

Advanced Readout Controller for PCI-based test systems of LHCb

Angel Guirao, Ken Wyllie, Francois Bal, Rui Pimenta, Hans Müller

CERN, 1211 Geneva 23, Switzerland
Angel.Guirao@cern.ch

Abstract

The Advanced ReadOut Controller (AROC) is a high-performance PCI card with a programmable FPGA node for two mezzanine-card based link ports. It is backwards compatible with the former PCI-FLIC which played an important role as the partition controller of NA60 for the readout of the legacy standards CAMAC and FERA. The success of the FLIC for the 40 MHz pixel chip test systems of LHCb and the need for more FLIC-like test systems resulted in a buffered, multi-link node architecture, which was inspired by the N*M node architecture of the former Readout Unit (RU) of LHCb. The Arco's host PC performs initialisation and configuration as standard PCI device. The AROC driver allows that high level applications, like Labview, can map directly into the Arco's data buffer or it's associated device registers. The status of the AROC and its first three applications are reviewed.

I. AROC OVERVIEW

Building up on experience with the programmable PCI-FLIC [1] controllers, the Aroc was designed as multi-link, programmable I/O node for advanced multi-link applications (Figure 1). The first target application is a test system for the Common L1 boards of LHCb where the Aroc operates as a data generator / receiver node, allowing to verify HLT and L1T data streams over Gigabit Ethernet based on the input data generated by the Aroc.

Previous FLIC-based test systems, such as the partition readout controller of NA60, or the 40 MHz pixel chip readout system [2] of LHCb, used Slink mezzanines to write detector data in a buffer which was read subsequently out via the PCI bus. Programmable I/O lines served as software controlled trigger. The new Aroc is backwards compatible with the FLIC but in addition allows building buffered I/O node architectures consisting of N input and M output links. PCI is used predominantly as the control bus, whilst data streams pass through a programmable N*M node which is implemented in an Altera Stratix FPGA (EP1S20F780C) device with 20k logical elements. An associated 128 Mbyte DDR memory serves as a quasi-simultaneous event buffer for both input and output.

The Aroc architecture was inspired by the FPGA-based Readout Unit (RU) [3] of LHCb which used multiple LVDS Slink mezzanines [4] for N-fold data input and a standard PMC mezzanine for transmitting event-packets towards an SCI-based trigger network [12]. Whilst the RU required a 9U crate and an embedded PMC card processor as PCI host,

the Aroc cards are inserted in standard PCs which serve both as "powered crate" and as PCI host that provides initialisation, control and data access for high level applications.

The Aroc's PMC/CMC and Slink connectors also provide backwards compatibility with former PCI-FLIC applications using ECL converters (RMH) or Slink LVDS I/O cards. Programmable LVDS signals for trigger applications and serial links for communication between Arocs are available via RJ45 connectors.

The system clocks for FPGA applications are several times faster than on the former FLIC cards. The 133 MHz DDR controller as well the 64 bit, 66 MHz PCI master/slave are implemented using commercial IP cores for the FPGA.

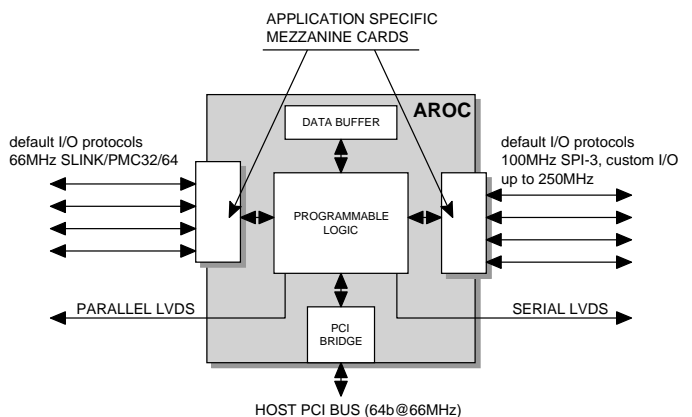


Figure 1: The Arco's Node Architecture

Three initial AROC applications are described.

