Development of Energy Recovery Algorithms in the CMS ECAL

Duncan Leggat

University of Southampton
Southampton
UK

May 2011

Abstract

Many beyond the Standard Model theories such as supersymmetry have proton-proton collision signatures that include missing transverse energy (MET) in the final state. To detect these at the LHC a hermetic detector must be maintained to a high degree of accuracy. Studies have shown that dead regions in the CMS ECAL ($\sim 1\%$ of the total detector, either masked in readout or RECO) contribute significantly to high MET signals in CMS. Different approaches to reducing the false MET signal were investigated. An investigation into the number of crystals that cannot be used for physics was carried out using beam splash and laser data. Following this an algorithm was developed to recover energy deposited in crystals that have no data-taking optical fibre from trigger information collected by the same crystals. The algorithm has been fully validated and is now a part of the standard reconstruction of CMS data. Finally a flagging algorithm was developed to use information captured by the ECAL preshower detector for use in areas with neither data read-out nor trigger information. The algorithm flags events that have high energy deposits in the preshower as having a high energy deposit, and so removing them from MET searches. The code is available for groups that wish to use it.

Acknowledgments

I would like to thank my supervisors David Petyt and Dave Cockerill and the immediate members of the ECAL DPG group for their help and support on any and all problems I have ecountered in my work on the ECAL. Without their help I would understand even less about the ECAL and the (sometimes very convoluted) software that keeps it running.

I would also like to thank Dave Barney and Chia-Ming Kyo for giving me a crash course in how the ECAL preshower and its related software work, as well as introducing me to the joys of Monte Carlo simulations.

I am also very grateful to the UK LTA PhD students who have made CERN feel like a second home, introduced me to many new activities and provided many entertaining evenings and sofas to sleep on.

As for new activities, I would like to thank Switzerland's hospitals for fixing my arm after I broke it snowboarding, and Southampton's insurance for ensuring that operation wasn't far beyond an affordable price tag.

I would especially like to thanks my housemates, Tom and Adam for putting up with my general messiness, lack of organisation and loud music at inappropriate times. It has been a pleasure living with them.

Finally I would like to thank Southampton University and Stefano Moretti for making this wonderful oppurtunity possible.

Contents

1 Introduction

The Compact Muon Solenoid (CMS) is a general-purpose detector operating at the Large Hadron Collider (LHC) at CERN [14]. It was built with the intention of filling in the gaps in our current understanding of particle physics and the Standard Model. In particular it is searching for the Higgs boson and evidence for physics beyond the standard model, such as supersymmetry and extra dimensions. In order to do this, it requires high precision tracking and energy measurements of the many different particles created in the proton-proton collisions provided by the LHC. It contains a silicon tracker that closely monitors the path of charged particles close to the beam-pipe, and an electromagnetic and hadronic calorimeter. The detector is surrounded by muon chambers that track the muon's position through 4 layers of iron and detector to determine its path and ultimately its momentum. The detector layout can be seen in 1.

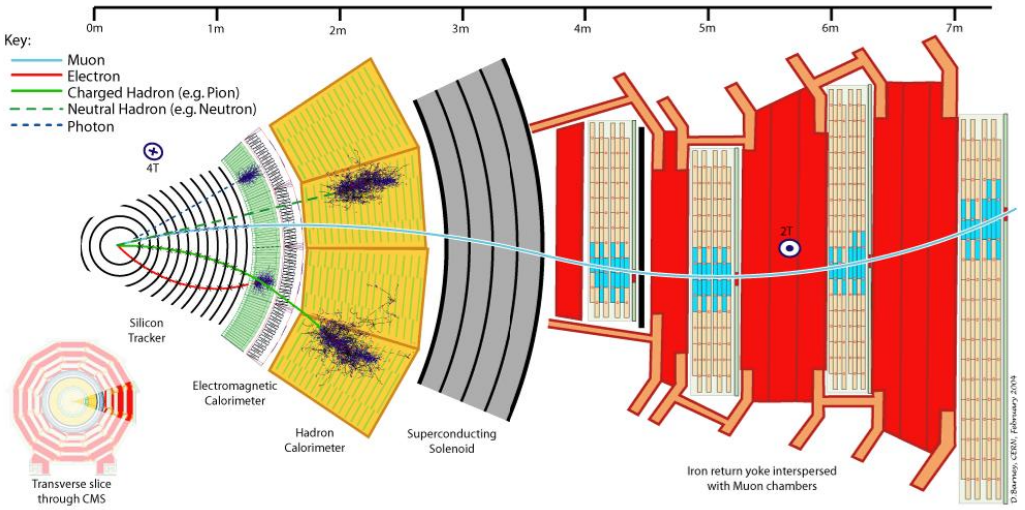

Figure 1: A transverse slice through the CMS detector showing the different subdetectors in the barrel.

The Standard Model is currently the best theory we have for understanding fundamental interactions between particles in the form of the three natural forces; the electromagnetic force and the weak and strong nuclear interactions. It has undergone rigorous tests and has given accurate predictions on the existence of the heavier quarks and W and Z gauge bosons. It also predicts the existence of the Higgs boson, which, if it exists at the predicted mass range, should be created at the LHC. However, there are many intrinsic problems with the Standard Model. Much like every particle

in the Standard Model has an anti-matter equivalent of opposite charge, it has been postulated that each particle in the standard model has a heavier 'superpartner'. This theory is known as supersymmetry (SUSY).

Whilst there is currently no experimental evidence for SUSY, it is possible that the record energy levels of the LHC will create the first instances of the lightest superpartners. Many potential SUSY signatures include jets and high missing transverse energy which can be easily distinguished from background collisions. The detection of this missing energy requires the full detector to be working as well as possible.

CMS contains an electromagnetic calorimeter (ECAL) [4] which is comprised three main sections, the barrel (abbreviated as EE) and endcap (EB) crystal scintillation detectors and the silicon preshower (ES). The barrel and endcap are Lead Tungstate ($PbWO_4$) scintillator crystal detectors that measure the energy of electromagnetic particles that enter it. The preshower covers a region in front of the endcap detector and is a silicon strip tracker detector with a high spatial resolution. It was designed with the intention of identifying energy deposits in the endcap as single or multiple particles, in particular discriminating between high-energy photons from potential Higgs boson decays and closely spaced double photon decays from the far more commonly produced pi^0.

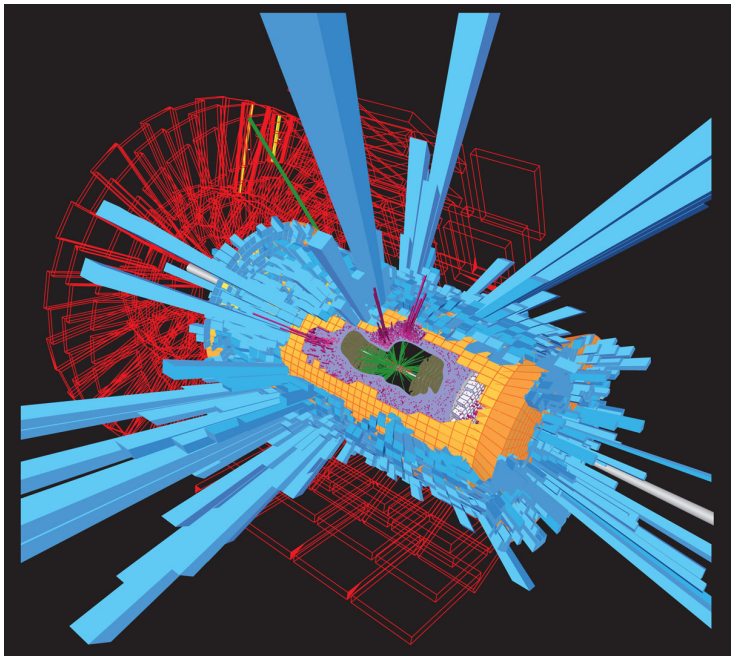

Figure 2: A top/anti-top pair-production event from CMS

The ECAL contains a total of 75848 crystals divided between the larger

barrel and two endcap detectors. Each crystal is a channel that the detector uses to take data. There are, as of the end of 2010 running, 150 inactive channels in the endcap (1.02%) and 409 inactive in the barrel (0.67%). The first aim of this project was to examine beam splash[1] events and laser calibration data for the ECAL to search for and classify any new problematic channels.

These inactive channels, or dead regions, lead to effective holes in the detector that electromagnetic particles can disappear through without having their energy registered. This gives rise to a false missing transverse energy signature. If a high-energy hadron jet caused by the hadronization of a quark ejected from its original nucleus coincides with one of these dead regions, large missing transverse energy is registered. Since the signature of SUSY particles involves large MET, the presence of these holes greatly impacts the sensitity of SUSY searches in CMS. For these reasons it is important to maintain CMS as a hermetic detector[2] and to attempt to lower the impact of any dead regions as far as possible.

In many cases both the crystal and front-end electronics of the problematic channels are fully functional but a faulty data read-out fibre optic cable stops information being relayed. In these cases, the front-end still supplies energy deposition information to the trigger system. By looking at the amount of energy registered in the trigger it is possible to estimate the energy that was deposited in the individual crystals. The second aim of this project was to develop an algorithm that would do this for the endcap.

The preshower detector gives another way to attempt to recover lost energy through dead ECAL regions. The preshower covers roughly half of the area of the endcap and gives high spatial but low energy resolution of incoming electromagnetic particles. By looking at the energy registered by the preshower it should be possible to create flags for the minimum energy deposited in a dead region. Whilst the resolution of these measurements will be relatively low it will help exclude a large amount of the false missing energy signal in the affected areas. The final aim of this project was to determine efficient and effective cut-off values for this.

[1] A large flux of beam-induced muons through the detector

[2] A hermetic detector is one designed to be able to observe as many particle decay options as possible by covering as large a viewing angle as possible. They are also known as 4π detectors because of the full angular range they cover.

2 The Standard Model, Supersymmetry and the Detection at the LHC

2.1 The Standard Model

The Standard Model is a very successful description of the workings of the electromagnetic, weak and strong interactions of subatomic particles that has been highly developed over the past 50 years. It has its origins in the unification of the electromagnetic and weak interactions into the so-called electroweak force by Sheldon Glashow in 1960 [10] and the subsequent integration of the Higgs mechanism by Steven Weinberg and Abdus Salam in 1967 [15]. The strong interaction became integrated into this in around 1974 when experiments showed that hadrons comprised fractionally charged quarks.

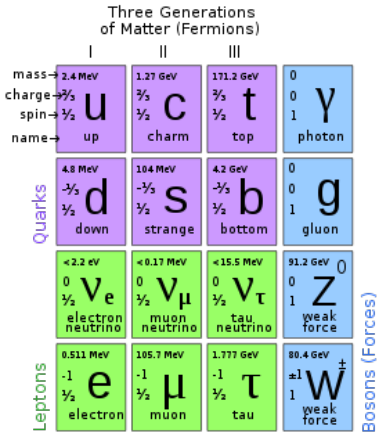

Figure 3: The Standard Model of particle physics. Contains 12 fermions divided between 6 quarks and 6 leptons divided into 2 generations and 4 force carrying gauge bosons. The Standard Model also predicts the existance of the Higgs boson.

2.1.1 Composition of the Standard Model

Figure 3 shows the composition of the standard model. It predicts 12 fermions[3] that are the fundamental particles that make up all matter. These are broken down into 6 quarks and 6 leptons. Quarks are the basic building blocks for nuclei whilst leptons include the electron and are responsible for the chemical properties of different atoms. The quarks are further broken down in 3 progressively more massive generations; 'up' and 'down' make up the first generation, 'charm' and 'strange' the second and 'top' and 'bottom' quarks

[3] A fermion is a fundamental particle with a half-integer spin whilst a boson is a particle with an integer spin.

make up the final third generation. The leptons comprise two different variants, the charged leptons and the neutral neutrinos, but these are also arranged into three generations; the electron and electron-neutrino, the muon and muon-neutrino and the tau and tau-neutrino. The electron, muon and tau are electrically charged and massive whilst the neutrinos are neutral and have a much smaller (but non-zero) mass.

These particles interact via the four fundamental forces of nature; gravity, electromagnetism and the weak and strong nuclear forces. We know that the last three of these interact via force-exchanging particles called gauge bosons. Photons (γ) carry the electromagnetic force, gluons the strong force and the W and Z bosons the weak nuclear force. W and Z bosons were predicted by the standard model and were later found experimentally at the mass predicted by the model. Gluons operate only over small distances and so cannot be observed directly, but their existence can be inferred from hadron jets (see section 5.1) in particle colliders. It is predicted that gravity also has a force-exchange particle known as a graviton, but no evidence has been seen for their existence. The standard model also predicts the existence of a fifth boson, the Higgs boson. This is the particle that gives mass to the other particles, and its discovery is one of the primary goals of the LHC.

