

## Chapter 4


# Electromagnetic calorimeter

The electromagnetic calorimeter of CMS (ECAL) is a hermetic homogeneous calorimeter made of 61 200 lead tungstate ( $\text{PbWO}_4$ ) crystals mounted in the central barrel part, closed by 7324 crystals in each of the two endcaps. A preshower detector is placed in front of the endcap crystals. Avalanche photodiodes (APDs) are used as photodetectors in the barrel and vacuum phototriodes (VPTs) in the endcaps. The use of high density crystals has allowed the design of a calorimeter which is fast, has fine granularity and is radiation resistant, all important characteristics in the LHC environment. One of the driving criteria in the design was the capability to detect the decay to two photons of the postulated Higgs boson. This capability is enhanced by the good energy resolution provided by a homogeneous crystal calorimeter.


### 4.1 Lead tungstate crystals

The characteristics [62] of the  $\text{PbWO}_4$  crystals make them an appropriate choice for operation at LHC. The high density ( $8.28 \text{ g/cm}^3$ ), short radiation length (0.89 cm) and small Molière radius (2.2 cm) result in a fine granularity and a compact calorimeter. In recent years,  $\text{PbWO}_4$  scintillation properties and other qualities have been progressively improved, leading to the mass production of optically clear, fast and radiation-hard crystals [63, 64]. The scintillation decay time of these production crystals 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 and varies with temperature ( $-2.1\% \text{C}^{-1}$  at  $18^\circ\text{C}$  [65]): at  $18^\circ\text{C}$  about 4.5 photoelectrons per MeV are collected in both APDs and VPTs. The crystals emit blue-green scintillation light with a broad maximum at 420–430 nm [64, 66]. Longitudinal optical transmission and radioluminescence spectra are shown in figure 4.1.

To exploit the total internal reflection for optimum light collection on the photodetector, the crystals are polished after machining. For fully polished crystals, the truncated pyramidal shape makes the light collection non-uniform along the crystal length. The effect is large because of the high refractive index ( $n = 2.29$  around the peak wavelength [67]) and the needed uniformity [68] is achieved by depolishing one lateral face. In the endcaps, the light collection is naturally more uniform because the crystal faces are nearly parallel. Pictures of barrel and endcap crystals with the photodetectors attached are shown in figure 4.2.



**Figure 4.1:** Longitudinal optical transmission (1, left scale) and radioluminescence intensity (2, right scale) for production  $\text{PbWO}_4$  crystals.



**Figure 4.2:**  $\text{PbWO}_4$  crystals with photodetectors attached. Left panel: A barrel crystal with the upper face depolished and the APD capsule. In the insert, a capsule with the two APDs. Right panel: An endcap crystal and VPT.

The crystals have to withstand the radiation levels and particle fluxes [69] 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 (section 4.9). The damage reaches a dose-rate dependent equilibrium level which results from a balance between damage and recovery at  $18^\circ\text{C}$  [64, 70].

To ensure an adequate performance throughout LHC operation, the crystals are required to exhibit radiation hardness properties quantified as an induced light attenuation length (at high dose rate) greater than approximately 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 conditions indicates that the damage will remain within the limits required for good ECAL performance [71, 72].

## 4.2 The ECAL layout and mechanics

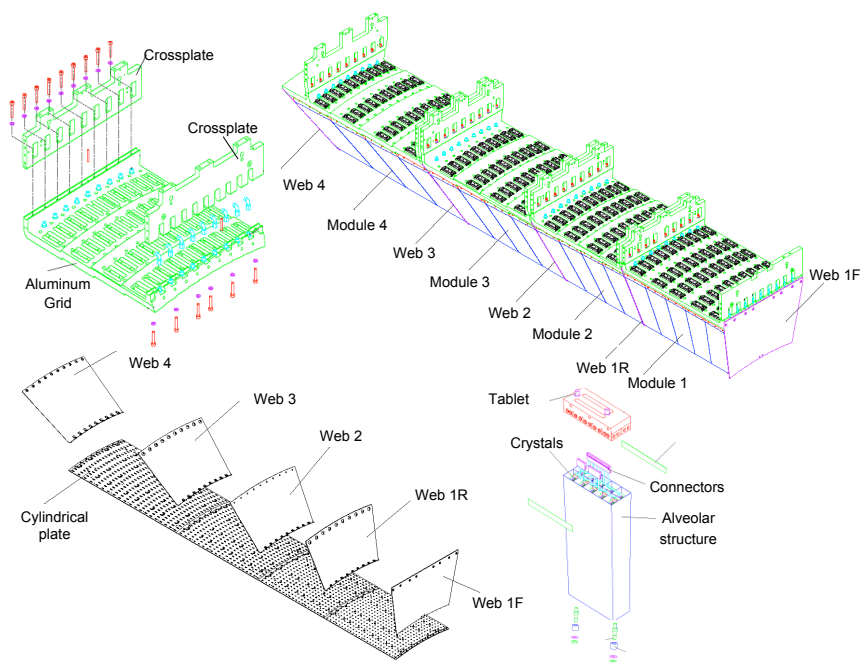
The barrel part of the ECAL (EB) covers the pseudorapidity range  $|\eta| < 1.479$ . The barrel granularity is 360-fold in  $\phi$  and  $(2 \times 85)$ -fold in  $\eta$ , resulting in a total of 61 200 crystals. The crystals have a tapered shape, slightly varying with position in  $\eta$ . They are mounted in a quasi-projective geometry to avoid cracks aligned with particle trajectories, so that their axes make a small angle ( $3^\circ$ ) with respect to the vector from the nominal interaction vertex, in both the  $\phi$  and  $\eta$  projections. The crystal cross-section corresponds to approximately  $0.0174 \times 0.0174$  in  $\eta$ - $\phi$  or  $22 \times 22$  mm<sup>2</sup> at the front face of crystal, and  $26 \times 26$  mm<sup>2</sup> at the rear face. The crystal length is 230 mm corresponding to  $25.8 X_0$ . The barrel crystal volume is  $8.14 \text{ m}^3$  and the weight is 67.4 t.

The centres of the front faces of the crystals are at a radius 1.29 m. The crystals are contained in a thin-walled alveolar structure (submodule). The alveolar wall is 0.1 mm thick and is made of an aluminium layer, facing the crystal, and two layers of glass fibre-epoxy resin. To avoid oxidation, a special coating is applied to the aluminium surface. The nominal crystal to crystal distance is 0.35 mm inside a submodule, and 0.5 mm between submodules. To reduce the number of different types of crystals, each submodule contains only a pair of shapes, left and right reflections of a single shape. In total, there are 17 such pairs of shapes. The submodules are assembled into modules of different types, according to the position in  $\eta$ , each containing 400 or 500 crystals. Four modules, separated by aluminium conical webs 4-mm thick, are assembled in a supermodule, which contains 1700 crystals (figures 4.3 and 4.4).

