The construction and testing of the silicon position sensitive modules for the LHCf experiment at CERN

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2010 JINST 5 P01012

(<http://iopscience.iop.org/1748-0221/5/01/P01012>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 137.138.124.142

The article was downloaded on 12/07/2011 at 08:34

Please note that [terms and conditions apply](#).

RECEIVED: September 30, 2009

REVISED: December 9, 2009

ACCEPTED: December 24, 2009

PUBLISHED: January 26, 2010

The construction and testing of the silicon position sensitive modules for the LHCf experiment at CERN

**O. Adriani,^{a,b} L. Bonechi,^{a,b} M. Bongi,^b G. Castellini,^{c,b} R. Ciaranfi,^b
R. D'Alessandro,^{a,b,1} M. Grandi,^b P. Papini,^b S. Ricciarini,^b A. Tricomi^{d,e}
and A. Viciani^b**

^a*Dipartimento di Fisica, Università degli Studi di Firenze,
Via G. Sansone 1, 50019 Sesto Fiorentino, Italy*

^b*INFN Sezione di Firenze,
Via G. Sansone 1, 50019 Sesto Fiorentino, Italy*

^c*IFAC-CNR,
Via Madonna del Piano 10, 50019 Sesto Fiorentino, Italy*

^d*Dipartimento di Fisica, Università degli Studi di Catania,
Viale Andrea Doria, 95125 Catania, Italy*

^e*INFN Sezione di Catania,
Via Santa Sofia 64, 95123 Catania, Italy*

E-mail: candi@fi.infn.it

ABSTRACT: In this paper the design criteria, construction and final performance of the silicon micro-strip modules installed in the LHCf experiment are described. LHCf is an experiment currently placed at CERN in the LHC tunnel. It consists of two small calorimeters each one placed 140 metres away from the ATLAS interaction point. Their purpose is to study very forward production of neutral particles in proton-proton collisions. The silicon modules are installed in one of the two calorimeters and provide precision information on the shower transverse profile.

KEYWORDS: Si microstrip and pad detectors; Calorimeters; Front-end electronics for detector readout; Data acquisition circuits

¹Corresponding author.

Contents

1	Introduction	1
2	Silicon detector modules	2
2.1	Silicon sensors	3
2.2	FE electronics: the PACE3 chip	5
2.3	FE hybrid	8
2.4	Module assembly	13
3	Front end readout and control electronics	17
3.1	Front end control	18
3.2	Readout and Data transmission	20
3.3	Low voltage powering, and biasing of the detectors	25
3.4	DAQ and slow control software	26
4	Measured performance	27
4.1	Laboratory set-up	28
4.2	Test beams	30
5	Conclusions	36

1 Introduction

LHCf [1, 2] will make high precision calorimetric measurements of photon and neutron showers taking data at the LHC collider at a centre of mass energy of up to 14 TeV in the very forward region. These data will be used to calibrate Monte Carlo codes that describe the development of showers caused by a cosmic ray hitting the Earth’s upper atmosphere, and are currently used by astroparticle physicists.

To accomplish these precision measurements, the LHCf calorimeters rely not only on a high quality calorimetric section (Scintillator tiles interleaved with tungsten absorber) but also on the use of precision position detectors which sample the shower transverse profile at different radiation lengths. One of the calorimeters (ARM1) accomplishes this task by using tightly packed 1mm thick scintillating fibres modules, the other (ARM2, see figure 1) uses instead silicon micro-strip detectors developed for tracking purposes. Only the latter system will be described in this paper while detailed descriptions of the calorimeters and other information on the scintillating fibres modules can be found in [1, 2].

During development many tests were made both in the laboratory and at test beam facilities (CERN). Four complete modules were built (with X and Y readout) and placed inside ARM2 for a total of 3072 readout channels. The modules are optically decoupled for what concerns data and

Figure 1. Photo of the ARM2 calorimeter. The outside dimensions are 90mm width, 300mm length and 600mm height. The light from the scintillator plates is guided through fibres to the photomultipliers (calorimetric part). The black rectangular boxes house the silicon sensors.

control signals, while the only remaining electrical connections to the outside world are the power supply lines. This configuration helps minimising interference with the calorimetric part of ARM2, and also helps implementing a robust grounding and shielding of the silicon detectors themselves. The LHCf apparatus has completed final installation and is now ready to take data at the start of LHC in 2009.

2 Silicon detector modules

Both LHCf calorimeters are capable of measuring photon energies up to a few TeV. The calorimeters are also capable of measuring accurately the shower position so as to reconstruct precisely the aperture angle of the two photons from π^0 decays, thus allowing a very precise calibration of the absolute energy scale. In figure 2 a schematic view of one of the calorimeters (ARM2) is shown. The calorimeter consists of a sandwich of tungsten plates interleaved with scintillator tiles, and the shower position measurement is performed by four pairs of silicon micro-strip detectors placed at

Figure 2. A schematic drawing of the LHCf ARM2 calorimeter. The two towers have transverse dimensions of $25 \times 25 \text{ mm}^2$ and $32 \times 32 \text{ mm}^2$.

6, 12, 30 and 42 X_0 from the front plate. The last two positions are used to help in distinguishing between electromagnetic and hadronic showers.

The peculiar geometry consisting of two separate mini towers, is dictated by the need of having minimum shower leakage from one tower to the adjacent one, and to allow an accurate measurement of the π^0 mass. The tower sizes are comparable to the Moliere radius of the electromagnetic showers, thus there will be a certain amount of side leakage. This is corrected for by applying the position measurements of the silicon layers. In fact the position sensitive detectors are also used to select very clean sample of events where only one photon has showered in one of the towers.

The overall dimensions of the silicon detector modules are $90\text{mm} \times 200\text{mm} \times 240\text{mm}$ WxHxT, including tungsten absorbers, scintillator tile, silicon sensors, and front end electronics. The dimensions were dictated by the overall constraints of the calorimeters themselves and had a considerable impact on the design of the electronics servicing the silicon components (see next sections).

2.1 Silicon sensors

A silicon detector module consists of two micro-strip sensors, coupled with kapton pitch adapters to two Front End Hybrids (FE) which host the amplifier and multiplexer electronics. The two silicon sensors (see figure 3) are used to measure both the x and y coordinates of the particles in the shower.

The x and y layers cover a region $6.4 \times 6.4 \text{ cm}^2$, using two square sensors, $285 \text{ }\mu\text{m}$ thick, with $80 \text{ }\mu\text{m}$ pitch strips. The sensors were made by Hamamatsu Corp. for the ATLAS collaboration to be used in their tracker [3]. The strips are made with standard p type implants on n type high resistivity bulk silicon and have integrated decoupling capacitors.

Figure 3. Micro strip sensor schematic drawing. The sensor employed has 768 strips at 80 μm pitch.

Acceptance tests of the sensors were performed on arrival. Current vs. voltage and capacitance vs. voltage curves were obtained. Typical plots are shown in figure 4. All sensors received passed our acceptance criteria of less than 100V depletion voltage and less than 1 μA current at 200V. For the other parameters, like number of broken capacitors or short circuited strips, the manufacturer data were used. Such a fine pitch (standard in tracking applications) is not needed in this case where the whole shower profile has a diameter of roughly 7.5 mm and where a 100 μm resolution on the shower centre position might seem sufficient. Thus it was decided right from the beginning to read out only a subset of the total number of strips. In order to decide the granularity, shower simulations for high energy photons were performed with the FLUKA package [4] so as to derive the total number of MIPS crossing a given silicon sensor and the consequent amount of charge released.

