

## FELIX: a PCIe based high-throughput approach for interfacing front-end and trigger electronics in the ATLAS Upgrade framework

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2016 JINST 11 C12023

(<http://iopscience.iop.org/1748-0221/11/12/C12023>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 128.141.192.122

This content was downloaded on 25/01/2017 at 16:20

Please note that [terms and conditions apply](#).

You may also be interested in:

[The Error Reporting in the ATLAS TDAQ System](#)

Serguei Kolos, Andrei Kazarov and Lykourgos Papaevgeniou

[Performance of the Demonstrator System for the Phase-I Upgrade of the Trigger Readout Electronics of the ATLAS Liquid Argon Calorimeters](#)

N. Dumont Dayot

[ATLAS TDAQ System Administration: evolution and re-design](#)

S Ballestrero, A Bogdanchikov, F Brasolin et al.

[Study of the VMM1 read-out chip in a neutron irradiation environment](#)

T. Alexopoulos, G. Fanourakis, T. Geralis et al.

[Development and implementation of optimal filtering in a Virtex FPGA for the upgrade of the ATLAS LAr calorimeter readout](#)

S Stärz

[ATLAS Level-1 Calorimeter Trigger upgrade for Phase-I](#)

W Qian

[A new approach to front-end electronics interfacing in the ATLAS experiment](#)

J. Anderson, A. Borga, H. Boterenbrood et al.

[The clock distribution system for the ATLAS Liquid Argon Calorimeter Phase-I Upgrade Demonstrator](#)

B. Deng, H. Chen, K. Chen et al.

[Front-end electronics for the FAZIA experiment](#)

F. Salomon, P. Edelbruck, G. Brulin et al.

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS,  
26–30 SEPTEMBER 2016,  
KARLSRUHE INSTITUTE OF TECHNOLOGY (KIT), KARLSRUHE, GERMANY

## FELIX: a PCIe based high-throughput approach for interfacing front-end and trigger electronics in the ATLAS Upgrade framework

J. Anderson,<sup>a</sup> K. Bauer,<sup>c</sup> A. Borga,<sup>f</sup> H. Boterenbrood,<sup>f</sup> H. Chen,<sup>b</sup> K. Chen,<sup>b,1</sup> G. Drake,<sup>a</sup>  
M. Dönszelmann,<sup>i</sup> D. Francis,<sup>d</sup> D. Guest,<sup>c</sup> B. Gorini,<sup>d</sup> M. Joos,<sup>d</sup> F. Lanni,<sup>b</sup>  
G. Lehmann Miotto,<sup>d</sup> L. Levinson,<sup>h</sup> J. Narevicius,<sup>h</sup> W. Panduro Vazquez,<sup>e</sup> A. Roich,<sup>h</sup>  
S. Ryu,<sup>a</sup> F. Schreuder,<sup>f</sup> J. Schumacher,<sup>d,g</sup> W. Vandelli,<sup>d</sup> J. Vermeulen,<sup>f</sup> D. Whiteson,<sup>c</sup>  
W. Wu<sup>b</sup> and J. Zhang<sup>a</sup>

<sup>a</sup>Argonne National Laboratory, 9700 South Cass Avenue B109, Lemont, IL 60439, U.S.A.

<sup>b</sup>Brookhaven National Laboratory, P.O. Box 5000, Upton, NY 11973-5000, U.S.A.

<sup>c</sup>University of California, Irvine, Irvine, CA 92697, U.S.A.

