


The Compact Muon Solenoid Experiment

**CMS Note**

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



06 September 2010 (v2, 08 September 2010)

## Second level of the CMS Drift Tubes read out system: the ROS (Read Out Server)

C.F. Bedoya, P. Arce, J.M. Cela, M.C. Fouz, M.I. Josa, J. Marin, J. Molina, J.C. Oller, I. Redondo, C. Willmott  
CIEMAT, Madrid, Spain

### Abstract

The ROS (Read Out Server) boards are the second level of the CMS (Compact Muon Solenoid) DT (Drift Tube) chambers data acquisition system. They are responsible for collecting the information from 25 ROBs (Read Out Boards) and performing data merging and quality monitoring, reducing data overhead to build a synchronized event fragment of one sector. Output data is sent from each ROS through an 800 Mbps optical link to the DDU (Device Dependent Units) that forward them to the central DAQ (Data Acquisition) system. The complete design of the final ROS boards is described in this note, together with several characterization and validation tests which guarantee not only the optimal functionality of the electronics but also its satisfactory operation under the CMS environmental conditions.

# 1 The CMS Drift Tube Read Out system

CMS is a general purpose detector designed to run at the highest luminosity at the LHC collider. The central feature of the CMS apparatus is a superconducting solenoid of 6 m diameter that generates a magnetic field of up to 4 T. Such a high field was chosen in order to allow the construction of a compact tracking system on its interior, and still perform good muon tracking on the exterior. In the barrel, where the magnitude of the residual magnetic field is of the order of 2 T in the iron return yoke, and the neutron background and muon rate are expected to be as low as a few Hz/cm<sup>2</sup>, drift tubes are used [1].

DT chambers are responsible for muon detection and precise momentum measurement over a wide range of energies. The DT system also provides a reliable and robust trigger system with precise bunch crossing assignment. DT chambers are integrated in the five wheels of the return yoke of the magnet (named YB-2, YB-1, YB0, YB+1 and YB+2). Each wheel is organized in four concentric rings around the beam line, namely MB1 to MB4, with 12 sectors per ring to allocate each station. As sectors 4 and 10 of ring MB4 contain two chambers each, the total amount of chambers is 250.

A drift tube chamber consists of three (or two in MB4) Superlayers (SL), each of them having four layers of drift cells, staggered by half a tube width. The two outer Superlayers perform muon  $\Phi$  coordinate measurement (angle on the bending plane), meanwhile the inner measures  $\theta$  coordinate (angle in the non-bending plane). An aluminium structure, the honeycomb, is glued in between the Superlayers to provide the required stiffness to the chamber and a lever arm between the two  $\Phi$  Superlayers for a better track fitting. Chamber read out and trigger electronics are located in an aluminium structure called Minicrate, which is attached to the C profile of the honeycomb.

The basic element of the DT chamber is the drift cell, which has cross section dimensions of 13 by 42 mm. The total number of sensitive cells is around 172,000. Any charged particle going through a cell volume will generate a signal (hit) in its anodic wire that will be amplified and discriminated by the front-end electronics before being sent to the read out electronics to perform time digitization. The position of the charged particle can be related to the time measurement since the drift velocity in the cell volume is constant. Each drift cell provides a resolution of 250  $\mu$ m, and the 100  $\mu$ m target chamber resolution is achieved by the 8 track points measured in the two ( $r$ - $\Phi$ ) SL.

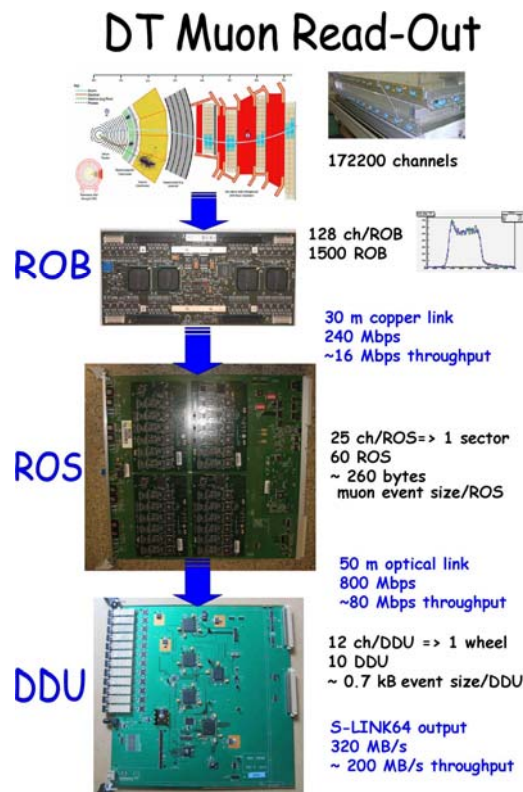


Figure 1: Schematic view of the DT read out chain.

The first element of the DT read out chain, the ROBs [2], digitizes up to 128 chamber signals and measures their time of arrival with respect to the Level 1 Accept trigger signal. ROBs are based in the HPTDC (High Performance Time to Digital Converter) [3] developed by the CERN EP/MIC group. ROBs are located inside the Minicrates. There are between 3 and 7 ROB boards per Minicrate to perform the full read out of one DT chamber, totalling 1500 ROBs in the system.

Output data is sent through 30 meter copper links at 240 Mbps to the 60 Read Out Server boards (ROS), located in the tower racks in the cavern. ROS boards are in charge of merging the information from one full sector and performing several tasks of data reduction and data quality monitoring, keeping the proper synchronization with the whole detector through the TTC (Timing Trigger and Control) system [4]. Each sector event is sent through a 50 m optical link at 800 Mbps to the DDU (Device Dependent Unit) [5] boards located in the underground service cavern (USC55). The DDU boards merge data from up to 12 ROS to build a consistent event fragment and send it to the global CMS DAQ through an S-LINK64 output at 320 MBps. These boards also perform error detection on data and send a fast feedback to the TTS (Trigger Throttling System) [6].

All this steps, detailed in figure 1, have to ensure a secure path from the chambers to the DAQ processing. The chosen read out architecture and buffer sizes, processing time and link bandwidths have been designed to guarantee the read out of the full DT detector at a Level-1 trigger rate of 100 kHz.

## 2.1 System requirements

The developed system must satisfy certain requirements imposed by LHC and CMS operation. On the functional side, even though the charged particle rate expected at the tower racks is foreseen to be low,  $10 \text{ Hz/cm}^2$ , the drift chambers require a system capable of performing continuous time digitization in an extended window as long as at least the maximum drift time (400 ns). Since the maximum drift time is much larger than the LHC bunch crossing period (25 ns), a read out system that manages overlapping triggers is required. Therefore, the actual event size depends not only on the hit rate but also on the L1A rate, since some hits may be read multiple times [2]. L1A trigger rate can reach 100 kHz due to LHC high luminosity, imposing requirements in terms of maximum allowed processing time.

On the environmental side, a remnant radiation is expected in the detector area as a result of the products of the successive interactions. This forces the employment of radiation tolerant devices that have to be tested. In 10 years of operation at an LHC luminosity of  $10^{34} \text{ cm}^{-2}\text{s}^{-1}$  the expected neutron fluence is of  $10^{10} \text{ cm}^{-2}$ , being charged particles flux of  $10 \text{ cm}^{-2}\text{s}^{-1}$  and the integrated dose of 0.4 Gy.

Furthermore, the high magnetic fields created by the CMS solenoid impose certain restrictions to the electronics power consumption. In particular, tower racks cooling capability is limited due to the need of using tangential fans, capable of operating under the actual magnetic field. Accordingly, power consumption in the ROS has to be minimized and thermal distribution has to be taken into account in the design of the board.

Finally, the operation of the CMS detector is foreseen to last about 10 years and during this time, the maintenance of the ROS will be restricted to LHC technical stops and shutdowns. Therefore, a robust and reliable system is demanded, not only requiring minimal interventions but minimizing as much as possible failures propagation. Accordingly, the system has been designed in such a way that an error in one part does not handicap the operation of the rest.

## 3 The ROS board

The ROS board is a 9U VME [7] board developed at CIEMAT. Each board is capable of reading 25 input channels (ROBs), i.e., one sector, and transmits merged digital data through an optical link to the DDU board. A view of the ROS board and its front panel can be seen in figure 2.




Figure 2: Image of the ROS board and of its front panel, with the 8 RJ-45 input connectors and LC optical output.

### 3.1 Integration in the Sector Collector crate

The 60 ROS boards are located in the so-called Sector Collector (SC) crates inside the cavern, in the towers on one side of the CMS wheels. There are two SC crates per wheel, where ROS is integrated with the Trigger Sector Collector (TSC) boards, sharing the power supply, slow control (VME) interface and the TTC signals that are distributed by the TIM (TTC Interface Module) board through a custom backplane (TIMBUS). The SC crates, the TIM board and the TIMBUS have been designed and produced at CIEMAT. A diagram of the SC crate is shown in figure 3.



Figure 3: Diagram of the Sector Collector crate with its boards VME base address.

The data from the ROB boards arrives to the ROS through two FTP (Foiled Twisted Pair) cables per Minicrate, 4 pairs each, and these differential pairs are routed to each ROS processing block as can be seen in figure 5. The mapping between ROB board and ROS input can be seen in [8]. An image of the SC crate fully installed can be seen in figure 4.




Figure 4: Image of a fully installed SC crate in the CMS detector.

