

University of California
Santa Barbara

Delayed Jets and the Dark Sector: Long-Lived Particles at the LHC and Beyond

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Physics

by

Ryan Thomas Schmitz

Committee in charge:

Professor David Stuart, Chair
Professor Claudio Campagnari
Professor Nathaniel Craig

September 2025

Acknowledgments

To my wife and partner, Lindsay: This dissertation is as much yours as it is mine. You have worked every bit as hard to bring out the best in me as I've worked during my PhD.

To the greatest advisor of all time, David: Thank you for allowing me the freedom to grow as a person and scientist through these many years. You've shaped how I view people and leadership as much as you've shaped my passion for physics.

To my many, many incredible colleagues in CMS, milliQan, and UCSB, but especially Matthew, Neha, Dariush, Mike, Claudio, Chris, Andy, Patrick, Bennett, and James: Your mentorship and friendship has been the absolute highlight of my PhD. I am so honored to be in this community with all of you.

To my mom: I owe you my intellectual curiosity. Someday, I will play “which is bigger” with my kids and hope to do as good a job as you did.

To my dad: I owe you my passion and confidence. I have always known I could chase my dreams knowing that, whatever happened, you'd be there to help.

To my friends from Elk River, the U, marching band, hockey, board games, baseball, and magic: No matter how busy I have gotten, no matter how many threads I've been pulled by, no matter how long it's been, you've all given me steadfast support. You've made me a more whole, well-rounded person, something that has also, ultimately, made me a better scientist.

Lastly, to my cat, Nova, who has been my companion for the many late nights I have spent working during my PhD: I hope you enjoy window-watching in London as much as you've enjoyed it in Santa Barbara. You've made enough late-night biscuits to earn it.

Abstract

Two searches for long-lived particles (LLPs) are presented. For the CMS experiment, the first timing-based LLP trigger was developed for Run 3, targeting LLPs of intermediate lifetime using the CMS ECAL. Trigger development and performance is discussed, and the expected results of the Run 3 analysis in a search for $H \rightarrow \chi\chi \rightarrow bbbb$ is presented. Data was recorded in 2023 and 2024 in $\sqrt{s} = 13.6$ TeV collisions, corresponding to a total integrated luminosity of 136.9 fb^{-1} . New sensitivity is expected for models with lifetimes between $0.4 \text{ m} < c\tau < 100 \text{ m}$ and masses between 400 GeV and 2 TeV. Peak sensitivity excludes twin Higgs branching ratios $\text{BR} < 10^{-3}$.

For the milliQan experiment, a complete description of the milliQan simulation framework is presented. Implementation and validation studies are discussed for each of the generation, propagation, and Geant4 detector simulation components of the overall framework. The Run 3 milliQan experiment and first analysis and results of the Run 3 bar detector are presented. Data was recorded in 2023 and 2024 in $\sqrt{s} = 13.6$ TeV collisions, corresponding to a total integrated luminosity of 124.7 fb^{-1} . New sensitivity is achieved to millicharged particles with charge $0.003 < q/e < 0.25$ and mass $0.45 \text{ GeV} < m < 20 \text{ GeV}$.

Contents

Acknowledgments	i
Abstract	ii
1 Introduction and Theory	1

1.1	The Standard Model	1
1.2	Long-Lived Particles	2
1.3	LLPs at CMS	4
1.4	mCPs at milliQan	8
2	The LHC and the CMS Detector	11
2.1	LHC Overview	11
2.2	The CMS Experiment	13
2.3	CMS Trigger	20
3	CMS Delayed-Jets Trigger	26
3.1	Trigger Strategy	26
3.2	Run 3 Delayed Jet Trigger	28
3.3	Trigger Implementation	32
3.4	Trigger Performance	35
4	CMS Delayed-Jets Analysis	38
4.1	Data and MC Samples	38
4.2	Event Selection	39
4.3	Timing Performance	47
4.4	Background Estimation	59
4.5	Interpretation	72
5	The milliQan Experiment: Detector, Commissioning, and Calibration	81
5.1	Scintillator	81
5.2	Photomultiplier Tubes	86
5.3	The milliQan Demonstrator	91
5.4	milliQan Readout Amplifiers	94
5.5	Run 3 milliQan Bar Construction	96

5.6	Run 3 milliQan Bar Detector	99
5.7	Run 3 Slab Detector	106
6	The milliQan Monte Carlo	112
6.1	Event Generation	113
6.2	Propagation	118
6.3	Geant4 Detector Simulation	124
6.4	Pulse Injection	134
6.5	Beam Muon Validation	138
6.6	Cosmic NPE Validation	144
7	The milliQan Experiment: Run 3 Bar Detector Analysis	147
7.1	Data and Trigger	147
7.2	Data Processing	149
7.3	Calibration	152
7.4	Search Design	161
7.5	Search Backgrounds	162
7.6	Selections and Validation	164
7.7	Results	175
7.8	Systematics	182
7.9	Sensitivity	185
8	Summary	188
A	Appendix	190
A.1	Delayed Jet Trigger Implementation	190
A.2	Complete List of Delayed Jet Triggers	192
A.3	ECAL Cell Timing: Ratio Method	192
A.4	ECAL Cell Amplitude: Multifit	193

A.5 Run 3 Delayed Jet Analysis Yields	195
A.6 Run 3 Delayed Jet N-1 Plots	204
A.7 Delayed Jet Analysis Event Flow	208
A.8 Simulation Class Description	208
A.9 Simulation Output Reference	211
A.10 Scintillator Surface Roughness	213
A.11 MilliQan Offline Tree Branches	214
A.12 Full MilliQan Result Cutflow	216
A.13 MilliQan Systematics	216

1 Introduction and Theory

1.1 The Standard Model

The Standard Model of particle physics is the most predictive, most rigorous theory of physics ever assembled. It predicts interactions over a wide range of energy scales with extreme precision, from eV-scale atomic physics to TeV-scale LHC interactions. The fermions and interactions of the Standard Model are summarized in Figure 1.1.




Figure 1.1: The fermions and interactions of the Standard Model [1].

In one recent example of the resilience of the Standard Model, the famous $g - 2$ anomaly,

considered for years to be one of the most significant, specific examples of hints towards BSM physics, was closed. The latest experimental and theoretical predictions yield [2, 3]:

$$a_{\mu}^{\text{exp}} = 0.001165920715 \pm 0.000000000145 \quad (1.1)$$

$$a_{\mu}^{\text{SM}} = 0.00116592033 \pm 0.000000000062 \quad (1.2)$$

$$a_{\mu}^{\text{exp}} - a_{\mu}^{\text{SM}} = 38(63) \times 10^{-11} \quad (1.3)$$

In spite of this spectacular agreement, persistent and irrefutable problems with the standard model still exist. It fails to include gravity. It cannot account for the matter-antimatter asymmetry in the universe. The hierarchy problem suggests the Higgs mass should be large due to GUT-scale corrections, but nevertheless the Higgs is very light. Neutrinos in the SM are massless in contradiction with neutrino oscillations. And most of all, it is overwhelmingly likely that the explanation for dark matter will require one or more BSM particle candidates. This is strongly suggested by the distinct and non-overlapping distributions of dark matter, such as in the galaxy cluster ZwCl0024+1652 in Figure 1.2 [4].

How can such tension exist between the lack of obvious electroweak anomalies and known problems with the Standard Model?

1.2 Long-Lived Particles

The original goal of BSM searches at the LHC was to target prompt, heavy new physics. This is consistent with the general trend we have observed in the Standard Model, where particle lifetime is a decreasing function of mass. However, a variety of SM lifetimes exist on timescales which are neither detector-prompt nor detector-stable, as in Figure 1.3.

Long-lived particle lifetimes can generally be generated through one of three mechanisms [7]:

1. The decay is mediated by a heavy, off-shell particle, introducing a suppression factor



Figure 1.2: The dark matter distribution measured via gravitational lensing of ZwCl0024+1652 overlaid onto its visible matter distribution. The observed ring structure is unique to the distribution of dark matter in this cluster [5].

$(m/M)^n$ for a particle of mass m and off-shell mediator M . In the Standard Model, decays of kaons are mediated through an off-shell W and are suppressed by a factor m^5/m_W^4 .

2. The decay is phase-space suppressed by the masses of the available decay products, such as for the neutron in the SM.
3. The decay is suppressed by a small coupling constant associated with an approximate symmetry. In the Standard Model, B hadron decays are suppressed by the small CKM matrix element V_{cb} .




Figure 1.3: Lifetimes of the Standard Model. SM lifetimes cover many orders of magnitude, including detector-prompt, metastable, and detector-stable states [6].

Under these mechanisms, there are few model-independent bounds on the possible lifetimes of new physics. This suggests that new physics with potentially minimal existing experimental constraints may be hiding in the lifetime phase space between heavy, prompt searches and high MET final states.

1.3 LLPs at CMS

Signatures for LLPs in CMS are varied and unusual compared to well-understood SM backgrounds. Accordingly, LLP signatures can appear as misreconstructed objects or noise and become rejected. As a result, LLP searches must be targeted and require tracking, timing, jet shape, or energy characteristics which are inconsistent with legitimate noise or prompt QCD backgrounds. A range of sample LLP searches in CMS is shown in Figure 1.4.




Figure 1.4: An overview of the types of LLP searches happening at CMS and ATLAS. LLP models require targeted analysis techniques to identify and separate signal from prompt backgrounds [8].




Figure 1.5: Sensitivity to delayed jets relies on the delay from slow-moving LLPs and the path length due to the displaced decays [9].

The CMS Electromagnetic Calorimeter (ECAL) provides an excellent opportunity to search for LLPs, as the timing capabilities of the ECAL can be used to identify nonprompt




Figure 1.6: (Left) The Twin Higgs signal model. We focus on the 4b final state for this search. (Right) A sample Twin Higgs signal event overlaid with the CMS detector dimensions.

or “delayed” jets produced by displaced decays of heavy LLPs inside the ECAL or tracker, as in Figure 1.5. I present a review of the CMS detector in Chapter 2.

Previous LLP searches in CMS using timing either targeted decays in the tracker with shorter lifetimes or model-dependent final states [10, 11]. In Run 2, ECAL timing was used to search for high MET final states, but the sensitivity to other models was limited by this MET requirement in the trigger [12]. In Run 3, new delayed jet triggers using ECAL timing were developed to expand the phase space to remove this MET requirement and allow targets with very low H_T , lower lifetimes, and single jet final states. The implementation and performance of these new ECAL timing triggers is discussed in Chapter 3. To exploit the advantages of these new triggers, we will use the Twin Higgs model as a low MET benchmarking model for this new Run 3 delayed jets search (Figure 1.6).

1.3.1 Twin Higgs

To understand the Twin Higgs model, let’s examine the hierarchy problem in more detail. If one computes the one loop corrections to the Higgs mass up to some energy Λ , the quadratic divergence terms are [13]:

$$\delta M_H^2 = \frac{\Lambda^2}{16\pi^2} \frac{3}{v^2} (M_H^2 + M_Z^2 + 2M_W^2 - 4M_t^2) \quad (1.4)$$

This highlights the problem: corrections to the Higgs mass are sensitive to the scale of new, arbitrarily heavy physics. At the Planck scale $\Lambda \approx 10^{19}$ GeV where quantum gravity is expected to emerge, it becomes clear that a very precise, fine-tuned cancellation of higher order corrections is necessary to produce the observed Higgs mass of $M_h = 125$ GeV.

One solution to this problem is to introduce a twin sector with gauge structure $[SU(3) \times SU(2) \times U(1)]_B$ as a mirror to the Standard Model with identical gauge structure $[SU(3) \times SU(2) \times U(1)]_A$ [14]. In addition, we introduce two Higgs doublets H_A and H_B under an approximate global symmetry such that the combined Higgs sector respects an $SU(4)$ symmetry. After this procedure, the one-loop contributions of the combined twin and SM sectors are given by:

$$\Delta V \propto \frac{\Lambda^2}{16\pi^2} (|H_A|^2 + |H_B|^2), \quad (1.5)$$

This contribution is now invariant under $SU(4)$ and does not contribute to the Higgs mass under a broken $SU(4)$ symmetry. The remaining contributions to the mass are now only logarithmically divergent in the new physics scale Λ [15].

A soft tuning of the energy scales of the SM and twin sectors can be set by the vev scales $\nu/f = \nu/\sqrt{\nu^2 + \nu'^2}$. If the twin vev ν' is heavier than the SM ν , the SM Higgs couplings to the twin sector are suppressed, a feature which is necessary to remain consistent with LHC Higgs measurements. Current constraints set $3 < f/\nu < 6$, bounded by experimental data from below and naturalness considerations from above [16]. This suggests a Twin Higgs mass between approximately 400 to 750 GeV.

Separate from model constraints, at what mass scale should we naturally expect this new physics to emerge? Returning to Eq. 1.4, we can compute the implied scale of new physics using the Higgs doublet mass parameter $M_H^2 = M_h^2/\sqrt{2} = 88$ GeV. Solving for Λ :

$$\Lambda^2 = \frac{16\pi^2 M_H^2 v^2}{3(M_H^2 + M_Z^2 + 2M_W^2 - 4M_t^2)} \quad (1.6)$$

Computing, we find $\Lambda = 550$ GeV. This is not a precise calculation, but it's a naturally motivated [17] and in agreement with Twin Higgs model arguments. In addition, heavy Higgs models are less constrained compared to SM portals, making this is a natural fit as a first benchmark of the new Run 3 delayed jet sensitivity.

There are other benefits to the model as well. The Twin Higgs has few constraints on LLP lifetime, low MET in the final state, and as the timing trigger sensitivity increases for heavier models¹, the model is both well-motivated and well-positioned to exploit the newly available sensitivity.

In Chapter 4, I discuss the analysis design and expected sensitivity for a new search targeting a heavy Twin Higgs. We explore Twin Higgs models of heavy Higgs mass $M_H \in [400, 2000]$ GeV and LLP lifetime $c\tau \in [0.1, 100]$ m decaying into $H \rightarrow \chi\chi \rightarrow bbbb$, a purely hadronic final state. We primarily focus on models in which the mass of the LLP is given by $m_X = m_H/2 - 25$ GeV, as highly boosted models suppress the measured jet delay times to which the trigger is most sensitive.

1.4 mCPs at milliQan

Let's now consider the problem of dark matter. The likely existence of a BSM particle explanation for dark matter has motivated a diverse array of searches which include both collider production and direct detection experiments. The so-called ‘‘WIMP miracle’’ in which supersymmetric theories apparently provided a viable dark matter candidate at accessible LHC energies led to significant focus on a single-particle explanation for dark matter [18]. However, no evidence for WIMP dark matter has been found at LHC or direct detection experiments to date, and this explanation for dark matter has become increasingly constrained

¹Prompt QCD backgrounds peak at low H_T , so heavier models which produce higher H_T are more distinguishable.

[19].

To account for this lack of evidence, an alternative explanation for dark matter which has become popular recently are “dark sector” models[20, 21]. In these models, one or more dark matter particles are embedded inside a larger dark sector which contain its own dark interactions.

One of the simplest dark sector extensions to the Standard Model involves the introduction of a massless abelian gauge field A' , satisfying $U(1)'$ and coupling to dark fermions with charge e' . A general Lagrangian for this model, written alongside the SM $U(1)$ hypercharge and including a charged dark fermion ψ , is given by:

$$\mathcal{L} = \mathcal{L}_{\text{SM}} - \frac{1}{4} A'_{\mu\nu} A'^{\mu\nu} - \frac{\kappa}{2} A'_{\mu\nu} A^{\mu\nu} + i\bar{\psi}(\not{\partial} + ie'A' + iM_{mCP})\psi \quad (1.7)$$

where κ sets the scale of “kinetic mixing” between A and A' [22]. If we perform a field redefinition $A' \rightarrow A' - \kappa A$, we diagonalize the interaction term and can now rewrite 1.7 as:

$$\mathcal{L} = \mathcal{L}_{\text{SM}} - \frac{1}{4} A'_{\mu\nu} A'^{\mu\nu} + i\bar{\psi}(\not{\partial} + ie'A' - i\kappa e'A + iM_{mCP})\psi \quad (1.8)$$

This introduces a SM hypercharge $\kappa e'$ to the dark fermion. If we consider a toy model where kinetic mixing is mediated by loops of particles M which are charged under both SM hypercharge and $U(1)'$ as (e, e') and M' which is charged as $(e, -e')$, the value of κ is naturally given by²:

$$\epsilon \sim \frac{e'e}{16\pi^2} \ln\left(\frac{M'}{M}\right) \simeq 6 \times 10^{-4} \left(\frac{e'}{e}\right) \ln\left(\frac{M'}{M}\right) \quad (1.9)$$

where M'/M is the mass ratio of the heavy mediators [22, 23, 24, 25]. This naturally implies a value $k \sim 0.001$, introducing a “millicharge” to fermions charged under the dark $U(1)'$.

This generic model strongly motivates the search for mCPs. However, there are other

²Note that we need two such particles to ensure that the charges trace to zero. Also note the substitution $\alpha = e^2/4\pi$.

model-independent justifications. mCP searches are a direct test of SM charge quantization, and SM constraints for heavy, fractionally charged particles are very weak at high mass, only excluded to $q/e > 0.33$ above 50 GeV [26, 27] and $q/e > 0.25$ above 5 GeV [28]. Fractionally charged particles of charge $q/e = 1/6$ exist in both SM and BSM interpretations and are not excluded by these collider and accelerator bounds [29]. Millicharge is also a generic feature of many models such as string theory and quantum gravity, and searches for mCPs provide model-independent tests of many mCP predictions [30].

In Chapters 6- 7, I discuss a search for mCPs using the milliQan experiment. I review the underlying MC framework of the experiment Chapter 6, discuss the experimental technology and commissioning in Chapter 5, then conclude by presenting a search using the Run 3 data in Chapter 7.

2 The LHC and the CMS Detector

2.1 LHC Overview




Figure 2.1: An aerial view of the Large Hadron Collider alongside its four major experiments [31].

The Large Hadron Collider (LHC) is the most powerful machine currently available as a probe of physics beyond the standard model (BSM). The LHC is a 27 km circular particle accelerator located outside Geneva, Switzerland. Located approximately 100 m underground, the LHC has four primary collision points: CMS (IP5) and ATLAS (IP1), which are general purpose experiments; ALICE (IP2), which studies heavy ion collisions; and LHCb (IP8), optimized for studying b-quarks.

In addition to the primary LHC experiments, the CERN accelerator complex supports a

variety of other experiments along the accelerator chain. A schematic layout of the CERN accelerator complex is shown in Figure 2.2.




Figure 2.2: The CERN Accelerator Complex. The LHC is supported by a chain of legacy accelerators, now used as injector stages for the LHC as well as proton sources for lower-energy experiments [32].

The LHC's primary mode of operation consists of two counter-rotating beams of 6.8 TeV protons which collide to produce a total center-of-mass energy of $\sqrt{s} = 13.6$ TeV. The protons are accelerated in several stages:

1. The current proton injector to the LHC is LINAC 4. First, an RF-plasma source generates H⁻ ions at 35 mA of beam current and 45 keV of beam energy. These H⁻ ions are injected into a series of accelerating stages including an RF Quadrupole (3 MeV), Drift-Tube Linac (50 MeV), Cell-Coupled Drift-Tube Linac (102 MeV), and Pi-mode structure (160 MeV), yielding a final H⁻ ion energy of 160 MeV [33].
2. These 160 MeV ions pass through a thin copper foil to strip them of their electrons

- before injection into the Proton Synchrotron Booster (PSB). The PSB accelerates beams to 2.0 GeV and injects them into the Proton Synchrotron (PS). The PS, built in 1959 as CERN's first flagship accelerator, now accelerates protons from 2.0 GeV up to 26 GeV where they are injected into the Super Proton Synchrotron (SPS) [34].
3. The SPS, initially built in 1976, now serves as the main injector to the LHC. Protons are accelerated to 450 GeV and injected into the LHC.
 4. Finally, in the LHC, protons are accelerated to 6.8 TeV. Proton bunches within each train are filled at a frequency of 40 MHz, or every 25 ns, with one bunch every 10 buckets of the LHC's 400 MHz RF cavities. Typical beam intensities in Run 3 involve instantaneous luminosities of $2 - 3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, or about 2×10^{11} protons per bunch.

2.2 The CMS Experiment




Figure 2.3: A schematic view of the Run 3 CMS detector with descriptions of each subsystem [35].

The CMS (Compact Muon Solenoid) experiment is a general-purpose LHC experiment designed to make precision measurements of the Standard Model as well as to search for exotic, beyond the Standard Model physics.

The CMS detector weighs 14 kilotonnes over a relatively “compact” dimension of 15 m x 28.7 m. The most distinguishing feature of the CMS detector is a 3.8 T solenoid at a 6 m inner diameter running down the axis of the cylindrical detector. This offers a strong, solenoidal magnetic field which allows for the precise identification and measurement of electrons, muons, photons, and hadrons.

The CMS detector consists of a series of subsystems, summarized below.

2.2.1 Pixel Detector

The CMS Inner Tracker consists of two primary detector modules. The innermost layer of the tracker is the silicon pixel detector. The pixel detector provides high precision 3D spatial measurements of charged particles nearest to the interaction point, allowing for high resolution particle tracking and vertex reconstruction.




Figure 2.4: The CMS Phase 1 Inner Tracker [36].

The current, Phase-1 pixel detector consists of 4 concentric barrel layers and 3 discs on each end, as in Figure 2.4. The pixel detector is made up of individual silicon pixels of size $100 \times 150 \mu\text{m}^2$, distributed over 1.85 m^2 of active silicon in total between $1184 + 672$ modules

in the barrel and forward pixel detector. This Phase-1 pixel detector is an upgrade over the original tracker, which consisted of only 3 barrel layers and 2 endcap layers. First deployed in 2017 near the end of Run 2, Run 3 is the first full physics run to take advantage of the upgraded pixel detector.

2.2.2 Silicon Strip Tracker

The outer silicon of the inner tracker is the Silicon Strip Tracker (SST), composed of 10 layers of silicon strip detectors in the barrel and 9 discs in the endcap. While the pixel detector is optimized for particle tracking and impact parameter resolution closest to the interaction point, the SST is optimized for tracking coverage. Consequently, the SST constitutes the majority of the total detector volume of the tracker, with 198 m^2 of active silicon distributed over 15146 modules. In total, the SST together with the pixel detector offers tracking coverage up to a full range of $|\eta| < 2.5$, as in Figure 2.6.




Figure 2.5: The CMS Silicon Strip Tracker. Single-sided modules are colored in red, double-sided modules in blue [37].

2.2.3 Electromagnetic Calorimeter (ECAL)

The CMS Electromagnetic Calorimeter (ECAL) is designed to identify and measure the energies and momenta of electrons and photons. Activity in the ECAL is also a significant component in jet reconstruction.

The Electromagnetic Calorimeter (ECAL) is made of an array of lead tungstate (PbWO_4)

crystals providing coverage up to $|\eta| < 3.0$. There are 61200 crystals in the barrel with dimension $2.2 \times 2.2 \times 23$ cm, segmented in units of approximately 0.087×0.087 in units of η, ϕ . There are 7324 crystals in each endcap with dimension $3 \times 3 \times 22$ cm, and together with the barrel form a total active volume of 9.7 m^3 . A preshower detector, composed of two planes of silicon interleaved with lead, is positioned before each endcap detector to distinguish between photons produced by the decays of neutral pions.




Figure 2.6: The CMS ECAL, with barrel, endcap, and component modules labeled [38].

Lead tungstate has several advantages as the material of choice for the CMS ECAL. The high-density material, with a short radiation length ($X_0 = 0.89 \text{ cm}$) and Molière radius ($R_M = 2.19 \text{ cm}$) enables for a compact calorimeter inside the CMS solenoid. Its fast scintillation time constant of approximately 10 ns allows for precision measurement of particle showers and enables the ECAL to uniquely¹ read out timing information to the High-Level Trigger (HLT) [41].

PbWO₄ is a radiation-hard material, though is subject to performance degradation under

¹Timing information is available in HCAL, but has poor resolution: HCAL analog signals are integrated over 25 ns intervals before pulse reconstruction[39]. HCAL timing is used in Run 3 at L1, however [40].

sustained radiation exposure; see Section 4.3. The light yield of lead tungstate is only 60-100 photons per MeV, compared to 10,000 photons per MeV in traditional plastic scintillator, so amplification is required from magnetic field-insensitive avalanche photodiodes (APDs) in the barrel and vacuum phototriodes (VPTs) in the endcaps, as in Figure 2.7.




Figure 2.7: Left: A CMS barrel crystal, coupled to two APDs. Right: A CMS endcap crystal, coupled to a VPT [38].

2.2.4 Hadronic Calorimeter (ECAL)

The CMS Hadron Calorimeter (HCAL) is used to identify and measure neutral and charged hadronic jets. Through jet reconstruction, the HCAL is also used to identify missing transverse energy (MET) and to assist in the identification of electrons and photons in the ECAL.

The HCAL consists of four subsystems: the hadron barrel (HB), the hadron endcap (HE), the hadron outer (HO), and the hadron forward (HF) calorimeters, as in Figure 2.13. The HB and HE are made from brass and steel absorbers interleaved with plastic scintillator tiles and overlaps in $|\eta| < 3.0$ coverage with the CMS barrel and endcap ECAL. The HF, constructed from steel and quartz tiles, covers the far forward region of $3.0 < |\eta| < 5.2$. Lastly, the HO is made of plastic scintillator and is the only barrel calorimeter layer located outside the radius of the CMS solenoid.




Figure 2.8: The CMS Hadron Calorimeter (HCAL). Layers of the same “depth” are optically summed using wavelength-shifting fibers and read out by the same photodetector [37].

2.2.5 CMS Solenoid




Figure 2.9: A schematic view of the CMS magnet. The return yoke is split into 12 azimuthal sections and 4 radial sections in the barrel and 3 sections in the endcaps [42].

Outside the HE and HB modules of the HCAL is the 6.8 T CMS superconducting NbTi solenoid. The solenoid is supported by a steel magnetic flux return yoke. At 12,500 tonnes, the magnet and return yoke constitute a majority of the mass of the CMS detector. The

return yoke also helps to shield the outer muon systems and calorimetry from jets, with 1.705 m of steel in total along the radius of the barrel and 1.541 m in each endcap, as in Figure 2.9.

2.2.6 Muon System

The muon system offers one of the central advantages of the CMS experiment: effectively identifying and tracking muons.




Figure 2.10: A view of the CMS detector in the R vs. Z plane, with muon stations highlighted in the barrel (MB) and endcaps (ME) [43].

Drift Tubes (DTs) are used in the barrel region of the muon system. The DTs are filled with a mixture of Ar and CO₂ with an anode wire in the center measuring ionization generated by the passage of muons.

DTs have best resolution when a charged particle passes perpendicular to the tube, so in the muon endcaps, cathode strip chambers (CSCs) are used. CSCs contain an orthogonal pattern of anode wires and copper cathode stripes in a gas chamber, providing a higher resolution position measurement in the forward region.

In both the barrel and endcaps there are resistive plate chambers (RPCs), a set of coupled gas chambers which offers improved timing precision at the cost of spatial resolution.

Finally, new for Run 3 is are gas-electron multiplier (GEM) chambers added to enhance forward track reconstruction in regions of high particle flux. These GEMs consisting of a set of layered polymer foils with holes punched in them which generate large electric fields when a potential is applied across the foils.



Figure 2.11: A picture taken of the author next to the CMS detector while open in late 2023. The exposed hadron endcap extends from the muon endcap detector on the right, while the red muon detectors interleaved with the iron return yokes are exposed on the left.

2.3 CMS Trigger

LHC proton collisions occur at a frequency of 40 MHz, or once every 25 ns. The largest standard event format recorded in CMS is the RAW format, containing all raw detector

information pre-reconstruction, and it is approximately 1 MB in filesize. This implies a data rate from CMS collisions alone of 40 TB/s. By comparison, according to a report in 2024 from Sandvine, the average rate of total global internet traffic is around 380 TB/s [44].




Figure 2.12: The CMS trigger system. Events are seeded using an initial trigger at L1, more complex trigger requirements are applied at HLT, and finally the trigger results are stored into datastreams depending on the trigger rates and data requirements [45].

To reduce these enormous data rates, we use a two-tiered trigger system as described in Figure 2.12.

2.3.1 L1 Trigger

The first stage is called the Level-1 (L1) Trigger. A custom hardware processor uses information from the ECAL, HCAL, and muon system and reduces the input 40 MHz collision rate down to about 100 kHz.

At L1, the basic objects generally available for selection are muons, $e\gamma$ (which are indistinguishable with no tracking information), jets, taus, and basic event kinematics such as transverse energy (HT) and missing transverse energy (MET). The L1 trigger must execute in a latency of 4 μ s, so L1 trigger paths generally consist of a single physics object sometimes combined with an event kinematic object to moderate total L1 trigger rates.

2.3.2 High-Level Trigger (HLT)

The second of the two-stage trigger is the HLT, or High-Level Trigger. At HLT, event rates are compressed from 100 kHz to a total output of a few kHz before data storage [46]. The HLT menu has a much larger latency budget compared to the L1 trigger, executing in around 260 ms, which allows for much more complex trigger algorithms [47].

The HLT contains a variety of trigger paths, generally broken into a few different categories:

1. **Dedicated analysis triggers**, such as `HT650_DisplacedDijet60_Inclusive`. These triggers are designed to target a particular final state or set of model parameters, and goal of the trigger is to collect data which was previously rejected by the existing trigger menu.
2. **General-purpose triggers**, such as `HT1050` and `PFMET120`. These triggers serve as general-purpose tools for validation, control samples, and for analyses which have no dedicated trigger.
3. **Backup and validation triggers**, such as `PFMET140`. These triggers don't collect any new data— in this case, a separate trigger with $MET > 120$ exists to already collect data satisfying $MET > 140$. However, the `MET140` trigger can be used to validate the performance of the `MET120` trigger by examining events which satisfy both trigger paths. Additionally, if the `MET120` trigger has a problem, the `MET140` trigger can be used as a backup with safer, higher thresholds to avoid loss of physics data.
4. **Prescaled triggers**, such as `HT350`. Some triggers are “prescaled” in which an automatic rate filter is applied to the trigger path. For example, a prescale of 100 records only 1% of triggers satisfying the trigger path. Key analysis triggers are typically unprescaled, but prescaled triggers allow the HLT to record data samples in regions of very high trigger rate, such as in events with low total energy or maximum jet P_T .

An HLT path begins by “seeding” on an L1 trigger path, offering basic selections and an initial dataset. After an HLT path satisfies its corresponding L1 trigger, it accesses key

event information such as particle tracks, energy, and ID. This information is reconstructed once per event and used to produce object collections such as jets, electrons, and muons for triggers to operate on. An HLT path uses this event and object information to perform its trigger algorithm. The implementation of the HLT path for the CMS analysis in this dissertation is described in detail in Section 3.1.

2.3.3 Data Streams

Depending on the trigger rates and data requirements, events triggered by the HLT can be recorded into one of a few different data streams.

The standard data stream is reconstructed immediately as newly triggered data at HLT is recorded. HLT paths populating this data are called “prompt” triggers. This data stream has the smallest overall rate, but has the advantage of providing immediate access to data as it’s being recorded.

An alternative to prompt data is the parking stream in which triggered data is reconstructed only at the end of the year when no new prompt data is being recorded. HLT paths populating this dataset are called “parking” triggers. This parking data is still recorded at the same time as prompt data, but as there is insufficient computing power to reconstruct all data promptly, this alternative stream is “parked” for later use. This delayed reconstruction allows for higher HLT rates and lower thresholds, increasing sensitivity to softer and more challenging models.

Lastly, a lightweight “scouting” datastream exists for events which don’t need full reconstruction. No raw detector information is saved for these events, only high level objects such as jets and leptons which are reconstructed online at the HLT. These extremely lightweight events allow for the lowest, most ambitious trigger thresholds, but are only appropriate for exploratory analyses or searches which are compatible with minimal low-level detector information.

2.3.4 Data Formats




Figure 2.13: A visualization of the different event sizes of CMS data formats [48].

The largest and most robust data format in CMS is the RAW format, containing all detector hit, energy, and timing information. However, few analyses access RAW information directly; the most detailed event format in CMS is called RECO, containing all the detector hit information in RAW as well as reconstructed high-level physics objects and event observables.

The level of information in RECO is unnecessary for most analyses, which may use objects like tracks but don't need individual tracking hits. For this reason, in Run 1, the Analysis Object Data (AOD) format was created, dropping a subset of the lowest level detector information but retaining all high level reconstructed physics objects.

A much more significant slimming came from the miniAOD format, standard for CMS Run 2 analyses. Rather than each analysis creating ntuples from scratch from AOD, miniAOD is a centrally produced ntuple containing only the most commonly used reconstructed object information such as leptons, photons, and jet collections with dedicated selections for general

use [49].

Similar to miniAOD is nanoAOD, an even slimmer, more optimally packed collection of only the most essential high level reconstructed objects. From the beginning of 40 TB/s—about 10% of the total rate of the entire internet— an analysis consisting of only nanoAOD samples running over the entirety of the prompt and parking HLT will accumulate only 40 MB/s, about five times slower than a typical home internet connection.

3 CMS Delayed-Jets Trigger

3.1 Trigger Strategy

LLP analyses have three primary handles to identify a long-lived particle, with each method's efficacy depending on lifetime:

- **Displaced Decays** ($0.1 \text{ mm} < c\tau < 0.1 \text{ m}$): If a long-lived particle decays inside the detector with a decay vertex or track displacement measurably far from the primary collision vertex, the displacement can be used to measure a decay lifetime. This method requires particle tracking information to identify the particle tracks and vertices, so it's limited to shorter lifetimes in which the LLP decays far enough inside the tracker to reconstruct a displacement, but high tracking precision can be used to target models with very low lifetimes.
- **Delayed Decays** ($0.1 \text{ m} < c\tau < 100 \text{ m}$): If tracking information is unavailable, such as if a particle decay occurs outside or near the edge of the tracker, an effective displacement can be computed as a result of the increased path due to a displaced decay, as in Figure 1.5. Additionally, there is an enhancement in sensitivity for slow-moving LLPs which is not present using purely tracking information; for example, a slow-moving LLP which produces two highly boosted decay products will show minimal track displacement, but can be measured using the delayed decay time, as in Figure 3.1.
- **Missing Energy** ($c\tau > 100 \text{ m}$): If the lifetime of an LLP is so large that it doesn't decay inside the detector, these signatures can be found using MET. However, as these searches are generic and indistinguishable from models producing stable exotic particles, the phase space for MET-only LLP signatures is highly constrained.




Figure 3.1: For LLPs with narrow mass splittings, jet time delay is dominated by the LLP’s low β . For the $M_H = 1$ TeV, $M_{LLP} = 475$ GeV shown here, $\beta = 0.3$ and $\beta\gamma c\tau \approx 0.3$ m. Figure credit: Dariush Imani

CMS analyses have been using dedicated displaced jet triggers to target LLPs since Run 1 [50]. MET triggers, including targeted triggers which require jets or leptons in the final state, have also been available since Run 1. However, although timing has been used for LLP searches since Run 1 [51], no dedicated delayed jet trigger at L1 or HLT was present in CMS for Run 1 or Run 2.

The first CMS analysis to use ECAL timing to search for delayed jets was EXO-19-001 in Run 2 [12]. With no online delayed jet trigger available, the analysis used a $MET > 120$ GeV trigger as its baseline dataset, with an offline requirement of $MET > 300$ GeV. Offline jet delay timing requirements of $t_{jet} > 3$ ns were applied to events satisfying this MET

requirement.




Figure 3.2: The Run 2 delayed jet analysis demonstrated that new sensitivity could be gained for lifetimes greater than 0.1 m [12].

This Run 2 delayed jet analysis demonstrated that ECAL timing could offer unique sensitivity to LLP models (Figure 3.2), but the trigger limitation constrained the analysis to models which contained high MET in the final state. This motivated the development of Run 3 delayed jet triggers with no MET requirement and much lower ECAL timing thresholds.

3.2 Run 3 Delayed Jet Trigger

New to Run 3, a slate of delayed jet triggers was developed to target new sensitivity for delayed jet models. The key properties used to separate each of eight primary signal triggers are summarized below:

- **Trigger Seed:** Triggers are seeded either with a threshold on H_T or on jet shape. For the H_T thresholds, an HLT seed with a threshold at $H_T > 280 \text{ GeV}$ (hltL1sHTT280er)

is combined with an online selection of $H_T > 430$ GeV to form the HT430 seeded triggers. Online H_T is calculated using a sum over all jets with $P_T > 40$ GeV and satisfying $|\eta| < 2.5$, with barrel jets ($|\eta| < 1.48$) and endcap jets ($1.48 < |\eta| < 2.5$) reconstructed separately. The H_T selection is placed above the H_T seed value to avoid turn-on effects and ensure the trigger is in the plateau region. Similarly, another HLT seed with the same threshold at $H_T > 280$ GeV (hltL1sHTT280er) is used for an online selection on $H_T > 350$ GeV. To get the lowest possible reach in H_T , an L1Tau seed (hltL1sTauVeryBigOR) is used. This approach exploits the fact that delayed jets appear narrower than prompt jets due to displaced decays that limit the spread of delayed jets. This is illustrated in Figure 3.3.

- **Jet Delay:** The ECAL jet time is computed online by taking the pt-weighted average of all cells within the timing window (± 12.5 ns). A more precise calculation of the time is performed later offline (see Section 4.3), but this initial time is used as a threshold to tune the rates of each trigger path. The rates for each trigger path are highly sensitive to the timing threshold used. For each signal trigger path, two backup paths with jet time thresholds increased by 0.25 ns and 0.5 ns were implemented to ensure that changes to timing calibrations between runs (as e.g. ECAL crystal transparency increases) could be accounted for by choosing different sets of triggers between the three timing threshold options.
- **Jet Multiplicity:** With jet delays serving as the main handle to reduce trigger rates, single jet triggers can be used at high jet delay times. This is an advantage of the delayed jet trigger design, as many other LLP searches rely on dijet signatures to reduce backgrounds and trigger rates. Dijet triggers also exist for this analysis, with significantly lowered timing thresholds.
- **Jet Isolation:** A “Trackless” requirement is included in some triggers to reduce the trigger’s timing threshold. This isolation cut requires that ≤ 2 prompt tracks of $P_T > 0.5$ GeV satisfy $IP_{xy} < 1$ mm. Triggers without this requirement are labeled as

“Inclusive”.




Figure 3.3: The displaced decay of a delayed jet limits the space the jet can spread into, giving it similar characteristics to narrow tau jets.

Additionally, there is a set of fixed requirements shared between triggers:

- **Jet $P_T > 40$ GeV:** This requirement helps reduce the effect of ECAL noise and is the same as the offline selection.
- **Jet track $\Delta R < 0.4$:** A standard requirement, all tracks within a ΔR of 0.4 are used as constituents for the jet.
- **ECAL Cell Energy > 0.5 GeV:** To reject ECAL noise, only ECAL cells with energy greater than 0.5 GeV are considered for online timing calculations.
- **Number of ECAL Cells > 5 :** To reject ECAL noise and spikes, delayed jets must consist of activity in greater than 5 ECAL cells passing the energy cut.

Finally, there are two slates of triggers used: prompt and parking triggers. The parking triggers operate with the most inclusive set of trigger requirements and at lower thresholds than the equivalent prompt triggers. These triggers were implemented for the 2023Cv3 dataset and onwards, representing a total luminosity of 129.9 fb^{-1} in 2023 and 2024. The first analysis presented in the dissertation focused solely on prompt triggers; parking triggers

will be used for a future analysis. A summary of all delayed jet triggers used in this analysis, as well as when the triggers were implemented, is given in Tables 3.1 and 3.2. A complete set of backup trigger paths, added to the HLT in addition to these signal triggers for validation and rate control, is shown in Appendix Table A.1.

Trigger	Since
HT430 seed, Double jet, trackless, 0.5 ns HLT_HT430_DelayedJet40_DoubleDelay0p5nsTrackless	2022F (360332)
HT430 seed, Double jet, inclusive, 1 ns HLT_HT430_DelayedJet40_DoubleDelay1nsInclusive	
HT430 seed, Single jet, trackless, 1 ns HLT_HT430_DelayedJet40_SingleDelay1nsTrackless	
HT430 seed, Single jet, inclusive, 2 ns HLT_HT430_DelayedJet40_SingleDelay2nsInclusive	
HT350 seed, Single jet, inclusive, 3 ns HLT_HT350_DelayedJet40_SingleDelay3nsInclusive	2023C-v3
L1Tau seed, Double jet, trackless, 1 ns HLT_L1Tau_DelayedJet40_DoubleDelay1nsTrackless	2022F (361971)
L1Tau seed, Double jet, inclusive, 1.25 ns HLT_L1Tau_DelayedJet40_DoubleDelay1p25nsInclusive	
L1Tau seed, Single jet, trackless, 2.5 ns HLT_L1Tau_DelayedJet40_SingleDelay2p5nsTrackless	
L1Tau seed, Single jet, inclusive, 3.5 ns HLT_L1Tau_DelayedJet40_SingleDelay3p5nsInclusive	

Table 3.1: Prompt Triggers (DisplacedJet dataset)

Trigger	Since
HT430 seed, Single jet, inclusive, 1 ns HLT_HT430_DelayedJet40_SingleDelay1nsInclusive	2023C-v3
HT350 seed, Single jet, inclusive, 1.5 ns HLT_HT350_DelayedJet40_SingleDelay1p5nsInclusive	
L1 Tau seed, Single jet, inclusive, 2.5 ns HLT_L1Tau_DelayedJet40_SingleDelay2p5nsInclusive	

Table 3.2: Parking Triggers (ParkingLLP dataset)




Figure 3.4: Flow of the single jet, 1 ns, trackless delayed jet HLT path. Colored nodes indicate modules unique to the trackless variant.