<sup>d</sup>CERN, CH-1211 Geneva 23, Switzerland

<sup>e</sup>Royal Holloway, University of London, Egham Hill, Egham TW20 0EX, United Kingdom

<sup>f</sup>Nikhef National Institute for Subatomic Physics / University of Amsterdam,  
Science Park 105, 1098 XG Amsterdam, Netherlands

<sup>g</sup>Department of Computer Science, University of Paderborn, Pohlweg 47, 33098 Paderborn, Germany

<sup>h</sup>Department of Particle Physics, The Weizmann Institute of Science, Rehovot 76100, Israel

<sup>i</sup>Radboud University Nijmegen, Comeniuslaan 4, 6525 HP Nijmegen, Netherlands

E-mail: [kchen@bnl.gov](mailto:kchen@bnl.gov)

**ABSTRACT:** The ATLAS Phase-I upgrade (2019) requires a Trigger and Data Acquisition (TDAQ) system able to trigger and record data from up to three times the nominal LHC instantaneous luminosity. The Front-End Link eXchange (FELIX) system provides an infrastructure to achieve this in a scalable, detector agnostic and easily upgradeable way. It is a PC-based gateway, interfacing custom radiation tolerant optical links from front-end electronics, via PCIe Gen3 cards, to a commodity switched Ethernet or InfiniBand network. FELIX enables reducing custom electronics in favour of software running on commercial servers. The FELIX system, the design of the PCIe prototype card and the integration test results are presented in this paper.

**KEYWORDS:** Data acquisition circuits; Data acquisition concepts

<sup>1</sup>Corresponding author.

---

## Contents

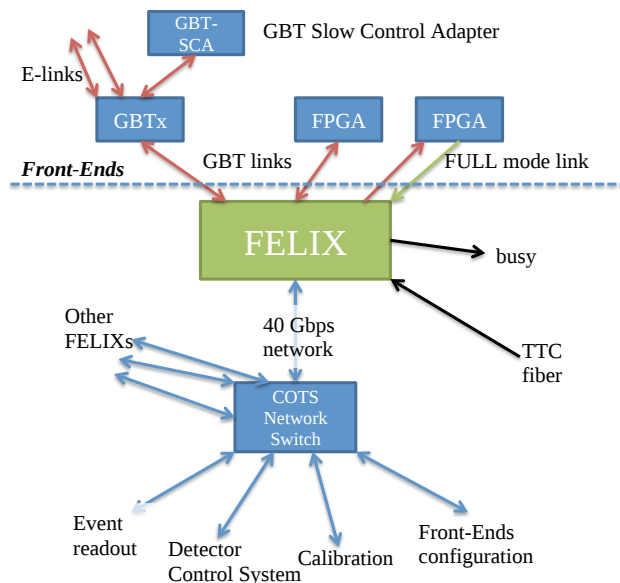
<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>FELIX card prototype</b>	<b>3</b>
<b>3</b>	<b>FELIX firmware</b>	<b>5</b>
<b>4</b>	<b>Integration test results</b>	<b>6</b>
4.1	Integration with LTDB	6
4.2	Integration with CaRIBOu system	6
<b>5</b>	<b>Conclusion</b>	<b>7</b>

---

## 1 Introduction

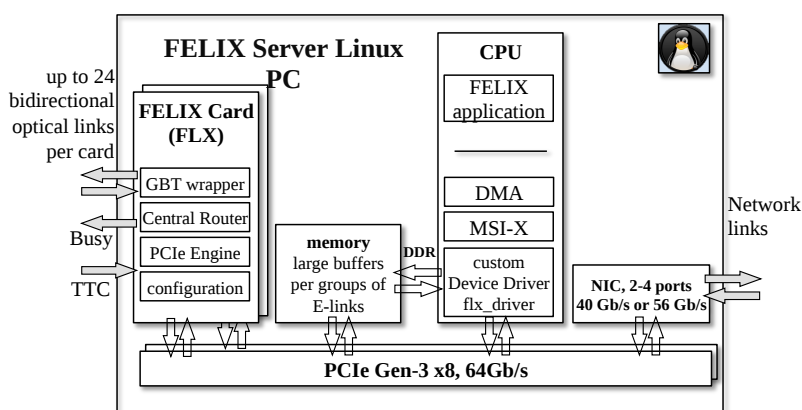
The ATLAS [1] Phase-I upgrade requires a Trigger and Data Acquisition (TDAQ) system able to trigger and record data from up to three times the nominal Large Hadron Collider (LHC) instantaneous luminosity. During the LHC Long Shutdown 2 (2019-2020), new ATLAS on-detector electronics for the Liquid Argon (LAr) Calorimeters and New Small Wheel (NSW) muon detectors will be installed. A new detector independent readout architecture, named Front-End Link eXchange (FELIX), will provide access to the TDAQ systems in a scalable, detector agnostic and easily upgradeable way, to the above mentioned FrontEnd systems. On one side of the FELIX system, the GigaBit Transceiver (GBT) [2] architecture and a protocol developed by CERN provides a high-speed (4.8 Gb/s) radiation-hard optical link for data transmission from the on-detector FrontEnd electronics. By means of time multiplexing, the GBT protocol provides up to 42 independent data links, however sharing the same fibre. As shown in figure 1, through such links FELIX receives and identifies different information streams, the so called “e-links”.