The charge released as a function of radiation length is shown in figure 5 for a typical 2 TeV photon shower, together with a transverse profile of the charge released at 12 X₀. Following these simulation results, it was decided to keep a relatively high granularity so as to avoid saturation effects on the electronics and not degrade the excellent position resolution for the shower centres. In fact the final choice for read out pitch was 160 μm (one every two), dictated more by the needs of the chosen FE amplifier than by the intrinsic resolution requirements.

With this pitch, linearity of the FE chain is assured (see later on), allowing the silicon detectors to not only provide a position measurement of the shower centre but also a complementary

Figure 4. Typical IV and CV curves of the silicon sensors used for the position sensitive modules.

Figure 5. Simulated charge release in the calorimeter silicon sensors as a function of X_0 for a 2 TeV photon impacting on the calorimeter (left); simulated charge release on a silicon sensor placed at $12 X_0$ as a function of strip position (160 μm readout pitch) for the same 2 TeV photon (right).

calorimetric measurement of the incident particle energy which can be useful to crosscheck the scintillator measurement itself.

2.2 FE electronics: the PACE3 chip

Once the readout pitch was chosen a FE Hybrid capable of dealing with 384 channels had to be developed along with the necessary ancillary electronics.

The choice of amplifier chip fell on the PACE3 [5] developed specifically for the CMS [6, 7] pre-shower detector (see figure 6). This is a synchronous 40 MHz analogue sampler (synchronous with the LHC machine clock) with 32 channels per chip, read through one multiplexed output. An added benefit of this choice is that the PACE3 is “rad-hard” and can thus be used even at higher luminosities than those envisaged for LHCf. In fact the silicon detector modules as a whole are by far the most radiation resistant active part of the LHCf calorimeters.

Figure 6. Photo of the DELTA (left) and PACE_AM (right) dies wire bonded together on an evaluation PCB. Once encapsulated together, they constitute a PACE3 chip.

Figure 7. Functional block diagram of the PACE3 chip assembly, the DELTA to the left and the PACE_AM to the right.

The PACE3 consists of two separate dies housed in a single package. The first die containing the amplifiers and shapers with calibration circuitry is called the DELTA, the second die, called PACE_AM, instead contains the analogue pipelines (32x192 cells deep) and the output multiplexer.

The functional block diagram is shown in figure 7. Here it suffices to say that both chips have various registers through which one can set the operating parameters of the DELTA and of the PACE_AM.

Through these registers one can influence the behaviour of the amplifier and thus optimise the response for a certain application. In our case one of the most important requests was linearity of response even for very large charge deposits. Thus a comprehensive study was begun where significant chip parameters were varied in appropriate ranges, and amplifier linearity, noise and gain were constantly monitored.

The end result of this study was a suitable set of parameters that allow the PACE3 to handle up to 2 pC of charge with a deviation from linearity of 6% (see figure 8). Thus only very few central strips and only for the most energetic showers will saturate the amplifier, but centre of gravity reconstruction can anyway be achieved with high precision by fitting the sidebands of the shower profile.

The analogue outputs of the DELTA itself are characterised by a very fast time formation as shown in figure 8 (right). The PACE_AM samples these outputs every 25 ns. The phase must be chosen so as to sample the signal at its maximum.

Figure 8. Voltage output of the PACE3 chip as a function of injected charge (left); voltage output of the PACE3 chip as a function of time for different injected charges (right).

Figure 9. Timing diagrams for the PACE3 chip digital and analogue outputs. The top frame is repeated three times (see bottom part of the picture) since three contiguous time samples are presented at the chip outputs at every trigger (LV1).

The 25 ns period samples, synchronous with the LHC clock, are stored continuously in analogue memory pipelines. Upon trigger assertion the output frame presents three contiguous samples, centred around the trigger latency value, for each one of the 32 channels. Together with the analogue output, other, separate, digital lines provide strobe signals (i.e. DataValid) and memory cell addresses. The timing diagram of the PACE3 chip outputs, is shown in figure 9.

The chip comes in a BGA (Ball Grid Array) package which was conceived for the CMS pre-shower application. Thus there are minor quirks (like the loss of addressing space on the I2C bus, due to the fact that only one “chip select” address line is actually available in the packaged version!) which have translated in more routed signal buses on the LHCf hybrid than what was originally expected.

Figure 10. Two half-hybrids assembly. Input signal bonding pads are placed lengthwise along the inner borders.

2.3 FE hybrid

Following what was said in the previous section, the design of the FE hybrid was a major undertaking in as much as many chips had to find place in a relatively small enclosure. Basically the hybrid has to service one of the eight Si sensors corresponding to 384 inputs from the read out micro-strips, thus a total of twelve PACE3 chips (32 inputs) are needed. Each chip needs individual I2C buses, and each one has independent analogue and digital outputs. Overall dimensions had already being fixed by the scintillator calorimetry requirements, and the silicon modules had to be made so that they would fit in the assigned envelope. Also power dissipation was an issue since the scintillators could be damaged by excessive temperatures (i.e. above 50 °C).

One of the main issues dealt when developing the FE hybrid was fitting all the 384 bonding pads needed for the silicon detector whilst maintaining enough separation between copper traces on the PCB. The total width of the hybrid and the silicon sensor were constrained by the calorimeter width and could not be broadened. So in order to have an adequate bonding area, the bonding pads were placed lengthwise where there were no significant space constraints, and a kapton fan-out was used to connect the sensor to the hybrid.

Since twelve chips are needed for the Si sensor readout, two symmetric half-hybrids were developed (Left and Right), each with six chips. The resulting assembly is shown in figure 10 with the bonding area in evidence in the middle between the two half-hybrids. The L shaped PCBs, when assembled together, leave a void area that provides easy and component free access to the bonding machine head. The surface copper traces are plated with an AuAl alloy, because of the aluminium wire used for bonding. Dividing the assembly in two independent PCBs also increases the yield with respect to failures caused by defective electronic components or interconnection faults.

The resulting assembly form factor, while very well adapted for the analogue inputs which come from the silicon sensor, is instead unfavourable for the digital and the analogue output lines. In fact each chip digital I/O signals travel in the same direction towards the connector side with the area underneath the last chip (to the left, in figure 10) being crossed by all the nets from the

Figure 11. Detail of the PCB traces underlying the PACE3 chip nearest to the connectors. Due to the high density of traces crossing the region where the bump bonding pads for the BGA packaged chips are placed, satisfactory routing was achieved only by using a 12 layer PCB.

Figure 12. Microcoax cable assembly, connected to a half-hybrid. Total thickness is less than 5 mm.

other adjacent PACE3 chips. These high density PCB areas with a relatively high number of vias have posed a serious challenge when routing the hybrid (figure 11), which was solved in the end by using a twelve layer design.

Because of space constraints, the readout electronics, as described later, sits outside the calorimeter enclosure at a distance of roughly one metre. Thus many signals, both digital and analogue have to travel between the FE hybrids and the outside. Many digital signals (i.e. I2C lines) are delivered singly to each chip, others (clock, trigger) are common to all. The total number of signals for each half-hybrid is 80 excluding supply rails that are powered with separate wires with relative sensing.

