

The Level 0 Trigger Decision Unit for the LHCb experiment

H. Chanal, O. Deschamps, R. Lefèvre, M. Magne, P. Perret

Laboratoire de Physique Corpusculaire de Clermont Ferrand
IN2P3/CNRS
24 avenue des Landais
63177 Aubièrre

chanal@clermont.in2p3.fr

Abstract

The Level 0 Decision Unit (L0DU) is one of the main components of the first trigger level (named level 0) of the LHCb experiment. This 16 layers custom board receives data from the calorimeter, muon and pile-up sub-triggers and computes the level 0 decision, reducing the rate from 40MHz to 1MHz. The processing is implemented in FPGA using a 40MHz synchronous pipelined architecture. The L0DU algorithm is fully configured via the Experiment Control System without any firmware reprogramming. An overall L0DU latency of less than 450ns has been achieved. The board was installed in the experimental area in April 2007 and since then has played a major role in the commissioning of the experiment.

I. INTRODUCTION

The LHCb experiment [1] is dedicated to b physics. It is installed at one interaction point of the Large Hadron Collider (LHC) at CERN. It is designed to exploit the large number of $b\bar{b}$ -pairs produced in pp interactions at $\sqrt{s}=14$ TeV at the LHC, in order to perform precise measurements and to search for new physics in CP asymmetries and rare decays in b -hadron systems. As the b and \bar{b} are produced at small angles and correlated, the detector has been designed as a single arm spectrometer. Figure 1 shows the layout of the experiment. The Vertex Locator and the tracking system (TT, T1-T3) provide very good vertexing and tracking capabilities while excellent particle identification is achieved thanks to two ring imaging Cherenkov detectors (RICH1 and RICH2), to the calorimeters (SPD/PS, ECAL and HCAL) and to five muon stations (M1-M5).




Figure 1: the LHCb detector

The interesting b decays account for a very small fraction of the 10MHz of visible interactions (around 1Hz for a branching ratio of 10^{-4}). In order to get an accurate selection of the events, a high performance versatile trigger has been developed.

This contribution will first introduce briefly the LHCb trigger with an emphasis on the level 0. The level 0 decision unit (L0DU) board and its internal processing will then be presented. The last part will finally focus on the project from the first prototype to the first data.

II. OVERVIEW OF LHCb TRIGGER

The whole LHCb detector runs with a 40MHz clock. It is not possible to store the data at such a high rate and most events are useless for the physics analysis (for example when there is no collision). Figure 2 presents the two stage trigger [2] which has been developed in order to reduce the rate from 40MHz to 2kHz for persistent storage. The first level is based on custom electronic and has to reduce the rate to 1MHz with a fixed latency of $4\mu s$. The second level (HLT) is a cluster of about 2000 PC which further reduces the rate to 2kHz. A lot of flexibility and good performances are needed at both stages.

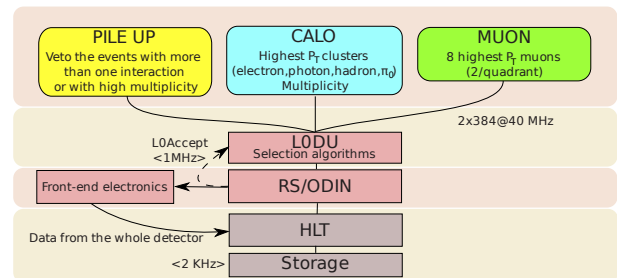


Figure 2: the LHCb trigger

III. THE L0 TRIGGER

Only the fastest sub-detectors can take part in the L0 event selection:

- The pile-up trigger which sends the reconstructed primary vertexes to be able to remove events with more than one interaction;
- The calorimeter trigger which sends the highest E_T γ , electron, π^0 , and hadron as well as the $\sum E_T$ and SPD multiplicity;
- The muon trigger which sends the two highest p_T muons per quadrant.

