LLP events present a difficulty for triggering, as the hadronic jets are relatively low energy and the H_T is low [21]. By designing a calorimeter trigger based on unique energy deposits and time delays and lowering the H_T threshold, we can increase sensitivity to LLP models

with long lifetimes. The delayed time signature will also allow for background rejection, as the LLP events are expected to have a much more significant time delay than the rest of the bunch crossing and background events.

5.3 Firmware Specifications

The uHTR firmware specifications for the HCAL-based LLP trigger are summarized here and further detailed in [61], and the firmware is tested on an HCAL teststand [62].

5.3.1 uHTR Pathway

In uHTR firmware version 2.5.0 and onward, six extended bits are used per trigger tower to transmit depth and timing information. The first four bits carry information relating to the arrival time and depth profile of the shower, while the last two are reserved:

- bit[0] depth flag (< 1 GeV in each of depths 1 and 2, and ≥ 5 GeV E_T in depth 3+)
- bit[1] at least one prompt cell (code 00 and above energy threshold of 4 GeV E_T)
- bit[2] at least one slightly delayed cell (code 01 and above energy threshold of 4 GeV E_T)
- bit[3] at least one very delayed cell (code 10 and above energy threshold of 4 GeV E_T)
- bit[4] reserved
- bit[5] reserved

The energy thresholds for delayed cells, prompt cells, and the depth flag are set in a configurable LUT. Currently, the thresholds for all three timing-based flags (prompt, slightly delayed, very delayed) are the same at 4 GeV, though they could be changed in the future to relax the delayed category energy selection while still requiring a higher prompt energy, for example.

Potential feature bit configurations are listed below as an example.

- 000001 means the depth flag is set
- 000010 means there is a prompt cell in the tower
- 000100 means at least one slightly (01) delayed cell is found
- 001000 means at least one very (10) delayed cell is found
- 110000 means the reserved muon bits are set

Multiple bits may be set simultaneously; the example above is simplified.

LUT Bit Settings

LUTs are read by the uHTR firmware to determine the energy thresholds required to set each feature bit. The LUTs may be reconfigured to adjust the energy requirements at each depth without changing the uHTR firmware. This preserves the option of having different energy requirements across the detector. The LUT bits are set based on linearized ADC energies, which are values from 0-255 after pedestal subtraction and gain corrections are applied and have a direct correlation transverse energy, E_T . The energy LUT is a 16-bit LUT configured for each depth individually.

The depth-based bit (feature `bit[0]`) uses `bit-12` and `bit-13` in the energy LUT to encode low energy in the lower depths and high energy in the higher depths, respectively. The algorithm is:

- `bit-12` is set to 1 in:
 - depth 1, if energy is < 1 GeV
 - depth 2, if energy is < 1 GeV
 - depths 3+, always set
- `bit-13` is set to 1 in:
 - depth 1, never
 - depth 2, never
 - depths 3+, if energy is ≥ 5 GeV

Finegrain `bit[0]` will be set if LUT `bit-12` is 1 for every depth, and `bit-13` is 1 for at least one depth.

The timing-based feature `bits[1:3]` use `bit-14` and `bit-15` in the energy LUT to encode the energy threshold for prompt and delayed cells. Since two bits in the energy LUT are available, setting different energy requirements for prompt and delayed cells is an option, though the same energy threshold for both was used in Run-3.

- `bit-14` is set to 1: if above prompt energy requirement (≥ 4 GeV)
- `bit-15` is set to 1: if above delayed energy requirement (≥ 4 GeV)

Each timing-based feature bit is set if both the relevant energy LUT bit is set and the expected TDC bit is set for the same cell within the tower. This results in:

- `bit[1] = Bit-14 AND TDC=00`
- `bit[2] = Bit-15 AND TDC=01`
- `bit[3] = Bit-15 AND TDC=10`

The feature bits do not indicate the multiplicity of cells — i.e., if two cells in the tower are both above the energy threshold with TDC=01, this cannot be distinguished from a single cell in the tower being above the energy threshold with TDC=01.

5.3.2 Level-1: Layer-1 and Layer-2 Pathway

The current trigger configuration allows six feature bits to be sent from the HCAL back-end to L1 Layer-1. However, the L1 Layer-1 to Layer-2 implementation restricts this to 1 bit per trigger tower. There is an available 6:1 LUT for HB feature bits from Calo Layer-1 to Layer-2 where the depth OR timing logic is implemented. Following this, one bit per trigger tower is transmitted to Layer-2, indicating if the tower is LLP flagged, and these bits are used to define delayed/displaced jets. A jet is flagged as delayed/displaced when at least two flagged trigger towers are within it.

In firmware, this is done via setting `bit[2]` of `tow.hwQual` in the tower word with the result of the 6:1 LUT. At Layer-2, the number of flagged trigger towers in an L1 jet region is summed (using the 9x9 towers around the center of an L1 jet). When at least two flagged trigger towers are within the 9x9 jet region, `jet.HwQual` is set to flag the jet. A full emulator in CMSSW was written and validated with firmware.

5.3.3 uGT Pathway

The upgraded Global Trigger processes the per jet LLP flag to use in the relevant Level-1 trigger pathways, including integration in the Trigger Menu Editor (TME) and Trigger Menu (UTM) library. The TME defines the trigger conditions for the LLP jet bit, while the UTM library generates the firmware needed for a new L1 Trigger menu.

uGT Jet Bits

At uGT, there are five reserved and quality bits for a jet object (`bits 27-31`). L1 jet `bit-27` (a quality bit) flags a jet as delayed based on HCAL timing and depth profiles indicative of an LLP

decay. This bit is referred to as “DISP” and stores the `jet.HwQual` described above. When this bit is set to 1, the jet has been tagged as an 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 5.1: Data structure of a jet object, modified from *Scales for inputs to uGT*, Table 11 [52], and *CaloL2 input/output specifications*, Page 4 [45].

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

Trigger Menu Seeds

The trigger menu defines a new LLP delayed/displaced jet seed 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_LL` = [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_LL]`.

In the menu, the AND of the LLP jet seed with a low threshold H_T seed (denoted as HTT in the L1 menu) is taken. For example, this seed is used in conditions such as:

$$\text{L1_HTT120_SingleLLPJet40} = \text{JET40}[\text{JET-DISP_LLP}] \text{ AND } \text{L1_HTT120}$$

This seed indicates that a single jet as low as 40 GeV will trigger if the displaced bit is set and the event H_T is over 120 GeV. The trigger menu now includes a selection of seeds with tighter jet and event energy requirements, along with a dijet LLP seed.

5.4 Trigger Commissioning

Early firmware testing on the Princeton HCAL teststand was valuable for firmware commissioning and debugging. The teststand hardware consists of a Vadatech crate, a Vadatech MCH (μ TCA

Carrier Hub module), an AMC13 (Advanced Mezzanine Card module), and a uHTR (μ TCA HCAL Trigger and Readout module) [62]. The AMC13 provides clock, timing, and DAQ for the HCAL, and the firmware for the trigger is loaded onto the FPGAs of the uHTR. Significant testing was performed with this setup on an early version of the HB uHTR firmware for the LLP trigger. “Patterns”, or sequences of ADC and TDC codes, are written to the front-end RAMs, which the uHTR then reads as data. The resulting feature bit patterns are analyzed to validate the firmware pathway [62].

In early Run-3 (2022), two main issues with the trigger pathway were identified. First, a $\pm 2i\phi$ swap occurred, where the feature bit for a tower was reported offset by two towers. This was fixed with a firmware update and expected to have minimal impact on the overall trigger performance, as flagged jets use information from the 81 towers within their neighboring region, and flagged towers are generally centrally located.

The second issue was the 6:1 LUT at Calo Layer-1. This LUT mistakenly only passed the first bit in 2022, so the timing information was not used at the trigger level. Only the depth bit is available at the jet trigger level in 2022. After extensive testing and validation using Monte Carlo data injection at Point 5, the LUT was updated to the expected 6:1 and functional in 2023. This issue significantly impacts the data, as all 2022 HCAL-flagged delayed jets are depth-based (no timing). The timing-based trigger was functioning for all of 2023, with the best performance after the HCAL time re-alignment in June 2023. Specific details about the time and run numbers at which these issues were fixed are listed in Appendix A.

The MC data injection sets specific patterns (ADC and TDC values for consecutive time slices) on the uHTR front end RAM buffers. This replaces the data typically read from the QIE11s, and then is read by the HCAL DAQ and trigger primitive paths, and sent to Level-1. Thus, the mock trigger primitives are a valuable method of testing the firmware, as localized hit patterns can be designed to mimic jets and long-lived particle decays.

5.5 Performance in Run-3 Collisions

The QIE phase scan is one of the best ways to demonstrate the delayed jet performance in collisions data and validate 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. The scan is taken relative to nominal phases, so it artificially produces delayed jets 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 fraction of each feature bit code during the phase scan illustrates the trigger performance on a tower-by-tower level. Figure 5.5 demonstrates that at the optimal alignment (0 ns, as determined in the TDC distributions in Figure 4.4), most of the timing-based codes are prompt. As the delay increases, a larger fraction of the feature bit codes are from delayed hits. Near optimal alignment (from -4 ns to 10 ns), the fraction of FG bits set based on depth is approximately constant, since the phase scan does not preferentially set the depth bit. However, at very low values (-10 ns), the fraction of feature bits set based on depth increases. This is due to the decreasing fraction of valid TDC codes. As more TDC codes are set to 3 (no TDC in Figure 4.4), fewer towers can be flagged based on timing information, and therefore most of the flagged towers come from depth.





Figure 5.5: Fraction of each feature bit (or fine-grain bit, FG) code during the 2023 HCAL phase scan. FG0 is the depth bit, FG1 is the prompt timing bit, FG2 is the slightly delayed (1 ns wide bin) timing bit, and FG3 is the very delayed timing bit. This is for $i\eta = -10$ in SOI (TS2) and is representative of other HB towers.

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 Figures 5.6a and b. The per-tower and per-jet turn-on curves both strongly depend on phase offset. Consistency across HB is seen with the three regions shown in Figure 5.6 a.

Five Level-1 trigger seeds are used in Run-3 and are listed in Table 5.2. These are designed as delayed jet triggers, with increasing event energy H_T thresholds and jet energy thresholds. A double jet trigger with low energy requirements is also used. The logical “OR” of these L1s is used to seed



(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). A small fraction of towers is flagged as delayed at low phase offset values 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. An L1 jet is LLP-flagged when at least two flagged HCAL LLP towers (depth or timing-based) are 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 JetMET dataset). The delayed jet fraction begins decreasing 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 5.6: Efficiency plots vs. QIE delay in 13.6 TeV collisions with 900 bunches during the HCAL phase scan, both published in [26].

HLT delayed and displaced jet triggers, discussed in Chapter 6. Thus, if any of the five HCAL based L1 triggers is set, the following HLT paths may also be set.

Figure 5.7 shows the rates of a selection of L1 triggers during the QIE phase scan. This illustrates the substantial 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 significantly 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.

L1 Trigger Path	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 5.2: L1 trigger paths in Run-3, with rates referenced from Run 368822 after the 2023 HCAL time alignment. The TDC-based time alignment from Chapter 4 decreased LLP rates by a factor of 3, as expected. The total rate is under 1 kHz.




Figure 5.7: 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 decrease 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 assumed to be constant. Published in [26].

5.6 Performance in Simulation


Samples of the process $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ (model described in Chapter 2) were produced in a private production Monte Carlo (MC) generation. In the MC generation, masses of the Higgs-like particle and LLP are selected, along with the lifetime of the LLP. Chosen values for the samples are listed in Table 5.3.

Plots for the timing trigger efficiency in simulation vs. LLP displacement and jet energy are shown in Figure 5.8. This demonstrates that efficiencies increase with LLP displacement until the end of

Higgs Mass (GeV)	LLP Mass (GeV)	LLP $c\tau$ (m)
$m_H = 125$	$m_S = 50$	$c\tau = 3, 30$
$m_H = 250$	$m_S = 120$	$c\tau = 0.5, 1, 10$
$m_H = 350$	$m_S = 160$	$c\tau = 0.5, 1, 10$
$m_H = 1000$	$m_S = 450$	$c\tau = 10, 100$

Table 5.3: Table of Higgs mass, LLP mass, and $c\tau$ values used in the MC samples.

the HCAL volume (up to 6 m). Jet efficiencies significantly increase after 40 GeV, motivating this as the lowest threshold Level-1 trigger seed choice. The efficiencies do not heavily depend on the $H_T \geq 120$ GeV selection.



(a) Delayed jet efficiency vs. LLP displacement from PV.

(b) Delayed jet efficiency vs. jet E_T .

Figure 5.8: Efficiency plots with LLP flagged L1 jets. Four Monte Carlo samples are combined in these plots ($m_H = 125$ GeV, $c\tau = 3$ m; $m_H = 250$ GeV, $c\tau = 10$ m; $m_H = 350$ GeV, $c\tau = 10$ m; $m_H = 1000$ GeV, $c\tau = 10$ m) to show trigger performance across a representative range of LLP displacements and jet efficiencies.

Integrated Luminosity Gains

Added efficiency is defined as the events that pass the delayed jet trigger but not a L1 HTT360 (or L1 HTT420) trigger. This comparison is chosen because HTT360 and HTT420 are available, un-prescaled, L1 seeds for Run-3. Thus:

$$\text{Added efficiency} = \frac{\text{Events passing delayed jet trigger with } \text{HTT} \leq 360}{\text{total events}} \quad (5.1)$$

$$= \frac{\text{Events passing delayed jet trigger OR HTT360} - \text{HTT360 events}}{\text{total events}} \quad (5.2)$$




Figure 5.9: Delayed jet efficiency vs. expected time delay, defined by $\Delta t = l_{\text{LLP}}/v_{\text{LLP}} + l_b/c - t_{\text{expected}}$. Efficiency plateaus after a TOF delay of ≈ 4 ns, and prompt LLP decays are suppressed due to no time delay present. At larger time delays, the efficiency decreases due to decreased event energy and the probability of prompt cells vetoing the towers due to PU.

The effective luminosity gain indicates how many times faster a discovery could happen with the timing trigger and is defined as:

$$\text{Effective luminosity gain} = \frac{\text{new efficiency}}{\text{old efficiency}} \quad (5.3)$$

$$= \frac{\text{added efficiency}}{\text{current } (H_T \text{ 360}) \text{ efficiency}} + 1 \quad (5.4)$$

$$= \frac{\text{Events passing delayed jet trigger OR } H_T \text{ 360}}{H_T \text{ 360 events}} \quad (5.5)$$

Results are tabulated in Table 5.5, showing a large integrated luminosity (2-5x) for the low mass sample in this final state.

5.6.1 Efficiency and Rate Tables

Efficiency tables report which sample is used, the absolute LLP jet efficiency, the absolute efficiency for L1_HTT360 and L1_HTT420, the LLP jet added efficiency compared to L1_HTT360 and L1_HTT420, the integrated luminosity as compared to L1_HTT360 and L1_HTT420, and the neutrino gun rates. All are evaluated on a private MC sample: 2k events for LLPs and 3.6k events for the neutrino gun.

By comparing the integrated luminosity and added efficiency results in Table 5.5, integrated lu-

minosity gains decrease more from increasing the L1 H_T from 120 to 160 or 200 GeV, than from increasing the jet energy from E_T 40 GeV to 60 GeV. Additionally, increasing jet energy is a powerful way to reduce the neutrino gun rates. On this limited sample, by jet energy of 60 GeV, the rates drop significantly.

The general trend we see is an increase in the signal efficiencies (focusing on $m_H = 125$ performance) by +20% for a 4x increase in the rate. As we want to remain under 1 kHz, these trade-offs are evaluated and in Table 5.4 a full span of rates is estimated, to aid in the selection of a seed without a prescale. The efficiencies shown in Figure 5.4 are the integrated luminosity gains, showing a range depending on a comparison between H_T 360 and H_T 420 GeV.

Combining the estimated rates and efficiencies from Table 5.5, we have a summary in Table 5.4, motivating the choice of:

- L1_HTT120_SingleLLPJet40 = JET40[JET-DISP_LL] AND L1_HTT120 with a prescale of 100
- L1_HTT160_SingleLLPJet50 = JET50[JET-DISP_LL] AND L1_HTT160 with a prescale of 10
- L1_HTT200_SingleLLPJet60 = JET60[JET-DISP_LL] AND L1_HTT200 as an un-prescaled seed
- L1_HTT240_SingleLLPJet70 as an un-prescaled seed
- L1_2SingleLLPJet40 as an un-prescaled seed (requires 2 LLP flagged jets and no event energy requirement)

Requirement (GeV)	HTT 120	HTT 160	HTT 200	HTT 240	HTT 0
Jet $E_T \geq 40$					
Rates (kHz)	23 kHz				
Efficiency factor	3.7-4.9				
Jet $E_T \geq 50$					
Rates (kHz)		6			
Efficiency factor		2.7-3.5			
Jet $E_T \geq 60$					
Rates (kHz)			1.4		
Efficiency factor			1.9-2.4		
Jet $E_T \geq 70$					
Rates (kHz)				0.6	
Efficiency factor				1.4-1.7	
2 LLP Jet $E_T \geq 40$					
Rates (kHz)					0.4
Efficiency factor					1.4-1.5

Table 5.4: Estimated L1 neutrino gun rates (kHz, pure rate) and signal efficiency gains ($m_H = 125$ GeV integrated luminosity, a range for comparing to un-prescaled HTT 360 or HTT 420) vs HTT and jet ET requirements.

Sample	Integrated luminosity gains (L1) compared to HTT 360	Integrated luminosity gains (L1) compared to HTT 420
$m_h = 125 \text{ GeV}, c\tau = 3 \text{ m}$	3.7	4.9
$m_h = 125 \text{ GeV}, c\tau = 30 \text{ m}$	2.5	3.5
$m_h = 250 \text{ GeV}, c\tau = 0.5 \text{ m}$	1.4	1.5
$m_h = 250 \text{ GeV}, c\tau = 1 \text{ m}$	1.7	2.2
$m_h = 250 \text{ GeV}, c\tau = 10 \text{ m}$	2.1	2.5
$m_h = 350 \text{ GeV}, c\tau = 0.5 \text{ m}$	1.3	1.5
$m_h = 350 \text{ GeV}, c\tau = 1 \text{ m}$	1.6	1.9
$m_h = 350 \text{ GeV}, c\tau = 10 \text{ m}$	2.1	2.6
$m_h = 1000 \text{ GeV}, c\tau = 10 \text{ m}$	1.3	1.4
$m_h = 1000 \text{ GeV}, c\tau = 100 \text{ m}$	1.2	1.3

Table 5.5: Level-1 trigger signal efficiencies and integrated luminosity gain, comparing $H_T = 120 \text{ GeV}$, jet energy = 40 GeV, with no triggerability requirements (2k events each), to $H_T = 360 \text{ GeV}$ or 420 GeV. Private MC L1TNtuple production, using the hwQual bit from HCAL to flag delayed or displaced jets.

Chapter 6

HCAL-based High Level LLP Triggers

The Level-1 LLP trigger is expected to provide a good signature at the high level trigger (HLT). After L1, the HLT can reconstruct displaced tracks and further refine the saved LLP candidate events. This reduces the bandwidth of data that must be passed through HLT and offline. At HLT with path reconstruction, SM backgrounds can be further reduced, as signals from Λ and K-short can be eliminated. Some other important backgrounds are prompt tails from $\pi^\pm$, beamspot spread, and late hadronic showers, as these all have signatures similar to LLPs. At the core, the L1 LLP trigger is a tool so HLT can identify displaced vertices or more detailed calorimeter signatures.

Displaced and delayed jet triggers using the HCAL significantly improve acceptance of a wide variety of signal models involving displaced and delayed jets. By lowering the HT thresholds at HLT, we are keeping events for offline analysis that were found to be important in various Run-2 LLP analyses when HLT HT was lower. These events have been cleaned significantly with tracking at HLT. The range 200 – 430 GeV H_T events are removed from the existing LLP HLTs due to the L1 H_T increase, and we can keep the HCAL-based L1 (and HLT H_T) thresholds low in Run-3 because of the very restrictive and selective L1 LLP seed.

6.1 HLT Paths

The HLT paths implemented are based on the successful LLP HLTs used in Run-2. Since these are seeded by the logical or of the five new L1 HCAL LLP seeds, the LLP coverage is expanded to lower masses and lower energy events. At HLT, additional tracking and energy requirements are required. There are 15 new HLT paths implemented, in four major groups.

- HLT_HT200_L1SingleLLPJet_DisplacedDijet35_Inclusive1PtrkShortSig5
- HLT_HT200_L1SingleLLPJet_DisplacedDijet40_Inclusive1PtrkShortSig5
- HLT_HT240_L1SingleLLPJet_DisplacedDijet40_Inclusive1PtrkShortSig5
- HLT_HT280_L1SingleLLPJet_DisplacedDijet40_Inclusive1PtrkShortSig5
- HLT_HT170_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack
- HLT_HT200_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack
- HLT_HT270_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack
- HLT_HT200_L1SingleLLPJet_DisplacedDijet60_DisplacedTrack
- HLT_HT320_L1SingleLLPJet_DisplacedDijet60_Inclusive
- HLT_HT420_L1SingleLLPJet_DisplacedDijet60_Inclusive
- HLT_HT200_L1SingleLLPJet_DelayedJet40_DoubleDelay0p5nsTrackless
- HLT_HT200_L1SingleLLPJet_DelayedJet40_DoubleDelay1nsInclusive
- HLT_HT200_L1SingleLLPJet_DelayedJet40_SingleDelay1nsTrackless
- HLT_HT200_L1SingleLLPJet_DelayedJet40_SingleDelay2nsInclusive
- HLT_L1SingleLLPJet

The first four triggers (`*_DisplacedDijet*_Inclusive1PtrkShortSig5`) require a displaced dijet, which is two calo jets above the energy threshold listed in the HLT path, each having at most one associated prompt track (impact parameter < 1.0 mm). The second group of four (`*_DisplacedDijet*_DisplacedTrack`) requires a displaced dijet with one displaced track. The displaced track requirement means that each of the two jets must have one associated track with impact parameter > 0.5 mm. The next two paths (`*_DisplacedDijet60_Inclusive`) have a displaced dijet, but allow for two associated prompt tracks (thus requiring the higher jet and HTT energy thresh-

olds). The next four (`*_DelayedJet40_DoubleDelay*ns*` and `*_DelayedJet40_SingleDelay*ns*`) require one or two jets with delayed ECAL jet timing. The final HLT `HLT_L1SingleLLPJet` is for monitoring. This highly prescaled path is a pass-through of the HCAL-based L1 LLP triggers and is used to evaluate the efficiency of the L1 jet triggers at HLT.

6.1.1 Added Acceptance

Compared to the previously existing HLT paths, the H_T requirements are significantly lowered for the paths seeded from the L1 LLP triggers. These are lowered by 230 GeV for each path, and lead to improvements in the integrated luminosity gains, up to a 3x increase for LLP efficiency. This is shown in Table 6.1. Table D.4 and D.5 show the integrated luminosity gains split by path, and separated by performance on each MC sample.

Sample	Integrated luminosity gains (HLT)
$m_h = 125 \text{ GeV}, c\tau = 3000 \text{ mm}$	2.8
$m_h = 125 \text{ GeV}, c\tau = 30000 \text{ mm}$	≥ 4
$m_h = 250 \text{ GeV}, c\tau = 500 \text{ mm}$	1.3
$m_h = 250 \text{ GeV}, c\tau = 1000 \text{ mm}$	1.7
$m_h = 250 \text{ GeV}, c\tau = 10000 \text{ mm}$	3.4
$m_h = 350 \text{ GeV}, c\tau = 500 \text{ mm}$	1.4
$m_h = 350 \text{ GeV}, c\tau = 1000 \text{ mm}$	1.8
$m_h = 350 \text{ GeV}, c\tau = 10000 \text{ mm}$	2.6
$m_h = 1000 \text{ GeV}, c\tau = 10000 \text{ mm}$	1.4
$m_h = 1000 \text{ GeV}, c\tau = 100000 \text{ mm}$	1.7

Table 6.1: HLT performance, using v3 config. Integrated luminosity is reported as (LLP L1 OR H_T L1 seed)/ H_T L1 seed to show the factor in physics acceptance added by the inclusion of the new HLT single LLP jet triggers.

