

A SEARCH IN THE MUON CHANNEL FOR
HEAVY RESONANCES DECAYING TO
LONG-LIVED NEUTRAL PARTICLES

ADAM PAUL HUNT

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

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

MAY 2016

© Copyright by Adam Paul Hunt, 2016.

All rights reserved.

Abstract

A search is performed for a heavy resonance decaying to two long-lived massive neutral particles that each decay to a pair of muons. The process is detected experimentally via a distinct topological signature consisting of a pair of oppositely charged leptons originating at a vertex significantly displaced from the LHC beam spot. Events were collected by the CMS detector at the LHC during pp collisions at $\sqrt{s} = 7$ TeV, and selected from data samples corresponding to 5.1 fb^{-1} of integrated luminosity in the muon channel. No significant excess is observed above standard model expectations, and an upper limit is set with 95% confidence level on the production cross section times the branching fraction to dimuons, as a function of the long-lived massive neutral particle lifetime.

Acknowledgements

First and foremost, I thank my original adviser, Valerie Halyo, for this incredible opportunity. Professor Halyo provided me with many great projects and a great forum. She encouraged me to take the spot light more often than I would have liked, and in the end, it paid off. She gave me the support and direction I needed even though, at times, I resisted. Her endless energy, curiosity and enthusiasm pushed me to achieve more than I thought was possible.

Next, I must thank my final adviser, Chris Tully, for not giving up. Thank you for helping me perform this final push to complete my dissertation. I am overwhelmed by your generosity. It was a pleasure working with you. Your insights were always spot on. Assisting you in the graduate electronics course was a highlight of my time at Princeton. I will never forget one of our tea time discussions during my first year. I was thoroughly confused by all the acronyms associated with CMS. When I asked if there was a glossary of all the acronyms, you promptly responded, "Oh, you can find those in the TDR", yet another acronym.

In addition to Valerie, the rest of the HF Luminosity Monitoring group at Princeton who performed the initial commissioning and calibration work deserve a special thanks: Dan Marlow, John Jones, Nadia Adam, and Jeremy Werner. It truly is amazing what we accomplished with such a small group in such a short time. I've learned so much from each of you.

Thank you to everyone in the Displaced Fermion group: E. Clement, K. Harder, I. Tomalin, V. Halyo, P. Jindal, T. Lou, P. Lujan, J. Werner, A. Zuranski, M. Jones, N. Leonardo, M. de Mattia, I. Shipsey, D. Silvers, Y. Zheng, A. Gay, C. Boulahouache. This analysis has been full of unexpected challenges. Thank you all for having the tenacity to keep it moving forward.

I need to thank all my friends CERN and FNAL. Without their companionship and support, I surely would have ended my graduate education prematurely. Although

we have never met and are not likely to, I have to thank George Chan, the author of Ph.D. Comics (<http://www.phdcomics.com>), for reminding me that the struggles I faced are all too common.

It has been an honor to work on the Compact Muon Solenoid. I will cherish this experience for the rest of my life. Thank you to all the scientists, engineers and technicians that made it possible. To be part of the commissioning effort is an incredible honor.

Lastly, I thank my wife, Kirsten. The long distance, late nights, never ending work, canceled plans and extreme pressure put an unbelievable strain on our relationship. Thank you for sticking by me and putting up with me. I owe you everything.

The work described in this dissertation is published in References [1, 2, 3].

To Kirsten, Madeline and Charles

Contents

Abstract	iii
Acknowledgements	iv
List of Tables	x
List of Figures	xiv
1 Introduction	1
2 Displaced Dilepton Signatures	3
2.1 Standard Model	3
2.2 Beyond the Standard Model	8
2.3 Experimental Search	11
3 The LHC and the CMS Detector	13
3.1 The Large Hadron Collider (LHC)	13
3.1.1 Accelerator Complex	15
3.2 The Compact Muon Solenoid - CMS	17
4 The CMS Online Luminosity System	32
4.1 Introduction	32
4.2 HF Lumi Calculation	33
4.3 HF and HLX operation	34
4.4 Luminosity Data Acquisition Software	36

4.4.1	Continuous Running	38
4.5	Luminosity Data Distributors	40
4.5.1	Luminosity TCP Repeater	40
4.5.2	DIP Distributor	41
4.5.3	Luminosity File Merger	41
4.5.4	SCAL Distributor	41
4.5.5	Lumi Server Monitoring	41
4.6	HLX Supervisor	42
4.7	Absolute Calibration using Van Der Meer Scans	44
4.7.1	Separation Scan Measurements	46
4.7.2	Results	49
4.7.3	Systematic Error	51
4.8	Lumi Operational Considerations	52
5	Data and Monte-Carlo Samples	56
6	Event Selection and Performance	61
6.1	Displaced track reconstruction	61
6.2	Displaced lepton identification	63
6.3	Selection of long-lived exotica	63
6.3.1	Selection efficiency	70
7	Background Estimation and Modeling	78
7.1	Background Normalization	78
7.2	Background shape	81
7.3	Statistical uncertainty on apriori background from counting method	81
8	Systematic Uncertainties	88
8.1	Luminosity	89

8.2	Effect of pileup	89
8.3	PDF, renormalisation and factorisation scale uncertainties	89
8.4	Track-finding efficiency	92
8.4.1	Tracking efficiency from cosmic muons	92
8.4.2	Tracking efficiency using embedding	100
8.4.3	Verification of displaced tracking efficiency using K_s^0	107
8.5	Trigger efficiency measurement	107
8.6	Transverse decay length correction	117
8.7	Effect of higher-order QCD corrections	119
8.8	Background uncertainty	120
9	Results	121
9.1	Upper limits	121
9.2	Limits on similar models with long-lived exotica	127
9.3	Summary	129

List of Tables

2.1	Standard Model Particle Properties - Isospin, hypercharge, electric charge	6
4.1	Results from the 2010 CMS calibration scans using the Van der Meer method. For the final calibration of the HF measurement, the results from the two scans are averaged.	50
4.2	Comparison of beam widths between online and offline methods. “Forward” and “backward” refer to the scan direction—i.e., whether the coordinate being scanned increases or decreases with time. The starred widths (**) come from a scan that was missing many data points for the offline methods. These numbers should not be considered accurate. All values are in mm.	51
4.3	Systematic errors on the 2010 CMS calibration scan measurements using the Van der Meer method.	52
5.1	Simulated signal samples used in the analysis. The masses of the H^0 and X bosons are given, as is the mean proper decay length of the X boson.	58
5.2	Simulated background samples used in the analysis. The weights shown are equal to the cross section divided by the number of simulated events used.	59

5.3	Simulated QCD background samples used in the analysis. The quoted cross sections have been scaled by the efficiency of the event filter in the case of the muon-enriched QCD samples. For the search, the normal QCD and Mu-enriched QCD samples are used. The weights shown are equal to the cross section divided by the number of simulated events used.	60
6.1	Table of selection cut efficiencies for candidates in the muon channel. Each cut efficiency is derived using candidates that pass all previous cuts. The efficiency of the first cut is determined relative a sample of events where an analysis trigger fired, and two tracks each with $p_T > 33$ GeV/c were present. The signal efficiency is illustrated for the case $M_{H^0} = 1000$ GeV/c ² , $M_X = 350$ GeV/c ²	74
6.2	The efficiency of $X \rightarrow \ell^+ \ell^-$ decays to pass all selection criteria, both for $H^0 \rightarrow XX$ events in which only one long-lived exotic decays to pairs of muons ϵ_1 and for the case where both decay to pairs of muon ϵ_2 . The uncertainties on these efficiencies are not shown in the table. They are dominated by the 20% relative uncertainty related to the tracking performance.	77
7.1	Comparison of various methods of estimating the background normalisation. The number of data events is also included for reference. . . .	80
7.2	Positive and negative one standard deviation statistical uncertainties on the apriori background estimated by counting simulated background events passing the normal selection cuts. The uncertainties are derived making use of the given ‘loose sample’ of candidates. The numbers shown in brackets indicate the estimated uncertainty if one removes the candidate with the largest weight.	87

8.1	Systematic uncertainties affecting the signal efficiency over the range of M_{H^0} and M_X values considered. In all cases, the uncertainty specified is a relative uncertainty. Note that the NLO uncertainty is only evaluated for the $M_{H^0} = 125 \text{ GeV}/c^2$ case. The relative uncertainty in the luminosity is taken to be 2.2%.	88
8.2	The fraction of X bosons, whose dimuon decay products both pass the kinematic cuts on the muon p_T and pseudorapidity. The uncertainties quoted on these numbers correspond to the systematic uncertainty on the PDF set.	91
8.3	Acceptance and efficiency for reconstructing muons with $p_T = 25 \text{ GeV}/c$ as a function of transverse impact parameter.	105
8.4	Acceptance and efficiency for reconstructing muons with $p_T = 25 \text{ GeV}/c$ as a function of the projected longitudinal impact parameter $d_{sz} = z_0 \sin \theta$	105
8.5	Efficiency for reconstructing muons with $p_T = 25 \text{ GeV}/c$ as a function of the number of reconstructed primary vertices.	106
8.6	Efficiency of the probe muon to fire a single leg of the displaced muon trigger paths, as estimated using the ‘Tag and Probe’ method, as applied to both data and simulated events. Value ranges correspond to estimates performed in different run periods. The values shown have a statistical uncertainty smaller than 1%.	113
9.1	This is identical to Table 8.2 except that it shows the results for $Z' \rightarrow XX'$ instead of $H^0 \rightarrow XX$, where X and X' are non-identical spin 0 particles of equal mass. The uncertainties quoted on the acceptances again correspond to the systematic uncertainty on the PDF set. . . .	129

9.2	This is identical to Table 8.2 except that it shows the results for $Z' \rightarrow Q\bar{Q}$, where Q is a spin-half particle, instead of $H^0 \rightarrow XX$. The uncertainties quoted on the acceptances again correspond to the systematic uncertainty on the PDF set.	130
-----	---	-----

List of Figures

2.1	Standard Model Particles	3
2.2	Standard Model Particle Interactions	4
2.3	Hidden Valley Potential [4]	9
3.1	The Large Hadron Collider beam orbit crossing configuration for the four interaction regions at ALICE, ATLAS, CMS and LHC-b.	14
3.2	LHC Dipole Cross-section. The Dipoles are placed in the curved regions of the accelerator tunnel and are used to steer the beams. This cross section shows the “two-in-one” design of super conducting magnets which was selected as a space saving measure due to the size of the LEP tunnel.	15
3.3	The LHC injector complex (not to scale) showing the path from Linac2 to the PSB, PSB to the PS, PS to SPS through transfer lines, and from the SPS to the LHC.	16
3.4	CMS Integrated Luminosity - Proton-Proton Collisions.	17
3.5	Daily peak luminosity at P5 (CMS) - 2011.	18
3.6	Compact Muon Solenoid	18
3.7	CMS Inner Tracker - single lines indicate single sided detectors and double lines indicate double sided detectors. From [5]	20
3.8	Electromagnetic Calorimeter - Side View.	22
3.9	Electromagnetic Calorimeter - Stereo View.	22

3.10	HCAL Longitudinal Cross Section.	24
3.11	Hadronic Calorimeter Barrel Module.	25
3.12	Hadronic Calorimeter End-Cap.	25
3.13	Hadronic Calorimeter Outer (HO) located in the gap between the magnet and the muon chambers.	26
3.14	Muon Detectors Longitudinal Cross Section	27
3.15	Number of nuclear interaction lengths in front of the muon stations as a function of pseudo-rapidity.	28
3.16	Schematic of the Level-1 Trigger System.	30
4.1	Luminosity Data Path	35
4.2	Luminosity DAQ	37
4.3	Diagram of continuous running operation. The first run starts on an OC0 and luminosity data is accumulated in 4096-orbit nibbles (purple). At some point a main CMS run is started, at which point another OC0 is sent and any partially-acumulated nibble is dropped (red). The data stream continues with a nibble aligned with the new run.	39
4.4	Luminosity Data Distribution	40
4.5	Luminosity Monitoring System Function Manager	43
4.6	VdM X scan	47
4.7	VdM Y scan	47
4.8	VdM X scan	48
4.9	VdM Y scan	48
4.10	Fill 1089 Zero Points	49
4.11	QIE Pedestals prior to LUT subtraction	53
4.12	HLX ET Sum Histogram showing QIE Reset effect	54

4.13	Afterglow from LED runs normalized to the peak height. E_T sum method... The vertical error bars indicate the total spread. The peaks after the main peaks are created by ion feedback.	54
5.1	An example of a simulated Monte Carlo event event with a H^0 mass of $400 \text{ GeV}/c^2$ and X boson mass of $150 \text{ GeV}/c^2$	57
6.1	Track reconstruction efficiency for single, isolated muons of $p_T = 50 \text{ GeV}/c$ as a function of the transverse impact parameter. The efficiency has little dependence on p_T	62
6.2	Distribution of $\sum p_T$ in an isolation cone around muons from simulated long-lived particle decay. The selection requires this to be less than $4 \text{ GeV}/c$	66
6.3	The mean $\sum p_T$ in an isolation cone around muons, as a function of the number of pileup interactions. The error bars show the RMS spread in $\sum p_T$ at each point. Tracks from pileup interactions increases the average momentum within the isolation cone.	67
6.4	Comparison of isolation cone momentum sum in data and simulation background events, for dimuon candidates. The candidates shown in these plots pass all selection criteria except isolation and lifetime-related cuts.	68
6.5	The efficiency to select $X \rightarrow \mu^+ \mu^-$ decay as a function of transverse decay length for dimuon candidates, shown for the case $M_{H^0} = 1000 \text{ GeV}/c^2$, $M_X = 150 \text{ GeV}/c^2$. The left (right) plot shows the efficiencies for events in which one (two) X bosons decay to dimuons. . .	69
6.6	The dimuon vertex χ^2/NDF distribution. This is required to be less than 5. The open histogram represents an example long-lived signal with a product of cross-section and branching fraction of 1 pb	70

6.7	The φ difference between the dimuon momentum vector and the vector between the primary and secondary vertices. This is required to be less than 0.2.	71
6.8	The number of tracker hits assigned in total, to the two muon candidates, which lie in front of (i.e. closer to the centre of CMS) the reconstructed vertex position. It is required to be no more than 1. . .	71
6.9	The cosine of the angle between the two muons that form a candidate. It must exceed -0.95 , so rejecting back-to-back tracks. This particular plot shows the distribution after all other cuts have been implied; in this case, no further events are rejected by this cut.	72
6.10	The transverse decay length significance of the candidates. It is required to be more than 5 standard deviations for dimuon candidates.	72
6.11	The minimum transverse impact parameter significance of the two leptons in the candidates for the dimuon channels. It is required to be more than 2 standard deviations.	73
6.12	The reconstructed dimuon mass after all selection cuts have been applied. The background in the dimuon channel is extremely small. . .	73
6.13	The reconstructed transverse decay length of the candidates in the dimuon channel, after all selection cuts have been applied.	74
6.14	The invariant mass distribution of dimuon candidates after applying all selection cuts except for the lifetime-related cuts. This predominantly selects prompt background such as Z bosons. The agreement of both shape and normalization between data and Monte Carlo simulation demonstrates a good understanding of the Standard Model backgrounds.	75

7.1	The transverse decay length significance of the candidates for the dimuon channel with loosened cuts in data and simulation. The vertical dashed line indicates the selection requirement used for signal events. There are no simulated QCD or $t\bar{t}$ events passing these selection requirements, so they are omitted.	80
7.2	Distribution of the dilepton mass and the fitted shape in a data sample with lifetime-related selection requirements removed, shown for the muon channel. The shape used is that of a Breit–Wigner distribution times a turn-on function, added to an exponential term.	82
7.3	Fit of the background shape to simulated background with lifetime-related selection requirements removed, shown for the muon channels. The function used is the same as that for Figure 7.2.	83
7.4	Alternative fits of the background to a sample with loosened versions of the signal selection requirements, for simulated events (left) and data (right), for the muon channel.	83
8.1	The number of reconstructed primary vertices per event in data and simulation, following the pileup reweighing procedure for the muons sample. The grey band in the simulation histograms shows the variation obtained following the procedure to estimate pileup systematics.	90
8.2	Cosmic muons reconstructed in the muon chambers alone. The p_T (a), η (b) and ϕ (c) distributions are shown. Data is in black, simulation in red.	94
8.3	Cosmic muons reconstructed in the muon chambers alone. The number of valid hits (a) and χ^2/ndof (b) distributions are shown. Data is in black, simulation in red.	94
8.4	Cosmic muons reconstructed in the muon chambers alone. The $ d_0 $ (a) and $ d_z $ (b) distributions are shown. Data is in black, simulation in red.	95

8.5	Distributions of the uncertainties on the transverse (a) and longitudinal (b) impact parameters for selected cosmic muons reconstructed in the muon chambers alone. The curves show the cuts applied on these uncertainties.	96
8.6	Selected cosmic muons reconstructed in the muon chambers alone. The p_T (a), η (b) and ϕ (c) distributions are shown. Data is in black, simulation in red.	96
8.7	Selected cosmic muons reconstructed in the muon chambers alone. The number of valid hits (a) and χ^2/ndof (b) distributions are shown. Data is in black, simulation in red.	97
8.8	Selected cosmic muons reconstructed in the muon chambers alone. The $ d_0 $ (a) and $ z_0 $ (b) distributions are shown. Data is in black, simulation in red.	97
8.9	ΔR between selected cosmic muons reconstructed in the muon chambers and the closest tracker track. Data is black, simulation red. . . .	98
8.10	Tracker tracks p_T (a), η (b) and ϕ (c) distributions. Data is in black, simulation in red.	99
8.11	Tracker tracks number of valid hits (a) and χ^2/ndof (b) distributions. Data is in black, simulation in red.	99
8.12	Tracker tracks $ d_0 $ (a) and $ d_z $ (b) distributions. Data is in black, simulation in red.	100
8.13	Efficiency of the tracker to find a track given a cosmic ray muon as a function of the transverse impact parameter of the muon. Data is in black, and simulation in red.	101

8.14	Efficiency of the tracker to find a track given a cosmic ray muon as a function of the longitudinal impact parameter of the muon when no cut is applied on this variable. Only muons with $ d_0 < 4$ cm are used. Data is in black, and simulation in red.	101
8.15	Ratio of the efficiencies of the tracker to find a track given a cosmic ray muon for data/simulation.	102
8.16	Transverse impact parameter of reconstructed tracker tracks originating from the decays of a long-lived exotic, shown for the case $M_{H^0} = 400$ GeV/ c^2 , $M_X = 150$ GeV/ c^2	102
8.17	Efficiency of the tracker to find a track, given a cosmic ray muon reconstructed in the muon chambers, as a function of the transverse (left) and longitudinal (right) impact parameters (with respect to the nominal interaction point of CMS). The efficiency is plotted in bins of 2 cm width. For the left plot, the longitudinal impact parameter $ z_0 $ is required to be less than 10 cm, and for the right plot, the transverse impact parameter $ d_0 $ must be less than 4 cm.	103
8.18	Tag-probe pair invariant mass distributions, per probe category, employed in the measurement of the <code>HLT_L2DoubleMu23_NoVertex</code> efficiency. 109	
8.19	Tag-probe pair invariant mass distributions, per probe category, employed in the measurement of the <code>HLT_L2DoubleMu30_NoVertex</code> efficiency. 110	
8.20	p_T dependence of the single-leg <code>HLT_L2DoubleMu23_NoVertex</code> (left) and <code>HLT_L2DoubleMu30_NoVertex</code> (right) efficiency, as obtained from the tag and probe method, both for data and simulation.	111
8.21	Pseudorapidity vs p_T efficiency maps, from Data (top left), MC (top right), absolute difference (bottom left) and fraction difference (bottom right) for the path <code>HLT_L2DoubleMu30_NoVertex</code>	112

8.22	Single leg HLT_DoubleMu3 efficiency for dimuon pairs versus the angular separation ΔR (calculated at the surface of the first muon station) of the two muons. Data measurements from J/ψ (black circles), ϕ (red squares) and ρ/ω (green triangles) are compared with simulation.	114
8.23	Distribution of the angle ΔR between two tracks originating from an X boson decay, simulated for various mass points. The distribution is shown for candidates passing all selection criteria, excluding the dimuon separation cut on ΔR . The top row, reading from left to right, shows the results for $M_{H^0} = 1000 \text{ GeV}/c^2$ with $M_X = 350$ and $150 \text{ GeV}/c^2$. The second row are the results for $M_{H^0} = 1000 \text{ GeV}/c^2$ with $M_X = 50$ and $20 \text{ GeV}/c^2$	115
8.24	Distribution of the angle ΔR between two tracks originating from an X boson decay, simulated for various mass points. The distribution is shown for candidates passing all selection criteria, excluding the dimuon separation cut on ΔR . The top plot shows the result for $M_{H^0} = 400 \text{ GeV}/c^2$ and $M_X = 150$. The second row, reading from left to right, shows the results for $M_{H^0} = 400 \text{ GeV}/c^2$ with $M_X = 50$ and $20 \text{ GeV}/c^2$ and the third row for $M_{H^0} = 200 \text{ GeV}/c^2$ with $M_X = 50$ and $20 \text{ GeV}/c^2$	116
8.25	Trigger efficiency as a function of mean lifetime for the muon triggers (right) used in our analysis.	117
8.26	The transverse decay length significance for the dimuon channel with the φ and d_0 cuts removed, before (left) and after (right) the smearing described is applied.	118
8.27	The transverse decay length significance of candidates for the dimuon channel with φ difference greater than $\pi/2$	119

9.1	The 95% CL upper limits on the mean number of X bosons that could pass the selection requirements in the electron (muon) channels are shown in the left (right) plot. A yellow shaded band shows the 95% quantile for the expected limits, but is almost entirely hidden by the observed limit curves.	124
9.2	The 95% CL upper limits on σ_B for the muon channel for a H^0 mass of 1000 GeV/ c^2 . Narrow yellow shaded bands show the 95% quantiles for the expected limits.	126
9.3	The 95% CL upper limits on σ_B for the muon channel for a H^0 mass of 400 GeV/ c^2 . Narrow yellow shaded bands show the 95% quantiles for the expected limits.	126
9.4	The 95% CL upper limits on σ_B for the muon channel for a H^0 mass of 200 GeV/ c^2 (left) and 125 GeV/ c^2 (right). Narrow yellow shaded bands show the 95% quantiles for the expected limits.	127

Chapter 1

Introduction

Particle accelerators have the ability to see billions of years into the past. They have the ability to create particles that have not existed in nature since the Big Bang. With the inauguration of the Large Hadron Collider in 2009, experimental particle physics has been given the opportunity to explore energies and intensities that have not been possible at any previous accelerator. Higher energies allow experimentalists to search for heavier particles that would otherwise have been impossible to create and at these intensities it is possible to collect enough data to search for very rare interactions. This is a very exciting time in particle physics.

In addition to being a very exciting time, the field of particle physics is at its most stressful point in history. With the last particle of the Standard Model having been discovered and stricter limits on Supersymmetry and many exotic theories within just three years of turning on the LHC, the world of particle physics is wondering what if this is it?. That is not to say there is no more work to be done or questions to be answered. It is a matter of which questions will be answered at the LHC and, better yet, what new questions will arise.

This dissertation describes the search for long-lived heavy neutral particles in the dimuon channels using the Compact Muon Solenoid (CMS) detector at the Large

Hadron Collider (LHC). Several models predict long-lived heavy particles including split SUSY models, Displaced Supersymmetry [6] and the Hidden Valley model [7]. Similar searches have been performed using the D0 detector at the Tevatron [8, 9] and the ATLAS detector at the LHC [10, 11].

In Chapter 2, I provide a brief description of the Standard Model followed by the Hidden Valley model and Displaced Supersymmetry. I then provide a description of the Large Hadron Collider and the Compact Muon Solenoid in Chapter 3. The Luminosity Monitoring System is described in detail in Chapter 4. Chapter 5 describes the data and Monte Carlo simulated events used in the analysis of long-lived neutral particles decaying to muon pairs. Chapter 6 describes the event reconstruction and selection. Chapter 7 describes the background estimate. Chapter 8 describes the systematic uncertainties. The limits on the production cross section are discussed in Chapter 9.

Chapter 2

Displaced Dilepton Signatures

2.1 Standard Model




Figure 2.1: Standard Model Particles




Figure 2.2: Standard Model Particle Interactions

In this section, I will briefly describe what has come to be known as the Standard Model of particle physics. The Standard Model of particle physics describes the interactions between the observed elementary particles, leptons and quarks, through the use of Quantum Field Theory. It explains the electromagnetic force, the weak force and the strong force. It describes the interactions of leptons, quarks and force carriers.

There are six types of leptons that are grouped into three generations: electrons, muons and taus and their corresponding neutrinos. Quarks are grouped into similar generations: up and down, charm and strange, top and bottom (sometimes referred to as truth and beauty). Quarks and leptons are both spin 1/2 particles which are classified as fermions. The Standard Model force carriers include the photon, the W boson, the Z boson, gluons and the Higgs boson. Figure 2.1 illustrates all the Standard Model particles grouped by type. Figure 2.2 indicates the possible interactions, which will be described in subsequent sections.

Quantum Electrodynamics

Quantum Electrodynamics (QED) is an Abelian gauge field theory with $U(1)$ symmetry that describes the interaction of light with charged particles according to the following Lagrangian density [12].

$$\mathcal{L} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + i\psi\hat{D}\psi \quad (2.1)$$

Quantum Chromodynamics

Quantum Chromodynamics (QCD) is the non-Abelian gauge field theory that describes the strong force. QCD is described by the $SU(3)_c$ symmetry. The non-Abelian nature of QCD comes from the fact that the phases are represented by matrices and no longer commute. The $SU(3)_c$ symmetry group contains eight generators corresponding to eight gluon color combinations.

$$\mathcal{L}_{QCD} = \sum_{flavor f} \bar{q}_{f,\alpha} \left(i\hat{D} - m_f \right)_{\alpha\beta} \hat{q}_{f,\beta} - \frac{1}{4} \hat{F}_{a\mu\nu} \hat{F}^{a\mu\nu} \quad (2.2)$$

Quarks have a charge of $2/3e$ or $-1/3e$, where e is the charge of the electron. Individual quarks have not been observed experimentally. Instead, quarks are bound together into hadrons by the strong force quanta called gluons. Two types of hadrons have been observed: mesons and baryons. Mesons are quark anti-quark pairs and have integer spin while (anti-)baryons are composed of three (anti-)quarks and have half integer spin. Due to the lifetime of the top, it does not participate in hadron production.

Electroweak

The Electroweak sector is described by a $SU(2)_L \times U(1)_Y$ gauge symmetry where L is isospin and Y is hypercharge. From the bosonic terms in the Standard Model

Table 2.1: Standard Model Particle Properties - Isospin, hypercharge, electric charge

	t	t_3	y	Q
$\nu_{eL}, \nu_{\mu L}, \nu_{\tau L}$	1/2	1/2	-1	0
$\nu_{eR}, \nu_{\mu R}, \nu_{\tau R}$	0	0	0	0
e_L, μ_L, τ_L	1/2	-1/2	-1	-1
e_R, μ_R, τ_R	0	0	-2	-1
u_L, c_L, t_L	1/2	1/2	1/3	2/3
u_R, c_R, t_R	0	0	4/3	2/3
d_L, s_L, b_L	1/2	-1/2	1/3	-1/3
d_R, s_R, b_R	0	0	4/3	-1/3

Lagrangian, one can arrive at the masses for the three gauge bosons and show that the photon is massless.

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2) \quad (2.3)$$

$$Z_\mu = c_w W_\mu^3 - s_w B_\mu \quad (2.4)$$

$$A_\mu = s_w W_\mu^3 + c_w B_\mu \quad (2.5)$$

Electric charge is related to weak isospin, t , and hypercharge, y ,

$$eQ = e(t_3 + y/2) \quad (2.6)$$

Higgs Mechanism

The LHC data have confirmed that the boson that was discovered at 125 GeV is described at the current level of measurement precision by the minimal standard mode Higgs boson. The properties of the boson agree with zero spin and positive parity [13].

The self-coupling of the Higgs boson and the resulting non-zero vacuum expectation value give rise to the standard model electroweak symmetry breaking (EWSB). With λ , the Higgs self-coupling parameter, and a vacuum expectation value $v = (\sqrt{2}G_F)^{-1} = 246$ GeV, the mass of the minimal Higgs boson is given by $m_H = \sqrt{\lambda/2}v$ at tree level.

