DDL, the ALICE data transmission protocol and its evolution from 2 to 6 Gb/s

2015 JINST 10 C04008

(http://iopscience.iop.org/1748-0221/10/04/C04008)

PUBLISHED BY IOP PUBLISHING FOR SISSA MEDIALAB

RECEIVED: *November 10, 2014*
REVISED: *January 16, 2015*
ACCEPTED: *January 22, 2015*
PUBLISHED: *April 10, 2015*

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS 2014,
22–26 SEPTEMBER 2014,
AIX EN PROVENCE, FRANCE

DDL, the ALICE data transmission protocol and its evolution from 2 to 6 Gb/s

F. Carena,[a] W. Carena,[a] V. Chibante Barroso,[a] F. Costa,[a,1] S. Chapeland,[a] C. Delort,[d] E. Dénes,[b] R. Divià,[a] U. Fuchs,[a] A. Grigore,[a] C. Ionita,[a] T. Kiss,[b] G. Simonetti,[c] C. Soós,[a] A. Telesca,[a] P. Vande Vyvre[a] and B. Von Haller[a]

[a] *CERN - European Organization for Nuclear Research, CERN CH-1211 Genève 23 Switzerland*
[b] *Research Institute for Particle and Nuclear Physics, Wigner Research Center, Budapest, Hungary*
[c] *Ludwig Maximiliams University of Munich, Munich, German*
[d] *Ministere des affaires etrangeres er europeennes, France*

E-mail: filippo.costa@cern.ch

ABSTRACT: ALICE (A Large Ion Collider Experiment) is the detector system at the LHC (Large Hadron Collider) that studies the behaviour of strongly interacting matter and the quark gluon plasma. The information sent by the sub-detectors composing ALICE are read out by DATE (Data Acquisition and Test Environment), the ALICE data acquisition software, using hundreds of multi-mode optical links called DDL (Detector Data Link). To cope with the higher luminosity of the LHC, the bandwidth of the DDL links will be upgraded in 2015. This paper will describe the evolution of the DDL protocol from 2 to 6 Gbit/s.

KEYWORDS: Optical detector readout concepts; Data acquisition concepts

[1] Corresponding author.

doi:10.1088/1748-0221/10/04/C04008

Contents

2015 JINST 10 C04008

1 Introduction

ALICE [1] is the heavy-ion experiment built to study the physics of strongly interacting matter and the Quark-Gluon Plasma in nucleus-nucleus collisions at the CERN LHC (Large Hadron Collider). ALICE consists of eighteen sub-detectors and its primary targets are heavy-ion lead-lead collisions (Pb-Pb). It also has a substantial physics program with proton-proton (pp) and proton-ion (pA) collisions. The experiment has been designed to cope with the highest particle multiplicities anticipated for Pb-Pb reactions.

1.1 The ALICE Detector Data Link

The ALICE Detector Data Link (DDL) [2] has been designed to address all the requirements for data transfer between the detectors and the ALICE Data Acquisition system [3]. The first version of the protocol, called DDL1, is a 2.125 Gb/s full duplex, multipurpose optical fibre link. It allows the bi-directional transmission of data between the front end electronics and the data acquisition system. During data taking the information flows from the front end electronics (FEE) to the DAQ, while at each start of a new run it is possible to send configuration control messages to the detector using the opposite communication channel. The DDL1 consists of three main components (figure 1):

- **Source Interface Unit (SIU)**: a mezzanine card connected to the Front End Electronics (FEE).
- **Destination Interface Unit (DIU)**: a VHDL block implemented in the FPGA of the Readout Receiver Card (RORC), a PCIX or PCIe custom card, installed in the DAQ farm that receives data from the detectors and stores information in the memory of the hosting PC.