3.3 Trigger Implementation

The representation of the single jet, 1ns trackless delayed jet trigger in the HLT menu is show below in Figure 3.4. The inclusive version of the trigger, discussed in Section 3.2, is constructed with the highlighted tracking modules removed. Below, I discuss how each module contributes to the trigger path.

3.3.1 Module Description

Preamble

HLTBEGINSequence + hltL1sVoHTT380 + hltPreHT430DelayedJet40SingleDelay1nsTrackless + hltPixelTrackerHVOn + hltStripTrackerHVOn + HLTAK4CaloJetsSequence

This preamble registers the trigger and initializes L1 and HLT objects, seeds the trigger with an L1 threshold of $H_T > 360$ GeV, registers HLT prescales (set to 1 for all signal triggers), checks if the pixel and strip trackers are powered on, and prepares the AK4 calo jets used downstream.

HLT energy and jet-quality filters

hltHtMht + hltHT430 + hltEmFraction0p01To0p99CaloJetSelector + hltSingleCentralCaloJetpt40 + hltCentralCaloJetptLowPtCollectionProducerSingle

This block is responsible for applying energy requirements to the event and jet objects. Here we compute H_T from jets with $P_T > 40$ GeV and $|\eta| < 2.5$; apply an online $H_T > 430$ GeV threshold; enforce jet $0.01 < EM_{frac} < 0.99$, $N_{Cells}^{90\%} \geq 3$, and $|\eta| < 2.0$ requirements to identify well-behaved central jets; applies an online $P_T > 40$ GeV cut; and produces the working calo jet collection for the rest of the path.

Trackless requirement

Trackless-only: HLTBTagPixelAndStripSetupForInclusiveDisplacedJets + HLTIterativeTrackingIteration0ForDisplaced + HLT2PromptTrackRequirementIter0DisplacedJetsLowPtSingle + hltDisplacedHLTCaloJetCollectionProducerLowPtSingle + HLTIterativeTrackingIter12ForDisplaced + HLT2PromptTrackRequirementIter12DisplacedJetsLowPtSingle

This segment applies a trackless requirement, enforcing a requirement on jets that $N_{track,prompt} \leq 2$ for prompt-identified tracks of $P_T^{track} > 0.5$ GeV and impact parameter

$IP_{xy} < 1$ mm. The block first prepares pixel/strip data and selects up to eight central jets. Then iterative tracking is performed, beginning with tight tracking requirements and iterating over looser collections only if the event passes tight selections to minimize runtime for CPU-intensive loose tracking algorithms. Finally, the module merges the iterative tracking collections and retains jets satisfying a requirement of ≤ 2 prompt tracks. For the full sub-module breakdown and flow diagrams, see Appendix A.1 (Figures A.1–A.5).




Figure 3.5: Shown is the impact on HLT menu runtime before and after adding ECAL timing. A reduction to 1% impact was achieved by adding an energy requirement to the ECAL cells used to compute jet time.

Timing filters

```
hltCaloJetTimingProducerSingle + hltCaloJetTimingFilterSingle1ns + HLTEndSequence
```

The final modules compute jet timing from ECAL cells satisfying $E > 0.5$ GeV, $|t| < 12.5$ ns, $\Delta R < 0.4$, and passing cell quality requirements, then applies the 1 ns selection for jets satisfying $N_{Cells} > 5$. Lastly, the result of the trigger is returned to the HLT.

The code responsible for these timing algorithms was new to the HLT for Run 3. A requirement of a lower bound on ECAL cells used to compute timing of

$E > 0.5$ GeV resulted in a 1% impact on HLT timing performance, shown as “partial timing” in Figure 3.5. Full code was merged into the central CMS software repository (CMSSW) and is accessible at `HLTrigger/JetMET/plugins/HLTJetTimingProducer.h` and `HLTrigger/JetMET/plugins/HLTJetTimingFilter.h`.

3.4 Trigger Performance




Figure 3.6: Trigger Efficiency for Run 2 (HT1050, black) and Run 3 (Inclusive, red; Trackless, blue) triggers. Efficiency is binned by event H_T .

We use simulation and data to benchmark the performance of the new Run 3 triggers. As detailed in Section 1.3, the simulation is based on a Twin Higgs model $H \rightarrow XX \rightarrow bbbb$. As this model has no MET in its final state, the best pre-existing, unprescaled trigger at the Run 2 HLT targeting this model was a generic $H_T > 1050$ GeV trigger. Figure 3.6 shows the significant gain in corrected H_T obtained by the new slate of Run 3 delayed jet triggers.

In addition to the triggers seeded on $H_T > 430$ GeV at L1, we obtain further gains in H_T reach by using L1 seeds on a reconstructed τ . At the cost of higher timing thresholds to ensure trigger rate control, this removes all H_T requirements from the final state as demonstrated in

Figure 3.7.




Figure 3.7: Trigger Efficiency for Run 3 HT430 and L1Tau triggers. For events with sufficiently large delays as in the $c\tau = 10$ m used for this validation, seeded on τ at L1 allows for improved reach in H_T without loss of efficiency in the plateau. Note the decrease of efficiency for high H_T due to strong H_T requirements enforcing prompt decays such that enough jets can be reconstructed to produce a large value of H_T . Jet requirements here are $P_T > 40$ GeV, $N_{Cells} > 5$, $|\eta| < 1.48$, and $t_{jet} > 2$ ns, mirroring online jet requirements for triggers at 2 ns.

Trigger performance in data has been measured for 2022 and 2023 when the triggers were first deployed to the HLT. The performance of these triggers in data is shown in Figure 4.25.

The trigger efficiency is calculated by applying offline jet and event requirements, then measuring the efficiency in which the online trigger fired for events satisfying those requirements. To access events with unbiased timing to measure this turn-on, an online trigger of $MET > 120$ was used with an offline requirement of $MET > 200$ GeV. Offline jet requirements include $P_T > 50$ GeV, $N_{Cells} > 8$, and $E_{EM} > 25$ GeV. For the H_T triggers, an offline requirement of $H_T > 430$ GeV was also used. Trigger turn-ons at 2 ns for H_T -seeded paths

and 3.5 ns for τ -seeded paths are observed, although differences between online and offline jet requirements generate a shallow turn-on and plateau inefficiencies. Studies measuring how differences between these online-offline requirements impact turn-on performance are ongoing.




Figure 3.8: Trigger efficiency measured in 2022 and 2023 for single jet, inclusive H_T -seeded triggers at 2 ns (left) and τ -seeded triggers at 3.5 ns (right).

In 2024, the total trigger rate for all new delayed jet triggers was 37 Hz for the prompt data stream and an additional 77 Hz in the parking data stream. Prompt and parking trigger pairs are designed with minimal timing overlap such that prompt and parking data reconstruction contains few repeat events; for example, the single inclusive trigger at 1 ns in parking and 1.5 ns in prompt actually contains a timing window of 1 – 1.5 ns in parking. New ideas are being explored towards the end of Run 3 about deploying triggers into the scouting data stream as well.

4 CMS Delayed-Jets Analysis

4.1 Data and MC Samples

The data samples were collected in 2023 and 2024 by the CMS detector in pp collisions at an energy of $\sqrt{s} = 13.6$ TeV, yielding a total integrated luminosity of 136.9 fb^{-1} . The 2023 dataset includes eras 2023B-D for a total luminosity of 27.9 fb^{-1} , and the 2024 dataset includes eras 2024B-I for a total luminosity of 109.1 fb^{-1} . The first Run 3 delayed jet triggers (on HT430 seeds) went live in run 360332 (2022F), the L1Tau-seeded triggers went live in run 361971 (2022F), and the HT350-seeded prompt triggers as well as all parking triggers went live at the start of 2023C-v3. The analysis used an inclusive union of all Run 3 delayed jet triggers, described in more detail in Section 3.1.

Cosmic backgrounds were examined by merging events skimmed over cosmic triggers in 2023 data with events skimmed over JetMET triggers in 2023 data, creating a cosmic-enriched background sample. The cosmic sample is from a dedicated cosmic run in Run2023F. The data sample used for the study was collected during 2023Cv1 in the JetMET dataset. Though these are not the same eras, since the purpose of the sample is to study combined cosmic data triggered under JetMET conditions, it forms an adequate representative sample for what a randomly overlapping cosmic may appear like in 2023 data.

Monte Carlo samples were generated privately in the Twin Higgs model using Pythia8 under the CP5 tune, decaying to four b-jets in the final state. The mass of the Higgslike moderator was varied from 125 GeV to 2 TeV, masses of the LLPs were varied from 50 GeV to $m_H/2 - 25$ GeV, and lifetimes of the LLPs were varied from $c\tau = 1$ mm to 100 m. The production model was gluon-gluon fusion, though alternative production modes in VBF were also explored. The nominal final state is four b-jets, and all MC studies have been performed using this final state, though 4-tau final states have also been produced.

4.2 Event Selection

Event selection followed a baseline approach motivated by EXO-19-001, but with a few key modifications. First, with no MET requirement at the trigger level, MET requirements are removed. Events are required to satisfy the following requirements:

- **CMS Event Filters:** Standard MET event filters are used to clean all events. These include:
 - Flag_globalSuperTightHalo2016Filter
 - Flag_EcalDeadCellTriggerPrimitiveFilter
 - Flag_BadPFMuonFilter
 - Flag_BadPFMuonDzFilter
 - Flag_hfNoisyHitsFilter
 - Flag_eeBadScFilter
- **Delayed Jet Trigger:** Events are selected using an inclusive OR over all delayed jet triggers, as summarized in Section 3.1.
- **Beam Halo Filter:** As detailed in Section 4.4.1, all events failing beam halo requirements are removed and placed into a beam halo control sample.
- **Cosmic Muon Filter:** Similarly, as detailed in Section 4.4.2, a cosmic muon filter is applied, and all events failing the DT and CSC activity requirements are removed and placed into a cosmic muon control sample.
- **Satellite Bunch Filter:** A veto is placed on events containing jets which otherwise satisfy all requirements but failing the PVDR jet isolation requirement, as described in Section 4.2.2. Events containing exactly one jet failing this prompt tagging requirement are used to form a satellite collision control sample. Background estimates using this sample are described in Section 4.4.3.

4.2.1 Event Bins

For events satisfying all standard event-level requirements, we apply three further bins.

First, we bin in event H_T . We compute H_T by summing all jets of $P_T > 40$ GeV within $|\eta| < 2.5$. This method is consistent with how H_T is computed in the trigger online. A comparison between this corrected HT40 and the standard calculation of H_T computed with no restriction on η is shown in Figure 4.1.




Figure 4.1: Summing calo jets in a restricted range of η ensures that out-of-time jets can be consistently included in the calculation of H_T , giving a more robust measure of the event H_T .

As the highest trigger threshold is set at $H_T = 430$ GeV, signal bins are chosen as $500 < H_T < 700$, $700 < H_T < 1000$ and $H_T > 1000$ in order to avoid effects in the plateau region of the trigger turn-on. A control bin is chosen as $H_T < 500$ GeV. A comparison of signal and background in H_T is shown in Figure 4.2. As the data H_T distribution peaks near 400 GeV, models with a BSM Higgs mass of 500 GeV and above gain sensitivity in higher H_T bins.




Figure 4.2: Comparing a variety of signal models to H_T , summed over $|\eta| < 2.5$, shows the benefit high signal masses gain from multiple H_T bins.

Another event-level bin is on a requirement or veto of an ISR jet, targeting large backgrounds generated from satellite bunch collisions near $t_{jet} = 5$ ns. An effective method to veto this background is to require a prompt ISR-like jet satisfying requirements summarized in Table 4.1.

ISR Selections
$ \eta < 2.5$
$P_T > 60$
$PVDR < 0.3$
$-1 < t_{jet} < 0.75$
$t_{jet}^{RMS} < 1.5$
Neutral EM Energy Fraction < 0.9
Charged EM Energy Fraction < 0.8
Charged Hadronic Energy Fraction > 0.2
Neutral Hadronic Energy Fraction < 0.99
EM Energy > 20
of ECAL Cells > 10
Hadronic Energy > 10

Table 4.1: Selections used to identify prompt, ISR-like jets.

$|\eta| < 2.5$ and $P_T > 60$ GeV are basic energy requirements and constrain the ISR jet to the barrel or endcap. The P_T requirement is placed higher than for standard jets, as the P_T distribution is shifted with respect to caloJets as in Figure 4.3.




Figure 4.3: (Left) A $P_T > 60$ GeV requirement for PFJets is equivalent to a $P_T > 40$ GeV requirement for caloJets, as is used for signal selections. (Right) This is supported by the gaussian distribution in the P_T difference centered at approximately -20 GeV.

$PVDR < 0.3$ and $-1 \text{ ns} < t_{jet} < 0.75 \text{ ns}$ constrain the ISR jet to the center of the prompt timing distribution, with the cut at $t_{jet} < 0.75 \text{ ns}$ placing the jet outside signal or validation timing regions. Charged and neutral PFJet energy fraction requirements are consistent with central recommendations for Jet ID. Finally, $EM \text{ Energy} > 10 \text{ GeV}$, $Number \text{ of ECAL Cells} > 10$, and $Hadronic \text{ Energy} > 10 \text{ GeV}$ veto noise and ensure jet activity is sufficient to reconstruct a reliable jet time.

4.2.2 Jet Selection

Once event selections are applied, signal jets are classified using a series of energy and timing requirements. Jet selections are summarized in Table 4.2.

We consider only barrel signal jets for this analysis. We require hadronic jets in the final state, so we use electromagnetic and hadronic energy selections to reject noise and soft, prompt jets. These requirements are jet $P_T > 40 \text{ GeV}$, $ECAL \text{ Energy} > 20 \text{ GeV}$,

HCAL Energy > 10 GeV, and we require that a jet is composed of greater than 10 ECAL cells with energy deposits above 0.5 GeV. These minimum energy requirements ensure that timing is well-measured in the ECAL, but we additionally apply a requirement of $t_{RMS} < 1.5$ ns to ensure good timing performance.

To remove beam halo, we apply a requirement that 90% of the total jet EM energy is distributed across greater than 7 ECAL cells. This targets beam halo, as most beam halo deposits result in the majority of energy concentrated in a small number of cells.

Jet Selection Criteria	Backgrounds Targeted
Jet $p_T > 40$ GeV	Prompt jets (1), ECAL noise (2)
$ \eta < 1.48$	Beam halo (4)
Total EM energy > 20 GeV	Prompt jets (1), ECAL noise (2)
Number of ECAL-cell hits > 10	ECAL noise (2)
Jet time RMS < 1.5 ns	Prompt jets (1), ECAL noise (2)
Hadronic energy > 10 GeV	ECAL noise (2), Beam halo (4)
Hadronic energy fraction > 0.2	ECAL noise (2), Beam halo (4)
$n_{90} > 7$	ECAL noise (2), Beam halo (4)
CSC EM-energy fraction < 0.6	Beam halo (4)
Jet η/ϕ shape ratio < 1	Beam halo (4)
Minimum PV $\Delta R > 0.3$	Prompt jets (1), Satellite bunches (3)

Table 4.2: Selection criteria and the dominant background types they reject. Background types: (1) Prompt jets, (2) ECAL noise, (3) Satellite bunches, (4) Beam halo

Beam halo is further rejected by applying requirements that there is significant hadronic energy (Hadronic Energy Fraction > 0.2), a maximum requirement on the ECAL energy which is matched to hits in the CSCs (CSC EM Energy Fraction < 0.6), and a shape requirement using the distribution of the ECAL cell hits in η and ϕ (Jet η/ϕ Shape Ratio < 1). The selection in hadronic energy fraction is consistent with the Run 2 delayed jet analysis, while the selection in CSC energy fraction is newly optimized using signal and

background distributions enriching in beam halo, as in Figure 4.4. The CSC Fraction and Shape Ratio requirements are used in the ABCD method to estimate beam halo backgrounds in Section 4.4.1.




Figure 4.4: Data enriched in beam halo possesses a CSC energy fraction distributed from 0.6 to 1.0, peaking at 1.0. Signal-like standard QCD jets peak near 0, with a flat component at higher CSC fraction values. To optimize signal efficiency, a threshold is chosen at 0.6, down from 0.8 in the Run 2 analysis.

Finally, to reject prompt jets, we define a new isolation variable called Primary Vertex ΔR (PVDR). We require that no track associated to a prompt vertex and satisfying track $P_T > 3$ GeV lies within a $\Delta R < 0.3$ cone of each jet, as in Figure 4.7. This width in ΔR , approximately the physical width of the jet, ensures that all hard tracks which compose the jet are not associated to any prompt primary vertex. This isolation requirement provides a separate handle using the tracker to identify prompt-like jets, even if the collisions are out-of-time. In Figure 4.5, we compare the performance of PVDR to P_T Track Fraction, a similar tracking variable used to tag prompt jets in the Run 2 delayed jet analysis. Figure 4.6 demonstrates the performance improvements gained using PVDR for any choice of threshold.

Backgrounds from prompt tails and satellite bunch collisions are estimated using an ABCD method in PVDR and jet time as described in Section 4.4.3.




Figure 4.5: Both the new PVDR isolation variable (left) and the Run 2 P_T fraction of tracks associated to primary vertices (right) provide discrimination between prompt QCD and signal jets. However, PVDR offers a more direct handle by isolating any high-energy pointing tracks within $dR < 0.3$ of the jet, ensuring best purity.




Figure 4.6: PVDR shows a modest improvement in signal efficiency and background rejection, a critical gain as prompt jet tails identified by PVDR form the leading background for the search.

In addition to binning in H_T and an ISR requirement, we bin our signal region in monojet and dijet categories. This binning reduces backgrounds from prompt tails by $O(10^2)$, allowing

for significant enhancement in sensitivity at low jet delay times. Dijet efficiencies for signal are shown compared to their monojet equivalent in Figure 4.8 and Table 4.3. A summary of all event- and jet-level requirements is shown in the event flow diagram in Appendix Figure A.16.




Figure 4.7: The tracker can be used as a handle independent from timing to identify prompt jets. Prompt Vertex ΔR (PVDR) is computed using the nearest prompt-tagged track with at least 3 GeV (red). Other softer pileup tracks are ignored (light blue).

M_H [GeV]	m_X [GeV]	$c\tau$ [m]	Efficiency [%]
1000	475	1	38
1000	475	0.1	21
1000	475	10	14
400	175	1	5.8
1000	225	1	8.5
Background			1.0

Table 4.3: Selection efficiency for the dijet selection as a fraction of monojet events, integrated above $t > 1$ ns. A dijet requirement reduces signal yield above 1 ns by around a factor of 10, but also reduces backgrounds by a factor of 100.




Figure 4.8: Requiring two signal jets condenses the width of the timing distribution, with the size of the effect increasing as the signal lifetime decreases. This effect is increased for models with lower signal efficiency.

4.3 Timing Performance

ECAL timing, the key variable driving analysis sensitivity, requires extensive calibration as a function of run, jet energy, and detector location. A variety of sources lead to timing drift: variations in clock distribution between different regions of the ECAL, uncertainty in the central timing calibration, and, most significantly, ECAL crystal transparency changes [52].

4.3.1 Run-based Timing Corrections

In the CMS ECAL, pulse timing T in the PbWO_4 crystals is computed using the following [52]:

$$T = T_{Max} + T_{Shape}(A_{pulse}) + T_{Correction} + T_{MC} \quad (4.1)$$

where:

- T_{Max} is the maximum pulse time extracted from the ratio method (see Appendix A.3);
- $T_{Shape}(A_{pulse})$ is a fixed timing correction as a function of pulse amplitude, derived from fixed pulse templates, where the pulse amplitude is determined using the multifit method (see Appendix A.4);
- $T_{Correction}$ is a central timing correction for each cell computed from cell timing distributions (updated a few times per week, sometimes daily near the start of new runs);
- T_{MC} is a fixed, global timing offset used to obtain better agreement with MC.




Figure 4.9: Narrower, faster pulses in crystals with radiation damage result in negative shifts to the value of T_{Max} [53].

The pulse templates are recalibrated around once per week during running [54]. However, as radiation damage accumulates during beam operation, ECAL crystals form “color centers”, radiation-damaged regions of increased opacity, decreasing the crystal transparency and

leading to lower overall light yield. Pulse width measurements in irradiated crystals show a faster, narrower pulse for highly radiation-damaged crystals with low transparency, as in Figure 4.9.




Figure 4.10: (Top) As ECAL crystal transparency decreases with increasing radiation damage, the pulse becomes faster and more narrow, changing the pulse shape [55]. (Bottom) Decreased crystal transparency reduces the total path length of the scintillation light, reducing the width of the pulse timing distribution [53].

This pulse narrowing can be interpreted by considering the reduced path length of scintillation photons in crystals of low transparency. In Figure 4.10, this effect is illustrated and the pulse width is demonstrated to be a function of the induced absorption coefficient μ_{ind} , a parameterized estimate of the transmission loss given by:

$$\mu_{ind} = \frac{1}{L} \ln \frac{T_{before}}{T_{after}} \quad (4.2)$$

where $L = 22$ cm is the length of the crystal and T_{before}, T_{after} are crystal longitudinal light transmission coefficients measured at 420 nm before and after irradiation.

The narrowing of the pulse shapes over time has the effect of biasing T_{Max} negative, causing a systematic shift of up to 2 ns per year before timing corrections are applied [52]. Over time, crystals also undergo “annealing”, a removal of these radiation-damaged color centers that restores crystal transparency. Therefore, during beam-off periods and some periods of low beam luminosity, crystal annealing causes another ECAL calibration drift opposite to the effect during normal operation.




Figure 4.11: Run-level timing calibrations are shown for eras 2023Cv1 through 2023Dv2. A decreasing trend in the timing correction later in runs suggests an increasing trend in the fitted timing mean.

To correct for these drifts, we select a sample enriched in prompt jets by requiring all standard jet selections from Table 4.2 except on PVDR, as to enhance standard prompt, pointing jet activity. We also require all event-level selections summarized in Section 4.2

except for the N_{jets} requirement, as this selection enforces delayed activity.

A gaussian fit is made to all jets passing event- and jet-level requirements, and the mean timing value extracted from this fit is used to derive a timing correction for each run. Results for timing correction means in 2023 are shown in Figure 4.11. The full fits are shown in Figure 4.12.




Figure 4.12: Run-based timing calibrations extracted from the era level are used to generate an era-level correction to signal acceptance. Intrarun timing variation is propagated as a systematic on timing.

In addition to timing drifting as a function of era, we also find that later runs in an era tend to have poorer timing resolution. This is illustrated by the increasing peak fit width in Figure 4.12 and Figure 4.13.




Figure 4.13: In addition to shifts in timing means, different run eras also show significant variation in timing resolution.

4.3.2 Spatial Timing Corrections

Separate from run-drifts in the timing, timing can also vary spatially across ECAL stations. This is due to a mixture of anisotropies in radiation exposure as a function of crystal η as well as systematic shifts in detector clock times [52]. To correct for these effects, a spatial timing correction is derived by fitting the timing means of each ECAL station. Stations are divided such that $i_\phi \in \{1, 2, \dots, 20\}$ and $i_\eta \in \{1, 2\}$, and a spatial correction is extracted from the timing fit for each station across all runs in each era. As the magnitude of spatial timing corrections is subdominant as compared to run corrections, spatial corrections are measured across full eras to ensure sufficient statistics can be collected from each crystal. An

example of the typical spatial variation extracted from an ECAL DPG laser transparency test is shown in Figure 4.14.




Figure 4.14: Spatial timing map collected by the CMS ECAL DQM in 2023.

4.3.3 Energy Corrections

In addition to spatial and temporal timing corrections, we perform global, year-level energy corrections to timing. Due to a timewalk effect in the ECAL where low energy deposits in ECAL cells are recorded later in time as compared to high energy deposits—likely due to pulse shape effects—jets with low ECAL energy show a positive timing bias. This effect is measured and corrected for uniformly across each run year, as shown for 2023 in Figure 4.15. This year-level approach for jet energy corrections to time is chosen to be consistent with how standard year-level jet energy corrections are applied to jet energy.




Figure 4.15: Energy-based corrections to timing remove the effects of timewalk, an effect which causes a delay for low-energy ECAL deposits.

4.3.4 Era-dependent backgrounds

Generally, as timing performance degrades due to ECAL transparency changes later in runs, timing resolution also increases. As the leading source of backgrounds are tails from the prompt jet timing distribution, eras with large timing resolution contain more background events passing all jet and event selections per lumi, as in Table 4.4. This shows the number of background events is not evenly distributed per era, and cleaner eras will provide a more sensitive handle for lower lifetime models.

Year+Era	Lumi (fb^{-1})	Events Passing Selections	Events per Lumi
2023Cv1	4.5	12034	2.62×10^3
2023Cv2	1.3	2665	2.04×10^3
2023Cv3	1.3	25537	1.91×10^4
2023Cv4	11.1	293414	2.61×10^4
2023Dv1	8.1	3356	4.09×10^2
2023Dv2	1.6	1098	6.67×10^2

Table 4.4: Events passing selections in each run era. The events per lumi concentration increases towards the end of an era, where the degraded resolution due to transparency changes causes broader prompt tails.

4.3.5 Corrections to MC Timing

While run- and spatial-timing corrections are applied to data to center the timing distribution at zero, MC timing distributions are inherently centered at zero with no artificial shifts to timing means. The only corrections measured in data and applied to MC are energy-based timing corrections, expected to correct for inherent effects of timewalk present in translating RAW MC energy deposits to an ECAL cell timing measurement.

However, since the online trigger time is a fixed value, the effective online timing threshold is not fixed, but instead given by

$$t_{\text{online}} = t_{\text{design}} - t_{\text{era}} \quad (4.3)$$

where t_{design} is the design target timing threshold and t_{era} is the mean era timing fit. For eras with a significant positive jet time mean, this online trigger threshold is effectively smaller, leading to higher backgrounds and higher signal efficiency. Similarly, for eras with negative jet time means, the online trigger threshold is higher, leading to lower backgrounds and lower signal efficiency.

To account for these corrections to online trigger efficiency, as illustrated in Figure 4.16, two adjustments to MC trigger timing thresholds are used.




Figure 4.16: To effectively model the impacts of online timing shifts on MC trigger efficiency, signal is generated separately for each data era with corresponding timing corrections applied to the trigger’s timing threshold.

First, there is an overall increase in the MC trigger timing threshold of 0.25 ns for dijet triggers and 0.5 ns for monojet triggers. This threshold helps to account for the turnon of timing triggers and ensures that trigger efficiency as a function of jet time operates in the trigger plateau region. In a subsequent analysis, this fixed threshold will be replaced with a more complete implementation of the trigger timing turn-on to target lower lifetime models.

Second, MC is generated for each era, and the effective post-correction online trigger timing threshold t_{online} is used as the baseline value for MC. In this way, MC will have a higher trigger efficiency for eras with a lower value of t_{online} and vice-versa.

The total effect of these two adjustments sets the effective MC trigger timing threshold at $t_{online} + 0.25$ ns for dijet triggers and $t_{online} + 0.5$ ns for monojet triggers. MC and data eras are evaluated as separate data samples and later combined in order to effectively evaluate the different sensitivity offered by eras of different values of t_{online} .

4.3.6 Offline Timing Reconstruction

As the online trigger timing algorithm is limited by CPU time at HLT, an optimized offline timing algorithm is used to more precisely reconstruct jet times. This offline timing algorithm works the following way:

1. Using all ECAL cells in a jet satisfying $E_{\text{cell}} > 0.5$ GeV, $|t_{\text{cell}}| < 12.5$ ns, and passing cell quality flags, identify the median cell time.
2. Remove all ECAL cells with $|t_{\text{cell}}^{\text{median}} - t_{\text{cell}}| > 3$ ns.
3. Compute the P_T -weighted jet time using the remaining cells:

$$t_{\text{jet}} = \frac{\sum_i P_{T,i} t_i}{\sum_i P_{T,i}}$$

A sample event showing cell energy as a function of time is shown in Figure 4.17.




Figure 4.17: ECAL cells from a representative, soft jet. Cells accumulate near the median, but some cells with times far from the jet median cause a disproportionate increase to jet time RMS.

This approach of using only cells in proximity to the median energy deposit to compute the jet time is an improvement over the Run 2 approach of only using the median cell time. Comparing reconstructed jet times with gen times as in Figure 4.18 and Table 4.5, we see

that a 3 ns window offers the best agreement with gen timing without sacrificing timing resolution.




Figure 4.18: A window of 3 ns around the jet mean time captures a majority of the energy physically associated with the jet.

Case	Mean (ns)	Std. dev. (ns)
t_{jet}^{Median}	-0.05	0.41
$ t - t^{Median} < 1$	-0.04	0.32
$ t - t^{Median} < 2$	-0.03	0.31
$ t - t^{Median} < 3$	-0.001	0.32
$ t - t^{Median} < 4$	0.02	0.33
$ t - t^{Median} < 5$	0.03	0.34

Table 4.5: Summary statistics comparing reconstructed and gen jet times for different jet timing windows around the median. In Run 2, t_{jet}^{Median} was used to compute the jet time, but it shows significantly worse performance compared to the improved P_T -weighted Run 3 algorithm.

One other conclusion we can draw from Table 4.5 is that all values of $|t - t^{Median}|$ outperform a simple choice of the median cell time. This is due to the increased resolution

gained by P_T -weighting, an algorithm which has been shown as in Figure 4.19 to give better agreement between generated and reconstructed jet times [56].




Figure 4.19: Weighting the time by P_T yields a significant improvement in timing resolution, as shown by the relative time difference between gen and reconstructed jet times for $|\eta| < 0.5$ (left) and $1 < |\eta| < 1.5$ (right) [56].

4.4 Background Estimation

The primary background sources for delayed jet signatures come from beam halo muons, cosmic ray muons, and out-of-time jets from prompt and satellite bunch collisions.

For beam halo and prompt backgrounds, we use an ABCD method to enrich control regions in background and predict contributions in the signal region. For cosmic muon backgrounds, we use an event-level veto. Predictions are first made in validation regions (VRs) that do not overlap with signal regions. The agreement of these predictions in the validation region is used to derive systematic uncertainties for the backgrounds in the signal region.

Prompt-like jets are the dominant background after all selections, with beam halo and cosmic ray muons suppressed through a series of targeted event and jet requirements. The details of each background estimation method are summarized below.

4.4.1 Beam Halo

Beam halo muons are produced from proton interactions in beam gas and collimator collisions at the fringes of each bunch. These interactions produce pions and kaons that decay to muons which are collinear with the beam momentum.




Figure 4.20: Beam halo travels parallel to each proton bunch. The time at which it interacts in the ECAL is a function of its position relative to the bunch.

Beam halo muons have a characteristic spatial and timing signature related to their production mechanism. As beam halo muons travel parallel to the beam, they produce energy deposits in the ECAL which are out-of-time with the bunch collisions, concentrated in ϕ , and spread across a wide range in η . The timing for energy deposits due to beam halo follow an explicit spatial relation given by:

$$t_{\text{halo}} = -\frac{Z + \sqrt{Z^2 + R^2}}{c} \quad (4.4)$$

Expressing in terms of η , we find:

$$t(\eta) = -\frac{R}{c} e^{\pm\eta} + \Delta t_{\text{bunch}} \quad (4.5)$$

where Δt_{bunch} is the time difference between primary and satellite bunch crossings. As beam halo precedes proton bunches in time, positive-time beam halo is produced by satellite

bunches in 2.5 ns intervals delayed with respect to the primary bunch crossing. Here, the CMS ECAL radius is used as $R = 1.4$ m. Solving in terms of η , we find:

$$\eta(t) = \pm \ln \left[\frac{c}{R} (\Delta t_{bunch} - t) \right] \quad (4.6)$$

Figure 4.21 illustrates a typical beam halo timing pattern with contours overlaid using Equation 4.6.




Figure 4.21: Beam halo follows contours in jet η and jet time set by the geometric relationship between the beam halo muons and the proton bunches.

When beam halo interacts in the CMS ECAL, its deposits follow a unique shape in which the distribution is much broader in η compared to its distribution in ϕ . An example event passing beam halo-enriched selections is shown in Figure 4.22.




Figure 4.22: As the distribution of ECAL cell hits is much broader in η as compared to ϕ for beam halo energy deposits, we can use a shape discriminant defined by a ratio of a jet's width in ϕ compared to η to reject beam halo.

To enrich in beam halo, we require that the fraction of ECAL energy matched to hits in the CSCs is greater than 0.6 and that the fraction of jet energy in the HCAL is less than 20% of the total energy in the HCAL or ECAL.

Using these jets enriched in beam halo, we developed a jet shape discriminant to offer further reductions in beam halo backgrounds. Since beam halo jets tend to deposit energy into ECAL cells with the same index position in ϕ (i_ϕ) but across a broad range of values of η (i_η), we developed a variance-like parameter to measure the shape of each jet:

$$M_\phi = \frac{1}{N_{\text{cells}}} \sum_{i_\phi} [N_\phi(i_\phi)]^2 \quad (4.7)$$

$$M_\eta = \frac{1}{N_{\text{cells}}} \sum_{i_\eta} [N_\eta(i_\eta)]^2 \quad (4.8)$$

Here, N_{cells} is the total number of ECAL cell constituents of the jet, $N_\phi(i_\phi)$ is the number of cells with azimuthal index i_ϕ , and $N_\eta(i_\eta)$ is the number of cells with pseudorapidity index

i_η .

The jet shape ratio parameter $S = M_\phi/M_\eta$ outperforms a simple ratio of the variance of the cell distribution in i_η and i_ϕ , as the jet's variance in η, ϕ is sensitive to single-cell energy deposits with a large spacing from the bulk of the jet. By comparison, this cell multiplicity value does not weight by distance from the jet, scoring each cell deposit equally. A plot showing signal and background distributions in jet shape, as well as a comparison of the performance to the η/ϕ variance ratio of the jet, are shown in Figures 4.23 and 4.24.




Figure 4.23: Using jet shape (left) provides a stronger discriminant for identifying beam halo backgrounds as compared to using the ratio of the η and ϕ variances (right).




Figure 4.24: A ROC plot highlights the much stronger efficiency of shape ratio used to identify beam halo rather than a ratio of the jet η/ϕ variance.

A requirement of $M_\phi/M_\eta < 1.5$ is 98.4% efficient on control jets but only 1.4% efficient

on beam halo-enriched jets, offering a powerful handle for background rejection. Figure 4.25 illustrates the efficient discrimination between a population of control jets and beam halo-enriched jets.




Figure 4.25: Control jets (left) peak at low values of M_ϕ and higher values of M_η , whereas in beam halo-enriched jets (right) the opposite is true.




Figure 4.26: Events enriched in beam halo (left) tend to have large CSC energy fraction and jet shape ratio, while signal samples (right) have low jet shape ratio and a range of values of CSC energy fraction which peak at zero.

We use the ABCD method in the jet shape parameter S and the CSC Energy Fraction. To measure the closure of this ABCD, we study a sideband by inverting the HCAL energy fraction and the ECAL N90 requirement and measuring beam halo event distributions in $t_{jet} < -2$ ns. Figure 4.26 illustrates the distribution in these variables of a beam halo-enriched sample compared with no prompt tracks ($PVDR > 0.3$) to a signal-like prompt jet-enriched

sample ($PVDR < 0.3$).




Figure 4.27: The ratio of events failing the shape requirement agrees in regions of a low, intermediate, and high CSC energy fraction.

A closure test is performed by binning the beam halo ABCD in CSC energy fraction and computing the ratio of events passing and failing the shape ratio requirement. The results in Figure 4.27 show good agreement between values of low, intermediate, and high CSC energy fraction.

With closure in the validation region, we proceed to estimate the beam halo backgrounds using our 2023 delayed jet trigger dataset. We will later scale by luminosity to obtain a background estimate for 2023+24.

We first measure predicted beam halo events using a sample with only partial jet and event requirements, removing those which significantly reject beam halo: hadronic energy, hadronic energy fraction, and ECAL N_{90} jet requirements, and the event-level satellite bunch collision veto. The event-level H_T requirement is also removed as beam halo has little H_T in general.¹

¹However, as this background is computed using the DJ skim, there are some implicit online HT or shape requirements to be present in the skim itself. Also note there is no additional, explicit beam halo jet- or event-level enrichment applied here; the enrichment comes from removing the selections targeting beam halo

The background is computed for $10 \text{ ns} > |t_{jet}| > 2 \text{ ns}$ to avoid prompt jet leakage. The ABCD prediction as shown in Table 4.6 under these selections yields an estimate of $N_{BH,H_T>0,partial} = 17.5 \pm 1.9$. This sample is used to predict the number of background beam halo events before applying the beam halo veto requirements.

Region	Selection	Events
A (predicted)	$S < 1.5$, CSC fraction < 0.6	17.5
B	$S < 1.5$, CSC fraction > 0.6	166
C	$S > 1.5$, CSC fraction < 0.6	86
D	$S > 1.5$, CSC fraction > 0.6	815

Table 4.6: Beam Halo, $H_T > 0$, 2023 delayed jet dataset, enriched by applying only partial region definitions and yields.

We estimate the cut efficiencies of each remaining beam halo veto requirement by measuring them in region D ($S < 1.5$, CSC fraction > 0.6). In this way, we estimate the efficiency for a beam halo to pass the remaining selections, then apply this to the beam halo background estimate. The number of events passing each selection are given in Table 4.7.

Selection	Events	Eff _{cut} (%)	Eff _{cum} (%)
Enriched	5039	—	100.0
Hadronic Energy > 10	899	17.8	17.8
Hadronic Energy Fraction > 0.2	14	1.6	0.3
$N_{90} > 7$	13	92.9	0.3
Satellite Veto	13	100.0	0.3

Table 4.7: Cut efficiencies when applying beam halo veto requirements to an enriched beam halo sample in JetMET.

Our total beam halo background is then given as the product of the total remaining beam halo cut efficiencies and the estimated beam halo background under partial selections:

$$N_{BH,H_T>0,partial} \times R_{hadE} \times R_{hadEF} \times R_{N_{90}} \times R_{SatelliteVeto} = N_{BH,H_T>0} \quad (4.9)$$

This procedure computes a beam halo background estimate of 0.23 ± 0.06 when no H_T requirement is applied. However, the lowest H_T requirement on signal is $H_T > 500$. Under this $H_T > 500$ requirement, there are zero events in all ABCD regions. We infer a total

background in 2023 of $N_{BH,H_T>0,partial} = 0$. A 95% constraint on the number of beam halo events in this region is obtained using an assumption of $N_{BH,H_T>500,partial} = 3$. Using this approach, we obtain a beam halo background estimate of $3 \times 13/5039 = 0.0077 \pm 0.0077$ events in 2023 per 8 ns. Rescaling this background per 9 ns and scaling from 2023 luminosity to 2023+24, we obtain a total background prediction of 0.043 ± 0.043 events. This background is treated as uniform in time for $t_{jet} > 1$ ns and is divided evenly over all timing bins.

4.4.2 Cosmic Muons




Figure 4.28: The maximum separation in ϕ in the DTs is greater than $\pi/2$ for this event, causing it to fail the cosmic veto requirement.

As cosmic rays appear in CMS collisions uniformly in time, they can fake an out-of-time jet and form a background for a delayed jet signature. To reject cosmic muons, we use an event-level filter on activity in the drift tubes (DTs) and resistive plate chambers (RPCs) in the CMS muon system. As cosmic muons must travel radially across the CMS detector to

form energy deposits in the ECAL, we use a requirement of $\max(\Delta\phi) < \pi/2$ between all DT and RPC hits in stations 3 and 4. We restrict this parameter to stations 3 and 4 as the outer stations are much purer in cosmic muon activity as compared to the inner stations, which are exposed to significant jet punch-through. An example of this parameter measured in a cosmic event is shown in Figure 4.28.

To assess the performance of this cosmic veto, we overlay the DT and RPC activity from a cosmic-enriched sample onto data from 2023Cv1, as outlined in Section 4.1. An event veto requirement of $\text{Max}(\Delta DT, \Delta RPC) < \pi/2$ yields an efficiency of 91.3% in control data and 13.4% in cosmic-enriched data, as in Figure 4.29. We will use this efficiency to estimate the number of background events due to cosmics from the number failing this cosmic veto requirement; however, as the cosmic-enriched data sample had only DT and RPC activity, we must use a $\text{MET} > 120$ GeV sample to study the efficiency of jet requirements.




Figure 4.29: Control events only fail the event-level cosmic veto when there is activity in the DT/RPC stations 3 and 4 on opposite sides of the detector. A cut placed on signal-like events of $\text{Max}(\Delta DT, \Delta RPC) < \pi/2$ has an efficiency of 91.3%, whereas for cosmic-enriched events the efficiency is 13.4%. For cosmic-enriched events, this veto is only failed in the cases where there was activity on one side of the DT/RPC system (0), or no deposits at all (-1). This veto inefficiency can appear for cosmics which enter through the barrel and exit through the endcaps, though these cosmics are unlikely to produce large hits in the barrel ECAL.

To estimate the number of cosmic background events using this $\text{MET} > 120$ GeV sample, we follow a similar procedure as for estimating beam halo. As before, we will estimate the cosmic background using 2023 data, then scale to the expected background for 2023+24 luminosity.

