Available on CMS information server

CMS NOTE -YYYY/N

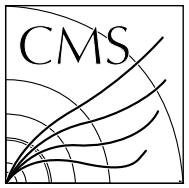

The Compact Muon Solenoid Experiment

CMS Note

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

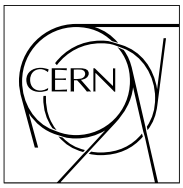

26 August 2010 (v2, 08 September 2010)

First level of the CMS Drift Tubes read out system: the ROB (Read Out Board)

C. F. Bedoya, J. M. Cela, J. Marin, J. C. Oller, C. Willmott

CIEMAT, Madrid, Spain

Abstract

The ROB (Read Out Boards) are the first level of the CMS (Compact Muon Solenoid) DT (Drift Tube) chambers data acquisition system. They are responsible for the time digitization of the incoming chamber signals in order to allow further reconstruction of charged particle tracks with a resolution of 250 um per cell. The complete design of the final boards is described in this paper, 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 Barrel Drift Tube 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 performing good muon tracking on the exterior.

Muons are measured in CMS by means of three different technologies of gaseous detectors. 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^2, 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, complemented by a set of Resistive Plate Chambers (RPC), which provides redundancy in the trigger.

The DT chambers are installed in the five wheels of the return yoke of the CMS magnet (named YB-2, YB-1, YB0, YB+1 and YB+2). Each wheel is divided in 12 sectors, each covering ~30º around the interaction point, and each sector is organized in four stations of DT chambers named MB1, MB2, MB3 and MB4, going from inside to outside, where MB stands for Muon Barrel. There are a total of 250 DT chambers in CMS. A schematic view of one CMS wheel is shown in figure 1.

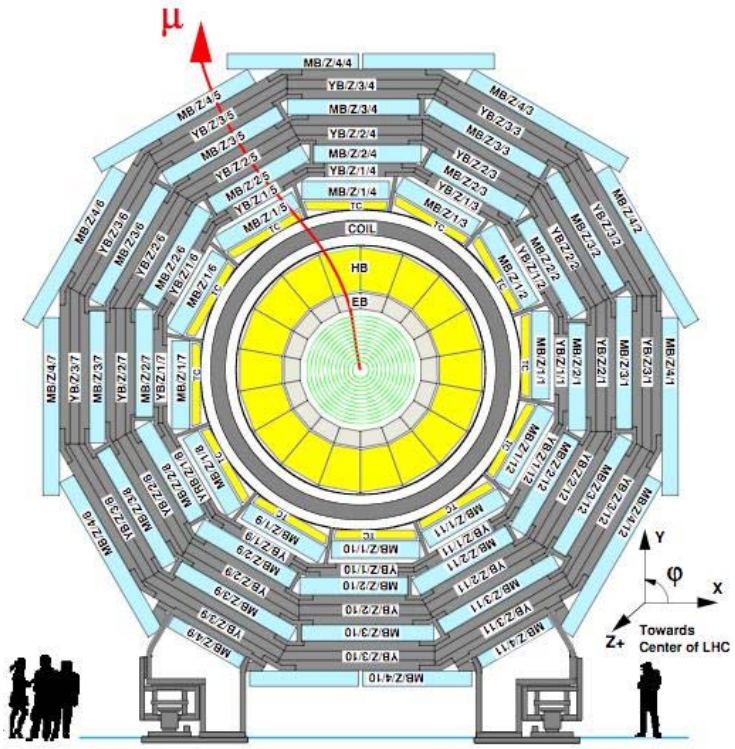

Figure 1: Transverse view of a CMS Barrel Yoke Wheel.

A DT chamber is made of three (or two in MB4) Superlayers (SL), each made up of four layers of rectangular drift cells staggered by half a tube width. The wires in the two inner and outer SLs are parallel to the beam line and provide the track measurement in the magnetic bending plane (r, Φ). In the central SL, the wires are orthogonal to the beam line and measure the θ position along the beam. The central θ-measuring SL is not present in the MB4 chambers, which therefore measure only the Φ coordinate.

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 ROBs in order 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 position resolution of 250 μm, and the 100 μm target chamber resolution is achieved by the 8 track points measured in the two (r- Φ) SLs.

2 DT Read Out Electronics

The DT read out system is staggered in several levels, merging data from all chamber channels to the CMS global DAQ system in the control room. A schematic view of the read out chain is shown in figure 2.

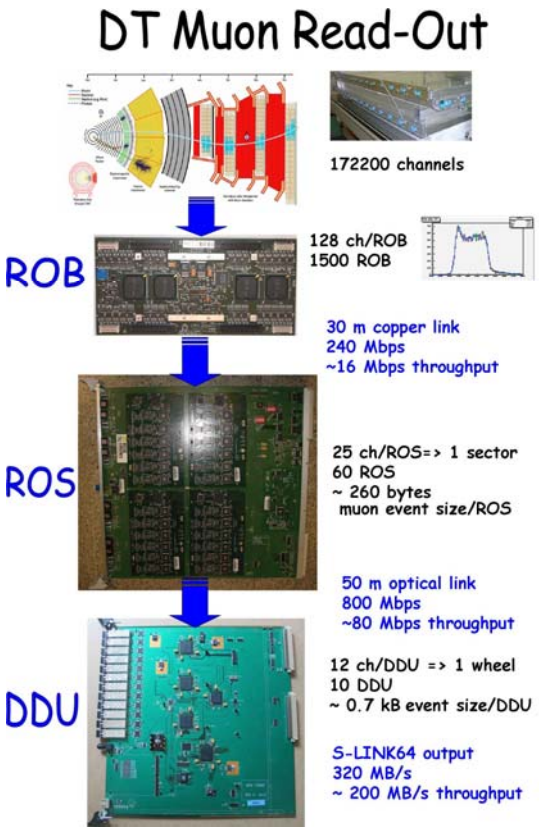

Figure 2: Schematic view of the DT read out chain.

First element in the chain after the front-end electronics is the ROB. The fundamental task of the ROBs is to digitize the chamber signals and measure their time of arrival with respect to the Level 1 Accept (L1A) trigger signal. This is done by means of the High Performance Time to Digital Converter (HPTDC) devices [2].

ROBs are located inside the so-called Minicrate, attached to the DT chambers. Each ROB performs the time digitization of 128 chamber channels. Between 3 and 7 ROB boards are installed per Minicrate in order to perform the full read out of one DT chamber, totalling 1500 ROBs in the system.

Next, the digital information 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 [3]. Each sector event is sent through a 50 m optical link at 800 Mbps to the DDU (Device Dependent Unit) [4] 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) [5].

All these steps 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 at the barrel is foreseen to be low, around 10 Hz/cm^2, the drift chambers require a system capable of performing continuous time digitization in an extended window at least the maximum drift time long (400 ns). The time resolution has to be around 1 ns in order to guarantee the required resolution of 250 μm.

Digitized chamber signals will have to be stored in memories until the L1A trigger is received and matching can be performed. Since the L1A latency can be up to 3.2 μs, buffers dimension should be properly sized.

L1A trigger rate can also reach 100 kHz due to LHC high luminosity, which imposes requirements in terms of maximum allowed processing time. Furthermore, 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.

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. The neutron fluence expected in 10 years of operation is 10^{10} cm^{-2}, being the charged particles flux 10 $cm^{-2}s^{-1}$ and the integrated dose 0.4 Gy.

Furthermore, the high magnetic fields created by the CMS solenoid and the limited space inside the CMS structure where the ROBs are located, force to dissipate the power of the electronics through a water cooling system. Accordingly, power consumption in the ROB 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 ROBs will be very restricted due to its inaccessibility. 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 ROB architecture

The ROB boards developed at CIEMAT receive the 50 ns LVDS signals coming from the front-end electronics and convert them to TTL to forward them to the TRB (Trigger Boards) connected on top. The time of arrival of hit´s rising edge is the magnitude to be measured.

ROBs are built around a 32-channel HPTDC ASIC developed by the CERN EP/MIC group. This device supplies the basic time elements to reconstruct muon tracks, that is, the relative time to a common trigger for every hit produced on chamber wires. One of its main advantages is its multi hit capability and its high programmability. On top of that, the adequate size and management of the buffers minimises the bottleneck of merging data from 32 channels per device into one common data path. In figure 3 a scheme of the buffers and registers in this TDC is shown.

The HPTDC is based on the Delay Locked Loop (DLL) principle, providing in low resolution mode, a time bin of 0.78 ns when it is clocked at the LHC 40.08 MHz frequency. Internal frequencies can be selected from a Phase Locked Loop (PLL) at the clock input, which allows higher resolution measurements if required, and also reduces the input jitter in the clock signal.

The digitized signals from each of the 32 channels are stored in the input “hit registers” and then are driven to the 256 words deep group memories common to every 8 channels (Level 1 buffers). The output of these group memories is merged into a common 256 deep memory once the trigger matching is performed. Until this operation can be done, the pending triggers are stored in a 16 deep FIFO. Despite its complexity, matching is done in a rather fast way, selecting from memory hits according to certain programmed parameters.

One of the main features of the HPTDC is the ability to handle overlapping triggers, as an individual hit may fall inside several time windows due to the high bunch crossing frequency that exceeds widely the cell drift time. In order to do that, the HPTDC does not erase hits from memory when matching is performed, but only when they are older than a programmed time which guarantees that they cannot belong to any further event. This mechanism will be explained in more detail in section 3.3.

Other remarkable features of this TDC are the JTAG programming capability, flexible signal interfaces and various read out modes. It also has internal counters for event number identification and for bunch crossing identification. This information can be attached to the output data with each event, together with programmable headers, error flags and debugging information.

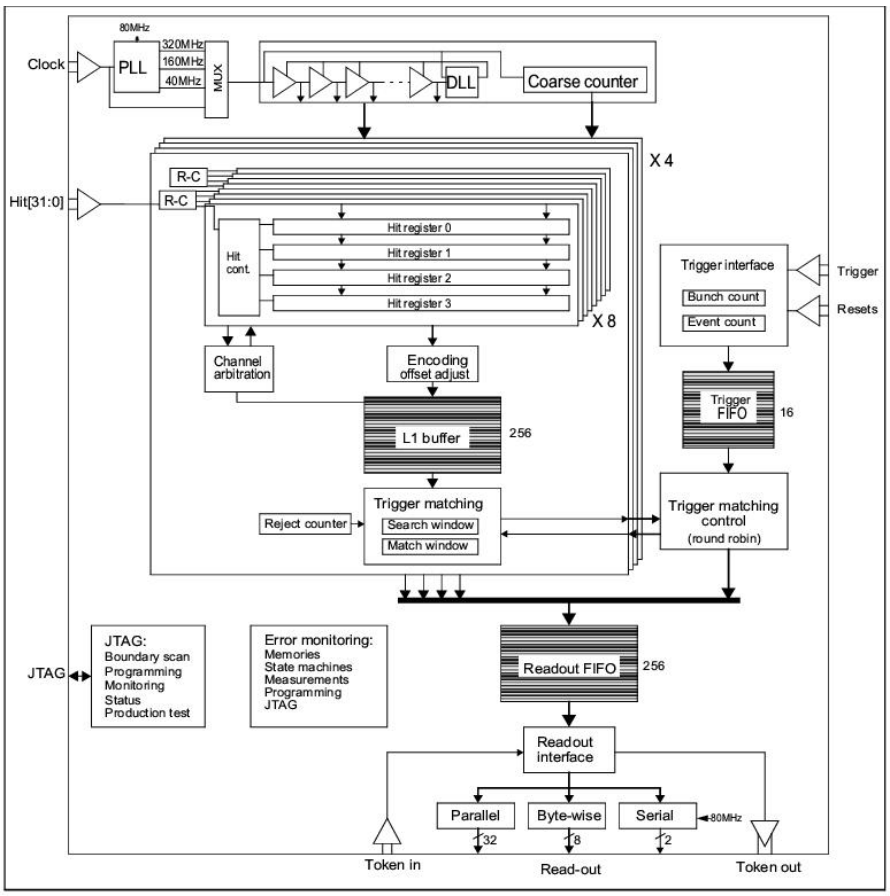

Figure 3: Architecture of the HPTDC.