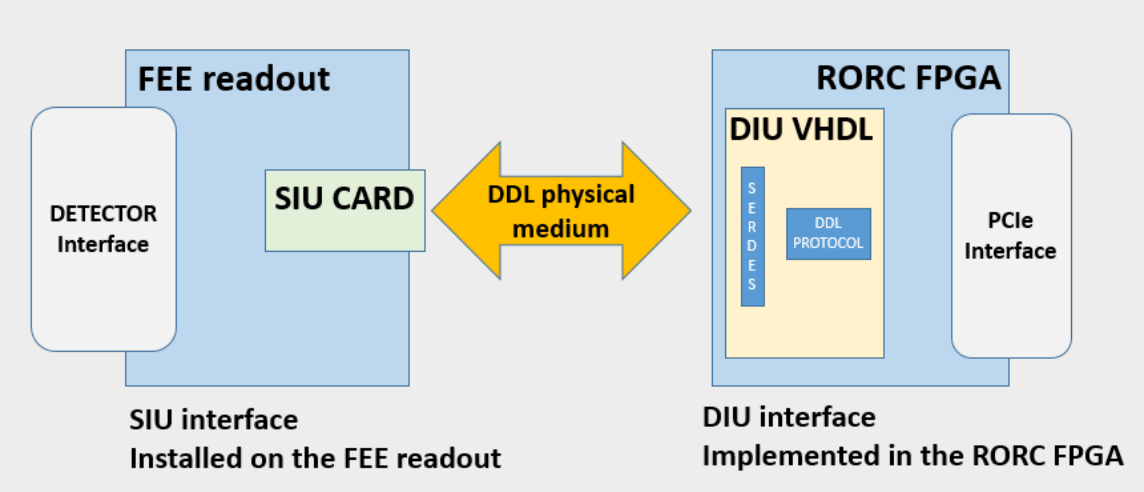

Figure 1. DDL1 components in the data taking configuration, FEE connected to the DAQ PC.

- **The physical medium**: a multi-mode OM2 optical fibre. The detectors are connected to the DAQ farm using different length of fibres up to 200 m.

The DDL1 is a point to point protocol, so each SIU is connected directly to one channel of the RORC through the optical fibre.

2 DDL firmware

The DDL provides a common data transfer protocol to the detectors that are connected to ALICE DAQ system. It encapsulates the data collected from the FEE in a DDL frame and serializes it through the optical fibre. Its main components are:

- **SERDES**: the serializer-deserializer transmits serialized data over high speed, differential pairs and it recovers and aligns the data on the receiving side.
- **TLK2501 RX/TX state machines**: the SIU and old release of the RORC cards are using a transceiver from Texas Instruments, TLK2501. This chip has state machines to recover the link and keep the data aligned. In the latest release of the RORC the transceiver of Texas Instruments has been replaced by the SERDES available in the FPGA. The state machines simulate the chip behaviour so the code is backward compatible with other cards that are equipped with the TLK2501.
- **CRC**: a cyclic redundancy check verifies the data to detect possible errors during the data transmission.
- **Framing**: the DDL works using 32 bit words, however the SERDES provides a data bus interface at 16 bit. This components receives 32 bit word and splits it into two 16 bit words.
- **DDL protocol and Command interface**: it encapsulates the data coming from the detector in DDL data blocks. When commands are sent from the RORC to the FEE it raises the control signals so the FEE can react to incoming commands or configuration data.

