Available on CMS information server **CMS IN -2009/014**

CMS Internal Note

The content of this note is intended for CMS internal use and distribution only

28 November 2008 (v3, 24 March 2009)

Simulation studies for the read-out electronics of the CMS Muon Drift Tubes detectors

P. Arce, M. Cerrada, C. Fernandez Bedoya, J. Molina, C. Willmott

Abstract

We present the studies made for the rates that the read-out electronics of the Muon Drift Tubes (DT) will have to face under luminosities of 10^{34}cm^{-2}s^{-1} and 10^{35}cm^{-2}s^{-1}. The studies were performed with particles coming from minimum bias events generated at the interaction point (IP) and propagated through the CMS detector until the creation of hits in the different DT chambers. We conclude that the rates expected for the Read-Out boards (ROBs) are not critical for normal operation in both scenarios while Read-Out Server boards (ROS) will probably need upgrade when running at ten times LHC luminosity.

CMS Internal Note

The content of this note is intended for CMS internal use and distribution only

Simulation studies of the read-out electronics for the CMS Muon Drift Tube detectors

P. Arce, M. Cerrada, C. Fernandez Bedoya, J. Molina, C. Willmott

CIEMAT - Centro de Investigaciones Energeticas, Medioambientales y Tecnologicas, Madrid, Spain

Abstract

We present the studies made for the rates that the read-out electronics of the Muon Drift Tubes (DT) will have to face under luminosities of $10^{34}\text{cm}^{-2}\text{s}^{-1}$ and $10^{35}\text{cm}^{-2}\text{s}^{-1}$. The studies were performed with particles coming from minimum bias events generated at the interaction point (IP) and propagated through the CMS detector until the creation of hits in the different DT chambers. We conclude that the rates expected for the Read-Out boards are not critical for normal operation in both scenarios, while the Read-Out Server boards will probably need to be upgraded for running at ten times LHC luminosity.

1 Introduction

The Compact Muon Solenoid is a general purpose detector installed in the Large Hadron Collider (LHC). It consists of a large superconducting solenoid with a radius of 3 m and a length of 13 m, which generates a magnetic field of 4 T. This large magnet allows all of the tracking devices and calorimetry to be placed inside the coil of the solenoid, while the muon system is located outside the coil, in the return yoke of the magnet. The muon system employs three different technologies to trigger and measure muons: Drift tubes in the barrel region, cathode strip chambers in the endcap region, and resistive plate chambers in both the barrel and endcap regions.

The barrel region is subdivided into five wheels (YB-2, YB-1, YB0, YB1, YB2) with 12 sectors per wheel. On each sector there are four levels of DT chambers, namely stations MB1 to MB4.

A DT chamber consists of three superlayers (SL), each with four layers of drift cells, separated by a light aluminum structure called the honeycomb. The two outer superlayers perform muon ϕ coordinate measurement (angle on the bending plane), while the inner one measures the θ coordinate (angle in the non-bending plane).

Any charged particle track going through a cell volume will generate a signal (hit) in its anodic wire that will be processed by the ROBs located in the so-called minicrate attached to the DT chamber. Each ROB allocates four High Performance Time to Digital Converters (HPTDC) [1] that will perform the time measurement of the hit with respect to the Level1 Accept (L1A) signal. With a drift velocity in the chamber cell of 55 μm/ns and a single wire resolution of 200 μm, HPTDC time resolution of 0.781 ns is more than enough to provide an accurate time measurement in order to allow further reconstruction of particle tracks.

As the maximum drift time inside the DT cell (400 ns) is much larger than the LHC bunch crossing period (25 ns), a read-out system that manages overlapping triggers is required. Even though the bunch crossing frequency is high, the complex CMS trigger system [2] will reduce the first level trigger (L1 Accept) rate to 100 KHz, having a L1A accept latency of around 3.2 μs. There are several other requirements for the read-out electronics imposed both from the functional operation of the LHC system and from the environmental conditions expected in the CMS detector. Accordingly, the read-out electronics has been designed to accomplish these requirements [3].