### 3.2 ROS architecture

ROS board is made of a motherboard to which several modules connect, namely GOLROS, ROSCTRL and four CEROS mezzanines, laying out as in figure 5. These mezzanines have been designed to allow proper cooling not blocking the vertical air flow. The subdivision in modules allows, on the one hand, parallel processing of different channel groups, optimizing processing speed, and on the other hand, fitting and routing the elevated number of devices that need to be used in a reasonable amount of space. It will also simplify the reparation of damaged boards.

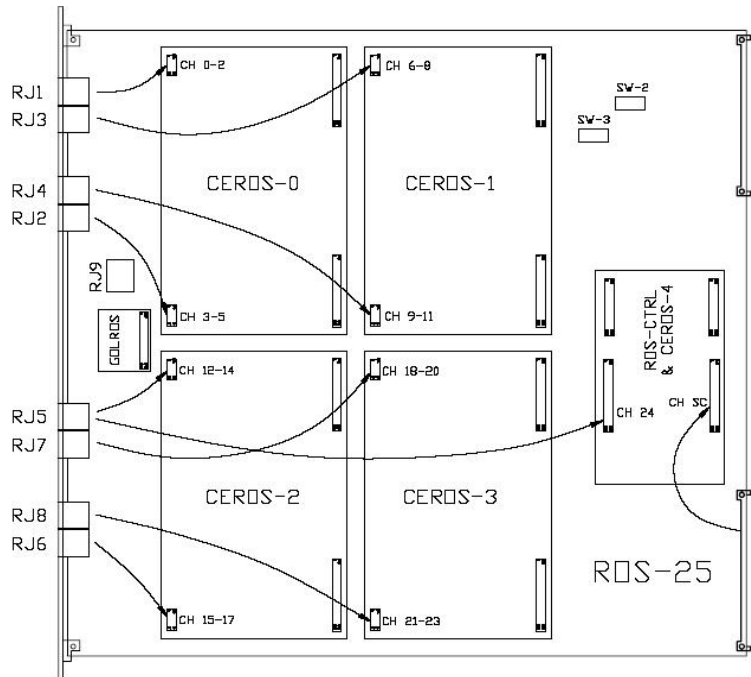


Figure 5: Diagram of the main components of the ROS board.

The main functionality of the different modules will be described here; more details will be given in the following sections:

- CEROS: each CEROS mezzanine is controlled by a CEROS FPGA (Field Programmable Gate Array), in charge of deserializing, storing and processing a group of 6 input channels, independently from the other mezzanines. Output data is transferred to a common parallel bus for further high speed optical transmission.



- **ROSCTRL**: this mezzanine and its corresponding FPGA have two main functionalities. The first one, called CEROS-4, is to perform the read out of the 25<sup>th</sup> channel. The second one is to coordinate the CEROS read out, interfacing the TTC system and insuring proper event building. It also connects to the neighbour TSC board to include local trigger data into the read out stream for debugging purposes.
- **GOLROS**: the parallel output bus is received at this mezzanine and sent to a GOL (Gigabit Optical Link) [9] device for serialization. Serial data is sent to a VCSEL (Vertical Cavity Surface Emitting Laser) transmitter that is connected to the output optical fiber.

The ROS board has been designed with high programmability in order to allow as much flexibility as possible in the operation mode. Moreover, several operation modes have been designed that permit easy debugging and possible functioning in test stands where main CMS systems are unavailable.

### 3.2.1 Normal operation mode

This is the operation mode during CMS data acquisition runs (figure 6). For each L1A, the ROSCTRL receives the event and bunch identification from the TIM board and activates the read out protocol of the different CEROS. Each CEROS will receive the corresponding event number and will process its 6 input channels in a token ring scheme. At each CEROS block, data is processed in parallel to speed up the read out and then, it is transmitted to a common bus to the GOLROS following a start token protocol controlled by ROSCTRL.

Data consistency, transmission errors and event synchronization are checked at each CEROS, and any integrity error is notified not only in the status registers but also within the data flow. Moreover, only valid information or error status is transmitted to further levels of the read out chain, creating a data reduction to achieve 8 kbytes per event in the whole detector.

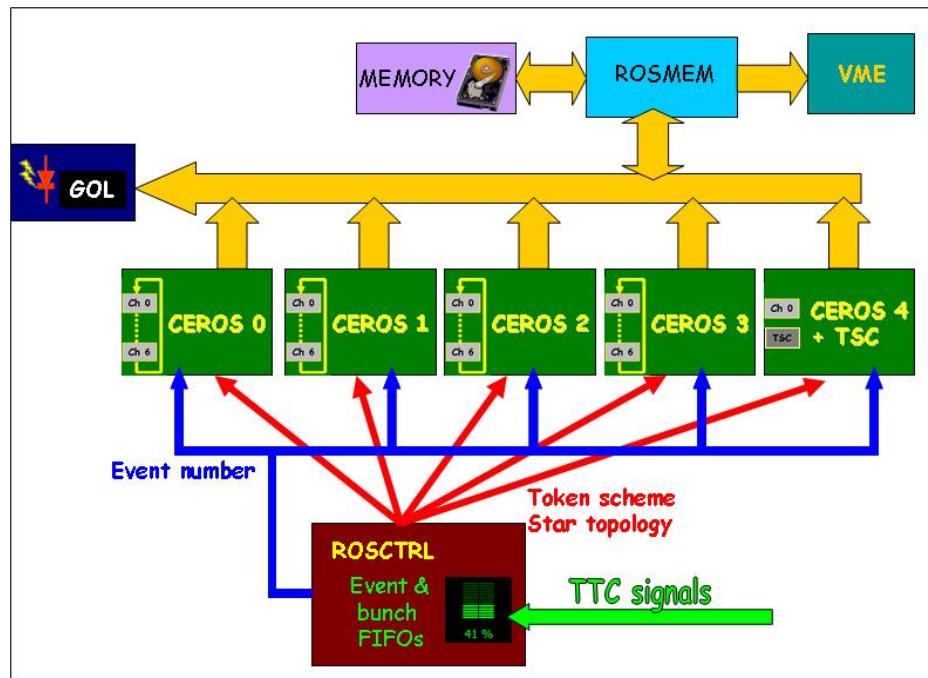


Figure 6: Diagram of the ROS modules interconnection in the normal operation mode.

### 3.2.2 Spy mode

This mode of operation can be used simultaneously with normal data acquisition. The ROS board can be programmed to make a snapshot of a programmable number of events or number of words in a 512 kbytes internal memory. This memory can be accessed through the VME interface without interfering regular data flow. Versatile interruption mechanisms have been designed to handle this spy mode. The fraction of events that can be read in this mode is limited to the VME bandwidth, but since it does not slow down ROS normal processing operation, it is very useful to scan data for debugging purposes.

### 3.2.3 Transmission test mode

In this operation mode the input channels processing is ignored and only the output transmission part remains functional. Data is written through VME into the ROS internal memory and the desired number of words are

sent to the optical transmitter. The number of cycles and the average bandwidth of the transmission is programmable, therefore, it is a very useful mode to study the reliability of the ROS-DDU link.

### 3.2.4 Direct FIFO read out

This mode of operation has been implemented for debugging purposes, where data can be read through VME directly from the channels input FIFOs. In this way, the data format is the same as the one provided by the ROB, since the ROS operation is almost transparent in this mode. As a main disadvantage, data read out has to be done channel by channel, with a large processing time overhead.

The ROS different operation modes are schematically represented in figure 7.




Figure 7: Diagram showing the different operation modes of the ROS board.

Finally, the ROS board has been designed to be able to work in tests stands lacking a TTC system. Automatic internal clock, trigger signals and event number generation have been implemented in the board to allow as much as possible to use it in stand alone mode.

## 3.3 CEROS module

The CEROS module performs the first stage of data merging. Each of the 6 channels (one channel in the CEROS-4) includes an equalizer (CLC014AJE), a deserializer (DS92LV1212A) and independent 4 kbytes FIFO (IDT72V243). These FIFOs include a PAF (Programmable Almost Full) flag to indicate when the FIFO occupancy is larger than a programmable limit.

A Xilinx XC2S50E-7FT256 FPGA manages the read out protocol that starts after ROSCTRL broadcasts the event number to process. A simplified diagram of the state machine that the FPGA follows to process each channel is shown in figure 8. The FPGA checks whether a channel is masked or blocked due to previous errors and, if not, it waits for its input data. Once arrived, it checks that the event number is synchronized to the one received from the TTC system and extracts the corresponding words from the channel FIFO. In the case that only headers and trailers are received from that ROB, the data is discarded and the FPGA starts to process the next channel. Otherwise, it waits until the ROSCTRL authorizes the transmission of the valid data to the output bus.

During data processing, the following checks are performed by the CEROS FPGA:

- Proper locking of the link and automatically block in case of unlock.
- Indicate any parity error during reception and marking it in the corresponding word. Total number of parity errors is counted and available through the monitoring registers.
- Verify proper format of input data, accepting only known data types and in the expected order.
- Check event synchronization between ROB input data and TTC information.
- Monitor the input FIFOs occupancy and signal when the programmed limit is exceeded. If the FIFO is completely full, the corresponding channel will be temporarily blocked to avoid data fragmenting.
- Limit the total number of words to be read from one single channel. Some problems, either in the transmission itself or in the ROB operation, could lead to a countless number of words before an event trailer is found. In this unlikely circumstance, the channel is blocked limiting the maximum time the ROS should be processing one single channel.

- If no words are found in the input FIFO, handle the correct timeout mechanism and limit waiting time to a programmable value.