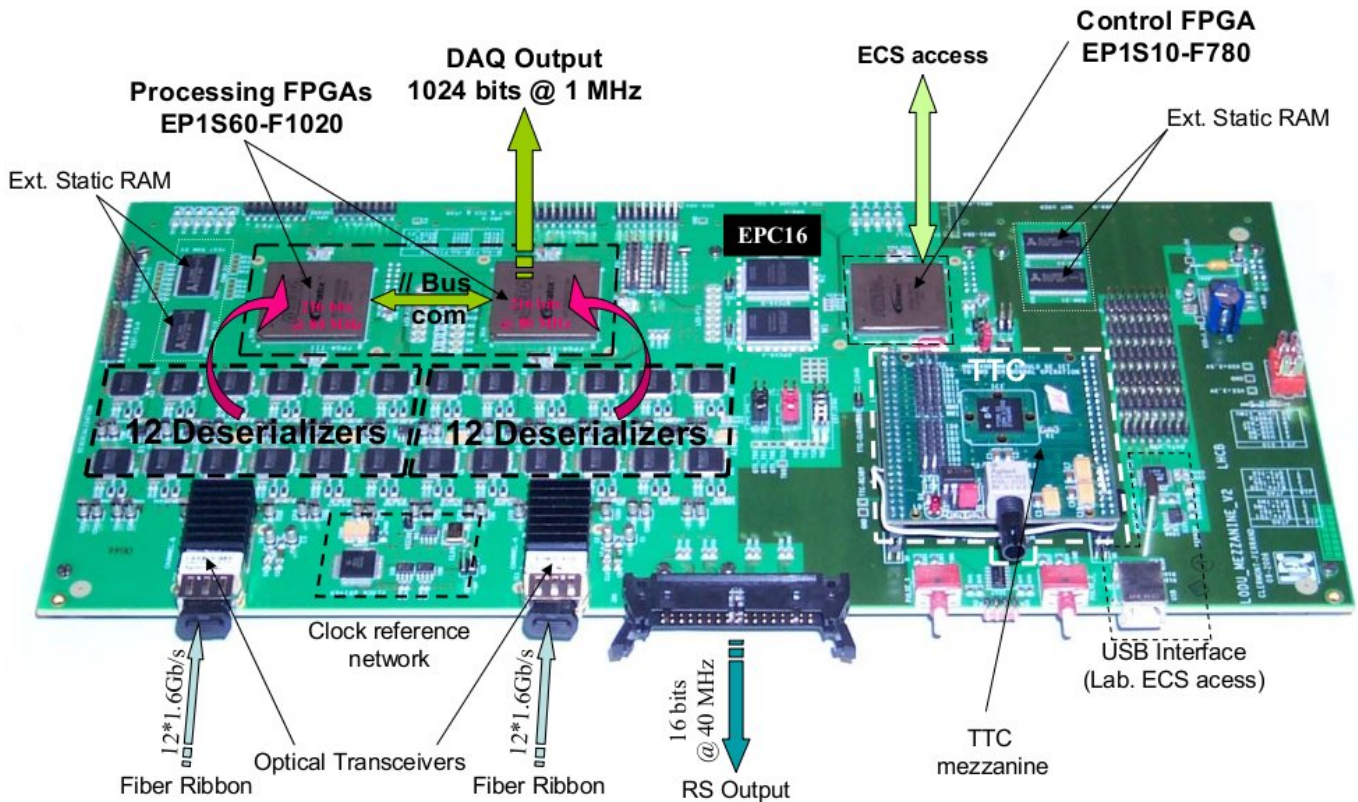


Figure 3: the L0DU mezzanine

The data are merged and processed in the L0DU. If the data fulfil a set of simple conditions the L0DU issues a validation signal sent to the Readout Supervisor (RS) [3][4] where it can be broadcasted (L0Accept signal) to the whole experiment. The data from the whole detector is then sent to the HLT farm for the processing of the next trigger level. To ensure a high flexibility of the L0 trigger, the conditions to be used in the L0DU are fully configurable.

IV. L0DU ARCHITECTURE

A. TELL1

As shown on Figure 4, the L0DU is a mezzanine of the TELL1 board [5][6]. The TELL1 has been designed for the LHCb experiment to handle the DAQ output and the common part of the interfaces. The data is sent to the DAQ via a Gigabyte Ethernet mezzanine. The TELL1 also provide an Experimental Control System (ECS) access using a small embedded Credit Card PC (CCPC) running a Linux system connected with standard Ethernet.

