EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

CERN–EP/98–20
27 January 1998

BEAM TESTS OF LEAD TUNGSTATE CRYSTAL MATRICES AND A SILICON STRIP PRESHOWER DETECTOR FOR THE CMS ELECTROMAGNETIC CALORIMETER

E. Auffray[5], P. Baillon[5], D. Barney[5], G. Bassompierre[2], Y. Benhammou[17], A. Blick[13], P. Bloch[5], P. Bonamy[6], J. Bourotte[11], L. Buiron[11], F. Cavallari[14], R. Chipaux[6], D. Cockerill[4], I. Dafinei[14], G. Davies[7], P. Depasse[17], K. Deiters[18], M. Diemoz[14], L. Dobrzynski[11], S. Donskov[13], H. El Mamouni[17], C. Ercoli[5], J. L. Faure[6], M. Felcini[19], F. Gautheron[2], M. Geleoc[6], A. Givernaud[6], S. Gninenko[10], N. Godinovic[15] D. J. Graham[7], J. P. Guillaud[2], E. Guschin[10], M. Haguenauer[11], H. Hillemanns[1], H. Hofer[19], B. Ille[17], A. Inyakin[13], S. Jaaskelainen[5], V. Katchanov[13], T. Kirn[1], K. Kloukinas[5], M. Korzhik[9], K. Lassila-Perini[19], P. Lebrun[17], P. Lecoq[5], G. Lecoeur[5], P. Lecomte[19], E. Leonardi[14], E. Locci[6], R. Loos[5], E. Longo[14], C. Mackay[16], E. Martin[7], F. Martin[17], J. P. Mendiburu[2], Y. Musienko[10], P. Nedelec[2], F. Nessi-Tedaldi[19], G. Organtini[14], S. Paoletti[5], J. P. Pansart[6], J. P. Peigneux[2], I. Puljak[15], S. Qian[6], E. Reid[3], D. Renker[18], A. Rosowsky[6], E. Rosso[5], R. Rusack[8], H. Rykaczewski[19], M. Schneegans[2], C. Seez[7], I. Semeniouk[10], P. Shagin[13], D. Sillou[2], A. Singovsky[13], V. Sougonyaev[13], I. Soric[15], P. Verrecchia[6], J. P. Vialle[2], T. S. Virdee[5]†, R. Y. Zhu[13]

Abstract

Tests of lead tungstate crystal matrices carried out in high-energy electron beams in 1996, using new crystals, new APDs and an improved test set-up, confirm that an energy resolution of better than 0.6% at 100 GeV can be obtained when the longitudinal uniformity of the struck crystal is adequate. Light loss measurements under low dose irradiation are reported. It is shown that there is no loss of energy resolution after irradiation and that the calibration change due to light loss can be tracked with a precision monitoring system. Finally, successful tests with a preshower device, equipped with silicon strip detector readout, are described.

(Submitted to Nuclear Instruments and Methods in Physics Research)

1 RWTH 1, Aachen, Germany
2 LAPP, Annecy, France
3 University of Bristol, Bristol, UK
4 Rutherford Appleton Laboratory, Didcot, UK
5 CERN, Geneva, Switzerland
6 Saclay, Gif-sur-Yvette, France
7 Imperial College of Science Technology and Medicine, London, UK
8 University of Minnesota, Minneapolis, USA
9 Institute of Nuclear Problems, Minsk, Belarus
10 INR, Moscow, Russia
11 LPNHE, Ecole Polytechnique, Palaiseau, France

12 California Institute of Technology, Pasadena, USA
13 IHEP, Protvino, Russia
14 University of Rome "La Sapienza", Rome, Italy
15 Technical University of Split, Split, Croatia
16 Brunel University, Uxbridge, UK
17 I. P. N. de Lyon, IN2P3-CNRS et Université C. Bernard, Villeurbanne, France
18 Paul Scherrer Institute, Villigen, Switzerland
19 ETH, Swiss Federal Institute of Technology, Zurich, Switzerland
† Also at Imperial College of Science Technology and Medicine, London

1 Introduction

In previous papers [1, 2] we have shown that an energy resolution of better than 0.6% at 100 GeV can be consistently achieved using an electromagnetic calorimeter (ECAL) of lead-tungstate crystals with avalanche photodiodes (APDs) as photodetectors, in the design configuration chosen for the CMS detector at the LHC [3]. We report here on further tests made in high-energy electron beams in 1996, using new crystals, new APDs, and an improved test set-up. We have also performed tests with a preshower device, equipped with silicon-strip detector readout, placed directly in front of the crystal matrix as is foreseen for the CMS electromagnetic calorimeter (ECAL) endcap at all times, and for the CMS ECAL barrel in the later, high luminosity, phase of LHC operation.

The two major issues for the crystals in the 1996 test-beam studies were longitudinal uniformity and radiation damage. The data presented here confirm that an energy resolution of better than 0.6% at 100 GeV can be obtained when the longitudinal uniformity of the struck crystal is adequate. The problem is to measure the uniformity of the crystal before the beam test and to uniformize it. We are beginning to be able to do this with sufficient precision and reliability. One aspect of the problem is the large spread of characteristics in the crystals used for the test beam, a consequence of the fact that the production parameters are currently in continuous development to improve the crystal properties. This will not be the case in the crystal production phase.

It was recognized in 1995 that radiation damage at LHC, causing loss of detected light, can pose a serious problem for the CMS ECAL even though it saturates after an integrated dose of a few kilorads. The damage anneals with relatively short time constants so that the calibration is subject to continual fluctuations (at room temperature the dominant time constant is of the order of a week). The strategy adopted by the CMS ECAL group for dealing with this problem is simple: crystal development is now being focused on developing crystals less sensitive to damage. Any residual damage must be followed and corrected for using a precision monitoring system. Here we report on light loss measurements comparing different crystals. We show that there is no loss of energy resolution after irradiation, and we demonstrate that the calibration change can be tracked with a precision monitoring system.

In addition to these crystal issues a 3-X_0 preshower device, equipped with silicon-strip detector readout and fast, LHC compatible, electronics was placed in front of the matrix and tests of position and energy resolution were carried out.

2 Test set-up for 1996

The test set-up used in 1996 was similar to the one already described in an earlier publication [2]. Here we give only a general overview and details of improvements that were made.

The crystals used were truncated pyramids with a length of 230 mm (25.8 X_0) and a 20.5×20.5 mm^2 front face projective to a point about 1430 mm from the front surface, and they closely approximate the crystals to be used near $|\eta| = 0$ in the baseline CMS design. They were placed in a light-tight wooden box, internally clad with grounded copper sheeting, on a remotely controlled moving table. The crystal axes were tilted by 3° to the beam direction in both transverse coordinates to correspond to the intended configuration in the CMS ECAL. The tilt prevents particles from passing through the cracks between crystals.

The data used in the analyses described here were taken with 7×7 crystal matrices in two 14-day periods in August and September. The first period was used to study radiation damage and the ability of the monitoring system to track it, and the second period was devoted to energy resolution studies.

