

# Recent Developments on the ALICE Central Trigger Processor

A. Bhasin<sup>1</sup>, D. Evans<sup>1</sup>, G.T. Jones<sup>1</sup>, P. Jovanović<sup>1</sup>, A. Jusko<sup>1</sup>, I. Králik<sup>2</sup>, M. Krivda<sup>2</sup>, R. Lietava<sup>1</sup>, B. Pastirčák<sup>2</sup>, L. Šándor<sup>2</sup>, J. Urbán<sup>3</sup> and O. Villalobos Baillie<sup>1§</sup>, for the ALICE collaboration.

<sup>1</sup>School of Physics and Astronomy, The University of Birmingham, Birmingham B15 2TT, UK

<sup>2</sup>The Institute of Experimental Physics of the Slovak Academy of Sciences, Košice, Slovakia

<sup>3</sup>P.J. Šafárik University, Faculty of Science, Košice, Slovakia

§ ovb@hep.ph.bham.ac.uk

## Abstract

The ALICE Central Trigger Processor has been constructed and tested, and will shortly be installed in the experimental area. In this review, we introduce the new developments in hardware and software, present a measurement of the minimum propagation time, and illustrate various trigger applications.

## I. INTRODUCTION

The ALICE detector has been designed to operate efficiently under a wide variety of different running conditions. The principal design aim is to allow the study of heavy ion interactions, which are characterized by very high multiplicities. The luminosities are modest compared to those for pp interactions. These two considerations determined the choice of detector technologies for ALICE. In particular, the principal tracking detector is a Time Projection Chamber (TPC), a relatively slow device with a drift time of 88 $\mu$ s that is capable of recording very high track densities. ALICE also contains other detectors which read out more rapidly. However, given a maximum interaction rate not exceeding 200 kHz, it was decided that detectors would not be required to employ pipelining. For this reason, a complex system of dead-time management has been developed to permit these detectors to read out more frequently than the TPC; in addition, the latency of the fastest level of the trigger must be short, to allow its use with track-and-hold electronics.

The study of heavy ions collisions requires collecting considerable statistics in minimum bias events and those selected according only to centrality. At the same time, new trigger signatures have been proposed that have much lower rates, notably those involving high  $p_T$  leptons or high  $p_T$  jets. Triggers of these types are supposed to run concurrently in ALICE, leading to a very heterogeneous loading of the trigger system. These considerations mean that it is particularly important in ALICE to control the use of resources by different trigger types, allowing sufficient flexibility to weight access to resources according to physics needs, by downscaling and selective disabling of common triggers.

In the past year the construction of the ALICE Central Trigger Processor (CTP) has been completed, with trigger systems in operation at Birmingham and at CERN. In this review, we shall focus on the practical issues related to the commissioning of the trigger system, after a brief overview of the system and its implementation.

## II. OVERVIEW OF THE ALICE CENTRAL TRIGGER PROCESSOR

The principal features of the ALICE CTP have been described in some detail elsewhere [1,2], and have been discussed at previous workshops [3,4]. A brief description is given here.

The CTP is designed for up to 24 different detectors.. At present there are in practice seventeen detectors in the experiment. These can be grouped into up to six different partitions, referred to as trigger *clusters*. These can each be sent independent trigger signals. Cluster definitions are programmable.

Owing to the low latency for some of the detectors, there are three levels in the ALICE trigger. The fastest level of the trigger, **L0**, has a latency of 1.2  $\mu$ s (from interaction to the reception of the L0 trigger at the detector). The typical functions of this trigger level are (i) to provide a strobe to detectors with fast track-and-hold electronics, and (ii) to initiate BUSY for all detectors in an affected cluster. Rate reductions at **L0** are not expected to be high as not all relevant trigger input signals are available at this time. The next trigger level, **L1**, arrives at the detector 6.5  $\mu$ s after the interaction takes place. At the time of the **L1** decision most of the trigger inputs are available, and therefore major rate reductions can be made. The typical function for this level of trigger would be to store event data in an appropriate multi-event buffer. However, a final decision as to whether to send the data to the data acquisition and High Level Trigger (HLT) system cannot be made at this stage as the final outcome for the past-future protection (see below) is not yet known. The detector with the longest sensitive period in ALICE, the TPC, will register hits up to 88  $\mu$ s from the time of an interaction. Therefore, the final level of the trigger (**L2**), cannot be given until after this period.