The standard model of particle physics can be summed up in the following equation

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + i\bar{\psi}\not{D}\psi + h.c. + \psi_i y_{ij} \psi_j \phi + h.c. + |D_\mu\psi|^2 - V(\phi) \quad (2.7)$$

described in References [14, 15, 12]. The gauge fields are those described by the spontaneous breaking of the Electroweak gauge group $SU(2)_L \times U(1)_y$ into QED to accompany QCD in a physical vacuum that hosts a non-zero Higgs vacuum expectation value, described by the minimum of the potential $V(\phi)$.

Fine tuning problem

The fine tuning problem arises from the one loop correction to the Higgs mass.

$$m_H^2 = m_H^2 + \frac{kg^2\Lambda^2}{16\pi^2} \quad (2.8)$$

If the cut off scale Λ is much larger than the electroweak scale, the cancellations required will be “unnatural”. The fine tuning problem implies that the Standard Model is incomplete.

The Higgs boson predicted by the Standard Model can fall into three mass regions: 1) stable 2) meta-stable 3) unstable. From the recent measurements of the Higgs boson, it is believed to be in the meta-stable region. Although the meta-stable state is not stable up the Planck scale, its stability is expected to be longer than the lifetime of the universe.

MSSM Higgs

Supersymmetry predicts five Higgs particles in order to explain away the unnatural difference between the electroweak scale and the Planck scale. The Minimal Supersymmetric model (MSSM) predicts the lightest Higgs boson to be less than 135 GeV. This region of interest is in agreement with what has been observed at the LHC. However, the lightest boson in the MSSM Higgs sector is largely degenerate with the properties of the standard model Higgs boson. Further investigation is needed to understand the physical origin of the fine tuning problem and whether Supersymmetry is a fundamental symmetry of the early universe.

2.2 Beyond the Standard Model

Several models of new physics predict the existence of massive, long-lived particles which could manifest themselves through their delayed decays to leptons. Such scenarios arise, for example, in various supersymmetric (SUSY) scenarios such as “split SUSY” [16] or SUSY with very weak R-parity violation [17], “hidden valley” models [6], and Z' models that contain long-lived neutrinos [18].

Hidden Valley Model

Models that can be classified as Hidden Valley Models are characterized by low mass (well below a TeV), neutral particles with long lifetimes [19, 7, 20, 4]. The models introduce a new confining non-Abelian gauge group G_v to the standard model. One example of Hidden Valley model, also known as a v -Model, adds a $U(1) \times SU(n_v)$ gauge group with couplings g' and g_v .

The Hidden Valley model is characterized by a set of low mass bound states below a TeV with a confining gauge interaction in a hidden sector. These bound states may communicate with the standard model through TeV suppressed gauge operators.




Figure 2.3: Hidden Valley Potential [4]

The name Hidden Valley comes from the shape of the potential (Figure 2.3) characterized by the barrier between the hidden sector and the Standard Model.

The Hidden Valley predicts a new set of quarks called v-quarks that can form bound states analogous to hadrons called v-hadrons. These v-sector particles are not charged under the standard model. For these particles to be detected at the LHC, there must be a set of mediator particles that are charged under both the Standard Model and the hidden sector.

The v-hadron mass is determined by the confinement scale or the v-quark mass, whichever is higher. The mass gap between the v-sector particles and the mediator particles is presumed to be large. The height of the barrier for which the mediator particles would have to overcome (or tunnel through) determines the strength of the interaction between the two sectors. The effective operator between the Standard Model and the Hidden Valley sector can be found by integrating out the mediating particles and is of the form

$$g_v g_{SM} \frac{O_v O_{SM}}{M^k} \quad (2.9)$$

where g_v is the coupling between the v-sector particles and the mediator, g_{SM} is the coupling constant between the mediator and the standard model particles, and k is the power to which the operator is suppressed by the mass of the mediating particle M . The mediator gauge group is described by a $U(1)_X$ symmetry. The Z' is a possible mediator.

Experimental Signature

The candidate events at the LHC are characterised as more spherical with lower thrust and a higher number of isolated leptons than events from Standard Model processes. Due to their long lifetime, these events will result in a very narrow low mass resonance with lepton pairs. For heavy v-hadrons, the decay to SM particles will be prompt. For light v-hadrons, the particles will be boosted enough such that

the experimental signature will include a displaced vertex. The analysis described in this thesis focuses on the latter [6].

2.3 Experimental Search

In what follows, the experimental setup and results are presented for the first search for new physics using data from the Compact Muon Solenoid (CMS) for the existence of massive, long-lived particles decaying to an oppositely charged pair of leptons. The search identifies events containing a pair of oppositely charged electrons or muons (dileptons) originating from a common secondary vertex within the volume of the CMS tracker, that is significantly transversely displaced from the event primary vertex. These leptons are assumed to originate from a 2-body decay of a long-lived particle, and so are required to form a narrow resonance in the dilepton mass spectrum. This topological signature has the potential to provide clear evidence for physics beyond the standard model (SM). It is also very powerful in suppressing backgrounds from standard model processes.

This signature is sensitive to a wide class of models. However, for the purpose of establishing a signal benchmark, a specific model of a long-lived, spinless, exotic particle X which has a non-zero branching fraction to dileptons is used. In this particular model, the X is pair-produced in the decay of a (non-SM) Higgs boson, i.e. $H^0 \rightarrow 2X$, $X \rightarrow \ell^+\ell^-$ [20], where the Higgs boson is produced through gluon-gluon fusion. This model predicts up to two displaced dilepton vertices in the tracking volume per event.

The D0 Collaboration has performed searches for leptons from delayed decays in its tracker volume [9, 8], but these searches are sensitive to a much smaller kinematic phase space region than CMS. The ATLAS Collaboration has performed searches

that are sensitive to decay lengths up to about 20 m by exploiting the ATLAS muon spectrometer [21, 10], using different decay channels from those considered here.

Chapter 3

The LHC and the CMS Detector

3.1 The Large Hadron Collider (LHC)

The Large Hadron Collider Figure 3.1 is the largest and most powerful particle accelerator in the world. The LHC is a proton-proton collider consisting of two counter-rotating rings beams that collide at four points around the ring. It uses the same tunnel and injector complex as the Large Electron-Positron Collider (LEP) which operated from 1989 until 2000. The LHC tunnel is approximately 27 km in circumference and 45 m to 170 m underground. The LHC tunnel has eight 528 m long straight sections connected by eight curved sections (Figure 3.1). Although there are eight interaction points, the beams only cross at the four that house the detectors.

The design energy of the LHC is 14 TeV and was operated at 7 TeV during the 2010 and 2011 physics runs and at 8 TeV during the 2012 run.

The LHC Dipole Magnets, shown in Figure 3.2, direct the proton beams around the ring. The magnet is a two-in-one design that operates at 1.9 K , from low pressure liquid helium-4 evaporate cooling. The nominal design field is 8.33 T corresponding to 7 TeV per beam. However, from 2010 to 2012 the runs were limited by the strength

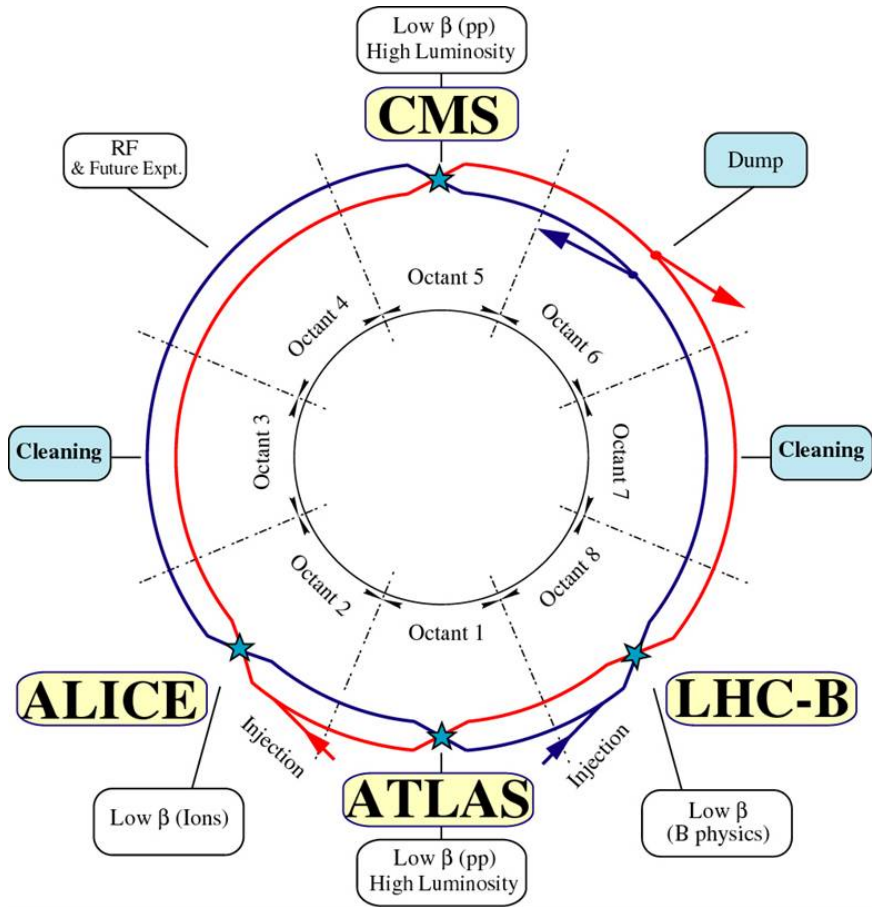


Figure 3.1: The Large Hadron Collider beam orbit crossing configuration for the four interaction regions at ALICE, ATLAS, CMS and LHC-b.

LHC DIPOLE : STANDARD CROSS-SECTION

CERN AC/DI/MM - HE107 - 30 04 1999

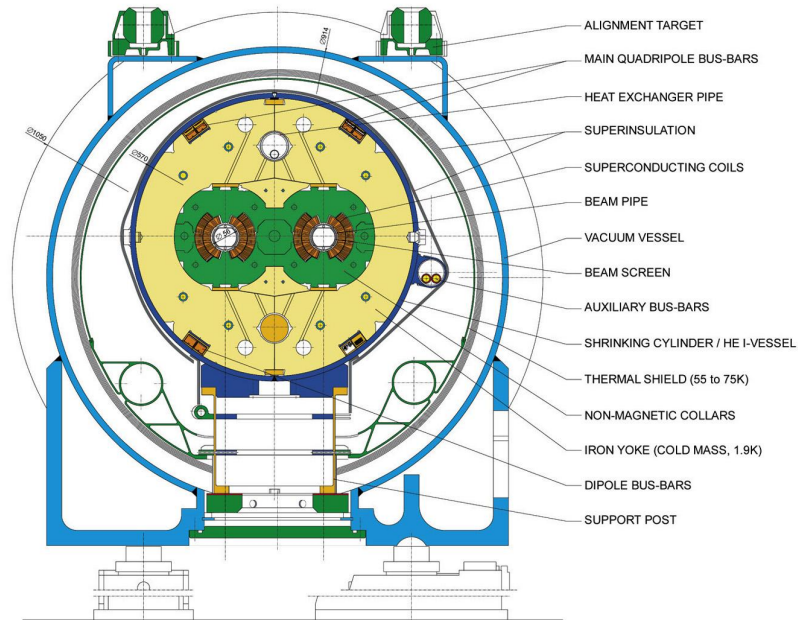


Figure 3.2: LHC Dipole Cross-section. The Dipoles are placed in the curved regions of the accelerator tunnel and are used to steer the beams. This cross section shows the “two-in-one” design of super conducting magnets which was selected as a space saving measure due to the size of the LEP tunnel.

and quality of welding at the bus bar junctions between the magnets. The focusing and defocusing of the beams are performed through the use of quadrupole magnets.

3.1.1 Accelerator Complex

The LHC injector complex includes: Linear Accelerator 2 (Linac2), Proton Synchrotron Booster (PSB), Proton Synchrotron (PS), Super Proton Synchrotron (SPS) and LHC (Figure 3.3). The Proton Linac2 accelerates the ionized hydrogen atoms up to 50 MeV and injects them into PS. PS accelerates the protons to 26 GeV. The PS also forms the bunches and provides the correct separation. SPS is the next stage where the protons reach an energy of 450 GeV in 4.3 seconds at which point they are injected into the LHC. In order to fill the LHC with the nominal 2808 bunches, this process has to be repeated twelve times. The total LHC beam current is 0.584 A.




Figure 3.3: The LHC injector complex (not to scale) showing the path from Linac2 to the PSB, PSB to the PS, PS to SPS through transfer lines, and from the SPS to the LHC.




Figure 3.4: CMS Integrated Luminosity - Proton-Proton Collisions.

At an injection energy of 450 GeV, the orbit frequency is 11,245.589 Hz . At 3.5 TeV, the orbit frequency shifts to 11,245.613 Hz . During 2010 and 2011, the LHC reached a maximum energy of 3.5 TeV per proton beam. During heavy ion collisions, where two beams of lead nuclei are used in place of the proton beams, Linac3 is used to accelerate the particles into the Low Energy Ion Ring (LEIR) where they can accumulate and then enters the PS.

The delivered integrated luminosities to the CMS experiment in 2010 to 2012 are shown in Figure 3.4 with the distribution of peak luminosities in 2011 plotted in Figure 3.5. A detailed description of the LHC can be found in Reference [22].

3.2 The Compact Muon Solenoid - CMS

This section briefly describes the major components of the Compact Muon Solenoid (CMS). A detailed description can be found in Reference [23].




Figure 3.5: Daily peak luminosity at P5 (CMS) - 2011.




Figure 3.6: Compact Muon Solenoid

The Compact Muon Solenoid (CMS) is a general purpose detector located at interaction point (IP) 5 of the LHC. The central feature of the CMS apparatus, shown in Figure 3.6, is a superconducting solenoid of 6 m internal diameter, providing a magnetic field of 3.8 T. Within the superconducting solenoid volume are a silicon pixel and strip tracker, a lead tungstate crystal electromagnetic calorimeter (ECAL), and a brass/scintillator hadron calorimeter (HCAL). Muons are measured in gas-ionization detectors embedded in the steel return yoke outside the solenoid. Extensive forward calorimetry complements the coverage provided by the barrel and endcap detectors.

Coordinate System

The x -axis of CMS points toward the center the accelerator ring, the y -axis points up and the z -axis points toward the Jura mountains or counterclockwise when viewed from above. The azimuthal angle, ϕ , is measured in the x - y plane from the positive x -axis. The origin is at the nominal interaction point. Particle production is roughly constant as a function of rapidity. For this reason, we use the spatial coordinate pseudorapidity, η , rather than polar angle, θ . Pseudorapidity, η , is defined in terms of the polar angle with respect to the beam axis, θ as

$$\eta = -\ln\left[\tan\left(\frac{\theta}{2}\right)\right]. \quad (3.1)$$

Magnet

The CMS magnet is a single solenoid positioned between the HB and HE. It provides a 3.8 T axial magnetic field with a 4-layer NbTi winding cooled to 4.2K. The strong magnetic field provides CMS with the ability to accurately measure the momentum of charged particles. The return yoke is made of 10,000 tonnes of common structural steel and has an inner diameter of 6 m and a length of 12.5 m. The CMS magnet has a stored energy of approximately 2.6 GJ when operated at 4 T.

Inner Tracker




Figure 3.7: CMS Inner Tracker - single lines indicate single sided detectors and double lines indicate double sided detectors. From [5]

In inner tracker measures charge particle trajectories with $p_t > 1 \text{ GeV}$ and reconstruction primary and secondary vertices [5]. The tracker is separated in the three pieces: the tracker inner barrel (TIB), the tracker outer barrel (TOB) and the silicon pixel detector. The inner tracker has a length of 5.8 m and a diameter of 2.6 m. This provides pseudo-rapidity coverage for charged particles of a range $|\eta| < 2.5$. With a total surface area of 200 m^2 , the inner track is the largest silicon tracker built.

Silicon Pixel Detector

The silicon pixel detector is made up of three barrel layers located at 4.4 cm, 7.3 cm and 10.2 cm and two end-cap disks on each end located at 49 cm on either side of the nominal interaction point. There are 1440 pixel modules with 66 million active elements each measuring $100 \times 150 \text{ } \mu\text{m}^2$ in $r - \phi$ and z . The number of pixels and size were chosen so such that the occupancy was limited to 10^{-4} .

Silicon Strip Detector

The silicon strip detector has 9.3 million active elements housed on 15,148 modules with 198 m^2 of active area.

The tracker inner barrel (TIB) uses four layers of silicon microstrip detectors each with a thickness of $320\text{ }\mu\text{m}$. The TIB spans from radii $20\text{ cm} < r < 55\text{ cm}$. The tracker inner disk (TID) has three layers extending radially to 55 cm on each end of the TIB.

The tracker outer barrel extends 118 cm along the z-axis on both sides of the nominal interaction point. The TOB spans radii $55\text{ cm} < r < 110\text{ cm}$ and uses silicon microstrip detectors with a thickness of $500\text{ }\mu\text{m}$. The tracker end cap (TEC) consists of 9 layers and extends out to 280 cm on either side of the nominal interaction point, complementing the TOB coverage for the forward region.

Track Reconstruction

Charged particle trajectories have 16 hits on average, where a hit is a spatial coordinate measured by an active sensor of the inner tracking detector. The basic steps of track reconstruction are seeding, trajectory building, trajectory cleaning and trajectory smoothing. During the seeding step, pairs or triples of hits are joined together to determine an initial direction and momentum of the track. The trajectory building then proceeds in the direction given by the seed to find compatible hits. The hits are found and fit using a Kalman filter. With iterative tracking, the selection criteria is loosened and the hits associated with tracks found in previous iterations are removed. For muon track reconstruction, there is a muon seeded step. “Out to in” seeds are provided by stand alone muons and “int to out” seeds are provided by tracker muons, described later in section 3.2.




Figure 3.8: Electromagnetic Calorimeter - Side View.




Figure 3.9: Electromagnetic Calorimeter - Stereo View.

Electromagnetic Calorimeter

The Electromagnetic Calorimeter (ECAL) is positioned just outside the silicon strip tracker and covers a pseudo-rapidity range of $|\eta| < 2.4$. The ECAL is made up of lead tungstate crystals ($PbWO_4$) and is the subdetector that is primarily responsible for detecting photons and electrons. The ECAL is made up of three major components: the ECAL barrel (EB), the ECAL end-cap (EE) and the Preshower (ES). The EB is made up of 61,200 pyramidal crystals mounted such that axis of each crystal makes a 3° angle with the nominal interaction vertex. The EB has a fiducial range is $|\eta| < 1.479$. The EE is constructed from 5 by 5 crystals to form a "supercrystal" structure and covers the fiducial region of $1.479 < |\eta| < 3.0$. The ES is positioned in front of the EE covering a fiducial region of $1.653 < |\eta| < 2.6$. The preshower is designed for π_0 rejection, where the two photons from the π_0 decay are spatially separated.

The lead tungstate crystals that make up the scintillating material of the EB and EE were selected because of the density, radiation length, Molière radius and scintillation time. The Molière radius determines the transverse containment of the electromagnetic shower. Within the $25ns$ LHC bunch crossing, the crystals scintillate 80% of the light and provide $25X_0$ radiation lengths of longitudinal shower containment. The light scintillated by the EB crystals is then collected by avalanche photodiodes (APDs). Two APDs are glued to the back of each crystal. The APDs are capable of measuring energies corresponding to a deposition of 100 MeV in to the crystals. Similarly, vacuum phototriodes (VTPs) to collect and measure the energy in the ECAL end-cap.

The ECAL energy resolution is characterized the function

$$\left(\frac{\sigma(E)}{E}\right)^2 = \left(\frac{2.8\%}{\sqrt{E}}\right)^2 + \left(\frac{0.12}{E}\right)^2 + (0.30\%)^2 \quad (3.2)$$

where S is the stochastic term, N is the noise term and C is the constant.

Hadronic Calorimeter

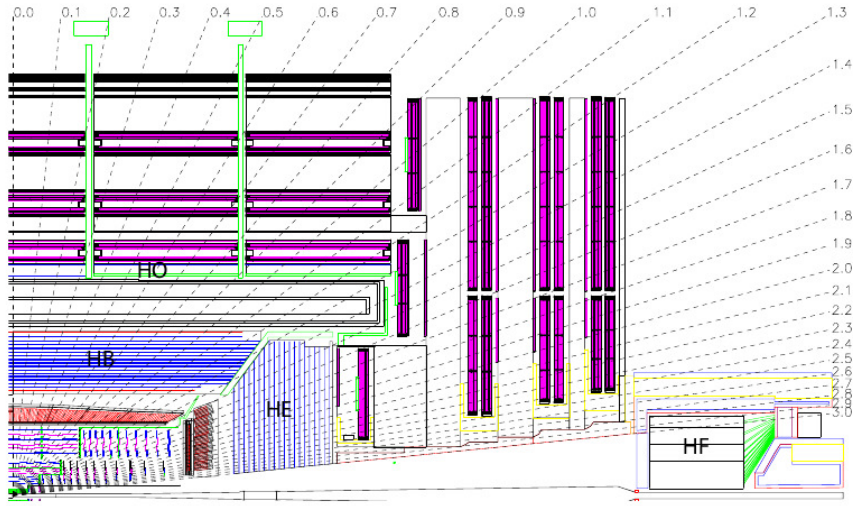


Figure 3.10: HCAL Longitudinal Cross Section.

The Hadronic Calorimeter (HCAL) is designed to measure the shower energy of hadronic particles. Many hadronic particles produced together in an angular cone are commonly referred to as a jet. Jets are measured in the HCAL as a superposition of overlapping hadronic showers. In addition, the HCAL is sensitive to individual minimum ionizing particles (MIPs) and is used to improve the efficiency of muons with low transverse momentum, where the momentum is poorly measured in the muon chambers or the muon stops in the HCAL. The HCAL consists of four sub-systems: the barrel (HB), end-cap (HE), outer (HO) and forward (HF) calorimeters.

The HE, HB and HO are segmented into towers of dimensions $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$. They are composed of brass absorber and plastic scintillator, and a light capturing wavelength shifting fiber readout. The absorber is a high density material that promotes hadronic interactions. The scintillator captures the energy from the charged particle and electromagnetic interactions in the plastic. The light from the plastic is captured and re-emitted by a wavelength shifting fiber and subsequently transported over clear fiber to a photodetector (hybrid photo-diode) that counts the number of photons from a bundle of clear fibers, known as a tower.




Figure 3.11: Hadronic Calorimeter Barrel Module.




Figure 3.12: Hadronic Calorimeter End-Cap.

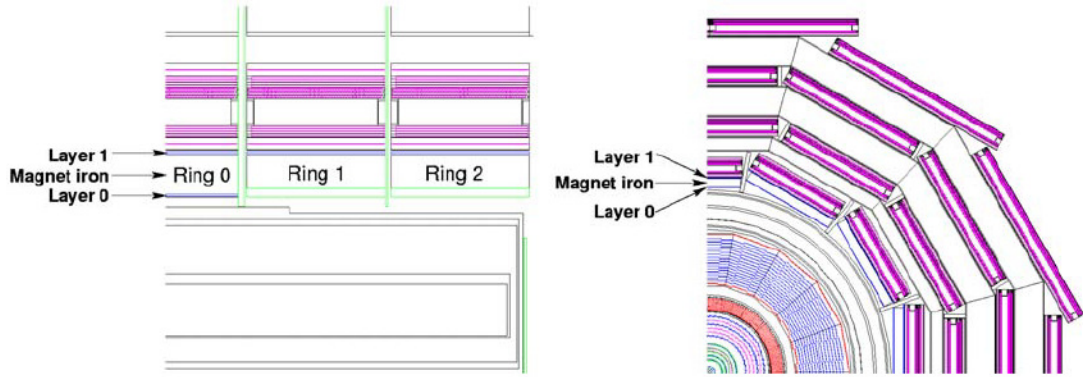


Figure 3.13: Hadronic Calorimeter Outer (HO) located in the gap between the magnet and the muon chambers.

The HB is located between the electromagnetic calorimeter and the superconducting magnet and covers a pseudo-rapidity of $|\eta| < 1.4$. HO is positioned on the outside of the magnet (Figure 3.13) and is used to improve the hadronic energy resolution by increasing the longitudinal shower containment and reduce the presence of fake muons by adding an additional interaction length in the barrel region. The HE covers the pseudo-rapidity range $1.3 < |\eta| < 3.0$ and is constructed from 600 tonnes of brass. A large portion of the brass came from melting down and recycling over one million Russian navy shells [24].

In order to improve the missing energy resolution, the HF covers a pseudo-rapidity range of $2.9 < |\eta| < 5.2$. Due to the particle flux and potential radiation damage in this region of the detector, the HF is constructed of steel absorber and quartz fibers. When highly relativistic particles pass through the quartz fibers, typically electrons and positrons from electromagnetic showers, Cherenkov radiation is produced. The light traveling down the fibers is collected by photomultiplier tubes (PMTs). In addition to tower segmentation, the HF is also segmented into depths by using two lengths of quartz fibers. The long fibers begin at the face of the absorber while the short fibers start at a depth of 220.0 mm.

The digitized signal is sent to the Hcal Trigger and Readout electronics (HTR) through fiber optic cables. On each HTR, there are two FPGAs. Once the signal arrives at the HTR, the non-linear 7-bit signal from the QIE is linearized using a 10-bit lookup table.

Muon System

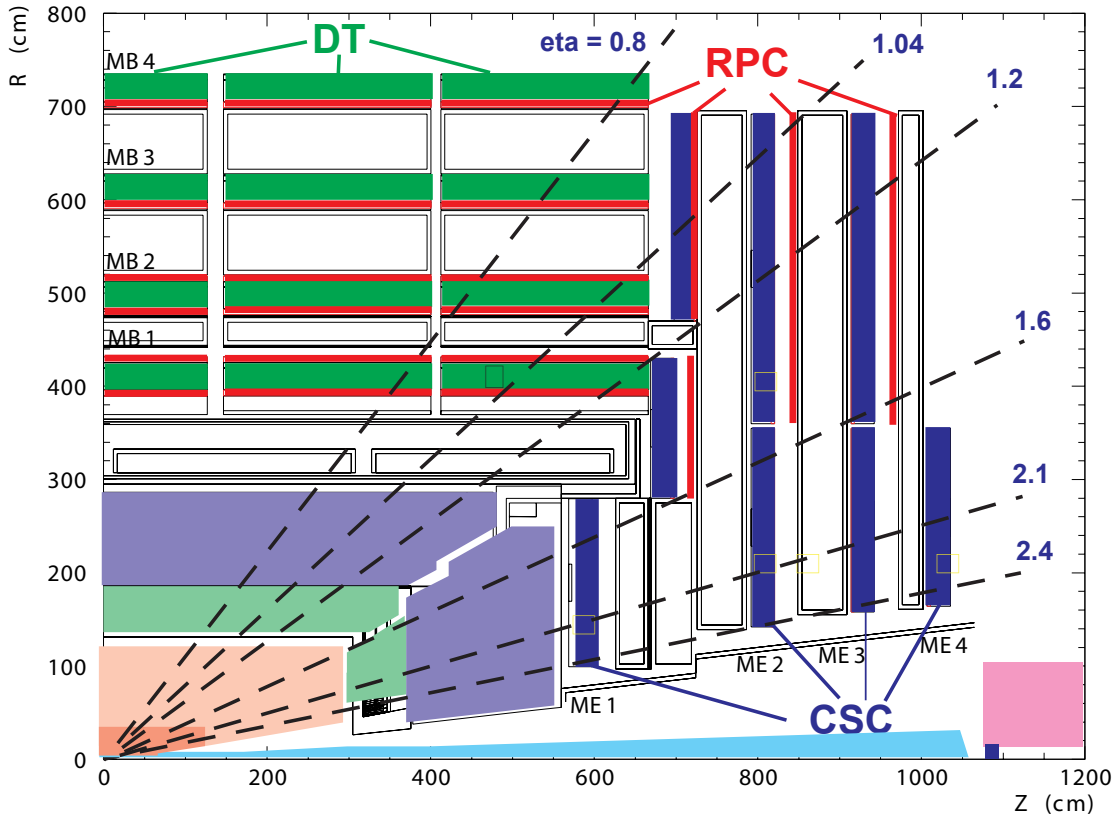


Figure 3.14: Muon Detectors Longitudinal Cross Section

The muon system is made up of three sub systems: drift tubes (DTs), cathode strip chambers (CSCs) and resistive plate chambers (RPCs). The drift tubes are located in the barrel at $|\eta| < 1.3$. The cathode strip chambers are located in the end cap at $0.9 < |\eta| < 2.4$. The resistive plate chambers are located in both the barrel and end cap at $|\eta| < 2.1$. Although the RPCs can provide some tracking information, the main function is to resolve ambiguities.