First, we enrich in cosmic activity by applying strict jet requirements of hadronic energy > 40 GeV and hadronic energy fraction > 0.5 . We apply all jet requirements except for a partial set of cosmic muon veto requirements of ECAL N_{cells} , ECAL $N_{90\%}$, and an event-level satellite bunch veto. Additionally, the background is computed for $10 \text{ ns} > |t_{\text{jet}}| > 2 \text{ ns}$ to avoid prompt jet leakage. Finally, we apply an H_T requirement of $H_T > 500$ GeV, integrating cosmic backgrounds over all signal H_T bins. Under these selections, we measure 7 events passing these partial, cosmic-enriched selections in 2023 data, but failing the event-level cosmic veto $\text{Max}(\Delta DT, RPC) > \pi/2$.

Using the cosmic-triggered sample, we measured an efficiency of 13.4% to pass $\text{Max}(\Delta DT, RPC) < \pi/2$, implying a ratio of pass-to-fail events of 0.155. The measurement of 7 events failing the $\text{Max}(\Delta DT, RPC)$ requirement therefore implies 1.1 ± 0.3 cosmic events in 2023 data passing partial selections.

We estimate the cut efficiencies of each remaining veto requirement by measuring them in the $\text{MET} > 120$ GeV sample for events satisfying $\text{Max}(\Delta DT, RPC) > \pi/2$, hadronic energy > 40 GeV, and hadronic energy fraction > 0.5 . As for beam muons, we also remove the H_T requirement, as muon backgrounds are generally low H_T and this enables us to study a sample most enriched in cosmics. The number of events passing each selection are given in Table 4.8.

Selection	Events	Eff_{cut} (%)	Eff_{cum} (%)
Enriched	9	—	100.0
$N_{\text{Cells}} > 10$	2	22.2	22.2
$n_{90} > 7$	1	50.0	11.1
Satellite Veto	1	100.0	11.1

Table 4.8: Cut efficiencies when applying cosmic veto requirements to an enriched cosmic sample in JetMET.

Our total cosmic background is then given as the product of the total remaining cosmic cut efficiencies and the estimated cosmic background under partial selections:

$$N_{Cosmic, H_T > 500, partial} \times R_{N_{Cells}} \times R_{N_{90}} \times R_{SatelliteVeto} = N_{Cosmic, H_T > 500} \quad (4.10)$$

This procedure computes a cosmic background estimate of 0.12 ± 0.03 events in 2023 per 8 ns. Scaling for luminosity to 2023+24 and scaling to a background per 9 ns, we measure a total cosmic background of 0.66 ± 0.17 events per 9 ns. This background is treated as uniform in time for $t_{jet} > 1$ ns and is divided evenly over all timing bins.

4.4.3 Prompt and Satellite Backgrounds

Out-of-time jets from prompt and satellite bunch collisions are the primary source of background for this search. To estimate these backgrounds, we use an ABCD method in t_{jet} and PVDR (see Section 4.2).




Figure 4.30: (Left) The ABCD used to estimate backgrounds from prompt jets, validated using negative timing. (Right) A closure test shows independence across a range of timing bins.

To validate the ABCD, we begin by using a sample triggered with PFMET120 in the JetMET dataset. Binning the ratio between events with high and low PVDR in t_{jet} , we see that there is good closure across a range of timing bins in Figure 4.30.




Figure 4.31: (Left) The ABCD in a sideband of low PVDR, used to test closure in positive timing. (Right) A closure test shows independence across a range of timing bins, matching the validating in negative timing.

To perform a validation of the ABCD in positive timing, we use a sideband in $0.2 < PVDR < 0.3$. As in Figure 4.31, we see there is also closure for this low-isolation sideband. A fit to this closure region yields a value of 0.0056 ± 0.0003 .




Figure 4.32: (Left) The ABCD in a sideband of low PVDR, this time using delayed jet triggers rather than PFMET120. (Right) A closure test shows independence across a range of timing bins, consistent with the other validation regions.

Finally, we can repeat this low-PVDR sideband using a skim over delayed jet triggers rather than using a MET trigger. In Figure 4.32, we see that we also obtain good closure in this sideband. A fit to this closure region yields a value of 0.0053 ± 0.0003 , showing good agreement with the previous low-PVDR closure test in the MET trigger. With closure

tests showing independence in three separate validation regions, we proceed to background estimates using this prompt jet ABCD method.

Though this ABCD method shows good robustness across a range of datasets and validation regions, large backgrounds are generated from satellite bunch collisions near $t_{jet} = 5$ ns. An effective method to veto this background is to require a prompt ISR-like jet satisfying requirements summarized in Table 4.1. The effects of applying this ISR requirement are shown in Figure 4.33. Finally, prompt backgrounds are determined using the ABCD method, with background yields binned in jet time.




Figure 4.33: (Left) The prompt ABCD over all 2023 data. (Right) Applying an event-level ISR requirement, we significantly reduce backgrounds induced by satellite bunch collisions near 5ns.

4.5 Interpretation

We use the Twin Higgs model to interpret our results. In this model, a heavy Higgs-like mediator pair produces two long lived particles which decay to four b-jets.

4.5.1 Systematic uncertainties

We consider the following uncertainties on signal yields:

- **Luminosity:** Lumi uncertainties are taken from central recommendations for 2023 and 2024, respectively
- **Trigger:** Signal jet times are varied by timing resolution measured in 2023 and 2024, and resulting impacts on signal trigger efficiency are used to derive a trigger systematic
- **MC statistics:** Signal MC is generated with 10^5 events per model point; statistical uncertainties are generated for each model depending on selection efficiency
- **Pileup reweighting:** MC events are reweighted using data pileup distributions, and an uncertainty is derived on production cross-section, taken as correlated between 2023 and 2024
- **Jet energy corrections:** Jet energies were varied following standard JEC prescriptions for 2023 and 2024 data, and uncertainties were measured using systematic variation of JECs on selections for jet P_T and energy.

4.5.2 Signal Efficiency

The selection efficiencies for a selection of Twin Higgs model parameters are shown in Table 4.9, 4.10, and 4.11. For these signal models, the LLP mass was taken to be $M_{LLP} = M_H/2 - 25$ GeV to assess similarly boosted final states across a range of twin Higgs masses, and the lifetime is taken to be $c\tau = 1$ m.

	MH400		MH1000		MH1500		Background	
	Yield	Eff%	Yield	Eff%	Yield	Eff%	Yield	Eff%
All events	14903.2	100.00%	201.5	100.00%	15.0	100.00%	372819190.0	100.00%
Passing Filters	14903.2	100.00%	201.5	100.00%	15.0	100.00%	371397475.0	99.62%
Delayed-Jet Trigger	3354.9	22.51%	162.2	80.48%	13.5	89.82%	99884305.0	26.89%
Event H_T Cut	252.6	7.53%	56.1	34.56%	2.6	18.96%	3903075.0	3.91%
NoCosmic	230.0	91.04%	48.6	86.65%	2.2	84.64%	3419125.0	87.60%
NoBeamHalo	230.0	100.00%	48.6	100.00%	2.2	100.00%	3419125.0	100.00%
NoSatellite	113.8	49.47%	34.0	69.93%	1.5	69.82%	1065740.0	31.17%
$P_T > 40$	113.8	100.00%	34.0	100.00%	1.5	100.00%	1065740.0	100.00%
$ \eta < 1.48$	90.7	79.77%	31.7	93.29%	1.4	95.30%	870430.0	81.67%
$E_{ECAL} > 20$	83.2	91.68%	30.8	97.16%	1.4	98.51%	765635.0	87.96%
# ECAL Cells in Jet > 10	81.6	98.05%	30.5	99.20%	1.4	99.51%	739465.0	96.58%
$t_{RMS} < 1.5$	81.2	99.60%	30.4	99.59%	1.4	99.90%	722690.0	97.73%
$E_{HCAL} > 10$	81.2	100.00%	30.4	99.96%	1.4	100.00%	720845.0	99.74%
$E_{HCAL}/E_{ECAL} > 0.2$	80.5	99.08%	30.2	99.33%	1.4	99.48%	685170.0	95.05%
# ECAL Cells, 90% > 7	80.4	99.91%	30.1	99.69%	1.4	99.56%	621415.0	90.70%
$E_{ECAL,CSC}/E_{ECAL} < 0.6$	76.3	94.82%	29.4	97.67%	1.4	97.74%	590145.0	94.97%
Shape Ratio < 1	76.3	100.00%	29.4	100.00%	1.4	100.00%	583475.0	98.87%
$dR > 0.3, t_{jet} > 1$	53.5	70.15%	26.3	89.31%	1.2	91.09%	4186.6	0.72%
$1 < t_{jet} < 1.5$	21.0	39.17%	4.0	15.26%	0.1	6.87%	3036.0	72.52%
$1.5 < t_{jet} < 2$	14.5	27.11%	3.8	14.46%	0.1	7.19%	1046.1	24.99%
$2 < t_{jet} < 2.5$	7.9	14.76%	3.4	13.10%	0.1	7.80%	78.9	1.88%
$2.5 < t_{jet} < 3$	5.4	10.05%	3.2	12.27%	0.1	8.71%	19.1	0.46%
$3 < t_{jet} < 4$	2.9	5.35%	4.8	18.43%	0.2	16.14%	2.2	0.05%
$t_{jet} > 4$	1.9	3.56%	7.0	26.48%	0.7	53.29%	4.3	0.10%

Table 4.9: 1-Jet Cutflow for H_T 500–700, signal scaled to the BSM Higgs cross-section, $B = 10^{-2}$.

	MH400		MH1000		MH1500		Background	
	Yield	Eff%	Yield	Eff%	Yield	Eff%	Yield	Eff%
All events	14903.2	100.00%	201.5	100.00%	15.0	100.00%	372819190.0	100.00%
Passing Filters	14903.2	100.00%	201.5	100.00%	15.0	100.00%	371397475.0	99.62%
Delayed-Jet Trigger	3354.9	22.51%	162.2	80.48%	13.5	89.82%	99884305.0	26.89%
Event H_T Cut	111.6	3.33%	29.9	18.44%	5.1	38.08%	1058805.0	1.06%
NoCosmic	102.0	91.37%	26.3	87.89%	4.4	84.67%	913500.0	86.28%
NoBeamHalo	102.0	100.00%	26.3	100.00%	4.4	100.00%	913500.0	100.00%
NoSatellite	46.2	45.33%	18.2	69.08%	3.1	70.37%	311410.0	34.09%
$P_T > 40$	46.2	100.00%	18.2	100.00%	3.1	100.00%	311410.0	100.00%
$ \eta < 1.48$	33.0	71.45%	16.7	91.88%	2.9	95.26%	257350.0	82.64%
$E_{ECAL} > 20$	29.8	90.12%	16.2	97.07%	2.9	98.53%	226875.0	88.16%
# ECAL Cells in Jet > 10	29.4	98.90%	16.0	99.05%	2.9	99.59%	219530.0	96.76%
$t_{RMS} < 1.5$	29.3	99.59%	16.0	99.55%	2.9	99.70%	214345.0	97.64%
$E_{HCAL} > 10$	29.3	99.86%	16.0	99.98%	2.9	99.96%	213765.0	99.73%
$E_{HCAL}/E_{ECAL} > 0.2$	29.0	99.13%	15.8	99.21%	2.8	99.14%	204140.0	95.50%
# ECAL Cells, 90% Energy > 7	28.8	99.37%	15.7	99.34%	2.8	99.54%	184710.0	90.48%
$E_{ECAL,CSC}/E_{ECAL} < 0.6$	27.6	95.86%	15.3	97.30%	2.8	97.87%	174875.0	94.68%
Shape Ratio < 1	27.6	100.00%	15.3	100.00%	2.8	100.00%	173045.0	98.95%
$dR > 0.3, t_{jet} > 1$	20.1	72.84%	13.4	87.73%	2.5	90.97%	1134.2	0.66%
$1 < t_{jet} < 1.5$	8.4	41.52%	2.6	19.39%	0.2	9.73%	836.3	73.73%
$1.5 < t_{jet} < 2$	5.1	25.12%	2.3	17.25%	0.3	10.24%	273.2	24.09%
$2 < t_{jet} < 2.5$	2.8	13.84%	2.0	15.21%	0.3	10.81%	19.6	1.73%
$2.5 < t_{jet} < 3$	2.1	10.45%	1.8	13.06%	0.3	10.34%	4.5	0.40%
$3 < t_{jet} < 4$	1.2	5.95%	2.3	17.34%	0.5	20.07%	0.2	0.02%
$t_{jet} > 4$	0.6	3.12%	2.4	17.75%	1.0	38.82%	0.4	0.04%

Table 4.10: 1-Jet Cutflow for H_T 700–1000, signal scaled to the BSM Higgs cross-section, $B = 10^{-2}$.

	MH400		MH1000		MH1500		Background	
	Yield	Eff%	Yield	Eff%	Yield	Eff%	Yield	Eff%
All events	14903.2	100.00%	201.5	100.00%	15.0	100.00%	372819190.0	100.00%
Passing Filters	14903.2	100.00%	201.5	100.00%	15.0	100.00%	371397475.0	99.62%
Delayed-Jet Trigger	3354.9	22.51%	162.2	80.48%	13.5	89.82%	99884305.0	26.89%
Event H_T Cut	40.2	1.20%	14.0	8.61%	4.6	34.14%	265595.0	0.27%
NoCosmic	38.1	94.69%	12.1	86.35%	3.9	84.50%	220865.0	83.16%
NoBeamHalo	38.1	100.00%	12.1	100.00%	3.9	100.00%	220865.0	100.00%
NoSatellite	15.7	41.20%	8.4	69.50%	2.7	70.24%	80070.0	36.25%
$P_T > 40$	15.7	100.00%	8.4	100.00%	2.7	100.00%	80070.0	100.00%
$ \eta < 1.48$	10.2	65.13%	7.6	90.36%	2.5	92.60%	67590.0	84.41%
$E_{ECAL} > 20$	9.3	90.91%	7.3	96.44%	2.5	98.40%	59865.0	88.57%
# ECAL Cells in Jet > 10	9.3	99.64%	7.3	99.25%	2.5	99.40%	58255.0	97.31%
$t_{RMS} < 1.5$	9.2	99.85%	7.2	99.31%	2.5	99.84%	56900.0	97.67%
$E_{HCAL} > 10$	9.2	100.00%	7.2	99.98%	2.5	100.00%	56760.0	99.75%
$E_{HCAL}/E_{ECAL} > 0.2$	9.2	100.00%	7.1	99.04%	2.5	99.23%	54570.0	96.14%
# ECAL Cells, 90% Energy > 7	9.1	98.05%	7.0	98.14%	2.4	98.71%	49650.0	90.98%
$E_{ECAL,CSC}/E_{ECAL} < 0.6$	8.2	90.95%	6.8	96.94%	2.4	97.72%	46655.0	93.97%
Shape Ratio < 1	8.2	100.00%	6.8	100.00%	2.4	100.00%	46285.0	99.21%
$dR > 0.3, t_{jet} > 1$	4.9	59.75%	5.9	86.26%	2.1	90.29%	261.3	0.56%
$1 < t_{jet} < 1.5$	3.2	64.03%	1.3	22.53%	0.4	18.00%	196.1	75.04%
$1.5 < t_{jet} < 2$	0.8	16.74%	1.1	18.85%	0.3	16.05%	59.7	22.83%
$2 < t_{jet} < 2.5$	0.5	10.42%	0.9	16.06%	0.3	15.14%	4.5	1.74%
$2.5 < t_{jet} < 3$	0.1	2.70%	0.8	13.41%	0.3	12.19%	0.8	0.29%
$3 < t_{jet} < 4$	0.2	3.12%	0.9	15.26%	0.4	16.60%	0.1	0.05%
$t_{jet} > 4$	0.1	2.99%	0.8	13.89%	0.5	22.01%	0.1	0.05%

Table 4.11: 1-Jet Cutflow for H_T 1000–10000, signal scaled to the BSM Higgs cross-section, $B = 10^{-2}$.

In Figure 4.34 we show the total signal efficiency and event yields including trigger and all selections for a luminosity of 136.9 fb^{-1} . Maximal efficiency is achieved for signal models with a lifetime near $c\tau = 1 \text{ m}$ and Twin Higgs mediator masses $M_H \geq 1 \text{ TeV}$. Sensitivity is optimal for LLP masses near $M_H/2$, and the sensitivities shown here assume $M_{LLP} = M_H/2 - 25 \text{ GeV}$. By contrast, signal yields and sensitivities for models with $M_{LLP} = M_H/4 - 25 \text{ GeV}$ are shown in Figure 4.34. Signal and event yields for the remaining analysis bins are given in Appendix A.5.




Figure 4.34: (Left) Total efficiency for signal models across a range of lifetimes and Twin Higgs masses. (Right) Event yields using a luminosity of 136.9 fb^{-1} .

4.5.3 Sensitivity

Signal and background yields are binned in H_T , N_{jet} , and ISR as shown in Table 4.12.

$\mathbf{H_T}$ bin [GeV]	$\mathbf{N_{jet}}$ bin	ISR selection
$500 \leq H_T < 700$	$N_{jet} = 1$	ISR required
$500 \leq H_T < 700$	$N_{jet} = 1$	ISR veto
$500 \leq H_T < 700$	$N_{jet} \geq 2$	ISR required
$500 \leq H_T < 700$	$N_{jet} \geq 2$	ISR veto
$700 \leq H_T < 1000$	$N_{jet} = 1$	ISR required
$700 \leq H_T < 1000$	$N_{jet} = 1$	ISR veto
$700 \leq H_T < 1000$	$N_{jet} \geq 2$	ISR required
$700 \leq H_T < 1000$	$N_{jet} \geq 2$	ISR veto
$H_T \geq 1000$	$N_{jet} = 1$	ISR required
$H_T \geq 1000$	$N_{jet} = 1$	ISR veto
$H_T \geq 1000$	$N_{jet} \geq 2$	ISR required
$H_T \geq 1000$	$N_{jet} \geq 2$	ISR veto

Table 4.12: Signal is binned in H_T , N_{jet} , and ISR, then compared to background using a template fit in jet time.

Background and signal yields are further compared using a shape analysis; total backgrounds and signal yields for a range of model parameters integrated over all H_T bins and shown in Figure 4.35.






Figure 4.35: Signal and background yields are compared as a function of jet time bin. Signal Twin Higgs mass, LLP mass are listed in the title, respectively. For models with a boosted LLP, the signal yields drop quickly in time. For less boosted models where $M_{\text{LLP}} \approx M_H/2$, the signal yields are flatter in lifetime.

Figure 4.36 shows 95% limits on $\sigma_H \times B^2$ in a plane of Twin Higgs mass vs lifetime. Sensitivity is enhanced for high lifetimes and heavy Higgs masses due to increased signal efficiency for less boosted LLPs. Replacing the value of σ_H with centrally produced BSM Higgs predictions in Figure 4.37 shows the sensitivity as a function of purely $B(H \rightarrow \chi\chi)$. Finally, Figure 4.38 shows the sensitivity for $M_H = 800$ GeV compared to other CMS LLP analyses.

Figure 4.36: Sensitivity as a function of M_H , M_{LLP} , and τ_{LLP} .Figure 4.37: Sensitivity expressed purely as a function of $B(H \rightarrow \chi\chi)$, using central values for BSM σ_H .




Figure 4.38: Sensitivity for a range of $H \rightarrow XX$ searches with purely hadronic final states. This analysis is more sensitive at longer lifetimes compared to previous analyses. Figure data for other searches originally from [57].

4.5.4 Outlook

The backgrounds and expected sensitivity for Chapter 4 were presented using 2023 data. Results and backgrounds were scaled from the 2023 luminosity of 27.9 fb⁻¹ to the full 2023+2024 luminosity of 109.1 fb⁻¹.

The final sensitivity and results are pending the processing of the 2024 data and an estimate of the final systematics. In particular, the dominant timing systematic will be derived from the online turn-on performance of the triggers with respect to jet time. These trigger turn-ons are currently being studied and 2024 data is currently being processed. Results are expected for the completed analysis in early 2026.

5 The milliQan Experiment: Detector, Commissioning, and Calibration

In this section, I will describe the milliQan experiment, then summarize the Run 3 bar detector analysis. I will largely focus on the bar detector, though many experimental features of milliQan are now present in the broader landscape of mCP direct detection experiments, including FORMOSA, SUBMET, and FLAME/LANSCE-mQ [58, 59, 60].

The fundamental technology of the milliQan experiment involves reading out energy deposits in organic plastic scintillator ¹, and read out with photomultiplier tubes (PMTs). These technologies are discussed in Section 5.1 and 5.2 below.

5.1 Scintillator

Scintillator is a material which emits light when excited by ionizing radiation. Scintillators generally come in two types: organic and inorganic.

5.1.1 Organic Scintillator

Organic Scintillator is composed of organic molecules containing carbon, hydrogen, nitrogen, and oxygen. Organic scintillator is composed of a polyvinyltoluene (PVT) or polystyrene (PS) base. The primary scintillator for the milliQan experiment is EJ-200, a PVT-based organic plastic scintillator produced by Eljen Technology. Some dark matter experiments use liquid organic scintillator, such as in the outer detector for LZ.

¹Uniformly referred to in this document as “plastic scintillator” or just “scintillator”, except in Section 5.1

As the typical emission spectrum for toluene or styrene is in the UV, organic scintillators are usually doped with wavelength-shifting fluors such as PPO (2,5-diphenyloxazole) or POPOP (1,4-bis(5-phenyloxazol-2-yl)benzene).

The mechanism for scintillation in an organic scintillator is given by the following process:

1. An ionizing particle deposits energy and excites the bulk of a plastic scintillator lattice. On the lattice, the energy deposit excites a local π -electron cloud on a benzene ring (such as toluene or styrene) to its first excited state S_1 .
2. The excited π -electrons induce a transition dipole moment in the donor benzene ring. Through dipole-dipole interactions, energy is transferred non-radiatively to an acceptor fluor, such as PPO. This process is known as Förster resonance energy transfer (FRET). The rate of energy transfer from the donor to the fluor acceptor W is given by:

$$W = \frac{1}{\tau_D} \left(\frac{R_0}{r} \right)^6 \quad (5.1)$$

where τ_D is the decay lifetime of the excited state in the donor, r is the donor-to-acceptor distance, and the Förster radius R_0 is given by:

$$R_0^6 = \frac{9\chi^2 n^4 c^4}{8\pi N_A} \int F_D(\omega) \mu_A(\omega) \frac{d\omega}{\omega^4} \quad (5.2)$$

where $F_D(\omega)$, $\mu_A(\omega)$ are the emission and absorption spectra of the donor and acceptor, ω is the frequency of the transition dipole resonance, n is the refractive index, N_A is the concentration of acceptors, and χ is a factor related to the spatial orientation of the dipoles, equal to $2/3$ for random orientations [61, 62]. If the emission spectrum of this transition dipole moment is well-matched to the absorption spectrum of the acceptor fluor, then $\int F_D(\omega) \mu_A(\omega) \frac{d\omega}{\omega^4}$ is maximized for highest light yield.

3. After the primary fluor is excited via FRET, secondary fluors may be excited to further shift the scintillation wavelength, either through re-emission of photon or further FRET. For a typical plastic scintillator with light yield 10,000 photons per MeV and a 400 nm

(3 eV) scintillation peak, the total energy conversion is around 3%.

A diagram of a typical organic scintillation process is shown in Figure 5.1.




Figure 5.1: The typical process of scintillation in an organic plastic scintillator. The concentration of the primary fluor must be high to engage in FRET, while secondary fluor concentrations may be lower as re-emission is usually radiative, involving absorption and emission of a scintillation photon [63].

5.1.2 Inorganic Scintillator

In inorganic scintillators, an electron-hole pair is created after an initial ionization event. Together, these form an “exciton”, a bound state of conduction electrons and holes. This exciton loses energy over time, then eventually recombines. This recombination results in the emission of scintillation light. The relaxation process is a function of the impurity doping in the crystal, providing energy levels in the forbidden region between valence and conduction to facilitate scintillation. This process is illustrated in Figure 5.2.

One inorganic scintillator of interest to mCP experiments is CeBr_3 , an expensive but high density and light yield material with approximately 30 times higher total light yield

for the same ionization length as organic plastic scintillator. The CMS ECAL, described in Section 2.2.3, uses PbWO_4 , an inorganic scintillator with low light yield but high radiation tolerance.




Figure 5.2: During inorganic scintillation, photons are produced by the relaxation of an electron-hole exciton, usually by doped impurities. Figure credit: A. Vasiliev via [64]

5.1.3 Cherenkov Light

Some experiments rely on Cherenkov light rather than charged particles, such as the Pb-glass calorimeters in the PHENIX detector [65]. Cherenkov light is also produced in scintillator, however it is generally subdominant compared to scintillation light yield.

The amount of Cherenkov light emitted by the passage of a charged particle through a medium is given by the Frank-Tamm formula:

$$\frac{\partial^2 E}{\partial x \partial \omega} = \frac{q^2}{4\pi} \mu(\omega) \omega \left(1 - \frac{c^2}{v^2 n^2(\omega)} \right) \quad (5.3)$$

For a wavelength range where the index of refraction is fixed, and for $\mu(\omega)/\mu_0 \approx 1$ in scintillator, we can rewrite this as:

$$\frac{\partial^2 N}{\partial x \partial \lambda} = \frac{2\pi\alpha z^2}{\lambda^2} \left(1 - \frac{1}{\beta^2 n^2(\lambda)}\right) \quad (5.4)$$

where the particle charge is given by $z = q/e$. Here we've used the fact that:

$$\mu_0 = \frac{2\alpha h}{e^2 c} = 4\pi \times \frac{\alpha \hbar}{e^2 c} \quad (5.5)$$

And we can rewrite total energy in the interval as the number of photons using $E = Nh\omega$.

Integrating over wavelength and the length of a bar L ,

$$N = 2\pi\alpha z^2 L \left(1 - \frac{1}{\beta^2 n^2}\right) \left(\frac{1}{\lambda_1} - \frac{1}{\lambda_2}\right) \quad (5.6)$$

where λ_1, λ_2 is the acceptance of the readout. Using an optical range of 300 nm to 550 nm for a typical PMT, $n = 1.58$ for EJ-200 scintillator, and $L = 60$ cm as in a milliQan bar, then at charge $z = 1$, we observe 2.5×10^4 Cherenkov photons with a distribution peaked at UV wavelengths.

Compare this to scintillation light: a typical MIP deposits 2 MeV/cm during ionization. At a datasheet specification of 10^4 photons per MeV, plastic scintillator will produce 1.2×10^6 photons in per 60 cm at charge $z = 1$. Therefore, Cherenkov light is only relevant to the sensitivity of very low light yield scintillator.²

5.1.4 Scintillator Properties

Several properties are important in the design considerations of plastic scintillator. Key properties are summarized below:

- **Density:** Higher density increases the energy deposit, but can degrade the path length photons can travel through the scintillator, and is strongly material-dependent.

²That doesn't mean it's useless, however! Unlike scintillation light, Cherenkov light has a very prompt, characteristic timing structure. Dual readout calorimetry uses both scintillation and Cherenkov light to improve timing resolution.

- **Light yield:** More light yield delivers more light per energy deposit. This is a good measure of the quality of a scintillator. Total photon yield is a function of the density times the light yield.
- **Emission Wavelength:** Depending on the fluors (organic) or impurity doping (inorganic), the scintillator will peak at different wavelengths. It's important to match scintillation emission wavelength with the photodetector peak radiant sensitivity.
- **Decay time:** Sometimes called the “speed” of the scintillator, this is how long it takes to relax after becoming ionized. Inorganic scintillator tends to be slower, while organic scintillator tends to be faster. Scintillator can have a fast and slow time constant depending on the scintillation scheme.
- **Radiation hardness:** For high-luminosity applications such as in the forward regions of ATLAS and CMS, inorganic scintillators are typically much more radiation-hard due to the strong electromagnetic field of the crystal lattice, shielding scintillation sites [64].

5.2 Photomultiplier Tubes




Figure 5.3: A photomultiplier tube diagram illustrating the electron avalanche after a photoelectron is generated from a scintillation photon [66].

PMTs are photodetectors with a large (typically. 1-2 in diameter) radius, good for applications where spatial resolution is not important. PMTs for the milliQan experiment are 2” in overall

diameter, well-matched to the scintillator cross-section of 5×5 cm.

The sensitive region of the PMT is called the photocathode. Here, the incident photon is converted to an electron via the photoelectric effect. The probability for this conversion is called the quantum efficiency, and for PMTs it is around 20 – 30%. In a R878 PMT, the photocathode material is bialkali, a mixture of two materials (Sb-Rb-Cs, Sb-K-Cs) [67]. Bialkali photocathodes are well-matched for the visual spectrum and are a common choice for PMTs which read out scintillation light [68].

The photocathode is deposited onto one side of a glass “window” which forms the direct interface to incident light. For R878 PMTs, the window material is borosilicate glass, chosen due to its high spectral transmittance above 300 nm, as in Figure 5.4.




Figure 5.4: The window material is used as a substrate for the photocathode. It must additionally match the spectrum targeted by the PMT; for optical wavelengths, borosilicate glass is a common choice [68].

After a photoelectron is created, it is focused and accelerated from the photocathode into the body of the PMT. Here, the accelerated electron strikes the first dynode, ejecting several

electrons depending on the type of dynode and primary electron energy (Figure 5.5).




Figure 5.5: The number of electrons produced at each dynode is a function of the dynode material and the accelerating voltage of the electron [68].

Dynodes generally consist of a thin antimonide, oxide, or phosphide layer deposited onto a steel, copper, or nickel substrate. Dynode material and layout affects the timing, gain, and linearity of the PMT. The R878 PMT is a box-and-grid dynode style, as in Figure 5.3, a dynode layout which maximizes collection efficiency of the photoelectron for high single photon counting efficiency at the cost of lower speed and higher magnetic susceptibility relative to other dynode layouts [67, 68].

Electrons ejected from the first dynode are accelerated into the second dynode, ejecting another 2-4 electrons each. Each dynode generates a new cascade of electrons, amplifying

the signal until a large current (gain up to 10^8) is read out at the anode. A sample output from a PMT (after amplification as described in Section 5.4) is shown below in Figure 5.6.




Figure 5.6: A sample output of a PMT, measured using an oscilloscope with a large timing window and persistent pulse display. The PMT output current is converted to a voltage, and the voltage is measured over time. After the primary pulse, afterpulses can be observed after residual gases cause additional avalanches.

PMTs experience “dark rate” due to thermal excitation in the photocathode. As electrons ejected thermally from the photocathode rather than due to the photoelectric effect are reconstructed identically, these dark counts appear identically to single photons. A typical dark rate for a R878 PMT is 1-2 kHz, although some PMTs such as R7725s can produce dark rates below 100 Hz.

Another form of dark rate are ejections of single electrons from the first dynode, a phenomenon we’ve called “dynode dark rate”. This causes a signal which is measurably smaller than a typical photoelectron, as the first dynode produces fewer initial electrons than is expected.

Another effect observed in Figure 5.6 is PMT afterpulsing. PMTs are vacuum tubes, but some residual gases linger inside the tubes. Molecules in these gases between can become ionized after collisions with electrons. These positive ions can be accelerated and strike the

photocathode or first dynode, producing many electrons.

Afterpulses have a characteristic timing structure due to the extra transit time for an ion of approximately zero starting energy compared to an electron with comparatively negligible transit time through the PMT [69]. The afterpulsing timing structure and frequency as a function of initial pulse amplitude is affected by the design of the PMT, as well as the variety, density, and distribution of residual gases. Notably, the afterpulsing timing structure is not purely a decaying exponential in the time since the original pulse; for a R7725 PMT, the afterpulsing rate peaks near $30\ \mu\text{s}$, as in Figure 5.7. This behavior requires applying an extra dead time window after the original pulse to ensure no correlated afterpulsing activity is measured.




Figure 5.7: Instantaneous dark rates for an R7725 PMT after exposure to a saturating, through-going, muon-like pulse. For the FORMOSA experiment, where saturating signals are frequent relative to the afterpulsing timing window and a long ($100\ \mu\text{s}$ – $1\ \text{ms}$) dead time is used, this afterpulsing-induced dead time can be a significant constraint.

One final effect relevant for high-gain PMTs is space-charge saturation. As very high electron densities are formed when a large initial signal is amplified at high gain, this current can influence electron trajectories and reduce the effective electric field, leading to reduced

current outputs and resulting in PMT non-linearity. In the milliQan experiment, this effect degrades measurement resolution for high charge deposits, such as resolving the difference between a $q/e = 0.1$ mCP and a $q = e$ through-going muon. Space-charge saturation can be measured and modeled, but best resolution is obtained by reducing PMT HV bias or reducing light collection.

5.3 The milliQan Demonstrator

The milliQan experiment began with the milliQan demonstrator [70]. The milliQan demonstrator was constructed of 3 layers of 3×2 scintillator bars of dimensions $5 \times 5 \times 80$ cm. Thin veto panels of dimension $0.25 \times 7 \times 40$ in ($0.7 \times 17.8 \times 101.6$ cm) covered the detector on the top and sides. Between each layer was a scintillator “slab” of dimension $1 \times 8 \times 12$ in ($2.5 \times 20 \times 30$ cm) and a lead block, originally placed to protect from any radiation, between layers.³ Finally, hodoscope bars of individual dimension $2 \times 2 \times 45$ cm were used to track beam and cosmic muon activity. The milliQan demonstrator is shown in Figure 5.8. Each scintillator channel was read out with a PMT; a majority of the PMTs were Hamamatsu R878s, but Hamamatsu R7725 and an Eljen ET9814B were also used for some channels.

³Later studies in MC showed that this lead block was not a significant benefit in reducing backgrounds, as radiation had minimal impact on total background.




Figure 5.8: (Left) A diagram of the Run 2 milliQan demonstrator. (Right) A photo of the demonstrator with the panels removed [70]. Note the underlying structure; it is reused for the Run 3 detector, but all other components were removed.

The demonstrator achieved a few key objectives:

- Measurement of backgrounds helped to deduce that, in order to achieve zero background, a fourth layer was required. The key background targeted by a fourth layer is PMT dark rate, as requiring a fourth dark count event in a narrow (< 15 ns) timing window while satisfying pointing, 4-in-a-row geometric requirements allows for negligible dark rate backgrounds for a dark rate of $O(1kHz)$. Cosmic backgrounds are also a strong function of layer number, as a more segmented detector requires one additional shower particle to deposit a small amount of MCP-like energy while also satisfying geometric requirements.
- The dataset allowed a successful demonstration of an analysis strategy that was largely

unchanged for the Run 3 detector. It achieved leading sensitivity, as shown in Figure 5.9, and the ABCD method used to project beam-off and beam-on backgrounds showed good closure, as in Tables 5.3 and 5.3.

- Measurements of our sensitivity indicated we were limited by acceptance in the bar detector design, leading to the proposal of the complementary Run 3 slab detector.
- A variety of validations between signal and MC allowed us to develop a thorough and well-validated mCP simulation, as discussed in Chapter 6.
- Initial operation of the R878 PMTs demonstrated that significant gain was necessary for consistent pulse reconstruction for single photons. Readout amplifiers were designed for this purpose, as discussed in Section 5.4.




Figure 5.9: The Run 2 demonstrator mCP sensitivity. Though the original purpose of the demonstrator was as a proof of concept, it also achieved leading mCP sensitivity at the time.

Region	N _{slab}	Min(NPE)	Prediction	Observation	Systematic
1	0	[2,20]	121.2 ^{+6.0} _{-5.9}	131	8%
2	0	> 20	47.4 ^{+5.2} _{-4.8}	45	5%
3	1	[5,30]	7.8 ^{+2.5} _{-1.8}	9	15%
4	1	> 30	2.7 ^{+2.1} _{-1.1}	4	48%
5	≥ 2	-	0.8 ^{+1.4} _{-0.4}	1	25%

Table 5.1: Measurements and predictions of the key signal regions for the Run 2 demonstrator using beam-off data. Agreement between measured and predicted values are used to derive a systematic for the beam-on data.

Region	N _{slab}	Min(NPE)	Prediction	Observation
1	0	[2,20]	124 ⁺¹¹ ₋₁₁	129
2	0	> 20	49.9 ^{+6.0} _{-5.4}	52
3	1	[5,30]	10.7 ^{+3.6} _{-2.6}	8
4	1	> 30	2.4 ^{+2.1} _{-1.1}	4
5	≥ 2	-	0.0 ^{+0.9} _{-0.0}	1

Table 5.2: Results of the beam-on measurements and predictions for the Run 2 demonstrator.

5.4 milliQan Readout Amplifiers

An important lesson from the milliQan demonstrator in Section 5.3 is that our pulses needed higher gain—a typical single photo-electron (SPE) signal was approximately 3–5 mV, just barely above the electronic noise threshold of $\sigma \approx 1$ mV. For Run 3, a new readout amplifier board was designed with the primary purpose of amplifying signals to ~ 30 –50 mV.

The main features of the readout amplifier are as follows:

- **Gain control:** The board includes two configurable gain switches with the listed attenuation factors (Table 5.3). In the highest-gain mode, the typical overall gain is around a factor of 10.
- **LED pulsing:** The base provides routing ports for an LED. These LEDs are seated

inside the 3D-printed PMT mount to the scintillator, allowing for tunable LED pulsing for the purpose of calibration or trigger validation. The DC offset to the LED is set either with a physical trimpot or via a digital trimpot, with the active setting selectable by a switch.

- **HV/Signal:** The readout amplifier is wired to the base amplifier of the PMT and offers an HV input directly to the PMT. Two SMA signal outputs at $50\ \Omega$ impedance are provided: OutB, which offers nominal gain, and OutA, which offers a $0.7\times$ gain compared to OutB.
- **LV power:** The readout amplifier is powered by an 8-pin, $\pm 5\ \text{V}$ low-voltage source. The ribbon cable provides the $+5\ \text{V}$ and $-5\ \text{V}$ power rails and a $+3.3\ \text{V}$ line for LED pulsing.

SW1	SW2	Effective Gain
0	0	$1\times$ (highest gain / lowest attenuation)
0	1	$\frac{1}{2}\times$
1	0	$\frac{1}{5}\times$
1	1	$\frac{1}{7}\times$

Table 5.3: Gain configurations of the PMT base readout amplifier.

Shown below in Figure 5.10 is an amplifier board connected to a mounted scintillator and PMT. This picture was taken during bench LED testing, so the exposed LED leads can be seen at the edge of the black electrical tape wrapping the base and mount.




Figure 5.10: A wired setup of a single milliQan bar, including scintillator, PMT, and base amplifier.

The readout scheme for the milliQan slab detector is largely similar compared to the bar detector. One additional feature for the slab detector is the presence of pairs of two-channel summing amplifiers, used to reduce the total number of readout channels on the milliQan slabs from 4 to 2. Similar to the readout amplifiers, these summing amplifiers offer configurable gain settings, primarily used for the matching of SPE peak amplitudes on summed channels such that a single photon measured in either of the summed channels offered a consistent pulse amplitude.

5.5 Run 3 milliQan Bar Construction

The wrapping procedure for the milliQan bars was carefully optimized for Run 3. milliQan scintillator bars are wrapped using Tyvek, a highly reflective ($R = 0.97$), relatively rigid material with the consistency of construction paper. The materials used for construction are

shown below in Figure 5.11 [71].




Figure 5.11: The header picture for the bar assembly guide, containing the materials list. The “creasers” were used to shape the Tyvek onto the bar, ensuring as tight a fit with as little deformation in the reflective Tyvek surface as possible. Two small cutouts of Tyvek were used to form the covering at the non-PMT face of the bar, and a 3M specular film was placed inside this endcap for maximum light collection efficiency.

In addition to Tyvek, MC studies of the performance of a scintillator bar on the bench (Section 6.3) indicated that the most critical surface to make as reflective as possible is the opposite end to the PMT.⁴ For this purpose, 3M Specular Film DF2000MA was chosen [72]. It has a greater than 99% reflectivity and reflects specularly, ensuring maximum transfer of photons from the end of the bar to the sensitive face of the PMT. A small cutout of this material was placed on the opposite face of each bar, shown in Figure 5.12.

⁴In fact, these MC studies suggested that a bar with no wrapping at all, or a highly nonreflective wrapping, around the long, bulk faces of the bar made almost no difference to the total light collection efficiency. This is due to the light guide effect of total internal reflection from a scintillator and air gap, and also due to the partially diffuse, non-focusing reflections of a photon on Tyvek.




Figure 5.12: A small specular film is inserted into the ends of milliQan bars to maximize light collection efficiency.

After a bar is effectively wrapped in Tyvek with the film insert, a visual inspection shows a kaleidoscopic pattern, as in Figure 5.13.





Figure 5.13: The specularly reflective end of the bar and lightguide properties of the scintillator produce a focusing of light. This phenomenon can be observed by looking directly down the axis of a completed bar.

To secure the PMT to the bar, a 3D-printed mount was designed, as in Figure 5.14. Exploiting the lack of perfect overlap between the 2-inch PMT and 5×5 cm scintillator, a hole is bored through the mount to provide a path for an LED to inject photons into the bar. The bars, PMTs, and mounts are wrapped in black electrical tape to prevent external light from entering the scintillator or PMT, ensuring that the only physical source of light measured by the PMTs is due to scintillation. A bar which is effectively wrapped this way is

said to be “light-tight”. Wrapping designs using black construction paper and aluminum foil were explored to provide additional light-tightening. However, early studies showed that black electrical tape was sufficient to shield the bar from external light, and additional materials added unnecessary extra bulk and complicated tolerances for later engineering designs.



Figure 5.14: (Left) A picture of an early design of the PMT mounts. An LED hole is bored nearby the PMT, offering an LED mount for light injection. The hole is not bored fully into the PMT, providing a small amount of natural attenuation for the injected photons to aid in the production of SPE-like signals. The mount itself is porous to light so must be fully wrapped. (Right) A scintillator, PMT, and LED mounted into a fully wrapped bar.