The trigger decisions are made using a total of 60 trigger inputs (24 **L0**; 24 **L1**; 12 **L2**). Owing to the number of inputs at each level, some restrictions are needed on the allowed logic functions. All inputs can be combined in AND. In addition for 6 selected classes the complementary signal can be selected. Finally, for a small subset of inputs (4 **L0** inputs) a look-up table can be used to allow any arbitrary logic function. Four such functions can be defined. Two are used in the past-future protection counters and in the interaction record [1], and two can be added to the list of inputs. In this way, for example it is possible to include an OR of different inputs in the minimum bias interaction definition in low multiplicity pp interactions, in order to improve trigger efficiency.

In order to protect the detector from recording events which have significant levels of overlap from other interactions before or after the one selected by the trigger, a system of past-future protection has been implemented. The details of the circuit have been described in some detail previously [4]. Here we illustrate the functions of the circuit in different conditions.

In heavy ion collisions the rate is such that pile-up has a significant probability in the TPC (>50%), but only a small probability in the other detectors. The overlap of even two full centrality interactions in an event readout would make reconstruction very difficult. In these circumstances, it makes sense to protect the detector over an interval equal to twice the drift time of the TPC, and to reject cases where two interactions having high multiplicity occur inside the protection interval. A different counter is used for low multiplicity events, since here a higher number of superimposed events can be tolerated. Figure 1 shows a schematic diagram of a simple case.




Figure 1: The top trace represents a series of interactions. The middle trace indicates when the past/future protection (with a threshold of 1 and a protection interval  $\Delta t = 3$  b.c. is asserted, and the bottom trace shows the final result for an L0 trigger, taking past-future protection into account.

In pp operation, the expected rate is such that pile-up in the TPC is inevitable, but also tolerable owing to the much lower multiplicities. In this case the past-future protection circuit serves to ensure that there is no pile-up in other, faster readout detectors, such as those of the Inner Tracking System (ITS). In this case, using a much smaller time window, it is desirable to check that there is no second

interaction accompanying a trigger, with no classification according to multiplicity. It is straightforward to reconfigure the parameters so as to cover both cases.

### III. ALICE CTP HARDWARE

Physically, the ALICE CTP is made up of seven different types of 6U board.. There is a board for each of the three levels of trigger (L0, L1 and L2), one for BUSY handling, one to give the interface to the DAQ (INT), some fan-out boards (FO) to transfer trigger signals to the Local Trigger Units, and a board I2C which hosts an I2C multiplexer collecting signals used to monitor voltages on the trigger boards. These are sent on special I<sup>2</sup>C lines on the J2 part of the backplane. The signals are transferred to the INT board via the front panel, where they are decoded by an FPGA; the decoded values are accessible via VME.. With the exception of the I2C board, the design for these boards is similar. In each case an ALTERA EPM3512 FPGA provides the VME interface, and also controls the loading of the main logic FPGA. In most cases, the main logic chip is an ALTERA CYCLONE EPIC20 FPGA, and its configuration data are kept in the on-board flash memory (Am29LV081. Additional features include the provision of a ‘snapshot memory’ (two CY7C-1382 ICs, giving a combined memory of 1 Mwords, 32 bits wide). The snapshot memories, present on every board, are used to record unbiased samples of trigger inputs and decisions for short periods of time (up to about 25 ms i.e. 300 LHC orbits. It is used principally for debugging and development purposes. In addition, a PCF8591T ADC is used to monitor all the supply voltages and the status of fuses on the board. The data are transmitted to the ALICE DCS (slow control system) using the independent I<sup>2</sup>C link, as discussed above.




Figure 2: Preliminary measurement of the propagation delay between the arrival of a trigger signal on the L0 board and the production of a trigger on an LTU. (20ns per division)

Figure 2 shows a preliminary measurement of the full propagation delay through the CTP system, from the arrival of a trigger signal at the L0 board to the output of an L0 signal from an LTU board. A minimum latency of 117 ns was found. The contribution of the CTP itself to the trigger latency turns out to be about one bunch crossing less.

