

The ALICE Data-Acquisition Read-out Receiver Card

F. Carena⁽¹⁾, W. Carena⁽¹⁾, S. Chapeland⁽¹⁾, E. Dénes⁽²⁾, R. Divià⁽¹⁾, T. Kiss⁽²⁾,
J.C. Marin⁽¹⁾, K. Schossmair⁽¹⁾, C. Soós⁽¹⁾, P. Vande Vyvre⁽¹⁾, A. Vascotto⁽¹⁾
(for the ALICE collaboration)

⁽¹⁾ CERN, 1211 Geneva 23, Switzerland

⁽²⁾ KFKI-RMKI, Budapest, Hungary

Abstract

The ALICE data-acquisition system will use more than 400 optical links, called Detector Data Links (DDLs) to transfer the data from the detector electronics directly into the PC memory through a PCI adapter: 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. Using the 64-bit PCI interface IP core, the D-RORC offers more than 400 MB/s bandwidth. This document describes the hardware and firmware architecture, the co-operation with the software running on the PC, as well as the performance of the D-RORC.

I. INTRODUCTION

In the data-acquisition (DAQ) system of the ALICE experiment [1], high-speed optical links, called Detector Data Link (DDL) [2], will carry detector data from the front-end electronics to the first layer of computers. The DDL can transfer event fragments of up to 2 GB size at 200 MB/s. The interface between the I/O bus of the computers and the DDL is realized by the DAQ Read-Out Receiver Card (D-RORC). Each D-RORC includes two integrated DDL Destination Interface Units (DIU).

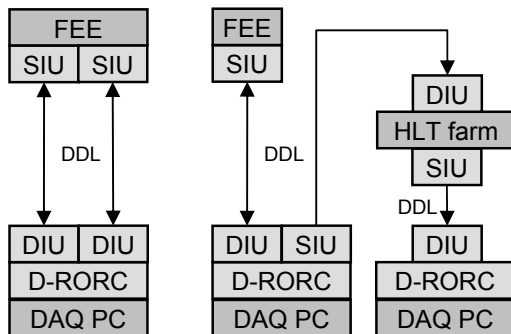


Figure 1: ALICE DAQ/HLT layout

The integrated DIU interfaces can be used in two different ways (see Figure 1). For detectors that do not require High-Level Trigger (HLT) processing, the DIUs will be connected

to the front-end electronics (FEE) via the DDL Source Interface Units (SIUs). For detectors that need the HLT system the D-RORC plays the role of a splitter. The raw data received on its first DIU port will be copied and transferred to the HLT farm using the second DIU port working as an SIU. This layout will ensure that the primary interface towards the DAQ is always provided by the DDL and the D-RORC.

II. HARDWARE

The hardware (Figure 2) is built around a programmable logic device (FPGA) from Altera APEX-E family (EP20K400). During power up, the device is loaded with the firmware stored in the EPC4 configuration memory. In-system programming (ISP) through PCI is possible using the JTAG port of the EPC4. For testing of the card, the firmware can also be uploaded directly into the programmable device via the JTAG interface of the FPGA.



Figure 2: D-RORC hardware

The card is connected to the host PC via the PCI interface, which is managed by an Intellectual Property (IP) core from PLD Applications [3]. The D-RORC complies with the 64-bit/66 MHz PCI standard. The I/O cells of the FPGA circuit can only operate in +3.3V compatible PCI or PCI-X environment.

The DIU ports are implemented by the on-board optical and electrical transceivers. The pluggable optical transceivers ease the maintenance and the system upgrade. The current version of the D-RORC is based on 2.125 Gb/s transceivers from Agilent. Using 850 nm VCSEL (Vertical Cavity Surface Emitting Laser) light source and 50/125 μm multimode

optical fibre, these optical transceivers can operate over a distance of up to 300 m.

The TLK2501 (Texas Instruments) electrical transceiver performs data conversion serial-to-parallel and parallel-to-serial. The device can operate at different frequencies providing a single-chip solution for different optical transceivers. The parallel data is encoded using the 8-bit/10-bit encoding format to provide DC balanced serial streams at the high-speed differential output. The incoming stream is decoded and the parallel data is synchronized to the extracted reference clock.

There are two input and output LVDS ports available on the board. These high-speed, general purpose serial ports can transmit and receive auxiliary data. They can be used to generate the detector busy signal for the detectors that do not produce the signal on their front-end electronics.

