

## Chapter 4

# Electromagnetic Calorimeter

### 4.1 Description of the ECAL

In this section, the layout, the crystals and the photodetectors of the Electromagnetic Calorimeter (ECAL) are described. The section ends with a description of the preshower detector which sits in front of the endcap crystals. Two important changes have occurred to the geometry and configuration since the ECAL TDR [5]. In the endcap the basic mechanical unit, the “supercrystal,” which was originally envisaged to hold  $6 \times 6$  crystals, is now a  $5 \times 5$  unit. The lateral dimensions of the endcap crystals have been increased such that the supercrystal remains little changed in size. This choice took advantage of the crystal producer’s ability to produce larger crystals, to reduce the channel count. Secondly, the option of a barrel preshower detector, envisaged for high-luminosity running only, has been dropped. This simplification allows more space to the tracker, but requires that the longitudinal vertices of  $H \rightarrow \gamma\gamma$  events be found with the reconstructed charged particle tracks in the event.

#### 4.1.1 The ECAL layout and geometry

The nominal geometry of the ECAL (the engineering specification) is simulated in detail in the GEANT4/OSCAR model. There are 36 identical supermodules, 18 in each half barrel, each covering  $20^\circ$  in  $\phi$ . The barrel is closed at each end by an endcap. In front of most of the fiducial region of each endcap is a preshower device. Figure 4.1 shows a transverse section through ECAL.

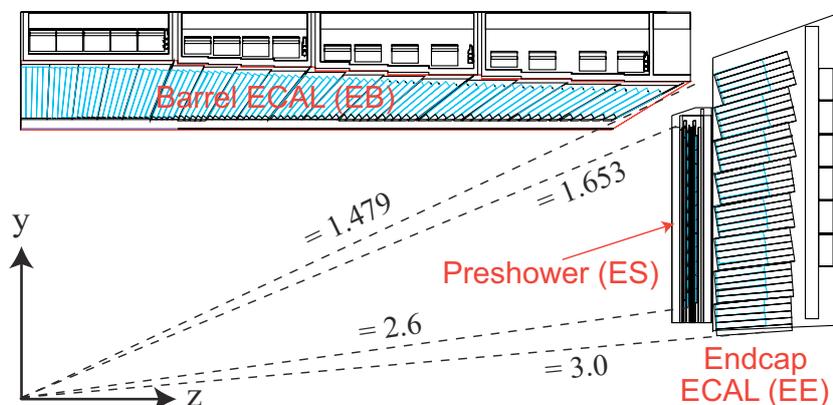


Figure 4.1: Transverse section through the ECAL, showing geometrical configuration.

The barrel part of the ECAL covers the pseudorapidity range  $|\eta| < 1.479$ . The barrel granularity is 360-fold in  $\phi$  and  $(2 \times 85)$ -fold in  $\eta$ , resulting in a total of 61 200 crystals. The truncated-pyramid shaped crystals are mounted in a quasi-projective geometry so that their axes make a small angle ( $3^\circ$ ) with the respect to the vector from the nominal interaction vertex, in both the  $\phi$  and  $\eta$  projections. The crystal cross-section corresponds to approximately  $0.0174 \times 0.0174^\circ$  in  $\eta$ - $\phi$  or  $22 \times 22 \text{ mm}^2$  at the front face of crystal, and  $26 \times 26 \text{ mm}^2$  at the rear face. The crystal length is 230 mm corresponding to  $25.8 X_0$ .

The centres of the front faces of the crystals in the supermodules are at a radius 1.29 m. The crystals are contained in a thin-walled glass-fibre alveola structures (“submodules,” as shown in Fig. CP 5) with 5 pairs of crystals (left and right reflections of a single shape) per submodule. The  $\eta$  extent of the submodule corresponds to a trigger tower. To reduce the number of different type of crystals, the crystals in each submodule have the same shape. There are 17 pairs of shapes. The submodules are assembled into modules and there are 4 modules in each supermodule separated by aluminium webs. The arrangement of the 4 modules in a supermodule can be seen in the photograph shown in Fig. 4.2.



Figure 4.2: Photograph of supermodule, showing modules.

The thermal screen and neutron moderator in front of the crystals are described in the model, as well as an approximate modelling of the electronics, thermal regulation system and mechanical structure behind the crystals.

The endcaps cover the rapidity range  $1.479 < |\eta| < 3.0$ . The longitudinal distance between the interaction point and the endcap envelop is 3144 mm in the simulation. This location takes account of the estimated shift toward the interaction point by 2.6 cm when the 4 T magnetic field is switched on. The endcap consists of identically shaped crystals grouped in mechanical units of  $5 \times 5$  crystals (supercrystals, or SCs) consisting of a carbon-fibre alveola structure. Each endcap is divided into 2 halves, or “Dees” (Fig. CP 6). Each Dee comprises 3662 crystals. These are contained in 138 standard SCs and 18 special partial supercrystals on the inner and outer circumference. The crystals and SCs are arranged in a rectangular

$x$ - $y$  grid, with the crystals pointing at a focus 1300 mm beyond the interaction point, so that the off-pointing angle varies with  $\eta$ . The crystals have a rear face cross section  $30 \times 30 \text{ mm}^2$ , a front face cross section  $28.62 \times 28.62 \text{ mm}^2$  and length of 220 mm ( $24.7 X_0$ ). The simulated end-cap geometry includes the aluminum backplate and approximate modelling of the support structure and electronics behind the crystals.

### 4.1.2 Lead tungstate crystals

Lead Tungstate crystals ( $\text{PbWO}_4$ ) are produced for CMS by the Bogoroditsk Techno-Chemical Plant in Russia and by the Shanghai Institute of Ceramics in China. The characteristics [119] of these production crystals make them an appropriate choice for operation at LHC. The high density ( $8.3 \text{ g/cm}^3$ ), short radiation length (0.89 cm) and small Molière radius (2.2 cm) results in a fine granularity and a compact calorimeter. The scintillation decay time is of the same order of magnitude as the LHC bunch crossing time: about 80% of the light is emitted in 25 ns. The light output is relatively low: about 4.5 photoelectrons per MeV are collected in both the avalanche photodiodes (APDs) and the vacuum phototriodes (VPTs), where the higher APD quantum efficiency is balanced by their smaller surface coverage on the back face of the crystal. The crystals emit blue-green scintillation light with a broad maximum at 420 nm [120, 121]. The light output variation with temperature,  $-1.9\%$  per  $^\circ\text{C}$  at  $18^\circ\text{C}$ , requires an ECAL cooling system capable of extracting the heat dissipated by the readout electronics and of keeping the crystal temperature stable within  $\pm 0.05^\circ\text{C}$  to preserve energy resolution. To exploit the total internal reflection for optimum light collection on the photodetector, the crystals are polished after machining. This is done on all but one side for EB crystals. For fully polished crystals, the truncated pyramidal shape makes the light collection non-uniform along the crystal length, and the needed uniformity [122] is achieved by depolishing one lateral face. In the EE, the light collection is naturally more uniform because the crystal geometry is nearly parallelepipedic, and just a mild tuning is being considered.

The crystals have to withstand the radiation levels and particle fluxes [5] anticipated throughout the duration of the experiment. Ionizing radiation produces absorption bands through the formation of colour centres due to oxygen vacancies and impurities in the lattice. The practical consequence is a wavelength-dependent loss of light transmission without changes to the scintillation mechanism, a damage which can be tracked and corrected for by monitoring the optical transparency with injected laser light. The damage reaches a dose-rate dependent equilibrium level which results from a balance between damage and recovery at  $18^\circ\text{C}$  [123, 120]. To ensure an adequate performance throughout LHC operation, the crystals are required to exhibit radiation hardness properties quantified as an induced light attenuation length always greater than 3 times the crystal length even when the damage is saturated. Hadrons have been measured to induce a specific, cumulative reduction of light transmission, but the extrapolation to LHC indicates that the damage will remain within limits required for good ECAL performance [124].

### 4.1.3 Photodetectors

#### 4.1.3.1 Barrel: avalanche photodiodes

In the barrel, the photodetectors are Hamamatsu type S8148 reverse structure (i.e., with the bulk n-type silicon behind the p-n junction) avalanche photodiodes (APDs) specially developed for the CMS ECAL. Each APD has an active area of  $5 \times 5 \text{ mm}^2$  and 2 are glued to the

back of each crystal. The APDs are sorted according to their operating voltage into bins 5 V wide, and then paired such that each pair has a mean gain of 50. The main properties of the APDs at gain 50 and 18°C are listed in Table 4.1.

Table 4.1: Properties of the APDs at gain 50 and 18°C.

Sensitive area	$5 \times 5 \text{ mm}^2$
Operating voltage	340–430 V
Breakdown voltage - operating voltage	$45 \pm 5 \text{ V}$
Quantum efficiency (430nm)	$75 \pm 2\%$
Capacitance	$80 \pm 2 \text{ pF}$
Excess noise factor	$2.1 \pm 0.2$
Effective thickness	$6 \pm 0.5 \mu\text{m}$
Series resistance	$< 10 \text{ Ohm}$
Voltage sensitivity of the gain (1/M.dM/dV)	$3.1 \pm 0.1\%/V$
Temperature sensitivity of the gain (1/M.dM/dT)	$-2.4 \pm 0.2\%/^\circ\text{C}$
Rise time	$< 2 \text{ ns}$
Dark current	$< 50 \text{ nA}$
Typical dark current	3 nA
Dark current after $2 \times 10^{12} \text{ n/cm}^2$	$5 \mu\text{A}$

The sensitivity to ionizing radiation traversing the APD (nuclear counter effect) is given by the effective thickness of  $6 \mu\text{m}$ , which translates into a signal from a minimum ionizing particle traversing an APD equivalent to about 100 MeV deposited in the  $\text{PbWO}_4$ .

For acceptance for the ECAL each APD was required to be fully depleted and to pass through a screening procedure involving 5 kGy of  $^{60}\text{Co}$  irradiation and 1 month of operation at 80°C. Each APD was tested to breakdown and required to show no significant noise increase up to a gain of 300. The screening and testing aimed to ensure reliable operation for 10 years under high luminosity LHC conditions for over 99% of the APDs installed in the ECAL [125]. Based on tests with hadron irradiations it is expected that the dark current after such operation will have risen to about  $5 \mu\text{A}$ , but that no other properties will have changed. Small samples of APDs were irradiated with a  $^{251}\text{Cf}$  source to monitor the effectiveness of the screening procedure selection of radiation resistant APDs.

#### 4.1.3.2 Endcap: vacuum phototriodes

In the endcaps, the photodetectors are vacuum phototriodes (VPTs) (type PMT188 from National Research Institute Electron in St. Petersburg). Vacuum phototriodes are photomultipliers having a single gain stage. These particular devices were developed specially for CMS and have an anode of very fine copper mesh ( $10 \mu\text{m}$  pitch) allowing them to operate in the 4 T magnetic field. Each VPT is 25 mm in diameter; one VPT is glued to the back of each crystal. The VPTs delivered to date have a mean quantum efficiency of the bialkali photocathode (SbKCs) of 22% at 430 nm, and a mean gain of 10.2 at 0 T.

When placed in a strong axial magnetic field, the response is slightly reduced and there is a modest variation of response with the angle of the VPT axis with respect to the field over the range of angles relevant to the CMS endcaps ( $6^\circ$  to  $26^\circ$ ). The mean response in a magnetic field of 4 T, with the VPT axis at  $15^\circ$  to the field direction, is 94.5% of that in zero field.

All VPTs are tested by the manufacturer before delivery, without an applied magnetic field. All VPTs are also tested on receipt by CMS to determine their response as a function of magnetic field up to 1.8 T. Each device is measured at a set of angles with respect to the applied field, spanning the range of angles covered by the endcaps. In addition, at least 10% of the tubes, selected at random, are also tested in a 4 T superconducting magnet, at a fixed angle of  $15^\circ$ , to verify satisfactory operation at the full field of CMS.

The estimated doses and particle fluences for 10 years of LHC operation are 0.5 kGy and  $5 \times 10^{13}$  n/cm<sup>2</sup> at the outer circumference of the endcaps and 20.0 kGy and  $7 \times 10^{14}$  n/cm<sup>2</sup> at  $|\eta| = 2.6$ . The VPTs are expected to be insensitive to such neutron fluences. The VPTs are required to show a loss of anode response of no more than 10% after a dose of 20.0 kGy.

#### 4.1.4 Preshower detector

The principal aim of the CMS Preshower detector (ES) is to identify neutral pions in the endcaps within a fiducial region  $1.653 < |\eta| < 2.6$ . It also helps the identification of electrons against minimum ionizing particles, and improves the position determination of electrons and photons with its superior granularity. The ES is a sampling calorimeter with 2 layers: lead radiators initiate electromagnetic showers from incoming photons/electrons whilst silicon strip sensors placed after each radiator measure the energy deposited and the transverse shower profiles.

The material thickness of the ES traversed at  $\eta = 1.653$  before reaching the first sensor plane is  $2 X_0$ , followed by a further  $1 X_0$  before reaching the second plane. Thus about 95% of single incident photons start showering before the second sensor plane. The orientation of the strips in the 2 planes is orthogonal. A major design consideration is that all lead is covered by silicon sensors, including the effects of shower spread, primary vertex spread etc. For optimum Level-1 trigger performance the profile of the outer edge of the lead should follow the shape of the ECAL crystals behind it. For the inner radius the effect of the exact profiling of the lead is far less critical, and thus a circular shape has been chosen. The lead planes are formed from 2 Dees that join close to the vertical axis.

Each silicon sensor measures  $63 \times 63$  mm<sup>2</sup>, with an active area of  $61 \times 61$  mm<sup>2</sup> divided into 32 strips (1.9 mm pitch). The nominal thickness of the silicon is  $320 \mu\text{m}$ ; a minimum ionizing particle (MIP) will deposit around 3.6 fC of charge in this thickness (normal incidence). The sensors are precisely glued to ceramic supports, which also support the front-end electronics assembly, and this is in turn glued to an aluminium tile that allows a 2 mm overlap of the active part of the sensors in the direction parallel to the strips.

The micromodules are placed on baseplates in groups of 7, 8 or 10 that, when coupled with an electronics motherboard placed above the micromodules, form a ladder. The spacing between silicon strips (at the edges) in adjacent micromodules within a ladder is 2.4 mm, whilst the spacing between strips in adjacent ladders is normally 2.5 mm; for the region where the 2 Dees join this spacing is increased to 3.0 mm.

The ladders are attached to the radiators in an  $x$ - $y$  configuration. Around 500 ladders are required, corresponding to a total of around 4300 micromodules and 137 000 individual read-out channels. Further details of the layout can be found in [126].

## 4.2 Readout electronics

This section describes the readout of the ECAL crystals and preshower detector. It also describes the generation and readout of the trigger data, the trigger primitives, passed to the Level-1 trigger; and the selection, by the Selective Readout Processor, of the channels for which the full precision data will be sent to the DAQ event builder. Finally the issue of the synchronization of the complete system is discussed.

The CMS-ECAL electronics can be divided into 2 subsystems. The on-detector electronics, composed of radiation-resistant circuits located just behind the crystals, and the off-detector electronics housed in underground counting rooms close to the experimental area. Both systems communicate through 90-m-long high-speed optical links, operated at 800 Mb/s.

### 4.2.1 Front end

The front end electronics of the ECAL must first amplify and shape the signal from the sensors, digitize the signal at 40 MHz, buffer the data until receipt of a Level-1 trigger, and then transmit the data to the off-detector electronics for insertion in the CMS data stream. In addition, the front end electronics uses the digitized data to calculate trigger primitives which are transmitted at 40 MHz to be used in the Level-1 Trigger decision.

The basic building block of the front end electronics is a group of 25 crystals (grouped in a  $5 \times 5$  geometry) a trigger tower in EB or supercrystal in EE. The trigger towers are composed of 4 different electronics boards. Each trigger tower contains a motherboard (MB), a Low Voltage Regulator Board (LVRB), 5 Very Front End (VFE) boards, and a Front End (FE) card.

The motherboards are located beneath the cooling system for the electronics and are used to route the signals from the photodetectors, APDs in EB or VPTs in EE, to the VFE cards, to distribute high voltage to the photodetectors, and to distribute Low Voltage to the VFE cards. Each supermodule contains 68 motherboards which each connect to the photodetectors of 25 crystals via kapton flexible-print cables. In addition, signals for temperature monitoring thermistors are routed from the sensor capsule to the VFE cards.

The LVRBs are connected directly to the external Low Voltage power supplies which sit in the CMS racks attached to the outside of the CMS iron yoke, approximately 20 metres from the supermodule. Each LVRB contains radiation-hard voltage regulators which provide the 2.5 V needed by the front end electronics. This regulated 2.5 V is distributed to the FE card by a small connector on the LVRB, and to the 5 VFE cards in a trigger tower via the motherboard.

Each VFE card contains amplification and digitization for the signals from 5 crystals. In order to achieve the low noise and high dynamic range requirements for the ECAL, 2 new radiation-hard ASICs were developed. The Multi Gain Pre-Amplifier (MGPA) contains 3 parallel gain stages which process the sensor signals. The signals from these 3 stages are routed to the AD41240, a custom designed Analog to Digital Converter which contains 4 channels each with 12 bits of information and an effective number of bits equal to 11. Three of the 4 ADCs on the AD41240 are used for each crystal. The ASIC digitizes these 3 inputs in parallel, determines whether each channel has saturated, and then outputs the data from the channel which has the highest gain, and was not saturated. The 3 MGPA gains are arranged so that the highest gain range has an amplification of a factor of 12 and a least significant bit of  $\simeq 35$  MeV. This range has a noise of around 40 MeV for the barrel and saturates at approximately 160 GeV. The other 2 gains are a factor of 6 and unity. The unity gain determines the

upper end of the dynamic range and saturates at approximately 1.7 TeV for the barrel, and 3.0 TeV for the endcap. The ADC is also designed so that once a range saturates, the ADC returns the next 5 samples without changing the gain. This prevents a second (return) gain change within the waveform of the digitized sample. The same chip is used for both the APDs and the VPTs with only a small change of external components on the VFE.

The signals from the 5 VFEs are collected on the FE card. Here they are buffered in a custom ASIC designed for the ECAL front end – the FENIX ASIC. Each FE card contains 7 FENIX ASICs. These are multipurpose and are used to

1. buffer the sampled data from the ADCs until a Level-1 trigger is received, and then transmit the data to the ECAL off-detector electronics. The amount of data which is transferred is configurable, but typically consists of the 10 ADC samples surrounding the beam crossing,
2. sum the samples from a group of 5 channels (called a strip) at 40 MHz,
3. sum the samples from the 5 strips (all 25 channels) for transmission to the calorimeter trigger.

The data are transmitted to the Level-1 trigger and the DAQ system using 2 opto-hybrids on each FE card. These hybrids contain radiation-tolerant laser diodes for electrical to optical conversion, and the CERN-developed radiation-hard GOL ASIC which provides parallel to serial conversion. The FE card also contains the clock distribution, and the control unit for allowing local configuration of all ASICs via I2C protocol.

Clock signals and configuration are distributed to the 68 FE cards using 8 independent electrical (i.e. the interconnect is not optical) token rings. These rings communicate optically to the off-detector electronics using a digital opto-hybrid (DOH) developed originally for the CMS tracker project. The DOHs are mounted on token ring link boards which then connect to the token rings. There are 2 rings which run in parallel and are designed to allow the recovery of the ring, in the event that one of the FE cards fails.

A fast test pulse, clocked by the ADC clock, can be injected into the input of each preamplifier by the MGPA. A simple 8 bit DAC allows injection of a sufficient range of charges to verify the functionality of the 3 gain ranges.

Fuller details of the front end electronics can be found in Refs. [127, 128, 129, 130].

### 4.2.2 Off-detector electronics

The ECAL off-detector readout and trigger architecture [131, 132] is illustrated schematically in Fig. 4.3. The system is composed of different electronic boards sitting in 18 VME-9U crates (the CCS, TCC and DCC modules) and in 1 VME-6U crate (the selective readout processor, SRP, system). The system serves both the DAQ and the trigger paths. In the DAQ path, the DCC performs data readout and data reduction based on the selective readout flags computed by the SRP system. In the trigger path, at each bunch crossing, trigger primitive generation started in the FE boards are finalized and synchronized in the TCC before transmission to the regional calorimeter trigger.