2015 JINST 10 C04008

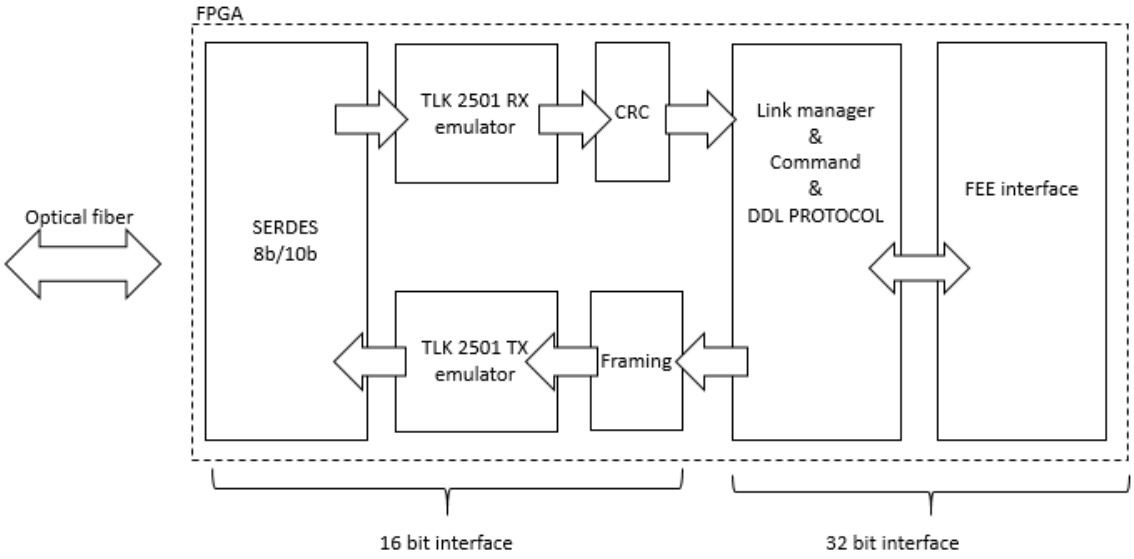

Figure 2. Main DDL firmware components and their communication paths.

- **FEE interface**: it is a 32 bit bi-directional bus used to interface the user logic to the DDL. Using this interface it is possible to send data and receive configuration instruction.

The connections between the different DDL blocks are showed in the figure 2.

3 DDL2: motivation and requirements

ALICE foresees a consolidation in 2015 to cope with the increased luminosity of the accelerator. Due to higher data throughput generated by two sub-detectors, the Time Projection Chamber (TPC) and the Transition Radiation Detector (TRD) the DDL protocol has been maintained but at a faster speed of up to 6 Gb/s.

An exact copy of the data received by the DAQ is sent to the ALICE High Level Trigger (HLT) to perform online tracking and data reduction. In the previous period run HLT was using PCI-X cards, called HRORC (HLT Readout Receiver Card). PCI-X is obsolete by now, so HLT upgraded their boards using the same hardware installed in the ALICE DAQ system upgrading as well the DDL protocol.

For the DDL2 the SIU card has been replaced by VHDL core to be implemented in the FPGA firmware of the readout electronics (figure 3).

Using a VHDL core to interface the detectors to the DAQ system it can be adapted to different cards, benefitting from the resources available on the newest FPGA, high speed clock and high speed serial links.

4 DDL2 performance

During Run2, ALICE will use different speeds of the DDL due to different detector electronics. The experiment will collect data using both generations of the data transmission protocol at the same time. The DDL1 is running at 2.125 Gb/s, while the DDL2 will run at higher speed. The

