The Compact Muon Solenoid Experiment
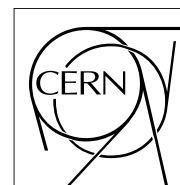
CMS Note

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland

13 October 2008 (v5, 23 April 2009)

Validation tests of the CMS TIB/TID structures

M. D'Alfonso[a], P. Azzurri[a], G. Balestri[a], A. Basti[a], L. Borrello[a], U. Cazzola[a], S. Dutta[a], P.G. Lenzi[b], F. Ligabue[a], S. Linari[a], S. Mersi[b], M. Meschini[b], F. Palla[a], A. Rizzi[a], A. Venturi[a,c], M. Vos[a,c,v]

a) INFN - sezione di Pisa, Italy
b) INFN - sezione di Firenze, Italy

Abstract

Tracker Inner Barrel half-cylinders and Tracker Inner Disks of the CMS tracker have been integrated in three INFN sites. Integrated structures are submitted to an extensive set of tests whose main aim is to validate the functioning of the structures in CMS-like conditions. The tests have furthermore proven to be a great opportunity to study several aspects of the performance in detail. In this note the tests are described in some detail and an overview of the results is presented.

1 Introduction

The Compact Muon Solenoid (CMS [1]) is one of the two large general-purpose experiments at the Large Hadron Collider (LHC). The innermost part of the CMS detector consists of a Tracker for the detection of charged particles. A cross-sectional view of the Tracker subsystems is shown in Figure 1.
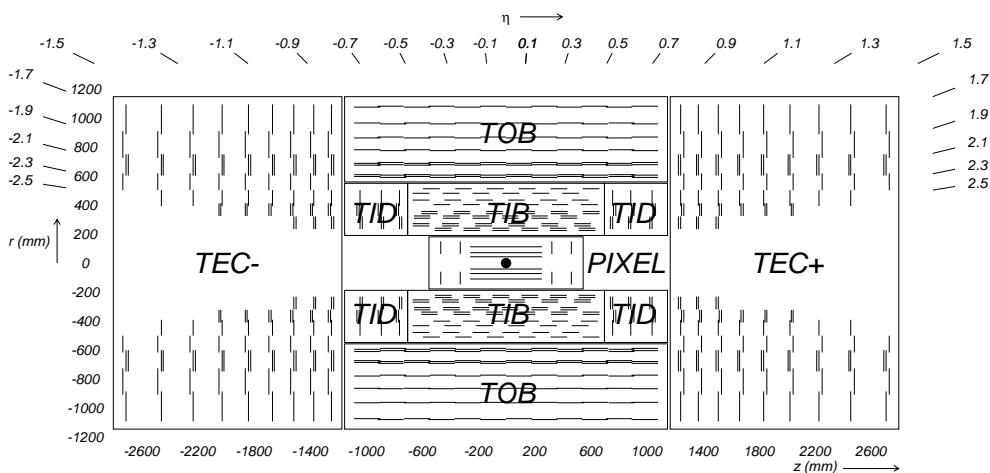

Figure 1: Side view of the CMS Tracker.

The innermost section of the CMS Tracker consists of a pixel detector. Three cylindrical layers are installed at radii of 4.4, 7, and 10.2 cm. The Pixel Barrel (PXB) is complemented by the pixel detector end-cap (PXEC). The remainder of the tracking volume is equipped with silicon micro-strip detectors. The Silicon Strip Tracker [2] consists of four subsystems. The four cylindrical layers of the Tracker Inner Barrel (TIB) extend to a radius of 55 cm. The Tracker Inner Disks (TID) provide coverage at large pseudorapidity ($|\eta| < 2.2$). Along the direction of the beam line (identified with the z-axis) the TIB/TID systems extend to ± 110 cm. The TIB/TID is surrounded by the Tracker Outer Barrel (TOB), which consists of six cylindrical layers. TOB extends the coverage in the central Tracker up to a radius of 116 cm, while spanning a z-range from $-118 < z < 118$ cm. The Tracker End Cap (TEC) completes the forward coverage. The TEC+ and TEC- each consist of 9 disks covering the region $124 < z < 282$ cm and $-282 < z < -124$ cm, respectively. Radially, TEC extends from an inner radius of 22.5 cm to the outer radius of the tracking volume ($r = 113.5$ cm).

Each of the barrel layers in the Silicon Strip Tracker provides a precise measurement of the $r - \phi$ coordinate of charged particles. In the first two layers of TIB and TOB, the first two disks of TID and disks 1,2, and 5 of TEC, a second coordinate (z in cylindrical layers, r on the disks) is measured. The layout and module parameters of the CMS Tracker subsystems are summarized in Table 1.

The building block of the TIB and TID, the module [3], is built around a single low-resistivity p-on-n Hamamatsu micro-strip sensor [4]. The dimensions of the TIB sensors are approximately 61.44×116.89 cm^2 with a thickness of 320 μm. Strips are implanted at a pitch of 80 μm for the two innermost TIB layers and 120 μm for the outermost TIB layers. The strips are oriented parallel to the beam line in the TIB. TID modules are built around a wedge shaped sensor of the same thickness with radial strips, with pitches ranging from 81 μm to 158 μm. The strip implants in the sensors are AC coupled to aluminum strips on the surface of the sensors.

The measurement of a second coordinate in the innermost TIB layers and TID rings is obtained by combining two single-coordinate measurements. These layers are equipped with back-to-back pairs of modules. The orientation of the strips in both modules differs by a stereo-angle of 100 mr.

Each of the readout strips on the sensor is wire-bonded to the Front-End electronics via a pitch adapter that matches the strip pitch to that of the Front-End inputs. The module contains four or six APV25 Front-End ICs (APV) [5]. Each of the APV chips processes the signals from 128 strips. A readout channel consists of a preamplifier and shaper with a 50 ns peaking time and an analog memory with 192 locations to store the samples during the latency of the first level trigger. On receipt of a trigger, three samples from the memory are processed. Two data taking modes are used throughout the validation tests. In *peak* mode the analog output sample corresponds to the peak amplitude. In *deconvolution* mode the APSP deconvolution filter re-filters the three samples with a shorter time

Table 1: Overview of the parameters of the CMS Tracker subsystems.

subsystem	R (cm)	z (cm)	granularity	# modules
PXB	4.4, 7.3, 10.2	$[-26.5, 26.5]$	$100 \times 150\ \mu\text{m}^2$ pixel	672 full 96 half-modules
PXEC	$[6,15]$	$z = \pm\, 34.5, 46.5$	$100 \times 150\ \mu\text{m}^2$ pixel	672
TIB	25.5, 33.9, 41.9, 49.8	$[-73.0\,, 73.0\,]$	80, 120 μm	2724 768 DS + 1188 SS
TID	$[23.5\text{-}33.5],$ $[32.6\text{ - }42.6], [37.7\text{-}50.2]$	$z= \pm\, 78.8, 91.8, 104.7$	81 μm to 158 μm	816 288 DS + 240 SS
TOB	60.8, 69.2, 78.0, 86.8, 96.5, 108.0	$[-118,118]$	122 μm, 183 μm	5208 1080 DS + 3048 SS
TEC	$[22.5,113.5]$	$[132.3, 266.8],$ $[-266.8, -132.3]$	96 - 188 μm	6400 1152 DS + 4096 SS

constant. In this case the output corresponds to the peak amplitude of the APSP filter. Thus, the signal is confined to a single beam crossing interval.

A multiplexer chip (MUX) gathers the information from two adjacent APVs and sends it to the Analog Opto-Hybrid (AOH) [6]. After electrical-to-optical conversion the signal is transmitted to the off-detector electronics over an analog optical link. The power of the laser transmitter can be optimized for each link using the *gain* and *bias* (i.e., offset) settings of the laser driver.

A PLL chip is mounted on each hybrid so that the trigger and clock delays can be optimized on a module by module basis. A Digital Control Unit (DCU) chip is used to monitor the hybrid and sensor temperature and several voltages.

Three to five modules (*strings*) are powered and controlled through a single flexible kapton cable known as the Mother Cable [7]. In the barrel structures the strings are parallel to the beam line, in TID they follow the rings of modules on the disk. Power to several strings can be supplied by a single power supply module; such a group of strings is referred to as a power group.

The cooling circuit has a similar hierarchy. Each string of modules is cooled by an individual aluminium pipe. Several cooling loops are fed in parallel by a cooling manifold installed along the flange of the cylinders (or the rim of the disk in TID). The group of modules cooled through a single manifold is referred to as a cooling loop.

Clock and control signals are distributed through the Digital Opto Hybrid Modules (DOHM [8]). This board converts the optical control signals to electrical signals and fans them out to a Clock and Control Unit (CCU) on each string. The CCU communicates with each of the Front-End ASICs (APV, DCU, PLL, MUX) using the I^2C [9] protocol. To this end the Mother Cable is equipped with one I^2C bus per module. A redundant token ring is implemented on the DOHM, so that the token can bypass a failing CCU. A group of modules controlled by a single DOHM is known as a control ring.

The mechanical structures for each TIB layer are divided in four parts: a forward and backward part (labeled TIB+ and TIB-), each composed of two quarter-cylinders (labeled UP and DOWN). The modules on the $3 + 3$ TID structures are arranged in three rings, each supported by a carbon fiber disk. The complete TIB and TID thus consists of 22 substructures: 16 TIB quarter-cylinders and 6 TID disks.

Between May 2005 and October 2006 the integration of TIB and TID structures [10] took place in three INFN sites: Firenze and Pisa for TIB and Torino for TID. During this phase detector modules were mounted on the carbon mechanical support structure. Each substructure was completed by integrating all auxiliary components (AOH, DOHM, cables, and fibers). The 22 integrated substructures were shipped to the San Piero site in Pisa for the final assembly into the TIB and TID.

Throughout the production chain of the Tracker, quality control of components was of utmost importance. All detector modules were tested thoroughly upon completion [11]. During the integration phase the correct functioning of components was verified after placement on the structure. The final step in the quality assurance of the detector was an exhaustive set of tests of assembled substructures upon arrival to San Piero. The main aim of these tests was to validate the correct functioning of the structures before assembly into TIB and TID. These tests represented the last opportunity to repair failures before shipment to CERN and integration into the experiment.

To validate the correct functioning of Tracker elements it was crucial that the operating conditions resembled

those in the CMS experiment as closely as possible. To this end, a large climate chamber and industrial chiller were installed at the Pisa San Piero site. An effort was furthermore made to ensure that all cables, fibers, offline electronics, and power supplies used in these tests corresponded to the design used in the experiment.

The validation tests have proven to be an important step towards the commissioning of the CMS Tracker. Invaluable experience was gained in the operation of detector elements with several hundred thousand channels. The large statistics allowed to study the uniformity and reliability of the detector components. Many system aspects of the Tracker operation could thus be studied for the first time. The validations tests allowed the verification of the performance of the cooling system in nominal CMS conditions and with a realistic load (power consumption of the modules). Similarly, the channel noise performance of large-scale system under standard operating conditions could be studied in detail. The configuration, control, and readout of a large number of modules was thoroughly exercised.

In this note an overview is given of the results accumulated during the validation tests of TIB and TID substructures at the San Piero site. In section 2 the experimental setup for the validation tests is introduced. The analysis software developed for the validation tests is described briefly in section 3. The different types of commissioning runs are also described in this section. In section 4 the standard tests performed on each substructure are outlined. Results representing the full statistics are presented. In section 5 two special validation runs with cosmic muons on TIB+ and TID+ are described. In section 6 the most important findings are summarized.

2 Setup

For the purpose of the validation tests of integrated structures a large test facility was installed in the clean room at the San Piero site of INFN Pisa. The heart of the setup was formed by a large climate chamber with a single-phase chiller to mimic conditions in the experiment. The setup furthermore consisted of the power supply, interlock, and data acquisition systems for operation of large integrated TIB and TID structures. Each of these subsystems is discussed briefly in the following subsections.

2.1 Climate chamber & liquid cooling system

A large climate chamber was installed in the San Piero clean room. With a floor area of over 10 m^2 and a height of 2.5 m it could host two detector substructures mounted on specially-made mobile support structures. The cables that supplied power to the substructure under test, the cooling pipes, and the control and readout fibers were passed through six holes in the climate chamber, which were then hermetically sealed. The temperature of the chamber was controlled to within one degree Celsius in a range between $+25°$C and $-30°$C. The climate chamber was flushed with dry air allowing a controlled humidity down to a dew point of $-25°$C.

Without cooling, the power dissipated by the front-end electronics would heat up the structure to beyond the temperature that is considered safe for the Analog Opto Hybrids ($40°$C). Active cooling was required for any test that involved powering up more than a single string of modules. Throughout the tests the temperature of the cooling manifold was continuously monitored and interlocked.

A chiller system was implemented for the validation tests. A large industrial chiller provided C_6F_{14} in liquid form at a temperature between $+25$ and $-25°$C and a pressure of up to 8 bar. A manifold outside the climate chamber divided the coolant flow into eight parallel lines, each of which was used to supply one manifold on the detector structure. The TIB manifolds were used to cool a maximum of five strings with six double-sided modules each, or up to eight strings with three single-sided modules. The cooling loops at this stage were not equivalent to the cooling loops in the experiment: as part of the integration at CERN, inner and outer manifolds were joined in the TIB. The same was true for the R1 and R2 manifolds in each TID disk and the R3 manifolds in all disks. With eight lines the structure with the most severe cooling requirements, a TIB layer 2 quarter-cylinder, could be tested at the same time.

To replicate the expected CMS environment as much as possible the following standard parameters were selected:

- a chamber ambient temperature of $-10°$C

- relative humidity controlled at a safe level (approximately 20 %)

- a coolant temperature of $-25°$C at the chiller

- *nominal* coolant flow (see section 4.2 for a complete description)

The input pressure (and thus coolant flow) was chosen in accordance with the hydraulic impedance of each cooling circuit, so as to approach the envisaged operating conditions as closely as possible. Although the cooling pipes were thermally insulated, a slight increase in coolant temperature was observed along the lines from chiller to the detector. The coolant temperature at the manifold on the detector layer varied from $-22°C$ to $-23°C$.

These standard settings did not allow fast access to the structure: the time to re-establish the safe humidity and cool down after the access was several hours. Therefore, in all cases a first set of tests was performed in conditions that allowed an access to be performed in a relatively short time. This second set of parameters was chosen as follows: ambient temperature of $+10°C$, dew point safely (at least $5°C$) below the coolant temperature, coolant temperature of $0°C$, and *nominal* coolant flow.