As full-loop test system for the LHCb Tell-1 cards [5] the Aroc will be equipped with two different mezzanines, an optical transmitter on one side and a Gigabit Ethernet receiver on the other side.

As development platform for the FPGA-based output stage logic on Tell-1, a first version of an IP-packet driver has been developed [6] to generate the HLT and L1T data streams of LHCb, consisting of IP formatted Multi-Event Packets (MEP's) [7] which are transmitted via the 100 MHz SPI-3 bus to the MAC chips on the 2-channel GBE mezzanines [8] of LHCb.

Also the TRU trigger card of the Alice PHOS trigger [9] is emulated on the Aroc before the final TRU card is available. For this purpose, a quad ADC mezzanine was designed, to be interfaced via the Aroc's high speed SAMTEC connector to the Aroc's FPGA.

II. AROC DESCRIPTION

The hardware architecture of the Aroc is based on experience gained with two previous designs: the PCI-FLIC of NA60 and the Readout-Unit (RU) of LHCb. The mechanical format is a long-format universal PCI card which can carry up to two mezzanine cards. The serial and parallel user I/Os are available via RJ45 connectors. The photo of the Aroc is shown in Figure.2




Figure 2: the photo of the AROC-PCI card shows a PCI bridge chip close to the 64 bit PCI host connector, above it an Altera Stratix FPGA and on their left side, a Flash Prom and DDR chips. The connectors in the centre are three rows of IEEE mezzanine card connectors for Slink/PMC (left) and one row of a SAMTEC connector (right) for high bandwidth (OC48) Gigabit links. Two of the four RJ45 connectors provide serial I/O via 16 bit LVDS serializes; the other two provide user-programmable I/O lines.

A. Buffered FPGA-node architectures of LHCb

The previous RU modules [3] of LHCb were optimised for high rate trigger data, received from up to four input Slink mezzanines with 4 channels each. These were read out by pre-processing FPGAs which assemble sub-events into a dual port data buffer, to be read-out by a FPGA node and outputting directly towards a 0.512 Gbyte/s PCI bus. All FPGAs on the RU contain 64-bit PCI master/slave cores, making configuration and data access via local PCI bus segments easy. An embedded processor PMC card served as Monitoring and Control Unit (MCU [10]) with remote boot capability for LINUX. Connected to the top level of bridged PCI bus segments on the RU board, the MCU card played the role as PCI monarch for the initialisation of all onboard PCI devices (FPGAs and NICs). The use of PCI as I/O bus enabled using commercial PMC cards as standard interfaces towards commercial network links like ATM, Gigabit Ethernet, Myrinet or SCI. A serial link, named TAGnet, with LVDS receivers/transmitter ports allowed interconnecting rows of RU modules within a serial TAGnet ring. A TAGnet master provided the destination assignment [11] for the FPGA-based hardware DMA engines. The resulting event-coherent data streams [12] were directed towards free columns of a 2-dimensional SCI torus network [13].

The Tell-1 architecture is similar to the RU, using up to four custom mezzanines for data input; each equipped with pre-processor FPGA's storing event fragments in a DDR memory which is used as a dual port event buffer. Events are read out by one central FPGA which assembles Multi-event Packets (MEP's) and transmits these via the point-point SPI-3

bus (0.416 Gbyte/s) to LHCb's custom multi-channel Gigabit Ethernet mezzanine. An optical TTC receiver port receives L1 trigger decisions and the destination addresses for the DMA engines, transmitting both L1T and HLT streams towards a Gigabit mezzanine. The initialisation and control is performed via a local 16 bit bus, derived from an embedded computer module from Digital Logic [14] which is used in LHCb as standard for embedded slow controls.

The purpose of the Aroc was to design a PCI card, similar to the buffered FPGA node architecture of the RU and Tell-1, however making use of the PCI bus and the excellent PCI support available in any standard PC. The Aroc's general purpose architecture has immediately found 3 applications.

B. New Aroc Node Architecture

The Aroc applies a similar, however more general architecture than the RU. A typical number of 3 Arocs may be used in modern PC with ATX chassis, the limiting factor being the power supply. As depicted in Figure 3, the data flow is directed from one mezzanine to the other one with intermediate processing in an FPGA and data buffering in DDR RAM, enabling quasi simultaneous data buffering for both mezzanines.