Another constraint of the total assembly is the module thickness, that cannot exceed that of the absorber plus silicon detector plus scintillator (roughly 25 mm). To connect all the signals

Figure 13. FE hybrid’s service chips are placed in the area between the two Samtec connectors.

a particular small pitch connector, developed by Samtec [12] for high speed applications, was used, coupled to a custom assembled ribbon cable formed with microcoax single wires each with a characteristic impedance of $50\ \Omega$. The resulting connector cable assembly is shown in figure 12. Another important aspect of this cable assembly is that the single coax cables minimise crosstalk between adjacent signals, which is especially important for the analogue outputs.

As depicted in figure 13 the FE half-hybrid hosts other ICs that serve as control and signal regeneration for the PACE3 chips. The DCU [13] (Digital Control Unit) chip collects a number of voltage and current outputs from the PACE3 chips that serve to indicate the operating conditions of internal parameters, like biasing of the preamplifier and shaper stages, and the purpose of which is to calibrate the DAC blocks of the PACE3. These signals are generated on each PACE as either current or voltage sources and are then put in parallel on the hybrid with an appropriate termination on the DCU inputs. With an appropriate I2C command to the DCU, each monitored signal can be multiplexed to the internal 12 bit ADC. Before selecting and reading the DCU ADC channel of course, the actively monitored PACE3 output has to be activated from its quiescent state. On the half-hybrid there are also four LVDS and unipolar CMOS buffers [14]. These buffers regenerate the main signals coming from the outside DAQ boards and drive the daisy-chained differential inputs of the six PACE3 chips. The buffers are necessary to avoid signal degradation caused by the input capacitance of the six chips.

Given the total chip count, all the output signals and the 192 (per half-hybrid) input pads, the resulting PCB was designed with a total of twelve layers. Some of these layers are ground planes which serve different purposes, like controlling the characteristic impedance of the differential and single ended nets that need to be $100\ \Omega$ and $50\ \Omega$ respectively and acting as a shield for the nets that run parallel to each other between adjacent layers. Due to the high density of signals per layer, special care was taken in routing nets so as to minimise crosstalk. For example, one such design criteria was to keep a minimum distance of five layers separating sensitive nets from each other.

Another issue to be tackled was power dissipation. In fact a fully loaded half-hybrid requires a total of $1.6\ \text{A}$ at 2.5V and so a total power dissipation of $4\ \text{W}$ is expected. In figure 14 an

Figure 14. Thermal imaging of a powered half-hybrid. Dark-blue corresponds to 18 °C and white-red to 60 °C.

Figure 15. X-ray image of the solder pads, underneath a PACE3 chip (left). 3D view of the solder paste bumps on the same pads (right).

infrared picture of the half-hybrid with all chips powered is shown. The picture was taken at a normal ambient temperature of 20 °C. Chips stay at 45 °C while the PCB board is at 38 °C in the hot region falling to 28 °C at the connector side. In order to reduce this thermal gradient and to facilitate heat dissipation within the calorimeter enclosure all the PCB ground planes are connected at the edges to the module aluminium support frame (see next section).

As mentioned above, for mechanical and bonding adaptability the total thickness of the PCB had to be 1.9 mm with a ± 0.1 mm tolerance, resulting in a layer separation of roughly 6 mils. Again this has implications on the copper traces widths which have to be roughly 4 mils to obtain 50 Ω single ended net or 100 Ω differential. With these technological parameters not all PCB manufacturers were capable of achieving good yields and reliability.

Finally, after a few trials, one firm [15] was chosen as a single supplier to deliver the complete set of 12+12 left(right) half-hybrids. Also at this point, it was decided to outsource the component

Figure 16. Clock signal (top) and LV1 trigger (bottom) monitored on the oscilloscope, during hybrid acceptance tests.

mounting because of the ball grid array packaging of the PACE3 chips which required special tooling not available in-house. Again, because of the scarcity of components, a very high yield was requested together with an extensive quality assurance procedure. In the end one firm [16] was chosen for all hybrid assembly related tasks.

Once finished the populated hybrids had to be screened for all possible defects. Before the functional test, described below, we inspected, by means of x-ray imaging, the half-hybrid PCBs ensuring that no subtle defect was present in the hybrid assembly. A typical x-ray picture is shown in figure 15 along with a 3D view of the solder paste ball grid array deposition. Then an extensive thermal cycling was conducted for each module between -20°C and $+60^{\circ}\text{C}$ with a 50% duty cycle for a total of 48 hours. Afterwards, a full functionality test of the module was performed, with a laboratory sub-set of the experiment read-out system. During these tests, in turn, each single electrical signal was also monitored with an oscilloscope (see figure 16) and each internal register was written and subsequently read with all allowed values. Also power supply parameters were verified and current absorption levels were monitored.

To connect the hybrid to the microstrip sensors a set of kapton fan-outs were made, together with very thin PCB pitch adapters, that adapt the sensor read-out pitch ($160\text{ }\mu\text{m}$) to the hybrid's (0.5 mm). A photo is shown in figure 17. The kapton fan-outs and pitch adapters were made at CERN with gold evaporated on the surface and subsequently etched with the desired pattern. The hybrid is glued to these kapton fan-outs which are subsequently glued again and bonded to the pitch adapters which are very flexible and also house the high voltage biasing network for the silicon sensor.

Figure 17. Kapton fan-outs (for left and right half hybrids) and pitch adapters (for X and Y oriented sensors). The green coloured pitch adapters have a large pad where conductive glue is placed. The pad provides biasing to the Si sensor backplane.

2.4 Module assembly

The assembly of the tracking modules has been done, together with the whole LHCf ARM2 calorimeter, at the Clean Room Facility located in the INFN-Firenze laboratories. Micro movements were used throughout under microscope supervision, and a DEA 6600 precision measuring machine was used to verify final assemblies. Also the wire bonding of the tracking modules was done in-house with a Delvotec 6400 wire bonder available at the Facility. In this section the tracking module assembly is described following a chronological order of the various steps.

A complete tracking module houses both the calorimetric part and the silicon micro-strip components. Two silicon assemblies are needed, one for the X view (strips parallel to the long side of the hybrid) and the other for the Y view (strips orthogonal to the long side of the hybrid). All the main pieces pertaining to the module are shown in figure 18. A black Delrin (Polyoxymethylene from Dupont) frame houses the tungsten absorbers with the corresponding scintillators and protects the silicon sensors from possible mechanical damages. The hybrid edges are held in contact with screws to the aluminium frame which is in thermal and electrical contact with the calorimeter walls. Aluminium was used in the hybrid zone because of its excellent heat conductivity, whilst Delrin was used in the active region of the calorimeter because it's a “light” material which helps avoid having particle showers initiated outside the volume of the tungsten absorbers. The thin (0.5 mm) aluminium plate on which the silicon sensor and hybrid lie also helps in the handling of the module during assembly and is at the basis of the whole assembly procedure.

Figure 18. Module assembly exploded view. Brown: tungsten absorber, Blue: scintillator tiles, Green: silicon sensor, Orange: Kapton fan-out and pitch adapter, Gold: thin (0.5 mm) aluminium plate, Light Grey: aluminium frame, Dark Grey: Delrin frame. Also visible (to the right) the back side of the second silicon and hybrid.

First the assembled hybrids are accurately placed and then glued to the kapton fan-out visible in the left part of figure 17. Care must be taken so that the bonding pads on the PCBs and on the kapton foil match closely (roughly 10 μm alignment tolerance). A finished hybrid pair glued to its kapton fan-out is the one shown already in figure 10. The hybrid assembly is placed in a custom made jig (only partly visible) which is used throughout the procedure as a precision holder.