In each module, the submodules are held in partial cantilever by an aluminium grid, which supports their weight from the rear. At the front the submodule free ends are connected together by pincers that cancel the relative tangential displacements. The submodule cantilever is reduced by the action of a 4-mm thick cylindrical plate where the front of the submodules are supported by setpins. Not all the submodules are connected to the cylindrical plate but only four rows in  $\phi$  from a total of ten. The portion of the submodule load taken at the front by the cylindrical plate is transmitted to the aluminium grids of the different modules via the conical webs interspaced between the modules [73]. Each module is supported and positioned in the supermodule at the rear end through the grid by a spine beam. The spine is provided with pads which slide into rails housed on the front face of the HCAL barrel, allowing the installation and support of each single supermodule. The cylindrical plate in front of the supermodule also provides the fixation of the monitoring system (see below) and the holes for its optical fibres.

All services, cooling manifolds and cables converge to a patch panel at the external end of the supermodule. Eighteen supermodules, each covering  $20^\circ$  in  $\phi$ , form a half barrel.


The endcaps (EE) cover the rapidity range  $1.479 < |\eta| < 3.0$ . The longitudinal distance between the interaction point and the endcap envelope is 315.4 cm, taking account of the estimated



**Figure 4.3:** Layout of the ECAL barrel mechanics.

shift toward the interaction point by 1.6 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*. Each Dee holds 3 662 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 1 300 mm beyond the interaction point, giving off-pointing angles ranging from 2 to 8 degrees. 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 a length of 220 mm ( $24.7 X_0$ ). The endcaps crystal volume is  $2.90 \text{ m}^3$  and the weight is 24.0 t. The layout of the calorimeter is shown in figure 4.5. Figure 4.6 shows the barrel already mounted inside the hadron calorimeter, while figure 4.7 shows a picture of a Dee.

The number of scintillation photons emitted by the crystals and the amplification of the APD are both temperature dependent. Both variations are negative with increasing temperature. The overall variation of the response to incident electrons with temperature has been measured in test beam [74] to be  $(-3.8 \pm 0.4) \% \text{ C}^{-1}$ . The temperature of the system has therefore to be maintained constant to high precision, requiring a cooling system capable of extracting the heat dissipated by the read-out electronics and of keeping the temperature of crystals and photodetectors stable within  $\pm 0.05^\circ \text{C}$  to preserve energy resolution. The nominal operating temperature of the CMS ECAL is  $18^\circ \text{C}$ . The cooling system has to comply with this severe thermal requirement. The system employs water flow to stabilise the detector. In the barrel, each supermodule is independently supplied




**Figure 4.4:** Front view of a module equipped with the crystals.


with water at 18°C. The water runs through a thermal screen placed in front of the crystals which thermally decouples them from the silicon tracker, and through pipes embedded in the aluminium grid, connected in parallel. Beyond the grid, a 9 mm thick layer of insulating foam (Armaflex) is placed to minimise the heat flowing from the read-out electronics towards the crystals. Return pipes distribute the water through a manifold to a set of aluminium cooling bars. These bars are in close contact with the very front end electronics (VFE) cards and absorb the heat dissipated by the components mounted on these cards. A thermally conductive paste (gap filler 2000, produced by Bergquist) is used to provide a good contact between the electronic components and a metal plate facing each board. This plate is coupled to the cooling bar by a conductive pad (ultrasoft gap pad, also produced by Bergquist). Both the gap pad and the gap filler have been irradiated with twice the dose expected in the ECAL endcaps after 10 years at the LHC and have shown no change in character or loss of performance.

Extended tests of the cooling system have been performed with good results [74]. Residual effects caused by a possible variation of the power dissipated by the electronics were measured in the extreme case of electronics switched on and off. The conclusion is that contributions to the constant term of the energy resolution due to thermal fluctuations will be negligible, even without temperature corrections.





**Figure 4.5:** Layout of the CMS electromagnetic calorimeter showing the arrangement of crystal modules, supermodules and endcaps, with the preshower in front.



**Figure 4.6:** The barrel positioned inside the hadron calorimeter.



**Figure 4.7:** An endcap Dee, fully equipped with supercrystals.

### 4.3 Photodetectors

The photodetectors need to be fast, radiation tolerant and be able to operate in the longitudinal 4-T magnetic field. In addition, because of the small light yield of the crystals, they should amplify and be insensitive to particles traversing them (*nuclear counter effect*). The configuration of the magnetic field and the expected level of radiation led to different choices: avalanche photodiodes in the barrel and vacuum phototriodes in the endcaps. The lower quantum efficiency and internal gain of the vacuum phototriodes, compared to the avalanche photodiodes, is offset by their larger surface coverage on the back face of the crystals.

#### 4.3.1 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 a pair is mounted on each crystal. They are operated at gain 50 and read out in parallel. The main properties of the APDs at gain 50 and  $18^\circ\text{C}$  are listed in table 4.1.

The sensitivity to the 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  $\text{PbWO}_4$ .

**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 (430 nm)	$75 \pm 2\%$
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	$< 10 \Omega$
Voltage sensitivity of the gain ( $1/M \cdot dM/dV$ )	$3.1 \pm 0.1\%/V$
Temperature sensitivity of the gain ( $1/M \cdot dM/dT$ )	$-2.4 \pm 0.2\%/^\circ\text{C}$
Rise time	$< 2 \text{ ns}$
Dark current	$< 50 \text{ nA}$
Typical dark current	3 nA
Dark current after $2 \times 10^{13} \text{ n/cm}^2$	$5 \mu\text{A}$

For ECAL acceptance 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 [75]. Based on tests with hadron irradiations [76] 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 in selecting radiation resistant APDs.

The gain stability directly affects the ECAL energy resolution. Since the APD gain has a high dependence on the bias voltage ( $\alpha_V = 1/M dM/dV \simeq 3.1\%/V$  at gain 50), to keep this contribution to the resolution at the level of per mille, the APDs require a very stable power supply system: the stability of the voltage has to be of the order of few tens of mV. This requirement applies to all the electrical system characteristics: noise, ripple, voltage regulation and absolute precision, for short and long term periods. A custom high voltage (HV) power supply system has been designed for the CMS ECAL in collaboration with the CAEN Company [77]. To remain far from high doses of radiation, the HV system is located in the CMS service cavern, some 120 m away from the detector. The HV channels are floating and use sense wires to correct for variations in the voltage drop on the leads. The system is based on a standard control crate (SY1527) hosting 8 boards expressly designed for this application (A1520E). The SY1527 integrate a PC capable of communicating with the board controller via an internal bus and different interfaces are available to integrate the SY1527 on the ECAL detector control system (DCS). The board design is based on a modular concept so that each HV channel is implemented on a separate module and up to 9 channels can be hosted on a single A1520E board. Each channel can give a bias voltage to 50 APD pairs from 0 to 500 V with maximum current of 15 mA. In total, there are 18 crates and 144 boards. Temperature



drift compensation is possible due to the presence on the crate of temperature probes that can be used to monitor the environment temperature for adjustments of the voltage setting.