Every possible failure during data processing will be signalled not only in the ROS internal registers that are accessed through the slow control, but also, within the data flow. In this way, offline analysis can easily discard events where data integrity errors may have happened.

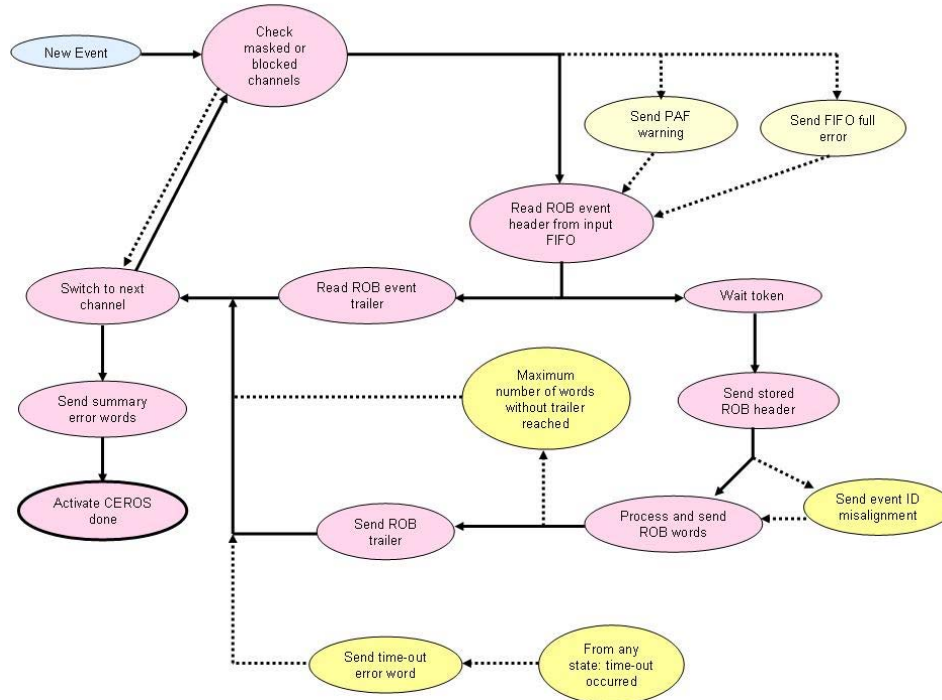


Figure 8: Diagram of the simplified CEROS state machine. In blue the initial state, in yellow, optional states depending on the input data.

### 3.4 ROSCTRL module

The ROSCTRL module includes a Xilinx XC2S100E-7FTE256 FPGA that manages the 25<sup>th</sup> input channel, an additional input channel for the TSC data and all the logic to control the ROS functionality in normal operation mode. This module receives the TTC signals from the TIMBUS backplane and activates the start token protocol that will permit the different CEROS to transmit their data to a 16 bit common bus (figure 6). Proper watchdogs have been implemented to insure that the event processing does not hang due to a CEROS malfunctioning. These watchdogs are only initialized due to a transition of the CEROS state machine that guarantees that the module is still processing valid data.

TTC event and bunch counter information is stored in 256 words FIFOs to allow handling of overlapping events. When the occupancy of these FIFOs exceeds a programmable value, a warning is signalled both in the registers and in the data flow so that the system can react in advance to loss of synchronization.

At each L1A, the ROS can also read the TSC information through some dedicated lines in the TIMBUS backplane. ROS and TSC boards are interleaved in the crate, so that each ROS reads the data of its twin TSC board. A data\_ready/get\_data handshake protocol has been implemented and TSC data will be stored in an internal ROS memory until it is included in the event fragment. This mode of operation is foreseen only during commissioning phases and no handling of overlapping of events has been implemented by default.

A summarized version of the state machine that describes the ROSCTRL functionality is shown in figure 9.



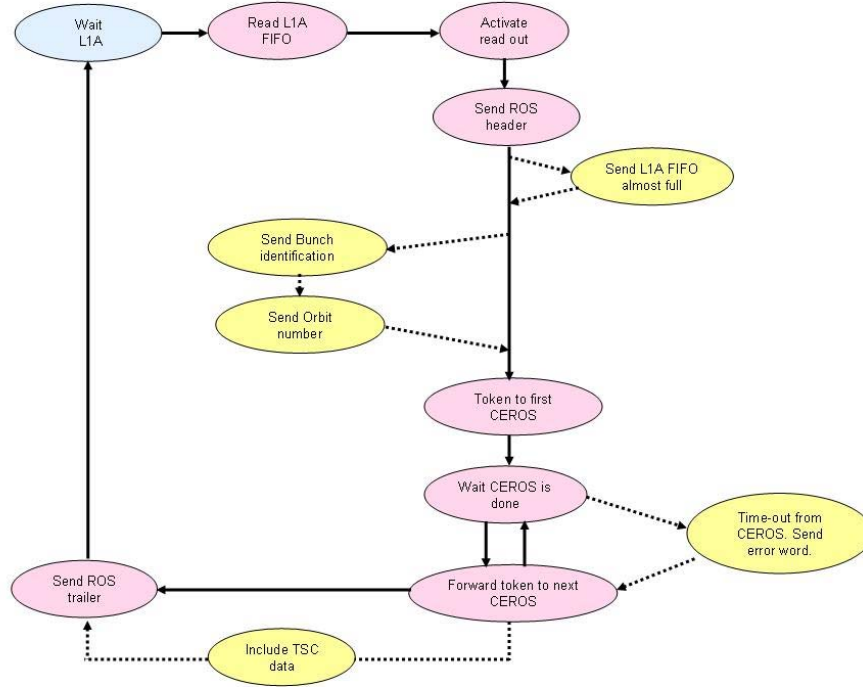


Figure 9: Diagram of the simplified ROSCTRL state machine. In blue the initial state, in yellow, optional states depending on the input data and on the programmed values.

### 3.5 GOLROS module

The GOLROS module is in charge of serializing the data and sending them to the DDU board in the underground service cavern. The GOL [10] ASIC is configured to use Gbit Ethernet protocol. Since 8B/10B codification is used, the effective bandwidth is 640 Mbps for a 800 Mbps link. The estimated throughput in the detector is around 80 Mbps.

Clock is fed to the GOL device through a QPLL [11] that filters the signal distributed by the TIM board. The GOL also includes an I<sup>2</sup>C protocol for slow control of internal parameters such as the laser bias current. The selected laser is an 850 nm VCSEL (HFE4190-541) capable of operating up to 2.5 Gbps. It includes a p-i-n photodiode to monitor optical output power. This optical power is measured through an ADC and allows performing remote calibration in case of optical loss due to aging or radiation effects of the fibers. In figure 10 the measured optical power of the VCSEL as a function of the programmed laser bias current in the GOL is shown. In the central region, where the values are more adequate for our device, the behaviour is pretty linear.




Figure 10: Optical output power of the VCSEL as a function of the programmed laser bias current at the GOL.

Also, the measured current from the ADC as a function of the programmed laser bias current in the GOL is shown in figure 11. Even though it is not completely linear, it is good enough to correlate programmed current, expected optical output power and current monitored by the photodiode.

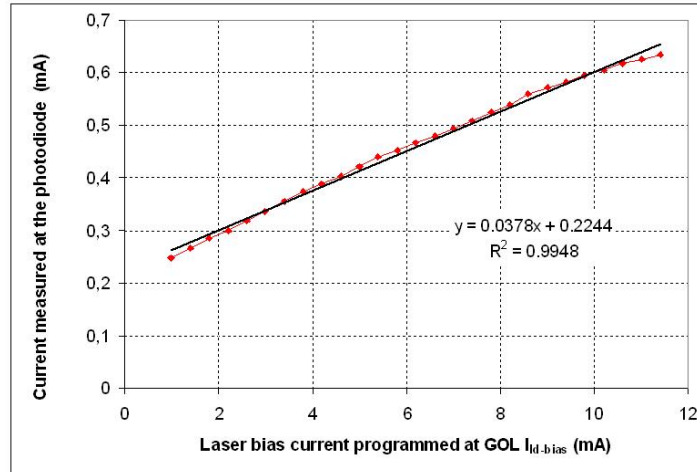


Figure 11: Current measured at the photodiode by the ADC as a function of the programmed laser bias current at the GOL.

Since in normal conditions, the optical output power of the VCSEL is -4.6 dBm and the optical input power of the DDU receiver is -17 dBm, the power budget of our link is around 12.4 dBm. The expected attenuation of our links is around 2.2 dB due to fibers attenuation factor and interconnection losses. Accordingly, there is a large margin for degradation of the link without losing reliability.

### 3.6 ROS motherboard

Besides interconnecting the different modules, the ROS motherboard contains additional functionality. First of all, it includes the voltage regulators for the different power supplies of the board with the corresponding temperature, current and voltage sensors. The ROS maximum power consumption is limited thanks to an automatic power supply protection circuitry with fast shut off capability. In case the different currents exceed the values specified in table 1, the regulators are switched off. Since over-currents could be short in time, due to radiation induced phenomena, after 700 ms the protection circuit will try to power on again the board. If the over-current persists, it will keep on repeating power on/off cycles guaranteeing a minimal current consumption and insuring the safety of the system. The total current consumption of a ROS board is 4.1 A, that is, 21.32 W.

Table 1: Maximum currents allowed for each of the ROS power supplies by the protection circuitry.

