
Clustering and Spike Killing Algorithms for CMS Electromagnetic Calorimeter Phase II Upgrade and Search for Non-Resonant New Physics in High Mass Dilepton Events in Association with b-tagged Jets

By

CHARLOTTE COOKE



School of Physics
UNIVERSITY OF BRISTOL

A dissertation submitted to the University of Bristol in accordance with the requirements of the degree of DOCTOR OF PHILOSOPHY in the Faculty of Science, School of Physics.

JULY 2023

Word count: 40,722



Abstract

This thesis details work completed in the search for physics beyond the Standard Model. The first half of the thesis focuses on the upcoming Phase II upgrade for the crystal electromagnetic calorimeter of the Compact Muon Solenoid. Test beam data was analysed using available Phase II prototype electronics and this study found that, in the absence of a magnetic field, a cluster of a 3×3 crystal matrix provides better energy resolution than a single crystal or a 3×5 crystal matrix. An energy correction method to recover energy when an electron impacts on the side of a crystal was also tested and found to return the resolution to close to that for electrons impacting on the centre of the crystal. Algorithms that will cluster energy and reject anomalous signals for the Phase II upgrade were also tested in software and on hardware using an existing board that was designed for the L1 trigger. The clustering algorithm was tested on 1000 events in hardware with no discrepancies, and both anomalous signal rejection algorithms were found to have the required 99.9% efficiency from 2 or 10 GeV, depending on the algorithm.

The second half of this thesis focuses on the ongoing search for a non-resonant Z' boson decaying to a dilepton pair in association with b-tagged jets, using 137 fb^{-1} of data collected by the Compact Muon Solenoid experiment, as well as a test of lepton flavour universality. Two models are considered for this search, both based on contact interactions. One considers the interaction between two b quarks and two leptons, where the b quarks can be in the initial and/or final state, and the second model considers the transition $b \rightarrow s\ell\ell$. The dilepton mass distribution of both of these models depends on the energy scale of the new physics.

The dielectron part of this search requires the use of a dedicated high-energy electron/positron identification criteria which was studied in detail to ensure that its behaviour was fully understood. For 2018, the efficiency of this dedicated criteria was found to be $86.94 \pm 0.05\%$ in data, compared to an efficiency of $89.32 \pm 0.07\%$ in the CMS detector simulation. The scale factor of $0.973 \pm 0.001 \text{ (stat)} \pm 0.004 \text{ (sys)}$ gives the ratio between the data and simulation and is used as a correction in the Z' boson search. The flavour ratio is measured which compares the differential cross sections as a function of dilepton mass for a quark-antiquark pair decaying to a dimuon or a dielectron pair. For a data sample with no b-tagged jets, the flavour ratio is in agreement with unity, which is consistent with lepton flavour universality as expected. The predicted lepton flavour ratio for a sample with 1 b-tagged jet shows that for some energy scales of new physics, there will be a deviation from unity at high mass. Once finished by the Z' group, this analysis is expected to probe higher energy scale ranges than previous analyses. The results are expected to either provide hints of physics beyond the Standard Model or allow more stringent limits on new physics to be set.

Dedication and acknowledgements

This endeavour would not have been possible without my supervisors Helen Heath and David Petyt who have provided unwavering support, advice and feedback throughout my PhD. Your knowledge and expertise have been immensely helpful to me, both with our weekly meetings and being able to knock on an office door for a chat whenever I had questions or was facing a problem. I couldn't have asked for better supervisors.

I am extremely grateful to Sam Harper for all of his help with the HEEP ID. I am deeply indebted to Thomas Reis, who has been instrumental in the ECAL trigger primitive algorithm development. Thanks should also go to the CMS Z' group who I had the pleasure of working with. I wish them the best of luck in completing our analysis.

Many thanks to the Bristol HEP group for being such fun to work with and for creating such a social and welcoming atmosphere. I would especially like to thank Rob, Mei, Katie, Rui, Ben, Chris, Jozie and Florian. I'll miss the Friday lunches and pub trips. A special thank you to Martisse who has performed magic to solve all sorts of problems over the years.

I would also like to thank my mum and my nana for their belief in me and for always being at the other end of the phone when I needed to talk. You both have always held me up when I needed support and no thanks will ever be enough. I'd also like to acknowledge Michelle, Allan and Dominic (MAD) for all their love and support. Thanks should also go to my dad, for providing a sunny place to decompress when I needed a break. I would like to extend my sincere thanks to Chloe for being there when I could not and always keeping me in the loop.

I would be remiss in not mentioning Abi, who was there for me when the world was crumbling and has been an amazing flatmate.

I am also grateful for the LTA crew who made my 18 months in Switzerland a time I will never forget. I would like to thank Clara and Will for being such great flatmates. I promise not to make Clara watch any more movies about racoon dogs! Also, I would like to thank Sammy, Gareth, Zak, Luke and Paul for all the wonderful evenings spent in one flat or another playing board games or watching films.

Lastly, I would like to thank my cat Yuki for her endless entertainment and emotional support.

Author's declaration

I declare that the work in this dissertation was carried out in accordance with the requirements of the University's Regulations and Code of Practice for Research Degree Programmes and that it has not been submitted for any other academic award. Except where indicated by specific reference in the text, the work is the candidate's own work. Work done in collaboration with, or with the assistance of, others, is indicated as such. Any views expressed in the dissertation are those of the author.

SIGNED: DATE:

Table of Contents

	Page
List of Tables	xi
List of Figures	xiii
1 Introduction	1
2 Theory	3
2.1 The Standard Model	3
2.1.1 Particles in the Standard Model	4
2.1.2 Quantum Electrodynamics	5
2.1.3 Quantum Chromodynamics	6
2.1.4 Electroweak Theory	6
2.1.5 Shortcomings in the Standard Model	9
2.2 Beyond the Standard Model	10
2.2.1 Supersymmetry	10
2.2.2 Sequential Standard Model	11
2.2.3 Grand Unified Theories	11
2.2.4 Contact Interactions	11
2.2.5 Extra Dimensions	12
2.3 Proton-Proton Collisions	13
2.4 Drell-Yan Production	13
2.5 Simulation	15
3 Compact Muon Solenoid and the Large Hadron Collider	17
3.1 Introduction	17
3.2 The Large Hadron Collider	17
3.3 Compact Muon Solenoid	20
3.3.1 Magnetic Solenoid	22
3.3.2 Tracker	22

TABLE OF CONTENTS

3.3.3	Electromagnetic Calorimeter	24
3.3.4	Hadron Calorimeter	25
3.3.5	Muon System	26
3.3.6	Trigger and Data Acquisition	28
3.3.7	Phase II Upgrades	29
4	Electromagnetic Calorimeter	33
4.1	Introduction	33
4.2	Current ECAL	33
4.2.1	Current Electron Clustering	35
4.2.2	Current Rejection of Anomalous APD Signals	37
4.2.3	Current Electron Reconstruction	38
4.2.4	Performance of the Current ECAL	45
4.3	ECAL Upgrade for Phase II	46
4.3.1	Barrel Upgrade	46
4.3.2	Endcap Upgrade	58
5	Test Beam Analysis for ECAL Phase II Upgrade	63
5.1	June 2018 Test Beam Setup	63
5.2	June 2018 Test Beam Analysis	69
5.2.1	Methods to Extract the Amplitude	69
5.2.2	Hodoscope Cuts	70
5.2.3	Incorrect Pedestal values	71
5.2.4	Problems with the 50 GeV Run	71
5.2.5	Calibration Constants	75
5.2.6	Energy Containment Correction	76
5.3	Resolution Performance	78
5.4	Conclusions	83
6	Testing Phase II Clustering and Spike Killing Algorithms	85
6.1	High Level Synthesis	86
6.2	Previous Work on the Linear Discriminant Algorithm	89
6.3	Comparison of HLS and C++ Algorithms	91
6.3.1	Clustering Algorithm Tests	92
6.3.2	Spike Killing Algorithm Tests	94
6.4	Hardware Test of HLS Algorithms	95
6.5	Possible Algorithm Improvements	101
6.6	Conclusion	104
7	Previous Searches for New Phenomena in Dilepton Final States	107

7.1	Backgrounds, Event Simulation and Event Selection	108
7.2	Resonant Searches	109
7.2.1	Resonant Search in Association with b tagged jets	110
7.3	Non-Resonant Searches	112
7.3.1	Non-Resonant Search in Association with b tagged jets	115
7.4	Lepton Flavour Universality	116
8	HEEP ID Using Ultra Legacy Data	121
8.1	Ultra Legacy Samples	122
8.2	Data and MC Samples	123
8.3	Method	124
8.3.1	HEEP ID	124
8.3.2	Tag and Probe Method	126
8.4	QCD Estimate	129
8.5	Results	130
8.5.1	Kinematic Distributions	130
8.5.2	HEEP ID Efficiency	132
8.5.3	N-1 Efficiencies	135
8.5.4	Comparison of EOY and UL Samples	144
8.6	Conclusion	144
9	Search for Non-Resonant New Physics in High Mass Dilepton Events in Association with b-tagged Jets	149
9.1	Theoretical motivation	149
9.2	Data, Background MC and Signal MC Samples	152
9.3	Baseline Event Selection	152
9.4	Deep Neural Net Based Discriminator	156
9.5	MC Corrections	161
9.6	Acceptance and Efficiency	161
9.7	Uncertainties	163
9.7.1	Electron Uncertainties	165
9.7.2	Muon Uncertainties	165
9.8	Lepton Flavour Ratio	165
9.9	Results	168
9.10	Conclusion	170
10	Summary	173
A	HEEP Efficiencies for 2016 and 2017	175

TABLE OF CONTENTS

Bibliography	179
---------------------	------------

List of Tables

Table	Page
2.1 Quarks and leptons of the SM. The masses listed are taken from [5].	3
2.2 Gauge bosons of the SM. The masses listed are taken from [5].	4
4.1 Cluster word definition for the Phase II ECAL upgrade. Proposed values from TDR [73].	56
5.1 Noise, constant and stochastic terms for each crystal matrix size considered in this analysis.	79
6.1 Post implementation timing estimates for the second order LD algorithm for 300 channels on a KU115.	90
6.2 Post implementation latency estimates for the second order LD algorithm for 300 channels on a KU115.	90
6.3 Post implementation resource usage for the second order LD algorithm for 300 channels on a KU115.	91
6.4 Post implementation timing estimates for the combined Swiss-cross and clustering algorithm for 120 channels.	99
6.5 Post implementation latency for the combined Swiss-cross and clustering algorithm for 120 channels.	99
6.6 Post implementation resource usage for the combined Swiss-cross and clustering algorithm for 120 channels.	100
6.7 Example input event to combined Swiss-cross and clustering algorithm. Rows 1-5 contain test beam data, with rows 6-8 padded with zeros. To convert from counts to GeV, all values should be divided by five. For example, the energy of 793 counts deposited in C3 corresponds to 158.6 GeV.	101

6.8	Example events clustered by the combined Swiss-cross and clustering algorithm. The * denotes the primary cluster, with † denoting the secondary cluster and $\triangle$ denoting a Swiss-cross cluster where applicable. The first event corresponds to the input shown in Table 6.7. The output for the first event is a single cluster of 948 counts (corresponding to 189.6 GeV) centred on C3 with 9 contributing crystals and a spike flag of zero. The second event has two EM clusters. The first cluster has an energy of 875 counts centred on C4 with nine contributing crystals and a spike flag of zero, while the second cluster has an energy of 29 counts centred on B1 with six contributing crystals and a spike flag of zero. In the third event, the first EM cluster has an energy of 72 counts centred on C5 with six contributing crystals and a second EM cluster with an energy of 24 counts with seven contributing crystals and a spike flag of zero. There is also a Swiss-cross cluster in this event, with an energy of 1 count centred on D2 with one contributing crystal and the spike flag set to one. To convert from counts to GeV, all values should be divided by five.	102
8.1	All the MC samples used in this analysis, the generators used and their cross-sections.	124
8.2	All the selections included in HEEP ID V7.0	124
8.3	HEEP ID efficiencies and scale factors for 2016, 2017 and 2018. Calculated for both the barrel and the endcap by comparing data to MC and by subtracting the non-DY MC from data and comparing that to DY MC.	139
9.1	The physics processes and generators for the background MC considered in this analysis.	155
9.2	χ^2 values in the smeared and unfolded phase space for both the dielectron and dimuon channels	168

List of Figures

Figure	Page
2.1 Higgs potential as a function of the Higgs field average value for $\mu = 90$ GeV and $\lambda=0.13$. From [18].	8
2.2 Next-to-leading order PDFs at two different factorisation scales: $Q^2 = 10$ GeV ² (left) and $Q^2 = 10^4$ GeV ² (right). From [33].	14
2.3 Feynman diagrams that contribute to the DY process: exchange of a virtual photon (top) and exchange of a Z^0 boson (bottom).	14
2.4 Illustration of a proton-proton collision as simulated using an MC event generator. The red circle at the centre depicts the hard collision and the Bremsstrahlung from parton showers is shown by the tree-like structure surrounding the centre. A secondary hard scatter event is depicted by the purple oval. Light green ovals depict hadronisation, while dark green ovals represent hadron decays. The yellow lines indicate soft photon radiation. From [36].	16
3.1 The CERN accelerator complex and experiments in 2022. From [45].	18
3.2 Total integrated luminosity from 2011 to 2022 recorded by the CMS experiment (in yellow) and delivered by the LHC to the CMS experiment (in blue). The flat periods represent times when no collisions took place, often for detector and accelerator upgrade purposes. From [47].	20
3.3 The CMS detector. From [48].	21
3.4 The coordinate system used by the CMS experiment. From [50].	22
3.5 Cross section of the CMS conductor. From [53].	23
3.6 Diagram showing one quarter of the CMS Tracker during Phase I. The green represents the pixel detector, while the red represents single-sided strip modules and the blue represents double-sided strip modules. From [54].	23
3.7 Schematic of the ECAL. From [58].	25
3.8 A schematic showing one quarter of the HB and HE layers of the CMS HCAL during Phase I. The segmentation in the top left is for the detector during Run 1 and Run 2, while the bottom right shows the segmentation of the detector in Run 3. The HO and the HF are not shown in this diagram. From [62].	27

3.9	Illustration of one quadrant of the CMS muon system showing all three muon subsystems. DTs are shown in beige, while CSCs are shown in green and RPCs are shown in blue. From [64].	28
3.10	Timeline showing both LHC and HL-LHC operations. The Phase II upgrade for CMS will occur during Long Shutdown 3. From [68].	30
4.1	Current electronics of the CMS ECAL. From [73].	34
4.2	Schematic showing how the sFGVB works in the case of an EM shower (left) and a spike (right). From [85].	38
4.3	Fraction of ECAL trigger primitives in the barrel that are due to spikes as a function of E_T . Data was collected from a single run in 2018 with a peak pileup of 50. The black line shows the result when the online pedestal values are used, while the dashed line shows the result with updated pedestal values. This is required as the pedestal values drift with time. From [86].	39
4.4	Illustration defining the Swiss-cross variables E_1 and E_4	39
4.5	The distribution of the Swiss-cross variable for 2010 data and MC events which do not include spikes. The data with values above one for $1 - E_4/E_1$ are due to noise, where the reconstruction method used in 2010 could give negative reconstructed amplitudes for noise pulses. From [85].	40
4.6	Multifit method during Phase I based on CMS simulation for an event with 20 pileup and a bunch spacing of 25 ns. The black circles represent the digitised samples. The red histogram shows the in-time pulse, while the yellow, green and light blue histograms represent three out-of-time pulses. The dark blue histogram is the sum of both the in-time and out-of-time contributions. From [89].	41
4.7	Relative response to laser light of the CMS ECAL over 8 years of data taking, given for different η ranges. The instantaneous luminosity of the LHC is also shown. From [91].	43
4.8	Intercalibration precision for the $Z \rightarrow ee$ calibrations, E/p calibrations, π^0 calibrations and the three method combined. From [95].	44
4.9	Left: the impact of energy scale and resolution for the ECAL barrel from the use of superclustering and energy corrections compared to the sum of energy deposits in a 5×5 crystal matrix. Right: impact of energy scale and resolution for the ECAL endcap from the use of superclustering, addition of energy from the preshower and energy corrections compared to the sum of deposits in a 5×5 crystal matrix. Both show the reconstructed Z mass. From [97].	45
4.10	Electron reconstruction efficiency as a function of η , with the data/MC ratio shown in the bottom panel for data collected during 2017. The error bars denote combined statistical and systematic uncertainties. The efficiency is measured using tag and probe with $Z \rightarrow ee$ events. From [98].	46

4.11	Electron resolution for Run 2 data as a function of η , split into 2016 (red), 2017 (green) and 2018 (blue). The grey region denotes the gap between the barrel and the endcap. From [95].	47
4.12	Invariant mass spectrum of the $H \rightarrow \gamma\gamma$ decay channel. Each event is weighted by the signal divided by the sum of the signal and background. The solid red line indicates the signal and background fit of the data, while the red dashed line indicates the background fit. The yellow and green bands indicate the ± 1 and ± 2 standard deviation uncertainties in the background estimate respectively. From [2].	48
4.13	Expected event rate in the CMS ECAL during Phase II (HL-LHC) as a function of E_T . No spike killing is applied and 200 pileup is assumed. The event rate is shown for an integrated luminosity of 0 (undamaged), 300, 1000, 3000 and 4500 fb^{-1} , with the Phase I (LHC) and Phase II (HL-LHC) L1 bandwidths also shown. From [73].	49
4.14	Left: pulse shapes of spikes (black) and EM showers (grey) in the Phase I CMS ECAL. Right: pulse shapes of spikes (red) and EM showers (blue) with the Phase II CMS ECAL. The Phase II pulse shapes are given by convoluting known pulse shapes of spikes and EM signals with the measured response of prototype Phase II electronics. From [73].	49
4.15	Expected light output of the ECAL barrel crystals during Phase II relative to conditions at the start of CMS running. The relative light output is shown as a function of η for a range of integrated luminosities from 300/ fb in red (at the end of Phase I) to 4500/ fb in blue (expected at the end of Phase II). Based on 50 GeV photon showers using Phase II electronics cooled to 9 °C. From [73].	50
4.16	Expected energy resolution as a function of incident particle energy for a range of integrated luminosities from 300/ fb in red (at the end of Phase I) to 4500/ fb in blue (expected at the end of Phase II). Shown for $ \eta = 0$ (left) and $ \eta = 1.45$ (right). From [73].	51
4.17	Schematic showing the ECAL barrel Phase II electronics for a 5×5 crystal matrix. From [99].	51
4.18	Illustration of the Phase II multifit algorithm at L1. The red curve illustrates the pulse in the bunch crossing under consideration, while the green curve shows a pulse from the previous bunch crossing and the orange shows a pulse from the subsequent bunch crossing. The grey lines indicate the nominal timing of each bunch crossing. The combined waveform from all three pulses is shown in blue, with the 12 samples considered by the multifit algorithm shown using blue diamonds. Y-axis units are arbitrary.	53
4.19	Expected performance of the Swiss-cross algorithm for different E_T thresholds and four integrated luminosities: 300 fb^{-1} (top left), 1000 fb^{-1} (top right), 3000 fb^{-1} (bottom left) and 4500 fb^{-1} (bottom right). From [73].	54

4.20	Expected signal efficiency of the LD algorithm in minimum bias interactions at 200 PU. From [73].	55
4.21	Distribution of the LD variable for spikes (red) and EM showers (blue). The left plot shows the distributions for plots with amplitudes 10 times greater than the noise, while the right plot shows the distribution for amplitudes 50 times greater than the noise. From [73]. The double peak structure of the APD signal on the right plot is due to the LD value varying with the phase of the spike as the spike simulation has a realistic energy and spike arrival time distribution.	56
4.22	Expected event rate in the CMS ECAL during Phase II (HL-LHC) as a function of E_T . One spike-killing algorithm is applied and 200 pileup is assumed. The event rate is shown for an integrated luminosity of 0 (undamaged), 300, 1000, 3000 and 4500 fb ⁻¹ . Top: LD spike killing algorithm is used. Bottom: Swiss-cross spike killing algorithm is used. From [73].	57
4.23	Top: reconstruction efficiency for electrons where the electrons were generated with a uniform transverse momentum distribution. Bottom: reconstruction efficiency for a prompt photon from $H \rightarrow \gamma\gamma$ events. Both are shown for several barrel ageing conditions from Run II conditions up to 4500 /fb. From [73].	59
4.24	$H \rightarrow \gamma\gamma$ mass resolution with different assumptions on vertex efficiency. Green represents no precise timing, blue represents precise timing in the calorimeter only, and red represents precise timing in the calorimeter and MIP timing. Conditions in Run 2 are shown in black. From [72].	60
4.25	The simulated absorbed dose in the HGCAL after an integrated luminosity of 3000 /fb. From [104].	61
4.26	Simulated hits in the HGCAL with a charge greater than the threshold on the time of arrival measurement (12 fC). Left: with no timing requirement. Right: with the requirement $ \Delta t < 90$ ps. From [104].	61
5.1	June 2018 setup for the ECAL test beam at CERN.	65
5.2	Inside the box used for the 2018 test beam run. The crystal matrix can be seen inside the copper coloured box, while there are bare APDs (i.e. not mounted to a lead tungstate crystal) above this. Though the crystal matrix itself is 5x6, only five of the six crystals per column could be read out as the VFE is limited to five channels. Photo provided by Simone Pigazzini.	66
5.3	Box with crystal matrix inside in situ at the H4 experimental area. This box is placed on a moving table that allows the box to be moved up, down, left and right remotely to allow the beam to be centred on different crystals. Photo provided by Simone Pigazzini.	67
5.4	Example pulse shape for a 200 GeV electron beam centred on C3, the central crystal in the crystal matrix.	68

5.5	Matrix scheme of the lead tungstate crystals for the June 2018 test beam. The four dark grey crystals - A4, A5, E4 and E5 - did not collect data during this test beam run due to broken VFE adaptors. Each column refers to one VFE.	68
5.6	Histogram showing the amplitude given by each of the three amplitude reconstruction methods in the test beam specific data format for a 3×3 crystal matrix at 200 GeV.	70
5.7	Histogram showing the second hodoscope plane with cuts of 3 mm in both coordinates on the first plane, accounting for the 2 mm offset between their x positions. The beam energy used was 200 GeV and the beam was aimed at the centre of C3, though the axes have been adjusted in this instance to place the centre of C3 at (0,0).	71
5.8	Histogram showing the effect of hodoscope cuts of increasing tightness on the single crystal amplitude for a 200 GeV electron beam centred on crystal C3.	72
5.9	Pulse shape for an event during a 150 GeV run which has two pulses, one at around 0 ns and the second at around 300 ns.	72
5.10	Amplitude of energy deposited in C3 during the 50 GeV run, measured in ADC counts.	73
5.11	2D Histogram showing the hodoscope occupancy for the 50 GeV run.	74
5.12	Histogram showing the effect of applying the correct calibration constant to the crystal C2 for a 200 GeV run.	75
5.13	Second order polynomial fits to the positive (left) and negative (right) regions of $\ln(E2/E1)$	77
5.14	Definitions of energy sums E1 and E2 in the case where $C2 < C4$ (left) and $C2 > C4$ (right). The black circle represents the electron impact position.	77
5.15	Uncorrected amplitude as function of $\ln(E2/E1)$ at 50 GeV. The bright spot around $\ln(E2/E1)=2$ is due to an increased proportion of events aimed at the top of C3.	78
5.16	Corrected amplitude as a function of $\ln(E2/E1)$ at 50 GeV for the same data as Figure 5.15. The bright spot around $\ln(E2/E1)=2$ is due to an increased proportion of events aimed at the top of C3.	79
5.17	Histograms showing Crystal Ball fits for 100 GeV electrons impacting below the crystal centre with (top) and without (bottom) energy corrections.	80
5.18	Resolution for a single crystal with hodoscope cuts (top left), a 3×3 matrix with hodoscope cuts (top right), a 3×3 matrix without hodoscope cuts (middle left), a 3×3 matrix with $\ln(E2/E1)$ energy corrections applied (middle right), a 3×5 matrix without hodoscope cuts (bottom left) and a 3×5 matrix with hodoscope cuts (bottom right).	81

6.1	Workflow for Vivado HLS, which is represented by the grey box. The C/C++/SystemC code of the algorithm itself, along with a test bench file and constraints/directives are given as inputs. The main steps in the workflow are C simulation, C synthesis and register transfer level (RTL) simulation (or cosimulation), with the firmware exported at the end of these to a programme such as Vivado Design Suite which allows the firmware to be loaded onto an FPGA. VHDL or verilog are generated from the C synthesis, while the RTL adapter is needed to use the test bench in the cosimulation. steps. RTL is used in HDLs to create a high-level representation of a circuit. From [115].	88
6.2	The MP7 board used for the hardware test of the LD algorithm and the later test of the combined Swiss-cross and clustering algorithm. The attached cable was irrelevant to the test described in this thesis as data was loaded and read out via the crate's backplane.	91
6.3	Energy spectra for a 3x3 cluster when clustering is done via the C++ algorithm (top) or via a static cluster with full test beam data precision (bottom). For ease of comparison, the results from the C++ algorithm have been returned to the same scale. Both energy spectra have been fitted with a Crystal Ball function.	93
6.4	Plot showing the difference in E_T between the HLS and C++ algorithms, where there was a disagreement between them with regard to overflow.	94
6.5	Plot showing the difference in E_T (arbitrary units) between the HLS and C++ algorithms once the overflow issue had been resolved.	95
6.6	Example spike pulse shape from the 2018 pion test beam campaign (top), and example energy deposit in ADC counts for a spike in channel C3 of the crystal matrix from the 2021 pion test beam campaign (bottom). Both were provided by David Petyt.	96
6.7	Difference between the HLS and C++ crystal spike flag for the LD algorithms. . .	97
6.8	Illustration of a bitonic sort of 16 numbers. In the first stage (comparison 1) of the algorithm, pairs of data are sorted into the correct order. The second stage (comparisons 2 and 3) then accordingly sorts two adjacent pairs into four sets of four ordered numbers. The third stage (comparisons 4-6) creates two sets of eight ordered numbers, and finally, the fourth stage (comparisons 7-10) creates a single array of 16 ordered numbers. From [120].	98
6.9	Energy spectra for a 3x3 cluster when clustering is done using the updated clustering algorithm.	99
6.10	The payload for the MP7 for 4x8 channels. The combined Swiss-cross algorithm is shown in yellow, while the deserialisation is shown in pink and the serialisation is shown in green. MP7 infrastructure logic is shown in cyan.	103
6.11	Expected number of 3x5 clusters given different energy thresholds. From [69] . . .	104

7.1	Feynman diagram for the s-channel production of a dilepton pair via a Z' resonance.	109
7.2	The upper limits at 95% CL obtained by CMS on the production cross section multiplied by the branching fraction for a Z' boson relative to that of the Z boson. The Z' is modelled as a spin-1 resonance with an intrinsic decay width of 0.6% of the resonance mass in both the dielectron and dimuon channels. The limit is shown for the dielectron channel (upper left), dimuon channel (upper right) and the two channels combined (bottom). The green shaded area denotes the 68% quantile for the expected limit, while the yellow area denotes the 95% quantile. The predictions for two proposed Z' bosons, Z'_{SSM} and Z'_ψ are also shown. From [129].	111
7.3	The upper limits at 95% CL obtained by CMS on the production cross section multiplied by the branching fraction for a spin-2 KK excitation of a graviton relative to that of the Z boson. The limit is shown for the dielectron channel (upper left), dimuon channel (upper right) and the two channels combined (bottom). The green shaded area denotes the 68% quantile for the expected limit, while the yellow area denotes the 95% quantile. The prediction for the spin-2 KK excitation is shown for coupling parameters, $k/\overline{M_{Pl}}$, of 0.01 (pink), 0.05 (red) and 0.10 (black). From [129].	112
7.4	The upper limits at 95% CL obtained by CMS on the masses of the DM particle and the DM mediator. Exclusion regions were found by comparing the cross section multiplied by the branching fraction for the decay to a Z boson to values given by the simplified model. The grey curves represent regions that are consistent with the observed DM relic density [146]. From [129].	113
7.5	Feynman diagrams for the decay of a Z' to a dimuon pair in association with at least one b-quark. From [149].	113
7.6	The upper limits at 95% CL obtained by ATLAS on the production cross section multiplied by the branching fraction for a Z' to that of Zh and the branching fraction $B(h \rightarrow b\bar{b}, c\bar{c})$. The limit is shown for the dimuon channel only. The green shaded area denotes the 68% quantile for the expected limit, while the yellow area denotes the 95% quantile. The prediction for HVT model A (dashed red) and HVT model B (solid pink) are also shown. From [150].	114
7.7	The expected and observed transverse mass distribution for the electron distribution in the search for a W' boson in final states with a lepton and missing transverse energy. From [152].	115
7.8	Lower limits at 95% CL obtained by CMS on the ultraviolet cutoff parameter, Λ , for the eight CI models considered. Shown for (upper left) the dielectron channel, (upper right) the dimuon channel, and (lower) their combination. Limits for all models are given for $m_{\ell\ell} > 400$ GeV. From [129].	116

7.9	Exclusion limits at 95% CL obtained by CMS on the ultraviolet cutoff for the three ADD models: GRW (first bin), Hewett (second bin) and HLZ (third to seventh bin). Shown for (upper left) the dielectron channel, (upper right) the dimuon channel, and (lower) their combination. The signal models are calculated at L), with an NNLO correction factor of 1.3 applied. From [129].	117
7.10	Upper limits at 95% CL obtained by ATLAS on the visible cross-section, σ_{vis} , for the (a) electron b-veto, (b) electron b-tag, (c) muon b-veto, and (d) muon b-tag categories. The green shaded area denotes the 68% quantile for the expected limit, while the yellow area denotes the 95% quantile. The red marker gives the strongest expected lower limit. From [161].	118
7.11	Ratio of the differential dilepton production cross section obtained by CMS as given in Equation 7.1, shown as a function of the dilepton pair mass for the case where both leptons are in the barrel (top left), at least one lepton is in the endcap (top right) and the combined result (bottom). The error bars represent both statistical and systematic uncertainties. From [129].	119
8.1	Dielectron mass resolution from $Z \rightarrow ee$ events as a function of η of the electron. UL (in this plot labelled as Data (Legacy)) data is shown in green with EOY data (in this plot simply labelled Data) is shown in pink. The grey banded region is the transition between the barrel and the endcap region of the ECAL at $1.44 < \eta < 1.57$ and is not included. Plot taken from [84].	123
8.2	Turn-on curve for a 32 GeV electron trigger in the barrel measured using a tag and probe method with $Z \rightarrow ee$ events using 2018 data.	129
8.3	E_T , ϕ and η of the failing same sign probe regions in the barrel (left) and endcap (right). The QCD contribution is calculated by subtracting the contribution of all MC backgrounds from the data.	131
8.4	Invariant mass of the tag and probe pair. Shown for the probe in the barrel (left) and the probe in the endcap (right) with all probes (top), passing probes (middle) and failing probes (bottom).	133
8.5	E_T of the probe. Shown for the probe in the barrel (left) and the probe in the endcap (right) with all probes (top), passing probes (middle) and failing probes (bottom).	134
8.6	η of the probe.	135
8.7	ϕ of the probe. Shown for the probe in the barrel (left) and the probe in the endcap (right) with all probes (top), passing probes (middle) and failing probes (bottom).	136
8.8	nPV of the probe. Shown for the probe in the barrel (left) and the probe in the endcap (right) with all probes (top), passing probes (middle) and failing probes (bottom).	137
8.9	Efficiency of the HEEP ID as a function of E_T , where the non-DY processes are included (left) and subtracted (right), also split into barrel (top) and endcap (bottom).	138

8.10	Efficiency of the HEEP ID as a function of η , where the non-DY processes are included (left) and subtracted (right).	139
8.11	Efficiency of the HEEP ID as a function of ϕ , where the non-DY processes are included (left) and subtracted (right), also split into barrel (top) and endcap (bottom).	140
8.12	N-1 efficiency of the H/E cut as a function of E_T (top), ϕ (middle) and η (bottom) where the non-DY processes are included in the MC and compared to data.	141
8.13	N-1 efficiency of the $ d_{xy} $ cut as a function of E_T (top), ϕ (middle) and η (bottom) where the non-DY processes are included in the MC and compared to data.	142
8.14	N-1 efficiency of the missing hit cut as a function of E_T (top), ϕ (middle) and η (bottom) where the non-DY processes are included in the MC and compared to data.	143
8.15	Efficiency of the HEEP ID as a function of ϕ in UL samples (left) and EOY samples (right) for both barrel (top) and endcap (bottom) for 2017 data. EOY plots are taken from [170].	145
8.16	Efficiency of the HEEP ID as a function of E_T in UL samples (left) and EOY samples (right) for both barrel (top) and endcap (bottom) for 2017 data. EOY plots are taken from [170].	146
9.1	Feynman diagrams for the production of dileptons in association with zero (left), one (centre) and two (right) b-tagged jets in the final state. From [171].	150
9.2	Generated mass distribution for the $bb\ell\ell$ model at three different values of Λ for constructive interference (left) and for different chirality combinations at $\Lambda = 10$ TeV (right). From [171].	150
9.3	Feynman diagram for the production of dileptons via the transition $b \rightarrow s\ell\ell$. From [171].	151
9.4	Generated mass distribution for the $bs\ell\ell$ model at three different values of Λ . From [171].	151
9.5	Leading b-tagged jet p_T (left) and η (right) for the $bb\ell\ell$ signal, shown for different values of Λ (top) and for different chirality combinations at $\Lambda = 10$ TeV (bottom). From [171].	153
9.6	Leading b-tagged jet p_T (left) and η (right) for the $bs\ell\ell$ signal for three different Λ values.	154
9.7	Illustration of a DNN. From [180].	157
9.8	Distributions for the following input variables for 2018 data in the dimuon final state: leading muon p_T (top left), minimum $M_{\ell b}$ (top right), missing transverse energy (bottom left) and minimum 3D angle between the lepton and the b-tagged jet (bottom right). Two mass points of the $bs\ell\ell$ signal model are also shown. From [171].	158
9.9	DNN score for the dimuon final state for two different mass points of the $bs\ell\ell$ signal samples. From [171].	159

9.10	Signal efficiency and the signal to signal plus background ratio (left) and ROC curve (right) for the DNN for the dimuon final state. The cut on the DNN discriminator at 0.6 is shown in black. From [171].	159
9.11	Distributions for minimum $M_{\ell b}$ (left) and the minimum 3D angle between the lepton and the b-tagged jet (right) after the DNN cut (> 0.6) for the dimuon final state in 2018. From [171].	160
9.12	Distributions for minimum $M_{\ell b}$ (left) and the minimum 3D angle between the lepton and the b-tagged jet (right) after the DNN cut (> 0.6) for the dielectron final state in 2018. From [171].	160
9.13	Acceptance $\times$ efficiency for the 1 b-tagged jet category for the dimuon channel (left) and the dielectron channel (right). From [171].	162
9.14	Acceptance $\times$ efficiency for the 0 (left) and 2+ (right) b-tagged jet categories for the dielectron channel. From [171].	162
9.15	Pileup uncertainty for DY + Jets (left) and for $bb\ell\ell$ signal (right) for the dimuon 0 b-tagged jets category for the BB region. From [171].	163
9.16	Uncertainties in the PDFs for the DY (green) and other backgrounds (red). From [171].	164
9.17	Uncertainty in b-tagging efficiency for DY + Jets (left) and for $bb\ell\ell$ signal (right) for the dimuon 0 b-tagged jets category for the BB region. From [171].	164
9.18	Response matrix for the 0 b-tagged jets category in the BB region in 2018 for the dimuon channel (left) and the dielectron channel (right). From [171].	167
9.19	Closure test for the category with 0 b-tagged jets for the dielectron category, split into the BB region (left) and BE region (right). From [171].	167
9.20	Flavour ratio for the combined BB and BE region for the 0 b-tagged jets category using the full Run 2 data. From [171].	169
9.21	Predicted lepton flavour ratio for the 1 b-tagged jet category in the BB region for different values of Λ . Data is blinded. From [192].	169
9.22	Expected upper limit on the cross-section of the $bb\ell\ell$ model in the left-left model for the dimuon channel in the BB region for the 0 b-tagged jet category. The dashed line shows the expected limit, with the green shaded area denoting the 68% quantile and the yellow area denoting the 95% quantile. The red line shows the theoretical prediction for the constructive LL $bb\ell\ell$ signal model. From [171].	171
A.1	Efficiency of the HEEP ID for 2017 as a function of η , where the MC includes the DY process.	176
A.2	Efficiency of the HEEP ID for 2016 as a function of E_T (top), ϕ (middle) and η (bottom), where the MC includes the DY process. For the E_T and ϕ plots, this is further split into the BB category (left) and BE category (right).	177

Chapter 1

Introduction

The Standard Model (SM) of particle physics is a theory that describes elementary particles and their interactions, which has predicted the existence of particles many years before they have been discovered experimentally. Examples of such particles include the Higgs boson discovered by the ATLAS and Compact Muon Solenoid (CMS) collaborations after a 40-year search [1, 2].

However, there are also shortcomings in the SM that suggest that it is not a complete theory. Extensions to the SM have been suggested, but so far, no proof of these extensions has been found. An upgrade is planned for the Large Hadron Collider (LHC), which provides collisions for the ATLAS and CMS experiments. Detector upgrades are planned to allow the detectors to make the best use of this accelerator upgrade, which is expected to provide a factor of 10 more data to both experiments than they have so far recorded. This massive increase in data will allow more stringent limits to be placed on these new theories, and hopefully lead to a discovery of physics beyond the SM (BSM).

Searches for BSM physics are also ongoing using data that has already been collected by the CMS Experiment. This data has been recently recalibrated to improve the reconstruction and identification of particles, with new analysis methods used that offer increased sensitivity to new physics. One such analysis is the search for a new Z' boson, where previous searches have searched for the decay to a dilepton pair. Now, this analysis is being expanded by considering the decay of a Z' boson to a dilepton pair in association with b-tagged jets, which is only possible with the recalibrated dataset due to the improved reconstruction of heavy flavour jets. This search is motivated by the increased sensitivity of the third generation of quarks to BSM physics due to their higher masses.

The outline of this thesis is as follows:

- Chapter 2 provides an overview of the SM and its shortcomings. Models that expand on the SM are also detailed, with a focus on models that propose a new Z' boson that could be observed by the CMS detector. Proton-proton collisions and their simulation are also explained, as well as Drell-Yan production, which forms the primary background to Z'

searches.

- Chapter 3 describes the LHC and the CMS experiment, both currently and after the upcoming Phase II upgrade.
- Chapter 4 goes into more detail on the CMS electromagnetic calorimeter (ECAL), including how energy deposited by electrons and photons is clustered, how the detector is calibrated to optimise its energy resolution and how anomalous signals are identified and rejected.
- Chapter 5 describes the results of tests with an electron beam carried out at the CERN SPS to investigate the resolution of the upgraded ECAL with the first versions of prototype Phase II electronics, and also presents the effect of an energy correction method designed to provide the energy of an incident electron when some energy has been lost to other parts of the ECAL due to the impact position.
- Chapter 6 details the work that has been completed to test three algorithms that will be implemented in firmware as part of the ECAL upgrade. Two of these algorithms reject anomalous signals, while the third creates basic clusters of localised energy. These tests have been completed in software and on hardware using test beam data.
- Chapter 7 outlines the previous experimental searches for a Z' boson, particularly the decay $Z' \rightarrow \ell\ell$.
- Chapter 8 describes the analysis of a dedicated high energy electron/positron ID using recalibrated CMS data that is needed for searches for $Z' \rightarrow ee$.
- Chapter 9 details the search for a Z' boson decaying to a dilepton pair in association with b-tagged jets. This search considers two BSM models and a test of lepton flavour universality.

Chapters 5, 6 and 8 are mostly my own work, with the full details of the work I undertook detailed at the start of each chapter. Chapter 9 was mainly completed by others in the CMS Z' group and requires the work I completed in Chapter 8 to search for the decay of a Z' boson to a dielectron pair.

