

Searches for new physics with di-electrons and data
management for the CMS detector at the LHC

James A. F. Jackson



*A thesis submitted to the University of Bristol in accordance with the
requirements of the degree of Doctor of Philosophy in the Faculty of Science*

32000 words
School of Physics
November 2010

Abstract

In 2010, the CMS detector at the CERN LHC started collecting pp collision data at a centre-of-mass energy $\sqrt{s} = 7$ TeV with a target integrated luminosity for the first run of 1 fb^{-1} .

Preparations for data taking with a particular emphasis on searches for new physics with the e^+e^- final state are presented. Phenomenological and detector effects relevant to searches with high E_t electrons were investigated. It is found that QCD and electroweak NLO corrections to the Drell-Yan process partially cancel, and should therefore be considered together when describing the di-electron high invariant mass spectrum. A detector configuration scheme to allow high E_t electrons to be correctly triggered is described.

An analysis to search for high p_t Z^0 s due to new s -channel physics processes was performed. It is shown that 3σ evidence of excited quark production with a resonant mass of $M_{q^*} = 1$ TeV could be observed with 600 pb^{-1} of integrated luminosity at $\sqrt{s} = 7$ TeV, making this a viable search channel for the first phase of LHC operation.

Alongside the machine commissioning, the LHC computing grid has been preparing for data taking operations. Investigations into the tape migration performance at the UK Tier-1 computing centre were undertaken. Through hardware configuration and algorithmic optimisations, the effective tape write rate per drive is shown to improve from 16 to 90 MBs^{-1} .

Acknowledgements

The past four years have been somewhat of a whirlwind tour of experimental physics, computing and phenomenology. I would like to thank all those who have supported and helped me - I have enjoyed myself immensely and that is largely due to the people I have worked (and gone to the pub) with. I'm bound to have forgotten somebody; I'll buy you a beer.

My sincere thanks to my supervisors Dave Newbold at Bristol and Andrew Sansum at RAL for being supportive and good-humoured throughout. Also thanks are due to Claire Shepherd-Themistocleous, who has provided invaluable help in suggesting and collaborating on many of the physics-related topics I have covered.

For collaborating on various physics topics, my thanks to Alexander Belyaev, Carlo Carloni, Samir Ferrag, Philippe Mine, Fulvio Piccinini, Stefano Pozzorini, Alessandro Vicini and Doreen Wackerath.

Helping me to battle the behemoth that is CASTOR were Martin Bly, Tim Folkes, Sebastien Ponce, Bonny Strong and Shaun De Witt. Thanks for the many discussions and support in changing things (...and often back again).

Then come the cast of many others whom I have worked with over the course of my PhD studies. In particular Chris Brew, Dave Evans, the other Dave Evans, Joel Goldstein, Sam Harper, Simon Metson and Monica Vazquez - thank you all.

To those who have enjoyed late nights by the fire at The Crown, looked after the menagerie or have otherwise been of great support, my thanks. In particular James Baird, Kelly Bowlt, Karen Crawford, Tom Greenham, Peter and Sally Hall, Georgie Herbert, Max Jeffries, Eve Pool and Jen Quayle.

Finally, thanks are due to my parents for their continued encouragement of my scientific curiosity and for supporting me throughout all of my studies.

Declaration

I declare that the work in this thesis was carried out in accordance with the requirements of the University's Regulations and Code of Practice for Research Degree Programmes and that it has not been submitted for any other academic award. While the work presented has been undertaken within the Compact Muon Solenoid collaboration, it is my own with the exception of the high order corrections chapter, and where referenced.

The high order corrections work was undertaken in the context of a Les Houches collaborative working group, within which I was responsible for the experimental interpretation and analysis, generator preparation, data production and final combination.

Any views expressed are those of the author.

SIGNED: DATE:.....

Contents

List of Figures	xv
List of Tables	xxiii
1 Introduction	1
1.1 Definitions	1
1.1.1 Coordinate systems	1
1.1.2 Units	2
1.1.3 Other definitions	2
1.1.4 Machine energies	2
2 Theoretical background	5
2.1 Groups, symmetries and invariance	6
2.1.1 $U(1)$ symmetry and quantum electrodynamics	6
2.1.2 Feynman rules, regularisation and renormalisation	7
2.1.3 Generalised $SU(n)$ groups	10
2.1.4 Spontaneous symmetry breaking	12
2.2 The Standard Model	13
2.2.1 Electroweak model of leptonic interactions	13
2.2.2 Electroweak symmetry breaking	15
2.2.3 Quantum chromodynamics	17
2.2.4 Fermion masses and gauge interactions	19
2.3 Problems with the Standard Model	21
2.3.1 Particle content and free parameters	21
2.3.2 The hierarchy problem	21
2.4 Beyond the Standard Model	22
2.4.1 Supersymmetry	22
2.4.2 Large extra dimensions	23
2.4.3 TeV^{-1} scale extra dimensions	24
2.4.4 Warped extra dimensions	24
3 The neutral-current Drell-Yan process in the high invariant mass region	27
3.1 Introduction	27
3.1.1 Experimental motivation	27

3.1.2	$\mathcal{O}(\alpha_{(s)})$ corrections and $\mathcal{O}(\alpha_{(s)}^2)$ significance	29
3.1.3	HORACE	31
3.1.4	MC@NLO	32
3.1.5	Generator configuration	33
3.2	Combination results	34
3.2.1	Corrections at $\sqrt{s} = 14$ TeV	35
3.2.2	Comparison with PDF effects	37
4	The LHC and the CMS Detector	41
4.1	The Large Hadron Collider	41
4.1.1	Machine commissioning	41
4.1.2	Magnet technology	42
4.1.3	Acceleration systems	42
4.1.4	Bunch structure and luminosity estimates	43
4.2	The Compact Muon Solenoid Detector	44
4.2.1	Superconducting Solenoid	46
4.2.2	Inner Tracking	46
4.2.3	Electromagnetic Calorimetry	48
4.2.4	Hadronic Calorimetry	50
4.2.5	Muon System	51
4.2.6	Level-1 Triggering	53
4.2.7	Data Acquisition and High Level Triggering	55
5	Level-1 triggering for high E_t electromagnetic objects	59
5.1	Introduction	59
5.1.1	Motivation	59
5.1.2	Electronics overview	60
5.2	ECAL TPG response to saturating signals	63
5.2.1	Experimental setup and validation	63
5.2.2	TPG simulation and configuration recommendations	64
5.3	Level-1 trigger choice for high E_t electron analyses	65
5.3.1	Electromagnetic Level-1 trigger algorithm	67
5.3.2	Trigger choice	68
5.3.3	Jet background rates at $L = 2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$, $\sqrt{s} = 14$ TeV	70
5.4	Conclusions	71
6	Computing for CMS	73
6.1	The Event Data Model	73
6.1.1	Data persistence	73
6.1.2	Event setup records	74
6.2	Data management	75
6.2.1	Data tiers	75

6.2.2	Data flow	75
6.2.3	Transfer technologies	76
6.3	Site requirements	77
6.3.1	Tier-0	77
6.3.2	Tier-1	77
6.3.3	Tier-2	78
6.4	RAL Tier-1 planning	78
6.4.1	Site requirements	78
6.4.2	Tape system performance	79
7	Optimising tape migration performance with the CASTOR HSM	81
7.1	RAL Tier-1 Hardware	81
7.1.1	Disk servers	81
7.1.2	Tape system	82
7.1.3	Network topology	83
7.1.4	Transfer performance	83
7.2	CASTOR	84
7.2.1	Architecture overview	84
7.2.2	Job scheduling	85
7.2.3	Tape migration	87
7.2.4	System monitoring	87
7.2.5	CMS configuration at RAL	89
7.3	Hardware optimisation	90
7.3.1	Linux I/O overview	91
7.3.2	Measurement tools	91
7.3.3	Baseline measurements and parameter choice	92
7.3.4	RAID controller and configuration parameters	93
7.3.5	Memory parameters	94
7.3.6	Block device and queue parameters	94
7.3.7	Configuration recommendations	97
7.4	CASTOR optimisation	104
7.4.1	Job scheduling	104
7.4.2	Migration candidate selection	105
7.4.3	Stream policies	105
7.5	Performance during STEP09	107
7.6	Conclusions	108
8	High p_t Z^0s as a probe for physics beyond the standard model	111
8.1	Phenomenological background	111
8.1.1	SUSY and Technicolour	111
8.1.2	Quark compositeness	112
8.1.3	Current experimental limits	113

8.2	Datasets considered	115
8.3	Triggering	115
8.3.1	Level-1 seed	117
8.3.2	HLT path	117
8.4	Electron reconstruction	117
8.4.1	Electron reconstruction algorithm	117
8.4.2	Close electron reconstruction	119
8.4.3	Quantification of energy loss	119
8.5	Electron identification cuts	123
8.5.1	Current HEEP electron selection	123
8.5.2	HEEP selection performance	126
8.6	Efficiency measurements	126
8.6.1	The Tag and Probe method	126
8.6.2	Background estimation	128
8.6.3	Efficiency factorisation scheme	129
8.6.4	Calorimeter clustering efficiency	130
8.6.5	GSF electron reconstruction efficiency	132
8.6.6	Electron identification efficiency	134
8.6.7	Trigger efficiency	136
8.6.8	Combination of integrated efficiencies	136
8.6.9	Cross check with $W^\pm \rightarrow e\nu$	138
8.6.10	Electron identification $N - 1$ efficiencies	140
8.7	Background estimation	140
8.7.1	General issues	140
8.7.2	$t\bar{t}$ using the b -tagging method	141
8.7.3	Jet backgrounds using the fake rate method	143
8.7.4	$Z^0 \rightarrow e^+e^-$ with Monte Carlo and W^\pm hadronic recoil	148
8.7.5	Combination of backgrounds	150
8.7.6	Control of backgrounds	151
8.8	Systematic Uncertainties	152
8.8.1	Electron identification	152
8.8.2	Calibration and alignment	153
8.8.3	PDF uncertainty	154
8.8.4	Factorisation and renormalisation scales	154
8.8.5	PDF choice	155
8.8.6	High order corrections	156
8.8.7	Combination	156
8.9	Search reach at $\sqrt{s} = 10$ TeV	157
8.9.1	The p-value scan technique	157
8.9.2	Data preparation and background fit	158
8.9.3	Search reach determination	158

8.10	Future prospects	160
8.10.1	Application of p-value technique to data	160
8.10.2	Analysis potential at $\sqrt{s} = 14$ and 7 TeV	160
8.10.3	Boosted W^{\pm} s	162
8.11	Conclusions	164
9	Conclusions	167
9.1	Work presented	167
9.1.1	Level-1 triggering for high E_t electrons	167
9.1.2	The neutral-current Drell-Yan process in the high invariant mass re- gion	167
9.1.3	Optimising tape migration performance with the CASTOR HSM . . .	168
9.1.4	High p_t Z^0 s as a probe for physics beyond the standard model . . .	168
9.2	Future prospects	169
A	Tag and Probe E_t, η binned efficiencies	171
B	$N - 1$ efficiency measurements	179
	Glossary	183
	Bibliography	187

List of Figures

1	Introduction	1
2	Theoretical background	5
2.1	QED electron-photon vertex and example one loop correction	8
2.2	Particle content of the Standard Model	14
2.3	The running of $\alpha_s(Q^2)$ at leading order	19
2.4	The running of gauge couplings in the Standard Model	23
2.5	The running of gauge couplings in the MSSM	23
2.6	Power-law running of gauge couplings with TeV^{-1} scale extra dimensions	24
3	The neutral-current Drell-Yan process in the high invariant mass region	27
3.1	Deviation from Standard Model Drell-Yan $d\sigma/dM_{ll}$ due to large and TeV^{-1} scale extra dimensions at $\sqrt{s} = 14 \text{ TeV}$	28
3.2	Tree level diagrams for the $q\bar{q}$ and $\gamma\gamma$ subprocesses	29
3.3	Estimated $\mathcal{O}(\alpha)$ corrections for $q\bar{q} \rightarrow e^+e^-$	30
3.4	Estimated $\mathcal{O}(\alpha^2)$ corrections for $q\bar{q} \rightarrow e^+e^-$	31
3.5	$\mathcal{O}(\alpha)$ virtual corrections	32
3.6	$\mathcal{O}(\alpha)$ real radiative corrections	32
3.7	$\mathcal{O}(\alpha)$ single photon-induced process corrections	32
3.8	High order QED photon showers	33
3.9	$\mathcal{O}(\alpha_s)$ real radiative corrections	33
3.10	$\mathcal{O}(\alpha_s)$ virtual corrections	33
3.11	M_{ee} differential cross sections at LO and NLO	35
3.12	Corrections to M_{ee} due to QCD and EWK NLO corrections	35
3.13	Electron p_t differential cross sections at LO and NLO	35
3.14	Corrections to the electron p_t due to QCD and EWK NLO corrections	35
3.15	De-electron p_t differential cross sections at LO and NLO	36
3.16	Corrections to the de-electron p_t due to QCD and EWK NLO corrections	36
3.17	Electron pseudorapidity differential cross sections at LO and NLO	36
3.18	Corrections to the electron pseudorapidity due to QCD and EWK NLO corrections	36
3.19	Di-electron pseudorapidity differential cross sections at LO and NLO	36

3.20	Corrections to the di-electron pseudorapidity due to QCD and EWK NLO corrections	36
3.21	Uncertainty on M_{ee} due to PDF fit	38
4	The LHC and the CMS Detector	41
4.1	Schematic layout of the LHC	42
4.2	Cross section of an LHC dipole	43
4.3	Nominal bunch structure for 25 ns LHC operation	44
4.4	Schematic layout of the CMS detector	45
4.5	Cross-section of the silicon tracking systems	47
4.6	Detailed layout of pixel detector	47
4.7	Silicon strip tracker geometry	47
4.8	Tracker momentum resolution	48
4.9	Tracker material budget	48
4.10	Overview ECAL layout	49
4.11	Evolution of ECAL energy resolution	50
4.12	Overview HCAL layout	51
4.13	Evolution of HCAL transverse energy resolution for each subsystem	52
4.14	Muon system layout	53
4.15	Evolution of muon transverse momentum resolution	54
4.16	Level-1 Trigger layout	54
4.17	Overview of the DAQ system	56
5	Level-1 triggering for high E_t electromagnetic objects	59
5.1	Schematic view of the ECAL front-end electronics	61
5.2	Example MGPA response	61
5.3	Response and linearity measurement of the MGPA experimental setup for a range of input charges	64
5.4	Verification of TPG simulation in the design range	65
5.5	TPG simulation for saturating signals	66
5.6	Example of early and late fake parasitic trigger with existing run-time configuration setting saturated channels to 0 after linearisation	67
5.7	Level-1 electromagnetic object trigger overview	68
5.8	Performance of L1_SingleIsoEG25 trigger	69
5.9	E_t and $ \eta $ detail of L1_SingleIsoEG25 performance	69
5.10	Performance of L1_SingleEG25 trigger	69
5.11	Jet rate and Z^0 efficiency of L1_SingleEG trigger	71
6	Computing for CMS	73
6.1	Illustration of EventSetup Interval of Validity behaviour	74
6.2	Overview of data flow and computing steps for CMS	76

7	Optimising tape migration performance with the CASTOR HSM	81
7.1	UK Tier-1 Network Topology	83
7.2	Switch stack network detail	84
7.3	CASTOR architecture	85
7.4	Example CASTOR monitoring	89
7.5	Data flow within the Tier-1	90
7.6	Summary of Linux I/O layers	92
7.7	Flowchart of disk-draining <code>bestTapeCopyForStream</code> algorithm	106
7.8	Example disk server network usage with job draining operational	106
7.9	STEP09 tape server network activity with stream policies: Dedicated CMS Tape Servers	109
7.10	STEP09 tape server network activity with stream policies: Shared Tape Servers	109
7.11	STEP09 tape server network activity with stream policies: Non-CMS Tape Servers	109
8	High p_t Z^0's as a probe for physics beyond the standard model	111
8.1	One of the representative Feynman diagrams for the cascade gluino decay leading to boosted Z^0 bosons	112
8.2	One of the representative Feynman diagrams for the Techni-rho ρ_{TC}^+ decay leading to boosted Z^0 bosons	112
8.3	Feynman diagrams for the single and double q^* production processes	113
8.4	Cross sections for the single and double q^* production processes	114
8.5	Z^0 decay kinematic probabilities in the (η, p_t) plane for $pp \rightarrow q^* \rightarrow qZ^0$	114
8.6	Z^0 decay kinematic probabilities in the (η, p_t) plane for $pp \rightarrow q^*q^* \rightarrow$ $qZ^0 + X$	114
8.7	The Bayesian limit on the differential cross section for anomalous $Z^0 \rightarrow$ e^+e^- production at 95% CL from CDF	115
8.8	$Z^0 \rightarrow e^+e^-$ p_t Data/Theory $- 1$ comparison from D0	115
8.9	Worst case $Z^0 \rightarrow e^+e^-$ decay angle from $X \rightarrow Z^0p$	120
8.10	Boosted Z^0 reconstruction efficiency with old algorithm	120
8.11	Boosted Z^0 reconstruction efficiency with the modified algorithm	120
8.12	Boosted Z^0 reconstruction efficiency with old algorithm for $-2 < \eta_Z < 2$	120
8.13	Boosted Z^0 reconstruction efficiency with the modified algorithm for $-2 <$ $\eta_Z < 2$	120
8.14	Example of close electron merging	121
8.15	Comparison of current and modified clustering algorithms for $\gamma_Z = 1, 5,$ 10 and 20	122
8.16	Reconstructed Z^0 invariant mass with modified algorithm	122
8.17	Breit-Wigner \otimes Gaussian fit to Z^0 invariant mass reconstructed with the old algorithm	123

8.18	Breit-Wigner \otimes Gaussian fit to Z^0 invariant mass reconstructed with the new algorithm	123
8.19	Modified reconstruction performance for Z^0 mass with di-electron p_t cut	123
8.20	Comparison of high p_t mass performance with standard reconstruction	123
8.21	Single-electron energy loss with modified reconstruction	124
8.22	Modified reconstruction performance for Z^0 mass with energy corrections	124
8.23	Reconstructed Z^0 invariant mass for $pp \rightarrow Z^0$ and $pp \rightarrow u^* \rightarrow Z^0 u$	124
8.24	Example of sideband subtraction technique	128
8.25	Tag + All Probes invariant mass spectrum for $\epsilon_{\text{Clustering}}$	130
8.26	Tag + Passing Probes invariant mass spectrum for $\epsilon_{\text{Clustering}}$	130
8.27	Tag and probe measurement of $\epsilon_{\text{Clustering}}$	131
8.28	Tag + All Probes invariant mass spectrum for $\epsilon_{\text{GSF Reco}}$	132
8.29	Tag + Passing Probes invariant mass spectrum for $\epsilon_{\text{GSF Reco}}$	132
8.30	Tag and probe measurement of $\epsilon_{\text{GSF Reco}}$	133
8.31	Tag + All Probes invariant mass spectrum for $\epsilon_{\text{Ele ID}}$	134
8.32	Tag + Passing Probes invariant mass spectrum for $\epsilon_{\text{Ele ID}}$	134
8.33	Tag and probe measurement of $\epsilon_{\text{Ele ID}}$	135
8.34	Tag + All Probes invariant mass spectrum for $\epsilon_{\text{Trigger}}$	136
8.35	Tag + Passing Probes invariant mass spectrum for $\epsilon_{\text{Trigger}}$	136
8.36	Tag and probe measurement of $\epsilon_{\text{Trigger}}$	137
8.37	$W^\pm \rightarrow e\nu$ cross check of clustering efficiency measurement	138
8.38	$W^\pm \rightarrow e\nu$ cross check of GSF electron reconstruction efficiency measurement	138
8.39	$W^\pm \rightarrow e\nu$ cross check of electron ID efficiency measurement	139
8.40	$W^\pm \rightarrow e\nu$ cross check of trigger efficiency measurement	139
8.41	Example $N - 1$ efficiency measurement of $ \Delta\eta_{\text{in}} $	140
8.42	Example $N - 1$ efficiency measurement of $E^{2 \times 5} / E^{5 \times 5}$	140
8.43	b -discriminant top identification efficiency as a function of cut value	142
8.44	n_1 selection for $t\bar{t}$ background estimation	143
8.45	$n_{1,\text{tight}}$ selection for $t\bar{t}$ background estimation	143
8.46	n_2 selection for $t\bar{t}$ background estimation	144
8.47	$n_{2,\text{tight}}$ selection for $t\bar{t}$ background estimation	144
8.48	$t\bar{t}$ estimate using $n_{1,\text{tight}}$ and $n_{2,\text{tight}}$ selections	144
8.49	Measured denominator jet spectrum in the barrel	146
8.50	Measured numerator jet spectrum in the barrel	146
8.51	Measured denominator jet spectrum in the endcap	146
8.52	Measured numerator jet spectrum in the endcap	146
8.53	Measured jet fake rate in the barrel	147
8.54	Measured jet fake rate in the endcap	147
8.55	Estimated jet background for $W^\pm + \text{Jets}$ and $\gamma + \text{Jets}$	148
8.56	Estimated jet background for $W^\pm + \text{Jets}$, $\gamma + \text{Jets}$ and QCD di-jets	148

8.57	Estimated jet fake rate for realistic data sample	148
8.58	Correlation between p_t of hadronic recoil system and Z^0	149
8.59	Background contamination to the $W^\pm p_t$ estimate	149
8.60	Performance of the W^\pm recoil estimate against $Z^0 \rightarrow e^+e^-$ from Monte Carlo	150
8.61	Combined background estimates without signal	150
8.62	Combined background estimates with 1 TeV u^*	150
8.63	Selected signal + background events without control cuts	151
8.64	Selected signal + background events with $60 < M_{ee} < 120$ GeV cut	151
8.65	Selected signal + background events with opposite sign electron cut	151
8.66	Selected signal + background events with opposite sign electron and $60 < M_{ee} < 120$ GeV cuts	151
8.67	Event yield change for 1 TeV u^* with both control cuts	152
8.68	Event yield change for 1.75 TeV u^* with both control cuts	152
8.69	PDF uncertainty on $p_{t,Z}$ for a 1 TeV u^* with the CTEQ6M PDF set	154
8.70	PDF uncertainty on $p_{t,Z}$ for a 2 TeV u^* with the CTEQ6M PDF set	154
8.71	p-value distribution for background only pseudoexperiments with 200 pb^{-1}	159
8.72	p-value distribution for signal + background pseudoexperiments with 200 pb^{-1}	159
8.73	3σ background p-value rejection criteria	161
8.74	5σ background p-value rejection criteria	161
8.75	Search reach determination	161
8.76	Expected 3σ evidence potential at $\sqrt{s} = 10$ TeV	161
8.77	Expected 5σ discovery potential at $\sqrt{s} = 10$ TeV	161
8.78	Example 200 pb^{-1} 5σ discovery pseudo-experiment background fit and p-value scan	162
8.79	Ratios of parton luminosities at $\sqrt{s} = 10$ and 14 TeV with the MSTW2008NLO PDF set	163
8.80	Ratios of parton luminosities at $\sqrt{s} = 7$ and 10 TeV with the MSTW2008NLO PDF set	163
8.81	Search reach determination at $\sqrt{s} = 14$ TeV	163
8.82	Estimated 3σ evidence potential at $\sqrt{s} = 14$ TeV	163
8.83	Estimated 5σ discovery potential at $\sqrt{s} = 14$ TeV	163
8.84	Discriminating boosted W^\pm s from background	164
9	Conclusions	167
A	Tag and Probe E_t, η binned efficiencies	171
B	$N - 1$ efficiency measurements	179
B.1	Barrel cuts $N - 1$ performance	179

B.2 Endcap cuts $N - 1$ performance 180

List of Tables

1	Introduction	1
2	Theoretical background	5
3	The neutral-current Drell-Yan process in the high invariant mass region	27
4	The LHC and the CMS Detector	41
5	Level-1 triggering for high E_t electromagnetic objects	59
5.1	ADC gain ID settings	61
5.2	Di-jet background rates for SingleEG25 Level-1 trigger	70
6	Computing for CMS	73
6.1	Projected resource requirements for the RAL Tier-1	79
7	Optimising tape migration performance with the CASTOR HSM	81
7.1	Viglen 2006 disk server hardware specifications	82
7.2	Clustervision 2005 disk server hardware specifications	82
7.3	Baseline hardware configuration	93
7.4	Disk I/O Rates - Baseline configuration	93
7.5	Disk I/O Rates - Multiple RAID Units (Write /dev/sdb, Read /dev/sdc)	95
7.6	Disk I/O Rates - Multiple RAID Units (Write /dev/sdc, Read /dev/sdb)	95
7.7	Effect on disk I/O rates of setting storsave to performance	95
7.8	Effect on disk I/O rates of varying memory parameters	95
7.9	Effect on disk I/O rates of varying max_sectors_kb	96
7.10	Effect on disk I/O rates of varying nr_requests	96
7.11	Effect on disk I/O rates of varying blockdev read ahead size: Completely Fair Queue scheduler	98
7.12	Effect on disk I/O rates of varying blockdev read ahead size: Anticipatory scheduler	99

7.13	Effect on disk I/O rates of varying blockdev read ahead size: Deadline scheduler	100
7.14	Effect on disk I/O rates of varying blockdev read ahead size: Noop scheduler	101
7.15	3Ware disk I/O performance per thread: CFQ Scheduler	102
7.16	3Ware disk I/O aggregate multithreaded performance: CFQ Scheduler	102
7.17	3Ware disk I/O performance per thread: AS Scheduler	103
7.18	3Ware disk I/O aggregate multithreaded performance: AS Scheduler	103
7.19	Tape write performance for CMS during STEP09	107
8	High p_t Z^0s as a probe for physics beyond the standard model	111
8.1	Background datasets	116
8.2	Cross-check datasets	116
8.3	Signal datasets generated	116
8.4	GSF electron preselection cuts	118
8.5	HEEP selection cuts v2.0	125
8.6	HEEP cuts event selection efficiency	127
8.7	Barrel HEEP selection performance	127
8.8	Endcap HEEP selection performance	127
8.9	Modified HEEP cuts event selection efficiency	127
8.10	Tag and Probe criteria for clustering efficiency measurements	130
8.11	Integrated clustering efficiency (barrel)	131
8.12	Integrated clustering efficiency (endcap)	131
8.13	Tag and Probe criteria for GSF electron reconstruction efficiency measurements	132
8.14	Integrated GSF electron reconstruction efficiency (barrel)	133
8.15	Integrated GSF electron reconstruction efficiency (endcap)	133
8.16	Tag and Probe criteria for electron ID efficiency measurements	134
8.17	Integrated electron ID efficiency (barrel)	135
8.18	Integrated electron ID efficiency (endcap)	135
8.19	Tag and Probe criteria for trigger efficiency measurements	136
8.20	Integrated trigger efficiency (barrel)	137
8.21	Integrated trigger efficiency (endcap)	137
8.22	Tag and Probe criteria for $N - 1$ efficiency measurements	140
8.23	Jet selection criteria	143
8.24	Measured b -tagging efficiency	143
8.25	Total number of estimated $t\bar{t}$ events	144
8.26	Jet triggers (1E31 v0.6 menu) used to construct the fake rate estimate	145
8.27	Loose jet fake rate selection	146
8.28	Tight jet fake rate selection	146
8.29	Systematic effect of electron ID cuts	153

8.30	Systematic effect of altering μ_f and μ_r	155
8.31	Systematic effect of PDF choice	156
8.32	Combination of systematic uncertainties to maximise search reach	157
8.33	Combination of systematic uncertainties to minimise search reach	157
9	Conclusions	167
A	Tag and Probe E_t, η binned efficiencies	171
A.1	Clustering efficiency ($Z^0 \rightarrow e^+e^-$ only)	172
A.2	Clustering efficiency (Background subtracted)	172
A.3	GSF electron reconstruction efficiency ($Z^0 \rightarrow e^+e^-$ only)	173
A.4	GSF electron reconstruction efficiency (Background subtracted)	173
A.5	Electron ID efficiency ($Z^0 \rightarrow e^+e^-$ only)	174
A.6	Electron ID efficiency (Background subtracted)	174
A.7	Trigger efficiency ($Z^0 \rightarrow e^+e^-$ only)	175
A.8	Trigger efficiency (Background subtracted)	175
A.9	Electron ID efficiency ($Z^0 \rightarrow e^+e^-$ only)	176
A.10	Electron ID efficiency ($W^\pm \rightarrow e\nu$)	176
B	$N - 1$ efficiency measurements	179

1. Introduction

1.1 Definitions

1.1.1 Coordinate systems

When discussing the physical dimensions of the detector, Cartesian coordinates are used, where z is aligned along the beampipe (pointing towards the Jura mountains), y is positive in the upwards vertical direction and x points inwards to the centre of the accelerator.

The kinematics of physical events and certain aspects of the detector are discussed in terms of the coordinate system (η, ϕ, z) . In this system, z is defined as in the Cartesian system and the azimuthal angle, ϕ , is given by

$$\phi = \arctan \frac{y}{x}.$$

The pseudorapidity, η , is defined as

$$\eta = -\ln \left(\tan \frac{\theta}{2} \right),$$

where the polar angle, θ , is given by

$$\theta = \tan^{-1} \frac{\sqrt{x^2 + y^2}}{z}.$$

The quantity ΔR is often used, which describes a separation in η and ϕ

$$\Delta R(\vec{v}_1, \vec{v}_2) = \sqrt{\Delta\eta(\vec{v}_1, \vec{v}_2)^2 + \Delta\phi(\vec{v}_1, \vec{v}_2)^2},$$

where $\vec{v}_{1,2}$ are vectors in the (η, ϕ, z) basis.

1.1.2 Units

All calculations are expressed in terms of natural units, where energy is measured in eV, and it is defined that

$$\frac{\hbar}{2\pi} = c = 1.$$

In this system, the units of mass, momentum and time are eV, eV and eV^{-1} respectively.

When discussing units of data, powers of 2 are indicated as KiB (1024 B), whereas powers of 10 are indicated as KB (1000 B).

1.1.3 Other definitions

Matrices are indicated by bold face (**T**), and three-vectors by an over-arrow (\vec{x}). Four vectors are indicated by Greek indices (x_μ), with summation over repeated indices assumed. The Minkowski metric $g_{\mu\nu} = g^{\mu\nu} = \text{diag}(+, -, -, -)$ is used throughout. Missing energy is denoted by \cancel{E}_t .

For reasons of brevity, unless otherwise stated, the term *electron* refers to both electrons (e^-) and positrons (e^+).

1.1.4 Machine energies

During the time this thesis was written, the expected startup beam energy was under discussion, and therefore the analyses contained within assume $\sqrt{s} = 14$ or 10 TeV.

2. Theoretical background

An accepted assumption about the universe is that the laws of physics should be independent of coordinate system. That is to say, if an experiment were performed at x_μ , the outcome should be the same as if the experiment were performed at x'_μ , where $x_\mu \rightarrow x'_\mu = \mathbf{T}x_\mu + \delta_\mu$, and \mathbf{T} is a transformation independent of x_μ . The case where \mathbf{T} is a rotation is analogous to the invariance of quantum mechanical observables under a phase transformation;

$$\hat{O}|\psi\rangle = \lambda|\psi\rangle = \lambda|\psi'\rangle \quad (2.1)$$

$$\text{where } |\psi\rangle \rightarrow |\psi'\rangle = e^{i\theta}|\psi\rangle. \quad (2.2)$$

Such invariance under a global transformation is termed a *global* symmetry.

This observation leads to an important conclusion. Special relativity limits the speed of information transfer in the universe to c . This implies that a transformation at point x_μ is, up to the light-cone limit, spacetime independent of a transformation at point x'_μ , and that therefore, following Eq. 2.1, there is no correct notion of absolute local phase for quantum wavefunctions.

Quantum mechanical systems must therefore be invariant under local transformations (yielding a *local* symmetry), where transformations become spacetime dependent. This is termed gauging the symmetry,

$$x_\mu \rightarrow x'_\mu = \mathbf{T}(x_\mu)x_\mu.$$

This chapter will explore the result of various symmetries as applied to Dirac fields, and how spontaneous symmetry breaking allows otherwise massless particles to acquire a mass term. This leads to the construction of the Standard Model, and discussion of problems and potential resolutions.

2.1 Groups, symmetries and invariance

2.1.1 $U(1)$ symmetry and quantum electrodynamics

Quantum electrodynamics (QED) is the prototypical gauge theory, and arises by requiring the Dirac free Lagrangian to be gauge invariant under local $U(1)$ transformations. Defining the gamma matrices, γ^μ , as having the fundamental property

$$\{\gamma^\mu, \gamma^\nu\} = 2g^{\mu\nu}\mathbf{I},$$

the Lagrangian density of a free electron,

$$\mathcal{L}_0 = \bar{\psi} (i\gamma^\mu \partial_\mu - m) \psi,$$

is invariant under a global transformation (Eq. 2.2). Promoting this global symmetry to a local symmetry requires replacing θ with $\theta(x_\mu)$. Such transformations form an Abelian Lie group, with generator 1, meaning any two elements commute,

$$[e^{i\theta(x_\mu)}, e^{i\theta(x'_\mu)}] = 0.$$

$\theta(x_\mu)$ therefore represents the arbitrary absolute phase of an electron wavefunction at every point in spacetime¹. As an arbitrary transformation can be created from the sum of many infinitesimal transformations, it is sufficient to only consider linear terms in the expansion of the group parameter,

$$e^{i\theta(x)} = 1 + i\theta(x) + \mathcal{O}(\theta^2(x)) \approx 1 + i\theta(x).$$

It therefore follows that

$$\begin{aligned} \psi(x) &\rightarrow e^{i\theta(x)}\psi(x) = (1 + i\theta(x))\psi(x), \\ \bar{\psi}(x) &\rightarrow e^{-i\theta(x)}\bar{\psi}(x) = (1 - i\theta(x))\bar{\psi}(x), \end{aligned} \quad (2.3)$$

and

$$\bar{\psi}(x)\partial_\mu\psi(x) \rightarrow \bar{\psi}(x)\partial_\mu\psi(x) + i\bar{\psi}[\partial_\mu\theta(x)]\psi(x). \quad (2.4)$$

It is clear from this behaviour that the Lagrangian is not invariant under local $U(1)$ gauge transformations. To restore the gauge invariance, the extra terms in Eq. 2.3 and 2.4 must be cancelled by performing a substitution of the derivative ∂_μ with the covariant derivative D_μ ,

$$D_\mu = \partial_\mu - ieA_\mu. \quad (2.5)$$

¹Spacetime indices will subsequently be dropped on x ($x_\mu \equiv x$) for clarity.

2.1 Groups, symmetries and invariance

A_μ is a new vector field which transforms as

$$A_\mu \rightarrow A_\mu + \frac{1}{e} [\partial_\mu \theta(x)].$$

To make A_μ a dynamical field, the field-strength tensor is defined as

$$F_{\mu\nu} = [\partial_\mu A_\nu] - [\partial_\nu A_\mu],$$

and a term for the free Lagrangian for A_μ is added to \mathcal{L}_0

$$\mathcal{L}_{A_\mu} = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu}.$$

This leads to the full gauge invariant QED Lagrangian²,

$$\mathcal{L}_{\text{QED}} = \bar{\psi} (i\gamma^\mu (\partial_\mu - ieA_\mu) - m) \psi - \frac{1}{4} F_{\mu\nu} F^{\mu\nu}, \quad (2.6)$$

where A_μ is identified as the photon field. When performing calculations in QED, it is usual to work in terms of the fine structure constant, α , rather than the gauge coupling, e . α is defined as

$$\alpha = \frac{e^2}{4\pi}, \quad (2.7)$$

where e is the charge of the electron.

In summary, imposing $U(1)$ local invariance directly results in the existence of the photon field. The covariant derivative contains the minimal coupling between the photon and electron fields, and is entirely determined by the gauge transformation properties of the electron field. Additionally, the photon does not carry a charge, and therefore there are no self-coupling terms for A_μ in \mathcal{L}_{QED} .

2.1.2 Feynman rules, regularisation and renormalisation

With a Lagrangian in hand, the Feynman rules for the theory can be constructed. Given two conjugate fields and a differential operator,

$$\bar{\psi}(x) \hat{O} \psi(y),$$

the propagator, D , for a particle excitation of the field $\psi(y)$ to travel to $\bar{\psi}(x)$, is given by the expression [1]

$$D(x-y) = \delta^4(x-y) \hat{O}^{-1}.$$

²The choice of $\frac{1}{4}$ in \mathcal{L}_{A_μ} is the only step motivated by a desired result; this choice ensures that the Euler-Lagrange equations of motion for A_μ result in the manifestly covariant form of Maxwell's equations ($\partial_\nu F^{\mu\nu} = e\bar{\psi}\gamma^\mu\psi$).




Figure 2.1: QED electron-photon vertex (a), and example one loop correction (b)

By moving into momentum space, with δ^4 given by

$$\delta^4(x-y) = \int \frac{d^4k}{(2\pi)^4} e^{ik(x-y)},$$

the momentum-space form of the propagator can be calculated. Applying this scheme [1, 2] to the QED Lagrangian, Eq. 2.6, results in the electron propagator

$$D_\psi = \left[\frac{i}{\gamma^\mu q_\mu - m + i\epsilon} \right],$$

and the photon propagator

$$D_{A_\mu} = \left[\frac{-ig^{\mu\nu}}{q^2 + i\epsilon} \right].$$

Vertex contributions to a scattering amplitude are simply determined by removing the fields from the relevant term. For example, the electron-photon interaction term $-ie\bar{\psi}A_\mu\psi$ contributes a factor of $-ie$ to a scattering amplitude.

With the Feynman rules in hand, the behaviour of certain processes can be considered. Taking the diagram for the electron-photon vertex, Fig. 2.1 (a), loop corrections permissible by the Feynman rules can be computed. One such correction is shown in Fig. 2.1 (b). This diagram contains, internally³, one photon propagator, behaving as $1/k^2$, and two electron propagators, behaving as $1/k$, where k is the internal momentum of the loop. As there are no constraints on the momentum carried in the loop, all momenta must be integrated over. This results in the contribution due to this diagram being

$$I \sim \int^\infty d^4k \frac{1}{k^4},$$

which diverges as $\lim_{k \rightarrow \infty} \ln k$.

The solution to these infinite corrections (termed ultraviolet divergences) is twofold. Firstly, the integral is regulated with some arbitrary cutoff, which can be taken to encode our ignorance of the high-energy behaviour of a theory. Secondly, the requirement that

³External fermion and boson lines carry spinor and polarisation indices, as required, but do not have associated propagators.

2.1 Groups, symmetries and invariance

all physical observables of the theory are invariant with respect to the energy scale at which it is being measured is enforced. This latter step is termed renormalisation.

Considering the vertex divergence described above, the vertex correction after regularisation is given by

$$I \sim \int^{\mu} d^4k \frac{1}{k^4} \sim \ln \mu. \quad (2.8)$$

The resulting vertex amplitude (ignoring all other possible diagrams that contribute) can therefore be written as

$$A \sim -ie(1 + f(\mu)), \quad (2.9)$$

where $f(\mu)$ is some function to be determined. Requiring the vertex amplitude to be invariant with respect to μ requires e to be dependent on the scale μ . This is enforced by replacing the measured quantity e with a bare and corrected term,

$$\begin{aligned} e &\rightarrow \hat{e}(\mu) = e + \delta_e(\mu), \\ A &\sim -i\hat{e}(\mu), \end{aligned}$$

and imposing the invariance condition

$$\begin{aligned} \frac{\partial}{\partial \ln \mu} A &= 0, \\ \frac{\partial}{\partial \ln \mu} [-i\hat{e}(\mu)] &= 0, \\ \frac{\partial}{\partial \ln \mu} [\hat{e}(\mu)] &= 0. \end{aligned}$$

The vertex amplitude correction, $f(\mu)$ in Eq. 2.9, is therefore absorbed into the redefinition of e .

In general terms, the behaviour of a coupling, g , with respect to the renormalisation scale, μ , is completely determined by the beta function, which is defined as

$$\beta(g) = \mu \frac{\partial g}{\partial \mu} = \frac{\partial g}{\partial \ln \mu}. \quad (2.10)$$

For the full QED case, the beta function is calculated (to one loop) as [2]

$$\beta(e) = \frac{e^3}{12\pi^3},$$

or in terms of the fine structure constant (using the relation in Eq. 2.7),

$$\beta(\alpha) = \frac{2\alpha^2}{3\pi}. \quad (2.11)$$

Solving Eq. 2.10 for Eq. 2.11 reveals the behaviour of the coupling α , which has the form

$$\alpha(\mu) = \frac{3\pi}{3c_0\pi - 2\ln\mu}, \quad (2.12)$$

where c_0 can be calculated by measuring α at some known scale. The remarkable thing about this result is that it describes exactly what is observed in nature - as energies increase (which is analogous to distance scales decreasing), the strength of the fine structure constant, and therefore the electromagnetic interaction, increases. In terms of regularisation and renormalisation, the combination of Eq. 2.8 with Eq. 2.12 allows calculations in the limit $\lim_{\mu \rightarrow \infty}$ to be performed; the renormalised value of e (or α) corrects for the emergent infinities in the loop integrals at any scale, regulated by μ .

In modern calculations, cutoff regularisation, as presented above, is very rarely used. The two most common schemes in use are Pauli–Villars regularisation, where new gauge fields are introduced to control loop divergencies, and dimensional regularisation, where the loop integral is performed in $d + \epsilon$ dimensions, and the limit $\lim_{\epsilon \rightarrow 0}$ is taken. There are also various renormalisation schemes, which define exactly which divergencies and constant terms are subtracted. The Modified Minimal Subtraction scheme, \overline{MS} , is often encountered, which absorbs both the divergent terms and a universal constant arising from the divergent diagrams.

In general, all physical terms (fields and couplings) in a theory must be renormalised in a consistent manner, with all diagrams at a given perturbation order considered. The example presented above only considers the coupling e . However, the intent is clear; regularisation and renormalisation allow calculations to be performed, and therefore cross sections for scattering processes computed, even when the theory appears to break down upon first inspection.

2.1.3 Generalised $SU(n)$ groups

By following a similar process to enforcing $U(1)$ gauge invariance on a free Dirac field, the results can be generalised to any $SU(n)$. A fundamental difference between $U(1)$ and $SU(n)$ gauge groups is that ψ becomes an n component multiplet of Dirac fields,

$$\psi = \begin{pmatrix} \psi_1 \\ \vdots \\ \psi_n \end{pmatrix}.$$

Transformations will have to operate in this multiplet space, and are therefore given by complex $n \times n$ matrices

$$\psi \rightarrow \mathbf{S}\psi, \quad \bar{\psi} \rightarrow \mathbf{S}^\dagger \bar{\psi}.$$

2.1 Groups, symmetries and invariance

An arbitrary transformation is given by

$$\mathbf{S} = e^{-i\theta^a(x)\mathbf{T}^a}, \quad (2.13)$$

where \mathbf{T}^a are the $n^2 - 1$ generators of the group, which are now $n \times n$ matrices, and $\theta^a(x)$ are the $SU(n)$ parameters. The $SU(n)$ label is due to enforcing the rules

$$\mathbf{S}\mathbf{S}^\dagger = \mathbf{I}, \quad \det[\mathbf{S}] = 1. \quad (2.14)$$

Group elements, and the generators themselves, no longer commute. The commutator of two generators is

$$[\mathbf{T}^a, \mathbf{T}^b] = if^{abc}\mathbf{T}^c,$$

where f^{abc} are the $SU(n)$ structure factors which entirely define the group algebra. The non-trivial commutator implies that such groups are non-Abelian. Analogous to Eq. 2.5, the covariant derivative is given by

$$\mathbf{D}_\mu = \mathbf{I}\partial_\mu - ig\mathbf{T}^a A_\mu^a.$$

There are now $n^2 - 1$ new gauge fields, one for each generator. These fields transform amongst themselves like the generators of the group

$$A_\mu^{a'} = A_\mu^a + f^{abc}\theta^b A_\mu^c - \frac{1}{g}[\partial_\mu\theta^a],$$

and the field strength tensor is defined as

$$F_{\mu\nu}^a = [\partial_\mu A_\nu^a] - [\partial_\nu A_\mu^a] + gf^{abc}A_\mu^b A_\nu^c.$$

The generalised Lagrangian is therefore given by

$$\mathcal{L}_{SU(n)} = \bar{\psi}_i (i\gamma^\mu (\mathbf{I}\partial_\mu - ig\mathbf{T}^a A_\mu^a) - \mathbf{I}m) \psi_i - \frac{1}{4}F_{\mu\nu}^a F^{a\mu\nu}, \quad (2.15)$$

where summation over component fields of ψ is assumed. Such a Lagrangian is invariant under a gauge transformation given by Eq. 2.13. The result of such a transformation is to rotate the fields amongst themselves; this is an internal symmetry and is fundamental in, for example, the weak interaction where fermions are transformed within generations, mediated by gauge interactions. Also, due to the non-Abelian nature of $SU(n)$ groups, the $-\frac{1}{4}F_{\mu\nu}^a F^{a\mu\nu}$ term in Eq. 2.15 is not a free Lagrangian as it contains terms of the form

$$gf^{abc}[\partial_\mu A_\nu^a]A^{b\mu}A^{c\nu} - g^2 f^{abc}f^{ade}A_\mu^b A_\nu^c A^{d\mu}A^{e\nu},$$

which implies that the gauge fields interact via three and four point couplings.

2.1.4 Spontaneous symmetry breaking

In both the $U(1)$ and $SU(n)$ cases, the Lagrangians do not contain terms of the form

$$-\frac{1}{2}m^2 A_\mu A^\mu$$

and therefore the gauge fields are massless. In the $U(1)$ case, where A_μ represents the photon, this is expected. However, when attempting to model electroweak interactions where not all of the gauge fields are experimentally observed to be massless, this is clearly problematic.

The explicit introduction of mass terms into a Lagrangian is not allowed as this breaks gauge invariance and renders the theory non-renormalisable. However, the spontaneous breaking of the gauge symmetry can result in massive interacting particles. As an example, a $U(1)$ gauge theory with a single gauge field is considered. The Lagrangian is given by

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu},$$

where

$$F_{\mu\nu} = [\partial_\mu A_\nu] - [\partial_\nu A_\mu].$$

To allow A_μ to acquire a mass, avoiding the addition of an explicit mass term, the model can be extended with the addition of a complex scalar field which couples to the gauge field. This results in the gauge invariant Lagrangian

$$\mathcal{L} = |D_\mu \Phi|^2 - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} - V(\Phi), \quad (2.16)$$

with

$$D_\mu = \partial_\mu - igA_\mu, \quad \text{and} \\ V(\Phi) = \lambda|\bar{\Phi}\Phi|^2 - \mu^2\bar{\Phi}\Phi.$$

Assuming that $\lambda > 0$ to provide a minimum energy bound, and taking $\mu^2 > 0$, the minimum of the potential can be seen to be not at $\Phi = 0$, but at

$$\Phi = e^{i\theta} \sqrt{\frac{\mu^2}{2\lambda}} = e^{i\theta} \frac{v}{\sqrt{2}},$$

for an arbitrary choice of θ between 0 and 2π for each x_μ . The vacuum is therefore degenerate, and the vacuum expectation value (VEV) is given by

$$\langle 0|\Phi|0\rangle = \frac{v}{\sqrt{2}} \neq 0.$$

By fixing the gauge such that $\theta(x)$ is chosen for which Φ is real, the field Φ may then be

2.2 The Standard Model

expanded about the physical vacuum as

$$\Phi = \frac{1}{\sqrt{2}}(v + H). \quad (2.17)$$

It is this arbitrary choice of θ which spontaneously breaks the symmetry. Substituting the potential after symmetry breaking back into the Lagrangian (Eq. 2.16) reveals separable free and interaction parts

$$\begin{aligned} \mathcal{L}_{\text{free}} &= \frac{1}{2}\partial_\mu H\delta^\mu H - \lambda v^2 H^2 + \frac{1}{2}g^2 v^2 A_\mu A^\mu - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}, \\ \mathcal{L}_{\text{int}} &= g^2 v A_\mu A^\mu H + \frac{1}{2}g^2 A_\mu A^\mu H^2 - \lambda v H^3 - \frac{\lambda}{4}H^4. \end{aligned}$$

From $\mathcal{L}_{\text{free}}$, a vector A_μ with mass gv and a scalar H with mass $v\sqrt{2\lambda}$ can be identified. \mathcal{L}_{int} reveals four vertices: AAH , $AAHH$, HHH and $HHHH$ with respective couplings $g^2 v$, $\frac{1}{2}g^2$, λv and $\frac{1}{4}\lambda$.

What is remarkable is that without the explicit addition of mass terms, the original massless gauge vector boson has acquired a mass term, and a new massive scalar H has appeared. This is termed the Higgs mechanism, and has been shown to be completely renormalisable.

2.2 The Standard Model

The Standard Model is the current best description of the physics of fundamental particles and their interactions. It is a gauge theory, with structure $SU(3) \otimes SU(2)_L \otimes U(1)_Y$. The particle content is shown in Fig. 2.2; how this particle content arises is partly by construction (the number of fermion families, for example) and partly due to the underlying gauge theory (the gauge bosons). The theoretical construction of the model for one fermion generation is detailed below. The theory is logically extendable to the three generations observed in nature.

2.2.1 Electroweak model of leptonic interactions

Following experimental observations, the symmetry group for electroweak interactions is chosen to be

$$SU(2)_L \otimes U(1)_Y,$$

where $U(1)_Y$ is no longer the gauge group of QED, but rather it is associated with the weak hypercharge, defined as

$$Y = Q - I_3, \quad (2.18)$$

where I_3 is the third component of weak isospin.

Quarks	$\begin{matrix} +2/3 & 1/2 \\ \mathbf{u} \\ \sim 3 \text{ MeV} \end{matrix}$	$\begin{matrix} +2/3 & 1/2 \\ \mathbf{c} \\ 1.27 \text{ GeV} \end{matrix}$	$\begin{matrix} +2/3 & 1/2 \\ \mathbf{t} \\ 171 \text{ GeV} \end{matrix}$	$\begin{matrix} 0 & 1 \\ \mathbf{Y} \\ 0 \end{matrix}$	$\begin{matrix} \pm 1 & 1 \\ \mathbf{W}^\pm \\ 80.4 \text{ GeV} \end{matrix}$	Gauge bosons	
	$\begin{matrix} -1/3 & 1/2 \\ \mathbf{d} \\ \sim 5 \text{ MeV} \end{matrix}$	$\begin{matrix} -1/3 & 1/2 \\ \mathbf{s} \\ 104 \text{ MeV} \end{matrix}$	$\begin{matrix} -1/3 & 1/2 \\ \mathbf{b} \\ 4.2 \text{ GeV} \end{matrix}$	$\begin{matrix} 0 & 1 \\ \mathbf{Z} \\ 91.2 \text{ GeV} \end{matrix}$	$\begin{matrix} 0 & 1 \\ \mathbf{\gamma} \\ 0 \end{matrix}$		
Leptons	$\begin{matrix} 0 & 1/2 \\ \mathbf{\nu}_e \\ < 2 \text{ eV} \end{matrix}$	$\begin{matrix} 0 & 1/2 \\ \mathbf{\nu}_\mu \\ < 2 \text{ eV} \end{matrix}$	$\begin{matrix} 0 & 1/2 \\ \mathbf{\nu}_\tau \\ < 2 \text{ eV} \end{matrix}$	$\begin{matrix} 0 & 0 \\ \mathbf{H} \\ > 114 \text{ GeV} \end{matrix}$			Scalar bosons
	$\begin{matrix} -1 & 1/2 \\ \mathbf{e} \\ 0.511 \text{ MeV} \end{matrix}$	$\begin{matrix} -1 & 1/2 \\ \mathbf{\mu} \\ 106 \text{ MeV} \end{matrix}$	$\begin{matrix} -1 & 1/2 \\ \mathbf{\tau} \\ 1.78 \text{ GeV} \end{matrix}$				

Figure 2.2: Particle content of the Standard Model, with associated charge (top left), spin (top right) and mass (bottom). The Higgs is unobserved.