	$I_{max}$ (A)
5 V	2.53
3.3 V	2.53
1.8 V	1

ROS motherboard is also in charge of the VME interface for the slow control, which is handled by two CPLDs (Complex Programmable Logic Device) ROSVME and ROSMEM. The first one controls the A16 access to the configuration and monitoring registers of the different modules, as well as the access to the 1-wire protocol for the on board sensors, the I<sup>2</sup>C interface for the GOL slow control and the JTAG interface for the remote reconfiguration of the FPGAs. ROSMEM is in charge of the A24 access to the ROS internal memory. It manages most of the debugging operation modes of the ROS board. The different registers of the ROS board are described in [8].

#### 3.6.1 Remote reconfiguration

All programmable logic devices in the ROS with the exception of ROSVME can be reprogrammed remotely through VME. This feature has great advantages since basically the complete functionality of the ROS board can be modified through a software program without arduous interventions.

Devices reconfiguration is done through a JTAG interface. In the case of the CEROS and ROSCTRL FPGAs, an external configuration memory (XC18V01) is accessed, while for the ROSMEM CPLD, the configuration

memory is internal to the device. The ROSVME CPLD firmware includes the necessary registers to access the JTAG lines through VME and it also contains some control logic to speed up the JTAG protocol for long instructions. ROSVME also controls the loading of a new program into the FPGAs in case the internal memory of the FPGAs is corrupted. A diagram of the main connections between the devices is shown in figure 12.

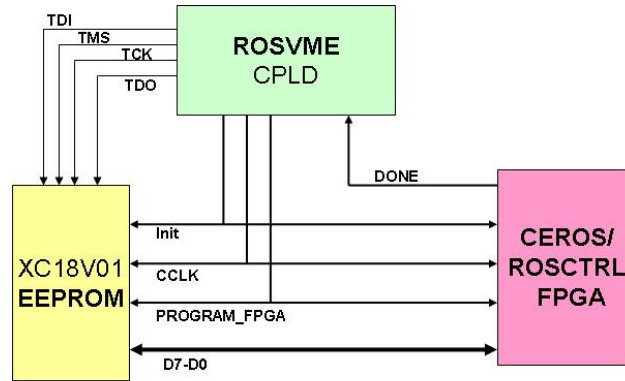


Figure 12: Diagram of the interconnections between the CEROS or ROSCTRL FPGA, its corresponding configuration memory and the ROSVME CPLD.

### 3.7 Data format

ROS words, likewise ROB words, are 32 bits long. Each ROS event is enclosed between a ROS header and ROS trailer which contain the event number information, the number of transmitted words and some flags to indicate the occupancy of the L1A FIFO. Additional words can be enabled to include the information of the bunch crossing and the orbit number of the corresponding event.

The ROB information for the 25 channels is sent in order. As mentioned before, only ROB events that contain time measurements or error information will be transmitted, in order to reduce the total throughput. At each CEROS, the ROB data format is modified to include certain information such as the ROS input channel the data correspond to or eventual parity errors. TSC data is also encapsulated with its corresponding headers and trailers as any other channel.

On top of that, there are two types of error words that are transmitted within the data flow. First type is a single 32 bit word per channel, sent only in the event where the error was found, that includes all the information regarding the error type and the status of the system at the time. Second type is a summary 32 bit error word per CEROS, which is sent on every event where any channel is kept blocked due to a malfunctioning. In this way, the status of the detector can be known at each event without knowledge of the past history, while maintaining a reduced overhead. The information of the ROS data format is detailed in [8].

### 3.8 The TTS interface

In CMS, the DDU boards have direct connection to the central trigger system through the TTS (Trigger Throttling System) interface. The TTS signals are used to slow down or stop the trigger rate in case of imminent buffers overflow or to require actions (TTC commands) in case of synchronization or hardware problems. Since the ROS does not communicate with the TTS directly, a protocol between ROS and DDU has been established in order to trigger the necessary actions.

In figure 13 the transitions between the different DDU TTS states and the corresponding TTS actions are shown. The TTS behaviour of the DDU is based on the information of internal FIFOs occupancy, lost of synchronization or error words transmitted from the ROS.

The DDU will be informed of any ROS malfunctioning through the ROS error words embedded in the data flow. There are two basic actions that can be triggered depending on the error words types: move to Out of Synch state or move to Warning Overflow (and possibly to Busy) state. When the central trigger detects an Out of Synch, it will generate a Resynch command [12]. When the state is Warning Overflow or Busy, the trigger rate will be slowed down or stopped, in order to avoid buffer overflows in the system.

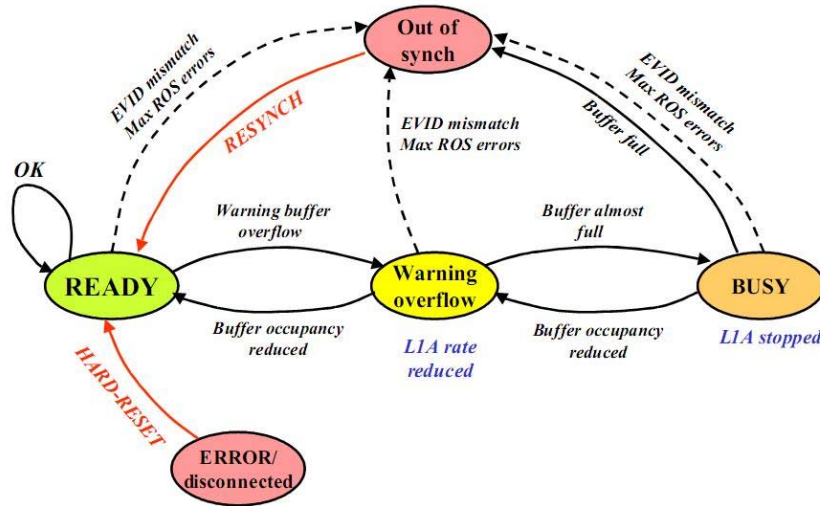


Figure 13: Transitions between different DDU TTS states and its corresponding TTS actions.

The ROS error words that may trigger an Out of Synch are described in table 2. These errors usually imply that the corresponding channel or group of channels is blocked. The Resynch procedure will empty all input FIFOs, unblock the channels and allow further data to be read out from the corresponding input. As can be seen, the number of elements that could generate one of these error words is very high; therefore, there is a programmable threshold in the DDU over which, the DDU will go into Out of Synch. Since these ROS error words are only generated once in the event where the error was found, there is no problem of double counting, and the threshold is directly translated into the fraction of the detector that is not being read out (In case of a CEROS timeout error word, the counter is incremented in 6 units). Also, there is a programmable flag in each ROS error word that can be disable to exclude that error word from DDU counting, in case of a known persistent problem.

Table 2: ROS error words that can trigger an Out of Synch state in the DDU TTS.

	Possible number of sources	Chamber channels affected on each error
Link timeout Error	1500	128
Event ID misalignment Error	1500	128
FIFO Full Error	1500	128
CEROS timeout Error	240	768
Max number words Error	1500	128

There are two possible warning words in the ROS that may trigger a DDU TTS transition to Warning Overflow: L1A FIFO imminent overflow or input FIFO almost full (table 3). The behaviour in the ROS for these warning words is slightly different than for the error words of table 2, since the corresponding channel is not blocked and the warning word is sent in every event where the occupancy of the FIFO is above a programmed value.

Table 3: ROS warning words that can trigger a Warning Overflow/Busy state in the DDU TTS.

	Possible number of sources	Chamber channels affected on each warning
ROS INPUT FIFO almost full	1500	128
ROS L1A FIFO imminent overflow	60	3200

When the warning corresponds to L1A FIFO imminent overflow, a DDU counter is incremented. This counter is decremented if no warning word is present in the event. If the counter exceeds a programmed threshold, the DDU moves to Warning Overflow state. If it goes over a second threshold, it will move to Busy state. An appropriate hysteresis mechanism has been designed to avoid toggling around the thresholds (figure 14).

When the warning word is of the type input FIFO almost full, the number of sources is very high and the affected area may be due to particular noisy region which effect does not want to be propagated. Therefore, two counters are used in the DDU. The first one is incremented with the number of warning almost full words received from all sources in that event. The second will count the number of consecutive events in which the first counter reached the threshold. Only if the second counter reaches a programmed value, a TTS action is

taken. A hysteresis mechanism similar to the one in figure 14 is also implemented.

As a final remark, several registers have been implemented in the ROS and in the DDU that will not be erased by a Resynch command and therefore, allow monitoring *a posteriori* any error circumstance.




Figure 14: Diagram of the DDU TTS transitions related to ROS L1A FIFO imminent overflow warnings.

## 4 Monte Carlo Simulations. Occupancy and processing time studies

A Monte Carlo simulation has been performed in order to study the expected hit rates in the DT Chambers and to estimate ROS event processing time during LHC operation [13]. Minimum bias events, produced at the IP (Interaction Point), have been generated and the resulting particles propagated through the CMS detector. These particles will produce particle showers in their interaction with the CMS material, and a fraction of them will eventually reach the muon chambers and produce a hit. We define punchthrough as those events without a muon in the trajectories of all the particles in the event.

Nearly one million minimum bias events were generated inside the CMSSW environment that includes the Pythia program and Geant4 simulation of the detector. The events were generated in the  $-3 \leq \eta \leq 3$  range and their passage through the detector was simulated with the 4T magnetic field effects included. High precision neutron physics were used (option QGSP BERT HP in Geant 4) in order to take into account all possible contributions from neutron scattering around the detector. We estimate that the error on neutron production can be a factor 2.