Figure 3.15: Number of nuclear interaction lengths in front of the muon stations as a function of pseudo-rapidity.

Muon Reconstruction

Part of my analysis has to do with modifying the muon reconstruction. The basic hit measurements can be summarized as follows [25]:

4 DT stations with 12 hits = 48 1D hits,

4 CSC stations with 6 hits = 24 2D hits, and

4 or 6 RPC stations with 1 hit = 4 or 6 ID hits.

Muons are reconstructed using two independent approaches: tracker muons and standalone muons. Tracker muons are tracks that are matched to a least one hit in the muon chambers. Tracker muon starts as a tracker track. These are best for low p_T . In contrast, standalone muon seeds are initial trajectory estimates from DT and CSC hit pairs forming at least 2 muon segments. The momentum estimate is from the bending in the magnetized return yoke, and no inner tracker information is used.

Global muons start as a standalone muon and find matching track. Then a refit takes all the stand alone muon hits and the tracker track to make a global muon. Using both muon chamber and tracker information improves the momentum resolution of muons above 100 GeV.

More information on the CMS muon system and reconstruction can be found in Reference [25].

Trigger and Data Acquisition System

Trigger and Data Acquisition System (TriDAS) is responsible for deciding which collisions contain interesting events and to record the detector measurements for these events for further analysis. The proton-proton cross section at the LHC design energy, $\sqrt{s} = 14 \text{ TeV}$, is approximately 100 mb. At the full luminosity of $\mathcal{L} = 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, each 25 ns bunch crossing will have on average 20 collisions which equates to a rate of 1 billion inelastic events per second. The job of the trigger system is to reduce the rate of recorded events down approximately 150–300 events per second. There are two reasons for this: 1) the subsystems are not able to record that many events per second and more importantly, most of the inelastic events that occur are from processes we understand well. If we can keep from recording “known” events, we will be able to study the more interesting ones where new physics processes are visible. A full description of the data acquisition system is found in Reference [23].

Level-1 Trigger

The Level-1 Trigger uses custom electronics to reduce the number of events consider to a maximum rate of 100 kHz. The muon trigger hardware consists of a Regional Muon Trigger (RMT) and a Global Muon Trigger (GMT) that combines track segment information from all three muon sub-detectors.

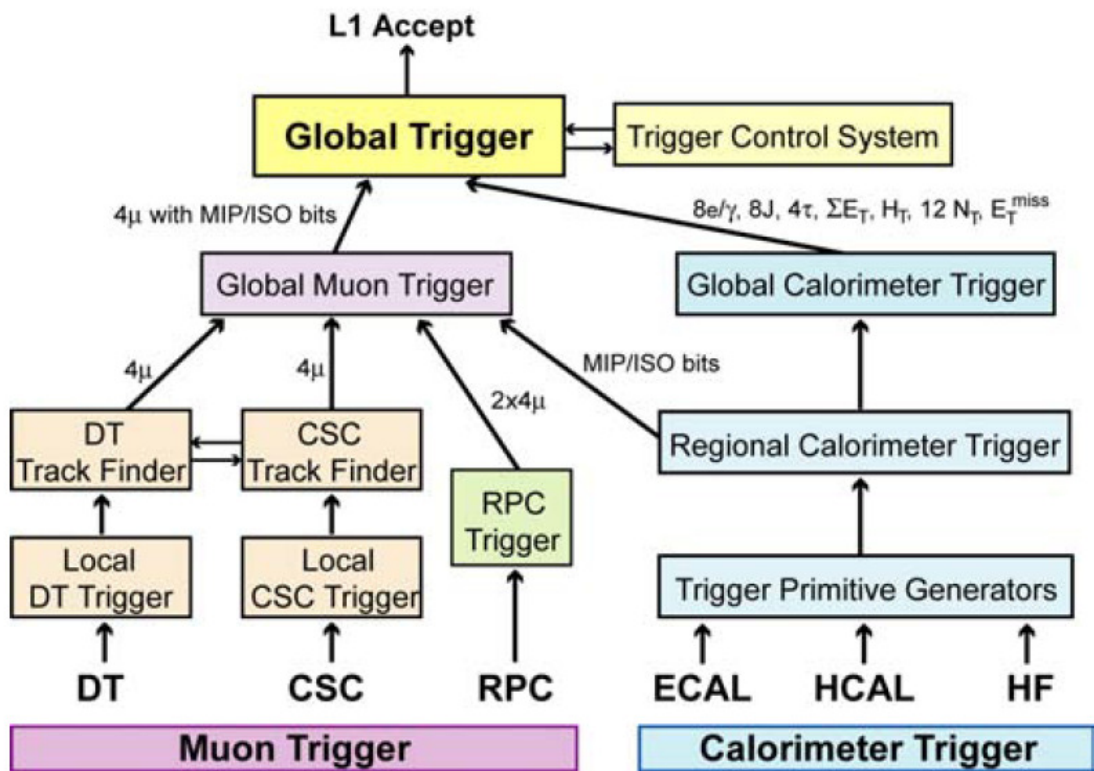


Figure 3.16: Schematic of the Level-1 Trigger System.

High-Level Trigger

Following an Level-1 Accept (L1A), the entire detector is read out into a trigger processes farm known as the High-Level Trigger (HLT). The HLT level trigger system applies over 400 individual trigger paths.

Chapter 4

The CMS Online Luminosity System

4.1 Introduction

The Luminosity Monitoring System fulfills two important services. It provides both CMS and LHC with a high precision real-time measurement of the performance of the LHC at interaction point 5 and it provides researchers with an overall normalization for physics cross section analyses. The Luminosity Monitoring System is not a stand-alone detector. Instead, it utilizes the signal from the Hadronic Forward Calorimeter (HF). Other systems like the BRAN (Beam Rate of Neutrals) and the Beam Scintillation Counters (BSC) have similar characteristics and are capable of providing online monitoring. It is expected that the Pixel Luminosity Telescope (PLT) will also be able to provide a high precision luminosity measurement in the future.

I will begin by describing the hardware and software components of the luminosity monitoring system followed by a description of the real-time calculation and calibration. I will conclude this chapter with a few noteworthy operational observations.

4.2 HF Lumi Calculation

In this section, I describe the two methods used to calculate luminosity from data acquired by the Forward Hadronic Calorimeter.

Luminosity is defined as the ratio of the number of events per unit time to the interaction cross-section.

$$\mathcal{L} = \frac{1}{\sigma_{event}} \frac{dN_{event}}{dt} \quad (4.1)$$

In a particle collider with counter-rotating beams, the luminosity

$$\mathcal{L} = \frac{N_p^2 n_b f_{rev} \gamma_r}{4\pi \epsilon_n \beta^*} F \quad (4.2)$$

where N_p is the number of particles per bunch, n_b number of bunches, f_{rev} revolution frequency, ϵ_n normalized beam emittance, β^* beta function at the crossing point, F crossing angle factor [26]. The crossing angle factor is defined as

$$F = \left(1 + \left(\frac{\theta_c \sigma_z}{2\sigma^*} \right)^2 \right) \quad (4.3)$$

where θ_c is the crossing angle to the beams, σ_z is the length of the bunch in the z and σ^* is the RMS of the beam size. From this, one can see that the luminosity can be written in terms of the average number of interactions per bunch crossing μ .

$$\mu = \frac{\sigma \mathcal{L}}{f_{rev} n_b} \quad (4.4)$$

Two standard methods for extracting a real-time relative instantaneous luminosity from the Forward Hadronic Calorimeter are used. The first is the zero counting method, in which the average fraction of empty towers is used to infer the mean number of interactions per bunch crossing.

The zero counting method uses the fact that the number number of interactions per bunch crossing follows a Poisson distribution

$$p(n; \mu) = \mu^n \frac{e^{-\mu}}{n!} \quad (4.5)$$

where n is the number of interactions in a given bunch crossing. By counting the number of bunch crossings with no interactions, we can calculate μ .

$$\mu = -\ln p(0) \quad (4.6)$$

where $p(0)$ is the fraction of empty bunch crossings.

The second method for measuring real-time luminosity uses the transverse energy sum deposited in the HF. The method relies on the fact that the transverse energy at the psuedo-rapidity range of the HF is linear with the number of interactions. Both methods provide relative luminosity measures. In order to report an absolute luminosity, a calibration scan is performed. I will discuss the results of one such scan in Section 4.7. A detailed description of the procedure employed can be found in [27, 28]. In the following sections, I will describe how these data are collected to provide a real-time luminosity measurement.

4.3 HF and HLX operation

The HLX is the primary hardware component of the CMS Luminosity Monitoring System. It is designed to construct histograms of physical quantities measured by the Hadronic Forward Calorimeters described in Section 3.2. These histograms are transmitted to a data acquisition (DAQ) system for processing, storage and monitoring.

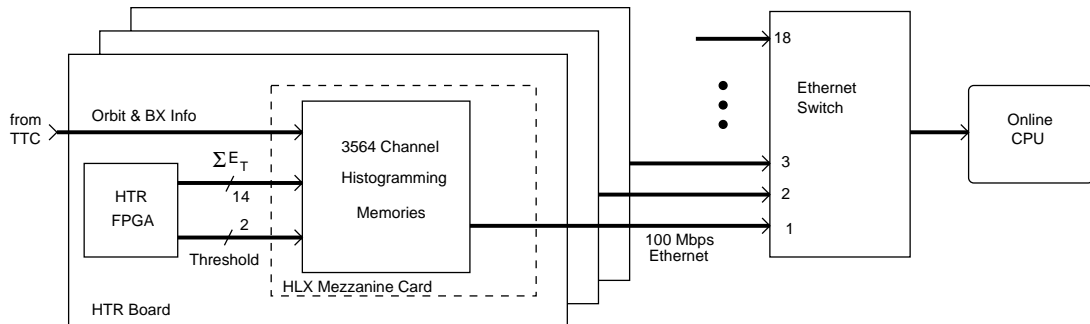


Figure 4.1: Luminosity Data Path

Each HLX is associated with a specific HF read out box (RBX) 20° wedge. Each wedge is segmented into 24 towers, with two PMTs per tower. HCal Trigger and Readout (HTR) firmware processes takes the non-linear signal from the QIE, linearizes the signal using a 10 bit look up table and passes the histogram entries to the HLXs. Each HLX accumulates eight separate histograms. Each bin in the histogram corresponds to one LHC bunch crossing of approximately 24.95 ns. These histogram entries are of two varieties: 1) transverse energy sum and 2) number of towers with energy relative to two configurable thresholds. These two thresholds separate the total number of towers occupied into three histograms. For all data taking to date, the lower threshold was set to one ADC and the upper threshold was set to five ADC. It is know that increasing the threshold from one ADC to five ADC increases the non-linearity of the zero counting method significantly while lowering the threshold increases the noise level including the afterglow. The HLX adds the entry to the histogram for one lumi nibble (2^{12} orbits). At the end of the lumi nibbles, the HLXs stream the histograms to the luminosity DAQ using UDP transport protocol. Although UDP can suffer from packet loss, it provides a fast non-blocking protocol with the ability to broadcast to any number of receivers on a single network subnet. This

allows multiple servers to receive packets in parallel for either testing or redundancy without interrupting the normal operation of the HLX and luminosity DAQ.

The HLXs use the signals from the central timing trigger and control (TTC) system to synchronize with the rest of CMS. In order to synchronize histogram buckets, the bunch counter is reset when the HLXs are sent a bunch counter zero (BC0) signal from the TTC system. In addition to running in parallel with the CMS DAQ, the HLX firmware has been designed to collect data while the remainder of CMS is stopped, not acquiring data. In the original firmware, the devices would start transmitting data once a TTC START command was received and stop upon a TTC STOP command. The HLXs now begin transmitting data as soon as they receive an orbit counter reset (OC0), which is ignored by the rest of the sub-detectors. When the TTC sends START and STOP commands to the other sub-detectors to initiate and terminate a CMS run, the HLXs remain in the running state and the state transition of CMS from stopped to running is recorded by the HLXs. The state information is then sent to the luminosity DAQ in the next nibble. Each time an OC0 is sent from the TTC to the detector and readout systems, it indicates that a new run has started. This generates a single-cycle reset pulse in the HLX that realigns the histograms with the new run and resets the orbit counter. This change of state is also indicated by each HLX to the luminosity DAQ by virtue of the fact the orbit counter has returned to zero.

4.4 Luminosity Data Acquisition Software

The Luminosity Monitoring System Data Acquisition Software is commonly referred to as the lumiDAQ. The lumiDAQ operates independent of the CMS DAQ. This is due to the fact that, unlike the CMS DAQ, the lumiDAQ is a real-time system whose data payload is time based rather than event based.

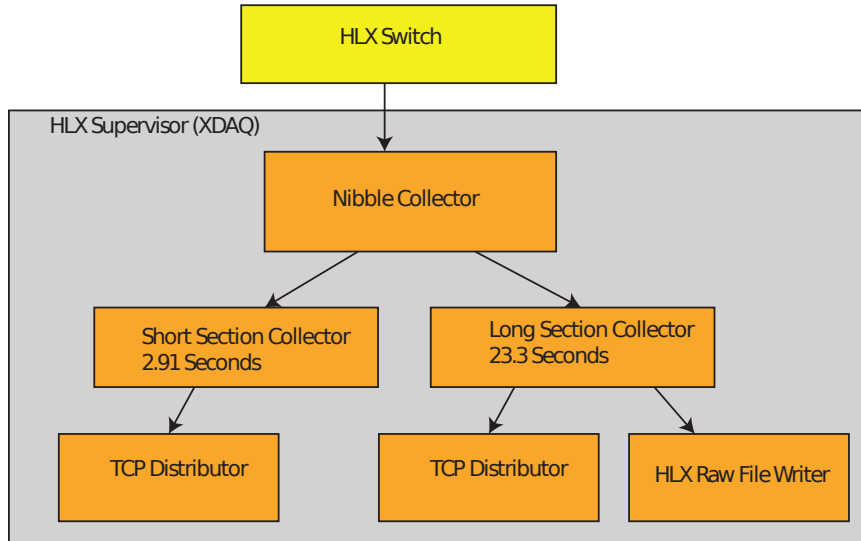


Figure 4.2: Luminosity DAQ

The lumiDAQ consists of three stages: the Nibble Collector, the Section Collectors and the Distributors. The Nibble Collector receives the UDP packets from the HLXs and merges the packets together into one data structure every lumi nibble. The data structure is then passed to the section collector where the per lumi nibble data is integrated for one lumi section. Once a lumi section is complete, the Section Collector passes the data structure to the Distributors to be processed. Because data collection and monitoring operate on different time scales, there are two Section Collectors in the lumiDAQ. The section collectors are called the Long Section Collector and the Short Section Collector. The Long Section Collector produces a lumi section approximately every 20 seconds (2^{18} orbits). This information is written to disk and later archived.

Previously, the luminosity DAQ would simply take the run number it received from RCMS and uses it as a reference for data sent to tier 0. This does not work in the case of continuous-running. Instead, the luminosity DAQ caches the next run number it receives from Run Control and waits to receive an OC0 flag in one of the luminosity nibbles from the HLXs. When it receives this signal, it updates the DAQ run number so that it equals the previously-cached value. The process repeats every time a new CMS run is started, without affecting the HLX behavior in any way.

The currently accumulated histograms will be immediately shipped as incomplete luminosity sections.

In order to achieve this goal, the implementation focused on reliability and performance. Features were abandoned when they were found to add unnecessary complexity. When the added functionality required an external library where the reliability was unknown or questionable, the task was externalized to one of the luminosity utilities.

4.4.1 Continuous Running

As mentioned previously, the Luminosity Monitoring System is also tasked with providing the LHC continuous status. The original specification of the luminosity system called for a continuous readout mode of the HF detector independent of CMS running. This, however, creates various complications due to the current nature of CMS operation. Firstly, the integration of luminosity data is quantised at the ‘lumi nibble’ boundary (4096 orbits). Ideally, CMS data taking should start and stop on this boundary. However, this was deemed to be too difficult to implement and would require changes in too many components of CMS. Therefore, a different approach was sought to avoid mis-aligned run or luminosity section boundaries. Secondly, the data from continuous operation that overlaps with a CMS data-taking run must be merged with the resulting data at tier 0, which requires an understanding of the relationship between the run data indices (orbit number and run number) and the indices in the data. Lastly, the HLX hardware in the HCAL readout that is responsible for providing luminosity data must be configured, started, and stopped in such a way that it is aware of changes in CMS running without stopping data taking or disrupting the rest of HCAL operations.

As described, these changes have been applied in a way that achieves effectively-continuous readout of luminosity data without disrupting normal CMS data taking.




Figure 4.3: Diagram of continuous running operation. The first run starts on an OC0 and luminosity data is accumulated in 4096-orbit nibbles (purple). At some point a main CMS run is started, at which point another OC0 is sent and any partially-accumulated nibble is dropped (red). The data stream continues with a nibble aligned with the new run.

- Pre-OC0 data loss (OC0 repetition period greater than 4096 orbits) - an OC0 will cause the HLX histogrammer system to restart, discarding any accumulated data. However, it will not stop a nibble currently shipping to the DAQ from being sent. In the worst case one will lose the 4095 orbits of data preceding the reset.
- Rapid multiple OC0s (OC0 repetition period less than 4096 orbits) - this will cause the HLX histogrammer to reset repeatedly, resulting in no data being sent to the DAQ as a full nibble is never completed.
- DAQ buffering limit - The online processing system takes more than a nibble-time to process data for the short (4-nibble) sections, and several seconds to complete the data processing for the longer (256-nibble) sections. In order to not lose partial sections in the advent of a run start, the DAQ software must ship on every OC0. If several follow in quick succession, this can cause a short-term overload of the luminosity DAQ. The buffers in place will tolerate up to a 10-section backlog, after which the DAQ will discard data and log the error. This is not expected to happen in typical scenarios.

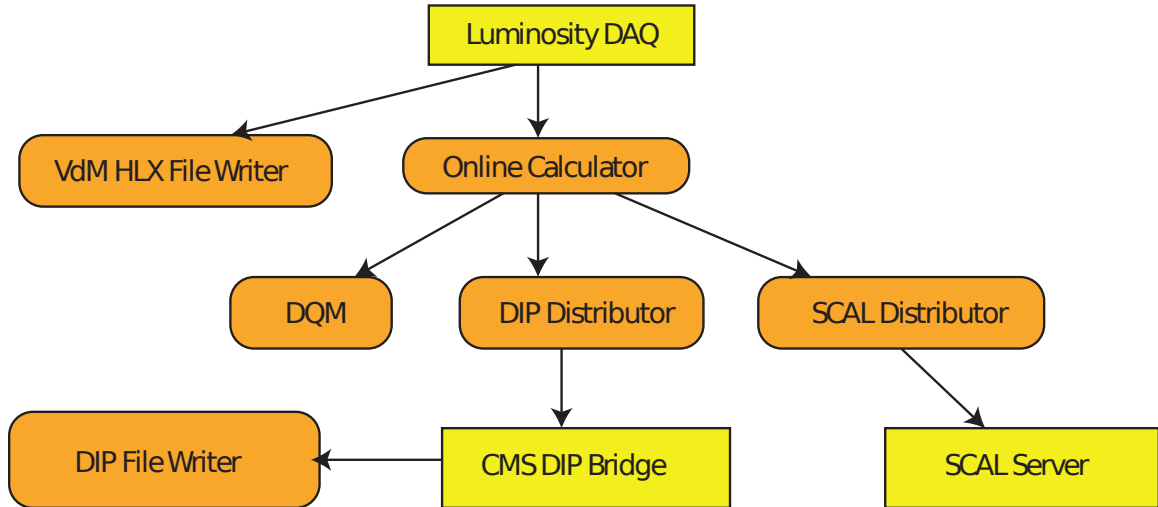


Figure 4.4: Luminosity Data Distribution

4.5 Luminosity Data Distributors

This section describes the utilities used to distribute data from the lumiDAQ to various consumers as illustrated in Figure 4.5

4.5.1 Luminosity TCP Repeater

The Luminosity TCP Repeater provides a buffer between the Luminosity DAQ and the online consumers of the luminosity information. It consists of two threads called the receiver and sender thread and a thread-safe concurrent queue provided by the open source version Intel Threading Building Blocks library [29]. The receiver thread connects to the Luminosity DAQ using the TCP Distributor class, receives the Luminosity payload described above and pushes it onto the queue. The sender thread pulls the payload off in order and distributes it to each of the consumers in series. Because of the simplicity of the receiver thread, the connection to the Luminosity DAQ has a very low risk of blocking.

4.5.2 DIP Distributor

Data Interchange Protocol is a standard communication protocol used by all of the LHC experiments to communicate with the LHC control room. The DIP Distributor publishes luminosity once every every lumi section.

4.5.3 Luminosity File Merger

The Luminosity File Merger combines the files produced by the Luminosity DAQ. The files asynchronously produced by the DIP File Writer are merged into one file per run. The File Merger runs once per hour, merging all CMS runs that completed since the last time it was run.

4.5.4 SCAL Distributor

As mentioned above, the raw uncompressed luminosity payload is just under 5MB. The SCAL Distributor was developed to combat network congestion and buffer overflow problems thought to be caused in part by the payload size. As requested by the SCAL Team, the SCAL Distributor provides integrated luminosity for run and per fill. The SCAL Distributor receives the trigger dead time from the DIP data stream mentioned above in order to provide an online measure of recorded luminosity in addition to delivered luminosity. This measure is precise enough to be used for physics analyses, but can be used to measure the operational efficiency during an LHC run.

4.5.5 Lumi Server Monitoring

The health of the luminosity servers is monitored once per day by a cron job running on each server. The scripts check the free disk space and the number of open file descriptors for each of the Lumi Utilities. An alert email is sent to system administrators when their attention is required.

4.6 HLX Supervisor

The HLX Supervisor is the XDAQ application responsible for controlling the state transitions of the luminosity DAQ (lumiDAQ), which include configuring, starting and stopping. These state transition instructions are sent by the HCAL function manager. The HLX Supervisor inherits from the `hcal::Application` class and defines the virtual functions used to transition between states. Because of the simplified state transitions of the lumi system, some of the functions only log the transition, but do not actually change the state of the lumiDAQ. Only three state transitions perform some function: `enable`, `disable` and `resume`. The other three, `pause`, `reset` and `coldInit`, are implemented, but they only send information to the logging server. They do not change the state of the lumiDAQ.

- **enable** stops the Nibble Collector if it is running, sets the run number and (re)starts the nibble collector. If fake mode is enabled, the HLX emulator will be created and started.
- **disable** stops the Nibble Collector from accepting any new packets from the HLXs. If the HLX emulator is enabled, the instance will be deleted.
- **resume** sets the next run number to be used by the lumiDAQ without stopping. When the section collector sees that the HLXs have received an OC0, the new run number is used and the lumi section counter is reset. If a new run number is not provided to the lumi DAQ before the HLXs receive an OC0, the section counter will not be reset and the OC0 will be ignored by the lumi DAQ. The basis for the design decision is due CMS Run Control operational procedures. During early running, the lumi team observed that the TTC system would send an OC0 to all the subsystems when it would resume from the paused state. This would reset the lumi section counters even though the run number did not change. This patch was implemented to prevent loss of data due to the fact that

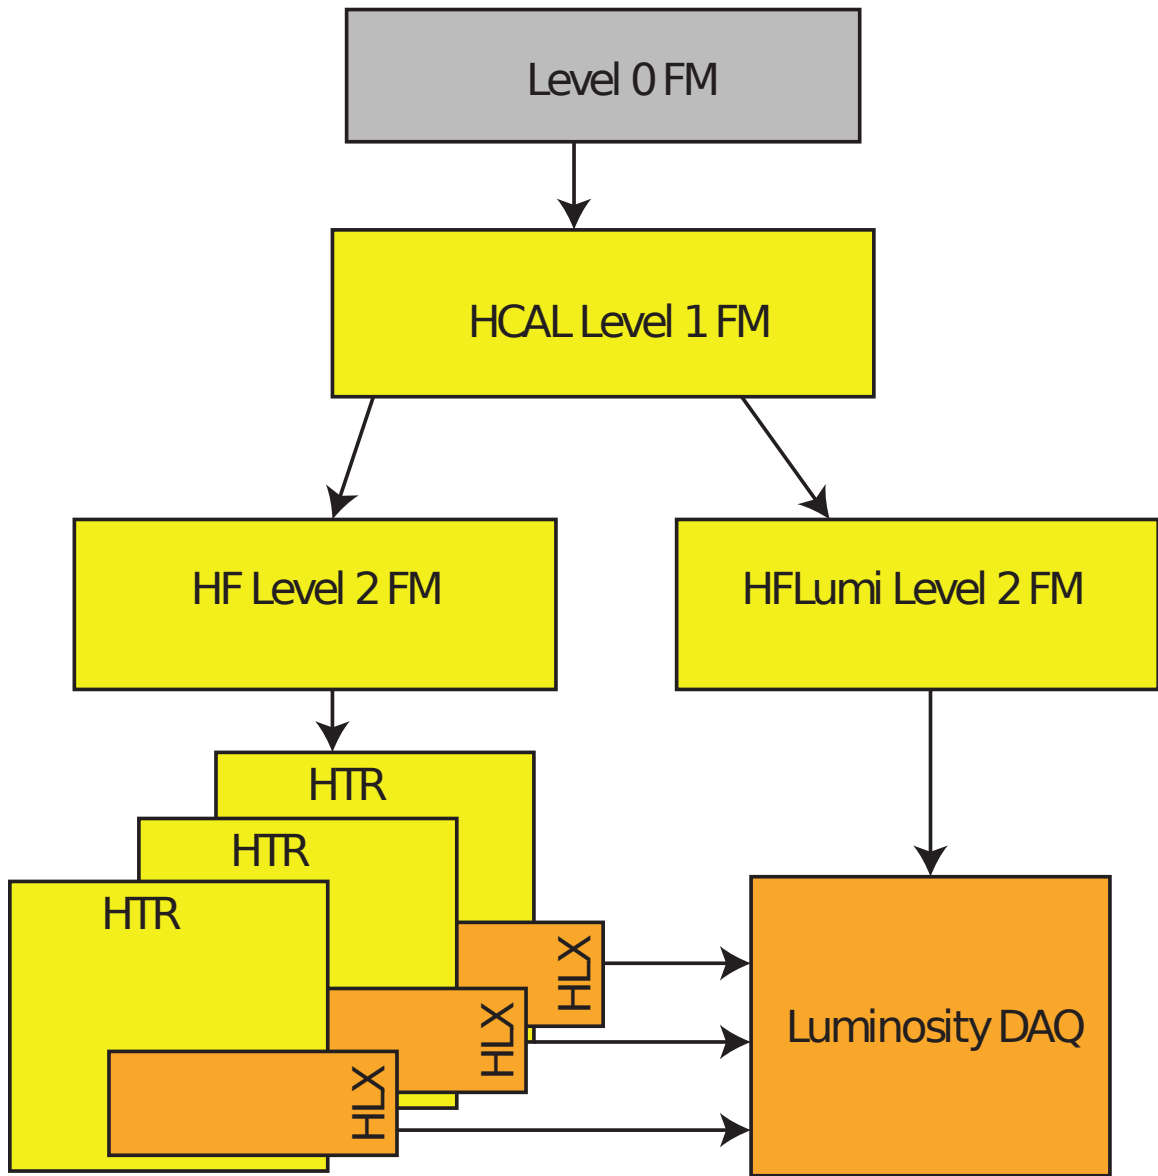


Figure 4.5: Luminosity Monitoring System Function Manager

the lumi section number and run number combination were not unique. The lumi file writer will not over write files. Instead, it will through an exception.

- **coldInit**, **reset** and **pause** are implemented, but they do not perform any function.

In order to avoid significant changes to the behavior of CMS Run Control, the Function Manager (FM) for HCAL has been extended to provide continuous operation of HF and the luminosity DAQ while hiding this behavior from top-level run control. When the HF function manager is first created, it initializes both the luminosity DAQ and HF detector readout partitions. Running then starts as normal with CMS. However, when the top-level run type is destroyed, the HF Function Manager is not, although it reports to the HCAL Function Manager that it has done so. As a result the luminosity DAQ continues to operate, as do the HLXs. When a new CMS run type is created, provided it is the same as previously used, the HF partition simply reports that it is ready and does nothing but continue to operate as before.