2 The Read Out Boards (ROB)

Figure 1 shows the architecture of the HPTDC, where the digitalized signals from each of the 32 channels are first stored in four input hit registers per channel and then are sent to 256-word-deep group memories common to every 8 channels (Level 1 buffers). There they are kept for 3.2 μs until the trigger matching is performed. At the arrival of the L1A, the trigger matching mechanism will look back in time the programmed trigger latency for all hits belonging to a programmed time window. We will assume in this discussion that we use a programmed window of 1 μs, wide enough to acommodate a maximum drift time of 400 ns. Those hits will be accepted and merged to a common 256-word-deep read-out memory.

Hits will only be rejected from the Level 1 buffer memory when they are older than a specified reject time in order to allow reading of overlapping events. Later track reconstruction will uniquely determine which hits belong to each L1A.

On every L1A, the ROBs will encapsulate accepted hits in an event data block that contains the event and bunch crossing numbers and channel identification along with the corresponding time measurement. These data will be transmitted to the Read-Out Server (ROS) boards through 200 Mb/s copper links.

3 The Read Out Server (ROS)

ROS boards are 9U VME boards located in the towers on one side of the CMS wheels. Each ROS has 25 input channels, one per ROB, performing the read-out of one complete sector.

Input data are deserialized and stored in 4 kbyte FIFO per channel. At each L1A, an ROS processes in parallel four groups of 6 channels (plus a group of one channel). Data are merged into a “sector event” that is transmitted through optical fiber to the next level in the data acquisition chain, the DDUs, located in the Control Room. Data consistency, transmission errors, and event synchronization are checked at the ROS and any integrity error is notified not only in 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, removing empty ROB blocks with only headers and trailers, this data reduction allow us to achieve around 8 kbytes event size in the whole detector.

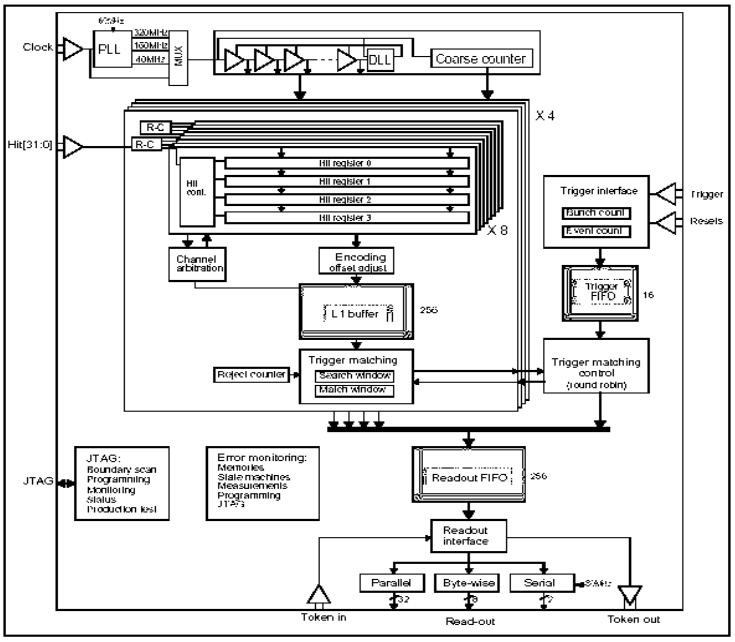

Figure 1: High Performance Time to Digital Converter architecture.

Every L1A trigger will generate one ROB event block with a master header and trailer information that includes the bunch event number, the number of words transmitted, and the time measurement words that include information about the channel and HPTDC number that generated the signal. These data are sent from each minicrate through two FTP cables to the ROS boards located in the towers on one side of the CMS wheels. Each ROS reads the information from 25 ROBs (one sector), and sends them to the DDUs after processing.

4 Monte Carlo simulations

In order to study the occupancy of the DT chambers we need to simulate events at the IP and then propagate them through the detector, where showers of particles will be created by the punchthrough process. We define punchthrough as those events that never had a muon in their trajectory, neither as parents or daughters.