The symmetry group has four generators (normalised according to $\text{Tr}(\mathbf{T}^a \mathbf{T}^b) = \frac{1}{2} \delta_{ab}$),

$$\mathbf{T}^1 = \frac{1}{2} \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \mathbf{T}^2 = \frac{1}{2} \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \mathbf{T}^3 = \frac{1}{2} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad \mathbf{Y} = \mathbf{Y} \mathbf{I}, \quad (2.19)$$

where I_3 can be identified with the generator \mathbf{T}^3 . \mathbf{T}^a are related to the Pauli matrices, such that $\mathbf{T}^a = \frac{\boldsymbol{\tau}^a}{2}$, and Y takes the values $Y(l_L) = -\frac{1}{2}$, $Y(e_R) = -1$. With these assignments, the charge operator, \mathbf{Q} , from Eq. 2.18 is given by

$$\mathbf{Q} = \mathbf{Y} + \mathbf{T}^3 = \begin{pmatrix} 0 & 0 \\ 0 & -1 \end{pmatrix}, \quad (2.20)$$

which has the expected eigenvalues as applied to the $SU(2)$ left handed lepton doublet and e_R singlet:

$$\begin{aligned} \mathbf{Q} \begin{pmatrix} \nu_L \\ 0 \end{pmatrix} &= 0 \begin{pmatrix} \nu_L \\ 0 \end{pmatrix}, \\ \mathbf{Q} \begin{pmatrix} 0 \\ e_L \end{pmatrix} &= -1 \begin{pmatrix} 0 \\ e_L \end{pmatrix}, \\ \mathbf{Q} \begin{pmatrix} e_R \end{pmatrix} &= -1 \begin{pmatrix} e_R \end{pmatrix}. \end{aligned}$$

The covariant derivative for the left-handed fields is defined as

$$\mathbf{D}_\mu^L = \mathbf{I} \partial_\mu - ig \frac{\boldsymbol{\tau}^a}{2} W_\mu^a - \mathbf{I} ig' Y(l_L) B_\mu, \quad (2.21)$$

where $a = 1, 2, 3$. As $SU(2)$ invariance is imposed on the right-handed electron, its co-

2.2 The Standard Model

variant derivative is

$$D_\mu^R = \partial_\mu - ig'Y(e_R)B_\mu.$$

The gauge fields therefore have standard kinetic energy terms

$$\mathcal{L}_{KE} = -\frac{1}{4}W_{\mu\nu}^a W^{a\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu},$$

where the field strength tensors are given by

$$\begin{aligned} W_{\mu\nu}^a &= [\partial_\mu W_\nu^a] - [\partial_\nu W_\mu^a] + g\epsilon^{abc}W_\mu^b W_\nu^c, \\ B_{\mu\nu} &= [\partial_\mu B_\nu] - [\partial_\nu B_\mu]. \end{aligned}$$

In the above, the gauge couplings g and g' have been introduced. These are related to the weak mixing angle and electron charge by

$$\begin{aligned} \tan \theta_W &= \frac{g'}{g}, \\ g &= \frac{e}{\sin \theta_W}, \\ g' &= \frac{e}{\cos \theta_W}. \end{aligned} \tag{2.22}$$

It is usual to see θ_W expressed as $\cos \theta_W = \frac{M_W}{M_Z}$ (the reason for which will become apparent from Eq. 2.27), giving $\sin^2 \theta_W = 0.231^4$.

The four gauge fields $W_\mu^{1,2,3}$, B_μ , do not correspond directly to physical gauge bosons, whereas linear superpositions of them do. The reasons for this will become apparent after electroweak symmetry breaking.

2.2.2 Electroweak symmetry breaking

Following from the Abelian Higgs model, in $SU(2) \otimes U(1)$ a scalar $SU(2)$ doublet, Φ is proposed,

$$\Phi = \begin{pmatrix} \phi^0 \\ \phi^+ \end{pmatrix}.$$

The scalar potential is again given by

$$V(\Phi) = \lambda |\bar{\Phi}\Phi|^2 - \mu^2 \bar{\Phi}\Phi, \tag{2.23}$$

where $\lambda > 0$ and $\mu^2 > 0$. With these conditions, the minimum of the potential is not at $\Phi = 0$. Making a global gauge transformation such that the vacuum expectation value

⁴Here the on-shell renormalisation scheme is implicit. In general, Eq. 2.22 defines the weak mixing angle. For instance, in the \overline{MS} scheme used in §3, $\sin^2 \theta_W(\mu = M_Z) = 0.231$.

(VEV) is real and in the first component yields

$$\langle 0|\Phi|0\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} v \\ 0 \end{pmatrix}.$$

Applying the $SU(2)$ form of the electromagnetic charge from Eq. 2.20 to this VEV gives

$$\mathbf{Q}\langle 0|\Phi|0\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 0 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} v \\ 0 \end{pmatrix} = 0,$$

which implies that $U(1)_{\text{EM}}$ remains unbroken. This results in the desirable symmetry breaking pattern

$$SU(2)_L \otimes U(1)_Y \rightarrow U(1)_{\text{EM}},$$

such that the photon remains massless. It is important to note that $U(1)_{\text{EM}}$ is not the gauge group associated with QED (§2.1.1), rather it is the unbroken component of the electroweak symmetry group after spontaneous symmetry breaking has occurred.

The physical implication of this is that while Dirac free Lagrangians are required to be gauge invariant under $SU(2)_L \otimes U(1)_Y$ transformations, the ground state of the universe is not entirely symmetric under such transformations, and the resulting spontaneously broken symmetry leaves an unbroken $U(1)_{\text{EM}}$ symmetry.

Expanding around the vacuum, as for the $U(1)$ case, Ψ is given by

$$\Psi = \frac{1}{\sqrt{2}} \begin{pmatrix} v + H \\ 0 \end{pmatrix}, \quad (2.24)$$

and contributes to the Lagrangian the terms

$$\mathcal{L}_{SB} = |\mathbf{D}_\mu \Phi|^2 - V(\Phi).$$

Defining W_μ^\pm as

$$W_\mu^\pm \equiv \frac{1}{\sqrt{2}} (W_\mu^1 \mp iW_\mu^2), \quad (2.25)$$

the term $\mathbf{D}_\mu \Phi$ is expanded as

$$\mathbf{D}_\mu \Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} \partial_\mu - \frac{i}{2} (gW_\mu^3 - g'B_\mu) & -\frac{ig}{\sqrt{2}} W_\mu^+ \\ -\frac{ig}{\sqrt{2}} W_\mu^- & \partial_\mu + \frac{i}{2} (gW_\mu^3 + g'B_\mu) \end{pmatrix} \begin{pmatrix} v + H \\ 0 \end{pmatrix}.$$

$|\mathbf{D}_\mu \Phi|^2$ therefore results in kinetic and interaction terms in the Lagrangian. The kinetic terms are given by

$$\mathcal{L}_{H,\text{kinetic}} = \frac{1}{2} (\partial_\mu H)^2 + \frac{g^2 v^2}{4} W_\mu^+ W^{-\mu} + \frac{v^2}{8} (gW_\mu^3 - g'B_\mu)^2, \quad (2.26)$$

which show that the fields W_μ^3 and B_μ mix. Separate physical gauge fields can be realised

2.2 The Standard Model

by defining Z_μ and A_μ as linear superpositions of W_μ^3 and B_μ ,

$$\begin{aligned} Z_\mu &\equiv \cos \theta_W W_\mu^3 - \sin \theta_W B_\mu, \\ A_\mu &\equiv \cos \theta_W B_\mu + \sin \theta_W W_\mu^3, \end{aligned}$$

where excitations of Z_μ can be identified as Z^0 bosons, and excitations of A_μ with photons, γ . The W_μ^\pm terms, from Eq. 2.25, already correspond to the physical W^+ and W^- bosons. In terms of these linear superpositions, Eq. 2.26 therefore contains mass terms for the physical gauge fields

$$\mathcal{L}_{H,\text{kinetic}} = \frac{1}{2} (\partial_\mu H)^2 + \frac{g^2 v^2}{4} W_\mu^+ W^{-\mu} + \frac{g^2 v^2}{8 \cos^2 \theta_W} Z_\mu Z^\mu,$$

with boson masses given by

$$M_{W^\pm} = \frac{1}{2} g v, \quad M_{Z^0} = \frac{1}{2} \frac{g v}{\cos \theta_W}. \quad (2.27)$$

The H mass term arises from the substitution of Eq. 2.24 into the potential (Eq. 2.23), yielding

$$\mu^2 H^2.$$

The H mass is, therefore,

$$M_H = \sqrt{2} \mu \equiv v \sqrt{2\lambda}. \quad (2.28)$$

The interaction terms involving the gauge fields W_μ^\pm , Z_μ and the scalar H , written in terms of the gauge boson masses, are given by

$$\mathcal{L}_{H,\text{int}} = \frac{2M_W^2}{v} W_\mu^+ W^{-\mu} H + \frac{M_W^2}{v^2} W_\mu^+ W^{-\mu} H H + \frac{M_{Z^0}^2}{v} Z_\mu Z^\mu H + \frac{M_{Z^0}^2}{2v^2} Z_\mu Z^\mu H H.$$

The photon field, A_μ , acquires no mass or interaction terms.

2.2.3 Quantum chromodynamics

It is apparent that in the quark sector, an additional charge is carried. By studying the evolution of the ratio R

$$R = \frac{\sigma(e^+e^- \rightarrow \mu^+\mu^-)}{\sigma(e^+e^- \rightarrow \text{hadrons})}$$

with collision energy, it is apparent that quarks are produced as unmixed states of three colour charges. This leads to quantum chromodynamics (QCD), the theory of strong interactions, being described by an additional $SU(3)$ non-Abelian gauge theory. The quarks

are therefore described by a three-component field

$$\psi_q = \begin{pmatrix} \psi_R \\ \psi_G \\ \psi_B \end{pmatrix}.$$

The generators of the group are related to the Gell-Mann matrices, $\mathbf{T}^a = \frac{\lambda^a}{2}$. This construction results in 8 gauge fields, representing the massless gluons. The gauge coupling constant, g , is taken to be the strong coupling g_s .

The particle content is extended to contain the quarks,

$$l_L \equiv \begin{pmatrix} \nu \\ e \end{pmatrix}_L, \quad l_R \equiv e_R, \quad q_L \equiv \begin{pmatrix} u \\ d \end{pmatrix}_L, \quad q_R \equiv u_R, d_R,$$

where under $SU(2)$ only the left-handed quarks are transformed by the gauge interactions. The Standard Model gauge group is now

$$SU(3) \otimes SU(2)_L \otimes U(1)_Y.$$

It should be noted that the leptons form singlets under $SU(3)$, and therefore do not couple to the gluons whereas the quarks form colour triplets, as above, and therefore couple to the new gluon fields. The covariant derivative for the left-handed quarks is expanded to

$$\mathbf{D}_\mu^{qL} = \mathbf{I}\partial_\mu - ig_s \frac{\lambda^a}{2} G_\mu^a - ig \frac{\tau^a}{2} W_\mu^a - \mathbf{I}ig'Y B_\mu, \quad (2.29)$$

and the usual kinetic terms are added

$$\begin{aligned} \mathcal{L}_{\text{QCD,Kinetic}} &= -\frac{1}{4} G_{\mu\nu}^a G^{a\mu\nu}, \quad \text{where} \\ G_{\mu\nu}^a &= [\partial_\nu G_\mu^a] - [\partial_\mu G_\nu^a] + gf^{abc} G_\mu^b G_\nu^c. \end{aligned}$$

In analog to the electroweak case, it is usual to work in terms of α_s , which is related to g_s by

$$\alpha_s(\mu) = \frac{g_s(\mu)^2}{4\pi}.$$

As α_s is large, it may be expected that perturbative calculations in the QCD sector would not be possible. However, they are possible due to behaviour of α_s with respect to $\mu = Q^2$, the squared momentum transfer scale. The one loop β function for n_f active quark flavours is given by [2]

$$\beta(\alpha_s) = - \left(11 - \frac{2n_f}{3} \right) \frac{\alpha_s^2}{2\pi}. \quad (2.30)$$

The running of α_s , as described by Eq. 2.30 is shown in Fig 2.3 to one loop, and assuming 5 quark flavours with mass threshold below the momentum scale Q . This is valid as following cutoff renormalisation, Q is essentially the upper momentum limit in the

2.2 The Standard Model

loop diagrams. Above momentum scales of around 2 GeV, α_s is small enough such that perturbative calculations should be reliable.




Figure 2.3: The running of $\alpha_s(Q^2)$ at leading order

It should also be noted that the running of α_s can be taken to explain why bare quarks or gluons have not been observed. As low energies imply large distances, α_s is strong at these large distances. This observation, coupled with Eq. 2.30, limits the number of quark flavours to $n_f \leq 16$.

2.2.4 Fermion masses and gauge interactions

The addition of explicit fermion mass terms is not allowed as such a term would mix left and right-handed fermions. However, the $SU(2)$ Higgs doublet (Eq. 2.24) allows a gauge invariant term to be introduced. The Yukawa interactions are given by

$$\mathcal{L}_{\text{Yukawa}} = -Y_e \epsilon_{ij} \Psi^i \bar{l}_L^j e_R + \text{h.c.}, \quad (2.31)$$

where

$$\epsilon_{ij} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

After symmetry breaking, Eq. 2.31 becomes

$$\mathcal{L}_{\text{Yukawa}} = -\frac{Y_e}{\sqrt{2}} (v + H) (\bar{e}_L e_R + \bar{e}_L^\dagger e_R^\dagger). \quad (2.32)$$

The electron mass is therefore given by

$$m_e = \frac{Y_e}{\sqrt{2}} v = \sqrt{2} \frac{Y_e M_W}{g},$$

and the neutrino remains massless. Neutrino masses can be accommodated in the Standard Model, but will not be discussed here.

Eq. 2.32 applied to the quark doublet results in a mass term for the d quark, but leaves u massless. This can be addressed by the addition of another invariant term,

$$\mathcal{L}_{\text{Yukawa},u} = -Y_u \Phi^\dagger \bar{q}_L u_R + \text{h.c.},$$

which leads to the u quark mass

$$m_u = \frac{Y_u}{\sqrt{2}} v = \sqrt{2} \frac{Y_u M_W}{g}.$$

The full expansion of Eq. 2.32 also yields an interaction term,

$$\frac{gm_e}{2M_W} H \bar{e}_L e_R, \quad (2.33)$$

and similar in the quark sector, where the coupling of the Higgs to any fermion is clearly proportional to the fermion mass.

Gauge interaction terms for the fermions arise from the expansion of the kinetic part of the Lagrangian with the full covariant derivatives for left- and right-handed fields. This gives

$$\mathcal{L}_{\text{Kinetic},f} = i \bar{l}_L \gamma^\mu \mathbf{D}_\mu l_L + i \bar{e}_R \gamma^\mu D_\mu e_R + i \bar{q}_L \gamma^\mu \mathbf{D}_\mu q_L + i \bar{d}_R \gamma^\mu \mathbf{D}_\mu d_R + i \bar{u}_R \gamma^\mu \mathbf{D}_\mu u_R, \quad (2.34)$$

where \mathbf{D}_μ (D_μ) contains the gauge fields for the allowed interactions in each case.

Defining the projection operators, which project out the left- and right-handed component fields, as

$$\begin{aligned} P_L &= \frac{1}{2} (1 - \gamma^5), \\ P_R &= \frac{1}{2} (1 + \gamma^5), \end{aligned}$$

allows the fermion fields to be decomposed into chiral components,

$$f_{\text{chiral}} = f_L + f_R = (P_L + P_R) f.$$

In terms of these projection operators, and taking the fermion fields to be chiral decompositions, the interaction terms from Eq. 2.34 between the fermion and Z_μ fields are therefore given by

$$\frac{g}{4 \cos \theta_W} \bar{e} (\gamma^\mu (1 - \gamma^5) - 4 \sin^2 \theta_W \gamma^\mu) e Z_\mu$$

for the left- and right-handed electron, and

$$- \frac{g}{2 \cos \theta_W} \bar{q}_i (T_i^3 \gamma^\mu (1 - \gamma^5) - 2Q_i \sin^2 \theta_W \gamma^\mu) q_i Z_\mu \quad (2.35)$$

2.3 Problems with the Standard Model

for the quarks, where T_i^3 and Q_i are the third component of weak isospin and the electric charge for quark q_i . These interaction terms (and similarly for those involving W^\pm) treat right- and left-handed components differently, which is the cause of parity violation in weak interactions.

In summary, it is possible to produce, through the annihilation of a quark with its anti-partner, a massive Z^0 which can subsequently decay to an e^+e^- pair. Such interactions, in the context of physics beyond the Standard Model, form the basis for this thesis.

2.3 Problems with the Standard Model

2.3.1 Particle content and free parameters

While the gauge sector particle content can be understood in terms of their generating groups, the number of fermions present in the Standard Model is entirely arbitrary. An outstanding problem is to determine if there are further fermion generations, or if the currently known generations can be explained as composite arrangements of currently unidentified particles. Some interactions between the Standard Model particles are also not currently understood. The most obvious missing component is that of gravitation. So far, all attempts to realise a renormalisable quantum field theory of gravity have failed.

While the Standard Model is highly predictive in terms of gauge boson masses, cross sections and decay rates, such predictions depend on 19 free parameters which must be measured. These free parameters, relating to the Higgs sector, fermion masses, quark mixing, gauge coupling strengths and a QCD phase, should ideally be predictions themselves.

2.3.2 The hierarchy problem

The Planck scale is the only mass scale that can be constructed from the fundamental constants,

$$M_P = \sqrt{\frac{\hbar c}{G_N}} = 1.2 \times 10^{19} \text{ GeV}, \quad (2.36)$$

where G_N is Newton's constant. From Eq. 2.27, it is clear that the electroweak breaking scale is of order $v \sim M_W \sim \mathcal{O}(100)$ GeV. This can be restated as the observation that the weak force is significantly stronger than gravity. With the fundamental scale for gravitational interactions given by Eq. 2.36, the equivalent for the weak interaction are the Fermi

constant and mass scale, defined as

$$G_F = \frac{\pi\alpha}{\sqrt{2} M_W^2 \sin^2 \theta_W},$$

$$M_F = \sqrt{\frac{(\hbar c)^3}{G_F}} = 2.9 \times 10^2 \text{ GeV}.$$

M_F is sensitive to corrections to the gauge boson mass, M_W , which become large at high momentum scales. This implies that G_F appears unnaturally large, unless there is a fine-tuning between the bare and renormalised gauge boson masses.

As the Higgs mechanism is responsible for the generation of mass terms, this hierarchy problem can be restated in terms of the Higgs self-coupling. As an example, fermion masses are logarithmically divergent [3],

$$\delta m_f \simeq \frac{3\alpha}{4\pi} m_f \ln(\Lambda^2/m_f^2),$$

where Λ is an ultraviolet cutoff at which new physics is expected. Such corrections are small, even for $\Lambda \sim M_P$. However, the 1-loop corrections to scalar masses are given by

$$\delta m_H^2 \simeq g_f^2, g^2, \lambda \int d^4k \frac{1}{k^2} \sim \mathcal{O}\left(\frac{\alpha}{4\pi}\right) \Lambda^2,$$

where g_f is the fermion coupling, g^2 the gauge boson loop coupling and λ the quartic scalar coupling. From Eq. 2.27 and Eq. 2.28, and assuming that λ is of order 1, it is expected that $m_H \sim M_W$. However, if new physics is present such that the relevant scale is $\Lambda \gg M_W$, the correction δm_H^2 becomes much larger than $\mathcal{O}(M_W)$.

One way around this problem is to tune the bare mass m_H such that it is large and negative, to completely cancel the 1-loop corrections. For $\Lambda \sim \mathcal{O}(M_W)$, this bare mass must be tuned to 32 significant figures. Additionally, the corrections of higher order must also be cancelled. While this is technically possible, it is not a desirable solution to the problem and therefore motivates searches for physical solutions.

2.4 Beyond the Standard Model

2.4.1 Supersymmetry

Supersymmetry (SUSY) is one proposed solution to the hierarchy problem which postulates that there are new particles with similar mass and equal couplings to those in the standard model [3, 4]. These particles are related by an operator Q , such that $Q|\text{fermion}\rangle = |\text{boson}\rangle$ and $Q|\text{boson}\rangle = |\text{fermion}\rangle$, and therefore the superpartners differ from their standard model counterparts by a half-unit of spin. Therefore, all fermions have bosonic SUSY partners, and vice-versa.

2.4 Beyond the Standard Model

As fermion loops possess a relative minus sign, the Higgs mass divergence is cancelled as

$$\delta m_H^2 \simeq \mathcal{O}\left(\frac{\alpha}{4\pi}\right) (\Lambda^2 + m_B^2) - \mathcal{O}\left(\frac{\alpha}{4\pi}\right) (\Lambda^2 + m_F^2) = \mathcal{O}\left(\frac{\alpha}{4\pi}\right) (m_B^2 - m_F^2).$$

As these supersymmetric particles have not been found, their masses must be higher than the known Standard Model particles. The mass divergence can be controlled if

$$|m_B^2 - m_F^2| \lesssim 1 \text{ TeV}^2,$$

suggesting that if SUSY exists, it may be accessible at the LHC.

In certain SUSY models, for example the Minimally Supersymmetric Standard Model (MSSM), the Lagrangian couplings unify at a scale of 10^{16} GeV (Fig. 2.5). In the Standard Model, these couplings do not unify (Fig. 2.4). This is perhaps a sign that SUSY may play a role in some unified theory.




Figure 2.4: The running of gauge couplings in the Standard Model [5, 6]




Figure 2.5: The running of gauge couplings in the MSSM [5, 6]

2.4.2 Large extra dimensions

The hierarchy between the Planck and weak scales ($G_F/G_N \approx 10^{-32}$) can be reduced if gravity is allowed to propagate in one or more extra spatial dimensions. The ADD model [7] suggests that the Planck scale is itself not fundamental, but rather is the consequence of the large size of the extra dimensions. The relationship between the effective 4-dimensional M_{Pl} and the proposed $(4+d)$ -dimensional M_{Pl} can be derived as

$$\left[M_{Pl}^{(4)}\right]^2 \sim R^d \left[M_{Pl}^{(4+d)}\right]^{(2+d)}.$$

If the fundamental scale is related to the break down of the Standard Model (~ 1 TeV), and assuming at least 2 extra dimensions⁵, the size of these dimensions, R , must be of the order 1 mm or smaller. Gravity at such scales has not yet been measured.

⁵The case of $d = 1$ is ruled out as gravity would be affected at scales already measured.

In such theories, gravitons can exist as excited Kaluza-Klein (KK) states in the extra dimensions. Such states form an effective 4-dimensional tower of masses given by [8]

$$M_j = \sqrt{m_0^2 + \frac{j^2}{R^2}}, \quad (2.37)$$

where m_0 is the unexcited mass. From Eq. 2.37, it is clear that a large continuum of states can exist when R is large. The exact phenomenology of the states depends on the production and decay cross sections, but in general leads to subtle changes in invariant mass and transverse momentum cross sections as compared to the Standard Model due to exchange of towers of virtual KK gravitons. Current limits set $M_{Pl}^{(d+4)} \gtrsim 1 \text{ TeV}$ for $d \geq 2$ [9].

2.4.3 TeV^{-1} scale extra dimensions

The unification of the running Lagrangian couplings is desirable in a unified theory. The TeV^{-1} scale extra dimensional model [10, 11], where the gauge fields are allowed to propagate in the bulk, was proposed as a mechanism whereby the coupling unification can happen at a scale $\ll M_{Pl}$. Above a scale of $M_C = 1/R$, the gauge KK states contribute to the running of the couplings, which modify them from a logarithmic running to a power-law running, as shown in Fig. 2.6.




Figure 2.6: Power-law running of gauge couplings with TeV^{-1} scale extra dimensions [11]

The KK states of the gauge bosons interfere with their Standard Model counterparts, which provides phenomenology accessible at the LHC. Current experimental limits set the scale $M_C \geq 6.8 \text{ TeV}$ [12].

2.4.4 Warped extra dimensions

One argument against the existence of large extra dimensions is that it introduces a new, albeit smaller, hierarchy between the compactification scale $M_C = 1/R$ and the weak scale. The RS family of models [13] attempt to generate the mass hierarchy from a small

2.4 Beyond the Standard Model

input parameter. This is accomplished through the existence of a single new extra dimension, where the 5-dimensional metric is

$$ds_5^2 = \frac{1}{k^2 z^2} (g_{\mu\nu} dx^\mu dx^\nu - dz^2).$$

The coordinate of the 5th dimension, z , is bounded at the Planck scale at $z = 1/k$ and at the TeV scale at $z = 1/Wk$. In this construction, $k \sim \mathcal{O}(M_{Pl})$, W describes the warping of the extra dimension, and $Wk \sim \mathcal{O}(1 \text{ TeV})$. The size of the extra dimension, $-\ln(W)/k$, is small. This AdS_5 space can be interpreted as containing a scaled Minkowski metric at every point along the new dimension.

In this set-up, and by requiring the effective 4-dimensional Planck scale at $z = 1/T$ to equal the weak scale, the hierarchy between the 5-dimensional Planck scale and the compactification scale, $M_c = 1/R$ is of order 50 as opposed to M_{Pl}/TeV . Such a model presents interesting phenomenology, namely TeV scale individual KK graviton excitations, which can decay to fermion pairs.

If the Standard Model gauge fields are allowed to propagate in the bulk, the Higgs quadratic divergence is cancelled not through the warping of the Planck scale, but due to the KK modes entering into the available loop diagrams and directly canceling the divergences. A further result is that the running of the gauge couplings is logarithmic, unlike in flat extra dimensional models. Additionally, these couplings are unified at $\Lambda = k$ [14].

3. The neutral-current Drell-Yan process in the high invariant mass region

The neutral-current Drell-Yan process [15] is an important background to spectroscopic searches for new physics in di-lepton channels. A detailed understanding of the best current theoretical knowledge, in both the QCD and electroweak (EWK) sectors, is required to ensure that any measured differences from the predicted Drell-Yan spectrum in the Standard Model are not due to high order contributions that have been neglected. This is of particular interest in searches for, for example, extra dimension models which predict subtle, broad changes to differential cross sections [9, 12].

A study combining the best available $\mathcal{O}(\alpha_s)$ QCD corrections with $\mathcal{O}(\alpha)$ EWK and higher order QED corrections due to photon showers is presented from a phenomenological viewpoint. Photon-induced processes, which contribute to the $pp \rightarrow Z^0 \rightarrow l^+l^-(n\gamma)+X$ signature, are also included. It is found that the QCD and EWK corrections partially cancel, resulting in a 10%–5% change in $\frac{d\sigma}{dM_{l^+l^-}}$ in the mass region 200–1500 GeV.

3.1 Introduction

3.1.1 Experimental motivation

The most general type of experiment at a collider is to measure differential cross sections and compare these to the predictions of the Standard Model (SM). Deviations from the SM could be indicative of new physics. As such, a detailed understanding of the SM predictions is vital.

One potential scenario, outlined in §2.4 is that of extra dimensions. To further illustrate the subtle changes such theories can cause, the case of Kaluza-Klein excitations of SM gauge bosons, when they are allowed to propagate in a TeV scale extra dimension, is considered. Following [12], the differential cross section for lepton pair production in the

presence of virtual KK Z/γ exchange due to such an extra dimension is given by

$$\frac{d\sigma}{dM_{ll}} = K \frac{M_{ll}^3}{72\pi s} \int \sum_q \left[f_q(x_1) f_{\bar{q}}(x_2) \sum_{\alpha=L,R} \sum_{\beta=L,R} \left(|M_{\alpha\beta}^{eq}(\hat{s})|^2 \right) \right] dy \quad (3.1)$$

where $\hat{s} = M_{ll}^2$, \sqrt{s} the center-of-mass collision energy, $x_{1,2} = (M_{ll}/\sqrt{s}) e^{\pm y}$ and $M_{\alpha\beta}^{eq}$ is the partial matrix element,

$$M_{\alpha\beta}^{eq}(\hat{s}) = e^2 \left(\frac{Q_e Q_q}{\hat{s}} + \frac{g_\alpha^e g_\beta^q}{\sin^2 \theta_W \cos^2 \theta_W} \frac{1}{\hat{s} - M_Z^2} + 2 \sum_{n=1}^{\infty} \left[\frac{Q_e Q_q}{\hat{s} - n^2 M_C^2} + \frac{g_\alpha^e g_\beta^q}{\sin^2 \theta_W \cos^2 \theta_W} \frac{1}{\hat{s} - n^2 M_C^2} \right] \right),$$

where $g_{L,R}^f = T_{3L,R}^f - Q_f \sin^2 \theta_W$, $T_{3L,R}^f$ is the third component of weak isospin for chiral fermion $f_{L,R}$, and M_C is the size of the extra dimension. The QCD coupling enters through K in Eq. 3.1, where $K = 1 + (\alpha_s(\hat{s})/2\pi)^{\frac{4}{3}} (1 + 4\pi^2/3)$. In the limit $M_C \rightarrow \infty$ the pure SM contribution can be calculated. Similar calculations are available in [9] for the modifications to the SM Drell-Yan process due to virtual graviton exchange in ADD large extra dimension models.

Evaluating Eq. 3.1, with the integration over the rapidity, y , performed numerically, allows the deviation from the Drell-Yan cross section to be determined. Fig. 3.1 shows this difference for various allowed sizes of extra dimension in both the TeV^{-1} and ADD scenarios. As can be seen, if the compactification scale of the extra dimension or the fundamental mass scale is $\mathcal{O}(10 \text{ TeV})$, the differences with respect to the SM are small.

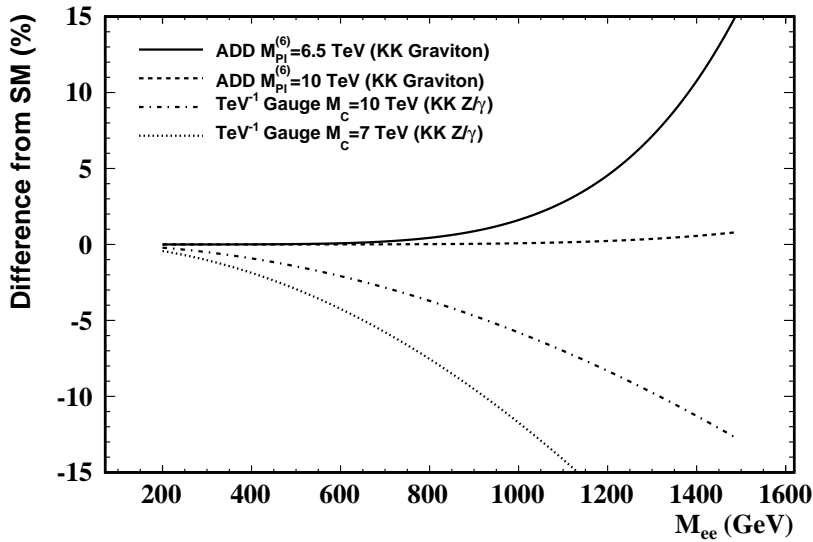


Figure 3.1: Deviation from Standard Model Drell-Yan $d\sigma/dM_{ll}$ due to large and TeV^{-1} scale extra dimensions at $\sqrt{s} = 14 \text{ TeV}$

3.1 Introduction

3.1.2 $\mathcal{O}(\alpha_{(s)})$ corrections and $\mathcal{O}(\alpha_{(s)}^2)$ significance

The leading order diagrams leading to an e^+e^- final state, assuming a photonic content within the proton, are shown in Fig. 3.2. The photon-induced subprocess (b) does not imply a correction to the γ, Z^0 resonance shown in (a), but is an irreducible background to the final state and should, therefore, be considered.




Figure 3.2: Tree level diagrams for the $q\bar{q}$ and $\gamma\gamma$ subprocesses

Next-to-leading order (NLO) QCD ($\mathcal{O}(\alpha_s)$) calculations of the Drell-Yan process have been available for some time [16], and were found to induce large corrections. Next-to-next-to-leading order ($\mathcal{O}(\alpha_s^2)$) cross sections have been calculated [17, 18], but fully differential distributions have only recently become available [19, 20, 21]. The stability of the $\mathcal{O}(\alpha_s^2)$ calculations to variations in the factorisation scale, and the small size relative to the $\mathcal{O}(\alpha_s)$ corrections naturally leads one to consider the effect of $\mathcal{O}(\alpha)$ corrections given that higher order QCD corrections are known to be small.

In the EWK sector, the dominant soft and collinear divergencies appear in diagrams where a virtual photon is exchanged between two charged external lines. Considering an s-channel process such as Fig. 3.2 (a), with $s = (p_q^\mu + p_{\bar{q}}^\mu)^2 = M_{e^+e^-}^2$ where $s \gg M_{W,Z}^2$, corrections can be expanded in terms of the small parameter $M^2/s \ll 1$. In addition, all terms of $\mathcal{O}(M^2/s)$ can be neglected, retaining only terms that are constant or divergent when the expansion parameter tends to zero. This expansion yields the double-logarithm correction [22, 23, 24, 25]

$$\mathcal{A} = \mathcal{A}_B(s) \left[1 + \sum_{n \geq 1} \left(\frac{\alpha}{4\pi} \right)^n \sum_{k=0}^{2n} C_{n,k} \ln^k \left(\frac{s}{M_W^2} \right) \right], \quad (3.2)$$

where $\mathcal{A}_B(s)$ is the tree level amplitude. Clearly this involves, at $\mathcal{O}(\alpha)$ ($n = 1$), terms of the order $\alpha \ln^2(s/M_W^2)$, which become important at large s . Similar formulae have been derived for the single logarithms [26]. Following Eqn. 3.2, approximations of the corrections due to the double (LL) and single (NLL) logarithms are given by [25, 26]

$$\delta_{DL} = \frac{\alpha}{4\pi s_W^2} C_{DL} \ln^2 \frac{s}{M_W^2}, \quad \delta_{SL} = \frac{\alpha}{4\pi s_W^2} C_{SL} \ln \frac{s}{M_W^2}, \quad (3.3)$$

where $s_W^2 = \sin^2 \theta_W(M_Z) = 0.231$, and $\alpha = 2.69 \times 10^{-3}$ (renormalised in the \overline{MS} scheme at the gauge boson mass). The approximation yields different results for the two quark sets u, c, t and d, s, b due to the differing couplings, given by the value of C_x in each




Figure 3.3: Estimated $\mathcal{O}(\alpha)$ corrections for $Q\bar{Q}, q\bar{q} \rightarrow e^+e^-$, $Q = u, c, t$ and $q = d, s, b$

case. The approximations are shown in Fig. 3.3 for the $M_{e^+e^-}$ range 200–1500 GeV. In these plots, the combination also includes a NNLL contribution constant in α from terms such as $C_{1,0} \ln^0(\dots)$ in Eqn. 3.2. At high invariant mass the combined approximation approaches -10% for $Q\bar{Q} \rightarrow e^+e^-$, which indicates that $\mathcal{O}(\alpha)$ corrections may be important relative to those at $\mathcal{O}(\alpha_s)$. Following the approximations above, the complete $\mathcal{O}(\alpha)$ corrections have been recently calculated [27, 28].

As $\mathcal{O}(\alpha)$ corrections have been shown to be important, it is required to check that $\mathcal{O}(\alpha^2)$ corrections are relatively small. In QED, the dominant logarithmic corrections of soft and collinear origin are known to exponentiate, meaning that as one-loop corrections have the form

$$\alpha K \mathcal{A}_B(s), \quad (3.4)$$

where K contains the logarithms, then the complete process amplitude, $\mathcal{A}(s)$, can be roughly approximated by

$$\mathcal{A}(s) \approx \mathcal{A}_B(s) e^{\alpha K} \quad (3.5)$$

$$\approx \mathcal{A}_B(s) \sum_{n=0}^k \frac{(\alpha K)^n}{n!} \quad (3.6)$$

$$\approx \mathcal{A}_B(s) \left(1 + \alpha K + \frac{1}{2}(\alpha K)^2 + \dots \right) \quad (3.7)$$

An active theoretical topic is to establish the exponential properties of EWK corrections, but it is by now clear that at least the leading EWK logarithms exponentiate as in QED [29]. This implies that, from Eqn. 3.7, the magnitude of $\mathcal{O}(\alpha^2)$ corrections can be estimated as

$$\mathcal{O}(\alpha^2) \approx \frac{1}{2} \mathcal{O}(\alpha)^2, \quad (3.8)$$

meaning, for instance, that in the presence of $\mathcal{O}(10\%)$ corrections at $\mathcal{O}(\alpha)$, the $\mathcal{O}(\alpha^2)$ corrections would be expected to contribute at the 0.5% level. Fig. 3.4 illustrates this estimation against the exact two-loop logarithmic calculation to N³LL. In the $Q\bar{Q}$ case the estimation is reasonable, whereas in the $q\bar{q}$ case the complete correction is negative

3.1 Introduction




Figure 3.4: Estimated $\mathcal{O}(\alpha^2)$ corrections for $Q\bar{Q}, q\bar{q} \rightarrow e^+e^-$, $Q = u, c, t$ and $q = d, s, b$

and therefore is not estimated as reliably by the exponentiating approximation. Both cases do indicate, however, that the $\mathcal{O}(\alpha^2)$ corrections are small relative to those at $\mathcal{O}(\alpha)$.

These approximations are only applicable to soft and collinear effects at high energy¹, but provide a clear motivation for the inclusion of $\mathcal{O}(\alpha)$ corrections to fully understand the Drell-Yan spectrum at high invariant mass. There is clearly less requirement to study $\mathcal{O}(\alpha_s^2)$ corrections as these have been shown to be small and stable in the QCD case, and are estimated to be small relative to $\mathcal{O}(\alpha)$ corrections in the EWK sector.

There is no available generator which incorporates both $\mathcal{O}(\alpha)$ and $\mathcal{O}(\alpha_s)$ corrections. Therefore, two generators have to be used, one for the EWK and one for the QCD sector, and the results combined. For the purpose of this study, HORACE [31, 32, 33, 34, 35, 36] was chosen as the best EWK generator, and MC@NLO [37, 38] as the best QCD generator. HERWIG [39] was used to model the parton shower following the initial hard process in each case. The various corrections and matching schemes each applies are described below.

3.1.3 HORACE

HORACE incorporates a variety of EWK and QED corrections to the Drell-Yan process, namely the complete $\mathcal{O}(\alpha)$ corrections, high order QED photon showers and, where the PDF has a photonic content, the introduction of photon-induced processes. A matching scheme is implemented to ensure the correct counting of all diagrams.

Some examples of virtual EWK corrections are shown in Fig. 3.5. In the scheme implemented by HORACE, soft and collinear divergencies are regulated with a small photon mass and loop divergencies are cancelled with renormalisation parameters in the on-shell scheme [40].

Real radiative corrections, such as those shown in Fig. 3.6 are given by all the Feyn-

¹The LL approximations have been shown to be valid for $M_{e^+e^-} > 200\text{GeV}$ [30].




Figure 3.5: $\mathcal{O}(\alpha)$ virtual corrections

man diagrams with the emission of one extra photon from the charged legs of the Born level diagrams. Note that HORACE does not include $\mathcal{O}(\alpha)$ corrections to the $\gamma\gamma \rightarrow e^+e^-$ subprocess as it was found that they are numerically negligible but require a delicate sampling of the phase-space to remove the risk of large cross-section fluctuations [41].




Figure 3.6: $\mathcal{O}(\alpha)$ real radiative corrections

In the $\mathcal{O}(\alpha)$ case, when using a PDF with a photonic content, diagrams such as those shown in Fig. 3.7 become viable. These single photon-induced processes result in an additional jet in the final state due to the outgoing quark.




Figure 3.7: $\mathcal{O}(\alpha)$ single photon-induced process corrections

HORACE also includes the effects of multiple photon emissions due to higher order QED terms, some examples of which are shown in Fig. 3.8. Multiple photon emission corrections are, as above, not applied to the $\gamma\gamma \rightarrow e^+e^-$ or $q\gamma \rightarrow e^+e^-q$ subprocesses.

3.1.4 MC@NLO

MC@NLO aims to evaluate $\mathcal{O}(\alpha_s)$ corrections and match these with the results of a leading-log (LL) parton shower model, in this case HERWIG. LL QCD parton showers readily simulate the emission of soft gluons in the initial and final state, however the addition of

3.1 Introduction




Figure 3.8: High order QED photon showers

$\mathcal{O}(\alpha_s)$ corrections introduces much harder emissions, which can drastically change the shape of transverse observable distributions.

The $\mathcal{O}(\alpha_s)$ real radiative corrections are exemplified in Fig. 3.9. The gq initial state diagram will be modelled to some degree by the parton shower at LO, which is driven by the PDF and therefore includes gluon splitting functions. However, this process is much more accurately modelled by exact $\mathcal{O}(\alpha_s)$ calculations.




Figure 3.9: $\mathcal{O}(\alpha_s)$ real radiative corrections

As MC@NLO includes complete $\mathcal{O}(\alpha_s)$ calculations, virtual corrections such as those shown in Fig. 3.10 are also included.




Figure 3.10: $\mathcal{O}(\alpha_s)$ virtual corrections

3.1.5 Generator configuration

HORACE, MC@NLO and HERWIG were configured with the following input parameters:

$$\begin{aligned}
 G_\mu &= 1.16639 \times 10^{-5} \text{ GeV}^{-2}, & \alpha &= 1/137.03599911, & \alpha_s &\equiv \alpha_s(M_Z^2) = 0.118, \\
 M_W &= 80.419 \text{ GeV}, & M_Z &= 91.188 \text{ GeV}, & \Gamma_Z &= 2.4952 \text{ GeV}, \\
 m_e &= 0.51099892 \text{ MeV}, & m_\mu &= 0.105658369 \text{ GeV}, & m_t &= 174.3 \text{ GeV}.
 \end{aligned}$$

The PDF set MRST2004QED [42] has been used throughout to describe the proton partonic content. This PDF set includes a photonic content which allows for the photon-induced processes discussed above to be included. The PDF factorisation scale has been

set equal to $\mu_F = \sqrt{(p_{t,Z})^2 + M_Z^2}$, where $p_{t,Z}$ and M_Z are the transverse momentum and invariant mass of the electron pair resulting from the Z^0 decay.

To ensure that any observed differences are due to physical reasons rather than generator misconfiguration, MCFM [43] (Monte Carlo for FeMtobarn processes) was used to validate the configurations of HORACE and MC@NLO. MCFM can compute the $pp \rightarrow Z^0 \rightarrow e^+e^-$ process at LO and NLO in QCD. It was configured with the above parameters, and run at both LO and NLO. The LO run was compared to HORACE (run at LO), and the NLO run to MC@NLO. In all cases, the differential cross sections agreed within statistical limits.

3.2 Combination results

The formula for the combination of QCD and EWK effects is given by [44, 45]

$$\left\{ \frac{d\sigma}{d\mathcal{O}} \right\}_{\text{QCD}\oplus\text{EW}} = \left\{ \frac{d\sigma}{d\mathcal{O}} \right\}_{\text{Best QCD}} + \left(\left\{ \frac{d\sigma}{d\mathcal{O}} \right\}_{\text{Best EWK}} - \left\{ \frac{d\sigma}{d\mathcal{O}} \right\}_{\text{LO}} \right)_{\text{HERWIG PS}} \quad (3.9)$$

where the differential cross-section with respect to any observable, \mathcal{O} , is given by three terms:

- The result of the best available QCD calculation (MC@NLO).
- The effects due to NLO EWK corrections and higher-order QED effects of multiple photon radiation (HORACE)
- The leading order distribution, subtracted to avoid double counting as this is included in both the QCD and EWK generators.

The EWK corrections are convoluted with a QCD parton shower, using HERWIG, and therefore include, in the collinear approximation, the bulk of the $\mathcal{O}(\alpha\alpha_s)$ corrections.

The results are obtained for an e^+e^- final state with the following cuts applied to select the events:

$$p_t^{e^\pm} > 25 \text{ GeV}, \quad |\eta^{e^\pm}| < 2.5, \quad M_{e^+e^-} > 200 \text{ GeV}. \quad (3.10)$$

The percentage corrections have been defined as

$$\delta = 100 \cdot \frac{(\sigma_{\text{NLO}} - \sigma_{\text{Born+PS}})}{\sigma_{\text{Born+PS}}} \quad (3.11)$$

As photons radiated from a final-state electron with a small opening angle can not be distinguished independently due to calorimeter granularity, a photon recombination scheme was included. For every electron and photon in the final state, the four-momentum vectors of the electron and photon are combined into an effective electron four-momentum

3.2 Combination results

vector if $\Delta R(e, \gamma) < 0.1$. Photons are excluded from this procedure if $|\eta_\gamma| > 2.5$. The event selection cuts (Eqn. 3.10) are applied after the recombination procedure.

3.2.1 Corrections at $\sqrt{s} = 14$ TeV

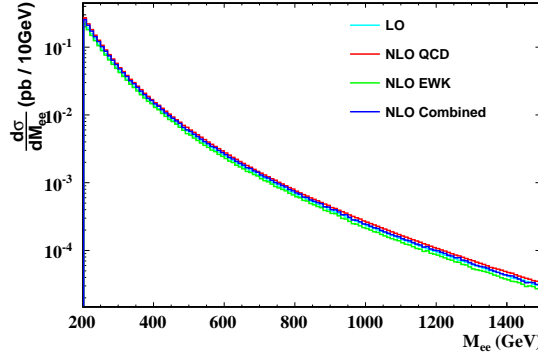


Figure 3.11: M_{ee} differential cross sections at LO and NLO

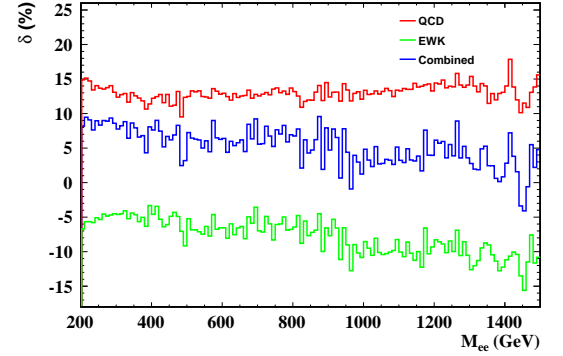


Figure 3.12: Corrections to M_{ee} due to QCD and EWK NLO corrections




Figure 3.13: Electron p_t differential cross sections at LO and NLO

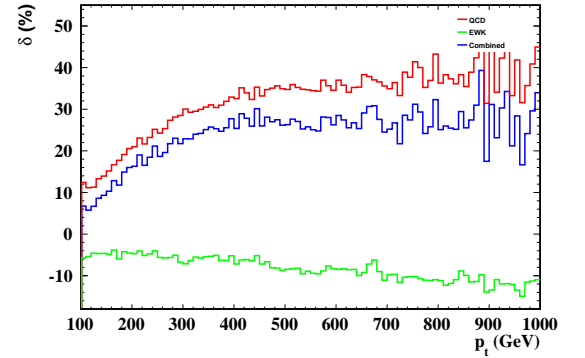


Figure 3.14: Corrections to the electron p_t due to QCD and EWK NLO corrections

Fig. 3.11 shows the interplay between the QCD and EW corrections for the di-lepton invariant mass differential cross section. The QCD corrections are quite flat and positive with a value of about 15% over the mass range 200–1500 GeV. The EW corrections are negative and vary from about -5% to -10% and thus partially cancel the NLO-QCD effect.

In Fig. 3.13 the lepton transverse-momentum distribution is shown. The NLO QCD corrections increase the lepton transverse momentum differential cross section from 10% to 35% in the interval considered (100–1000 GeV). The NLO EWK corrections are negative and fall from -5% to -10% over the same range. The corrections to the di-electron transverse momentum differential cross section, shown in Fig. 3.15, clearly show that the QCD corrections are responsible for large corrections to the $Z p_t$, but the EWK corrections do not contribute.

The neutral-current Drell-Yan process in the high invariant mass region




Figure 3.15: De-electron p_t differential cross sections at LO and NLO

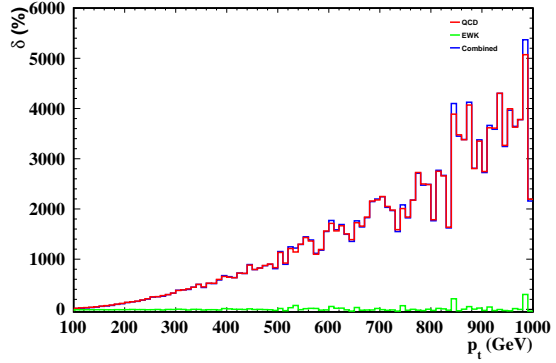


Figure 3.16: Corrections to the de-electron p_t due to QCD and EWK NLO corrections




Figure 3.17: Electron pseudorapidity differential cross sections at LO and NLO

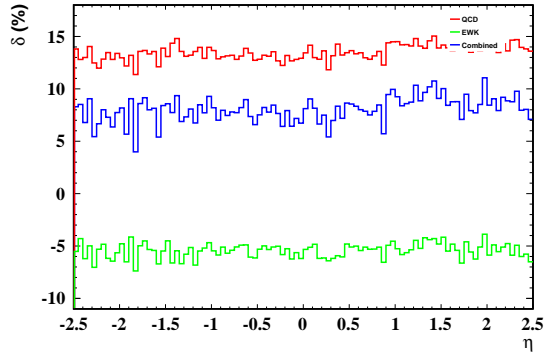


Figure 3.18: Corrections to the electron pseudorapidity due to QCD and EWK NLO corrections




Figure 3.19: Di-electron pseudorapidity differential cross sections at LO and NLO




Figure 3.20: Corrections to the di-electron pseudorapidity due to QCD and EWK NLO corrections

3.2 Combination results

3.2.2 Comparison with PDF effects

The Parton Density Functions (PDFs) describing the effective parton content per proton are parameterised by the parton flavour, f_i , the energy scale of an interaction, $Q = \sqrt{s}$, and the momentum fraction of each parton, x_i . As the PDFs are derived from fits to data, there is a degree of uncertainty in their fitted values.

The Hessian method [46] provides a way to both fit the available data and explore the resulting uncertainties. The result of the method is a set of parameterised values, a_i , with $i = 1..N$, which represent the position of the global minimum χ^2 of the fit parameters, and the Hessian matrix, H_{ij} , which contains all of the second partial derivatives of a_i in the region of the local minima. By varying the global χ^2 according to

$$\Delta\chi^2 = \chi^2 - \chi_0^2 = \sum_{i=1}^N \sum_{j=1}^N H_{ij}(a_i - a_i^0)(a_j - a_j^0),$$

such that $\Delta\chi^2$ is equal to some pre-defined tolerance, T^2 , N PDF uncertainty eigensets representing the observed tolerances in the Hessian eigenvector basis are produced. T is typically of the order 10 to 15, and is chosen such that all datasets used in the fit can be self-consistently described by the central PDF fit and each uncertainty eigenset.

By construction, the error sets are symmetric about the central fit, however the effect of the errors propagated through to an experimental observable may not be. The modified tolerance method takes into account maximal asymmetric variations on a physical observable due to the PDF fit uncertainties [47]. The variation of a given observable X is defined as

$$\begin{aligned} \Delta X^+ &= \sqrt{\sum_{i=1}^N [\max(X_i^+ - X_0, X_i^- - X_0, 0)]^2} \\ \Delta X^- &= \sqrt{\sum_{i=1}^N [\max(X_0 - X_i^+, X_0 - X_i^-, 0)]^2}, \end{aligned}$$

where X_0 is the observable with respect to the central PDF set and X_i^\pm the observable with respect to the i th positive or negative uncertainty eigenset. This implies that the data will have to be simulated N times, in order to measure X for each PDF uncertainty eigenset.

For an event generated by a Monte Carlo simulation, the differential cross-section can be expressed as

$$\frac{d\sigma}{dY}[\text{pp} \rightarrow X] \approx \sum_i \sum_j (f_i^n(x_1, Q) f_j^n(x_2, Q)) \hat{\sigma}, \quad (3.12)$$

where i, j are incoming parton types, $\hat{\sigma}$ is the cross section for the process $ij \rightarrow X$, $x_{1,2}$ the momentum fraction carried by each parton, Q the event scale and $f_i(x_{1,2})$ the probability

to find the parton i with momentum fraction $x_{1,2}$ in the proton according to PDF set n . This expression allows the PDF to be factored out of a given event, and a differential cross-section to be re-weighted on an event by event basis, with no need for re-simulation. The event weight for a new PDF set S is given by

$$w_E = \frac{f_i^S(x_1, Q) f_j^S(x_2, Q)}{f_i^0(x_1, Q) f_j^0(x_2, Q)},$$

the result of which is clear when combined with Eq. 3.12; the old PDF contribution is removed and a new contribution applied. This removes the need for re-simulation of data if the parton information is available.

Applying this method to the tree level Drell-Yan process yields the uncertainty shown in Fig. 3.21. The uncertainty is of a similar order of magnitude to that of the higher order corrections and the example motivating physics case discussed in §3.1.1. Therefore, accurate determination of the PDFs from LHC data will be vital to drive down this uncertainty.

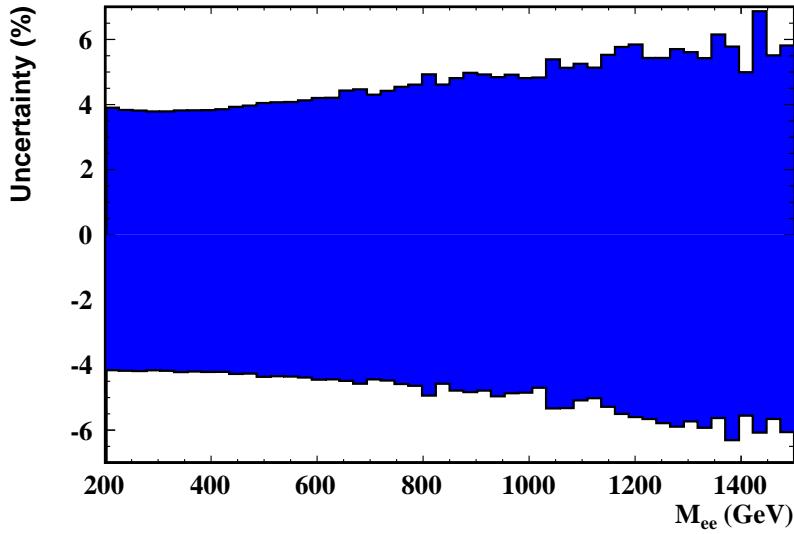


Figure 3.21: Uncertainty on M_{ee} due to PDF fit

4. The LHC and the CMS Detector

4.1 The Large Hadron Collider

The LHC [48] is a two-ring, 27 km circumference, superconducting proton accelerator and collider, situated in Geneva, Switzerland. It occupies the tunnel formerly housing the LEP machine. The tunnel lies between 45 m and 170 m below the surface on a plane inclined at 1.4%, sloping towards Lake Geneva.