Each ROB houses 4 HPTDC in a 10 layers 22.6 x 9.8 cm PCB. The number of HPTDCs per board has been optimized according to the available space, number of existing channels per chamber and estimated bandwidth. The board includes several ground and power planes to ease thermal conduction, which will be done through the metallic strips that can be seen in figure 4, which are in contact with the Minicrate structure.

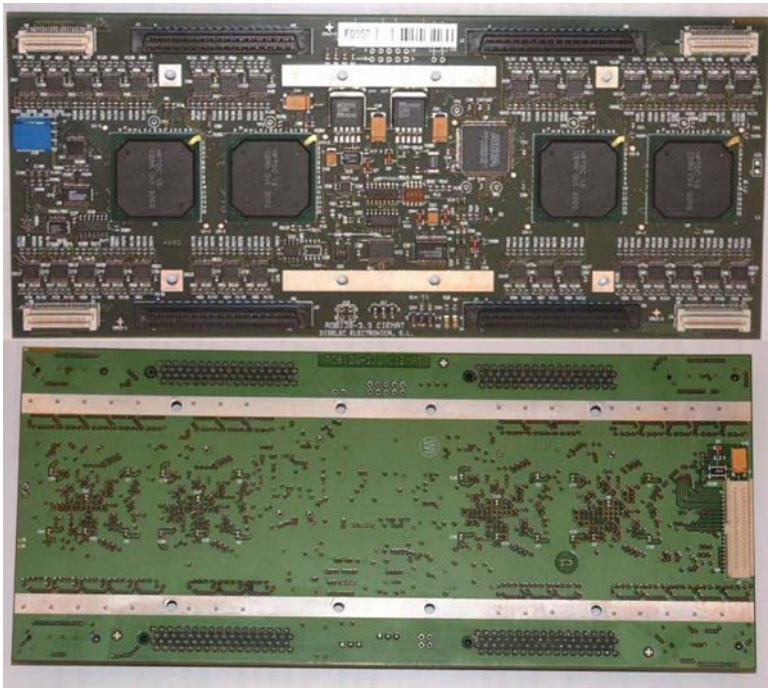

Figure 4: Front and back image of the Read Out Board.

The ROB board includes a 3.3 V and a 2.5 V regulator, reversed power supply protection and a power supply protection circuitry with fast shut off capability to avoid over-currents. These over-currents may be generated not only by electrical failures, but also by radiation induced phenomena. This protection circuitry is completely automatic and in case the currents exceed the values specified in table 1, both regulators are switched off and a fault flag is signalled to the slow control system. In case of no action from the slow control, the circuit will automatically try to power on again after 700 ms. If the over-current persists, it will keep on repeating the power on/off cycles. In such way, the safety of the system is guaranteed, obtaining an average current consumption of 42 mA in case of short circuit in the ROB. Current consumptions in normal operation are detailed in table 1, being the average power consumption of each ROB of 3.7 W.

Table 1: ROB current consumption for the different power supplies (first column) and maximum current allowed by the automatic protection circuitry.

	I (A)	I_{max} (A)
3.3 V	0.5	1.74
2.5 V	0.5	1.8

ROB boards also include an on-board sensor that allows to measure the 3.3V and 2.5 V power supply, the 2.5 V current and the board temperature. This sensor, which includes a serial number to uniquely identify each board remotely, is read through a *1-wire* interface through the slow control system.

3.1 Read out protocol

The four HPTDCs are connected inside the ROB in a clock synchronous token ring scheme for the read out of the digital data (figure 5). One of the HPTDCs is configured as master, controlling the token that authorizes data transmission to a common bus connected to a 240 Mbps serializer (DS92LV1021). A failsafe bypassing mechanism has also been implemented in order to guarantee operation of remaining TDCs in case of one failing.

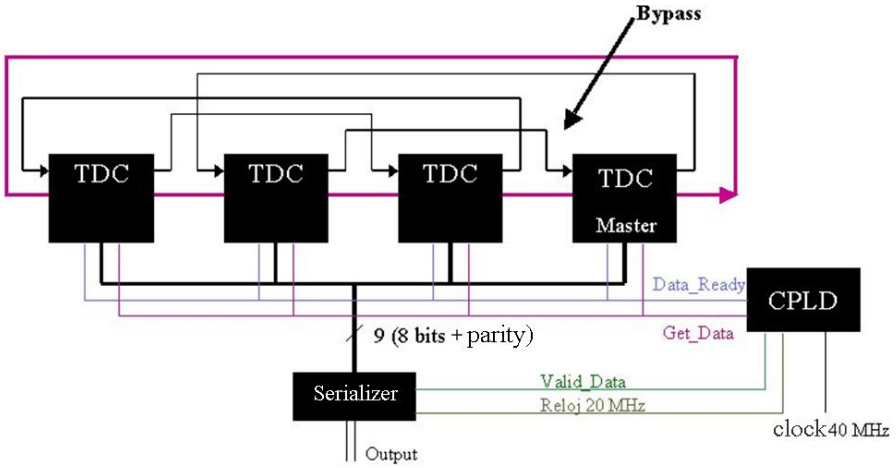

Figure 5: Diagram of the HPTDCs read out interconnections in the ROB.

An Altera EPM7128AE CPLD manages the data_ready/get_data read out protocol, controlling the handshake with each of the HPTDCs and the serializer. HPTDC is operated in byte-wise mode at 20 MHz, being the effective bandwidth of the link 160 Mbps for an estimated throughput of 16 Mbps. Data is sent through an approximately 30 meters copper link to the next level of the read-out system, the ROS boards. The reliability of this link has been measured and the tests showed less than 10^{-15} bit error rate. Parity checks are performed on each byte at the receiver, as part of the data integrity checks when building the HPTDC 32 bit words. In normal operation, a ROB event consists on the following:

- one header from the master HPTDC which contains event and bunch crossing identification,
- one time measurement word for each of the leading edges detected that satisfy the trigger matching requirements. This word includes the identification of the HPTDC and the channel that received the hit. In low resolution mode, the time measurement is obtained from the first 17 bits of the leading time field in TDC bins (clock period divided by 32).
- in case any HPTDC is in error state, an error word is included describing main error flags.

- one trailer from the master HPTDC which contains again the event identification and a word count of the number of 32 bit words transmitted in that event.

HPTDC header 32 bits word.

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	TDC				Event ID												Bunch ID											

HPTDC time measurement 32 bits word.

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	1	0	0	TDC				Channel					Leading time																		

HPTDC error 32 bits word.

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	1	1	0	TDC													Error flags														

HPTDC trailer 32 bits word.

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	1	TDC				Event ID												Word count											

Figure 6: Main HPTDC 32 bit words and its data format.

Besides controlling the read out protocol, as well as other minor tasks, the Altera CPLD includes the logic to operate the ROB in a special operation mode called Test Pulses which is a basic aspect of the DT calibration mechanism. In this mode of operation, simultaneous signals are injected at front-end level simulating vertical tracks orthogonal to the chamber. This procedure allows performing inter-channel synchronization by comparing the different time measured in all channels in the same chamber. Moreover, it is also a useful tool to scan for dead channels in all the electronics chain, independently from gas or high voltage problems. The Altera CPLD receives the Test Pulse control signals and controls the state machine that allows correct channels masking for each Test Pulse event, sweeping all chamber channels. In the following figure it can be seen an example of the single channel time histogram in one Test Pulse run, together with an event display.

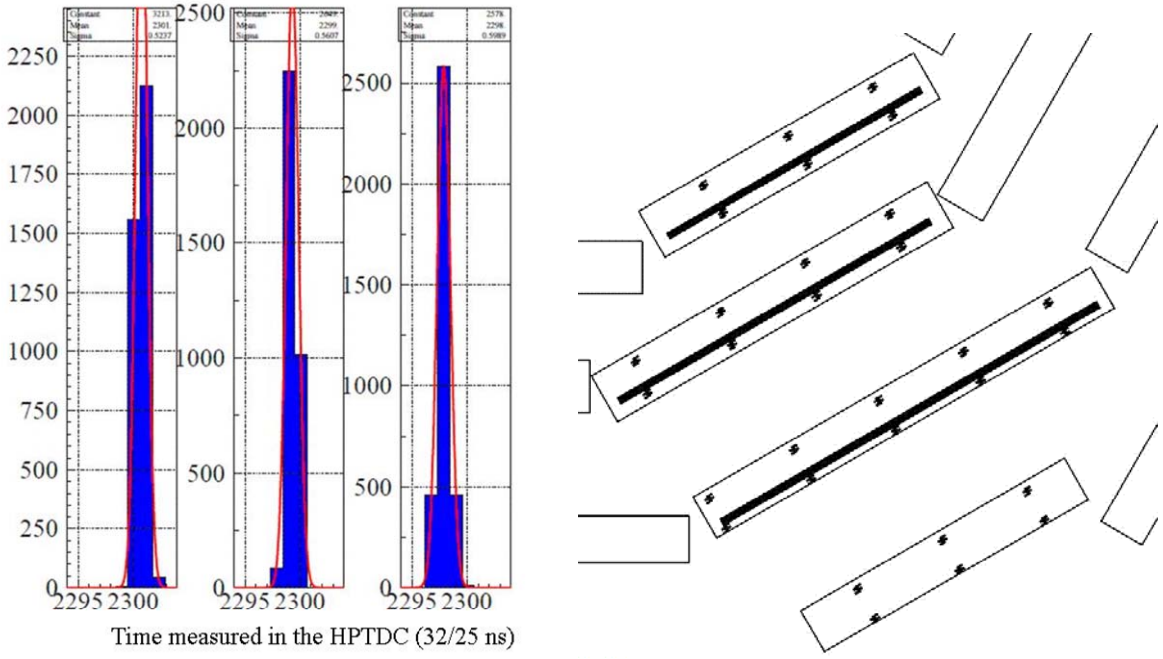

Figure 7: On the left, time histogram of the Test Pulse signals in three HPTDC channels. On the right, event display of one Test Pulse event.

3.2 Integration in the Minicrate

Chamber read out and trigger electronics are lodged in the Minicrates, mounted on the front side of each chamber, inside the C profile surrounding the honeycomb panel which provides stiffness to the chamber and a lever arm between the two Φ Superlayers for a better track fitting.

The Minicrate is an aluminium profile with large number of support structures to provide sound construction and thermal conduction for refrigeration of the electronics through a water cooling system. Minicrates are 10.5 cm wide and 5.5 cm high; their length varies from 107.1 cm to 207.5 cm. The weight of a complete Minicrate also varies between 8 and 16 kg.

In the Minicrate, the ROBs are integrated together with the muon trigger and slow control electronics. They share, among others, power supplies, water cooling, wire chamber signals and TTC signals (LHC clock, L1A, bunch reset, etc.). Inside each Minicrate, a layer of trigger boards (TRB) (indicated in figure 8) is mounted on top of a layer of ROBs while a central Server (SB) collects the trigger information. A Chamber Control Board (CCB) is located below the SB for communication with the external slow control system. A ROlink board is placed on one side of the Minicrate to serve as patch panel between ROBs output and the cables to the ROS.