5.6 Run 3 milliQan Bar Detector

The Run 3 milliQan bar detector consists of 64 bars of dimension $5 \times 5 \times 60$ cm, configured across 4 layers of 16 channels (4×4) each. The bars were shortened from 80 cm to 60 cm in Run 3 to accommodate space for a fourth layer. The side and top veto panels were made longer and thicker to improve active and passive veto performance, with new dimensions

$5 \times 25 \times 170$ cm. As space is constrained fitting the detector into the cavern, the inter-layer panels and lead blocks were removed, and instead two “endcap” veto panels of dimension $5 \times 25 \times 25$ cm were placed at the front and back of the detector. A diagram is given in Figure 5.15.




Figure 5.15: A diagram of the Run 3 bar detector. This image comes from the Run 3 milliQan sensitivity projection; in the final detector, the orientation of the PMTs is inverted [73].

5.6.1 Detector Assembly

The bar detector layers of the detector are physically constructed as four “supermodules”, aluminum shells containing four layers of 2×2 “units” of bars. The supermodules ensure that each path remains rigid and aligned, and enables the alignment of the supermodules to be modified slightly to better account for path-length differences from particles traveling straight from the CMS IP.⁵

⁵To be more specific, if you imagine a particle traveling straight from the IP, a particle hitting a bar on the outside of milliQan will have a small angle imparted to it. The supermodules need to be angled slightly

A picture of a bar unit is shown in Figure 5.16. A plastic brace provides slots for each of the 2×2 bars, while fasteners on each side of the unit secure the bars to the brace. Grooves are cut into the top of the brace to provide a path for HV and signal cabling.




Figure 5.16: A bar unit before μ -metal wrapping. On the right, a finished unit is seated inside a supermodule.

A supermodule is assembled using four sets of units. Cables are fed through individual, circular slots at the end of the supermodule and laid through the grooves on top of the bar units. This ensures the supermodules can be tightly packed and fit onto the detector mount. A fully cabled supermodule is shown in Figure 5.17.⁶

such that this effect is accounted for. If unaccounted for, the size of the effect is approximately a 10-20% efficiency loss.

⁶These pictures are also taken during different trips, on the left was an intervention in May 2023 while the



Figure 5.17: A fully assembled, cabled supermodule. HV and signal BNC cables are laid across the top tray of the supermodule, while low voltage ribbon cables weaved through the side. (Left) UCD graduate student Jacob Steenis working on the supermodule. (Right) The author, analyzing data after assembly.

Supermodules are mounted to a “cage”, an aluminum mounting unit designed to store all supermodules and scintillator veto panels. An image of the cage when the veto panels were first mounted is shown in Figure 5.18. Side, top, and end panels are secured independently of the bar supermodules, important during periods of maintenance on the detector.

right was during the initial detector installation in Summer 2022; the right picture was probably taken by Sam Alcott.




Figure 5.18: The cage provides mounting points for all the detector components of the Run 3 bar detector. Here it is partially loaded, with the front, top, and side panels for the front half attached.

The cage is mounted to a large, rotating arm which was laser-aligned such that an angle of 43.1° is perfectly aligned to the CMS IP. The fully assembled, cabled, mounted bar detector is shown in Figure 5.21.



Figure 5.19: The Run 3 bar detector. One of the hardest challenges of the bar detector is cable management, as each of the 72 total scintillator channels is supplied by 3 cables, and all cables enter the detector from a similar point very close to the ground. The LV supply and LV splitter board can be seen near the bottom of the figure.

The bar detector is supplied by a CAEN HV module, with each channel connected directly to the HV supply at 1450 V. Low voltage is supplied by a custom LV splitter board and powered by a conventional bench LV source, as seen in Figure 5.21.

Signal outputs are read using a set of five fast, CAEN 16-channel V1743 digitizers, with clock signals synchronized between each board. The digitizer samples the output waveforms at 400 MHz over 1024 samples in a 2.56 μ s measurement window; this sampling rate is necessary to ensure pulse times can be reconstructed accurately inside a pulse timing window of 20 ns. The LVDS digital outputs for each digitizer board are chained and sent to an

external, custom trigger board, designed by Andy Haas, which handles the firmware for milliQan trigger.

The first supermodule of the Run 3 bar detector was assembled in August 2022. The bar detector began collecting commissioning data in October 2022, and began collecting physics data in June of 2023.

5.6.2 μ -Metal Shielding

One important lesson learned about bar detector performance during assembly was the PMT response to magnetic field during operation. R878 PMTs are susceptible to magnetic fields, as their gain mechanism relies on the propagation of electrons. If a magnetic field is applied across the PMT, two adverse effects occur: the initially emitted photoelectron from the photocathode may miss the first dynode, resulting in a decrease in measured quantum efficiency; and the electrons during amplification have trajectories modified by the magnetic field, causing fewer electrons to strike each dynode and reducing overall gain.

Initially, two layers of μ -metal shielding were applied to each bar: First, individual μ -metal sheaths were wrapped around each PMT during initial mounting construction before being inserted into the scintillator bar PMT mounts. Second, larger sheets of μ -metal were wrapped around the PMTs in each 2×2 unit.

However, in situ bench tests showed that performance of the bar units still showed significant dependence as a function of orientation in the ≈ 20 Gauss magnetic field in the milliQan cavern. As a remedy, new μ -metal shields were designed for each unit, with an increased thickness of 0.05 mm to 0.5 mm (Figure 5.20). These thicker μ -metal sheets were significantly less malleable and required machine operation to bend into the proper shape. However, performance with these thicker shields was much stronger, so all bar units in the detector were replaced with upgraded shields. This upgrade was performed in May 2023, immediately prior to the start of physics data collection in June 2023.




Figure 5.20: (Left) A picture of a standard bar unit with original μ -metal wrapping. (Right) A unit after an upgraded μ -metal shield was applied.

5.7 Run 3 Slab Detector

In addition to the bar detector, the Run 3 milliQan experiment also contains a slab detector. The slab detector consists of 48 scintillator slabs of dimension $5 \times 40 \times 60$ cm, arranged in a 3×4 configuration in each of four layers as in Figure 5.23.




Figure 5.21: The Run 3 slab detector.

The primary design goal of the slab detector is to maximize acceptance to mCPs. If mCPs are rarely produced but relatively higher in charge, detector acceptance rather than detector light yield becomes the dominant factor in mCP sensitivity.



Figure 5.22: The author testing a single milliQan slab in late 2023.

The slab geometry does not offer the same light guide properties as the long, narrow bar geometry. This means that in order to get comparable total light yield, four PMTs read out the same slab rather than one. These PMTs are placed on opposite ends of the 40 cm face of the slab to maximize the minor light guide effects of the rectangular geometry. A wrapped, mounted slab is shown in Figure 5.22. Given the poor leverage of mounting a PMT to the face of one of the slabs, PMTs are mounted to the slabs using rods to pull each PMT from opposite ends of the slab together towards the scintillator face. The underlying component scheme for the slabs is otherwise similar to the bars, with Tyvek wrapping, EJ200 scintillator, and R878 PMTs.




Figure 5.23: Mounting slabs onto the first layer. Helping mount is UCD postdoc Juan Tafoya.

The slab detector structure is a freestanding, non-rotating aluminum frame. As the structure is fixed, slabs are individually mounted to the structure directly without the need of additional support components such as a cage or supermodule.

With 48 slabs and 4 PMTs per slab, each needing signal, HV, and LV cabling, a significant challenge is reducing the total number of power inputs and signal outputs. HV splitter boxes are used to split a single HV supply to four channels. Similarly, a single LV ribbon cable with several cable heads supplies the LV power for the 6 channels sharing the same side of the slab for each row and layer.




Figure 5.24: The summing amplifier reduces the total number of readout channels from 192 to 96, a significant cost and complexity savings. It is powered by the same low voltage ribbon cable that powers the base readout amplifiers. Later versions were adjusted to more easily allow for gain tuning.

Signal outputs are combined using a summing amplifier. One challenge of summing separate PMTs into the same readout channel is the SPE peak calibration; if the pulse height for an SPE in channel 1 is 40 mV and channel 2 is 80 mV, an 80 mV reading is degenerate between two photons in channel 1 or one photon in channel 2. To account for this, summed

channels are “gain matched” using gain tuning handles in the summing amplifier as well as on each PMT readout amplifier base as discussed in Section 5.4. A summing amplifier, originally designed by UCSB REU student Emily Pottebaum, is shown in Figure 5.24.

5.7.1 Slab Detector Cd109 Calibration

One further complication of the slab detector as compared to the bar detector is the light calibration. MC studies of the bar detector have shown that light collection efficiency is approximately uniform along the length of a bar due to the light guide effect of the scintillator. However, the slab light collection efficiency geometry is nontrivial, so a significant number of calibration measurements are required to fully characterize it. In addition, the lower light collection efficiency of the slab even at positions near the front of a PMT face results in a poor source peak when measured using the typical source calibration isotope, Cd109.

To address these challenges, I designed a calibration procedure consisting of 7 measurements across the slab, to be executed after the PMTs on each side of the slab were summed and gain-matched. This summing allowed for an increase in light collection efficiency, resulting in a more distinguishable source peak. The slab calibration setup is shown in Figure 5.25.




Figure 5.25: (Left) A Cd109 source injects 22 keV gammas into the slab, generating a fixed energy deposit which is used to study the light collection efficiency of the slab at different points. (Right) Positioning the source at several locations is necessary to extract the total optical response of the slab.

The slab detector was fully mounted in late 2024 (Figure 5.26) and began collecting physics data in 2025.⁷ While some analysis of Run 3 slab data has occurred, I will focus the following chapter on only the completed bar detector analysis.




Figure 5.26: The fully mounted, cabled, and calibrated slab detector, completed in late 2024.

⁷It's still to be determined exactly when the first physics quality data will begin, but some time early in the 2025 dataset.

6 The milliQan Monte Carlo

In this chapter, I briefly describe the motivations and techniques for simulation of the MilliQan experiment. I'll outline the general structure of the simulation, describing its basic purpose and how it functions. I'll describe the relationships between the generation, propagation, and detector simulation components. I'll also source the various spectra and properties implemented in the simulation at various stages. Finally, I'll discuss results used from the milliQan simulation for both the Run 2 and Run 3 experiments.

The milliQan simulation framework contains three primary components:

1. **Event Generation:** mCPs are generated in CMS collisions using standard production modes for dilepton pairs, except with production couplings scaled by q^2 and mass factors adjusted to reflect the simulated mCP. Muons generated in CMS collisions (“Beam Muons”) are also generated in this stage. This is described in Section 6.1.
2. **Signal and Muon Propagation:** mCPs and muons are propagated from the CMS IP, through a detailed model of the CMS detector, then through 17 m of rock to the milliQan detector. 2 m before the face of the detector, four-momenta of the particles are saved. These values, as well the event weights, are passed to the detector simulation. This is described in Section 6.2.
3. **Detector Simulation:** Using Geant4, a detailed model of the milliQan detector is used for simulation of particle interactions from muons and mCPs. The energy and optical response of the simulation have been rigorously modeled against bench performance, and channels are individually calibrated to accurately reflect experimental performance. This is described in Section 6.3.
4. **Pulse Injection:** Outputs from the Geant4 simulation are converted into waveforms, after which they are analyzed using the same tools as for experimental data. This

procedure is described in Section 6.4.

The following sections will discuss performance for the milliQan experiment. However, in addition, the Geant4 detector simulation has been used to produce signal and background estimates for the FORMOSA, SUBMET, and FLAME/LANSCE-mQ experiments.

6.1 Event Generation

All standard model processes which produce dilepton pairs can produce millicharged particles. The full spectrum of production modes for $q/e = 1$ are shown in Figure 6.1, and listed below.




Figure 6.1: MilliQan mCP production modes and cross-sections as a function of mass.

At lowest mass, Dalitz decays from the π and η mesons dominate. At intermediate mass, meson decays of the ω , ρ , ϕ , J/ψ , and $\Upsilon(1S)$ dominate. At highest masses, Drell-Yan production dominates. Production in all modes scales as q^2 , with mass factors for each production mode as described in the following Sections 6.1.1-6.1.3. I will summarize in each case where mCP generation differs from SM generation; a more complete description is available at [74].

6.1.1 Dalitz decays

The differential distribution for the Dalitz decay is given by [75]:

$$\frac{d\Gamma(\pi^0 \rightarrow e^+e^-\gamma)}{ds} = \frac{2\alpha}{3\pi s} \left(1 - \frac{s}{m_\pi^2}\right)^3 \left(1 + \frac{2m_e^2}{s}\right) \sqrt{1 - \frac{4m_e^2}{s}} |F(s, 0)|^2 \Gamma(\pi^0 \rightarrow \gamma\gamma) \quad (6.1)$$

In a Dalitz production of a mCP, the electron mass m_e is replaced by the mCP mass, and the lepton coupling α is rescaled as $q^2\alpha$. The equations are then integrated over the invariant mass distribution to obtain decay branching ratios. Note that the diphoton decay width, described by $\Gamma(\pi^0 \rightarrow \gamma\gamma) \propto \alpha^2$, is not charge-rescaled, as the only the lepton vertex is impacted by the mCP charge coupling.




Figure 6.2: ALICE measurements of the Monash 2013 tune shows good agreement, except overproduction at low P_T [76].

Dalitz decays are generated in the minimum bias Pythia8 [77] event generator under the Monash 2013 tune [78]. For η and π^0 production, we use ALICE measurements of the Pythia8.2, Monash 2013 tune to validate generation, as in Figure 6.2 [76]. We find that this

tune agrees well but overpredicts production at $P_T < 3.0$ GeV, so we divide by a rescaling function from $[0.5, 3.0]$ GeV which takes the value 2.0 at 0.5 and 1.0 at 3.0.

6.1.2 Vector meson decays




Figure 6.3: ALICE measurements at 2.76 TeV for the ρ (left, [79]) and 7 TeV for the ω (right, [80]), compared to the Monash 2013 tune. Agreement for ω is taken to match ρ at lowest P_T , where data is not available.

According to the Van Royen-Weisskopf formula, the decay width of vector mesons into two leptons is well described by:

$$\Gamma(V \rightarrow ll) \propto \alpha^2 \sqrt{1 - \frac{4m_l^2}{m_V^2}} \left(1 + \frac{2m_l^2}{m_V^2}\right) \quad (6.2)$$

Replacing the lepton mass m_l with the mCP mass m_χ and scaling $\alpha^2 \rightarrow q^4 \alpha^2$ for mCPs, we find that:

$$\frac{\Gamma(V \rightarrow \chi\chi)}{\Gamma(V \rightarrow ll)} = q^4 \frac{\sqrt{1 - 4x_\chi} (1 + 2x_\chi)}{\sqrt{1 - 4x_l} (1 + 2x_l)} \quad (6.3)$$

where $x_\chi = m_\chi/m_V$ and $x_l = m_l/m_V$.

We use the Pythia8, Monash 2013 tune to produce vector mesons. For the ρ and ω we use ALICE measurements to validate the tune, as in Section 6.1.1 [79, 80]. Good agreement is shown except at lowest P_T , as in Figure 6.3. We use a similar reweighting scheme as before, except using a function which takes the value 2.0 at 0.5 GeV and 1.0 at 1 GeV.

For the ϕ meson, ATLAS comparisons under the Pythia8 Monash tune suggests underproduction at all values of P_T [81], whereas the Pythia6 DW tune show good agreement [82]. Uniquely for the ϕ , we use Pythia6.

Charmonium direct production spectra for the J/ψ and $\psi(2S)$ are taken from theoretical calculations [83, 84, 85]. Production via B-decay is estimated using the semi-analytical framework FONLL [86, 87], using the NNPDF30_nlo_as_0118 as PDFs.

Finally, Υ production was measured using 7 TeV data collected by ATLAS at $P_T < 20$ GeV and 13 TeV data collected by CMS at $P_T > 20$ GeV [88, 89]. To scale the two distributions into agreement, production cross section measurements of the Υ from CMS at $|\eta| < 2.0$ and from LHCb at $2.0 < y < 2.5$ were used [89, 90].

6.1.3 Drell-Yan

Drell-Yan production is achieved by modifying electron couplings in `MadGraph5_aMC@NLO` to the γ and Z to $e \rightarrow ee$ and $-\epsilon \tan\theta_W \rightarrow -\epsilon \tan\theta_W$, respectively [91]. As Drell-Yan generation is the most CPU-intensive of any generation mode, it was only considered above invariant mass of 2 GeV, below which it is subdominant to decay modes.

6.1.4 mCP production

To model mCP production, for each production mode, we use the following procedure:

1. For decays, produce the SM particle, sampling from a differential cross-section distribution $d\sigma/dP_t$ to generate the SM kinematics
2. Compute the mCP event cross-section as a function of mass and charge using Equa-

tion 6.4

3. Produce the mCP using the SM decay modified by mCP mass kinematics

$$\sigma_{mCP} = \frac{q^2 \times \sigma \times BR \times L}{N} \quad (6.4)$$

Here q^2 is the mCP charge, σ is the SM cross-section, BR is the production branching ratio, L is the luminosity, and N is the number of generated events. Most modes also have an η, ϕ requirement to reduce CPU runtime; events which do not satisfy this requirement are still recorded in N . In Run 3, the acceptance requirements were $-0.1 < \eta < 0.3$ and $-0.1 < \phi < 2.0$.

6.1.5 Muon Generation

In addition to mCPs, muons are generated in CMS collisions, as muons produced from the IP serve as a good SM validation of detector alignment, performance, and triggers. We consider four production modes in decreasing order of event weight: heavy flavor decays, π/K decays in flight, W production, and Drell-Yan.




Figure 6.4: Beam muon P_T distribution for each production mode (left). FONLL $b\bar{b}$ production compared to CMS measurements in the low- P_T range underestimates measurement by around a factor of 1.25 (right, [92]).

Generation for heavy flavor (b,c) decays is performed using FONLL at 13.6 TeV, including an additional production correction of 1.25 ± 0.25 to the heavy flavor decay mode to account for known underproduction of $b\bar{b}$ in FONLL compared to CMS measurements [92]. The P_T distribution and FONLL underestimates are illustrated in Figure 6.4. Heavy flavor production represents 82% of the total beam muon flux through milliQan.

Decays from π, K are generated using Pythia 8 in the Monash 2013 tune. Decays are generated in a cylinder of radius 2 m and total length 8 m, accounting for central values of η which are within the radius of the CMS ECAL. Decays in flight represent 12% of the total beam muon flux through milliQan.

W and DY production of muons are generated using MadGraph5_aMC@NL0. W production represents 5% of the total beam muon flux, and DY production represents $< 1\%$.

6.2 Propagation

6.2.1 RK4 Integration

After event generation, a fourth-order Runge-Kutta (RK4) integrator propagates the particles through a simplified (but precise) model of the CMS detector, the CMS magnetic field, and the 17 m of rock between CMS and milliQan [93].

The RK4 integration method is used to estimate the position $\vec{x}$ and momentum $\vec{p}$ under an initial condition described by $d\vec{p}/dt = f(t, \vec{p}(t))$. The integration method is given by [94]:

$$\begin{aligned}
 w_{k+1} &= w_k + \frac{1}{6} \left(K_1 + 2K_2 + 2K_3 + K_4 \right) h, \\
 K_1 &= f(x_k, w_k), \\
 K_2 &= f\left(x_k + \frac{1}{2}h, w_k + \frac{1}{2}K_1h\right), \\
 K_3 &= f\left(x_k + \frac{1}{2}h, w_k + \frac{1}{2}K_2h\right), \\
 K_4 &= f(x_k + h, w_k + K_3h).
 \end{aligned} \tag{6.5}$$

The equation of motion used during propagation are given by:

$$\frac{d\vec{x}}{dt} = \frac{\vec{p}c}{E} = \frac{\vec{p}c^2}{\sqrt{((\vec{p}c)^2 + (mc^2)^2)}} \quad (6.6)$$

$$\frac{d(\vec{p}c)}{dt} = qc\vec{v} \times \vec{B} \quad (6.7)$$

The magnetic field is constructed using a map of the CMS magnetic field sampled in a grid of 10 cm by 10 cm points.

6.2.2 Interactions in Matter

dE/dx energy loss and multiple scattering are implemented to simulate particle interactions in matter during propagation.

The Bethe-Bloch formula is used to model energy loss by moderately relativistic charged heavy particles [95]:

$$\left\langle -\frac{dE}{dx} \right\rangle = Kz^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \quad (6.8)$$

After dE/dx is computed, we multiply by $dx = vdt$ to obtain the energy loss for each step.

The method used to model multiple scattering is described in Refs [93, 95, 96]. If we define:

$$\theta_0 = \theta_{plane}^{rms} = \frac{1}{\sqrt{2}} \theta_{space}^{rms} \quad (6.9)$$

then we can model the θ_0 as:

$$\theta_0 = \frac{13MeV}{\beta cp} z \sqrt{\frac{x}{X_0}} \left[1 + 0.038 \ln \left(\frac{xz^2}{X_0 \beta^2} \right) \right] \quad (6.10)$$

where p , βc , and z are the momentum, velocity, and charge number of the incident particle.

As described in Figure 6.5, in addition to angular deflection θ_{plane} , there is also a correlated transverse deviation y_{plane} described by the correlation coefficient $\rho_{y\theta} = \sqrt{3}/2$.




Figure 6.5: A particle incident in the plane of the figure is subjected to multiple scatters as it travels through matter [95].

To implement these deflections in the simulation, two independent $\text{Gaus}(0, 1)$ random variables z_1, z_2 are chosen, related to $\theta_{plane}, y_{plane}$ by:

$$\begin{aligned}
 y_{plane} &= z_1 x \theta_0 \sqrt{1 - \rho_{y\theta}^2} / \sqrt{3} + z_2 \rho_{y\theta} x \theta_0 / \sqrt{3} \\
 &= z_1 x \theta_0 / \sqrt{12} + z_2 x \theta_0 / 2; \\
 \theta_{plane} &= z_2 \theta_0.
 \end{aligned} \tag{6.11}$$

This procedure is repeated for both directions orthogonal to the direction of motion, $\theta_{plane,x}, \theta_{plane,y}$, which are independent and identically distributed.

6.2.3 Model of CMS

The material model of CMS was updated for Run 3 to more precisely describe the material composition and distribution during propagation. This is visualized in Figure 6.6 and summarized below:

Radial range $r(\text{m})$	Density	Material	Subsystem
$1.29 < r < 1.52$	$8.3\text{g}/\text{cm}^3$	PbWO_4	ECAL
$1.8 < r < 2.76$	$7.87\text{g}/\text{cm}^3$	70/30 Cu/Zn	HCAL
$3.423 < r < 3.736$	$7.87\text{g}/\text{cm}^3$	Fe	Solenoid
$3.868 < r < 4.048$	$7.87\text{g}/\text{cm}^3$	Fe	Yoke, TC
$4.615 < r < 4.9$	$7.87\text{g}/\text{cm}^3$	Fe	Yoke, L1
$5.35 < r < 5.97$	$7.87\text{g}/\text{cm}^3$	Fe	Yoke, L2
$6.375 < r < 6.995$	$7.87\text{g}/\text{cm}^3$	Fe	Yoke, L3
$13.2 < r < 14.4$	$2.3\text{g}/\text{cm}^3$	Concrete	Support
$14.4 < r < 31.4$	$2.514\text{g}/\text{cm}^3$	Rock	Rock
$31.4 < r < 32.8$	$2.3\text{g}/\text{cm}^3$	Concrete	Support




Figure 6.6: The CMS detector is modeled using precise thicknesses and densities according to the material composition of each layer.

The ECAL is modeled as 0.23 m of PbWO_4 , consistent with the crystal length in the barrel where milliQan is located. The HCAL is modeled as a mixture of $8.53 \text{ g}/\text{cm}^3$ Cu/Zn absorbers and low-density plastic scintillator ($\sim 1.0 \text{ g}/\text{cm}^3$) [38]. A typical radial absorber thickness is 0.858 m, with the remainder of the 0.96 m modeled HCAL thickness filled with plastic scintillator tiles, but there are segments of varying thickness depending on the HCAL region. An average density of this configuration is $(1/0.96 \text{ m})(0.858 \text{ m} \times 8.53 \text{ g}/\text{cm}^3 + 0.102 \text{ m} \times \text{g}/\text{cm}^3) = 7.43 \text{ g}/\text{cm}^3$. Since this density is close to the material composition of the surrounding steel yokes, the density of iron was used as $7.87 \text{ g}/\text{cm}^3$ and the uncertainty on the density is used as a material systematic.

The magnet return yoke thicknesses and radii are extracted from a description of the CMS magnetic system [97]. In addition to the thickness of the yoke, an additional 5 cm of

steel is added to account for all other subsystems, cooling, and outside scaffolding between the CMS detector and the rock of the surrounding cavern.

The thickness of the intervening rock is set to 17 ± 0.5 m, a value taken from an internal model of the CMS cavern and drainage gallery produced for the milliQan LOI, as in Figure 6.7 [98].




Figure 6.7: A model of the drainage gallery’s position with respect to the CMS cavern estimated a rock thickness of 17 m [98].

The material composition of the rock is estimated using a geotechnical survey of a borehole extracted near ATLAS, shown in Figure 6.8 [99].

Formation	Depth [m]	Rock Unit	w [%]	ρ [t/m ³]	v_p [m/s]	I_s [MPa]	σ_{ci} [MPa]	RQD	Description
	28.20							62	grey, massive, hard
	85.00		6.5	2.44	3691	0.20			97 Fine to medium-fine-grained sandstone, grey, massive, hard Fine to medium-fine-grained sandstone, grey, massive, soft Sandy marl, brown, stratified, hard 92 Moltey marl, grey, redish stains, slickensided and fissured
	86.70								
	87.10		4.8	2.57	4219	0.98			
	87.60								
	88.20		6.0	2.54	3050	0.30			
	88.50								
	90.40		8.7	2.45	2975	0.14			
	90.70		8.9	2.41	2062				

Figure 6.8: A borehole with material densities and classifications, extracted near ATLAS [99].

We assume this is approximately representative of the material composition and density near CMS for the same depth, as the geotechnical survey identified the primary difference between ATLAS and CMS geologic conditions as different bedrock depths and material compositions at or above 50 m, significantly above the depth of milliQan. However, we derive systematic uncertainties on the rock composition and density to reflect possible geological differences between the measured borehole and our position near CMS.

With segment lengths ℓ_i (m) and rock species densities ρ_i (g/cm³), the weighted average of the densities is given by:

$$\bar{\rho} = \frac{\sum_i \ell_i \rho_i}{\sum_i \ell_i} = \frac{(1.7)(2.44) + (1.8)(2.57) + (1.9)(2.54) + (0.3)(2.43)}{1.7 + 1.8 + 1.9 + 0.3} = 2.514 \text{ g/cm}^3. \quad (6.12)$$

Taking the spread 2.44–2.57 g/cm³ in rock species ρ_i , we assign a systematic uncertainty on the rock species of 0.065 g/cm³.

A further systematic is derived on the rock species density. As described in Figure 6.9, uncertainty on the density of a species is generally ± 0.04 for Med.-strong sandstone and all kinds of marls, the species present at 85–90 m depth near ATLAS. Therefore, we use

this value as the density measurement uncertainty. The resulting total rock density with systematic uncertainty is then given as $2.514 \pm 0.076 \text{ g/cm}^3$.

Variable	Units	MARLS			SANDSTONES		
		1. Very weak	2. Weak	3. Med.-weak	4. Weak	5. Med.-strong	6. Strong ¹
w	[%]	8.6 ± 1.2	6.7 ± 1.2	5.0 ± 1.21	7.8 ± 1.8	4.1 ± 1.02	3.0 ± 1.09
ρ	[t/m ³]	2.44 ± 0.04	2.5 ± 0.04	2.55 ± 0.04	2.38 ± 0.06	2.56 ± 0.04	2.60 ± 0.03
w_L	[%]	51.4 ± 6.7	48 ± 5.5	(34.8 ± 3.7)	-	-	-
w_P	[%]	25 ± 2.6	23 ± 1.7	(20.7 ± 1.9)	-	-	-
σ_{ci}	[MPa]	3.7 ± 1.5	8.0 ± 2.87	15 ± 4.9	10.8 ± 3.2	22.9 ± 5.5	46.1 ± 14.2
E_s	[MPa]	340 ± 241	690 ± 534	1960 ± 1508	1230 ± 624	3420 ± 1490	9417 ± 4750
$\sigma_{i,b}$	[MPa]	0.58 ± 0.31	1 ± 0.43	1.76 ± 0.7	0.8 ± 3.5	2.11 ± 0.77	3.3 ± 1.5
I_s	[MPa]	0.22 ± 0.13	0.33 ± 0.17	0.57 ± 0.27	0.26 ± 0.13	0.7 ± 0.27	1.59 ± 0.6
V_{pl}	[m/s]	2240 ± 425	2500 ± 444	3040 ± 536	2455 ± 456	3340 ± 547	3955 ± 524
RQD	[-]	71 ± 27	91 ± 13	97 ± 8	96 ± 8	98 ± 4	99 ± 2
ISRM grade		$R1$	$R2$	$R2$	$R2$	$R2 - R3$	$R3$

Figure 6.9: Density and uncertainty for rock species near CMS [99].

Finally, though the original model of the CMS and drainage gallery included the rock thickness, an engineering description of the construction of the CMS gallery suggests an additional 40 cm of primary shotcrete support in the drainage gallery and CMS cavern and 80 cm of secondary concrete lining. This adds a total of 1.6 m of concrete of approximate density 2.3 m.

After propagation, mCPs and beam muons which pass through a $1 \text{ m} \times 1 \text{ m}$ plane located 2 m before the milliQan detector are accepted into the Geant4 full detector simulation.

6.3 Geant4 Detector Simulation

In this section, I detail the implementation and validation of this full detector simulation.

6.3.1 Geant4

Geant4 is a C++-based toolkit used to simulate the passage of particles through matter [100]. Released originally in 1998, Geant4 was designed as a general purpose detector simulation framework to simulate the passage of particles and radiation through matter, providing geometry, tracking, physics models, and detector response. Geant4 has since seen

the development of other dedicated packages that expands its physics coverage to other particle physics applications, such as optical photon tracking, as well as new applications such as in medicine and astrophysics [101, 102]. Geant4 models all kinds of physical extremes, from sub-eV thermal neutron capture to TeV-scale LHC interactions [103].

The guiding principle of Geant4 is the abstraction of all physical processes, as all processes are treated as the same from a tracking point of view [104]. Because of this abstraction, new physical processes can be included just by assigning the details of interaction and cross section to any relevant particles.

During particle “tracking” (the process of propagating a particle through the environment), each possible physics process proposes an interaction length to the Geant4 process manager, sampling from a distribution of possible interaction lengths according to its cross-section. The shortest proposed interaction length is selected from this list of candidate processes (including transportation, the lack of an interaction). The particle is then propagated along its space-time point by a calculated step length and the selected process is invoked, changing the parent particle’s position and momentum and producing secondary tracks as necessary.

6.3.2 Simulation Implementation

A typical Geant4 simulation consists of the following key components:

- **G4VUserDetectorConstruction:** Here, the detector and all material components and properties are defined. Detector primitives are composed of:
 - **G4VSolids** (e.g. G4Box), raw shapes which can be added and subtracted as boolean objects and which define the core geometry of each object
 - **G4LogicalVolumes**, which are assigned to G4VSolids and contain material properties. Logical volumes track interactions and can be assigned particle hit and energy deposit information.
 - **G4PVPlacement:** The final “Physical Volume Placement” assigns a position and

rotation of the logical volume and underlying G4VSolid into the detector world. Placements are made inside other logical volumes; for example, in the milliQan bar detector, there are two nested placements: the outer logical volume contains the wrapping; the intermediate volume contains an air gap, and the inner-most volume contains the scintillator.

- **G4VUserPhysicsList:** Described in Section 6.3.3, a physics list describes all interactions being considered in a Geant4 simulation.
- **G4VUserPrimaryGeneratorAction:** These are the particles which are the inputs to the simulation. The milliQan signal simulation uses a custom generator to assign mCP mass and charge and load in custom mCP four-momenta. However, for other background studies such as of cosmic ray muons, a parameterized distribution can be specified using General Particle Source commands in an initialization macro, such as “/gps/particle mu-” or “/gps/ene/mono 4 GeV”. GPS commands are discussed in greater detail here [105].

Several additional custom classes are defined for the milliQan simulation. These include data structures that store tracking information for muons, gammas, electrons, and mCPs, and hit information for PMTs and scintillator volumes (the latter score energy). A complete description of the user files and their functions is described in Appendix A.8.

Simulation outputs are stored in a custom ROOT output which includes class information such that, after a shared library is loaded, methods from the Geant4 classes can be used to access information about tracks, hits, and event quantities. More information about the content of the simulation output is described in Appendix A.9.

6.3.3 Physics Implementation

The Geant4 simulation used for milliQan is configured with a set of specialized models used to optimize electromagnetic, hadronic, and optical physics accuracy.

The base physics list for the Geant4 milliQan simulation is the `ShieldingPhysicsList` [106]. This list was chosen to accurately model interactions in passive and active shielding, including scintillator panels, the rock of the surrounding cavern, and the Pb blocks between layers in Run 2. In addition, the high precision neutron interactions and radioactive decays modeled in the `Shielding` list have allowed the milliQan simulation to be adapted for a diverse range of other applications: the FLAME detector in the high-intensity neutron environments at LANSCE, high energy neutron scatter background modeling for the S-particle search in the milliQan slab detector, and activated x-ray detectors in a radioactive environment for beamline diagnostics devices. `G4EmStandardPhysics_option4`, the highest precision Geant4 EM physics library, is used to further ensure well-modeled EM interactions.

mCP physics are implemented by adapting the `G4Monopole` class, which offers methods to assign a configurable mass and electric charge. Ionization (`G4hIonisation`) and multiple scattering (`G4hMultipleScattering` and `G4WentzelVIModel`) are assigned to mCP objects as candidate processes during tracking.

6.3.4 Optics implementation




Figure 6.10: Photon detection and scintillation spectra for the (left, [107]) R878 PMT and (right, [108]) EJ200 scintillator. The peaks are overlapping, suggesting maximal efficiency for scintillation photons.

The optical physics list used for surface reflections is the `UNIFIED` model. In this model, surface interface and material are configured for each interface. In the milliQan simulation,

the “groundteflonair” option is used to model the angular distributions for reflection off Tyvek. One characteristic of note of Tyvek and Teflon that is accounted for using this approach is their angular dependence on diffuse vs. specular reflections. That is, for sharper angles of incidence, light is more likely to be reflected specularly (i.e. geometry preserving) rather than diffusely (isotropic) [109]. Meanwhile, surface interfaces are set to dielectric_metal such that reflectivity values can be set directly. For the Tyvek wrapping, a representative reflectivity is $R = 0.97$, a value which accurately described bench measurements and is consistent with literature [110].

Material optical properties for the scintillator light yield (10000 photons/MeV), refractive index (1.58), bulk attenuation length (3.8 m), and scintillation time constant (2.1 ns) were taken from the datasheet specifications for EJ200 scintillator, a scintillator with “a long attenuation length and a scintillation spectrum well-matched to common PMTs” used in milliQan and related mCP experiments [108].

R878 PMT quantum efficiency is modeled after a wavelength-dependent quantum efficiency scan [107]. The scintillation spectrum of EJ200 is taken from the datasheet, and both PMT and scintillation spectra are shown in Figure 6.10. The scintillation peak is in the central region of the radiant sensitivity of the R878, suggesting good optical coupling.

Propagation of optical photons is modeled in the G4SteppingAction under the following procedure:

- Scintillation photons generated from energy deposits in the scintillator are emitted isotropically.
- A scintillation photon is propagated through a detector volume. The photon is tracked through the scintillator until it is either absorbed or reaches a boundary according to the G4ProcessManager. If a photon is absorbed in the bulk, its track terminates.
- A photon can reach three kinds of boundaries in the milliQan simulation: scintillator-air, air-Tyvek, and air/scintillator-PMT.

1. At the scintillator-air interface, the photon undergoes total internal reflection or

transmission according to its angle of incidence. If the photon undergoes total internal reflection, its track is removed from propagation at a 1% probability according to a 1% scintillator surface roughness parameter. This parameter showed good agreement with bench measurements and was motivated by the typical angular deviation at a surface feature size of 400nm, as described in Appendix A.10.

2. At the air-Tyvek interface, the photon is rescattered in a semi-diffuse pattern according to the modeled angular distribution of teflon in the UNIFIED groundteflonair finish. A reflectivity parameter of Tyvek of 0.97 absorbs photons which undergo reflection at a 3% probability.
3. At the air/scintillator-PMT interface, the PMT quantum efficiency is applied. When a photon encounters the PMT, the simulation kills the photon track and applies the (wavelength-dependent) quantum efficiency. Whether the photon is successfully detected or not, the photon track is removed the moment this quantum efficiency is applied. Additionally, this quantum efficiency is modified by an angular-dependent quantum efficiency spectrum, derived from measurements of the angular dependence of QE in the Hamamatsu PMT Handbook and shown in Figure 6.11 [68].




Figure 6.11: Angular Detection Efficiency of Hamamatsu PMTs. In order to interact with the photocathode, the photon must pass through this angular detection efficiency, which is implemented in the `mqSteppingAction` before the application of the PMT quantum efficiency. Note that the angular detection efficiency is approximately described by a $\cos(\theta)$ distribution [68].

6.3.5 Air Gap Modeling

A key lesson in modeling optical propagation for the milliQan bars was the inclusion of a small, < 1 mm air gap between the Tyvek wrapping and the plastic scintillator. Without an air gap, all reflections apply the reflectivity value for Tyvek and scatter according to the Tyvek semi-diffuse reflection spectrum. If an airgap is included, an air-scintillator interface appears which generates total internal reflection without absorption at the Tyvek interface and without significant semi-diffuse rescattering.

Early studies showed that the air gap actually models the majority of the “light guide” effect of the scintillator bar, wherein light with an angle emitted under the critical angle of transmission of the bar is repeatedly internally reflected, causing significant light collection

to the coupled PMT.

Additionally, this air gap modeling is physically reasonable, as wrapping with Tyvek is done without extra mechanical or chemical treatment to ensure no air gap, such as a spray or glue treatment to the scintillator.

An air gap may also appear between the PMT and scintillator in poorly optically coupled PMTs. This effect is treated as a loss in overall transmission at the PMT interface, a constant which is absorbed into the quantum efficiency scaling of the PMT during calibration.

6.3.6 Optical Validation

To validate the optical modeling performance of the Geant4 detector simulation, we have performed a variety of sim-data comparisons of light collected along different points along the scintillator bars.¹ One comparison was made using a muon telescope, explicitly recreated in simulation including a truncated $\cos^2(\theta)$ distribution and modeled trigger scintillator paddles. The results are shown in Figure 6.12. Data measurements were converted to NPE using the area-NPE scaling value of the point nearest to the PMT as it is the least exposed to the bulk light properties of the bar and wrapping. The slope of a fit to each measurement point along the scintillator was -16.5 ± 1.6 in data and -15.3 ± 1.6 in MC, showing good agreement.

¹Studies repeated by Dariush Imani for the slab detector with good results.




Figure 6.12: A comparison of cosmic telescope measurements of light yield at 5 different positions in simulation and data. The two distributions show good agreement.

6.3.7 milliQan Full Simulation Construction

The milliQan simulation geometry is constructed using a parameterization of individual bars, panels, and slabs. Scintillator is coupled to PMTs by cutting a hole in the simulated wrapping and coupling the PMT directly to the scintillator material.

milliQan bars are placed inside of simulated supermodules. This supermodule has a 4 mm-thick Al frame with measurements reflective of the physical supermodule the milliQan bars are mounted in. As this supermodule frame is the passive material closest to each milliQan bar, it is the most relevant to simulate to generate any showers which may be created by the passage of charged particles. This is the only aspect of the detector frame which is explicitly simulated in the milliQan detector simulation. A simulated μ -Metal shield around the PMTs was explored, as well as the placement of a nearby dense object representing the support structure, but neither had a significant impact on generated signal- or cosmic-like

background activity.

The four supermodules of the bar detector are placed into a horizontal detector volume at positions consistent with in-situ measurements, then rotated at $\theta = 43.1^\circ$, consistent with the milliQan geometry. The parent detector volume is constructed as a trapezoid to carefully avoid overlaps with the cavern floor.² Active shielding scintillator slabs are constructed in a similar manner as for the bars, including an air gap, and are appropriately placed according to in-situ measurements around the bar detector.

The cavern is modeled as a cylinder with a rectangular prism cut out of the floor, inspired by pictures and measurements of the milliQan drainage gallery dimensions. The cylinder is given an inner diameter of $2.9 \text{ m} \times 15 \text{ m}$, and the cutout is given width of 2.5 m , height 0.27 m , and length 15 m . The cavern has an overall wall thickness of 1 m . This wall thickness is chosen to accurately model EM and hadronic interactions in the cavern walls which penetrate into the gallery and interact with the milliQan detector. These interactions are extremely relevant to milliQan; they form showers from cosmic muons, beam muons, and other penetrating particles which can fake signal-like events. The cavern is made from standard concrete of density 2.3 g/cm^3 .

The milliQan slab detector is constructed similarly to the milliQan bar detector. It is placed in a cavern of identical dimension as the bar detector, and the PMTs are coupled to the scintillator in the same way. However, the slabs are placed via parameterization directly into the cavern, rather than into a parent detector volume.