4.7 Absolute Calibration using Van Der Meer Scans

The separation scan method for absolute luminosity determination was pioneered by S. Van Der Meer (VdM) at the ISR [30]. The size and shape of the interaction region is measured by recording the relative interaction rate as a function of the transverse beam separations. If the beam profile (in terms of proton density) in x and y is given by the function $F(x, y) = f_x(x)f_y(y)$ then the instantaneous luminosity can be written as

$$\mathcal{L}_0 = \frac{N_1 N_2 \nu_{\text{orb}} N_b F(0, 0)}{\int f_x(\Delta x) d\Delta x \int f_y(\Delta y) d\Delta y}, \quad (4.7)$$

where \mathcal{L}_0 is the peak instantaneous luminosity, N_i is the bunch intensity in beam i , ν_{orb} is the orbit frequency, N_b is the number of colliding bunches per beam, and Δx and Δy are the beam separations in the x (horizontal) and y (vertical) planes.

The beam intensities are measured using Fast Beam Current Transformers (FBCT), which measure the current in each 25-ns LHC bunch [31]. The FBCT measurements, which provide accurate bunch-to-bunch values, are normalized to a low-bandwidth measurement of the total circulating current, made by DC current transformers.

In order to fit the tails of the distributions observed in CMS, it is necessary to use double-Gaussian distributions for the functions f_x and f_y . Namely we have

$$f_x(x) = \frac{h_x}{\sqrt{2\pi}\sigma_{1x}} e^{-\frac{x^2}{2\sigma_{1x}^2}} + \frac{(1-h_x)}{\sqrt{2\pi}\sigma_{2x}} e^{-\frac{x^2}{2\sigma_{2x}^2}}, \quad (4.8)$$

where h_x is the fraction (by area) of the Gaussian with width σ_{1x} . Nominally the Gaussian with width σ_{1x} is considered the core Gaussian, while the Gaussian of width σ_{2x} fits the tails of the distribution. A similar equation may be written for f_y . Inserting these functional forms into Equation 4.7 we find

$$\mathcal{L}_0 \equiv \frac{N_1 N_2 \nu_{\text{orb}} N_b}{2\pi \sigma_{\text{eff}}(x) \sigma_{\text{eff}}(y)} \equiv \frac{\dot{N}}{\sigma}, \quad (4.9)$$

where the effective beam size¹ $\sigma_{\text{eff}}(j)$ for each scan plane j is given by

$$\sigma_{\text{eff}}(j) \equiv \left(\frac{\sigma_{1j} \sigma_{2j}}{h_j \sigma_{2j} + (1-h_j) \sigma_{1j}} \right). \quad (4.10)$$

In general the luminosity as a function of the beam separation d , can then be written as

$$\mathcal{L} = \mathcal{L}_0 \left(\frac{h_j}{\sqrt{2\pi}\sigma_{1j}} \exp\left(-\frac{d^2}{2\sigma_{1j}^2}\right) + \frac{(1-h_j)}{\sqrt{2\pi}\sigma_{2j}} \exp\left(-\frac{d^2}{2\sigma_{2j}^2}\right) \right), \quad (4.11)$$

¹The beam size here is the convolution of both beams.

for separation of the beams in plane j .

Fits of the (uncalibrated) interaction rates as function of the separation taken during the first CMS calibration scan are illustrated in Figures 4.6 – 4.7. The fits to the scan results allow determination of the effective beam size as well as the absolute luminosity.

4.7.1 Separation Scan Measurements

In 2010 CMS had two 25-point calibration scans, in LHC fills 1058 and 1089. In fill number 1058 there were two colliding bunches in CMS, bunches 1 and 1786, and one scan for each plane (horizontal and vertical) was taken. The results for the horizontal and vertical scans in fill 1058 are shown on linear scales in Figures 4.6 and 4.7. The results are shown for the two colliding bunches combined. Separate fits to the individual bunches agree in both x and y .

In fill number 1089 there was one colliding bunch and two scans in each plane were taken. The 25 points were first scanned from negative to positive offsets, and then from positive to negative offsets. This allowed any effects due to hysteresis to be studied. Results for two scans in fill 1089 are shown on a semi-log scale in Figures 4.8 and 4.9.

The resulting effective widths from the two scans and the instantaneous luminosity as calculated using Equation 4.7, are given in Table 4.1. The calibration of the raw CMS luminosity for the two scans is calculated using the peak luminosity from the zero points taken during the scan.

The two offline luminosity methods were used to cross check the scans based on the online HF luminosity method. The results are shown in Table 4.2. With the exception of the y -scan of Fill 1058, where the CMS DAQ was not working during part of the scan, the agreement among the methods is extremely good.

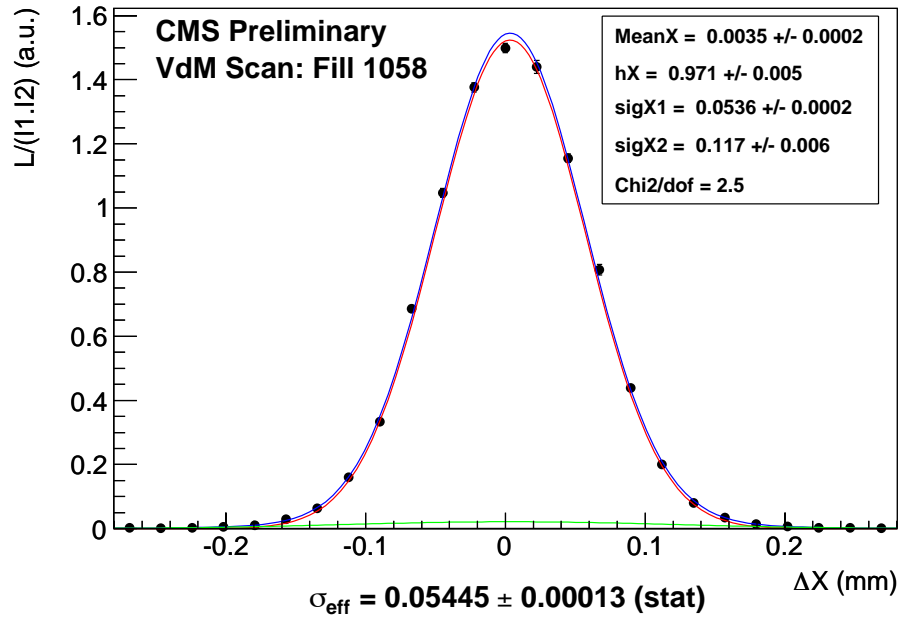


Figure 4.6: Van der Meer scan results in x , from LHC fill 1058. The blue curve is the total double-Gaussian, the red curve is the core Gaussian (σ_1), and the green curve is the Gaussian for the tails (σ_2).

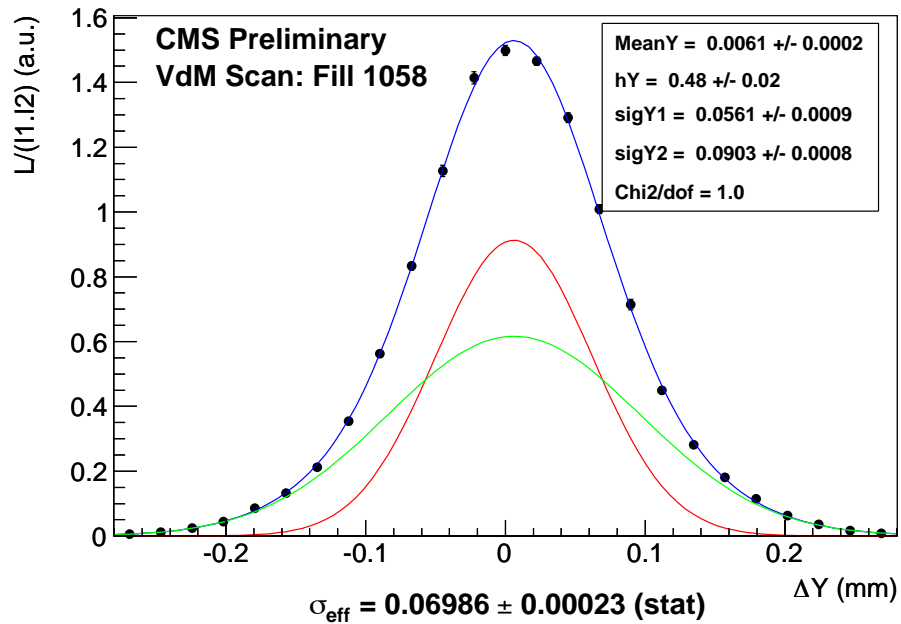


Figure 4.7: Van der Meer scan results in y , from LHC fill 1058. The blue curve is the total double-Gaussian, the red curve is the core Gaussian (σ_1), and the green curve is the Gaussian for the tails (σ_2).




Figure 4.8: Van der Meer scan results in x , from LHC fill 1089. The blue curve is the total double-Gaussian, the red curve is the core Gaussian (σ_1), and the green curve is the Gaussian for the tails (σ_2).




Figure 4.9: Van der Meer scan results in y , from LHC fill 1089. The blue curve is the total double-Gaussian, the red curve is the core Gaussian (σ_1), and the green curve is the Gaussian for the tails (σ_2).

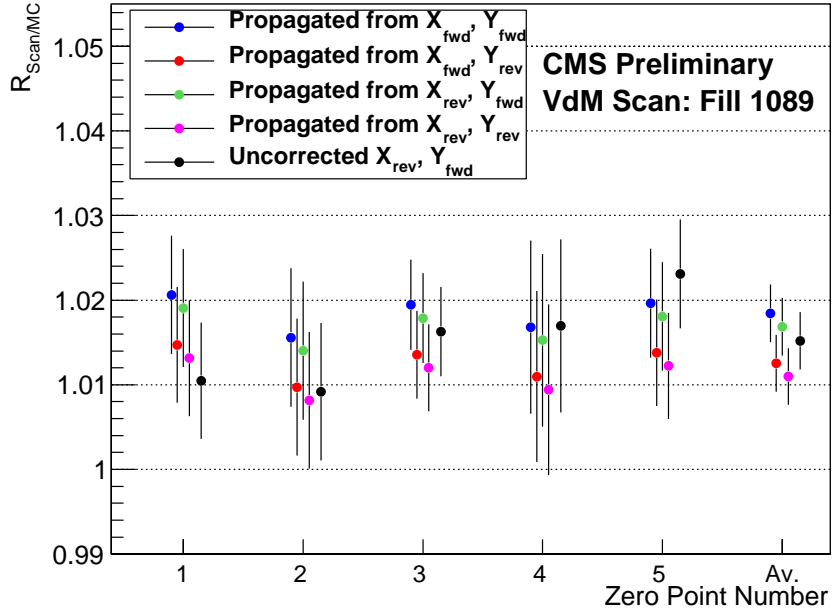


Figure 4.10: Calibration relative to MC normalization at the five central zero points in Fill 1089. The different sets of points show the corrections as propagated from the four different pairs of measured x and y beam widths. The uncorrected points for the reverse x and forward y pair are also shown in black.

4.7.2 Results

From the separation scans taken during fill 1058 and fill 1089, a fit of the (uncalibrated) specific luminosity as function of the separation between the beams in each plane allowed determination of the effective beam size as well as the peak luminosity. The calibration of the raw CMS luminosity for the two scans given in Table 4.1 is calculated using the peak luminosity from the zero points (points where the beam offsets were set to zero) taken during the scan.

For each set of scans, we use the average of five zero points, which were evenly distributed throughout the scan interval. To each zero point a correction for changes in the beam emittance, ϵ (elliptical area of the beam in phase-space) is applied. Increasing beam emittance over the duration of the fill causes a corresponding increase in the beam widths since, $\epsilon \propto \sigma^2$. The beam width measured at a particular zero point therefore, will not be correct at a later point in time. The rate at which the beam

Table 4.1: Results from the 2010 CMS calibration scans using the Van der Meer method. For the final calibration of the HF measurement, the results from the two scans are averaged.

	Fill 1058	Fill 1089
$\sigma_{\text{eff}}(x)$ (mm)	0.05445 ± 0.00013	0.05525 ± 0.00012
$\sigma_{\text{eff}}(y)$ (mm)	0.06986 ± 0.00023	0.05926 ± 0.00012
N_1	1.02×10^{10}	1.91×10^{10}
N_2	0.919×10^{10}	2.24×10^{10}
\mathcal{L} from Scan ($\text{cm}^{-2}\text{s}^{-1}$)	$(8.79 \pm 0.02) \times 10^{27}$	$(23.36 \pm 0.03) \times 10^{27}$
\mathcal{L} from MC Calibration ($\text{cm}^{-2}\text{s}^{-1}$)	$(9.05 \pm 0.05) \times 10^{27}$	$(23.00 \pm 0.06) \times 10^{27}$
Ratio Scan Calib. to MC Calib.	0.969 ± 0.006	1.017 ± 0.003

width increases, however, can be obtained from the slope of a linear fit to the beam widths as calculated from emittances measured during the fill. The corrected beam width at any zero point is given by propagating from one of the fitted widths using this slope. The uncorrected and emittance corrected ratios at each of the zero points are shown in Figure 4.10. For the final calibration we use the result as propagated from the x and y scans closest together in time, the positive to negative (reverse) x scan and the negative to positive (forward) y scan.

In some cases, the beams were not perfectly centered, meaning that the luminosity at the zero points is not the true maximum. In such cases, the raw luminosity used for the calibration is corrected using the probability density functions from the beam profile fits. This correction was as large as 1.8%.

The final calibration results from the two scans are also given in Table 4.1. The bottom line in the table shows the factors by which the interim MC calibration must be corrected to bring it in line with the calibration from the scans. There is a 5% difference between the normalization correction factors for the two scans. Combining them in a weighted average gives an overall normalization factor of 1.007 ± 0.003 .

Table 4.2: Comparison of beam widths between online and offline methods. “Forward” and “backward” refer to the scan direction—i.e., whether the coordinate being scanned increases or decreases with time. The starred widths (**) come from a scan that was missing many data points for the offline methods. These numbers should not be considered accurate. All values are in mm.

Fill	Scan	$\sigma_{\text{eff}}(i)$ HF Offline	$\sigma_{\text{eff}}(i)$ Vtx Offline	$\sigma_{\text{eff}}(i)$ HF Online
1058	X	0.05388	0.05412	0.05445 ± 0.00013
	Y	0.06491**	0.06464**	0.06986 ± 0.00023
1089	X forward	0.05513	0.05534	0.05503 ± 0.00012
	X backward	0.05531	0.05534	0.05525 ± 0.00012
	Y forward	0.05906	0.05940	0.05926 ± 0.00012
	Y backward	0.06001	0.06007	0.05985 ± 0.00010

4.7.3 Systematic Error

A summary of the systematic errors and total estimated systematic error on the measurement can be found in Table 4.3. The total systematic uncertainty is 11%. This error is dominated by the measurement of the beam currents, which have an RMS measurement error of 5% per beam. At this writing, possible correlations between the beam intensity measures are still under study. We therefore conservatively assume that the beam intensity measurements are completely correlated and add the uncertainties for the two beams linearly.

Another source of systematic error is the beam-shape uncertainty, which is associated with the possibility that the double gaussian is not a perfect match to the actual beam shape. A systematic uncertainty of 3% from this source was estimated by replacing the double-gaussain fit with a spline fit.

Other sources of systematic error include a fit systematics term, which is based on variations in the fit parameters observed fitting the offline distributions (see Table 4.2); a scale-calibration error associated with possible inaccuracies in the methods used to determine the beam offsets; and a zero-point uncertainty, related to the variation of the beam size (emittance growth) during the scans.

Table 4.3: Systematic errors on the 2010 CMS calibration scan measurements using the Van der Meer method.

Error	Value (%)
Beam Background	0.1
Fit Systematics	1.0
Beam Shape	3.0
Scale Calibration	2.0
Zero Point Uncertainty	2.0
Beam Current Measurement	10.0
Total	11.0

We anticipate that as the LHC beam currents increase and further efforts are made to understand the bunch-current measurements, the systematic error from this source should drop significantly, at which point the other errors listed in Table 4.3 will gain in relative importance. We further anticipate that uncertainties not related to the current will also decrease, as we carry out additional scans and gain a better understanding of the sources of error.

One additional sort of systematic uncertainty, not accounted for in this analysis, is the “calibration drift” error. Specifically, the results summarized in Tables 4.1 and 4.3 refer to data taken at the time of the scans. Drifts in gains and pedestals can result in a change in the calibration constant. These effects are expected to be small (at the percent level), with the precise value depending on the details of the data sample under study.

4.8 Lumi Operational Considerations

Due to the integrations capabilities of the HLXs, several signal artifacts were more prominent. In this section, I will describe these features in slight detail.

The first of these features is the QIE noise pedestal. The QIE contains four capacitor banks that are used in a round robin order. From figure 4.11, we can clearly see the separation of the pedestal means. Although, at production luminosity




Figure 4.11: QIE Pedestals prior to LUT subtraction

levels, this is a small number, correcting this effect by setting the QIE pedestal to zero improved the noise level. This also exposed other noise sources.

It was believed that resetting the QIE capacitory banks every orbit was required for proper operation. This produced a periodic effect on the order of 11kHz with an amplitude much larger than the random noise as shown in Figure 4.12. When the QIE reset was reduced to once every 103 orbits, the periodic effect dropped below the Gaussian noise. This irreducible noise is on the order of $10^{25} \text{ cm}^{-2} \text{ s}^{-1}$.

The final source of noise I will mention was first observed during laser and LED runs. The PMTs have the advantage of containing most of the signal from an event in the first 25 ns bunch crossing. However, the PMTs exhibit a long tail that last for more than 100 bunch crossings after falling below the noise as shown in Figure 4.13. In addition to the tail, there are several peaks in the signal that are due to heavy ions being released from dynodes and then striking the photocathode releasing additional electrons. This was verified by adjusting the PMT voltage and observing the proportional movement of the ion feedback peaks away from the main peak when the voltage decreased and toward the main peak when the voltage increased.

During a series of runs, we pulsed a set of LEDs in from of the PMTs in order to simulate events. We integrated the signal from each HLX for 15 93 second lumi

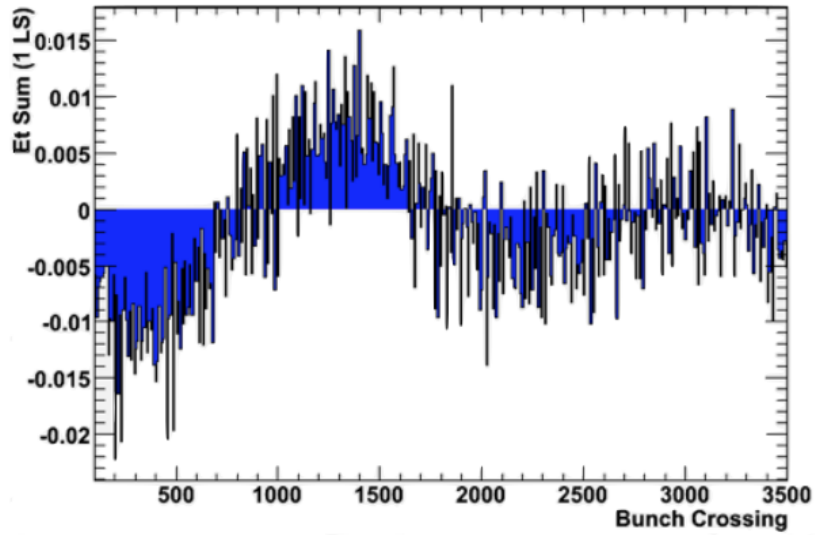


Figure 4.12: HLX ET Sum Histogram showing QIE Reset effect




Figure 4.13: Afterglow from LED runs normalized to the peak height. E_T sum method... The vertical error bars indicate the total spread. The peaks after the main peaks are created by ion feedback.

sections. Because the distance between the LEDs and the PMTs is not equal for all PMTs, the magnitude of the signal is reduced for PMTs farther away. In order to correct for this effect, we normalized the signal from each HLX such that the peak signal is equal to one (Figure 4.13). We found that every HLX exhibits the same effect in approximately the same proportions.

Chapter 5

Data and Monte-Carlo Samples

For this analysis, pp collision data at a centre-of-mass energy of 7 TeV corresponding to an integrated luminosity of $5.1 \pm 0.1 \text{ fb}^{-1}$ are used in muon channels. A parallel study of the electron channel was performed, but it is not covered in this dissertation. For the muon channel, the trigger requires two muons, each reconstructed without using any primary vertex constraint and having $p_T > 30 \text{ GeV}/c$ and made from at least two track segments in the muon systems. The tracker information is not used in either trigger.

The data are collected with the double muon trigger (`HLT_L2DoubleMu23_NoVertex` or `HLT_L2DoubleMu30_NoVertex`). The track-finding algorithm used in the HLT is very fast and simple, but is only able to reconstruct tracks produced near the beam spot. For this reason, the muon channel uses a trigger that demands only that the muons be found in the muon chambers. In addition, these muons are reconstructed in the HLT without using any beam spot constraint when fitting their trajectories. (Such a constraint is applied in all other muon triggers that rely solely on the muon chambers). A detailed study of the performance of these triggers for the displaced leptons search can be found in References [32, 33].

Events are selected by requiring two identified muons in the muon systems, not necessarily originating from the beam spot, with $p_T > 33$ GeV/ c each.

Signal Monte Carlo (MC) samples were generated using PYTHIA V6.424 [34] to simulate H^0 production through gluon fusion ($gg \rightarrow H^0$). Using PYTHIA PYUPDA cards, subsequently the H^0 was forced to decay to two long-lived spin 0 exotic particles ($H^0 \rightarrow XX$), which each then decayed to dileptons ($X \rightarrow \ell^+ \ell^-$). Samples with different combinations of H^0 masses ($M_{H^0} = 125, 200, 400, 1000$ GeV/ c^2) and X boson masses ($M_X = 20, 50, 150, 350$ GeV/ c^2) were generated. These are listed in Table 5.1. The X lifetime used in these samples was chosen to give a mean transverse decay length of approximately 20 cm in the laboratory frame. These datasets have names of the form HTo2LongLivedTo4F_MH-120_MFF-20_CTau-130_7TeV-pythia6 and each contain 50 thousand events. Several simulated background samples were used. They are listed in Table 5.2 and Table 5.3. All data and simulated Monte Carlo samples were reconstructed with CMSSW 4_2_7 or CMSSW 4_2_8. The simulated samples correspond to the ‘Summer11’ production. An example of a generated event, for a H^0 mass of 400 GeV/ c^2 and X boson mass of 150 GeV/ c^2 , is shown in Figure 5.1. In this event, one X boson decays to dimuons and the other to dielectrons.

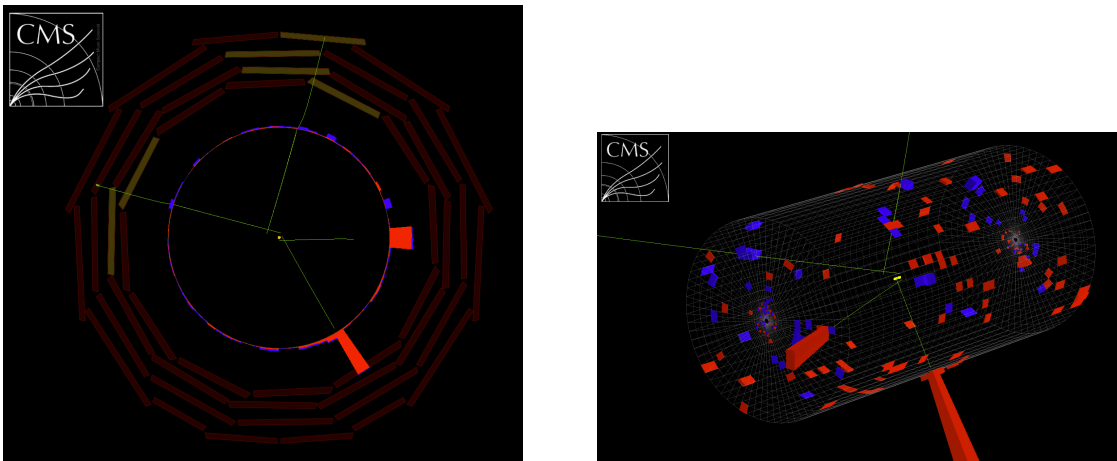


Figure 5.1: An example of a simulated Monte Carlo event with a H^0 mass of 400 GeV/ c^2 and X boson mass of 150 GeV/ c^2 .

Table 5.1: Simulated signal samples used in the analysis. The masses of the H^0 and X bosons are given, as is the mean proper decay length of the X boson.

M_{H^0} (GeV/ c^2)	M_X (GeV/ c^2)	$c\tau$ (cm)
1000	350	35.0
1000	150	10.0
1000	50	4.0
1000	20	1.5
400	150	40.0
400	50	8.0
400	20	4.0
200	50	20.0
200	20	7.0
125	50	50.0
125	20	13.0

Table 5.2: Simulated background samples used in the analysis. The weights shown are equal to the cross section divided by the number of simulated events used.

Dataset name	Cross section (pb)	Weight factor (fb)
TT_TuneZ2_7TeV-pythia6-tauola	94.0	8.63×10^{-2}
WWto2L2Nu_TuneZ2_7TeV_pythia6_tauola	2.93	1.39×10^{-2}
ZZto2L2Nu_TuneZ2_7TeV_pythia6_tauola	0.193	8.76×10^{-4}
WZto3LNu_TuneZ2_7TeV_pythia6_tauola	0.339	1.66×10^{-3}
DYtoEE_M-10to20_TuneZ2_7TeV-pythia6	2659	2.83
DYtoEE_M-20_TuneZ2_7TeV-pythia6	1300	5.92×10^{-1}
DYtoMuMu_M-10to20_TuneZ2_7TeV-pythia6	2659	1.36
DYtoMuMu_M-20_TuneZ2_7TeV-pythia6	1300	6.43×10^{-1}
DYtoTauTau_M-20_TuneZ2_7TeV-pythia6-tauola	1300	6.43×10^{-1}
ZJetToEE_Pt-20to30_TuneZ2_7TeV_pythia6	131	7.94×10^{-1}
ZJetToEE_Pt-30to50_TuneZ2_7TeV_pythia6	84.0	5.09×10^{-1}
ZJetToEE_Pt-50to80_TuneZ2_7TeV_pythia6	32.3	3.03×10^{-1}
ZJetToEE_Pt-80to120_TuneZ2_7TeV_pythia6	9.99	9.08×10^{-2}
ZJetToEE_Pt-120to170_TuneZ2_7TeV_pythia6	2.74	2.55×10^{-2}
ZJetToEE_Pt-170to230_TuneZ2_7TeV_pythia6	0.722	6.56×10^{-3}
ZJetToEE_Pt-230to300_TuneZ2_7TeV_pythia6	0.194	2.16×10^{-3}
ZJetToEE_Pt-300_TuneZ2_7TeV_pythia6	0.0758	9.48×10^{-4}
ZJetToMuMu_Pt-20to30_TuneZ2_7TeV_pythia6	131	1.59
ZJetToMuMu_Pt-30to50_TuneZ2_7TeV_pythia6	84.0	5.78×10^{-1}
ZJetToMuMu_Pt-50to80_TuneZ2_7TeV_pythia6	32.3	3.03×10^{-1}
ZJetToMuMu_Pt-80to120_TuneZ2_7TeV_pythia6	9.99	2.50×10^{-1}
ZJetToMuMu_Pt-120to170_TuneZ2_7TeV_pythia6	2.74	1.37×10^{-1}
ZJetToMuMu_Pt-170to230_TuneZ2_7TeV_pythia6	0.722	4.87×10^{-2}
ZJetToMuMu_Pt-230to300_TuneZ2_7TeV_pythia6	0.194	6.63×10^{-2}
ZJetToMuMu_Pt-300_TuneZ2_7TeV_pythia6	0.0758	1.49×10^{-3}

Table 5.3: Simulated QCD background samples used in the analysis. The quoted cross sections have been scaled by the efficiency of the event filter in the case of the muon-enriched QCD samples. For the search, the normal QCD and Mu-enriched QCD samples are used. The weights shown are equal to the cross section divided by the number of simulated events used.