The LHC is designed to provide proton-proton collisions at a centre of mass energy of $\sqrt{s} = 14$ TeV, at high luminosity ($L = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$). $p\bar{p}$ collisions are ruled out by the high beam intensity required for the specified luminosity. As such, the LHC requires a two-ring design to allow for the two opposite magnetic dipole bending fields. The nominal field strength is 8.33 T, and the total stored energy (magnets and beams) is approximately 1.2 GJ, placing stringent demands on beam containment and safety to stop beam-induced damage to the accelerator and detectors.

There are four interaction regions, two high-luminosity points (for CMS and ATLAS), and two lower-luminosity (for LHCb and ALICE). The overall machine layout is shown in Fig. 4.1.

4.1.1 Machine commissioning

The LHC first successfully circulated proton beams on the 10th September 2008. However, a fault soon after put the machine out of operation for a year. On the 20th November 2009, after the machine was repaired and a new magnet protection system installed, beams were circulated once again. Problems found during the shutdown resulted in the decision to limit the machine energy to 3.5 TeV per beam during the first run.

After shutdown in 2011, the machine will be out of commission for a further year while modifications are performed. These modifications to the magnet bus bar connections will allow the machine to reach the design energy of 7 TeV per beam.

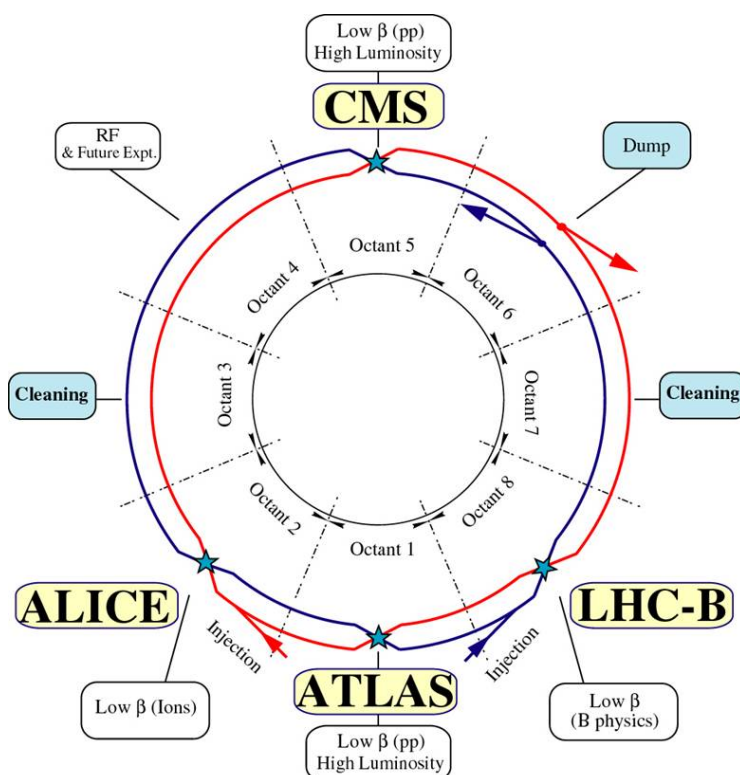


Figure 4.1: Schematic layout of the LHC [48]

4.1.2 Magnet technology

The LHC relies on magnets based on NbTi superconductors cooled to 1.9 K using liquid helium. The ring is constructed from 1232 bending dipoles, all with the same design, illustrated in Fig. 4.2. The dipoles are 16.5 m long and have a mass of 27.5 t. At injection energy, the magnets have a current of 763 A (corresponding to a field of 0.54 T), which ramps up to 11850 A at nominal 7 TeV energy.

Various other magnets are used in the LHC, namely about 4800 corrector magnets, focusing quadrupoles and kickers for beam injection and extraction.

4.1.3 Acceleration systems

The capture, acceleration and maintenance of the beams is handled by a 400.8 MHz superconducting RF system. The system provides 16 MV per beam during cruising at high energy, and provides an energy gain of 485 keV per turn during acceleration. Due to the high proton mass, the synchrotron radiation loss per turn is small with respect to the beam energy (7 keV at 7 TeV).

The beams in the machine are separated by 194 mm, which does not allow enough space for individual RF systems per-beam. Therefore, the beams are separated to 420 mm by special superconducting dipoles either side of the RF systems. Each RF system has

4.1 The Large Hadron Collider




Figure 4.2: Cross section of an LHC dipole [48]

eight cavities, each providing 2 MV of accelerating voltage. The required 4800 kW of RF power is generated by sixteen 300 kW klystrons.

4.1.4 Bunch structure and luminosity estimates

The LHC bunch structure is prepared by the Proton Synchrotron (PS) and Super Proton Synchrotron (SPS). At design luminosity, each beam will consist of bunches spaced by 25 ns. However, not all bunches are filled; there are gaps due to the pre-accelerator to allow injection and extraction kicker magnets to ramp up and down. The nominal bunch pattern is shown in Fig. 4.3. There are three main bunch-spacing patterns; 12 missing bunches between each 72-bunch batch (to allow for the SPS injection kickers), 38 between each batch of 3 or 4 set of bunch batches (to allow for the LHC injection kickers), and a large 119 bunch gap at the end of the structure to allow for the LHC beam-dump kicker ramp-up time. The experiments may use these un-filled bunches for calibration, pedestal and alignment measurements.

Throughout a physics run, the instantaneous luminosity will decrease with time. This is due to a number of effects but is dominated by beam loss due to collisions. The effective beam lifetime for such loss is given by

$$\tau_L = \frac{N_{\text{tot},0}}{L_0 \sigma_{\text{tot}} k},$$

where $N_{\text{tot},0}$ is the initial beam intensity, L_0 the initial luminosity, σ_{tot} the total proton-proton cross section and k the number of interaction points. Assuming two experiments (LHCb and ALICE require lower luminosity collisions and therefore do not dominate the above calculation), a decay time of $\tau_{\text{nuclear}} = 44.85$ h can be determined. The decay in

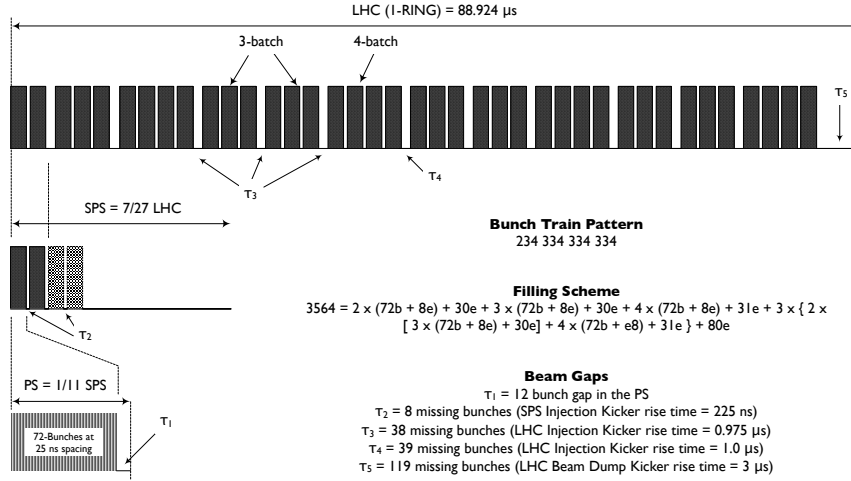


Figure 4.3: Nominal bunch structure for 25 ns LHC operation [49]

collision luminosity as a function of time is given by

$$L(t) = \frac{L_0}{(1 + t/\tau_{\text{nuclear}})^2},$$

with the time to reach $1/e$ of the initial luminosity given by

$$t_{1/e} = (\sqrt{e} - 1) \tau,$$

which yields a luminosity decay time of $\tau_{\text{nuclear},1/e} = 29$ h. Factoring in other losses due to intra-beam scattering and collisions with residual particles in the vacuum leads to an estimation of the total luminosity decay time to reach L_0/e of $\tau_{L,1/e} = 14.9$ h.

The integrated luminosity over one run is given by

$$L_{\text{int}} = L_0 \tau_L (1 - e^{-T_{\text{run}}/\tau_L}).$$

With an expected machine turnaround time between fills (T_{ta}) expected to take a minimum of 1.15 h (scaled to 7 h given past collider experience), the optimum run time per fill, assuming 200 days operation per year, can be determined as 5.5 h (12 h), leading to a total annual integrated luminosity of 80 fb^{-1} to 120 fb^{-1} . In the first year of running, it is expected that $\mathcal{O}(100 \text{ pb}^{-1})$ of integrated luminosity will be delivered, while the machine and experiments are commissioned and understood.

4.2 The Compact Muon Solenoid Detector

The Compact Muon Solenoid (CMS) detector is designed as a general-purpose detector. At design energy and luminosity, the LHC will deliver approximately 10^9 inelastic events

4.2 The Compact Muon Solenoid Detector




Figure 4.4: Schematic layout of the CMS detector [50]

per second, with around 20 such collisions superimposed on each event of interest. This places huge demands on the triggering, detection and readout systems to cope with the high interaction rate (event selection, storage and analysis), high occupancy and associated synchronisation.

While CMS is a general purpose detector, its primary aim is to determine the mechanism behind electroweak symmetry breaking, and explore physics at the TeV scale. The requirements for such an aim can be expressed as [50]:

- Good muon identification and momentum resolution over a wide range of momenta and angles, good di-muon mass resolution ($\approx 1\%$ at 100 GeV), and the ability to determine unambiguously the charge of muons with $p < 1$ TeV
- Good charged-particle momentum resolution and reconstruction efficiency in the inner tracker. Efficient triggering and offline tagging of τ s and b -jets, requiring pixel detectors close to the interaction region
- Good electromagnetic energy resolution, good di-photon and di-electron mass resolution ($\approx 1\%$ at 100 GeV invariant mass), wide geometric coverage, π^0 rejection, and efficient photon and lepton isolation at high luminosities
- Good missing transverse energy and di-jet mass resolution, requiring hadronic calorimeters with a large hermetic geometric coverage with fine lateral segmentation.

CMS consists of a silicon tracking system (pixel and strips), an electromagnetic calorimeter (ECAL) and a hadronic calorimeter (HCAL), all contained within a 4 T superconducting solenoid. Surrounding the solenoid is the iron magnetic return yoke, instrumented with the muon systems (drift tubes (DT), resistive plate chambers (RPC) and cathode strip chambers (CSC)). Additionally there is very forward hadronic calorimetry (HF). The overall cut-away design is shown in Fig. 4.4. A two-stage triggering strategy is used, with the Level-1 (L1) trigger performing the first event selection step, and the High Level Trigger (HLT) subsequently reducing the event rate further.

4.2.1 Superconducting Solenoid

The superconducting magnet [51] for CMS has a length of 12.5 m, diameter of 6.3 m, mass of 220 t, and is designed to sustain a uniform 4 T field within the 6 m diameter free bore (the radial thickness of the cold mass is 312 mm, corresponding to a radiation length of $3.8\chi_0$). The stored energy at full field is 2.6 GJ.

The magnet is constructed from four layers of an NbTi Rutherford-type cable, co-extruded with aluminium. The conductor itself therefore becomes part of the mechanical support structure, and takes 70% of the magnetic hoop stress (130 MPa at full field). The remaining 30% is taken by the support mandrel.

The cold mass is cooled to 4.5 K with liquid He. When cooled, the cold mass shrinks by 26 mm longitudinally, and 14 mm radially. Following a fast discharge (i.e. triggered by a quench), it can take up to 3 days to re-cool the coil.

4.2.2 Inner Tracking

The inner tracking system is needed to measure the momentum and charge of charged particles produced in LHC collisions. The CMS tracking system [52] consists of an inner pixel detector and an enclosing silicon strip tracker. In total, it has a length of 5.8 m and diameter of 2.5 m.

The overall layout in cross-section is shown in Fig. 4.5, detailing the location of the inner pixel detector, inner barrel and disks (TIB and TID), outer barrel (TOB) and endcaps (TEC). The geometrical acceptance of the tracking system extends up to $|\eta| < 2.5$.

The pixel detector has three layers in the barrel, at radii of 4.4, 7.3 and 10.2 cm. Two disks at ± 34.5 and ± 46.5 cm complete forward coverage, detailed in Fig. 4.6. The pixel detector has 66 million pixels, each $100 \times 150 \mu\text{m}^2$, and occupies an active area of 1 m^2 . It is expected to have a mean occupancy of 10^{-4} / pixel in each bunch crossing [53].

The silicon strip tracker covers the radial region from 20 to 116 cm. The TIB / TID cover the region up to 55 cm in radius, and are composed of 4 barrel layers and 3 endcap

4.2 The Compact Muon Solenoid Detector




Figure 4.5: Cross-section of the silicon tracking systems [50]




Figure 4.6: Detailed layout of pixel detector [50, 52]

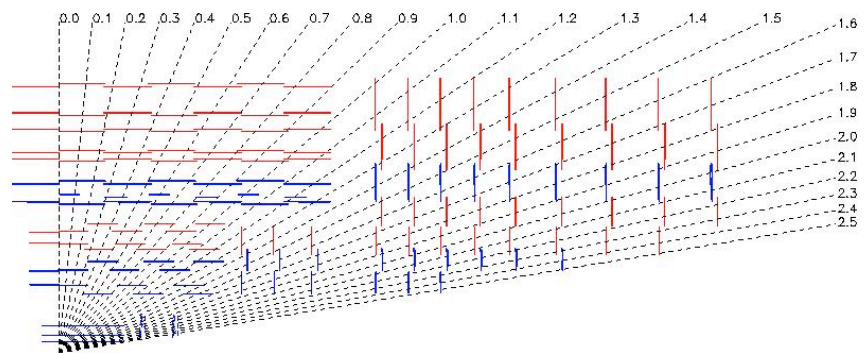


Figure 4.7: Silicon strip tracker geometry [52]

disks on either side. It is expected to have an occupancy of $\approx 2 - 3\%$ per bunch crossing. The TOB consists of 6 barrel layers, extending up to ± 118 cm in z . The TECs, each consisting of 9 disks, fully complete the tracking coverage. In this outermost region, the occupancy is expected to be $\approx 1\%$. The geometry is shown in detail in Fig. 4.7.

The silicon microstrip sensors which make up the TIB have a minimum size of $10 \text{ cm} \times 80 \mu\text{m}$. The TOB, being further away from the interaction point, is constructed with larger sensors, up to $25 \text{ cm} \times 180 \mu\text{m}$.

This complete coverage, and number of active sensors, allows very good momentum resolution, $\mathcal{O}(2\%)$ for 100 GeV tracks, as shown in Fig. 4.8. However, the large number of active elements, and associated powering, cooling and readout electronics pushes the material budget to $1.8\chi_0$ over a restricted area, as shown in Fig 4.9. The result of this is particles showering in the tracker, which must be taken into account in the reconstruction algorithms.

To control the damage due to incident radiation, the tracker must be kept below -10°C . The total active power to dissipate is around 60 kW. The cooling system uses C_6F_{14} due to its thermal and mechanical properties. The cooling system provides a total cooling capacity of 128 kW, delivering up to $77 \text{ m}^3\text{h}^{-1}$ of cooling fluid down to -35°C .




Figure 4.8: Tracker momentum resolution [52]




Figure 4.9: Tracker material budget [50]

4.2.3 Electromagnetic Calorimetry

The CMS ECAL [54] design was driven by the desire to be able to probe the low-mass Higgs sector. As such, it must be sensitive to the $\gamma\gamma$ decay channel of the postulated Higgs boson, and be able to detect the e^+e^- decay modes of intermediate Z^0 and W^\pm bosons resulting from Higgs decays. This requires a calorimeter with good energy resolution. Such requirements, along with the LHC operating conditions, naturally lead themselves

4.2 The Compact Muon Solenoid Detector

to the choice of a homogenous crystal calorimeter.

PbWO₄ crystals were chosen as they have a short radiation length (0.89 cm) and small Molière radius (2.2 cm), which results in a calorimeter with fine granularity. The crystals emit 80% of their scintillation light in 25 ns, however the light yield is low. At 18 °C, about 4.5 photoelectrons per MeV are collected, meaning that electronics noise must be kept to an absolute minimum.

The ECAL is made of 61200 such crystals in the barrel section, with 7423 crystals in each of the endcaps. Photodetection is performed by avalanche photodiodes (APDs) in the barrel, and vacuum phototriodes (VPTs) in the endcaps. The overall layout is shown in Fig. 4.10.




Figure 4.10: Overview ECAL layout [54]

The barrel has geometric coverage up to $|\eta| < 1.479$, and consists of 360 crystals in ϕ , and two modules of 85 crystals in η . The crystals have a tapered shape, with dimensions of $22 \times 22 \text{ mm}^2$ at the front face and $26 \times 26 \text{ mm}^2$ at the rear. The length of the crystals (230 mm) corresponds to an interaction length of $25.8\chi_0$. To reduce energy loss in cracks, the crystals are arranged to off-point the interaction region by about 3° in both η and ϕ .

The endcap covers the geometric range $1.479 < |\eta| < 3.0$, and each is 315.4 cm from the interaction point. In the endcap, crystals are grouped into 5×5 arrays (supercrystals). At the inner and outer edges of the calorimeter, partial supercrystals are used to keep constant radial coverage. Each crystal is identical, with front face dimensions $28.62 \times 28.62 \text{ mm}^2$, rear face $30 \times 30 \text{ mm}^2$, and a length of $24.7\chi_0$ (220 mm).

The energy-dependent resolution of the calorimeter can be parameterised as

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2,$$

where S is the stochastic term, N the noise term, and C the constant term. Results from a test-beam run in 2004 show typical figures of $S = 2.8\%$, $N = 120$ MeV and $C = 0.3\%$. Later tests with the final front end electronics reduced the noise term by 10%. This corresponds to a measured resolution of 0.5% for 120 GeV electrons. The evolution of resolution with energy is shown in Fig. 4.11




Figure 4.11: Evolution of ECAL energy resolution [54]

The calorimeter must be kept cool, to remove heat due to the readout and triggering electronics, but also for the stability of light yield. The variation in light yield, due to crystal and APD / VPT effects has been measured at 18°C to be $-3.8 \pm 0.4\% \text{ } ^\circ\text{C}^{-1}$ [55]. The required temperature and allowed variation is $18 \pm 0.05^\circ\text{C}$, to keep the energy resolution stable. To achieve this, the detector is equipped with a circulating water cooling system.

4.2.4 Hadronic Calorimetry

The CMS HCAL [56] consists of barrel (HB), outer (HO), endcap (HE) and forward (HF) subdetectors. Their locations are shown in Fig. 4.12. The HB covers the geometric range $|\eta| < 1.3$, The HE extends this to $1.3 < |\eta| < 3$, and the HF further still to the region $3 < |\eta| < 5.2$.

The HB is a sampling calorimeter constructed of 36 wedges, each constructed from 14 layers of 50.5 mm and 56.5 mm brass absorber plates, interspaced with 3.7 mm thick plastic scintillator tiles. The effective absorber thickness at $\eta = 0$ is $5.82 \lambda_I$, increasing as $1/\sin\theta$ to $10.6 \lambda_I$ at $|\eta| = 1.3$. In terms of hadronic shower development, the ECAL adds approximately $1.1 \lambda_I$ of material in front of the HB. The scintillator is segmented by 0.087 in both η and ϕ . The scintillation light is guided by wavelength shifting fibers, and measured with hybrid photodiodes (HPDs).

4.2 The Compact Muon Solenoid Detector

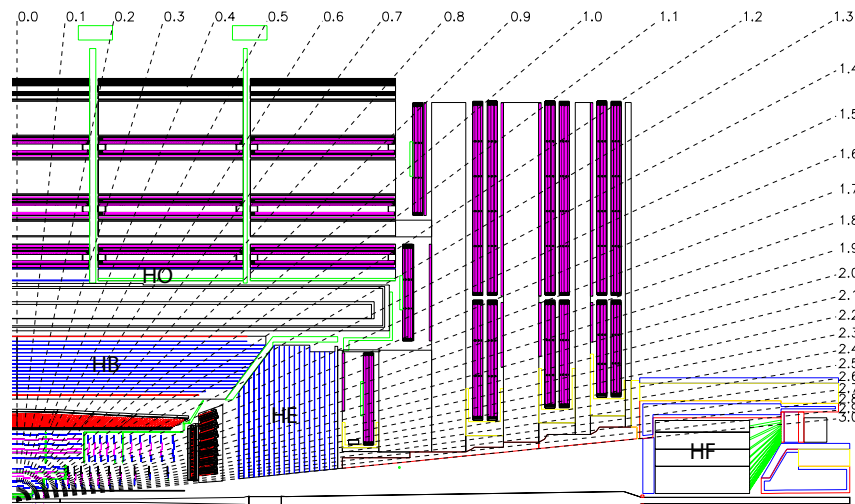


Figure 4.12: Overview HCAL layout [50]

In the central regions, there will be insufficient containment of hadronic showers in the EB and HB. For this reason, the hadronic calorimetry is extended outside of the solenoid with the HO. Using the solenoid coil as an additional absorber adds $1.4/\sin\theta_{\lambda_I}$ to the effective hadronic stopping power. The HO is divided into 5 rings in η , to fit with the muon system. The central ring has two layers of scintillating tiles, whereas the remaining layers have only one. The segmentation is designed to mirror that of the HB such that consistent hadronic towers can be constructed.

The HE is constructed from the same materials as the HB. It longitudinally consists of 19 absorber tiles, each 79 mm thick with a 9 mm gap for the scintillation tiles. The detector granularity is $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$ for $|\eta| < 1.6$ and $\Delta\eta \times \Delta\phi \approx 0.17 \times 0.17$ for $|\eta| \geq 1.6$. The absorber plates are staggered to ensure all incident particles will shower.

The HF will reside in an extreme radiation environment, with a dose of ≈ 10 MGy at $|\eta| = 5$ after 10 years of LHC operation. To withstand this environment for a decade, quartz fibres were chosen as the scintillation medium. They are expected to have a loss in optical transmission of a half after such a radiation dose. The calorimeter is constructed from a steel absorber, composed of 5 mm plates, with grooves for the fibres. The absorber is 165 cm ($10\lambda_I$) long, and is instrumented with two sets of fibres. One set runs the full length, and the second set starts at a depth of 22 cm, allowing electromagnetic and hadronic showers to be separated.

The transverse energy resolution of each subdetector is shown in Fig. 4.13, and is typically of order 10% for high E_t jets.

4.2.5 Muon System

The muon systems of CMS are designed to be sensitive to the whole kinematic region accessible in LHC collisions. Of particular interest is the so-called 'golden mode' decay




Figure 4.13: Evolution of HCAL transverse energy resolution for each subsystem [56]

$H \rightarrow Z^0 Z^{0(*)} \rightarrow 4\mu$. As such full geometric coverage for the muon system is essential. To cope with the varying magnetic and radiation environments present, the muon system is constructed from three different technologies [57]. The region $|\eta| < 1.2$ is covered by drift tubes (DT), cathode strip chambers (CSC) cover the region $0.9 < |\eta| < 2.4$. Resistive plate chambers (RPC) complement the muon system to provide a fast trigger capability in the region $|\eta| < 1.6$. The overall layout is shown in Fig. 4.14.

The drift tubes, which cover the central region, consist of four stations surrounding the beam, arranged in concentric cylinders interspersed with the iron return yoke for the superconducting solenoid. The inner three cylinders have a total of 60 drift chambers, and the outer cylinder 70, with a total of about 172000 sensitive wires. Drift chambers can be used in the barrel due to the low rate, backgrounds, and relatively low magnetic field strength.

In the endcap region, with a high magnetic field and large particle fluence, CSCs are used. These have a fast response time, ideal for use in triggering. The CSCs are multiwire proportional chambers, each with 6 anode wire planes between 7 cathode panels. The CSCs are arranged into four disks in each endcap.

To enable the muon trigger system to unambiguously assign a bunch crossing to each muon track, an independent fast system is required. RPCs are gaseous parallel-plate detectors that have a fast response time (measured at 1.26 ns for the CMS design [58]) suitable for such a task. The barrel and endcap region are instrumented with RPCs. At start up, the coverage will extend to $|\eta| < 1.6$, and will be extended to $|\eta| < 2.1$ for high-luminosity running.

The combined transverse momentum resolution of the stand-alone muon systems,

4.2 The Compact Muon Solenoid Detector

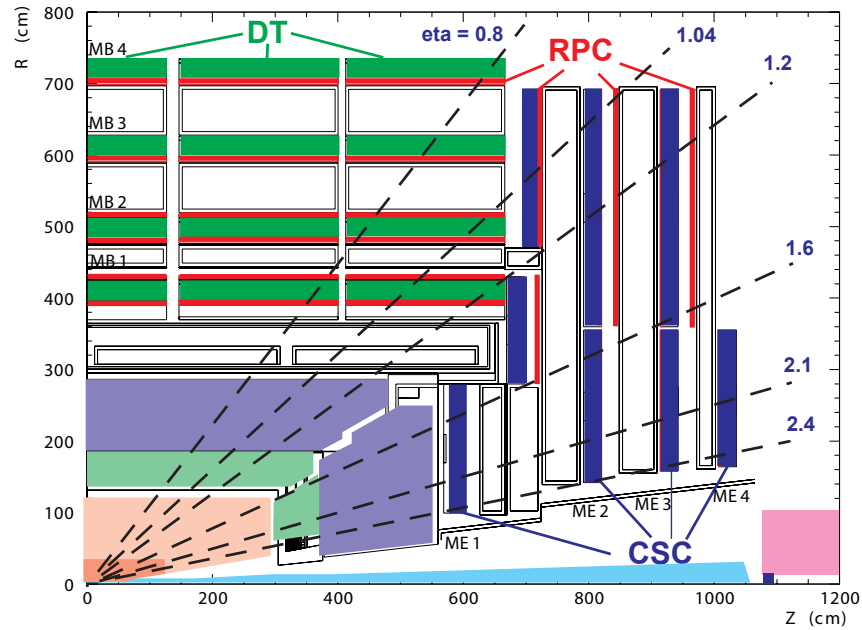


Figure 4.14: Muon system layout [57]

and when combined with inner tracking information is shown in Fig. 4.15. The region $0.8 < |\eta| < 1.2$ is excluded as the barrel and endcap systems overlap.

4.2.6 Level-1 Triggering

To read out and store the result of every bunch crossing would require a bandwidth and storage capacity of approximately 40 TBs^{-1} , clearly far outside of the available technology. Also, due to the rarity of the processes of interest at the LHC, most bunch crossings do not result in an interesting event. The L1 trigger [59] performs the first stage of selecting, in real time, those interesting events to store.

The L1 trigger is implemented in custom electronics, residing on-detector and in the counting room, situated approximately 90 m from the experiment. It is designed to reduce the rate to 100 kHz, using information from the calorimeter and muon systems, with a maximum latency between a collision and the L1 accept signal being received by the front-end electronics of $3.1 \mu\text{s}$. The overall pipelined layout is shown in Fig. 4.16.

All triggering decisions start with the calculation of trigger primitives. These contain coarse-grained information on potential trigger object kinematics (e.g. E_t , η , ϕ) and quality (e.g. number of good muon hits, or E_t ranking) for each subsystem. For triggering purposes, the calorimeters are segmented into trigger towers. In the central region ($|\eta| < 1.74$), the towers have granularity of 0.087 in both η and ϕ . In the forward region the towers are larger. The trigger primitives for the ECAL are generated on-detector in the case of the barrel, and with a combination of on-detector and near-detector processing for the endcaps. HCAL trigger primitives are generated on the HCAL Trigger and




Figure 4.15: Evolution of muon transverse momentum resolution [57]

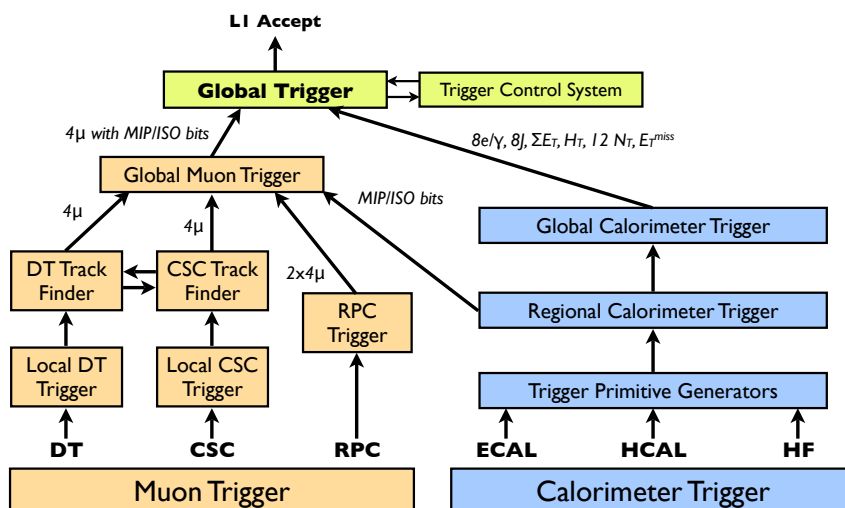


Figure 4.16: Level-1 Trigger layout [59], detailing the flow of trigger primitives and derived quantities through the various trigger subsystems

4.2 The Compact Muon Solenoid Detector

Readout boards, situated off-detector. These trigger primitives are sent to the Regional Calorimeter Trigger (RCT).

The RCT is responsible for identifying electron / photon candidates and transverse energy sums for each calorimeter region. A region is defined as 4×4 trigger towers. The RCT classifies the candidates in each region into isolated and non-isolated collections, based on the energy depositions surrounding a maximum tower, and the Hadronic / EM fraction of the candidate. The regional energy sums, for jet energy measurement, are also assigned a τ -veto bit, to distinguish narrow τ -jets from quark / gluon jets. There are 18 RCT crates, each covering a region of $\Delta\eta \times \Delta\phi = 5.0 \times 0.7$.

The Global Calorimeter Trigger (GCT) collates the information from the RCT crates, and determines the highest-ranked jets, electron / photon candidates and also provides total E_t , E_t^{miss} , H_t (the scalar transverse energy sum of all jets over threshold), and jet counts. These candidates are then sent to the Global Trigger (GT).

The muon trigger combines all three muon systems to achieve a momentum resolution and efficiency much improved from using just the stand-alone systems. The barrel DTs provide track segments in the ϕ projection, and hit patterns in the η projection. The endcap CSCs deliver 3-dimensional track segments. The Regional Muon Trigger combines track segments from these two subsystems to produce complete tracks. The RPC chambers produce completely stand-alone track candidates.

For every bunch crossing, the Global Muon Trigger (GMT) receives four muon candidates from the DTs, barrel RPCs, CSCs and endcap RPCs. Each candidate is assigned a p_t , charge, η , ϕ and a quality code. The RCT provides calorimeter isolation information in regions of 0.35 in η and ϕ . The GMT combines the candidates, taking into account the isolation criteria, and provides 4 muon candidates to the GT.

The GT, on receipt of the final object candidates from the GCT and GMT, makes a decision whether to accept or reject the event. Each object is characterised by its p_t or E_t , (η , ϕ) coordinates, charge (for muons), and quality. The GT contains completely programmable logic allowing trigger algorithms to be designed based on topological event conditions. Up to 128 trigger algorithms can be executed in parallel.

4.2.7 Data Acquisition and High Level Triggering

Following a L1 accept, the entire detector must be read out such that the next trigger step can occur. Due to the varying distance and latencies in data read out from different subdetectors, the DAQ [60] must cope with asynchronous event fragments arriving in the counting room and assemble them into single events on dedicated Filter Units (FU). The DAQ design to perform this is shown in Fig. 4.17.

The detector front ends, which read out a total of ≈ 55 M channels, are connected to

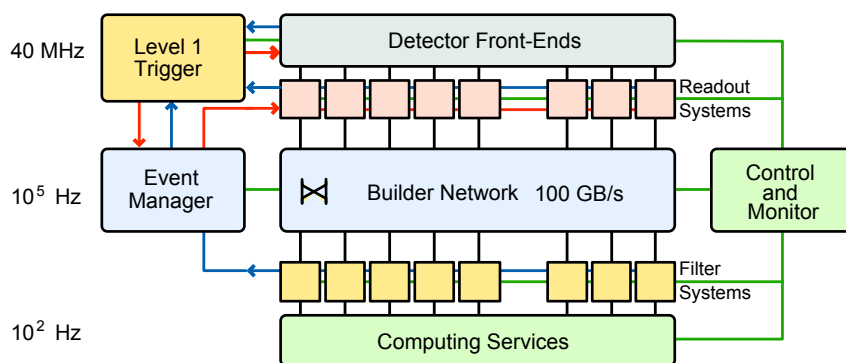


Figure 4.17: Overview of the DAQ system [60]. Upon a L1 accept, all Detector Front-Ends are read out. The Builder Network assembles each whole event onto a Filter Unit. Data to be kept is written to central Computing Services.

the event builder. The event builder is based on a Myrinet network, which allows an event fragment from any front end driver (FED) to be routed to any FU. The sustainable aggregate throughput of the network is $\approx 1.4 \text{ Tbs}^{-1}$.

With an event assembled on a FU, the High Level Trigger (HLT) algorithms are run. This software-based triggering step performs high-level reconstruction algorithms, using all detector data. At design luminosity, there will be ≈ 2500 CPU cores to keep up with the 100 kHz L1 rate. The time budget per event is 50 ms.

The HLT algorithms are designed to reduce the event rate down to 150–300 Hz, resulting in $\mathcal{O}(250 \text{ MBs}^{-1})$ final data rate. This is stored on a large disk cache at the experimental site (the Storage Manager) before being transferred to CERN for further processing.

5. Level-1 triggering for high E_t electromagnetic objects

Measurements of the ECAL Very Front End (VFE) electronics are presented, which are used as input to a simulation of the L1 Trigger Primitive Generator (TPG) on-detector electronics. A run-time configuration for the digital processing steps within the TPG is described which allows the correct bunch-crossing identification and triggering for saturating signals in the ECAL.

The response of the complete L1 trigger, as described by the bit-level emulation available in `CMSSW_1_3_1_HLT6` and `CMSSW_2_2_9`, to high p_t electrons is investigated. It is shown that a non-isolated trigger is required to maintain efficiency at high p_t , and that jet background rates are under control for an instantaneous luminosity of $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$, making the modified trigger configuration suitable for use during the first lower-luminosity phase of LHC running.

5.1 Introduction

5.1.1 Motivation

Certain extensions to the Standard Model allow for the production of high transverse momentum electrons, either directly through the decay of new heavy particles, or through the production of W^\pm and Z^0 bosons with high transverse momentum. As a general purpose experiment, CMS is searching for any and all accessible signs of new physics, and care should be taken to ensure that all reasonably allowable decay modes of interest can be investigated.

New physics signatures can only be analysed if events of interest can be triggered, as the data for an untriggered event can not be recovered. Due to the behaviour of the ECAL trigger electronics, it was feared that very high energy (TeV scale) electrons would cause the trigger to fire on an incorrect bunch crossing [53]. Due to the L1 trigger rules, designed to protect the tracker readout, this would result in the tracker data being lost for the event. It was therefore crucial to understand the behaviour of the ECAL trigger

electronics in the high energy regime, and determine how high transverse momentum electrons could be triggered on with high efficiency.

5.1.2 Electronics overview

Coarse-grained information from the calorimeter is used in the Level-1 triggering step. Reading out the entire calorimeter would place stringent demands on available bandwidth, power and cooling. As such, the ECAL on-detector front-end electronics performs the first stages of digitisation and triggering.

Fig. 5.1 shows an overview of the on-detector electronics. Each photodetector feeds a 3-channel signal-shaping Multi Gain Pre-Amplifier (MGPA [61]), which consists of a common pre-amplifier stage followed by three independent amplifier stages. This amplification and signal shaping stage is followed by a 4-channel, 40 MHz 12-bit ADC AD41240 [62] which selects the lowest unsaturated gain output from the MGPA. Digital logic internal to the ADC determines whether each gain channel is saturated, and then outputs the data from the channel with the highest gain which was not saturated. The ADC is also designed so that once an input gain range saturates, the ADC returns the next 5 samples without changing the gain. This corresponds to the decay time of the MGPA response, shown in Fig. 5.2.

One Very Front End (VFE) board contains 5 MGPA and ADCs to handle the digitisation of 5 channels. Each set of 5 channels is termed a strip (in the barrel), or a pseudostrip (in the endcap). In the barrel, the strips are aligned in ϕ . In the endcap, the pseudostrips follow a geometry designed to follow the (η, ϕ) segmentation of the HCAL towers and L1 trigger processing. The digital output of the VFE boards (12 bits corresponding to the actual sample and 2 bits representing the gain range chosen) follow the Low Voltage Differential Signalling (LVDS) specification.

5 VFEs (25 channels in a super crystal) are connected to a Front End (FE) board, which consists of parallel data pipelines for all readout channels and the first stages of triggering. A FE board consists of a sequence of custom ASICs designed specifically for the ECAL front end - the FENIX (Front End Intermediate data eXtractor) ASICs. For each VFE channel the data goes into two independent pipelines; the trigger path and the DAQ path. In the DAQ path the raw data are buffered until a L1 trigger accept is received. In the trigger path, each channel is converted from E to E_t , and the transverse energy sum of all 25 channels is computed every 25 ns. Digital filtering and peak finding is then performed, to search for bunch-crossing aligned sample peaks, corresponding to an incident particle. This is termed Trigger Primitive (TP) generation.

These TPs (E_t , Bunch Crossing Identification (BXID) and the Fine Grain Veto Bit (FGVB)) are sent to the Regional Calorimeter Trigger (RCT) every 25 ns, whereas the full data frame (consisting of 10 consecutive ADC samples - 3 pedestal, 2 rising edge, 1 peak and

5.1 Introduction

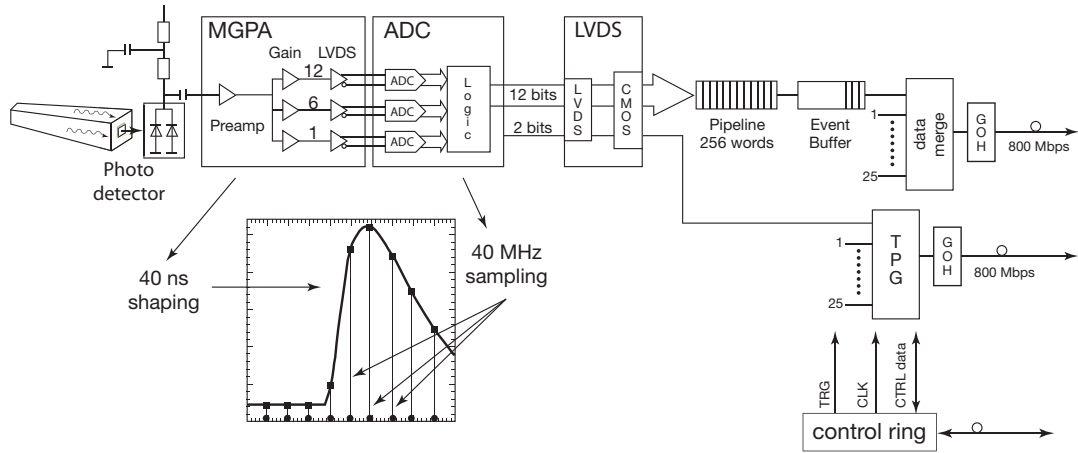


Figure 5.1: Schematic view of the ECAL front-end electronics [50]




Figure 5.2: Example MGPA response

Table 5.1: ADC gain ID settings

Gain ID	Channel
01	$\times 12$
10	$\times 6$
11	$\times 1$
00	$\times 1$ (saturated)

4 decay) for each channel is only read out on receipt of a Level-1 trigger from the Global Trigger. All processing occurs for every bunch crossing at the LHC machine clock frequency of 40MHz, with a latency of 200 ns (8 clock cycles). This technique has been applied to test beam data recorded in 2004 where it has been shown that a timing resolution of less than 1 ns can be achieved for signals with an amplitude greater than 2 GeV [53].

The use of the MGPA and gain switching ADC allows the full resolution of the ADC to be used for a wide range of input signals, roughly equivalent to energies between 500 MeV and 1.5 (3) TeV deposited in the barrel (endcap), with a constant noise term of ± 3 ADC counts. The MGPA has been designed to respond linearly to signals within this input range. For signals above this range, the situation changes dramatically. The MGPA is no longer guaranteed to respond in a linear fashion, and all of the ADC inputs can saturate.

The algorithms implemented in the front end electronics are immutable. Flexibility is given by the choice of parameters, definable per channel:

- Sampling phase
- Pedestal subtraction constant per gain range (P)
- Multiplicative constant per gain range (G)
- Shift coefficient per gain range (S)

The first parameter is used to align all samples across the calorimeter in order to account for particle time-of-flight from the nominal interaction vertex. The remaining three parameters are used to unpack the dynamically compressed ADC output (12 sample bits and 2 gain ID bits) into a linear scale. The output of the per-channel lineariser is given by

$$\frac{G(ADC - P)}{2^{S+2}}.$$

This is implemented in the FENIX by the digital operation

$$\text{Out} = (G * (ADC - P)) \gg (S + 2).$$

These constants can be individually programmed at run-time for each channel and (pseudo)strip. It is interesting to note that in this manner, the conversion from E to E_t is implicit in the algorithms at the linearisation stage through the multiplicative and shift constants, and therefore requires no computationally expensive trigonometric operations.

The behaviour in the saturation regime was not considered fully when the electronics was designed, indeed it was feared that saturating signals would destroy the ability of the BXID within the TPG [53], either due to incorrect lineariser output, or the filter (which

5.2 ECAL TPG response to saturating signals

acts as a differentiator) misidentifying the peak position when presented with sequential saturated samples. An incorrect BXID would result in the entire detector being read out for a collision either before or after the actual collision of interest. Due to the L1 trigger rules, designed to protect the tracker readout, the subsequent bunch crossing could never be read out even if another physics object in the event resulted in a correctly timed trigger. For most sub-detectors, which have a response time integrated over a number of bunch crossings, this would be manageable in the reconstruction. However, the tracker data would be irrecoverable.

5.2 ECAL TPG response to saturating signals

5.2.1 Experimental setup and validation

A RAL endcap MGPA test board was used to measure the saturation of the unity gain channel of an MGPA configured for endcap operation. A software simulation of the ADC gain switching and saturation behaviour was written, which was used to create a set of samples from the real MGPA responses. The FENIX algorithms were applied to find the BXID. A LeCroy Wavesurfer 454 digital oscilloscope was used to record the MGPA response at resolutions of 1 mV and 1 ns.

The MGPA is a charge sensitive amplifier, so charge was injected through a 180 pF capacitor. A function generator delivered square signals varying between 0 and 190 mV. The oscilloscope was triggered on the falling edge of the injection signal. For the VFE used in the endcap, the design saturation value is 16 pC, corresponding to 3 TeV. The gains from each VPT vary significantly depending on their pseudorapidity with respect to the nominal interaction vertex position, and the capacitor tolerance was unknown, so the energies mentioned are approximate. Assuming a calibrated capacitor, the injected charge varies between roughly 3 and 42 pC. The MGPA was therefore tested with injected charges which represent the photodetector output for incident electrons and photons with twice the energy of the electronics design limits.

The filter is realised using a sliding weighted sum of five consecutive samples. These sums are used as inputs of the peak finder. The peak finder is permanently scrutinising a set of three consecutive input samples and making two comparisons in parallel between the middle sample value and the two other neighbouring samples in the sliding window. It outputs 1 when the middle sample value is greater than the others, otherwise it outputs 0.

The three gains ($\times 1$, $\times 6$ and $\times 12$) and the switching of the ADC were simulated. The gain latching behaviour after gain range saturation was also simulated. The functions within the FENIX ASICs (filter and peak finder) were reproduced by the algorithm code taken from the `CMSSW_1_2_0` software package:




Figure 5.3: Response and linearity measurement of the MGPA experimental setup for a range of input charges

```
SimCalorimetry/EcalTrigPrimAlgos/src/EcalFenixAmplitudeFilter.cpp
SimCalorimetry/EcalTrigPrimAlgos/src/EcalFenixPeakFinder.cpp
```

A stand-alone version was written, using the default set of filter weights:

```
weights[0]=0XFFFFFFC1; weights[1]=0XFFFFFFF1; weights[2]=0X28;
weights[3]=0X1B; weights[4]=0X9;
```

The pulse shape exhibits a non linear amplitude behaviour and peak shifting when the input charge is larger than 16 pC (Fig. 5.3). A linearity test performed on the endcap VFE, with injected charges up to 43 pC, shows that the MGPA behaves well in the design region. Each point in the plot is an average of 30 run responses. This test indicates that the charge injection capacitor is fairly accurate and it validates the experimental rig.

5.2.2 TPG simulation and configuration recommendations

Following sampling, the MGPA responses had an ADC saturation simulation applied. The saturation simulation first re-times the signals, ensuring the response is presented to the FENIX algorithms with the timing that would be expected from timing-calibrated VFE electronics. The timing of this is determined from the average peak sample from all responses within the MGPA design range. The timing-corrected signal is then saturated, taking an MGPA output of 1.3 V as representing the saturation point. This voltage corresponds to an arbitrary ADC count of 262143. Finally, samples are taken every 25 ns, representing the clocking period for all the ECAL (V)FE electronics.

Sample sets across the linear range were used to ensure the simulation was behaving as expected in the design range. Fig. 5.4 shows two such results for input charges of 3.4 and 13.9 pC. In both cases the correct bunch crossing is identified, as expected.

For the saturating case, all sample sets up to a representative energy of 5 TeV were

5.3 Level-1 trigger choice for high E_t electron analyses

tested. Fig. 5.5 details the simulation results for injected charges of 15.6 pC (1.46 TeV), 18.9 pC (1.77 TeV), 22.3 pC (2.09 TeV), 25.7 pC (2.41 TeV), 29.1 pC (2.73 TeV) and 42.4 pC (3.98 TeV). In all cases the correct bunch crossing is identified, even in the cases where the saturation occurs on the leading-edge sample rather than at the expected peak itself.

These results assume that for the 4th ADC gain range (all inputs saturated), the subsequent FENIX linearisation behaviour is identical to that in the 3rd gain range. At the time of testing, it was verified that this gain range had not been considered as being relevant to run-time operation, and all lineariser constants were set to 0. For a channel in saturation, the lineariser output would therefore be 0, which would result in an early initial trigger, and a fake parasitic second trigger, as shown in Fig. 5.6. The run-time setup has now been modified such that all VFE settings for the 4th range are automatically identical to those in the 3rd range. This behaviour was also absent from the detector simulation, and has subsequently been added in line with the run-time configuration. With the constants correctly in place, the correct behaviour is assured for real detector running as the 4th ADC input is tied high on the VFE board. This has been confirmed by direct measurement of final production VFE boards.

5.3 Level-1 trigger choice for high E_t electron analyses

With the front end electronics behaviour understood and under control, it is required to determine the most suitable electron trigger to use for high E_t electron analyses. The electron / photon trigger algorithm is described, and its performance investigated.




Figure 5.4: Verification of TPG simulation in the design range. 3.4 pC is equivalent to 320 GeV incident energy and 13.9 pC to 1.3 TeV.

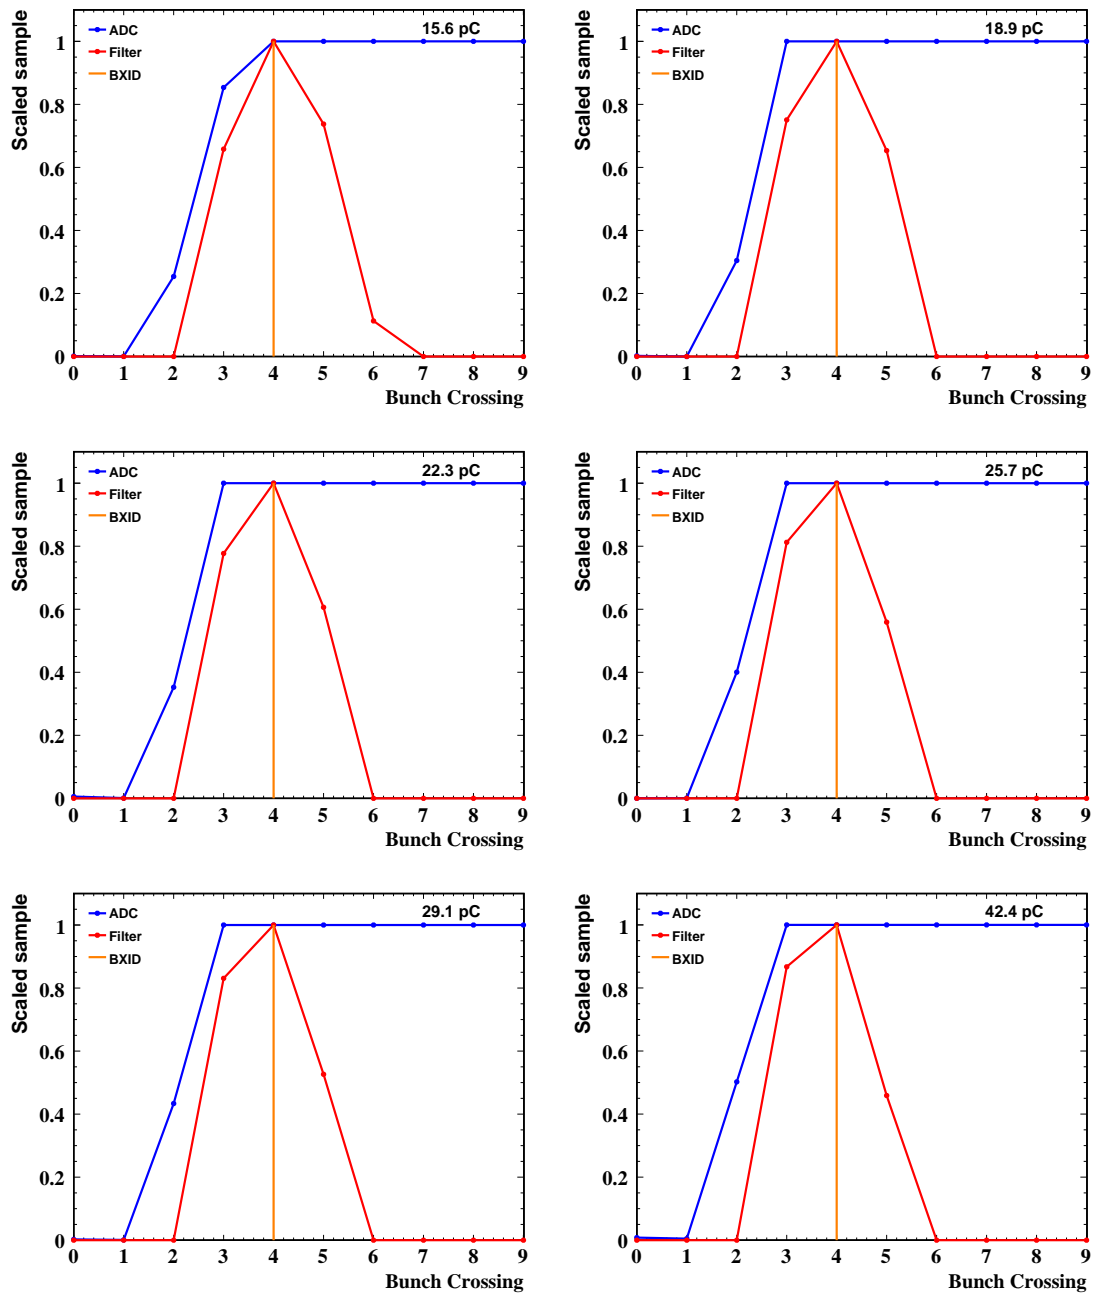


Figure 5.5: TPG simulation for saturating signals. In all cases, the ADC (blue curve) can be seen to saturate and latch. Despite saturation, the filter (red curve) and peak finder (orange curve) still identify the correct bunch crossing.




Figure 5.6: Example of early and late fake parasitic trigger with existing run-time configuration setting saturated channels to 0 after linearisation

5.3.1 Electromagnetic Level-1 trigger algorithm

The L1 EG (Electron / Gamma) ID algorithm takes the trigger primitives for each trigger tower from the ECAL TPG and the equivalent subsystem of the HCAL, and performs EM identification and jet rejection. The algorithm, which is run by the RCT, is based around a sliding window of 3×3 trigger towers, shown in Fig. 5.7, and is defined as [59]:

Electron Candidate if:

- $E_t^{\text{Hit}} > E_t^{N,E,S,W}$
- HEVeto^{Hit} not set

Isolated electron if:

- At least one 5 tower corner of ECAL neighbours is quiet ($E_t < x$)
- HEVeto^{Neighbours} not set

where a 5 tower corner is defined as five neighbouring towers adjacent to two edges of the central hit tower (e.g. NE, N, NW, W and SW would be one of the four possible 5 tower corners).

The HEVeto algorithm is defined as a lookup table based on the ECAL E_t ($E_{t,E}$), HCAL E_t ($E_{t,H}$) and the ECAL Fine Grain Veto Bit (FGVB). The current lookup algorithm is:

- If $E_{t,E} > \text{maxForFGCut}$ then Veto if $E_{t,H}/E_{t,E} > 5\%$
- Else if $E_{t,E} > \text{minActivity}$ then Veto if $E_{t,H}/E_{t,E} > 5\%$ or FGVB set
- Else if $E_{t,E}$ saturated, do not Veto

where x , maxForFGCut and minActivity are noise cuts, currently set to 3 GeV.