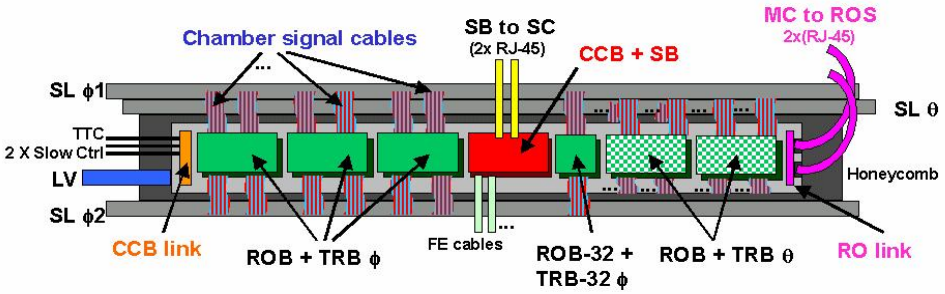

Figure 8: Diagram of a Minicrate with its main components.

With a total number of 250 Minicrates, there are up to 20 types, depending on the station type and sector. Accordingly, the number of ROBs varies, as can be seen in figure 9, for the different types of Minicrates. ROB boards are equivalent for Φ or θ chamber channels, and the only distinction is the ROB-32 in the MB1 Minicrates, that houses only one TDC due to mechanical constrains.

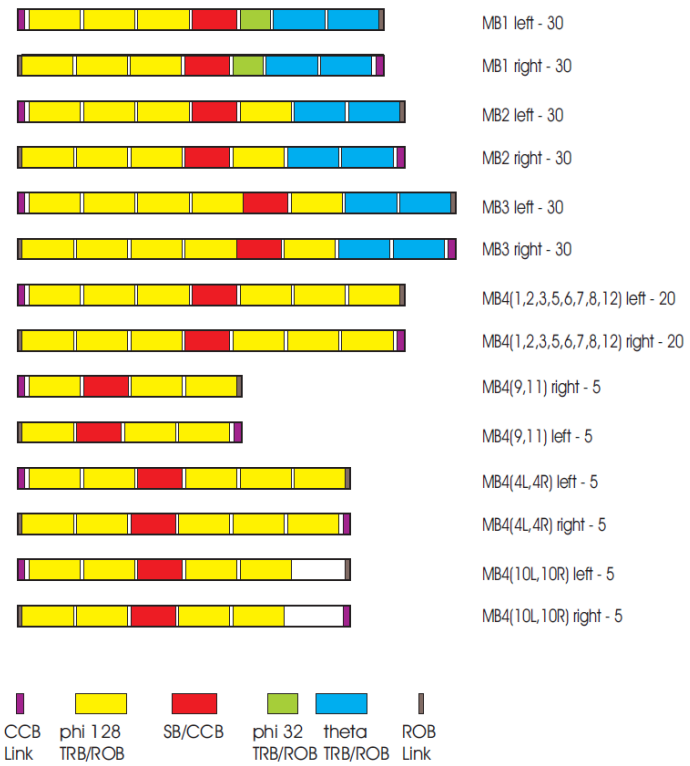

Figure 9: Diagram of the different types of Minicrates and the different distribution of the boards.

ROBs are connected to the CCB for configuration and monitoring through a parallel cable called ROBUS. This bus carries independent power on signals and address lines for each board, the JTAG interface lines for communication between CCB and ROBs, fast signals such as the L1A, bunch reset or event reset, *1-wire* signal for temperature, power and current monitoring, a power fault line and the Test Pulse signals. ROB internal errors, as well as buffer overflows, are reported to the CCB like other system faults and, additionally, they are also notified within the read out data flow.

LHC clock is distributed through a separate line from the CCB to each ROB to minimize interferences. ROB cable clock lengths have been selected to guarantee proper sampling of the signals throughout the ROBUS length. TRB clock distribution also follows independent lines. In this case, cable lengths have been optimized to maximize trigger efficiency, compensating front-end cables lengths.

Minicrate power consumption is detailed in table 2. The 5 V power is distributed directly to the CCB boards, while the 3.3V that feeds all ROBs and TRBs is distributed through copper bars fixed in the bottom part of the Minicrate. The total power consumption in one CMS wheel is around 4 kW. The Minicrates are cooled by water flowing in tubes extruded in the Minicrates´s aluminium profiles along the full length. As can be seen in figure 10, the cooling system is effective enough and the ROBs installed in the final CMS system maintain a constant temperature of around 24 ºC, with minimal variations between different Minicrates.

Table 2: Current and total power consumption for the different types of Minicrates.

	I (A) 5 V	I (A) 3,3 V	P (W)
MB1	1.4	21	76
MB2	1.4	23.5	85
MB3	1.4	26.7	95
MB4 (1,2,3,5,6,7,8,12)	1.4	22	80
MB4 (4)	1.4	19.3	71
MB4 (10)	1.4	15	57
MB4 (9,11)	1.4	12.3	48

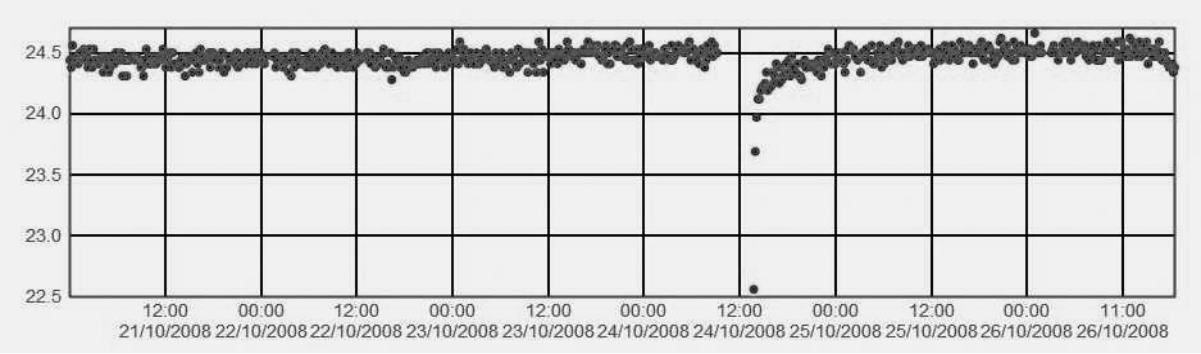

Figure 10: Temperature measured in a ROB installed in the final CMS detector. The drop in temperature shows the effect of switching off and on the Minicrate.

The production and assembly of the 250 Minicrates was a shared task carried out by CIEMAT, INFN Padova and INFN Bologna. All parts of the DT read out and trigger electronics were extensively tested before and after the installation in the chambers. The assembly and testing of the slow control and read out part of the Minicrates took place at CIEMAT between May 2004 and March 2006. A Minicrate during the testing procedure can be seen in figure 11.

3.3 Operation of the ROB board

In order to reconstruct a muon track, the following aspects should be taken into consideration:

- Interesting muons will be produced in LHC bunch crossings and therefore, synchronously with LHC clock that is distributed to the whole electronics through the TTC system.
- Different TTC fibers lengths are compensated on each DT chamber at the TTCrx device [6] inside the CCB. Therefore, L1A signals will arrive to each chamber simultaneously.
- The position of each wire and cell inside the CMS detector is known with a precision of 150 μm both

from construction and from the CMS alignment system measurements. Therefore, the measurement to be provided by the ROB, the drift time, corresponds to the distance between the muon track and the cell anodic wire.

- The L1A latency is fixed with respect to the corresponding collision for any event, independently from the part of the detector that generated the trigger, i.e., the L1A signal always arrives to the ROBs after the same number of clock cycles after the bunch crossing.

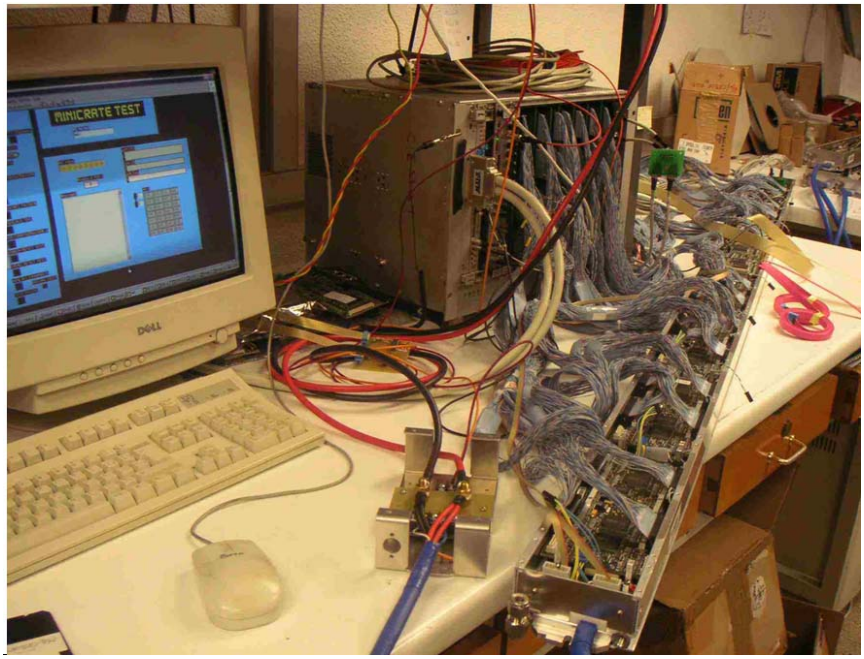

Figure 11: One Minicrate during the test procedure after assembly at CIEMAT.

Accordingly, the trigger matching mechanism consists in searching for stored hits that fall inside the appropriate time window for each L1A. There are two basic parameters to configure: the match window, which establishes the time interval within which hits will be related to a particular L1A, and the effective L1A latency, which represents the time that will be subtracted in the HPTDC to the L1A arrival time in order to determine the beginning of the match window. These parameters are configured through some particular fields in the HPTDC registers as indicated in what follows for our operation mode:

- Match window: chosen to accommodate maximum drift time.
- Search window: since hits are not stored in strict temporal order inside the HPTDC buffers, search mechanism will not stop until hits outside a search window, wider than the match window, are encountered.
- Roll Over: HPTDC bunch crossing counter is 12 bits, while each LHC orbit has 3564 clock cycles. Bunch crossing counter roll over value should be configured properly (3563) to allow proper subtraction of L1A arrival time at the beginning of the orbits.
- Trigger offset: can be used to compensate a different arrival time between bunch counter reset and L1A signals. In our system it should be equal to the Roll Over value.
- Coarse offset: with all other values configured as described, the coarse offset will represent the effective trigger latency in the HPTDC.
- Reject offset: it will determine when hits found in memories are older enough to belong to any further L1A and can be erased.

The diagram in figure 12 describes the trigger matching mechanism in the HPTDC. The different time contributions to be taken into account for the calibration of the time measurements are also indicated.

