Improving the Robustness of the ATLAS Calorimeter Software Trigger

by

Mark Alexander Baker
B.Sc., Mount Allison University, 2007

A Thesis Submitted in Partial Fulfillment of the

Requirements for the Degree of
MASTER OF SCIENCE
in the Department of Physics and Astronomy

Improving the Robustness of the ATLAS Calorimeter Software Trigger

by

Mark Alexander Baker
B.Sc., Mount Allison University, 2007

Supervisory Committee

Dr. R. V. Kowalewski, Supervisor (Department of Physics and Astronomy)

Dr. J. Albert, Member (Department of Physics and Astronomy)

Supervisory Committee

Dr. R. V. Kowalewski, Supervisor (Department of Physics and Astronomy)

Dr. J. Albert, Member (Department of Physics and Astronomy)

Abstract

The ATLAS experiment pushes the leading edge of experimental particle physics. Increasingly complex hardware, however, brings increasingly complex problems which manifest themselves not only in the detector, but also within the software which drives the detector. The magnitude of the expected interaction rate, too, adds enormous stress to the detector system and the software trigger. In order to prepare the software for these challenges, various detector quantities are considered which may provide debugging handles and robustness against detector problems arising in the ATLAS calorimeter trigger. The effect of electronics noise suppression on these quantities is studied and a brief study of the software trigger performance is followed by recommendations for the implementation of robustness checks.

Table of Contents

List of Tables

List of Figures

Acknowledgements

This research was funded in part by a grant from NSERC. This thesis could not have been constructed without the continued patience and assistance of Diego Casadei, Ryan Taylor, and Bob Kowalewski.

Chapter 1

Introduction

Spanning the border between France and Switzerland just outside of Geneva, the Large Hadron Collider (LHC) marks the forefront of experimental high energy physics. Situated 100 m underground and measuring 27 km in circumference, the LHC is designed to accelerate and collide opposing beams of protons in the search for the source of electroweak symmetry breaking and the means by which fundamental particles acquire mass. The 14 TeV center-of-mass energy and $10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ design luminosity make way for searches at unprecedented energy scales. The widely publicized Higgs mechanism [1], which awaits its potential discovery at the LHC, is one possible solution to the electroweak symmetry problem.

Along the path of the colliding beams are situated four large-scale detectors in which these proton-proton (p-p) collisions occur. A Toroidal LHC ApparatuS (ATLAS), which concerns the study of this paper, is one such general-purpose detector. A three-level trigger system spanning from custom-designed electronics to complex software algorithms provides ATLAS with the necessary rejection power to meet the stringent recording limitations, reducing the 40 MHz bunch-crossing rate to a manageable 200 Hz data acquisition rate. The trigger must be streamlined and robust; an early-rejection philosophy ensures that uninteresting events are eliminated with a

minimum of processing. This necessitates a good understanding of the detector and algorithms which are robust against any of a series of hardware and software faults.

This document concerns the implementation of robustness checks in a certain part of the software-based trigger. Chapter 2 describes the detector and the relevant readout electronics, leading naturally into Chapter 3 in which the three levels of the trigger are explained. An analysis of robustness quantities and techniques follows in Chapter 4 and the conclusions and recommendations drawn as a result of these studies completes the discussion in Chapters 5 and 6.

Chapter 2

ATLAS Detector

Cylindrically shaped, the ATLAS detector [2] is effectively symmetric about the plane which intersects the cylinder at its center – the two sides of the detector are referred to as side A and side C respectively. The ATLAS detector is aligned along the beam axis which is assigned to the z-coordinate. The plane transverse to the z-axis is hence the x-y plane, with the positive x direction defined to be pointing to the center of the LHC ring and the positive y direction pointing upwards. Positive z is defined from the nominal interaction point (the place at which the two beams collide, effectively located in the center of the detector) pointing towards side-A of the detector. The azimuthal angle ϕ is measured about the beam axis and the polar angle θ gives the angle subtended from the z-axis. A convenient alternative measure for θ is given by $\eta \equiv -\ln\tan(\theta/2)$, called the pseudorapidity, over which particle production at the LHC will be roughly constant because of the high momentum scale of the particles.

The detector can be described in terms of its functionality by beginning at the interaction point and moving outwards (begin at the center of Figure 2.1, inside of the pixel detector and move outwards). Just outside of the center is found the inner detector system, immersed in a 2 T solenoidal field, which provides tracking and vertex resolution for charged tracks with p_{T} as low as 0.5 GeV and $|\eta| < 2.5$.

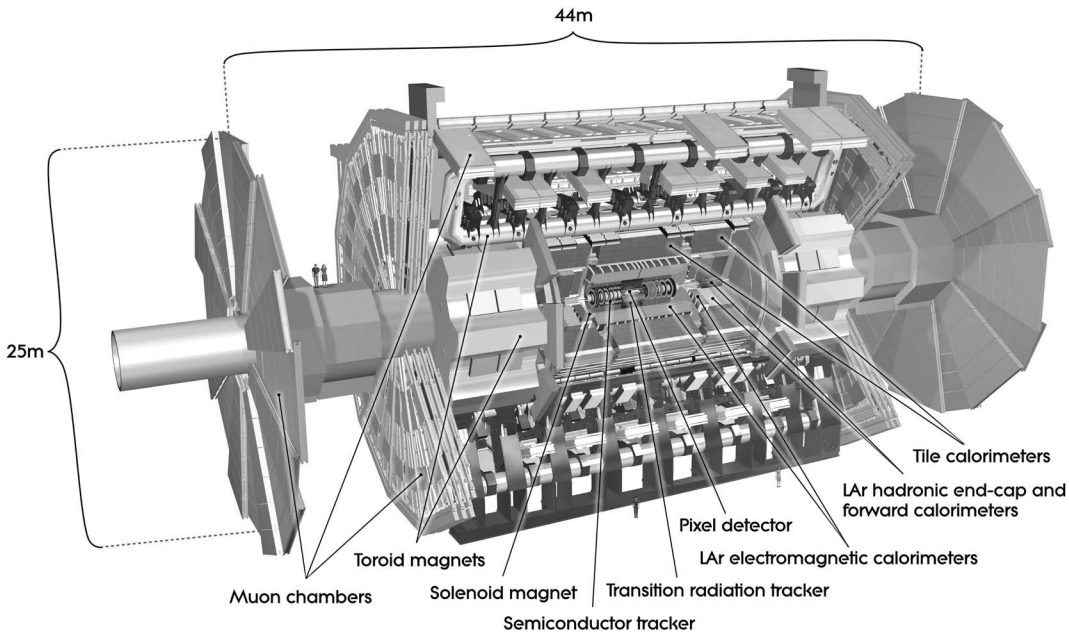

Figure 2.1: Cross-section of the entire ATLAS detector. Diagram taken from [2].

Following further from the point of interaction and beyond the inner detector is found the electromagnetic (EM) calorimeter in which energy from photons and charged particles is sampled. Moving another layer beyond is the hadronic calorimeter which samples energy from hadrons and jets. Finally, the muon spectrometer and a system of toroid magnets encompasses the region $|\eta| < 2.7$ and makes up the majority of the volume of the ATLAS detector.

While these pieces together make up the whole of the detector, only the portions relevant to the trigger robustness work contained in the analysis section will be described in detail.

2.1 Calorimetry

ATLAS calorimeters cover the pseudorapidity range $|\eta| < 4.9$ with varying granularity suited to meet the physics demands for sensitive measurements such as missing transverse energy (MET). Depth in the ATLAS calorimeter is a measurement from

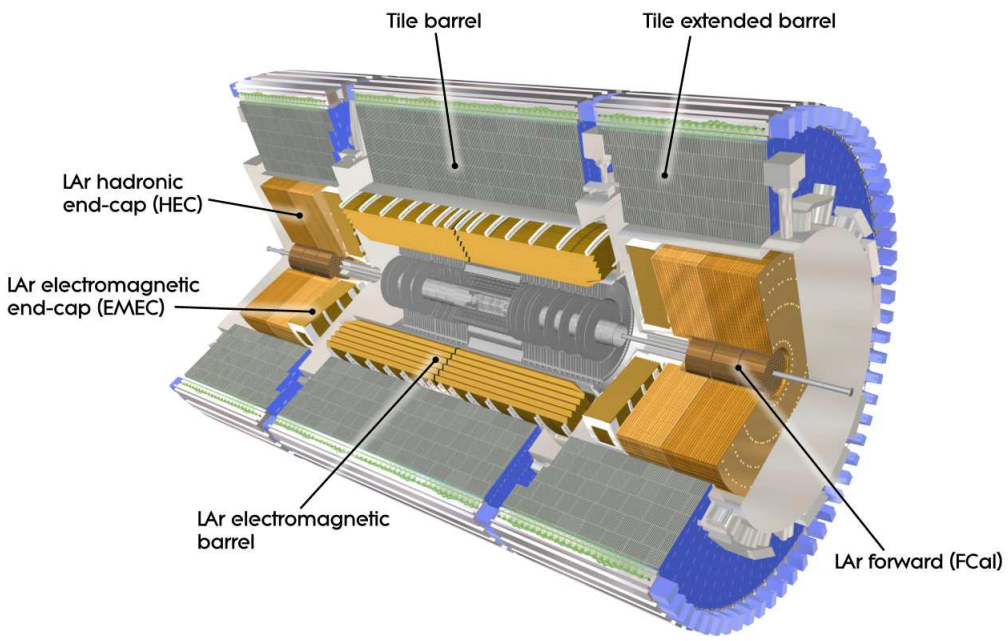

Figure 2.2: Cross-section of the ATLAS calorimeter. Diagram taken from [2].

the nominal interaction point which describes the number of radiation and interaction lengths through which particles must pass following a p-p collision. In the same way that the granularity varies to meet physics requirements, so too does the depth of the ATLAS calorimeter, starting at 22 radiation lengths (X_0) in the barrel and 24 X_0 in the end-caps.

2.1.1 Electromagnetic Calorimetry

The electromagnetic calorimetry is divided into two barrel segments (EM barrel *EMB*), two end-caps (the EM end-cap *EMEC*), and a portion of the forward calorimeter (FCal), as pictured in Figure 2.2. Each EMB segment measures 3.2 m in length and has inner and outer radii of 2.8 m and 4 m respectively. Modules of the barrel are grouped into 'trigger towers' as pictured in Figure 2.4. Each wheel of the EMEC measures 0.63 m in thickness and has inner and outer radii of 0.330 m and 2.098 m

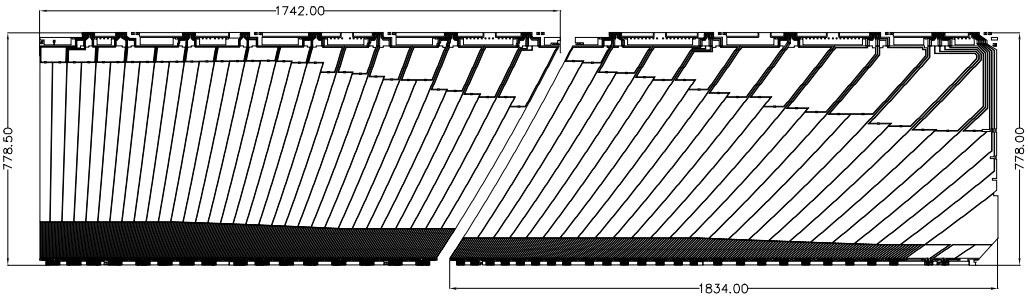

(a) barrel

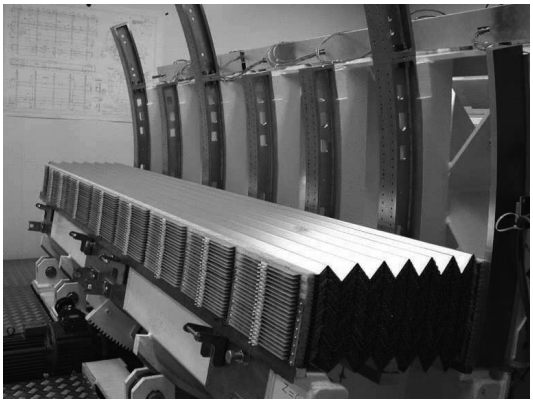

(b) folded barrel segment

(c) folded end-cap segment

Figure 2.3: Illustrated here are the three layers of the EMB (2.3a) before folding of the electrodes (units are in mm). 2.3b shows a portion of the EMB after folding and 2.3c shows a folded end-cap segment. Diagrams taken from [2].

respectively. In both the barrel and the end-caps, lead absorber plates sandwich layers of liquid Argon (LAr), the active detector medium. Liquid argon is so chosen for its radiation hardness and linearity over a large energy range, necessary to meet the physics demands for the products of 14 TeV center-of-mass p-p collisions. In order to eliminate the existence of cracks between segments and provide complete ϕ coverage, the absorber geometry is accordion shaped (see Figures 2.3b and 2.3c). Suspended between these accordion-shaped absorber plates by honeycomb spacers, readout electrodes operate with a drift gap of 2.1 mm, corresponding to a total drift time of roughly 450 ns in the presence of a 2000 V potential difference. Additionally,

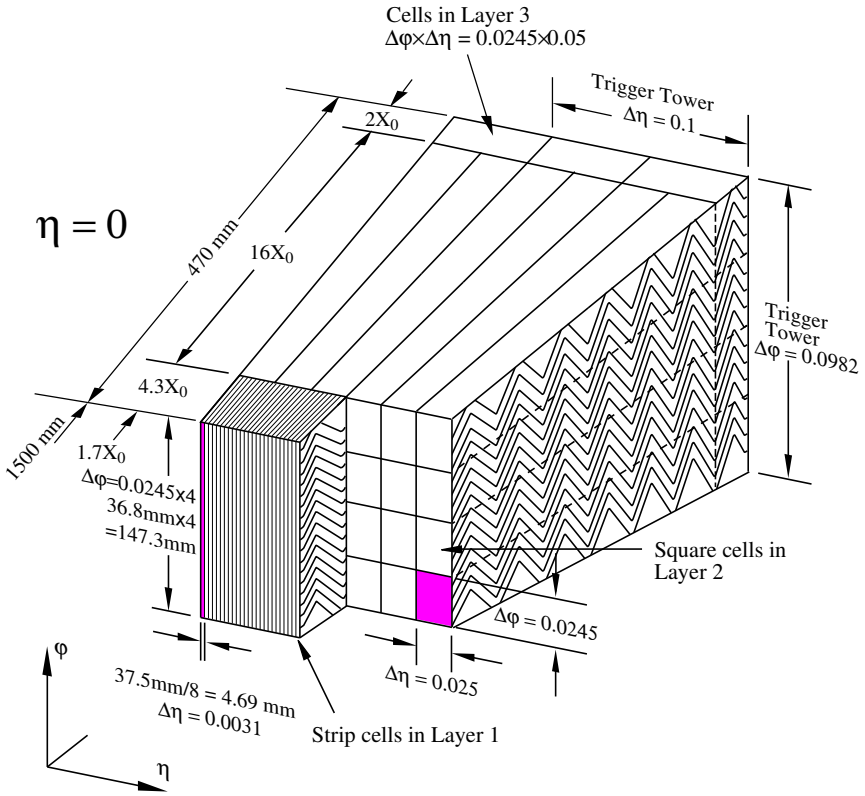

Figure 2.4: Typical module cross section of the EMB. Three distinct layers, the designation of 'trigger tower,' and the accordion geometry are clearly visible. Diagram taken from [2].

this geometry lends itself to the existence of three separate layers within the barrel and two or three layers in the end-caps (see Figure 2.3a). Among the three layers, the layer nearest to the interaction point has the finest granularity where the outermost layer, which receives only the tail of the electromagnetic shower, has the coarsest granularity. The number of channels by layer is given in Table 2.1 and the relative depth of the layers is shown in Figures 2.5a and 2.5b. Trigger towers span the three layers of depth in the barrel. In addition to three separate layers in the barrel, an 11 mm deep LAr presampler sits in front of the first layer to provide shower sampling before the active EM calorimeter.

Surrounding the beam line over the range $3.1 < |\eta| < 4.9$, FCal cells are subject

	PreSampl	First	Second	Third
EMB	7808	57216	28672	13824
EMEC	1536	28544	23424	10240
FCalA	—	1008	500	254
FCalB	—	1008	500	254
HECA	—	1536	1344	—
HECB	—	1472	1280	—
TileBar	Between all layers: 5980			
TileExt	Between all layers: 3640			

Table 2.1: Number of signal channels by calorimeter layer.