Figure 3 shows a recent trigger setup, involving a full set of CTP boards, as used at CERN to transmit trigger signals

to two detector partitions in the DAQ laboratory. This setup is discussed below.

The CTP communicates with each detector independently, since in general detectors may have different triggering patterns. The interface to each detector comes via another VME card, the Local Trigger Unit (LTU). This was described in detail in the 2004 workshop [3]. There is one LTU per detector. The LTU receives signals from the CTP and generates the signals necessary to

send to the detectors, via LVDS cables in the case of the L0 signal, and using the RD-12 TTC system for the other signals. It can also be used to emulate the CTP. Extensive facilities exist to emulate different trigger sequences starting in the LTU, including sequences with intentional errors. The LTU can also be used with an external trigger signal for test beam work.



Figure 3: Trigger setup at CERN, showing the full set of CTP boards, together with two sets of TTC interfaces, which transmit to detector emulators about 40m away.

All the VME cards were designed at the University of Birmingham and produced through the Rutherford Appleton Laboratory. 53 LTUs were produced, together with varying numbers of each type of CTP board, so as to allow for an adequate stock of spares. Production is now completed. Many LTUs have been delivered to ALICE detector groups to allow them to test their electronics.

After final tests of the system, three full trigger systems were shipped to CERN in July 2006. Two systems are being used for integrated trigger-DAQ-ECS tests on the CERN Meyrin site, and one is awaiting installation in the ALICE pit. In addition, two complete setups remain in Birmingham, where tests and training sessions are continuing.

The CERN tests use a set of cables and fibres equivalent to those used for connections with detectors in order to connect the trigger and DAQ laboratories. In a first test it was demonstrated that a detector simulator board (a DDG board), connected to the trigger system via a TTC

and DDL link, can successfully register trigger words and messages, transferring trigger data correctly to the event headers. The second test is discussed in the next section.

#### IV. ALICE TRIGGER SOFTWARE

The operation of the ALICE CTP and the LTUs requires the preparation of a large body of software, for example to

- (i) configure and load CTP parameters,
- (ii) execute, run control functions under the control of the ALICE ECS (Experiment Control System)
- (iii) provide monitoring facilities to check for the correct functioning of the trigger, and
- (iv) provide debugging facilities for when faults occur.

The software framework for the trigger was decided in 2003. Coding is written mainly in C, and SMI++[5], and either Python/Tk or Tcl/Tk is used for the graphical interface. DIM[6] is used for communication with the DCS and ECS systems. The operating system is LINUX, running on a VME processor (Concurrent Technologies VP315).

Flash memory is used for the configuration of the main FPGA, making it significantly easier to distribute firmware updates to users. (This is important, for example, in the case of LTUs, which are currently being used in ALICE institutes in several different countries.) The user can download the latest firmware versions from the trigger web pages, and load them via VME to the flash memory.

All the CTP boards can be configured and read out with existing software. For example, figure 4 shows a graphical interface used for configuration of trigger classes. Active classes (only) are shown as rows, with the selected trigger inputs and vetos shown according to the column used. This panel combines parameters affecting the three boards L0, L1 and L2. Output clusters are indicated by colours, with a text explanation available by hovering a mouse over the required box. Similar graphical menus exist for all the trigger functions, allowing trigger configurations to be defined and made. Help facilities are also provided. The storage of trigger configurations is being studied; it is expected that a database using MySQL, similar to that being developed by the ALICE ECS group, will be employed.




Figure 4: Graphical interface for configuration of trigger classes. In this example three trigger classes (numbers 1, 28 and 49) are in use. Colours are used to indicate the trigger condition; the colours on the far left column of the panel indicate the selected output trigger cluster.

