

2 The C-RORC PCIe Card and its Application in the 3 ALICE and ATLAS Experiments

**A. Borga^a, F. Costa^b, G. J. Crone^c, H. Engel^{d*}, D. Eschweiler^e, D. Francis^b, B. Green^f,
M. Joos^b, U. Kebschull^d, T. Kiss^g, A. Kugel^h, J. G. Panduro Vazquez^f, C. Soos^b, P.
Teixeira-Dias^f, L. Tremblet^b, P. Vande Vyvre^b, W. Vandelli^b, J. C. Vermeulen^{a*}, P.
Werner^b, and F. J. Wickensⁱ for the ALICE and ATLAS Collaborations**

^a*Nikhef National Institute for Subatomic Physics and University of Amsterdam, Amsterdam,
Netherlands*

^b*CERN, Geneva, Switzerland*

^c*Department of Physics and Astronomy, University College London, London, United Kingdom*

^d*Institut für Informatik, Johann Wolfgang Goethe-Universität Frankfurt, Frankfurt, Germany*

^e*Frankfurt Institute for Advanced Studies, Johann Wolfgang Goethe-Universität Frankfurt,
Frankfurt, Germany*

^f*Department of Physics, Royal Holloway University of London, Surrey, United Kingdom*

^g*Wigner Research Centre for Physics, Hungarian Academy of Sciences, Budapest, Hungary*

^h*ZITI Institut für technische Informatik, Ruprecht-Karls-Universität Heidelberg, Mannheim,
Germany*

ⁱ*Particle Physics Department, Rutherford Appleton Laboratory, Didcot, United Kingdom*

4 *E-mail: henggel@cern.ch, j.vermeulen@nikhef.nl*

5 **ABSTRACT:** The ALICE and ATLAS DAQ systems read out detector data via point-to-point serial
links into custom hardware modules, the ALICE RORC and ATLAS ROBIN. To meet the increase
in operational requirements both experiments are replacing their respective modules with a new
common module, the C-RORC. This card, developed by ALICE, implements a PCIe Gen 2 x8
interface and interfaces to twelve optical links via three QSFP transceivers. This paper presents the
design of the C-RORC, its performance and its application in the ALICE and ATLAS experiments.

6 **KEYWORDS:** Data acquisition circuits; Data acquisition concepts; Digital electronic circuits;
7 Online farms and online filtering; Optical detector readout concepts.

*Corresponding authors.

9 **Contents**

10	1. Introduction	1
11	1.1 ALICE Online Architecture in Run 1 and Run 2	1
12	1.2 ATLAS: Upgrade of the ReadOut System	2
13	2. The Common Read-Out Receiver Card (C-RORC)	4
14	3. Applications of the C-RORC in ALICE and ATLAS	5
15	3.1 ALICE Data Acquisition	5
16	3.2 ALICE High-Level Trigger	6
17	3.3 ATLAS Readout System	7
18	4. Conclusion & Outlook	8

20 **1. Introduction**21 **1.1 ALICE Online Architecture in Run 1 and Run 2**

22 ALICE [1] is the heavy-ion experiment at the CERN LHC dedicated to the study of the physics of
 23 strongly interacting matter. It has been designed to cope with the high particle densities produced
 24 in central Pb-Pb collisions. The data captured from all 18 subdetectors are read out by the ALICE
 25 Data Acquisition (DAQ) system via around 500 serial optical links called Detector Data Links
 26 (DDLs) [2]. The data sent via DDLs from the cavern to the counting rooms is received in custom
 27 FPGA based DAQ Read-Out Receiver Cards (D-RORCs). These boards are installed in servers
 28 acting as Local Data Concentrators (LDCs). For each DDL an exact copy of the incoming data is
 29 forwarded within the D-RORC FPGA to another DDL towards the High-Level Trigger (HLT). A
 30 simplified overview of the read-out architecture is shown in figure 1.