6.1.2 HLT Rates

The time alignment of HCAL greatly impacts the HCAL-seeded LLP HLT rates, as this impacts the L1 seeds. The HCAL time re-alignment was expected to reduce HCAL-based LLP HLT rates by a factor of 3-4, as was seen from comparing runs before and after the time alignment. After HCAL re-alignment, the displaced and delayed jet triggers using the HCAL information have a rate in 2023 Run-3 of 30 Hz at an instantaneous luminosity of $2 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, as shown in Figure 6.2.

In addition, the time alignment minimally impacted other triggers (jet-based, H_T -based, etc.), as the compensating energy response corrections balance this. Standard triggers are minimally affected and have a rate reduction factor of 1 - 1.08. Low p_T jet triggers are least affected, while high p_T jet triggers were reduced by 4–8%. However, it should be noted that L1 deployed jet energy corrections

(JEC) during the same run as the new time alignment was deployed, complicating the efforts to understand the impact of HCAL time alignment and energy response corrections independently. L1 JEC are expected to increase jet rates by $\approx 10\%$.

HLT Path	Rate Before Alignment (Hz, Run 368753)	Rate After Alignment (Hz, Run 368822)	Rate Reduction Factor
HLT_HT200_L1SingleLLPJet_DisplacedDijet40_Inclusive1PtrkShortSig5_v	59.26	19.31	3.07
HLT_HT240_L1SingleLLPJet_DisplacedDijet40_Inclusive1PtrkShortSig5_v	45.77	14.08	3.25
HLT_HT280_L1SingleLLPJet_DisplacedDijet40_Inclusive1PtrkShortSig5_v	31.76	9.52	3.33
HLT_HT170_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack_v	13.31	4.34	3.06
HLT_HT200_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack_v	12.58	2.97	4.24
HLT_HT200_L1SingleLLPJet_DisplacedDijet60_DisplacedTrack_v	5.01	1.25	4.02
HLT_HT270_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack_v	7.14	1.44	4.97
HLT_HT320_L1SingleLLPJet_DisplacedDijet60_Inclusive_v	11.41	3.31	3.45
HLT_HT420_L1SingleLLPJet_DisplacedDijet60_Inclusive_v	5.33	1.23	4.33
HLT_HT200_L1SingleLLPJet_DelayedJet40_SingleDelay1nsTrackless_v	12.43	4.63	2.68
HLT_HT200_L1SingleLLPJet_DelayedJet40_SingleDelay2nsInclusive_v	1.38	0.14	9.58
HLT_HT200_L1SingleLLPJet_DelayedJet40_DoubleDelay0p5nsTrackless_v	4.99	2.11	2.37
HLT_HT200_L1SingleLLPJet_DelayedJet40_DoubleDelay1nsInclusive_v	[Low Stats]	[Low Stats]	[Low Stats]

Table 6.2: HLT path rates before and after the HCAL alignment, at a pileup of 54 and 2400 colliding bunches. Rates are calculated using the RateMon tool.

6.2 HLT Efficiencies


The efficiencies are evaluated on a MC LLP sample such that the HLT efficiency vs. truth variables can be evaluated. HLT efficiencies are plotted for all events passing the baseline selection criterion (outlined below), monitoring HLT, and HCAL LLP-based HLTs. The monitoring HLT is a “pass-through” HLT, where any event passing the HCAL LLP-based L1 triggers satisfies this monitoring path. Figures 6.1 through 6.8a show the fraction of events selected by the L1 and HLT paths, while Figures 6.1 through 6.8b show the fraction of events selected by each category of HCAL LLP-based HLTs. The four groups of HLTs are described in Section 6.1 and are displaced dijet with ≤ 1 prompt track, displaced dijet with ≤ 2 prompt tracks, displaced dijet with at least one prompt track, and ECAL timing-based delayed jet. The efficiencies are generally similar for each HLT, with the highest contribution from displaced dijet paths allowing for prompt tracks (which also have the highest rate).

For the HLT efficiency plots vs. LLP truth variables (decay r , decay z , η , or travel time, shown in Figures 6.1, 6.2, 6.3, and 6.4), the requirements are:

- One of the three leading jets in the event is matched to the LLP of interest
- The matched jet $p_T > 100$ GeV
- Event $H_T > 170$ GeV (informed by lowest threshold HLT path)

The choice of jet $p_T > 100$ GeV and event $H_T > 170$ GeV are motivated by a hard jet cut and the

lowest HLT HT cut. Cuts of jet $p_T > 60$ GeV and event $H_T \geq 200$ GeV, motivated by the L1 trigger thresholds, were also studied, showing similar distributions. When studied on a $m_H = 125$ GeV sample, the statistics were improved with the lower jet cut due to the softer jet distribution present in this sample. The plots shown in Figures 6.1 through 6.8 are on a $m_H = 350$, $m_S = 80$ GeV, $c\tau = 0.5$ m sample.



(a) HLT efficiencies for all events, HCAL LLP L1 triggers, and HCAL LLP HLTs passed.


(b) HLT efficiencies for HCAL LLP L1 triggers passed, HCAL LLP HLTs passed, and then split by four categories of HLTs.

Figure 6.1: HLT efficiencies vs. LLP decay radius, with peak efficiency for LLPs decaying in HCAL barrel depths 2-4 (214 – 295 cm). This LLP r region providing best efficiency is used as a selection cut for the HLT efficiency plots vs. jet and event kinematics.


Figure 6.3 demonstrates that HLT efficiency peaks between $|\eta_{LLP}| < 1.26$. This is used as a cut in the LLP-jet truth matching and is further physically motivated by the HCAL tower layout, as $|\eta| < 1.26$ is the center of tower $i\eta = 15$, the last full tower in HB. In Figure 6.3, some events with $|\eta_{LLP}| > 1.26$ remain, as a jet can be matched to LLP decay products (if LLP does not decay in HCAL barrel).

For the HLT efficiency plots vs. jet and event kinematics (leading jet E , p_T , η , event H_T , shown in Figures 6.5, 6.6, 6.7, and 6.8), the requirements are:

- LLP is matched to the leading jet
- LLP truth matching requires $|\eta_{LLP}| \leq 1.26$ (middle of HB $i\eta = 15$)
- LLP decay r is in HCAL depth 2-4 (informed by best efficiency region in HLT efficiency study)




(a) HLT efficiencies for all events, HCAL LLP L1 triggers, and HCAL LLP HLTs passed.




(b) HLT efficiencies for HCAL LLP L1 triggers passed, HCAL LLP HLTs passed, and then split by four categories of HLTs.

Figure 6.2: HLT efficiencies vs. LLP decay z , showing that triggers gain efficiency as z increases.




(a) HLT efficiencies for all events, HCAL LLP L1 triggers, and HCAL LLP HLTs passed.




(b) HLT efficiencies for HCAL LLP L1 triggers passed, HCAL LLP HLTs passed, and then split by four categories of HLTs.

Figure 6.3: HLT efficiencies vs. LLP η , showing flat performance for $|\eta| < 1.26$.




(a) HLT efficiencies for all events, HCAL LLP L1 triggers, and HCAL LLP HLTs passed.




(b) HLT efficiencies for HCAL LLP L1 triggers passed, HCAL LLP HLTs passed, and then split by four categories of HLTs.

Figure 6.4: HLT efficiencies vs. LLP travel time, with highest efficiencies over 6 ns, as expected based on the prompt-delayed boundary encoded in the TDC LUT.

vs. LLP truth variables)




(a) HLT efficiencies for all events, HCAL LLP L1 triggers, and HCAL LLP HLTs passed.




(b) HLT efficiencies for HCAL LLP L1 triggers passed, HCAL LLP HLTs passed, and then split by four categories of HLTs.

Figure 6.5: HLT efficiencies vs. event H_T , showing a plateau over 160 GeV (L1) and 280 GeV (HLT).




(a) HLT efficiencies for all events, HCAL LLP L1 triggers, and HCAL LLP HLTs passed.




(b) HLT efficiencies for HCAL LLP L1 triggers passed, HCAL LLP HLTs passed, and then split by four categories of HLTs.

Figure 6.6: HLT efficiencies vs. jet energy, showing a plateau above 80 GeV.




(a) HLT efficiencies for all events, HCAL LLP L1 triggers, and HCAL LLP HLTs passed.




(b) HLT efficiencies for HCAL LLP L1 triggers passed, HCAL LLP HLTs passed, and then split by four categories of HLTs.

Figure 6.7: HLT efficiencies vs. jet p_T , showing a plateau above 70 GeV, except for the displaced dijet path allowing for 2 prompt tracks. This trigger plateaus higher due to the less restrictive track requirement.



(a) HLT efficiencies for all events, HCAL LLP L1 triggers, and HCAL LLP HLTs passed.



(b) HLT efficiencies for HCAL LLP L1 triggers passed, HCAL LLP HLTs passed, and then split by four categories of HLTs.

Figure 6.8: HLT efficiencies vs. jet η , showing high efficiency for $|\eta| < 1.26$.

Overall, high HLT efficiency is seen for jets in the barrel region, $p_T \geq 60$ GeV, and H_T of 200 GeV. Evaluating the efficiency of the HLTs with respect to LLP truth variables shows high efficiency for LLPs that decay in the barrel region, particularly HCAL depth 3 through 4 ($r = 214\text{--}295$ cm).

6.2.1 Improvements for 2024: CalRatio Trigger

At the offline analysis level (Section 8.3.2), jet neutral hadron energy fraction was determined to be a variable providing significant discrimination power between LLP jets and background jets. Building off of this, a “CalRatio” HLT is implemented for 2024 data-taking that uses the HCAL-based L1 single jet LLP triggers to seed a HLT that identifies jets with a high neutral hadron energy fraction. A similar CalRatio trigger has been used in ATLAS to search for LLP decaying in the calorimeter with trackless jets and a low energy fraction in the ECAL. In contrast to the CMS neutral hadron fraction approach, ATLAS uses the ratio of ECAL and HCAL energies. Using the neutral hadron energy fraction lets the new trigger benefit from the particle flow algorithm to remove HCAL energy from charged (and thus track-matched) particles.

6.3 HLT Performance

Sensitivity to lower energy events is enabled with the displaced jet HLT paths seeded by the dedicated HCAL-based L1 LLP jet triggers. Final state sensitivity of $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ is increased for a range of mass and lifetime parameters, particularly providing improvements at the low to mid ($m_H = 125 - 350$ GeV) mass range due to the benefit of lowering the event H_T threshold. This is the main model used for benchmarking studies; however, sensitivity to other models producing delayed or displaced hadronic jets is expected. The limitations of the currently implemented HLTs are that additional detailed calorimeter information is not utilized beyond the L1 seed, however, the CalRatio triggers in place starting in 2024 address this gap. Additionally, the L1 trigger performance impacts the HLT capabilities, and the improvements in time alignment and optimization of the prompt timing window discussed in Chapters 4 and 5 lead to enhanced HLT performance as well.

Chapter 7

Datasets and Monte Carlo Samples

7.1 Run-3 Datasets

The search for long-lived particles decaying in the HCAL (CMS AN-2023/174, [60]) uses data collected with the CMS detector from proton-proton collisions at $\sqrt{s} = 13.6$ TeV in 2022 and 2023 during Run-3 of the LHC. The integrated luminosity delivered by the LHC and recorded by CMS is shown in Figure 7.1, with 41.5 fb^{-1} collected in 2022 and 31.9 fb^{-1} collected in 2023 [33]. The delivered luminosity in Figure 7.1 includes all luminosity from when stable beams are reached until the LHC dumps the beam, and CMS recorded luminosity is slightly lower than delivered luminosity due to DAQ inefficiency or other detector impacts [33].

A slightly lower fraction of the luminosity recorded by CMS is “certified” or considered good for use in physics analysis. Data is certified when all subdetectors, triggers, and physics objects perform as expected. The certified luminosity sections are saved in a “Golden JSON”, and these data are used for analysis. This comprises 34.9 fb^{-1} in 2022 and 27.8 fb^{-1} in 2023.

Data collected with the Run-3 displaced and delayed jet HLTs seeded with the HCAL-based L1 LLP jet triggers are analyzed. The triggers are detailed in Chapters 5 and 6, for the L1 and HLTs, respectively.




Figure 7.1: Cumulative luminosity delivered by the LHC to CMS (blue) and recorded by CMS (orange) for Run-3 proton-proton collisions at $\sqrt{s} = 13.6$ TeV. The startup runs at the injection energy of 450 GeV are not included here. Figure from [33].

7.1.1 Skim for LLP Analysis

The HCAL displaced jet dataset from 2023 is saved at AOD (analysis object dataset) level. However, the HB reconstructed hits are needed for this analysis, though they are not saved in standard AOD. This low-level information is necessary, as the LLP analysis variables require access to the per-channel HCAL timing and energy information. The solution is to create a “skim” dataset, specifying that only events passing the HCAL-based LLP triggers will be saved, while also recording the HBHE reconstructed hits (rechits) for those events.

Thus, a dedicated skim is used for the Run-3 LLP analysis. This allows a specific data collection to be saved for offline analysis when a given set of triggers are passed. The skim is called “EXOLLP-JetHCAL”, and the datatier is “AOD + RAW + HBHERechits”. When the HCAL LLP seeded HLTs are passed, all the AOD (analysis object data), RAW (low-level hit information), and HBHE rechits from RECO (essentially a copy of HBHE digitized hits) are saved. The skim is under the DisplacedJet dataset, identified by:

```
dataset=/DisplacedJet/*Run2023*EXOLLPJethCAL-PromptReco-v*/AOD
```

The skim was in place starting in 2023 and collected about 80 million events from 15 HLTs seeded by 5 L1 paths. This skim contains about 30% of the events in the DisplacedJet dataset, which includes many additional tracking, ECAL, and muon system-based triggers not included in the HCAL specific

skim.

7.2 LLP Monte Carlo Samples

Monte Carlo simulation is performed using a set of tools developed by the particle physics community to generate events from proton-proton collisions, and then describe their evolution. MADGRAPH5_AMC@NLO is a framework that generates hard scattering events, computes cross sections, and provides tools to simulate both SM and BSM phenomenology [5]. Next-to-leading-order (NLO) accuracy is achieved for QCD corrections to the SM process, while leading-order (LO) is simulated for any input Lagrangian. After the initial hard scattering event, the propagation of particles is simulated using PYTHIA, which models the cross sections, parton distribution functions, parton showering, hadronization, and decays for an event [8]. PYTHIA has been used at most colliders in the last four decades, including LEP (e^+e^-), the Tevatron (pp), and the LHC (pp and lead-ion collisions).

Once collision and subsequent particle propagation are simulated, the next step is to simulate the interaction of the particles with the CMS detector. GEANT4 is a toolkit that describes how particles interact in various materials (such as the silicon, brass, and scintillators of CMS) through electromagnetic, strong, and weak interactions such as bremsstrahlung, photon conversion, and showering [1]. DELPHES is another framework that simulates the detector response and outputs observables (leptons, jets, MET) useful at the analysis level [36]. GEANT4 performs “full simulation” to perform detailed studies while fully accounting for detector effects, whereas DELPHES performs a “parametric simulation” geared toward phenomenological studies with a given detector.

7.2.1 Masses and Lifetimes

The model used in the analysis development is a Higgs-like particle decaying to a pair of long-lived particles, each decaying to $b\bar{b}$, as described in Section 2.3. A range of exotic Higgs masses, LLP masses, and LLP $c\tau$ s are used in the analysis to scan the parameter space and determine where sensitivity is the highest. LLP masses are set to be between 0.12 and 0.48 of the Higgs mass, allowing for an on-shell decay and a moderately boosted LLP (at lower LLP to Higgs mass ratios) or slow moving LLP (at higher mass ratios). Table 7.1 lists the parameters used for the available samples and the expected full signal grid.

Higgs Mass (GeV)	LLP Mass (GeV)	LLP $c\tau$ (m)	Number of Events Generated
$m_H = 125$	$m_{LLP} = 15$	$c\tau = 1$	2 million
$m_H = 125$	$m_{LLP} = 50$	$c\tau = 3$	4 million
$m_H = 250$	$m_{LLP} = 120$	$c\tau = 10$	4 million
$m_H = 350$	$m_{LLP} = 80$	$c\tau = 0.5$	0.5 million
$m_H = 350$	$m_{LLP} = 160$	$c\tau = 10$	4 million
$m_H = 125$	$m_{LLP} = 25, 50$	$c\tau = 10$	-
$m_H = 200$	$m_{LLP} = 45, 90$	$c\tau = 10$	-
$m_H = 250$	$m_{LLP} = 60, 120$	$c\tau = 10$	-
$m_H = 350$	$m_{LLP} = 80, 160$	$c\tau = 10$	-
$m_H = 500$	$m_{LLP} = 110, 220$	$c\tau = 10$	-
$m_H = 750$	$m_{LLP} = 180, 360$	$c\tau = 10$	-
$m_H = 1000$	$m_{LLP} = 225, 450$	$c\tau = 10$	-

Table 7.1: Table of Higgs mass, LLP mass, and $c\tau$ values used in the MC samples for the analysis. The full analysis will use various samples at a wide variety of m_H, m_S , and $c\tau$ values. A range of points we expect sensitivity to are shown in the lower portion of the table, and a selection of these has been generated, highlighted in blue.

7.2.2 Production of Samples

In the Monte Carlo production, events are saved at RECO level. This is necessary to access the HCAL reconstructed hit collection needed for the analysis. Thus, the data tiers for the analysis are determined: dedicated HCAL LLP skim in 2023 collisions and RECO (reconstructed detector hits) in MC. This means that large event sizes are necessary to work with; however, it is beneficial in developing new calorimeter-based variables.

7.2.3 LLP Lifetime Re-weighting

Long-lived particle Monte Carlo samples are generated at a few specific lifetime points, typically 0.5, 1, 3, 10 m. To interpolate between these lifetime points, event re-weighting is used. Intermediate samples are generated by event re-weighting according to the formula:

$$w_{\tau_{\text{old}} \rightarrow \tau_{\text{new}}} = \frac{\tau_{\text{old}}}{\tau_{\text{new}}} \cdot e^{-t_{\text{event}} \left(\frac{1}{\tau_{\text{new}}} - \frac{1}{\tau_{\text{old}}} \right)} \quad (7.1)$$

where t_{event} is the decay time for the LLP in the initial event, given by

$$t_{\text{event}} = \frac{\text{LLP travel distance (lab frame)}}{\gamma \cdot \beta} \quad (7.2)$$

For a sample with two long-lived particles, the re-weighting equation becomes:

$$w_{\tau_{\text{old}} \rightarrow \tau_{\text{new}}} = \left(\frac{\tau_{\text{old}}}{\tau_{\text{new}}} \right)^2 \cdot e^{-(t_{\text{LLP } 1} + t_{\text{LLP } 2}) \cdot \left(\frac{1}{\tau_{\text{new}}} - \frac{1}{\tau_{\text{old}}} \right)} \quad (7.3)$$

Effectively, re-weighting works by taking a given $c\tau$ for the LLP in an individual event and scaling it higher (to produce a sample with a larger $c\tau$) or lower (to produce a sample with a smaller $c\tau$). Thus, it is more effective to re-weight a high $c\tau$ sample to a low $c\tau$, as there are higher statistics at each decay point, and the statistical uncertainty of the reweighted sample is minimized. For the re-weighted samples used in the analysis, a closure test will be performed by re-weighting two samples to the same lifetime (such as re-weighting samples with $c\tau = 10$ m and 3 m both to $c\tau = 1$ m) and confirming that the distributions are the same between the two. The re-weighting procedure allows for coverage of a wide range of $c\tau$ values without producing individual Monte Carlo at each point, saving significant computational power in the simulation step.

Chapter 8

Analysis Strategy

In a LLP search, a central aim is to demonstrate sensitivity to rare decays and have a much higher sensitivity than that imposed by coupling constraints. For the LLP models of interest, these rare decays may be a small fraction of the Higgs decays (branching ratio $\leq 1\%$), and the integrated luminosity, trigger efficiencies, and analysis efficiencies further limit the data we access. Nonetheless, in this thesis we demonstrate that we have sensitivity to potential LLP decays, which have been improved by implementing a dedicated LLP trigger. Relying on the very distinct signatures of the LLP events identifies them from Standard Model backgrounds (important for the event selection and background rejection). The sensitivity to a small branching fraction is present in many other searches, such as $H \rightarrow \gamma\gamma$ with a BR of 0.2% and $H \rightarrow ZZ \rightarrow 4l$ with a BR of 0.1-3%, both of which are greatly helped by the very distinct final state signature.

Utilizing the dedicated HCAL timing and depth-based long-lived particle triggers, a search for long-lived particles decaying before or in the HCAL is performed using the data collected by CMS in 2022 and 2023. The analysis strives to identify signatures and characteristics of LLP events through variables characterizing the jet shape, spread, and lateral distribution and is interpreted in an exotic Higgs model. In this model, a Higgs-like particle decays into two long-lived particles, each decaying to $b\bar{b}$ ($H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$), as described in Chapter 2. Chapter 7 details the datasets and Monte Carlo samples, and the full list of MC samples used is in Appendix C.

8.1 Object Selection

We start by validating the objects used. As the Higgs portal model produces hadronic jets, we use standard anti- k_T (AK4) jets, as described in Section 3.4.1. The use of particle flow jets is validated, and then the reconstruction efficiency of LLPs is evaluated, considering the effects of the LLP decay location.

8.1.1 Particle Flow Jets

CMS has multiple methods to reconstruct jet objects, calorimeter-based approach (“calo jets”), jet + track approach, and particle flow approach (“PF jets”). PF jets are advantageous as the particle reconstruction typically leads to improved energy resolution, and are typically the preferred jet type for analysis. As LLPs can produce non-standard jets, the performance of calo jets and PF jets is compared to validate the choice of jet object. The reconstruction efficiency of particle flow jets and calorimeter jets is evaluated, determining that the reconstruction efficiency of PF jets is very high, and thus these are used in the analysis.

Events with the LLP decaying within the HB volume, $LLP |\eta| \leq 1.5$, and LLP decay matched to a generator level jet are evaluated. The ΔR for each generator and reconstructed jet pair (split by PF AK4 jets and calo AK4 jets) peaks with $\Delta R < 0.1$, indicating a high LLP reconstruction efficiency. The reconstruction efficiency for events with at least one gen-reco jet pair with $\Delta R < 0.4$ is 85.9% (PF jets) and 85.8% (calo jets). When requiring the reconstructed jets have $p_T > 30$ GeV, the efficiencies become 63% (PF jets) and 59% (calo jets) [57]. Signal reconstruction efficiency is improved at higher jet p_T values, for both calo and PF jets.

Though both calo and PF jets have similar reconstruction efficiency, the PF jet reconstruction is cleaner than the calo jet reconstruction, as $\Delta R(\text{gen,calo})$ has an extended tail with a significant fraction of events above $\Delta R > 1$. Thus, AK4 PF jets are used for this analysis.

8.1.2 LLP-Jet Matching Efficiencies

The efficiency of long-lived particle matching to a reconstructed jet is evaluated in simulation. This is necessary to determine if relying on standard jet objects is valid for this analysis. LLPs often have unique decay topologies (displaced vertices, kinked path length, and non-standard incident angles), making it very important to validate the assumptions and use of standard objects.

