

**AN EXAMINATION OF BACKGROUNDS TO EARLY-RUN MINIMUM-BIAS
EVENTS IN ATLAS AT THE LHC**

by

Alden Reid Stradling

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

Doctor of Philosophy

(Physics)

at the

UNIVERSITY OF WISCONSIN–MADISON

2008

© Copyright by Alden Reid Stradling 2008
All Rights Reserved

To Dorothy, Reid, Aubrey, Marian and Liam. You paid the most, and prayed the hardest. I love you without reservation, more than I can ever say.

ACKNOWLEDGMENTS

There are many thanks due to many, here. These few words can hardly express – and do not come close to quitting claim to — the debt of gratitude I owe you all.

To Profs. Sau Lan Wu and Bruce Mellado, my thanks, both for the guidance and the opportunity. Also, by extension, to the funding agencies who support their great contribution to ATLAS and the LHC – thanks for the marvelous chance to live and work at the heart of the worldwide high-energy physics community.

To Luis Flores, Trevor Vickey, Yibin Pan, and Gary and Rebecca Stradling, intrepid and dedicated proofreaders. My thanks for your care and excellent comments.

To Neng Xu, thanks for harnessing the necessary computing and not getting impatient with continual support demands.

To William Bell, Arthur Moraes, Regina Kwee, Michael Leyton, Craig Buttar and Peter Steinberg, gratitude for the teaching and support. Again thanks to William, for the detailed explanations of technical issues and moral support, e outra vez à Arthur pela ajuda com o código e a teoria.

To Vadim Taranov, for the extensive help with both accelerator halo samples and understanding the machine behavior itself.

To Paolo Calafiura, Andrea Dell'Acqua, Giulio Usai, Ketevi Assamagan and Sven Vahsen, for working enthusiastically and uncomplainingly to modify Athena to suit the needs of one user, and for shouldering the support load which I incurred.

To Amir Farbin, for patience and support, my abiding appreciation.

Again to my parents, Gary and Rebecca Stradling, for their great examples, love, patience, discipline and friendship – and the gifts (to the extent that I possess them) of curiosity, focus,

rigor, diplomacy, and eloquence. I know that the outlook was bleak, sometimes – thanks for sticking with it, during those years of sacrifice.

To Dorothy. I can never repay to her the debt I daily incur, for her work, support, valiance, and persistence in doing things right. I could not do what I do without the strengths and habits she has helped me to create – the gaps in my capacity that she has bridged and fortified. Her approval is my greatest incentive – the admiration of one so admirable has great worth. I cannot express more than the smallest part of my appreciation and love in these few words. To her and our children – what an adventure.

DISCARD THIS PAGE

TABLE OF CONTENTS

	Page
LIST OF TABLES	viii
LIST OF FIGURES	ix
NOMENCLATURE	xii
ABSTRACT	xvi
1 Introduction	1
2 Motivation	2
3 Standard Model	4
3.1 Historical	4
3.2 Horizons	6
3.3 Measurements of the Standard Model at the LHC	8
4 The Large Hadron Collider	11
4.1 Fundamentals of the LHC	11
4.2 LHC Physics Goals	13
4.3 Magnets	15
4.3.1 Cooling the Magnets	16
4.4 Beam Operations	17
4.4.1 Loading the Beam	17
4.4.2 Maintaining the Beam	17
4.4.3 Dumping the Beam	18
4.5 LHC Performance	18
4.5.1 Initial	18
4.5.2 Nominal	20
4.5.3 Limitations	20

	Page
5 The ATLAS Experiment	22
5.1 Coordinate System	23
5.2 Conception	23
5.3 Tracking	27
5.3.1 Pixel	29
5.3.2 SCT	34
5.3.3 Transition Radiation Tracker (TRT)	37
5.4 Solenoid Magnet	39
5.5 Calorimetry	39
5.5.1 Electromagnetic Calorimeter	40
5.5.2 Minimum Bias Trigger Scintillators	45
5.5.3 Hadronic Calorimeter	48
5.6 Muon Detection	52
5.6.1 Monitored Drift Tube (MDT) Chambers	53
5.6.2 Cathode Strip Chambers (CSC)	53
5.6.3 Thin Gap Chambers (TGC), Resistive Plate Chambers (RPCs) and Triggering	53
5.6.4 Toroid Magnets	54
5.7 Shielding	56
6 Data Handling	57
6.1 The ATLAS Trigger	58
6.1.1 Level-1 Trigger (LVL1)	59
6.1.2 Level-2 Trigger (LVL2)	60
6.1.3 Event Filter (EF)	61
6.2 Computing in ATLAS	61
6.3 Batch Processing and High-Performance Computing	62
6.4 ATLAS Grid Computing	63
6.4.1 Development	63
6.4.2 Production and Analysis on the Grid	64
7 Effective Simulation of ATLAS Physics	66
7.1 Event Generation	66
7.2 Simulation of the ATLAS Detector	67
7.2.1 Geant4	67
7.3 Hits and Digits	68
7.3.1 Pileup	69

Appendix

	Page
7.4 Online and offline reconstruction of ATLAS Events	70
7.5 Athena	71
7.6 The Athena Framework	71
7.7 Generating Minimum-bias Backgrounds	72
7.7.1 pp collisions	72
7.7.2 Beam-Gas Events	74
7.7.3 Beam Halo Events	75
8 Hardware and computing contributions	77
8.1 High-performance Computing in ATLAS	77
8.2 The SCT/Pixel Read-Out Drivers	78
8.2.1 Silicon Read-Out Drivers (RODs)	79
9 Minimum Bias	86
9.1 Introduction	86
9.2 Minimum-bias events	87
9.3 Backgrounds to minimum-bias events	89
9.3.1 Beam Gas	89
9.3.2 Beam Halo	93
9.3.3 Triggering	106
10 Results	109
10.1 Triggering and MBTS performance	112
10.1.1 MBTS and Beam Halo	114
10.1.2 MBTS and Beam-Gas Collisions	116
10.1.3 MBTS and Minimum-bias Pileup	118
10.2 Tracking and Event Characteristics	123
10.2.1 Background Track Behaviors	127
10.2.2 Minimum-Bias Pileup Track Behaviors	131
10.2.3 Track Parameter Cuts – Backgrounds	133
10.2.4 Track Parameter Cuts – Minimum-bias Pileup	137
11 Conclusions and Outlook	141
LIST OF REFERENCES	144
APPENDICES	

Appendix

	Page
Appendix A: Minimum bias reference datasets	153
Appendix B: Details of beam gas and halo run	154

DISCARD THIS PAGE

LIST OF TABLES

Table	Page
3.1 The Standard Model fermion families and generations.	5
3.2 Standard Model bosons.	6
4.1 LHC Parameters	14
9.1 Stage-1 LHC beam parameters.	88
9.2 Beam-gas rates in ATLAS – startup.	91
9.3 Beam-gas rates in ATLAS – run.	92
9.4 Beam gas proportions in ATLAS	93
9.5 Halo Particle Species	103
10.1 Track selection cuts.	134
10.2 Beam Halo with no minimum-bias pileup.	135
10.3 Beam Gas with no minimum-bias pileup.	136
10.4 Minimum-bias pileup (control).	138
10.5 Beam halo with minimum-bias pileup.	139
10.6 Beam Gas with minimum-bias pileup.	140
Appendix	
Table	

DISCARD THIS PAGE

LIST OF FIGURES

Figure	Page
3.1 LHC pp scattering modes.	9
3.2 pp and $\bar{p}p$ scattering data.	10
4.1 Schematic layout of the LHC.	12
4.2 A cross section of an LHC dipole magnet.	16
5.1 Cutaway of the ATLAS detector (perspective view).	22
5.2 A perspective cutaway of the ATLAS Inner Detector.	27
5.3 Material distribution in the ATLAS ID.	28
5.4 The pixel subsystem, barrel and endcaps, including its support frame.	30
5.5 An exploded perspective view of a pixel module.	31
5.6 An SCT Module	34
5.7 An assembled SCT barrel.	35
5.8 TRT cutaway.	37
5.9 The ATLAS inner detector solenoid, before insertion into its cryostat.	39
5.10 Calorimeter systems.	40
5.11 Calorimeter material distribution.	42
5.12 LAr Electrode.	43
5.13 EM calorimeter cell structure	44

Figure	Page
5.14 ATLAS MBTS cutaway view.	47
5.15 Endcap cryostat cutaway.	48
5.16 Tile calorimeter module.	49
5.17 Hadronic endcap module.	50
5.18 Tungsten forward calorimeter detail.	51
5.19 Muon perspective cutaway.	52
5.20 Muon trigger schematic.	54
5.21 Photo of endcap toroid A being moved.	55
5.22 Toroid system winding geometry.	55
5.23 ATLAS shielding.	56
6.1 The Trigger/DAQ system.	58
6.2 LVL1 trigger decision flow diagram.	59
7.1 ATLAS simulation flow.	67
7.2 The Athena component model.	71
8.1 An ATLAS Pixel/SCT ROD.	80
9.1 The Phase-I LHC collimator layout.	95
9.2 Photo – Tertiary Collimator.	96
9.3 Halo particle distributions at the scoring plane	98
9.4 Beam halo scoring plane schematic	99
9.5 Halo particle p_T w.r.t. x	100
9.6 Halo particle dx w.r.t. x	101
9.7 Multiplicity of beam halo events.	104

Figure	Page
9.8 Beam halo p_T vs. multiplicity plots.	104
9.9 Radial density of halo particles	105
9.10 The ATLAS MBTS trigger strategy.	107
10.1 Threshold scans with the old MBTS threshold.	113
10.2 New MBTS threshold fit.	114
10.3 Beam halo trigger efficiency scans.	116
10.4 Trigger efficiency scans for beam-gas events.	117
10.5 Trigger efficiency – MBTS_1.	120
10.6 Trigger efficiency – MBTS_2.	122
10.7 Trigger efficiency – MBTS_1_1.	123
10.8 Canonical minimum bias charged particle plots.	126
10.9 $dN_{ch}/d\eta$ – no pileup.	127
10.10 dN_{ch}/dp_T – no pileup.	129
10.11 Charged particle multiplicity – no pileup.	130
10.12 $dN_{ch}/d\eta$ for normal pileup.	131
10.13 dN_{ch}/dp_T for normal pileup.	132
10.14 Charged particle multiplicity for normal pileup.	133
Appendix	
Figure	

DISCARD THIS PAGE

NOMENCLATURE

ϕ	The azimuthal angle, where $\tan(\phi) \equiv \frac{p_y}{p_x}$. It is measured around the beam axis.
θ	The angle measured from the beam axis.
η	The pseudorapidity is defined as $\eta = -\ln(\tan \frac{\theta}{2})$.
σ	This variable usually refers to the standard deviation of some implicit Gaussian distribution. It is used in several ways. In the context of detector performance, σ refers to the resolution of a reconstructed quantity. It is also the symbol that is used for the cross-section of a given particle interaction.
ADC	Analog to Digital Converter.
ADD	The Arkani-Hamed-Dimopoulos-Dvali brane-world model.
ALICE	A Large Ion Collider Experiment at CERN LHC – one of the four main experiments at the LHC.
AOD	Analysis Object Data
ATLAS	A Toroidal LHC ApparatuS – one of the four main experiments at the LHC.
Axes	The beam direction defines the z-axis, and the x-y plane is the plane transverse to the beam direction. The positive x-axis is defined as pointing from the interaction point to the center of the LHC ring, and the positive y-axis is pointing upwards.
BPM	Beam Position Monitor
Calo	Short for Calorimeter.

CM	Center of Mass energy – \sqrt{s}
CMS	Compact Muon Solenoid – one of the four main experiments at the LHC.
CSC	Computing Service Challenge
DAQ	Data AcQuisition
DFM	Data Flow Monitor
DSP	Digital Signal Processor
ECAL	Electromagnetic Calorimeter
EF	Event Filter – a part of the trigger system.
EFB	Event Fragment Builder – a part of the silicon ROD.
EM	Electromagnetic
ESD	Event Summary Data
FE	Front End (on-detector electronics)
FNAL	Fermi National Accelerator Laboratory
FPGA	Field-Programmable Gate Array
GB	gigabyte
HCAL	Hadronic Calorimeter
ID	Inner Detector
IP	Interaction Point – the beam crossing region at the center of an experiment.
LAr	Liquid Argon
LCG	LHC Computing Grid

LEP	Large Electron-Positron collider, LHC's predecessor.
LHC	Large Hadron Collider
LHCb	LHC-beauty – one of the four main experiments at the LHC.
LHe	Liquid Helium
LN ₂	Liquid Nitrogen
LSP	Least Supersymmetric Particle
LVDS	Low-Voltage Differential Signaling
LVL1	Level-1 Trigger
LVL2	Level-2 Trigger
MB	megabyte
MBTS	Minimum Bias Trigger Scintillators
MCC	Module Control Chip – part of a Pixel module.
MDT	Monitored Drift Tube – a part of the muon spectrometer.
MIP	Minimal Ionizing Particle – relevant to the TRT.
MSSM	Minimal Supersymmetric Standard Model
NIM	Nuclear Instrumentation Module
PS	Proton Synchrotron – one of the CERN accelerator complex.
PSB	Proton Synchrotron Booster
QCD	Quantum Chromodynamics
QED	Quantum Electrodynamics

RDO	Raw Data Object – output from Athena digitization.
RF	Radio Frequency – referring to the beam driving method used in the LHC.
RMS	Root Mean Squared
ROB	Read-Out Buffer – part of the Trigger/DAQ system.
ROD	Read-Out Driver – part of the Trigger/DAQ system.
RoI	Region of Interest – area for further examination by the LVL2 trigger.
ROS	Read-Out Subsystem – a trigger component.
RPC	Resistive Plate Chamber – part of the muon spectrometer.
SCT	SemiConductor Tracker
SFI	Sub-Farm Input – a trigger component.
SM	Standard Model
SPS	Super Proton Synchrotron – a part of the CERN accelerator complex.
SUSY	SUperSYmmetry
TB	terabyte
TGC	Thin-Gap Chamber – part of the muon spectrometer.
TOT	Time over Threshold – specific to the pixel subsystem.
TRT	Transition Radiation Tracker
WIMP	Weakly Interacting Massive Particle

AN EXAMINATION OF BACKGROUNDS TO EARLY-RUN MINIMUM-BIAS EVENTS IN ATLAS AT THE LHC

Alden Reid Stradling

Under the supervision of Professors Sau Lan Wu and Bruce Mellado
At the University of Wisconsin-Madison

The initial stages of the Large Hadron Collider (LHC) run will be a source of critical information – about the ATLAS detector and about the physics of pp collisions at $\sqrt{s} = 14$ TeV, including parton distribution evolution and the cross-section of σ_{pp} . The accelerator itself will be the source of some detector interest, as we have a first look at what have so far been speculations on the quality of the vacuum in the experimental insertion, and the cleanliness of the beam from the accelerator. The shakedown period, with its low beam squeeze, low luminosity, and undemanding trigger menus, will be very useful in addressing some of these questions, as it lacks the pileup and radiation levels that will arrive with higher luminosity – making it an important opportunity to investigate minimum-bias events in relative isolation. For the short lifetime of the Minimum Bias Trigger Scintillators (MBTS), which are expected to fail within a few months of running, they will aid in discriminating the minimum bias signal of inelastic non-single-diffractive pp collisions. Using single- or double-coincidence signatures, the MBTS system and other trigger and analysis strategies attempt to avoid triggering on otherwise empty bunch crossings and eliminate the effects of beam-gas collisions and beam-halo effects which would lead these spurious triggers and reduce the general minimum-bias trigger efficiency. An examination of the effects of beam halo and beam-gas interactions on the minimum-bias trigger response is made. The signatures of the beam halo and beam gas are examined from the standard ATLAS tracking reconstruction.

Sau Lan Wu and Bruce Mellado

ABSTRACT

The initial stages of the Large Hadron Collider (LHC) run will be a source of critical information – about the ATLAS detector and about the physics of pp collisions at $\sqrt{s} = 14$ TeV, including parton distribution evolution and the cross-section of σ_{pp} . The accelerator itself will be the source of some detector interest, as we have a first look at what have so far been speculations on the quality of the vacuum in the experimental insertion, and the cleanliness of the beam from the accelerator. The shakedown period, with its low beam squeeze, low luminosity, and undemanding trigger menus, will be very useful in addressing some of these questions, as it lacks the pileup and radiation levels that will arrive with higher luminosity – making it an important opportunity to investigate minimum-bias events in relative isolation. For the short lifetime of the Minimum Bias Trigger Scintillators (MBTS), which are expected to fail within a few months of running, they will aid in discriminating the minimum bias signal of inelastic non-single-diffractive pp collisions. Using single- or double-coincidence signatures, the MBTS system and other trigger and analysis strategies attempt to avoid triggering on otherwise empty bunch crossings and eliminate the effects of beam-gas collisions and beam-halo effects which would lead these spurious triggers and reduce the general minimum-bias trigger efficiency. An examination of the effects of beam halo and beam-gas interactions on the minimum-bias trigger response is made. The signatures of the beam halo and beam gas are examined from the standard ATLAS tracking reconstruction.

Chapter 1

Introduction

This dissertation will begin with an exploration of the physics motivations behind the Large Hadron Collider (LHC) in Chapters 2 and 3, and then discuss some of the particulars of constructing such a machine – addressing the accelerator ring in Chapter 4, then moving to the specifics of the design and construction of the ATLAS detector and computing, looking at the hardware and software requirements in Chapters 5, 6 and 8, then dwelling a bit on the software and Monte Carlo techniques used to simulate the physics and detector response in Chapter 7. We will finally come back to the accelerator and beampipe, investigating the small interface between the beam and the experimental conditions and examining the backgrounds that are inherent to the accelerator itself in Chapter 9.

The detector is something that fascinates me in and of itself, but I will not describe it in as much detail as I would like, except insofar as it is necessary to illustrate the parts I *did* work on – specifically the SCT/Pixel readout drivers, the Grid computing effort, and the cabling. For this reason (and because of its relevance to the minimum bias study and related topics), the inner detector silicon components will get more than their share of detail, and efforts will be made to illustrate the details of LHC operation to clarify certain aspects of the beam halo background. For more detail, please refer to the forthcoming detector description paper in [15].

Chapter 2

Motivation

The LHC exceeds any previous accelerator's attainable center-of-mass energy by a factor of seven or more. As we examine the events created in this regime, we must have an understanding of their basic behavior before we can look at more subtle effects. While the fundamental operations of the Standard Model are well-accepted and understood, there are difficulties in applying it to some parts of a practical experiment. The practical difficulties in operating such an immensely complex and demanding experiment stem from many sources. Principal among these are time-dependent effects, calibration of individual detectors, alignment issues, data-collection requirements, LHC-experiment interaction, and fault handling. Even in the most ideal circumstances, however, when the practicalities are well in hand, the experimental program is far from clear. The detector has physical and temporal resolution limits inherent in its design, due to compromises that must be made for engineering and funding reasons. Ambiguities and degenerate signatures abound, and a full understanding of the detector's response, the experimental and statistical error, and the basic physics involved are fundamental to the process of sifting the signal for something new. While ambiguities and alternate explanations for signals and signatures exist, no discovery can be claimed.

Some of the fundamental complications faced by ATLAS (A Toroidal LHC ApparatuS) stem from our present inability to make *ab-initio* calculations of some non-perturbative physical processes. Specifically, the problem of infrared slavery in quantum chromodynamics (QCD) [16] makes it impractical to directly calculate higher-order corrections to low-energy QCD processes, because of the gluons' self-interaction. Only when energies are high enough do the perturbative loop calculations stop diverging. To permit the calculation of cross-sections in a

hadron collider, however, one has to know the distribution of sub-hadronic particles. Since this low-energy bound state falls firmly in the non-perturbative part of QCD calculation, there is no analytical way to calculate the distribution of partons in, for example, the protons which will be collided in the center of the ATLAS experiment. The parton distribution functions (PDFs) must be determined by direct measurement (by definition – they are the encapsulation of the non-perturbative QCD behavior at p_T of less than 2 GeV), along with the pp cross-section. While evolution schemes like DGLAP [17, 18, 82] describe the energy ranges successfully reached to this point, the jump to $\sqrt{s} = 14$ TeV will be a test of that scheme’s completeness, with inclusive jet distribution measurements. The soft end of the QCD scale, however, is not as easily treated. Tuned simulations like PYTHIA [57] and PHOJET [86, 102] contain extrapolations of high-energy behavior of hadronizing struck quarks - but there are no data to tune them properly yet. Furthermore, such processes cross the threshold between the high- (> 1 GeV) and low- Q^2 regimes, making them inaccessible to calculation. Measurements of these behaviors have been made in various experiments, and are available for comparison.

The $dN_{ch}/d\eta$ and dN_{ch}/dp_T behavior of the minimum-bias (inelastic, non-single-diffractive pp , for ATLAS)[90] events can be used in such comparisons, and give us a handle on important parts of the hadronization simulation [90] and its scaling, allowing us to compare our QCD background behavior directly with previous experiments.¹

The corrections that arise from this will propagate up through the rest of the physics analysis done in ATLAS, in the understanding of QCD behaviors and backgrounds. Successfully determining the real $dN_{ch}/d\eta$ and dN_{ch}/dp_T distributions requires setting appropriate triggers for the minimum-bias events, and filtering backgrounds (such as beam-gas collisions and beam halo) from the data, giving a clean look at the distributions. This dissertation will explore the processes by which these backgrounds are created, some of their characteristics, and the means of their removal from the minimum-bias signal.

¹ N_{ch} – number of charged particles, and p_T is the transverse momentum of a particle – the part of its momentum directed perpendicular to the beamline

Chapter 3

Standard Model

3.1 Historical

Significant progress has been made in the last century in understanding the mechanisms that lie beneath the observed physical world. Exploiting a few effects at first, it gradually became clear that the behavior of matter on an atomic scale differed substantially from that which is observed on the everyday macro scale. In the intervening years, a formalism has grown up around the observations made and the theories that explain them best, and it has become known as the Standard Model (SM) [19].

As ambiguities have been resolved, the picture has become clearer. We find that the material world as we know it is built of several fundamental particles, and the fields that moderate their interactions one with another. These fields constitute the four basic forces observed in the universe, of which three are encompassed in the SM – the strong and weak internuclear forces and the electromagnetic (EM) force. The susceptibility of a particle to interaction with a given field can be expressed as “charge” – but charge becomes more complicated than the familiar EM case of positive and negative – the strong force, for example, admits three charge types (known as red, green and blue), each of which has two poles (for example, blue and anti-blue). Gravitation is notably absent from the SM formulations.

Particles are regarded, for our purposes, as *fundamental* rather than *composite*. They can be grouped into several families and categories. First, all the half-integer spin particles of which ordinary baryonic matter is made are called “fermions”. Within this class, particles are further broken into “quark” or “lepton” designations, based on whether or not they participate in strong

interactions (quarks do, leptons don't). Particle “families”, such as electron – muon – tau, correspond to the rows in Table 3.1 and have identical quantum numbers, differing only in mass and lepton number(between quarks and leptons) the flavor quantum number in the other two generations. Each generation comprises a lepton/lepton neutrino doublet, and an up-type and down-type quark.

	Generation 1		Generation 2		Generation 3	
Quarks	Up	u	Charm	c	Top	t
	Down	d	Strange	s	Bottom	b
Leptons	Electron Neutrino	ν_e	Muon Neutrino	ν_μ	Tau Neutrino	ν_τ
	Electron	e	Muon	μ	Tau	τ

Table 3.1: The Standard Model fermion families and generations.

The fields that these particles interact in can be expressed in terms of field quanta called “bosons” – integer-spin carriers, or propagators, of energy and quantum numbers from one particle to another. The bosons for the EM and strong forces are massless, a requirement for long-distance interactions at low energies. In principle, in fact, all gauge bosons “should” be massless. The question immediately arises why, if the strong force has a massless gauge boson, do we not feel its effects outside the confines of the atomic nucleus, and why even there it should be for the most part residual rather than direct. As it happens, the gluon (unlike a photon, for example) has a self-interaction term (the bosons themselves interact with each other) – making the force between color-charged particles increase unboundedly as their distance increases.

The weak force is also short range, but for a different reason – its carrier bosons are quite massive, and consequently have very short potential path lengths at any energies other than the highest (due to Heisenberg Uncertainty – $\Delta E \Delta t \geq \hbar/2$ limits the distance of a massive virtual particle). The massive vector bosons (Z^0 , W^\pm) are a notable deviation from the normal tendency for bosons to be massless. In fact, the initial predictions of an intermediate vector boson arose from the need to restore unitarity to preliminary models of weak interactions, at least up to a

certain energy [19]. Later, the W^\pm and Z^0 bosons, fully fleshed out, served as the mechanism to provide spontaneous symmetry breaking to the model without introducing unphysical (unobserved) massless Goldstone bosons. The successful prediction of the W^\pm and Z^0 masses by Weinberg [87] (16 years before their detection) was a strong indicator that dependence on gauge invariance, unitarity and spontaneous symmetry breaking was a fruitful avenue to pursue. An offshoot of symmetry breaking is the emergence of the Higgs mechanism, which proposes that the weak force $SU(2)_L$ symmetry group underwent a phase change as the universe cooled after its high-energy origins, and that the W^\pm and Z^0 bosons' masses (and therefore their longitudinal polarization) are results of that phase change. The remaining degrees of freedom in the neutral doublet would yield a massive, spin-0 particle with a strongly mass-dependent coupling to all quarks and leptons – and to its siblings in the weak force symmetry group.

Electromagnetic Force		Weak Force		Strong Nuclear Force	
Photon	γ	W^\pm, Z Gauge Bosons	W^+, W^-, Z^0	Gluons	g

Table 3.2: The Standard Model bosons.

3.2 Horizons

While the SM addresses all known particles and their interactions, and encapsulates the present knowledge of quantum electrodynamics (QED), weak force interactions and quantum chromodynamics (QCD), there remain areas left unaddressed for various reasons. Indeed, the goals of the LHC (mentioned briefly in Section 4.2) center around some of these blind spots.

One of the underpinnings of the SM, QCD, is the source of the first problem, and the most germane to this dissertation. The aforementioned self-interaction of gluons increases the force between separated quarks without limit, meaning that there is no useful concept of a free quark outside, for example, the fireball of a nuclear collider's interaction point. The situation is outside of our ability to calculate using perturbation theory (the self-interaction terms diverge, so there

is no good place to truncate the series) – making *ab initio* calculations impossible. A better understanding of low- Q^2 QCD in pp collisions is one of the goals of the minimum-bias event studies.

Another primary difficulty with the Standard Model is its inability to derive all its own parameters from first principles. It is a powerful predictive system, but only when correctly tuned. Such constants as the masses of all the fermions, coupling constants for each of the three forces (which evolve with interaction energy), the four mass- to electroweak-eigenstate angles (CKM matrix) to predict flavor mixing, and the Higgs free parameter must be included by hand. None of the proposed courses of research at the LHC makes any significant headway in reducing this problem – indeed, supersymmetry (SUSY) can aggravate the issue substantially by doubling the number of particle masses to measure and plug in, plus other additions. However, SUSY can address several major concerns, including the hierarchy problem – the fact that the vast divergence between the Higgs scale and gravitation introduced a divergent Higgs mass unless precisely balanced through fine tuning.

Supersymmetry implies an extension of the idea of symmetry breaking, suggesting that perhaps all particles have “partners” – complementary particles with identical quantum numbers and a spin difference of 1/2. If such particles existed, their coupling to a Higgs propagator would cancel the divergent lepton couplings at higher orders that make the fine tuning so suspect.

SUSY particles are also important in the possible explanation of cold dark matter – the “lightest supersymmetric particles” (LSPs) are very weakly interacting, and could account for the mass apart from the observable fermionic matter we observe. SUSY particles have not been observed in any previous experiments – but if they exist, their upper mass bounds are well within the reach of the LHC. It may also be able to pull gravity into the picture after all. This final weakness, as mentioned before, is the inability of the SM to speak to the topic of gravity. At 10^{-36} times weaker than the EM force, there is good reason that it is harder to integrate – for all the efforts of large-baseline interferometry, no direct measurement of a gravitational wave has yet been made – but making quantum predictions about gravity is problematic.

Though it is unlikely that any predictions of string theory will be addressable at LHC energies (or for that matter, at any accelerator within human reach), there are some possibilities. If energy leaks from the local brane from collisions, it will be measurable as missing energy, and if micro black holes form, their presence and signature will be of great interest as probes of the links between general relativity and quantum field theory. A final possibility lies in the realm of string theory, where some anti-De Sitter space models of the universe have been shown to correspond to a conformal field theory, possibly putting some observables within the reach of accelerator physics.

3.3 Measurements of the Standard Model at the LHC

Most of the previous section (and Section 4.2) has been included to give an idea of the kinds of primary measurements that will be made at the LHC, and the techniques used to make them. Many measurements will look for missing energy from a collision. Others will be primarily interested in seeing an excess of one sort or another, perhaps with a particular topology.

All of the measurements will require a good understanding of the pp cross-section evolution, and all of them will be looking at high- p_T signatures buried in a background of simultaneous pp collisions (centered around 23 pp interactions per bunch crossing at design luminosity) and multiple interactions in the signal pp collision itself. Most of these pp collisions will be soft – that is, the struck partons will acquire only a small p_T of ≤ 2 GeV. These collisions will compose the majority of both the nonsignal collisions in each bunch crossing (one definition of the overloaded term “pileup”), and as an underlying component of even the signal events being sought in a given analysis.

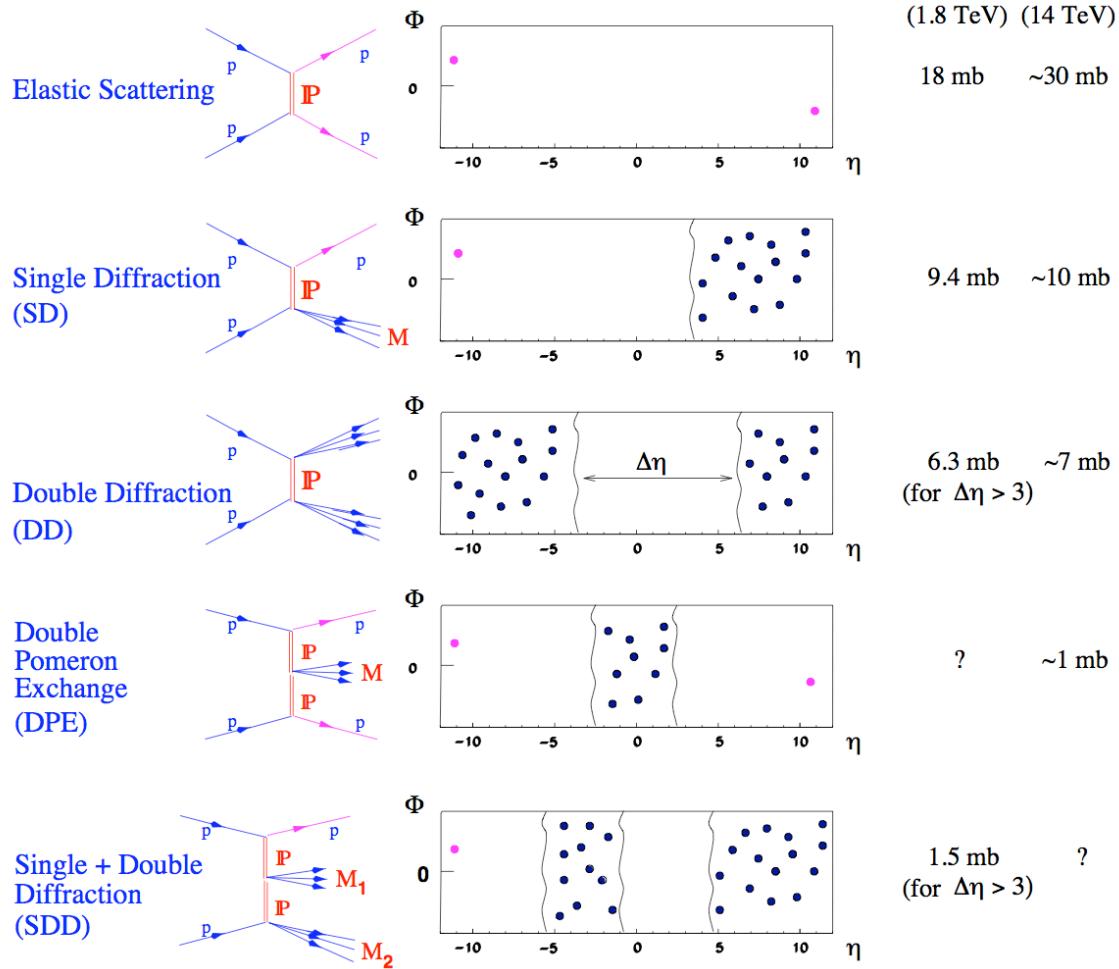


Figure 3.1: Examples of various possible scattering modes at the LHC, and their Tevatron (1.96 TeV) and LHC (14 TeV) cross-sections. The characteristic single-diffractive (one side empty) and double-diffractive (large rapidity gap) signatures are reflected in the triggering mechanisms used for minimum-bias triggers (see Chapter 9). Figure from [21].

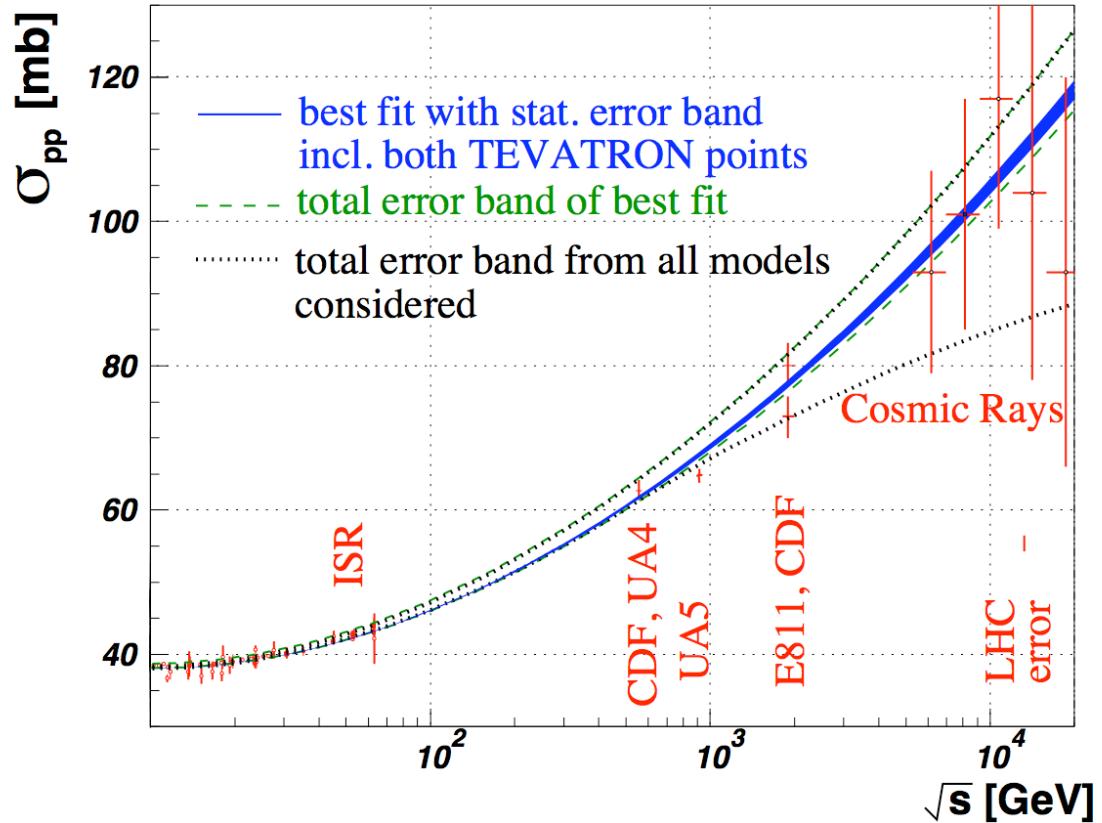


Figure 3.2: COMPETE [22] fits to all available pp and $\bar{p}p$ scattering data with statistical (blue solid) and total (green dashed) error bands, the latter taking into account the Tevatron ambiguity. The outermost curves (dotted) give the total error band from all parameterizations considered. Figure from [23].

Chapter 4

The Large Hadron Collider

This chapter will summarize the characteristics and challenges associated with the construction and operation of the Large Hadron Collider (LHC) and examine the scope of the project.

4.1 Fundamentals of the LHC

The LHC has been the culmination of more than two decades' research and construction [24]. Under the border area between Geneva, Switzerland and neighboring France runs a 27 km annular tunnel. Originally prepared for the Large Electron-Positron Collider (LEP) machine, the tunnel has now been reused to host the new accelerator.

In contrast to the LEP [25] accelerator, whose beams were made up of accelerated electrons and positrons, the LHC will accelerate and collide protons in normal physics mode, with the capacity to accelerate charged Pb ion in specific runs for the nuclear physics community.

The use of pp collisions has its great advantages. Because of its mass/charge ratio, a proton radiates much less of its energy in the accelerator than an electron would do in its place. The disadvantages of using protons (or any composite particle) lie in the complexity of the collision. In LEP, the events were relatively clean, arising from the interactions of point particles – backgrounds were generally limited to beam halo and scattering from gas atoms in the imperfect vacuum around the collision point. In the LHC, collisions will take place between any of a number of partons (the quarks, virtual quarks and gluons) within the proton, or perhaps between more than one per proton.

In addition to this complication, the accelerator is substantially upgraded. Where LEP started with four bunches of electrons at a time in the accelerator, the LHC will have a peak design field of 8.33 T, 2,808 bunches per beam, and a crossing rate of 40 MHz. The detectors will have to handle much higher final state particle multiplicities, much more often, and at much higher energies. The design Center-of-Mass energy (CM, expressed hereafter as \sqrt{s}), is 14 TeV.) [24]

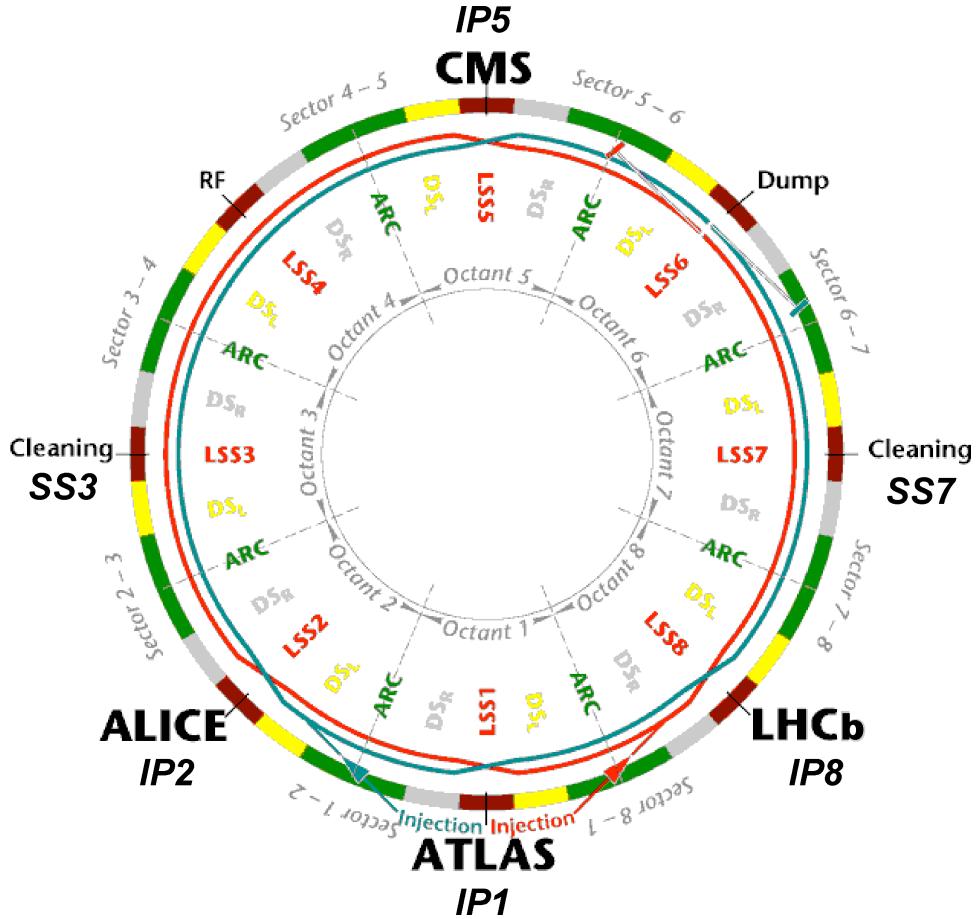


Figure 4.1: Schematic layout of the LHC.

Two high-luminosity experimental insertions at Point 1 (ATLAS) and Point 5 (CMS) [26] house general-purpose detectors, while Point 2 (ALICE) [81] and Point 8 (LHCb) [99] are low-luminosity special-purpose insertions (for Pb ion collisions and B-physics, respectively). Point

4 contains the RF (radio-frequency) acceleration sections, Point 6 is the beam dump, and Points 3 and 7 are beam cleaning insertions.

4.2 LHC Physics Goals

- **Search for a Standard Model (SM) Higgs boson from the LEPII low-mass limit (114.1 GeV) up to the theoretical upper bound of 1 TeV [27]**

Examining symmetry breaking in the electroweak sector ($SU(2) \times U(1)$) of the Standard Model (SM) is the key physics motivation for the LHC, which will either find the Higgs or invalidate the SM symmetry-breaking mechanism.

- **Precision Standard Model measurements**

In the process of providing the necessary statistics for Higgs searches, the LHC will serve as a B-meson, W, Z and top quark factory, with a $t\bar{t}$ cross-section (for example) of ~ 1 nb. In contrast to the top-quark discovery at FNAL, the LHC will allow the mass to be measured to within theoretical (rather than statistical) errors.

- **Beauty physics and CP violation**

Further Standard Model tests include investigation of CP violation, which can also have consequences to the validity of the SM and various of its proposed extensions. Measurement of CP violation in the B_d^0 system involves determining the three interior angles of the unitarity triangle corresponding to the Cabibbo-Kobayashi-Maskawa (CKM) matrix. This goal is largely handled by the special-purpose LHCb detector, but is very attractive at the low initial pileups in ATLAS as well, where vertexing is not as challenging.

- **Search for supersymmetry (SUSY), exotics, black holes, monopoles and strangelets**

The LHC experiments will be able to make a definitive statement on the existence of a SUSY signal. Hermeticity is key to the measurement, since the presence of a Weakly Interacting Massive Particle (WIMP) [28] (assuming this is a Lightest Supersymmetric Particle (LSP), the signature will be missing E_T) predicted in the Minimal Supersymmetric Standard Model (MSSM) [28], the first of the mass hierarchy expected, would

Parameter	Units	Injection	Collision
Proton energy	[GeV]	450	7000
Relativistic gamma		479.6	7461
Number of particles per bunch			$1.15 \cdot 10^{11}$
Number of bunches			2808
Circulating beam current	[A]		0.582
Stored energy per beam	[MJ]	23.3	362
RMS bunch length	[cm]	11.24	7.55
RMS beam size (IP1, IP5)	μm	375.2	16.7
RMS beam size (IP2, IP8)	μm	279.6	70.9
Peak luminosity (IP1, IP5)	$\text{cm}^{-2}\text{sec}^{-1}$	–	$1.0 \cdot 10^{34}$
Inelastic cross section	[mb]		60.0
Total cross section	[mb]		100.0
Events per bunch crossing		–	19.02
RMS bunch length	[cm]	11.24	7.55
Luminosity lifetime	[hours]	–	29.1

Table 4.1: LHC Parameters

likely be a neutralino, and its signature would be that of missing energy in the collision. Other missing energy channels might indicate extra dimensions, helping with the supergravity and ADD space questions [126]. Black holes might also make an appearance, with a dramatic isotropic decay product distribution [73].

- **Quark-Gluon Plasma/Nuclear**

During Pb-Pb runs, ATLAS will work alongside ALICE and CMS to explore these collisions at 105 TeV.