2.1 Temperature control

Both the scintillation light yield of the crystal and the APD gain are temperature dependent ($\approx -2\%/^\circ$ for the crystals, and $\approx -4\%/^\circ$ for the APDs), so the temperature of the test set-up had to be carefully stabilized. The four sides of the crystal matrix parallel to the beam were surrounded by water-cooled copper plates. The preamplifier and bias-control circuit boards were also in direct

contact with water-cooled copper bars. In addition to this cooling, which was similar to the previously described set-up, heat screens covered most of the six interior surfaces of the surrounding box. These heat screens consisted of two 0.5-mm-thick aluminium sheets sandwiching a 4-mm-thick sheet of aluminium cut with a zig-zag pattern to form a water channel with a section of about 200 mm^2 through which water was passed. Two independently temperature controlled water baths were used: one for crystal regulation and the other for the electronics and heat screens. For the front screen the 4-mm-thick aluminium sheet was replaced by a 2-mm-thick sheet of light-weight fibre-reinforced epoxy. For the data-taking periods discussed in detail below a temperature of 18˚C was set.

The temperatures of the front and rear surfaces of the crystals were continuously monitored and written to tape with the other data. Fig. 1 shows the temperature measured on the back of a crystal near the centre of the matrix over a period of five days. It can be seen that with the improved temperature stabilization system the daily temperature fluctuation was reduced to about 0.15˚C peak-to-peak (i.e. about 0.05˚C r.m.s. variation) on the back face of the crystals. (Note that the absolute calibration of the sensors, used for the vertical scale, has an uncertainty of about 0.5˚C.) This represents a reduction by more than a factor of three when compared to 1995 when the largest peak-to-peak variation (then on the front of the crystals) was about 0.5˚C [2]. It is assumed that the temperature of the APDs closely followed that of the back face of the crystals, since the APDs were in better thermal contact with the crystals than with anything else.

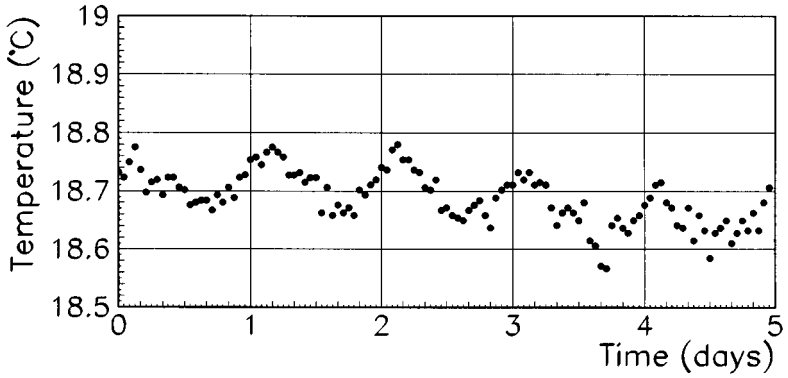

Figure 1: Temperature monitored on the back of a crystal near the centre of the September matrix over a period of five days.

On the front face of the crystals the measured temperature variation was about three times smaller, presumably owing to the absence of the thermal gradients from the preamplifier and line driver, which are present at the back. The temperature variation in the bulk of the crystals is assumed to be similar to (or smaller than) that at the front of the crystals. Thus gain variations caused by temperature changes should be almost entirely due to the sensitivity of the APD gain to temperature, rather than to the temperature-dependent light output of the crystals. The results are entirely consistent with this interpretation.

2.2 Readout and electronics

The APDs used for light detection continue to be improved. The first version of the EG&G APD used was the same as that used in 1995, but the second version and the Hamamatsu device were developments on the previous versions, achieved by modifications to the depth structure of the drift and gain regions [1]. The main improvement in the EG&G device was a reduced excess noise factor for a small penalty in the increased effective thickness of silicon sensitive to direct signals from shower leakage. The main improvement in the Hamamatsu device was the reduction of both the capacitance and the effective thickness at the same time.

The signals from these APDs went into fast, low-noise preamplifiers designed around a high-performance JFET. These amplifiers had a peaking time of 35 ns and a measured electronics noise of about 1200 electrons when connected to the EG&G APD and 2500 electrons when con-

nected to the Hamamatsu device. These figures correspond, using the average light collected per unit energy for the 1996 crystals, to an energy equivalent noise of about 12 MeV per channel for the EG&G APDs and 25 MeV for the Hamamatsu devices. In practice this figure was not achieved. The ground loop and/or pick-up noise present in 1995 and associated with the 80 m cable from the beam area to the control room was eliminated using an isolation interface to the ADCs, but degradation of the electronic shielding around the preamplifiers — an unforeseen consequence of improvements to the cooling system — resulted in significant coherent pick-up noise inside the box. Owing, presumably, to the details of the pin-out and cabling characteristics, the Hamamatsu APDs were particularly prone to this problem. In practice the observed energy equivalent noise on a sum of nine channels was between 120 and 180 MeV for the EG&G APDs with which most of the data were taken.

2.3 Monitoring the calibration

Much attention was paid to light monitoring during the tests. Optical fibres fed light pulses to the front of the crystals. During the year both red (λ = 676 nm) and green (λ = 532 nm) lasers, a pulsed xenon system using a diffraction grating to provide four different wavelengths, and a LED[1] light source (pulse length of 30 ns) were used. The light monitoring signals were viewed by silicon photodiodes, as well as being injected into the crystals. In the case of the LED the photodiode readout was used in a feedback loop which stabilized the magnitude of the light pulse. The other photodiode monitoring signals were digitized and written to tape along with the other data, to be used offline.

Monitoring with various wavelengths provided useful information for the crystal manufacturers, but only the LED system (λ = 660 nm) achieved sufficient stability to allow tracking of the calibration changes due to radiation damage with the required precision. Fig. 2 shows the response to the LED monitoring pulses in a channel far from the beam over a period of about 25 hours. The peak-to-peak variation of 0.4% (0.15% r.m.s.) is consistent with the variation of APD gain due to the daily temperature cycle.

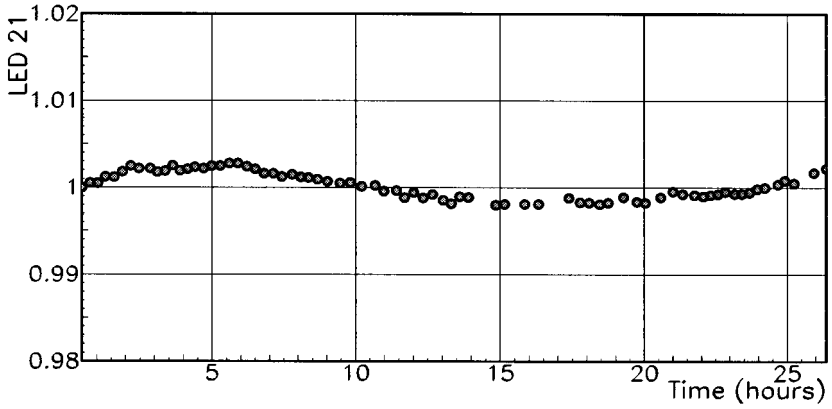

Figure 2: LED monitoring signal seen in a channel far from the beam over a period of 24 hours.

2.4 Crystal preparation

On delivery to CERN crystals were visually inspected and measured geometrically before being given an extensive series of quality checks, including measurements of light yield, optical absorption as a function of wavelength, and longitudinal uniformity.

