The MROD

The Read Out Driver for the ATLAS MDT Muon Precision Chambers

Design Review Report
Overview

Marcello Barisonzi, Henk Boterenbrood, Rutger van der Eijk, Peter Jansweijer,
Gerard Kieft, Jos Vermeulen
NIKHEF, Amsterdam, The Netherlands

Adriaan König, Charles Timmermans, Thei Wijnen
University of Nijmegen and NIKHEF, Nijmegen, The Netherlands

Scope and Summary.

This note is the Intermediate Design Report of the MROD (MDT Read Out Driver), currently being developed at NIKHEF and the University of Nijmegen for the read out of the MDT muon precision detector of the Atlas experiment at the LHC. It gives an overview the MROD functionality and the design of the MROD-1 prototype and it presents preliminary results on the MROD-1 performance.

In general the MROD-1 prototype behaves as expected. Final throughput numbers are still in the works. The MROD-1 performance may or may not turn out to be just sufficient. In any case MROD-1 spare capacity will be virtually non-existent. Alternative MROD solutions, which are basically an upgrade of the current MROD-1 design, have been explored and will be presented together with pros and cons.

1. Introduction.

The Monitored Drift Tubes (MDT chambers) form the main component of the Atlas muon precision detector. Three concentric layers of MDT chambers in the toroidal magnetic field form the muon spectrometer. The MDT system consists of some 1200 chambers and contains a total of about 300k individual drift tubes. The largest MDT chambers contain 432 drift tubes.

Directly on the MDT chambers, front-end cards equipped with 24-channel TDC-chips (Atlas Muon TDC, or AMT) digitize the wire signals. Each front-end card thus serves a maximum of 24 drift tubes and the largest MDT chambers have 18 front-end cards with as many TDCs. Upon receipt of a level-1 trigger signal, each TDC of a chamber sends its data over a serial 40 Mbit/s point-to-point connection to the on-chamber service module, the CSM [1]. The CSM in turn deserialises the 32-bit TDC data words and mutiplexes them into a 1 Gbit/s optical link (GOL) [2] to the MROD. The CSM acts as a time division multiplexer where the TDCs take turns in a circular fashion. Each TDC has its fixed time slot. If a TDC happens to have no data, a zero word is inserted by the CSM as a place holder. Note that the GOL is a strictly one way connection.

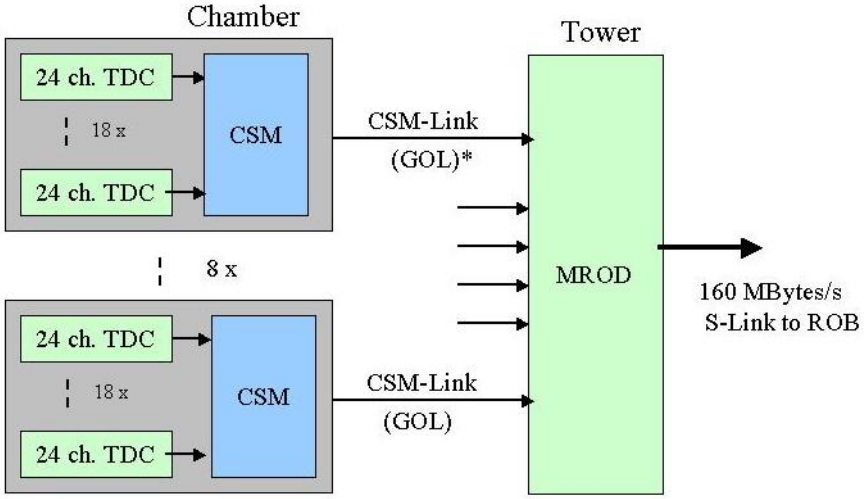

Figure 1. Overview of MDT read-out chain.

The task of the MROD is to receive the data streams from five to eight MDT chambers. The MROD then builds the event fragments from the incoming data and sends these over an S-Link connection to a Read Out Buffer (ROB) for further processing. A schematic overview of the MDT read-out chain is given in figure 1.

2. MROD Functionality.

The main task of the MROD, located in the radiation shielded USA15 underground area, is to receive and demultiplex the data streams from the CSMs. One MROD receives data coming

from five to eight MDT chambers which together form a “tower”: a set of chambers, which cover a predetermined ΔφΔη bite of the full 4π solid angle. The MROD builds the event fragments from the incoming data and sends these over an S-Link connection to a Read Out Buffer (ROB). In addition, the MROD detects and reports errors and inconsistencies in the incoming data streams (and where possible initiates corrective action), it collects statistics and it allows to “spy” on the data. Moreover, zero suppression and data reduction schemes may be executed by the MROD.