To characterize the efficiency of matching a LLP decay product to a PF AK4 jet, the fraction of LLPs

matched is evaluated vs. various LLP truth quantities. The following procedure is followed:

- If the LLP decays in the HB volume (decay radius $177 \leq r < 295$ cm and $|\eta_{LLP}| \leq 1.26$), it can be directly matched to a jet. If $\Delta R(\text{jet}, \text{LLP}) \leq 0.4$, then the LLP is matched to the jet.
- If the LLP decays before the HCAL volume, then the LLP decay products are used for matching instead. For each decay product (two b quarks in the case of $H \rightarrow ss \rightarrow 4b$), the ΔR between the jet direction and the LLP decay direction is evaluated. This is done by taking the jet direction (p_T, η, ϕ) , correcting for the LLP direction (jet vector - LLP vector), and comparing the resulting η, ϕ from the new vector to the η, ϕ of the LLP decay product.
- *Note:* In either case, the jet energy is cut on, typically at 40 GeV. The initial matching does not impose this requirement, so the efficiencies can be evaluated at various jet energies.

Figure 8.1 requires leading jet $p_T \geq 40$ GeV, $\eta_{jet} \leq 1.26$, and that this jet is matched to the leading LLP. For jets matched to LLPs that decay in the tracker, the fractional p_T is slightly less than 1. This means that not all of the LLP p_T was captured in the jet. This is due to these LLPs decaying into two b quarks relatively early on, and the two b quarks are not necessarily found in the same jet. When the LLP decays in the calorimeters, the decay products are contained within one jet (a more columnated shower), and the jet p_T matches the LLP p_T with more accuracy. Sometimes, the jet p_T is larger than LLP p_T , as other particles (from pile-up, other backgrounds) may also be picked up along with the LLP-associated shower. This indicates that a single jet captures the decay when an LLP decays in the calorimeter.





Figure 8.1: For leading jets matched to the leading LLP, the ratio of $\frac{\text{jet } p_T}{\text{LLP } p_T}$ is plotted. The peak at 1 means that the jet matches the LLP p_T well, as expected.

The jet matching efficiency to LLPs is evaluated vs. jet η to understand positional impacts and vs. LLP decay radius r to understand sensitivity regions. Matching requires LLP $p_T \geq 40$ GeV, LLP $\eta \leq 1.26$, and $\Delta R < 0.4$ to an AK4 jet with p_T over the specified value. In Figure 8.2, the jet momentum requirement is scanned. The lower ratio plots show the matching efficiency, with the denominator as all LLPs within the acceptance region.



(a) Fraction of LLPs matched to a jet, by LLP decay radius.

(b) η of jets matched to LLP.

Figure 8.2: Matching efficiencies of LLPs to jets, split by jet position and LLP decay r . The jet p_T is scanned, showing the efficiency drop as higher jet energies are required. The lowest un-prescaled Level-1 HCAL LLP trigger requires a jet energy of 60 GeV. Plots are for $m_H = 125$ GeV.

In Figure 8.2, a matching efficiency of about 90% for jets over 40 GeV is achieved for LLPs decaying in the calorimeter (177 - 295 cm). The matching efficiency is approximately flat by jet η , as shown in Figure 8.2b. There is a drop of efficiency above $|\eta| = 1.26$, as expected due to the position of required LLP decay. A cut on jet $|\eta| \leq 1.26$ is also imposed in the analysis, further motivated by Figure 8.6 and the associated discussion. Trigger efficiencies vs. jet kinematics and LLP truth variables are presented in Chapter 6. The HLT paths also have high efficiency for LLP decays in the HCAL, with LLP $|\eta| \leq 1.26$, and reach peak efficiency over a jet energy of 100 GeV.

8.2 Event Selection

Collision events are selected from the HCAL-based LLP displaced jet skim described in Chapter 7, meaning they are required to pass the HLT triggers described in Chapter 6. Further offline jet requirements are also imposed, requiring at least one jet per event to have $p_T \geq 40$ GeV and $|\eta| \leq 1.26$. Events in LLP Monte Carlo samples also have truth matching requirements imposed. The LLP is required to decay within the HCAL barrel (radial extent of $177 \leq r \leq 295$ cm) and with $|\eta_{LLP}| \leq 1.26$, and match to a PF AK4 jet.

8.2.1 Cut-Flow Table

Cutflow tables quantify how many events remain after each selection cut. The cuts in Table 8.1 are on truth variables, considering the position of the LLP and if it is triggerable, though without requiring LLP to jet matching. This is useful to separate acceptance (“is this an event we can possibly trigger?”) and efficiencies (“how good is the trigger at selecting events of interest?”). The main region of interest for the hadronic analysis is the HCAL volume to capture LLPs decaying within the calorimeter.

The selected cuts are motivated by the lowest un-prescaled L1 trigger thresholds, particularly informing the LLP energy (60 GeV) and event H_T (200 GeV) requirements. Tables 8.1 and E.2 show the number of events and fractions passing each selection for $m_H = 350$ GeV, while Tables E.1 and E.3 show the same for $m_H = 125$ GeV. The tables focus on radial selection cuts and timing selection cuts, respectively.

Selection	Number of Weighted Entries		Fraction of Weighted Entries	
	LLP 0	LLP 1	LLP 0	LLP 1
All	465995	465995	1.0	1.0
LLP r in HCAL depth 3 or 4	21439	21384	0.046	0.0459
LLP $\eta \leq 1.26$	14391	14307	0.0309	0.0307
LLP $E > 60$ GeV	14391	14307	0.0309	0.0307
LLP 0 OR LLP 1	28155		0.0604	
LLP 0 AND LLP 1	543		0.0012	
LLP OR, with Event $HT > 200$ GeV	24712		0.053	
LLP AND, with Event $HT > 200$ GeV	516		0.0011	
L1 passed (monitoring path)	34324		0.0737	
L1 + Leading jet neutral hadron frac $> 70\%$	25358		0.0544	
L1 + Leading jet neutral hadron frac $> 80\%$	24877		0.0534	
HLT dijet, ≤ 1 prompt track	16195		0.0348	
HLT dijet, 1 displaced track	2449		0.0053	
HLT dijet, ≤ 2 prompt track	11887		0.0255	
HLT delayed jet, ECAL timing	3707		0.008	

Table 8.1: Depth-based cutflow table, showing the number and fraction of events that pass the displaced LLP decay selections. The last portion of the table shows the fraction of events passing the L1 and HLT HCAL LLP triggers. $m_H = 350$, $m_S = 80$ GeV, $c\tau = 0.5$ m.

8.2.2 W+Jets Selection

From the displaced jet skim, W+jets are selected to use as a background sample, providing a large sample of unbiased prompt jets for use in comparing the custom jet object variables (Section 8.3) and in a data-driven background estimation (Section 9.1.2). The processes that fake LLP-like events are not necessarily modeled well in MC, so the sample of W+jets events contained in the displaced jet skim is used instead. As the W+jets from the displaced jet skim are used as a background sample, they are vetoed from the signal region.

The primary decay modes of the $W^\pm$ boson (conjugate for W^-) are [79]:

$$W^+ \rightarrow e^+ + \nu \quad (10.7\%) \quad (8.1)$$

$$W^+ \rightarrow \mu^+ + \nu \quad (10.6\%) \quad (8.2)$$

$$W^+ \rightarrow \tau^+ + \nu \quad (11.4\%) \quad (8.3)$$

$$W^+ \rightarrow \text{hadrons} \quad (67.4\%) \quad (8.4)$$

For W+jets events, those with an electron or muon in the final state are selected (comprising about 21% of the W boson decays). A recent CMS search for $W \rightarrow \ell + \nu$ in association with a charm quark presents a standard selection for W+jets [24], which we modify for use in our data-driven background estimation in the hadronic LLP analysis. The requirements on the lepton, jet, and event used in this LLP analysis are:

- Require exactly one lepton (muon or electron) satisfying the requirements:
 - $p_T^\ell > 20$ GeV and $|\eta^\ell| < 2.4$
 - Passes tight muon or electron ID (cut-based, as recommended by the muon and EGamma POGs)
 - Isolated with $I_{comb}/p_T^\ell < 0.15$, using the value recommended from [24]
 - * Tighter requirements are used for the electron: in the barrel, $I_{comb}/p_T^\ell < 0.695$ and in the endcap $I_{comb}/p_T^\ell < 0.0821$
 - Transverse mass of lepton $M_T \geq 55$ GeV to reduce QCD multijet background
 - Reject event if contains another isolated lepton with $p_T^\ell > 20$ GeV (reduce Z + jets and $t\bar{t}$ backgrounds)

- Jet requirements:

- $p_T^{jet} > 30 \text{ GeV}$ and $|\eta_{jet}| < 1.26$

- $\Delta\phi > 2$ between jet and ϕ direction vector sum of lepton and MET. This is stricter than the $\Delta R(jet, \ell) > 0.4$ required in [24]

- MET sum $E_T > 30 \text{ GeV}$

M_T , the transverse mass of the lepton, is defined as:

$$M_T = \sqrt{2 \cdot p_T^\ell \cdot p_T^{miss} [1 - \cos(\phi_\ell - \phi_{p_T^{miss}})]} \quad (8.5)$$

The combined isolation variable I_{comb} accounts for the charged, neutral, photon, and pileup isolation:

$$I_{comb} = I_{charged} + I_\gamma + I_{neutral \text{ hadron}} - [0.5 \cdot I_{\mu \text{ pileup}} \text{ OR } e_{\text{effective area}}] \quad (8.6)$$

Out of 80 million displaced HCAL LLP jet skim events in 2023, 50,000 are identified as W+jet candidate events (0.06%). At least one jet must satisfy the above requirements when an event is identified as a W+jets event. If multiple jets pass the requirements ($\Delta\phi > 2$ from the lepton), then all jets passing that requirement are used as W+jets background.

8.2.3 Minimal Signal Contamination

The W+jets selection provides a region depleted in the LLP signal. The signal contamination is very low, with $\approx 10^{-5}$ of the LLP MC signal samples identified as W+jets events, shown in Table 8.2. Therefore, as there is essentially no signal in the W+jets region, the W+jets events from the 2023 displaced jet skim can be analyzed without concern of unintentionally unblinding the data. W+jets events are plotted along with LLP signal MC to develop custom jet object variables that are used in a machine learning classifier with W+jets as the background training sample.

Signal Sample Parameters			Contamination in W+jets region
$m_H = 125$	$m_S = 15$	$c\tau = 1 \text{ m}$	$3.08 \cdot 10^{-5}$
$m_H = 125$	$m_S = 50$	$c\tau = 3 \text{ m}$	$2.45 \cdot 10^{-5}$
$m_H = 250$	$m_S = 120$	$c\tau = 10 \text{ m}$	$4.90 \cdot 10^{-5}$
$m_H = 350$	$m_S = 80$	$c\tau = 0.5 \text{ m}$	$9.07 \cdot 10^{-5}$
$m_H = 350$	$m_S = 160$	$c\tau = 10 \text{ m}$	$5.93 \cdot 10^{-5}$

Table 8.2: Signal contamination in the W+jets event selection, demonstrating that the W+jets region is depleted in signal, to the level of 10^{-5} and higher.

8.3 Sensitive Jet Object Variables

A set of calorimeter-based and track-based variables are developed to discriminate between LLP and prompt jets. To understand the distribution of the analysis variables in signal and background, the variables of interest are compared for MC LLP (with five mass and lifetime points from Table 7.1) and W+jets events selected from the 2023 displaced-jet skim dataset. The full selection for the W+jets events is described in Section 8.2.2, and this dataset provides a large set of prompt, Standard Model jet processes to compare the signal jets with. In the analysis variable plots, the MC LLP events are split by LLP decay position to identify observable results of LLPs decaying further into the calorimeters. LLPs are required to match to a PF jet, following the procedure in Section 8.1.2.

8.3.1 HCAL Rechit Based Variables

HCAL reconstructed hits (rechits) are used to define several custom variables to distinguish LLP jets from prompt jets. The HCAL rechits used are within $\Delta R \leq 0.4$ of the jet center (motivated by the use of AK4 jets) and pass the HCAL recommended PF thresholds. A rechit is considered associated with a jet when it satisfies $\Delta R(\text{jet}, \text{rechit}) = \sqrt{\Delta\eta^2 + \Delta\phi^2} \leq 0.4$, with $\Delta\eta = \eta_{\text{rechit}} - \eta_{\text{jet}}$ and $\Delta\phi = \phi_{\text{rechit}} - \phi_{\text{jet}}$.

PF Thresholds

The HCAL group recommends thresholds be applied to each rechit before the rechit is used in the particle flow algorithm. Many analysis variables are constructed from HB rechits, and only rechits above the PF energy thresholds are considered. Table 8.3 details the recommended PF cuts for Run-3 data.

Time Period	Depth 1 (GeV)	Depth 2 (GeV)	Depth 3 (GeV)	Depth 4 (GeV)
2022 initial	0.1	0.2	0.3	0.3
2022 re-reco	0.3	0.25	0.3	0.3
2023	0.4	0.3	0.3	0.3

Table 8.3: PF energy thresholds in GeV on rechits recommended by HCAL and in CMSSW. These are derived from the HB mean energy of rechits in neutrino + PU 2023 MC sample and used as the threshold for rechits in the analysis.

Custom HCAL Rechit-based Variables

The fractional HCAL energy deposited in each depth is plotted in Figure 8.3. This is computed by considering all HB rechits within $\Delta R \leq 0.4$ of a jet, and computing the fraction of the total associated rechit energy contained in rechits in each depth layer (by definition, for a single jet, the fraction in depth 1, 2, 3, and 4 will sum to 1). As the LLP decays further into the calorimeters, there is less energy deposited in depth 1 and more deposited in depths 3 and 4, and this distribution evolves slightly differently depending on the LLP decay location. This variable is, therefore, an interesting qualifier of LLP decay location and useful in a machine learning classifier approach.





Figure 8.3: The fraction of the total HCAL energy deposited in each depth is determined from HCAL hits associated with the jet within $\Delta R \leq 0.4$ (for signal, also requiring the jet is matched to a LLP). This differs considerably for LLPs decaying at various radii and is a helpful discrimination variable. Plots are for $m_H = 125$, $m_S = 15$ GeV, $c\tau = 3$ m.

The fraction of the jet energy contained in the most energetic associated HB rechit is shown in Figure 8.4a. For non-long-lived samples (W+jets and LLP decaying in the tracker), this is generally below 10%. As the LLP decays further into the calorimeter system, this fraction increases, as more

of the jet energy is contained in the HCAL rechits.



(a) Ratio of leading HB rechit energy to the total jet energy. This is higher for jets associated with a LLP decaying in the calorimeters since more energy is deposited in the calorimeter, and the fraction increases as the LLP decays with a larger radius.

(b) $S_{\phi\phi}$ for a jet. Jets from LLPs decaying in the calorimeters have less spread in ϕ , and therefore smaller $S_{\phi\phi}$ values.

Figure 8.4: Leading rechit fractional energy and jet $S_{\phi\phi}$, variables that use the HB rechit energy information. Plots are for $m_H = 125$, $m_S = 15$ GeV.

Figure 8.4b shows the HCAL hit-based variable $S_{\phi\phi}$, defined by the second moment of the spatial energy distribution (variance):

$$S_{\phi\phi} = \frac{\sum \Delta\phi^2 \cdot E_{\text{rechit}}}{E_{\text{total}}} \quad (8.7)$$

For each HCAL rechit associated with the jet, $\Delta\phi = \phi_{\text{jet center}} - \phi_{\text{HCAL rechit}}$ is weighted by the rechit energy, and normalized by the energy of all HCAL rechits matched to the jet. $S_{\phi\phi}$ has the best separation as compared to $S_{\eta\phi}$ and $S_{\eta\eta}$. Composite variables of S_{major} and S_{minor} (the major and minor axis of an ellipse, computed from $S_{\phi\phi}$, $S_{\eta\phi}$, $S_{\eta\eta}$) were also considered, but the ϕ direction spread was found to be a more powerful discriminating variable.

$S_{\phi\phi}$ encapsulates the minimal transverse spread of the LLP jets, as they have little time for the decay products (as the only charged particles present) to spread in the magnetic field before being stopped by the calorimeter. Jets resulting from LLPs decaying in the calorimeters are narrower in the ϕ direction (low value of $S_{\phi\phi}$) than background jets, as illustrated in Figure 8.5.

Low-Energy HB Rechits

In early studies of the rechit-based discrimination variables, investigations identified significant unexpected contributions from low-energy rechits at the edge of the barrel. Specifically, peaks at low leading rechit energies (such as in Figure 8.4a for LLPs decaying in HCAL depth 4) are primarily




Figure 8.5: Illustration of the spread of calorimeter energy in a prompt jet compared to a LLP jet. The LLP jet has a more clustered energy spread due to the decay in the calorimeter, creating a narrow jet cone. This spread is described by $S_{\phi\phi}$.

contributed to from rechits near the edge of the barrel. Before the fix outlined below, the low-energy rechit peak was present in LLPs decaying in the ECAL and all depths of the HCAL.

Events with a leading rechit energy < 8 GeV tend to have high energy jets, high energy fractions in depth 1, low energy fractions in depth 4, leading rechits in depth 1, wide energy spread within jet, result from LLPs with a large (290 cm) decay radius at high LLP η , and a low total number of HB rechits. These low energy leading rechit events are investigated in Figure 8.6: for jets resulting from LLP decays in HCAL, the jet η for those with leading rechit energy under and over 8 GeV is compared. A pattern where the lower energy rechits are primarily from jets near the barrel-endcap transition region of HCAL (above $|\eta| = 1.3$, or $|i\eta| = 15, 16$) is clear. Jets near the start of HE match to HB rechits based on a $\Delta R \leq 0.4$ cone. However, these jets tend to have a considerable contribution from endcap rechits, so the HB rechits do not accurately represent these HE jets. To address this, a cut on the jets used in the analysis is placed at $|\eta|_{\text{jet}} \leq 1.26$. Applying this jet positional cut significantly reduces the peaks at low rechit energies, and the remaining low energy rechit contribution is preferentially from LLPs decaying at high η and high decay R (depths 3 and 4), naturally creating a jet with fewer HB rechits.

The specific choice of $|\eta|_{\text{jet}} \leq 1.26$ cut is based on the HCAL layout (note that for $i\eta = 16$, depth 4




Figure 8.6: Leading jet η for jets matched to LLPs and split in categories based on the energy of the leading associated HB rechit. The low energy rechit peaks above $|\eta| = 1.3$ motivate the analysis level jet cut of $|\eta| \leq 1.26$ based on the physical HCAL tower geometry.

is in HE):

$$i\eta = 15 \text{ is } \eta = 1.26 \qquad i\eta = 16 \text{ is } \eta = 1.35 \qquad i\eta = 17 \text{ is } \eta = 1.44 \qquad (8.8)$$

The $|\eta|_{\text{jet}} \leq 1.26$ cut is physically motivated by the HCAL tower layout, as this is the center of the last full tower in HB ($i\eta = 15$). Furthermore, as seen in Chapter 6, the displaced and delayed jet trigger efficiency also peaks between $|\eta|_{\text{jet}} \leq 1.26$, so selecting this region with jet and LLP cuts aligns with the peak trigger efficiency. For jets in the region $-1.26 \leq \eta \leq 1.26$, the jet η distribution is flat and shows no differences between 2023 LLP skim data and 2023 LLP MC, as shown in Figure 8.7.

8.3.2 Track Based Variables

Figure 8.8 plots the fraction of jet energy reconstructed as neutral hadrons, charged hadrons, photons, electrons, or muons (by definition, these will add up to 1 for an individual jet). LLPs decaying in ECAL or HCAL have high neutral hadron fractions, as these are trackless jets and, therefore, cannot be reconstructed as charged particles in particle flow. A neutral hadron fraction cut at > 0.8 preserves over 95% of the LLP signal and reduces the background by $> 10^3$. LLP associated jets tend to have low charged hadron energy fractions and photon energy fractions below 20%. Both the electron and muon energy fractions are peaked at 0 for the signal, with a small portion of events having non-zero values for these, for which the distributions are shown in Figure 8.8d and e.




Figure 8.7: Jet η for background W+jets and LLP signal, split by LLP decay location. Distribution is flat with respect to η , except for a small decline at $|\eta| \geq 1.2$, as expected due to the LLP matching restriction of $|\eta|_{LLP} \leq 1.26$.


Figures 8.9 and 8.10 demonstrate the trackless quality of LLP jets. For LLP signal events, the ΔR between the jet and its leading track is approximately flat, while for backgrounds, $\Delta R(\text{jet}, \text{track})$ peaks below 0.1, as the most energetic track tends to be located with the jet center. This large ΔR value for signal events indicates that they are essentially “trackless,” i.e., the associated tracks are effectively from random tracks near the jet direction. Due to the trackless quality of signal events, the fraction of the jet p_T contained in the leading track is low. A cut of fractional track $p_T < 0.05$ provides almost 10^2 background rejection while preserving over 90% of the LLP signal.

Similar distributions for the sub-leading track are in Figure 8.10, further illustrating the trackless quality of these jets. The background distributions are more signal-like for the subleading track, as the ΔR separation is less clustered than the leading track.


8.3.3 Discussion of Sensitive Jet Object Variables

The distributions highlighted illustrate two main approaches to distinguishing signal and background: rechit-based and track-based. A combination of these is used to separate the signal and background jets. Correlations between track and rechit-based variables are further discussed in Section 8.4.3, while a summary of the jet characteristics identified from track and rechit-based variables is outlined here.


The rechit-based variables show that a significant fraction of the LLP jet energy is in the leading HB rechits, with characteristic energy spread between HCAL depths. Jets matched to LLPs decaying




(a) Neutral hadron fraction for background and LLP matched jets, with a high (> 0.7) fraction for jets associated with LLPs decaying in the calorimeters.




(b) Charged hadron fraction for background and LLP matched jets, with a low (< 0.2) fraction for jets associated with LLPs decaying in the calorimeters.



(c) Photon energy fraction for background and LLP matched jets, with a high (< 0.2) fraction for jets associated with LLPs decaying in the calorimeters.




(d) Electron energy fraction for background and LLP matched jets, with a very low (< 0.1) fraction for jets associated with LLPs decaying in the calorimeters.




(e) Muon energy fraction for background and LLP matched jets, with most jets having 0 muon energy fraction, and otherwise a widely spread distribution.

Figure 8.8: PF jet energy fraction variables: neutral hadron, charged hadron, photon, electron, and muon energy fraction. Jets from LLPs decaying in the calorimeters have particularly high neutral hadron fractions. Plots are for $m_H = 125$, $m_S = 15$ GeV, $c\tau = 3$ m.



(a) ΔR between the jet center and the leading track. This peaked below 0.1 for the W+jets background, while for LLPs with significant displacement (in either calorimeter), this rises toward 0.4.

(b) Ratio of the leading track to jet transverse momentum. The peak below 0.1 for jets associated with LLPs decaying in the calorimeter illustrates the trackless quality of these jets, as the leading tracks have a low momentum compared to the jet.




(c) $\Delta\eta$ between the jet center and the leading track.

(d) $\Delta\phi$ between the jet center and the leading track.


Figure 8.9: Variables illustrating the trackless quality of the LLP matched jets. The large ΔR values indicate that jets are matched to effectively random tracks, making the fractional track p_T low for signal. For jets with no matched tracks (more likely in signal), the values of the default ΔR , $\Delta\eta$, $\Delta\phi$ are 0.5.

within the calorimeter tend to have HCAL energy deposits located centrally around the jet (low $S_{\phi\phi}$), and a lower energy fraction in depth 1 and higher fraction in depths 3 and 4.


The track-based variables show that LLP jets are trackless when the LLP decays in the calorimeters. $\Delta\eta$ and $\Delta\phi$ between the jet and the associated leading track are approximately flat, while they are much lower for background or LLPs decaying in the tracker. Thus, ΔR is high for signal, as few tracks correlate well with the jet. Similarly, the ΔR between the leading and sub-leading track is approximately flat for the signal and peaked lower for the background. These distributions support the conclusion that the signal is mostly trackless, and track matching starts associating




(a) ΔR between the jet center and the subleading track. This peaked below 0.1 for the W +jets background, while for LLPs with significant displacement (in either calorimeter), this rises toward 0.4.



(b) Ratio of the subleading track to jet transverse momentum. The peak below 0.1 for jets associated with LLPs decaying in the calorimeter illustrates the trackless quality of these jets, as the subleading tracks have a low momentum compared to the jet.



(c) $\Delta\eta$ between the jet center and the subleading track.