The operating gain of 50 requires a voltage between 340 and 430 V. 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. Each pair is mounted in parallel in a *capsule*, a moulded receptacle with foam, which is then glued on the back of each crystal. The capsules are connected to the read-out electronics by Kapton flexible printed circuit boards of variable length, dictated by the capsule's position within the submodule. Each capsule receives the bias voltage through an RC filter network and a protection resistor.

One 100 k $\Omega$  negative temperature coefficient thermistor from Betatherm, used as temperature sensor, is embedded in every tenth APD capsule. There are twenty-two different types of capsules, differing by the Kapton length and by the presence of the thermistor.

### 4.3.2 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 [78] 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, with an active area of approximately 280 mm<sup>2</sup>; one VPT is glued to the back of each crystal. One Betatherm thermistor is embedded into each supercrystal. The VPTs have a mean quantum efficiency of the bialkali photocathode (SbKCs) of 22% at 430 nm, and a mean gain of 10.2 at zero field. 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° to 26°). The mean response in a magnetic field of 4 T, with the VPT axis at 15° to the field direction, is typically > 90% of that in zero field [79].


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°, 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 kGy and  $7 \times 10^{14}$  n/cm<sup>2</sup> at  $|\eta| = 2.6$ . Sample faceplates from every glass production batch were irradiated with a <sup>60</sup>Co source to 20 kGy. The faceplates were required to show a transmission loss, integrated over the wavelength range corresponding to PbWO<sub>4</sub> emission, of less than 10%. Irradiation of VPTs in a nuclear reactor to  $7 \times 10^{14}$  n/cm<sup>2</sup> showed a loss in anode sensitivity entirely consistent with discolouration of the faceplate caused by the accompanying gamma dose (100 kGy) [80]. Irradiations of tubes biased to the working voltage, with both gammas and neutrons showed no adverse effects, apart from an increase in anode current, attributable to the production of Cerenkov light in the faceplates.

The VPTs are operated with the photocathode at ground potential and the dynode and anode biased at +600 V and +800 V respectively. The high voltage system is based (like the APD system) on CAEN SY1527 standard control crates, although for the VPTs, the crates are equipped with standard 12-channel A1735P boards, each channel rated at 1.5 kV and 7 mA. At the operating bias, the VPT gain is close to saturation thus the voltages do not have to be controlled very precisely. However, care must be taken to minimise ripple and noise, since these would feed directly into the input of the sensitive preamplifier that is connected to the anode. Filtering is achieved with RC networks mounted inside the supercrystals (SC), close to the VPTs. An entire endcap is biased using one SY1527 crate equipped with just two A1735P boards. On each board, only eight of the twelve output channels will initially be used, leaving four spare channels. The spare outputs may be used at a later stage, if noisy channels develop which can be recovered by operating at a lower bias voltage. The HV from the CAEN power supplies is transmitted to the SCs via a custom designed HV distribution system which provides hard-wired protection against over-voltage and over-current, and sensitive current monitoring. For each endcap, this system is housed in five crates. Each crate hosts up to five input cards, receiving the HV from the power supplies, and up to six output cards, with each output card serving up to twelve SCs. The HV supplies and distribution system are mounted in two racks (one for each endcap) located in the Service Cavern. Each SC is served by two coaxial cables (one for the anode, one for the dynode) running from the Service Cavern to the detector, via intermediate patch panels. The total cable length is approximately 120 m and the cable capacitance forms part of the filter network. Inside an SC the HV is distributed to the VPTs via five filter cards, each serving five VPTs. The spread in anode sensitivity among the VPTs is 25% (RMS). They are therefore sorted into six groups which are distributed on the endcaps with the highest sensitivities at the outer circumference grading to the lowest sensitivities at the inner circumference. This arrangement provides a roughly constant sensitivity to the transverse energy across the endcaps.

The anode sensitivity of a VPT may show a dependence on count rate (anode current) under certain conditions. For example, in the absence of a magnetic field, if the count rate falls to a few Hz, following a period of high rate operation, the anode sensitivity may rise suddenly and take several hours to return to the nominal value. The magnitude of the effect may vary from a few percent to a few tens of percent. In the presence of a strong magnetic field (as in normal CMS operation), the effect is strongly suppressed or absent. Nevertheless, it has been judged prudent to incorporate a light pulser system on the ECAL endcaps. This delivers a constant background rate of at least 100 Hz of pulses of approximately 50 GeV equivalent energy to all VPTs, thus ensuring that they are kept “active”, even in the absence of LHC interactions.

The system consists of a control and trigger unit located in the Service Cavern, and sets of pulsed light sources mounted on the circumference of each Dee. The light is produced by Luxeon III light emitting diodes (type LXHL-PR09), whose peak emission wavelength is 455 nm. The LEDs are driven by high output current op-amps (LT6300 from Linear Technology). The drive pulses have amplitudes of 1.2 A and a widths of 80 ns. A single light source consists of a cluster of seven LEDs and associated drive-circuits. These are configured singly or in pairs, with the drive-circuits and LEDs mounted on double-sided printed circuit boards housed within metal enclosures. There are four such enclosures distributed around the circumference of each Dee, housing 19 light sources. A schematic representation of the system for distributing the light pulses is shown in figure 4.8.



**Figure 4.8:** Distribution system for VPT stabilisation light pulses.

An all-silica optical fibre (CF01493-43 from OFS (Furukawa)) is inserted into a hole drilled into the lens of each LED and collects light by *proximity focusing*. The seven fibres from a given light source are combined into a single bundle that transports light to a diffusing sphere which has a dual role, acting also as part of the distribution network of the laser monitoring system. Light from each diffusing sphere is distributed to up to 220 individual detector channels through the set of optical fibres that also carry the laser monitoring pulses. Light is injected via the rear face of a crystal, which carries the VPT, and reaches the VPT via reflection from the front of the crystal. The system is synchronized to pulse during a fraction of the  $3 \mu\text{s}$  abort gaps that occur during every  $89 \mu\text{s}$  cycle of the LHC circulating beams.

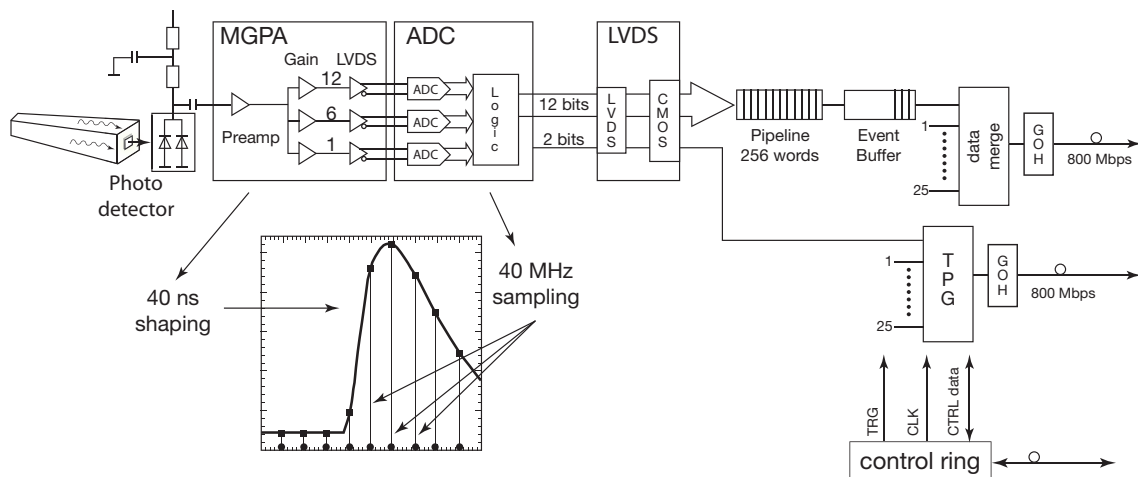
#### 4.4 On-detector electronics

The ECAL read-out has to acquire the small signals of the photo-detectors with high speed and precision. Every bunch crossing digital sums representing the energy deposit in a trigger tower are generated and sent to the trigger system. The digitized data are stored during the Level-1 trigger latency of  $\approx 3 \mu\text{s}$ .

