Upgrade of the MDT Electronics for SLHC using Selective Readout

O. Biebel[*], J. Dubbert[+], S. Horvat[+], O. Kortner[+], H. Kroha[+], R. Richter[+], D. Schaile[*]

The MDT data volume at high luminosities

The data volume read out by the MDT electronics at high luminosity is completely dominated by the background hits caused by converted neutrons and gammas. At the nominal luminosity of 10^{34} cm^{-2}s^{-1} simulation of the radiation background predicts hit rates of up to 70 kHz per tube. At a LVL1 trigger rate of 100 kHz this leads to about 10% saturation of the readout bandwidth of the electronics chain. If a “safety factor” of 5 was applicable, rates would be 5 times higher and saturation 50%, leaving little headroom for luminosity upgrades. These numbers apply to the “hottest” chambers, located in the endcap region, i.e. the Small Wheel and the Inner rings of the EM chambers, while the rates in the EO Wheel and in the Barrel are up to a factor five lower and therefore less critical[1].

In summary, the readout bandwidth of the MDT chain in the forward region is not sufficient for SLHC luminosities, even if the “safety factor” of 5 needs not fully to be applied. Therefore, a concept for the upgrade of the readout chain should be developed.[2]

Options for an improved readout

An obvious solution would be the enlargement of the readout bandwidth via faster transmission lines and processors, using more modern technologies than were available at the times, when the present system was designed. This option is discussed in a separate EoI, already presented to the Upgrade Steering Group (“Upgrade of the MDT Readout Chain for the SLHC”).

In the present EoI we want to discuss an alternative approach, trying to minimize the required changes to the existing readout system and therewith labour and cost. The idea is to reduce the data volume to be transferred to the off-detector electronics as early as possible in the data path. The amount of “useful” data in the MDT stream, in fact, is minute; with only about 1,5 tracks above trigger thresholds per event, the number of tracks in the Muon spectrometer is very low, even at SLHC (see appendix), and only these tracks, passing at least one trigger threshold are assigned a RoI, are presented to the LVL2 and have a chance to survive to later trigger stages. The LVL2’s decision to retain or discard MDT information is entirely based on the RoI information received from the Muon trigger system.[3] In other words, the large majority of the MDTs containing no tracks above threshold (about 99% of the 1200 MDTs) and not having been assigned a RoI, will be discarded from the data stream by the LVL2. If those chambers without RoIs, containing only background hits, could be identified at a early stage, their data would not

[*] Ludwig-Maximilian-Universität, München
[+] Max-Planck-Institute für Physik, München

[1] If the deadtime in the preamplifier stage, which is now programmed to the standard value of 700 ns, was reduced to increase efficiency at high background rates, an additional factor of about two is needed for the readout bandwidth to cope with the additional hits.

[2] As working hypothesis we assume that the „safety factor“, i.e. the ratio between background at LHC and simulated value will be 5. The consequences of a Beryllium beam pipe on the background rates are not taken into account.

[3] As an exception to this, isolated tracks from the inner detector may be extrapolated to the Muon system, looking for a match, even if there was not RoI provided from the Muon trigger in this region. This procedure, however, is only foreseen for low luminosity runs around 10^{33} cm^{-2}s^{-1} s and is of no relevance for the high-luminosity situation discussed here.

have to be transferred to the MRODs, and a reduction of the required readout bandwidth by more than an order of magnitude would be the consequence.

Concept for the Selective Readout scheme

Identification of MDT chambers without tracks

In the present trigger system the MuCTPI condenses the RoI information received from the trigger logic (i.e. the RPC and TGC subsystems) in two different ways. First, a short list is prepared for the LVL1 decision, containing only the *multiplicities* of tracks having passed the different p_T thresholds. If a LVL1 confirmation takes place, a second list of RoIs is prepared for the LVL2 system, containing detailed information on location and momentum of each track.

The concept of selective readout is to use this list of RoIs (anyhow prepared for the LVL2) for the selection of MDT chambers to be read out. MDTs not corresponding to an RoI will be discarded from the data stream already at the level of the CSM, i.e. will not be transferred to the MROD. This way an average of nine chambers[4] would have to be read out per relevant track (i.e. per RoI). For an average event with 1.5 tracks, this would result in about 15 chambers instead of 1200 chambers in the present scheme, resulting in a reduction of the transferred data volume of about two orders of magnitude.

On the chamber side this readout scheme could function in the following way. After reception of a LVL1 trigger from the TTC, the CSM[5] would wait for a second signal from the Muon trigger confirming or vetoing the readout of the data belonging to this trigger. On a veto, the CSM would sent an empty event notification to the MROD, as is already now foreseen for events without hits.