Dataset name	Cross section (pb)	Weight factor (fb)
QCD_Pt-1800_TuneZ2_7TeV_pythia6	3.58×10^{-4}	1.23×10^{-5}
QCD_Pt-1400to1800_TuneZ2_7TeV_pythia6	1.09×10^{-2}	6.21×10^{-5}
QCD_Pt-1000to1400_TuneZ2_7TeV_pythia6	0.332	1.59×10^{-2}
QCD_Pt-800to1000_TuneZ2_7TeV_pythia6	1.84	4.56×10^{-2}
QCD_Pt-600to800_TuneZ2_7TeV_pythia6	15.6	4.27×10^{-3}
QCD_Pt-470to600_TuneZ2_7TeV_pythia6	70.2	1.87×10^{-2}
QCD_Pt-300to470_TuneZ2_7TeV_pythia6	1170	1.88×10^{-1}
QCD_Pt-170to300_TuneZ2_7TeV_pythia6	2.43×10^4	4.14
QCD_Pt-150_MuPt5Enriched_TuneZ2_7TeV-pythia6	2.87×10^3	7.26×10^{-1}
QCD_Pt-120to150_MuPt5Enriched_TuneZ2_7TeV-pythia6	5.52×10^3	5.66×10^{-1}
QCD_Pt-80to120_MuPt5Enriched_TuneZ2_7TeV-pythia6	2.85×10^4	2.53×10^1
QCD_Pt-50to80_MuPt5Enriched_TuneZ2_7TeV-pythia6	1.40×10^5	1.29×10^1
QCD_Pt-30to50_MuPt5Enriched_TuneZ2_7TeV-pythia6	5.97×10^5	5.49×10^1
QCD_Pt-20to30_MuPt5Enriched_TuneZ2_7TeV-pythia6	1.34×10^6	1.33×10^2
QCD_Pt-15to20_MuPt5Enriched_TuneZ2_7TeV-pythia6	1.67×10^6	8.03×10^2

Chapter 6

Event Selection and Performance

6.1 Displaced track reconstruction

CMS exploits a so-called ‘iterative tracking’ algorithm to reconstruct tracks in the tracker [35, 36]. The first tracking iteration is dedicated to finding tracks originating near the primary vertex, since these are easiest to reconstruct. It seeds the tracks with pairs or triplets of hits from the pixel tracker, which provides high resolution 3D position measurements. It then extrapolates these seeds outwards, assigning to the track additional hits from the pixel or strip trackers using the Kalman filter algorithm. Hits assigned to these tracks are excluded from further searches, so simplifying the task for subsequent iterations. A total of five iterations is used, with some of the additional ones being dedicated to finding very low momentum tracks, and others being dedicated to finding highly displaced tracks. The reconstruction of very displaced tracks uses seeds produced from hits in pairs of strip tracker stereo layers (since these provide 3D hit position measurements). As these have a very poor resolution in z , these seeds are of poorer quality than those produced from pixel tracker hits. A beam spot constraint is used to help find these tracks, as is the case for the earlier tracking iterations. However, the constraint is extremely loose (tens of centimetres).

The displaced tracking algorithms, whose initial development was largely motivated by this analysis, are described in detail in Reference [37]. The performance of the track reconstruction algorithms has been studied with data [5].




Figure 6.1: Track reconstruction efficiency for single, isolated muons of $p_T = 50$ GeV/ c as a function of the transverse impact parameter. The efficiency has little dependence on p_T .

The tracking efficiency is very small for tracks whose impact parameters exceed about 25 cm. This is a result of the loose beam spot constraint used during track finding. It is also zero for tracks produced more than 50 cm from the beam line, because of the location of the seeding layers. Figure 6.1 shows the tracking efficiency for isolated particles, (obtained with a particle gun simulation), as a function of their transverse impact parameter. When making this plot, only tracks classified as *high purity* are used. This requirement, which is imposed on all tracks used in this paper, is defined in Reference [5], but in essence selects good quality tracks by imposing cuts on the track's χ^2 , number of assigned hits, etc.

6.2 Displaced lepton identification

The standard muon reconstruction algorithms [25] offer the possibility of reconstructing muons either with the muon chambers alone (*stand-alone muons*) or in combination with the tracker (*global muons*). The former should be particularly useful for a displaced lepton search, since they offer the possibility to find highly displaced muons that did not produce enough hits in the tracker to be reconstructed there. Unfortunately, the fitted trajectories of these stand-alone muons exhibit a bias towards the beam spot. (This arises because a weak beam spot constraint is used when seeding these muons. In principle, this constraint should be removed by a final fit to obtain their trajectories, known as ‘smoothing’. However, this smoothing step is disabled during stand-alone muon reconstruction, as it has been observed to cause inefficiencies). Although this bias could be removed by refitting the muon trajectories, and indeed was successfully done for an earlier (unpublished) analysis of the 2010 data. However, rerunning the reconstruction requires using RECO instead of AOD datasets, which is not practical for the large 2011 Monte Carlo samples.

In order to allow the analysis to be run on the AOD, no specific offline lepton identification is therefore applied in this analysis. Instead, tracks are considered to be *identified as leptons* if they can be matched to ‘trigger objects’ from the HLT_`L2DoubleMu23_NoVertex` or HLT_`L2DoubleMu30_NoVertex` triggers, within a cone of size $\Delta R < 0.1$ (where $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$).

6.3 Selection of long-lived exotica

For the online selection, dedicated triggers were used, designed to ensure good trigger efficiency for displaced leptons. For the muon channel, the trigger requires two muons reconstructed in the Muon Detector with $p_T > 30$ GeV/ c . Further details of these are given in Chapter 5.

Events are required to contain a primary vertex[38] with at least four associated tracks whose position is displaced from the nominal interaction point by no more than 2 cm in the direction transverse to the beam and no more than 24 cm in the direction along the beam. Furthermore, to reject events produced by the interaction of beam-related protons with the LHC collimators, the fraction of tracks classified as “high purity”, as defined in Reference [5], must exceed 25% in any event with at least 10 tracks. These are standard CMS event cleaning cuts.

The selection of muon candidates from displaced secondary vertices begins by searching for high-purity tracks with transverse momenta $p_T > 33$ GeV/c. Each track must have at least 6 hits, of which at least two must be 3D hits. These criteria are slightly higher than the corresponding trigger thresholds, to minimise dependence on the trigger inefficiency in the p_T turn-on region. The tracks are required to have pseudorapidity $|\eta| < 2$, as the efficiency for finding tracks from displaced secondary vertices falls off at large $|\eta|$. To reject promptly produced particles, the tracks must have a transverse impact parameter significance with respect to the beam line of $|d_0/\sigma_d| > 2$ in the muon channel. Tracks are considered to be identified as leptons if $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$ is less than 0.1. Here, $\Delta\eta$ and $\Delta\phi$ are the differences between the track and a lepton trigger object in pseudorapidity and ϕ , the azimuthal angle about the anticlockwise-beam direction. Standard CMS offline muon identification algorithms are not applied, since they are inefficient for leptons from highly displaced vertices. The muon candidates reconstructed by the offline muon identification algorithms were found to be biased towards the beam spot. However these algorithms are not needed to suppress the very low backgrounds present in this analysis.

The X boson candidates are formed from pairs of oppositely-charged lepton candidates. The two corresponding tracks are fitted to a common vertex, which is required to have a chisquare per degree of freedom $\chi^2/\text{dof} < 5$. For events in the muon channel, the vertex must lie at a distance of more than 5 standard deviations from the

primary vertex in the transverse plane. (Here, the beam line is assumed to be parallel to the z -axis and to pass through the principal primary vertex in the event). If either track has more than one hit closer to the centre of CMS than their common vertex, the event is rejected.

Both muon candidates are required to be isolated, to reject background from jets. A hollow isolation cone of radius $0.03 < \Delta R < 0.3$ is constructed around each candidate. Within this isolation cone, the $\sum p_T$ of all tracks with $p_T > 1$ GeV/ c , excluding the other muon candidate, must be less than 4 GeV/ c . This requirement has very little effect on the signal efficiency, which is relatively insensitive to the number of primary vertices in each event. According to simulation, the mean $\sum p_T$ in the isolation cone increases from 0.6 to 1.2 GeV/ c as the number of additional primary vertices per event increases from 0 to 20.

Figure 6.2 shows the $\sum p_T$ distribution in the isolation cone for simulated long-lived particles. This shows that the chosen cut will give good signal efficiency. One would expect $\sum p_T$ to increase with the number of pileup interactions per event. This dependence is indeed seen in simulated signal events, as shown in Figure 6.3, but is small. Figure 6.4 shows that the $\sum p_T$ distribution agrees well between data and simulated background.

Cosmic ray muons may be reconstructed as back-to-back tracks. To reject them, a requirement of $\cos(\alpha) > -0.95$ is applied, where α is the opening angle between the two tracks. Background from misidentified leptons is reduced by requiring that the two muon candidates are not both matched to the same trigger object. Additional background rejection is achieved by requiring that, projected into the plane perpendicular to the beam line, the reconstructed momentum vector of the X boson candidate is collinear with the vector from the primary vertex to the secondary vertex. The collinearity angle is required to be less than 0.2 radians in the muon channel. Owing to the difficulty of modelling the trigger efficiency for closely spaced muon

pairs, the two tracks in muon channel candidates must be separated by $\Delta R > 0.2$. To eliminate background from J/ψ and Υ decays and from γ conversions, X boson candidates are required to have dimuon invariant masses larger than $15 \text{ GeV}/c^2$. If more than one X boson candidate is identified in a given event, all the selected candidates are retained.




Figure 6.2: Distribution of $\sum p_T$ in an isolation cone around muons from simulated long-lived particle decay. The selection requires this to be less than $4 \text{ GeV}/c$.

Figure 6.5 shows how the efficiency to select a given $X \rightarrow \mu^+ \mu^-$ decay varies as a function of the decay length measured transverse to the beam axis. It illustrates this for the case $M_{H^0} = 1000 \text{ GeV}/c^2$, $M_X = 150 \text{ GeV}/c^2$. It also compares this efficiency in events with only one and events with two X bosons decaying to the desired lepton species. The two results are very similar, suggesting that the efficiency to reconstruct an X boson does not depend strongly upon whether there is a second one in the same event. (One possible way in which such a dependence may arise, is because the double muon trigger is inefficient for pairs of closely spaced muons, as explained in Section 8.5. As a result it may fail to trigger on a single low mass X boson decaying to

CMS Simulation




Figure 6.3: The mean $\sum p_T$ in an isolation cone around muons, as a function of the number of pileup interactions. The error bars show the RMS spread in $\sum p_T$ at each point. Tracks from pileup interactions increases the average momentum within the isolation cone.




Figure 6.4: Comparison of isolation cone momentum sum in data and simulation background events, for dimuon candidates. The candidates shown in these plots pass all selection criteria except isolation and lifetime-related cuts.

dimuons, but should trigger if the event contains a second low mass X boson decaying to dimuons, since it can use one muon from each X boson).

Figures 6.6 to 6.11 show the distributions of the principal variables that are used to select the signal. These figures compare the distributions seen in data with those in the simulated background. Superimposed on these figures is also the expected distribution from the signal process $H^0 \rightarrow 2X$, $X \rightarrow \mu^+ \mu^-$, with $M_{H^0} = 1000 \text{ GeV}/c^2$, $M_X = 350 \text{ GeV}/c^2$, illustrated for a cross section of 1 pb. When plotting these distributions, all cuts are applied except for the cut on the variable being plotted.

Figure 6.12 shows the reconstructed dimuon mass after all selection cuts. Zero candidates are seen in the dimuon channel.

Figure 6.13 shows the reconstructed transverse decay length of the selected candidates.




Figure 6.5: The efficiency to select $X \rightarrow \mu^+\mu^-$ decay as a function of transverse decay length for dimuon candidates, shown for the case $M_{H^0} = 1000 \text{ GeV}/c^2$, $M_X = 150 \text{ GeV}/c^2$. The left (right) plot shows the efficiencies for events in which one (two) X bosons decay to dimuons.

By removing the lifetime-related cuts (lepton d_0/σ , transverse decay length significance, and $\Delta\varphi$), one can obtain a control sample, which is dominated by promptly produced muon pairs. The dimuon mass spectrum obtained with this selection is shown in Figure 6.14. Good agreement is seen in both shape and normalization between data and Monte Carlo simulation.




Figure 6.6: The dimuon vertex χ^2/NDF distribution. This is required to be less than 5. The open histogram represents an example long-lived signal with a product of cross-section and branching fraction of 1 pb.

6.3.1 Selection efficiency

Tables 6.1 show the candidate selection efficiency of the individual cuts for the muon channel. Reasonable agreement between data and the background simulation is seen (compared with the assumed systematic uncertainties, discussed in Chapter 8). The discrepancies towards the bottom of the table are statistical fluctuations caused by the very small number of data events passing all cuts.




Figure 6.7: The φ difference between the dimuon momentum vector and the vector between the primary and secondary vertices. This is required to be less than 0.2.




Figure 6.8: The number of tracker hits assigned in total, to the two muon candidates, which lie in front of (i.e. closer to the centre of CMS) the reconstructed vertex position. It is required to be no more than 1.




Figure 6.9: The cosine of the angle between the two muons that form a candidate. It must exceed -0.95 , so rejecting back-to-back tracks. This particular plot shows the distribution after all other cuts have been implied; in this case, no further events are rejected by this cut.




Figure 6.10: The transverse decay length significance of the candidates. It is required to be more than 5 standard deviations for dimuon candidates.




Figure 6.11: The minimum transverse impact parameter significance of the two leptons in the candidates for the dimuon channels. It is required to be more than 2 standard deviations.




Figure 6.12: The reconstructed dimuon mass after all selection cuts have been applied. The background in the dimuon channel is extremely small.




Figure 6.13: The reconstructed transverse decay length of the candidates in the dimuon channel, after all selection cuts have been applied.

Table 6.1: Table of selection cut efficiencies for candidates in the muon channel. Each cut efficiency is derived using candidates that pass all previous cuts. The efficiency of the first cut is determined relative a sample of events where an analysis trigger fired, and two tracks each with $p_T > 33$ GeV/c were present. The signal efficiency is illustrated for the case $M_{H^0} = 1000$ GeV/c², $M_X = 350$ GeV/c².

selection cut	dimuon cut efficiency		
	data	background MC	signal MC (1000/350)
primary vertex filter	99.8%	99.8%	99.9%
lepton $ \eta $	90.0%	90.6%	95.8%
lepton d_0/σ	2.4%	2.6%	99.4%
lepton track isolation	39.2%	47.2%	90.3%
opposite charge	98.4%	97.9%	99.7%
lepton back-to-back veto	89.7%	92.0%	97.4%
lepton ΔR	94.6%	96.6%	99.9%
lepton trigger match	97.2%	98.3%	97.0%
vertex χ^2	98.1%	99.2%	97.5%
vertex $\Delta\varphi$	5.8%	6.7%	98.3%
max hits before vertex	100%	100%	96.6%
vertex L/σ	0.0%	0.0%	99.9%
dilepton mass	100.0%	100%	100%




Figure 6.14: The invariant mass distribution of dimuon candidates after applying all selection cuts except for the lifetime-related cuts. This predominantly selects prompt background such as Z bosons. The agreement of both shape and normalization between data and Monte Carlo simulation demonstrates a good understanding of the Standard Model backgrounds.

The efficiency of the full set of selection criteria for $X \rightarrow \mu^+ \mu^-$ decays is defined on a per-candidate basis, not on a per-event basis. It is determined from the simulation for the muon channel, for two different classes of events: first for $H^0 \rightarrow XX$ events in which only one long-lived exotic particle decays to the chosen lepton species, defining efficiency ϵ_1 , and second for events in which both long-lived exotic particles decay to chosen lepton pairs, defining efficiency ϵ_2 . The efficiencies ϵ_1 and ϵ_2 are usually almost identical, indicating that the efficiency to select an X boson candidate is not strongly affected by whether or not the second X boson in the event decays to the same lepton channel. The only exception is in the case of small M_X/M_{H^0} , where the dimuon trigger is inefficient for the two nearly collinear muons from the decay of the same X boson, but the trigger requirement can still be satisfied by muons from separate X bosons. The efficiencies are estimated for a range of X boson lifetimes, corresponding to mean transverse decay lengths of $\approx 0.7 - 600$ cm, by re-weighting the simulated signal events. A subset of these efficiencies are shown in Table 6.2. The maximum efficiency (for $M_{H^0} = 1000$ GeV/ c^2 , $M_X = 150$ GeV/ c^2 , $c\tau = 1$ cm) is approximately 52% in the muon channel, but becomes significantly smaller at lower H^0 masses or longer lifetimes. For extremely long lifetimes, such that most X bosons decay too far from the primary vertex to be reconstructed, the efficiency will scale in inverse proportion to the proper lifetime τ of the X boson.

Table 6.2: The efficiency of $X \rightarrow \ell^+ \ell^-$ decays to pass all selection criteria, both for $H^0 \rightarrow XX$ events in which only one long-lived exotic decays to pairs of muons ϵ_1 and for the case where both decay to pairs of muon ϵ_2 . The uncertainties on these efficiencies are not shown in the table. They are dominated by the 20% relative uncertainty related to the tracking performance.

M_{H^0} (GeV/ c^2)	M_X (GeV/ c^2)	$c\tau$ (cm)	Dimuon channel	
			ϵ_1	ϵ_2
1000	350	3.5	0.56	0.58
		35.0	0.25	0.27
		350.0	0.019	0.021
1000	150	1.0	0.52	0.50
		10.0	0.35	0.37
		100.0	0.035	0.035
1000	50	0.4	0.35	0.46
		4.0	0.25	0.33
		40.0	0.041	0.034
1000	20	0.2	0.010	0.021
		1.5	0.012	0.020
		15.0	0.003	0.004
400	150	4.0	0.39	0.49
		40.0	0.18	0.20
		400.0	0.015	0.022
400	50	0.8	0.36	0.41
		8.0	0.24	0.26
		80.0	0.026	0.022
400	20	0.4	0.18	0.30
		4.0	0.14	0.19
		40.0	0.016	0.019
200	50	2.0	0.095	0.13
		20.0	0.055	0.065
		200.0	0.006	0.004
200	20	0.7	0.12	0.15
		7.0	0.069	0.083
		70.0	0.007	0.008
125	50	5.0	0.011	0.011
		50.0	0.003	0.004
		500.0	0.000	0.000
125	20	1.3	0.017	0.020
		13.0	0.008	0.010
		130.0	0.001	0.002

Chapter 7

Background Estimation and Modeling

An interpretation of the observed dilepton mass spectrum requires an estimate of the background normalization and a parametrization of the background shape as a function of M_X . However, as discussed in Chapter 9, the results presented have very little dependence on these quantities.

7.1 Background Normalization

The number of background events passing all the selection criteria for X boson candidates is estimated from simulated samples using the distribution of the transverse decay length significance L_{xy}/σ_{xy} . We parameterize this distribution with the sum of two falling exponentials. By integrating the fitted curve over the signal region, defined by $L_{xy}/\sigma_{xy} > 5$ for the muon channel, an estimate of the mean total background in the mass spectrum is obtained. The estimate gives $0.02_{-0.02}^{+0.09}$ candidates. This estimate of the mean total background is used to derive the results in Chapter 9. To verify that the simulation correctly describes the L_{xy}/σ_{xy} spectrum, data and simulation distributions are compared in Figure 7.1, after relaxing the lifetime-

related selection requirements to increase the number of events and ensure that the plots are dominated by background. Specifically, the threshold on the decay length significance L_{xy}/σ_{xy} was decreased to 3 and the requirement on the individual muon d_0/σ_d was removed entirely. Figure 7.1 also shows that the main background to this search consists of prompt dimuons that have been reconstructed with large decay length significance.

As a check, one can perform the fit to the L_{xy}/σ distribution only in the background region $L_{xy}/\sigma < 5$. This yields a larger statistical uncertainty, but since it would be unaffected by any signal, it can be done using data too. The result obtained from simulated events is $0.02_{-0.02}^{+0.09}$ candidates in the muon channel, whilst from data one obtains $0.01_{-0.01}^{+0.01}$ candidates. Alternatively, one can simply count the number of simulated background events passing the selection requirements, which is $0.01_{-0.01}^{+3.65}$ events, as seen in Figure 6.12. The statistical uncertainties on these numbers of simulated background events are derived using the calculation described in Section 7.3. While this is the simplest way of deriving a background estimate, it also requires a degree of confidence in the ability of the simulation to model the tails of the distribution in which our signal lies, so fitting to the decay length significance distribution may provide a more robust estimate. The corresponding number of candidates seen in data (which must be greater than or equal to the number of background candidates, depending on whether any signal is present) is 0 events in the muon channel. This is compatible with the simulated result, so proving that any systematic uncertainties on the latter do not lead to a statistically significant underestimate of the background level.

Table 7.1 summarizes these methods of estimating the background level, which are consistent within (large) statistical uncertainties. It is important to realize that these are estimates of the total number of background candidates in the background region. The probable background under any signal mass peak should be approximately two

orders of magnitude smaller (depending on the shape and mass resolution). Because this background is so small, the limits are very insensitive to the the background shape.




Figure 7.1: The transverse decay length significance of the candidates for the dimuon channel with loosened cuts in data and simulation. The vertical dashed line indicates the selection requirement used for signal events. There are no simulated QCD or $t\bar{t}$ events passing these selection requirements, so they are omitted.

Background estimation technique	$\mu\mu$ estimate
Simulation: fit to full region	$0.02^{+0.09}_{-0.02}$
Simulation: fit to background region only	$0.02^{+0.09}_{-0.02}$
Data: fit to background region only	$0.01^{+0.01}_{-0.01}$
Count of simulated events passing selection	$0.01^{+3.65}_{-0.01}$
Count of data events passing selection	0

Table 7.1: Comparison of various methods of estimating the background normalisation. The number of data events is also included for reference.

7.2 Background shape

An estimate of the background shape can be obtained from the M_X distribution of a background sample. However, after applying all selection requirements there are too few events to measure its shape accurately. Since the dimuon mass distribution and lifetime-related variables are only weakly correlated in simulated background candidates, the shape of the mass distribution is instead obtained by fitting a parameterized function to data samples with the lifetime-related selection requirements removed. Namely, no selections are made on the individual muon d_0/σ_d , the transverse decay length significance L_{xy}/σ_{xy} , or the collinearity angle $\Delta\varphi$. (This is the prompt sample shown in Figure 6.14.) Figure 7.2 shows the results of these fits to the muon data sample. (This sample should be background dominated, so can be fit to data without worrying that any signal present might distort the spectrum). The background is represented by the sum of two functions: a Breit–Wigner function, to represent the Z resonance, multiplied by a Gaussian error function, to approximate the effect on the selection efficiency of the lepton p_T thresholds; and an additional function, given by an exponential term, to represent the non-Z background. The fits give the fraction of the background from the Z as 0.994 ± 0.001 for the muon channels.

7.3 Statistical uncertainty on a priori background from counting method

Section 7.1 described how an a priori estimate of the background normalisation can be obtained by simply counting the number of candidates N that pass the selection cuts in simulated background events. This section describes how the statistical uncertainty on this estimate is calculated.




Figure 7.2: Distribution of the dilepton mass and the fitted shape in a data sample with lifetime-related selection requirements removed, shown for the muon channel. The shape used is that of a Breit–Wigner distribution times a turn-on function, added to an exponential term.




Figure 7.3: Fit of the background shape to simulated background with lifetime-related selection requirements removed, shown for the muon channels. The function used is the same as that for Figure 7.2.




Figure 7.4: Alternative fits of the background to a sample with loosened versions of the signal selection requirements, for simulated events (left) and data (right), for the muon channel.

The candidates in the Monte Carlo events have weights used to adjust the number of events to the correct cross section, the amount of pile-up etc. If one expects a mean of μ_i candidates with weight w_i , then the mean expected total number of weighted candidates μ is:

$$\mu = \sum_i \mu_i w_i$$

The RMS difference between this and the observed number of weighted candidates N is given by the statistical uncertainty σ_N :

$$\sigma_N^2 = \sum_i \mu_i w_i^2$$

Since the μ_i are unknown, often an estimate for σ_N is performed by using the observed number of candidates N_i with weight w_i in place of the expected number μ_i in this equation. However, if μ_i is a small number, as is the case here, then this is a poor approximation. e.g. If one expected a mean of 1.0 candidates with a weight of 100, then these candidates should contribute ± 100 to the statistical uncertainty. But if one observes zero candidates of this weight, then one making the above approximation would infer that they contributed nothing to the statistical uncertainty. One would therefore seriously underestimate σ_N . The uncertainties on the Monte Carlo distributions shown in the figures in Chapter 6 are derived in this conventional way. In consequence, these estimated uncertainties are probably too small, and exaggerate the size of any discrepancy between data and simulation.

As explained in Section 7.1, there exist additional samples of candidates, selected with loose lifetime-related cuts or no lifetime-related cuts at all. Such a loose sample has N_i^L candidates of weight w_i . Suppose that the efficiency of background candidates to pass the normal cuts with respect to a looser set is ϵ . Hypothesise that this

efficiency is the same for candidates of any weight. This assumption corresponds, above all, to assuming that different MC samples have similar efficiencies to pass the lifetime-related cuts. Then, a good estimate of the mean number of candidates expected with normal cuts w_i is:

$$\mu_i \approx N_i^L \epsilon$$

Hence the expected mean number of candidates passing the normal cuts μ and the statistical uncertainty on the number of candidates passing the normal cuts σ_N are given by:

$$\mu \approx \sum_i (N_i^L \epsilon) w_i = \epsilon N^L$$

$$\sigma_N^2 \approx \sum_i (N_i^L \epsilon) w_i^2 = \epsilon \sigma_{N^L}^2$$

where one defines $N^L = \sum_i N_i^L w_i$ and $\sigma_{N^L}^2 = \sum_i N_i^L w_i^2$.

The most likely value of ϵ is given by:

$$N = \epsilon N^L \tag{7.1}$$

The maximum and minimum values of ϵ are consistent at one standard deviation with the observed number of candidates N passing the normal cuts are given by:

$$N = \epsilon_{\pm} N^L \mp \sqrt{\epsilon_{\pm}} \sigma_{N^L}$$

which is:

$$\epsilon_{\pm} N^L = N \left[1 + \frac{\sigma_{N^L}^2}{2NN^L} \right] \pm \sigma_{N^L} \sqrt{\frac{N}{N^L}} \sqrt{1 + \frac{\sigma_{N^L}^2}{4NN^L}} \quad (7.2)$$

Equations 7.1 and 7.2 are respectively, estimates of the number candidates passing the normal cuts and the one standard deviation upper/lower limits on that number. (Note that the mean of the upper and lower limits is larger than the central value, as expected for small statistics). The differences between the upper/lower limits and the estimated number of candidates give the positive and negative uncertainties. These are shown in Table 7.2. Results are given for the loose set of lifetime-related cuts, described in Section 7.1, an alternative set of loose cuts, and for a set with all lifetime-related cuts removed. These are all reasonably consistent with each other, supporting the assumption made above that the efficiency for background of the normal cuts relative to the loose ones is approximately the same for candidates of any weight. The only exception is in the dielectron channel, for the ‘no lifetime cuts’ case, where the estimated uncertainty is larger. Investigation shows that this is due to a single candidate with a very large weight. If this is removed, the estimated uncertainty in the ‘no lifetime cuts’ case becomes comparable with that obtained using the loose cuts. This indicates that the statistical uncertainty on the statistical uncertainty in the dielectron channel is large. Also shown in the table is the effect of using the normal cuts in the place of the loose cuts (i.e. Setting $N^L = N$ and $\sigma_{N^L} = \sigma_N$ in the equations). The estimated uncertainties in this case are much smaller, (and presumably very inaccurate), which is not surprising as the expected number of candidates for any given weight is often less than one.

Table 7.2: Positive and negative one standard deviation statistical uncertainties on the apriori background estimated by counting simulated background events passing the normal selection cuts. The uncertainties are derived making use of the given ‘loose sample’ of candidates. The numbers shown in brackets indicate the estimated uncertainty if one removes the candidate with the largest weight.

Selection cuts	Dimuon Candidates	
	+ve uncertainty	-ve uncertainty
Normal	0.007	0.003
Loose	3.25	0.005
Alternative Loose	3.65	0.005
No Lifetime Cuts	4.04 (4.03)	0.005

Chapter 8

Systematic Uncertainties

The primary systematic uncertainty comes from the efficiency in detecting and reconstructing signal events. This uncertainty derives from uncertainties in the efficiency of reconstructing tracks from displaced vertices, the trigger efficiency, the modelling of pileup in the simulation, the parton distribution function sets, the renormalisation and factorisation scales used in generating simulated events, and the effect of higher order QCD corrections. In addition, systematic uncertainties in the integrated luminosity, and the background estimate are considered.

Table 8.1 summarises the sources of systematic uncertainty affecting the signal efficiency. The relative uncertainty in the luminosity is taken to be 2.2% [39].

Table 8.1: Systematic uncertainties affecting the signal efficiency over the range of M_{H^0} and M_X values considered. In all cases, the uncertainty specified is a relative uncertainty. Note that the NLO uncertainty is only evaluated for the $M_{H^0} = 125 \text{ GeV}/c^2$ case. The relative uncertainty in the luminosity is taken to be 2.2%.