Finally, there is also accompanying magnetic field in the cavern which is given by $B(x, y, z) = (-1.61, -1.73, 21.57) \text{ Gauss}$, where $+x$ points vertically upwards, $+z$ points to the right with respect to Figure 6.13, and $+y$ points out of the page with respect to Figure 6.13. This magnetic field was disabled for Run 3 milliQan studies, as early studies

²One important component the detector construction is the consideration of material overlaps: G4PVPlacements must be done carefully to ensure that daughter volumes inside larger parent volumes do not extend beyond the boundaries of the parent volume. When this happens, internally tracked particles such as optical photons encounter undefined behavior and can e.g. transmit directly through bar wrapping, ejecting unphysically into the cavern. This is visually reflected in milliQan detector visualizations by the presence of cyan particles exiting scintillator volumes freely.

showed little impact on signal or muon trajectories.³

The fully constructed milliQan bar detector is shown below in Figure 6.13, displaying the result of a sample cosmic muon shower event and a sample mCP signal simulation.




Figure 6.13: A cosmic muon generates several gammas in a secondary EM shower, some of which strike the detector (left). A particle of charge $0.1e$ passes through the milliQan detector (right). As this particle is of relatively high charge, it deposits photons in both end panels and all four layers.

6.4 Pulse Injection

To remove purely electronic differences between MC and data, we use a “pulse injection” framework to explicitly convert simulated NPE to pulse waveforms.

First, a pulse template and SPE area distribution was determined by pulsing a low-intensity LED on an R878 PMT. This waveform template was then stretched in the time dimension such that the number of rise and fall samples for a pulse of fixed amplitude was equivalent to in-situ measurements of the milliQan detector; the result is in Figure 6.14. This sample was also used to measure the typical SPE area distribution by performing a 2-Gaussian fit and extracting the parameters of the 1-NPE fitted distribution. These results are shown in Figure 6.15.

³However, recent discussions of the shape of beam muon tails suggest most beam muons may be low momentum, as in Figure 6.4, and therefore more susceptible to magnetic field effects. Future studies of beam muon rate validations, which generally observe an overestimate in MC, should re-examine the impacts of the magnetic field on beam muon efficiency in MC.




Figure 6.14: This waveform is used as the basic template when converting simulated NPE into pulses during injection.




Figure 6.15: A low-intensity LED exposes a R878 PMT to 1- and 2-photon peaks. A double Gaussian fit is performed, and the 1-SPE Gaussian is used to generate the area distribution sampled from during pulse injection.

To perform pulse injection, each simulated photoelectron is first subject to a detector calibration to apply realistic quantum efficiencies and detector effects to match each physical channel. These values were obtained by measuring the peak location of a simulated Cd109 sample with an intentionally high quantum efficiency, then scaling the mean of the distribution down to match an equivalent measurement performed for each experimental channel. The

data scale factors are described by:

$$S = \frac{\langle NPE_{Cd109,Data} \rangle}{\langle NPE_{Cd109,MC} \rangle} \quad (6.13)$$

The simulated Cd109 peak was located at approximately $\langle NPE_{Cd109,MC} \rangle = 11$, and the average data value is given by $\langle NPE_{Cd109,Data} \rangle = 7.65$. The distribution of data calibrations is shown in Figure 6.16. The MC value is intentionally chosen to be higher than the data distribution such that data can be scaled down by randomly rejecting simulated NPE in a channel according to its data scale factor.




Figure 6.16: The distribution of NPE peak values for the 64 milliQan bar detector channels. MC, set at an equivalent NPE peak value of 11, is rescaled to match the performance of the appropriate channel.

Photoelectrons passing data rescaling are converted to the representative pulse template from Figure 6.14, then scaled by sampling from the distribution in Figure 6.15 such that its pulse area matches a randomly sampled pulse area from the SPE area distribution. In this way, each SPE passing data rescaling becomes a typical waveform with a typical area

distribution. The resulting pulse is then shifted according to the time delay at which the photon was detected in MC.

A pulse is produced from each channel after all simulated photoelectrons have been injected, with each photon's waveform summed after injection. After, 1 mV of Gaussian noise is overlaid on top of the pulse to simulate electronic noise; this is representative of the typical scale of electronic noise in the milliQan channels.

To avoid large noise contributions in the sidebands when many photons are pulse injected in the same channel, only the 400 ns segment centered on the peak value of the pulse template is used. The average pulse template includes small fluctuations (< 0.1 mV) in regions far before the start of the pulse. These don't impact pulse reconstruction for small ($NPE < 100$) pulses, but for very large energy deposits these values can cross the pulse reconstruction threshold as many templates with the same small fluctuations are laid on top of each other.

A truncation of the window considered for the average waveform template partially solves this problem, but for channels with very large NPE, the number of rise samples of the pulse becomes sensitive to how early the truncation occurs, as in Figure 6.17.

This sets a maximum value on the amount of area that can be reconstructed from pulse injection, but as nonlinear PMT effects such as space-charge saturation aren't encoded into the injection, very high energy deposits with long saturation times are poorly modeled regardless. Accordingly, this pulse injection process is most effectively used in studies of signal sensitivity and background effects of small ($NPE < 100$) energy deposits.




Figure 6.17: For signals of very large amplitudes, the number of rise samples for the pulse becomes a strong function of the early behavior of the pulse template. A PMT slew rate modeling, not included in this implementation, would be required to more properly model these effects for large pulses. In this implementation, a truncation was chosen that best reproduced rise sample behavior when compared to data.

6.5 Beam Muon Validation

An inclusive validation of the simulation framework at all stages— including generation, propagation, detector simulation, and injection— can be obtained from a validation study of the rate and angular distribution of beam muons.

6.5.1 Beam Muon ID

First, thresholds used to identify through-going muons are determined for the front and back panels. Events satisfying $\text{area} > 300\text{k pVs}$ in each bar layer are identified using the beam muon skim; an event with a large deposit in each layer is consistent with a through-going beam muon path through the central region of the detector. The unbiased area distributions

of the front and back panels are then plotted in Figure 6.18. An area requirement of 300k pVs in the endcap panels has high efficiency for these tagged beam muon events.

Similarly, in Figure 6.19, the maximum area distribution of each layer is plotted for events in which each other layer and both endcaps satisfies area > 300 k pVs. A selection of 1000k pVs in the bars is fully efficient for through-going muons in these distributions.⁴ However, to be sensitive to clipping trajectories in which a muon travels through only a portion of a bar, a reduced threshold of 300k pVs is used. This threshold is chosen to be above the tail of the muon shower area distribution, observed near 200k pVs in Figure 6.19, while otherwise being set as low as possible to maximize beam muon acceptance.




Figure 6.18: Area of pulses in front and back panels for a selection of 300k pVs in each of the layers.

⁴Note the different peak locations in the area distribution correspond to fully saturated pulses, which become saturated at $\sim t = 1200$ ns and remain saturated until the end of the window. The peak structure is related to an effect of “early saturation”, in which some channels have different saturation points compared to the expected 1250 mV digitizer readout limit.




Figure 6.19: Area of pulses in each layer for a selection of 300k pVs in each of the other layers and front/back panels. Events are also members of the beam muon skim and must satisfy a maximum pulse area $> 100\text{k pVs}$ in layers 0, 1, and 2.

6.5.2 Beam Muon Flux in Data

Using the results from Section 6.5.1, we define a reconstructed beam muon as an event with at least one pulse satisfying a 300k pVs area threshold in a bar in each layer and a calibrated pulse timing difference between layer 0 and layer 3 satisfying $|\Delta t| < 15\text{ ns}$.

To remove contributions from cosmic muons, we perform a Gaussian fit in the Δt calibrated timing difference using a cosmic-enriched beam-off beam muon skim, as in Figure 6.20.




Figure 6.20: A fit of the layer 3 - layer 0 timing distribution is obtained using the cosmic-rich beam-off beam muon skims. Cosmics in this distribution are normalized to the beam on dataset time, then subtracted from the beam-on beam muon dataset.

This produces a Gaussian fit peaking near -17 ns. The fit is normalized to the integral of the beam-off calibrated timing histogram and scaled to the beam-on dataset time. We then integrate this fit over the interval $t \in (-15, 15)$ ns to estimate the number of cosmic muons satisfying all beam muon selections in the beam-on dataset. This results in an estimated cosmic muon rate of 6.3×10^{-4} Hz. This rate is subtracted run-by-run from the beam muon estimates according to the run duration. After removing these cosmics, the corrected beam muon estimate results in a total of 19955 measured beam muons in 124.7 fb^{-1} , yielding a cross section for reconstructed beam muons of $0.160 \pm 0.010 \text{ pb}$. Figure 6.21 demonstrates that the beam muon flux estimate shows good linearity with both run-level and cumulative luminosity, indicating little contamination.




Figure 6.21: The muon rate as a function of run (left) and cumulative luminosity (right). At a bar area threshold of 300k pVs, beam muon events have high purity while retaining a linear relationship with integrated luminosity.

6.5.3 Beam Muon Flux in MC

Beam muon rates in data are compared to MC as a validation of rate estimates and detector performance. Muons are generated in MC in four modes: Heavy flavor decays, pion/kaon decays in flight, W decays, and Drell-Yan. Generation is performed using FONLL at 13.6 TeV, including an additional production correction of 1.25 ± 0.25 to the heavy flavor decay mode to account for known underproduction of $b\bar{b}$ in FONLL compared to CMS measurements. The P_T distribution and FONLL underestimates are illustrated in Figure 6.4.

Generated beam muons in MC were propagated to the milliQan detector and simulated in the same method used for signal MC; see Section 6.2. The detector response is pulse-injected following the approach of Section 6.4 to achieve data-like waveforms, and the same pulse area and timing requirements as used for data discussed above were applied to reconstruct beam muons in MC.

The total cross section computed for reconstructed beam muons in MC as a result of this process is 0.22 ± 0.06 pb. A distribution of event weights and yields passing beam muon reconstruction in MC is shown in Figure 6.22. The MC uncertainty is dominated by a 25%

systematic uncertainty in the leading $b\bar{b}$ mode, with a sub-leading contribution of 14% over all decay modes due to uncertainty in the thickness and density of the rock between CMS and milliQan during propagation as in Figure 6.4. This simulated result agrees within 1.0σ of the data cross section measurement of 0.16 ± 0.01 pb.




Figure 6.22: MC event weights for beam muons are dominated by heavy flavor production. A constant number of events were generated in each mode; higher event yields for W and DY production illustrate that beam muon reconstruction efficiency is higher for high- P_T muons.




Figure 6.23: The detector is well aligned horizontally (left), but shows a slight vertical misalignment (right), indicated by the excess of events with positive y-slopes compared to negative y-slopes. In both cases, MC slightly overestimates the fraction of straight paths with zero slope, and underestimates the fraction with large slopes.

We use the beam muon skim to measure detector alignment in the horizontal (x) and

vertical (y) planes compared to MC. Figure 6.23 shows a symmetric distribution in x and a mildly asymmetric distribution in y.

This asymmetry is used to derive a correction on MC signal yield and a systematic on detector alignment, as discussed in the analysis Section 7.8.2.

6.6 Cosmic NPE Validation

Another important aspect of detector response modeling is the shape of the muon shower. Secondary electromagnetic showers from cosmic and beam muons, as in Figure 6.24 below, can create low-energy, signal-like deposits in bars separate from the original muon. This can occur even when the original muon misses the detector, as in Figure 6.13.




Figure 6.24: An early version of the “full” milliQan detector design consisted of many bars. Here the transverse passage of a cosmic muon generates gammas in a secondary muon shower as it passes through the detector, causing hits in other bars in the layer.

To study the modeling of this important background, we use a sample of through-going beam muons and study the NPE distribution of “non-neighboring bars”, bars which are not

in or adjacent to the bars containing the muon path. These bars will only be populated with shower products which are spatially separated from the initial muon and high enough energy to penetrate through at least 5 cm of scintillator of the neighboring bar to the muon.

The resulting NPE distributions are shown for the demonstrator and the Run 3 milliQan bar detector in Figure 6.25 [70]. We see good overall agreement for a range of NPE, supporting effective modeling of NPE from shower products.




Figure 6.25: (Left) Beam muon NPE validation using the Run 2 demonstrator detector [70]. (Right) Beam muon NPE validation using Run 3 bar detector. As signal amplification saturates near $NPE = 100$ in Run 3, the bar detector range is much more constrained. In both cases, good agreement is observed for a wide range of NPE.

Another validation of shower modeling can be obtained in studying showers from cosmic muons, a convenient background source as it still exists during beam-off periods. As this background source was relevant for the demonstrator, studies of cosmic shower modeling validations were undertaken using a dedicated version of the demonstrator, rotated such that it was positioned flat along the axis of the tunnel rather than at the usual 43.1° angle [73]. Max/min NPE ratio and a measurement of the number of bars hit in a cosmic shower event are shown in Figure 6.26. We see good agreement, demonstrating that not only the individual energy deposits are well modeled in NPE agreement, but also the shape of the shower.




Figure 6.26: (Left) Max/Min NPE in each cosmic shower event. (Right) the number of bars hit in a cosmic shower. Note that the value extends up to 18: 6 bars per layer and 2 in a fourth layer for the horizontal, four-layer-demonstrator configuration [73]. Good agreement is seen for low bar counts; the disagreement at very high bar counts corresponds to very large, energetic showers and is unimportant for milliQan.

These validations support good overall agreement between sim and data in detector response modeling. However, one area of significant, sustained disagreement between MC and data is in the rate of showering: for a fixed number of muons, we see a significantly higher number of showers in data compared to MC, measured at 3.0 ± 0.6 higher rates in data using cosmic showers. For showers which are generated, we see good physical agreement, so this value is used to modify the overall normalization for MC when comparing to data.

7 The milliQan Experiment: Run 3

Bar Detector Analysis

In this section, I will discuss the data sample, analysis strategy, and results of the Run 3 milliQan bar detector.

7.1 Data and Trigger

The bar detector data samples were collected from June 2023 to 2024 in pp collisions at an energy of $\sqrt{s} = 13.6$ TeV, yielding a total integrated luminosity of $124.7 \pm 3.8 \text{ fb}^{-1}$. Data is collected using all runs with no issues: all channels are active and are correctly mapped to the proper readout channels. These runs represent 87% of the total 142.7 fb^{-1} delivered luminosity during running. The majority of this data was collected in 2024, as in Figure 7.1.




Figure 7.1: MilliQan recorded luminosity as a function of time. Around 8 fb^{-1} was collected in 2023, with the remainder collected in 2024.

The trigger was handled by a custom trigger board. The V1743 CAEN digitizers each produce 8 LVDS groups which provide inputs to the trigger. An LVDS group is a pair of two channels in the digitizer, with each group following a (1,2), (3,4), (5,6),... pairing scheme. Each LVDS group consisted of two horizontally adjacent bars in the same 2×2 bar unit, with two LVDS groups per unit. An LVDS signal digitizes the analog channel input and sets the output to high if either channel is above a preset voltage threshold. Most channels were set to a common voltage of 15 mV, but due to DC fluctuations and noise differences between different channels, some individual thresholds were adjusted low (12 mV) or high (18 mV). Panel trigger thresholds were set to 700 mV to target large energy deposits made by background muons or showers.

Triggers are generated by requiring coincident signals between sets of LVDS inputs. The coincidence window was set using the trigger board's 125 MHz clock to 160 ns in 2023 and reduced to 80 ns in 2024. However, a 70 ns delay is introduced to each window due to internal signal processing delay in the trigger board, setting the effective coincidence windows to 230 ns in 2023 and 150 ns in 2024. Additionally, whenever an LVDS signal is set to high, it remains high for an additional 120 ns. These trigger performance parameters have been validated using online-offline measurements of trigger efficiency.

After a trigger is recorded, a 125 μ s deadtime is applied using a mixture of an explicit ≈ 1 μ s deadtime in the trigger board and a 125 μ s deadtime in the digitizers. This ensures that no correlated, delayed activity such as afterpulsing (Section 5.2) or DC offsets in the PMTs due to very large energy deposits.

The full trigger menu is listed in Table 7.1. Two primary signal triggers are in the menu: 3-in-a-row, requiring three hits in different layers which are in the same layer position; and 4 layers hit, requiring at least one hit in each layer in any position, targeting through-going, straight mCPs. In addition to the primary signal triggers, the menu contains a backup trigger (4-in-a-row), a zero bias trigger, a dedicated beam muon trigger (front and back panel), and a variety of background and validation triggers (Panel triggers, N hits, Separate/Adjacent

Trigger	Prescale	Average Rate (Hz)
4 Layers Hit (1)	1.0	0.13
3 Hits in Line (2)	1.0	0.75
Separate Layers (3)	2.5×10^{-4}	0.25
Adjacent Layers (4)	2.0×10^{-4}	0.26
N Layers Hit (5)	0.02	0.42
External (6)	1.0	0.0
N Hits (7)	2.0×10^{-4}	0.34
Internal (8)	1.0	0.0
Top Panels (9)	2.0×10^{-4}	0.13
Top Panels + Bottom Bar (10)	0.01	0.35
Front & Back Panel (11)	1.0	0.03
— (12)	1.0	0.0
Zero Bias (13)	1.0	0.1 (0.5)
2 Hits in Line (14)	1.0×10^{-5}	0.003
4 Hits in Line (15)	1.0	0.009

Table 7.1: The bar detector trigger menu with prescales and measured rates listed. The zero bias trigger had an explicit rate set originally at 0.5 Hz and decreased to 0.1 Hz in March 2024. Total rate budget is around 5 Hz due to data readout limitations.

layers).

7.2 Data Processing

Data processing is handled by “milliQan Offline”, a set of software tools which performs pulse reconstruction, calibration, and event classification.

7.2.1 Pulse Reconstruction

Recorded data for triggered events consist of waveforms for each readout channel. The beginning of a pulse is identified by recording three samples above 10 mV without any sample recorded under 1 mV. Similarly, the end of the pulse is identified by recording three consecutive samples below 10 mV. Note that the first three samples above 10 mV need not be consecutive. The start of the pulse is recorded as the lower bin edge of the last sample below 10 mV before the start, and the end is recorded as the lower bin edge of the first sample after the pulse has ended (i.e. the lower bin edge of the first sample after the three consecutive

Pickup Cut Selections	Pickup	Pickup Tight
Fall samples < 2	or	or
Rise samples < 2	or	or
$17 < \text{height} < 100$ and $\frac{\text{area}}{1000} < 0.04 \times (\text{height} - 17)$	or	or
$\text{height} < 25$ and $\text{duration} < 4.6 \times (\text{height} - 12)$	or	or
$\text{height} \geq 25$ and $\text{duration} < 60$ and $\text{fallSamples} < 6$	or	or
$\frac{\text{area}}{1000} < 0.4$ and $\text{duration} < \frac{\text{area} \times 150}{1000}$	or	or
$\frac{\text{area}}{1000} \geq 0.4$ and $\text{duration} < 60$	or	or
Rise samples $\geq$ fall samples	-	and
$70 < \text{pulse start time} < 2400$	-	and
< 3 Pulses in channel	-	and
Duration < 100	-	and

Table 7.2: Two flags are used to identify pickup. If any pickup “or” selection is satisfied the flag is set to true. In addition, the pickupTight flag requires a more stringent, targeted set of requirements.

low samples).

When a pulse has been identified, several pulse characteristics are recorded. A full list of all branches in the processed “offline trees” is given in Appendix A.11. A few key pulse characteristics are height (pulse amplitude), duration (time between start and end of pulse), area, and NPE.

Pulse area is measured after the start and end of a pulse is identified by subtracting the sideband mean from the each pulse sample, then integrating over the duration of the pulse. The sideband mean and RMS are recorded over the first 50 samples (125 ns) of an event. Pulse area is converted to pulse NPE using a per-channel calibration detailed in Section 7.3.

One additional pulse requirement is a flag rejecting pickup. Pickup is a purely electronic signal induced in the readout due to activity in nearby channels, ground or power fluctuations, or any other source of unwanted cross-talk between electrical components. The consecutive samples requirements at the beginning and end of a pulse are targeted at rejecting pickup and other sources of fluctuation, but further pickup rejection is achieved using an explicit pickup flag with selections summarized in Table 7.2.

7.2.2 Datasets

Triggered data is processed into a set of skims: Signal, beam muon, cosmic muon, and zero bias data. The requirements in each data skim selection are below in Tables 7.3– 7.5. Data skim selections are loose; later offline requirements are used to further identify signal, beam muon, and cosmic muon events. A zero bias skim is produced using the explicitly fired zero bias trigger at the specified trigger rate from Table 7.1.

Beam Muon Skim Selections	
Scope	Selection
Event-level	≥ 3 layers hit
Pulse-level	pulse area $\geq 1 \times 10^5$

Table 7.3: A loose beam muon skim on a large hit in each of 3 layers is possible as beam muon rates are low.

Cosmic Muon Skim Selections	
Scope	Selection
Pulse-level	pickup cut (tight)
Pulse-level	first pulse in channel
Pulse-level	$900 \leq t \leq 1500$
Pulse-level	height ≥ 15
Pulse-level	area $\geq 1 \times 10^5$
Event-level	≥ 3 pulses in same column
Event-level	top/side panel hit in same layer as column hit

Table 7.4: Cosmic muons are identified as three large pulses in the same column of a single layer, with panel hits over the associated layer. This identifies either direct activity from a cosmic muon or a related shower.

Signal Skim Selections	
Scope	Selection
Event-level	Veto: top/side panel area $> 5 \times 10^5$
Pulse-level	first pulse in channel
Pulse-level	$900 \leq t \leq 1500$
Pulse-level	height ≥ 15
Event-level	≥ 3 pulses in straight line

Table 7.5: Signal selections require three, centrally timed pulses in a straight line with a veto on panel activity. A first-pulse requirement is a veto against afterpulsing and ringing.

7.3 Calibration

7.3.1 NPE Calibration

Pulse areas can be converted to number of photoelectrons (NPE) by measuring the single photoelectron (SPE) peak during normal operation. Most signal analysis uses NPE rather than pulse area as the defining variable to represent energy deposits in a bar. The exception is in beam and cosmic muon analysis, where saturating (pulse height > 1200 mV) deposits cannot be easily converted to NPE as the number of photoelectrons is no longer a linear function of pulse area for a saturated pulse.




Figure 7.2: An SPE peak fit using a Gaussian + Landau.

To extract the SPE peak, a combined Gaussian + Landau fit is used. The Gaussian component reflects the area distribution of standard single photon events from PMT dark rate events, while the Landau component reflects a mixture of noise and sub-SPE dark rate as discussed in Section 5.2. An example SPE fit is shown in Figure 7.2. To minimize noise

SPE Calibration Selections
rise samples ≥ 3
fall samples ≥ 3
first pulse
pickup cut tight

Table 7.6: Selections applied to data to remove noise from a distribution of sPE pulses.

contributions, selections are applied from Table 7.6 before applying SPE peak fits.

Each channel has an SPE fit extracted using this method. The distribution of SPE area peaks as a function of channel is shown below in Figure 7.3.




Figure 7.3: SPE calibrations are peaked around 3000 pVs; this approximately corresponds to a single photon pulse amplitude of 30 mV.

7.3.2 Energy Calibration

A Cd109 calibration is used in each channel to extract an energy calibration, important for MC estimates of signal efficiency. The Cd109 calibration is used primarily to convert pulse area into an equivalent NPE. In addition, it provides a measure of the energy sensitivity of the detector.

Cd109 decays via electron capture of an inner shell electron to Ag109. This electron

capture produces an excited nuclear state. In a small number of decays, this nucleus de-excites through direct emission of an 88 keV gamma. In most decays the de-excitation proceeds through internal conversion, ejecting a bound electron. This conversion results in up to two vacancies in the K- or L-shells— one from the electron capture decay and the other from internal conversion.¹ These vacancies are filled through $K\alpha$, $K\beta$, or L-transitions as electrons from outer shells fill the vacancy and emit a characteristic x-ray. The emission spectrum of Cd109 is listed in Table 7.7.

Process	Energy (keV)	Photons Per Disintegration
Ag L	2.63–3.75	0.104
Ag $K\alpha_2$	21.99	0.292
Ag $K\alpha_1$	22.16	0.551
Ag $K\beta_1$	24.91–25.15	0.153
Ag $K\beta_2$	25.46–25.51	0.027
γ (Ag)	88.03	0.0366

Table 7.7: Combined Cd109 X-ray and γ emissions from electron capture decay to Ag109. Probabilities are photons per 100 disintegrations. Full data at [111].

The dominant x-ray emission (84%) is a $K\alpha$ emission centered at 22.1 keV, with a secondary emission near 25 keV. Low-energy Compton scatters of the 88 keV gamma and the L x-ray emission lead to a measurable enhancement in low-energy SPE deposits during calibration.

In Cd109 calibration of a bar, the source is placed at the center of the bar’s 60 cm length. The 22 keV $K\alpha$ x-ray results in a characteristic energy deposit into the scintillator, appearing as a Gaussian distribution in reconstructed pulse areas. A fit is made to this Cd109-induced pulse area distribution. This fit is compared to the SPE pulse area distribution for the bar, yielding an energy calibration. A distribution of the NPE location of the Cd109 calibration peak is given in Figure 7.4.

¹Note that not all decays produce two vacancies, as holes may be produced in outer shells, or shell vacancies may be filled with the emission of an Auger electron carrying the transition energy rather than an x-ray.




Figure 7.4: NPE values are produced by dividing the Cd109 source peak area value by the area of the single photon peak. This produces a value which represents the average number of photons detected by the PMT for an energy deposit of 22 keV.

7.3.3 Timing Calibration

An important component of background reduction is a narrow timing window. Since backgrounds which are uniform scale as $R_{bkg} \propto T^3$, where T is the width of the coincidence window in a four-channel coincidence (see Section 7.5.1), an effective timing calibration is important to make this window as narrow as possible.

The objective of timing calibration is to adjust the center of each channel timing peak to zero, such that particles propagating at $v = c$ from the CMS IP will be measured at $t = 0$ in each channel.

Channel times can vary depending on cable length and different intrinsic electronic processing times throughout the readout chain. To measure these timing calibrations, beam muon events are used. Beam muon events are selected from the beam muon skim using

Timing Calibration Selections	Event/Pulse Based
Digitizer boards synchronized	Event
Pickup cut tight	Pulse
First pulse in channel	Pulse
nPE > 200	Pulse
Pulse in trigger window ($1100 \leq t \leq 1400$)	Pulse
Hits per layer ≥ 1	Event
Front panel hit	Event
4 hits in straight line	Event

Table 7.8: Beam muon selections used for timing calibration. Selections require a clean beam muon with a straight line path through the detector, a hit in the front panel, and central timing values.

selections given in Table 7.8.

The hit times in each bar layer are calibrated to the hit time in the front panel. Gaussian fits are made using this timing difference, and the resulting peak value is subtracted from the channel time. A set of sample timing differences fit for four channels is shown in Figure 7.5.




Figure 7.5: Timing distributions for four sample channels are shown. These raw timing differences are used to derive the timing calibration, with the peak values for each Gaussian fit subtracted from the raw channel time.

Timing calibrations are also required for MC. The modeled spacing between layers is

86 cm, implying a speed of light timing difference between consecutive layers of 2.87 ns; with no electronic modeling, this is the only relevant timing difference requiring calibration. This is confirmed through explicit beam muon simulations, with ΔT between Layer 0 and Layers 1,2,3 shown in Figure 7.5. Timing resolution in MC is a pure function of the number of photons and the path of the throughgoing MIP, so very sharp peaks are observed for beam muons of $q/e = 1$ which deposit $O(10^6)$ photons.




Figure 7.6: MC timing calibrations. Sharp peaks are observed at the time of flight timing difference of 2.87 ns.

7.3.4 Timing Validation

Timing is validated using the cosmic and beam muon skim. Time differences between layers in the beam muon skim are calibrated through the process in Section 7.3.3 to peak at $t = 0$. Meanwhile, cosmic muons which travel from the opposite direction through the detector peak at approximately $2\Delta T(L_0, L_3)$ after purely electronic timing differences are calibrated

between layer 1 and layer 4.²

In Figure 7.7, the calibrated timing distribution $\Delta T_{calib}(L_0, L_3)$ is shown for beam and cosmic muons. A small population of cosmic muons appears in the beam muon sample, so the beam-off timing distribution of cosmic muons can be compared to the beam-on timing distribution of beam and cosmic muons. Selections in the beam muon timing plot are stricter, requiring hits in both endcap panels satisfying pulse area $> 10^5$ pVs. Otherwise, both plots require pulse area $> 3 \times 10^5$ pVs in each layer. Good separation between each distribution is observed, with a centering of the beam muon distribution near $\Delta T_{calib} = 0$.




Figure 7.7: (Left) The cosmic muon timing distribution measured between the first and last layers. A 2-Gaussian fit is performed for consistency. (Right) The beam muon timing distribution. Under these tight selections with activity required in four layers and both endcaps, throughgoing cosmic muons are present but highly suppressed.

7.3.5 Timewalk

In addition to the standard timing calibrations performed using beam muons, another timing calibration is applied to remove the effect of “timewalk”: a systematic timing shift in the measured pulse start time as a function of the pulse size.

For unsaturated pulses, timewalk can be corrected for by fitting to the halfmax region between 20% and 80% of the maximum pulse amplitude. A toy model shows good agreement

²This is because the calibration procedure applies a factor of $\Delta T(L_0, L_3)$ to account for the propagation time of particle incident from the CMS IP, while cosmic muons have an additional propagation delay time of $\Delta T(L_0, L_3)$ since they have to physically propagate from the fourth layer to the first.

for the pulse start times using this approach in Figure 7.8. However, for large, saturating pulses, the pulse time is recorded earlier and becomes a function of the total pulse area. The value of this fitted rising edge time is called `timeFit`.




Figure 7.8: (Left) A toy model of pulse Landaus are modeled with different amplitudes. (Right) Using a fit to the central values of the pulse rising edge, good agreement is seen for unsaturated pulses. However, saturated pulses are recorded earlier, producing a dependence for pulse timing on amplitude.




Figure 7.9: (Left) Timewalk as measured using a second, reference pulse in the layer as a function of pulse area. (Right) A profile of the timewalk effect. This profile is used to derive a timewalk correction.

In addition to this standard timewalk behavior, there seems to be a second timewalk effect. Plotting the pulse time against a second, reference time in the same layer as in Figure 7.9, we

observe a effect on the scale of $O(10 \text{ ns})$, significant relative to the overall timing resolution of the experiment.

To correct for this second timewalk effect, we use an explicit model of the timewalk as measured in Figure 7.9. That model is chosen as:

$$\Delta t = \sigma(s[x - x_p]) \cdot [e^{a_2(b_2 - x)} + c_2 + d_2 x] + \sigma(s[x_p - x]) \cdot [c_1 + d_1 x] \quad (7.1)$$

for $\text{NPE} > 4$, and

$$\Delta t = \exp(a_1[b_1 - x]) \quad (7.2)$$

for $\text{NPE} < 4$. Here, x is the pulse area, x_p is the transition region between the flat “knee” and the decaying exponential, σ is the sigmoid function, and a_i, b_i, c_i, d_i are fit parameters. The model is chosen to fit to the distinct regions at low, intermediate, and high values of NPE. The result of the model fit is shown in Figure 7.10 below, with good agreement to the sample timewalk data. This fit is repeated and applied individually for each channel.




Figure 7.10: The full timewalk fit shows good agreement to the data profile.

7.4 Search Design

The key signature for a millicharged particle in the bar detector is a straight trajectory with energy deposits in all four layers. The trajectory should be straight as the low charge and ionization of mCPs means magnetic field deflection will be minimal compared to its momentum. The deposits should be in all four layers in order to reject backgrounds. A sample simulated event display for a $q/e = 0.003$ mCP is shown below.




Figure 7.11: An mCP passes through all four layers of the detector, depositing energy which is converted to scintillation light in each layer.

For very low charge such as $q/e = 0.003$, we expect no photons to be detected in the front and back slabs. However, for higher charge such as for $q/e = 0.1$ below, we expect to consistently measure light as it passes through an endcap slab.




Figure 7.12: At higher mCP charges, endcap slabs become sensitive detector volumes with consistent mCP deposits. The size of the deposit in the endcap slab can be used to distinguish between a high-charge mCP and a through-going beam muon.

To distinguish between these two classes of mCP, two signal regions are designed:

- **SR1**, in which an event-level veto is placed on hits recorded in either endcap panel
- **SR2**, in which an event-level requirement is placed on hits in at least one endcap panel

7.5 Search Backgrounds

The primary backgrounds for the milliQan experiment are dark rate pulses from the PMTs, showers from cosmic muons, and beam muons.

7.5.1 Dark rate overlap background

PMT dark rates are generated due to thermionic emission in the photocathode. For R878 PMTs in the bar detector, this generates a single photon-like signal at a rate of approximately

1 kHz.

To compute the dark rate background contribution for the bar detector, we model this dark rate as a Poisson process for each of the 16 straight-line paths in the detector.

First, suppose one of the four bars triggers a dark rate single photoelectron (SPE) pulse, occurring at rate R_1 . The probability that the remaining bars independently produce a dark pulse in the trigger window of length T is given by:

$$\text{Poisson}(k > 0, R_2T) \times \text{Poisson}(k > 0, R_3T) \times \text{Poisson}(k > 0, R_4T) \quad (7.3)$$

Since $(R_N T) \approx 1.5 \times 10^{-5} \ll 1$, we can expand each Poisson term as:

$$\text{Poisson}(k > 0, R_N T) = 1 - \text{Poisson}(k = 0, R_N T) = 1 - e^{-(R_N T)} \approx R_N T \quad (7.4)$$

This means the final rate, assuming PMT 1 initiates the first dark pulse, is:

$$R_1 \times (R_2 T) \times (R_3 T) \times (R_4 T) = R_1 R_2 R_3 R_4 T^3 \quad (7.5)$$

Since this process occurs for each of the four PMTs, the total four-fold coincidence dark rate per path is therefore given by:

$$4R_1 R_2 R_3 R_4 T^3 \quad (7.6)$$

For typical dark rates of $R_N = 1$ kHz and a coincident timing window of $T = 20$ ns, the expected total dark rate is 3.2×10^{-11} Hz. For an integrated livetime of approximately 6 months, this yields a total expected dark rate for 2023+2024 of 0.001 events.

7.5.2 Projections Paper: Shower Backgrounds

In the Run 3 analysis, backgrounds are data-driven. However, it's worth noting explicitly that shower backgrounds are rejected by the design of a four-layer detector, as the width and

Selection	Run 3
≥ 1 per layer	8.1×10^5
$= 1$ per layer	6.0×10^3
Cosmic panel veto	1.1×10^3
Beam muon panel veto	780
Four in line	0.19
NPE Max/Min < 10	0.061
$-15 \text{ ns} < \Delta t_{\text{max}} < 15 \text{ ns}$	0.012

Table 7.9: Initial predictions for the Run 3 bar detector due to cosmic showers. More details on the study in the Run 3 milliQan projections paper [73].

segmentation of the shower required to satisfy requirements of exactly one hit per layer in a pointing path produce substantial background attenuation. Results of a study of cosmic background estimates from the original Run 3 bar detector projections paper are given in Table 7.9. More details of this study are given in the original projections paper [73]. Note that showers emerging from the cavern walls due to beam muons are less likely than those from cosmic muons, as cosmic muons are much more common (10^9 per year in the vicinity of the detector) than beam muons (10^5 per year, depending on luminosity).

Though cosmic shower backgrounds sources are expected to be negligible for the bar detector, this does not eliminate the possibility of in-time backgrounds from beam muons, particularly in search regions of high charge.

7.6 Selections and Validation

We start by applying event- and pulse-level quality criteria to all signal regions as listed in Table 7.10.

Cut	Selections	Logic	Event/Pulse
Pickup	Rise Samples ≥ 3	OR	Pulse
	Fall Samples ≥ 5	OR	Pulse
	Height ≥ 450 mV	OR	Pulse
Noise	Pre-Pulse Mean ≤ 5	AND	Pulse
	Pre-Pulse RMS ≤ 5	AND	Pulse
	Sideband Mean ≤ 5	AND	Pulse
	Sideband RMS ≤ 5	AND	Pulse
Sub-SPE Dark Rate	Duration ≥ 50	AND	Pulse
	nPE ≥ 0.3	AND	Pulse
Digitizers Synchronized	Digitizers outputs should be well-matched	—	Event
First Pulse in Channel	Keep only the earliest pulse per channel	—	Pulse
Pulse Time in Trigger Window	$1100 \leq t \leq 1400$	—	Pulse
Top/Side Panel Veto	No coincident hits in top or side panels	—	Event
Hits Per Layer ≥ 1	At least one hit in each layer	—	Event

Table 7.10: Quality cuts applied to all signal regions. Pickup, noise, and sub-SPE dark rate are applied per pulse with the indicated logic.

These quality cuts required a centrally timed pulse in each layer with a veto on top/side panel activity to reject showers. We require the first pulse in the window to reject ringing and afterpulsing, and we additionally apply targeted selections to veto pickup, noise, and sub-SPE dark rate.³

7.6.1 SR1 Selections

As discussed in Section 7.4, in SR1 we target low-charge mCPs with an event veto on activity in the front and back endcap panel. A summary of the SR1 signal selections is given in Table 7.11.

³Note the the requirements for this pickup requirement are different than discussed in Table 7.2, with more stringent requirements on the number of rise and fall samples.

Signal Region 1 Selections	Event/Pulse Based
4 Bars Hit	Event
$E_{max}/E_{min} < 10(5)$	Event
Energy Deposit $\leq \text{Min}(1000 \text{ keV}, E_{\mu} - 50 \text{ keV})$	Pulse
Front/Back Panel Hits = 0	Event
Straight Line Path	Event
$t_{max} - t_{min} < 20 \text{ ns}$	Event

Table 7.11: SR1 selections, targeting low charge ($q/e < 10^{-2}$). Selections are applied sequentially following the quality selections defined in Table 7.10. At energy deposits of $> 50 \text{ keV}$, the max/min energy requirement is reduced to 5 from 10.

The SR1 signal selections are justified by the following:

- Exactly 4 Bars Hit: This requirement significantly reduces backgrounds from dark rates and cosmic showers, requiring four physically distinct measured energy deposits.
- Front/Back Panel Veto: For low charges with $\langle NPE \rangle = 1$ in the 60 cm bars, the 5 cm thickness and smaller light collection efficiency of the endcap panels ensures that this veto requirement is highly efficient for signal while rejecting beam-induced backgrounds from beam muons.
- Energy Deposit $\leq \text{Min}(1000 \text{ keV}, E_{beam.Muon} - 50 \text{ keV})$: Using a rescaling of NPE under our Cd109 calibration, exploiting the fact that the calibration is generated by a 22 keV characteristic $K\alpha$, we can express the NPE as an energy deposit. This allows us to set a physical background requirement which is independent of the optical response in a given channel. One limitation of this procedure is that the upper bound of an energy deposit, corresponding to a pulse which saturates at the center of the window and remains saturating until the end of the window, is less than 1000 keV for channels. This is possible for channels with small values of $(\text{keV}/\text{NPE}) \times 1/(\text{area}/\text{NPE})$, given by the Cd109 energy scaling divided by the value of the SPE area peak. To account for this, a maximum energy requirement is set as the lesser of 1000 keV or 50 keV below the characteristic energy peak associated with a fully saturating, beam-muon-like deposit in the bar. The distribution of these beam-muon-like energies as a function of channel is shown in Figure 7.13.




Figure 7.13: Measured energy as a function of channel for through-going muon events. For channels with $E_{beam,Muon} < 1000$ keV, $E_{beam,Muon}$ is used instead to set an energy maximum in SR1.

- $Energy_{max}/Energy_{min} < 10(5)$: As signal energy deposits are from mCP ionization, they should be approximately consistent in their measured energy deposit. In comparison, physically induced backgrounds from showers or radiation will be approximately flat in energy deposit. Therefore, a max/min requirement targets backgrounds due to ionization. At higher mCP charges corresponding to energies greater than 50 keV, this requirement is tightened from $10 \rightarrow 5$ as the max/min energy distribution for signal tightens when not subject to low-NPE statistical fluctuations.
- Straight Line Path: Requirement of a straight line path significantly reduces backgrounds as dark rate and shower background sources are uncorrelated in position within a layer. Requiring a straight line path is not only consistent with a mCP incident from the CMS IP, but it induces a factor 16^3 reduction in these uniformly distributed backgrounds.
- $|t_{max} - t_{min}| < 20$ ns: Requiring a tight timing window ensures signal activity is consistent with a particle trajectory in-time with activity from the IP. This is most important for the $R \propto T^3$ dependence on dark rate backgrounds.

7.6.2 SR1 MC and Validation

A selection of N-1 plots for SR1 are given below. For the other selections, signal efficiency is $\sim 100\%$. In Figure 7.14, we see energy deposit is an increasing function of charge, but Max/Min is a decreasing function of charge; this justifies the adjustment to the Max/Min requirement at higher energies.




Figure 7.14: (Left) Energy deposited and (Right) Max/Min energy ratio in SR1 with all other cuts applied.




Figure 7.15: (Left) Bar drift and (Right) $\Delta t = |t_{\max} - t_{\min}|$ in SR1 with all other cuts applied.

Meanwhile, in Figure 7.16, we see two interesting effects. First, the signal angular distribution is primarily straight, but with lighter models showing a broader angular spread.

And second, both the mean and the spread of the signal Δt decreases as a function of charge, an effect consistent with timing resolution improving with energy deposit.

Signal efficiency in the target charge range $q/e < 0.02$ is $> 90\%$ for all SR1 selections. A summary of the N-1 SR1 signal yields, with all but the indicated selection applied, is shown in Table 7.12.