Figure 3: Data path architecture

The FPGA node is also interfaced to a local PCI bus segment on the Aroc mezzanine card, allowing the use of standard PMC mezzanines, which become part of the PCI device hierarchy of the host system during initialisation. In order not to compete with the PCI bandwidth on the host's PCI bus, a PCI bridge separates the local PCI segment, such that the FPGA and the PMC card can exchange data at full 64 bit@66 MHz PCI bandwidth like on the RU.

C. Front-side PCI or Slink mezzanines

The front-side mezzanine can accommodate either CERN Slink32/64 or IEEE-standard PMC 32/64 mezzanines, which may also use the optional PN4 connector for user defined I/O.

The Slink I/O controller logic in the FPGA is a 64 bit extension of the previous 32 bit Slink controller on the FLIC cards.

A transparent PCI bridge isolates the local PCI bus from the main PCI bus segment such that 64bit@66MHz PCI

transactions between a PMC mezzanine and the FPGA do not influence the host PCI bus. A full interrupt service can be provided to any multi-function PMC mezzanines or to the FPGA.

The 64-bit PCI bus may be clocked up to 66 MHz, though a 33 MHz operation is usually sufficient for host initialisation and for PCI-driver controlled access operations to AROC resources. The driver was generated under W2000/XP using a licensed commercial package [15].

D. Back-side Gigabit mezzanines

The backside mezzanine connector is to accommodate the LHCb Gigabit Ethernet mezzanines via SAMTEC QTS/QSS 50Ω connectors [16]. The SPI-3 industry protocol for the GBE mezzanines of LHCb was implemented as a custom VHDL module for the FPGA on the Aroc or Tell-1 cards.

The initialisation of the GEB cards requires a multiplexed 16-bit microprocessor bus, which is mapped from PCI via FPGA logic. In this way, high-level test software on the host PC can configure the MAC chip registers on the Gigabit mezzanine like a normal PCI peripheral.

E. Custom mezzanines

The front-side connector pins add up to a total of 120 programmable FPGA pins for custom use of the two 64 pin Slink and the PN4 connectors. An example is the custom transmitter card for testing the Tell-1 test system, under design for transmitting pre-assembled test data via 12-fold optical ribbon fibres: An FPGA-driven 16bit@80 MHz bus is connected to a 1.6 Gbit/s serializer chip on the mezzanine, using the 64 pin IEEE connectors (otherwise used for Slink mezzanines). The serial stream is passed through 2.5 Gbit/s programmable switch chips in order to drive all or a subset of the 12-way optical laser module.

The SAMTEC high speed connector of the back-end can also be used as a custom connector, offering a total of 143 pins to the user, including 2 clock lines. The quad ADC mezzanine of the ALICE PHOS trigger is an example.

F. Software for Aroc systems

The software for Windows 2000/XP is based on the Windriver tool from Jungo [15] for creating PCI drivers. The preferred user-level software is National Instrument's Labview. The dynamic load library (Aroc.dll) was written in Visual C [17]. As software interface between Labview and the driver, it provides a simple mechanism to access resources on the Aroc card.

LINUX applications may use the simple "mmap" memory access mechanism of Linux which maps PCI memory into user program space. This method has been used extensively on the PCI-FLIC. Alternatively, for better real-time performance, Linux kernel modules can be written which communicate directly with the kernel. This method has been used for the realtime TAGnet driver [11] on the PCI-FLIC under Linux.

The PLDA PCI core [18] used in the Stratix FPGA provides the standard PCI 2.2 configuration space which allows implementing registers and memory via PCI's standard Base Address Register (BAR) conventions. The PCI memory structure of the Aroc in terms of BAR registers is shown in table 1.

Table 1: AROC PCI memory map

| | |
|------------|------------------------------|
| BAR0 | Main DDR buffer |
| BAR1, BAR2 | Reserved for future use |
| BAR3 | Direct access to SPI-3 ports |
| BAR4 | GE MAC configuration |
| BAR5 | S-LINK configuration |

BAR 3, 4 and 5 are assigned here to baseline Slink or Gigabit cards used on the Aroc. Custom cards will use these BAR's depending on the chosen interface connector.