4.3 Magnets

The LHC magnets consist of beam direction, focusing and correction elements. The backbone 1,232 dipole magnets, as seen in Figure 4.2, form the ring itself, deflecting the counter-rotating proton beams in separate parallel beampipes. The magnet is designed to have the \vec{B} in one pipe mirror that of the other using the same magnetic coils, saving both energy and materiel. The coils themselves are cooled to 1.9 K using liquid helium II, making the niobium-titanium windings superconducting up to the critical design field, and a bit beyond for safety. Control systems monitor the dipole voltage in realtime to prevent quench¹ damage. In the event of a quench condition, dipole voltages start to rise (because of increased resistance) and a passthrough diode comes into play after a certain threshold, taking the magnet out of the string of series-linked dipoles and initiating a discharge of the rest of the string of 154 dipoles over a period of about two minutes. This automated action prevents the current from the entire string being dumped into a single dipole and melting it.

¹A “quench” is an unanticipated loss of superconductivity in a magnet due to localized heating. This can arise from friction between moving conductors as the magnet changes temperature or field geometry, or from energy losses from the beam. The quench quickly spreads as regions around the suddenly-resistive region are heated in turn.

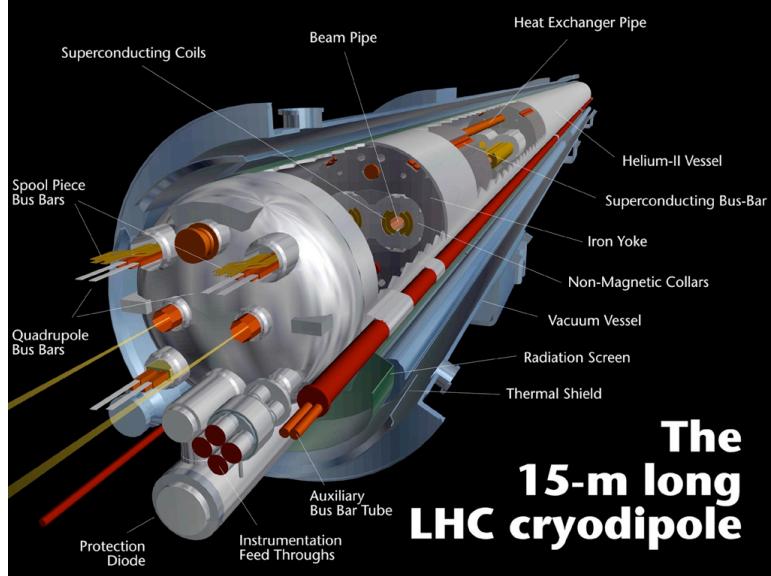


Figure 4.2: A cross section of an LHC dipole magnet.

Other magnets, such as correcting magnets in the dipoles, quadrupoles, sextupoles and octopoles, round out the complement of approximately 9,300 magnets in the accelerator, many embedded within the dipole shells themselves. The large number of correcting magnets is necessary to create beam stability in what is essentially a confined plasma.

4.3.1 Cooling the Magnets

Regarding cooling, the LHC is a challenge. The facilities have to be in place to cool 40,000 tonnes of cold mass to between 4.5 K and 1.9 K over a circumference of 27 km. The initial cooldown [29] will require the equivalent of $3.17 \cdot 10^6$ gallons of liquid helium (LHe) for its initial cooling, though much of it will be done with truckloads of liquid nitrogen, until the temperatures get down to the liquid nitrogen cooling limit (85 K). This will require $\sim 1,260$ tonnes of LN₂, each truckload running to 40 m³ (or about 32.2 tonnes), coming 48 times/day at each of the 5 cryogenic points along the accelerator ring for about 5 days. When the helium compressors finally take over, they finish the cooling to 4.5 K for the warmer components. A

final cooling stage (where the LHe is evaporated at 10 millibars) takes the main dipoles down to the final temperature of 1.9 K.

4.4 Beam Operations

4.4.1 Loading the Beam

The beam is brought to its injection energy in a cascade of older, refitted accelerators. The initial acceleration of the proton packets is done by a 50 GeV linear accelerator to boost the particles to an initial energy sufficient to run in the synchrotron chain, and they are handed off to the Proton Synchrotron (PS) through the Proton Synchrotron Booster (PSB).

While the linac is capable of producing a continuous stream of particle bunches, the transfer of bunches from the PS to the next accelerator in line requires an intervention. “Kicker” magnets are used to extract the bunches once they are brought to the appropriate energy – and during the kicker rise time, no bunches can be expected to make the transfer correctly. There are, therefore, unavoidable gaps in the bunch pattern.

From the PS to the SPS (Super Proton Synchrotron), another set of gaps are introduced, with a different rise time, and the energy is increased to its required LHC injection value of 450 GeV. The kicker magnets have a higher required field at this energy, and the rise time is longer. Finally, an artificial gap is required at the end of a filling pattern to allow for the dump kicker magnets in the LHC to have time to correctly redirect the beam before having particle buckets to handle. At the enormous energies of the LHC, the beam is quite capable of doing mechanical damage to the accelerator, and misdirected bunches during a kicker magnet rise time would be problematic.

4.4.2 Maintaining the Beam

Once the beam has been injected into the LHC ring, it is accelerated to 7 TeV by a series of RF cavities, eight per beam, running at 400.8 MHz. They bring the beam up to design energy, then maintain it against the 6 keV beam losses per turn expected from synchrotron radiation.

The beam itself has an anticipated luminosity lifetime of 20-30 hours, after which a dump and fill are necessary.

4.4.3 Dumping the Beam

Because the beam is very energetic (equivalent to a 1,000 kg station wagon traveling at approximately 1,870 mi/h) and concentrated (area on the order of tens of square microns), the question of where and how to dump it is a very serious one – especially at design energy.

An initial consideration is that at any time, a quench of a dipole or corrector magnet might degrade the beam trajectory control. The beam itself can handily destroy the vacuum capability of the beampipe or other critical components, including dipole magnets or delicate detectors. A number of collimators are installed in the beamline that both clean up the beam in several stages and absorb the damage of a serious beam deviation. The sound of a beam striking a collimator jaw in tests is macroscopically audible, and resembles that of a large piece of thick steel being struck by a heavy hammer.

If the beam behaves for long enough to dump it in a controlled way (or at the end of a fill's lifetime), the beam dumps at Point 6 are used to absorb the energy. The dumps consist of two targets (one for each beam direction) and a pair of beam sweepers to shift it. Each dump target consists of a 7.7 m cylinder of graphite, with a radius of 32 cm. The beam sweepers take the diverted beam from the beam septum magnet and sweep it in a “pretzel” pattern over the face of the target block to keep temperatures reasonably low – in the range of 1800 K.

4.5 LHC Performance

4.5.1 Initial

Since the nominal beam width at the interaction point (for design luminosity) is $7 \mu\text{m}$, beam parameters tunes are very sensitive. In addition to the geometrical problems with creating a 27-km magnetic lattice with the alignment and electrical stability to create a consistent path for the bunches in the accelerator, one has to deal with the nature of the beam itself – as a plasma, it has inherent instabilities from its uniform positive charge. By tuning the currents in the bending

and corrector magnets, and positioning the collimators correctly, the accelerator operators will gradually increase the accuracy of the beam overlap, minimize the beam losses (which reduce beam lifetime and can cause magnet quenches beyond a certain point).

Beam commissioning has been scheduled to begin the third week in May of 2008, and the pilot physics run (43 bunches, no crossing angle and no beam squeezing) at the end of July 2008. This is, at present, a best-case scenario [30]. Pilot physics will produce less than one event per crossing, with minimal interaction point beam tightness. With the start of 75 ns operations planned for the end of 2008, the beam squeeze will increase and the crossing angle will start to be implemented, but with moderate intensities. The event rate will at less than one event per bunch crossing, and the nominal luminosity will be around $10^{32} \text{ cm}^{-2} \text{ sec}^{-1}$

In this case, luminosity is defined [31] as

$$\mathcal{L} = \frac{(N_{b1}N_{b2}) \cdot k_b \cdot f_b}{2\pi \sqrt{(\sigma_{x1}^2 + \sigma_{x2}^2)(\sigma_{y1}^2 + \sigma_{y2}^2)}} \exp\left(-\frac{(\bar{x}_1 - \bar{x}_2)^2}{2(\sigma_{x1}^2 + \sigma_{x2}^2)} - \frac{(\bar{y}_1 - \bar{y}_2)^2}{2(\sigma_{y1}^2 + \sigma_{y2}^2)}\right) \quad (4.1)$$

\mathcal{L} is the interaction rate of particles per unit cross-section. N_{b1} and N_{b2} are the number of particles per bunch per beam, σ_{xn} and σ_{yn} are the beam sizes in the transverse directions cross-section specific to each beam n , k_b is the number of bunches and f_b is the bunch circuit frequency. The coordinates x and y represent the beam position offset at the interaction point. Obviously, we want to minimize beam offset, minimize beam cross-sections and maximize the number and frequency of bunches to increase the instantaneous luminosity.

In the pilot physics run (the second part of what is known as Stage I of commissioning), the accelerator will run at design energy, but with significantly decreased intensity – 10^9 particles per bunch, as opposed to 10^{11} for normal operation. The beam plan will consist of 43 bunches in circulation per beam [32], each running without crossing angle or squeeze. This “pilot beam” is a standard part of normal run [101], providing a verification of the machine state before accepting the more potent and dangerous physics load. The machine cannot accept a physics beam without a pilot beam in place.

The pilot beam is not, however, a source of physics, and it will be run without squeeze or crossing angle for the pilot run. Any events produced will be an afterthought for the majority of

the LHC physics program. The next stage of commissioning will pump the number of bunches coming from the Proton Synchrotron (PS) from 1 to 3 or 4 per PS injection, putting the beam load at 156 bunches per beam – another foreseen low-intensity “safe” beam for machine testing. This will be pushed in intensity until the end of 2008.

4.5.2 Nominal

During Stage II, the LHC will run (it is hoped) with 75 ns bunch spacing (a third of the design frequency, achieved by loading 24 bunches per PS load) and ramps up the squeeze and crossing angle. Expected performance is $10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ and approximately one event per crossing. The maximum performance possible for this stage is $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.

After the 75 ns operation has yielded good results, the accelerator will go to full 25 ns fills (2808 bunches in circulation total) and push the squeeze and crossing angle to yield 50% of nominal intensity. The 25 ns Ops I stage will have a maximum luminosity of $2 \cdot 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$.

After the 25 ns Ops I stage, the accelerator will stop for a second stage installation of components to allow it to progress in luminosity (by increasing current and refining the spot further). Additional collimators will be installed to allow better beam cleaning and protection (necessary to prevent quenches, as clamping the beam tighter is a source of instability in and of itself) and installing further dilution kickers at the beam dump, allowing a longer sweep length on the face of the beam dump to prevent damage to both the vacuum window on the end of the beamline (a stainless-steel-clad carbon-carbon plate) and the dump itself [33]. Once these protections are in place, the LHC will be able to safely achieve nominal luminosity.

4.5.3 Limitations

Beam and beam energy will continue to be limited by current capacity and stability in the magnet systems, and beam-protection system load limits. The magnets are run close to their maxima, skirting the saturation field for the NbTi current elements (where they revert to normal conductors). The beam’s losses must stay within quench limits (8.5 W/m) [34], so collimators must be able to handle the inevitable damage to protect the magnets. Finally, the dilution system

must be able to prevent damage to to the beam dump carbon block – the limit with the present foreseen final system is a beam at full intensity and 7.5 TeV [106].

Chapter 5

The ATLAS Experiment

In observing any physical quantity, knowledge of the instrument one uses to make the measurement is of primary importance. ATLAS is here presented, with an overview of its capabilities, potential, and the tradeoffs made in its conception and construction. A cutaway of the experiment is shown in Figure 5.1

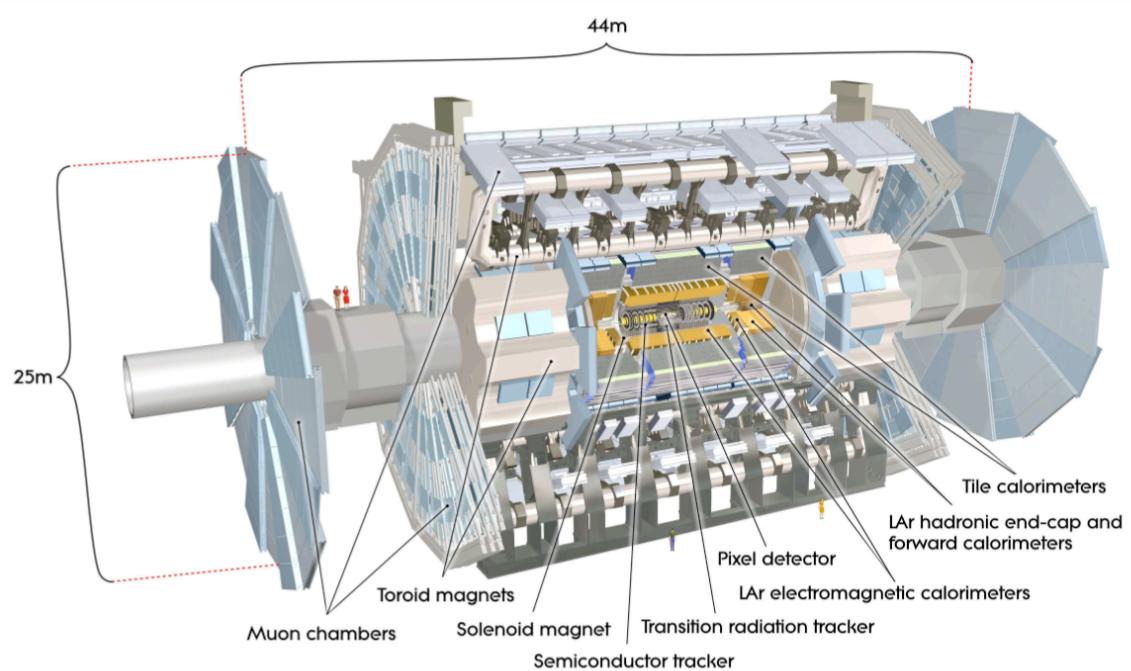


Figure 5.1: Cutaway of the ATLAS detector (perspective view), showing all the major subsystems. Figure from [15].

5.1 Coordinate System

The ATLAS experiment is cylindrical, with its axis of symmetry oriented along the beamline. The accelerator ring itself is tilted 1.4° from the horizontal – the original construction required the deviation to avoid the bedrock of the Jura mountains. ATLAS’s orientation reflects that tilt.

The primary axes are:

z : oriented along the tangent to the beamline circle. Positive z is associated with the “A” side of the cavern (the side closest to the airport and Geneva). Where useful, the coordinates x and y are also cited – positive x is directed toward the center of the LHC ring, and y upward.

ϕ : 2π radians in the plane perpendicular to the z axis.

η : $\eta = -\ln \left[\tan\left(\frac{\theta}{2}\right) \right]$, centered at the point of rotational symmetry in ϕ and the point of bilateral symmetry in z – the halfway point in the cylinder. η is a preferred coordinate because the difference in the rapidity of any two particles is independent of their Lorentz boosts along the z axis. η of 0 corresponds to 90° in θ . $\theta=45^\circ$ corresponds to η of 0.88, and $\theta=10^\circ$ corresponds to η of 2.44. Approaching $\theta=0$, η approaches infinity.

The beams of the LHC are designated as “1” and “2”. Beam 1 enters the A side of the cavern – that is, it rotates clockwise when the LHC tunnel is viewed from above.

5.2 Conception

The ATLAS detector is designed with several principal requirements. First and foremost, it must be feasible and reliable. This has called for proven techniques and technologies to be used as basic components in the design. The principle of proven technology is, however, balanced with the need to use these technologies to the very limit of their possible performance, without sacrificing their durability. The issue of cost is also a weighty one – however attractive a detector technology is, it must also be affordable both in development time and monetary cost.

In physics terms, the detector is optimized for some specific goals. One of the most compelling is a final determination of the status of the Standard Model Higgs particle (or its supersymmetric cousins) – whether to their discovery or exclusion. Subordinate goals include investigation of supersymmetric (SUSY) theories in general, and a deeper probe of the Standard Model in some of its more difficult areas, including the questions surrounding nonperturbative QCD. For these researches, the following characteristics were considered crucial:

- Hermeticity of the detector must be able to account for the missing transverse energy (hereafter E_T^{miss}) of an escaped neutrino (or other unknown, non-interacting particle) in an event to balance the conservation of energy. This is especially critical for some SUSY and Higgs searches. Losses due to detector coverage must be minimized. The colliding particles themselves (protons) complicate the picture – the energy of the partons (quarks or gluons) that actually collide is unknown for any individual collision, though the sum of the transverse energies of the partons is assumed to be zero. We must therefore capture as much of the transverse energy (E_T) as is practical, while keeping the detector intact. In the high- η region of the detector, both beam losses and high particle multiplicities can do significant damage to components placed there. The ATLAS detector covers out to about $|\eta| = 5.0$.
- Tracking must be able to pinpoint the vertices of particles and jets to within the smallest possible precision and error. This will allow the charged particles emanating from multiple collisions at a time (one definition of pileup), each with high multiplicity (or a high number of produced particles per event) that are expected during LHC operation to be parsed correctly and assembled into accurate tracks. Triggering on the tracking is impractical at this point in time because of the enormous data rate from the tracker and substantial processing that must be done on the data before they are usable.
- Calorimeters are not required to be compensating, but do have to maintain an excellent resolution for low- η electromagnetic events to complement the tracking system. There must also be calorimeter components that react in the appropriate timescale to allow an

effective trigger. Coverage of the calorimeter system runs out to $|\eta| = 5$, but the higher- η regions suffer from smaller resolution and higher background. As η increases, so too does the rate at which beam losses, beam backgrounds and collision remnants impact the detector, and the design compromises that allow greater radiation resistance also reduce resolution within the same cost constraints. For the barrel region, the specification is to have a $\sigma/E \sim \frac{10\%}{\sqrt{E}}$.

- Muon spectrometry has accounted for much of the engineering effort in ATLAS. Its signature toroidal magnets are the largest subsystems of the experiment, and act solely on muons which escape the calorimeters, measuring their energy by watching their progress through the magnetic field. There must be subsystems of the muon spectrometer that are capable of responding with data within the trigger allowances as well.
- Correct particle identification is a major factor in the detector's success. The sequence of the detector layers can reveal the different particles that are being tracked and stopped, allowing separation of electrons and photons and pions and muons. Precision in the tracking system can improve identification rates for more ephemeral particles like τ s and b-mesons and top quarks, of which we see only a signature.
- Triggering is the linchpin of an experiment where the data flow far exceeds the tape speed for final data storage. More detail about the trigger scheme in ATLAS can be found in Section 6.1

In practical terms, all of these goals relate to the event rate anticipated from the LHC, which has been designed to produce bunch crossings every 25 ns. Components must be able to resolve and transmit their signals within that period of time. Once assembled, the data must be evaluated for usefulness in a specially-designed trigger system that will separate mundane events from “interesting” quantities – events that contain relevant data to the searches underway. The trigger must be flexible enough to adapt as the real events from the detector dictate. It must also remain fast enough (a competing requirement) to both collect the relevant data from widespread detector components within the time limitations of buffer memories installed on each detector

or readout system, and analyze those data within the 25 ns limitation to avoid backing up the system with events that cannot be analyzed or stored.

The event rate from the LHC also engenders a huge flux of damaging particles around the interaction point [35]. The real total flux has been modeled to range from 10^9 neutrons \cdot cm $^{-2}$ s $^{-1}$ near the beampipe to a few kilohertz in most of the muon system. Charged hadrons of greater than 20 MeV occur at similar rates (in the same order of magnitude). For components close to the beamline, the design must either allow replaceability or invulnerability to these particle fluxes, and must also minimize activation of components to allow safer servicing during shutdowns.

As we move on to the subsystems of the detector, it is important to note that all of them affect at least their nearest neighbors. In many cases the effects of one subdetector's or magnet's operation will have a far-reaching effect on other subdetectors.

5.3 Tracking

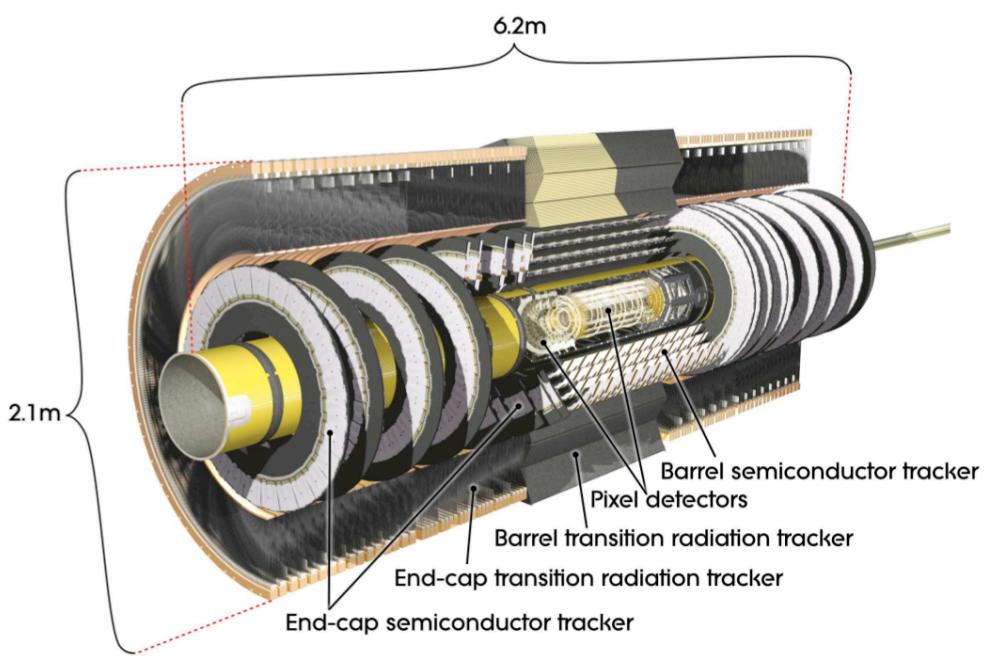


Figure 5.2: A perspective cutaway of the ATLAS Inner Detector, comprising the pixels, SCT and TRT. Figure from [15].

The inner part of the detector serves in the key role of tracking particle trajectories and getting the first part of an event’s “fingerprint” by measuring the charge and mass/energy ratio of its daughter particles. Ideally, the tracker would be completely transparent to these particles and have infinite resolution, being able to perfectly parametrize the particle’s path and starting point, which can tell us what the lifetime of its parent particle was.

In fact, there are (as always) tradeoffs to be made. The “material budget” of the tracker quantizes the deleterious effects of the material we place in the way of the particle on its lifetime. Whereas a photon might speed through empty space without change, it is far more likely to “convert” – become an electron-positron pair – in the strong fields of the nucleus of a high-Z

atom. Charged particles like electrons undergo multiple scattering and *bremsstrahlung*¹ as they pass through matter, and these losses hinder measurement. Figure 5.3 shows the radiation and interaction lengths of the inner detector as installed.

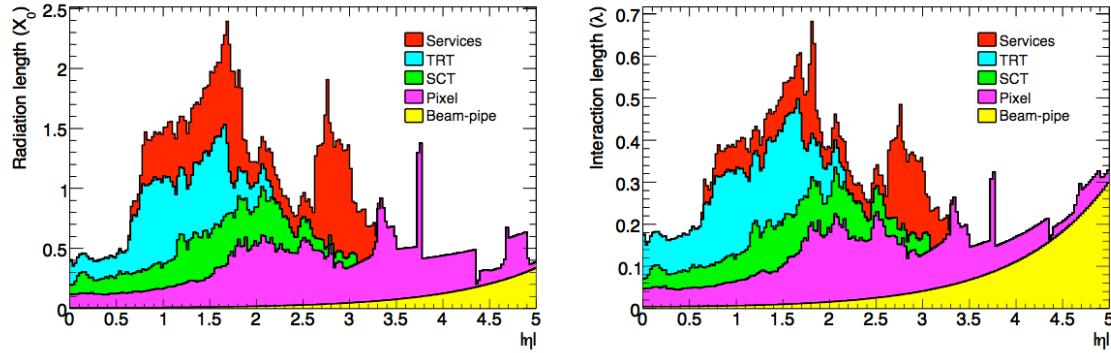


Figure 5.3: Material distribution in the ATLAS ID. On the left, the radiation lengths of the various ID components, and their interaction lengths on the right. Figure from [15].

The loss of energy per distance in matter is related to a quantity called the “radiation length”, a scaling variable for the probability of scattering or *bremsstrahlung*, which is defined as follows: [36]

$$1/X_0 = \frac{4\alpha N_A Z(Z+1) r_e^2 \log(183Z^{-\frac{1}{3}})}{A} \quad (5.1)$$

where

- α = the fine structure constant ($\approx \frac{1}{137}$);
- N_A = Avogadro’s Number, $6.022 \cdot 10^{23}$ /mole;
- Z = atomic number of the traversed material ;
- A = atomic weight of the traversed material; and
- r_e = classical electron radius ($2.818 \cdot 10^{-13}$ cm);

¹Electromagnetic radiation produced by the acceleration of a charged particle after passing through the electric and magnetic fields of a nucleus.

Measures clearly indicated to avoid squandering the material budget are:

- Avoid high-Z materials assiduously (and with them, we manage to handle the A factor as well); and
- Minimize material used in sensors, supports, cabling, cooling and insulation.

The X_0 designation also applies to photons – the likelihood of a photon undergoing pair production is $7/9X_0$. Photons, however, do not undergo parasitic energy loss like electrons, which convert at an average rate of

$$-(dE/dx)_{brem} \doteq E/X_0. \quad (5.2)$$

With a material budget for the inner detector that varies from $0.3X_0$ to over $1.5X_0$ within the envelope (115 cm in radius) of the calorimeter cryostat (averaging under 1.0 – the peak is located at the deepest part of the cylinder as seen from the interaction point, where the services are also concentrated to be able to exit the detector by way of the endcap crack), photon conversion will be common and compensation must be made.

5.3.1 Pixel

The pixel detector [37], pictured in Figure 5.4 is at once the most sensitive and most vulnerable detectors in the ATLAS ensemble. It sits 5 cm from the interaction point, and experiences the harshest particle fluences in the experiment’s barrel section. It is designed for a 10-year operation in the outermost layers, but the inner layer (known as the B-layer, for its value in B-vertex discrimination) will only survive 3 years at luminosities of $10^{33} \text{ cm}^{-2}\text{s}^{-1}$, plus one year at $10^{34} \text{ cm}^{-2}\text{s}^{-1}$, before the necessary depletion voltage (described below) on the silicon sensors exceeds design limitations. At that point, it becomes a material liability until such time as the inner detector can be upgraded or the pixel B-layer replaced. This is less of a loss at high luminosities, however, where the charged particle multiplicity limits possible resolution in any case. The pixel structure is also carefully designed to minimize expenditure of material budget – the cabling, support and cooling infrastructure has been built of low-Z materials

with high rigidity and predictability. Services have also been minimized by mounting chains of pixel modules onto staves² and disks, and using the stave or disk itself as a mechanically rigid alignment unit and cooling platform at the same time.

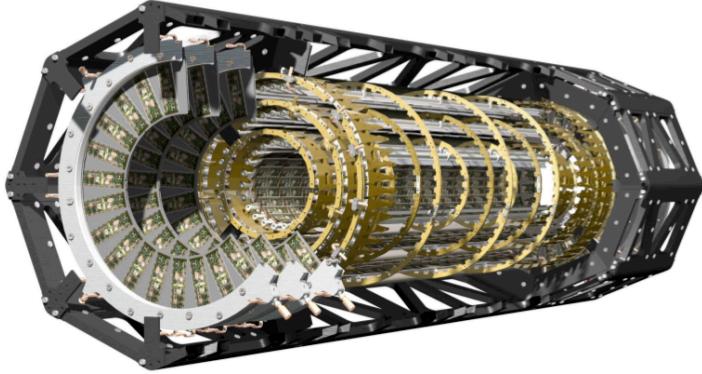


Figure 5.4: The pixel subsystem, barrel and endcaps, including its support frame. Figure from [15]

Construction Beginning with the modules themselves, we trace the interlocking solutions to the competing requirements of rad-hardness, transparency and sensitivity. When high-energy particles traverse the silicon bulk, the regular lattice of the crystal undergoes distortions as atoms are knocked out of place, compromising the insulating properties of the n bulk – effectively an increase in doping in areas that must not be doped for correct performance. The depleted state of the sensor becomes progressively harder to maintain without unreasonable voltages, and the detector loses discrimination power, eventually becoming altogether useless. After multiple iterations with prototype processes, the decision was taken to use the n^+ on n technique for the sensor, which allows operation even under partial-depletion circumstances. This process is more expensive and difficult because it requires dual-side processing. Diffusion oxygenation is also applied to the sensors to reduce the effective doping, optimize the annealing process

²A “stave” is a mechanical unit of shared mechanical structure and infrastructure for 13 pixel modules.

(which compensates for the lattice damage done by simply warming and cooling during detector accesses), and minimize the “reverse annealing” process.

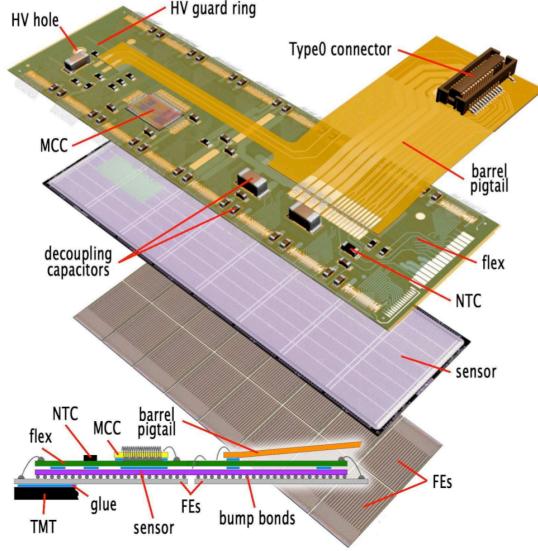


Figure 5.5: An exploded perspective view of a pixel module. A side view of the module illustrates the bump-bonding technique used to join the hybrid to the FE chips. Figure from [15].

The modules themselves, once built, consist of 16 individual chips, each containing 2,880 pixels of (usually) $50 \times 400 \mu\text{m}$ in size – see Figure 5.5 for an illustration. There are some variations in the interchip regions, where long or “ganged” pixels are used to stitch together the gaps. Each chip has its own Front-End (FE) chip for readout, command decoding, clock, Level-1 trigger signals and synchronization timing information. The FE chip is attached via bump-bonding, where indium or solder balls with a $25 \mu\text{m}$ pitch are deposited in a penetration through the passivation layer on the silicon, then used to bond the pixels themselves directly to the FE chips on every channel. This is necessary to fulfill the speed requirement – with 25 ns to read the whole pixel system, chained readout would be prohibitively slow. Once the data are collected on a clock cycle, FE chip buffers store the data (a timestamp and Time Over Threshold (TOT) value) for up to $6.4 \mu\text{s}$, until the Level-1 (LVL1) decision can be made. The 16 FE chips are all operated by a Module Control Chip (MCC) [38] on a star topology (increasing fault tolerance in individual FE chips), which takes care of all module-level operations, does data

compression and event building at the module level, and feeds the optical link events that have been marked by the LVL1 system for further processing. The modules are themselves arranged in a shingled configuration (20° from normal) in barrels positioned at 50.5, 88 and 122.5 mm from the IP [39]. The total number of modules in the barrel is 1,456. The endcaps are mounted in a disc pattern, in eight sectors per disc, six modules per sector, for a total of 288 endcap modules and 1744 modules total.

Cooling The power use of the pixel system (about 4 W per module) sits around 3500 W under normal operation. There are several levels of operation, and the electrical consumption is clearly different for a configured module as opposed to one with no processing load, but 3500 W is the usual heat profile. All of this heat is being produced in a volume the size of a 12 year old child, whose average heat production might be closer to 80 W. The detector is also in an insulated chamber, both to keep changes in its temperature from affecting other detectors, and vice-versa. The modules are least sensitive to effective doping at a temperature of -7° C, which is their planned operating temperature. When planning the system, therefore, a low-Z, high-volume cooling system had to be brought into play. The cooling plant will have a capacity of 10 kW, and will feed coolant (C_3F_8) at a temperature of -25° C through cooling circuits shared between pixel staves. The cooling pipes are integrated directly into the carbon-carbon staves, taking advantage of the mechanical stiffness of the material and its high thermal conductivity. After some experience with the cooling pipes, it was discovered [37] that there were corrosion issues due to a galvanic reaction between the aluminum of the cooling pipes and their nickel plating which were corroding when combined with atmospheric moisture. 43 loaded staves were repaired with a glued pipe insert, which reduces the thermal efficiency of the stave by approximately 10%. Such repaired staves are inserted in the outer barrel (where radiation damage is lower and operating voltages will grow less quickly, and the occupancy is lower, leading to less electronics heating) and paired in the bi-stave cooling loop with a full-capacity stave. The rest were repaired before loading, using a new aluminum compound and laser welding.

Performance Performance in the pixel system can be broken down into the categories of maximum occupancy, resolution and alignment. Occupancy is a principal design concern for a vertex tracker in the LHC. At smaller luminosities and pileup levels, the likelihood of hit degeneracies is small, and the data rate (derived from the absolute number of particles traversing the planes of the detector) is well within reason. When the luminosity is sitting at nominal, however, and the LVL1 trigger starts demanding a 100 kHz event rate, real limitations exist on the ability of the system to read data off the front end [38]. The detector occupancy cannot be reduced – there will be 500-1,000 [115] charged tracks per bunch crossing, and with at least three hits per track, the requirement to get the data off before the buffers fill is a difficult one. The links from each module run at a selectable rate of 40, 80 or 160 MHz [38], and at full luminosity the hit loss rates from the B-layer are necessarily high for the lowest transfer rates. Occupancy of the pixel system [38] will average 0.004% – it’s the readout system that will have to adapt to adverse circumstances. The track and impact parameter resolution of the pixel system is determined by several factors – its known position, the noise and threshold performance, the pixel size itself, and the Lorentz angle of the modules in the 2 T solenoidal magnetic field (which modifies charge migration paths in the sensor) [40]. With increasing depletion voltage (and therefore increasing electric field across the sensor), the drift effect decreases [37]. Final track resolution in on the fine (x) axis is expected to be $7\text{ }\mu\text{m}$, and $62 - 67\text{ }\mu\text{m}$ on the long axis (y), depending on angle of incidence. The y resolution is enhanced if a cluster of hits is formed rather than a single hit, because the “bricked” structure of the pixel layout allows correlations to be drawn (through charge sharing effects) about the hit location. Alignment will be performed in collaboration with the SCT, taking the residuals (the difference between the track position and the hit position) from the detectors and the produced tracks in the reconstruction software. These residuals can be calculated either with or without the hit included in the track reconstruction (biased and unbiased residuals, respectively) [41]. Overlapping residuals from two different modules allow correlations between a module and its neighbors. The calculations for the displacements of individual modules are sensitive to various details of their construction. 200,000 events suffice for a good set of alignment coefficients, and take about five hours to run

across a group of ten machines[41]. Final coefficients will have an uncertainty of $1 \mu\text{m}$ in local x and $40 \mu\text{m}$ in local y . This data sample should be available within the first seconds of LHC running, and preliminary work is already underway using cosmics rays that traverse the cavern.

Readout This topic is treated at length in Section 8.2.

5.3.2 SCT

The SemiConductor Tracker (SCT) follows a similar design philosophy to that seen in the pixel subdetector, but is conceived to give more space points in the outer silicon tracking volume with reduced cost and dataflow requirements.

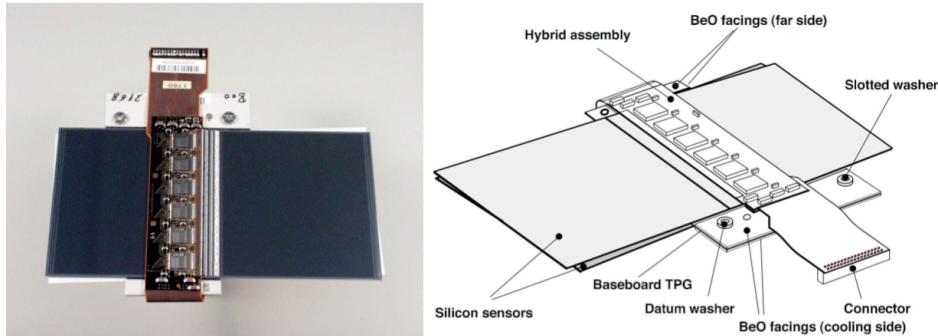


Figure 5.6: An SCT module. The TPG (thermal pyrolytic graphite) baseboard serves as a low-Z structural material of great stability, and as an excellent heat conductor for uniform cooling. Figure from [15].

Construction To simplify the sensors themselves, reduce processing costs and increase yield, the choice of 768 strips per module, arranged p-in-n substrate, allowed single-sided processing and reasonable radiation tolerance. A maximum bias of 500 V was chosen to allow operation through the foreseen ten years of ATLAS operation [42]. Because strips give only a 2D line in space, the sensors were arranged back-to-back and glued at a 40 mrad angle to each other to form a module, allowing the formation of space points with a reasonable resolution in the

ATLAS z -axis (for the barrel). The modules themselves (seen in Figure 5.6 are 64 mm wide and 128 mm long (each strip is divided in two, and ganged together directly). Each individual strip has a pitch of $18 \mu\text{m}$ and is spaced $80 \mu\text{m}$ from its neighbor. The endcap modules are made in a keystone shape to allow better tiling coverage when endcap discs are assembled, but the overall dimensions are similar to the barrel modules [75]. There are 2,112 barrel modules and 1,976 endcap modules, arranged in four barrels at radii of 300 mm, 373 mm, 447 mm and 514 mm (seen in Figure 5.7), and endcap discs at z positions of 853.8, 934, 1,091.5, 1,299.9, 1,399.7, 1,771.4, 2,115.2, 2,505 and 2,720.2 (all \pm , and all in mm from the IP) [75].

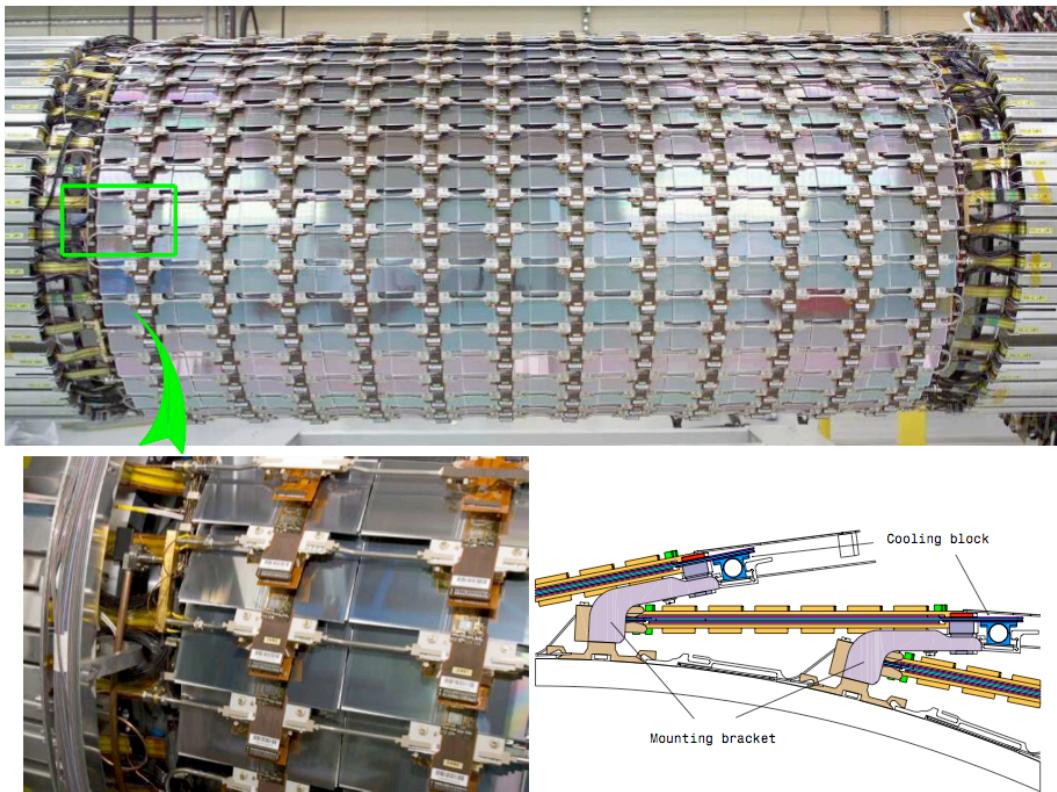


Figure 5.7: An assembled SCT barrel. Details of the module mounting are shown below, illustrating the “shingled” nature of the module mounting and the integration of the cooling system. Figure from [15]

The stereo angle is made small (reducing the possible resolution along the strips) to avoid the formation of “ghost” space points, which arise from ambiguities in the crossing strips. Let’s take, for example, strips placed at 90° to each other. Each “front” strip crosses 768 other strips on the module behind it, meaning that each intersection between a “back” strip and our front strip can form a space point. If there are two hits in the same module, given the right-angle configuration, it is guaranteed that for possible spacepoints will be formed, and that there will be no way of distinguishing the spurious from the valid. At a 40 mrad crossing angle, however, any given strip crosses a maximum of six others ($512 \mu\text{m}$), reducing the number of possible ghosts dramatically – they can only form if there is a second hit within the six overlapped strips.

The SCT is designed to have a small material budget and small cooling footprint. The on-detector processing is minimized, and there are no TOT or analog measurements possible. The cooling block on the carbon-carbon mounting substrate is connected directly to the chip hybrid to minimize heating on the module itself, and there are no thermal bridges other than the bond wires that connect the chips to the silicon.

Like the pixels, the SCT is sensitive to temperature fluctuations, and needs to be cooled to -7° C to maximize its operating lifetime. The cooling system, again similar to the pixels, uses C_3F_8 in liquid form to cool the modules. By the end of its loop, much of the liquid has gone gaseous – extra fluid is allowed to prevent temperature gradients that arise from different phases of coolant in different parts of the detector. When the fluid has passed out of the SCT, it is vaporized in a heater system to allow continued flow.

Readout This topic is treated at length in Section 8.2.

5.3.3 Transition Radiation Tracker (TRT)

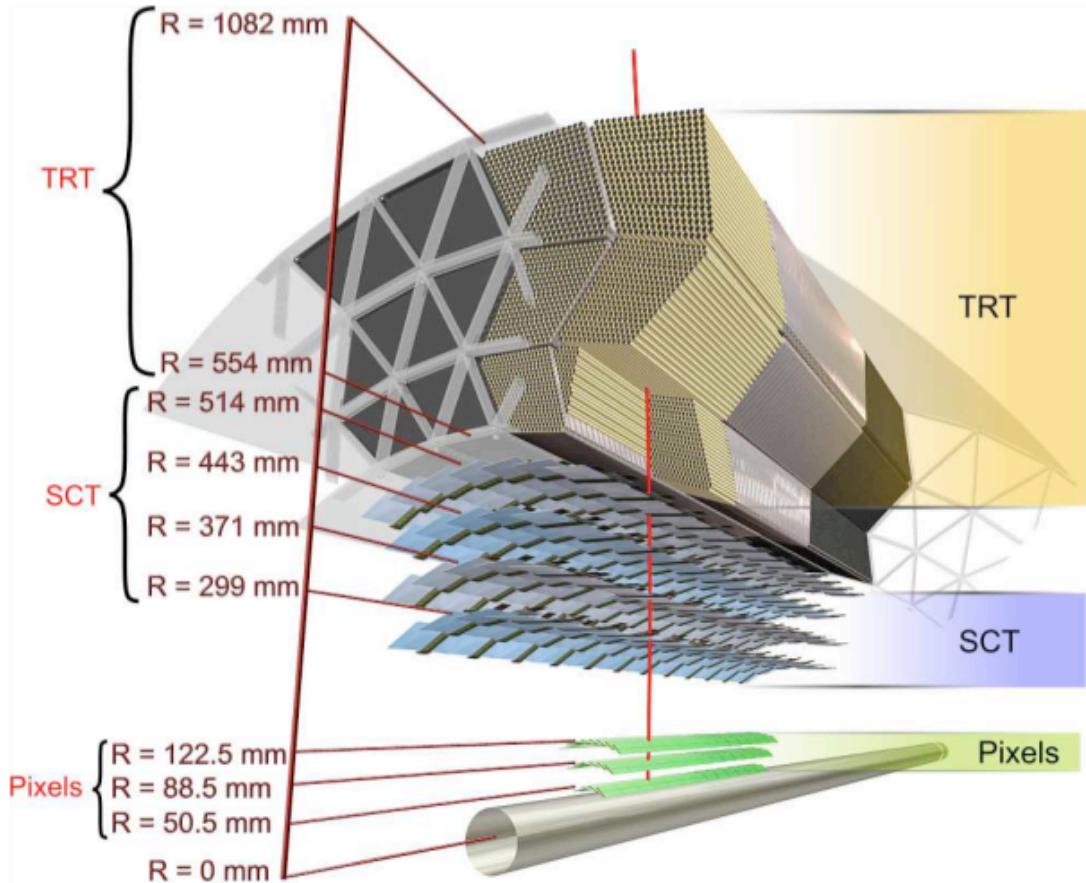


Figure 5.8: A 10 GeV p_T charged particle track traverses the various ID components. This figure illustrates the relative position, modular configuration, and straw orientation of the TRT barrel relative to other Inner Detector (ID) components. Figure from [15].

