

# FELIX: the new detector interface for the ATLAS experiment

*Nikolina Ilic*<sup>1,\*</sup>, *Jos Vermeulen*<sup>2</sup> and *Serguei Kolos*<sup>3</sup>, on behalf of the ATLAS TDAQ Collaboration

<sup>1</sup>Radboud University Nijmegen, Houtlaan 4, 6525 XZ Nijmegen, Netherlands

<sup>2</sup>Nikhef, Science Park 105, 1098 XG Amsterdam, Netherlands

<sup>3</sup>University of California Irvine, Irvine, CA 92697, USA

**Abstract.** During the next major shutdown from 2019-2020, the ATLAS experiment at the LHC at CERN will adopt the Front-End Link eXchange (FELIX) system as the interface between the data acquisition, detector control and TTC (Timing, Trigger and Control) systems and new or updated trigger and detector front-end electronics. FELIX will function as a router between custom serial links from front end ASICs and FPGAs to data collection and processing components via a commodity switched network. Links may aggregate many slower links or be a single high bandwidth link. FELIX will also forward the LHC bunch-crossing clock, fixed latency trigger accepts and resets received from the TTC system to front-end electronics. The FELIX system uses commodity server technology in combination with FPGA-based PCIe I/O cards. The FELIX servers will run a software routing platform serving data to network clients. Commodity servers connected to FELIX systems via the same network will run the new Software Readout Driver (SW ROD) infrastructure for event fragment building and buffering, with support for detector or trigger specific data processing, and will serve the data upon request to the ATLAS High-Level Trigger for Event Building and Selection. This proceeding will cover the design and status of FELIX and the SW ROD.

## 1 Introduction

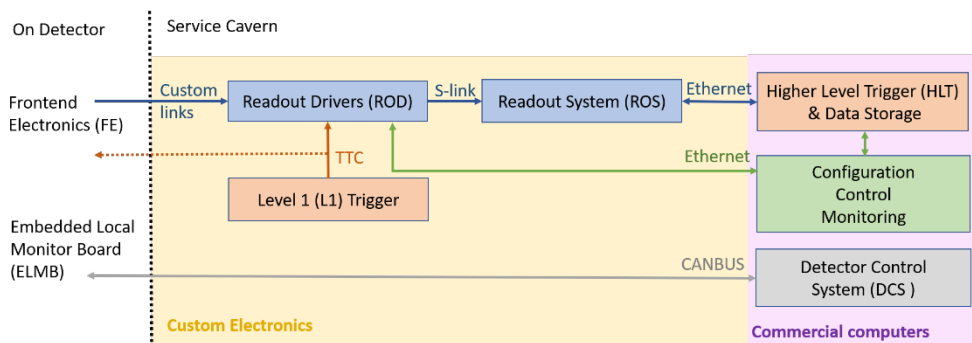
The ATLAS detector [1] is one of the two general purpose detectors located at the Large Hadron Collider (LHC) at CERN. The LHC collides bunches of particles at a rate of 40 MHz. Since ATLAS is capable of permanently storing data at a rate of 1.5 kHz, it uses a two-level trigger system to identify which events to retain. The hardware-based Level-1 (L1) trigger reduces the 40 MHz collision rate to a maximum of 100 kHz. The L1 trigger makes its decision on the basis of coarse sums of calorimeter energy deposits and of coarse transverse momentum measurements of muon candidates to identify events with interesting signatures. The software-based High-Level Trigger (HLT) further reduces the L1 rate to 1.5 kHz. The HLT decides which events to retain by using higher granularity calorimeter and muon detector information, together with inner detector tracking information.

\*Corresponding author: [Nikolina.Ilic@cern.ch](mailto:Nikolina.Ilic@cern.ch)

## 2 Planned Upgrades

The ATLAS detector will be upgraded during the second 2-year Long Shutdown (LS2) starting in 2019 [2]. An upgrade to the Liquid Argon (LAr) electromagnetic calorimeter electronics will allow for higher granularity energy measurements to be used by the L1 trigger [3]. The small wheels equipped with muon chambers, installed in the forward direction, will be replaced by New Small Wheels (NSWs) [4]. The NSW consists of eight layers of Micromegas detectors (MM) and small-strip Thin Gap Chambers (sTGC), which will provide more accurate and higher resolution tracking abilities in the forward regions of the ATLAS detector. The new muon chambers will also be able to handle the high hit rates associated with the high LHC luminosities that are expected after the Long Shutdown 2 (LS3) upgrades in 2024-2026 [5]. In addition to the NSW upgrades, additional resistive plate muon chambers, called BIS78, will be installed in the transition region between the ATLAS barrel and endcap. BIS78 will allow the L1 trigger to more efficiently reject signatures that could fake muons in this region. The calorimeter electronics of the L1 trigger system will also be upgraded to utilize the improved electromagnetic and muon measurements provided by the LS2 upgrades, resulting in improved selectivity and background rejection.