The MDT detector consists of 1172 MDT chambers, which form 192 ΔφΔη towers. Of these towers, 164 comprise five or six MDT chambers, the other 28 towers consist of seven or eight MDT chambers. Clearly a total of 192 MRODs is needed. Since 85% of the towers consist of five or six chambers, it is an option to build separate 6-input and 8-input MRODs. One could also consider the model of all 6-input MRODs with optional 2-channel extensions. It might however be more economical to have all identical (8-input) MROD units. A final decision on this issue will be postponed to the mass production stage. The MRODs will be put in 9U VME64x crates in USA15. With 12 MRODs per crate, a total of 16 MROD crates is foreseen. (see figure 2.) For a detailed description of the tower partitioning and the allocation of crates in USA15 see [3].

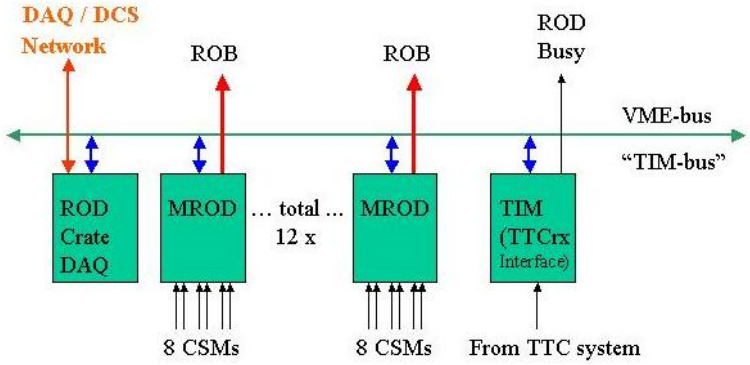

Figure 2: Schematic drawing of MROD crate.

3. Data rates.

Based on the most recent background calculations (assuming the full LHC design luminosity) the maximum output data rate for any single MDT tower (and hence for a single MROD) has been estimated to be 0.8 Gbit/s. Taking into account the canonical factor five uncertainty in these background calculations, this number grows to 1.4 Gbit/s [4]. This estimate has been based on a level-1 trigger rate of 100 kHz, a maximum drift time of 750 ns and a single TDC word for each wire hit.

4. Data format.

The event format for the MDT data has been described in detail [4]. The event format can best be visualized as consisting of three nested levels of “envelopes”. The lowest of the three levels is the TDC level, the next higher level is formed by the MDT chamber (≡CSM) level whereas the highest level corresponds to a tower of chambers (≡ MROD). Each envelope has

the same basic structure: one or more unique header words followed by a number of data words and closed with a unique trailer containing the word count for the envelope as a whole. For the TDC and chamber (or CSM) levels, the respective trailers include a 12-bit event number. Envelopes may be empty, i.e. contain zero data words. At each level the envelopes of the lower level are integrally included as data words in the envelope of the current level. The TDC envelopes are generated by the TDCs in the form of header and trailer words. In the absence of hits, the TDCs outputs an empty envelope for each level-1 trigger. Since the CSM is virtually transparent to the data, the chamber (or CSM) level envelope is generated by the MROD, as is the tower (or MROD) level envelope.

5. MROD-1 Prototype.

The MROD-1 was conceived as the first full size, full speed MROD. It provides for six CSM inputs. The MROD-1 makes extensive use of Analog Devices ADSP-21160 “SHARC” DSPs. Important feature of this DSP is the presence of six 40/80 MByte/s half duplex SHARC-links which allows fast point-to-point interconnections between the SHARCs. The data transfers take place under DMA control. The MROD-1 design relies heavily on the use of these SHARC-links.

The MROD-1 has a modular design with well separated input and output sections. Physically the MROD-1 prototype is built as a 9U VME64x module with three dual-input daughter boards, called the MROD-in, on a motherboard, called the MROD-out. Each MROD-in serves two CSM inputs and contains two Altera APEX 20K200 FPGAs (one for each input) and one SHARC. In addition the MROD-in has two dual ported 1 MByte memories; one for each input channel. The MROD-out contains two SHARC DSPs and one APEX 20K100 FPGA which connects to the common external bus of the two SHARCs, to the VME64x bus and to the S-Link interface. All SHARCs in the MROD-1 prototype are heavily interconnected by means of their SHARC-links. Towers of seven or eight chambers are accommodated by using two MROD-1 modules in a master-slave configuration. The slave’s MROD-out transparently connects one of its MROD-ins to the master’s MROD-out over the LVDS SHARC-link via the P3 connectors. This allows the use of two extra inputs. A schematic overview of the MROD-1 is given in figure 3.