At the same time, the other PCB circuit, the pitch adapter, is glued on the back of a silicon micro-strip detector. Two types of circuits are used depending on the orientation of the strips (parallel to the long side of the hybrid or orthogonal). This operation is done with the help of a jig with vacuum suction holes to hold the pitch adapter in place. The large central bias pad (see left part of figure 19) is then covered with a thin (30 μm) layer of conductive glue (E-SOLDER 3025, Epoxy Produkte GmbH), while the rest of the surface is covered with a thin layer of Araldite 2011. The glue was then cured at 60 $^{\circ}\text{C}$ for 12 hours.

The sensor is then placed by hand on top of the PCB circuit and aligned under a microscope using reference marks etched on the circuit and on the sensor. To ensure that the conductive glue obtains its best electrical properties, a slight pressure is applied to the silicon sensor by placing a 5 Kg brass weight (clothed with a dust free tissue) on top of the sensor.

Figure 19. Partially assembled module held in the precision assembly jig. Glued pieces are, from the left: hybrid assembly, PCB pitch adapter and the silicon sensor (strips parallel to the hybrid).

Table 1. Measured X,Y coordinates of the silicon sensor reference marks. Variations for both the four X and the four Y oriented sensors are within $\pm 200 \mu\text{m}$.

Sensor	X(mm)	Y(mm)
Nominal	-19.000(mm)	9.500(mm)
1X	-18.864	9.305
2X	-18.930	9.509
3X	-19.280	9.367
4X	-18.799	9.556
Nominal	44.760(mm)	73.260(mm)
1Y	44.969	73.338
2Y	45.069	73.295
3Y	44.953	73.185
4Y	44.659	73.253

Once the hybrid is assembled, and the pitch adapter with the Si sensor is ready, the 0.5 mm thin aluminium plate is placed on the precision jig. The thin plate is then held in position with the use of precision dowel pins. These pins are then used to guide the hybrid assembly in the correct position on top of the plate where it is then glued. At the same time, the pitch adapter, with its glued silicon sensor, is also placed on the thin aluminium plate, aligned to the hybrid and the fiducial marks of the plate and then glued (see figure 19). The glue used in both cases was Araldite 2011. Great care was taken in ensuring that an even amount of glue was placed underneath the bonding pads with no visible air bubbles. Alignment was ensured also by microscope using the various bonding pads as reference marks. The glue was cured at room temperature to avoid stresses due to the different thermal coefficients of the materials.

After the glue is set, the jig with the hybrid and sensor components is ready for bonding. Since each sensor has 768 strips and only 384 electronic channels are available, only 1 strip out of two is actually connected to the pitch adapter, while the other is left floating. The wire used was 25 μm aluminium bonding wire for hybrid to kapton fan-out, for sensor to pitch adapter bonding and for kapton fan-out to pitch adapter PCB bonding. Trial runs were performed to ensure that a set of optimal bonding parameters were used.

Figure 20. Assembled module, consisting of two joined pieces, for X and Y coordinate readout, alongside the two precision jigs used in the assembly. The end connectors and the power cables of the hybrid are visible. The black Delrin frame also houses the tungsten absorbers and the scintillator plates which are read out through a fibre wave guide.

During these runs the bonds were subsequently tested with a bond puller strength gauge and were verified to exceed 7 grams. Nearly 10000 wire bonds were made without any significant failure or reworks.

Once the bonding is done, the same jig is used to fix both Delrin and aluminium frames. Dowel pins are used to guide the frames in position. A thin layer of Dow Corning 340 heat sink compound is spread on the inside edges of the aluminium frame, after which the screws are tightened. The Delrin and aluminium frame are then screwed together while keeping the assembly in place on the jig with vacuum suction, thus keeping the pieces aligned (figure 20).

Tolerances achieved on the assembled modules were of the order of a few 100 μm maximum displacement (see table 1) as measured between the centre of the sensor and the centre of the aluminium/Delrin frame using the through holes for reference marks. These measurements were made with our DEA 6600 measuring machine which has an intrinsic accuracy of about 3 μm .

At this point the tungsten absorbers and scintillator plates are inserted in the black Delrin frame completing the X (Y) module. Two such assemblies, with differently oriented sensors, are then joined together back to back to obtain the final module with X and Y coordinate readout. As stated before, in order to avoid reaching temperatures capable of damaging the scintillators, the aluminium frames are fixed with screws to the calorimeter sidewall which dissipates the heat produced by the electronics. During this final part of the calorimeter assembly another thin layer of heat sink compound is applied.

All finished modules (position sensitive silicon and plain calorimeter ones) are stacked together (figure 21) using precision high rigidity G-10 bars which traverse the calorimeter from front to back. These bars, together with the bottom plate of the calorimeter, ensure the relative alignment of all the modules. Each module has precisely machined through holes for the passage of these bars.

Figure 21. Partially assembled calorimeter. The four silicon modules are in position together with the G-10 frames containing the other scintillators/absorber sandwiches. In addition to the many fibres used as wave guides, the micro-coax ribbon cables are also visible (blue) together with the modules power cables (black and red). The modules aluminium frames have threaded holes on the sides for the screws that hold them in contact to the calorimeter external wall, an 8mm thick copper panel.

3 Front end readout and control electronics

LHCf has a FE chip developed precisely for the purpose of working with the LHC machine. Thus the DAQ system has been based on a synchronous architecture designed to work at the nominal clock frequency of 40.08 MHz (matching the LHC clock frequency).

The DAQ can be subdivided in two main elements. One comprises the clock and trigger distribution plus the handling of the various control registers (Front End Control), the other is related to digital to analogue conversion and data transmission and storage (Front End Readout).

From the counting rooms, where the computers and machine interfaces are located, to the LHC tunnel, where the experimental apparatus lies, there is a distance (cable length wise) of more than 200 metres. Thus a considerable amount of electronics was placed as close as possible (one metre) to the detectors so as to minimise the number of links needed. Digitisation is performed in situ, while data communications and control has been implemented with optical fibres effectively minimising the number of links while at the same time electrically decoupling the detector in the tunnel from the electronics in the control room.

An overview of the whole electronic chain is given in the following sections. In figure 22 a block diagram depicts the various items with the most significant connections in evidence. Detailed descriptions of the LHCf read out and control can be found in [17, 18].

Figure 22. Block diagram of the electronics servicing the detector modules. One detector module (two sensors) is connected to a MDAQ board which hosts four ADC boards. Clock and trigger signals are distributed through digital opto-hybrids. Output data from the ADCs is sent via optical links (FOXI) to the control room computers.

3.1 Front end control

The front end control is in charge of distributing the machine clock and the trigger (arriving from the calorimeter section) to the FE hybrids and the acquisition boards in the tunnel. It also allows the user to write and read all the registers of the PACE3 chips and other parameters used in the

Table 2. 3-bit code sequences. T1 is generated by the calorimeter and is synchronous with the LHC clock. CalPulse triggers an internal charge injection on the PACE3 inputs used for calibration purposes. The Re-Sync sequence actually resets the pipe line pointers to the programmed latency value. BC0 is a special signal from the machine indicating an empty orbit.

Pattern	Command	Function
100	T1 (LV1)	First Level Trigger
110	CalPulse	Calibration Trigger
101	Re-Sync	Resets pipe line
111	BC0	Bunch Crossing Zero

Figure 23. Encoded stream with the LHC clock and calorimetric trigger. The 40 MHz clock is fully recovered by the PLLs in the MDAQ boards.