There is a 1 MWord memory on each board, which is used to collect data on a bunch-crossing-by-bunch-crossing basis, creating an unbiased 'snapshot' of the trigger operation. The snapshots are read out through the VME backplane, and are used by online monitoring processors for the CTP; snapshots do not appear in the normal DAQ datastream. They will be heavily used for debugging purposes while the system is being commissioned, as they allow all aspects of the trigger operation to be followed. For many of these, a software simulation of the hardware already exists, allowing detailed checks. They can also be used for timing (alignment) checks, as described by Roman Lietava [7]

As an example of the use of a snapshot, we return to the example of past-future protection, introduced in Section II. Figure 5 shows a snapshot using events on the L0 board. Timings are given in bunch crossing numbers, i.e. in 25 ns increments. The top

line shows a randomly occurring trigger input which is included in the interaction definition but does not produce a trigger. The next line shows another interaction contribution which does produce a trigger. The interaction definition used for the past-future protection is the OR of these two contributions. In order to see the action of the past-future protection, two classes are activated, both with the same input conditions. In one case a past-future protection spanning 15 BCs is set; in the other there is no past-future protection requirement. The pulse at bunch crossing 1038 makes a potential trigger in both classes. However, the burst of activity in the random trigger immediately before (BCs 1031-1033) fires the past-future protection, shown in the bottom line in the snapshot. Thus the class with past-future protection (10c1st1) gives no trigger, as the trigger is vetoed by the past-future protection, while the class without past-future protection (10c1st4) does give a trigger.



Figure 5: L0 Board snapshot showing the action of past-future protection. The past-future protection circuit is activated by the OR of activity in the signals shown in the first two rows, but only the second row (*scaledbc2*) is used in the trigger. There are two classes with different output clusters. In one (*l0c1st1*) there is a past-future protection interval of 15 BC, while in the other (*l0c1st4*) there is not. The interval when the past-future protection is asserted is shown in the bottom row (*pf1*), and prevents a trigger from being generated in *l0c1st1*.

Software is also needed to provide the interface to the ECS system. This is under development, using the combined trigger-DAQ-ECS system at CERN for tests. The ECS system foresees splitting the ALICE detector into partitions, which are non-overlapping (no detector is allowed to be in more than one partition). Partitions have independent DCS, and can be operated independently as if they were separate sub-experiments. (This will be particularly useful during the commissioning phase.) Trigger clusters have a different logic, as they can be overlapping, so the trigger-ECS interface has had to be written so as to respect the constraints imposed by the ECS partitions. This has been tested. In a separate development, the start-of-run sequence, which is co-ordinated by the ECS, has been implemented. At start-of-run, the ECS ensures the detectors are ready (using the DCS system), the DAQ is ready and the trigger is set. When all is ready, a signal is sent to the CTP to generate a special “start-of-run” event, which is specially treated by the DAQ. A corresponding “end-of-run” event is sent to signal the end of data-taking. At present, runs can be started and stopped in the test system, and partitions of arbitrary nature can be configured, with the corresponding constraints taken into account in the allowed CTP cluster definitions.

## V. SUMMARY

Construction and testing of the ALICE CTP has been completed. Five full setups exist, allowing tests and training to proceed in parallel at CERN and Birmingham. The minimum L0 latency of the CTP has been found to be about 90ns, i.e. within the system requirements [8].

Software exists allowing the full configuration of the system, and monitoring tasks are under development. In particular, the use of ‘snapshots’, illustrated in this report, allows very detailed study of CP operation. The CERN tests are being used for integration of the CTP with the other ALICE control systems, and installation in the ALICE pit is expected to start by the end of October 2006.

## REFERENCES

- [1] ALICE Technical Design Report on Trigger, Data Acquisition, High Level Trigger and Control System, CERN-LHCC-2003-062, 7 January 2004.
- [2] The ALICE Central Trigger Processor web site: <http://www.ep.ph.bham.ac.uk/user/pedja/alice/>.
- [3] A. Jusko *et al.*, Proc 10th Workshop on Electronics for LHC and Future Experiments, Boston, MA, U.S.A. September 2004. CERN 2004-010, p. 277.
- [4] O. Villalobos Baillie *et al.*, Proc. 11th Workshop on Electronics for LHC and Future Experiments, Heidelberg, Germany, September 2005. CERN 2005-038, 284.
- [5] B.J. Franek and C. Gaspar, IEEE Transactions on Nuclear Science, **45**, 1998, 1946.
- [6] C. Gaspar, M. Dönszelmann and P. Charpentier, Comput. Phys. Commun. **140**, 2001, 102
- [7] R. Lietava, *Timing in the ALICE Trigger System*, these proceedings.
- [8] D. Evans and P. Jovanović, ALICE User Requirement Document, ALICE-INT-2003-055.