2.2 Power supplies

The Tracker structures were powered by the same CAEN power supply (PS) system based on the SY1527 controller and 4601 floating PS modules to be used in the experiment [12]. Each PS module supplied two low voltages (2.5 V and 1.25 V) and two separate High Voltage lines (up to 600 V).

Power to the modules was supplied through the same set of cables to be used in the experiment; the long (50 m) shielded LIC60 cables that bridge the distance between the CMS experiment and the power supplies, a lighter CAB60 cable used for the 6 m inside the Tracker, and a third *medusa* cable that connects the *margherita* patch panel to the TIB end flange and the TID disks.

The Digital Opto Hybrid Module and up to 15 CCUs were supplied through a single 2.5 V line. As the PS modules of type 4602 and the final design control cables were not available at the time of these tests, 4601 modules and standard LIC60 cables were used to power the control rings as well.

On the first and second TIB layers a power group corresponds to a string of three double-sided six chip modules powered by the same Mother Cable (i.e., a total of $3 \times 2 \times 6 = 36$ APVs). For the third and fourth TIB layers, which have single-sided modules with four APVs each, a power group is made up of three or four strings (i.e., $4 \times 3 \times 4 = 48$ or $3 \times 3 \times 4 = 36$ APVs). For TID disks a power group corresponds to the string which has either three double-sided modules (36 APVs) or ten single-sided modules (40 APVs).

For strings of double-sided modules the sense lines were connected on the Mother Cable. For single-sided modules, several strings were powered by a single cable. In this case the sense lines were connected on a small board in the medusa cable at a distance of approximately 10-20 cm from the Mother Cable connector.

2.3 Interlock

To guarantee the safety of the structures under test, a scaled-down version of the CMS interlock system was installed. The most important functionality was based on the temperature of the cooling manifolds of the layers under test. The temperature was monitored by reading out a number of pt1000 temperature probes glued on each of the cooling manifolds. Furthermore, four C_6F_{14} detectors were installed in the test chamber. These were connected to the interlock system to alert the operator in case of major coolant leakage inside the chamber.

The interlock was managed through a Graphical User Interface implemented in PVSS [13]. This software allowed continuous monitoring of the manifold temperature. The interlock system was active throughout all tests, i.e., for over one year. Fortunately, the maximum alert situation never occurred.

2.4 Optical links

During the validation tests the analog data from the modules was transmitted to the off-detector electronics over the same analog optical links used in the experiment. The AOH devices, mounted on the structure in the integration phase, are equipped with two or three fibers of 2 m length. Each of these fibers was connected to one of 12 fibers in a 6 meter ribbon. Eight such ribbons were grouped together and connected to a 96 fiber 60 meter long Ericsson multi-ribbon. The ribbons and multi-ribbons were identical to those used in the experiment. For all structures the same set of multi-ribbon cables and ribbons were used.

2.5 Control & Data acquisition

The data acquisition system used for the validation tests was based on hardware and software developed for the DAQ that is used in the CMS experiment.

2.5.1 Hardware

The principal hardware components of the DAQ used for the validation tests were:

- 1 VME9U Front-End Controller (FEC [14])

- 8 VME9U Front-End Driver (FED [15])

- 1 Trigger Sequence Card (TSC [16]) PCI card + VME6U card

- 1 VME6U Timing Trigger and Control ex (TTCex)

- 1 SBS VME/PCI controller

In the CMS Tracker DAQ a Front-End Controller (FEC) is used to configure the Front-End electronics. The FEC controls the settings of the devices on the structure, setting parameters such as the APV operation mode, the PLL delays to the trigger and clock signal, the AOH laser gain and bias. It also reads back the results from the Digital Control Unit used to monitor temperatures and voltages. Each FEC channel communicates over an optical link with a Digital Opto Hybrid Module (DOHM). The DOHM converts the optical commands to electrical and distributes the signal to the Clock and Control Units (CCUs) in the control ring. Each CCU is connected to several modules via an I^2C bus. Layer 1 and Layer 2 control rings contain 3 to 5 CCUs, Layer 3 and 4 contain up to 15 CCUs, TID control rings always have 4 CCUs. Each CCU controls 3 to 6 modules. Therefore, one control ring consists of 18 to 45 modules. The VME9U FEC board provides 8 FEC channels or mFECs.

The analog data from the tracker are processed by the Front-End Driver. The optical signal is electrically converted and digitized with a 10-bit ADC. Two operating modes of the FED were used in this test. In the so called *scope* mode the FED behaves like an oscilloscope and simply samples the digitized optical signal coming from the modules in 25 ns steps. In *physics* mode the FED recognizes the digital header preceding the analog signal. Throughout the validation tests the digitized signal was communicated to the DAQ PC without further processing. Specifically, the FED pedestal and common mode subtraction and zero-suppression capabilities were not used. Each VME9U FED contains 8 FPGAs, each of which processes the data from 12 fibers (i.e., one ribbon). One 9U FED can read up to 96 fibers, each containing multiplexed data from 256 strips.

Trigger signals and the system clock are generated by the TSC card and distributed to FEDs via the TTCex board and to the modules via the FEC. It is possible to feed an external trigger to the TSC. This allows noise studies with a random trigger (the internally generated one is clock based) or particle signal studies with cosmic rays. All communication between the VME crate and the DAQ running on a PC, including the transmission of the FED data, comes about through an SBS controller. A schematic view of the connections is shown in Figure 2.

The DAQ at the San Piero site was designed for the largest substructure to be tested. With a setup consisting of one 9U VME FEC and 8 FEDs, a quarter shell of TIB layer 2 could be read out at the same time. This complex system, with 8 control rings and over 160 000 channels, corresponds to nearly 8 % of the completed TIB/TID.

2.5.2 Software

The DAQ software was developed in the XDAQ environment[17] by the TkDAQ team [18]. The main pieces of this software are (XDAQ applications):

- FecSupervisor: drives the Tracker FEC

- Fed9USupervisor: controls and reads out the FEDs

- TrackerSupervisor: controls the Tracker-specific run parameters, starts and stops runs, etc.

- RootAnalyzer: processes FED data, ROOT I/O.

Before starting data taking for a TIB/TID substructure, a preliminary step was needed to register all the cabling information. Each fiber or cable connection was registered by scanning its 2D barcode. The connection was uploaded to the TIB/TID integration data base via a Java application. The combination of the integration DB information and the cabling map allowed the identication of the power supply and FEC and FED channel that were used to power, control, and read out each module. Before starting the data taking the operator verified that all this
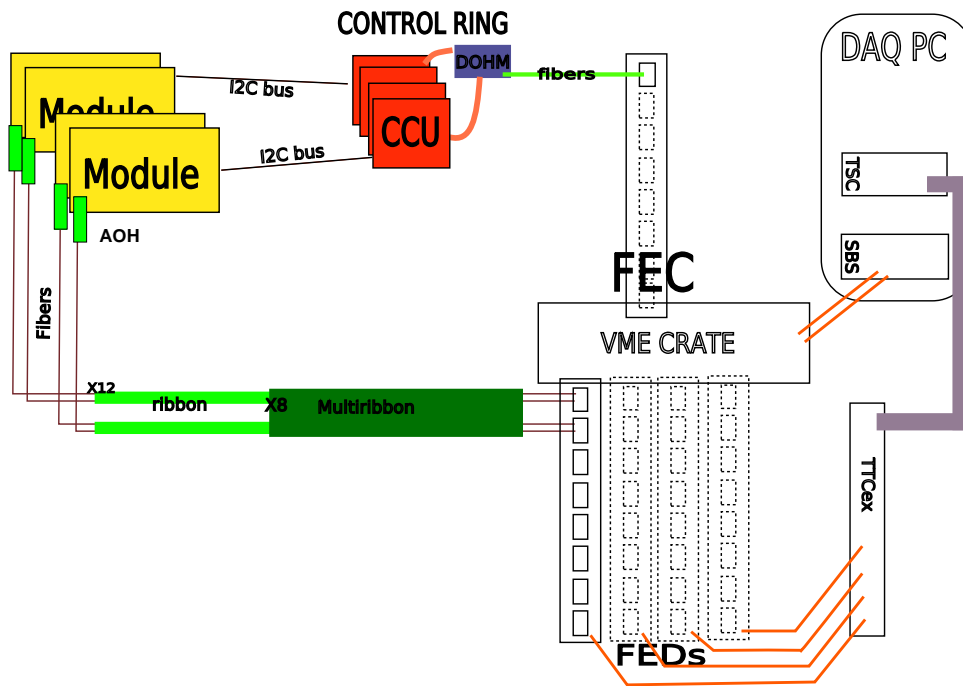
Figure 2: Schematic representation of the control signals and data flow for readout of the CMS Tracker.

information was available (i.e., all connections were registered). If this check was successful a new *Configuration* (i.e., a snapshot of cabling connection) was created and made available for the analysis of the subsequent runs.

The first test on any structure consisted of a scan of the control rings. The FecProfiler utility was used to scan all possible I^2C addresses on all FEC channels that present a *closed* control ring. The list of responding devices was compared to the configuration declared in the integration database. For each device an entry was created in the *fec.xml* file to be used by the DAQ in the remaining tests. Default properties were assigned to each type of device on the basis of a template file. The FecProfiler program was also used to read the temperature, voltage, and current measurement from the Digital Control Units (DCU).

A typical *run* was performed in three steps, which corresponded to three different states of a finite state machine in TrackerSupervisor:

- Configure : parameters such as run type, number of events, and optimal values are set by the operator and sent to the XDAQ applications. In the configuration phase the parameters of all Tracker devices are set by FEC using the *fec.xml* file as source.

- Enable: trigger is activated and data taking starts

- Halt: data taking is stopped and an automatic data processing script is started.

Several different run types were used to test different parts of the Tracker. The different types of runs are explained in some detail in section 3. For each run the RootAnalyzer application in the DAQ software produced a (*TBMonitor*) ROOT file. The RootAnalyzer histograms were used by the analysis programs for several run types. In runs where the analog signal was analysed a second (*sir*) ROOT file was produced that contains the event by event, strip by strip signal. A separate database was set up to contain summary information from each run. The run type and number, an identifier specifying the configuration used, and the climate chamber status were stored in the run DB. Then the run data (xml, ROOT, and log files) were packed in a tar file and moved to an analysis server, which automatically ran the appropriate analysis program for that run type.

3   The Analysis Client

The analysis client was responsible for the analysis of the data taken in different operation modes of the DAQ system. It was used to gather all information to determine the structure performance. It also played an important
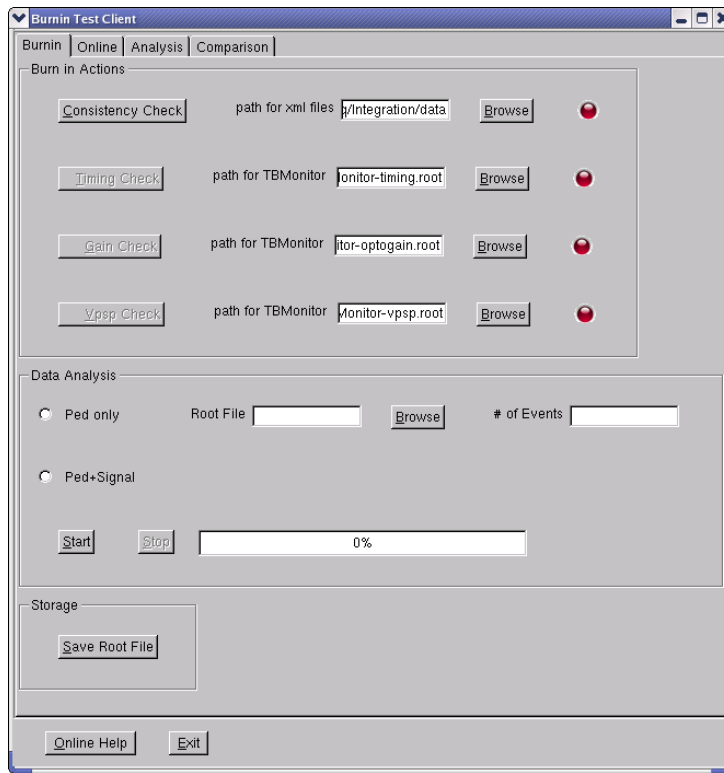
7

Figure 3: Graphical User Interface of the San Piero burn-in Test Client.

role in the monitoring of the data taking procedure. The summary information presented to the operator after each run was very helpful in detecting errors and failures at an early stage.

Important tools of the client included:

- An XML parser based on Xerces [21]

- A multi-threaded analysis environment, which was necessary for the interactive mode of the client

- A tool to access the MySQL [22] run information and integration databases

The analysis client had two running modes: interactive and batch. The interactive mode was based on the Graphical User Interface (GUI) package of ROOT [19]. A screen capture of the GUI is shown in Figure 3. The interactive mode was especially useful at the beginning of the test activity and to debug failures. Various distributions with results are updated in the plotting area of the GUI. During most of the routine testing phase only the batch mode was used to gain speed and achieve optimized bookkeeping.

The resulting distributions were saved in files and a web page was updated accordingly for each run. During the validation tests all results were made available on dynamic web pages. PHP-scripts retrieved the requested information from the integration database, the history of the configuration of the experimental setup, the run database, and the server holding the results of the analyses. These web pages contained quite detailed summary information (a histogram of the channel noise for the complete structure). Very detailed results (the noise on a single APV channel) were also kept. The results were presented in a hierarchical structure that allowed easy navigation from the concise data sheet for a structure to the full detail of the single-module or single-channel result. A typical folder structure as presented by the analysis client is shown in Figure 4.

This representation of the data proved to be an invaluable tool throughout the validation tests. Often, problems could be resolved at an early stage, thus saving an enormous amount of time, thanks to the fast detection and identification of problems.

The standard test procedure consisted of a series of runs: a connection run, a timing scan, an opto-gain scan, a second timing scan to update the FED thresholds, a VPSP scan, and finally the pedestal and noise runs. For each
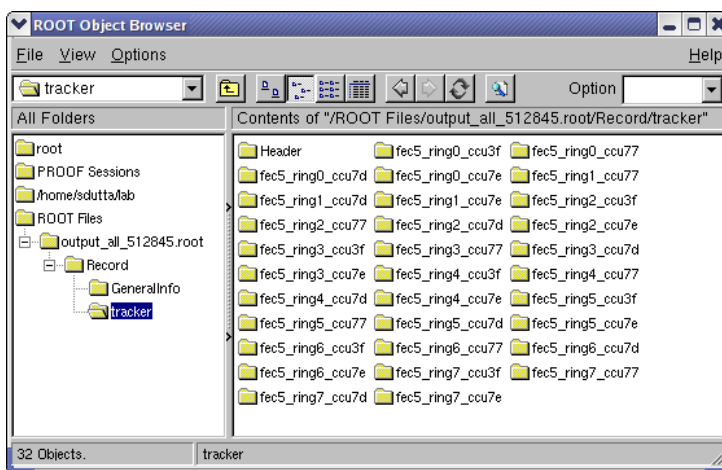
Figure 4: The folder structure that allowed the navigation from the high-level results in the data sheet down to the level of individual modules.

noise run, a ROOT file containing all the relevant results was created. This included the information on the timing, optolink performance, and VPSP obtained in the preceding runs.