Via a Field Programmable Gate Array (FPGA) based PCIe card, the detailed description of which follows in this paper, the GBT link data is funneled to the host PC memory. From there on



**Figure 1.** Role of the FELIX system in the ATLAS Trigger and DAQ system.

data packets will be routed via a commercial switched network. In the opposite direction, FELIX receives packets from the network and forwards them to specific on-detector electronic modules. The GBT frame consists of 116 bit data (payload) and a 4 bit header. The header is used to align the frame at the receiver side, and it can either be DATA (0101) or IDLE (0110). Four of the 116 bit data are used for slow control and monitoring, based on the High-Level Data Link Control (HDLC) [3] protocol. By means of HDLC, FELIX will support the configuration of GBTx ASICs [4] via the IC (Internal Control) slow control (2-bits). The 2-bits for the EC (External Control) or any regular 2-bit e-link can be used to communicate with the GBT-SCA (Slow Control Adapter) ASIC [5]. The SCA is used to control and monitor devices on on-detector Front-End boards. The SCA supports user interface ports like I2C, SPI, JTAG and GPIO. Finally, a FELIX system has to handle the input from the Time, Trigger and Control (TTC) system [6], by recovering the LHC clock and forwarding the machine-synchronous trigger information. This information will be distributed to on-detector electronics over low-and-fixed-latency GBT links, and also to new and upgraded off-detector first-level trigger systems. For readout of the latter a lightweight protocol, the so called “FULL mode”, with higher throughput (9.6 Gb/s) than the GBT protocol is envisaged to be used.

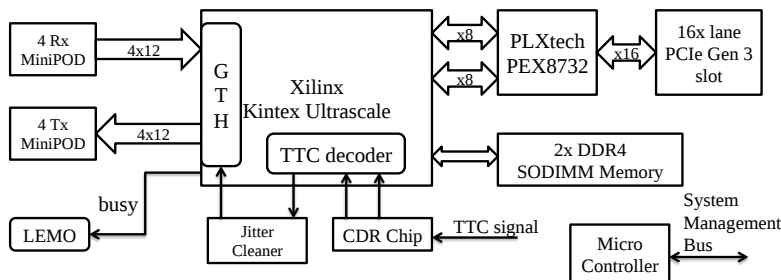


**Figure 2.** Diagram of the FELIX system.

All functions described above are implemented in FPGAs mounted on PCIe cards, the so called “FLX cards”. As a system, FELIX consists of a PC running a Linux based OS (SLC6), an Ethernet or InfiniBand Network Interface Card (NIC), and up to two FLX cards, as depicted in figure 2. The Hitech Global HTG-710 is used as an FLX demonstrator card. It is equipped with an 8-lane PCIe Gen3 (64 Gb/s) interface and with two CXP transceivers providing interfaces for 24 bidirectional optical links (max. 13.1 Gb/s). Moreover a custom mezzanine was designed to receive and decode the TTC clock and data information. The FLX demonstrator card firmware is also ported to a Xilinx VC-709 evaluation board, which has the same type of FPGA and PCIe interface as the HTG-710, but less optical interfaces. This second card targets detector and trigger system test setups. As none of the above options matches completely the final requirements in terms of FLX card, a third board, known as the FLX-711 (figure 5) had been developed and adopted as a candidate prototype of the final FLX card. Drivers and software tools have been developed for control and monitoring as well benchmarking of these boards. Data routing and the connection to the COTS (Commercial Off-The-Shelf) network is implemented in a software pipeline running on the FELIX host PC. The