Figure 15 shows the number of hits detected by the five different wheels of the DT muon detector. As expected, external wheels show a higher occupancy. It can be seen that 99123 hits were detected in the 994883 events generated at the IP. From those 99123 hits, approximately 88% were produced by punchthrough particles, and are shown by the red dashed line.




Figure 15: Occupancy of particles in the five wheels of the CMS DT muon detector. The black continue line represents all the hits produced, while the red dashed line shows the hits produced by the punchthrough process. They represent around 88% of the total number of hits detected.



Most of these background particles have an arrival time different from the muons generated at the IP. Figure 16 shows the time of flight for all the events together with the ones corresponding to punchthrough. Clearly we can see that the punchthrough process takes longer times to reach the detector (in fact, 28% of them will reach the detectors after a period time of 1  $\mu$ s).




Figure 16: Arrival time (ns) of particles to the chambers during the 1 us time window. The peak around 20 ns is attributed to the time of flight of particles originated at the IP and reaching the first chamber. The red dashed lines are the hits caused by punchthrough particles.

Figure 17 shows the event hit multiplicity produced by the particles that go through the detector. Peaks for 12, 24, 36 and 44 hits correspond to events where a charged particle track leaves signals in 1, 2, 3 or 4 muon chambers. It can be seen that in most of the events we expect very few hits.

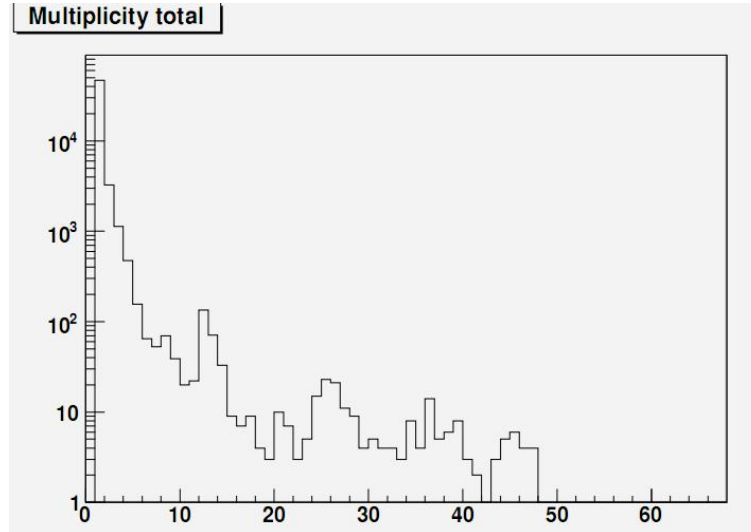


Figure 17: Histogram of the hit multiplicity of the simulated events.

Even though the punchthrough events do not contain useful physics information, they will contribute to the occupancy of the memories in the read out system and to the average event size; therefore, we are taking those hits into consideration to study the hit rates in the detector.

Since we are interested in the maximum number of hits that the ROB and ROS have to deal with, we have to look for the most populated sector in the most populated wheel. In figure 18 we can observe that the maximum amount of hits are in sector 6 of wheel 2, which has an occupancy of around 2800 hits. We consider this worst case for the rest of the discussion.

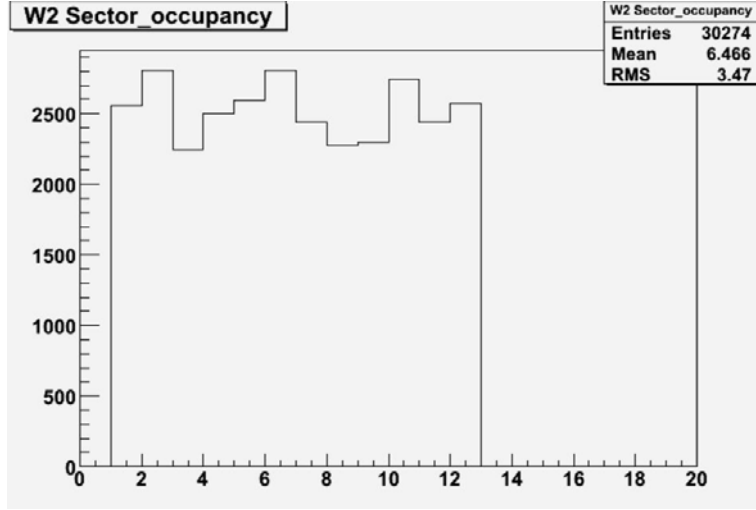


Figure 18: Occupancy of particles by sector in wheel 2 of the DT muon detector. The big differences among sectors are due to multiples interactions.

#### 4.1 ROB buffers occupancy

For a luminosity of  $10^{34} \text{ cm}^{-2}\text{s}^{-1}$  and using the total cross section value  $\sigma = 80 \text{ mb}$ , we have that  $80 \cdot 10^7$  events per second are produced at the interaction point. Figure 19(a) shows that from the total of particles arriving to each sector, 79% of them will give a signal in the first station, 17% in the second, 3% in the third, and only 1% in the fourth one. Therefore, maximum buffers occupancy will happen in the HPTDCs of MB1. Therefore, from the largest occupancy of 2800 hits, we have that 2212 (79%) are in the first station. As the event rate in the LHC is  $80 \cdot 10^7$ , the frequency of hits that we may expect in the MB1 is 1.78 MHz.

In figure 19(b) the occupancy of the three SLs in MB1 is shown. Since each  $\Phi$  ROB is connected to both  $\Phi_1$  and  $\Phi_2$  SLs which contain 37% and 30% of the hits respectively, and there are 12 HPTDCs reading  $\Phi$  SLs, we expect that each HPTDC will receive a hit rate of 99.4 kHz, which translates in a chamber channel rate of 3.1 kHz. Comparing this result with other studies [14] that obtain 1-2 kHz per channel, we conclude that our higher rate is due to considering the worst case scenario.

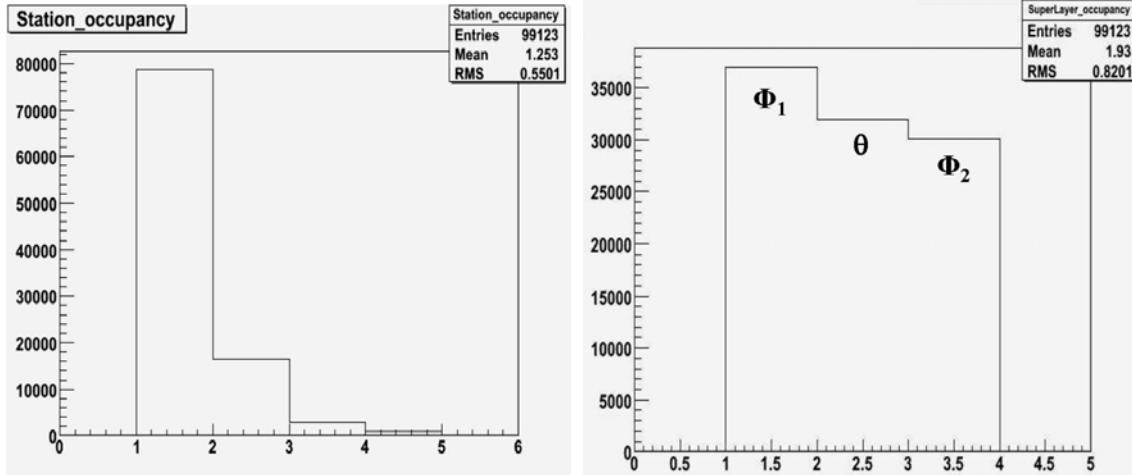


Figure 19: (a) Particles detected in each station of the CMS DT muon detector. (b) Distribution of the hits of the chamber among the SLs.

These 3.1 kHz hits will remain in the HPTDC Level 1 buffer until the trigger matching is performed, after a latency of  $3.2 \mu\text{s}$ . From tests performed with the HPTDC [3], it can be seen in figure 20 that the Level 1 buffers of the HPTDC are capable of standing without problem 1 MHz of hit rates in all channels with a latency of  $10 \mu\text{s}$ . Note that in the previous calculation, the rate obtained was 3.1 kHz per channel, therefore, one can conclude that the hit rates that we obtain from minimum bias events are much smaller than what the HPTDC Level 1 buffers can stand.

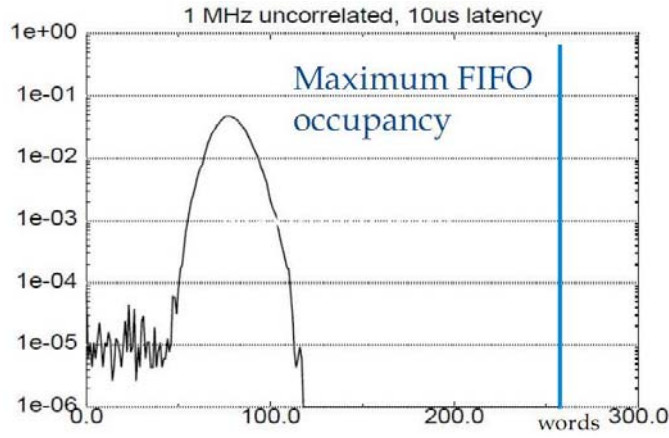


Figure 20: L1 buffer occupancy at 1 MHz hit rate for 10  $\mu$ s window.