The Standard Model is what is as known as a quantum field theory, in that is uses the idea of quantum fields permeating space as a basis for the fundamental interactions. A field is a quantity associated to every point in space time. A classical field has a fixed number of parameters[4], but a quantum field can have an infinite number allowing it to parameterise the entirety of the standard model given the right set of restrictions. A particle passing through space also passes through and - depending on how it interacts - disturbs this field; oscillations in the field give rise to the gauge bosons and the fundamental interactions of nature[5]. Particles interacting with the theoretical Higgs field are slowed down in their movement through spacetime via the Higgs boson - massless particles like the photon do not interact with the Higgs field and so are free to move at the speed of light.

2.2 Problems with the Standard Model

Although the Standard Model has been very successful in both explaining known phenomena and correctly predicting the existence and mass of further particles, it is not a complete description of nature. Gravity is not included

[4]For example, the electromagnetic field has two parameters, the electric and magnetic componenets.

[5]An oscillation in the electromagnetic field is a photon.

and any current attempts to reconcile general relativity with a quantum field description give unphysical results and break down before the Planck scale[6].

Another major blow to the Standard Model was the experimental evidence that neutrinos have a finite mass. Naturally the Standard Model predicts that neutrinos should be massless, but evidence from distant supernovae[7] and solar neutrino observations[8] imply that they must have a small mass. It is possible to reconcile the standard model with a neutrino with mass, but it adds further constraints to an already very fine-tuned theory. There are 19 unrelated and arbitrary constants in the Standard Model and the addition of a neutrino mass adds a further 7 or 8 further unrelated constants; many consider a theory based around such arbitrary measurements inelegant and unlikely to be the true description of nature. The problem of fine-tuning is further compounded by the Higgs mass quadratic divergence problem. For a Higgs boson at a light mass predicted by the Standard Model the mass diverges quadratically, meaning that whatever mass the Higgs is measured at requires heavy fine tuning of the model. The need for such fine-tuning is known as the 'Hierarchy Problem' of particle physics.

The Standard Model also offers no description for observed 'dark matter' in the universe. The observed rotation speed of observed galaxies (calculated from the red-shift of stars in different parts of the galaxy) appears to be much higher than that predicted for the amount of mass we observe in them. This apparent extra mass is also observed in gravitational lensing of distant light sources[9] and the temperature distribution of hot gas in other galaxies. Our current observations of these anomalies imply that only 17% of matter in the universe can be described by ordinary matter and the rest is only present in the form of 'dark matter'; a form of matter that only interacts through the gravitational force. The standard model cannot offer any particle that could act in this way.

[6]The Planck scale is the energy scale at which quantum effects of gravity become strong, $\sim 1.22 \times 10^{19}$ GeV.

[7]A supernova gives off a large flux of photons and neutrinos. Observations of distant supernovae show that the neutrinos arrive slightly after the photons from the explosiong; implying that they travel slower than the speed of light and so must have a non-zero mass.

[8]The detected neutrino flux from the sun is a third less than that predicted by theory; this is due to neutrino oscillation that can only occur if the neutrino has mass.

[9]The mass of galaxies curves spacetime and focuses light from distant sources. The effect is larger than it should be for the amount of mass in observable matter in the galaxies.

2.3 Beyond the Standard Model - Supersymmetry

Many potential solutions to these problems have been proposed, but one of the most promising is supersymmetry (SUSY). This theory states that in the same way that every particle has an associated anti-matter equivalent, each particle of the Standard Model has a heavier 'superpartner' with spin a 1/2 unit different. This means that all fermions have a bosonic superpartner, and all the standard model bosons have a fermionic particle associated with them.

Supersymmetry is a very appealing proposal for a number of reasons. It allows the unification of the strong, weak and electromagnetic forces into one force at $\sim 10^{16}$ GeV, provides plausible Dark Matter candidates in the form of stable and massive yet electrically neutral superpartners, and the contributions of heavier superpartners remove the quadratic divergence and hence hierarchy problems. Although the theory is very appealing there is currently no experimental evidence to support it and the precise properties (such as mass or charge) of the superpartners are unknown. Whilst the search for them is a priority, the lack of information about superpartners makes the search very difficult at any high energy particle detector.

2.4 Detection of Supersymmetry at CMS

In a proton-proton collision, the initial state contains two protons being accelerated along the beam pipe towards the collision point - this means that there is no transverse momentum in the initial state. Conservation of linear momentum means that in the final state the sum of transverse momentum must therefore be zero overall. Particles such as neutrinos interact very weakly and so do not leave any information on their momentum in the detector; by looking for an imbalance in the momentum of the final state of a p-p collision the existence of a non-interacting particle can be inferred. At the high energy of collisions at the LHC the mass of particles is negligible compared to their energy, so from Einstein's equation $E = \sqrt{(pc)^2 + (mc^2)^2}$ momentum is equivalent to energy. This gives rise to the search for missing transverse energy, or missing Et or MET.

In SUSY models where the lightest superpartners may be created at the LHC many are electrically neutral and so do not interact with the particle detectors. Many final states representative of SUSY have large MET signals. Other potential signals for SUSY events include hadron jets (explained in section 5.1) and distinctive combinations of multiple leptons in the final state [8]. Because these searches are highly sensitive to the amount of missing energy in the final state, it is important to have high precision detectors

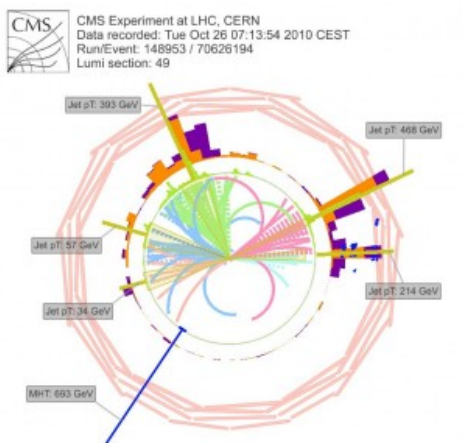

Figure 4: An event display from October 2010 showing an event with a high number of high energy jets and a large amount of missing E_t represented by the blue line in the bottom left.

covering as large an angle around the interaction point as possible. Any gaps or blind spots in the detector can give rise to false missing energy signatures, so maintaining the detector and coming up with new ways to recover loss in performance is crucial to the overall running of CMS.

3 The CMS ECAL

The CMS ECAL consists of three main detectors - the barrel, two endcaps and preshower detectors. The barrel detector sits parallel to the beam-pipe and covers the range $|\eta| < 1.48$[10] whilst the two endcaps close the detector covering the region $1.48 < |\eta| < 3.0$. The preshower sits in front of the endcap and covers the region $1.65 < |\eta| < 2.6$. The endcap and barrel detectors are made up of Lead Tungstate scintillation crystals connected to photodetectors.

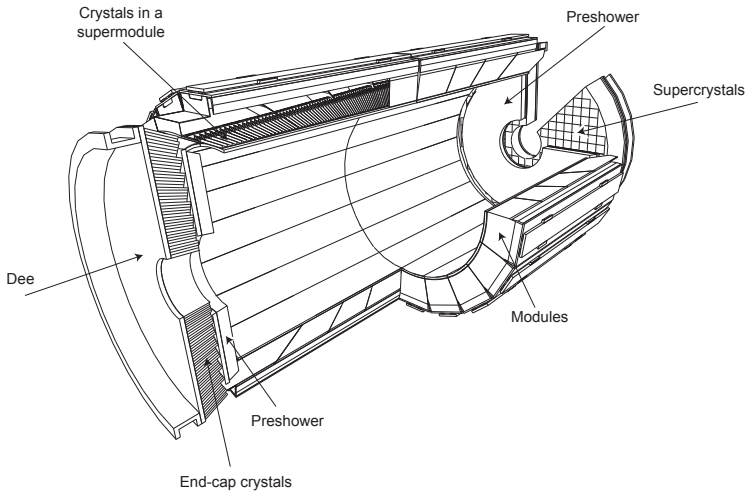

Figure 5: The CMS ECAL. The different sections of the calorimeter are labelled.

3.1 Lead Tungstate Scintillation Detectors

The endcap and barrel detectors are made of Lead Tungstate crystals ($PbWO_4$), which produce scintillation light. When a high-energy electron (or other charged particle) enters the scintillator it begins an electromagnetic shower. At energies expected at the LHC, electrons primarily lose energy through bremsstrahlung radiation - the process of emission of a photon from the acceleration of the electron in the electromagnetic field due to interaction and scattering off atomic nuclei. These emitted photons are also of high energy

[10] η is known as the pseudorapidity and is a spatial co-ordinate used in high-energy particle physics. It represents the angle of a particle (or detector element) relative to the beam axis. It is defined as $\eta = -ln[tan(\frac{\theta}{2})]$ where θ is the angle between the beam axis and the particle.

and so mainly interact through electron-positron pair production. When the shower has lost sufficient energy it ionises the scintillation crystals which re-emit the energy as photons that the attached photodetectors measure.

Lead tungstate was chosen for the specific needs of the detector as it is very dense (8.28g/cm^3) and has a short decay length[11] (8.8mm) allowing a compact design whilst the small Molière radius[12] in comparison to other crystal scintillators (2.19cm) allows for a high granularity[13] and resolution on the detector. Lead tungstate also has a very fast decay time (80% within 25ns) meeting the criteria of fast read-out required for the large number of collisions (luminosity) provided by the LHC. The draw-backs of $PbWO_4$ are its relatively low light yield and high temperature dependence ($-2\%/^\circ$C)[12]. This requires both a high quantum efficiency photodetector and a very stable cooling system.

The barrel consists of 61200 crystals split into 36 supermodules each connected to a pair of avalanche photodiodes (APD) to detect the scintillation light. The endcap contains 14648 crystals divided between two endcaps, each of which is split into two 'dee's for construction purposes. In the endcap the higher radiation levels would damage APDs, so vacuum photo triodes (VPT) are used instead.

3.2 The Trigger and Laser Calibration Systems

3.2.1 The ECAL Trigger

The nominal LHC design luminosity is 10^{34} cm^{-2}s^{-1} giving the equivalent of 10^9 interactions per second[14]. The maximum rate that the on-line computer farm[15] can archive is 100Hz so a suppression of 10^7 is required in the data. This is handled by the CMS trigger system, a full description of which can be found in the technical design report [3]. The crystals in both sections of the detector are arranged in 5x5 'supercrystals' connected to front-end electronics. The front-end (FE) cards receive information from the crystals, process it and send information to the trigger system. The data is only read-out if a trigger signal above a programmable threshold is received.

Groups of crystals in the ECAL are associated with trigger towers that collect information from the FE electronics regarding the energy and timing

[11]The length over which a high-energy electron loses $1/e$ of its energy.

[12]The Molière radius is a material constant that indicates the radius of fully contained electromagnetic showers from high energy photons or electrons. The constant is defined as the radius of a cylinder within which 90% of a particle's energy is deposited.

[13]The amount the detector is broken down into smaller subsections

[14]The machine currently operates at $10^{33} cm^{-2} s^{-1}$ [2]

[15]The off-machine electronics recording the data

of hits in the crystals. The tower creates a 'trigger primitive' (TP) that contains the sum of the transverse energies in the linked crystals, the timing (with 25ns precision) for which the collision occurred and a fine-grain veto bit[16] (FGVB)[6]. The TP stores the transverse energy as a single byte which saturates at 64GeV. (This was increased to 128GeV for the beginning of 2011 running). The so-called Level-1 trigger[3] reduces the data rate to below 100kHz before passing this information to the High Level Trigger (HLT) which makes the final reduction to 100Hz which is written to disk.

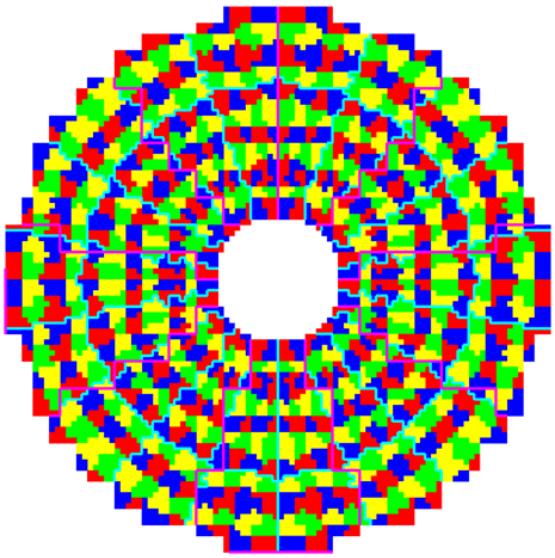

Figure 6: The trigger towers in the CMS ECAL endcap. The towers are arrange in *ieta* rings of 72 towers each. Note that at $|ieta| > 26$ there are half as many towers in each ring.

