

The Level-0 Muon Trigger for the LHCb Experiment

E. Aslanides ^a, J.P. Cachemiche ^a, J. Cogan ^a, P.Y. Duval ^a, R. Le Gac ^a,
O. Leroy ^a, P.L. Liotard ^a, F. Marin ^a, S. Favard ^a and A. Tsaregorodtsev ^a

^aCentre de Physique des Particules de Marseille (CPPM)
CNRS/IN2P3 – Aix-Marseille Université, France

cachemiche@cppm.in2p3.fr
legac@cppm.in2p3.fr

Abstract

The Level-0 Muon Trigger looks for straight tracks crossing the five muon stations of the LHCb muon detector and measures their transverse momentum. The tracking uses a road algorithm relying on the projectivity of the muon detector. The architecture of the Level-0 muon trigger is pipeline and massively parallel. Receiving 130 GBytes/s of input data, it reconstructs muon candidates for each bunch crossing (25 ns) in less than 1.2 μ s. It relies on an intensive use of high speed multigigabit serial links where high speed serializers/deserializers are embedded in Field Programmable Gate Arrays (FPGAs).

I. OVERVIEW

The muon system has been designed to look for muons with a high transverse momentum, a typical signature of a b-hadron decay. It is composed of the muon detector and the level-0 muon trigger.




Figure 2: The layout of the LHCb spectrometer [1] showing the electromagnetic (ECAL) and hadronic (HCAL) calorimeters as well as the five muon stations M1-M5 interleaved with iron filters.

A. Muon detector

The muon detector [2], shown in Figure 1 consists of five muon stations interleaved with muon filters. The filter is composed of the electromagnetic and hadronic calorimeters and three iron absorbers. Stations M2-M3 are devoted to the muon track finding while stations M4-M5 confirm the muon identification. The first station M1 is placed in front of the

calorimeter and plays an important role for the transverse momentum of the muon track.

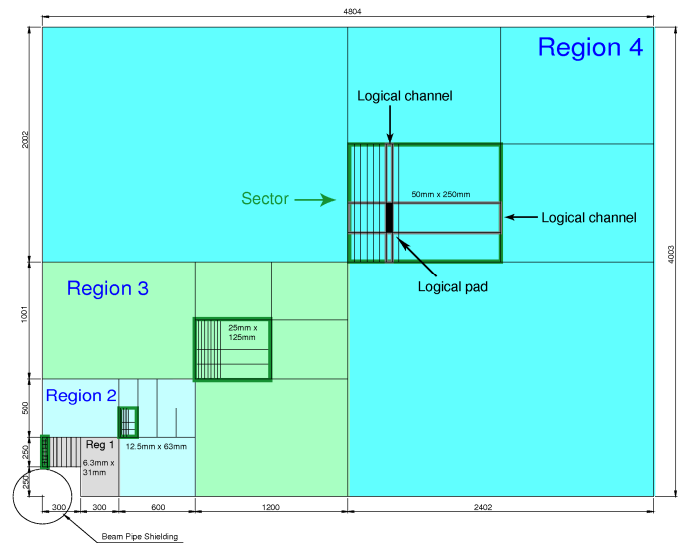


Figure 1: Front view of one quadrant of muon station M2 showing the dimension of regions.

Each station has two detector layers with independent readout. A detector layer contains two gaps in stations M2-M5. To achieve the high detection efficiency of 99% per station and to ensure redundancy, the signal of corresponding physical channels in the two gaps and two layers are logically OR-ed in the chamber to form a logical channel. The total number of physical channels in the system is about 120,000 while the number of logical channels is 25,920.

Each station is subdivided into four regions with different logical pad dimensions as shown in Figure 2. Region and pad sizes scale by a factor two from one region to the next. The logical layout in the five muon station is projective in y to the interaction point. It is also projective in x when the bending in the horizontal introduced by the magnetic field is ignored.

Pads are obtained by the crossing of horizontal and vertical strips whenever possible. Strips are employed in stations M2-M5 while station M1 and region 1 of station M4-M5 are equipped with pads.

Strips allow a reduction in the number of logical channels to be transferred to the muon trigger. The processor receives 25,920 bits every 25 ns forming 55,296 logical pads by crossing strips.

Each station is subdivided into *trigger sectors* as shown in Figure 2. They are defined by the size of the horizontal and

vertical strips and match the dimension of underlying chambers.

B. Level-0 Muon Trigger

The Level-0 muon trigger looks for muon tracks with a large transverse momentum, p_T . The track finding is performed on the logical pad layout. It searches for hits defining a straight line through the five muon stations towards the interaction point as shown in Figure 3. The position of the track in the first two stations allows the determination of its p_T .