The new electromagnetic calorimeter electronics, the NSW muon detectors, BIS78, and L1 calorimeter trigger electronics will require a new infrastructure for reading out data and sending control information to the electronics. The current readout system is schematically shown in Figure 1. Once events are flagged as interesting by the L1 trigger, they are routed to custom Readout Drivers (RODs). The RODs process and check the data before sending it to the Readout System (ROS), where it is temporarily stored. The ROS consists of server PCs equipped with custom I/O based on PCIe expansion cards to receive the data from the RODs. The HLT, also consisting of server PCs, initially retrieves the data it needs for a trigger decision from the ROS via a commodity ethernet network connecting both systems. After the HLT decides which events to retain for permanent storage, the full event information is retrieved, reconstructed and stored. Software running on other commodity servers configures, controls, monitors detector electronics, and controls infrastructure through the Detector Control System (DCS). The current system contains many customized components and communication protocols (custom links, TTC [6], S-LINK [7], CANBUS) which reduce the flexibility with which information can be communicated to and from the detector.

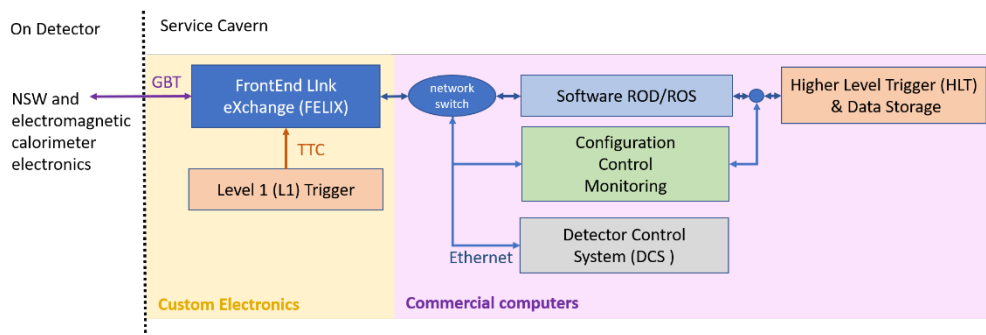


**Figure 1.** The current ATLAS Readout System. Custom designed Readout Drivers (ROD) and Readout System (ROS) forwards data from the front-end electronics of the detector to commercial computers running dedicated software. The dedicated software includes functionalities such as Data Storage, configuration, control and monitoring of the detector as well as High-Lever Trigger (HLT).

The Detector Control System (DCS) is software used for control and monitoring and receives information from on detector Embedded Local Monitoring Boards (ELMB) via dedicated CANBUS lines.

The new approach for the readout and control of the upgraded systems, shown schematically in Figure 2, allows for a reduction of the number of different types of custom electronics and protocols. The functionality that is currently provided by dedicated electronics and firmware will be moved to software running on commodity servers. This is made possible by advancements in computer technology and a radiation hard GigaBit Transceiver (GBT) link technology, developed by CERN [8]. The GBT link will be used to transmit the Timing, Trigger and Control (TTC) information to the front-end electronics. It will also be used to configure and control the front-end electronics through slow control. The protocol is bi-directional, enabling data to be passed from the front-end electronics to the data collection system and back up to the front-end. The multiplexing of relatively slow links (80, 160 or 320 Mb/s per link) by the standard GBT protocol makes it possible to transfer different independent data streams via the same physical link (raw transfer speed: 4.8 Gb/s). The slow links are known as E-links.