These trigger towers are numbered by their position in the detector, labelled by '*iphi*' and '*ieta*' [11]. *ieta* is similar to the pseudorapidity of the tower, but is based around the relative positions in the detector rather than the absolute angular value. *iphi* is the azimuthal angle of the tower. In the barrel the geometry is simple and the FE cards and read-out correspond to the same 5x5 blocks of crystals as the trigger towers. In the endcap the

[16]The fine-grain veto bit is a single bit which is a measure of how spread-out the energy is within the trigger tower. The tower looks for the highest energy deposit in its region, sums it with its next highest neighbour and finds the ratio of this energy to the total deposited in the tower. If this ratio is above a certain threshold (typically 95%), the bit is set. As electrons and photons both deposit their energy in very few crystals (roughly 80% in a single crystal), different physics objects can be quickly identified in this way.[5]

geometry is more complicated; the need to assign the trigger towers angular identification means that FE cards and trigger towers have overlaps between them.

3.2.2 The Laser Monitoring System

During the 10 year expected lifetime of the ECAL the crystals will be exposed to up to 4 kGy in the barrel and 200 kGy in the endcaps[12]. Prolonged exposure to radiation causes the crystals to darken, transmitting less light and thus giving a lower gain for the detector. Whilst heating the crystals can partially reverse the damage in a process called 'crystal annealing', it is important to monitor the changes for calibration purposes. Radiation damage changes to crystal transparency can take place in much shorter time scales than can be corrected through full detector calibration - hours as opposed to days - and so require an on-the-fly way to monitor and correct the data. This is achieved by a laser monitoring system [1]. Every crystal is connected by fibre optic cables to a red and blue laser, whilst the endcap crystals are also connected to a 600nm LED for monitoring purposes. During standard running the laser cycles through and illuminates different sections of the ECAL, providing information regarding the timing and gain of each channel.

3.3 The ECAL Preshower

The ECAL preshower is a silicon strip detector built to improve pi^0/γ exclusion. One of the most distinguishable Higgs boson decays is into high-energy photons and the ECAL has been designed with this decay in mind. However, a problem arises from neutral pion (pi^0) decays in the endcap region of the detector. pi^0s (which are commonly produced at the LHC) can decay into two lower energy but very closely spaced photons. In the high $|\eta|$ range that the endcaps cover it is possible that these two photons will hit in a highly localised area and appear as a high-energy photon mimicking that of a Higgs decay. The preshower contains modules of 6.3 x 6.3 cm^2 containing 32 strips of 1.9 mm each. Since the endcap crystals have an area of 3 x 3 cm^2 the preshower provides a much higher granularity to distinguish between the two events.

The preshower contains two planes of orthogonal silicon strip detectors each fronted by a layer of lead. The lead layers, 2 and 1 radiation length[17] in thickness respectively, cause the photons to shower creating electron-positron

[17]One radiation length is the distance within a material that a high-energy electron loses all but $1/e$ of its energy due to Bremsstrahlung radiation

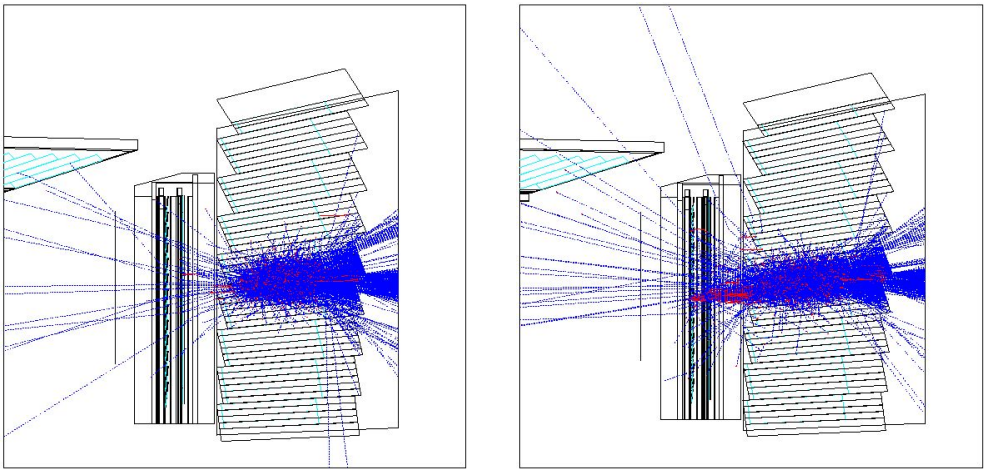

(a) Higgs single-photon decay (b) pi^0 two-photon decay

Figure 7: The two figures show a single high-energy Higgs photon and a two-photon pi^0 decay depositing energy in the ECAL endcap. In the endcap crystals the events look very similar but the higher spatial resolution of the preshower allows the two different events to be identified.

pairs which the silicon sensors detect. The detector works by doping thin strips of silicon to make them into diodes which are then reverse biased. Charged particles passing through this cause small ionization currents which can be measured. By looking at photon energy deposits in the endcap and projecting back towards the interaction point the preshower is examined to determine the nature of the hit - whether it be a true single high-energy photon or a photon pair.

4 Beam Splash and Laser Data Analysis

A beam splash event occurs when the accelerated beam of protons are intentionally directed into a collimator block[18], causing a shower of pions that quickly decay into muons and other high energy particles. This large flux of particles passes through the detector and causes every channel to register large energy. The resulting read-out information is used to check calibration, timing information and the status of various parts of the detector.

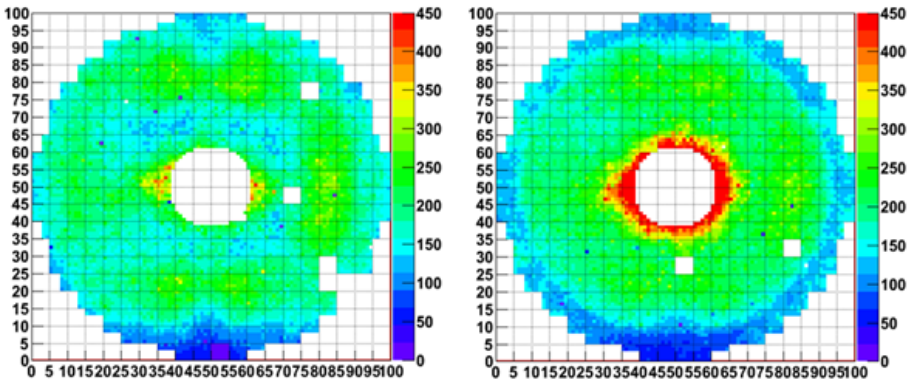

Figure 8: A beam splash event from run 144980 in the CMS ECAL endcaps. The energy deposition is non-uniform due to features of the cavern and the detector acting as shielding.

The energy deposited is not uniform throughout the detector; the splash particles are generated predominantly along the beam axis and as such a larger energy deposition is clear in higher pseudorapidity areas of the detector. The layout of the cavern containing CMS and various parts of the detector also affect the energy deposition in these events, for example there is a noteable section at the bottom of the detector lower in energy that is attributed to the fact that the detector is slightly set into the floor. There is also a ring around one of the endcaps where the barrel detector acts as shielding from the incoming particles.

A channel with zero or low amplitude implies a problem, but it could be as simple as a broken laser fibre. By looking at both laser data and beam splash events problematic channels can be categorised as an intrinsic problem with the crystal or a laser problem. The initial aim of this project was to analyse and compare beam splash events with contemporary laser monitoring

[18] A collimator is usually used to focus the proton beam.

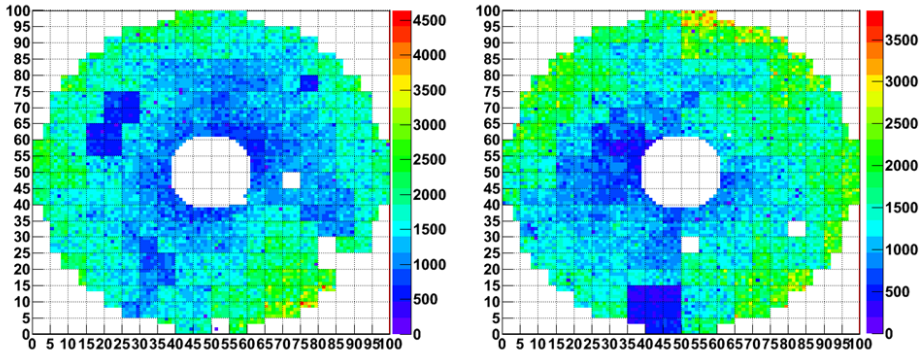

Figure 9: Data from the laser monitoring system during run 144839. There is a large spead in amplitudes in the read-out so normalisation must be applied to determine low response channels.

data to investigate the status of data channels in the ECAL and specifically to check that the current masking map was still relevant.

4.1 Analysis of Endcap and Comparison with 2009 Study

Data was taken from the beam splash run[19] 144890 from summer 2010 whilst laser data was obtained from a calibration run taken around the same time, run 144839. The aim of this analysis was to search for channels with low energy response to the splash or laser data. These are considered to be problematic channels. Information from crystal read-out is stored in the CMS computing software framework (CMSSW) as a collection of reconstructed hits (recHits). A recHit is a c++ data structure containing information on the energy, timing, pulse shape and various flags relating to the status of the channel and the hit.[maybe go into more detail?]

The non-uniformity in energy deposition meant that a precise definition of low response channel needed to be defined. From looking at the splash energy maps in figure 8 we can see a local uniformity in energy distribution. The laser map in figure 9 also displays uniformity in local 5x5 supercrystals, so for both data sets the amplitude of each channel is normalised over its local supercrystal using the formula:

$$r_i = \frac{a_i}{\frac{1}{(N_{xtal}-1)}\sum_{j=1}^{N_{xtal}} a_j}(j \neq i) \tag{1}$$

[19] A run is a numeric label given to the collection of events that occur during a single operation of the beams.

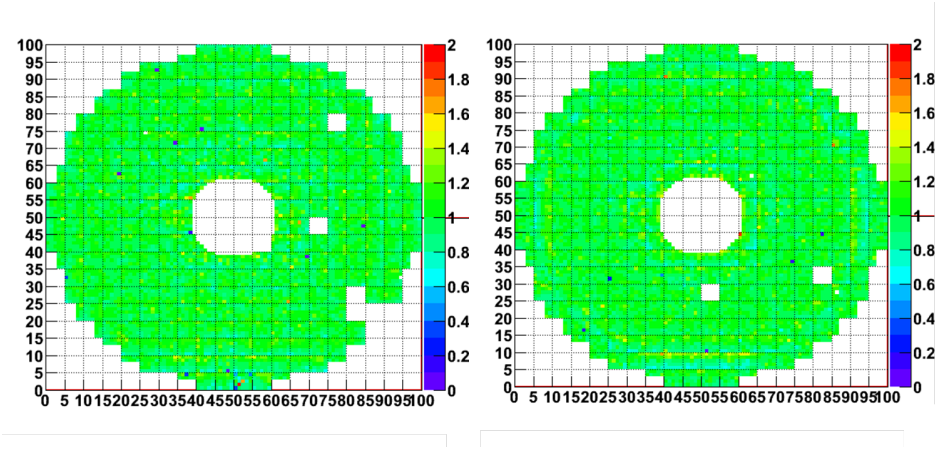

Figure 10: The splash events of figure 8 have been normalised. It is now possible to pick out channels with low response.

Channels with a normalised pulse height of less than 5% in both splash and laser data were flagged, whilst channels with low laser response only are labelled as a monitoring problem and so are still used for studying physics. Figure 12 shows the difference in pulse response between a responding but low gain channel and a malfunctioning crystal. New dead channel candidates had their pulse shapes checked and were then flagged as either dead or low gain channels. Comparing results with a study on the endcap using the same method carried out in 2009 revealed 3 new dead channels. There was also one channel from the 2009 study that has appeared to recover; the pulse height and shape are now nominal, as can be seen in figure 13.

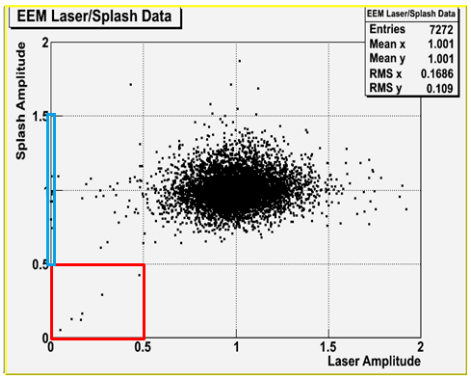

Figure 11: Comparison of the laser and splash amplitudes over the two runs. The problematic channels are those within the red box with low amplitude in both data sets. Those in the blue box could indicate a broken laser fibre.