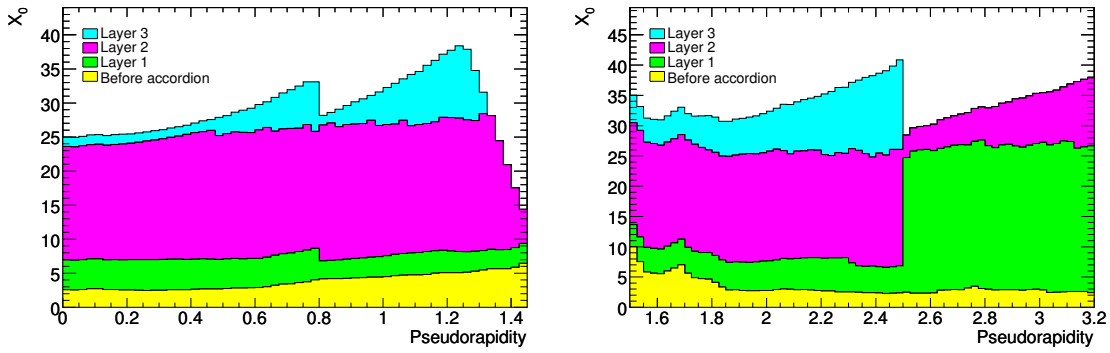

(a) barrel (b) end-cap

Figure 2.5: Cumulative material in radiation lengths of the EMB (2.5a) and the EMEC (2.5b). The jump at $\eta = 0.8$ is the point at which the thickness of the EMB lead absorber plates drops from 1.53 mm to 1.13 mm, limiting the decrease of the sampling fraction with increasing η. Diagrams taken from [2].

to extreme particle flux. For this reason, copper absorbers were chosen to optimize heat removal, each with 12,260 precision holes into which electrodes are inserted. A potential difference of 250 V offers a drift time of 60 ns. Only a small portion of the FCal cells are electromagnetic (FCal1); the hadronic portions (FCal2 and FCal3) are positioned behind FCal1.

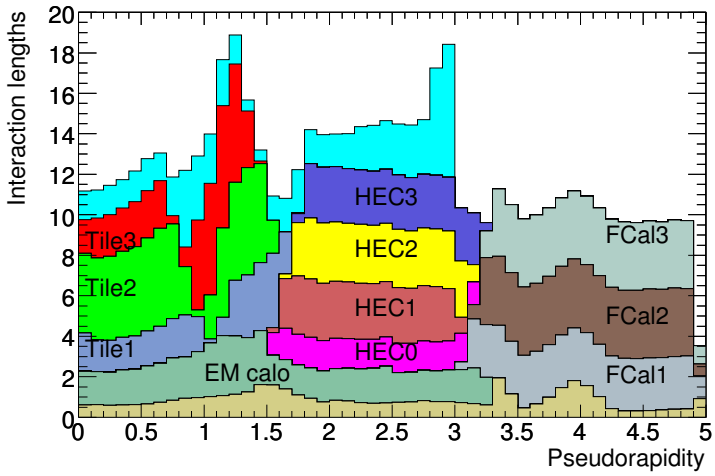

Figure 2.6: Cumulative material in interaction lengths in front of the EM calorimeters, within the EM calorimeters, within each hadronic sampling, and the total amount outside of the active calorimetry. Diagram taken from [2].

2.1.2 Hadronic Calorimetry

Hadronic calorimeters are divided into the tile calorimeter, the FCal, and the LAr hadronic end-cap, known as the HEC (see Figure 2.2).

The 5.8 m long tile barrel encapsulates the EMB while the 2.6 m long tile extended barrel reaches to surround the end-caps with inner and outer radii of 2.28 m and 4.25 m respectively. Together, the tile barrel and tile extended barrel cover the range $|\eta| < 1.7$. Throughout the tile calorimeter, scintillator is used as the active medium, sandwiched between layers of steel absorber plates as pictured in Figure 2.7. Cracks between the barrel and the end-cap cryostats are instrumented with specialized modules which resemble the tile calorimeter in an effort to provide corrections for energy lost to inactive material. Modules in this place cover the range $1.0 < |\eta| < 1.2$ and fall under the umbrella of TileGap, pictured in Figure 2.8.

The HEC is positioned on the outside of the EMEC, constructed from flat plates of copper sandwiching the LAr active medium. Each end-cap is built from two wheels

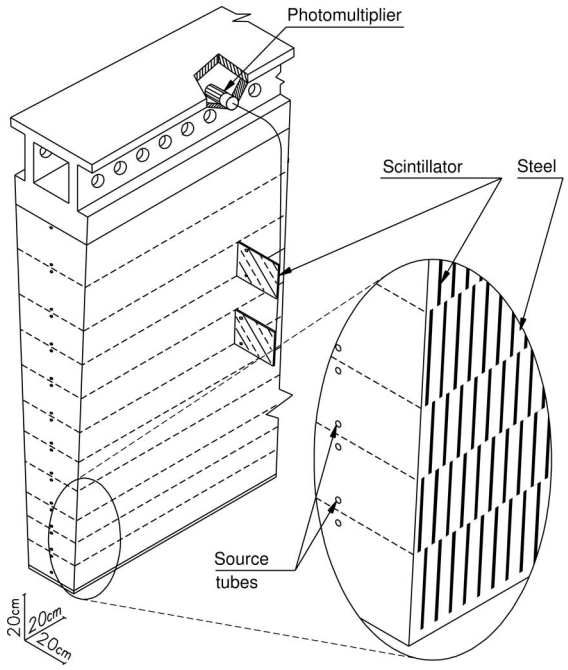

Figure 2.7: Layers of scintillator and steel are tightly arranged in the modules of the tile calorimeter. Diagram taken from [2].

and every plate of each of these wheels has an outer radius of 2.030 m. In order to accommodate the positioning of the FCal, the inner radius increases from 0.372 m for the first nine plates of the first wheel to 0.475 m for the rest of the plates (see Figure 2.9). Together these wheels cover the range $1.5 < |\eta| < 3.2$. Drift gaps of 1.8 mm are maintained by honeycomb spacers and equate to roughly 430 ns drift times for electrons across the 1800 V potential difference. In addition to its duties as a hadronic detector, the HEC serves to detect muons.

The close proximity of the FCal with the HEC offers a nearly hermetic design, minimizing energy loss through cracks between these systems. The highly dense FCal2 and FCal3 employ tungsten to optimize the absorption length of the detector and maintain drift gaps as small as 0.25 mm to avoid issues owing to ion buildup.

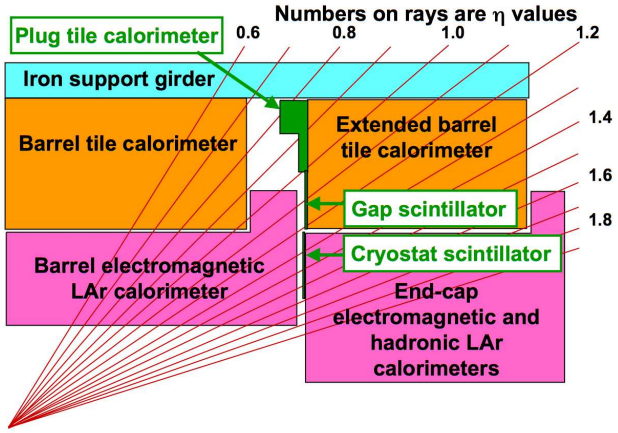

Figure 2.8: Transition region between barrel and end-cap cryostats. Gap scintillators (TileGap) provide corrections for energy lost in inactive material. Diagram taken from [2].

2.2 Muon Spectrometer

The muon spectrometer uses magnetic deflection provided by large toroid magnets and high-precision tracking to detect muons in the largest volume of the ATLAS detector. Two air-core toroid magnets positioned at the ends of a main barrel toroid provide a field which is mostly orthogonal to the muon trajectories, covering the range $|\eta| < 2.7$. In the transition region, $1.4 < |\eta| < 1.6$, the bending field is composed of contributions from each of the barrel and the end-cap toroids.

After muons enter the fields they are detected in either large planar chambers perpendicular to the beam axis ($1.4 < |\eta| < 2.7$) or in the chambers arranged in cylindrical layers about the barrel ($|\eta| < 1.4$). To give a sense of the scale, the chambers positioned about the barrel fall at radii of approximately 5, 7.5, and 10 m from the point of interaction. Each of the end-cap and barrel regions has two distinct chamber types which serve the functions of tracking and triggering. Concentric to the barrel, Monitored Drift Tubes provide precision tracking while Resistive Plate

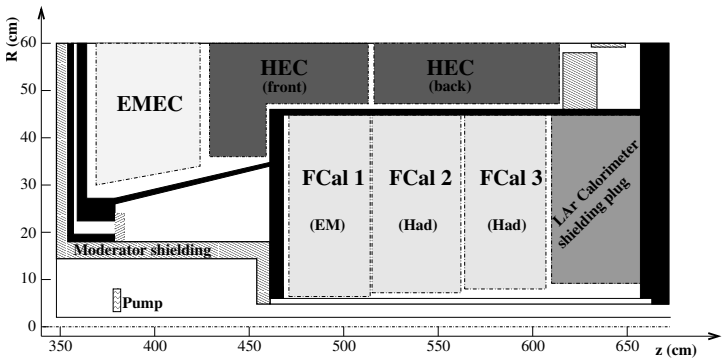

Figure 2.9: The end-cap sub-detectors are closely packed, offering a nearly hermetic design. Note also the increase of the HEC inner radius as the FCal1 is introduced into the end-cap assembly. Diagram taken from [2].

Chambers provide triggering in the range $|\eta| < 1.05$. The analogous chambers aligned perpendicular to the beam are the Cathode Strip Chambers for tracking and Thin Gap Chambers for triggering in the range $1.05 < |\eta| < 2.4$. Triggering chambers also provide well-defined p_{T} thresholds and measure a second muon coordinate – a coordinate orthogonal to that given by the precision tracking chambers.

2.3 Readout Hardware

Readout electronics are the interface between the calorimetry and the Level-1 (L1) trigger. These provide the trigger towers (0.1×0.1 in $\Delta\eta \times \Delta\phi$) for the L1 trigger system. Front-end electronics are housed on the detector systems themselves while back-end electronics are housed off-detector in a cavern 70 m away. Commercial (and therefore, less expensive) electronics are appropriate for back-end electronics since they are not exposed to the immense radiation levels produced by collisions in the detector. However, in the interest of speed, front-end electronics are made to be radiation-tolerant and serve the function of amplifying, processing, and digitizing the analogue signals right on the detector. The logical division of the electronics in the readout of the LAr calorimeter system is the front-end *crate*, of which there are 58

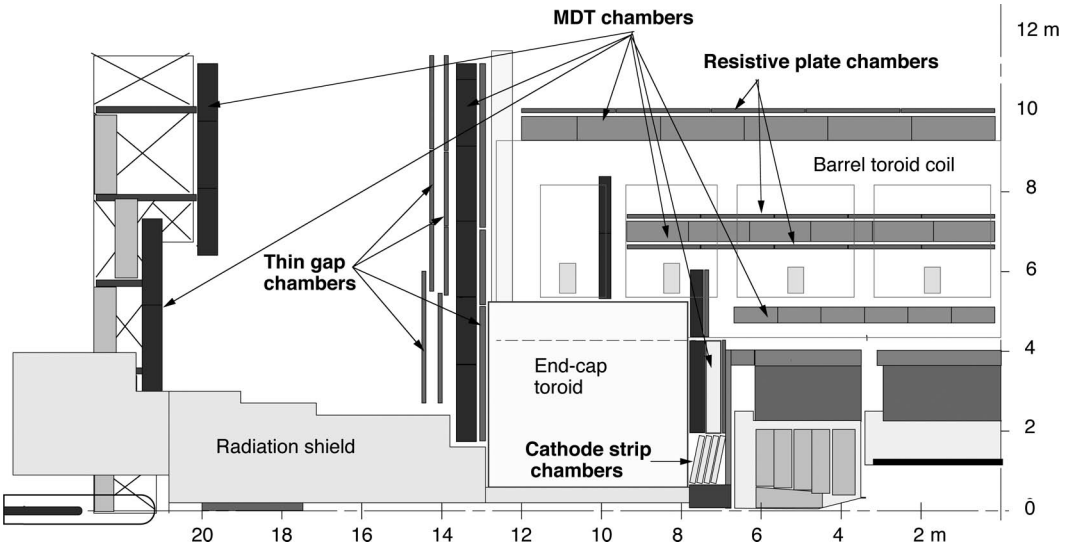

Figure 2.10: Side view of one quadrant of the muon spectrometer. Diagram taken from [3].

in total. Each crate possesses about 28 Front End Boards, known as FEBs (the front end electronic boards which receive the raw calorimeter signal) to which there are connected up to 128 calorimeter channels; each of the 182,468 LAr channels are read out for each event.

FEBs receive the raw signal from the detector and amplify it and shape it. Each signal is split into three overlapping gain scales whose ratios are 1/10/100 with respect to the first, lowest energy scale. The gain scale registering the highest energy without saturation is selected for readout and digitized by a special gain selector chip. This three-scale system allows for a larger dynamic range with an optimal energy resolution. These readout electronics cover the three TeV upper limit to the lower $O(10)$ MeV thermal noise limit of the calorimeters using 12-bit ADCs, equating to 4096 ADC channels.

A typical amplified and shaped signal which is read from the FEB is sampled and is pictured in Figure 2.11, overlaid with an unshaped signal. Notice the shape of

the pulse which is produced by the FEB: the integral is roughly zero and the peak is narrow. These features help to reduce the effects of pile-up[1] owing to abundant QCD backgrounds by having the integrated pulse shapes approximately cancel each other out. Liberated electrons in the LAr medium drift to their electrodes in the presence of the electric field to generate this signal which is sampled at a rate of 40 MHz, the LHC bunch-crossing frequency.

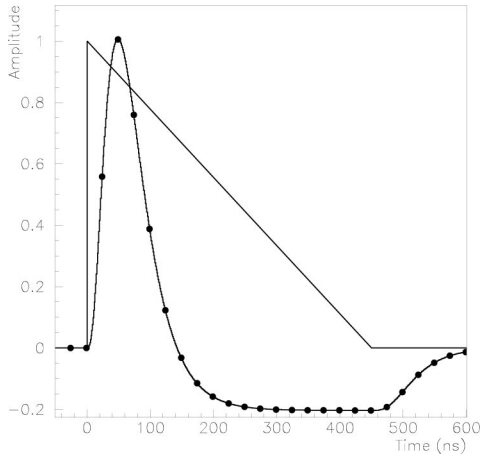

Figure 2.11: Amplitude versus time for a sample pulse before and after shaping through a FEB. The triangular curve represents the signal in a LAr barrel EM cell before the shaping. Diagram taken from [2].

Shaped signals from the FEBs are read by the off-detector trigger hardware and digitally processed, formatted, checked for integrity, and made available for subsequent layers of the trigger and data acquisition system. Readout Driver (ROD) modules are connected to processing-unit cards which house two Digital Signal Processors (DSPs) each. A buffer preceding the DSPs performs quality checks and formats the signal for use in the DSP while a buffer following the DSP holds the processed result. The DSP itself performs the crucial task of calculating the energy and signal timing

[1]Pile-up has multiple definitions. In this sense, pile-up refers to the overlap of signals from one bunch crossing to another. As the signals overlap, the shaped results tend to, on average, cancel out the background energy from minimum bias events such as soft QCD processes.

of each calorimeter cell using Optimal Filtering Coefficients (OFCs) [4]. Calibration circuits located in the front-end electronics provide some of the values used to construct the OFCs and are used to re-measure these values during calibration runs of the detector. The format of the data coming from the Readout System (ROS) is the format which is recorded to tape if the event is accepted and is referred to as bytestream format.

Chapter 3

Trigger

The LHC design luminosity is $10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, offering statistics necessary to probe deeper into high energy, rare physics processes. This substantial rate is accompanied by a substantial volume of data and, since storing recorded data from the experiment is crucial to the existence of offline analysis, requires a powerful rejection factor. With an expected proton-proton interaction rate of 1 GHz, meeting the 200 Hz rate of the financially and technologically limited data acquisition system requires a rejection rate of approximately 5×10^6 [2]. It is the ATLAS trigger which is responsible for such a feat.