The on-detector electronics has been designed to read a complete trigger tower ( $5 \times 5$  crystals in  $\eta \times \phi$ ) or a super-crystal for EB and EE respectively. It consists of five Very Front End (VFE) boards, one Front End (FE) board, two (EB) or six (EE) Gigabit Optical Hybrids (GOH), one Low Voltage Regulator card (LVR) and a motherboard.

The motherboard is located in front of the cooling bars. It connects to 25 photo-detectors and to the temperature sensors using Kapton flexible printed circuit boards and coaxial cables for EB and EE respectively. In the case of the EB the motherboard distributes and filters the APD bias voltage. Two motherboards are connected to one CAEN HV supply located at a distance of about 120m with remote sensing. In the case of the EE the operating voltages for the VPTs are distributed and filtered by a separate HV filter card, hosting as well the decoupling capacitor for the anode signals. Five of these cards serving five VPTs each are installed into each super-crystal. One LVR and five VFE cards plug into the motherboard.

Each LVR card [81] uses 11 radiation-hard low voltage regulators (LHC4913) developed by ST-microelectronics and the RD49 project at CERN. The regulators have built in over-temperature



**Figure 4.9:** Schematic view of the on-detector electronics: the scintillation light is collected by photodetectors (in the figure the case of APD is presented), the signal is shaped by a Multi-Gain Pre-Amplifier and digitized by 40-MHz ADC; a radiation-hard buffer (LVDS) adapts the ADC output to the FE card, where data pipeline and Trigger Primitives Generation (TPG) are performed; trigger words are sent at 25 ns rate, while data are transmitted on receipt of a Level-1 trigger; GOHs provide in both cases the data serializer and the laser diode, sending the signals on a fibre to the off-detector electronics over a distance of about 100 m. A control token ring connects groups of FE cards, providing Level-1 trigger (TRG) and clock (CLK) signals, together with control data in and out (CTRL data).

protection, output current limitation and an inhibit input. The output voltages of 2.5 V are distributed to the FE card and via the motherboard to the VFE cards. Three Detector Control Unit (DCU) ASICs on each LVR card, interfaced to the FE card, monitor all input and output voltages. All regulators, excluding the one providing power to the control interface of the FE card, can be powered down remotely by an external inhibit. Four LVR cards are connected by a passive low voltage distribution (LVD) block to one radiation and magnetic field tolerant Wiener low voltage power supply located about 30 m away in racks attached to the magnet yoke.

The signals are pre-amplified and shaped and then amplified by three amplifiers with nominal gains of 1, 6 and 12. This functionality is built into the Multi Gain Pre-Amplifier (MGPA) [82], an ASIC developed in  $0.25\ \mu\text{m}$  technology. The full scale signals of the APDs and VPTs are 60 pC and 12.8 pC corresponding to  $\approx 1.5\ \text{TeV}$  and 1.6–3.1 TeV for EB and EE respectively. The shaping is done by a CR-RC network with a shaping time of  $\approx 40\ \text{ns}$ . The MGPA has a power consumption of 580 mW at 2.5 V. The output pulse non-linearity is less than 1%. The noise for gain 12 is about  $8000e^-$  for the APD configuration and about  $4000e^-$  for the VPT configuration. The MGPA contains three programmable 8-bit DACs to adjust the baseline to the ADC inputs. An integrated test-pulse generator with an amplitude adjustable by means of an 8-bit DAC allows a test of the read-out electronics over the full dynamic range.

A schematic view of the signal read-out is given in figure 4.9. The 3 analog output signals of the MGPA are digitized in parallel by a multi-channel, 40-MHz, 12-bit ADC, the AD41240 [83], developed in  $0.25\ \mu\text{m}$  technology. It has an effective number of bits of 10.9. An integrated logic



selects the highest non-saturated signal as output and reports the 12 bits of the corresponding ADC together with two bits coding the ADC number.

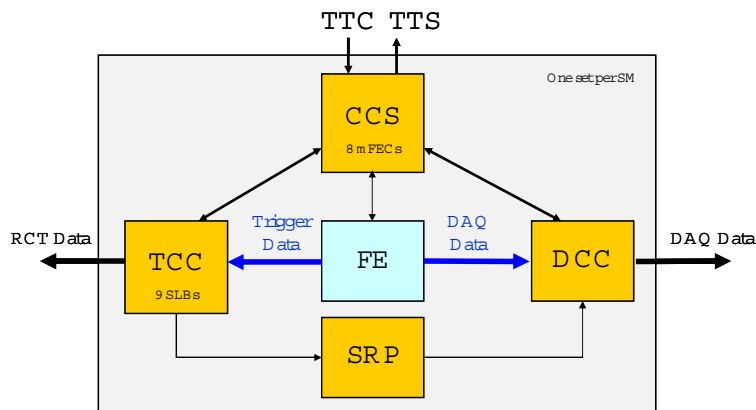
If the read-out switches to a lower gain as the pulse grows, it is prevented from immediately reverting to the higher gain when the pulse falls: once the pulse has declined to the point where it could be read out at the higher gain again, the read-out is then forced to continue reading out at the lower gain for the next five samples.

A radiation-hard buffer (LVDS\_RX) developed in  $0.25\ \mu\text{m}$  technology, adapts the low voltage differential output signals of the AD41240 to the single ended CMOS inputs on the FE card. Five identical read-out channels are integrated into a VFE card, together with a Detector Control Unit (DCU) for the measurement of the APD leakage currents and the read-out of the thermistors. The noise obtained with the VFE cards installed into supermodules is typically 1.1, 0.75 and 0.6 ADC counts for gains 12, 6 and 1 respectively. This corresponds to  $\approx 40\ \text{MeV}$  for gain 12.