4.2 Extention of Analysis to Barrel

The 2009 study only looked at the endcap; this analysis extended the method for the first time to the barrel detector. In the barrel the normalisation process is slightly more complicated; there are several 5x5 supercrystals with low but non-zero response which would appear to be working normally if normalised in the same way as the endcap. To avoid this, the barrel was normalised over 10x10 regions.

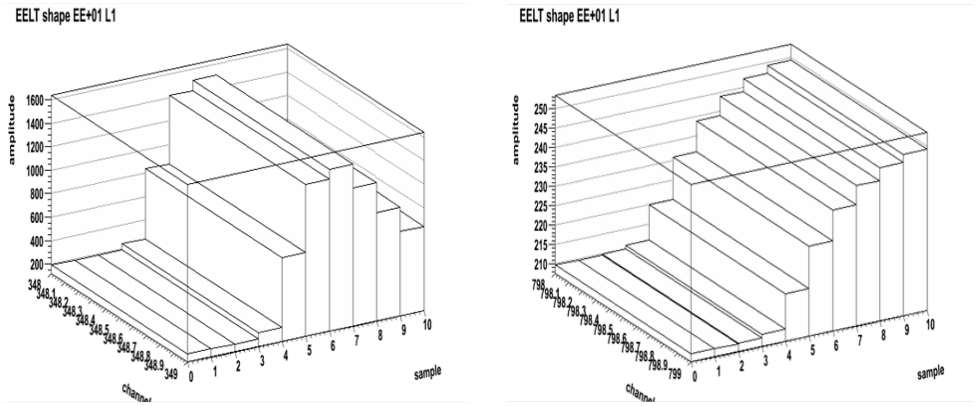

(a) Good Pulse Shape (b) Bad Pulse Shape

Figure 12: 12a shows how the photodetectors should respond to a hit; a well-defined hit at around 150ns (the timing axis is divided into 10 samples divided by 25 ns) with a tail that quickly disipates. 12b shows the characteristic of a malfunctioning channel; no well-defined hit and increasing in energy after the 10 measurements. Photodetectors with this response cannot be used in read-out. The two displayed pulses are from the same EE supermodule, EE+1.

As the 2009 study was not carried out in the barrel the list of bad channels was compared with the then-current list of known dead channels. Every channel is assigned a value that details its current status and instructs the computing software whether to use the data, to attempt to recover the energy from a dead channel or to ignore a bad channel. This list is updated regularly to reflect the constantly changing nature of the channels in the ECAL. Channels with status 0 are normal crystals, 2 is a channel with no laser data but okay elsewhere, labels 3-7 are designate noisy channels, 8-10 represent gain issues and 11-14 represent dead channels. For this study the

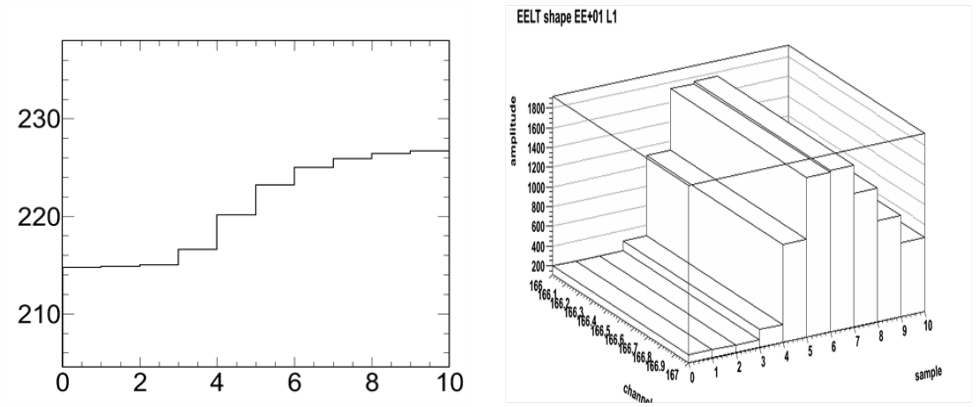

(a) 2009 Response (b) 2010 Response

Figure 13: Channel 45,90 in EE+. The channel appears to have recovered between the 2009 and 2010 analysis.

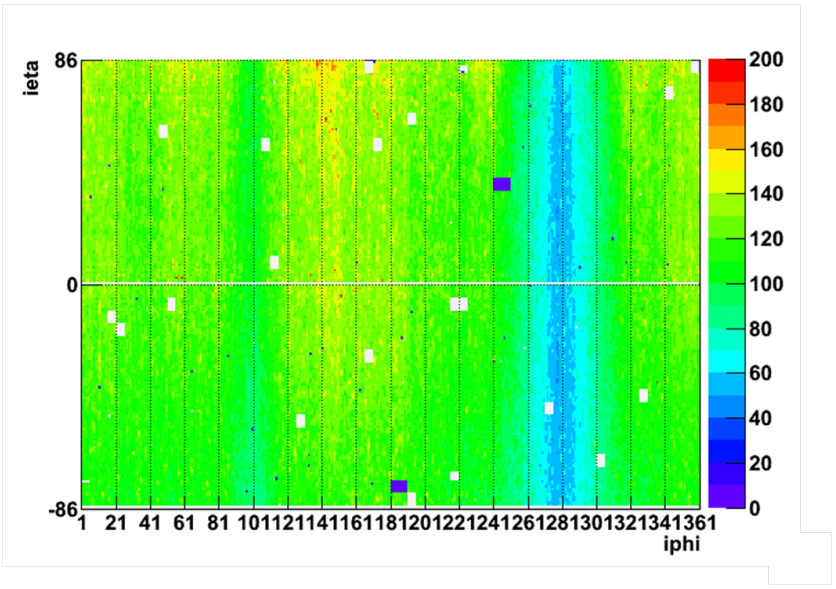

Figure 14: Beam splash event from run 144980 in the CMS ECAL barrel detector. The barrel displays some of the same splash characteristics as the barrel; the shadowing caused by the floor of the cavern can be seen at $iphi = 280$ and energy deposits are greater in larger $ieta$ values because of the direction of the beam.

channels of interest are those assigned 0,2,11[20], 12[21], 13[22] and 14[23].

[20]Non-responding isolated channel.
[21]Dead 5x1 read out strip; a Very-Front-End (VFE) card.
[22]Trigger receives information but data read-out is broken.
[23]A completely dead channel.

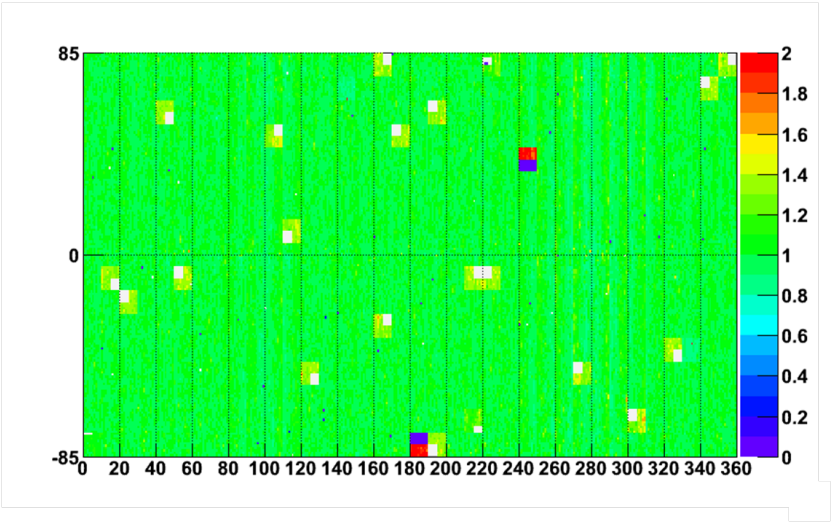

Figure 15: The same event as figure 14 but normalised over 10x10 crystals. As dead regions are now included in the normalisation good channels near a dead CCU appear to

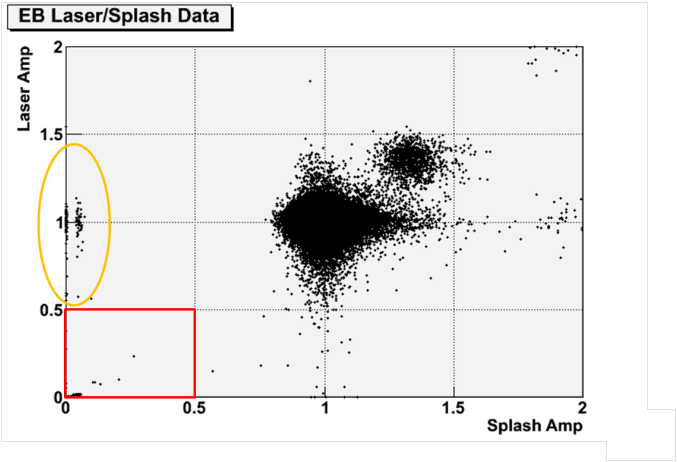

Figure 16: Comparison between the normalised laser and splash amplitudes for the ECAL barrel detector. The group of points to the top right of the main ball is caused by the shadowing seen in figure 15. The points in the orange oval are data with zero splash but nominal laser amplitude. These are caused by reconstruction masking on some channels.[maybe add more on this]

Comparing the results of the study with the collisions2010_7iov [13] masked channel list (figure 17), which was valid between runs 142630 and 147113, three new dead channels were identified. There were three channels marked

as bad that showed a normal amplitude response in the study. Of these, one was very noisy during a few runs and is being monitored to see whether it can be unmasked. Another still has a poor pulse response and so remains masked. The final channel appears to have fully recovered and as a result of this study is now being used to take data. All of the mentioned changes were updated on the subsequent masked channel list, used in the reconstruction of ECAL data.

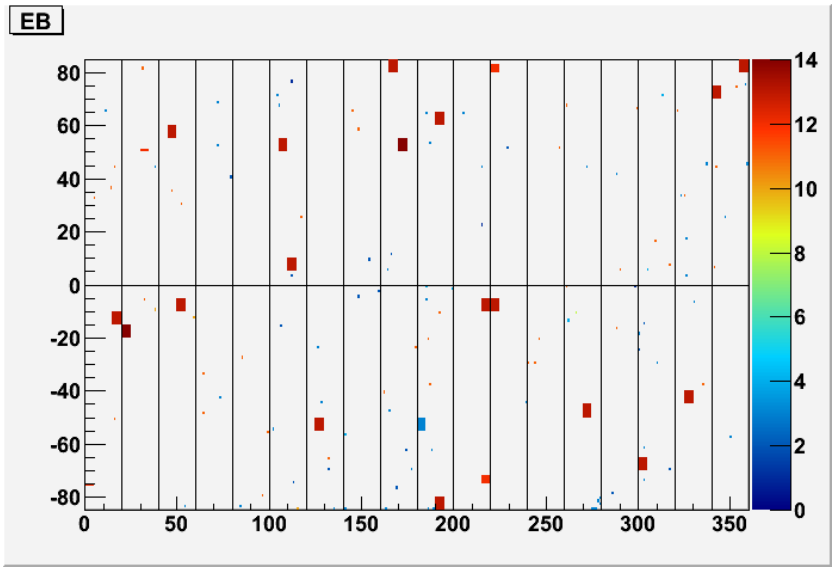

Figure 17: Channel status map collisions2010_7iov for the ECAL barrel detector. This channel map was current at the time of the analysis.

5 Energy Recovery Using Trigger Primitive Information

5.1 Reasons for Study

Of the channels flagged as dead in the ECAL, the majority [find a statistic] are in the form of dead Communication and Control Units (CCUs); a device that is associated with every 5x5 supercrystal node and handles the transmission of data from the front-end electronics. When a CCU is masked it means that the entire supercrystal is not read out. This creates a dead region in the detector - a hole in which energy from electromagnetic particles can disappear. Many of the physics signatures that are of interest to those analysing data from the LHC - such as the search for extra dimensions or supersymmetry - rely on an accurate estimation of missing momentum carried away by exotic or otherwise invisible particles. Many of these searches include jets in the final state and large missing transverse energy (MET).