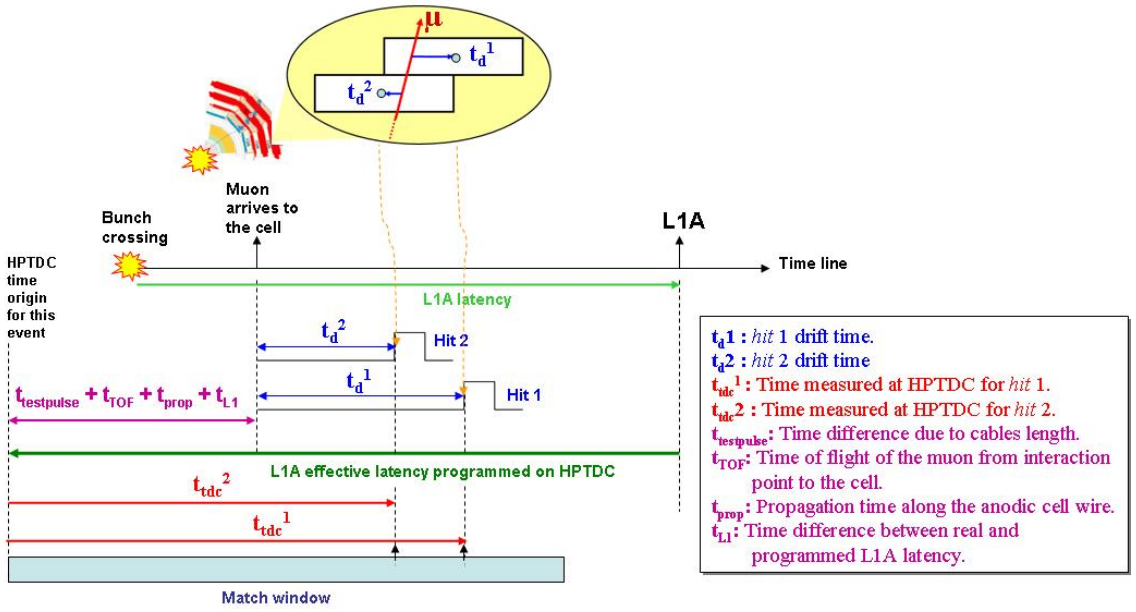

Figure 12: Diagram of the matching mechanism in the HPTDC and the different time contributions in the measurements.

The trigger overlapping mechanism is described in figure 13, where two muons from two different bunch crossings go through neighbouring cells inside the same HPTDC. When the first L1A signal arrives (L1A μ_A) both hits from the first and from the second muon will satisfy the matching criteria and thus, all of them will be read out. However, they will not be erased from the memories to allow the second L1A (L1A μ_B) to encounter the hits that belong to its muon. Hits will only be erased from the memories when they are older than the reject time, ensuring that they cannot belong to any further L1A. At the arrival of the second L1A (L1A μ_B), the 8 hits will be matched again and read out. Hence, hits measurements may be duplicated, but that will guarantee that no information is lost in any event.

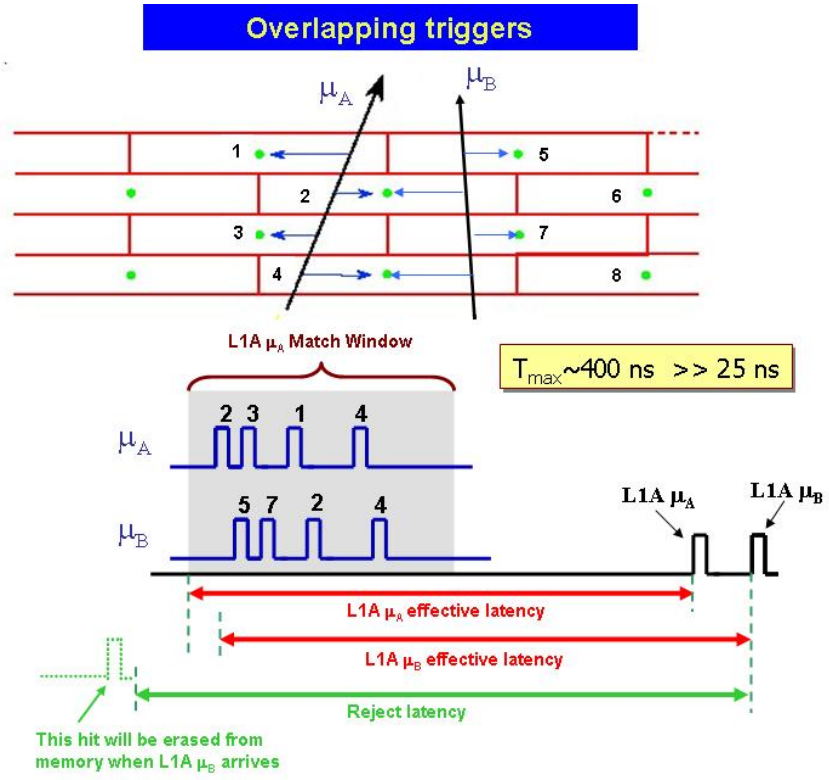

Figure 13: Diagram showing how the DT read out system handles overlapping triggers.

Duplicated information will not create a problem during track reconstruction since only hits that belong to the proper L1A will satisfy the Mean Timer algorithm [7]. An example of this is shown in the following figures for one muon crossing one Superlayer that is read out twice, one for the correct L1A (figure 14) and one for a L1A that arrived 100 ns later (figure 15). On the other hand, this extra read out information has to be taken into account for the design of the links bandwidth, since it will increase significantly the data throughput.

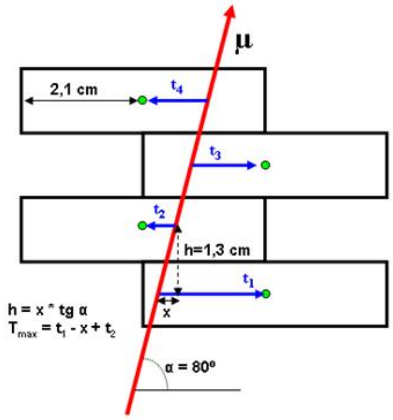

	Time (ns)	Distance to the wire (cm)
t_1	294,66	1,60
t_2	134,30	0,73
t_3	210,23	1,14
t_4	218,72	1,19

$$MT_1 = \frac{t_1 + t_3 + 2 \cdot t_2}{2} = 386{,}74ns$$

$$MT_2 = \frac{t_2 + t_4 + 2 \cdot t_3}{2} = 386{,}74ns$$

Figure 14: Representation of one muon track crossing one Superlayer. The calibrated time measurements and the corresponding distance to the wire for a drift velocity of 54.3 μm/s are indicated in the table. Both calculated Mean Timers are correct and equal to the maximum drift time.

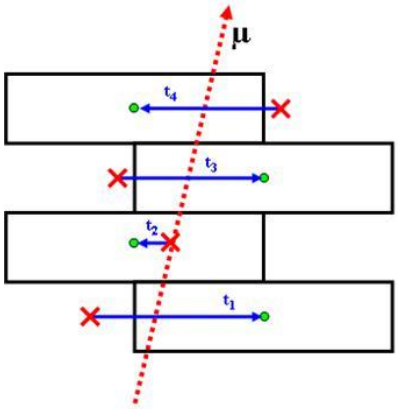

	Time (ns)	Distance to the wire (cm)
t_1	194,66	3,58
t_2	34,30	0,63
t_3	110,23	2,03
t_4	118,72	2,19

$$MT_1 = \frac{t_1 + t_3 + 2 \cdot t_2}{2} = 186{,}74ns$$

$$MT_2 = \frac{t_2 + t_4 + 2 \cdot t_3}{2} = 186{,}74ns$$

Figure 15: The same muon hits from the event presented in figure 14 are read again by a L1A arriving 100 ns later. In this case, the calculated distance to the wire does not produce a reconstructed track. Similarly, the Mean Timer values calculated provide an incorrect value. These hits can be discarded from this event.

4 ROB characterization and validation tests

4.1 Time measurement tests of the ROB

Various tests have been carried out in the laboratory to characterize the time measurement performance of the ROB and to guarantee it is adequate for operation in the CMS DT chambers.

4.1.1 ROB linearity measurements

The differential and integral non linearity (DNL and INL) of the HPTDC assembled in the ROB has been measured using the CDT (Code Density Test) method [8]. A set-up was mounted as shown in figure 16: TTCvi and TTCex boards [9] to generate the clock, L1A signal and other TTC commands, a CCB board for the slow control, a ROS board to read out the data through a VME interface and a Tektronix AFG3252 pulse generator whose output was injected in one of the ROB channels after proper level conversion. The generated pulse had a

frequency of 2.789 MHz, low enough not to have more than one pulse per event but high enough to maximize sampling efficiency. A slow rising edge of the pulse (~10 ns) that created an indetermination in the sampling point, and the own drift of quartz crystals, guaranteed that the pulse was asynchronous with respect to the ROB clock.

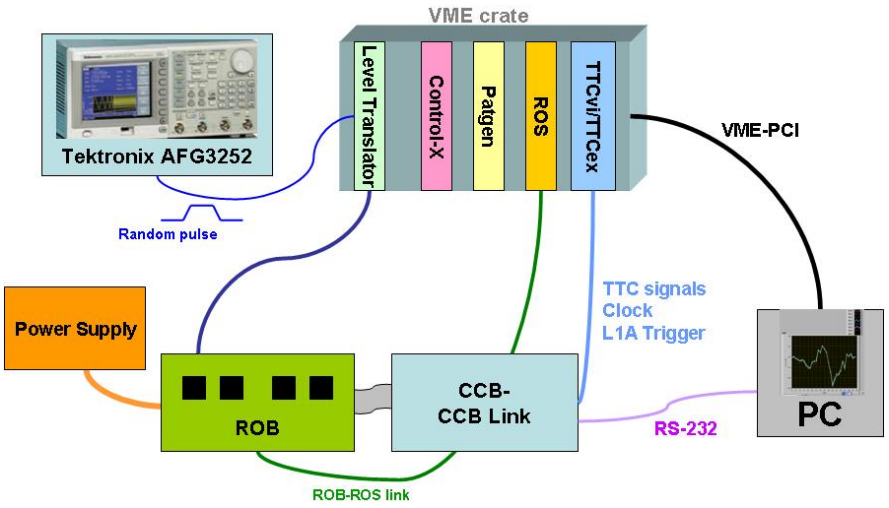

Figure 16: Set-up built to perform the DNL and INL measurements in the ROB.

In figure 17 the time histogram for 690,000 measurements per bin show the stability all along the matching window programmed in the HPTDC, which depends mainly on the stability of the clock signal to the ROB, since the coarse measurement comes from a counter. The various structures corresponding to the 32 delay elements of the HPTDC DLL can be clearly seen. The total difference among these 32 bin structures is smaller than 0.41%.

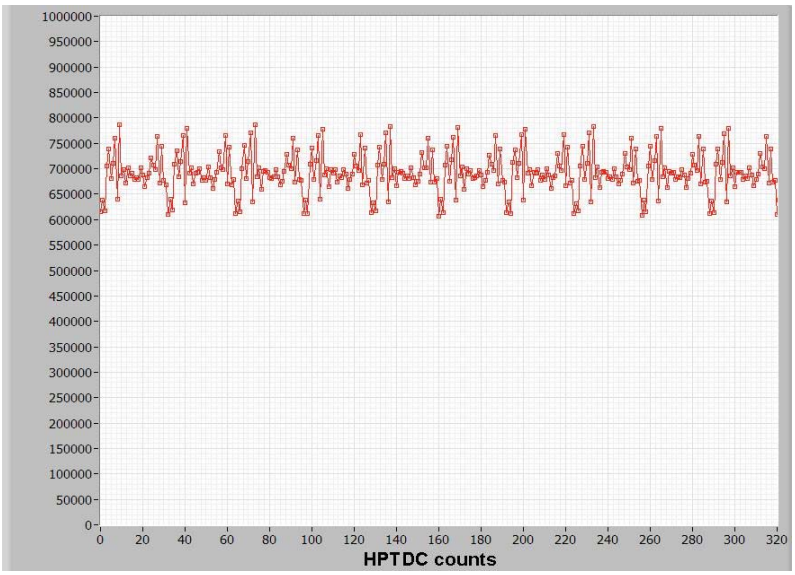

Figure 17: Time histogram for 690,000 samples per bin of a random signal measured in the HPTDC in the ROB.