The design philosophy behind the MROD-in is that, in the absence of any errors, the FPGA processes the data without intervention of the MROD-in SHARC. The FPGA demultiplexes the CSM data and thus reconstructs the data streams of the individual TDCs. The data are subsequently stored in the associated memory. On the fly the FPGA recognizes the TDC trailer words and records the event numbers. Once the FPGA has recognized a complete event (i.e. the trailer words of all active TDCs have been received), the data is transferred to the MROD-in SHARC. In this transfer, empty TDC envelopes may optionally be skipped as a means of zero suppression. The chamber (or CSM) level envelope is generated in this TDC data are then surrounded by a CSM level envelope which is generated. The MROD-in passes the data on to the MROD-out over one of its SHARC-links. The MROD-out immediately encloses the TDC data with the next envelope level, the chamber level.

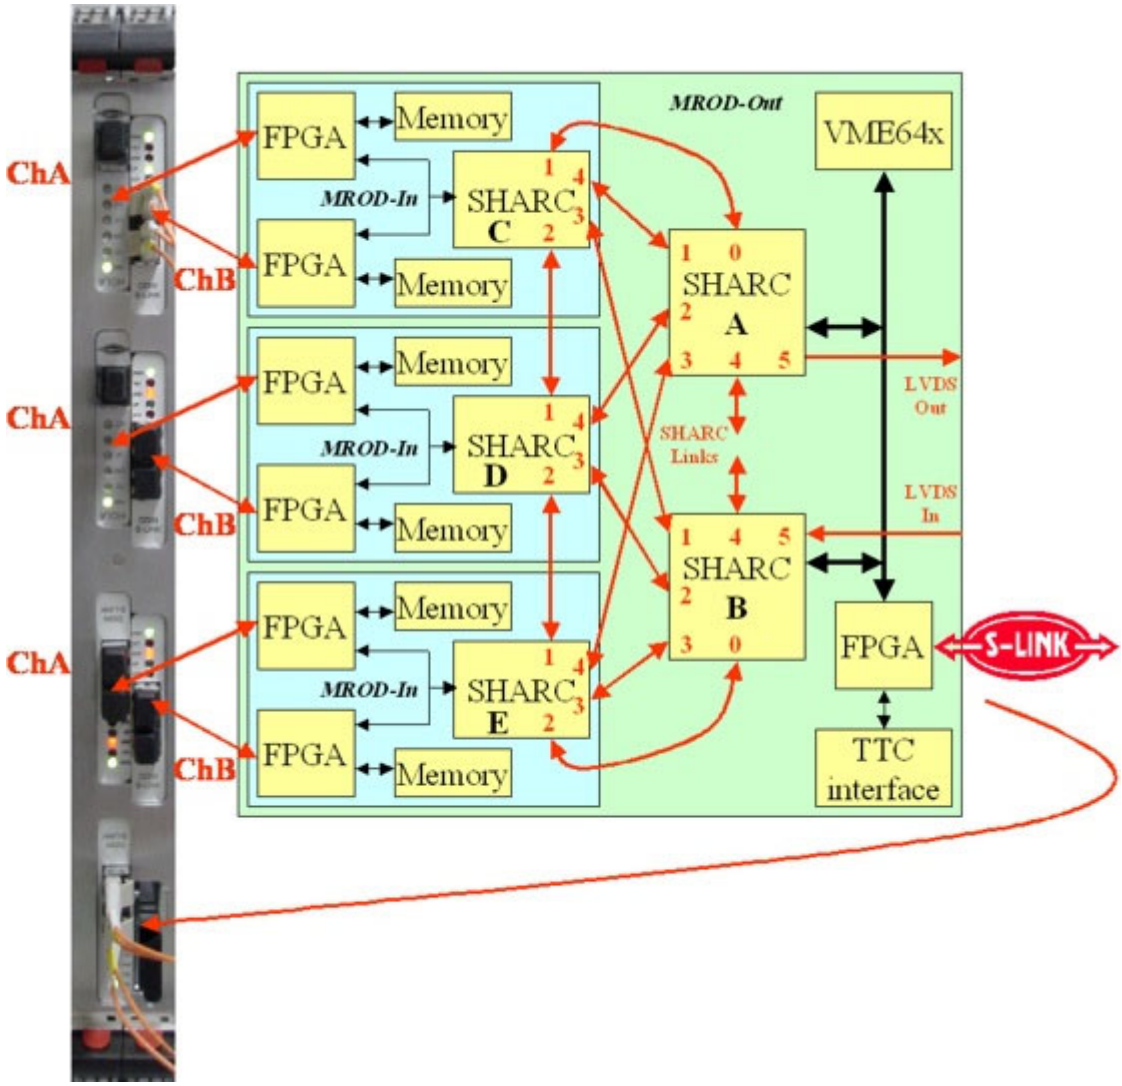

Figure 3. Schematic overview of the MROD-1 prototype.