The clock and control system (CCS) board distributes the system clock, trigger and broadcast commands, configures the FE electronics and provides an interface to the trigger throttling

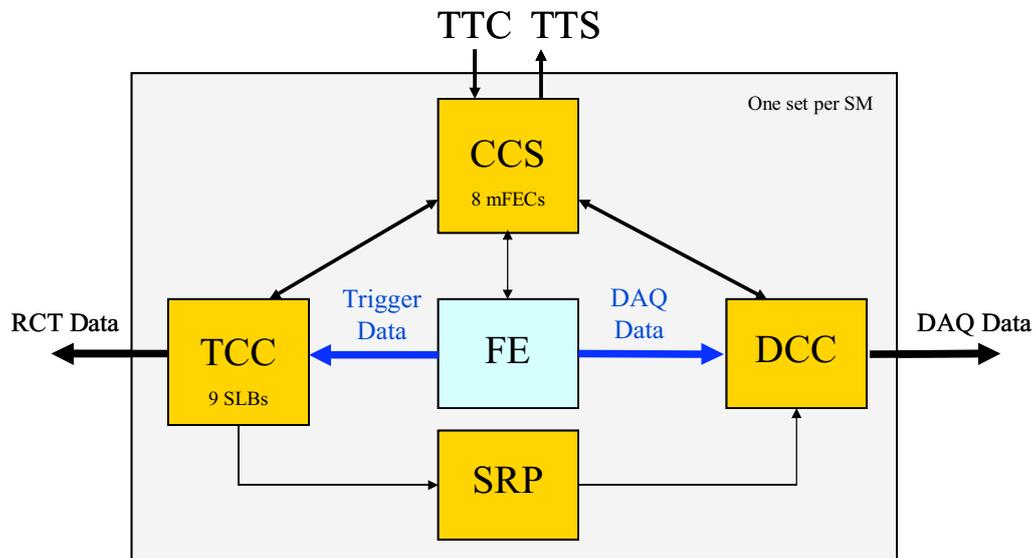


Figure 4.3: Schematic view of ECAL off-detector electronics.

system. The TTC signals are translated and encoded by suppression of clock edges and sent to the mFEC mezzanine cards. The mFEC interfaces optically with a FE token ring. The 8 mFECs of the CCS board control a supermodule. The TCC and the DCC cards in the off-detector crate receive the encoded TTC signals from the CCS card through a dedicated backplane.

The trigger concentration card (TCC) [133] main functionalities include the completion of the trigger primitive generation and their transmission to the synchronization and link board (SLB) mezzanines [134] at each bunch crossing, the classification of each trigger tower and its transmission to the Selective Readout Processor at each Level-1 trigger accept signal, and the storage of the trigger primitives during the Level-1 latency for subsequent reading by the DCC.

Each TCC collects trigger data from 68 FE boards in the barrel, corresponding to a supermodule, and from 48 FE boards in the endcaps corresponding to the inner or outer part of a  $20^\circ$  sector. In the endcaps, trigger primitive computation is completed in the TCCs, which must perform a mapping between the collected pseudo-strips trigger data from the different supercrystals and the associated trigger towers. The encoded trigger primitives (8 bits for the nonlinear representation of the trigger tower  $E_T$  plus 1 bit for the fine grain veto) are time aligned and sent to the regional trigger processors by the SLB. The trigger primitives are stored in the TCC during the Level-1 latency for subsequent reading by the DCC. In the barrel region a single TCC is interfaced with 1 DCC. In the endcap region, a DCC serves 4 TCCs covering a  $40^\circ$  sector

The data concentration card (DCC) [135, 136] is responsible for collecting crystal data from up to 68 FE boards. Two extra FE links are dedicated to the readout of laser monitoring data (pin diodes). The DCC also collects trigger data transmitted from the TCC modules and the selective readout flags transmitted from the SRP system. A data suppression factor near 20 is attained using a programmable selective readout algorithms. When operating in the selective readout mode the SRP flags indicate the level of suppression that must be

applied to the crystal data of a given FE readout. For the application of zero suppression, time samples pass through a finite impulse response filter with 6 consecutive positions and the result is compared to a threshold. If any time sample of the 6 has been digitized at a gain other than the maximum, then zero suppression is not applied to the channel.

Data integrity is checked, including verification of the event-fragment header, in particular the data synchronization check, verification of the event-fragment word count and verification of the event-fragment parity bits. Identified error conditions, triggered by input event-fragment checks, link errors, data timeouts or buffer memory overflows are flagged in the DCC error registers and incremented in associated error counters. Error conditions are flagged in the DCC event header.

Input and output memory occupancy is monitored to prevent buffer overflows. If a first occupancy level is reached, the Trigger Throttling System (TTS) signal “Warning Overflow” is issued, requesting a reduction of the trigger rate. In a second level a TTS signal “Busy” inhibits new triggers and empty events (events with just the header words and trailer) are stored. DCC events are transmitted to the central CMS DAQ using the S-LINK64 [41, 42] at a maximum data rate of 528 MB/s, while an average transmission data flow of 200 MB/s is expected after ECAL data reduction. Laser triggers (for crystal transparency monitoring) will occur with a programmable frequency and synchronously with the LHC gap. No data reduction is applied for these events, which are readout following a TTC test enable command. A VME memory is used for local DAQ, allowing VME access to physics events and laser events in spy mode.

The selective readout processor (SRP) [137] is responsible for the implementation of the selective readout algorithm. The system is composed by a single 6U-VME crate with twelve identical algorithm boards (AB). The AB computes the selective readout flags in different calorimeter partitions. The flags are composed by 3 bits, indicating the suppression level that must be applied to the associated readout units.

### 4.2.3 Trigger and readout

The ECAL data, in the form of trigger primitives, are sent to the Level-1 calorimeter trigger processor, for each bunch crossing. The trigger primitives each refer to a single trigger tower and consist of the summed transverse energy deposited in the tower, and a compactness bit, which characterizes the lateral extension of the electromagnetic shower. The accept signal, for accepted events, is returned from the global trigger in about 3  $\mu$ s. The selected events are read out through the data acquisition system to the Filter Farm where further rate reduction is performed using the full detector data.

The readout system is structured into sets of  $5 \times 5$  crystals. The FE card stores the data, in 256-clock cycles deep memory banks, awaiting a Level-1 trigger decision during at most 128 bunch crossings after the collision occurred. It implements most of the Trigger Primitives Generation (TPG) pipeline (Section 4.2.4). In the barrel, each FE is served by 3 optical links: 2 dedicated fibres for sending the data and trigger primitives respectively, and a third link which transmits the clock, control and Level-1 trigger signals.

In the barrel, these  $5 \times 5$  crystal sets correspond to the trigger towers. Each trigger tower is divided into 5  $\phi$ -oriented strips, whose energy deposits are summed by the FE board trigger pipeline to give the total transverse energy of the tower, called the main trigger primitive.

In the endcaps, the readout modularity maps onto the  $5 \times 5$  mechanical units (supercrystals). However the sizes of the trigger towers vary in order to approximately follow the  $\eta, \phi$  geometry of the HCAL and Level-1 trigger processor. The supercrystals are divided into groups of 5 contiguous crystal. These groups are of variable shape and referred to as pseudo-strips. The trigger towers are composed of several pseudo-strips and may extend over more than one supercrystal. Since the readout structure does not match the trigger structure, only the pseudo-strip summations are performed on the detector. The total transverse energy of the trigger tower is computed by the off-detector electronics. Hence, each endcap FE board is served by 7 optical links, 5 of them being used to transmit the trigger primitives.

After time alignment (Section 4.2.7) the ECAL trigger primitives are sent to the regional calorimeter trigger, via 10-m-long 1.2 Gb/s electrical cables, where together with HCAL trigger primitives, the electron/photon and jets candidates are computed as well as the total transverse energy.

#### 4.2.4 Trigger primitive generation

The TPG logic implemented on the FE boards combines the digitized samples delivered by the VFE boards to determine the trigger primitives and the bunch crossing to which they should be assigned. The logic must reconstruct the signal amplitude to be assigned to each bunch-crossing from the continuous stream of successive digitizations.

The TPG logic is implemented as a pipeline, operated at the LHC bunch crossing frequency. The trigger primitives are delivered to the regional calorimeter trigger after a constant latency of 52 clock cycles, of which 22 are used for the transmission over the optical fibres and cables. The signal processing performed in the VFE and FE barrel electronics has a total duration of only 17 clock cycles. The remaining part of the latency is mainly due to formatting and time alignment of the digital signals. Ideally, the output of this processing should be a stream of zeroes, unless there is a signal in the tower resulting from a bunch crossing exactly 17 clock cycles before. In this case the output is a word encoding the summed transverse energy in the tower together with the compactness bit. The endcap pipeline is split between the on-detector and off-detector electronics and implements very similar algorithms. The trigger primitives are expected to be delivered to the regional calorimeter trigger in 50 clock cycles in the endcap case.

The trigger primitive generation was studied using data taken in a test beam in November 2004. Electron beams with energies ranging from 15 to 120 GeV were used. In the test-beam experiment, electrons hit the detector at random times with respect to the 40 MHz clock used for sampling the analogue signals. Moreover, the parameters loaded in the chips (intercalibration constants and amplitude filter coefficients) were not optimized and only a limited amount of data were recorded. In order to assess the quality of the trigger primitives in the full range of energy and in LHC-like conditions, a functional model of the trigger primitives electronics and a Monte-Carlo simulation, based on GEANT4, has been developed to reproduce the online processing of electromagnetic showers. The actual noise characteristics of the VFE electronics were introduced by superimposing onto the simulated digital responses the recorded electronics responses of channels when the beam was off. The full simulation chain was tuned on the recorded data events in order to obtain a perfect matching of the simulation results with respect to the trigger primitives recorded online.

The response of the trigger primitives generation was first studied in 2 different geomet-

rical configurations, namely with the electron impact in the centre and in the corner of a trigger tower. For both configurations, the TPG shows good linearity versus the energy of the impinging electrons. Using the full simulation, the resolution expected for the ECAL electronics chain processing the trigger primitives was estimated. Figure 4.4 shows the relative resolution of the total transverse energy of an electromagnetic shower at the level of the ECAL trigger primitive generation. Two sets of points corresponding to the 2 geometrical configurations are presented in this figure. There is very little difference in the ECAL trigger primitives generated in these 2 extreme configurations.

Figure 4.5 shows the bunch crossing identification efficiency (BCID) versus the total transverse energy of the electromagnetic shower. The BCID for electromagnetic showers with energy  $\geq 1$  GeV is almost 100%. These characteristics will enable the time alignment of the trigger primitives (Section 4.2.7) to be performed without difficulty.

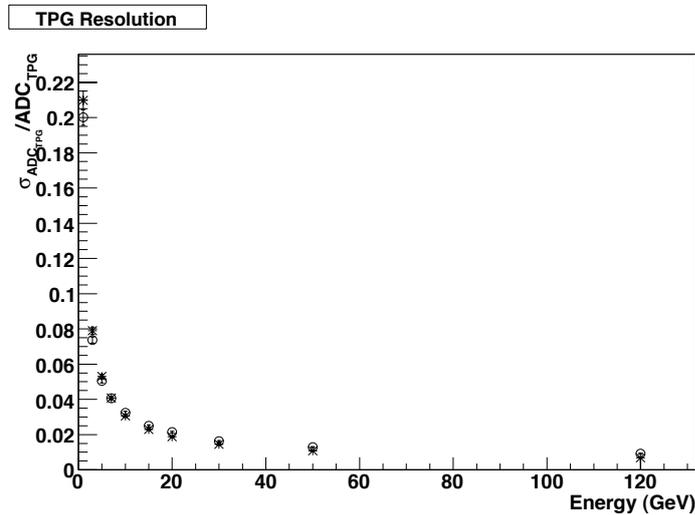


Figure 4.4: Trigger primitive resolution function (Monte Carlo simulation tuned using experimental data). Star markers refer to electrons impacting the centre of a trigger tower, while the circle markers refer to impacts at the corner of a trigger tower.

### 4.2.5 Selective readout

About 100 kB per event has been allocated for ECAL data. The full ECAL data for an event, if all channels are read out, exceeds this target by a factor of nearly 20. Reduction of the data volume, “selective readout”, can be performed by the Selective Readout Processor [132, 137] so that the suppression applied to a channel takes account of energy deposits in the vicinity. For the measure of the energy in a region, the trigger tower sums are used. In the barrel the readout modularity corresponds exactly to the  $5 \times 5$ -crystal trigger towers. In the endcap, the situation is more complex. The simplified and illustrative description below is given for the barrel case.

The selective readout algorithm classifies the trigger towers of the ECAL into 3 classes using the Level-1 trigger primitives. The energy deposited in each trigger tower is compared to 2

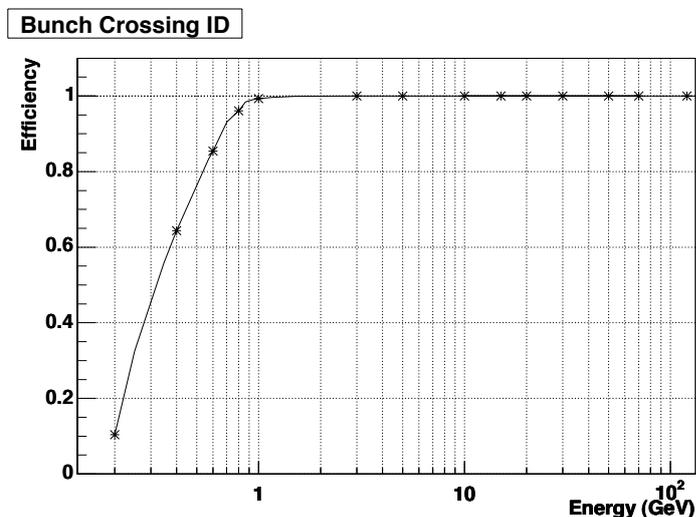


Figure 4.5: Bunch crossing assignment efficiency (Monte Carlo simulation tuned using experimental data).

thresholds. Trigger towers with an energy above the higher threshold are classified as high interest trigger towers, those with an energy between the 2 thresholds as medium interest, and those with an energy below the lower threshold as low interest trigger towers.

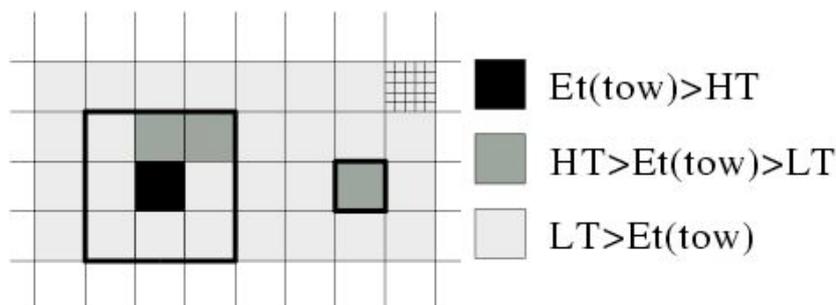


Figure 4.6: Selective readout regions. The figure illustrates the case of one trigger tower with a high transverse energy deposit (above the higher threshold HT, in black): the crystals of the  $3 \times 3$  trigger-tower matrix around this trigger tower are read without zero suppression threshold. The crystals of the trigger tower with a medium transverse energy deposit (between the higher and lower thresholds, HT and LT, in dark grey) are also read without zero suppression. The remaining towers, shown in grey, are read out with zero suppression

These classifications, illustrated in Fig. 4.6, can be used flexibly to implement a range of algorithms by using different thresholds to define the classes, and different suppression levels for the readout of the channels within each class. The algorithm currently used in the simulation provides adequate data reduction even at high luminosity. The algorithm functions as follows: if a trigger tower belongs to the high interest class ( $E_T > 5$  GeV) then the crystals of this trigger tower and of its neighbour trigger towers (225 crystals in the barrel case) are read with no zero suppression. If a trigger tower belongs to the medium interest

class ( $E_T > 2.5$  GeV), then the crystals of this trigger tower (25 crystals in the barrel case) are read with no suppression. If a trigger tower belongs to the low interest class and it is not the neighbour of a high interest trigger tower, then the crystals in it are read with zero suppression at about  $3\sigma_{\text{noise}}$ .

For debugging purpose, the selective readout can be deactivated and either a global zero suppression (same threshold for every channel) or no zero suppression applied. Even when the selective readout is not applied the selective readout flags are inserted into the data stream and can be used offline for debugging purpose.

#### 4.2.6 Preshower electronics

The preshower electronics can best be described in 2 parts: on-detector and off-detector, both of which are described in more detail in [138]. Figure 4.7 shows a schematic view of the complete electronics chain with this division marked.

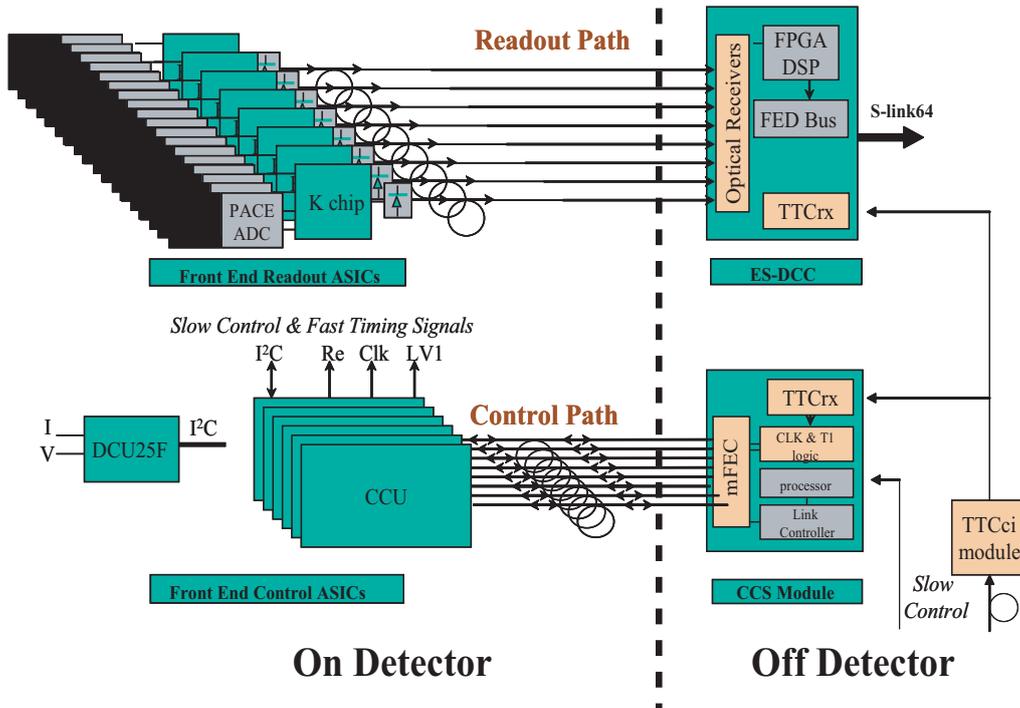


Figure 4.7: Preshower electronics chain.

##### 4.2.6.1 On-detector electronics

Each of the approximately 4300 32-channel silicon sensors in the Preshower detector is DC-coupled to a front-end ASIC (PACE3 [139]) that performs preamplification, signal shaping and voltage sampling into a high dynamic range 192-cell deep analogue memory at 40 MHz. For each Level-1 trigger received, 3 consecutive cells of the memory, corresponding to samples on the baseline, near the peak and after the peak, are read-out for all 32 channels through a 20 MHz multiplexer. The PACE3 has a switchable gain:

- Low gain: For normal physics running with a high dynamic range (0–1600 fC<sup>1</sup>) with a S/N of around 3 for a single MIP,
- High gain: For calibration purposes, with a reduced dynamic range (0–200 fC) but with a S/N approaching 10 for a single MIP.

The PACE3 has many programmable features, including an internal injection pulse that is used to intercalibrate the 2 gains, programmable biases and currents through 8-bit DACs, variable latency, etc. The DACs can be calibrated via multiplexed analogue outputs to DCU25F chips mounted on the same front-end hybrids as the PACE3. The sensor and hybrid are mounted on a ceramic support and an aluminium “tile” that allows overlapping of the sensors in one dimension, to form a “micromodule”. The micromodules are then mounted on the absorber plates. To avoid DC-coupling (and associated noise problems) of the sensor/electronics to the absorber planes, the aluminium tiles are constructed in 2 parts, with a glass-fibre insulation layer between.