In what follows, unless otherwise specified, the results correspond to 20,000 samples per HPTDC bin which provides a precision of 0.02 LSB with a 99% confidence in the DNL calculation. The results obtained (figures 18 and 19) are very good, with DNL inside the range ±0.2 LSB and INL within ±0.25 LSB. The root mean square values are reported in table 3, being smaller than the ones stated in the HPTDC specifications. The measurements were repeated for different channels (figure 20), obtaining very similar results, with maximum differences smaller than 2.5%.

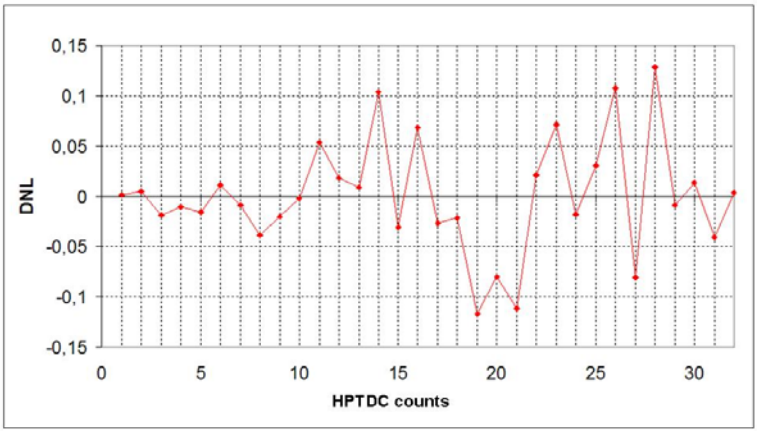

Figure 18: DNL (Differential Non Linearity) measured in the HPTDC in the ROB.

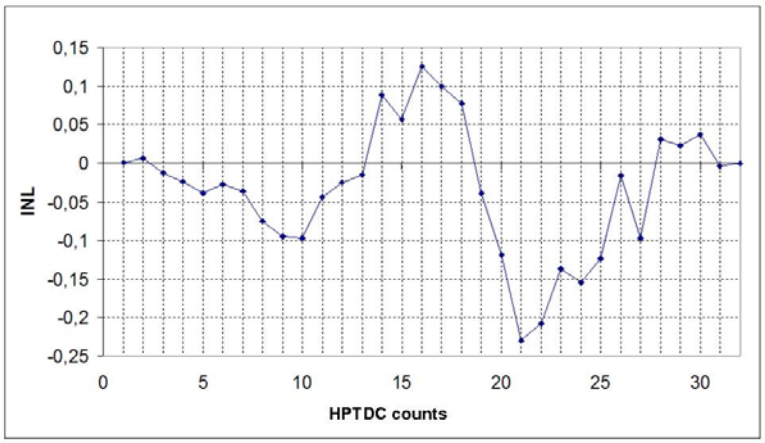

Figure 19: INL (Integral Non Linearity) measured in the HPTDC in the ROB.

Table 3: Comparison between the DNL and INL measured in the ROB and the ones indicated in the HPTDC specifications.

	Measured in the ROB	HPTDC specifications
DNL rms	0.06 LSB	0.08 LSB
INL rms	0.08 LSB	0.11 LSB

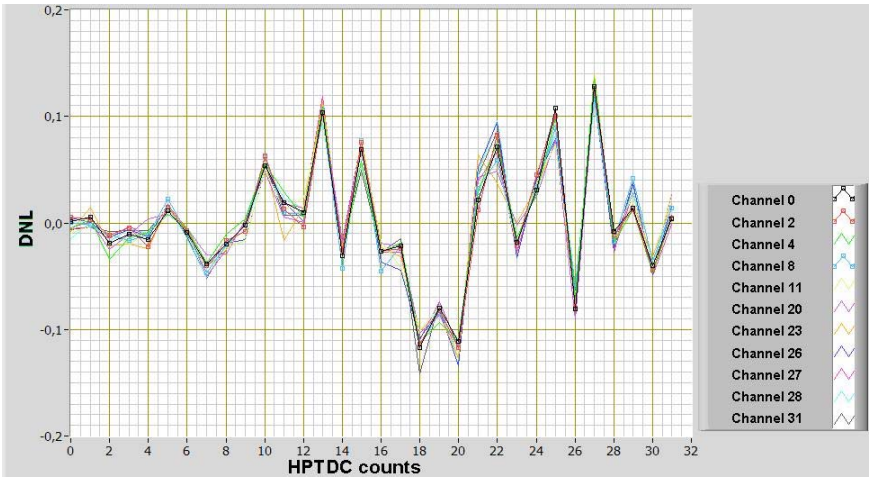

Figure 20: Comparison of the measured DNL for several ROB channels.

4.1.2 ROB channels crosstalk

The channel crosstalk was also measured to validate ROB PCB design, which was done with special care to minimize interferences. The set-up (figure 21) was similar to that of the previous test, though in this case a custom board developed at CIEMAT, Control-X, was used to generate a T0 signal, which was injected in one of the ROB channels, and a L1A signal at a fixed time with respect to T0 in order to insure its detection in the ROB. A random pulse T1 coming from the AFG3252 was being sampled as before in a nearby channel. When T0 and T1 are injected in channels far away from each other there is no substantial effect in the measured DNL (accumulated differences in both DNLs through all the DLL range is smaller than 0.41%).

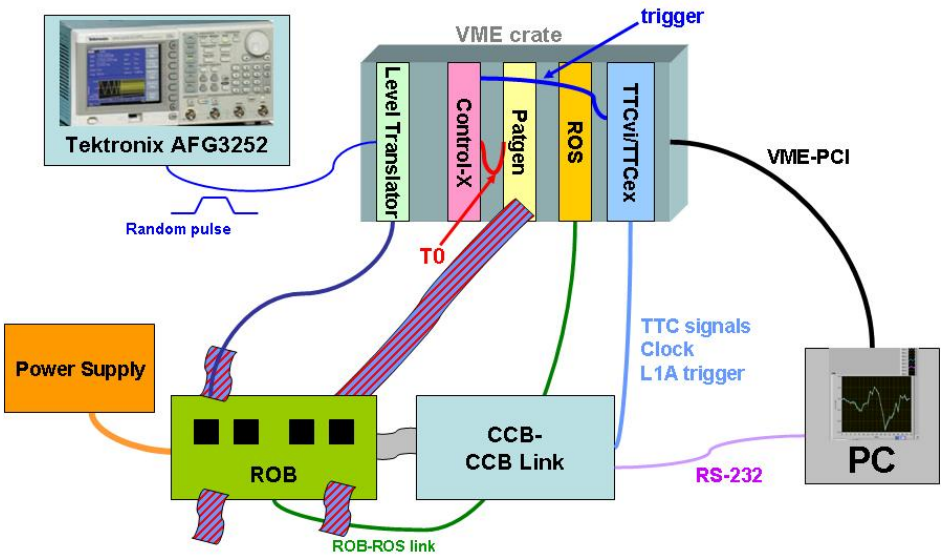

Figure 21: Set-up built to perform the channels crosstalk measurements in the ROB.

In figure 22 the DNL of T1 when T0 is injected in channel 3 and T1 in channel 4 of HPTDC 0 is shown. These channels are neighbours in the HPTDC but go through different LVDS receivers in the ROB. The DNL has been calculated along the matching window and it has been folded every 32 bins to highlight the affected region. As it can be seen, in the blue line, corresponding to the region where T0 was also being sampled, some effect is present due to the interference with T0. The DNL variation is shown on the right, showing maximum differences of 0.095 LSB which corresponds to 74 ps.

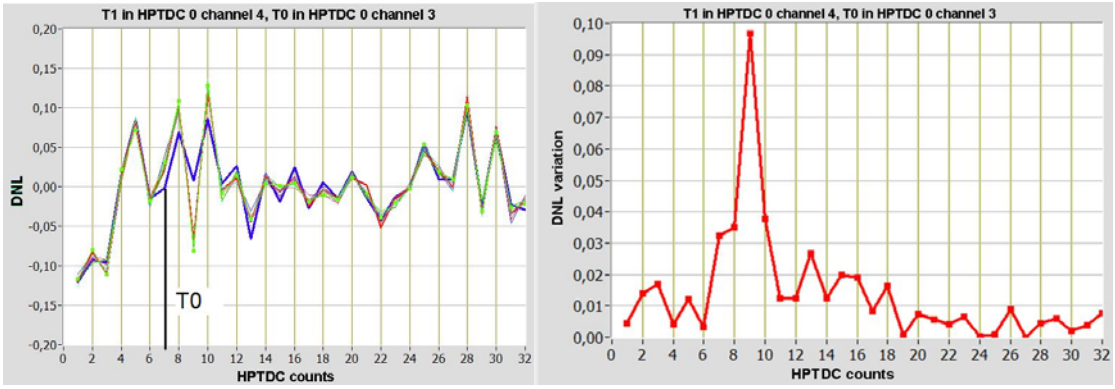

Figure 22: Crosstalk effect of T0 (channel 3) in the measured DNL in channel 4 of HPTDC 0. The DNL over the full matching window has been folded each 32 bins to show the effect in the T0 region (left). The DNL variation of that region with respect to non affected regions is shown on the right plot. These two channels go through different LVDS receivers.

In figure 23 the effect in the DNL when both signals are injected in neighbour channels that go through the same LVDS receiver is shown. It can be seen that in this case (red line), the effect is much larger, around 0.13 LSB, that is, 101 ps. The result is identical if T0 is masked inside the HPTDC. Thus, most of the interference happens before entering the HPTDC ASIC.

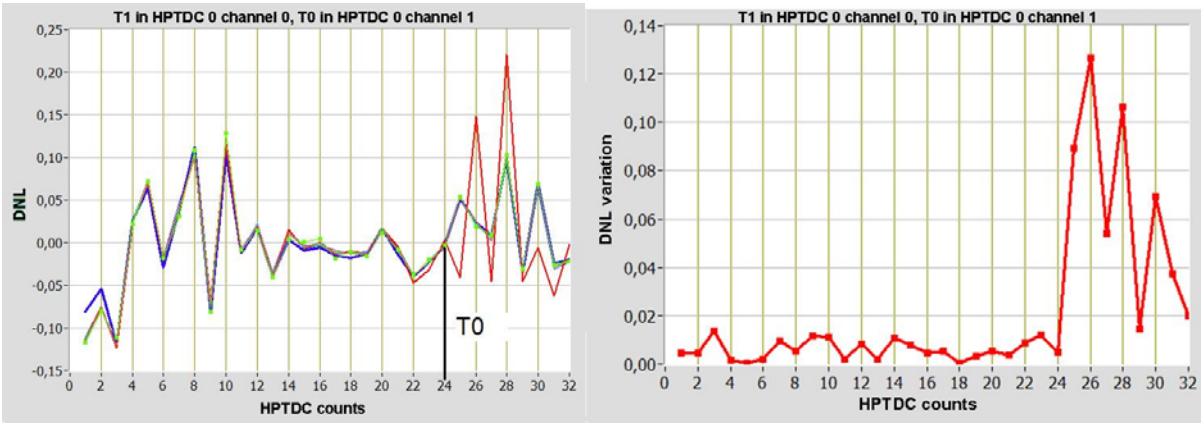

Figure 23: Crosstalk effect of T0 (channel 1) in the measured DNL in channel 0 of HPTDC 0. The DNL over the full matching window has been folded each 32 bins to show the effect in the T0 region (left). The DNL variation of that region with respect to non affected regions is shown on the right plot. These two channels go through the same LVDS receiver.