Figure 4: Track finding by the Level-0 muon trigger.

To simplify the processing and to hide the complex layout of stations, we subdivided the muon detector in 192 towers pointing toward the interaction point as shown in Figure 5. A tower contains logical pads with the same layout: 48 pads from M1, 2×96 pads from M2 and M3, 2×24 pads from M4 and M5. Therefore the same algorithm can be executed in each tower. Each tower is connected to a processing element, the key element of the trigger processor.

The architecture of the Level-0 muon trigger is pipeline and massively parallel. Receiving 130 Gbytes/s of input data, it reconstructs muon candidates for each bunch crossing (25 ns) in less than 1.2 μ s.

C. Track Finding Algorithm

The track finding algorithm is illustrated in Figure 4. For each logical-pad hit in M3 (track seed), an extrapolated position is set in M2, M4 and M5 along a straight line passing through the hit and the interaction point. Hits are looked for in these stations in search windows termed Field Of Interest (FOI), approximately centred on the extrapolated positions. FOIs are open along the x -axis for all stations and along the y -axis only for stations M4 and M5. The size of a FOI depends on the station considered, the level of background and the minimum-bias retention required. When at least one hit is found inside the FOI for each station M2, M4 and M5, a muon track is flagged and the pad hit in M2 closest to the extrapolation from M3 is selected for a subsequent use.

The track position in station M1 is determined by making a straight-line extrapolation from M3 and M2, and identifying in the M1 FOI the pad hit closest to the extrapolation point.




Figure 3: The track finding algorithm.

Since the logical layout is projective, there is one-to-one mapping from pads in M3 to pads in M2, M4 and M5. There is also a one-to-one mapping from pairs of pads in M2 and M3 to pads in M1. This allows the track finding algorithm to be implemented using only logical operations.

D. A complex system

The implementation of the track finding algorithm is complex: huge number of logical channels distributed in a large volume; mixture of pads and strips; granularity of logical channels varying between regions and stations. Moreover, there is one-to-one correspondence between towers and trigger sectors except for station M2-M3 region R1 and R2. In region R1, a trigger sector is shared by two towers while in region R2 a tower maps two sectors, as shown in Figure 2.

Finally, processing elements have to exchange a huge number of logical channels to avoid inefficiency on the border of a tower. The topology of data exchange depends strongly on the granularity and on location of the tower.

II. ARCHITECTURE

An overview of the Level-0 muon architecture is given in Figure 5. Each quadrant of the muon detector is connected to a level-0 muon processor via 312 optical links grouped in 36 ribbons containing 12 optical fibres each. An optical fibre transmits serialized data at 1.6 Gbits/s on a distance of approximately 100 meters.

To collect data coming from a tower spread over five stations and to send them to a unique processing element we use a switch panel located close to the muon processor.

The processing elements run in parallel 96 tracking algorithms on logical channels coming from a tower. It is

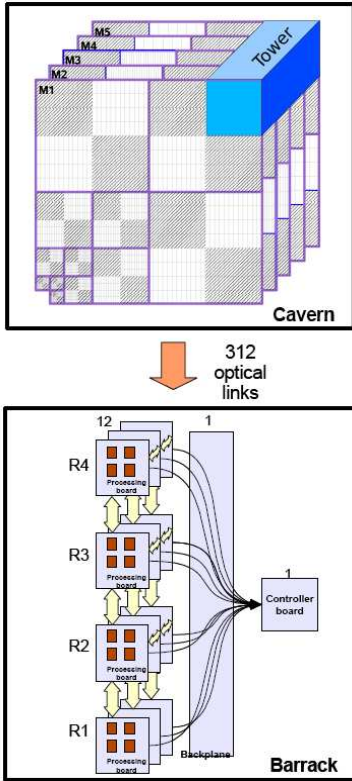


Figure 5: Overview of the Level-0 muon architecture for a quadrant of the muon detector.

implemented in a FPGA named processing unit (PU).

A processing board contains 4 PUs and an additional FPGA to select the two muons with the highest transverse momentum within the board, named BCSU (Best Candidate Selection Unit).

A muon processor houses 12 processing boards, a custom backplane and a controller board. The custom backplane is mandatory to exchange logical channels between PUs located on different processing boards.