Groups of 7, 8 or 10 micromodules are connected, via polyimide cables embedded in the front-end hybrids, to “system motherboards” (SMBs). The SMBs contain AD41240 12-bit ADCs for digitizing the analogue data from the PACE (1 ADC is used for 1 or 2 PACE3). The digital data from up to 4 PACE3 are then formatted and packaged by a second Preshower ASIC called the K-chip [140]. The K-chip also performs synchronization checks on the data, adds bunch/event counter information to the data packets, and transmits the data to the off-detector VME electronics via gigabit optical hybrids (GOH). The SMB also contains an implementation of the CMS Tracker control system, based on the CCU25 chip and its associated PLL25, QPLL, LVDSmux4p and LVDSbuf. The CCU25 provides a method to communicate to the K-chips and PACE3s via the I<sup>2</sup>C protocol, to program registers, etc. The combination of micromodules and an SMB, along with heatsinks and an aluminium baseplate, forms a “ladder”.

Groups of up to 12 ladders are connected via polyimide cables to form “control rings”. Each of these control rings communicates to the off-detector board “Clock and Control System” (CCS) modules via digital opto-hybrids (DOH) mounted on 2 of the SMBs in the control ring (for redundancy purposes). Each Preshower plane contains 12 control rings. There are thus 48 control rings for the complete Preshower system.

#### 4.2.6.2 Off-Detector electronics

The Preshower does not contribute to the Level-1 trigger. Consequently there are only 2 main parts to the off-detector electronics: the CCS module and the preshower data concentrator card - ES-DCC.

The CCS module is identical to that used by the EB and EE, except that for the ES only 3 or 4 of the 8 FEC mezzanines are mounted. One CCS module communicates with 3 or 4 control rings.

The ES-DCC reads the raw data from multiple GOH and performs bunch-crossing assignment, pedestal subtraction, common-mode rejection, charge reconstruction and zero suppression (threshold application), before formatting and sending sparsified data to the central DAQ via the S-LINK64 interface.

<sup>1</sup>A single minimum ionizing particle deposits around 3.6 fC in the 320 μm silicon sensors.

A “spy mode” will also be implemented, allowing full (non-zero-suppressed) data to be transferred to a local DAQ via the VME backplane or via the S-LINK64 (special condition, not during normal physics running). This will be used for startup conditions, pedestal runs, etc.

The design of the ES-DCC is ongoing, but will be largely based upon the ECAL DCC. The number of optical inputs may be reduced for the ES due to the larger amount of internal processing required in comparison to the EB and EE. The number of Slink64 connections reserved for the ES is 56, setting an upper limit to the number of ES-DCC.

## 4.2.7 Synchronization

In this section the procedures of synchronization of individual readout channels, as well as of the trigger primitives are described. A further issue related to timing, a procedure to verify the constancy of the phase of the sampling clock with respect to the signals is also discussed.

### 4.2.7.1 Readout synchronization

After digitization, data are held in a buffer pipeline awaiting a possible Level-1 Accept decision. If such a decision is given, the data are read out. The Level-1 Accept corresponding to a given bunch crossing has to match the data from the same bunch crossing. For the ECAL this means, in particular, that the sample corresponding to the maximum has to be propagated to the DAQ at a fixed position in the 10-sample timeframe which is read out. This can be achieved by setting the readout pointers in the pipeline memories inside the FE boards. The most important source of variation between the pointer values among the channels of a supermodule is due to the Level-1 signal distribution and is in the range of 2 to 3 units of the 40 MHz clock. The values of the readout pointers will be obtained during the commissioning period of the supermodules prior to their installation in the CMS detector. Laser pulses are sent to every channel and the corresponding 10-sample frames are analysed offline. The readout pointer values can be easily deduced from the position of the maximum of the digitized pulses inside the 10-sample streams. Correction of these settings will be calculated using the *a priori* knowledge of the different mean time of flight of the particles reaching the different parts of ECAL.

Because ECAL data is also sent to the Level-1 trigger, 2 possibilities exist for monitoring of the synchronization of the readout pipelines. They both are based on the generation of the Trigger Primitives. In the first method, a full readout is made of all trigger towers above a threshold of 1-2 GeV and a stream of 16 samples is extracted from the pipeline memory. An offline comparison of the recorded trigger primitives with the results of an emulation of the TPG (based on the recorded 16-sample stream) allows the readout synchronization to be verified. This is achieved by requesting that the results of the emulation match perfectly the online values of the TPG in the same trigger tower. A second method based on the same technique developed for the trigger primitives synchronization is described below.

### 4.2.7.2 Trigger Primitives synchronization

The CMS trigger system is a pipelined and synchronous system working at the LHC reference clock frequency. At each processing stage trigger data must be synchronized. The trigger primitives generated by the front-end electronics are collected by the Trigger Concentra-

tor Cards (TCC) which are part of the off-detector electronics. They are sent to the Regional Calorimeter Trigger which combines them in order to extract the total transverse energy in a pair of trigger towers. Non-negligible differences on the trigger primitives arrival time to the processors are introduced by different particle flight paths, different optical transmission fibre lengths and different phase lock delays in the electronic serializers. Thus a synchronization procedure is necessary. The synchronization is performed in the Synchronization and Link Board (SLB), a mezzanine board of the TCC. Trigger channel synchronization relies on the TTC Bunch Crossing Zero (BC0) broadcast command that can be adjusted relative to the LHC orbit signal, on a synchronization FIFO and on an accumulator histogram that reflects the LHC bunch crossing structure. A common hardwired control signal distribution guarantees aligned trigger data transmission through all high speed (1.2 Gb/s) transmission links. The SLB trigger channel accumulator histograms are analysed by an online readout program which is responsible for setting-up the trigger channel synchronization before data taking.

Trigger channel occupancy at each bunch crossing is dominated by minimum bias events. Simulations with a minimum-bias event rate corresponding to low luminosity ( $\mathcal{L} = 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ ) give a trigger channel occupancy of  $0.8 \times 10^{-4}$  at transverse energies higher than 1.0 GeV. Although lower transverse energies imply higher channel occupancy and faster accumulation histogramming, energy deposits near the tower noise have lower BCID efficiency. For energies higher than 1.0 GeV the BCID filter can be considered fully (100%) efficient. A Poisson generator was used to simulate particle collisions along the LHC orbit for this level of occupancy.

Channel synchronization relies on a correct identification of the LHC bunch structure. The channel alignment can be investigated using the discrete correlation function between the ideal bunch structure and the structure acquired in the SLB trigger channel accumulator. In Fig. 4.8 are shown the correlation curves after 2 minutes and 25 minutes of a LHC run for the barrel case. As expected the maximum of the correlation function is obtained for a correct alignment ( $\Delta BX=0$ ). The errors associated with the correlation points decrease with the acquisition time since higher statistics are acquired, therefore the correlation peak become more significant for long acquisition periods. The plot in Fig. 4.9 represents the correlation significance as a function of the accumulation time.

Correct alignments can be considered to be obtained if the significance is greater than 5. This is achieved after 5 minutes of LHC run. For high luminosity ( $\mathcal{L} = 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ ), the same performance would be obtained a hundred times faster, i.e. in 3 seconds. Using the same time alignment procedure for the endcap part would require an acquisition time longer because in this case the procedure is applied to the pseudo-strip signals. Calorimeter trigger channel synchronization using identification of the LHC bunch structure can be considered an efficient and fast process.

#### 4.2.7.3 Phase adjustment monitoring

As described in Section 4.2.1 each individual signal is sampled with a 40 MHz clock and digitized by an ADC after amplification. From the recorded set of sample values the amplitude of the corresponding signal can be reconstructed. The amplitude reconstruction procedure presented in Section 4.3 assumes the phase of the sampling clock with respect to the signal to be fixed and constant. A verification procedure to ensure this will be necessary.

The clock distribution system allows adjustment of the phase of the sampling clock by steps

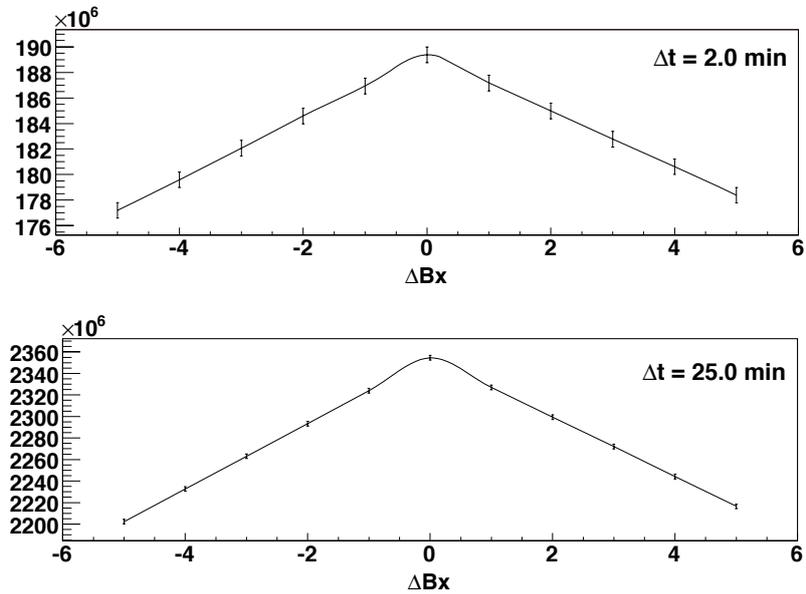


Figure 4.8: Correlation curves between ideal LHC bunch structure and the bunch structure acquired in the SLB trigger channel accumulator (after 2 and 25 minutes) as a function of the bunch crossing phase between the 2 structures.

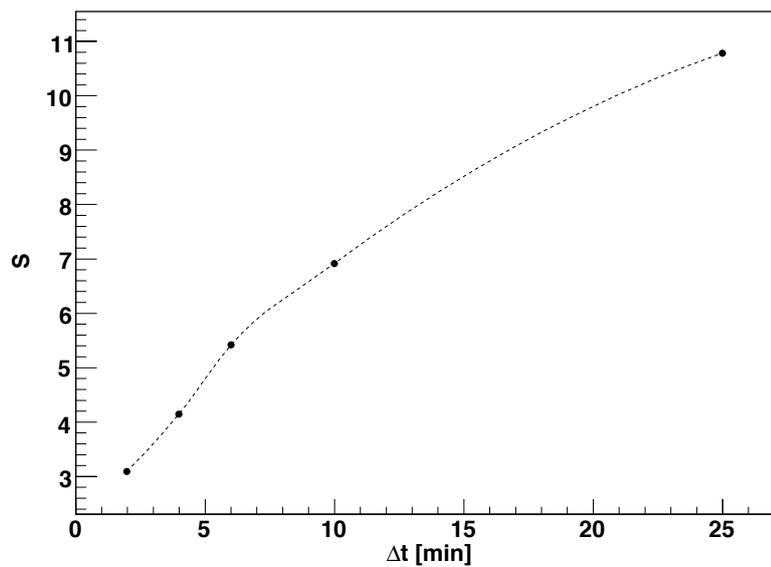


Figure 4.9: Peak significance for the correlation function after a given accumulated period (barrel case).

of 1 ns in a range of 25 ns for groups of 25 channels (corresponding to a trigger tower in the barrel and a supercrystal in the endcap).

The phase of signals in the installed detector, before any adjustment, will not be constant across the detector because of the geometry and architecture of the clock distribution, and because of the different time of flight from the interaction point to the detector. Laboratory measurements, performed on a large fraction of the final VFE electronics channels, show that the total spread of the peaking time distribution is less than 1 ns. The time of flight ranges from 3.9 ns for  $\eta = 0$  to 10.2 ns at  $\eta = 1.5$ .

Event-to-event phase variation caused by fluctuations on the time of flight because the longitudinal spread in the position of the primary vertex is less than  $\pm 0.2$  ns. Such fluctuations should not be of any consequence.

In order to adjust the sampling phase and subsequently verify that it remains constant, it is proposed to measure the phase using a digital filtering technique. This technique has been tested on beam data recorded in 2004. Figure 4.10 shows the distribution of the resolution of the difference between the time extracted using the digital filtering technique and the test beam trigger time (given by TDC) as a function of signal amplitude. The signal amplitude is expressed in units of the noise (i.e. around 40 MeV). It can be seen that a resolution less than 1 ns can be achieved for signals with an amplitude greater than 2 GeV. It has been shown that this time resolution is dominated by the test beam trigger timing measurement, which amounts to 0.6–0.7 ns.

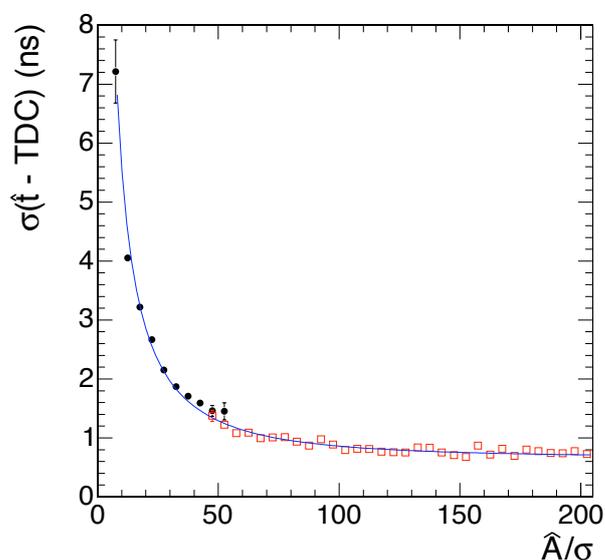


Figure 4.10: Resolution of the digital filtering technique versus the amplitude of the pulse expressed in units of noise.

### 4.3 Detector performance

In November 2004 a complete barrel supermodule (SM10) fully equipped with near final electronics was tested with high energy electrons ( $20 \leq E \leq 250$  GeV) in the CERN H4 beam line. The data taken have allowed verification of the performance of a completed element of

the final detector. Detailed studies have been made of amplitude reconstruction, electronics noise and energy resolution. The results of these studies is summarized in this section.

The ECAL performance for unconverted photons or for electrons in the test beam, is entirely dominated by contributions from effects which have to be added to the shower simulation. This is done when the electronics response is simulated—the “digitization” simulation. The results presented below make it clear that the values used for two of the most important of these contributions, noise and photo-statistics, are consistent with what is measured in data. The results also verify that the algorithm used to reconstruct the signal magnitude from a time frame of consecutive digitizations does not itself add a contribution to the energy resolution. The most important contribution to the intrinsic ECAL performance, intercalibration, is the subject of Section 4.4.

### 4.3.1 Amplitude reconstruction

The raw data for a single channel consists of a series of consecutive digitizations of the signal making up a time frame. The number of samples is adjustable ( $2+4n$ ) with a default of 10. The digitizations are made at the bunch crossing frequency of 40 MHz using an ADC clock that is locked to the LHC bunch structure. The timing of the signal will be adjusted in LHC running so that the signal pulse maximum corresponds to one of the samplings. The channel to channel variation, after this adjustment, is expected to be less than 1 ns.

The simplest method of reconstructing the amplitude is to take the sampling on the maximum as the measurement of the signal. However, one of the reasons for reading out a larger number of samples is to allow more sophisticated digital processing of the signal to reduce the noise contribution. The other reason is to enable identification of out-of-time (other bunch-crossing) pile-up. The signal amplitude is computed as a linear combination of discrete time samples as shown in equation 4.1.

$$\hat{A} = \sum_{i=0}^{i=N} w_i \times S_i \quad (4.1)$$

where  $w_i$  are the weights,  $S_i$  the time sample values in ADC counts and  $N$  is the number of samples used in the filtering. The weights are determined to minimize the noise contribution. Details of the mathematical formalism of the optimization procedure can be found in [141].

#### 4.3.1.1 Test beam studies

The shape of the signal pulse as a function of time is shown in Fig. 4.11. If there are no correlations of noise between time samples, the total noise level decreases with the number of samples used in the reconstruction. With the shape shown in Fig. 4.11, the use of 5 samples to reconstruct the amplitude should give noise level about 40% lower than the single sample noise. Considering more than 5 samples does not significantly improve the reduction of the noise.

Test beam data have been used to study the choice of amplitude reconstruction algorithm [142]. It is observed that the correlations between the samplings prevent the 40% noise reduction factor being achieved. The noise level obtained using 5 samples is very similar to that obtained from a single sample. There is also some noise correlation between nearby channels,

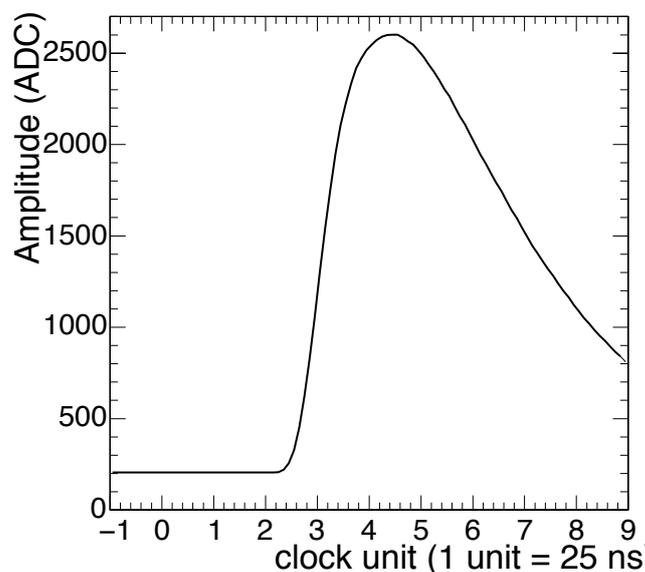


Figure 4.11: Profile of the signal pulse from a crystal of the SM10 supermodule using an electron beam of 120 GeV.

so that the noise measured in a sum of 25 channels (a typical cluster size used for ECAL shower reconstruction) is, in fact, slightly worse when the reconstruction uses 5 samples rather than a single sample.

This correlation is due to the presence of a small level of low frequency noise (i.e., noise at a much lower frequency than the 40 MHz digitization) that remains despite the grounding and screening. Because of this, the signal pulse sits on a baseline which varies from event to event. The weights can be determined so as to remove this varying pedestal. This is best done by additionally using samplings taken before the signal pulse together with the signal samplings. Using 3 samplings taken before the signal pulse, and 5 samplings on the signal pulse reduces the total noise seen in a sum of 25 channels by 20% as compared to what is measured by reconstruction with a single sampling followed by pedestal subtraction using an average pedestal value. With this reconstruction the total noise seen in a sum of 25 channels is almost exactly 5 times the noise seen in a single channel showing that the coherent noise has been very effectively removed. The average value of the noise seen in 1000 channels in supermodule tested in the H4 test beam in October 2004 is slightly less than 40 MeV. The implementation of a noise covariance matrix, and its use in the derivation of the weights does not improve the noise performance.

The optimization of the weights to minimize the noise contribution is insensitive to the small variation of signal pulse shapes present in the supermodule tested. Thus it is possible to use a single set of weights, which can be determined before LHC startup, for the amplitude reconstruction of all channels.

Small variations of the pulse maximum timing from channel to channel generate a small bias in the reconstructed amplitude which, in the absence of any precise overall absolute calibration, are simply swallowed by the intercalibration of the calorimeter. Systematic variation of the timing of maximum might result in an unacceptable variation of the channel response with time. A shift of 1 ns corresponds to a response variation of about 0.02%, but a shift of

3 ns corresponds to 0.2%. To prevent degradation of the resolution induced by such variation, the time of maximum must be carefully monitored.

The reconstruction of the signal is affected by jitter on the timing. The expected jitter in LHC running is about 0.2 ns, dominated by the longitudinal spread of the interaction vertex. The resolution is not degraded by such a small uncertainty on the signal timing.

#### 4.3.1.2 Use of other MGPA gain ranges

For signal pulses resulting from very large energy deposits in a single crystal ( $E > 160$  GeV in the barrel, and  $E > 300$  GeV in the endcap) the data-frame contains samples in different MGPA gain ranges. Reconstruction using the peak sample alone greatly simplifies the treatment of such data-frames, and the noise is negligible at these energies. The relative calibration of the different ranges must be determined precisely to avoid any degradation of the resolution in gain transition and introducing a nonlinearity of the response.

Further details concerning the amplitude reconstruction, the determination of gain ratios, the effect of pile-up, as well as the complications of reconstructing test beam data where the ADC clock is asynchronous with respect to the signal pulse, are given in [142].