All crystals were wrapped with Tyvek[2] sheets of ~150 ± 50 μm thickness. On the two end faces, holes were left for the APD and a crystal temperature sensor on the back, and another

1. Stanley Super Bright SBR5501, peak emission at 660 nm.
2. *Tyvek* is a trademark of Dupont. We used type 1056, with a nominal thickness of 160 μm.

temperature sensor and three optical fibres on the front. The longitudinal non-uniformity of light collection of all crystals was first measured in the laboratory using a ^{22}Na source and a two-inch Philips XP2262B photomultiplier coupled to the back of the tapered crystals with optical grease. One of the 511-keV gamma rays is tagged by a plastic scintillator equipped with a photomultiplier, forcing the second gamma to be inside the crystal. The tagging counter with source is displaced along the 23-cm length of the crystal and the signal, gated (200 ns length) by the tag, is digitized in a charge integration ADC.

The light-yield curves obtained for crystals with all six faces optically polished depend on the attenuation length (λ_{att}):

• Some older crystals with poor attenuation length ($\lambda_{att} \approx 30–50$ cm) have curves with continuous decrease from PM to the small end. In some cases, the curves are nearly flat for the first 10–12 X_0, resulting in a 'natural uniformity'.

• In crystals with a longer attenuation length ($\lambda_{att} \approx 50–80$ cm), the focusing effect due to the tapered shape of the crystals results in a U-shaped curve.

• For crystal of good optical quality ($\lambda_{att} > 80$ cm), the focusing effect dominates, giving a light-yield curve rising towards the small end.

Most of the crystals received in 1996 were of the second or third type. Some crystals of the first and second types were uniformized in the 3–10 X_0 region by a local action: stripes of black ink on the inside of the Tyvek wrapping. The position, length and width of the stripes had to be optimized to obtain the desired effect. A significant loss of light is observed with this method (10–20%). For the August and September runs, crystals of all types were uniformized by depolishing the coupling face which reduces the focusing effect by randomizing the escape angles of the photons, or/and depolishing a side face which reduces the total internal reflection. The depolishing was performed on glass with a diamond powder (4–12 μm granularity) in an alcohol suspension. This treatment results in a tilt of the total non-uniformity curve towards larger light response near to the photodetector. If the yields in the shower maximum region are compared before and after the treatment, a small light-yield loss (~7% on average) is observed.

The crystal uniformity was checked using APDs rather than PMs as photodetectors in a low-momentum proton beam (405 MeV, kinetic energy 84 MeV) at the Paul Scherrer Institute (Villigen, Switzerland). Some systematic differences in the curves were observed, presumably due to geometrical and quantum efficiency differences between the photodetectors.

2.5 Preshower detector

The preshower device consisted of a single silicon detector plane[1] measuring 60 mm × 60 mm with 32 strips, placed downstream of a lead absorber, and directly in front of the crystal matrix. Two absorber thicknesses were studied: 2.5 and 3.0 X_0. The distance between the absorber and the silicon detector was approximately 5 mm, whilst the distance between the silicon and the crystals was about 80 mm. The mother board holding the silicon-strip detector and its associated electronics was mounted on a large printed circuit board and attached to the box containing the crystals.

The preamplifier and analogue pipeline used are designed for operation at the LHC. The charge collected in the detector is integrated over 25 ns time samples. The intrinsically fast response of the silicon detectors and a fast, DC-coupled preamplifier (FCICON 18 ns rise time), ensure that the total charge is collected within two time samples. The preamplifier is followed by an analogue memory which allows retrieval of the charges for trigger latency up to 3.2 μs (128 time slots). The preamplifiers and memory form the front-end analogue chip, DYN_{LDR}. Each DYN_{LDR} chip has 32 channels, and receives input from the strips on a single silicon detector unit. A schematic view of the DYN_{LDR} chip is shown in Fig. 3.

For each measurement, four 25 ns samples (three for the full charge collection and one for

1. Produced at ELMA (Zelenograd, near Moscow, Russia)

baseline subtraction) were transmitted via a 2.5 MHz multiplexer to an ADC (Pentek with 12 bit dynamic range).

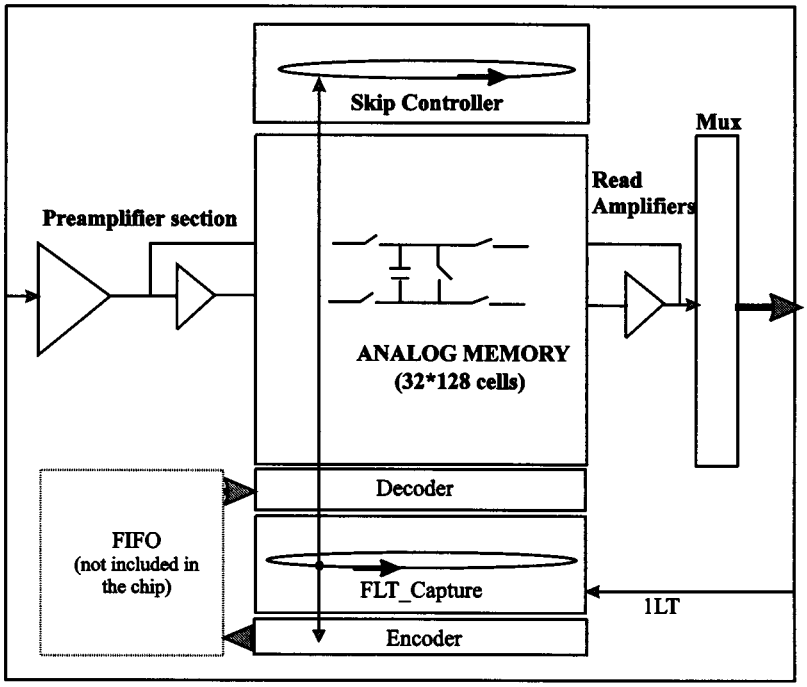

Figure 3: Schematic diagram of the DYN_{LDR} chip.

3 Parametrization of the energy resolution

Data-taking for energy resolution studies was generally preceded by calibration. Using a 50 GeV electron beam, runs of 40k events were taken with a 20×20 mm^2 trigger counter aligned with each of the towers in turn. This was then followed by an energy scan — runs of 100k events were taken with the beam momentum set at each of six momenta in turn (15, 35, 50, 80, 120 and 150 GeV) using the same trigger counter aligned on each of the towers to be studied, in turn.

The energy reconstruction algorithm used consists of summing the signals from 3×3 crystals centred on the struck crystal. For most of the results described below only events where the electron is incident in a 4×4 mm^2 region in the centre of the 3×3 array are used. Since 80% of the energy is in the central crystal the contributions to the energy resolution from photostatistics and longitudinal non-uniformity are completely dominated by its properties, and the measured energy resolution thus characterizes it. More precisely the GEANT [4] Monte Carlo simulation predicts that 80.7% of the energy found in a sum of nine crystals will be in the central crystal when the electron impact point is restricted to a 4×4 mm^2 region in the centre of the crystal. This represents 75.0% of the total incident energy since 93.0% of the incident energy is contained in the 3×3 array. This result is for 120 GeV electrons but is, to very good approximation, energy independent.

The noise term in the energy resolution was calculated run by run using events taken with random triggers within the beam spill. The width of the noise distribution for the sum of nine towers was determined by a Gaussian fit to a region $\pm 1.5\sigma$ about the peak. The noise was subtracted quadratically from the energy resolutions before fitting to obtain the constant and stochastic terms.