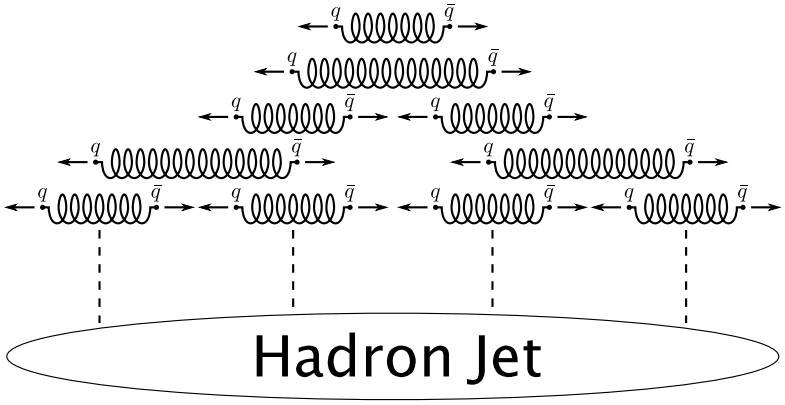

Figure 18: An example of colour confinement leading to hadron jets. The increasing strength of the strong force at greater separations means that at some separation it is more energetically favourable to create a quark/anti-quark pair than continue pulling the quarks apart. The process is similar to an elastic band stretching - eventually it becomes favourable for the band to snap and create two smaller elastic bands.

A jet is the resulting cone of hadrons and other particles that is produced when a single quark or gluon undergoes hadronization. Unlike the electromagnetic potential the strong force potential from a quark increases greatly as distance increases. This means that in order to move quarks further apart requires large amounts of energy to be put into the system - at a certain point it becomes more energetically viable to create a quark anti-quark

pair from the vacuum and create a new nucleus. This leads to a principle known as colour confinement and means that quarks cannot exist in isolation. Hadronization is the process whereby the isolated particle generates further quarks and starts a hadron jet, such as that seen in figure 18. The exact mechanism by which jets are created is not well understood and hence their study is of great importance to quantum chromodynamic (QCD) theories.

A gap in the ECAL means that electromagnetic particles can deposit their energy in the detector but not be registered, giving rise to a false missing momentum signature. In particular one group searching for supersymmetry from jet signatures have expressed concerns with hadron jets overlapping dead ECAL regions, creating a sizeable false signature. Figure 19 shows the impact of the dead region; 50% of high jet MET events overlap with dead regions in the ECAL.

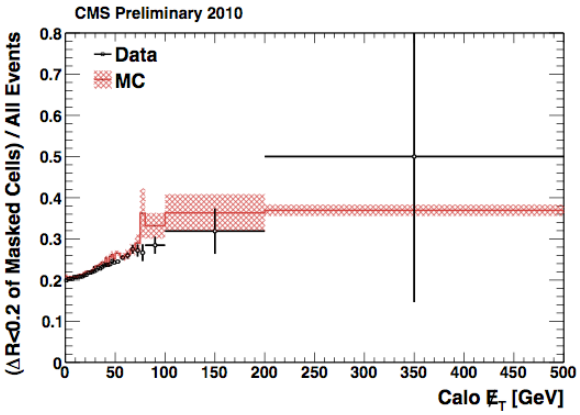

Figure 19: This plot shows the ratio of jets incident on dead regions in the ECAL to the total number of events as a function of the jet energy. At low MET the ratio is 20% increasing to 50% at high MET. This shows the substantial effect of dead ECAL regions.

The FE electronics pass information to the data acquisition and trigger systems along different fibre optic pathways. In some of the cases of dead CCUs the data-read out optical fibre is damaged but the trigger information is still processed as normal. In these areas it is possible to estimate the amount of energy deposited in the trigger tower's associated crystals from the information in the trigger primitive. The recoverable areas of the ECAL are shown in figure 20.

The main drawbacks of this method are that trigger primitives store only the sum of the energy in their related crystals and saturate at a relatively low energy level of 64 GeV. The resulting information is therefore of lower granularity than the full crystal information, and any high-energy deposits

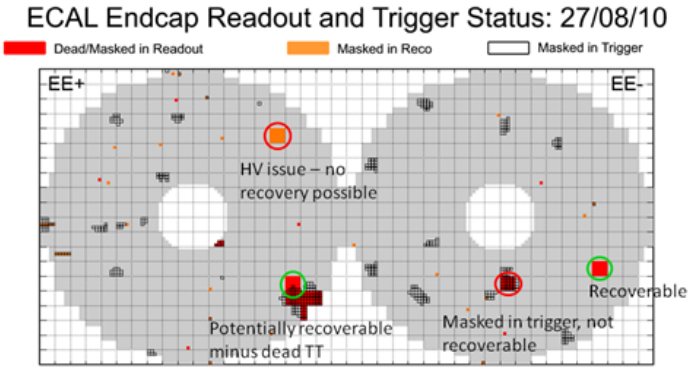

Figure 20: The regions in the ECAL endcap that can be recovered using trigger primitive information. Areas masked in trigger or with voltage supply problems that effect the trigger cannot be recovered using trigger information.

saturate the TP scale. In the latter case, the TP information provides a lower limit on the energy deposited in the crystals.

5.2 Existing Code and Required Alterations

In the ECAL barrel there already exists an algorithm to estimate the crystal energies in dead CCUs from an available TP. This is relatively straightforward in the barrel detector because each 5x5 supercrystal is associated to one trigger tower, and so a dead CCUs' energy directly correlates to the TP energy of a single tower. The front-end electronics of the endcap maintain the same 5x5 supercrystal layout as the barrel, but the trigger towers are arranged by angle, as seen in figure 6; this leads to complications with overlaps between trigger towers and CCUs. Figure 21 shows an example of the mismatched geometries; the energy from the associated towers must be summed and any known channels then removed from this in order to estimate the total deposited energy.

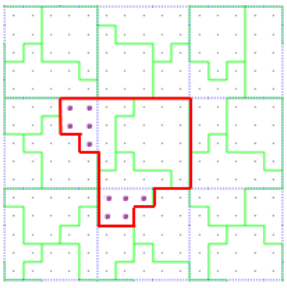

Figure 21: An example of mis-matched FE and trigger towers in the endcap. The central supercrystal is associated with all the trigger towers within the red outline. If it has a broken data fibre the purple crystals will still send data and must be subtracted from the total E_t of the red TPs.

There was already code meant to recover energy from trigger informa-

tion in the endcap, but it failed validation because it returned negative and [seemingly unreasonable] energy values. An example of the code being run over splash events can be seen in figure 22. To understand why this would happen, the study began by simply equating the energy parameters of the TPs and the crystal recHits associated with the crystals. The results can be seen in figure 23a, and do not give the 1:1 correlation one would expect from two systems being given the same information.

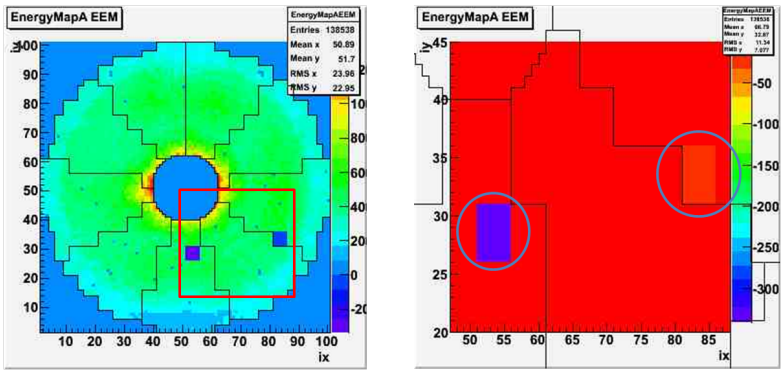

Figure 22: The original recovery algorithm acting on the splash events of run 144980. In regions that should be recovered from trigger information there is a large negative energy returned.

5.2.1 Energy Equivalence

The reason for this spread is that the type of energy stored in the two sources is different. The crystals record the total energy deposited in them by the electromagnetic particle, but the trigger is only passed the transverse energy. This gives rise to the dispersion seen in the energy comparison. The gap between the first two bands is attributed to the small gap in eta between the barrel and endcap crystals.

The energy mismatch is corrected in the barrel's recovery algorithm but due to the complicated numbering system in the endcap it was largely ignored. Because it is the actual crystal energy that is of interest to most analyses it was originally planned to convert the trigger primitive transverse energy to the full energy and run the recovery using this, but the position of the trigger towers meant that their exact pseudorapidity was ambiguous. The original code ran into this problem, but the bug was never fixed; instead an arbitrary (and incorrect) constant was attached to the conversion. For this reason it was decided that the energy of the towers would be summed as transverse energy, and any known channels would be converted to Et before

subtraction. Once the process is complete, the individual recHits are converted back to E and inserted into the collection of recHits. The conversion factor between E and Et is $1/cosh(\eta)$ where η is the pseudorapidity of the hit.

Once this was corrected, a more subtle problem appeared. The trigger towers are arranged in ieta rings of 72 towers each. This continues in the endcap, giving the towers the layout seen in figure 6. For the rings with $|ieta| > 26$ there is not enough space to physically fit 72 towers, so the number is halved to 36. For the purpose of read-out the existing towers are mirrored to give the complete 72 tower ring, but in doing so the TP energy in each is halved so that the total energy remains the same. This can clearly be seen as the third band in figure 23a. This must be taken into account when reconstructing energy in the very high presudorapidity regions of the detector. Once these two factors were included the comparison between crystal and TP energy was run again, this time giving a clear 1:1 correspondance with a slight spread in resolution, as seen in figure 23b.

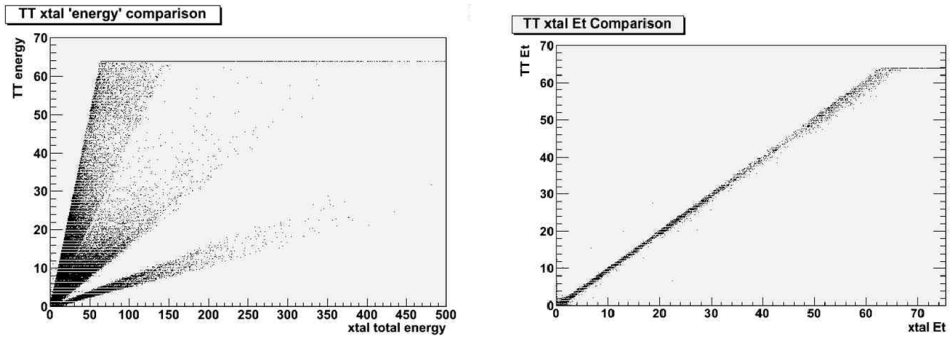

(a) Before correction (b) After correction

Figure 23: The two plots show the total trigger energy against the total energy registered by the crystals before and after correction for transverse energy. The three bands in 23a come from the barrel, endcap and $|ieta| > 26$ regions respectively.

5.2.2 Trigger Primitive Saturation

A further major issue with the original recovery algorithm was the saturation of the TPs. The TP compresses Et information into one byte - there are therefore a maximum of 256 possible energies that they can register. During the 2010 running this byte scaled linearly with the energy in the crystals with a constant of 0.25 GeV. This means that the largest value that can be

stored in a TP is 63.75 GeV any energy deposited in the crystals higher than this is truncated. This can be seen clearly in figure 23.

The saturation of the TPs causes two major problems in the recovery algorithm. The first is the loss of important physics signatures at higher energies for recovered items. For 2011 running the compressed energy constant has been changed to 0.5 GeV, giving a higher saturation of 127.5 GeV but lower resolution - a particularly important aspect at the far more common low energy events. There are ongoing discussions into the feasibility and benefit of changing to a non-linear energy scale which would give higher saturation and better resolutions at lower energies. The main barrier blocking this implementation is that it would require a change to the trigger firmware, a time-consuming and difficult process [9].

The second problem related to saturation only manifests itself in the endcap as it is a problem related to the overlapping geometries of various elements of the endcap detector. In the barrel the energy from one TP is divided between the 25 associated crystals directly. In the endcap any crystals with functioning data links are subtracted from the total of the TP energy. If the TPs are saturated then it is possible that the total of all known channels is higher than the sum of tower energies over the whole region. This results in a negative energy being computed for the recovered recHit, something that is unphysical.

A negative energy can also be returned from low energy CCUs. The TP E_t has a low resolution at low energies, so very small energy deposits or noise in the channels will not be registered by the trigger. It follows that a tower with $E_t = 0$ may be flagged for recovery with associated crystals that have a non-zero energy; either from a small hit or noise in the detector. This would lead to the recovery subtracting energy from the zero TP sum and returning a negative energy.

Originally this was overcome by only enacting the recovery on a region with a TP energy sum over a given value (for example 1.5 GeV), but this meant that any low energy recovery would not occur. It was replaced with a piece of code that removed crystals associated with a trigger tower containing zero energy from the recovery. This has the added benefit of more accurately determining the location of the hit; before, areas with no energy in the TP would still be assigned energy if there was a hit in any of the associated towers. This also has the effect of removing masked trigger towers from the recovery.

A completely new algorithm was created for estimating the energy in supercrystals containing saturated TPs. If the saturated TP has any known channels their energy is summed. The algorithm assumes a homogenous distribution of energy and estimates how much should be in the dead region.