Main Data Acquisition boards (MDAQ, see next section). It is implemented through the use of a PCI controller board (FEC [19]) developed for the CMS tracker which interfaces optically to a digital optical hybrid module (DOHM [20, 21]) which is placed in proximity of the detector. The trigger, fast re-sync and a special calibration trigger used by the FE electronics are encoded in a single stream together with the LHC 40 MHz clock. A sequence of three bits defines what kind of operation is requested of the FE (see table 2).

The trigger appears as a single missing clock (figure 23), the calibration trigger as two consecutive missing clocks, while the re-sync is given by two missing clocks separated by a single clock transition. The slow control parameters instead travel on a separate data line, synchronous with the transmitted clock. Thus 2+2 fibres are needed for a control link, two to transmit the encoded clock and data and two to receive the return signals from the control ring.

The DOHM distributes the signals to eight Communication and Control Units (CCU [22]), one for each hybrid. The distribution is organised as a token ring architecture (figure 24) with each CCU answering only when addressed. The CCUs provide up to 16 I2C buses and a certain number of individually addressable I/O. These I2C buses and I/O are used for setting PACE3 parameters and acquisition modes of the MDAQ.

Operations, in a ring architecture, can be jeopardised by the failure of one of the CCU nodes. Since access to the LHC tunnel during machine operations could be problematic and of short duration anyway, it was decided to implement a redundancy scheme. This issue had already been addressed in the CMS tracker with the use of a scheme where clock and data lines are duplicated. This redundancy is wired in the DOHM and in the FEC module which actually handle both a primary circuit and a secondary one; details of this scheme can be found in [21]. Here it suffices to

Figure 24. Simplified scheme of the CCU control rings (left). The FEC sits in the control room while the DOHM (photo on the right) and the MDAQs boards (that host the CCUs) are in proximity of the detector. The DOHM boards (Main + auxiliary) can connect up to 15 CCUs providing redundancy. The fibres (visible on the right) are connected directly to the FEC, 200 metres away. From the connectors, twisted pair flat cables transport the signals to and from the CCUs.

say that the ring is resilient to single point failures and also to two point failures as long as they are not caused by two consecutive elements in the ring.

Referring to figure 22 the Front End Control hardware is distributed on various boards. The FEC and the trigger encoder are PCI boards in the counting room. The DOHM (figure 24 at the right) is a board placed next to the MDAQ boards. Finally the CCUs and the PLLs [23] (needed for clock recovery and trigger decoding) are hosted on the MDAQ boards themselves.

The control part of the DAQ is not only involved in the distribution of “fast” signals, and of setting the PACE3 registers, but also monitors slow control parameters like temperatures and voltages through the DCUs present both on the FE hybrids and on the MDAQ boards. In fact the DCUs have on board a temperature sensor which is read out through the internal ADCs. Other temperature probes are wired directly to the counting room, in order to monitor detector ambient conditions without having to power on the electronics.

3.2 Readout and Data transmission

The Main Data Acquisition Boards (MDAQ), while hosting the clock and trigger distribution (CCU and PLL) as described in the previous section, has the main purpose of digitising the analogue data from the PACE3 chips, and send it via optical transmitters to a receiver board in the counting room.

The MDAQ board (see figure 25 left) was designed and developed at our INFN-Firenze laboratories and is based on a ALTERA Cyclone II FPGA (EP1C6Q240C7) which handles the various acquisition sequences with the use of a Finite State Machine (FSM). Each board measures 430 mm X 190 mm, and consists of eight layers of which six needed for the signal routing and two for power distribution. Like the hybrid, heat from the electronic components is carried by the ground plane to the edges of the boards, where thermal vias put in contact the PCB copper layer with the metal cage hosting all the detector electronics. Signals are mainly LVDS and so all lines have a controlled impedance of 110Ω. Feature size for the copper traces are of the order of 7 mils width and 2.8 mils thickness, while isolation has been kept at a minimum of 8 mils.

The ADCs needed for the digitisation are hosted on a smaller PCB (see figure 25 right) of which four are connected to one MDAQ board. The ADC boards were designed using only six

Figure 25. The MDAQ circuit board (left). The FPGA sits in the centre, the two CCUs are at the top and bottom edges while the optical transmitter is placed to the right of the FPGA. The empty spaces at the left are for the ADC daughter boards (photo at the right). The ADC daughter board takes care of both analogue to digital conversion and data storage. The three chips at the top are dual ADCs and the four at the bottom 18-bit wide FIFOs, connected to a parallel 72 bit wide bus.

planes (two of which for power distribution). Like the MDAQ boards, here too some of the surface was used for heat dissipation with bare copper pads in thermal contact with the electronics.

Each ADC board can handle the analogue stream from a half-hybrid. Thus one MDAQ services two sensors which may or may not belong to the same module.

In fact the final cabling was done so as to have some further redundancy in case one MDAQ board fails. Thus sensor 1X for example, is connected to the same MDAQ board as sensor 3Y. In case this particular MDAQ board should fail only one view of the first and third module are lost, but the other views are still capable of sampling the shower profile.

Referring to figure 22, data sourced by one half-hybrid (six PACE3 chips) is sent to an ADC board which hosts three dual ADCs from Analog Device (model 9238BSTZ-40, 12 bit, 40MSPs). The ADCs are clocked continuously, but only on trigger assertion are the converted values stored in the FIFOs (TI model SN74V225). Thus the triggered event is stored for each half-hybrid of the

Figure 26. Flow diagram of the FSM (Finite State Machine) implemented in the control FPGA.

calorimeter locally on an ADC board. Once the PACE3 has finished multiplexing the analogue data out, the digitised data are stored in the FIFOs; the total time needed is 92×3 clock periods (25ns) = $6.9\text{ }\mu\text{s}$, after which the next phase can begin.

The read-out of the four ADC daughter boards (4 half-hybrids, for a total of $4 \times 6 = 24$ PACE chips) is carried out in parallel by the MDAQ board. Each ADC board is individually read on a 72 bit wide data bus, which is then serialised in a sequence of bytes called a “Data Frame”, with at the end a CRC control code calculated by the controller FPGA that controls this parallel read-out. The “Data Frame” is transmitted to the counting room by the Fiber Optical Transmitter/Receiver Interface (FOXI) transmitters. These are 100 Mbit/s links which are plugged on the MDAQ board and can be replaced, should the need arise, by faster (1Gbit/s) links quite easily.

The time needed for trigger processing (read-out + transmission) is $370\text{ }\mu\text{s}$. This is the dead time necessary for trigger processing, during which the BUSY output line is kept asserted by the controller FPGA and effectively inhibits further triggers from the calorimeter. The system can be reconfigured so as not to inhibit the trigger in such a “hard” manner. The FIFO buffers allow the

storage of up to 10 events, also the PACE3 chips have an internal FIFO allowing the chips to receive triggers even while they are multiplexing out data from previous events. In reality, since the ARM1 and ARM2 calorimeters maximum DAQ acquisition rate is 1 KHz, the default configuration for the silicon electronics is the one with “hard” inhibit which is simpler to control.

Figure 26 shows in more detail the main MDAQ operations with the aid of a flow chart diagram. The default state of the FSM (Finite State Machine) is idle (after power on). When in idle state, the first LV1 trigger is always accepted and processed; after processing, as last operation, the board returns to idle state.

For each accepted LV1 trigger (event) the MDAQ controller FPGA performs a fixed sequence of operations:

-START-

- step 1-** it immediately sends a LV1 pulse to all 24 PACE3 chips and sets to 1 the BUSY output line;
- step 2-** it transfers event data from PACE3 chips to the FIFO chips, by reading from ADCs and writing to the FIFOs.