Analysis of an electron run is begun by locating, in terms of the drift chamber coordinates, the centre of the crystal covered by the trigger. This is done by finding the point of maximum signal in that crystal. This procedure is reproducible with a precision of better than 0.1 mm. It should be noted that because of the 3° tilt the point found is a notional or effective 'centre' rather than the geometrical centre of the crystal. A calibration is made with the 50 GeV runs.

The data taken in September have been analysed to extract an energy resolution function of the form:

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{a}{\sqrt{E}}\right)^2 + (c)^2$$

after quadratic subtraction of the measured noise (generally between 120 and 180 MeV for a sum of 9 crystals). Fig. 4 shows fits to the energy resolution as a function of energy for beam incident on two different crystals in the September matrix. A linear fit is made possible by plotting the squared variables of the above equation. Each of the central 25 crystals of the 7×7 matrix can be the centre of a sum of nine crystals, and data were taken in all of them. There are thus 25 values for the parameters a and c.

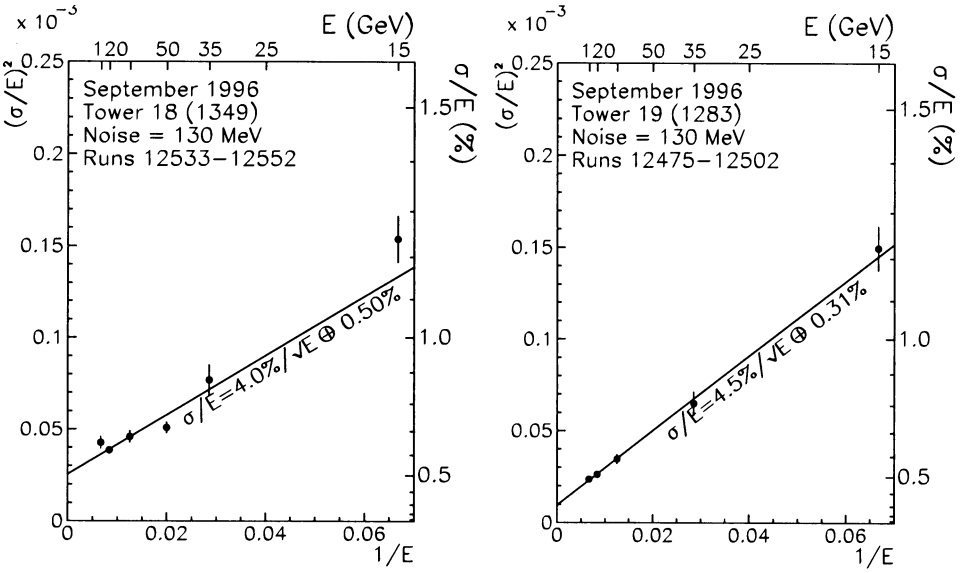

Figure 4: Fits to the energy resolution (noise subtracted) as a function of energy for beam incident in two adjacent crystals in the September matrix.

3.1 The stochastic term

The stochastic term of the energy resolution is composed of a contribution from shower containment and a contribution from photostatistics. The containment contribution, predicted by shower Monte Carlo simulation for a sum of nine crystals, is 2%/√E. The photostatistics contribution is given by

$$a_{pe} = \sqrt{\frac{F}{N_{pe}}},$$

where N_{pe} is the number of primary photoelectrons released in the APD per GeV, and F is the APD excess noise factor. In September EG&G APDs were used. The majority of these APDs were of the older type that has an excess noise factor of 2.8 at a gain of 50. The newer type APDs have an excess noise factor of 2.2 at the same gain. With N_{pe} ranging between about 1000 and 2500 photoelectrons per GeV, we expect and observe contributions to the stochastic term between 3.5% and 5.5%, giving a total stochastic term of between 4% and 6%. There is a tail stretching to higher values with three values beyond 7%. This is largely due to a batch of APDs for which the voltage to give a gain of 50 was incorrectly measured before the beam test, with the consequence that the gain was set too high. The excess noise factor rises steeply with gain for EG&G APDs of the old type, reaching 4.2 at a gain of 100.

It should be mentioned that it is planned to increase the APD area in the final CMS ECAL by a factor of about two. This would increase N_{pe} by almost the same factor and reduce the contribution of photostatistics to the stochastic term.

The photostatistics contribution can be measured using the LED signal width:

$$a_{led} = \sqrt{\frac{\sigma^2_{led} - \sigma^2 noise}{\mu_{led}}}$$

(where σ_{led} and μ_{led} are the fitted Gaussian width and mean of the LED signal, respectively, and σ_{noise}, which is in fact negligible, is the fitted noise width, all in GeV). The values of the photostatistics contribution measured by the LED can be added quadratically to the 2% predicted for shower containment and compared with the values of the stochastic term obtained from fitting the energy resolution. Fig. 5 shows the fractional difference between the predicted and fitted stochastic terms for the 25 cases. The error on the fractional difference is about 10%, giving confidence that the main error in the resolution fitting, the partition of the resolution between a and c, is well under control.

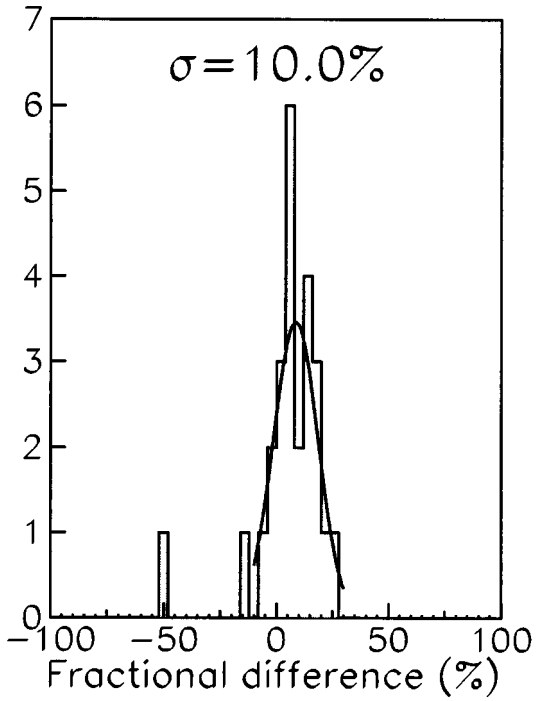

Figure 5: The fractional difference between the stochastic term predicted by the LED peak width measurement and the fitted stochastic term.

3.2 The constant term

The remaining term in the energy resolution is the constant term. The values of the constant term given by the energy resolution fit have a minimum value of about 0.3%. The value predicted by the shower Monte Carlo simulation about 0.15% for a crystal with uniform longitudinal response. The difference is made up from beam σ/p (calculated to be 0.17% with the collimator settings we used for p = 150 GeV, and 0.15% for the other momenta), the effect of shower leakage into the APD, and residual miscalibration effects.

Crystals delivered in 1995 generally had a longitudinal response curve with a parabolic shape. For many of the crystals the flat region at the minimum of the parabola was near the shower maximum. Such a response curve allows a good energy resolution to be achieved. For further details see Ref. [2].

Crystals delivered in 1996 were of higher quality and tended to have a longer attenuation length than those available previously. In these crystals the focusing effect, due to the projective geometry of the crystal, tends to dominate, resulting in longitudinal response curves that decrease monotonically from front to back by 15% or so. Such response curves give contributions of 0.7% or more to the energy resolution [5]. In order to obtain good energy resolution the response curves must be modified. The crystals were uniformized as has been described in an earlier section.