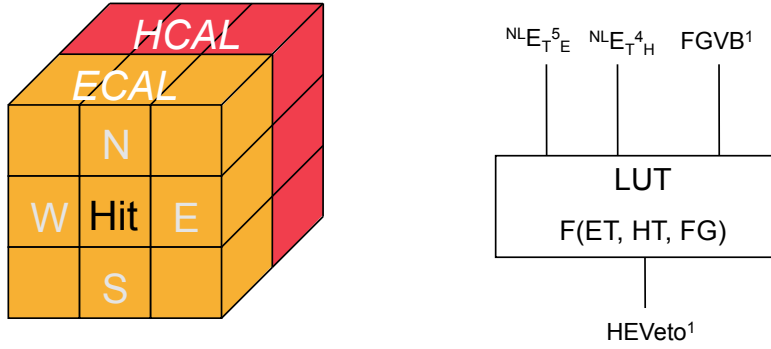


Figure 5.7: Level-1 electromagnetic object trigger overview

5.3.2 Trigger choice

For high E_t electron analyses it is desirable to trigger on events in a wide energy range with a single trigger which has known, stable efficiency. This allows control and efficiency measurements performed using, for example, Z^0 decays to be extrapolated to the signal region without worrying about changes in trigger strategy.

While the isolation criteria in the L1 EM algorithms are designed to reduce jet backgrounds, the inclusion of an H/E cut is cause for concern, as high E_t electrons can be expected to have significant hadronic punch-through. This could cause high-energy isolated electrons to be classified in the non-isolated set of candidates, rendering a naive isolated trigger choice inefficient. To investigate this, a sample of 1.6M single electrons and photons in the kinematic range $0 < E < 1000$ GeV and $-3 < \eta < 3$ was generated, with the full detector and L1 trigger simulation in CMSSW_2_2_9.

The fraction of events passing the trigger, as a function of the generated electron E_t for the $E_t > 25$ GeV isolated trigger (L1_SingleIsoEG25¹) is shown in Fig. 5.8. Electrons in the barrel-encap transition region were excluded to remove geometric acceptance effects from the efficiency measurement. There is rapidly decreasing performance after the turn on region, and this trigger is clearly unsuitable for high E_t electron analyses. The trigger efficiency in the $E_t, |\eta|$ plane is shown in Fig. 5.9.

The same procedure was repeated for the $E_t > 25$ GeV non-isolated trigger (L1_SingleEG25), the results of which are shown in Fig. 5.10. The behaviour is now stable throughout the entire energy regime tested. It is therefore recommended that all analyses with high E_t electrons use a non-isolated L1 EG trigger. These do not require towers adjacent to the central hit tower to pass the HEVeto criteria, as described in §5.3.1. The trigger efficiency can be measured from data using the Tag and Probe techniques discussed in §8.6.1.

¹L1 triggers are named by convention L1.<Multiplicity><Isolation><Object><Threshold>. For example, L1_SingleIsoEG25 represents a single isolated EGamma object trigger with $E_t > 25$ GeV.

5.3 Level-1 trigger choice for high E_t electron analyses




Figure 5.8: Performance of L1_SingleIsoEG25 trigger




Figure 5.9: E_t and $|\eta|$ detail of L1_SingleIsoEG25 performance




Figure 5.10: Performance of L1_SingleEG25 trigger

5.3.3 Jet background rates at $L = 2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$, $\sqrt{s} = 14 \text{ TeV}$

As L1 trigger bandwidth is limited ($\mathcal{O}(100 \text{ kHz})$ at design luminosity), it is required to check that background rates with non-isolated triggers are acceptable with respect to the allowed EGamma trigger bandwidth (nominally 30 kHz).

For the calculation of the jet background rates, a QCD di-jet Monte Carlo simulation was used. These samples are binned by the parton-level p_t of the leading jet, and production cross-sections calculated. The data used was from the `mc-onse1-120` series of generation, which was privately re-reconstructed, and had the trigger emulator run with `CMSSW_1_3_1`. Roughly 20000 events in the p_t bins 20–30, 30–50, 50–80 and 80–120 GeV were used.

Table 5.2 shows the calculated L1 acceptance and rate, assuming an instantaneous luminosity of $\mathcal{L} = 10^{32} \text{ cm}^{-2}\text{s}^{-1}$. The total measured rate is $270 \pm 22 \text{ Hz}$, which is in agreement with previous measurements using minimum bias samples.

The calculation was repeated, varying the L1 trigger p_t threshold, to provide information on the background rate variation to allow suitable choices of the trigger threshold to be made at higher luminosity. This is shown in Fig. 5.11(a), including a fit to the function $\alpha e^{-\beta(x-\gamma)}$, where α, β, γ are the fit parameters. Given the fitted function, the empirical relation between L1 trigger threshold T and di-jet background rate \mathcal{R} is

$$\mathcal{R} = 11.12 \cdot 10^{-7} \cdot e^{-0.12(T-183.7)}, \quad (5.1)$$

Increasing the L1 threshold will, as well as reducing the di-jet background rate, also reduce the efficiency of the trigger to events around the Z^0 resonance, which are used for calibration and efficiency measurements. To quantify this effect, Fig. 5.11(b) shows the relation between the di-jet rate and Z^0 efficiency for a number of L1 threshold settings. The largest cut available is 63 GeV, above which the L1 energy scale saturates.

It is clear that jet backgrounds are under control with non-isolated L1 EG triggers at startup. However, when luminosity increases the trigger threshold will need to be

Table 5.2: Di-jet background rates for SingleEG25 Level-1 trigger

Leading Jet p_t (GeV)	Cross Section (mb)	Acceptance (%)	Rate (Hz)
20 – 30	6.23×10^{-1}	0.042 ± 0.01	26 ± 9
30 – 50	1.63×10^{-1}	0.66 ± 0.05	107 ± 9
50 – 80	2.16×10^{-2}	4.5 ± 0.1	97 ± 3
80 – 120	3.08×10^{-3}	13.0 ± 0.2	40 ± 1

5.4 Conclusions




Figure 5.11: Jet rate (a) and Z^0 efficiency against jet rate (b) of L1_SingleEG trigger

increased to keep the rates under control. This can result in a significant decrease in the number of Z^0 events which are triggered. The Z^0 production cross section is large, such that the reduction in efficiency will still allow accurate efficiency measurements to be made with the triggered events.

5.4 Conclusions

With the existing detector configuration scheme, it is shown that high E_t electrons could cause the detector to be read out on an incorrect bunch crossing due to the ECAL Front End electronics zeroing the E_t of all saturated towers. A detector configuration scheme is proposed whereby the linearisation parameters for the saturating gain range are set equal to those of the unity gain channel. It is shown that this results in correct triggering for single channels with up to at least 4 TeV (8 TeV) incident energy in the barrel (endcap).

To ensure efficient triggering for high E_t electrons, it is shown that a non-isolated L1 EGamma trigger is required. The efficiency for these triggers is constant and near-unity for all electron energies above the trigger threshold. Jet background rates are shown to be acceptable in the LHC startup luminosity regime.

These results are now used in the running on-detector configuration scheme, and the experiment electronics emulation was modified to incorporate the better understanding of the ECAL electronics. All analyses at CMS utilising high E_t electrons rely on these results to ensure efficient and uniform triggering of their analysis datasets.

6. Computing for CMS

At the nominal LHC luminosity of $L = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ at $\sqrt{s} = 14 \text{ TeV}$, CMS is expected to produce $\mathcal{O}(6 \text{ PB})$ of raw data per year [63]. Including reconstructed data, analysis skims and Monte Carlo simulations, the total annual figure is $\mathcal{O}(10 \text{ PB})$. A worldwide computing grid (WLCG) has been implemented to allow the distributed transfer, storage and processing of this data to occur.

Using a custom data persistence layer, a modular event format (the Event Data Model) has been defined. This allows one data format to hold raw, reconstructed and simulated data. Event metadata and conditions data are also handled in a content-agnostic fashion by the CMS software framework.

6.1 The Event Data Model

6.1.1 Data persistence

The CMS Event Data Model [64] (EDM) is built around the concept of an event. Physically, an event is the result of one collision crossing in the LHC. In EDM terms, an event is a C++ object container which stores all raw and reconstructed data for one physical event. Examples of such data are low-level detector hits (i.e. in the tracking system and ECAL), reconstructed subdetector-based elements (i.e. tracks and clusters), and final physics object (i.e. electron candidates).

Using ROOT and Reflex, any C++ class can be stored in the event. The software framework also provides EDM-aware pointers such that classes can store persistable references to other EDM data members in an event. Data within events are uniquely identified by four quantities:

- C++ *type of the data*: e.g. `reco::Track`
- *Module label*: The name of the module which created the data
- *Product instance label*: A user-defined name to allow one module to register many data types

- *Process name*: A label which is common across all modules in one execution of the software framework, defined in the CMSSW configuration file

The persistence layer supports both in-memory and on-disk storage. As such, data persisted to the event by one module can subsequently be read by other modules in the same execution path. The user can choose drop and keep rules at the time of storage, meaning that some data products can be considered transient during a software framework execution run.

In addition to the data itself, the software framework also persists provenance data to each output file. This contains details of configuration parameters used in each run. As data products may be added in each execution run, multiple provenance descriptions may exist in an individual file. As each process name must be unique in an output file, this label is also used to identify the associated provenance.

6.1.2 Event setup records

Along with the actual physics event data, the detector conditions (such as calibration constants, temperatures and alignments) at the time an event was recorded are vital for accurate and reproducible reconstruction and analysis. CMS have implemented a generic relational store of conditions data, based around the concept of an Interval of Validity (IOV). As with the EDM, any C++ class can be persisted to the conditions database, however for performance reasons small, flat data structures are preferred.

Each object registered to the conditions database is tagged as belonging to a certain record type, and having a certain time period (the IOV) for which it is valid. Fig. 6.1 illustrates this for three record types. When the framework reads an event, the EventSetup machinery provides the correct EventSetup records corresponding to a vertical slice in IOV space. For example, events in the timestamp region 1 – 5 will have IOV 1 for Record 1, IOV 2 for Record 2 and IOV 4 for Record 3, whereas events in the timestamp region 10 – 17 will have IOV 1 for Record 1, IOV 3 for Record 2 and IOV 5 for Record 3.




Figure 6.1: Illustration of EventSetup Interval of Validity behaviour

6.2 Data management

6.2.1 Data tiers

Data for CMS is arranged in a number of tiers. While all data for offline analysis is in a common format (described by the EDM), the contents of each event vary depending on the processing stage. For detector data, there are three main formats:

- RAW: Detector data after online formatting, containing the Level-1 and HLT trigger results and, when required for debugging purposes, high-level objects created during HLT processing
- RECO: Reconstructed detector hits / clusters and physics objects (tracks, electrons, etc)
- AOD: Reconstructed physics objects required for analyses

As the EDM is modular, it is also possible to construct FEVT data, which consists of RAW and RECO.

The primary role of RAW data is as input to the Tier-0 and Tier-1 reconstruction passes. It is also invaluable for detector commissioning and understanding. As such, a subset of RAW data is transferred to the CMS Analysis Facility¹ (CAF) to allow rapid analysis. The RAW event size is expected to be 1.5 MB, although this may be larger during startup, where sampling windows and zero suppression will be being tuned. It is estimated that there will be $\mathcal{O}(5 \text{ PB})$ of RAW data produced per year of nominal LHC running.

RECO is the first analysis-ready data format. The event content is such that many reconstruction algorithms can be re-run (for example re-fitting tracks and re-running calorimeter clustering algorithms). The RECO event size is designed to be 250 kB. Allowing for 2 copies of the 1st pass reconstruction and 3 reprocessing runs per year (with updated calibration, alignment and algorithms) yields a volume of 2.1 PB per year.

AOD can be thought of as a stripped-down RECO format. In the long term, it is designed to be the main analysis data format. The event size is designed to be 50 kB, resulting in a volume of 2.6 PB per year due to dataset overlaps, 4 reprocessing versions per year, and storing a full AOD copy at each Tier-1.

6.2.2 Data flow

First step processing occurs at CERN, at the Tier-0. Here, data are converted into a format suitable for offline analysis, sorted into datasets and reconstruction algorithms run.

¹The CAF consists of a 1 PB disk pool and 1000 dedicated batch jobs slots located at CERN. It is designed to facilitate rapid-turnaround calibration, commissioning and analysis tasks.

A tape copy of all RAW data is kept. After initial processing, the RAW and first pass reconstructed data are transferred to a Tier-1 site for custodial storage on tape, and further processing. The final step comes when data that is subsequently processed at a Tier-1 is transferred to a Tier-2 site, and at this stage users may perform analysis. The overall data flow in this tiered architecture is illustrated in Fig. 6.2.

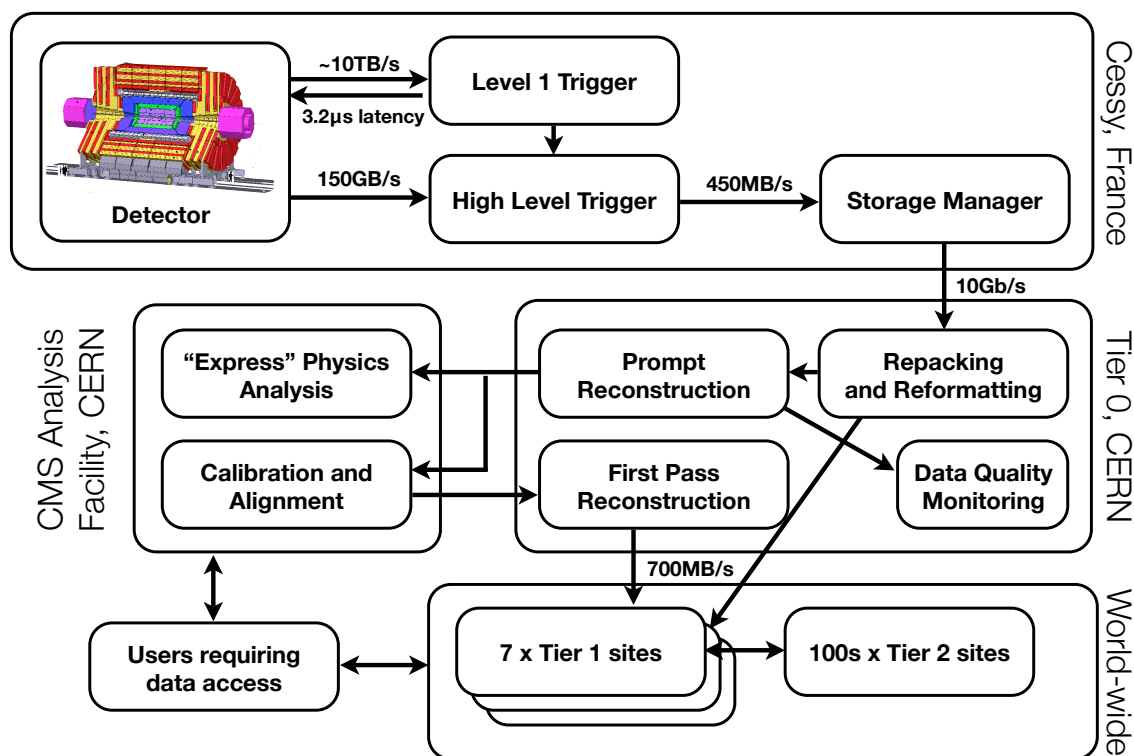


Figure 6.2: Overview of data flow and computing steps for CMS

6.2.3 Transfer technologies

As every site is free to choose a storage technology, a way of mediating transfers between sites running differing technologies is required. All sites expose a Storage Resource Manager (SRM), which provides a common gateway through which transfers can be managed. To provide a hands-off approach to file transfer, all production transfer requests are handled by the File Transfer Service (FTS). While polling of transfer status is available, once a request is made to FTS, no further client interaction is required. FTS handles the mediation of the SRMs at both sites, and manages the transfer once it has started.

When an SRM receives a transfer request, an appropriate transfer handler will be launched. Generally, GridFTP [65] is used, although other protocols are available. GridFTP is a grid-enabled extension to classic FTP, which adds functionality for parallel transfers, renegotiation of TCP parameters depending on a transfer channel performance, and full security within the grid model.

As a GridFTP client, and SRM implementation are native to a particular storage tech-

6.3 Site requirements

nology, but expose a common interface to the world, this setup allows any site to transfer to or from any other site. In practice, the experiments add their own data management and cataloging layer to allow experiment specific requirements (such as transfer channel commissioning) to be enforced. CMS uses PhEDEx [66] (Physics Experiment Data Export) for this purpose, which has now managed petabytes of data transfers for CMS.

6.3 Site requirements

It is evident from the data requirements that a large amount of data must be transferred and stored. The requirements at each tier of computing are designed such that event data can be efficiently transferred, stored and processed to facilitate rapid physics analysis.

6.3.1 Tier-0

The Tier-0 at CERN is a central CMS resource, used only for controlled production purposes. It has four primary roles:

- Acceptance of RAW data
- Reconstruction of RAW data
- Mass storage
- Distribution of RAW and RECO data

In order to meet the requirements of nominal LHC running, the resource requirements of the T0 are 5 Gbs^{-1} external network connectivity (to the Tier-1 sites), 4.6 MSI2K^2 of CPU resource (corresponding to roughly 3500 CPU cores with current hardware), 400 TB of disk cache and 5 PB of tape storage per year.

6.3.2 Tier-1

The Tier-1 centres, of which CMS uses 7, have three main roles:

- Receive and store experimental RAW and RECO data
- (Re)process RAW and RECO data, skim RAW to produce AOD
- Distribute RECO and AOD data to other Tier-1 and associated Tier-2 sites.

²Mega SPECint 2000, a measure of integer calculation performance [67].

Tier-1s are primarily for production activities. However, for 'hot channel' analyses, physicists may be granted limited access to run at Tier-1s.

The resource requirements at a Tier-1 site are 7.2 Gbs^{-1} external network connectivity for imports, and 3.5 Gbs^{-1} for exports, 2.5 MSI2K of CPU resource (roughly 1800 CPU cores, currently), 1.2 PB of disk and 2.8 PB of tape per year. The storage systems needs to support an aggregate throughput of roughly 800 MBs^{-1} .

All CMS Tier-1 sites are now connected to the LHC Optical Private Network (OPN). This provides bi-directional transfer capacity at 10 Gbs^{-1} , with all data routed via CERN. It should be noted that, due to the incoming transfer and local processing requirements, the mass storage system at a Tier-1 needs to be able to sustain data write rates of hundreds of megabytes per second.

6.3.3 Tier-2

The Tier-2 centres have three main roles:

- Receive and store experimental RECO and AOD data
- Run user analysis jobs on the locally stored data
- Generate Monte Carlo events

The resources required for a Tier-2 are variable. However, the nominal CMS requirements are at least 1 Gbs^{-1} external network connectivity for data import and export, 10 Gbs^{-1} aggregate internal bandwidth for disk to CPU for analysis jobs, 0.9 MSI2K of CPU resource (roughly 650 cores) and 200 TB of disk.

Any CMS user may run jobs at any Tier-2, as allowed by local job scheduling (i.e. one user should not be able to block an entire site). Tier-2s may also provide local queue submission for institutional users.

Tier-2s may be federated; that is to say a nominally sized Tier-2 may be declared as the combined resources from two or more physical sites. These may appear to the user as separate sites, or as a single site.

6.4 RAL Tier-1 planning

6.4.1 Site requirements

During the design and planning phase of the CMS computing model, a transfer and storage model was constructed [68]. This matches model parameters, such as event sizes,

6.4 RAL Tier-1 planning

with Tier-1 and Tier-2 allocated resource fractions. The result of this model for RAL during the startup year, and after a year's nominal running (including continuing data taking) is shown in Table 6.1.

Table 6.1: Projected resource requirements for the RAL Tier-1

<i>Resource</i>	<i>Startup requirements</i>	<i>Year 2 requirements</i>
Disk capacity	604.38 TB	1408.41 TB
Tape capacity	1072.43 TB	2937.58 TB
Storage read	201.31 MBs ⁻¹	434.33 MBs ⁻¹
Storage write	280.81 MBs ⁻¹	345.73 MBs ⁻¹
WAN inbound rate	263.34 MBs ⁻¹	327.23 MBs ⁻¹
WAN outbound rate	126.99 MBs ⁻¹	356.45 MBs ⁻¹

6.4.2 Tape system performance

Taking the model figures from §6.4.1 allows the required performance of the tape storage system at the site to be determined. Assuming that CMS has a share of 6 drives at startup, and 8 during nominal running, and that in ideal conditions a drive can recall data at 100 MBs⁻¹ ³, results in a figure of 4 drives available for writing during startup and nominal running. This means that each drive must be able to perform at around 90 MBs⁻¹ in realistic write conditions.

Such a demanding write requirement, including time for mounting, seeking, and unmounting tapes, poses significant problems to the operation of the tape system at RAL. These problems, and solutions to them, will be addressed in the following chapter.

³Such performance has been observed during STEP09 testing [69]. However, correct pre-stage management of data will be required to replicate these results in a production setting.

7. Optimising tape migration performance with the CASTOR HSM

The UK Tier-1 uses the CERN Advanced STORage Manager Hierarchical Storage Manager (CASTOR HSM [70]) to handle all data import, custodial storage, local access and export. The ability to move incoming data to tape (termed tape migration) at a rate equal to or greater than the import rate is a vital requirement of a Tier-1 computing centre to ensure no transfer queues build and that all custodial data is safe. The custodial tape component of CASTOR in use at the UK Tier-1 was not performing as expected. While, as shown in §6.4.2, per-drive performance of 90 MBs^{-1} is required for CMS workflows, the drives were only averaging 16 MBs^{-1} . This under-performance was of severe concern.

A program of work was put in place to optimise the drive performance. This consisted of four major components; the creation of a system monitoring framework, writing a custom I/O bandwidth test suite, optimisation of disk server performance and the optimisation of CASTOR scheduling policies.

The most useful metric of tape rate is the integrated write rate, which is defined as the total volume of data transferred per tape mount divided by the entire time the drive is in use, including tape mount, seek and unmount. The integrated write rates have been found to increase from 16 MBs^{-1} to 90 MBs^{-1} per drive on equally loaded systems as a result of these changes, which are now also used at CERN for all experimental Tier 0 and user analysis CASTOR instances.

7.1 RAL Tier-1 Hardware

7.1.1 Disk servers

Two classes of disk hardware were in use at the time of the study. The properties are detailed in Table 7.1 for the poorly performing hardware, and Table 7.2 for the higher performance hardware. The system specifications were very similar, differing only in the RAID controller used, and the drive configuration.

Table 7.1: Viglen 2006 disk server hardware specifications

<i>Component</i>	<i>Type</i>
Motherboard	Supermicro H8DA8
CPU	2 x AMD Opteron 275
RAM	4GB (4 x 1GB) ECC REG + 4GB Swap
Network	Broadcom Corporation NetXtreme BCM5704 Gigabit
System disk	2 x 250GB SATA WDC (RAID1)
RAID controller	3ware 9550SX 16port PCI-X
Data disk	14 x 500GB WDC WD5000YS (RAID5 + Hot spare)

Table 7.2: Clustervision 2005 disk server hardware specifications

<i>Component</i>	<i>Type</i>
Motherboard	Supermicro H8DA8
CPU	2 x AMD Opteron 275
RAM	4GB (4 x 1GB) ECC REG + 4GB Swap
Network	Broadcom Corporation NetXtreme BCM5704 Gigabit
System disk	2 x 250GB SATA WDC WD2500YS-01SHB0 (RAID1)
RAID controller	Areca 1170 PCI-X
Data disk	22 x 500GB SATA WDC WD5000YS-01MPB0 (RAID6)

7.1.2 Tape system

The tape system is based around two Sun StorageTek SL8500 tape robots, housing a total of 18 Sun StorageTek T10000A Tape Drives. These drives use a 500 GB capacity tape cartridge, and each robot can house up to 10,000 tapes. The maximum read / write rate per drive is 120 MBs^{-1} , which throttles back to 50 MBs^{-1} if the incoming rate is low. Each tape drive is connected by 2 or 4 Gbs^{-1} fibre channel networking to a dedicated tape server. These servers are dual core (single CPU) Intel machines with 4 GB of RAM, and 1 Gbs^{-1} networking.

Of the 18 tape drives, 16 are in production, while two are used for test and certification. Of the production drives, CMS have 2 dedicated (ATLAS and LHCb also have 2 dedicated, with 1 dedicated to other users). All remaining drives are available on-demand.

T10000B drives are now available, and these use the same media, formatted to 1 TB capacity, but with the same R/W rates as the T10000A drives. These are currently being purchased by the site.

7.1.3 Network topology

At the time of development, the overall network topology was as shown in Fig. 7.1. The main element of the network is a switch stack, detailed in Fig. 7.2. Each stack consists of a number of Nortel 5510 switches, which provide $48 \times 1 \text{ Gbs}^{-1}$ ports (for disk servers, worker nodes and tape servers) and one Nortel 5530 switch, which provides a further $48 \times 1 \text{ Gbs}^{-1}$ ports and $2 \times 10 \text{ Gbs}^{-1}$ ports for up/down links. The switches in each stack are interconnected with a circular 40 Gbs^{-1} backbone.

Each stack is connected to a central stack of 5530 switches. Currently, only one 10 Gbs^{-1} link is used per switch, but these could be duplexed if bandwidth became a problem. The central stack has now been replaced with a Force10 10 Gbs^{-1} switch.

External connectivity is provided by two routes, the Optical Private Network (OPN) to CERN, which also routes traffic to the other 10 Tier-1 centres worldwide (of which 6 are used by CMS), and SuperJanet, via the RAL site router. All traffic via the site router bypasses the site firewall through dedicated network routing implemented on the OPN router.

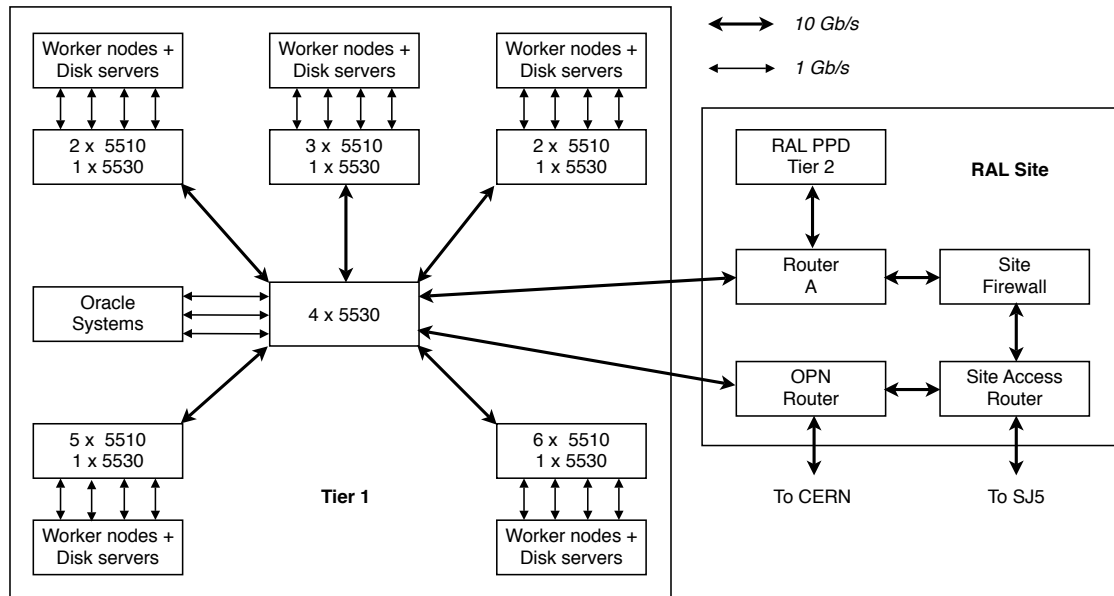


Figure 7.1: UK Tier-1 Network Topology

7.1.4 Transfer performance

External transfers use GridFTP, and each transfer requires a GridFTP process to be running on the disk server. This process itself is not native to CASTOR, rather it uses the RFIO protocol to access the disk server resources. External import transfers therefore potentially perform two network hops internally; first to the GridFTP process and then to the target disk server (or vice-versa for exports). The CASTOR scheduling ensures that

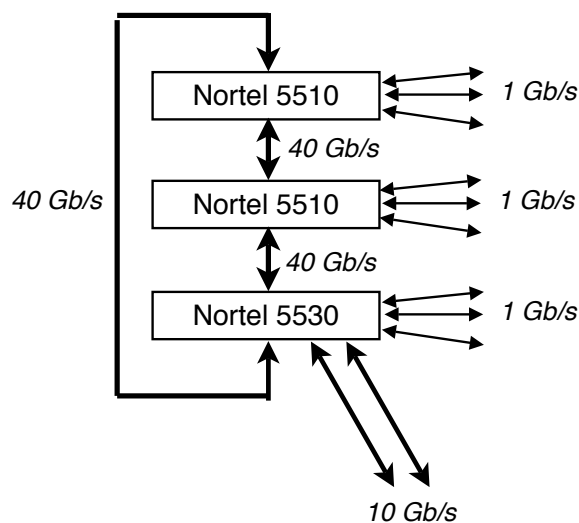


Figure 7.2: Switch stack network detail

the GridFTP process is running on the target disk server to minimise the impact of this behaviour. Despite this, only $\mathcal{O}(6)$ GridFTP processes can be run on one disk server, due to CPU load and memory usage problems¹. Rates per process in the region of $5\text{--}10\text{ MBs}^{-1}$ are observed.

7.2 CASTOR

7.2.1 Architecture overview

CASTOR is a collection of software components that run in different contexts to create a fully functional storage manager. The overall architecture is shown in Fig. 7.3. It has a stateless, database driven architecture. There are four main database schemas:

- *Stager*: Tracks data on disk, and requests for access to it
- *Name Server*: Stores information about each file (permissions, ownership, size)
- *CUPV (CASTOR User Privilege Validation)*: Stores system access permissions
- *VMGR (Volume ManaGeR)*: Stores information about the tape system

The architecture can be summarised as five main components which communicate with the databases, and each other, to realise the working system:

- *Stager*: Runs request handling, request tracking, tape migration candidate allocation and co-ordinates tape movements with the tape system

¹CASTOR 2.1.7 and above have an internal version of GridFTP that rectifies the problems with the external implementation.

7.2 CASTOR

- *Name server*: Handles communication with the Name Server database in response to new or updated files, also handles requests for file location on tape and tape contents
- *Scheduler (LSF)*: Co-ordinates the scheduling of all schedulable jobs, responding to instantaneous system load and limits
- *Disk Server*: Runs copy jobs to handle file I/O, garbage collection to maintain available disk space and communicates with the Stager and Scheduler to execute jobs
- *Tape system*: Responsible for queueing tape requests, managing tape migration and recalls, and tracking the location of tape media

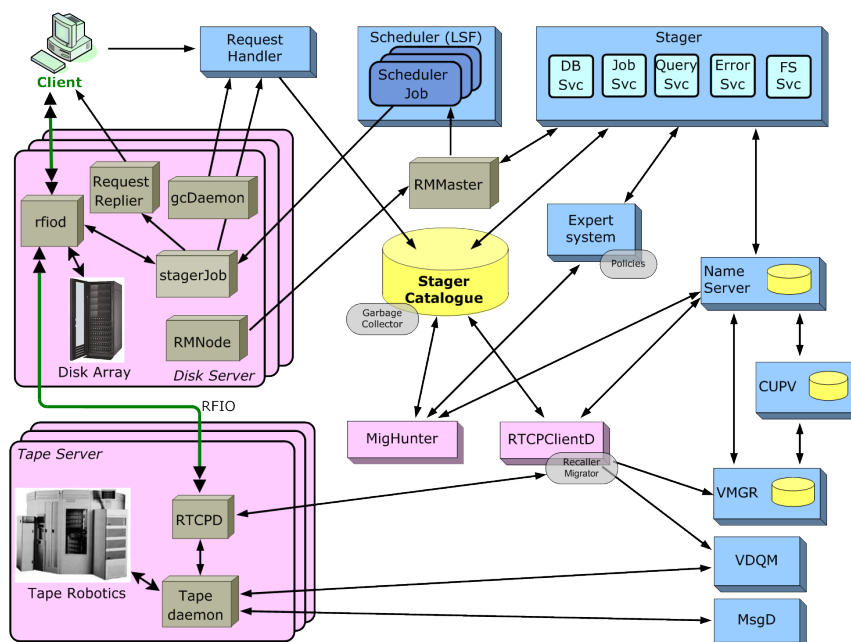


Figure 7.3: CASTOR architecture

Disk servers are grouped into Disk Pools. Disk Pools themselves are members of Service Classes. Requests are handled at the Service Class level. In general, each Disk Pool is a member of one Service Class, however the capability exists for a Disk Pool to belong to more than one Service Class.

7.2.2 Job scheduling

The required behaviour of the CASTOR system imposes requirements on the scheduling system. All possible scheduling scenarios are presented below. Note that not all of these are scheduled by the job scheduler, anything related to the tape system is handled by the Stager (MigHunter and RTCPClient processes) and the tape system (Volume Device Queue Manager - VDQM).

Read from a file in a pool, disk copy exists in pool

Select the best available server which possesses a disk copy in the given pool as the read source.

Read from a file in a pool, disk copy exists in another pool

The read request will be scheduled onto the least loaded server in the required Disk Pool. The file will be synchronously copied from the least loaded server in other pools which possess a disk copy. When this copy operation completes, the file can be read. Note that this operation is scheduled on both the source and destination disk servers.

Read from a file in a pool, no disk copy exists

A tape recall request is sent to the tape system. This recall happens asynchronously and the original read request waits in a pending state. When the recall is completed to an arbitrary server, the original request is allowed to be scheduled as above. Due to this, read requests pending a tape recall do not use up a job slot on the disk server until the recall has taken place.

Read from a file in a pool, disk copy is currently being staged to required pool

The read request will not be scheduled until the disk copy is staged on the Disk Pool. It is important to note that no further disk copies will be created on the Disk Pool. The request waits in a pending state until the stage operation has completed, and therefore does not use up a job slot whilst waiting. When the stage is complete, the request is scheduled as a normal read request, described above.

Read from a file in a pool, disk copy is currently being staged to another pool

A read request will be scheduled onto the least loaded server in the required Disk Pool. This request will then block until a disk to disk copy from another Disk Pool completes. Note that this disk to disk copy itself may have to wait on a tape recall itself. At all times this operation uses the job slot taken by the original read request.

Write a file to a pool

Select the best available server in the given pool as the write destination.

At any stage, if a best available server is not available for a user operation, the job will wait in the queue until the next scheduling iteration. This does not apply for system operations such as tape system disk writes and reads. In all cases, if a file is being staged to a pool, the file is marked as not available for garbage collection until the original request has been serviced. This prevents cyclic recall / delete behaviour with the original request pending for ever in the scenario where the system is run at the garbage collection limit.

The choice of algorithm to choose the best available server in each case is user-definable per Service Class, based on Python scripts that have access to metrics such as server load, jobs running, disk space available and the number of migration streams.

7.2.3 Tape migration

Tape migration in CASTOR follows the concept of streams. A stream corresponds to a route to a tape drive for a given set of tapes (a tape family), which allows data to be grouped to optimise read performance. A stream can write to any tape drive, but only to the relevant tape family. Following a new file being written to CASTOR, the tape migration workflow consists of six stages:

- Initial migration candidate selection and stream allocation
- Stream startup
- Tape queue request
- Tape mount and seek
- Final migration candidate selection and copy until tape full
- Tape rewind and unmount

The initial candidate selection and stream allocation is handled by the `MigHunter` process on the stager, and runs for each Service Class which needs to source tape migrations. Migration policies allow files to be placed into relevant tape families based on Name Server metadata such as file class² and file name. The stream startup process is also handled by the `MigHunter`, and policies can be used to restrict streams starting. This is useful to ensure a reasonable amount of data will be written in one tape mount. The tape queue requests are handled by the `VDQM`, which also manages the queues and the tape robot.

The final candidate selection is performed by a stored procedure in the stager database (`bestTapeCopyForStream`), which selects the next candidate from a stream to write to tape.

The migration and stream policies were not used at RAL at the start of this work, but are key to improving migration performance, as detailed in §7.4.

7.2.4 System monitoring

In order to monitor improvements in the tape system it was required to collect observational data of the storage system in use. Currently, no unified monitoring system for CASTOR exists and therefore a monitoring system was developed. A collection of Python scripts collects and graphs data on the following metrics once per minute per disk server:

²A file class in CASTOR is a numeric identifier attached to directories and files. New files inherit the file class of their parent directory.

- Network I/O rates
- CPU usage (user, system and WAIT I/O)
- Load average
- Running LSF jobs
- Migration queue
- Items pending disk to disk copy and
- Total and free disk space

The system also monitors the instantaneous drive use (read / write per experiment), and integrated tape space used and available.

For collection of LSF job, migration queue, tape and pending disk to disk data the monitoring system remotely executes commands on the CASTOR instance LSF and stager hosts respectively. The results of these commands are parsed into meaningful values. The other metrics are obtained by querying a web-based interface to the Tier-1 Ganglia³ instance. This accepts keywords detailing which hosts and metrics to report, and returns an XML document which is easy to parse.

All metrics are aggregated by Disk Pool to provide an averaged overview of all disk servers within a Disk Pool, and behaviour of shared services such as the tape system. The monitoring system periodically queries CASTOR for what disk servers are present and in which Disk Pools. This ensures that data aggregation is always up to date.

Fig. 7.4 shows an aggregated overview of a Disk Pool importing data at around 170 MBs^{-1} where the tape system is not capable of migrating the imported data fast enough, hence the build-up in the migration queue. Additionally, the CPU load plot shows that about 20% of CPU time is spent in WaitIO. This is indicative of storage system contention. At 04:00, inbound transfers were stopped, and the outgoing network rate can be seen to rise to around 250 MBs^{-1} while the migration queue quickly clears. At this point, three drives were being used for migration from this Disk Pool. The garbage collection system can also be seen to be effective during this rapid migration phase, as available disk space remains constant, rather than decreasing, as candidates migrated to tape become available for removal from disk.

As well as providing historical and instantaneous data for development purposes, the monitoring framework has proved itself indispensable in the day-to-day management of the CMS CASTOR instance at RAL and has been made available to LHCb and ATLAS to allow the Tier-1 team and the experiment contacts to monitor their CASTOR instances.

³Ganglia is a monitoring framework which collects and aggregates system-level information, such as CPU load, network traffic and disk space, from a cluster of hosts.

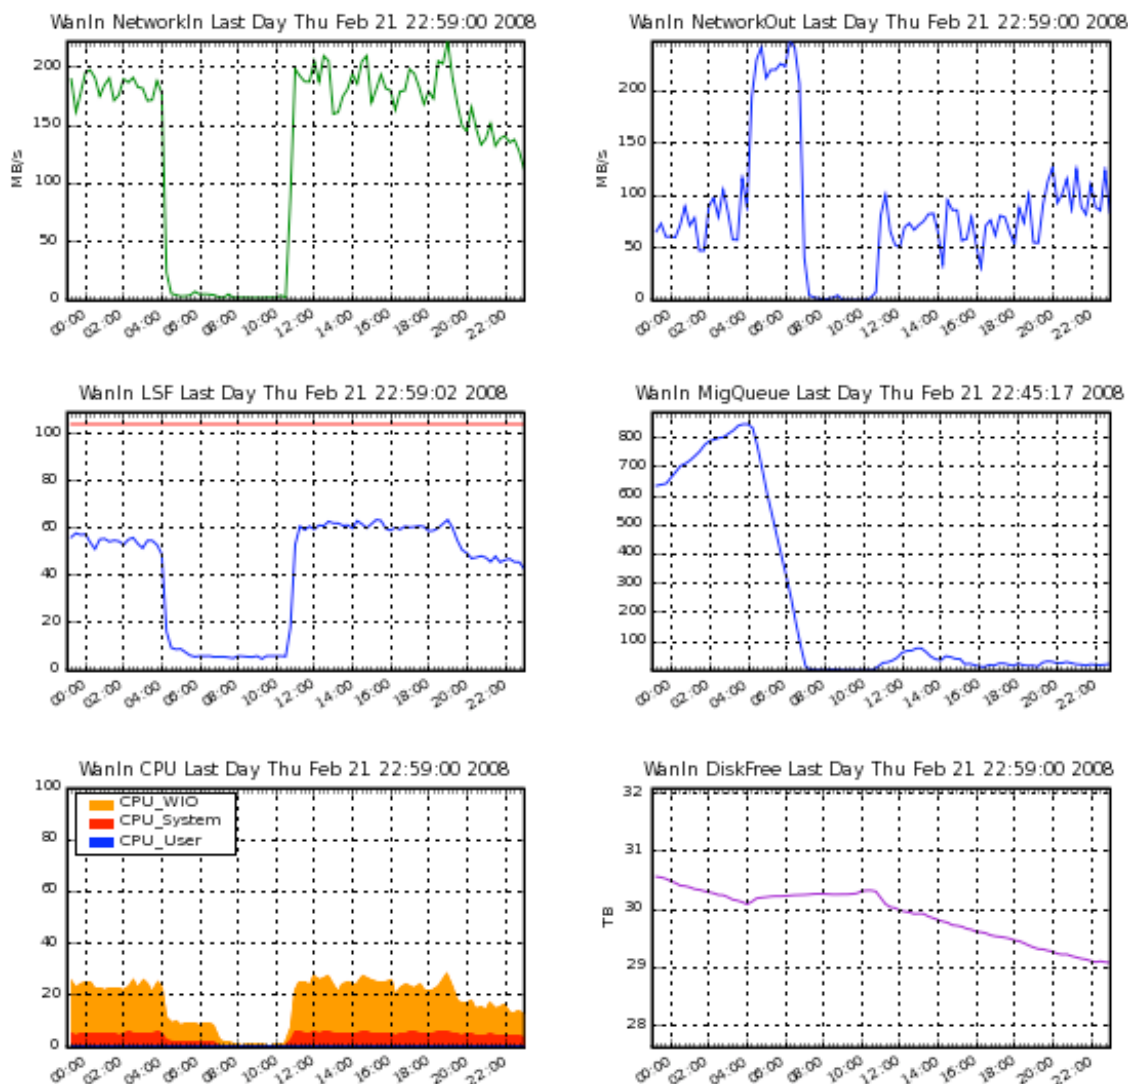


Figure 7.4: Example CASTOR monitoring for WanIn Disk Pool. Top to bottom, left to right: Inbound network rate, outbound network rate, running scheduled jobs, number of files pending tape migration, average CPU load (Orange: WaitIO, Red: system processes, Blue: user processes), remaining disk capacity.

7.2.5 CMS configuration at RAL

Fig. 7.5 provides an overview of the data flow for CMS within the Tier-1. All incoming traffic is routed to the WanIn Disk Pool. This is intended purely as a volatile transfer buffer; data is written in and then migrated to tape as quickly as possible. Data can also be copied directly from this pool to both the WanOut and FarmRead Disk Pools. The FarmRead Disk Pool is accessed by the batch worker nodes for read / write operations during processing jobs. The FarmRead pool can also stage data from tape if it is not already resident on the pool itself, or on the import pool. Data can be copied from all pools, or staged from tape, onto the WanOut pool. This is, analagous with the WanIn pool, purely an export transfer buffer.

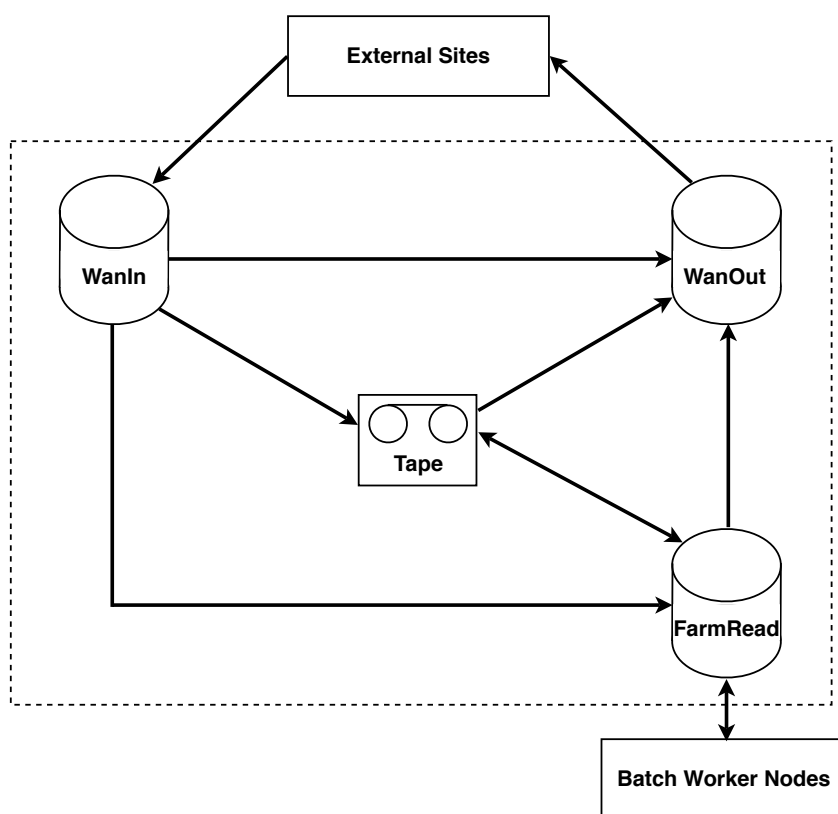


Figure 7.5: Data flow within the Tier-1

The motivation for splitting the pools is that each operation (import, export, farm read / write) has different characteristic I/O patterns. The varying machine configurations and number of running jobs per server required to keep a responsive service can therefore be tuned appropriately. More advanced scheduling options have become available, and at the time of writing CMS are investigating merging the FarmRead and WanOut pools such that disk to disk copies between the two are removed.

7.3 Hardware optimisation

All data being written to tape must be sourced from a disk server. It was clear that the disk server performance for one class of hardware was a major cause of the low tape migration rates. Measurement and optimisation of disk-based I/O systems is notoriously difficult; there are many layers of optimisation and scheduling occurring, namely at the disk drive, RAID controller, block device and filesystem level. The performance of one stage is dependent of the behaviour of those above and below it, and the number of tunable parameters makes for an exceptionally large parameter space to explore if an exhaustive search were to be performed.

As an exhaustive search for an optimal configuration would take an intractable amount of time, a factorised approach was used. A base configuration was prepared, with one

7.3 Hardware optimisation

parameter varied at a time in order to reveal which parameters are responsible for the best gain in I/O rates. This method fails in two cases; where the I/O patterns used are not representative of production dataflow and where parameter variations are correlated.

7.3.1 Linux I/O overview

The I/O systems in Linux are based around two abstraction layers; the Filesystem and Block device. Filesystems are the highest level of abstraction, and are used by applications for the majority of file operations through POSIX calls. Filesystems track files and their associated metadata. Underlying a filesystem is a block device. This is an abstract collection of storage chunks (typically 4 kB in size), which are assumed to be contiguous by the filesystems.

The filesystems do not actually have direct access to the block devices, instead they communicate through a queue provided by the kernel. This allows access to the underlying hardware to be optimised to meet the specific I/O demands of the server. There are four queues available:

- *Complete Fair Queue (CFQ)*: Maintains a per-process I/O queue, and attempts to distribute the available I/O bandwidth equally among all I/O requests
- *Anticipatory Scheduler (AS)*: Introduces a configurable delay before handling requests to aggregate or re-order to improve data access locality
- *Deadline Scheduler (DL)*: Aims to provide near real-time access to the underlying storage, where all processes get an equal time slice in a cyclic fashion
- *NOOP*: Provides a simple first in first out queue and uses the minimal amount of CPU instructions to accomplish requests merging and sorting

The block devices are built on top of the actual physical storage provided by N disk drives and an associated controller. Controllers and drives can also queue, schedule and re-order requests.

7.3.2 Measurement tools

The disk servers are used in a multi-threaded I/O situation, where multiple threads can be reading and writing from disk simultaneously. By measuring at the OS level, effects of network hardware and limitations are not convoluted with raw disk I/O measurements. Freely available I/O profiling tools such as IOZone [71] and Bonny++ [72] were evaluated with a view to emulate this simultaneous I/O behaviour, however none were configurable enough to be useful. For this reason a benchmarking tool was written with the following features:




Figure 7.6: Summary of Linux I/O layers

- Variable read / write block size
- Multiple concurrent or sequential read / write threads
- Ability to read / write from / to multiple filesystems

As Linux utilises all unused RAM as a filesystem cache, to speed up small I/O operations, the tool discards the first few GB of data (configurable), after which the OS is forced to write to disk. Filesystems perform optimally when empty, and therefore all tests were performed on filesystems that were initially 50% full. A filesize of 2 GB was used, with a read / write block size of 1 MiB. These figures roughly correspond to the file sizes CMS regularly uses, and the I/O patterns of GridFTP, which is used to transfer data, and the tape migrator processes. In all cases, at least 100 GB was written / read.

The tool is now used for measurement and acceptance tests for newly procured hardware at the T1.

7.3.3 Baseline measurements and parameter choice

There are two classes of systems in use at RAL, based on Areca and 3Ware RAID controllers respectively. The tape migration rate from the 3Ware systems was of concern, and therefore investigations concentrated on these systems. A baseline configuration was taken as the currently used system configuration, shown in Table 7.3. The meaning of the settings are described in the following sections. The baseline measurements for 0–1 readers and writers are shown in Table 7.4. As can be clearly seen, the read rate deteriorates significantly in the presence of one write thread. This was significantly reducing the tape migration rate from disk servers which also had running write processes.

Due to machine preparation and test run time, each test iteration took a working day to complete. Therefore, due to time constraints, the number of parameters explored was

7.3 Hardware optimisation

limited to those considered to give the maximal gain in performance. The final set, described in the following sections, were decided upon by referring to the Red Hat Kernel Internals guide [73], and by consulting staff within the Tier-1 who have Linux performance tuning experience. In all cases, it should be remembered that the final aim was not to produce optimum theoretical performance, but to produce good enough performance. Essentially, this corresponds to being able to read from disk at network line speed ($\approx 1 \text{ Gbs}^{-1}$).

Table 7.3: Baseline hardware configuration

<i>Parameter</i>	<i>Value</i>
Scheduler	CFQ
blockdev ra	512 blocks
max_sectors_kb	128 kB
nr_requests	128
dirty_background_ratio	6
dirt_ratio	27
RAID array storsav	Protection

Table 7.4: Disk I/O Rates - Baseline configuration

<i>Writers</i>	0		1	
	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	132 ± 52
1	124 ± 19	0	2.5 ± 0.3	127 ± 34

7.3.4 RAID controller and configuration parameters

It was suspected that the I/O rate reduction was due to disk head contention. This is where physical movement of the disk heads between regions being written to, and regions required for writing, results in a lower I/O throughput. The existing RAID configuration was to use all N drives as one physical RAID unit, which was then partitioned into three EXT3 filesystems. This was modified to allocate three RAID units, such that reading and writing could be attempted from different physical sets of disks. The downside of this configuration is that an extra two disks are required for parity, reducing the array volume by 1 TB.

The results of this configuration are shown in Tables 7.5 and 7.6, and show no improvement over the baseline. It was therefore concluded that disk head contention is not

the cause of the read rate problem.

There is one other relevant parameter on the RAID controller itself that can be changed. The 3Ware controllers have the option to locally cache data in a battery-backed cache such that atomic write operations (a single unit of work which is executed by the RAID controller) can complete after power failure. By default these are set to the full safety mode. To check the impact of this on read / write performance, Table 7.7 shows the read / write rates after disabling the feature. No increase in the read rates are observed, but there is a small increase (albeit within errors) in the write rates.

7.3.5 Memory parameters

With the disk head contention theory ruled out, indications in the RAID controller firmware release notes showed that the controller was very poor at context-switching between reads and writes. In communication with the vendor, it was clear that this was not going to be fixed in a useful timeframe. Therefore, ways to mitigate the problem were investigated. The general aim with all the parameters which were modified was to limit the number of absolute write requests getting to the RAID controller, while making the most use of any scheduled read.

Two memory parameters were investigated, with the aim of keeping more data to write in local cache, thus resulting in fewer larger writes as opposed to many small writes.

- `dirty_background_ratio`: The percentage of total dirty (used for I/O caching) system memory at which data will start to be written to disk by a background kernel process (`pdflush`)
- `dirty_ratio`: The percentage of total dirty system memory at which data will start to be written to disk by the process itself

Table 7.8 shows the results of the changes, which negatively impact both read and write rates.

7.3.6 Block device and queue parameters

It was expected that with more work to optimise, the kernel I/O scheduler could improve performance by grouping contiguous requests, reducing read / write context switching in the RAID controller. Additionally, if a read request were scheduled, the opportunity to read a large amount of data during that scheduling iteration should be taken. Three parameters concerning the block device and queue behaviour were tested. These are

- `max_sectors_kb`: Controls what size I/O requests are handed to the controller

7.3 Hardware optimisation

Table 7.5: Disk I/O Rates - Multiple RAID Units (Write /dev/sdb, Read /dev/sdc)

<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	135 ± 49
1	128 ± 14	0	2.2 ± 0.1	117 ± 30

Table 7.6: Disk I/O Rates - Multiple RAID Units (Write /dev/sdc, Read /dev/sdb)

<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	133 ± 43
1	126 ± 22	0	2.6 ± 0.2	122 ± 27

Table 7.7: Effect on disk I/O rates of setting RAID array storsav to performance: Completely Fair Queue scheduler, blockdev ra 512, max_sectors_kb 128, nr_requests 128, dirty_background_ratio 6, dirty_ratio 27. All values in MB/s

<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	172 ± 62
1	124 ± 19	0	2.8 ± 0.2	163 ± 52

Table 7.8: Effect on disk I/O rates of varying memory parameters: Completely Fair Queue scheduler, blockdev ra 512, max_sectors_kb 128, nr_requests 128, dirty_background_ratio 20, dirty_ratio 60. All values in MB/s

<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	39 ± 20
1	123 ± 20	0	1.2 ± 0.2	71 ± 31

- `nr_requests`: Controls the maximum number of requests allowed in a queue
- `blockdev ra` (*Block device read-ahead*): Controls how much data to prospectively read and cache on each read operation from the block device

In addition, the actual I/O scheduler used was varied along with the block device read-ahead to ascertain whether the default, `CFQ`, is the best choice.