We also have to take into account the reading speed of the four HPTDCs inside one ROB. We know that the actual number of hits that will be stored in the HPTDC read out buffer will be reduced by the trigger matching mechanism. For a matching window of 1  $\mu$ s, the hit rates will be reduced to 0.31 kHz if one considers an uniform distribution of the hits. A better approximation is to consider that 72% of the hits will enter in the time window, as obtained from the time of flight figure (figure 16). In such a case, we will be expecting 2.4 kHz hits that are read out per HPTDC channel. In figure 21 an example is shown in which the HPTDC can stand 100 kHz hit rate per channel (uncorrelated within the channels), 50 kHz trigger rate, and four HPTDCs sharing a 40 MHz serial read out protocol. Although in the present discussion the trigger rate is 50 kHz, a faster 20 MHz byte-wise read out protocol is used in the ROB. Either case, the occupancies we obtain (including the estimated error in the neutron production) are far from saturating the HPTDC buffers. It is also worth saying that the HPTDC can be configured to reject hits whenever the read out FIFO is full, so that it can continue operating even though some hits are lost.

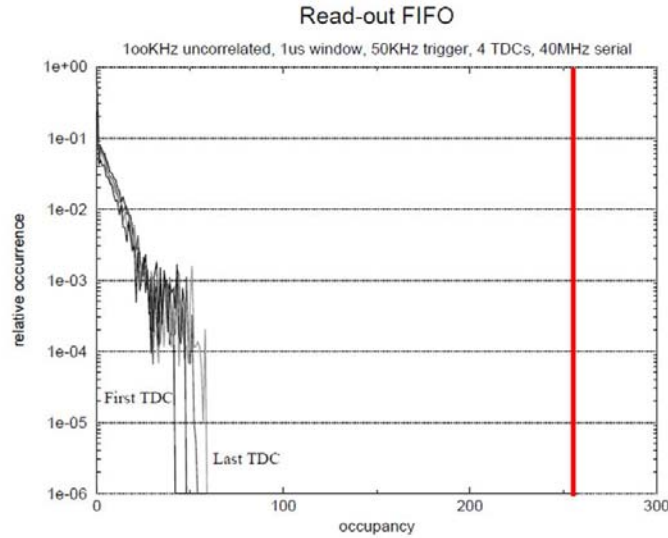


Figure 21: Read out FIFO occupancies for four HPTDCs sharing a slow serial read out.

## 4.2 ROS event processing time

Specific software has been developed to simulate the behaviour of the ROS board, i.e., the different firmware of the CEROS and ROSCTRL FPGAs and their interconnection. This software allows calculating the ROS processing time of an event for a given pattern of number of words in its ROB input channels. The time for the ROS to process an event depends on the particular ROS channel that received the hit, as can be seen in figure 22, where the time is measured in terms of the bunch crossing (bx), equivalent to 25 ns.

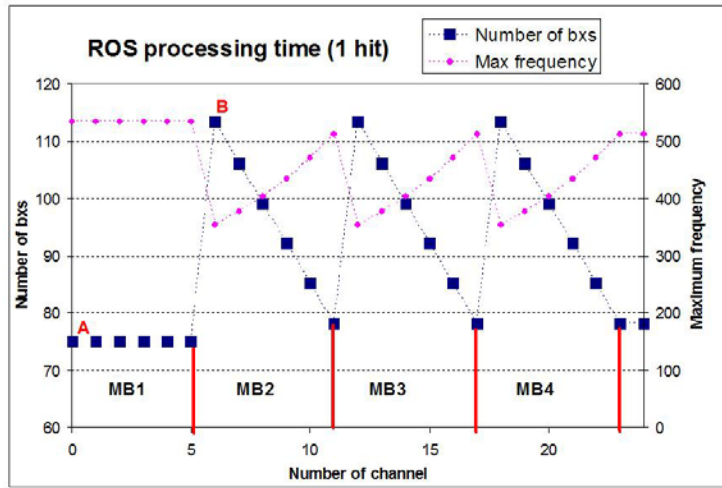


Figure 22: Time for the ROS to process 1 hit versus channel number (ROB) that received that hit. Maximum processing frequency as explained in the text is also plotted.

Channels corresponding to MB1 stations are faster, which accommodates to the higher occupancy expected there. In the previous figure, maximum frequency is calculated just by dividing the time to process an event (i.e. black squares in figure 22), that is, it does not mean that this is the maximum L1A rate that it can stand. We are not taking into account the input and trigger FIFOs from the ROS board that will allow storing events until the previous one has been processed. Therefore, the real maximum allowed L1A frequency is larger than what is shown in this paper and depends on the particular event size and on the actual time between L1As.

For events with a larger number of hits but all in the same channel, the time to process an event increases according to the formula:

$$\text{Time to process an event} = 73 + \text{offset} + 2 * H$$

where H is the number of hits in the same ROS channel and offset is the extra number of bxs that a particular channel needs to process an event with respect to what a channel in MB1 does, as can be seen in figure 22.

Figure 23 plots the maximum frequency as explained before versus the number of hits in the same ROS input channel both in the best case (hits in MB1, point A in figure 22) and in the worst case (hits in first channel of MB2, point B in figure 22). It can be seen that the ROS handles easily very noisy regions contained in the same input channel (equivalent to a group of 128 chamber channels). Note that 1 hit per event is equivalent to having one chamber cell with 1 MHz noise or all the 128 cells with 8 kHz noise.

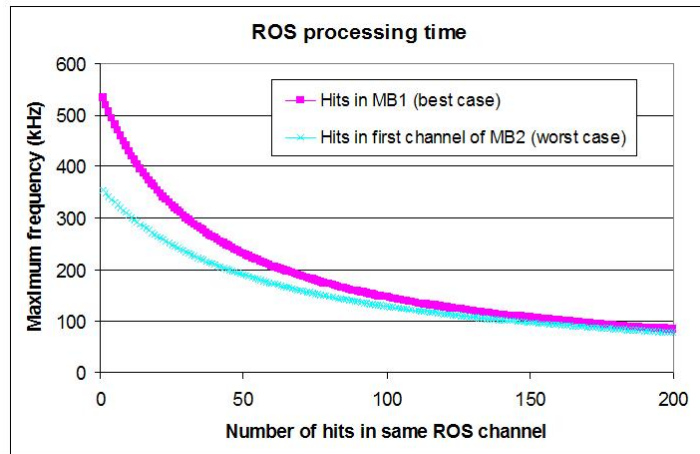


Figure 23: Maximum ROS processing frequency (assuming no FIFOs) versus number of hits with all hits in the same input channel.

As shown, the way the ROS has been designed, it takes much more time to process an event when hits are spread between different ROS input channels compared to what we obtain if all the hits are in the same ROS input channel. This effect can be seen in figure 24, where an average frequency has been calculated for the different processing times depending on the ROS input channel that gets the hit. The lower curve is the extrapolation of

the processing time that we would get if all the hits were distributed along all the ROS input channels. Square dots, corresponding to multiples of the number of channels (25) represent realistic situations. Figure 24 also shows a track-like curve where an extrapolation has been made for the actual processing time that would involve processing a muon-like track. A muon track will produce around 44 hits distributed in a particular configuration through the ROS input channels. Without storing any event, ROS can handle one muon per L1A at 134 kHz, and two muons in the same sector per L1A at 103 kHz. Needless to say, this rate is much higher than what we expect at the LHC. The bottom line of the x axis represents the equivalent noise frequency that one should have in all cells of all the chambers in that sector to obtain the corresponding number of hits. Note that there are 2900 chamber channels per sector.

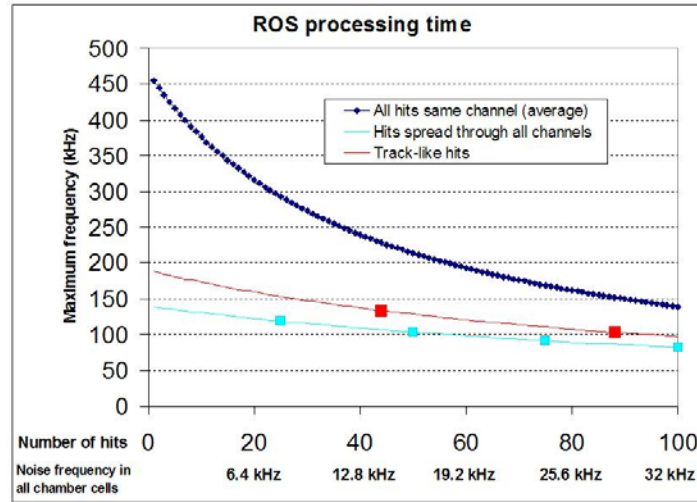


Figure 24: ROS processing time without storing events versus number of hits in one sector for three different hit distributions.

It is worth saying that at present, the level of electronic noise present while running the detector in normal operating conditions (high voltage ON, no cells masked, etc.) is very low as can be seen in figure 25 for different data acquisition campaigns both with magnetic field on and off. We declare a cell as noisy when the hit rate is over 500 Hz. Note that the expected background during LHC will provide in the order of 1 to 10 kHz of noise per cell. The number of noisy cells in the system is very low, in the order of 20 or 30 out of 172200 and is pretty stable in different run conditions.