Nearly one million minimum bias events where generated inside the CMSSW enviroment 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. Even with this option, we estimate that the error on neutron production is a factor of 2.

Figure 2 shows the number of hits detected by the five different wheels of the muon detectors. It can be seen that 99123 hits were detected, which represents around 10% of the 994883 events generated at the IP. From those 99123 hits, approximately 88% were produced by particles created by punchthrough processes shown by the red dashed line.

Most of these background particles have an arrival time different from the muons generated at the IP. Figure 3 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 get created and to reach the detectors (in fact, 35% of them will reach the detectors after this period time of 1 μs).

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 on the most populated wheel. In Figure 4 we can observe that the maximun 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

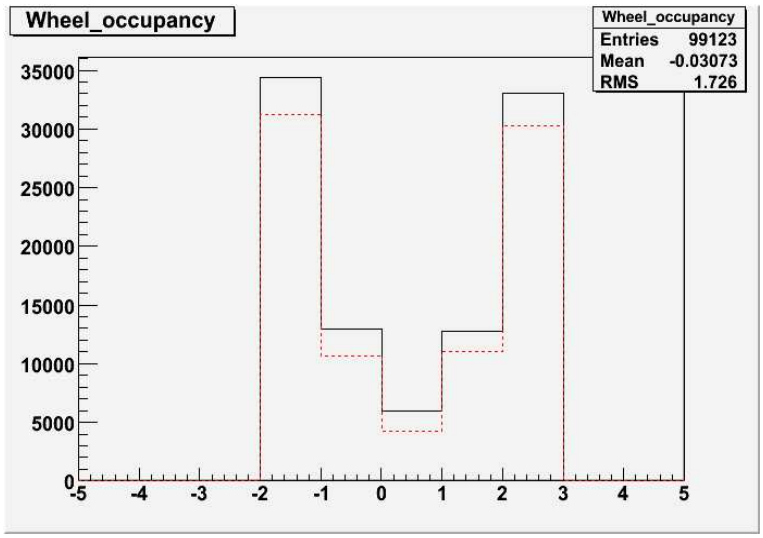

Figure 2: Occupancy of particles in the five wheels of the CMS Muon detectors. The black continue line represents all the hits produced, while the red dashed line shows the hits produced by the punchthrough process. They represent around 88% of the total hits detected.

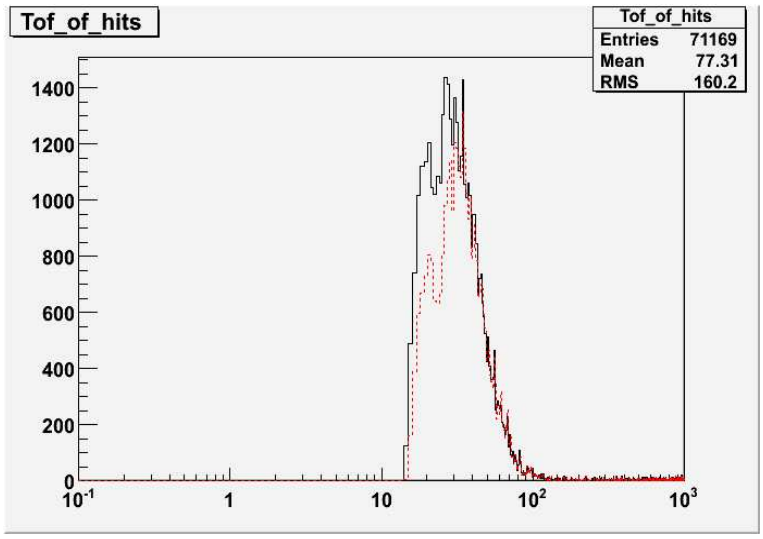

Figure 3: Arrival time of particles to the chambers during the 1 μs time window. The peak around 20 ns is attributed to the time of flight for particles originated at the IP and reaching the first chamber. The red dashed lines are the hits caused by punchthrough particles.

rest of the discussion.