(d) $\Delta\phi$ between the jet center and the subleading track.

Figure 8.10: Variables illustrating the trackless quality of the LLP matched jets, shown for the sub-leading track.


essentially random tracks to the jet. This is further supported by the high neutral hadron energy fraction observed for signal events. With a high energy deposit in the HCAL but without many high energy tracks that are associated with the jet, much of the energy is determined to be from neutral hadrons.

Event displays provide a visual way to understand the energy distribution resulting from a LLP decay. Figure 8.11 shows an event where a LLP decayed in the HCAL (using a skim of MC LLP events), first plotting the energy deposits in HCAL near the jet center (a) and then the overall visualization of the whole event (b). From Figure 8.11a it is clear that most of the HCAL energy is deposited centrally in depths 3 and 4. This is what variables such as $S_{\phi\phi}$ and energy fractions

within HCAL depths capture. Based on these event displays, a convolutional neural network may be a promising direction in developing a strong classifier for LLP selection based on calorimeter “images”.



(a) Calorimeter focused event display of the HB energy deposits (scale from yellow to blue), tracks (orange circles), and LLP position (red circle). HCAL depths, from the top left, are depth 1, 2 (top row), depth 3, 4 (bottom row).



(b) Global event display of the same event, with one LLP decaying in the HB and creating a cluster of calorimeter hits.

Figure 8.11: Event display of a LLP event, with one LLP decaying in the calorimeter and producing a jet with most energy deposited in HCAL depths 3 and 4. Created from code in the HCAL LLP analysis EventDisplay GitLab repo [58], with an event from $m_H = 250$, $m_S = 120$ GeV, with $c\tau = 10$ m.

Additional Variables Investigated

Ongoing studies explore the separation power of additional track-based variables such as the track hit multiplicity, the number of missing tracker hits, and impact parameter significance. Impact parameters describe how close to the primary vertex the track approaches, and are separated into the xy and z directions. Typically, the significance of impact parameters is used, where significance is d_{xy}/error_{xy} (and similarly for z). Thus, a high significance means the track is far from the primary vertex and has a low error.

As for rechit-based variables, jet mass, elliptical shower shape, and HCAL timing are investigated. These are additional candidate variables for separating the signal region. HCAL timing is not in place in the L1 trigger for 2022, and this analysis focused on a depth-based approach to identify LLPs decaying in the calorimeters. Though timing variables are not currently a key variable in this search, we aim to validate the timing variables for offline analysis use.


8.3.4 Dependence on Jet Energy

Understanding the dependence of the analysis variables on jet energy is crucial, as different signals and samples have different energy scales. Ideally, the analysis variables are robust against changes in energy scales – i.e., if a jet has higher energy, the above trends will still identify LLP jet characteristics from standard jets.


A selection of track- and rechit-based analysis variables for LLP MC and blinded data, split by jet p_T , is shown in Figures 8.12 and 8.13. The slight dependencies seen on jet energy move in the expected direction – higher energy jets are more collimated (lower $S_{\phi\phi}$), tend to have a lower neutral hadron fraction, and an energy fraction in depth 1 that is less broadly spread due to less pile-up contribution. The correlations between each variable are investigated further in Section 8.4.3.

8.4 Discriminator: Boosted Decision Tree


A boosted decision tree (BDT) is a machine learning method to separate categories of images, data, and, in this case, physics events. A decision tree is a simple classifier with several “nodes”, that make a binary decision about the data. This could be of the form “Is the jet energy greater or less than 60 GeV?”, and the answer to that query determines which path down the tree is followed next. A boosted decision tree uses gradient boosting to combine each of these trees into an ensemble, with each tree weighted according to its performance or accuracy to minimize errors and improve the




(a) Energy fraction in depth 1, split by jet energies.



(b) Neutral hadron fraction, split by jet energies.



(c) $S_{\phi\phi}$ split by jet energies.



(d) Fractional track p_T , split by jet energies.


Figure 8.12: Discriminating variables, split by jet p_T bins in LLP MC, with $m_H = 125$, $m_S = 15$ GeV. Most variables are relatively independent of jet energy, though higher jet energies correspond to lower $S_{\phi\phi}$ values.

final result.


The BDT used here was developed with TMVA, a library in ROOT that performs multivariate analysis. Various LLP MC datasets are input as signals, and W+jets selected from the 2023 HCAL displaced jet skim are the background samples.

8.4.1 Training Variables


The variables input into TMVA for the BDT training are those identified in the preceding section to provide separation between signal (LLPs decaying in the calorimeters) and background (prompt decays and W+jets). The variables input to the BDT describe where the jet energy is deposited, the trackless quality of signal jets, and the specific spread (lateral and radial) of energy in the HCAL, as shown in Figures 8.3 through 8.10. The full list of BDT input variables is:




(a) Energy fraction in depth 1, split by jet energies.



(b) Neutral hadron fraction, split by jet energies.



(c) $S_{\phi\phi}$ split by jet energies.



(d) Fractional track p_T , split by jet energies.

Figure 8.13: Discriminating variables, split by jet p_T bins with blinded data (blinded with neutral hadron fraction ≤ 0.6).

- Jet η
- Jet neutral hadron energy fraction
- Jet charged hadron energy fraction
- Jet photon energy fraction
- Jet electron energy fraction
- Jet muon energy fraction
- Fraction of jet HCAL energy deposited in depths 1, 2, 3, and 4
- Jet $S_{\phi\phi}$
- Fractional leading rechit energy (fraction of jet energy contained in most energetic HCAL rechit)

- Jet to associated track $\Delta\eta$ and $\Delta\phi$ for the first two leading tracks
- Fractional p_T of the two leading tracks associated with the jet (fractional p_T is $\frac{\text{track } p_T}{\text{jet } p_T}$)

The jet muon energy fraction and the fraction of HCAL energy deposited in depth 4 are not input to the BDT. TMVA does not allow input variables to be redundant, and in both of those cases, the information is already given by other input variables:

$$\sum_{i=1}^4 \frac{E_{\text{depth } i}}{E_{\text{total}}} = 1 \quad (8.9)$$

and

$$E_{\text{electron}} + E_{\text{photon}} + E_{\text{charged hadron}} + E_{\text{neutral hadron}} + E_{\text{muon}} = 1 \quad (8.10)$$

Leading jet ϕ was removed from the training as it is uneven in the 2023 data, with differences in $\pm\pi$. The leading jet energy and momentum are also not trained on, to avoid biasing the training on jet energy and p_T . The jet energy is used to normalize the leading rechit energy (BDT input is fractional leading rechit energy), as the leading rechit energy scales with jet energy; while the jet momentum is used to normalize the track momentum to give fractional track p_T .

8.4.2 BDT Setup

Event selections described in Section 8.2 are applied to the signal and background, with additional requirements that energy fractions in each depth are between 0 and 1 and $S_{\phi\phi} > 0$ to avoid any cases where there are no HB hits associated with the selected jet. Events passing these selections are used in the BDT training and testing, with the W+jets sample used as the background.

8.4.3 BDT Performance


Receiver operating characteristic (ROC) curves show the performance of a classifier by plotting signal and background acceptance. The output of a classifier, such as a BDT, is a score assigned to each event or object, from -1 to 1 . Ideally, the background scores are near -1 , while the signal scores are close to 1 . A certain amount of signal and background will pass a given cut on the BDT score, and as the BDT score cut threshold is scanned, the full range of the ROC curve is determined. Figure 8.14 shows the fraction of signal and background passed at various BDT discriminators (a) and the ROC curve found by scanning the BDT discriminator (b).

ROC curves are created by evaluating the performance of “testing” samples, while the BDT is created

based on “training” samples. Using different data sets for training and testing is required to ensure that the data validation procedure can be completed and avoid overtraining (learning the specifics of the input events to the extent that the BDT cannot separate a new set of signal and background events well). TMVA avoids overtraining by comparing the testing and training efficiency. The split between training and testing datasets is 25% (training) and 75% (testing), except for samples with lower statistics, where a split of 40% (training) and 60% (testing) is used. The 40-60 train-test split is used for samples in Table 7.1 with fewer than 4 million events.

The BDT discriminator score is shown in Figure 8.14a, with signal (blue) and background (red) samples overlayed. The limiting factor on performance is not signal acceptance, as the signal scores are highly clustered above 0.75. Rather, the limiting factor is rejecting the remaining 50 events with very high BDT scores (> 0.95). Sufficiently rejecting the background that appears LLP signal-like is one of the challenges of this analysis.

Figure 8.14b demonstrates the signal and background separation of the trained classifiers in a ROC curve. A classifier is trained for each signal sample and for a combination of signals (six BDTs), and these are all tested on a given signal sample, for example, $m_H = 350$, $m_S = 80$ GeV, $c\tau = 0.5$ m as shown in the plot. The best performing classifier is that trained on $m_H = 350$, $m_S = 80$ GeV. As this classifier also performs well on most other signal samples (shown in Figure E.1), it is used in the following background estimation studies in Chapter 9.



(a) BDT discriminator scores for signal and background.

(b) ROC curve for signal, tested on BDTs trained on each signal sample.

Figure 8.14: BDT discriminator distribution and ROC curve, for $m_H = 350$, $m_S = 80$ GeV, $c\tau = 0.5$ m. Reference points are shown for comparison to a cut-based approach (blue) and the point at which the discriminator score is 0.9995 (black), with about 80% signal efficiency and almost 10^4 background rejection.

The BDT trained on $m_H = 350$, $m_S = 80$ GeV, with a $c\tau = 0.5$ m demonstrates the best background

suppression while maintaining high signal efficiency. Input variables to the classifier are plotted to compare different signals, and a selection is shown in Figure 8.15. Note that the two samples with LLP mass close to 25% of the Higgs mass tend to have analysis variable distributions that are more extreme – i.e., either pushed toward 0 or 1, in the direction that has been identified as more signal-like. This is shown in the low $S_{\phi\phi}$ value and HCAL depth 1 energy fraction, and the high neutral hadron energy fraction and leading rehit energy. This trend explains why the BDT trained on $m_H = 350$, $m_S = 80$ GeV is one of the best performing classifiers.




Figure 8.15: Discriminating analysis variables for different LLP signals overlayed. Note that the mass ratio between the Higgs and LLP affects the result, splitting the plots into two groups (one with LLP mass close to 25% of the Higgs mass, the other with LLP mass close to 50% of the Higgs mass).

Some variables are highly correlated in the BDT, such as neutral hadron fraction and $S_{\phi\phi}$, with a correlation of about -50. A negative correlation means lower $S_{\phi\phi}$ (more signal-like) corresponds with a higher neutral hadron fraction (also more signal-like). Neutral hadron fraction is also correlated

with leading HB reconstructed hit energy, and inversely with fractional track p_T , fractional energy in depth 1, photon energy fraction, and charged hadron energy fraction. Figure 8.16 illustrates the correlation between each input variable.




Figure 8.16: Correlation plot from the BDT training, showing the degree of correlation or anti-correlation between all input variables.

8.4.4 Performance of Triggered Events

To demonstrate that the L1 and HLT HCAL-based LLP triggers select events that perform well at the analysis level, ROC curves are evaluated for events with a LLP decay in the HCAL, separately for events passing any of the HCAL based HLTs, and events not passing the triggers. As shown in Figure 8.17, events passing the HLTs score higher in the BDT as compared with events not passing the HLTs. For $m_H = 125$, $m_S = 15$ GeV, $c\tau = 3$ m the difference between HLT triggered events and all events is also pronounced, though the overall BDT performance is lower.

The performance of triggers on events selected for the analysis can also be determined from the fraction of LLP MC events with a jet passing a given BDT cut that also has the HCAL-based HLT set. For $m_H = 350$, $m_S = 80$ GeV, $c\tau = 0.5$ m, 8.5% of the events have 1+ jets with BDT score over 0.9995 (requiring jet energy ≥ 40 GeV and event $H_T \geq 200$ GeV). Of those events, 56% have an HCAL-based HLT set. Similarly, of the events passing any of the HLTs (7.4%), 71% of them have at least 1 jet with a BDT score over 0.9995.




Figure 8.17: ROC curve for $m_H = 350, m_S = 80$ GeV, $c\tau = 0.5$ m, requiring that LLP decays in the HCAL. Events that pass the HCAL-based HLTs perform slightly better than the sample of all LLP events not passing the selected HLTs, demonstrating that the trigger selects events that perform well at the analysis level. This ROC curve is evaluated on a per-event basis, evaluating the performance of the highest p_T jet.

This indicates that the LLP triggers flag a large fraction of jets that score highly in the BDT. The events that are selected by the BDT but not by the HLTs may be due to the trigger selecting dijets, whereas the BDT is trained on single jets from LLPs decaying in the calorimeter volume. The tracking inefficiencies (imperfect selection of displaced tracks) present in the HLTs potentially contribute as well.

8.4.5 BDT Cutflow Comparison

A jet-based cutflow is performed to compare with the BDT results. The cutflow results for signal are presented in Table 8.4 for $m_H = 350, m_S = 80$ GeV, and in Table 8.5 for the W+jets background. Additional signal points are in Tables E.5, E.6, E.7, and E.8. The signal efficiencies for the BDT comparison cutflow are 66% ($m_H = 125, m_S = 15$), 58% ($m_H = 125, m_S = 50$), 65% ($m_H = 250, m_S = 120$), 77% ($m_H = 350, m_S = 80$), 68% ($m_H = 350, m_S = 160$ GeV). This is compared with background efficiency of 0.08% (Table 8.5). These cutflow comparisons are added to the ROC curves, as shown in Figure 8.14 by the blue cross mark.

Selection	Number of Entries	Fraction of Entries	
		Of all entries	Of jet matched entries
All (jet $ \eta \leq 1.26$ and $p_T \geq 40$)	357901	1.0	
Jet matched to LLP in HCAL	31907	0.0892	
Neutral hadron fraction $> 80\%$	30967	0.0865	0.9705
Charged hadron fraction $< 10\%$	30468	0.0851	0.9549
$S_{\phi\phi} < 0.02$	30181	0.0843	0.9459
Track $ \Delta\eta > 0.05$	27229	0.0761	0.8534
Track $ \Delta\phi > 0.05$	24628	0.0688	0.7719
Leading rehit energy > 8 GeV	24591	0.0687	0.7707

Table 8.4: Per jet cutflow table for $m_H = 350$, $m_S = 80\text{GeV}$, $c\tau = 10$ m to compare with BDT results.

Selection	Number of Entries	Fraction of Entries
All (jet $\eta \leq 1.26$ and $p_T \geq 40$)	81201	1.0
Neutral hadron fraction $> 80\%$	176	0.0022
Charged hadron fraction $< 10\%$	131	0.0016
$S_{\phi\phi} < 0.02$	123	0.0015
Track $ \Delta\eta > 0.05$	86	0.0011
Track $ \Delta\phi > 0.05$	72	0.0009
Leading rehit energy > 8 GeV	69	0.0008

Table 8.5: Per jet cutflow table for W+jets background to compare with the BDT performance.

Chapter 9

Backgrounds and Sensitivities

The trackless characteristics and calorimeter energy deposition patterns of LLP decays within the calorimeter have been identified and used in a classifier to distinguish signal from background, achieving a background rejection of about 10^4 based on a single jet tag approach. The next step in the analysis is to predict the backgrounds expected in the LLP signal region and estimate the expected sensitivities.

9.1 Background Estimations

In many analyses, Monte Carlo events are used to simulate and estimate background contributions; however, in LLP analyses, the backgrounds are often where Standard Model particles have edge cases that look like LLPs, which is difficult to model well. In these types of searches, misidentified long-lived particle jets can result from detector resolution effects, and this contribution is measured in a data control region. A data-driven background estimation such as the ABCD or fake-rate method is commonly used in LLP searches. This analysis utilizes the fake-rate method described in Section 9.1.2, where one estimates the background contributions in the signal region by measuring the mis-tag rate in control regions.

Similar analyses, such as a Run-2 ECAL-based LLP search, have found that the dominant backgrounds are jets from Standard Model processes that are mistagged as delayed or displaced, potentially due to tracking inefficiency, causing jets to appear trackless [25]. Similar to the approach used in the ECAL-based analysis, we assume the mistag efficiencies of each jet are independent and, therefore, can apply a data-driven fake-rate background estimation method. The backgrounds used

for this estimate are W+jets events in data, selected as described in Section 8.2.2, and are used to identify the probability of misidentifying a prompt jet as a BDT-flagged LLP candidate jet.

9.1.1 Sources of Background

The kaon (K-long, a long-lived neutral kaon, with a mean proper lifetime of $c\tau = 15.3$ m) is a source of irreducible background for this search. The primary decays are to pions, or a pion plus two leptons. The K-long hadronic decay provides a significant source of neutral hadron energy, an important discriminating variable for this analysis. The primary hadronic decay modes of the neutral kaon are [79]:

$$K_L^0 \rightarrow 3\pi^0 \quad (19.5\%) \quad (9.1)$$

$$K_L^0 \rightarrow \pi^+\pi^-\pi^0 \quad (12.5\%) \quad (9.2)$$

For completeness, the primary semileptonic decay modes are:

$$K_L^0 \rightarrow \pi^\pm e^\pm \nu_e \quad (40.6\%) \quad (9.3)$$

$$K_L^0 \rightarrow \pi^\pm \mu^\pm \nu_\mu \quad (27.0\%) \quad (9.4)$$

Decays of hadronic taus are also important backgrounds to consider. The main decay modes of the tau are [79]:

$$\tau^\pm \rightarrow \rho^\pm \nu_\tau \rightarrow \pi^\pm \pi^0 \nu_\tau \quad (25.5\%) \quad (9.5)$$

$$\tau^\pm \rightarrow \pi^\pm 2\pi^0 \nu_\tau \quad (9.3\%) \quad (9.6)$$

In the hadronic decays of the K-long and the $\tau^\pm$, the resulting neutral pion deposits energy in the ECAL (due to photons from the neutral pion decay) and is thus reconstructed as a non-isolated photon. However, due to the long lifetime of the neutral kaon, the neutral pion could be deposited in the HCAL and reconstructed as neutral hadron energy. The charged pion will produce an HCAL deposit. The significant neutral hadron energy fraction makes this likely to contribute as a background.

Background contributions from lambda particles are also valuable to consider. The decay modes of

the neutral lambda baryon are [79]:

$$\Lambda^0 \rightarrow p^+ \pi^- \quad (63.9\%) \quad (9.7)$$

$$\Lambda^0 \rightarrow n^0 \pi^0 \quad (35.8\%) \quad (9.8)$$

The neutron, proton, and pions can contribute to hadronic energy, though only the neutron and π^0 are neutral. However, Λ^0 and K_S have shorter lifetimes, and if they decay in the tracker, they would be excluded based on the tracks.

Backgrounds from cosmic rays or beam halo

Non-collision backgrounds may include cosmic muons or beam halo. Cosmic muons may hit the back of the HCAL, causing the depth trigger to fire, as a significant amount of energy can be deposited in the later depth layers. The contribution of cosmic muons will be studied with a dataset requiring at least one cosmic muon (identified from hits in the muon system).

Beam halo backgrounds are machine-induced backgrounds from proton collisions with gas in the beam pipe, causing a muon parallel to the beam direction [25]. These muons may be observed in the calorimeters, producing a relatively high energy cluster (up to 100s of GeV). Dedicated filters are available to identify beam halo events from calorimeter deposits. As a check, we will compare the number of BDT-tagged jets in a control region with the beam halo filter applied and inverted to understand if beam halo events impact the signal region.

9.1.2 Fake Rate Method


LLP backgrounds are difficult to model, so a data-driven background estimation is employed. For the background estimation method, the mistag rate is evaluated in a signal-depleted region. The signal-depleted region used is W+jets events produced in 2023 collisions, as the LLP signal does not involve W-boson production. The W+jets region is depleted in signal to the level of 10^{-5} , as demonstrated in Table 8.2. The mistag rate is the rate of non-signal jets identified as LLP candidates, and is found by:

- Numerator: jets above a selected p_T threshold (40 GeV) with $|\eta| \leq 1.26$ above a selected BDT discriminator threshold (considered a LLP candidate jet)
- Denominator: jets above a selected p_T threshold (40 GeV) with $|\eta| \leq 1.26$

The mistag rate is evaluated vs. kinematic variables and then used to estimate the number of background events in a signal region (requiring event pre-selections and at least one tagged jet).

Fraction of Mis-tagged Events in W+Jets

The jet mistag rate vs. jet p_T and η is shown in Figure 9.1. There is a higher probability of a high BDT score at lower jet p_T . Cross-checking also demonstrated that the mistag rate is flat vs. jet ϕ . Note that these are more relaxed BDT cuts than those that will be used to select the signal.



(a) Number of mistagged jets (top) and fraction (bottom) vs. jet η .

(b) Number of mistagged jets (top) and fraction (bottom) vs. jet p_T .

Figure 9.1: Jets from a W+jets sample that are tagged with a high BDT score ($> 0.5, 0.9, 0.99, 0.999$) for the classifier trained on $m_H = 350$, $m_S = 80$ GeV, $c\tau = 0.5$ m.

The best performing BDT is the one trained on $m_H = 350, m_S = 80$ GeV, $c\tau = 0.5$ m. This BDT is run over each “valid” jet in events from 2023 data and the LLP samples, where a valid jet is considered one with $p_T \geq 40$ GeV and $|\eta| \leq 1.26$. Up to 6 valid jets per event are saved per event for both data and signal. The background estimation is then performed by determining how many W+jets events have no tagged jets, 1 tagged jet, and 2+ tagged jets. Results are presented in Tables 9.1 and 9.2.

Jet cuts of 60 GeV were also tested, however, the reduction of W+jets background in Table 9.1 is relatively small, minimally decreasing the single jet tagged region. The jet cut of 40 GeV is currently

used for all tables since the impact of increasing the jet cut was less than a 20% reduction on the W+jets background region.

Selection	Number of Entries	Fraction of Entries
All events (from skim)	48574	1.0
Jet ≥ 40 GeV p_T (one of 6 leading)	48354	0.9955
0 jets with BDT score ≥ 0.9999	48352	0.9954
1 jet with BDT score ≥ 0.9999	2	$4.117 \cdot 10^{-5}$
2+ jets with BDT scores ≥ 0.9999	0	0.0

Table 9.1: W+jets events in 2023 data, showing the fraction of events passing jet kinematic cuts and single or double jet tags, with a jet p_T threshold of 40 GeV.

Background Prediction from 1-tag Measurement

The fraction of events in the W+jets selection passing the one jet selection is used to predict the one jet tag region in the 2023 displaced jet skim. In the simplest approach, the fraction of W+jets events passing the 1 jet tag from Table 9.1 would be used as the probability for an event having a single jet mistagged, thus $p_{\text{mistag}} = 4.117 \cdot 10^{-5}$. However, this overlooks an expected higher mistag rate on triggered jets. To be considered in the W+jets selection, the event as a whole must pass the HCAL-based LLP triggers (as these events are selected from the displaced jet skim), and individual jets can be split by jets that have passed the L1 jet trigger, and those that have not. Separate mis-tag efficiencies are measured for each of these.

To account for the bias from triggered vs. non-triggered jets, the efficiency measurement is split into two regions: **trigger matched**, using jets that did fire the trigger; and **not triggered**, using jets that did not fire the trigger (but in an event where another jet is confirmed to fire the trigger). As the LLP flag is saved as part of the L1 jet word (Section 5.3.2), a determination can be made for each reconstructed jet if the Level 1 jet it matches is a triggered jet or not. For trigger matching, up to six “valid” reconstructed jets are considered, and the closest L1 jet to the reconstructed jet is identified. A $\Delta R(\text{reco}, \text{L1}) \leq 0.4$ is required. The ΔR distribution of both signal and data events supports the choice of $\Delta R \leq 0.4$, which is further motivated by using AK4 PF jets. Finally, the background estimation cutflow is split into two categories to find the probabilities that a triggered or non-triggered jet is mistagged.