A. Generic design

A single generic processing board was designed. The size of each connection between PUs has been maximized to accommodate all configurations. We use 40 MHz and 80 MHz parallel links as well as 1.6 GHz serial links. The resulting topology of the data exchange is shown in Figure 6. In such a framework, firmwares loaded in FPGAs differs according to the area covered by the board. We handle 48 configurations, one per tower of a quadrant.

B. Data flow within a processor

Logical channels are received over 312 optical links. The trigger aligns them in time. Neighbouring information is exchanged between PUs. Tracking algorithms are launched in parallel for the 96 seeds of a PU, which makes 4608 tracking algorithms on each quarter of detector. Each PU delivers a

maximum of two candidates to the BCSU. There is one BCSU per processing board. Each BCSU receives eight candidates and selects the two candidates with the highest p_T . At this stage only 24 candidates remain in a processor. The BCSUs send their candidates to the *control unit* (CU) and to the *slave unit* (SU) located on the controller board. The CU and the SU select the two candidates with the highest p_T and send them through two high speed serial links to the L0 Decision Unit.




Figure 6: Interconnections between processing units.

III. IMPLEMENTATION

Our implementation relies on the massive use of multigigabit serial links deserialized inside FPGAs. Processors are interfaced to the outside world via optical links while processing elements are interconnect with high speed copier serial links.

The number of pins available on standard high density connectors is not sufficient to transfer at 40 MHz the huge amount of logical channels required to run the track finding algorithm. Multiplexing the data at 80 MHz divides the number of connections by a factor two, but the routing density is very high and therefore sensitive to cross talk. For this reason we decided to mainly use 1.6 Gbits/s serial links for the interconnection between FPGAs.

A. Massive use of multigigabit serial links

By serializing most of the data exchanges at 1.6 Gbits/s we divide the number of connections by a factor 16. Sensitivity to cross-talk and to noise is decreased by a large factor since links are routed on differential lines. However, routing requires a lot of care since the geometry of the tracks must be totally controlled to guaranty a good impedance matching and to minimize electromagnetic emissions to the environment as well as sensitivity to electromagnetic perturbations from the environment.

A processing board embeds 92 high speed serial links while the backplane assuring the connectivity between the processing elements uses mixed technologies: 288 single-

ended links at 40 MHz and 110 differential serial links at 1.6 Gbits/s.

B. Processing board

The synoptic of the processing board is shown in Figure 7. It contains five FPGAs from the Stratix GX family [3]: four PUs and one BCSU. High speed serializer/deserializer are embedded in FPGAs. The board sends data to the data acquisition system through a classical L0-buffer/L0-derandomizer mechanism [4]. The processing board is remotely controlled via Ethernet by an embedded PC [5] running Linux.

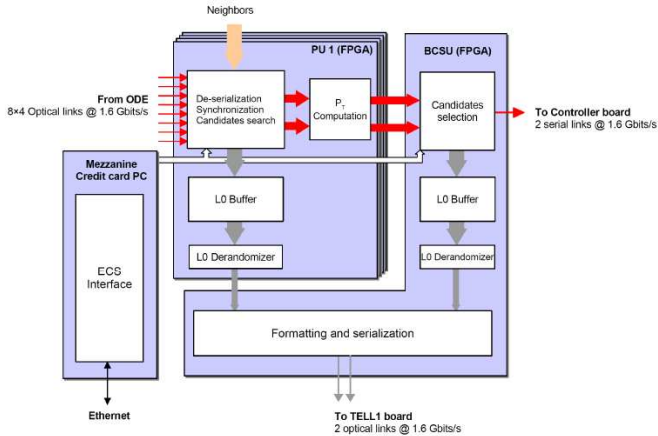


Figure 7: Synoptic of the processing board.

The board shown in Figure 8 is implemented in a 366.7×220 mm format. The printed circuit board is




Figure 8: Top view of a processing board.

composed of 18 layers where 1512 components are mounted. The power consumption is lower than 60 W.

C. Controller board

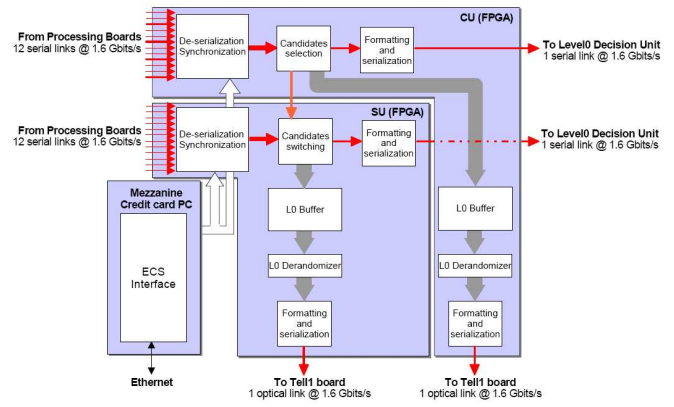


Figure 9: Synoptic of the controller board.