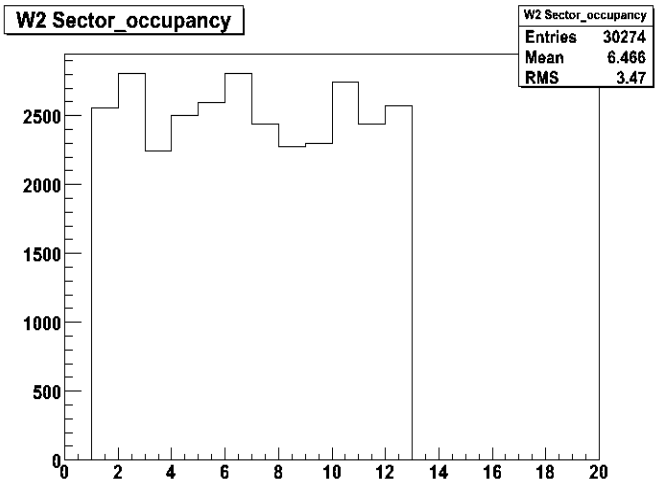

Figure 4: Occupancy of particles by sector in wheel 2 of the DT Muon detector. The big differences among sectors are due to multiples interactions, where we can have as much as 70 hits per event.

5 Occupancy and hit rates

5.1 ROB L1 buffer occupancy

For a luminosity of $10^{34}\text{cm}^{-2}\text{s}^{-1}$ and using the total cross section value σ = 80 mb, we have that 80×10^7 events per second are produced at the interaction point.

Figure 5(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 occupancies will happen in the HPTDCs of MB1.

So, 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 80×10^7, the frequency of hits that we may expect in the MB1 is 1.78 MHz.

Now we want to know the occupancy among the superlayers inside the MB1, since we are interested in the rate of the SL1. This is shown in Figure 5(b), where we can see that 37% and 30% of the hits falls in SL 1 and 3 respectively, which indicates that the rate received by these two SLs is 1.27 MHz (we have to take into account the SL3 as well since the ROBs reads both superlayers at the same time).

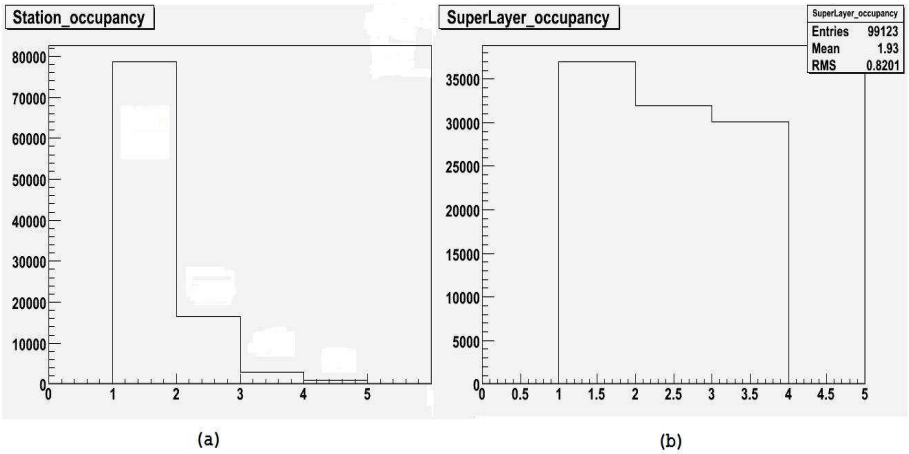

Figure 5: (a) Particles detected in each station of the CMS muon detectors. (b) Distribution of the hits of the chamber among the SLs.

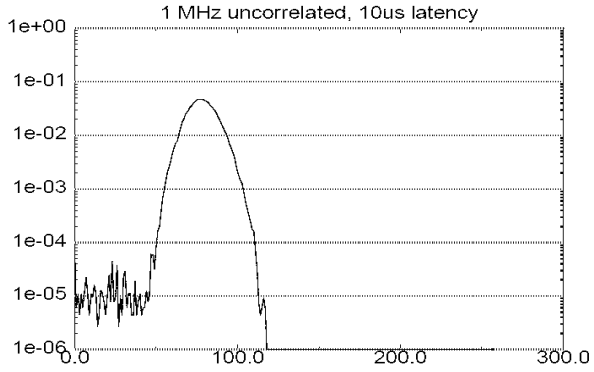

Figure 6: L1 buffer occupancy at 1MHz hit rate for 1 μs window.