Table 7.9 shows the results of varying `max_sectors_kb`. Reducing this setting was intended to schedule smaller write requests, allowing more reads to occur. As can be seen, there was no change from the baseline measurements.

Table 7.10 shows the results of varying `nr_requests`. Increasing this setting was intended to allow the controller to schedule pro-actively, as it will have more information available as to pending work. Anecdotal evidence with previous 3Ware hardware suggested that making this at least twice as big as the controller queue length (254 requests) would have positive results. Clearly, this setting does not impact on read / write rates with the hardware under test.

Table 7.9: Effect on disk I/O rates of varying `max_sectors_kb`: Completely Fair Queue scheduler, `blockdev ra 512`, `max_sectors_kb 64`, `nr_requests 128`, `dirty_background_ratio 6`, `dirty_ratio 27`. All values in MB/s

<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	133 ± 51
1	129 ± 19	0	2.6 ± 0.1	126 ± 36

Table 7.10: Effect on disk I/O rates of varying `nr_requests`: Completely Fair Queue scheduler, `blockdev ra 512`, `max_sectors_kb 128`, `nr_requests 512`, `dirty_background_ratio 6`, `dirty_ratio 27`. All values in MB/s

<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	131 ± 49
1	126 ± 20	0	2.7 ± 0.3	127 ± 33

The results of varying `blockdev ra` are shown in Tables 7.11 (`CFQ`), 7.12 (`AS`), 7.13 (`Deadline`) and 7.14 (`NOOP`). On the whole, the schedulers perform equally well, favouring a higher read-ahead to improve read rates in a multi-threaded environment.

There is one anomalous point in the `AS` testing, at the smallest read-ahead size (256 blocks). To see if the improvement were artificial in the test environment, all disk servers

7.3 Hardware optimisation

in the WanIn pool were re-configured with the AS scheduler and `blockdev ra` of 256 blocks. With production traffic flows, the average throughput on the Disk Pool halved. Such a result shows the difficulty of testing for production I/O patterns in an artificial environment. The main difference is that the GridFTP processes on the WanIn can write data out of order, whereas the test suite writes in order, which allows the AS to optimise the write patterns, giving reading a chance to obtain controller time.

7.3.7 Configuration recommendations

It is recommended that all 3Ware based systems have the `blockdev ra` set to 16 KiB. This provides a significant performance gain by allowing a large amount of data to be read when an I/O read request does get scheduled.

The interaction of the I/O scheduler with the data patterns seen in production is not fully understood. It was hoped, from the performance measurements, that changing to the Anticipatory scheduler would improve read rates, however this was not the case. It is recommended to use the Linux default scheduler, CFQ, due to the apparent instability of AS with respect to tuning parameters.

It was found that changing the Linux memory protection parameters negatively impact on both read and write performance, and therefore the current settings are recommended to be kept.

Changing the RAID controller cache settings also does not impact performance, therefore the secure mode whereby any outstanding data can be flushed to disk automatically in the event of server failure is recommended.

Following these recommendations, the 3Ware hardware was tested with both the CFQ and AS schedulers in a variety of threading environments. The results are shown in Tables 7.15 and 7.16 using CFQ and 7.17 and 7.18 using AS. These tables are used to determine read and write job scheduling limits.

Optimising tape migration performance with the CASTOR HSM

Table 7.11: Effect on disk I/O rates of varying blockdev read ahead size: Completely Fair Queue scheduler, max_sectors_kb 128, /sys/block/sda/queue/nr_requests 128, dirty_background_ratio 6, dirty_ratio 27. All values in MB/s

<i>Read ahead 256 blocks</i>				
Writers	0		1	
Readers	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	130 ± 47
1	110 ± 12	0	1.5 ± 0.1	129 ± 35

<i>Read ahead 512 blocks</i>				
Writers	0		1	
Readers	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	128 ± 47
1	130 ± 17	0	2.6 ± 0.3	126 ± 33

<i>Read ahead 1024 blocks</i>				
Writers	0		1	
Readers	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	130 ± 48
1	144 ± 18	0	4.5 ± 0.2	123 ± 32

<i>Read ahead 2048 blocks</i>				
Writers	0		1	
Readers	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	128 ± 46
1	219 ± 23	0	14 ± 2	107 ± 25

<i>Read ahead 4096 blocks</i>				
Writers	0		1	
Readers	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	128 ± 43
1	296 ± 27	0	24 ± 7	101 ± 28

<i>Read ahead 8192 blocks</i>				
Writers	0		1	
Readers	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	127 ± 46
1	340 ± 29	0	37 ± 12	91 ± 32

<i>Read ahead 16384 blocks</i>				
Writers	0		1	
Readers	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	128 ± 49
1	321 ± 37	0	40 ± 14	81 ± 19

7.3 Hardware optimisation

Table 7.12: Effect on disk I/O rates of varying blockdev read ahead size: Anticipatory scheduler, max_sectors_kb 128, /sys/block/sda/queue/nr_requests 128, dirty_background_ratio 6, dirty_ratio 27. All values in MB/s

<i>Read ahead 256 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	134 ± 41
1	106 ± 15	0	45 ± 8	31 ± 17

<i>Read ahead 512 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	133 ± 41
1	131 ± 17	0	4.9 ± 0.8	120 ± 40

<i>Read ahead 1024 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	130 ± 41
1	138 ± 18	0	4.7 ± 0.7	114 ± 27

<i>Read ahead 2048 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	128 ± 38
1	218 ± 23	0	10 ± 1	105 ± 24

<i>Read ahead 4096 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	127 ± 39
1	293 ± 24	0	21 ± 2	95 ± 21

<i>Read ahead 8192 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	132 ± 41
1	331 ± 34	0	27 ± 5	92 ± 23

<i>Read ahead 16384 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	130 ± 38
1	329 ± 37	0	29 ± 6	98 ± 36

Optimising tape migration performance with the CASTOR HSM

Table 7.13: Effect on disk I/O rates of varying blockdev read ahead size: Deadline scheduler, max_sectors_kb 128, /sys/block/sda/queue/nr_requests 128, dirty_background_ratio 6, dirty_ratio 27. All values in MB/s

<i>Read ahead 256 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	132 ± 43
1	103 ± 15	0	1.5 ± 0.1	125 ± 32

<i>Read ahead 512 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	130 ± 40
1	124 ± 19	0	2.6 ± 0.1	122 ± 29

<i>Read ahead 1024 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	131 ± 40
1	138 ± 18	0	4.6 ± 0.2	119 ± 29

<i>Read ahead 2048 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	132 ± 41
1	217 ± 23	0	14 ± 3	103 ± 24

<i>Read ahead 4096 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	120 ± 44
1	301 ± 116	0	26 ± 6	96 ± 23

<i>Read ahead 8192 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	133 ± 42
1	340 ± 32	0	41 ± 9	81 ± 16

<i>Read ahead 16384 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	132 ± 42
1	321 ± 41	0	44 ± 13	78 ± 17

7.3 Hardware optimisation

Table 7.14: Effect on disk I/O rates of varying blockdev read ahead size: Noop scheduler, max_sectors_kb 128, /sys/block/sda/queue/nr_requests 128, dirty_background_ratio 6, dirty_ratio 27. All values in MB/s

<i>Read ahead 256 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	136 ± 40
1	103 ± 16	0	1.7 ± 0.1	118 ± 31

<i>Read ahead 512 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	131 ± 40
1	296 ± 30	0	3.1 ± 0.2	117 ± 26

<i>Read ahead 1024 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	137 ± 44
1	141 ± 18	0	5.3 ± 0.2	111 ± 24

<i>Read ahead 2048 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	134 ± 42
1	221 ± 23	0	18 ± 3	94 ± 19

<i>Read ahead 4096 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	131 ± 40
1	296 ± 30	0	31 ± 5	84 ± 14

<i>Read ahead 8192 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	134 ± 40
1	339 ± 39	0	50 ± 11	78 ± 23

<i>Read ahead 16384 blocks</i>				
<i>Writers</i>	0		1	
<i>Readers</i>	<i>Read</i>	<i>Write</i>	<i>Read</i>	<i>Write</i>
0	0	0	0	134 ± 42
1	332 ± 42	0	56 ± 16	70 ± 12

Table 7.15: 3Ware disk I/O performance per thread: CFQ Scheduler

Writers	0		1		2		3		4		5		6		7		8		10		25		40		
	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	
0	0.0	0.0	0.0	110.7	0.0	55.3	0.0	22.0	0.0	21.6	0.0	23.4	0.0	17.9	0.0	0.0	0.0	0.0	10.0	0.0	0.0	3.2	0.0	1.6	
1	295.8	0.0	35.9	79.3	31.5	45.5	29.1	19.3	25.6	13.3	20.3	13.1	25.5	8.2	25.0	6.8	24.4	5.8	22.4	4.4	19.3	1.5	18.9	0.9	
2	195.3	0.0	26.2	73.9	20.8	29.0	20.9	16.0	19.4	11.0	14.1	9.8	18.9	6.7	18.5	5.5	18.4	4.8							
3	120.4	0.0	15.2	75.8	16.8	28.5			14.9	11.1	17.4	9.1	15.3	7.6	20.9	6.6	17.6	4.6							
4	90.4	0.0	19.1	67.3	11.7	32.0	11.1	15.9	16.5	12.5			15.6	7.3			18.8	5.0							
5	82.8	0.0	13.4	70.0					6.8	9.6							18.0	5.3	6.9	4.0	8.5	1.6	9.6	1.0	
6	57.4	0.0							12.0	12.6															
7		0.0					11.5	16.4																	
8		0.0	19.3	67.9																					
10	23.0	0.0																							
25	7.5	0.0																			1.3	1.4			

Table 7.16: 3Ware disk I/O aggregate multithreaded performance: CFQ Scheduler

Writers	0		1		2		3		4		5		6		7		8		10		25		40	
	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W
0	0.0	0.0	0.0	110.7	0.0	110.5	0.0	65.9	0.0	86.3	0.0	116.8	0.0	107.7	0.0	0.0	0.0	0.0	0.0	100.3	0.0	79.7	0.0	63.2
1	295.8	0.0	35.9	79.3	31.5	91.0	29.1	58.0	25.6	53.4	20.3	65.5	25.5	49.1	25.0	47.3	24.4	46.2	22.4	43.7	19.3	37.1	18.9	36.2
2	390.7	0.0	52.5	73.9	41.5	58.1	41.8	48.0	38.9	44.1	28.2	49.2	37.8	40.2	37.0	38.6	36.8	38.4						
3	361.3	0.0	45.6	75.8	50.4	57.0			44.8	44.4	52.1	45.6	45.5		62.7	46.0	52.9	37.1						
4	361.7	0.0	76.5	67.3	46.8	63.9	44.4	47.6	66.0	50.1			78.2	44.0			75.2	39.9						
5	414.1	0.0	67.0	70.0					34.0	38.5							90.1	42.7	34.7	40.2	42.4	38.9	47.9	40.1
6	344.5	0.0					80.7	49.3	72.1	50.6														
7		0.0																						
8		0.0	154.3	67.9																				
10	230.1	0.0																						
25	187.2	0.0																			33.3	34.7		

Table 7.17: 3Ware disk I/O performance per thread: AS Scheduler

3Ware disk I/O performance per thread, AS Scheduler, blockdev ra 256, max-sectors.kb 128, nr_requests 128, dirty_background_ratio 6, dirty_ratio 27	0		1		2		3		4		5		6		7		8		10		25		40	
	Writers	Readers	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W
0	0.0	0.0	0.0	117.8	0.0	54.6	0.0	33.2	0.0	22.4	0.0	17.5	0.0	14.7	0.0	12.7	0.0	10.2	0.0	8.6	0.0	2.6	0.0	1.4
1	102.0	0.0	42.4	26.7	38.7	12.5	37.0	8.2	29.9	6.8	40.2	3.9	33.7	3.6	37.9	2.2	39.1	1.8	35.6	1.6	37.5	0.4	37.1	0.3
2	32.2	0.0	19.2	24.5	18.3	10.5	18.3	5.9	18.3	3.8	18.3	2.7	18.6	2.1	18.9	1.8	18.7	1.5						
3	17.2	0.0	12.5	24.2	11.5	10.0	11.7	5.0	11.8	3.8	12.0	2.8	11.9	2.1	12.0	1.8	12.2	1.4						
4	13.4	0.0	9.4	24.5	8.7	9.9	8.7	6.1	8.7	3.9	8.7	2.6	8.9	2.1	9.1	1.7	9.3	1.3						
5	11.9	0.0	7.8	24.3	6.8	10.7	7.2	5.9	7.1	3.8	7.3	2.7	7.5	2.1	7.4	1.8	7.6	1.3	7.2	1.0	7.3	0.3	7.4	0.2
6	8.2	0.0	6.3	24.1	5.9	10.5	5.9	5.0	6.0	3.9	6.0	2.6	6.4	2.0	6.3	1.7	6.4	1.4						
7	7.7	0.0	5.8	24.8	5.5	11.0	5.3	6.2	5.6	3.7	5.6	3.0	5.4	2.2	5.6	1.8	5.4	1.7						
8	6.9	0.0	5.0	23.6	4.7	10.7	4.6	6.1	4.7	3.9	4.7	2.9	4.7	2.3	4.8	1.7	4.7	1.6						
10	4.3	0.0																						
25	1.6	0.0																						

Table 7.18: 3Ware disk I/O aggregate multithreaded performance: AS Scheduler

3Ware disk I/O aggregate multithreaded performance, AS Scheduler, blockdev ra 256, max-sectors.kb 128, nr_requests 128, dirty_background_ratio 6, dirty_ratio 27	0		1		2		3		4		5		6		7		8		10		25		40	
	Writers	Readers	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W	R	W
0	0.0	0.0	0.0	109.3	0.0	99.7	0.0	89.4	0.0	89.4	0.0	87.5	0.0	88.1	0.0	89.1	0.0	81.4	0.0	85.9	0.0	65.8	0.0	54.3
1	102.0	0.0	42.4	26.7	38.7	25.0	37.0	24.6	29.9	27.2	40.2	19.5	33.7	21.7	37.9	15.5	39.1	14.7	35.6	15.9	37.5	10.9	37.1	13.3
2	64.4	0.0	38.4	24.5	36.5	20.9	36.6	17.8	36.7	15.1	36.6	13.6	37.2	12.6	37.9	12.3	37.4	11.7						
3	51.7	0.0	37.4	24.2	34.4	20.1	35.2	15.1	35.4	15.0	36.0	14.2	35.6	12.9	35.9	12.7	36.7	11.2						
4	53.5	0.0	37.6	24.5	34.8	19.8	34.8	18.2	34.9	15.6	35.7	13.1	36.8	12.5	36.4	11.8	37.1	10.3						
5	59.6	0.0	39.1	24.3	34.2	21.4	35.8	17.6	35.3	15.2	36.5	13.6	37.4	12.4	36.9	12.3	37.8	10.5	36.0	9.8	36.7	8.1	37.1	9.5
60	49.4	0.0	37.7	24.1	35.5	20.9	35.3	14.9	35.8	15.4	36.1	13.1	38.5	12.3	38.0	11.7	38.6	11.4						
7	54.1	0.0	40.3	24.8	38.5	22.1	37.2	18.5	39.0	14.9	39.0	14.9	37.6	13.5	39.0	12.9	38.0	13.4						
8	54.9	0.0	40.0	23.6	37.8	21.5	36.7	18.2	37.8	15.7	37.7	14.7	37.8	13.7	38.0	12.2	37.4	12.6						
10	42.6	0.0																						
25	39.3	0.0																						

7.4 CASTOR optimisation

Having optimised the disk servers for read performance, the next stage was to investigate what optimisations could be performed within CASTOR itself. As the tape and job scheduling operations are independent, the main aim was to reactively schedule in response to tape behaviour, stopping CASTOR scheduling write jobs to a disk server which was sourcing files for tape migration. Two optimisations were required; firstly modifying how LSF schedules write jobs to disk servers, and secondly changing the algorithm used to select migration candidates.

7.4.1 Job scheduling

CASTOR schedules jobs at the filesystem level, as disk servers generally have more than one filesystem due to limitations in filesystem sizes when using EXT3. While this is a legacy issue, with more modern filesystems supporting very large capacities, EXT3 continues to be used. Filesystems may be placed into one of three statuses; `PRODUCTION`, `DRAINING` or `OFFLINE`. Those in `PRODUCTION` may have any jobs scheduled to them, `DRAINING` only tape migration and source disk to disk copy jobs, while nothing may be scheduled to an `OFFLINE` filesystem.

The default scheduling algorithm attempts to evenly distribute the number of running jobs between all available filesystems in a Disk Pool. This was easily modified such that no jobs would be scheduled to a filesystem if any migration streams were running on that filesystem, or other filesystems on the same disk server.

It is important to note that currently running jobs can not be cancelled, and disk servers can not be pro-actively cleared of running jobs before starting migration. The aim was, therefore, to allow the modified scheduling policy to operate in a self-optimising fashion, whereby running write jobs will drain, allowing the read operations to gain improved I/O bandwidth. Due to removing the ability to schedule new write jobs to a migrating disk server, the total import capacity of the Disk Pool will be reduced. However, this must be balanced with the T1 requirement of storage of custodial data on tape in an efficient manner.

Despite the new scheduling policy being put in place, no improvement in migration rates was observed. On investigation, it was determined that migration streams were only present on a disk server for a short while, even in the presence of a large number of migration candidates. As the reactive scheduling requires migration streams to be continuously present in order to allow existing write jobs to finish, the inconsistent migration stream numbers were not allowing the disk servers to drain of write jobs as intended.

7.4.2 Migration candidate selection

As described in §7.2.3, the choice of which migration candidates to copy to tape at a given time is made by the `bestTapeCopyForStream` stored procedure in the Stager database. The pre-existing algorithm copies alternately from two distinct filesystems on separate disk servers into one tape server. The motivation for this is that disk reads may be slow, therefore copying two files at a time into local cache will allow the files to stream sequentially to tape at a reasonable rate. However, this breaks down when the files are larger in size than half of the available tape server cache. With the systems at RAL, this means that files of 2 GB or over will fail, as the tape copy system will stop a migration stream if there is not enough space on the tape server to copy the file.

This previous optimisation was therefore identified as the cause of the unstable stream count responsible for the failure of the job scheduling changes to improve migration rates. A new procedure was developed, with the aim of draining first a filesystem, and then any other filesystems, on a given disk server. This would allow the modified job scheduling algorithms to take advantage of a stable migration count on a disk server.

The modified algorithm is shown in Fig. 7.7. With the new algorithm in place, disk servers were observed draining of write jobs, while read performance increased to nearly 100 MBs^{-1} , as shown in Fig. 7.8.

It should be noted that the algorithm will fail in the case where there are more running migrations streams than disk servers. However, this is not the case for LHC-scale production systems. In addition, the algorithm assumes there are candidates available to migrate on the Disk Pool. This is ensured by the calling logic, which only executes `bestTapeCopyForStream` if there are migration candidates available in a stream.

7.4.3 Stream policies

The algorithms and scheduling policies as presented operate in a self-optimising fashion, whereby building migration queues allow the job scheduling operations to drain jobs, and therefore the tape migration rate increases until files are drained. However, while the peak tape write rate after draining will be high, the integrated rate will be lower. The result of this is that tape drives are stopped from performing other useful work (such as recalls, and read / write for other users).

CASTOR provides a solution to this in the form of stream policies. These allow control over whether to attach a migration stream to the tape drive based on the number of queued migration candidates and the volume of data to migrate. A stream policy was written which only starts up a stream when a user-definable volume of data is queued.

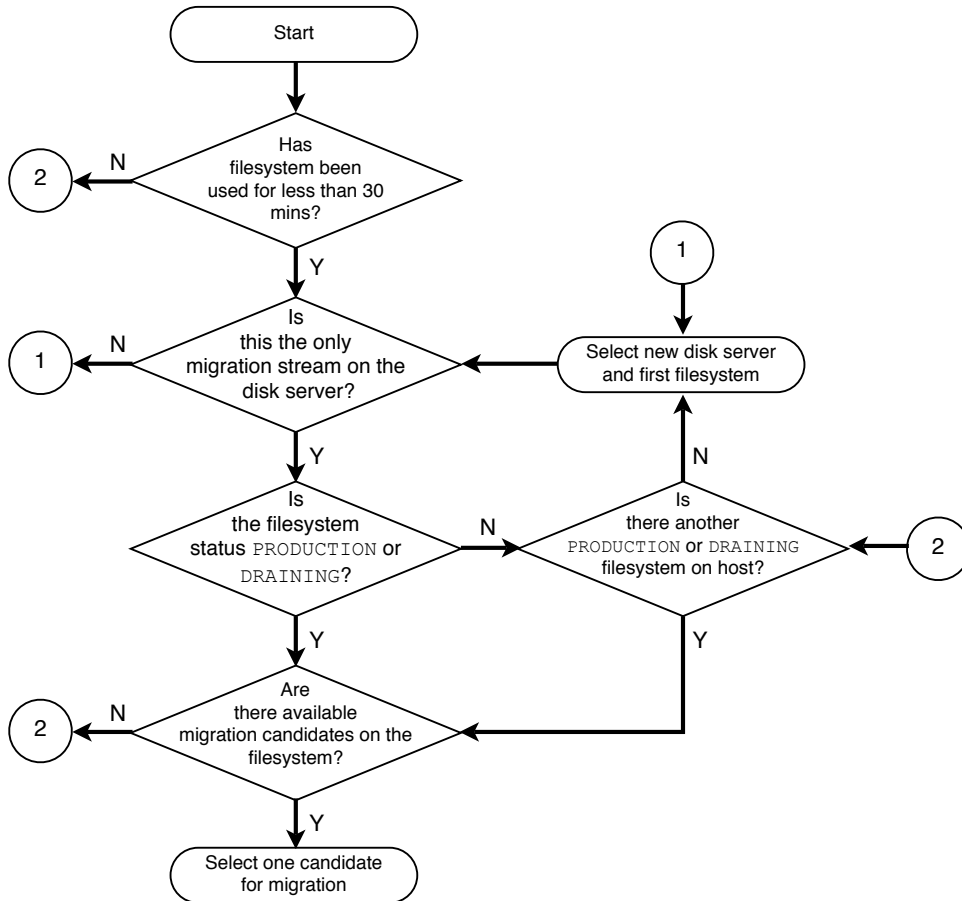


Figure 7.7: Flowchart of disk-draining bestTapeCopyForStream algorithm




Figure 7.8: Example disk server network usage with job draining operational

7.5 Performance during STEP09

Table 7.19: Tape write performance for CMS during STEP09

<i>Day</i>	<i>Rate (MBs⁻¹)</i>	<i>Write mounts</i>	<i>Volume per mount (GB)</i>
1	16 ± 5	160	3
2	12 ± 3	115	3
3	8 ± 1	72	2
4	11 ± 1	361	2
5	18 ± 1	439	4
6	21 ± 2	650	4
7	15 ± 3	369	3
8	13 ± 3	383	17
9	47 ± 7	249	37
10	35 ± 5	186	11
11	45 ± 6	140	21
12	62 ± 4	71	44
13	50 ± 11	44	114
14	84 ± 1	22	359
15	90 ± 1	21	464

In stable streaming operation, this value should be equal to

$$V_{\text{Start}} = \frac{n_{\text{Drives}} C_{\text{Tape}}}{n_{\text{Families}}},$$

where n_{Drives} is the number of drives assigned to the Service Class, C_{Tape} the capacity of each tape and n_{Families} the number of tape families in use. Following this queuing strategy, all mounted tapes will be full at the end of a migration round, reducing the impact of tape handling time on the effective tape system migration rate.

7.5 Performance during STEP09

In June 2009, a cross-experiment test of the global LHC computing systems was performed. This provided the perfect opportunity to measure the performance of the tape system with the optimisations in place.

The measured values for each day are given in Table 7.19. The initial measurements were disappointing, with daily integrated tape rates of around 15 MBs⁻¹, as had been previously seen. However, this was found to be due to a bug in CASTOR which was not respecting the stream policy.

After this was fixed (on the evening of day 8), with the queued volume threshold set to 500 GB, rates increased to $\mathcal{O}(50 \text{ MBs}^{-1})$. When this threshold was reached, CASTOR was configured to mount 4 tapes, resulting in 125 GB per tape being written. The threshold was increased to 2 TB on the evening of day 13, which allowed the integrated tape write

rate to reach nearly 90 MBs^{-1} for the last two days.

The network behaviour verifies that the optimisations were working as expected. Figs. 7.9–7.11 show the network traffic for 12 of the tape servers available for use in a 12 hour window. Fig. 7.9 shows drives dedicated to CMS, and validates that the disk servers are providing files at 100 MBs^{-1} to the tape servers. Fig. 7.10 shows shared drive usage, and the difference in write rates between CMS and non-CMS use is clearly visible. It was verified from tape logs that the peaks at 100 MBs^{-1} are due to CMS usage only. Fig. 7.11 shows tape drives dedicated to experiments other than CMS. The write rates are unstable and lower with respect to CMS activity.

7.6 Conclusions

Through hardware configuration optimisations and algorithm development, a significant increase in the rate at which data can be written to tape has been made possible. The integrated write rates have been found to increase from 16 MBs^{-1} to 90 MBs^{-1} per drive as a result of the developments presented. It has therefore been shown that the tape requirements in the RAL Tier-1 plan (§6.4.2) are achievable, and no extra hardware is required to make up for a lack of tape bandwidth.

Disk writes were found to be significantly limiting the read behaviour of certain RAID controllers. This resulted in very slow disk reads when sourcing files for tape migration, when the disk servers were under bi-directional load. By limiting the number of disk writes (algorithmically), and allowing reads, when scheduled, to source a reasonable amount of data (hardware configuration), the overall performance of the disk servers improved in the context of tape migrations.

The optimisations presented are now used by the other LHC experiments at RAL, and on all CASTOR instances in use at CERN.

7.6 Conclusions

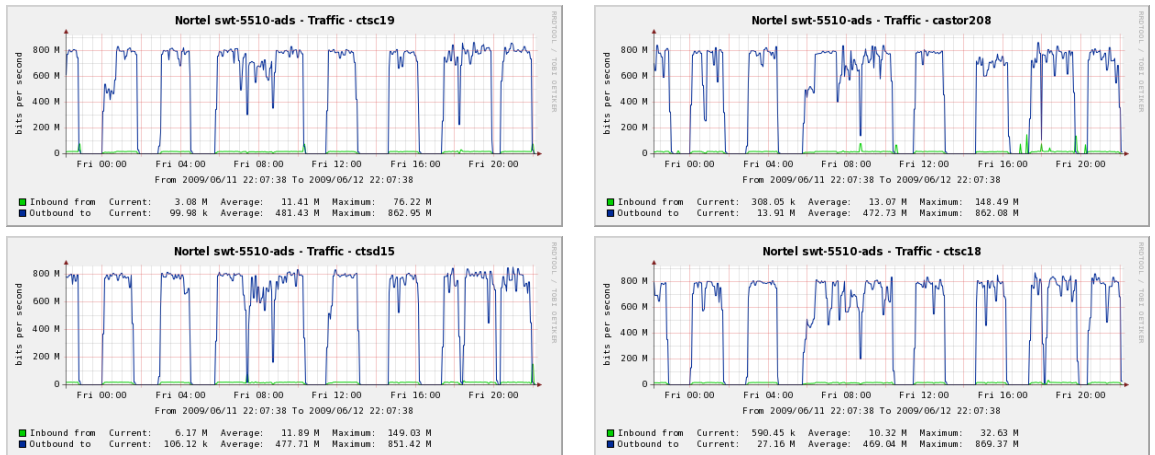


Figure 7.9: STEP09 tape server network activity with stream policies: Dedicated CMS Tape Servers. All data transferring at 100 MBs^{-1} .

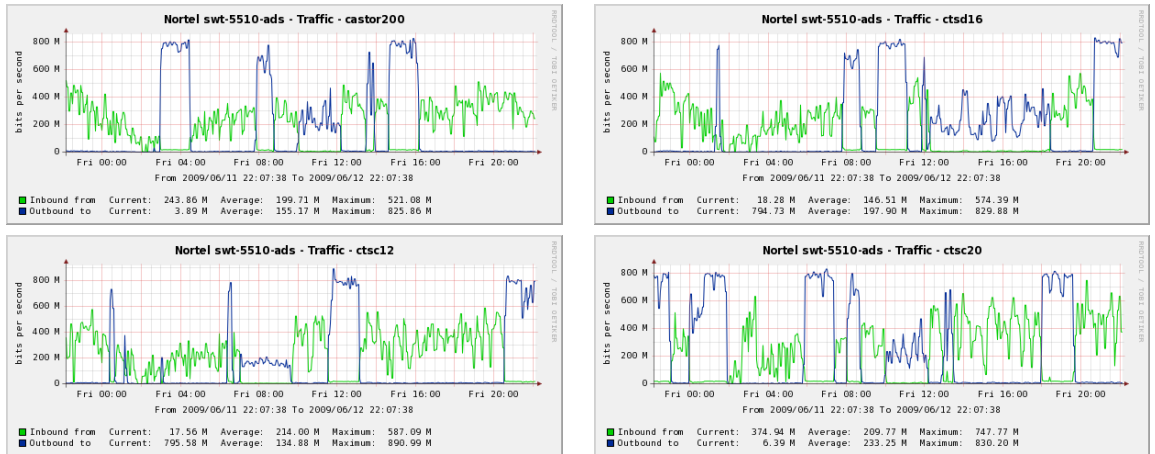


Figure 7.10: STEP09 tape server network activity with stream policies: Shared Tape Servers. Data from CMS Disk Servers transferring at 100 MBs^{-1} , data from other experiment Disk Servers substantially slower.

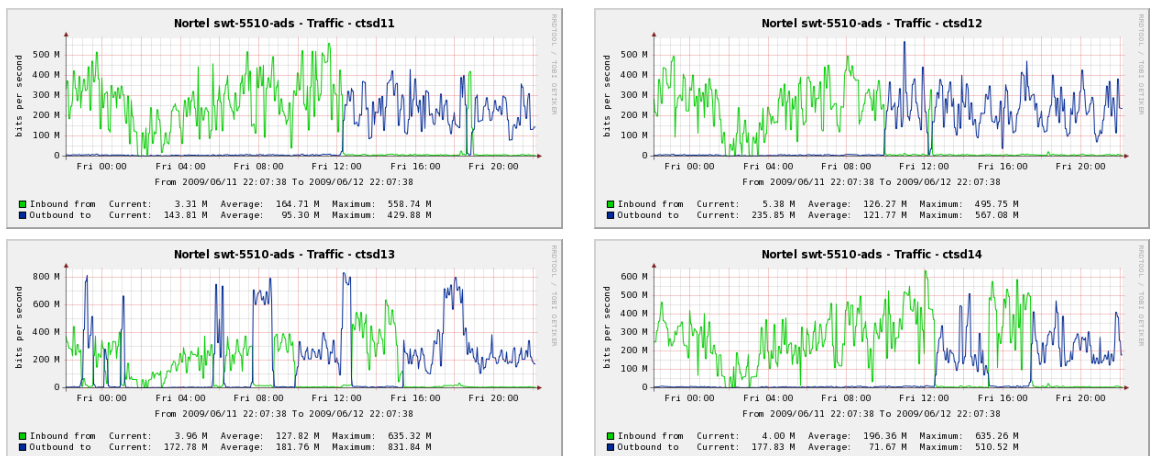


Figure 7.11: STEP09 tape server network activity with stream policies: Non-CMS Tape Servers. Majority of data transferring substantially slower than 100 MBs^{-1} .

8. High p_t Z^0 s as a probe for physics beyond the standard model

For potential discovery channels, it is desirable to perform a model-independent analysis to search for excesses to standard model backgrounds. A search is proposed looking for e^+e^- final states from the decays of standard model Z^0 s, which are themselves boosted due to being a decay product of a new heavy object.

To benchmark the channel, an excited quark model is used. It is found that 3σ evidence for new physics at a mass of 1 TeV could be found with 200 pb^{-1} of integrated luminosity at $\sqrt{s} = 10 \text{ TeV}$. The analysis is presented as if performed on 200 pb^{-1} of data¹.

8.1 Phenomenological background

Most known theories – candidates for the new physics beyond the Standard Model – predict new heavy particles participating in the electroweak interactions and, therefore, decaying into Z^0 bosons. Many of these candidate models for physics beyond the Standard Model therefore predict boosted Z^0 signatures. Among these are various SUSY models, technicolour and compositeness models [3, 74, 75].

8.1.1 SUSY and Technicolour

A possible source due to SUSY is where boosted Z^0 bosons are produced from the decays of charginos or neutralinos. When charginos or neutralinos originate from initial gluino decays, their production rate could be quite large, even at 7 TeV with the LHC at its early stage. The example of such cascade decay is shown in Fig. 8.1. Another example of models which can give rise to boosted Z^0 bosons are various models of Walking Technicolor, versions of which became especially popular recently [74]. Representative Feynman diagram for Techni-rho ρ_{TC}^+ decay leading to boosted Z^0 bosons is shown in Fig. 8.2.

¹At the time this analysis was performed, the LHC was expected to start up with a 10 TeV, $\int L dt =$




Figure 8.1: One of the representative Feynman diagrams for the cascade gluino decay leading to boosted Z^0 bosons




Figure 8.2: One of the representative Feynman diagrams for the Techni-rho ρ_{TC}^+ decay leading to boosted Z^0 bosons

8.1.2 Quark compositeness

In this study, a compositeness model with excited quarks decaying to quark and boosted Z^0 boson pairs [75] is used. For definiteness, the excited fermions are taken to be spin $\frac{1}{2}$, isospin $\frac{1}{2}$ partners. It is also assumed that an excited fermion has acquired a mass before electroweak symmetry breaking. The left and right-handed components are therefore considered in isodoublets. The first fermion generation therefore can be written as:

$$l_L \equiv \begin{pmatrix} \nu_e \\ e \end{pmatrix}_L, \quad l_R \equiv e_R; \quad l_L^* \equiv \begin{pmatrix} \nu_e^* \\ e^* \end{pmatrix}_L, \quad l_R^* \equiv \begin{pmatrix} \nu_e^* \\ e^* \end{pmatrix}_R$$

$$q_L \equiv \begin{pmatrix} u \\ d \end{pmatrix}_L, \quad q_R \equiv u_R, d_R; \quad q_L^* \equiv \begin{pmatrix} u^* \\ d^* \end{pmatrix}_L, \quad q_R^* \equiv \begin{pmatrix} u^* \\ d^* \end{pmatrix}_R$$

Constraints from $g - 2$ measurements lead to the requirement that excited fermions have chiral form interactions with the standard model fermions [76]. The couplings of excited fermions to gauge bosons are vector like,

$$\mathcal{L}_{\text{eff}}^1 = \bar{f}^* \gamma^\mu \left(f_s g_s \frac{\lambda^a}{2} G_\mu^a + f g \frac{\tau^a}{2} W_\mu^a + f' g' \frac{\mathbf{Y}}{2} B_\mu \right) f^*,$$

while transitions between standard model and excited fermion states are given by magnetic-moment type interactions [77],

200 pb $^{-1}$ pilot run, and as such the analysis is focussed around those conditions.

8.1 Phenomenological background

$$\mathcal{L}_{\text{eff}}^2 = \frac{1}{2\Lambda} \bar{f}_R^* \sigma^{\mu\nu} \left(f_s g_s \frac{\lambda^a}{2} G_{\mu\nu}^a + f g \frac{\tau^a}{2} W_{\mu\nu}^a + f' g' \frac{Y}{2} B_{\mu\nu} \right) f_L, \quad (8.1)$$

where $G_{\mu\nu}^a$, $W_{\mu\nu}^a$ and $B_{\mu\nu}$ are the $SU(3)$, $SU(2)$ and $U(1)$ field strength tensors, f_s , f and f' are parameters representing the underlying dynamics, and Λ is the compositeness scale. The couplings are taken to be of order unity, $f_s = f = f' = 1$, and a further simplifying assumption is made by equating the mass of the excited fermion and the compositeness scale $\Lambda = m_{q^*}$.

It is interesting to note that the vertices $\bar{f}^* g f^*$ are not suppressed by Λ , whereas $\bar{f}_R^* g f_L$ vertices are. This would prove significant by allowing the process $pp \rightarrow q^* q^*$ (Fig. 8.3(b)) to be an available discovery channel if $pp \rightarrow q^*$ (Fig. 8.3(a)) were heavily suppressed by a large compositeness scale. The production cross sections for these two scenarios as a function of mass are given in Figs. 8.4(a) and (b).

Figs. 8.5 and 8.6 show the Z^0 decay kinematics of the composite quark, for m_{q^*} of 1 and 3 TeV. The decay around the resonance is generally central ($|\eta| < 2$) and therefore the Z^0 decay products will be well within the detector calorimetric acceptance.

For the purposes of this study, q^* masses of 1.0, 1.25, 1.5, 1.75 and 2.0 TeV will be considered. These correspond, in the q^* rest frame, to Z^0 boosts, given in terms of the Lorentz γ factor, where $\gamma_Z = \sqrt{1 + p_Z^2/m_Z^2}$, of 5.5, 6.9, 8.2, 9.6 and 11.0.

8.1.3 Current experimental limits

Both CDF and D0 have investigated the production of high p_t Z^0 s [78, 79]. Fig. 8.7 illustrates the differential cross section for anomalous Z^0 production at the 95% CL from CDF, while Fig. 8.8 shows the very good agreement between data and theory for the D0 measurement. The filled points represent a theoretical prediction which controls the divergence at $p_t = 0$ and includes soft gluon emission [80], both of which are not included in the Fixed Order $\mathcal{O}(\alpha_s^2)$ calculation.




Figure 8.3: Feynman diagrams for the $pp \rightarrow q^* \rightarrow qe^+e^-$ (a) and $pp \rightarrow q^*q^* \rightarrow qe^+e^- + X$ (b) processes




Figure 8.4: Cross sections for the $pp \rightarrow q^* \rightarrow qe^+e^-$ (a) and $pp \rightarrow q^*q^* \rightarrow qe^+e^- + X$ (b) processes, $\Lambda = M_{q^*}$




Figure 8.5: Z^0 decay kinematic probabilities in the (η, p_t) plane for $pp \rightarrow q^* \rightarrow qZ^0$, $M_{q^*} = 1$ TeV (a) and 3 TeV (b). Total boxed area normalised to 1.

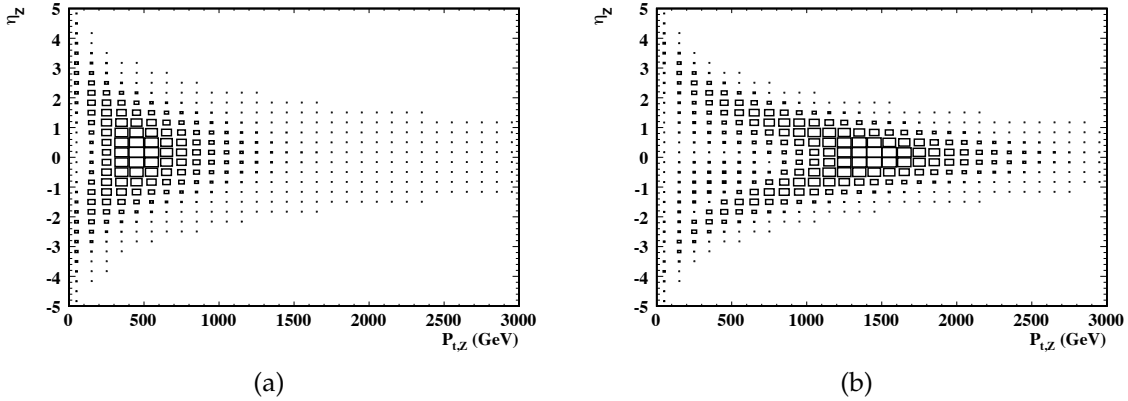


Figure 8.6: Z^0 decay kinematic probabilities in the (η, p_t) plane for $pp \rightarrow q^*q^* \rightarrow qZ^0 + X$, $M_{q^*} = 1$ TeV (a) and 3 TeV (b). Total boxed area normalised to 1.

8.2 Datasets considered




Figure 8.7: The Bayesian limit on the differential cross section for anomalous $Z^0 \rightarrow e^+e^-$ production at 95% CL from CDF [78]




Figure 8.8: $Z^0 \rightarrow e^+e^- p_t$ Data/Theory $- 1$ comparison from D0 [79, 80].

Current experimental limits from these analyses show agreement with the Standard Model for $Z^0 p_t < 350$ GeV.

8.2 Datasets considered

Any Standard Model process which results in either a real Z^0 in the final state, or combinations of real and fake electrons producing a Z^0 candidate constitute the backgrounds to the excited quark search. The backgrounds considered are, therefore, collective di-boson production ($W^\pm W^\pm$, $W^\pm Z^0$, $Z^0 Z^0$, $W^\pm \gamma$, $Z^0 \gamma$), $t\bar{t}$ +Jets, W^\pm +Jets, Z^0 +Jets, γ +Jets and QCD di-jets.

Background and cross-check datasets used were from official CMS production, and are detailed in Tables 8.1 and 8.2. The GEN-SIM-RECO data tier was used throughout, and the p_t cuts refer to the p_t of the leading generator-level hadronic parton. These datasets were produced in the 2.1.X series of CMSSW releases, and rereconstructed with 2.2.4.

Signal datasets were generated with CalCHEP 2.5.2 [81], hadronised with Pythia 6 [82]. The detector response was simulated and events reconstructed with CMSSW 2.2.4. The CTEQ6m [83] PDF set was used. The generated signal sets are detailed in Table 8.3. Note that the cross sections do not include $BR(Z^0 \rightarrow e^+e^-)$.

8.3 Triggering

Triggering for CMS follows a two stage process. Firstly, the Level-1 (L1) trigger is run on coarse-grained data from every bunch crossing. If an algorithm implemented in the L1 Global Trigger passes, the entire detector is read out and further triggering performed in

High p_t Z^0 s as a probe for physics beyond the standard model

Table 8.1: Background datasets

<i>Dataset name</i>	<i>Process</i>	<i>Cross section (pb)</i>
/Zee/Summer08_IDEAL.V9_v1	$Z \rightarrow e^+e^-$	864.3
/WJets-madgraph/Fall08_IDEAL.V9_v1	$W + \text{Jets}$	35550.0
/TTJets-madgraph/Fall08_IDEAL.V9_v2	$t\bar{t} + \text{Jets}$	375.0
/WW_2l/Summer08_IDEAL.V11_redigi_v1	$WW \rightarrow ll\nu\nu$	7.4
/WZ_incl/Summer08_IDEAL.V9_v2	WZ , inclusive decay	32.4
/ZZ_2l2n/Summer08_IDEAL.V11_redigi_v1	$ZZ \rightarrow ll\nu\nu$	0.42
/ZZ_4l/Summer08_IDEAL.V11_redigi_v1	$ZZ \rightarrow l^+l^-l^+l^-$	0.105
/Wgamma/Summer08_IDEAL.V11_redigi_v1	$W\gamma$, inclusive decay	36.6
/Zgamma/Summer08_IDEAL.V11_redigi_v1	$Z\gamma$, inclusive decay	11.0
/PhotonJetPt15/Summer08_IDEAL.V9_RECO-v4	$\gamma + \text{Jets}$, $p_t > 15$ GeV	288813.0
/PhotonJetPt30/Summer08_IDEAL.V9_v1	$\gamma + \text{Jets}$, $p_t > 30$ GeV	32203.8
/PhotonJetPt80/Summer08_IDEAL.V9_RECO-v4	$\gamma + \text{Jets}$, $p_t > 80$ GeV	1012.08
/PhotonJetPt170/Summer08_IDEAL.V9_RECO-v1	$\gamma + \text{Jets}$, $p_t > 170$ GeV	51.36
/PhotonJetPt300/Summer08_IDEAL.V9_v1	$\gamma + \text{Jets}$, $p_t > 300$ GeV	4.193
/PhotonJetPt470/Summer08_IDEAL.V9_RECO-v4	$\gamma + \text{Jets}$, $p_t > 470$ GeV	0.45125
/PhotonJetPt800/Summer08_IDEAL.V9_v1	$\gamma + \text{Jets}$, $p_t > 800$ GeV	0.02
/QCDpt15/Summer08_IDEAL.V9_v3	QCD Di-Jets, $p_t > 15$ GeV	1457159248
/QCDpt30/Summer08_IDEAL.V9_v4	QCD Di-Jets, $p_t > 30$ GeV	109057220.4
/QCDpt80/Summer08_IDEAL.V9_v2	QCD Di-Jets, $p_t > 80$ GeV	1934639.5
/QCDpt170/Summer08_IDEAL.V9_v3	QCD Di-Jets, $p_t > 170$ GeV	62562.9
/QCDpt300/Summer08_IDEAL.V9_v1	QCD Di-Jets, $p_t > 300$ GeV	3664.6
/QCDpt470/Summer08_IDEAL.V9_v1	QCD Di-Jets, $p_t > 470$ GeV	315.5
/QCDpt800/Summer08_IDEAL.V9_RECO-v5	QCD Di-Jets, $p_t > 800$ GeV	11.9

Table 8.2: Cross-check datasets

<i>Dataset name</i>	<i>Process</i>	<i>Cross section (pb)</i>
/ZJets-madgraph/Fall08_IDEAL.V9_reco-v2	$Z + \text{Jets}$, inclusive decay	3700.0
/Wenu/Summer08_IDEAL.V9_v1	$W \rightarrow e\nu$	n/a

Table 8.3: Signal datasets generated

<i>u^* mass (TeV)</i>	<i>Cross section (pb)</i>
1.0	10.3
1.25	3.4
1.5	1.3
1.75	0.53
2.0	0.24

8.4 Electron reconstruction

the High Level Trigger (HLT).

The HLT consists of the full `CMSSW` software framework, running a number of trigger paths. Despite which L1 triggers pass, all HLT paths are run on the event. To reduce backgrounds, all physics object HLT paths require a L1 seed match. For example, in the HLT photon triggers, which trigger on ECAL SuperClusters, each reconstructed SuperCluster is only considered in the HLT path if the SuperCluster location matches (in ΔR) one of the L1 e/γ candidates.

8.3.1 Level-1 seed

Following §5.3, a non-isolated L1 trigger is required. In addition to the issues for isolated high-energy electrons (namely hadronic punch-through and lateral shower-shape), for a boosted Z^0 analysis the closeness of Z^0 decay products could result in the failure of L1 isolation criteria. Therefore, an HLT path seeded by a `L1_SingleEGX` (see §5.3 for naming conventions) trigger is required, where X is the lowest E_t un-prescaled trigger threshold available in the running menu.

8.3.2 HLT path

Following L1 triggering, a suitable HLT path is required. The collaboration currently has prepared two menus for early luminosity running at 8×10^{29} and $10^{31} \text{ cm}^{-2}\text{s}^{-1}$. As the modified electron reconstruction (described in §8.4.2) will not be run in the HLT, it is desirable for the HLT selection to be as loose as possible, and as such the recommended HLT paths are `HLT_Photon15.L1R` and `HLT_Photon25.L1R` for the two menus respectively. These are non-isolated photon triggers which require a single SuperCluster with E_t greater than the given threshold, matching a L1 EGamma candidate with $\Delta R < 0.2$. By using a photon trigger, problems due to incorrect track matching cuts will not be present. §8.6.7 validates this choice of trigger, with near-unity efficiency for events containing two selected electrons.

8.4 Electron reconstruction

8.4.1 Electron reconstruction algorithm

The electron reconstruction algorithm in CMS has three distinct stages; seeding, track finding and preselection. The algorithms are described in detail in [84].

Electron candidates start with a seed SuperCluster in the ECAL. A SuperCluster is a collection of locally clustered energy maxima (BasicClusters), merged to account for bremsstrahlung radiation in the ϕ direction. The energy-weighted average position of all

calorimeter hits is, in the limit of perfect calorimeter and clustering efficiency, identical to that of the original electron had no radiation occurred. This position is used to back-propagate for both charge hypotheses to the inner two pixel layers. If two matching (in ϕ and z) pixel hits are found, they are associated to the SuperCluster and marked as a seed.

Following seeding, a Gaussian Sum Filter (GSF) [85] method is used to reconstruct the full track parameters. This algorithm uses a Bethe-Heitler model of electron energy losses at each tracker layer, and can therefore track electrons that have a large bremsstrahlung energy fraction. The algorithm iterates over tracker layers in the outwards direction until the last tracker layer. If no compatible hits are found in two subsequent layers the track is abandoned. If there are many candidate hits at a given layer, they are all extrapolated consecutively. A χ^2 statistic is computed for each track, which measures how compatible each set of candidate track hits is. At the end of the algorithm, the best two candidates are kept, each of which must have at least 5 hits. The final track parameters are computed with a further backward fit of the trajectory.

All electron candidates are then tested against a loose set of cuts, shown in Table. 8.4, and only those passing are finally stored as true electron candidates. The efficiency of the preselection cuts reaches 99% in the central region, dropping to 92% at high $|\eta|$ [84]. §8.5.1 describes the cuts in detail.

Table 8.4: GSF electron preselection cuts

<i>Cut</i>	<i>EB</i>	<i>EE</i>
E_t	$> 4 \text{ GeV}$	$> 4 \text{ GeV}$
$ \Delta\eta_{\text{in}} $	< 0.02	< 0.02
$ \Delta\phi_{\text{in}} $	$< 0.15 \text{ rad}$	$< 0.15 \text{ rad}$
H/E (Depth 1)	< 0.1	< 0.1
H/E (Depth 2)	< 0.1	< 0.1

The reconstruction algorithm applies a correction for the momentum, based on the value of $E_{\text{SuperCluster}}/p_{\text{Track}}$. If E/p is more than 2.5σ from 1, where σ is the error on E/p , the magnitude of the momentum is taken as the energy as measured in the calorimeter. For high energy electrons, where the very straight track is likely to result in a badly measured momentum, the result is such that the momentum and energy assignment is dominated by the calorimeter information. The track information is still important such that an electron can be reconstructed in the first place, and its position determined to high accuracy.

8.4.2 Close electron reconstruction

The standard reconstruction algorithm used at CMS is extremely efficient for single electrons (99% in the central region). However, it begins to break down when two electrons are in close proximity to each other. Fig. 8.9 shows the worst case opening angle from two electrons in the $X \rightarrow Z^0 p$, $Z^0 \rightarrow e^+ e^-$ system, where the decay electrons are produced perpendicular to the Z^0 boost direction. For increasing M_X , the closeness soon becomes a problem for the standard reconstruction algorithms. Fig. 8.10 shows the efficiency for a Z^0 to be reconstructed (with no electron ID cuts) from two reconstructed electrons from a custom Monte Carlo which fires di-electron pairs simulating the kinematics of isotropically decaying Z^0 s with arbitrary boost. Above γ_Z of about 6 (corresponding to a $Z^0 p_t$ of ≈ 540 GeV), the reconstruction efficiency can be seen to drop significantly.

The efficiency drop is due to the brehmstrahlung correction algorithm performed as part of the supercluster step detailed in §8.4.1. When the two electrons are aligned in ϕ and geometrically close, the clustering algorithm treats the two basic clusters formed as belonging to a primary electron and a radiated photon. The result of this is one supercluster, each consisting of two well-formed basic clusters, with two basic tracks pointing to it. As the GSF electron algorithm is SuperCluster seeded, only one electron is reconstructed.

Fig. 8.14 shows an example of this for a Z^0 with $\gamma_Z = 20$. The red tower is the SuperCluster seed BasicCluster from one electron, and the yellow tower, aligned in ϕ , is the BasicCluster from the second electron which has been merged as being due to brems. The red cross shows the outer track extrapolation as reported by the GSF track associated with the single reconstructed electron. The event also contains two BasicTracks, each aligned perfectly with the two BasicClusters.

The solution to this problem is twofold. Firstly, a simple clustering algorithm is used to avoid a basic cluster merging the two electrons. A suitable clustering algorithm, `FixedMatrix5x5`, is already available [86]. This limits the cluster size to a 5×5 matrix of crystals around a maximal-energy crystal. This algorithm is already used by default in the endcap, but must be enabled in the barrel to allow the proposed method to work. Secondly, no brehmstrahlung correction is applied at the superclustering step. That is to say, all BasicClusters are promoted to SuperClusters, with a minimal E_T cut of 15 GeV.

The result of this modified algorithm is shown in Fig. 8.11, and a comparison given with the old algorithm in slices of γ_Z in Fig. 8.15. The efficiency is now constant up to at least $\gamma_Z = 20$, corresponding to a Z^0 emitted from the decay of a 3.6 TeV object.