Each PACE3 chip, after receiving the synchronous LV1 pulse, generates a fixed sequence of data on its analogue output line corresponding to the information stored on 3 consecutive columns for all its 32 internal channels (figure 9). Before transmitting each sequence of 32 analogue read-out values, the PACE chips send the corresponding 8-bit column address on a dedicated COL_ADDR digital LVDS output line; also whenever a valid read-out value is present on the analogue output line, another auxiliary digital LVDS output line (DATA_VALID) is kept asserted. The 24 analogue PACE3 output lines are simultaneously and continuously digitised by 24 ADC sections and each ADC 12-bit parallel digital output is input to the FIFO chips.

The FPGA generates the synchronisation signal for the ADC chips (ADC_CLOCK) in such a way that the sampling of the PACE3 chip multiplexed analogue output is properly performed. The proper sequence of control signals for the FIFO chips is generated by the FPGA in such a way that each ADC word is immediately and correctly written in a FIFO location.

During this phase the FPGA also reads the content of the COL_ADDR lines and writes the values in internal temporary registers (also known as event registers), it then checks that the 3 consecutive column addresses (which refer to the three consecutive time samples that the PACE3 chip provides) are effectively consecutive numbers as expected during proper operation, and also that all PACE3 chips are transmitting the same sequence of column addresses, which is the expected behaviour after they have been properly initialised at the beginning of the run.

Assertion of the DATA_VALID lines for each PACE3 chip is constantly monitored while the analogue data are being sampled by the on-board ADCs. Additional checks are performed on the FIFO_EMPTY and FIFO_FULL outputs from each FIFO chip, to ensure that each FIFO memory is behaving as expected, i.e. it is empty before event data writing is started and it does not become full while event data are being written. If an error condition is found for

Figure 27. Data Frame sent to the counting room by the MDAQ FPGA.

any of the previously described consistency checks, a corresponding alarm code is written in the event registers.

-step 3- after finishing the first read-out phase, the FPGA transmits the event header, the event data buffer (i.e. the event data previously stored on FIFO chips) and the event trailer through the optical fibre FOXI channel. This is the “Data Frame” (see figure 27), which will be received and decoded by the control room electronics.

The proper sequence of control signals for the read-out of the FIFO chips is generated during this phase.

The event header contains auxiliary information on the event, in particular an event counter (increased by one with each new accepted LV1 trigger) and all the information written on the temporary event registers during step 2. The event trailer contains a cyclic redundancy check (CRC) code which is calculated over all the event header and data buffer, with the purpose of verifying that subsequent handling of the data has not introduced any data corruption.

-step 4- finally, the FPGA resets to 0 the BUSY output line and returns to the idle state ready for a new LV1 trigger.

-END-

Four MDAQ boards are needed for the four silicon XY modules. The data sent are received, in the counting room by a custom built VME receiver board. The board acts as a deserializer and presents the data from the MDAQ boards as a sequence of 32 bit words. CRC checks are implemented in hardware (like in the MDAQ) and also the header of the event is decoded to check in real time for eventual alarm signals from the detectors.

Figure 28. The power distribution board. The two LV power connectors are at the right. The three side connectors are used for (from left to right): Filtering of the hardwired temperature sensors, Silicon sensor biasing filters, LV sensing. The LV regulators on board (visible in the bottom part) provide the stabilised 2.5 Volt to the silicon module hybrids.

3.3 Low voltage powering, and biasing of the detectors

All power supplies, for reasons of accessibility are placed in the control room and not inside the tunnel. To keep the detector electrically decoupled from the control room electronics, floating power supplies were used for powering the detector and biasing the silicon sensors.

The powering system of the LHCf detectors is installed inside two racks which are located in the counting room in the USA-15 cavern. The low voltage levels for the silicon tracking system of the ARM2 detector are generated by means of four Agilent N6700 LV supplies, each provided with four independent fully programmable output sections (Model N6732B-ATO DC Power Module 8V, 6.25A, 50W). Because the power lines are common to each couple of half-hybrids connected to a silicon sensor (i.e. either the X or Y section of a whole module), eight independent lines are sufficient for powering all four ARM2 tracking modules. Two fully loaded N6700 (eight lines in total) are therefore used to generate the 2.5 V for the front-end electronics. The third N6700 generate the four 5 V sections which are needed for the MDAQ motherboards which manage data transfer from the front-end circuits to the VME digital acquisition board. Finally the first section of the fourth Agilent supply is used to generate the 5 V model for powering the DOHM board while the other three sections are kept as spares. For silicon sensor biasing, a CAEN A1519 HV board provides the necessary 250 V.

All the LV and HV supplies are placed at an accessible position far away from the experiment. The power lines which connect the supplies to the detector in the TAN region are screened multi-polar cables more than 200 m long. All the conductors are dimensioned in such a way as to limit the voltage drop between the power supply and the load to a few volts. Each of the N6700 output

Figure 29. Connection diagram for the LHCf slow control.

lines are provided with two remote sense connections, which are used for the remote regulation of the power levels directly at the load instead of at the output terminals, thus allowing the automatic compensation of the voltage drop along the cables. All the supply modules (including the HV ones) are electrically floating and the grounding point of the return lines and of the cable shields has been placed on the distribution board (figure 28) where filtering and low voltage drop regulators are used to provide a stable supply to the local electronics. This was found to minimise the noise level on the front-end and read-out electronics.

3.4 DAQ and slow control software

A custom Slow Control software has been implemented in C++ language to allow controlling, switching on/off and setting up all the LHCf devices. This software is based on a server/client system which can manage multiple connections from users inside the CERN network domain.

The Slow Control Server (SCS) (figure 29) runs on a dedicate machine in the LHC counting room located in the USA-15 cavern. Its main tasks are verifying the hardware status, setting some alarm flags true in case some subsystem is not working properly and logging all the relevant information of the run. Typically the information logged consists of FE and run parameters and a record of all operations performed either by the user or by background tasks. Working parameters are read and stored to a file every few seconds, in such a way as to allow a detailed reconstruction of the run-time status and configuration of the whole system and to have a fast identification of a possible anomalous condition. A Slow Control Client (SCC) software can be launched from any PC within the CERN intranet which has permission to connect to the SCS main process.

The SCC can be run in two different modes: operative mode and monitoring mode. In the operative mode two different user levels are defined: standard and expert. Standard user can only

run a sub-set of safe commands and a certain number of pre-defined procedures, that are safe sequences of commands used to perform complex operations, like switching on all the low voltages of a sub-detector in the correct way or switching off one of the LHCf detectors completely. On the other side the expert user has a complete access to the resources and can execute any implemented command. A password is required to enter the expert mode. The SCC operative mode allows managing all the LV and HV of all the sub-systems, including scintillators and scintillating fibres for detector I and scintillators and silicon layers for detector II, reading temperatures from eleven PT1000 sensors installed inside detector II, and moving the detectors from the “garage” (protected) position to the “running” position or vice versa.

A simple command interpreter has been implemented with the aim of having a common syntax when calling commands for the different sub-systems. An on-line help makes it possible to recall the usage of all commands. If the monitoring mode is activated, the SCC enters a text-based monitoring status, where parameters for all the relevant sub-systems can be kept under control. Unsafe conditions are highlighted by following a colour scheme, in such a way that the operator in shift during data taking can promptly recognise the system where the problem has arisen.

