

# The ALICE Data-Acquisition system

T. Anticic<sup>(1)</sup>, F. Carena<sup>(2)</sup>, W. Carena<sup>(2)</sup>, S. Chapeland<sup>(2)</sup>, O. Cobanoglu<sup>(3)</sup>, E. Dénes<sup>(4)</sup>, R. Divià<sup>(2)</sup>, U. Fuchs<sup>(2)</sup>,  
 T. Kiss<sup>(4)</sup>, J.C. Marin<sup>(2)</sup>, I. Makhlyueva<sup>(2)</sup>, F. Ozok<sup>(3)</sup>, K. Schossmaier<sup>(2)</sup>, C. Soós<sup>(2)</sup>, P. Vande Vyvre<sup>(2)</sup>,  
 A. Vascotto<sup>(2)</sup>, S. Vergara<sup>(5)</sup>  
 (for the ALICE collaboration)

<sup>(1)</sup> Ruder Boškovi Institute, Zagreb, Croatia

<sup>(2)</sup> CERN, Geneva, Switzerland

<sup>(3)</sup> INFN Turin, on leave from Istanbul University, Istanbul, Turkey

<sup>(4)</sup> KFKI-RMKI, Budapest, Hungary

<sup>(5)</sup> Benemerita Universidad Autonoma de Puebla, Mexico

## Abstract

ALICE (A Large Ion Collider Experiment) is a general-purpose, heavy-ion detector designed to study the physics of strongly interacting matter and the quark-gluon plasma in nucleus-nucleus collisions at the CERN Large Hadron Collider (LHC). A large bandwidth and flexible DAQ is required to collect sufficient statistics in the short running time available per year for heavy ion and to accommodate very different requirements originated from the large set of detectors and the different beams used.

## I. INTRODUCTION

The ALICE experiment [1, 2] is designed to cope with the highest particle multiplicities anticipated for Pb-Pb reactions. In addition to heavy systems, the ALICE Collaboration will study collisions of lower-mass ions and protons. ALICE will operate in several different running modes with significantly different characteristics.

The experiment has been primarily designed to run with heavy-ion beams, which are characterized by relatively low rates (interaction rate lower than 10 kHz at design luminosity  $L=10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ ), relatively short running time (order of few weeks per year) but very high multiplicity and correspondingly large event size. The requirements on the lower level trigger selectivity are therefore modest compared to the other LHC experiments, whereas the trigger complexity is considerable and requires partial or full event reconstruction in the High-Level Trigger (HLT). A large bandwidth DAQ is required to collect sufficient statistics in the short running time available [3].

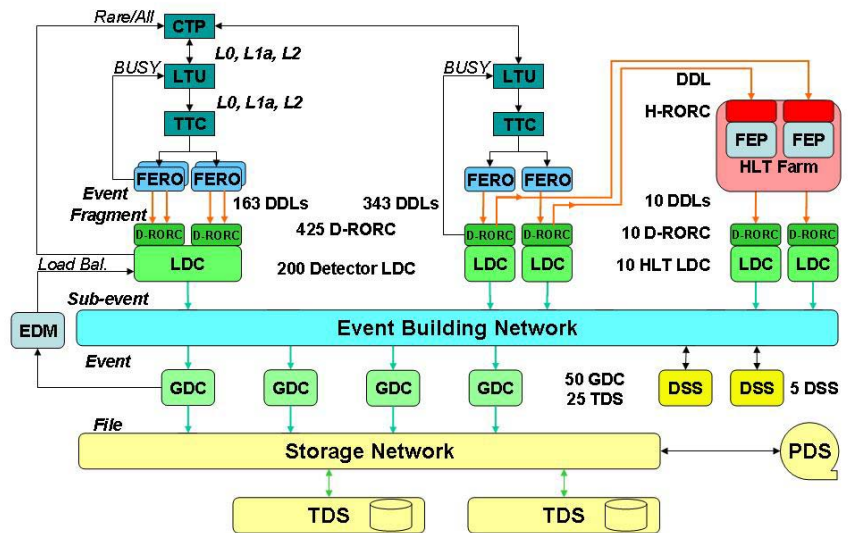
In pp or PA running modes, the interaction rates are much higher than in heavy-ion runs (up to 200 kHz, limited by pile-up in the TPC detector), whereas the event is small and the running time is of typically several

months per year in pp mode. Therefore the requirement on trigger selectivity is increased while requirements on trigger complexity and the DAQ bandwidth are much reduced.

## II. THE ALICE DATA ACQUISITION SYSTEM

The architecture of the ALICE data-acquisition system (DAQ) and High-Level Trigger (HLT) is shown in Figure 1. The detectors receive the trigger signals and the associated information from the Central Trigger Processor (CTP), through a dedicated Local Trigger Unit (LTU) interfaced to the Timing, Trigger and Control (TTC) system. The readout electronics of the detectors is interfaced to the ALICE-standard optical link, the Detector Data Link (DDL).

Figure 1: ALICE DAQ and HLT architecture



### III. DATA TRANSFER FROM THE DETECTORS

The DAQ system will use more than 450 DDLs to carry detector data from the front-end electronics directly into the memory of the first layer of computers called Local Data Concentrators (LDCs).

The interface between the I/O bus of the computers and the DDL consists of the DAQ Read-Out Receiver Card (D-RORC). The D-RORC includes two DDL interfaces, which can either receive detector data, or copy and transfer them to the High-Level Trigger system. The 64-bit/66 MHz PCI interface is realized by an IP core which is responsible for the PCI transaction management, as well as for the DMA arbitration on the local PCI buss(es).