31 The HLT is the first system in ALICE where data from all detectors is combined and recon-
 32 structed. This compute cluster is comparable in size to the DAQ cluster and additionally contains
 33 Graphics Processing Units (GPUs). The interface nodes are equipped with custom FPGA based
 34 HLT Read-Out Receiver Cards (H-RORCs), receiving the detector data via DDLs and performing
 35 first reconstruction steps. In addition to software based data processing on the nodes, the computing
 36 power of the HLT could significantly be enhanced by implementing pre-processing algorithms in
 37 the H-RORC firmware and offloading computations to GPUs [3]. Output nodes pass the processed
 38 data back to the DAQ system via H-RORCs and DDLs.

39 The HLT decisions for each event are readout by the DAQ, using the DDLs as for any other
 40 detector. The sub-events from the detector LDCs and the HLT decision are then sent over the Event
 41 Building Network for global processing and finally into long term storage.

42 The Read-Out Receiver Cards for DAQ and HLT have similar requirements, however they have
 43 been developed and maintained as independent projects. The H-RORC contains a Xilinx Virtex-4
 44 FPGA and connects to DDLs via pluggable add-on boards hosting the optical links. The interface
 45 to the host machine is implemented with PCI-X. The D-RORCs have been used in two different
 46 revisions: one with PCI-X and one with PCIe interfacing to the host machine. These boards use
 47 Altera APEX or Stratix II FPGAs and have two optical interfaces per board. During Run 1 around
 48 400 D-RORCs and around 240 H-RORCs were used in the DAQ and HLT systems.

49 The read-out architecture described will remain the same for Run 2. LHC luminosities after
 50 Long Shutdown 1 are expected to be in the range of $1 - 4 \times 10^{27} \text{ cm}^{-2}\text{s}^{-1}$ with a center-of-mass
 51 energy of 5.1 TeV for Pb-Pb collisions. The expected data rates require that the read-out system as
 52 deployed during Run 1 is upgraded. The Time Projection Chamber (TPC) is replacing its Readout
 53 Control Unit with a redesign for higher detector bandwidth and increased output link rate (RCU2).
 54 The Transition Radiation Detector (TRD) is implementing a higher read-out link rate with the
 55 existing Global Tracking Unit (GTU) hardware. Therefore the original version of the DDL (also
 56 referred to as DDL1) has been upgraded to the DDL2 [4], which supports higher link rates. The
 57 increasing data rates and read-out changes also affect the systems of DAQ and HLT and in particular
 58 the Read-Out Receiver Cards.

59 Both types of RORCs used during Run 1 are limited in their optical read-out capabilities by
 60 the DDL1 link rates. Additionally, the PCI-X host interface is obsolete and increasingly rare in
 61 recent server PCs. These facts require a replacement of the Run 1 Read-Out Receiver Cards.

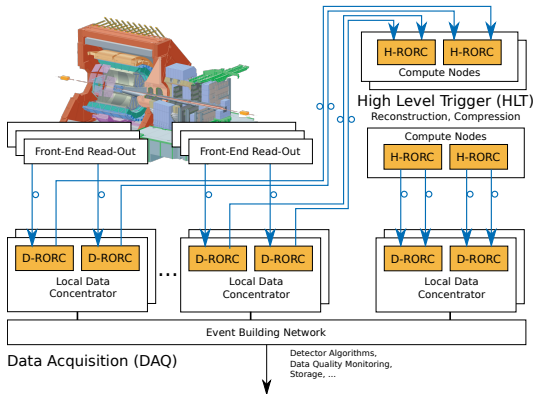


Figure 1: The ALICE online architecture with focus on the RORCs in DAQ and HLT.

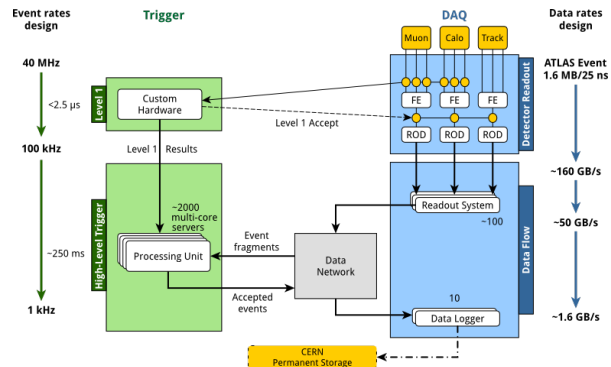


Figure 2: The ATLAS TDAQ system in Run 2.