Chapter 2

Theory

2.1 The Standard Model

The Standard Model (SM) of particle physics describes elementary particles and their interactions. These elementary particles are split into fermions and bosons. Fermions are matter particles with non-integer spin (all fermions in the SM are spin $1/2$), while bosons are force-carrying particles with integer spin. Quarks and leptons are fermions, with six quarks and six leptons in the SM. Each fermion has its own antifermion. Due to the spin $1/2$ nature of fermions, they must obey the Pauli exclusion principle. This means that two identical fermions cannot occupy the same quantum state.

This thesis uses natural units throughout, meaning that

$$(2.1) \quad c = \hbar = \epsilon_0 = \mu_0 = 1$$

where c is the speed of light, $\hbar$ is the reduced Planck's constant, ϵ_0 is the permittivity of free space and μ_0 is the permeability of free space.

This section is largely based on Thomson's Modern Particle Physics [3] and Joel Goldstein's lectures on theoretical particle physics [4].

Generation	Quarks			Leptons		
	Flavour	Electric Charge (e)	Mass	Flavour	Electric Charge	Mass
First	up (u)	+2/3	$2.16^{+0.49}_{-0.26}$ MeV	electron (e)	-1	511 keV
	down (d)	-1/3	$4.67^{+0.48}_{-0.17}$ MeV	electron neutrino (ν_e)	0	< 1.1 eV
Second	charm	+2/3	1.27 ± 0.02 MeV	muon (μ)	-1	106 MeV
	strange (s)	-1/3	93^{+1}_{-5} MeV	muon neutrino (ν_μ)	0	< 0.19 MeV
Third	top (t)	+2/3	172.9 ± 0.4 GeV	tau (τ)	-1	1.78 GeV
	bottom (b)	-1/3	$4.18^{+0.03}_{-0.02}$ GeV	tau neutrino	0	< 18.2 MeV

Table 2.1: Quarks and leptons of the SM. The masses listed are taken from [5].

Gauge boson	Force	Electric Charge	Mass (GeV)	Spin	Range (m)
Photon (γ)	electromagnetic	0	0	1	∞
$W^\pm$	weak	± 1	80.377 ± 0.012	1	10^{-18}
Z	weak	± 1	91.1876 ± 0.0021	1	10^{-18}
gluon (g)	strong	0	0	1	10^{-15}

Table 2.2: Gauge bosons of the SM. The masses listed are taken from [5].

2.1.1 Particles in the Standard Model

Table 2.1 lists all the quarks and leptons that comprise the SM. Quarks are particles that interact via the electromagnetic, weak and strong interactions. All quarks have an electric charge. For up (u), charm (c) and top (t) quarks this charge is $+\frac{2}{3}e$, while for down (d), strange (s) and bottom (b) quarks this charge is $-\frac{1}{3}e$. Quarks are also arranged into generations in order of their mass, with a pair of quarks of different charges in each generation. The first generation of quarks comprises of up and down quarks, the lightest quarks in the SM. The second generation of quarks includes the strange and charm quarks, while the third generation is made up of the top and bottom quarks, the heaviest quarks in the SM. The first generation of quarks and electrons make up all normal matter observed on Earth. Heavier quarks and leptons are unstable and rapidly decay to their lighter counterparts.

In addition to electric charge, quarks also have a colour charge. This colour charge has three possible states: red, green and blue. All three colours combined create a colour-singlet state, as is the case with baryons. In mesons, a colour anti-colour combination also gives a colour-singlet state. Due to colour confinement, quarks can only be found in colour-singlet states.

Leptons are colourless particles that interact via electromagnetic and weak interactions. As with quarks, leptons are also split into three generations. The first generation is the electron and electron neutrino, while the second generation comprises the muon and the muon neutrino and finally the third generation includes the tau and the tau neutrino. Again, these generations are in mass order of the charged leptons. The electron, muon and tau are charged leptons, while their corresponding neutrinos are neutral leptons. The rest mass of neutrinos is significantly smaller than that of any of the other matter particles in the SM (around a million times lighter than an electron). As neutrinos are chargeless, their only SM interactions are via the weak force.

Bosons are split into two subgroups: gauge bosons and scalar bosons. Gauge bosons are integer spin particles, with all the gauge bosons in the SM summarised in Table 2.2. First, there is the photon, a massless, chargeless particle that mediates electromagnetism. Then, there are eight types of gluon, which mediate the strong interaction. Gluons are massless and have no electric charge, though they do have colour charge. They can interact with themselves and with quarks. There are three bosons that mediate the weak interaction: W^+ , W^- and Z^0 (often simply referred to as Z).

There is significant variation in the strength and range of force carried by gauge bosons. As electromagnetism is carried by massless photons, it has an infinite range and its strength is governed by the electromagnetic coupling constant. By contrast, the weak force is mediated by massive particles, which give it a range of $\approx 10^{-18}$ m. The effective strength of the weak force is described by the Fermi constant and it is 10^{-6} times weaker than the strong force. Though the strong force is mediated by massless gluons, it has a limited range of 10^{-15} m due to colour confinement. This will be discussed in Subsection 2.1.3.

The Higgs boson is the only scalar boson in the SM. It is a massive particle with zero spin, electric charge or colour charge, and gives mass to particles in the SM via the Higgs mechanism. Details on this can be found in Subsection 2.1.4. In the SM, the mass of the Higgs boson is a free parameter. The decay width, branching fractions and production cross sections depend on the Higgs boson mass [6].

2.1.2 Quantum Electrodynamics

The behaviour of free massive fermions can be described by the Dirac Lagrangian

$$(2.2) \quad \mathcal{L}_{Dirac} = \bar{\psi}(i\partial\!\!\!/ - m)\psi$$

where ψ is a fermionic field with a mass m and $\partial\!\!\!/ = \gamma_\mu \partial^\mu$. When expanded, the first term of this equation is the kinetic term and the second term is the mass term.

The Dirac Lagrangian is invariant under the U(1) transformation

$$(2.3) \quad \psi(x) \rightarrow e^{-iq\theta}\psi(x)$$

which is essentially a change of phase. This invariance is known as local gauge invariance, which is similar to gauge invariance, but allows the phase $q\theta(x)$ to vary at different points in space-time. This gives rise to a spin-1 vector field A_μ and its transform A'_μ which is given by

$$(2.4) \quad A_\mu \rightarrow A'_\mu = A_\mu - \partial_\mu \theta$$

From this, it is possible to define the covariant derivative, D_μ , as

$$(2.5) \quad D_\mu = \partial_\mu + iqA_\mu$$

where q is a coupling constant. It is also possible to define the field strength tensor, $F_{\mu\nu}$, by

$$(2.6) \quad F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$$

The field strength tensor is also local gauge invariant, leading to the QED Lagrangian

$$(2.7) \quad \mathcal{L}_{QED} = \bar{\psi}(i\partial\!\!\!/ + q\mathcal{A} - m)\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}$$

where the $\mathcal{A}$ term is the interaction term, the $F_{\mu\nu}$ term is the kinetic term and the other terms come from Equation 2.3.

2.1.3 Quantum Chromodynamics

While QED has a U(1) gauge invariance, quantum chromodynamics (QCD) is SU(3) gauge invariant. This is because of the three orthogonal colour components of quarks which can be represented as

$$(2.8) \quad q = \begin{pmatrix} r \\ g \\ b \end{pmatrix}$$

where r, g and b stand for the three colour components red, green and blue. As QCD is SU(3) invariant, there are eight generators and therefore eight gluons which carry a combination of colour and anticolour. The QCD Lagrangian is given as

$$(2.9) \quad \mathcal{L}_{QCD} = \bar{\psi} \left(i\not{\partial} + g_s \not{A}_i + \frac{\lambda_i}{2} - m \right) \psi - \frac{1}{4} F_i^{\mu\nu} F_{i\mu\nu}$$

where λ_i are the Gell-Mann matrices which represent the generators and g_s is the strong coupling constant [7].

The non-Abelian (elements do not commute) nature of SU(3) leads to gluon self-interactions at first order. At scales greater than a femtometre, it becomes more energetically favourable for two separating quarks to produce a quark-antiquark pair than to continue to increase the distance between the original pair. Therefore, quarks cannot be separated from their parent meson or hadron without producing new hadrons. This is known as colour confinement and is the same process that produces jets of particles that can be detected in particle detectors.

2.1.4 Electroweak Theory

Weak isospin can also be described using a gauge theory, specifically an SU(2) interaction between the fermion doublets $u \leftrightarrow d$, $e \leftrightarrow \nu_e$, and similarly for the second and third generations. The three weak isospin gauge bosons, $W^\pm$ and Z, are collectively referred to as W_i .

To describe the weak force, it is necessary to first discuss helicity and chirality. Helicity is the projection of spin in the direction of motion and is frame dependent, so therefore not Lorentz invariant. Chirality is a fundamental property of fermions which gives the eigenstates of the γ^5 matrix, which is defined as

$$(2.10) \quad \gamma^5 \equiv i\gamma^0\gamma^1\gamma^2\gamma^3 = \begin{pmatrix} 0 & I_2 \\ I_2 & 0 \end{pmatrix}$$

where I_2 is the 2×2 identity matrix. An eigenvalue of $+\frac{1}{2}$ corresponds to a right-handed particle while an eigenvalue of $-\frac{1}{2}$ corresponds to a left-handed particle. In the massless limit, the helicity and chirality of a particle will be the same.

The weak force only interacts with left-handed chiral particles, which have a charge known as weak isospin, I_3 . Left-handed fermions have a weak isospin of $+1/2$, while the $W^\pm$ bosons

have a weak isospin of ± 1 and all other particles have a weak isospin of zero. Weak interactions can violate parity conservation [8, 9], as well as the combination of charge and parity, this is known as charge-parity (CP) violation [10, 11].

One consequence of parity violation is helicity suppression in decays, where the decay rate is lower than would otherwise be expected. This can occur in the case of a massive scalar particle decaying to a fermion-antifermion pair. One would expect the decay rate to be proportional to the momentum of the final state particles, therefore lower fermion masses would be expected to give larger decay rates. However, this is not what is observed. The scalar boson does not have any angular momentum, so, to conserve this quantity, the spins of the two fermions must be opposite, and therefore the helicity of the two must be the same. For light fermions, there is a small overlap between a left-handed chiral state and $+1/2$ particle helicity, and vice versa for the antifermion.

The weak and electromagnetic forces were unified by Glashow [12], Weinberg [13] and Salam [14]. As may be expected given the groups for the electromagnetic and weak forces, the electroweak forces are an $SU(2) \times U(1)$ gauge group. The generator of the $U(1)$ component is known as weak hypercharge, Y_W and is defined as

$$(2.11) \quad Y_W = 2(q - I_3)$$

where q is the electromagnetic charge and the factor of two comes from convention but could be absorbed elsewhere.

The generators of the $SU(2)$ component, responsible for the weak force, are $\mathbf{T} = \sigma/2$ where σ are the Pauli matrices. As mentioned previously, the weak force interacts differently with left- and right-handed fields. The left-handed fields, ψ_L are represented as doublets, while right-handed fields, ψ_R are represented as singlets. This can be written as

$$(2.12) \quad \psi_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix}, \begin{pmatrix} \nu_L \\ \ell_L \end{pmatrix}$$

$$\psi_R = u_R, d_R, \ell_R$$

where u_L represents up, charm and top quarks and d_L represents down, strange and bottom quarks, with each u_L and d_L pair being of the same generation.

The left and right-handed components of fermionic fields can be obtained using projection operators

$$(2.13) \quad P_{L,R} = \frac{1}{2}(1 \pm \gamma^5)$$

As before, local gauge invariance is required which leads to defining the covariant derivative as

$$(2.14) \quad D_\mu = \partial_\mu + ig\mathbf{T} \cdot \mathbf{W}_\mu + i\frac{g'}{2}Y_W B_\mu$$




Figure 2.1: Higgs potential as a function of the Higgs field average value for $\mu = 90$ GeV and $\lambda=0.13$. From [18].

where g' is another coupling constant and W_μ and B_μ are gauge fields.

From this, the electroweak Lagrangian can be written as

$$(2.15) \quad \mathcal{L}_{EWK} = \bar{\psi} i \gamma^\mu D_\mu \psi - \frac{1}{4} W_{\mu\nu}^a W_a^{\mu\nu} - \frac{1}{4} B_{\mu\nu} B^{\mu\nu}$$

The SM gauge bosons are recovered by a linear combination of the gauge fields W_μ^a and B_μ^a .

However, Equation 2.15 only allows massless particles due to gauge invariance, which does not agree with the massive particles that have been observed experimentally. This leads to the prediction of the Higgs boson. In the SM, particles have mass due to their interactions with the Higgs field. The Higgs mechanism comes from the spontaneous symmetry breaking of the $SU(2) \times U(1)$ symmetry [15–17]. This is achieved by giving the field a non-zero vacuum expectation value, v , by using a potential of the form

$$(2.16) \quad V(\phi) = -\mu^2 |\phi|^2 + \lambda |\phi|^4$$

where ϕ is a complex scalar field and with $\mu^2 > 0$. The shape of this potential can be seen in Figure 2.1.

With this new potential, it is possible to write the Lagrangian as

$$(2.17) \quad \mathcal{L}_{scalar} = D_\mu \phi^\dagger D^\mu \phi - V(\phi)$$

This Lagrangian gives rise to massive gauge bosons. Gauge invariance is reconciled with the observed massive fermions by the introduction of a Yukawa interaction between the scalar field ϕ and the fermions. For example, for the electron, the Yukawa term is

$$(2.18) \quad \mathcal{L}_e = -G_e (\bar{l}_L \phi e_R + \bar{e}_R \phi^\dagger l_L)$$

where G_e is an arbitrary coupling constant and l_L is defined as

$$(2.19) \quad l_L = \begin{pmatrix} \nu_{eL} \\ e_L \end{pmatrix}$$

2.1.5 Shortcomings in the Standard Model

The SM is a well-tested theory that is expected to be an effective theory, meaning a low-energy approximation of a higher-energy scale theory. The indications of this come from shortcomings in the SM, some of which are detailed below.

- **Gravity.** The only one of the four fundamental forces that is missing from the SM is gravity. The combination of the classical physics theory of general relativity with the SM would require a quantum theory of gravity, which could arise from the addition of a spin-2 gauge boson for gravity, known as a graviton. However, this approach leads to a theory that cannot be normalised and gives diverging predictions of observable quantities.
- **Matter antimatter asymmetry,** also known as baryon asymmetry. There are more baryons in the universe than anti-baryons, as evidenced by the fact that there is matter in the universe today. There is a small amount of baryon asymmetry in the SM through CP violation, but it is insufficient to describe the amount of matter currently observed in the universe.
- **Dark matter.** Astronomical observations such as galaxy rotation curves, gravitational lensing and redshift-space distortions have found that the visible content of the universe only makes up around 5% of the total energy content of the universe. Approximately 25% of the total energy content is expected to come from dark matter, which is matter that does not interact electromagnetically or via strong interactions, and therefore it has not yet been detected [19]. There are no candidates for dark matter within the SM. The remaining 70% of the energy content is attributed to dark energy which has been proposed to explain the accelerated rate of expansion of the universe.
- **Hierarchy problem.** There is a significant difference between the energy scale at which the electroweak forces unify of around 10^2 GeV and the Planck scale at 10^{19} GeV. The Planck scale is the point at which the quantum effects of gravity dominate and the SM breaks down. This leads to the question of why the EW scale is so small compared to the Planck scale, and why the two scales are not closer in value.
- **Fine-tuning problem.** This is related to the hierarchy problem above. In the SM, parameters have to cancel out to 11 digits to produce a Higgs mass around the observed value of 125 GeV. This level of cancellation without any apparent reason is called fine-tuning.

- **Number of free parameters in the SM.** There are 25 free parameters in the SM to be determined experimentally (26 if the strong CP phase is not taken to be zero). These free parameters include particle masses and also parameters that determine the transitions between quarks and neutrinos. There is significant variation in particle masses from < 1 eV for neutrinos to $> 10^7$ eV for top quarks, along with a mass hierarchy amongst different generations and within generations. It is possible that these hierarchies ultimately arise from another model with fewer free parameters.

The first three shortcomings listed (gravity, matter-antimatter asymmetry and dark matter) are significant discrepancies between the SM and what we have observed in nature. The other problems listed are more aesthetic in nature, that is unsatisfying parts of the SM which raise questions but are not necessarily wrong.

2.2 Beyond the Standard Model

2.2.1 Supersymmetry

Supersymmetry (SUSY) is one of many beyond the SM (BSM) theories that explains some of the shortcomings listed above. For each fermion in the SM, SUSY predicts a supersymmetric partner that is a spin zero boson, while for each boson in the SM, SUSY predicts a spin $\frac{1}{2}$ fermion.

Both baryon number and lepton number are conserved in the SM, while this is not usually true in SUSY. Instead, a new quantum number is introduced called R-parity, which is defined as

$$(2.20) \quad R_p = (-1)^{3(B-L)+2s}$$

where B is the baryon number ($B = (n_q - n_{\bar{q}})/3$, where n_q is the number of quarks and $n_{\bar{q}}$ is the number of antiquarks), L is the lepton number ($L = n_l - n_{\bar{l}}$ where n_l is the number of leptons and $n_{\bar{l}}$ is the number of antileptons) and s is the spin of the particle. All SM particles have $R_p = +1$ while all SUSY particles have $R_p = -1$. SUSY particles in the TeV range would solve the fine-tuning problem by the addition of loop corrections involving SUSY particles that cancel those in the SM. This energy range is accessible with the Large Hadron Collider (LHC).

Searches for supersymmetry have been undertaken by several particle physics experiments, including the Compact Muon Solenoid (CMS) experiment based at the LHC. For data taken between 2016 and 2018, the CMS experiment was able to place lower limits on the gluino mass between 2000 and 2310 GeV at 95% confidence level, and lower limits on the squark mass of 1190 to 1630 GeV, depending on the production scenario [20].

2.2.2 Sequential Standard Model

The sequential standard model (SSM) [21] predicts a Z'_{SSM} , which is a heavy version of the Z boson found in the SM. In experimental searches for new heavy resonances, the Z'_{SSM} is often used as a benchmark model. Details of previous experimental searches for new massive resonances can be found in Chapter 7.

2.2.3 Grand Unified Theories

Other possible extensions to the SM are grand unified theories (GUTs), where the strong, weak and electromagnetic forces are unified in a single gauge interaction. This is expected to occur at a high energy scale known as the GUT scale, and at lower energies, the symmetry is expected to break to recover the SM. The introduction of a new gauge group suggests that there is at least one new neutral gauge boson.

The minimal GUT extension of the SM is to a $SU(5)$ group [22], which gives a symmetry breaking scale of $\mathcal{O}(10^{16})$ GeV. However, an $SU(5)$ group gives rise to predictions of the proton lifetime that are incompatible with current measurements, and therefore an $SU(5)$ group cannot be an extension to the SM.

Instead of an $SU(5)$ group, another possibility is an $SO(10)$ group, which contains the $SU(5)$ group. The symmetry breaking of this proposed group is as follows

$$(2.21) \quad SO(10) \rightarrow SU(5) \times U(1)_\chi \rightarrow SU(3) \times SU(2) \times U(1) \times U(1)_\chi$$

where $SU(3) \times SU(2) \times U(1)$ is the SM and the additional $U(1)_\chi$ term would provide a new particle of charge χ . As with the SUSY particles mentioned above, the mass of this new boson would be in the TeV range.

The $SO(10)$ group can also be extended to another unitary group known as the exceptional E_6 group. In this case, the symmetry breaking is as follows

$$(2.22) \quad \begin{aligned} E_6 &\rightarrow SO(10) \times U(1)_\psi \rightarrow SU(5) \times U(1)_\psi \times U(1)_\chi \\ &\rightarrow SU(3) \times SU(2) \times U(1) \times U(1)_\psi \times U(1)_\chi \end{aligned}$$

The groups $U(1)_\chi$ and $U(1)_\psi$ can be linearly combined to give $U(1)_\theta = U(1)_\chi \cos \theta - U(1)_\psi \sin \theta$. In this case, only one of the new bosons is expected to be at the TeV scale, with the other significantly heavier. Four mixing angles, specifically $\theta = 0, -\pi/2, \sin^{-1} \sqrt{3/8}$ and $-\sin^{-1} \sqrt{3/8}$, give rise to four possible Z' bosons that couple differently to quarks and leptons. These four bosons are the $Z'_\psi, Z'_\chi, Z'_\eta$ and Z'_I bosons, which are often considered in experimental Z' searches and will be further described in Chapter 7.

2.2.4 Contact Interactions

When the energy scale of the new physics phenomena exceeds the accessible energy scale provided by a collider, then it is not possible to search directly for a resonance and different

models must be considered to search for BSM physics. One such model is four-fermion contact interaction (CI) theory, where four-fermions are considered to interact at a single point rather than via a mediator such as a W boson.

One possible extension to the SM is the existence of preons, which would be the constituents of quarks and leptons [23, 24]. These preons would be bound by a new gauge interaction similar in behaviour to QCD, and, if the collision energy exceeded the binding energy of this gauge interaction, then preons could be directly observed via production processes. If this condition is not satisfied, then it may still be possible to observe deviations from the SM in the high energy tails of distributions such as dilepton mass.

The term of the Lagrangian of CI theory for a quark-antiquark pair producing a dilepton pair is given by:

$$(2.23) \quad \mathcal{L} = \frac{g^2}{\Lambda^2} \sum_{i,j=L,R} \eta_{ij} (\bar{q}_i \gamma^\mu q_i) (\bar{\ell}_j \gamma_\mu \ell_j)$$

where g is a coupling constant defined such that $g^2 = 4\pi$, Λ is the compositeness scale, η_{ij} are helicity parameters, q_L is a left-handed quark doublet, q_R is a sum over the right-handed quark singlets. The compositeness scale is the scale at which quarks and leptons are composite particles. The helicity parameters are +1 or -1. If the helicity current is positive, then there is destructive interference and the cross-section observed will be reduced compared to the SM value. The converse is also true, with a negative helicity current causing constructive interference. The compositeness scale may vary for each helicity current, so each helicity current and its associated scale is usually considered separately in any search for BSM physics.

The BSM search I contributed to, which is detailed in Chapter 9, uses two CI models to interpret results in the search for new physics in a dilepton final state.

2.2.5 Extra Dimensions

Theories including extra spatial dimensions have been proposed to help solve the hierarchy problem. There are multiple models for these extra dimensions, the most notable ones with relevance to this thesis include one proposed by Arkani-Hamed, Dimopoulos and Dvali (ADD) [25, 26], in which SM particles are confined to the four traditional dimensions of space and time, and another model proposed by Randall and Sundrum (RS) [27], which adds one more spatial dimension and predicts that gravity can propagate through all five dimensions in the model. Each model provides a different experimental signature, with the Randall–Sundrum (RS) model predicting a distinct spin-2 high-mass resonance, while the ADD model predicts a non-resonant excess of events at high mass due to a series of nearly mass-degenerate excitations.

There are three different conventions to parameterise the cross-section in the ADD model, which were presented by Hewett [28]; Giudice, Rattazzi and Wells (GRZ) [29] and Han, Lykken and Zhang (HLZ) [30]. The Hewett and GRW conventions do not depend on the number of

extra dimensions, while the HLZ convention does depend on this. Because of this, experimental limits on the HLZ model are quoted for a range of extra dimensions. Experimental limits for the Hewett convention are quoted with respect to an energy cutoff scale, Λ_T , while in the GRZ and HLZ conventions, limits are quoted with respect to the string scale, M_S .

Both CI theory and the ADD model are valid below a certain energy scale, with the point at which they break down known as the ultraviolet cutoff [31]. In experimental searches for new phenomena involving these models, lower limits are often placed on the ultraviolet cutoff parameter, Λ . Experimental searches have been completed and placed exclusion limits of between 2 and 6 TeV, depending on the model used [32].

2.3 Proton-Proton Collisions

The collisions of interest in this thesis are proton-proton collisions at the LHC. Protons are baryons comprised of two up quarks and one down quark, with an overall charge of +1. These three quarks in a proton are referred to as valence quarks as they are the quarks which give the quantum numbers of the hadron. Valence quarks interact with each other and exchange gluons, which in turn produce more gluons or sea quarks. Sea quarks are quark-antiquark pairs that are produced by gluon interactions within the proton. The combination of quarks and gluons within a proton are referred to as partons.

Each parton carries a fraction of the total momentum of the proton. A parton distribution function (PDF) gives the probability for a parton to be carrying a given fraction of the proton's momentum. It should be noted that PDFs depend on the factorisation scale which corresponds to the resolution with which hadrons are probed, so measurements from many different experiments are combined for PDF measurements. PDFs at two different factorisation scales can be seen in Figure 2.2.

2.4 Drell-Yan Production

The production of an electron-positron pair from proton-proton collisions, known as the Drell-Yan (DY) process [34], is of particular interest in this thesis as the major background process for the analysis presented in Chapter 8. The leading order Feynman diagrams for the DY process can be seen in Figure 2.3, where a quark-antiquark pair annihilate to produce a γ^* or a Z^0 , with the mediator then creating an electron-positron pair. The differential DY cross-section has been measured experimentally by the CMS collaboration for the dilepton invariant mass range $15 < m < 3000$ GeV [35].




Figure 2.2: Next-to-leading order PDFs at two different factorisation scales: $Q^2 = 10 \text{ GeV}^2$ (left) and $Q^2 = 10^4 \text{ GeV}^2$ (right). From [33].




Figure 2.3: Feynman diagrams that contribute to the DY process: exchange of a virtual photon (top) and exchange of a Z^0 boson (bottom).

2.5 Simulation

The simulation of physics processes and detector response is essential to allow comparison between experimental data and the SM. In particle physics, Monte Carlo (MC) simulation is used. This simulation includes four steps: generation, detector simulation, digitisation and reconstruction. The first stage, generation, involves simulating the hard scatter between partons in the colliding protons and the particles generated from this scatter. Detector simulation takes the particles from the generation step and simulates the path of these particles through the detector, including the interaction with detector elements, while the digitisation step calculates the expected response of the detector electronics to the particles, recording hits in the active volumes of the detector. Finally, the reconstruction step takes the detector hits and uses them to construct the physics objects, such as electrons or photons, and measure their associated properties such as energy or momentum.

An illustration of a proton-proton collision is shown in Figure 2.4. This highlights the different steps of the collision that must be considered during the generation step, from the initial hard scatter to the parton shower and then hadronisation of the partons. Several different MC event generators are in existence, and the MC event generators used in this thesis are described below.

Both aMC@NLO [37] and positive weight hardest emission generator (POWHEG) [38–40] are next-to-leading order (NLO) MC event generators that calculate the hard scatter. As NLO generators, they consider tree-level and single-loop Feynman diagrams. POWHEG starts by calculating the hardest process in the event to reduce the chance of generating events with negative weights for terms that interfere destructively, though this is dependent on the PDFs. This is contrasted with aMC@NLO which regularly generates events with negative weights. Negative weights can cause unphysical behaviour in low statistics cases, such as when calculating the efficiency of particle identification criteria. In this case, the particles that fail the identification criteria may have a negative weight, giving an efficiency greater than 100%.

The hard scatter events generated by aMC@NLO or POWHEG can then be fed into the Pythia generator [41, 42]. While Pythia is capable of computing the entire particle generation step, as it is only LO it is often used to generate the initial and final state radiation and hadronisation of events based on the information given by another hard scatter generator. After this, events are propagated through the detector using the GEANT4 program [43].

GEANT4 is provided with a detailed model of the detector and simulates the hits, particle showers and hence response of the detector electronics. Finally, this simulated response and digitisation is reconstructed using the same software as the recorded data to reconstruct physics objects.

The simulation of signals and backgrounds is an essential part of the analysis presented in Chapters 8 and 9.




Figure 2.4: Illustration of a proton-proton collision as simulated using an MC event generator. The red circle at the centre depicts the hard collision and the Bremsstrahlung from parton showers is shown by the tree-like structure surrounding the centre. A secondary hard scatter event is depicted by the purple oval. Light green ovals depict hadronisation, while dark green ovals represent hadron decays. The yellow lines indicate soft photon radiation. From [36].

Chapter 3

Compact Muon Solenoid and the Large Hadron Collider

3.1 Introduction

This chapter details the Large Hadron Collider (LHC), an accelerator based at the European Organisation for Nuclear Research (CERN), as well as the Compact Muon Solenoid (CMS) Experiment. This thesis covers both analysis of CMS data and the Phase II upgrade of the detector.

3.2 The Large Hadron Collider

The LHC is a 27 km long circular proton-proton accelerator based at the CERN accelerator complex [44]. It is the last accelerator in the complex, preceded by other smaller accelerators that accelerate particles up to a given energy before the particles move on to the next accelerator in the complex. The LHC uses superconducting magnets to guide particles around the ring and boost them up to the desired energy.

The operating periods of the LHC are referred to as runs, with Run 1 occurring between 2010 and 2012 at a centre of mass energy of up to 8 TeV, while Run 2 occurred between 2015 and 2018 with a centre of mass energy of 13 TeV. During Run 3 of the LHC (2022-2025), the beam time is largely dedicated to protons, with about one month a year dedicated to heavy ion physics. The centre of mass energy for Run 3 is 13.6 TeV. Proton physics is the main focus of CMS and this thesis, so although the LHC can also generate heavy ion collisions, these will not be discussed in any detail.

The full CERN accelerator complex can be seen in Figure 3.1. In Run 3, hydrogen gas is injected into the ion source and a strong electric field is used to produce H^+ . These H^+ are attracted towards a Caesium cathode which donates electrons to the ions producing H^- .

These H^- ions are accelerated up to 160 MeV using linear accelerator 4 (LINAC4). After LINAC4, the H^- are stripped of their electrons to give protons. In previous running periods, LINAC2 was used to accelerate protons to 50 MeV. Next, the protons are injected into the proton synchrotron booster (PSB), a circular accelerator, which accelerates the protons up to an energy of 1.4 GeV. From here, the protons are injected into another circular accelerator, the proton synchrotron (PS), which accelerates the protons to 25 GeV. After this the protons are injected into the super proton synchrotron (SPS), the final accelerator before the LHC, where they will reach energies of 450 GeV. In Run 3, the protons reach an energy of 6.8 TeV in each direction to give a centre of mass energy of 13.6 TeV. This is an increase of 600 GeV on the centre of mass energy in Run 2.




Figure 3.1: The CERN accelerator complex and experiments in 2022. From [45].

There are two proton beams in the LHC, with each beam consisting of bunches of around 1.5×10^{11} protons and with approximately 2800 bunches per beam. The two proton beams are accelerated in opposite directions and guided around the LHC by a series of magnets. The

LHC uses a combination of 1232 dipole magnets and 540 quadrupole magnets. Each dipole magnet has an 8.4 T magnetic field while the each quadrupole magnet has a 7.5 T magnetic field [46]. The dipole magnets are responsible for deflecting the two beams along the trajectory of the LHC beam pipes, while the quadrupole magnets squeeze the beam. There are also eight superconducting radio-frequency cavities that accelerate the beam.

The two beams can be brought to collide at specific interaction points around the ring. Four of these interaction points coincide with the four main experiments on the LHC ring:

- ALICE (A Large Ion Collider Experiment): a dedicated heavy-ion (lead-lead) physics experiment. The focus of this experiment is to study quark-gluon plasma, a phase of matter that existed shortly after the Big Bang.
- ATLAS (A Toroidal LHC ApparatuS): a general-purpose experiment designed to investigate phenomena beyond the Standard Model (SM) and perform precision SM measurements.
- CMS (Compact Muon Solenoid): a second general-purpose experiment with the same experimental aims as ATLAS. It is described in detail in Section 3.3.
- LHCb (LHC beauty): an experiment that focuses on b-physics, including CP violation in charm and beauty hadron decays.

At these interaction points, proton bunches collide in a bunch crossing. While most of the protons will miss each other and continue to travel around the ring, there will be several collisions per bunch crossing. The background of particles from the collisions other than the hard scattering collision of interest is called pileup (PU). This background can be split into in-time PU, which is other collisions from the same bunch crossing, and out-of-time PU due to collisions from a previous or subsequent bunch crossing.

During Run 3, the collision rate is 40 MHz with a bunch crossing every 25 ns. The instantaneous luminosity, $\mathcal{L}$ is a measure of collision rate and is defined as

$$(3.1) \quad \mathcal{L} = \frac{f N_1 N_2}{4\pi \sigma_x \sigma_y}$$

where f is the frequency of the bunch crossings, N_1 and N_2 are the number of particles per bunch, and σ_x and σ_y are the widths of the beam in the x and y directions respectively. Assuming a Gaussian beam profile, then $\sigma_x = \sigma_y$. The integrated luminosity for CMS for proton-proton collisions from 2011 to 2022 is shown in Figure 3.2.




Figure 3.2: Total integrated luminosity from 2011 to 2022 recorded by the CMS experiment (in yellow) and delivered by the LHC to the CMS experiment (in blue). The flat periods represent times when no collisions took place, often for detector and accelerator upgrade purposes. From [47].

3.3 Compact Muon Solenoid

The CMS experiment is a hermetic detector with a nearly 4 T magnetic field which causes charged particles to follow a curved path as they travel through the detector. A schematic outlining the whole detector can be seen in Figure 3.3, with the subdetectors explained in the following subsections.

As previously mentioned, the CMS detector is designed as a general purpose detector capable of making discoveries at the high luminosities provided by the LHC. The design goal of CMS is to efficiently identify and precisely measure muons, electrons and photons with an energy resolution of around 1% over a large range of momenta [49]. This requires a high precision electromagnetic calorimeter (ECAL) and muon system, as well as high quality tracking.

Relative to the ATLAS detector, the CMS detector is compact with a length of 22 m and a diameter of 15 m (the ATLAS detector is 46 m long and 25 m in diameter). However, the CMS detector is twice the weight at approximately 12,500 t. The central solenoid magnet of




Figure 3.3: The CMS detector. From [48].

the ATLAS detector provides a 2 T magnetic field, while the CMS magnetic field is 3.8 T. The CMS detector is cylindrical, with the interaction point at the centre. It can be split into two parts: a barrel section that is parallel to the beamline, and two endcaps that close the detector at each end.

The CMS detector uses two coordinate systems, which are both illustrated in Figure 3.4. The first of which is the Cartesian coordinate system, where the x-axis points towards the centre of the LHC ring, the y-axis points upwards and the z-axis points in the anti-clockwise direction of the LHC beam trajectory when viewed from the top. The second coordinate system is a spherical coordinate system, with the origin defined as the centre of the detector. In the spherical coordinate system, the three coordinates are the distance from the origin, the azimuthal angle ϕ which is the angle in the transverse (x-y) plane relative to the x-axis and finally the polar angle θ which is defined as the angle in the longitudinal (y-z) plane relative to the z-axis. Instead of polar angle, pseudorapidity η is often used. Pseudorapidity is defined as:

$$(3.2) \quad \eta = \log(\tan(\theta/2)) = \frac{1}{2} \ln \frac{|\vec{p}| + p_z}{|\vec{p}| - p_z}$$

For particles with a mass significantly smaller than their momentum, pseudorapidity is equivalent to rapidity, y , which is defined as




Figure 3.4: The coordinate system used by the CMS experiment. From [50].

$$(3.3) \quad y = \frac{1}{2} \ln \frac{E + p_z}{E - p_z}$$

3.3.1 Magnetic Solenoid

The superconducting solenoid in the CMS detector provides an axial magnetic field of 3.8 T, achieved using a current of around 20 kA. It is one of the largest superconducting magnets that have been built to date. The solenoid has a length of 12.5 m and an inner diameter of 6 m [51]. The superconductor is a Rutherford cable, a twisted multifilament-strand cable design common in particle physics [52]. The cables are NbTi and coated with copper. On either side of the Rutherford cable is pure aluminium conductor which acts as a magnetic stabiliser, with further aluminium alloy used as a mechanical stabiliser as shown in Figure 3.5.

The magnet is cooled to 4 K using liquid helium. The iron return yoke controls the magnetic field outside of the solenoid and consists of two main components: a barrel yoke and two endcap yokes. The barrel yoke contains five rings, while each endcap yoke contains three disks. Outside the solenoid, the magnetic field strength is approximately 2 T. The tracker and calorimeters are inside the solenoid, while the muon chambers are interspersed with the return yoke.

3.3.2 Tracker

The CMS tracker is the subdetector closest to the interaction point, lying inside the solenoid. It is cylindrical with a length of 5.8 m and a diameter of 2.5 m, covering a pseudorapidity region of $|\eta| < 2.5$. The tracker is designed to reconstruct both primary (proton-proton interaction)




Figure 3.5: Cross section of the CMS conductor. From [53].

and secondary (particle decay) vertices and also to provide measurements of the momenta of electrically charged particles using the magnetic field provided by the solenoid. It therefore requires high granularity and a fast response. The tracker is made up of two main sections: a pixel tracker (closest to the interaction point) and a strip tracker (which surrounds the pixel tracker). This can be seen in Figure 3.6. Both are made of silicon which provides the required granularity and response characteristics.




Figure 3.6: Diagram showing one quarter of the CMS Tracker during Phase I. The green represents the pixel detector, while the red represents single-sided strip modules and the blue represents double-sided strip modules. From [54].

The tracker is designed with a high radiation tolerance so that it is able to cope with the harsh radiation environment of the LHC. To minimise the interactions of particles (such as multiple scattering, photon conversion or bremsstrahlung) within the tracker, the minimum amount of material was used. To reduce the effect of thermal noise in the tracker and to avoid

thermal runaway, the tracker undergoes cooling. Thermal runaway is when, due to insufficient or interrupted cooling, a positive feedback loop between leakage current and temperature causes a sensor to break. The pixel tracker is kept at -20°C , while the temperature of the strip tracker has been reduced from -15°C in earlier data taking years to -20°C in Run 3.

The pixel detector is the closest part of the detector to the interaction point starting at 2.9 cm away, therefore experiencing the highest particle flux. The current pixel detector consists of four layers in barrel and three layers in the endcap regions. Early in Phase I (2010-2025), there was one fewer layer in each region [55]. The pixel detector offers excellent granularity with 120 million pixels.

The strip detector surrounds the pixel detector. As the strip detector is further away from the interaction point, it does not have to be as granular as the pixel detector and so contains nearly 10 million channels [56]. Some of the modules in the strip detector are single-sided, while others are double-sided. The single-sided modules can only give a two dimensional position measurement ($r - \phi$ in the barrel, $z - \phi$ in the endcap), so the purpose of the double-sided modules is to recover the third dimension of the position measurement. This is done by placing a second strip module that is tilted by 100 mrad with respect to the first. The efficiency of the layers in the strip detector are $> 99\%$, while in the pixel detector the efficiencies are $> 99.5\%$ [57].

3.3.3 Electromagnetic Calorimeter

The CMS ECAL is a hermetic, homogeneous and fine grained detector made of lead tungstate (PbWO_4) crystals. There are 61,200 crystals in the barrel region split into 36 supermodules, while the endcaps contain 14,648 crystals grouped into 4 half disk “Dees”. The barrel gives the ECAL coverage of $|\eta| < 1.479$, with the endcap extending this coverage up to $|\eta| < 3$. In the barrel, the crystals are 23 cm long which corresponds to 26 radiation lengths, while in the endcap they are 22 cm long. The front face of the crystal in the barrel is $2.2 \times 2.2 \text{ cm}^2$ in the barrel, and in the endcap it is $2.86 \times 2.86 \text{ cm}^2$.

The barrel uses avalanche photodiodes (APDs) as photodetectors, with the endcaps using vacuum phototriodes (VPTs). Two different photodetectors are used as the APDs do not have sufficient radiation tolerance to be used in the endcaps (radiation damage increases with η), while the VPTs will not work in the barrel due to the longitudinal direction of the magnetic field there. The purpose of the ECAL is to accurately measure the energy and position of electrons, positrons and photons, as well as to measure the EM component of jets and contribute towards the contribution of energy sums (such as missing transverse energy).