3.1 Connection run

After cabling each structure a connection run was performed to create a map of the connections. To this end, the AOH laser drivers were switched on one by one. In each event the FED channel receiving light was registered. In the analysis process this map was compared to the one expected from the integration database (i.e., the *Configuration*). At the end of a connection run the *module.xml* file was created, containing the list of connections of modules (identified by their logical position on the FEC control rings and the APV pair I^2C address) to FED channels. The analysis client used an xml parsing tool to scan the *fec.xml* and *module.xml* files to obtain the FEC, Ring, CCU, I^2C Channel, FED and FED channel, and the DCU numbers. The expected position of the detector was then obtained from the integration database using the DCU number.

A very simple, but highly efficient cross-check was performed at this stage. Whenever the number of modules did not match the number of devices found in the connection run, or the number of fibers per module was different from what was declared in the integration database, the operator was alerted. Thus, connectivity problems were caught immediately. At the offline analysis stage a detailed comparison of the mapping between FEC addresses and FED channels found in the connection run and that expected from the integration database was performed.

After gathering all information the client created a logical structure of folders following the electrical connections of the setup. One folder per detector module was created. The analysis results corresponding to each module were inserted in this folder structure during the analysis step, in the form of histograms or numbers. Finally, the complete structure was written to a ROOT file.

3.2 Timing Run

During the timing run the FED was used in scope mode to reconstruct the *tick mark*: a digital "1" emitted periodically by the APV. The PLL delay was scanned with a 1 ns step to map out the contour of the tick mark. The optimal PLL values to align all tick marks in time were chosen and saved in the *fec.xml* file.

The position in time of the tick mark was analyzed for each detector module and stored. The optimal delay values were chosen by the XDAQ RootAnalyzer and written to the *fec.xml* file to be used in the subsequent runs. Two summary plots were created: (i) the positions of the tick marks in time from all the fibers and (ii) the fine tuned delays of all detector modules versus logical position in the control ring. A typical example is presented in Figure 5.

Any problems in the time alignment of individual modules immediately would show up in the uppermost figure as a horizontally displaced tick mark. Fibers with very low light yield could be identified due to the reduced tick mark height. A mouse click on a given tick mark in the GUI window showed the address of the module that produced the tick mark. With this information the operator could decide whether an access to the setup was required (i.e., to clean the fiber connection). The lowermost histogram presents the optimal PLL delays chosen by the XDAQ
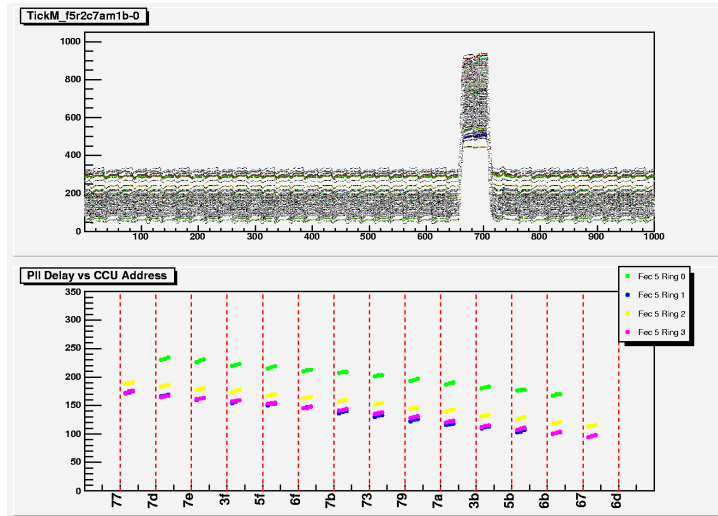
Figure 5: The Graphical User Interface of the analysis client for a timing run. In the upper figure, the reconstructed tick marks of all fibers are overlaid in a single figure. In the lower figure, the delay setting, chosen on the basis of a fit to the tick mark, is plotted versus the module's logical position in the control ring.

algorithm versus the logical position of the module in the control ring. The linear dependence of the delays with CCU address reflects the delay accumulated as the signals are passed along the token ring. Each of the modules controlled by one CCU is drawn as a separate marker in the same column. The modules controlled by a single CCU have the same delay to within 1 ns. Any deviation from the nominal behavior was taken to indicate a problem in the synchronization or a difference in the control ring from what was declared in the integration database.

The timing run was also used to determine the FED thresholds for the header finding algorithm. The tick mark height was measured and the FED zero-one thresholds for each channel were updated accordingly. Therefore, the timing run was repeated after the optimization of the optical gain described in the following section.

3.3 Gain Run

During a gain scan run the four gain settings and all possible bias settings of the laser transmitters (AOH) were tested. For each combination the digital baseline and tick mark height were measured and stored in a number of histograms. The optimal gain setting was chosen such that the height of the tick mark was closest to 640 ADC counts. The bias setting was chosen so as to center the tick mark in the dynamic range.

After each XDAQ opto-gain scan the analysis client GUI presented monitoring information, an example of which is shown in Figure 6. To quickly identify deviations from the optimal performance three quality markers were determined for each link :

- the slope of baseline versus laser bias current

- the final gain setting

- the tick mark height for optimized settings

An optical link was considered to function correctly if all of the following three conditions were met: the slope of baseline versus bias current should be greater than 40 counts/step, the algorithm should not choose the highest gain setting, and the resulting tick mark height should be greater than 400 counts. All three markers are sensitive to links with a low light yield. For links that failed any of the above criteria a flag was raised. The address of the link in question was displayed in the on-line GUI or on the run web page together with the geographical location of the fiber extracted from the integration database. On the basis of this information, the operator decided whether an intervention was required.

Thanks to the careful handling of the fibers throughout the integration, operator intervention was relatively rare: typically only one or two opto-links per substructure failed one or more of the above criteria. The cleaning procedure cured the problem in about 90 % of the cases. In the remaining 10 % of cases, the fiber was investigated
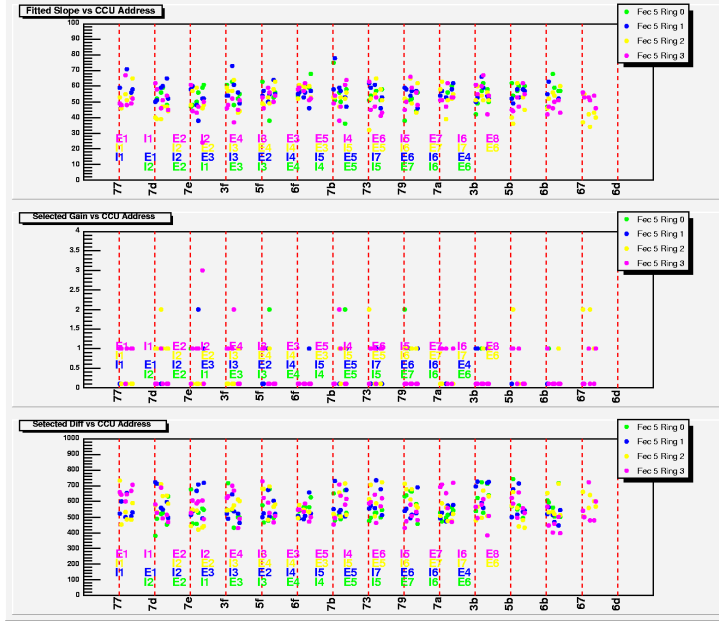
10

Figure 6: The Graphical User Interface of the analysis client for a gain scan. In the upper histogram, the slope of the digital baseline dependence on the laser driver bias current is shown. In the central figure, the chosen gain setting is presented. In the lower figure the tick mark height for the chosen optimal settings is shown. This last marker is directly proportional to the optical gain of the link.

using a different setup. In all but one case (out of 8800 fibers) the problem resided in the ribbon or multi-ribbon, while the fiber on the structure was fully functional. In only one case a fiber was found to be damaged so severely that the AOH had to be replaced.

3.4 VPSP Run

In the VPSP run the APV register that controls the height of the analog baseline was scanned through all possible values. The value was chosen so as to set the analog baseline at one third of distance between digital baseline and digital "1". The analysis client required the VPSP register setting to be in the range from 30 to 45.

3.5 Pedestal Run

In the noise or pedestal runs the internal trigger of the TSC was used to take data. The raw analog response was processed by the analysis client to evaluate the noise performance of the detector modules. Pedestal and Noise of individual strips and the Common Mode correction of a group of strips were determined for each detector module present in the system. The different steps of the analysis algorithm are described below.

3.5.1 Initial Pedestal and Noise

The N_{init} events taken were used solely to determine initial pedestal and noise values for the individual strips. Throughout the validation tests, a value of $N_{init} = 500$ events was used.

For a given strip the initial pedestal was defined as the average of the ADC counts:

$$\text{Pedestal}_{init} = \sum_{i=1}^{N_{init}} (\text{ADC}_i)/N_{init}$$

and the noise was defined as the variance of the ADC counts

$$(\text{Noise}_{init})^2 = < (\text{ADC}_i)^2 > - < \text{ADC}_i >^2$$

In order to avoid contamination from true signals (e.g., from cosmic rays), the average and RMS were calculated after truncation. That is, the N_{init} events in each strip were sorted in ADC counts and the 5 % of events with the smallest and largest ADC counts were removed before doing the average and RMS determinations.

3.5.2 Updated Pedestal and Raw Noise

For the remaining events the pedestal and the raw noise were calculated using groups of N_{iter} events. Throughout the validation tests N_{iter} was taken to be 200 events, while the total number of events in each noise run was at least 2000. With this frequent recalculation of pedestals, drift of the pedestals due to temperature changes was taken into account.

$\text{Pedestal} = \sum_{i=1}^{N_{iter}} (ADC_i)/N_{iter}$

$(\text{RawNoise})^2 = <(ADC_i)^2> - <ADC_i>^2$

During this calculation possible signal was rejected using a cut

$(ADC - \text{Pedestal}) > n_{signal} \times \text{RawNoise}$

where $n_{signal} = 3$ and Pedestal and RawNoise were taken from the previous iteration.

3.5.3 Common Mode

The common mode noise correction was calculated for a group of N_{strip} consecutive strips on the same chip. In the validation tests, as in the algorithms envisaged for the FEDs in the experiment, the group of strips was was taken to be the full APV (128 channels).

The Analysis Client calculated the Common Mode correction as the weighted average over a group of strips, where the weight was defined on the basis of the noise of the strip under consideration :

$\text{CommonMode} = (\sum_{j=1}^{N_{strip}} (ADC_j - \text{Pedestal}_j)/\text{Noise}_j^2)/(\sum_{j=1}^{N_{strip}} 1/\text{Noise}_j^2)$

Contamination from possible real signals was avoided in this calculation by requiring:

$(ADC - \text{Pedestal}) > n_{signal} \times \text{Noise}$

3.5.4 Common Mode Subtracted Noise

The Common Mode subtracted noise Noise_{CMS} was defined by the variance of the quantity ADC - Pedestal - Common Mode. Unless stated otherwise, all noise results presented in this note, refer to the common-mode subtracted noise.

3.5.5 Noisy and Dead Strips

The faulty strips were classified as noisy or dead accordingly to their noise behavior. A strip was considered to be 'dead' when its noise was smaller than a given fraction of the average noise:

$\text{Noise}_{CMS} < F_{dead} \times \langle \text{Noise}_{CMS} \rangle$

where $\langle \text{Noise}_{CMS} \rangle$ is the average noise over all strips connected to a given readout chip. The cut on the fraction F_{dead} of the average noise was taken to be 75 %.

Noisy strips were identified in the analysis client using an RMS cut : $\text{Noise}_{CMS} > \langle \text{Noise}_{CMS} \rangle + F_{noisy} \times \sigma_{noise}$, where σ_{noise} is the RMS spread of the noise over the strips in a given readout chip. For the validation tests a five sigma cut was used ($F_{noisy} = 5$).

The strips identified as noisy or dead according to the above definitions were masked. The strip mask was defined dynamically, i.e., at each iteration of the analysis the mask was recalculated. The final noisy and dead strip list contains those channels that were masked in more than 50 % of the iterations.

3.5.6 Cluster Finding

Clusters were searched in the pedestal and common mode subtracted data. A cluster was reconstructed using a three threshold algorithm. Any strip with a Signal-to-Noise ratio greater than the *seed threshold* was used to seed cluster finding. For each seed, the neighboring strips were searched and strips with a Signal-to-Noise ratio greater than the *neighbor threshold* were added to the cluster. There was no explicit maximum value for the number of strips allowed in a cluster. Consecutive strips with signal above threshold were added to the cluster until the algorithm encountered a signal below threshold. Only if a single strip with signal below threshold was followed by a strip with signal above threshold the cluster finding continued. One such *hole* was allowed per cluster. Finally,

an overall *cluster threshold* was used to validate the total cluster. Typical values for the three thresholds were 3σ, 2σ, and 5σ for *seed*, *neighbor* , and *cluster* thresholds, respectively.

4   Results

In this section, an overview is given of the most important results from over year and a half of intensive testing. Many different aspects of the performance were studied throughout the validation tests of 22 separate substructures, with over 3500 modules, 8000 fibers, and over 2 million readout channels. In separate subsections an overview is given of the most important lessons that were learned on the power consumption and cooling performance, on the long term operation experience, and on the exercise of the control of the structures. Finally, results are presented from a study of the optical link performance and the noise behavior.

4.1   Power consumption

The voltage at each of the Power Supply connectors, the sensed voltage and the current were monitored and logged throughout the validation tests.
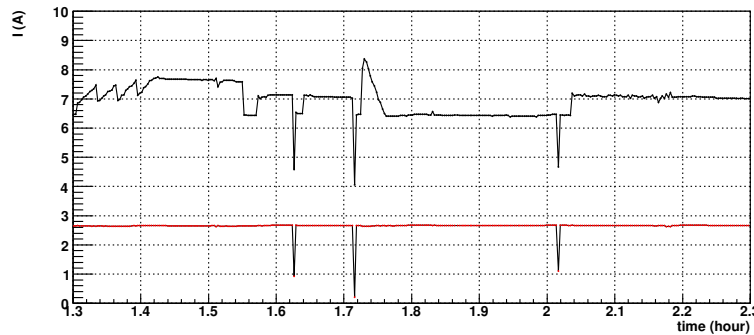

Figure 7: Current of a Layer 3 power group with 12 modules versus time during the configuration phase. The upper points represent the current on the digital (2.5 V) line, the lower points that on the analog (1.25 V) line. See text for an explanation of the features observed in the current.