3.1 Event Trigger

The ATLAS trigger is divided into three hierarchical levels: L1, Level-2 (L2), and the Event Filter (EF). L2 and the EF together are commonly referred to as the High Level Trigger (HLT). Events accepted at L1 are passed onto the higher levels to further reduce the event rate. In the trigger, the aim is to reject early. In order to meet the stringent minimum-bias rejection factor, the trigger is subdivided into hardware and software levels, each refining the trigger decision made at the previous level. First, specialized hardware at L1 makes a decision in under $2.5\,\mu$s by using

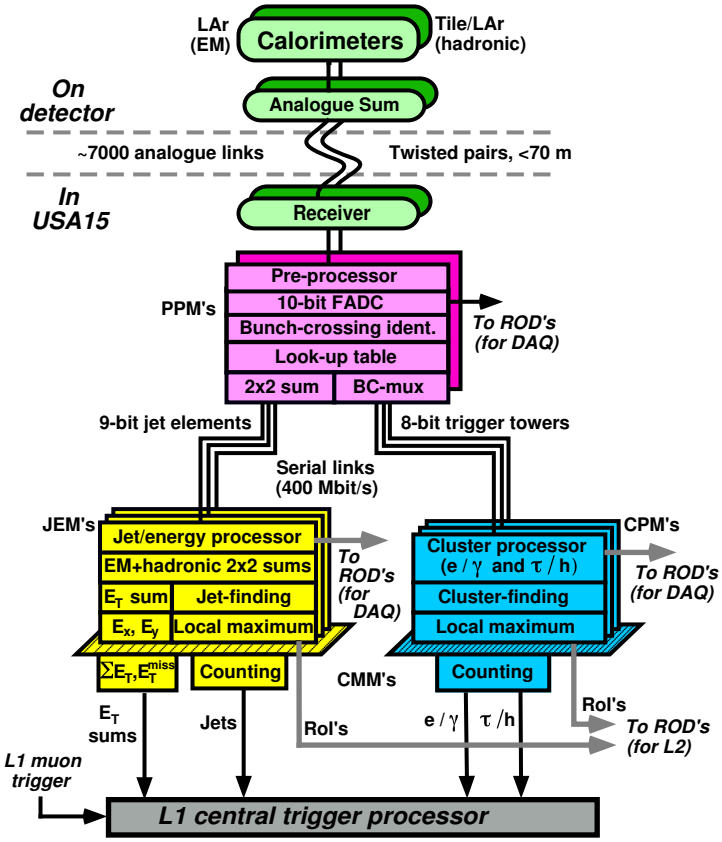

Figure 3.1: Flow diagram of the L1Calo trigger. Analogue input is received from the calorimeters, digitized, and processed. Trigger object multiplicities and energy sums are returned to the CTP. Diagram taken from [2].

only a subset of the full detector information, reducing the trigger rate from 1 GHz to 75 kHz[1]. Software at L2 and the EF offer decisions based on the entirety of the detector and reduce the trigger rate to its final 200 Hz. The following subsections will describe briefly the functioning of the individual trigger levels.

3.1.1 L1 and the L1Calo Trigger

With reduced granularity information and custom electronics at L1, the L1 acceptance rate is capped at 75 kHz (upgradable to 100 kHz). Information from the barrel,

[1]Upgrading the readout system bandwidth would allow data acquisition to operate at 100 kHz. An upgrade will likely commence once budgetary constraints are lifted [2].

end-cap, and forward calorimeters identify high-E_{T} signatures such as electrons, photons, and jets. Multiplicity of the detected trigger objects and global flag information is sent from the calorimeters to the Central Trigger Processor (CTP) where a final L1 decision is made. If the event is accepted, finer-granularity information about the geometry is packaged into ROIs (Regions of Interest) and is buffered for readout by a higher level. The trigger systems are designed such that the period of time from the p-p collision to the decision of the L1 trigger – the trigger latency – may be as high as 2.5 μs.

The data samples of interest in this paper have passed the Level-1 Calorimeter Trigger (L1Calo) trigger, making them interesting candidates for robustness studies since this trigger requires sizeable energy deposits in the calorimeter. Following from the top of Figure 3.1, 7000 reduced granularity (0.1×0.1 in $\Delta\eta \times \Delta\phi$) trigger towers from EM and hadronic calorimeters send information to the off-detector L1Calo system. Analogue signals are digitized by the pre-processor (which are found on the FEBs of the respective front-end readout electronics as explained in Section 2.3), then are matched to their appropriate bunch-crossing by a digital filter. A look-up table produces the E_{T} values which are handed off, along with the signal, to two separate systems: the Cluster Processor (CP) and the Jet/Energy-sum Processor (JEP). Where the CP produces electron and photon multiplicities, the JEP produces jet multiplicities and scalar energy sums of these quantities are sent to the CTP. At the same time, information required for building ROIs is kept in the FEBs until the CTP emits an acceptance signal. The entire process described above finishes execution about 1.5 μs after the collision occurs.

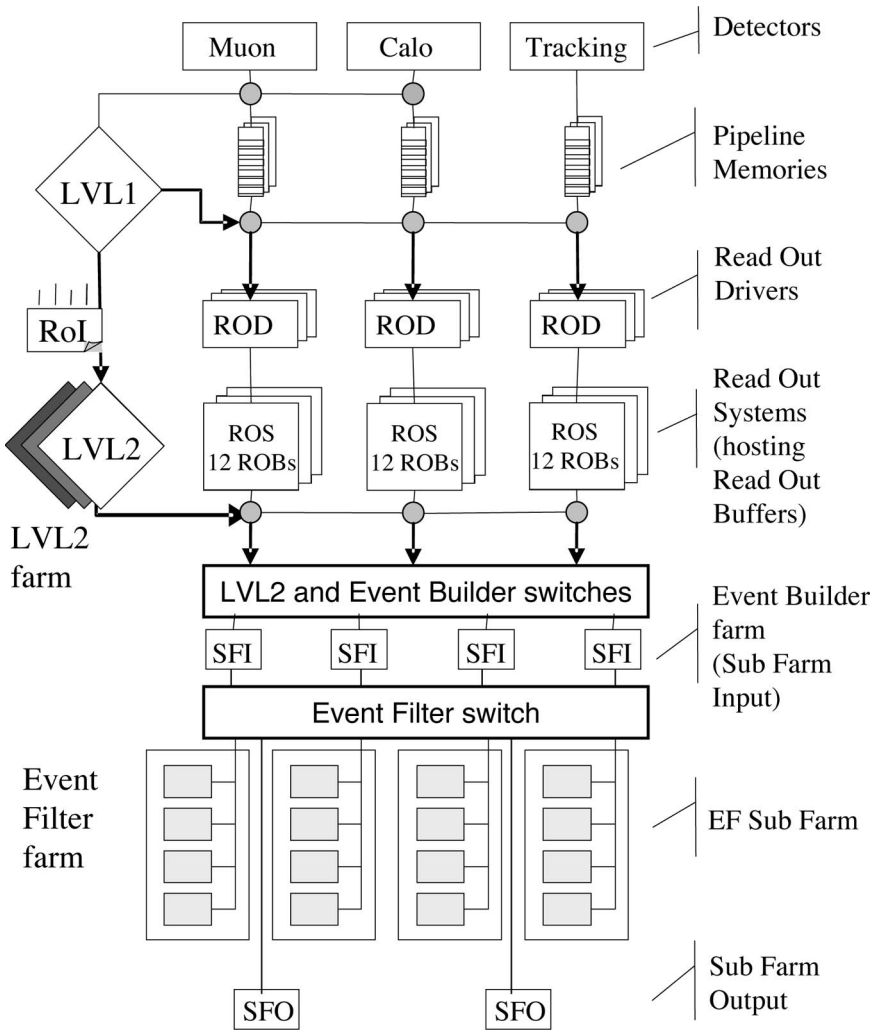

Figure 3.2: Flow diagram of the three levels of the ATLAS trigger. Diagram taken from [5].

3.1.2 L2 and the Event Filter

During the trigger latency at L1, event data buffers in front-end electronics until a L1 accept signal is received. Once this signal is received, buffered event data is transmitted through RODs which perform error checking and event-building tasks, thus packaging the data into L1 trigger objects for use in L2. This packaged data is subsequently loaded into Readout Buffers (ROBs) where it remains available on

request in the readout system. A list of the ROIs travels to L2, providing L2 with the mapping between the geometric location in η and ϕ with the appropriate ROB in which the scalar energy information is being stored. The decisions of the L2 trigger are thus made using only the information found in the ROIs, about 2% of the full event information [5]. Performing more elaborate analysis of this small amount of data is sufficient to reduce the incoming L1 accept rate from 75 kHz to 3.5 kHz. Processing at L2 is performed on a farm of roughly 500 nodes [6] each with 2.5 GHz dual quad-cores, allowing for a per-event processing time of approximately 40 ms.

As events are accepted by L2, data is gathered from the ROS by event builder nodes known as SFIs and processed into fully reconstructed events for use in the EF. If the events are accepted by the EF, they are sent to the SFOs until they are pulled and recorded to permanent storage. At this reconstructed level, the data from calorimeters take the form of software objects called CaloCells. It is from these CaloCell objects that algorithms at the EF level procure information about the cell's energy, geometric position, and sampling. The EF is made up of about 2100 commercial nodes and is given the luxury to perform offline-like analysis decisions within the trigger. The final trigger accept rate after the EF is 200 Hz, with each event receiving up to 4 s of processing time.

Both in L2 and the EF (the HLT), trigger processing is made up of two building blocks: feature extraction (Fex) and hypothesis. If it is desired to trigger on a certain quantity, that 'feature' is extracted using the Fex and the hypothesis is tested, either accepting or rejecting the event. *Chains* of interleaved Fex algorithms and hypotheses are seeded at L1, tested on the first hypothesis, and then passed to the next 'link' in the chain. At the next link (perhaps in L2, for example) a feature may be extracted and then another hypothesis is tested. In order for the event to be accepted and ultimately recorded, it must satisfy all of the hypotheses in the chain; it is chains

which define acceptance. Under this scheme, feature-computing algorithms can be re-used which minimizes redundant code development. Chains which are ultimately designed to trigger on a certain physics object such as an electron, photon, or τ are separated into *slices* to facilitate the organization of the trigger software effort. Additionally, many Fex employ caching to avoid duplicate calculations, inevitable given that a single Fex may be run multiple times owing to its presence in chains within various slices.

3.2 MET Trigger Slice

Theoretically, interactions must necessarily conserve energy and momentum. Therefore, if the sum of the energy in the detector appears to favor a certain direction for a particular event, it may be the case that an undetected particle has carried away some portion of the energy. For this reason, missing energy is a signature of great interest for SUSY [7] models of particle physics. Furthermore, since a fine-granularity sum of the energy in all calorimeter cells cannot be achieved at L2 owing to time constraints, the MET trigger slice is based almost entirely at the EF level.

Practically, however, the z component of missing energy cannot be known because the original momentum in this direction is not well known. For, although the energy of the colliding particles is well defined, the parton distribution functions which describe the composition of the protons allow for collisions in which the net z momentum is non-zero. Additionally, much of the energy in an event may slip through small angles along the beam line. This drives the generalized measurement of missing energy to a restricted form, one in which only the energy in the x-y plane is considered – that is, only missing *transverse* energy. Hence, MET is given by a two-dimensional vector constructed from the negative of the energy sum over the whole detector and offers

a signature of great physical importance. Formally, the vector sum of MET is given by

$$\not\boldsymbol{E}_{\mathrm{T}} \equiv -\sum \boldsymbol{E}_{\mathrm{T}} = -\sum_i (E_{xi}\,\hat{\mathbf{x}} + E_{yi}\,\hat{\mathbf{y}}),$$

with i taken over all CaloCells[2]. In the MET slice, the components $\not E_x$ and $\not E_y$ of this vector, along with the scalar sum

$$\Sigma E_{\mathrm{T}} \equiv \sum |\boldsymbol{E}_{\mathrm{T}}| = \sum_i \sqrt{E_{xi}^2 + E_{yi}^2},$$

the total transverse energy, form the basis for features used in MET slice trigger chains. In particular, for first data taking at the LHC, inclusive thresholds on ΣE_{T} will be 150, 250, 360, and 650 GeV and inclusive thresholds on $\not E_{\mathrm{T}}$ will be 15, 20, 25, 30, 40, 50, 70, and 80 GeV. These chains will be prescaled[3] such that their final acceptance rates are equal.

[2]For other purposes, the sum may be taken over the low-granularity FEBs instead of high-level CaloCells at the EF. The information from the FEBs is available faster, allowing for faster trigger decisions with reduced granularity.

[3]Prescaling a chain refers to reducing the outgoing rate of the chain after processing events. For example, if a trigger rate on a certain chain were too high for the ATLAS data acquisition system, a prescaling could be applied which offered only one in every ten acceptable events to the next level the chain.

Chapter 4

Analysis

Robustness checks in $E_{\mathrm{T}}^{\mathrm{miss}}$ provide handles for further investigation of an issue after the data has been recorded. Such offline analysis should be able to identify, for example, the cause of a spike in the trigger rate originating from within the $E_{\mathrm{T}}^{\mathrm{miss}}$ slice. Identifying which quantities provide the best handles when computed within a reasonable time budget is the purpose of the analysis. Following is a description of the $E_{\mathrm{T}}^{\mathrm{miss}}$ Fex, the data samples used in this study, a lengthy consideration of candidate quantities, and a brief code performance study. From this information, recommendations concerning robustness are made in Chapter 6.

4.1 $E_{\mathrm{T}}^{\mathrm{miss}}$ Feature Extraction

The EF $E_{\mathrm{T}}^{\mathrm{miss}}$ Fex makes use of the full detector information, looping over each cell in the LAr and tile calorimeters. Additionally, the energy contribution from muons is calculated using information retrieved from the muon spectrometer. In software, this procedure is broken into four tools:

- `EFMissingETFromMuons("TheMuonTool")`
- `EFMissingETFromCells("TheCellTool")`

Group	Components
EMB	PreSamplB
	EMB1, EMB2, EMB3
EMEC	PreSamplE
	EME1, EME2, EME3
HEC	HEC0, HEC1, HEC2, HEC3
Tile	TileBar0, TileBar1, TileBar2
	TileGap1, TileGap2, TileGap3
	TileExt0, TileExt1, TileExt2
FCal	FCalEM
	FCalHad1, FCalHad2
Muons	

Table 4.1: The 25 samplings (24 calorimeter sampling and one for muons) maintained in the EF $E_{\mathrm{T}}^{\mathrm{miss}}$ Fex.

- `EFMissingETFlags("TheFlagsTool")`
- `EFMissingETFromHelper("TheHelperTool")`

`TheMuonTool` calculates the muon contribution, `TheCellTool` makes a loop over all calorimeter cells and calculates the calorimeter quantities, `TheFlagTool` makes quality checks and sets flags, and `TheHelperTool` combines the information from the previous tools and translates the result into a persistent `TrigMissingET` object. During the running of this tool sequence, the contributions from sub-detector portions are stored as separate components known as *samplings* in a transient software object called the `TrigEFMissingEtHelper`. Table 4.1 lists all 25 samplings in logical groupings. This mode of running with the four tools described above, although it is the default, is not the only possible configuration. In fact, the $E_{\mathrm{T}}^{\mathrm{miss}}$ slice has implemented the functionality to run three different configurations of the $E_{\mathrm{T}}^{\mathrm{miss}}$ Fex in parallel: the default configuration described above, a configuration including a noise suppression feature, and one which computes the calorimeter based quantities ($\not\!\!E_{\mathrm{T}}$,

ΣE_{T}, etc.) using information from the FEB headers instead of CaloCells[1]. Once the event quantities are calculated, the information from the transient helper object is used to calculate the global $E_{\mathrm{T}}^{\mathrm{miss}}$ quantities.

The structure of the algorithm here is significant: the loop over all cells provided by `TheCellTool` allows for the intermediate discrimination of any *cell-based* quantity. Quality flags set earlier in the chain, noise suppression, or additional error checking may serve as criteria on which to eliminate problematic cells. Noise suppression, for example, is easily applied to each cell, removing the contribution of low-energy CaloCells to the computation of calorimeter-based quantities. Noise suppression is further explained in section 4.3.1.

4.1.1 DK Calibration

Following the Fex loop over all cells, `TheHelperTool` combines the results from individual sub-detector samplings into the final desired quantities such as $\not{E}_{\mathrm{T}}$ and ΣE_{T}. In the same way that a feature such as noise suppression may be applied over all cells, the loop over the samplings of the detector offers the chance to apply a calibration or perform some quality-based monitoring[2] on a sampling-by-sampling basis. It is precisely a calibration which is currently applied by default, known as the Djilkibaev-Konoplich (DK) calibration. Although the calibration is done on a sampling-based level, it utilizes cell-based quantities. In the loop over all cells within `TheCellTool` the quantity

$$s_j = \sum_i \mathrm{sgn}(E_{ji})$$

is computed and stored for each $j \in \{\text{all 24 calorimeter samplings}\}$. The sgn function in this notation returns either 1 for a positive number or -1 for a negative; the sum

[1] Which produces a lower-granularity result in a much shorter time.