### 4.3.2 Energy resolution in test beam

#### 4.3.2.1 Resolution for central impact

Showers in the ECAL are reconstructed by building clusters of crystals. In the test beam the best performance is obtained using a simple sliding window centred on the crystal having the maximum energy, summing  $3 \times 3$  or  $5 \times 5$  crystals. The energy contained in such a cluster varies with the shower position. The simplest verification of the performance of the ECAL is provided by a measurement of the energy resolution when the incident electrons are restricted, using the position measured in a set of beam hodoscope detectors, to a small region so that this variation is negligible. Figure 4.12 shows examples of the distribution of the reconstructed energy, fitted by a Gaussian to obtain the energy resolution  $\sigma_E/E$ .

#### 4.3.2.2 Resolution for uniform impact

The test beam data was taken using a trigger covering a  $20 \times 20$  mm<sup>2</sup> area, roughly matching the beam size ( $\sigma \approx 10$  mm) and only slightly smaller than the crystal granularity. When reconstructing showers from incident electrons spread over this area the energy resolution can be improved by making a correction for the varying containment as a function of the shower position as measured by the ECAL. In the analysis reported here, the position was measured independently in the 2 lateral coordinates using the logarithm of the ratio of crystal energies on either side of a reference line defined by the crystal edge. A single correction function, parameterized from the data, was used for all energies, and for all regions in the supermodule. Excellent performance is obtained even, for example, for runs where the supermodule was positioned such that the beam (and the  $20 \times 20$  mm<sup>2</sup> trigger) was centred on the corner of a crystal. Figure 4.13 shows an example of the energy reconstructed in such a case, together with the energy distribution before applying the correction.

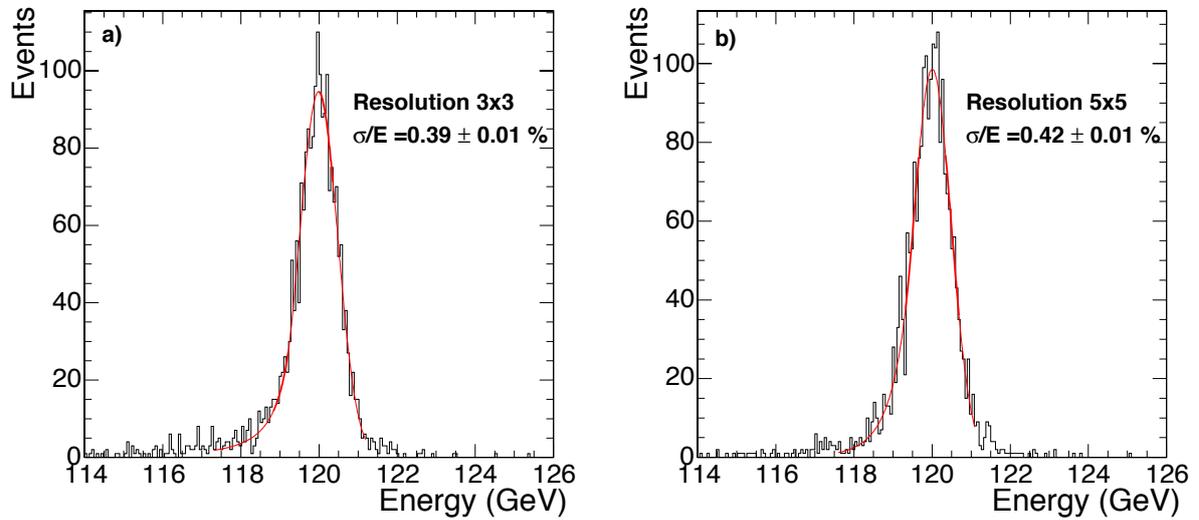


Figure 4.12: Distribution of the energy reconstructed in (a) a  $3 \times 3$  cluster and (b) a  $5 \times 5$  cluster, when 120 GeV electrons are incident in a  $4 \times 4 \text{ mm}^2$  region.

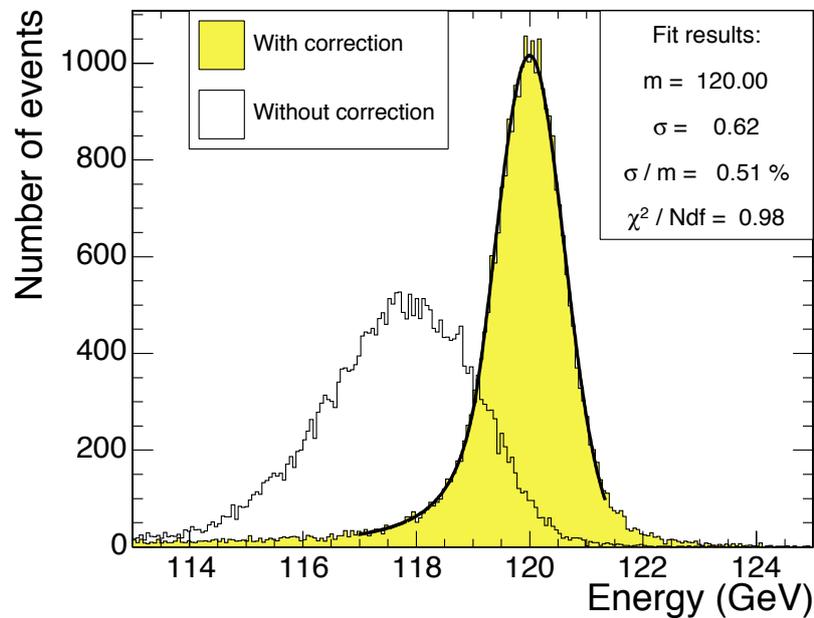


Figure 4.13: Distribution of the energy reconstructed in a  $3 \times 3$  cluster when the beam of 120 GeV electrons is centred on a crystal corner. The outline histograms show the energy reconstructed before applying the energy containment function.

### 4.3.2.3 Energy resolution as a function of energy

The energy resolution, measured by fitting a Gaussian function to the reconstructed energy distributions, has been parameterized as a function of energy,

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2, \quad (4.2)$$

where  $S$  is the stochastic term,  $N$  the noise and  $C$  the constant term. As presented earlier, Fig. 1.7 shows an example of the energy resolution as function of energy, together with the fitted function. Similar results are obtained for sets of data (energy scans) taken throughout the supermodule. The upper series of points are the resolution values obtained from fits to the energy distribution obtained for unselected events taken with a  $20 \times 20 \text{ mm}^2$  trigger and reconstructed using the containment correction described in the previous section. The lower series of points are obtained from events selected to fall within a  $4 \times 4 \text{ mm}^2$  region. Due to the tight selection, the statistical error on the measured energy resolution in the lower set is rather large for some data points.

## 4.4 Calibration and alignment

Calibration is a severe technical challenge for the operation of the CMS ECAL. It is naturally seen as composed of a global component, giving the absolute energy scale, and a channel-to-channel relative component, which is thereafter referred to as intercalibration.

Many small effects which are negligible at low precision need to be treated with care as the level of precision of few per mil is approached. The essential issue is stability in both time and space, so that showers in different locations in the ECAL in data recorded at different times are accurately related to each other.

The main source of channel-to-channel response variation in the barrel is the crystal-to-crystal variation of scintillation light yield which has an RMS of  $\approx 8\%$ . In the endcap the VPT signal yield, the product of the gain, quantum efficiency and photocathode area, has an RMS variation of almost 25%. The nature and technology of the ECAL provides no convenient or *a priori* way of intercalibrating the channels and the target precision can only be achieved using physics events. Over the period of time in which the physics events used to provide an intercalibration are taken the response must remain stable and constant to high precision. Where there is a source of significant variation, as in the case of the changes in crystal transparency caused by irradiation and subsequent annealing, the variation must be precisely tracked by an independent measurement. The changes in crystal transparency are tracked and corrected using the laser monitoring system.

The final goal of the calibration strategy is to achieve the most accurate energy measurement for electron and photons. Schematically the reconstructed energy might be decomposed to 3 factors,

$$E_{e,\gamma} = G \times \mathcal{F} \times \sum_i c_i \times A_i, \quad (4.3)$$

where  $G$  is a global absolute scale. The function  $\mathcal{F}$  is a correction function depending on the type of particle, its position, its momentum and of the clustering algorithm used. The

$c_i$  factors are the intercalibration coefficients while the  $A_i$  are the signal amplitudes, in ADC counts, which are summed over the clustered crystals.

The correction function  $\mathcal{F}$ , discussed in more detail in Chapter 10, corrects for energy loss due to bremsstrahlung and various containment variations. In the current working model,  $\mathcal{F} = 1$  is chosen for the reference  $5 \times 5$  crystal shower reconstruction algorithm, used for unconverted photons, or for electrons in the test beam. Different reconstruction algorithms are used to estimate the energy of different electromagnetic objects, i.e., unconverted photons, electrons and converted photons, each of them having their own correction functions. At present these “algorithmic” corrections are obtained from the simulated data by accessing the generated parameters of the Monte Carlo simulation. For some of the corrections, for example the containment corrections, this is an acceptable procedure provided that test beam data is used to verify the simulation, so that, in effect, the simulation is being used only as a means of interpolating and extrapolating from data taken in the test beam. In other cases, where the test beam provides no useful information, for example in issues related to conversions and bremsstrahlung radiation in the tracker material, it is necessary to use information that can be obtained from data taken *in situ* with the running detector. Two particularly useful channels which can be used to obtain such information are under investigation:  $Z \rightarrow ee$ , and  $Z \rightarrow \mu\mu\gamma$  (the photon coming from inner bremsstrahlung). The latter, in the case of unconverted photons, is also able to set the global scale  $G$ .

Preliminary estimates of the intercalibration coefficients  $c_i$  are obtained from laboratory measurements of crystal light yield, test beam precalibration of some supermodules, and the commissioning of further supermodules with cosmic rays. Imposing the  $\phi$ -independence of energy deposited in the calorimeter can be used to rapidly improve on this start-up intercalibration for fixed  $\eta$  regions. The method of intercalibration with physics events that has been investigated in the most detail, uses the momentum of the abundant electrons as measured in the tracker, mainly from  $W \rightarrow e\nu$ , which have a similar  $p_T$  to the photons of the benchmark channel  $H \rightarrow \gamma\gamma$ . A complementary method, not relying on the momentum measurement, is based on  $\pi^0 \rightarrow \gamma\gamma$  and  $\eta \rightarrow \gamma\gamma$  mass reconstruction.

The achievable precision, the methods used, and the path to full precision will be a function of time and available luminosity.

#### 4.4.1 Intercalibration at start-up

##### 4.4.1.1 Laboratory measurements

The calorimeter is being assembled in 2 regional centres: at CERN and at INFN-ENEA Casaccia near Rome. During the assembly phase, all the detector components are characterized [143, 144] and the data are saved in the construction database. It is thus possible to predict the calibration  $c_i$  of each channel  $i$  using the laboratory measurements as

$$\frac{1}{c_i} \propto LY \cdot \varepsilon_Q \cdot c_{ele} \cdot M, \quad (4.4)$$

where LY is the Light Yield of the crystals, M and  $\varepsilon_Q$  are respectively the gain and quantum efficiency of the photo-detectors and  $c_{ele}$  is the calibration of the electronics chain. The crystal LY is measured in the laboratory with a photo-multiplier tube, exciting the crystal with a  $^{60}\text{Co}$  source. This gives an average  $LY_{\text{PMT}}$  for the  $\text{PbWO}_4$  crystals of 10 pe/MeV at 18°C.

It is a delicate and difficult measurement because the  $\text{PbWO}_4$  crystals have a rather low LY and the energy of the  $^{60}\text{Co}$   $\gamma$  is only 1.2 MeV. Due to the different percentage of the crystal rear face covered by the PMT with respect to the VPT or APD and to the quantum efficiency, different values are measured in equipped ECAL units:  $\approx 4.5$  pe/MeV for the barrel channels read out with APDs and a similar number in the endcap channels read out by VPTs (the larger surface area of the VPT photocathodes being balanced by the higher quantum efficiency of the APDs).

The measurements span about 7 years of crystal production, so the stability of the LY bench calibration is crucial and is constantly controlled using reference crystals. The LY measurement of each crystal is rescaled using the daily reference crystal measurements [145].

The laboratory LY measurement can be improved exploiting the correlation between the crystal LY and the Longitudinal Transmission at 360 nm ( $\text{LTO}_{360}$ ). The latter can be measured with better precision and stability [146, 147].

Test-beam intercalibration constants, as described in Section 4.4.1.2, can be used to measure the precision achieved by laboratory estimations of intercalibration. Figure 4.14a shows the intercalibration coefficients calculated from laboratory measurements plotted against the test-beam measurements. Figure 4.14b shows that a resolution of about 4% is obtained with this method.

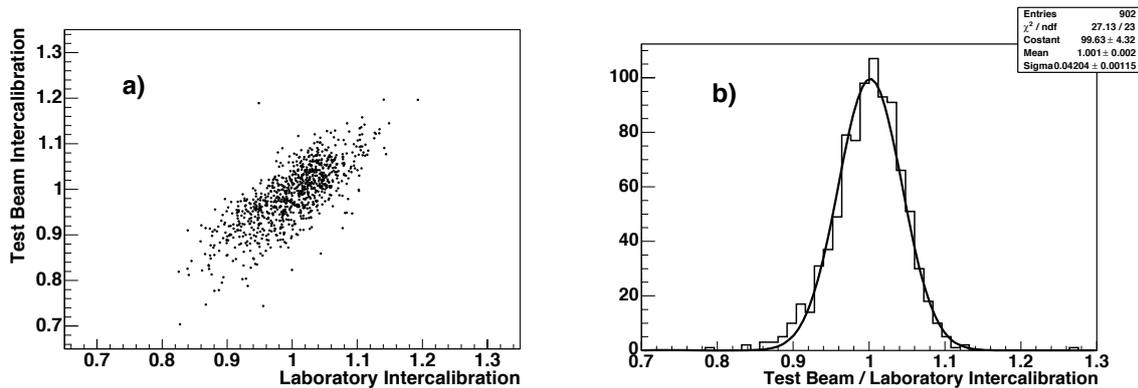


Figure 4.14: (a) Intercalibration coefficients obtained with supermodule 10 at the 2004 test-beam with high energy electrons versus intercalibration coefficients calculated from laboratory measurements and (b) distribution of the fractional difference between the laboratory estimated coefficients and coefficients measured in the test-beam.

#### 4.4.1.2 Testbeam precalibration

The intercalibration method presented here was developed using test beam data taken at CERN between year 2000 and 2004. In the test beam, supermodules are mounted on a rotating table which allows rotation in both the  $\eta$  and  $\phi$  coordinates and a full scan of the supermodules with high energy electron beams. The electron position is measured with a set of hodoscopes.

The response of a single crystal to electrons depends on the electron impact position. The dependence in the 2 lateral coordinates can be factorized and fitted with a 4<sup>th</sup> order polynomial separately. The corrected response of the single crystal  $S_{\text{CORR}}$  can be obtained from the

measured amplitude  $S_{\text{meas}}$  as:

$$S_{\text{corr}} = S_{\text{meas}} \frac{P_x^{\text{max}} P_y^{\text{max}}}{P_x(x) P_y(y)}, \quad (4.5)$$

where  $x$  and  $y$  are the measured positions of the incident electron in the two lateral coordinates, and  $P_{x,y}^{\text{max}}$  is the maximum of the polynomial. Figure 4.15 shows the fit and the crystal response after correction. Only the events impinging in a central  $7 \times 7 \text{ mm}^2$  window are used, (about 25% of events taken with a  $20 \times 20 \text{ mm}^2$  trigger). The intercalibration coefficients  $c_i$  for crystal  $i$  are defined as the ratio of the mean value of the corrected response with respect to a reference value.

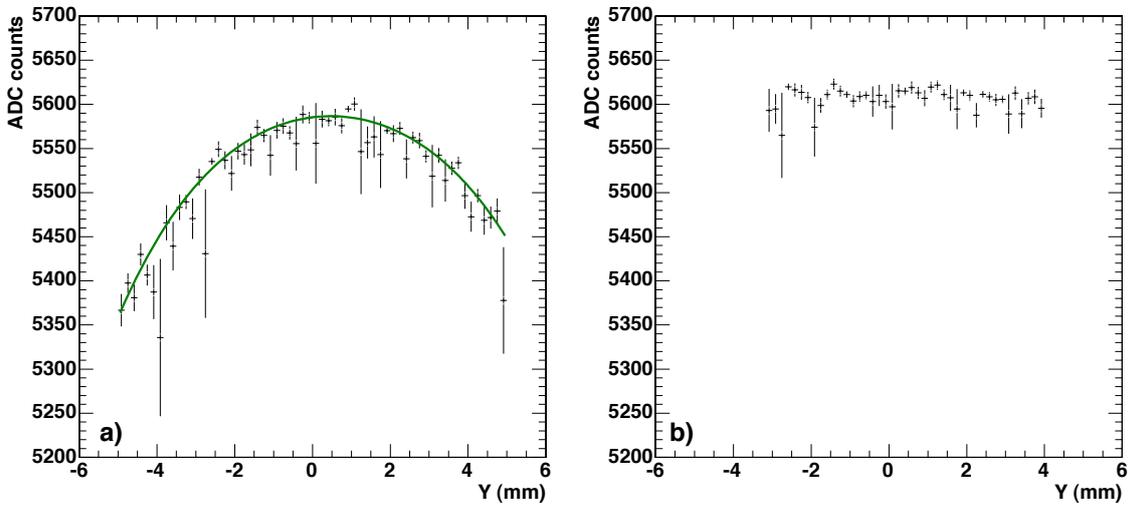


Figure 4.15: a) Polynomial fit to energy in a single crystal as a function of a lateral coordinate; b) energy in a single crystal as a function of a lateral coordinate after application of polynomial corrections for both lateral coordinates.

The statistical uncertainty remains negligible (less than 0.1%) provided that at least 1000 events are taken per crystal. Several tests have been done to evaluate the robustness of this method: using a different polynomial for each crystal or the same polynomial for all crystals. The intercalibration precision, when these constants are used *in situ*, is expected to be limited by variations occurring in the time between their determination in the test beam and their utilization in the installed detector.

At least 5 supermodules will be precalibrated in the H4 beam during the summer of 2006. The same supermodules will also be calibrated with cosmic rays, as described in Section 4.4.1.3 in order to compare the response of the 2 methods. It is foreseen to repeat the precalibration of at least one supermodule in order to quantify the reproducibility of the procedure, and the transportation of the constants over a period of a couple of months.

#### 4.4.1.3 Measurements with cosmic ray muons

Intercalibration coefficients for barrel supermodules are also obtained using cosmic muons which are well aligned with the crystal axes [148]. For this measurement the APD gain is

increased by a factor 4 with respect to the gain to be used during normal data taking by increasing the bias voltage.

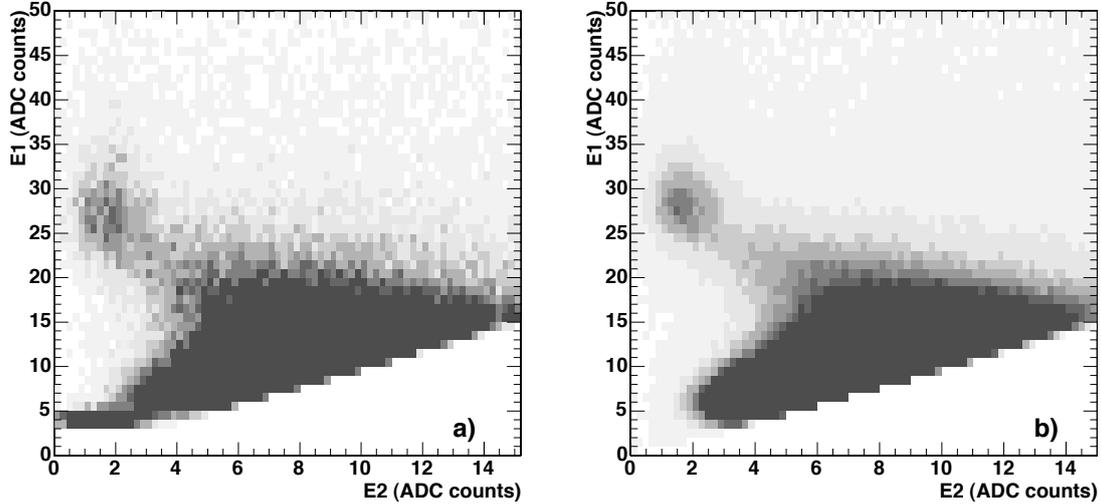


Figure 4.16: Scatter plot of  $E_1$  vs  $E_2$  in real cosmic muon data (a) and in simulation (b). The number of events in the plots corresponds to 41 hours for the data and to 1 week for the simulation. Edge crystals have been excluded.