The Front-End Link eXchange (FELIX) will use the GBT protocol over standard optical links and transceivers to route data and commands (including TTC) between electronics on the ATLAS detector and commodity servers. FELIX will also use the GBT protocol to route control data (including TTC), to non-radiation-hard electronics installed in an underground service area beside the ATLAS detector (off-detector electronics). For transfers from off-detector electronics towards FELIX, a protocol called FULL Mode will be used and consists of simple 8B/10B encoding with a raw transfer speed of 9.6 Gb/s. In contrast to the pre-upgrade system, the functionality of the RODs and of the ROS will be implemented by software running on commodity servers. The computers containing the ROD/ROS functionality are referred to as Software RODs (SW RODs). The detector-specific data processing, which previously occurred on custom RODs, will be implemented in customizable Data Handler software applications on the SW RODs. The HLT, the system used for control/configuration/monitoring and DCS will continue running on commodity servers after the upgrade and will communicate via a commodity Ethernet network.



**Figure 2.** The upgraded ATLAS readout system. Data is routed by a custom designed FELIX board from the front-end detector electronics onto commercial computers. The commercial computers contain software to perform the data processing in the Software ROD/ROS, as well as detector control, monitoring, configuration, High-Level Trigger and Data Storage.

FELIX is planned to be detector agnostic and will act as a router between the custom links and a general commodity network switch. Due to this flexibility, the FELIX system is

planned to be used to read out all the ATLAS subdetectors after the LS3 upgrades. The largest LS3 upgrade project is the installation of a new tracker, called the Inner Tracker (ITk), whose links will be read out by FELIX. The ITk will consist of silicon strip and pixel trackers and cover a larger pseudorapidity range compared to what the current inner detector covers. ATLAS will also upgrade the readout electronics of the other subdetectors, and all the readout and a large fraction of the control will be done via FELIX. The legacy and new readouts will exist after LS2 and will be replaced by the FELIX only readout after LS3.

## 3 The FELIX System

### 3.1 Protocols

For the LS2 upgrades, the optical link connected to the front-end electronics is the Versatile link [1]. Links from FELIX towards detector or trigger system, called downlinks, use the standard GBT protocol. The type of E-link to be used (80, 160 or 320 Mb/s) and the type of protocol to be used for different E-links is configurable. Each E-link can be configured to forward either TTC signals or control data to the detector electronics received by FELIX via dedicated inputs of FELIX. Alternatively, it is possible to configure E-links for passing information sent to FELIX via the network and to select whether 8B/10B, HDLC or no encoding should be applied.

Optical links from the detector to FELIX, called uplinks, use the GBT protocol. For each E-link it is possible to configure whether 8B/10B, HDLC or no decoding is applied to the incoming data. For 8B/10B encoded event data transferred across E-links or across FULL mode links, control symbols indicate event fragment boundaries. Flow control is used for FULL mode links to indicate when the front-end system is temporarily overloaded. It is implemented with the help of control symbols.

The LS2 FELIX design will be expanded to meet additional requirements such as support of the upgraded TTC system, higher complexity of configuration of control signals, vertical and/or horizontal aggregation of data across links and handling a higher link density at rates up to 1 MHz. The uplink protocols supported for the LS3 system will include the lpGBT protocol and Aurora [4], in addition to the FULL mode and standard GBT protocols. For the downlinks towards the ITk detector, E-links of the lpGBT protocol will transfer the strip and pixel-specific coded data for control and trigger signaling.

### 3.2 Hardware

The FELIX system will consist of FPGA PCIe cards and Network Interface Cards (NICs) placed inside server PCs. The baseline choice of the FPGA card (Figure 3), known as the FLX-712, has a 16-lane PCIe Gen3 interface, a Xilinx Kintex Ultrascale XCKU115 FPGA, 8 miniPODs interfacing to 48 bi-directional optical links, a TTC interface and a LEMO connector for a "BUSY" signal output. The server PCs will run Linux and have a chassis that can be mounted onto racks located in the underground cavern beside the ATLAS detector. Either one or two FLX-712 cards will be installed per PC, each connecting to either 12, 24 or 48 links, depending on the type of front-end electronics to which the boards will form the interface.



**Figure 3.** The FELIX card (FLX-712).

For testing and development, lower priced, commercially available Xilinx VC709 boards are often used instead of the FLX-712. These boards are equipped with a separate small mezzanine, called the TTCfx, which interfaces to the TTC system. A VC709 interfaces to four bi-directional links and has an 8-lane PCIe Gen3 interface.