Source	Uncertainty
Pileup modelling	2%
Parton distribution functions	< 1%
Renormalisation and factorisation scales	< 0.5%
Tracking efficiency	20%
Trigger efficiency	2.6% (e), 11% (μ)
NLO effects ($M_{H^0} = 125 \text{ GeV}/c^2$ only)	4-12%

8.1 Luminosity

For the running period corresponding to this analysis, CMS estimates the relative uncertainty on the luminosity to be 2.2% [3]. We use this uncertainty in calculating the final cross-section estimates.

8.2 Effect of pileup

The likelihood of a given number of pileup events occurring in the data can be calculated from the distribution of the instantaneous luminosity during the 2011 run. The number of true pileup events in the simulation is also known. The simulation can therefore be reweighed to match the data. This is done using the procedure in [40]. Figure 8.1 compares the number of reconstructed primary vertices per event in data and in simulation following this reweighing.

The systematic uncertainty in this procedure is estimated by adjusting the reweighting so as to vary the mean number of primary vertices in the simulation by $\pm 5\%$, following the recommendation of [41]. The results of the reweighing including the systematic variations are also shown in Figure 8.1. This uncertainty gives rise to a relative systematic uncertainty in the signal selection efficiency of less than 2% for all mass points.

8.3 PDF, renormalisation and factorisation scale uncertainties

A generator-level study is performed, in which the fraction of X bosons is determined, whose dilepton decay products both pass the kinematic cuts on p_T and rapidity described in Section 6.3. This fraction a , defined to be the *acceptance*, is presented in Table 8.2 for each of the simulated signal samples. It is apparent that with the

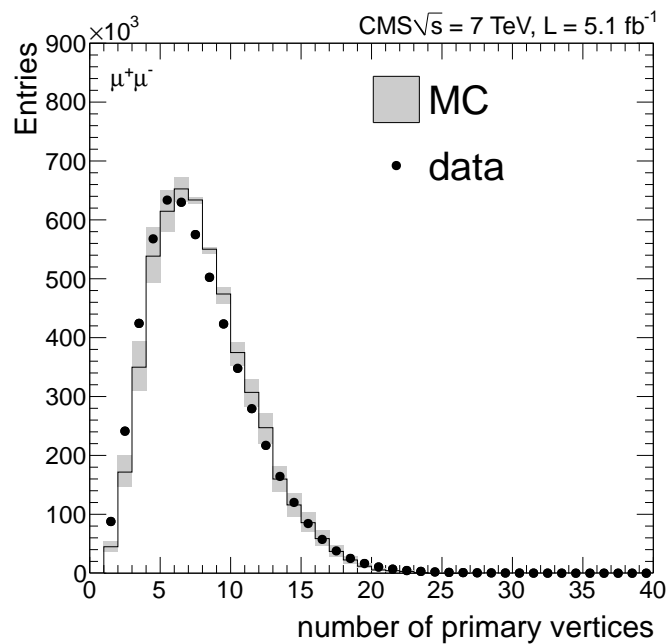


Figure 8.1: The number of reconstructed primary vertices per event in data and simulation, following the pileup reweighting procedure for the muons sample. The grey band in the simulation histograms shows the variation obtained following the procedure to estimate pileup systematics.

current kinematic cuts, which are determined by the trigger thresholds, the sensitivity is good only for the higher H^0 masses.

All simulated samples were generated using the CTEQ6L1 [42] PDF set. Systematic uncertainties on the acceptances due to PDF variations are evaluated using the PDF4LHC procedure [43]. These acceptance uncertainties are also shown in Table 8.2.

Table 8.2: The fraction of X bosons, whose dimuon decay products both pass the kinematic cuts on the muon p_T and pseudorapidity. The uncertainties quoted on these numbers correspond to the systematic uncertainty on the PDF set.

M_{H^0} (GeV/ c^2)	M_X (GeV/ c^2)	Dimuon channel
1000	350	0.866 ± 0.005
1000	150	0.760 ± 0.004
1000	50	0.764 ± 0.003
1000	20	0.762 ± 0.003
400	150	0.673 ± 0.008
400	50	0.501 ± 0.004
400	20	0.509 ± 0.004
200	50	0.179 ± 0.002
200	20	0.197 ± 0.002
125	50	0.02930 ± 0.00008
125	20	0.0397 ± 0.0003

By default, for the Higgs production with the PYTHIA Monte Carlo generator, the QCD renormalisation μ_r and factorisation μ_f scales are chosen to be equal $\mu_r = \mu_f = \mu = M_{H^0}$. The dependence of the acceptance on the choice of scales was estimated by varying μ by a factors of 0.5 and 2. It was found to be well below 0.5%, whilst the production cross section changed by typically 10%. One would expect the factorisation scale to affect the acceptance more than the renormalisation scale, since the former modifies the resolved parton densities and the final leptons boost, whereas the latter impacts significantly only the Initial State Radiation.

8.4 Track-finding efficiency

8.4.1 Tracking efficiency from cosmic muons

Three methods have been explored to understand if the efficiency to reconstruct displaced tracks is correctly modelled by the CMS detector simulation. The first exploits cosmic rays and checks the efficiency to reconstruct isolated particles. The second embeds individual simulated displaced tracks in real data events and so determines the reconstruction efficiency in a high-occupancy environment. The third, which merely provides a cross-check, uses K_s^0 decays. Since the physics analysis pursued here searches for isolated leptons, the first of these three methods is considered to be more pertinent. It is therefore used to measure the systematic uncertainty on the tracking efficiency, which is used as input to the physics analysis. The other two methods merely provide a check.

Cosmic muons provide an abundant source of very displaced tracks. Since the alignment of the tracker relies heavily on them, they are collected frequently, both during proton-proton collisions and when there is no beam in the LHC. By reconstructing the cosmic rays in the muon chambers and measuring the fraction of the events in which associated tracks are found in the Tracker, one can estimate the track finding efficiency for isolated tracks.

The timing of cosmic muons can be an issue, as the readout of the tracker (which can only occur at multiples of 25 ns) can not be perfectly synchronized with the cosmic rays. When there is no beam, the silicon strip tracker is therefore read out in a special data taking mode, known as *peak-mode* [5], in which its time resolution is significantly worse than 25 ns. (It is not expected that changing to ‘peak mode’ should have any significant effect on the tracking efficiency, which, for isolated tracks, is mainly influenced by the fraction of bad silicon modules).

This study uses events from good cosmic runs taken during 2011 without LHC beam and reconstructed with CMSSW 4.2. The main trigger used in cosmic runs is the `L1_SingleMuOpen`. Only events where the tracker high voltage is on are used. The data are compared with a dedicated simulation sample of simulated cosmic rays. It should be noted that the latter only includes cosmic rays traversing the tracker (within 80 cm of the central z -axis of CMS), so differs from the data.

The cosmic ray muons are reconstructed using the muon chambers alone. Those cosmic rays which traverse the tracker will give rise to two ‘stand-alone’ muons in the top and bottom halves of CMS. These two ‘stand-alone’ muons are combined into a single reconstructed track crossing the entire detector, by the dedicated ‘cosmic-Muons1Leg’ algorithm [44].

Tracker tracks are reconstructed using the standard algorithm used in pp collisions. They are matched in a ΔR cone to the reconstructed cosmic muon track. For each cosmic ray muon, two tracker tracks are thus searched for.

Figures 8.2-8.4 show the properties of the cosmic muons reconstructed in the muon chambers, comparing data and simulation. Some differences can be seen, as a result of the cuts used in producing the simulated sample. The asymmetric pseudorapidity distribution arises from uneven distribution mass above the CMS experiment (notably the presence of the shaft to the surface at one end). The poor χ^2 distribution in the simulated events is related to inadequacies in the simulation of the cosmic muon arrival time, and is not thought to influence the tracking performance. Here and in all that follows in this subsection, all impact parameters are calculated with respect to the centre of CMS.

In order to select a sample of cosmic muons that are well reconstructed in the muon chambers, the following cuts are used. The cuts on impact parameter improve the agreement between data and simulation by rejecting large impact parameter cosmics, which were not generated in the simulation.




Figure 8.2: Cosmic muons reconstructed in the muon chambers alone. The p_T (a), η (b) and ϕ (c) distributions are shown. Data is in black, simulation in red.

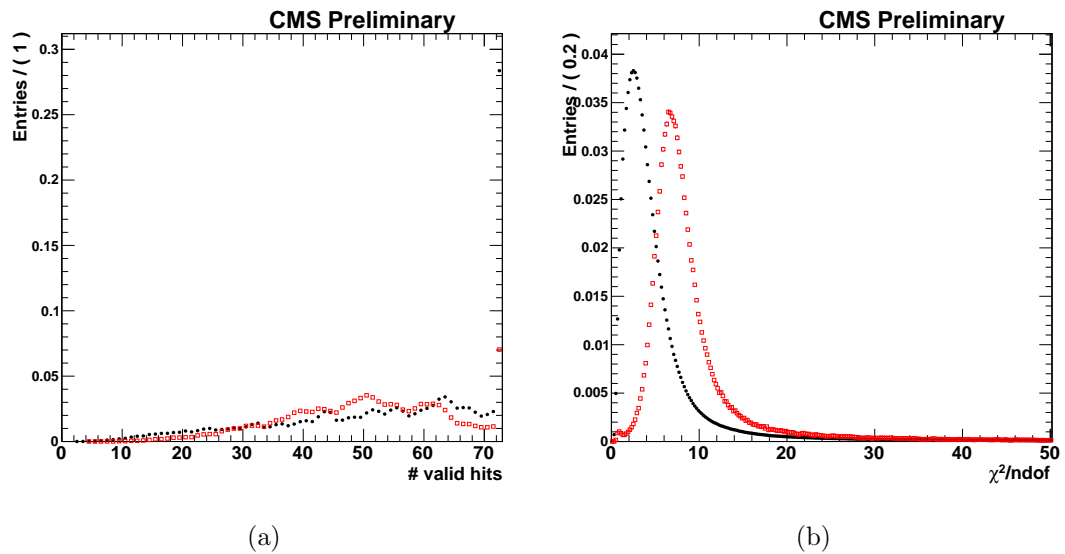


Figure 8.3: Cosmic muons reconstructed in the muon chambers alone. The number of valid hits (a) and χ^2/ndof (b) distributions are shown. Data is in black, simulation in red.

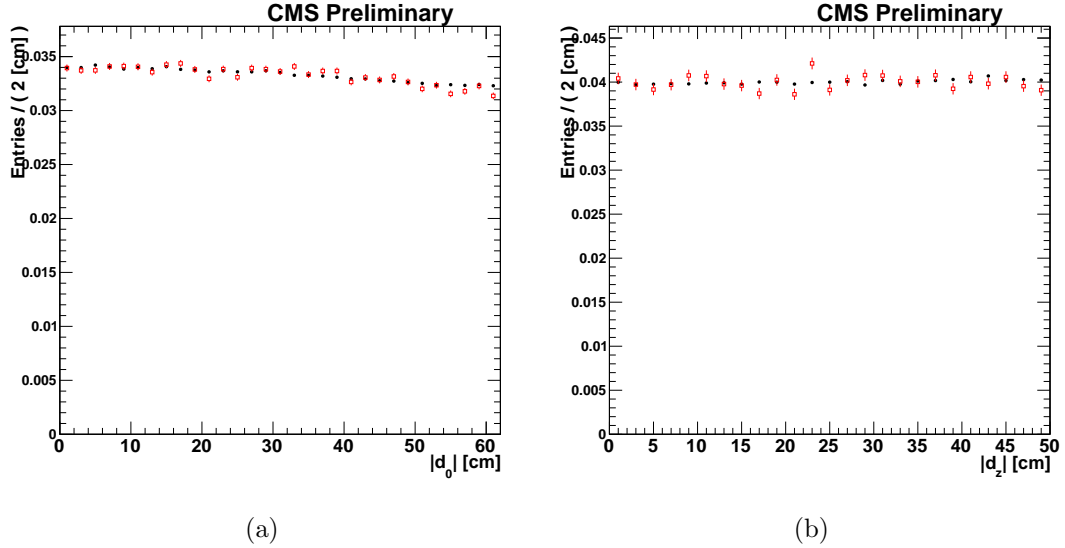


Figure 8.4: Cosmic muons reconstructed in the muon chambers alone. The $|d_0|$ (a) and $|d_z|$ (b) distributions are shown. Data is in black, simulation in red.

- Exactly one cosmic muon in the event.
- $p_T > 35$ GeV/c
- $|\eta| < 2$
- Transverse impact parameter $|d_0| < 50$ cm.

In addition, to ensure that the cosmic muon is well measured, the estimated uncertainties on the transverse (σ_{d_0}) and longitudinal (σ_{z_0}) impact parameters are both required to be less than 1 cm. They must both also be less than $f(\min(p_T, 200))$, where $f(p_T)$ is a third order polynomial fitted on the σ_{d_0} distribution and shifted by 0.1 along the positive direction of the y -axis. Its purpose is to reject outliers from the main distribution.

The dependences of the impact parameter uncertainties on the muon p_T are shown in Figure 8.5, together with a curve indicating the cut placed on these uncertainties.

After applying the cosmic muon selection cuts, the agreement between data and simulation is improved as shown in Figures 8.6 and 8.8.

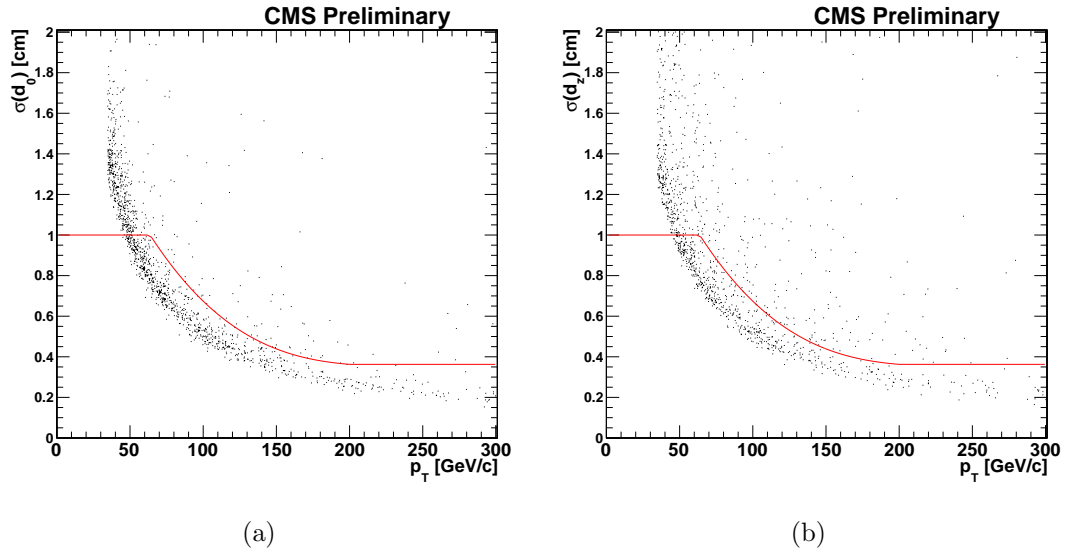


Figure 8.5: Distributions of the uncertainties on the transverse (a) and longitudinal (b) impact parameters for selected cosmic muons reconstructed in the muon chambers alone. The curves show the cuts applied on these uncertainties.




Figure 8.6: Selected cosmic muons reconstructed in the muon chambers alone. The p_T (a), η (b) and ϕ (c) distributions are shown. Data is in black, simulation in red.

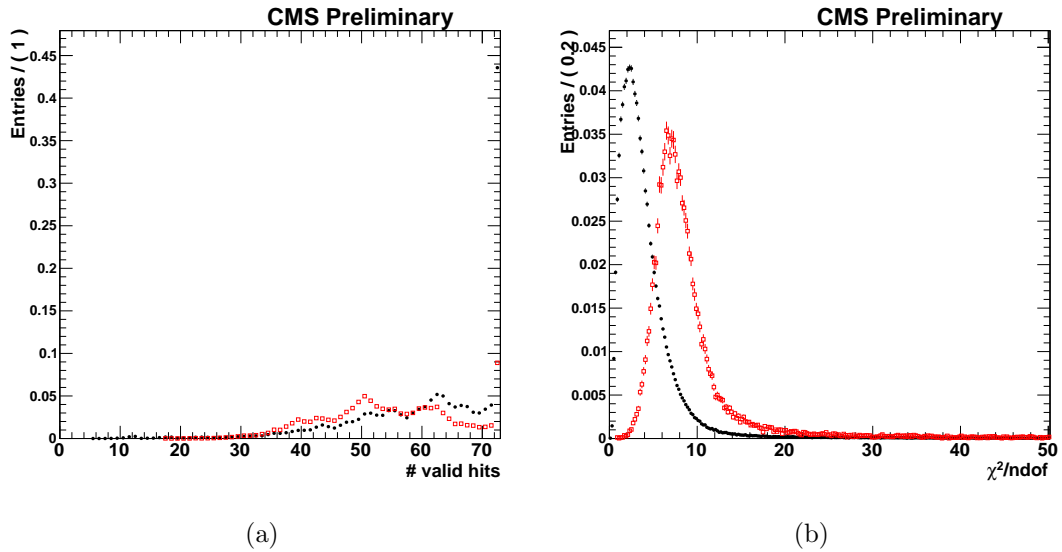


Figure 8.7: Selected cosmic muons reconstructed in the muon chambers alone. The number of valid hits (a) and χ^2/ndof (b) distributions are shown. Data is in black, simulation in red.




Figure 8.8: Selected cosmic muons reconstructed in the muon chambers alone. The $|d_0|$ (a) and $|z_0|$ (b) distributions are shown. Data is in black, simulation in red.

A matching Tracker track is searched in a cone of radius $\Delta R < 1$ (allowing for the fact that reconstructed track may go in the opposite direction to the cosmic). This wide cut is sufficient to identify the track as shown in Figure 8.9.

Tracker tracks are selected with the same cuts used for the physics analysis, described in Section 6.3:

- $p_T > 25 \text{ GeV}/c$
- $|\eta| < 2$
- ‘high purity’ requirement
- at least 6 valid hits.

The kinematic characteristics of these tracks in data and simulation are shown in Figures 8.10-8.12.




Figure 8.9: ΔR between selected cosmic muons reconstructed in the muon chambers and the closest tracker track. Data is black, simulation red.

The efficiency to reconstruct a tracker track associated to a cosmic muon reconstructed in the muon chambers, as a function of the transverse and longitudinal impact parameters, is shown in Figures 8.13 and Figure 8.14. When making the first




Figure 8.10: Tracker tracks p_T (a), η (b) and ϕ (c) distributions. Data is in black, simulation in red.

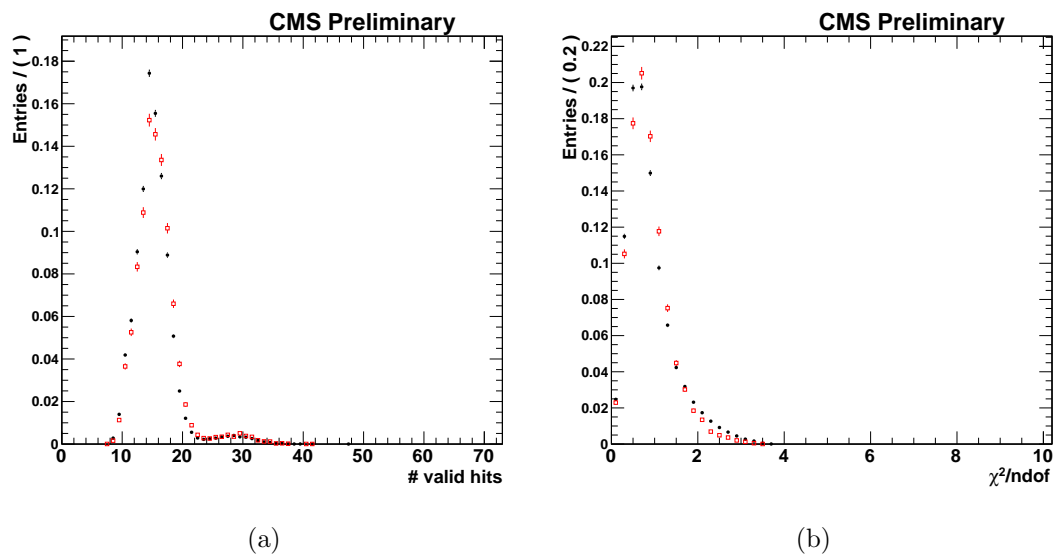


Figure 8.11: Tracker tracks number of valid hits (a) and $\chi^2/ndof$ (b) distributions. Data is in black, simulation in red.

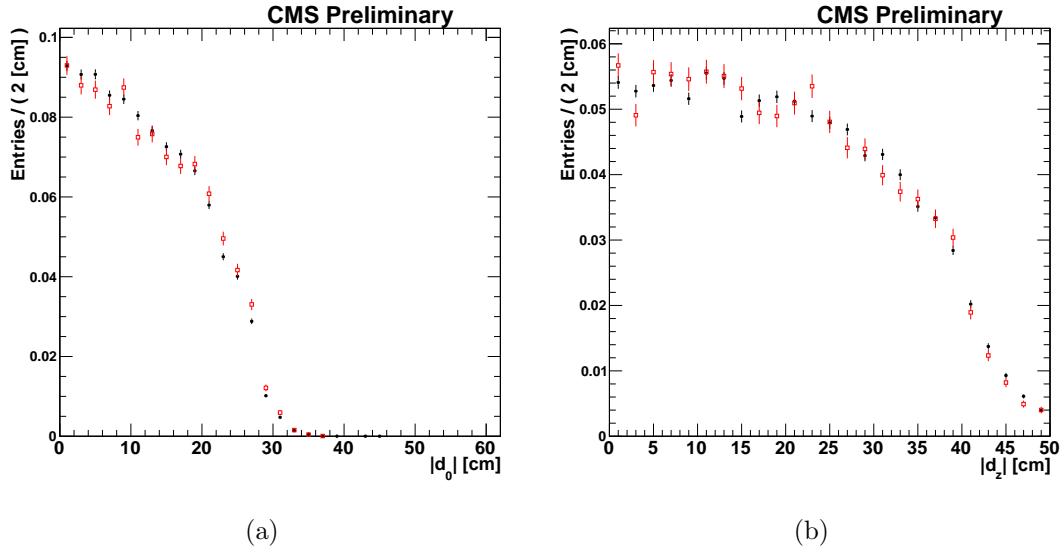


Figure 8.12: Tracker tracks $|d_0|$ (a) and $|d_z|$ (b) distributions. Data is in black, simulation in red.

plot a cut on $|z_0| < 10$ cm was applied, while for the second plot a cut on $|d_0| < 4$ cm is applied. The ratio $\text{efficiency}(\text{data})/\text{efficiency}(\text{simulation})$ is shown in Figure 8.15. From this plot, it is concluded that the tracking efficiency for displaced, isolated leptons is understood to better than 10% relative in the region of interest $|d_0| < 20$ cm. Figure 8.16 shows that very few tracks originating from the simulated signal events are reconstructed with $|d_0| > 20$ cm. The corresponding relative systematic uncertainty on the efficiency to reconstruct the dilepton candidates is 20%.

8.4.2 Tracking efficiency using embedding

The efficiency with which displaced tracks are reconstructed is also estimated by embedding hits from pp collision data into simulated events containing single displaced tracks. The absolute tracking efficiency is defined as the probability with which the simulated track is reconstructed in the presence of realistic event occupancies, given that it was reconstructed in isolation when no additional hits were present.




Figure 8.13: Efficiency of the tracker to find a track given a cosmic ray muon as a function of the transverse impact parameter of the muon. Data is in black, and simulation in red.




Figure 8.14: Efficiency of the tracker to find a track given a cosmic ray muon as a function of the longitudinal impact parameter of the muon when no cut is applied on this variable. Only muons with $|d_0| < 4$ cm are used. Data is in black, and simulation in red.




Figure 8.15: Ratio of the efficiencies of the tracker to find a track given a cosmic ray muon for data/simulation.




Figure 8.16: Transverse impact parameter of reconstructed tracker tracks originating from the decays of a long-lived exotic, shown for the case $M_{H^0} = 400 \text{ GeV}/c^2$, $M_X = 150 \text{ GeV}/c^2$.




Figure 8.17: Efficiency of the tracker to find a track, given a cosmic ray muon reconstructed in the muon chambers, as a function of the transverse (left) and longitudinal (right) impact parameters (with respect to the nominal interaction point of CMS). The efficiency is plotted in bins of 2 cm width. For the left plot, the longitudinal impact parameter $|z_0|$ is required to be less than 10 cm, and for the right plot, the transverse impact parameter $|d_0|$ must be less than 4 cm.

Strictly speaking, since the physics selection used here requires the presence of isolated leptons, the efficiency to reconstruct tracks in a high occupancy event is not very important. So the results presented here are simply a cross check.

The embedding method used here is similar to that developed for a previous study [45]. It allows one to measure the absolute efficiency for reconstructing tracks with any chosen trajectory and orientation with respect to other objects in an event. However, because the tracker alignment differs in data and simulated events, the embedded hits will in general not be reconstructed as tracks.

The simulated events used each consist of single muon of $p_T = 25$ GeV/ c , whose trajectory is randomly distributed in ϕ and η . The muons were generated with their production point distributed in a uniform radial distribution up to 50 cm from the primary vertex, and randomly distributed in ϕ and η , independently of the muon direction.

A track is classified as being *correctly reconstructed* if it passes the same selection criteria used in the physics analysis. In addition, at least 75% of the hits assigned to the track must correspond to the simulated muon. Furthermore, its reconstructed z_0 must be consistent (within 2 cm) with that of the simulated muon.

The track *acceptance* is defined as the fraction of the muons that are ‘correctly reconstructed’ in these original single particle events. This acceptance is studied as a function of the transverse and longitudinal impact parameters and is shown in Tables 8.3 and 8.4. It drops rapidly with increasing impact parameter.

Hits from the pp collision dataset `/HT/Run2011A-05AugReReco-v1` and `/HT/-Run2011A-PromptReco-v6` are then embedded into these single particle simulated events. For comparison, the effect of embedding hits from the simulated pp collision dataset `/QCD_Pt-120to170_TuneZ2_7TeV_pythia6` is also studied.

The efficiency to ‘correctly reconstruct’ the muons after the hits have been embedded is also shown in Tables 8.3 and 8.4. It is defined using only the subset of tracks which were ‘correctly reconstructed’ in the single particle events. It is thus a measure of the drop in tracking efficiency caused by the high occupancy environment.

The reconstruction efficiency defined in this way decreases as the impact parameters increase, but remains above 90%. The difference between data and simulation is less than 2%, which is an indication of the systematic uncertainty on the tracking efficiency deduced with this technique.

Table 8.5 shows the efficiency as a function of the number of reconstructed primary vertices in the event. For tracks within acceptance, the efficiency is observed to be high, with no strong dependence on the amount of pileup when systematic uncertainties are considered.

Table 8.3: Acceptance and efficiency for reconstructing muons with $p_T = 25$ GeV/ c as a function of transverse impact parameter.

$ d_0 $ (cm)	Acceptance	MC+data efficiency	MC+MC efficiency
0-5	0.480 ± 0.002	0.975 ± 0.001	0.973 ± 0.001
5-10	0.365 ± 0.003	0.953 ± 0.002	0.948 ± 0.002
10-15	0.292 ± 0.003	0.938 ± 0.002	0.938 ± 0.003
15-20	0.229 ± 0.003	0.918 ± 0.004	0.913 ± 0.005
20-25	0.170 ± 0.003	0.902 ± 0.005	0.900 ± 0.006
25-30	0.094 ± 0.003	0.887 ± 0.008	0.885 ± 0.010
30-35	0.032 ± 0.002	0.905 ± 0.016	0.906 ± 0.019
35-40	0.006 ± 0.001	0.852 ± 0.068	0.852 ± 0.068

Table 8.4: Acceptance and efficiency for reconstructing muons with $p_T = 25$ GeV/ c as a function of the projected longitudinal impact parameter $d_{sz} = z_0 \sin \theta$.