Lead tungstate crystals were chosen due to their high density (8.28 g/cm^3), radiation hardness (up to 10 Mrad), short radiation length (0.85 cm), Moliere radius (2.19 cm) and scintillation decay time. The Moliere radius is defined as the radius of a cylinder that contains 90% of the energy deposited in an EM shower. This short Moliere radius means that if an




Figure 3.7: Schematic of the ECAL. From [58].

electron strikes the centre of a crystal, most of its energy will be deposited in that same crystal. The scintillation decay time of lead tungstate crystals means that 80% of the light is emitted within 25 ns [59]. This is compatible with the 40 MHz interaction rate of the current LHC and the high luminosity LHC (HL-LHC). The short radiation length and Moliere radius allows the ECAL to be compact while also providing sufficient material for the electrons and photons to interact with before exiting the ECAL, while radiation tolerance is required for the longevity of the detector in the harsh radiation environment.

The preshower of the ECAL is placed in front of the endcaps and is primarily to improve the discrimination between prompt photons (e.g. from $H \rightarrow \gamma\gamma$ decays) and photons from π^0 decays. This is achieved by measuring the transverse profile of EM showers after about three radiation lengths. A single incident photon will have a single peak in its transverse shower profile, while two photons that are close together in space will show a main peak with a secondary peak in the tail of the first. It is a sampling calorimeter with two thin layers of lead to generate the EM showers which are then detected by the following layer of silicon strip sensors.

For more detail on the ECAL please see Section 4.

3.3.4 Hadron Calorimeter

The main purpose of the hadron calorimeter (HCAL) is to measure hadronic energy and position. Like the ECAL, the HCAL is also a hermetic calorimeter, though its coverage extends up to $|\eta| = 5$ [60]. This is particularly important in the measurement of missing energy, which is an important signature of physics beyond the standard model.

The HCAL is comprised of 4 sections: the HCAL barrel (HB), HCAL endcap (HE), HCAL outer (HO) and HCAL forward (HF). The location of each of these sections can be seen in Figure 3.8. Both the HB and HE are sampling calorimeters, which alternate between layers of active and absorber material. This calorimeter design can be used for both electromagnetic and

hadronic calorimetry, but means that only part of the energy is measured, hence the addition of the HO and HF to measure energy from hadrons that travel beyond the HB or HE. Both the HB and HE use plastic scintillator as the active material and brass as the absorber. The light emitted from the plastic scintillator is then captured by wavelength shifting fibres. These fibres are fused to clear optical fibres, which transport the light to the photodetectors and front-end electronics. During Phase I, the HCAL photodetectors were upgraded from hybrid photodiodes to silicon photomultipliers [61].

The HO uses the return yoke and magnet as the absorber, while using the same plastic scintillators as the active material and fibres that are used in the HB and HE.

The HF is a Cherenkov calorimeter which uses a steel absorber and has quartz fibres which run longitudinally through the steel to collect Cherenkov light. This light primarily comes from the electromagnetic component of showers caused by particles interacting with the detector. There is a 5 mm spacing between quartz fibres in the HF, with fibres associated a particular region of the detector bundled together and photomultiplier tubes used to convert the optical signal into an electrical one.

All of the HCAL calorimeters share a common electronics chain after the phototransducers. The signals from the phototransducers are integrated over a period of 25 ns and digitised before being sent to the backend electronics. The front-end electronics uses a continuous readout, with readings for every bunch crossing transmitted to the back end. There, a HCAL trigger and readout board calculates trigger primitives and sends them to the calorimeter trigger system. This card retains the data samples while waiting for a decision from the level 1 (L1) trigger, the first trigger level in CMS. If the trigger accepts the event, the data is sent to a data concentrator card which will then send the information to the central DAQ. Otherwise, the data is discarded.

The energy resolution of the HCAL is $120\%/\sqrt{E(\text{GeV})}$ for neutral hadronic particles [63]. Unlike with charged particles, it is not possible to gather information on neutral particles from the tracker, making the energy resolution of neutral hadrons in the HCAL particularly important.

3.3.5 Muon System

Muons are less affected by the Coulomb force than electrons when they pass near nuclei due to their higher mass. This means that a muon with the same momentum as an electron will travel much further through a dense material, therefore the muon systems of CMS are placed further away from the interaction point than the ECAL and HCAL, with the subsystems interspersed between the plates of the iron return yoke which acts as a hadron absorber. The CMS muon system is based on gas ionisation and is capable of precise momentum and position reconstruction, having been a key focus in the design of CMS.

The technologies that make up the muon system are: drift tubes (DT), cathode strip chambers (CSC), and resistive plate chambers (RPC). DTs are used in the barrel, while CSCs




Figure 3.8: A schematic showing one quarter of the HB and HE layers of the CMS HCAL during Phase I. The segmentation in the top left is for the detector during Run 1 and Run 2, while the bottom right shows the segmentation of the detector in Run 3. The HO and the HF are not shown in this diagram. From [62].

are used in the endcaps. Both the CSCs and DTs provide excellent energy and timing resolution. The muon systems collectively cover a range of $|\eta| < 2.4$, with their relative positions shown in Figure 3.9. RPCs are used in both the barrel and the endcap. For CMS to measure muons in the magnetic field, the muon chambers must detect the muons at several points along their tracks.

In the barrel, the rate of muons is small and the neutrino background is usually low. Here, the magnetic field is mostly uniform with a strength of 0.4 T or less between the yoke segments. However in the endcaps, both the muon and background rates are higher. The magnetic field is also stronger and non uniform. CSCs are used here as they have a short drift path and hence a fast response time, tolerate the variations in the magnetic field and can also be finely segmented.

The RPCs provide excellent time resolution, which is required to correctly identify the bunch crossing of an event. This complementary system can provide a fast, independent trigger with a lower p_T threshold than the DT and CSC up to $|\eta| < 1.6$. One drawback of the RPCs is that their position resolution is somewhat limited, hence why RPCs are used in combination with DTs and CSCs which offer increased position resolution.




Figure 3.9: Illustration of one quadrant of the CMS muon system showing all three muon subsystems. DTs are shown in beige, while CSCs are shown in green and RPCs are shown in blue. From [64].

3.3.6 Trigger and Data Acquisition

CMS produces significantly more information than it can store, so it uses triggers to decide whether to keep or discard events. The CMS trigger is split into two parts: the L1 trigger and the high level trigger (HLT). The L1 trigger is the first stage and is made of a combination of FPGAs and custom ASICs. As the L1 trigger must provide fast decisions, it only considers information from the ECAL, HCAL and muon systems. The Phase I L1 trigger does not include information from the tracker. This is due to the impracticality of reading out the tracker data using the high speed link technology at the time the detector was built.

The L1 trigger considers a variety of information from the calorimeters and the muon detectors to search for physics objects of interest. These physics objects are reconstructed using algorithms which check events against sets of criteria such as minimum p_T and isolation requirements to form object candidates such as electrons, jets and muons. If an event meets at least one set of criteria, then the whole L1 trigger system sends the L1 accept signal to the trigger control and distribution system. This accept signal causes all the sub-detector information to be read-out to the HLT. Detector readout is handled by the data acquisition (DAQ) system. The L1 trigger reduces data rates from the 40 MHz interaction rate of the LHC to 100 kHz.

The HLT is a software based trigger which can access information from all the CMS subdetectors. In Run 1 and Run 2, it consisted of approximately 30,000 CPU cores and reduced the rate from 100 kHz to around 1 kHz [65]. In Run 3, the HLT uses a combination of CPU and GPU cores, with GPUs planned for at least 50% capacity of the HLT at the beginning of Phase II [66]. As with the L1 trigger, there are many different trigger paths in the HLT. The HLT uses simplified versions of offline reconstruction methods to reconstruct physics objects within the time constraints of the HLT, which has an average computing time of approximately 300 ms. Events are reconstructed in parallel and the steps are ordered such that the least complex filters occur first and events that would not pass any of the approximately 500 HLT trigger paths are rejected as swiftly as possible to minimise resource usage.

When an event is accepted by the HLT, the data is sent to the CERN Tier 0 computing centre. From here, the data can be reconstructed offline and saved into CMS data analysis formats.

3.3.7 Phase II Upgrades

In 2029, after the current Run 3 of data taking, the LHC will start running as the upgraded HL-LHC. This upgrade is designed to provide a peak luminosity of $5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ and an integrated luminosity of 3000 fb^{-1} by the end of HL-LHC, a ten times increase on the predicted luminosity at the end of the LHC [67]. The current schedule of the HL-LHC can be seen in Figure 3.10. This accelerator upgrade will provide the CMS experiment with an opportunity to collect significantly more data for both BSM searches and precision physics measurements, but will also require a significant detector upgrade. This detector upgrade will take place during long shutdown 3 (LS3) from 2026 to 2028.

This detector upgrade, known as the Phase II upgrade, aims to complete the minimum required upgrades to maintain performance with the most cost effective design. By maintaining the current physics acceptance for leptons, photons and other trigger objects, the CMS experiment will maintain the physics performance achieved in 2012. The requirements placed on the trigger to maintain this performance are the main drivers for the electronics upgrades. Radiation tolerance is also a consideration due to the significant increase in integrated luminosity for HL-LHC compared to the LHC.

To keep the physics acceptance as it was in 2012, the CMS experiment will add a L1 track trigger to identify tracks associated with calorimeter and muon tracker objects. The trigger will also use finer granularity data from both the muon and the calorimeter triggers. This will improve the transverse momentum resolution, particle matching and isolation. Doing this will require increasing the rate and latency of the L1 trigger, and also increasing the HLT output rate. Specifically, the L1 rate will increase to up to 1 MHz, with latency up to $12.5 \mu\text{s}$ [69], while the HLT output rate will be increased to 10 kHz [70]. This preserves the Phase I reduction in data rate between the L1 trigger and the HLT.



Figure 3.10: Timeline showing both LHC and HL-LHC operations. The Phase II upgrade for CMS will occur during Long Shutdown 3. From [68].

Without the L1 track trigger, the L1 rates for a given threshold would be significantly higher. Either a rate above 1 MHz would have to be supported or higher thresholds would need to be introduced, which would lead to a reduction in physics events recorded. The addition of the L1 track trigger also requires an increase in the latency to $12.5 \mu\text{s}$.

In Phase II, both the pixel and strip trackers will need to be replaced as the total dose provided by the HL-LHC is significantly larger than their designed radiation tolerance. The new tracker will have increased radiation tolerance and also increase in detector acceptance up to $|\eta| = 4$ [71].

A new sub-detector, called the minimum-ionising particle (MIP) timing detector (MTD) will be added to provide precise timing information to reconstructed tracks. This sub-detector will aid the tracker by reducing the probability of misidentified tracks or vertex merging, therefore increasing the reconstruction performance. As detailed later in Section 4.3.1.3, this MIP timing will also help to maintain the mass resolution provided by the ECAL for $H \rightarrow \gamma\gamma$ decays. At the start of Phase II, the MTD is expected to provide a timing resolution of about 30 ps, though this resolution is expected to increase to up to 60 ps by the end of Phase II due to radiation damage [72].

The upgrade of the ECAL barrel is detailed later in Section 4.3.1, while the upgrade of the ECAL and HCAL endcap is detailed in 4.3.2. In brief, the ECAL barrel will not suffer significantly from the increase in luminosity, so both the lead tungstate crystals and the APDs will remain while all the electronics will be replaced to satisfy the Phase II trigger rate and latency requirements. The endcap will suffer more significantly from radiation damage and

therefore the current ECAL and HCAL endcaps will be replaced by a new high-granularity calorimeter (HGCAL).

As with the ECAL barrel, the brass absorber and plastic scintillators of the HCAL barrel will work well enough in Phase II and therefore will not be replaced. The front-end electronics include silicon photomultipliers that were only installed recently during LS2 and will remain for Phase II. Like the ECAL, the current HCAL back-end electronics require upgrading to meet the new L1 trigger rate, and their design will be based on the design used for the ECAL barrel [73].

The muon systems will also require upgrades to maintain performance during Phase II. The current DTs, CSCs and RPCs will be upgraded to cope with the increased L1 latency and read-out rates. As with the other sub-detectors, this includes a replacement of the read-out electronics. In addition to this, two new muon subsystems will be added to extend the muon acceptance up to $|\eta| = 4$ to match the upgraded tracker (currently $|\eta| = 2.4$), as well as to improve granularity, resolution and redundancy at high η . The first of these new systems is the gas electron multiplier, which is a technology that has previously been used in other particle physics experiments such as LHCb. It is proven to work at hit rates of a few MHz/cm², which is compatible with the rate expected during Phase II. The second new system is the improved resistive plate chambers which have reduced recovery time and detector ageing effects compared to the current RPCs [74].

Chapter 4

Electromagnetic Calorimeter

4.1 Introduction

As already discussed in Section 3.3.3, the CMS ECAL is a fine-grained, hermetic and homogenous detector split into a barrel and two endcap regions. This chapter details the design and performance of the ECAL in both Phase I and Phase II. Phase I refers to the current detector, while Phase II refers to the upgraded detector needed during the high luminosity LHC (HL-LHC) era. Section 4.2 covers the Phase I ECAL, including the electronics used, current clustering methods, mitigation of anomalous signals (spikes), electron reconstruction and performance. Then, Section 4.3 details the upgrades necessary for the Phase II ECAL, as well as the expected performance. Particular attention is paid to clustering and spikes as these are important to describe for Chapters 5 and 6.

4.2 Current ECAL

The current ECAL readout system can be seen in Figure 4.1. When an electron or photon impacts on a lead tungstate crystal, there is an electromagnetic (EM) shower and light is produced. This light is converted into a photocurrent by the photodetector in that section of the detector. In the barrel, avalanche photodiodes (APDs) are used while in the endcaps vacuum phototriodes (VPTs) are used. These photodetectors were chosen as CMS required photodetectors with gain and a high quantum efficiency due to the low intrinsic light yield of the lead tungstate crystals.

The photocurrent is amplified by the multi-gain pre-amplifier (MPGA) ASIC. This is a charge-sensitive amplifier with three output gain values: $\times 1$, $\times 6$ and $\times 12$. The gain value used depends on the energy of the electron or photon that caused the shower, with the three different gain values allowing greater precision for lower energy particles while also avoiding saturation for high energy particles. In the barrel, the dynamic range is 40 MeV to 1.7 TeV.




Figure 4.1: Current electronics of the CMS ECAL. From [73].

A multi-channel analogue-to-digital converter (ADC) converts the analogue signal to a digital signal. It has a resolution of 12 bits per gain value and a sampling frequency of 40 MS/s. The final on-detector step is the front-end (FE) card, which controls the data pipeline and trigger primitive generation. The FE card calculates and transmits trigger primitive energies at 40 MHz and holds crystal data in a buffer until an L1 accept signal is received. Triggering on electrons and photons occurs through the two levels of the CMS triggering system, as described in Subsection 3.3.6. The first level of the trigger is the level 1 (L1) trigger, with the second level being the high-level trigger (HLT). The current L1 trigger data granularity is in blocks of 5×5 crystals (i.e. one trigger tower). The FE card generates energy sums for each trigger tower, stores data until an L1 trigger decision has been made, and sends data that has been accepted by the trigger to the data acquisition system.

Optical fibres then read out data to the off-detector readout cards. There are three separate off-detector cards: one for trigger functions, one for DAQ functions and finally one for clock and control functions. Comparatively, in Phase II there will be a single off-detector board that will combine all three of these functions. In the current detector, data is read out with a selective readout algorithm, which only reads out information from crystals in high interest regions or those that meet a given energy threshold of $E_T > 2$ GeV to reduce output data bandwidth. This contrasts with Phase II during which all crystal data can be read out.

4.2.1 Current Electron Clustering

In the ECAL, a combination of photon conversion and electron bremsstrahlung causes EM showers to have an increased spread in the ϕ direction due to the solenoidal magnetic field in CMS. The clustering methods used have changed over time, only the methods used for Run 2 are described here as these are relevant to the analysis based on Run 2 data described in Chapters 8 and 9.

The steps involved in the offline clustering are:

- Clusters are formed by grouping crystals with energies over a given threshold. The threshold is set to be two or three times bigger than the electronic noise level expected for the crystals [75]. In 2018, the noise level in the barrel was around 100 MeV at $|\eta| = 0$, increasing to approximately 130 MeV at $|\eta| \approx 1.5$. In the endcap, the noise was approximately 0.1 GeV at $|\eta| \approx 1.5$, increasing to around 3 GeV at $|\eta| = 3$ [76]. The crystal that contains the highest energy in a given region and also has $E_T > 1$ GeV is defined as the seed crystal around which the cluster is formed.
- A supercluster is created by adding energy deposits in a given geometric area around the seed cluster. This includes energy from photon conversions and bresmsstrahlung [77].
- Trajectory seeds from the pixel detector are considered to determine whether they are compatible with the position of the supercluster, which would indicate an electron. If they are compatible, then this information is used to seed the Gaussian sum filter (GSF) tracking step.
 - While a Kalman filter (KF) is linear least-squares recursive estimator [78], the GSF extends this to a non-linear method. Rather than approximating the energy loss distribution with a single Gaussian (as in the KF), the distribution in the GSF is comprised of multiple Gaussians. The GSF requires that all of the probabilities that are involved in the reconstruction of tracks are described by a sum of Gaussians [79, 80].
- Generic tracks with $p_T > 2$ GeV that are not specifically linked to an electron are also reconstructed.
- Tracks from photons that undergo pair production of e^+e^- are identified using a dedicated algorithm [81].
- All the above information is imported into a particle flow (PF) algorithm. The PF algorithm is designed for global event reconstruction, where it takes information from all the subdetectors (i.e. tracks and clusters) and uses this information to identify all of the particles in an event and their properties [82].

- The PF algorithm identifies electron and photon objects in the event, with electrons starting from GSF tracks and photons starting from superclusters. The refined superclusters are defined at this step as the final list of linked ECAL clusters for each electron or photon.
- Loose selection criteria are applied to create an unbiased dataset, which is used as the starting point for most analyses that involve electrons and/or photons. This includes the analyses detailed in Chapters 8 and 9. Objects that fail these criteria are then considered by the PF framework as a potential neutral hadron, charged hadron or nonisolated photon.

In the L1 trigger, electrons and photons are clustered starting from 5×5 crystal arrays (one trigger tower) in the barrel, with a different pattern in the endcap due to the crystal layout there. The seed tower is defined as the tower with the largest transverse energy deposit that is above the 2 GeV threshold. Again, steps are taken to account for energy losses due to bremsstrahlung. This is done by adding surrounding trigger towers to the cluster if the energy in that particular array is at least 1 GeV within a 2×5 trigger tower region in η and ϕ [83]. As there is no tracker information available to the L1 trigger, it is not possible to distinguish between electrons and photons.

To improve the identification of electrons and photons, within the limits of information, resources and time available to the L1 trigger, requirements are placed on:

- Isolation of the energy deposited
- A shape variable, which considers the energy split between the seed tower and the rest of the cluster
- The energy deposited in the HCAL downstream from the seed tower
- Variables such as $E_{2 \times 5}/E_{5 \times 5} > 90\%$, which requires that the energy in the cluster is mostly deposited in a small window around the seed crystal (with an extension in ϕ on account of the bending due to the magnet). See Section 4.2.2 for more details on one of these variables.

Clustering in the HLT begins with the L1 candidate, with the HLT using the same method as the offline reconstruction to build superclusters. The processing time is limited in the HLT so clusters are only created around crystals that seeded clusters in the L1 trigger. These crystals are known as seed crystals and are defined by having the greatest energy deposit in that region and also having an $E_T > 1$ GeV. As tracking information is available to the HLT, it is possible to distinguish between electrons and photons here. Electron candidates are defined as superclusters with a compatible reconstructed track. Identification and isolation criteria (such as $E_{2 \times 5}/E_{5 \times 5}$ mentioned above) are applied for both electrons and photons using information

from the ECAL, HCAL and the tracker. Thresholds are also set on the transverse energy in the ECAL as measured at the HLT.

The objects used by the HLT and the offline clustering are similar but there are differences in their calibrations, which can result in an object having a different energy offline than it did in the HLT. This is due to the online HLT laser corrections (see Section 4.2.3.2 for details) being updated less frequently than they are offline. Also, all HLT electrons are ECAL driven, while offline electrons may also be tracker driven, which is useful for low energy or nonisolated electrons [84].

4.2.2 Current Rejection of Anomalous APD Signals

Spikes are caused when the silicon of the APDs is directly ionised by a heavily ionising particle such as a proton or a pion, rather than from light from the lead tungstate crystals as in an EM shower. The rate of spikes increases with the intensity of the proton beams. There are two main methods to reject these anomalous signals in the current detector.

Firstly, there is online rejection, which stops spikes from triggering CMS. Without online rejection, spikes would contribute considerably to the L1 bandwidth at high luminosity as many spikes meet the energy threshold criteria set by the L1 electron and photon triggers, meaning that electron and photon trigger thresholds would have to be raised significantly which would compromise physics results. When spike rejection was first used in CMS in 2010, it reduced the rate of a 20 GeV trigger by a factor of five. In the current running of the detector, to reduce rates by a factor of five without spike rejection the trigger thresholds would have to be increased by at least 15 GeV.

In the L1 trigger, a strip fine-grained veto bit (sFGVB) is used to determine whether an energy deposit is EM shower-like or spike-like. A strip refers to a 5×1 crystal array extended in ϕ , and for the sFGVB, the number of hits above a threshold in a strip are summed up. For an EM shower, we expect to see multiple crystals with an energy above the threshold as some energy will also be deposited in the surrounding crystals, while for a spike we only expect one crystal to be above the threshold. Accordingly, the strip gives an output of one if there are multiple crystals above the threshold (EM shower-like), or zero if no crystals or only one crystal is above the threshold (spike-like). The crystals above the threshold do not have to be adjacent. This is illustrated in Figure 4.2. The results from the 5 strips that make up the 5×5 crystal matrix are combined together (with a logical OR) to give a single result for the sFGVB to send on to the L1 trigger, energy deposits above 16 GeV are currently rejected if the sFGVB is zero (spike-like).

This sFGVB method has been used in CMS since 2011, with optimisations to the threshold as needed to offer optimal performance under the given detector conditions (e.g. increased noise and pileup). In 2011 and 2012, the sFGVB rejected over 95% of spikes with $E_T > 8$ GeV. Over 99.5% of well-identified electrons with $E_T > 20$ GeV passed the sFGVB, so it only has a small




Figure 4.2: Schematic showing how the sFGVB works in the case of an EM shower (left) and a spike (right). From [85].

effect on electron trigger efficiency [85]. It is important to note that the spike contamination is still sizeable, and that it is not possible to kill low energy spikes due to the difficulty in distinguishing them from EM signals with the current electronics. The fraction of ECAL trigger primitives that were due to spikes in one run in 2018 can be seen in Figure 4.3.

Spikes that have not been rejected by the sFGVB can be dealt with offline or in the HLT. In both cases, this is done using the Swiss-cross algorithm. As with the sFGVB, the Swiss-cross utilises the difference between spikes and EM showers in energy deposits across the surrounding crystals. The Swiss-cross algorithm works on the variable $1 - E_4/E_1$, where E_1 is the energy deposited in the crystal under consideration, and E_4 is the sum of the energy deposited in the four crystals that share a side with the central crystal. This is illustrated in Figure 4.4, while Figure 4.5 shows the difference in the distribution of the Swiss-cross variable for data and MC that does not include spikes. A cut on this variable at 0.95 rejects over 99% of spikes with $E_T > 10$ GeV and has a negligible impact on EM shower selection efficiency [85]. This is due to the fact that a spike will deposit all of its energy in E_1 , with noise in the surrounding crystals, while an EM shower impacting on the centre of a crystal will only deposit 80% of its energy in E_1 , with the remaining 20% deposited in the remaining crystals.

In addition to the Swiss-cross algorithm, a timing cut is added to remove any remaining spikes offline. This is based on the fact that spike timing only depends on the electronics shaping time and does not include the response time of scintillation light as seen for an EM shower. Therefore, a cut on the signal timing is placed to only keep events that are ± 3 ns which is over 90% efficient at rejecting spikes, again with negligible impact on EM shower efficiency [87]. For EM showers with $E_T > 1$ GeV, the timing resolution is less than 1 ns. The timing cut is not used online due to concern regarding the possibility of a hardware issue causing an unexpected timing shift at the HLT, which would then cause real events not to be recorded.

4.2.3 Current Electron Reconstruction

The offline reconstruction of the energy deposited by an EM shower is:

$$(4.1) \quad E_{e,\gamma} = \sum_i [A_i(t) \cdot L_i(t) \cdot C_i(t)] \cdot G(\eta) \cdot F_{e,\gamma} + E_{preshower}$$




Figure 4.3: Fraction of ECAL trigger primitives in the barrel that are due to spikes as a function of E_T . Data was collected from a single run in 2018 with a peak pileup of 50. The black line shows the result when the online pedestal values are used, while the dashed line shows the result with updated pedestal values. This is required as the pedestal values drift with time. From [86].




Figure 4.4: Illustration defining the Swiss-cross variables E1 and E4.




Figure 4.5: The distribution of the Swiss-cross variable for 2010 data and MC events which do not include spikes. The data with values above one for $1 - E_4/E_1$ are due to noise, where the reconstruction method used in 2010 could give negative reconstructed amplitudes for noise pulses. From [85].

where i indicates the sum of crystals involved in the supercluster, t is time, $A_i(t)$ is the signal amplitude in ADC counts, $L_i(t)$ is a correction for crystal transparency loss, $C_i(t)$ are channel-to-channel intercalibrations that ensure a homogeneous energy measurement through the detector, $G(\eta)$ converts from ADC counts to GeV, $F_{e,\gamma}$ are clustering corrections that differ depending on whether the particle is an electron or a photon, and finally $E_{\text{preshower}}$ is the energy estimate from the preshower. $G(\eta)$ is derived by calibrating the ECAL energy measurements using a Standard Model candle, specifically Z decays.

The electron reconstruction can be split into three sections: signal amplitude reconstruction, crystal ageing corrections and crystal inter-calibrations. The strategy for each of these calibration steps is explained below.

4.2.3.1 Signal amplitude reconstruction

Since Run 2, amplitude reconstruction is performed using a multifit algorithm that fits the signal pulse by using a single in-time pulse and up to nine out-of-time pulses. The pulse shape was fixed, while the signal amplitudes and the time offset of the out-of-time pulses were the outputs of the algorithm. This was changed from the Run 1 weights algorithm that had worsening performance with increasing pileup. By contrast, the multifit algorithm performance was constant with respect to out-of-time pileup for conditions experienced in Run 2 [88]. An illustration of the multifit method is shown in Figure 4.6.

To extract the amplitude of a signal pulse and generate a single crystal energy measurement, both the signal shape and the pedestal (baseline of the signal from which the amplitude is




Figure 4.6: Multifit method during Phase I based on CMS simulation for an event with 20 pileup and a bunch spacing of 25 ns. The black circles represent the digitised samples. The red histogram shows the in-time pulse, while the yellow, green and light blue histograms represent three out-of-time pulses. The dark blue histogram is the sum of both the in-time and out-of-time contributions. From [89].

measured) are needed. Both of these variables are monitored and corrected continuously to ensure proper modelling. Radiation damage to the lead tungstate crystals in the ECAL causes crystal light transmission to decrease, therefore causing changes in the signal shape. To account for this, the signal pulse shapes were recomputed every 3 to 4 fb^{-1} during Run 2.

By contrast, the pedestal is mostly dependent on the FE ASICs and the APDs. Both the pedestal mean and the pedestal RMS (noise) change with time. The changes in mean pedestal value are due to the ageing of very front-end (VFE) application-specific integrated circuits (ASICs) and occur in both the barrel and the endcap of the ECAL. These values are monitored by pedestal data recorded every 40 minutes using the same laser events that are used to monitor radiation damage (see Section 4.2.3.2 below). The updates of the pedestal mean and RMS values are necessary to maintain the performance of the template fit used for amplitude reconstruction, as well as to maintain performance in mitigating spikes in the sFGVB algorithm.

4.2.3.2 Crystal ageing corrections

The CMS ECAL suffers from significant radiation rates during proton-proton collisions. This radiation rate is not constant throughout the detector, and is more significant at high η values. Though this radiation does not affect the scintillation mechanism, it does cause a reduction in crystal light transmission. As the crystals in the endcap are subject to higher radiation rates, they also suffer from an increased loss in crystal light transmission relative to the barrel. Some recovery of the transparency occurs via partial annealing of the lead tungstate crystals at room temperature, though over time there will be a net loss in transparency due to LHC radiation.

The transparency loss is monitored on a per-crystal basis using a dedicated laser monitoring system which injects light during the LHC abort gap, a large proton-free gap between the bunches of the LHC. Every 40 minutes, a measurement of the relative response of every individual crystal is obtained, which allows time and crystal-dependent transparency corrections to be computed. The relationship between the crystal's response to laser light compared to an electromagnetic shower is modelled as:

$$(4.2) \quad \frac{S}{S_0} = \left(\frac{L}{L_0} \right)^\alpha$$

where S is the measured electromagnetic response, S_0 is the reference electromagnetic response, L is the measured laser response, L_0 is the reference laser response before LHC irradiation, and α is an η dependent variable. In the barrel, α is measured to be around 1.5, while in the endcap the value varies between 0.6 and 1.1 [90]. The measured laser response and the value of α are used to correct the measured electron energy to the reference value. The relative response of different regions of the ECAL to laser light over time can be seen in Figure 4.7.

This correction is not linear because the light measured from the laser or scintillation light depends on the transmission coefficients of the crystal. This gives an exponential dependence on the path length which is different between the two, as laser light has to travel the full length of the crystal to reach the photodetectors. Additionally, the transmission will be reduced by irradiation [92].

4.2.3.3 Crystal inter-calibrations

As different crystals within the ECAL have different intrinsic light outputs, inter-calibration coefficients are used to ensure a homogeneous response across the detector. These calibration coefficients are calculated per crystal by using standard candles, information from the tracker and by exploiting the azimuthal symmetry of minimum bias events in the detector. The standard candles are well-known resonances that are used to calibrate the energy response, with the resonances $\pi^0 \rightarrow \gamma\gamma$ and $Z \rightarrow ee$ specifically used for these ECAL calibrations.

In $\pi^0 \rightarrow \gamma\gamma$ inter-calibrations, all diphoton candidates are used to create an invariant mass distribution. The π^0 mass from this measured spectrum is compared with the expected value [5]




Figure 4.7: Relative response to laser light of the CMS ECAL over 8 years of data taking, given for different η ranges. The instantaneous luminosity of the LHC is also shown. From [91].

and the mass shift between the two is used to generate an energy correction for each individual crystal. As this is done by looking at multiple photons, this calibration algorithm iteratively corrects each crystal output according to the shift in the π^0 mass.

Similarly, the $Z \rightarrow ee$ calibration uses dielectron pairs to create an invariant mass spectrum and compares the measured Z mass with the expected value [5]. As this resonance has a smaller cross-section than $\pi^0 \rightarrow \gamma\gamma$, it was only introduced during Run 2 as an inter-calibration method due to the increased integrated luminosity per year. In this instance, the single crystal calibration constants are calculated by maximising a likelihood, which compares the predicted MC mass distribution to the one measured in the data.

The symmetry of the detector in the ϕ direction is also used to calibrate the detector, as there should be no bias in events towards a particular region of ϕ . This calibration method is accurate to a few percent, with the accuracy limited by systematics. This is specifically due to the variation in the distribution of the material in front of the ECAL [93].

The final calibration method exploits the independent measurement of transverse momentum for electrons. To combine the transverse momentum measurement from the tracker with transverse energy from the ECAL, high energy electrons from $W^\pm$ and Z decays are selected via identification, isolation and kinematic cuts. This combination of cuts gives electrons with energy approximately equal to their momentum due to the negligible effect of the electron




Figure 4.8: Intercalibration precision for the $Z \rightarrow ee$ calibrations, E/p calibrations, π^0 calibrations and the three method combined. From [95].

mass at high energy, therefore the calibration algorithm iteratively assigns coefficients to each individual crystal so that the ratio of energy to momentum for this carefully selected dataset is equal to 1 [94].

These calibrations are then combined to give an overall calibration. The precision of this combined calibration and the calibration for each individual method can be seen in Figure 4.8.

4.2.3.4 Preshower Calibration

In the endcaps, some fraction of the shower energy is deposited in the preshower, which also requires precise calibration. The fraction of energy deposited in the preshower decreases with increasing electron or photon energy but is usually around 7%. To keep the combined energy resolution of the endcap and the preshower below 0.4%, the preshower must have a channel-to-channel calibration precision of 5%.

All of the preshower sensors were calibrated using cosmic rays prior to installation, while in situ they are calibrated using charged pions and muons with $p_T > 1$ GeV. This momentum requirement gives a signal-to-noise ratio greater than 10. To calibrate each channel, the pulse height distribution is fitted with the convolution of a Landau distribution and a Gaussian function, with the position of the peak from this fit giving the calibration value. The cosmic calibrations have a precision of 2.5%, while the in-situ calibrations have a precision of 2.2%. These calibration values are below the required 5% calibration precision for the desired energy resolution [96].




Figure 4.9: Left: the impact of energy scale and resolution for the ECAL barrel from the use of superclustering and energy corrections compared to the sum of energy deposits in a 5×5 crystal matrix. Right: impact of energy scale and resolution for the ECAL endcap from the use of superclustering, addition of energy from the preshower and energy corrections compared to the sum of deposits in a 5×5 crystal matrix. Both show the reconstructed Z mass. From [97].

4.2.4 Performance of the Current ECAL

Offline data analysis has found that a single photon trigger used during 2012 that has a threshold of $E_T = 20$ GeV and no pre-scaling (i.e. all events passing this trigger were recorded) was more than 99% efficient for transverse energies greater than 40 GeV [97]. This trigger provided full efficiency for the $H \rightarrow \gamma\gamma$ search in that year, with similar triggers and efficiencies in subsequent data taking years.

The effect of superclustering on the energy scale and resolution for $Z \rightarrow ee$ can be clearly seen in Figure 4.9. When the energy from a 5×5 crystal matrix is used, the peak in the mass spectrum is at noticeably lower reconstructed mass than the Z mass reconstructed from the supercluster energy, with a much larger low-mass tail. The energy corrections applied to the supercluster shifts the mass spectrum closer to the review of particle physics mass of the Z boson, with the width of the spectrum also decreasing. In the endcap, the addition of energy recorded by the preshower (ES) makes a considerable improvement to the energy scale and resolution.

Figure 4.10 shows the electron reconstruction efficiency as a function of η in 2017. There were no significant detector issues in this year, so 2017 can be considered to be representative of the whole of Run 2. The reconstruction efficiency is a largely stable value of over 95%, with reductions in efficiency only occurring in the transition region between the barrel and the endcap at $1.44 < |\eta| < 1.57$ and is not considered in physics analyses. The efficiency also increases slightly with increasing p_T .

Figure 4.11 shows the electron resolution as a function of η for Run 2. The central barrel




Figure 4.10: Electron reconstruction efficiency as a function of η , with the data/MC ratio shown in the bottom panel for data collected during 2017. The error bars denote combined statistical and systematic uncertainties. The efficiency is measured using tag and probe with $Z \rightarrow ee$ events. From [98].

region offers the best resolution for all years, with the resolution generally worsening with increasing η . 2016 offers the best resolution, while values for 2017 and 2018 are similar in the barrel. The differences in the endcap are more pronounced, with 2018 having a poorer resolution as expected due to detector ageing.

The performance of the ECAL played an important role in the discovery of the Higgs boson. Five different decay channels were considered in the analysis, with the combination of two channels which rely on ECAL measurements, namely $H \rightarrow \gamma\gamma$ and $H \rightarrow ZZ \rightarrow 4\ell$, corresponding to a 5σ observation. These two decay modes offered the best mass resolution, leading to a resolution of around 1% on the Higgs mass in the best $H \rightarrow \gamma\gamma$ category [2]. The mass distribution from the diphoton decay channel can be seen in Figure 4.12.

4.3 ECAL Upgrade for Phase II

4.3.1 Barrel Upgrade

The main motivation for the upgrade of the CMS ECAL in Phase II is to maintain the current physics performance in the more challenging conditions presented by the HL-LHC, particularly in terms of higher luminosity and increased pileup. In terms of the ECAL barrel specifically,




Figure 4.11: Electron resolution for Run 2 data as a function of η , split into 2016 (red), 2017 (green) and 2018 (blue). The grey region denotes the gap between the barrel and the endcap. From [95].

Electrons from $Z \rightarrow ee$ decays are used which includes both low and high bremsstrahlung electrons. The resolution is extracted from an unbinned likelihood fit using a Voigtian as the signal model.

the upgrade requirements can be split into the following:

- Accommodate the increased HL-LHC trigger requirements
 - Trigger latency will increase from $4 \mu\text{s}$ to $12.5 \mu\text{s}$
 - Level-1 trigger rate will increase by a factor of 7.5 to 750 kHz
- Provide single crystal information to the trigger to improve clustering and spike rejection
- Limit noise due to radiation damage in the photodetectors by cooling to 9°C (was 18°C in Phase I)
- Provide shorter pulse shaping by increasing the sampling rate to 160 MHz (a four fold increase) to improve timing resolution for out-of-time PU mitigation using the multfit algorithm (see Subsection 4.3.1.2) and spike rejection.

To meet the Phase II trigger requirements, all the current on and off detector electronics must be replaced. While the Phase I ECAL provides L1 trigger information in the form of trigger primitives in blocks of 5×5 crystals, the Phase II system will provide single crystal information to the L1 trigger. The upgrade will also provide tracking information to the L1




Figure 4.12: Invariant mass spectrum of the $H \rightarrow \gamma\gamma$ decay channel. Each event is weighted by the signal divided by the sum of the signal and background. The solid red line indicates the signal and background fit of the data, while the red dashed line indicates the background fit. The yellow and green bands indicate the ± 1 and ± 2 standard deviation uncertainties in the background estimate respectively. From [2].

trigger, while the Phase I L1 trigger does not receive any information on tracking. This increased granularity and tracking information will allow EM showers to be matched to tracks and will also improve isolation compared to the current detector. This combination of trigger upgrades will significantly reduce the proportion of background events recorded at the calorimeter trigger thresholds required for precision Higgs boson studies.

To meet the improved timing resolution and noise filtering requirements, the VFE card must be replaced. The VFE card provides pulse amplification, shaping and digitisation. The increased sampling rate will enable spikes to be rejected based on their pulse shape compared to an EM signal. As shown in Figure 4.14, the EM and spike pulse shapes are similar in Phase I. However, with the timing improvements in Phase II, the spikes will have a pulse shape that is shorter and starts sooner than that of an EM shower. If the spikes are not rejected at an efficiency of $> 99\%$ in the trigger in Phase II, then they would saturate the trigger at the energy thresholds required by CMS for many physics analyses. The expected rate of events without any spike rejection is shown in Figure 4.13, with the L1 bandwidth shown for comparison.

Every component of the Phase I ECAL has been studied to determine whether or not it will be able to maintain the required performance during Phase II. For the lead tungstate crystals,




Figure 4.13: Expected event rate in the CMS ECAL during Phase II (HL-LHC) as a function of E_T . No spike killing is applied and 200 pileup is assumed. The event rate is shown for an integrated luminosity of 0 (undamaged), 300, 1000, 3000 and 4500 fb^{-1} , with the Phase I (LHC) and Phase II (HL-LHC) L1 bandwidths also shown. From [73].




Figure 4.14: Left: pulse shapes of spikes (black) and EM showers (grey) in the Phase I CMS ECAL. Right: pulse shapes of spikes (red) and EM showers (blue) with the Phase II CMS ECAL. The Phase II pulse shapes are given by convoluting known pulse shapes of spikes and EM signals with the measured response of prototype Phase II electronics. From [73].