Figure 25 Distribution of the cell noise rate for different conditions of data taking (CRAFT08)

After describing ROS processing capacities, we examine the impact of the number of hits in the simulated LHC events. From the  $80 \cdot 10^7$  events per second that we will have at LHC luminosity, and the worst case sector (2800 hits), the maximum hit rate per sector translates to 2.25 MHz. Considering a matching window of  $1 \mu\text{s}$  and 100 kHz L1A rate, we obtain that the number of hits that the ROS will have to process per L1A is 2.25. Again, one could take into account the time of flight distribution of all hits, and that would produce 16.2 hits to process



per L1A.

As shown in figure 17, most of the simulated events produce 1 or 2 hits, therefore, it seems reasonable to consider an average of 3 hits as the background that the ROS will have to handle in every L1A besides the actual number of muons. In such a case, the maximum frequency at which the ROS can process events without storing any data in FIFOs varies between 459 kHz if all the hits are spread in different MB1 ROS input channels and 199 kHz in the worst case, when all the hits are in the first ROB of MB2, MB3, and MB4. There is not much difference if one considers the 16.2 hits per event, which in the worst case will lead to a processing frequency of 182 kHz. Therefore, ROS will process events much faster than the expected L1A rate, so most of the time ROS FIFOs will be empty.

As a remark, these values obtained of the ROS processing time have been compared with real operation of the board at the laboratory injecting known patterns, obtaining a total match.

The largest event size per ROS input channel that we may expect including background and muon is around 52 bytes. Since the ROS input FIFOs have a capacity of 4 kbytes, up to 77 consecutive events can be stored assuming that all the events have the muon in the same ROS input channel. When muons are spread throughout all chambers and not always in the same ROB, up to 142 events can be stored in the input FIFOs before they fill-up. Therefore, there is quite a big margin for bursts of L1As before ROS FIFOs fill-up. It is also worth saying that an L1A input FIFO allows storing 256 L1As in case one L1A is received before finishing processing the previous event.

In summary, rates expected at LHC are low enough for comfortable operation of DT read out electronics, which will easily handle this data throughput, including the electronic noise measured in the DT chambers seen in the fully installed CMS detector.

## 5 ROS irradiation tests

The ROS boards are installed inside the CMS cavern in an environment that is not severely hostile in terms of radiation dose (maximum 0.4 Gy in 10 years of LHC operation), but subject to a substantial probability of Single Event Effects (expected fluence  $10^9$  p cm<sup>-2</sup> in 10 years of operation). Since radiation hard devices were not going to be employed, irradiation tests have been performed in the selected devices in order to verify that they are tolerant to such doses.

Irradiation tests were performed at the Cyclotron Research Centre at the University of Louvain (UCL) with a total fluence of  $5 \cdot 10^{10}$  cm<sup>-2</sup> protons at 60 MeV in two different campaigns in June 2003 with a ROS prototype (ROS-8) and in March 2007 with the final prototype. Irradiating with 60 MeV protons is considered an effective method of testing the electronics for total dose effects, displacement damage and Single Event Upsets (SEU) [15]. An image of the ROS test stand for the irradiation campaign can be seen in figure 26.




Figure 26: Image of the ROS test stand for the irradiation campaign.

In the first irradiation campaign the input stage of the ROS channels was tested with a ROS prototype that could read 8 channels (figure 27). The irradiated devices were the CLC014AJE equalizer, the DS92LV1212A deserializer and the IDT72V263 FIFO. For each run, a collimator allowed illuminating two devices of the same kind. A custom set-up was built in which known data was transmitted between a ROB and the ROS and any type of data corruption was constantly monitored.




Figure 27: Image of the ROS-8 prototype used during the first irradiation campaign.

The results obtained are presented in table 4, where  $N_r$  is the number of devices irradiated,  $N_F$  the number of failures observed,  $N_T$  the total number of devices in the final CMS system,  $\lambda$  the estimated failure rate for a 90% confidence level and MTBF the mean time between failures. The 18 failures observed in the equalizer were due to locking problems. When irradiating the deserializer we found 11 parity errors in the data and 18 errors due to unlock of the links. In the case of the FIFOs, 135 parity errors were found in the data. All these errors are consistent with Single Event Effects. Power consumption was also monitored and no malfunction or degradation in the irradiated devices was observed during the tests.

Table 4: Estimated mean time between failures at 90% confidence level due to radiation in the ROS-8 board for 10 years of operation in CMS.

	CLC014	DS92LV1212A	IDT72LV263
$N_T$	1500	1500	1500
$N_F$	18	29	135
$N_r$	16	16	16
$\Phi_{CMS} (cm^{-2}s^{-1})$	20	20	20
$\Phi_{LOV} (cm^{-2}s^{-1})$	$5 \cdot 10^{10}$	$5 \cdot 10^{10}$	$5 \cdot 10^{10}$
$\lambda$	$6.75 \cdot 10^{-7}$	$1 \cdot 10^{-6}$	$5 \cdot 10^{-6}$
MTBF (hours)	411.5	255.4	54.9
MTBF (days)	17.1	10.6	2.3

In the second campaign one full ROS was already available and hence, the remaining devices were irradiated. Dedicated tests were done depending on the region to be irradiated, in total, 9 areas were tested with different collimators depending on the size of the devices and the number of them to cover. Some areas include several devices with common functionality that can be tested with the same software; therefore, they were irradiated simultaneously. In table 5 the actual devices that were irradiated in each region is summarized.

Table 5: Summary of the different devices irradiated in each of the ROS regions.

Power supply region	Clock region	VME region	TTC signals receivers
MAX4375TEUB DS2438Z MIC39301-2.5BU MIC29301-3.3BU BTS612N1	DS90LV048A DS1100L DS90LV018A IXO71-40MHz CY2309ZC-1H	74LVC16245 74LVCH16244 74LVCH244 74ALS642 74ALS688	SN65LVDM1676 DS90LV017A DS90CP22MT DS90LV048A

The software developed for these tests includes read-back of the configuration memories, continuous VME write and read accesses, operation of the ROS in the different modes including data acquisition at 100 kHz trigger rate with input data in all 25 channels, monitoring of the power consumption, stability of the clock, etc.

The results obtained are indicated in tables 6 and 7, where  $N_T$  is the number of devices irradiated,  $N_F$  the number of failures observed,  $N_r$  the total number of devices in the final CMS system,  $\lambda$  the estimated failure rate for a 90% confidence level and MTBF the mean time between failures.

Table 6: Estimated mean time between failures at 90% confidence level due to radiation in some of the devices of the ROS boards for 10 years of operation in CMS.

	CPLD XC2C384	CPLD XC2C512	FPGA XC2S50E	PROM XC18V01	CY7C1041CV33 Memory
$N_T$	60	60	240	240	60
$N_F$	6	46	1614	0	20
$N_r$	2	1	2	2	1
$\Phi_{CMS} (cm^{-2}s^{-1})$	20	20	20	20	20
$\Phi_{LOV} (cm^{-2}s^{-1})$	$5 \cdot 10^{10}$	$5 \cdot 10^{10}$	$5 \cdot 10^{10}$	$5 \cdot 10^{10}$	$5 \cdot 10^{10}$
$\lambda$	$7.2 \cdot 10^{-8}$	$1.1 \cdot 10^{-6}$	$2.95 \cdot 10^{-5}$	$4.8 \cdot 10^{-8}$	$4.8 \cdot 10^{-7}$
MTBF (hours)	3858.02	251.61	3.59	2513.28	578.70
MTBF (days)	160.75	10.48	0.15	104.72	24.11

Table 7: Estimated mean time between failures at 90% confidence level due to radiation in some of the devices of the ROS boards for 10 years of operation in CMS.

	Power supply region	Clock region	VME region	TTC signals receivers
$N_T$	60	60	60	60
$N_F$	0	40	12	16
$N_r$	2	1	2	2
$\Phi_{CMS} (cm^{-2}s^{-1})$	20	20	20	20
$\Phi_{LOV} (cm^{-2}s^{-1})$	$5 \cdot 10^{10}$	$5 \cdot 10^{10}$	$5 \cdot 10^{10}$	$5 \cdot 10^{10}$
$\lambda$	$1.2 \cdot 10^{-8}$	$9.6 \cdot 10^{-7}$	$1.44 \cdot 10^{-7}$	$1.92 \cdot 10^{-7}$
MTBF (hours)	10053.11	289.35	1929.01	1446.76
MTBF (days)	418.88	12.06	80.38	60.28

Going into the detail of some of the errors found:

- When irradiating the CY7C1041CV33 and the XC2C512 CPLD the errors found were due to single bit upsets of the data transferred. For the XC2C384 CPLD 5 of the errors were also due to single bit upsets, though, in one case, the crate needed to be power cycled due to an incorrect VME access.
- During the irradiation of the XC2S50E FPGA there were many upsets in the output data, but the read out mechanism remained operational. A couple of times it was necessary to reload the FPGA configuration and, once, a power cycle of the crate was needed.
- In the power supply region the only observed effect was variations of the voltage and current smaller than 1%, which is compatible with the previous radiation tests performed in those devices in the ROB boards [2].
- In the clock region 5 errors were due to QPLL unlocks and the rest due to incorrect sampling of the data

which is related to incorrect phase of the clock.

- In the VME region all the failures were due to data read with one bit upset, so the most sensitive device seems to be the 74LVC16245.
- When irradiating the TTC region the failures observed were 2 upsets of the orbit number, 1 upset of the bunch number, 10 QPLL unlocks (related with operation of DS90CP22MT device) and 3 incorrect sampling of the TSC data.

Summarizing, the selected components for the ROS are adequate for operation under the radiation doses expected during operation in CMS. The smallest time between failures for the XC2S50E device is still compatible with the expected operation mode since FPGA reconfiguration is foreseen frequently in the system.