The measured longitudinal response curves have been put into the shower Monte Carlo simulation to predict the expected contributions from longitudinal non-uniformity. The correlation between the observed constant terms and those predicted (i.e. the predicted contribution from longitudinal uniformity plus 0.3%) (Fig. 6) confirms our assumption that longitudinal non-uniformity is the dominant contribution to the observed constant term. The energy resolution at 100 GeV for those crystals where the predicted contribution from longitudinal non-uniformity is less than 0.3%, is shown in Fig. 7.

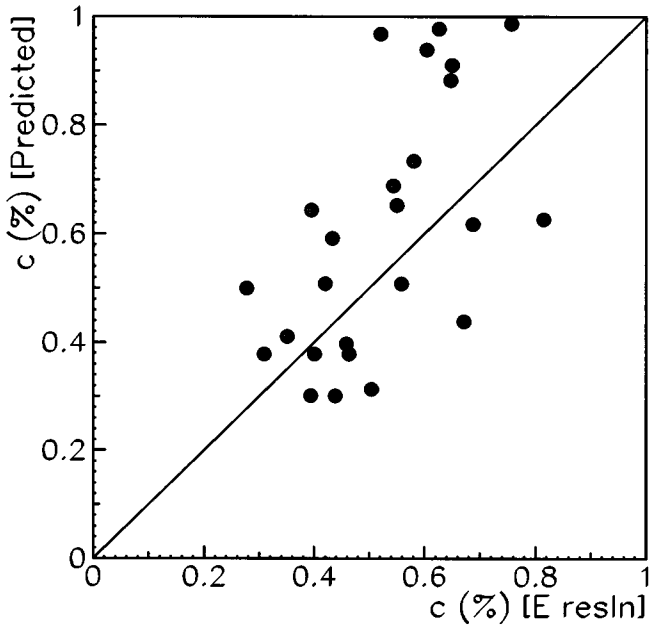

Figure 6: The constant term c, predicted by Monte Carlo simulation using the measured longitudinal response plotted against the fitted constant term. The line of unity gradient is drawn to guide the eye.

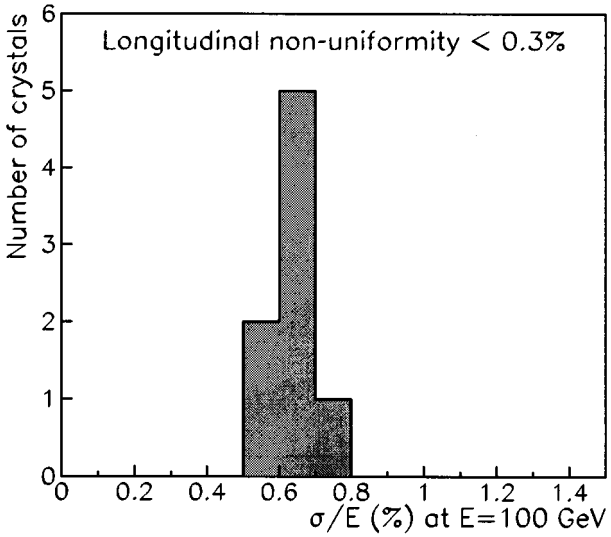

Figure 7: Energy resolution at 100 GeV for those crystals with a predicted contribution from longitudinal non-uniformity of less than 0.3%

It is clear that work remains to be done to improve our ability both to measure precisely, and to correct the longitudinal response. This work will be greatly simplified by the production of crystals with more consistent characteristics, which is to be expected when the present crystal R&D phase, involving the continual deliberate variation of production parameters, draws to a close.

4 Crystal behaviour under irradiation

The beam intensity, at an energy of 120 GeV, can be increased so as to give about 17k electrons per 2.4 s burst (14.4 s cycle time) into a 20×20 mm^2 trigger covering the front face of a crystal. This corresponds to a dose, at shower maximum, of about 25 rad/h, similar to the barrel dose rate for the CMS ECAL at an LHC luminosity of 10^{34} cm^{-2} s^{-1}. We studied the signal losses, with irradiations lasting at least 8 h, of 10 crystals. This provided useful information for crystal producers. Three crystals were studied for longer periods. Table 1 summarizes the results. Some progress in the development of more radiation-hard crystals is evident. Irradiation studies using radioactive sources suggest that the observed signal loss is due to the formation of colour centres, which shorten the attenuation length for the observable scintillation light, and that the intrinsic scintillation mechanism is not damaged [6]. This conclusion is also supported by our ability to follow the light loss using a monitoring system that measures optical transmission (see below). Such behaviour is also familiar in other crystals used for calorimetry in high energy physics [7].

Crystal ID	Producer	Total dose (rad)	Signal loss
1266	BTCP	850	20%
1283	SIC	900	9%
		650	8%
1318	BTCP	500	4%

Table 1: Crystals subjected to prolonged irradiation in the beam

Crystal 1283 was irradiated twice. Between the two irradiations it was annealed by heating it and maintaining it at 200˚C for 12 h. This restores the light yield observed before irradiation. Fig. 8 shows the beam signal as a function of dose (at shower maximum) for the two irradiations. The loss follows the same curve during both irradiations.

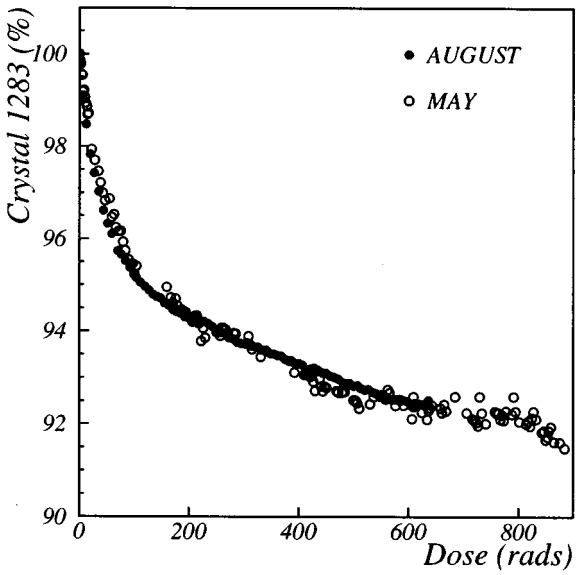

Figure 8: Beam signal as a function of dose for the two irradiations of crystal 1283. The signal has been normalized to 100% at the start of each irradiation.

4.1 Energy resolution after irradiation

Since the radiation damage causes the formation of colour centres, which decrease the light attenuation length, and does not damage the intrinsic scintillation mechanism, it is expected that the longitudinal uniformity of the crystal will not be significantly modified by the damage provided that the initial light attenuation length is sufficiently large, and the loss is not too large. Thus it is expected that the energy resolution will not be significantly degraded by the radiation damage. This has been verified. Fig. 9 shows the energy distribution in a sum of nine crystals for 120 GeV electrons incident in a 4×4 mm^2 area centred on the central crystal (1283) observed before and after irradiation with 650 rad. Appropriate, and different, calibration constants are used for the two runs. The resolution is excellent in both cases, and unchanged within the measurement errors.