The firmware also contains an embedded data generator, which can produce formatted data blocks of configurable length for testing purposes.

The DDL software consists of two parts: a Linux kernel device driver and an API library. Another Linux kernel device driver has been designed to access the physical memory of the PC. During the data-acquisition phase neither the D-RORC firmware nor the LDC software do polling via PCI during data transfer. Hence, the PCI bus is not occupied by dummy operations and is left available for data transfers.

### IV. DAQ SOFTWARE FRAMEWORK

All the hardware elements of the DAQ system are driven and controlled by the DATE software framework. DATE performs the dataflow and the control of the DAQ system.

#### A. Data-flow

The dataflow include the sub-event building, the event building and the load balancing described hereafter.

Each LDC can handle one or more D-RORCs. The D-RORCs perform concurrent and autonomous DMA transfers into the LDCs' memory with minimal software intervention. The event fragments originated by the various D-RORCs are logically assembled into sub-events by the LDCs. The role of the LDCs is then to ship the sub-events to a farm of machines called the Global Data Collectors (GDCs), where the whole events are built from all the sub-events pertaining to the same trigger. The event-building network is a standard communication network (commodity equipment) supporting the well-established TCP/IP protocol. The role of the GDCs is to collect the sub-events and assemble them into whole events. The GDCs also feed the recording system with the events that eventually end up in Permanent Data Storage (PDS).

#### B. Control and Configuration

The ALICE DAQ run-control is in charge of the configuration and the control of all the processes needed by the DAQ. It starts and stop these processes, configures and synchronizes them and detects error conditions. It also includes an interactive application for the operator.

Moreover, within an experiment, several data acquisitions can be performed at the same time: this is the case, for

example, of detectors being tested or calibrated while others collect data. The run control system handles the configuration and synchronization issues of complex systems.

### V. PERFORMANCE MEASUREMENTS

The performance of the DDL and D-RORC hardware, firmware and software has been measured on a test bed. This development system consists of high-end server PCs with up to six PCI-X slots and dual Xeon processors running at 2.4 GHz. The PCI slots were assigned to four PCI segments, which were attached to two controllers on the motherboard. Data blocks were generated by either the internal data generator of the D-RORC card or the external data generator, which was implemented in the DDL sender.

When adding data generators, the bandwidth increases steadily until it reaches its maximum. In case of the internal pattern generator, the 264 MB/s maximum bandwidth is defined by the PCI clock frequency (66 MHz) and the width of the internal data path (32-bit). In case of the DDL sender, the 206 MB/s maximum is determined by the DDL bandwidth. The results have shown that the bandwidth scales with the number of senders. The highest bandwidth of 1045 MB/s was achieved with four D-RORC cards installed on separate PCI segments.

Another application of the D-RORC card is to be the motherboard for the DDL Data Generator (DDG). It will be interfaced to the D-RORC using a daughter card attached to its CMC interface. The DDG has been designed to be a realistic data source for the DDL. It generates realistic detector events stored in the memory of a host computer by using the D-RORC DMA and it is triggered by the TTC.

Since several years, the ALICE team and the CERN/IT department have jointly executed several ALICE Data Challenges: periodic, large-scale, high throughput, and distributed computing exercises, whose goals are to prototype the ALICE DAQ system, to test hardware and software components, and to accomplish an early integration of the overall ALICE computing infrastructure. The highest event-building throughput obtained so far is shown in Figure 2.




Figure 2: DATE event-building throughput shown by AFFAIR.

### VI. MONITORING

The performance of a system as large as the ALICE DAQ system, including several processes distributed on hundreds of processors needs to be continuously and closely monitored.

The AFFAIR package gathers performance metrics from the LDCs as well as GDCs and performs the centralized handling of them. Statistics and trends are stored in a database and can be viewed using any Web browser. AFFAIR is used extensively during the Data Challenges as shown in Figure 2.

In order to guarantee the best possible quality of the output data stream, the physics data are monitored with the MOOD software package as shown in Figure 3.

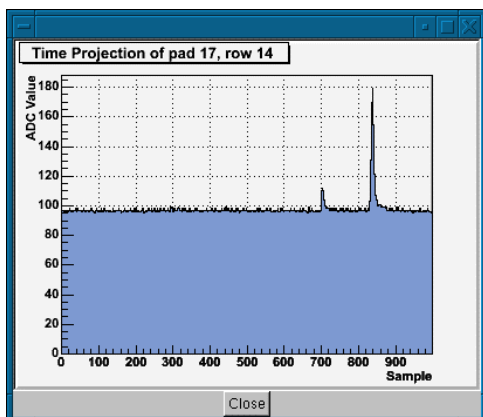


Figure 3: MOOD plot showing the time projection of 1 TPC channel

## VII. CONCLUSION

The LHC is scheduled to start colliding beams in June 2007. The ALICE experiment will be operational at the start-up of the LHC.

The DAQ system has developed all the required hardware and software elements and has tested them during the test beams of several detectors and during data challenges.

## VIII. REFERENCES

- 1) ALICE Collaboration, *ALICE – Technical Proposal for A Large Ion Collider Experiment at the CERN LHC*, CERN-LHCC-1995-71, December 1995
- 2) ALICE Collaboration, *ALICE – Physics Performance Report Vol. I*, CERN-LHCC-2003-049, December 2003.
- 3) ALICE Collaboration, *Trigger, Data Acquisition, High Level Trigger, Control System Technical Design Report*, CERN-LHCC-2003-062, ALICE TDR 010, January 2004.