Table 9.2 shows the fraction of events with 1 jet scoring highly in the BDT, split by all events, events where the jet is triggered, and events where the jet is not triggered. A BDT cut of 0.9999 is imposed, and at higher BDT cuts, the mistag probability for non-triggered jets decreases faster than the mistag probability for triggered jets. At a BDT cut of 0.99, the mistag probability for

non-triggered jets is 0.5 that of triggered jets, while at 0.9995, the probability for non-triggered jets is 0.25 that of triggered jets.

Selection	Number of Entries	Fraction of Entries	
		All	Valid Jet
All events (from skim)	48574	1.0	
Jet ≥ 40 GeV p_T (one of 6 leading)	48354	0.9955	
Jet $ \eta \leq 1.26$	48354	0.9955	1.0
0 jets with BDT score ≥ 0.9999	48352	0.9954	1.0
1 jet with BDT score ≥ 0.9999	2	$4.117 \cdot 10^{-5}$	$4.136 \cdot 10^{-5}$
2+ jets with BDT scores ≥ 0.9999	0	0.0	0.0
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, triggered	21907	0.451	1.0
0 jets with BDT score ≥ 0.9999	21906	0.451	1.0
1 jet with BDT score ≥ 0.9999	1	$2.059 \cdot 10^{-5}$	$4.565 \cdot 10^{-5}$
2+ jets with BDT scores ≥ 0.9999	0	0.0	0.0
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, not triggered	20484	0.4217	1.0
0 jets with BDT score ≥ 0.9999	20484	0.4217	1.0
1 jet with BDT score ≥ 0.9999	0	0.0	0.0
2+ jets with BDT scores ≥ 0.9999	0	0.0	0.0

Table 9.2: W+jets events cutflow, split by trigger matched and unmatched jets that score highly in the BDT. A BDT score cut of 0.9999 is imposed. The mistag probability is higher for the trigger-matched jets, as expected.

From Table 9.2, take $1/21907 = 4.565 \cdot 10^{-5}$ as the probability that an event containing at least one triggered jet over 40 GeV also scores over 0.9999 in the BDT discriminator. The number of events predicted to have one tag in 2023 data (in a control region, not signal) is shown in Table 9.3. The estimation of the number of events in the double jet tagged region is estimated as $p_{2 \text{ tagged jets}} = p_{1 \text{ tagged jet}}^2$, making the assumption that the probability of mistagging jets within the same event is independent of each other. Note that the probability used for the background prediction is based on one remaining W+jets event, and therefore has a large uncertainty (± 1 event).

Selection	Number of Entries	Fraction of Entries
All events (2023 LLP skim)	82312351	1.0
Jet ≥ 40 GeV p_T (one of 6 valid jets)	80358669	0.9763
One tagged jet prediction ($p = 0.000046$)	3668	0.000046
Two tagged jet prediction ($p = 0.000046$)	0.17	0.000046^2
Jet ≥ 60 GeV p_T (one of 6 valid jets)	72985190	0.8867
Jet $ \eta \leq 1.26$	72985190	0.8867
One tagged jet prediction ($p = ?$)	not enough stats	-
Two tagged jet prediction ($p = ?$)	-	-

Table 9.3: Predicted backgrounds for the single tag and double tag regions in the full 2023 skim dataset. Predictions are made from the fraction of W+jets events with a triggered jet falling in the single tag region, from Table 9.2 (and a similar one with a jet cut of 60 GeV) with a BDT cut of 0.9999.

Fine-tuning of the BDT working point selection is ongoing, with results at a cut of 0.9999 presented here. Shown in Appendix E.3, the results at a lower working point of 0.9995 allow for more signal acceptance, though predict 13 events in the double jet tagged region with a jet energy of 40 GeV (Table E.14), as compared to < 1 event with a BDT cut of 0.9999 in Table 9.3.

Background Estimation in the Signal Region

Table 9.4 presents the number of events with a single jet passing a BDT score of 0.9999 in 2023 data. The 0 jet tag and 2+ jet tag regions are blinded to keep the analysis unbiased. The number of events seen in the 1 jet tag bin is higher than predicted, indicating further iterations on the background estimation are necessary. Specifically, there are four times the predicted number of events from the W+jets approach in the 1 jet tag bin. For the W+jets prediction to match this result, 4 events in the double tag region would need to be predicted. Considering a Poisson distribution with a mean of 4, the probability that the background estimation predicts 1 event or less occurs as a fluctuation 9% of the time (≤ 2 events would be more likely, occurring 24% of the time). This motivates further work on the background estimation to achieve closure in the single jet tag region. Potential contributions are enhanced backgrounds or the inclusion of other backgrounds in the triggered sample. Further studies of the characteristics of the W+jets events (with more statistics, and at a looser BDT cut) as well as studies with other control samples (Z+jets) are underway, and the ongoing work is discussed in Section 9.1.3.

Selection	Number of Entries	Fraction of Entries	
		All	Valid Jet
All events (from skim)	82312351	1.0	
Jet ≥ 40 GeV p_T (one of 6 leading)	80358669	0.9763	
Jet $ \eta \leq 1.26$	80358669	0.9763	1.0
1 jet with BDT score ≥ 0.9999	15575	0.0002	0.0002
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, triggered	35360640	0.4296	1.0
1 jet with BDT score ≥ 0.9999	6821	0.0001	0.0002
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, not triggered	29505082	0.3585	1.0
1 jet with BDT score ≥ 0.9999	2811	0.0	0.0001

Table 9.4: Actual 1 tag regions from 2023 skim data, at a BDT cut of 0.9999.

The double tag requirement will increase the background suppression by a factor of $2 \cdot 10^4 - 5 \cdot 10^5$ (precise values will result from improvements in the background estimation methods, validated by the ongoing closure tests in Section 9.1.3). Given the ratios in Table 9.4, a first order estimate predicts 1.3 events in the double tagged bin ($6821 \cdot 0.0002$), when requiring that the highly scoring jet is trigger matched. Though further background estimation work is still underway, this demonstrates

that a very low background level in the signal region is expected.

Signal Efficiencies at Selected BDT Cuts

Table 9.5 evaluates the signal efficiency at a BDT discriminator cut of 0.9999, both for events with jet kinematic cuts and HLT acceptance cuts. Tables for each of the other four signal samples are in Tables E.9 through E.12. Similar results for a lower BDT cut of 0.9995 are in Appendix E.3, Tables E.15 through E.19.

Selection	Number of Entries	Fraction of Entries	
		All	Valid Jet
All	299110	1.0	
Jet ≥ 40 GeV p_T and $ \eta \leq 1.26$	249804	0.8352	1.0
0 jets with BDT score ≥ 0.9999	232685	0.7779	0.9315
1 jet with BDT score ≥ 0.9999	16745	0.056	0.067
2+ jets with BDT scores ≥ 0.9999	374	0.0013	0.0015
LLP HLT passed	21982	0.0735	0.088
HLT and 0 jets with BDT score ≥ 0.9999	12447	0.0416	0.0498
HLT and 1 jet with BDT score ≥ 0.9999	9242	0.0309	0.037
HLT and 2+ jets with BDT scores ≥ 0.9999	293	$9.796 \cdot 10^{-4}$	0.0012
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, triggered	21530	0.072	1.0
LLP HLT passed	21396	0.0715	0.9938
HLT and 0 jets with BDT score ≥ 0.9999	12096	0.0404	0.5618
HLT and 1 jet with BDT score ≥ 0.9999	9201	0.0308	0.4274
HLT and 2+ jets with BDT scores ≥ 0.9999	99	0.0003	0.0046
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, not triggered	242792	0.8117	1.0
LLP HLT passed	17316	0.0579	0.0713
HLT and 0 jets with BDT score ≥ 0.9999	16919	0.0566	0.0697
HLT and 1 jet with BDT score ≥ 0.9999	397	0.0013	0.0016
HLT and 2+ jets with BDT scores ≥ 0.9999	0	0.0	0.0

Table 9.5: Fraction of signal events passing jet kinematic selections and split in each of the three jet tagging bins (no jets tagged, 1 tagged, 2+ tagged), with a jet p_T threshold of 40 GeV and a BDT score of 0.9999. The fraction that also passes any of the relevant HCAL-based HLTs is also shown, indicating that the events with two jets scoring highly in the BDT also pass the LLP HLTs 78% of the time. From $m_H = 350$, $m_S = 80$ GeV, and $c\tau = 0.5$ m.

40% of the events selected by the HLTs have at least one jet with a BDT score over 0.9999. Additionally, of all events selected by the 2+ jet tagged requirement, 78% of these pass the LLP HLTs, indicating that the triggered events perform well in the analysis. Of the events with a trigger matched jet, nearly all of them pass the HLT. Furthermore, in 99.6% of the events that pass the HLTs and have a jet scoring over 0.9999, the high-scoring jet is identified as trigger matched. Both of these metrics identify that the triggered jets align well with the selections applied in the analysis.

Figure 9.2 compares the number of events in the 0, 1, and 2+ jet tag bins for W+jets events identified in 2023 data with LLP events with $m_H = 350$, $m_S = 80$ GeV, $c\tau = 0.5$ m. With the dedicated

HLTs required, the fraction of signal events in the 0 tag bin is reduced, while the fraction in the 2+ tag bin is minimally impacted. This further illustrates that the triggers select events that pass the analysis selections, which allows for a further reduction of backgrounds offline while maintaining signal efficiency.




Figure 9.2: Fraction of events in each jet tag category (0 tag, 1 tag, 2+ tag) from W+jets events in 2023 and LLP MC, as in Tables 9.2 and 9.5, evaluated with respect to triggered matched events (W+jets) or events passing the HLT selection (LLP) with a jet passing the energy and η requirements. A jet threshold of 40 GeV is required, and the BDT score required is 0.9999.

9.1.3 Closure Tests

Multiple approaches are underway to further improve the background estimation method. We will continue developing and improving the background estimation method until agreement in the single-tagged jet region is achieved between the predictions and data. Currently, the prediction is lower than the observed number of single tagged events (Tables 9.3 and 9.4).

Background Estimation Improvements

The current approach treats the jet mistag probability inclusively, and to improve upon this method, the background estimation will be split into bins of jet kinematic variables (η , p_T , etc.) to account for the impact of jet kinematics in the fake-rate background estimation. Initial studies in Figure 9.1 indicate that lower p_T jets are more likely to have higher BDT scores, while the η distribution is relatively flat.

Additionally, the mistag rate will be measured in other samples, such as Z+jets and QCD multijet, as a cross-check. This will extend the current W+jets based method and provide additional statistics to evaluate the mistag probabilities, as well as account for the contributions of other backgrounds in the triggered sample. The selections on 2023 data will also be compared with Monte Carlo (for example, W+jets MC will be compared with W+jets events selected in 2023 data) to validate the data and MC agreement. This will also provide an indirect validation of the LLP signal MC.

9.2 Systematic Uncertainties

A description of the expected approach to systematic uncertainties in the analysis is presented. The main uncertainties are expected to be attributed to standard jet energy uncertainties and specific uncertainties on the custom HB reconstructed hit variables and HCAL-based triggers.

9.2.1 Jet Energy Corrections

After the proton-proton collision, the quarks and gluons produced hadronize into stable particles that form jets. However, the response of the CMS detector (and each sub-detector!) is not linear in the jet's energy. Thus, the problem of translating measured jet energy into the actual energy of particles is non-trivial. The calibration of the jet energy measurement impacts the sensitivity of searches in CMS, and the Jet Energy Resolution and Correction (JER and JEC) group aims to keep the systematic uncertainty under 0.1%. JECs are effectively subsequently applied scale factors to the jets, addressing L1 pileup, MC-truth corrections, and residuals [50]. Jet energy resolution (JER) accounts for the fact that the resolution of jets in MC is higher than in data, requiring that MC jets are smeared to reflect data better.

In an analysis, both JES (jet energy scale corrections) and JER uncertainties are accounted for and are assumed to be uncorrelated. Current JER corrections are not dependent on p_T . Jet veto regions will also be accounted for. This is to accommodate regions of the calorimeters that record anomalous high or low-energy jets, and events with jets that fall within the veto region are typically rejected to prevent the inclusion of non-physical MET into the analysis [50].

9.2.2 Custom Variable Contributions

Custom uncertainties due to the triggers or custom variables may be needed in the analysis. A number of custom variables, such as HCAL-based energy fractions and $S_{\phi\phi}$ are key to the BDT, and

further work will determine their contributions to the overall systematic uncertainties.

9.3 Expected Sensitivities

Additional statistical tools and CMS Combine will be used to compute this analysis's expected sensitivities and limits. While studies are ongoing on optimizing the search strategy and discriminator choice, an initial sensitivity calculation is made before the full predictions are available to provide an estimation of the capabilities of this analysis. No systematic or statistical uncertainties are accounted for in the expected sensitivity calculations. We expect the analysis to be dominated by statistical uncertainties, as it is a low background search, and we will account for both this and systematic uncertainties (expected to be low). As the expectation is that statistical uncertainties limit the search, limit improvements will result from the increased data collected throughout Run-3 and further developments at the trigger level. In particular, the proposed expansion of the Level-1 LLP trigger to include the HCAL endcap will further improve LLP signal acceptance.

With two years of Run-3 data providing 67 fb^{-1} of data and the LLP acceptance of 0.09796% from Table 9.5 (using a double jet tag with a BDT threshold of 0.9999), and assuming there is very low background in the signal region, the Higgs to LLP branching ratio that would produce 3 observed events can be determined. The observation of these events (or lack of observation) would be used to define a discovery significance (or set limits on the $H \rightarrow \text{LLP}$ branching ratio).

$$\text{integrated luminosity} \cdot \text{LLP acceptance} \tag{9.9}$$

$$\cdot \text{Higgs production cross section} \cdot H \rightarrow \text{LLP branching ratio} = \text{observed events} \tag{9.10}$$

Inverting this and inputting the LLP acceptance number and Run-3 integrated luminosity gives:

$$\frac{\text{observed events}}{\text{integrated luminosity collected} \cdot \text{LLP acceptance}} = \sigma(pp \rightarrow H) \cdot \text{BR}(H \rightarrow \text{LLP}) \tag{9.11}$$

$$\approx \frac{3 \text{ events}}{67 \text{ fb}^{-1} \cdot 9.769 \cdot 10^{-4}} = 45.8 \text{ fb} \tag{9.12}$$

Thus, at $\sigma(pp \rightarrow H) \cdot \text{BR}(H \rightarrow \text{LLP})$ of 45.8 fb, 3 events would be expected to be observed in the double-tagged region. The trigger efficiency contribution is already accounted for due to the fact that HLT triggered events are used in the LLP acceptance calculation. With this strategy, the analysis could set limits on the Higgs to LLP branching ratio at $\text{BR} \leq 45.8 \text{ fb} / 48,520 \text{ fb} = 0.094\%$,

with a Higgs boson cross-section of 48.52 pb [51].

Expected sensitivities for each LLP sample evaluated are shown in Table 9.6 and Figure 9.3, with the full calculations shown in Appendix E.3. For all samples except $m_H = 125$, $m_S = 15$ GeV (the Standard Model Higgs mass decaying to two low mass LLP particles, which has reduced sensitivity due to limitations at the trigger acceptance level), we expect sensitivity to branching ratios under 1%.

LLP Sample Parameters	Expected $H \rightarrow$ LLP BR Sensitivity
$m_H = 125$ $m_S = 15$ $c\tau = 1$ m	13.5%
$m_H = 125$ $m_S = 50$ $c\tau = 3$ m	0.71%
$m_H = 250$ $m_S = 120$ $c\tau = 10$ m	0.56%
$m_H = 350$ $m_S = 80$ $c\tau = 0.5$ m	0.094%
$m_H = 350$ $m_S = 160$ $c\tau = 10$ m	0.26%

Table 9.6: Expected sensitivity to the branching ratio for each LLP sample in the analysis, based on a double jet tag with a BDT score of 0.9999. Note that systematic and statistical uncertainties are not yet accounted for in this estimation.




Figure 9.3: Expected limits for each LLP model evaluated in the analysis.

9.4 Summary

When requiring multiple jets scoring highly in the displaced hadronic jet BDT, the backgrounds are suppressed to above 10^7 . Additional background studies are underway and will account for the possibilities that the trigger selection includes significant other background contributions not

yet included in the estimation, or that the analysis selection is background enhanced (for example, selecting W+jets events that have signal-like characteristics).

The calorimeter-based LLP analysis demonstrates sensitivity to a wide range of $H \rightarrow SS \rightarrow b\bar{b}b\bar{b}$ mass and lifetime parameters, resulting in sensitivity to Higgs to LLP branching ratios of 10^{-3} . This discovery potential is significant, as the branching fraction for Higgs to invisible decays is currently constrained at $\text{BR}_{\text{inv}} < 0.16$ [23]. Thus, this search probes possible Higgs decays to LLPs, with sensitivity to branching ratios that are allowed by the current limits.

Chapter 10

Conclusions and Outlook

A novel timing trigger using the Hadron Calorimeter in the Compact Muon Solenoid (CMS) experiment at the Large Hadron Collider (LHC) is designed and deployed for Run-3 of the LHC, with data taking beginning in 2022. The timing trigger increases the experiment's sensitivity to long-lived particles (LLPs), by adding a low-level hardware trigger identifying jets arriving at delayed times or with unique depth signatures, both of which are characteristics of LLPs. The trigger relies on the Phase-1 Hadron Calorimeter upgrade, which provides greater segmentation granularity and an increased number of readout channels, allowing for timing information to be used in the hardware trigger system. This dissertation details the trigger algorithm, implementation, and performance, while focusing on the timing capabilities of the calorimeter and the use of these capabilities for higher precision detector alignment. The dedicated Level-1 and High Level Triggers take the event rate from 40 MHz to under 30 Hz (a reduction of almost 10^7 !) by identifying unique characteristics of LLP events.

A search is conducted for long-lived particles using events from the 2022 and 2023 Run-3 datasets from the CMS experiment that pass the HCAL-based LLP triggers. LLPs are predicted in many beyond the Standard Model theories, with long lifetimes due to small couplings, suppressed phase space, and scale suppression, and their unique signatures provide an excellent way of probing physics beyond the Standard Model. This dissertation presents a dedicated LLP search using Run-3 data, focusing on using the unique energy signatures and trackless jets characteristic of LLP events to reject backgrounds to the level of 10^7 , while maintaining signal efficiency for exotic Higgs models. The combination of the trigger and analysis selections provides a powerful background rejection

technique while maintaining expected signal sensitivity to branching ratios $< 1\%$.

Further work on the calorimeter-based LLP search is ongoing, with the full analysis expected to use the 2022 and 2023 Run-3 datasets. Future searches can benefit from the improved displaced jet HLTs implemented for 2024 and the TDC-based time alignment, as jet timing information is expected to be a key variable in a future search after validation in this first HCAL-based search. Expanding the trigger and analysis coverage to the HCAL endcap can also improve signal acceptance. A full Run-3 HCAL-based LLP search utilizing both the timing and depth capabilities from the implemented triggers will contribute to the calorimeter-based LLP coverage. In addition, future detectors such as the MIP timing detector and the High-Granularity Calorimeter (HGCAL, replacing the current calorimeter endcaps) in Phase-2 are exciting avenues to further explore the use of segmented depth and timing information in long-lived particle searches.

Appendix A

Run-3 Notable Runs and Significant Changes

Run-3 of the LHC started on July 5, 2022, with stable beams at 6.8 TeV, and a center of mass collision energy of 13.6 TeV. The L1 trigger menu included the single jet LLP triggers at the start of running, followed by the HLT triggers seeded with these L1 triggers. A few issues were present and debugged during 2022 and start of 2023 data taking, as detailed below.

The TDC BX positioning was off by one BX at the start of Run-3, placing the recorded TDC one sample too early. uHTR firmware 2.7.1 fixed this BX offset, and was installed on July 15, 2022. uHTR firmware 2.7.2 was installed on August 3, 2022, to fix TP BX pipeline offset issue.

November 23-26, 2022, Runs 362597 to 362696, L1 had a “mistaken” LUT, which was the requested 6:1 LUT. The rest of the 2022 runs were with a LUT just passing the first bit (depth only), while this was fixed for 2023 runs where the 6:1 performed the full depth OR timing. The 2023 6:1 LUT was verified with pattern testing from the uHTRs.

There was a $\pm 2i\phi$ swap where the fine grain bits were reported in a different tower than the ADC and TDC measurements used to compute them. This swap created a pattern where finegrain bits were reported two towers offset of expected, in groups of $2i\phi$. This has minimal impact on the flagged LLP jets, as the jets use 9×9 towers. uHTR firmware 2.7.4 fixes the $\pm 2i\phi$ swaps but has an offset between the timeslice in which the TP digi and fine grain bits are reported. This was not an issue in 2022 data due to a “patch” fix. The offset fix is in firmware 2.7.5 (updated on April 19, 2023,

before 2023 13.6 TeV pp physics) and applies a delay to both the DAQ and trigger paths.

Important changes, with run numbers and data taking eras¹ are itemized:

- Aug 5, 2022: QIE phase scan in 13.6 TeV collisions, Run 356815 (Run2022C). This scan covered -2 ns to +8 ns in 2 ns steps, and was taken around the 2022 operating QIE phases.
- Oct 12: new TDC LUTs tested in local LED Run 360234
- Oct 18: TDC LUTs updated for collisions, and L1 seeds prescaled before this. Run 360794 (Run2022F) is a run used to analyze TDC LUT updates.
- Oct 25: L1T unprescaled highest threshold seeds, and OR of all 5 L1 seeds are added to HLT paths (Run2022F).
- Nov 15: QIE phase scan in 13.6 TeV collisions, Runs 362085 and 362087 (Run2022F), with 2.5 hours of data taking. This scan covered -4 ns to 10 ns in 2 ns steps, and was taken around the 2023 recommended QIE phases.
- Nov 16: HB QIE phases changed to $i\phi$ dependent phases, with CWT (not TDC alignment). Run 362102 (end of Run2022F, change in conditions and response corrections as well).
- Nov 23-26: Runs 362597-362696 the L1 6:1 LUT was correct, and passed both timing and depth information. All other runs during 2022 only had depth information passed to the trigger due to a mistake in the LUT.
- Feb 21, 2023: QIE scan in a local LED run for CWT validation, local Run 363663.
- April 2: Monitoring path HLT_L1SingleLLPJet added to the HLT v1.0 menu and CMSSW releases.
- May 1: QIE phase scan in 900 GeV collisions, Runs 366873, 366874, 366876, 366891 - 366895 (Run2023B), with 18 hours of data taking. This scan covered -10 ns to +20 ns, in 1 ns steps, and is around the 2023 recommended QIE phases, that are TDC-based.
- May 17: v1.1 HLT menu deployed online after luminosity reached $10^{34} \text{ cm}^{-2} \text{ s}^{-2}$. This disabled the path HLT_HT200_L1SingleLLPJet_DisplacedDijet35_Inclusive1PtrkShortSig5, and was deployed in the middle of Run2023Cv2 at Run 367661.
- June 12: Trigger key changed to Run3May2023_v2.0, HB QIE phase delays to phaseTuning_HB_2022-04-27, and HB TDC LUTs to /nfshome0/gkopp/QIEscan/HBTDCLUT_prompt_delayed_

¹https://twiki.cern.ch/twiki/bin/viewauth/CMS/PdmVRun3Analysis#2023_Era_definition

9Nov2022_Flat12_P5.xml, in Run 368775 (Run2023C4). The first stable beams after the update are in Run 368822 (end of 2023 Cv4). This is the full update of HB QIE phases to the recommended 2023 ones based on the TDC analysis. HCAL gains and response corrections are also updated.

Appendix B

HCAL QIE Phases and TDC-LUTs

B.1 TDC LUTs

Tables B.1, B.2, and B.3 document the three TDC LUTs used during 2022 and 2023 datataking. LUTs were adjusted to compensate for the timing mis-alignment, and again once the TDC-based HCAL QIE phases were deployed. Table B.4 is the TDC LUT used for 2024 datataking, and is the only LUT where the slightly delayed window is larger than 1 ns.

Table B.5 documents the available delayed range in 2022 and early 2023 (before the HCAL TDC-based time alignment). This is found by determining where the prompt peak is currently (how far from SOI-1 to SOI boundary) and how much of the remaining 25 ns range is available to flag the hit as delayed.