Construction The parts of the ID beyond the SCT is called the Transition Radiation Tracker (TRT). They are closely spaced, but have a thermal separation due to very different temperature requirements. The TRT consists of layers of 4 mm carbon-polyamide straws, each with a fine ($30 \mu\text{m}$) wire strung precisely down its central axis, supported in the middle by a guide, and

surrounded by xenon gas (70%), CO_2 (27%) and O_2 (3%). The xenon is a good x-ray absorber, and the oxygen and CO_2 are useful photon quenchers, and increase the electron drift speed.

Performance Electron drift times range from 20-70 ns – or from one to three bunch crossings. While this disqualifies the TRT from the LVL1 trigger, it will be included in the LVL2 trigger decision. The number of straws (370,000 in the detector as designed) allow a significant number of hits per track – 36 on average – and give a radial positional precision of $170 \mu\text{m}$ per straw hit, as drift time measurements give a location within the straw for the points of closest and furthest approach.

e/π Separation Not only does the TRT allow tracking enhancements, it offers enhanced particle identification capabilities. Since electrons, as strongly ionizing particles, create a wake of transition radiation (TR) photons as they pass through material boundaries, they can be distinguished from minimally-ionizing particles (MIPs) such as pions (which do the same thing to a much lesser extent). To emphasize the difference, foils of polypropylene $15 \mu\text{m}$ thick are spaced $200\text{-}300 \mu\text{m}$ apart by either regular or randomized polyethelyne fiber nets. Both the spacer and the foil enhance the number of TR boundaries that the particles will pass through, allowing an electron's signature to be more distinct while minimally affecting a MIP.

Occupancy One major concern with the TRT is its high-luminosity usefulness. Charged particle multiplicity during a high-lumi collision may be very high (13% to 38%, depending on the TRT region), and lead to saturation of the TRT because of its coarse granularity. The preparatory TDR studies have concluded that while there will be some loss of drift-time (and therefore positional) accuracy under these circumstances, the trigger and pattern-recognition algorithms are unlikely to suffer serious degradation at design luminosity [43, 76, 128].

5.4 Solenoid Magnet

The tracking volume is held at a magnetic field of 2 T by a single-layer NbTi solenoid magnet, shown in Figure 5.9, which shares a cryostat with the LAr detector (to reduce the amount of material in front of the calorimeter). It is relatively lightweight (7.4 kJ/kg), stores 40 MJ, and is quench-safe – the cold mass is designed to be able to absorb all the energy of a quench without safety systems, and come back to operating temperature within a day.



Figure 5.9: The ATLAS inner detector solenoid, before insertion into its cryostat.

5.5 Calorimetry

Once the results of the collision have been tracked and identified to the extent possible, the next step is to stop what can be stopped and evaluate its energy (and if one can do so by inference, its particle ID as well). Two classes of calorimeter are used in ATLAS – the

Electromagnetic Calorimeters (ECAL) and the Hadronic Calorimeters (HCAL), both of which are non-compensating calorimeter implementations [44, 48] (they respond more strongly to electromagnetic (EM) particles than to hadrons. The layers serve, rather, as filters, removing the EM particles first, then the hadrons. Depending on the η position, and therefore the radiation environment, of the calorimeter, various technologies are used – but the fundamental nature of the two calorimeters is to present a selective barrier to various kinds of particle. Their relative positions and geometries are seen in Figure 5.10

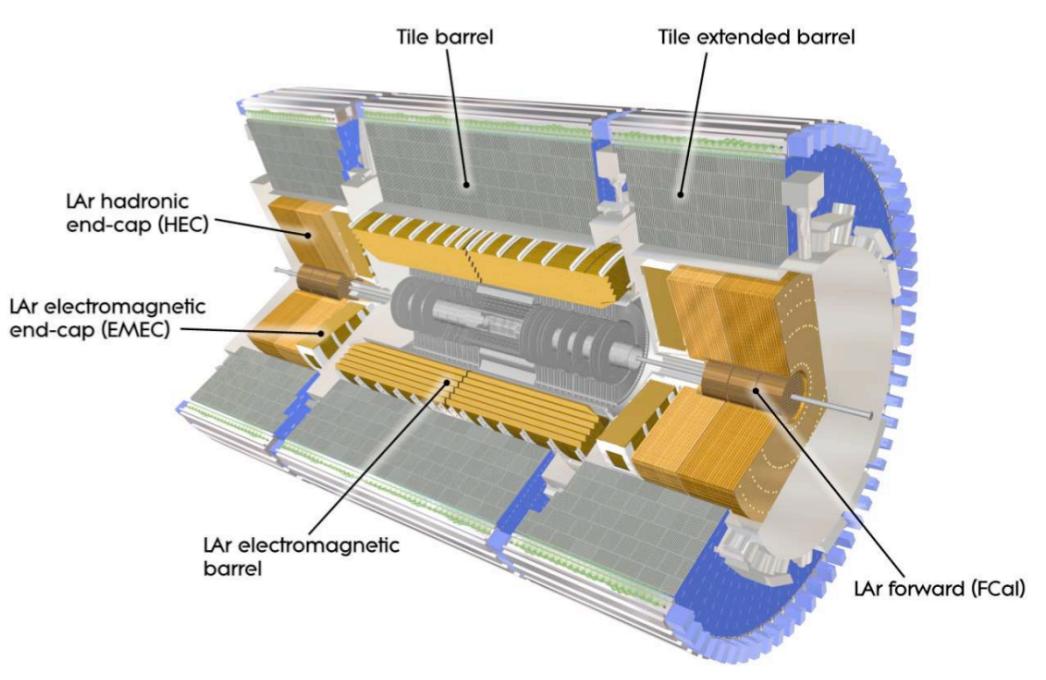


Figure 5.10: The calorimeter subsystems, comprising the LAr electromagnetic calorimeter (barrel and endcaps), the tile and endcap hadronic calorimeters, and the forward calorimeters. Figure from [15].

5.5.1 Electromagnetic Calorimeter

The ECAL is key to various particle searches in ATLAS, and is therefore of great interest. The design allows for construction of a calorimeter with minimal cracks and excellent energy

resolution. It is also technically challenging, expensive, and very difficult to simulate due to the complex geometry of its volumes.

ECAL Construction The aforementioned complex geometry is one of the primary features of the ATLAS Liquid Argon (LAr) EM calorimeter. One challenge in constructing a hermetic detector is the problem of joining modules together into an integral whole. In the silicon tracker, where surface is the real issue, the detector modules are shingled, one overlapping another, to close all the gaps. In building a calorimeter, the gaps themselves become a problem, as there will be energy recovery asymmetries to any particle shower that winds up passing through dead material – the energy is lost, and with it efficiency. The “accordion” design of the LAr EM calorimeter is a powerful way to remove that inefficiency, building a gapless module arrangement. Since this is a sampling calorimeter, a passive absorber material will be incorporated to add radiation lengths ($24 X_0$ in the barrel, $26 X_0$ in the endcap) in a consistent way. Active material, permeating the gaps in the absorber, allows cascades from incident particles to be created and detected.

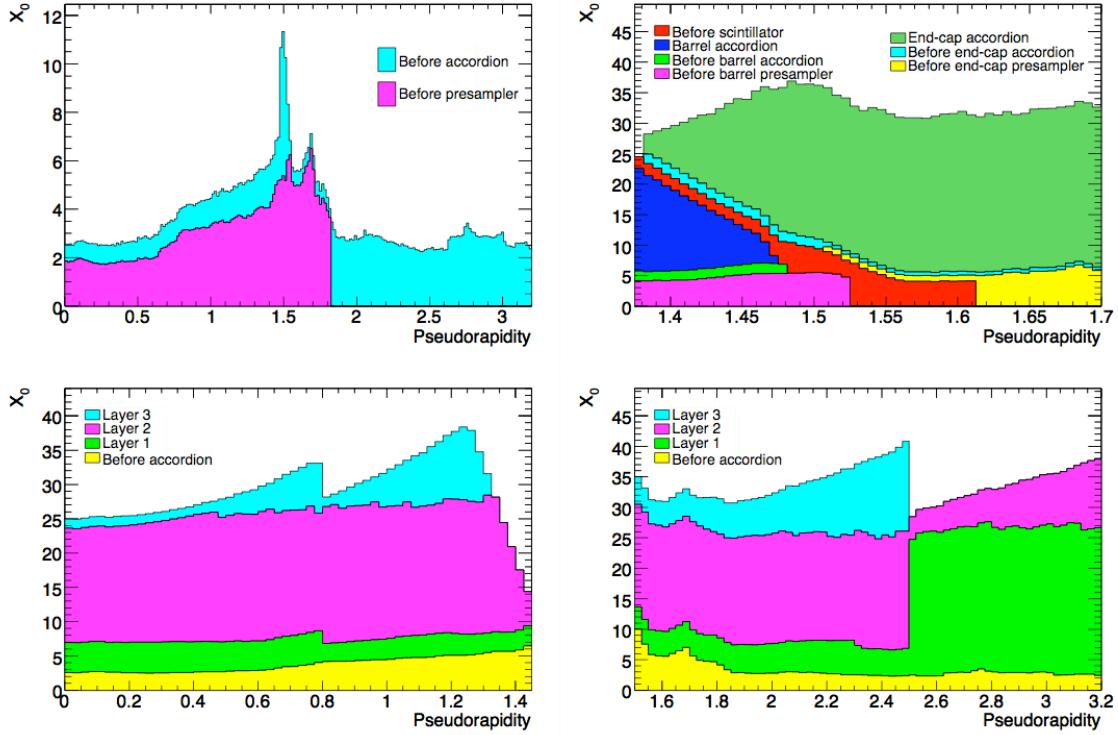


Figure 5.11: Calorimeter material distribution as a function of η . On the left, thickness is expressed in radiation lengths – on the right, in absorption lengths. Figure from [15]

The absorbers, 1.5 mm accordion-folded sheets of lead, (1.1 mm after $|\eta| = 0.8$) are clad in 0.2 mm of steel for rigidity. They are interleaved with segmented copper-Kapton electrodes (separated by honeycomb-style spacers). The electrode segmentation allows various parts of the calorimeter volume to be read in layers (called “samplings”) of decreasing resolution, as seen in Figures 5.12 and 5.13. It is also segmented in consistent units of η .

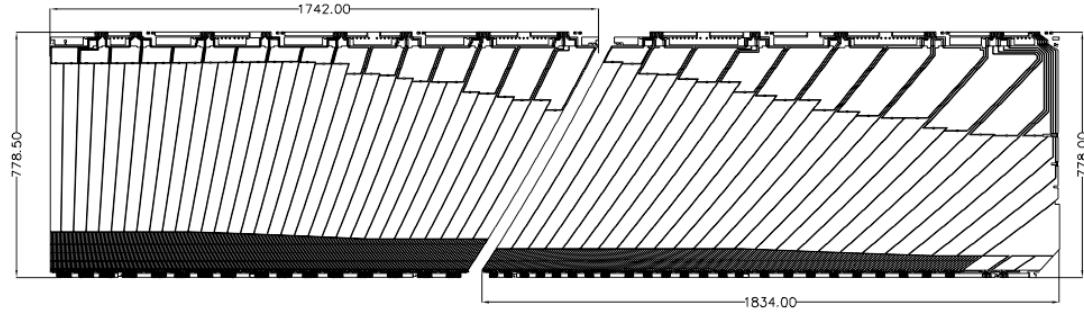


Figure 5.12: The electrodes from the barrel and extended barrel EM calorimeters. The samplings of different granularities are clearly visible in the up-down direction, and the consistent- η segmentation is seen from side to side. Separate segments of these electrodes are ganged with their counterparts on other electrode sheets to form the calorimeter cell structure, illustrated in Figure 5.13. Figure from [15].

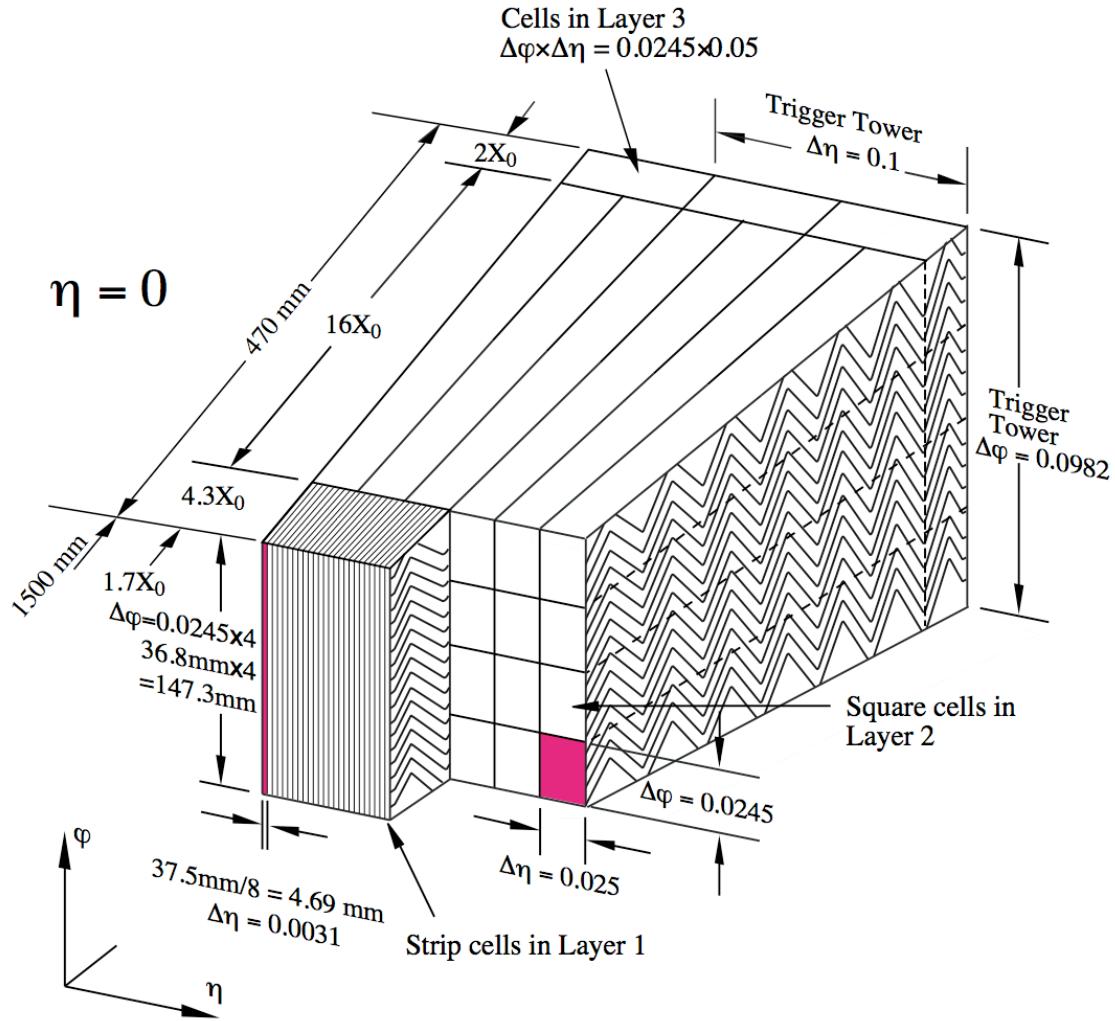


Figure 5.13: Illustration of the cell structure in the LAr EM calorimeter. Also seen is the very coarse ganging of cells into trigger towers to speed trigger response. Figure from [15].

The modules are layered together into a ring of fully-overlapping absorbers and electrodes, mechanically stabilized and connected to services. Placed in a cryostat, it is then bathed in high-purity liquid argon, which serves as the active material – something easy to renew after radiation damage, and which has shown itself to be a good showering medium. The barrel and endcap EM calorimeters are built along the same lines, and the only cracks in the system are at $\eta = 0$ and $1.4 < |\eta| < 1.6$, where the calorimeter endcaps and barrel meet (and through which

the inner detector and LAr service pass). Presamplers are also placed in front of the calorimeters (especially important in the crack) to estimate losses from showering the ID and crack/services material. The LAr calorimeter’s performance is not a primary focus of the present work, but it is worth mentioning that the energy resolution of the LAr calorimeter is expected to be $\sim 10\%/\sqrt{E}$, with an expected maximum resolvable energy of 3 TeV [45]. Another salient issue is that because of the speed of electron propagation in the LAr, integration time for most of the LAr system is ~ 450 ns, corresponding to 18 bunch crossings. Data collection is therefore done using advanced pulse-shaping techniques, and sampling the leading edge of the signal pulse.

5.5.2 Minimum Bias Trigger Scintillators

The minimum bias trigger scintillators (MBTSs) are part of the tilecal system, wedged between the endcap cryostats and the inner detector. Its specific geometry can be seen in Figure 5.14. Constructed of polystyrene doped with a fluorescing agent, the 16 trapezoidal wedges (two segments in η , eight azimuthally) on each side of the ID provide high- η ($2.12 < |\eta| < 3.85$) response to diffractive jet signatures, as illustrated in Figure 3.1. The 3-cm tiles are covered with a boronated polyethylene neutron moderator, and clad with 3 mm of aluminum. Readout is done through fiberoptic cables connecting to the standard photomultiplier tubes and fast readout electronics used in the Tile Calorimeter (see 5.5.3), which communicate the MBTS response to the LVL1 trigger Central Trigger processor (see 6.1.1) for inclusion in the global LVL1 trigger. The trigger pulse itself is carried by several long cables to the patch panel. These Low-Voltage Differential Signaling (LVDS) pulses are then converted into single ended signals, as transmission distance is no longer an issue. 32 discriminator channels will then be used to produce 32 logic signals. These signals will be first fed into the CTP as 32 NIM (Nuclear Instrumentation Module) individual inputs. After commissioning, this will change to a pair of 5-bit multiplicities (0-16 each side). Two basic triggering strategies are expected to be used in the MBTS system – double-coincidence anywhere in the system (MBTS_2) or opposing coincidence – at least one MBTS segment over threshold on each side. An examination of the trigger efficiencies for

normal minimum bias events, including the inherent trigger bias caused by response and acceptance limitations, can be seen in [46]. A detailed look at their behavior under conditions with beam contamination can be found in Chapter 10.

Recent improvements to the system were made during a period of refurbishment to the Tile Calorimeter, which allowed the MBTS readout 3-in-1 cards to use high-gain mode in communication with the LVL1 trigger, rather than the low-gain LVL1 which is the Tile Calorimeter default. The standard readout automatically switches between high-gain and low-gain depending on signal strength. The resulting signal-to-noise ratio in the LVL1 trigger response is expected to improve by a factor of five, though this remains untested as of the date of this writing. A more detailed treatment of the subsystem, including detailed response testing and system characterization, can be found in [47].

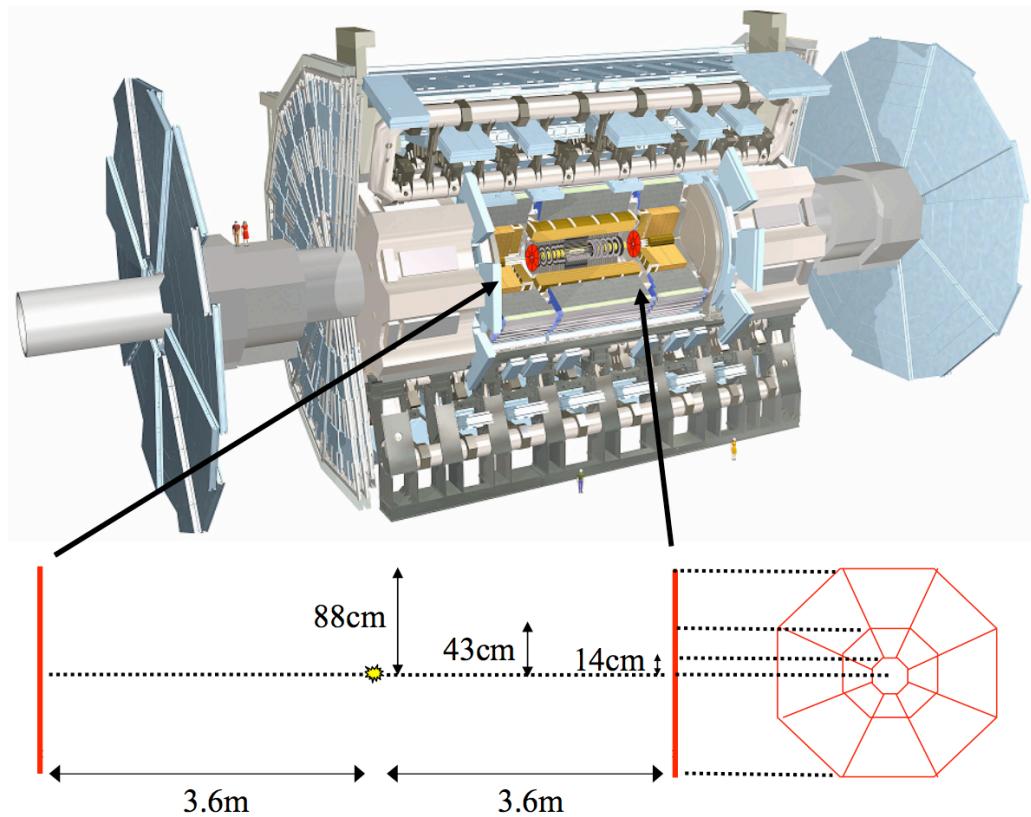


Figure 5.14: Perspective cutaway of ATLAS, showing the position and geometry of the MBTS scintillators, shown in red at either end of the inner detector, mounted on the faces of the calorimeter cryostats. Figure from [47]

5.5.3 Hadronic Calorimeter

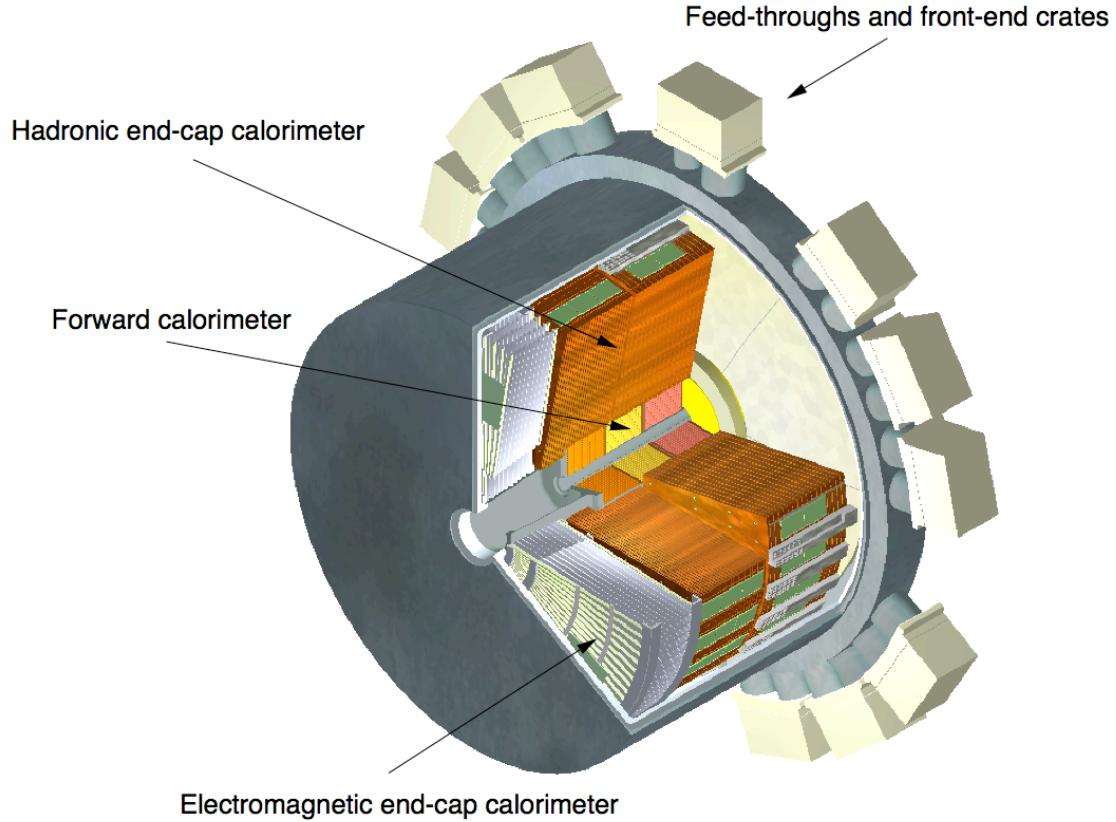


Figure 5.15: Perspective cutaway of the endcap calorimeter cryostat. Figure from [15].

The Hadronic Calorimeter catches whatever passes through the EM calorimeter (barring muons, neutrinos and any weakly-interacting particles) – usually some small spillover from the ECAL or hadrons from the IP. The barrel HCAL [48] ($|\eta| < 1.7$), with its distance from the IP and minimal exposure to the hard forward particles, is built of a less expensive and easier-to-fabricate combination of iron and plastic scintillator tiles. The doped scintillator passes its signal photons up fiber optics to the readout electronics on the outer layer of each calorimeter section. The calorimeter is built in 64 sections to allow modular assembly in the ATLAS cavern

- both the barrel and endcap hadronic calorimeters would be far too heavy for the cranes used to lower detector components.

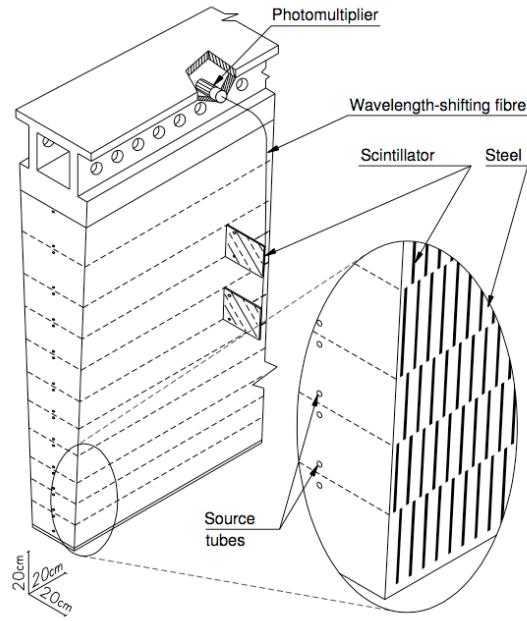


Figure 5.16: A single tile module (one of 64 in each of the barrel sections). Figure from [15].

Inside the barrels of the tile calorimeter lie the cryostats for the entire LAr system. We have examined some components of the LAr calorimeters already – the barrel and endcap EM calorimeters. In the same cryostats as the endcap EM calorimeter (nestled inside and behind it) are the LAr hadronic calorimeter and the forward calorimeters (see Figure 5.15).

The engineering constraints on detectors in this region have been discussed previously, in Section 5.3. To compensate for active material radiation damage in the detector, the same LAr technology is used. The hadronic LAr endcap calorimeters run from $1.5 < |\eta| < 3.2$. The absorbers, in this case, are composed of finely-machined copper plates, and EST (electrostatic transformer) readout structures are placed in the gaps, just as in the EM calorimeter, and spaced using similar honeycomb composites – a cutaway of which is seen in Figure 5.17.

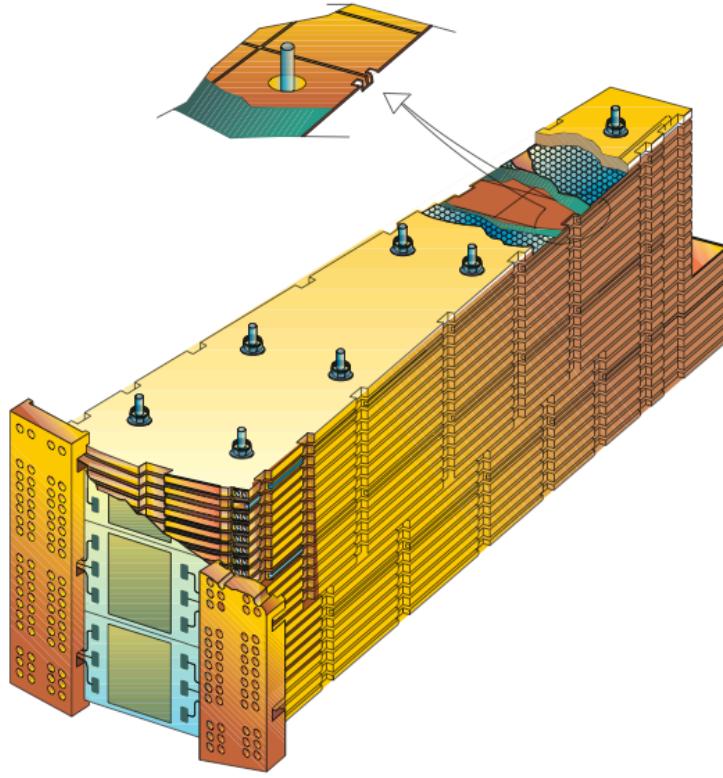


Figure 5.17: A HEC module in schematic view, with a cutaway showing the readout structure, honeycomb spacing, and pad electronics. Figure from [15].

The forward calorimeters (FCAL) are also LAr calorimeters. The integration of these three subsystems into a single cryostat (because of their shared active material) is a significant advantage. Dead material and gaps can be minimized because there is no need for environmental separation between the three systems. One disadvantage to LAr calorimeters in general, however, is that cryostat access is extremely difficult – the subsystems need to be extremely reliable and resistant to damage. The FCAL is placed in the most demanding environment in the experiment, and as such is built in a simple and robust manner. The materials list is brief, the density and Z of its components high, and it maximizes its radiation thickness with respect to the IP. The first layer of FCAL is made up of a copper matrix and copper sense electrodes, while the other two are tungsten matrix/copper electrode structures. The sensing electrodes themselves

are copper tubes with a copper rod down the center. Because the environment in which the sensing is done is itself very noisy, one design constraint is to maximize absorber material. The central rod, therefore, is almost the same diameter as the electrode tube, allowing a $250 \mu\text{m}$ gap for the LAr active material. The gap is maintained by a wound fiber (quartz for radiation resistance). The copper matrix is made up of stacked and drilled copper plates, and the tungsten sections are made of an array of formed and sintered slugs (Figure 5.18, held in place by the electrode array itself).

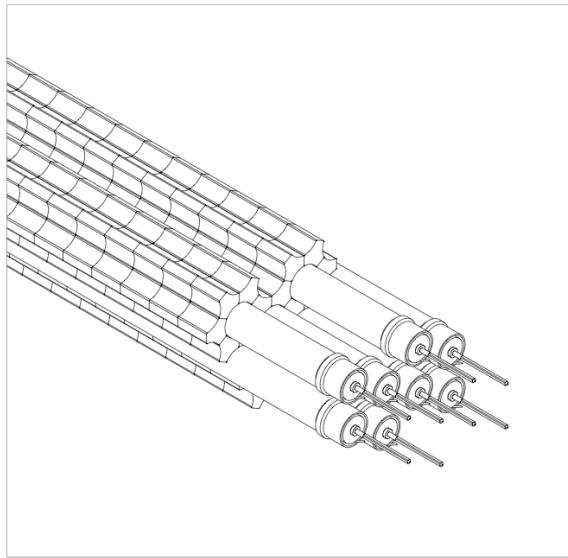


Figure 5.18: Detail from the tungsten FCAL structure, showing the detector tubes (with embedded rods) in their matrix of tungsten support/absorber. Figure from [15].

The energy resolution (in the most simplistic approximation) of the hadronic calorimeters is expected to be around $\sim 50\%/\sqrt{E}$ plus some corrections. The FCAL response time is substantially faster than that of the other LAr calorimeters (due to its $250 \mu\text{m}$ gap), cutting drift times substantially, to ~ 50 ns.

5.6 Muon Detection

Anything that passes the $12 - 22$ absorption lengths' worth of material in the calorimeter system falls into one of three categories, all long lived – muons, neutrinos and unknown weakly-interacting particles. The first is the only one we can count on reliably detecting, so the calorimeter system is succeeded by a muon spectrometer. The basic components of the muon system itself are tracking chambers layered inside and outside a toroidal magnet system – a barrel toroid ($|\eta| < 1$) and two endcap toroids ($1.4 < |\eta| < 2.7$). Muons in the intervening gap are handled by a superposition of the two fields. Since the muons are charged particles, they leave tracks in the inner tracker as well, allowing precision vertex measurement.

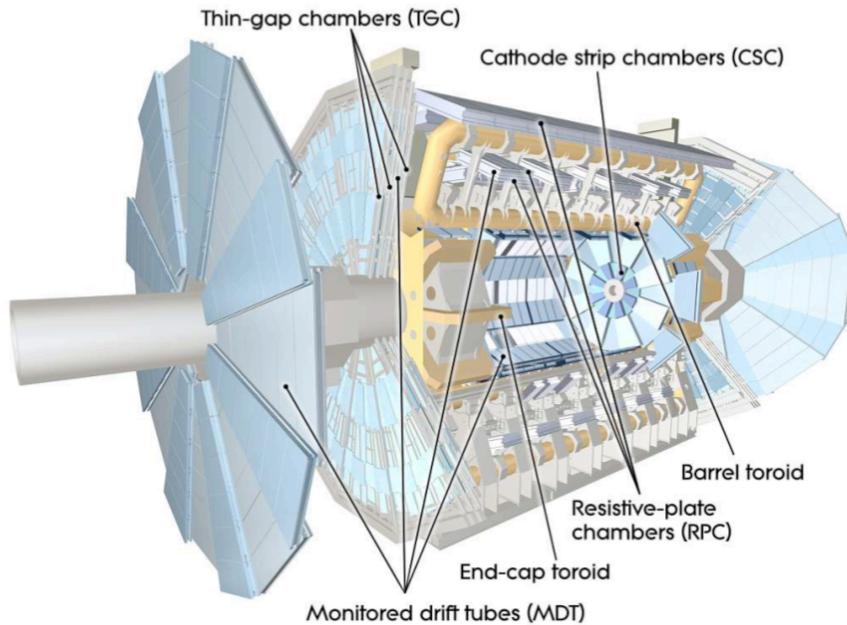


Figure 5.19: Perspective cutaway of the muon spectrometer. Figure from [15].

The muon spectrometer is made up of various different technologies; we'll take each in turn. The first two are designed for resolution in η , to measure the curvature of the tracks.

An overview of the placement and geometry of the muon subsystems is seen in cutaway in Figure 5.19

5.6.1 Monitored Drift Tube (MDT) Chambers

The MDT chambers serve as the most common component of the muon system. Comprising 370,000 channels in 1194 modules, and covering 5500 m², it is clear that price and ease of manufacturing these modules are crucial issues. Coordinates other than the bending radius are considered lower priorities. Those are aided by the Resistive-Plate Chamber (RPC) and Thin-Gap Chamber (TGC) measurements, which are fast but spatially imprecise, providing a good timing measurement and correlating to the precise radius measurements from other subsystems. The MDTs are built of pairs of 3- or 4-layer stacks of 30 mm aluminum tubes, with a 50 μ m sensor wire down the center, and filled with a Ar-CH₄-N₂ gas mixture. Precision on-line alignment data for each MDT is provided by an in-plane laser geometry monitoring system, allowing their relative position to be monitored to a precision of 10 μ m or better.

5.6.2 Cathode Strip Chambers (CSC)

The CSCs provide the same service as the MDTs – correct curvature measurement – but with higher radiation resistance, and higher price. They are built around the multiwire proportional counter principle, with anode wires strung through a gas environment (Ar-CO₂-CF₄), and cathodes laid in a network along the base surface of the chamber. It's designed to give fast response, low deadtime and to handle high rates without being damaged.

5.6.3 Thin Gap Chambers (TGC), Resistive Plate Chambers (RPCs) and Triggering

The TGCs and RPCs are important for spatial measurements, of course, but temporal measurements are the really important ones they produce. The MDTs and CSCs, because of their long drift times and distance from the IP, are difficult to correlate to a specific bunch crossing. A combination of the two allow both precision curvature and timing measurements.

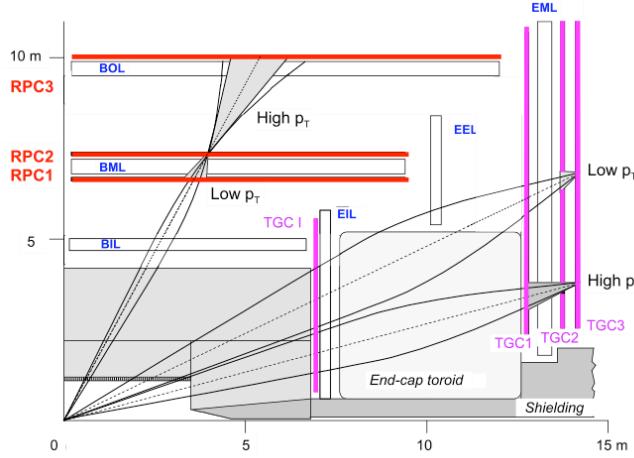


Figure 5.20: The muon trigger system, showing the relative positions of MDTs (such as the BOL) and CSCs and RPC/TGC trigger system. The figure also illustrates the differing track behavior for low- and high- p_T tracks, which allow the trigger to discriminate by p_T at high rates. Figure from [15].

5.6.4 Toroid Magnets

A major portion of the detector's volume (40%) is occupied by the ATLAS Barrel Toroid. Together with its endcaps, it produces a peak field of 4 T for the muon spectrometer, allowing muons of up to 2 GeV to be resolved [49]. A schematic of the conductor placement in the barrel is seen in Figure 5.22. To maintain a consistent and high field in the “fringe” regions at the ends of the barrel, corrector magnets are also needed. The 240 tonne endcap toroids serve this purpose, mounted just outside the CSCs and the calorimeter endcaps. With eight conductors each, nestled all the way into the gaps of the barrel toroid elements, they provide high and predictable field for the crack and high- η regions of the detectors. Each side's set of eight conductors is contained in a single cryostat, as seen in Figures 5.21 and 5.22



Figure 5.21: Endcap Toroid A, left, being moved from its assembly hall at CERN. Beside it on the right is the inner structure of Endcap C, being readied for placement in its own cryostat.

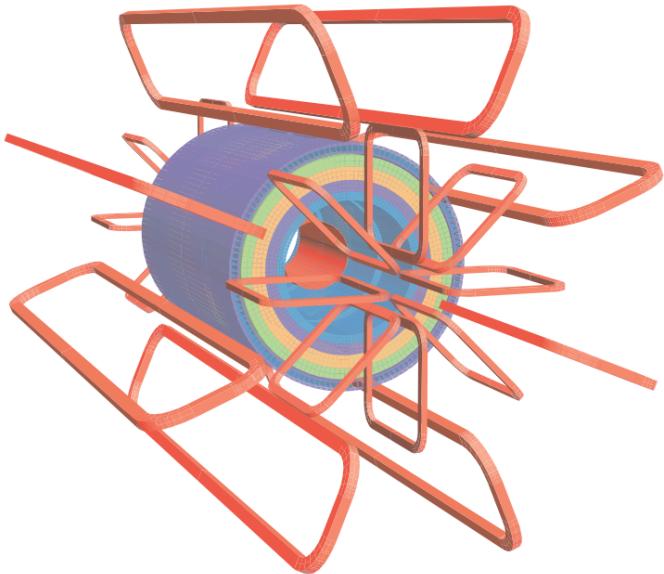


Figure 5.22: Geometry of the ATLAS barrel and endcap toroid conductive elements as they are positioned in the toroid magnet superstructure.

5.7 Shielding

While the shielding is both crucial to the success of the detector and to the specific rates of beam halo which will be seen in the detector, little space will be given herein to a discussion of its characteristics. Its purpose is to reduce the rate of neutrons, ionizing charged particles, and muons from the accelerator tunnel. It has been optimized for both performance and mass. Particles are met with cast iron or copper absorbers, with a second layer (in many cases) of boronated polyethylene clad in stainless steel to stop neutron flux and its associated high-energy gamma byproducts. Any joint in the shielding requires “chicanes” – short 90-degree changes in direction of any joining surface, to prevent any line-of-sight paths through which particles (especially neutrons) could escape. An overview illustration of the shielding is seen in Figure 5.23.

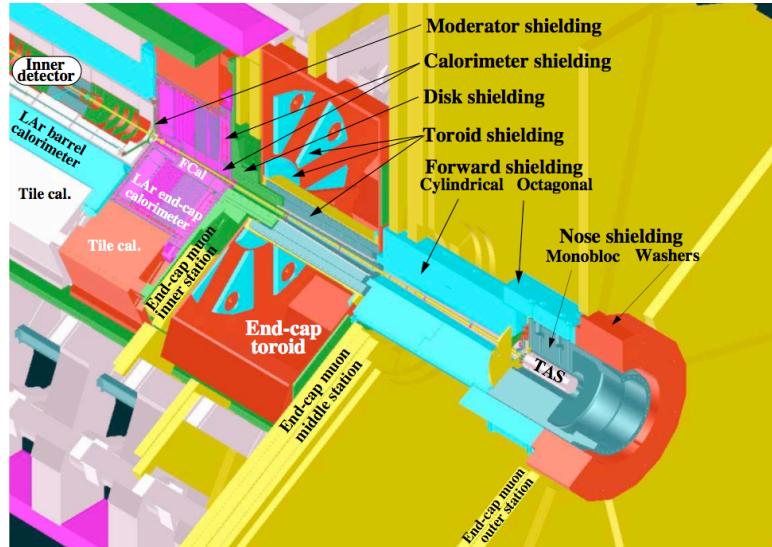


Figure 5.23: The ATLAS shielding layout, mirrored on both sides of the cavern.

Chapter 6

Data Handling

The data rate from ATLAS would be unmanageable in its raw form. At ~ 25 MB/event before zero suppression (~ 1.6 MB after), with a peak of 40 million events per second for a luminosity of $10^{34} cm^{-2}s^{-1}$, the detector produces a petabyte per second (or 64 TB/s with zero-suppression). To be written to tape in a reliable fashion, the data stream must be on the order of 100 MB/s, to allow for variations in rate and some reprocessing runs. The reduction from 64 TB/s to 100 MB/s requires that the data be analyzed and stepped down in stages, eliminating uninteresting events in the least costly way possible. This is done by implementing several layers of trigger, which do increasingly detailed looks at the flow of events and compare them to preset criteria to determine whether they will be included or let fall. The trigger system is tightly integrated with the data acquisition (DAQ) systems of each of the subdetectors. Either the output from the subdetector is an input to a trigger process, or it must be stored until a trigger decision is made. As such, the first component of the trigger is really the output buffers integrated into the readout of every detector component. The front-end DAQ of every detector subsystem contains a pipelined buffer, allowing the data read from that component to be stored for a period of $\sim 2.5 \mu s$, as the first stage of the trigger decides whether to accept the BC's data in the first place. The data are replicated off of some of the DAQ channels to make a trigger decision. From this point on, the DAQ systems serves as a blackboard on which all the trigger's decision are marked, and from which the trigger's final results are copied.