The baseline choice for LS2 FELIX PCs is to be determined. The current options are a 2U high single CPU server with an Intel Xeon E5-1660V4 (8 core, 3.2 GHz) CPU with 32 GB of memory, or alternatively with an Intel Xeon Gold 5118 (12 core, 2.3 GHz) CPU with 48 GB of memory. The PC includes a Network Interface Card (NIC). The current baseline choice for the software RODs is a 1U high server with 2 Intel Xeon Gold 5118 CPUs, 2 x 48 GB memory, a dual port Mellanox Connect X-5 100 GbE interface and a dual port 40 GbE interface for connecting to the HLT.

The number of FELIX links, cards and FELIX PCs is shown in Table 1 for the LS2 system. In Table 1, the L1 calorimeter trigger upgrades needed for reading out the upgraded on-detector LAr calorimeter electronics are referred to as “calorimeter off-detector trigger electronics”. The L1 trigger electronics upgrades needed for reading out the NSW are included under the “New Small Wheel muon detectors” section of Table 1.

**Table 1.** Number of FELIX links, cards and FELIX PCs that will be installed during LS2.

	# FE Links	# PCs	# of FLX-712 per PC	# of links per FLX-712	Link protocol
LAr Calorimeter on-detector electronics	620	7	2	48	GBT
Calorimeter off-detector trigger electronics uplinks/downlinks	295/71	25	1	12	FULL mode/ GBT
New Small Wheel muon detectors	1152	30	2	24	GBT
BIS78 muon detector	16	1	1	16	GBT

After LS3, it is expected that each FELIX card will be able to handle either 24 or 48 links, depending on the required speed and type of protocol. For the ITk Pixel detector 48 links are needed, which is expected to be achievable since the planned Aurora links will resemble FULL mode links which require fewer FPGA resources than GBT links. To estimate the number of FELIX PCs needed after LS3, it is assumed that 2 FELIX boards per PC will be installed, and PCIe Gen4 interfacing will be utilized. The PCIe Gen4 standard has been

released and offers double the bandwidth of the current PCIe Gen3. It is estimated that 17093 links from the detector, 545 FELIX cards, 279 FELIX PCs and 545 Data Handler PCs will be required. The latter implement the Software ROD functionality.

### 3.3 Software

The FelixCore application is a multi-threaded application running on the server PCs of the FELIX system that controls the FPGA cards via PCIe registers. These are accessible via memory mapped I/O, set up with the help of a dedicated driver. During transfer of data to the SW RODs, the data is buffered in circular buffers on the FELIX PC. The contiguous memory needed for the circular buffers to or from which data is transferred under DMA control is allocated by another driver.

FelixCore reconstructs complete chunks of the data written to the circular buffers and forwards these to any destination that is registered as receiving data with FelixCore. The data are forwarded via the network, using NetIO, a network protocol-agnostic software layer. FelixCore also forwards data, received from the network by means of NetIO, to the FPGAs.

## 4 Summary

The FELIX system is a detector interface for readout and control, encapsulating functionalities common to all ATLAS subdetectors. FELIX will be initially deployed for the LS2 upgrades, which include newer LAr calorimeter electronics, the New Small Wheel, additional BIS78 muon chambers, and upgraded L1 trigger calorimeter electronics. The FELIX system will be expanded to meet post-LS3 requirements and act as the interface with all subdetectors, including the new Inner Tracker (ITk).

## References

1. ATLAS Collaboration, The ATLAS Experiment at the CERN Large Hadron Collider, JINST 3 S08003 (2008).
2. ATLAS Collaboration, Letter of Intent for the Phase-I Upgrade of the ATLAS Experiment, CERN-LHCC-2011-012 (2011).
3. ATLAS Collaboration, ATLAS Liquid Argon Calorimeter Phase-I Upgrade Technical Design Report, CERN-LHCC-2013-017 (2013).
4. ATLAS Collaboration, New Small Wheel Technical Design Report, CERN-LHCC-2013-006 (2013).
5. ATLAS Collaboration, Technical Design Report for the Phase II Upgrade of the ATLAS Trigger and Data Acquisition System, CERN-LHCC-2017-020 (2017).
6. B. G. Taylor, Timing Distribution at the LHC, Proc. 8th Workshop on Electronics for LHC Experiments, CERN-2002-003 (2002).
7. H. C. van der Bij, R. A. McLaren, O. Boyle and G. Rubin, S-LINK, a data link interface specification for the LHC era, IEEE Trans. Nucl. Sci., **44** no. 3, 398-402 (1997).
8. P. Moreira. et al. The GBT, a Proposed Architecture for Multi-Gb/s Data Transmission in High Energy Physics, 10.5170/CERN-2007-007.332 (2007).