B.2 Timing Adjustments

Table B.6 are the QIE adjustments implemented in June 2023, in half ns steps. After this, the HCAL is aligned to within 0.5 ns.

$ i\eta $	Depth 1	Depth 2	Depth 3	Depth 4
1	8	14	15	17
2	8	14	15	17
3	8	14	14	17
4	8	14	14	17
5	8	13	14	16
6	8	13	14	16
7	8	12	14	15
8	8	12	14	15
9	7	12	13	15
10	7	12	13	15
11	7	12	13	15
12	7	12	13	15
13	7	11	12	14
14	7	11	12	14
15	7	11	12	14
16	7	11	12	-

Table B.1: v1 TDC LUT: Boundary between prompt and first delayed bin (00 and 01) in HB, with the last TDC value in the prompt bin listed. This encoding is used for the entire HB, and the slightly delayed window is 1 ns wide. Early 2022 HB TDC LUTs.

$ i\eta $	Depth 1	Depth 2	Depth 3	Depth 4
1	22	14	7	7
2	22	14	7	7
3	22	12	6	5
4	22	12	8	5
5	22	9	6	5
6	22	9	6	4
7	22	8	6	1
8	20	8	4	1
9	19	8	1	-1
10	15	4	-1	-1
11	15	4	-1	-1
12	15	4	-1	-1
13	11	1	-1	-1
14	11	-1	-1	-1
15	7	-1	-1	-1
16	11	-1	-1	-

Table B.2: v2 TDC LUT: Based on August 2022 QIE phase scan, the LUT is adjusted to account for the alignment differences across HB. The boundary between prompt and first delayed bin (00 and 01) in HB, with the last TDC value in the prompt bin listed. This encoding is used for the entire HB. Late 2022 HB TDC LUTs, used to compensate for the varying time alignment of the detector. -1 means the prompt range has been completely removed, and the first 1 ns of the range will be the slightly delayed window.

B.3 Compensating Response Corrections

The compensating response correction factors deployed along with the June 2023 TDC-based HCAL alignment are documented in Table B.7. The corrections are generally less than 10%.

$ i\eta $	Depth 1	Depth 2	Depth 3	Depth 4
$1 - 16, t_p$	12	12	12	12
$1 - 16, t_d$	14	14	14	14

Table B.3: v3 TDC LUT: After the 2023 QIE phases for consistent time alignment across the HB are deployed, the TDC LUTs are updated such that the prompt range encompasses 0-6 ns everywhere.

$ i\eta $	Depth 1	Depth 2	Depth 3	Depth 4
$1 - 16, t_p$	18	12	12	12
$1 - 16, t_d$	24	18	18	18

Table B.4: v4 TDC LUT: For 2024 datataking, the TDC LUTs are updated such that the prompt range encompasses 0-9 ns in depth 1 to account for the wider spread of arrival times in depth 1. The slightly delayed window is widened for all depths equally.

$ i\eta $	Depth 1	Depth 2	Depth 3	Depth 4
1	0-17	0-23	0-25	2-27
2	0-17	0-23	1-26	2-27
3	0-17	0-24	1-26	3-28
4	0-17	0-24	1-26	4-29
5	0-17	0-24	2-27	4-29
6	0-17	0-25	2-27	4-29
7	0-17	0-25	4-29	4-29
8	0-19	0-25	4-29	5-30
9	0-19	1-26	4-29	6-31
10	0-19	2-27	4-29	8-33
11	0-20	3-28	6-31	8-33
12	0-20	3-28	6-31	8-33
13	0-20	3-28	7-32	8-33
14	0-20	5-30	8-33	8-33
15	0-21	5-30	8-33	8-33
16	0-21	5-30	8-33	-

Table B.5: Available delayed range (ns). Blue indicates that the range starts at time delay of 0, while green indicates the range starts at a non-zero value and extends beyond 25ns (due to the fact that the pulse arrives in SOI-1, but timing sensitivity is only in SOI).

$i\eta$	Depth 1	Depth 2	Depth 3	Depth 4
-16	-2	-10	-16	-
-15	2	-12	-19	-20
-14	2	-11	-16	-22
-13	3	-8	-15	-21
-12	6	-6	-13	-18
-11	6	-6	-10	-15
-10	7	-5	-9	-14
-9	9	-3	-8	-13
-8	11	0	-6	-9
-7	10	0	-5	-8
-6	12	1	-3	-6
-5	13	3	-2	-4
-4	14	3	0	-4
-3	14	3	-1	-3
-2	14	4	0	-2
-1	14	6	2	0
1	16	5	1	0
2	14	3	-1	-2
3	13	2	-2	-2
4	13	3	-1	-3
5	14	2	-3	-5
6	14	0	-4	-7
7	11	-1	-5	-9
8	12	-1	-6	-10
9	11	-3	-8	-12
10	9	-4	-10	-14
11	7	-6	-12	-17
12	8	-7	-13	-17
13	6	-9	-16	-21
14	4	-11	-17	-22
15	4	-12	-20	-20
16	-2	-10	-16	-

Table B.6: QIE phase adjustment from October 2022 values to TDC-based June 2023 values. Adjustments are in half ns steps, with positive values meaning that the pulse is moved earlier in SOI; negative values meaning that the pulse is moved later in SOI.

$ 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 B.7: 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.

Appendix C

Monte Carlo Simulation Generation

C.1 Monte Carlo for Trigger Development

For the trigger development, private production Monte Carlo is used, following the CMS Run-3 recommendations. MC production is done in CMSSW_11_2_0 or above, to include the updated HCAL TDC simulation, following a bug fix implemented in January 2021 to update the simulation to reflect the flat current threshold applied for the TDC measurement. The data tier GEN-SIM-DIGI-RAW is primarily used for trigger development. A few specific customization commands are used:

- `--customise_commands 'process.hcalRawDataauHTR.packHBTDC = cms.bool(False)'`
- `--conditions auto:phase1_2021_realistic`
- `--era Run3`
- `--pileup Run3_Flat55To75_Poisson00TPU`
- `--pileup_input das:/RelValMinBias_14TeV/CMSSW_11_2_0_pre8-112X_mcRun3_2021_realistic_v10-v1/GEN-SIM`

Table 5.3 lists the values of Higgs mass, LLP mass, and $c\tau$ used for trigger development and evaluation.

C.2 Monte Carlo for LLP Analysis

Exotic Higgs decays to two LLPs that decay to two b quarks are generated in CMSSW_13_0_13, with the TDC LUT used in 2023 data taking¹. The resulting Monte Carlo datasets are:

- dataset=/ggH_HToSSTobbbb_MH-125_MS-15_CTau1000_13p6TeV/lpclonglived-crab_PrivateProduction_Summer22_DR_step2_RECOSIM_ggH_HToSSTobbbb_MH-125_MS-15_CTau1000_13p6TeV_batch1_v1-59a22edf0600a784f6c900595d24e883/USER
- dataset=/HToSSTo4B_MH350_MS80_CTau500/lpclonglived-crab_PrivateProduction_Summer23BPix_DR_step2_RECOSIM_HToSSTo4B_MH350_MS80_CTau500_batch1_v1-6c03a81f0d97498cab5c296ab3fa9a76/USER/
- dataset=/HToSSTo4B_MH125_MS50_CTau3000/lpclonglived-crab_PrivateProduction_Summer23BPix_DR_step2_RECOSIM_HToSSTo4B_MH125_MS50_CTau3000_batch1_v1-6c03a81f0d97498cab5c296ab3fa9a76/USER
- dataset=/HToSSTo4B_MH125_MS50_CTau3000/lpclonglived-crab_PrivateProduction_Summer23BPix_DR_step2_RECOSIM_HToSSTo4B_MH125_MS50_CTau3000_batch2_v1-6c03a81f0d97498cab5c296ab3fa9a76/USER
- dataset=/HToSSTo4B_MH250_MS120_CTau10000/lpclonglived-crab_PrivateProduction_Summer23BPix_DR_step2_RECOSIM_HToSSTo4B_MH250_MS120_CTau10000_batch1_v1-6c03a81f0d97498cab5c296ab3fa9a76/USER
- dataset=/HToSSTo4B_MH250_MS120_CTau10000/lpclonglived-crab_PrivateProduction_Summer23BPix_DR_step2_RECOSIM_HToSSTo4B_MH250_MS120_CTau10000_batch2_v1-6c03a81f0d97498cab5c296ab3fa9a76/USER
- dataset=/HToSSTo4B_MH350_MS160_CTau10000/lpclonglived-crab_PrivateProduction_Summer23BPix_DR_step2_RECOSIM_HToSSTo4B_MH350_MS160_CTau10000_batch1_v1-6c03a81f0d97498cab5c296ab3fa9a76/USER
- dataset=/HToSSTo4B_MH350_MS160_CTau10000/lpclonglived-crab_PrivateProduction_Summer23BPix_DR_step2_RECOSIM_HToSSTo4B_MH350_MS160_CTau10000_batch2_v1-6c03a81f0d97498cab5c296ab3fa9a76/USER

¹<https://github.com/cms-sw/cms-sw/pull/42023/files>

C.2.1 Potential Future Extensions

Vector-like lepton (VLL) models are also promising for this type of search. VLL models pair-produce LLPs that decay to $H + \text{lepton}$ or $W + \text{lepton}$. Previous VLL searches focused on relatively prompt VLLs, where the sensitivity does not extend significantly beyond the tracker.

Appendix D

Trigger Efficiencies and Integrated Luminosity Gains

Rate and efficiency tables are presented for the L1 and HLT triggers developed, using CMSSW_12_3_0 just before the start of Run-3. Efficiencies are shown at various event H_T and jet energy thresholds, corresponding to the thresholds applied for the L1 triggers that seed these HLT paths. Comparisons are made to event H_T cuts of 360 and 420 GeV, as these were the lowest un-prescaled H_T cuts available at HLT in Run-2, and thus, what many displaced jet triggers use. By comparing events passed to those passing these thresholds, the integrated luminosity gains (efficiency gains) provided by adding the new L1 and HLT pats are computed, showing gains of 3-5 times, depending on the mass point.

Integrated luminosities with H_T and jet thresholds of 120/40 GeV (event H_T /jet E_T) are in Table D.1, 160/50 GeV in Table D.2, 200/60 GeV in Table D.3. Table D.4 shows efficiency gains split by each HLT path averaged over all signal samples, while Table D.5 is split up by each signal sample.

Results are shown for LLP MC samples listed in Table 5.3.

Sample	% w/LLP jet, $H_T=120$	% H_T	% Added Eff	Integrated Lumi Gain
$m_h = 125$ GeV, $c\tau = 3000$ m	7.75			
HTT _{comp} = 360		2.5	6.7	3.7
HTT _{comp} = 420		1.8	6.9	4.9
$m_h = 125$ GeV, $c\tau = 30000$ m	1.4			
HTT _{comp} = 360		0.95	1.4	2.5
HTT _{comp} = 420		0.55	1.4	3.5
$m_h = 250$ GeV, $c\tau = 500$ m	7.05			
HTT _{comp} = 360		13	4.8	1.4
HTT _{comp} = 420		9.3	5	1.5
$m_h = 250$ GeV, $c\tau = 1000$ m	12.2			
HTT _{comp} = 360		12	8.6	1.7
HTT _{comp} = 420		8.2	9.4	2.2
$m_h = 250$ GeV, $c\tau = 10000$ m	8.85			
HTT _{comp} = 360		5.9	6.7	2.1
HTT _{comp} = 420		4.5	7	2.5
$m_h = 350$ GeV, $c\tau = 500$ m	11.5			
HTT _{comp} = 360		25	6.7	1.3
HTT _{comp} = 420		16	8.1	1.5
$m_h = 350$ GeV, $c\tau = 1000$ m	19.6			
HTT _{comp} = 360		23	13	1.6
HTT _{comp} = 420		15	14	1.9
$m_h = 350$ GeV, $c\tau = 10000$ m	14.4			
HTT _{comp} = 360		9.9	11	2.1
HTT _{comp} = 420		7.4	12	2.6
$m_h = 1000$ GeV, $c\tau = 10000$ m	40.8			
HTT _{comp} = 360		46	12	1.3
HTT _{comp} = 420		39	15	1.4
$m_h = 1000$ GeV, $c\tau = 100000$ m	7.91			
HTT _{comp} = 360		18	3.7	1.2
HTT _{comp} = 420		15	4.5	1.3
Neutrino gun rates at $H_T=120$	23kHz total, 17kHz pure (unprescaled)			
Neutrino gun rates at $H_T=120$	167 Hz total, 23 Hz pure (prescale 50)			
Neutrino gun rates at $H_T=120$	82 Hz total, 7 Hz pure (prescale 100)			

Table D.1: Level 1 trigger signal efficiencies and integrated luminosity gain, April 2022 using hwQual flagged L1TNTuples (123X, private). $H_T = 120$ GeV, jet energy = 40 GeV, with no triggerability requirements (2k events each).

Sample	% w/LLP jet, $H_T=160$	% H_T	% Added Eff	Integrated Lumi Gain
$m_h = 125$ GeV, $c\tau = 3000$ m	5.15			
HTT _{comp} = 360		2.5	4.1	2.7
HTT _{comp} = 420		1.8	4.3	3.5
$m_h = 125$ GeV, $c\tau = 30000$ m	0.85			
HTT _{comp} = 360		0.95	0.85	1.9
HTT _{comp} = 420		0.55	0.85	2.5
$m_h = 250$ GeV, $c\tau = 500$ m	5.7			
HTT _{comp} = 360		13	3.5	1.3
HTT _{comp} = 420		9.3	3.7	1.4
$m_h = 250$ GeV, $c\tau = 1000$ m	9.7			
HTT _{comp} = 360		12	6.2	1.5
HTT _{comp} = 420		8.2	7	1.9
$m_h = 250$ GeV, $c\tau = 10000$ m	6.65			
HTT _{comp} = 360		5.9	4.6	1.8
HTT _{comp} = 420		4.5	5	2.1
$m_h = 350$ GeV, $c\tau = 500$ m	9.6			
HTT _{comp} = 360		25	5.2	1.2
HTT _{comp} = 420		16	6.5	1.4
$m_h = 350$ GeV, $c\tau = 1000$ m	17.3			
HTT _{comp} = 360		23	11	1.5
HTT _{comp} = 420		15	12	1.8
$m_h = 350$ GeV, $c\tau = 10000$ m	11.4			
HTT _{comp} = 360		9.9	8.2	1.8
HTT _{comp} = 420		7.4	8.9	2.2
$m_h = 1000$ GeV, $c\tau = 10000$ m	39.5			
HTT _{comp} = 360		46	11	1.2
HTT _{comp} = 420		39	14	1.4
$m_h = 1000$ GeV, $c\tau = 100000$ m	7.1			
HTT _{comp} = 360		18	3	1.2
HTT _{comp} = 420		15	3.8	1.3
Neutrino gun rates at $H_T=160$	kHz total, 0 pure (unprescaled)			
Neutrino gun rates at $H_T=160$	398 Hz total, 104 Hz pure (prescale 10)			
Neutrino gun rates at $H_T=160$	28 Hz total, 28 Hz pure (prescale 50)			

Table D.2: Level 1 trigger signal efficiencies and integrated luminosity gain, April 2022 using hwQual flagged L1TNTuples (123X, private). $H_T = 160$ GeV, jet energy = 50 GeV, with no triggerability requirements (2k events each).

Sample	% w/LLP jet, $H_T=200$	% H_T	% Added Eff	Integrated Lumi Gain
$m_h = 125$ GeV, $c\tau = 300$ m	3.35			
HTT _{comp} = 360		2.5	2.3	1.9
HTT _{comp} = 420		1.8	2.5	2.4
$m_h = 125$ GeV, $c\tau = 30000$ m	0.45			
HTT _{comp} = 360		0.95	0.45	1.5
HTT _{comp} = 420		0.55	0.45	1.8
Neutrino gun rates	0			
$m_h = 250$ GeV, $c\tau = 500$ m	4.25			
HTT _{comp} = 360		13	2.1	1.2
HTT _{comp} = 420		9.3	2.3	1.2
$m_h = 250$ GeV, $c\tau = 1000$ m	7.7			
HTT _{comp} = 360		12	4.2	1.4
HTT _{comp} = 420		8.2	5	1.6
$m_h = 250$ GeV, $c\tau = 10000$ m	4.75			
HTT _{comp} = 360		5.9	2.8	1.5
HTT _{comp} = 420		4.5	3.1	1.7
$m_h = 350$ GeV, $c\tau = 500$ m	8.2			
HTT _{comp} = 360		25	4	1.2
HTT _{comp} = 420		16	5.2	1.3
$m_h = 350$ GeV, $c\tau = 1000$ m	15.1			
HTT _{comp} = 360		23	8.4	1.4
HTT _{comp} = 420		15	9.9	1.7
$m_h = 350$ GeV, $c\tau = 1000$ m	8.45			
HTT _{comp} = 360		9.9	5.4	1.5
HTT _{comp} = 420		7.4	6	1.8
$m_h = 1000$ GeV, $c\tau = 10000$ m	37.5			
HTT _{comp} = 360		46	8.8	1.2
HTT _{comp} = 420		39	12	1.3
$m_h = 1000$ GeV, $c\tau = 100000$ m	6.55			
HTT _{comp} = 360		18	2.5	1.1
HTT _{comp} = 420		15	3.3	1.2
Neutrino gun rates at $H_T=200$	1.4 kHz total, 0.4 kHz pure (unprescaled)			

Table D.3: Level 1 trigger signal efficiencies and integrated luminosity gain, April 2022 using hwQual flagged L1TNTuples (123X, private). $H_T = 200$ GeV, jet energy = 60 GeV, with no triggerability requirements (2k events each).

HLT Path	Integrated luminosity gains (HLT)
HLT_HT200_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack_v1	1.46
HLT_HT270_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack_v1	1.56
HLT_HT200_L1SingleLLPJet_DisplacedDijet60_DisplacedTrack_v1	1.28
HLT_HT170_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack_v1	1.42
HLT_HT420_L1SingleLLPJet_DisplacedDijet60_Inclusive_v1	1.64
HLT_HT320_L1SingleLLPJet_DisplacedDijet60_Inclusive_v1	1.56
HLT_HT200_L1SingleLLPJet_DelayedJet40_DoubleDelay0p5nsTrackless_v1	1.41
HLT_HT200_L1SingleLLPJet_DelayedJet40_DoubleDelay1nsInclusive_v1	1.46
HLT_HT200_L1SingleLLPJet_DelayedJet40_SingleDelay1nsTrackless_v1	1.58
HLT_HT200_L1SingleLLPJet_DelayedJet40_SingleDelay2nsInclusive_v1	1.67
HLT_HT200_L1SingleLLPJet_DisplacedDijet30_Inclusive1PtrkShortSig5_v1	1.58
HLT_HT200_L1SingleLLPJet_DisplacedDijet35_Inclusive1PtrkShortSig5_v1	1.54
HLT_HT200_L1SingleLLPJet_DisplacedDijet40_Inclusive1PtrkShortSig5_v1	1.51
All paths, combined	1.56

Table D.4: HLT performance, using v3 config. Integrated luminosity is reported as $\frac{LLP \text{ L1 OR } H_T \text{ L1 seed}}{H_T \text{ L1 seed}}$ to show the factor in physics acceptance added by inclusion of the new HLT single LLP jet triggers. This is averaged over 9 signal samples.

	Integrated luminosity gains (HLT) by $m_H, c\tau$									
m_H (GeV)	125	250	250	250	350	350	350	1000	1000	
$c\tau$ (mm)	3000	500	1000	10000	500	1000	10000	10000	100000	Overall
HLT Path										
HLT_HT200_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack_v1	4.00	1.35	1.71	4.50	1.39	1.80	2.67	1.23	1.42	1.46
HLT_HT270_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack_v1	4.00	1.45	1.56	6.00	1.45	2.00	2.20	1.40	1.75	1.56
HLT_HT200_L1SingleLLPJet_DisplacedDijet60_DisplacedTrack_v1	nan	1.24	1.17	6.00	1.26	1.60	2.00	1.15	1.10	1.28
HLT_HT170_L1SingleLLPJet_DisplacedDijet40_DisplacedTrack_v1	4.00	1.39	1.79	3.00	1.31	1.65	2.38	1.21	1.38	1.42
HLT_HT420_L1SingleLLPJet_DisplacedDijet60_Inclusive_v1	nan	1.25	1.54	nan	1.35	1.45	1.89	1.80	3.00	1.64
HLT_HT320_L1SingleLLPJet_DisplacedDijet60_Inclusive_v1	6.00	1.31	1.41	16.00	1.36	1.60	2.57	1.54	2.11	1.56
HLT_HT200_L1SingleLLPJet_DelayedJet40_DoubleDelay0p5nsTrackless_v1	nan	1.40	1.82	2.67	1.36	1.71	2.67	1.24	1.54	1.41
HLT_HT200_L1SingleLLPJet_DelayedJet40_DoubleDelay1nsInclusive_v1	nan	1.71	2.83	nan	1.48	1.90	4.00	1.27	1.79	1.46
HLT_HT200_L1SingleLLPJet_DelayedJet40_SingleDelay1nsTrackless_v1	nan	1.43	2.44	7.50	1.51	1.93	2.70	1.33	1.76	1.58
HLT_HT200_L1SingleLLPJet_DelayedJet40_SingleDelay2nsInclusive_v1	nan	2.00	3.63	nan	1.84	2.26	4.33	1.37	1.95	1.67
HLT_HT200_L1SingleLLPJet_DisplacedDijet30_Inclusive1PtrkShortSig5_v1	2.78	1.28	1.67	3.27	1.45	1.84	2.59	1.39	1.63	1.58
HLT_HT200_L1SingleLLPJet_DisplacedDijet35_Inclusive1PtrkShortSig5_v1	3.60	1.25	1.63	3.23	1.41	1.82	2.49	1.34	1.59	1.54
HLT_HT200_L1SingleLLPJet_DisplacedDijet40_Inclusive1PtrkShortSig5_v1	4.67	1.27	1.59	3.08	1.39	1.78	2.53	1.32	1.53	1.51
All paths, combined	2.78	1.30	1.69	3.44	1.42	1.78	2.60	1.37	1.73	1.56

Table D.5: HLT performance, using v3 config. Integrated luminosity is reported as $\frac{LLP \text{ L1 OR } H_T \text{ L1 seed}}{H_T \text{ L1 seed}}$ to show the factor in physics acceptance added by inclusion of the new HLT single LLP jet triggers.

Appendix E

Analysis Cutflow Tables

E.1 Event Selection

Table E.1 shows the event selection cutflow similar to Table 8.1, with cuts on LLP decay position and HLT requirements. These are similar cuts to those imposed at the analysis level.

Selection	Number of Entries		Fraction of Entries	
	LLP 0	LLP 1	LLP 0	LLP 1
All	462918	462918	1.0	1.0
LLP r in HCAL depth 3 or 4	47636	48334	0.1029	0.1044
LLP $\eta \leq 1.26$	24866	25431	0.0537	0.0549
LLP $E > 60$ GeV	15442	15841	0.0334	0.0342
LLP 0 OR LLP 1	30554		0.066	
LLP 0 AND LLP 1	729		0.0016	
LLP OR, with Event HT > 200 GeV	4609		0.01	
LLP AND, with Event HT > 200 GeV	190		0.0004	
Leading jet neutral hadron frac $> 70\%$	109660		0.2369	
Leading jet neutral hadron frac $> 80\%$	93772		0.2026	
L1 passed (monitoring path)	0		0.0	
HLT dijet, ≤ 1 prompt track	885		0.0019	
HLT dijet, 1 displaced track	59		0.0001	
HLT dijet, ≤ 2 prompt track	167		0.0004	
HLT delayed jet, ECAL timing	86		0.0002	

Table E.1: Cutflow table, showing the number and fraction of events that pass the displaced LLP decay selections. The last portion of the table shows the fraction of events passing the L1 and HLT HCAL LLP triggers. Evaluated on $m_H = 125$, $m_S = 15$, $c\tau = 1$ m.