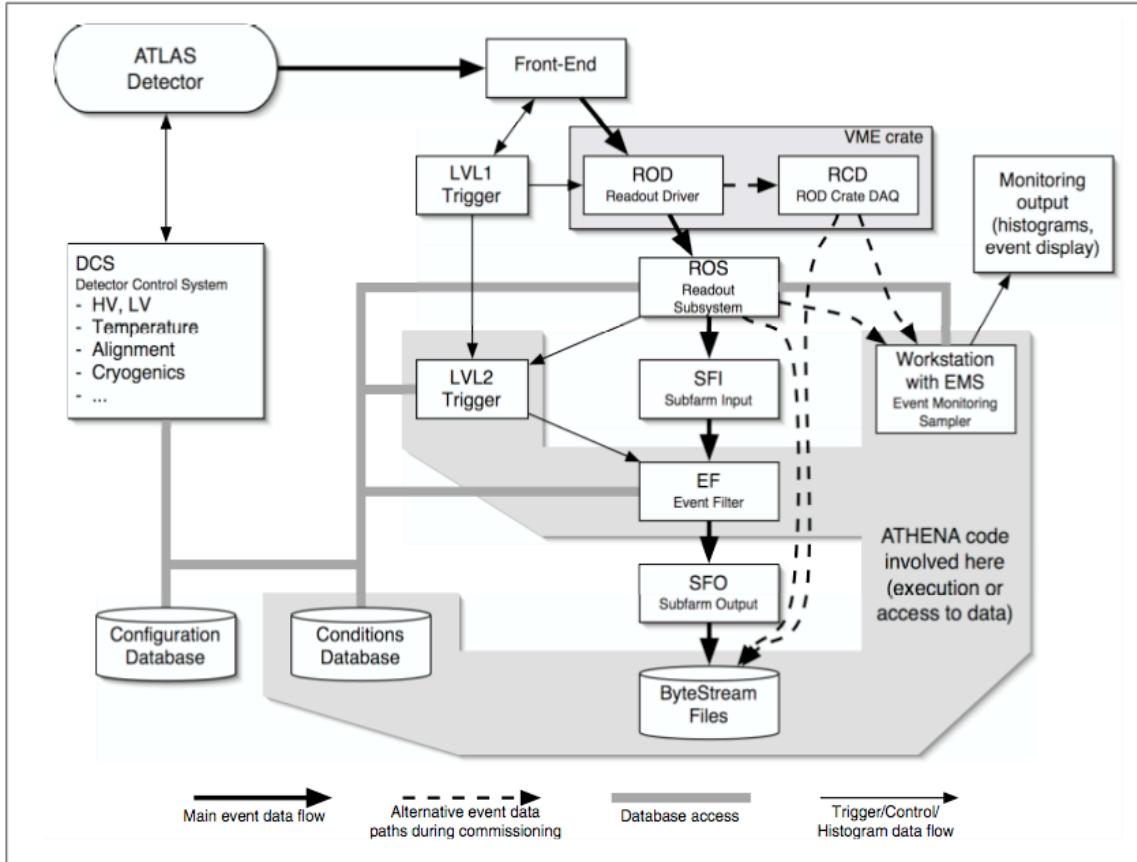


Figure 6.1: The Trigger/DAQ system. Figure from [15].

6.1 The ATLAS Trigger

This description of the ATLAS trigger system here is given somewhat short shrift, serving as only the lightest of summaries to draw the outlines of the system, due to the extreme complexity of the trigger hardware and software. Any attempt to write about the system beyond the superficial ramifies quickly, as the decisions made about all components and architectures inform large parts of the system. A thorough treatment of the trigger can be had in several authoritative references [50, 51, 52, 72].

6.1.1 Level-1 Trigger (LVL1)

The initial stage of the trigger system, LVL1, also comprehends the greatest data reduction. Since the data rate must reduce from 40 MHz to 75 kHz (a factor of 533x), and the raw data stream is at its greatest in this phase, the trigger has no time to do any flexible preprocessing of the data before a decision. Rather, sums of some calorimeter components (trigger towers) are done in hardware and sent to the LVL1 Calo subsystem. Data from the RPCs and TGCs in the muon spectrometer are sent to the LVL1 Muon subsystem. These are then combined in the Central Trigger Processor (CTP) with other inputs – from, for example, temporary cosmics triggers, or the Minimum Bias Trigger Scintillators (MBTSSs), or from the Beam Pickup Monitors (BPMs). There are also random triggers and prescales based on certain BC intervals.

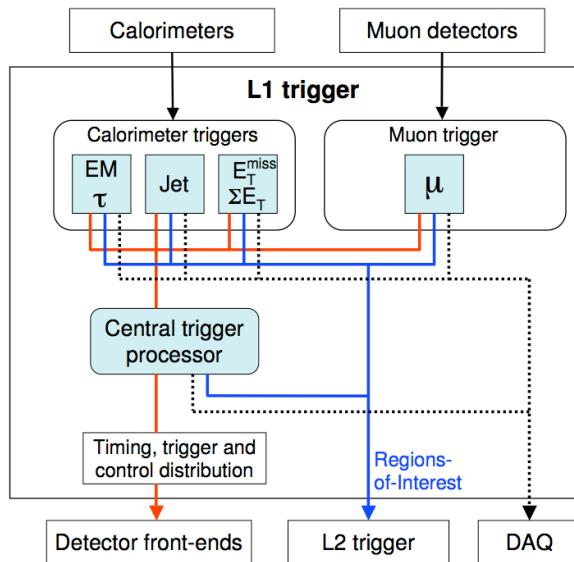


Figure 6.2: LVL1 trigger decision flow diagram. Figure from [15]

All these inputs are processed in the CTP, which issues (or not) a Level-1 accept (L1A) signal, dumping the data for that given BC into the ReadOut Drivers (RODs), such as those mentioned in Section 8.2.1. The LVL1 trigger can be configured on the fly, allowing up to

256 criteria to be added to a “trigger menu”, each with its own prescale setting and priority (either LOW or HIGH), and the results of these criteria determine the rate of L1A signals that are sent. The maximum accept rate for ATLAS nominal run is 75 kHz – after that point, throttling is necessary to prevent problems downstream. An upgrade to a 100 kHz L1A rate is foreseen for future ATLAS upgrades. Once the L1A is sent, data are streamed through the subsystem-specific RODs, which do checking and formatting, and then passed into the Read-Out Subsystem (ROS), populated by individual Read-Out Buffer machines (ROBs) that handle the data during the next decision phase.

6.1.2 Level-2 Trigger (LVL2)

The LVL1 decision is made based on specific criteria, and the regions of the detector that satisfied those criteria are called Regions of Interest (RoIs). The LVL2 trigger performs a limited reconstruction of the particles in this region of interest using commodity PCs linked together on a gigabit Ethernet backbone, and determines from this rough reconstruction whether the ROI was actually interesting for the reasons the LVL1 trigger kept it. This is called a “seeded” approach, in contrast to the “stepped” approach of the LVL1 – the processing requirements are substantially reduced, allowing the LVL2 to make its decision in \sim 10 ms. The network between the LVL2 and the ROS transfers just those data necessary to give the trigger the relevant objects to reconstruct the RoIs. If an event is rejected, the LVL2 Data Flow Manager (DFM) informs the ROS that it needs to dump the contents of that particular BC. Trigger menus are very flexible at this stage – the LVL2 and higher trigger systems can run some of the same algorithms as the offline reconstruction system (depending on time constraints), using an online version of Athena. This is an advantage for trigger flexibility and implementation, allowing faster optimization of the trigger acceptance.

Up until this point, the event has been fragmented across several machines within the ROS, due to the fact that reassembling the event from the various subsystems would be a waste of

precious time. However, once the event has been accepted by the LVL2, evaluation of its characteristics will require that the whole event be reconstructed and examined, so the event is defragmented by a sub-farm input node (SFI) and passed on to the event filter.

6.1.3 Event Filter (EF)

The event filter has all the data from a given BC, and is able to make a final decision on the event based on even the slowest of the subsystems' output, within the LVL2 RoIs – and to build those RoIs with a more complete detector description, including material data and updated calibration and alignment data. After 1-2 s of processing, a final decision is made on whether to accept or reject the event (reducing the data stream by a further factor of 10, to \sim 200 Hz), and the event and any results the EF has built are written to mass storage.

The EF is the first place that a whole picture of the inner detector is available. An example of a trigger subset that might require such completeness is an empty event filter [53], called in the aforementioned a “random-based track trigger”. Clearly, empty events are not desirable parts of the data stream. Events with a minimum of trigger-induced selection bias (as described in Chapter 9), will be sought in the initial phases of ATLAS run. One useful way to avoid bias is to trigger randomly at a certain rate – but there is a very large chance (during the initial stages of LHC commissioning, as discussed in Section 4.5.1) that any chosen event will be empty except for beam-gas or beam halo particles during initial run. To reduce this possibility, a trigger based on the number of SCT spacepoints [46] in a given event can be used to filter NSD events from empties, and substantially increase the effectiveness of a random-based minimum-bias trigger.

6.2 Computing in ATLAS

Once the data have passed the trigger and are stored, they must be reconstructed and analyzed (see Section 7.4). The data undergo several calibration passes, where interpretations of the detector output are optimized with more detail in the light of the conditions data from the run. This is one of several computation-intensive tasks that are required of the ATLAS computing system, requiring unprecedented resources to keep up with the data flow.

Large-scale computing paradigms have evolved from the mainframe to massively parallel systems and midrange machines, to the present working model of batch systems built of commodity processors in large numbers, meshed together in a fast network. ATLAS will be using this concept extensively (especially for the higher-level triggers, as mentioned in Sections 6.1.2 and 6.1.3), and is working to move beyond some of its limitations.

6.3 Batch Processing and High-Performance Computing

Classic batch processing has been around since limited computing resources needed to be shared among groups of people – or since the first computing systems. In a heavily-networked environment, physical access to the machines is sharply limited, and a cluster of computation nodes is usually accessible through a remote sign-on, whether to a gateway “head” node or directly to the compute node itself. Access requirements are crucial to the success of this approach, as is a “fair-share” system to allocate resources efficiently. ATLAS computing resources have historically comprised a patchwork of individual clusters. The primary CERN cluster, 1xplus, and its mass storage backing system, CASTOR [54], are among the largest employed in ATLAS, though universities and national laboratories such as Brookhaven National Laboratory (BNL) are ramping up disk and processor resources as the experiment prepares to start up. To use the resources in these disparate and balkanized clusters requires a substantial amount of individual effort and preparation. Movement of datasets between computing resources is time-consuming, especially from overseas locations. While bandwidth has grown substantially in recent years, dataset sizes have more than kept pace. Maintaining secure access between these resources is also crucial – and individual users must cope with password change and service change requirements among many separate bureaucracies.

Fair shares are mediated by batch software sitting in the head node of the cluster and doing job scheduling and finalization. Such systems as PBS, LSF and Condor are in current use as batch computation handlers, each with its own strengths and weaknesses.

6.4 ATLAS Grid Computing

The ATLAS computation requirements [55] dwarf the capacity of any practical local resource. To meet these requirements, the ensemble of computing clusters, past, present and future, need to be accessible to any ATLAS user anywhere in the world. This also has a balancing effect on the economic aspects of HEP research, making resources accessible to smaller groups and more constrained national funding agencies that would be ruinous to maintain on their own.

The ideal of Grid computing is for an individual user with minimal specific training to be able to use the vast resources of the ATLAS collaboration (and other Grids that may be shared from other collaborations and projects) as if it were a single machine, or at worst a local cluster running a batch system. In practice, this is a difficult goal.

The conceptual problems associated with this ideal are tied to the fundamental needs of collective computing, mentioned above – security and efficient resource usage. These resources are not just the computers, storage and networks in play, but also the end users and the collaboration’s ability, in sum, to achieve its research goals.

If system security is too lax, compromised systems are inevitable, and resources are lost both in computation and liability. If it is too tight, each individual user spends, in principle, a significant part of the productive hours of the day jumping through the hoops. If the systems’ common protocols are efficient in publishing the instantaneous state and capabilities of a resource, it can be utilized correctly and not waste throughput. If failures, availability, or latent problems are not detectable and infrastructure is not capable of good uptime and performance with the human resources at our disposal, there is a financial and productivity penalty accrued by the collaboration.

6.4.1 Development

ATLAS has requirements that do not fit well into commercial grid solutions, both for price and architecture reasons, so development of Grid tools has largely been done by members of the

collaboration as part of the experiment. The initial grid development teams (NorduGrid, from the Scandinavian countries, Grid3/OSG from the US, and LCG, centered at CERN) created competing products, and in the process refined techniques in parallel and investigated problems and assumptions from various specific implementation strategies. This cross-pollination approach has gone through several iterations, and some previous protocols and concepts for Grid implementation have been abandoned in the face of the realities of international, intercontinental and interhemispheric communications. One crucial concept is that of efficient and reliable data tracking and transfer. The assumption that data could be moved to wherever one might find a free processor ran into the aforementioned data size problems – the process of moving the data, even locally over an intranet, can be more resource intensive than the computation itself. In addition to lag times from communications, the asynchronous nature of the dispersed network of clusters requires that any data transactions be tracked and mediated, verified, and if necessary redone, in a completely reliable way. These data are the fundamental currency of the experiment, and data loss or inaccessibility has a direct and lasting impact on the return on investment that ATLAS can give.

As a result, one Distributed Data Management (DDM) system is in development that catalogs the balkanized storage systems, wrapping their individual infrastructure choices out of sight to the average user and, through client software and backend tracking databases distributed and synced throughout the world, attempts to allow unified data access. This choice was made to mitigate the Babel of different file transfer protocols and their unreliability, and create a scalable worldwide data tracking system for the petabytes of physics being produced at the LHC.

6.4.2 Production and Analysis on the Grid

The commencement of data taking and analysis will bring with it a number of challenges, and will in all likelihood break many existent models of ATLAS data processing strategies – so I will not address the present sets of Grid Monte Carlo production and physics analysis, other than to state that the intent (and realistically, the requirement) of the collaboration is to use Grid resources, centrally managed but democratically accessible, to produce datasets and run several

stages of analysis – sending the jobs to the dataset rather than having to transport the dataset to a local node for the user to manipulate. The distillation of these steps would then be returned to the requester, where final operations and publication plots will be made. The debate about where in the process this separation should occur is ongoing.

Chapter 7

Effective Simulation of ATLAS Physics

The ATLAS experiment remains unfinished. Even so, the efforts to connect, characterize and calibrate its subsystems is ongoing. From the very beginning of the ATLAS design phase, it has been necessary to attempt to predict the detector’s response to various particles’ passage. As data on the real detector’s performance arrives, it allows the simulations to be tuned, more closely mirroring the final reality.

7.1 Event Generation

The initial stages of event simulation in ATLAS are done in *generators*, or software packages that codify a particular physics process (or set of physics processes). They tend to simulate the collision of a pair of protons in the center of the ATLAS detector, process the physics of the event, and output the particles with their parent particles, positions, directions and momenta stored in 4-vectors. In general, the reach of a final-state particle from a generator is very small – the final vertex will sit just outside the interaction point. For some generators, such as the ones used in cosmic ray simulation, these assumptions are not valid – the particles may originate at a cavern wall, for example.

To interface these final state particles with the ATLAS simulation framework (called Athena [55]), the 4-vectors must first be put in a few specific format, called `GenParticle` and `GenVertex` which comprise `GenEvents`, and stacks them together into “persistified” objects called `POOL` containers, which are stored in a `pool.root` file. These events can be made transient again in the Athena framework – this transient/persistent model is the medium of object

exchange for most ATHENA jobs, and allows the restoration of whole persistified C++ objects and classes, rather than just data members.

7.2 Simulation of the ATLAS Detector

Once the HepMC are in place, we move to the step of propagating them through the detector. The Geant4 [63] package is responsible for this part of the event simulation across the ATLAS experiment

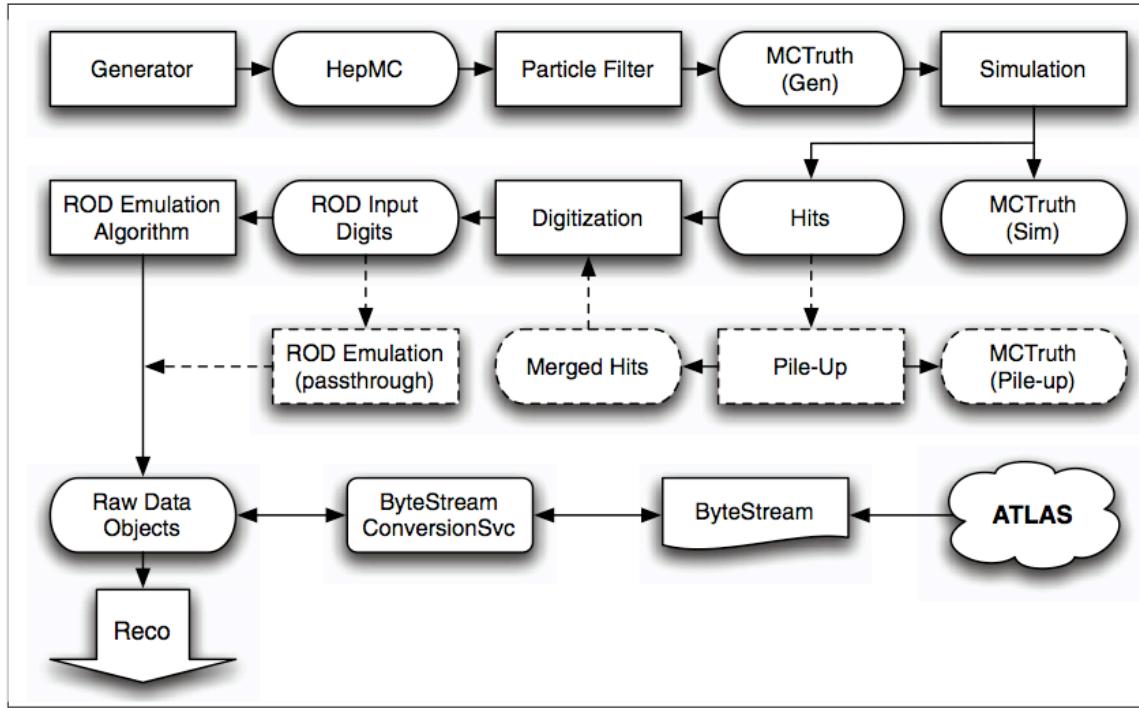


Figure 7.1: The ATLAS simulation flow. Figure from [15]

7.2.1 Geant4

Geant4 (G4), as integrated into Athena, serves as the present simulation engine for the ATLAS detector. Initial studies for ATLAS were done with either the GEANT3/DICE or fast parameterized simulators such as ATLFAST [64], which are legacy FORTRAN code. G4 takes

the simple structures of `GenEvent`, `GenParticle` and `GenVertex` out of the POOL persistified format into the `StoreGate` transient data store, and propagates them through a model of the detector’s volumes in discrete steps. Each volume, boundary and transition in the detector that is traversed records its time offset and energy deposition, and modifies the particle that passed through. Simulation takes as input a description of the detector geometry, either with or without misalignments and deformations, and the magnetic field generated by the toroid, endcaps and solenoid – and their interaction with the iron in the detector and support structure. Dead material is also taken into account – details such as wiring, cooling pipes, tape, tags, support structures and shielding are incorporated into the geometric model of ATLAS. `Geant4` tends to look at events as always originating at time $t = 0$ with respect to the BC. For the purposes of beam-gas interaction simulation, a modification needed to be made to the interface between `G4` and `Athena` to allow `G4` to calculate the time offset of a given vertex from its z position in the detector, and propagate that time offset to all daughter particles of the vertex. A. Dell’Acqua provided that functionality, available from `Athena` release 13.0.10, in the `Simulation/G4Extensions/VertexTimeOffset` package. For releases before 13.X.0, the package must be imported and built by hand.

7.3 Hits and Digits

The final output of the detector (called the “hits” file) contains timing and energy deposition for the detector, but needs to be further processed to actually look like detector output. This is done in the digitization step – translating a timed hit from a notation about a detector volume into the sort of digital response one would get from that detector component in reality. Two kinds of output – digits and bytestream – correspond to different use cases. For producers of MC simulation, the events are simply distinguishable and are destined to be stored in a number of small files, and the maintenance of the Truth particle information is important. The ATLAS bytestream, in contrast, is the realtime output of the detector elements as requested by the trigger (discussed in Section 6.1), and comes in a somewhat different format, with trigger and run-specific information. The digitization, once done, represents the real output of the detector

as a Raw Data Object (RDO), organized into containers corresponding to the detector elements that were involved, including the noise they produce. The standard digitization, however, does not account for the latent effects of previous collisions' particles that remain in the detector. Remembering that a particle at c travels only 7.5 m in the time between bunch crossings, and that the integration times for several of the drift-style detector subsystems are substantially longer than 25 ns, and it becomes clear that a method is needed to bring in effects from before the present bunch crossing for a reasonable approximation of what we will see in the real detector output. Furthermore, there will be background particles (beam halo, beam-gas collisions and cosmics) that enter the detector with the imminent set of bunch crossings, and so also need to be included. Furthermore, at higher luminosities the ATLAS experiment will see more than one pp collisions per BC. Several minimum-bias events overlaid with the vertex, and others further out, to correctly describe run conditions. Specific pileup variables allow these variables to be tuned to the specific luminosity and run condition needed.

7.3.1 Pileup

To take care of all these effects, the pileup options in the digitization software must be enabled. The pileup takes background hits files as inputs, and offsets each volume's timestamp to correspond to a previous bunch crossing – or even to inter-BC time slices, down to a resolution of 1 ns. This history of the detector is built up, then the contribution of each hits signature to a given detector volume at any given time is summed up, and processed as normal digitization. The process requires a large amount of memory for correct running. From the Athena 13.0.10 release, P. Calafiura included, at my request, new hooks for running beam-gas interactions and beam halo in the pileup. Another needed feature was the ability to tune the rates of any given pileup contribution to account for both rare and common backgrounds in the same event. With these modifications (and the aforementioned timing modifications to simulation), the correct behavior of the BG events and the halo can be simulated and digitized correctly. To scale either the BG or BH samples to higher or lower rates, a simple change in the center of the Poisson

distribution (one of the pileup variables) is needed, and this is done in the jobOptions files that configure the run.

7.4 Online and offline reconstruction of ATLAS Events

Whether arising from MC or the detector, the digits need to be reconstructed to be useful in a physics study. This implies that the hits in tracking and muon systems need to be matched into tracks, the cell energies in the calorimeters need to be summed and grouped into clusters that correspond to real particle hits. The full reconstruction involves hundreds of packages, each taking a look at the RDOs or bytestream in a different way. The reconstruction results therefore necessarily contain many overlaps, as each particle can be picked up in several ways. The full reconstruction is stored in Event Summary Data (ESD) files, which are further processed and thinned to provide Analysis Object Data (AOD) files appropriate for end-user analysis, and which can be further transformed into user-specific ntuples for ROOT analysis.

7.5 Athena

7.6 The Athena Framework

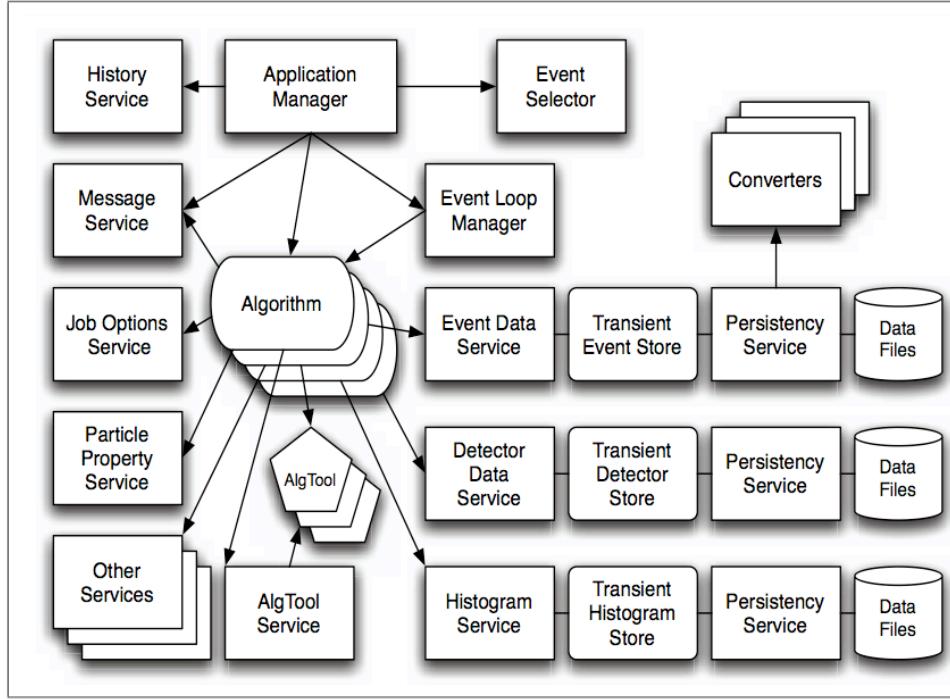


Figure 7.2: The Athena component model. It is clear from the diagram that the Athena framework provides services of interest to ATLAS users, and allows its customers to implement the actual physics functionality – plugging in various algorithms (such as generators, simulation packages and reconstruction tools) and storing their results in a consistent set of formats. Figure from [15]

To allow that simulation to be compared quickly and accurately to the results from the detector (to either support or falsify the hypotheses we are testing), a software framework was proposed near the beginning of ATLAS. The fundamental purpose of this framework was to provide services and algorithms in an orderly and central manner, allowing both the real and simulated detector data streams to be accepted, manipulated, and stored in the same manner and with the same tools – reducing the number of reconciliation steps and discrepancies one would

have to account for in a real analysis. Athena is, at this juncture, an enormous undertaking, and addressing its details is not within the scope of this work. Those aspects of it that *are* relevant are the following:

- The basic data units of Athena are based on the concept of events. Because of the nature of ATLAS and the LHC, there is no need to (in general) look for correlations between separate bunch crossings. All processes in Athena and its underlying Gaudi [56] framework operate on one event at a time.
- Externally developed generator packages are wrapped with appropriate interfaces and added to the Athena framework to produce these events – Athena is extensible by design.
- Athena’s event generation framework is therefore limited by the generators that have been imported – and their stability and quality within Athena.
- All data are handled by the Athena StoreGate service, which allows the event objects to be used in memory, then stored as retrievable objects on disk – states called “transient” and “persistent” data stores, as mentioned earlier.

7.7 Generating Minimum-bias Backgrounds

7.7.1 pp collisions

The starting point for the experiment, as with any other, is to find out where we are – to get our bearings. There are several ways of doing so. “Standard candles” of various sorts are sought to quantitatively compare the response of a given subsystem with previous experiments – resonances like the J/ψ , or Z bosons decaying to two electrons or two muons, which have known masses and can be well-reconstructed, allow us to calibrate the detector response. Rates of these and other processes can tell us about the differences between this energy regime that we are entering and those that have been probed previously. pp collisions are by no means simple or clean. As mentioned in Chapter 9, there are many partons involved, and their participation and energy are only describable statistically for the moment. The “interesting” events that lie at

the focus of ATLAS are always present on a backdrop of further collisions in the same bunch crossing and possibly multiple parton interactions within the pp collision itself. Because we are probing the boundaries of our understanding of the structure of a proton as part of the experiment itself, it is nontrivial to describe the cross-sections for various kinds of parton collisions (including soft and double-Pomeron collisions). Two primary event generators make attempts to correctly describe pp collisions with different models. PYTHIA, a MC event generator developed by the Lund group, and is widely used within ATLAS. For more details on this very complex and configurable package, please see the citation at [57]. PYTHIA has been evolving rapidly over the last several versions. The baseline tuning of PYTHIA as a minimum-bias generator was investigated in [90] in some detail, optimizing it and comparing it to the PHOJET package [86, 102] from 1997. The primary optimizations were in tuning the p_T cutoff for the transition between hard and soft QCD physics, and in varying the parton distributions between the “core” of each proton and its outliers. Since the paper, PYTHIA has incorporated the concept of color resummation, which has proven to be as important an effect on the minimum bias (and underlying event) physics simulation as either of the other two. For the interim, it has been determined that minimum bias MC production will follow the same parameters as the underlying event tune, to simplify production and await an infusion of real data at 14 TeV to allow realistic tuning. Because the generator itself is largely in flux as of the date of this writing, it is not therefore fruitful to discuss the specific settings used in PYTHIA during the runs that produced the baseline MC for this study. The datasets are standard Computing Service Challenge (CSC) datasets for ATLAS, and are listed in Table A.1 – which means that their parameters are generally available. It suffices to note that the datasets represent the state of the art at this time – and that the effects we seek to understand in the present document are differences in behavior in reference to a standard. These investigations should be revisited in the light of initial data with new MC generator tunes to better understand the interplay of the beam gas and beam halo with the minimum bias triggers and events.

The beam halo and beam-gas events are not generated in the same way as the minimum-bias events – their sources are fundamentally different. More on the nature and characteristics of the beam-gas and beam halo events can be found in Section 9.3.

7.7.2 Beam-Gas Events

Beam-gas events (BG) are simply $p - A$ collisions – the impact of a relativistic proton on an essentially stationary atomic nucleus. For the purposes of this study, it was determined that a generator designed for nuclear physics might be more appropriate. HIJING (Heavy Ion Jet INteraction Generator) [58] was chosen, for both convenience and practicality. The ATLAS Heavy Ions community had already integrated HIJING into Athena, placing the interaction point at $x, y, z = 0, 0, 0$ and giving the generator a “sense of direction” – enough geometry to make incoming Pb nuclear pairs arrive with the correct momentum vector. Missing from the efforts was a means of placing a stationary target nucleus and bombarding it with a probe particle from a given direction. The appropriate modifications to the generator interface in Athena were made, including some that allow the interaction vertex to be shifted, *post facto*, to another arbitrary location within the bounds of the beampipe. Later, those efforts were also abstracted into the Generators/TruthExamples/EvtShift package for general use with HepMC generator output. With directionality and location implemented, samples for various expected target nuclei were produced – 500,000 H interactions (250,000 from each side), 100,000 each of C and O, and 10,000 each of Be, Cu, Cr, Fe, N, Ti, V, W and Zr. These smaller element sets represent very unlikely contamination possibilities, though they are all present in the materials of the beampipe, Non-Evaporative Getter (NEG) beampipe coating, the collimators in the accelerator and the TAS shielding. They were simulated as an exercise in looking at the signatures of higher-Z contaminants, so that the effects of such events would be recognizable – should something so unlikely occur.

7.7.3 Beam Halo Events

While W and Cu proton-nucleus collisions are likely to be vanishingly rare in the ATLAS experiment beampipe, they will be the dominant part of the beam halo. Beam halo arises from a number of possible interactions in the accelerator, but the ones that matter to the experiment are the parts of the halo that continue past (or originate at) the final collimators before the IP, and any BG interactions that occur in the beampipe before the ATLAS cavern. The original estimates of ATLAS beam halo were made in 2001-2003 in a series of LHC Project Notes [59, 60], which included various accelerator effects, primarily the interaction of the beam with residual gases in the beampipe. This simulation was done with the MARS [85, 91] package, and was destined primarily as a dataset for shielding and construction purposes. The MARS simulation contains no time structure – it composes a list of notional particles, with their energies and trajectories, and assigns each a probability weight. There is also no correlation between the particles themselves – it is impossible to assemble a realistic shower of simultaneous particles from the dataset. Further complicating the issue, the beam optics version used in the study (V.6.4) did not contain the collimators for the final machine, nor were proton loss maps available for reference to implement collimator loss profiles correctly. The study, therefore, included the weaker part of the machine background – its beam-gas interactions – while excluding the very important point source of beam proton losses on the tertiary collimators. To rectify the weaknesses of the simulation, I approached one of the authors of the previous study, Vadim Talanov, about the possibility of a new round of beam simulation. The resulting dataset was produced in a substantially different manner. The FLUKA [61, 62] generator was used, and with it came the ability to correlate final-state particles to their parents – allowing showers to be reassembled and simulated as an ensemble in Athena. The source of the dataset is the impact of tertiary beam halo on the tertiary collimators (TCTs) – specifically the horizontal TCT, which bears the brunt of the tertiary halo (the final TCT before the ATLAS cavern is vertical). Because of its distance – 148.26 m from the IP – the proton impacts on the collimator jaws is effectively a point source, and is simulated as such. This simplifies the process substantially – simulating diffuse sources (such as all the gas distributed in the curved and nonuniform beampipe of the accelerator) is

substantially more complicated and was left aside for the present study. Other sources, such as collisions at other IPs, are also not included. Most beam contamination that manages to survive a substantial distance around the ring will be collinear and contemporaneous with the beam packet itself, and unlikely to impinge on the detector. The FLUKA simulation was run with the optics V.6.500, in the collision optics mode – the collimators in their normal position for a physics run. Particles are placed in positions on the strike plane at 23 m from the interaction point – just after the cavern wall, behind the shielding. The particles were translated into the HepMC format for inclusion in Athena using a modified version of `CosmicsGenerator`, and momentum translations were done with `Generators/TruthExamples/EvtShift` to allow some of the sample to come from the C side of the cavern as well. There was no need for any x or y mirroring of the particle positions, because the accelerator tunnel itself is a mirror image from one side of the cavern to the other.

Chapter 8

Hardware and computing contributions

For a major portion of my time working at CERN, I participated in the development of hardware and computing necessary for ATLAS operation. This chapter will take a small break (between the description of the detector and the central thesis topic – the effects of accelerator backgrounds on the minimum-bias trigger) to describe some of those contributions.

8.1 High-performance Computing in ATLAS

My initial activity, upon arriving at CERN in June 2003, was an attempt (with Athena v.6 and thereabouts) to run MC and simulation for the Wisconsin group. Up until then, the group's primary mechanism for such simulation was running standalone Pythia and Atlfast. Full simulation had been far too computationally expensive to use regularly. I started running the full Athena chain, from generation through reconstruction, as the standard method of creating MC datasets.

Full simulation of things like $\text{Higgs} \rightarrow 4\ell$ were becoming important to the group's research direction, so I generated a set by hand and submitted it to the official ATLAS production group for processing. Since the return speed was excellent, the group requested more MC samples through that channel. This arrangement, though very convenient for the end users, became increasingly difficult for the production team. I had by then written a framework that automated the MC generation production and bookkeeping, and was running larger and larger sets with the help of Armin Nairz and Luc Goosens, the production team liaisons. Eventually, I became a *de facto* member of the production team, if not formally.

Further work was done in progress of time on both the Grid and batch systems, providing MC to the group and providing end-user evaluation and feedback to various infrastructure groups and production systems in ATLAS through the summer of 2006.

8.2 The SCT/Pixel Read-Out Drivers

After the Grid efforts, I was retasked to responsibility for new development on the control code for the Digital Signal Processors (DSPs) on the ATLAS Silicon RODs. I will first discuss some of the challenges of the ID readout system, then address my role in developing the software.

Readout Reading the pixel detector requires a chain of hardware and software to smooth the flow of data and minimize downtime due to overloads and error conditions. The pixel modules' Module Control Chips (MCC) communicate across a duplex optical line at up to 160 Mb/s, with a single 80 Mb/s incoming line for control. The initial fibers from the IP are rad-hard, but become conventional fibers at an outer patch panel. The signals are received in the Back-of-Crate Card (BCC) connected via a VME bus to the Read-Out Drivers (RODs), which handle the interaction with the LVL1 trigger, the LHC timing system, and the Read-out Buffers (ROBs). The load is significant – at 1% occupancy, the trigger maximum is 7.5 kHz to avoid deadtime [39].

The SCT is read out in a fashion similar to (though much simpler than) the pixel system. Onboard chips (known as the ABCD chips [65]), fashioned as ASICs in a rad-hard process, read and compress the hit pattern on 128 strips per chip, and contains discriminators, a binary pipeline, a derandomizing buffer, data compression and readout control logic, and calibration circuitry. The buffers allow the data to remain at the front end until a LVL1 decision is made, the data compression reduces the transmission load and therefore increases the maximum trigger rate, and the calibration circuitry allows charge injection for module calibration. The six chips per module (twelve per module pair) are connected to each other in a fault-tolerant pattern, allowing faulty chips to be bypassed – even (in most cases) allowing the chips from one side

of the module to communicate through the other side if the principal chip is malfunctioning. Each side of the module has its own optical link, but both sides can be routed through one link at need. The data are read off at 80 MHz (80 Mb/s), and control data come in on one fiber per module pair at the same rate.

8.2.1 Silicon Read-Out Drivers (RODs)

The pixel and SCT modules share many characteristics. In an effort to reduce cost and simplify the hardware bill of materials, the RODs for both detectors are based on the same hardware, designed to be flexible in implementation. Since I have spent some time with the ATLAS Silicon RODs, I will take more than a passing interest in this section. There are several reasons to put as much of the readout chain off-detector as possible. Principal concerns are radiation damage, heat load, expense and material budget. While small, light and radiation-hard components are not out of the question, their expense, development time and accessibility in case of failure are serious considerations. The concern over heat, however, is the most serious – computation at the front end is costly, and successful designs will minimize it. There are unavoidable complications to the idea of offloading the data preparation stage from the detector front-end. Principal among them is the demanding LHC environment. At nominal run, the light path (reaching up to 100 meters from the IP) to the readout system means that just the signal transmission time takes up 13 bunch crossings. Timing is therefore all-important. The data stream undergoes some compression on the SCT ABCD chips and the pixel MCC, but the construction of events, time offset handling, data validation, sampling, run control and calibration runs are all handled in the RODs. The data stream itself is taken from an optical to electronic datapath in a module called the Back-of-Crate Card (BCC), one of which is mounted opposite each ROD in the VME crates they occupy. The BOC handles the details of the optical links (including the data output to the ReadOut Buffer (ROB), the next step in the data chain, and the input of the LHC system clock, fanned out to the BOCs from the timing board mounted in the middle of the crate. A ROD-Crate Controller (RCC) single-board computer sits in the first slot

of the crate and handles communications to the rest of the DAQ/DCS control systems, whose gatekeeper is the Read-Out Sub-system (ROS) cluster.

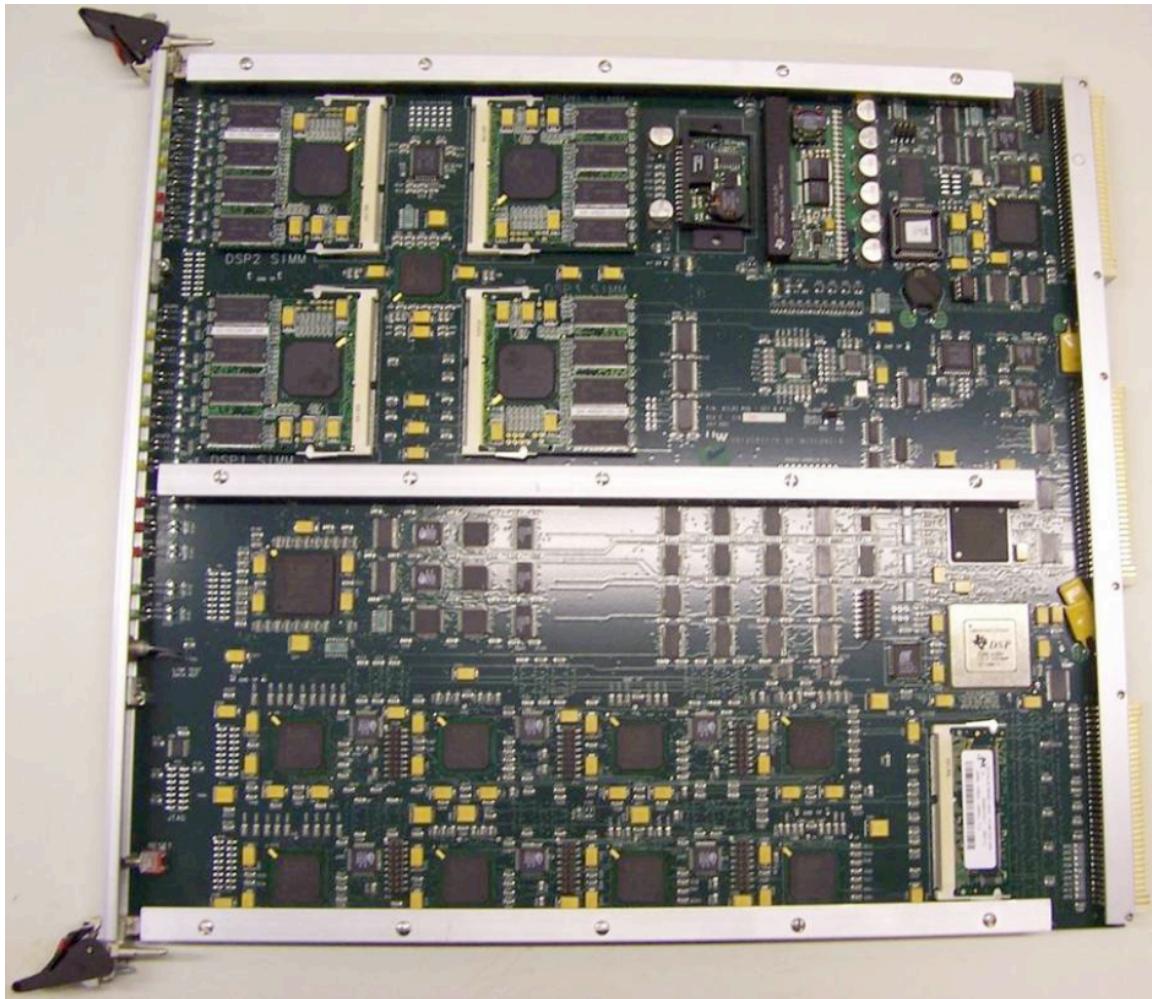


Figure 8.1: An ATLAS Pixel/SCT ROD. Visible in the top half of the board, on the left, are the four SDSP daughtercards and their associated RAM, and the router FPGA centered among them. Along the bottom, the eight formatter FPGAs are visible, above them is the EFB (event fragment builder), and to their right, the MDSP (silver) with its RAM. Above the MDSP, we see the controller FPGA. Figure from [66].