The last important task of the SCS concerns the exchange of information between the experiment and the LHC. On one side the LHC machine publishes continuously some useful variables that can be acquired by all the experiments. These variables, which contain information about the beam status, the operations in progress and the machine parameters, can be accessed through the network by connecting to the DIP (Data Interchange Protocol) service, a server/client system implemented by LHC. In case some operation is done on the beams, or the machine is in the injection phase, the LHCf detectors are automatically moved to their safe positions to avoid possible damages. On the other side the LHC requires some information from the LHCf detectors to be published in real-time on the DIP service, including the status of each detector, the trigger rate, and estimates of the beam position updated every few seconds on the basis of partially reconstructed data samples. This information is used by the LHC to optimise the beam parameters and the collider conditions.

4 Measured performance

Many tests were performed on the silicon modules, both during production and afterwards when they were integrated in the calorimeter. These tests not only served to verify that the assembled pieces were fully functional but also allowed the development of a suitable set of software tools to analyse the data.

The modules were fully characterised in the laboratory for dead/faulty strips, gain, and noise. The assembled calorimeters were placed at CERN test beams, and bombarded with electrons, hadrons and muons of various energies to characterise their energy and impact point resolution. In the following a sample of the many results obtained will be shown.

In figure 30 and in figure 31 typical distributions of pedestals and pedestal RMS values are shown for all the silicon detector channels (3072). These plots show the excellent quality achieved in the design and assembly. They were taken during the 2007 test beam run in full operating conditions in a “noisy” environment.

Figure 30. Pedestals for the four Y views (left) and the four X views (right). There is very good homogeneity between different sensors and also between different chips in a same module.

4.1 Laboratory set-up

The laboratory set up which was used for module characterisation is a small subset of the final DAQ. The final FEC and a single MDAQ board were used, coupled to a prototype VME receiver board.

Figure 31. Pedestal RMS values for the four Y views (left) and the four X views (right). Only one bad channel is evident on a total of 3072.

This set up was also used during hybrid production for the acceptance tests. For each assembled module a complete test was made with all relevant data written to a table. For example using the

Figure 32. Typical calibration curves for the registers of one PACE3 chip (left) and Vshaper register setting dispersion (as derived from the calibration procedure) for a sample of PACE3 chips mounted on the half hybrids (right). These calibration curves are all obtained against the DCUs internal 14 bit ADC which is used as reference.

DCUs internal ADCs all the register values for the various PACE3 parameters were calibrated. In figure 32 a typical curve obtained for one of the calibration data sets is shown, together with a dispersion plot for one of the most relevant parameters affecting the gain of the chips. In general there is a substantial homogeneity between different chips. Nonetheless these tables provide us with the means of setting all the PACE3 in exactly the same working conditions. This setup was also used to obtain the linearity data shown in the preceding sections and the scope images used for hybrid qualification (see figure 16).

Once the modules were integrated in the calorimeter, only a functional test was performed to ensure that no damage had occurred during assembly. Also a HV biasing test was performed to ensure the integrity of the sensors.

4.2 Test beams

Three test beam campaigns have been performed, in 2004, 2006 and in 2007. The last two were very similar for what concerns beam conditions and experimental set up. The results shown in this section refer to the 2007 set up, with the final calorimeter assembly.

Typically the test beam set up (see figure 33) used a precision telescope tracker taken from the ADAMO experiment ([24]) to determine the precise impact point of the incoming particle on the calorimeter surface.

A couple of scintillators provided the particle trigger and data were taken mainly with electrons (up to 200 GeV) and protons (up to 350 GeV). Also some muon data were taken to verify if the silicon modules could detect the passage of MIPs through the calorimeter. In figure 34, 35 and 36 the response is shown to three types of incident particles of the silicon modules (placed as stated in the first section at 6, 12, 30 and 42 X_0). The first one is a high energy muon which behaves as a MIP releasing very little charge in the silicon detector; PACE3 chips were set in High Gain mode for this measurement. The other two are the shower profiles of a hadronic and electromagnetic particle, with PACE3 chips set in Low Gain mode.

The shower centre position is the most important quantity that the silicon modules must provide for the correct operation of the LHCf experiment. As shown in figure 37, the measured

Figure 33. Photo of the test beam set up. At the front behind a thin bronze shield sits the ADAMO telescope. Right behind there's the aluminium enclosure of the LHCf calorimeter (ARM2). The blue micro-ribbon cables together with the LV power cables can be seen protruding from the calorimeter, and reaching to where the MDAQ and DOHM boards are placed.

resolution for the shower centre position on the first silicon module is better than $40\text{ }\mu\text{m}$ after the alignment procedure ($\sim 50\text{ }\mu\text{m}$ before).

The alignment procedure, used as a prerequisite to improve the test beam data analysis, is performed in two subsequent steps: first the five layers of the ADAMO tracker are internally aligned, then the X and Y silicon sensors of the LHCf Arm2 detector are aligned with respect to the ADAMO telescope. For the first step a specifically defined χ^2 , constructed with the sum

Figure 34. Muon passing through the calorimeter as seen in the four X views (left) and four Y views (right) with the silicon modules in ARM2, with the PACE3 chips set in High Gain. The silicon modules can track single MIPS (once the preamplifiers have been set in High Gain), as well as pC signals originating from dense electromagnetic shower cores (in Low Gain).

over some thousands events of the difference between the “real” x-y ADAMO measured points and the extrapolated ones from a linear fit of all the reconstructed points, is minimised. This χ^2 is essentially function of the x and y offsets and θ_{xy} rotation angle of the various ADAMO modules.

Figure 35. Proton (350 GeV) interacting in the calorimeter as seen in the four X views (left) and four Y views (right) with the silicon modules in ARM2 (Low Gain). The shower starts after the first silicon module which as a consequence does not show any signal. Also the shower has a wide lateral extension with secondary emissions visible in the last layer.

Once the ADAMO tracker has been internally aligned, the reconstructed track is extrapolated towards the LHCf detector. The next step is to align the Arm2 silicon sensors with respect to the ADAMO defined reference system; this is done by evaluating a χ^2 of the differences between the

Figure 36. Electron (200 GeV) interacting in the calorimeter as seen in the four X views (left) and four Y views (right) with the silicon modules in ARM2 (Low Gain). The shower starts immediately but extinguishes itself before the last two layers which as a consequence do not show any signal. The shower has a narrow lateral extension.

extrapolated tracks and the measured positions of the reconstructed shower centres. The χ^2 is then minimised by translating (x-y offsets) and rotating (θ_{xy} rotation) the silicon sensors. In this way the residuals between the ADAMO track and the reconstructed shower centres are optimally evaluated, maximising the measured spatial resolution.

Figure 37. X residuals distribution of the shower centre as reconstructed by the first silicon module ($6X_0$), referred to the extrapolated ADAMO telescope track, after alignment. Data taken for 200 GeV electrons.

Figure 38. X and Y residuals spread as a function of incident electron energy (first module data only) compared to MC predictions.

As expected the resolution depends on the energy of the incident particle (see figure 38). In general the Y view has lower performance. This can be attributed partly, to the longer connections (fan outs) between the Y sensor strips and the half hybrid assembly, which slightly worsens the S/N ratio. Interestingly though, the main cause of this performance loss, which has been verified in our simulations, is due to the greater distance between the tungsten absorber and the Y sensor (1.5 mm) with respect to the X one.

Figure 39. Calorimetry measurement of the electron energy using only the first and second module data (for 200 GeV electrons).