The measured currents versus time throughout several commissioning runs for a power group of TIB layer 3 are shown in Figure 7. The characteristic features of the different types of runs clearly stand out. Hard resets of the Front-End are visible as sudden dips in the current at $t = 1.62$, 1.72, and 2.02. In the initial phase, between $t = 1.3$ and $t = 1.4$, a gain scan, with the gradual increase in the opto-hybrid bias current, can be observed. The period from $t = 1.63$ to $t = 1.7$ corresponds to a timing run. The most significant current variation occurs at $t = 1.73$ hours during the scan of the VPSP setting on the APV that regulates the baseline of the analog output. As the VPSP setting was increased from 1 to 50, the baseline decreased by over 800 mV, leading to a change in the current consumption of 2 A or nearly 30 %. During the noise run ($t > 2.2$) the current was stable to within 40 mA on the 2.5 V line, and 10 mA on the 1.25 V line.

In the validation tests, the same types of cables were used to supply power to the modules as in the CMS experiment, including 50 m of LIC60 cable, 6 m of CAB60 cable, and approximately 1 m of *medusa* cable. On the TIB and TID structures the power is distributed by the Mother Cables. The voltage drop over the supply lines is corrected for by a sense circuit. On TIB layers 1 and 2 the sense wires are connected on the Mother Cable, while on layers 3 and 4 the sense connection is made on the *medusa* cable, approximately 10 cm before the Mother Cable connection.

The resistance for the full length of a supply cable was measured to be 0.36 ± 0.04 Ω on the 2.5 V line and 0.77 ± 0.10 Ω on the 1.25 V line. At nominal load this implies that the voltage on the power supply connector needs to exceed the required module voltage by several volts. Measurements of the voltages close to the sense connection confirmed that the voltage drop on the cables is effectively compensated. There is, however, an additional, uncompensated voltage drop on the Mother Cable, the MC-hybrid connector and module kapton tail. For the 2.5 V line and nominal load this additional voltage drop amounts to approximately 50 mV on single-sided modules and up to 150 mV on double-sided modules.

The differences between voltages on the three module positions along the Mother Cable were found to be very

small, of the order of 20 mV. It is therefore believed that the voltage drop on the 2.5 V line in the MC is negligible. The difference between single-sided and double-sided modules, on the other hand, is quite large. The current through the kapton pigtail between module and MC increases by 50 % between 6 chip modules and 4 chip modules, while the total current flowing in the MC increases by a factor three (36 versus 12 APVs). These observations indicate that the main voltage drop occurred at the end of the Mother Cable. Compensation of this voltage drop in the case of varying power consumption is clearly non-trivial. As the effect of the undervoltage seemed to be small, no attempt was made to correct for it [1].

The High Voltage was set to 400 V on both lines of the power group. For all structures under test the total leakage current on each line at nominal CMS temperature was too low to be measured using the CAEN monitoring. The same was true for the leakage current readout of the Digital Control Unit (DCU) chip on each module. Even for measurements in a relatively warm environment ($10°C$ air temperature, $0°C$ coolant) the resolution of $0.2\ \mu A$ (for the CAEN monitoring) or $0.7\ \mu A$ (for the DCU) was found to be too coarse to yield a meaningful measurement. Therefore, in the validation tests, the only requirement on the leakage current was that it should be less than 10 μA. All of the structures met this requirement.

4.2 Cooling tests

For all tests described in this note, the structures were cooled by flowing C_6F_{14} in liquid form in the manifolds. This was the first time that the TIB and TID cooling system was operated under CMS-like conditions.

To mimic CMS conditions as faithfully as possible, the coolant temperature was set to the value specified for the experiment ($-25°C$). The thermal load presented by the structure depends strongly on the configuration of the Front-End electronics. Therefore, the full system was configured in the operational mode envisaged for CMS physics runs. In particular, the APV Front-End settings were adjusted for each temperature setting following the prescriptions in reference [23]. The coolant flow rate was set in agreement with the power dissipated by the electronics associated to each cooling loop. In Table 2 the coolant flow for the different substructures is indicated.

Table 2: Coolant flow in each cooling manifold during the validation tests.

structure	module type	coolant flow (l / min)
TIB	double-sided	4.0
TIB	single-sided	3.2
TID	double-sided	1.4
TID	single-sided	0.9

The ambient temperature was actively controlled to the temperature expected in the experiment ($-10°C$). Heat exchange with the cold air was found to contribute significantly to the cooling of the sensors.

The temperatures were measured using the two systems present on TIB and TID structures. The two pt1000 sensors mounted on all cooling manifolds were read out by the interlock system. Much more detailed results were obtained by reading out the the sensors installed on the sensor and hybrid of each module. The DCU chip on the hybrid converts the analog level measured on both into digital readings. These two measurements, plus a third one that is internal to the DCU, are made available through the slow-control protocol. A global calibration of the thermistor measurements was used, thus ignoring the individual offset and gain of each DCU. Nevertheless, a resolution of approximately $1.5°C$ was obtained [2]. The spread in the internal DCU temperature measurement was much larger and this DCU channel was not used for absolute measurements.

The distribution of hybrid and sensor temperatures in nominal conditions are shown in Figures 8 and 9. The mean values and standard deviations of the distributions of Figures 8 and 9 are listed in Table 3. Generally, the cooling performance of TIB/TID structures as determined in these tests was found to be very satisfactory. There is one caveat: after many years of operation at the LHC the sensors on the inner layers (rings) of the TIB (TID) will have been heavily irradiated and have leakage currents well above the typical validation currents of hundreds of nanoamperes per module. The validation tests performed on non-irradiated sensors therefore do not measure the performance of the cooling system under the full heat load.

As expected, the average hybrid temperatures for double-sided modules on TIB and TID, with six APV chips per

---

[1] The one exception was the control ring test, detailed in section 4.4

[2] The calibration of the temperature probes available at the time of the validation tests has since been replaced. The current baseline calibration differs from the one used here.
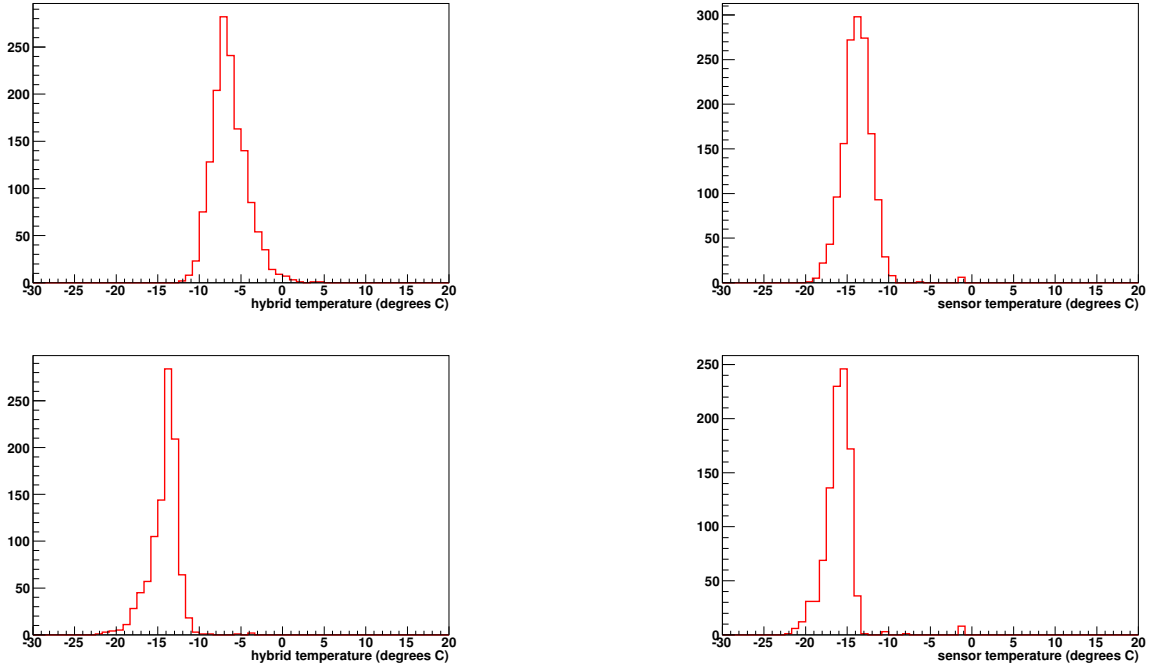
Figure 8: Hybrid (leftmost figures) and sensor (rightmost figures) temperature distribution for all 1476 half-modules on double-sided TIB layer 1 and 2 (upper figures), and for 986 modules on single-sided TIB layer 3 and 4 (lower figures). These measurements were performed before mounting of a final string in the TIB layers. The cooling performance of TIB+ layer 4 down was not tested under the same conditions and is not included in the figures. The temperature measurements were taken under conditions very similar to those expected in CMS: the chamber temperature was set to $-10°$C, and the nominal flow of coolant at $-25°$C was circulated in the manifolds.
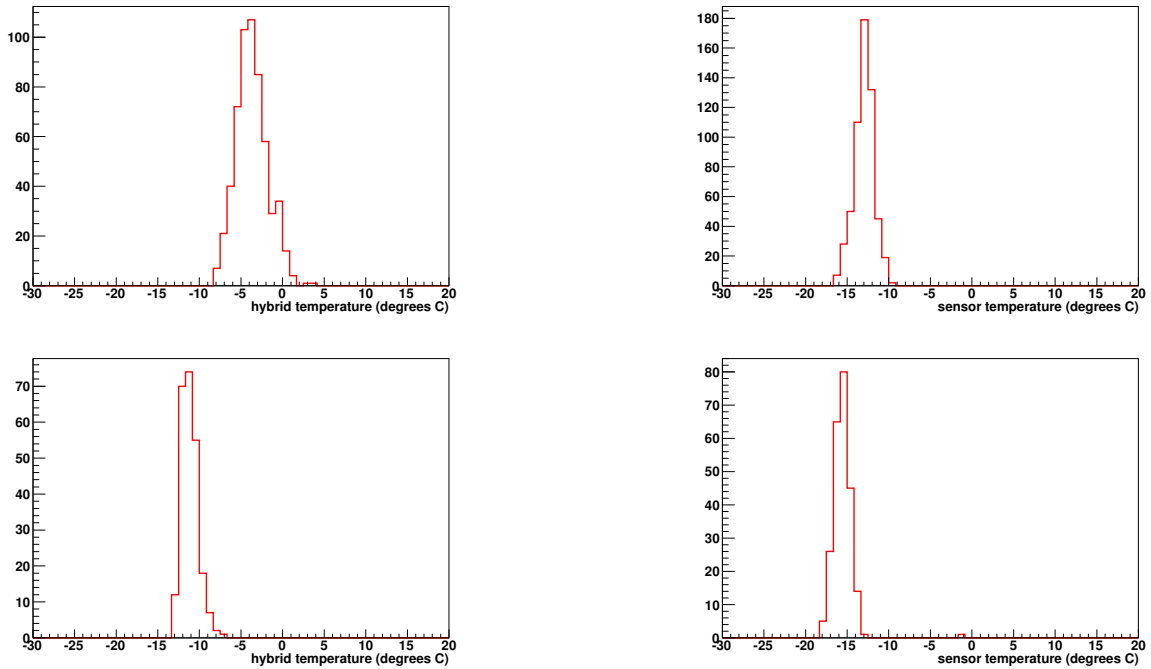

Figure 9: Hybrid (leftmost figures) and sensor (rightmost figures) temperature distribution for all 576 half-modules on double-sided rings 1 and 2 on the TID disks (upper figures) and 239 modules on single-sided TID ring 3 (lower figures). The temperature measurements were taken under conditions very similar to those expected in CMS: the chamber temperature was set to $-10°$C, and the nominal flow of coolant at $-25°$C was circulated in the manifolds.

15

Table 3: Mean and standard deviation of the temperature distributions measured on the hybrid and sensor thermistors of the various substructures.

module type		hybrid temperature		sensor temperature	
structure	*type*	*mean (°C)*	*RMS (°C)*	*mean (°C)*	*RMS (°C)*
TIB	double-sided	− 6.5	2.1	− 13.7	1.8
TIB	single-sided	− 14.2	1.8	− 16.0	2.0
TID	double-sided	− 3.7	1.9	− 13.0	1.2
TID	single-sided	− 11.2	1.0	− 15.6	1.3

hybrid and two hybrids mounted back-to-back, are considerably higher ($-6.5°$C and $-3.7°$C, on TIB and TID respectively) than single-sided modules of TIB and TID geometry ($-14.2°$C and $-11.2°$C, respectively).

All hybrid temperatures are contained in a narrow, nearly Gaussian distribution. The RMS of all distributions is less than $2°$C. This RMS contains contributions from the individual DCU calibrations and coolant and climate chamber temperature variations over the multiple measurements. The temperature variations across the structure were much smaller.

In the validation test conditions, the sensor leakage current is a negligible heat source. However, the relatively high ambient temperature ($-10°$C) compared with the coolant temperature ($-25°$C) causes the sensors to be warmed by the environment. Another potential heat source is the nearby hybrid. The resulting average sensor temperature for single-sided modules was $-16.0°$C for TIB modules and $-15.6°$C for TID modules. For double-sided modules the sensors were only slightly warmer: $-13.7°$C for TIB and $-13.0°$C for TID. The low sensor temperatures (and the small difference between double- and single-sided modules) show that the module design was quite successful in decoupling the main heat source (the hybrid) from the sensors.

The temperature distributions for the sensors were similarly found to be quite narrow: the RMS was always less than $2°$C. Again, the measurements were taken at different times and the actual spread in TID hybrid temperatures at any given time is much smaller.