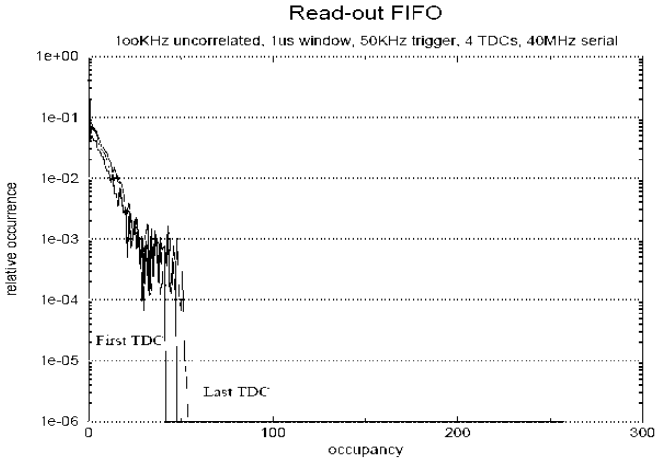

Figure 7: Read-out FIFO occupancies for four TDCs sharing a slow serial read-out.

If we distribute this rate equally among the 12 HPTDCs that read the two SLs, we have that each TDC will have to work at a 105 kHz rate, which translates to a channel rate of 3.3 kHz.

These signals will remain in the HPTDC Level 1 buffer until the trigger matching is performed, with a latency of 3.2 μs. From tests performed with the HPTDC [1], it can be seen in Figure 6 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 μs. Note that in the previous calculus, the rate obtained was 3.3 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.

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. If we assume that we will read 1 μs of time window with a trigger rate of around 100 kHz, that means that we will be reading 10% of the time, so the hit rates will be reduced accordingly. In Figure 7 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 100 kHz, a faster 20 MHz byte-wise read-out protocol is used in the ROBs. In either case, the occupancies we obtain (including the estimated error in the neutron production) are far from saturating the HPTDC buffers, even if we consider an increase in luminosity of 10 times more.

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.

5.2 ROS occupancy

ROS boards receive the timing information of all the HPTDCs in one sector, once these hits have been accepted by the trigger matching mechanism. As said before, an ROS has 25 input channels, each of which corresponds to 1 ROB, i.e., 4 HPTDCs. 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 8, where the time is measured in terms of the bunch crossing (bx), equivalent to 25 ns.

As it is done, channels corresponding to MB1 stations are faster, which relates to the higher occupancy expected there. In the previous figure, the maximum frequency is calculated just by dividing the time to process an event, 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 frequency is larger than what is shown in this paper and depends on the particular event size and on the actual time between events.

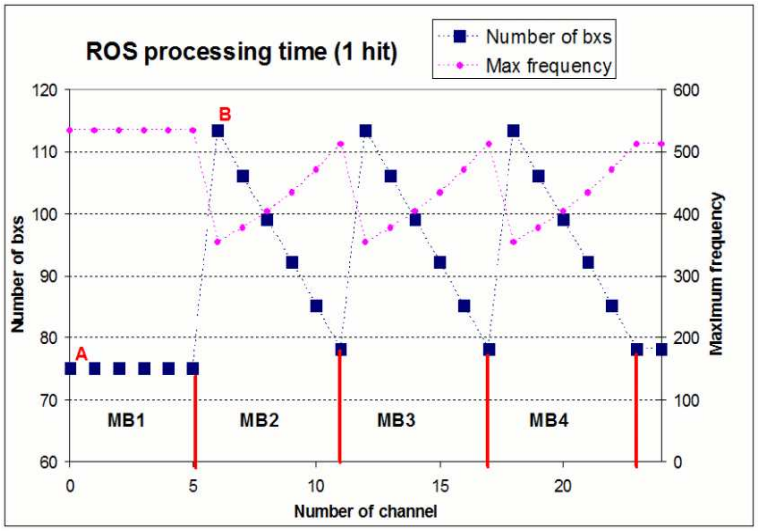

Figure 8: Time for the ROS to process 1 hit versus channel number that received that hit. Maximum event frequency as explained in the text is also plotted.