Figure 4.15: Expected light output of the ECAL barrel crystals during Phase II relative to conditions at the start of CMS running. The relative light output is shown as a function of η for a range of integrated luminosities from 300/fb in red (at the end of Phase I) to 4500/fb in blue (expected at the end of Phase II). Based on 50 GeV photon showers using Phase II electronics cooled to 9 °C. From [73].

the primary concern is radiation-induced ageing. As mentioned in Section 4.2.3.2, while the scintillation mechanism of the crystals is not affected by radiation, radiation damage causes defects in the crystals, which in turn diminishes the crystal transparency, therefore reducing the light output. In the current ECAL, this effect is carefully monitored and corrected as discussed in Section 4.2.3.2.

The predicted Phase II light output seen in Figure 4.15 was calculated using Monte Carlo simulations and test beam data. Figure 4.15 shows that even after 4500 fb⁻¹, the highest η regions of the barrel (which are most affected by radiation damage) will still have around 35% of the light output that they did at the start of data taking. This is higher than the light output currently found in the highest η regions of the endcap, as shown in Figure 4.7, giving confidence that the crystals in the barrel can remain for the entirety of Phase II.

The predicted Phase II energy resolution can be seen in Figure 4.16. The energy resolution decreases with increasing integrated luminosity. At any given integrated luminosity, the $|\eta| = 0$ region offers a better resolution than the $|\eta| = 1.45$ region as the low η region will have suffered less radiation damage. For showers with energies above 50 GeV, the constant term dominates the energy resolution.

There are two ways in which the APDs can be damaged by radiation. The first is when gamma




Figure 4.16: Expected energy resolution as a function of incident particle energy for a range of integrated luminosities from 300/fb (at the end of Phase I) to 4500/fb (expected at the end of Phase II). Shown for $|\eta| = 0$ (left) and $|\eta| = 1.45$ (right). From [73].




Figure 4.17: Schematic showing the ECAL barrel Phase II electronics for a 5×5 crystal matrix. From [99].

rays cause surface defects, which in turn increase surface current and reduce the quantum efficiency of the APDs. The second is radiation damage from hadrons, chiefly neutrons, which causes bulk damage, increasing the bulk current. It is this second form of radiation damage that is the main concern during Phase II as the electronic noise depends on the square root of the bulk current. Reducing the operating temperature from 18°C to 9°C will keep the noise to a manageable level of around $200 \mu\text{A}$ in the highest η region of the barrel. Without this reduction in temperature, the electronic noise would be a factor of 10 larger, corresponding to 1 GeV after 3000/fb.

4.3.1.1 Readout Electronics

The Phase II readout electronics can be seen in Figure 4.17. As previously mentioned, both the lead tungstate crystals and the APDs will remain for Phase II so the very first part of the readout chain is unchanged. The new pre-amplifier ASIC is the calorimeter trans-impedance amplifier (CATIA). It provides minimal pulse shaping and has two output gain values: $\times 1$ and

$\times 10$. Next is the Lisbon-Torino ECAL data transmission unit (LiTE-DTU), a data conversion, compression and transmission ASIC. It uses two 12-bit ADCs with a sampling frequency of 160 MS/s (4 times higher than currently). The LiTE-DTU also provides lossless compression, as well as utilising a look-ahead algorithm which prevents samples from the two different gains from being mixed in a single time frame.

The combination of 5 CATIA chips and 5 LiTE-DTU chips makes a single VFE board. In addition to the roles carried out by the CATIA and LiTE-DTU, the VFE also calibrates the ADC values using a reference voltage provided by the CATIA, as well as providing readout of temperature sensors for components such as CATIAs and APDs. There is also a low-voltage regulator which provides the correct voltages to the ASICs and the FE board. It is a radiation tolerant, magnetic field tolerant and low-noise design that will readout APD leakage current. APD leakage current is the sum of the dark current, noise and photocurrent which limits the performance of the APDs.

The FE is designed to stream single crystal granularity data at 40 MHz to the off-detector electronics via high-speed optical links. This is not possible with the architecture of the current electronics, which streams data in 5×5 crystal blocks to the trigger. The new FE will send clock signals directly to the VFE and will also be responsible for monitoring APD dark current (electric current generated by the APDs in the absence of light from the crystals).

After the FE, there is the off-detector barrel calorimeter processor (BCP) board. This board is responsible for data acquisition and trigger functions, including the generation of trigger primitives. Each BCP board receives signals from 600 crystals and has one commercial FPGA which will run all of the trigger primitive algorithms [99].

4.3.1.2 BCP Algorithms

The planned features of the BCP are:

- Spike rejection
- Conversion of digitised pulse data into transverse energy
- Precise timing measurement
- Basic clustering of localised energy
- Encoding of trigger primitives and sending them to the L1 trigger
- Generation and transmission of clock and control signals to the FE boards
- DAQ functions

Of these, the spike rejection, transverse energy conversion, timing measurement and clustering will be completed via high-level synthesis (HLS) algorithms. High-level synthesis refers to




Figure 4.18: Illustration of the Phase II multifit algorithm at L1. The red curve illustrates the pulse in the bunch crossing under consideration, while the green curve shows a pulse from the previous bunch crossing and the orange shows a pulse from the subsequent bunch crossing. The grey lines indicate the nominal timing of each bunch crossing. The combined waveform from all three pulses is shown in blue, with the 12 samples considered by the multifit algorithm shown using blue diamonds. Y-axis units are arbitrary.

algorithms that are written in a high-level language such as C++ and then converted to register transfer level (RTL) languages through a tool such as Vivado HLS [100]. HLS is discussed in more detail in Section 6.1. The remaining features will be directly coded in an RTL.

The conversion from digitised pulse data into transverse energy and the precise timing will be combined into one multifit algorithm, similar to the multifit algorithm that was used in Run 2 (see Section 4.2.3.1). The multifit algorithm uses a predetermined pulse shape to determine the amplitude and time of the pulse in any given bunch crossing by considering 12 consecutive samples. An illustration showing three pulses in a single channel can be seen in Figure 4.18. The multifit algorithm also considers contributions from the bunch crossing immediately before the one in question, as well as the bunch crossing immediately after [101].

There are two spike-killing algorithms under consideration, with the intention of fitting both algorithms on the FPGA to provide redundancy. The first spike-killing algorithm is the Swiss-cross algorithm, which has previously been used in CMS both offline and in the HLT. Details of the Swiss-cross algorithm can be found in Section 4.2.2. The Swiss-cross algorithm does not require an increase in sampling rate to work, but it is sensitive to noise in the surrounding crystals for low-energy deposits. The single crystal granularity offered by the Phase II upgrade means that this algorithm can now be used in the L1 trigger. The expected performance of the




Figure 4.19: Expected performance of the Swiss-cross algorithm for different E_T thresholds and four integrated luminosities: 300 fb^{-1} (top left), 1000 fb^{-1} (top right), 3000 fb^{-1} (bottom left) and 4500 fb^{-1} (bottom right). From [73].

Swiss-cross algorithm can be seen in Figure 4.19

The second spike-killing algorithm is the linear discriminant (LD) algorithm. Unlike the Swiss-cross algorithm, the LD method only requires information from the crystal under consideration. It uses a second-order polynomial to calculate the LD value which is used to distinguish between spikes and EM showers. The LD works well with lower energy particles, but it requires 160 MHz sampling and its performance will be degraded for any noisy channel that must be digitised at a lower sampling rate of 80 MHz. The LD is defined as

$$(4.3) \quad LD = R_+ - \sum_{i=0}^O p_i R_-^i$$




Figure 4.20: Expected signal efficiency of the LD algorithm in minimum bias interactions at 200 PU. From [73].

where

$$(4.4) \quad \begin{aligned} R_+ &= \frac{a_{n+1}}{a_n} \\ R_- &= \frac{a_{n-1}}{a_n} \end{aligned}$$

The variables a_{n-1} , a_n and a_{n+1} are defined as the three consecutive samples around the peak (a_n is the peak as shown in Figure 4.14 (right)). The weights p_i are chosen to give an LD value of -0.5 for a spike and 0 for an EM signal. Figure 4.21 shows the distribution of the LD variable for spikes and EM showers for two different amplitudes. The order of the polynomial is denoted by $\mathcal{O}$ [102]. The expected signal efficiency of the LD algorithm is nearly 100%, as shown in Figure 4.20.

The event rate after applying one of the two spike killing algorithms can be seen in Figure 4.22. There is a significant reduction in the event rate compared to the case without any spike killing applied, as was shown in Figure 4.13.

Finally, the clustering algorithm searches the 600-crystal region served by a single BCP looking for local maxima to seed a cluster. A seed crystal must be above a given noise threshold, have an energy greater than the neighbouring crystals and the seed must not be flagged by either spike-killing algorithm. The number of bits allocated to each cluster can be seen in Table 4.1. The difference between the original bit distribution and the proposed one lies in the number of bits allocated for the number of crystals and the spike flag. The crystal matrix size is not expected to exceed 3×5 , so only four bits are required for the number of crystals in the matrix.

Chapter 5 details a test beam study regarding clustering, while Chapter 6 gives results on HLS testing of both of the spike-killing algorithms and the clustering algorithms. In particular,




Figure 4.21: Distribution of the LD variable for spikes (red) and EM showers (blue). The left plot shows the distributions for plots with amplitudes 10 times greater than the noise, while the right plot shows the distribution for amplitudes 50 times greater than the noise. From [73]. The double peak structure of the APD signal on the right plot is due to the LD value varying with the phase of the spike as the spike simulation has a realistic energy and spike arrival time distribution.

Quantity	Original N bits	Proposed N bits
E_T	10	10
Time	5	5
η	8	8
ϕ	8	8
$N_{crystal}$	8	4
Spike flag	1	2
Spare bits	0	3
Total	40	40

Table 4.1: Cluster word definition for the Phase II ECAL upgrade. Proposed values from TDR [73].




Figure 4.22: Expected event rate in the CMS ECAL during Phase II (HL-LHC) as a function of E_T . One spike-killing algorithm is applied and 200 pileup is assumed. The event rate is shown for an integrated luminosity of 0 (undamaged), 300, 1000, 3000 and 4500 fb^{-1} . Top: LD spike killing algorithm is used. Bottom: Swiss-cross spike killing algorithm is used. From [73].

at 10 GeV the LD algorithm will provide a spike event rate of a few hundred Hz after an integrated luminosity of 4500 fb^{-1} , while the event rate without any spike killing would be in excess of 10^7 Hz .

4.3.1.3 Expected Performance

As previously mentioned, radiation damage to both the lead tungstate crystals and APDs will affect the performance of the ECAL barrel somewhat. Figure 4.23 shows the reconstruction efficiency for electrons and photons in Run II and for ageing conditions throughout Phase II. For electrons, performance is maintained at high p_T , while at lower p_T there is some efficiency loss which could be mitigated through careful tuning of cluster parameters [103]. Photon reconstruction efficiency and energy resolution is also reduced by ageing, although as with electrons, this could be improved by optimising clustering algorithms [81].

Improved timing will be important in Phase II, particularly because of the challenges presented by maintaining reconstruction performance with up to 200 pileup events per bunch crossing. The primary vertex efficiency of $H \rightarrow \gamma\gamma$ will decrease from 75% to 30% without any improvements in timing compared to the Phase I detector. However, the planned 30 ps timing will allow improved vertex localisation. As illustrated in Figure 4.24, precision timing from both the ECAL and minimum ionising particle (MIP) timing is expected to give similar results to those seen in Run 2. The expected sensitivity gain is approximately 10% for the $H \rightarrow \gamma\gamma$ resolution relative to the case with no precise timing with up to 140 pileup events per bunch crossing, with details on the proposed MIP timing detector given in [72].

4.3.2 Endcap Upgrade

As with the ECAL barrel, the components of the ECAL endcap must be considered for Phase II to determine whether or not they are capable of maintaining the required performance. As shown in Figure 4.7, the endcap has suffered significantly more than the barrel from radiation damage to date, with damage in regions of the endcap already considerably exceeding what will be seen in the barrel at the end of Phase II. Therefore, the plan for the endcap is to completely remove the current subdetector and replace it with a high granularity calorimeter (HGCAL). This new subdetector will replace both the ECAL and the HCAL endcaps.

The design requirements of the new HGCAL are as follows:

- Radiation tolerance
- Fine granularity for energy resolution and discrimination between detected objects
- Dense calorimeter for compact showers
- Precision timing




Figure 4.23: Top: reconstruction efficiency for electrons where the electrons were generated with a uniform transverse momentum distribution. Bottom: reconstruction efficiency for a prompt photon from $H \rightarrow \gamma\gamma$ events. Both are shown for several barrel ageing conditions from Run II conditions up to 4500 /fb. From [73].




Figure 4.24: $H \rightarrow \gamma\gamma$ mass resolution with different assumptions on vertex efficiency. Green represents no precise timing, blue represents precise timing in the calorimeter only, and red represents precise timing in the calorimeter and MIP timing. Conditions in Run 2 are shown in black. From [72].

- Contribute to L1 trigger

Extensive research by groups within the CMS collaboration demonstrated that silicon sensors have the radiation tolerance required to withstand the radiation environment provided by the HL-LHC. After exposure to fluences over 50% higher than that expected in Phase II, the silicon sensors tested demonstrated sufficient charge collection for use in the endcap [104]. Therefore, silicon sensors were chosen as the active material in the front section of the HGCal. The rear section of the detector, which will be subject to lower radiation damage than the front section, will utilise plastic scintillator tiles with the light read out by silicon photomultipliers. Simulations have been carried out to investigate the radiation tolerance of this design, with Figure 4.25 showing the accumulated dose after 3000 fb⁻¹.

The granularity and density of the HGCal will provide a good energy resolution and a higher efficiency for reconstructing narrow jets compared to the current endcaps. The segmentation of the detector will give a $H \rightarrow \gamma\gamma$ mass resolution that is at worst comparable to the resolution of the Phase I ECAL endcap when the Higgs boson was discovered. As in the barrel, precise timing capabilities are essential for pileup mitigation. As shown in Figure 4.26, the addition of a timing requirement significantly reduces the number of hits projected to the front face of the




Figure 4.25: The simulated absorbed dose in the HGCal after an integrated luminosity of 3000 /fb. From [104].




Figure 4.26: Simulated hits in the HGCal with a charge greater than the threshold on the time of arrival measurement (12 fC). Left: with no timing requirement. Right: with the requirement $|\Delta t| < 90$ ps. From [104].

HGCal.

Chapter 5

Test Beam Analysis for ECAL Phase II Upgrade

The aim of the work described in this chapter is to use test beam data collected with prototype ECAL Phase II readout electronics to improve understanding of amplitude reconstruction and energy clustering for the Phase II upgrade. In particular, I explore appropriate data quality cuts and the optimal crystal matrix size for the clustering algorithm to obtain the best energy resolution for electrons and photons in Phase II.

To date, there have been three test beam runs using the H4 beam line of the CERN SPS which provides a high-purity electron beam. The analysis presented below uses data from a June 2018 test beam run, which was chosen due to the availability and completeness of the data and reconstruction at the time this analysis was completed.

Section 5.1 details the setup of the 2018 test beam. Then, Section 5.2 explains the steps taken throughout the analysis. This is divided into Subsection 5.2.1 on how the signal amplitudes were extracted from the test beam data, then Subsections 5.2.2 and 5.2.3 which explain the data quality cuts used in the analysis. Subsection 5.2.4 details the particular problems that were encountered in the 50 GeV run and how they were mitigated. Next, Subsection 5.2.5 explains the energy corrections needed to equalise the crystal response, before Subsection 5.2.6 details the energy corrections for electrons impacting on the side of a crystal. Finally, Section 5.3 shows the resolution performance for three different sizes of crystal matrix found in this analysis and based on this infers the optimal crystal matrix size for clustering.

All work detailed in this chapter is my own, with the results presented based on an internal CMS analysis note I wrote at the time the work was completed [105].

5.1 June 2018 Test Beam Setup

The June 2018 tests used an electron beam with the beam energy varied from 25 to 250 GeV and a small ($\sim 1\%$) pion contamination. The electron beam was provided by the H4 beam line

of the CERN SPS, and the beam passed through two hodoscope planes before impacting on a 5x5 crystal matrix. The energy spread of the beam is assumed to be small enough to be neglected in the following analysis. Hodoscopes are segmented detectors used to determine the trajectories of charged particles. These segments are usually made of a scintillating material, with the scintillation light converted into an electrical signal by a photomultiplier tube. A diagram of the experimental setup can be seen in Figure 5.1, with photos of the setup itself shown in Figures 5.2 and 5.3.

This test beam run investigated all of the available first prototypes of the ECAL Phase II application-specific integrated circuits (ASICs). The first of these is the calorimeter trans-impedance amplifier (CATIA), one of the two Very Front End (VFE) components, which amplifies and shapes the signals from the APDs. The second half of the VFE, the Lisbon-Torino ECAL data transmission unit (LiTE-DTU), was not tested as there was no prototype available at the time. The LiTE-DTU will perform analog to digital conversion as well as selecting, compressing and transmitting data. Instead, a commercial ASIC with comparable performance for the limited number of channels used during this test beam was used. A detailed explanation of the Phase II ASICs can be found in Section 4.3.1. This analysis uses data collected from an energy scan from 25 to 250 GeV that was sampled at 160 MHz, with narrow beam collimators that resulted in a small beam width of 10 mm. For each energy, approximately 50,000 events were recorded.

The data collection is triggered by a coincidence in the two hodoscope planes. An example pulse shape can be seen in Figure 5.4. When data sampling occurs, the first 20 samples are used to calculate the pedestal value for that particular pulse. This is simply done by calculating the mean value in this window, though variables such as the sum of the values in this window are also calculated and saved for debugging purposes.

There are two microchannel plate detectors (MCPs) in the test beam setup. MCPs are made from plates of a resistive material (such as glass) that contain a regular array of microchannels, which run parallel to each other. They are used to detect single particles and, in the particular case of this test beam, are used to detect electrons and provide timing information.

To help mirror the conditions found in the CMS detector, the test beam hardware was housed in an insulated box which used a water cooling circuit to maintain a constant 18°C. This is the current operating temperature of the CMS ECAL. A dedicated test was also performed at 9°C, the expected operating temperature for the first run in Phase II, though there is only a limited amount of data available at this temperature so it is not used in this analysis.

The crystal matrix scheme for the June 2018 test beam can be seen in Figure 5.5. It should be noted that there was no magnetic field in this test beam setup, while the ECAL in CMS is subject to a 3.8 T magnetic field. For an electron impacting on the centre of a crystal, about 80% of the electron's energy is deposited in that crystal, with nearly all the rest of the energy deposited in the surrounding 3x3 crystal matrix. Roughly 4% of the electron's energy is




Figure 5.1: June 2018 setup for the ECAL test beam at CERN.



Figure 5.2: Inside the box used for the 2018 test beam run. The crystal matrix can be seen inside the copper coloured box, while there are bare APDs (i.e. not mounted to a lead tungstate crystal) above this. Though the crystal matrix itself is 5x6, only five of the six crystals per column could be read out as the VFE is limited to five channels. Photo provided by Simone Pigazzini.




Figure 5.3: Box with crystal matrix inside in situ at the H4 experimental area. This box is placed on a moving table that allows the box to be moved up, down, left and right remotely to allow the beam to be centred on different crystals. Photo provided by Simone Pigazzini.

deposited in crystals that share a side with the central crystal, while around 1% of the energy is deposited in crystals that share a corner with the central crystal. A small proportion of the electron's energy is also lost in cracks between the crystals.

In addition to this test beam run, there was a subsequent test beam period in 2021 which used a full 5×5 crystal matrix and updated Phase II prototypes. The main added feature of the 2021 test beam run compared to 2018 is the addition of gain switching, which comes into play at around 150 GeV. There are two gain levels (1 and 10) in the CATIA, which allows a more precise measurement of the ADC counts at lower energy by allocating them more bits while limiting the probability of saturating ADC counts at higher energies.

As gain switching has the possibility to change the pulse shape, there was the potential for the linear discriminant spike killing algorithm to require some tuning. Neither the swiss cross nor clustering algorithms consider pulse shape, so they are not expected to change compared to the 2018 results.

In 2022, there was a larger-scale test of the latest prototype electronics using a spare ECAL supermodule. This test beam time was largely devoted to set-up and testing of the supermodule




Figure 5.4: Example pulse shape for a 200 GeV electron beam centred on C3, the central crystal in the crystal matrix.



Figure 5.5: Matrix scheme of the lead tungstate crystals for the June 2018 test beam. The four dark grey crystals - A4, A5, E4 and E5 - did not collect data during this test beam run due to broken VFE adaptors. Each column refers to one VFE.

with the full Phase II ECAL readout electronics, with limited time available to record data for analysis purposes. However, the next scheduled test beam in 2023 will be focused on collecting data with electrons and pions for analysis.

In addition to the 2018 electron test beam campaign, there was a test utilising the same experimental setup at the H2 pion beam line in September 2018. The beam had an energy of 120 GeV, with lead absorbers used to filter out electrons and obtain a pure pion beam. Over 1.3 million events were recorded in the crystal matrix, with an additional 3.8 million events recorded that impacted on the bare APDs seen at the top of Figure 5.2 [106]. Due to issues with the VFEs, only 16 channels were functioning for this test.

5.2 June 2018 Test Beam Analysis

5.2.1 Methods to Extract the Amplitude

The raw test beam data was converted into a test beam specific data format for analysis, which used the samples in each recorded pulse shape to calculate information such as the time and amplitude of the recorded pulse. This data format contains three different measurements of the pulse amplitude. These are called maximum, `amp_max` and `fit_ampl`. Maximum simply returns the pedestal subtracted maximum sample value in the pulse window. `Amp_max` interpolates between different samples in the pulse using a second-order polynomial to determine the amplitude of the pulse. This method of amplitude reconstruction struggles in situations where one sample has a positive value and the next sample goes below the pedestal value. Finally, `fit_ampl` extracts a value for the amplitude using a template fit to a reference pulse shape. As with the template fit used in the CMS detector, the free parameters of this fit are the amplitude and time. The results of these three methods are compared in Figure 5.6. This shows that maximum gives a poor resolution, while `amp_max` creates a second artificial peak at around 10000 ADC counts due to occasional interpolation issues when one or more samples drop below the pedestal. Therefore, `fit_ampl` was chosen to reconstruct the amplitude for each event in the analysis presented below. As mentioned in Section 4.2.3.1, a template fitting technique is also used for amplitude reconstruction of signals in the CMS ECAL [88].

Though `fit_ampl` is the best available option for amplitude reconstruction, some pulses are not well described by the template used by `fit_ampl`. This is usually due to multiple pulses within the sampling window. These events were removed by placing a χ^2 cut on the data. As there was a correlation between χ^2 and the electron beam energy, a cut of $\chi^2/E < 5$ (where E is measured in ADC counts) was set, where χ^2 is

$$(5.1) \quad \chi^2 = \sum \frac{(O_i - E_i)^2}{E_i}$$

where O_i is the observed value and E_i is the expected value. This correlation is thought to be due to an imperfect pulse shape and the exact value of the cut was chosen to balance between




Figure 5.6: Histogram showing the amplitude given by each of the three amplitude reconstruction methods in the test beam specific data format for a 3×3 crystal matrix at 200 GeV.

removing poorly reconstructed events while retaining as many events as possible.

5.2.2 Hodoscope Cuts

Hodoscope cuts can be placed on this test beam data in a range between 2 and 10 mm, with analysis of energy resolution usually selecting events with electrons hitting the crystal in a position inside a square of $6 \times 6 \text{ mm}^2$ around the crystal centre. As the location of the centre varies from crystal to crystal, hodoscope cuts need to be correctly centred to ensure that only data from electrons hitting the desired region are kept.

Figure 5.1 shows that there are two hodoscope planes for the test beam data providing independent x-y position measurements. While they have approximately the same position in y, there is an offset of $2 \pm 0.5 \text{ mm}$ between their positions in x. As shown in Figure 5.7, almost all the electrons that pass within $3 \pm 0.5 \text{ mm}$ of the crystal centre of C3 in the first hodoscope plane also pass within 3 mm in the second hodoscope plane once the offset between the two planes has been accounted for. Because of this, almost all of the hits on the second hodoscope plane are within 3 mm of (0,0). As such, the two planes are in good agreement and either could be used to select events near the crystal centre. For this analysis, just the first hodoscope plane is used.

As the hodoscope cuts are tightened, fewer events pass the selection criteria. This means that the peaks reduce in height and width, with the mean of the Gaussian also shifting towards higher energies due to better containment. More stringent hodoscope cuts leave a smaller fraction of the original data due to the width of the beam profile. Without any cuts, the single crystal




Figure 5.7: Histogram showing the second hodoscope plane with cuts of 3 mm in both coordinates on the first plane, accounting for the 2 mm offset between their x positions. The beam energy used was 200 GeV and the beam was aimed at the centre of C3, though the axes have been adjusted in this instance to place the centre of C3 at (0,0).

amplitude has a noticeable exponential tail. This can be seen in Figure 5.8. The exponential tail comes from electrons impacting near the edge of the crystal, therefore depositing more energy in the surrounding crystals. The tail is less pronounced for a 3×3 crystal matrix due to better shower containment.

5.2.3 Incorrect Pedestal values

As shown in Figure 5.9, there are some recorded instances of pulses with an incorrect pedestal value. This occurs when there are two pulses in the sampling window, with the first occurring at the start of the window which is used to calculate the pedestal. Such events are excluded by placing a cut on the sum of samples in the baseline sample window. This cut was tested on a single run and found to remove all the incorrect pedestal events.

5.2.4 Problems with the 50 GeV Run

The 50 GeV run has a considerable peak in the distribution of the energy deposited in the central crystal, C3, around 0 in addition to the expected peak of around 1500 ADC counts. This can be seen in Figure 5.10 which shows the energy deposited in the C3. If ± 3 mm hodoscope




Figure 5.8: Histogram showing the effect of hodoscope cuts of increasing tightness on the single crystal amplitude for a 200 GeV electron beam centred on crystal C3.




Figure 5.9: Pulse shape for an event during a 150 GeV run which has two pulses, one at around 0 ns and the second at around 300 ns.




Figure 5.10: Amplitude of energy deposited in C3 during the 50 GeV run, measured in ADC counts.

cuts are used then this peak at zero is removed, though this also causes a factor of two reduction in events compared to other energies with the same hodoscope cuts applied. I investigated a variety of possible causes for this second peak.

Firstly, I considered whether an issue had occurred partway through the run. If this had occurred then I would expect the amplitude to have a dependency on spill number, where a spill is the removal of a controllable number of protons from the H4 beamline, which are subsequently collided onto a target and steered down a short beamline to provide the electron beam used in this test beam. The spill number is a counter and is used to see which spill corresponds to the data collected. If there was a time-dependent issue with this run, I would be able to place a cut so only the unaffected spills were used. However, the amplitude in C3, the central crystal, does not vary considerably with spill number.

Next, a number of events were studied to see if there were anomalous pulse shapes. As expected, the pedestal was around 0 before and after the event, with the peak at around 1500 ADC counts as expected. As such, a non-zero pedestal or atypical pulse shape is not causing the peak near 0.

I also investigated the amplitude of C3 with regards to the hodoscope position in both x and y. As expected from this run, the width of the beam is approximately 10 mm, with almost all events with around 1500 ADC counts happening here and a sharp reduction in events outside




Figure 5.11: 2D Histogram showing the hodoscope occupancy for the 50 GeV run.

of this region.

Then I looked at the energy deposited in all 21 crystals to see if there was any deviation from the expected energy distribution. A larger than expected amount of energy was deposited in the crystal C5, showing an energy similar to that found in C2 and C4 which are expected to contain about 4% of the electron's energy. This increased spread of crystals with high counts suggests a beam misalignment. Logbook entries also state that there was a slow collection rate during this run and that afterwards the magnets were tuned which led to a factor of four increase in the collection rate. This slower collection rate would be expected from a misaligned beam and is consistent with hodoscope cuts cleaning up the low energy contributions.

Finally, the position of the beam at the hodoscope was considered as this indicates where the electrons impact on the crystal matrix. As shown in Figure 5.11, there is a clear preference for the upper region of the hodoscope and some preference specifically towards the top right corner. For a correctly aligned beam, this should show a uniform distribution. This confirms that the 50 GeV run had a misaligned beam.

Fits can still be made to the peak to find the standard deviation, σ , for the single crystal, 3×3 and 3×5 matrix resolution. Fits were completed using a Crystal Ball function that consists of a Gaussian core portion and an exponential low-end tail [107]. The completion of these fits allowed the extraction of the mean and standard deviation to determine the resolution. Due to




Figure 5.12: Histogram showing the effect of applying the correct calibration constant to the crystal C2 for a 200 GeV run.

the channels that could not collect data (see Figure 5.1), the 3×5 crystal matrix is 3 crystals wide in x and 5 crystals wide in y , keeping C3 as the centre. To find σ in the cases without hodoscope cuts, the Crystal Ball fit was also restricted such that it excluded the lowest energy part of the exponential below 1000 ADC counts in the 1×1 and below 1750 ADC counts in the 3×3 and 3×5 as this was caused by the beam misalignment and is not present at any other energy. However, due to the beam misalignment, there are fewer events, and hence larger errors, for all 50 GeV σ values. Similar numbers of events were recorded for all other energies.

5.2.5 Calibration Constants

Each crystal has a slightly different light yield from electrons or photons. Therefore, each crystal has a calibration constant, which must be considered, particularly when summing the energy from multiple crystals. This is similar to what is described in Section 4.2.3.3. The crystals for this test beam were calibrated in a previous analysis [108], and the same crystal calibration constants are used in this analysis. The calibration constant in the previous analysis set the calibration constant to 1 for C3, with the calibration constants for all other crystals set relative to this using test beam data centred on each crystal. Unlike the CMS barrel, these calibration constants are not set such that the mean calibration constant is 1. As shown in Figure 5.12, the peak in ADC counts for C2 is higher than that for C3 without the calibration constant, but when calibrated the two are in good agreement. This energy agreement after calibration gives confidence that the corrections are working as intended.

5.2.6 Energy Containment Correction

While many runs in the test beam have electrons impacting on the centre of crystals, in CMS electrons may impact on any part of the crystal. If the impact occurs near an edge or corner of a crystal, then some energy will be lost due to cracks between the crystals. This energy loss amounts to $\sim 2\%$ on average, which can be recovered via the use of energy corrections. As mentioned in Section 5.2.6, there are multiple ways to correct the measured energy. As this method is intended to be applied online in the L1 clustering algorithms, I am only looking at containment corrections as this is simplest to implement online. Other correction options include empirically derived corrections which use variables such as particle energy, η , shower shape and particle type, or multivariate analysis-based corrections, which use multiple detector-based quantities and are trained on MC truth. The width of the Gaussian, σ , increases by $\sim 20\%$ for electrons impacting on the sides of the crystal, this width can also be reduced to around the value for electrons impacting on the centre of the crystal by using corrections. The correction used here is based on the ratio of the energy deposited in certain rows (or columns) of crystals compared to the energy deposited in the remaining crystals of the matrix [109].

Events were selected where C3, the central crystal, always had the highest energy deposit in the crystal matrix. Three runs were combined to effectively form a scan from the top to the bottom of the crystal. In the first run, the beam was directed at the centre of the crystal, while in the second run, the beam was directed at the top of the crystal and in the last run the beam was directed at the bottom of the crystal. As there are no runs impacting on the left or right of the crystal in this test beam data, it was only possible to find the energy correction function in the y-direction. However, the method described below has been used previously in CMS to correct in 2D with good results [110].

Firstly, the energy deposited in C2 and C4 was compared for each event. If a higher energy was measured in C4, then two energy sums, E1 and E2, are defined such that E1 is the sum of the energies in crystals B4, C4 and D4, while E2 is the sum of the energy in the remaining six crystals. Conversely, if C2 has the greater energy deposit, then E2 is defined as the sum of the energies in crystals B2, C2 and D2, while E1 is the sum of the energies in the remaining six crystals as shown in Figure 5.14. Due to the limited number of crystals in the test beam data, Cartesian coordinates are used to describe relative position rather than η and ϕ .

The energy measured in the 3×3 crystal matrix was then evaluated as a function of $\ln(E2/E1)$. The data was divided into 100 bins of $\ln(E2/E1)$, and Crystal Ball functions were fitted to the energy distribution in each bin, allowing for a measurement of the mean. The results were then normalised such that the mean measured energy had a value of 1 for $\ln(E2/E1) = -3$. This corresponds to a situation where very little energy is deposited in E2 and the majority is deposited in E1, as expected for an electron impacting on the centre of the crystal. Next, second-order polynomials were independently fitted to the positive and negative regions of $\ln(E2/E1)$ as shown in Figure 5.13 to obtain correction functions. The correction functions




Figure 5.13: Second order polynomial fits to the positive (left) and negative (right) regions of $\ln(E2/E1)$.




Figure 5.14: Definitions of energy sums E1 and E2 in the case where $C2 < C4$ (left) and $C2 > C4$ (right). The black circle represents the electron impact position.

were chosen such that the correction function multiplied by the second-order polynomial fit would give a straight line at $\ln(E2/E1)=1$. Finally, these correction functions were applied on an event-by-event basis.

As shown in Figure 5.15, the amplitude measured in the 3×3 crystal matrix before containment corrections has a dependence on the parameter $\ln(E2/E1)$. This dependence can be removed by applying the energy correction as shown in Figure 5.16. The energy correction also improves the mean by 1% (i.e. moves the mean closer to the value measured for runs centred on the centre of the crystal) and improves σ by 20% when the electrons are impacting on the bottom half of the crystal, as shown in Figure 5.17. Similarly, for electrons impacting on the




Figure 5.15: Uncorrected amplitude as function of $\ln(E2/E1)$ at 50 GeV. The bright spot around $\ln(E2/E1)=2$ is due to an increased proportion of events aimed at the top of C3.

top half of the crystal, applying the energy correction improves the mean by 1% and σ by 24%. For electrons impacting near the centre of the crystal, both the mean and σ remain close to their original values.

5.3 Resolution Performance

The energy resolution of three different crystal matrix sizes was investigated using the energy spectra of that crystal matrix size for a range of energies which are known to two decimal places. For the 3×3 crystal matrix, the effect of the $\ln(E2/E1)$ energy correction on the energy resolution is also studied. All energy spectra were fitted with a Crystal Ball function to account for the exponential tails in the Gaussian distributions. Both σ_{CB} and the mean energy, $\langle E \rangle$, were extracted from these fits and used to plot $\sigma/\langle E \rangle$ vs beam energy for each size of crystal matrix with and without hodoscope cuts.

Each resolution plot was subsequently fitted with

$$(5.2) \quad \frac{\sigma}{\langle E \rangle} = \sqrt{\left(\frac{N}{E}\right)^2 + \left(\frac{S}{\sqrt{E}}\right)^2 + C^2}$$




Figure 5.16: Corrected amplitude as a function of $\ln(E2/E1)$ at 50 GeV for the same data as Figure 5.15. The bright spot around $\ln(E2/E1)=2$ is due to an increased proportion of events aimed at the top of C3.

Crystal Matrix Size	Noise term	Stochastic term	Constant term
1×1 hodoscope cuts	0.124 ± 0.071	0.0852 ± 0.0042	0.00339 ± 0.00060
3×3 no hodoscope cuts	0.43 ± 0.022	0.0478 ± 0.0075	0.00288 ± 0.00068
3×3 hodoscope cuts	0.414 ± 0.0012	0.0314 ± 0.0062	0.00341 ± 0.00031
3×3 containment correction	0.404 ± 0.017	0.0496 ± 0.0049	0.00301 ± 0.00043
3×5 no hodoscope cuts	0.5036 ± 0.0067	0.0478 ± 0.0030	0.00343 ± 0.00021
3×5 hodoscope cuts	0.4736 ± 0.015	0.0399 ± 0.0069	0.00329 ± 0.00045

Table 5.1: Noise, constant and stochastic terms for each crystal matrix size considered in this analysis.

. The fit parameters are N, S and T, where N stands for the noise term, S stands for the stochastic term and C stands for the constant term. The noise term is important at low energies and includes contributions such as the ADC pedestal widths. The stochastic term includes all processes that are governed by Poisson statistics and therefore includes contributions due to fluctuations in the number of photoelectrons released at the face of the APDs as well as fluctuations in sampling (energy loss). Finally, the constant term affects the resolution at all energies and includes sources such as errors in the crystal intercalibration [111].




Figure 5.17: Histograms showing Crystal Ball fits for 100 GeV electrons impacting below the crystal centre with (top) and without (bottom) energy corrections.




Figure 5.18: Resolution for a single crystal with hodoscope cuts (top left), a 3×3 matrix with hodoscope cuts (top right), a 3×3 matrix without hodoscope cuts (middle left), a 3×3 matrix with $\ln(E2/E1)$ energy corrections applied (middle right), a 3×5 matrix without hodoscope cuts (bottom left) and a 3×5 matrix with hodoscope cuts (bottom right).

The fits for each resolution plot can be seen in Figure 5.18 and the fit values are given in Table 5.1. There is no plot for a single crystal resolution without hodoscope cuts as this fit did not converge. As expected, the single crystal provides worse resolution than larger crystal matrices due to the poor shower containment, with $\sim 20\%$ of the electrons' energy lost to the surrounding crystals. The exponential tail in the Crystal Ball fits is most pronounced for a single crystal and reduced in the 3×3 and 3×5 matrices due to the addition of the energy deposited in the surrounding crystals.

The noise term for the 3×3 matrix with hodoscope cuts is approximately 3 times that for the single crystal, and that for the 3×5 matrix is around 4 times as large as the single crystal term. With the fit not converging for the 1×1 crystal matrix with no hodoscope cuts, it is not possible to confirm whether the noise term follows this same pattern when hodoscope cuts are not used, but both the 3×3 and 3×5 have slightly larger N and S parameters than the same matrix size with hodoscope cuts. This will be due to the addition of events that are less well-contained and impact closer to the edge of the crystal. This roughly follows the expected $\sqrt{n}$ scaling for the noise term, where n is the number of crystals in the crystal matrix, however as the crystals have different noise values it is not exact.

During a previous test beam with Phase I electronics the mean value of the stochastic term was 2.8% while the constant term was 0.30% [110]. Errors are not quoted. These measurements were taken by considering a 3×3 crystal matrix with 2 mm hodoscope cuts, so the closest comparison using Phase II electronics is the 3×3 with 3 mm hodoscope cuts. The stochastic terms of both Phase I and Phase II are comparable, with respective values of 2.8% and $3.4\pm 0.6\%$. The two constant terms are also in agreement at 0.30%.

The performance of a 3×5 crystal matrix was also explored. At low energies, it provides a resolution comparable to a single crystal due to increased noise sensitivity, while at higher energies resolution is close to that for a 3×3 . With a 3×5 matrix having 6 additional crystals compared to a 3×3 , if an electron impacts at the centre of the matrix, then only a tiny fraction of the deposited energy will be contained in the additional crystals, with the majority being deposited in the closest 9 crystals. This means that the additional crystals are contributing to the noise in the crystal matrix while only contributing a small fraction of the energy deposited by the electrons. However, the relative performance of the 3×5 cluster may improve in the CMS magnetic field as electron clusters are extended in ϕ due to the bending in the field, and a 3×5 cluster may have less shower leakage than a 3×3 cluster.

The S and C values for the $\ln(E2/E1)$ corrected resolution are within the error bounds for the uncorrected values, while the corrected noise term is just outside the error bounds of the uncorrected fit without hodoscope cuts. This shows that the $\ln(E2/E1)$ correction maintains the resolution for electrons impacting on or close to the centre of the crystal, where we expect the correction to only make small adjustments to the energy of each event.

Also, it should be noted that all resolution plots are from electron beams aimed at the centre

of the crystal, while in CMS electrons can impact on any point of the crystal. For electrons impacting on the side of a crystal, the resolution will decrease considerably without energy corrections. As shown in Figure 5.17, the $\ln(E2/E1)$ correction provides a 20% improvement in σ and a 1% improvement in E when the beam is centred on the side of a crystal. Accordingly, the resolution for impacting on the side of a crystal is $\sim 20\%$ worse without energy corrections.