The basic Labview programs developed for the test system of LHCb provides control the configuration of the PMC-Sierra MAC chip on the Gigabit mezzanine. In order to monitor the Ethernet packets generated by the Aroc, the public domain network analyser Ethereal was used. Further versions of the Aroc-based Gigabit Ethernet platform will centralize the Labview control over the complete test system (figure 4).

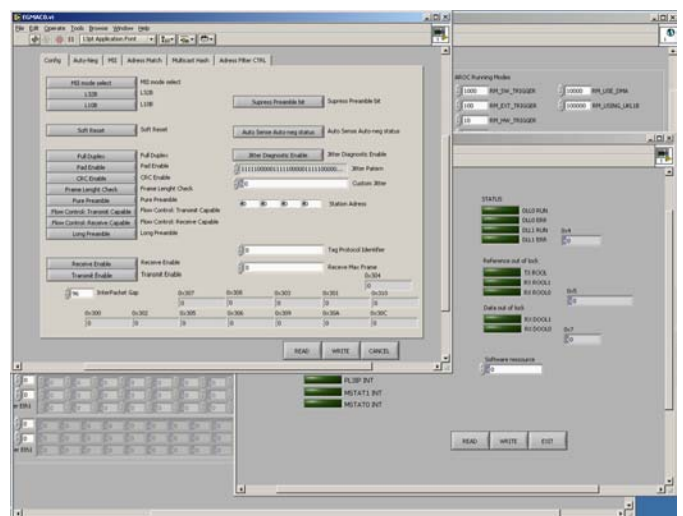


Figure 4: Labview control screen of the Gigabit mezzanine test system on the AROC

III. FIRST AROC APPLICATIONS

A. Gigabit links for LHCb

The LIT and High-Level Trigger (HTL) data streams of LHCb are based in the Gigabit Ethernet (GE) technology. The TELL-1 boards [5] send the LIT data to computer farms by driving Ethernet packets over a MAC chip to the physical links. There are two mezzanine types, providing two or four GE duplex channels depending on the detector requirement.

B. Test system for L1 electronics of LHCb

A test system for the LHCb Tell-1 electronics is under design in which the Aroc produces data fragments like an LHCb front-end chip. In order to verify the operation of the Tell-1 electronics (i.e. production of a stream of formatted Gigabit packets), the Aroc's data buffer is pre-loaded with software-generated data fragments and with their resulting Gigabit packets (formatted as IPv4 packets). During the test, fragments are transmitted via a 12-fold optical ribbon cable to the inputs of the Tell-1 cards. A custom, parallel fibre transmitter mezzanine using the Slink connectors is under way. For the Aroc's receiving port, the bi-directional Gigabit Ethernet cards are used as data receiver. In test mode, the Aroc's hardware will transmit, receive and compare data in a full loop at transmit rates in the 1 MHz range. The test system is conceived to be operated by dedicated Labview VI's which also allow to configuring MAC registers and read/modify the DDR memory of the Aroc.

The test system is a full-loop concept where the Aroc is mastering the Tell-1 board like a Device-Under-Test (DUT). For this purpose, the Aroc is equipped with two mezzanines, one of which generates and transmits data to the DUT while the other mezzanine receives the generated Gigabit output stream for comparison with the input fragments (figure 5).