If the new sum is greater than the saturation value then this value is returned; if not, the standard approach of subtracting the known channels from the TP energy is resumed. A flowchart of the complete recovery is given in A.

Before final insertion into the recHit collection, all recovered channels have a 'kRecovered' flag set. This tells anyone that wants to use the data for analysis that the channel has been recovered and to use it accordingly. Some analyses, such as the search for Higgs$\rightarrow \gamma\gamma$ require a high energy resolution and so may not wish to use recovered recHits. Searches that are highly sensitive to MET (such as SUSY and jet physics) will use the recovered information.

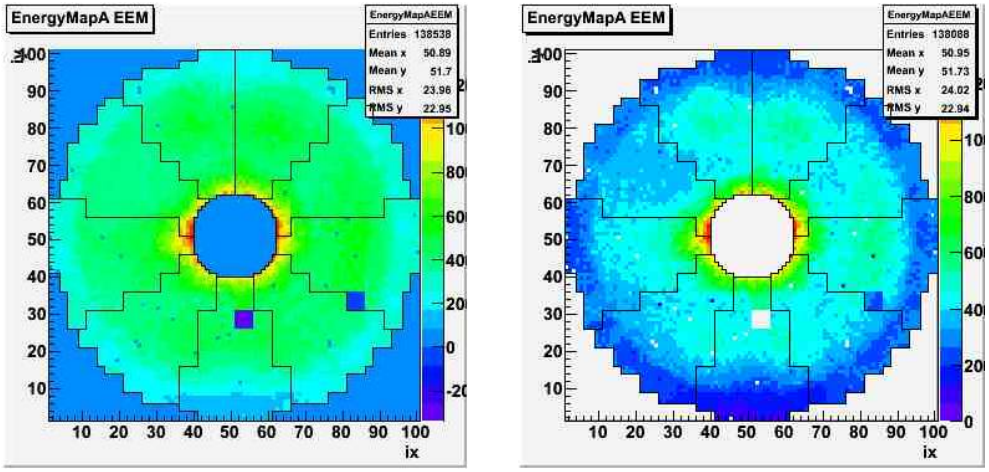

(a) Original recovery (b) New recovery

Figure 24: Two plots showing the recovery algorithm before and after the applied changes. The data is the 144980 splash run.

5.3 Validation and Implementation

Once the recovery code was fully implemented, a series of tests and validations were carried out. The first test was to turn on the recovery for a set of known data; the summer 2010 beam splash run 144980. The results of the recovery algorithm before and after the alterations can be seen in figure 24. The recovered CCU is noticeably lower in energy than its surroundings because there is one trigger tower completely contained within the FE card. Splash events lead to very high (above saturation level) energy depositions in the towers, and with no known channels the energy in the isolated tower can only be returned as 63.75 GeV. Because of the removal of masked trigger

towers the second dead CCU in the EE- endcap now returns zero after the recovery, preferable to the unphysical negative value returned by the original.

The recovery was then run over a large amount of fully functioning CCUs with both data and trigger information. The difference between crystal energy and the energy returned by the recovery was plotted and can be seen in figure 25a. There is a notable tail into the negative that is partly attributed to the problem with saturated TPs, and there is an overall offset in the figure of around -5 in the mean. To investigate this, the same data was plotted against the real crystal energy of the CCU. The resulting plot, seen in figure 25b, seems to imply that there is a calibration mismatch between the TPs and the crystal energies. A potential improvement to the recovery algorithm would be to test this over a large number of events and use the resulting information to correct for this offset.

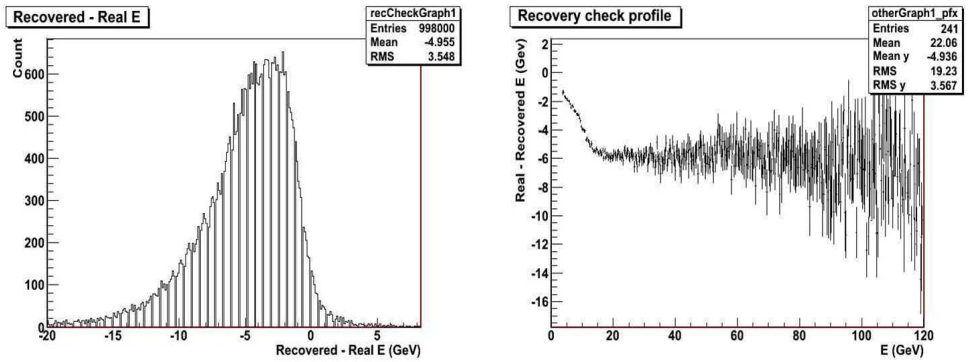

(a) Raw Results (b) Results plotted against Energy

Figure 25: The two tests carried out on the recovery algorithm. The recovery algorithm was tested on working channels and the difference between actual and recovered energy is plotted in 25a. 25b shows the same results plotted against the energy of the hit. The offset between the two implies a difference in calibration between the crystals and the trigger.

The code has now been fully validated and is fully implemented into the standard reconstruction procedure for CMS. An example of the recovery in action can be seen in figure 26 for recent beam splash events. The recovery algorithm now fully recovers the splash events; the hit in the recovered regime is the same colour as the crystals around it, indicative of a fully recovered area. Figure 27 shows how the recovery of dead FEs in the ECAL improves the tail on MET jet events.

Recent running has shown potential problems with more FE cards; new areas are having intermittent problems in data acquisition. This means the

trigger recovery algorithm may become more important in the near future for maintaining optimum performance of the ECAL.

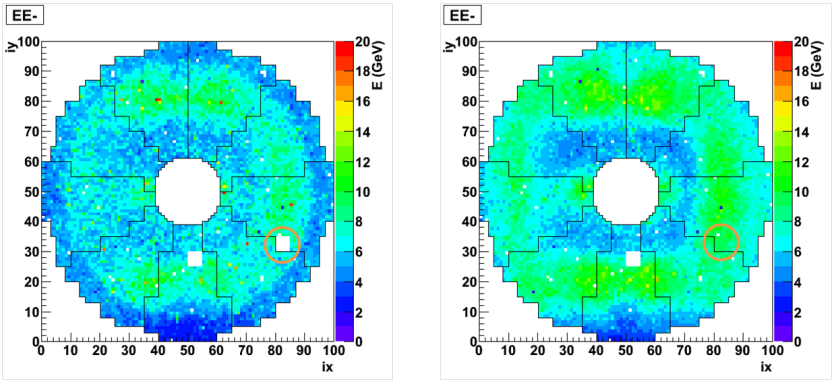

Figure 26: 2011 splash events both without and with the recovery turned on. Thanks to the developed algorithm the formerly dead CCU is recovered fully.

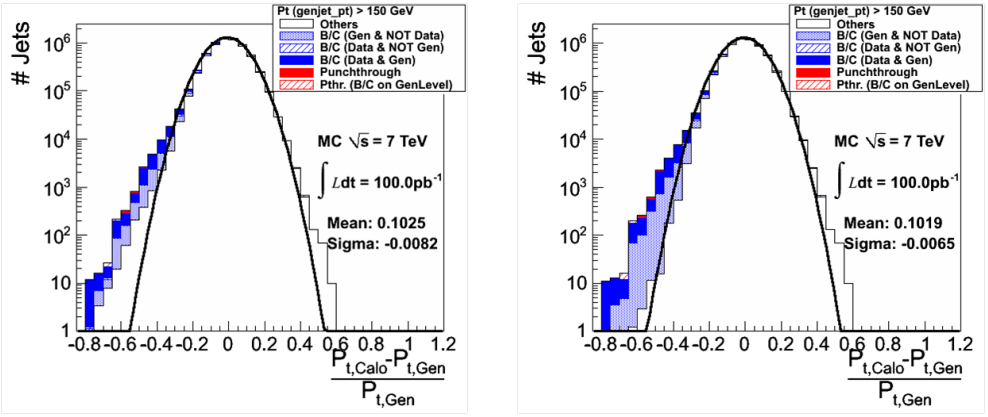

(a) Without Recovery (b) With Recovery

Figure 27: The figures show the number of MC jets versus the difference between the real (generated) energy of the jet and the energy that the detector reads off with and without a recovery for dead CCUs turned on. With the recovery turned off there is a larger tail on the number of events, meaning a larger false MET signal. Images courtesy of [9].

6 Use of the CMS ECAL Preshower in Dead Channel Recovery

6.1 The ECAL Preshower

Upon completion of the trigger recovery, it was suggested that the ECAL preshower detector could be used to make estimations of energy in dead regions in the EE with no trigger information. There are 100 channels in the endcap with no data or trigger read-out information, roughly half of the dead channels in the endcap. The ECAL preshower detector (ES) is a silicon tracker based detector that is positioned in front of the ECAL endcap. It covers the range $1.65 < |\eta| < 2.6$, meaning it covers [some percentage that I will work out?] of the endcap detector. The total detector is 20cm thick; comprising of two planes of orthogonal silicon strip detectors each fronted by a plane of lead. The lead absorbers, 2 and 1 radiation length in thickness respectively, begin electromagnetic showers in photons and the resulting electron-positron pairs can be measured in the silicon strip detectors. It is primarily used for pi^0/γ exclusion, as described in section 3.3.

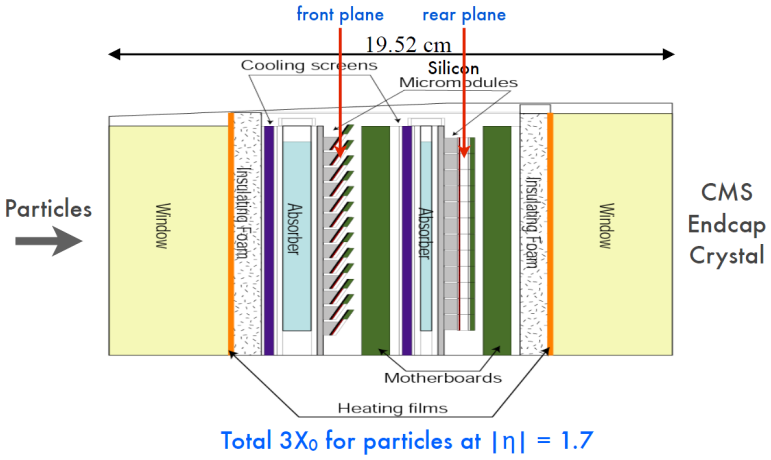

Figure 28: The internal structure of the ECAL preshower. The two planes of lead absorber of differing thickness can be seen in front of the crossed silicon strip detectors. As a silicon detector the preshower must be kept at a temperature between -10 to -15^0C. The lead tungstate of the ECAL crystal detector must be kept at a higher temperature and is very susceptible to slight changes - hence the internal cooling structure and outer heating layers of the preshower.

The preshower is only needed for pi^0/γ exclusion at high eta values and

so is only present in front of the endcap. At lower η values pion decays are of lower energy and the resulting photons have a greater angle of separation. In the barrel the crystals alone give a high enough granularity that they can distinguish the two types of events.

As the LHC increases its luminosity, collisions are likely to occur closer together as more protons are included in each proton bunch and the bunches are crossed more frequently; at high luminosity running CMS expects to find up to 20 primary vertices (or points of interaction) per bunch crossing. At such high values of pile-up it can be difficult to determine the primary vertex of all the individual particles. Originally it was envisaged to have a preshower for the barrel to measure the direction of photons. As they are electrically neutral photons do not leave information in the tracker, meaning that the only measurement of their position is in the ECAL. The preshower would have given a second tracking position allowing more accurate primary vertex location. Through extensive simulation it was determined that Higgs primary vertex location from the barrel alone would be sufficiently efficient that the heavy investment both financially and mechanically in a barrel preshower was not necessary.

Although the preshower is a silicon tracker detector and therefore designed with a high spatial resolution as its primary goal, it also takes a measure of the energy of the particles it tracks. Each layer of lead begins an electromagnetic shower, reducing the energy of the photon[24]. Produced electron-positron pairs cause small ionization currents in the silicon strip detectors which are detected and counted. As can be seen in figure 29 the native energy axis of the preshower is the number of Minimum Ionizing Particles (MIPs) that it detects. The number of MIPs produced in a shower is directly proportional to the energy of the shower's instigator, allowing the energy of the incident particle to be estimated. The energy of the produced pairs varies greatly, leading to a poor resolution on energy from the preshower. The different thickness of the lead planes means that the total energy deposited in the preshower is $E_{ES} = \alpha(E_{P1} + 0.7E_{P2})$, where α is a scaling constant and E_{P1} and E_{P2} are the energies registered in the first and second plane of the preshower respectively.