On the other side of the VME backplane, the ROD takes the incoming data stream and compresses it, reads potential error codes, looks for bit-flip errors in transmission and Bunch-Crossing ID (BCID) or Level-1 Trigger ID (L1ID) offsets between the LHC and the front-end modules, formats the data for transmission, and sends it upstream to the Data AQuisition system (DAQ), all in a chain of Field-Programmable Gate Arrays (FPGAs). The programming is flexible – the chips can be reconfigured *in situ* as needed – but their behavior is much less like the CPU of a PC, and much more like a large array of gate logic. This gives them the virtue of full predictability – their behavior is very deterministic susceptible to simulation in all available states – and their programmability allow the same hardware configuration to be programmed for pixel or SCT use. While more flexible processors are also “deterministic”, their number of available states is far more like a Turing machine, and the order of the simulation problem is much higher. Since the data stream cannot afford any avoidable hiccups, the present design was chosen because there are no “soft spots” in its data stream. Eight Xilinx FPGAs handle the formatting, one handles the event fragment building, one the programming of the others, one the control of the FPGAs in general, and one serves as a router for the data stream. The router is also capable of various sampling modes, sipping a copy of some part of the data torrent to the other part of the ROD – the DSP (Digital Signal Processor) farm. The DSPs serve as control and configuration devices for the ROD, and allow advanced operation of the ROD and FE electronics from higher-level systems. Several pieces of necessary functionality include:

- Module configuration
- Error tracking
- Data histogramming and validation
- Calibration and threshold setting
- Module verification

The flexibility of the DSP permits several codebases to be used (supporting for example, the pixel or SCT modules with their substantially different operating principles), and allows functionality to be added as the experiment comes online and the DAQ system evolves. Some of these items are discussed in more detail below. DSPs are small and flexible general-purpose CPUs. Depending on the implementation, they have a varying number of cores, varying instruction sets, and may or may not operate in floating-point mode. They are also programmable under the more flexible auspices of C or C++, allowing more robust and higher-level tools to be used, but can be programmed in assembler for greater speed and control over the processor's internals. The DSPs communicate with the rest of the electronics (including each other) over a Host-Port Interface (HPI) built into the DSPs themselves. The Master DSP (a Texas Instruments C6201) takes over the operation of the ROD once it's been booted, handling state changes and instructions from the Rod Crate Controller (RCC), a single-board computer located in one of the half-height VME slots to the left of the crate. Its primary task is trigger distribution to the front-end electronics, but it also does housekeeping tasks, allowing, for example, access to various registers and memories on the ROD and allowing EEPROM reflashes. It also serves as the primary (in fact, usually the sole) access to its four subordinate Slave DSPs, implemented on Texas Instruments C6713 (floating-point) units. Each of these DSPs has its own bank of RAM on a 256 MB SO-DIMM, and each DSP and its RAM inhabit a daughter card mounted on the motherboard, allowing easy replacement of defective components or (in principle) upgrades. The master and slaves work from data that have been DMA-transferred from the data stream in some proportion to the full flow – there is not enough transfer bandwidth or memory on the slaves to handle the complete stream. The slaves can be requested to do histogramming and data analysis tasks, passing results up to higher-level monitoring systems that can make decisions regarding dead modules and chips, malfunctioning strips, or other kinds of error conditions from the front end. They also process the data that come from the various kinds of calibration scans that are specific to the pixel and SCT modules – for example, making a Time over Threshold (TOT) scan of a pixel module, or instructing the ABCD chips on an SCT module to inject varying amounts of charge into the module's strips, then reading them out to measure

charge response. Coding for the DSPs has presented its own challenges. Especially in the pixel system, the DSPs are heavily tasked. It was determined early on in the development process that the supplied BIOS (the underlying services upon which compiled code usually depend, such as memory management, arithmetic operations, I/O systems and others) was too resource-intensive. Subsequent efforts to replicate those services have resulted in two codebases for the DSP project, one for the SCT and one for the Pixel. While their underlying mechanics are dissimilar, their basic interfaces to the DAQ system follow the same general lines. The host computer (in this case, the RCC) runs an instance of the DAQ software. Among the interfaces to the DAQ is an implementation of certain data structures called primitives, and it is through these primitives that the ROD DSPs are controlled. The DAQ software's copy of the primitives must be identical to the copy running on the DSPs, because there is no predefined protocol for data passing between these systems other than the primitives themselves. The technique is, therefore, to use the HPI to write both a copy of a necessary primitive (or primitive list, containing a sequence of primitives to be executed) to a known clock of memory, then to flip a bit in the communications block of the MDSP's memory. The communications memory map (depending on the code version) is either hard-coded as memory offsets from a base point, or reflected in a specific struct defined between the two communicators. This bit flip informs the MDSP that new instructions are available, and the MDSP parses the blocks of data that have been deposited in its memory by casting them as single or listed primitives, then calling its own copies of those primitives' code. Version control between any given "master" and "slave" is therefore crucial to correct operation, and checks are in place to make sure the correlation is good. Other, more complex (and longer-running) operations are created as "tasks", and are distinct from primitives in that they have no corresponding member in the "master's" code. Rather, they are configured and initiated via primitives, and then run independently until closed down or until they finish. Things that are well-adapted to task implementation are things like error monitoring and reporting, detector monitoring and histogramming. In the silicon ROD codebase, "master" and "slave" are actually general concepts that transcend the M- and SDSP designation. The RCC communicates with the MDSP in the same way that the MDSP communicates with the SDSP –

all primitives that can be reused, are. There are substantial divergences in functionality between the M- and SDSPs, but the communications protocols are the same. In the same sense, the RCC does not communicate directly with the SDSPs – the master is passed a primitive whose function is to pass a primitive list to a given slave. Likewise, the communication back to the DAQ system is strongly tiered as well. Each SDSP maintains a pair of text buffers, for error or normal output, which are polled regularly by the MDSP. The MDSP, in its turn, has its own text buffers, and also contains text buffers for each of the SDSPs – it acts as a communications clearing-house. The RCC then polls the MDSP buffers periodically, and collects their contents to pass up the chain. It is notable that in this process, “slaves” never initiate contact – it is always the “master” in the chain that both initiates action and collects the results. Primitives can also pass back status messages and text through their own interconnects, allowing runtime control. This model allows a degree of deadlock protection. Internal consistency is much easier to maintain if only one system is modifying any given piece of memory at a time. An excellent example of this design paradigm is the text buffers. Without the text buffers, coding and debugging are problematic. The RCC is capable of polling the main memory of the MDSP itself (and likewise the master can poll the slaves), and debugging requires sorting through the memory contents of the running DSP to determine its state and any faults that may be occurring. Since this is not an efficient or desirable communications technique, the text buffers are necessarily simple and robust. Supposing that the DSP has generated some quantity of return text, it appends the text to a buffer, and moves the tail pointer, whose location is known from the aforementioned memory map. In the meantime, the host (in an asynchronous fashion) has arrived at a moment where it is ready to look at the text buffers. It checks the difference between the head and tail of the text buffer (which loops) to see if data are ready for transfer. If yes, it takes the data between head and tail pointers, and sets the head pointer to the last location it looked. Neither side is modifying something in a way that could deadlock, and there is no need for problematic realtime synchronicity. Of course, if the host fails to read the text quickly enough from the buffer, data are lost either by overwriting or failing to write because the buffer is full. However, compared

to a lockup in any part of the control system (for something as crucial as the RODs), this is a small loss, as long as the data stream is allowed to continue uninterrupted.

Chapter 9

Minimum Bias

9.1 Introduction

As discussed in Chapter 1, there are a number of physics motivations in ATLAS. One thing they have in common is that all of them will arise from the scattering of quarks and gluons. Quantum Chromodynamics (QCD), the mechanism that we use to describe these interactions, is extremely useful in describing the higher-transverse momentum ($p_T > 2$ GeV) varieties of these collisions. Transverse momentum is used as a quantity of interest because while conservation of energy is valid in a pp collision, there is no guarantee that the specific partons interacting in any given event will have momenta that cancel in z – in fact, the likelihood of such an event is small. In the r - ϕ plane, however, the total p_T of the partons themselves should be well below the momenta of the products of the event. We therefore take p_T as the quantity of interest in produced particles. Signal events in ATLAS and its sister detectors will generally populate the low- η regions of the detector with high p_T particles. These areas have been instrumented with the best of the aforementioned tracking and calorimetry components of the detector. The underlying event and QCD backgrounds, however, will place particles in all parts of the solid angle, at energies that range much lower than the effective resolvable track p_T (around 500 MeV is the limit). This implies that the majority of the QCD background is not even resolvable within ATLAS. That which can be resolved, however, needs to be observed. Some techniques for low- p_T tracking (for tracks below 500 MeV) are under investigation – see [46] for more details. Low- p_T tracking will not be included in this treatment, due to its preliminary nature and substantial resource requirements.

The focus of this thesis is the backgrounds that will arise in the machine, and which will affect the whole chain, from readouts to results. The minimum bias events are an excellent place to study these backgrounds, because minimum bias studies start at the very beginning of detector run, when event rates are small and there is a substantial likelihood of having bunch crossings in which there is no pp collision at all. The minimum bias studies are therefore incorporated as a backdrop upon which to examine the beam backgrounds – and one can evaluate the effect that the beam halo and beam-gas events have on the minimum-bias triggers at the same time.

9.2 Minimum-bias events

As mentioned in Chapter 2, the $dN_{ch}/d\eta$ and dN_{ch}/dp_T behavior of the minimum-bias (inelastic, non-single-diffractive pp , for ATLAS)[90] events can be used to examine important parts of hadronization simulation [90] and its scaling, allowing us to compare our QCD background behavior directly with previous experiments. Minimum-bias events are a subset of the pp collisions in the LHC. The total cross-section for pp interactions can be written as:

$$\sigma_{tot} = \sigma_{elas} + \sigma_{sd} + \sigma_{dd} + \sigma_{nd}$$

where these cross-sections are elastic (σ_{elas}), single diffractive (σ_{sd}), double diffractive (σ_{dd}), and non-diffractive (σ_{nd}), respectively.

Since a primary source of bias in a detector such as ATLAS is the trigger itself, there are compelling reasons to take a high-statistics sample that minimizes the trigger’s inherent bias. There will still be inevitable biases determined by the detector’s imperfect geometric acceptance and energy resolution, but a trigger that selects all non-empty events would be a “perfect” minimum-bias trigger.

For reasons of compatibility with the results of experiments in previous hadronic colliders (further discussed later on in Section 9.3.3), the definition of a minimum-bias trigger has remained fairly fixed over time. One common feature of these experiments was the presence of a two-armed coincidence triggers. In UA5, for example, the high-acceptance-angle

streamer chambers [10, 9] were used to look for the differential rates of forward-only to forward-backward jets as indicators of the rates of single- and double-diffractive events. The triggers generally exclude the elastic cross-section completely, and it is ignored in minimum-bias studies. Although there is some variation between definitions of “minimum-bias”, depending on the experiment, a useful definition might be $\sigma_{nsd} = \sigma_{tot} - \sigma_{elas} - \sigma_{sd}$ – a minimum-bias trigger looks for inelastic, non-single-diffractive events.

Preliminary studies have shown that determining the number of inelastic pp interactions in a bunch crossing should be possible [67]. Typically the distance between pairs of adjacent pp interactions is much greater than the resolution on the z_0 impact parameter of a reconstructed track, so vertex finding can be achieved from a simple clustering of track z_0 values. For each non-diffractive inelastic pp collision it is predicted that there will be on average ~ 15 and ~ 35 charged particles produced within the inner detector acceptance for $p_T > 500$ MeV and $p_T > 150$ MeV, respectively. The probability of finding no particles per collision is small: $\sim 3\%$ for the case of $p_T > 500$ MeV and 0.4% for the case of $p_T > 150$ MeV.

k_b	N	β^*	Luminosity	Events/BC	P(0)	P(1)	P(> 1)
1	10^{10}	18	10^{27}	< 1	0.99	0.01	< 1
43	10^{10}	18	$4 \cdot 10^{28}$	< 1	0.99	0.01	< 1
43	$4 \cdot 10^{10}$	18	$2 \cdot 10^{29}$	< 1	0.99	0.01	< 1
43	$4 \cdot 10^{10}$	2	$6 \cdot 10^{30}$	0.76	0.47	0.36	0.18
156	$4 \cdot 10^{10}$	2	$2 \cdot 10^{31}$	0.76	0.47	0.36	0.18
156	$9 \cdot 10^{10}$	2	10^{32}	3.9	0.03	0.10	0.87

Table 9.1: Beam parameters for the Stage 1 physics run in 2008.

9.3 Backgrounds to minimum-bias events

9.3.1 Beam Gas

Beam gas and halo rates define the beam conditions. For various stages of accelerator operation, rate changes of several orders of magnitude in these effects are anticipated.

The beam gas interaction rate postulated in this thesis derives from simulations made by A. Rossi [68], regarding the rates of gas absorption and desorption by the Non-Evaporative Getter (NEG) beampipe coating – a layer of zirconium, titanium and vanadium, vapor-deposited on the inside of the pipe to pump away small gas remnants, increasing the vacuum quality. In various stages of run, and at various times after the initial bakeout, the NEG is closer or further from saturation, and its efficiency directly affects the beam gas particle density.

Other factors that affect the beam-gas “luminosity” are beam current and other debris in the beampipe. Current is relatively easy to factor in, as it is a critical variable and its progression is well-planned. Debris (and indeed the densities and species of the gases in the pipe) are not easily anticipated, and all statements we make here about rates should be taken as speculative – the density numbers may be off by a certain factor (though the proportions should be relatively reliable) and the contaminants question (any gases or debris other than H₂, CO, CO₂ and CH₄) is not addressed in [68]. The familiar components of atmospheric gas proportions, however, are absent – the source of contaminant gases in a UHV system is usually adsorbed gases (H₂, CO, CO₂, etc) desorbing from the surface.

As such, we have simulated small, arbitrary contributions from other contaminants that might reasonably be found in the beampipe at some point. Without run experience, there is no concrete basis from which to cite any specific rates for any of these debris channels. Their *raison d'être* is to be part of an extended toolkit that anticipates possible user needs.

From the beampipe itself, we can imagine having beryllium (Be), copper and iron (Cu and Fe), and other possible leakage indicators, such as nitrogen (N). From the NEG, we take Ti, Zr and V, and from the tertiary collimator jaws, Cu and tungsten (W). While these samples were produced, no investigation of the “exotic” beam-gas interactions will be made in the current

scope of this document – there is no solid expectation that such debris will ever actually be present, as the vacuum will be impressive and the beampipe walls very cold in the interaction region, increasing the getter effectiveness. After seeing the preliminary results of this study, it was decided to neglect these higher-Z contaminant channels for the time being, as it is not expected that they will show materially different behaviors.

To determine the beam gas interaction rate, we take the standard beam gas interaction formula [69] and adapt it to our circumstances. The head-on collision of 7 TeV proton on a stationary target has a a centre of mass energy of 115 GeV, giving a pp cross-section $\sigma_{pp} \approx 40$ mb. σ_{pA} is the cross-section for a collision of a proton with an atom of atomic mass A, and is approximated as $\sigma_{pA} \cong \sigma_{pp} \cdot A^{0.7}$. The rate of beam gas interactions can then be calculated using:

$$R_{pA} = (f_{LHC} n_b N_p) \cdot (n_{gas} \ell_{ATLAS}) \cdot (\sigma_{pp} A^{0.7}) \quad (9.1)$$

where f_{LHC} is the LHC revolution frequency, n_b is the number of bunches, N_p is the number of protons in a bunch, n_{gas} is the target gas density and ℓ_{ATLAS} is the length of the interaction region of interest. Estimates for beam gas rates for the gases given in [68] and other impurities during start-up and after machine conditioning are given in Tables 9.2 and 9.3, for initial run beam conditions corresponding to 43 and 156 bunches, and for normal running with 2,808 bunches at nominal beam current. As noted in the captions, each bunch configuration has its own bunch population – the number of protons expected per bunch. All these calculations correspond to beam-gas interactions along the full 23 meters of the beampipe.

Some clear trends are seen – hydrogen dominates after beam conditioning has happened, but the other contaminants are more important by far at the beginning of a run – especially methane. There are actually more interactions after the beam conditioning than before. One may guess that the events will be less spectacular from H_2 than from CH_4 , at least when the carbon nucleus is involved – but this guess is not supported by simulation. See Figure 10.11 for an illustration of charged-particle multiplicity rates from the beam-gas samples.

The beam gas events can occur anywhere along the length of ATLAS cavern at $z = \pm 23$ m, meaning that there can be up to six points along z , (± 7.5 m, ± 15 m, ± 22.5 m), where the beam

Gas	Bunches		
	43	156	2808
N_p	$4 \cdot 10^{10}$	$9 \cdot 10^{10}$	$1.15 \cdot 10^{11}$
H_2	$4.97 \cdot 10^{-2}$	$5.66 \cdot 10^{-1}$	65.1
CH_4	$5.70 \cdot 10^{-2}$	$7.08 \cdot 10^{-1}$	65.1
CO	$2.10 \cdot 10^{-3}$	$1.92 \cdot 10^{-2}$	1.00
CO_2	$3.90 \cdot 10^{-3}$	$4.11 \cdot 10^{-2}$	3.31
Total	$1.14 \cdot 10^{-1}$	1.34	135

Table 9.2: Beam-gas interaction rates expected in ATLAS – machine startup (Hz). N_p represents the number of protons per bunch for the calculation, reflecting the bunch population changes relating to the various stages of LHC operation. Results derived with Equation 9.1.

gas events are in time with the bunch crossings at $z = 0$. Other interactions along the beam-line will be out of time by up to 12.5 ns. The effect of the in-time and out-of-time interactions still need to be evaluated with full detector simulation and reconstruction. Tables 9.2 and 9.3, however, show that the expected rates during start-up are very low, even allowing for large uncertainties – on the order of a few Hz.

Beam gas events were generated using the HIJING [58] event generator, version 1.383, appropriately modified for use in Athena and with the collider characteristics of the LHC. Based on the LHC machine specifications [68], several targets which are expected to be present in the residual gas density for pp runs were used to generate beam gas samples. The principal component of the beam gas will be desorbed hydrogen (H_2), at between 10^{10} and 10^{11} molecules/m³. Other components such as CO, CO_2 and CH_4 will also be present, with a combined effect of $\sim 20\%$ of the total rate. What really matters for the purposes of the simulation, in the particular case of the residual gas, is the total number of atomic nuclei of each species – a relativistic proton will not be affected by any molecular-level structure.

Gas	Bunches		
	43	156	2808
N_p	$4 \cdot 10^{10}$	$9 \cdot 10^{10}$	$1.15 \cdot 10^{11}$
H_2	$3.87 \cdot 10^{-1}$	3.30	130
CH_4	$4.71 \cdot 10^{-2}$	$4.05 \cdot 10^{-1}$	18.6
CO	$9.90 \cdot 10^{-3}$	$8.08 \cdot 10^{-2}$	2.34
CO_2	$1.70 \cdot 10^{-2}$	$1.40 \cdot 10^{-1}$	4.06
Total	$4.62 \cdot 10^{-1}$	3.94	156

Table 9.3: Beam-gas interaction rates expected in ATLAS – after machine conditioning (Hz). N_p represents the number of protons per bunch for the calculation, reflecting the bunch population changes relating to the various stages of LHC operation. Results derived with Equation 9.1.

	Hydrogen	Carbon	Oxygen
Machine Startup	74%	23%	3%
After Conditioning	90%	7%	3%

Table 9.4: Beam-gas proportions anticipated for ATLAS – before and after machine conditioning.

Breaking down the various proportions of composite gases, we derive a proportion for nuclei in the beampipe of ATLAS (during normal run, after machine conditioning) as seen in Table 9.4.

Simulation of the beam-gas events is problematic for reasons of synchronicity. Beam gas events are not tied to any particular bunch crossing. They can happen anywhere in the beam volume, and can therefore arrive strongly out-of-time – an effect of interest in studying trigger behaviors, and in detectors with extended integration times.

Simulation using the ATLAS implementation of the Geant4 framework sees events as happening in-time with the bunch crossing, and as being a single initial vertex. To include greater timing complications, modifications had to be made to Geant4’s Athena interface, allowing it to notice the z position of the initial vertex and calculate a time offset $\Delta t = z/c$, where in this case the z position is from the bunch’s point of view, being from before (negative) or after (positive) crossing the interaction point (IP). These modifications were first introduced in Athena 13.0.10, with the `VertexTimeOffset` package in `AtlasG4Eng`. Substantial recompilation of the `Simulation` packages is required to use these options before the 13.X.0 releases.

9.3.2 Beam Halo

The beam halo simulation has recently become available, and with it the opportunity to add realistic beam backgrounds to the Athena simulation. Proton loss rates on the tertiary collimators (TCTs) can rise from $2 \cdot 10^6/\text{s}$ to $2 \cdot 10^8/\text{s}$ – going from one halo event in every 20 bunches to several per bunch. The high limitation arises from the energy capacity of the inner triplet magnets – past this loss level, the magnets quench and the accelerator shuts down. Beam

lifetime under these circumstances is 0.2 hours. As such, no reasonable rate estimation can be made before the accelerator actually begins operation. Rather, we will examine the effects of an exaggerated loss rate, so as to see a relatively pure sample of beam halo, and its effects outside the primary BG, clearly.

The pileup was set, independent of “foreground” event, to a rate of 0.8 – that is, the number of events chosen per bunch crossing in the pileup (which comprises 36 events before the primary BC and 32 after) will be a Poisson distribution centered around 0.8. This gives a 45% chance, in any given bunch, of no halo being chosen – giving a reasonable chance of having halo in the ID region at the time of an event. Halo particles will not reach the ID or the MBTSs from their origin at the 23 m scoring plane at the cavern entrance for three bunch crossings, so to have enough of them in flight is necessary to be able to have them interact with the detector at all.

The remaining 55% of the bunches will contain one or more halo events. More than one event per BC is certainly possible, and will certainly be the case in low-beam-lifetime operation (as the LHC operators figure out how to correctly tune the accelerator). Of those halo events, not all will make it to the detector – the events are generated behind the shielding, which has been specifically designed to block a significant portion of the halo to reduce in-detector noise and irradiation. Indeed, the MARS simulation of the halo was primarily destined for detector dose studies.

A new run of beam halo production was done, with FLUKA [61, 62] being used as the generator, appropriately modified to carry the parent particle identification in all of the downstream particles, allowing them to be associated into discrete bunches. In addition, the simulation included the presence of the tertiary collimators (see Figures 9.1 and 9.2), something that was not included in the beam optics (Version 6.4) used in the previous simulation. With the parent particle information, appropriate shower groupings can be made, better approximating the shower that might result from the loss of a proton against a gas molecule or collimator jaw inside the accelerator.

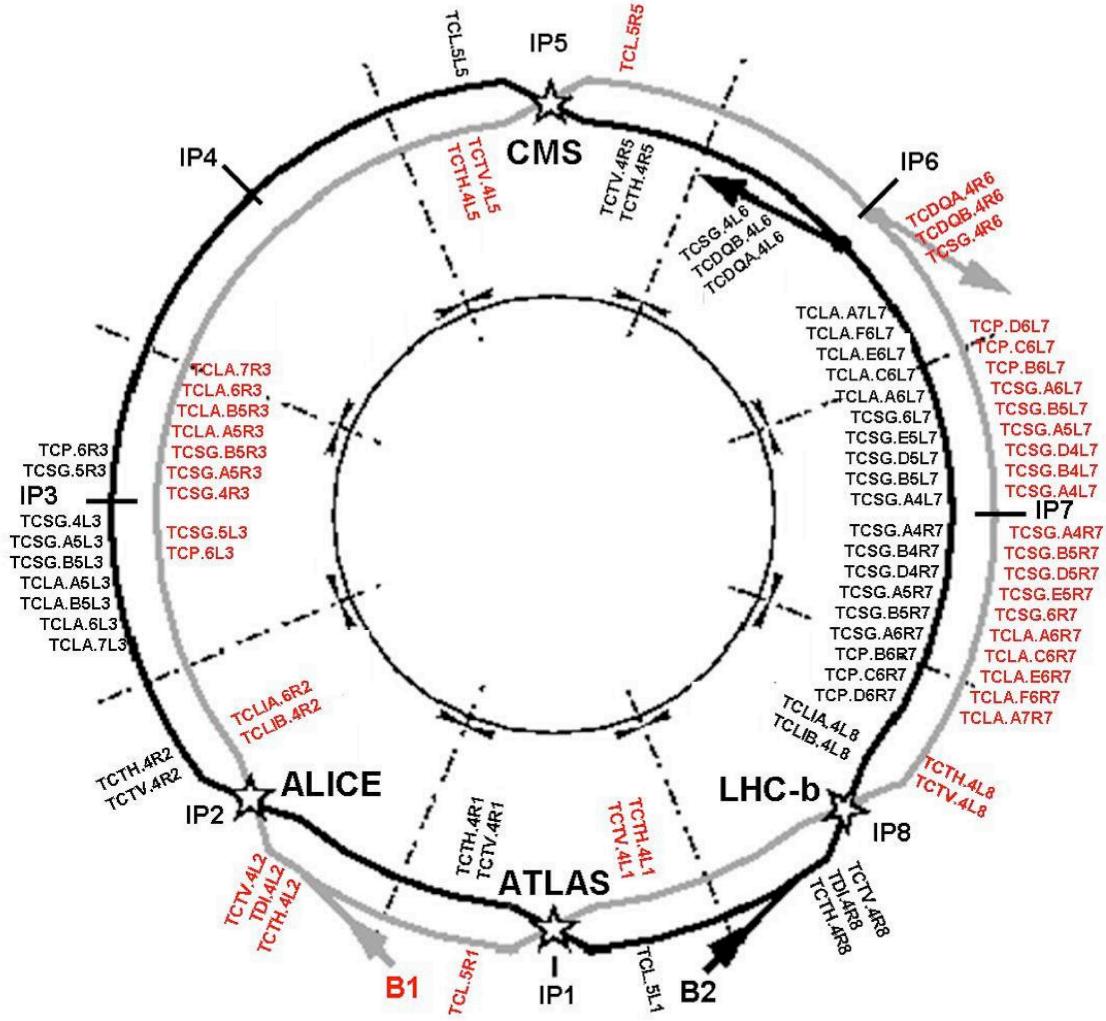


Figure 9.1: The Phase-I LHC collimator layout. On the left of the diagram are the collimators in place for momentum cleaning (catching protons with a momentum mismatch that are slowed by synchrotron radiation) – there are far fewer of them because of the less demanding application. On the right, a much larger number of collimators for each beam, which handle the betatron cleaning throughout the accelerator run, handling the inevitable proton losses arising from beam instabilities. The tertiary collimators (labeled TCTH1, for example) are the last line of protection from machine damage for each of the experiments, and come in pairs – one horizontal (TCTH) and one vertical (TCTV)

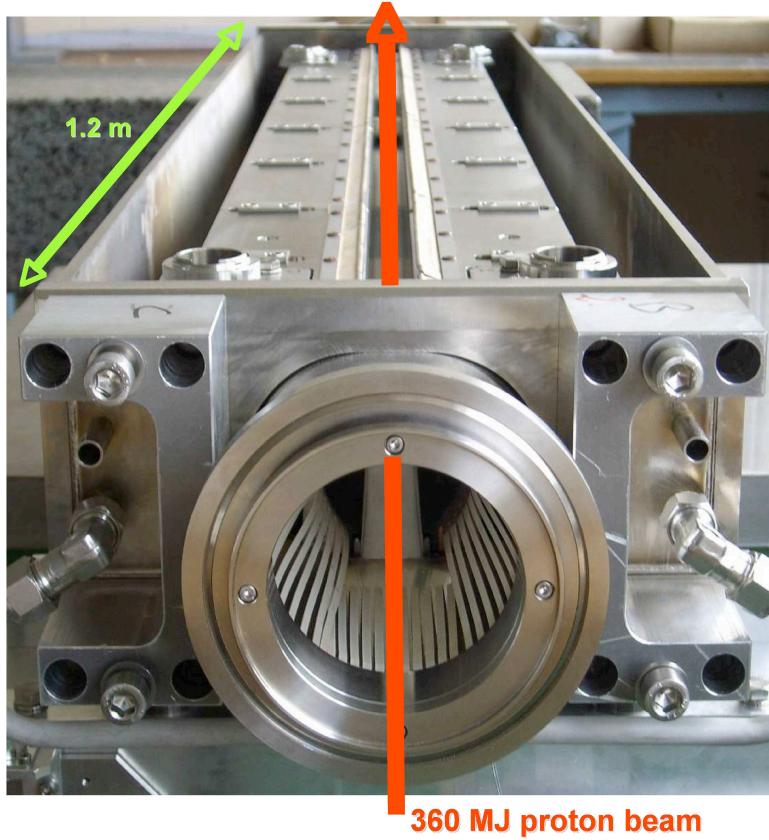


Figure 9.2: An example of a tertiary collimator – notable are the horizontal jaws, 1.2 m long, that open only under power as a machine protection feature in case of power failure. Also visible in the aperture are the fingers of the expansion joint that connects the collimator beampipe to the accelerator beampipe.

The halo rate is unpredictable by its nature, presenting possible variations of two orders of magnitude or more, depending on collimator settings and beam stability. For a normal LHC beam lifetime (~ 20 h) [34], beam proton losses will stand around $4 \cdot 10^9$ /s. All but one in 2,000 will be picked off by the primary and secondary collimators. Those protons that shower from the primary and secondary collimators will be stopped by the tertiary collimators (TCTs). It is the showering from the TCTs, in addition to the beam-gas interactions in the accelerator,

that produce the halo. In LHCb, collimator losses are anticipated to be 97% of the halo [78] – similar numbers will apply to ATLAS.

For this reason, and because of the relative difficulty of including it again, beam-gas interactions in the accelerator were not added to this run. Since the residual gas distribution in the accelerator varies substantially, especially between warm and cold sectors, the simulation has to run over hundreds of meters' worth of beampipe, with varying gas densities. The rates are anticipated to be much less than the collimator contribution, so it can be neglected for the moment.

During peak losses, however, proton loss rates on the TCTs can rise from $2 \cdot 10^6/\text{s}$ to $2 \cdot 10^8/\text{s}$ – going from one halo event in every 20 bunches to several per bunch. This variability requires a flexible approach to study, which will be implemented using the pileup mechanism.

In the case of both the beam halo and the beam gas, the events have nonstandard points of origin, and can remain in the detector volume for a significant period – on the order of six bunch crossings (or 150 ns) maximum. At high halo rates, there can be several halo bunches transiting the detector at a time from each side. Using pileup digitizations takes into account timing offsets, with a maximum granularity of 1 ns (30 cm) between primary BCs (event vertices). The pileup mechanism has been modified to allow absolute rates of gas and halo particles to be injected as of Release 13.0.10 of Athena.

The halo displays some strong asymmetries in both the plane of the accelerator and in the detector's exposure from each side [70, 71]. These effects are not unexpected. For the asymmetries in x vs. y , the structure of the TCT pair is clearly implicated – the vertical halo from the first TCT is sharply suppressed by the horizontal TCT in its turn, while the horizontal TCT has nothing following it to create a similar effect (see Figures 9.3 and 9.5). The asymmetries carry further, we see in Figure 9.6 that the direction cosine dx of the particles in the tunnel area behaves differently in x and y .

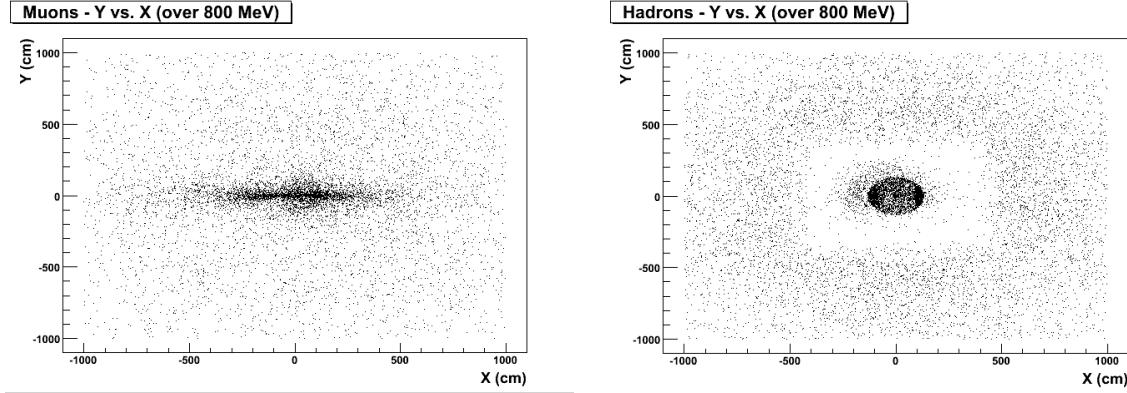


Figure 9.3: On the left, a plot of muons over 800 MeV on the x - y plane of the 23 m ATLAS cavern wall (the scoring plane.) The $x - y$ asymmetry is noticeable. On the right, the same plot for all non-muon components of the halo – note the visible “shadow” of the additional concrete shielding on which the beampipe shield is mounted, and the concentrations on the left and right side of the tunnel aperture. In the center, one sees a very strong dot – hadrons in this region tend to run parallel to the beampipe.

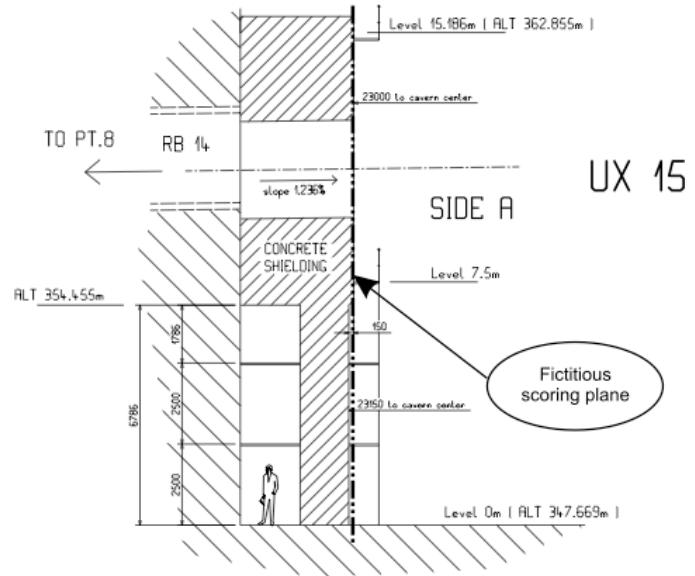


Figure 9.4: The scoring plane at 23 m in the ATLAS cavern (taken from drawing LHCJUX150003). Note the shielding nose support, a concrete structure, whose “shadow” can be seen in later plots.

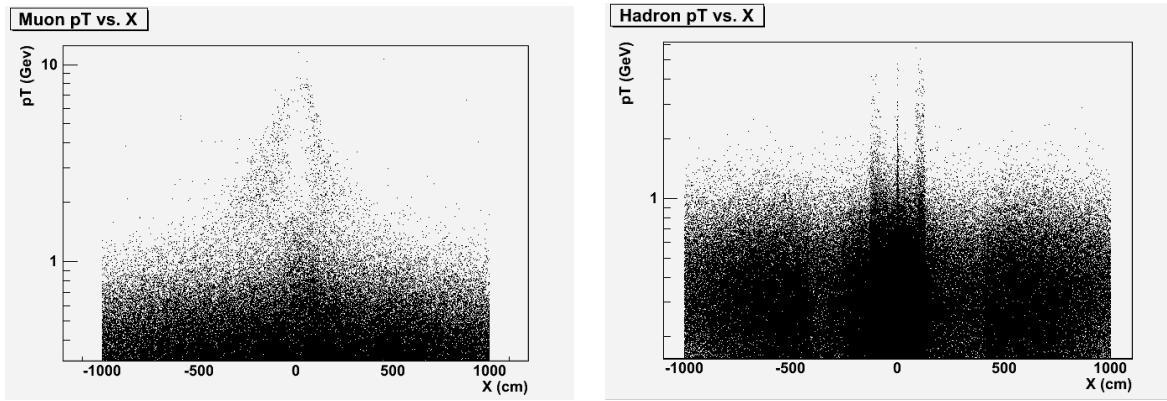


Figure 9.5: On the left, a plot of muon p_T with respect to x . The x plots were chosen to further illustrate the asymmetries mentioned in Figure 9.3 – the p_T vs. x plots show a simple spike at $y = 0$, with no further structure. The difference between the muon plots in Figure 9.9 is clearly illustrated here as well – high p_T is correlated with a clear path to the cavern.

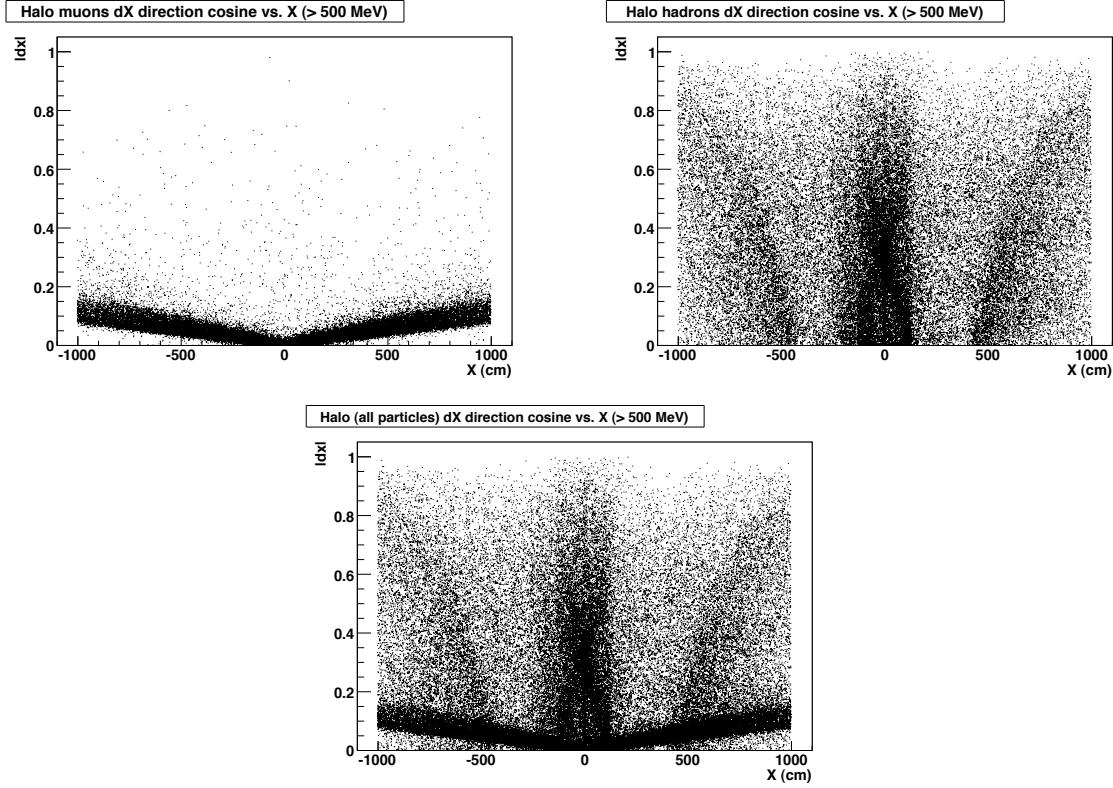


Figure 9.6: On the left, a plot of muon dx with respect to $x - dx$ being the x direction cosine of the particle as it emerges from the scoring plane. The muons' signature, in particular, is unmistakable, like the divergence of flashlight beam underwater – the muons are following very straight paths from their origin, as shown by the very small and consistent deviations from $dx = 0$. The hadrons, in contrast, have their dx widely distributed across the possible values. The interesting structure in the central two meters (the width of the tunnel aperture) is not present in the equivalent y plot – the two side-tunnel bright points mentioned in Figure 9.3 are in fact very collinear, and that collinearity reduces toward the center of the tunnel and the beampipe. The “wings” on the plot represent the sides of the concrete shielding support. Below, with the two signals together, it is quite clear that they are separate effects overlaid.

The halo itself breaks down into particles species as shown in Table 9.5. In contrast to the 2004 AT Division simulation, [59, 60], while the large majority of the particles now produced

are neutrons, muons compose 40% of the sample, rather than the 6% mentioned in those citations. The 2004 sample was made using the beam-gas interaction as its primary source – it's not surprising that this new sample should present a somewhat different profile. It is clear from Figure 9.6 that while there is some scattering of the secondary hadrons from the collimator collision that obscures the picture, the beam halo in this case really derives from a point source.

Some preliminary samples, terminating before the ATLAS cavern (for more widespread applicability) also contained particles such as deuterons and tritons, but these were absent from the final sample due to the material they traversed (or failed to traverse) in the remaining interval.

The particle species rates for the 1 s sample, seen in Table 9.5, corresponds to a total of 1,230,118 proton losses against the collimators, of which events 402,720 actually reach the ATLAS scoring plane (before the shielding). The number of daughter particles (the total of the columns in Table 9.5, 1,940,617), gives an average particle multiplicity of 4.2 particles per event arriving at the cavern. The majority of the events (from the perspective of the tertiary collimator source) don't even make it to the cavern. Of those that do, the vast majority have only one particle. There are, however, those whose multiplicities are thousands of times higher, clearly indicating that random choice of particles following, say, a Poisson distribution would poorly approximate a significant part of the beam halo. This can be seen in Figure 9.7. It is also clear that the total p_T of the high-multiplicity particles is not insignificant, but that the low-multiplicity events tend to contain the particles with the highest p_T . Examples of this behavior are seen in Figure 9.8.

Species	Total Particles	% of Sample
n	1,041,965	53.69 %
μ^+	394,222	20.31 %
μ^-	377,647	19.46 %
p	45,116	2.324 %
π^+	35,774	1.843 %
π^-	35,641	1.836 %
k^+	3,755	0.193 %
k^0	1,737	0.08951 %
\bar{k}^0	1,734	0.08935 %
k^-	1,556	0.08018 %
\bar{n}	743	0.03829 %
p^-	604	0.03112 %
λ	75	0.00386 %
k_L	28	0.00144 %
$\bar{\lambda}$	10	0.00052 %
σ^-	4	0.00021 %
ξ^-	3	0.00015 %
σ^+	2	0.00010 %
ξ^-0	1	0.00005 %
Total	1,940,617	

Table 9.5: Beam halo particles anticipated for ATLAS at the 23-m scoring plane on the A side of the cavern during 1 s of full-luminosity running, nominal β , Phase I optics.

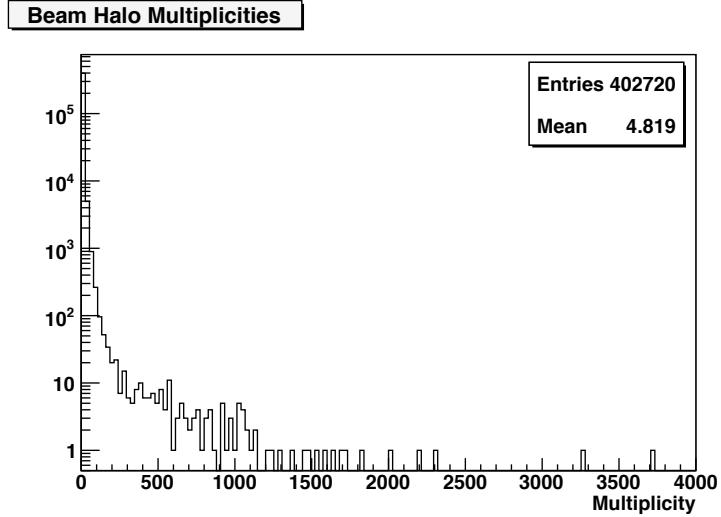


Figure 9.7: Multiplicity of beam halo events.

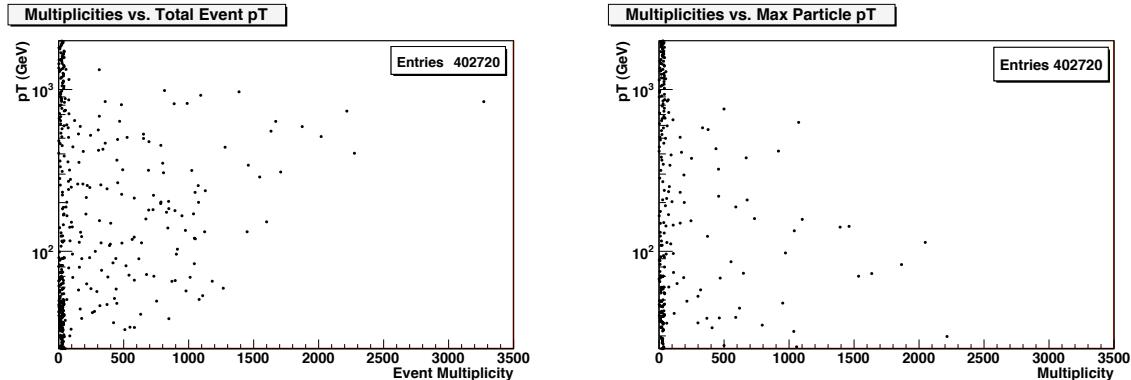


Figure 9.8: On the left, the transverse energy of the summed particles in an event plotted against the event's multiplicity. It is notable that low- p_T high-multiplicity events are not present. The same plot, but with the maximum single particle p_T in the event plotted, and the reason is clarified – most of the high- p_T single particles are part of low-multiplicity showers.

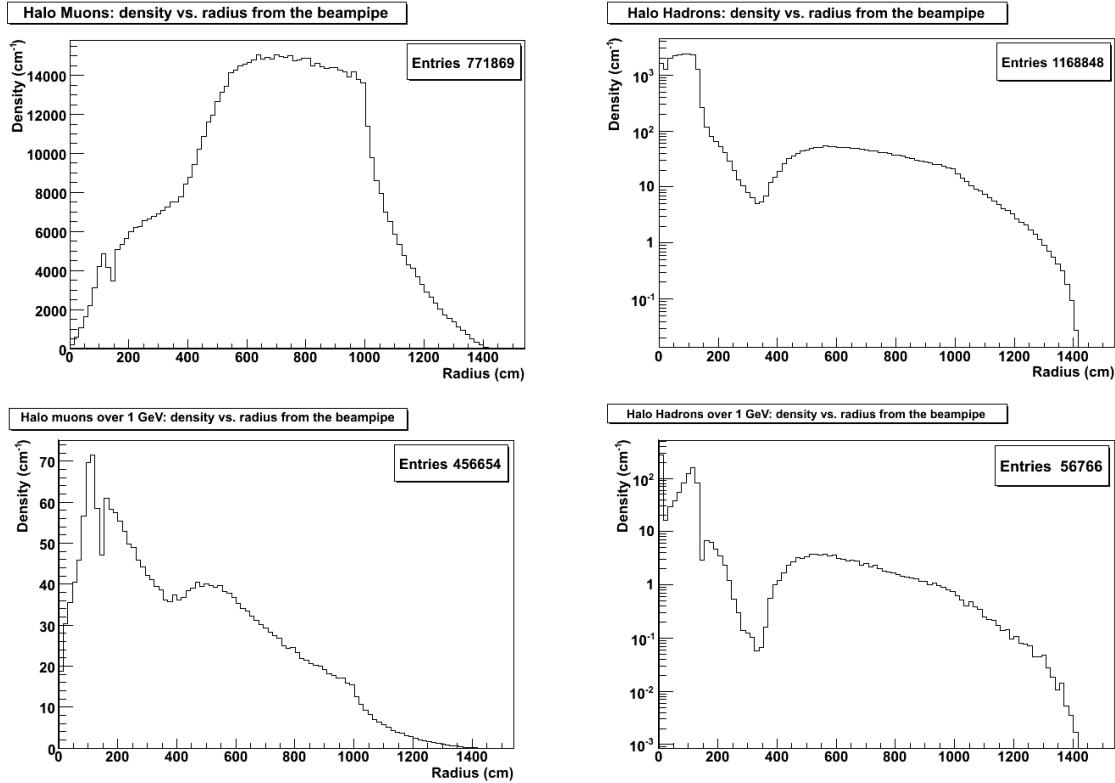


Figure 9.9: More plots clarifying the shape of the halo, and the differences between the muon and hadron components. In contrast to the whole particle set, the high-energy muons are much more centered around the beampipe, while the hadrons reflect a similar distribution, except for the strikingly high number of high- p_T hadrons in the first bin. The striking dip at between 200 and 400 cm in the halo hadron density (and for that matter, the shoulder in the total halo muon density) delineates the presence of extra shielding material in the shielding nose support structure, a roughly square block of concrete (illustrated in Figure 9.4), whose effect can be seen in Figure 9.3.