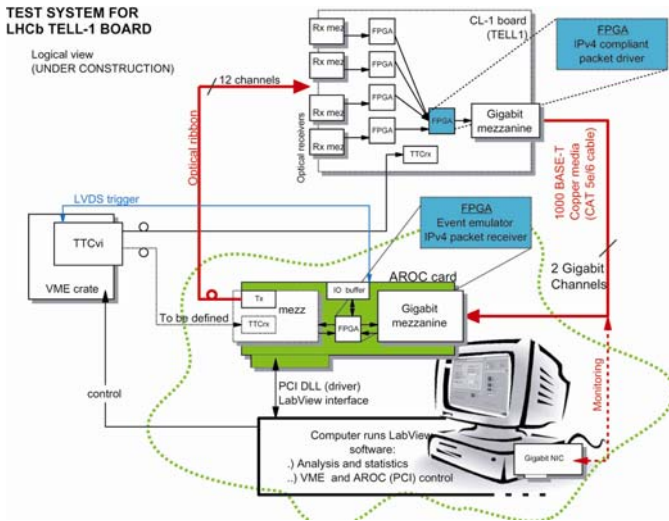


Figure 5: Test System for LHCb Tell-1 Electronics (simplified)

C. The LHCb Gigabit Ethernet card

For the test and development of FPGA-based Gigabit stream drivers for LHCb, dual /quad Gigabit mezzanines (Figure. 6) are interfaced via the 100 MHz SPI-3 bus connector on the Aroc's back-side. SPI-3 is a FIFO like bus, suited for bi-directional master-write protocols.

The MAC chip on the mezzanine can either drive optical transceivers or PHYSical chips serving 2 copper links. The first Gigabit card under test has two channels, each having its own egress or ingress FIFO. The channels share the same 416 Mbyte/s SPI-3 bus and are selected dynamically via a simple address mechanism.

D. The Alice PHOS Trigger emulator

Another application example is the FPGA-based PHOS L0/L1 trigger processor TRU [9] of the Alice experiment.




Figure 6: Dual channel Gigabit mezzanine for LHCb

The AROC-based TRU setup emulates the ALICE PHOS calorimeter trigger electronics which will have 112 analogue-summed input channels corresponding to $2 \times 16 \times 14$ crystals of an ALICE PHOS module. In order to develop the trigger algorithms before the final TRU trigger board for 112 channels becomes available, a custom ADC mezzanine has been designed to fit on the Aroc's high speed SAMTEC connector. It contains a subset of 4 channels of the Alto-16 chip [19]. Each analogue differential input channel is amplified /attenuated through a digitally programmable amplifier, with different settings for high- p_T and low- p_T physics. The ADC's digitise the input signals with 10 bit resolution at a 20 MHz sampling rate.

The serial links of the Aroc can also be used to combine data of adjacent TRU processors into a "Super-TRU". No output mezzanine is needed for the trigger signals since the Aroc's direct LVDS input/outputs can be used to transmit the NRZ encoded trigger decisions at 40 MHz.

The TRU emulator setup with ADC mezzanine is depicted in figure 7. The interface between the AROC and the mezzanine uses the high-speed SAMTEC connector that delivers 40 bit of digitised data (4 channels), clocks and control signals, including the gain control bits of the AD8369 digitally controlled amplifiers.




Figure 7: AROC-based PHOS trigger development platform

All ADC mezzanine signals are connected to the Aroc's FPGA which is of the same type as on the final TRU card. The L0 trigger is based on charge summing over all combination of 4*4 crystal areas with the requirement to output NRZ-encoded trigger decisions at 40 MHz. Hence a very fast peak finding algorithm has to be applied to a very low number of ADC samples. The Aroc is used to develop and test the algorithms with four ADC test channels and input test pulses. All gathered data and relevant status registers of the FPGA can be monitored through the PCI bus which on this case also emulates the Readout Control Unit (RCU) of Alice.

IV. SUMMARY

The AROC-PCI card is backwards compatible with the former PCI-FLIC but offering an enhanced N*M node architecture similar to the LHCb Readout Unit. The number of links, N or M, is defined by the mezzanine cards used on the Aroc. A large variety of standard or custom mezzanines can be used and three specific ones (dual Gigabit, quad ADC and 12-fold optical) are close to finalization. The Aroc has found immediately three applications, demonstrating its role as a general-purpose PCI platform for data processing and data acquisition. Its target application, the full loop test system for the LHCb Level-1 electronics, is under way whilst first results with a Gigabit Ethernet stream driver on the Aroc were reported successful. The Emulator for the PHOS trigger with a custom-designed quad ADC mezzanine is equally in progress. First versions of Labview-based, high level control software for Aroc based test systems is operational using an Aroc specific PCI driver and library for Windows OS.