The feasibility of this method was verified by measurements on a supermodule which was exposed to cosmic ray muons for 41 hours in November 2004, and with a detailed GEANT4 simulation [149]. Good agreement was found between data and simulation. Well aligned cosmic rays, giving a large signal in the crystal they pass through ( $E_1$ ), are selected by vetoing on signals above a rather low threshold in the adjacent crystal with highest signal ( $E_2$ ), obviating the need for external tracking (Fig. 4.16).

In the region covered by the trigger, which roughly corresponds to Module 1, an agreement of about 3% was achieved with respect to the testbeam calibration as shown in Fig. 4.17. The statistical contribution to the overall uncertainty was estimated to be 2%.

Some difficulties emerge when calibrating edge crystals (for which the veto based on neighbouring channels is inefficient) and the extreme crystals of Module 4 (due to the unfavorable angle with respect to the dominant cosmic ray direction). However, by using a more selective trigger setup in the cosmic telescope and by inclining the supermodule by  $\approx 10^\circ$ , an overall precision of 3% should be achievable in 1 week of data taking for channels in Modules 1–3, and 3.5% for channels in Module 4.

#### 4.4.2 Phi independence

A method taking advantage of the  $\phi$ -symmetry of deposited energy to intercalibrate crystals within rings at constant  $\eta$  has been studied. Intercalibration is performed by comparing the total energy deposited in each crystal with the mean of the distribution of total energies for all crystals at that pseudorapidity. Two choices of event trigger have been investigated: random bunch crossings [150], and Level-1 jet triggers [151]. The results shown in the following are based on a Level-1 jet-triggers simulated sample.

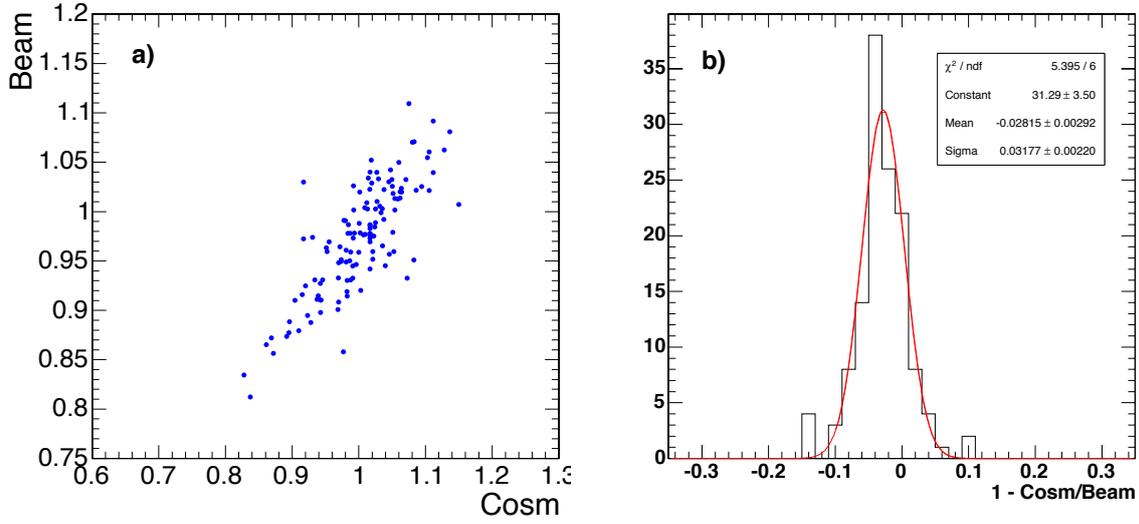


Figure 4.17: Correlation between (a) the testbeam calibration coefficients and cosmic ray muon coefficients, and (b) relative precision of the cosmic muon calibration constants. Only those channels with more than 50 cosmic muon events are considered.

The value of the total transverse energy,  $\Sigma E_T$ , deposited in each crystal from all selected events is determined. In the case of jet triggers, only crystals with transverse energy deposits in the range  $1 < E_T < 6$  GeV (barrel), and  $1 < E_T < 4$  GeV (endcaps) contribute to the energy sums. The lower threshold excludes noise, and the upper threshold improves the stability of the sums. In addition, in order to avoid trigger bias, crystals associated with the jet which has triggered the event (the highest- $E_T$  jet in the event) are excluded by requiring that crystals are separated from the position of the triggering jet by more than 1 rad. The intercalibration precision for a given  $\eta$  is obtained from the Gaussian width of the distribution of  $\Sigma E_T$  for the pair of rings of crystals at that absolute value of  $\eta$ . The precision is determined separately for each of 85 pairs of rings in the barrel and 39 pairs of rings in the endcaps.

With perfect  $\phi$ -symmetry the intercalibration precision is expected to vary as  $1/\sqrt{N}$ , where  $N$  is the number of events. A limit on the precision arises due to non-uniformities in  $\phi$ , primarily from the inhomogeneity of tracker material, but also from geometrical asymmetries such as the varying off-pointing angle of endcap crystals, and the boundaries between barrel supermodules. These non-uniformities result in a spread in the  $\Sigma E_T$  values which cannot be reduced by increasing the statistics of the event sample. The limit  $s$  on the precision is determined by calculating the precision for different numbers of events,  $N$ , and fitting a two parameter function

$$f(N) = \sqrt{s^2 + (m/\sqrt{N})^2}.$$

to the measured points. The parameter,  $s$ , measures the limit on the attainable precision. The technique has been tested directly by applying a set of miscalibration factors, chosen randomly from a Gaussian distribution of width 4.5%, to crystals in a pair of rings, and determining the residual miscalibration after correction. Figure 4.18 shows the distribution of residual miscalibrations  $\epsilon_R$  after 2 iterations of the correction procedure, obtained with  $1.1 \times 10^7$  Level-1 jet trigger events. The Gaussian width of the distribution is  $(2.7 \pm 0.1)\%$ .

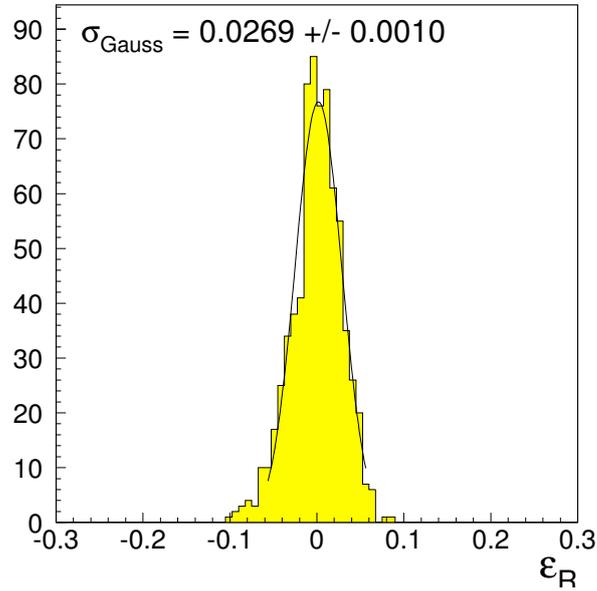


Figure 4.18: Distribution of residuals  $\epsilon_R$ , for the 720 crystals in the pair of rings at  $|\eta| = 0.23$ , after 2 iterations of the correction procedure, for  $1.1 \times 10^7$  Level-1 jet trigger events.

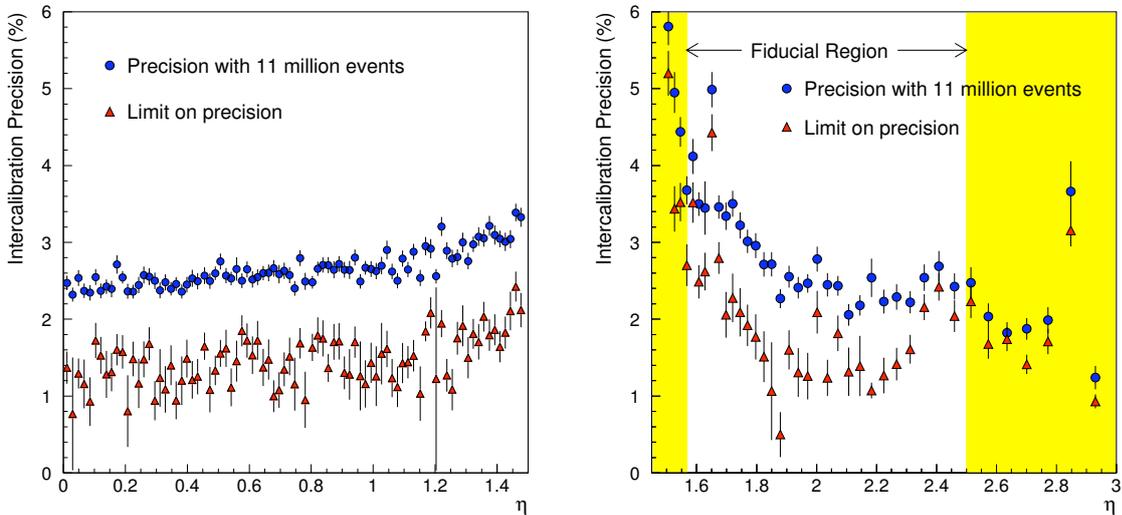


Figure 4.19: Intercalibration precision which can be obtained with 11 million Level-1 jet trigger events and the limit on the intercalibration precision due to tracker material inhomogeneity as a function of  $\eta$ .

It can be seen in Fig. 4.19 that without using any knowledge about the material distribution in the tracker, the limit on the precision is close to 1.5% throughout the barrel and between 3.0% and 1.0% for the fiducial region of endcaps. It can be expected that the limit on the precision will be closely approached with a few tens of millions of events.

This is equivalent to about 10 hours of data taking, under the assumption that 1 kHz of Level-1 bandwidth is allocated to single jet triggers, and that the calibration software has

access to this rate, either running on the Filter Farm, or more probably, running offline on a highly compacted data stream (a few tens of channels stored per event). With increasing knowledge of the material distribution in the tracker, after the start of LHC running, the attainable precision of the method is expected to increase, with the potential of providing rapid and repeated calibration of the ECAL. Intercalibration of different  $\eta$  rings can be done with one of the following methods.

### 4.4.3 Intercalibration using single electrons

Once the Tracker is fully operational and well aligned, intercalibration of different crystals within a single module can be performed using the momentum measurement of isolated electrons. The main difficulty in using electrons for intercalibration is that they radiate in the tracker material in front of the ECAL, and both the energy and the momentum measurement are affected. Moreover the average amount of bremsstrahlung varies with tracker material thickness.

Detailed Monte Carlo simulation studies of calibration using electrons have been made using fully simulated  $W \rightarrow e\nu$  events, digitized with pile-up corresponding to low luminosity ( $\mathcal{L} = 2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ ) [152]. While the methodology and techniques used are applicable to the whole ECAL, the characteristics of the events, the material budget and the geometry of the detector differs between the barrel and the endcap regions, and the 2 regions have been studied separately.

For these studies the ECAL energy was measured by summing the  $5 \times 5$  array of crystals around the crystal with the maximum signal ( $S_{25}$ ). In the endcap the energy measured in the preshower and associated with the electron cluster is added to the energy summed in the crystals. The energy in a cluster was used for intercalibration, rather than the single crystal calibration used in the ECAL test beam. The cluster energy does not require the complexity of a single crystal containment correction, which would always be a potential source of error. Moreover, it makes better use of the available information by using also the energy deposited in crystals other than the one containing most of the energy. The choice to use  $S_{25}$  rather than the “superclustering” algorithms (Section 10.1), which are generally used for electron reconstruction, was motivated by the wish to cleanly separate the intercalibration from the algorithmic corrections required for the superclustering algorithms.

Each energy measurement contains the contribution of many crystals, each with its own calibration constant. In order to extract those constants the individual crystal contributions must be unfolded, while minimizing the difference between the energy and momentum measurements. Two algorithms to achieve this minimization have been considered: an iterative technique which was used for the in-situ calibration of the BGO crystals in the L3/LEP experiment [153] and a matrix inversion algorithm. The results, both in terms of precision and in terms of speed of algorithm, are similar, and show no dependence on the technique used.

#### 4.4.3.1 The barrel case

The event selection was based on variables which are sensitive to the amount of bremsstrahlung emission, and consequently measure the quality of the energy and momentum reconstruction. These variables are the number of valid hits and the  $\chi^2/\text{n.d.f.}$  of the electron track,  $S_9/S_{25}$  measured by the ECAL (where  $S_9$  is the energy contained in a  $3 \times 3$  array of crystals).

tals around the crystal with the maximum signal), and the ratio  $S_{25}/P$ , where  $P$  is the track momentum.

The selection cuts were chosen by scanning the four-dimensional phase space of the selection variables and choosing the point giving the best precision (global minimum) in the determination of the calibration constants. The optimal point is dependent on both the number and the quality of the input electron measurements, both of which depend on the average amount of bremsstrahlung in the region considered. The sensitivity of the results to the selection cuts was extensively investigated. The calibration precision attained shows a rather wide global minimum with respect to variations of the selection cuts, which guarantees the stability of the selection process and allows for potential small differences between simulated events, used to derive the selection cuts, and data.

Due to the variation of the average value of  $S_{25}/P$  with pseudorapidity, caused by the variation of the amount of material in front of the ECAL, the calibration task was divided into 2 steps. In the first step crystals in small regions in  $\eta$ , over which the average value of the  $S_{25}/P$  is rather constant, were intercalibrated. In the second step the small regions are intercalibrated with each other.

To test the calibration algorithms, a Gaussian miscalibration has been applied crystal by crystal, with a spread of 4% around 1. This represents a conservative estimate of the residual miscalibration expected from the precalibration, and early intercalibration of the crystals at start-up (Section 4.4.1). For each test many Monte Carlo experiments ( $\geq 50$ ) were performed with different randomly chosen sets of miscalibration constants. The results given are the averages of the results of the experiments.

The calibration precision versus  $\eta$  achievable for a fixed integrated luminosity follows the tracker material budget distribution (Fig. 4.20). There are fewer electrons that have radiated only a little at higher  $\eta$ . Fewer events are selected by the High Level Trigger (HLT) here and those selected are still, on average, less well measured electrons. The Monte Carlo simulation data used to obtain these results correspond to about  $5 \text{ fb}^{-1}$ . This estimation uses the PYTHIA cross section for  $W$ -production with no k-factor.

The calibration precision was also extensively studied in different  $\phi$  regions keeping the same  $\eta$  interval. There is no evidence of any  $\phi$  dependence.

For electrons impacting crystals on the module boundaries, there is a loss in the collected energy. The calibration procedure artificially increases the coefficients for these crystals to compensate for energy loss. A suitable method must be used to deal with this problem. This may involve applying a correction for clusters spanning intermodule boundaries. Such a correction can be derived from test beam data.

It is found that extending the intercalibration area in  $\phi$  direction, keeping the same  $\eta$  region, does not affect the intercalibration accuracy once a module boundary correction is applied.

After crystals within regions are intercalibrated, the regions have to be calibrated among themselves. This task is accomplished by selecting electrons with minimum energy loss due to bremsstrahlung. After this selection, the resulting values of the peaks of the  $S_{25}/P$  distributions that are found are consistent with the pseudorapidity dependence of shower containment.

### Dependence on integrated luminosity

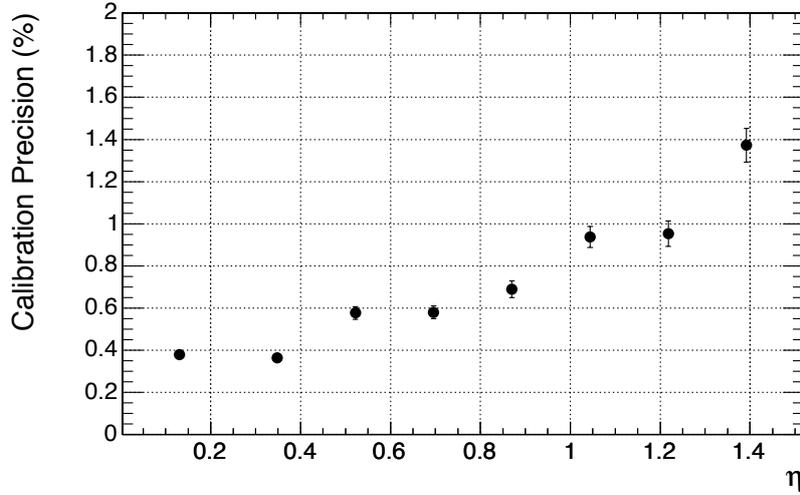


Figure 4.20: Calibration precision versus  $\eta$  for  $5 \text{ fb}^{-1}$  of integrated luminosity.

The calibration precision achievable is strongly dependent on the available statistics, as shown in Fig. 4.21. The dependence of the precision on the number of electrons per crystal output by the HLT, i.e. the number of electrons available to be used for calibration, can be fitted by the function:

$$\text{resolution} = \frac{A}{\sqrt{\text{HLT events per Crystal}}} \oplus C \quad (4.6)$$

The parameters  $A$  and  $C$  depend on  $\eta$  and have greater values as  $\eta$  increases. The constant term,  $C$ , sets the limit on the precision that can be reached with infinite statistics. In Table 4.2 the fitted values of  $A$  and  $C$  are listed for the different  $\eta$  regions. The corresponding fitted curves are shown in Fig. 4.21.

Table 4.2:  $A$  and  $C$  parameters for different  $\eta$  regions in the barrel.

$\eta$	$A$ (%)	$C$ (%)
0.000–0.261	6.19	0.12
0.783–0.957	10.7	0.27
1.305–1.479	15.0	0.42

#### 4.4.3.2 The endcap case

For the study of the endcap GSF tracking was used for the electron track reconstruction (Section 10.4.2), and the event selection parameters for tracks are consequently different. As in the barrel case, the event selection uses variables which are sensitive to the amount of bremsstrahlung emission, and thus measure the quality of the energy and momentum reconstruction. The variables used were:  $S_9/S_{25}$  and the ratio  $P_{\text{out}}/P_{\text{in}}$  between the momentum measured at the end of the electron track and that measured at the vertex. As in the barrel case a cut was also made on the ratio  $S_{25}/P$  itself. The selection efficiency varied between 10% and 30%. The sensitivity of the calibration precision shows only a small dependence on the selection cuts.

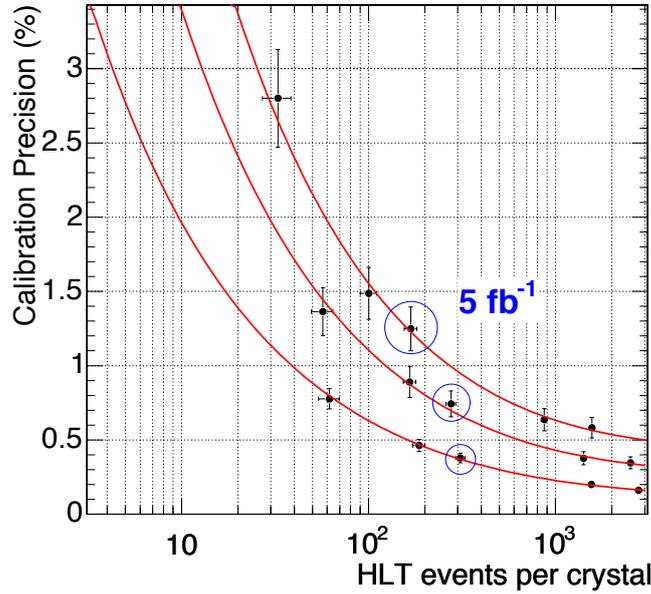


Figure 4.21: Calibration precision versus HLT events per crystal for different  $\eta$  regions. Upper curve: the last 10 crystals in the ECAL Barrel ( $1.305 < \eta < 1.479$ ), Middle curve: 10 crystals in the middle of ECAL Barrel ( $0.783 < |\eta| < 0.957$ ), Lower curve: the first 15 crystals in the ECAL Barrel ( $0.0 < |\eta| < 0.261$ ). The third point along each line gives the precision for  $5 \text{ fb}^{-1}$  of integrated luminosity.