8.4.3 Quantification of energy loss

As the bremsstrahlung recovery algorithms are present to ensure minimal energy loss during electron reconstruction, it is required to check how the removal of this affects




Figure 8.9: Worst case $Z^0 \rightarrow e^+e^-$ decay angle from $X \rightarrow Z^0 p$




Figure 8.10: Boosted Z^0 reconstruction efficiency with old algorithm




Figure 8.11: Boosted Z^0 reconstruction efficiency with the modified algorithm




Figure 8.12: Boosted Z^0 reconstruction efficiency with old algorithm for $-2 < \eta_Z < 2$




Figure 8.13: Boosted Z^0 reconstruction efficiency with the modified algorithm for $-2 < \eta_Z < 2$




Figure 8.14: Example of close electron merging

the energy measurements of high p_t electrons. Fig. 8.16 shows the effect of the modified algorithm to the reconstruction of Z^0 invariant mass, using the standard model $Z^0 \rightarrow e^+e^-$ sample. The low sideband shape is typical of missing energy due to the lack of brehmstrahlung correction.

To check that this problem is limited to low p_t Z^0 s, Fig. 8.19 shows the modified reconstruction performance where various cuts on the reconstructed Z^0 p_t are applied. Fig. 8.20 compares the $p_t > 120$ GeV cut with the standard reconstruction to show the sideband behaviour is restored. A Gaussian convoluted with a Breit-Wigner distribution fitted to the two histograms (Figs. 8.17 and 8.18) gives the standard reconstructed invariant mass of 90.42 ± 0.14 GeV, whereas with the modified algorithm the mass is 84.45 ± 0.14 GeV. While the fits clearly do not describe the data accurately, they describe the peak region sufficiently well to accurately determine the peak position.

To confirm that the low sideband behaviour is due to missing energy, an energy correction was determined from a sample of 1M single electrons and positrons as

$$F(E, \eta) = \frac{E_{MC} - E_{Reco}}{E_{MC}}. \quad (8.2)$$

Each reconstructed electron 4-vector is then scaled by

$$\frac{E_{Reco}}{1 - F(E, \eta)}.$$

The resulting corrections and mass distribution compared to the standard reconstruction are shown in Figs. 8.21 and 8.22. The energy loss is largest at low energy and in the tracker barrel to endcap transition region, where the material effects are large. The mass plot confirms that the sideband behaviour is restored. The reconstructed mass in this case is 92.4 ± 3.0 GeV.

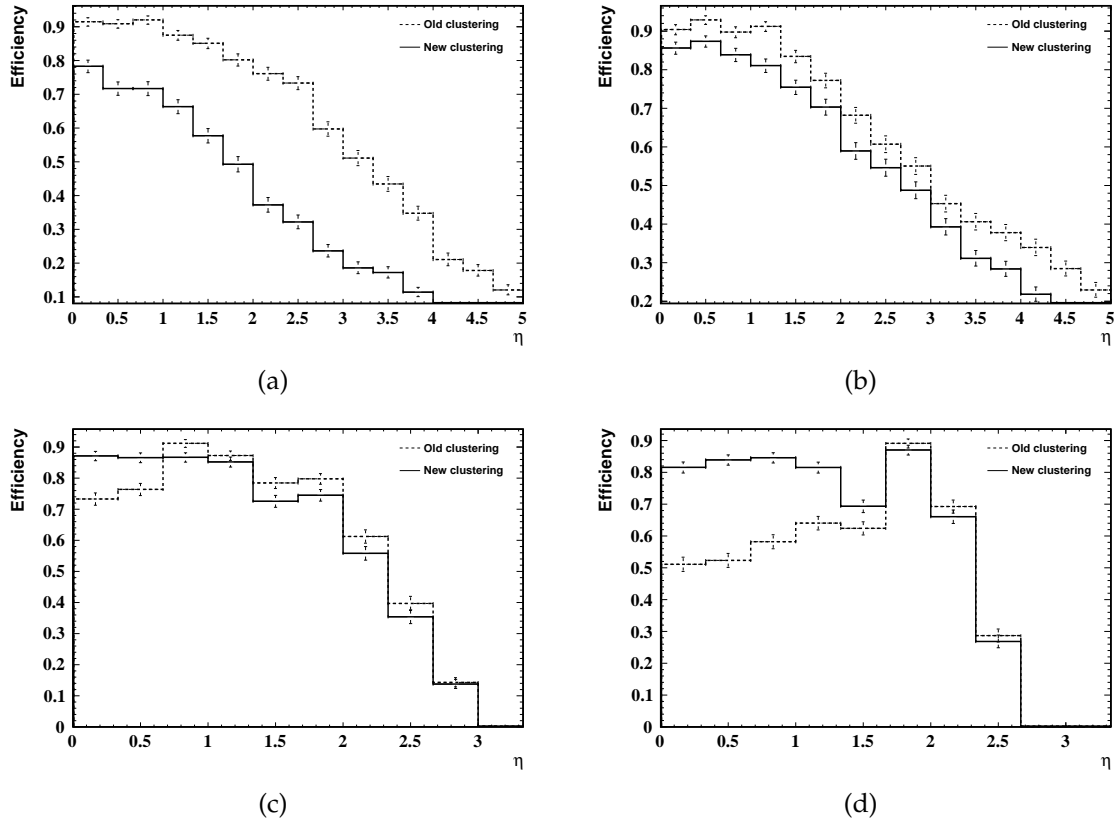


Figure 8.15: Comparison of old and modified clustering algorithms for $\gamma_Z = 1$ (a), 5 (b), 10 (c) and 20 (d)




Figure 8.16: Reconstructed Z^0 invariant mass with modified algorithm

8.5 Electron identification cuts




Figure 8.17: Breit-Wigner \otimes Gaussian fit to Z^0 invariant mass reconstructed with the old algorithm




Figure 8.18: Breit-Wigner \otimes Gaussian fit to Z^0 invariant mass reconstructed with the new algorithm




Figure 8.19: Modified reconstruction performance for Z^0 mass with di-electron p_t cut




Figure 8.20: Comparison of high p_t mass performance with standard reconstruction

As the energy loss is large in the region where efficiencies are measured (around 50 GeV, near the dominant energy for electrons from Z^0 decays), but low at higher energy, no energy correction will be applied in the analysis. The effect of energy loss on the analysis result will be taken into account by the systematic uncertainty on electron reconstruction and identification efficiencies. The reconstructed Z^0 lineshapes for highly boosted Z^0 's from u^* decays compared to those from the standard model process are shown in Fig. 8.23.

8.5 Electron identification cuts

8.5.1 Current HEEP electron selection

The High Energy Electron Pairs (HEEP) group maintains a recommended set of cuts to use in di-electron analyses. For reasons of reproducibility and consistency it is desirable to use a similar set of cuts. The cuts are listed in Table 8.5, and detailed below.

E_t is defined as $E_c \sin \theta_{\text{trk}}$, where E_c is the corrected electron energy, and θ_{trk} the polar




Figure 8.21: Single-electron energy loss with modified electron reconstruction. z -axis is the fractional energy loss as defined in Eq. 8.2.




Figure 8.22: Modified reconstruction performance for Z^0 mass with energy corrections




Figure 8.23: Reconstructed Z^0 invariant mass for $pp \rightarrow Z^0$ and $pp \rightarrow u^* \rightarrow Z^0 u$

angle of the electron track measured at the inner layer of the tracker extrapolated to the interaction vertex. Note that E_c is an inverse error squared-weighted sum of the electron supercluster energy and track momentum. For high energy electrons it is dominated by the supercluster energy measurement.

$|\eta_{\text{SC}}|$ is the pseudorapidity of the electron's supercluster, with respect to $(x, y, z) = (0, 0, 0)$. As such it is suitable for fiducial cuts but not for use in physics results.

$|\Delta\eta_{\text{in}}|$ is the difference in η between the track position at the inner tracker layer, extrapolated back to the interaction vertex and then out to the calorimeter surface, and the η of the electron supercluster.

$|\Delta\phi_{\text{in}}|$ is the difference in ϕ between the track position at the inner tracker layer, extrapolated back to the interaction vertex and then out to the calorimeter surface, and the ϕ of the electron supercluster.

8.5 Electron identification cuts

Table 8.5: HEEP selection cuts v2.0

<i>Variable</i>	<i>Barrel</i>	<i>Endcap</i>
E_t	$> 25 \text{ GeV}$	$> 25 \text{ GeV}$
$ \eta_{\text{SC}} $	< 1.422	$1.560 < \eta_{\text{SC}} < 2.5$
$ \Delta\eta_{\text{in}} $	< 0.005	< 0.007
$ \Delta\phi_{\text{in}} $	$< 0.09 \text{ rad}$	$< 0.09 \text{ rad}$
H/E	< 0.05	< 0.05
$\sigma_{i\eta i\eta}$	n/a	< 0.0275
$E^{2\times 5}/E^{5\times 5}$	$> 0.94 \text{ OR } E^{1\times 5}/E^{5\times 5} > 0.83$	n/a
EM + Had Depth 1	$< 3 + 0.002E_t \text{ GeV}$	$< 5.5 \text{ GeV}$ for $E_t < 50 \text{ GeV}$ else
Isolation		$< 5.5 + 0.05(E_t - 50) \text{ GeV}$
Had Depth 2 Isolation	n/a	$< 0.5 \text{ GeV}$
Track p_t Isolation	$< 7.5 \text{ GeV}$	$< 15 \text{ GeV}$

H/E is the ratio of hadronic energy in the HCAL RecHit closest to the electron's position in the calorimeter, to the energy of the electron super cluster.

$\sigma_{i\eta i\eta}$ measures the spread in η in units of crystals of the electron energy in the 5×5 matrix of crystals centred on the seed crystal. It is insensitive to bremsstrahlung radiation, which is expected in ϕ only.

Electromagnetic Isolation is defined as the transverse energy in the ECAL with $E > 0.08 \text{ GeV}$ ($E > 0.03 \text{ GeV}$ in the endcap), where E is the energy in a single crystal, in a cone of radius $\Delta R < 0.3$, excluding the inner region of $\Delta R < 0.045$ (0.070 in the endcap), and an η strip of ± 0.02 . The variable is used in combination with the Hadronic Depth 1, defined below.

Hadronic Depth 1 Isolation is defined as the HCAL transverse energy in the first layer, deposited in a cone of radius $\Delta R < 0.3$ centred on the electron position in the ECAL. An inner exclusion cone of $\Delta R < 0.15$ removes energy due to EM punch-through.

Hadronic Depth 2 Isolation is defined as the HCAL transverse energy in the second layer, deposited in a cone of radius $\Delta R < 0.3$ centred on the electron position in the ECAL. An inner exclusion cone of $\Delta R < 0.15$ also applies.

Track p_t Isolation is defined as the scalar sum p_t of GeneralTracks² in a cone of $0.015 < \Delta R < 0.2$, with $p_t > 1.0$, a transverse impact parameter (z_0) of $\pm 0.2 \text{ cm}$ from the electron GSF Tracks, and a maximum transverse distance of 0.1 cm from the nominal beamspot position in the x, y plane.

The selection cuts are designed to be maximally efficient for high-energy electrons, but allow enough low-energy electrons to pass to be used in efficiency measurements.

²These are tracker-only tracks, reconstructed using a Kalman Filter algorithm [52].

8.5.2 HEEP selection performance

As a test of the standard HEEP selection performance, the event selection efficiency, defined as

$$\epsilon_E = \frac{\text{Events with } \geq 2 \text{ fiducial electrons passing HEEP cuts}}{\text{Total number of events with } \geq 2 \text{ fiducial electrons}},$$

was calculated for each signal dataset. The results, shown in Table 8.6, clearly show a low efficiency and drop-off with increasing invariant mass.

The efficiency of each cut was measured to determine where the problem lay. Each reconstructed electron was matched to a Monte Carlo truth electron, and the efficiency of each individual cut calculated. These efficiencies are shown in Tables 8.7 and 8.8 for electrons in the barrel and endcap respectively³, from a u^* decay for five u^* masses. The *EM + Had D1* and *Track p_t Isolation* cuts are clearly responsible for the efficiency drop-off due to the closeness of the electrons mimicking a single non-isolated electron.

The event selection was re-measured with these two cuts removed. The results, in Table 8.9, show that the efficiency is now much greater and stable. The increase of the efficiency with increasing m_{u^*} is due to the E_t cut acceptance increasing in the presence of higher p_t leptons from Z^0 s with higher boosts.

8.6 Efficiency measurements

8.6.1 The Tag and Probe method

It is important to be able to measure the efficiency for each analysis step from data such that it can be checked that the detector and reconstruction software are performing well. The efficiency can thus only be measured in a kinematically accessible region, and the evolution of the measured efficiency to high energy will have to be determined from simulation.

The Tag and Probe method allows efficiencies to be measured from data. By constructing an object pair from one object which passes tight electron selection (the tag) and one loose object (the probe), where the pair invariant mass is close to that of the Z^0 ($85 < M < 95$ GeV), a high-purity sample of probe objects from $Z^0 \rightarrow e^+e^-$ decays can be constructed with which to measure the efficiency of an analysis step. This method has previously been used by both CDF and D0 at the Tevatron [87], and has been studied in detail for use in CMS [88].

The method is unbiased due to two requirements. Firstly, the tag object must pass the trigger, removing trigger bias from the probe object selection, and the probe object

³Errors are not shown as the measurements were performed on an extremely large dataset and are for illustration only. Errors can be taken to be at the 0.0005 level.

8.6 Efficiency measurements

Table 8.6: HEEP cuts event selection efficiency

m_{a^*} (TeV)	Efficiency
1	0.676 ± 0.004
1.25	0.566 ± 0.004
1.5	0.434 ± 0.004
1.75	0.332 ± 0.003
2.0	0.264 ± 0.002

Table 8.7: Barrel HEEP selection performance

Cut	1 TeV	1.25 TeV	1.5 TeV	1.75 TeV	2 TeV
E_t	0.96	0.97	0.98	0.98	0.98
$ \eta_{\text{SC}} $	0.98	0.99	0.99	0.99	0.99
$ \Delta\eta_{\text{in}} $	0.99	0.99	0.99	0.99	0.99
$ \Delta\phi_{\text{in}} $	1.00	1.00	1.00	0.99	1.00
H/E	0.99	0.99	0.99	0.99	0.99
$E^{2 \times 5} / E^{5 \times 5}$	0.96	0.96	0.96	0.96	0.97
EM + Had D1 Isolation	0.90	0.75	0.57	0.45	0.36
Track p_t Isolation	0.96	0.89	0.71	0.56	0.46

Table 8.8: Endcap HEEP selection performance

Cut	1 TeV	1.25 TeV	1.5 TeV	1.75 TeV	2 TeV
E_t	0.96	0.96	0.96	0.97	0.97
$ \eta_{\text{SC}} $	0.93	0.92	0.92	0.92	0.92
$ \Delta\eta_{\text{in}} $	0.98	0.99	0.99	0.99	0.99
$ \Delta\phi_{\text{in}} $	1.00	1.00	1.00	1.00	1.00
H/E	0.99	0.99	0.99	0.99	0.99
$\sigma_{\text{in}\eta}$	1.00	0.99	0.99	0.99	0.99
EM + Had D1 Isolation	0.96	0.93	0.87	0.83	0.79
Had D2 Isolation	0.97	0.97	0.97	0.97	0.96
Track p_t Isolation	0.99	0.98	0.94	0.90	0.88

Table 8.9: Modified HEEP cuts event selection efficiency

m_{a^*} (TeV)	Efficiency
1	0.797 ± 0.005
1.25	0.823 ± 0.005
1.5	0.842 ± 0.006
1.75	0.853 ± 0.006
2.0	0.864 ± 0.005

definition must be unbiased with respect to the step under test (for example, using a tracker-seeded track to measure calorimeter clustering efficiency).

8.6.2 Background estimation

Despite the use of the Tag and Probe method to produce a high-purity sample with which to measure an efficiency, there will be significant backgrounds to the measurements, especially when the probe object is of low purity. For early analyses, a robust method to determine and correct for this background contamination is required.

A sideband subtraction method is proposed. The invariant mass of each tag and probe pair is histogrammed, for both the denominator and numerator selections. The number of events in the lower ($40 < M < 60$ GeV) and upper ($120 < M < 140$ GeV) sidebands are counted, as illustrated in Fig. 8.24. A linear extrapolation between the two is performed to extract the expected background contamination in the signal region, which is subsequently subtracted from the numerator and denominator in the efficiency calculation. For binned efficiencies, individual numerator and denominator sideband histograms for each bin are produced to allow bin-by-bin subtraction of backgrounds. The background shapes are not entirely linear, but the systematic effect is small and easily quantifiable in each case.




Figure 8.24: Example of sideband subtraction technique

This technique is robust and easy to understand, making it ideal for use in initial conditions. The resulting systematic bias is also easy to understand (and will be discussed for each measurement below). In general, and most importantly for those measurements performed with low-purity probes, the background will be over-estimated by this technique. Additionally, the efficiency measurement will be dominated by Z^0 s with low p_t . This is a limitation of the technique in the limit of low integrated luminosity. As more

8.6 Efficiency measurements

data becomes available, a p_t cut on the tag-probe pair invariant mass could allow the efficiency measurements to be made with boosted Z^0 s.

8.6.3 Efficiency factorisation scheme

Given a number of efficiency measurements for an analysis, the total efficiency is given by

$$\epsilon_T = \prod_{i=1}^n \epsilon_i.$$

This combination assumes that all the efficiencies are uncorrelated, which in general may not be true. This requirement can be enforced by measuring efficiencies in a given sequence, with each efficiency measured with respect to the previous. The efficiencies of interest for this analysis are those associated with electron reconstruction, electron identification and trigger efficiency. They are combined to give an event yield efficiency (integrated over all electron p_t) as

$$\epsilon_T = \epsilon_{\text{Offline}}^2 \times \epsilon_{\text{Trigger}|\text{Offline}},$$

where

$$\begin{aligned} \epsilon_{\text{Offline}} = & \epsilon_{\text{Clustering}} \times \epsilon_{\text{Tracking}|\text{Clustering}} \times \epsilon_{\text{Reconstruction}|\text{(Tracking} \cap \text{Clustering)}} \\ & \times \epsilon_{\text{ID}|\text{(Tracking} \cap \text{Clustering} \cap \text{Reconstruction)}} \end{aligned}$$

and $\epsilon_{\text{Trigger}}$ is defined as the probability that at least one of the electrons didn't fail the trigger,

$$\epsilon_{\text{Trigger}} = 1 - (1 - \epsilon_{\text{Online}})^2.$$

Following this factorisation scheme, ϵ_{Online} is the efficiency of the L1 + HLT to pass given that an electron in the event passes selection cuts. This factors out the geometric acceptance from the trigger efficiency measurement. Additionally, the tracking and reconstruction steps can be considered together to give the final event efficiency defined as

$$\begin{aligned} \epsilon_T = & \left(\epsilon_{\text{Clustering}} \times \epsilon_{\text{GSF Reconstruction}|\text{Clustering}} \times \epsilon_{\text{ID}|\text{(GSF Reconstruction} \cap \text{Clustering)}} \right)^2 \\ & \times \left[1 - (1 - \epsilon_{\text{Online}})^2 \right]. \quad (8.3) \end{aligned}$$

It should be noted that the efficiencies, as measured in the following sections, are not used as input to the analysis search reach or, in case of data, for the statistical interpretation of an initial discovery. However, their measurement is vital to ensure that the detector and software behaviour is understood and well behaved.

8.6.4 Calorimeter clustering efficiency

The Tag and Probe criteria used for the calorimeter clustering efficiency are detailed in Table 8.10. The globalTrack collection is not calorimeter seeded, and as such the probe is unbiased with respect to the clustering efficiency measurement. The invariant mass plots of Tag + All Probes (all possible pairs of objects passing the tag and probe selection criteria) and Tag + Passing Probes (all possible pairs of objects passing the tag and probe selection criteria, where the probe object must also pass the object match required by the efficiency measurement) used to compute the background estimation are shown in Figs. 8.25 and 8.26.

To remove turn on effects, a cut on the probe p_t of 20 GeV is applied for the efficiency measurement with respect to $|\eta|$, and for the integrated efficiencies. The resulting efficiencies with respect to probe p_t and η are shown in Fig. 8.27, with the integrated efficiencies shown in Tables 8.11 and 8.12. The effect of the background estimation is to underestimate the true measured efficiency by 0.8% in both the barrel and endcap.

Table 8.10: Tag and Probe criteria for clustering efficiency measurements

<i>Tag Criteria</i>	<i>Probe Criteria</i>
GSF Electron	GeneralTrack
$ \eta < 1.44$ or $1.56 < \eta < 2.5$	$ \eta < 1.44$ or $1.56 < \eta < 2.5$
$E_t > 15$ GeV	$p_t > 10$ GeV
Passes modified HEEP cuts	Significance > 4
	ValidHits > 4
	$\Delta Z_{vtz}(T, P) < 0.01$ cm
	$\Delta R(\text{Probe}, \text{SuperCluster}) < 0.1$




Figure 8.25: Tag + All Probes invariant mass spectrum for $\epsilon_{\text{Clustering}}$




Figure 8.26: Tag + Passing Probes invariant mass spectrum for $\epsilon_{\text{Clustering}}$

8.6 Efficiency measurements

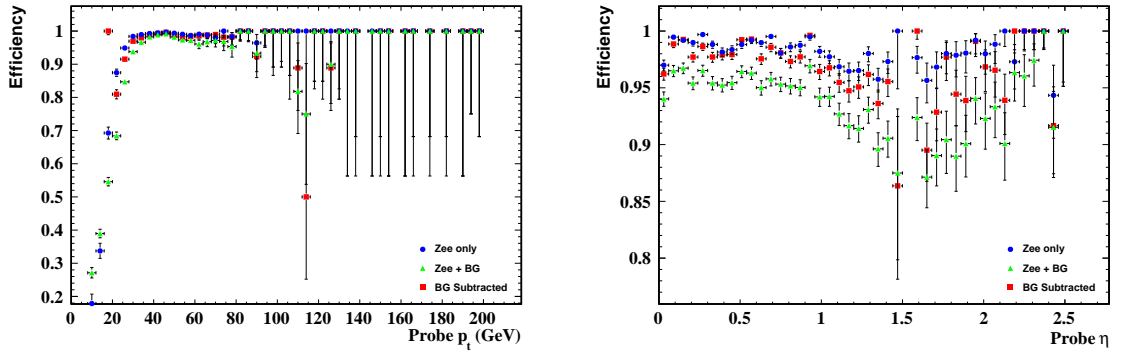


Figure 8.27: Tag and probe measurement of $\epsilon_{\text{Clustering}}$

Table 8.11: Integrated clustering efficiency (barrel)

<i>Dataset</i>	<i>Pass</i>	<i>Total</i>	<i>Efficiency</i>
$Z^0 \rightarrow ee$	22307	22641	$0.9852^{+0.0008}_{-0.0008}$
BG	953	1808	
$Z^0 \rightarrow ee + \text{BG}$	23260	24449	$0.951^{+0.001}_{-0.001}$
BG Estimate	1037.5	1712	
BG Subtracted	22222.5	22737	$0.977^{+0.001}_{-0.001}$

Table 8.12: Integrated clustering efficiency (endcap)

<i>Dataset</i>	<i>Pass</i>	<i>Total</i>	<i>Efficiency</i>
$Z^0 \rightarrow ee$	1446	1477	$0.979^{+0.004}_{-0.004}$
BG	140	245	
$Z^0 \rightarrow ee + \text{BG}$	1586	1722	$0.921^{+0.006}_{-0.007}$
BG Estimate	220.25	315.25	
BG Subtracted	1365.75	1406.75	$0.971^{+0.004}_{-0.005}$

8.6.5 GSF electron reconstruction efficiency

Following the efficiency factorisation scheme, the GSF electron reconstruction efficiency is defined as the probability to reconstruct an electron given that a SuperCluster due to a real electron was reconstructed. The Tag and Probe criteria used for the measurement are detailed in Table 8.13. The invariant mass plots of Tag + All Probes and Tag + Passing Probes used to compute the background estimation are shown in Figs. 8.28 and 8.29.

The resulting efficiency with respect to probe E_t and η are shown in Fig. 8.30, with the integrated efficiencies shown in Tables 8.14 and 8.15. The effect of the background estimation is to underestimate the true measured efficiency by 0.6% in the barrel and 1% in the endcap.

Table 8.13: Tag and Probe criteria for GSF electron reconstruction efficiency measurements

<i>Tag Criteria</i>	<i>Probe Criteria</i>
GSF Electron	superCluster
$ \eta < 1.44$ or $1.56 < \eta < 2.5$	$ \eta < 1.44$ or $1.56 < \eta < 2.5$
$E_t > 15$ GeV	$E_t > 20$ GeV
Passes modified HEEP cuts	$\Delta R(\text{Probe, GSF Ele}) < 0.1$




Figure 8.28: Tag + All Probes invariant mass spectrum for $\epsilon_{\text{GSF Reco}}$




Figure 8.29: Tag + Passing Probes invariant mass spectrum for $\epsilon_{\text{GSF Reco}}$

8.6 Efficiency measurements

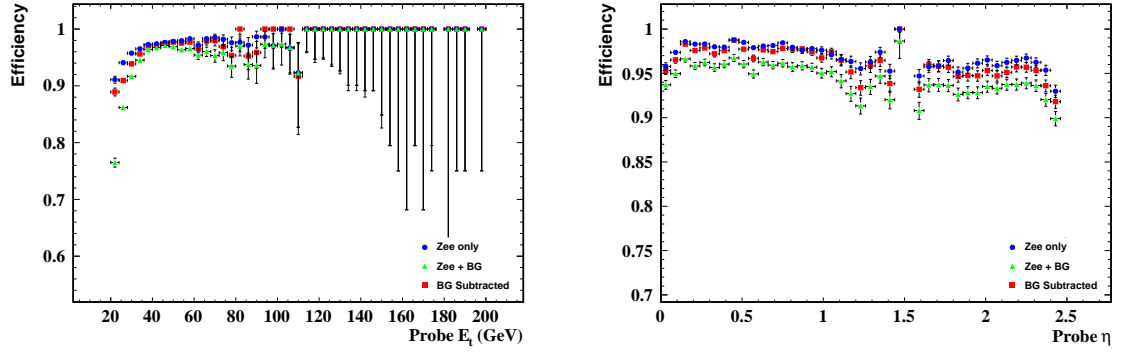


Figure 8.30: Tag and probe measurement of $\epsilon_{\text{GSF Reco}}$

Table 8.14: Integrated GSF electron reconstruction efficiency (barrel)

<i>Dataset</i>	<i>Pass</i>	<i>Total</i>	<i>Efficiency</i>
$Z^0 \rightarrow ee$	41805	42770	$0.9774^{+0.0007}_{-0.0007}$
BG	655	1731	
$Z^0 \rightarrow ee + \text{BG}$	42460	44501	$0.954^{+0.001}_{-0.001}$
BG Estimate	904.5	1725.75	
BG Subtracted	41555.5	42775.2	$0.9715^{+0.0008}_{-0.0008}$

Table 8.15: Integrated GSF electron reconstruction efficiency (endcap)

<i>Dataset</i>	<i>Pass</i>	<i>Total</i>	<i>Efficiency</i>
$Z^0 \rightarrow ee$	19248	20271	$0.950^{+0.002}_{-0.002}$
BG	224	849	
$Z^0 \rightarrow ee + \text{BG}$	19472	21120	$0.922^{+0.002}_{-0.002}$
BG Estimate	316.5	738	
BG Subtracted	19155.5	20382	$0.940^{+0.002}_{-0.002}$

8.6.6 Electron identification efficiency

Following the efficiency factorisation scheme, the electron identification efficiency is defined as the probability of an electron to pass offline selection cuts given that a GSF electron was reconstructed from a real electron. The Tag and Probe criteria used for the measurement are detailed in Table 8.16. The invariant mass plots of Tag + All Probes and Tag + Passing Probes used to compute the background estimation are shown in Figs. 8.31 and 8.32.

The resulting efficiency with respect to probe E_t and η are shown in Fig. 8.33, with the integrated efficiencies shown in Tables 8.17 and 8.18. The effect of the background estimation is to overestimate the true measured efficiency by 1.3% in the barrel and 0.7% in the endcap.

Table 8.16: Tag and Probe criteria for electron ID efficiency measurements

Tag Criteria	Probe Criteria
GSF Electron	GSF Electron
$ \eta < 1.44$ or $1.56 < \eta < 2.5$	$ \eta < 1.44$ or $1.56 < \eta < 2.5$
$E_t > 15$ GeV	$E_t > 20$ GeV
Passes modified HEEP cuts	Passes modified HEEP cuts




Figure 8.31: Tag + All Probes invariant mass spectrum for $\epsilon_{\text{Ele ID}}$




Figure 8.32: Tag + Passing Probes invariant mass spectrum for $\epsilon_{\text{Ele ID}}$

8.6 Efficiency measurements

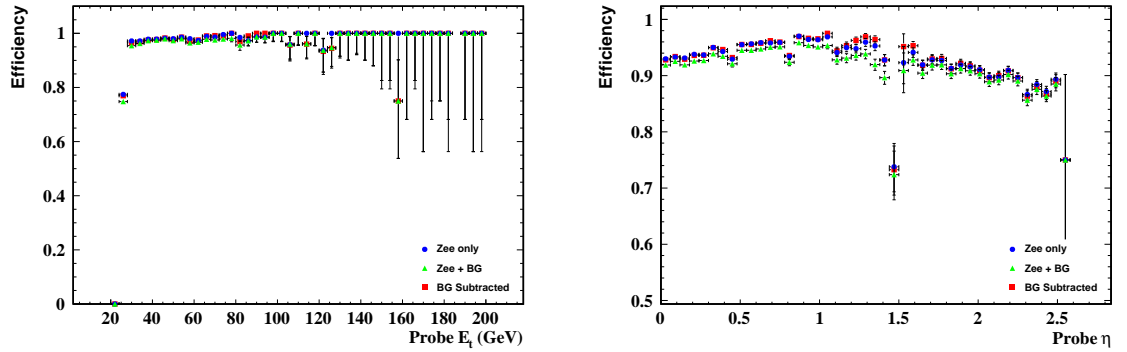


Figure 8.33: Tag and probe measurement of $\epsilon_{\text{Ele ID}}$

Table 8.17: Integrated electron ID efficiency (barrel)

<i>Dataset</i>	<i>Pass</i>	<i>Total</i>	<i>Efficiency</i>
$Z^0 \rightarrow ee$	44679	47616	$0.938^{+0.001}_{-0.001}$
BG	862	1950	
$Z^0 \rightarrow ee + \text{BG}$	45541	49566	$0.919^{+0.001}_{-0.001}$
BG Estimate	1443	3185	
BG Subtracted	44098	46381	$0.951^{+0.001}_{-0.001}$

Table 8.18: Integrated electron ID efficiency (endcap)

<i>Dataset</i>	<i>Pass</i>	<i>Total</i>	<i>Efficiency</i>
$Z^0 \rightarrow ee$	19297	21992	$0.877^{+0.002}_{-0.002}$
BG	368	737	
$Z^0 \rightarrow ee + \text{BG}$	19665	22729	$0.865^{+0.002}_{-0.002}$
BG Estimate	580.25	1141.25	
BG Subtracted	19084.8	21587.8	$0.884^{+0.002}_{-0.002}$

8.6.7 Trigger efficiency

Following the efficiency factorisation scheme, the trigger efficiency is defined as the probability of an electron to pass online selection given that a GSF electron which passes offline selection cuts was reconstructed from a real electron. The Tag and Probe criteria used for the measurement are detailed in Table 8.19. The invariant mass plots of Tag + All Probes and Tag + Passing Probes used to compute the background estimation are shown in Figs. 8.34 and 8.35.

The resulting efficiency with respect to probe E_t and η are shown in Fig. 8.36, with the integrated efficiencies shown in Tables 8.20 and 8.21. The effect of the background estimation is to overestimate the true measured efficiency by 0.04% in the barrel and underestimate by 0.05% in the endcap.

Table 8.19: Tag and Probe criteria for trigger efficiency measurements

<i>Tag Criteria</i>	<i>Probe Criteria</i>
GSF Electron	GSF Electron
$ \eta < 1.44$ or $1.56 < \eta < 2.5$	$ \eta < 1.44$ or $1.56 < \eta < 2.5$
$E_t > 15$ GeV	$E_t > 20$ GeV
Passes modified HEEP cuts	Passes modified HEEP cuts
$\Delta R(\text{Ele.}, \text{HLT Cand.}) < 0.1$	$\Delta R(\text{Ele.}, \text{HLT Cand.}) < 0.1$




Figure 8.34: Tag + All Probes invariant mass spectrum for $\epsilon_{\text{Trigger}}$




Figure 8.35: Tag + Passing Probes invariant mass spectrum for $\epsilon_{\text{Trigger}}$

8.6.8 Combination of integrated efficiencies

With the integrated efficiencies measured, the total event efficiency for the three possible topologies (Barrel - Barrel, Barrel - Endcap and Endcap - Endcap) can be computed from Eq. 8.3. The combined efficiencies are 0.81 (EB-EB), 0.73 (EB-EE) and 0.65 (EE-EE).

8.6 Efficiency measurements

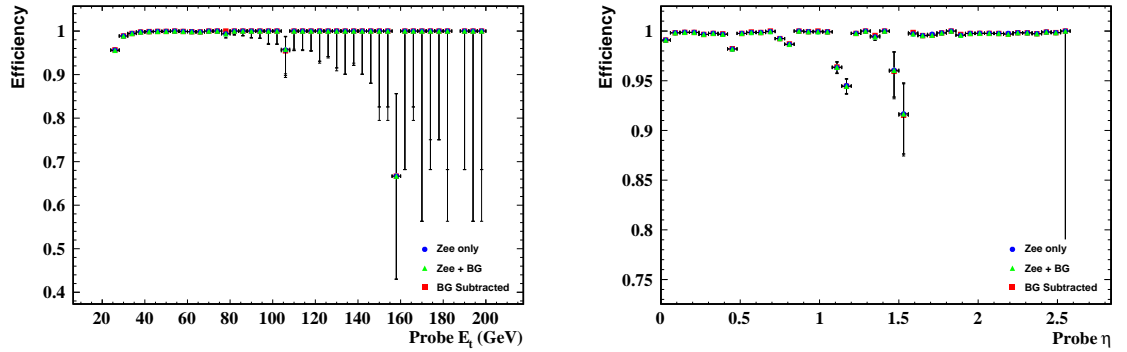


Figure 8.36: Tag and probe measurement of $\epsilon_{\text{Trigger}}$

Table 8.20: Integrated trigger efficiency (barrel)

<i>Dataset</i>	<i>Pass</i>	<i>Total</i>	<i>Efficiency</i>
$Z^0 \rightarrow ee$	40983	41223	$0.9942^{+0.0004}_{-0.0004}$
BG	758	759	
$Z^0 \rightarrow ee + \text{BG}$	41741	41982	$0.9943^{+0.0004}_{-0.0004}$
BG Estimate	453.5	468.5	
BG Subtracted	41287.5	41513.5	$0.9946^{+0.0004}_{-0.0004}$

Table 8.21: Integrated trigger efficiency (endcap)

<i>Dataset</i>	<i>Pass</i>	<i>Total</i>	<i>Efficiency</i>
$Z^0 \rightarrow ee$	17497	17545	$0.9973^{+0.0004}_{-0.0004}$
BG	312	325	
$Z^0 \rightarrow ee + \text{BG}$	17809	17870	$0.9966^{+0.0004}_{-0.0005}$
BG Estimate	200.25	204.5	
BG Subtracted	17608.8	17665.5	$0.9968^{+0.0004}_{-0.0004}$

8.6.9 Cross check with $W^\pm \rightarrow e\nu$

With the efficiencies measured from data, it is desirable to compare the results to an unbiased Monte Carlo sample. This is for two reasons. For the purposes of this analysis, a cross-check will validate the choice of factorisation scheme and measurement method. For real data, the cross-check ensures the detector performance is well understood and modeled.

To perform the cross check, a sample of $W^\pm \rightarrow e\nu$ decays was used. For each efficiency measurement, only the probe criteria were considered. Each object passing the relevant probe criteria was also required to match (with $\Delta R < 0.1$) the Monte Carlo truth electron from the W decay. The resulting efficiencies are shown in Figs. 8.37 to 8.40.

In general, the two methods agree very well. However, the electron ID efficiency measurements (Fig. 8.39) show a disagreement at the % level with respect to η . As only a single probe object match is performed, this cross-check method does not take into account the factorisation of efficiency measurements implicit in the Tag and Probe method. Additionally, the Monte Carlo object match results in an object sample of higher purity than even high-quality probe objects in the Tag and Probe method. These systematic differences between the two measurements result in the small disagreements observed.

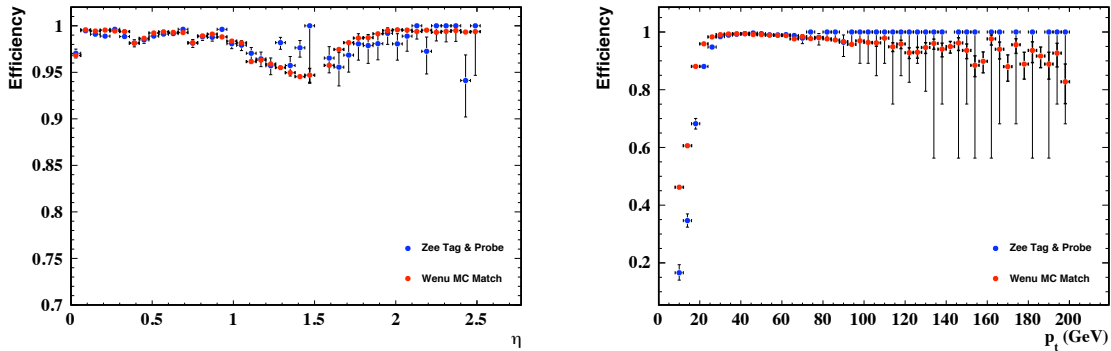


Figure 8.37: $W^\pm \rightarrow e\nu$ cross check of clustering efficiency measurement

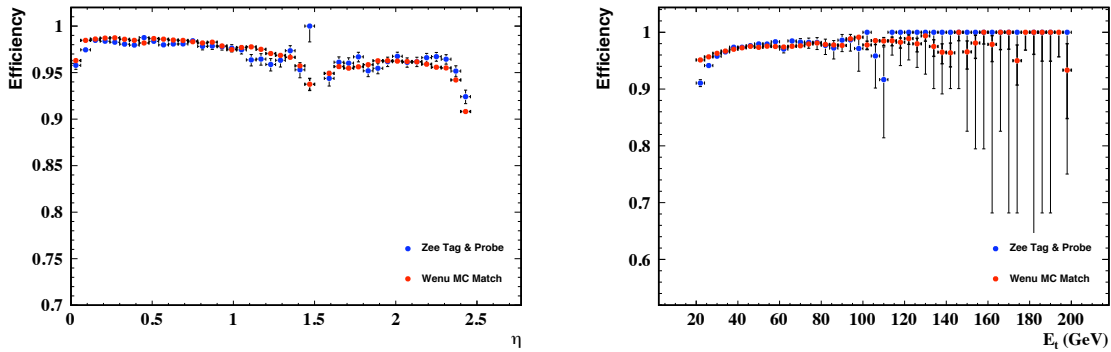


Figure 8.38: $W^\pm \rightarrow e\nu$ cross check of GSF electron reconstruction efficiency measurement

8.6 Efficiency measurements

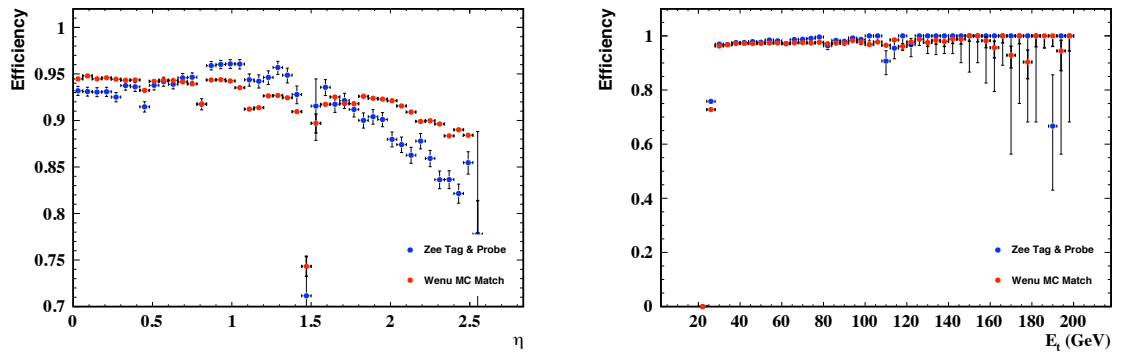


Figure 8.39: $W^\pm \rightarrow e\nu$ cross check of electron ID efficiency measurement




Figure 8.40: $W^\pm \rightarrow e\nu$ cross check of trigger efficiency measurement

8.6.10 Electron identification $N - 1$ efficiencies

While the Tag and Probe method provides a high-purity sample of electrons with which to measure efficiencies, the purity can be improved further to allow the efficiency of each electron identification cut to be measured. The $N - 1$ method applies all cuts except the one of interest to a candidate electron. If all cuts pass, that electron is then used to measure the cut of interest.

The Tag and Probe criteria used are listed in Table 8.22. Figs. 8.41 and 8.42 show that the efficiencies measured with the $N - 1$ method are stable in the presence of background for two cuts, $|\Delta\eta_{in}|$ and $E^{2\times 5}/E^{5\times 5}$. The plots for all cuts in the barrel and endcap are shown in Appendix B.

Table 8.22: Tag and Probe criteria for $N - 1$ efficiency measurements

<i>Tag Criteria</i>	<i>Probe Criteria</i>
GSF Electron	GSF Electron
$ \eta < 1.44$ or $1.56 < \eta < 2.5$	$ \eta < 1.44$ or $1.56 < \eta < 2.5$
$E_t > 15$ GeV	$E_t > 20$ GeV
Passes modified HEEP cuts	Passes $N - 1$ modified HEEP cuts




Figure 8.41: Example $N - 1$ efficiency measurement of $|\Delta\eta_{in}|$




Figure 8.42: Example $N - 1$ efficiency measurement of $E^{2\times 5}/E^{5\times 5}$

8.7 Background estimation

8.7.1 General issues

The background datasets considered for this analysis are detailed in §8.2. This list was produced by considering which physics processes with cross sections of equal order to, or greater than, the excited quark production cross section ($\mathcal{O}(1 \text{ pb})$), could result in one or more electrons (real or mis-identified) in the final state.

8.7 Background estimation

All solid filled histograms are stacked. Where the term ‘Realistic sample’ is used, this corresponds to the combination of all of the analysed datasets.

8.7.2 $t\bar{t}$ using the b -tagging method

As there is a $t\bar{t}$ contribution to the $Z^0 p_t$ spectra after all cuts have been applied, a method to estimate it is required. The b -tagging method, described in [86], is robust for start-up and can be applied on top of the existing event selection. As the branching ratio $\text{BR}(t \rightarrow Wb)$ is close to 100% [89], the $t\bar{t}$ background to this search channel is due to the case where the two W bosons decay to $e\nu$.

The jets used in this analysis are reconstructed with the iterativeCone5CaloJets algorithm [90]. Jets are identified as resulting from the decay of a b -quark due to the long lifetime of the b , which leads to a decay vertex displaced from the primary interaction vertex. Jets which are identified as being due to the decay of a b -quark are termed b -tagged jets. Various algorithms for b -tagging are available, with the jetBProbabilityJetTag algorithm used here being one of those implemented in the CMS software framework [91].

The b -tagging method described here is run on top of the existing event selection, and as such identifies $t\bar{t}$ events which have two high quality electrons and one or two identified b -jets from the $t\bar{t}$ decays. The technique is not sensitive to direct $b\bar{b}$ production as such events are highly unlikely to also contain two high-quality isolated electrons. Additionally, the jet collections are explicitly cleaned to remove jets which are also identified as electrons passing the selection criteria.

The observed number of events with exactly one and two b -tags are given by n_1 and n_2 . These are related to the actual number of events with one and two b -jets within detector acceptance (N_1, N_2) by

$$n_1 = N_1\epsilon_b + 2N_2\epsilon_b(1 - \epsilon_b), \quad (8.4)$$

$$n_2 = N_2\epsilon_b^2, \quad (8.5)$$

where ϵ_b is the b -tagging efficiency. N_1 and N_2 are related to the true number of $t\bar{t}$ events by

$$N_1 = NA_1, \quad (8.6)$$

$$N_2 = NA_2, \quad (8.7)$$

where A_1 (A_2) is the geometric acceptance for exactly one (two) b -jets from a $t\bar{t}$ event, determined from Monte Carlo simulation. To determine the b -tagging efficiency from

data, Eqs. 8.6 and 8.7 are substituted into Eqs. 8.4 and 8.5 which yields

$$n_1 = NA_1\epsilon_b + 2NA_2\epsilon_b(1 - \epsilon_b), \quad (8.8)$$

$$n_2 = NA_2\epsilon_b^2. \quad (8.9)$$

Rearranging 8.9 for N and substituting into 8.8 results in an expression for ϵ_b in terms of measurable quantities,

$$\epsilon_b = \frac{A_1/A_2 + 2}{n_1/n_2 + 2}.$$

With this measurement of ϵ_b performed, the true number of $t\bar{t}$ events, N , can be determined independently from either the n_1 or n_2 samples by rearranging Eqs. 8.8 and 8.9 for N , giving

$$N = \frac{n_1}{\epsilon_b(A_1 + 2A_2(1 - \epsilon_b))},$$

$$N = \frac{n_2}{A_2\epsilon_b^2}.$$

By binning the n_1 and n_2 samples as a function of di-electron p_t , N calculated for each bin will provide an estimate of the $t\bar{t}$ contamination on a bin-by-bin basis.

The jet selection criteria used are detailed in Table 8.23. The discriminant is chosen to err on the side of selection purity against the number of jets selected. The discriminant efficiency as a function of cut value is shown in Fig. 8.43, with the chosen b -jet efficiency indicated.




Figure 8.43: b -discriminant top identification efficiency as a function of cut value

For the above to be valid, the n_1 and n_2 samples must be of equivalent purity. The n_2 sample will be inherently more pure than the n_1 sample due to the requirement of two b -tagged jets. Fig. 8.44 shows that a large source of contamination is $Z^0 \rightarrow ee$ events

8.7 Background estimation

Table 8.23: Jet selection criteria

Jet Algorithm	iterativeCone5CaloJets
Jet E_t	> 20 GeV
Jet $ \eta $	< 2.4
B Discriminant	jetBProbabilityJetTags
Discriminant cut	> 4.0

Table 8.24: Measured b -tagging efficiency

Data sample	Efficiency
$t\bar{t}$ only	0.377 ± 0.058
$t\bar{t}$ only (tight)	0.382 ± 0.067
Realistic (tight)	0.356 ± 0.062

with an associated b -jet (real or faked). To veto these, $n_{1,\text{tight}}$ and $n_{2,\text{tight}}$ selections are defined where events with one or two b -jets are vetoed if the di-electron invariant mass is between 70 and 110 GeV. Figs. 8.45 and 8.47 illustrate the power of this veto to remove contamination whilst having a minimal impact on the true $t\bar{t}$ sample.

The acceptances as measured from Monte Carlo are $A_1 = 0.146 \pm 0.005$ and $A_2 = 0.79 \pm 0.01$. The measured b -tagging efficiencies are given in Table 8.24, and the total number of estimated $t\bar{t}$ event in Table 8.25. In these tables, ' $t\bar{t}$ only' refers to the technique run only on the $t\bar{t}$ sample and ' $t\bar{t}$ true' to the value calculated from Monte Carlo truth.




Figure 8.44: Selected events containing exactly 1 b -tagged jet for loose selection (n_1)




Figure 8.45: Selected events containing exactly 1 b -tagged jet for tight selection ($n_{1,\text{tight}}$)

8.7.3 Jet backgrounds using the fake rate method

The remaining jet backgrounds (γ +Jets, W^\pm +Jets and QCD Di-Jets) can be estimated using the fake rate method. This method makes use of the fact that events with one selected




Figure 8.46: Selected events containing exactly 2 b -tagged jets for loose selection (n_2)




Figure 8.47: Selected events containing exactly 2 b -tagged jet for tight selection ($n_{2,tight}$)

Table 8.25: Total number of estimated $t\bar{t}$ events

Data sample	Events with $70 < M_{ee} < 110(\text{GeV})$
$t\bar{t}$ true	378
$t\bar{t}$ only (n_1)	374 ± 95
$t\bar{t}$ only (n_2)	381 ± 101
Realistic (n_1)	428 ± 119
Realistic (n_2)	428 ± 124




Figure 8.48: $t\bar{t}$ estimate using $n_{1,tight}$ and $n_{2,tight}$ selections

8.7 Background estimation

signal electron are much more likely than events containing two selected signal electrons. There are two stages to the method: Firstly, a sample unbiased with respect to the signal selection is used to measure the probability that a jet fakes a selected signal electron. Secondly, this measurement is applied to all the jets in events with only one reconstructed signal electron to estimate the fake background.

The unbiased measurement sample is selected using jet triggers. As the current HLT menu was not available in the software version used for this study, and wanting to make the most use of the available simulated data, a pseudo-HLT reweighting scheme was used. The triggers used, and the prescales applied, are listed in Table 8.26. The *JetX* trigger requires at least one reconstructed jet in the event to have $p_T > X$ GeV. As for standard offline jets, online jets are reconstructed with the iterativeCone5CaloJets algorithm [90]. No further cleaning is applied. The leading offline jet in each event was taken to trigger the highest E_t trigger available, and the event weight was then scaled by the inverse of the appropriate trigger prescale. As the 1E31 HLT menu applies the jet trigger criteria to energy-corrected jets, this treatment on reconstructed jets is consistent with the actual HLT.

Table 8.26: Jet triggers (1E31 v0.6 menu) used to construct the fake rate estimate

<i>Trigger</i>	<i>L1 Prescale</i>	<i>HLT Prescale</i>	<i>Total Prescale</i>
Jet30	1000	5	5000
Jet50	100	2	200
Jet80	10	2	20
Jet110	1	1	1

The fake rate measurement is performed in bins of jet E_t , and is defined for each bin as

$$F(E_t) = \frac{\sum \text{Jet objects passing tight electron selection}}{\sum \text{Jet objects passing loose electron selection}}. \quad (8.10)$$

The loose (denominator) selection is detailed in Table 8.27, and the tight (numerator) selection in Table 8.28. The ΔR requirement between the triggered and candidate jet is to ensure the set of measurement objects is not enhanced by the initial trigger selection.

The measured numerator and denominator jet spectra are shown in Figs. 8.49 and 8.50 for the barrel, and in Figs. 8.51 and 8.52 for the endcap. Dividing the numerator spectrum by the denominator spectrum, on a bin-by-bin basis, yields the fake rate as a function of E_t , as shown in Fig. 8.53 for the barrel and Fig. 8.54 for the endcap. There is some contamination to the fake rate from events which are not from the di-jet sample, but the effect is within statistical errors at 200 pb^{-1} .

With the fake rate measurement performed, it is then possible to estimate the jet contribution from all events which contain one selected signal lepton. For all events, the

Table 8.27: Loose jet fake rate selection

Cut	Value
$\Delta R(\text{Trig.}, \text{Cand.})$	> 0.2
$ \eta $	< 2.5
E_t	$> 20 \text{ GeV}$
Had / EM	< 0.2

Table 8.28: Tight jet fake rate selection

Cut	Value
$\Delta R(\text{Trig.}, \text{Cand.})$	> 0.2
$ \eta $	< 2.5
E_t	$> 20 \text{ GeV}$
Modified HEEP selection cuts	Must pass




Figure 8.49: Measured denominator jet spectrum in the barrel




Figure 8.50: Measured numerator jet spectrum in the barrel




Figure 8.51: Measured denominator jet spectrum in the endcap




Figure 8.52: Measured numerator jet spectrum in the endcap

8.7 Background estimation




Figure 8.53: Measured jet fake rate in the barrel




Figure 8.54: Measured jet fake rate in the endcap

highest E_t electron is taken as the triggered lepton⁴, and must pass the tight selection criteria. All other objects passing the loose selection but not the tight selection (to remove signal bias) are then histogrammed by trigger-fake pair p_t , weighted according to the fake rate as given by the loose object ,

$$N(p_t) = \sum_{\text{loose}} w_e \frac{F(E_t)}{1 - F(E_t)},$$

where $F(E_t)$ is given by Eq. 8.10 and w_e is the event weight. Pairs are excluded from this process if their invariant mass lies within the range $70 < M_{ee} < 110$ GeV to reject $Z^0 \rightarrow ee$ events with one low-quality electron. Correlated errors are added linearly, and uncorrelated errors quadratically, by tracking the error-weighted event weights per fake object E_t in a separate histogram. The bin error can therefore be expressed as

$$\Delta N(p_t) = \sqrt{\sum_{E_t \text{ bins}} \left(\sum_{\text{loose}} w_e \Delta F(E_t) \right)^2}$$

where the *loose* sum runs over all objects in the given p_t bin and enclosing E_t bin, and the E_t bins sum runs over all possible fake rate bins.

The results of the estimation as applied to W +Jets, γ +Jets and QCD di-jet events are shown in Figs. 8.55 and 8.56, along with the actual number of events passing the full analysis selection. The method can be seen to match the true selection well, with the power of the available statistics particularly visible for the sample containing QCD di-jets.