5.4 Conclusions

The energy resolution for a single crystal, a 3×3 crystal matrix and a 3×5 crystal matrix has been measured for an energy range of 25 to 250 GeV, both with and without hodoscope cuts. In test beam conditions, a 3×3 matrix with hodoscope cuts provided the best energy resolution performance. The obtained constant term of around 0.3% for the 3×3 matrix with hodoscope cuts is consistent with expectations for Phase II. For all sizes of crystal matrix considered, the resolution was improved with the addition of hodoscope cuts. An energy correction based on $\ln(E2/E1)$ was also tested and found to retain the energy resolution for electrons impacting on the centre of a crystal. The resolution was improved with the $\ln(E2/E1)$ correction for electrons impacting on the side of a crystal by both increasing the mean energy and reducing σ compared to the uncorrected resolution for electrons impacting on the side of a crystal, with the mean and σ close to that for electrons impacting on the centre of the crystal.

Chapter 6

Testing Phase II Clustering and Spike Killing Algorithms

This chapter details the work that has been undertaken to produce and test some of the algorithms that are required to produce trigger primitives for the Phase II ECAL upgrade for HL-LHC. All of the trigger primitive algorithms must fit on the FPGA of the barrel calorimeter processor (BCP) board. As previously mentioned in Subsection 4.3.1.2, two of the planned features of the BCP firmware are anomalous APD signal (spike) rejection and basic clustering of energy deposited in the ECAL barrel, and these are the algorithms covered in this chapter.

There is a planned redundancy with regard to spike killing by the use of two separate algorithms. The Swiss-cross algorithm struggles to distinguish between EM showers and spikes for low-energy deposits due to noise in the surrounding channel. However, it does not require the increased 160 MHz sampling rate of the Phase II upgrade to provide spike rejection. The linear discriminant (LD) algorithm works well with lower-energy particles as it does not use information from surrounding channels to determine whether an energy deposit is a spike or an EM shower, but it does require 160 MHz sampling. The combination of these two algorithms allows for spike killing at low energy and provides a backup for channels that must be digitised at a lower sampling rate due to noise. Full details of the Swiss-cross algorithm used in the current ECAL can be found in Subsection 4.2.2, while the LD method is described in Subsection 4.3.1.2.

The timing and amplitude reconstruction will be done by a separate multifit algorithm (also described in Subsection 4.3.1.2) and it is expected that the multifit algorithm will provide amplitude and timing information to the Swiss-cross and clustering algorithms, while the LD algorithm only needs the linearised samples. As all of these algorithms will have to fit on a single field-programmable gate array (FPGA), they will have to share resources. When the algorithms are more mature, optimisation of their resource usage to ensure that they all fit on the FPGA will be important. However, that work is beyond the scope of this thesis.

The work outlined in this chapter has been completed over the course of several years, so the algorithms have been updated between the sections presented. During the software

comparison in Section 6.3, all three algorithms (LD, Swiss-cross and clustering) were separately coded firmware, while during the hardware test in Section 6.4, the Swiss-cross and clustering algorithms were combined into one. Both the Swiss-cross and clustering algorithms scan across every crystal in the matrix, which will contain 600 crystals in the final version, and compare the energy in the central crystal with that of its neighbours, giving them a common starting point which means that combining the two into a single piece of firmware will make more efficient use of available FPGA resources.

Section 6.1 gives an overview of high-level synthesis (HLS) and motivates why it has been used to write the firmware for these algorithms. Then, Section 6.2 covers previous work on the LD algorithm including a hardware test that I was not personally involved in but is included for completeness. In Section 6.3, the HLS versions of the three algorithms are compared to independently written C++ counterparts. Finally, Section 6.4 details the hardware test of the combined Swiss-cross and clustering algorithm. Much of this work has been done in close collaboration with Thomas Reis at Rutherford Appleton Laboratory. The workload has been split between us with Thomas mostly writing the algorithms, while I focused on testing the algorithms. The exception to this is the clustering HLS algorithm tested in Section 6.3, which I wrote using the Swiss cross HLS code as a starting point.

6.1 High Level Synthesis

An FPGA is a programmable, flexible device that can be programmed via firmware to complete a variety of tasks including providing timing signals and aiding in triggering of particle physics experiments. Each FPGA has some common components:

- Flip-flops (FF): a small memory used to store a bit.
- Look-up table (LUT): a small memory that holds a custom truth table. This can perform operations such as “AND”, “OR”, addition and subtraction.
- Digital signal processor (DSP): a small processor to quickly perform operations such as addition and multiplication.
- Block RAM (BRAM): a larger memory used to store data, though slower than a FF.

Traditionally, firmware for trigger applications was written using a hardware description language (HDL) such as VHDL or Verilog [112]. Writing firmware in an HDL requires electronics knowledge to correctly code, while HLS only requires knowledge of the hardware limitations to code correctly and efficiently. Additionally, it can be difficult to verify firmware written with an HDL in C++ as an emulator is needed. With HLS, it is possible to write a native C++ version of an algorithm at the same time as writing a HLS C++ version which will be used to generate

the firmware. This parallel development also improves the ease of integration with the CMS software framework used for simulation and reconstruction [113, 114].

For HLS to be used instead of an HDL, it must result in similar resource usage on the FPGA. For trigger applications in particular, it is essential that the firmware meets the strict latency requirements of the trigger system ($12.5\ \mu\text{s}$ in the case of CMS Phase II). When these criteria are met, then HLS generally provides a shorter development time and code that is more easily adaptable. The similarity of HLS to C++ also allows for a larger cohort of physicists to contribute towards firmware development than if then firmware was developed in an HDL. This increased accessibility can also allow algorithms to be changed quickly and easily if required during data taking, e.g. to account for new pileup conditions or increased radiation damage. It is this combination of features that has led to HLS being chosen over an HDL for developing the trigger primitive algorithms.

The algorithms detailed in this Chapter were developed using Vivado HLS (particularly v2019.1 and v2019.2). The normal Vivado HLS workflow is shown in Figure 6.1. A C++ design is converted into an HDL implementation using Vivado HLS, with the developer also providing constraints such as the clock frequency, clock uncertainty and the specific FPGA model to be used. Alongside this C++ design, the developer also provides a C++ test bench that is used to confirm in a simulation that the algorithm works as intended [115].

In the test bench, the HLS function is executed with a chosen set of inputs and the result is then compared to a predetermined output. For example, for the clustering algorithm considered in this thesis, individual crystal energies and spike flags are given as inputs and the output clusters are compared with the reference values for E_T , position and number of crystals. This test bench step is called C simulation. Once the design is working in the test bench, the next step is synthesis, which synthesises an HDL code from the given C++ code according to the given timing and FPGA constraints. This synthesis process produces a report on the performance of the design specifically: a clock estimate that shows whether the design satisfies the required clock frequency; a latency estimate which provides the number of clock cycles required to finish processing and the interval at which the algorithm accepts data; and a utilisation estimate of the FPGA resources. The clock frequency determines the frequency at which the FPGA components are activated, and therefore how fast any algorithms run. The interval at which the algorithm accepts data can also be adjusted in HLS to ensure that the data is processed within the required timeframe, but if this requires the algorithm to accept more data before the previously accepted data has been analysed, then there will be an associated FPGA resource cost.

At this stage, the code can be optimised. This can be done via any one, or a combination of, three methods:

- Pipelining: increase throughput by allowing the next iteration of a function or a loop to begin before the previous one has finished.




Figure 6.1: Workflow for Vivado HLS, which is represented by the grey box. The C/C++/SystemC code of the algorithm itself, along with a test bench file and constraints/directives are given as inputs. The main steps in the workflow are C simulation, C synthesis and register transfer level (RTL) simulation (or cosimulation), with the firmware exported at the end of these to a programme such as Vivado Design Suite which allows the firmware to be loaded onto an FPGA. VHDL or verilog are generated from the C synthesis, while the RTL adapter is needed to use the test bench in the cosimulation. steps. RTL is used in HDLs to create a high-level representation of a circuit. From [115].

- Unrolling: reduce latency and increase throughput by creating multiple instances of a for loop that run in parallel. This only works if the loop iterations are independent and it will impact resource usage because multiple copies of the loop are generated in firmware.
- Memory partitioning: allows faster memory access for arrays by switching arrays from being saved in the BRAM by default into single registers or ports.

This synthesised design can then be tested using the HDL simulator in a stage known as co-simulation, again compared to the C test bench. Situations where the C simulation is passed but the co-simulation fails are due to hardware limitations, such as the target FPGA having insufficient memory to run the algorithm. Once all the above steps give good results, then the design can be exported into a firmware format that can then be loaded onto the FPGA for a hardware test [116].

6.2 Previous Work on the Linear Discriminant Algorithm

Details of how the LD algorithm works can be found in Subsection 4.3.1.2. In previous work, both a second and third-order LD have been tested using HLS and their performances evaluated. Full details of this work are given in an internal CMS analysis note [102], with the important findings summarised below. As division is resource intensive in an FPGA, the LD value given in Equation 4.3 was rewritten to avoid denominators. For the third-order LD, this gives the expression

$$(6.1) \quad 100 * a_{n+1}a_n^2 - LD_{thr}a_n^3 - p_0a_n^3 < p_1a_{n-1}a_n^2 + p_2a_{n-1}^2a_n + p_3a_{n-1}^3$$

As a reminder, the variables a_{n-1} , a_n and a_{n+1} are defined as the three consecutive samples around the peak (a_n is the peak). The weights p_i are chosen to give an LD value centred at zero for an EM signal and values below zero for spikes. For this test, LD_{thr} was set at -10 which corresponds to a cut of -0.1 in Equation 4.3. The factor of 100 is added to allow the use of integer weights in HLS while retaining a direct conversion for the cut at -0.1 without rounding as would be the case when using a factor of 128. All weights were calculated using electron beam data from the 2018 ECAL test beam.

The second-order expression is

$$(6.2) \quad 100 * a_{n+1}a_n - LD_{thr}a_n^2 < p_1a_{n-1}a_n + p_2a_{n-1}^2$$

To minimise resource usage on the FPGA, the number of bits for each variable was optimised to fit the maximum expected values. Similarly, in an FPGA, multiplication uses digital signal processor (DSP) blocks, which are limited in number. To minimise DSP block usage, the multiplication by 100 was implemented by multiplying by 4, 32 and 64 and then summing. These multiplications by powers of two were performed using bit shifts instead of using DSP blocks.

Synthesis of 300 channels was performed to give resource usage estimates on a Xilinx KU115 FPGA for both the second- and third-order LD algorithms. The clock frequency used during synthesis was chosen to be 160 MHz to match the clock frequency for the Phase II ECAL. Table 6.1 gives the timing estimates, while Table 6.2 gives the latency estimates and finally Table 6.3 gives the resource estimates. The third-order algorithm used half of the available DSPs, that would severely limit the available space for the clustering and multifit algorithms which need to fit on the same chip. However, the second-order LD algorithm used just over 20% of the available DSPs on the FPGA, leaving a good fraction for the remaining algorithms. Both algorithms have comparable spike-killing performance and meet the timing requirements.

The detector note [102] was written in early 2020, and at this time the design of the BCP used two KU115 FPGAs, though this has since been changed to one VU13P FPGA. This FPGA upgrade will provide nearly three times the FFs, 30% more LUTs, 25% more BRAM and 11% more DSPs. The increased resource usage will allow the use of more sophisticated algorithms and for spike-killing redundancy with the inclusion of the LD and Swiss-cross algorithms.

As there were no BCP boards available at the time, the firmware was tested on an existing board that was designed for the L1 trigger, known as an MP7 [117] which can be seen in Figure 6.2. This existing board allowed a 72-channel hardware test. However, the test beam data used only 16 channels of data, so the remaining channels were filled with zero amplitude. The hardware test used 13,400 samples from the test beam and found perfect agreement between the hardware results and the expected output. The latency from the input buffer to the output buffer was found to be within specifications for both the second and third-order algorithms.

Based on the above results, the second-order LD algorithm was chosen over the third-order algorithm. The second-order algorithm has shown that it is capable of rejecting spikes at the required level of $>99.9\%$, while also freeing up more resources on the FPGA to be used by the other trigger primitive algorithms.

Target	Estimated	Uncertainty
6.25 ns	4.773 ns	0.78 ns

Table 6.1: Post implementation timing estimates for the second order LD algorithm for 300 channels on a KU115.

Latency		Interval	
Min	Max	Min	Max
3	3	1	1

Table 6.2: Post implementation latency estimates for the second order LD algorithm for 300 channels on a KU115.



Figure 6.2: The MP7 board used for the hardware test of the LD algorithm and the later test of the combined Swiss-cross and clustering algorithm. The attached cable was irrelevant to the test described in this thesis as data was loaded and read out via the crate’s backplane.

Resource	Used	Available	Utilisation %
LUT	56789	663360	8.56
FF	100399	13266720	7.57
BRAM	314	2160	14.54
DSP	1200	5520	21.74

Table 6.3: Post implementation resource usage for the second order LD algorithm for 300 channels on a KU115.

6.3 Comparison of HLS and C++ Algorithms

We have developed a software framework to test the three HLS algorithms in the CMS simulation and reconstruction framework against C++ counterparts. These C++ counterparts have been independently coded so are not written in exactly the same way as the HLS code, though the same outputs are expected for both. As neither the Swiss-cross nor the clustering algorithms had been tested before, they were the main focus of this test, though the LD algorithm was also tested in the same framework.

6.3.1 Clustering Algorithm Tests

The criteria for the initial stand-alone version of the 3x3 clustering algorithm are as follows:

- Seed has a greater E_T than any crystals it shares a side with (corners are not considered)
- The seed must have $E_T > 0.25$ GeV. This threshold can be easily tuned at a later date
- Seed is not an LD spike. Later this requirement will be changed such that the seed cannot be an LD or a Swiss-cross spike, but this version of the algorithm was not provided with any information from the Swiss-cross algorithm.

A subset of the data used in Chapter 5 was used to compare the HLS and C++ clustering algorithms, specifically a 200 GeV electron run centred on the central crystal C3 with over 50,000 events. As the data format specified in the ECAL Phase II technical design report [73] allows 12 bits for E_T while the test beam used 14 bits, the energy originally recorded is rounded to the nearest integer and then divided by 10 before it is given to either the C++ or the HLS algorithm.

To confirm the validity of this method, the results from the C++ algorithm were compared to a static cluster centred on C3 with the full test beam data precision. The static cluster is expected to give a slightly different output to the clustering algorithm, with the expectation being that the mean and standard deviation of the two methods will be within a couple of per cent of each other as the test beam data was centred on C3, so a static cluster there should be a reasonable approximation of the output of the clustering algorithm for this data. The two energy spectra can be seen in Figure 6.3. The increased number of clusters for the C++ algorithm is caused by instances where there is a second cluster in a single event that meets the clustering requirements, and this choice was intentionally made to minimise the possibility of missing overlapping clusters. Some cases were seen in the clustering data where there were two potential seeds in close proximity, with the second seed sharing a corner with the first.

As in Chapter 5, both of these energy spectra were fitted using a Crystal Ball function and the results of that fit are given in Figure 6.3. There is a 0.05% difference in the mean of the two fits and a 2% change in the standard deviation. The clustering algorithm, where clusters are seeded around the crystal with the greatest energy deposit, gives a slight improvement on the standard deviation compared to the fixed clustering method, where clusters are always created around C3 as a seed. This improvement is expected due to the small number of events in the dataset that deposit most of their energy in one of the crystals next to C3.

To make a comparison between the HLS and C++ algorithms for clustering and spike killing, a number of validation plots were generated that allowed for quick and easy comparison. The variables checked were:

- LD spike flag




Figure 6.3: Energy spectra for a 3x3 cluster when clustering is done via the C++ algorithm (top) or via a static cluster with full test beam data precision (bottom). For ease of comparison, the results from the C++ algorithm have been returned to the same scale. Both energy spectra have been fitted with a Crystal Ball function.




Figure 6.4: Plot showing the difference in E_T between the HLS and C++ algorithms, where there was a disagreement between them with regard to overflow.

- Swiss-cross spike flag
- Number of clusters
- Cluster η
- Cluster ϕ
- Cluster E_T
- Number of crystals in a cluster

The clustering algorithms were tested and after some adjustments, both the C++ and HLS algorithms were in complete agreement for all the above variables for all the test beam events used. During testing, disagreements with respect to the number of crystals were found, where one algorithm counted the seed crystal and the other did not. Another disagreement was caused by overflow in E_T , and this example plot is shown in Figure 6.4. In this case, the C++ algorithm was not limited in the number of bits it could give to E_T , while the HLS algorithm was limited in the number of bits and overflow occurred when the value was above 1023 counts. Figure 6.5 shows the agreement in E_T after this bug was fixed.

6.3.2 Spike Killing Algorithm Tests

After the clustering algorithms were tested, both spike-killing algorithms were tested. As there would be limited use in testing them using electron beam data, pion events from the 2018 H2 test beam were used, as pions impacting on APDs cause spikes. An example of the pulse shape and energy deposit for spike events recorded in the crystal matrix can be seen in Figure 6.6.




Figure 6.5: Plot showing the difference in E_T (arbitrary units) between the HLS and C++ algorithms once the overflow issue had been resolved.

The Swiss-cross algorithm test had no discrepancies between the C++ and the HLS algorithms in over 2800 events.

For the LD algorithms, a discrepancy in the spike flag was found in approximately 1 in 2650 samples and is shown in Figure 6.7. Over 44,000 samples were tested in total for the LD, which corresponds to 2750 events. This was investigated and found to occur at the exact threshold that distinguishes between an EM signal and a spike. The working hypothesis is that this discrepancy is due to the limited precision available to the HLS algorithm, with tests planned to confirm this once the firmware has been updated and it is possible to test on a BCP. Even with this minor disagreement, the design is expected to meet the required 99.9% spike rejection as the LD algorithm is nearly 100% efficient (see Figure 4.20).

6.4 Hardware Test of HLS Algorithms

After the work detailed in Section 6.3, the clustering and Swiss-cross algorithms were combined into a single HLS algorithm. The combined algorithm works by scanning across the whole crystal matrix, considering whether each crystal meets the criteria for a Swiss-cross spike or an EM cluster. If one of these thresholds is met, then the Swiss-cross or clustering information is recorded, with all of the outputs of the algorithm sorted after every crystal has been considered. The recorded information is the E_T , η , ϕ , number of crystals in the cluster and the spike flag. The clustering criteria were updated with respect to the previous algorithm to:

- The E_T of any of the surrounding crystals (corners are considered) cannot be greater than the E_T of the seed.




Figure 6.6: Example spike pulse shape from the 2018 pion test beam campaign (top), and example energy deposit in ADC counts for a spike in channel C3 of the crystal matrix from the 2021 pion test beam campaign (bottom). Both were provided by David Petyt.

- The seed must have $E_T \geq 1$ arbitrary count (approximately 0.125 GeV). This threshold can be easily tuned at a later date.
- Seed is not an LD or Swiss-cross spike.

The Swiss-cross algorithm will create Swiss-cross clusters for seeds that are Swiss-cross spikes, which is simply the sum of the energy in the seed crystal and the four crystals it shares a side with. This is useful for debugging purposes to confirm that the Swiss-cross algorithm is working as intended and could be used in the future to tune the thresholds of the algorithm.

As it may be possible to create more clusters than it is possible to send to the L1 trigger due to the limited bandwidth, a sorting algorithm was also added to prioritise sending clusters




Figure 6.7: Difference between the HLS and C++ crystal spike flag for the LD algorithms.

with the highest energy. The Swiss-cross clusters are given a penalty when sorting so all EM clusters are sorted at the top, giving EM clusters priority when sending to the L1 trigger.

The sorting algorithm chosen for this version of the algorithm was a bitonic sort. A bitonic sort is a parallel sorting algorithm that compares and swaps data in an array in a predetermined order to create an ordered array. It has previously been used in other FPGA algorithms with good results [118, 119]. An illustration of the bitonic sort algorithm representation used can be seen in Figure 6.8. In this version of the sorting algorithm, comparisons are always made in the same direction such that the data with the highest E_T is at the top of the array and data with the lowest E_T is sorted to the bottom of the array.

In cases where the array length is not a power of two, then the next highest power of two is used with the remaining inputs set as zeros. As we only expect positive values, then this allows the additional zeros to be placed at the bottom of the array and the comparison logic for that section can be removed, reducing the resource usage.

A small-scale test was performed using the Vivado C simulation to confirm that the algorithm gave the expected output for all events tested. As mentioned before in Section 6.3, the test beam data has 14 bits for E_T rather than the 12 bits ADC input to the BCP, so the number of bits must be reduced. As in Section 6.3, this is done by a combination of rounding and bit-shifting (in this case, essentially a division by eight), this allows the occasional event to saturate and thus allows for this part of the algorithm to be tested also.

As the small-scale C simulation test gave the expected result, the Vivado C simulation was then used to cluster events from a full test beam run of over 50,000 events. The results of this can be seen in Figure 6.9. This is the same test beam run that was analysed in Figure 6.3.




Figure 6.8: Illustration of a bitonic sort of 16 numbers. In the first stage (comparison 1) of the algorithm, pairs of data are sorted into the correct order. The second stage (comparisons 2 and 3) then accordingly sorts two adjacent pairs into four sets of four ordered numbers. The third stage (comparisons 4-6) creates two sets of eight ordered numbers, and finally, the fourth stage (comparisons 7-10) creates a single array of 16 ordered numbers. From [120].

As expected, the newest version of the algorithm creates more clusters than the static C3 cluster method, as some events contain multiple clusters. The E_T distribution given by the newest version of the algorithm was also fitted using a Crystal Ball function, with a mean and a standard deviation compatible with the previous version and the static clustering method. As the clustering criteria have changed slightly between algorithm versions and the data rounding method was also updated, the outputs are not expected to be identical. A by-hand comparison of clusters was completed for around 50 events and the algorithm gave the expected result.

As the simulation provided the expected energy resolution, the design was synthesised into VHDL. The simulation results are given for 120 channels, with the resource estimates given compared to one of the super logic regions (SLR). The VU13P FPGA on the most recent version of the BCP contains four SLRs, which each contain a fourth of the resources. Accordingly, if a 150-channel implementation of an algorithm fits within the required resource usage of a single SLR, then the full 600-channel implementation should fit on the whole board, providing that the SLRs can work independently. Synthesising 150 channels is significantly faster than synthesising 600 channels in Vivado, with the synthesis time correlated to resource usage. For the current algorithm, it takes 1 minute to synthesise 16 channels, 4 minutes to synthesise 32 channels, 35 minutes to synthesise 64 channels, and 6 hours to synthesise 120 channels. The time required to synthesise also depends on the build machine specification (in this case a Dell R740 with 2 x Xeon Golden 6132 and 160 GB memory). This increase in time is mostly driven




Figure 6.9: Energy spectra for a 3x3 cluster when clustering is done using the updated clustering algorithm.

by the non-linear increase in the sorting algorithm.

Target	Estimated	Uncertainty
6.25 ns	5.818 ns	0.78 ns

Table 6.4: Post implementation timing estimates for the combined Swiss-cross and clustering algorithm for 120 channels.

Latency		Interval	
Min	Max	Min	Max
15	15	4	4

Table 6.5: Post implementation latency for the combined Swiss-cross and clustering algorithm for 120 channels.

Resource	Used SLR	Available SLR	Utilisation SLR %
LUT	277627	432000	64
FF	27264	864000	3
BRAM	0	1344	0
DSP	30	3072	~ 0

Table 6.6: Post implementation resource usage for the combined Swiss-cross and clustering algorithm for 120 channels.

The timing estimate meets requirements as shown in Table 6.4. This is achieved by using unrolling where possible in the algorithm. Furthermore, the HLS code uses pipelining to begin each iteration every four clock cycles, as shown by the interval in Table 6.5. This is because, unlike the LD algorithm which considers pulse sample information, the clustering algorithm uses the reconstructed amplitude of pulses. In the final version of the algorithm, this information will be provided by the multifit algorithm every four clock cycles.

As shown in Table 6.6, the combined algorithm uses no BRAM, very limited DSP and only 3% of the available FF in one SLR. The lack of BRAM use is due to the use of memory partitioning, which instead causes data to be stored in the faster FFs. However, even for 120 channels, the algorithm uses 64% of the LUT in the SLR. This value does not scale linearly either, with a 64-channel implementation only requiring 26% of the SLR LUT. This nonlinearity is driven by the bitonic sort, while the resource usage for the Swiss-cross and clustering parts of the algorithm have a linear dependence on the number of channels. As such, a 150-channel implementation will not fit on a single SLR, and therefore the algorithm in its current form is too resource intensive to fit on the BCP alone. Further optimisation is needed for all three algorithms to fit on the BCP. Possible ways to reduce the resource usage of the algorithm are detailed in Section 6.5.

The combined Swiss-cross and clustering algorithm has been tested on hardware using 1000 events from the test beam data detailed in Chapter 5. As there was no BCP board available, the hardware test was completed using an MP7 board instead, which can be seen in Figure 6.2. The MP7 has a native clock speed of 240 MHz, with the links also running at 240 MHz, instead of the 160 MHz sampling rate of the BCP. It also has more limited resources due to the FPGA on the board, and only 72 links. Because of this, the algorithm was synthesised at this increased clock speed of 240 MHz and the input was chosen to be 4×8 channels. Algorithms can run at simple fractions of the native board clock, so it would have been possible to run at 180 MHz, while 160 MHz would have been more complicated. The test beam data itself is based on a 5×5 crystal matrix, though four channels were not operational, so this was reduced to 5×4 (though only losing two channels of data) and the remaining 3×4 was padded with zeros. An example input matrix is shown in Table 6.7.

In addition to the algorithm itself, VHDL code was written to take data in from an input buffer, and similarly store the algorithm output in a buffer. The input data is deserialised from

	A	B	C	D
1	1	2	5	2
2	2	15	47	11
3	2	27	793	24
4	0	7	16	8
5	0	1	1	3
6	0	0	0	0
7	0	0	0	0
8	0	0	0	0

Table 6.7: Example input event to combined Swiss-cross and clustering algorithm. Rows 1-5 contain test beam data, with rows 6-8 padded with zeros. To convert from counts to GeV, all values should be divided by five. For example, the energy of 793 counts deposited in C3 corresponds to 158.6 GeV.

the buffer, and the output data is serialised and written to the output buffer. The buffer of the MP7 can contain a maximum of 256 events (or 1024 frames), and it is also necessary to account for the delay between the input and the output of the algorithm (usually 36 for the payload of this algorithm). Therefore, the test beam data was given to the MP7 in five chunks of 200 events to ensure that none were missed in the buffer. The resource usage of the firmware tested on the MP7 FPGA can be seen in Figure 6.10.

For the 1000 events tested, most were 200 GeV electrons impacting on C3 (see Figure 5.5 for the test beam matrix), with others impacting on one of the surrounding crystals and some events just containing noise. The 1000 events tested mostly contained a single cluster, with 17 events containing multiple clusters for a total of 1020 clusters. The hardware test correctly clustered every single event, including one event where the energy saturated and a few events in which the Swiss-cross flag was set correctly, which gives confidence that both the Swiss-cross and the clustering parts of the algorithm are working well. A few examples of events and their clusters, as given by the algorithm, are shown in Table 6.8.

6.5 Possible Algorithm Improvements

As previously mentioned, the current version of the combined Swiss-cross and clustering algorithm is too resource intensive and this must be reduced in the future. Though the bitonic sort has been used in other HLS algorithms with success, it requires too many resources to be used in this case, with the algorithm sorting every single cluster to only send information on a maximum of $\mathcal{O}(10)$ clusters and spikes per event.

The clusters will be used to aid the L1 trigger in deciding whether an event should be recorded. A study has shown that an energy threshold of 2 GeV gives less than 10 clusters across the whole detector regardless of whether a 3×3 , 3×5 or 5×5 cluster is used [121]. In the current L1 trigger, the seed tower is defined as the tower with the largest transverse energy

	A	B	C	D
1	1	2	5	2
2	2	15	47	11
3	2	27	793*	24
4	0	7	16	8
5	0	1	1	3

	A	B	C	D
1	2	11 [†]	1	1
2	3	9	3	1
3	2	7	18	5
4	0	31	698*	22
5	0	11	66	17

	A	B	C	D
1	0	1	1	0
2	1	2	0	1 [△]
3	1	15 [†]	2	0
4	0	2	1	4
5	0	24	32*	9

Table 6.8: Example events clustered by the combined Swiss-cross and clustering algorithm. The * denotes the primary cluster, with [†] denoting the secondary cluster and [△] denoting a Swiss-cross cluster where applicable. The first event corresponds to the input shown in Table 6.7. The output for the first event is a single cluster of 948 counts (corresponding to 189.6 GeV) centred on C3 with 9 contributing crystals and a spike flag of zero. The second event has two EM clusters. The first cluster has an energy of 875 counts centred on C4 with nine contributing crystals and a spike flag of zero, while the second cluster has an energy of 29 counts centred on B1 with six contributing crystals and a spike flag of zero. In the third event, the first EM cluster has an energy of 72 counts centred on C5 with six contributing crystals and a second EM cluster with an energy of 24 counts with seven contributing crystals and a spike flag of zero. There is also a Swiss-cross cluster in this event, with an energy of 1 count centred on D2 with one contributing crystal and the spike flag set to one. To convert from counts to GeV, all values should be divided by five.




Figure 6.10: The payload for the MP7 for 4x8 channels. The combined Swiss-cross algorithm is shown in yellow, while the deserialisation is shown in pink and the serialisation is shown in green. MP7 infrastructure logic is shown in cyan.

deposit above a threshold of 2 GeV, so an energy threshold of 2 GeV for Phase II is a reasonable value. In normal operations, Swiss-cross clusters will not be used downstream but they will be useful in special debugging runs.

One idea for producing an ordered list of clusters is to start by creating a list of e.g. 10 clusters initialised to zero energy. Then, each found cluster would be compared to this list and if the energy of the cluster is greater than the energy of any of the other clusters in the list, then the new cluster is placed in the position of that cluster instead, with all other clusters moving one down the list and the last cluster being removed to keep the list at 10 clusters. This method should significantly reduce the resource usage of the sorting step. Priority would be given to EM clusters in this list over spikes tagged by the Swiss-cross algorithm as in the current algorithm.




Figure 6.11: Expected number of 3×5 clusters given different energy thresholds. From [69]

Once the resource usage is at an acceptable level, it may also be possible to investigate switching from a 3×3 crystal matrix to a 3×5 crystal matrix to increase the containment of an EM shower in the CMS magnetic field. The energy threshold of the clustering algorithm should also be tuned as the algorithm progresses towards the final version. The expected number of clusters for a 3×5 crystal matrix for some likely energy thresholds can be seen in Figure 6.11.

Future hardware tests will also be important to retest the combined algorithm once these improvements have been made. In addition to testing with electron data as done in Section 6.4, the combined algorithm could also be tested using pion data to confirm that the Swiss-cross algorithm is working as expected using dedicated spike events. Once the multifit algorithm is ready and it is possible to fit both algorithms on a BCP, then tests should also be undertaken to confirm that the algorithm chain behaves as expected. The C++ counterparts to the combined algorithm should also be further developed and used in the CMS simulation to further study the trigger performance, especially for the expected Phase II detector conditions (noise and pileup).

6.6 Conclusion

In conclusion, the two spike-killing algorithms and the clustering algorithm have been tested in software by comparing HLS and C++ counterparts. The clustering algorithm was tested on over

50,000 events from an electron test beam with no discrepancies between the two implementations of the algorithm found. The Swiss-cross algorithm was tested using 2800 events from a pion test beam and no discrepancies were found between the HLS and C++ implementations. The LD algorithm was tested, again using events from a pion test beam, with discrepancies between the two implementations of the algorithm found in approximately 1 in 2650 samples. This was investigated and found to occur at the threshold value that distinguishes between a spike and an EM shower and is thought to be caused by differences in precision between the two algorithms. A test is planned to confirm this hypothesis in the future when the firmware has been updated and a BCP is available for testing. Both spike-killing algorithms work at the required 99.9% efficiency from 2 GeV in the case of the LD algorithm and from 10 GeV for the Swiss-cross algorithm.

The LD algorithm has been tested on hardware using an MP7 board. The results of this test, using 13,400 samples, which correspond to over 800 events, were exactly as expected. Resource estimates were also obtained for the algorithm based on the version of the BCP at the time this test was completed. Based on this test, a second-order LD was chosen over the third-order LD as it gave the required spike-killing efficiency with significantly less resources required.

The Swiss-cross and clustering algorithms were combined and also tested on hardware using an MP7 board. This test found no discrepancies in the output for 1000 events from the same electron test beam used in the software test. However, the bitonic sort used to sort the output EM and Swiss-cross clusters uses a significant amount of FPGA resources, particularly LUTs, to the point where the combined algorithm for 600 channels would not fit on the BCP. As such, further optimisation is required in the future, likely with a more optimised sorting method, consistent with the physics requirements of the algorithms. Development of the algorithms is on track to produce viable versions by the start of LS3 as required, with refinement possible over the course of LS3 and during HL-LHC to optimise based on evolving detector conditions.

Chapter 7

Previous Searches for New Phenomena in Dilepton Final States

As discussed in Section 2.2, there are many BSM models proposed to explain shortcomings in the SM (see Subsection 2.1.5). The most pertinent BSM models to this thesis are ones that propose new physics with a dilepton final state, such as a new Z' gauge boson with a mass at, or beyond that accessible at the LHC. Though not discussed in any detail in this thesis, there are also ongoing searches for other BSM particles such as the W' gauge boson, in final states consisting of a lepton and a neutrino [122].

At the LHC, searches for Z' bosons using proton-proton collisions have been undertaken for several years. The CMS Collaboration has performed searches using Run 1 data recorded at $\sqrt{s} = 7$ TeV [123, 124] and $\sqrt{s} = 8$ TeV [125, 126] and Run 2 data at $\sqrt{s} = 13$ TeV [127–129]. The ATLAS Collaboration has also published results based on data collected at $\sqrt{s} = 7$ TeV [130, 131], $\sqrt{s} = 8$ TeV [132] and $\sqrt{s} = 13$ TeV [133].

At a hadron collider such as the LHC, the decay mode mainly used in searches for the Z' gauge boson is $Z' \rightarrow \ell\ell$, where ℓ is an electron or a muon [134, 135]. Searches do not consider $Z' \rightarrow \tau\tau$ due to the limited tau identification efficiency when minimising the probability of jets, electrons and muons being misidentified as taus [136]. Channels such as the decay of a Z' to two jets [137], $Z' \rightarrow t\bar{t}$ [138] and $Z' \rightarrow e\mu$ [139] can also be used. As the analysis in Chapter 9 is a $Z' \rightarrow \ell\ell$ search, this chapter will focus on the dilepton decay mode.

The dilepton search can be split into two categories: resonant and non-resonant. If the Z' boson mass is accessible at the LHC, then a resonant analysis would be capable of making a discovery, while non-resonant analyses search for Z' boson masses beyond the energy provided by the LHC. Previous analyses in these two categories are described in Section 7.2 and Section 7.3 respectively. Both of these sections focus on results from the most recently published CMS analysis [129], and as such Section 7.1 also gives the criteria used by CMS in that analysis. This analysis also considered a test for lepton flavour universality to complement both the resonant and non-resonant Z' searches [129]. This is discussed in more detail in Section 7.4.

7.1 Backgrounds, Event Simulation and Event Selection

For Z' searches in CMS, background processes are estimated from MC with the exception of the background from jets misidentified as electrons, which is instead estimated using a control region in the data. The dominant source of background is the Drell-Yan (DY) process, which was described in Section 2.4. This is also used in the analysis detailed in Chapter 8. The background is normalised independently for the dielectron and dimuon channels by considering a window around the Z boson peak of $60 < M(\ell\ell) < 120$ GeV and normalising according to the number of events observed in the data in that region.

To model the DY process as accurately as possible in the most recent published CMS results [129], NLO Powheg is used, with NNLO corrections applied, and with parton showers and hadronisation processed by Pythia. More information on the simulation of events can be found in Section 2.5. This combination of NLO Powheg, Pythia and NNLO corrections is also used for the background processes $t\bar{t}$, tW (production of a single top in association with a W), and WW. The background processes WZ and ZZ are also studied, with the samples used for 2016 data generated using LO Pythia, and samples used for 2017 and 2018 data generated using a combination of NLO Powheg and aMC@NLO. Finally, W production in association with jets is simulated using aMC@NLO to LO. Both the resonant and the non-resonant signal samples are generated to LO using Pythia.

Electrons are required to have $p_T > 35$ GeV and have $|\eta_{SC}| < 1.442$ or $1.566 < |\eta_{SC}| < 2.5$, where η_{SC} is the pseudorapidity of the ECAL supercluster. More information on how superclusters are defined can be found in Subsection 4.2.1. This η_{SC} requirement excludes the overlap region between the barrel and endcaps which has a reduced efficiency. Electrons must also pass the high energy electron/positron identification criteria (HEEP ID), a detector-based set of selection criteria designed to reject poorly isolated or misreconstructed electrons and electrons originating from jets. Further details of the HEEP ID can be found in Chapter 8. Events must have at least one electron in the barrel because of the high rate of jets reconstructed as electrons in the case where both electrons are detected in the endcap. There is no requirement on the charges of the electrons as the probability of sign misidentification increases with energy, with the misidentification probability unacceptably high at TeV energy scale considered in Z' searches.

Muons are required to have $p_T > 53$ GeV, $|\eta| < 2.4$ and to pass a high energy muon ID criterium. The muon p_T requirement is larger than the electron p_T requirement because of the differences in the p_T requirements of their respective triggers. Unlike electrons, the dimuon pair are required to have charges of opposite sign as a misreconstructed charge is likely to be caused by the poor resolution of the track fit which would in turn give an inaccurate muon p_T measurement. The probability of assigning the incorrect sign to the charge of a 2 TeV muon is around 10^{-4} [140]. This low misidentification probability has a negligible effect on signal efficiency.




Figure 7.1: Feynman diagram for the s-channel production of a dilepton pair via a Z' resonance.

In cases where there are three or more leptons in a given event, then the pair with the mass closest to that of the Z boson is chosen if the mass is within 20 GeV of the Z mass. This is done to suppress the diboson background at high mass. Otherwise, the two highest p_T leptons are chosen.

7.2 Resonant Searches

In resonant searches, the expected physics signature is a narrow Z' boson resonance produced in the s channel and decaying to a dilepton pair. A Feynman diagram for this resonance is shown in Figure 7.1. These analyses focus on a mass window centred around the assumed Z' mass. A Z' boson in the few TeV range solves the fine-tuning problem and resonant searches to date have placed limits at around this energy, though the higher end of this range has not yet been excluded. The size of the mass window depends on both the resolution of the detector and the intrinsic decay width of the Z' model under investigation. As mentioned in Section 2.2, there are several different Z' models. To account for this, resonant searches investigate a range of masses and decay widths. Unbinned maximum likelihood fits are used to obtain the background normalisation for each mass window.

To reduce uncertainties from both theoretical and experimental sources, limits are not set directly on the cross-section and branching fraction of the Z' , but rather on the ratio of the Z' cross-section and branching fraction to those of the Z boson. Uncertainties that are a function of dilepton mass must still be considered.

The signal is modelled by the convolution of a Breit-Wigner function and a double-sided Crystal Ball function [107], where the Breit-Wigner function models the decay width of the resonance and the Crystal Ball function models the instrumental resolution, which is dependent on the invariant mass of the dilepton pair.

Three types of model are expected to produce a resonant signature in such an analysis are considered. The first is a new Z' gauge boson, with more detail on possible models given in Subsections 2.2.1 to 2.2.2. Next is a spin-2 graviton which is predicted by the Randall–Sundrum (RS) model of extra dimensions, detailed in Subsection 2.2.5 [27, 38]. The experimentally expected signature is the resonant production of spin-2 Kaluza-Klein (KK) excitations of

gravitons, G_{KK} , which will be detected by their subsequent decay to dileptons. KK refers to a unified gravitational and electromagnetic theory based on five dimensions which predicts the existence of gravitons [141]. These spin-2 resonances have a coupling $k/\overline{M_{pl}}$, where k is the warp factor of the five-dimensional space predicted by the RS model and $\overline{M_{pl}}$ is the reduced Planck mass. The results are quoted with regard to both this coupling and the resonance mass.