packet processing performance satisfies the requirement of FELIX [7]. As the FELIX data handling and FELIX software aspects have been extensively discussed in previous documents [7, 8], the progress of the FELIX card prototype (FLX-711) design and testing will be the main focus of this paper. An update on the FELIX firmware development and the integration test results of FELIX with FrontEnds will also be presented.


## 2 FELIX card prototype



**Figure 3.** Diagram of FELIX prototype FLX-711.


Figure 3 shows a block diagram of the 16-lane Gen3 PCIe board: the FLX-711. The on-board FPGA is a Xilinx Kintex UltraScale XCKU115FLVF1924 with 64 high-speed GTH transceivers [9]. Since the Integrated Block for PCI Express supports at most 8-lane operation, a PCIe switch PEX8732 is used to connect two 8-lane endpoints to the 16-lane PCIe slot. This solution ensures the nominal bandwidth (128 Gbps) required by the specifications of the FLX-711. The other 48 GTH transceivers are connected to 8 MiniPODs [10], 4 used as transmitters and 4 for as receivers. The speed of these 48 optical links can be up to 14 Gb/s which is limited by the MiniPODs. An ADN2814 [11] is used to recover the 160 MHz LHC TTC clock and data. An on-board jitter cleaner chip is then used to clean the TTC clock, and provide clean reference clock [12] for transceivers. As this board is also used as part of the test setup for the LAr Trigger Digitizer Board (LTDB) in the ATLAS LAr Phase-I upgrade [13], where buffering of 320 channels Analog-to-Digital Converter (ADC) data are required, two DDR4 small outline dual in-line memory modules (total capacity 16 GB) have been added on-board. Lastly, a micro-controller (ATMEGA256A) is used to program the FPGA from selectable bitfiles stored in a flash memory [14]. Software in the PC communicates and controls the reconfiguration process via the System Management Bus. Figure 4 shows the complexity of the Printed Circuit Board (PCB) stackup.

Due to the board size limit, the FLX-711 uses two types of blind vias to achieve the complete routing: one is for the high-speed signals connected with MiniPODs, the other is for the dense



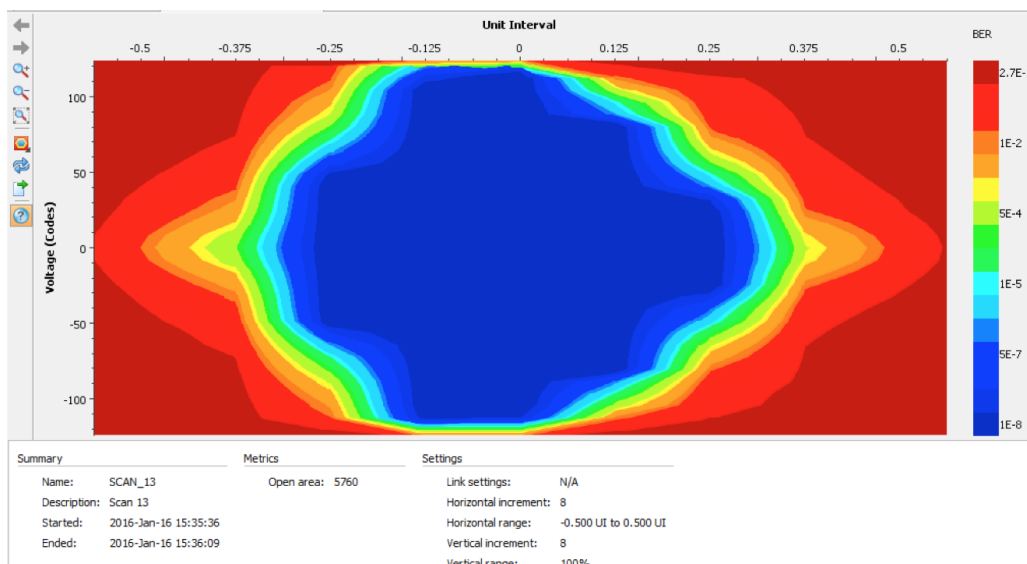
**Figure 4.** Stackup of FLX-711 board.