In our case, the TELL1 allows to have a remote access to the board registers with a dedicated I^2C bus and to remotely reprogram the 3 L0DU FPGA with a JTAG bus.

B. The L0DU mezzanine

Figure 3 shows the L0DU mezzanine. It is a 16 layers 9U board. It relies on three FPGA. The smaller (EP1S10 from Al-

tera) is used to access to the registers and for synchronization tasks while the two bigger (EP1S40 from Altera) are doing the processing. One of the processing FPGA (FPGA1) deals with the calorimeter and pile-up inputs. The other (FPGA2) copes with the muon trigger inputs. The core of the L0 algorithm is executed on FPGA1 where all the relevant data are centralized. FPGA1 is the most heavily used of the two processing FPGA: more than 65% of its logical resources are used.




Figure 4: the L0DU mezzanine plugged on a TELL1 board




Figure 5: functional schematic of the L0DU

The optical part is composed of 24 deserializers (TLK2501 from Texas Instrument) and two optical transceivers (HFBR-782BE from Agilent). It allows the connection of two fiber ribbons of 12 optical fibers each. 7 single optical fibers are used by the calorimeter trigger, 2 by the pile-up trigger and 8 by the muon trigger. There are 7 spares. The links between the optical transceivers and the deserializers are running at 1.6GHz. Between the deserializers and the FPGA, a 384 data bit bus is running at 80MHz. This part of the PCB has required special care and an accurate simulation with the SpectraQuest software from Cadence.

The clock and the L0Accept signals are received by an embedded TTC mezzanine [7]. The clock is broadcasted to the three FPGA using a dedicated LVDS network while the synchronization signals are sent to the control FPGA where they are treated and sent to the two processing FPGA.

The L0DU is linked to the TELL1 using 200 pins connectors. Only two of the four processing FPGA of the TELL1 are used as the data is sent by the two L0DU processing FPGA.

V. L0DU PROCESSING

The processing done in FPGA1 and FPGA2 can be decomposed in several blocks as shown on Figure 5. First the data coming from the sub-trigger systems are treated by pre-processing block which include the time alignment and some

data sorting. The decision making block computes the decision and constitutes the core of the L0 algorithm. It is highly configurable and can be easily changed remotely without any change in the FPGA firmware.

VI. PRE-PROCESSING

A. Time alignment

The time alignment can be decomposed in two independent parts. In the first stage, the data coming at 80MHz from the 24 optical deserializers have to be demultiplexed and put in the same 40MHz clock domain. In our case, we have a lot of incoming clocks (24) which can not be routed in our FPGA clock networks. Figure 6 presents the acquisition principles. We use a single 160MHz internal clock to acquire at both raising and falling edges the incoming data from the sub-detector inputs. A multiplexer allows the selection of a given edge. In a second step, the 16 bits LSB and MSB are acquired using two enable signals, one being delayed by 12.5ns with respect to the other. The result is then resynchronized with the local 40MHz clock.

To get the right phase for the acquisition of the incoming data, the corresponding clock is acquired by steps of 3.125ns with the local 160MHz clock (using both edges). Each value is digitalized 256 times and averaged to get an accurate representation of the incoming clock cycle. According to the digitalization

results, a LUT allows an automatic choice of the edge and of the enable signal used in the data acquisition module.

In the second stage, all the sub-detector data are aligned on the same event. 24 dual port RAM with a depth of 256 are used to introduce the necessary configurable delays.




Figure 6: clock synchronization principle

B. Data sorting

In FPGA2, the muons are sorted using a merge sort algorithm in three steps. It allows the processing of the data in one clock cycle using simple comparators and an associated multiplexer. Only the three largest muons are sent to FPGA1.

VII. DECISION BUILDING

The decision building flow is given on Figure 7.

A. Compound data

Compound data are created by combining the sub-trigger inputs. It can be either the sum or the difference of two elementary data, such as the E_T of two calorimeter candidates or the p_T of two muon candidates, or a mask applied on the address of

a candidate. 36 pre-synthesized blocks containing one of these operations are available in the latest LODU firmware.




Figure 7: decision building flow

A 8 bit wide and 3564 depth RAM has been introduced to eventually apply different L0 conditions according to the position of the event within the LHC cycle.