2015 JINST 10 C04008

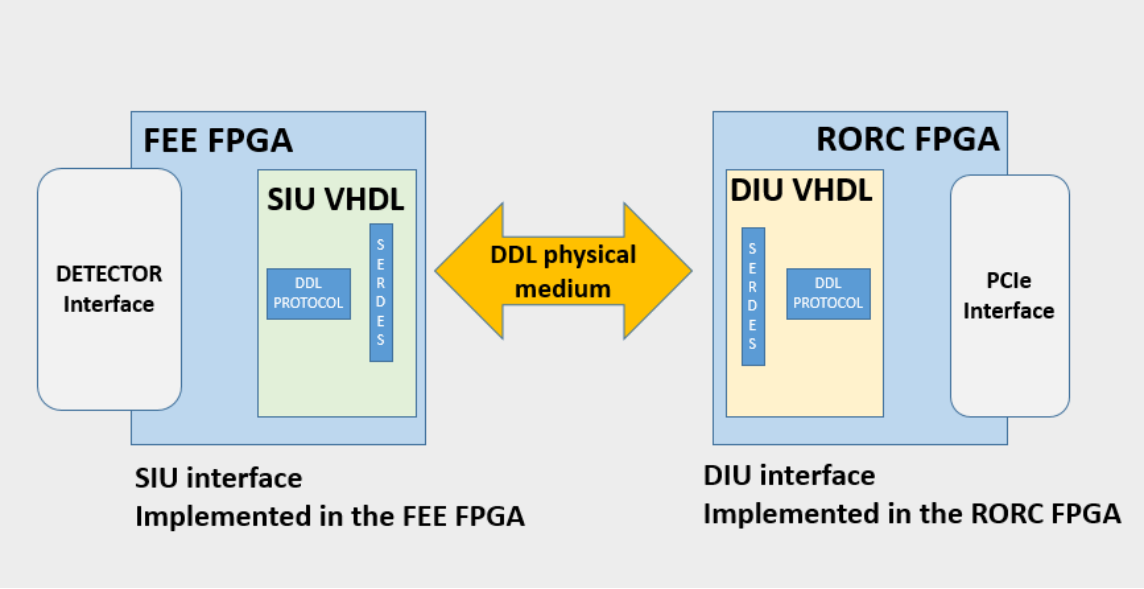

Figure 3. DDL2 components in the data taking configuration, FEE connected to the DAQ PC. The SIU card has been replaced by a VHDL component.

TPC is aiming at 4.25 Gb/s, while TRD will use 4 Gb/s. The different link speeds are defined by the detector hardware where the DDL protocol is implemented.

The ALICE DAQ system and its readout receiver card have been thoroughly tested with the FEE using DDL2 at readout speeds, from 2.125 Gb/s up to 5.3125 Gb/s.

5 Implementation in different FPGAs

The biggest challenge in this project has been the implementation of the code in different FPGA families. TPC, TRD, HLT and DAQ are using electronics equipped with different FPGAs to collect data, so the DDL2 protocol has been implemented in the following devices:

- TPC will use MICROSEMI SmartFusion2.
- TRD will use XILINX Virtex4.
- The DAQ and HLT system will use XILINX Virtex6.
- The other detectors will use the DDL1, however some of them are planning to upgrade their electronics to implement the DDL2 for the RUN3 in 2018. For this reason the DDL code has been ported also for XILINX Kintex7 and Virtex7.

Although the DDL protocol is mainly VHDL custom code, the implementation in different families requires particular attention to the timing constraints and the instantiation of specific Intellectual Property blocks, or IP cores. The user, using these "FPGA functions", can access specific features of the devices, as the SERDES (SERializer DESerializer) and the FIFO (First-in, First-out) both used in the code. IP cores of different FPGA differ in user interface and sometimes in behavior. These differences increased the development time and introduced an additional effort in maintaining the same code for different devices.