$ d_{sz} $ (cm)	Acceptance	MC+data efficiency	MC+MC efficiency
0-5	0.542 ± 0.002	0.966 ± 0.001	0.966 ± 0.001
5-10	0.459 ± 0.003	0.955 ± 0.002	0.950 ± 0.002
10-15	0.366 ± 0.003	0.953 ± 0.002	0.944 ± 0.003
15-20	0.291 ± 0.004	0.945 ± 0.003	0.937 ± 0.004
20-25	0.224 ± 0.004	0.934 ± 0.004	0.938 ± 0.005
25-30	0.148 ± 0.004	0.921 ± 0.006	0.921 ± 0.007
30-35	0.102 ± 0.004	0.927 ± 0.008	0.920 ± 0.010
35-40	0.059 ± 0.003	0.930 ± 0.011	0.933 ± 0.013

Table 8.5: Efficiency for reconstructing muons with $p_T = 25 \text{ GeV}/c$ as a function of the number of reconstructed primary vertices.

n_{PV}	MC+data efficiency	MC+MC efficiency
1	0.983 ± 0.004	0.980 ± 0.006
2	0.977 ± 0.003	0.980 ± 0.002
3	0.971 ± 0.002	0.973 ± 0.002
4	0.967 ± 0.002	0.972 ± 0.002
5	0.961 ± 0.002	0.964 ± 0.002
6	0.955 ± 0.002	0.958 ± 0.002
7	0.952 ± 0.002	0.954 ± 0.002
8	0.948 ± 0.003	0.945 ± 0.003
9	0.942 ± 0.003	0.947 ± 0.003
10	0.941 ± 0.004	0.946 ± 0.003
11	0.938 ± 0.005	0.941 ± 0.003
12	0.923 ± 0.007	0.943 ± 0.003
13	0.923 ± 0.009	0.944 ± 0.003
14	0.921 ± 0.012	0.938 ± 0.004
15	0.928 ± 0.015	0.934 ± 0.004
16	0.899 ± 0.028	0.927 ± 0.005
17	0.958 ± 0.024	0.911 ± 0.005
18	0.895 ± 0.050	0.924 ± 0.007

8.4.3 Verification of displaced tracking efficiency using K_s^0

CMS has previously used displaced tracking for a measurement of K_s^0 production [46]. In that paper, it was verified that K_s^0 reconstructed in CMS can be used to measure the K_s^0 lifetime, and that the result is within 1% of its world average value. Although this agreement can not be translated into a measurement of the displaced tracking efficiency, it does provide some additional evidence that its dependence on decay length is well modelled in the simulation. (However, since [46] only studied K_s^0 with $p_T < 8$ GeV/ c and typical transverse decay length less than 5 cm, this check is only sensitive to problems in a limited region of phase space.)

8.5 Trigger efficiency measurement

The trigger efficiency is measured using the *Tag and Probe* method. In this method, one looks for pairs of leptons coming from the decay of a resonance, chosen here to be the Z boson. One of the two leptons (referred to as the *Tag*) is required to pass tight lepton identification selection criteria, including the requirement to match within $\Delta R < 0.5$ with a single-lepton trigger object at a particular trigger threshold.

The other candidate (referred to as the *Probe*) is used to estimate the efficiency that a lepton passing the offline selection cuts of the analysis (Section 6.3), would also match within $\Delta R < 0.5$ with one of the dilepton trigger objects used to select events in this analysis.

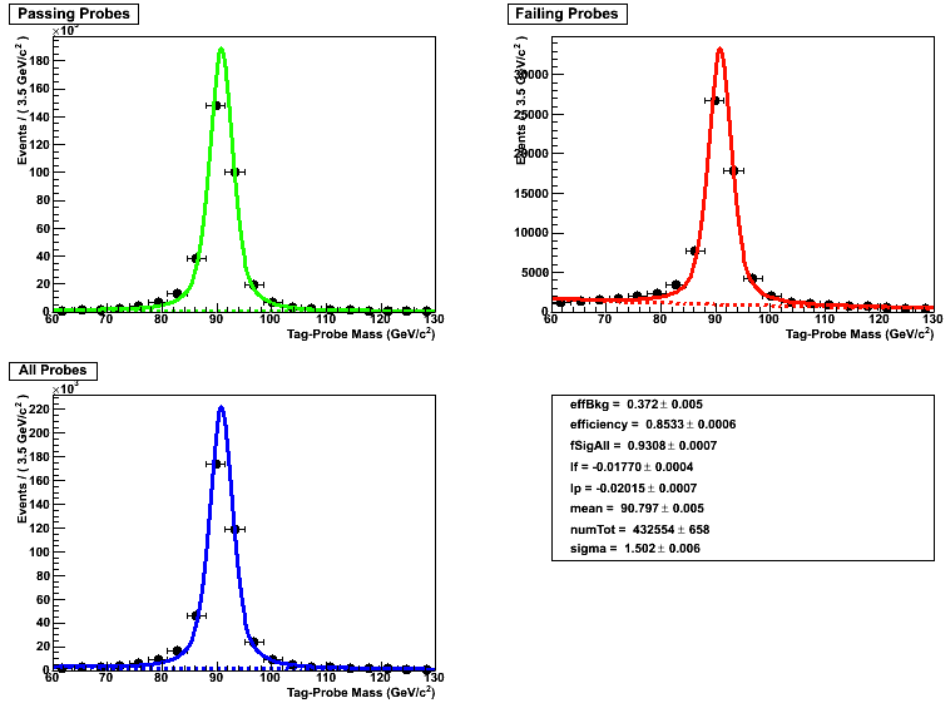
For measuring the efficiency for the muon triggers `HLT_L2DoubleMu23_NoVertex` and `HLT_L2DoubleMu30_NoVertex`, the Z boson events are selected from the `SingleMu` primary dataset, taking events that pass any of the unrescaled single muon triggers, `HLT_IsoMu17`, `HLT_Mu24`, `HLT_Mu30`, `HLT_IsoMu17`, `HLT_IsoMu20_eta2p1`. The tag must be identified as a ‘global muon’, and is required to match within $\Delta R < 0.5$ with a trigger object from one of the single muon triggers. The probe is a ‘high

‘purity’ track that passes the track selection criteria and kinematic cuts used in the physics analysis, as listed in Section 6.3. When combined with the tag it must form an invariant mass that is compatible with that of a Z boson.

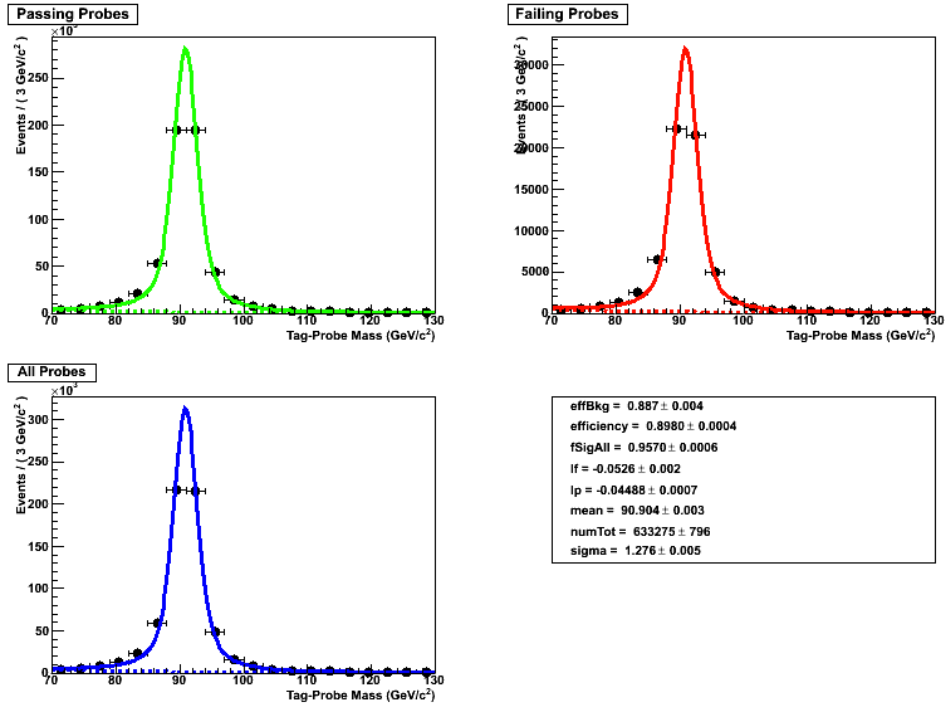
The efficiency of a single leg of the `HLT_L2DoubleMu30_NoVertex` trigger can be estimated from the fraction of the probes that match, within a $\Delta R < 0.5$ with a trigger object from the `HLT_L2DoubleMu30_NoVertex` trigger. It is also required that the tag and the probe are separated by $\Delta R > 0.2$ at the innermost muon chamber layer, in order to avoid the complications arising from nearby muons.

The invariant mass distributions of the tag plus probe in the vicinity of the Z boson mass is shown in Figures 8.18 and 8.19, for both data and simulation. The efficiency is extracted as a parameter from a simultaneous fit to the two ‘tag and probe’ categories, where the probe is and where it is not matched to a single-leg of the `HLT_L2DoubleMu23(30)_NoVertex` path. The trigger efficiency for paths `HLT_L2DoubleMu23_NoVertex` and `HLT_L2DoubleMu30_NoVertex` are shown, integrated in pseudorapidity η and as a function of the p_T of the probe in Figure 8.20. Figure 8.21 displays the (p_T, η) 2D efficiency maps for data and simulation, and observed differences. In the barrel the efficiency for both data and simulation in the last p_T bin is about 90% or higher, while at high η it drops to less than 80%.

The efficiency of a muon to fire a single leg of the `HLT_L2DoubleMu30_NoVertex` trigger, as estimated using this technique, for the range $p_T > 33 \text{ GeV}/c$, is presented in Table 8.6. The table compares the results obtained from data with those from simulation. The difference between the two is taken as the systematic uncertainty. This relative difference, based on the estimates performed for the two trigger paths and in different run periods, and integrating over the muon kinematic range used in the physics analysis, ranges from 4.0% to 5.7% per single-muon leg. To be conservative, we take the largest difference observed; this is also consistent with the data-simulation variations estimated as a function of p_T in Figure 8.20. Since the

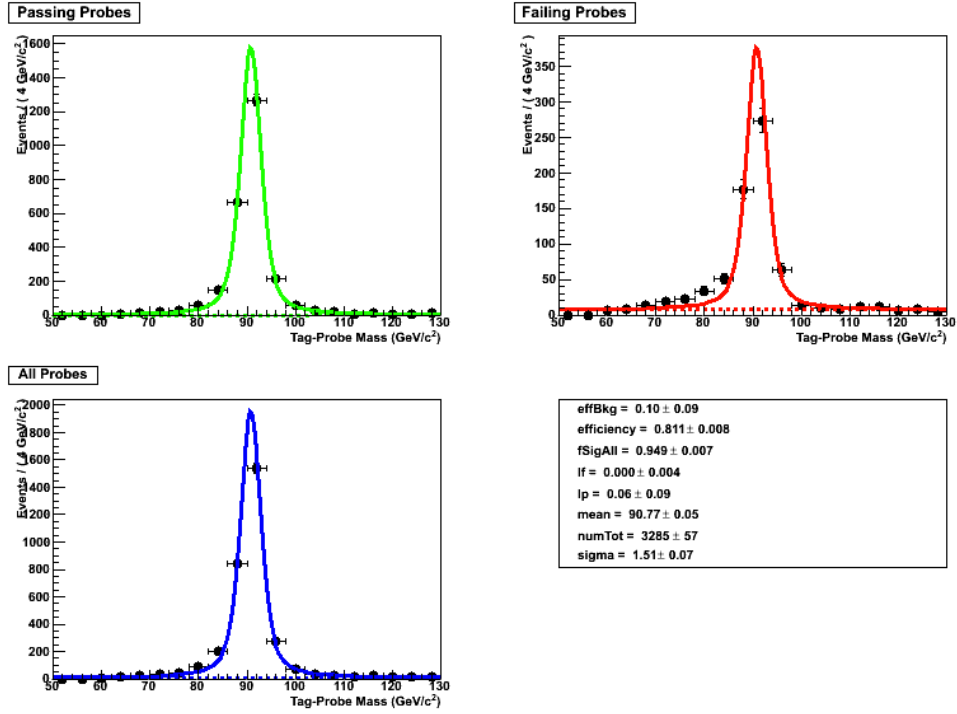


(a) Data

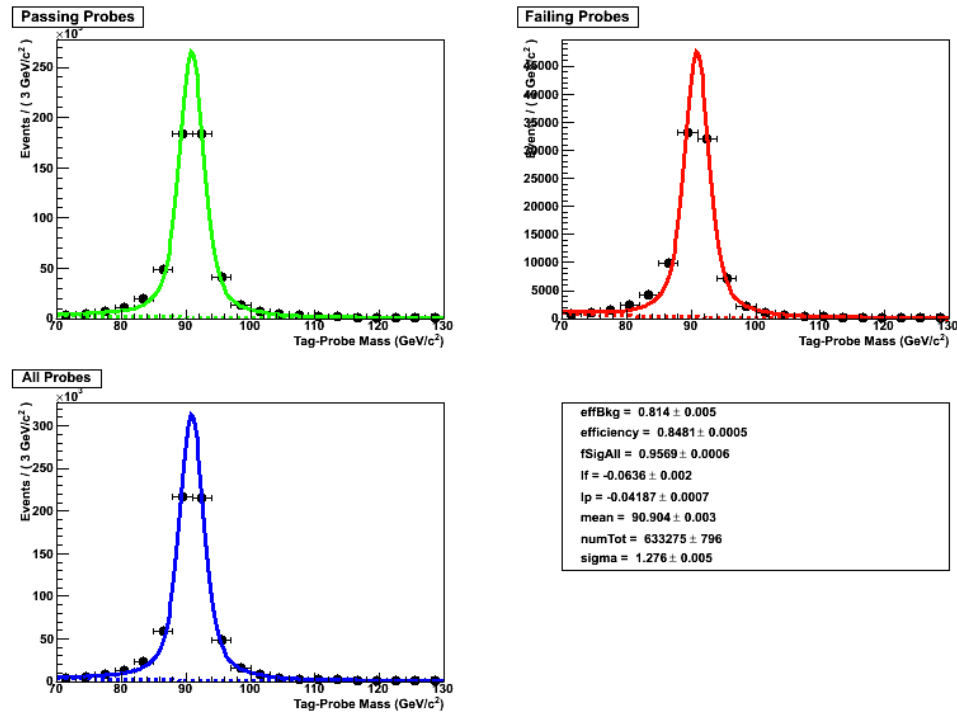


(b) Monte Carlo

Figure 8.18: Tag-probe pair invariant mass distributions, per probe category, employed in the measurement of the HLT_L2DoubleMu23_NoVertex efficiency.




(a) Data




(b) Monte Carlo

Figure 8.19: Tag-probe pair invariant mass distributions, per probe category, employed in the measurement of the HLT_L2DoubleMu30_NoVertex efficiency.



(a) HLT_L2DoubleMu23_NoVertex



(b) HLT_L2DoubleMu30_NoVertex

Figure 8.20: p_T dependence of the single-leg HLT_L2DoubleMu23_NoVertex (left) and HLT_L2DoubleMu30_NoVertex (right) efficiency, as obtained from the tag and probe method, both for data and simulation.

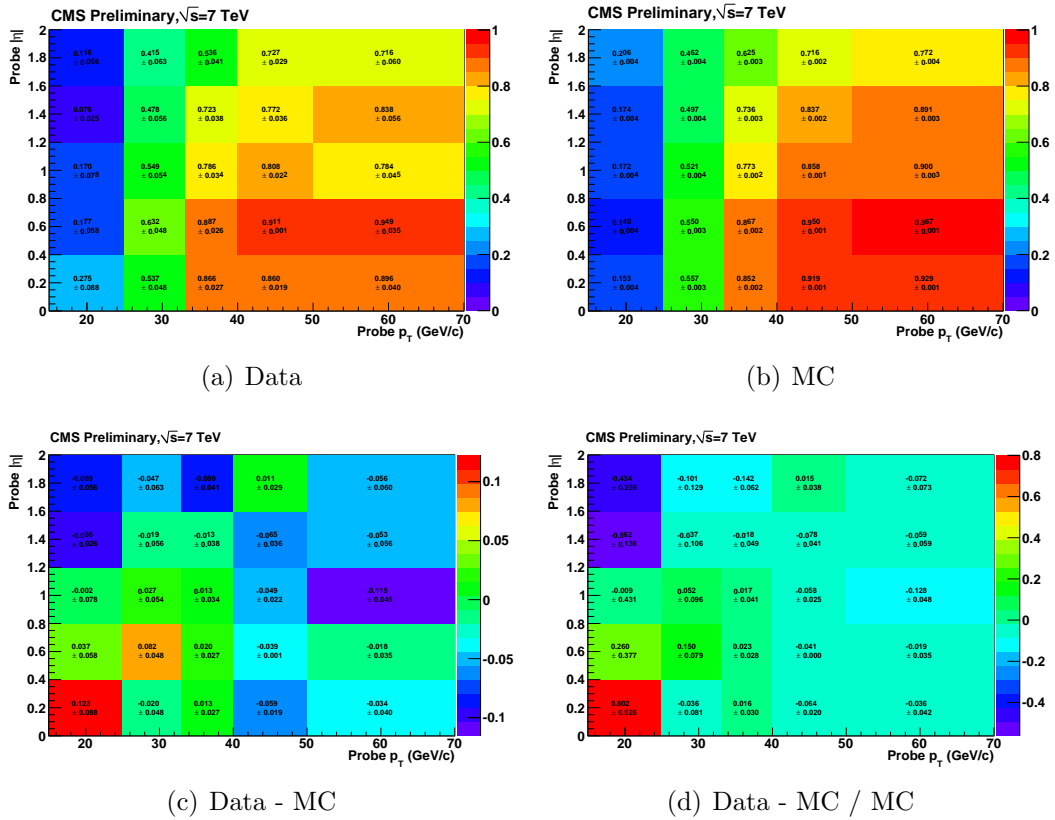


Figure 8.21: Pseudorapidity vs p_T efficiency maps, from Data (top left), MC (top right), absolute difference (bottom left) and fraction difference (bottom right) for the path `HLT_L2DoubleMu30_NoVertex`.

HLT_L2DoubleMu23(30)_NoVertex triggers demand the presence of two muons, the relative systematic uncertainty on its efficiency is taken as twice that on the single leg, namely $\pm 11\%$. Repeating the tag and probe procedure by requiring the tag to pass a higher signal muon p_T threshold of 40 GeV/ c yields a 1% difference, demonstrating the robustness of the method.

Table 8.6: Efficiency of the probe muon to fire a single leg of the displaced muon trigger paths, as estimated using the ‘Tag and Probe’ method, as applied to both data and simulated events. Value ranges correspond to estimates performed in different run periods. The values shown have a statistical uncertainty smaller than 1%.

displaced muon trigger path	p_T^μ (GeV/ c) range	single leg efficiency	
		Data	Simulation
HLT_L2DoubleMu23_NoVertex	> 33	0.85	0.90
HLT_L2DoubleMu30_NoVertex	> 33	0.81	0.85

The double muon triggers impose some requirements on the separation of the two muons in the muon chambers, in order to avoid spurious double-muon signatures generated by a single muon. This leads to inefficiency for closely spaced pairs of muons. If this inefficiency differs in data and simulation, this can give rise to an additional source of systematic uncertainty.

The angular separation of dileptons from Z boson decays is too large to allow them to be used to study this effect. Instead therefore, low mass resonances, such as J/ψ are used. As these also have low momenta, the dependence of the trigger efficiency on ΔR is studied using a dimuon trigger with a lower 3 GeV/ c p_T threshold HLT_DoubleMu3, and correspondingly lower offline p_T thresholds for defining the tag and probe. Figure 8.22 shows how the single leg efficiency of this trigger varies as a function of ΔR . A clear drop in efficiency is seen for $\Delta R < 0.2$, which justifies the cut used in the analysis.




Figure 8.22: Single leg HLT_DoubleMu3 efficiency for dimuon pairs versus the angular separation ΔR (calculated at the surface of the first muon station) of the two muons. Data measurements from J/ψ (black circles), ϕ (red squares) and ρ/ω (green triangles) are compared with simulation.

For comparison, the ΔR separation of selected dilepton candidates from the simulated signal events is shown in Figures 8.23-8.24 (with the $\Delta R > 0.2$ cut applied during the physics selection for the dimuon channel removed).

The efficiency of the trigger is determined relative to the offline selection as a function of the offline E_T threshold using a *Tag and Probe* method. One leg (the tag) must be matched within a cone of aperture $\Delta R < 0.1$ to a triggered object with an $E_T > 38$ GeV, while the other leg is probed to assess whether it passes this same trigger matching criterion, and hence determine the trigger efficiency for a single leg. The offline E_T threshold is varied between 25-50 GeV to study the turn-on behavior of the trigger. The single leg trigger efficiency is plotted versus the offline E_T threshold

To ensure that the efficiencies obtained from the Z sample, which contains prompt electrons and muons, are valid for our signal sample as well, we also examine the trigger efficiency in simulated signal events as a function of the mean lifetime of the X boson. This study is performed at Monte Carlo generator level using a sample with

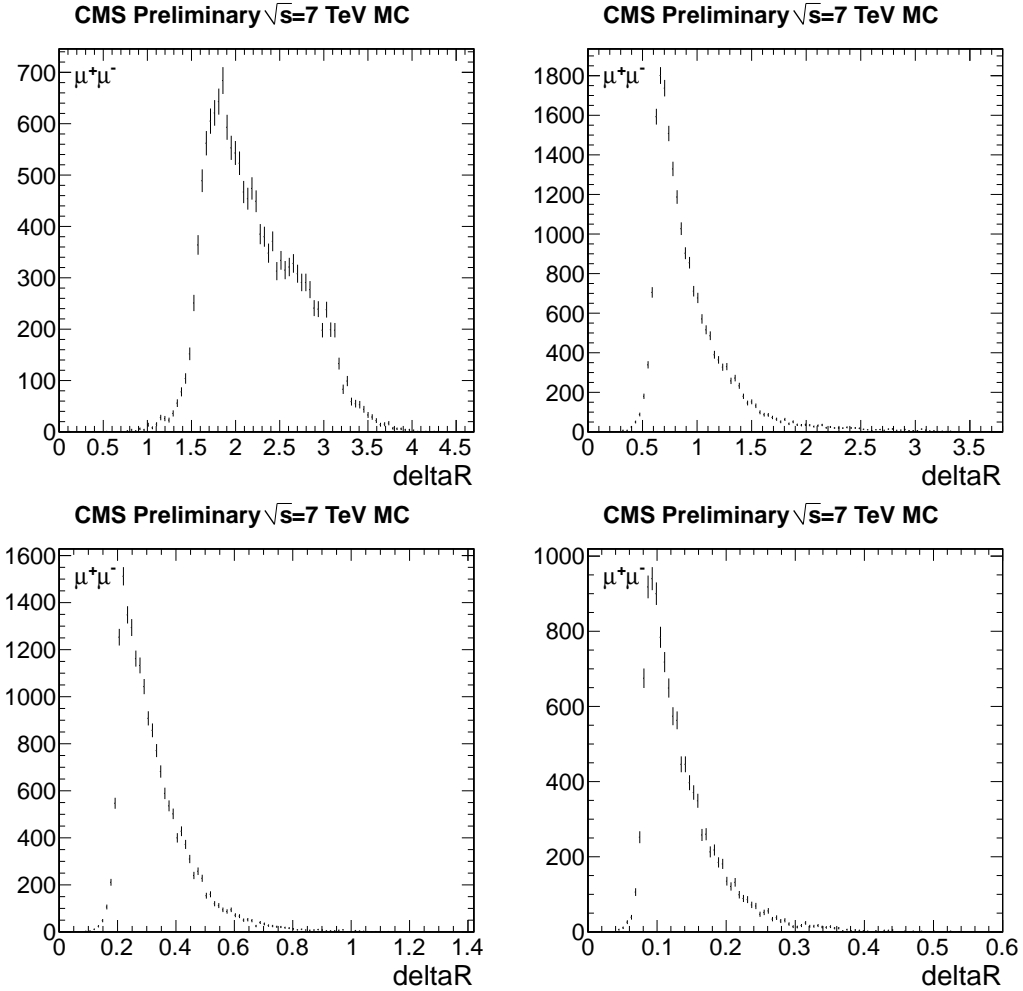


Figure 8.23: Distribution of the angle ΔR between two tracks originating from an X boson decay, simulated for various mass points. The distribution is shown for candidates passing all selection criteria, excluding the dimuon separation cut on ΔR . The top row, reading from left to right, shows the results for $M_{H^0} = 1000$ GeV/ c^2 with $M_X = 350$ and 150 GeV/ c^2 . The second row are the results for $M_{H^0} = 1000$ GeV/ c^2 with $M_X = 50$ and 20 GeV/ c^2 .

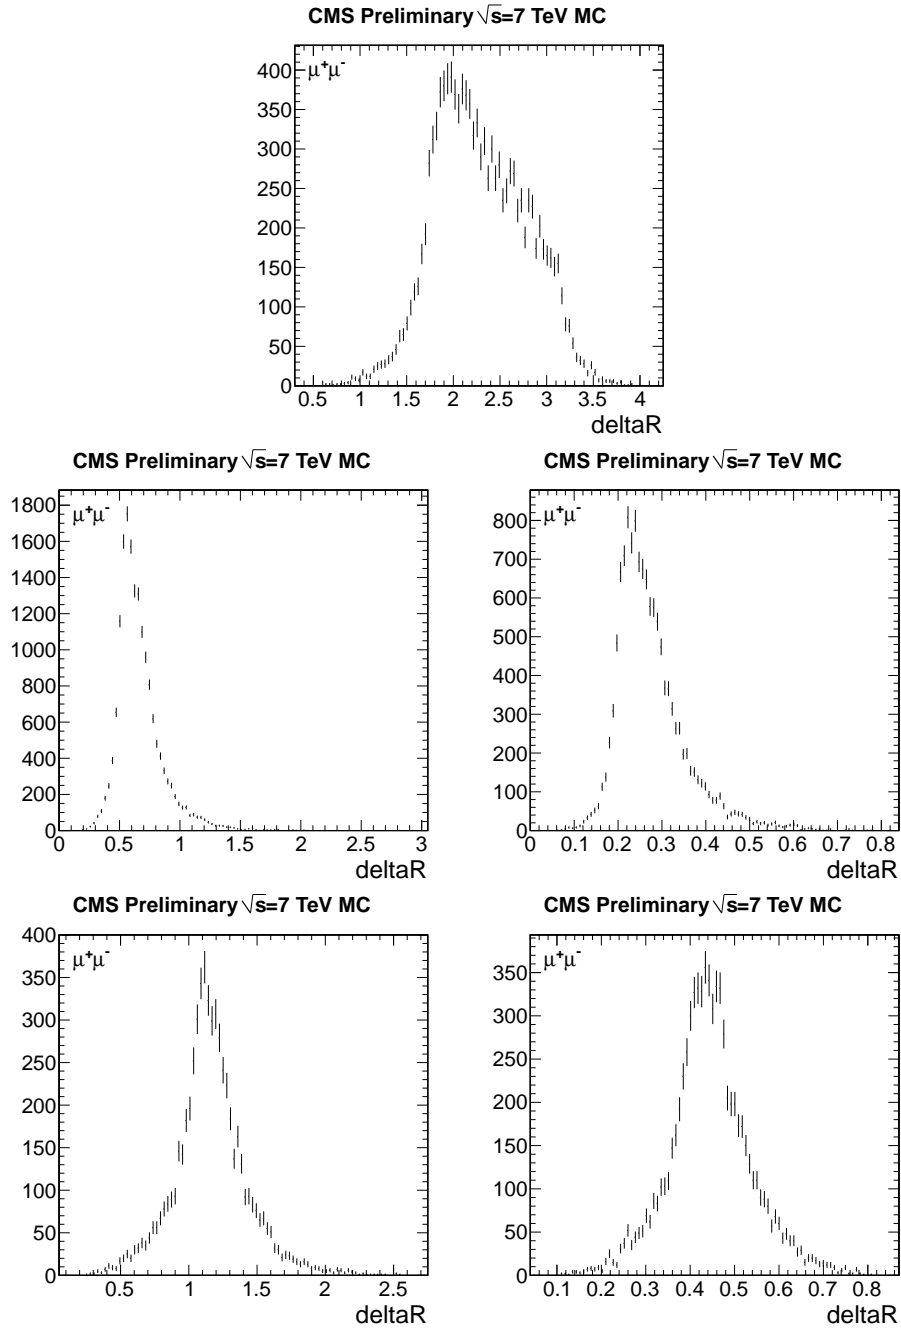


Figure 8.24: Distribution of the angle ΔR between two tracks originating from an X boson decay, simulated for various mass points. The distribution is shown for candidates passing all selection criteria, excluding the dimuon separation cut on ΔR . The top plot shows the result for $M_{H^0} = 400$ GeV/c^2 and $M_X = 150$. The second row, reading from left to right, shows the results for $M_{H^0} = 400$ GeV/c^2 with $M_X = 50$ and 20 GeV/c^2 and the third row for $M_{H^0} = 200$ GeV/c^2 with $M_X = 50$ and 20 GeV/c^2 .