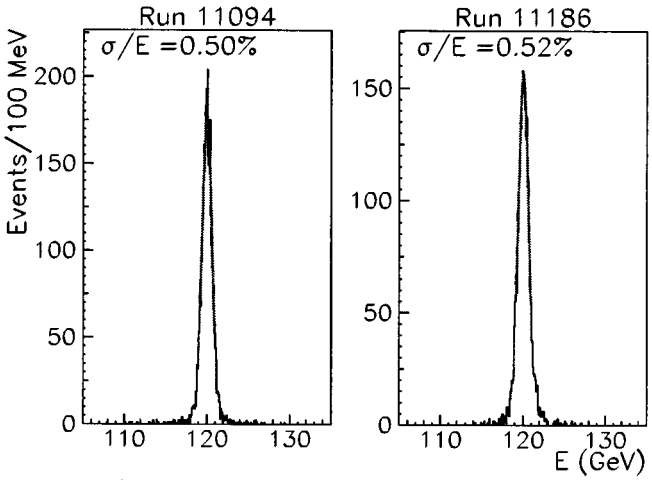

Figure 9: Energy distribution seen in a sum of nine crystals for 120 GeV electrons incident in a 4×4 mm^2 area centred on the central crystal (1283) before and after irradiation to 650 rad.

During irradiation, when the beam rate is raised to about 17k events per burst the measured width of the energy distribution is increased by the addition of a pile-up contribution. This contribution has been measured very precisely by averaging results over many runs and amounts to 0.34% added quadratically. Random triggers show a corresponding broadening of the pedestal which also suffers a significant negative shift. These effects are believed to be due to the ADC or the ground-loop isolation interface used to couple to the ADC and do not reflect what is expected with LHC compatible electronics. It should be noted that in these studies the energy flow during the burst corresponds to that which would be expected at an instantaneous LHC luminosity in excess of 5×10^{34} cm^{-2} s^{-1}.

4.2 Tracking the changing calibration

In order to use the light monitoring system to track the calibration change due to radiation damage, it is necessary to determine the constant of proportionality between the change observed by the monitoring system and the change in the beam signal. Because of the different optical paths taken by the injected monitoring light as compared to the scintillation light this constant cannot be expected to be 1. Furthermore the LED system is monitoring the transparency of the crystal at a wavelength far from the scintillation peak. The relation between the change in transparency seen by the LED light and the change seen by the scintillation light varies from crystal to crystal. It is for this reason that much effort is being devoted to the development of a precision light monitoring system at a wavelength close to the scintillation peak ($\lambda \approx 500$ nm).

The constant of proportionality, Δbeam/ΔLED has been determined for the nine crystals centred on crystal 1283 during the irradiation detailed above. This was achieved for each of the nine crystals on a run-by-run basis (each run lasted approximately 20 minutes) by comparing the

average LED to the average beam signal in the crystal when the incident electron was incident in a tightly restricted region (typically 2×2 mm^2). Restricting the point of incidence in this way gives a reasonably Gaussian distribution of energy even in the eight crystals surrounding the struck crystal, so that the changes in the beam signal can be observed with sufficient precision.

Fig. 10 shows the correspondence between the beam signal and the LED signal for the cen tral crystal. Each point is derived from a single run. The starting value for the LED has been normalized to 1, and the beam values have been normalized so that the straight-line fit passes through the point (1,1). It can be seen that the straight line describes the relationship quite well — the inset histogram shows the residual difference from a straight line as a percentage. The residual difference has an r.m.s. of 0.18%. This residual difference results from temperature variation, as can be seen in Fig. 11 where the difference is plotted as a function of time and compared to the temperature difference. The variation is dominated by the gain variation of the APD with temperature, and implies a gain variation with temperature, dM/dT ≈ 4%/°C.

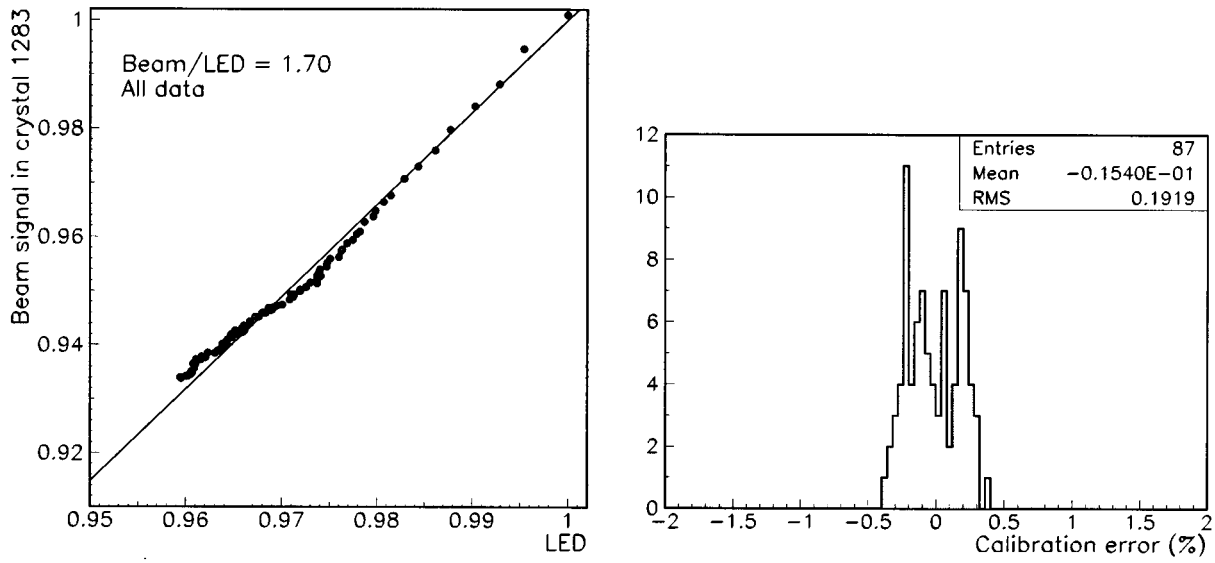

Figure 10 : (a) Correspondence between monitoring LED signals and beam signals during irradiation to 650 rad. (b) histogram of the vertical distances of the points in (a) from the straight-line fit.

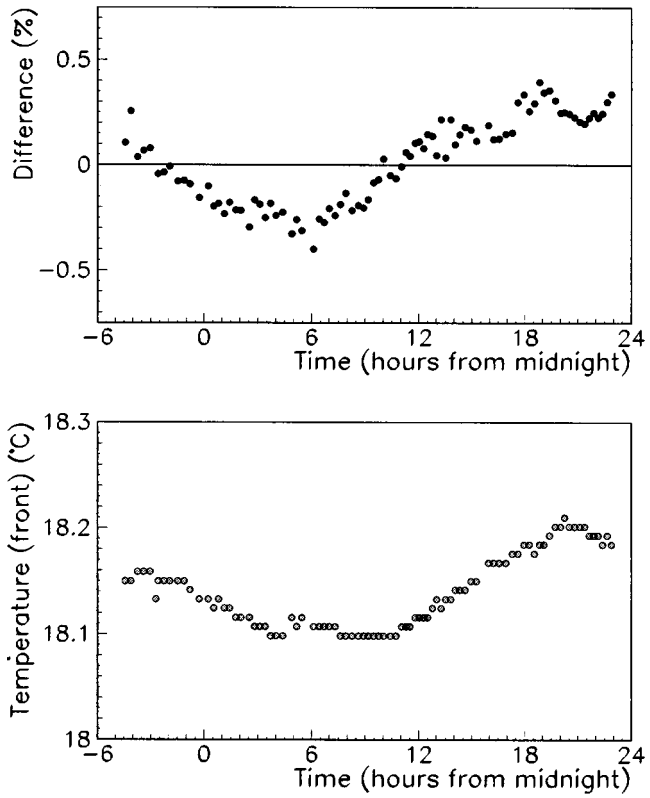

Figure 11: (a) Residual difference between straight-line prediction (Δbeam/ΔLED = 1.70) and observed beam signal during irradiation to 650 rad; (b) temperature difference observed at the back of the crystal matrix during the irradiation.