DDR4 traces. Because of this layer stackup three laminations are required for manufacturing the PCB. The 1078LR, 1078MR and 1078HR are different prepreg constructions of Megtron6 material [15].



**Figure 5.** The FLX-711 card.

The first version of the FLX-711 board is shown in figure 5. All of the hardware features have been successfully tested. To test the PCIe interface two Wupper [16] Direct Memory Access (DMA) engines are implemented in the FPGA, while counter data is used to test the throughput to the server. The total measured throughput of these two 8-lane PCIe endpoints can be up to 101.7 Gb/s, in agreement with the PCIe specifications. To test the optical links the IBERT [17] from XILINX is used to perform BER (Bit Error Rate) testing at link speeds of 12.8 Gb/s, for all of the 48 fiber optical links. All of the 48 channels are connected as loopback. The result shows that the BER is smaller than  $10^{-15}$ ; an eye diagram of a typical channel is shown as in figure 6.



**Figure 6.** The eye diagram of one channel at 12.8 Gb/s.

On the remote reconfiguration interface: the i2c-tools [18] can successfully communicate with the micro-controller through the System Management Bus. The micro-controller is in turn able to set the FPGA configuration pins to initiate the FPGA programming, loading the image from the



target segment in the Flash memory. The highest two address pins of the Flash are controlled by the ATMEGA256A, thus splitting the Flash into 4 segments.

### 3 FELIX firmware

Details of the FELIX firmware are described in [8]. The main modules are the Wupper PCIe Engine, the Central Router for internal data multiplexing, the TTC decoder and the optimized GBT-FPGA core. A block diagram of the firmware for FLX-711 is shown in figure 7. The overall occupancy in a KCU115 FPGA is 22% LUTs for 4 channels. The Central Router takes about 16% of the total, and scales linearly with the channel number. Since there are two PCIe endpoints, two distinct Wupper PCIe Engines are implemented. Half of the channels are therefore connected to each engine. As introduced earlier, a customized lightweight protocol, called “FULL mode”, is defined for the links between FPGA-based FrontEnds and FELIX systems; the goal being to provide a higher maximum payload. Figure 8 shows a block diagram of both the FrontEnd and FELIX ends of a FULL mode link. The link speed is 9.6 Gb/s, but as data is encoded in 8b/10b a maximum user payload of 7.68 Gb/s can be achieved. The packet size is in units of 32 bit. As an upper limit estimation, eight channels, each with a maximum payload of 7.68 Gb/s could be transferred within the PCIe Gen3 8-lane bandwidth (maximum 64 Gb/s). In the FLX-711 case, as the PCIe interface is 16-lanes, up to 16 channels can be supported.

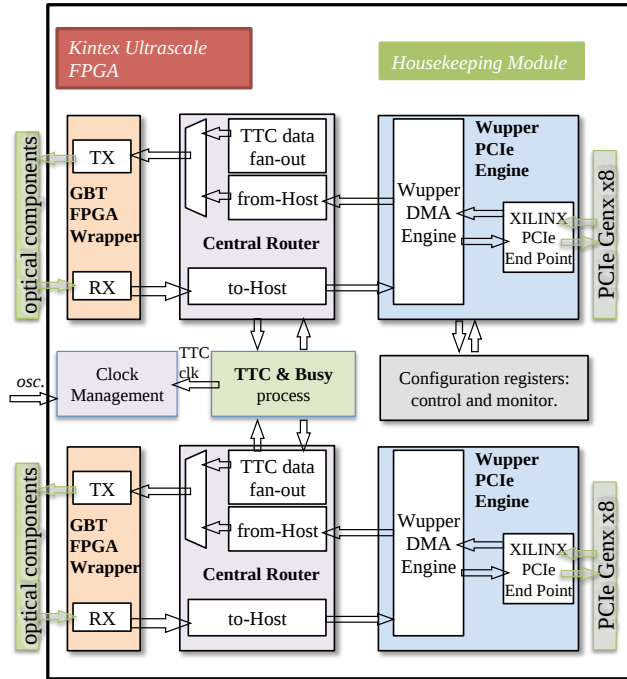


Figure 7. The firmware diagram of FLX-711.