[2] While the monitoring which was once implemented is no longer available, in principle it illustrates the access to sampling-level information.

of signs returns the net surplus of cell signs in either a positive or negative direction. Then a set of constants a_j and b_j are constructed such that the calibration for a quantity X (where X could be E_y, ΣE_{T}, ΣE, etc.) is given by

$$X_j = \sum_j (a_j X_j + b_j s_j)$$

Currently, these untuned constants take the forms $b_j = 0 \; \forall \; j$ and

$$a_j = \begin{cases} 1.05 & : j \in \{\text{EM samplings}\} \\ 1.40 & : j \in \{\text{Hadronic samplings}\} \\ 1.00 & : j \in \{\text{Muons}\} \end{cases}$$

This calibration method is desirable not simply because it is tunable, but because it is *unfoldable*. Calibration constants a_j, b_j, and s_j are recorded to disk during runtime and can therefore be unapplied to resulting quantities, changed, and re-applied entirely offline without re-running the HLT. All of the studies which follow employ the DK calibration.

4.2 Data Samples

A combination of Monte Carlo (MC) and Cosmic Ray (CR) samples are used to test the quantities of interest in this study and their parameters are listed in table 4.2. Samples used in this study include a highly energetic dijet sample (J8) generated using Pythia [8] and a fully hadronic $t\bar{t}$ sample generated in Mc@NLO [9]. These generated events based on theoretical models are then processed through a detector simulation[3] [10] which emulates the passage of the particles through the matter of

[3]Detector simulations are only as good as the geometry which is constructed in software. Since this geometric software description improves with an improved understanding of the actual detector, geometry *tags* track the versions as they are developed. For the purpose of this study, the geometry tags are nearly indistinguishable. For completeness the geometry tags are listed in Table 4.2.

the detector, effectively translating the generated particles into energy deposits in hardware.

The JX samples consist of dijet events each with a specific p_{T} range: J0, $p_{\mathrm{T}} <$ 17 GeV; J1, $17 < p_{\mathrm{T}} < 35$ GeV; J6, $560 < p_{\mathrm{T}} < 1120$ GeV; J8, $p_{\mathrm{T}} > 2240$ GeV. The $t\bar{t}$ sample consists of ∼100,000 events undergoing the process $t\bar{t} \rightarrow$ hadrons. Cosmics

Dataset	Type	Athena Version	Generator	Geometry
89955	CR	15.2.0.4	—	ATLAS-GEO-03-00-00
90272	CR	15.2.0.4	—	ATLAS-GEO-03-00-00
CosSim	MC	15.2.0.4	Pythia	ATLAS-GEO-07-00-00
JX	MC	14.1.0/15.2.0.4	Pythia	ATLAS-CSC-05-01-00
TTbar	MC	14.2.25	Mc@NLO	ATLAS-GEO-02-01-00

Table 4.2: Summary table of data samples used in this study.

run 90272 ran on September 28, 2009 for 15 hours and 42 minutes until the run was stopped owing to busy LAr software and the subsequent dropping of trigger rates to zero. Run 89955 ran for four hours and 26 minutes and was stopped owing to a faulty HLT trigger chain. Both were run with functioning solenoid and toroid magnets. In these studies only events passing the L1Calo trigger are included; 349,024 L1Calo events are available in 90272 and 40,819 from 89955.

4.3 Robustness in $E_{\mathrm{T}}^{\mathrm{miss}}$

Theoretical expectations that $\not{E}_{\mathrm{T}}$ should be zero are combatted by intrinsic sources of p_{T} in ATLAS: parton distribution functions intrinsic to the colliding protons carry several hundred MeV of p_{T} [11]; the transverse spread of the beam owing to beam divergence contributes about 500 MeV [12]; and the non-zero crossing angle of the two colliding beams account for about 5 GeV [13]. While these contributions from known detector effects are significant, they do not nearly match the scale of electronics problems and coherent noise to the total *fake* $E_{\mathrm{T}}^{\mathrm{miss}}$ [14]. Failure in a front-end

module, deformed pulse shapes, systematic cell noise, and noise bursts are but a few among many studied problems, all to which the trigger is sensitive. Eliminating such problems is a near impossibility; making simple checks for reasonable values of detector quantities, however, is easy. Using simple checks to diagnose and characterize problems in online running provides the algorithm with handles to study these effects. Once the exploratory stages of running have ended and sufficient information from these diagnostic tools has been collected, the Fex can be redeveloped such that it is less sensitive to common detector problems and provides for handles with strong rejection power over events with symptoms of bad data quality.

Quantities which have close correlations with faulty hardware, software problems, or a host of possible detector effects may be utilized in the $E_{\mathrm{T}}^{\mathrm{miss}}$ trigger slice in any one of three levels: the cell level, the sampling level, and the event level. Cell level checks, while offering the finest granularity and lowest level of information, may be computationally cumbersome since the ATLAS calorimeter contains in excess of 180,000 individual channels. Sampling-level checks are much faster and may be more sophisticated since each of the final energy quantities for each sampling may be combined or compared with other samplings to form useful discriminants; however, the diagnostic value of the sampling-level checks is limited in its precision. Finally, event-level checks are meant only to give the end user hints as to any peculiar behaviour which occurs in a given event; once the hint is dropped, offline analysis may clarify the source of the problem and suggest implementations which better identify the error encountered in online running. Checks between these three levels are intended to make the algorithms more robust against common problems and to provide the framework for event rejection at the trigger level should this be necessary.

4.3.1 Electronics Noise

At the EF level of the trigger, electronics noise in the detector is expected to be Gaussian with a mean of zero and varies in magnitude by sampling. Figure 4.2 shows the average noise values by sampling in the EMB, Tile, FCal and HEC. MC samples listed in Table 4.2 follow a noise distribution which is approximately the same. For the purpose of triggering, electronics noise is utterly uninteresting and, unfortunately, perhaps occasionally misleading. As has been stated, the noise is distributed along a Gaussian and may therefore occasionally take on very high or very low values within the distribution. This is allowed, but it can be misleading for a trigger algorithm which is sensitive to the individual energies of cells. Such is the concern for candidate robustness checks in the $E_{\mathrm{T}}^{\mathrm{miss}}$ Fex whose decisions rely upon the cell energies themselves. For this reason, suppressing noise – that is, ignoring cells with low signal-to-noise ratios – becomes an ideal method for reducing the presence of uninteresting cell energies.

Figure 4.1 shows the effect of applying noise suppression on cosmics run 90272 using $2\sigma_{\mathrm{noise}}$ noise suppression[4]. There are some noteworthy features here: the mean energy in the noise-suppressed plot is positive while the mean is approximately zero for the unsuppressed plot. This is owing to the one-sided suppression in which only cells *above* a $2\sigma_{\mathrm{noise}}$ threshold are kept. One might choose a two-sided cut which excludes all cells below a certain magnitude, thereby eliminating cells which satisfy $|E\sigma_{\mathrm{noise}}| < 2$, where E is the energy of a given cell. Such a two-sided cut on a sample containing only noise would have a mean energy of approximately zero; the noise suppression is said to be symmetric. One might think that the energy bias introduced through the asymmetric cut is problematic; the next section shows how

[4]The quantity σ_{noise} refers to the standard deviation of the Normal along which noise values in a particular sampling are distributed.

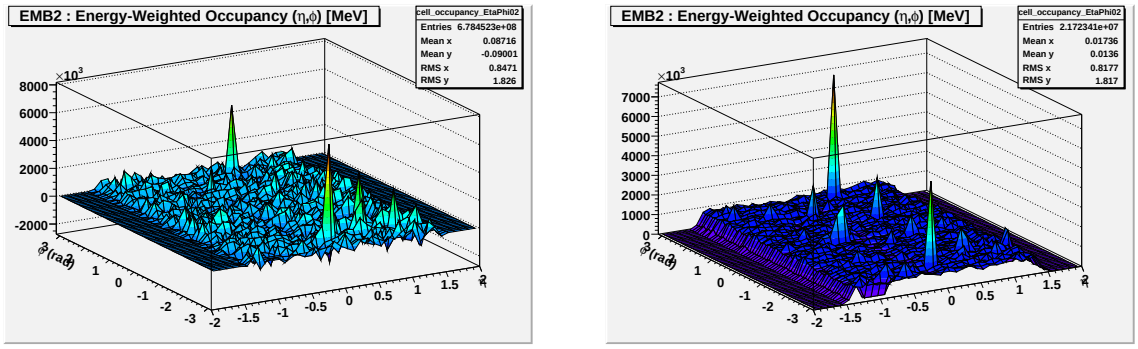

(a) without 2σ suppression

(b) with 2σ suppression

Figure 4.1: Energy distribution in (η, ϕ) of EMB2 for cosmics run 90272 before and after a one-sided 2σ noise cut on all cells.

the bias can simply be subtracted.

The tallest peaks in both plots of Figure 4.1 are strong candidates for hot channels[5] while the smaller peaks in the noise-suppressed plot are good candidates for cosmic ray events. Without noise suppression, these same peaks are washed away over time by noise as illustrated in the unsuppressed plot.

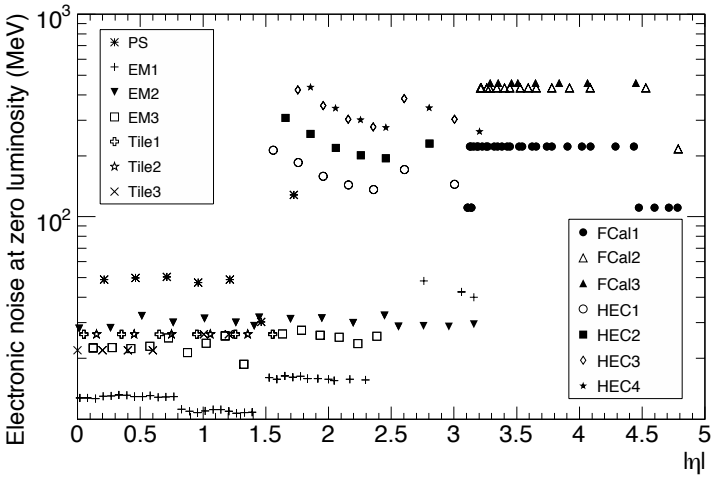

Figure 4.2: The average electronics noise by sampling measured from cosmics runs in 2008. Diagram taken from [2].

[5]If one investigates the surrounding samplings corresponding to the region around the peaks, it is found that there are no corresponding energy deposits. Hence it is unlikely that a cosmic, which is most likely to deposit energy in multiple samplings, is the source of this isolated signature.

In the studies that follow, noise suppression may come in two forms: a cut on the energy value of a cell and a cut on the Signal-to-Noise-Ratio (SNR) of the cell. The noise value of a cell is retrieved from a database based on the cell's gain. Since noise fluctuations give rise to both positive and negative energies, the question of two- and one-sided[6] forms of these cuts makes room for debate. Large negative energy deposits are unphysical and, if they are ignored owing to a one-sided cut, the algorithm may be ignoring a problem inherent in the event. By the same token, making no cuts on cell energy values allows for statistical electronics noise to dominate; energy sums may assume negative values and cause certain fractions to take unphysical values. While offline analysis has the luxury of rerunning the trigger with any noise suppression as often as is desired, the online trigger has only one opportunity; identification of suspicious behaviour in online running may be greatly facilitated by the application of noise suppression in some form.

Making a unilateral energy cut across all cells seems unbalanced: Figure 4.2 shows that a low 11 MeV in the EMB is contrasted by the ~ 500 MeV noise value in the HEC. An algorithm which makes use of one value for a cut across the cells of all samplings biases the resulting distributions: real energy deposits in the barrel disappear while noise in the HEC remains. Figure 4.3 shows the fraction of cells in each sampling which survive two-sided energy cuts against the number of cells which survive a two-sided SNR cut. The SNR cut is proportional to the noise level of the sampling, offering a more democratic distribution of cells for advanced checks. Both cuts allow for the passage of high energy cells which are considered of greatest interest for robustness checks.

For cell-level checks, a one-sided cut seems a poor choice since the goal is to

[6]One-sided cuts are understood to exclude all cells *below* a certain threshold; cells with negative cut quantities are therefore eliminated. Two-sided cuts exclude all cells inside a symmetric cut centered at zero.

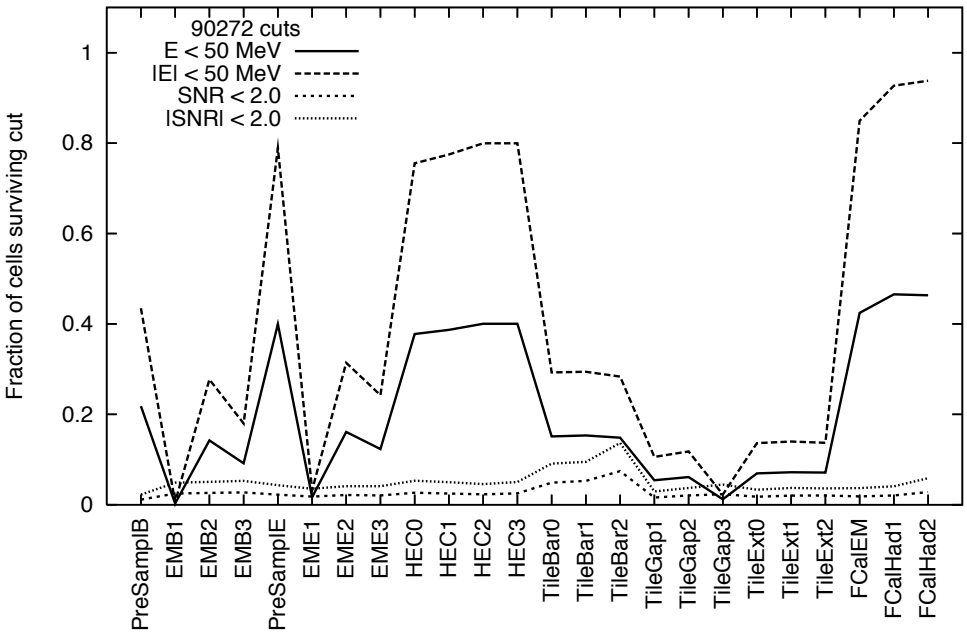

Figure 4.3: Applying one- and two-sided cuts on cosmics run 90272. The cut on the SNR offers a more democratic distribution without discarding high energy cells of interest.

identify problems based solely on the properties of the one cell; ignoring cells with large negative energy, for instance, eliminates the chance to identify these cells as problematic. Sampling- and event-level checks are concerned with the aggregate quantities formed from the cell energies and therefore aim to flag problems based on spurious values of these aggregates. If the presence of unmasked noise distributions form aggregate quantities which are often unphysical, however, then these checks lose their discrimination power. For these reasons and the arguments that follow, one-sided noise suppression on the SNR is applied over all cells when computing these aggregate quantities. For the simplest robustness checks on cells themselves, a two-sided cut allows for the detection of very large negative energy values. The thresholds for these cuts have some flexibility; since $1\sigma_{\mathrm{noise}}$ is considered too inclusive, $2\sigma_{\mathrm{noise}}$ serves as the lowest threshold for construction of aggregate quantities while stronger

suppressions are explored.

4.3.2 Cell-Level Checks

Making cell-level checks through calibration runs and systematic hardware studies is complimented by online data quality monitoring; some implementation must be made in $E_{\mathrm{T}}^{\mathrm{miss}}$ which gives online shifters the ability to diagnose subtle but influential problems. From the time at which a signal is generated in a calorimeter cell to its amplification, shaping, digitization, and eventual retrieval by the EF, many problems have the opportunity to occur. Cell-level checks which are simple and fast are an elementary level of safety which also provide the finest granularity of information for troubleshooting.

Knowing that performing cell-level checks becomes computationally cumbersome owing to the number of cells in the calorimeter, making checks only on *significant* cells becomes a natural goal: that is, making quality checks on cells which contribute large terms to the energy sums. In the $E_{\mathrm{T}}^{\mathrm{miss}}$ slice, such a cell quantity might take the form of E_x, E_{T}, or even the cell energy E or a noise value σ_{noise}. This approach is consistent with the early rejection philosophy of the trigger, offering advanced quality checks only after determining that they may be important. The first approach to deciding which cells are significant is identifying whether the cell contains a real energy deposit owing to some reaction or simply electronics noise. Choosing the value of a SNR cut to eliminate low-significance cells from robustness checks allows for the tuning of the fraction of cells to be checked. A CR simulation and a low-energy dijet sample are tested for various cut values in Figure 4.4. The resolution of this measurement is sensitive the number of cells in the sampling; refer to Table 2.1 for the number of channels in each sampling. If the number of cells passing the cut is too high and leads to an overburden of computations, the SNR cut ought to be increased. Therefore, let

$2\sigma_{\mathrm{noise}}$ – corresponding to a survival rate of about 2.25% of cells – be selected with the understanding that the value is tunable in order to suit time constraints. When such a cut is applied, the energy distribution of the noise artificially increases. This bias can be trivially removed by subtracting the new mean of the noise distribution: with one-sided $2\sigma_{\mathrm{noise}}$ suppression, 2.25% of cells survive the cut and their expectation value for energy is given by $2.37\sigma_{\mathrm{noise}}$ so that, for the EM calorimeter, the bias should be roughly

$$180000\,\mathrm{cells} \times 2.25\% \times 20\,\mathrm{MeV} \times 2.37 \approx 210\,\mathrm{GeV}.$$

The bias is not problematic – it is a symptom of the asymmetric energy cut. What is relevant is the proportion of the total energy made up by the noise contribution.