The synoptic of the controller board is shown in Figure 9. It contains two FPGAs from the Stratix GX family: one CU and one SU. High speed serializer/deserializer are embedded in FPGAs. The board shares a lot of common functionalities with the processing board: same embedded PC, same L0-buffer/L0-derandomizer mechanism. The role of the board is (i) to broadcast all the control, clock and reset signals received from the TTCrx [6], (ii) to collect the best candidates coming from the processing boards, and (iii) to select the two with the highest p_T and to send them to the L0 Decision Unit.




Figure 10: Top view of the controller board.

The board shown in Figure 10 is implemented in a 366.7×220 mm format. Its printed circuit board is composed of 14 layers where 948 components are mounted. The power consumption is lower than 50 W.

D. Backplane

The backplane contains 15 slots: 12 for the processing boards, one for the controller board and two for test. The latter allow to check a processing board by looping its outputs on its inputs.

The backplane shown in Figure 11 is implemented in a 395.4×426.72 mm format. Its printed circuit board is composed of 18 layers.




Figure 11: Top view of the backplane.

IV. DEBUGGING AND VERIFICATION TOOLS

Each board embeds a PC running Linux which is interfaced to FPGAs by a custom 16 bits bus. Therefore, we can control the operation of any FPGA of the system and monitor its status through error detection mechanisms, error counters, spy and snooping mechanisms.

A. Testability

To test a complete processor in a stand alone mode, we implement data injection buffer which can be substituted to the input data. Results of the processing can be read back by the embedded PC at the output of the L0-derandomizer buffer as shown in Figure 12.

B. Verification

The level-0 muon trigger is a very complicated system. Any malfunction can therefore be difficult to understand and interpret. At each stage we log the input and results of the processing. A PU produces 512 bits per bunch crossing resulting to 2,300 bits for a board and to 30,000 bits for a

processor. To handle such a quantity of information, we have developed a software emulator reproducing the behaviour of the hardware on a bit to bit basis. By comparing hardware reports with outputs of the emulator run on the same input data, we can point out any faulty components.

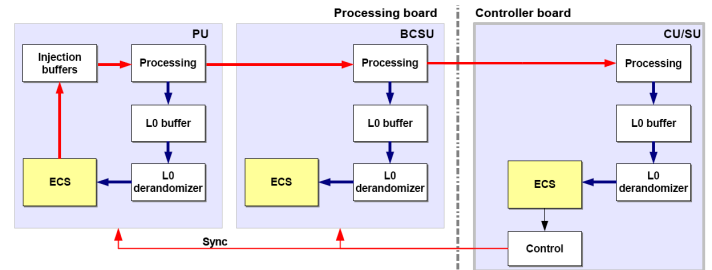


Figure 12: Data injection buffer and read back system.

V. CONCLUSION

The level-0 muon trigger is a very complex system. We have developed a very challenging architecture based on 248 large FPGAs exchanging data on 2272 high speed links at 1.6 Gbits/s among which 920 component-to-component or board-to-board copper links. 18432 tracking algorithms are permanently running in the system and deliver a result every 25 ns.

All tests demonstrate that the hardware is functional. We have checked, exhaustively, data path integrity by programming FPGAs in a mode where every possible output is toggled at maximum rate. The large number of multigigabit serial links proves to be very reliable.

A very good visibility on each internal node of such a complex system is mandatory. We have developed powerful hardware/software test mechanisms. Their purpose is to help us to understand and to validate the trigger behaviour during the debug and commissioning phases.

VI. REFERENCES

- [1] *LHCb Reoptimized Detector Design and Performances Technical Design Report*, LHCb, CERN, LHCC 2003-030
- [2] *LHCb Muon System Technical Design Report*, LHCb, CERN, LHCC 2001-010
- [3] *Stratix GX Device Handbook – Volumes 1, 2 and 3 – June 2006*
- [4] *Requirements to the L0 front-end electronics*, J. Christiansen, LHCb 2001-014
- [5] *SmartModule SM520PC Data sheet v1.1 from Digital Logic*
- [6] *TTCrx reference manual – A timing, Trigger and Control Receiver ASIC for LHC detectors – J. Christiansen, A. Marchioro, P. Moreira and T. Toifl – December 2005 – version 3.11*