B. Elementary conditions

An elementary condition block has been designed to define simple cuts on sub-trigger data. Each block is the combination of a data input, an operator ($>$, $<$, $=$, \neq), and a threshold. 128 elementary condition blocks are available in the LODU.

C. Trigger channel

Elementary conditions are combined in so called trigger channels, each trigger channel being an "and" of any of the 128 elementary conditions. Up to 32 trigger channels can be defined.

D. Decision

The L0 decision is defined as an "or" of any of the 32 trigger channels.

E. Special trigger bit

Two flags are implemented :

- A force trigger bit (FTB) which indicates a problem in the data time alignment or an error in the L0 processing either from the LODU itself or from a sub-trigger. This flag may be used to force the storage on disk of the event for further analysis.
- A timing trigger bit (TTB) which is used to flag special se-

quences in a time window of ± 2 bunch crossings around the current bunch crossing, for instance to get isolated events. Here two modes are possible. The first option is to base the flag on the L0 decisions obtained for any of the 5 bunch crossings. The second is to look at the results, for any of the 5 bunch crossings, of simple comparisons of the $\sum E_T$ input coming from the calorimeter with two programmable thresholds.

VIII. TEST BENCH

We developed at the laboratory a dedicated test bench in order to test the firmware and stress the various links such as the optical fibers. A synoptic of the test bench is given on Figure 8.




Figure 8: the L0DU test bench

This test bench use two PC to control the boards: the first one runs the user interface while the second one provide access to the various registers and handles the booting of the CCPC. A RS board broadcasts the clock and the synchronization signals to the whole test bench. A specifically developed board, the GPL (Figure 9), is used to send known patterns in optical format and to receive the decision via a 16 bit LVDS link.




Figure 9: the GPL board

IX. PROTOTYPING AND COMMISSIONING

Three evolutions of the board have been designed. The first prototype [8] was produced in 2001 to validate the various con-

cepts. The second prototype [9] was built in 2005 and had all the required functionalities. To cope with the expected algorithmic flexibility, the FPGA size has been increased in the final boards which were received in 2007.

The final L0DU board was installed in the experimental area and connected to the RS in February 2007. The first combined tests were made with the calorimeter trigger in April 2007. The calorimeter system together with the L0DU triggered on their first cosmics at the end of 2007. Cosmics involving both the muon trigger and the calorimeter trigger have been recorded in April 2008. The L0 system then provides the triggers on the first beam induced particles in August 2008. Lastly, the pile-up joined the L0 trigger path in December 2008.

X. CONCLUSION

A very flexible L0 trigger board has been developed. It is in use in the experimental area since 2007. Specific algorithms have been extensively used to commission the detectors of LHCb and to record millions of cosmics and even thousands of VELO tracks during test of the transfer line from the SPS to the LHC.

The L0DU board is ready for the beam restart in November 2009.

REFERENCES

- [1] The LHCb Collaboration, The LHCb detector at the LHC, JINST 3 (2008) S08005
- [2] The LHCb Collaboration, Trigger Technical Design Report, CERN-LHCC/2003-31 (2003)
- [3] R. Jacobsson et al., Readout Supervisor Design Specifications, LHCb Note 2001-12 (2001)
- [4] R. Jacobsson et al., Timing and Fast Control, LHCb Note 2001-16 (2001)
- [5] G. Haefeli, Contribution to the development of the acquisition electronics for the LHCb experiment, PhD Thesis, Ecole Polytechnique Federale de Lausanne, 2004
- [6] G. Haefeli et al., The LHCb DAQ interface board TELL1, Nucl. Instr. and Meth. A 560, 494 (2006)
- [7] J. Christiansen et al., TTCrx Reference Manual, CERN-EP/MIC (http://ttc.web.cern.ch/TTC/TTCrx_manual3.10.pdf)
- [8] R. Cornat, Conception et réalisation de l'électronique frontale du détecteur de pied de gerbe et de l'unité de décision du système du premier niveau de déclenchement de l'expérience LHCb, PhD Thesis, Université Blaise Pascal Clermont-Ferrand 2, 2002
- [9] J. Laubser, Conception et réalisation de l'unité de décision du système de déclenchement de premier niveau du détecteur LHCb au LHC, PhD Thesis, Université Blaise Pascal Clermont-Ferrand 2, 2007