Cell Energy Saturation

Since the $E_{\mathrm{T}}^{\mathrm{miss}}$ Fex constructs an energy 'vector' based on the net energy in the x-y plane, a single large energy deposit in a cell could easily shift the direction of the alleged missing energy. In the case of a faulty cell with a large energy value – positive or negative – the rate of fake $E_{\mathrm{T}}^{\mathrm{miss}}$ is certain to spike even for chains with a large transverse energy requirement. For this reason, Chapter 6 proposes flagging cells with SNR $< -6\sigma_{\mathrm{noise}}$. Additionally, cells of the detector have physical limits on the energy they are able to detect. Hence, the implementation of a check for energy-saturated cells is not unreasonable.

Although constructing a useful set of thresholds requires the input of real beam data and a comprehensive understanding of the ATLAS detector, a naïve model should well illustrate the aim of such a cut. First, a set of maximum expected energy values should be constructed from representative MC samples. In this example, a J8 sample is selected for its high energy EM depositions, a $t\bar{t}$ sample for its high energy hadronic depositions and an SU4 sample. All samples are processed and

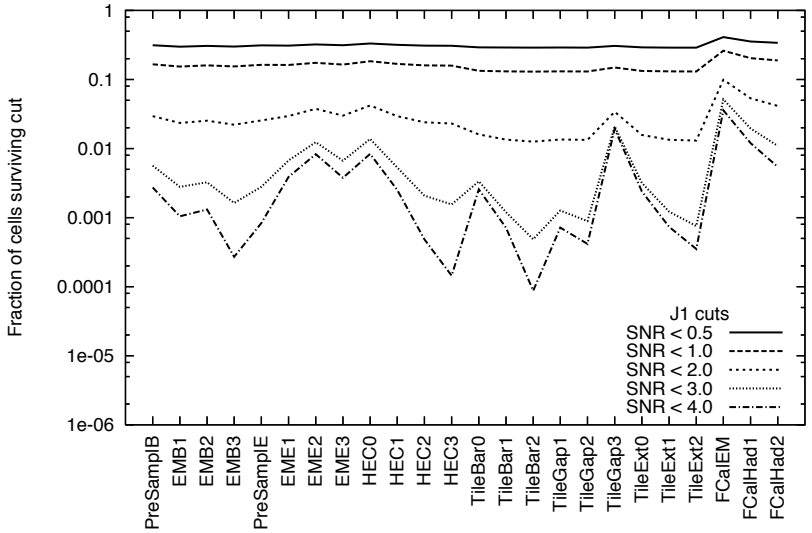

(a) J1

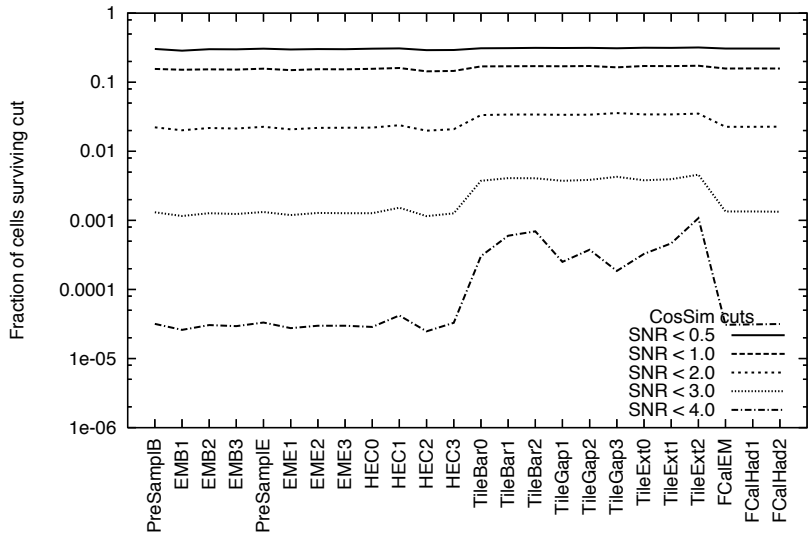

(b) CosSim

Figure 4.4: Fraction of cells surviving SNR cuts for J1 and CR simulation samples. Since a flat distribution is expected for a sample containing only noise, the shapes of the distributions are characteristic of the data: note the favouring of cells in the the tile calorimeter of the CosSim in which most of the energy is deposited.

a representative 'approximate maximum' is selected for each sampling. As default values before tuning, the minimum of two times the value given in Table 4.3 and the three TeV maximum from the readout electronics [2] should serve as a reasonable

	J8	$t\bar{t}$	SU4		J8	$t\bar{t}$	SU4
PreSamplB	120	70	80	TileBar0	1800	150	200
EMB1	240	70	85	TileBar1	2400	180	200
EMB2	2100	250	350	TileBar2	1000	85	80
EMB3	450	80	90	TileGap1	300	40	50
PreSamplE	65	70	60	TileGap2	340	45	60
EME1	120	120	100	TileGap3	320	130	120
EME2	1000	1100	500	TileExt0	360	200	220
EME3	120	180	150	TileExt1	1100	130	150
HEC0	300	400	320	TileExt2	1350	150	120
HEC1	380	420	330	FCalEM	1000	1200	1100
HEC2	180	250	200	FCalHad1	380	800	500
HEC3	120	140	120	FCalHad2	220	300	180

Table 4.3: 'Approximate maximum' cell energies for each sub-detector sampling are given in GeV for three MC samples: J8, SU4, and $t\bar{t}$.

benchmark.

It is found that applying the cuts in Table 4.3 produces no difference in a cosmics run (89955, in this case) as expected. The check for cell saturation is a basic one, and its implementation for a conservative number of cells is computationally negligible, as discussed in Section 5.1.

4.3.3 Sampling-Level Checks

Sampling-level checks have access to the summary information composed from groups of cells, offering the ability to compare any single component of energy, $\not\!\!E_{\mathrm{T}}$, or the scalar ΣE_{T} between two or more samplings. These checks need only be done once per event so they consume very little computational time and may therefore be more complicated.

Transverse Energy to Sum Transverse Energy

The ratio between the transverse energy in a sampling and the cell-based sum of all the transverse contributions loosely measures the significance of the energy imbalance

in the transverse plane. While the ratio usually small, this is not necessary in order to be considered physical: consider the usual case in which the event's overall E_{T} balance is achieved by contributions from a series of different samplings. In principle,

$$\left(\frac{E_{\mathrm{T}}}{\Sigma E_{\mathrm{T}}}\right)_{\mathrm{sampling}} \in [0,1],$$

but negative energy values spread the ratio beyond this interval. In the presence of statistical noise which does not dominate the signal, the distribution is smeared at its limits. When statistical noise does dominate the signal, any ratio may be produced. So too is the case for large, unbalanced, negative energy values: ratios much greater than one or much less than zero are not uncommon. Figure 4.5a shows such a distribution for cosmics run 90272 and two dijet samples without noise suppression of any kind.

When noise suppression is applied, a high ratio indicates that the magnitude of the transverse energy is on the same scale as all of the energy deposited in this sampling. Cosmics typically demonstrate high values for this ratio since their origin is not from the nominal interaction point and, hence, their energy deposition occurs along a non-projective path. Figure 4.6 illustrates an event in which this is not the case. Applying one-sided noise suppression at the cell-level eliminates values beyond the interval $[0,1]$ and prevents the artificial deflation of the denominator owing to negative cell noise. While it can be argued that making a one-sided cut discards 'interesting' events with large negative energy, it should be noted that large negative energy deposits *are* unphysical and should be treated at an earlier time such as in cell-level robustness checks. Note the progression toward higher values of the ratio as harder noise suppression is applied. Eliminating additional positive noise from the sample will always reduce ΣE_{T} while E_{T} will tend to increase on average since there

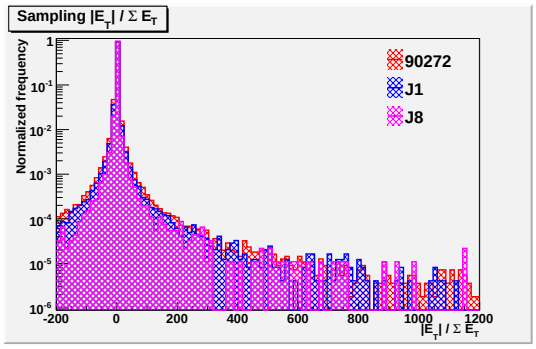

(a) without suppression

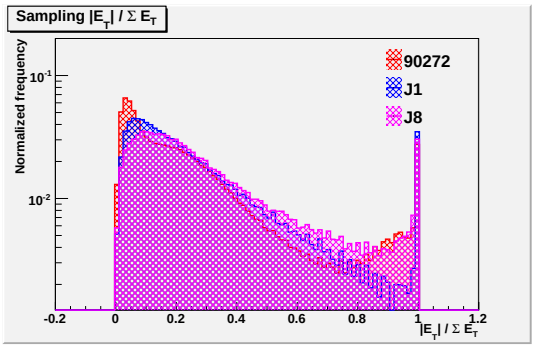

(b) with $2\sigma_{\mathrm{noise}}$ noise suppression

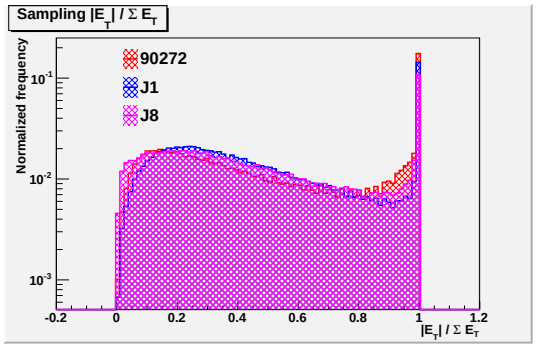

(c) with $3\sigma_{\mathrm{noise}}$ noise suppression

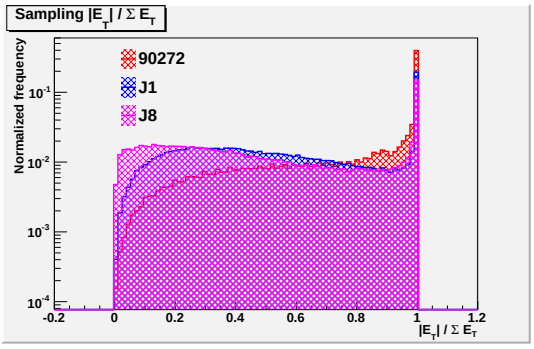

(d) with $4\sigma_{\mathrm{noise}}$ noise suppression

Figure 4.5: The ratio $E_{\mathrm{T}}/\Sigma E_{\mathrm{T}}$ filled for each sampling in each event of cosmics run 90272 and dijet samples J1 and J8. Plots are made for each of two, three, and four σ_{noise} one-sided noise suppression.

are fewer energy contributions likely to cancel out the signal[7].

Transverse energy as a robustness check is particularly relevant for $E_{\mathrm{T}}^{\mathrm{miss}}$ since high ratios naturally correspond to interesting events in this trigger slice. However, as a sampling-level check, this quantity proves to be not very useful. Figure 4.5 shows that the distributions for this sampling-level ratio between cosmics and dijet MC are barely distinguishable; therefore, although this ratio is simple and computationally inexpensive, nothing can be said at this time merely by identifying a threshold on this quantity. Perhaps this quantity may be sensitive to some telltale symptom of

[7]Though it is assumed that the energy deposits arising from signal are greater than those of the noise, this need not be the case. Low-energy interactions which still reach the calorimetry may well be excluded along with the noise for a given noise suppression.

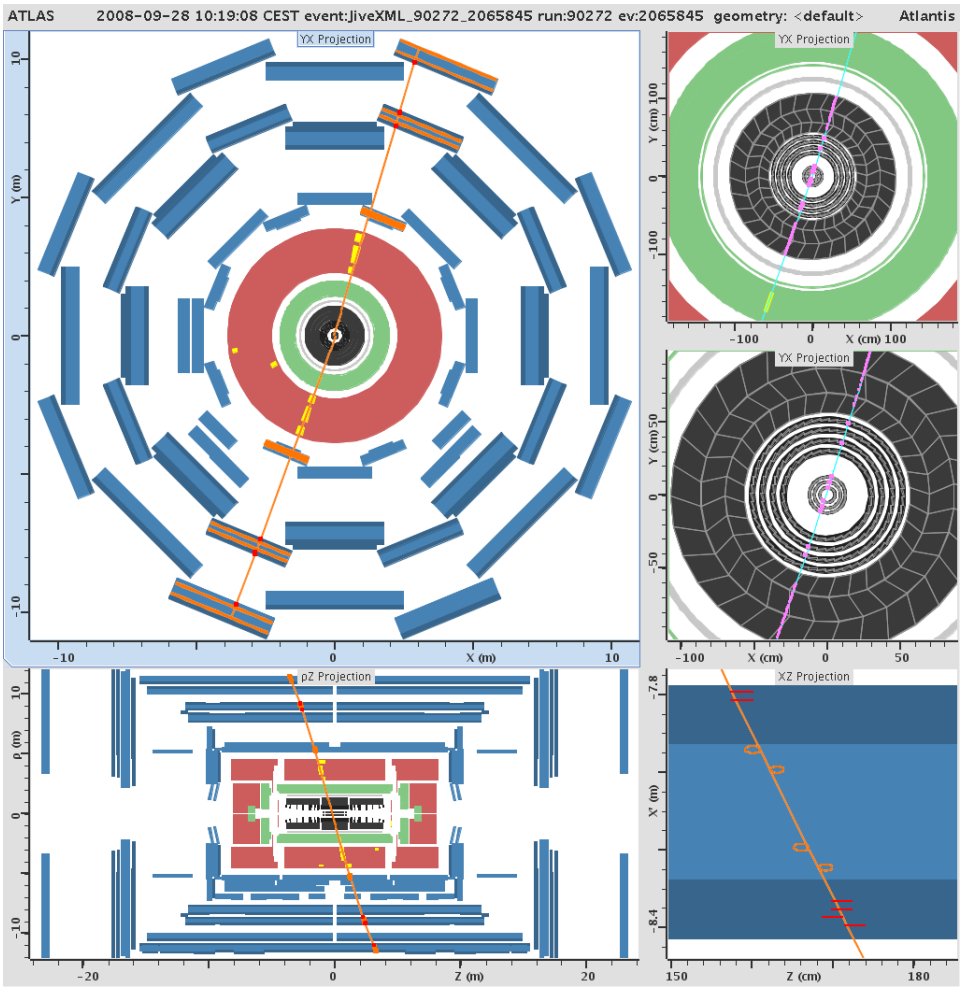

Figure 4.6: Cut-away views of a single event in cosmics run 90272 displayed by Atlantis [15]. An instance of a cosmic with a projective path.

bad data quality, but this has yet to be seen. Section 4.3.4 shows that this ratio is very successful when considered over an entire event.

Longitudinal Energy Profiles

Layers within a sub-detector have energy depositions which are dependent on the thickness of the layer. Figure 2.3 illustrates the existence of separate layers of varying thickness in the EMB. It is not unreasonable to expect that, among the set of all possible energy distributions, certain longitudinal profiles must be unphysical. For

example, suppose that a large energy deposit is found in the first layer of the EMB. If in the same event the second layer of the EMB and the Presampler contain an insignificant amount of energy, then there exists reason to suspect that a malfunction (either software or hardware) has occurred in the first layer. Since such quantities are available and readily accessible from within the EF, simple checks on an event-by-event basis seem reasonable.