Tables E.2 and E.3 apply timing requirements to the LLP instead of decay position requirements.

This highlights the contribution of both the timing and depth trigger contributions.

LLP time is calculated as:

$$t = \frac{\sqrt{\Delta x_{LLP}^2 + \Delta y_{LLP}^2 + \Delta z_{LLP}^2}}{\beta_{LLP} \cdot c} \quad (\text{E.1})$$

This accounts for the time delay due to a slow-moving (slower than the speed of light) LLP, with $\beta = \frac{v}{c}$. In this approach, the geometric path length from the LLP travel is not accounted for but could contribute to an additional time delay.

Selection	Number of Entries		Fraction of Entries	
	LLP 0	LLP 1	LLP 0	LLP 1
All	465995	465995	1.0	1.0
LLP $t > 6$ ns and $r < 295$ cm	136770	137476	0.2935	0.295
LLP $\eta \leq 1.26$	50724	51020	0.1089	0.1095
LLP $E > 60$ GeV	50724	51020	0.1089	0.1095
LLP 0 OR LLP 1	95245		0.2044	
LLP 0 AND LLP 1	6499		0.0139	
LLP OR, with Event HT > 200 GeV	80567		0.1729	
LLP AND, with Event HT > 200 GeV	6372		0.0137	

Table E.2: Cutflow table, showing the number and fraction of events that pass the delayed (timing-based) LLP decay selections. Evaluated on $m_H = 350$, $m_S = 80$, $c\tau = 0.5$ m.

Selection	Number of Entries		Fraction of Entries	
	LLP 0	LLP 1	LLP 0	LLP 1
All	462918	462918	1.0	1.0
LLP $t > 6$ ns and $r < 295$ cm	160328	161666	0.3463	0.3492
LLP $\eta \leq 1.26$	55057	56065	0.1189	0.1211
LLP $E > 60$ GeV	33941	34635	0.0733	0.0748
LLP 0 OR LLP 1	65211		0.1409	
LLP 0 AND LLP 1	3365		0.0073	
LLP OR, with Event HT > 200 GeV	9381		0.0203	
LLP AND, with Event HT > 200 GeV	863		0.0019	

Table E.3: Cutflow table, showing the number and fraction of events that pass the delayed (timing-based) LLP decay selections. Evaluated on $m_H = 125$, $m_S = 15$, $c\tau = 1000$ mm.

The event selection cutflow applies simple selections on the trigger and jet requirements to the events from the 2023 displaced jet skim data and LLP MC to understand what fraction of events in the skim are cut by these initial requirements. The results for the data are in Table E.4, while those for signal are included in the tables in Section E.3.

Selection	Number of Entries	Fraction of Entries
All events (skim)	29976438	1.0
Jet ≥ 40 GeV p_T (one of 3 leading)	29811024	0.9945
Jet $ \eta \leq 1.26$	29398519	0.9807

Table E.4: Event selection cutflow, for 25% of the 2023 displaced jet skim. The jet p_T and η requirements only remove 2% of the dataset, a minimal cut compared with the later selections.

E.2 BDT Comparison

Jet-based cutflow results are in Tables E.6 through E.8 for four of the LLP signal points, with one additional signal point and the W+jets background in Tables 8.4 and 8.5, respectively. Results of the jet-based overlaid with the BDT results plotted in Figure E.1.

Selection	Number of Weighted Entries	Fraction of Weighted Entries	
		Of all entries	Of jet matched entries
All (jet $ \eta \leq 1.26$ and $p_T \geq 40$)	92777	1.0	
Jet matched to LLP in HCAL	35331	0.3808	
Neutral hadron fraction $> 80\%$	31955	0.3444	0.9044
Charged hadron fraction $< 10\%$	28311	0.3052	0.8013
$S_{\phi\phi} < 0.02$	28143	0.3033	0.7966
Track $ \Delta\eta > 0.05$	25615	0.2761	0.725
Track $ \Delta\phi > 0.05$	23436	0.2526	0.6633
Leading rehit energy > 8 GeV	23330	0.2515	0.6603

Table E.5: Per jet cutflow table to compare with BDT results for $m_H = 125$, $m_S = 15$, $c\tau = 1$ m.

Selection	Number of Weighted Entries	Fraction of Weighted Entries	
		Of all entries	Of jet matched entries
All (jet $ \eta \leq 1.26$ and $p_T \geq 40$)	453581	1.0	
Jet matched to LLP in HCAL	224507	0.495	
Neutral hadron fraction $> 80\%$	192880	0.4252	0.8591
Charged hadron fraction $< 10\%$	163660	0.3608	0.729
$S_{\phi\phi} < 0.02$	160239	0.3533	0.7137
Track $ \Delta\eta > 0.05$	145098	0.3199	0.6463
Track $ \Delta\phi > 0.05$	131868	0.2907	0.5874
Leading rehit energy > 8 GeV	130976	0.2888	0.5834

Table E.6: Per jet cutflow table to compare with BDT results for $m_H = 125$, $m_S = 50$, $c\tau = 3$ m.

Selection	Number of Weighted Entries	Fraction of Weighted Entries	
		Of all entries	Of jet matched entries
All (jet $ \eta \leq 1.26$ and $p_T \geq 40$)	488860	1.0	
Jet matched to LLP in HCAL	127385	0.2606	
Neutral hadron fraction $> 80\%$	115191	0.2356	0.9043
Charged hadron fraction $< 10\%$	105690	0.2162	0.8297
$S_{\phi\phi} < 0.02$	100942	0.2065	0.7924
Track $ \Delta\eta > 0.05$	91386	0.1869	0.7174
Track $ \Delta\phi > 0.05$	82930	0.1696	0.651
Leading rehit energy > 8 GeV	82563	0.1689	0.6481

Table E.7: Per jet cutflow table to compare with BDT results for $m_H = 250$, $m_S = 120$, $c\tau = 10$ m.

E.2.1 Additional ROC Curves

Six BDTs are developed and trained on each signal and a combination of signals. These six classifiers are then tested on each of the five signals. Figure E.1 shows ROC curves for each signal point. The best performing classifiers are trained and tested on the same signal sample. Overall, $m_H = 350$, $m_S = 80$ GeV (trained and tested on this sample) has the best signal and background separation.

Selection	Number of Weighted Entries	Fraction of Weighted Entries	
		Of all entries	Of jet matched entries
All (jet $ \eta \leq 1.26$ and $p_T \geq 40$)	865800	1.0	
Jet matched to LLP in HCAL	204134	0.2358	
Neutral hadron fraction $> 80\%$	190636	0.2202	0.9339
Charged hadron fraction $< 10\%$	177614	0.2051	0.8701
$S_{\phi\phi} < 0.02$	169507	0.1958	0.8304
Track $ \Delta\eta > 0.05$	153092	0.1768	0.75
Track $ \Delta\phi > 0.05$	138971	0.1605	0.6808
Leading rehit energy > 8 GeV	138507	0.16	0.6785

Table E.8: Per jet cutflow table to compare with BDT results for $m_H = 350$, $m_S = 160$, $c\tau = 10$ m.

Note that the combination (grey) separates the signal and background well for other samples. In all cases, the BDT performance surpasses the cut-based approach.

E.3 0, 1, and 2+ Tagged Jet Regions

E.3.1 BDT Discriminator Threshold 0.9999

Tables E.9 through E.12 show the fraction of each signal sample passing single and double tagged jet BDT scores above 0.9999, with and without the HLT requirement. The fifth signal sample is in Table 9.5.

E.3.2 Expected Sensitivities

Using Equation E.2, the expected sensitivities for each mass point can be determined. The LLP acceptance from Tables E.9 through E.12 is used. The resulting cross-section times branching fraction is compared to the Higgs boson cross-section of 48.5 pb to determine the limits expected on each exotic Higgs decay mass point with the double jet tag analysis approach.


$$\frac{\text{observed events}}{\text{integrated luminosity collected} \cdot \text{LLP acceptance}} = \sigma(pp \rightarrow H) \cdot BR(H \rightarrow \text{LLP}) \quad (\text{E.2})$$

$m_H = 125, m_S = 15$ GeV, $c\tau = 1$ m:

$$\frac{3 \text{ events}}{67 \text{ fb}^{-1} \cdot 6.84 \cdot 10^{-6}} = 6546 \text{ fb} \quad BR \leq 13.5\% \quad (\text{E.3})$$


$m_H = 125, m_S = 50$ GeV, $c\tau = 3$ m:

$$\frac{3 \text{ events}}{67 \text{ fb}^{-1} \cdot 1.30 \cdot 10^{-4}} = 344 \text{ fb} \quad BR \leq 0.71\% \quad (\text{E.4})$$




(a) ROC curve for $m_H = 125, m_S = 15$ GeV, and $c\tau = 1$ m.

(b) ROC curve for $m_H = 125, m_S = 50$ GeV, and $c\tau = 3$ m.



(c) ROC curve for $m_H = 250, m_S = 120$ GeV, and $c\tau = 10$ m.

(d) ROC curve for $m_H = 350, m_S = 80$ GeV, and $c\tau = 0.5$ m. The black point is at a BDT discriminator cut of 0.9995.



(e) ROC curve for $m_H = 350, m_S = 160$ GeV, and $c\tau = 10$ m.

Figure E.1: ROC curves, with reference points shown for comparison to a cut-based approach (blue), as determined from Tables E.6 through E.8 and 8.4 for the signal, and Table 8.5 for the background.

Selection	Number of Entries	Fraction of Entries	
		All	Valid Jet
All	292437	1.0	
Jet ≥ 40 GeV p_T and $ \eta \leq 1.26$	121612	0.4159	1.0
0 jets with BDT score ≥ 0.9999	114749	0.3924	0.9436
1 jet with BDT score ≥ 0.9999	6810	0.0233	0.056
2+ jets with BDT scores ≥ 0.9999	53	0.0002	0.0004
LLP HLT passed	577	0.002	0.0047
HLT and 0 jets with BDT score ≥ 0.9999	492	0.0017	0.004
HLT and 1 jet with BDT score ≥ 0.9999	83	0.0003	0.0007
HLT and 2+ jets with BDT scores ≥ 0.9999	2	$6.84 \cdot 10^{-6}$	0.0
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, triggered	14092	0.0482	1.0
LLP HLT passed	539	0.0018	0.0382
HLT and 0 jets with BDT score ≥ 0.9999	463	0.0016	0.0329
HLT and 1 jet with BDT score ≥ 0.9999	76	0.0003	0.0054
HLT and 2+ jets with BDT scores ≥ 0.9999	0	0.0	0.0
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, not triggered	103383	0.3535	1.0
LLP HLT passed	463	0.0016	0.0045
HLT and 0 jets with BDT score ≥ 0.9999	452	0.0015	0.0044
HLT and 1 jet with BDT score ≥ 0.9999	11	$3.761 \cdot 10^{-5}$	0.0001
HLT and 2+ jets with BDT scores ≥ 0.9999	0	0.0	0.0

Table E.9: Fraction of signal events passing jet kinematic selections and split in each of the three jet tagging bins (no jets tagged, 1 tagged, 2+ tagged), with a jet p_T threshold of 40 GeV and a BDT score of 0.9999. The fraction that also passes any of the relevant HCAL-based HLTs is also shown, with a very low fraction passing both the HLT and the BDT cut. From $m_H = 125$, $m_S = 15$ GeV, and $c\tau = 1$ m.

$m_H = 250, m_S = 120$ GeV, $c\tau = 10$ m:

$$\frac{3 \text{ events}}{67 \text{ fb}^{-1} \cdot 1.66 \cdot 10^{-4}} = 270 \text{ fb} \quad BR \leq 0.56\% \quad (\text{E.5})$$

$m_H = 350, m_S = 80$ GeV, $c\tau = 0.5$ m:

$$\frac{3 \text{ events}}{67 \text{ fb}^{-1} \cdot 9.77 \cdot 10^{-4}} = 45.8 \text{ fb} \quad BR \leq 0.094\% \quad (\text{E.6})$$

$m_H = 350, m_S = 160$ GeV, $c\tau = 10$ m:

$$\frac{3 \text{ events}}{67 \text{ fb}^{-1} \cdot 3.54 \cdot 10^{-4}} = 126 \text{ fb} \quad BR \leq 0.26\% \quad (\text{E.7})$$

Except for $m_H = 125, m_S = 15$ GeV with a $c\tau = 1$ m, all samples have an expected sensitivity to a Higgs to LLP branching ratio of less than 1%. For other LLP samples, about 75-90% of the events passing the double jet tag also pass the HLTs, while this is about 4% for this particular sample. This HLT inefficiency impacts the final result, leading to lower sensitivity for the 125 GeV Higgs

Selection	Number of Entries	Fraction of Entries	
		All	Valid Jet
All	2943719	1.0	
Jet ≥ 40 GeV p_T and $ \eta \leq 1.26$	930601	0.3161	1.0
0 jets with BDT score ≥ 0.9999	884022	0.3003	0.9499
1 jet with BDT score ≥ 0.9999	46075	0.0157	0.0495
2+ jets with BDT scores ≥ 0.9999	504	0.0002	0.0005
LLP HLT passed	73689	0.025	0.0792
HLT and 0 jets with BDT score ≥ 0.9999	55477	0.0188	0.0596
HLT and 1 jet with BDT score ≥ 0.9999	17830	0.0061	0.0192
HLT and 2+ jets with BDT scores ≥ 0.9999	382	$1.30 \cdot 10^{-4}$	0.0004
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, triggered	115963	0.0394	1.0
LLP HLT passed	68488	0.0233	0.5906
HLT and 0 jets with BDT score ≥ 0.9999	50951	0.0173	0.4394
HLT and 1 jet with BDT score ≥ 0.9999	17392	0.0059	0.15
HLT and 2+ jets with BDT scores ≥ 0.9999	145	0.0	0.0013
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, not triggered	803135	0.2728	1.0
LLP HLT passed	42277	0.0144	0.0526
HLT and 0 jets with BDT score ≥ 0.9999	41481	0.0141	0.0516
HLT and 1 jet with BDT score ≥ 0.9999	795	0.0003	0.001
HLT and 2+ jets with BDT scores ≥ 0.9999	1	0.0	0.0

Table E.10: Fraction of signal events passing jet kinematic selections and split in each of the three jet tagging bins (no jets tagged, 1 tagged, 2+ tagged), with a jet p_T threshold of 40 GeV and a BDT score of 0.9999. The fraction that also passes any of the relevant HCAL-based HLTs is also shown, indicating that the events with two jets scoring highly in the BDT also pass the LLP HLTs 76% of the time. From $m_H = 125$, $m_S = 50$ GeV, and $c\tau = 3$ m.

decaying into a low mass LLP.

E.3.3 BDT Discriminator Threshold 0.9995

Table E.13 shows the W+jets background estimation with a BDT cut of 0.9995.

From Table E.13, take 0.0004 as the probability that an event containing at least one triggered jet over 40 GeV also scores over 0.9995 in the BDT discriminator. The number of events predicted to have one tag in 2023 data (in a control region, not signal) is shown in Table E.14. The estimation of the number of events in the double jet tagged region is estimated as $p_{2 \text{ tagged jets}} = p_{1 \text{ tagged jet}}^2$, making the assumption that the probability of mistagging jets within the same event is independent of each other.

Tables E.15 through E.19 are cutflow tables showing the fraction of events in each signal sample that pass the basic event selection cuts (jet p_T and η) and subsequently fall in the 0, 1, or 2+ jet tagged bins, with and without the HLT requirements. The corresponding table on the W+jets background is Table E.13. For $m_H = 350$, $m_S = 80$ GeV (Table E.18), 70% of the events selected by the HLTs

Selection	Number of Entries	Fraction of Entries	
		All	Valid Jet
All	2937728	1.0	
Jet ≥ 40 GeV p_T and $ \eta \leq 1.26$	1122827	0.3822	1.0
0 jets with BDT score ≥ 0.9999	1078214	0.367	0.9603
1 jet with BDT score ≥ 0.9999	44061	0.015	0.0392
2+ jets with BDT scores ≥ 0.9999	552	0.0002	0.0005
LLP HLT passed	77847	0.0265	0.0693
HLT and 0 jets with BDT score ≥ 0.9999	51229	0.0174	0.0456
HLT and 1 jet with BDT score ≥ 0.9999	26131	0.0089	0.0233
HLT and 2+ jets with BDT scores ≥ 0.9999	487	$1.66 \cdot 10^{-4}$	0.0004
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, triggered	90775	0.0309	1.0
LLP HLT passed	73529	0.025	0.81
HLT and 0 jets with BDT score ≥ 0.9999	47341	0.0161	0.5215
HLT and 1 jet with BDT score ≥ 0.9999	25940	0.0088	0.2858
HLT and 2+ jets with BDT scores ≥ 0.9999	248	0.0001	0.0027
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, not triggered	1018830	0.3468	1.0
LLP HLT passed	47193	0.0161	0.0463
HLT and 0 jets with BDT score ≥ 0.9999	46761	0.0159	0.0459
HLT and 1 jet with BDT score ≥ 0.9999	432	0.0001	0.0004
HLT and 2+ jets with BDT scores ≥ 0.9999	0	0.0	0.0

Table E.11: Fraction of signal events passing jet kinematic selections and split in each of the three jet tagging bins (no jets tagged, 1 tagged, 2+ tagged), with a jet p_T threshold of 40 GeV and a BDT score of 0.9999. The fraction that also passes any of the relevant HCAL-based HLTs is also shown, indicating that the events with two jets scoring highly in the BDT also pass the LLP HLTs 88% of the time. From $m_H = 250$, $m_S = 120$ GeV, and $c\tau = 10$ m.

have at least one jet with a BDT score over 0.9995. Additionally, of all events selected by the 2+ jet tagged requirement, 80% of these pass the LLP HLTs, indicating that the triggered events perform well in the analysis.

An expected sensitivity calculation is performed for the lower BDT threshold of 0.9995 on the $m_H = 350$, $m_S = 80$ GeV signal, following the approach used in Section 9.3. At the lower BDT threshold, the LLP acceptance is 0.24% from Table E.18 (using a double jet tag with a BDT threshold of 0.9995), resulting in:

$$\frac{\text{observed events}}{\text{integrated luminosity collected} \cdot \text{LLP acceptance}} = \sigma(pp \rightarrow H) \cdot BR(H \rightarrow \text{LLP}) \quad (\text{E.8})$$

$$\approx \frac{3 \text{ events}}{67 \text{ fb}^{-1} \cdot 0.0024} = 18.7 \text{ fb} \quad (\text{E.9})$$

With this strategy, the analysis could set limits on the Higgs to LLP branching ratio at $BR \leq \frac{18.7 \text{ fb}}{48,520 \text{ fb}} = 0.038\%$, with a Higgs boson production cross-section of 48.52 pb [51].

Selection	Number of Entries	Fraction of Entries	
		All	Valid Jet
All	2946400	1.0	
Jet ≥ 40 GeV p_T and $ \eta \leq 1.26$	1462810	0.4965	1.0
0 jets with BDT score ≥ 0.9999	1379648	0.4682	0.9431
1 jet with BDT score ≥ 0.9999	81969	0.0278	0.056
2+ jets with BDT scores ≥ 0.9999	1193	0.0004	0.0008
LLP HLT passed	141103	0.0479	0.0965
HLT and 0 jets with BDT score ≥ 0.9999	87911	0.0298	0.0601
HLT and 1 jet with BDT score ≥ 0.9999	52149	0.0177	0.0356
HLT and 2+ jets with BDT scores ≥ 0.9999	1043	$3.54 \cdot 10^{-4}$	0.0007
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, triggered	160464	0.0545	1.0
LLP HLT passed	133596	0.0453	0.8326
HLT and 0 jets with BDT score ≥ 0.9999	81361	0.0276	0.507
HLT and 1 jet with BDT score ≥ 0.9999	51627	0.0175	0.3217
HLT and 2+ jets with BDT scores ≥ 0.9999	608	0.0002	0.0038
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, not triggered	1316911	0.447	1.0
LLP HLT passed	84637	0.0287	0.0643
HLT and 0 jets with BDT score ≥ 0.9999	83668	0.0284	0.0635
HLT and 1 jet with BDT score ≥ 0.9999	968	0.0003	0.0007
HLT and 2+ jets with BDT scores ≥ 0.9999	1	0.0	0.0

Table E.12: Fraction of signal events passing jet kinematic selections and split in each of the three jet tagging bins (no jets tagged, 1 tagged, 2+ tagged), with a jet p_T threshold of 40 GeV and a BDT score of 0.9999. The fraction that also passes any of the relevant HCAL-based HLTs is also shown, indicating that the events with two jets scoring highly in the BDT also pass the LLP HLTs 87% of the time. From $m_H = 350$, $m_S = 160$ GeV, and $c\tau = 10$ m.

Selection	Number of Entries	Fraction of Entries	
		All	Valid Jet
All events (from skim)	48574	1.0	
Jet ≥ 40 GeV p_T (one of 6 leading)	48354	0.9955	
0 jets with BDT score ≥ 0.9995	48335	0.9951	0.9996
1 jet with BDT score ≥ 0.9995	19	0.0004	0.0004
2+ jets with BDT scores ≥ 0.9995	0	0.0	0.0
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, triggered	21907	0.451	1.0
0 jets with BDT score ≥ 0.9995	21898	0.4508	0.9996
1 jet with BDT score ≥ 0.9995	9	0.0002	0.0004
2+ jets with BDT scores ≥ 0.9995	0	0.0	0.0
Jet ≥ 40 GeV p_T , $ \eta \leq 1.26$, not triggered	20484	0.4217	1.0
0 jets with BDT score ≥ 0.9995	20482	0.4217	0.9999
1 jet with BDT score ≥ 0.9995	2	0.0	0.0001
2+ jets with BDT scores ≥ 0.9995	0	0.0	0.0

Table E.13: W+jets events cutflow, split by trigger matched and unmatched jets that score highly in the BDT. A BDT score cut of 0.9995 is imposed. The mistag probability is higher for the trigger-matched jets, as expected.

Selection	Number of Entries	Fraction of Entries
All events (2023 LLP skim)	82312351	1.0
Jet ≥ 40 GeV p_T (one of 6 valid jets)	80358669	0.9763
One tagged jet prediction ($p = 0.0004$)	32143	0.0004
Two tagged jet prediction ($p = 0.0004$)	13	0.0004^2
Jet ≥ 60 GeV p_T (one of 6 valid jets)	72985190	0.8867
Jet $ \eta \leq 1.26$	72985190	0.8867
One tagged jet prediction ($p = 0.0003$)	21895	0.0003
Two tagged jet prediction ($p = 0.0003$)	7	0.0003^2

Table E.14: Predicted backgrounds for the single tag and double tag regions in the full 2023 skim dataset. Predictions are made from the fraction of W+jets events with a triggered jet falling in the single tag region, from Table 9.2 (and a similar one with a jet cut of 60 GeV) with a BDT cut of 0.9995.

Selection	Number of Entries	Fraction of Entries
All	292437	1.0
Jet ≥ 40 GeV p_T and $ \eta \leq 1.26$	121612	0.4159
0 jets with BDT score ≥ 0.9995	105895	0.3621
1 jet with BDT score ≥ 0.9995	15391	0.0526
2+ jets with BDT scores ≥ 0.9995	326	0.0011
LLP HLT passed	577	0.002
HLT and 0 jets with BDT score ≥ 0.9995	314	0.0011
HLT and 1 jet with BDT score ≥ 0.9995	248	0.0008
HLT and 2+ jets with BDT scores ≥ 0.9995	15	0.0001

Table E.15: Fraction of signal events passing jet kinematic selections and split in each of the three jet tagging bins (no jets tagged, 1 tagged, 2+ tagged), with a jet p_T threshold of 40 GeV and a BDT score of 0.9995. The fraction that also passes any of the relevant HCAL-based HLTs is also shown, with a very low fraction both passing the BDT and HLT. From $m_H = 125$, $m_S = 15$ GeV, and $c\tau = 1$ m.