The calibration precision versus  $\eta$  achievable for a fixed integrated luminosity is shown in Fig. 4.22. No evidence of any  $\phi$  dependence was found. The precision is limited by the momentum resolution, which is worse than in the barrel. No significant effect on the precision achieved was seen for crystals at the border of the supercrystal (Section 4.1) structure.

The calibration precision achievable is strongly dependent on the available statistics. The dependence of the precision on the number of electrons per crystal output by the HLT was fitted using the parametrization in Equation 4.6 for the region at  $\eta \approx 2.0$ . The data points and fitted curve is shown in Fig. 4.23.

#### 4.4.3.3 Background contamination

The rate of the dijet background in the single electron trigger stream (HLT output) is estimated to be 27 Hz at low luminosity (Section 10.2.2) out of which 16 Hz are expected in the barrel. The residual background has been investigated for the barrel case. After the selection described above is applied the surviving background corresponds to a rate of 2.3 Hz. One third of this rate comes from  $b/c \rightarrow e$  semileptonic decays. Such decays might be useful in the calibration process, increasing the overall calibration statistics. In fact their  $S_{25}/P$  distribution is very similar to the  $W \rightarrow e\nu$  signal electrons. If required, the background can easily be further reduced by a factor 10 using isolation cuts, with only a small effect on the signal efficiency.

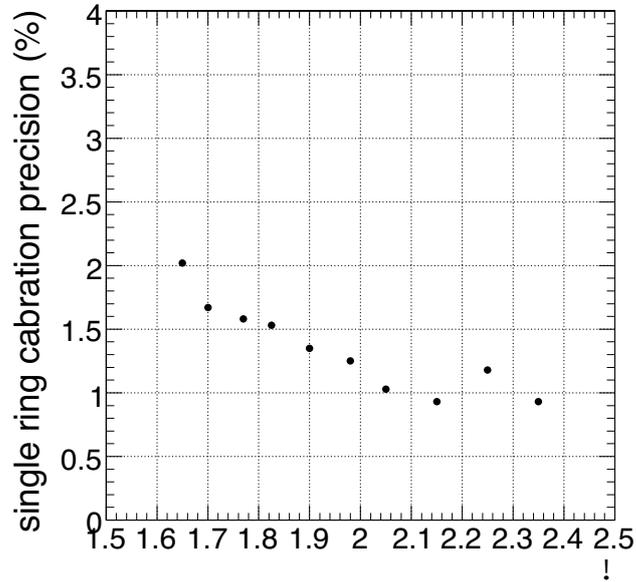


Figure 4.22: Calibration precision versus  $\eta$  for  $7 \text{ fb}^{-1}$  of integrated luminosity.

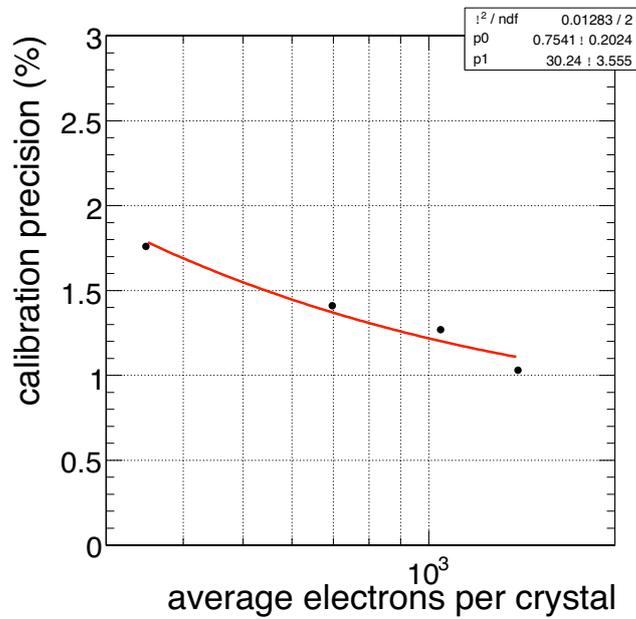


Figure 4.23: Calibration precision versus HLT events per crystal for a region at  $\eta \approx 2.0$ . The last point gives the precision for  $7 \text{ fb}^{-1}$  of integrated luminosity.

#### 4.4.3.4 Systematic Uncertainties

It is assumed that the tracker can be well aligned with  $5 \text{ fb}^{-1}$  of data. Nevertheless, in order to study the effect of a misaligned/miscalibrated tracker, the inverse of the track momentum was smeared by 2%. In that case, the calibration precision achievable using  $5 \text{ fb}^{-1}$  of data worsens by a relative 20%.

If the initial miscalibration error is increased to 8%, the calibration precision achieved, for a given number of events, is only very slightly worsened. Such an initial high value corresponds to the calibration uncertainty for completely unmeasured crystals in the barrel.

It is inevitable that there will be a few non-functioning channels. As in the case of crystals at the module boundaries, the calibration procedure artificially increases the coefficients of the crystals in a  $3 \times 3$  window around the dead one, in order to compensate for the energy loss caused the missing energy. A comprehensive correction procedure for dead channels might be to parameterize the energy content of the crystals in the  $5 \times 5$  matrix as function of energy,  $\eta$ ,  $\phi$ , and particle impact point. This might be achieved by looking in areas without dead crystals and making a one-to-one correspondence of the energy deposition pattern using a technique similar to the one described in Section 10.1.4.

#### 4.4.3.5 Summary of intercalibration using single electrons

Although further study remains to be done, the preliminary results from studies of intercalibration using single electrons are quite encouraging. The intercalibration precision attainable with  $10 \text{ fb}^{-1}$ , averaged over the barrel, is about 0.6%. The single electron rates have been calculated rather conservatively using the PYTHIA cross sections and no k factors, and not including additional contributions from electrons in  $b$  decays. Further optimization may involve more complete separation of the task of intercalibration of local regions where the amount of tracker material is fairly constant, from the task of intercalibrating those local regions to each other. This would allow the use of clustering algorithms which are less sensitive to bremsstrahlung than the  $5 \times 5$  clusters, hence allow the use of a significantly larger fraction of the available electrons.

#### 4.4.4 Use of $Z \rightarrow ee$

The  $Z$  mass constraint in  $Z \rightarrow ee$  decays is a powerful tool for ECAL calibration. A number of different uses are envisaged, from the tuning of the algorithmic corrections for electron reconstruction, to the intercalibration of regions of the ECAL, for example as a complement to the  $\phi$  symmetry method at the start-up.

A study of the use of the  $Z \rightarrow ee$  has been performed with a fully simulated data sample digitized with low luminosity pile-up, corresponding to an integrated luminosity of  $2.4 \text{ fb}^{-1}$  [154]. This study was restricted to  $Z$  decays with both electrons in the barrel, but the method can be extended to the whole ECAL.

An iterative method has been developed to tune the algorithmic corrections and to extract intercalibration constants of regions or individual crystals from samples of  $Z \rightarrow ee$  events. Starting from the 2 electron invariant mass reconstruction, the average quantity  $\epsilon$  in event  $i$  is defined as

$$\bar{\epsilon}^i = \frac{1}{2} \cdot \left[ \left( \frac{M_{inv}^i}{M_Z} \right)^2 - 1 \right], \quad (4.7)$$

The calibration constants are obtained from the peak of the  $\bar{e}^i$  distribution, so as to reduce the sensitivity to tails. Convergence is reached after 10–15 iterations. The intercalibration precision reached is insensitive to the magnitude of the initial miscalibration.

#### 4.4.4.1 Tuning of algorithmic corrections

The  $Z \rightarrow ee$  sample also allows algorithmic corrections ( $\mathcal{F}$  in Equation 4.3) to be determined or cross-checked. This was tested for the case of the functions used to correct the supercluster  $\eta$  dependence for different electron classes. It was assumed that the crystals were intercalibrated with a precision of only 2%. Miscalibration constants, with an RMS spread of 2% were applied to the crystals.

The electron energy was reconstructed with the superclustering algorithms described in Section 10.1, and the electron classification described in Section 10.4 was used. Events are selected where either both electrons belong to the golden class, or to the showering class (the “best” and “worst” classes, respectively). The energy is reconstructed in superclusters but without applying the algorithmic correction for variation in  $\eta$ . Figure 4.24 compares the correction factors obtained from  $Z \rightarrow ee$  events used as real events can be used, with those obtained from Monte Carlo studies where the true values of the generated kinematics were used. The agreement in the shape is excellent, but small overall rescaling values are needed which are different for different classes of electrons and which have been used in the figure (1.005 for the golden electrons, and 1.008 for the showering electrons). The origin of these rescaling factors is under investigation, one suggestion is that the difference is due to the presence of pile-up in the  $Z \rightarrow ee$  sample, and the absence of pile-up in the electron sample used to extract original correction factors.

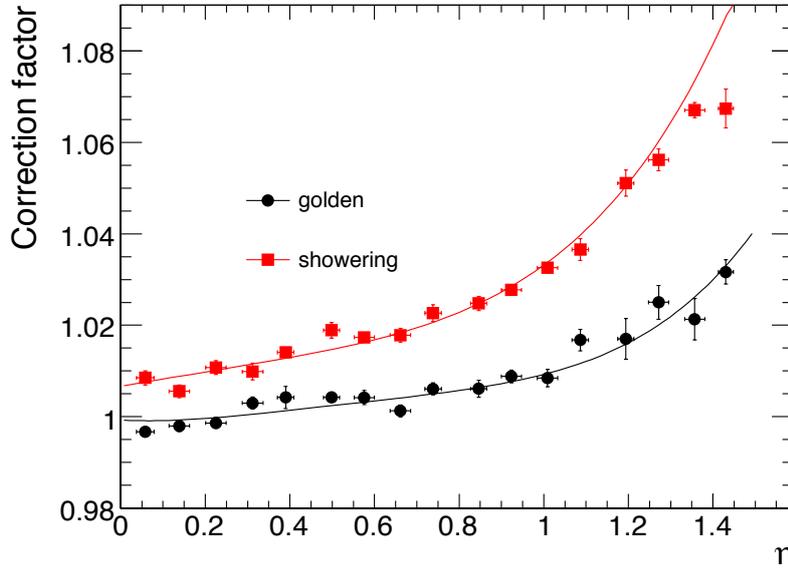


Figure 4.24: Correction  $f(\eta)$  obtained using the Monte Carlo simulation compared to that obtained from  $2 \text{ fb}^{-1}$  of  $Z \rightarrow ee$  with intercalibration precision at 2%.

#### 4.4.4.2 Intercalibration of crystal rings

In a start-up scenario, where the algorithmic correction factors are taken from Monte Carlo simulation,  $Z \rightarrow ee$  events might be used to obtain a preliminary estimate of the intercalibration factors between rings. The golden electron class from the classification described in Section 10.4 is chosen since these electrons radiate little, and their reconstructed energy shows the least dependence on the tracker material, and hence  $\eta$ . A 4% algorithmic correction is still needed to equalize the supercluster energy measurement over the length of the barrel. Only events with two golden electrons are selected.

The method has been tested taking the calorimeter regions as rings of crystals (at fixed  $\eta$ ) in the ECAL barrel. There are 170 such rings of 360 crystals each. The miscalibration constants can be set up to have a different RMS variation between crystals within a ring to the variation of the average miscalibration value between rings. To reduce the sensitivity to tails, the peak of the mass distribution is used to obtain the intercalibration coefficients. The peak is obtained from an iterative Gaussian fit in a region  $[-2.5\sigma, 2.5\sigma]$  around the peak. Crystals on the module borders are excluded.

The performance of the method was measured by comparing the applied miscalibration factors with the calibration coefficient determined by the method. The results given here are obtained when starting from a 5% miscalibration between rings and a 2% miscalibration between crystals within a ring. Using events corresponding to an integrated luminosity of  $2.0 \text{ fb}^{-1}$  the distribution of the residual miscalibration is shown in Fig. 4.25. The RMS spread of this distribution, corresponding to 0.6%, gives the achieved ring intercalibration precision.

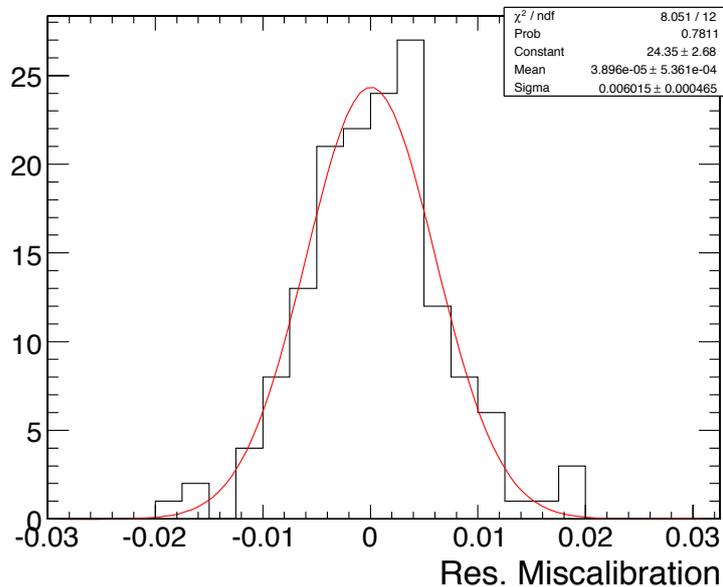


Figure 4.25: Distribution of the residual miscalibration using  $2.0 \text{ fb}^{-1}$ .

The ring intercalibration precision as a function of the number of events per ring is shown in Fig. 4.26. Table 4.3 shows the accuracy reached for each module type for an integrated luminosity of  $2.0 \text{ fb}^{-1}$ .

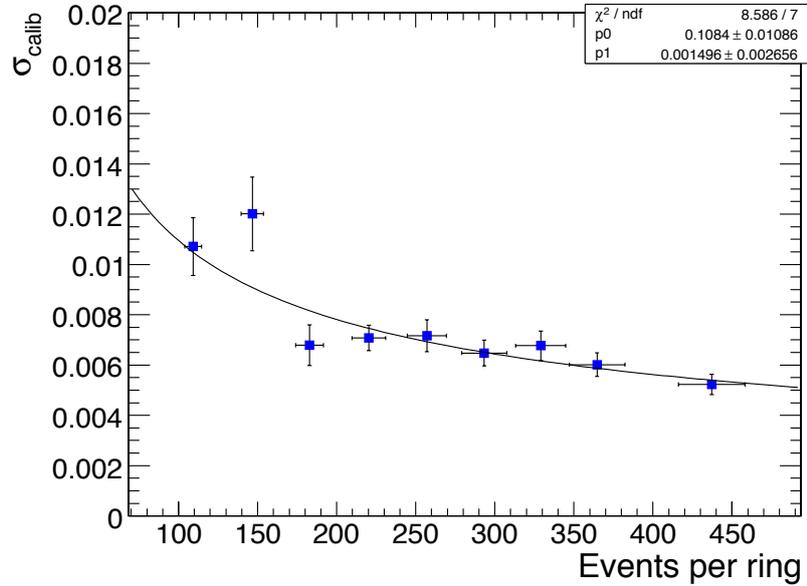


Figure 4.26: Ring-to-ring calibration precision achieved versus average number of events per ring. An average of 370 events per ring corresponds to an integrated luminosity of  $2.0 \text{ fb}^{-1}$ . The full curve is a fit to the sum of a statistical term and a constant term.

Table 4.3: Intercalibration precision between rings using a data sample equivalent to  $2.0 \text{ fb}^{-1}$ .

Module 1	$0.32\% \pm 0.05\%$
Module 2	$0.5\% \pm 0.1\%$
Module 3	$0.6\% \pm 0.1\%$
Module 4	$1.3\% \pm 0.4\%$

#### 4.4.5 Intercalibration from $\pi^0$ and $\eta \rightarrow \gamma\gamma$ events

The possibility of intercalibrating the ECAL using mass reconstruction of  $\pi^0$  and  $\eta \rightarrow \gamma\gamma$  is being investigated. Experiments such as CLEO and L3 have successfully used low energy signals to calibrate crystal calorimeters to high accuracy. In the case of CMS, these low mass particles could provide an important additional calibration tool which is useful for

- relatively rapid intercalibration of all crystals;
- study of the effects of crystal transparency corrections from the laser monitor;
- and rapid check-out and monitoring of detector performance.

The intercalibration obtained from low-energy  $\pi^0$ s is not sensitive to tracker material as much as the intercalibration obtained from electrons is. First,  $\pi^0$  decays into photons that do not convert are unaffected. Converted low-energy photons give rise to low-energy electrons, which reach the ECAL far from the expected photon impact point because of the magnetic field. As a consequence, a selection of  $\pi^0$ s based on the selection of pairs of closeby ECAL clusters retains mostly either unconverted photons or photons which convert just in front of the calorimeter. For this reason, the energy resolution does not deteriorate and no energy bias is introduced at high  $\eta$ . The only effect of the tracker material is a rate loss at larger  $\eta$  values.

The minimum separation for 2 photons in the barrel from a  $\pi^0$  with  $E_T$  of 5 GeV is about 65 mm, corresponding to nearly 3 times the crystal granularity. Such  $\pi^0$ s occur in almost any kind of event. It has been shown that  $\pi^0$ s useful for calibration can be located within events using the ECAL Level-1 trigger information. They can then be reconstructed in a small region identified by the trigger. This operation would be similar to a normal first step of the High Level Trigger, requiring very little processing time to extract the small amount of information relevant for calibration.

To obtain an estimate of how many  $\pi^0$  can be reconstructed with little background using data read out with all Level-1 triggers, a large data set of high  $E_T$  jet events was studied. It was found that about 1.4% of them have a usable  $\pi^0$  in the barrel and that almost all of these are tagged by the isolated electron Level-1 trigger information. The selection applied to get this rate is rather simple. First, some fairly stringent shower shape cuts were applied to the individual photon candidates:  $0.5 < R_1 < 0.9$ ,  $R_4 > 0.2$  and  $0.9 < R_9 < 1$ , where  $R_n$  is the fraction of the cluster energy in a square array of  $n$  crystals centered on the cluster. Energy cuts were also made on the photon candidates. These cuts,  $1.5 < E_\gamma < 5$  GeV, are slightly tightened at high  $\eta$ . To reduce the combinatorial background, it is important to only consider  $\pi^0$  candidates with small opening angles. The cuts used correspond to photon separations at the ECAL front face between 60 and 90 mm.

The reconstructed mass of the candidates selected in this way is shown in Fig. 4.27 in 3 bins in  $\eta$  in the ECAL barrel. There is a  $\pi^0$  mass peak with relatively little background in all 3 regions, and the mass resolution is about 8% in each case. Reconstructed  $\pi^0$ s are found coming from both the jet event and from pile-up events in the same bunch-crossing.

This result suggests that it will be possible to intercalibrate the ECAL using the reconstructed mass of  $\pi^0$ . With 1000 events per crystal, a statistical precision of 0.5% can be estimated for the intercalibration constants. This needs to be demonstrated, and sources of systematic error must be investigated. Such a demonstration is in progress but extremely large numbers of Monte Carlo simulated events are required to probe a precision of 0.5%.