## 6 Commissioning and operation in magnetic field campaigns

The first data taking exercise with the CMS magnetic field was the Magnet Test and Cosmic Challenge (MTCC) during summer-autumn 2006 [16]. This challenge was the very first global exercise for CMS in which 5% of the full system was installed and operated together on the surface hall and was the predecessor to the CRAFT (Cosmic Run at Four Tesla) exercises that will be described in more detail.

CRAFT exercises were two extended global data taking periods conducted in the experimental cavern with the magnetic field of the CMS detector on and with all the final systems in place: CRAFT08 ended on November 11th 2008 and CRAFT09 ended on September 1st 2009. These month-long data-taking challenges had the following goals:

- Test the solenoid magnet at nominal field (3.8 T) in-situ with the CMS experiment in its final installed configuration underground.
- Gain experience operating CMS continuously for one month.
- Collect more than 300 million cosmic triggers for performance studies of the CMS detectors.

These goals were successfully met and the cosmic muon dataset collected has proven invaluable for understanding the performance of the CMS experiment as a whole. During these campaigns, the 100% of the DT system was operational and due to its favourable location for cosmic detection, the contribution of the DT system to the global campaign was extremely relevant. During CRAFT08 370 million cosmics were collected, and 83% of these events were triggered and read out by the DT system. In CRAFT09 the collection increased to 523 million cosmics, 92% of them acquired by the DT system.

As cosmics cross the detector from top to bottom a dedicated muon configuration and synchronization was set up in the DT trigger chain. It required the coincidence of at least two chambers in the same or nearby sector without requiring that the muon tracks point to the nominal interaction point. Also the upper sectors were delayed with respect to the bottom ones to take into account the time of flight and trigger at the same bunch crossing cosmic muons crossing both top and bottom sectors.

During these data taking periods we could confirm that the DT trigger rates were very stable with time. Since the cosmic rate in the cavern underground is low, random triggers were also injected in order to stress the system to the 100 kHz maximum expected during LHC running. No problems were seen in the DT read out system when running at high trigger rate; data integrity was not affected and no backpressure or bottlenecks were detected in the read out path.

Calibration events (around 100 Hz) were also injected during data taking. In the DT system the calibration mechanism works through the so-called Test Pulses, in which signals are injected at front-end level simulating vertical tracks orthogonal to the chamber. This procedure allows performing inter-channel synchronization and it is also a useful tool to scan for dead channels in all the electronics chain. One of the goals in these campaigns was to verify that the calibration mechanism can be implemented in the around 2.5  $\mu$ s of the LHC orbit abort gap, so that no dedicated running period would be needed. After a few corrections in the timing configuration to avoid leaks outside the orbit gap, the calibration stream was operated very satisfactorily in both CRAFT exercises.

The DT read out system demonstrated high reliability and stability during this long data taking periods in which it has been operated continuously. Very few problems were seen during these data taking exercises. By CRAFT08, after one year of operation of the detector in many local and global data taking campaigns, only 1,2 % of the detector was lost due to various types of problems, which were fixed during the 2008-2009 shutdown. Main activities during this shutdown included improvement of the secondary back-up copper

connection to the Minicrate and reinforcement of the DT safety system in order to move toward a centrally supervised operation scheme.

No issues have been found in the DT read out electronics for running with the magnetic field on. The only unexpected effect observed were some problems while reading the 1-wire temperature sensors in a few ROS boards while ramping up the magnet, which was easily solved with a power cycle of the crate.

The data integrity provided by the DT read out system during these campaigns has been excellent. The number of events in which some inconsistency has been found is very low: 15 events out of 460 million. The configuration time of the DT DAQ system is below 1 minute, and very rarely (twice in CRAFT08 and twice in CRAFT09) any error in the DT read out forced to stop the data acquisition. During these campaigns it was also possible to verify that the TTS mechanism worked satisfactorily and that the DT system recovers smoothly from sporadic errors.

Figure 28 shows the percentage of errors versus run number in CRAFT09 as detected in the ROB/ROS system. These errors can be due to parts being off, lack of communication with Minicrates, transmission problems, etc. It can be seen that the number of errors is very low and that there is no dependency with the magnetic field.

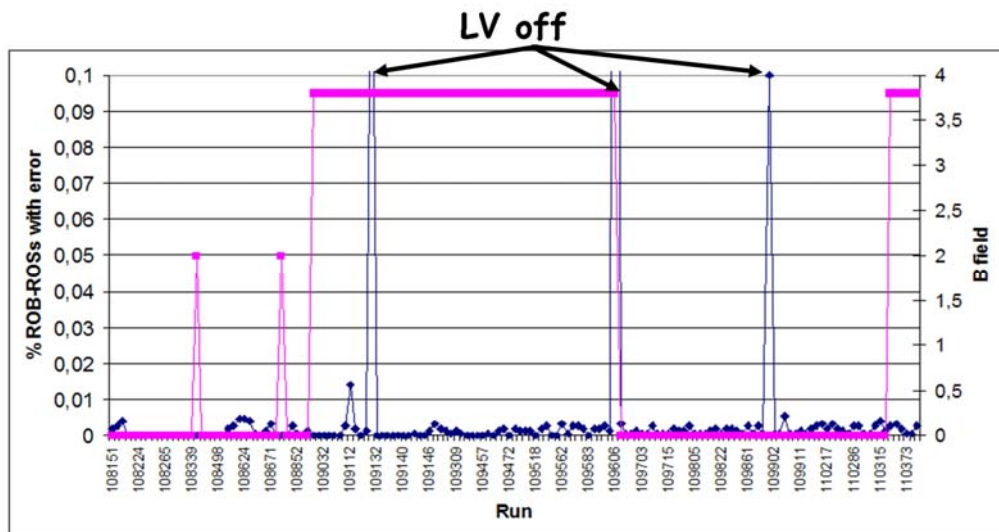


Figure 28: Percentage of errors versus run number in CRAFT09 as seen in the DT read out chain. The pink line shows the value of the magnetic field.

## 5 Conclusions

The second level of the CMS DT read out system, the ROS, is operating successfully after several years of design, production and installation. The different prototypes and the final boards have been extensively tested in laboratory and commissioned in the final system. ROS boards have shown adequate behaviour under the magnetic field and radiation environment expected during the 10 years of operation of CMS in LHC.

Additionally, Monte Carlo simulations of the DT detector occupancy and ROS processing time during LHC operation have been performed, showing that rates expected are low enough for comfortable operation of the designed electronics.

Moreover, ROS boards have behaved very efficiently and stable during long data acquisition campaigns in the last years. The quality of the data acquired during the CRAFT exercises is very good and the data integrity problems are extremely low. The cosmic data collected through this period have been very invaluable to study the performance of the detector.

Summarizing, we conclude that the developed ROS boards are reliable for operation for the expected CMS environmental conditions.

## 6 Acknowledgements

Giorga Mila, from the DT collaboration, for producing the noise studies.



## References

- [1] CMS Collaboration, *The Muon Project. Technical Design Report*, CERN-LHCC 97-32 (1997).
- [2] C.F. Bedoya et al, *First level of the CMS Drift Tubes read out system: the ROB (Read Out Board)*, CMS Note pending publication.
- [3] J. Christiansen, *High Performance Time to Digital Converter. Version 2.1*. CERN/EP-MIC (2002).
- [4] B.G. Taylor, *TTC distribution for LHC detectors*, IEEE Trans. Nucl. Sci. 45 (1998) 82.  
<http://www.cern.ch/TTC/intro.html>
- [5] J.M. Cela et al, *CMS DT Chambers Read-Out Electronics*, CMS CR 2008/018. CERN, 2008.
- [6] A. Racz, *Trigger Throttling System for CMS DAQ*, Proceedings of the sixth Workshop on electronics for LHC experiments, Cracow, 11-15 September, 2000.
- [7] *VME-BUS in Physics Conference*, CERN 86/01, Yellow Book 1986.
- [8] C.F. Bedoya, *ROS user manual v 2.0.*, [http://wwwae.ciemat.es/cms/DTE/i\\_SC.htm#ROS25](http://wwwae.ciemat.es/cms/DTE/i_SC.htm#ROS25)
- [9] GOL User Manual. [http://proj-gol.web.cern.ch/proj-gol/gol\\_manual.pdf](http://proj-gol.web.cern.ch/proj-gol/gol_manual.pdf)
- [10] QPLL User Manual. <http://proj-qpll.web.cern.ch/proj-qpll/>
- [11] IEEE Computer Society, *IEEE Standard Test Access Port and Boundary Scan Architecture*. IEEE 1993.
- [12] The CMS Trigger/DAQ Group, *CMS LI Trigger Control System*, CMS NOTE 2002/033.
- [13] P. Arce, M. Cerrada, C.F. Bedoya, J. Molina, C. Willmott, *Simulation studies for the read-out electronics of the CMS Muon Drift Tubes detectors*, CMS IN-2009/014.
- [14] A. Benvenuti, V. Genchev, *Barrel Muon System Background Simulations of the CMS TDR design*, CMS NOTE 1998/052. nota de Alberto simulación 234 tesis
- [15] M. Huhtinen and F. Faccio, Nucl. Instr. and Methods Phys. Res. A, Vol. 450 (2000) 155.
- [16] The CMS Collaboration, *The CMS Magnet Test and Cosmic Challenge (MTCC Phase I and II)*, CMS NOTE 2007-005.