Implementation on the MuCTPI side

As discussed above, the decision which MDT chambers to read is based on the *full* RoI list, produced by the MuCTPI, which is available ≤ 100 μs after the LVL1 trigger. For each RoI it contains the sector and subsector location, the flag for location close to the sector boundary and a coarse momentum estimate. A detailed description of RoI information for the LVL2 is given in ref. 1 (sect. 13.4.4, p. 359).

In order to use this information for selective MDT readout, a processor would have to be added to the present Muon trigger system, which would translate the RoI list into a list of chambers to be read out, sending the corresponding flags to the CSMs, e.g. via a system of fibers.[6] The algorithm for the chamber selection would have to be controlled by a supervisor, which could impose a more or less restrictive readout strategy, depending on different experimental conditions. As the supervisor may also decide for a complete readout, the present readout scheme could be considered a particular option under the selective readout scheme.

[4] In most cases a track only traverses three chambers. Close to the boundary, however, a track may migrate into the adjacent chamber (in the bending η-plane); so as a precaution the information of these chambers should also be retained, leading to nine chambers.- The possibility of losses due to tracks crossing boundaries in Φ has to be studied in more detail. This crossing could happen because of non-radial tracks (due to decays) or because of Φ-components in the magnetic field.

[5] The Chamber Service Module is collecting the hit information from the TDCs, sending the formatted data to the corresponding MROD via an optical link.

[6] It may not be necessary to implement one link per MDT, as chambers might be grouped according to trigger sectors, requiring only one link per group.

Implementation on the chamber side

In the present system the readout of a MDT is started when the CSM receives a LVL1 trigger. The LVL1 signal is passed to the TDCs which transfer the hits recorded in a certain time frame ("drift time window") to the CSM, where the data are formatted and stored into an internal buffer, waiting for readout to the CSM. This buffer is structured as a derandomizer in order to cope with fluctuations in LVL1 frequency and event size. As a consequence data transfer from the CSM to the MROD had to be asynchroneous:data from a given event may be transferred at different times for different chambers, depending on the length of the derandomizer queue and data volume.

In the selective readout scheme, the readout of the derandomizing buffer towards the MROD would only start after the confirmation signal has been received. In most cases a veto will be received, the readout of the data being replaced by a notification to the MROD.

It should be noted that in the selective scheme MDT chambers are read out or skipped *completely*. No attempt is made to select data in certain regions of a chamber. The reduction of the overall data volume is due to the fact that the large majority of the MDTs (> 98%) do not contain any useful information (i.e. no RoIs). Skipping the readout of these MDTs leads to a large reduction of the overall data volume, while the data volume of the *selected* MDTs is not reduced. A further data reduction is only done in the LVL2 processors, where the location of the track is narrowed down to a road size of a few tubes.