As it is done, channels corresponding to MB1 stations are faster, which relates to the higher occupancy expected there. In the previous figure, maximum frequency is calculated just by dividing the time to process an event, 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.

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

Time to process an event = 73 + 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 8.

Figure 9 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 8) and in the worst case (hits in first channel of MB2, point B in Figure 9). 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.

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 10. The upper curve corresponds to the processing time when all the hits are in the same channel. We have done an average taking into account 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 real situations.

Figure 10 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

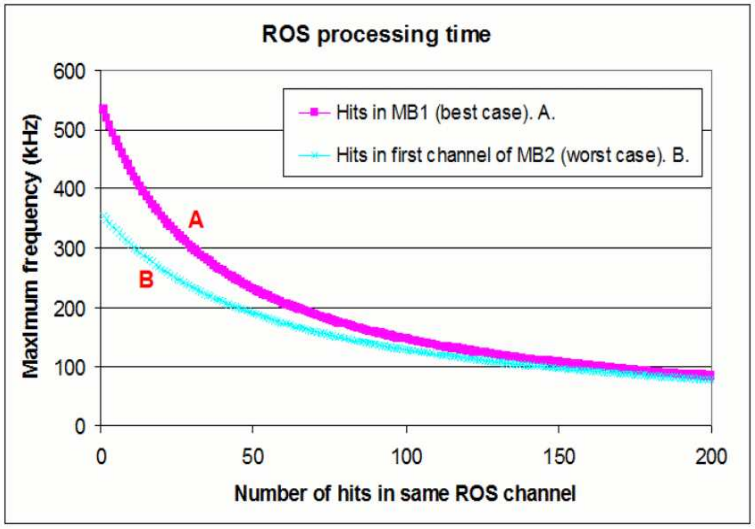

Figure 9: Maximum ROS processing frequency assuming no FIFOs versus number of hits with all hits in the same input channel.

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.

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.) gives around 10^5 cells with noise below 20 Hz, 10^4 below 100 Hz, and around 20 cells with noise rates above kHz in all the 60 sectors of the detector. Therefore, the levels of this noise are very low.

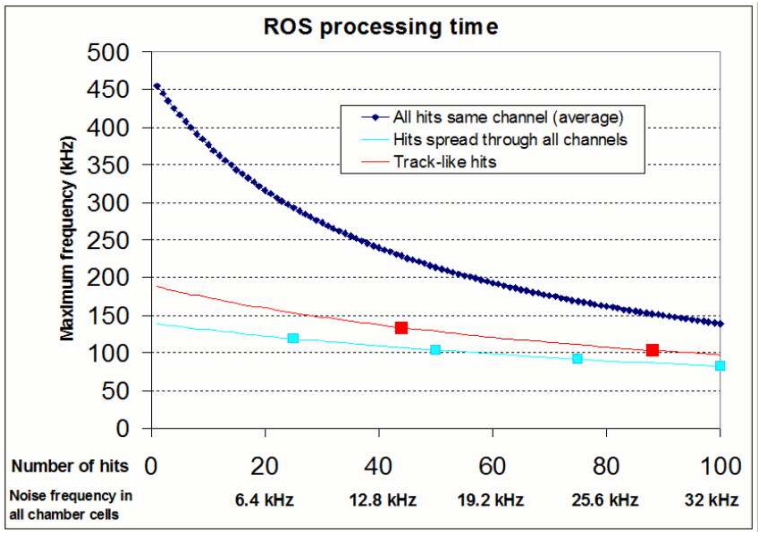

Figure 10: ROS processing time without storing events versus number of hits in one sector for three different hit distributions.