The beam gas and halo samples having been correctly produced and verified, the next step is to interleave them with empty events and inelastic pp events in the pileup mechanism as a part of the digitization. This step is particularly important in the case of the beam gas, where the time offset is really taken into account. The pileup tracks the state of the detector across a

number of bunch crossings, allowing the effects of previous events to leave their traces on the long-integration parts of the detector, such as the LAr Calorimeter and various drift tubes.

9.3.3 Triggering

The minimum bias signal, as a concept, is largely defined by its trigger – as mentioned before, its nominal reason for being is to produce a data stream with a minimum of trigger-induced bias. To make this work in an environment like the LHC is challenging. The first requirement is that the trigger find valid pp collisions with some physically interesting content – not just elastic collisions, but single-diffractive or double-diffractive events, as illustrated in Figure 3.1. It is intuitive to put the triggering system in the way of the jets seen in the figure – at moderate η to avoid damage from the hard forward QCD jets, but still be in the range of the jets. This principle was seen in UA5, where the high-acceptance-angle streamer chambers [10, 9] were used to look for the differential rates of forward-only to forward-backward jets as indicators of the rates of single- and double-diffractive events. It is also used at the CERN ISR [13] and at CDF at Fermilab [14, 8, 7, 6].

Much of what is done in the ATLAS minimum bias trigger system is consistent with those that have gone before, making results comparable. The basic principles are the same – collect events which exhibit strong QCD behavior (such as inelastic nondiffractive events, or single- and double-diffractive collisions) and examine them in comparison to previous results to understand the underlying physics and how they've changed in the energy interval that has been traversed.

A perfect minimum bias system would select any event at all where something happened, and completely exclude empty BCs. For initial ATLAS run, however, as seen in Table 9.1, there will be substantially less than one event per bunch crossing – in fact, the probability for initial run conditions is less than 1% that an event will be produced. The beams will be unsqueezed, and the beam current will be less than 1% of its nominal value – and there will be many fewer bunches. The ATLAS minimum bias trigger employs a two-arm system – the MBTSs described in Section 5.5.2, located one either side of the detector, between the inner

detector and the LAr endcap calorimeter cryostat, as an initial minimum bias trigger. It will allow reliable LVL1 triggering of useful minimum-bias events, and fully exclude empty events which would otherwise swamp the trigger system with useless readout. The MBTS and random trigger work together as illustrated in Figure 9.10. As the LHC beam runs, the random trigger fires at irregular intervals, using the beam pickup monitors (BPMs) [2] to verify a valid bunch crossing is present before sending its LVL1 signal. In parallel, the MBTS can also request a trigger, based on one of three conditions – at least one scintillator over threshold on at least one side (MBTS_1), at least two scintillators above threshold in the whole MBTS system (MBTS_2), or the two-armed requirement – one scintillator on each side over threshold (MBTS_1_1). All of these options introduce both a geometric and energy bias, but the less biased of the two trigger options is the MBTS_2 [46].

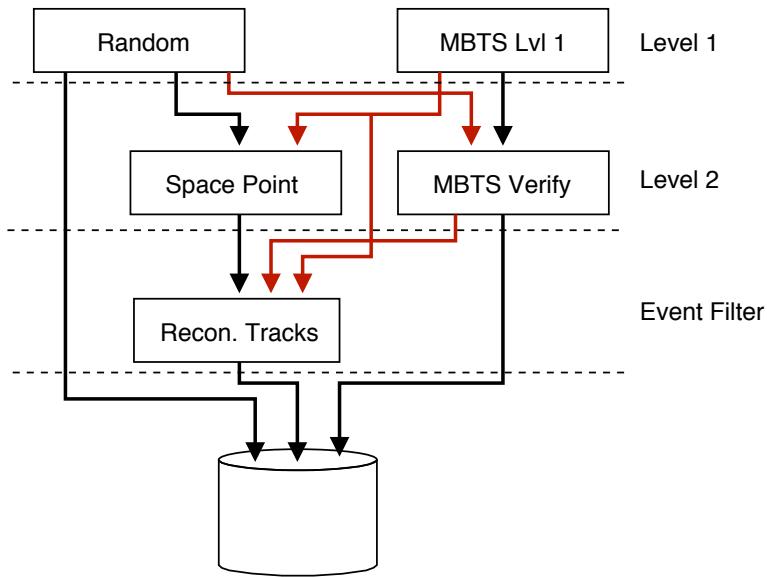


Figure 9.10: The ATLAS MBTS trigger strategy.

From the LVL1 trigger, the events are passed through a filter for number of SCT space points (shown to be an effective way of reducing empty event and synchronized beam-gas interactions at the IP [53]), then passing on to be reconstructed at the event filter. There are also direct paths to storage if so chosen. Past the LVL2 trigger are some Event Filter track quality cuts, which

further clean out machine backgrounds and cosmic rays. These triggers will be flags rather than eliminations, allowing for offline refinement of the dataset to avoid possible biases that may arise from the track cut itself.

The initial ATLAS run is expected to produce an event rate of 0.06 minimum bias collisions per bunch-crossing. At those rates, there should be very few double-collision events (a requirement for minimum bias analyses), and a need (as mentioned before) for reliably catching those events that do occur, rather than depending on the random trigger. The present proposed trigger menu for the minimum bias trigger is 2 Hz for the random + BC without any filtering, 4 Hz for random + BC passed through the space point filter, and 4 Hz for the MBTS LVL1trigger – a total input of 10 Hz. If rates are higher or lower than anticipated, these figures will be modified, favoring the random + BC trigger if there is a higher (~ 1) chance of pp collisions per BC, and favoring the MBTS LVL1 if rates are substantially lower than expected.

Chapter 10

Results

The minimum bias trigger was extensively investigated in [46], using certain early-run parameters for the anticipated collision rates at a luminosity of $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$, with a bunch spacing of 75 ns – as mentioned in the Chapter 9, a rate centered around 0.06 Hz – meaning that using Poisson statistics to determine the distribution, the rate for empty events will be about 94%, for one event will be about 5.7%, and so on. Background rates in pileup will be cited in this way from here on. These studies were done without pileup – that is, the only event included in the event (or any events prior or post-collision) was the primary signal event.

Since the primary goal of the present work is to come to an understanding of the physics behavior of the beam-gas and beam halo backgrounds to ATLAS for a range of possible uses, some additional samples and differing rates were added.

The first issue (as seen in Tables 9.2 and 9.3) is that the beam gas rate is expected to be very small. If realistic rates were applied in running pileup on a minimum-bias dataset, or even a set of empty events, it is likely that the necessary computational cost would be substantial, with little return on investment. The beam gas pileup used in these results is uniformly set to arrive at a rate following the Poisson distribution, centered at 0.3 events per BC (10.4 MHz) – giving a 74% chance of having no beam-gas event in a bunch crossing. This number was chosen to keep the rates high enough for study, while allowing the time necessary for the effects of any adjacent beam-gas collisions to clear from parts of the detector that are concerned.

If the LAr calorimeter, TRT, or muon system responses were addressed in the study, the rate would have been set lower still, to around around 0.075 events per BC (3.2 MHz), giving a 1:12 chance of having a collision for any given BC. This would allow both the maximal 7 BCs

necessary for a worst-case beam-gas event to clear the physical volume of the detector (though not the more pileup-sensitive systems like the LAr calorimeter) and a safety margin.

The beam-gas event rates used in this study will be unattainable under normal run conditions. Supposing that the vacuum is a factor of 100 worse than the estimates [68], and growing the interaction rate linearly with n_{gas} as shown in Equation 9.1, the background would only be seen at a rate of 15 kHz – beam-gas collisions in 0.04% of the BCs.

The beam halo is a different matter. As mentioned in Section 9.3.2, the halo’s rate dependencies lie in the tuning of the accelerator and the beam lifetime. Even at low event rates, the beam halo can be very high. At any given luminosity, it is possible to approach 200 MHz – corresponding to five tertiary-collimator proton losses per BC. The reason for this is that beam losses are not luminosity-dependent – but actually determine beam lifetime.

This raw rate is somewhat mitigated by both the shielding and earth between the TCTs and ATLAS, taking the event rate from 1.23 MHz (for nominal optics and beam lifetime) to a third of that ($\sim 400,000$ events/s), and the shielding on the ATLAS detector is expected to further reduce particle flux. Clearly, however, the expected loss rate can be safely assumed to be noticeable, so the inflated halo rates in this study are appropriate to its ends.

I chose a rate of 0.8 events per bunch crossing (around 22 MHz, in global terms) to give a high chance of one or more halo particles per bunch crossing – and to inflate the rate of the Beam 2 beam halo, whose 2% contribution remains obscure even so. This is not an unreasonable rate setting, considering that peak rates after culling the sample through the tunnel and cavern wall can peak at 4 MHz, with no reason to doubt that multiple overlapping halo will happen regularly.

Besides computing efficiency and visibility, one final reason to use artificially inflated background rates is to account for the worst case scenario. One aspect of this study is to examine the most difficult possible situations, and apply existing corrective techniques to rate their efficacy under these circumstances. The “pure” samples are also addressed in a subset of this analysis – but all of the datasets were also prepared with pileup set consonant with luminosities of between

10^{30} and $10^{31} \text{ cm}^{-2} \text{ s}^{-1}$ as well, to look at how the beam halo and gas will show up later on in the initial year of LHC operation.

Another point of difference between this study and previous minimum bias papers lies in the tracking reconstruction algorithms used to build the samples. The minimum bias analysis requires as large an acceptance, both in geometry and energy, as possible. It is inevitable that many charged particle tracks will not be accounted. Significant efforts, however, were put into refining the ATLAS standard reconstruction’s track resolution as much as possible. Track reconstruction becomes increasingly difficult with decreasing p_T – as the track radius decreases with momentum, the particles can begin to curve back in on themselves in the magnetic field of the solenoid. The ATLAS tracking software relies on the assumption that the particle will not curve beyond a certain point, and looks for track/hit matches within a defined “road” – and if the particle curves out of that road, the algorithm fails to find a realistic track.

This analysis does not include the low- p_T tracking optimizations that were included in [46] for several reasons. The first relates to the physics goals of this investigation – an understanding of the backgrounds to these processes will not benefit quite as much from detailed tracking performance as the analysis itself does. If there are indications that there might be merit in examining the issue further, they must be balanced against the new and experimental nature of the tracking modifications made. There is already one software tool in this study that is still in its shakedown process, and introducing more free variables is counterindicated. Finally, the low- p_T tracking modifications can be computationally expensive.

In place of the *de facto* 150 MeV cut that went into the minimum bias analysis of [46], the standard Athena reconstruction cut of 500 MeV is used in the tracking reconstruction. This does not exclude those particles from participation in other parts of the dataset, however – they are still there and capable of affecting the MBTSs, for example.

A final note is that this study was done with the ATLAS-CSC-01-02-00 geometry – nominally a geometry that includes all of the distortions and asymmetries known in ATLAS in its present form. The reconstruction, however, was also fed the same geometry model – so there

is no included geometric ambiguity and attendant errors. In essence, these datasets represent a perfectly aligned detector, in which all corrections have been perfectly applied.

10.1 Triggering and MBTS performance

Starting out with the samples, an important first sanity check was to plot the MC Truth “foreground” event’s distributions against the tracks that were reconstructed in the event. In this way, one can detect obvious problems in the production mechanism. This was especially crucial when using the new version of the pileup mechanism for its first analysis use. When it became clear by examination of the events and comparison to other distributions that there were no obvious problems with the dataset, trigger studies began.

The primary concept in a trigger – of any sort – is the idea of a threshold. This may be part of a whole menu of possible options, but it is important to understand the effects of each individual constraint in an experimental HEP trigger system, or one can lose both accelerator runtime and events.

The MBTS system, addressed in Sections 5.5.2 and 9.3.3, is a relatively simple triggering mechanism. When one or more energetic particles passes through the 3 cm of scintillator plate, it interacts with the dopants that have been added to the plastic and creates a small flash. This flash is passed along a fiberoptic cable and read off by a photosensor. In the most basic terms, the threshold should be cited in terms of light – but that’s not as useful as the output voltage of the photosensor that is translated into a signal for the readout system. The threshold we will use is therefore expressed in terms of electrical potential (mV). Before the upgrades cited in Section 5.5.2, the MBTS scintillators were attached to conventional Tile Calorimeter readout 3-in-1 cards, set to the low-gain setting universal to the TileCal. The threshold chosen to discriminate between real and spurious signal in an MBTS readout was chosen after simulating the response of the MBTS readout to empty events, and then choosing a value for the voltage that lay 3σ from the central value of the peak that formed, after fitting it to a Gaussian [46]. This threshold was between 7.5 and 7.75 mV, which corresponded to a noise acceptance (or spurious trigger rate) of 0.05%. This noise arises primarily from electronic variations in the photosensor

and in its readout electronics, and it is inescapable. The beginning of every plot in this section starts with a trigger acceptance (or number of events that the trigger considers to have passed its criteria).

The plots in 10.1 show the results published in for trigger threshold scans from 1 mV to 100 mV, for each of the four primary datasets used in the analysis. The nondiffractive set has very high efficiency all the way along for both triggers, and both exclude the empty events efficiently. The single- and double-diffractive acceptances are somewhat lower for the two-sided requirement. This serves as an illustrative baseline for the trigger threshold scans that will be done with the beam halo and beam gas samples, but some changes come into play. There is no pileup incorporated, either – these are events with no background.

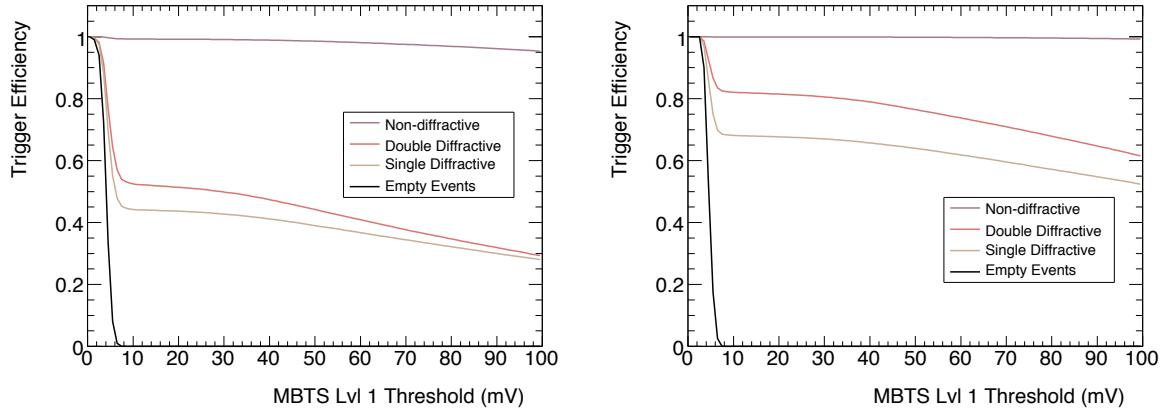


Figure 10.1: Trigger efficiency scans of the several diffractive samples for the MBTS_1_1 (left) and MBTS_2 (right) triggers, using the pre-Athena 12.0.65 digitization model [46].

Between release 12.0.6 and 12.0.65 of Athena, the software model was brought up to date, as previously mentioned. The new hookups for the MBTS system are expected to give better performance with a much better signal-to-noise ratio – but this is a change that needs to be taken into account. The new noise Gaussian fit is seen in Figure 10.2. Clearly the old threshold would be well within the noise band on this plot – so another one is chosen, further out at 65 mV (to be clear of the tails shown in Figure 10.2. These changes introduce a scaling difference on the threshold axis of any plots made after the change by a factor of 6.21, to account for the

increased range covered by the more capable electronics. This factor was chosen by contrasting the Gaussian in Figure 10.2 with its counterpart for Athena Release 12, found in [46].

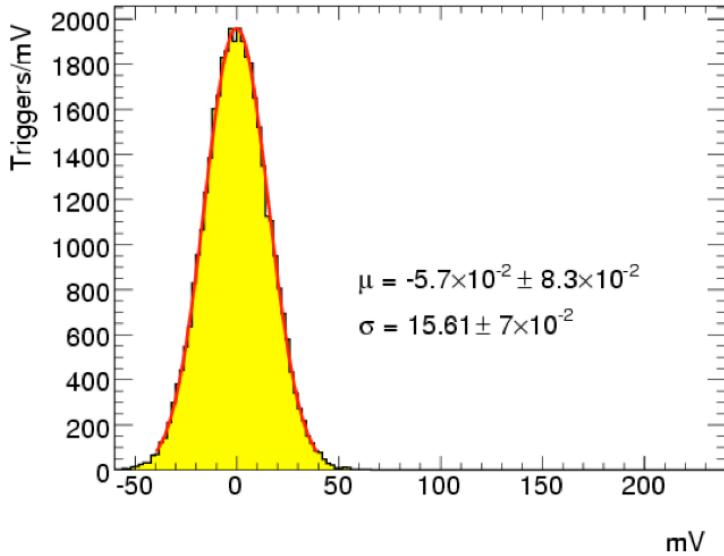


Figure 10.2: New MBTS threshold fit. Thanks to William Bell for the timely plot.

The closest this background analysis gets to having similar results is in the response of the MBTS trigger to the beam halo sample produced with no minimum bias pileup. The halo samples require some special procedures, however.

10.1.1 MBTS and Beam Halo

When reconstructing the raw halo samples, the simulated and digitized particles are started from the 23 m scoring plane at the far wall of ATLAS. The time necessary at c to reach the IP from that wall is 76 ns, or three bunch crossings. Nothing apart from some parts of the muon system are sensitive to anything contained in the event. Rather than depend on the basic procedure of taking the raw event and digitizing it, it is necessary to give the time in which to propagate – so they are placed into the pileup system without other components (the minimum bias pileup fraction, for example, set to 0), and set to an appropriately high event rate (0.8 events per BC, or 22 MHz).

What results (seen in Figure 10.3) is similar to that which was seen in Figure 10.1. For the first 45-50 mV of the plot, the system noise dominates, and if the system were run with in this threshold range, trigger rates would be untrustworthy. Before threshold (56 mV) the plot stabilizes at between 70% and 80% trigger efficiency for the single- and double-diffractive samples. For the “empty” events – the only dataset in this analysis that is truly empty of anything, as far as the MBTS response is concerned – there are no detectable signals in the MBTS system above threshold. Looking at the differences between the responses of the various trigger configurations in Figure 10.3, we see that the efficiency drop for the MBTS_1_1 configuration is the steepest of them all – very few halo events manage to trigger both sides – but even the MBTS_1 trigger (the most permissive) never manages to have a nonzero efficiency above 60 mV for the halo/empty events .

To avoid confusion, I’ll mention here that all other references to “empty” in this section are referring to the foreground event in pileup, not to the background processes with which it was “piled up”. This will also become clearer as we look at the trigger and tracking efficiencies of the various “empty” samples.

Figure 10.3 also serves as a good set of reference plots (especially the MBTS_2 plot), due to the fact that by itself, it appears that the beam halo is incapable of bringing the MBTSs above threshold by themselves. The events in the MBTS_1_1 sample are therefore almost unadulterated minimum-bias responses as individual, pileup-free events – and can be used to contrast to the plots of MBTS response in [46], with the previous MBTS electronics.

It is instructive to look at the foreground events to these halo samples, and compare them to the results seen in Figure 10.1. Other than the diminishing curve in the older plots, the responses of the MBTS_1_1 and the MBTS_2 triggers are much the same (at the accepted threshold of 65 mV) as the threshold response of the old MBTS_2 trigger at 7.5 mV – 100% for the nondiffractive sample, about 80% for the double-diffractive events, and about 70% for the single-diffractive sample. The MBTS_1_1 is also comparable in the same ranges. The efficiencies of the new electronics simulation are far more stable, but these plots indicate that the datasets are consistent – a good starting point.

We will see later on that the halo can have some effect in combination with other types of events – but that when there is no “foreground” event to enhance, the trigger response to the halo drops to zero.

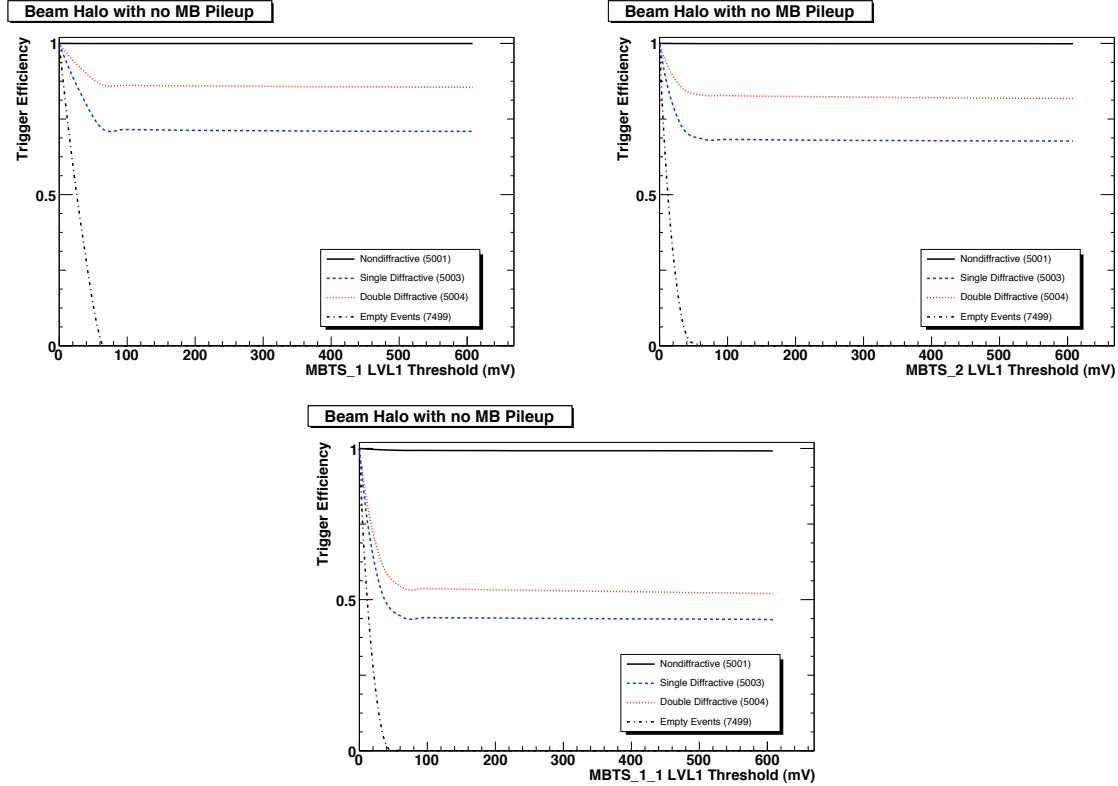


Figure 10.3: Trigger efficiency scans for the beam halo. Minimum-bias rates for pileup were set at zero to allow examination of the beam halo effects alone. Note the y scale difference in relation to the subsequent plots – from 1 to 0 rather than 1 to 0.8

10.1.2 MBTS and Beam-Gas Collisions

From the pure beam halo sample, we move on to the beam-gas samples, seen in Figure 10.4. There is no special need to run the pure gas samples through pileup before reconstruction – it is not anticipated that the beam gas will ever run under conditions where coincident beam gas events would ever be a concern, and the time offset of the beam gas location from the IP has already been taken into account. Essentially, the particles’ timer has been wound back to before

the collision, and its collinear daughter particles will arrive at the IP in sync with the beam. All that is done to this sample, therefore, is to digitize it and perform a trigger scan on the resulting MBTS values.

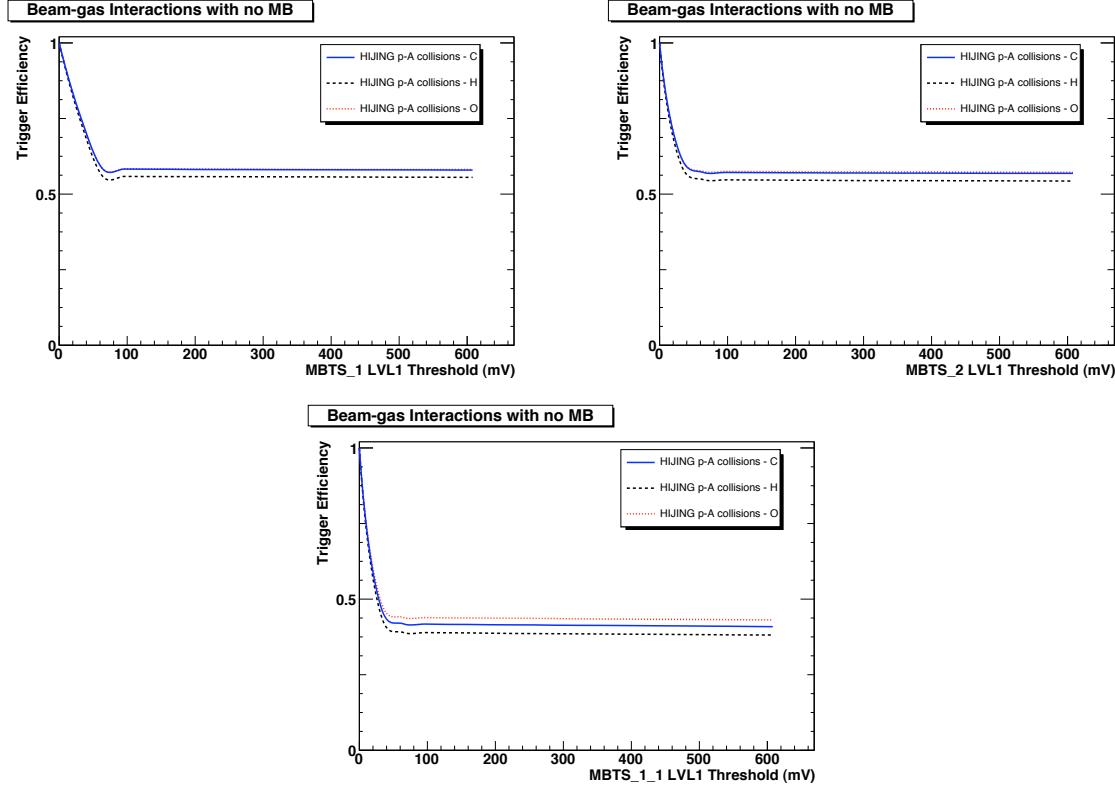


Figure 10.4: Trigger efficiency scans for beam-gas events. These events are not piled up, but they are randomized along the beamline. One can again note the y scale difference in relation to the subsequent plots – this covers the gamut from 1 to 0.

One of the interesting questions about the beam-gas events is what difference the Z of the target nucleus will make in the collision. There is a small but noticeable difference, here, between the H ($Z=1$) and the O ($Z=8$) and C ($Z=6$), and as one might expect, the higher- Z elements increase the trigger efficiency by a small margin. In either case, their trigger efficiency is about 0.57, remaining stable across the whole dynamic range of the scintillator readout, for both the MBTS_1 and the MBTS_2 settings. The MBTS_1_1 trigger manages to separate the responses a bit more, but it is clear that the difference in trigger efficiency between hydrogen

and oxygen is on the order of 3%. The differences in event structure between the different gas species will be examined in 10.2.1.

One notable feature of these plots, which will remain consistent through the next section, is that the best trigger to select for single- and double-diffractive events is the MBTS_2 (as seen in Figure 10.1). The best for eliminating these backgrounds, however, is clearly the MBTS_1_1, which is substantially less responsive to these generally single-sided contaminants.

10.1.3 MBTS and Minimum-bias Pileup

In implementing the threshold scan over the machine-background-enriched samples, the scan is run over each of the datasets in turn, as was seen in Figure 10.1. There are, however, more variations on the datasets to take into account. The beam gas and beam halo enrichment have to be run separately, since their frequency scales are so dramatically different that to overlay them would give meaningless results. We parse the data, rather, as an average look at how an individual beam gas or beam halo event will affect the trigger under various circumstances.

Figures 10.5, 10.6 and 10.7 scan the thresholds of, respectively, the MBTS_1, MBTS_2, and MBTS_1_1 triggers. In the pure minimum bias sample scans, an “empty event” really was nothing. In these plots there is a reasonable chance (as mentioned in the chapter’s introduction) that a minimum-bias event of one of the three channels we use will be in the event anyhow as a part of the pileup stream, along with the beam gas or halo.

It is worth a momentary aside to discuss this stability seen in both Figure 10.4 and Figure 10.3. Since the foreground events are well-defined and are not swamped by any inconvenient backgrounds, they are relatively “clean” – there are no out-of-time sources of trigger particles that would artificially inflate the trigger response. In the case of the plots in Figure 10.1, one notes that even after threshold, where one might expect stability, the curve still settles to a final value. This can be interpreted as an indication of the quality of the low-gain TileCal’s discrimination – even at the chosen gain setting, the trigger response remains fairly strongly voltage-dependent. The present high-gain system appears to be sensitive enough through its

whole voltage range to correctly distinguish events with regularity once the minimum threshold has been passed.

Contrariwise, there is the fact that the beam gas and beam halo events in Figures 10.5, 10.6 and 10.7 (which are the same set of plots repeated for each of the possible MBTS triggers) demonstrate that there are still circumstances where the trigger efficiency is not constant after threshold. This indicates that there are properties of the pileup component of the event stream that boost the amount of signal coming from the MBTS tiles at a wide range of input values – as the trigger threshold steps over the sample, it keeps finding more and more events to eliminate. In the minimum bias primary event samples, there is not so much to cut away – but in the “empty event” sample, these broad-spectrum backgrounds are all there is, and can continue to be cut arbitrarily (within the range of these plots) without reaching a core of real signal events that must trigger the MBTSs in any case.

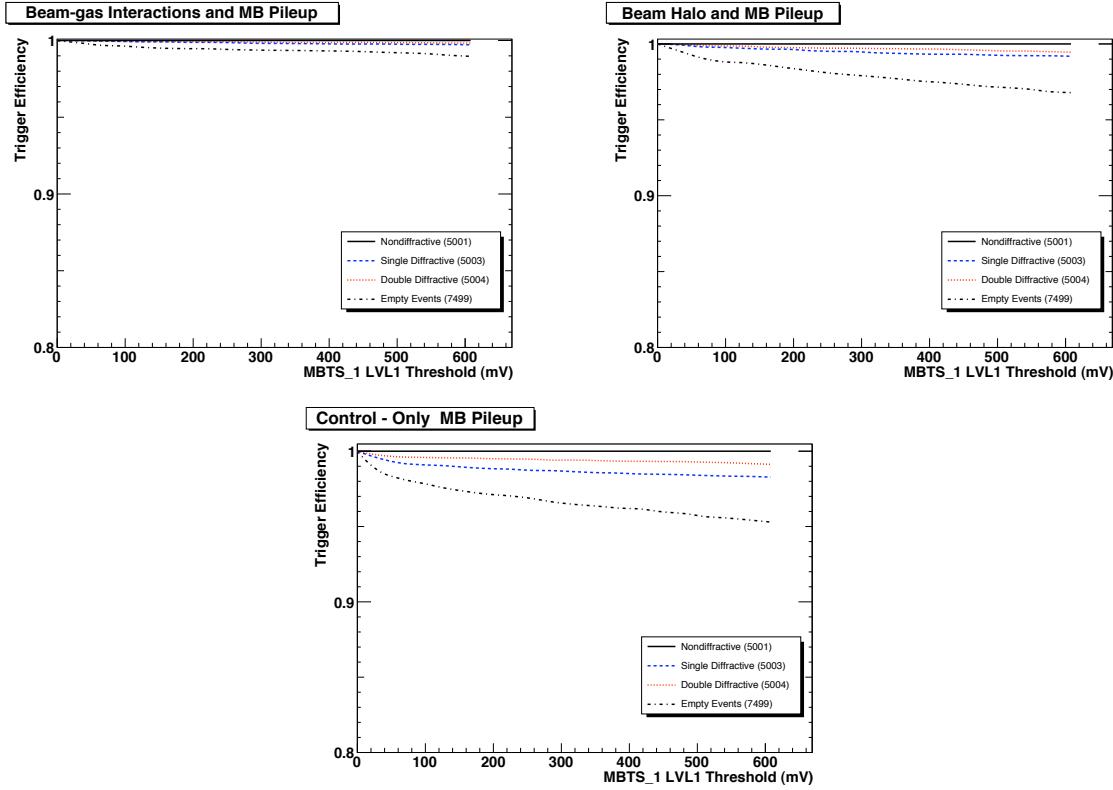


Figure 10.5: Trigger efficiency scans of the several diffractive samples for the MBTS_1 trigger. Top left is the scan for included beam-gas events with a Poisson centered at 0.3 events/BC, top right is the beam halo, with a Poisson centered at 0.8 events/BC, and bottom center is a scan of the MBTS_1 scan without added beam gas or beam halo.). The minimum bias event rate was set at 0.47 events/BC.

From these trigger response plots for the MBTS options, several things are immediately clear. The first – MBTS pileup at any significant rate will have a pronounced effect on the MBTS efficiency. Contrasting to Figure 10.1, the efficiencies shown in these plots are never below 0.88 – even for empty foreground events. Pileup at rates corresponding to a Poisson distribution around 0.47 events/BC will trigger a minimum of 94% of all BCs, no matter what the foreground event is. Certainly none of the desirable sets of minimum-bias trigger rates are diminished – but the number of faked empties will be substantially increased. Cleanup with other methods will be necessary.

The trigger response to the MBTS_1 option with beam gas and pileup incorporated is above 99%. Any event with beam gas and any MB pileup will trigger the MBTS_1 option. Even the best-case MBTS_1_1 trigger accepts 98% of the empty events it sees – this clearly illustrates the MBTS system’s limitations, and indicates that when pileup rates significantly surpass the ranges explored in [46], their utility decreases sharply. This is, however, when the random trigger comes into its own and becomes more useful in any case – here we are using the MBTS system to illustrate trigger responses as a means of characterizing the different backgrounds.

The implications here are that the beam-gas background can significantly enhance the false trigger rate for the MBTS system. It is notable that the beam-gas rate in these scans is substantially smaller than that of the beam halo, and yet its effect are far greater. Indications of why that is the case will be seen in 10.2.1.

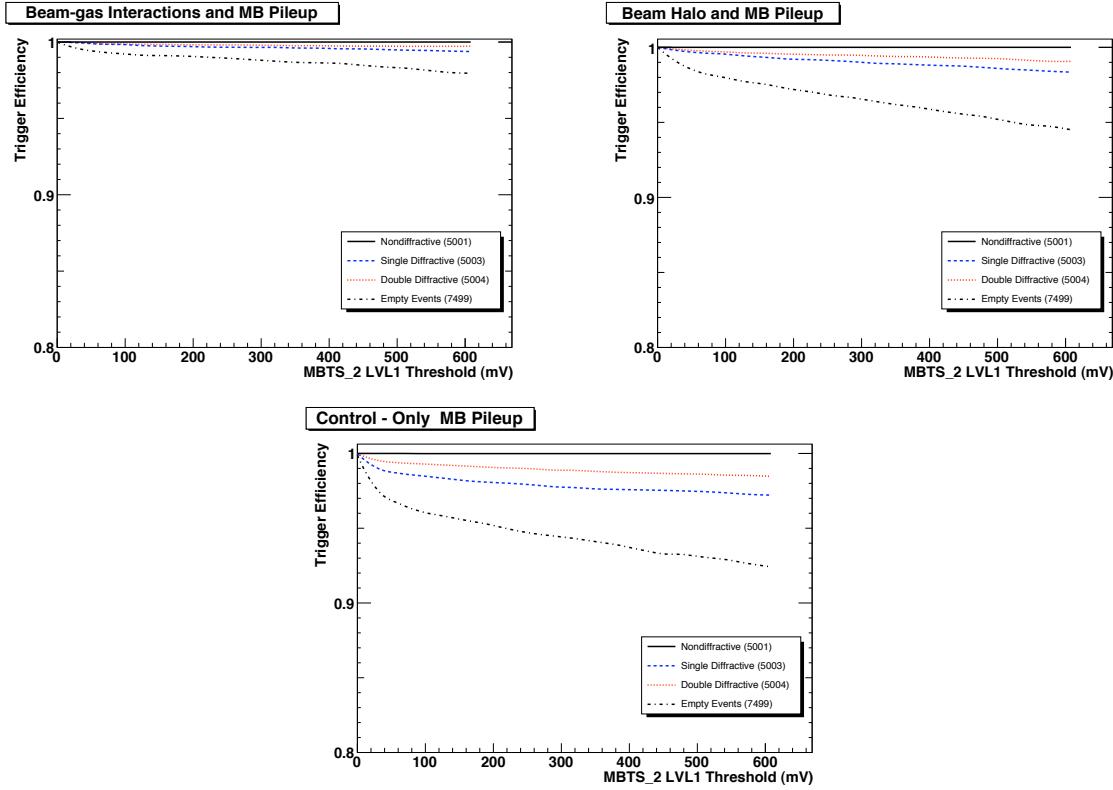


Figure 10.6: Trigger efficiency scans of the several diffractive samples for the MBTS_2 trigger. Top left is the scan for included beam-gas events with a Poisson centered at 0.3 events/BC, top right is the beam halo, with a Poisson centered at 0.8 events/BC, and bottom center is a scan of the MBTS_2 scan without added beam gas or beam halo.

While we saw in 10.1.1 that the halo response on its own was not a significant contributor to the MBTS response, it is clear from these figures that the halo can contribute to the trigger rate, raising the empty event rate by 2.5% for MBTS_2 and 16% for MBTS_1_1 at threshold. This is not a tremendous effect, but is much more noticeable than the effects from the halo pileup alone.

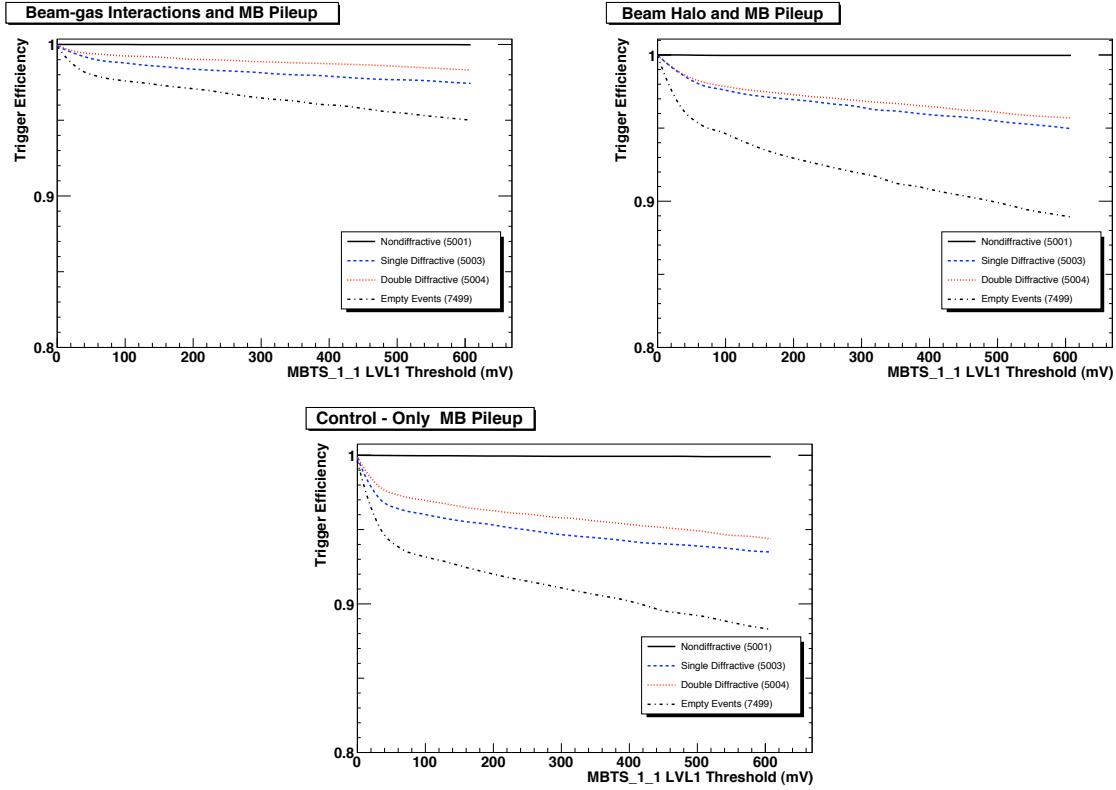


Figure 10.7: Trigger efficiency scans of the several diffractive samples for the MBTS_1_1 trigger. Top left is the scan for included beam-gas events with a Poisson centered at 0.3 events/BC, top right is the beam halo, with a Poisson centered at 0.8 events/BC, and bottom center is a scan of the MBTS_1_1 scan without added beam gas or beam halo.

Once again, we see that the MBTS_1_1 trigger gives the best background elimination, even at these high pileup rates. The choice of MBTS_1_1 and MBTS_2 trigger options will depend on whether run conditions dictate that it is more important to preserve signal events or eliminate backgrounds.

10.2 Tracking and Event Characteristics

Moving from using the MBTS system as a window into the behavior of the beam gas and beam halo backgrounds, we now look at their effects on the behavior of certain canonical plots

used as standard reference points in comparing minimum-bias event behavior across detectors. It is certainly useful to compare the ATLAS detector under various background conditions to the detector characteristics and physics behavior that we see in [46].

To make a legitimate comparison, of course, the same datasets that were used in [46] were used as the foregrounds and pileup sources for this study, as has already been mentioned. In addition, the same geometry (ATLAS-CSC-01-02-00) was used, and similar plots produced. Some technical variations arise from several sources, however. Primary among them is the use of the pileup mechanism, and its requirement that the analysis and event generation be done with release 13.0.30 of Athena, rather than with one of the 12.0.x series of releases which were used in [46]. In the same sense, the pileup machinery itself is still undergoing its shakedown in Athena release 13.0.x, though significant validation has been done, especially for the components of the detector that are crucial to this analysis. Finally, the simulation of the beam gas events was done using a slightly modified Geant4 interface to Athena, allowing the aforementioned beam gas time offsets. While bugs and malfunctions in these different tools have been identified and corrected in the course of this analysis, there is no reason to believe at this point that there are any significant underlying differences to be noted between these different toolsets, allowing direct comparison between the the tracking analyses of this study and the ATLAS Minimum Bias Computing Service Challenge (CSC) internal note.

The distribution of charged particle tracks in a minimum bias event can be looked at in various ways, but it is important to divorce the physics of the event from the detector, conditions and accelerator considerations as much as possible. To do this with the minimum-bias events, a standard set of plots includes a normalized plot of the charged particle densities as a function of η (also known as a $dN_{ch}/d\eta$ plot) and as a function of p_T (a dN_{ch}/dp_T plot). These quantities are important because multiple parton interactions are likely to shed part of their energy as low- p_T particles in the low- η regions, accessible to out tracking, in addition to losing a large part of that energy into the less-sensitive forward regions of the detector, or into the uninstrumented beampipe [90]. Another plot one commonly looks at is the charged particle multiplicity one sees in events.

These plots have been made for the normal, unadulterated minimum-bias events in the CSC note, and they can be readily seen there. They serve as a backdrop to the effects of the beam gas and beam halo, from which we can see the differences induced by particles that are almost certainly coming from trajectories that do not intersect the IP of ATLAS.