Table 2 shows the loss of beam signal for the nine crystals, and the constants of proportionality, Δbeam/ΔLED, that were extracted. It can be seen that the losses for some of the surrounding crystals are quite large, despite the fact that the dose in the four crystals directly adjacent to the central crystal is about one half of that in the central crystal. The corner crystals receive about one third of the dose in the central crystal.

–14% 1.45	–9% 1.70	–3% 2.20
–27% 1.15	–8% 1.70	–6% 1.75
–6% 1.60	–3% 1.80	–4% 2.50

Table 2: Loss of beam signal and the constants of proportionality, Δbeam/ΔLED, for the nine crystals centred on crystal 1283

Using these values the entire set of data taken during the irradiation was reanalysed, correcting for the changing calibration on a run-by-run basis. Fig. 12 shows the corrected energy distribution for all events with an electron incident in the 15×15 mm^2 region centred on crystal 1283. When the pile-up contribution of 0.34% has been subtracted, the energy resolution is 0.55%. This compares to 0.53% for the 15×15 mm^2 resolution of single runs before the irradiation (or 0.63%

with pile-up during irradiation). Thus less than 0.2% has been added to the resolution by the use of the monitoring, and this 0.2% seems almost entirely due to the temperature-variation effect previously discussed. A similar result is obtained for electrons incident in a restricted 4×4 mm^2 area, but the 15×15 mm^2 result is shown here since it is much more demanding on the intercalibration precision.

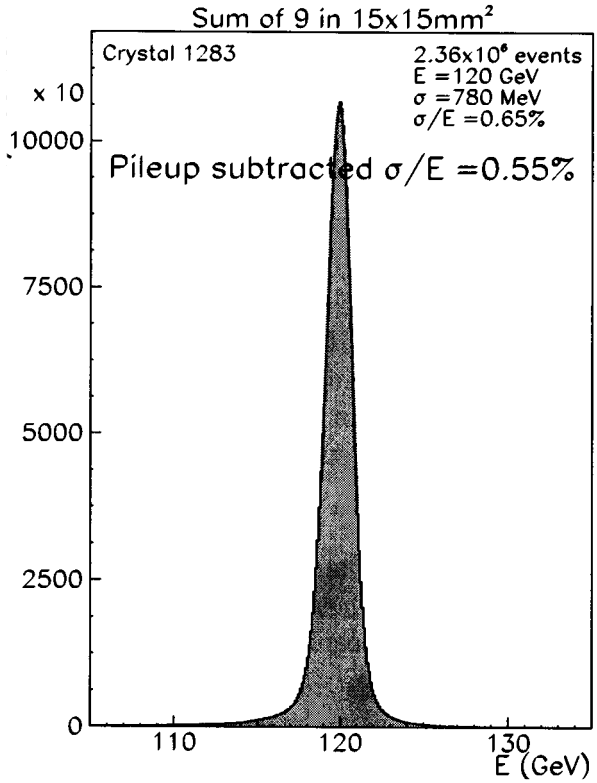

Figure 12: Energy distribution in a sum of nine crystals obtained using LED monitoring data to follow the calibration, for all electrons incident in a 15×15 mm^2 area centred on crystal 1283 taken throughout the irradiation to 650 rad.

5 Test with preshower detector

Tests of the bare silicon detector and associated electronics were made with 50 GeV electrons in the X3 beam at the CERN SPS prior to the tests with the crystals. Fig. 13 shows histograms of the signal and the noise. The signal-to-noise ratio (S/N) is about 3.2, which corresponds to S/N $\approx$ 2.9 for a true minimum-ionizing particle (mip). However, because of a large cell-to-cell pedestal variation in the analogue memory (about 6 mV r.m.s.), the histograms in Fig. 14 were obtained by subtracting individual pedestal values for each memory cell. Our goal is to achieve S/N $\geq$ 5 for mips, without using individual memory-cell pedestal values. We are optimistic that this can be achieved – a similar type of memory designed for the ATLAS SCT in 1996 (using the radiation-hard DMILL process) has given much better results (r.m.s. pedestal uniformity around 1 mV), and we will take advantage of this work. It should also be possible to increase the gain of the preamplifier slightly, which will improve the signal-to-noise ratio.

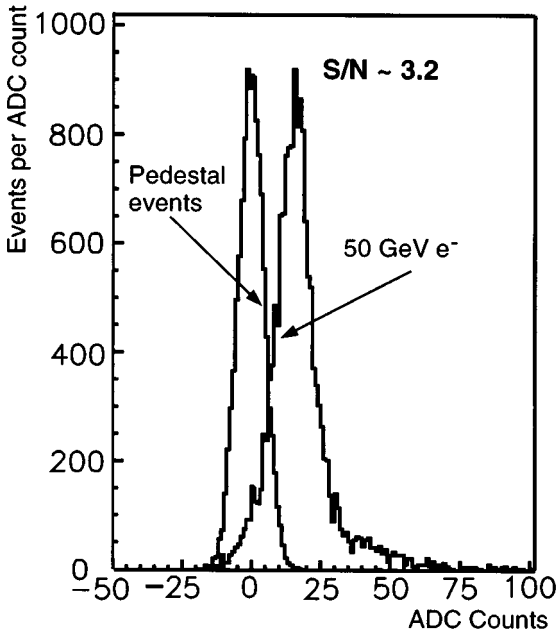

Figure 13: The 50 GeV electron signal and the noise (empty triggers) seen in silicon strips.

5.1 Energy resolution with preshower detector present

Data were taken in the H4 beam with high-energy electrons. The signal in the silicon strips was spread over about three time slots because the system is operating asynchronously with the beam trigger, unlike the intended operation in LHC. The signals seen in the preshower silicon strips were used to measure the energy deposited in the lead absorber. The reconstructed energy, E_{rec}, is obtained by summing the energy deposited in the crystals together with the energy deposited in the preshower weighted by a calibration factor, α:

$$E_{rec} = E_{nxn} + \alpha E_{pre} ,$$

where E_{nxn} is the energy reconstructed in an $n \times n$ array of crystals (both 3×3 and 5×5 were tested), centred on the crystal with the highest energy deposit, and αE_{pre} is the weighted sum of signals in five preshower strips (highest + four nearest neighbours) in three time slots.

The parameter α is found to be largely independent of incident electron energy, and only an energy-independent value was used in the analysis: 16.8 MeV/mip for a 3×3 crystal sum, and 16.0 MeV/mip for a 5×5 crystal sum.