V. ACKNOWLEDGEMENTS

The authors would like to thank Antoine Junique for his professional PCB layout of the TRU mezzanine, Luciano Musa for his kind assistance for using the ALTRO-16 chip on a mezzanine card, Antonis Bonos for generating the PCI driver and providing the first Labview software for the AROC and Guido Haefeli for excellent collaboration in the co-design of the Aroc and the TELL-1 boards.

VI. REFERENCES

[1] *PCI-FLIC: A plug & play approach to Data Acquisition*, Toledo, J.; Muller, H.; Buytaert, J.; Bal, F.; David, A.; Guirao, A.; Mora, F.J. IEEE Trans. on Nucl. Science Vol 49, Issue 3, Part: 2, June 2002 Page(s): 1190 -1194

- [2] *Readout of High Speed S-link data via a buffered PCI card*, A.Guirao et al. PCaPAC 4th International Workshop on Personal Computers and Particle Accelerator Controls, 14-17 October 2002 Frascati (Italy)
- [3] *A Readout Unit for high rate applications*, J.Toledo, F.Bal, D.Dominguez, A.Guirao, H.Müller, Proc. 12th Real Time Congress on Nucl. and Plasma Sciences, Valencia June 2001, p 230 ff. and Transactions on Nucl. Science, VOL 49, N02, April 2002, p 448-454
- [4] *LVDS link cards (Slink) for multiplexers, Velo trigger and other applications*, Francois Bal, H.Muller, LHCb Technical note LHCb DAQ 2002-005
- [5] *Tell1: Common L1 read out board for LHCb* (this conference) G.Haefeli et al. (University of Lausanne)
- [6] *Common Gigabit Ethernet interfaces for HLT and L1-trigger links of LHCb* (this conference) H.Müller, F.Bal, A.Guirao
- [7] *Requirements to the L1 front-end electronics*, LHCb Technical Note LHCb 2003- 078
- [8] Gigabit Ethernet mezzanines for DAQ and Trigger links of LHCb, H.Muller, F.Bal, A.Guirao, LHCb note DAQ 2003-021, 28 April 2003
- [9] *Trigger Electronics for the Alice PHOS detector*, H.Müller et al., 9th Pisa meeting on advanced detectors, May 2003, Elba Italy (to appear in Nucl. Instr. Methods special edition)
- [10] *A networked mezzanine Linux processor card*, A.Guirao et al. 12th Proceedings IEEE Real Time Conf, RT 2001, Valencia June 200, p 81-83
- [11] *Logiciel haute performance pour carte TAGnet*, Sebastien Gonzalve, Stagiaire Francais CERN April-Sept 2002, Stage 3eme annee ISIMA (french only)
- [12] *TAGnet, a Twisted Pair Protocol for Event-Coherent DMA transfers in Trigger Farms*, H.Muller, F.Bal, S.Gonzalve, A.Guirao, F. Vinci dos Santos, Proceedings 8th Workshop on Electronics for LHC, Colmar 2002 p. 289 ff.
- [13] *A scalable 1 MHz Trigger Farm Prototype with Event-Coherent DMA input*, Realtime Conference 2003, Montreal, May 18-23, 2003 Ivan Kisel et al..
- [14] Digital Logic AG <http://www.digitallogic.com>
- [15] Jungo Windriver <http://www.jungo.com>
- [16] SAMTEC high speed QSS connectors <http://www.samtec.com>
- [17] *Developpement de logiciel de configuration d'un chip MAC via logique FPGA sur port PCI*, Antonis Bonos, Stagiaire Francais CERN April -Sept 03 Stage 2eme annee ISIMA (french only) <http://cern.ch/ep-div-ed/Documents/Mini-these-Antonis.pdf>
- [18] PCI IP core: PLD Applications <http://www.plda.com>
- [19] *The ALTRO Chip: A 16-channel A/D Converter and Digital Processor for Gas Detectors*, R. Esteve Bosch et al., Proc. IEEE NSS/MIC, November 2002, Norfolk Virginia