Finally, resonances at high dilepton masses may be caused by spin-1 dark matter (DM) mediators [142]. The model considered consists of a single DM particle at the TeV scale with a single spin-1 mediator particle that is either a vector or an axial-vector boson [143].

If a discovery is made, it would be possible to distinguish between these models by analysing the width of the resonance and the forward-backward asymmetry of candidate events. Forward-backward asymmetry compares the number of candidate events in the forward region of $\cos\theta > 0$ to the number of candidate events in the backwards region of $\cos\theta < 0$, where θ is the scattering angle. It is also possible to use forward-backward asymmetry to search for Z' bosons [144].

The most recent CMS results using data from 2016, 2017 and 2018 for the resonant production of the Z' boson are shown in Figure 7.2. For the combined dielectron and dimuon channel, the expected limit on the Z'_{SSM} is 5.14 TeV with an observed limit of 5.15 TeV. Similarly, the combined channels place an expected limit of 4.55 TeV on the Z'_ψ and an observed limit of 4.56 TeV.

Results using the same data for spin-2 KK excitations of gravitons are shown in Figure 7.3. The limits for these spin-2 excitations vary depending on the coupling, with a coupling of $\overline{M_{pl}} = 0.01$ giving an expected limit of 2.53 TeV and an observed limit of 2.47 TeV, and a coupling of $\overline{M_{pl}} = 0.1$ giving an expected limit of 4.78 TeV and an observed limit of 4.81 TeV.

The upper limits on the masses of the DM particle are shown in Figure 7.4. A variety of decay widths were considered from 0.5% to 3.5%, with the narrowest decay widths providing the strongest limits. In the vector model, mediators with a mass below 1.92 TeV are excluded at 95% CL. In the axial-vector model, the exclusion limit for the mediator mass is set at 4.64 TeV at 95% CL. These limits are similar to those set by ATLAS [145].

7.2.1 Resonant Search in Association with b tagged jets

The search for $Z' \rightarrow \ell\ell$ has recently been extended to consider b-jets produced in association with the dilepton final state. This search is motivated by the increased sensitivity of the third generation of quarks to BSM physics due to their higher masses. Experiments such as LHCb have observed anomalies in b-hadron decays that suggested lepton flavour violation [147], though more recent measurements of the ratio of these decays have been consistent with the SM [148]. As previously mentioned, in the inclusive dilepton final state the background is dominated by the DY process. By requiring b-jets in the final state, there are additional initial states that can produce a Z' boson to then decay to a dilepton pair and b-tagged jets. These additional initial states can therefore reduce the DY background. The Feynman diagrams for the production of a




Figure 7.2: The upper limits at 95% CL obtained by CMS on the production cross section multiplied by the branching fraction for a Z' boson relative to that of the Z boson. The Z' is modelled as a spin-1 resonance with an intrinsic decay width of 0.6% of the resonance mass in both the dielectron and dimuon channels. The limit is shown for the dielectron channel (upper left), dimuon channel (upper right) and the two channels combined (bottom). The green shaded area denotes the 68% quantile for the expected limit, while the yellow area denotes the 95% quantile. The predictions for two proposed Z' bosons, Z'_{SSM} and Z'_ψ are also shown. From [129].

dimuon pair are shown in Figure 7.5, there are equivalent diagrams for the production of a dielectron pair.

ATLAS has published results on this dilepton and b-jet final state, in the channel $Z' \rightarrow Zh \rightarrow \ell\ell b\bar{b}$, where h is the Higgs boson [150]. The benchmark model is a simplified model known as the heavy vector triplets (HVT) model, which adds an $SU(2)$ triplet for the W' and Z' bosons, with additional parameters in the model to describe the coupling of these new bosons to fermions, gauge bosons and the Higgs boson. The HVT model is split into Model A, where the branching fractions to fermions and gauge bosons are similar, and Model B in which the gauge boson branching fraction is higher than the fermionic branching fraction. The results of the ATLAS search can be seen in Figure 7.6. The ATLAS analysis excludes Z' bosons at masses below 2.67 TeV for HVT model A, and below 2.83 TeV for HVT model B [150].

CMS is expecting to publish results soon using Run 2 data from a resonant search for the




Figure 7.3: The upper limits at 95% CL obtained by CMS on the production cross section multiplied by the branching fraction for a spin-2 KK excitation of a graviton relative to that of the Z boson. The limit is shown for the dielectron channel (upper left), dimuon channel (upper right) and the two channels combined (bottom). The green shaded area denotes the 68% quantile for the expected limit, while the yellow area denotes the 95% quantile. The prediction for the spin-2 KK excitation is shown for coupling parameters, $k/\overline{M}_{Pl}$, of 0.01 (pink), 0.05 (red) and 0.10 (black). From [129].

decay of a Z' resonance via $Z' \rightarrow \mu\mu$ associated with at least one b-tagged jet [149].

7.3 Non-Resonant Searches

In non-resonant searches, the expected physics signature is an excess of events in the high-mass dilepton tails. This is illustrated for a W' boson search in Figure 7.7. To perform this type of search, the samples are split into different bins according to their invariant mass. More recent analyses in CMS also bin events according to their scattering angle in the Collins-Soper frame [129, 151]. The results from all of these bins are combined to give the final results. The bins with the highest mass values provide the largest signal-to-background ratios and are therefore the bins that are most sensitive to BSM physics. The bins with lower mass values can provide insight into interference between a Z' and SM physics. Unlike in the resonant search, the cross-section of the Z boson is not used to normalise the signal, nor can the background estimate be obtained




Figure 7.4: The upper limits at 95% CL obtained by CMS on the masses of the DM particle and the DM mediator. Exclusion regions were found by comparing the cross section multiplied by the branching fraction for the decay to a Z boson to values given by the simplified model. The grey curves represent regions that are consistent with the observed DM relic density [146]. From [129].




Figure 7.5: Feynman diagrams for the decay of a Z' to a dimuon pair in association with at least one b-quark. From [149].




Figure 7.6: The upper limits at 95% CL obtained by ATLAS on the production cross section multiplied by the branching fraction for a Z' to that of Zh and the branching fraction $B(h \rightarrow b\bar{b}, c\bar{c})$. The limit is shown for the dimuon channel only. The green shaded area denotes the 68% quantile for the expected limit, while the yellow area denotes the 95% quantile. The prediction for HVT model A (dashed red) and HVT model B (solid pink) are also shown. From [150].

from the data. Because of this, non-resonant searches are more sensitive to uncertainties in the modelling of signal and background.

The models most often tested in non-resonant searches are the contact interaction (CI) model and the Arkani-Hamed, Dimopoulos and Dvali (ADD) model, which are described in Subsections 2.2.4 and 2.2.5 respectively. The three conventions used to parameterise the cross-section in the ADD model are Hewett [28]; Giudice, Rattazzi and Wells (GRZ) [29] and Han, Lykken and Zhang (HLZ) [30]. For these signatures, the CMS collaboration has published results at 8 [126] and 13 TeV [153]. The ATLAS collaboration has published results for the ADD model at 8 TeV [154] and for the CI model at 13 TeV [155]. In addition to the dilepton search, constraints for both models have been provided by dijet final states from both the ATLAS and CMS Collaborations, with the CMS results additionally considering diphoton final states [156–158].

Figure 7.8 shows the exclusion limits at 95% CL for the CI model. As detailed in Subsection 2.2.4, the CI model is split into eight models, each with their own limit. For the combined dilepton channels, the limit on the ultraviolet cutoff parameter, Λ , is set between 23.9 and 36.4 TeV, dependent on the exact CI model under consideration [129].

The exclusion limits at 95% CL for the ADD models are shown in Figure 7.9. The observed limit for the combined dielectron and dimuon channel for the GRW convention is 7.5 TeV, while for the Hewett convention, the observed limit is 6.7 TeV. In the HLZ convention, extra




Figure 7.7: The expected and observed transverse mass distribution for the electron distribution in the search for a W' boson in final states with a lepton and missing transverse energy. From [152].

dimensions from $n = 3$ to $n = 7$ are considered. The $n = 2$ case is not considered in the analysis due to gravitational and astrophysical experiments which have provided limits on this case that are beyond the accessible energy range [159, 160]. For the $n = 3$ case, exclusion limits at 95% are placed at 8.9 TeV, for the combined channels. This limit decreases with an increasing number of extra dimensions to the $n = 7$ case, which places exclusion limits at 5.9 TeV.

7.3.1 Non-Resonant Search in Association with b tagged jets

As with the resonant searches, it is also possible to perform a non-resonant search for new physics with a dilepton final state in association with b-tagged jets. ATLAS has published results at $\sqrt{s} = 13$ TeV on the transition $b \rightarrow s\ell\ell$ mediated by a Z' boson in the context of the CI model and placed exclusion limits on Λ at 2.0 TeV for the dielectron channel and 2.4 TeV for the dimuon channel [161]. The upper limit on the cross-section for electrons and muons is shown in Figure 7.10. CMS is also searching for new phenomena in this final state, with results presented in Chapter 9.




Figure 7.8: Lower limits at 95% CL obtained by CMS on the ultraviolet cutoff parameter, Λ , for the eight CI models considered. Shown for (upper left) the dielectron channel, (upper right) the dimuon channel, and (lower) their combination. Limits for all models are given for $m_{\ell\ell} > 400$ GeV. From [129].

7.4 Lepton Flavour Universality

It is also possible to search for lepton flavour universality violation by considering the ratio

$$(7.1) \quad R_{\mu^+\mu^-/e^+e^-} = \frac{d\sigma(q\bar{q} \rightarrow \mu^+\mu^-)/dm_{\ell\ell}}{d\sigma(q\bar{q} \rightarrow e^+e^-)/dm_{\ell\ell}}$$

which is predicted to deviate from unity at high dilepton mass in some models. Models that cause this deviation include leptoquarks [162] and heavy neutral gauge bosons [163]. The ratio for CMS Run 2 data can be seen in Figure 7.11. This ratio is in good agreement with unity as expected until 1.5 TeV, after which there are large statistical uncertainties due to the limited number of events in this energy range. A χ^2 test was performed with a null hypothesis that lepton flavour universality was not violated, giving a χ^2 per degree of freedom value of 17.9/7 for the combination of the two search regions, corresponding to a one-sided p-value of 0.012 [129] ($\sigma = 2.3$). Only the region from 400 GeV onwards was considered as the models predict negligible deviations from the SM prediction of unity in the lower mass region [164].




Figure 7.9: Exclusion limits at 95% CL obtained by CMS on the ultraviolet cutoff for the three ADD models: GRW (first bin), Hewett (second bin) and HLZ (third to seventh bin). Shown for (upper left) the dielectron channel, (upper right) the dimuon channel, and (lower) their combination. The signal models are calculated at L), with an NNLO correction factor of 1.3 applied. From [129].




Figure 7.10: Upper limits at 95% CL obtained by ATLAS on the visible cross-section, σ_{vis} , for the (a) electron b-veto, (b) electron b-tag, (c) muon b-veto, and (d) muon b-tag categories. The green shaded area denotes the 68% quantile for the expected limit, while the yellow area denotes the 95% quantile. The red marker gives the strongest expected lower limit. From [161].




Figure 7.11: Ratio of the differential dilepton production cross section obtained by CMS as given in Equation 7.1, shown as a function of the dilepton pair mass for the case where both leptons are in the barrel (top left), at least one lepton is in the endcap (top right) and the combined result (bottom). The error bars represent both statistical and systematic uncertainties. From [129].

Chapter 8

HEEP ID Using Ultra Legacy Data

In this chapter, I investigate the efficiency of the high-energy electron-positron identification (HEEP ID) criteria used by CMS for Run 2 with recalibrated data sets. These recalibrated data sets are referred to as Ultra Legacy (UL) data sets and are better-calibrated samples than so-called end-of-year (EOY) samples, with MC that is more reflective of the conditions in the detector at the time. This recalibrated data offers improved resolution and reconstruction of particles, including electrons [84], which means that new physics searches will have increased sensitivity when using this data compared to the EOY Run 2 data samples.

The HEEP ID is used in analyses such as the search for the process $Z' \rightarrow ee$ and $W' \rightarrow e\nu$ [165]. Because the HEEP ID is designed for use with high energy electrons where there are a low number of events, it is a simple cut-based ID. The aim of the analysis presented in this chapter is to check that the energy evolution seen in data is well-modelled in Monte Carlo.

The HEEP ID places cuts on well-understood detector-based variables and is designed to be robust to changes in detector conditions, including changes in pileup. The cuts of the HEEP ID are chosen such that the ID has a high efficiency of 90% or more for electrons with $E_T > 50$ GeV from Z decays, and the cuts are tuned to give a flat efficiency with respect to E_T . This allows analyses to extrapolate to > 1 TeV electrons and still expect the HEEP ID to have the same efficiency as it did below 1 TeV with some level of confidence.

In Section 8.1, the difference between EOY and UL samples is explained. In Section 8.2, all relevant information regarding the data and MC used for this analysis is listed. Section 8.3 covers the methods used in this analysis, with the HEEP ID detailed in 8.3.1 and the tag and probe method that is used to calculate the efficiency of the HEEP ID explained in 8.3.2. Data-driven estimation of the QCD background is explained in Section 8.4. In Section 8.5, all the results from this analysis are detailed. This section is split into Section 8.5.1 which covers the kinematic distributions, followed by Section 8.5.2 which shows the efficiencies and scale factors (the ratio between the data and MC efficiencies) of the HEEP ID. Section 8.5.3 illustrates the effect of individual cuts within the HEEP ID and finally, Section 8.5.4 compares UL with EOY results from a previous analysis. All work in this chapter is mine except for the

previously mentioned EOY results.

8.1 Ultra Legacy Samples

Shortly after every data-taking year at CMS, an EOY data set is produced. The EOY MC samples have to be created in advance for them to be ready in time. As such, EOY MC samples do not include imperfections caused by issues that developed during the course of data taking that particular year. After Run 2, CMS recalibrated the full data set to create the UL samples. This process involved recalibrating every subdetector of the CMS experiment and all physics objects. The simulation of each year was also redone using updated detector conditions, including the addition of detector problems. The modelling of dynamic inefficiencies, radiation damage and detector noise is also improved.

The updated measurements in the ECAL improve the reconstruction and identification of both electrons and photons, giving both a better resolution and improved data/MC agreement. The ECAL improvements are due to more granular calibration using the azimuthal symmetry of minimum bias events and standard model (SM) candles, where well-known resonances such as π^0 and Z decays are used to calibrate the energy response. See Section 4.2.3 for more details on ECAL calibration. An updated correction for the ECAL crystal transparency loss that is caused by radiation damage also improved ECAL reconstruction and identification. As shown in Figure 8.1, UL samples provide a relative resolution improvement of over 50% for high η values. For 2017, UL samples have also been shown to give improved data/MC agreement for electron ID variables [84].

Another area which has seen significant improvement in the UL dataset is the identification of heavy flavour jets, which is an essential component of searches for new physics beyond the SM, such as the search described in Chapter 9, as well as for measurements of SM processes. Heavy flavour jets are jets that contain b or c hadrons. As with the improvements seen in UL samples for electrons, this recalibrated data also improves the reconstruction of heavy flavour jets. Heavy flavour jet identification is performed using multi-classification deep-neural-network algorithms [75, 166, 167]. These algorithms give better agreement between data and MC using UL samples and a reduced light jet mistag rate in data over a wide p_T range compared to EOY samples. The b-tagged jet scale factors are also closer to 1 at high jet p_T [168]. This improvement is of particular interest to searches for models that involve b-tagging, such as the search for $Z' \rightarrow \ell\ell$ in association with b-tagged jets as the addition of b-tagged jets offers improved sensitivity compared to an inclusive search. With EOY data, it was only possible to complete an inclusive search for $Z' \rightarrow \ell\ell$.




Figure 8.1: Dielectron mass resolution from $Z \rightarrow ee$ events as a function of η of the electron. UL (in this plot labelled as Data (Legacy)) data is shown in green with EOY data (in this plot simply labelled Data) is shown in pink. The grey banded region is the transition between the barrel and the endcap region of the ECAL at $1.44 < |\eta| < 1.57$ and is not included. Plot taken from [84].

8.2 Data and MC Samples

This analysis uses data from three years of proton-proton collisions at the CMS Experiment between 2016 and 2018, known as Run 2. The centre of mass energy for all this data is 13 TeV. The integrated luminosity for 2016 was 35.93 fb^{-1} , while for 2017 it was 41.48 fb^{-1} and finally for 2018 it was 59.83 fb^{-1} [169]. Scale factors are calculated on a per-year basis to allow them to be applied on a yearly basis in line with the MC that is generated for each year of data taking.

The MC samples used in this analysis can be seen in Table 8.1. The dominant physics process is Drell Yan (DY) production of dilepton pairs, with other physics processes such as W +Jets, the production of a W boson in association with jets, also contributing. DY production occurs when a quark in one hadron annihilates with the antiquark of another hadron, producing a photon or Z boson which then decays to a dilepton pair. In our case, the DY background is predominantly dielectron pairs, though dimuon pairs also contribute a very small amount. All MC samples were centrally produced and are recommended for use in UL analyses by the Physics Data and MC Validation group of CMS. The MC samples were chosen to match as closely as possible those used in the EOY analysis, which is largely reflected in the choice of

Physics Process	Generator	Cross section (pb^{-1})
Drell Yan	MC@NLO	6077
W+Jets to LNu	madgraphMLM-pythia8	53550
$t\bar{t}$ to Semi Leptonic	powheg-pythia8	364.4
$t\bar{t}$ to 2L2Nu	powheg-pythia8	87.3
ST top	powheg-pythia8	35.6
ST anti top	powheg-pythia8	35.6
GJets HT 40 to 100	madgraphMLM-pythia8	20790
GJets HT 100 to 200	madgraphMLM-pythia8	5036
GJets HT 200 to 400	madgraphMLM-pythia8	1128
GJets HT 400 to 600	madgraphMLM-pythia8	126.4
GJets HT 600 to inf	madgraphMLM-pythia8	41.24
WW	pythia8	118.7
WZ	pythia8	27.56
ZZ	pythia8	12.14

Table 8.1: All the MC samples used in this analysis, the generators used and their cross-sections.

generators shown.

8.3 Method

8.3.1 HEEP ID

Variable	Barrel	Endcap
E_T	$> 35 \text{ GeV}$	$> 35 \text{ GeV}$
η range	$ \eta_{SC} < 1.4442$	$1.566 < \eta_{SC} < 2.5$
ECAL Driven	yes	yes
$ \Delta\eta_{in}^{seed} $	< 0.004	< 0.006
$ \Delta\phi_{in} $	< 0.06	< 0.06
H/E	$< 1/E + 0.05$	$< 5/E + 0.05$
full 5x5 $\sigma_{i\eta i\eta}$	not applicable	< 0.03
full 5x5 E^{2x5}/E^{5x5}	$> 0.94 \text{ OR } E^{1x5}/E^{5x5} > 0.83$	not applicable
ECAL Isolation and Hadronic Depth Isolation 1	$< 2 + 0.03E_T + 0.28\rho$	$< 2.5 + 0.28E_T + 0.28\rho$ for $E_T < 50$ else $< 2.5 + 0.03E_T + 0.28\rho$
Track Isolation: Track P_T	< 5	< 5
Inner Layer Lost Hits	≤ 1	≤ 1
$ d_{xy} $	< 0.02	< 0.05

Table 8.2: All the selections included in HEEP ID V7.0

The full criteria for the HEEP ID used in this analysis (HEEP ID V7.0) can be seen in Table 8.2. I will now explain the purpose of all the cuts that comprise the HEEP ID. The criteria can be split into three sections: acceptance, identification and isolation. The acceptance criteria place cuts on η_{SC} and E_T , while the identification criteria focus on shower shape variables such

as H/E . As the name suggests, the isolation criteria focus on variables such as track p_T for track isolation.

First, the E_T and η cuts in the HEEP ID simply cut out low energy electrons and remove the gap region between the barrel and the endcap, as well as setting a limit at $|\eta| < 2.5$ above which the tracker coverage is too limited.

In CMS, electrons can either be classified as ECAL driven or tracker driven (determined by whether the tracker was seeded by a trajectory from a pixel match and supercluster or a standard track). The ECAL driven selection refers to the case where the supercluster led the electron reconstruction. The HEEP ID requires electrons to be ECAL driven as this electron classification method is designed for high-energy isolated electrons, while particle flow is useful for low-energy non-isolated electrons and recovering regions where the pixels were not fully functional. ECAL driven is highly efficient for high-energy electrons, while the tracker driven selection is largely irrelevant at high energies.

The next variable used in the ID, $|\Delta\eta_{in}^{seed}|$, is the difference in η between the track position as measured by the inner tracker layer extrapolated out to the surface of the ECAL compared to that measured by the ECAL. Similarly, $|\Delta\phi_{in}|$ is the difference in ϕ between the position measured by the inner tracker extrapolated out to the surface of the ECAL and that measured by the ECAL. These cuts are specifically placed on the seed because the supercluster position will be biased due to bending in the magnetic field, particularly for ϕ .

Both of these variables are sensitive to the relative alignment of the ECAL and the tracker, which is also improved in the UL samples.

A cut is also applied on H/E , i.e. the ratio of energy deposited in the HCAL to that deposited in the ECAL. For electrons, we expect almost all of the energy deposited to be in the ECAL, so a large fraction in the HCAL suggests a different particle or a poorly measured electron, neither of which we want.

Next is $\sigma_{i\eta i\eta}$, a measure of the spread in η of the electron's energy in the 5×5 crystal matrix with the seed crystal at its centre. This spread is measured in units of crystals and the measurement is solely used in the endcap. In the barrel, the full 5×5 $E^{2\times 5}/E^{5\times 5}$ is used instead. This is simply the ratio between the energy in the energy deposited in a 2×5 crystal matrix centred around the seed, compared to a full 5×5 crystal matrix centred around the seed. A 2×5 crystal matrix is chosen in case the electron impacts on the side of a crystal and the bulk of the energy is split across multiple crystals, but there is an alternative 1×5 possibility for electrons that impact more centrally on the seed crystal. For a well-isolated electron, we expect the majority of its energy to be deposited in the crystals directly next to the seed.

A measure of ECAL isolation comes from considering the transverse electromagnetic energy of the hits close to the electron under consideration. A requirement is placed on the transverse electromagnetic energy of all reconstructed hits with an energy > 0.08 GeV in the barrel and > 0.1 GeV in the endcap in a cone of radius 0.3 centred on the electron's position in the

calorimeter, but excluding those in an inner cone of radius three crystals and eta strip of total width of three crystals. This is done to consider all the energy in the vicinity of the electron while excluding the energy deposits directly due to the electron. This is only used in conjunction with hadronic depth isolation 1, which is defined as the hadronic energy deposited in the first transverse depth of the HCAL (there are three depths in Run 2) in a cone of radius 0.3 centred on the electron's position in the calorimeter. As with ECAL isolation, the energy within a cone of radius 0.15 is excluded from the sum. The average momentum per unit area, ρ , is used as a measure of pileup. As there is a linear term involving ρ in this cut, the requirement is looser when there is higher pileup and hence more activity from other interactions in the event.

Track p_T is defined as the p_T sum of tracks in a cone of radius 0.3 with $p_T > 0.7$ GeV/c, excluding a strip in η of width 0.03. HEEP candidates fail this cut due to insufficient tracker isolation. An additional exclusion criterion is also applied using the variable z_0 , where z_0 is defined as the minimum distance in z from the primary vertex. This cut is ± 0.1 cm in the barrel and ± 0.5 cm in the endcap.

As the name suggests, inner layer lost hits is a requirement on the number of missing hits in the tracker. Lost hits in the tracker are primarily used to reject photon conversions, as missing hits in multiple layers suggest that the track did not start from the interaction point but rather started later due to photon conversion within the tracker. As the tracker has a small inefficiency, a single missing hit is allowed and non-functional tracker modules are not considered. Finally, the transverse impact parameter, d_{xy} , is defined as the distance in the transverse plane of the track from the interaction point and again a large value of this parameter is sensitive to photon conversions. If this distance is too large, the detected electron is unlikely to come from the interaction point.

8.3.2 Tag and Probe Method

The tag and probe method is used to evaluate the efficiency of the HEEP ID using the resonance $Z \rightarrow e\bar{e}$. In this instance, the tag is a well-identified electron that meets the selection criteria given below, and the probe is another electron in the event that is then used to measure the efficiency of the ID.

These are the tag and probe criteria used for evaluating the HEEP ID efficiency:

- Tag passes HEEP ID V7.0
- Tag is a barrel electron (i.e. $|\eta_{SC}| < 1.4442$)
- Tag is matched to a 32 GeV electron trigger in data (i.e. the tag must have passed the selection criteria of this particular trigger).

- Note that in CMS a trigger of a given energy is designed such that the trigger is approximately 50% efficient at that uncorrected energy. Please see Section 3.3.6 for more details on triggering in CMS.
 - In MC, a turn-on curve of this particular electron trigger from data is applied to emulate this.
 - A data-driven method is used to account for any imperfections in the modelling of this efficiency in the MC. The trigger is not a step function due to the differences between the online and offline electron identification and reconstruction.
 - A 27 GeV trigger is used in 2016 which is not available in later years due to changes in the trigger menu to reoptimise for higher pileup. This is not expected to have any appreciable effect on results due to the 35 GeV E_T cut in the HEEP ID.
- Probe passes the acceptance criteria (η_{SC} and E_T) of HEEP ID V7.0
 - The invariant mass of the tag and probe is in the window $70 < M(ee) < 110$ GeV/ c^2 (Note the Z mass is ≈ 91 GeV)
 - All possible combinations of tag and probe in an event are selected

As can be seen in Figure 8.4, this tag and probe criteria gives a clear peak around the Z boson mass, giving confidence that I am selecting real electrons from the Z boson decay as intended with limited contamination from other SM processes.

The efficiency is defined as

$$(8.1) \quad \epsilon = \frac{N_{passingprobes}}{N_{allprobes}},$$

where $N_{passingprobes}$ is the number of probes that pass the HEEP ID V7.0, while $N_{allprobes}$ is the total number of probes (i.e. the sum of the passing and the failing probes). This efficiency is measured in both data and MC as functions of the following electron variables: η_{SC} , ϕ_{SC} , E_T and nPV, where nPV is the number of primary vertices and an estimate of pileup. The primary method for measuring the efficiency is by subtraction of all the non DY MC and the QCD contribution from the data, which is then compared to the DY MC. All the MC used can be seen in Table 8.1 and the data-driven method to calculate the QCD contribution is detailed in Section 8.4. A crosscheck is done by comparing the data to the combined MC as this should also provide the same scale factor, with a measurement of the scale factor being the primary aim of this study.

Due to imperfections in the MC simulation, there are small differences between the data and MC electron resolution and response. The data also has conditions that vary over time, causing a slight change in response while the MC conditions are static though designed to model average conditions for the given data-taking period. To correct for both of these effects,

energy corrections known as scale and smearing corrections are applied. Scaling is where the energy of electrons in data is scaled to match the response seen in MC samples while smearing is where MC events are smeared to match the resolution seen in data. Scaling is dependent on lumisection, η and shower shape, while smearing is only dependent on η and shower shape (as the MC does not have a lumisection). The electron energy used throughout this analysis is the energy deposited in the ECAL.

Due to concerns regarding inaccuracies in the modelling of the turn-on region of the HLT trigger in MC, the MC events are weighted using the trigger turn-on curve from the data. A turn-on curve that is only 1% different from that seen in data will cause all the distributions to have worse data/MC agreement. This is particularly pronounced in the E_T distribution where there is an offset between the peak in data and the peak in MC that causes kinks in the tails of the $M(ee)$ distributions. These distribution shifts come directly from reweighting the MC with an incorrect turn-on curve. For example, a HLT turn-on curve that is shifted towards lower E_T will give a reweighted MC curve that is similarly weighted towards lower E_T , with more events passing than expected also once all the other reweightings have been applied. These effects are minimised by applying a data-driven HLT turn-on curve to MC events.

The HLT turn-on curve is also generated using the tag and probe method. In this instance, the tag is required to be a barrel electron that passes a 32 GeV electron trigger (27 GeV for 2016). This time, the probe is required to be in the barrel (as it will be applied to the tag in the HEEP tag and probe method, which also has this requirement). The turn-on curve is made by plotting the efficiency as defined in Equation 8.1 as a function of E_T . The turn-on curve measured from data is applied to MC as follows. A random number between 0 and 1 is generated for each potential tag in the MC, and if the random number generated is less than the measured trigger efficiency at the energy of the tag then the tag is saved. Otherwise, the tag is removed.

The turn-on curve for 2018 data can be seen in Figure 8.2. As expected, this trigger does not accept any data below ≈ 30 GeV with a sharp turn-on region between 30 and 40 GeV, then smooths out to approximately 90% efficient. While this is a 32 GeV trigger, the trigger places cuts on the uncorrected energy, while the corrected energy as plotted is usually 1-2% higher. As such, the 32 GeV trigger being 50% efficient around 35 GeV is expected.

All the MC events were generated with a pileup distribution similar to that seen in data. However, the MC pileup distribution is reweighted to account for the difference between the data and MC pileup distributions. This is done by applying a weight to each event where the weight is calculated from the ratio between the normalised pileup distributions of the data and the MC. With this construction of the weight, the overall number of events is kept constant but the reweighted distribution in MC more closely matches that seen in the data.

When generating MC events as many events as possible are computed, which is often more events than recorded in the data the MC is compared to. This provides better statistics for the




Figure 8.2: Turn-on curve for a 32 GeV electron trigger in the barrel measured using a tag and probe method with $Z \rightarrow ee$ events using 2018 data.

MC but does therefore require that the events are reweighted accordingly. For a MC sample where all events have a weight of 1, this reweighting is simply

$$(8.2) \quad weight_{X_{sec}} = \frac{\sigma \times \mathcal{L}}{N_{gen}}$$

where $\mathcal{L}$ is the luminosity and N_{gen} is the number of generated events.

However, in MC where weights also have to be accounted for, this formula has to be reworked as the number of generated events needs to be replaced with the sum of the weights of all the generated events, where w_i is the weight of an event. For example, the 2018 DY sample had to be reweighted by a factor of approximately 0.00007 to account for all the events that were generated and their associated weights.

$$(8.3) \quad weight_{X_{sec}} = \frac{\sigma_i^{N_{gen}} \times \mathcal{L}}{\sum w_i}$$

8.4 QCD Estimate

MC does exist for QCD events, where both the tag and probe are jets misidentified as electrons or positrons. However, the number of events for this process is very low because of the low fake rate, so a fake rate measurement from MC would require the generation of many events. Instead of using MC, the QCD background is calculated using a data-driven method as described below. At the sensitivity level of this study, we assume that the probability of a jet being misidentified

as an electron is equal to the probability of it being misidentified as a positron. This equal probability means that, for many variables, both same-sign (SS) lepton and opposite-sign (OS) lepton tag and probe pairs have the same number of events and topographical features. This means that the shape and normalisation of the SS region can be used to calculate the shape and normalisation of the distributions in the SS+OS region.

The steps in the same sign QCD estimation are as follows:

- Start from tag and probe pairs in data with the same requirements as detailed in the section above
- Require same sign pairs
- Require that probe fails the HEEP ID
- Subtract all MC contributions from the data, with the MC also having the same sign and probe failing HEEP ID requirement. In any cases where the MC is greater than the data, set that bin value to 0. This assumes that any remaining events are due to QCD backgrounds that are not modelled by the MC
- Multiply this result by two to account for both the same sign and the opposite sign regions
- Attribute this result to jet backgrounds from QCD events

This background estimate will affect the efficiency in MC by increasing the number of all probes, therefore reducing the efficiency in regions where these jets are found.

As seen in Figure 8.3 which is plotted for the same-sign failing probes, this QCD construction gives perfect data MC+Jets agreement in all cases where the MC is smaller than the data. Any bins where the MC is greater than the data will not have a perfect agreement. By construction, this method will not give a data/MC ratio greater than 1 in the same sign failing region.

8.5 Results

In this section, I will show the results from 2018 UL data. The same plots from 2016 and 2017 have been analysed but are not shown here, though the final efficiencies and scale factors are quoted. The efficiency plots for both 2016 and 2017 can be found in the appendix.

Due to differences in the detector between the barrel and the endcap, the HEEP ID also has different cuts in these two regions and so all plots of all variables (except η) are also separated according to whether the probe was in the barrel or the endcap.

8.5.1 Kinematic Distributions

Figures 8.4 to 8.8 show the kinematic distribution of the variables of interest in the HEEP ID. The QCD contribution is assumed to be zero in the passing probes region, while it does




Figure 8.3: E_T , ϕ and η of the failing same sign probe regions in the barrel (left) and endcap (right). The QCD contribution is calculated by subtracting the contribution of all MC backgrounds from the data.

have an important contribution in the failing probes region. In Figure 8.4, there is a clear peak around the mass of the Z boson (≈ 91 GeV) in the passing probes region in both the barrel and the endcap. This gives confidence that the electrons selected by the cuts are in fact products of a Z boson decay and are real electrons. This figure also shows some contamination from other standard backgrounds such as W+Jets.

Figure 8.5 shows that the E_T distribution of the selected probes is largely concentrated below 100 GeV as expected from a Z boson decay. Due to the limited number of events at higher energies, the bin width is increased in this region. As there are fewer events in the endcap than in the barrel, the E_T range is reduced to 600 GeV rather than 1 TeV and the binning is also different between the two detector regions. The difference in the number of events in the two regions can be seen in Figure 8.6, with the endcap being the far left and the far right regions and the barrel being the central region. Note that the binning is finer for the central barrel region in this plot.

Due to the symmetry of the detector, the distribution of ϕ is largely flat as expected and shown in Figure 8.7. The reduction in the number of events at some ϕ values is well-modelled in MC and is largely due to additional material before the ECAL in that section of the detector. In the region around $\phi = 1.5$, the reduction in the number of events in the passing probes region is due to dead pixel modules.

Figure 8.8 shows that the bulk of the data and MC distributions, from 10 to 30 nPV have a flat ratio. This does not hold at particularly low or high pileup, but does show that the pileup corrections work as intended in the central region. This effect is partially due to approximations used by the reweighting algorithm. Every bunch has a different pileup, with variations of up to 10%. To avoid saving 2440 numbers per lumisection, an average over all bunches is saved instead. The reweighting also assumes Gaussian tails, but they are non-Gaussian. This does not affect the bulk but it does mean that the tails are not correct. This effect has also been seen with EOY analyses so is not a new feature.

8.5.2 HEED ID Efficiency

The HEED ID efficiencies for the two calculation methods (with and without non-DY subtraction) are shown as functions of E_T , η and ϕ in Figures 8.9, 8.10 and 8.11 respectively. To check the accuracy of the non-DY contributions, they are subtracted from the data and compared to DY. The two methods give scale factors that are in good agreement for all years of data analysed as seen in Table 8.3.

The main sources of systematic uncertainties in the scale factor in this analysis are in the non-DY processes, specifically QCD, $t\bar{t}$ and W+Jets. To account for this, the QCD and W+Jets contributions are assigned an uncertainty of 50% while $t\bar{t}$ is better understood so it is assigned an uncertainty of 10%. For the two MC samples, their cross sections are adjusted by this uncertainty while for QCD the number of events is scaled up or down by 50%. As all three




Figure 8.4: Invariant mass of the tag and probe pair. Shown for the probe in the barrel (left) and the probe in the endcap (right) with all probes (top), passing probes (middle) and failing probes (bottom).





Figure 8.5: E_T of the probe. Shown for the probe in the barrel (left) and the probe in the endcap (right) with all probes (top), passing probes (middle) and failing probes (bottom).

Figure 8.6: η of the probe.

of these processes affect the scale factor in the same direction, all three are increased and then decreased to give the systematic uncertainty quoted. This is a conservative approach as the three systematics are probably not correlated.

8.5.3 N-1 Efficiencies

This section checks the robustness of each cut in the HEEP ID to investigate if there are any selection biases between data and MC for any of the cut variables. This is done by looking at N-1 plots, which remove one of the cuts in the ID. An N-1 efficiency is defined as

$$(8.4) \quad \epsilon_{N-1} = \frac{N_{\text{passingHEEP}}}{N_{\text{passingN-1}}}$$

where $N_{\text{passingHEEP}}$ is the number of events passing the full HEEP ID and $N_{\text{passingN-1}}$ is the number of events that pass all of the cuts in the HEEP ID except for the cut in question.




Figure 8.7: ϕ of the probe. Shown for the probe in the barrel (left) and the probe in the endcap (right) with all probes (top), passing probes (middle) and failing probes (bottom).




Figure 8.8: nPV of the probe. Shown for the probe in the barrel (left) and the probe in the endcap (right) with all probes (top), passing probes (middle) and failing probes (bottom).




Figure 8.9: Efficiency of the HEEP ID as a function of E_T , where the non-DY processes are included (left) and subtracted (right), also split into barrel (top) and endcap (bottom).

All N-1 plots have been studied and none revealed any significant dependence on E_T or any other undesirable characteristics such as poor data/MC agreement. Three N-1 plots are shown here to illustrate their typical features. Figure 8.12 shows the N-1 efficiencies for H/E. This is one of the cuts that has a small energy dependence at low E_T and then becomes flat. Importantly, this is well-modelled in MC. One MC bin at high E_T has an efficiency greater than one, which is simply due to the combination of low numbers of events and the DY sample being generated using aMC@NLO that allows negative weights. This has been checked by removing the weights from the DY sample in which case the efficiency drops to a physical value of below one. As expected, the H/E cut does not show any particular ϕ dependence in the barrel, with small drops in the efficiency in some regions of the endcap due to issues with the HCAL in 2018. There is a small variation in η , which is most prominent at high η and at the change from the barrel to the endcap.

Figure 8.13 shows the N-1 efficiencies for $|d_{xy}|$. In contrast with H/E, the E_T plots here




Figure 8.10: Efficiency of the HEEP ID as a function of η , where the non-DY processes are included (left) and subtracted (right).

	Year	Barrel	Endcap
Data	2016	0.8706 ± 0.0004 (stat)	0.833 ± 0.001 (stat)
MC (DY + non DY)	2016	0.8881 ± 0.0012 (stat)	0.842 ± 0.003 (stat)
Scale Factor	2016	0.985 ± 0.001 (stat) ± 0.004 (sys)	0.990 ± 0.003 (stat) ± 0.007 (sys)
Data - non DY	2016	0.8868 ± 0.0006 (stat)	0.855 ± 0.002 (stat)
MC (DY)	2016	0.9045 ± 0.0011 (stat)	0.864 ± 0.003 (stat)
Scale Factor	2016	0.985 ± 0.002 (stat) ± 0.004 (sys)	0.990 ± 0.003 (stat) ± 0.007 (sys)
Data	2017	0.8668 ± 0.0004 (stat)	0.846 ± 0.001 (stat)
MC (DY + non DY)	2017	0.8887 ± 0.0007 (stat)	0.858 ± 0.002 (stat)
Scale factor	2017	0.979 ± 0.001 (stat) ± 0.005 (sys)	0.987 ± 0.002 (stat) ± 0.010 (sys)
Data - non DY	2017	0.8746 ± 0.0005 (stat)	0.846 ± 0.001 (stat)
MC (DY)	2017	0.8968 ± 0.0007 (stat)	0.879 ± 0.002 (stat)
Scale Factor	2017	0.978 ± 0.001 (stat) ± 0.005 (sys)	0.989 ± 0.002 (stat) ± 0.011 (sys)
Data	2018	0.8546 ± 0.0003 (stat)	0.8354 ± 0.0008 (stat)
MC (DY + non DY)	2018	0.8785 ± 0.0007 (stat)	0.8521 ± 0.002 (stat)
Scale factor	2018	0.973 ± 0.001 (stat) ± 0.004 (sys)	0.980 ± 0.002 (stat) ± 0.011 (sys)
Data - non DY	2018	0.8694 ± 0.0005 (stat)	0.8650 ± 0.001 (stat)
MC (DY)	2018	0.8932 ± 0.0007 (stat)	0.8815 ± 0.002 (stat)
Scale factor	2018	0.973 ± 0.001 (stat) ± 0.004 (sys)	0.981 ± 0.002 (stat) ± 0.010 (sys)

Table 8.3: HEEP ID efficiencies and scale factors for 2016, 2017 and 2018. Calculated for both the barrel and the endcap by comparing data to MC and by subtracting the non-DY MC from data and comparing that to DY MC.