In order to implement extra features that are not provided by the D-RORC, one can use an appropriate extension board attached to the Common Mezzanine Card Family (CMC) [4] compatible interface of the card. The interface consisting of four connectors provide about 180 user I/O pins, as well as power pins for the extension board.

III. FIRMWARE

The firmware (Figure 3.) of the D-RORC can be divided into five building blocks: the PCI interface, the receiver and transmitter data paths, the control registers and the DDL interface.

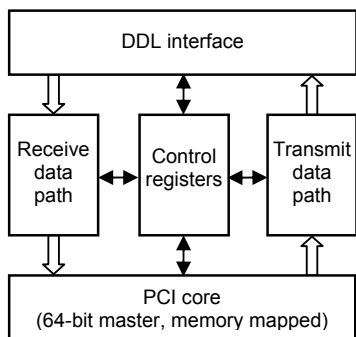


Figure 3: D-RORC firmware architecture

The main task, the autonomous data transfer across the PCI bus, is carried out by the firmware in co-operation with the software. DMA descriptors, containing a buffer address and length, are prepared by the PC software and then fed into the D-RORC. They are stored in a FIFO, called Receive Address FIFO (RAF).

Upon receiving data from the DDL, the receiver DMA manager of the D-RORC fetches one descriptor and starts the transfer without any further CPU assistance. At the end of each data block marked by the DDL-generated Data Transmission Status Word (DTSTW), the D-RORC signals the end of the transfer by writing a status directly into the Receive Report FIFO (RRF) in the PC memory.

The DMA manager can also handle data blocks which span over several allocated memory buffers. In this case, a special status is written into the RRF and the transfer is continued in the next free memory buffer referenced by the next descriptor in the RAF.

If detector data need to be copied to the HLT system, the D-RORC can be switched into data splitter mode. In this case, the raw data is duplicated in the receiver and the copy is fed into the transmitter, which transfers it to the HLT receiver using the second DDL channel. The firmware takes care of the flow-control generated by either the local (DAQ) or remote (HLT) computer ensuring that no data is lost during the copy process.

Using the backward channel of the DDL accessible through the output data path of the D-RORC, configuration data can be downloaded to the FEE. The data path is managed by the transmitter DMA of the firmware, which is similar to the receiver described above. The descriptors are stored in the Transmit Address FIFO (TAF), and the report is written to the Transmit Report FIFO (TRF).

The firmware also contains an embedded data generator, which can produce formatted data blocks of configurable length. The data blocks from the generator can be directed to the DDL, or to the receiver block using the software controlled loop-back in the D-RORC.

The 64-bit/66 MHz PCI interface is realized by the IP core from PLD Applications. The PCI core is responsible for the PCI transaction management, as well as for the DMA arbitration on the local bus. The core supports PCI mastering on four independent DMA channels. Two of these DMA channels are used per one DDL link.

The PC software can control the operation and read the status of the D-RORC using hardware registers. The registers are implemented in the firmware and accessible through the PCI. Each of them can be addressed using the appropriate base address and offset of the register. An independent control and status register set is assigned to each DDL channel.

IV. SOFTWARE

The DDL software [5] has been designed in accordance with the ALICE data-acquisition software (DATE) [6], which is being developed in CERN. It consists of two parts: a Linux kernel module (also called device driver) and an API library. The role of the device driver is to find the D-RORC on the PCI bus and to map its registers to the user memory space.

The API routines ensure that only one process can access a DDL channel at any given moment, using a suitable locking mechanism. They provide a simple interface for the execution of the DDL specific transactions, as well as for the control of the card.

As it is described in the previous section, during the data acquisition phase the software fills the RAF with descriptors and polls the RRF, which is located in the PC memory. It is an important feature of the firmware and software that neither polling via PCI, nor PCI interrupts are necessary for data

transfer. Hence, the PCI bus is not occupied by dummy operations and is left available for data transfers.

In order to access the physical memory of the PC, the DDL software and DATE use the PHYMEM package, which is part of the DATE distribution.

Several command-line executable programs have been developed using the above mentioned API library. They can perform simple tasks, as well as complex operations to receive or send data blocks. The most important features are listed hereafter:

- Identifying hardware components.
- Resetting components.
- Sending commands and reading status/registers.
- Sending data blocks.
- Receiving and checking (optional) data blocks.

V. PERFORMANCE MEASUREMENTS

The performance of the 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 motherboard controllers. 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.