The sample from which the energy fractions are constructed intimately determines the distribution; only real beam data is able to determine which thresholds discriminate against problems in the detector. However, a study of the sensitivity of this quantity to various samples is instructive. The quantity studied is

$$F_{\mathrm{EMB}X} = \frac{E_{\mathrm{EMB}X}}{\sum_i E_{\mathrm{EMB}i}}$$

and describes the fraction of total energy deposited in the EMB (given by $\sum_i E_{\mathrm{EMB}i}$) deposited in layer X (with $E_{\mathrm{EMB}X}$ giving the energy in the layer of choice). Figure 4.7 shows the results for this study over samples of $t\bar{t}$ and cosmics run 90272. It is found that energetic dijet samples provide the narrowest distributions per layer in the EMB, though beam data is likely to resemble some admixture with a noisy distribution such as the cosmics run. Applying one-sided noise suppression as in Figure 4.7f narrows the cosmics distribution, from which reasonable thresholds for $F_{\mathrm{EMB}X}$ (given a sample containing mostly noise) can be read. As the noise suppression is increased, the four distinct noise peaks in cosmics run 90272 are eliminated and the cosmics signal surfaces. Unlike the $t\bar{t}$ and dijet samples, the EMB layers are not clearly distinguishable in run 90272.

Given $2\sigma_{\mathrm{noise}}$ noise suppression, fractions above 0.97 or below 0.1 for EMB2 and above 0.6 for EMB1, EMB3, and PreSamplB, if at the appropriate energy scale, should be investigated. Selecting appropriate energy scales ultimately requires the

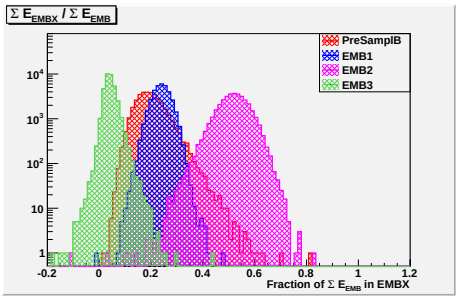

(a) $t\bar{t}$, no noise suppression

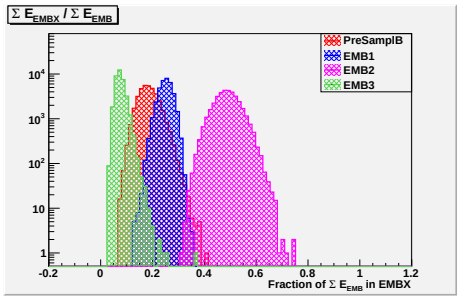

(b) $t\bar{t}$, $2\sigma_{\mathrm{noise}}$ noise suppression

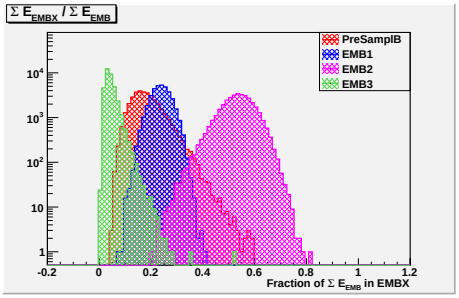

(c) $t\bar{t}$, $3\sigma_{\mathrm{noise}}$ noise suppression

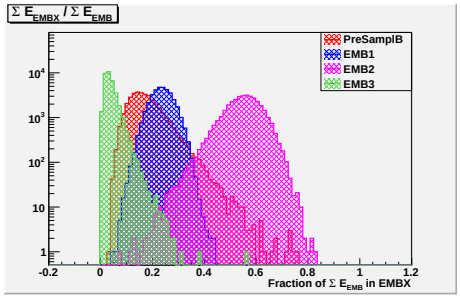

(d) $t\bar{t}$, $4\sigma_{\mathrm{noise}}$ noise suppression

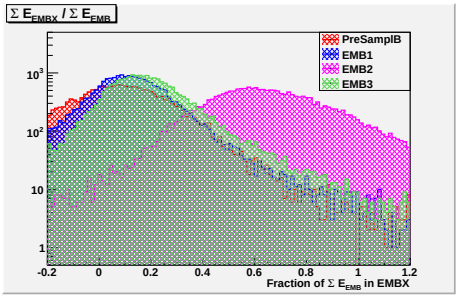

(e) 90272, no noise suppression

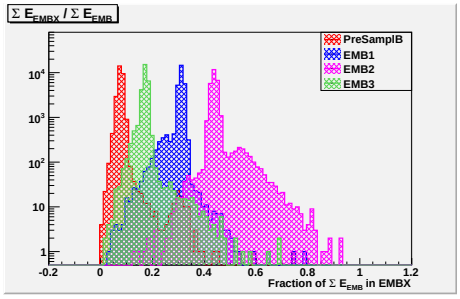

(f) 90272, 2σ noise suppression

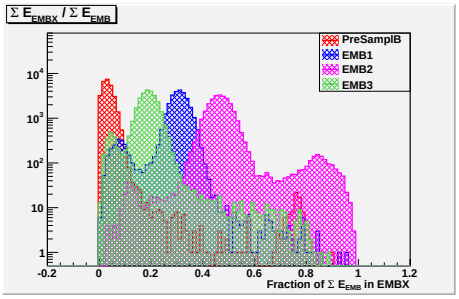

(g) 90272, 3σ noise suppression

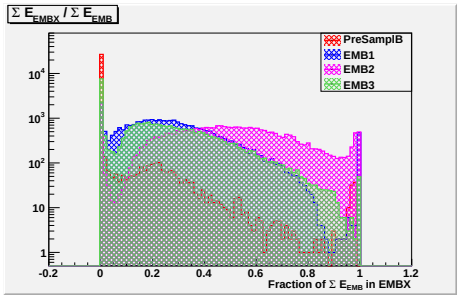

(h) 90272, 4σ noise suppression

Figure 4.7: Fraction of the energy in the EMB which falls into each layer of the EMB (including the presampler). The $t\bar{t}$ sample is not sensitive to the noise suppression since its real energy deposits fall far above the energy scale of the noise. J1, J6, and J8 samples (not pictured here) are also very insensitive to these cuts.

input of beam data, but a simple low estimate can be given by the contribution from cells with positive noise in the smallest sample. With $2\sigma_{\mathrm{noise}}$ suppression, 2.25% of cells survive the cut and their expectation value for energy is given by $2.37\sigma_{\mathrm{noise}}$ so that, for EMB2, given an average noise of 30 MeV, this estimated low-energy scale should begin around 59 GeV.

$$33000\,\mathrm{cells} \times 2.25\% \times 30\,\mathrm{MeV} \times 2.37 \approx 59\,\mathrm{GeV}$$

Compare this with the scale of the suppressed noise contribution from all EM calorimeter cells, on the order of 200 GeV.

4.3.4 Event-Level Checks

Missing Transverse Energy to Sum Transverse Energy

In the same way that the ratio $E_{\mathrm{T}}/\Sigma E_{\mathrm{T}}$ is considered in Section 4.3.3 for sampling-level checks, it is possible to construct a check on the detector as a whole. While a single sampling may well have a transverse energy on the order of the total energy in that sampling, the same is not the case for the detector as a whole. As the contributions to ΣE_{T} grow and the contributions to $\not{E}_{\mathrm{T}}$ average out (that is, shrink) the ratio is driven well under unity. It is for this reason that a large ratio is particularly striking.

Large ratios, however, are not sufficient to qualify an event as suspicious. Considering figure 4.8a, for example, an algorithm seeking only high $\not{E}_{\mathrm{T}}/\Sigma E_{\mathrm{T}}$ ratios selects the top few events at the tail around 0.6. Looking to the corresponding plot which breaks the ratio into its constituents (Figure 4.8b) it is clear that the energy scale is very large and, indeed, the events corresponding to the large ratios contribute significantly to $\not{E}_{\mathrm{T}}$. Now, for a lower energy J0 sample pictured in Figure 4.8c, the algorithm again selects the top few events corresponding to large ratios. In this case,

however, Figure 4.8d shows that the ratio is high but the event's contribution to $\not\!\!E_{\mathrm{T}}$ is not significant; since beam data will be a combination of many processes whose contributions to $\not\!\!E_{\mathrm{T}}$ and ΣE_{T} will vary depending on the energy scale of the interaction, it is valuable to first place a minimum requirement on the size of the $\not\!\!E_{\mathrm{T}}$. Whether or not a minimum threshold is necessary, Figure 4.9b demonstrates that a large fraction of noise-dominated events from cosmic ray sample 90272 illustrate distinctly different behaviour and may be cut away from dijet samples J1 and J8 with a simple and very low $E_{\mathrm{T}}/\Sigma E_{\mathrm{T}}$ threshold. Figure 4.9a highlights the strong correlation found in the cosmics sample which is not present in the J0 or J6 data and shows an approximate 200 GeV bias in the ΣE_{T} consistent with $2\sigma_{\mathrm{noise}}$ noise suppression.

Electromagnetic Fraction

The fraction of the total energy deposited in the EM calorimetry is given by

$$F_{\mathrm{EM}} = \frac{E_{\mathrm{EM}}}{E_{\mathrm{EM}} + E_{\mathrm{HAD}}}.$$

Components in the EM calorimetry include the EMB, EMEC, and the one EM component of the forward calorimeters, FCalEM. Events with F_{EM} near zero are associated with energy deposition arising from energetic cosmic muons since they typically leave very little energy in the EM portions of the detector. In the less-frequent case that a cosmic muon produces bremsstrahlung radiation in the EM calorimeters, the F_{EM} jumps to nearly one; cuts on very high and very low values of the F_{EM} are employed to clean data offline, removing obvious sources of fake $E_{\mathrm{T}}^{\mathrm{miss}}$. For the purposes of robustness, the goal is to begin by finding physical limits on F_{EM} and then studying samples with bad data quality in an attempt to identify a relation between F_{EM} and the source of the poor data quality.

Figure 4.10 overlays the F_{EM} for three samples before and after applying $2\sigma_{\mathrm{noise}}$

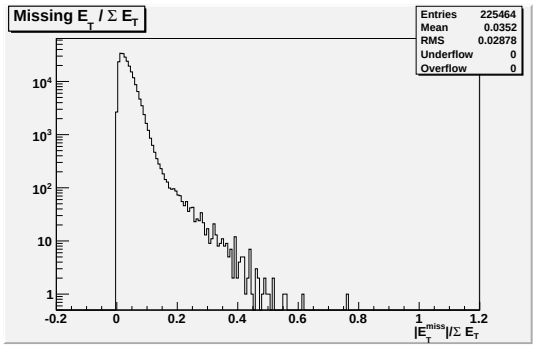

(a) J6

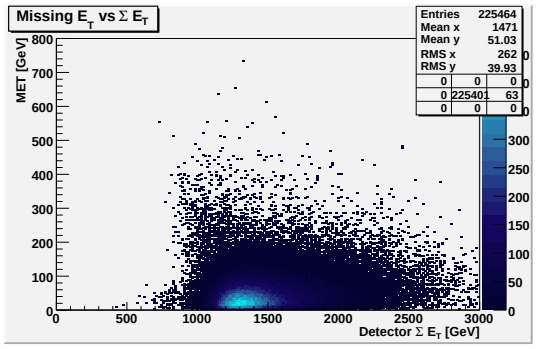

(b) J6

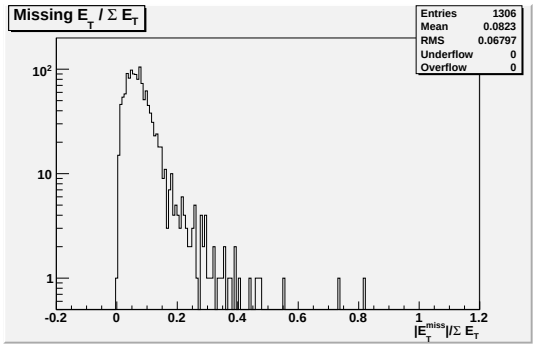

(c) J0

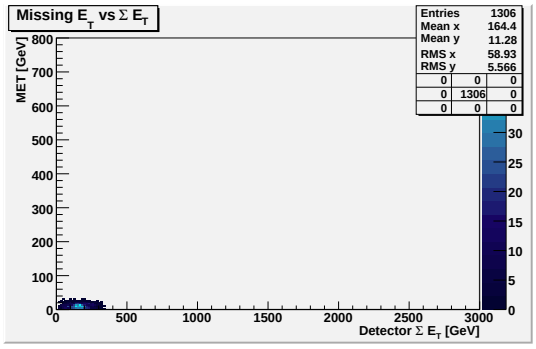

(d) J0

Figure 4.8: Comparison of the ratio $\not{E}_{\mathrm{T}}/\Sigma E_{\mathrm{T}}$ to its components for dijet samples J6 and J0, respectively. While the distributions for the ratio are similar, their constituents are dissimilar. A minimum energy requirement accompanying the ratio would ensure the flagging of only events which contribute significantly to $\not{E}_{\mathrm{T}}$.

noise suppression: cosmics run 90272 and two dijet samples. The soft J1 sample maintains the highest F_{EM} since the low-energy events are unlikely to penetrate and deposit the bulk of the energy beyond the EM calorimeters. The same cannot be said for the J8 sample whose F_{EM} illustrates that the deposition of energy is nearly split between hadronic and EM components. The difference between these and cosmics run 90272, however, is of the greatest importance. It is only in 90272 that the vast majority of the energy falls into hadronic samplings, characteristic of non-projective paths along which energy is deposited. Applying noise suppression eliminates the unphysical values for F_{EM} and allows for the assertion of a cut on F_{EM}: values below

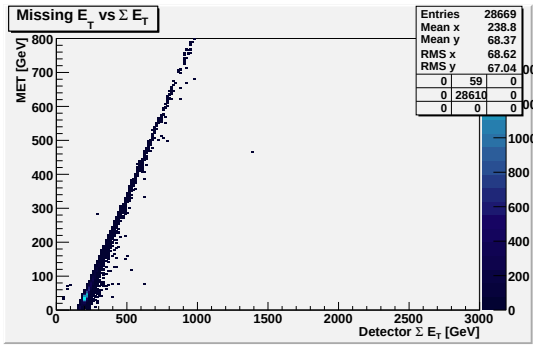

(a) 90272, $2\sigma_{\mathrm{noise}}$ noise suppression

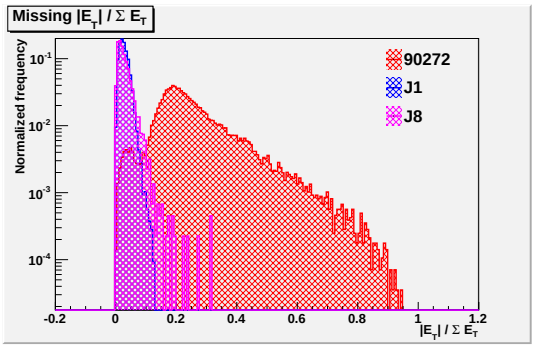

(b) 90272, J1, J8 with $2\sigma_{\mathrm{noise}}$ suppression

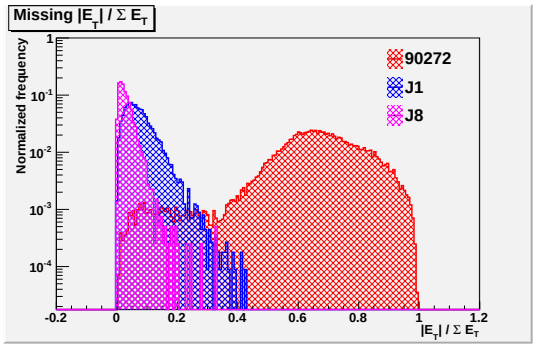

(c) 90272, J1, J8 with $3\sigma_{\mathrm{noise}}$ suppression

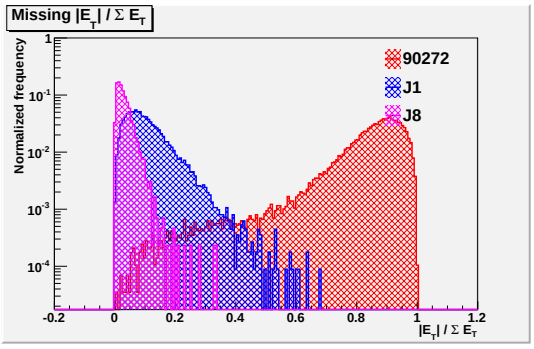

(d) 90272, J1, J8 with $4\sigma_{\mathrm{noise}}$ suppression

Figure 4.9: Run 90272 has a strongly correlated relation between $\not{E}_{\mathrm{T}}$ and ΣE_{T} since cosmics deposit energy along non-projective paths. The remaining plots overlay cosmics run 90272 with dijet samples J1 and J8 under varying noise suppression.

a rough estimate of 0.2 should be considered for investigation. Approximately 0.57% of events in 90272 with $2\sigma_{\mathrm{noise}}$ noise suppression fall within this cut. Table 4.4 lists lower thresholds which keep 99.9% of J8 events. Upper bounds are dependent on the severity of the noise suppression; a $2\sigma_{\mathrm{noise}}$ suppression calls for a maximum around 0.9 while 3 and $4\sigma_{\mathrm{noise}}$ suppression expect ratios up to the physical limit.