62 1.2 ATLAS: Upgrade of the ReadOut System

63 The focus of the ATLAS experiment [5] at the LHC is the study of high-energy proton-proton collisions
 64 at high luminosities. The experiment makes use of a trigger system consisting of three levels
 65 to reduce the event rate to a manageable level. The first level consists of dedicated hardware. Data
 66 from events accepted by this level are transferred from the front-end electronics to the ReadOut
 67 Drivers (RODs). These are sub-detector specific modules, located in an underground service area
 68 adjacent to the cavern in which the experiment is installed. An important task of the RODs is to
 69 build event fragments and output these to the ReadOut System (ROS). For each first-level trig-

70 ger accept each ROD outputs one event fragment. Each fragment contains an identifier, the L1Id,
71 which is, apart from resets, monotonically increasing for consecutive fragments. A supervisor re-
72 ceives and forwards the same L1Id and additional information to a core of a higher-level trigger
73 processing node selected by the supervisor for handling the event. For the second-level trigger the
74 nodes request only part of the event data from the ROS, initially using the additional information
75 provided by the first-level trigger. The L1Id is forwarded as part of each request for data associated
76 with that L1Id via the Ethernet network connecting the nodes and the ROS. The ROS responds
77 by sending the data requested. For Run 1 the second level of triggering was implemented using
78 a dedicated set of server PCs. Upon acceptance by this level, full event building was performed
79 by another dedicated set of server PCs known as the Event Builder¹, which like the second-level
80 trigger processors requested the event data from the ROS, but instead of a fraction all data were
81 requested. Full events were then built and forwarded to the highest trigger level, known as the
82 Event Filter and running on another dedicated set of server PCs. For Run 2 the same approach
83 will be used, but all processing of an event, i.e. second-level processing, event building and Event
84 Filter processing will be done on the same processing node. As in Run 1 event fragments will
85 be discarded in the ROS upon delete requests that are broadcast to the ROS after a second-level
86 trigger reject or after successful building of the full event (or of a partial event in case of certain
87 types of events, in particular calibration events). A diagram of the structure of the Trigger and DAQ
88 (TDAQ) system for Run 2, with data volumes and trigger rates indicated, is presented in figure 2.

89 The event fragments are transferred from the RODs to the ROS via dedicated point-to-point
90 links in the form of optical fibers, using the S-link protocol [6] and running at either 160 MB/s or
91 200 MB/s maximum throughput. For Run 1 about 1600 of these links were deployed, this number
92 increases to about 1800 for Run 2. The ROS as deployed during Run 1 was built from about 150
93 server PCs, with typically 4 ROBINS [7] installed in a PC. ROBINS are PCI plug-in cards with three
94 inputs for the point-to-point links via which the RODs output their data. Each PC was also equipped
95 with a PCIe plugin card connecting via two ports to the data collection network, implemented with
96 1 Gb Ethernet technology. Each ROBIN contained a 64 MB paged memory buffer for each of the
97 three inputs, a Xilinx Virtex-II FPGA, a PCI interface chip and a PowerPC processor keeping track,
98 together with the FPGA, of the association between page number and L1Id of a fragment stored
99 in a buffer memory. Requests were forwarded by the PC to a ROBIN via its 64-bit 66 MHz PCI
100 interface, data requested was written to the memory of the host via DMA.

101 The increase of the number of ROD to ROS links for Run 2 made a reduction of the rack space
102 used per ROD to ROS link desirable. Furthermore 64-bit PCI technology is becoming obsolete,
103 motherboards with the four slots of the ROS PCs used in Run 1 are not readily available for the
104 current generation of CPUs (Ivy Bridge or Haswell architecture), a PCIe solution was therefore
105 required. In addition higher level trigger conditions adapted to the higher luminosity and collision
106 energies of Run 2 and the higher maximum average level-1 accept rate of 100 kHz (instead of about
107 70 kHz for Run 1) will result in more data being requested from the ROS. In view of this it was
108 decided to replace the ROS used in Run 1 by a more compact ROS with PCIe based ROBINS and
109 capable of handling readout of at least 50% of the data received via the ROD to ROS links. With the
110 CPU power available in modern server PCs it was considered feasible to move the tasks of the on-

¹the PCs are also referred to as SFIs (SubFarm Inputs).