ATLAS tracking assumes that any track will originate within a certain volume around the IP, and bases all of its track-finding algorithms (in conventional reconstruction) upon that assumption. The tracks are formed by minimizing undesirable characteristics in a track parameter covariance matrix, which contains several track quality parameters. If the parameters diverge too much from a given ideal even after minimization, the track is likely to be spurious, and can be discarded by applying quality cuts. Such a set of cuts was proposed in the CSC note as a means of eliminating track fakes and tracks from secondary particles.

One important feature of these cuts (whose properties and effects are detailed below in 10.2.3 and 10.2.4) is that for events with heavy minimum-bias pileup, beam-gas or beam halo backgrounds, these cuts eliminate all tracks in the events. This complete elimination is further discussed in the appropriate subsections, but is relevant to these charged particle plots because they have not been cleaned using the CSC track quality selection. The simple reason for this is that any plots made after even a partial cleaning would be empty – few or none of the tracks would pass any of the substantiative cuts. This will account for some of the differences in features between the CSC plots and those hereafter.

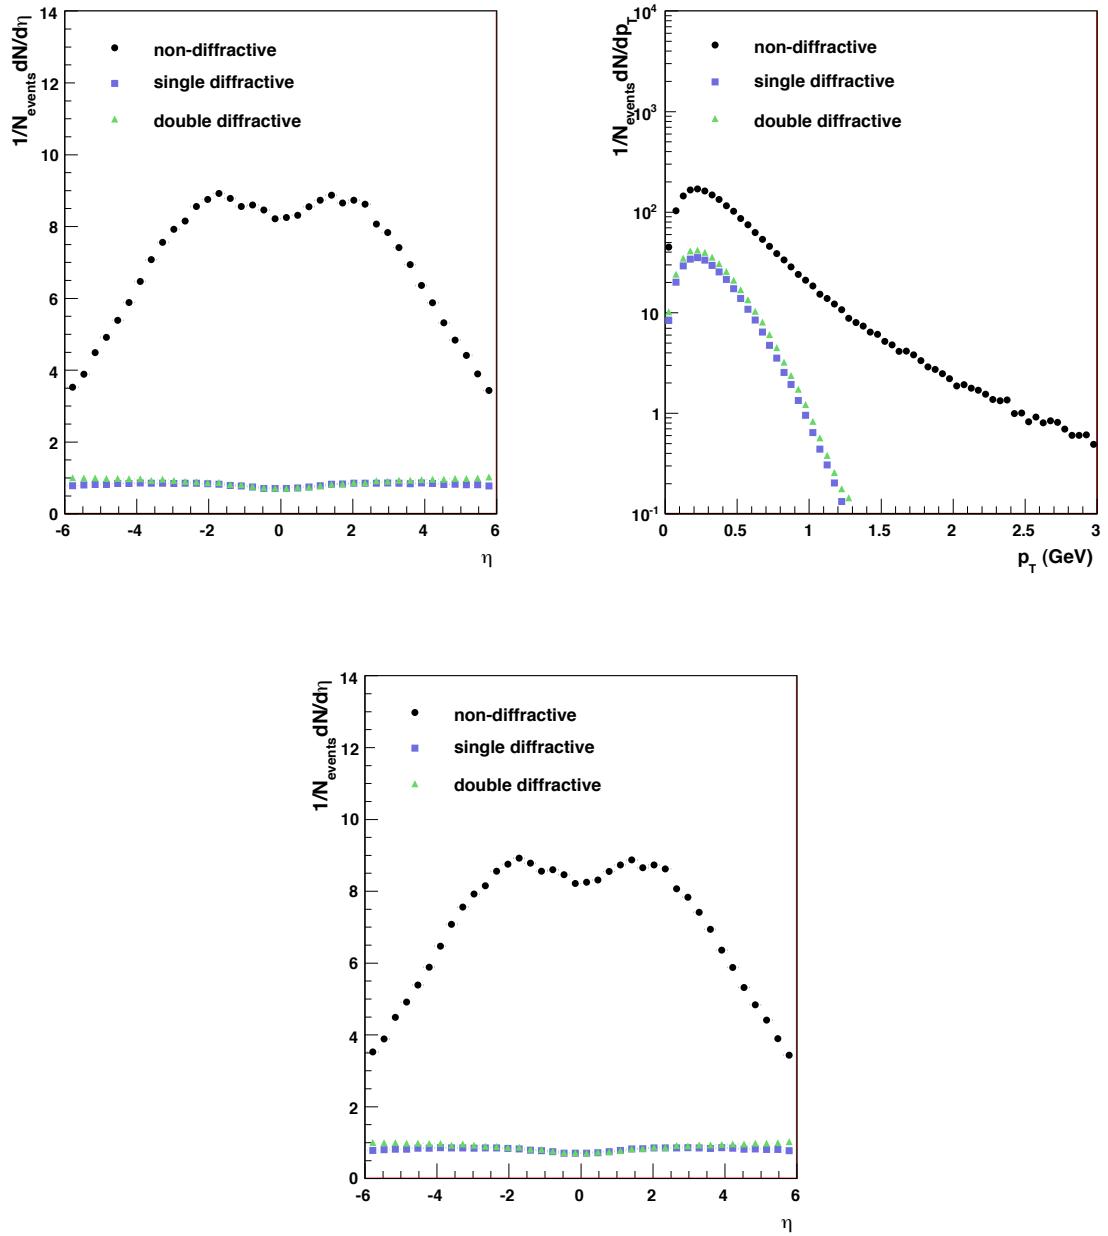


Figure 10.8: Canonical minimum bias charged particle plots. Taken from [46]. These plots represent generated particles that have not been taken through simulation and reconstruction.

10.2.1 Background Track Behaviors

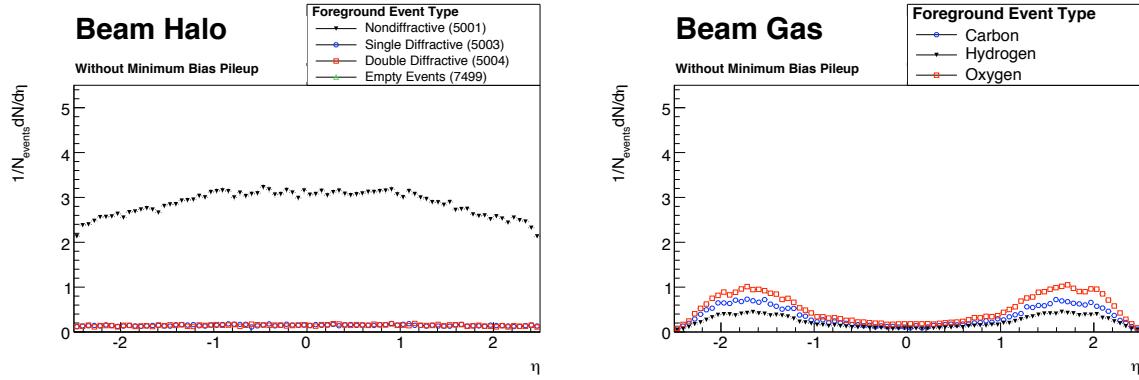


Figure 10.9: $dN_{ch}/d\eta$ – no pileup. On the left-hand curve, the empty-event plot has no tracks at all, and is not plotted. Please note for comparison purposes that the binning in this plot is identical to that in Figure 10.12.

We begin with the backgrounds themselves, to clarify their tracking behaviors and see their effects on simple foreground minimum-bias samples. The first plots in this section (in Figure 10.9) are of the beam halo on the left, and the beam gas on the right. The halo plot is the more complex one – the halo sample is here shown piled up over the three minimum-bias sample foreground events (without any minimum-bias pileup background), and piled up over empty events as well.

It is notable that the beam halo over empties produced no tracks at all – there are no points on the plot for that dataset; though it was run and incorporated, no bins were filled in the plotting. This is unsurprising – the beam halo particles are only capable of producing pointing tracks in the tracker endcaps, and very few halo particles have the necessary dx to spoof the tracking system. It would be surprising to see anything from the empty/halo combination. The same paucity of halo/empty tracks is reflected in Figures 10.10 and 10.11.

The other datasets (5001, 5003 and 5004) in the plot behave very similarly to the $dN_{ch}/d\eta$ plot seen its top left plot, once some issues are addressed.

The first and foremost of these issues is that the plots in Figure 10.8 are derived from PYTHIA-generated particles, and have not been through simulation and reconstruction - they are there to give a reference point for what an ideal detector would be able to discriminate. The η range stretches to six, and the central region shows the characteristic depression arising from the pseudorapidity coordinates.

The $dN_{ch}/d\eta$ plot for beam halo in Figure 10.9 has proportionately smaller (though directly comparable) behaviors in the 5003 and 5004 (single- and double-diffractive) samples. The drop in multiplicity rides on a couple of factors. One may be that the beam halo was responsible for corrupting the tracks with extraneous spacepoints, making the track fitter unable to fit some of the charged particle tracks in the detector.

The other arises from the simple fact that a large part of the particles that were produced by PYTHIA for the plots in Figure 10.8 are in large part below the 500 MeV cutoff for standard Athena reconstruction, and as such were lost – this can be seen clearly in the top-right dN_{ch}/dp_T plot. We can also infer that many of these lost particles were in the $|\eta|$ range of between one and two. Further losses stem from tracking inefficiencies. The losses constitute roughly a factor of four in charged particle track density with respect to $|\eta|$. Because of the number of possible sources for this degradation and the fact that deconvoluting them is neither practical, nor within the scope of this work, I will make no further attempt at this time to further refine the track efficiency.

The beam-gas collision plot on the right of Figure 10.9 is not directly comparable to anything in the CSC note, in contrast. As mentioned before, the beam-gas events needed no pileup for this step in the analysis, and are shown as individual events. Clearly, the majority of the tracks created by the beam-gas events have very few reconstructable charged tracks, and those that can be reconstructed are largely found in the portions of the ID where the endcaps are most prevalent – indicating (unsurprisingly) that they are more susceptible to track reconstruction. The number of IP-centered beam-gas events will be very small, but the particles from anything off-center will likely strike the endcaps at a reasonable angle, making fake tracks far more likely in that part of the detector than in the barrel, where beam-gas particle trajectories from off-center

events will be near-parallel to the ID sensing planes. The rates of track reconstruction are also likely to be suppressed by the out-of-time nature of the beam halo, which greatly increases the likelihood that particles from the event will not arrive in-time with a bunch crossing, and will therefore not be counted in the trigger timing windows of the faster components of the ID.

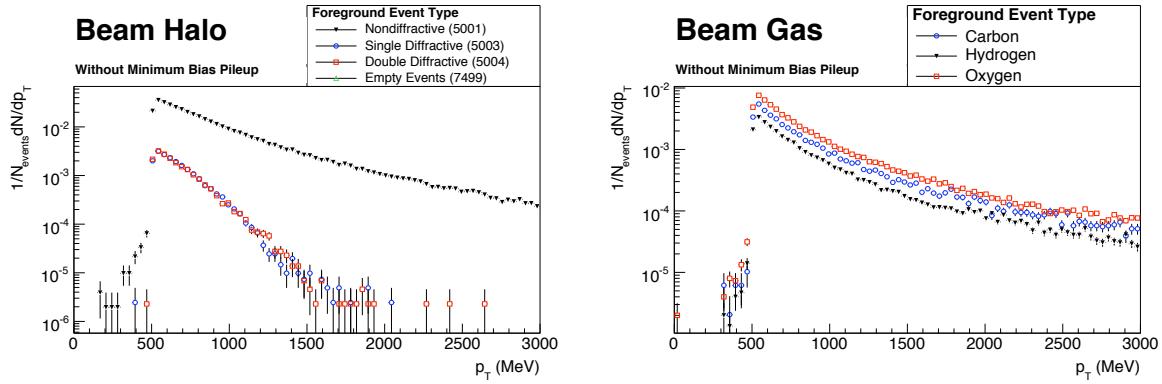


Figure 10.10: dN_{ch}/dp_T – no pileup. It is notable in the left-hand plot that no tracks were created by the empty events with a halo overlay.

When accounting for the 500 MeV track reconstruction cut, the beam-halo dN_{ch}/dp_T on the left in Figure 10.10 looks very similar to its counterpart in Figure 10.8 – except for the charged-particle densities themselves. After reconstruction and contamination, p_T cuts and (especially) the geometric cuts that remove the lion’s share of the charged particles in the PYTHIA generated set, the normalized density of particles drops from the hundreds to the hundredths. The relative proportions of non-diffractive to single- and double-diffractive (SD and DD) tracks remains similar, however. On the right, the beam-gas events produce a curve similar in magnitude to the single- and double-diffractive events, but which falls off much more slowly, showing a much higher rate of high- p_T particles.

In all of the p_T plot in this and subsequent sections, there is a proportionately small number of tracks shown with a p_T smaller than 500 MeV. Such sub-threshold particles were elided in the CSC note by a hard cut (though in that case the threshold was 150 MeV) – I leave it in place to illustrate the nature of the cut as arising from tracking constraints rather than analytical fiat.

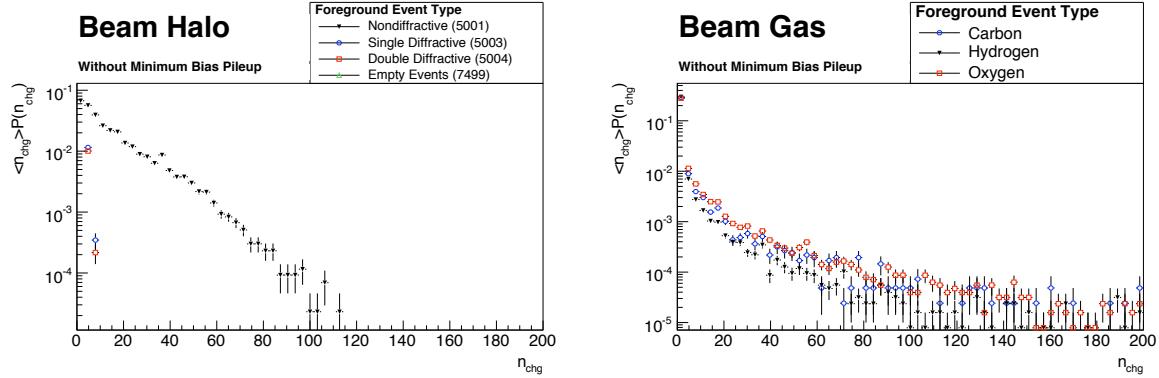


Figure 10.11: Charged particle multiplicity – no pileup. It is notable in the left-hand plot that no tracks were created by the empty events with a halo overlay.

The charged-particle multiplicity plots in Figure 10.11 are strongly limited by the lack of viable tracks with which to work. For the SD and DD datasets, the plot follows the same lines as one sees in the bottom-center plot in Figure 10.8, but the number of available tracks is simply too low to plot. For the beam-gas events of the right of Figure 10.11, the lion’s share have low multiplicities of reconstructed particles (unsurprisingly), but occasional high-multiplicity events represent beam-gas interactions within the inner detector that provide reconstructable tracks in higher numbers.

Across all of these plots of beam-gas interactions, one notable feature is that though the variations between the high- and low-Z targets are visible in the plots, they are not large, and are all attributable to differences in number of charged particles produced.

10.2.2 Minimum-Bias Pileup Track Behaviors

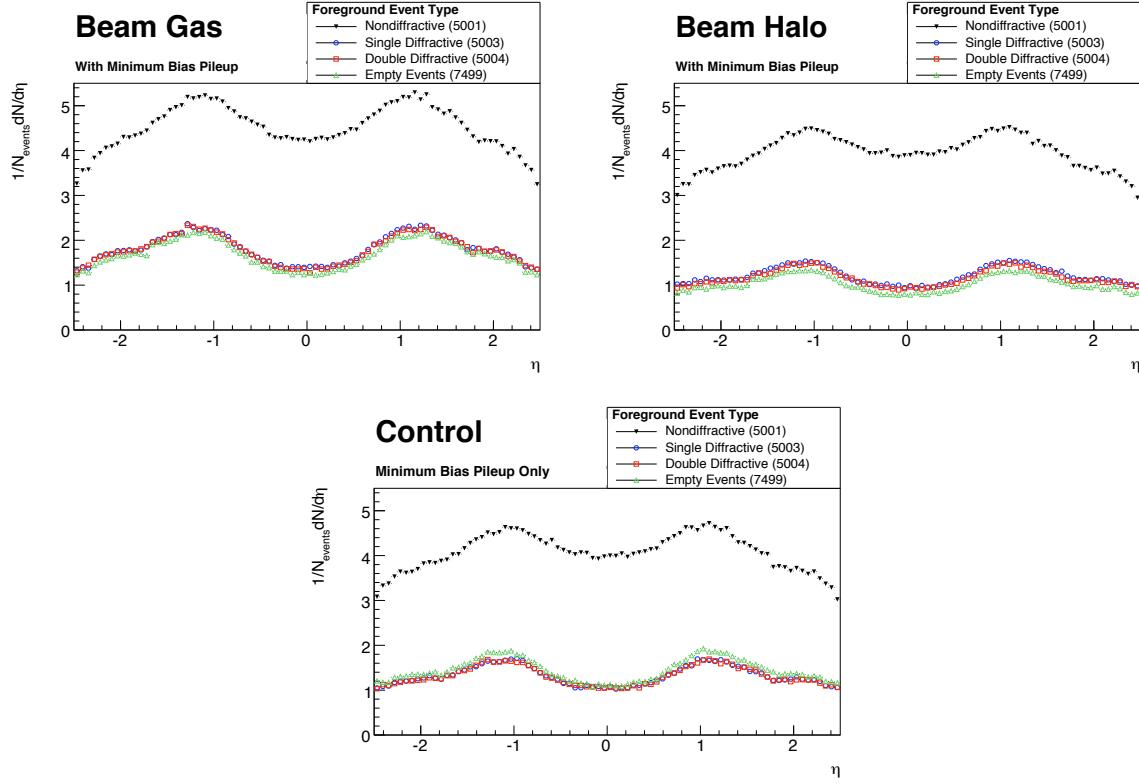


Figure 10.12: $dN_{ch}/d\eta$ for normal pileup. Please note for comparison purposes that the binning in this plot is identical to that in Figure 10.9.

With the introduction of minimum-bias pileup background, the picture changes substantially. In comparison with the halo plot in Figure 10.9, the charged particle density w.r.t η nearly doubles for the ND dataset, and rises between seven and ten times for the SD and DD samples. The clean cause of all this can be seen by looking at the empty event $dN_{ch}/d\eta$ plot in Figure 10.12, where the baseline minimum-bias pileup effect is seen. That rising tide lifts and forms all the curves in the plot, providing humps at $\eta= 1$. The beam halo has a minimal effect, in comparison, on the plots – in fact, it slightly depresses the densities of the SD, DD and empty samples. The beam gas events augment the track densities at the $\eta=1$ humps and

slightly elevate the curve outside the central region – clearly, this is an additive effect between the bottom-center plot in Figure 10.12 and the right plot in Figure 10.9.

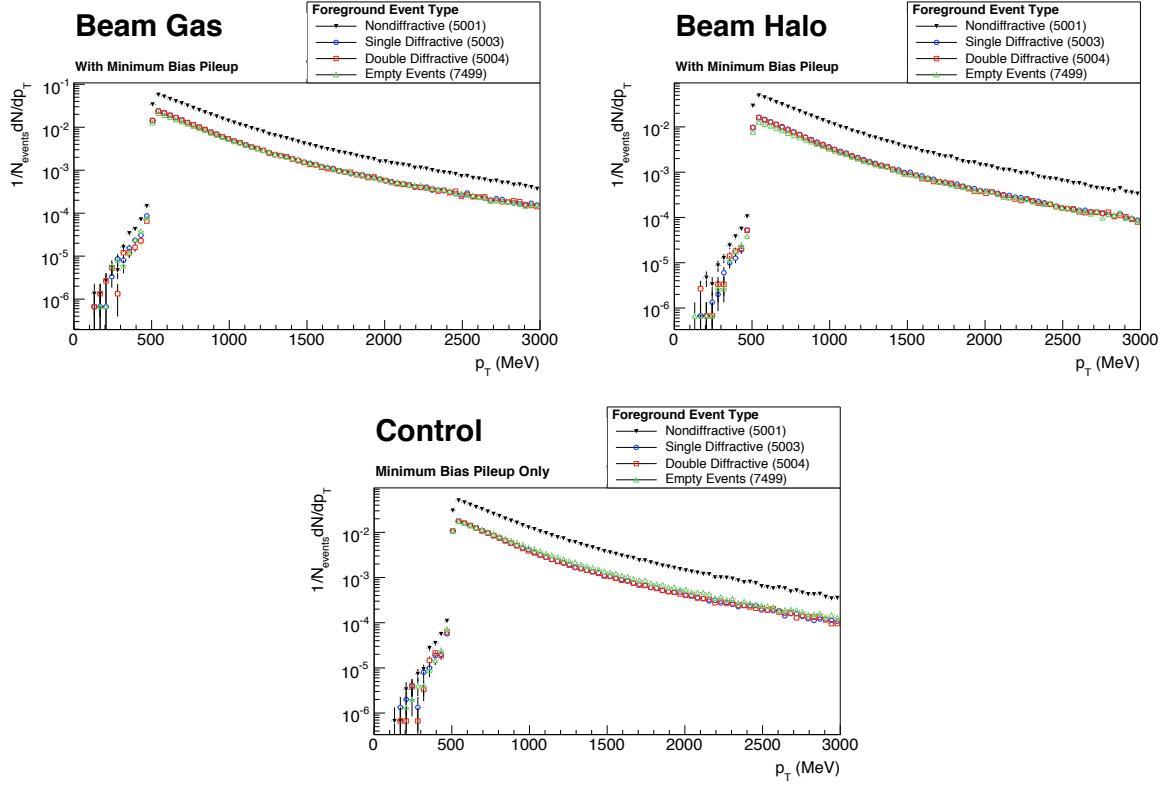


Figure 10.13: dN_{ch}/dp_T for normal pileup.

In Figure 10.13, the same trend holds - the curves are nearly identical (and again, the ND dataset is consistent with Figure 10.8), but the SD and DD datasets follow the same curve, only lower. The tracks from the minimum-bias pileup are composed mostly of the 5001 ND dataset, as discussed in [46], and will therefore overlay its characteristics on the SD, DD and empty dataset behaviors. The beam-gas events boost the dN_{ch}/dp_T densities by roughly a factor of two.

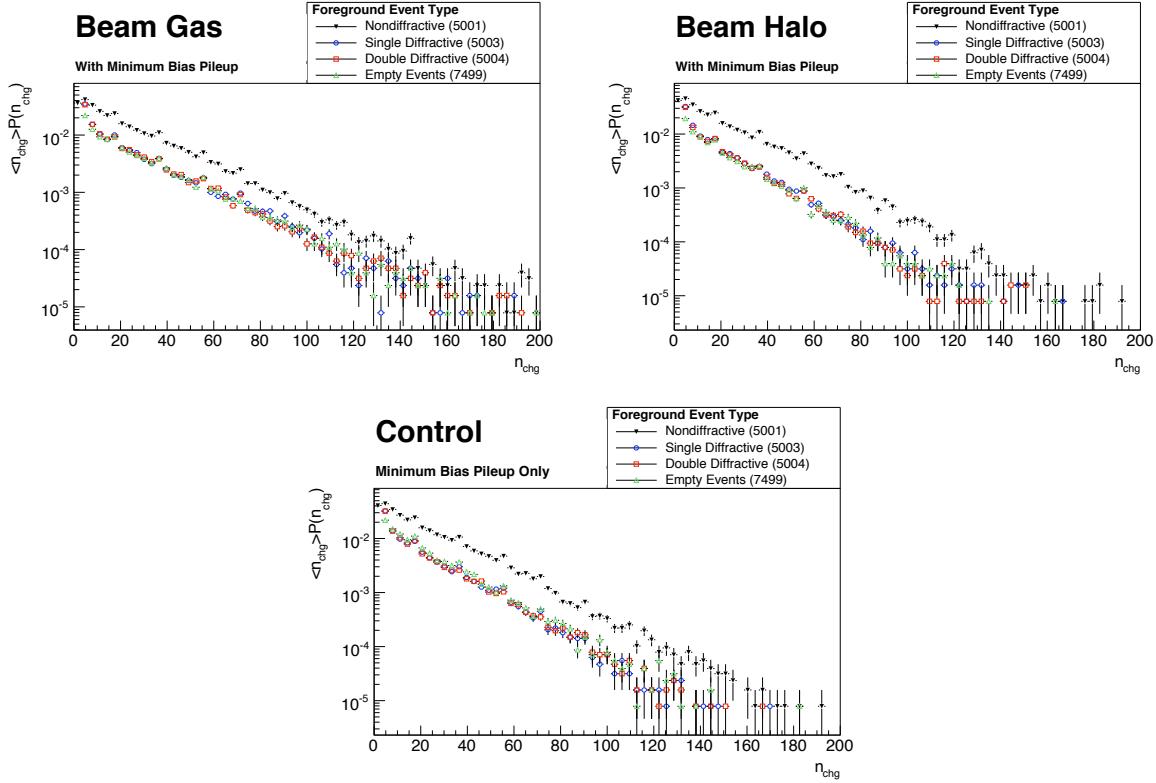


Figure 10.14: Charged particle multiplicity for normal pileup.

Charged particle multiplicities, as seen in Figure 10.14, are almost identical across the three plots, indicating that the majority of the charged particle tracks are coming from the primary event, in the case of the ND dataset, or from the ND components of the minimum-bias pileup, in the case of the SD, DD and empty datasets.

In sum – the beam gas, even with a substantially smaller inclusion rate than the beam halo in these datasets, had a much larger effect on the $dN_{\text{ch}}/d\eta$ and dN_{ch}/dp_T plots than did the halo. The minimum-bias pileup background will inform any charged-particle density plot into which it is admitted. Removal of these contaminated events will be explored in the next section.

10.2.3 Track Parameter Cuts – Backgrounds

The track quality cuts mentioned at the beginning of Section 10.2 are defined to clean out tracks whose reconstruction was problematic – whose track quality parameters indicated that

they were suspect, whose p_T is out of range, or whose constituent spacepoints are too few to give confidence that the track is valid.

The CSC note track quality cuts are listed in Table 10.1.

Quality cuts	No. of B-layer hits ≥ 1 No. of Silicon hits ≥ 5 $P(\chi^2) > 0.0$
Resolution cuts	$ \sigma_{d_0} < 1.6$ mm $ \sigma_{z_0} < 6.0$ mm $ \sigma_\phi < 0.03$ $ \sigma_\theta < 0.015$ $ \sigma_{q/p_T} < 0.0003$ GeV $^{-1}$
Track-to-vertex cut	$N_\sigma < 3$

Table 10.1: Track selection cuts used in the Minimum Bias CSC note. The resolutions $\sigma_{d_0}^2$, $\sigma_{z_0}^2$, σ_ϕ^2 , σ_θ^2 and σ_{q/p_T}^2 are the five diagonal elements in the track parameter covariance matrix. Table from [46]. These cuts are applied in the following tables.

When these cuts are applied to the beam halo and gas samples, none of the tracks pass the cuts. As seen in the Tables 10.2 and 10.3, track passes the θ cut of 0.015. This is the case for the entire dataset – not a single track passes. Other track parameters give similarly high cut rates. All the numbers in the tables that follow are independent, rather than cascading, cuts. The percentages represent the percent of tracks failing the cuts.

The implications of this are that the beam halo creates no tracks that can pass even a fairly loose quality selection cut – unsurprising – and that it will therefore be easily filtered from a given dataset. That is to say, any halo-contaminated events can be identified by high rates of bad tracks, and eliminated.

The beam gas, likewise, can be completely identified as a primary event, and completely removed using track quality cuts.

	ND – 5001	SD – 5003	DD – 5004
p_T	0.0%	0.0%	0.0%
η	1.1%	1.2%	1.4%
B-layer	12.6%	10.3%	10.1%
SCT	4.7%	4.6%	4.1%
χ^2	0.0%	0.0%	0.0%
d_0	14.7%	12.5%	12.5%
z_0	91.6%	92.2%	91.8%
θ	100.0%	100.0%	100.0%
ϕ	99.0%	98.9%	99.1%
Q over P	82.0%	91.2%	91.4%

Table 10.2: Beam Halo with no minimum-bias pileup, for the ND, SD, DD and empty foreground sets. Cuts as in Table 10.1. All the numbers in the tables that follow are independent, rather than cascading, cuts. The percentages represent the percent of tracks failing the cuts.

	H	O	C
p_T	0.0%	0.0%	0.0%
η	0.2%	0.1%	0.1%
χ^2	0.0%	0.0%	0.0%
B-layer	47.7%	47.8%	47.5%
SCT	52.5%	54.0%	53.3%
d_0	90.9%	91.0%	90.7%
z_0	97.6%	97.6%	97.6%
ϕ	99.2%	99.2%	99.3%
θ	100.0%	100.0%	100.0%
Q over P	77.0%	77.0%	76.9%

Table 10.3: Beam Gas with no minimum-bias pileup. Cuts as in Table 10.1. All the numbers in the tables that follow are independent, rather than cascading, cuts. The percentages represent the percent of tracks failing the cuts.

10.2.4 Track Parameter Cuts – Minimum-bias Pileup

The same set of cuts were also applied to the datasets that contained minimum-bias pileup. The most important of these is the “control” set in Table 10.4, which displays the same impenetrability to minimum-bias background contamination as the halo and gas datasets of the previous section. Tables 10.5 and 10.6 show the same cut effects for minimum-bias pileup overlaid with beam gas and beam halo), but they are only included for completeness – if all the inputs to the dataset are individually eliminated by the cuts, their sum will be as well.

	ND – 5001	SD – 5003	DD – 5004	Empty – 7499
p_T	0.0%	0.0%	0.0%	0.0%
η	1.1%	1.1%	1.1%	1.1%
B-layer	21.5%	38.2%	38.4%	42.1%
SCT	5.6%	5.1%	5.2%	5.4%
χ^2	0.0%	0.0%	0.0%	0.0%
d_0	20.4%	28.6%	29.2%	30.8%
z_0	91.9%	92.6%	92.1%	92.6%
ϕ	99.0%	99.1%	99.1%	99.0%
θ	100.0%	100.0%	100.0%	100.0%
Q over P	82.0%	82.6%	82.5%	80.7%

Table 10.4: Minimum-bias pileup (control), for the ND, SD, DD and empty foreground sets. Cuts as in Table 10.1. All the numbers in the tables that follow are independent, rather than cascading, cuts. The percentages represent the percent of tracks failing the cuts.

	ND – 5001	SD – 5003	DD – 5004	Empty – 7499
p_T	0.0%	0.0%	0.0%	0.0%
η	1.1%	1.1%	1.0%	1.1%
B-layer	21.1%	37.7%	38.4%	43.0%
d_0	20.0%	27.7%	28.3%	30.5%
SCT	5.4%	4.9%	4.9%	4.8%
χ^2	0.0%	0.0%	0.0%	0.0%
z_0	91.8%	92.1%	92.1%	92.3%
ϕ	99.0%	99.0%	99.1%	99.1%
θ	100.0%	100.0%	100.0%	100.0%
Q over P	81.9%	82.6%	82.7%	81.0%

Table 10.5: Beam halo with minimum-bias pileup, for the ND, SD, DD and empty foreground sets. Cuts as in Table 10.1. All the numbers in the tables that follow are independent, rather than cascading, cuts. The percentages represent the percent of tracks failing the cuts.

	ND – 5001	SD – 5003	DD – 5004	Empty – 7499
p_T	0.0%	0.0%	0.0%	0.0%
η	1.1%	0.9%	1.0%	1.0%
B-layer	23.9%	39.7%	39.9%	41.9%
SCT	7.4%	8.8%	8.5%	8.4%
χ^2	0.0%	0.0%	0.0%	0.0%
d_0	24.1%	34.2%	34.2%	34.8%
z_0	92.0%	92.0%	92.1%	92.3%
ϕ	99.0%	99.1%	99.0%	99.0%
θ	100.0%	100.0%	100.0%	100.0%
Q over P	81.9%	82.2%	82.0%	81.1%

Table 10.6: Beam Gas with minimum-bias pileup, for the ND, SD, DD and empty foreground sets. Cuts as in Table 10.1. All the numbers in the tables that follow are independent, rather than cascading, cuts. The percentages represent the percent of tracks failing the cuts.

Chapter 11

Conclusions and Outlook

Initial forecasts of the beam gas rate of 60 kHz [1] and beam halo compositions have been revisited and substantially revised.

The beam halo will be quite variable. Under ideal circumstances, halo rates may be less than one halo event per hundred events or less. These rates can vary widely, however, possibly peaking at a halo shower for each bunch crossing, though this correlates to small beam lifetimes. Halo diminution is something that can be expected as the LHC comes under progressively better control, but at no point is there a guarantee that halo will remain low-rate and manageable at any given time. The collimators have been taken into account, and it has been demonstrated that they are the dominant source of halo. The nature of the halo's composition has also been substantially revised from the 2004 dataset, showing that neutrons make up around half of the particles (from the collimators) that reach the ATLAS cavern, and that muons and hadrons have very different and distinct arrival trajectories in the cavern due to their particular properties and sources. Treatment of the beam halo without taking into account particle multiplicity (or the number of simultaneous particles produced by a given event) and correlations (especially in regards to triggers and large-surface detectors) has been found wanting.

It seems possible, though, by looking at the correlations in the 2007 halo dataset, to approximate the same sorts of particle distributions (large-multiplicities of low energy, or single high-energy halo particles – matching the energy profile of the ad-hoc assembled events to that of the 2007 sample). It would be interesting, in the future, to integrate the admittedly imperfect 2004 halo simulation into the pileup as well, to investigate any differences.

The halo has been demonstrated to have a minor effect on the MBTS system, even at high rates, and no practical effect on the ATLAS tracking other than a suspected small track-spoiling rate.

Beam-gas events have been calculated to be a much smaller rate than previously suspected, but can substantially disrupt proper tracking in an event in spite of their out-of-time nature. Their effects reside mostly in the $1.0 < |\eta| < 2.7$ range, and gas target species matters to some small degree regarding the number of tracks produced, to a small degree. Future investigations should include some of the more “exotic” beam-gas possibilities as well, to know that their behavior tracks with the lower-Z elements examined in this study.

The minimum-bias triggering should remain unaffected by beam gas to any serious degree at the beginning of ATLAS and LHC operation, and though halo is unpredictable, it should not have a serious impact on minimum-bias measurements. The same is not necessarily true for muon-intensive studies, and the next logical step in the investigation of both beam gas and beam halo is to look at some of these analyses with appropriate beam gas and beam halo pileup included. The muon component of the beam halo is not shieldable, and will therefore not be attenuated to the degree to which the hadron flux is.

There are several possible improvements to the procedure used in this study. One of the most important will be to create customized pileup samples with the correct event proportions contained within one overarching file. Perhaps divisions into smaller sets (minimum bias, correctly proportioned; beam gas and beam halo, correctly proportioned) could also be done to allow variable scales as well. The problems that the pileup mechanism has been suffering will be taken care of by the next major release of Athena, at which point pileup event recycling can see use again.

Other possibilities for further research include redoing many of these investigations with higher statistics (perhaps with runs in the ATLAS Grid production system), and at different minimum-bias pileup rates, to allow some of the effects to be deconvoluted.

ATLAS and the LHC, upon startup, will exhibit behaviors that make measurement of precision quantities and detection of rare processes very difficult. These behaviors, however, contain

the keys for understanding the workings of the Standard Model at the 14 TeV energy scale, and must be carefully investigated and analyzed. Some of these unknowns arise from uncertainties in our understanding of the physics, while others just come from the difficulty of executing such an ambitious experiment in an imperfect world. Coming to grips with these challenges is satisfying, at times thrilling, but will be an ongoing endeavor for a long time to come.

LIST OF REFERENCES

- [1] M Boonekamp, F Gianotti, R A McPherson, Marzio Nessi, and P Nevksi. Cosmic ray, beam-halo and beam-gas rate studies for ATLAS commissioning. Technical Report ATL-GEN-2004-001, CERN, Geneva, Feb 2004.
- [2] Ralf Spiwoks et al. The ATLAS Level-1 Central Trigger Processor (CTP). *ATL-DAQ-CONF-2005-030; CERN-ATL-DAQ-CONF-2005-030*, 2005.
- [3] F. Abe et al. *Phys. Rev. D*, 56:3811, 1997.
- [4] F. Abe et al. *Phys. Rev. Lett.*, 79:584, 1997.
- [5] F. Abe et al. *Phys. Rev. D*, 47:4857, 1993.
- [6] S. G. Matinyan and W. D. Walker. *Phys. Rev. D*, 59:034022, 1999.
- [7] T. et al. Alexopoulos. *Phys. Lett. B*, 435:453, 1998.
- [8] F. et al. Abe. Pseudorapidity distributions of charged particles produced in $\bar{p}p$ interactions at $\sqrt{s} = 630$ GeV and 1800 GeV. *Phys. Rev.*, D41:2330, 1990.
- [9] R. E. et al. Ansorge. Charged particle correlations in $\bar{p}p$ collisions at c.m. energies of 200 GeV, 546 GeV and 900 GeV. *Z. Phys.*, C37:191–213, 1988.
- [10] R. E. et al. Ansorge. Diffraction dissociation at the CERN pulsed collider at cm energies of 900 GeV and 200 GeV. *Z. Phys.*, C33:175, 1986.
- [11] R. E. et al. Ansorge. *Z. Phys. C*, 43:357, 1989.
- [12] G. J. Alner. *Phys. Rep.*, 154:247, 1987.
- [13] A. Breakstone. *Phys. Rev. D*, 30:528, 1984.
- [14] Valeria Tano. The underlying event in jet and minimum bias events at the Tevatron. arXiv:hep-ph/0011141v1, November 2001.
- [15] ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *Journal of Instrumentation*, forthcoming, 2008.

- [16] Vernon D. Barger and R. J. N. Phillips. *Collider physics*. REDWOOD CITY, USA: ADDISON-WESLEY (1987) 592 P. (FRONTIERS IN PHYSICS, 71).
- [17] G. Altarelli and G. Parisi. *Nucl. Phys. B*, 126(298), 1977.
- [18] Yu. L. Dokshitzer. *Sov.Phys. JETP*, 46:641, 1977.
- [19] S. F. Novaes. Standard Model : An introduction. *Particles and Fields*, arXiv:hep-ph/0001238v1, January 2000.
- [20] J.M. Campbell, J.W. Huston, and W.J. Stirling. Hard interactions of quarks and gluons: a primer for LHC physics. *Rept.Prog.Phys.*, 70:89, 2007.
- [21] M. Diele et al. The First Year at the LHC: Diffractive Physics. *Czech . J. Physics*, 55:B757–B767, July 2005.
- [22] J. R. Cudell et al. Benchmarks for the forward observables at RHIC, the Tevatron Run II and the LHC. *Physical Review Letters*, 89, 2002.
- [23] M. Diele et al. Diffraction and Total Cross-Section at the Tevatron and the LHC. *Hadron Collider Physics Symposium 2005, Les Diablerets, Switzerland*, July 2005.
- [24] The LHC Study Group. The Large Hadron Collider. Conceptual Design Report. Technical Report CERN/AC/95-05, CERN, 1995.
- [25] Stephen Myers. *The LEP Collider, from design to approval and commissioning*. John Adams' Memorial Lecture. CERN, Geneva, 1991. Delivered at CERN, 26 Nov 1990.
- [26] Austin Ball, Michel Della Negra, L Foa, and Achille Petrilli. *CMS Physics Technical Design Report*. Technical Design Report CMS. CERN, Geneva, 2006. revised version submitted on 2006-09-22 17:44:47.
- [27] Search for the Standard Model Higgs boson at LEP. Technical Report DELPHI-2001-113-CONF-536. CERN-DELPHI-2001-113-CONF-536. L3-Note-2699. OPAL-Physics-Note-PN-479. LHWG-Note-2001-03. CERN-ALEPH-CONF-2001-046. CERN-ALEPH-2001-066, CERN, Geneva, Jul 2001.
- [28] Jonathan Richard Ellis. Supersymmetry for Alp Hikers. (hep-ph/0203114. CERN-TH-2002-052):45 p, Mar 2002.
- [29] G Riddone. Cooling down a whole machine. 2004.
- [30] L Evans. Current status of LHC project. CERN EDMS 849946, June 2007.
- [31] Mike Lamont. Luminosity optimisation. *Chamonix XII*, (CERN AB 2003 008), 2003.

- [32] Paul Collier. Proton bunch pattern in the LHC for Totem operation and for initial collision commissioning. Technical Report <https://edms.cern.ch/document/406238/1>, CERN, 2003.
- [33] A. Presland, D. Ramos, B. Goddard, J-M. Jimenez, and R. Veness. A large diameter entrance window for the LHC beam dump line. Knoxville, Tennessee, 2005. Proceedings of 2005 Particle Accelerator Conference, IEEE Computer Society.
- [34] R. Assmann. Beam halo and beam losses in IR1 and IR5, 2007. <http://indico.cern.ch/conferenceDisplay.py?confId=17450>.
- [35] S Baranov, M Bosman, I Dawson, V Hedberg, A Nisati, and M Shupe. Estimation of radiation background, impact on detectors, activation and shielding optimization in ATLAS. Technical Report ATL-GEN-2005-001. ATL-COM-GEN-2005-001. CERN-ATL-GEN-2005-001, CERN, Geneva, Jan 2005.
- [36] Particle Data Group. Review of Particle Physics. *Physics Review D*, 2004.
- [37] Markus Cristinziani. The ATLAS pixel detector, September 2006.
- [38] R. Beccherle et al. MCC: the Module Controller Chip for the ATLAS pixel detector. *ATL-INDET-2002-002*, 2002.
- [39] Reiner Klingenberg. The ATLAS pixel detector. *Nucl. Instrum. Meth.*, A579:664–668, 2007.
- [40] C. Troncon et al. A measurement of Lorentz angle and spatial resolution of radiation hard silicon pixel sensors. *Nuc. Inst. Meth. A*, 481:204–221, May 2002.
- [41] F. Heinemann. Track based alignment of the ATLAS silicon detectors with the robust alignment algorithm. *COM-INDET-2007-014*, 2007.
- [42] A. Ahmad et al. The silicon microstrip sensors of the ATLAS semiconductor tracker. *Nucl. Instrum. Meth.*, A578:98–118, 2007.
- [43] The ATLAS Collaboration. Inner Detector Technical Design Report. Technical Report CERN/LHCC/97-17, CERN, 1997.
- [44] The ATLAS Collaboration. Liquid Argon Calorimeter Technical Design Report. Technical Report CERN/LHCC/96-41, CERN, 1996.
- [45] The ATLAS Collaboration. Detector and Physics Performance Technical Design Report. Technical Report CERN/LHCC/99-14, CERN, 1999.
- [46] ATLAS Collaboration. A study of minimum bias events at $\sqrt{s}=14$ TeV. ATL-PHYS-PUB-2007-XXX [NOT YET PUBLISHED], December 2007.

- [47] Eric Feng and James Pilcher. Triggering ATLAS with minimum bias trigger scintillators. *ATLAS Internal Note*, (TILECAL-NOTE-07-XXX), March 2007.
- [48] The ATLAS Collaboration. Tile Calorimeter Technical Design Report. Technical Report CERN/LHCC/96-42, CERN, 1996.
- [49] The ATLAS Collaboration. Muon Spectrometer Technical Design Report. Technical Report CERN/LHCC/97-22, CERN, 1997.
- [50] The ATLAS Collaboration. ATLAS DAQ, EF, LVL2 and DCS Technical Progress Report. Technical Report CERN/LHCC/98-16, CERN, 1998.
- [51] The ATLAS Level-1 Trigger Group. Level-1 Technical Design Report. Technical Report ATLAS TDR-12, CERN, 1998.
- [52] Ralf Spiwoks et al. The ATLAS Level-1 Central Trigger Processor (CTP). *ATL-DAQ-CONF-2005-030; CERN-ATL-DAQ-CONF-2005-030*, 2005.
- [53] Regina Kwee. Diploma thesis - studies for the development of a minimum bias trigger at low luminosities in the ATLAS-experiment at the Large Hadron Collider. July 2007.
- [54] Giuseppe Lo Presti, Olof Barring, Alasdair Earl, Rosa Maria Garcia Rioja, Sebastien Ponce, Giulia Taurelli, Dennis Waldron, and Miguel Coelho Dos Santos. CASTOR: A distributed storage resource facility for high performance data processing at CERN. In *MSST '07: Proceedings of the 24th IEEE Conference on Mass Storage Systems and Technologies (MSST 2007)*, pages 275–280, Washington, DC, USA, 2007. IEEE Computer Society.
- [55] The ATLAS Collaboration. Computing Technical Design Report. Technical Report CERN/LHCC/2005-022, CERN, 2005.
- [56] P Mato. GAUDI - architecture design document. Technical Report LHCb-98-064, CERN, Geneva, Nov 1998.
- [57] T. Sjostrand, S. Mrenna, and P. Skands. Pythia 6.4 physics and manual. *JHEP*, 05:026, 2006.
- [58] M. Gyulassy and X. N. Wang. Hijing 1.0: A monte carlo program for parton and particle production in high-energy hadronic and nuclear collisions. *Comput. Phys. Commun.*, 83:307, 1994.
- [59] I. Azhgirey, I. Baishev, K. M. Potter, and V. Taranov. Methodical study of the machine induced background formation in the IR8 of LHC. *LHC Project Note 258*, June 2001.
- [60] I. Azhgirey, I. Baishev, K. M. Potter, and V. Taranov. Cascade simulations for the machine induced background study in the IR1 of the LHC. *LHC Project Note 324*, September 2003.