TileGap Fraction

The TileGap chambers instrument the space between the barrel and the end-caps in which energy is lost to inactive material. Their relatively small size (see Figure 2.8) suggests a proportionately small fraction of energy deposition and it would signal

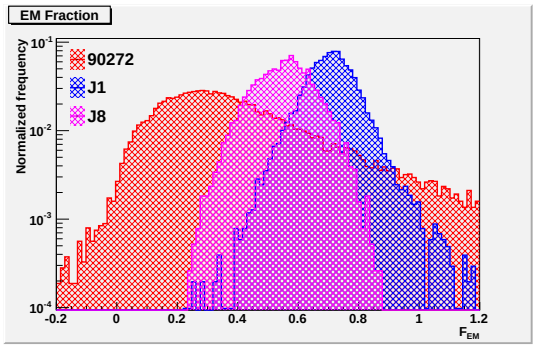

(a) without noise suppression

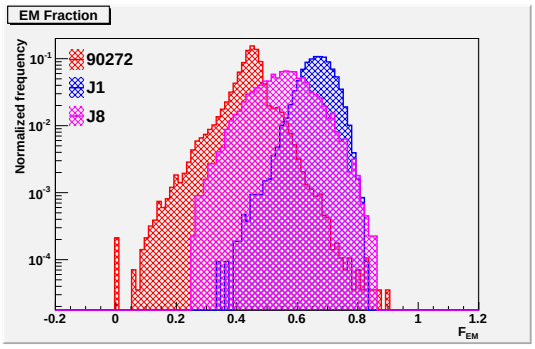

(b) $2\sigma_{\mathrm{noise}}$ noise suppression

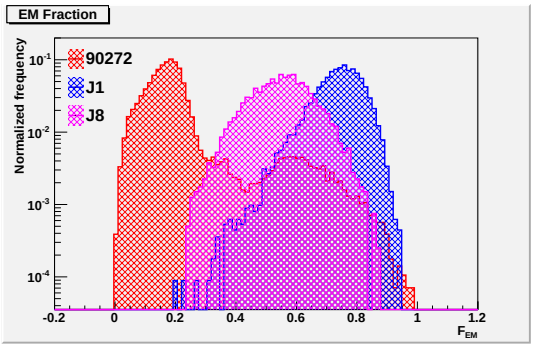

(c) $3\sigma_{\mathrm{noise}}$ noise suppression

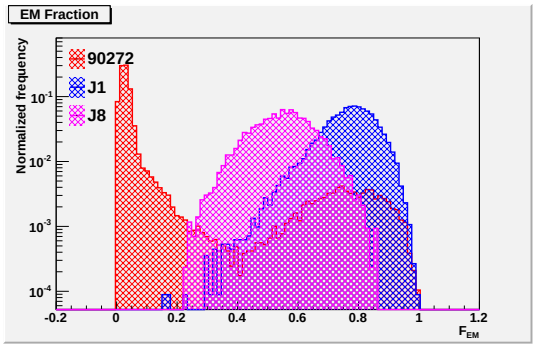

(d) $4\sigma_{\mathrm{noise}}$ noise suppression

Figure 4.10: Comparison of F_{EM} before and after 2, 3, and $4\sigma_{\mathrm{noise}}$ noise suppression for cosmics run 90272 and dijet samples J1 and J8. Although not pictured, $t\bar{t}$ samples exhibit nearly the same behaviour as the J1 data but with a less sensitive upper value for F_{EM} which tends to remain around 0.9.

some suspicion if a large portion of the total calorimeter energy were found within them. This instrumentation is also very limited, leading to a larger uncertainty in the measured energy; if a large fraction of the event energy is discovered here, the uncertainty alone may qualify the event worthy of flagging. For this reason the

suppression	F_{EM} cut	J8 events kept	90272 events kept
none	0.264	99.90%	66.01%
$2\sigma_{\mathrm{noise}}$	0.274	99.91%	97.76%
$3\sigma_{\mathrm{noise}}$	0.250	99.90%	14.43%
$4\sigma_{\mathrm{noise}}$	0.243	99.91%	8.89%

Table 4.4: One-sided F_{EM} cut efficiencies for cosmics run 90272 given a J8 efficiency of 99.9% taken over various noise suppressions.

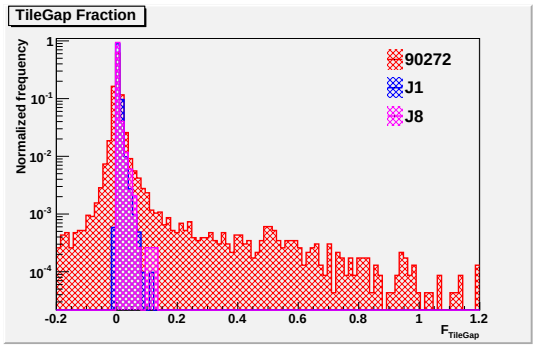

(a) without noise suppression

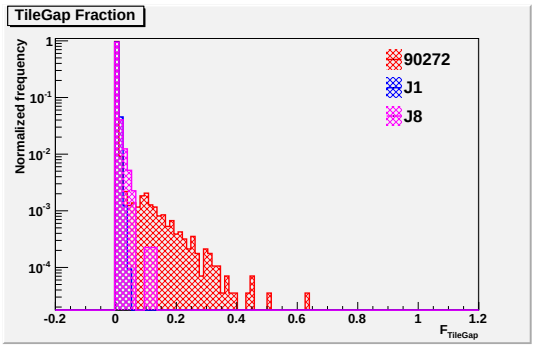

(b) $2\sigma_{\mathrm{noise}}$ noise suppression

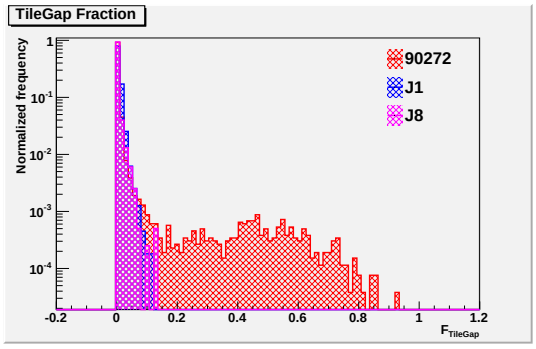

(c) $3\sigma_{\mathrm{noise}}$ noise suppression

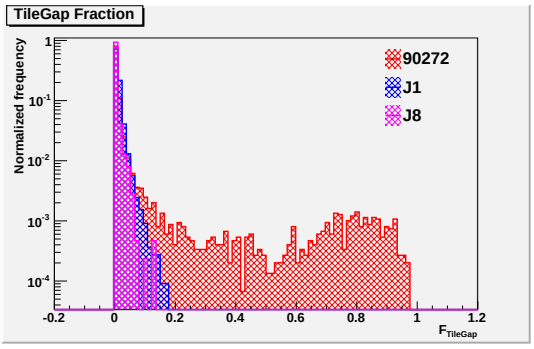

(d) $4\sigma_{\mathrm{noise}}$ noise suppression

Figure 4.11: F_{TG} for cosmics run 90272 and two dijet samples, before and after various noise suppressions. High values are suspect since energy deposition in the TileGap calorimeters is unlikely to compete with the total calorimeter energy.

TileGap fraction,

$$F_{\mathrm{TG}} = \frac{E_{\mathrm{TileGap}}}{\Sigma E},$$

becomes a good measure for events with unusually high TileGap energy deposits. Cosmics runs have an atypically high F_{TG} distribution owing to the reduced overall energy sum found in the EM components and a disproportionate exposure to cosmic muons; this effect is evident especially in Figure 4.4b in which the fraction of cells surviving cuts in the TileGap exceeds the fraction in the EMB, EMEC, HEC, and FCal. Figure 4.11 shows the F_{TG} for cosmics run 90272 and two dijet samples. Applying $2\sigma_{\mathrm{noise}}$ noise suppression over all cells eliminates the unphysical values of the fraction, demonstrating reasonable limits for F_{TG}. Requiring $F_{\mathrm{TG}} < 0.2$ safely

contains the signal from the dijet samples. Increasing severity of noise suppression demonstrates that signals from J1 and J8 (and $t\bar{t}$, though not pictured here) lead to consistent F_{TG} values while noise-sensitive samples such as cosmics run 90272 transform dramatically, taking on increasing values of F_{TG}.

Chapter 5

Implementation

Robustness checks in the $E_{\mathrm{T}}^{\mathrm{miss}}$ slice are still in their early stages as developers get a sense of what is computationally feasible and what checks are effective. Therefore, before a sufficient understanding of the detector is obtained (and, hence, before rejection of cells or events in the EF $E_{\mathrm{T}}^{\mathrm{miss}}$ code can take place) rejection manifests itself in the setting of designated bits in a '`status`' bit string. Each of the 24 calorimeter samplings listed in Table 4.1 has its own set of 16 status bits while the event as a whole has 32 bits: when the event checks are finished, the first 16 bits of the event bit string are set by performing a logical OR over all sampling `status` flags. The `status` is recorded to permanent storage along with the rest of the event information and can therefore be analyzed offline along with all relevant detector quantities. Bits in these strings are shared between L2 and the EF, so only the EF-relevant flags among all 32 bits are shown in Table 5.1.

At the cell-level a parameter `makeRobustness`, which is true by default, toggles the cell-based checks. With this parameter, cells are provided additional quality checks if the absolute value of their SNR is above a certain threshold. This threshold is configurable for each sampling and stored in an array, `MinCellSNratio[24]`, with default values of 2.0 for all samplings. Additionally, noise suppression may be applied

leading to the rejection of all cells under a certain SNR. This behaviour is disabled by default. In code, the checks proceed as follows[1]:

```
if (m_makeRobustness &&
    fabs(E) > m_MinCellSNratio[cDDE->getSampling()] *
    m_noiseTool->getNoise(*m_it, ICalorimeterNoiseTool::TOTALNOISE)
   ) {
  float time = (*m_it)->time() * 1e-3;  // ns
  float quality = (*m_it)->quality();
  if (time < metComp->m_minTime) metComp->m_minTime = time;
  if (time > metComp->m_maxTime) metComp->m_maxTime = time;
  if (quality > metComp->m_maxQlty) metComp->m_maxQlty = quality;
  if (E < metComp->m_minE) metComp->m_minE = E;
  if (E > metComp->m_maxE) metComp->m_maxE = E;
}
```

First the time and quality measurements for this cell are retrieved, then checks are made which determine if this cell's time, quality, or energy values are maxima or minima for this sampling in this event. If so, the value is recorded as that sampling's

[1]This implementation is recent as of the time of writing, July 16, 2009.

Bit	Name	Level	Description
5	OverflowSumET	cell	E overflow
8	BadCellQuality	cell	cell chi-square too high
9	BadCellEnergy	cell	non-reasonable cell energy
10	BadCellTime	cell	non-reasonable cell time
12	BadEnergyRatio	comp	non-reasonable sampling E fraction
13	BadCompEnergy	comp	non-reasonable sampling SumE
18	BadEMfraction	event	non-reasonable EM energy fraction
19	ObjInCrack	event	jet/track in crack
27	GlobMaxMEtSumEtRatio	event	non-reasonable event $E_{\mathrm{T}}^{\mathrm{miss}}$ ratio

Table 5.1: Meaning of the status bits which are relevant to the EF $E_{\mathrm{T}}^{\mathrm{miss}}$ code. Bits zero through 15 in the event-level bit strings are set according to the OR of bits zero through 15 for each sampling.

Level	Check	Configurables	Bit
comp	$\lvert\not\!\!E_{\mathrm{T}}/\Sigma E_{\mathrm{T}}\rvert > R$	$R =$ `CompMaxMEtSumEtRatio`	13
comp	$\Sigma E < T$	$T =$ `MinCompE`	12
comp	$\Sigma E > T$	$T =$ `MaxCompE`	12
comp	$\min(E_{\mathrm{cell}}) < T$	$T =$ `MinCellEnergy`	9
comp	$\max(E_{\mathrm{cell}}) > T$	$T =$ `MaxCellEnergy`	9
comp	$\min(\mathrm{Time}_{\mathrm{cell}}) < T$	$T =$ `MinCellTime`	10
comp	$\max(\mathrm{Time}_{\mathrm{cell}}) > T$	$T =$ `MaxCellTime`	10
comp	$\max(\mathrm{Quality}_{\mathrm{cell}}) > T$	$T =$ `WorstCellQlty`	8
	$F_{\mathrm{EMB}X} > T,\ \ F_{\mathrm{EMB}X} < U$	$T, U =$ `Max-`, `MinSumEratioInEMB`	11
	$F_{\mathrm{EME}X} > T,\ \ F_{\mathrm{EME}X} < U$	$T, U =$ `Max-`, `MinSumEratioInEME`	11
$\uparrow$	$F_{\mathrm{HEC}X} > T,\ \ F_{\mathrm{HEC}X} < U$	$T, U =$ `Max-`, `MinSumEratioInHEC`	11
event	$F_{\mathrm{TileBar}X} > T,\ \ F_{\mathrm{TileBar}X} < U$	$T, U =$ `Max-`, `MinSumEratioInTileBar`	11
$\downarrow$	$F_{\mathrm{TileGap}X} > T,\ \ F_{\mathrm{TileGap}X} < U$	$T, U =$ `Max-`, `MinSumEratioInTileGap`	11
	$F_{\mathrm{TileExt}X} > T,\ \ F_{\mathrm{TileExt}X} < U$	$T, U =$ `Max-`, `MinSumEratioInTileExt`	11
	$F_{\mathrm{FCal}X} > T,\ \ F_{\mathrm{FCal}X} < U$	$T, U =$ `Max-`, `MinSumEratioInFCal`	11
event	$F_{\mathrm{EM}} > T,\ \ F_{\mathrm{EM}} < U$	$T, U =$ `Max-`, `MinEMfraction`	18
event	$F_{\mathrm{TG}} > T,\ \ F_{\mathrm{TG}} < U$	$T, U =$ `Max-`, `MinTileGapEratio`	19
event	$\lvert\not\!\!E_{\mathrm{T}}/\Sigma E_{\mathrm{T}}\rvert > R$	$R =$ `GlobMaxMEtSumEtRatio`	27

Table 5.2: Checks performed by `TheFlagsTool`, their corresponding configurable parameters, and the appropriate status bits which are set as a result of the check. A level given by 'comp' indicates that the check is done for each of the 24 samplings and, therefore, there are distinct corresponding thresholds for each sampling.

maximum or minimum time, energy, or quality, as appropriate. Following the end of the event, these cell-based quantities are carried through to the flagging code in which all of the results from the run are checked and the appropriate `status` bits are set accordingly. Table 5.2 shows which checks take place, what variables are used in the checks, and which bits are set accordingly. At the time of this writing, the default values for the configurable parameters are untuned, taking values which are unlikely to lead to many status bit changes[2].

[2]In principle, this is correct. However, the nature of the data and the chosen noise suppression truly determine reasonable ranges for parameters. As an examples, the default values for `MinEMfraction` and `MaxEMfraction` are zero and one, respectively – Figure 4.10 clearly shows that a cosmics run or J1 sample without noise suppression take on values outside of these seemingly benign thresholds.

5.1 Code Performance

Performance of `TheCellTool`, in which the loop over all calorimeter cells is performed, is measured using Valgrind [16]. Valgrind offers consistent measurements of performance through a host of separate tools, of which *callgrind* is employed in this study. While the details of the callgrind tool are complex, it is sufficient for this purpose to state that the profiling measurements are strongly correlated to the actual execution time of the algorithms. The real benefit in callgrind is its simplified output, given in the form of a *callgraph*. Figure 5.1 is one such example in which seven events in cosmic run 90272 are profiled. The parent algorithm located to the left, `addFullLArCellsToHelper()`,[3] has a total 'cost' of $\sim 331 \times 10^6$, of which most is attributed to the basic operations inside of that function itself. However, each time another function is called from within `addFullLArCellsToHelper()`, callgrind tracks the cost of that function and assigns it to a node in the graph. In Figure 5.1 the most expensive functions are `getSampling()` and `GetComponent()`, making up 22.0% of the total cost of `TheCellTool`.

Enabling robustness checks leads to a significant increase in computational cost: Figure 5.2 shows that the overall cost of `addFullLArCellsToHelper()` more than doubles. Indeed, applying the cell-level checks indiscriminately over all cells without first rejecting cells based on a SNR leads to a significant computational cost. Applying $2\sigma_{\text{noise}}$ noise suppression reduces the cost increase from 2.060 to 1.557 times the cost of the algorithm without cell-level robustness checks. More can be done, however: since the major contributors to these increases have already been called

[3]The point of interest is to profile the efficiency of the loop over all cells and measure how the computational cost changes as extra cell checks are applied. The loop over all cells is located in two functions: `addFullLArCellsToHelper()` and `addFullTileCellsToHelper()`. The former loops over LAr cells and the latter over tile cells. For this reason, it is necessary only to monitor the behaviour of one since they require the same functions and therefore exhibit the same changes; in this study, only the LAr function will be studied.