Figure 8. The FELIX FULL mode.

## 4 Integration test results

FELIX will be used to interface several FrontEnds, such as the LAr calorimeter trigger electronics and the muon system New Small Wheel (NSW) in the ATLAS Phase-I upgrade. For the Phase-II upgrade of HL-LHC (High-Luminosity LHC), the plan is to adopt FELIX for interfacing all the FrontEnds. A series of integration tests with the LTDB and the Control and Readout ITk (Inner Tracker) Board (CaRIBOu) have been done in the course of this year, and will be summarized in the following paragraphs.

### 4.1 Integration with LTDB

In the LAr Phase-I upgrade, the LAr Trigger Digitizer Board (LTDB) is used to digitize the input analog signals, and transmit them to the back-end [13]. On the LTDB prototype, there are five GBTx and five GBT-SCA chips. GBT-SCA chips are used to control the power, the I2C slaves and perform the on-board temperature measurement. Besides the interface to EC links with the GBT-SCA chip, each GBTx on the LTDB provides the recovered 40 MHz TTC clock to the ASICs NEVIS ADC [19] and serializers LOCx2 [20], and sends the BCR (Bunch Crossing Reset) signal to LOCx2. Both FLX-709 and FLX-711 have been used to demonstrate the interfacing to the LTDB board. Test results show that the LTDB works according to specifications, and can cope with the recovered clock and BCR signal sent from the FELIX GBT links. Communication with the GBTx and GBT-SCA on the LTDB can also be performed by using the IC and EC bits in the bidirectional GBT frame.

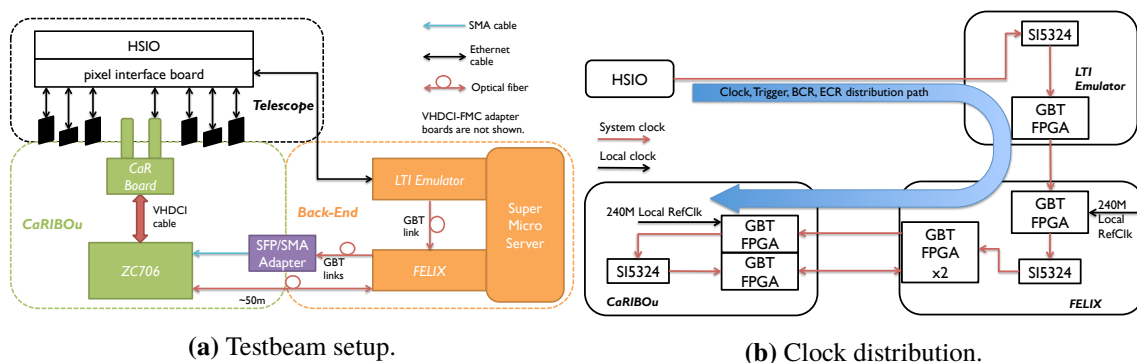
### 4.2 Integration with CaRIBOu system

CaRIBOu is a modular test system for silicon sensor research and development for the ATLAS upgrade [21]. It consists of several boards: the control and readout board, the Xilinx ZC-706 evaluation kit, and several Front-End chip carrier boards. The FELIX demonstrator VC-709 has been used as back-end to interface to CaRIBOu systems. The integration beam test has been successfully carried out at CERN in August 2016.