The FE card [84] stores the digitized data during the Level-1 trigger latency in 256-word-deep dual-ported memories, so called pipelines. Five such pipelines and the logic to calculate the energy sum of the 5 channels once every bunch crossing are integrated into an ASIC developed in  $0.25\ \mu\text{m}$  technology called FENIX. Each VFE card is serviced by a FENIX chip. Thus the energy is summed in strips of 5 crystals along  $\phi$ . In the case of the EE the five strip sums are transmitted by five GOHs (see below) to the off-detector electronics Trigger Concentrator Card (TCC), while in the case of the EB a sixth FENIX sums the five strip sums and calculates the “fine-grain” electromagnetic bit, set to identify electromagnetic shower candidates on the basis of the energy profile of the trigger tower. The trigger tower energy sum together with the fine-grain bit is transmitted using one GOH to the TCC. On receipt of a Level-1 trigger the corresponding data, ten 40-MHz samples per channel, are transmitted in  $\approx 7.5\ \mu\text{s}$  to the off-detector electronics Data Concentrator Card (DCC) using an identical GOH. The Clock and Control Unit (CCU) ASIC together with the LVDS\_MUX ASIC provide the interface to the token rings.

The ECAL serial digital data links are based on the technology developed for the CMS Tracker analog links (section 3.3). The GOH consists of a data serializer and laser driver chip, the GOL, and a laser diode with an attached fibre pigtail. Fibres, fibre interconnections and a 12-channel NGK receiver module complete the optical link system. It uses single mode fibres operating at 1310 nm wavelength over a distance of about 100 m. The fibre attenuation of  $\approx 0.04\ \text{dB}$  is negligible. The optical links are operated at 800 Mbit/s.

The VFE and FE electronics are controlled using a 40-MHz digital optical link system, controlled by the off-detector Clock and Control System (CCS) boards. A 12-fibre ribbon is connected to the token ring link board, generating an electrical control ring, the token ring. Each supermodule has 8 token rings which connect to groups of eight to ten FE cards including the two FE cards of the laser monitoring electronics module (MEM). The system has redundancy, as long as there are no two consecutive FE cards malfunctioning, by means of two independent bi-directional optical links, using 4 fibres each. It provides fast and slow control functions. While the fast control transmits the level one trigger information and the 40-MHz clock, the slow control comprises the configuration of the FE and VFE electronics as well as the read-out of status information, temperatures, voltages and APD leakage currents.



**Figure 4.10:** Schematic view of ECAL off-detector electronics.

## 4.5 Off-detector electronics

### 4.5.1 Global architecture

The ECAL off-detector read-out and trigger architecture [85, 86] is illustrated schematically in figure 4.10. 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 read-out processor, SRP, system). The system serves both the DAQ and the trigger paths. In the DAQ path, the DCC performs data read-out and data reduction based on the selective read-out flags computed by the SRP system. In the trigger path, at each bunch crossing, trigger primitive generation started in the FE boards is 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 system. The TTC signals are translated and encoded by suppression of clock edges and sent to the *mezzanine* Front End Controller cards (mFEC). The mFEC interfaces optically with a FE token ring. The 8 mFECs of the CCS board control a supermodule.

The trigger concentration card (TCC) [87] main functionalities include the completion of the trigger primitive generation and their transmission to the synchronization and link board (SLB) mezzanines [88] at each bunch crossing, the classification of each trigger tower and its transmission to the Selective Read-out 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  $20^\circ \phi$  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 the fine-grain bit) 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) [89, 90] is responsible for collecting crystal data from up to 68 FE boards. Two extra FE links are dedicated to the read-out of laser monitoring data (PN diodes). The DCC also collects trigger data transmitted from the TCC modules and the selective read-out flags transmitted from the SRP system. A data suppression factor near 20 is attained using a programmable selective read-out algorithm. When operating in the selective read-out mode the SRP flags indicate the level of suppression that must be applied to the crystal data of a given FE read-out. 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 a S-LINK64 data link interface 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 read-out 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 read-out processor (SRP) [91] is responsible for the implementation of the selective read-out algorithm. The system is composed by a single 6U-VME crate with twelve identical algorithm boards (AB). The AB computes the selective read-out flags in different calorimeter partitions. The flags are composed of 3 bits, indicating the suppression level that must be applied to the associated read-out units.

#### 4.5.2 The trigger and read-out paths

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 the fine-grain bit, which characterizes the lateral profile 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 read-out system is structured in 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 4.5.3).

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. Each FE is served by two optical links for sending the data and trigger primitives respectively and a third electrical serial link which transmits the clock, control and Level-1 trigger signals.

In the endcaps, the read-out 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 crystals. 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 read-out 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 6 optical links, 5 of them being used to transmit the trigger primitives. As in the barrel an electrical serial link transmits the clock, control and Level-1 trigger signals.

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

### 4.5.3 Algorithms performed by the 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 bunchcrossing 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 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 fine-grain 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.

### 4.5.4 Classification performed by the selective read-out

About 100kB 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 read-out*, can be performed by the Selective Read-out Processor [86, 91] 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 read-out



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 read-out 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 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.

These classifications can be used flexibly to implement a range of algorithms by using different thresholds to define the classes, and different suppression levels for the read-out 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 ( $E_T > 5 \text{ GeV}$ ) then the crystals of this trigger tower and of its neighbour trigger towers (225 crystals in total in the barrel case) are read with no zero suppression. If a trigger tower belongs to the medium interest class ( $E_T > 2.5 \text{ GeV}$ ), 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 purposes, the selective read-out can be deactivated and either a global zero suppression (same threshold for every channel) or no zero suppression applied. Even when the selective read-out is not applied the selective read-out flags are inserted into the data stream and can be used offline for debugging purposes.


## 4.6 Preshower detector

The principal aim of the CMS Preshower detector 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 high granularity.

### 4.6.1 Geometry

The Preshower is a sampling calorimeter with two layers: lead radiators initiate electromagnetic showers from incoming photons/electrons whilst silicon strip sensors placed after each radiator measure the deposited energy and the transverse shower profiles. The total thickness of the Preshower is 20 cm.

The material thickness of the Preshower 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 two 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



**Figure 4.11:** Photograph of a complete type-1 ladder, with an inset showing details of a micro-module.

critical, and thus a circular shape has been chosen. The lead planes are arranged in two Dees, one on each side of the beam pipe, with the same orientation as the crystal Dees.

Each silicon sensor measures  $63 \times 63 \text{ mm}^2$ , with an active area of  $61 \times 61 \text{ mm}^2$  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 3.6 fC of charge in this thickness (at normal incidence). The sensors are precisely glued to ceramic supports, which also support the front-end electronics assembly (see below), 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. In order to improve noise performance the tile is constructed in two parts, with a glass fibre insulation in between. The combination of sensor + front-end electronics + supports is known as a micromodule.

The micromodules are placed on baseplates in groups of 7, 8 or 10 that, when coupled to an electronics system motherboard (SMB) 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 two Dees join this spacing is increased to 3.0 mm.

Figure 4.11 shows a complete ladder (*Type-1* for 8 micromodules) and an inset shows the micromodule.

The ladders are attached to the radiators in an  $x$ - $y$  configuration. Around 500 ladders are required, corresponding to a total of around 4 300 micromodules and 137 000 individual read-out channels. Further details of the layout can be found in [92].