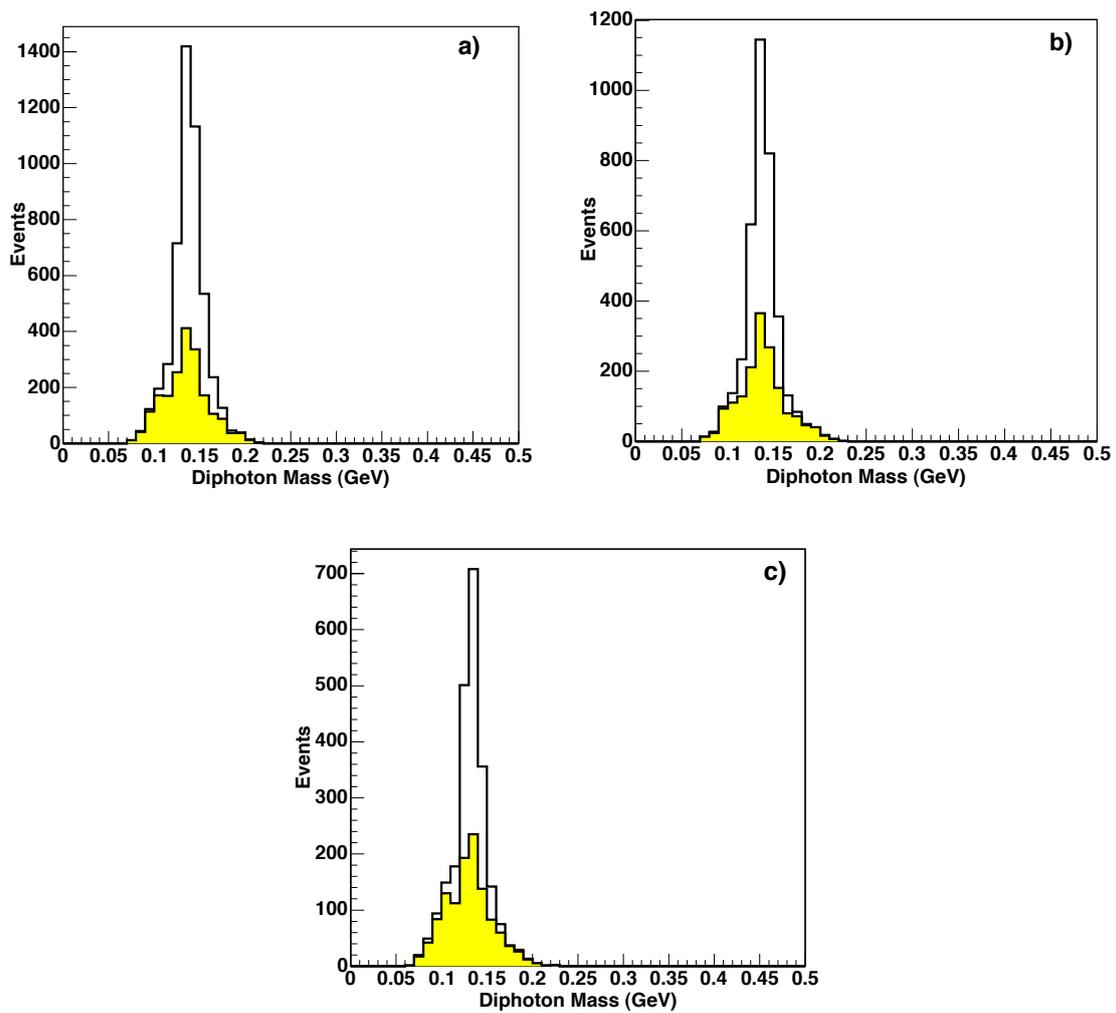


Figure 4.27: Reconstructed mass of  $\pi^0$  candidates in 3 regions of the barrel: a)  $|\eta| < 0.5$ , b)  $0.5 < |\eta| < 1.0$ , and c)  $|\eta| > 1.0$ . The unshaded histograms represent the signal while the shaded ones the background.

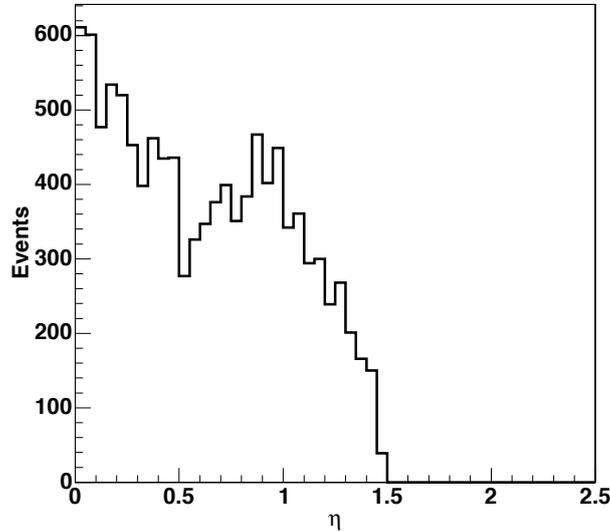


Figure 4.28: The  $\eta$  distribution of  $\pi^0$  candidates in the barrel.

As can be seen from the number of  $\pi^0$ s in the 3 histograms of Fig. 4.27, the rate decreases with  $\eta$ , but there is still large rate in the last bin at the edge of the barrel. The  $\eta$  distribution of the  $\pi^0$  candidates is shown in Fig. 4.28. With an assumed a Level-1 global trigger rate of 25 kHz, and with 1.4% of these triggers assumed to have usable  $\pi^0$ s, about 100  $\pi^0$ s per crystal can be obtained in a running period of less than 5 hours.

Events from  $\eta \rightarrow \gamma\gamma$  are also being studied. The signal has a much lower rate once the background is reduced sufficiently, but the mass resolution is about 3%. The  $\eta$  should be a useful calibration tool at higher energy, although it will take longer, and may prove very useful in the endcap.

#### 4.4.6 Inner bremsstrahlung photons in Z boson decays to muons

A significant rate of high- $p_T$  photons with very little background and an energy which can be known independently of the ECAL is available in radiative decays of  $Z \rightarrow \mu\mu$ . These photons are being investigated as a valuable tool for various calibration related tasks, as well as a probe for measuring photon reconstruction efficiency. They can be used, for example, to intercalibrate different regions of the ECAL (coefficient  $c_i$  of Equation 4.3), and to tune the various cluster correction algorithms (coefficient  $\mathcal{F}$ ) and the overall energy scale (coefficient  $G$ ). They can also be used to relate the energy scale of unconverted photons to that of electrons (from converted photons).

The available rates have been investigated using a sample of events generated using a full matrix element calculation of radiative decays. The background from Z bosons produced with additional jets has been studied also. Events are selected by placing a loose cut on the dimuon mass ( $40 < M_{\mu\mu} < 80 \text{ GeV}/c^2$ ) and searching for a photon with  $p_T > 15 \text{ GeV}/c$  within a radius of  $\Delta R < 0.8$  of either muon, where  $\Delta R$  is defined as  $\sqrt{\Delta\phi^2 + \Delta\eta^2}$ . Figure 4.29 shows the signal and background rates as a function of photon  $p_T$ . The rate is calculated after a mass window of  $87.2 < M < 95.2 \text{ GeV}/c^2$  has been placed on the three-

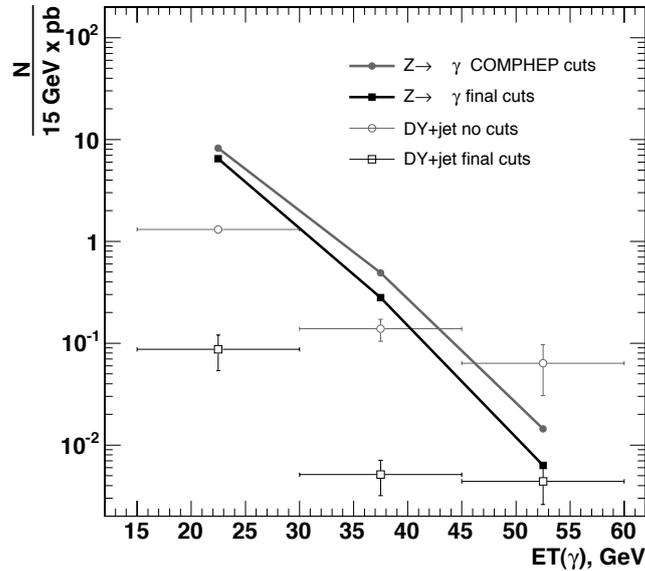


Figure 4.29: Signal and background yields, as a function of photon  $E_T$ , for three 15 GeV bins. The rates are shown before and after selection cuts.

body mass. With the selection cuts described, the signal-to-background ratio is  $\approx 80$  in the bin  $15 < E_T^\gamma < 30$  GeV, with a rate of  $65 \text{ pb}^{-1}$  over the region  $|\eta| < 2.5$ . For an integrated luminosity of only  $1 \text{ fb}^{-1}$ , an average of nearly 1 such photon per crystal will be collected. Therefore, a statistical precision better than 0.1% would be achieved with unconverted photons for the intercalibration of 10-crystal-wide rings of crystals (i.e. groups of 3600 crystals) in the barrel.

#### 4.4.7 Preshower calibration

##### 4.4.7.1 Overview

The energy deposited in the lead of the preshower will be estimated using the charge measurements made by 2 layers of silicon sensors. To preserve the overall photon energy resolution, a 5% charge measurement is necessary over a very large dynamic range (1–400 MIPs equivalent), resulting in a required strip calibration accuracy of a few per cent. There are 2 principle sources of response variation (sensor-to-sensor and channel-to-channel) at start-up: sensor thickness (RMS of  $\approx 1\%$ ) and gain uniformity of the electronics (RMS  $\approx 5\%$ ). Radiation damage to the sensors (decrease in charge collection efficiency by  $\approx 17\%$  over 10 years) and the electronics (decrease in gain by  $\approx 2\%$  over 10 years) will need to be assessed by periodic in-situ calibrations.

The principle difficulty in calibrating the preshower is that the only physics reference points are “zero” (zero output for zero input) and one MIP. The full dynamic range can only be explored systematically using an internal calibration circuit (ICC). Single MIPs with a S/N ratio suitable for calibration can only be measured when the front-end electronics is in high gain mode (HG). Consequently, the ICC must be used to transfer this information from HG to low gain (LG)—i.e. intercalibrate the 2 gains.

There are thus 2 parts to the full preshower calibration procedure:

- Absolute calibration using real MIPs
- Relative calibration between gains and between strips using an internal electronic charge generation system

Each of these parts is described in the sections below, followed by a summary of the procedures to execute before preshower installation in CMS, at start-up of CMS and subsequent periodic calibrations.

#### 4.4.7.2 Determination of the MIP scale

The absolute calibration scale can be found by detecting single MIPs in HG mode. Prior to installation of the preshower in CMS, stacks of up to 12 ladders will be built, sandwiched between scintillators to allow triggering and detection of cosmic rays. The cosmic rate is approximately 1 per minute per  $\text{cm}^2$  and are mostly muons with a mean energy of around 4 GeV [119] equating to 1 MIP per strip per minute. Running the ladders for 8–10 hours will thus give sufficient statistics per strip for an accurate absolute measurement of the MIP scale for each strip.

In CMS MIPs can be approximated with high energy muons and/or charged pions. The occupancy in the preshower is rather low, around 0.2% on average at low luminosity. Noise is thus the dominant “signal” in the strips, necessitating the use of the tracker and/or muon chambers to point to strips that have been traversed.

Studies of full jet and muon events at low luminosity have shown [155] that indeed both muons and charged pions can be used for calibration. Muon events are cleaner, but the muon chamber coverage only extends to  $|\eta| < 2.4$ . Pions are plentiful in both jet and minimum bias (i.e. pile-up) events and the tracker extends to  $|\eta| < 2.5$ .

The overall efficiency for muons and high  $p_T$  pions is around 15%, taking into account geometrical factors and tracking efficiencies. Using a combination of both types of event, a sensor-by-sensor calibration to the required few percent accuracy can be made in the order of a week, with the inter-strip calibration being achieved with the ICC. This takes into account the HLT rates and the LHC duty cycle. As the in-situ calibrations will be mostly following the drop in the charge collection efficiency (which in turn is a linear function of the accumulated radiation damage [156]), geometrical factors (forward-backward endcap symmetry; use of “rings” of sensors at similar  $\eta$  etc.) can be used to decrease the calibration time substantially.

These MIP calibrations must be taken with the front-end in HG mode. The dynamic range of the preshower will thus be reduced during the calibration periods.

#### 4.4.7.3 The internal calibration circuit

The ICC inside the front-end electronics provides precise charge injections into the pre-amplifier, mimicking real signals. The amplitude of the injected signal is controlled by an on-board 8-bit DAC coupled with a switchable *precision*: high precision produces injection signals between  $-16$  and  $12$  MIPs equivalent, in  $\approx 0.1$  MIP steps, whilst low precision allows signals between  $-120$  and  $500$  MIPs in  $\approx 2.4$  MIP steps.

The ICC uses a voltage step applied to a fixed “injection” capacitor. Each of the 32 channels in a front-end board (corresponding to the 32 strips of the silicon sensor) has an individual

injection capacitor, with a variation in value significantly less than 1%.

The ICC can be used to find an intercalibration between the 2 gain modes and serves to intercalibrate the electronics channels within a single micromodule if necessary (e.g. to check the uniformity of the decrease in preamplifier gain with radiation damage). It is also used to explore the linearity of response of the electronics over the full dynamic range.

Before it can be used for these purposes, the ICC itself must be calibrated. There are several steps to calibrate the ICC.

- With the ICC set to high precision and the front-end in HG, the internal DAC is set to provide an output comparable to the real MIP signal. This, coupled with the zero point, gives a “MIPs per DAC step in high precision” value.
- The DAC is then set to a high value (e.g. 220) and the corresponding number of MIPs,  $M_1$ , measured.
- The ICC is switched to low precision and the DAC adjusted to provide again  $M_1$  on the output. The “MIPs per DAC step in low precision” can then be estimated

Switching the front-end to LG mode then allows the intercalibration of the 2 gains.

The time required for the calibration of the ICC (including the intercalibration of the 2 gains) and the intercalibration of the strips should be of the order of a few hours maximum for the complete preshower and can be done between LHC fills.

#### 4.4.7.4 The preshower calibration schedule

The absolute MIP calibration with cosmic rays, together with a first calibration between the 2 gains and the examination of the full dynamic range, will be done before preshower installation in 2007/2008.

A first in-situ calibration (with MIPs) will be performed during the preshower commissioning phase in 2008 at low luminosity. As mentioned above, a sensor-by-sensor calibration can be achieved in the order of a week. It is likely that this commissioning time will be coupled with that of the endcap, resulting in a longer period being available, and thus a strip-by-strip calibration of the preshower can be performed. In either scenario, this first in-situ calibration will be a check of the calibration with cosmic rays.

Subsequent in-situ calibrations should be performed at the beginning and end of each year, and are essentially to follow the effects of radiation damage. A few days maximum will allow a sufficiently accurate calibration to be achieved.

#### 4.4.8 Correction for crystal transparency changes

##### 4.4.8.1 Crystal behaviour under irradiation

Although radiation resistant, ECAL crystals show rapid loss of optical transmission under irradiation due to the production of colour centres within  $\text{PbWO}_4$  which absorb a fraction of the transmitted light. This process is accompanied by transmission recovery due to self annealing processes, leading to cyclic transparency behavior between LHC collision runs and machine refills (Fig.4.30). The magnitude of the changes is dose-rate dependent, and is expected to range between 1 or 2 per cent at low luminosity in the barrel, to tens of per cent

in the high  $\eta$  regions of the endcap at high luminosity. The inter calibration of the crystals would be unacceptably degraded by these radiation induced transparency changes were they not measured and corrected for.

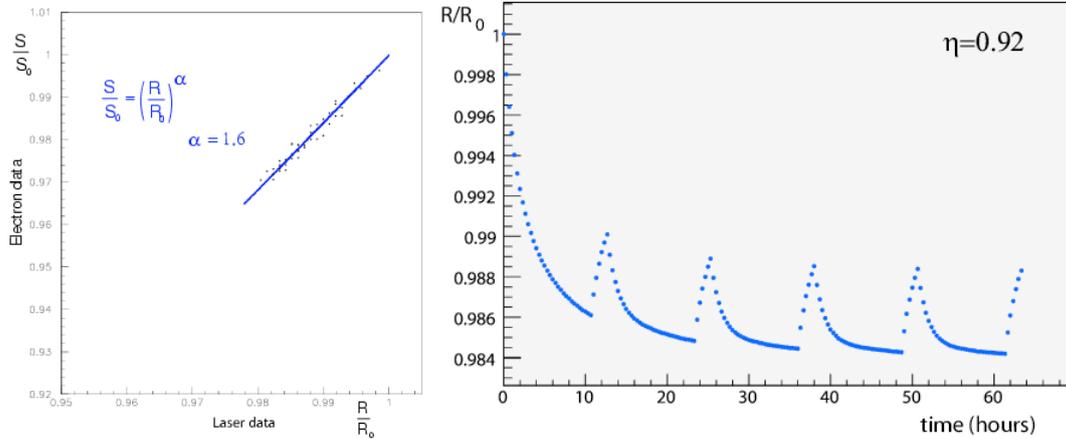


Figure 4.30: (a) Relation between the transmission losses for scintillation light and for laser light; (b) simulation of crystal transparency evolution at LHC based on test-beam results. For this illustrative example a luminosity of  $\mathcal{L} = 2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$  was assumed, together with a machine cycle consisting of a 10 hour coast followed by 2 hours filling time. The crystal behaviour under irradiation was modeled on data taken during a crystal irradiation in the test beam.

The evolution of the crystal transparency is measured using laser pulses injected into the crystals via optical fibres. The response is normalized by the laser pulse magnitude measured using silicon PN diodes. Thus  $R(t) = APD(t)/PN(t)$  is used as the measure of the crystal transparency. The laser monitoring system [5] performing this task is briefly outlined in the next section. Because of the different optical paths and spectra of the injected laser pulses and the scintillation light, the change in response to the laser light is not equal to the change in response to scintillation light. For attenuations  $< 10\%$  the relationship between the changes can be expressed by a power law,

$$S(t)/S(t_0) = [R(t)/R(t_0)]^\alpha, \quad (4.8)$$

where  $S(t)$  represents the response to scintillation light and  $\alpha$  is characteristic of the crystal ( $\alpha \approx 1.53$  for BCTP crystals, and  $\alpha \approx 1.0$  for SIC crystals). This power law describes well the behaviour of all the crystals that have been evaluated in the test beam, and this formula is expected to be valid in the barrel for both low and high luminosity at LHC.

#### 4.4.8.2 Laser-monitoring system overview

Figure 4.31 shows the basic components of the laser-monitoring system: Two laser wavelengths are used for the basic source. One, blue, at  $\lambda=440 \text{ nm}$ , very close to the scintillation emission peak, which is used to follow the changes in transparency due to radiation, and the other, infra-red, at  $\lambda=796 \text{ nm}$  far from the emission peak, hence very little affected by changes in transparency, which can be used to verify the stability of other elements in the system. The spectral contamination is less than  $10^{-3}$ . The lasers are operated such that the full width at half maximum of the pulses is  $\approx 30 \text{ ns}$ . The lasers can be pulsed at a rate of  $\approx 80 \text{ Hz}$ , and the

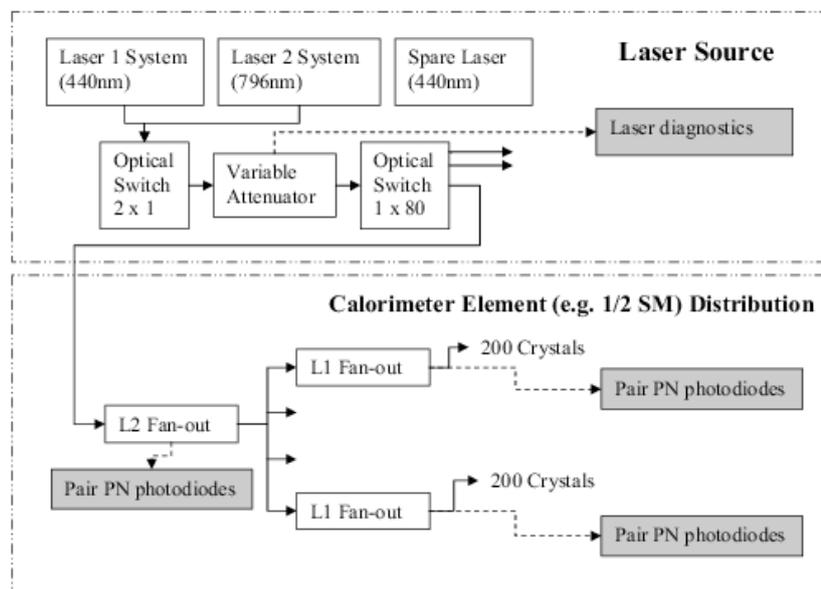


Figure 4.31: The components of the laser monitoring system.

pulse timing jitter is less than 3 ns to allow adequate trigger synchronization with the LHC bunch train and ECAL ADC clock.

The pulse energy of 1 mJ/pulse at the principal monitoring wavelength corresponds to  $\approx 1.3$  TeV, and a linear attenuator allows 1% steps down to 13 GeV. The pulse intensity instability is less than 10% which guarantees a monitoring precision of 0.1% by using the PN silicon photodiode normalization.