The calorimeter has of course been simulated completely and so have the silicon modules. Full details of the simulation program are given in [1]. For the X and Y residuals there is some discordance between what is expected from the MC and what the data show. In general the MC predictions are more optimistic by roughly 20% which probably derive from an imprecise description of the charge sharing behaviour of the sensors. Work is ongoing to improve the agreement for the start of the LHC data taking.

Another very interesting result is the actual energy measurement using only the silicon data. This is indeed possible as long as the PACE3 chips do not saturate, a condition that happens in our estimations, only for a small number of cases, and only for very (above 2 TeV) energetic photons. In figure 39 the energy resolution for 200 GeV electrons is shown. The result obtained (summing the first and second modules) is of the order of 10%. In figure 40 the linearity from 50 to 200 GeV, corresponding to the extent of the available electron energies, is shown.

This capability of the silicon modules will be used as the experiment ages to cross check the scintillator response. In fact the organic scintillators are very sensitive to radiation dose, while the silicon sensors and FE hybrids are relatively radiation hard having been developed specifically for LHC operations (tracker/calorimetry).

5 Conclusions

The final installation of LHCf-Arm1 and LHCf-Arm2 inside the LHC tunnel has been completed in January 2008 (see figure 41) and the first signals were recorded during the LHC start-up test on 10th September 2008. The synchronisation with the LHC clock and bunch crossing signals was performed without glitches, and generally the DAQ and slow control worked flawlessly.

Figure 40. Linear behaviour of the silicon modules response as a function of electron energy. Calorimetry measurement of the electron energy using only first and second module data.

Unfortunately the machine broke down before any proton proton collisions could be established but beam gas data was recorded. Initial data looks very promising. Beam gas background seems to be well within predictions made by us, taking into account the expected vacuum conditions. LHCf will start its measurements during the beam-commissioning phase at 3.5+3.5 TeV energy which is planned in the fall of 2009, and will complete its first part of the physics program with the 7+7 TeV runs.

Acknowledgments

For the design and production of the kapton fan-outs and pitch adapters we wish to acknowledge the help and support of Manuel Sanchez Suarez of CERN who also provided helpful inputs to the module design, and of Rui de Oliveira in charge of the production.

We are also indebted to David Barney and Paul Aspell, from the CMS experiment, for the precise answers to our many questions on the PACE3 chip and for their enormous patience and goodwill.

Special thanks go to a former student of ours, Vladimiro Noce who graduated while working on this detector. He now pursues at another Department his interests in engineering; we wish him the best of luck.

Finally we wish to acknowledge the help of Mirko Brianzi who bonded the modules, of Enrico Scarlini responsible for the Clean Room Facility and the equipment within, of Florentina Manolescu who helped in the assembly of the modules, and of Nunzio Guardone, who contributed to the final cables installation at CERN.

Figure 41. Photo of ARM2 in its final installation position inside the TAN, 140 metres from the ATLAS interaction point. The calorimeter sits in the middle, the cage with the silicon detector electronics is at the right, while the calorimeter phototube amplifiers are at the left. Behind is the luminosity monitor.

References

- [1] O. Adriani et al., *The LHCf detector at the CERN Large Hadron Collider*, [2008 JINST 3 S08006](#).
- [2] LHCf collaboration, O. Adriani et al., *Technical design report of the LHCf experiment: Measurement of photons and neutral pions in the very forward region of LHC*, [CERN-LHCC-2006-004](#) (2006).
- [3] ATLAS collaboration, *Inner detector technical design report*, CERN-LHCC-99-16 and CERN-LHCC-99-17 (1997).
- [4] A. Fasso, A. Ferrari, J. Ranft and P.R. Sala, *FLUKA: a multi-particle transport code*, online at <http://www.fluka.org>, CERN-2005-10 (2005).
- [5] P. Aspell, *The design and development of the front-end electronics for the CMS preshower detector*, Ph.D. Thesis, Lyon, France (2001).
- [6] CMS collaboration, *The CMS experiment at the CERN LHC*, [2008 JINST 3 S08004](#).
- [7] P. Aspell et al., *PACE3: a large dynamic range analogue memory ASIC assembly designed for the readout of silicon sensors in the LHC CMS preshower*, in proceedings of the 10th workshop on electronics for LHC experiments (LECC2004), September 13–17, Boston, Massachusetts, U.S.A. (2004).
- [8] S. Roesler et al., *DPMJET3*, online at <http://home.cern.ch/sroe-sler/dpmjet3.html>.
- [9] S. Ostapchenko, *QGSJET-II: towards reliable description of very high energy hadronic interactions*, *Nucl. Phys. Proc. Suppl.* **151** (2006) 143 [[hep-ph/0412332](#)]; *Non-linear effects in high energy hadronic interactions*, [hep-ph/0501093](#).

- [10] K. Kasahara, *The Epics home page*, online at <http://cosmos.n.kanagawa-u.ac.jp>.
- [11] H. Drescher, *Remnant break-up and muon production in cosmic ray air showers*, *Phys. Rev. D* **77** (2008) 056003 [[arXiv:0712.1517](#)].
- [12] *Connector SAMTEC EQCD QTE*, online at http://www.samtec.com/technical_specifications/overview.aspx?series=eqcd.
- [13] G. Magazzú, et al., *The detector control unit: an ASIC for the monitoring of the CMS silicon tracker*, *IEEE Trans. Nucl. Sci.* **51** (2004) 1333.
- [14] G. Cervelli, *CMS tracker LVDS buffer*, http://cmstrackercontrol.web.cern.ch/cmstrackercontrol/documents/LVDSBUF_datasheet_D.pdf.
- [15] *ALBA ELETTRONICA*, online at <http://www.albait.com>, via Leonardo da Vinci 3, Mogliano Veneto, Treviso, Italy.
- [16] *SMTechnology Srl*, Via del Colle 78/A, Calenzano, Firenze, Italy.
- [17] A. Viciani, *Elettronica di lettura per il sistema tracciante di LHCf*, Graduate Thesis, University of Florence, Firenze, Italy (2007), online at <http://hep.fi.infn.it/LHCf/index.html>.
- [18] V. Noce, *Progettazione e realizzazione di una scheda di digitalizzazione dei dati per i sensori al silicio di LHCf*, Graduate Thesis, University of Florence, Florence, Italy (2006), online at <http://hep.fi.infn.it/LHCf/index.html>.
- [19] K. Kloukinas et al., *FEC-CCS: a common front-end controller card for the CMS detector electronics*, in the proceedings of the 12th Workshop on Electronics for LHC and Future Experiments (LECC2006), September 25–29, Valencia, Spain (2006).
- [20] *Digital Optohybrid Pages*, online at <http://cms-tk-opto.web.cern.ch/cms-tk-opto/control/components/doh.html>.
- [21] F. Benotto et al., *Design and test of the digital opto hybrid module for the, CMS tracker inner barrel and disks*, *CERN-CMS-NOTE-2008-013* (2008).
- [22] C. Paillard et al., *The CCU25: a network oriented coommunication and control unit integrated circuit in a 0.25 μ -m CMOS technology*, published in the proceedings of the 8th Workshop on Electronics for LHC experiments, September 9–13, Colmar, France (2002) .
- [23] P. Placidi et al., *CMS tracker PLL reference manual. Version 2.1*, *CERN - EP/MIC* (2000).
- [24] L. Bonechi et al., *LHCf: a calibration tool for cosmic ray physics at LHC*, published in the proceedings of the 29th International Cosmic Ray Conference, August 3–10, Pune, India (2005).