The FPGA checks for a number of error and exception conditions:

- Parity errors on the TDC to CSM link (this condition is encoded in the data by the CSM)
- Link and/or parity errors on the CSM to MROD link (GOL)
- No data (each TDC outputs at least an empty envelope per level-1 trigger)
- Incorrect or too long event fragments from a TDC indicative of erratic behavior.

In case an error is detected, the FPGA may interrupt the SHARC which is then assumed to intervene appropriately. For very serious conditions (too large event fragment or a memory buffer overrun) the FPGA will independently decide to ignore (shut off) an individual TDC channel. In any case the SHARC will be notified. With this approach, the MROD's speed and throughput requirements are guaranteed by the FPGA. At the same time the SHARC, which can be programmed in the high level languages C or C++, allows for adaptability, flexibility and testability in dealing with errors and exception conditions.

The task of the MROD-out is to build the event fragments from the data received from the three MROD-ins and to transfer the fragments over the Read Out Link (S-Link) to the Read

Out Buffer (ROB). In the initial – low luminosity – phases of LHC running, the MROD-out may also perform monitoring tasks, collect statistics, copy events to the VME interface for spying purposes etc. Moreover the MROD interfaces to the TTC system and the Busy logic. The TTC and Busy functionality is realized through the TTC Interface Module (TIM) together with a customized P3 backplane in the VME crate.

The MROD-1 prototype has been used in the 2003 run of the H8 test beam. After a period of fierce debugging of the MROD's FPGA code and SHARC software, the MROD-1 has been used on a routine basis. In total two MROD-1 modules were used to read out 12 MDT chambers. In the test beam environment the event rate was very low. Rate tests of the current MROD-1 module are therefore performed in three principally different ways:

- Internal emulation: MROD-in SHARCs can been programmed to generate data for the MROD-out.
- External emulation: Independent external SHARC acts as a data generator which is then optically connected to an MROD input. Depending on the model of SHARC used, input event rates from 75 kHz up to 130 kHz can be realised.
- Hardware generator: A CSM like device (with a reprogrammed FPGA) is optically connected to the MROD input and generates TDC data with event fragment sizes according to a Poisson distribution. Very high input rates (well over 100 kHz) can be realized this way.

Combinations of the above three methods are also possible. Preliminary tests show single channel throughput rates of 90 kHz. The tests also indicate that the details of organization of the MROD software are extremely important. In particular the way DMAs are sequenced and initiated has large effects on the overall MROD performance. There is still room for further optimization. This includes the implementation of certain workarounds to counteract some of the SHARC's hardware shortcomings. Also a still more efficient organization of DMA in the MROD-out will likely improve the performance. The results of the rate tests are not final but the canonical 100 kHz event rate may be within reach. Tests are being continued. It is anticipated that by January 1st, 2004 final results will be available.

Note that at the full data rate, the MROD-out SHARCs are fully occupied with moving around the data and have very little processing power left for tasks like monitoring, spying and the collection of statistics.

6. Implementation choices.

In the current MROD-1 a number of implicit and explicit choices have been made which may need some justification:

- Application of Gigabit Optical Link (GOL) between CSM and MROD.
- Extensive use of SHARC DSPs.
- Use of TTC Interface Module (TIM).

The application of the Gigabit Optical Link is motivated by the fact that one wants to avoid any galvanic connection between the on-chamber CSM and the off-chamber MROD, which are some 100 meters apart. The GOL offers an optical connection with a bandwidth (1 Gbit/s) which nicely fits the maximum bandwidth of an MDT chamber equipped with 18 front-end cards (18 × 40 Mbit/s). Compared to an S-Link connection, the GOL has the advantage that it requires only half the number of optical fibers and transmitter/receiver pairs. Relative disadvantage of the GOL is that it is strictly uni-directional. The choice of the GOL therefore

precludes the use of an Xon/Xoff protocol or similar and implies that the MROD must be designed in such a way that in normal running conditions the data from the CSM will always be consumed. Since no direct communication from the MROD to the CSM is possible, any feedback has to go through the DCS and/or DAQ system and as a consequence will be rather slow. The CSM may in principle send control and or status information to the MROD but no protocol has been defined (yet ?).