111 board processor of the ROBIN to the CPU of the ROS PC, simplifying the design of the ROBIN and
 112 also simplifying support, as software and the development environment for the on-board processor
 113 no longer have to be maintained. This new version of the ROBIN is known as the RobinNP, "NP"
 114 refers to "No Processor". The custom board developed by the ALICE collaboration, the C-RORC,
 115 described in the next section, provides all functionality required for the RobinNP, as is discussed in
 116 Section 3.3.

117 **2. The Common Read-Out Receiver Card (C-RORC)**




Figure 3: Photo of the C-RORC board with the major components and features annotated.

118 The lack of suitable commercial platforms to replace the Run 1 Read-Out Receiver Cards
 119 deployed in ALICE led to the development of a custom board. Even though the development was
 120 driven by ALICE requirements, the target platform was kept as generic as possible. A photo of the
 121 final board with the major components annotated is shown in figure 3. The board is a full-width,
 122 full-height PCIe card according to the PCIe specification. The height of the components is kept
 123 within the specification to allow installation of boards into adjacent PCIe slots. The boards are
 124 powered from 6-pin GPU power cables.

125 The central component on the board is a Xilinx Virtex-6 FPGA. This FPGA already comes
 126 with a PCIe hard block for up to eight lane PCIe generation 2 (8x 5.0 Gbps). A PCIe throughput
 127 close to the theoretical limit has been observed. The board interfaces to 12 serial full duplex optical
 128 links via three QSFP modules, with each QSFP module connecting to four optical links. Break-
 129 out fibers are available to connect to the existing fiber installations. The serial links are directly
 130 connected to the transceivers of the FPGA (GTX), which limits the maximum serial link rate to
 131 around 6.6 Gbps. An on-board configurable reference clock oscillator makes it possible to use
 132 almost any link rate within the supported range. On-board DDR3 memory can be installed in two
 133 SO-DIMM sockets. The required memory controllers can be implemented in the FPGA and allow
 134 operation of single ranked modules up to 1066 Mbps and dual ranked modules up to 606 Mbps.
 135 Both interfaces have been tested with a variety of different modules up to 2x 8 GB total capacity.
 136 FPGA configuration files can be stored in on-board synchronous flash memories for fast auto-
 137 configuration of the board upon power-on. Additionally, there is enough memory to store multiple

138 FPGA configurations. A configuration microcontroller can be accessed by the host machine even
139 if the PCIe link is down. This allows implementation of a safe firmware upgrade procedure by
140 always keeping a known-to-be-working configuration in the flash memory.

141 The large scale production of the boards was organized as a common effort between ALICE
142 and ATLAS. Extensive hardware tests have already been conducted by the contractor. More appli-
143 cation specific tests have been done by ALICE and ATLAS at CERN. At the time of this writing
144 359 boards have successfully been produced, tested and delivered to CERN, of which most have
145 been installed in the ALICE DAQ and HLT and ATLAS DAQ systems.

146 **3. Applications of the C-RORC in ALICE and ATLAS**

147 With the C-RORC there is now a common hardware platform for three applications in two LHC
148 experiments: ALICE Data Acquisition, ALICE High-Level Trigger and ATLAS TDAQ Read-
149 Out System. Even though the platform is the same, each application has to interface to existing
150 application-specific hardware and software infrastructure. For this reason firmware for each of the
151 three applications is developed independently. Nevertheless, common building blocks are reused
152 and approaches are shared. The following sections describe the applications in more detail.

153 **3.1 ALICE Data Acquisition**

154 The ALICE DAQ system handles the data flow from the detector to permanent data storage in the
155 CERN computing center and is responsible for uploading configuration data to the detectors [8].
156 The interface to the DDLs in the DAQ Read-Out Receiver Card firmware is therefore providing
157 two operating modes: *data taking* and *detector configuration*.

158 In *data taking* mode the receiving channel of each read-out link is used to transfer event data
159 from the detector electronics to the DAQ farm. The transmitting channel is used for flow control. In
160 *detector configuration* mode the transmitting channel is used to send configuration data to the front
161 end electronics. The receiving channel is used for acknowledgments from the front end electronics.