N-1 Cut	Events Passing			
	m=0.1, q=0.004	m=0.5, q=0.007	m=1.0, q=0.008	m=1.7, q=0.02
Total Events	368.18	88.69	69.61	12.82
All Cuts	8.68 (2.36)	14.52 (16.37)	14.13 (20.3)	1.47 (11.48)
Top/Side Panel Veto	8.68 (2.36)	14.52 (16.37)	14.13 (20.3)	1.47 (11.5)
≤ 4 Bars Hit	8.68 (2.36)	14.52 (16.37)	14.13 (20.3)	1.48 (11.51)
Front/Back Panel Veto	8.79 (2.39)	15.90 (17.92)	16.27 (23.38)	3.24 (25.27)
Energy < 1000 keV	8.68 (2.36)	14.52 (16.37)	14.13 (20.3)	1.47 (11.48)
Energy Max/Min $< 10(5)$	10.02 (2.72)	17.06 (19.23)	16.51 (23.71)	1.54 (12.02)
Straight Line Path	9.44 (2.56)	14.68 (16.55)	14.15 (20.33)	1.47 (11.48)
$\Delta T(\text{max-min}) < 20$ ns	9.22 (2.5)	14.78 (16.66)	14.27 (20.5)	1.47 (11.5)

Table 7.12: N-1 Event yields in SR1. In each line, the indicated selection has been removed and all other selections applied.




Figure 7.16: (Left) NPE and (right) Δt for beam-off events for pointing and non-pointing path requirements. The distributions agree well, suggesting independence.

To estimate backgrounds, we use the data-driven ABCD method in the variables Δt and a “pointing path” requirement, enforcing that a signal-like hit is positioned in the same row and column in each layer. Before estimating backgrounds, two validations were performed in the beam-off dataset to assess the independence of energy deposit and Δt on whether the path was pointing or non-pointing. Results are shown in Figure 7.16; the pointing and

non-pointing distributions appear consistent in each variable, suggesting no dependence of NPE or Δt on apparent signal trajectory.

To test closure of the ABCD method for background estimation, a sideband in beam-off data is used. The sideband is defined by a modification to the pointing path requirement such that the total deviation of a path $\Delta N = \Delta N_{Row} + \Delta N_{Col}$ is given by $\Delta N = 1$, such that a single hit in one layer is required to deviate by a single row or column. The results are shown in Figure 7.17, with an ABCD estimate of $0.31^{+0.28}_{-0.18}$. With a measured value in the signal region $A = 1$, this suggests good closure.




Figure 7.17: A sideband ABCD in $\Delta N = 1$ for almost-pointing paths in SR1. Agreement between the prediction of $0.31^{+0.28}_{-0.18}$ and observation of 1 shows good closure.

7.6.3 SR2 Selections

Complementary to SR1, SR2 targets higher charge generally satisfying $q/e > 10^{-2}$. As with SR1, quality cuts from Table 7.10 are first applied, followed by the selections in Table 7.13.

Signal Region 2 Selections	Event/Pulse Based
4 Bars Hit	Event
$E_{max}/E_{min} < 10$	Event
$\text{Max}(\text{Front/Back Panel nPE}) < 70$	Event
$\text{Front/Back Panel Hits} \geq 1$	Event
Straight Line Path	Event
$t_{max} - t_{min} < 20$ ns	Event

Table 7.13: SR2 selections, targeting high charge ($q/e > 10^{-2}$). Selections are applied sequentially following the quality selections defined in Table 7.10.




Figure 7.18: Panel Hits for signal models with all other SR2 selections applied. At $q/e = 0.03$, the average number of panels hit becomes greater than 1.

The motivation for most selections follows from the discussion in Section 7.6.1. The two selections in an endcap panel requirement and max/min NPE endcap panel requirement are justified as follows:

- **Front/Back Panel Hits** While panels can be vetoed on low charge signal in SR1, for sufficiently high charge a measured deposit in the endcap panels becomes likely. As illustrated in Figure 7.18, for a charge of approximately $q/e \sim 0.03$, it becomes more likely to observe one or more endcap panel deposits than it is to observe zero. This complementary requirement allows for best sensitivity for both low and high charges.
- **Max(Front/Back Panel nPE) < 70**: For large charges, particularly at the upper end of the charge sensitivity of SR2 $q/e \sim 0.1$, the bars saturate and produce a poor measurement of the actual energy deposit. As the endcap panels saturate at much higher values of charge, they can be used to distinguish between mCPs of high charge and throughgoing muons. A max NPE requirement of 70 provides high efficiency for signal up to $q/e < 0.2$ and low efficiency for higher charge, as in Figure 7.19.




Figure 7.19: Panel NPE for signal models with all other SR2 selections applied. A maximum requirement of $NPE < 70$ is efficient for $q/e = 0.2$ but becomes quickly inefficient for higher charges.

7.6.4 SR2 MC and Validation

As before, a selection of N-1 plots for SR2 are given below; selections not depicted are highly efficient. In Figure 7.20, while the bar drift suggests that the angular distribution of the mCP continues to tighten at high mass, the N_{hits} plot suggest more bars are being hit at heavy, high charge model points. This demonstrates that as charge approaches $q/e = 1$, mCPs begin to form muon-like EM showers as the charge coupling approaches the SM coupling.




Figure 7.20: (Left) Bar drift and (Right) the number of distinct bars hit per event in SR2 with all other cuts applied.

Meanwhile, in Figure 7.21, most model points see a tightening in E_{Max}/E_{Min} and Δt with increased charge as we observed in SR1. However, for very high masses such as the model point at $m = 10$ GeV, $q/e = 0.2$, the mCP can be produced at very low-momentum, measurably increasing its time of flight and shifting the distribution to peak at higher values of Δt . The timing requirement $\Delta t < 20$ ns is still highly efficient for this model point, however.




Figure 7.21: (Left) Max/min energy ratio and (Right) $\Delta t = |t_{max} - t_{min}|$ in SR2 with all other cuts applied.

Signal efficiency in the target charge range $q/e > 0.01$ is $> 90\%$ for all SR2 selections except the requirement of an endcap panel hit, which is a function of charge. A summary of the N-1 SR2 signal yields, with all but the indicated selection applied, is shown in Table 7.14.

N-1 Cut	Events Passing			
	m=1.3, q=0.009	m=1.7, q=0.03	m=6.0, q=0.1	m=10.0, q=0.2
Total Events	71.95	30.68	15.29	42.32
All Cuts	3.18 (4.42)	5.77 (18.8)	2.99 (19.55)	6.99 (16.53)
Top/Side Panel Veto	3.18 (4.42)	5.77 (18.8)	3.00 (19.61)	7.08 (16.73)
Front/Back Panel Required	17.26 (23.99)	7.47 (24.33)	3.22 (21.09)	7.44 (17.58)
Energy Max/Min < 10(5)	3.80 (5.28)	5.82 (18.98)	2.99 (19.58)	6.99 (16.53)
Straight Line Path	3.23 (4.49)	5.79 (18.86)	3.00 (19.61)	7.02 (16.58)
$\Delta T(\text{max-min}) < 20$ ns	3.30 (4.58)	5.80 (18.92)	3.04 (19.87)	7.59 (17.94)
≤ 4 Bars Hit	3.18 (4.42)	5.78 (18.82)	3.05 (19.93)	7.84 (18.53)
Max(Front/Back Panel nPE) < 70	3.18 (4.42)	5.77 (18.8)	2.99 (19.58)	7.30 (17.25)

Table 7.14: N-1 Event yields in SR2. In each line, the indicated selection has been removed and all other selections applied.

A closure test is also performed on SR2, using the same almost-pointing requirement ΔN as was used for SR1. The results are shown in Figure 7.22, with an ABCD estimate of $3.40^{+1.69}_{-1.20}$. With a measured value in the signal region $A = 5$, this suggests good closure.




Figure 7.22: A sideband ABCD in $\Delta N = 1$ for almost-pointing paths in SR2. Agreement between the prediction of $3.40^{+1.69}_{-1.20}$ and observation of 5 shows good closure.

7.7 Results

7.7.1 SR1

With good closure and the key ABCD variables showing independence in validations, we proceed to measure the signal and background results of our beam-off and beam-on datasets in SR1. The results are shown in Figures 7.23 and 7.24. The results are also summarized in Table 7.15.




Figure 7.23: ABCD background estimates and observations for beam-off data. A prediction of $0.32^{+0.24}_{-0.16}$ agrees well with an observation of 0.




Figure 7.24: ABCD background estimates and observations for beam-on data. A prediction of $0.10^{+0.12}_{-0.07}$ agrees well with an observation of 0.

LHC-Beam	Time (s)	Lumi (fb^{-1})	A	B	C	D	Est
Beam-On	12214771	124.7	0	5	152	3	$0.10^{+0.12}_{-0.07}$
Beam-Off	16096891	0	0	8	175	7	$0.32^{+0.24}_{-0.16}$

Table 7.15: SR1 ABCD background estimates and signal observations for beam off and beam on data. 0 events are observed in both backgrounds in good agreement with the background prediction.

In beam-on data measured in 124.7 fb^{-1} , 0 events are observed.⁴ This is consistent with the ABCD background prediction of $0.10^{+0.12}_{-0.07}$. SR1 cut efficiencies are shown in Figure 7.25, and the cutflow yields are given in Table 7.16.




Figure 7.25: SR1 cumulative (top) and individual (bottom) cut efficiencies for beam-off and beam-on data.

⁴In the first unblinding of the analysis, 1 event was observed. But this event was trivially identified as a beam muon, as one of the channels was maximally saturated to the end of the readout window. It turned out this channel was incapable of failing the $E_{dep} < 1000 \text{ keV}$ requirement, which required a rework of this requirement to be consistent with Figure 7.13. After that selection was updated, the results are as shown, with zero observed signal events.

Selection	Beam-Off	Beam-On	m=0.1 q=0.004	m=1.0 q=0.008
Total Events	1002647.0 (100.0)	790776.0 (100.0)	324.0 (88.0)	61.26 (88.0)
Digitizers Synchronized	1002617.0 (100.0)	790772.0 (100.0)	324.0 (88.0)	61.26 (88.0)
Pickup	1002617.0 (100.0)	790772.0 (100.0)	324.0 (88.0)	61.26 (88.0)
Dark Rate	1002612.0 (100.0)	790772.0 (100.0)	324.0 (88.0)	61.26 (88.0)
≤ 6 Bars	669318.0 (66.76)	506770.0 (64.09)	324.0 (88.0)	61.26 (88.0)
First Pulse	669318.0 (66.76)	506770.0 (64.09)	324.0 (88.0)	61.26 (88.0)
Trigger Window	668781.0 (66.7)	506417.0 (64.04)	323.89 (87.97)	61.26 (88.0)
Top/Side Panel Veto	377523.0 (37.65)	287811.0 (36.4)	266.91 (72.49)	38.91 (55.9)
4 Layers	2360.0 (0.24)	3369.0 (0.43)	11.97 (3.25)	19.34 (27.78)
≤ 4 Bars	921.0 (0.09)	985.0 (0.12)	11.65 (3.16)	19.34 (27.78)
Noise	921.0 (0.09)	985.0 (0.12)	11.65 (3.16)	19.34 (27.78)
Energy Max/Min $\leq 10(5)$	262.0 (0.03)	336.0 (0.04)	10.3 (2.8)	16.52 (23.73)
Energy ≤ 1000 keV	262.0 (0.03)	331.0 (0.04)	10.3 (2.8)	16.52 (23.73)
Front/Back Panel Veto	258.0 (0.03)	209.0 (0.03)	10.09 (2.74)	14.32 (20.57)
Straight Line	7.0 (0.0)	3.0 (0.0)	9.22 (2.5)	14.27 (20.5)
$\Delta T(\text{max-min}) \leq 20$ ns	0.0 (0.0)	0.0 (0.0)	8.68 (2.36)	14.13 (20.3)

Table 7.16: Cutflow for number of events passing each cut in the SR1 cutflow.

7.7.2 SR2

The results of the SR2 beam-on background ABCD estimate and signal observation is shown in 7.26⁵. The results are also summarized in Table 7.17.

LHC-Beam	Time (s)	Lumi (fb^{-1})	A	B	C	D	Est
Beam-On	12214771	124.7	2	82	1223	13	$0.87^{+0.33}_{-0.26}$
Beam-Off	16096891	0	0	1	49	$0^{+1.84}_{-0.0}$	$0.0^{+1.88}_{-0.0}$

Table 7.17: SR2 ABCD background estimates and signal observations for beam off and beam on data. 2 events are observed in beam-on data, yielding a 1.23σ agreement with the background prediction.

⁵Only the beam-on ABCD is shown due to poor statistics in the beam-off sample. Beam-induced backgrounds are much more significant in SR2 compared to SR1, as a direct hit from a throughgoing muon may potentially satisfy all selections except the maximum panel NPE requirement




Figure 7.26: ABCD background estimates and observations for beam-on data. 2 events are observed for a prediction of $0.87^{+0.33}_{-0.26}$. This observation is consistent with background.




Figure 7.27: Event displays of the two events satisfying all signal selections in SR2. Both events were present in the same 4-bar path, and in both all associated bar channels were maximally saturated for the entire length of the readout.

In beam-on data measured in 124.7 fb^{-1} , 2 events are observed. This agrees with the ABCD background prediction of $0.87^{+0.33}_{-0.26}$ to within 1.23σ . However, upon further inspection

of the two signal events, it is likely these two events are both through-going beam muons. Event displays of each event are shown in Figure 7.27.

These 2 events were predicted in good agreement with the ABCD background estimation, but the structure of the events reveals two features of the experiment and SR2 search design.

First, it was known that the front and back panels were not fully hermetic: the dimensions of the front and back panels are 25×25 cm, while the overall dimensions of the bar detector layers including structural spacing is 29.8×27.1 cm. This allows for the possibility of some beam muon paths to pass through the detector without directly interacting with the front or end panels. In principle, these events would be vetoed by the endcap panel requirement in SR2, but showers associated with the muon as it emerges from the cavern can generate small energy deposits in the front panel. Indeed, in both events, the front endcap panel showed an energy deposit. After these events were observed, the signal was quickly replicated in MC under a pathological placement of a beam muon near the corner of the exposed bars [112].

The fraction of the non-hermetically sealed region of the detector is approximately 1%; for the reconstructed beam muon cross section of 0.16 pb measured from Section 6.5, this suggests $O(100)$ muons may have passed through exposed bar layers with no corresponding panel coverage. As muons frequently shower or travel through the detector with a significant transverse displacement, it's reasonable to suggest that 2 of these muons satisfy all other signal selections and constitute the signal observed in SR2.

Secondly, separately from the issue of coverage, these muons could have been excluded from their energy character in the bars: A fully saturating deposit in the bars with only a very small deposit in the panels is not consistent with an mCP, as the energy deposit per length is very different in the two scintillator volumes, even considering the much smaller size of the endcap panels. An improvement to the SR2 search design would include a max/min energy requirement which includes the endcap panel in addition to the bars. This requirement will be used in the next generation of this analysis at the end of Run 3 and should significantly reduce expected backgrounds, regardless of whether endcap coverage was perfectly hermetic.

Finally, SR1 cut efficiencies are shown in Figure 7.25, and the cutflow yields are given in Table 7.18.




Figure 7.28: SR2 cumulative (top) and individual (bottom) cut efficiencies for beam-off and beam-on data.

Selection	Beam-Off	Beam-On	m=1.7 q=0.03	m=10.0 q=0.2
Total Events	1002647.0 (100.0)	790776.0 (100.0)	27.0 (88.0)	37.24 (87.99)
Digitizers Synchronized	1002617.0 (100.0)	790772.0 (100.0)	27.0 (88.0)	37.24 (87.99)
Pickup	1002617.0 (100.0)	790772.0 (100.0)	27.0 (88.0)	37.24 (87.99)
Dark Rate	1002612.0 (100.0)	790772.0 (100.0)	27.0 (88.0)	37.24 (87.99)
First Pulse	1002612.0 (100.0)	790772.0 (100.0)	27.0 (88.0)	37.24 (87.99)
Trigger Window	1001980.0 (99.93)	790383.0 (99.95)	26.99 (87.96)	37.24 (87.99)
Top/Side Panel Veto	445185.0 (44.4)	342109.0 (43.26)	16.61 (54.13)	22.71 (53.66)
4 Layers	4627.0 (0.46)	9152.0 (1.16)	7.65 (24.93)	9.54 (22.54)
Noise	4482.0 (0.45)	9113.0 (1.15)	7.65 (24.93)	9.54 (22.54)
Energy Max/Min $\leq 10(5)$	532.0 (0.05)	1827.0 (0.23)	7.56 (24.64)	9.5 (22.45)
Straight Line	83.0 (0.01)	1372.0 (0.17)	7.52 (24.51)	9.43 (22.28)
$\Delta T(\text{max-min}) \leq 20$ ns	55.0 (0.01)	1355.0 (0.17)	7.47 (24.35)	8.63 (20.39)
Front/Back Panel Required	50.0 (0.0)	1320.0 (0.17)	5.78 (18.84)	8.15 (19.26)
≤ 4 Bars	2.0 (0.0)	84.0 (0.01)	5.77 (18.81)	7.3 (17.25)
Max(Front/Back Panel nPE) < 70	0.0 (0.0)	2.0 (0.0)	5.77 (18.81)	6.99 (16.52)

Table 7.18: Cutflow for number of events passing each cut in the SR2 cutflow.

7.7.3 Combined

Simulated and observed signal yields for SR1 and SR2 are shown in Table 7.19, with select MC model points for each signal region. Cuts for the following lines have been combined:

- **Cosmic Muon Veto:** Signal skim selections (Table 7.5)
- **Pulse/Event Quality:** Digitizers synchronized, pickup, dark rate, first pulse, trigger window, (≤ 6 bars for SR1)
- **Shower Veto:** Top/Side panel veto, 4 layers hit
- **Beam Muon Veto:** Energy < 1000 keV

Selection Criteria	Signal Region 1			Signal Region 2		
	Data	Signal MC	Signal	Data	Signal	Signal
	Beam-On t=3393h	m=0.1 GeV Q/e=0.004	m=1.0 GeV Q/e=0.008	Beam-On t=3393h	m=1.7 GeV Q/e=0.03	m=10.0 GeV Q/e=0.2
Triggered Events	26864552	324.0	61.3	26864552	27.0	37.2
Cosmic Muon Veto	790776	324.0	61.3	790776	27.0	37.2
Pulse/Event Quality	506417	323.9	61.3	790383	27.0	37.2
Shower Veto	3369	12.0	19.3	9152	7.7	9.5
SR1: ≤ 4 Bars	985	11.7	19.3	—	—	—
Noise Filter	985	11.7	19.3	9113	7.7	9.5
Energy Max/Min	336	10.3	16.5	1827	7.6	9.5
SR1: Beam Muon Veto	331	10.3	16.5	—	—	—
SR1: End Panel Veto	209	10.1	14.3	—	—	—
Straight Line	3	9.2	14.3	1372	7.5	9.4
$\Delta T(\text{max-min}) \leq 20$ ns	0	8.7	14.1	1355	7.5	8.6
SR2: End Panel Required	—	—	—	1320	5.8	8.2
SR2: ≤ 4 Bars	—	—	—	84	5.8	7.3
SR2: Max(Front/Back Panel nPE) < 70	—	—	—	2	5.8	7.0

Table 7.19: Cutflows for SR1 and SR2 as well as several signal MC points. Definitions of cuts in table are provided in the text. A full version of this table with the condensed cuts expanded is available in Appendix A.12.

7.8 Systematics

In this section, I briefly summarize the primary sources of systematic uncertainty considered for the analysis.

Other sources of systematic uncertainty from trigger efficiency and material density and composition during signal propagation were assessed and found to be negligible.

Most signal for Run 3 was propagated through material at a set of similar, representative charges: 0.001, 0.01, and 0.1. The particle was then assigned its proper charge during Geant4 detector simulation. As mCP interactions in matter are weak, especially at low charge, these systematics arising from charge differences were ignored.

7.8.1 MC signal production systematics

Systematics on signal MC at production are the leading source of systematic uncertainty in the Run 3 analysis.

Signal yields are assigned a generic lognormal systematic uncertainty derived purely from 1σ fluctuations on the total production cross section. In addition, a 20% uncertainty on Drell-Yan production was used, reflecting the uncertainty from Madgraph at LO near $\eta \sim 0.1$, as in Figure 7.29.




Figure 7.29: Uncertainties on Drell-Yan production in Madgraph. Though we run MadGraph5_aMC@NLO, we take a conservative uncertainty using the LO band [113].

The total uncertainty on production can be seen in the shaded band in the plot of cross section vs. mass given by Figure 6.1.

7.8.2 Alignment

From studies of beam muon alignment in MC and data (Section 6.5), we identify an asymmetry in the vertical (y) detector alignment plane compared to the horizontal (x) plane. We use the asymmetry in x to derive a signal inefficiency and uncertainty due to misalignment.

To measure the degree of misalignment in data, we identify the largest hit in the first and last layer of the detector in beam muon events. Then, we measure the decrease in the event fractions with the same entry and exit position in the misaligned y-dimension as compared to the aligned x-dimension, as shown in Figure 7.30. The misaligned y-dimension has a straight line path event fraction of 0.88 ± 0.01 compared to the aligned x-dimension. We apply this factor to all MC signal yields to account for this misalignment.




Figure 7.30: The 1D correlation plot shows the misalignment in the y-axis (left). The ratio of the content in the 0-bin in the x-axis and y-axis can be used to derive a misalignment for signal. The 2D correlation shows the independence of each axis (right).

7.8.3 Luminosity

For each milliQan run, the start and end time of the run are recorded. The duration of the run is then compared to the fill luminosity and total beam time of the run, and the luminosity of the run is computed as:

$$L_{file} = L_{fill} \times \frac{t_{file}}{t_{fill}} \quad (7.7)$$

The total milliQan luminosity is then given by the sum over all files, $L = \sum_i L_{file}^i$. However, this procedure implies a fixed run luminosity, and the luminosity in a CMS fill is a decreasing function of the time since the fill began.

To account for this, a more granular measurement of the luminosity is taken for a sample

set of runs. This yielded a measurement of 8521 pb^{-1} compared to an initial luminosity estimate of 8031 pb^{-1} , a 6.1% total error. This value is taken to be a symmetric, systematic uncertainty on the central value of the two measurements. As a result, the total luminosity is given by $124.7 \pm 3.8 \text{ fb}^{-1}$. This luminosity uncertainty is subdominant compared to MC production uncertainties.

7.8.4 Timing and Energy Response

Systematics are extracted from the width of the timing distribution measured in each channel, as given by:

$$\Delta T_{\text{max-min}}^{\text{corrected}} = \left| (t_{\text{maxLayer}} - t_{\text{minLayer}}) \pm \sqrt{(t_{\text{maxLayer},\sigma})^2 + (t_{\text{minLayer},\sigma})^2} \right| \quad (7.8)$$

Energy systematics are similarly extracted from the width of Cd109 22 keV gamma peak, with energy uncertainties implemented as:

$$n_{PE}^{\text{corrected}} = n_{PE} \pm \frac{n_{PE}}{\mu_{\text{source}}} * \left(\sigma_{\text{source}} - \frac{\sigma_{\text{sim}} \mu_{\text{sim}}}{\mu_{\text{source}}} \right) \quad (7.9)$$

The impact of the timing and energy response systematics are given in Appendix A.13. Timing systematics are found to have a negligible on signal efficiency, which energy systematics are subdominant with an effect size peaking near 7% for heavier models.

7.9 Sensitivity

Simulated, pulse-injected signal yields are combined with the observed results and background estimates in SR1 and SR2 to produce a sensitivity estimate. Background were uncertainties from ABCD estimates were treated as Γ distributions with rate parameter $N=D\alpha$, with $\alpha = B/C$ for the number of events in ABCD regions B,C,D. The observed limit is computed using Combine, using the asymptotic likelihood method to assign a 95% CL.

New sensitivity is achieved for millicharged particles with charge $0.003 < q/e < 0.25$ and mass $0.45 \text{ GeV} < m < 20 \text{ GeV}$, as in Figure 7.31.




Figure 7.31: The observed limit for the Run 3 bar detector over 124.7 fb⁻¹ collected in 2023 and 2024.

During the remainder of Run 3, another 180 fb⁻¹ is expected to be recorded for the bar detector. Scaling expected signal and background yields by the new luminosity, but reducing total combined background yields in SR2 by a factor of 10 following the discussion in Section 7.7.2, Figure 7.32 projects the sensitivity for the full Run 3 dataset. The 2021 slab detector projection is included, but a full Run 3 analysis of the slab detector including its expected $\sim 150 \text{ fb}^{-1}$ dataset will follow in the subsequent full Run 3 analysis.




Figure 7.32: The observed limit for the Run 3 bar detector over 124.7 fb⁻¹ collected in 2023 and 2024 combined with an expected 180 fb⁻¹ collected in 2025 and 2026. The 2021 slab detector limit is shown for comparison.

8 Summary

In summary, I've presented two analyses: a search for long-lived particles using CMS, and a search for millicharged particles using milliQan.

For the CMS experiment, I developed the first trigger using ECAL timing (Chapter 3). A menu of these triggers was deployed during Run 3, and I led the first analysis targeting a Twin Higgs model (Chapter 4). We expect leading sensitivity for twin Higgs mass $M_H > 400$ GeV and lifetime $0.4 \text{ m} < c\tau < 100 \text{ m}$, to be completed in early 2026 after a final processing and analysis of the remaining 2024 data. Beyond the completion of this analysis, data collected using this new ECAL delayed jet trigger will be used for further analyses targeting LLPs of intermediate lifetimes.

Lessons learned from this analysis will help improve current and future timing triggers. Negative-time delayed jet triggers can provide robust validation samples for the currently implemented positive-time counterparts. Timing deployed in scouting at the HLT can be used complementary to the existing prompt and parking datasets to enhance rates, provide faster measurements of online performance, and target new models. And new offline timing algorithms developed for the analysis provide a roadmap for the improvement of online timing at the HLT.

For the milliQan experiment, I presented a complete description of the milliQan simulation framework (Chapter 6) and Run 3 milliQan bar detector analysis (Chapter 7). New sensitivity was achieved for millicharged particles with charge $0.003 < q/e < 0.25$ and mass $0.45 \text{ GeV} < m < 20 \text{ GeV}$.

The new sensitivity will be arguably the least important long-term impact of milliQan. The lessons learned during design, commissioning, operation, and analysis of the milliQan detector (Chapter 5) have already formed the intellectual basis of the SUBMET, FORMOSA, and LANSCE-mQ experiments, as well as for the milliQan slab detector. The mCP simulation

framework I have developed was used to create background and signal sensitivity estimates for each of those experiments, and the robust optical and energy modeling of the Geant4 detector simulation I developed has been and is being used to test new ideas about mCP experimental design. While we are approaching the end of data-taking for the LHC Run 3, for many of these mCP experiments, the data-taking has just begun.

As we approach the second half of the LHC era, we must confront the challenge of the lack of recent breakthrough particle discoveries with the reality that new, accessible breakthroughs must exist. It will take fundamentally new approaches, new technologies, new kinds of experiments to find what we have been missing. Through new triggers at CMS, new dedicated experiments at milliQan, and new tools for simulation and analysis along the way, I hope this work is able to contribute to the start of an agile and innovative new era in particle physics.

A Appendix

A.1 Delayed Jet Trigger Implementation




Figure A.1: HLTBTagPixelAndStripSetupForInclusiveDisplacedJets sub-sequence (trackless only).




Figure A.2: HLTIterativeTrackingIteration0ForDisplaced sub-sequence (trackless only).




Figure A.3: HLT2PromptTrackRequirementIter0DisplacedJetsLowPtSingle sub-sequence (trackless only).




Figure A.4: HLTIterativeTrackingIter12ForDisplaced sub-sequence (trackless only).




Figure A.5: HLT2PromptTrackRequirementIter12DisplacedJetsLowPtSingle sub-sequence (trackless only).

A.2 Complete List of Delayed Jet Triggers

Trigger	Parking	Prescale	Type
HLT_L1Tau_DelayedJet40_SingleDelay4nsInclusive_v6	No	1	Backup
HLT_L1Tau_DelayedJet40_SingleDelay3p75nsInclusive_v6	No	1	Backup
HLT_L1Tau_DelayedJet40_SingleDelay3p5nsInclusive_v8	No	1	Signal
HLT_L1Tau_DelayedJet40_SingleDelay3nsTrackless_v6	No	1	Backup
HLT_L1Tau_DelayedJet40_SingleDelay2p75nsTrackless_v6	No	1	Backup
HLT_L1Tau_DelayedJet40_SingleDelay2p75To4nsInclusive_v6	Yes	1	Signal
HLT_L1Tau_DelayedJet40_SingleDelay2p6To4nsInclusive_v6	Yes	0	Backup
HLT_L1Tau_DelayedJet40_SingleDelay2p5nsTrackless_v8	No	1	Signal
HLT_L1Tau_DelayedJet40_SingleDelay2p5To4nsInclusive_v6	Yes	0	Backup
HLT_L1Tau_DelayedJet40_DoubleDelay1p75nsInclusive_v6	No	1	Backup
HLT_L1Tau_DelayedJet40_DoubleDelay1p5nsTrackless_v6	No	1	Backup
HLT_L1Tau_DelayedJet40_DoubleDelay1p5nsInclusive_v6	No	1	Signal
HLT_L1Tau_DelayedJet40_DoubleDelay1p25nsTrackless_v6	No	1	Signal
HLT_L1Tau_DelayedJet40_DoubleDelay1p25nsInclusive_v8	No	50	Backup
HLT_L1Tau_DelayedJet40_DoubleDelay1nsTrackless_v8	No	50	Backup
HLT_L1Tau_DelayedJet40_DoubleDelay0p75nsInclusive_v8	No	1000	Backup
HLT_L1Tau_DelayedJet40_DoubleDelay0p5nsTrackless_v8	No	500	Backup
HLT_\$H_T\$430_DelayedJet40_SingleDelay2p5nsInclusive_v6	No	1	Backup
HLT_\$H_T\$430_DelayedJet40_SingleDelay2p25nsInclusive_v6	No	1	Backup
HLT_\$H_T\$430_DelayedJet40_SingleDelay2nsInclusive_v10	No	1	Signal
HLT_\$H_T\$430_DelayedJet40_SingleDelay1p5nsTrackless_v6	No	1	Backup
HLT_\$H_T\$430_DelayedJet40_SingleDelay1p5nsInclusive_v8	No	100	Backup
HLT_\$H_T\$430_DelayedJet40_SingleDelay1p25nsTrackless_v6	No	1	Backup
HLT_\$H_T\$430_DelayedJet40_SingleDelay1p25To1p75nsInclusive_v6	Yes	1	Signal
HLT_\$H_T\$430_DelayedJet40_SingleDelay1p1To1p6nsInclusive_v6	Yes	0	Backup
HLT_\$H_T\$430_DelayedJet40_SingleDelay1nsTrackless_v10	No	1	Signal
HLT_\$H_T\$430_DelayedJet40_SingleDelay1nsInclusive_v8	No	500	Backup
HLT_\$H_T\$430_DelayedJet40_SingleDelay1To1p5nsInclusive_v6	Yes	0	Backup
HLT_\$H_T\$430_DelayedJet40_SingleDelay0p5nsTrackless_v8	No	100	Backup
HLT_\$H_T\$430_DelayedJet40_SingleDelay0p5nsInclusive_v8	No	2500	Backup
HLT_\$H_T\$430_DelayedJet40_DoubleDelay1p5nsInclusive_v6	No	1	Backup
HLT_\$H_T\$430_DelayedJet40_DoubleDelay1p25nsInclusive_v6	No	1	Signal
HLT_\$H_T\$430_DelayedJet40_DoubleDelay1nsTrackless_v6	No	1	Backup
HLT_\$H_T\$430_DelayedJet40_DoubleDelay1nsInclusive_v10	No	50	Backup
HLT_\$H_T\$430_DelayedJet40_DoubleDelay0p75nsTrackless_v6	No	1	Backup
HLT_\$H_T\$430_DelayedJet40_DoubleDelay0p5nsTrackless_v10	No	1	Signal
HLT_\$H_T\$430_DelayedJet40_DoubleDelay0p5nsInclusive_v9	No	500	Backup
HLT_\$H_T\$350_DelayedJet40_SingleDelay3p5nsInclusive_v6	No	1	Backup
HLT_\$H_T\$350_DelayedJet40_SingleDelay3p25nsInclusive_v6	No	1	Backup
HLT_\$H_T\$350_DelayedJet40_SingleDelay3nsInclusive_v6	No	1	Signal
HLT_\$H_T\$350_DelayedJet40_SingleDelay1p75To3p5nsInclusive_v6	Yes	0	Signal
HLT_\$H_T\$350_DelayedJet40_SingleDelay1p6To3p5nsInclusive_v6	Yes	0	Backup
HLT_\$H_T\$350_DelayedJet40_SingleDelay1p5To3p5nsInclusive_v6	Yes	1	Backup

Table A.1: Full delayed-jet trigger paths used in Run 3 with Parking label (Yes = ParkingLLP, No = prompt/DisplacedJet), prescales, and classification.

A.3 ECAL Cell Timing: Ratio Method

The CMS ECAL extracts the time of the maximum pulse amplitude T_{Max} using the ratio method [114]. The object of this fitting method is to identify the time at which the pulse attains its maximum value, despite the fact that the sampling rate of 40 MHz (a 25 ns period)

is slow compared to a typical pulse width of approximately 100 ns. The basic strategy is to use the ratio of consecutive 25 ns samples and a precise measurement of the typical pulse shape to estimate the value of T_{Max} . A typical pulse template and the resulting fit is shown in Figure A.6.




Figure A.6: (Left) A typical ECAL pulse template with 10 samples at 25 ns intervals. (Right) Ratios of consecutive samples form a polynomial. The value of each sample ratio forms an estimate for $T - T_{Max}$ which can be used to predict the pulse T_{Max} [114].

The uncertainty of each ratio $R = A_i/A_{i+1}$ estimate of T_{Max} from the full pulse sample is given by:

$$\sigma_{T_{Max}} = \frac{dT(R)}{dR} \sigma_R \quad (\text{A.1})$$

where $T(R)$ is the polynomial fit to the T_{Max} obtained from the sample ratios, and σ_R is the uncertainty on the ratio. The final value for T_{Max} is obtained using a weighted estimate for each R_i satisfying $0.2 < R < 1.4$, with weights derived as the inverse of the uncertainty. The weights are highest in the rising component of the waveform, where the derivative of the polynomial fit is highest.

A.4 ECAL Cell Amplitude: Multifit

The CMS ECAL extracts the amplitude of a pulse in a crystal using a fitting method called “multifit” [54]. To manage the challenge of pileup contributions from early, out-of-time pulses, the multifit is computed as the sum of an in-time pulse with up to 9 out-of-time pulses: up to 5 occurring in earlier bunch crossings (BX), and up to 4 in later BX. The specific minimization is given below in Equation A.2.

$$\chi^2 = (\sum_{j=0}^{N_{BX}} A_j \vec{p}_j - \vec{S})^T \mathbf{C}^{-1} (\sum_{j=0}^{N_{BX}} A_j \vec{p}_j - \vec{S}) \quad (\text{A.2})$$

where the number of pulse templates considered is $N_{BX} \in [1, 10]$, the vector $\vec{S}$ contains the 10 digitized readout samples after pedestal subtraction, $\vec{p}_j$ contain the pulse templates for each BX, and A_j are the pulse amplitudes obtained by the fit in the 10 bunch crossings. The pulse templates p_j each have the same shape, but are shifted by $j \times 25$ ns to center on each bunch crossing. Additionally, the templates are extrapolated to a total of 15 samples, 5 samples beyond the last digitized sample, to properly model contributions from very early out-of-time pulses in later bunch crossings. This extrapolation is shown in Figure A.7.




Figure A.7: Extrapolating the pulse beyond the 10 samples used to form the template allows for a better estimate of out-of-time overlaps with later bunch crossings [54].

A.5 Run 3 Delayed Jet Analysis Yields

	MH400		MH1000		MH1500		Background	
	Yield	Eff%	Yield	Eff%	Yield	Eff%	Yield	Eff%
All events	14903.2	100.00%	201.5	100.00%	15.0	100.00%	372819190.0	100.00%
Passing Filters	14903.2	100.00%	201.5	100.00%	15.0	100.00%	371397475.0	99.62%
Delayed-Jet Trigger	3354.9	22.51%	162.2	80.48%	13.5	89.82%	99884305.0	26.89%
Event H_T Cut	252.6	7.53%	56.1	34.56%	2.6	18.96%	3903075.0	3.91%
NoCosmic	230.0	91.04%	48.6	86.65%	2.2	84.64%	3419125.0	87.60%
NoBeamHalo	230.0	100.00%	48.6	100.00%	2.2	100.00%	3419125.0	100.00%
NoSatellite	113.8	49.47%	34.0	69.93%	1.5	69.82%	1065740.0	31.17%
HasISR	83.9	73.75%	8.1	23.71%	0.2	15.94%	853515.0	80.09%
$P_T > 40$	83.9	100.00%	8.1	100.00%	0.2	100.00%	853515.0	100.00%
$ \eta < 1.48$	66.6	79.43%	7.4	92.00%	0.2	95.49%	704765.0	82.57%
$E_{ECAL} > 20$	61.3	91.92%	7.2	96.78%	0.2	98.35%	621015.0	88.12%
# ECAL Cells in Jet > 10	59.9	97.73%	7.1	99.34%	0.2	98.95%	602120.0	96.96%
$t_{RMS} < 1.5$	59.7	99.78%	7.1	99.59%	0.2	100.00%	590540.0	98.08%
$E_{HCAL} > 10$	59.7	100.00%	7.1	99.94%	0.2	100.00%	589190.0	99.77%
$E_{HCAL}/E_{ECAL} > 0.2$	59.3	99.30%	7.0	99.03%	0.2	100.00%	562175.0	95.41%
# ECAL Cells, 90% Energy > 7	59.2	99.88%	7.0	99.80%	0.2	99.66%	511815.0	91.04%
$E_{ECAL,CSC}/E_{ECAL} < 0.6$	56.3	95.10%	6.8	96.80%	0.2	96.87%	486260.0	95.01%
Shape Ratio < 1	56.3	100.00%	6.8	100.00%	0.2	100.00%	480570.0	98.83%
$dR > 0.3, t_{jet} > 1$	40.1	71.20%	5.9	86.95%	0.2	89.49%	3203.7	0.67%
$1 < t_{jet} < 1.5$	15.4	38.34%	0.9	14.54%	0.0	7.04%	2332.3	72.80%
$1.5 < t_{jet} < 2$	10.7	26.79%	0.8	13.55%	0.0	8.71%	793.0	24.75%
$2 < t_{jet} < 2.5$	6.3	15.66%	0.7	11.64%	0.0	9.03%	61.0	1.91%
$2.5 < t_{jet} < 3$	4.0	9.94%	0.6	10.78%	0.0	8.05%	15.7	0.49%
$3 < t_{jet} < 4$	2.3	5.65%	1.0	17.68%	0.0	15.40%	1.7	0.05%
$t_{jet} > 4$	1.5	3.62%	1.9	31.81%	0.1	51.77%	0.0	0.00%

Table A.2: 1-Jet Cutflow for H_T 500–700, signal scaled to the BSM Higgs cross-section, $B = 10^{-2}$.

	MH400		MH1000		MH1500		Background	
	Yield	Eff%	Yield	Eff%	Yield	Eff%	Yield	Eff%
All events	14903.2	100.00%	201.5	100.00%	15.0	100.00%	372819190.0	100.00%
Passing Filters	14903.2	100.00%	201.5	100.00%	15.0	100.00%	371397475.0	99.62%
Delayed-Jet Trigger	3354.9	22.51%	162.2	80.48%	13.5	89.82%	99884305.0	26.89%
Event H_T Cut	111.6	3.33%	29.9	18.44%	5.1	38.08%	1058805.0	1.06%
NoCosmic	102.0	91.37%	26.3	87.89%	4.4	84.67%	913500.0	86.28%
NoBeamHalo	102.0	100.00%	26.3	100.00%	4.4	100.00%	913500.0	100.00%
NoSatellite	46.2	45.33%	18.2	69.08%	3.1	70.37%	311410.0	34.09%
HasISR	40.6	87.73%	9.8	53.75%	0.8	25.71%	279635.0	89.80%
$P_T > 40$	40.6	100.00%	9.8	100.00%	0.8	100.00%	279635.0	100.00%
$ \eta < 1.48$	28.9	71.39%	9.0	91.91%	0.7	94.64%	232435.0	83.12%
$E_{ECAL} > 20$	26.2	90.37%	8.7	97.08%	0.7	98.77%	204665.0	88.05%
# ECAL Cells in Jet > 10	25.8	98.75%	8.6	99.03%	0.7	99.49%	198495.0	96.99%
$t_{RMS} < 1.5$	25.7	99.53%	8.6	99.59%	0.7	99.74%	194235.0	97.85%
$E_{HCAL} > 10$	25.7	99.84%	8.6	99.99%	0.7	99.89%	193715.0	99.73%
$E_{HCAL}/E_{ECAL} > 0.2$	25.5	99.30%	8.5	99.04%	0.7	98.53%	185355.0	95.68%
# ECAL Cells, 90% Energy > 7	25.3	99.28%	8.4	99.32%	0.7	99.23%	168155.0	90.72%
$E_{ECAL,CSC}/E_{ECAL} < 0.6$	24.1	95.28%	8.2	96.76%	0.7	96.87%	159225.0	94.69%
Shape Ratio < 1	24.1	100.00%	8.2	100.00%	0.7	100.00%	157545.0	98.94%
$dR > 0.3, t_{jet} > 1$	17.9	74.03%	7.2	87.52%	0.6	89.35%	1004.2	0.64%
$1 < t_{jet} < 1.5$	6.9	38.92%	1.3	17.48%	0.0	7.33%	744.4	74.13%
$1.5 < t_{jet} < 2$	4.6	25.59%	1.1	15.89%	0.1	8.86%	237.6	23.66%
$2 < t_{jet} < 2.5$	2.6	14.41%	1.0	14.25%	0.1	9.79%	17.6	1.76%
$2.5 < t_{jet} < 3$	2.0	11.36%	0.9	12.67%	0.1	10.02%	4.3	0.43%
$3 < t_{jet} < 4$	1.1	6.06%	1.3	18.35%	0.1	17.90%	0.1	0.01%
$t_{jet} > 4$	0.7	3.66%	1.5	21.35%	0.3	46.09%	0.0	0.00%

Table A.3: 1-Jet Cutflow for H_T 700–1000, signal scaled to the BSM Higgs cross-section, $B = 10^{-2}$.