For all structures a second temperature measurement was made under nominal conditions, with the exception of the coolant temperature, which was set to $-10°$C. Taking the difference between both temperature measurements a number of uncertainties are significantly reduced. First of all, offset variations between different DCUs cancel. Secondly, changing the coolant temperature, while leaving the air temperature unchanged isolates the effect of the coolant. Thus, the ΔT provides a sensitive measure of the thermal coupling between the sensor or hybrid and the cooling pipes. In Figure 10 the result for the single-sided TIB modules is shown.
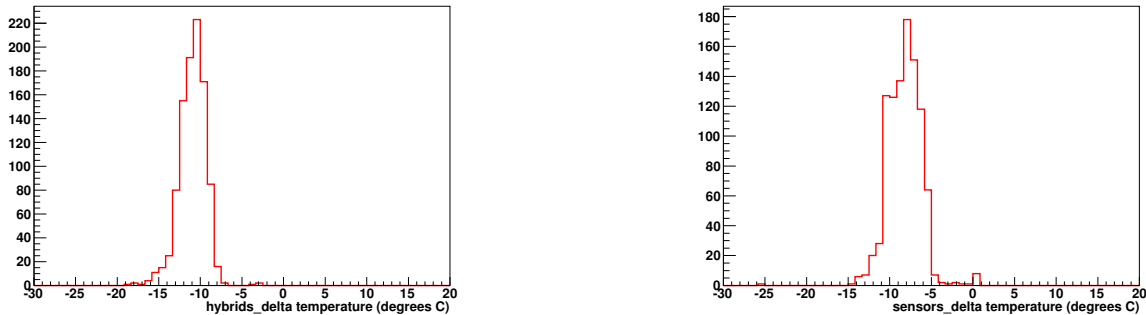

Figure 10: Hybrid (leftmost figure) and sensor (rightmost figure) ΔT distribution for all modules on single-sided TIB layers 3 and 4. The figures show the temperature difference between two measurements for coolant temperatures of -10 and $-25°$C. The coolant flow and chamber temperature were maintained at their nominal values, i.e. similar to those expected in CMS.

Throughout the tests of TIB/TID structures, maps of absolute temperature and ΔT versus geographical position on the structure were found to be very efficient in quickly identifying modules or cooling loops with reduced cooling performance. In two cases, cooling loops with severely reduced perfromance were found. For TID disk 1, the entire ring was rejected and the manifold replaced. A string consisting of three single-sided modules in TIB layer 3 was also found to have a much reduced coolant capacity. Simulations showed that the sensors will operate at an acceptable temperature provided the ambient temperature stays close to the expected $-10°$C. Therefore, in this case the risks involved in a repair were judged to outweigh the potential gain.

Several cases of single modules or half-modules with reduced thermal coupling to the cooling system were found. In the TID disks, the glue contact between hybrids of many double-sided modules was improved. In TIB double-sided layers, a number of modules with increased temperatures were found. At closer inspection this problem turned out to be due to imperfections in the mechanical mounting. For several modules the cooling performance was restored by tightening the screws that hold the hybrid on the cooling ledge. In other cases there was an air gap between hybrid and ledge due to glue on the module precision pin.

4.3   Long term stability

During the validation tests all TIB/TID structures were fully powered and clocked for at least two 12 hour periods, one of these with the structure at low temperature. For a large fraction of this time, the system was read out as well. Several TIB structures were tested for much longer times and underwent a large number of thermal cycles. Thus, the validation tests included a burn-in of the components.

Only a single case of component failure occurred during the validation tests: one AOH stopped responding after 10 hours of data-taking. Optical inspection of the device revealed no underlying reason for the cause of this failure. One Digital Opto Hybrid was found to be broken immediately after the test session. Inspection revealed a mechanical detachment of one of the signal fibers. As this failure is probably due to mechanical stress/mishandling, this instance was not counted under infant mortality.

High and Low Voltage trips of single CAEN PS channels occurred occasionally. The rate of such problems was very low and they were always limited to a single channel. In several cases the problem was traced to faulty LIC60-CAB60 connections [3]. In one case a double-sided string developed a short on the Mother Cable during the first temperature ramp-down. Inspection of the Mother Cable revealed a contamination of the MC surface with traces of conducting material, probably solder. Replacement of the Mother Cable effectively eliminated the problem.

In another case the DOHM low voltage power supply tripped after 10 hours of stable operation. It was found that due to the reduced space in the dense layer 1, one of the cards holding the CCU device could make electrical contact with the cooling circuit. It was decided to provide additional isolation for all CCUs in this and a few other positions.

The validation tests provided important operational experience with the readout of TIB/TID structures using the CMS DAQ system. In one and a half year of daily data taking the system proved to be quite robust. The DAQ system was active for over 3000 hours and took data during about half of this time. The off-detector electronics (FEC, FED, TSC, power supplies) demonstrated similar reliability.

4.4   Control ring tests

Slow control in the CMS Tracker involves many components: the Front-End Contoller (FEC), the Digital Opto Hybrid Module (DOHM), the Clock and Control Units (CCUs), and finally the Front-End devices (APV, AOH, PLL, DCU) on the modules. The communication along this chain was tested extensively. Several standalone programs and scripts were used to validate the correct functioning of the token ring, fully testing the redundancy of the chain.

FecProfiler, a standalone program, was used to scan the entire control ring for CCUs and I^2C devices. The resulting list of devices was compared to the integration database. For all discovered devices the communication was tested by writing a pre-defined set of values into the registers and reading back the result. Generally the I^2C communication proved very robust, but a number of isolated problems were encountered.

On several occasions, defective control rings were found on TIB structures. These were invariably due to faulty connections between the flat cable connecting the DOHM to the end of the Mother Cable or to badly seated CCUs. On three occasions, entire modules were found missing due to faulty insertion of the Kapton tail that connects the module to the Mother Cable. Simply reconnecting the cable or the CCU was sufficient in these cases to ensure correct functioning. In only two cases the components, a CCU and a Mother Cable, were replaced. No connectivity problems occurred on TID disks, where the Mother Cable is held in place on the carbon fiber surface of the disk.

A special effort was made to investigate the robustness of the TIB and TID against a specific I^2C error (code "230"). This failure was investigated extensively by the TOB community [24]. In the I^2C protocol a "low" data

---

[3] For these tests this connection relied on the friction in the connector as the final clip that should hold the connectors in place was not yet available.

line in combination with a "high" clock line indicates the "end of transmission". Therefore, during communication devices hold the data line "high" for a certain time after the clock has been sent "low". In some cases, this time is not sufficient and the data line "takes over" the clock line while descending. The device then interprets this event as "end of transmission" and the "230" error results.

For TIB structures, these errors were found to occur only in very specific conditions. The failure affected a limited set of modules, all of them from a specific 4-chip TIB hybrid batch [4]. Moreover, these modules functioned correctly at room temperature. For some modules the onset of the problem occured at a coolant temperature of $-10°C$, but most modules worked correctly down to coolant temperatures of $-20°C$. For a large majority of failures, modules were mounted in the "0x1a" position, i.e., the Mother Cable position farthest from the CCU. Finally, the failure was found to depend strongly on the voltage levels on the module. Due to the distance (and connectors) between the connection of the sense lines and the module, the actual voltage on the hybrid was lower than the nominal voltage required from the power supply: numbers for this uncompensated voltage drop are given in section 4.1. The failure was efficiently remedied by slightly raising the supply voltage on the 2.5 V line. All modules tested were found to function correctly with a set voltage of 2.6 V, except for one that required 2.65 V. Under normal load conditions these settings correspond to approximately nominal voltage on the hybrid LV capacitors [25].

Modules that showed a failure in any of the configuration tests were routinely replaced by modules from hybrid batches not showing the I^2C weakness. A total of 31 modules was thus replaced. A break-down of this number among the different layers is presented in Table 4.4. As there were sufficient spares to complete TIB integration, the replaced modules were not redeployed on the TIB.

Table 4: The number of replaced modules on the TIB structures.

structure	TIB+				TIB-			
quarter-layer	L4 up	L4 down	L3 up	L3 down	L4 up	L4 down	L3 up	L3 down
number	4	4	3	5	2	4	6	3

To summarize, the validation tests included an extensive exercise of the full control chain. The global performance was found to be robust. Rare failures, corresponding to a specific hybrid batch, were studied in great detail. While not fully understood, an effective remedy was found.

4.5   Analog baseline setting VPSP

The analog baseline of the APV output signal is controlled by the VPSP setting. Throughout the validation tests, VPSP runs were performed on all structures under nominal CMS conditions (i.e., a coolant temperature of $-25°C$ and an ambient temperature of $-10°C$). The distribution of the optimal VPSP settings is shown in Figure 11.

The distribution corresponds to a subset of modules (a total of 9428 APV chips) for which this information is available. While these modules included TIB and TID modules located on different TIB layers, the optimal VPSP setting was found to be very uniform. The mean of a Gaussian fit to the distribution obtained is 35.96 ± 0.02, the standard deviation is 1.67 ± 0.02.

In the analysis client a nominal VPSP range from register settings of 30 to 45 was defined. The software flagged any APV for which the optimal VPSP value was outside this range. No anomalous VPSP values were found throughout the validation tests.

Figure 11: Distribution of the optimal VPSP register setting under nominal CMS conditions.

The VPSP optimization was repeated under warm conditions, with a coolant temperature of $0°C$ and a chamber temperature of $10°C$. The APV parameters were adjusted for each temperature setting. The distribution of optimum VPSP values varies only slightly with temperature: the warmer settings yield a mean of 35.19 ± 0.02 and a standard deviation of 1.7 ± 0.02.
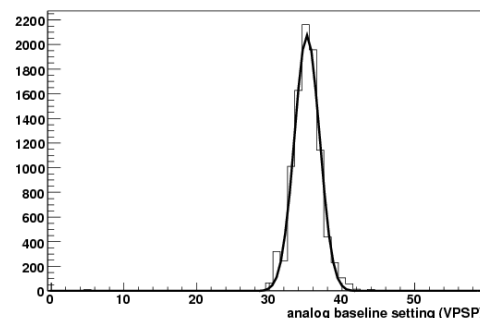
---

[4] Hybrid batch number 031215_TIB4PD produced on December 15, 2003.

4.6 Optical link performance

The laser driver on the Analog Opto Hybrid has four different gain settings, where the highest gain is approximately 2.5 times larger than the smallest. A 50 step register allows the adjustment of the laser bias current.

The tick mark produced by the APV has a height of 800 mV. The FED Analog to Digital Converter is calibrated to yield full scale (1024 counts) for a 1.024 V input signal. Therefore, the desired optical link gain of 0.8 corresponds to a tick height of 640 counts.
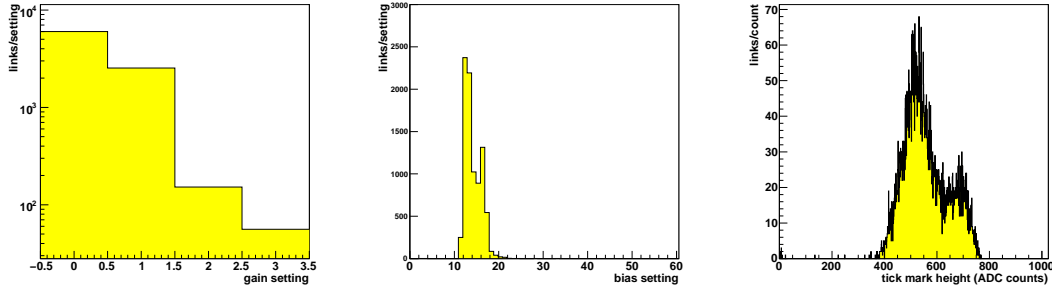

Figure 12: Distribution of the laser driver gain (leftmost figure) and bias current (central figure) settings chosen to reach nominal optical gain. In the rightmost figure the distribution of the measured tick mark height after choosing the optimum gain and bias setting for each link is shown.

To optimize the optical link gain and baseline the gain and bias current were scanned. For each combination of values the digital baseline and the height of the tick mark were measured. A combination of parameters was then chosen to provide an optical link gain as close as possible to the nominal gain of 0.8, while centering the tick mark in the dynamic range of the FED.
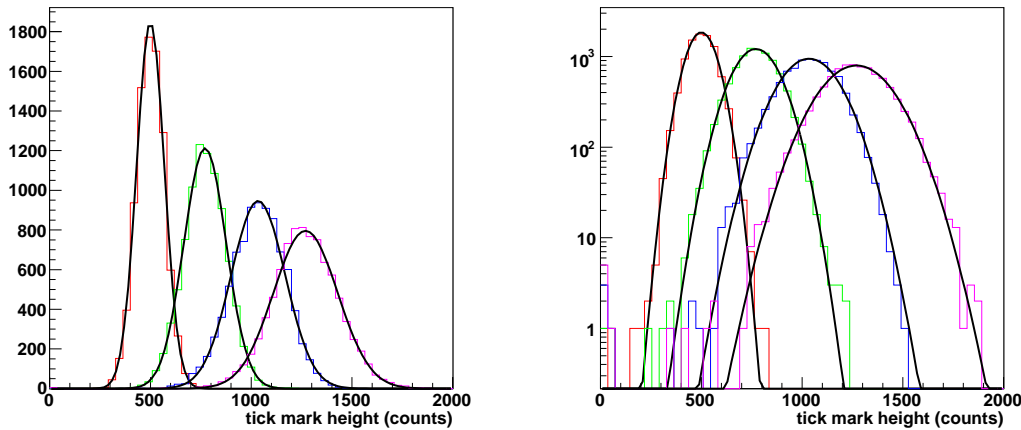

Figure 13: Distribution of the tick mark height for the optical links in the TIB/TID. The four distributions from left to right correspond to gain 0,1,2, and 3 of the AOH laser driver. The laser bias was optimized for each gain setting. Measurements were obtained under nominal CMS conditions ($-25°$C on the chiller, $-10°$C ambient temperature). Tick marks beyond the dynamical range of the FED are reconstructed using the algorithm described in the text. The rightmost figure represents the same data in logarithmic scale.

.

The optical link performance and the result of the optimization procedure were checked carefully, as explained in section 3.3. In case of less than nominal performance, operator intervention (cleaning of the several optical interfaces) generally brought the link performance within the nominal band.

The results from opto-gain optimization are shown in Figure 12. The three histograms represent the chosen gain setting (leftmost histogram), the distribution of the optimal bias current (central histogram) and the resulting tick mark height (rightmost histogram). For a large majority of the nearly 9000 fibers the optimal tick mark height was obtained when operated with the lowest gain settings. For approximately 69 % of fibers gain setting 0 was chosen, while 29 % of the fibers was operated with gain setting 1. The fraction of fibers operated with gain settings 2 (3)

is 1.7 % (0.6 %). The two peaks that are clearly visible in the distributions of the bias current and the tick mark height reflect the large populations associated with the two lowest gain settings.