⁴This is valid as the L1 trigger orders electron candidates by E_t and the single photon / electron HLT paths are seeded by the highest-ranked object.




Figure 8.55: Estimated jet background for W^\pm +Jets and γ +Jets




Figure 8.56: Estimated jet background for W^\pm +Jets, γ +Jets and QCD di-jets




Figure 8.57: Estimated jet fake rate for realistic data sample (no QCD Di-Jet). The $t\bar{t}$ contribution estimated in §8.7.2 has been subtracted.

8.7.4 $Z^0 \rightarrow e^+e^-$ with Monte Carlo and W^\pm hadronic recoil

The main remaining background is due to irreducible $pp \rightarrow Z^0 \rightarrow e^+e^-$ production. Two methods have been devised to account for this. Firstly the p_t spectrum of the selected Z^0 candidates will be determined from Monte Carlo simulation, scaled to fit the data in the region $0 - 250$ GeV after the other backgrounds have been considered. This method allows a check of the background slope shape. However, this relies on the Monte Carlo description of the high p_t region to be correct in the region where the signal is expected, as such a derivation of the high p_t shape from data is desired to be able to cross-check the Monte Carlo behaviour.

In the kinematic region above the W^\pm and Z^0 masses, the two can be considered to have identical production and decay kinematics. As the W^\pm production cross section is roughly 3 times that of the Z^0 , and the branching ratio $W^\pm \rightarrow e\nu$ is roughly 3 times that

8.7 Background estimation

of $Z^0 \rightarrow e^+e^-$, a factor of 10 more W^\pm than Z^0 events are expected which can be used to estimate the $Z^0 p_t$ spectrum.

As the observable of interest is the boson p_t , complete reconstruction of the W^\pm decay products is not required. Instead, events containing exactly one well isolated (passing the full HEEP selection) electron are selected⁵. The four-vectors of all jets which are separated from the electron by $\Delta R > 0.4$ with loose selection cuts ($E_t > 20$ GeV and $|\eta| < 2.5$) are then summed to yield the recoil system p_t . The p_t correlation between the recoil system and the W^\pm can be checked by performing the same summation in events with high-quality Z^0 candidates, and comparing the recoil and leptonic p_t . Such a plot also allows the jet energy scale to be determined, and a correction derived if required. The resulting correlation is shown in Fig. 8.58. The backgrounds to the W^\pm estimate are shown in Fig. 8.59, and are dominated by QCD below 150 GeV.

With the $W^\pm p_t$ spectrum measured, it must be scaled to correct for the cross section and branching ratio differences. To remove the effect of the QCD contamination, this is performed in the 150 – 250 GeV region. The comparison of the recoil method with the Monte Carlo method is shown in Fig. 8.60⁶. As it was available, a $Z^0 \rightarrow ee$ sample statistically independent from that used to create the pseudo-data was used for the scaled estimate. Above the boson mass scale ($\mathcal{O}(100$ GeV)), the methods agree very well, with the W^\pm extending well into the low-statistics tail as expected. The methods do not agree at very low p_t , which is expected due to the $E_t > 20$ GeV cut on recoil jets.




Figure 8.58: Correlation between p_t of hadronic recoil system and Z^0




Figure 8.59: Background contamination to the $W^\pm p_t$ estimate

While it is desirable to estimate the $Z^0 \rightarrow e^+e^-$ background from data, this may not be possible in the presence of a signal. An object which decays to a Z^0 would also be likely to decay to W^\pm . Indeed, the cross section for $pp \rightarrow u^* \rightarrow W^\pm d \rightarrow e\nu_e d$ ($M_{u^*} = 1$ TeV) is 3.7 pb, compared to 0.37 pb for the $Z^0 \rightarrow e^+e^-$ channel. However, preliminary investigations (§8.10.3) show that it may be possible to separate $W^\pm + \text{Jets}$ from $u^* \rightarrow$

⁵As HEEP selected electrons will have passed the HLT_Photon15.L1R trigger (§8.3.2), such events need no further trigger requirements for inclusion in the analysis dataset.

⁶QCD di-jets were not considered in the $Z^0 \rightarrow e^+e^-$ estimation due to large event weights biasing the normalisation.




Figure 8.60: Performance of the W^\pm recoil estimate against $Z^0 \rightarrow e^+e^-$ from Monte Carlo

$W^\pm d$. The $\gamma + \text{jets}$ channel may also be considered, as the photon production kinematics are similar in the high p_t , central η range to that of the Z^0 . However, the same problem remains; for the u^* considered here, the $u \rightarrow \gamma u$ cross section is 6.6 pb. It may transpire that the only available method is that of estimation from Monte Carlo, for which a large-statistics, inclusive jet sample will be required.

8.7.5 Combination of backgrounds

The complete combined background estimation for a 200 pb^{-1} pseudoexperiment is shown in Fig. 8.61, with and without a 1 TeV u^* . The decay $u^* \rightarrow Wd$ is explicitly excluded.




Figure 8.61: Combined background estimates without signal




Figure 8.62: Combined background estimates with 1 TeV u^*

8.7 Background estimation

8.7.6 Control of backgrounds

Motivated by a desire to be insensitive to uncertainties in calorimeter inter-calibration and energy scales in the high energy region, no mass cut on the e^+e^- invariant mass is used. However, as understanding of the high-energy regime develops, and the energy scales are determined, a cut on the lepton pair invariant mass could be reliably applied. Figs. 8.63 and 8.64 show the effect a cut of $60 < M_{e^+e^-} < 120$ GeV has on the analysis selection. It is clearly desirable to include this cut as soon as possible due to the effect on the QCD di-jet background.

A further cut can be made by requiring both selected electrons to have opposite charge. The effect of this cut is shown in Fig. 8.65, and is of equal power to the invariant mass cut for the QCD di-jet background, but less powerful for the other backgrounds. The combination of both cuts is shown in Fig. 8.66.




Figure 8.63: Selected signal + background events without control cuts




Figure 8.64: Selected signal + background events with $60 < M_{ee} < 120$ GeV cut




Figure 8.65: Selected signal + background events with opposite sign electron cut




Figure 8.66: Selected signal + background events with opposite sign electron and $60 < M_{ee} < 120$ GeV cuts

While the control cuts appear powerful, they must be taken in the context of event yield change for signal samples. Figs. 8.67 and 8.68 show the percentage change in the event yield for u^* masses of 1 and 1.75 TeV. With real data, the 10% loss will have to

be taken into consideration along with the actual observed jet background spectrum to determine which cuts to make.




Figure 8.67: Event yield change for 1 TeV u^* with both control cuts




Figure 8.68: Event yield change for 1.75 TeV u^* with both control cuts

8.8 Systematic Uncertainties

To calculate the effect of systematic uncertainties on the search reach determination, the uncertainty due to electron identification cut selection, PDF choice, PDF uncertainty, PDF factorisation scale, and the renormalisation scale on each signal and background sample was determined. The high-statistics histograms from which search determinations are run were scaled to minimise or maximise the search potential to determine upper and lower systematic limits.

To determine the lower systematic limit of the search potential, the expected event count in each background histogram was scaled up on a bin-by-bin basis, while the signal component was scaled down. The reverse was performed when attempting to maximise the search potential. Therefore, all uncertainties are presented in terms of the uncertainty in the selected di-electron event yield.

8.8.1 Electron identification

Following the factorisation scheme given in §8.6.3, the number of events passing final electron ID selection is

$$N_P = N_T \epsilon_{ID}^2,$$

where N_T is the number of events with at least two reconstructed electron candidates. To investigate the difference in electron identification efficiency between Monte Carlo simulation and data, each selection variable (detailed in Table 8.5) was smeared by an arbitrary conservative 10% Gaussian to measure the change in event selection efficiency.

8.8 Systematic Uncertainties

Given the new efficiencies, the expected number of events is given by

$$N_{\text{scaled}} = N_P \frac{\epsilon_{\text{new}}^2}{\epsilon_{\text{old}}^2},$$

and the uncertainty is therefore given by

$$\frac{N_{\text{scaled}} - N_P}{N_P} = \frac{\epsilon_{\text{new}}^2}{\epsilon_{\text{old}}^2} - 1. \quad (8.11)$$

Given a change in efficiency of

$$\Delta\epsilon = \frac{\epsilon_{\text{new}} - \epsilon_{\text{old}}}{\epsilon_{\text{old}}},$$

this can be rearranged to substitute back into Eq. 8.11, giving the event yield uncertainty in terms of a variation in efficiency due to smearing as

$$\frac{N_{\text{scaled}} - N_P}{N_P} = (\Delta\epsilon - 1)^2 - 1.$$

The changes in efficiency are shown in Table 8.29. The global uncertainty due to electron ID cuts is therefore taken to be 5%. After this analysis was completed, it was shown that the arbitrary 10% is indeed very conservative; the data to MC agreement has been measured to be $97.9 \pm 0.6\%$ (EB) and 99.3 ± 1.1 (EE) [92]. However, this does not significantly modify the search reach as this is not the leading systematic.

Table 8.29: Systematic effect of electron ID cuts

<i>Cut</i>	<i>EB</i>	<i>EE</i>
$ \Delta\eta_{\text{in}} $	-0.18%	-0.009%
$ \Delta\phi_{\text{in}} $	+0.001%	-0.02%
$\sigma_{i\eta j\eta}$	n/a	-2.6%
H/E	-0.02%	+0.2%
$E^{2\times 5}/E^{5\times 5}$	-2.4%	n/a
Had D2 Isolation	n/a	-0.02%
Combined	-2.4%	-2.6%
Event uncertainty	-4.7%	-5.1%

8.8.2 Calibration and alignment

All samples were reconstructed with the `IDEAL_V12:ALL` global tag, which describes a perfectly aligned and calibrated detector. To quantify the effect of a mis-calibrated and mis-aligned detector, the 1, 1.5 and 2 TeV signal samples were reconstructed with the `STARTUP_V12:ALL` tag.

For the subdetectors of interest, this set of conditions corresponds to the tracker aligned with 2008 cosmic muon data, and the ECAL pre-calibrations determined from cosmic muon data in the barrel, and laboratory measurements of crystal light yields and photo-detector gains in the endcaps. The resulting precisions are 1.5% ($|\eta| < 1$), 1.5 \rightarrow 2.2% ($1 < |\eta| < 1.479$) and 10% ($|\eta| > 1.479$). It should be noted that 9 of the 27 barrel super-modules are actually calibrated to 0.3% with test beam data.

For the 1 (2) TeV u^* sample, the effect of startup conditions was to lower the number of accepted events after all selection criteria had been applied by 2.6% (2.4%). The $Z^0 \rightarrow e^+e^-$ sample, reconstructed with startup conditions, resulted in a 2.4% reduction in selected events. The systematic effect of startup misalignment is therefore taken to be 2.5%.

8.8.3 PDF uncertainty

The Hessian method, described in §3.2.2 can be used to determine the uncertainty in the $Z^0 p_t$ spectrum due to the PDF fit. The results of this method as applied to the p_t of a Z^0 from the decay of a u^* with masses 1 and 2 TeV are shown in Figs. 8.69 and 8.70. As the majority of the events fall below the $m_{u^*}/2$ region due to kinematic limits, an uncertainty of between 3% and 10% will be taken for the signal samples. All backgrounds will be subject to a 5% systematic error due to PDF uncertainty.




Figure 8.69: PDF uncertainty on $p_{t,Z}$ for a 1 TeV u^* with the CTEQ6M PDF set




Figure 8.70: PDF uncertainty on $p_{t,Z}$ for a 2 TeV u^* with the CTEQ6M PDF set

8.8.4 Factorisation and renormalisation scales

Production cross sections at leading order are sensitive to variations in the factorisation scale, μ_f , at which the incoming parton PDFs are probed. At next to leading order in QCD, new processes which introduce a factor of $\alpha_s(\mu_r)$ become available, where μ_r is the renormalisation scale at which the strong coupling is evaluated. In general, the NLO cross section cancels the LO dependence on μ_f . The same arguments hold for higher calculation orders (e.g. [19]).

8.8 Systematic Uncertainties

It is therefore desirable to determine the effect of the choice of μ_f and μ_r on the assumed production cross-sections with an NLO calculation where available. The dependence on the two scales is generally tested by setting $\mu_f = \mu_r$ equal to some relevant event scale. MC@NLO was used, with μ_f and μ_r set to the scaled geometric mean of the transverse mass of all final-state particles,

$$\mu_f = \mu_r = S \left(\prod_{i=1}^N (M_i^2 + p_{t,i}^2) \right)^{\frac{1}{N}},$$

where S is the required scaling and i runs over all N final-state particles. S was set to 2 and 0.5 to get upper and lower boundaries. The uncertainty is given by

$$\frac{\sigma_{S_x} - \sigma_{S_1}}{\sigma_{S_1}},$$

where σ_{S_1} is the reference cross section with $S = 1$.

Not all channels of interest are available in MC@NLO, and therefore MadGraph was used for $\gamma + \text{Jets}$, and CALCHEP for the q^* dependence. In all cases the PDF set CTEQ6M was used. The resulting uncertainties are given in Table 8.30.

Table 8.30: Systematic effect of altering μ_f and μ_r

<i>Channel</i>	$S = 2$	$S = 0.5$
Di-boson	+2%	-2%
$t\bar{t} + \text{Jets}$	-13%	+14%
$W^\pm + \text{Jets}$	+2%	-2%
$Z^0 + \text{Jets}$	+4%	-6%
$\gamma + \text{Jets}$	-6%	+1%
u^* (1 TeV)	-8%	+10%
u^* (1.25 TeV)	-7%	+10%
u^* (1.5 TeV)	-8%	+10%
u^* (1.75 TeV)	-8%	+10%
u^* (2 TeV)	-8%	+10%

8.8.5 PDF choice

The effect of the PDF choice on global production cross sections was quantified by running the MadGraph simulations, as prepared above, with three different PDF sets. The difference in cross section from the baseline CTEQ6M set (with $S = 1$) was calculated, and these differences summed in quadrature. The uncertainties, shown in Table 8.31, are defined as

$$\frac{\sigma_{\text{New}} - \sigma_{\text{CTEQ6M}}}{\sigma_{\text{CTEQ6M}}}.$$

Table 8.31: Systematic effect of PDF choice

<i>Channel</i>	CTEQ6L1	CTEQ5M	MRST2002NLO
Di-boson	+5%	+7%	+2%
$t\bar{t}$ + Jets	+5%	+5%	+8%
W^\pm + Jets	+10%	+3%	0%
Z^0 + Jets	-5%	+5%	+1%
γ + Jets	+20%	0%	-3%
u^* (1 TeV)	-3%	+4%	+3%
u^* (1.25 TeV)	-4%	+6%	-3%
u^* (1.5 TeV)	-6%	+5%	+4%
u^* (1.75 TeV)	-7%	+5%	+4%
u^* (2 TeV)	-8%	+4%	+3%

8.8.6 High order corrections

High order corrections are explicitly accounted for in the simulation used as the dominant background, $Z^0 \rightarrow ee$, is produced with either Pythia ISR (for the Z_{ee} sample), which is tuned to accurately simulate existing observed p_t spectra or with complete real QCD NLO and higher order corrections from Madgraph (for the Z_{Jets} sample). While the virtual effects may be negative, they are not included as §3 shows the real effect on the boson p_t dominates. The EWK correction to the di-electron transverse momentum is not relevant relative to the QCD corrections, and is therefore not included.

8.8.7 Combination

The PDF effects are correlated, and therefore are added linearly. The remaining uncertainties are treated in quadrature. The expressions for the combined systematic variation for each dataset are, therefore,

$$(\Delta S_{\text{sys}})^2 = (\Delta P_{\text{Choice}} + \Delta P_{\text{Uncert.}})^2 + (\Delta\sigma_\mu)^2 + (\Delta\epsilon_{\text{Ele ID}})^2 + (\Delta C)^2,$$

where the uncertainties are given by ΔP_{Choice} for the PDF choice, $\Delta P_{\text{Uncert.}}$ for the PDF fit, $\Delta\sigma_\mu$ for the factorisation and renormalisation scale, $\Delta\epsilon_{\text{Ele ID}}$ for the electron ID and ΔC for calibration and alignment.

The effects need to be considered in respect to a claimed search reach. As such, each parameter is chosen to increase signal and reduce background for an upper limit on the search reach, and the reverse to determine the lower limit. The combined uncertainties are shown in Tables 8.32 and 8.33.

8.9 Search reach at $\sqrt{s} = 10$ TeV

Table 8.32: Combination of systematic uncertainties to maximise search reach

<i>Channel</i>	<i>PDF uncert.</i>	<i>PDF choice</i>	<i>μ scale</i>	<i>Ele. ID</i>	<i>Cal. & Align.</i>	<i>Combination</i>
Di-boson	-5%	-4.5%	-2%	-5%	-2.5%	-11%
$t\bar{t}$ + Jets	-5%	-5.5%	-13%	-5%	-2.5%	-18%
W^\pm + Jets	-5%	-5%	-2%	-5%	-2.5%	-12%
Z^0 + Jets	-5%	-3.5%	-6%	-5%	-2.5%	-12%
γ + Jets	-5%	-10%	-6%	-5%	-2.5%	-17%
u^* (1 TeV)	+3%	+3%	+10%	0%	0%	+12%
u^* (1.25 TeV)	+4.8%	+4%	+10%	0%	0%	+13%
u^* (1.5 TeV)	+6.5%	+4.5%	+10%	0%	0%	+15%
u^* (1.75 TeV)	+8.3%	+4.5%	+10%	0%	0%	+16%
u^* (2 TeV)	+10%	+4.5%	+10%	0%	0%	+18%

Table 8.33: Combination of systematic uncertainties to minimise search reach

<i>Channel</i>	<i>PDF uncert.</i>	<i>PDF choice</i>	<i>μ scale</i>	<i>Ele. ID</i>	<i>Cal. & Align.</i>	<i>Combination</i>
Di-boson	+5%	+4.5%	+2%	0%	0%	+9%
$t\bar{t}$ + Jets	+5%	+5.5%	+14%	0%	0%	+18%
W^\pm + Jets	+5%	+5%	+2%	0%	0%	+10%
Z^0 + Jets	+5%	+3.5%	+4%	0%	0%	+9%
γ + Jets	+5%	+10%	+1%	0%	0%	+15%
u^* (1 TeV)	-3%	-3%	-8%	-5%	-2.5%	-11%
u^* (1.25 TeV)	-4.8%	-4%	-7%	-5%	-2.5%	-13%
u^* (1.5 TeV)	-6.5%	-4.5%	-8%	-5%	-2.5%	-15%
u^* (1.75 TeV)	-8.3%	-4.5%	-8%	-5%	-2.5%	-16%
u^* (2 TeV)	-10%	-4.5%	-8%	-5%	-2.5%	-17%

8.9 Search reach at $\sqrt{s} = 10$ TeV

At the time this analysis was performed, the LHC was expected to start up in the autumn of 2009 with the aim of delivering 200 pb^{-1} of integrated luminosity over its first run at $\sqrt{s} = 10$ TeV. It was not known what the profile of energy and luminosity would be with respect to time, and as such, it was considered instructive to consider the search reach for excited quark production in the di-electron channel as a function of luminosity, assuming \sqrt{s} stays constant at 10 TeV.

8.9.1 The p-value scan technique

A p-value scan is a technique used in generic searches for new physics, for example in previous high-mass di-electron searches at the Tevatron [93]. Given a binned data sample and background estimation, a sliding window of n bins is moved along the selected events sample, and the sum of bins in the window calculated. Following Poisson statis-

tics, the probability for the data sample to fluctuate from the background estimation to the observed value, or higher, is calculated. This probability is the bin p-value. For convenience, results will be presented in terms of $-\log_{10}(p)$.

A further step must be taken to quantify the fact that the background may fluctuate anywhere along the scan range (the so-called *look elsewhere effect*), and that due to the sliding window summation, all the bin p-values have intercorrelations. Scans are therefore performed on background-only pseudoexperiments to determine the significance of such effects.

8.9.2 Data preparation and background fit

The data used in the search reach determination is constructed by running all available events for each dataset through the full event selection procedure. Each event is weighted such that each dataset histogram represents the expected number of events for an integrated luminosity of 200 pb^{-1} . The expected per-bin event yield for a given luminosity, L , can therefore be determined by multiplying each bin by $L/200 \text{ pb}^{-1}$.

The background estimation is taken from a fit to the data in regions that are subject to signal exclusions from the Tevatron. The data is fitted to the expression $e^{-\alpha p_t} p_t^{-\beta}$, where α and β are the fit parameters, in the region $50 \leq p_t \leq 250 \text{ GeV}$. Typical fit values are $\alpha = 0.0131 \pm 0.002$ and $\beta = 1.8 \pm 0.4$. To factor the fit uncertainty into the search reach determination, the +ve (-ve) errors were taken as a further systematic to maximise (minimise) the search reach.

8.9.3 Search reach determination

Following the above, the search reach is determined in three phases. Firstly, 2×10^9 pseudoexperiments are thrown on the background-only hypothesis for a given luminosity, and the minimum p-value for the experiment determined. The resulting minimum p-values for each pseudoexperiment are histogrammed. In practice, 100000 bins between $-\log_{10} p = 0$ and 10 are used. 3σ and 5σ background rejection criteria can be determined by identifying the p-value, p , for which the ratio

$$\frac{\int_p^\infty f(p) d[-\log_{10} p]}{\int_0^\infty f(p) d[-\log_{10} p]}$$

equals 0.27% (3σ) or $5.733 \times 10^{-5}\%$ (5σ), where $f(p)$ represents the p-value histogram value for bin p . The p-value distribution is shown in Fig. 8.71, and the calculated background rejection criteria are shown in Figs. 8.73 and 8.74.

Although high statistics event samples were used to calculate the expected number of events for the pseudoexperiments, the dominant samples correspond to a luminosity




Figure 8.71: p-value distribution for background only pseudoexperiments with 200 pb^{-1}




Figure 8.72: p-value distribution for signal + background pseudoexperiments with 200 pb^{-1}

of $\ll 3000 \text{ pb}^{-1}$. To ensure this was not biasing the results when constructing pseudoexperiments, the technique was run in two ways. The first method was as described above, varying the bin content by $L/200 \text{ pb}^{-1}$, before running the fit. The second method performed a fit on the 200 pb^{-1} background-only histogram and projected this onto a new histogram, therefore producing a smooth background estimation. In both cases, the pseudoexperiments were then constructed by allowing the bin contents to vary following Poisson statistics. Both methods produce consistent results.

10000 pseudoexperiments are then thrown based on the expected number of background and signal events for a given luminosity, the background fit run, and the minimum p-value determined. An example p-value distribution is shown in Fig. 8.72. A weighted average of the minimum p-values is calculated as

$$p_{\text{likely}} = \frac{1}{\sum_{m=0}^N B(m)} \sum_{n=0}^N B(n)M(n),$$

where $B(n)$ and $M(n)$ are the count and midpoint of bin n . This represents the mean p-value to be observed at a given luminosity in the presence of a signal.

In all cases, the systematic errors derived in §8.8 are incorporated by varying the expected event yield up or down by the given amount per dataset, and re-running the complete statistical tools on the modified pseudodata.

The background only and signal + background p-value criteria can then be plotted as a function of luminosity, as shown in Fig. 8.75. The required luminosity for evidence (3σ) or discovery (5σ) can be determined from the intersection of the signal and background curves for each u^* mass. The derived 3σ and 5σ discovery potential curves are shown in Figs. 8.76 and 8.77.

There is an implicit signal width assumption in the p-value technique, where a given number of bins is summed over in the sliding window. As long as the background scans are performed with the same binning and window width, any window size can be used.

In the results presented here, a window of 3 bins (60 GeV) is used.

8.10 Future prospects

8.10.1 Application of p-value technique to data

The p-value technique described in §8.9.1 may be applied to a single data sample to determine the significance of any fluctuations from the background estimate. This is equivalent to looking at an infinitesimal slice in $\int L dt$ of the search reach determination plot (Fig. 8.75). Fig. 8.78 shows a 200 pb^{-1} pseudoexperiment which results in a 5σ background rejection p-value scan.

While the systematic uncertainty on integrated luminosity should be considered when applying the technique to data, the effect will be small. This is because the integrated luminosity is only used to scale the di-boson simulated background, which is a minor contributor to the background spectrum, and due to the background fluctuation p-values being well behaved with respect to changes in luminosity.

8.10.2 Analysis potential at $\sqrt{s} = 14$ and 7 TeV

Given $f_p(x, Q^2)$ from the PDF for a given initial state parton, p , the production cross section can be written as

$$\sigma = \sum_{p_1} \sum_{p_2} \int_{-\log \frac{\sqrt{s}}{M_{ll}}}^{\log \frac{\sqrt{s}}{M_{ll}}} f_{p_1}(x_1, M_{ll}^2) f_{p_2}(x_2, M_{ll}^2) \hat{\sigma} dy, \quad (8.12)$$

where $x_1, x_2 = \frac{M_{ll}}{\sqrt{s}} e^{\pm y}$, y is the rapidity of the produced system, \sqrt{s} the centre-of-mass energy of the collision, and the summation runs over all allowable initial state parton pairs. The ratio with Eq. 8.12 computed for two different centre-of-mass energies is equal to how the hard scattering cross section will change. Note that the $\hat{\sigma}$ terms cancel, so no knowledge of the matrix element is required. This relies on the assumption that production edge effects are not relevant (i.e. the significantly contributing part of M_{q^*} phase space is available in both cases).

Fig. 8.79 shows the calculation performed with the MSTW2008NLO PDF set, with $\sqrt{s} = 14$ and 10 TeV. At an intermediate mass of 1 TeV, roughly twice as many events are expected at $\sqrt{s} = 14$ TeV than at 10 TeV. In order to conservatively estimate the effect this has on the search reach, all event yields were scaled by 2, and the complete p-value scan re-evaluated. This ignores two effects; firstly that most of the background is in a regime where the event yield scaling is not as drastic, and secondly that pileup and the structure of the underlying event are likely to change.

8.10 Future prospects




Figure 8.73: 3σ background p-value rejection criteria




Figure 8.74: 5σ background p-value rejection criteria




Figure 8.75: Search reach determination




Figure 8.76: Expected 3σ evidence potential at $\sqrt{s} = 10$ TeV




Figure 8.77: Expected 5σ discovery potential at $\sqrt{s} = 10$ TeV




Figure 8.78: Example 200 pb^{-1} 5σ discovery pseudo-experiment background fit and p-value scan

The expected background and signal p-values are shown in Fig. 8.81, with the expected 3σ and 5σ curves shown in Figs. 8.82 and 8.83. As the background expectation values are well behaved with respect to luminosity, the statistical tools validate the naive assumption that twice the event yield requires half the luminosity to gain the same statistical significance. Discovery potential is therefore greatly enhanced by running at $\sqrt{s} = 14 \text{ TeV}$.

As of 2010, the LHC is running at $\sqrt{s} = 7 \text{ TeV}$, with the aim of delivering 1 fb^{-1} of integrated luminosity. It is instructive to apply the naive scaling, validated in the 14 to 10 TeV scaling, to the 7 TeV scenario. Fig. 8.80 shows the parton luminosity ratio between $\sqrt{s} = 10$ and 7 TeV. For a resonant mass of 1 TeV, the parton luminosities are reduced by a factor of 3. Therefore, 3σ evidence potential exists at an integrated luminosity of around 600 pb^{-1} , indicating the the channel remains viable in the first LHC run if the machine delivers the planned integrated luminosity.

8.10.3 Boosted W^\pm s

New physics with electroweak couplings would be expected to couple to W^\pm as well as to Z^0 . The relative coupling strengths depend on the underlying dynamics, but nonetheless the $q^* \rightarrow W^\pm q$ channel should be pursued to confirm a discovery or reject the signal hypothesis.

Following the boosted Z^0 analysis, the channel of interest is that where the W^\pm decays leptonically, $q^* \rightarrow Wq, W \rightarrow e\nu_e$. Such a decay leads to significant \cancel{E}_t , correlated with the electron direction. This allows the neutrino three-vector to be approximated in the collinear limit,

$$\vec{p}_{\nu_e} = (\cancel{E}_x, \cancel{E}_y, \frac{\sqrt{\cancel{E}_x^2 + \cancel{E}_y^2}}{\sqrt{p_{x,e}^2 + p_{y,e}^2}} p_{z,e}),$$

where \vec{p}_e is the electron momentum. The neutrino four-vector is defined as $p_\mu^{\nu_e} = (\vec{p}_{\nu_e}, |\vec{p}_{\nu_e}|)$.

Plotting the electron-neutrino invariant mass against the angle in ϕ between the elec-

8.10 Future prospects




Figure 8.79: Ratios of parton luminosities at $\sqrt{s} = 10$ and 14 TeV with the MSTW2008NLO PDF set




Figure 8.80: Ratios of parton luminosities at $\sqrt{s} = 7$ and 10 TeV with the MSTW2008NLO PDF set




Figure 8.81: Search reach determination at $\sqrt{s} = 14$ TeV




Figure 8.82: Estimated 3σ evidence potential at $\sqrt{s} = 14$ TeV




Figure 8.83: Estimated 5σ discovery potential at $\sqrt{s} = 14$ TeV




Figure 8.84: Discriminating boosted W^\pm s from background for signal (a), $W^\pm + \text{Jets}$ (b), $t\bar{t}$ (c) and $Z^0 \rightarrow e^+e^-$ (d)

tron and \cancel{E}_t provides a powerful discriminant between signal and background, as shown in Fig. 8.84. Such a discriminant may potentially allow the Z^0 background estimation technique, described in §8.7.4, to be used in the presence of signal, with a cut optimised to reject signal.

8.11 Conclusions

It is found that 3σ evidence for excited quark production at a mass of 1 TeV could be found with 200 pb^{-1} of integrated luminosity at $\sqrt{s} = 10 \text{ TeV}$. This requires modifications to the standard electron reconstruction, in order to correctly identify electrons which are very closely spaced in the electromagnetic calorimeter.

For the benchmark model, the expected search reach has been computed as a function of q^* mass, and scaled to the LHC first run beam parameters ($\sqrt{s} = 7 \text{ TeV}$). For a resonant mass of 1 TeV, 3σ evidence could be discovered with 600 pb^{-1} of data. As the LHC is aiming to deliver 1 fb^{-1} in the 2010–2011 running period at this energy, this analysis has shown that the q^* search in the di-electron channel is a viable first data search for CMS.

9. Conclusions

9.1 Work presented

A complete data collection, storage and processing chain is key to any particle physics analysis. At the experimental site, data must be captured on-detector, triggered and read out. Global transfer and processing of the data must then occur before it is available in a form suitable for end-user analysis. Additionally, accurate theoretical and phenomenological understanding of physical processes is vital to ensure that all known physics is accounted for at the same precision as current experimental measurements. Preparations for data taking across these areas with a particular emphasis on searches for new physics with the e^+e^- final state have been presented.

9.1.1 Level-1 triggering for high E_t electrons

Studies into Level-1 triggering for high E_t electrons have been performed. It is shown that the current detector configuration scheme would result in TeV scale electrons causing the trigger to fire on the incorrect bunch crossing, losing the tracker data for the actual event of interest. This is due to the configuration resulting in the trigger primitive generation electronics zeroing the effective E_t of a saturating ECAL channel.

A configuration scheme which correctly treats saturating channels has been presented. Simulation studies which validate the scheme for channels with incident energy of up to 4 TeV have been performed. This configuration scheme is now used on-detector, and the electronics simulation, a key part of the Level-1 trigger emulator, has been updated to reflect the new knowledge gained about the operation of the ECAL Front End electronics. Without this work, CMS would be unable to trigger on TeV scale electrons, which would leave potentially exciting new physics processes inaccessible.

9.1.2 The neutral-current Drell-Yan process in the high invariant mass region

While some physics models predict striking resonant signatures in the e^+e^- invariant mass spectrum, other models predict more subtle changes to observable distributions.

For example, models with TeV^{-1} scale extra dimensions, where the size of the extra dimension is of order 10 TeV , predict changes to the $M_{e^+e^-}$ differential cross section at the percent level. Therefore, an accurate description of the Standard Model processes contributing to the underlying Drell-Yan spectrum is required to ensure any observed deviations from the Standard Model are due to new physics, rather than poorly described Standard Model processes.

A study combining the best-available $\mathcal{O}(\alpha_s)$ QCD, $\mathcal{O}(\alpha)$ electroweak and higher order QED photon shower calculations has been performed. Processes involving photons in the initial state have also been included. It is shown that the combined QCD and EWK corrections partially cancel, resulting in a 10%–5% change in $\frac{d\sigma}{dM_{e^+e^-}}$ in the invariant mass region 200–1500 GeV. Previous studies of the Drell-Yan process in phenomenological and data analysis settings have only considered the QCD corrections as important; this work shows that the QCD and EWK corrections must both be considered as they are of equal order but opposite sign.

9.1.3 Optimising tape migration performance with the CASTOR HSM

The custodial storage of physics data is one of the main roles of the Tier-1 computing centres. Due to cost, space and power constraints, the medium of choice is magnetic tape. In the CMS computing model, all data at a site is stored on tape; disk copies are used as a transient cache of recently accessed (or soon to be accessed) files. To ensure that disk buffers do not completely fill with incoming data, either from WAN transfers or the output of local batch jobs, the ability to write files to tape at a rate greater than or equal to the incoming rate is key to the operation of a Tier-1 computing centre.

The UK Tier-1 resource planning calls for a tape write rate per drive of 90 MBs^{-1} . In data-challenge testing, write rates of only 16 MBs^{-1} were observed. Development of optimisations in disk server configuration and tape migration scheduling algorithms in the CASTOR HSM were undertaken. By optimising Linux filesystem parameters, and implementing new job scheduling and candidate file selection algorithms, the tape write rate per drive, including tape mount, seek and unmount time, has been shown to reach the required 90 MBs^{-1} in realistic operation. These optimisations are now used by all experiments at RAL, and in all CASTOR instances in use for experiment Tier-0s and user analysis at CERN.

9.1.4 High p_t Z^0 s as a probe for physics beyond the standard model

Most candidate theories for physics beyond the Standard Model predict new heavy particles which participate in electroweak interactions. As a result, these new heavy particles can decay into Z^0 bosons. As the scale of new physics ($\mathcal{O}(1 \text{ TeV})$) is much greater than the Z^0 mass, these candidate theories predict boosted Z^0 signatures. Due to low back-

9.2 Future prospects

grounds, and the fact that ECAL energy resolution improves with respect to increasing E_t , a search for boosted Z^0 s in the e^+e^- channel should be carried out at the LHC.

An analysis strategy has been presented, using an excited quark model ($pp \rightarrow u^* \rightarrow Z^0 u \rightarrow e^+e^-u$) to benchmark the search potential. The standard CMS electron reconstruction algorithms begin to break down for Z^0 s with $p_t \gtrsim 450$ GeV. Therefore, techniques to reconstruct electrons from boosted Z^0 s, where the decay products are close together, are presented. The new reconstruction is shown to be efficient up to at least a Z^0 p_t of 1.8 TeV, corresponding to the decay of a ≈ 4 TeV object. Using the new reconstruction techniques, it is shown that 3σ evidence of 1 TeV excited quark production could be found with 600 pb^{-1} of integrated luminosity at $\sqrt{s} = 7$ TeV, making this a viable search channel for early LHC running. At $\sqrt{s} = 14$ TeV, 3σ evidence could be found with 100 pb^{-1} , and 5σ with 200 pb^{-1} . 1 fb^{-1} would make masses in excess of 1.5 TeV accessible at 5σ .

9.2 Future prospects

With the LHC now operational with a target integrated luminosity for the first run of 1 fb^{-1} at $\sqrt{s} = 7$ TeV, physics analysis and analysis support (detector operations and computing) are now underway. The response of the Level-1 trigger to high E_t electrons must be studied to ensure that the configuration scheme presented in this thesis performs as expected on-detector. This will be possible as the integrated luminosity (and indeed collision energy) increases, yielding real high E_t electrons with which to perform the study.

Tape drive technology continues to progress. CMS is now using the T10000B drives at the UK Tier-1, which have a capacity of 1 TB and maximum write rate of 120 MBs^{-1} . The next generation of drives, the T10000C, will have a capacity of 5 TB and write rate of 240 MBs^{-1} . The UK Tier-1 is expected to migrate to these new drives as the T10000B drives are phased out. This will require infrastructure upgrades to support the increased write rate (the current maximum rate of 100 MBs^{-1} is limited by the 1 Gbs^{-1} networking throughout the Tier-1), by upgrading to 10 Gbs^{-1} networking for all disk servers, and potentially further development within CASTOR if individual disk servers can not support 240 MBs^{-1} throughput when sourcing files to migrate.

The boosted Z^0 analysis is being pursued with data within the CMS collaboration, with the aim of significantly improving on the existing Tevatron limits. It would be recommended to also perform a search for boosted W^\pm , as §8.7.4 shows that the production cross section could be higher than that of boosted Z^0 s for physics processes which couple to Z^0 and W^\pm .

A. Tag and Probe $E_t, |\eta|$ binned efficiencies

Table A.1: Clustering efficiency ($Z^0 \rightarrow e^+e^-$ only)

p_t (GeV) η	0.0 – 0.3	0.3 – 0.6	0.6 – 0.9	0.9 – 1.2	1.2 – 1.5	1.5 – 1.8	1.8 – 2.1	2.1 – 2.4	2.4 – 2.7	2.7 – 3.0
0 – 10	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
10 – 20	$0.57^{+0.03}_{-0.03}$	$0.55^{+0.03}_{-0.03}$	$0.45^{+0.03}_{-0.03}$	$0.36^{+0.03}_{-0.03}$	$0.34^{+0.04}_{-0.04}$	$0.43^{+0.06}_{-0.06}$	$0.55^{+0.05}_{-0.05}$	$0.71^{+0.05}_{-0.06}$	$0.8^{+0.1}_{-0.1}$	$0.5^{+0.5}_{-0.5}$
20 – 30	$0.959^{+0.007}_{-0.008}$	$0.940^{+0.009}_{-0.009}$	$0.964^{+0.007}_{-0.008}$	$0.91^{+0.01}_{-0.01}$	$0.90^{+0.01}_{-0.01}$	$0.91^{+0.02}_{-0.02}$	$0.97^{+0.01}_{-0.02}$	$1.00^{+0.01}_{-0.01}$	$1.00^{+0.00}_{-0.00}$	$0.5^{+0.5}_{-0.5}$
30 – 40	$0.991^{+0.002}_{-0.002}$	$0.989^{+0.002}_{-0.002}$	$0.989^{+0.002}_{-0.002}$	$0.992^{+0.002}_{-0.003}$	$0.990^{+0.003}_{-0.004}$	$0.99^{+0.01}_{-0.01}$	$0.98^{+0.01}_{-0.01}$	$0.99^{+0.01}_{-0.01}$	$1.00^{+0.00}_{-0.00}$	$0.5^{+0.5}_{-0.5}$
40 – 50	$0.993^{+0.001}_{-0.002}$	$0.997^{+0.001}_{-0.001}$	$0.996^{+0.001}_{-0.002}$	$0.996^{+0.002}_{-0.002}$	$1.000^{+0.000}_{-0.003}$	$1.00^{+0.00}_{-0.01}$	$1.00^{+0.00}_{-0.01}$	$0.99^{+0.01}_{-0.02}$	$0.97^{+0.02}_{-0.04}$	$0.5^{+0.5}_{-0.5}$
50 – 60	$0.992^{+0.003}_{-0.005}$	$0.991^{+0.004}_{-0.005}$	$0.991^{+0.005}_{-0.007}$	$0.995^{+0.004}_{-0.007}$	$0.95^{+0.02}_{-0.03}$	$1.00^{+0.00}_{-0.04}$	$1.00^{+0.00}_{-0.04}$	$1.00^{+0.00}_{-0.06}$	$0.5^{+0.2}_{-0.2}$	$0.5^{+0.5}_{-0.5}$
60 – 70	$0.99^{+0.01}_{-0.01}$	$0.99^{+0.01}_{-0.01}$	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.03}$	$0.96^{+0.03}_{-0.05}$	$1.00^{+0.00}_{-0.09}$	$1.0^{+0.0}_{-0.2}$	$1.0^{+0.0}_{-0.3}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
70 – 80	$0.98^{+0.02}_{-0.03}$	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.07}$	$0.8^{+0.1}_{-0.2}$	$1.0^{+0.0}_{-0.1}$	$1.0^{+0.0}_{-0.3}$	$1.0^{+0.0}_{-0.3}$	$1.0^{+0.0}_{-0.4}$	$0.5^{+0.5}_{-0.5}$
80 – 90	$1.00^{+0.00}_{-0.04}$	$0.96^{+0.03}_{-0.03}$	$1.00^{+0.00}_{-0.06}$	$1.00^{+0.00}_{-0.09}$	$1.0^{+0.0}_{-0.2}$	$0.5^{+0.5}_{-0.5}$	$1.0^{+0.0}_{-0.3}$	$1.0^{+0.0}_{-0.3}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
90 – 100	$1.00^{+0.00}_{-0.06}$	$1.00^{+0.00}_{-0.06}$	$1.0^{+0.0}_{-0.1}$	$1.0^{+0.0}_{-0.2}$	$1.0^{+0.0}_{-0.4}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$

Table A.2: Clustering efficiency (Background subtracted)

p_t (GeV) η	0.0 – 0.3	0.3 – 0.6	0.6 – 0.9	0.9 – 1.2	1.2 – 1.5	1.5 – 1.8	1.8 – 2.1	2.1 – 2.4	2.4 – 2.7	2.7 – 3.0
0 – 10	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
10 – 20	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
20 – 30	$0.91^{+0.01}_{-0.01}$	$0.953^{+0.008}_{-0.009}$	$0.92^{+0.01}_{-0.01}$	$0.88^{+0.01}_{-0.01}$	$0.85^{+0.02}_{-0.02}$	$0.87^{+0.03}_{-0.03}$	$0.93^{+0.02}_{-0.03}$	$0.97^{+0.01}_{-0.02}$	$1.00^{+0.00}_{-0.07}$	$0.5^{+0.5}_{-0.5}$
30 – 40	$0.986^{+0.002}_{-0.003}$	$0.985^{+0.002}_{-0.003}$	$0.974^{+0.003}_{-0.004}$	$0.983^{+0.003}_{-0.004}$	$0.984^{+0.004}_{-0.005}$	$0.98^{+0.01}_{-0.01}$	$0.96^{+0.01}_{-0.02}$	$1.00^{+0.00}_{-0.01}$	$0.98^{+0}_{-0.08}$	$0.5^{+0.5}_{-0.5}$
40 – 50	$0.993^{+0.001}_{-0.002}$	$0.994^{+0.001}_{-0.002}$	$0.996^{+0.001}_{-0.002}$	$0.991^{+0.002}_{-0.002}$	$1.000^{+0.000}_{-0.003}$	$0.99^{+0.01}_{-0.01}$	$1.000^{+0.000}_{-0.009}$	$1.00^{+0.00}_{-0.01}$	$0.97^{+0.02}_{-0.05}$	$0.5^{+0.5}_{-0.5}$
50 – 60	$0.980^{+0.006}_{-0.007}$	$0.986^{+0.005}_{-0.007}$	$0.995^{+0.002}_{-0.005}$	$0.999^{+0.000}_{-0.006}$	$0.93^{+0.02}_{-0.03}$	$1.00^{+0.00}_{-0.04}$	$1.00^{+0.00}_{-0.05}$	$0.96^{+0.04}_{-0.07}$	$0.6^{+0.2}_{-0.2}$	$0.5^{+0.5}_{-0.5}$
60 – 70	$0.99^{+0.01}_{-0.01}$	$1.000^{+0.000}_{-0.009}$	$0.94^{+0.02}_{-0.03}$	$0.99^{+0.01}_{-0.03}$	$0.99^{+0.00}_{-0.05}$	$0.9^{+0.1}_{-0.1}$	$0.9^{+0.0}_{-0.2}$	$1.0^{+0.0}_{-0.4}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
70 – 80	$0.94^{+0.03}_{-0.04}$	$0.96^{+0.02}_{-0.04}$	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.09}$	$0.7^{+0.1}_{-0.2}$	$1.0^{+0.0}_{-0.2}$	$1.0^{+0.0}_{-0.3}$	$1.0^{+0.0}_{-0.4}$	$1.0^{+0.0}_{-0.4}$	$0.5^{+0.5}_{-0.5}$
80 – 90	$0.98^{+0.02}_{-0.04}$	$0.98^{+0.02}_{-0.05}$	$1.00^{+0.00}_{-0.07}$	$1.0^{+0.0}_{-0.1}$	$1.0^{+0.0}_{-0.2}$	$0.5^{+0.5}_{-0.5}$	$1.0^{+0.0}_{-0.4}$	$1.0^{+0.0}_{-0.3}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
90 – 100	$1.00^{+0.00}_{-0.07}$	$1.00^{+0.00}_{-0.06}$	$0.9^{+0.1}_{-0.1}$	$1.0^{+0.0}_{-0.2}$	$1.0^{+0.0}_{-0.4}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$

Table A.3: GSF electron reconstruction efficiency ($Z^0 \rightarrow e^+e^-$ only)

E_t (GeV) η	0.0 - 0.3	0.3 - 0.6	0.6 - 0.9	0.9 - 1.2	1.2 - 1.5	1.5 - 1.8	1.8 - 2.1	2.1 - 2.4	2.4 - 2.7	2.7 - 3.0
0 - 10	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
10 - 20	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
20 - 30	$0.957^{+0.005}_{-0.005}$	$0.956^{+0.006}_{-0.006}$	$0.952^{+0.007}_{-0.007}$	$0.93^{+0.01}_{-0.01}$	$0.92^{+0.01}_{-0.01}$	$0.92^{+0.01}_{-0.01}$	$0.931^{+0.007}_{-0.008}$	$0.947^{+0.006}_{-0.007}$	$0.84^{+0.02}_{-0.02}$	$0.5^{+0.5}_{-0.5}$
30 - 40	$0.976^{+0.002}_{-0.003}$	$0.982^{+0.002}_{-0.002}$	$0.978^{+0.002}_{-0.002}$	$0.976^{+0.003}_{-0.003}$	$0.965^{+0.005}_{-0.005}$	$0.962^{+0.005}_{-0.005}$	$0.958^{+0.004}_{-0.004}$	$0.962^{+0.004}_{-0.004}$	$0.84^{+0.01}_{-0.01}$	$0.5^{+0.5}_{-0.5}$
40 - 50	$0.983^{+0.002}_{-0.002}$	$0.988^{+0.002}_{-0.002}$	$0.986^{+0.002}_{-0.002}$	$0.978^{+0.003}_{-0.003}$	$0.972^{+0.005}_{-0.005}$	$0.967^{+0.005}_{-0.005}$	$0.971^{+0.003}_{-0.003}$	$0.967^{+0.003}_{-0.004}$	$0.86^{+0.01}_{-0.01}$	$0.5^{+0.5}_{-0.5}$
50 - 60	$0.985^{+0.003}_{-0.004}$	$0.984^{+0.004}_{-0.004}$	$0.995^{+0.002}_{-0.003}$	$0.983^{+0.005}_{-0.007}$	$0.98^{+0.01}_{-0.01}$	$0.980^{+0.007}_{-0.009}$	$0.969^{+0.007}_{-0.008}$	$0.976^{+0.006}_{-0.007}$	$0.91^{+0.02}_{-0.02}$	$0.5^{+0.5}_{-0.5}$
60 - 70	$0.97^{+0.01}_{-0.01}$	$0.993^{+0.004}_{-0.007}$	$0.992^{+0.005}_{-0.007}$	$0.98^{+0.01}_{-0.01}$	$0.95^{+0.02}_{-0.03}$	$0.98^{+0.01}_{-0.02}$	$0.96^{+0.01}_{-0.02}$	$0.977^{+0.01}_{-0.02}$	$0.91^{+0.04}_{-0.05}$	$0.5^{+0.5}_{-0.5}$
70 - 80	$1.00^{+0.00}_{-0.01}$	$1.00^{+0.00}_{-0.01}$	$0.99^{+0.01}_{-0.02}$	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.04}$	$0.95^{+0.03}_{-0.04}$	$0.98^{+0.01}_{-0.03}$	$0.94^{+0.03}_{-0.04}$	$0.90^{+0.05}_{-0.08}$	$0.5^{+0.5}_{-0.5}$
80 - 90	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.02}$	$0.98^{+0.01}_{-0.03}$	$1.00^{+0.00}_{-0.05}$	$1.0^{+0.0}_{-0.1}$	$0.95^{+0.03}_{-0.06}$	$0.97^{+0.02}_{-0.04}$	$0.93^{+0.04}_{-0.06}$	$0.9^{+0.1}_{-0.1}$	$0.5^{+0.5}_{-0.5}$
90 - 100	$0.91^{+0.04}_{-0.06}$	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.06}$	$1.00^{+0.00}_{-0.06}$	$1.0^{+0.0}_{-0.1}$	$1.0^{+0.0}_{-0.1}$	$1.00^{+0.00}_{-0.09}$	$1.0^{+0.0}_{-0.1}$	$1.0^{+0.0}_{-0.4}$	$0.5^{+0.5}_{-0.5}$

Table A.4: GSF electron reconstruction efficiency (Background subtracted)

E_t (GeV) η	0.0 - 0.3	0.3 - 0.6	0.6 - 0.9	0.9 - 1.2	1.2 - 1.5	1.5 - 1.8	1.8 - 2.1	2.1 - 2.4	2.4 - 2.7	2.7 - 3.0
0 - 10	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
10 - 20	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
20 - 30	$0.936^{+0.006}_{-0.006}$	$0.929^{+0.007}_{-0.007}$	$0.945^{+0.007}_{-0.008}$	$0.92^{+0.01}_{-0.01}$	$0.87^{+0.02}_{-0.02}$	$0.91^{+0.01}_{-0.01}$	$0.901^{+0.008}_{-0.009}$	$0.901^{+0.008}_{-0.008}$	$0.78^{+0.02}_{-0.02}$	$0.5^{+0.5}_{-0.5}$
30 - 40	$0.966^{+0.003}_{-0.003}$	$0.973^{+0.003}_{-0.003}$	$0.973^{+0.003}_{-0.003}$	$0.967^{+0.003}_{-0.004}$	$0.951^{+0.006}_{-0.006}$	$0.956^{+0.005}_{-0.005}$	$0.950^{+0.004}_{-0.004}$	$0.957^{+0.004}_{-0.004}$	$0.83^{+0.01}_{-0.01}$	$0.5^{+0.5}_{-0.5}$
40 - 50	$0.983^{+0.002}_{-0.002}$	$0.986^{+0.002}_{-0.002}$	$0.984^{+0.002}_{-0.002}$	$0.976^{+0.003}_{-0.003}$	$0.970^{+0.005}_{-0.006}$	$0.965^{+0.005}_{-0.005}$	$0.970^{+0.003}_{-0.003}$	$0.968^{+0.003}_{-0.004}$	$0.86^{+0.01}_{-0.01}$	$0.5^{+0.5}_{-0.5}$
50 - 60	$0.980^{+0.004}_{-0.005}$	$0.975^{+0.005}_{-0.006}$	$0.990^{+0.003}_{-0.004}$	$0.986^{+0.005}_{-0.006}$	$0.99^{+0.01}_{-0.01}$	$0.984^{+0.006}_{-0.009}$	$0.965^{+0.008}_{-0.009}$	$0.970^{+0.007}_{-0.008}$	$0.91^{+0.02}_{-0.02}$	$0.5^{+0.5}_{-0.5}$
60 - 70	$0.96^{+0.01}_{-0.01}$	$1.00^{+0.00}_{-0.04}$	$0.987^{+0.006}_{-0.009}$	$0.98^{+0.01}_{-0.02}$	$0.95^{+0.02}_{-0.03}$	$0.96^{+0.02}_{-0.02}$	$0.92^{+0.02}_{-0.02}$	$0.98^{+0.01}_{-0.02}$	$0.92^{+0.04}_{-0.05}$	$0.5^{+0.5}_{-0.5}$
70 - 80	$1.0^{+0.0}_{-0.01}$	$1.00^{+0.00}_{-0.01}$	$0.97^{+0.01}_{-0.02}$	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.05}$	$0.93^{+0.03}_{-0.05}$	$0.97^{+0.02}_{-0.03}$	$0.88^{+0.04}_{-0.05}$	$0.91^{+0.06}_{-0.08}$	$0.5^{+0.5}_{-0.5}$
80 - 90	$1^{+0}_{-0.02}$	$1^{+0}_{-0.02}$	$0.95^{+0.02}_{-0.04}$	$1.00^{+0.00}_{-0.05}$	$1.0^{+0.0}_{-0.1}$	$0.96^{+0.04}_{-0.07}$	$0.97^{+0.02}_{-0.05}$	$0.92^{+0.02}_{-0.06}$	$0.94^{+0.05}_{-0.09}$	$0.5^{+0.5}_{-0.5}$
90 - 100	$0.90^{+0.04}_{-0.06}$	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.06}$	$0.92^{+0.06}_{-0.09}$	$1.0^{+0.0}_{-0.1}$	$1.0^{+0.0}_{-0.1}$	$1.0^{+0.0}_{-0.1}$	$1.0^{+0.0}_{-0.1}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$

Table A.5: Electron ID efficiency ($Z^0 \rightarrow e^+e^-$ only)