162 The ALICE DAQ Run 2 setup is a mixed installation consisting of C-RORCs for all TPC, TRD
163 and HLT-to-DAQ links. The previous D-RORC boards are still in use with the remaining detectors.
164 The C-RORCs use six optical links to receive detector data and the other six links to send a copy
165 of the data to the HLT. The copy process between the links is directly implemented in the RORC
166 firmware. The DDL protocol has been ported to the higher DDL2 rates to support the detectors
167 that upgrade their read-out for Run 2. The firmware interface to the host server via PCIe is based
168 on a PLDA DMA engine [9] for six data channels. This is the same interface as already used for
169 the D-RORC boards, which allows a common device driver and software interface for both types
170 of boards.

171 The host memory for DMA operations is managed with the *physmem* driver and divided into
172 page-like segments with known physical start addresses and lengths. These buffer descriptors are
173 pushed into a FIFO in the RORC firmware and then used as start addresses for DMA transfers.
174 For each descriptor used for a DMA transfer, the RORC writes an entry into a second DMA buffer
175 in the host memory to inform the software of new data. The DAQ farm for Run 2 will consist
176 of a cluster of around 130 servers with 10 Gb Ethernet interconnect, in which 59 C-RORCs are
177 installed.

178 3.2 ALICE High-Level Trigger

179 In the ALICE HLT one C-RORC replaces three to six of the previous H-RORC boards, thus allow-
 180 ing a much denser integration of the optical links into the cluster. Up to 12 links per board are used
 181 to receive data from the DAQ system. The optical link protocol is identical to that used for ALICE
 182 DAQ: DDL at different link rates depending on the detector. For Run 2, 74 C-RORCs have been
 183 installed into 2U dual socket IvyBridge servers together with GPUs and 56 Gb InfiniBand inter-
 184 connect. The overall HLT for Run 2 consists of 180 compute nodes, each with two 12-core CPUs
 185 and a GPU, and some infrastructure machines. A schematic picture of the node configuration and
 186 an overview of the dataflow inside the HLT C-RORC firmware is shown in figure 4.




Figure 4: C-RORC Installation in the ALICE HLT for Run 2 and schematic drawing of the dataflow in the firmware.

187 The existing HLT data transport framework assumes one process per DDL. With 12 links per
 188 board this requires DMA engine firmware that is able to operate 12 DMA channels independently.
 189 This was not possible with any available commercial PCIe DMA core for the given FPGA archi-
 190 tecture, so a custom DMA engine was developed. This DMA engine handles scatter-gather DMA
 191 descriptor lists provided by the host system and thus allows the standard Linux memory subsystem
 192 to be used for buffer allocation and mapping. The possibly scattered physical memory fragments
 193 are mapped into a contiguous virtual memory region by a user space device driver library. The
 194 DMA buffers are used as ring buffers, with each DMA channel using two: *EventBuffer* and *Re-*
 195 *portBuffer*. Detector data from the optical links are directly written into the *EventBuffer*. Once an
 196 event is fully transferred, an entry is written into the *ReportBuffer* containing the offset and length of
 197 the event in the *EventBuffer*. All hardware access is performed from user space using the Portable
 198 Driver Architecture (PDA, [10]) library together with a user space device driver. The PDA allows
 199 memory-mapping the DMA buffer twice to consecutive virtual memory addresses, which allows a
 200 transparent handling of the wrap-around effects of the ring buffers.

201 An essential part of the HLT firmware is the FastClusterFinder online pre-processing algo-
 202 rithm [11], which can be integrated into the dataflow to extract features of the raw TPC data while
 203 passing through the RORC. The FastClusterFinder can handle the full bandwidth of the DDL link
 204 and induces only a marginal additional readout latency of a few microseconds while saving a sig-
 205 nificant amount of CPU resources compared to the same processing steps in software. This was
 206 developed for DDL1 speed and and has now been tuned to support the higher optical link rates of
 207 the DDL2 protocol.

208 The ALICE HLT uses the DDR3 memory on the C-RORC only to replay previously recorded
 209 detector data into the system. Six DMA channels share one DDR3 SO-DIMM module. This allows
 210 the full HLT chain to be tested with real detector data without requiring DAQ or detector resources.
 211 The on-board DDR3 memory is not used during physics runs.