The large body of data accumulated during these tests, together with the precisely controlled environment, allow to characterize the performance of the optical links in detail. The essential distributions are shown in Figure 13. For each of the four gain settings, the laser driver bias current was varied and the optimal value (as described in section 3.3) is chosen. The height of the tick mark measured in ADC counts for the given gain setting and optimal bias current was filled into the histogram. This procedure was repeated for all four gain settings and for all fibers on TIB and TID.

Table 5: The mean and standard deviation of the Gaussian fits to the opto distributions in Figure 13. Measurements were obtained under nominal CMS conditions (i.e., a coolant temperature of $-25°$C and an ambient temperature of $-10°$C).

	gain 0	*gain 1*	*gain 2*	*gain 3*
mean	500 ± 1	771 ± 1	1035 ± 1	1270 ± 1
RMS	69 ± 1	105 ± 1	134 ± 1	160 ± 1
gain	0.625 ± 0.001	0.964 ± 0.001	1.294 ± 0.002	1.587 ± 0.003
gain/gain0	1	1.54	2.07	2.54

For unit optical gain the FED dynamic range corresponds to a 1.024 V electrical signal. Therefore, for the highest gain setting(s) the FED reading saturates for the 800 mV tick mark signal. The tick mark height can nevertheless be reconstructed using an algorithm developed by Stefanos Dris [26]. It uses the result of the laser driver bias current scan to measure the slope of the analog baseline vs. the laser driver bias current and the bias currents at which tick mark and baseline cross a threshold value. The tick mark height is determined as the product of the measured slope and the difference of the crossing points of tick mark and baseline.

The optical gain distributions for each gain settings are described very well by a Gaussian, reflecting the high quality of the optical connections. The mean and RMS of the distributions, obtained from a Gaussian fit to the data, are listed in Table 5.

Table 6: The center and width of the Gaussian fits to the opto distributions corresponding to data taken at a coolant temperature of $0°$C.

	gain 0	*gain 1*	*gain 2*	*gain 3*
mean	451 ± 1	688 ± 1	937 ± 1	1171 ± 2
RMS	67 ± 1	102 ± 1	138 ± 1	160 ± 1
gain	0.564 ± 0.001	0.860 ± 0.001	1.171 ± 0.002	1.464 ± 0.003
warm/cold	0.90	0.89	0.91	0.92
gain/gain0	1	1.53	2.08	2.60

At higher temperatures, the optical gain was found to be slightly lower. Results are given in Table 6. A 10 % reduction in light yield was observed when comparing data taken at a coolant temperature of $0°$C with those taken with a coolant temperature of $-25°$C. Considering that the effective coolant temperature in the cooling manifold is closer to $-22°$C a temperature dependence of -0.5 ± 0.1 % /$°$C was inferred, where the error includes the uncertainty in the temperature difference and the differences found between the different gain settings.

4.7 Noise

The noise performance of the TIB and TID structures was studied carefully. Noise runs of the structures under nominal CMS conditions represent the central measurement of the validation procedure. Moreover, in an effort to optimize the grounding and shielding of the TIB and TID, extensive test programs were performed on selected structures. In the following the grounding scheme is introduced briefly. For a more detailed description the reader is referred to reference [27]. An overview of the most important noise results is presented.

4.7.1 Grounding

The main ground mesh in the TIB/TID system is provided by the cooling circuit. The ground of each Mother Cable is tightly coupled to the cooling system through a solder connection between a ground wire (HV return) in the power supply *medusa* cable and a *corona*, a copper-beryllium ring that goes around the full end-flange of each detector layer and is connected to the cooling manifolds.

On the modules a gold-plated copper clip (*cometa*) is soldered to the ground side of 1.25 V and 2.5 V capacitors. One edge of the clip scratches the carbon fiber module frame, providing a solid ground for the module frame. A picture of the hybrid with the *cometa* is shown in Figure 14.

The noise performance in the very first TIB burn-in runs was in general quite satisfactory, although some modules showed some excess (non-reproducible) noise on occasion. While most modules had a perfectly flat noise profile (versus strip number), in some cases a clear peak was seen. For single-sided modules the largest noise excess was found on the boundary between the first and second APVs and between the third and fourth APVs. For double-sided TIB modules, the noise was largest between the second and third and between the fourth and fifth APVs. Thus, the shape of the noise reflects the position of the cooling pipes under the sensor (and pitch adapter).

The only reproducible feature seemed to be that modules close to the end-flange (where the DOHM and CCUs are located) were more badly affected. Otherwise, the noise peak could vanish or double from one run to the next.
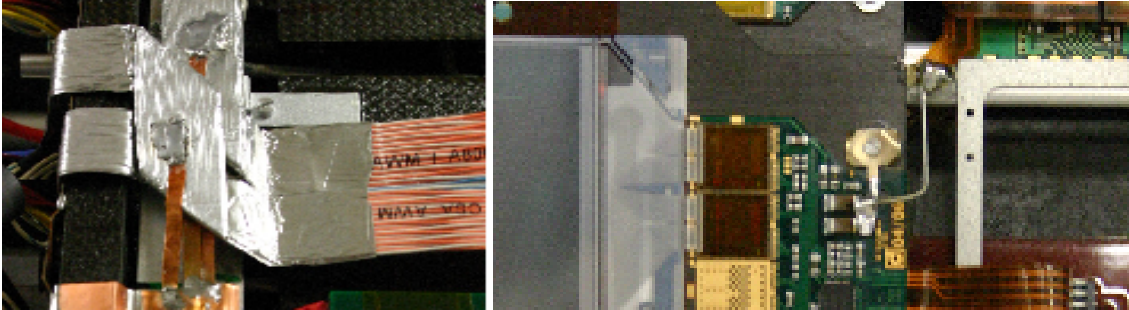

Figure 14: Leftmost figure: shielding of the cables fanning out the control signals from the DOHM to the Mother Cables. Rightmost figure: close-up view of a TIB hybrid. The *cometa* is clearly visible sitting just above the hybrid on the carbon fiber module frame. The wire soldered between the hybrid capacitors (AC coupling 2.5 V and 1.25 V to ground) and the cooling pipe can be distinguished.

.

Initially, the answer was to carefully shield all digital activity on the DOHM and, at a later stage, the final part of the flat signal cable that carries the signals from the DOHM to the CCU on the head of the Mother Cable. In Figure 14 a picture of the shielded flat cables is shown.

Further investigation of the excess noise revealed a curious periodicity. Shown in Figure 15 are the pedestal-subtracted ADC counts of one APV chip versus channel number (geographically ordered) and versus event number (time). The data show a clear "see-saw" effect: even events show a net positive slope: the ADC counts increases for higher channel numbers. In odd events the slope is inverted. To estimate the magnitude of the effect, the difference between the ADC counts of the first and and last channels of the chip (determined as the average of channels 1-16 and 112-128) is shown in the rightmost panel of Figure 15. Large oscillations of more than 10 counts occur. The magnitude of the oscillation is very clearly correlated with the excess noise structures observed.
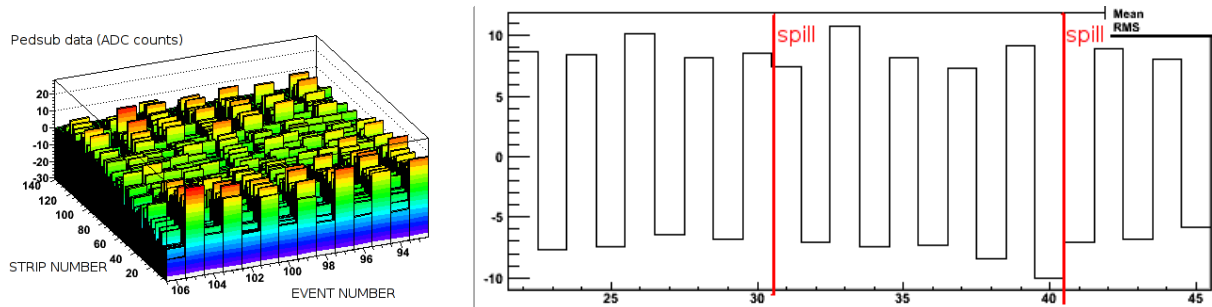

Figure 15: Pedestal-subtracted ADC counts for 128 channels of one APV chip versus event number (leftmost figure). A clear odd-even structure is seen: in even events there is a net slope (the ADC counts increase towards higher channel numbers), for odd events the slope is inverted. The difference of the average of the last and first 16 channels of the chip versus event number is shown in the rightmost figure.

During noise runs, typically 20 events were triggered at a modest, regular frequency of 10 kHz in a burst or spill. After the spill the FED buffers were transmitted through the VME bus. Depending on the total load on the DAQ the inter-spill period could take up to several seconds. Importantly, however, the interval between two consecutive

triggers corresponds to a fixed and odd number of (40 MHz) clock cycles. Therefore, although there may be a large time interval between consecutive events, the correspondence of even and odd event numbers to different phases of a 20 MHz signal is maintained within the readout spill.

Probing with a sensitive oscilloscope (active 1pF, 1MOhm probes) a clear 20 MHz structure, on top of the expected 40 MHz and higher harmonics, was indeed observed in the potential difference between the hybrid ground (taken on the hybrid capacitor) and the cooling pipes. In Figure 16 a time-series of the potential difference (leftmost figure) and its Fourier transform are shown (rightmost figure).

The excess noise on a certain string was found to be clearly correlated to the amplitude of the distortion on the cooling pipe: in runs where a large 20 MHz pick-up was observed, the excess was predictably high.
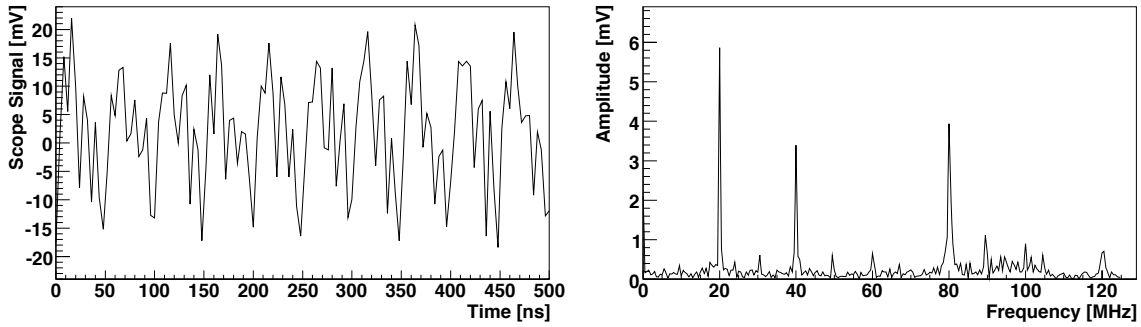

Figure 16: Oscilloscope time-series (leftmost figure) and Fourier transform (rightmost figure) of the potential difference between cooling pipe and hybrid ground.

The source of this pick-up turned out to be the DATA line of the control ring (which runs at this frequency). An idle line transmits a periodic 20 MHz signal. As part of the redundancy scheme, the ring is duplicated. The relative phase of both rings cannot be controlled by the DAQ: every time the ring is configured the idle signal on both rings may be in phase or in opposite phases. This feature explains the irreproducibility of the noise from one run to the next and the relative stability of the noise within the run.

Several configurations of the cooling circuit to module ground connections were tested for TIB structures equipped with single and double-sided modules. Both the *corona* and the ground connection through the DOHM seem to reduce the sensitivity to the 20 MHz pick-up for single-sided layers, where the power groups span up to four different Mother Cables. In double-sided layers the situation is more complex: the noise performance was best when the ground connections between Mother Cable and cooling circuit were removed.

Finally, it was found that a very robust noise performance could be achieved by creating a second ground connection between the Mother Cable and the cooling circuit at the far end of cooling loop. Therefore, on all TIB layers a wire was soldered between the ground on the hybrid capacitors and the cooling pipe (the wire is visible in Figure 14). It must be noted that much of this was learned while the validation tests of TIB+ were well advanced. Therefore, for some TIB single-sided layers the modifications took place only after the validation tests. In these cases no noise results are available in the final ground configuration.

The situation in the TID was found to be more subtle. The ground connections between the cooling circuit and the Mother Cable ground plane are achieved through a cable that runs around the full perimeter of the disks. A second, slightly less direct, connection is provided through the DOHM ground that is connected to the cooling circuit as well. While the qualitative features of the excess noise observed in TIB were found in TID as well, its effect was much less pronounced. This reduced sensitivity is believed to be due to the different orientation of the cooling pipes with respect to the readout strips (the pipes under the sensor are nearly perpendicular to the strips in TID). Also, the net of grounding connections between pipes, Mother Cables and DOHMs is such that *near* and *far* are not identified as readily as along a TIB string. Therefore, the additional grounding contact between cooling pipe and hybrid ground was not adopted for the TID.

4.7.2 Noise results

For each of the 2 million strips in the TIB and TID the common mode subtracted noise was calculated along the lines described in section 3.5. The noise results of each optical link were corrected for the variation of the optical link gain using the measurement of the tick mark height. The ADC counts were rescaled to a nominal optical gain of 0.8 (i.e., to yield a tick mark height of 640 ADC counts). To facilitate comparison with other subdetectors a

standard conversion factor of $25000. \times 8/640 = 312.5e^-$ / ADC count was introduced. This scale approximately establishes an equivalence between the noise and the input signal of the amplifier in electrons. The conversion factor is based on two assumptions: a minimum ionzing particle traversing 320 μm of silicon is expected to create 25 000 electron/hole pairs. Furthermore, the height of the tick mark at nominal APV gain corresponds to 8 MIP signals. The validity of these assumptions is discussed in Section 5.