A block diagram of the test setup is shown in figure 9a. The pixel interface board outputs the system clock, and the commands to all of the FE-I4 [22] boards via RJ-45 connectors. Commands are encoded in the appropriate FE-I4 format. One Ethernet cable is used to connect the clock and the commands to a VC-709 which emulates the function of an LTI (Local Trigger Interface) board. Firmware on this board decodes the FE-I4 command format, and extracts the TTC signals, such as the trigger, BCR and Event Counter Reset (ECR) commands. These TTC signals are then sent to FELIX, which then distributes them to the FrontEnd CaRIBOu system. The links from the LTI emulator to FELIX and between FELIX and CaRIBOu are all GBT links. Figure 9b shows the clock distribution of the system. SI5324 [23] on the LTI emulator cleans clock from the Ethernet cable, and generates synchronized reference clock for GTH transceivers. FELIX recovers the system clock from one link of the LTI emulator, cleans it with SI5324, uses it as its system clock and also as a reference clock for other transceivers. Similarly CaRIBOu recovers the system clock and uses it as reference clock for the data and command transmission link with FELIX. This clock scheme guarantees that all system clocks of all the different boards are synchronized. The data



from CaRIBOu to FELIX, the commands and status information transferred between FELIX and CaRIBOu, are thus all synchronized to the same system clock.



**Figure 9.** Diagrams for the testbeam.

The test shows that FELIX can be used to do the calibration of the CMOS sensor AMS180V4 [24] and the readout ASIC FE-I4B. The data from FE-I4B is encoded inside the ZC-706, in the format specified by FELIX. One 8-bit e-link is used for the data transmission. The whole system works according to specifications, and showed no glitches for more than 12 hours runs. The software provided by FELIX can stream data continuously to disk.

## 5 Conclusion

FELIX is a system that interfaces the custom links used in the ATLAS readout system to standard commercial networks. For the links toward FrontEnds, it supports the GBT protocol. For the FrontEnds which needs a higher throughput toward FELIX, a customized lightweight FULL mode is defined. The development of the FELIX system is still ongoing, but integration tests with several FrontEnd systems described in this paper, show that the FELIX system with both the demonstrator and prototype cards can: (i) stably distribute TTC information, (ii) receive data from FrontEnd, (iii) save it to disk through the PCIe interface. A clear sign that FELIX is steadily heading toward its Final Design Review scheduled for the End of 2016.

## References

- [1] ATLAS collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, 2008 *JINST* **3** S08003.
- [2] P. Moreira et al., *The GBT, a Proposed Architecture for Multi-Gb/s Data Transmission in High Energy Physics*, in proceedings of *Topical workshop on electronics for particle physics (TWEPP 2007)*, Prague, Czech Republic, 3–7 September 2007, pp. 332–336.
- [3] ISO/EIC 13239: 2002(E), *Information technology - Telecommunications and information exchange between systems - High-Level Data Link Control (HDLC) procedures. HDLC standard*.
- [4] *CERN GBT Project: GBTX Manual*.