212 **3.3 ATLAS Readout System**

213 As mentioned in Section 1.2, the C-RORC provides all the functionality required for the RobinNP:
 214 12 ROD to ROS links can be connected to a single board, four times as many as to the ROBIN, and
 215 the PCIe interface has a throughput of at maximum about 15 times that of the PCI interface of a
 216 ROBIN. The resource requirements for implementation of RobinNP functionality are fully satisfied
 217 by the FPGA, which is also capable of handling the requirements with respect to data throughput.
 218 Furthermore slots for on-board memory allow up to 16 GB of buffer memory, while a ROBIN has
 219 192 MByte of buffer memory. Higher speeds than the current 160 or 200 MB/s are also possible for
 220 the input links. The C-RORC therefore made it possible to build a new compact ROS that makes
 221 use of 98 2U high server PCs, with two C-RORCs installed in most server PCs, therefore 24 ROD
 222 to ROS links can be connected to a single ROS PC (to be compared to 12 links connected to a 4U
 223 high ROS PC for Run 1). Because of the factor of two increase of the number of links connected
 224 to a single PC, and because of the higher request fractions the networking infrastructure also had
 225 to be upgraded: instead of having two 1 Gb Ethernet links a ROS PC is now connected with four
 226 10 Gb Ethernet links to the data collection network.

227 A schematic diagram of the RobinNP
 228 firmware and its interactions with the host PC is
 229 presented in figure 5. The firmware consists of
 230 two identical parts, referred to as ROBGROUPs,
 231 each connecting to six ROD to ROS links (labeled
 232 as ROL (ReadOut Link) in the diagram) and a
 233 common part implementing an eight lane Gen 1
 234 PCIe interface and the DMA engine. The latter is
 235 the engine available from PLDA [9]. Each ROBG-
 236 Group has one shared buffer memory, consisting
 237 of a 4 GByte DDR3 SO-DIMM module, which is
 238 logically subdivided in six partitions, one for each
 239 ROD to ROS link. Pages in the buffer memories
 240 are managed by multi-threaded software running
 241 on the ROS PC, a typical page size is 2 kByte. For
 242 each memory partition the PC provides informa-
 243 tion on free memory pages, via FIFOs implemented in firmware, to each of the 12 input handlers.
 244 Incoming fragments are stored in free pages. For every page used, information on the page num-
 245 ber, L1Id and length of the fragment stored is entered in the Used Page FIFO of the input handler
 246 that handled the fragment. Per ROBGROUP the information from each of these FIFOs flows into the
 247 "Combined Used Page FIFO", and is subsequently transferred to the memory of the PC by means of
 248 DMA by the "FIFO duplicator". The information is used by a dedicated thread for "indexing", i.e.
 249 information is stored on the relation between L1Id and the page (or pages if the fragment is larger
 250 than the page size) in which a fragment is stored as well as on the length of the fragment. Data
 251 requests received via the network cause a look-up of this information and forwarding of requests
 252 for reading data from the pages concerned. These data are then read by the FPGA from the DDR3
 253 memory and passed to the DMA engine for transfer to the memory of the PC. For each ROBGROUP

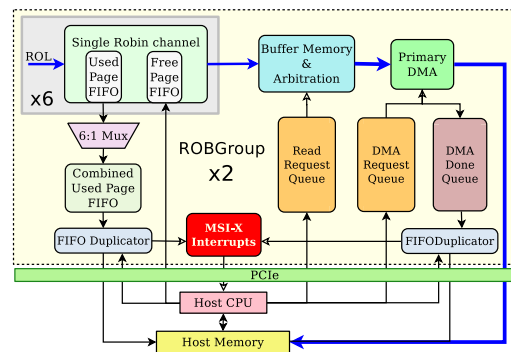


Figure 5: RobinNP firmware organization and flow of data from host CPU to the firmware (by means of programmed I/O) and from the firmware to the host memory (by means of DMA).