	MH400		MH1000		MH1500		Background	
	Yield	Eff%	Yield	Eff%	Yield	Eff%	Yield	Eff%
All events	14903.2	100.00%	201.5	100.00%	15.0	100.00%	372819190.0	100.00%
Passing Filters	14903.2	100.00%	201.5	100.00%	15.0	100.00%	371397475.0	99.62%
Delayed-Jet Trigger	3354.9	22.51%	162.2	80.48%	13.5	89.89%	99884305.0	26.89%
Event H_T Cut	40.2	1.20%	14.0	8.61%	4.6	34.15%	265595.0	0.27%
NoCosmic	38.1	94.69%	12.1	86.35%	3.9	84.49%	220865.0	83.16%
NoBeamHalo	38.1	100.00%	12.1	100.00%	3.9	100.00%	220865.0	100.00%
NoSatellite	15.7	41.20%	8.4	69.50%	2.7	70.25%	80070.0	36.25%
HasISR	14.9	94.74%	7.5	89.80%	1.7	63.68%	76760.0	95.87%
$P_T > 40$	14.9	100.00%	7.5	100.00%	1.7	100.00%	76760.0	100.00%
$ \eta < 1.48$	9.7	65.04%	6.8	90.22%	1.6	92.71%	65000.0	84.68%
$E_{ECAL} > 20$	8.8	90.74%	6.6	96.49%	1.6	98.30%	57540.0	88.52%
# ECAL Cells in Jet > 10	8.7	99.62%	6.5	99.28%	1.6	99.45%	55970.0	97.27%
$t_{RMS} < 1.5$	8.7	99.84%	6.5	99.33%	1.6	99.89%	54735.0	97.79%
$E_{HCAL} > 10$	8.7	100.00%	6.5	99.98%	1.6	100.00%	54600.0	99.75%
$E_{HCAL}/E_{ECAL} > 0.2$	8.7	100.00%	6.4	98.98%	1.6	99.37%	52525.0	96.20%
# ECAL Cells, 90% Energy > 7	8.5	97.93%	6.3	98.18%	1.5	98.36%	47835.0	91.07%
$E_{ECAL,CSC}/E_{ECAL} < 0.6$	7.8	91.26%	6.1	96.97%	1.5	97.69%	44925.0	93.92%
Shape Ratio < 1	7.8	100.00%	6.1	100.00%	1.5	100.00%	44560.0	99.19%
$dR > 0.3, t_{jet} > 1$	4.8	61.78%	5.4	88.95%	1.4	91.55%	243.9	0.55%
$1 < t_{jet} < 1.5$	2.7	57.12%	0.9	16.55%	0.2	12.15%	183.3	75.14%
$1.5 < t_{jet} < 2$	0.9	18.23%	1.0	18.77%	0.2	14.06%	55.7	22.81%
$2 < t_{jet} < 2.5$	0.7	14.56%	0.9	16.51%	0.2	14.74%	4.1	1.67%
$2.5 < t_{jet} < 3$	0.2	3.85%	0.8	15.02%	0.2	11.69%	0.7	0.30%
$3 < t_{jet} < 4$	0.0	1.01%	0.9	17.15%	0.3	18.42%	0.1	0.05%
$t_{jet} > 4$	0.3	5.23%	0.9	16.01%	0.4	28.94%	0.0	0.02%

Table A.4: 1-Jet Cutflow for H_T 1000–10000, signal scaled to the BSM Higgs cross-section, $B = 10^{-2}$.

	MH400		MH1000		MH1500		Background	
	Yield	Eff%	Yield	Eff%	Yield	Eff%	Yield	Eff%
All events	14903.2	100.00%	201.5	100.00%	15.0	100.00%	372819190.0	100.00%
Passing Filters	14903.2	100.00%	201.5	100.00%	15.0	100.00%	371397475.0	99.62%
Delayed-Jet Trigger	3354.5	22.51%	162.2	80.48%	13.5	89.86%	99884305.0	26.89%
Event H_T Cut	252.4	7.53%	56.1	34.57%	2.6	18.96%	3903075.0	3.91%
NoCosmic	229.9	91.07%	48.6	86.65%	2.2	84.68%	3419125.0	87.60%
NoBeamHalo	229.9	100.00%	48.6	100.00%	2.2	100.00%	3419125.0	100.00%
NoSatellite	21.1	9.19%	20.8	42.84%	1.0	45.46%	242440.0	7.09%
$P_T > 40$	21.1	100.00%	20.8	100.00%	1.0	100.00%	242440.0	100.00%
$ \eta < 1.48$	18.8	89.09%	20.2	96.90%	1.0	98.93%	219060.0	90.36%
$E_{ECAL} > 20$	18.2	96.91%	20.0	99.02%	1.0	99.52%	207765.0	94.84%
# ECAL Cells in Jet > 10	18.1	99.33%	19.9	99.71%	1.0	99.91%	205115.0	98.72%
$t_{RMS} < 1.5$	18.1	100.00%	19.9	99.90%	1.0	99.96%	203715.0	99.32%
$E_{HCAL} > 10$	18.1	100.00%	19.9	99.97%	1.0	100.00%	203505.0	99.90%
$E_{HCAL}/E_{ECAL} > 0.2$	18.1	100.00%	19.8	99.74%	1.0	99.89%	197485.0	97.04%
# ECAL Cells, 90% Energy > 7	18.1	100.00%	19.8	99.90%	1.0	99.97%	185165.0	93.76%
$E_{ECAL,CSC}/E_{ECAL} < 0.6$	17.8	98.28%	19.6	99.00%	1.0	99.25%	178690.0	96.50%
Shape Ratio < 1	17.8	100.00%	19.6	100.00%	1.0	100.00%	177330.0	99.24%
$dR > 0.3, t_{jet} > 1$	15.2	85.40%	18.7	95.16%	0.9	95.14%	34.4	0.02%
$1 < t_{jet} < 1.5$	5.3	34.79%	2.1	11.21%	0.0	2.33%	23.5	68.32%
$1.5 < t_{jet} < 2$	4.5	29.82%	2.4	12.83%	0.0	5.18%	9.7	28.13%
$2 < t_{jet} < 2.5$	2.9	19.10%	2.4	12.88%	0.1	5.85%	0.9	2.58%
$2.5 < t_{jet} < 3$	1.1	7.42%	2.3	12.58%	0.1	7.34%	0.0	0.11%
$3 < t_{jet} < 4$	0.9	5.86%	3.7	20.02%	0.1	15.21%	0.0	0.00%
$t_{jet} > 4$	0.5	3.01%	5.7	30.48%	0.6	64.10%	0.3	0.87%

Table A.5: 2-Jet Cutflow for H_T 500–700, signal scaled to the BSM Higgs cross-section, $B = 10^{-2}$.

	MH400		MH1000		MH1500		Background	
	Yield	Eff%	Yield	Eff%	Yield	Eff%	Yield	Eff%
All events	14903.2	100.00%	201.5	100.00%	15.0	100.00%	372819190.0	100.00%
Passing Filters	14903.2	100.00%	201.5	100.00%	15.0	100.00%	371397475.0	99.62%
Delayed-Jet Trigger	3362.6	22.56%	162.2	80.50%	13.5	89.86%	99884305.0	26.89%
Event H_T Cut	111.8	3.32%	29.9	18.44%	5.2	38.14%	1058805.0	1.06%
NoCosmic	102.0	91.23%	26.3	87.87%	4.4	84.67%	913500.0	86.28%
NoBeamHalo	102.0	100.00%	26.3	100.00%	4.4	100.00%	913500.0	100.00%
NoSatellite	8.0	7.84%	11.3	42.95%	2.2	50.00%	86650.0	9.49%
$P_T > 40$	8.0	100.00%	11.3	100.00%	2.2	100.00%	86650.0	100.00%
$ \eta < 1.48$	7.0	87.59%	10.8	95.89%	2.1	97.78%	78235.0	90.29%
$E_{ECAL} > 20$	6.9	98.85%	10.7	98.97%	2.1	99.45%	73710.0	94.22%
# ECAL Cells in Jet > 10	6.9	100.00%	10.7	99.66%	2.1	99.74%	72635.0	98.54%
$t_{RMS} < 1.5$	6.8	98.83%	10.7	99.81%	2.1	99.89%	71970.0	99.08%
$E_{HCAL} > 10$	6.8	100.00%	10.7	99.98%	2.1	99.98%	71875.0	99.87%
$E_{HCAL}/E_{ECAL} > 0.2$	6.8	100.00%	10.6	99.67%	2.1	99.55%	69965.0	97.34%
# ECAL Cells, 90% Energy > 7	6.8	99.49%	10.6	99.83%	2.1	99.87%	65500.0	93.62%
$E_{ECAL,CSC}/E_{ECAL} < 0.6$	6.7	98.81%	10.5	98.65%	2.1	98.98%	63090.0	96.32%
Shape Ratio < 1	6.7	100.00%	10.5	100.00%	2.1	100.00%	62630.0	99.27%
$dR > 0.3, t_{jet} > 1$	5.8	85.53%	10.0	96.03%	2.0	96.57%	15.5	0.02%
$1 < t_{jet} < 1.5$	1.6	28.33%	1.1	10.50%	0.1	4.77%	10.9	70.64%
$1.5 < t_{jet} < 2$	1.2	21.07%	1.5	14.72%	0.2	8.73%	4.2	26.88%
$2 < t_{jet} < 2.5$	1.5	25.49%	1.6	15.92%	0.2	9.13%	0.3	1.84%
$2.5 < t_{jet} < 3$	0.6	9.80%	1.5	14.71%	0.2	10.15%	0.0	0.19%
$3 < t_{jet} < 4$	0.6	10.25%	2.2	21.73%	0.4	22.25%	0.0	0.04%
$t_{jet} > 4$	0.3	5.06%	2.3	22.42%	0.9	44.96%	0.1	0.42%

Table A.6: 2-Jet Cutflow for H_T 700–1000, signal scaled to the BSM Higgs cross-section, $B = 10^{-2}$.

	MH400		MH1000		MH1500		Background	
	Yield	Eff%	Yield	Eff%	Yield	Eff%	Yield	Eff%
All events	14903.2	100.00%	201.5	100.00%	15.0	100.00%	372819190.0	100.00%
Passing Filters	14903.2	100.00%	201.5	100.00%	15.0	100.00%	371397475.0	99.62%
Delayed-Jet Trigger	3362.6	22.56%	162.2	80.50%	13.5	89.86%	99884305.0	26.89%
Event H_T Cut	40.3	1.20%	14.0	8.61%	4.6	34.08%	265595.0	0.27%
NoCosmic	38.3	94.88%	12.1	86.35%	3.9	84.49%	220865.0	83.16%
NoBeamHalo	38.3	100.00%	12.1	100.00%	3.9	100.00%	220865.0	100.00%
NoSatellite	2.4	6.36%	4.9	40.63%	2.0	51.24%	26455.0	11.98%
$P_T > 40$	2.4	100.00%	4.9	100.00%	2.0	100.00%	26455.0	100.00%
$ \eta < 1.48$	1.8	74.84%	4.7	95.50%	1.9	95.91%	23985.0	90.66%
$E_{ECAL} > 20$	1.7	95.56%	4.6	98.53%	1.9	99.39%	22510.0	93.85%
# ECAL Cells in Jet > 10	1.7	100.00%	4.6	99.83%	1.9	99.72%	22195.0	98.60%
$t_{RMS} < 1.5$	1.7	100.00%	4.6	99.82%	1.9	99.96%	21980.0	99.03%
$E_{HCAL} > 10$	1.7	100.00%	4.6	99.99%	1.9	100.00%	21955.0	99.89%
$E_{HCAL}/E_{ECAL} > 0.2$	1.7	100.00%	4.6	99.49%	1.9	99.50%	21535.0	98.09%
# ECAL Cells, 90% Energy > 7	1.7	100.00%	4.5	99.55%	1.9	99.46%	20250.0	94.03%
$E_{ECAL,CSC}/E_{ECAL} < 0.6$	1.7	94.92%	4.5	98.44%	1.9	98.91%	19385.0	95.73%
Shape Ratio < 1	1.7	100.00%	4.5	100.00%	1.9	100.00%	19285.0	99.48%
$dR > 0.3, t_{jet} > 1$	1.3	76.53%	4.3	95.03%	1.8	96.36%	0.0	0.00%
$1 < t_{jet} < 1.5$	0.7	55.94%	0.5	12.25%	0.2	10.24%	0.0	0.00%
$1.5 < t_{jet} < 2$	0.5	36.47%	0.8	17.70%	0.3	14.62%	0.0	0.00%
$2 < t_{jet} < 2.5$	0.0	2.04%	0.7	16.06%	0.3	16.14%	0.0	0.00%
$2.5 < t_{jet} < 3$	0.1	5.55%	0.7	16.11%	0.2	13.73%	0.0	0.00%
$3 < t_{jet} < 4$	0.0	0.00%	0.9	20.18%	0.3	19.19%	0.0	0.00%
$t_{jet} > 4$	0.0	0.00%	0.8	17.72%	0.5	26.08%	0.0	0.00%

Table A.7: 2-Jet Cutflow for H_T 1000–10000, signal scaled to the BSM Higgs cross-section, $B = 10^{-2}$.

	MH400		MH1000		MH1500		Background	
	Yield	Eff%	Yield	Eff%	Yield	Eff%	Yield	Eff%
All events	14903.2	100.00%	201.5	100.00%	15.0	100.00%	372819190.0	100.00%
Passing Filters	14903.2	100.00%	201.5	100.00%	15.0	100.00%	371397475.0	99.62%
Delayed-Jet Trigger	3354.9	22.51%	162.2	80.48%	13.5	89.76%	99884305.0	26.89%
Event H_T Cut	252.6	7.53%	56.1	34.56%	2.6	18.97%	3903075.0	3.91%
NoCosmic	230.0	91.04%	48.6	86.65%	2.2	84.64%	3419125.0	87.60%
NoBeamHalo	230.0	100.00%	48.6	100.00%	2.2	100.00%	3419125.0	100.00%
NoSatellite	21.1	9.17%	20.8	42.84%	1.0	45.48%	242440.0	7.09%
HasISR	15.8	75.09%	3.9	18.67%	0.1	12.79%	146885.0	60.59%
$P_T > 40$	15.8	100.00%	3.9	100.00%	0.1	100.00%	146885.0	100.00%
$ \eta < 1.48$	13.9	87.88%	3.8	97.00%	0.1	99.62%	132460.0	90.18%
$E_{ECAL} > 20$	13.3	95.82%	3.7	99.49%	0.1	100.00%	124925.0	94.31%
# ECAL Cells in Jet > 10	13.2	99.10%	3.7	99.79%	0.1	100.00%	123270.0	98.68%
$t_{RMS} < 1.5$	13.2	100.00%	3.7	99.96%	0.1	100.00%	122410.0	99.30%
$E_{HCAL} > 10$	13.2	100.00%	3.7	99.94%	0.1	100.00%	122270.0	99.89%
$E_{HCAL}/E_{ECAL} > 0.2$	13.2	100.00%	3.7	99.79%	0.1	100.00%	119305.0	97.58%
# ECAL Cells, 90% Energy > 7	13.2	100.00%	3.7	99.94%	0.1	100.00%	112695.0	94.46%
$E_{ECAL,CSC}/E_{ECAL} < 0.6$	13.0	98.26%	3.7	98.86%	0.1	99.34%	109175.0	96.88%
Shape Ratio < 1	13.0	100.00%	3.7	100.00%	0.1	100.00%	108305.0	99.20%
$dR > 0.3, t_{jet} > 1$	11.1	85.34%	3.5	94.02%	0.1	91.30%	12.4	0.01%
$1 < t_{jet} < 1.5$	4.3	38.89%	0.4	10.19%	0.0	3.38%	9.1	73.52%
$1.5 < t_{jet} < 2$	3.1	28.14%	0.4	10.94%	0.0	5.53%	3.0	24.07%
$2 < t_{jet} < 2.5$	2.0	17.72%	0.4	10.73%	0.0	6.03%	0.3	2.35%
$2.5 < t_{jet} < 3$	0.7	6.73%	0.4	10.22%	0.0	4.35%	0.0	0.06%
$3 < t_{jet} < 4$	0.6	5.81%	0.7	19.65%	0.0	13.99%	0.0	0.00%
$t_{jet} > 4$	0.3	2.71%	1.3	38.27%	0.1	66.72%	0.0	0.00%

Table A.8: 2-Jet Cutflow for H_T 500–700, signal scaled to the BSM Higgs cross-section, $B = 10^{-2}$.

	MH400		MH1000		MH1500		Background	
	Yield	Eff%	Yield	Eff%	Yield	Eff%	Yield	Eff%
All events	14903.2	100.00%	201.5	100.00%	15.0	100.00%	372819190.0	100.00%
Passing Filters	14903.2	100.00%	201.5	100.00%	15.0	100.00%	371397475.0	99.62%
Delayed-Jet Trigger	3347.8	22.46%	162.0	80.40%	13.5	89.76%	99884305.0	26.89%
Event H_T Cut	111.3	3.33%	29.9	18.43%	5.1	38.08%	1058805.0	1.06%
NoCosmic	101.7	91.38%	26.2	87.89%	4.4	84.67%	913500.0	86.28%
NoBeamHalo	101.7	100.00%	26.2	100.00%	4.4	100.00%	913500.0	100.00%
NoSatellite	8.0	7.82%	11.3	42.94%	2.2	50.04%	86650.0	9.49%
HasISR	6.6	83.14%	5.6	49.35%	0.5	22.94%	64590.0	74.54%
$P_T > 40$	6.6	100.00%	5.6	100.00%	0.5	100.00%	64590.0	100.00%
$ \eta < 1.48$	5.9	89.13%	5.4	96.40%	0.5	97.16%	58000.0	89.80%
$E_{ECAL} > 20$	5.8	98.63%	5.3	99.03%	0.5	99.77%	54300.0	93.62%
# ECAL Cells in Jet > 10	5.8	100.00%	5.3	99.65%	0.5	99.60%	53480.0	98.49%
$t_{RMS} < 1.5$	5.7	98.62%	5.3	99.82%	0.5	100.00%	52965.0	99.04%
$E_{HCAL} > 10$	5.7	100.00%	5.3	99.98%	0.5	100.00%	52900.0	99.88%
$E_{HCAL}/E_{ECAL} > 0.2$	5.7	100.00%	5.3	99.62%	0.5	99.21%	51590.0	97.52%
# ECAL Cells, 90% Energy > 7	5.7	99.39%	5.3	99.83%	0.5	99.71%	48655.0	94.31%
$E_{ECAL,CSC}/E_{ECAL} < 0.6$	5.6	98.59%	5.2	98.49%	0.5	99.10%	47055.0	96.71%
Shape Ratio < 1	5.6	100.00%	5.2	100.00%	0.5	100.00%	46690.0	99.22%
$dR > 0.3, t_{jet} > 1$	4.7	83.50%	4.9	93.90%	0.4	94.28%	12.4	0.03%
$1 < t_{jet} < 1.5$	1.4	30.59%	0.7	13.64%	0.0	8.36%	9.2	74.03%
$1.5 < t_{jet} < 2$	1.1	23.47%	0.7	14.04%	0.0	8.29%	3.0	23.98%
$2 < t_{jet} < 2.5$	0.9	19.86%	0.7	14.54%	0.0	10.34%	0.2	1.78%
$2.5 < t_{jet} < 3$	0.5	11.68%	0.6	13.28%	0.0	9.45%	0.0	0.15%
$3 < t_{jet} < 4$	0.5	9.70%	1.0	20.74%	0.1	17.14%	0.0	0.05%
$t_{jet} > 4$	0.2	4.69%	1.2	23.76%	0.2	46.43%	0.0	0.00%

Table A.9: 2-Jet Cutflow for H_T 700–1000, signal scaled to the BSM Higgs cross-section, $B = 10^{-2}$.

	MH400		MH1000		MH1500		Background	
	Yield	Eff%	Yield	Eff%	Yield	Eff%	Yield	Eff%
All events	14903.2	100.00%	201.5	100.00%	15.0	100.00%	372819190.0	100.00%
Passing Filters	14903.2	100.00%	201.5	100.00%	15.0	100.00%	371397475.0	99.62%
Delayed-Jet Trigger	3347.8	22.46%	162.0	80.40%	13.5	89.76%	99884305.0	26.89%
Event H_T Cut	40.0	1.20%	13.9	8.61%	4.6	34.12%	265595.0	0.27%
NoCosmic	37.9	94.69%	12.0	86.35%	3.9	84.50%	220865.0	83.16%
NoBeamHalo	37.9	100.00%	12.0	100.00%	3.9	100.00%	220865.0	100.00%
NoSatellite	2.4	6.38%	4.9	40.68%	2.0	51.23%	26455.0	11.98%
HasISR	2.4	100.00%	4.4	89.20%	1.2	61.55%	22725.0	85.90%
$P_T > 40$	2.4	100.00%	4.4	100.00%	1.2	100.00%	22725.0	100.00%
$ \eta < 1.48$	1.8	74.83%	4.2	95.05%	1.2	95.70%	20515.0	90.28%
$E_{ECAL} > 20$	1.7	92.52%	4.1	98.54%	1.2	99.35%	19190.0	93.54%
# ECAL Cells in Jet > 10	1.7	100.00%	4.1	99.84%	1.2	99.74%	18935.0	98.67%
$t_{RMS} < 1.5$	1.7	100.00%	4.1	99.82%	1.2	100.00%	18755.0	99.05%
$E_{HCAL} > 10$	1.7	100.00%	4.1	99.98%	1.2	100.00%	18735.0	99.89%
$E_{HCAL}/E_{ECAL} > 0.2$	1.7	100.00%	4.1	99.46%	1.2	99.42%	18390.0	98.16%
# ECAL Cells, 90% Energy > 7	1.7	100.00%	4.0	99.56%	1.2	99.46%	17345.0	94.32%
$E_{ECAL,CSC}/E_{ECAL} < 0.6$	1.6	95.19%	4.0	98.37%	1.1	98.97%	16660.0	96.05%
Shape Ratio < 1	1.6	100.00%	4.0	100.00%	1.1	100.00%	16570.0	99.46%
$dR > 0.3, t_{jet} > 1$	1.1	67.92%	3.6	91.61%	1.1	93.93%	0.0	0.00%
$1 < t_{jet} < 1.5$	0.7	61.02%	0.7	17.97%	0.2	14.18%	0.0	0.00%
$1.5 < t_{jet} < 2$	0.3	29.32%	0.7	18.00%	0.1	13.23%	0.0	0.00%
$2 < t_{jet} < 2.5$	0.0	3.21%	0.6	16.04%	0.1	13.49%	0.0	0.00%
$2.5 < t_{jet} < 3$	0.1	6.44%	0.5	14.57%	0.1	11.58%	0.0	0.00%
$3 < t_{jet} < 4$	0.0	0.00%	0.6	17.35%	0.2	18.70%	0.0	0.00%
$t_{jet} > 4$	0.0	0.00%	0.6	16.07%	0.3	28.82%	0.0	0.00%

Table A.10: 2-Jet Cutflow for H_T 1000–10000, signal scaled to the BSM Higgs cross-section, $B = 10^{-2}$.

A.6 Run 3 Delayed Jet N-1 Plots




Figure A.8: N-1 plot in Jet P_T with a cut at $P_T < 40$.




Figure A.9: N-1 plot in Jet CSCFrac with a cut at $\text{CSCFrac}_i < 0.6$.




Figure A.10: N-1 plot in Jet totalEmEnergy with a cut at $E_{EM} > 20$.




Figure A.11: N-1 plot in Jet abs(eta) with a cut at $|\eta| < 1.48$.




Figure A.12: N-1 plot in Jet hadEnergyFrac with a cut at hadEnergyFrac > 0.2 .




Figure A.13: N-1 plot in Jet hadEnergy with a cut at hadEnergy > 10 .




Figure A.14: N-1 plot in Jet N_{90} with a cut at $N_{90} > 7$.




Figure A.15: N-1 plot in Jet N_{Cells} with a cut at $N_{Cells} > 10$.

A.7 Delayed Jet Analysis Event Flow




Figure A.16: Events flow into signal samples, control samples, or are rejected according to event level requirements. Events with a jet passing all requirements are then binned in an ISR requirement, H_T , and jet multiplicity.

A.8 Simulation Class Description

This section summarizes the main milliQan Geant4 detector simulation source files and their responsibilities. File names are shown in **bold**; each entry describes what data/behavior it controls.

A.8.1 Core Geometry, Materials, and Optics

mqDetectorConstruction Assembles the detector geometry and material properties. Modify here the number of bars per layer, the number of layers, the bar material(s), bar shape, and related geometric composition. Also defines optical interfaces between volumes (indices of refraction, reflectivities), the scintillation spectrum for the scintillator, and PMT definitions including quantum efficiency (QE). This is the file most frequently edited.

Types (geometry presets)

- **default**: 4×4 milliQan Run-3 bar detector.
- **defaultNoCavern**: Same as **default**, with the cavern disabled.
- **NaIOnly**: Minimal setup with a PMT, NaI crystal, and plastic mount.
- **superMod**: A milliQan supermodule with an explicit aluminum shell.

mqBarParameterisation, mqPMTParameterisation Parameterisations used to place/construct bars and PMTs. Modify the structure of the detector parameterisation here.

A.8.2 Stepping, Sensitive Detectors, and Hits

mqSteppingAction Handles particle propagation through the geometry at the step level. Access information about each step, the current particle, and what it interacted with. Example use cases: inspect neutron captures and their target nuclei; check if a photon underwent Compton scattering; and perform first-pass photon detection at the PMT boundary before handing off to **mqPMTSD**. Also treats edge cases such as photons approaching the PMT at sharp angles that should not count as detected.

mqPMTSD, mqScintSD Sensitive detectors. For PMTs, extract and record properties of incoming particles at detection (energy, position, momentum, hit time, etc.) and store them in the global PMT-hit collection.

mqPMTHit/mqScintHit and mqPMTRHit/mqScintRHit The **Hit** classes define what per-hit information exists (e.g., what fields a PMT photon hit stores). Add new per-hit details here. The corresponding **RHit** classes map those hits into ROOT objects for output.

A.8.3 Tracking and Event/Run Hooks

mqMuonTrack/mqNeutronTrack/mq...Track Definitions for explicitly tracked particle types. Geant4 can transport many particle species, but to save initial/final positions, momenta, etc., for a given species you must opt in here.

mqTrackingAction Collects and stores information for explicitly tracked particles.

mqUserEventInformation Defines event-level parameters that are global to the event, independent of any specific hit or track.

mqUserTrackInformation Track-level user information for all particle types.

mqEventAction Aggregates top-level event information, including access to hit collections and tracked-particle collections.

mqRunAction and associated messenger Entry/exit hooks at the run level. Typically where run-level parameters would be adjusted. In this codebase they are used minimally

(explicit RNG seeding, passing a histogram manager).

A.8.4 Physics Lists and Primary Generation

mqShieldingPhysicsList.cc Selects and configures the physics list for the simulation. Edit here to switch to a different physics list if desired, or to turn on/off different processes. The most typical process to disable/enable is optical photon tracking, as it slows the simulation down by a lot when turned on and sometimes energy tracking is sufficient.

MilliQEMPhysics, MilliQMuonPhysics Explicit process definitions for mCP physics (legacy; overridden by current physics list).

MilliQMonopole mCP particle definition constructor, deriving from G4Monopole.

MilliQMonopoleEquation Equation of motion for mCP propagation in a magnetic field. This given explicitly below:

$$\mathbf{y} \equiv (x, y, z, p_x, p_y, p_z, -, t), \quad (\text{A.3})$$

$$E = \sqrt{\mathbf{p}^2 + m^2 c^2}, \quad p \equiv \|\mathbf{p}\|, \quad (\text{A.4})$$

$$\frac{d\mathbf{r}}{ds} = \frac{\mathbf{p}}{p}, \quad (\text{A.5})$$

$$\frac{d\mathbf{p}}{ds} = \frac{g E}{p} \mathbf{B} + \frac{q}{p} (\mathbf{p} \times \mathbf{B}), \quad (\text{A.6})$$

$$\frac{dt}{ds} = \frac{E}{p c}. \quad (\text{A.7})$$

Here q is the electric charge (fElCharge) and g is the magnetic charge (fMagCharge). For mCPs the magnetic charge is set to zero, but as the particle type inherits from G4Monopole the full equation of motion is provided. Notably, the magnetic field was turned off for signal mCP simulations; studies showed a minimal impact as $q/p \ll 1$ for most models, and disabling the B-field significantly reduced runtime.

MilliQMonopoleFieldSetup and messenger Magnetic-field setup and wiring to the mCP equation of motion.

MilliQMonopolePhysics and messenger Implements mCP ionization and multiple scattering.

MilliQPrimaryGeneratorAction and messenger Reads particle charge, mass, energy, position, and momentum from predefined input files. Used for outputs of mCP and beam muon propagation.

mqPrimaryGeneratorAction and messenger References GPS commands for particle initialization. Used for bench tests (Cd109) and cosmic muon tests.

A.8.5 Output, Filenames, and Logging

mqHistoManager Creates and writes the ROOT output file(s) and manages histograms/ntuples.

mqGenFileName Utility to generate output filenames (timestamps, run tags, etc.).

mqSession and associated messenger Manages high-level logging/output and allows simple user commands during a session. No custom implementation but provided for ease of use.

mqSteppingVerbose Step-level verbose logger for detailed per-step diagnostics.

A.9 Simulation Output Reference

This appendix documents the parameters written to the ROOT outputs from the simulation. The example files shown here concern millicharged particles (mCPs), but for runs with muons the **Gamma/Electron/Muon** data structures will likewise populate with the relevant data.

A.9.1 Track-Level Information (e.g., MuonTrack)

trackID Unique (per-event) identifier for a single track. Primaries are always **1** with **parentID = 0**.

initialProcessName, finalProcessName Generating/terminating process names (as **TString**).

initialVolumeName, finalVolumeName Volume names at creation/destruction.

initialCopyNo, finalCopyNo Copy numbers at creation/destruction.

initialTime, finalTime Geant4 event times of track creation/destruction.

initialEnergy, finalEnergy Track energies (MeV) at creation/destruction.

energyDeposit_MeV Scored deposited energy for the track.¹

parentID trackID of the parent particle.

initialPosition, finalPosition 3D positions at creation/destruction.

totalTrackLength Path length between creation and destruction; particles exiting the simulation without absorption typically have lengths of order $\gtrsim 20$ m.

Xposition, Yposition, Zposition Entry position of muons into the milliQan cavern.

Xfposition, Yfposition, Zfposition Exit position of muons from the milliQan cavern.

nbOfElastics, nbOfInteractions Currently unused interaction counters.

energyDiff Difference between initial and final track energy.

muonTrigUp, muonTrigLow Currently unused muon-telescope trigger flags.

A.9.2 Hit-Level Concepts

There are two hit classes:

- **PMTHit** — Populates when an optical photon reaches a PMT photocathode. At photocathode entry, the photon track is terminated and a detection is decided according to the PMT quantum efficiency (QE). The track is killed regardless of outcome. This class pertains to photon detection only.

¹In Geant4, “deposited energy” is not simply “energy lost.” To avoid EM divergences at low energies, would-be low-range secondaries (shorter than the step length in the current material, ~ 0.7 mm by default in this setup) are not spawned explicitly; instead, their total energy is subtracted from the parent and recorded as a deposit. To obtain the total energy loss, add the deposited energy to the energies carried by any explicitly generated secondaries.

- **ScintHit** — Populates when a user-specified particle (typically a muon or an mCP, depending on the primary) enters and exits a scintillator volume. Temporary values are recorded on entry; a hit is finalized if the particle exits the same volume. No detection efficiency is applied here; this serves as an MC-truth record of particle traversal.

A.9.3 PMT Hit Fields

pmtNumber Copy number of the PMT. For bars, this equals the bar number; for slabs (4 PMTs per slab), numbering differs (see detector numbering notes).

initialHitTime Time of the hit in the PMT.

trackID Identifier of the track that generated the hit.

parentID Identifier of the parent of the track that generated the hit.

A.9.4 Scintillator Hit Fields

Whenever a particle enters or exits a scintillator, a hit is generated containing the particle's energy information over that traversal. Summing all hits for a given scintillator yields that scintillator's total energy deposit for the event.

trackID Track ID of the particle traversing the scintillator.

parentID Track ID of the parent of the traversing particle.

EDep_MeV Total energy loss of the particle from entry to exit (initial minus final).

trackLength Path length inside the scintillator from entry to exit.

copyNo Copy number of the scintillator volume (see detector numbering notes for bars/slabs).

hitTime Time at which the particle entered the scintillator.

exitTime Time at which the particle exited the scintillator.

particleName Name of the particle (e.g., **e+**, **neutron**) that produced the hit.

hitPosition, exitPosition Entry and exit positions (x, y, z) of the hit.

A.9.5 Event-Level Information

These quantities appear once per event. For a run of N events, each parameter below contains N entries.

EventID Event index $\{0, 1, 2, \dots, N-1\}$.

NbGamma/Muon/Photon/Electron/... Tracks For each explicitly tracked particle type, counts the number of tracks in the event. Empty for untracked species (e.g., if photons are disabled) or for species not produced (e.g., mCP counters in a pure muon run).

NbPMTHits, NbScintHits Total number of hits recorded in the event. For PMTs, this is the total number of detected photons across all PMTs; for Scint, this is the total number of successful bar/panel/slab traversals by the primary of interest.

NbOfCerenkovPhotons, NbOfScintillationPhotons Total numbers of optical photons generated by the corresponding processes.

absorptionCount, **boundaryAbsorptionCount** The former counts absorptions by **OpAbsorption**; the latter counts absorptions at boundaries (**OpBoundary** with **boundaryStatus = Absorbed**).

pmtsAboveTrigger Currently empty; can be used to apply a trigger threshold and process triggers at runtime.

Edep_MeV Sum of deposited energies (implicit losses) over the event (by particle type or globally, as configured).

barHit, **slabHit**, **panelHit** Counters of how many times the primary of interest entered bars/slabs/panels, respectively. Currently unused (historically for the demonstrator); can be re-enabled if needed.

muonTrig Not used at present; may be used to define triggers based on geometry (e.g., passage) rather than PE thresholds at runtime.

eventWeight Event weight from the generator (used when injecting four-vectors).

ProcessID Generator process type identifier (used when injecting four-vectors).

Detector Numbering. Bar and slab (panel) copy-number schemes differ. For bars, pmtNumber generally matches the bar number; for slabs with four PMTs per slab, a distinct numbering convention is used. Refer to the bar and slab numbering documentation convention in the milliQan simulation README.

A.10 Scintillator Surface Roughness

A study of similarly polished scintillator to EJ200 found the typical scintillator surface roughness feature had an angular size of 3° , as in Figure A.17 [115].




Figure A.17: (Left) A measurement of the surface roughness of a machine-polished plastic scintillator sample. (Right) The widths of the distributions of angles as a function of the feature size [115].

For a critical angle of 39.3° , the value for the EJ200-air interface for a refractive index

$n = 1.58$, this can be modeled as a random deflection in the surface normal by 3° . If we evaluate:

$$\int_{37.8^\circ}^{40.8^\circ} \sin \theta d\theta = 0.033$$

So up to 3.3% of trajectories for an isotropically emitted scintillation photon will be potentially transmitted rather than internally reflected. However, half of the features will make the angle of incidence shallower rather than steeper, so only 1.65% of photons will experience transmission when they would have otherwise reflected.

To first model this effect, for 1% of photons which totally internally reflect, we instead randomly rescattered their momentum direction. After studies showed the efficiency to detect these photons dropped dramatically upon rescatter, we instead killed these photon tracks to save computationally expensive optical photon tracking time without a significant impact to total light collection efficiency modeling. The total modeling of the 1% scintillator roughness showed good agreement with the estimates of the bulk attenuation length compared to experimental validation studies.

A.11 MilliQan Offline Tree Branches

Table A.11: Run/DAQ, trigger, channel, and pulse-level variables.

Name	Type	Description
event	int	Event number
runNumber	long int	The number of the run
fileNumber	long int	Each run is split into multiple files (1k events each). This is the index of one of those files
boardsMatched	bool	Are all digitizers matched together
DAQEventNumber	int	Event number stored in DAQ
daqFileOpen	ulong	Time raw DAQ file was opened (Unix time)
daqFileClose	ulong	Time raw DAQ file was closed (Unix time)
maxPulseIndex	int	Index of the max pulse in an event
nConsecSamples	vector<int>	Number of consecutive samples above threshold required to define a pulse
nConsecSamplesEnd	vector<int>	Number of consecutive samples after a pulse required before defining a new pulse
highThreshold	vector<int>	Threshold voltage must be over to define a pulse
lowThreshold	vector<int>	Threshold voltage must dip below to end pulse
totalFillLumi	float	Total luminosity of LHC fill (μb^{-1})
fillId	int	List of LHC fill IDs this file contains
beamType	TString	Beam types (null, PROTONS, IONS)
beamEnergy	float	Energy of beam (GeV)
betaStar	float	β^* of beam (cm)

Continued on next page

Name	Type	Description
beamOn	bool	LHC beam on (True), beam off (False)
fillStart	ulong	Time of the LHC fill start (Unix time)
fillEnd	ulong	Time of the LHC fill end (Unix time)
beamInFill	bool	Beam present at any point in fill
goodRunLoose	bool	Run passes goodRunLoose requirements (use goodRunsList to check)
goodRunMedium	bool	Run passes goodRunMedium requirements (see note above)
goodRunTight	bool	Run passes goodRunTight requirements (see note above)
goodSingleTrigger	bool	Single-trigger run passing goodSingleTrigger requirements (see note above)
triggerThreshold	vector<float>	Threshold in mV set for each trigger
triggerEnable	vector<bool>	Status of each trigger (0 = disabled, 1 = enabled)
triggerMajority	vector<int>	Majority level of digitizer (how many groups must fire to issue trigger) (don't use)
triggerPolarity	vector<int>	Trigger firing on rising (0) or falling (1) edge
triggerLogic	vector<int>	Trigger logic of groups (0 = OR, 1 = AND, 2 = MAJORITY)
dynamicPedestal	vector<float>	Dynamic voltage correction applied to pedestal (mV)
sidebandMean	vector<float>	Mean of sideband (sideband defined 0–50)
sidebandRMS	vector<float>	RMS of sideband
sidebandMeanRaw	vector<float>	Mean of sideband before dynamic pedestal correction (sideband defined 0–50)
sidebandRMSRaw	vector<float>	RMS of sideband before dynamic pedestal correction
maxThreeConsecutive	vector<float>	Max 3 bins in a row of a pulse
chan	vector<int>	Digitizer channel associated with this index (map)
chanWithinBoard	vector<int>	Digitizer channel within one board (0–15)
row	vector<int>	Row of bar
column	vector<int>	Column of bar
layer	vector<int>	Layer of bar
type	vector<int>	Type of channel (0 = bar, 1 = slab, 2 = panel)
board	vector<int>	Digitizer board channel is connected to
height	vector<float>	Height of pulse (mV)
area	vector<float>	Area of pulse (pV s)
pickupFlag	vector<bool>	Is pulse flagged as pickup
pickupFlagTight	vector<bool>	Tighter pickup criteria
nPE	vector<float>	Number of photoelectrons pulse is equal to
riseSamples	vector<int>	Number of bins between 20% and 80% maximum before peak
fallSamples	vector<int>	Number of bins between 20% and 80% maximum after peak
ipulse	vector<int>	Index of pulses in event per channel
npulses	vector<int>	Number of pulses in event per channel

Continued on next page

Name	Type	Description
pulseIndex	vector<int>	Global pulse index of pulses in event
time	vector<float>	Time of pulse within 2.56 ms window (μ s)
timeFit	vector<float>	Time of fitted pulse
time_module_calibrated	vector<float>	Time of pulse after timing calibration
timeFit_module_calibrated	vector<float>	Time of fitted pulse after timing calibration
duration	vector<float>	Width of pulse
delay	vector<float>	Time since end of previous pulse (9999 if no previous pulse in channel)
max	vector<float>	Max pulse height
iMaxPulseLayer	vector<int>	Index of max pulse in each layer per event
maxPulseTime	vector<float>	Time of max pulse in each channel per event
present	vector<Long64_t>	Is there data present in digitizers
event_trigger_time_tag	vector<Long64_t>	Trigger Time Tag from CAEN digitizers (ns)
event_time	vector<Long64_t>	Time since standard Unix epoch (s)
event_time_fromTDC	double	TDC time from digitizers (per group) (Time-to-Digital Converter)
v_groupTDC_gN	vector<Long64_t>	TDC time from digitizers (per group), should all be the same (saved as <code>event_time_fromTDC</code>)
totalFileLumi	float	Total luminosity in file (pb^{-1})
tClockCycles	float	Number of trigger board clock cycles
tTime	float	Trigger board time (clock cycles $\times$ 8 ns)
tStartTime	float	Start time of trigger board
tTrigger	float	Trigger bit string; see list of triggers in TWiki
tTimeDiff	float	Time difference between trigger and DAQ matching
tMatchingTimeCut	float	Allowed time difference between matched trigger and DAQ events
tEvtNum	int	Event number in <code>MatchedEvents</code> files
tRunNum	int	Run number in <code>MatchedEvents</code> files
tTBEvent	int	Event number in <code>TriggerBoard</code> files

A.12 Full MilliQan Result Cutflow

A.13 MilliQan Systematics

A.13.1 Timing Systematic

The results of applying the timing systematic from Section 7.8 are given below:

Selection Criteria	Signal Region 1			Signal Region 2		
	Data	Signal MC	Signal	Data	Signal	Signal
	Beam-On t=3393h	m=0.1 GeV Q/e=0.004	m=1.0 GeV Q/e=0.008	Beam-On t=3393h	m=1.7 GeV Q/e=0.03	m=10.0 GeV Q/e=0.2
Triggered Events	26864552	324.0	61.3	26864552	27.0	37.2
Skim Events	790776	324.0	61.3	790776	27.0	37.2
Digitizer Boards Matched	790772	324.0	61.3	790772	27.0	37.2
Pickup Cut	790772	324.0	61.3	790772	27.0	37.2
Dark Rate Cut	790772	324.0	61.3	790772	27.0	37.2
≤ 6 Bars	506770	324.0	61.3	—	—	—
First Pulse	506770	324.0	61.3	790772	27.0	37.2
Trigger Window	506417	323.9	61.3	790383	27.0	37.2
Top/Side Panel Veto	287811	266.9	38.9	342109	16.6	22.7
4 Layers	3369	12.0	19.3	9152	7.7	9.5
SR1 : ≤ 4 Bars	985	11.7	19.3	—	—	—
Noise Filter	985	11.7	19.3	9113	7.7	9.5
Energy Max/Min	336	10.3	16.5	1827	7.6	9.5
SR1 : Beam Muon Veto	331	10.3	16.5	—	—	—
SR1 : End Panel Veto	209	10.1	14.3	—	—	—
StraightLine	3	9.2	14.3	1372	7.5	9.4
$\Delta T_{\text{max-min}} \leq 20$ ns	0	8.7	14.1	1355	7.5	8.6
SR2 : End Panel Required	—	—	—	1320	5.8	8.2
SR2 : ≤ 4 Bars	—	—	—	84	5.8	7.3
SR2 : Max(Front/Back Panel nPE) < 70	—	—	—	2	5.8	7.0

Table A.12: Cutflows for SR1 and SR2 as well as several signal MC points, with all cuts expanded.