## 4.6.2 Preshower electronics

Each silicon sensor is DC-coupled to a front-end ASIC (PACE3 [93]) that performs preamplification, signal shaping and voltage sampling. Data is clocked into an on-chip 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 time 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 large dynamic range (0-1600 fC) with a S/N of around 3 for a single MIP;
- High gain: For MIP calibration purposes [94], with a reduced dynamic range (0-200 fC) but with a S/N approaching 10 for a single MIP.

The PACE3 are soldered to front-end hybrids that contain embedded polyimide cables to connect to the SMBs. The SMBs contain AD41240 12-bit ADCs that digitize the data from 1 or 2 PACE3. The digital data are then formatted and packaged by a second Preshower ASIC known as the K-chip [95]. 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 Preshower-DCC (see below) via gigabit optical hybrids (GOH). The SMB also contains an implementation of the CMS tracker control system.

Groups of up to 12 ladders are connected via polyimide cables to form control rings. Off-detector CCS cards (identical to those of the ECAL except not all FEC mezzanines are mounted for the Preshower) communicate via digital optical hybrids (DOH) mounted on 2 of the SMBs in each control ring. The full Preshower comprises 4 planes of 12 control rings each.

The Preshower-DCC [96] is based on the DCC of the ECAL except it is a modular design incorporating a VME host board mounted with optoRx12 [97] mezzanines. The modular design has allowed a development collaboration with the TOTEM experiment which uses the same components but in a different manner. The optoRx12 incorporates an NGK 12-way optical receiver and an Altera Stratix GX FPGA that performs data deserialization, pedestal subtraction, common-mode noise reduction, bunch crossing assignment, charge reconstruction and zero suppression [98]. The sparsified data from up to 3 optoRx12 are merged by another FPGA on the host board that then transmits data packets to the event builder via an Slink64 interface. The host board also provides data spying as well as TTC and VME interfaces. A provision has been made on the host board to allow the plug-in of an additional mezzanine board mounted with FPGAs/processors that could provide more data reduction power if necessary in the future.

## 4.7 ECAL detector control system

The ECAL Detector Control System (DCS) comprises the monitoring of the detector status, in particular various kinds of environmental parameters, as well as the ECAL safety system (ESS), which will generate alarms and hardwired interlocks in case of situations which could lead to damaging the detector hardware. It consists of the following sub-systems: ECAL Safety System (ESS), Precision Temperature Monitoring (PTM), Humidity Monitoring (HM), 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 LVR boards. Further details on the ECAL DCS are available [99].

The whole DCS software is based on the commercial SCADA package PVSS II (chapter 9). A distributed system is built out of several applications dedicated to the DCS sub-systems. 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 sub-systems.

#### 4.7.1 Safety system

The purpose of the ESS [100] is to monitor the air temperature of the VFE and FE environment (expected to be around 25–30°C) and the water leakage detection cable, which is routed inside the electronics compartment, to control the proper functioning of the cooling system and 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 read-out system, with full built-in redundancy, is independent of the DAQ and control links and based on a Programmable Logic Controller (PLC) situated in the Service Cavern. In case of any critical reading hardwired interlock signals will be routed to the relevant crates in order to switch off the HV and LV and/or 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.

#### 4.7.2 Temperature

The number of scintillation photons emitted by the crystals and the amplification of the APD are both temperature dependent, as described in section 4.2. Therefore a major task for the ECAL DCS is the monitoring of the system's temperature and the verification that the required temperature stability of  $(18 \pm 0.05)^\circ\text{C}$  of the crystal volume and the APDs is achieved. The PTM is designed to read out thermistors, placed on both sides of the crystal volume, 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 read-out is based on the Embedded Local Monitoring Board (ELMB) developed by ATLAS which functions completely independently of the DAQ and control links. In addition, sensors fixed to the back surface of every tenth crystal in the barrel, and one in 25 crystals in the endcap, are read out by the DCUs placed on the VFE boards. With this temperature monitoring it has been shown that the water cooling system can indeed ensure the required temperature stability [74].

#### 4.7.3 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 in 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.



#### 4.7.4 HV and LV

The DCS system operates the CAEN HV system via an OPC server. The functionalities include the independent configuration of the HV channels with various set of voltages, the monitor of the voltage and the current delivered by each channel and the database recording of the settings. The ECAL Safety System can switch off the HV via the individual board interlocks.

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 Regulator Boards are equipped with DCUs that measure the voltages and these measurements are read via the Token Ring. Overall the power is supplied by MARATON crates (WIENER), which are operated and monitored by the DCS.


### 4.8 Detector calibration

Calibration is a severe technical challenge for the operation of the CMS ECAL. Many small effects which are negligible at low precision need to be treated with care as the level of precision of a few per mille is approached. ECAL calibration is naturally seen as composed of a global component, giving the absolute energy scale, and a channel-to-channel relative component, which is referred to as intercalibration. The essential issues are uniformity over the whole ECAL and stability, 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  $\approx 8\%$  within most supermodules, although the total variation among all barrel crystals is  $\approx 15\%$ . In the endcap the VPT signal yield, the product of the gain, quantum efficiency and photocathode area, has an RMS variation of almost 25%. Preliminary estimates of the intercalibration coefficients are obtained from laboratory measurements of crystal light yield and photodetector/electronics response [101]. Applying this information reduces the channel-to-channel variation to less than 5% in the barrel and less than 10% in the endcaps.

All 36 supermodules were commissioned in turn by operating them on a cosmic ray stand for a period of about one week. A muon traversing the full length of a crystal deposits an energy of approximately 250 MeV, permitting intercalibration information to be obtained for the barrel ECAL [102]. In 2006, nine supermodules were intercalibrated with high energy electrons (90 and 120 GeV), in a geometrical configuration that reproduced the incidence of particles during CMS operation. One of the supermodules was exposed to the beam on two occasions, separated by an interval of one month. The resulting sets of inter-calibration coefficients are in close agreement, the distribution of differences having an RMS spread of 0.27%, indicating a reproducibility within the statistical precision of the individual measurements (figure 4.12).

A comparison of the cosmic ray and high energy electron data demonstrates that the precision of the cosmic ray inter-calibration is better than 1.5% over most of the volume of a supermodule, rising to just above 2% at the outer end (corresponding to  $\eta \approx 1.5$ ). The mean value of the precision




**Figure 4.12:** Distribution of differences of inter-calibration coefficients from a supermodule exposed to a high energy electron beam on two occasions, separated by a period of one month. The reproducibility of the intercalibration coefficients ( $\text{RMS}/\sqrt{2}$ ) is measured to be 0.2%.

of the cosmic intercalibration, averaged over all the channels in the nine supermodules for which a comparison with electrons can be made, is 1.5% (figure 4.13).