Figure 8.11: Efficiency of the HEEP ID as a function of ϕ , where the non-DY processes are included (left) and subtracted (right), also split into barrel (top) and endcap (bottom).

have no energy dependence and are flat throughout. The efficiency with respect to ϕ is largely flat, but there is a noticeable drop in efficiency around 0.6 due to issues with the tracker. As this is well modelled in the MC this is not of concern. As with H/E, the efficiency with respect to η is flat in the central barrel region, with small reductions in efficiency in the outer barrel region and in the endcaps.

Finally, the N-1 efficiencies for missing hits can be seen in Figure 8.14. In contrast to H/E and $|d_{xy}|$, the missing hit cut has a very small effect on the efficiency of HEEP ID as evidenced by the N-1 plots which show nearly 100% efficiency throughout. Again, this cut is slightly less efficient in the endcap and at the edges of the barrel region.




Figure 8.12: N-1 efficiency of the H/E cut as a function of E_T (top), ϕ (middle) and η (bottom) where the non-DY processes are included in the MC and compared to data.




Figure 8.13: N-1 efficiency of the $|d_{xy}|$ cut as a function of E_T (top), ϕ (middle) and η (bottom) where the non-DY processes are included in the MC and compared to data.




Figure 8.14: N-1 efficiency of the missing hit cut as a function of E_T (top), ϕ (middle) and η (bottom) where the non-DY processes are included in the MC and compared to data.

8.5.4 Comparison of EOY and UL Samples

For this subsection, I have chosen to switch to 2017 data and MC to better illustrate the improvements from EOY samples to UL samples. This is because there was a detector issue that occurred part way through 2017 where some pixels stopped functioning, in turn causing a drop in efficiency in that region of the detector. As this could not be foreseen, the EOY MC samples do not show this feature as they were generated in advance of data collection. Also, while I have seen the EOY plots for 2018, this was only in the form of a presentation and the work has not been fully written up in a dedicated analysis note and therefore cannot be shown here.

The EOY results presented here were the work of B. Clerbaux, W. Fang and R. Goldouzian with full details available at [170]. For conciseness and due to the availability of EOY plots I am allowed to show, only data VS MC is shown here but the same observations also hold for the non-DY subtraction.

As shown in Figure 8.15, there is a noticeable improvement in the UL data MC agreement around $\phi = 1.5$ in the barrel compared to that seen with the EOY samples. This is due to the detector issue mentioned previously which means that the EOY MC does not accurately reflect the conditions in EOY data. Both the EOY and UL data have similar efficiency drops in these affected regions.

Figure 8.16 compares E_T for the two different datasets. In the barrel, the UL data looks to have a more constant efficiency while EOY has a small upwards tendency. There are some differences between EOY and UL MC though it should be noted that there are very limited numbers of events and therefore considerably larger errors associated with this region. EOY and UL MC also use different generators for the DY samples (powheg vs aMC@NLO) which could cause some changes. Both could be considered to match the data within errors. In the endcap, both EOY and UL samples look worse than their barrel counterparts. For UL, the efficiency tends upwards towards high E_T as it does in EOY samples. The EOY results were sufficiently flat with respect to E_T to use in physics analyses, as is UL. The endcap necessarily has to place lower confidence limits than the barrel on the HEEP ID efficiency being the same beyond 1 TeV due to larger statistical errors. It may be possible to tune the parameters of the HEEP ID to make UL endcap samples results flatter like those seen in UL barrel samples, but this is beyond the scope of this study.

8.6 Conclusion

The measurement of the HEEP ID efficiency and scale factor for all Run 2 UL samples has been performed on a per-year basis for a full integrated luminosity of 138 fb^{-1} at $\sqrt{s} = 13 \text{ TeV}$. All results are close to those found previously using EOY data, with small improvements noticeable particularly where the updated MC more accurately reflects conditions seen in data after a




Figure 8.15: Efficiency of the HEEP ID as a function of ϕ in UL samples (left) and EOY samples (right) for both barrel (top) and endcap (bottom) for 2017 data. EOY plots are taken from [170].




Figure 8.16: Efficiency of the HEEP ID as a function of E_T in UL samples (left) and EOY samples (right) for both barrel (top) and endcap (bottom) for 2017 data. EOY plots are taken from [170].

detector issue during data taking, which is particularly well illustrated by the distributions as functions of ϕ in 2017 as shown in Figure 8.15. The results from the two methods (data vs MC and data - non-DY vs DY) give the same scale factors within errors as expected. The efficiency for each year also looks as expected and is flat within errors with respect to E_T as required. As such, these results have been used in the $Z' \rightarrow \ell\ell$ analysis as detailed in the next section.

Chapter 9

Search for Non-Resonant New Physics in High Mass Dilepton Events in Association with b-tagged Jets

This chapter details the search for beyond the Standard Model (BSM) physics in final states with a dilepton and b-tagged jets. This search considers both the dimuon and dielectron channels, with the dielectron channel using the high energy electron/positron identification criteria (HEEP ID) detailed in Chapter 8. The measurements of the efficiency and scale factors for the HEEP ID are my contribution to this analysis, with the rest of the work discussed in this chapter completed by other researchers in the analysis team. This analysis has not yet been unblinded so it is not possible to show final results, with all progress to date detailed in this chapter. As in Chapter 8, in cases where results have not been combined for the whole of Run 2, preference is given to showing preliminary results with 2018 ultra legacy data.

9.1 Theoretical motivation

The search for a final state that produces a dilepton pair in association with b-tagged jets is motivated by the increased sensitivity of the third generation of quarks to BSM physics, with more information given in Subsection 7.2.1. As mentioned in that section, CMS has performed a resonant search for this final state, while the non-resonant search for this final state is presented below.

As with previous CMS non-resonant analyses (see Section 7.3), one of the models in this search is based on the four-fermion contact interaction model, which is detailed in Subsection 2.2.4. Feynman diagrams that produce a dilepton final state in association with b-tagged




Figure 9.1: Feynman diagrams for the production of dileptons in association with zero (left), one (centre) and two (right) b-tagged jets in the final state. From [171].




Figure 9.2: Generated mass distribution for the $bbl\ell$ model at three different values of Λ for constructive interference (left) and for different chirality combinations at $\Lambda = 10$ TeV (right). From [171].

jets are shown in Figure 9.1. As shown in Figure 9.2, the mass distribution in the contact interaction (CI) model depends on both the CI energy scale, Λ , and the different chirality combinations. As this CI model involves the interaction of two b quarks and two leptons, it is referred to as the $bbl\ell$ signal model in this analysis.

In addition to the $bbl\ell$ signal model as given in Equation 2.23, a second model is considered which proposes that the transition $b \rightarrow s\ell\ell$ will be mediated by a Z' boson [172] or leptoquarks [173, 174]. The Feynman diagram for this production mechanism for a dilepton final state is shown in Figure 9.3. This second model, referred to as the $bs\ell\ell$ model, can also be described with a Lagrangian:

$$(9.1) \quad \mathcal{L} = \frac{C_{ij}^{U\ell}}{v^2} (\bar{u}_L^i \gamma_\ell u_L^j) (\bar{u}_L \gamma_\ell u_L) + \frac{C_{ij}^{D\ell}}{v^2} (\bar{d}_L^i \gamma_\ell d_L^j) (\bar{d}_L \gamma_\ell d_L)$$

where v is the vacuum expectation value of the Higgs (see Subsection 2.1.4 for more details), $C_{ij}^{U\ell}$ and $C_{ij}^{D\ell}$ are matrices that define the flavour structure of the interaction and are given in Equation 9.2. The first term in the Lagrangian describes the transition of up-type quarks to a




Figure 9.3: Feynman diagram for the production of dileptons via the transition $b \rightarrow s \ell \ell$. From [171].




Figure 9.4: Generated mass distribution for the $bs\ell\ell$ model at three different values of Λ . From [171].

dilepton pair, while the second term describes the transition of down-type quarks to a dilepton pair.

$$(9.2) \quad C_{ij}^{U\ell} = \begin{pmatrix} C_{u\mu} & 0 & 0 \\ 0 & C_{c\mu} & 0 \\ 0 & 0 & C_{t\mu} \end{pmatrix}, C_{ij}^{D\ell} = \begin{pmatrix} C_{d\mu} & 0 & 0 \\ 0 & C_{s\mu} & C_{bs\mu}^* \\ 0 & C_{bs\mu} & C_{t\mu} \end{pmatrix}$$

In Equation 9.2, the $C_{bs\mu}$ matrix elements define the BSM physics with the energy scale given by $C_{bs\mu} = g_*^2 v^2 / \Lambda^2$, where the coupling g_* is assumed to be 1. The mass distribution for the $bs\ell\ell$ process is dependent on the energy scale, with the MC-generated mass distribution for three energy scales shown in Figure 9.4.

The six final states considered in this analysis are a dimuon pair associated with 0 b-tagged jets, a dimuon pair associated with 1 b-tagged jet, a dimuon pair associated with at least 2

b-tagged jets, a dielectron pair associated with 0 b-tagged jets, a dielectron pair associated with 1 b-tagged jet, and a dielectron pair associated with at least 2 b-tagged jets.

The kinematics of the leading b-tagged jet were studied for both the $bb\ell\ell$ and the $bs\ell\ell$ signal models. For the $bb\ell\ell$ signal model, different energy scales and different chirality combinations were considered. The p_T and η distributions for the $bb\ell\ell$ signal model are shown in Figure 9.5. Only the energy scale was varied for the $bs\ell\ell$ signal model, with p_T and η distributions shown in Figure 9.6. To allow comparison amongst the different energy scales, all distributions are normalised to unity. For the $bb\ell\ell$ signal model, the kinematics of the leading b-tagged jet were found to be consistent across different chiralities at a fixed energy scale. In the $bs\ell\ell$ signal model, the η distribution does not noticeably change for different energy scales as expected, while in the p_T distribution, high energy scales have a tendency to produce more b-tagged jets at high p_T .

9.2 Data, Background MC and Signal MC Samples

As in the analysis described in Chapter 8, this analysis uses Run 2 data collected by the CMS experiment, which corresponds to a total integrated luminosity of 137 fb^{-1} . The full list of background MC samples used is given in Table 9.1. To give sufficient statistics in the high-mass tails, the MC samples used are binned by their mass.

For the $bs\ell\ell$ signal, MC samples are generated in steps of 1 TeV between 1 and 6 TeV. For the $bb\ell\ell$ signal, MC samples are generated in steps of 4 TeV between 6 and 26 TeV. The $bb\ell\ell$ samples are generated for all chirality combinations, so both constructive and destructive interference with the DY process are considered. As with the background MC, the signal MC is also mass binned to ensure sufficient statistics at high mass. The MC generators used for both signal models are MadGraph and Pythia8.

9.3 Baseline Event Selection

The baseline event selection criteria used in the signal region are motivated by the standard Z' selection criteria detailed in Section 7.1. The additional criteria relevant to b-tagged jets were chosen based on the best signal significance.

Specifically, the selection criteria for muons are:

- The two muons must have opposite signs
- Pass a 50 GeV HLT muon trigger
- Have $p_T > 53 \text{ GeV}$ and $|\eta| < 2.4$
- Pass dedicated high p_T muon ID criteria and loose tracker isolation requirements. The muon ID places requirements on the number of valid pixel hits, valid muon hits, number




Figure 9.5: Leading b-tagged jet p_T (left) and η (right) for the $bbl\ell$ signal, shown for different values of Λ (top) and for different chirality combinations at $\Lambda = 10$ TeV (bottom). From [171].

of tracker layers with valid measurement, the number of matched muon stations, the relative error in the muon p_T measurement and the distance from the primary vertex.

- The 3D angle between the two muons must be less than $\pi - 0.02$. This suppresses the cosmic muon background
- Events are rejected if they contain an electron with $p_T > 10$ GeV and $|\eta| < 2.5$ that is within a cone of radius $\Delta R = 0.3$ around a muon.

For electrons, the selection criteria are:

- Pass a double electron trigger, where both electrons have at least 25 GeV for 2018, or in 2016 and 2017 both electrons have at least 33 GeV




Figure 9.6: Leading b-tagged jet p_T (left) and η (right) for the $bs\ell\ell$ signal for three different Λ values.

- Pass the HEEP ID (as shown in Table 8.2, this places a requirement of $E_T > 35$ GeV and $|\eta_{SC}| < 1.4442$ or $1.566 < |\eta_{SC}| < 2.5$)
- Events are rejected that contain a muon with $p_T > 10$ GeV and $|\eta| < 2.5$ that is within a cone of radius $\Delta R = 0.3$ around an electron. This suppresses the WZ background

The choice of a double electron trigger rather than a single electron trigger is motivated by the double electron trigger being highly efficient for the HEEP selection, while a single electron trigger with isolation requirements would be less efficient. The lowest energy trigger without prescaling or an isolation requirement has a $p_T > 100$ GeV requirement. A trigger without isolation requirements also makes the selection of sidebands to investigate fake rates easier.

If an event with more than two leptons passes the above criteria, then the pair with the mass closest to the Z boson mass is chosen if the mass is within 20 GeV of the Z boson mass. This is done to suppress the diboson background at high mass, specifically the WZ and ZZ contributions. For these backgrounds, a high mass is possible by combining an electron from a W or Z with another electron from the Z boson, while in the case of two electrons from the same Z boson, the mass is usually around 90 GeV. Therefore, the mass window requirement suppresses the background in triple electron events by looking for a pair of electrons that are near the Z boson mass and therefore more likely to come from a single Z boson rather than WZ or ZZ. This is similar to the selection described in Section 8.3.2, which placed a requirement on the dielectron mass of $70 < M(ee) < 110$ GeV/ c^2 . If there is no dilepton pair within this window, then the pair with the highest lepton p_T is chosen.

For selections involving b-tagged jets (i.e. not the 0 b-tagged jet category), the following selection criteria are used:

Physics Process	Generator
DY	aMC@NLO
TT2L2Nu	powheg-pythia8
TT fully leptonic	powheg-pythia8
Single top	powheg-pythia8
ST anti	powheg-pythia8
TTWJets to L Nu	aMC@NLO
TTZ to LLNuNu	aMC@NLO
ttHJet to bb	aMC@NLO
WW to 2L 2Nu	powheg-pythia8
WZ to 2L 2Nu	aMC@NLO
WZ to 1L 1Nu 2Q	aMC@NLO
WZ to 2Q 2L	aMC@NLO
ZZ to 2L 2Nu	powheg-pythia8
ZZ to 2Q 2L	powheg-pythia8
ZZ to 4l	powheg-pythia8
Triboson (e.g. WWZ)	aMC@NLO
GluGluH to ZZ to 4L	powheg-pythia8
Vector boson fusion HtoBB Z to LL	powheg-pythia8
ZH HtoBB ZtoLL	powheg-pythia8
ggZH HtoBB ZtoLL	powheg-pythia8
ttH HoZZ 4L	powheg-pythia8
GluGluHToZZTo2L2Q	powheg-pythia8
ZH HtoZZ 4L	powheg-pythia8

Table 9.1: The physics processes and generators for the background MC considered in this analysis.

- Require at least one jet with $p_T > 30$ GeV and $|\eta| < 2.4$
- Pass a set of particle flow general jet identification criteria, which places requirements on variables such as the charged hadron fraction, neutral hadron fraction and neutral EM fraction to remove jets originating from calorimetric noise [175].
- Pass a set of b-tagged jet identification criteria [176]. The criteria used change slightly depending on the mass range and the number of b-tagged jets. This change is based on the criteria with the highest signal significance, $s/\sqrt{B}$. In most cases, the more stringent criteria are used (misidentification probability of 0.1%), while in two mass ranges in the 2+ b-tagged jets category, looser criteria are used (misidentification probability of 1%).

The b-tagged jets are identified using heavy-flavour jet identification algorithms that are 68% efficient at identifying b-tagged jets when using simulated $t\bar{t}$ events, with a 1% probability of misidentifying a light-flavour jet. Identification algorithms exploit the features of b-tagged jets. The lifetime of B hadrons gives rise to a secondary vertex that is displaced from the primary vertex. The distance between the primary and secondary vertex at closest approach

is then used to help distinguish between b-tagged and other jets. The p_T of the jet is also used as the decay products of heavy flavour hadrons usually have a larger p_T than other jets. Additionally, around 20% of b-tagged jets produce an electron or muon, so the presence of one of these particles is also considered in heavy jet identification [177].

As in Chapter 8, the analysis considers two pseudorapidity regions separately. The first is where both leptons are in the barrel, which for electrons places a requirement that $|\eta| < 1.4442$ and for muons places a slightly different requirement that $|\eta| < 1.2$ owing to the different geometries of the two subdetectors. This is known as the barrel-barrel (BB) region. The second region is called the barrel-endcap and endcap-endcap (BE+EE) region, where at least one lepton is in the endcap. In the muon channel, both BE and EE events are considered. In the electron channel only BE pairs are considered as higher noise levels in the endcap and a high rate of misidentifying jets as electrons means that EE pairs are not considered.

9.4 Deep Neural Net Based Discriminator

With the baseline selection criteria detailed above, the Drell-Yan process is the dominant background in the zero b-tagged jet category. In the one b-tagged jet and two b-tagged jet categories, there are significant contributions from the $t\bar{t}$ process. To suppress the $t\bar{t}$ background in the signal region, several kinematic variables have been studied. A deep neural network (DNN) based discriminator has been used to suppress this background contribution.

A DNN is a class of machine learning algorithm that contain multiple hidden layers, as illustrated in Figure 9.7. The individual nodes in one layer receive information from the previous node layer. Each node can apply a functional transformation. The DNN uses the nodes jointly to implement a complex nonlinear mapping between the input and the output [178]. To give the mapping, a technique called error backpropagation is used, which uses the gradient of the loss function to adapt the weights of each node starting at the output and working backwards to the input [179]. DNNs require large amounts of labelled data to train the network.

The DNN used to suppress the $t\bar{t}$ process considers the following kinematic variables:

- Lepton p_T
- Lepton η
- b-tagged jet p_T
- b-tagged jet η
- Number of b-tagged jets
- Missing transverse energy, E_T^{miss}
- Minimum mass of a lepton and a b-tagged jet, $M_{\ell b}$




Figure 9.7: Illustration of a DNN. From [180].

- Minimum 3D angle between the lepton and b-tagged jet

These kinematic variables have different shape distributions for signal events and the $t\bar{t}$ background. The distributions for some of these variables are shown in Figure 9.8, with all input distributions to the DNN showing good data MC agreement. To avoid bias in the network, dilepton mass is not used as a variable.

To train the DNN, single top and $t\bar{t}$ MC samples were used as the background, and signal samples with a mass of at least 200 GeV were used. The DNN is trained by minimising a loss function, with loss functions outputting smaller numbers the closer the algorithm's result is to the correct result. For this DNN, the loss function used is binary cross-entropy. This is a well-established loss function that is used for classification [181, 182]. The binary cross-entropy loss function is defined as

$$(9.3) \quad L((x, y), h) = -y \log h(x) - (1 - y) \log(1 - h(x))$$

where x is the feature, y is the label and h is the probability of x giving y .

To train the network, 80% of the total samples were used, with the remaining 20% of the samples used to validate the DNN. There are four hidden layers in the network, with 128 nodes in one layer, and 64, 32 and 16 nodes in the remaining layers. The model was shown the training data 20 times, at which point the loss function using the validation dataset met the required threshold.

The trained model has been evaluated for all the MC background samples as well as for signal samples at different mass points. The output of the DNN is shown in Figure 9.9. The effect of choosing different DNN scores has also been studied for the signal efficiency, background efficiency and the signal to signal plus background ratio, $S/(S + B)$. This is shown in Figure 9.10.




Figure 9.8: Distributions for the following input variables for 2018 data in the dimuon final state: leading muon p_T (top left), minimum $M_{\ell b}$ (top right), missing transverse energy (bottom left) and minimum 3D angle between the lepton and the b-tagged jet (bottom right). Two mass points of the $bs\ell\ell$ signal model are also shown. From [171].




Figure 9.9: DNN score for the dimuon final state for two different mass points of the $bs\ell\ell$ signal samples. From [171].




Figure 9.10: Signal efficiency and the signal to signal plus background ratio (left) and ROC curve (right) for the DNN for the dimuon final state. The cut on the DNN discriminator at 0.6 is shown in black. From [171].




Figure 9.11: Distributions for minimum $M_{\ell b}$ (left) and the minimum 3D angle between the lepton and the b-tagged jet (right) after the DNN cut (> 0.6) for the dimuon final state in 2018. From [171].




Figure 9.12: Distributions for minimum $M_{\ell b}$ (left) and the minimum 3D angle between the lepton and the b-tagged jet (right) after the DNN cut (> 0.6) for the dielectron final state in 2018. From [171].

To achieve a signal efficiency of 70%, a cut on the DNN discriminator was placed at 0.6. A cut at this discriminator value suppresses the background by 96% and gives a signal to signal plus background ratio of 0.7. Distributions for some variables after this DNN cut are shown in Figure 9.11 for the dimuon channel, and in Figure 9.12 for the dielectron channel. Reasonable data MC agreement is seen after the DNN cut for all variables in all years.

9.5 MC Corrections

To account for well-understood differences between the data and MC, the following corrections are applied to the MC:

- Trigger efficiency weights are applied for both the muon and electron channels
- Scale factors are applied for b-tagging to account for differences in b-tagging efficiency between data and MC
- Pileup reweighting is applied similar to what is described in Chapter 8
- HEEP ID scale factors are applied in the electron channel. Though the scale factors are single values, the uncertainty increases with increasing E_T so the uncertainty used for a given E_T is either the stated uncertainty in the scale factor in Table 8.3, or $1 + (E_T - 90) \times 0.0022$ in the BB region or $1 + (E_T - 90) \times 0.0143$ in the BE region, whichever is larger. These values come from fitting the measured scale factor as a function of E_T with a polynomial of order 0.
- High p_T muon ID scale factors are applied in the muon channel
- Jet energy resolution scale factors and jet energy corrections are applied
- Mis-reconstructed events are removed by placing a cut on E_T^{miss} . This considers events that were effected by events such as dead cells in the ECAL, poor calibration or noisy hits in various subdetectors.
- Events that are affected by an issue in 2018 with the HCAL endcaps are removed if they contain jets. A region of the HCAL had to be turned off part way through the year due to damage to one of the power supplies. It could not be repaired until the detector could be opened at the end of year technical stop.

9.6 Acceptance and Efficiency

To account for physics objects that escape the detector, fail reconstruction or fail the selection conditions, the acceptance and efficiency for both channels and the different categories are measured. This measurement was performed using the DY + Jets MC samples and the $b\bar{b}\ell\ell$ signal samples. The acceptance $\times$ efficiency is shown for the dielectron and dimuon channels in the BB region for one b-tagged jet in Figure 9.13. The dimuon efficiency is slightly higher than the dielectron efficiency due to efficiency losses in pixel matching for the electron trigger. The dielectron channel for 0 and 2+ b-tagged jets is shown in Figure 9.14. The acceptance $\times$ efficiency is consistent for signals with different Λ values for both the dimuon and the dielectron channels.




Figure 9.13: Acceptance $\times$ efficiency for the 1 b-tagged jet category for the dimuon channel (left) and the dielectron channel (right). From [171].




Figure 9.14: Acceptance $\times$ efficiency for the 0 (left) and 2+ (right) b-tagged jet categories for the dielectron channel. From [171].




Figure 9.15: Pileup uncertainty for DY + Jets (left) and for $bb\ell\ell$ signal (right) for the dimuon 0 b-tagged jets category for the BB region. From [171].

9.7 Uncertainties

There are many uncertainties that affect the estimation of both background and signal yields in this analysis. Some uncertainties are common to both channels, while others must be considered independently for electrons and muons.

One uncertainty that affects both channels is uncertainty in the luminosity measurement. As this analysis normalises the MC according to the number of events in a 30 GeV window around the Z peak, this uncertainty cancels out for everything except the signal samples. For the signal, an uncertainty of 1.2% will be applied to the 2016 luminosity measurement, with an uncertainty of 2.3% for 2017 and 2.5% for 2018 [183–185].

Another uncertainty comes from the pileup reweighting, as this method is dependent on the minimum bias cross-section. The minimum bias cross-section value has an uncertainty of 4.6%. To quantify the effect of this uncertainty, the impact on the event yield is estimated by shifting the central value up and down by 4.6%. The effect of this uncertainty can be seen in Figure 9.15.

Uncertainties in the PDFs must also be considered. The uncertainties in the PDFs increase with the fraction of the proton momentum carried by the parton. This uncertainty is split into two categories, with one containing the DY MC and the other containing all non-DY MC, with the PDF uncertainty as a function of dilepton mass shown in Figure 9.16. The non-DY MC has a higher uncertainty than the DY MC due to the inclusion of processes such as gluon fusion where the PDFs are not as well known, particularly at high mass. In the high mass region, the DY uncertainty is around 20%, while for non-DY processes the uncertainty is nearly 50%. Determination of the PDF uncertainties for the signal MC is currently in progress.




Figure 9.16: Uncertainties in the PDFs for the DY (green) and other backgrounds (red). From [171].




Figure 9.17: Uncertainty in b-tagging efficiency for DY + Jets (left) and for $bb\ell\ell$ signal (right) for the dimuon 0 b-tagged jets category for the BB region. From [171].

There is also uncertainty in the efficiency of b-tagging. This is quantified by scaling the b-tagging scale factors up and down by 1σ , where σ is the measured uncertainty of the scale factor. This can be seen for the 1 b-tagged jet category in Figure 9.17. Based on this, a conservative estimate of 10-20% uncertainty is derived.

Uncertainties that come from the jet energy scale and resolution are yet to be quantified.

9.7.1 Electron Uncertainties

For the electron channel, an overall uncertainty of 6% in the BB region, and 8% in the BE+EE region, is assigned to account for uncertainties related to the trigger, electron reconstruction and electron identification efficiency. This uncertainty is the same for all three data taking years. The uncertainty in the transverse energy resolution of the electrons is considered to be negligible. An uncertainty of 1% in the BB region and 2% in the BE+EE region is assigned to the electron transverse energy scale [98]. Again this is the same for all three years of data taking.

An additional systematic uncertainty of 20% is allocated to the MC for 2016 and 2017 to account for an issue in the L1 trigger at that time. This issue is known as prefiring and occurred when a timing offset due to the transparency losses of the lead tungstate crystals was applied offline but not online. This caused an increasing fraction of ECAL trigger primitives to associate an energy deposit with the previous bunch crossing. This online problem was fixed for 2018 [83].

9.7.2 Muon Uncertainties

In the muon channel, the uncertainty in the muon trigger was assigned to be 1% for 2017 and 2018 data. For 2016, the uncertainty was 0.3% in the BB region and 0.7% in the BE+EE region. The uncertainty for muon reconstruction and identification is also quantified. In 2017 and 2018, this uncertainty is given as a 1% for identification and 1% in reconstruction for the BB region. The BE+EE region has some dependence on mass so is not quoted as a single value, the maximum value is 2%. For 2016, both the BB and BE have a mass dependence and at 4 TeV, the uncertainty is 1.5% for the BB region and 6.5% for the BE+EE region. These values were determined by previous analyses looking for massive dilepton pairs with data collected over the same time period [186, 187].

The muon transverse momentum resolution uncertainty was determined to be 10%. This was done by considering dimuon pairs from boosted Z bosons. The transverse muon momentum scale must also be considered, which gives an uncertainty of 1% for the BB region and 3% for the BE+EE region.

Quantification of the impact of the above uncertainties on the mass spectrum is ongoing.

9.8 Lepton Flavour Ratio

The second goal of this analysis is to measure the lepton flavour ratio in all final states for the different b-tagged jet categories. This expands on the previous analysis shown in Section 7.4, where the lepton flavour ratio was measured without placing any requirements on b-tagged jets. The approach detailed below takes the measured mass distributions and corrects them to obtain the underlying “true” distribution in a process known as unfolding. Without unfolding,

a comparison of the muon and electron events would not be possible due to their differing resolution and acceptance.

Unfolding takes experimental data and transforms them according to the following principles:

- Detector response is modelled as a matrix of transition probabilities, known as the response matrix
- Expected background counts are subtracted from experimental data
- The response is inverted and applied to experimental data to give distributions of the underlying distribution
- The unfolded result of the underlying distribution can then be compared directly with theoretical predictions

Bin-by-bin correction factors can have biases and give the wrong normalisation for the unfolded distribution, hence the choice to invert the matrix response [188, 189]. Matrix inversion is very sensitive to variations in the variable under consideration, which can cause large uncertainties and nonphysical values in the unfolded distribution. To quantify this effect, the condition number is calculated during the unfolding process, which is the ratio between the largest and the smallest values in the response matrix. If the condition number is less than 10, then no further corrections are needed.

In this analysis, unfolding was performed using a ROOT package, TUnfold [190], with each year of data unfolded separately. Unfolding was performed for the electron and muon events after subtracting all non-DY backgrounds. The response matrix was quantified using the bin-to-bin migration, which gives the probability that an event with an underlying mass value in one bin is reconstructed into another mass bin. To get the unfolded mass spectra, the response matrix was inverted and applied to the measured mass spectra. The response matrix for the dielectron and dimuon channels for 2018 data in the barrel-barrel region is shown in Figure 9.18. Similar response matrices were also generated for the BE+EE region, as well as for 2016 and 2017 data. The bremsstrahlung effect in muons gives the muon channel a different response matrix to the electron channel. The condition numbers vary from 1.08 in the 2018 BB dielectron channel, to 1.46 in the 2016 BB dimuon channel. As all of these are well below 10, no further corrections are necessary.

A second approach for this analysis is also planned, which will use the profile likelihood approach. In this approach, the ratio is measured as the best-fit value of a profile likelihood ratio test by using the invariant mass distributions. MC is used to estimate the PDFs of different processes.

To confirm that the unfolding has been performed correctly, a closure test is performed. This is done by comparing the unfolded mass spectra to the generated DY MC mass spectra, with the ratio of the two expected to be 1. The closure test for the BB region in the category




Figure 9.18: Response matrix for the 0 b-tagged jets category in the BB region in 2018 for the dimuon channel (left) and the dielectron channel (right). From [171].




Figure 9.19: Closure test for the category with 0 b-tagged jets for the dielectron category, split into the BB region (left) and the BE region (right). From [171].

with no b-tagged jets for both the dielectron and dimuon channels can be seen in Figure 9.19. As expected, the ratio is 1, giving confidence that the closure test has worked as expected. The ratio is also 1 for the BE+EE region in the category with no b-tagged jets for both the dielectron and dimuon channels. The categories with b-tagged jets have not been analysed at this point.

A second test is also performed to check the unfolding procedure, known as a bottom-line test. This test considers the χ^2 values for both the folded and unfolded spectra, with a comparison of the two values to make to ensure that agreement between the data and the model under consideration has not deteriorated after unfolding, which would enhance the ability of

Category	2016		2017		2018	
	$\chi^2_{smeared}$	$\chi^2_{unfolded}$	$\chi^2_{smeared}$	$\chi^2_{unfolded}$	$\chi^2_{smeared}$	$\chi^2_{unfolded}$
Dielectron BB	0.0659918	0.0092014	0.0211259	0.00237607	0.00254675	0.000273692
Dielectron BE	0.0152565	0.00168724	0.0171893	0.0049471	4.38273e-10	4.92358e-13
Dimuon BB	0.47399	0.177559	0.845618	0.56082	0.0151119	0.000984
Dimuon BE+EE	0.493206	0.0787145	0.871936	0.579026	0.286054	0.00627264

Table 9.2: χ^2 values in the smeared and unfolded phase space for both the dielectron and dimuon channels

the unfolded data to reject a model [191]. The two χ^2 values used in the bottom-line test are:

(9.4)

$$\chi^2_{smeared} = (data_{reco} - bkg_{reco} - prediction_{reco})V_{reco}^{-1}(data_{reco} - bkg_{reco} - prediction_{reco})$$

$$\chi^2_{unfolded} = (data_{unfolded} - prediction_{gen})V_{unfolded}^{-1}(data_{unfolded} - prediction_{gen})$$

where V is the covariance matrix.

The results of the bottom-line test can be seen in Table 9.2. The reason for the χ^2 values being so small is not currently understood and is currently under investigation. As expected, all values of $\chi^2_{smeared}$ and $\chi^2_{unfolded}$ are of the same order, so the unfolded data is not more discriminative than the data recorded by the CMS detector. The χ^2 values are the smallest for 2018 which is expected due to it being the year with the fewest detector issues. Additionally, there was a tracker upgrade at the end of 2017.

The unfolded distribution is also corrected using the acceptance and efficiency of electrons and muons, as well as by the application of a scale factor estimated by fitting a polynomial of degree 0 to the low mass region of 200-400 GeV. To obtain the flavour ratio for the underlying event, the unfolded invariant mass spectra are scaled to the inverse of the ratio from the unfolded DY MC. The flavour ratio for the 0 b-tagged jets category can be seen in Figure 9.20. The lepton flavour ratio is in agreement with unity for this 0 b-tagged jet category, which suggests lepton flavour universality.

As the data has not yet been unblinded it is not possible to show results for the 1 or 2+ b-tagged jets channels yet. The predicted lepton flavour ratio for the 1 b-tagged jet channel is shown in Figure 9.21. This shows that for some potential values of Λ , there will be a deviation from unity at high mass. At low mass, lepton flavour universality will be conserved regardless of the value of Λ .

9.9 Results

At this point in the analysis, only results for the $bb\ell\ell$ signal model are mature enough to share. Results for the $bs\ell\ell$ signal model are a work in progress. For the $bb\ell\ell$ process, limits will be set individually for each of the eight possible chiralities. To do this, both the background and the signal are normalised to the Z boson peak to absorb any constant scaling terms.




Figure 9.20: Flavour ratio for the combined BB and BE region for the 0 b-tagged jets category using the full Run 2 data. From [171].




Figure 9.21: Predicted lepton flavour ratio for the 1 b-tagged jet category in the BB region for different values of Λ . Data is blinded. From [192].

Limits are set using a multi-bin counting experiment, with the probability of observing a given number of events, N_{obs} given by

$$(9.5) \quad P(N_{obs}) = \frac{a^{N_{obs}}}{N_{obs}!} \cdot e^{-a}$$

where a is the assumed mean and the distribution is assumed to be Poisson. The background contributions are accounted for using

$$(9.6) \quad a = \mathcal{L} \cdot \prod_i \epsilon_s \cdot v_{i,s} \cdot \sigma_s + \prod_j \epsilon_b \cdot v_{j,b} \cdot \sigma_b \frac{N_{obs,Z-peak}}{\sigma_{b,Z-peak}}$$

where s denotes signal, b denotes background, $Z - peak$ denotes a variable in the normalisation mass window, ϵ is the efficiency, σ is the cross-section, and v are nuisance parameters.

Limits are set using a Bayesian approach with confidence limits set calculated from the posterior distribution given by

$$(9.7) \quad \Pi_{post}(\sigma|obs) = \int d\epsilon_s d\epsilon_b d\sigma_b \prod_{i,j,k} dv_{i,s} dv_{j,b} \frac{a_k^{N_{obs,k}}}{N_{obs,k}!} \cdot e^{-a_k} \cdot \pi(\epsilon_s, \epsilon_b, v_{i,s}, v_{j,b}, \sigma_b) \cdot \pi_{poi}(\sigma_s)$$

where k indicates the product over the $\cos(\theta^*)$ bins and range of masses, while $\pi(\epsilon_s, \epsilon_b, v_{i,s}, v_{j,b}, \sigma_b)$ refers to the combined function for nuisance parameters. The π_{poi} term is a prior function for the parameter of interest and is a flat distribution.

The 95% confidence limit on the cross section upper limit can be obtained by requiring

$$(9.8) \quad 0.95 \stackrel{!}{=} \int_0^{\sigma_{lim}} \Pi_{post}(\sigma_s|obs) d\sigma_s$$

The counting experiment model was implemented using the Higgs Combine Tool and limits were calculated using this tool also.

The expected upper limit for the left-left model with constructive interference in the case of zero b-tagged jets is shown in Figure 9.22. As expected for this case, there is only sensitivity to this model for low values of Λ . The 1 and 2+ b-tagged jet cases are expected to have increased sensitivity. Expected limits for all categories will be calculated soon.

9.10 Conclusion

The search for non-resonant new physics in high mass dilepton events in association with b-tagged jets is in progress and approval for unblinding is expected soon once the analysis has been finalised. The channels involving b-tagged jets are expected to be more sensitive to new physics than the channels without b-tagged jets due to the higher masses of the third generation of quarks. This analysis is expected to probe higher Λ ranges than the previous CMS analysis also using Run 2 data, with the expectation that this will either provide hints of BSM physics or will allow more stringent limits to be set. The lepton flavour ratio will also




Figure 9.22: Expected upper limit on the cross-section of the $bbl\bar{l}$ model in the left-left model for the dimuon channel in the BB region for the 0 b-tagged jet category. The dashed line shows the expected limit, with the green shaded area denoting the 68% quantile and the yellow area denoting the 95% quantile. The red line shows the theoretical prediction for the constructive LL $bbl\bar{l}$ signal model. From [171].

be probed, with certain values of Λ expected to show deviations from unity at high mass in the case of dilepton production in association with b-tagged jets. The lepton flavour ratio for the case of dilepton production with no b-tagged jets is in agreement with unity as expected. Without the HEEP ID, it would not have been possible to study high mass dielectron events, which would have limited this physics search to the search for non-resonant high mass dimuon events, which would provide reduced sensitivity compared to the combined dilepton channel. Additionally, it would not be possible to study the lepton flavour ratio.

Chapter 10

Summary

The work in this thesis details studies undertaken in the development of Phase II upgrades to the CMS ECAL. To maintain the detector's performance, both the on- and off-detector electronics will be replaced. Available prototype electronics have been tested in a test beam environment using an electron beam to analyse the response to EM signals. The tests found that the resolution for a 3×3 crystal matrix was consistent with expectations for Phase II. In test beam conditions without a magnetic field, a 3×3 crystal matrix was found to provide a better energy resolution than a 3×5 crystal matrix or a single crystal. An energy correction method was also tested for electrons that impact on the side of a crystal, with this method found to correct the energy resolution to close to that for electrons impacting on the centre of the crystal.

I investigated the performance of three trigger algorithms for the ECAL Phase II upgrade. The three algorithms are the clustering algorithm, which performs basic clustering of localised energy, the Swiss-cross algorithm, which rejects spikes based on the energy deposits in the surrounding crystals, and the linear discriminant (LD) algorithm, which rejects spikes based on their pulse shape. These three algorithms have been written using high-level synthesis (HLS) and tested against C++ counterparts. The clustering algorithm was tested on over 50,000 events from an electron test beam with no discrepancies between the two algorithms found. The Swiss-cross algorithm was tested using 2800 events from a pion test beam and no discrepancies were found. The LD algorithm was tested, again using events from a pion test beam, with discrepancies between the two found in approximately 1 in 2650 samples. Both spike-killing algorithms work at the required 99.9% efficiency from 2 GeV in the case of the LD algorithm and from 10 GeV for the Swiss-cross algorithm.

The Swiss-cross and clustering algorithms were combined and I tested the combined algorithm on hardware, finding no discrepancies in the 1000 tested events from electron test beam data. The current sorting algorithm in this combined algorithm is too resource intensive and the plan is to replace it with a different sorting algorithm in the future. The LD algorithm was also tested on hardware and gave the expected results for over 800 events.