E_t (GeV) η	0.0 – 0.3	0.3 – 0.6	0.6 – 0.9	0.9 – 1.2	1.2 – 1.5	1.5 – 1.8	1.8 – 2.1	2.1 – 2.4	2.4 – 2.7	2.7 – 3.0
0 – 10	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
10 – 20	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
20 – 30	$0.63^{+0.01}_{-0.01}$	$0.62^{+0.01}_{-0.01}$	$0.59^{+0.01}_{-0.01}$	$0.70^{+0.02}_{-0.02}$	$0.67^{+0.02}_{-0.02}$	$0.65^{+0.02}_{-0.02}$	$0.57^{+0.01}_{-0.01}$	$0.55^{+0.01}_{-0.01}$	$0.50^{+0.02}_{-0.02}$	$0.5^{+0.5}_{-0.5}$
30 – 40	$0.986^{+0.002}_{-0.002}$	$0.979^{+0.002}_{-0.002}$	$0.983^{+0.002}_{-0.002}$	$0.987^{+0.002}_{-0.002}$	$0.973^{+0.004}_{-0.005}$	$0.969^{+0.004}_{-0.005}$	$0.959^{+0.004}_{-0.004}$	$0.929^{+0.005}_{-0.005}$	$0.91^{+0.01}_{-0.01}$	$0.5^{+0.5}_{-0.5}$
40 – 50	$0.986^{+0.002}_{-0.002}$	$0.983^{+0.002}_{-0.002}$	$0.984^{+0.002}_{-0.002}$	$0.987^{+0.002}_{-0.002}$	$0.979^{+0.004}_{-0.005}$	$0.978^{+0.004}_{-0.004}$	$0.968^{+0.003}_{-0.004}$	$0.943^{+0.004}_{-0.004}$	$0.947^{+0.008}_{-0.008}$	$0.5^{+0.5}_{-0.5}$
50 – 60	$0.987^{+0.003}_{-0.004}$	$0.984^{+0.004}_{-0.004}$	$0.984^{+0.004}_{-0.005}$	$0.994^{+0.003}_{-0.004}$	$0.98^{+0.01}_{-0.01}$	$0.97^{+0.01}_{-0.01}$	$0.983^{+0.005}_{-0.006}$	$0.962^{+0.008}_{-0.009}$	$0.96^{+0.01}_{-0.02}$	$0.5^{+0.5}_{-0.5}$
60 – 70	$0.986^{+0.005}_{-0.007}$	$0.985^{+0.006}_{-0.008}$	$0.988^{+0.005}_{-0.007}$	$0.98^{+0.01}_{-0.01}$	$0.99^{+0.01}_{-0.02}$	$0.97^{+0.01}_{-0.02}$	$0.96^{+0.01}_{-0.02}$	$0.96^{+0.01}_{-0.02}$	$1.00^{+0.00}_{-0.02}$	$0.5^{+0.5}_{-0.5}$
70 – 80	$1.000^{+0.000}_{-0.008}$	$1.000^{+0.000}_{-0.008}$	$0.99^{+0.01}_{-0.01}$	$0.98^{+0.01}_{-0.02}$	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.02}$	$0.98^{+0.01}_{-0.02}$	$0.97^{+0.02}_{-0.02}$	$0.95^{+0.03}_{-0.06}$	$0.5^{+0.5}_{-0.5}$
80 – 90	$1.00^{+0.00}_{-0.01}$	$0.99^{+0.01}_{-0.02}$	$0.98^{+0.01}_{-0.02}$	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.04}$	$0.94^{+0.03}_{-0.05}$	$0.98^{+0.02}_{-0.03}$	$0.91^{+0.04}_{-0.06}$	$0.9^{+0.1}_{-0.1}$	$0.5^{+0.5}_{-0.5}$
90 – 100	$1.00^{+0.00}_{-0.02}$	$0.98^{+0.01}_{-0.02}$	$0.93^{+0.03}_{-0.04}$	$0.97^{+0.02}_{-0.04}$	$1.00^{+0.00}_{-0.07}$	$1.00^{+0.00}_{-0.06}$	$1.00^{+0.00}_{-0.04}$	$1.00^{+0.00}_{-0.05}$	$1.0^{+0.0}_{-0.2}$	$0.5^{+0.5}_{-0.5}$

Table A.6: Electron ID efficiency (Background subtracted)

E_t (GeV) η	0.0 – 0.3	0.3 – 0.6	0.6 – 0.9	0.9 – 1.2	1.2 – 1.5	1.5 – 1.8	1.8 – 2.1	2.1 – 2.4	2.4 – 2.7	2.7 – 3.0
0 – 10	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
10 – 20	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
20 – 30	$0.67^{+0.01}_{-0.01}$	$0.67^{+0.01}_{-0.01}$	$0.67^{+0.02}_{-0.02}$	$0.82^{+0.02}_{-0.02}$	$0.94^{+0.01}_{-0.02}$	$0.72^{+0.02}_{-0.02}$	$0.59^{+0.01}_{-0.01}$	$0.56^{+0.01}_{-0.01}$	$0.51^{+0.02}_{-0.02}$	$0.5^{+0.5}_{-0.5}$
30 – 40	$0.983^{+0.002}_{-0.002}$	$0.980^{+0.002}_{-0.002}$	$0.984^{+0.002}_{-0.002}$	$0.985^{+0.002}_{-0.002}$	$0.961^{+0.005}_{-0.005}$	$0.967^{+0.004}_{-0.005}$	$0.957^{+0.004}_{-0.004}$	$0.926^{+0.005}_{-0.005}$	$0.91^{+0.01}_{-0.01}$	$0.5^{+0.5}_{-0.5}$
40 – 50	$0.985^{+0.002}_{-0.002}$	$0.980^{+0.002}_{-0.002}$	$0.986^{+0.002}_{-0.002}$	$0.984^{+0.002}_{-0.002}$	$0.977^{+0.004}_{-0.004}$	$0.970^{+0.004}_{-0.005}$	$0.967^{+0.003}_{-0.004}$	$0.942^{+0.004}_{-0.005}$	$0.944^{+0.008}_{-0.009}$	$0.5^{+0.5}_{-0.5}$
50 – 60	$0.981^{+0.004}_{-0.004}$	$0.973^{+0.005}_{-0.005}$	$0.979^{+0.004}_{-0.005}$	$0.997^{+0.003}_{-0.003}$	$0.97^{+0.01}_{-0.01}$	$0.96^{+0.01}_{-0.01}$	$0.983^{+0.005}_{-0.005}$	$0.95^{+0.01}_{-0.01}$	$0.93^{+0.02}_{-0.02}$	$0.5^{+0.5}_{-0.5}$
60 – 70	$0.976^{+0.007}_{-0.009}$	$0.983^{+0.006}_{-0.008}$	$0.998^{+0.003}_{-0.003}$	$0.97^{+0.01}_{-0.01}$	$1.00^{+0.00}_{-0.01}$	$0.96^{+0.02}_{-0.02}$	$0.93^{+0.01}_{-0.01}$	$0.96^{+0.01}_{-0.01}$	$0.99^{+0.01}_{-0.02}$	$0.5^{+0.5}_{-0.5}$
70 – 80	$0.98^{+0.01}_{-0.01}$	$0.98^{+0.01}_{-0.01}$	$1.00^{+0.00}_{-0.01}$	$0.94^{+0.02}_{-0.03}$	$0.93^{+0.03}_{-0.04}$	$0.96^{+0.03}_{-0.04}$	$0.97^{+0.02}_{-0.02}$	$0.96^{+0.02}_{-0.02}$	$0.99^{+0.01}_{-0.05}$	$0.5^{+0.5}_{-0.5}$
80 – 90	$1.00^{+0.00}_{-0.01}$	$0.98^{+0.02}_{-0.02}$	$0.96^{+0.02}_{-0.02}$	$0.96^{+0.03}_{-0.03}$	$0.93^{+0.04}_{-0.06}$	$0.89^{+0.04}_{-0.06}$	$0.99^{+0.01}_{-0.02}$	$0.86^{+0.05}_{-0.06}$	$0.90^{+0.06}_{-0.09}$	$0.5^{+0.5}_{-0.5}$
90 – 100	$1.00^{+0.00}_{-0.02}$	$0.96^{+0.02}_{-0.03}$	$0.90^{+0.04}_{-0.05}$	$0.95^{+0.03}_{-0.05}$	$0.8^{+0.1}_{-0.1}$	$1.00^{+0.00}_{-0.06}$	$0.99^{+0.01}_{-0.05}$	$1.00^{+0.00}_{-0.04}$	$1.0^{+0.0}_{-0.3}$	$0.5^{+0.5}_{-0.5}$

Table A.7: Trigger efficiency ($Z^0 \rightarrow e^+e^-$ only)

	0.0 - 0.3	0.3 - 0.6	0.6 - 0.9	0.9 - 1.2	1.2 - 1.5	1.5 - 1.8	1.8 - 2.1	2.1 - 2.4	2.4 - 2.7	2.7 - 3.0
0 - 10	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
10 - 20	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
20 - 30	$0.980^{+0.004}_{-0.005}$	$0.957^{+0.007}_{-0.007}$	$0.963^{+0.007}_{-0.008}$	$0.94^{+0.01}_{-0.01}$	$0.97^{+0.01}_{-0.01}$	$0.973^{+0.007}_{-0.009}$	$0.981^{+0.005}_{-0.006}$	$0.984^{+0.004}_{-0.005}$	$0.990^{+0.005}_{-0.009}$	$0.990^{+0.005}_{-0.009}$
30 - 40	$0.997^{+0.001}_{-0.001}$	$0.995^{+0.001}_{-0.001}$	$0.996^{+0.001}_{-0.001}$	$0.982^{+0.002}_{-0.003}$	$1.000^{+0.000}_{-0.008}$	$0.998^{+0.001}_{-0.002}$	$1.000^{+0.000}_{-0.005}$	$1.000^{+0.000}_{-0.006}$	$1.000^{+0.000}_{-0.002}$	$1.000^{+0.000}_{-0.002}$
40 - 50	$0.9992^{+0.0003}_{-0.0005}$	$0.9993^{+0.0003}_{-0.0005}$	$0.9992^{+0.0004}_{-0.0005}$	$0.996^{+0.001}_{-0.001}$	$0.999^{+0.001}_{-0.001}$	$0.999^{+0.001}_{-0.001}$	$0.9996^{+0.0003}_{-0.0006}$	$0.9996^{+0.0003}_{-0.0006}$	$1.000^{+0.000}_{-0.002}$	$1.000^{+0.000}_{-0.002}$
50 - 60	$0.999^{+0.001}_{-0.001}$	$0.999^{+0.001}_{-0.001}$	$0.999^{+0.001}_{-0.001}$	$1.000^{+0.000}_{-0.003}$	$1.000^{+0.000}_{-0.005}$	$1.000^{+0.000}_{-0.004}$	$1.000^{+0.000}_{-0.002}$	$1.000^{+0.000}_{-0.002}$	$1.000^{+0.000}_{-0.008}$	$1.000^{+0.000}_{-0.008}$
60 - 70	$1.000^{+0.000}_{-0.004}$	$1.000^{+0.000}_{-0.004}$	$1.000^{+0.000}_{-0.005}$	$1.000^{+0.000}_{-0.008}$	$0.98^{+0.01}_{-0.02}$	$0.99^{+0.01}_{-0.01}$	$1.000^{+0.000}_{-0.009}$	$1.000^{+0.000}_{-0.008}$	$1.000^{+0.000}_{-0.02}$	$1.000^{+0.000}_{-0.02}$
70 - 80	$0.99^{+0.01}_{-0.01}$	$1.00^{+0.00}_{-0.01}$	$1.00^{+0.00}_{-0.01}$	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.04}$	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.05}$	$1.00^{+0.00}_{-0.05}$
80 - 90	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.04}$	$1.00^{+0.00}_{-0.08}$	$1.00^{+0.00}_{-0.07}$	$1.00^{+0.00}_{-0.04}$	$1.00^{+0.00}_{-0.05}$	$1.0^{+0.0}_{-0.1}$	$1.0^{+0.0}_{-0.1}$
90 - 100	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.05}$	$1.00^{+0.00}_{-0.07}$	$1.0^{+0.0}_{-0.1}$	$1.00^{+0.00}_{-0.09}$	$1.00^{+0.00}_{-0.07}$	$1.0^{+0.0}_{-0.1}$	$1.0^{+0.0}_{-0.3}$	$1.0^{+0.0}_{-0.3}$

Table A.8: Trigger efficiency (Background subtracted)

	0.0 - 0.3	0.3 - 0.6	0.6 - 0.9	0.9 - 1.2	1.2 - 1.5	1.5 - 1.8	1.8 - 2.1	2.1 - 2.4	2.4 - 2.7	2.7 - 3.0
0 - 10	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
10 - 20	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
20 - 30	$0.981^{+0.004}_{-0.005}$	$0.958^{+0.007}_{-0.007}$	$0.964^{+0.007}_{-0.008}$	$0.94^{+0.01}_{-0.01}$	$0.98^{+0.01}_{-0.01}$	$0.973^{+0.008}_{-0.009}$	$0.982^{+0.006}_{-0.006}$	$0.983^{+0.004}_{-0.005}$	$0.990^{+0.005}_{-0.009}$	$0.990^{+0.005}_{-0.009}$
30 - 40	$0.997^{+0.001}_{-0.001}$	$0.995^{+0.001}_{-0.001}$	$0.996^{+0.001}_{-0.001}$	$0.983^{+0.002}_{-0.003}$	$1.000^{+0.000}_{-0.009}$	$0.997^{+0.001}_{-0.002}$	$1.000^{+0.000}_{-0.005}$	$1.000^{+0.000}_{-0.006}$	$1.000^{+0.000}_{-0.002}$	$1.000^{+0.000}_{-0.002}$
40 - 50	$0.9993^{+0.0003}_{-0.0005}$	$0.9993^{+0.0003}_{-0.0005}$	$0.9992^{+0.0004}_{-0.0005}$	$0.996^{+0.001}_{-0.001}$	$0.999^{+0.001}_{-0.001}$	$0.999^{+0.001}_{-0.001}$	$0.9996^{+0.0003}_{-0.0006}$	$0.9996^{+0.0003}_{-0.0006}$	$1.000^{+0.000}_{-0.002}$	$1.000^{+0.000}_{-0.002}$
50 - 60	$0.999^{+0.001}_{-0.001}$	$0.999^{+0.001}_{-0.001}$	$0.999^{+0.001}_{-0.002}$	$1.000^{+0.000}_{-0.003}$	$1.000^{+0.000}_{-0.005}$	$1.000^{+0.000}_{-0.004}$	$1.000^{+0.000}_{-0.002}$	$1.000^{+0.000}_{-0.002}$	$1.000^{+0.000}_{-0.008}$	$1.000^{+0.000}_{-0.008}$
60 - 70	$1.000^{+0.000}_{-0.004}$	$1.000^{+0.000}_{-0.004}$	$1.000^{+0.000}_{-0.005}$	$1.000^{+0.000}_{-0.008}$	$0.98^{+0.01}_{-0.02}$	$0.99^{+0.01}_{-0.02}$	$1.000^{+0.000}_{-0.009}$	$1.000^{+0.000}_{-0.009}$	$1.000^{+0.000}_{-0.03}$	$1.000^{+0.000}_{-0.03}$
70 - 80	$0.99^{+0.01}_{-0.01}$	$1.00^{+0.00}_{-0.01}$	$1.00^{+0.00}_{-0.01}$	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.04}$	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.05}$	$1.00^{+0.00}_{-0.05}$
80 - 90	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.05}$	$1.00^{+0.00}_{-0.08}$	$1.00^{+0.00}_{-0.07}$	$1.00^{+0.00}_{-0.04}$	$1.00^{+0.00}_{-0.05}$	$1.0^{+0.0}_{-0.1}$	$1.0^{+0.0}_{-0.1}$
90 - 100	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.06}$	$1.00^{+0.00}_{-0.07}$	$1.0^{+0.0}_{-0.1}$	$1.00^{+0.00}_{-0.1}$	$1.00^{+0.00}_{-0.07}$	$1.0^{+0.0}_{-0.1}$	$1.0^{+0.0}_{-0.3}$	$1.0^{+0.0}_{-0.3}$

Table A.9: Electron ID efficiency ($Z^0 \rightarrow e^+e^-$ only)

0-10	0.0-0.3	0.3-0.6	0.6-0.9	0.9-1.2	1.2-1.5	1.5-1.8	1.8-2.1	2.1-2.4	2.4-2.7	2.7-3.0
0-10	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
10-20	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
20-30	$0.63^{+0.01}_{-0.01}$	$0.62^{+0.01}_{-0.01}$	$0.60^{+0.01}_{-0.01}$	$0.70^{+0.02}_{-0.02}$	$0.67^{+0.02}_{-0.02}$	$0.65^{+0.02}_{-0.02}$	$0.57^{+0.01}_{-0.01}$	$0.55^{+0.01}_{-0.01}$	$0.50^{+0.02}_{-0.02}$	$0.5^{+0.5}_{-0.5}$
30-40	$0.986^{+0.002}_{-0.002}$	$0.979^{+0.002}_{-0.002}$	$0.983^{+0.002}_{-0.002}$	$0.987^{+0.002}_{-0.002}$	$0.973^{+0.004}_{-0.005}$	$0.969^{+0.004}_{-0.005}$	$0.959^{+0.004}_{-0.004}$	$0.929^{+0.005}_{-0.005}$	$0.91^{+0.01}_{-0.01}$	$0.5^{+0.5}_{-0.5}$
40-50	$0.986^{+0.002}_{-0.002}$	$0.983^{+0.002}_{-0.002}$	$0.984^{+0.002}_{-0.002}$	$0.987^{+0.002}_{-0.002}$	$0.980^{+0.004}_{-0.005}$	$0.978^{+0.004}_{-0.004}$	$0.968^{+0.003}_{-0.004}$	$0.943^{+0.004}_{-0.004}$	$0.947^{+0.008}_{-0.008}$	$0.5^{+0.5}_{-0.5}$
50-60	$0.987^{+0.003}_{-0.004}$	$0.984^{+0.004}_{-0.004}$	$0.984^{+0.005}_{-0.005}$	$0.994^{+0.003}_{-0.004}$	$0.98^{+0.01}_{-0.01}$	$0.97^{+0.01}_{-0.01}$	$0.983^{+0.005}_{-0.006}$	$0.962^{+0.008}_{-0.009}$	$0.96^{+0.01}_{-0.01}$	$0.5^{+0.5}_{-0.5}$
60-70	$0.986^{+0.005}_{-0.007}$	$0.985^{+0.006}_{-0.008}$	$0.988^{+0.005}_{-0.007}$	$0.98^{+0.01}_{-0.01}$	$0.99^{+0.01}_{-0.02}$	$0.97^{+0.01}_{-0.02}$	$0.96^{+0.01}_{-0.02}$	$0.96^{+0.01}_{-0.02}$	$1.00^{+0.00}_{-0.00}$	$0.5^{+0.5}_{-0.5}$
70-80	$1.00^{+0.000}_{-0.000}$	$1.000^{+0.000}_{-0.008}$	$0.99^{+0.01}_{-0.01}$	$0.98^{+0.01}_{-0.02}$	$1.00^{+0.00}_{-0.03}$	$1.00^{+0.00}_{-0.03}$	$0.98^{+0.01}_{-0.02}$	$0.97^{+0.02}_{-0.02}$	$0.95^{+0.03}_{-0.06}$	$0.5^{+0.5}_{-0.5}$
80-90	$1.00^{+0.01}_{-0.01}$	$0.99^{+0.01}_{-0.02}$	$0.98^{+0.01}_{-0.02}$	$1.00^{+0.00}_{-0.02}$	$1.00^{+0.00}_{-0.04}$	$0.94^{+0.03}_{-0.03}$	$0.98^{+0.02}_{-0.03}$	$0.91^{+0.04}_{-0.06}$	$0.88^{+0.07}_{-0.10}$	$0.5^{+0.5}_{-0.5}$
90-100	$1.00^{+0.00}_{-0.02}$	$0.98^{+0.01}_{-0.02}$	$0.93^{+0.03}_{-0.04}$	$0.97^{+0.02}_{-0.04}$	$1.00^{+0.00}_{-0.07}$	$1.00^{+0.00}_{-0.06}$	$1.00^{+0.00}_{-0.04}$	$1.00^{+0.00}_{-0.05}$	$1.0^{+0.0}_{-0.2}$	$0.5^{+0.5}_{-0.5}$

 Table A.10: Electron ID efficiency ($W^\pm \rightarrow e\nu$)

0-10	0.0-0.3	0.3-0.6	0.6-0.9	0.9-1.2	1.2-1.5	1.5-1.8	1.8-2.1	2.1-2.4	2.4-2.7	2.7-3.0
0-10	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
10-20	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$	$0.5^{+0.5}_{-0.5}$
20-30	$0.526^{+0.003}_{-0.003}$	$0.530^{+0.003}_{-0.003}$	$0.528^{+0.003}_{-0.003}$	$0.521^{+0.003}_{-0.003}$	$0.488^{+0.003}_{-0.003}$	$0.510^{+0.003}_{-0.003}$	$0.523^{+0.003}_{-0.003}$	$0.514^{+0.003}_{-0.003}$	$0.522^{+0.006}_{-0.006}$	$0.5^{+0.5}_{-0.5}$
30-40	$0.9841^{+0.0006}_{-0.0006}$	$0.9805^{+0.0007}_{-0.0007}$	$0.9800^{+0.0007}_{-0.0007}$	$0.9787^{+0.0008}_{-0.0008}$	$0.976^{+0.001}_{-0.001}$	$0.973^{+0.001}_{-0.001}$	$0.9619^{+0.0010}_{-0.0010}$	$0.932^{+0.001}_{-0.001}$	$0.923^{+0.002}_{-0.002}$	$0.5^{+0.5}_{-0.5}$
40-50	$0.9878^{+0.0009}_{-0.0009}$	$0.985^{+0.001}_{-0.001}$	$0.982^{+0.001}_{-0.001}$	$0.980^{+0.001}_{-0.001}$	$0.982^{+0.002}_{-0.002}$	$0.974^{+0.002}_{-0.002}$	$0.966^{+0.002}_{-0.002}$	$0.942^{+0.002}_{-0.002}$	$0.928^{+0.004}_{-0.004}$	$0.5^{+0.5}_{-0.5}$
50-60	$0.988^{+0.002}_{-0.002}$	$0.983^{+0.002}_{-0.002}$	$0.977^{+0.002}_{-0.003}$	$0.979^{+0.003}_{-0.003}$	$0.975^{+0.004}_{-0.004}$	$0.976^{+0.003}_{-0.003}$	$0.972^{+0.003}_{-0.003}$	$0.946^{+0.004}_{-0.004}$	$0.943^{+0.007}_{-0.008}$	$0.5^{+0.5}_{-0.5}$
60-70	$0.984^{+0.003}_{-0.003}$	$0.982^{+0.003}_{-0.004}$	$0.985^{+0.003}_{-0.004}$	$0.981^{+0.004}_{-0.004}$	$0.986^{+0.004}_{-0.005}$	$0.977^{+0.005}_{-0.005}$	$0.962^{+0.005}_{-0.005}$	$0.949^{+0.006}_{-0.006}$	$0.92^{+0.01}_{-0.01}$	$0.5^{+0.5}_{-0.5}$
70-80	$0.981^{+0.005}_{-0.006}$	$0.980^{+0.005}_{-0.005}$	$0.982^{+0.005}_{-0.006}$	$0.978^{+0.006}_{-0.006}$	$0.973^{+0.007}_{-0.009}$	$0.983^{+0.006}_{-0.006}$	$0.957^{+0.008}_{-0.009}$	$0.971^{+0.006}_{-0.007}$	$0.95^{+0.01}_{-0.02}$	$0.5^{+0.5}_{-0.5}$
80-90	$0.987^{+0.005}_{-0.007}$	$0.986^{+0.005}_{-0.007}$	$0.982^{+0.006}_{-0.008}$	$0.981^{+0.006}_{-0.008}$	$0.96^{+0.01}_{-0.01}$	$0.96^{+0.01}_{-0.01}$	$0.96^{+0.01}_{-0.01}$	$0.96^{+0.01}_{-0.01}$	$0.95^{+0.02}_{-0.02}$	$0.5^{+0.5}_{-0.5}$
90-100	$0.991^{+0.005}_{-0.008}$	$0.987^{+0.006}_{-0.009}$	$0.996^{+0.003}_{-0.006}$	$0.98^{+0.01}_{-0.01}$	$0.98^{+0.01}_{-0.02}$	$0.97^{+0.01}_{-0.02}$	$0.98^{+0.01}_{-0.01}$	$0.94^{+0.01}_{-0.02}$	$0.94^{+0.03}_{-0.04}$	$0.5^{+0.5}_{-0.5}$

B. $N - 1$ efficiency measurements

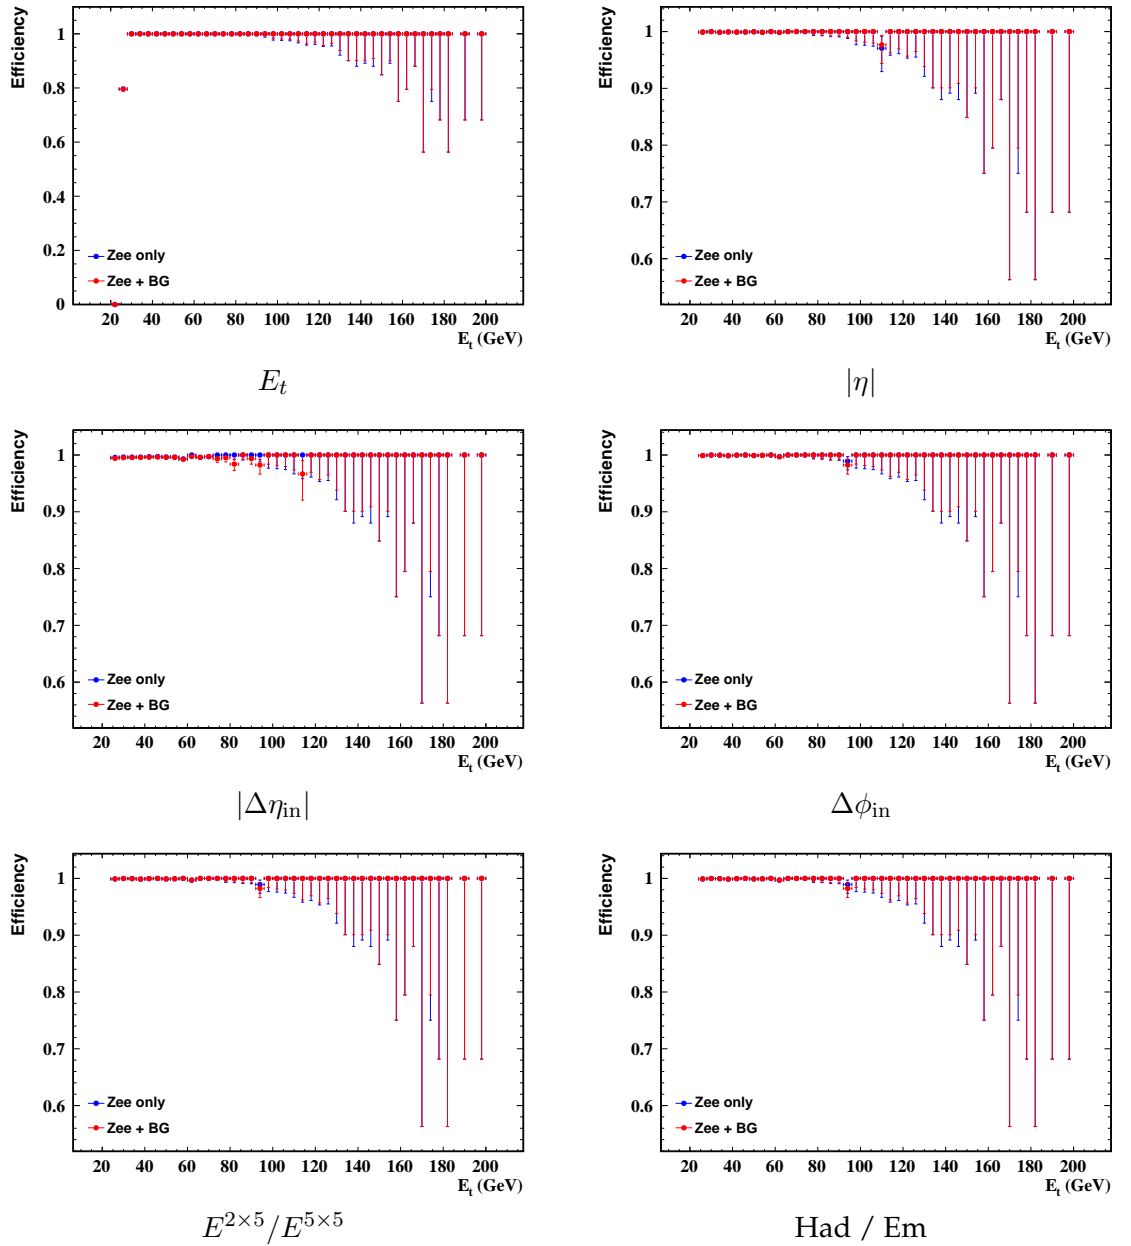


Figure B.1: Barrel cuts $N - 1$ performance

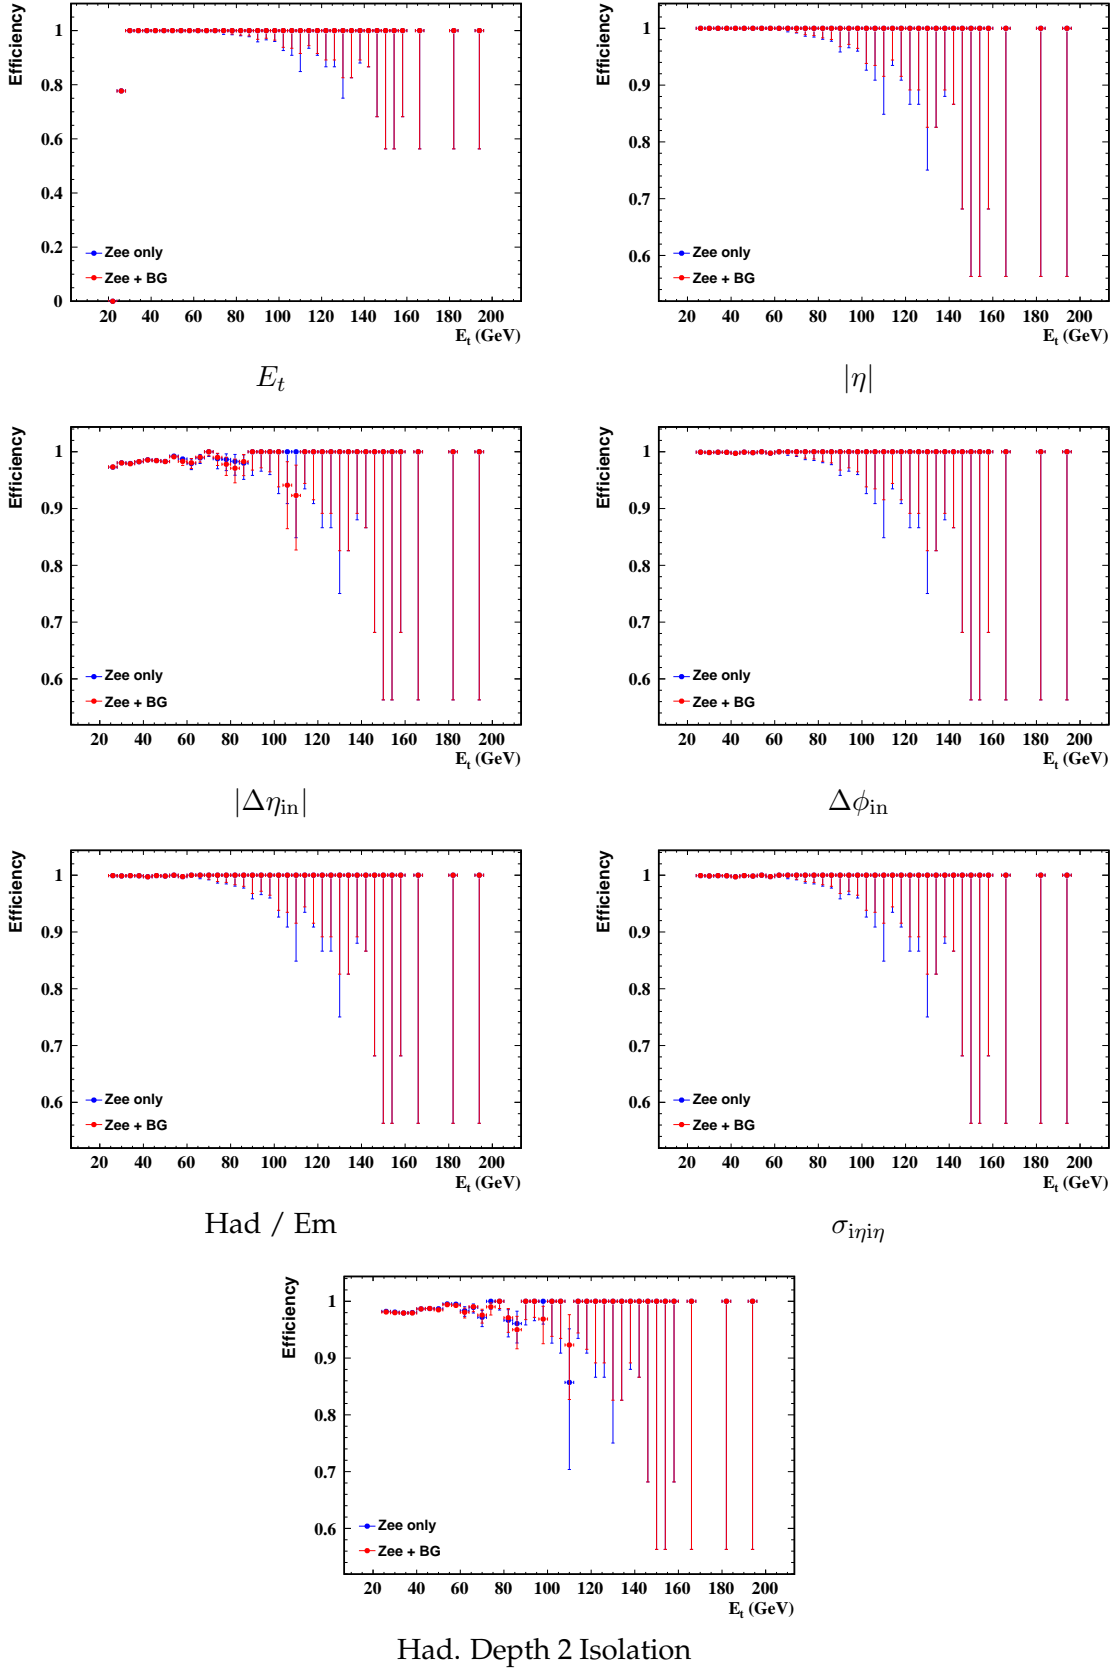


Figure B.2: Endcap cuts $N - 1$ performance

Glossary

ADC: Analog to Digital Converter
ALICE: A Large Ion Collider Experiment
APD: Avalanche Photo Diode
ASIC: Application Specific Integrated Circuit
ATLAS: A Toroidal LHC ApparatuS
BXID: Bunch crossing identification
CAF: CMS Analysis Facility
CASTOR: CERN Advanced STORAge Manager
CMS: Compact Muon Solenoid
CSC: Cathode Strip Chamber
CUPV: CASTOR User Privilege Validation
DAQ: Data acquisition
DT: Drift Tube
EB: ECAL Barrel
ECAL: Electromagnetic calorimeter
EDM: Event Data Model
EE: ECAL Endcap
EG: Electron / Gamma
EM: Electromagnetic
EWK: Electroweak
FE: Front End (ECAL electronics)
FED: Front-End Driver
FENIX: Front End Intermediate data eXtractor
FGVB: Fine Grain Veto Bit
FTP: File Transfer Protocol
FTS: File Transfer Service
FU: Filter Unit
GCT: Global Calorimeter Trigger
GMT: Global Muon Trigger
GSF: Gaussian Sum Filter
GT: Global Trigger
HB: HCAL Barrel

HCAL: Hadronic calorimeter
HE: HCAL Endcap
HEEP: High Energy Electron Pairs
HF: HCAL Forward
HLT: High Level Trigger
HO: HCAL Outer
HPD: Hybrid Photo Diode
HSM: Hierarchical Storage Manager
IOV: Interval Of Validity
KK: Kaluza-Klein
LHC: Large Hadron Collider
LHCb: Large Hadron Collider Beauty
LL: Leading Logarithm
LO: Leading Order
LVDS: Low Voltage Digital Signalling
MGPA: Multi Gain Pre-Amplifier
MSSM: Minimally Supersymmetric Standard Model
NLL: Next-to-Leading Logarithm
NLO: Next-to-Leading Order
OPN: Optical Private Network
PDF: Parton Density Function
PS: Proton Synchrotron
QCD: Quantum chromodynamics
QED: Quantum electrodynamics
RAID: Redundant Array of Independent Disks
RCT: Regional Calorimeter Trigger
RF: Radio Frequency
RFIO: Remote File I/O
RPC: Resistive Plate Chamber
RS: Randall-Sundrum
RTCP: Remote tape copy
SM: Standard Model
SPS: Super Proton Synchrotron
SRM: Storage Resource Manager
SUSY: Supersymmetry
TEC: Tracker Endcap
TIB: Tracker Inner Barrel
TID: Tracker Inner Disk
TOB: Tracker Outer Barrel
TP: Trigger Primitive
TPG: Trigger Primitive Generator

GLOSSARY

VDQM: Volume Device Queue Manager

VEV: Vacuum Expectation Value

VFE: Very Front End (ECAL electronics)

VMGR: Volume manager

VPT: Vacuum Photo Triode

Bibliography

- [1] A. Zee. *Quantum Field Theory in a Nutshell*. Princeton University Press, 2003.
- [2] M. Srednicki. *Quantum Field Theory*. Cambridge University Press, 2007.
- [3] K. A. Olive. Introduction to Supersymmetry: Astrophysical and Phenomenological Constraints. *Les Houches Summer School*, 1999.
- [4] J. Ellis, G. Ridolfi, and F. Zwirner. Higgs Boson Properties in the Standard Model and its Supersymmetric Extensions. *Comptes Rendus Physique*, 8:999–1012, 2007.
- [5] J. Ellis, S. Kelley, and D. V. Nanopoulos. Precision LEP data, supersymmetric GUTs and string unification. *Phys. Lett. B*, 249:441–448, 1990.
- [6] P. Langacker and M. Luo. Implications of precision electroweak experiments for m_t , ρ_0 , $\sin^2 \theta_W$, and grand unification. *Phys. Rev. D*, 44:817–822, 1991.
- [7] N. Arkani-Hamed, S. Dimopoulos, and G. Dvali. The hierarchy problem and new dimensions at a millimeter. *Phys. Lett. B*, 429:263–272, 1998.
- [8] B. C. Allanach, C. Grojean, P. Skands, et al. Les Houches “Physics at TeV Colliders 2005” Beyond The Standard Model Working Group: Summary Report. *hep-ph/0602198*, 2005.
- [9] K. Cheung and G. Landsberg. Drell-Yan and diphoton production at hadron colliders and low scale gravity model. *Phys. Rev. D*, 62:076003, 2000.
- [10] K. R. Dienes, E. Dudas, and T. Gherghetta. Grand unification at intermediate mass scales through extra dimensions. *Nucl. Phys. B*, 537:47–108, 1999.
- [11] K. R. Dienes, E. Dudas, and T. Gherghetta. Extra spacetime dimensions and unification. *Phys. Lett. B*, 436:55–65, 1998.
- [12] K. Cheung and G. Landsberg. Kaluza-Klein states of the standard model gauge bosons: Constraints from high energy experiments. *Phys. Rev. D*, 65:076003, 2002.
- [13] L. Randall and R. Sundrum. Large Mass Hierarchy from a Small Extra Dimension. *Phys. Rev. Lett.*, 83:3370–3373, 1999.

-
- [14] L. Randall and M. D. Schwartz. Unification and Hierarchy from 5D Anti-de Sitter Space. *Phys. Rev. Lett.*, 88:081801, 2002.
- [15] S.D. Drell and Tung-Mow Yan. Massive Lepton-Pair Production in Hadron-Hadron Collisions at High Energies. *Phys. Rev. Lett.*, 25:316–320, 1970.
- [16] G. Altarelli, R. K. Ellis, and G. Martinelli. Large perturbative corrections to the Drell-Yan process in QCD. *Nucl. Phys. B*, 157:461, 1979.
- [17] R. Hamberg, W. L. van Neerven, and T. Matsuura. A complete calculation of the order α_s^2 correction to the Drell-Yan K-factor. *Nucl. Phys. B*, 359:343–405, 1991.
- [18] R. Hamberg, W. L. van Neerven, and T. Matsuura. Erratum to: “A complete calculation of the order α_s^2 correction to the Drell-Yan K-factor”. *Nucl. Phys. B*, 644:403–404, 2002.
- [19] C. Anastasiou, L. J. Dixon, K. Melnikov, and F. Petriello. Dilepton rapidity distribution in the Drell-Yan process at NNLO in QCD. *Phys. Rev. D*, 91:182002, 2003.
- [20] C. Anastasiou, L. J. Dixon, K. Melnikov, and F. Petriello. High-precision QCD at hadron colliders: Electroweak gauge boson rapidity distributions at next-to-next-to leading order. *Phys. Rev. D*, 69:094008, 2004.
- [21] K. Melnikov and F. Petriello. Electroweak gauge boson production at hadron colliders through $O(\alpha_s^2)$. *Phys. Rev. D*, 74:114017, 2006.
- [22] Johann H. Kühn, A. A. Penin, and Vladimir A. Smirnov. Subleading Sudakov logarithms in electroweak processes. *Nucl. Phys. Proc. Suppl.*, 89(94–99), 2000.
- [23] Johann H. Kühn, S. Moch, A. A. Penin, and Vladimir A. Smirnov. Next-to-next-to-leading logarithms in four-fermion electroweak processes at high energy. *Nucl. Phys. B*, 616:286–306, 2001.
- [24] Bernd Jantzen, Johann H. Kühn, Alexander A. Penin, and Vladimir A. Smirnov. Two-loop electroweak logarithms. *Phys. Rev. D*, 72:051301, 2005.
- [25] Bernd Jantzen, Johann H. Kühn, Alexander A. Penin, and Vladimir A. Smirnov. Two-loop electroweak logarithms in four-fermion processes at high energy. *Nucl. Phys. B*, 731:188–212, 2005.
- [26] S. Pozzorini and A. Denner. One-loop leading logarithms in electroweak radiative corrections. *Eur. Phys. J. C*, 18:461, 2001.
- [27] U. Baur, O. Brein, W. Hollik, C. Schappacher, and D. Wackerroth. Electroweak Radiative Corrections to Neutral-Current Drell-Yan Processes at Hadron Colliders. *Phys. Rev. D*, 65:033007, 2002.
- [28] V. A. Zykunov. Weak radiative corrections to the Drell-Yan process for large invariant mass of a dilepton pair. *Phys. Rev. D*, 75:073019, 2007.

-
- [29] Stefano Pozzorini. Private communication, 19-03-2008.
- [30] The Les Houches Standard Model Handles and Candles Working Group. Tools and Jets Summary Report. *hep-ph/0803.0678v1*, 2008.
- [31] C.M. Carloni Calame, G. Montagna, O. Nicrosini, and A. Vicini. Precision electroweak calculation of the production of a high transverse-momentum lepton pair at hadron colliders. *JHEP*, 0710:109, 2007.
- [32] C.M. Carloni Calame, G. Montagna, O. Nicrosini, and A. Vicini. In the proceedings of Tevatron-for-LHC "Tevatron-for-LHC Report: Top and Electroweak Physics". *hep-ph/0705.3251*, 2007.
- [33] C.M. Carloni Calame, G. Montagna, O. Nicrosini, and A. Vicini. Precision electroweak calculation of the charged current Drell-Yan process. *JHEP*, 0612:016, 2006.
- [34] C.M. Carloni Calame, G. Montagna, O. Nicrosini, and A. Vicini. in the proceedings of Les Houches 2005 "Physics at TeV colliders. The QCD, EW and Higgs working group: Summary Report". *hep-ph/0604120*.
- [35] C.M. Carloni Calame, G. Montagna, O. Nicrosini, and M. Treccani. Multiple photon corrections to the neutral-current Drell-Yan process. *JHEP*, 0505:019, 2005.
- [36] C.M. Carloni Calame, G. Montagna, O. Nicrosini, and M. Treccani. Higher-order QED corrections to W-boson mass determination at hadron colliders. *Phys. Rev. D*, 69:037301, 2004.
- [37] S. Frixione and B.R. Webber. Matching NLO QCD computations and parton shower simulations. *JHEP*, 0206:029, 2002.
- [38] S. Frixione, P. Nason, and B.R. Webber. Matching NLO QCD and parton showers in heavy flavour production. *JHEP*, 0308:007, 2003.
- [39] G. Corcerlla, I.G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson, M.H. Seymour, and B.R. Webber. HERWIG 6: an event generator for Hadron Emission Reactions With Interfering Gluons (including supersymmetric processes). *hep-ph/0011363*.
- [40] M. Böhm, W. Hollik, and H Spiesberger. On the 1-Loop Renormalization of the Electroweak Standard Model and its Application to Leptonic Processes. *Fortschr. Phys.*, 34:687–751, 1986.
- [41] A. Vicini. Private communication, 19-03-2008.
- [42] A. D. Martin, R. G. Roberts, W. J. Stirling, and R. S Thorne. Parton distributions incorporating QED contributions. *Eur. Phys. J. C*, 39:155–161, 2005.
- [43] J. Campbell, R. K. Ellis, and D. L. Rainwater. Next-to-leading order QCD predictions for $W + 2\text{jet}$ and $Z + 2\text{jet}$ production at the CERN LHC. *Phys. Rev. D*, 68(094021), 2003.

-
- [44] Giovanni Balossini et al. Combining QCD and electroweak corrections to W-boson production at hadron colliders. *Prepared for 15th International Workshop on Deep-Inelastic Scattering and Related Subjects (DIS2007), Munich, Germany, 16-20 Apr 2007.*
- [45] G. Balossini et al. Precision predictions and tools for weak boson production at the LHC. *Acta Phys. Polon. B*, 38:2347–2355, 2007.
- [46] J. Pumplin, D. Stump, R. Brock, D. Casey, J. Huston, J. Kalk, H.L. Lai, and W.K. Tung. Uncertainties of predictions from parton distribution functions II: the Hessian method. *Phys. Rev.*, 65, 2002.
- [47] D. Bourilkov, R. C. Group, and M. R. Whalley. LHAPDF: PDF Use from the Tevatron to the LHC. *hep-ph/0605240*, 2005.
- [48] L. Evans and P. Bryant. LHC Machine. *JINST*, 3(S08001), 2008.
- [49] R. Bailey and P. Collier. Standard filling schemes for various LHC operation modes. *CERN LHC Project Note 323*, 2003.
- [50] The CMS Collaboration. The CMS experiment at the CERN LHC. *JINST*, 3(S08005), 2008.
- [51] The CMS Collaboration. The Magnet Project Technical Design Report. *CERN/LHCC 1997 - 010*, 1997.
- [52] The CMS Collaboration. The Tracker Project Technical Design Report. *CERN/LHCC 1998 - 006*, 1998.
- [53] The CMS Collaboration. CMS physics technical design report, Volume 1 - Detector performance and software. *CERN/LHCC 2006 - 001*, 2006.
- [54] The CMS Collaboration. The Electromagnetic Calorimeter Technical Design Report. *CERN/LHCC 1997 - 033*, 1997.
- [55] The CMS Electromagnetic Calorimeter group. Results of the first performance tests of the CMS electromagnetic calorimeter. *Eur. Phys. J. C*, 44(S1.1), 2006.
- [56] The CMS Collaboration. The Hadronic Calorimeter Project Technical Design Report. *CERN/LHCC 1997 - 031*, 1997.
- [57] The CMS Collaboration. The Muon Project Technical Design Report. *CERN/LHCC 1997 - 032*, 1997.
- [58] I. Ahmed et al. Testbeam results of endcap RPCs. *CMS Note 2006/044*, 2006.
- [59] The CMS Collaboration. The TriDAS Project Technical Design Report, Volume I: The Trigger Systems. *CERN/LHCC 2000 - 038*, 2000.
- [60] The CMS Collaboration. The Trigger and Data Acquisition project, Volume II: Data Acquisition & High-Level Trigger. *CERN/LHCC 2002 - 026*, 2002.

-
- [61] M. Raymond, J. Crooks, M. French, and G Hall. The MGPA electromagnetic calorimeter readout chip for CMS. *Nuclear Science, IEEE Transactions on*, 52:756–763, 2005.
- [62] G. Minderico et al. A CMOS low power, quad channel, 12bit, 4MS/s pipelined ADC for applications in particle calorimetry. *Proceedings of the 9th workshop on electronics for LHC experiments*, LECC2003:88–91, 2003.
- [63] The CMS Collaboration. The Computing Project Technical Design Report. *CERN/LHCC 2005 - 023*, 2005.
- [64] P. Elmer et al. Experience with the CMS event data model. *J. Phys.: Conf. Ser.*, 219(032022), 2010.
- [65] W. Allcock, J. Bresnahan, R. Kettimuthu, M. Link, C. Dumitrescu, I. Raicu, and I. Foster. The Globus striped GridFTP framework and server. *SuperComputing 05*, ACM Press, 2005.
- [66] T. Barras, D. Newbold, and L. Tuura. The CMS PhEDEx System: a Novel Approach to Robust Grid Data Distribution. *GridPP All-Hands Meeting*, 2005.
- [67] <http://www.spec.org/>.
- [68] D. Newbold. Private communication, 30-06-2010.
- [69] D. Bonacorsi. Experience with the CMS Computing Model from commissioning to collisions. *CMS CR 2011/015*, 2011.
- [70] <http://castor.web.cern.ch/castor>.
- [71] <http://www.iozone.org>.
- [72] <http://www.coker.com.au/bonnie++>.
- [73] Red Hat Kernel Internals. RHD236-RHEL5-EN-20070801.
- [74] R. Foadi, M. T. Frandsen, T. A. Rytto, and F. Sannino. Minimal walking technicolor: Setup for collider physics. *Phys. Rev. D*, 76(5):055005, 2007.
- [75] U. Baur, M. Spira, and P. M. Zerwas. Excited-quark and -lepton production at hadron colliders. *Phys. Rev. D*, 42(3):815, 1990.
- [76] F. M. Renard. Limits on masses and couplings of excited electrons and muon. *Phys. Lett. B*, 116(4):264, 1982.
- [77] N. Cabibbo, L. Maiani, and Y. Srivastava. Anomalous Z decays: excited leptons? *Phys. Lett. B*, 139(5-6):459, 1984.
- [78] The CDF Collaboration. Searching for the anomalous production of Z-bosons with high transverse momentum in 0.94 fb^{-1} at the Tevatron. *CDF Note 8452*, 2006.

-
- [79] The D0 Collaboration. Measurement of the inclusive differential cross section for Z bosons as a function of transverse momentum in $\bar{p}p$ collisions at $\sqrt{s} = 1.8$ TeV. *Phys. Rev. D*, 61(032004), 2000.
- [80] G. A. Ladinsky and C.-P. Yuan. Nonperturbative regime in QCD resummation for gauge boson production at hadron colliders. *Phys. Rev. D*, 50(4239), 1994.
- [81] A. Pukhov. CalcHEP 2.3: MSSM, structure functions, event generation, batchs, and generation of matrix elements for other packages. *hep-ph/0412191*, 2009.
- [82] T. Sjöstrand, S. Mrenna, and P. Skands. PYTHIA 6.4 physics and manual. *JHEP*, 05(026), 2006.
- [83] J. Pumplin, D. Stump, J. Huston, H.L. Lai, P. Nadolsky, and W.K. Tung. New Generation of Parton Distributions with Uncertainties from Global QCD Analysis. *JHEP*, 07(012), 2002.
- [84] W. Adam et al. Electron Reconstruction in CMS. *CMS AN 2009/164*.
- [85] R. Frühwirth. Track fitting with non-Gaussian noise. *Comput. Phys. Commun.*, 100, 1997.
- [86] D. Evans. Studies of Electrons in the CMS Detectors and a Proposed Search for High Invariant Mass e^+e^- Pairs. *Ph.D. Thesis, University of Bristol*, 2008.
- [87] C. Hill et al. Electron identification in offline release 5.3. *CDF Note 7309*, 2005.
- [88] G. Daskalakis et al. Measuring electron efficiencies at CMS with early data. *CMS AN 2007/019*, 2007.
- [89] Particle Data Group. Review of Particle Physics. *Journal of Physics G*, 37(7A):075021, 2010.
- [90] A. Heister, O. Kodolova, V. Konopliyanikov, S. Petrushanko, J. Rohlf, C. Tully, and A. Ulyanov. Measurement of Jets with the CMS Detector at the LHC. *CMS Note 2006/036*, 2006.
- [91] The CMS Collaboration. Algorithms for b Jet Identification in CMS. *CMS PAS BTV-09-001*, 2009.
- [92] The CMS Collaboration. Search for Resonances in the Dilepton Mass Distribution in pp Collisions at $\sqrt{s} = 7$ TeV. *CMS PAPER EXO-10-013*, 2011.
- [93] The CDF Collaboration. The search for new physics in high mass di-electron events in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV. *CDF Note 8421*.