In the first test series the single channel bandwidth of the D-RORC has been measured. The two curves depicted in Figure 4 show the bandwidth measured using the internal data generator of the D-RORC (Internal) and the data generator implemented in the DDL sender (External).

In both cases, 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.




Figure 4: Single channel performance

In the second test two inputs of the D-RORC were used to measure the combined bandwidth. The results in Figure 5 show that the bandwidth scales with the number of external

DDL senders used (412 MB/s maximum). Using the internal data generator, however, small bandwidth loss has been seen due to the arbitration between the DMA channels (484 MB/s maximum).

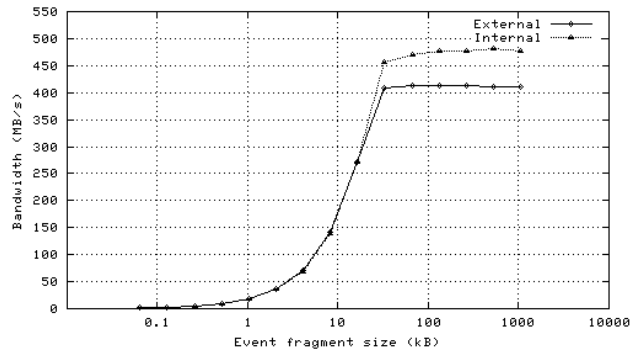


Figure 5: Dual channel performance

In order to see the behaviour and the maximum input bandwidth of the fully populated system, several tests have been carried out using six D-RORC cards in the same PC. No error was observed during these tests, which proves that several D-RORC cards can be used in one PC without any difficulty. The highest bandwidth of 1045 MB/s was achieved with four D-RORC cards installed on separate PCI segments.

Finally, the D-RORC working as data splitter has been tested. Data blocks produced by an external source were transferred to the first D-RORC, where they were copied and transferred to the second D-RORC playing the role of the HLT receiver. The content of the data blocks has been checked on each node. The same performance as in the first test has been measured, and the operation of the data splitter including the flow control management was completely transparent to the receiver nodes.

VI. APPLICATIONS

The D-RORC cards have been tested thoroughly in our lab (see previous section). They have then been produced in larger quantities to be used in real applications. Some of these interesting applications are described hereafter.

The first important challenge for the D-RORC card was the test beam of one complete Inner Read-out Chamber of the ALICE Time Projection Chamber detector [7], which took place in May 2004. In this setup two D-RORC cards played the role of the data splitters and copied the incoming data to the HLT system using their integrated DDL channels. The data blocks produced by the HLT were received by the DAQ system using a third D-RORC card. The total amount of data recorded on tapes reaches about 500 GB.

Another application of the D-RORC card is the DDL Data Generator (DDG) [8], which is foreseen to be used by the ALICE trigger system. It will be interfaced to the D-RORC using a daughter card attached to its CMC interface. The DDG will be able to generate realistic detector events and can be used as a calibrated on-line reference.

VII. SUMMARY

The D-RORC card has been developed as the interface between the DDLs and the front-end computers of the ALICE Data-Acquisition System. The card is connected to the 64-bit/66 MHz PCI interface, in order to provide the required bandwidth for the two on-board DDL channel. These DDL channels can be used in two different ways: either they are both connected to the detector, or only one is connected to the detector, while the other transfers the copy of the raw data to the HLT system.

The optimal utilization of the PCI bus during the data transfer is achieved by the co-operation between firmware and software. The DMA is fully controlled by the firmware leaving the CPU resources available for other tasks. Intensive laboratory tests, as well as field applications have proved that the D-RORC meets the ALICE requirements.

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] G. Rubin, P.V. Vyvre, C. Soós, ALICE Detector Data Link (DDL) – Interface Control Document, ALICE Internal Note, ALICE-INT-2004-018, July 2004.
- [3] PLD Applications, PCI-X and PCI Core User’s Guide, <http://www.plda.com>.
- [4] IEEE 1386-2001 and IEEE 1386.1-2001, IEEE Standard for a Common Mezzanine Card Family, <http://www.ieee.org>, 2001.
- [5] ALICE DAQ, ALICE DDL RORC Library User’s Manual, <http://cern.ch/ddl>.
- [6] ALICE DAQ, DATE V4 User’s Guide, Internal Note ALICE-2002-036.
- [7] ALICE Technical Design Report of the Time Projection Chamber, CERN/LHCC 2000-001, January 2000
- [8] S. Vergara-Limon *et al.*, DDL Data Generator, ICN/LD/04-01, June 2004.