- [5] A. Caratelli, S. Bonacini, K. Kloukinas, A. Marchioro, P. Moreira, R. De Oliveira et al., *The GBT-SCA, a radiation tolerant ASIC for detector control and monitoring applications in HEP experiments*, 2015 *JINST* **10** C03034.
- [6] RD12 PROJECT collaboration, B.G. Taylor, *TTC distribution for LHC detectors*, *IEEE Trans. Nucl. Sci.* **45** (1998) 821.
- [7] J. Schumacher et al., *Improving packet processing performance in the ATLAS FELIX project: analysis and optimization of a memory-bounded algorithm*, in *Proceedings of the 9<sup>th</sup> ACM International Conference on Distributed Event-Based Systems (DEBS '15)*, 2015, pp. 174–180.
- [8] J. Anderson et al., *A new approach to front-end electronics interfacing in the ATLAS experiment*, 2016 *JINST* **11** C01055.
- [9] XILINX, *UltraScale Architecture GTH Transceivers*, [http://www.xilinx.com/support/documentation/user\\_guides/ug576-ultrascale-gth-transceivers.pdf](http://www.xilinx.com/support/documentation/user_guides/ug576-ultrascale-gth-transceivers.pdf).
- [10] Avago Technologies, *MiniPOD Embedded Optical Modules*, <http://www.fit-foxconn.com/Product/ProductDetail?topClassID=&&PN=AFBR-814VXYZ>.
- [11] Analog Devices, *Continuous Rate 10 Mb/s to 675 Mb/s Clock and Data Recovery IC with Integrated Limiting Amp*, <http://www.analog.com/media/en/technical-documentation/data-sheets/ADN2814.pdf>.
- [12] Xilinx, *Kintex UltraScale FPGAs Data Sheet: DC and AC Switching Characteristics*, [http://www.xilinx.com/support/documentation/data\\_sheets/ds892-kintex-ultrascale-data-sheet.pdf](http://www.xilinx.com/support/documentation/data_sheets/ds892-kintex-ultrascale-data-sheet.pdf).
- [13] ATLAS collaboration, *ATLAS Liquid Argon Calorimeter Phase-I Upgrade Technical Design Report*, CERN-LHCC-2013-017.
- [14] A. Borga et al., *The C-RORC PCIe card and its application in the ALICE and ATLAS experiments*, 2015 *JINST* **10** C02022.
- [15] Panasonic, *High Speed, Low Loss Multi-layer Materials MEGTRON6*, [https://industrial.panasonic.com/content/data/EM/PDF/ipcdatasheet\\_1611\\_R-5775.pdf](https://industrial.panasonic.com/content/data/EM/PDF/ipcdatasheet_1611_R-5775.pdf).
- [16] *PCIe Gen3x8 DMA for virtex7*, [http://www.opencores.org/project,virtex7\\_pcie\\_dma](http://www.opencores.org/project,virtex7_pcie_dma).
- [17] XILINX, *IBERT for UltraScale GTH Transceivers v1.3*, [http://www.xilinx.com/support/documentation/ip\\_documentation/ibert\\_ultrascale\\_gth/v1\\_3/pg173-ibert-ultrascale-gth.pdf](http://www.xilinx.com/support/documentation/ip_documentation/ibert_ultrascale_gth/v1_3/pg173-ibert-ultrascale-gth.pdf).
- [18] *I2CTools*, <http://www.lm-sensors.org/wiki/I2CTools>.
- [19] J. Kuppambatti, J. Ban, T. Andeen, P. Kinget and G. Brooijmans, *A radiation-hard dual channel 4-bit pipeline for a 12-bit 40 MS/s ADC prototypewith extended dynamic range for the ATLAS Liquid Argon Calorimeter readout electronics upgradeat the CERN LHC*, 2013 *JINST* **8** P09008 [[arXiv:1308.0028](https://arxiv.org/abs/1308.0028)].
- [20] L. Xiao et al., *LOCx2, a low-latency, low-overhead, 2 × 5.12-GBPS transmitter ASIC for the ATLAS Liquid Argon Calorimeter trigger upgrade*, 2016 *JINST* **11** C02013.
- [21] H. Liu et al., *Development of a modular test system for the silicon sensor R&D of the ATLAS Upgrade*, [arXiv:1603.07950](https://arxiv.org/abs/1603.07950).
- [22] M. Garcia-Sciveres et al., *The FE-I4 pixel readout integrated circuit*, in *Proceedings of the Seventh International “Hiroshima” Symposium on the Development and Application of Semiconductor Tracking Detectors*, Hiroshima, Japan, August 29 – September 1 2009, *Nucl. Instrum. Meth.* **636** (2011) S155.

- [23] Silicon Labs, *Si5324 ANY-FREQUENCY PRECISION CLOCK MULTIPLIER/JITTER ATTENUATOR*, <http://www.silabs.com/Support Documents/TechnicalDocs/Si5324.pdf>.
- [24] HVC MOS collaboration, I. Peric, *Active pixel sensors in high-voltage CMOS technologies for ATLAS*, 2012 *JINST* **7** C08002.