(that is, the appearance of new functions in the call graph are not dominating the increased cost), savings can be made merely by anticipating the usage of the variable and storing the result after the first instance of the call. In code, these optimizations take the following form, transforming this:

```
TrigEFMissingEtComponent *metComp = metHelper->GetComponent(cDDE->getSampling());
...
// 3. noise suppression (optional)
  if (m_doCellNoiseSupp &&
      E < m_noiseRMS*(getNoise( *m_it, ICalorimeterNoiseTool::TOTALNOISE)))
...
if (m_makeRobustness && fabs(E) > m_MinCellSNratio[cDDE->getSampling()] *
    getNoise( *m_it, ICalorimeterNoiseTool::TOTALNOISE)
   ) {
```

into this:

```
unsigned int mysampling = cDDE->getSampling();
TrigEFMissingEtComponent *metComp = metHelper->GetComponent( mysampling );
...
double mynoise = 0;
if(m_doCellNoiseSupp || m_makeRobustness){
  mynoise = m_noiseTool->getNoise( *m_it, ICalorimeterNoiseTool::TOTALNOISE);
}
// 3. noise suppression (optional)
  if (m_doCellNoiseSupp && E < m_noiseRMS*(mynoise))
...
if (m_makeRobustness && fabs(E) > m_MinCellSNratio[mysampling] * mynoise
   ) {
```

With these simple optimizations in place, the relative cost increase when includ-

ing robustness checks is 1.913 times without noise suppression and 1.565 times with $2\sigma_{\text{noise}}$ noise suppression. This is an improvement of 7.1% for the case without noise suppression but an increased cost of 0.5% with noise suppression. These miniscule gains and losses are owing to the fact that the bulk of the computation comes from merely requesting noise values for the robustness checks – cutting more than 97% of the cells could not be possible without requesting their noise values.

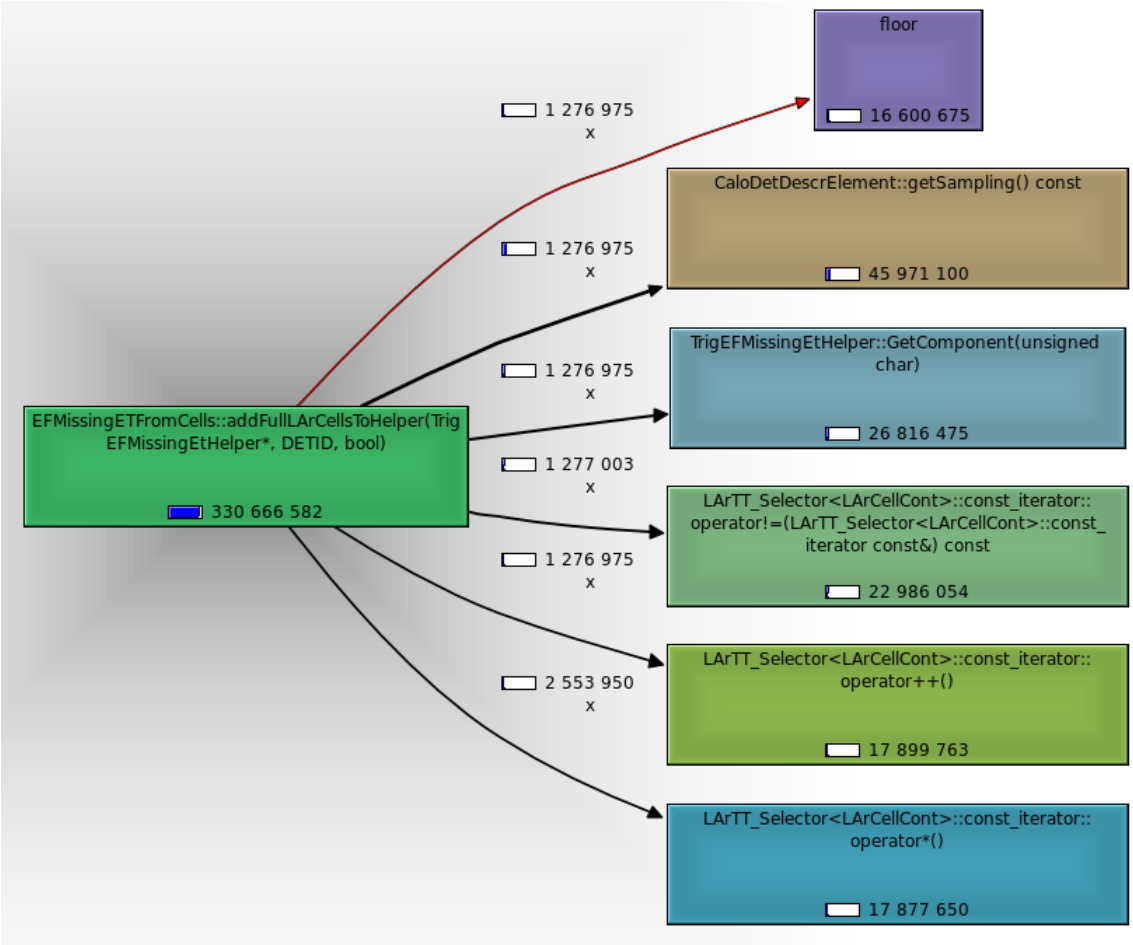

Figure 5.1: Callgraph for seven events of cosmics run 90272 using the callgrind tool. No extra robustness checks are performed and no noise suppression is applied. The units of the callgraph are a combined measure of the 'cost' of a function call, including data reads, cache misses, etc.

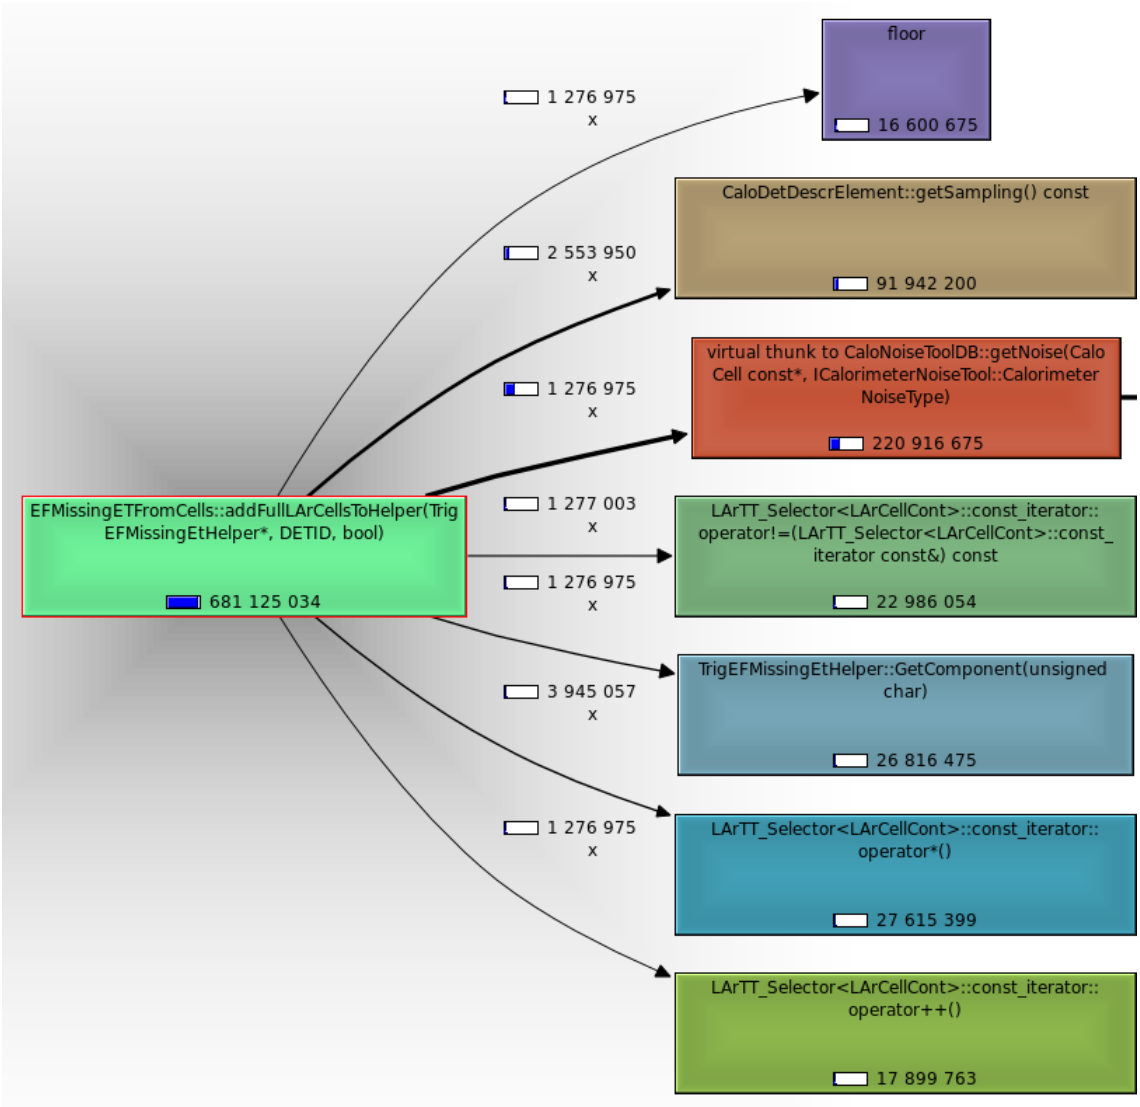

Figure 5.2: Callgraph for seven events of cosmics run 90272 using the callgrind tool. Robustness checks are performed and no noise suppression is applied. The cost of the parent function `addFullLArCellsToHelper()` doubles when robustness is applied.

Chapter 6

Conclusion

In order to prepare the MET software trigger to deal with unexpected input data due to hardware or software glitches, a number of observables have been studied and the cost of computing these quantities on code performance has been evaluated. Two issues have come into play during the analysis: the selection of which cells to apply cell-level robustness checks and which sampling- and event-level checks have good discrimination power. Though cell-level robustness checks ideally would be performed on *all* cells, the cost of the loop over all cells doubles and it is generally agreed that not all cells indeed require checking. Additionally, adding sampling- and event-level robustness without suppressing noise leaves certain quantities without any discriminating power as the possible ranges are spread beyond the limits of what is physical. Although a noise cut optimization is beyond the scope of this study, plots comparing noise suppressions between 2 and $4\sigma_{\text{noise}}$ suggest that noise dominates useful quantities below $2\sigma_{\text{noise}}$ and that signals are very clean beyond $4\sigma_{\text{noise}}$. For these reasons, it is recommended that the current implementation of the cell robustness described in Section 5 be accompanied by a cell-level one-sided noise suppression of at least $2\sigma_{\text{noise}}$. In order to optimize the noise suppression applied, a

study must maximize the number of interesting physics events selected while rejecting a sufficient fraction of the background. With this optimal suppression applied, the quality variables examined in this thesis should once again be analyzed. For events consisting primarily of noise, the computational savings are on the order of 20% when applying noise suppression, with diminishing returns as the number of cells above the noise threshold approaches the total number of cells in the calorimeters. With this recommendation in mind, the following additional points are provided:

1. `MinCellEnergy` – Since large negative energy is unphysical, the lower limit on cell energies being suspicious should be taken as some number of cells N with energy less than $-6\sigma_{\text{noise}}$. For $\sim$195,000 cells in the calorimeter, this equates to a 0.019% chance that a cell of this nature will appear in a given event. As a first trial, N may be taken as one cell and increased to accommodate the increasingly understood behaviour of the detector.

2. `MaxCellEnergy` – As a safe first estimate these values should take on the minimum of either 3 TeV or twice the value in Table 4.3.

3. `CompMaxMEtSumEtRatio` – Cosmics run 90272 and various dijet samples provide evidence suggesting that this ratio offers reasonable values for the entire physical range (see Figure 4.5b). Furthermore, the construction of the ratio guarantees that the distribution will stay within physical limits; unless real beam data can show otherwise, these checks may be safely ignored.

4. `MinSumEratioInEMB` – Reasonable minima with noise suppression can be as low as zero for PreSamplB, EMB1, and EMB3. Until further investigation, EMB2 may take a minimum around 0.1.

5. `MaxSumEratioInEMB` – Maxima for PreSamplB, EMB1, and EMB3 should be

set at 0.8. Energetic dijet samples suggest a maximum around 0.95 for EMB2.

6. `GlobMaxMEtSumEtRatio` – This ratio has successfully separated cosmic events from dijet samples and a threshold as low as 0.4 may distinguish this boundary (see Figure 4.9). As a first step, however, a conservative 0.6 is recommended.

7. `MinEMfraction` – Dijets are separated from cosmics at the minimum around 0.25 (see Figure 4.10b).

8. `MinEMfraction` – No samples exceed a maximum around 0.9 (see Figure 4.10b).

9. `MinTileGapEratio` – A lower bound of zero contains the signal from dijet samples (with the added bonus of excluding some cosmics in the case that no noise suppression is applied, see Figure 4.11a).

10. `MaxTileGapEratio` – Fractions above 0.2 are so few that they are worth checking until beam data is available (see Figure 4.11b).

The work described in this thesis has informed the implementation of the production software used for the ATLAS MET trigger, which is expected to collect its first data from proton-proton collisions within a month. The studies performed here have guided the choice of noise suppression cut; a 3sigma one-sided suppression is now the default. The data quality assessment criteria described in recommendations four (4) through ten (10) are implemented in production code, and result in the setting of quality flags in the trigger output. Although in-depth checks at the cell level are deemed to be too computationally expensive, the `MinCellEnergy` and `MaxCellEnergy` (items one and two from the above list) are also monitored. Further tuning the parameters of the quantities above requires repeating these studies with well-understood beam data.

Bibliography

[1] Peter W. Higgs. Spontaneous symmetry breakdown without massless bosons. *Phys. Rev.*, 145(4):1156–1163, May 1966.

[2] The ATLAS Collaboration. The atlas experiment at the cern large hadron collider. *Journal of Instrumentation*, 3(08):S08003, 2008.

[3] P Laurelli et al. ATLAS muon spectrometer: Technical design report. *Internal Technical Design Report*, 1997. CERN-LHCC-97-22.

[4] W E Cleland and E G Stern. Signal processing considerations for liquid ionization calorimeters in a high rate environment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 338(2-3):467–497, 1 1994.

[5] K Kordas et al. The atlas data acquisition and trigger: concept, design and status. Technical Report ATL-DAQ-CONF-2007-022. ATL-COM-DAQ-2007-015, CERN, Geneva, Nov 2006.

[6] I Riu et al. Integration of the trigger and data acquisition systems in atlas. *Nuclear Science, IEEE Transactions on*, 55(1):106–112, Feb. 2008.

[7] T Lari. Search for supersymmetry with early atlas data. *IFAE 2006*, pages 207–210, 2007.

[8] T Sjostrand, S Mrenna, and P Skands. PYTHIA 6.4 Physics and Manual. *JHEP*, 05:026, 2006.

[9] S Frixione and B R Webber. Matching NLO QCD computations and parton shower simulations. *JHEP*, 06:029, 2002.

[10] M Goossens. *GEANT: Detector Description and Simulation Tool, long writeup W5013; March 1994.* CERN Program Library. CERN, Geneva, 1993.

[11] Xin-Nian Wang. Systematic study of high *pt* hadron spectra in *pp*, *pa*, and *aa* collisions at ultrarelativistic energies. *Phys. Rev. C*, 61(6):064910, May 2000.

[12] *ATLAS detector and physics performance: Technical Design Report, 1.* Technical Design Report ATLAS. CERN, Geneva, 1999.

[13] O Buning et al. LHC Design Report. 2. The LHC infrastructure and general services. CERN-2004-003-V-2.

[14] R Teuscher et al. Fake missing transverse energy from atlas calorimeter cosmic ray data. Technical Report ATL-CAL-INT-2008-001. ATL-COM-CAL-2008-002, CERN, Geneva, Feb 2008.

[15] Atlantis Event Display. http://cern.ch/atlantis/.

[16] Nicholas Nethercote and Julian Seward. Valgrind: a framework for heavyweight dynamic binary instrumentation. In *PLDI '07: Proceedings of the 2007 ACM SIGPLAN conference on Programming language design and implementation*, pages 89–100, New York, NY, USA, 2007. ACM.