254 a second FIFO duplicator transfers information concerning completed DMA transfers from a FIFO
255 to the memory of the PC. This information is used for collecting the data requested, which is output
256 via the network. Clear requests are also sent to the ROS via the network. These requests result in
257 the identifiers of the pages concerned being recycled onto a free page stack and eventually back
258 onto the Free Page FIFOs, thus allowing the data in memory to be overwritten. The communication
259 between RobinNP and the PC is interrupt driven: the indexer thread is woken upon storage of new
260 event data and the thread used for data collection is woken upon the completion of DMA transfers.
261 Interrupt coalescence has been implemented in an innovative way: an interrupt only occurs if the
262 buffer to which data is transferred from the FIFO with which the interrupt is associated is empty
263 upon arrival of new data. During normal operation the PC does not need to read any data via PCIe
264 from FIFOs in the FPGA, as all data is written under DMA control to the memory of the PC. In
265 this way optimum utilisation of the available PCIe bandwidth is achieved.

266 At the time of writing the installation of the new ROS has just been completed. Each of the
267 98 installed ROS PCs has a single CPU motherboard equipped with an Intel E5-1650v2 six-core
268 3.5 GHz CPU and 16 GB of memory. The CPU connects directly to 40 PCIe Gen3 lanes, 32
269 lanes are connected to a riser card with four 8 lane connectors. Two connectors are used for two
270 C-RORCs, the other two for two dual-port 10 Gb Ethernet NICs with optical transceivers. The
271 operating system of the PCs is Linux (SLC6). This configuration has been shown to be able to
272 satisfy the 50% readout fraction requirement at 100 kHz first-level trigger accept rate with two
273 C-RORCs with RobinNP firmware installed [12].

274 **4. Conclusion & Outlook**

275 This paper presents the C-RORC, a PCIe-based FPGA read-out board, which will be used in two
276 of the major LHC experiments for three applications in data taking for Run 2 . All parties strongly
277 profited from the collaboration. The significant increase in production volume with respect to
278 to deployment restricted to ALICE led to cost savings per board for both experiments. Usage
279 experience, implementation methods and partly even source code could be shared between the
280 developers of the different applications reducing the overall development time. All boards required
281 for Run 2 have been successfully produced, tested, delivered and installed in the ALICE DAQ and
282 HLT systems and in the ATLAS DAQ system.

283 **Acknowledgments**

284 Supported by the German Federal Ministry of Education and Research BMBF 05P12RFCAA.

285 **References**

- 286 [1] ALICE Collaboration, *The ALICE Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008)
287 S08002.
- 288 [2] ALICE Collaboration, *The Technical Design Report of the Trigger, Data-Acquisition, High Level*
289 *Trigger, and Control System*, tech. rep., CERN-LHCC-2003-062.
- 290 [3] S. Gorbunov and D. Rohr, on behalf of the ALICE Collaboration, *ALICE HLT high speed tracking on*
291 *GPU*, *IEEE Trans. Nucl. Sci.* **58** (2011), no. 4 1845–1851.

- 292 [4] F. Costa, *DDL, the ALICE Data Transmission Protocol and its Evolution from 2 to 6 Gb/s*, submitted
293 for publication in these conference proceedings.
- 294 [5] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008)
295 S08003.
- 296 [6] H. C. van der Bij, R. A. McLaren, O. Boyle and G. Rubin, *S-LINK, a data link interface specification*
297 *for the LHC era*, *IEEE Trans. Nucl. Sci.* **44** (1997), no. 3 398–402.
- 298 [7] R. Cranfield et al., *The ATLAS ROBIN*, *JINST* **3** (2008) T01002.
- 299 [8] F. Carena et al., *The ALICE data acquisition system*, *Nucl. Instr. Meth. Phys. Res. A* **741** (2014) 130 –
300 162.
- 301 [9] PLDA PCIe EZDMA IP core for Xilinx FPGAs, <http://www.plda.com>.
- 302 [10] D. Eschweiler and V. Lindenstruth, *The Portable Driver Architecture*, in *Proceedings of the 16th*
303 *Real-Time Linux Workshop*, Open Source Automation Development Lab (OSADL), October, 2014.
- 304 [11] T. Alt, *A FPGA based pre-processor for the ALICE High-Level Trigger*. PhD thesis,
305 Goethe-University Frankfurt, to be published.
- 306 [12] A. Borga et al., *Evolution of the ReadOut System of the ATLAS experiment*,
307 *ATL-DAQ-PROC-2014-012*, <https://cds.cern.ch/record/1710776> (2014).