There are 3 light sources, 2 blue and 1 infrared. The duplication of the blue source provides fault tolerance and allows maintenance of one while the other is in use, ensuring that a source at the wavelength used to track changes in transparency is always available. Each source consists of an Nd:YLF pump laser, its power supply and cooler unit and corresponding transformer, a Ti:Sapphire laser and its controller, and a NESLAB cooler for an LBO crystal in the Ti:S laser. Each pair of the YLF and Ti:S lasers and their corresponding optics are mounted on an optical table. Each source has its own diagnostics, 2 fibre-optic switches, internal monitors and corresponding PC based controllers. Further details can be found in [157].

The monitoring light pulses are distributed via a system of optical fibres. A fibre optic switch at the laser directs the laser pulses to 1 of 88 calorimeter regions (72 half supermodules in the barrel and 8 regions in each endcap). A two-stage distribution system mounted on each calorimeter region delivers the light to each crystal.

To provide continuous monitoring, about 1% of the  $3.17 \mu\text{s}$  beam gap in every  $88.924 \mu\text{s}$  LHC beam cycle will be used to inject monitoring light pulses into crystals. The time needed to scan the entire ECAL is expected to be about 30 minutes.

#### 4.4.8.3 Operational procedures

Each measurement of the crystal transparency on 1 of the 80 calorimeter elements is envisaged to consist of 600 laser pulses (7.5 s at 80 Hz). The data will consist of the time frames read

out from the illuminated channels ( $\approx 1000$  channels), the time frames read out from the illuminated PN diodes, and the data read from the 2 Gs/s flash ADC. The latter samples the laser pulse viewed by a fast PIN diode located at the source to allow correction for variation of the laser pulse width.

Conceptually the monitoring and correction of changes in crystal transparency involves 3 steps:

- detailed data quality monitoring of the laser calibration data,
- the extraction of the  $R(t) = APD(t)/PN(t)$  values from the laser data,
- the setting and tuning of the reference starting point  $R(t_0)$  and the value(s) of  $\alpha$  used for the crystals

It is intended that the first 2 steps will be performed online. The data will be rapidly processed in a small farm of processors to extract the mean value  $R(t) = APD(t)/PN(t)$  for each measurement time,  $t$ . This processing is expected to involve correction for the non-linearity of the PN diode system, and may involve correction for variation of the laser pulse width.

The output of this processing will be a value of  $R(t)$  for each crystal, measuring the crystal transparency, which will be stored in the online database. Since the scan over the entire ECAL is expected to take about 30 minutes, a new value for each crystal will be produced every 30 minutes. It is intended that this initial processing will be made sufficiently reliable and robust so that it will not need to be repeated. However, the raw data will also be stored to allow reprocessing if necessary.

#### 4.4.8.4 Crystal transparency correction and systematic uncertainties

The transmission corrections may have systematic uncertainties, which result in intercalibration errors, which in turn contribute to the ECAL energy resolution. The magnitude of these contributions can be estimated based on experience from beam tests [158, 159]. The 4 main sources of error are 1) the dispersion of  $\alpha$  between crystals, 2) the stability and reproducibility of the monitoring system itself, 3) laser pulsewidth changes, and 4) photodetector gain changes.

##### 1. Dispersion of $\alpha$

The value of  $\alpha$  is determined by the optical properties of the crystals, and hence by the details of the crystal production. Using data from the test beam,  $\alpha$  has been measured to be 1.53 (1.0) with a relative uncertainty 5% (10%) for crystals produced by BCTP (SIC); Although the relative dispersion is larger for the SIC crystals, the absolute contribution to the systematic uncertainty is similar to the BCTP crystals. The magnitude of this uncertainty is proportional to the magnitude of the changes being corrected for, and is larger for a group of crystals under a higher or a wider range of radiation exposure rates. Unfortunately it will be not possible to measure the parameter  $\alpha$  for all crystals before LHC operation. There is good reason to believe, however, that a direct evaluation for a group of crystals can be obtained using physics calibration events (for example,  $W \rightarrow e\nu$ ) taken *in situ*.

##### 2. Monitoring stability

The light pulses injected into the crystals are simultaneously monitored by a local pair of PN photodiodes mounted close to the point of injection. This monitoring is done

for groups of 200 crystals. These reference photodiodes are stable, radiation hard and insensitive to the magnetic field. The laser monitoring stability is thus controlled by the stability of the relative calibration of injected laser light on each group of crystals. The redundancy of normalization (there are 10 PN photodiodes per supermodule) is used to detect relative miscalibration of PN reference between groups of crystals sharing the same normalization.

### 3. Laser pulsewidth

When illuminated with laser pulses from the monitoring system, the pulse shape at the

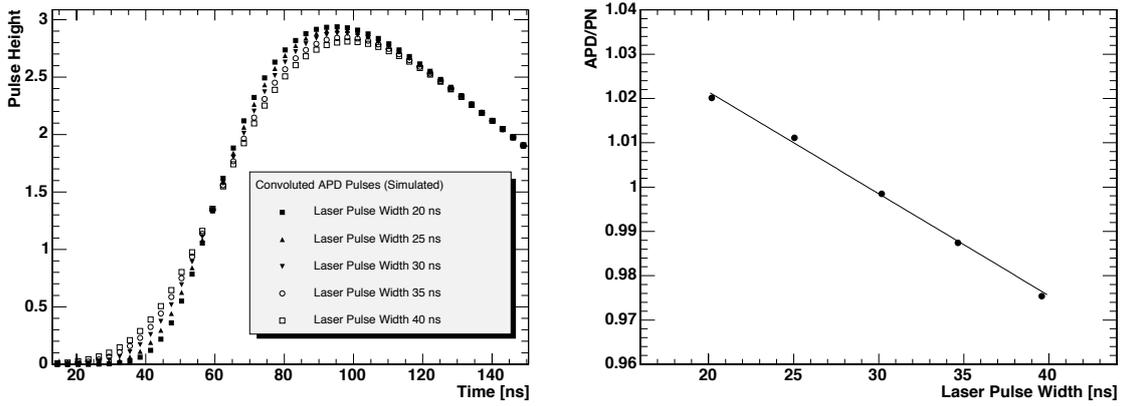


Figure 4.32: (a) Convolutions of the preamplifier shaping and Gaussian functions simulating the laser pulse widths of 20 ns, 25 ns, 30 ns, 35 ns and 40 ns respectively. The PN diode response is stable due to its long shaping time. (b) Variation of the ratio APD/PN, which is used to monitor the change in the crystal transparency, as a function of the laser pulse width.

output of the APD and the PN readout electronics can be described as a convolution of the laser pulse shape and the shaping function of the respective readout systems. The laser pulse shape, measured in the test beam, is Gaussian and the electronics shaping is well described by the formula

$$A(t) = A_0 \left( \frac{x - t_0}{t_{\text{rise}}} \right)^a e^{-a \left( \frac{x - t_0 + t_{\text{rise}}}{t_{\text{rise}}} \right)}, \quad (4.9)$$

where  $a$  and  $t_{\text{rise}}$  are measured on experimental pulse shapes. For the APD electronics  $t_{\text{rise}} \approx 50$  ns and for the PN diode electronics  $t_{\text{rise}} \approx 800$  ns, so that the sensitivity to the laser pulse width is very different between the two.

To investigate the impact of the variation of the laser pulse width this convolution has been simulated for different laser pulse widths. As can be seen in Fig. 4.32a, the convolution results in a significant variation of the maximum height of the signal in the APD. The effect in the PN is negligible due to the long shaping time. Thus the crystal transparency measurement is very sensitive to changes of the laser pulse width as shown in Fig.4.32b. With a laser pulse width of about 30 ns, the sensitivity of the ratio APD/PN is 0.23%/ns.

This effect can be corrected by measuring the laser pulse width event by event with a precision better than 1 ns using the flash ADC at 2 Gs/sec.

### 4. APD gain

A change in the gain of any part of the electronics chain will affect equally the scintil-

lation response and the response to laser pulses. Since the constant of proportionality,  $\alpha$ , between the magnitude of response variation of the laser signal and the scintillation signal is not 1, any such change of gain will lead to systematic errors in the transparency correction if not independently monitored and corrected for. The most significant change of gain expected is for the APD, where the gain is directly related to the high voltage applied. This depends on the HV power supply and the dark current of the APD. Such effects can be controlled by the monitoring of the high voltage and the dark current of each APD (the LSB for this measurement is 322 nA which is sufficient to control the APD gain with a precision better than 0.1%). Such systematic contributions can also be identified by comparing the crystal response to blue and infrared laser light.

#### 4.4.9 Alignment

The ECAL is capable of providing quite precise measurements of shower position, and a resolution of about 0.5 mm is achievable for very high energy showers [160]. Because of the 3° off-pointing angle of the crystal axes the shower-to-shower variation of shower depth results in an uncertainty in lateral shower position approaching 0.5 mm. This sets a limit to the useful precision of ECAL alignment with respect to the tracker.

The key step of the electron HLT selection matches shower position to pixel detector hits. The full width of the selected region is  $\Delta\phi \approx 0.04$  rad, corresponding to about 60 mm in the ECAL, and thus places only very mild demands on the alignment precision required.

Precision measurements made on a number of supermodules during construction show that the variations of the crystal lateral positions with respect to their nominal positions within the supermodule are only a few hundred microns. It thus seems likely that the nominal construction geometry can be taken for all relative positions within a supermodule. Similarly, it is expected that nominal positions will be adequately precise within endcap Dees. Alignment for the ECAL then consists of measuring the position of 36 supermodules and 4 Dees with respect to the tracker. The initial placement errors are not expected to be larger than a few mm. These errors will be rapidly reduced using track/cluster matching of electrons.

#### 4.4.10 Conditions/calibration database

The conditions database stores data which relates to and describes the condition of the detector. For the ECAL, it will contain data from many sources. Much effort has been put into listing and summarizing the type and amount of ECAL data in the conditions database, and it has been estimated that about 1 GB per day of such data will be produced. The Detector Control System (DCS [8]) will store the relevant values of the High Voltages, Low Voltages, Temperatures and other environment sensors and the ECAL Safety System. The temperature sensors which are located on the crystal and the near the photodetectors will be read via the Token Ring, as well as all other temperature sensors on the Very Front End cards (VFE) and Low Voltage Regulator Boards (LVRB) and the APD dark currents. These data will go in the condition database to track changes and produce history plots. The ECAL OD and DAQ will monitor the status of the read-out and will periodically produce summary information that will also be stored to track problems in the read-out and trigger system.

The standard reconstruction requires a few parameters per channel (of the order of 10–20). Many of these parameters have to do with the electronics chain and are unlikely to change

often. Some parameters (like the pedestals) may need to be monitored on a daily basis. The changes in crystal transparency, monitored using the laser system, will require a new value for each crystal every 30 minutes, and interpolation may be needed between 2 consecutive values. Of all values used in the reconstruction, the latter are expected to have the shortest interval of validity.

The detector control system monitors the stability of environment variables (like temperature, high voltage, low voltage, etc.) to verify that they remain constant to a precision which guarantees a stable response. From beam-test experience, it can be expected that these variables will not change significantly over a short time-scale. Thus it is expected that corrections will not need to be applied. If it is necessary the environment variables data will be copied from the online to the offline database, otherwise a simple channel status indicator can be computed to summarize the stable status of each channel.

## 4.5 Monitoring

In this section the different elements which contribute to the ECAL monitoring are described: the detector monitoring and the data quality monitoring. The task of the former is to continuously inspect the operational conditions of the hardware, and that of the latter is to check the detector performances on different categories of collected events. They will allow the detection of either hardware problems or software configuration errors. Any unexpected behaviour needs to be identified and investigated rapidly.

### 4.5.1 Monitoring of operational conditions

#### 4.5.1.1 ECAL detector and safety control system

The ECAL detector control system (DCS) monitors the detector status, in particular various environmental parameters, as well as the detector safety system, which will generate alarms and hardwired interlocks in situations which could lead to damage of the detector hardware. It consists of the following subsystems: Precision Temperature Monitoring (PTM), Humidity Monitoring (HM), ECAL Safety System (ESS), High Voltage (HV), Low Voltage (LV) and monitoring of the laser operation, the cooling system and of the parameters (temperatures in capsules, temperatures on the printed circuit boards, APD leakage currents) read out by the DCUs on the VFE and LVRBs.

A major aspect for the ECAL detector control is the monitoring of the system temperature and the verification that the required temperature stability of the crystal volume and the APDs, expected to be  $(18 \pm 0.05)^\circ\text{C}$ , is achieved. The PTM is designed to read out thermistors, placed on both the front and back of the crystals, with a relative precision better than  $0.01^\circ\text{C}$ . In total there are ten sensors per supermodule. Two immersion probes measure the temperature of the incoming and outgoing cooling water, whereas two sensors per module, one on the grid and one on the thermal screen side of the crystal volume, monitor the crystal temperature. The readout is based on the Embedded Local Monitoring Board (ELMB) developed by ATLAS and completely independent of the DAQ and control links. However, the latter are used to read out further thermistors, mounted in the capsules, via the DCUs placed on the VFE boards.

The purpose of the ESS is to monitor 1) the air temperature of the VFE and FE environment (expected to be around 25–30°C) and 2) the water leakage detection cable, which is routed inside the electronics compartment, 3) to control the proper functioning of the cooling system and 4) to automatically perform pre-defined safety actions and generate interlocks in case of any alarm situation. One pair of temperature sensors is placed at the centre of each module. The completely redundant readout system is independent of the DAQ and control links and based on a Programmable Logic Controller (PLC) situated in the underground service cavern (USC55). It evaluates the data from the temperature sensors, from the water leakage detection system and the signals received from the cooling system, which has its own PLC. If any critical value is detected hardwired interlock signals will be routed to the relevant crates in order to switch off the LV and/or to the cooling PLC in order to stop the water flow on a certain cooling line. The proper functioning of the ESS PLC itself is monitored by the general CMS detector safety system. In addition to this very robust, independent and continuously running system, the temperature of the VFE and LVRBs can be monitored with the DCUs placed on them. These data are read out via the control link and the local DAQ. They will not be used to trigger automatic actions based on hardwired links, but warnings and alarms might be triggered through software for notification of the operators. This is also true for the HM system, which monitors the humidity level in the electronics compartment and consists of one humidity sensor per module and an ELMB-based readout.

The whole DCS software is based on the commercial SCADA package PVSS II. A distributed system is built out of several applications dedicated to the PTM, HM, ESS, the monitoring of the status of the cooling system and the laser operation, as well as the monitoring and control of the HV and LV systems. The DCU data read out via the local DAQ will be written directly to the conditions database, but can be retrieved with user interfaces based on PVSS and/or ROOT. Every application is implemented as a Finite State Machine (FSM) and linked to a supervisory level, which summarizes the overall ECAL DCS status and itself incorporates a FSM. Finally, this ECAL DCS supervisor is linked to the general CMS DCS supervisory node, in order to communicate the status and alarms and to receive commands which are propagated down to the relevant subsystems. In general it is the CMS DCS supervisor which will communicate with the CMS Run Control, but it is possible that the Run Control and/or the local ECAL DAQ system communicate with the ECAL DCS supervisor via the exchange of SOAP (simple object access protocol) messages.

#### 4.5.1.2 Temperature

The number of scintillation photons emitted by the crystals and the amplification of the APD are both temperature dependent. Both variations are negative when the temperature increases and in both cases the magnitude is measured to be about  $-2\%/^{\circ}\text{C}$  [161]. Much effort has been put into ensuring the thermal stability of the ECAL, and measurements indicate that no significant or measurable variation of response due to temperature variation is to be expected during ECAL operation [161]. During operation numerous temperature probes will monitor the temperature of the ECAL, and, in particular, precision measurements from sensors fixed to the back surface of every tenth crystal in the barrel, and one in 25 crystals in the endcap, will be made frequently and regularly.

### 4.5.1.3 Dark current

The APD dark current will increase during CMS operation due to bulk damage of the silicon structure by neutrons. Part of this damage anneals, but the overall effect will be an increase of electronics noise, due to an increasing dark current, over the lifetime of the detector. The dark current of all APD channels will be continuously monitored.

### 4.5.1.4 HV and LV

The APD gain dependence on high voltage is  $\approx 3\%/V$  at the chosen operating gain (50). The High Voltage system has been designed and qualified to maintain the voltage stable over about 30 days with a maximum deviation of 66 mV. This limits the effect on the calibration to below 0.2%. Such stability is obtained using sense wires that measure the voltage at the load, correcting for instability or varying voltage drop over the cables from the HV crates to the experiment. The high voltage boards will require periodic calibration in order to insure the stability over longer periods. The DCS system will monitor the voltage and the current delivered by each channel.

The ECAL amplification and digitization electronics located on the VFE electronics cards require a very stable low voltage to maintain constant signal amplification. The system uses Low Voltage Regulators that guarantee the required stability of the signal amplification. The Low Voltage Regulators Boards are equipped with DCUs that measure the voltages and these measurements are read via the Token Ring.

## 4.5.2 Monitoring using events

The purpose of the ECAL Data Quality Monitor is to follow the time evolution of relevant quantities describing the running conditions of the detector and its performance. The aim is twofold, first, to spot possible problems or system instabilities, and promptly correct for them during the data-taking, and second to monitor the high-level quantities that are needed by the offline reconstruction. The data to be monitored consist not only of standard physics trigger data, but also of special ECAL-only triggers acquired during the normal physics data-taking by triggering during the LHC machine orbit gap. Some of these special triggers will need to be taken outside of normal physics data-taking periods, and the stand-alone VME readout option may be needed for this purpose. These special ECAL-only triggers are

- laser triggers;
- test-pulse triggers;
- empty bunch crossing triggers.

Test-pulse triggers, and empty bunch triggers used to measure the pedestal in the gain ranges other than the highest, require special settings to be downloaded to the FE. So also do laser triggers with the MGPA gain range forced, which will be needed in order to allow the study of the relative gains of the different gain ranges. In normal circumstances special settings will not be downloaded to the FE during physics data taking periods.

The histograms of the “monitorable quantities” will all be made available through a common DQM histogram framework using the Physics and Data Quality Monitoring package (Section 2.9), and the histograms and the results of the monitoring will be available also offline from the ECAL Conditions Database (Section 4.4). A prototype system has already been

commissioned and will be tested at the end of 2005. It will include a *source* and a first set of *clients* as defined in Section 2.9.2.3. They will allow to visualize the histograms, to save them to ordinary ROOT files, and to save the relevant informations to the conditions database for future needs.

From the experience gained to date, they can be classified in different categories:

- Raw data quantities: they include error flags from the DAQ, bad synchronization of the different FE streams, problems in the decoding of the raw data, data integrity, correct read-out order of towers and active channels, reasonable read-out values.
- Single channel raw quantities: the number and location of dead, noisy and in general malfunctioning channels will need to be monitored. This category includes the monitoring of
  - signal pulse phase: the time of signal pulse maximum ( $t_{\max}$ ) needs to be monitored for both physics and laser events.
  - pedestal events: the mean values and widths of the pedestals will be monitored, for example using trace plots. Given the large number of channels, it will be useful to monitor the average deviation of pedestals and widths of all channels, comparing them to the last high statistics pedestal run taken before the current fill. For some channels, correlations between samples and between channels will be computed and monitored.
  - laser events: the response of each channel to the injection of a laser light (normalized to the laser light detected by reference PN diodes) will be monitored as a function of time.
- Global raw quantities: the operation of selective readout will need to be monitored and validated. The monitoring of the zero suppression algorithms will be possible using a map of active channels, averaged over several events.
- Reconstructed quantities: after pedestal subtraction, gain correction, and energy reconstruction, it will be possible to identify problematic channels using a map of the average energy deposition.
- Physics quantities: several useful higher level quantities will need to be monitored in standard physics triggers. A short and incomplete list includes:
  - average number of crystals above a given threshold per event
  - average number of energy clusters per supermodule
  - average number of crystals per energy cluster
  - average number of energy clusters per event
  - raw/calibrated energy per trigger tower
  - average energy per event in ( $\phi$ ,  $\eta$ ) bins
  - average energy per event
  - pedestal values (only pedestals in the lowest gain range are available in normal running conditions)

The final set of quantities to be monitored is expected to be fully verified, tested and possibly further extended during the ECAL precalibration with cosmic muons, the magnet test, and the test beam data taking in 2006.