[24]The energy deposition in lead lowers the resolution of the whole ECAL but since energy measured in the preshower is proportional to that deposited in the lead the overall resolution is only affected slightly [7].

6.2 Preshower Recovery

6.2.1 Strip Association

The first problem encountered with any preshower recovery is the association between endcap crystals and preshower strips. There is a method available in CMSSW calorimeter geometry that allows any point in the calorimeter to have its location extrapolated back to the primary vertex, and another that will find the preshower strip that intersects this extended line. However, this only locates the one strip in either plane that the line directly intersects. In order to do analysis of energy measured in the two detectors a method for summing over all strips associated with the relevant crystal is necessary. This is done by finding the strips associated with all neighbouring crystals and then working out averages to find the highest and lowest strips necessary for summing in each plane. Once the beginning and end strips are located the method sums the energy across the range. Because the preshower does not cover the entire endcap detector there are areas where there are no associated strips to a crystal, causing further problems with an association method. By checking that the strips exist before attempting to sum over the area the algorithm avoids potential crashes.

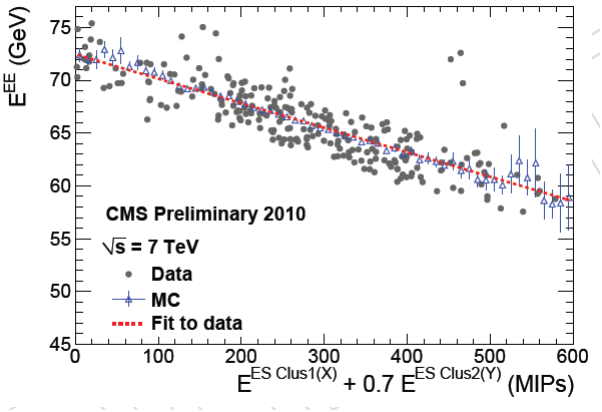

Figure 29: A plot showing the correspondence between energy deposition in EE and associated deposition in the EE for electron candidates with total energy between 70 and 75 GeV. The plot shows both MC and data. There is no 1:1 correspondence between the detectors, so knowing information in one cannot be used to directly estimate it in the other.

6.2.2 Energy Resolution

Figure 29 shows how the energy of a 70-75 GeV electron can be distributed between the preshower and the endcap. Any electron of a given energy has a distribution that looks like this; a negative correlation between detectors that sums to the total of the electron. By looking at a large number of electrons of varying energy a full distribution of E_{EE} can be found for a given value of E_{ES}. In a region where there is no information from the endcap - for example, a dead CCU - this distribution could then be used to estimate the energy that should appear in the preshower.

The algorithm used to sum energy was generalised so that it could be used to sum over 5x5 supercrystals.[25] From an electron gun Monte Carlo (MC) sample[26] the highest hit in the endcap was found and the 5x5 region around it located. This energy in this area was then summed for both the endcap crystals and the associated preshower strips. Figure [] shows the energy in the preshower plotted against the crystal energy in the endcap. It is possible to see the shape of the electron distribution in this plot, and as such there is no 1:1 correspondence as with the trigger information. A direct recovery is therefore not possible.

Figure [] is useful for examining how an electron hit correlates in the endcap and preshower, but it does not give much information on how the preshower responds to low energy hits or whether it registers energy when there is no electromagnetic particle incident. For this reason a large group of CCUs were selected in the endcap and examined throughout the MC sample for both total crystal and preshower energy[]. This gives a more representative view of how the preshower acts in normal running and more specifically how the information from one preshower area (i.e. the strips associated with a dead CCU) could be used to estimate the endcap energy deposits.

By taking large bin sizes for E_{ES} it is possible to create a 2D distribution for a given range of values of E_{ES}. Figure 31 gives an example of such a distribution. The problem with this method is that this distribution has a very low resolution - the fitted gaussian in figure 31 has a mean of 37.97 GeV with $\sigma = 17.91$. This means that for an ES hit of 0.75GeV only 50% of hits are in the range 25.9 to 50.05 GeV. This low accuracy makes a direct recovery like that with the TPs a less viable option. Instead it was suggested that

[25] As the preshower does not cover the entirety of the endcap it is possible that a desired supercrystal is only partially covered. In this case the algorithm finds the edges of the preshower and sums as much information as possible.

[26] A Monte Carlo sample is a randomly generated simulation of the system. An electron gun is a simulation of single electrons passing through CMS.

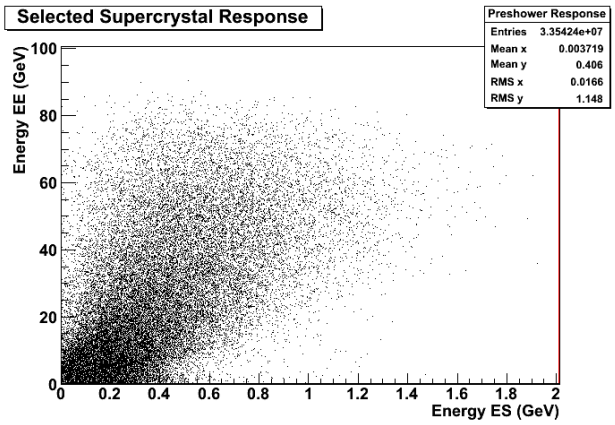

Figure 30: A plot showing the energy deposition in selected supercrystals in both the ES and EE for an electron gun MC sample. There is a correlation between the two sets of energy but a large spread.

a series of flags be implemented; if the preshower registers above a certain level in energy a flag is set which indicates that it is likely a large energy deposit existed in the associated dead endcap region. Whilst this will not provide direct information on the lost energy, it will allow a higher efficiency of rejection of false missing Et events. Any additional information that can be obtained regarding these events is considered very useful by those analysing high energy jets and missing transverse energy.

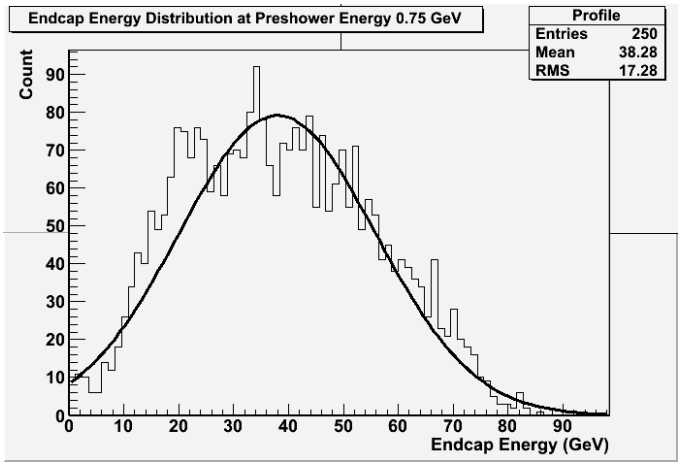

Figure 31: A profile of the distibution of endcap energies for ES hits of 0.75 GeV. A gaussian is fitted to the results.

6.3 Preshower Energy Flagging

For each point in the [infodump] histogram the efficiency and purity of a potential flag were calculated. As demonstrated in figure 32 the flag acts as a pivot on this histogram; the interesting information are the points with a higher E_{EE} (areas A and B), as these are more likely to be interesting physics events, whilst the flag only identifies points with preshower energy over E_{ES} (areas B and C). The efficiency of the flag is the fraction of hits above the given E_{EE} value that are flagged, i.e. $\frac{B}{A+B}$. The purity is the fraction of flagged hits that are above the desired E_{EE}, i.e. $\frac{B}{B+C}$. Together these give an indication of how well a flag placed at and above E_{ES} detects events above the desired E_{EE} threshold.

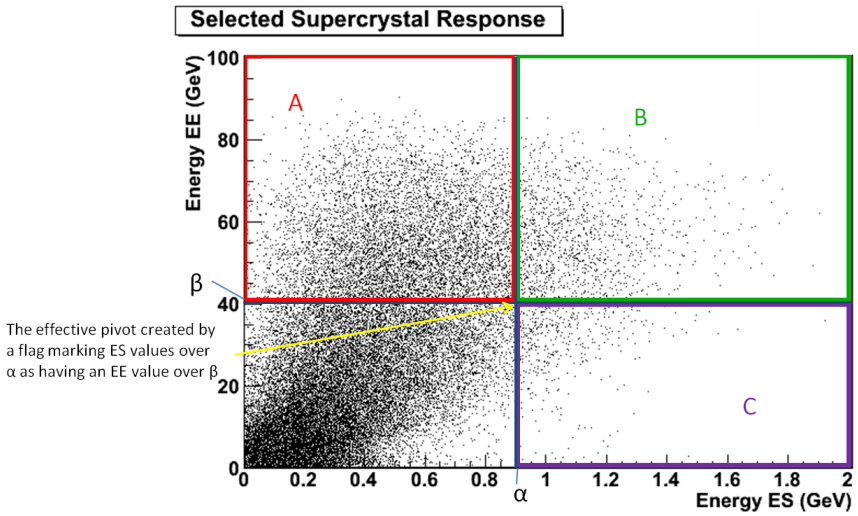

Figure 32: A figure demonstrating how the efficiency and purity of a flag marking ES deposits above α as EE hits above β were determined.

Efficiency and purity plots for a large electron MC sample are shown in figures 33a and 33d. The combination of the two in an efficiency $\times$ purity plot can be seen in figure 33g. By looking at the maximum value of this graph for varying values of E_{ES} suitable flag locations can be identified. Figure 34 is a cross-section of the three graphs for a fixed E_{EE} value of 25 GeV. It shows how the improvement in purity of a higher cut is mirrored by a drop in the efficiency. Figure 35 shows the optimal efficiency $\times$ purity for a flag at each value of E_{ES}, along with the efficiency and purity values for that cut. Finally figure 36 shows scatter plots of the optimal values of E_{EE} and E_{ES} for different cut parameters.

The above analysis was run over MC data for electrons of transverse

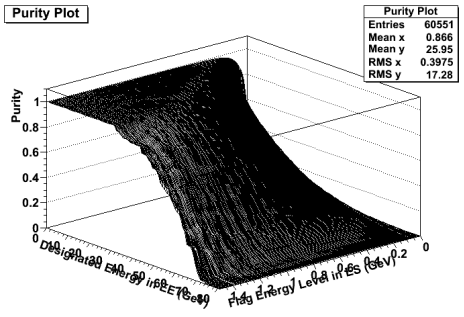

(b) 3D Plot

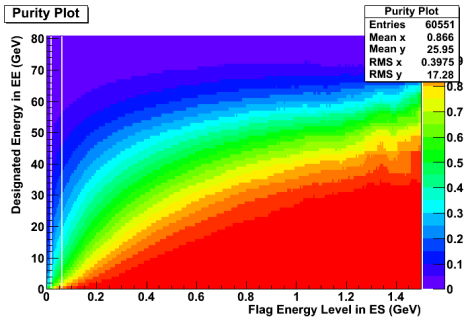

(c) 2D Plot

(c) Purity Graphs

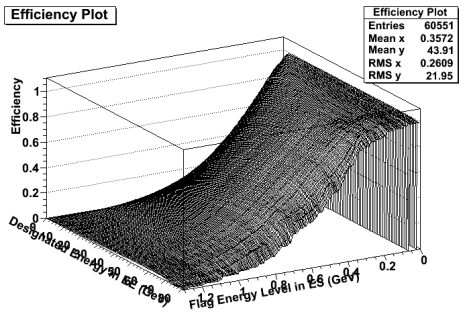

(e) 3D Plot

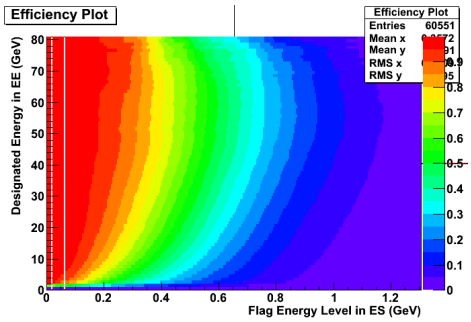

(f) 2D Plot

(f) Efficiency Graphs

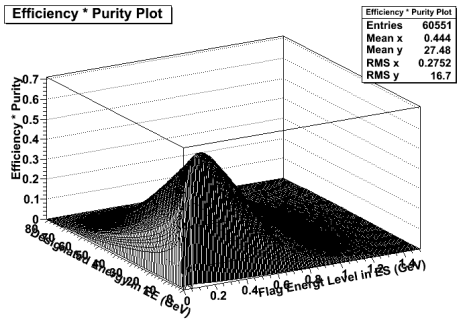

(h) 3D Plot

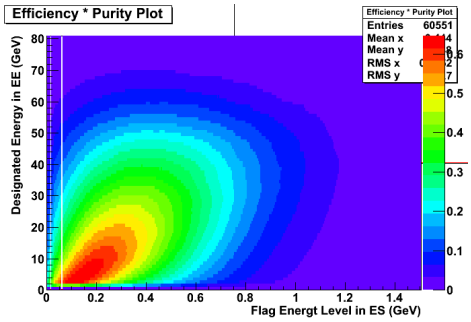

(i) 2D Plot

(i) Efficiency * Purity Graphs

Figure 33: Purity, efficiency and combined purity × efficiency plots for a potential flag for ES hits. Each point on the graph represents a flag for hits above the value of ES flagging hits in the endcap of higher energy deposition than EE.