Selection	Number of Entries	Fraction of Entries
All	2943719	1.0
Jet ≥ 40 GeV p_T and $ \eta \leq 1.26$	930601	0.3161
0 jets with BDT score ≥ 0.9995	837003	0.2843
1 jet with BDT score ≥ 0.9995	91975	0.0312
2+ jets with BDT scores ≥ 0.9995	1623	0.0006
LLP HLT passed	73689	0.025
HLT and 0 jets with BDT score ≥ 0.9995	38842	0.0132
HLT and 1 jet with BDT score ≥ 0.9995	33631	0.0114
HLT and 2+ jets with BDT scores ≥ 0.9995	1216	0.0004

Table E.16: Fraction of signal events passing jet kinematic selections and split in each of the three jet tagging bins (no jets tagged, 1 tagged, 2+ tagged), with a jet p_T threshold of 40 GeV and a BDT score of 0.9995. The fraction that also passes any of the relevant HCAL-based HLTs is also shown, indicating that the events with two jets scoring highly in the BDT also pass the LLP HLTs 75% of the time. From $m_H = 125$, $m_S = 50$ GeV, and $c\tau = 3$ m.

Selection	Number of Entries	Fraction of Entries
All	2937728	1.0
Jet ≥ 40 GeV p_T and $ \eta \leq 1.26$	1122827	0.3822
0 jets with BDT score ≥ 0.9995	1045330	0.3558
1 jet with BDT score ≥ 0.9995	76030	0.0259
2+ jets with BDT scores ≥ 0.9995	1467	0.0005
LLP HLT passed	77847	0.0265
HLT and 0 jets with BDT score ≥ 0.9995	35013	0.0119
HLT and 1 jet with BDT score ≥ 0.9995	41615	0.0142
HLT and 2+ jets with BDT scores ≥ 0.9995	1219	0.0004

Table E.17: Fraction of signal events passing jet kinematic selections and split in each of the three jet tagging bins (no jets tagged, 1 tagged, 2+ tagged), with a jet p_T threshold of 40 GeV and a BDT score of 0.9995. The fraction that also passes any of the relevant HCAL-based HLTs is also shown, indicating that the events with two jets scoring highly in the BDT also pass the LLP HLTs 83% of the time. From $m_H = 250$, $m_S = 120$ GeV, and $c\tau = 10$ m.

Selection	Number of Entries	Fraction of Entries
All	299110	1.0
Jet ≥ 40 GeV p_T and $ \eta \leq 1.26$	249804	0.8352
0 jets with BDT score ≥ 0.9995	221904	0.7419
1 jet with BDT score ≥ 0.9995	26984	0.0902
2+ jets with BDT scores ≥ 0.9995	916	0.0031
LLP HLT passed	21982	0.0735
HLT and 0 jets with BDT score ≥ 0.9995	6351	0.0212
HLT and 1 jet with BDT score ≥ 0.9995	14901	0.0498
HLT and 2+ jets with BDT scores ≥ 0.9995	730	0.0024

Table E.18: Fraction of signal events passing jet kinematic selections and split in each of the three jet tagging bins (no jets tagged, 1 tagged, 2+ tagged), with a jet p_T threshold of 40 GeV and a BDT score of 0.9995. The fraction that also passes any of the relevant HCAL-based HLTs is also shown, indicating that the events with two jets scoring highly in the BDT also pass the LLP HLTs 80% of the time. From $m_H = 350$, $m_S = 80$ GeV, and $c\tau = 0.5$ m.

Selection	Number of Entries	Fraction of Entries
All	2946400	1.0
Jet ≥ 40 GeV p_T and $ \eta \leq 1.26$	1462810	0.4965
0 jets with BDT score ≥ 0.9995	1323857	0.4493
1 jet with BDT score ≥ 0.9995	135553	0.046
2+ jets with BDT scores ≥ 0.9995	3400	0.0012
LLP HLT passed	141103	0.0479
HLT and 0 jets with BDT score ≥ 0.9995	59587	0.0202
HLT and 1 jet with BDT score ≥ 0.9995	78702	0.0267
HLT and 2+ jets with BDT scores ≥ 0.9995	2814	0.001

Table E.19: Fraction of signal events passing jet kinematic selections and split in each of the three jet tagging bins (no jets tagged, 1 tagged, 2+ tagged), with a jet p_T threshold of 40 GeV and a BDT score of 0.9995. The fraction that also passes any of the relevant HCAL-based HLTs is also shown, indicating that the events with two jets scoring highly in the BDT also pass the LLP HLTs 83% of the time. From $m_H = 350$, $m_S = 160$ GeV, and $c\tau = 10$ m.

Appendix F

MIP Timing Detector and Future LLP Searches

The MIP Timing Detector (MTD) is a detector to measure minimum ionizing particles (MIPs) and will be added to CMS during the Phase-2 upgrade for the High Luminosity LHC [14]. The MTD is designed to measure the time of particles, which is incredibly beneficial in a high pileup scenario, such as the expected pileup of 200 interactions at HL-LHC.

The MTD has the potential to contribute significantly to LLP searches. Though the timing information from MTD will not be available at the Level-1 trigger, the 30-40 ps precision expands the accuracy with which CMS can identify LLPs or their decay products within the barrel region $|\eta| < 3.0$ [65]. This means that the delay from LLPs with a displaced vertex could be identified by measuring $\beta_{\text{LLP}} = v_{\text{LLP}}/c$, and potentially allowing for the LLP mass to be reconstructed [14]. The MTD has two subdetectors: the barrel timing layer (BTL) and the endcap timing layer (ETL). The BTL will be comprised of LYSO:Ce scintillator crystals, readout with SiPMs. A testbeam at the Fermilab Test Beam Facility was performed in May 2019 and February 2020 to characterize the BDT module prototypes and documented in CMS DN-2020/017 [7]. Highlights of the testbeam analysis are included here.

F.1 Barrel Timing Layer Testbeam


Two setups were tested: individual LYSO bars with SiPM readouts at each end and a full module (12 LYSO bars, read out with an array of SiPMs on each end). The SiPM signals are then directed to a TOFPET2 chip for digitization and data acquisition [7]. The results here focus on the response uniformity and cross-talk in the module setup. After applying preselection cuts, the MIP response is found from the energy distribution and fit to a Landau distribution. The most probable value of the fit corresponds to the channel's response to a MIP.

F.1.1 Response Uniformity


Intercalibration (IC) coefficients are then defined to quantify the differences in response to a MIP:

$$IC_{coeff}[i] = \frac{MIP_{peak}[i]}{\frac{\sum_{j=1}^N MIP_{peak}[j]}{N}} \quad (F.1)$$

where $\sum_{j=1}^N MIP_{peak}[j]/N$ averages over the MIP peak values of each channel in an array.



(a) IC coefficients, pin connected module



(b) IC coefficients, flex connected module

Figure F.1: IC distributions for the pin and flex connected modules at 6 V overvoltage. The spread of IC coefficients is below 10% across the pin-connected module and slightly wider on the flex module. Figures from [7].

Plots of the IC coefficient distribution are in Figure F.1, with SiPMs on the top and bottom of each array indicated. The spread of the ICs is representative of various effects defining the response of a given channel to MIPs [7]:

- Differences in the light output of a given crystal bar (including both intrinsic light yield and light transport efficiency effects);

- Quality of optical coupling of each end of the bar with the corresponding SiPMs (top and bottom);
- Thickness of the protective silicon resin layer on top of a SiPM array;
- Response of the SiPM: mainly the gain and PDE, which depend on the breakdown voltage of each SiPM (since they all receive a common bias voltage).

F.1.2 Efficiency

For each bar within the array, we determine its position in the coordinate frame of the tracker by plotting the efficiency of MIP detection in a certain bar as a function of the tracker x coordinates. We define efficiency as the percent of hits at an x position that are above the MIP energy for an individual channel. These histograms are then fit with a two-sided error function, leaving the position (center) as a free parameter.




Figure F.2: Efficiency of MIP peak detection for the pin-connected module for each channel across the top of the array. Figure from [7].

The efficiency fit results are used to define the positions of each bar. A position cut is applied around the center of each bar to avoid the optical cross-talk due to MIPs crossing the boundaries between adjacent bars. The measured positions of each bar match the expected pitch between adjacent crystals of 3.2 mm as in the geometrical specifications of the array from the manufacturer. The fit in Figure F.2 determines the positions. Additionally, the fit results show that the tracker resolution is approximately 0.42 for both arrays.

The expected efficiency for MIP detection in a given sensor is 100% since a MIP that crosses a few mm of LYSO will produce a large amount of light that the SiPM detects. However, the measured efficiency is around 70% for both arrays, as shown in Figure F.2. The observed inefficiencies are attributed to the scattering of beam particles in the material upstream of the arrays. To verify this hypothesis, we measured the efficiency for events where the MIP also passed through the small

reference crystal positioned downstream of the arrays. We require that the energy measured in the pixel crystal is compatible with the energy of a MIP. In this case, the efficiency is $\approx 98\%$ when the array crystal bar is aligned with the pixel crystal [7].

F.1.3 Cross-talk

Optical cross-talk effects are studied by measuring how much light is detected by the SiPMs coupled to the bar hit by a MIP with respect to the amount of light detected by the SiPMs coupled to the neighboring bars (which are not traversed by a MIP). Light may be detected in neighboring SiPMs mainly because of two reasons [7]:

- Photons cross the boundary between adjacent crystal bars due to the imperfections of the lateral crystal surfaces and to non-perfect ESR (Enhanced Specular Reflector by 3M ©) wrapping (either partial coverage of the crystal surface and limited reflectivity);
- Photons that exit the crystal end face are refracted at a wider angle due to the change of refractive index from the crystal ($n = 1.81$) to the glue and silicon resin layers ($n \sim 1.4 - 1.5$) and can thus be detected in neighboring SiPMs as well as travel within the glue/resin layer and be detected further away in a different SiPM of the array.

F.2 Timing in Future Detectors

Improvements in timing detectors are valuable at the HL-LHC in handling the high pile-up environment. In addition, timing calorimeters will be vital in a muon collider to reduce beam induced backgrounds (BIBs) resulting from the continuous decays of muons along the collider [32]. The BIBs are created from muons decaying (lifetime of $c\tau = 2.2 \mu s$) and creating secondary particles from synchrotron photons or induced electromagnetic showers with machine components that produce muons, hadrons, and electrons. To reduce these backgrounds, excellent energy and time resolution and required for muon collider detectors. Current timing-focused work, such as the HCAL timing-based trigger, and future timing detectors, such as the MTD, are invaluable precursors to the types of detectors required for a successful muon collider.

Bibliography

- [1] S. Agostinelli et al., *GEANT4—a simulation toolkit*, Nucl. Instrum. Meth. A **506** (2003), 250–303.
- [2] 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.
- [3] ———, *Review of opportunities for new long-lived particle triggers in Run 3 of the Large Hadron Collider*, 2021.
- [4] Alexandre Alves, Alex Gomes Dias, and Roberto da Silva, *The 7% rule: A maximum entropy prediction on new decays of the higgs boson*, Nuclear Physics B **959** (2020), 115137.
- [5] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, Journal of High Energy Physics **2014** (2014), no. 7, 1–134.
- [6] David Bailin and Alexander Love, *Supersymmetric gauge field theory and string theory*, CRC Press, 1994.
- [7] A. Benaglia, F. De Guio, G. Kopp, M. Lucchini, M. Malberti, B. Marzocchi, P. Meridiani, and A. Mohammadi, *Characterization of BTL module prototypes readout with TOFPET2 ASIC using proton beams*, The Compact Muon Solenoid Experiment Detector Note (2020), 1–34.
- [8] Christian Bierlich, Smita Chakraborty, Nishita Desai, Leif Gellersen, Ilkka Helenius, Philip Ilten, Leif Lönnblad, Stephen Mrenna, Stefan Prestel, Christian T. Preuss, Torbjörn Sjöstrand, Peter Skands, Marius Uthmeim, and Rob Verheyen, *A comprehensive guide to the physics and usage of PYTHIA 8.3*, 2022.

- [9] Dr. Aaron Bundock, *The CMS Calorimeter Trigger for LHC Run II*, Journal of Physics: Conference Series **928** (2017), no. 1, 012004.
- [10] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez, *The anti-kt jet clustering algorithm*, Journal of High Energy Physics **2008** (2008), no. 04, 063.
- [11] CERN, *CERN’s Accelerator Complex*, 2024, [Online; accessed 16-February-2024].
- [12] Somnath Choudhury, *Performance of the High-Level Trigger System at CMS in LHC Run-2*, IEEE Transactions on Nuclear Science **68** (2021), no. 8, 2035–2042.
- [13] Sergio Cittolin, Attila Rácz, and Paris Sphicas, *CMS The TriDAS Project: Technical Design Report, Volume 2: Data Acquisition and High-Level Trigger. CMS trigger and data-acquisition project*, Technical Design Report CMS, CERN, Geneva, 2002.
- [14] Collaboration CMS, *A MIP Timing Detector for the CMS Phase-2 Upgrade*, Tech. report, CERN, Geneva, 2019.
- [15] CMS Collaboration, *CMS Physics Technical Design Report. Volume I: Detector Performance and Software*, Tech. Report CERN-LHCC-2006-001. CMS TDR 8.1, CERN, Geneva, February 2006, Final version.
- [16] ———, *The CMS experiment at the CERN LHC*, Journal of Instrumentation **3** (2008), no. 08, S08004–S08004.
- [17] ———, *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.
- [18] ———, *Shape, transverse size, and charged hadron multiplicity of jets in pp collisions at $\sqrt{s} = 7$ TeV*, JHEP **06** (2012), 160, Comments: Supplementary data tables available.
- [19] ———, *Particle-flow reconstruction and global event description with the CMS detector*, Journal of Instrumentation **12** (2017), no. 10, P10003.
- [20] ———, *Search for long-lived particles decaying into displaced jets*, Tech. Report CMS-PAS-EXO-19-021, CERN, Geneva, 2020.
- [21] ———, *The Phase-2 Upgrade of the CMS Level-1 Trigger*, Tech. Report CERN-LHCC-2020-004. CMS-TDR-021, CERN, Geneva, Apr 2020, Final version.

- [22] ———, *Out-of-Time Pileup Subtraction in HCAL at Level-1 (CMS DP-2021/012)*, Tech. report, CERN, cms-dpg-conveners-hcal@cern.ch, May 2021, Available at <https://twiki.cern.ch/twiki/bin/view/CMSPublic/HcalDPGResultsCMSDPS2021012>.
- [23] ———, *A portrait of the Higgs boson by the CMS experiment ten years after the discovery*, Nature **607** (2022), no. 1, 60–68.
- [24] ———, *Measurement of the production cross section of a W boson in association with a charm quark in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Tech. report, CERN, Geneva, 2022.
- [25] ———, *Search for long-lived particles in pp collision events with delayed trackless jets at $\sqrt{s} = 13$ TeV*, Tech. report, CERN, Geneva, 2022.
- [26] ———, *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>.
- [27] ———, *Search for long-lived particles decaying in the CMS muon detectors in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Tech. report, CERN, Geneva, 2023.
- [28] CMSSW, *Github commit implementing LUT for prompt - delayed boundary in simulation*, Available at <https://github.com/cms-sw/cmssw/commit/43e2e9b3c6e006e86202c488f1fa6db7db5e3a6c#diff-d2373454a91b97bf3479bad0e948f8dcca8e77e449a73565537a17c9460ae08f>.
- [29] ———, *HcalTestAnalysis.cc in CMSSW*, Available at <https://cmssdt.cern.ch/lxr/source/SimG4CMS/Calo/plugins/HcalTestAnalysis.cc#0663>.
- [30] CMS HCAL Collaboration, *Design, Performance, and Calibration of CMS Hadron-Barrel Calorimeter Wedges*, Tech. Report 1, CERN, Geneva, 2008.
- [31] ———, *HCAL uTCA Trigger and Readout Module (uHTR)*, CMS Public DocDB, July 2020, Available at https://cms-docdb.cern.ch/cgi-bin/PublicDocDB/RetrieveFile?docid=12306&filename=uhtr_spec.pdf&version=21.
- [32] F. Collamati, C. Curatolo, D. Lucchesi, A. Mereghetti, N. Mokhov, M. Palmer, and P. Sala, *Advanced assessment of beam-induced background at a muon collider*, Journal of Instrumentation **16** (2021), no. 11, P11009.
- [33] Luminosity POG Conveners, *CMS Luminosity – Public Results*, Twiki, 2023.

- [34] 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).
- [35] David Curtin, *Discovering Neutral Naturalness*, LHC Searches for Long-Lived BSM Particles: Theory Meets Experiment, 13 November 2015, Amherst Center for Fundamental Interactions.
- [36] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi, *Delphes 3: a modular framework for fast simulation of a generic collider experiment*, Journal of High Energy Physics **2014** (2014), no. 2, 1–20.
- [37] D. de Florian et al., *Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector*, CERN Yellow Reports: Monographs, CERN, Oct 2016, 869 pages, 295 figures, 248 tables and 1645 citations. Working Group web page: <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/LHCHXSWG>.
- [38] 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>.
- [39] Laurent Thomas (L1 DPG), *L1T prefiring of TeV jets*, PdmV General Meeting: HCAL prefire follow-up, 7 November 2023, Available at <https://indico.cern.ch/event/1337844/#6-trigger-prefire-correction>.
- [40] Manuel Drees, *An Introduction to Supersymmetry*, Current topics in physics. Proceedings, Inauguration Conference of the Asia-Pacific Center for Theoretical Physics (APCTP), Seoul, Korea, June 4-10, 1996. Vol. 1, 2, 1996.
- [41] Manuel Drees, Rohini Godbole, and Probir Roy, *Theory and phenomenology of sparticles: An account of four-dimensional $N=1$ supersymmetry in high energy physics*, World Scientific, 2004.
- [42] Sonia El Hedri, Ann E Nelson, and Devin GE Walker, *Reducing the quadratic divergence in the higgs boson mass squared without top partners*, Physical Review D **98** (2018), no. 3, 035029.
- [43] EXO Group, on behalf of the CMS Collaboration, *CMS Exotica Summary plots for 13 TeV data*, CMS Twiki, May 26, 2021.
- [44] PdmV Group: Adriano Di Florio, Mikko Voutilainen, and Sunil Bansal, *PdmV General Meeting: HCAL prefire follow-up*, PdmV General Meeting, 7 November 2023, Available at <https://indico.cern.ch/event/1337844/>.

- [45] 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.
- [46] Tullio Grassi, *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.
- [47] 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>.
- [48] ———, *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.
- [49] David Griffiths, *Introduction to Elementary Particles*, Wiley-VCH, Weinheim, DE, 2008.
- [50] CMS JERC Group, *CMS Jet Energy Resolution and Corrections*, Online, Available at <https://cms-jerc.web.cern.ch>.
- [51] LHC Cross Section Working Group, *SM Higgs production cross sections at $\sqrt{s} = 13$ TeV (update in CERN Report 4 2016)*, CERN Twiki (2023).
- [52] 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.
- [53] Werner Herr and B Muratori, *Concept of luminosity*, CAS - CERN Accelerator School: Intermediate Accelerator Physics (2006), 361–377.
- [54] Elliot Hughes, *A New Generation Of Charge Integrating ADC For The CMS HCAL Upgrade: The QIE10/11*, TIPP 2014:140604, 2014, Available at <https://home.fnal.gov/~chlebana/CMS/Phase1/Elliot-Hughes.pdf>.

- [55] I. Neutelings and J. Riebesell and A. Tsagkaropoulos and E. Soto Apolinar and J. and A. Ostadrahimi and others, *Graphics with TikZ in LaTeX*, Available at <https://tikz.net/category/physics/particle-physics/>.
- [56] Harjot Kaur, Jyoti Babbar, Suman Dasgupta, and Sunanda Banerjee, *IsoTrack Report: 2023D Data Response Corrections*, HCAL DPG Meeting, 1 September 2023, Available at <https://indico.cern.ch/event/1320801/#4-isotrack-report>.
- [57] Kiley Kennedy, *Update: Truth Matching Study*, Run 3 HCAL LLP Working Meeting, 6 June 2023, Available at <https://indico.cern.ch/event/1294292/>.
- [58] Kiley Elizabeth Kennedy, *Event Display*, GitLab repository, <https://gitlab.cern.ch/kikenned/EventDisplay>, 2024.
- [59] Simon Knapen and Steven Lowette, *A Guide to Hunting Long-Lived Particles at the LHC*, Annual Review of Nuclear and Particle Science **73** (2023), no. 1, 421–449.
- [60] Gillian Kopp, Kiley Kennedy, Cristián Peña Chris Tully, Si Xie, and Adi Bornheim, *Search for long-lived particles using hadronic jets in proton-proton collisions at $\sqrt{s} = 13.6$ TeV*, The Compact Muon Solenoid Experiment Analysis Note (2023), 1–2.
- [61] 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.
- [62] Gillian Kopp, Chris Tully, and Kelvin Mei, *HCAL uHTR Teststand Set-up and Maintenance Instructions*, CMS DocDB, September 2021, Available at <https://cms-docdb.cern.ch/cgi-bin/DocDB/ShowDocument?docid=14328>.
- [63] Gillian Baron Kopp, Christopher Tully, Kiley Elizabeth Kennedy, Wonyong Paul Chung, Svitlana Hoienko, Jeremiah Michael Mans, Michael David Krohn, Bryan James Crossman, Joshua Hiltbrand, Andris Skuja, and Long Wang, *A Novel Timing Trigger with the CMS Hadron Calorimeter*, CERN Document Server (2023).
- [64] Thomas S. Kuhn, *The structure of scientific revolutions*, University of Chicago Press, Chicago, 1962.

- [65] Lawrence Lee, Christian Ohm, Abner Soffer, and Tien-Tien Yu, *Collider searches for long-lived particles beyond the Standard Model*, Progress in Particle and Nuclear Physics **106** (2019), 210–255.
- [66] Stephen P. Martin, *A supersymmetry primer*, Adv. Ser. Direct. High Energy Phys.18,1(1998) (1997, 2016), 1–98.
- [67] 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_HBUUpgrade_LHCC2019.pdf.
- [68] Laurent Thomas (on behalf of the L1 DPG and JME groups), *HCAL time misalignment in 2022/2023: impact for the L1 trigger and for JME*, PPD General Meeting, 12 October 2023, Available at <https://indico.cern.ch/event/1333510/#25-run-3-dpg-impact-of-hcal-ti>.
- [69] Maurizio Pierini, *The CMS Physics Program: Run-3 and HL-LHC*, L1T Workshop in Athens, 11-15 September 2023, Available at <https://indico.cern.ch/event/1288569/contributions/5491063/>.
- [70] Tai Sakuma and Thomas McCauley, *Detector and Event Visualization with SketchUp at the CMS Experiment*, Journal of Physics: Conference Series **513** (2014), no. 2, 022032.
- [71] Brian Shuve, *Theory Overview of LLPs*, 3rd Workshop of the LLP Community, May 16, 2018.
- [72] ———, *How the LLP Landscape has Changed in 6 Years (theory)*, 11th Workshop of the LLP Community, May 30, 2022.
- [73] A.M. Sirunyan et al., *Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Journal of Instrumentation **15** (2020), no. 10, P10017–P10017.
- [74] Steven Schramm, *Jet Reconstruction (Experimental)*, PREFIT School, DESY, March 2020, Available at <https://indico.cern.ch/event/817757/contributions/3712530/attachments/2001044/3340282/StevenSchramm-JetReco-Lecture1.pdf>.
- [75] Laurent Thomas, *Overview of the Level-1 Trigger for Run-3*, L1T Workshop at CERN, 27-29 April 2022.
- [76] Mark Thomson, *Modern Particle Physics*, Cambridge University Press, New York, 2013.

- [77] Christopher G. Tully, *Elementary Particle Physics in a Nutshell*, Princeton University Press, 2011.
- [78] Wikipedia contributors, *Physics beyond the standard model — Wikipedia, the free encyclopedia*, 2017, [Online; accessed 31-January-2024].
- [79] R. L. Workman and Others, *Review of Particle Physics*, PTEP **2022** (2022), 083C01.