Figure 8.25: Trigger efficiency as a function of mean lifetime for the muon triggers (right) used in our analysis.

$M_{H^0} = 400 \text{ GeV}/c^2$ and $M_X = 50 \text{ GeV}/c^2$; some cuts are also applied at the generator level to mimic the offline cuts used in our analysis. Specifically, we require two muons with $p_T > 33 \text{ GeV}$, and that the X boson has a transverse decay length less than 50 cm, so that it still decays within the CMS tracker. The sample is reweighted to obtain a variety of mean lifetimes following the same procedure as used in the rest of the analysis. The results of this study are shown in Figure 8.25. We observe that the muon efficiency changes by less than 4% over the range of $c\tau$ values considered in our analysis, but this change is small enough that any systematic effects should be much smaller than the value assigned from the tag and probe procedure.

8.6 Transverse decay length correction

If one looks at the distribution of the transverse decay length significance, after removing the cuts on the lepton d_0 significance and $\Delta\varphi$, one observes a discrepancy in the muon channel between simulated events and data, as illustrated in Figure 8.26. Since the discrepancy is visible with loosened cuts, even at small decay length significance,

where the plots are background dominated, it must be due to some inadequacy in the simulation. To correct this difference, we apply a scale factor of 1.2 to the transverse decay length significance in simulated background events in the dimuon channel; the plots shown in Section 6.2 have this correction applied.

By taking the standard selection cuts and requiring the collinearity angle of candidates to be greater than $\pi/2$, rather than less than 0.2 for the muon channel, one gets a sample of candidates that is dominated by background. The transverse decay length significance of these candidates is shown in Figure 8.27 and further demonstrates that the applied smearing factor is reasonable. These plots also show that detector resolution and Bremsstrahlung effects are not expected to yield any background beyond a transverse decay length significance of 20.

The systematic uncertainty resulting from this correction is discussed in Section 8.8.




Figure 8.26: The transverse decay length significance for the dimuon channel with the φ and d_0 cuts removed, before (left) and after (right) the smearing described is applied.




Figure 8.27: The transverse decay length significance of candidates for the dimuon channel with φ difference greater than $\pi/2$.

8.7 Effect of higher-order QCD corrections

For $M_{H^0} = 125 \text{ GeV}/c^2$, the leptons from the X boson decay have a combined efficiency of only a few percent for passing the lepton p_T requirements. For this reason the signal efficiency at this mass is sensitive to the modelling of the Higgs p_T spectrum, which may in turn be influenced by higher order QCD corrections. To study this effect, we reweight the LO H^0 p_T spectrum from our signal sample to match the corresponding Higgs p_T spectrum evaluated at NLO. For $M_{H^0} = 125 \text{ GeV}/c^2$ and $M_X = 20$ (50) GeV/c^2 the signal efficiency changes by 4% (12%). This change is taken as an additional systematic uncertainty in the efficiency for the $M_{H^0} = 125 \text{ GeV}/c^2$ case. For larger H^0 masses, the corresponding systematic uncertainty is below 0.5%, and hence neglected.

8.8 Background uncertainty

As discussed in Chapter 7, there are two primary sources of uncertainty in the background estimation. The uncertainty on the overall background normalisation is used as a systematic uncertainty when determining the expected limits. In addition, different shapes used to fit the background shape are considered and used in the final fit; this results in a negligible change in the observed limits.

In addition, as mentioned in Section 6.2, the transverse decay length significance observed in the dimuon channel is corrected by a factor of 1.2. To assess a systematic uncertainty due to this correction, we perform the L_{xy}/σ fit in Section 7.1 on the unsmeared distribution. This results in a value of 0.01 ± 0.01 events, so we take the difference between this and our nominal value of 0.02 as an additional systematic uncertainty.

Figures 7.3 and 7.4 show the resulting distributions for these crosschecks. For the sample with loosened cuts, the function used for the fit is changed. In the muon channel, because there are relatively few events overall, the distribution is fitted with a Breit-Wigner function (without the error function).

Chapter 9

Results

This section provides a summary of the results for establishing an upper limit of the production cross-section times the branching fraction as a function of the long-lived particle lifetime.

9.1 Upper limits

After all selection requirements are applied, no candidates survive in the muon channel, consistent with the expected mean number of $0.02_{-0.02}^{+0.09}$. We set 95% confidence level (CL) upper limits on the signal process using the statistics software package developed by the CMS Higgs Group [47], which is based upon `Roostats`. It employs the CL_s method [48, 49], which makes use of an unbinned likelihood fit to the dilepton mass spectrum.

This fit to the mass spectrum uses the following functions:

- A Gaussian signal function to represent the signal's mass distribution. For each H^0 mass, the mass resolution used in the Gaussian is obtained from the simulated signal samples, as a function of the X boson mass and lifetime, in-

terpolating between the generated X boson masses with a smooth curve when necessary.

- The sum of two background functions, one distribution representing the background from the Z peak and another more slowly varying distribution representing the non-Z background. These functions are obtained as described in Chapter 7.

The limit calculation takes into account the systematic uncertainties described in Chapter 8 by introducing a nuisance parameter for each uncertainty, marginalized by a log-normal prior distribution. In particular, the normalisation of the Z (non-Z) background can be constrained by the a priori estimate of the total background normalisation presented in Section 7.1, multiplied by the estimated Z (non-Z) background fraction from Section 7.2. However, the fits to the mass spectrum performed as part of the limit calculation strongly constrain the normalisation of the non-Z background, even though the normalisation of a signal is unknown. As a result, it is not necessary to use the a priori estimate of the non-Z background normalisation when calculating the observed limits. (In practice, this allows one to set the assumed uncertainty on the non-Z background normalization to a very large number in the limit calculation software, thereby eliminating any systematic uncertainties related to it). In the case of the Z background, the a priori estimate of the normalisation is used, but it affects only the limits for X bosons whose mass is close to that of the Z. To calculate expected limits, one must have a prediction of the background normalisation, which is taken from Chapter 7. (If the background is predicted to be large, then the expected limit will be poor, whereas if it is predicted to be small, then the expected limit will be good). If one wished to determine the significance of a discovery, instead of setting limits, one could still set the assumed a priori uncertainty on the non-Z background normalisation to a very large number. The background normalisation estimates of Chapter 7 thus have no real importance, except near the Z resonance.

As a first step, upper limits are placed on the mean number N_X of X bosons that could pass the selection requirements, as a function of the X boson mass. The resulting upper limits on N_X at 95% CL for the electron and muon channels are presented in Figure 9.1. These limits are independent of the particular model assumed for X boson production, except for the mass resolution assumed for the signal, which affects the width of the resulting peaks in the observed limit. The mass resolution used for these limits is derived from the Monte Carlo simulation for the hypothesis $M_H = 1000 \text{ GeV}/c^2$, which has the largest mass resolution of the signal points studied and hence yields the most conservative limits. The limits on N_X are close to 3.0 at most masses, as one would expect from Poisson statistics with zero observed signal, but are larger at mass points near the masses of dilepton candidates observed in the data. Since the fitted background levels under the signal peak are extremely small, the limits are not expected to depend on the background shape, and indeed using the alternatives described in Chapter 7 give negligible changes in the results. This figure also shows the 95% CL expected limit band. Except near the Z resonance, the a priori predictions of the background normalisation are very small, so the expected limit is close to 3.0 and the expected limit band is extremely narrow; the median value of the expected limit is in fact equal to 3.0 everywhere.

This can be understood by bearing in mind that expected limits are calculated by generating lots of toy Monte Carlo datasets, by making Poisson fluctuations around the mean a priori background prediction. If the predicted mean background is so small, that within a mass bin of width comparable to the dilepton mass resolution, an average of substantially less than 0.05 candidates are expected, then more than 95% of the toy Monte Carlo samples will have zero candidates in that bin. They will all give an identical limit (close to 3.0), so the expected limit band will be of negligible width. One might also wonder why the observed limit deviates so far outside the expected limit band at certain mass points in Figure 9.1. This does

not mean that a significant signal has been seen! The definition of the expected limit band is that 95% of the time, it should contain the observed limit. Since Figure 9.1 has a great many bins, it is natural that the observed limit fluctuates outside the expected limit band in some of them. – That it does so in much less than the 5% of bins that one might naively expect, is simply a result of the fact that the observed number of candidates must be an integer. – The masses at which these fluctuations occur each correspond to those at which a single dilepton candidate was seen in the data, and obviously a single candidate at any given mass does not indicate that a discovery has been made.




Figure 9.1: The 95% CL upper limits on the mean number of X bosons that could pass the selection requirements in the electron (muon) channels are shown in the left (right) plot. A yellow shaded band shows the 95% quantile for the expected limits, but is almost entirely hidden by the observed limit curves.

The expected number of signal dilepton candidates passing the selection cuts can be expressed as:

$$N_X = 2\mathcal{L}\epsilon_1\sigma B [1 + B(\epsilon_2/\epsilon_1 - 1)] \quad (9.1)$$

where \mathcal{L} is the integrated luminosity, $\epsilon_{(1,2)}$ are the efficiencies defined in Section 6.3.1, σ is the production cross section of the heavy resonance decaying to $X X$, and B is the branching fraction for the decay $X \rightarrow \ell^+ \ell^-$. This expression takes into account that either one or both X bosons in an event may decay to the chosen lepton species, and that, as mentioned in Section 6.3.1, the efficiency to select such an X boson is slightly different in the two cases. Using this equation, the likelihood function can be expressed in terms of σB , thus allowing upper limits to be placed on this quantity. Since N_X in Equation (9.1) depends not only on σB , but also on B , the upper limits depend on the assumed value of B . However, the factor $(\epsilon_2/\epsilon_1 - 1)$ is in practice always positive or very small. Hence if one assumes infinitesimally small B when calculating the limits on σB , such that the factor in square brackets in Equation (9.1) is equal to 1, the resulting limits will be valid, and in some cases conservative, for any value of B .

For each combination of the H^0 and X boson masses listed in Table 5.1, and for a range of X boson lifetimes, the 95% CL upper limits on σB are calculated. The observed limits are shown in Figures 9.2 – 9.4. (No results are shown for $M_{H^0} \leq 200 \text{ GeV}/c^2$ in the electron channel, since the high trigger thresholds result in a very low signal efficiency.) Note that for the muon channel in the $M_{H^0} = 1000 \text{ GeV}/c^2$, $M_X = 20 \text{ GeV}/c^2$ case, the efficiency is significantly reduced because the muons are produced very close together, which causes trigger inefficiencies. Since the observed dilepton candidates do not have masses close to those of the X bosons considered in these plots, they have no effect on the limits. The bands show the 95% quantile for the expected limits.

For H^0 or X boson masses other than those plotted in Figures 9.2 – 9.4, exact limits are not computed, since no simulated signal samples are available with which to determine the signal selection efficiency. However, since the observed limits appear to be monotonic functions of the H^0 and X boson masses, one can infer approximate




Figure 9.2: The 95% CL upper limits on σB for the muon channel for a H^0 mass of $1000 \text{ GeV}/c^2$. Narrow yellow shaded bands show the 95% quantiles for the expected limits.




Figure 9.3: The 95% CL upper limits on σB for the muon channel for a H^0 mass of $400 \text{ GeV}/c^2$. Narrow yellow shaded bands show the 95% quantiles for the expected limits.

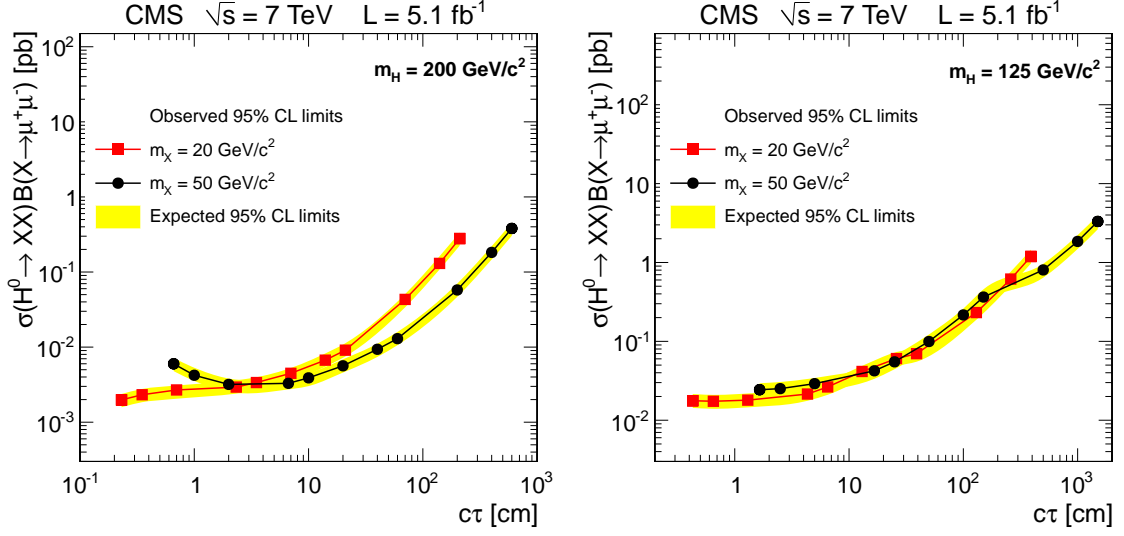


Figure 9.4: The 95% CL upper limits on σB for the muon channel for a H^0 mass of 200 GeV/c^2 (left) and 125 GeV/c^2 (right). Narrow yellow shaded bands show the 95% quantiles for the expected limits.

limits for other masses, provided the latter lie within the range of those shown in the figures. For example, for $M_{H^0} = 1000 \text{ GeV}/c^2$, it should be safe to assume that the limits for $M_X = 170 \text{ GeV}/c^2$ would be at least as good as the weaker of the limits for $M_X = 150 \text{ GeV}/c^2$ and $M_X = 350 \text{ GeV}/c^2$. However, for X bosons that are close in mass to the candidates seen in data, the limits would be worse than this. For these particular masses, Figure 9.1 gives an indication of the factors by which the limits would be degraded.

9.2 Limits on similar models with long-lived exotica

In an alternative signal model, in which $H^0 \rightarrow XY$, where the Y boson does not decay to dileptons, the expected number of signal candidates N_X passing the selection cuts can be expressed as $N_X = \mathcal{L} \epsilon'_1 \sigma B$, where ϵ'_1 is the efficiency to select signal candidates, in this scenario. If the X and Y bosons have identical masses, then $\epsilon'_1 = \epsilon_1$, and

comparison with Equation (9.1) shows that the limits would be a factor of 2 worse than those presented in Figures 9.2 – 9.4.

The limits quoted above are for H^0 bosons produced through gluon-gluon fusion. If the H^0 bosons were instead produced by the sum of all standard model production mechanisms, their momentum spectra would be different. For $M_{H^0} = 125 \text{ GeV}/c^2$, the selection efficiency would then be larger by a factor of approximately 1.18 (1.08) for $M_X = 20$ (50) GeV/c^2 and there would be a corresponding improvement in the limits. This change in efficiency is estimated by reweighting the H^0 boson p_T spectrum to that of a standard model H^0 boson.

If the initial resonance were a Z' with spin 1 instead of a H^0 boson, the acceptance for a dilepton pair to pass the p_T and rapidity selection cuts would be slightly different, mainly because Z' are produced by $q\bar{q}$ annihilation, whilst H^0 are produced predominantly through gluon-gluon fusion. The difference in the parton distribution functions changes the acceptance. If the Z' is to decay to a pair of long-lived spin 0 particles, these cannot be identical if CP is to be conserved, but they are assumed in what follows to have equal mass. Their angular distributions would differ depending on whether they were fundamental spin 0 bosons or spin $\frac{1}{2}$ ‘hidden valley’ quarks that have hadronised in the dark sector into spin 0 bosons [6]. Studies with simulated events (performed at generator level, by forcing Pythia to simulate $Z' \rightarrow H^0 A^0$ or $Z' \rightarrow q\bar{q}$), show that the change in the acceptance for these two Z' models relative to the original H^0 model is less than 3% for a Z' mass of 1000 GeV/c^2 , less than 11% for a 400 GeV/c^2 Z' mass, and below 25% (in the muon channel) for a 200 GeV/c^2 Z' mass. The change in acceptance for a 125 GeV/c^2 Z' mass can be seen to be much larger, ranging from 40% to 80%. As such, this simple generator level study cannot be assumed to give an accurate estimate of the limits for a 125 GeV/c^2 Z' . The exact changes in the acceptance can be evaluated by comparing the results in Tables 9.1 and 9.2 with Table 8.2. The original limits quoted above for the process $H^0 \rightarrow XX$

could be converted to approximate limits on these two Z' models simply by scaling them in inverse proportion to their relative acceptances. A similar generator-level correction could be used to convert these limits to approximate limits on any other process giving rise to the production of two spin 0 long-lived particles.

Table 9.1: This is identical to Table 8.2 except that it shows the results for $Z' \rightarrow XX'$ instead of $H^0 \rightarrow XX$, where X and X' are non-identical spin 0 particles of equal mass. The uncertainties quoted on the acceptances again correspond to the systematic uncertainty on the PDF set.

M'_Z (GeV/ c^2)	M_X	Dimuon channel
1000	350	0.836 ± 0.002
1000	150	0.748 ± 0.001
1000	50	0.738 ± 0.001
1000	20	0.741 ± 0.001
400	150	0.602 ± 0.002
400	50	0.448 ± 0.001
400	20	0.464 ± 0.001
200	50	0.134 ± 0.001
200	20	0.149 ± 0.001
125	50	0.01796 ± 0.00010
125	20	0.01170 ± 0.00004

9.3 Summary

A search for long-lived neutral particles, X, produced in pp collisions at $\sqrt{s} = 7$ TeV and decaying to $\mu^+\mu^-$ has been performed. In the $\mu^+\mu^-$ channel no candidates are observed. These results are consistent with standard model expectations and are used to derive upper limits on the product of cross section times branching fraction for a Higgs boson, in the mass range 200 – 1000 GeV/ c^2 , decaying into a pair of X bosons, in the mass range 20 – 350 GeV, which each decay to $\mu^+\mu^-$. The limits are typically in the range 0.7 – 10 fb, for X bosons with lifetimes in the range $0.1 < c\tau < 200$ cm. For a Higgs mass of 125 GeV/ c^2 , the corresponding limits are in the range 10 – 100 fb.

Table 9.2: This is identical to Table 8.2 except that it shows the results for $Z' \rightarrow Q\bar{Q}$, where Q is a spin-half particle, instead of $H^0 \rightarrow XX$. The uncertainties quoted on the acceptances again correspond to the systematic uncertainty on the PDF set.

M'_Z (GeV/ c^2)	M_Q	Dimuon channel
1000	350	0.834 ± 0.002
1000	150	0.745 ± 0.001
1000	50	0.741 ± 0.001
1000	20	0.745 ± 0.001
400	150	0.596 ± 0.003
400	50	0.449 ± 0.001
400	20	0.465 ± 0.001
200	50	0.135 ± 0.001
200	20	0.153 ± 0.001
125	50	0.005799 ± 0.00003
125	20	0.010682 ± 0.00005

These are the most stringent limits in these channels to date. These results also allow approximate limits to be inferred on the production of singly or doubly produced $X \rightarrow \ell^+\ell^-$ from other resonances, such as the Z' .

Bibliography

- [1] CMS Collaboration, “Measurement of CMS Luminosity ”, Technical Report CMS-PAS-EWK-10-004, CERN, Geneva, (2010).
- [2] CMS Collaboration, “Search in leptonic channels for heavy resonances decaying to long-lived neutral particles”, *JHEP* **1302** (2013) 085, [arXiv:1211.2472](#).
- [3] CMS Collaboration, “Absolute Calibration of the CMS Luminosity Measurement”, *CMS PAS EWK-11-001* (2011).
- [4] M. J. Strassler, “Why Unparticle Models with Mass Gaps are Examples of Hidden Valleys”, [arXiv:0801.0629](#).
- [5] CMS Collaboration, “CMS Tracking Performance Results from early LHC Operation”, *Eur. Phys. J.* **C70** (2010) 1165–1192, [doi:10.1140/epjc/s10052-010-1491-3](#), [arXiv:1007.1988](#).
- [6] T. Han et al., “Phenomenology of hidden valleys at hadron colliders”, *JHEP* **0807** (2008) 008, [arXiv:0712.2041](#).
- [7] M. J. Strassler and K. M. Zurek, “Echoes of a hidden valley at hadron colliders”, *Phys. Lett.* **B651** (2007) 374–379, [doi:10.1016/j.physletb.2007.06.055](#), [arXiv:hep-ph/0604261](#).

- [8] D0 Collaboration Collaboration, “Search for long-lived particles decaying into electron or photon pairs with the D0 detector”, *Phys.Rev.Lett.* **101** (2008) 111802, [arXiv:0806.2223](#).
- [9] D0 Collaboration Collaboration, “Search for neutral, long-lived particles decaying into two muons in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ -TeV”, *Phys.Rev.Lett.* **97** (2006) 161802, [arXiv:hep-ex/0607028](#).
- [10] ATLAS Collaboration, “Search for displaced vertices arising from decays of new heavy particles in 7 TeV pp collisions at ATLAS”, *Phys.Lett.* **B707** (2012) 478–496, [doi:10.1016/j.physletb.2011.12.057](#), [arXiv:1109.2242](#).
- [11] ATLAS Collaboration, “Search for a light Higgs boson decaying to long-lived weakly-interacting particles in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector”, [arXiv:1203.1303](#).
- [12] I. Aitchison and A. Hey, “Gauge Theories in Particle Physics: Volume I: From Relativistic Quantum Mechanics to QED, Third Edition”. Graduate Student Series in Physics. Taylor & Francis, 2002.
- [13] Particle Data Group Collaboration, “Review of Particle Physics”, *Phys. Rev. D* **86** (Jul, 2012) 010001.
- [14] P. Langacker, “The standard model and beyond”. Series in High Energy Physics, Cosmology and Gravitation. Taylor and Francis, Boca Raton, FL, 2010.
- [15] C. C. Tully, “Elementary particle physics in a nutshell”. Princeton Univ. Press, Princeton, NJ, 2011.
- [16] J. L. Hewett et al., “Signatures of long-lived gluinos in split supersymmetry”, *JHEP* **0409** (2004) 070, [arXiv:hep-ph/0408248](#).

- [17] R. Barbier et al., “R-parity violating supersymmetry”, *Phys. Rept.* **420** (2005) 1–202, [arXiv:hep-ph/0406039](#).
- [18] L. Basso et al., “Phenomenology of the minimal B-L extension of the Standard model: Z’ and neutrinos”, *Phys.Rev.* **D80** (2009) 055030, [arXiv:0812.4313](#).
- [19] M. J. Strassler and I. Tomalin, “Long-lived exotica production at the LHC/Tevatron”, in *Physics at TeV colliders. Proceedings, 6th Workshop, dedicated to Thomas Binoth, Les Houches, France, June 8-26, 2009*, pp. 160–170. 2009.
- [20] M. J. Strassler and K. M. Zurek, “Discovering the Higgs through highly-displaced vertices”, *Physics Letters B* **661** (2008), no. 4, 263–267, [arXiv:hep-ph/0605193](#).
- [21] ATLAS Collaboration, “Search for a light Higgs boson decaying to long-lived weakly-interacting particles in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector”, *Phys.Rev.Lett.* **108** (2012) 251801, [doi:10.1103/PhysRevLett.108.251801](#), [arXiv:1203.1303](#).
- [22] e. Evans, Lyndon and e. Bryant, Philip, “LHC Machine”, *JINST* **3** (2008) S08001.
- [23] CMS Collaboration, “The CMS experiment at the CERN LHC”, *JINST* **03** (2008) S08004.
- [24] L. Taylor, “Using Russian navy shells”.
<http://cms.web.cern.ch/news/using-russian-navy-shells>, 2011.
- [25] CMS Collaboration, “Performance of CMS muon reconstruction in pp collision events at $\sqrt{s} = 7$ TeV”, *JINST* **7** (2012) P10002, [arXiv:1206.4071](#).

- [26] CMS Collaboration, “CMS Physics: Technical Design Report Volume 1: Detector Performance and Software”. Technical Design Report CMS. CERN, Geneva, 2006.
- [27] CMS Collaboration, “Measurement of the Luminosity”, *CMS PAS EWK-10-004* (2010).
- [28] CMS Collaboration, “Absolute luminosity normalization”, *CMS DP-2011-002* (2011).
- [29] J. Reinders, “Intel Threading Building Blocks”. O’Reilly & Associates, Inc., Sebastopol, CA, USA, first edition, 2007.
- [30] S. van der Meer, “Calibration of the Effective Beam Height in the ISR”,.
- [31] D. B. et al., “The LHC Fast BCT System: A comparison of Design Parameters with Initial Performance”, *Proceedings of BIW10, Santa Fe, New Mexico, US* (2010).
- [32] S. Harper et al., “Triggers to find Displaced Fermions from Exotica in the 2011 Run”, *CMS AN-Note 2011/112* (2011).
- [33] A. Gay, A. Belyaev, and I. Tomalin, “Triggering on events with highly displaced vertices”, *CMS AN-Note 2009/082* (2009).
- [34] S. M. T. Sjostrand and P. Z. Skands, “PYTHIA 6.4 Physics and Manual”, *JHEP 0605* (2006) 576, [arXiv:0603175](https://arxiv.org/abs/0603175).
- [35] M. Pioppi, “Iterative Tracking”, *CMS IN-Note 2007/065* (2007).
- [36] W. Adam et al., “Track Reconstruction in the CMS Tracker”, *CMS NOTE 2006/041* (2006).

- [37] P. Azzuri et al., “Very Large Impact Parameter Reconstruction”, *CMS AN-Note* **2009/107** (2009).
- [38] CMS Collaboration, “Tracking and Vertexing Results from First Collisions”, *CMS PAS TRK-10-001* (2010).
- [39] CMS Collaboration, “Absolute Calibration of the Luminosity Measurement at CMS: Winter 2012 Update”, Technical Report CMS-PAS-SMP-12-008, CERN, Geneva, (2012).
- [40] CMS Collaboration, “CMS reweighting procedure for pile-up”, (2012).
<https://twiki.cern.ch/twiki/bin/view/PileupReweighting/>.
- [41] CMS Collaboration, “CMS estimating systematic errors due to pile-up modeling”, (2012).
<https://twiki.cern.ch/twiki/bin/viewauth/CMS/PileupSystematicErrors?rev=6>.
- [42] J. Pumplin et al., “New generation of parton distributions with uncertainties from global QCD analysis”, *JHEP* **07** (2002) 012, [arXiv:hep-ph/0201195](https://arxiv.org/abs/hep-ph/0201195).
- [43] D. Bourilkov, R. C. Group, and M. R. Whalley, “LHAPDF: PDF use from the Tevatron to the LHC”, [arXiv:hep-ph/0605240](https://arxiv.org/abs/hep-ph/0605240).
- [44] N. N. C. Liu, “Reconstruction of Cosmic and Beam-Halo Muons”, *CMS NOTE-2008/001* (2008).
- [45] M. Jones and D. Silvers, “A Measurement of the Absolute Track Reconstruction Efficiency Using a Track Embedding Method”, *CMS AN-Note* **2010/209** (2010).
- [46] CMS Collaboration, “Strange Particle Production in pp Collisions at $\sqrt{s} = 0.9$ and 7 TeV”, *JHEP* **05** (2011) 064, [doi:10.1007/JHEP05\(2011\)064](https://doi.org/10.1007/JHEP05(2011)064),
[arXiv:1102.4282](https://arxiv.org/abs/1102.4282).

- [47] CMS Collaboration, “RooStats-based statistics tools for Higgs PAG”, (2012).
<https://twiki.cern.ch/twiki/bin/view/CMS/SWGuideHiggsAnalysisCombinedLimit>.
- [48] A. L. Read, “Presentation of search results: The CL(s) technique”, *J. Phys.*
G28 (2002) 2693–2704, doi:10.1088/0954-3899/28/10/313.
- [49] T. Junk, “Confidence level computation for combining searches with small
statistics”, *Nucl. Instrum. Meth. A* **434** (1999) 435, arXiv:hep-ex/9902006.