All results in this section correspond to data taken under nominal CMS conditions as defined in previous sections:

- the cooling coolant temperature was set to $-25°$C, with *nominal* flow

- the chamber temperature was controlled at $-10°$C, the dew point is at a considerably lower temperature

- the APV chips were operated in deconvolution mode, with FE parameters as suggested by [23] for the hybrid temperature

- the sensors were biased at a voltage of 400 V

The data obtained for the 22 substructures is combined into larger structures by superposing the noise distributions. In Figure 17 the noise result for TIB+ (leftmost figure) and TIB- (rightmost figure) are shown. The distribution represents all 860 000 channels in each of the structures. In both figures, the x-axis represent the common mode subtracted noise in ADC counts after rescaling to the nominal tick mark height. A second scale on top of the figure shows the equivalent noise in electrons.
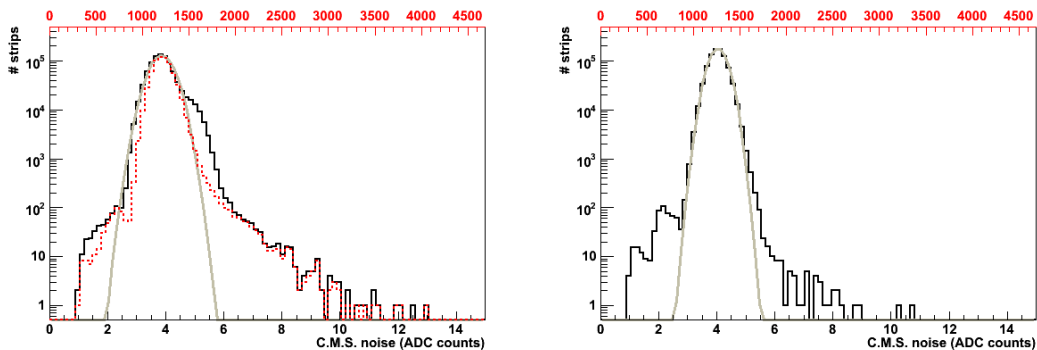

Figure 17: Noise distribution for all channels in TIB+ (leftmost figure) and TIB- (rightmost figure). The upper scale indicates the equivalent charge in electrons (see text). The dotted line in the leftmost figures is the distribution of TIB+ excluding the downward halve of layer3 and the upward halve of layer2. Data taken in *nominal* CMS conditions.
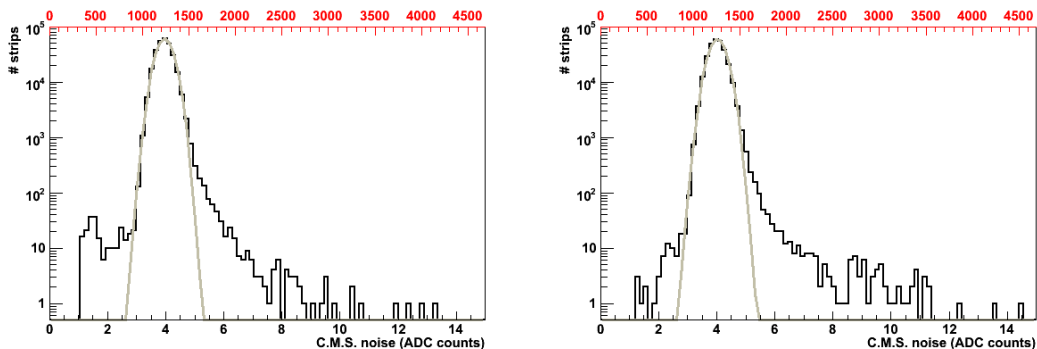
.

Figure 18: Noise distribution for all channels in TID+ (leftmost figure) and TID- (rightmost figure). The upper scale indicates the equivalent charge in electrons (see text). Data taken in nominal CMS conditions.

.

The uniformity in TIB+ results was clearly affected by an error in the start-up phase. On one TIB+ layer the wrong Front-End settings were used, leading to an increased mean of the noise distribution (visible as a slight shoulder in the distribution for TIB+). Also, as already mentioned, at the time of data taking the final grounding scheme was

not implemented for layer 4 of TIB+: this led to a significant tail on the high side of the TIB+ distribution. In the next section the noise results after implementation of the final grounding scheme are presented.

In Figure 18 the strip noise distribution is shown for all channels on the three disks of TID+ (leftmost plot) and TID- (rightmost plot).

The running conditions were carefully kept uniform throughout the year spanned by the validation tests. Indeed, the average temperature on layers of the same type agrees very closely, and the temperature variation between modules within each layer is quite small. There are, however, a number of systematic differences between structures of different types. Due to the increased power density on 6-chip hybrids, the hybrid temperature on double-sided TIB layers is greater than for single-sided modules in the same conditions by approximately $10°$C. Moreover, the readout pitch is 80 μm in layers 1 and 2 and 120 μm in layers 3 and 4, while TID wafers have radial strips with a pitch varying from 81 to 158 μm.

The average noise results for TIB and TID structures are listed in Table 7. In the first column the number of channels taken into account in the fit is given. Especially for TIB layers the number of channels was considerably less than the full number of channels in the finalized layer. For all TIB structures one string of modules was left open, to be mounted only after assembly of the cylinder. For this reason, approximately 25 000 channels are missing in TIB+ and TIB-. In all structures a small remaining difference is caused by masked channels and fibers.

A comparison of the noise results from the different types of layers in TIB yields no significant differences between TIB single- and double-sided modules or between TIB and TID modules. The noise results for TIB- double-sided modules yield an average noise of 4.03 ADC counts. The four values for the average noise on each half-layer lie within 0.13 ADC counts (the RMS is 0.07 ADC counts). For the four single-sided half layers in TIB- an average noise of the structure of 4.09 ADC counts, with an RMS of 0.06 ADC counts.

The Front-End noise is known to be dominated by the contribution from the load capacitance on the input to the APV. In reference [28] the load capacitance was shown to be proportional to the width over pitch ratio of the sensors, that is identical for all sensors used in TIB and TID and in other Tracker systems as well. Therefore, the observed identical noise results (within errors) is indeed expected.

Table 7: Gaussian fit results from noise distributions. Results are given for deconvolution mode and in nominal conditions, for peak mode and nominal conditions and for deconvolution runs in *warm* conditions (see text). The statistical error of the fit is below 0.001 ADC counts in all cases. For TIB+ the results are shown for the complete statistics and separately, marked with (*), for all layers except the lower half of layer 3 and the upper half of layer 1, whose data were taken with the wrong Front-End settings.

structure	channels	μ (ADC counts)	σ (ADC counts)
deconvolution, nominal conditions			
TIB+	864 749	3.86	0.38
TIB+ (*)	636 098	3.91	0.31
TIB-	857 215	4.05	0.29
TID+	279 808	3.97	0.27
TID-	281 481	4.03	0.28
peak, nominal conditions			
TIB+	866 129	2.77	0.22
TIB+ (*)	637 478	2.78	0.21
TIB-	858 688	2.87	0.21
TID+	279 810	2.79	0.20
TID-	281 481	2.81	0.21
deconvolution, warm conditions			
TIB+ (*)	784 787	4.79	0.35
TIB-	868 115	4.84	0.32
TID+	282 202	4.70	0.31
TID-	281 782	4.75	0.32

TIB and TID results were found to be compatible within the typical spread of measurements on layers of the same type. The three TID- structures have an average noise of 4.03 ADC counts and an RMS of 0.02 ADC counts.

Apart from the principal result under CMS conditions, as described earlier in this section, two other sets of measurements are given in Table 7. The second set corresponds to nominal temperature, but in this case the APV was

operated in peak mode. On average, the noise in peak mode is a factor of 1.44 smaller, as expected. The third set of measurements is taken in deconvolution mode, like the first one, but at relatively high temperature: the coolant temperature is $0°C$, the ambient temperature was kept at $10°C$. The APV Front-End parameters were changed accordingly, following the prescription by reference [23]. The combined effect of temperature and parameter settings is an increase in the noise level by 20 %.

4.7.3 Robustness of the performance

Throughout the validation tests large structures (a TIB semi-cylinder with up to 210 modules or a full TID disk) were powered, clocked, cooled, and read out in one continuous process. Thus, the tests provide a measurement of the noise performance in the presence of other modules and the digital control chain. After optimizing the grounding scheme the noise performance was found to be excellent.

In principle, electrical interference or grounding issues in the assembled TIB and TID systems could lead to sources of noise that were not present in the validation test. In the TIB, for example, the space between the four cylinders is very small. Especially, the Digital Opto-Hybrid Module of layer N is very close to the inner modules of layer N+1. A deterioration of the noise performance of the assembled TIB with respect to the single layer can therefore not be excluded *a priori*.

To explore the noise performance in conditions closely resembling those in the CMS collision hall, a dedicated test was performed after subassembly of the layers in which a full ϕ slice of four TIB+ layers was cabled and connected to the cooling plant. The availability of only 8 FECs and FEDs limited the scope of this test to one control ring per layer (two for double-sided layers), thus covering an angle of approximately 45 degrees.

The result in Figure 19 shows the excellent noise performance of the full system. There is no sign of interference or pick-up due to the presence of the other layers in the closely-stacked TIB layout. Note that this run was taken at relatively high temperature: $0°C$ coolant, $10°C$ air temperature. The front-end settings were changed accordingly. The average noise and RMS are fully compatible with the individual structure's result

Figure 19: Noise distribution for a ϕ slice of TIB+ in *warm* conditions (see text).

in the same conditions. They cannot be compared directly to the results obtained at nominal CMS temperatures presented above.

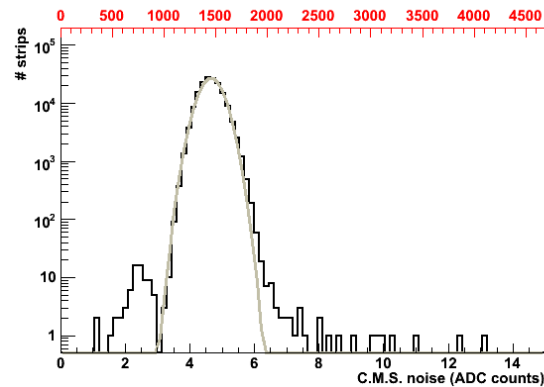The validation test noise runs have been performed using software triggers at fixed and rather long intervals. Thus, noise induced by, for example, data transmission of the previous event could be under- or overestimated systematically. For several layers the trigger rate was therefore varied in the accessible range (up to 100 kHz). Also, the performance with random triggers was tested. No significant change in noise performance was observed in either case.

4.7.4 Defect count

The burn-in noise runs at nominal temperature in peak mode are used to determine an estimate of the defective channels in each layer. A channel is considered dead (i.e., unbonded) if its noise level is below 75 % of the average noise level on the chip. A channel is considered noisy if its noise level is greater than the average noise on the chip by more than five times the (truncated) root-mean-square of the noise levels on that chip. The defect count thus obtained is listed in Table 8.

The noisy channel counts should be interpreted with great care. The definition of the noisy channel threshold is largely arbitrary. Moreover, the results depend quite strongly on external factors like the grounding and the power supplies. For example, the noisy channel count changed significantly after the service pack against common mode oscillations was installed on the CAEN system. It moreover depends on the operating mode of the chip (peak, deconvolution) and the data taking environment (warm/cold). The results listed here represent data taken at cold (CMS-like conditions) in deconvolution mode.

Table 8: Defect counts for TIB and TID.

structure	total # channels	# noisy channels	# noisy (no edge)	total # dead channels	# extra
TIB+	890 880	653	214	470	162
TID+	282 624	238	34	419	96
TIB-	890 880	372	105	99	89
TID-	282 624	293	80	115	62
TOTAL	234 7008	1556 (0.07 %)	433 (0.02 %)	1103 (0.05 %)	409 (0.02 %)

During the module construction phase in the different sites every module was characterized in the module production test [11], using a setup where a single module was read out using the ARC system [29] (referred to as single-module test in the following). A summary of the results is available from the construction database and a cross-check with the validation test results was routinely made. A comparison of the noisy channel lists from the structure validation tests and those from the single-module tests yields virtually no overlap. There was very little overlap between channels flagged as noisy in in the individual modules tests and channels flagged as noisy in the validation tests.

A large fraction of the strips labeled as noisy was found to be associated to the first or last of the 128 APV channels. These channels are excluded in the results listed in the fourth column of Table 8. This effect was observed also in the single-module tests, but edge channels were not uploaded into the construction database. Therefore, a more quantitative comparison is not possible.

The foregoing arguments suggest that the noisy channels found in either of the tests cannot simply be interpreted as intrinsically noisy APV channels (i.e., due to some defect in the chip). Whether a channel is noisy depends critically on the setup. The definition of noisy channel is moreover largely arbitrary. Therefore a channel mask for the TIB/TID based on excess noise cannot be extracted from the integration database.

The dead channel count is evaluated using the criterion that the noise level of the channel is less than 75 % of the chip average. As the peak mode runs turn out to be more sensitive, the dead channel is evaluated in peak mode, and nominal CMS conditions. In the case of dead and unbonded channels the single-module test results and the validation tests agree quite well. Typically, 90 % of the dead strips as determined by the single-module test were also found in the validation test. With the current data it was not possible to determine whether the remaining 10 % was due to "false positive" results of the single-module test or due to the inefficiency of the validation test. The fifth column in the table contains all dead channels, found in the single-module test, by the validation test or both. The sixth column lists those dead channels that were only found in the validation test. Again, part of these may be due to "false positive" results of the validation tests or inefficiency of the single-module test. The remainder is due to defects created during the integration (i.e., broken bonds). The total number of dead channels is very small. While the fraction of dead channels is similar for TIB and TID, the fraction of "new" dead channels is significantly larger in the TID.

5 Cosmic tests

As part of the qualification of the TIB/TID, cosmic runs were taken for several of the substructures. A significant fraction of TID+ was read out while triggering on the coincidence of the signal from two scintillators. Cosmic data were furthermore collected on a ϕ slice of the assembled TIB+.

The tests were performed using the standard burn-in DAQ and power system. The TIB+ structure was cooled with coolant at $0°$C using the standard burn-in cooling plant. The ambient temperature was set to $10°$C. For the TID+ test cooling was by ventilation (the San Piero climate chamber temperature was lowered to $-15°$C to keep the hybrids at approximately $30 - 35°$C).

Two large scintillators (120 cm by 20 cm) were installed to provide a trigger signal. In the TID+ both were mounted below the structure at a distance of approximately 1 meter. For TIB+ one scintillator was mounted on top of the cylinder and the other in the diametrically opposite position. The configuration of the trigger system allowed muons with a quite large range of incidence angles (up to 50 degrees in TID+, up to approximately 30 degrees in TIB+) to trigger the readout. The longer path length through the silicon leads to an increase of the average deposited signal.

The signal from the photomultipliers was discriminated using a NIM discriminator module. Additional NIM modules were used to implement a basic trigger logic (requiring the AND of both scintillator signals, implementing a rough hold-off scheme). The final trigger signal was sent to the TSC. To synchronize the readout of the modules

with the trigger signal, several time alignment steps were performed. The first step was the correction of the module-FED ribbon and multi-ribbon lengths (more specifically, two different types of Ericsson multi-ribbon cables were used with lengths of 50 and 60 meters). The time difference between both multi-ribbon types was measured using the results of the timing scan and the FED delays were corrected accordingly.

In the second step the standard XDAQ Tracker timing scan was performed to correct for the different delays in the clock and trigger signals across the structure. Finally, a latency scan was performed to select the correct position in the APV analog pipeline. The standard XDAQ Tracker latency scan shifts the coarse delay (multiples of 25 ns) of the global system trigger. In Figure 20 the average cluster charge (expressed as a signal to noise ratio) versus latency (each units represents a 25 ns step) is shown. The latency that maximizes the average cluster charge was chosen.