The ultimate intercalibration precision will be achieved *in situ* with physics events. As a first step, imposing the  $\phi$ -independence of the energy deposited in the calorimeter can be used to rapidly confirm, and possibly improve on, the start-up intercalibration within fixed  $\eta$  regions. The intercalibration method that has been investigated in the most detail uses the momentum of isolated electrons measured in the tracker. These electrons, mainly from  $W \rightarrow e\nu$ , are abundant ( $\sigma \approx 20$  nb) and have a similar  $p_T$  to the photons of the benchmark channel  $H \rightarrow \gamma\gamma$ . A complementary method, not relying on the tracker momentum measurement, is based on  $\pi^0 \rightarrow \gamma\gamma$  and  $\eta \rightarrow \gamma\gamma$  mass reconstruction. Most methods of intercalibration will be local to a region of the ECAL, and a further step intercalibrating these regions to one another will be needed. This is a consequence of the significant systematic variations that occur as a function of pseudorapidity such as (or including): the large variation of the thickness of the tracker material, the variation of the structure of the ECAL (both the major differences between the barrel and endcap, and the small continuous variation of the geometry along the length of the barrel), and the variation of background characteristics for  $\pi^0 \rightarrow \gamma\gamma$ .

Over the period of time in which the physics events used to provide an intercalibration are taken the response must remain stable 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.



**Figure 4.13:** Distribution of the relative differences between the inter-calibration coefficients measured with high energy electrons and those obtained from cosmic ray muons.

The final goal of calibration is to achieve the most accurate energy measurements for electrons and photons. 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 will be important to find ways of using 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, and also assist in the step of intercalibrating regions of the ECAL to one another, are under investigation:  $Z \rightarrow e^+e^-$ , and  $Z \rightarrow \mu^+\mu^-\gamma$  (the photon coming from inner bremsstrahlung).




**Figure 4.14:** 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.

## 4.9 Laser monitor system

Although radiation resistant, ECAL  $\text{PbWO}_4$  crystals show a limited but rapid loss of optical transmission under irradiation due to the production of colour centres which absorb a fraction of the transmitted light. At the ECAL working temperature ( $18^\circ\text{C}$ ) the damage anneals and the balance between damage and annealing results in a dose-rate dependent equilibrium of the optical transmission, if the dose rate is constant. In the varying conditions of LHC running the result is a cyclic transparency behaviour between LHC collision runs and machine refills (figure 4.14). The magnitude of the changes is dose-rate dependent, and is expected to range from 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 performance of the calorimeter would be unacceptably degraded by these radiation induced transparency changes were they not measured and corrected for.

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 photodiodes. PN type photodiodes were chosen because of their very narrow depletion zone ( $\approx 7 \mu\text{m}$  with +4 V reverse bias), making them much less sensitive to type inversion than the faster PIN photodiodes. Thus  $R(t) = APD(t)/PN(t)$  is used as the measure of the crystal transparency. The laser monitoring system [69] 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 changes in crystal transparency cause a change in response to the laser light which is not necessarily equal to the change in response to scintillation light. For attenuations  $< 10\%$  the





**Figure 4.15:** Relation between the transmission losses for scintillation light and for laser light for a given crystal. The signals are followed during the irradiation and the recovery.

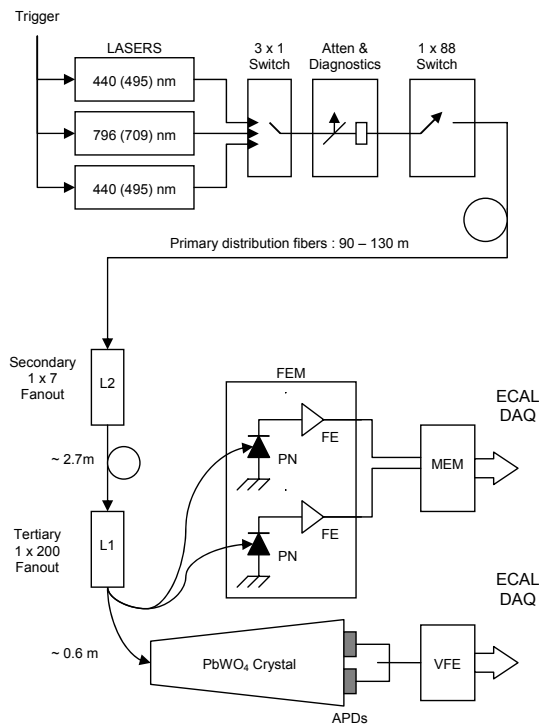
relationship between the changes can be expressed by a power law,

$$\frac{S(t)}{S(t_0)} = \left[ \frac{R(t)}{R(t_0)} \right]^\alpha, \quad (4.1)$$

where  $S(t)$  represents the response to scintillation light and  $\alpha$  is characteristic of the crystal which depends on the production method ( $\alpha \approx 1.53$  for BCTP crystals, and  $\alpha \approx 1.0$  for SIC crystals). An example of this relationship is given in figure 4.15. 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.

#### 4.9.1 Laser-monitoring system overview

Figure 4.16 shows the basic components of the laser-monitoring system: two laser wavelengths are used for the basic source. One, blue, at  $\lambda=440$  nm, is very close to the scintillation emission peak, which is used to follow the changes in transparency due to radiation, and the other, near infra-red, at  $\lambda=796$  nm, is far from the emission peak, and 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  $\approx 30$  ns. The lasers can be pulsed at a rate of  $\approx 80$  Hz, and the pulse timing jitter is less than 3 ns which allows adequate trigger synchronization with the LHC bunch train and ECAL ADC clock.



**Figure 4.16:** The components of the laser monitoring system.


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 near 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 [103].

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.

The first laser system was installed in the CERN H4 test beam site in August 2001. The other two laser systems were installed at H4 in August, 2003. All three laser systems have been used in



**Figure 4.17:** Relative stability between a pair of reference PN photodiodes monitoring 200 crystals measured in autumn 2004 at the CERN test beam facility.

the ECAL test beam program since their installation, and more than 10 000 laser hours have been cumulated.

The relative stability between a pair of reference PN photodiodes monitoring the same group of 200 crystals is shown in figure 4.17. The system achieves 0.0074% RMS over 7.5 days operation.

The response to injected laser light (normalized by the reference PN photodiodes) is presented in figure 4.18 for a group of 200 crystals measured for 11.5 days at the wavelength of 440 nm, showing that a stability of 0.068% is achieved at the scintillation wavelength.


The effect of the monitor correction procedure is presented in figure 4.19, showing that electron signals taken during an irradiation test at H4 are effectively corrected using laser monitor runs taken during the same data-taking period, providing an equalisation of the corrected response at the level of few per mille [104].

## 4.10 Energy resolution


For energies below about 500 GeV, where shower leakage from the rear of the calorimeter starts to become significant, the energy resolution can be parametrized as in equation (1.1) (chapter 1.1), that is repeated for convenience here:

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

where  $S$  is the stochastic term,  $N$  the noise term, and  $C$  the constant term. The individual contributions are discussed below.



**Figure 4.18:** Stability of crystal transmission measurements at 440 nm (blue laser) over 11.5 days operation for a module of 200 crystals.



**Figure 4.19:** Effect of the monitor correction procedure on test beam data: full black points refer to signals measured during test beam irradiation, open red points are the same after the monitor corrections.