The current 6-channel MROD-1 uses 7 large FPGAs in conjunction with 5 SHARC DSPs. The question has been raised whether this relatively large number of DSPs is required. The answer is threefold:

- The MROD heavily utilizes the SHARC links for the internal data forwarding.
- The SHARC can be programmed in a much more flexible way when dealing with errors and exception conditions than any practical FPGA code would be able to do. In practice any attempted "FPGA only" solution would not only be much less flexible but also require still larger – and hence more expensive – FPGAs.
- The DSP solution thus yields adaptability, flexibility and testability at no significant cost penalty.

The MROD makes use of the TTC Interface Module (TIM, designed by UCL) for the following reasons:

- The TIM offers all of the TTC functionality required by the MROD.
- One TIM serves the full MROD crate through a tailor made P3 backplane and thus eliminates the need for TTC receivers on each individual MRODs. Use of the TIM effectively reduces the number of required TTC fibers from 192 to 16.
- As a bonus, the TIM also implements the ROD Busy logic.

7. Preliminary MROD-1 performance evaluation.

The current, preliminary, results on MROD-1 performance indicate that a 100 kHz event rate may be realized. Although nominally sufficient, even this value does not make one particularly confident. At 100 kHz the MROD will not have much headroom for any additional tasks. Moreover, if the events turn out to be significantly larger then now expected, (eg. due to higher backgrounds, longer maximum drift time etc.) the MROD has hardly any spare capacity to cope with the increased data volumes.

We have therefore looked at alternatives. In particular we have looked at faster SHARC DSPs. In the current SHARC family (ADSP21160) only limited improvements are possible in the form of modest (~20%) increases in clock frequency. One alternative in the SHARC tradition is the so-called TigerSHARC (ADSP-TS20X) family, announced by Analog Devices. De differences between the parameters of the ADSP21160 SHARC and of the newly announced TigerSHARC DSP are substantial:

- More than 6 times higher clock frequency
- Up to 6 times more internal memory
- 4 full duplex SHARC links at nominally 5 times higher speed (The old and new SHARC links are not compatible, rendering any hybrid solutions useless).
- Internal buses are 128 bits wide (64 bits in ADSP21160)
- More DMA sources/destinations possible.

Usage of the TigerSHARC would have the additional advantage that it is a newer and more modern DSP and although there are no indications that Analog Devices wants to discontinue the ADSP21160, the TigerSHARC may be expected to have a longer lifetime than the ADSP21160. NIKHEF has ordered a TigerSHARC development kit for evaluation, which is due for delivery in December 2003.

Clearly there is also a number of disadvantages when one decides to migrate to the TigerSHARC:

- Interfacing to the pipelined external bus requires a non-trivial modification in the design of the FPGAs.
- Very likely an extra prototype cycle is required (some 9 to 12 extra months).
- New devices may (and as experience learns probably will) have undocumented features (hardware shortcomings which are euphemistically called “anomalies” by Analog Devices) which slow down the development work.

8. The next generation: MROD-2.

Pending final results on the performance, it is still too early to make a final proposal for the MROD-2 which was foreseen to be the so-called production module zero. Irrespective of the outcome of a possible decision to go for a TigerSHARC implementation, the MROD-2 will evolve from the current MROD-1 in the following rather straightforward ways:

- The MROD-2 will be one single PCB VME64x module, i.e. without any daughter boards (the only expection being the S-Link output board). This full integration eliminates a large number of connectors and therfore will reduce cost and increase reliability.
- The GOL receiver board (which currently rides piggy-back on the MROD-in with its own FPGA) will also be fully integrated on the main PCB and the GOL logic will be integrated in the MROD-in FPGA . This integrated FPGA code has already been designed.
- The MROD-2 will get four MROD-ins, thus accomodating up to eight MDT chambers per tower. The option of equipping only three MROD-ins with components allows the production of 6-input MRODs.

9. Conclusions.

In general the MROD-1 prototype behaves as expected. Final throughput numbers are still in the works. In the en the MROD-1 performance may or may not turn out to be sufficient. Spare capacity is virtually non-existent. Alternative MROD solutions which are in line with the MROD-1 design, have been explored and presented together with pros and cons. The MROD design team would appreciate advice from the Review Board on how to proceed from here.