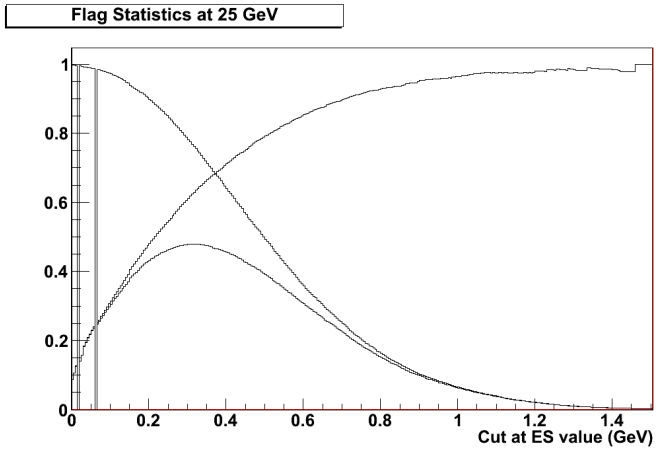

Figure 34: A slice through the plots as seen in figure 33 for an E_{EE} value of 25 GeV. For a higher value of E_{ES} the flag picks up less of the high endcap deposits, so the efficiency of the flag decreases, but the number of low endcap deposits falls faster so purity increases. The performance of the flag peaks between the two.

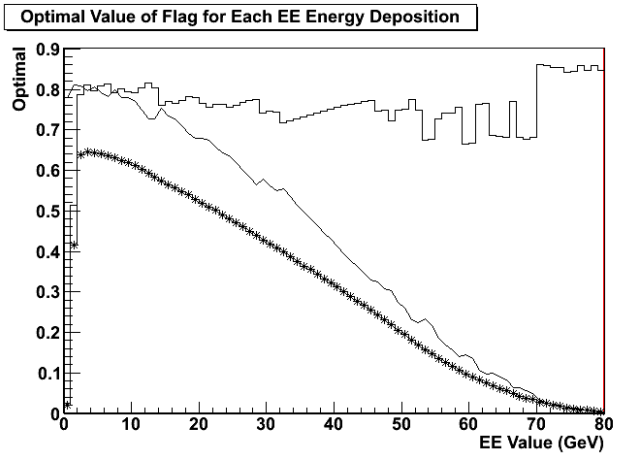

Figure 35: The plot shows the optimal value of efficiency × purity for each value of E_{EE}. For each of these values the related values of the efficiency and purity were found and also plotted. The efficiency of the optimal remains fairly constant at around 0.9 but the purity of flagging drops at higher values of E_{EE}.

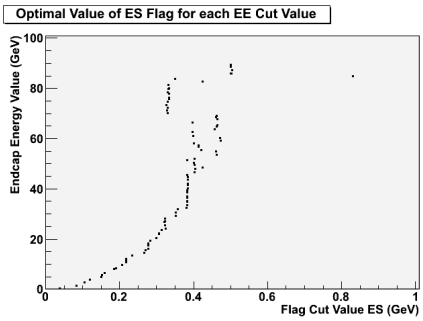

(a) Optimal ES flag for a given EE cut-off

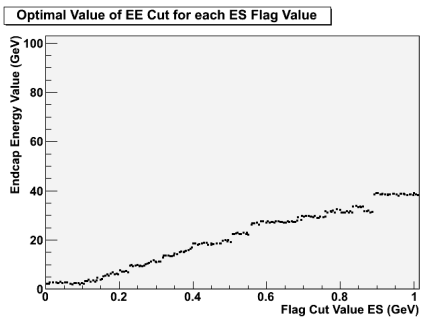

(b) Optimal EE cut-off for a given ES Flag

Figure 36: Figure 36a shows where a cut in E_{ES} should be made for optimal performance at that value of E_{EE}, whilst 36b shows which value in E_{EE} gives the best results for a flag at a value of E_{ES}.

momentum ranging from 2 to 20 GeV. Whilst this gives a good indication of how the preshower responds to electron hits a sample of MC data with a larger distribution of energies would be more suitable for a final determination of the potential flags. The flagging code now exists, is documented and is available to any analyses groups that wish to use it. Similarly to the trigger recovery, any analysis that requires a high energy resolution will not use the preshower flagging methods. For searches such as SUSY high MET jets however, the flags will provide a way to reject a large amount of the false signal produced by the dead regions of the endcap. As around half of the dead channels in the endcap have no usable trigger information this is a very useful alternative method of flagging events that are not genuine MET events.

7 Conclusions

The project examined various aspects of the CMS ECAL to determine and improve performance. Beam splash events were compared with contemporary laser monitoring data to identify any new dead channels, and to investigate the general status of the ECAL channels. 6 new dead channels were identified whilst 2 have appeared to recover. One channel with low laser response was also identified. The channels were added to the channel status map.

Methods for recovering energy in dead regions of the ECAL were then investigated. Where available, information stored in the trigger primitives can be used to estimate the amount of energy that would have been deposited in their associated crystals and an algorithm was developed for the endcap to do this. Problems caused by saturation were overcome and the recovery returns realistic, non-negative energies. It is now fully integrated in current CMSSW releases.

Finally the use of the ECAL preshower in energy recover was investigated. Whilst the poor energy resolution of the tracker detector means it is not useful in a direct recovery it can be used to flag events that are likely to have high energy depositions in dead regions and is therefore useful in cutting down the false missing transverse energy signal created generated by jets coinciding with dead regions in the ECAL.

A Trigger Recovery Flowchart

References

[1] A. Bornheim. The CMS ECAL Laser Monitoring System. *11th Topical Seminar On Innovative Particle and Radiation Detectors*, 2008.

[2] CMS Outreach. CMS Times. `http://cms.web.cern.ch/cms/Media/Publications/CMStimes/2011/05_02/index.html`, May 2011.

[3] CMS Collaboration. The Trigger and Data Acquisition project, Volume 1, Technical Design Report. Technical report, CERN-LHCC 2000-038, 2000.

[4] The CMS Collaboration. CMS ECAL TDR. *CERN/LHCC 97-33*, 1997.

[5] The CMS Collaboration. CMS Trigger TDR, 2000.

[6] The CMS Collaboration. Performance of the cms level-1 trigger during commissioning with cosmic ray muons and lhc beams. *JINST*, 5, 2010.

[7] E. Auffray et al. Beam tests of lead tungstate crystal matrices and a silicon strip preshower detector for the CMS electromagnetic calorimeter. *Nuclear Instructions and Methods*, A 412:223=237, 1998.

[8] P. Konar et al. How to look for supersymmetry under the LHC lamppost. *Physical Review Letters*, 105(22), 2010.

[9] U. Gebbert. Impact of TP saturation threshold on jets and MET. *DPG/Ph Egamma Meeting*, January 2011.

[10] S. L. Glashow. Partial-symmetries of weak interactions. *Nuclear Physics*, 22:579–588, 1961.

[11] H. Heath. CMS ECAL Endcap Channel Numbering. *CMS Internal Note 2006/027*, 2006.

[12] Q. Ingram. The lead tungstate electromagnetic calorimeter of CMS. `http://www.panic05.lanl.gov/abstracts/272/proc_ingram_272_20_Dec_05.pdf`.

[13] D. Petyt. ECAL Channel Status. `https://twiki.cern.ch/twiki/bin/viewauth/CMS/EcalChannelStatus`.

[14] The CMS Collaboration. *CMS Physics TDR, Volume I*, volume CERN-LHCC-2006-001. LHCC Printing, February 2006.

[15] S. Weinberg. A model of leptons. *Physical Review Letters*, 19:1264–1266, 1967.

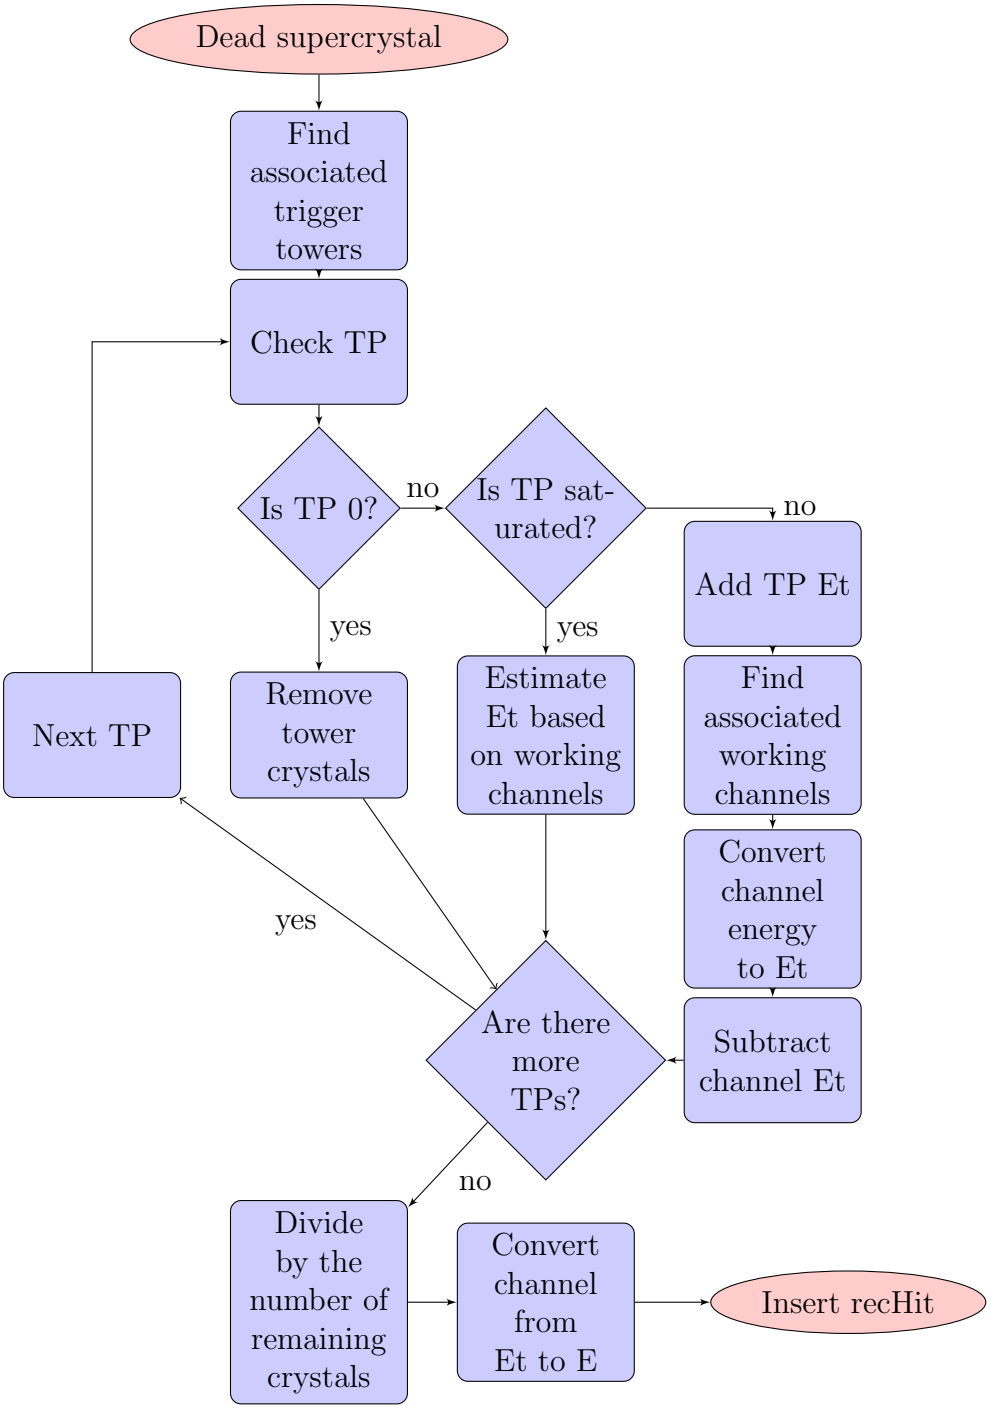

Figure 37: A flowchart showing the flow of the recovery program for a dead supercrystal.