No attempt was made to further optimize the timing of the system: particularly the gate of the TSC, that allows to select triggers that arrive in a narrow interval within the 25 ns clock cycle, was left open. Therefore, the TID+ results represent the signal height averaged over the full 25 ns interval. To minimize the impact of the non-synchronous trigger all runs were performed in peak mode. The TIB+ time alignment was significantly worse than that of TID. Therefore, these data are not included in the quantitative analysis.

Figure 20: Result of the latency scan (peak mode) in the TID cosmic run.

To determine the signal height, the TID data were processed using the analysis client described in section 3. In Figure 21 the cluster signal results from one of the cosmic runs on TID+ are shown. All results from the different disks, and from different module types in each disk, are compatible within errors. In all cases the most probable signal is 87.6 ± 0.5 counts.
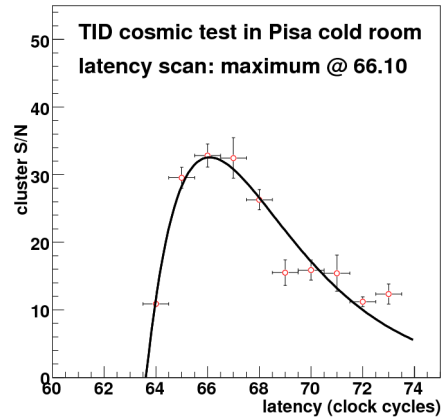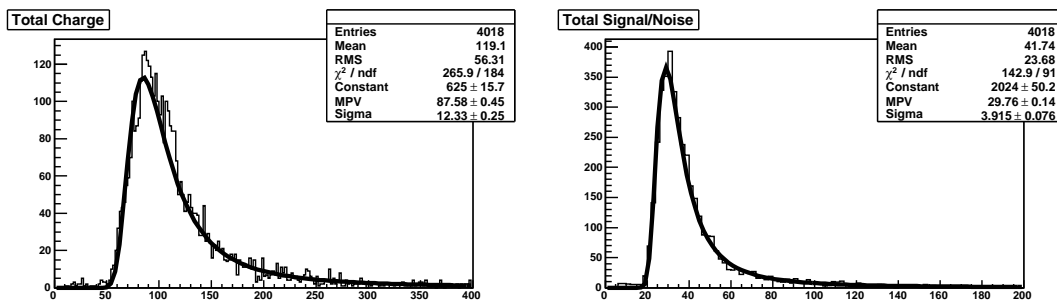

Figure 21: Cluster signal (leftmost figure) and Signal-to-Noise ratio (rightmost figure) distributions from cosmic muons in TID+. Data taken in peak mode.

Due to the non-perpendicular incidence of cosmic rays the Signal-to-Noise ratio observed in the cosmic test is considerably larger than that observed in beam tests under otherwise similar conditions [30]. In beam tests of TIB modules in peak mode a median S/N of about 25.5-26.1 was observed. The S/N ratio in the TID cosmics test was found to be 29.7 (an increase of approximately 15 %). The increase in the path length can be estimated using a simple geometrical model of the scintillator trigger acceptance. If the cosmic rays are assumed to have an isotropic distribution (this case can be considered as an upper limit) the median deposited charge is expected to increase by up to ~ 25 % (~ 7 %) for TID+ (TIB+) compared to a perpendicularly incident minimum ionizing particles.

A slight reduction of the observed median signal is expected due to the asynchronous trigger. Further uncertainties include the energy spectrum of cosmic muons and the details of temperature control and the Front-End parameter settings. With the resulting (large) uncertainty the obtained result is considered to be fully compatible with expectations from test beams. This measurement forms an important confirmation that the value used for the conversions factor to convert the noise measurement in ADC counts is correct to approximately 10 %.

In Figure 22 a cosmic track reconstructed in TIB using the full CMSSW chain, from pedestal subtraction to Kalman
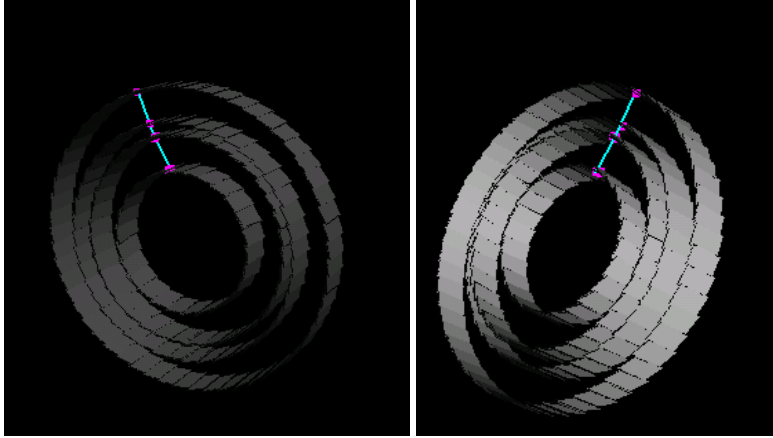
Filter track fit [5], is shown.

Figure 22: Reconstructed tracks in four TIB+ layers. The two double-sided layers (1 and 2) yield 3D space points, while layers 3 and 4 contribute only $R\phi$ coordinate. The full reconstruction chain was performed with standard CMSSW modules.

.

6   Conclusions

The TIB/TID integration phase was successfully completed. An extensive set of validation tests of the substructures was performed. The main aim of these tests was to validate the assembled structures in conditions that mimic the CMS environment as faithfully as possible. All readout conditions were chosen to resemble our current best guess of the CMS conditions: to this end the ambient temperature was carefully controlled, while the layers were cooled by a nominal C_6F_{14} flow through the manifolds. Final design CMS power supplies, cables, and DAQ system were used throughout the tests.

For each substructure, the devices were mapped onto DAQ channels using an automated procedure. Results were compared in detail to the integration database. Then, a number of commissioning runs was performed to optimize the settings of the PLL delay, the optical data transmission link and the analog baseline level. Once the system parameters were set, a number of noise runs was taken.

Throughout the validation tests many small problems appeared. All of these were solved in situ. The system operation proved to be very stable. Communication along the control chain was found to be reliable. The performance of the optical links for transmission of the analog data from the Front-End was found to be very uniform and in close agreement with expectations.

The power consumption of TIB and TID structures agreed closely with expectations. The voltage drop of several Volts in the long supply cables is effectively compensated by the sense circuit on the power lines. The uncompensated voltage drop was found to be in the range from 50 to 150 mV for the 2.5 V line under nominal load.

The TIB/TID liquid cooling scheme was tested for the first time under nominal CMS conditions. While the thermal load from the Front-End electronics was representative of that in the CMS experiment, the leakage current in the non-irradiated sensors was a negligible heat source. Under these conditions the hybrid temperature was found to be less than $0°C$ on 99 % of double-sided modules. On single-sided modules the hybrid temperature was significantly lower: on 99 % of modules the temperature is less than $-8°C$. With only five exceptions, the sensor temperature was well below $-10°C$ on all module types.

After careful optimization of the grounding scheme in TIB, the noise performance was excellent. The deconvolution mode single-strip noise in CMS conditions was found to be approximately 4.0 ADC counts, roughly equivalent to 1200 electrons. The noise performance remained unchanged when several layers in the final configuration were read out simultaneously. Random triggers and a higher trigger rate (up to 100 kHz) were not found to affect the noise results.

An analysis of the validation test data allows the determination of a defect count. Overall, less than 0.1 % of

---

[5] M. Mangano, private communication

channels were found to be defective. Noisy channels, defined as those channels where the noise is more than 5 σ above the chip average, made up 0.07 % of the total. Most of these were found on the edges of APVs. The fraction of unbonded channels, detected through their lower noise, was 0.05 % of the total. For a large majority of these channels, this result merely corresponded to the result of the single-module test. The remaining 0.02 % was observed solely in the validation tests and corresponds (at least in part) to damage during the integration phase.

A number of dedicated runs with triggers from a pair of scintillators mounted on top and below the assembled TIB and TID structure were taken. An analysis of these data showed Signal-to-Noise results were consistent with those obtained in earlier beam tests [30] and the magnet test [31]. These data provide an absolute scale for the noise results, although with limited precision. The observed conversion factor was found to be in agreement with the noise rescaling based on the tick mark height. These data were furthermore used to exercise the CMSSW reconstruction software. Tracks with measurements on four TIB layers were reconstructed with the default CMS track reconstruction algorithm.

References

[1] The CMS Letter of Intent, CERN/LHCC 92-3, the CMS Technical Proposal, CERN/LHCC 94-38, the CMS physics Technical Design Reports, CERN/LHCC/2006-001 and CERN/LHCC/2006-021.

[2] Tracker Technical Design Report, CERN/LHCC 1998/006 (1998) and addendum to the CMS Tracker TDR, CERN/LHCC 2000/016 (2000). More recent information is found in the Trigger-DAQ TDR, CERN/LHCC/2002-26 and in the Physics TDRs (previous reference)

[3] G. Sguazzoni, CMS Inner Tracker Detector Modules, CMS CR-2004/059

[4] M. Krammer for the CMS Tracker Collaboration, The silicon sensors for the Inner Tracker of the Compact Muon Solenoid experiment, Nucl. Instrum. and Methods A531 (2004) 238-245; L. Borello, E. Focardi, A. Macchiolo, A. Messineo, Sensor design for the CMS silicon strip tracker, CMS Note 2003/020; M. Krammer, The Silicon Sensors for the Compact Muon Solenoid Tracker - Design and Qualification Procedure, CMS Note 2003/015

[5] M.J. French et al., Design and results from the APV25, a deep sub-micron CMOS front-end chip for the CMS tracker, Nucl. Instrum. and Methods A466:359-365,2001; P. Barrillon et al., Production Testing and Quality Assurance of CMS Silicon Microstrip Tracker Readout Chips, CMS Note 2004/016

[6] J. Troska et al., Optical readout and control systems for the CMS tracker, IEEE Trans. Nucl. Sci.50 (2003) 1067-1072; D. Ricci, Experience in the Development of the CMS Inner Tracker Analog Optohybrid Circuits: Project, Qualification, Volume Production, Quality Assurance and Final Performance, CMS Note 2005/015; M.-T. Brunetti, B. Checcucci, V. Postolache, D. Ricci, Electrical Qualification of the Pre-production of Analogue Opto-Hybrid Circuits for the CMS Tracker Inner Barrel and Inner Disks. CMS Note 2003/027

[7] R. d'Alessandro et al., Signal and power distribution in the CMS inner tracker, CMS note in preparation

[8] G. de Robertis et al., The Digital Opto Hybrid Module (DOHM) for the Tracker Inner Barrel and Disks (TIB/TID): CCU interconnection and redundancy, CMS Note 2008/013

[9] Philips Semiconductors, The II C-Bus Specification, Version 2.1, document order number: 9398 39340011, January 2000

[10] C. Civinini, A. Messineo, G. Sguazzoni for the INFN TIB/TID consortium, The Tracker Inner Barrel Integration, CMS Note in preparation

[11] S. Albergo et al., Test of the inner tracker silicon microstrip modules, CMS Note 2009/001

[12] Power supplies are provided by CAEN, Viareggio, Italy. http:www.caen.it

[13] PVSS (for Prozevisualisierungs- und Steuerungssystem) is a Supervisory Control and Data Acquisition (SCADA) package developed by the Austrian company ETM Professional Control and recommended by CERN. For more information, consult the CERN IT webpage: http:itcobe.web.cern.chitcobeServicesPvsswelcome.html

[14] F. Drouhin et al., The CERN CMS Silicon Strip Tracker Control System, CMS CR 2004/032

[15] C. Foudas et al., The CMS tracker readout front end driver, IEEE Trans. Nucl. Sci. 52 (2005) 2836-2840; J. Coughlan et al., Design of the Front-End Driver Card for CMS Silicon Micro-strip Tracker Readout, Proceedings of 6th Workshop on Electronics for LHC Experiments, Krakow, CERN/LHCC/2000-041, p444

[16] M. Ageron, Trigger Sequencer Card User Manual, IPNL Lyon, 2001

[17] V. Brigljevic et al.,Using XDAQ in Application Scenarios of the CMS Experiment, CMS CR-2003/007

[18] L. Mirabito, Tracker Data Acquisition user guide Version 1.0, CMS IN-2004/019; L. Mirabito et al., Tracker Data Acquisition for Beamtest and Integration, CMS IN-2003/021

[19] Rene Brun, Fons Rademakers et al., *ROOT Users Guide.*, http://root.cern.ch/root/doc/RootDoc.html

[20] *CMS Tracker TIB Integration Test Software*, http://savannah.cern.ch/projects/cmstibintsft

[21] *Xerces C++ Parser*, http://http://xml.apache.org/xerces-c

[22] MySQL is an open source relational database management system. For further information consult: http:www.mysql.com

[23] M. Raymond, et al., Final Results from the APV25 Production Wafer Testing, 11th Workshop on Electronics for LHC Experiments, vol. CERN/LHCC/2005-038, pp. 453-457, 2005

[24] I^2C Signal propagation on TOB cards: where do we stand, A. Marchioro, 13/7/2003 v.2, hep.ucsb.edu/cms/rod/I2C_Status_middle_July_2005_for_TOB_meeting.pdf

[25] M. D'Alfonso, M.Vos, TIB $I^2$2C Summary of measurements so far, presentation in CMS week, 21st September 2005; M. Meschini, What we have learnt from I^2C errors, presentation in CMS week, 7th December 2005

[26] S. Dris, K. Gill, J. Troska, F. Vasey, Performance Verification of the CMS Tracker FED Analog Front-End Electronics, CMS Note 2006/026; S. Dris et al., Predicting the Gain Spread of the CMS Tracker Analog Readout Optical Links, CMS Note 2006/145

[27] R. d'Alessandro et al., Grounding and shielding of TIB and TID, CMS Note in preparation

[28] S. Albergo et al., Optimization of the silicon sensors for the CMS tracker, Nucl. Instrum. and Methods A 466 (2001) 300-307

[29] M. Axer et al., A test setup for quality assurance of front end hybrid, CMS-Note 2001/046

[30] T. Boccali et al., The 2003 Tracker Inner Barrel Beam Test, CMS Note 2008/014

[31] D. Abbaneo et al., Tracker operation and performance at the Magnet Test and Cosmic Challenge, CMS Note 2007/029; D. Benedetti et al., Tracking and alignment with the silicon strip tracker at the CMS Magnet Test Cosmic Challenge, CMS Note 2007/030