I analysed a high energy electron/positron (HEEP) ID using 137 fb^{-1} of data collected between 2016 and 2018 at $\sqrt{s} = 13\text{ TeV}$ by the CMS experiment to aid in the search for physics beyond the SM. The HEEP ID places cuts on well-understood detector-based variables and is designed to be robust to changes in detector conditions, including changes in pileup. This data has been recalibrated to improve the reconstruction and identification of particles, giving both a better resolution and improved data/MC agreement. The data was analysed on a per-year basis using a tag and probe method, with non-DY subtraction used to crosscheck results. The efficiency and scale factor for the ID was measured, with results compared to the previous results from the dataset with the old calibrations applied. The robustness of each cut in the HEEP ID was also checked. In 2018, the data efficiency of the HEEP ID was $86.94 \pm 0.05\%$, with an MC efficiency of $89.32 \pm 0.07\%$ and a scale factor of $0.973 \pm 0.001\text{ (stat)} \pm 0.004\text{ (sys)}$. This result, and the results for the other years, are close to those found previously using EOY data, with small improvements noticeable particularly with better modelling of a detector issue.

The search for physics beyond the SM using that the HEEP ID has focused on new physics with a final state of dileptons produced in association with b-tagged jets. The dilepton states considered are dimuons or dielectrons, with ditau production not considered due to the limited tau identification efficiency when minimising the probability of jets, electrons and muons being misidentified as taus. The requirement for b-tagged jets is motivated by the increased sensitivity of the third generation of quarks to BSM physics due to their higher masses.

Two signal models are considered during this analysis. The first signal model is based on the four-fermion contact interaction model, where four fermions are considered to interact at a single point rather than via a mediator such as a W boson. This model predicts a compositeness scale, Λ , and also depends on helicity current which can cause constructive or destructive interference. Accordingly, signal models for different Λ values and different helicity currents are considered. The second signal model considers the transition $b \rightarrow s\ell\ell$ mediated by a Z' boson or scalar leptoquarks. In this model, different Λ values must be considered.

Events are selected based on passing an appropriate lepton trigger, meeting p_T and η requirements and passing the HEEP ID or a dedicated high energy muon ID. A deep neural net based discriminator is used to suppress the $t\bar{t}$ background in the final state categories that include b-tagged jets.

Lepton flavour universality is also considered in this analysis. The flavour ratio for the 0 b-tagged jets category is in agreement with unity, which suggests lepton flavour universality as expected. The predicted lepton flavour ratio for the 1 b-tagged jet channel shows that for some potential energy scales, there will be a deviation from unity at high mass.

The analysis is in progress and approval for unblinding is expected once the analysis has been finalised, and is expected to probe higher energy scale ranges than the previous CMS analysis also using Run 2 data. The results are expected to either provide hints of BSM physics or allow more stringent limits to be set.

Appendix A

HEEP Efficiencies for 2016 and 2017

This appendix contains the efficiency plots that were not shown in Chapter 8. Figure A.1 shows the HEEP ID efficiency as a function of η for 2017 data compared to 2017 MC, while the efficiency as a function of E_T was shown in Figure 8.16 and as a function of ϕ in Figure 8.15. All efficiency plots considered for 2016 data and MC are shown in Figure A.2.

The drop in efficiency around $\eta = 0$ in Figure A.2 is due to imperfect modelling of a detector issue in 2016, where high energy deposits in the silicon strip sensors caused significant downtime due to saturation of one of the front-end chips. The MC sample for this year is split into two, before and after the detector issue was resolved, with the MC sample after this issue was resolved showing the same ratio between the data and MC across all regions of the detector. While the MC for the detector before the fix was applied does have a drop in efficiency in this $\eta \approx 0$ region, it is not as significant as the reduction in efficiency observed in data. The N-1 efficiencies were also checked and it was found that the ECAL driven and H/E cuts cause a drop in efficiency in this region for both data and MC. The $|\Delta\eta_{in}^{seed}|$ cut causes a drop in both the MC and the data, but the data efficiency decrease is larger than that of the MC. There is also a small contribution from the EM+HD1 cut, but the primary cause is the $|\Delta\eta_{in}^{seed}|$ cut which is consistent with the known detector issue.




Figure A.1: Efficiency of the HEEP ID for 2017 as a function of η , where the MC includes the DY process.




Figure A.2: Efficiency of the HEEP ID for 2016 as a function of E_T (top), ϕ (middle) and η (bottom), where the MC includes the DY process. For the E_T and ϕ plots, this is further split into the BB category (left) and BE category (right).

Bibliography

- [1] ATLAS Collaboration, “Observation of a new particle in the search for the standard model higgs boson with the ATLAS detector at the LHC,” *Physics Letters B*, vol. 716, no. 1, pp. 1–29, 2012.
- [2] CMS Collaboration, “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC,” *Physics Letters B*, vol. 716, no. 1, pp. 30–61, 2012.
- [3] M. Thomson, *Modern particle physics*. Cambridge, United Kingdom: Cambridge University Press, 2013.
- [4] J. Goldstein, “Theoretical particle physics 408 course notes,” 2019. Lecture course at the University of Bristol.
- [5] Particle Data Group, “Review of Particle Physics,” *Progress of Theoretical and Experimental Physics*, vol. 2022, no. 8, p. 083C01, 2022.
- [6] A. Djouadi, “Higgs phenomenology: a short review,” 1996.
- [7] M. Gell-Mann, “The Eightfold Way: A Theory of strong interaction symmetry,” 3 1961.
- [8] T. D. Lee and C.-N. Yang, “Question of Parity Conservation in Weak Interactions,” *Phys. Rev.*, vol. 104, pp. 254–258, 1956.
- [9] C. S. Wu *et al.*, “Experimental Test of Parity Conservation in β Decay,” *Phys. Rev.*, vol. 105, pp. 1413–1414, 1957.
- [10] J. W. Cronin, “The discovery of CP violation,” *Eur. Phys. J. H*, vol. 36, pp. 487–508, 2012.
- [11] J. H. Christenson *et al.*, “Evidence for the 2π Decay of the K_2^0 Meson,” *Phys. Rev. Lett.*, vol. 13, pp. 138–140, 1964.
- [12] S. L. Glashow, “Partial-symmetries of weak interactions,” *Nuclear Physics*, vol. 22, no. 4, pp. 579–588, 1961.
- [13] S. Weinberg, “A Model of Leptons,” *Phys. Rev. Lett.*, vol. 19, pp. 1264–1266, 1967.

- [14] A. Salam, “Weak and Electromagnetic Interactions,” *Conf. Proc. C*, vol. 680519, pp. 367–377, 1968.
- [15] P. W. Higgs, “Broken symmetries, massless particles and gauge fields,” *Physics Letters*, vol. 12, no. 2, pp. 132–133, 1964.
- [16] P. W. Higgs, “Broken symmetries and the masses of gauge bosons,” *Phys. Rev. Lett.*, vol. 13, pp. 508–509, 1964.
- [17] F. Englert and R. Brout, “Broken symmetry and the mass of gauge vector mesons,” *Phys. Rev. Lett.*, vol. 13, pp. 321–323, 1964.
- [18] I. Melo, “Higgs potential and fundamental physics,” *European Journal of Physics*, vol. 38, no. 6, p. 065404, 2017.
- [19] Planck Collaboration, “Planck 2018 results. vi. cosmological parameters,” *Astronomy and Astrophysics*, vol. 641, p. A6, 2020.
- [20] CMS Collaboration, “Search for supersymmetry in proton-proton collisions at 13 TeV in final states with jets and missing transverse momentum,” *Journal of High Energy Physics*, vol. 2019, no. 10, 2019.
- [21] G. Altarelli, B. Melé, and M. Ruiz-Altaba, “Searching for new heavy vector bosons in $p\bar{p}$ colliders,” *Z. Phys. C*, vol. 45, pp. 109–121, 1989.
- [22] H. Georgi and S. L. Glashow, “Unity of all elementary-particle forces,” *Phys. Rev. Lett.*, vol. 32, pp. 438–441, 1974.
- [23] E. Eichten, I. Hinchliffe, K. Lane, and C. Quigg, “Supercollider physics,” *Rev. Mod. Phys.*, vol. 56, pp. 579–707, 1984.
- [24] I. A. D’Souza and C. S. Kalman, *Preons: Models of leptons, quarks and gauge bosons as composite objects*.
World Scientific, 1992.
- [25] N. Arkani-Hamed, S. Dimopoulos, and G. Dvali, “Phenomenology, astrophysics, and cosmology of theories with submillimeter dimensions and TeV scale quantum gravity,” *Phys. Rev. D*, vol. 59, p. 086004, 1999.
- [26] N. Arkani-Hamed, S. Dimopoulos, and G. Dvali, “The hierarchy problem and new dimensions at a millimeter,” *Physics Letters B*, vol. 429, no. 3-4, pp. 263–272, 1998.
- [27] L. Randall and R. Sundrum, “Large mass hierarchy from a small extra dimension,” *Phys. Rev. Lett.*, vol. 83, pp. 3370–3373, 1999.

- [28] J. L. Hewett, “Indirect collider signals for extra dimensions,” *Physical Review Letters*, vol. 82, no. 24, pp. 4765–4768, 1999.
- [29] G. F. Giudice, R. Rattazzi, and J. D. Wells, “Quantum gravity and extra dimensions at high-energy colliders,” *Nuclear Physics B*, vol. 544, no. 1-2, pp. 3–38, 1999.
- [30] T. Han, J. D. Lykken, and R.-J. Zhang, “Kaluza-klein states from large extra dimensions,” *Physical Review D*, vol. 59, no. 10, p. 105006, 1999.
- [31] J. M. Carmona and J. L. Cortés, “Infrared and ultraviolet cutoffs of quantum field theory,” *Physical Review D*, vol. 65, no. 2, 2001.
- [32] J. Kretzschmar, “Searches for extra dimensions with the ATLAS and CMS detectors,” tech. rep., CERN, Geneva, 2016.
- [33] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt, “Parton distributions for the LHC,” *The European Physical Journal C*, vol. 63, no. 2, pp. 189–285, 2009.
- [34] I. R. Kenyon, “The drell-yan process,” *Reports on Progress in Physics*, vol. 45, no. 11, p. 1261–1315, 1982.
- [35] CMS Collaboration, “Measurements of differential z boson production cross sections in proton-proton collisions at $\sqrt{s} = 13$ TeV,” *Journal of High Energy Physics*, vol. 2019, no. 12, 2019.
- [36] S. Höche, “Introduction to parton-shower event generators,” 2015. Lectures presented at TASI 2014. <https://doi.org/10.48550/arXiv.1411.4085>.
- [37] J. Alwall *et al.*, “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations,” *Journal of High Energy Physics*, vol. 07, no. 79, 2014.
- [38] S. Frixione, G. Ridolfi, and P. Nason, “A positive-weight next-to-leading-order monte carlo for heavy flavour hadroproduction,” *Journal of High Energy Physics*, vol. 2007, no. 09, pp. 126–126, 2007.
- [39] S. Alioli, P. Nason, C. Oleari, and E. Re, “A general framework for implementing NLO calculations in shower monte carlo programs: the POWHEG BOX,” *Journal of High Energy Physics*, vol. 2010, no. 6, 2010.
- [40] P. Nason, “A new method for combining NLO QCD with shower monte carlo algorithms,” *Journal of High Energy Physics*, vol. 2004, no. 11, pp. 040–040, 2004.
- [41] T. Sjöstrand, S. Mrenna, and P. Skands, “PYTHIA 6.4 physics and manual,” *Journal of High Energy Physics*, vol. 2006, no. 05, pp. 026–026, 2006.

- [42] T. Sjöstrand *et al.*, “An introduction to PYTHIA 8.2,” *Computer Physics Communications*, vol. 191, pp. 159–177, 2015.
- [43] S. A. et al, “Geant4—a simulation toolkit,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, no. 3, pp. 250–303, 2003.
- [44] O. S. Brüning *et al.*, *LHC Design Report*.
CERN Yellow Reports: Monographs, Geneva: CERN, 2004.
- [45] E. Mobs, “The CERN accelerator complex in 2019. Complexe des accélérateurs du CERN en 2019,” 2019.
CERN-GRAPHICS-2019-002-1.
- [46] L. Rossi, “The LHC superconducting magnets,” *Conf. Proc. C*, vol. 030512, p. 141, 2003.
- [47] CMS Collaboration, “Cms luminosity - public results.”
Available at <https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults>.
- [48] T. Sakuma and T. McCauley, “Detector and event visualization with sketchup at the cms experiment,” *Journal of Physics: Conference Series*, vol. 513, no. 2, p. 022032, 2014.
- [49] CMS Collaboration, “CMS: letter of intent by the CMS Collaboration for a general purpose detector at LHC,” tech. rep., CERN, Geneva, 1992.
Open presentation to the LHCC 5 November 1992, M. Della Negra/CERN, CMS Spokesman.
- [50] I. Neutelings, “Cms coordinate system.” https://tikz.net/axis3d_cms/, 2021.
Accessed: 10/03/2023.
- [51] A. Hervé, “The CMS detector magnet,” *IEEE Trans. Appl. Supercond.*, vol. 10, no. 1, pp. 389–94, 2000.
- [52] L. Hoddeson, A. W. Kolb, and C. Westfall, *Fermilab*.
Chicago: University of Chicago Press, 2008.
- [53] A. Hervé *et al.*, “Status of the construction of the CMS magnet,” *IEEE Trans. Appl. Supercond.*, vol. 14, no. 2, pp. 542–547, 2004.
- [54] CMS Collaboration, “Cms tracker detector performance results.”
Avaiable at <https://twiki.cern.ch/twiki/bin/view/CMSPublic/DPGResultsTRK>.
- [55] W. Adam *et al.*, “The cms phase-1 pixel detector upgrade,” *Journal of Instrumentation*, vol. 16, no. 02, p. P02027, 2021.

- [56] P. Azzurri, “The CMS Silicon Strip Tracker,” *J. Phys.: Conf. Ser.*, vol. 41, pp. 127–134, 2006.
- [57] V. Veszpremi, “Operation and performance of the cms tracker,” *Journal of Instrumentation*, vol. 9, no. 03, p. C03005, 2014.
- [58] C. Biino, “The CMS Electromagnetic Calorimeter: overview, lessons learned during Run 1 and future projections,” *Journal of Physics: Conference Series*, vol. 587, no. 1, p. 012001, 2015.
- [59] CMS Collaboration, *CMS Physics: Technical Design Report Volume 1: Detector Performance and Software*. Technical design report. CMS, Geneva: CERN, 2006.
- [60] CMS Collaboration, “The CMS hadron calorimeter project: Technical Design Report,” tech. rep., Geneva, 1997.
- [61] J. Mans *et al.*, “CMS Technical Design Report for the Phase 1 Upgrade of the Hadron Calorimeter,” tech. rep., 2012.
- [62] C. Isik, “Phase 1 upgrade of the cms hadron barrel calorimeter,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1042, p. 167389, 2022.
- [63] F. Beaudette, “The CMS Particle Flow Algorithm,” *International Conference on Calorimetry for the High Energy Frontier Proceedings*, 2013.
- [64] CMS Collaboration, “Performance of the CMS Muon Detectors in 2016 collision runs,” 2016.
- [65] H. Sert, “CMS Run 2 High Level Trigger Performance,” tech. rep., CERN, Geneva, 2020.
- [66] A. Bocci and on behalf of the CMS Collaboration, “Cms high level trigger performance comparison on cpus and gpus,” *Journal of Physics: Conference Series*, vol. 2438, no. 1, p. 012016, 2023.
- [67] O. Aberle *et al.*, “High-Luminosity Large Hadron Collider (HL-LHC): Technical design report,” tech. rep., Geneva, 2020.
- [68] The HL-LHC Collaboration, “Ls3 schedule change.” <https://hilumilhc.web.cern.ch/article/ls3-schedule-change>, 2022. Accessed: 13/04/2023.
- [69] CMS Collaboration, “The Phase-2 Upgrade of the CMS Level-1 Trigger,” tech. rep., CERN, Geneva, 2020.

- [70] CMS Collaboration, “The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger,” tech. rep., CERN, Geneva, 2021.
- [71] CMS Collaboration, “The Phase-2 Upgrade of the CMS Tracker,” tech. rep., CERN, Geneva, 2017.
- [72] CMS Collaboration, “A MIP Timing Detector for the CMS Phase-2 Upgrade,” tech. rep., CERN, Geneva, 2019.
- [73] CMS Collaboration, “The Phase-2 Upgrade of the CMS Barrel Calorimeters,” tech. rep., CERN, Geneva, 2017.
- [74] CMS Collaboration, “The Phase-2 Upgrade of the CMS Muon Detectors,” tech. rep., CERN, Geneva, 2017.
- [75] CMS Collaboration, “Particle-flow reconstruction and global event description with the CMS detector,” *Journal of Instrumentation*, vol. 12, no. 10, p. P10003, 2017.
- [76] CMS Collaboration, “Performance of the cms electromagnetic calorimeter in pp collisions at $\sqrt{s} = 13$ TeV,” *Paper in preparation*, vol. EGM-18-002.
- [77] D. Valsecchi for the CMS Collaboration, “Deep learning techniques for energy clustering in the CMS ECAL,” *Journal of Physics: Conference Series*, vol. 2438, no. 1, p. 012077, 2023.
- [78] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Transactions of the ASME—Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.
- [79] R. Frühwirth, “A Gaussian-mixture approximation of the Bethe–Heitler model of electron energy loss by bremsstrahlung,” *Computer Physics Communications*, vol. 154, no. 2, pp. 131–142, 2003.
- [80] W. Adam, R. Frühwirth, A. Strandlie, and T. Todor, “Reconstruction of Electrons with the Gaussian-Sum Filter in the CMS Tracker at the LHC,” tech. rep., CERN, Geneva, 2005.
- [81] CMS Collaboration, “Performance of photon reconstruction and identification with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV,” *Journal of Instrumentation*, vol. 10, no. 08, pp. P08010–P08010, 2015.
- [82] M. Dordevic, “The CMS Particle Flow Algorithm,” *EPJ Web of Conferences*, vol. 191, p. 02016, 2018.

-
- [83] CMS Collaboration, “Performance of the CMS level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV,” *Journal of Instrumentation*, vol. 15, no. 10, pp. P10017–P10017, 2020.
- [84] CMS Collaboration, “Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC,” *CMS Paper*, no. EGM-17-001, 2020.
- [85] D. A. Petyt for the CMS Collaboration, “Mitigation of Anomalous APD signals in the CMS Electromagnetic Calorimeter,” *Journal of Physics: Conference Series*, vol. 404, no. 1, p. 012043, 2012.
- [86] CMS Collaboration, “CMS ECAL trigger plots,” 2019.
Available at <https://cds.cern.ch/record/2690933>.
- [87] CMS Collaboration, “Electromagnetic calorimeter commissioning and first results with 7 TeV data,” tech. rep., CERN, Geneva, 2010.
- [88] CMS Collaboration, “Reconstruction of signal amplitudes in the CMS electromagnetic calorimeter in the presence of overlapping proton-proton interactions,” *Journal of Instrumentation*, vol. 15, no. 10, pp. P10002–P10002, 2020.
- [89] A. Zghiche, “Optimizing the performance of the CMS Electromagnetic Calorimeter to measure Higgs properties during Phase I and Phase II of the LHC,” *International Journal of Modern Physics A*, vol. 35, no. 34n35, p. 2044011, 2020.
Publisher: World Scientific Publishing Co.
- [90] R. Tramontano, “The ultimate CMS ECAL calibration and performance for the legacy reprocessing of LHC Run 2 data,” in *Proceedings of 41st International Conference on High Energy physics — PoS(ICHEP2022)*, (Bologna, Italy), p. 643, Sissa Medialab, 2023.
- [91] CMS Collaboration, “CMS ECAL Response to Laser Light,” 2019.
Available at <http://cds.cern.ch/record/2668200>.
- [92] P. Bonamy *et al.*, “The ecal calibration: Use of the light monitoring system,” *CMS Note 1998/013*, 1998.
- [93] T. T. d. Fatis, “Calibration of the electromagnetic calorimeter of the CMS experiment,” *Journal of Physics: Conference Series*, vol. 160, no. 1, p. 012051, 2009.
- [94] M. Cipriani on behalf of the CMS Collaboration, “Performance of the CMS electromagnetic calorimeter during the LHC Run II and its role in precision physics measurements,” in *Proceedings of Sixth Annual Conference on Large Hadron Collider Physics — PoS(LHCP2018)*, vol. 321, p. 084, 2018.

- [95] CMS Collaboration, “ECAL 2016 refined calibration and Run2 summary plots,” *CMS Detector Performance Summary*, no. CMS DP 2020 021, 2020.
- [96] CMS Collaboration, “Energy calibration and resolution of the CMS electromagnetic calorimeter in pp collisions at $\sqrt{s} = 7$ TeV,” *Journal of Instrumentation*, vol. 8, no. 09, pp. P09009–P09009, 2013.
- [97] A. Benaglia, “The CMS ECAL performance with examples,” *Journal of Instrumentation*, vol. 9, no. 02, pp. C02008–C02008, 2014.
- [98] CMS Collaboration, “Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC,” *Journal of Instrumentation*, vol. 16, no. 05, p. P05014, 2021.
- [99] C. Cooke on behalf of the CMS Collaboration, “Upgrade of the CMS Barrel Electromagnetic Calorimeter for the High Luminosity LHC,” *Instruments*, vol. 6, no. 3, p. 29, 2022.
- [100] F. Winterstein, S. Bayliss, and G. A. Constantinides, “High-level synthesis of dynamic data structures: A case study using Vivado HLS,” in *2013 International Conference on Field-Programmable Technology (FPT)*, pp. 362–365, 2013.
- [101] J. Hakala, N. Loukas, L. Lutton, and N. Marinelli, “Amplitude and time reconstruction for phase-2 ecal barrel trigger primitives: studies and hls implementation,” *CMS Detector Note*, vol. CMS DN-2020/012, 2020.
- [102] D. A. Petyt and T. Reis, “CMS ECAL barrel Phase-2 APD spike tagger firmware implementation studies using Vivado HLS,” *CMS Detector Note*, no. CMS DN 2020/011, 2020.
- [103] CMS Collaboration, “Performance of electron reconstruction and selection with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV,” *Journal of Instrumentation*, vol. 10, no. 06, pp. P06005–P06005, 2015.
- [104] CMS Collaboration, “The Phase-2 Upgrade of the CMS Endcap Calorimeter,” tech. rep., CERN, Geneva, 2017.
- [105] C. Cooke, H. F. Heath, and D. A. Petyt, “ECAL test beam resolution analysis with prototype phase ii readout,” *CMS Detector Note*, no. CMS DN 2020/036, 2020.
- [106] F. Micheli, D. A. Petyt, S. Pigazzini, and T. Reis, “APD spike studies with prototype ECAL barrel Phase-2 electronics using 2018 ECAL test beam data,” *CMS Detector Note*, vol. CMS DN-2019/018, 2019.

- [107] M. Oreglia, “A Study of the Reactions $\psi' \rightarrow \gamma\gamma\psi$,” Master’s thesis, 1980. Stanford University.
- [108] N. Marinelli, “Ecal barrel trigger primitives from clusters.” Internal CMS ECAL Upgrade Meeting, 2018.
- [109] J. Descamps and P. Jarry, “Periodic position dependence of the energy measured in the CMS electromagnetic calorimeter,” tech. rep., CERN, Geneva, 2006.
- [110] P. Adzic, “Energy resolution of the barrel of the CMS Electromagnetic Calorimeter,” *Journal of Instrumentation*, vol. 2, no. 04, p. P04004, 2007.
- [111] R. C. Fernow, *Introduction to Experimental Particle Physics*. Cambridge University Press, 1986.
- [112] M. Shahdad, “An overview of vhdl language and technology,” in *23rd ACM/IEEE Design Automation Conference*, pp. 320–326, 1986.
- [113] N. P. Ghanathe *et al.*, “Software and firmware co-development using high-level synthesis,” *Journal of Instrumentation*, vol. 12, no. 01, p. C01083, 2017.
- [114] M. Husejko, J. Evans, and J. C. Rasteiro Da Silva, “Investigation of High-Level Synthesis tools’ applicability to data acquisition systems design based on the CMS ECAL Data Concentrator Card example,” *Journal of Physics: Conference Series*, vol. 664, no. 8, p. 082019, 2015.
- [115] Xilinx, *Vivado Design Suite User Guide: High-Level Synthesis*. 2021.
- [116] S. Bologna, “Introduction to HLS,” 2019.
- [117] K. Compton, S. Dasu, A. Farmahini-Farahani, S. Fayer, R. Fobes, R. Frazier, T. Gorski, G. Hall, G. Iles, J. Jones, P. Klabbers, D. Newbold, A. Perugupalli, S. Rader, A. Rose, D. Seemuth, W. Smith, and J. Tikalsky, “The mp7 and ctp-6: multi-hundred gbps processing boards for calorimeter trigger upgrades at cms,” *Journal of Instrumentation*, vol. 7, no. 12, p. C12024, 2012.
- [118] D. Prasad, M. Y. M. Yusof, S. S. Palai, and A. H. Nawi, “Sorting networks on fpga,” in *Proceedings of the WSEAS International Conference on Telecommunications and Informatics*, pp. 29–31, 2011.
- [119] R. Chen, S. Siriyal, and V. Prasanna, “Energy and memory efficient mapping of bitonic sorting on fpga,” in *Proceedings of the 2015 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pp. 240–249, 2015.

- [120] K. Wang *et al.*, “Optimized sorting network for successive cancellation list decoding of polar codes,” *IEICE Electronics Express*, vol. 14, no. 18, pp. 20170735–20170735, 2017.
- [121] L. Lutton and L. Zygala, “Energy resolution for 160 mhz.” Internal CMS ECAL Upgrade Meeting, 2018.
- [122] A. Tumasyan *et al.*, “Search for new physics in the lepton plus missing transverse momentum final state in proton-proton collisions at $\sqrt{s} = 13$ TeV,” *JHEP*, vol. 07, p. 067, 2022.
- [123] CMS Collaboration, “Search for resonances in the dilepton mass distribution in pp collisions at $\sqrt{s} = 7$ TeV,” *Journal of High Energy Physics*, vol. 2011, no. 5, 2011.
- [124] CMS Collaboration, “Search for narrow resonances in dilepton mass spectra in pp collisions at $\sqrt{s} = 7$ TeV,” *Physics Letters B*, vol. 714, no. 2, pp. 158–179, 2012.
- [125] CMS Collaboration, “Search for heavy narrow dilepton resonances in pp collisions at $\sqrt{s} = 7$ TeV and $\sqrt{s} = 8$ TeV,” *Physics Letters B*, vol. 720, no. 1-3, pp. 63–82, 2013.
- [126] CMS Collaboration, “Search for physics beyond the standard model in dilepton mass spectra in proton-proton collisions at $\sqrt{s} = 8$ TeV,” *Journal of High Energy Physics*, vol. 2015, no. 4, 2015.
- [127] CMS Collaboration, “Search for narrow resonances in dilepton mass spectra in proton–proton collisions at $\sqrt{s}=13$ TeV and combination with 8 TeV data,” *Physics Letters B*, vol. 768, pp. 57–80, 2017.
- [128] CMS Collaboration, “Search for high-mass resonances in dilepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV,” *Journal of High Energy Physics*, vol. 2018, no. 6, 2018.
- [129] CMS Collaboration, “Search for resonant and nonresonant new phenomena in high-mass dilepton final states at $\sqrt{s} = 13$ TeV,” *Journal of High Energy Physics*, vol. 2021, no. 7, 2021.
- [130] ATLAS Collaboration, “Search for high mass dilepton resonances in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS experiment,” *Physics Letters B*, vol. 700, no. 3-4, pp. 163–180, 2011.
- [131] ATLAS Collaboration, “Search for high-mass resonances decaying to dilepton final states in pp collisions at $\sqrt{7}$ TeV with the ATLAS detector,” *Journal of High Energy Physics*, vol. 2012, no. 11, 2012.

-
- [132] ATLAS Collaboration, “Search for high-mass dilepton resonances in pp collisions at $\sqrt{s} = 8$ TeV with the atlas detector,” *Phys. Rev. D*, vol. 90, p. 052005, 2014.
- [133] ATLAS Collaboration, “Search for new high-mass phenomena in the dilepton final state using 36 fb^{-1} of proton-proton collision data at $\sqrt{13}$ TeV with the ATLAS detector,” *Journal of High Energy Physics*, vol. 2017, no. 10, 2017.
- [134] V. Barger, N. G. Deshpande, J. L. Rosner, and K. Whisnant, “Production, decays, and forward-backward asymmetries of extra gauge bosons in e_6 ,” *Phys. Rev. D*, vol. 35, pp. 2893–2896, 1987.
- [135] CDF Collaboration, “Search for new physics in high-mass electron-positron events in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV,” *Phys. Rev. Lett.*, vol. 99, p. 171802, 2007.
- [136] CMS Collaboration, “Performance of reconstruction and identification of τ leptons decaying to hadrons and ν_τ at $\sqrt{s} = 13$ TeV,” *Journal of Instrumentation*, vol. 13, no. 10, pp. P10005–P10005, 2018.
- [137] G. Weiglein *et al.*, “Physics interplay of the LHC and the ILC,” *Physics Reports*, vol. 426, no. 2-6, pp. 47–358, 2006.
- [138] T. Han, G. Valencia, and Y. Wang, “Hadron collider signatures for new interactions of top and bottom quarks,” *Phys. Rev. D*, vol. 70, p. 034002, 2004.
- [139] CDF Collaboration, “Search for high-mass resonances decaying to $e\mu$ in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV,” *Phys. Rev. Lett.*, vol. 96, p. 211802, 2006.
- [140] CMS Collaboration, “Performance of the reconstruction and identification of high-momentum muons in proton-proton collisions at $\sqrt{s} = 13$ TeV,” *Journal of Instrumentation*, vol. 15, no. 02, p. P02027, 2020.
- [141] V. H. S. Kumar and P. K. Suresh, “Gravitons in kaluza-klein theory,” 2006.
Available at <https://arxiv.org/abs/gr-qc/0605016>.
- [142] D. Alves *et al.*, “Simplified models for lhc new physics searches,” *Journal of Physics G: Nuclear and Particle Physics*, vol. 39, no. 10, p. 105005, 2012.
- [143] A. Albert *et al.*, “Recommendations of the LHC dark matter working group: Comparing LHC searches for heavy mediators of dark matter production in visible and invisible decay channels,” 2017.
Available at <https://arxiv.org/abs/1703.05703>.
- [144] E. Accomando, A. Belyaev, J. Fiaschi, K. Mimasu, S. Moretti, and C. Shepherd-Themistocleous, “Forward-backward asymmetry as a discovery tool for z' bosons at the LHC,” *Journal of High Energy Physics*, vol. 2016, no. 1, 2016.

- [145] “Dark matter summary plots for s-channel mediators,” tech. rep., CERN, Geneva, 2020.
- [146] P. A. R. Ade *et al.*, “Planck 2015 results. xiii. cosmological parameters,” *Astronomy and Astrophysics*, vol. 594, p. A13, 2016.
- [147] J. Albrecht, D. van Dyk, and C. Langenbruch, “Flavour anomalies in heavy quark decays,” *Progress in Particle and Nuclear Physics*, vol. 120, p. 103885, 2021.
- [148] LHCb Collaboration, “Measurement of lepton universality parameters in $B^+ \rightarrow K^+ \ell^+ \ell^-$ and $B^0 \rightarrow K^{*0} \ell^+ \ell^-$ decays,” tech. rep., CERN, Geneva, 2022.
- [149] CMS Collaboration, “Search for a high mass dimuon resonance associated with b quark jets at $\sqrt{s} = 13$ TeV,” 2022.
EXO-22-016, to be published.
- [150] ATLAS Collaboration, “Search for heavy resonances decaying into a w or z boson and a higgs boson in final states with leptons and b-jets in 36 fb^{-1} of $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector,” *Journal of High Energy Physics*, vol. 2018, no. 3, 2018.
- [151] J. C. Collins and D. E. Soper, “Angular distribution of dileptons in high-energy hadron collisions,” *Phys. Rev. D*, vol. 16, pp. 2219–2225, 1977.
- [152] “Search for leptonic decays of W’ bosons in pp collisions at $\sqrt{s} = 8$ TeV,” tech. rep., CERN, Geneva, 2013.
- [153] CMS Collaboration, “Search for contact interactions and large extra dimensions in the dilepton mass spectra from proton-proton collisions at $\sqrt{s} = 13$ TeV,” *Journal of High Energy Physics*, vol. 2019, no. 4, 2019.
- [154] ATLAS Collaboration, “Search for contact interactions and large extra dimensions in the dilepton channel using proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector,” *The European Physical Journal C*, vol. 74, no. 12, 2014.
- [155] ATLAS Collaboration, “Search for new non-resonant phenomena in high-mass dilepton final states with the ATLAS detector,” *JHEP*, vol. 2020, no. 5, 2020.
- [156] CMS Collaboration, “Search for physics beyond the standard model in high-mass diphoton events from proton-proton collisions at $\sqrt{s} = 13$ TeV,” *Phys. Rev. D*, vol. 98, p. 092001, 2018.
- [157] CMS Collaboration, “Search for new physics in dijet angular distributions using proton–proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter and other models,” *The European Physical Journal C*, vol. 78, no. 9, 2018.

-
- [158] ATLAS Collaboration, “Search for new phenomena in dijet events using 37 fb^{-1} of pp collision data collected at $\sqrt{s} = 13\text{ TeV}$ with the atlas detector,” *Phys. Rev. D*, vol. 96, p. 052004, 2017.
- [159] E. Adelberger, J. Gundlach, B. Heckel, S. Hoedl, and S. Schlamminger, “Torsion balance experiments: A low-energy frontier of particle physics,” *Progress in Particle and Nuclear Physics*, vol. 62, no. 1, pp. 102–134, 2009.
- [160] S. Hannestad and G. G. Raffelt, “Supernova and neutron-star limits on large extra dimensions reexamined,” *Phys. Rev. D*, vol. 67, p. 125008, 2003.
- [161] ATLAS Collaboration, “Search for new phenomena in final states with two leptons and one or no b -tagged jets at $\sqrt{s} = 13\text{ TeV}$ using the atlas detector,” *Phys. Rev. Lett.*, vol. 127, p. 141801, 2021.
- [162] D. Bečirević, N. Košnik, O. Sumensari, and R. Z. Funchal, “Palatable leptoquark scenarios for lepton flavor violation in exclusive $b \rightarrow s\ell_1\ell_2$ modes,” *J. High Energ. Phys.*, no. 35, 2016.
- [163] P. Langacker, “The physics of heavy z' gauge bosons,” *Reviews of Modern Physics*, vol. 81, no. 3, pp. 1199–1228, 2009.
- [164] A. Greljo and D. Marzocca, “High- p_t dilepton tails and flavor physics,” *The European Physical Journal C*, vol. 77, no. 8, 2017.
- [165] CMS Collaboration, “Search for heavy gauge w' bosons in events with an energetic lepton and large missing transverse momentum at $\sqrt{s} = 13\text{ TeV}$,” *Physics Letters B*, vol. 770, pp. 278–301, 2017.
- [166] E. Bols, J. Kieseler, M. Verzetti, M. Stoye, and A. Stakia, “Jet Flavour Classification Using DeepJet,” *Journal of Instrumentation*, vol. 15, no. 12, pp. P12012–P12012, 2020.
- [167] CMS Collaboration, “Identification of b -quark jets with the CMS experiment,” *Journal of Instrumentation*, vol. 8, no. 04, pp. P04013–P04013, 2013.
- [168] CMS Collaboration, “B-tagging performance of the cms legacy dataset 2018,” *CMS Performance Note*, vol. CMS DP-2021/004, 2021.
- [169] A. Giraldi, “Precision luminosity measurement with proton-proton collisions at the cms experiment in Run 2,” 2022.
Available at <https://arxiv.org/abs/2208.08214>.

- [170] B. Clerbaux, W. Fang, and R. Goldouzian, “Heep selection efficiency scale factor study for cms with the full 2017 data sample,” *CMS Detector Note*, vol. CMS AN-2018/143, 2018.
- [171] S. Jeon *et al.*, “Search for nonresonant new physics in high-mass dilepton events in association with b-tagged jets,” *CMS Analysis Note*, vol. CMS AN-2022/058, 2022.
- [172] Y. Afik, J. Cohen, E. Gozani, E. Kajomovitz, and Y. Rozen, “Establishing a search for $b \rightarrow s\ell\ell$ anomalies at the LHC,” *Journal of High Energy Physics*, vol. 2018, no. 8, 2018.
- [173] S. Patnaik and R. Singh, “A light shed on lepton flavor universality in b decays,” *Universe*, vol. 9, no. 3, 2023.
- [174] B. Gripaios, M. Nardecchia, and S. A. Renner, “Composite leptoquarks and anomalies in b -meson decays,” 2015.
- [175] “Jet algorithms performance in 13 TeV data,” tech. rep., CERN, Geneva, 2017.
- [176] E. Bols, J. Kieseler, M. Verzetti, M. Stoye, and A. Stakia, “Jet flavour classification using deepjet,” *Journal of Instrumentation*, vol. 15, no. 12, p. P12012, 2020.
- [177] CMS Collaboration, “Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV,” *Journal of Instrumentation*, vol. 13, no. 05, pp. P05011–P05011, 2018.
- [178] G. Montavon, W. Samek, and K.-R. Müller, “Methods for interpreting and understanding deep neural networks,” *Digital Signal Processing*, vol. 73, pp. 1–15, 2018.
- [179] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning representations by back-propagating errors,” *Nature*, vol. 323, no. 6088, pp. 533–536, 1986.
- [180] T. X. Vu *et al.*, “Machine learning-enabled joint antenna selection and precoding design: From offline complexity to online performance,” *IEEE Transactions on Wireless Communications*, vol. 20, no. 6, pp. 3710–3722, 2021.
- [181] G. Kasieczka, S. Marzani, G. Soyez, and G. Stagnitto, “Towards machine learning analytics for jet substructure,” *Journal of High Energy Physics*, vol. 2020, no. 9, 2020.
- [182] P. Keicher, “Machine learning in top physics in the atlas and cms collaborations,” 2023. Available at <https://arxiv.org/abs/2301.09534>.
- [183] CMS Collaboration, “CMS Luminosity Measurements for the 2016 Data Taking Period,” tech. rep., CERN, Geneva, 2017.
- [184] CMS Collaboration, “CMS luminosity measurement for the 2017 data-taking period at $\sqrt{s} = 13$ TeV,” tech. rep., CERN, Geneva, 2018.

- [185] CMS Collaboration, “CMS luminosity measurement for the 2018 data-taking period at $\sqrt{s} = 13$ TeV,” tech. rep., CERN, Geneva, 2019.
- [186] G. Abbiendi *et al.*, “Search for high-mass resonances decaying to muon pairs in pp collisions at $\sqrt{s} = 13$ TeV with the full 2016 data set of 37 fb and combination with 2015 result,” *CMS Analysis Note*, no. CMS AN 2016/391, 2016.
- [187] B. Clerbaux *et al.*, “Search for high-mass resonances in the di-electron final state with 2017 data,” *CMS Analysis Note*, no. CMS AN 2018/021, 2016.
- [188] P. Vischia, “Intensive course on statistics for hep,” 2022.
Lecture course at CP3 UCLouvain.
- [189] CERN, *Proceedings of the PHYSTAT 2011 Workshop on Statistical Issues Related to Discovery Claims in Search Experiments and Unfolding: CERN, Geneva, Switzerland 17 - 20 Jan 2011. PHYSTAT 2011 Workshop on Statistical Issues Related to Discovery Claims in Search Experiments and Unfolding*, (Geneva), CERN, 2011.
- [190] S. Schmitt, “TUnfold, an algorithm for correcting migration effects in high energy physics,” *Journal of Instrumentation*, vol. 7, no. 10, pp. T10003–T10003, 2012.
- [191] R. D. Cousins, S. J. May, and Y. Sun, “Should unfolded histograms be used to test hypotheses?,” 2016.
Available at <https://arxiv.org/pdf/1607.07038.pdf>.
- [192] A. Kaur, “Search for nonresonant new physics in high-mass dilepton events in association with b-tagged jets.” Internal CMS Weekly General Meeting, 2023.