The presence of the preshower results in an additional contribution to the energy resolution. This contribution can be extracted by quadratic subtraction of the resolution obtained with crystals, when the preshower was not present, from the resolution obtained when the preshower was present. This preshower contribution, for an absorber thickness of 2.5 X_0, is plotted as a function of beam energy in Fig. 14. The resolution with 3 X_0 of lead absorber is poorer than with 2.5 X_0. The CMS barrel preshower design has a 2.5 X_0 absorber and one silicon plane, whilst the design for the endcap has a 3.0 X_0 absorber and two silicon planes (at depths of 2 and 3 X_0). Detailed simulation shows that the energy resolution with 3 X_0 and two silicon planes in the CMS endcap configuration is similar to that with 2.5 X_0 and one silicon plane. As can be seen in Fig. 14 the preshower gives a negligible contribution to the energy resolution for energies above 50 GeV. Corresponding results from detailed simulations are also shown in Fig. 14. There is good agreement between data and simulation.

5.2 Position resolution of preshower detector

A centre-of-gravity method was used to calculate the position of the electron in the preshower. This measurement is compared to the position as given by beam chambers. Plotting the difference between these two measurements as a function of the preshower position measurement yields the characteristic S-curve (shown for the case of position measurement in the crystals in Ref. [1]). This curve is used to obtain a correction to the centre-of-gravity estimate. Fig. 15 shows the spatial precision of the preshower position measurement, as a function of energy, for a 2.5 X_0 absorber. The solid points are results from the test beam, whilst the open points were obtained from a detailed simulation. There is a good agreement between data and simulation, except for the low-energy points where the test-beam results are worse than expected. The dashed line in Fig. 15 shows the fitted position resolution:

$$\sigma = 1325/\sqrt{E} \oplus 260\ \mu m$$

(where E is in GeV and $\oplus$ denotes quadratic summation).

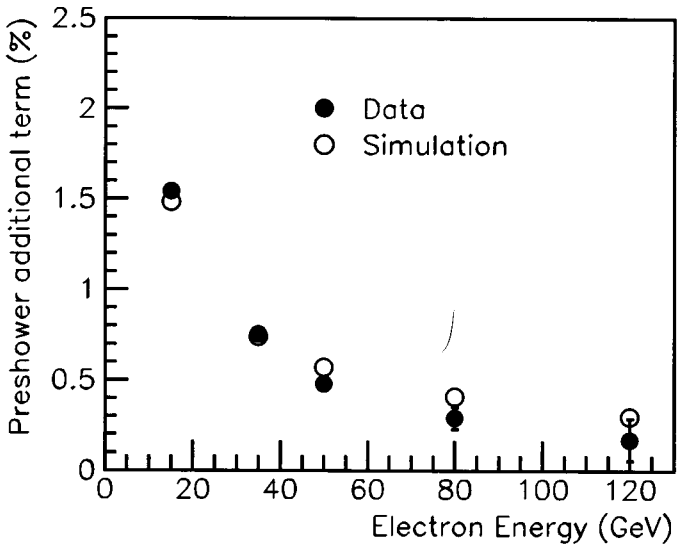

Figure 14: The preshower contribution to the energy resolution as a function of energy. Simulation results are also shown.

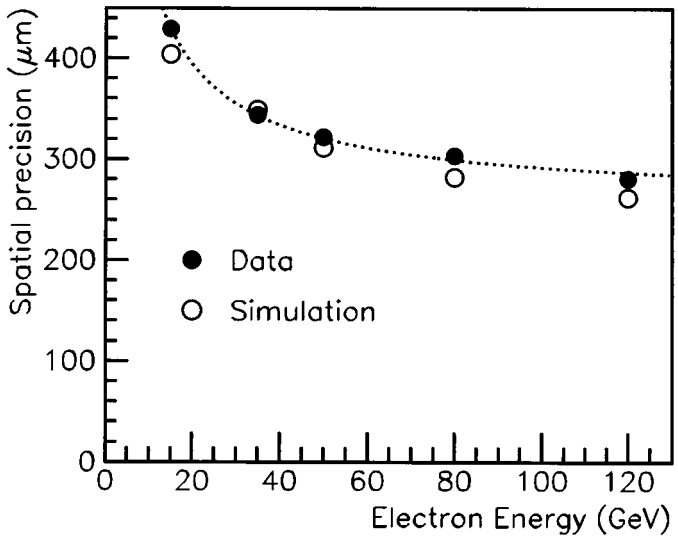

Figure 15: The spatial precision of the preshower position measurement, as a function of energy, for a 2.5 X_0 absorber. The solid points are results from the test beam, whilst the open points were obtained from a detailed simulation. The line is a fit to the position resolution: $\sigma = 1325/\sqrt{E} \oplus 260$ μm (see text).

6 Summary

We obtain excellent energy resolution with lead-tungstate crystals when the longitudinal uniformity of the struck crystal is adequate. Techniques for achieving longitudinal uniformity are now better understood, as is the relation between the longitudinal response curves measured by a photomultiplier tube when the crystal is excited by a ^{22}Na source in the laboratory and the curves measured with APDs in a beam. Uniformization of the crystals for future beam tests will take advantage of this understanding.

We have shown that there is no irrecoverable loss of energy resolution after radiation damage which induces a light output decrease of 8%, and we have demonstrated that the calibration change caused by the loss can be tracked with a precision monitoring system. Future work will aim at perfecting a laser system suitable for the full-size CMS calorimeter, working at a wavelength close to the scintillation peak.

A precision comparison of the shower lateral width with that predicted by GEANT shows that the simulated showers are wider, by about 6%, than those observed.

A preshower device, equipped with silicon-strip detector readout and fast, LHC compatible, electronics has been tested in front of the matrix. The preshower position resolution and the energy resolution of the combined system are largely as predicted by simulation.

Acknowledgments

We would like to thank all the engineers and technicians from the collaborating laboratories who have helped with the work described in this paper. We are also grateful for the cooperation of other members of the CMS collaboration, and in particular those of the CMS ECAL group.

References

[1] J. P. Peigneux *et al.*, Results from tests on matrices of lead tungstate crystals using high energy beams, Nucl. Instr. and Meth. A 378 (1996) 410.

[2] G. Alexeev *et al.*, Studies of lead tungstate crystal matrices in high energy beams for the CMS electromagnetic calorimeter at the LHC, Nucl. Instr. and Meth. A385 (1997) 425.

[3] CMS, the Compact Muon Solenoid, Technical Proposal, CERN/LHCC 94–38, LHCC/P1, 15 December 1994.

[4] GEANT3, version 3.21/07 (release 191296), Detector Description and Simulation Tool, CERN Program Library Long Writeup Q123.

[5] D. Graham and C. Seez, CMS Note 96/002 (1996).

[6] Ren-Yuan Zhu *et al.*, Lead tungstate radiation damage and cure, VI Int. Conf. on Calorimetry in HEP, Frascati, 8–14 June 1996, INFN (1996) 577.

A. N. Annenkov *et al.*, Systematic study of the short-term instability of $PbWO_4$ scintillator parameters under irradiation, CMS Note 97/055 (1997), submitted to Radiation Measurements.

[7] Ren-Yuan Zhu, On quality requirements to the barium fluoride crystals, Nucl. Instr. and Meth. A340 (1994) 442.

K. Kazui *et al.*, Study of the radiation hardness of CsI(Tl) crystals for the BELLE detector, KEK-PREPRINT-96-164 (1997), submitted to Nucl. Instr. and Meth.