2015 JINST 10 C04008

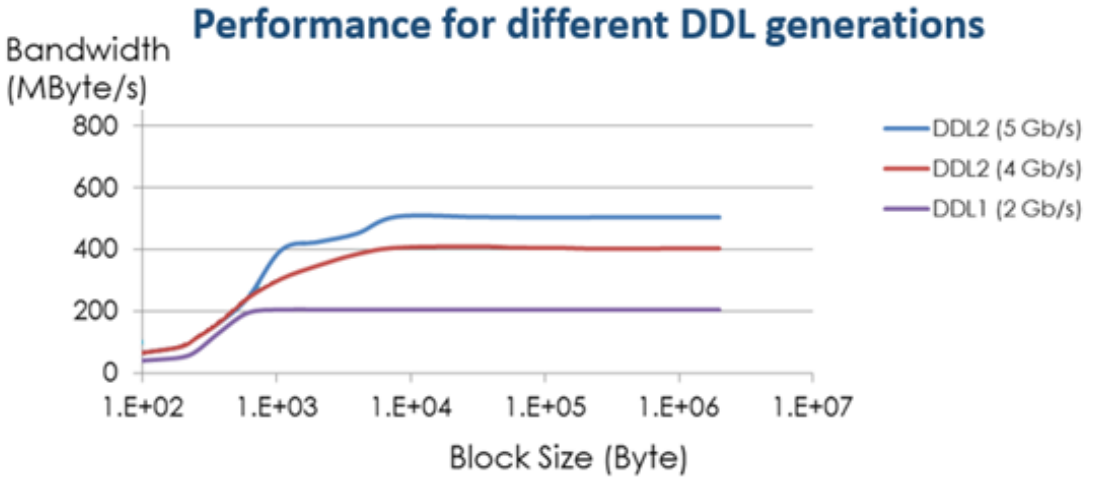

Figure 4. DDL throughput, using different link speed and varying the event size.

The critical part in the DDL protocol is the implementation of the SERDES. XILINX delivers strong and flexible solutions in their FPGA to access the high speed serial links, while the latest product of MICROSEMI, Smart Fusion2, is not as robust as its counterpart and it requires special attention during the implementation. Delay lines and synchronization FIFOs are necessary and they have to be inserted between the user logic and the SERDES, for both receiving and transmitting channels.

5.1 DDL2 verification

The different implementations of the DDL2 protocol have been tested using the ALICE DATE software framework that provides software tools to check the DDL status. These programs are used to verify the different components of the DDL (SIU, optical fibre and DIU) in case there is a problem during data taking.

The SIU-loopback program runs on the machine that hosts the RORC. It generates different pattern of data (32 bit word), incremental or alternated and sends it to the SIU, using the DDL protocol. From the SIU data is sent back to the RORC and the program checks if there are bit flip or other errors in the data transmission. For each FPGA implementation the DDL has been stressed collecting more than 80 TB of data without errors.

In addition to that a data generator has been implemented in the SIU VHDL component. In this configuration it has been possible to collect data using the DATE software. The data generator sends data to the RORC using the DDL protocol. The DATE software collects it and perform sanity check to verify that the information are well formatted.

For each implementation different speeds of the DDL has been successfully verified:

- XILINX Virtex4,Virtex6, Kintex7, Virtex7: 2.125 Gb/s, 3.125 Gb/s, 4.25 Gb/s 5.3125 Gb/s
- MICROSEMI SmartFusion2: 2.125 Gb/s. More development is needed to run at higher speed.

2015 JINST 10 C04008

6 Conclusions

The DDL2 protocol has been successfully ported and tested in different FPGA families, XILINX and MICROSEMI. Different speeds of the link have been tested, from 2.125 up to 5.3125 Gb/s in XILINX devices. One of the main challenges in the development has been the implementation of the DDL protocol in MICROSEMI FPGA, SmartFusion2. Different coding technique, following MICROSEMI reference design, had to be applied to use properly the new technology. The protocol has been tested in a stable way at 2.125 Gb/s. The targeted speed for the TPC, 4.25 Gb/s, is still under development. During Run2 the number of links that will use the DDL2 will be a total of 240, connecting to the DAQ system TPC, TRD and HLT.

References

[1] ALICE collaboration, *The ALICE experiment at the CERN LHC*, 2008 *JINST* **3** S08002.

[2] C. Soós et al., *The ALICE detector data link*, *5th Conference on Electronics for LHC Experiments*, Snowmass, CO, U.S.A., 20–24 September 1999.

[3] ALICE collaboration, *The ALICE data acquisition system*, *Nucl. Instrum. Meth.* **A 741** (2014) 130.

[4] H. Engel et al., *The C-RORC PCIe card and its application in the ALICE and ATLAS experiments*, proceedings of the *Topical Workshop on Electronics for Particle Physics 2014*, Aix en Provence, France, 22–26 September 2014.

2015 JINST 10 C04008