Finally, the crosstalk effect of 31 channels to one was measured injecting T0 in all HPTDC channels except channel 4, where T1 was injected. The results can be seen on figure 24. In this case the effect is large, up to 0.37 LSB (289 ps) and it disturbs the time measurement for up to 23 bins until the DLL feedback is capable of restoring the original situation. This value obtained is larger than the 150 ps indicated in the HPTDC specification, but either case it is smaller than the required resolution for the DT chambers and, also, the probability of getting 32 simultaneous signals in CMS negligible.

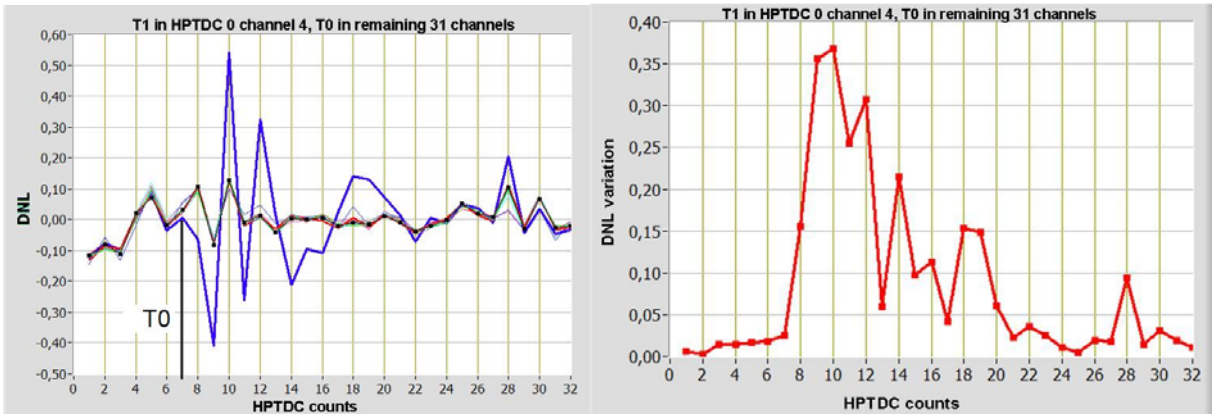

Figure 24: Crosstalk effect of T0 (all other 31 HPTDC channels) in the measured DNL in channel 4 of HPTDC 0. The DNL over the full matching window has been folded each 32 bins to show the effect in the T0 region (left). The DNL variation of that region with respect to non affected regions is shown on the right plot.

4.1.3 ROB time resolution

The time resolution of the ROB board was also measured with a very similar set-up as the one in figure 21. Two signals asynchronous with the ROB clock were injected in two different HPTDC channels where no crosstalk effect was foreseen. The two signals were generated with the AFG3252 and their separation could be controlled with a precision of 100 ps (a minimum separation of 100 ns was chosen). Since the generation of these signals was uncorrelated with the L1A, only those events in which both signals were present were selected. The position within the matching window of both signals is random, but their time difference is a fixed value whose dispersion depends only on the HPTDC resolution.

For each 100 ps delay step, 2000 measurements were taken until the 25 ns range was covered. By taking the most likely value, the staircase can be built as seen in figure 25. The difference with respect to the linear fit corresponds to the HPTDC quantization errors.

For a single measurement, the uncertainty is determined by this quantization error; however, when sampling 2000 times the same signal, statistical fluctuations will provide higher resolution. Representing the standard deviations of each of those samples, the diagram of figure 26 is obtained. In the ideal case, this diagram appears as a replica of parabolas that assume null values for time intervals multiple of 1 LSB and maximum value

(0.5 LSB) just in the middle between the LSB multiples. Calculating the mean value of these standard deviations, the HPTDC resolution is obtained. The values compared with the ones in the HPTDC specifications and with the theoretical limit are shown in table 4.

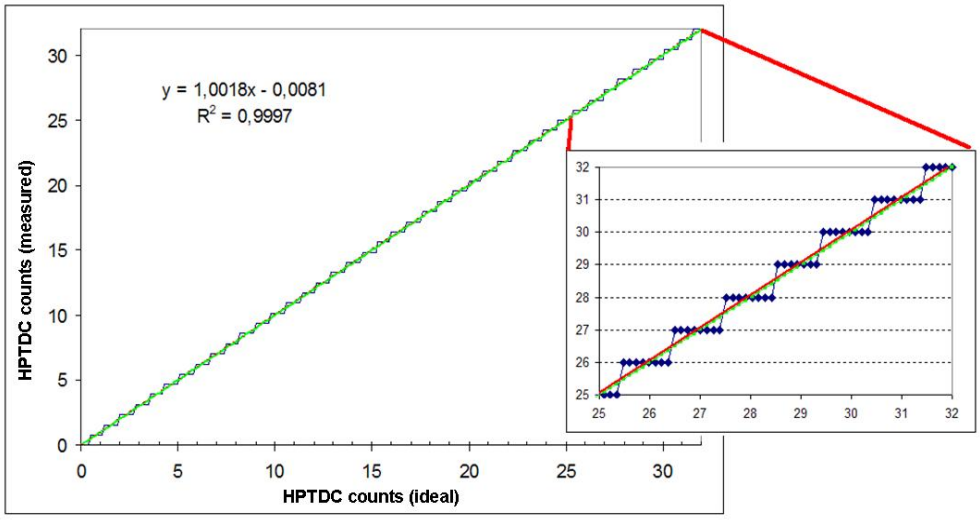

Figure 25: Measured staircase of the HPTDC in the ROB.

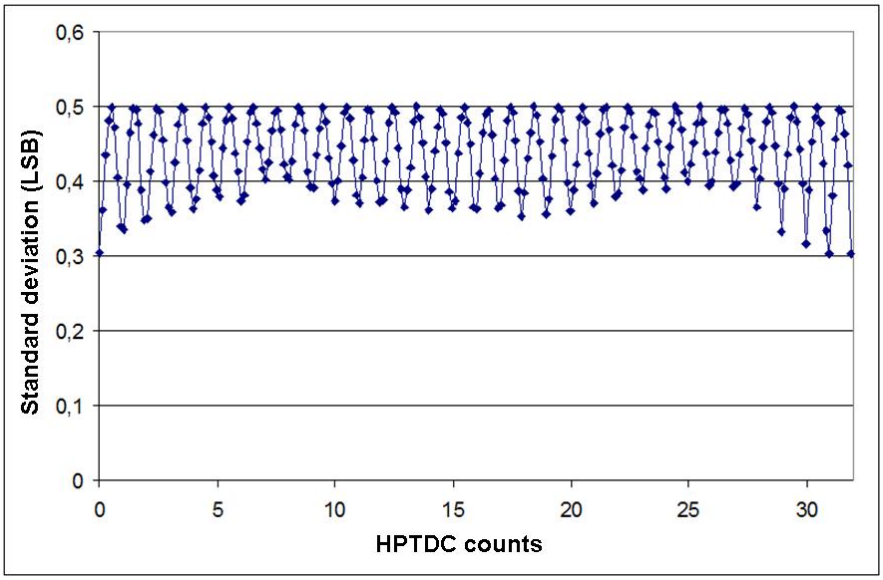

Figure 26: Experimental error standard deviation curve as a function of the measured time interval in the HPTDC.

Table 4: Time resolution of the HPTDC in the ROB.

	Measured in the ROB	HPTDC specifications	Theoretical Limit
Resolution (LSB)	0.309 LSB	0.34 LSB	0.289 LSB
Resolution (ns)	0.241 ns	0.265 ns	0.225 ns

4.2 ROB temperature tests: time measurement variation, life-time tests, burn-in

With the aim of analyzing the stability of the time measurement in the ROB, temperature tests have been performed in an environmental chamber where a ROB could be fully operated and monitored (figure 27).

Figure 27: Environmental chamber employed for thermal cycles and lifetime studies with the test stand for ROB operation.

Temperature cycles (0.2ºC/min) between 0º and 70ºC, which are the temperature limits of some of the devices, have been performed (figure 28), showing that the maximum time shift is 1 ns between 0º and 70ºC, with a maximum slope of 45 ps/ºC around 40ºC.

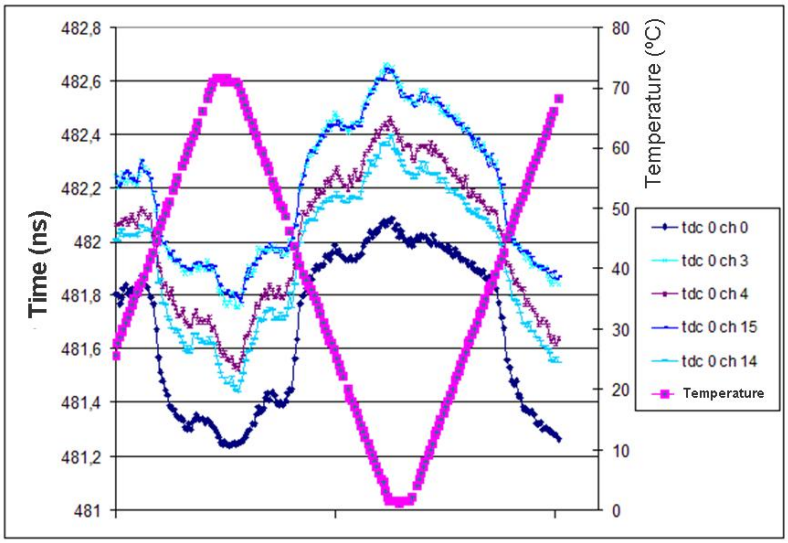

Figure 28: Variation of the time measurement of a fixed signal entering different channels in the HPTDC as a function of the board temperature.

After a careful study, it was determined that 30% of that timing measurement variation is due to the LVDS receivers (DS90LV048) themselves. Figure 29 shows the dependency of the DS90LV048 propagation delay with respect to the temperature.

In this test, the stability of the voltage regulators has also been measured, obtaining maximum voltages shifts below 0.2 mV/ºC and current variations around 0.4 mA/ºC.

Lifetime tests have also been performed with a similar set-up to a fully operating ROB at 105ºC ambient temperature. The goal of this test was to spot weakest devices or failure modes in the ROB by means of an accelerated stress test. During 4 months of operation no failure was found, which will mean a worst case failure rate per year below 0.026%. This failure rate has been calculated considering a low activation energy failure

mechanism, such as solder bonding (0.4 eV) [10].

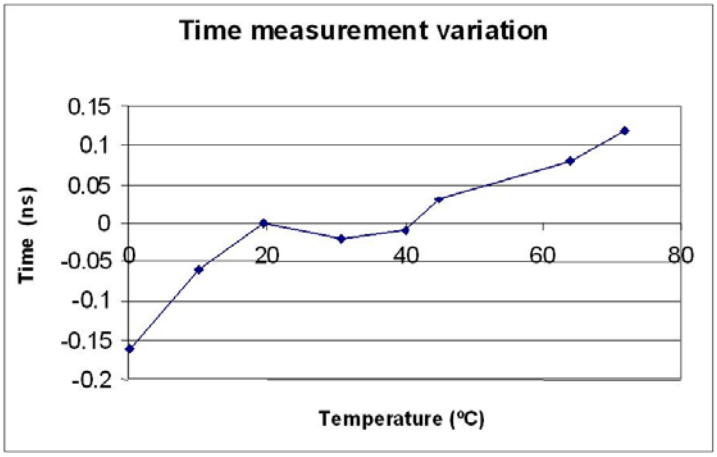

Figure 29: Time measurement deviation with respect to the value obtained at 20ºC due to the DS90LV048 LVDS receivers at the ROB.