- [61] P. A. et al. Aarnio. Fluka86: User's guide. *CERN TIS-RP/168*, 1986.
- [62] et al. Fasso, A. Fluka: performances and applications in the intermediate energy range. *Meeting on Shielding Aspects of Accelerators, Targets and Irradiation Facilities, Arlington, TX*, 1994.
- [63] S. et al. Agostinelli. Geant4: A simulation toolkit. *Nucl. Instrum. Meth.*, A506:250–303, 2003.
- [64] E. Richter-Was, D. Froidevaux, and L. Poggioli. ATLFAST 2.0, A fast simulation package for ATLAS. ATL-PHYS-98-131, 1998.
- [65] W Dabrowski, Francis Anghinolfi, C M Buttar, V Cindro, A G Clarks, I Dawson, D Dorfan, T Dubbs, N Falconer, M French, A Greenall, A A Grillo, R Happer, Pierre Jarron, J Kaplon, J Kudlaty, G Kramberger, C Lacasta, D La Marra, Daniela Macina, I Mandic, M Mikuz, G Meddeler, O Milgrome, H Niggli, P W Phillips, S Roe, A Smith, H Spieler, E Spencer, R Szczygiel, Peter Weilhammer, M Wolter, and A Zsenei. Design and performance of the ABCD chip for the binary readout of silicon strip detectors in the ATLAS semiconductor tracker. *IEEE Trans. Nucl. Sci.*, 47(6, pt.1):1843–50, 2000.
- [66] Trevor Vickey. A Read-Out Driver for silicon detectors in ATLAS. *12th LECC Workshop Proceedings*, September 2006.
- [67] I. Dawson and K. K. Prokofiev. Luminosity measurement in the interaction region of the LHC using proton proton collision multiplicities. *Nucl. Instr. and Meth.*, A578:480–484, 2007.
- [68] A. Rossi. Residual gas density estimations in the LHC insertion regions IR1 and IR5 and the experimental regions of ATLAS and CMS for different beam operations. *LHC Project Report 783*, 2004.
- [69] M. Ferro-Luzzi. Proposal for an absolute luminosity determination in colliding beam experiments using vertex detection of beam-gas interaction. *Nuclear Instruments and Methods in Physics Research, Section A*, 553:388–399, November 2005.
- [70] Olivier Sim Bruning, Paul Collier, P Lebrun, Stephen Myers, Ranko Ostojic, John Poole, and Paul Proudlock. *LHC Design Report*. CERN, 2004.
- [71] 2007.
- [72] ATLAS Collaboration. Atlas trigger performance status report. Technical Report CERN-LHCC-98-015, CERN, Jun 1998.
- [73] B Betz, M Bleicher, U Harbach, T Humanic, B Koch, and H Stocker. Mini black holes at the LHC: Discovery through di-jet suppression, mono-jet emission and a supersonic boom in the quark-gluon plasma in ALICE, ATLAS and CMS. *J. Phys. G*, 32(hep-ph/0606193):S429–S438. 16 p, Jun 2006.

- [74] P. F. et al. Åkesson. Atlas tracking event data model, available at <http://cdsweb.cern.ch/record/973401>. *ATL-SOFT-PUB-2006-004*, 2006.
- [75] A. Abdesselam et al. The ATLAS semiconductor tracker end-cap module. *Nucl. Instrum. Meth.*, A575:353–389, 2007.
- [76] S Mehlhase and T C Petersen. A probability based approach to PID in the TRT detector of ATLAS. Technical Report ATL-COM-INDET-2006-017, CERN, Geneva, Oct 2006.
- [77] I. Wingerter-Seez. Study of energy reconstruction using optimal filtering with the LAr electromagnetic calorimeter. *ATLAS Internal Note LARG-NO-19*, 1995.
- [78] M. Lieng. Background simulations for the LHCb beam condition monitor, 2007. <http://indico.cern.ch/conferenceDisplay.py?confId=17450>.
- [79] P. Schacht et al. Hadronic calibration of the ATLAS liquid argon end-cap calorimeter in the pseudorapidity region $1.6 < |\eta| < 1.8$ in beam tests. *Nuc. Inst. Meth. A*, 531:481–514, May 2004.
- [80] F. Akesson, T. Atkinson, M.J. Costa, M. Elsing, S. Fleischmann, A. Gaponenko, W. Liebig, E. Moyse, A. Salzburger, and M. Siebel. Atlas tracking event data model. Technical report, CERN, 2006.
- [81] Christian Wolfgang Fabjan, Lodovico Riccati, Karel Safarik, and Hans de Groot. *ALICE physics performanceTechnical Design Report*. Technical Design Report ALICE. CERN, Geneva, 2005. Revised version submitted on 2006-05-29 15:15:40.
- [82] V.N. Gribov and L.N. Lipatov. *Sov. J. Nucl. Phys.*, 15(438), 1972.
- [83] E. Feng and J. E. Pilcher. Triggering ATLAS with minimum bias trigger scintillators. *ATL-COM-TILECAL-2007-013*, 2007.
- [84] Distributed analysis using GANGA. web-address: <https://twiki.cern.ch/twiki/bin/view/Atlas/DistributedAnalysisUsingGanga>.
- [85] I. Azhgirey, I. Kurochkin, and V. Talanov. Development of MARS code package for radiation aspects of electronuclear installation design. *Proc. 15th Conf. on Charged particles Accelerators, Protvino*, page 74, 1996.
- [86] Phojet manual (program version 1.05c, june 96). <http://physik.uni-leipzig.de/~eng/phojet.html>.
- [87] S. Weinberg. *Phys. Rev. Lett.*, 19:1264, 1967.
- [88] J. Lu and D. M. Gingrich. Calorimeter calibration using ep of hadrons from τ decay in $z \rightarrow \tau\tau$ events using CSC data. March 2007.

- [89] Oliver Kortner, Hubert Kroha, Jens Schmaler, and Chrysostomos Valderanis. Alignment of the ATLAS muon spectrometer with tracks and muon identification at high background rates. *Nuc. Inst. Meth. A*, August 2007.
- [90] Arthur Moraes, Craig Buttar, and Ian Dawson. Prediction for minimum bias and the underlying event at LHC energies. *Europ. Phys. Journal C*, 50(2), April 2007.
- [91] I. Azhgirey, I. Kurochkin, A. Sannikov, and E. Savitskaya. *Nuclear Instruments and Methods in Physics Research, Section A*, A408, 1998.
- [92] The ATLAS Collaboration. Detector and Physics Performance Technical Design Report. Technical Report CERN/LHCC/99-15, CERN, 1999.
- [93] B. Aubert et al. Performance of the ATLAS electromagnetic calorimeter barrel module 0. *Nuc. Inst. Meth. A*, A500:202–231, 2003.
- [94] J. Lu and D. M. Gingrich. Update of calorimeter calibration using ep of hadrons from τ decay in $z \rightarrow \tau\tau$ events using 12.0.6 CSC data. August 2007.
- [95] R.M. Barnett et al. Review of particle properties. *Phys. Rev. D54 1*, 1996.
- [96] S. Behrens and A.C. Melissinos. *Univ. of Rochester Preprint*, 1981.
- [97] N. Benekos et al. Muon detector description as built and its simulation for the ATLAS experiment. *Nuc. Inst. Meth. A*, 572, 2007.
- [98] L3 ALEPH, DELPHI, OPAL Collaborations, and the LEP working group for Higgs boson searches. Search for the Standard Model Higgs Boson at LEP. *ALEPH 2001-066 CONF 2001-046*, 2001.
- [99] R Antunes-Nobrega et al. *LHCb reoptimized detector design and performance Technical Design Report*. Technical Design Report LHCb. CERN, Geneva, 2003.
- [100] The ATLAS Electromagnetic Liquid Argon Calorimeter Group. Performance of the ATLAS electromagnetic calorimeter barrel module 0. *CERN-EP/2002-087*, 2002.
- [101] E. Bravin. Bringing the first LHC beams into collision at all 4 IPs. *Chamonix XV*, 2007.
- [102] R. Engel. *Z. Phys. C*, 66:203, 1995.
- [103] F. James. *MINUIT: Function Minimization and Error Analysis*. CERN Program Library Long Writeup D506, 1994.
- [104] D.H. Wilkinson. Ionization Energy Loss by Charged Particles Part I. The Landau Distribution. *Nucl. Inst. Meth.*, A383:513, 1996.
- [105] U. Blumenschein. Installation, commissioning and operation of the tile hadron calorimeter of ATLAS. *Nuc. Inst. Meth. A*, 572:24–25, November 2007.

- [106] B Goddard. Beam dynamics and intensity challenges - beam dump, November 2004.
- [107] M. Aharrouche. The ATLAS liquid argon calorimeter: Construction, integration, commissioning and combined test beam results. *Nuc. Inst. Meth. A*, August 2007.
- [108] J. Lu and D. M. Gingrich. Investigation of hadronic calorimeter calibrations using ep of single charged hadrons from τ decay in $z \rightarrow \tau\tau$ events with ATLAS using Rome Monte Carlo production. April 2006.
- [109] F. Djama. Using $z^0 \rightarrow e^+e^-$ for Electromagnetic Calorimeter Calibration. *ATL-LARG-2004-008*, 2004.
- [110] Alfred M. Asner et al. Large Hadron Collider in the LEP tunnel : A Feasibility Study of Possible Options. *CERN-DIR-TECH-84-01*, 1984.
- [111] Atlas computing technical design. *ATL-TDR-017* (2005).
- [112] S. Schuh. Precision aspects of ATLAS muon chamber design and construction. *Nuc. Inst. Meth. A*, 572:135–138, October 2007.
- [113] W. Seligman. ParticleGenerator: An Athena Monte-Carlo event generator for simple particles, August 2005.
- [114] M. Nessi. Minimum bias trigger scintillator counters (MBTS) for early ATLAS running. *ATLAS Overview Week, Freiburg, Germany*, 2004.
- [115] Kevin Einsweiler. The ATLAS pixel detector, April 2005.
- [116] J. Dubbert. First experience with the ATLAS muon spectrometer. *Nuc. Inst. Meth. A*, August 2007.
- [117] Ilias Efthymiopoulos. The readout driver (ROD) for the ATLAS liquid argon calorimeters. *Nuc. Inst. Meth. A*, 461:481–482, 2001.
- [118] The PanDA Production and Distributed Analysis System. web-address: <https://twiki.cern.ch/twiki/bin/view/Atlas/PanDA>.
- [119] The ATLAS Collaboration. Magnet System Technical Design Report. Technical Report CERN/LHCC/97-18, CERN, 1997.
- [120] A. Moraes, C. Buttar, and I. Dawson. *Eur. Phys. J. C*, 50:435, 2007.
- [121] A. Artikov, Chokheli D., Huston J., Miller B., and Nessi M. Minimum bias scintillator counter geometry. *Technical Report AT-GE-ES-0001*, 2004.
- [122] F. Hubaut et al. Test beam Measurement of the Crosstalk in the EM Barrel Module 0. *ATL-LARG-2000-007*, 2000.

- [123] R. Sacco et al. Position resolution and particle identification with the ATLAS EM calorimeter. *ATLAS Note submitted to Elsevier Science*, 2005.
- [124] G. Sauvage et al. Construction, assembly and tests of the ATLAS electromagnetic barrel calorimeter. *Nuc. Inst. Meth. A*, 558:388–418, November 2006.
- [125] R. Wigmans. *Calorimetry: Energy Measurements in Particle Physics*. Oxford University Press, Clarendon, 2000.
- [126] J Weng, C Saout, A de Roeck, G Quast, and M Spiropulu. Search for add direct graviton emission in photon plus missing transverse energy final state at CMS. Technical Report CMS-NOTE-2006-129. CERN-CMS-NOTE-2006-129, CERN, Geneva, Jun 2006.
- [127] The ATLAS Collaboration. Inner Detector Technical Design Report. Technical Report CERN/LHCC/97-16, CERN, 1997.
- [128] Vasiliki Mitsou. The ATLAS Transition Radiation Tracker. *arXiv:hep-ex/0311058v1*, November 2003.
- [129] G. Schlager et al. The status and performance of the ATLAS hadronic tile calorimeter. *Nuc. Inst. Meth. A*, August 2007.

Appendix A: Minimum bias reference datasets

The datasets in Table A.1 were used as a backdrop on which to overlay the beam gas and beam halo samples.

Dataset Name
misal1_csc11.005001.pythia_minbias.simul.HITS.v12003104_tid004278
users.williambell.ganga.datafiles.mc12.005003.pythia_sdiff.simul.HITS.v12000318
users.williambell.ganga.datafiles.mc12.005004.pythia_ddiff.simul.HITS.v12000318
users.williambell.ganga.datafiles.mc12.007499.singlepart_empty.simul.HITS.v12000318

Table A.1: Datasets used for minimum bias studies

Appendix B: Details of beam gas and halo run

Since the job configuration code itself makes an enormous difference in the final results derived from a simulation program, I will include in this appendix a summary of the various kinds of run made to get the gas and halo samples simulated, digitized and reconstructed.

B.1 Beam gas generation

Beam gas was simulated in Athena release 13.0.10, with a version of HIJING that incorporated my modifications in the official code, allowing vertex shifting and randomization, and allowing the side of the cavern to be chosen for the duration of the run. In the future, it may be worthwhile for me to modify things further, allowing randomization of the incoming proton's sign on the z -axis.

jobOptions for this run are as follows:

```
#####
# Job options file
#=====
#-----
# General Application Configuration options
#-----
#theApp.setup( MONTECARLO )
include( "AthenaCommon/Atlas_Gen.UnixStandardJob.py" )
include( "PartPropSvc/PartPropSvc.py" )
include( "AthenaPoolCnvSvc/WriteAthenaPool_jobOptions.py" )
print "Generating Hydrogen Beam Gas, 7 TeV, Standard p"
#-----
# Private Application Configuration options
#-----
theApp.Dlls  +=  [ "Hijing_i" ]
```

```

theApp.Dlls    += [ "TruthExamples" ]
theApp.Dlls    += [ "GeneratorObjectsAthenaPoolPoolCnv" ]
theApp.TopAlg =  [ "Hijing", "EvtShift" ]
theApp.ExtSvc += [ "AtRndmGenSvc" ]
AthenaPoolCnvSvc = Service( "AthenaPoolCnvSvc" )
AthenaPoolCnvSvc.CommitInterval = 100
MessageSvc = Service( "MessageSvc" )
MessageSvc.OutputLevel      = INFO
MessageSvc.debugLimit       = 100000
MessageSvc.infoLimit        = 100000
theApp.EvtMax = 5000
Hijing = Algorithm( "Hijing" )
Hijing.randomizeVertices = TRUE
Hijing.wide = FALSE
EvtShift = Algorithm( "EvtShift" )
EvtShift.mirrorP = FALSE # True changes to beam 2
Stream1 = Algorithm( "Stream1" )
Stream1.ItemList += [ "EventInfo#*", "McEventCollection#*" ]
AtRndmGenSvc = Service( "AtRndmGenSvc" )

#####
# Algorithms Private Options
#####
#Random numbers must be RANDOM! If you want, you can do the following:
# import random
# AtRndmGenSvc.Seeds = ["HIJING \%\d \%\d" \% (random.randint(),
#                         random.randint()),
#                         "HIJING_INIT \%\d \%\d" \% (random.randint(),

```

```

        random.randint())
AtRndmGenSvc.Seeds = ["HIJING 344878 885060",
                      "HIJING_INIT 303544 369585"]
#Sets Hijing to p-A mode, with appropriate Z and A for oxygen
Hijing.Initialize = ["efrm 7000", "frame LAB",
                     "targ A", "iat 16", "izt 8"]
Stream1.OutputFile = "hijing.0.evgen._00001.pool.root"
# End jobOptions

```

Notable settings are the randomization of the vertex location in z (and in r , if desired), side shifting so that the gas interaction can come from one beam or the other, and settings for the specific contaminant gas species.

B.2 Beam gas simulation

Beam gas requires a somewhat modified G4 version - the modifications are part of the Athena 13.X.0 series, but were not available as a built0in option in Athena 13.0.30, where all the simulation was done. The following jobOptions take the result of the last script and simulate it, taking into account the vertex (and therefore time) offset.

```

#=====
# Job options file for Geant4 Simulations
__version__="$Revision: 1.39 $"
#=====

#--- Detector flags -----
from AthenaCommon.DetFlags import DetFlags
# - Select detectors
DetFlags.ID_setOn()
DetFlags.Calo_setOn()

```

```

DetFlags.Muon_setOn()

DetFlags.simulate.Truth_setOn()

#--- AthenaCommon flags -----
from AthenaCommon.AthenaCommonFlags import athenaCommonFlags
athenaCommonFlags.PoolHitsOutput='hijing.0.simul._00001.pool.root'
fname=['/xdata04/stradling/Hijing/run/Elements/
hijing.0.evgen/hijing.0.evgen._00001.pool.root']
athenaCommonFlags.PoolEvgenInput=fname
athenaCommonFlags.EvtMax=200
athenaCommonFlags.SkipEvents=0

#--- Simulation flags -----
from G4AtlasApps.SimFlags import SimFlags
# Look into SimFlags.SimLayout for other possible values
SimFlags.SimLayout='ATLAS-CSC-01-02-00' # specific value
SimFlags.SimLayout.set_On() # uses the default value
SimFlags.CalibrationRun.set_Off()

# sets the EtaPhi, VertexSpread and VertexRange checks on
#SimFlags.EventFilter.set_On()

# sets LArParametrization
#SimFlags.LArParameterization=True

# No magnetic field
#SimFlags.MagneticField.set_Off()

# - reads events already generated
SimFlags.KinematicsMode='ReadGeneratedEvents'

# (the input file name is athenaCommonFlags.PoolEvgenInput]

#--- Output printout level -----
#output threshold (2=DEBUG, 3=INFO, 4=WARNING, 5=ERROR, 6=FATAL)
#you can override this for individual modules if necessary

```

```

MessageSvc = Service( "MessageSvc" )
MessageSvc.OutputLevel = 3
from G4AtlasApps import AtlasG4Eng

AtlasG4Eng.G4Eng.load_Dict('VertexTimeOffsetDict')
vFilter=AtlasG4Eng.G4Eng.gbl.VertexTimeOffset('VertexTimeOffset')
vFilter.switchOn()

#=====
# Job configuration
# ***>> Do not add flags or simulation options below this line
#=====

from G4AtlasApps import SimKernel
#--- End jobOptions.G4Atlas_Sim.py file -----

```

Here, the nonstandard behaviors center around the `VertexTimeOffset` under `AtlasG4Eng.G4Eng` package, which takes care of the mechanics of calculating the time offset and applying it to the `GenVertex` and all its children in the simulation.

B.3 Beam halo generation

The generation for the beam halo physics was done by Vadim Talanov of the AT Division at CERN, using the most recent optics and proton-collimator loss maps for the LHC. The job output came in the form of a series of columns of raw ASCII text. A sample is shown below, with the following column definitions: 1. Parent particle barcode, 2. FLUKA particle ID, 3. Number of generations from parent, 4. Energy (GeV), 5-7: x, y, z , 8-10: d_x, d_y, d_z - the three direction cosines of the particles. The columns are split in this context to avoid spilling over the margin.

```
2 11      4  0.173153586014E+01 -0.262777151859E+03 -0.789360888433E+02
```

```

0.000000000000E+00 -0.241746993430E-01
0.748533307296E-01 -0.996901480985E+00
2 10      5  0.592921454410E+00  0.151791230023E+03  0.821341368467E+03
0.000000000000E+00  0.406395073357E-01
0.488459615878E-01 -0.997979209443E+00
4 10      5  0.462541897408E+01 -0.826964400291E+03 -0.603478976291E+03
0.000000000000E+00 -0.994155802002E-01
-0.367479067876E-01 -0.994367202677E+00
6 14      7  0.374591078777E+01 -0.371218189483E+03 -0.743243249620E+03
0.000000000000E+00 -0.115769591734E+00
-0.467085271352E-01 -0.992177259930E+00

```

Some comments on the structure are appropriate at this point. It is clear that there are strong variations in multiplicity for the individual showers. Missing parent barcodes represent events that showered, but did not reach the 23 m strike plane in the ATLAS cavern (represented by the 0s in the z column). The FLUKA barcodes were converted to PDG IDs, the directions, species and energies combined (looking up the correct masses) to give a correct p vector, and the input locations correlated to GenVertex objects with the appropriate coordinates.

The conversion from this prepared state to HepMC happened in a modified version of the `CosmicsGenerator`. The following code was inserted into the `callGenerator` method by Giulio Usai:

```

while(!m_ffile.eof()&& !eventdone)
{
  int dummy;
  int id(0),parent(0),generation(0);
  double weight(0);
  double kin_energy(0);
  double d_x(0),d_y(0),d_z(0);

```

```
double v_x(0),v_y(0),v_z(0);

Hep3Vector vert,mom;
HepLorentzVector Vertex,Momentum;

m_ffile>>parent;

if (first) { //initialize on first line
    hold_parent = parent;
    first = false;

    m_ffile >> id
    >> generation >> kin_energy
    >> v_x
    >> v_y
    >> v_z
    >> d_x
    >> d_y
    >> d_z;

    std::cout <<"pippo1 \t"
    << parent << '\t'
    << id << '\t'
    << generation << '\t'
    << kin_energy<< '\t'
    << v_x << '\t'
    << v_y << '\t'
    << d_x << '\t'
    << d_y
    << std::endl;
}
```

```
else { // all other lines

    if (parent == hold_parent) { // same event
        m_ffile >> id
        >> generation
        >> kin_energy
        >> v_x
        >> v_y
        >> v_z
        >> d_x
        >> d_y
        >> d_z;

        std::cout << "pippo2 \t"
        << parent << '\t'
        << id << '\t'
        << kin_energy << '\t'
        << v_x << '\t'
        << v_y << '\t'
        << d_x << '\t'
        << d_y
        << std::endl;
    }

    else { // new event
        eventdone=true;
        hold_parent=parent;
        int ndig = (int)log10 (parent) + 1;
```

```

    std::cout << "parent= " << parent << "char=" << ndig << std::endl;
    for (int i=0; i<ndig; ++i) {
        m_ffile.unget();
    }
    break;
}

v_z= 2300;
if( m_swap ){
    v_z = -v_z;
    d_z = -d_z;
}

// convert to mm units
vert.setX(v_x * 10);
vert.setY(v_y * 10);
vert.setZ(v_z * 10);

// convert to MeV's
kin_energy = kin_energy * 1000;
const HepPDT::ParticleData* particle = m_particleTable-> \
    particle(HepPDT::ParticleID(abs(id)));
double mass = particle->mass().value();
double energy = kin_energy + mass;
double pmodule = sqrt(pow(energy,2)-pow(mass,2));
mom.setX(pmodule * d_x);
mom.setY(pmodule * d_y);

```

```

mom.setZ(pmodule * d_z);

Vertex.setVect(vert);

Vertex.setE(m_time);

Momentum.setVect(mom);

Momentum.setE(energy);

// end modification

    double polx = 0;

    double poly = 0;

    double polz = 0;

HepMC::Polarization thePolarization;

    thePolarization.set_normal3d(HepNormal3D(polx,poly,polz));

    m_polarization.push_back(thePolarization);

// units are already converted to MeV's and mm.

    m_fourPos.push_back(Vertex);

    m_fourMom.push_back(Momentum);

    m_pdgCode.push_back(id);

}

if (m_ffile.eof())
{
    log << MSG::FATAL << "End of file reached - stop " << endreq;
    return StatusCode::FAILURE;
}
}

```

This code works in Athena 12.0.6. All necessary units conversions are done, and reversal of z and momentum is done to replicate the dataset for beam 2, if the appropriate flag is given. The jobOptions that correspond to this code (to allow it to produce the correct results) are:

```

#=====

theApp.setup( MONTECARLO )

include( "PartPropSvc/PartPropSvc.py" )

#-----

# Private Application Configuration options

#-----

#load relevant libraries

theApp.Dlls  += [ "CosmicGenerator", "TruthExamples" ]

theApp.TopAlg = [ "CosmicGenerator", "DumpMC" ]

# The following is needed to load the Athena Random

# Number Generation Service.

theApp.ExtSvc += [ "AtRndmGenSvc" ]

AtRndmGenSvc = Service( "AtRndmGenSvc" )

AtRndmGenSvc.Seeds = [ "SINGLE 2040160768 443921183" ]


include( "AthenaPoolCnvSvc/WriteAthenaPool_jobOptions.py" )

#-----

# Private Application Configuration options

#-----

theApp.Dlls  += [ "TruthExamples" ]

theApp.Dlls  += [ "GeneratorObjectsAthenaPoolPoolCnv" ]

AthenaPoolCnvSvc = Service( "AthenaPoolCnvSvc" )

# AtRndmGenSvc.ReadFromFile = true;

#-----

MessageSvc.OutputLevel      = 1

#-----

# Event related parameters

#-----

```

```

# Number of events to be processed (default is 10)
theApp.EvtMax = -1

EventSelector = Service( "EventSelector" )
EventSelector.RunNumber = 8851

StoreGateSvc = Service( "StoreGateSvc" )
StoreGateSvc.Dump = FALSE

#-----
# CosmicGenerator parameters
#-----

CosmicGenerator = Algorithm( "CosmicGenerator" )
CosmicGenerator.eventfile = \
'pdgID_tcth.411.b1_lowb.coll_hor_b1.fluka.txt'
CosmicGenerator.haloswap = True
AthenaPoolCnvSvc.CommitInterval = 10000

Stream1 = Algorithm( "Stream1" )
Stream1.ItemList += [ "EventInfo#*", "McEventCollection#*" ]
Stream1.OutputFile = "halo.C.evgen._00001.pool.root"
#=====

# End of job options file
#####

```

The outputs of these jobs are POOL files, ready for ATLAS simulation. The beam halo side A has been assigned dataset 8850, and side C is 8851

B.4 Beam halo simulation

Since the halo simulation requires no special modifications to the ATLAS/Geant4 interface, it is run as a standard CSC 13.0.30 simulation job, with appropriate filenames inserted by a bash script:

```
csc_atlasG4_trf.py \
    inputEvgenFile=${EVGEN}\
    outputHitsFile=${SIMUL}\
    maxEvents=${NUM_EVTs}\
    skipEvents=${skip}\
    randomSeed=$RANDOM\
    geometryVersion=ATLAS-CSC-01-02-00\
```

I also used the CSC transforms and the ATLAS-CSC-01-02-00 geometry for the plain digitization and reconstruction of the samples. The results of this simulation are direct inputs to the pileup step.

B.5 Pileup and its variations

The Athena 13.0.30 release does not include the more exotic forms of pileup, such as halo and beam gas, that I needed for this study. Thanks to Paolo Calafiura, preliminary modifications to the pileup system were made available for me, allowing both inclusion of different pileup streams and fine-grained settings on the rates at which the streams would seed events into the piled-up event. Two means of determining the pileup insertion rate per event are available - Fixed, meaning there are n pileup events from the given cache that will go into each bunch crossing, and Poisson, meaning that events will be inserted consistent with a Poisson distribution, centered around the value being passed it.

In principle, the cache of pileup events can be far smaller than the actual number needed for a given set of events, each with 69 BCs being examined (-36 to 32). each event is read a given number of times at random intervals, and then recycled. What happens in this particular case (patched 13.0.30) is rather less reliable. From the datasets produced, event reuse at very high rates made the distributions unreliable. The most important thing, however, is that the cache, when loaded from a set of files, is always loaded sequentially. If there is a mix of files from various sources to reflect certain proportions, for example, care must be taken to always

randomize the file load order (expressed in a python list) so that all the files are exposed to the front of the line at some point.

For these runs, I set the event reuse rate to 1 - that is, any event will only be consumed once in a run of pileup. Randomization of the files was done with a python preflight script that scanned the files available, chose correct proportions, populated a list, and randomized it.

Another important aspect of pileup runs is the system memory and data IO rate in the computing resources one is using. A given pileup event is memory-intesive - most of my observed runs required above 1.2 GB of RAM per running core to reliably run the events. The IO necessary for access to the pileup source collections was the primary bottleneck, however. After using several NFS shared disks to store the event collections, and finding that the IO rates swamped the disk servers and made them unresponsive and unstable, I moved the source to CASTOR, counting on its superior throughput to compute nodes at the CERN site. Like the local NFS disks, though, CASTOR had latency issues and download speed problems that took the events, whose time on a Core 2 Duo Xeon processor at 2.33 GHz came in at approximately 45 s, would routinely vary between 5 min. and several hours per event.

The solution of merit came by taking sections of the pileup sources and copying random, proportional reduced datasets to cache disks on all of the compute nodes, and sourcing/randomizing from those file lists. Event production ran smoothly after that.

Making pileup work in 13.0.30 required several packages to be compiled with specific tags, and the ones I used are listed below:

- Generators/GeneratorModules-01-03-32-02
- Control/PileUpComps-00-03-00
- Control/PileUpTools-01-02-00
- Reconstruction/Jet/JetRec-00-04-52-18

Other useful tags for those interested in muons are below:

- MuonSpectrometer/MuonIdHelpers-00-07-13-02

- MuonSpectrometer/MuonGeoModel-00-05-01-01
- MuonSpectrometer/ MuonCnv/MuonRdoToPrepData-00-04-02-25
- MuonSpectrometer/ MuonCalib/MdtCalib/MdtCalibSvc-01-06-03-01
- MuonSpectrometer/MuonReconstruction/

MuonRIO_OnTrackCreators/

MdtDriftCircleOnTrackCreator-02-00-04-02

To add halo collections in correct proportions to the pileup, I defined an input collection as a list with the following python script:

```
import os, commands
from random import randint
def randomize(l):
    ''' Randomizes a list by swapping its members '''
    for i in range(len(l)*10):
        r,rr=randint(0,len(l)-1),randint(0,len(l)-1)
        l[r],l[rr]=l[rr],l[r]
    return l
path='/castor/cern.ch/user/s/stradlin/BeamHalo/simul'
short='pup.008850.halo_A_fluka'
cmd='nsls -l %s/%s |grep -v " 0 "|grep HITS|
awk \'{print $9}\' % (path,short)
haloAColl = ['rfio:%s/%s/%s' % (path,short,i)
for i in commands.getoutput(cmd).split('\n')]
short='pup.008851.halo_C_fluka'
cmd='nsls -l %s/%s |grep -v " 0 "|grep HITS|
awk \'{print $9}\' % (path,short)
haloCColl = ['rfio:%s/%s/%s' % (path,short,i)
```

```

for i in commands.getoutput(cmd).split('\n')]

# Handles the A/C imbalance - factor of 57 over 403 files
beamHaloCollections=randomize(randomize(haloAColl)+
randomize(haloCColl)[:7])

xdir='/xrootd/stradling/BeamHalo/'

atemp=[]
ctemp=[]

try:
    os.stat(xdir)

    ltemp=os.listdir(xdir)

    for i in ltemp:
        if i.startswith('pup.008850'):
            atemp.append(i)
        if i.startswith('pup.008851'):
            ctemp.append(i)

    if len(ltemp) > 150:
        ctemp=randomize(ctemp)[:len(atemp)/57]
        atemp=randomize(atemp)

        beamHaloCollections=randomize(atemp+ctemp)

        beamHaloCollections=[xdir+i for i in beamHaloCollections]

    else:
        pass
except OSError:
    pass

```

which is then imported at the head of the created python jobOptions, seen under the cat EOF below.

```
#!/bin/bash
```

```

#===== Test the input =====
if test \$# -ne 5; then
  echo "USAGE: pup.sh <Job number> <Physics Short> <stage>
<events> <mb rate>"
  echo "EXAMPLE: pup.sh 1 5001.pythia_minbias halo 200 0.47"
  echo "Something wrong with your input, please re-run it."
  echo "Exit....."
  exit 0
fi
export JOB_INT=$1
export PHY_SHORT=$2
export STAGE=$3
export EVENTS=$4
export MB_RATE=$5
export PHY_SHORT=misal1_mc12.00${PHY_SHORT}
export OUT_PATH=/castor/cern.ch/user/s/stradlin/Pileup
export JOB_NUMBER='users/stradling/Scripts/zpad.py ${JOB_INT} 5'
export PHY_NUMBER=${PHY_SHORT:12:6}
shift 5;
#===== Common Environment =====
ln -s /afs/cern.ch/user/a/atlcond/coolrep/sqlite130
ln -s /afs/cern.ch/user/a/atlcond/coolrep/sqlite200
ln -s /afs/cern.ch/atlas/project/database/DBREL/
packaging/DBRelease/current/geomDB
source /users/stradling/cmthome/setup.sh -tag=13.0.30
export base=/users/stradling/Athena/13.0.30:$base
export SEAL_PLUGINS=/users/stradling/Athena/13.0.30/
InstallArea/i686-slc4-gcc34-opt/lib/modules:$SEAL_PLUGINS

```

```
export JOBOPTSEARCHPATH=/users/stradling/Athena/13.0.30/
InstallArea/jobOptions:$JOBOPTSEARCHPATH
export LD_LIBRARY_PATH=/users/stradling/Athena/13.0.30/
InstallArea/i686-slc4-gcc34-opt/lib:$LD_LIBRARY_PATH
export CMTPATH=/users/stradling/Athena/13.0.30:$CMTPATH
export run=/users/stradling/Athena/13.0.30/PhysicsAnalysis/
AnalysisCommon/UserAnalysis/run:$run
export PATH=/users/stradling/Athena/13.0.30/InstallArea/
share/bin:/users/stradling/Athena/13.0.30/InstallArea/
i686-slc4-gcc34-opt/bin:$PATH
export DATAPATH=/users/stradling/Athena/13.0.30/InstallArea/
share:$DATAPATH
export XMLPATH=/users/stradling/Athena/13.0.30/
InstallArea/XML:$XMLPATH
export PYTHONPATH=/users/stradling/Athena/13.0.30/
InstallArea/python:$PYTHONPATH
export TestArea=/users/stradling/Athena/13.0.30:$TestArea
export ROOTMAPSEARCHPATH=/users/stradling/Athena/13.0.30/
InstallArea/rootmap:$ROOTMAPSEARCHPATH
#===== set up athena =====
export HALO_TRUE='False'
export BGAS_TRUE='False'
export LOGDIR=/xdata04/stradling/puplogs
export PUP_NAME=\${PHY_SHORT}.pileup.RDO.\$STAGE._\${JOB_NUMBER}.pool.root
export PUP_LOG=\${PHY_SHORT}.pileup.RDO.\$STAGE._\${JOB_NUMBER}.log.txt
# Turn on beam gas or halo for the right flags.
```

```

if [ $STAGE == 'halo' ]; then export HALO_TRUE='True'; fi
if [ $STAGE == 'bgas' ]; then export BGAS_TRUE='True'; fi
ln -s ~stradlin/Scripts/PileupBase/*.py .
for i in `seq 0 50 $(( ${EVENTS}-1 ))` ; do
  cat <<EOF > myPileup${i}.py
#####
# Job options file to run PileUp. Need AtlasDigitization.pyx
#=====
#-----
# Get Configuration flags:
#-----
import AthenaCommon.Include
AthenaCommon.Include.excludeTracePattern += [
  '/beamGasFiles.py',
  '/beamHaloFiles.py',
  '/minBiasFiles.py',
  '/ioParameters.py'
]
include( "Digitization/Digitization_flags.py" )
include("beamHaloFiles.py")
include("beamGasFiles.py")
include("minBiasFiles.py")
include("ioParameters.py")
EvtMax = 50
SkipEvents = skipDict['${PHY_NUMBER}'][int('${JOB_NUMBER}')]+${i}
PoolHitsInput= runDict['${PHY_NUMBER}'][int('${JOB_NUMBER}')]
PoolRDOOutput="${PUP_NAME}"
#-----

```

```
# Pileup configuration
#-----

initialBunchCrossing = -36
finalBunchCrossing = 32
doMinimumBias=True
numberOfCollisions = ${MB_RATE}
minBiasInputCols = minBiasCollections[:10000]
doCavern = False
#Double-sided, normal proportions
doBeamGas = ${BGAS_TRUE}
numberOfBeamGas = 0.3 #Keeps the beam gas rate up for study.
beamGasInputCols = beamGasCollections
doBeamHalo =${HALO_TRUE}
numberOfBeamHalo = 0.8 #Keeps the beam halo rate up for study
beamHaloInputCols = beamHaloCollections
#-----

# Set some of the global flags. Like eg the DD version:
#-----

DetDescrVersion="ATLAS-CSC-01-02-00"
#select detectors
DetFlags.all_setOn()
DetFlags.Truth_setOn()
#print job configuration
DetFlags.Print()
include("Digitization/AtlasDigitization.py")
# Fixes for 13.0.30
from GeneratorModules.GeneratorModulesConf
import MergeMcEventCollection
```

```

MergeMcEventCollection.DoSlimming = False

from AthenaCommon.AthenaCommonFlags import jobproperties
jobproperties.Digitization.rndmSeedOffset1.set_Value($RANDOM)
jobproperties.Digitization.rndmSeedOffset2.set_Value($RANDOM)

from AthenaCommon.AppMgr import ServiceMgr
ServiceMgr.MessageSvc.OutputLevel = INFO
ServiceMgr.MessageSvc.setVerbose += [ "PileUpStream" ]
pupElm = ServiceMgr.PileUpEventLoopMgr
pupElm.OutputLevel = VERBOSE
pupElm.bkgCaches["MinBiasCache"].OutputLevel = DEBUG
pupElm.bkgCaches["MinBiasCache"].ReadDownscaleFactor = 1
pupElm.bkgCaches["MinBiasCache"].CollDistribution = "Poisson"
if doBeamGas:
    pupElm.bkgCaches["BeamGasCache"].OutputLevel = DEBUG
    pupElm.bkgCaches["BeamGasCache"].ReadDownscaleFactor = 1
    pupElm.bkgCaches["BeamGasCache"].CollDistribution = "Poisson"
if doBeamHalo:
    pupElm.bkgCaches["BeamHaloCache"].OutputLevel = DEBUG
    pupElm.bkgCaches["BeamHaloCache"].ReadDownscaleFactor = 1
    pupElm.bkgCaches["BeamHaloCache"].CollDistribution = "Poisson"
#Open file in "update" mode
flag=${i}
if flag:
    ServiceMgr.PoolSvc.FileOpen = "update";
#####
EOF
time athena.py -s -p Digitization/pileUpBootstrap.py \
myPileup${i}.py &> ${PUP_LOG}.${i}

```

```

cat ${PUP_LOG}.${i} >> ${PUP_LOG}

done

checkFile.py ${PUP_NAME} >> ${PUP_LOG}

rfmkdir -p ${OUT_PATH}/${PHY_SHORT}/${STAGE}

rfcp ${PUP_NAME} ${OUT_PATH}/${PHY_SHORT}/\
${STAGE}/${PUP_NAME}

#End Script

```

Note that the ReadDownscaleFactors are set to 1 - this controls the buggy event recycle rates until further attention can be given to the problem. Another kludge is the need to run events in blocks of 50. It was observed that during pileup runs, there were consistent problems with writing the

- CaloCalibrationHitContainer_LArCalibrationHitActive,
- CaloCalibrationHitContainer_LArCalibrationHitInactive and
- CaloCalibrationHitContainer_LArCalibrationHitDeadMaterial

containers - consistently on the 83rd, 126th or 193rd event in each run, failing 33% of the total runs. To keep the file numbers down, the POOL file was reopened in update mode for all subsequent runs.

B.5.1 Adding beam gas

One can derive the necessary beam gas settings from above:

```

doBeamGas = True

numberOfBeamGas = 0.3 #Keeps the beam gas rate up for study.

beamGasInputCols = beamGasCollections

pupElm.bkgCaches ["BeamGasCache"] .OutputLevel = DEBUG

pupElm.bkgCaches ["BeamGasCache"] .ReadDownscaleFactor = 1

```

```
pupElm.bkgCaches["BeamGasCache"].CollDistribution = "Poisson"
```

The real beam gas rate for any normal study will, of course, vary from the rate I set here, justified for sample enrichment. Rates vary linearly with gas partial pressure and beam current, and can be calculated from the starting points given in Tables 9.2 and 9.3.

B.5.2 Adding halo

Similarly with the halo, the following settings are valid. The halo rate is dependent only on beam lifetime (or rather, the beam lifetime depends on how much halo it sheds), so the ranges mentioned in Section 9.3 are germane. In the case of comparative halo and gas rates, beam gas is negligible compared to the halo rate under most circumstances.

```
doBeamHalo =True
numberOfBeamHalo = 0.8 #Keeps the beam halo rate up for study
beamHaloInputCols = beamHaloCollections
pupElm.bkgCaches["BeamHaloCache"].OutputLevel = DEBUG
pupElm.bkgCaches["BeamHaloCache"].ReadDownscaleFactor = 1
pupElm.bkgCaches["BeamHaloCache"].CollDistribution = "Poisson"
```

B.5.3 Adding correctly proportioned minimum bias

The discussion of minimum bias proportions in Chapter 9 directly informs the choices one makes in realistic minimum bias pileup samples. The proportions of 70%:17.5%:10.2% for nondiffractive:single diffractive: double diffractive are reflected in the randomized file set chosen for the run.

```
#####
# Major changes - no anticipation of corrupt files,
# but it is clear that the input dataset to the pileup
# needs to be randomized by hand. The file list
```

```
# will therefore be read in each time, combined,
# and randomized in the appropriate proportions.

#####
import os, commands
from random import randint
def randomize(l):
    ''' Randomizes a list by swapping its members '''
    for i in range(len(l)*10):
        r,rr=randint(0,len(l)-1),randint(0,len(l)-1)
        l[r],l[rr]=l[rr],l[r]
    return l
path='/castor/cern.ch/user/s/stradlin/MinBias/'
misal1_mc12.005001.pythia_minbias.simul.HITS.v12000701'
ND5001=['rfio:%s/%s' % (path,i) for i in
commands.getoutput('nsls %s' % path).split('\n')]
path='/castor/cern.ch/user/s/stradlin/MinBias/'
misal1_mc12.005003.pythia_sdiff.simul.HITS.v12000318'
SD5003=['rfio:%s/%s' % (path,i) for i in
commands.getoutput('nsls %s' % path).split('\n')]
path='/castor/cern.ch/user/s/stradlin/MinBias/'
misal1_mc12.005004.pythia_ddiff.simul.HITS.v12000318'
DD5004=['rfio:%s/%s' % (path,i) for i in
commands.getoutput('nsls %s' % path).split('\n')]
prop5001=7635
prop5003=500
prop5004=354
minBiasCollections=randomize(
randomize(ND5001) [:prop5001]+\
```

```

randomize(SD5003) [:prop5003]+
randomize(DD5004) [:prop5004]) [:400]
xdir='/xrootd/stradling/MinBias/'
try:
    os.stat(xdir)
    ltemp=os.listdir(xdir)
    if len(ltemp) > 200:
        minBiasCollections=randomize(ltemp)
        minBiasCollections=[xdir+i for i in minBiasCollections] [:400]
    else:
        pass
except OSError:
    pass

```

B.6 Reconstruction

Reconstruction was done with the standard CSC reco transformation for 13.0.30, with ESD and AOD outputs.

```

time csc_reco_trf.py \
    inputRDOFile=${RDO_NAME} \
    outputESDFile=${ESD_NAME} \
    outputAODFile=${AOD_NAME} \
    ntupleFile=${NTUP_NAME} \
    maxEvents=-1 \
    skipEvents=0 \
    geometryVersion=ATLAS-CSC-01-02-00 \
    triggerConfig=DEFAULT \
    DBRelease=/xdata02/stradling/db/DBRelease-4.3.1.tar.gz
    # This file was copied locally for speed and permission reasons.

```