Mass (GeV)	Charge (Q/e)	Events Nominal	Events Up	Events Down	Sys Up (%)	Sys Down (%)
1.3	0.009	3.178	3.162	3.193	0.5	-0.5
1.7	0.01	0.126	0.124	0.128	1.6	-1.6
1.7	0.03	5.767	5.742	5.792	0.4	-0.4
3.0	0.03	1.809	1.807	1.815	0.1	-0.3
6.0	0.1	2.989	2.974	2.969	0.5	0.7
10.0	0.1	2.080	2.077	2.080	0.1	0.1
10.0	0.2	6.995	7.256	6.559	-3.7	6.2

Table A.13: Weighted event counts for masses/charges in SR2 with and without applying timing shifts.

Mass (GeV)	Charge (Q/e)	Events Nominal	Events Up	Events Down	Sys Up (%)	Sys Down (%)
0.1	0.004	8.682	8.567	8.896	1.3	-2.5
0.3	0.004	3.823	3.729	3.791	2.5	0.8
0.5	0.007	14.517	14.356	14.492	1.1	0.2
1.0	0.005	2.411	2.378	2.432	1.4	-0.9
1.0	0.008	14.134	14.116	14.154	0.1	-0.1
1.3	0.005	2.122	2.062	2.132	2.8	-0.5
1.7	0.02	1.471	1.469	1.471	0.1	0.0

Table A.14: Weighted event counts for masses/charges in SR1 with and without applying timing shifts.

Timing distributions used to derive these systematics are given by:




Figure A.18: Timing calibrations for all channels in layer 1.




Figure A.19: Timing calibrations for all channels in layer 2.




Figure A.20: Timing calibrations for all channels in layer 3.




Figure A.21: Timing calibrations for all channels in layer 4.

Mass (GeV)	Charge (Q/e)	Events Nominal	Events Up	Events Down	Sys Up (%)	Sys Down (%)
0.1	0.004	8.682	8.682	8.682	0.0	0.0
0.3	0.004	3.823	3.823	3.823	0.0	0.0
0.5	0.007	14.517	14.517	14.517	0.0	0.0
1.0	0.005	2.411	2.411	2.411	0.0	0.0
1.0	0.008	14.134	14.134	14.134	0.0	0.0
1.3	0.005	2.122	2.122	2.122	0.0	0.0
1.7	0.02	1.471	1.471	1.471	0.0	0.0

Table A.17: Weighted event counts for masses/charges in SR1 with and without applying PMT response corrections to nPE/energy for the Energy < 1000 keV selection.

A.13.2 Energy Systematic

The results of applying the energy systematic from Section 7.8 to both the bar and panel channels are given below:

Mass (GeV)	Charge (Q/e)	Events Nominal	Events Up	Events Down	Sys Up (%)	Sys Down (%)
1.3	0.009	3.178	3.114	3.152	2.0	0.8
1.7	0.01	0.126	0.116	0.121	7.9	4.0
1.7	0.03	5.767	5.760	5.362	0.1	7.0
3.0	0.03	1.809	1.808	1.708	0.1	5.6
6.0	0.1	2.989	2.989	2.788	0.0	6.7
10.0	0.1	2.080	2.080	1.944	0.0	6.5
10.0	0.2	6.995	6.995	6.527	0.0	6.7

Table A.15: Weighted event counts for masses/charges in SR2 with and without applying PMT response corrections to nPE/energy for the energy max/min selection.

Mass (GeV)	Charge (Q/e)	Events Nominal	Events Up	Events Down	Sys Up (%)	Sys Down (%)
1.3	0.009	3.178	3.178	3.178	0.0	0.0
1.7	0.01	0.126	0.126	0.126	0.0	0.0
1.7	0.03	5.767	5.765	5.767	0.0	0.0
3.0	0.03	1.809	1.806	1.810	0.2	-0.1
6.0	0.1	2.989	2.969	2.993	0.7	-0.1
10.0	0.1	2.080	2.075	2.086	0.2	-0.3
10.0	0.2	6.995	6.581	7.212	5.9	-3.1

Table A.16: Weighted event counts for masses/charges in SR2 with and without applying PMT response corrections to NPE/energy for the front/back panel NPE < 70 selection.

Energy calibrations are extracted from bar channels using the Cd109 source fits given below:

Energy calibrations were extracted from front and back endcap panels using a beam muon sample, given below:

Mass (GeV)	Charge (Q/e)	Events Nominal	Events Up	Events Down	Sys Up (%)	Sys Down (%)
0.1	0.004	8.682	8.682	8.634	0.0	0.6
0.3	0.004	3.823	3.761	3.604	1.6	5.7
0.5	0.007	14.517	14.214	13.946	2.1	3.9
1.0	0.005	2.411	2.357	2.330	2.2	3.4
1.0	0.008	14.134	13.518	13.890	4.4	1.7
1.3	0.005	2.122	2.049	2.041	3.4	3.8
1.7	0.02	1.471	1.444	1.365	1.8	7.2

Table A.18: Weighted event counts for masses/charges in SR1 with and without applying PMT response corrections to nPE/energy for the energy max/min selection.




Figure A.22: Cd 109 source fits for channels in layer 1.




Figure A.23: Cd 109 source fits for channels in layer 2.




Figure A.24: Cd 109 source fits for channels in layer 3.




Figure A.25: Cd 109 source fits for channels in layer 4.




Figure A.26: The width of the gaussian fit to the front (left) and back (right) panels is used to derive an energy uncertainty. Beam muons are used for the endcap panels in place of Cd109, as no dedicated sim-data comparison was made to extract the energy response of veto panels.

Bibliography

- [1] *Standard Model of Elementary Particles (SVG)*. Wikimedia Commons. 2019. URL: https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg.
- [2] D. P. A. et al. *Measurement of the Positive Muon Anomalous Magnetic Moment to 127 ppb*. 2025. arXiv: 2506.03069 [hep-ex]. URL: <https://arxiv.org/abs/2506.03069>.
- [3] R. A. et al. *The anomalous magnetic moment of the muon in the Standard Model: an update*. 2025. arXiv: 2505.21476 [hep-ph]. URL: <https://arxiv.org/abs/2505.21476>.
- [4] M. J. Jee et al. “Discovery of a Ringlike Dark Matter Structure in the Core of the Galaxy Cluster Cl 0024+17”. In: *The Astrophysical Journal* 661.2 (June 2007), p. 728. DOI: 10.1086/517498. URL: <https://dx.doi.org/10.1086/517498>.
- [5] NASA/ESA; M. J. Jee and H. Ford (JHU). *Dark matter ring in galaxy cluster Cl 0024+17 (ZwCl 0024+1652)*. ESA/Hubble image heic0709a. Credit: NASA, ESA, M. J. Jee and H. Ford (JHU); Release date: 15 May 2007. 2007. URL: <https://esahubble.org/images/heic0709a/>.
- [6] L. Lee, C. Ohm, A. Soffer, and T.-T. Yu. “Collider searches for long-lived particles beyond the Standard Model”. In: *Progress in Particle and Nuclear Physics* 106 (2019), pp. 210–255. ISSN: 0146-6410. DOI: <https://doi.org/10.1016/j.pnpnp.2019.02.006>. URL: <https://www.sciencedirect.com/science/article/pii/S0146641019300ndidate109>.
- [7] S. Knapen and S. Lowette. “A Guide to Hunting Long-Lived Particles at the LHC”. In: *Ann. Rev. Nucl. Part. Sci.* 73 (2023), pp. 421–449. DOI: 10.1146/annurev-nucl-101920-013011. arXiv: 2212.03883 [hep-ph].
- [8] H. Russell. *An experimental introduction to long-lived particle searches at the LHC*. 2017. URL: https://indico.cern.ch/event/607314/contributions/2542309/attachments/1447873/2231444/20170424_LLPs.pdf.
- [9] J. Liu, Z. Liu, and L.-T. Wang. “Enhancing Long-Lived Particles Searches at the LHC with Precision Timing Information”. In: *Phys. Rev. Lett.* 122 (13 Apr. 2019), p. 131801. DOI: 10.1103/PhysRevLett.122.131801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.122.131801>.
- [10] “Search for long-lived particles using out-of-time trackless jets in proton-proton collisions at $\sqrt{s} = 13\text{TeV}$ ”. In: *Journal of High Energy Physics* 2023.7 (July 2023). ISSN: 1029-8479. DOI: 10.1007/jhep07(2023)210. URL: [http://dx.doi.org/10.1007/JHEP07\(2023\)210](http://dx.doi.org/10.1007/JHEP07(2023)210).
- [11] “Search for long-lived particles decaying into displaced jets in proton-proton collisions at $\sqrt{s} = 13\text{ TeV}$ ”. In: *Phys. Rev. D* 99 (3 Feb. 2019), p. 032011. DOI: 10.1103/PhysRevD.99.032011. URL: <https://link.aps.org/doi/10.1103/PhysRevD.99.032011>.

-
- [12] “Search for long-lived particles using nonprompt jets and missing transverse momentum with proton-proton collisions at $s=13\text{TeV}$ ”. In: *Physics Letters B* 797 (2019), p. 134876. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2019.134876>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269319305908>.
- [13] I. Masina and M. Quiros. “On the Veltman condition, the hierarchy problem, and high-scale supersymmetry”. In: *Physical Review D* 88.9 (Nov. 2013). ISSN: 1550-2368. DOI: 10.1103/physrevd.88.093003. URL: <http://dx.doi.org/10.1103/PhysRevD.88.093003>.
- [14] B. Batell, M. Low, E. T. Neil, and C. B. Verhaaren. *Review of Neutral Naturalness*. 2022. arXiv: 2203.05531 [hep-ph]. URL: <https://ndidatearxiv.org/abs/2203.05531>.
- [15] Z. Chacko, H.-S. Goh, and R. Harnik. “Natural Electroweak Breaking from a Mirror Symmetry”. In: *Physical Review Letters* 96.23 (June 2006). ISSN: 1079-7114. DOI: 10.1103/physrevlett.96.231802. URL: <http://dx.doi.org/10.1103/PhysRevLett.96.231802>.
- [16] Z. Chacko, C. Kilic, S. Najjari, and C. B. Verhaaren. “Collider signals of the Mirror Twin Higgs boson through the hypercharge portal”. In: *Phys. Rev. D* 100 (3 Aug. 2019), p. 035037. DOI: 10.1103/PhysRevD.100.035037. URL: <https://link.aps.org/doi/10.1103/PhysRevD.100.035037>.
- [17] N. Craig. “Naturalness: past, present, and future”. In: *Eur. Phys. J. C* 83.9 (2023), p. 825. DOI: 10.1140/epjc/s10052-023-11928-7. arXiv: 2205.05708 [hep-ph].
- [18] G. Bertone, D. Hooper, and J. Silk. “Particle dark matter: evidence, candidates and constraints”. In: *Physics Reports* 405.5–6 (Jan. 2005), pp. 279–390. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2004.08.031. URL: <http://dx.doi.org/10.1016/j.physrep.2004.08.031>.
- [19] G. Arcadi, D. Cabo-Almeida, M. Dutra, P. Ghosh, M. Lindner, Y. Mambrini, J. P. Neto, M. Pierre, S. Profumo, and F. S. Queiroz. “The waning of the WIMP: endgame?”. In: *The European Physical Journal C* 85.2 (Feb. 2025). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-024-13672-y. URL: <http://dx.doi.org/10.1140/epjc/s10052-024-13672-y>.
- [20] M. B. et al. *US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report*. 2017. arXiv: 1707.04591 [hep-ph]. URL: <https://arxiv.org/abs/1707.04591>.
- [21] In: *Physics Reports* 1115 (Apr. 2025), pp. 448–569. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2024.09.013. URL: <http://dx.doi.org/10.1016/j.physrep.2024.09.013>.
- [22] B. Holdom. “Two $U(1)$ ’s and charge shifts”. In: *Physics Letters B* 166.2 (1986), pp. 196–198. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(86\)91377-8](https://doi.org/10.1016/0370-2693(86)91377-8). URL: <https://www.sciencedirect.com/science/article/pii/0370269386913778>.
- [23] A. Berlin and D. Hooper. *High-Energy Neutrinos From Millicharged Dark Matter Annihilation in the Sun*. 2024. arXiv: 2407.04768 [hep-ph]. URL: <https://arxiv.org/abs/2407.04768>.

- [24] M. Baumgart, C. Cheung, J. T. Ruderman, L.-T. Wang, and I. Yavin. “Non-abelian dark sectors and their collider signatures”. In: *Journal of High Energy Physics* 2009.04 (Apr. 2009), pp. 014–014. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2009/04/014. URL: <http://dx.doi.org/10.1088/1126-6708/2009/04/014>.
- [25] T. Cohen, D. J. Phalen, A. Pierce, and K. M. Zurek. “Asymmetric dark matter from a GeV hidden sector”. In: *Physical Review D* 82.5 (Sept. 2010). ISSN: 1550-2368. DOI: 10.1103/physrevd.82.056001. URL: <http://dx.doi.org/10.1103/PhysRevD.82.056001>.
- [26] S. e. a. Chatrchyan. “Search for fractionally charged particles in pp collisions at $\sqrt{s} = 7$ TeV”. In: *Phys. Rev. D* 87 (9 May 2013), p. 092008. DOI: 10.1103/PhysRevD.87.092008. URL: <https://link.aps.org/doi/10.1103/PhysRevD.87.092008>.
- [27] A. e. a. Hayrapetyan. “Search for Fractionally Charged Particles in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. Lett.* 134 (13 Apr. 2025), p. 131802. DOI: 10.1103/PhysRevLett.134.131802. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.134.131802>.
- [28] A. A. Prinz et al. “Search for Millicharged Particles at SLAC”. In: *Physical Review Letters* 81.6 (Aug. 1998), pp. 1175–1178. ISSN: 1079-7114. DOI: 10.1103/physrevlett.81.1175. URL: <http://dx.doi.org/10.1103/PhysRevLett.81.1175>.
- [29] S. Koren and A. Martin. *Phenomenology of Fractionally Charged Particles: Two Reqs Are Better Than One*. 2025. arXiv: 2507.16900 [hep-ph]. URL: <https://arxiv.org/abs/2507.16900>.
- [30] G. Shiu, P. Soler, and F. Ye. “Millicharged Dark Matter in Quantum Gravity and String Theory”. In: *Phys. Rev. Lett.* 110 (24 June 2013), p. 241304. DOI: 10.1103/PhysRevLett.110.241304. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.110.241304>.
- [31] A. Del Rosso. “Aerial view of the LHC and the four major experiments”. In: (2017). General Photo. URL: <https://cds.cern.ch/record/2253966>.
- [32] E. Lopienska. “The CERN accelerator complex, layout in 2022”. In: (2022). General Photo. URL: <https://cds.cern.ch/record/2800984>.
- [33] E. Sargsyan et al. “Linac4 Source and Low Energy Experience and Challenges”. In: 2024. URL: <https://jacow.org/hb2023/papers/wea4i2.pdf>.
- [34] H. Bartosik and G. Rumolo. *Performance of the LHC injector chain after the upgrade and potential development*. 2022. arXiv: 2203.09202 [physics.acc-ph]. URL: <https://arxiv.org/abs/2203.09202>.
- [35] T. Sakuma. “Cutaway diagrams of CMS detector”. In: (2019). URL: <https://cds.cern.ch/record/2665537>.
- [36] T. T. G. of the CMS Collaboration. *The CMS Phase-1 Pixel Detector Upgrade*. 2020. arXiv: 2012.14304 [physics.ins-det]. URL: <https://arxiv.org/abs/2012.14304>.

-
- [37] “Development of the CMS detector for the CERN LHC Run 3”. In: *Journal of Instrumentation* 19.05 (May 2024), P05064. ISSN: 1748-0221. DOI: 10.1088/1748-0221/19/05/p05064. URL: <http://dx.doi.org/10.1088/1748-0221/19/05/P05064>.
- [38] “The CMS experiment at the CERN LHC”. In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08004. DOI: 10.1088/1748-0221/3/08/S08004. URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08004>.
- [39] A. e. a. Tumasyan. “Performance of the local reconstruction algorithms for the CMS hadron calorimeter with Run 2 data”. In: *Journal of Instrumentation* 18.11 (Nov. 2023), P11017. ISSN: 1748-0221. DOI: 10.1088/1748-0221/18/11/p11017. URL: <http://dx.doi.org/10.1088/1748-0221/18/11/P11017>.
- [40] G. Kopp. “Long-Lived Particle Triggering with the CMS Hadron Calorimeter”. In: Dec. 2024, p. 895. DOI: 10.22323/1.476.0895.
- [41] J. P. Ernenwein. *Overview of the CMS electromagnetic calorimeter*. Tech. rep. Geneva: CERN, 1999. DOI: 10.1016/S0920-5632(99)00544-7. URL: <https://cds.cern.ch/record/421977>.
- [42] V. Klyukhin. “Design and Description of the CMS Magnetic System Model”. In: *Symmetry* 13.6 (2021). 17 pages, 9 figures, 67 references, p. 1052. DOI: 10.3390/sym13061052. arXiv: 2201.10478. URL: <https://cds.cern.ch/record/2773274>.
- [43] A. e. a. Sirunyan. “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $s=13$ TeV”. In: *Journal of Instrumentation* 13.06 (June 2018), P06015. DOI: 10.1088/1748-0221/13/06/P06015. URL: <https://dx.doi.org/10.1088/1748-0221/13/06/P06015>.
- [44] Sandvine. “Global Internet Phenomena Report”. In: (2024). URL: https://www.applogicnetworks.com/hubfs/Sandvine_Redesign_2019/Downloads/2024/GIPR/GIPR%202024.pdf.
- [45] “Enriching the physics program of the CMS experiment via data scouting and data parking”. In: *Physics Reports* 1115 (2025). CMS physics results from the first decade of LHC data, pp. 678–772. ISSN: 0370-1573. DOI: <https://doi.org/10.1016/j.physrep.2024.09.006>. URL: <https://www.sciencedirect.com/science/article/pii/S0370157324003247>.
- [46] V. Khachatryan. “The CMS trigger system”. In: *Journal of Instrumentation* 12.01 (Jan. 2017), P01020. DOI: 10.1088/1748-0221/12/01/P01020. URL: <https://dx.doi.org/10.1088/1748-0221/12/01/P01020>.
- [47] Donato, Silvio. “CMS trigger performance”. In: *EPJ Web Conf.* 182 (2018), p. 02037. DOI: 10.1051/epjconf/201818202037. URL: <https://doi.org/10.1051/epjconf/201818202037>.

- [48] C. P. Dengra, J. Flix, and A. Sikora. “Content delivery network solutions for the CMS experiment: The evolution towards HL-LHC”. In: *Journal of Parallel and Distributed Computing* 197 (2025), p. 105014. ISSN: 0743-7315. DOI: <https://doi.org/10.1016/j.jpdc.2024.105014>. URL: <https://www.sciencedirect.com/science/article/pii/S0743731524001783>.
- [49] G. Petrucciani, A. Rizzi, and C. Vuosalo. “Mini-AOD: A New Analysis Data Format for CMS”. In: *Journal of Physics: Conference Series* 664.7 (Dec. 2015), p. 072052. ISSN: 1742-6596. DOI: [10.1088/1742-6596/664/7/072052](https://doi.org/10.1088/1742-6596/664/7/072052). URL: <http://dx.doi.org/10.1088/1742-6596/664/7/072052>.
- [50] “Search for long-lived neutral particles decaying to quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 8$ TeV”. In: *Phys. Rev. D* 91 (1 Jan. 2015), p. 012007. DOI: [10.1103/PhysRevD.91.012007](https://doi.org/10.1103/PhysRevD.91.012007). URL: <https://link.aps.org/doi/10.1103/PhysRevD.91.012007>.
- [51] “Search for long-lived particles in events with photons and missing energy in proton-proton collisions at $s=7\text{TeV}$ ”. In: *Physics Letters B* 722.4 (2013), pp. 273–294. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2013.04.027>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269313003067>.
- [52] “Performance of the CMS electromagnetic calorimeter in pp collisions at $s = 13$ TeV”. In: *Journal of Instrumentation* 19.09 (Sept. 2024), P09004. DOI: [10.1088/1748-0221/19/09/P09004](https://doi.org/10.1088/1748-0221/19/09/P09004). URL: <https://dx.doi.org/10.1088/1748-0221/19/09/P09004>.
- [53] M. T. Lucchini. *Evolution of the response of the CMS ECAL and possible design options for calorimetry at the High Luminosity LHC*. Tech. rep. CMS-CR-2013/401. Presented at IEEE NSS/MIC 2013. 2013. URL: https://cds.cern.ch/record/1631037/files/CR2013_401.pdf.
- [54] “Reconstruction of signal amplitudes in the CMS electromagnetic calorimeter in the presence of overlapping proton-proton interactions”. In: *Journal of Instrumentation* 15.10 (Oct. 2020), P10002. DOI: [10.1088/1748-0221/15/10/P10002](https://doi.org/10.1088/1748-0221/15/10/P10002). URL: <https://dx.doi.org/10.1088/1748-0221/15/10/P10002>.
- [55] T. e. a. Adams. “Beam test evaluation of electromagnetic calorimeter modules made from proton-damaged PbWO₄ crystals”. In: *Journal of Instrumentation* 11.04 (Apr. 2016), P04012. DOI: [10.1088/1748-0221/11/04/P04012](https://doi.org/10.1088/1748-0221/11/04/P04012). URL: <https://dx.doi.org/10.1088/1748-0221/11/04/P04012>.
- [56] W. H. Chiu, Z. Liu, M. Low, and L.-T. Wang. “Jet timing”. In: *Journal of High Energy Physics* 2022.1 (Jan. 2022). ISSN: 1029-8479. DOI: [10.1007/jhep01\(2022\)014](https://doi.org/10.1007/jhep01(2022)014). URL: [http://dx.doi.org/10.1007/JHEP01\(2022\)014](http://dx.doi.org/10.1007/JHEP01(2022)014).
- [57] “Dark sector searches with the CMS experiment”. In: *Physics Reports* 1115 (2025). CMS physics results from the first decade of LHC data, pp. 448–569. ISSN: 0370-1573. DOI: <https://doi.org/10.1016/j.physrep.2024.09.013>. URL: <https://www.sciencedirect.com/science/article/pii/S0370157324003363>.
- [58] M. Citron et al. *Input to the ESPPU 2026 update: Searching for millicharged particles with the FORMOSA experiment at the CERN LHC*. 2025. arXiv: 2504.12973 [hep-ex]. URL: <https://arxiv.org/abs/2504.12973>.

- [59] C. Campagnari et al. *Design and Mechanical Integration of Scintillation Modules for SUB-Millicharge Experiment (SUBMET)*. 2025. arXiv: 2507.19311 [physics.ins-det]. URL: <https://arxiv.org/abs/2507.19311>.
- [60] Y.-D. Tsai, I. Hwang, R. Schmitz, M. Citron, K. Gunthoti, J. Steenis, H. Jeong, H. Moon, J. H. Yoo, and M. X. Liu. *LANSCE-mQ: Dedicated search for milli/fractionally charged particles at LANL*. 2024. arXiv: 2407.07142 [hep-ph]. URL: <https://arxiv.org/abs/2407.07142>.
- [61] M. Koshimizu. “Fundamental processes and recent development of organic scintillators”. In: *Journal of Luminescence* 278 (2025), p. 121008. ISSN: 0022-2313. DOI: <https://doi.org/10.1016/j.jlumin.2024.121008>. URL: <https://www.sciencedirect.com/science/article/pii/S0022231324005726>.
- [62] J. R. Lakowicz. “Energy Transfer”. In: *Principles of Fluorescence Spectroscopy*. Boston, MA: Springer US, 1999, pp. 367–394. ISBN: 978-1-4757-3061-6. DOI: 10.1007/978-1-4757-3061-6_13. URL: https://doi.org/10.1007/978-1-4757-3061-6_13.
- [63] M. Tanabashi et al. “Review of Particle Physics”. In: *Phys. Rev. D* 98 (2018), p. 030001. DOI: 10.1103/PhysRevD.98.030001. URL: <https://pdg.lbl.gov/2018/reviews/>.
- [64] P. Lecoq. “Scintillation Detectors for Charged Particles and Photons”. In: *Particle Physics Reference Library: Volume 2: Detectors for Particles and Radiation*. Ed. by C. W. Fabjan and H. Schopper. Cham: Springer International Publishing, 2020, pp. 45–89. ISBN: 978-3-030-35318-6. DOI: 10.1007/978-3-030-35318-6_3. URL: https://doi.org/10.1007/978-3-030-35318-6_3.
- [65] “PHENIX calorimeter”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 499.2 (2003). The Relativistic Heavy Ion Collider Project: RHIC and its Detectors, pp. 521–536. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(02\)01954-X](https://doi.org/10.1016/S0168-9002(02)01954-X). URL: <https://www.sciencedirect.com/science/article/pii/S016890020201954X>.
- [66] W. Commons. *PhotoMultiplierTubeAndScintillator.svg* — *Wikimedia Commons*. 2025. URL: <https://commons.wikimedia.org/w/index.php?title=File:PhotoMultiplierTubeAndScintillator.svg&oldid=1073135538>.
- [67] Hamamatsu Photonics K.K. *R878 — Photomultiplier Tube (51 mm head-on, bialkali)*. Product page / datasheet. URL: https://www.hamamatsu.com/us/en/product/optical-sensors/pmt/pmt_tube-alone/head-on-type/R878.html.
- [68] Hamamatsu Photonics K.K. *Photomultiplier Tubes: Basics and Applications*. 4th ed. Hamamatsu City, Japan: Hamamatsu Photonics K.K., 2017. URL: https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/etd/PMT_handbook_v4E.pdf.
- [69] P. B. Coates. “The origins of afterpulses in photomultipliers”. In: *Journal of Physics D: Applied Physics* 6.10 (June 1973), p. 1159. DOI: 10.1088/0022-3727/6/10/301. URL: <https://dx.doi.org/10.1088/0022-3727/6/10/301>.

- [70] A. Ball et al. “Search for millicharged particles in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. D* 102 (3 Aug. 2020), p. 032002. DOI: 10.1103/PhysRevD.102.032002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.102.032002>.
- [71] R. Schmitz. “Bar Assembly Guide”. June 14, 2021. 2021. URL: https://docs.google.com/document/d/1W2YnOULDzqqBZt1dRrdpYagNk_V0H4P0gn9Wzymxq2o/.
- [72] 3M Company. *3M Specular Film DF2000MA — Product Bulletin DF2000*. Technical datasheet, Revision C. Jan. 2017. URL: <https://multimedia.3m.com/mws/media/9824490/3mtm-specular-film-df2000ma-technical-data-sheet.pdf>.
- [73] A. Ball et al. “Sensitivity to millicharged particles in future proton-proton collisions at the LHC with the milliQan detector”. In: *Physical Review D* 104.3 (Aug. 2021). ISSN: 2470-0029. DOI: 10.1103/physrevd.104.032002. URL: <http://dx.doi.org/10.1103/PhysRevD.104.032002>.
- [74] C. Campagnari, B. Marsh, F. Golf, and F. Setti. *Generation of milliCharged particles at the LHC*. Apr. 2023. URL: https://github.com/rtschmitz/milliq_mcggen/blob/master/docs/milliqanGeneration.pdf.
- [75] L. Bergström. “Some rare decays of light mesons”. In: *Physica Scripta* 48.1 (July 1993), p. 117. DOI: 10.1088/0031-8949/48/1/021. URL: <https://dx.doi.org/10.1088/0031-8949/48/1/021>.
- [76] “0 and meson production in proton-proton collisions at $\sqrt{s} = 8\text{TeV}$ ”. In: *The European Physical Journal C* 78.3 (Mar. 2018). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-018-5612-8. URL: <http://dx.doi.org/10.1140/epjc/s10052-018-5612-8>.
- [77] T. Sjöstrand, S. Mrenna, and P. Skands. “A brief introduction to PYTHIA 8.1”. In: *Computer Physics Communications* 178.11 (2008), pp. 852–867. ISSN: 0010-4655. DOI: <https://doi.org/10.1016/j.cpc.2008.01.036>. URL: <https://www.sciencedirect.com/science/article/pii/S0010465508000441>.
- [78] P. Skands, S. Carrazza, and J. Rojo. “Tuning PYTHIA 8.1: the Monash 2013 tune”. In: *The European Physical Journal C* 74.8 (Aug. 2014). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-014-3024-y. URL: <http://dx.doi.org/10.1140/epjc/s10052-014-3024-y>.
- [79] “Production of the $\rho(770)^0$ meson in pp and Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV”. In: *Phys. Rev. C* 99 (6 June 2019), p. 064901. DOI: 10.1103/PhysRevC.99.064901. URL: <https://link.aps.org/doi/10.1103/PhysRevC.99.064901>.
- [80] “Production of $\omega(782)$ in pp collisions at $\sqrt{s} = 7$ TeV”. In: (2018). URL: <https://cds.cern.ch/record/2316785>.
- [81] “The differential production cross section of the (1020) meson in $\sqrt{s} = 7$ TeV pp collisions measured with the ATLAS detector”. In: *The European Physical Journal C* 74.7 (July 2014). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-014-2895-2. URL: <http://dx.doi.org/10.1140/epjc/s10052-014-2895-2>.

- [82] S. A. Maier, S. R. Andrews, L. Martín-Moreno, and F. J. García-Vidal. “Terahertz Surface Plasmon-Polariton Propagation and Focusing on Periodically Corrugated Metal Wires”. In: *Physical Review Letters* 97.17 (Oct. 2006). ISSN: 1079-7114. DOI: 10.1103/physrevlett.97.176805. URL: <http://dx.doi.org/10.1103/PhysRevLett.97.176805>.
- [83] Y.-Q. Ma, K. Wang, and K.-T. Chao. “ $J/\psi(\psi')$ Production at the Tevatron and LHC at $\mathcal{O}(\alpha_s^4 v^4)$ in Nonrelativistic QCD”. In: *Phys. Rev. Lett.* 106 (4 Jan. 2011), p. 042002. DOI: 10.1103/PhysRevLett.106.042002. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.106.042002>.
- [84] Y.-Q. Ma and R. Venugopalan. “Comprehensive Description of J/ψ Production in Proton-Proton Collisions at Collider Energies”. In: *Phys. Rev. Lett.* 113 (19 Nov. 2014), p. 192301. DOI: 10.1103/PhysRevLett.113.192301. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.113.192301>.
- [85] Y.-Q. Ma, K. Wang, and K.-T. Chao. “Complete next-to-leading order calculation of the J/ψ and ψ' production at hadron colliders”. In: *Phys. Rev. D* 84 (11 Dec. 2011), p. 114001. DOI: 10.1103/PhysRevD.84.114001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.84.114001>.
- [86] M. Cacciari, S. Frixione, N. Houdeau, M. L. Mangano, P. Nason, and G. Ridolfi. “Theoretical predictions for charm and bottom production at the LHC”. In: *Journal of High Energy Physics* 2012.10 (Oct. 2012). ISSN: 1029-8479. DOI: 10.1007/jhep10(2012)137. URL: [http://dx.doi.org/10.1007/JHEP10\(2012\)137](http://dx.doi.org/10.1007/JHEP10(2012)137).
- [87] M. Cacciari, M. L. Mangano, and P. Nason. *Gluon PDF constraints from the ratio of forward heavy quark production at the LHC at $\sqrt{S} = 7$ and 13 TeV*. 2015. arXiv: 1507.06197 [hep-ph]. URL: <https://arxiv.org/abs/1507.06197>.
- [88] “Measurement of upsilon production in 7 TeV pp collisions at ATLAS”. In: *Phys. Rev. D* 87 (5 Mar. 2013), p. 052004. DOI: 10.1103/PhysRevD.87.052004. URL: <https://link.aps.org/doi/10.1103/PhysRevD.87.052004>.
- [89] “Measurement of quarkonium production cross sections in pp collisions at $s=13$ TeV”. In: *Physics Letters B* 780 (2018), pp. 251–272. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2018.02.033>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269318301357>.
- [90] “Measurement of Υ production in pp collisions at $\sqrt{s} = 13$ TeV”. In: *Journal of High Energy Physics* 2018.7 (July 2018). ISSN: 1029-8479. DOI: 10.1007/jhep07(2018)134. URL: [http://dx.doi.org/10.1007/JHEP07\(2018\)134](http://dx.doi.org/10.1007/JHEP07(2018)134).
- [91] 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”. In: *Journal of High Energy Physics* 2014.7 (July 2014). ISSN: 1029-8479. DOI: 10.1007/jhep07(2014)079. URL: [http://dx.doi.org/10.1007/JHEP07\(2014\)079](http://dx.doi.org/10.1007/JHEP07(2014)079).

- [92] “Measurement of the differential inclusive B+ hadron cross sections in pp collisions at s=13 TeV”. In: *Physics Letters B* 771 (2017), pp. 435–456. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2017.05.074>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269317304379>.
- [93] C. Campagnari, B. Marsh, and F. Golf. *Propagation of simulated millicharged particles to the milliQan detector*. Sept. 2019. URL: https://github.com/rtschmitz/milliq_mcgen/blob/master/docs/milliqanPropagation.pdf.
- [94] D. Tan and Z. Chen. “On A General Formula of Fourth Order Runge-Kutta Method”. In: *Journal of Mathematical Sciences Mathematics Education* 7.2 (2012), pp. 1–10. URL: <https://msme.us/2012-2-1.pdf>.
- [95] M. Tanabashi et al. “Review of Particle Physics”. In: *Phys. Rev. D* 98 (2018), p. 030001. DOI: 10.1103/PhysRevD.98.030001. URL: <https://pdg.lbl.gov/2018/reviews/>.
- [96] G. R. Lynch and O. I. Dahl. “Approximations to multiple Coulomb scattering”. In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 58.1 (1991), pp. 6–10. ISSN: 0168-583X. DOI: [https://doi.org/10.1016/0168-583X\(91\)95671-Y](https://doi.org/10.1016/0168-583X(91)95671-Y). URL: <https://www.sciencedirect.com/science/article/pii/0168583X9195671Y>.
- [97] V. Klyukhin. “Design and Description of the CMS Magnetic System Model”. In: *Symmetry* 13.6 (2021). ISSN: 2073-8994. DOI: 10.3390/sym13061052. URL: <https://www.mdpi.com/2073-8994/13/6/1052>.
- [98] A. Ball et al. *A Letter of Intent to Install a milli-charged Particle Detector at LHC P5*. 2016. arXiv: 1607.04669 [physics.ins-det]. URL: <https://arxiv.org/abs/1607.04669>.
- [99] E. J. Fern, V. Di Murro, K. Soga, Z. Li, L. Scibile, and J. A. Osborne. “Geotechnical characterisation of a weak sedimentary rock mass at CERN, Geneva”. In: *Tunnelling and Underground Space Technology* 77 (2018), pp. 249–260. ISSN: 0886-7798. DOI: <https://doi.org/10.1016/j.tust.2018.04.003>. URL: <https://www.sciencedirect.com/science/article/pii/S0886779817310349>.
- [100] T. Basaglia, Z. W. Bell, D. D’Agostino, P. V. Dressendorfer, S. Giani, M. G. Pia, and P. Saracco. *Geant4: a Game Changer in High Energy Physics and Related Applicative Fields*. 2024. arXiv: 2405.12159 [physics.comp-ph]. URL: <https://arxiv.org/abs/2405.12159>.
- [101] *Medical Applications*. Geant4 website. URL: https://geant4-internal.web.cern.ch/applications/medical_applications.
- [102] *Space Applications*. Geant4 website. URL: https://geant4-internal.web.cern.ch/applications/space_applications.
- [103] S. A. et al. “Geant4—a simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8). URL: <https://www.sciencedirect.com/science/article/pii/S0168900203013688>.

-
- [104] *Physics Processes — Book For Application Developers: Tracking and Physics*. Geant4 Handbook / User Documentation. URL: https://geant4.web.cern.ch/documentation/pipelines/master/bfad_html/ForApplicationDevelopers/TrackingAndPhysics/physicsProcess.html.
- [105] *Geant4 General Particle Source (GPS) — User Manual*. User manual / tutorial PDF. URL: https://www.fe.infn.it/u/paterno/Geant4_tutorial/slides_further/GPS/GPS_manual.pdf.
- [106] SLAC National Accelerator Laboratory. *Shielding Physics List Description*. SLAC Geant4 Physics Lists. 2012. URL: https://www.slac.stanford.edu/comp/physics/geant4/slac_physics_lists/shielding/physlistdoc.html.
- [107] R. Mirzoyan, M. Laatiaoui, and M. Teshima. “Very high quantum efficiency PMTs with bialkali photo-cathode”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 567 (Nov. 2006), pp. 230–232. DOI: 10.1016/j.nima.2006.05.094.
- [108] Eljen Technology. *EJ-200, EJ-204, EJ-208, EJ-212 — Plastic Scintillators (General Purpose)*. Product datasheet. URL: <https://eljentechnology.com/products/plastic-scintillators/ej-200-ej-204-ej-208-ej-212>.
- [109] M. Janecek and W. Moses. “Optical Reflectance Measurements for Commonly Used Reflectors”. In: *Nuclear Science, IEEE Transactions on* 55 (Sept. 2008), pp. 2432–2437. DOI: 10.1109/TNS.2008.2001408.
- [110] M. Janecek. “Reflectivity Spectra for Commonly Used Reflectors”. In: *IEEE Transactions on Nuclear Science* 59.3 (2012), pp. 490–497. DOI: 10.1109/TNS.2012.2183385.
- [111] LNE-LNHB/CEA. *Table of Radionuclides: Cd-109 (tables)*. LNHB online database. URL: https://www.lnhb.fr/nuclides/Cd-109_tables.pdf.
- [112] R. Schmitz. “SR2 Pathological Beam Muon Event Display”. URL: <https://drive.google.com/file/d/1EHikQ1GLa5QS3yrYaFXxNHbLmeDfyDFI>.
- [113] C. Anastasiou, L. Dixon, K. Melnikov, and F. Petriello. “High-precision QCD at hadron colliders: Electroweak gauge boson rapidity distributions at next-to-next-to leading order”. In: *Physical Review D* 69.9 (May 2004). ISSN: 1550-2368. DOI: 10.1103/physrevd.69.094008. URL: <http://dx.doi.org/10.1103/PhysRevD.69.094008>.
- [114] “Time reconstruction and performance of the CMS electromagnetic calorimeter”. In: *Journal of Instrumentation* 5.03 (Mar. 2010), T03011. DOI: 10.1088/1748-0221/5/03/T03011. URL: <https://dx.doi.org/10.1088/1748-0221/5/03/T03011>.
- [115] P. Papacz. “Optimisation of the Particle Detection Efficiency for Scintillation Detectors with SiPM Readout”. MA thesis. RWTH Aachen U., Mar. 2010.