Finally, non destructive burn-in tests have been performed to all the produced ROBs in order to screen for latent defects in the semiconductors and discard those devices with infant mortality [11]. Batches of up to 48 powered and clocked ROBs were stored in *ad-hoc* oven (figure 30) at a maximum of 60ºC for up to two weeks, equivalent to 3 months of operation in normal conditions. These tests allowed discarding 14 boards out of the 1600 ROBs produced. Main failures obtained were due to assembly defects which have activation energies in the order of 0.5 to 0.7 eV. Accordingly, the acceleration factors obtained in these tests are up to 17, and therefore, the expected failure rate due to this type of defect is reduced in the same factor.

Figure 30 Image of the ROB burn-in test stand.

4.3 ROB irradiation tests

The ROB boards are installed close to the detector 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 5 10^{10} p/cm^{-2} in 10 years of operation). Since radiation hard devices were not going to be used, 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 10^{10} cm^{-1} protons at 60 MeV in two different campaigns in June and December 2001. 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) [12]. An image of a ROB during the irradiation campaign can be seen in figure 31.

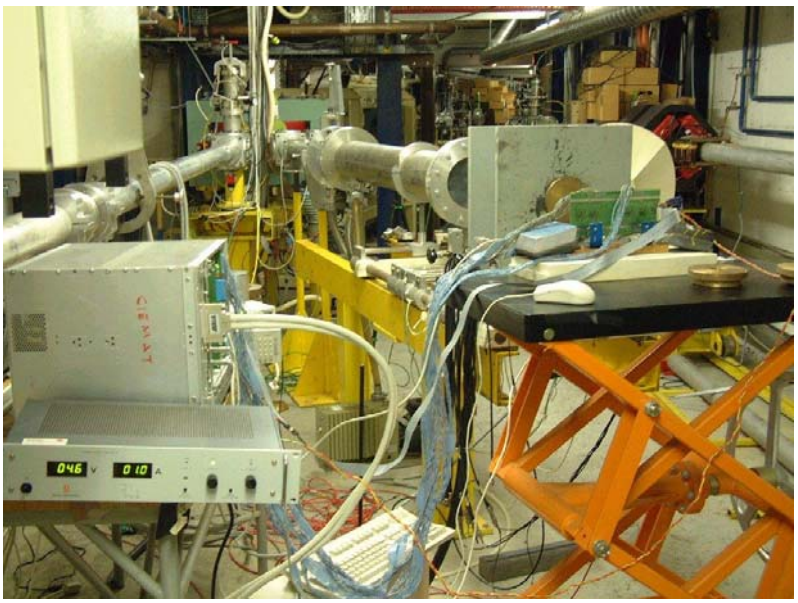

Figure 31: Image of the ROB during the irradiation campaign.

In the first campaign the LVDS receivers DS90LV048A and the regulators were irradiated. In the first case, eight DS90LV048A were mounted in a PCB powered and with 0 V inputs. Output signals and power consumption were monitored and neither malfunction was observed nor any current increase. In the second case, the regulators were also mounted on a custom PCB where the regulated voltage and the ground current were monitored. The results (figures 32 and 33) show that the regulated voltage varies less than 20 mV (0.6 %) for the MIC29151-3.3BU and the ground current increases less than 6 mA (40%), also for the MIC29151-3.3BU, which is acceptable for operation in CMS.

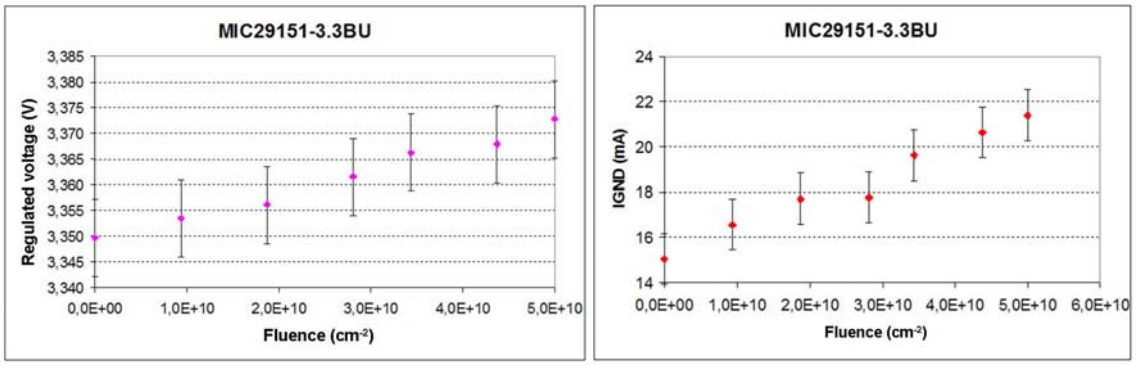

Figure 32: Variation of the regulated voltage and ground current of the MIC29151-3.3BU regulator as a function of the proton fluence.

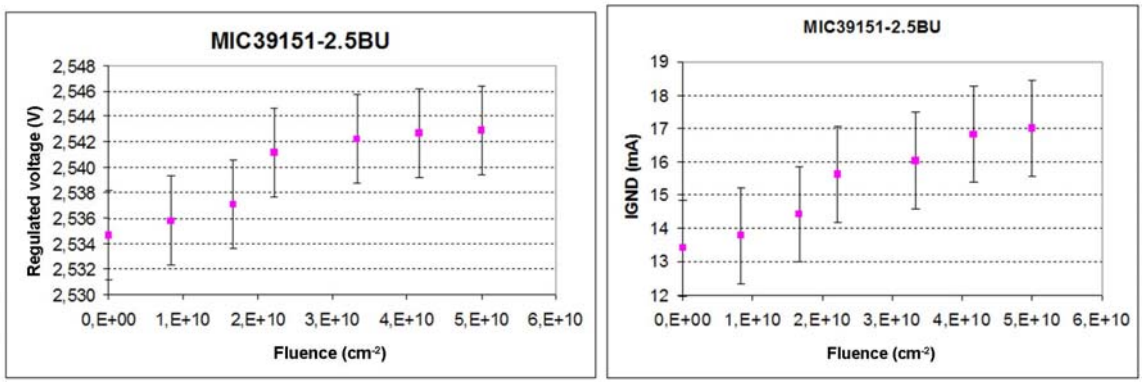

Figure 33: Variation of the regulated voltage and ground current of the MIC39151-2.5BU regulator as a function of the proton fluence.

In the second campaign one full ROB was already available and hence, the HPTDCs, the Altera EPM7128AE and all the other devices, which are located in the ROB central region, were irradiated. During irradiation, ROB was operated in normal mode and in Test Pulse mode, injecting known input patterns and transmitting the time digital data to a ROS prototype board that was read through a VME interface. ROB operation was continuously monitored (current consumption, sensor reading, status registers through JTAG interface, etc) in order to detect any malfunction. It is worth mentioning that a triple redundancy mechanism that solves 1 bit upsets has been implemented in the Altera firmware, and one output monitoring any upset was read with an external counter during the irradiation tests.

The results obtained are indicated in table 5, 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, λ the estimated failure rate for a 90% confidence level and MTBF the mean time between failures. As can be seen, only one malfunction, a SEU, was observed in one HPTDC. No latch-up or any other significant effect was observed.

Table 5: Estimated mean time between failures at 90% confidence level due to radiation in the ROB boards for 10 years of operation in CMS.

	HPTDC	ROB central region	Altera EPM7128AE
N_T	5820	1500	1500
N_F	1	0	0
N_r	8	2	4
Φ_{CMS} ($cm^{-2}s^{-1}$)	40	40	40
Φ_{LOV} ($cm^{-2}s^{-1}$)	$5 \cdot 10^{10}$	$5 \cdot 10^{10}$	$5 \cdot 10^{10}$
λ	$5.8 \cdot 10^{-7}$	$6 \cdot 10^{-7}$	$3 \cdot 10^{-7}$
MTBF (hours)	477.3	201.1	402.1
MTBF (days)	19.9	8.4	16.8

Summarizing, the selected components for the ROB are tolerant to the radiation doses expected during operation in CMS, validating the system design. The estimated mean time between failures is of ~20 days for all the HPTDCs in the detector.

4.4 Testbeam

ROB design has also been validated in several test beams at CERN: two of them in October and November 2001 with an MB2 chamber [13], one in May 2003 with one MB3 chamber equipped with a full Minicrate [14] and one in May 2004 with two chambers MB1 and MB3 and their corresponding Minicrates [15]. Since the operation of the read out electronics was very satisfactory in all the data taking campaigns, we will focus in the details of the first two test beams, which were the ones where the ROB design was validated for the first time.

The test beams took place at the GIF (Gamma Irradiation Facility) at CERN where the first produced MB2 chamber was placed under a muon beam and operated under normal gas and voltage conditions (figure 34). An RPC chamber was attached to the top of the DT chamber, and, for the first time, both detectors were operated coupled together. The performance of the DT chamber was studied for several operating conditions. Background rates, similar to the ones expected at LHC, were produced by a ^{137}Cs gamma source located 4 m upstream of the DT chamber, which could be switched on and off and its flux regulated by a system of lead filters.

Chamber signals were read out simultaneously by multi-hit CAEN TDCs and by one ROB. The collimated beam covered an area of 10 x 10 cm^2 with 2 mrad angular spread, that is, around ±0.8 cells in Φ and ±1.3 cells in θ, which is equivalent to 40 ROB channels.

Trigger signal was provided by a set of scintillators, which were placed in front of the chamber with their corresponding electronics, delayed to simulate trigger latency as in normal CMS operation. Trigger signal was synchronised with the system clock for L1A generation but also, a non-synchronised trigger was fed into two ROB channels for relative time measurement and redundant information of the trigger arrival time.

An experimental set-up was developed to implement ROB control and data acquisition. The ROB was connected through a set of custom boards and a VME interface to a PC for HPTDC programming and monitoring and data storage. Control of the read out system, data unpacking and spills management was done by software.

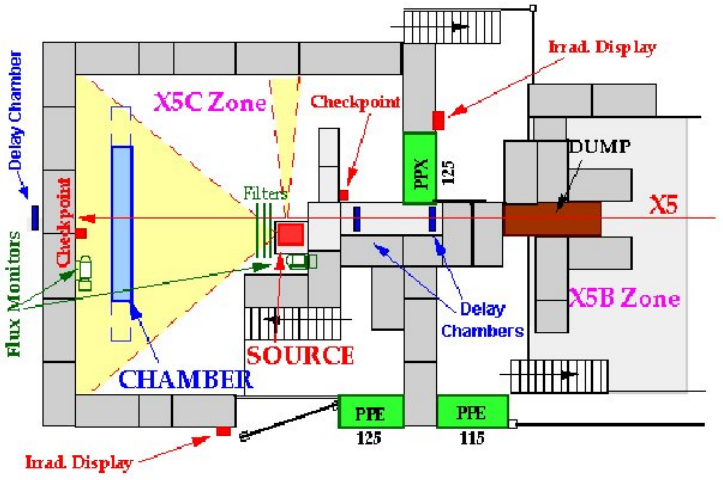

Figure 34: Schematic representation of the GIF test set-up at CERN.