After describing ROS processing capacities, let us now examine the impact of the number of hits in the simulated LHC events. From the 80×10^7 events per second that we will have at LHC luminosity, the number of events in a time window of 1 μs corresponds to 800. We saw in Figure 4 that the maximum occupancy in one sector is around 2800 hits, so, from the 998433 events simulated we obtain that the number of hits that the ROS will have to process per L1A is 2.25.

Figure 11 shows the hit multiplicity produced by the particles that go through the different stations. Peaks for 12,

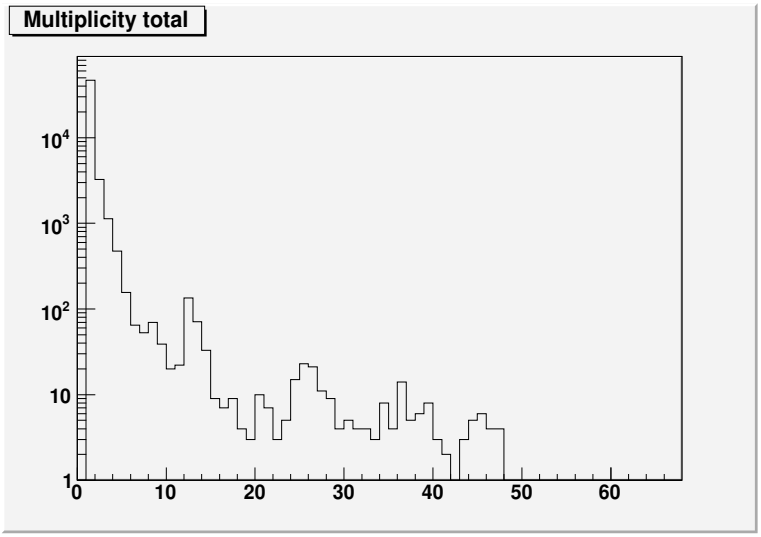

Figure 11: Number of hits created by the particles as they goes through the differents DT chambers.

24, and 44 hits correspond to events where a charged particle track leaves signals in 2, 3, or 4 muon chambers. Nevertheless, it can be seen that 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. Therefore, for background noise events, ROS will process events much faster than the expected L1A rate, so most of the time ROS FIFOs will be empty.

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. If muons are spread throughout all the chambers and not always in the same phi coordinate, 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 for processing overlapping triggers.

In summary, the ROS processing time is fast enough for LHC running.

For the sLHC case, ROS will have to process 22 hits, but if we take into consideration the factor of 2 due to the uncertainty in the neutron generation, we can have 44 hits per event in each sector from background. In such a case, ROS processing time is too large to be able to operate with a reasonable safety margin. Therefore, several approaches are being studied at present in order to reduce processing time: discard events with only one hit per ROB (a muon will produce at least 4 hits), speed up the handshake protocol between the different blocks in the ROS, or a complete redesign of the ROS board with higher frequency clocks.

6 Conclusion

A simulation with minimum bias events has been performed in order to study the hit rates in the DT chambers during LHC operation. We made our calculation independently of bunch crossing time, so any change in this parameter to extrapolate for SLHC will not affect the results. Rates expected are low enough for comfortable operation of DT read-out electronics, which will easily handle this data throughput, including present electronic noise seen in the fully installed CMS detector. ROB boards will also be adequate for operating in future scenarios with 10 times more luminosity, though ROS boards will probably need further upgrades to speed-up their processing time.

7 Aknowledments

One of the authors (J. Molina) is sponsored by the CIEMAT as part of the contribution of the International HELEN (High Energy Physics LatinAmerican-European Network), under contract AML/19.0902/97/0666/II-0485-FC-FA-

FCD-FI (105215).

References

[1] J. Christiansen. High Performance Time to Digital Converter. Version 2.1. CERN/EP-MIC. July 2002.

[2] CMS Collaboration. (2000, Dec.) CMS. The TriDAS Project. Technical Design Report. Volume I. The Trigger Systems. CERN/LHCC 2000-38.

[3] J.M Cela et al. CMS Drift Tube Chambers Read-Out Electronics, Proceedings of the TWEPP-07. Prague, September 2007.