### The stochastic term

There are three basic contributions to the stochastic term:

1. event-to-event fluctuations in the lateral shower containment,

2. a photostatistics contribution of 2.1%,
3. fluctuations in the energy deposited in the preshower absorber (where present) with respect to what is measured in the preshower silicon detector.

The contribution to the stochastic term coming from fluctuations in the lateral containment is expected to be about 1.5% when energy is reconstructed by summing an array of  $5 \times 5$  crystals, and about 2% when using  $3 \times 3$  crystals.

The photostatistics contribution is given by:

$$a_{\text{pe}} = \sqrt{\frac{F}{N_{\text{pe}}}} \quad (4.3)$$

where  $N_{\text{pe}}$  is the number of primary photoelectrons released in the photodetector per GeV, and  $F$  is the excess noise factor which parametrizes fluctuations in the gain process. This factor has a value close to 2 for the APDs, and is about 2.5 for the VPTs. A value of  $N_{\text{pe}} \approx 4500$  pe/GeV is found for the barrel, giving  $\approx 2.1\%$  for the photostatistics contribution to the stochastic term. In the endcap the photostatistics contribution is similar, since the larger collection area of the VPT largely compensates for the reduced quantum efficiency of the photocathode.

The contribution to the energy resolution from the preshower device can be approximately parametrized as a stochastic term with a value of  $5\%/\sqrt{E}$ , where  $E$  is in GeV. But, because it samples only the beginning of the shower, the resolution is, in fact, predicted to vary like  $\sigma/E \propto 1/E^{0.75}$ . A beam test in 1999 [105] verified this prediction.

### The constant term

The most important contributions to the constant term may be listed as follows:

1. non-uniformity of the longitudinal light collection,
2. intercalibration errors,
3. leakage of energy from the back of the crystal.

The effects of the longitudinal light collection curve have been studied in detail. Quite stringent requirements are made on the crystal longitudinal uniformity. Requiring the constant term contribution due to non-uniformity be less than 0.3%, sets a limit on the slope of the longitudinal light collection curve in the region of the shower maximum of  $\approx 0.35\%$  per radiation length. A small increase in response towards the rear of the crystal helps to compensate the rear leakage from late developing showers, which would otherwise cause a low energy tail. The required response is achieved in the barrel by depolishing one long face of the crystals to a designated roughness. This surface treatment is incorporated into the crystal production process.

The effect of rear leakage is very small. Charged particles leaking from the back of the crystals can also give a direct signal in the APDs (nuclear counter effect), but test beam data show that this effect is negligible for isolated electromagnetic showers: no tails on the high side of the energy distribution are observed even at the highest electron energy tested (280 GeV).

## The noise term

There are three contributions to the noise term:

1. electronics noise,
2. digitization noise,
3. pileup noise.

The signal amplitude in the test beam is reconstructed using a simple digital filter. The noise measured, after this amplitude reconstruction, for channels in barrel supermodules is  $\approx 40$  MeV/channel in the highest gain range. This noise includes both electronics and digitization noise. The amplitude reconstruction makes use of an event-by-event baseline subtraction using 3 digitization samples taken directly before the signal pulse. This procedure removes the small channel-to-channel correlated noise. Its success is evidenced by the fact that, after this procedure, the noise in the sum of 25 channels is almost exactly 5 times the noise in a single channel [106].

In the endcap it is intended to sort the VPTs in bins of overall signal yield, which includes the photocathode area, the quantum efficiency and the VPT gain. The VPTs with higher overall signal yield are used for the larger radius regions of the endcap. This has the result that the transverse energy equivalent of the noise will be more or less constant, with a value of  $\sigma_{\text{ET}} \approx 50$  MeV.

Neutron irradiation of the APDs in the barrel induces a leakage current which contributes to the electronics noise. The evolution of the leakage current and induced noise over the lifetime of the experiment has been extensively studied. The expected contribution is equivalent to 8 MeV/channel after one year of operation at  $\mathcal{L} = 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ , and 30 MeV/channel at the end of the first year of operation at  $\mathcal{L} = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$  [69].


The shaped signals from the preamplifier output will extend over several LHC bunch crossings. When using a multi-weights method to reconstruct the signal amplitude [106], up to 8 time samples are used. Pileup noise will occur if additional particles reaching the calorimeter cause signals which overlap these samples.

The magnitude of pileup noise expected at low luminosity ( $\mathcal{L} = 2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ ) has been studied using detailed simulation of minimum bias events generated between  $-5$  and  $+3$  bunch crossings before and after the signal. The average number of minimum bias events per bunch crossing was 3.5. Figure 4.20 shows the reconstructed amplitude observed with and without pileup in the absence of any signal. The fraction of events with a signal beyond the Gaussian distribution of the electronics noise is small, showing that at low luminosity the pileup contribution to noise is small.

## Energy resolution in the test beam

In 2004 a fully equipped barrel supermodule was tested in the CERN H4 beam. The energy resolution measured with electron beams having momenta between 20 and 250 GeV/c confirmed the expectations described above [107]. Since the electron shower energy contained in a finite crystal matrix depends on the particle impact position with respect to the matrix boundaries, the intrinsic performance of the calorimeter was studied by using events where the electron was limited to a  $4 \times 4 \text{ mm}^2$  region around the point of maximum containment (*central impact*). Figure 1.3 shows





**Figure 4.20:** Reconstructed amplitude in ECAL barrel channels in the absence of a signal, without pileup (dashed histogram) and with pileup (solid histogram). A Gaussian of width 40 MeV is superimposed on the dashed histogram.


the resolution as a function of energy when the incident electrons were restricted in this way. The energy is reconstructed by summing  $3 \times 3$  crystals. A typical energy resolution was found to be:

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{2.8\%}{\sqrt{E}}\right)^2 + \left(\frac{0.12}{E}\right)^2 + (0.30\%)^2,$$

where  $E$  is in GeV. This result is in agreement with the expected contributions detailed in the earlier part of this section. (Results from beam-test runs taken in 2006, using the final VFE card, show a 10% improvement of the noise performance.)

The energy resolution was also measured with no restriction on the lateral position of the incident electrons except that provided by the  $20 \times 20$  mm<sup>2</sup> trigger. The trigger was roughly centred ( $\pm 3$  mm) on the point of maximum response of a crystals. In this case a shower containment correction was made as a function of incident position, as measured from the distribution of energies in the crystal, to account for the variation of the amount of energy contained in the matrix. For energy reconstruction in either a  $3 \times 3$  or a  $5 \times 5$  matrix an energy resolution of better than 0.45% is found for 120 GeV electrons after correction for containment. Figure 4.21 shows an example of the energy distributions before and after correction for the case of reconstruction in a  $5 \times 5$  matrix, where the correction is smaller than for the  $3 \times 3$  case.

The energy resolution has also been measured for a series of 25 runs where the beam was directed at locations uniformly covering a  $3 \times 3$  array of crystals. In this case a resolution of 0.5% was measured for 120 GeV electrons.



**Figure 4.21:** Distribution of energy reconstructed in a  $5 \times 5$  matrix, before and after correction for containment, when 120 GeV electrons are incident over a  $20 \times 20\text{mm}^2$  area.