Tests were made under two different beam conditions. First, at P2B (Oct. 4 to Oct. 23, 2001) with a non structured beam of about 6000 triggers/spill, that is, 1.2 kHz trigger rate, with spill duration of 5.1s in periods of 16.8 s. Second, during P2C period (Oct. 25 to Nov. 4, 2001) with a 25 ns structured beam of ~26000 triggers/spill, i.e. 5 kHz trigger rate.

Some of the basic TDC configuration parameters are the following: 40 MHz clock, PLL to reduce jitter, leading measurements, 0.78125 ns bin resolution, 8 bits parallel read-out at 20 MHz synchronous bus, 1.1μs latency, 1.3μs reject window and various search and matching windows around 1μs depending on the run.

4.4.1 P2B test beam results

During the first test beam period in October 2001 7 runs of 5 10^5 events/run were registered. The correct operation of the system was confirmed, with no TDC errors, no FIFO overflows and only a few overlapping triggers. Considering that the matching window was set to 800 ns and we had about 1 trigger/ms, the probability of having overlapping triggers was only 8 10^{-4}. Since we were measuring also the non-delayed trigger signals as hits, overlapping triggers appear as multiple trigger signals in the same event, so they can easily be recognised. The correct operation of the TDC in overlapping triggers was confirmed, as well, as the offline capability to disentangle those events.

During this exercise, the proper operation of the TDC in the vicinities of the bunch counter roll over was also confirmed. In figure 35 an image of a time histogram obtained during these tests is shown. Data were taken also with gamma irradiation over the whole chamber, to simulate background under normal LHC operating conditions. The effective noise increase was observed but had no effect on HPTDC behaviour.

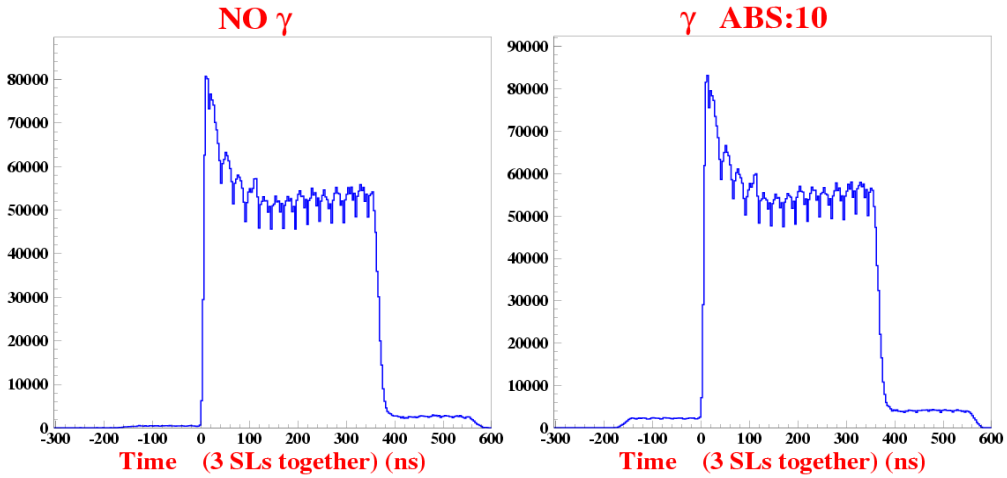

Figure 35 Typical drift time-box distribution after relative T0 correction for a run without gamma irradiation (left) and for a gamma background using an attenuation filter of factor 10. The baseline increase can be observed due to the effect of the induced background noise.

4.4.2 P2C test beam results

During the second beam period the muon beam had a 25 ns structure similar to LHC beams. In this exercise 9 runs were taken (5 10^6 events), with basically the same set-up and configuration as in the first test beam.

The HPTDC was configured to send its internal buffers occupancy in some debug words within the data stream. In figure 36 an example of the occupancy of those memories during a run can be seen. The effect of two very noisy channels (16 and 19) was affecting TDC 1 occupancies. It was confirmed that the HPTDC flagged correctly errors due to occasional buffer overflows corresponding to these MHz noisy channels. Moreover, in case of buffer overflow, only the 8 channels sharing the same buffer were affected by the loss of hits, without compromising the correct operation of the rest of the channels. The arbitration mechanism between buffers worked very satisfactorily.

L1 buffer (max)	TDC 0	TDC 1	TDC 2	TDC 3
group 0 (0-7)	6	16	14	11
group 1 (8-15)	17	17	21	12
group 2 (16-23)	16	254	5	5
group 3 (24-31)	13	15	0	0

•L1 buffer minimum occupancy = 0; except TDC 1 group 2 = 8

Trigger FIFO	TDC 0	TDC 1	TDC 2	TDC 3
max occupancy	3	3	3	2

•Trigger FIFO minimum occupancy = 0

Read out FIFO	TDC 0	TDC 1	TDC 2	TDC 3
max occupancy	56	120	79	59

•Read out FIFO minimum occupancy = 3

Figure 36: Occupancy of the HPTDC internal buffers. 96 channels were connected to the chamber, group 3 of TDC 2 and 3 were left disconnected and group 2 of TDC 2 and TDC 3 were receiving the asynchronous trigger signal for reference.

From the debug information, it could also be seen that the trigger FIFO was far from overflowing, thus matching is done much faster than the trigger rate. The read out FIFO has a fairly high occupancy in the TDC with noisy channels, but we could see that it could stand this amount of noise at the given transmission bandwidth.

With the given beam structure the number of overlapping events was much larger and up to 5 overlapping events were detected. The hits rejection and reconstruction mechanism were tested successfully. In figure 37 the 25 ns structure of the muon beam as obtained from the time differences between the ROB asynchronous trigger inputs can be observed.

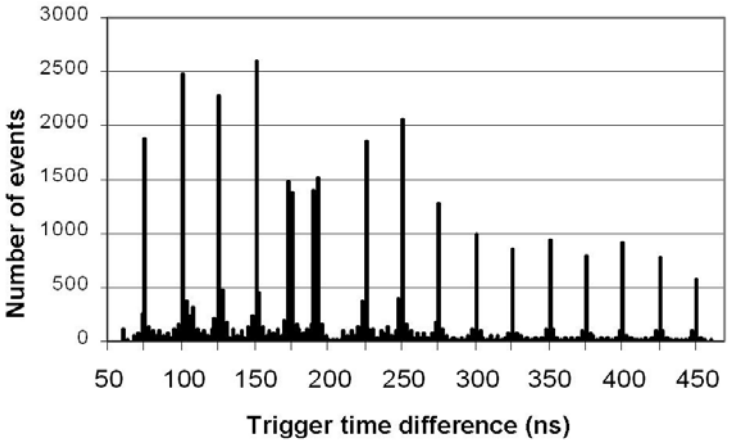

Figure 37: Histogram of the time differences between the asynchronous trigger input signals in the ROB that show the structure of the 25 ns muon beam.

Finally, the basic parameters that show the proper operation of the DT chamber were analyzed both from CAEN TDC and ROB data. Results of the Mean Timer, drift velocity and single wire efficiency are shown in table 6. The values have been obtained for different statistics but the results are fully compatible.

Table 6: Comparison of some DT chamber parameters during 2001 test beam in a data acquisition with CAEN TDC board and ROB board.

	CAEN data acquisition			ROB data acquisition		
	SLθ	SLΦ1	SLΦ2	SLθ	SLΦ1	SLΦ2
Average Mean Timer (ns)	380.4	381.3	380.6	379.2	379.9	381.5
Drift velocity (± 0.3) (μm/ns)	55.21	55.07	55.18	55.38	55.28	55.04
Efficiency	99.92 ± 0.01	99.97 ± 0.01	99.97 ± 0.01	99.97 ± 0.01	99.97 ± 0.01	99.97 ± 0.01

Summarizing, the proper operation of the ROB board with a real DT chamber and similar LHC conditions was confirmed during these exercises. Although maximum trigger rate (5 kHz) was lower than what expected at LHC (100 kHz), the muon rate was much higher (500 Hz/cm^2 versus 10 Hz/cm^2). It was verified that the ROB can handle without bottlenecks the expected data rate with higher background levels and also in the case of noisy channels.

5 Conclusions

After several years of design, production and installation the first level of the CMS DT read out system is ready for operation. The different prototypes and the final boards have been extensively tested in laboratory, where its time measurements capabilities have been characterized and validated. ROB functionality has also been validated in different test beams where it has been operated with a DT chamber in conditions similar to the ones expected in the CMS detector.

Moreover, lifetime and burn-in tests have allowed the study of limited maintenance problems, which are minimized by the employment of independent functional units, sensoring elements and other failure detection mechanisms.

Radiation tests of the ROB have also shown a very good behaviour of the selected devices, with expected mean time between failures due to Single Event Upsets of 20 days for the whole CMS detector.

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

References

[1] CMS Collaboration, *The Muon Project. Technical Design Report*, CERN-LHCC 97-32 (1997).

[2] J. Christiansen, *High Performance Time to Digital Converter. Version 2.1.* CERN/EP-MIC (2002).

[3] B.G. Taylor, *TTC distribution for LHC detectors*, IEEE Trans. Nucl. Sci. 45 (1998) 82. http://www.cern.ch/TTC/intro.html

[4] J.M. Cela et al, *CMS DT Chambers Read-Out Electronics,* CMS CR 2008/018. CERN, 2008.

[5] A. Racz, *Trigger Throttling System for CMS DAQ,* Proceedings of the sixth Workshop on electronics for LHC experiments, Cracow, 11-15 September, 2000.

[6] *TTCrx reference manual*, available at: http://ttc.web.cern.ch/TTC/TTCrx_manual3.11.pdf

[7] RD5 Collaboration, F. Gasparini et al, *Bunch crossing identification at LHC using a mean-timer technique*, Nucl. Inst Methods Phys. Res. A, Vol. 336 (1993) 91.

[8] V. Liberali, F. Maloberti, and M. Stramesi, *ADC Characterization Using the Code Density Test Method With Deterministic Sampling,* Proc. Int'l Mixed Signal Testing Workshop, (1996), pp. 113-118.

[9] Ph. Farthouat, P. Gällnö, *TTC-VMEbus Interface TTCvi-MkII*, RD12 Project, available at: http://ttc.web.cern.ch/TTC/TTCviSpec.pdf

[10] V.Lakshminarayanan, *Environmental-stress screening improves electronic-design reliability,* Center for developlment of telematics (2001) EDN.

[11] U.S. Department of Defense, *The Test Method Standard Microcircuits. MIL-STD-883E, method 1015.9*,

available at: http://scipp.ucsc.edu/groups/glast/electronics/mil-std-883.pdf

[12] M. Huhtinen and F. Faccio, Nucl. Instr. and Methods Phys. Res. A, Vol. 450 (2000) 155.

[13] C. Albajar et al, *Test beam analysis of the first CMS drift tube muon chamber,* Nucl. Instrum. Methods Phys. Res. A 525 (2004) 465-484.

[14] P. Arce et al, *Bunched Beam Test of the CMS Drift Tubes Local Muon Trigger,* Nucl. Instrum. Methods Phys. Res. A, Vol. 534, Issue 3, pp. 441-485, (2004).

[15] M. Aldaya et al, *Results of the first integration test of the CMS drift tubes muon trigger*, Nucl. Instrum. Methods Phys. Res. A, Vol. 579, pp. 951-960, (2007).