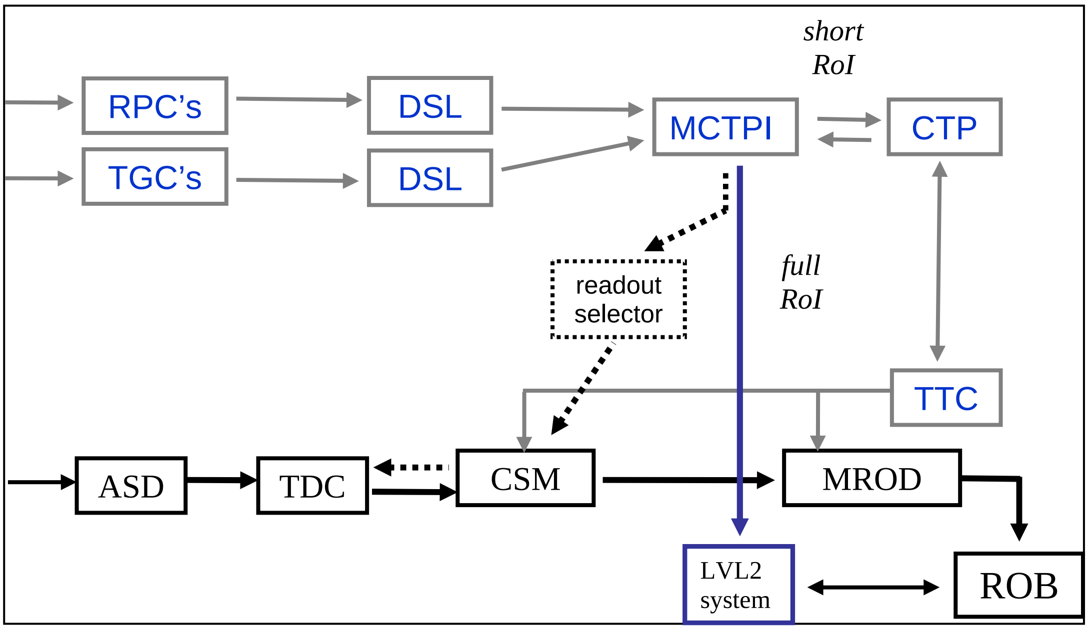

Fig. 1: Schematic diagram of the Muon trigger and the MDT readout system. In the selectic readout scheme the *full* RoI list, prepared by the MuCTPI for the LVL2, would be used to determine the MDT chambers belonging to these RoI, enabling them for data readout. The information in this list would have to converted into confirmation resp. veto signals, directly transmitted to the chambers via a system of fast links.

To also reduce the data transfer rate on the links between TDCs and CSM a further possiblility may be explored. The LVL1 transmission from CSM to TDCs, triggering the data transfer in the TDCs, might be withheld until reception of a confirmation signal. Therefore, in most cases, the TDCs would send no data, and the CSM would not fill up the derandomizer. However, this would probably require modification of the TDC, as data may have to be stored for a longer period, due to the delayed arrival of the confirmation signal w.r.t. the LVL1 signal (up to 100 μs). The extended storage time may exceed the storage capacity of the TDC and call for a new, more performant device.

Summary

In view of a SLHC upgrade the selective readout scheme has the potential for a large reduction of transfer rates and data volumes, resulting in considerable cost reduction. Many components of the present readout system could remain unchanged, in particular the off-detector components, i.e. MROD and ROB, which otherwise would not have to be upgraded for more storage space and processor speed. The fiber links from CSM to MROD and from MROD to ROB could also remain unchanged, having to transmit only < 5% of what is necessary with full readout.

At the frontend, the readout logics can be adapted to the new scheme because of the FPGA-based design of the CSM. However, an additional (optical) input will have to be implemented, which probably leads to a considerable modification of the existing CSM hardware. It might be possible to avoid a separate input, using the B-channel of the TTC to transmit the confirmation signal. Whether this is a realistic option in the present TTC scheme will have to be explored.

As mentioned above, the TDCs most probably need replacement, for reasons of storage capacity and radiation tolerance.

A number of important performance figures and operation parameters for the selective readout, like efficiency, timing, protocol would have to be worked out in detail. Using simulated events, the data reduction potential of different selection algorithms could be analysed.- In order to optimize the MDT readout, a detailed simulation of timing and data flow has to be done.

A close collaboration with the TDAQ community would be a necessity for a detailed definition of the readout concept and the final design of a selection processor. The total amount of resources, labour and cost, would have to be analysed.

Appendix

Muon rates in the ATLAS detector

This is a very simplified estimate of the muon multiplicities at SLHC, based on the muon rates given in Ref. 1 (Table 14-4, p. 391).

At a luminosity of 10^{33} $cm^{-2}s^{-1}$ and a total trigger rate of 100 kHz, the trigger rate for muons with $p_T > 6$ GeV is 23 kHz, corresponding to a probability of about 0.06% to find such a muon in any given beam crossing (23kHz/40MHz).- If an event was triggered by *one* muon, the probability to find a *second one* above 6 GeV would thus be 0.06%, which leads to an average muon multiplicity of 1.0006, assuming that these 2 muons are uncorrelated.- At 100 times this luminosity the probability to find any muon above 6 GeV would be 5% and the muon multiplicity of events triggered by a muon 1.06. More detailed estimates yield values of about 1.5 muons per event. For the motivation of the selective readout scheme, the exact value is not very important.

It should be noted that the actual threshold used in the LVL1 menu for the single muon trigger has no strong influence on the muon *multiplicity*, as several thresholds *below* the single muon trigger threshold are used in the RPCs and TGCs to define RoIs. On a LVL1 *all* RoI are transmitted to the LVL2, including those which would not have been able to trigger the event on their own. E.g.: if the single muon trigger threshold at SLHC was set to 30 GeV, the triggering muon as well as the other muons (corresponding e.g. to thresholds of 6 or 10 GeV) are assigned RoIs and are retained in the *full* list presented to the LVL2 (see e.g. Ref. 1, section 13.2).

Ref. 1: First-Level Trigger TDR, CERN/LHCC/98-14, june, 30, 1998