

FELIX: The New Approach for Interfacing to Front-end Electronics for the ATLAS Experiment

J. Anderson, A. Borga, H. Boterenbrood, H. Chen, K. Chen, G. Drake, M. Donszelmann, D. Francis, B. Gorini, D. Guest, F. Lanni, G. Lehmann Miotto, L. Levinson, J. Narevicius, A. Roich, S. Ryu, F. Schreuder, J. Schumacher, W. Vandelli, J. Vermeulen, W. Wu, J. Zhang

Abstract— From the ATLAS Phase-I upgrade and onward, new or upgraded detectors and trigger systems will be interfaced to the data acquisition, detector control and timing (TTC) systems by the Front-End Link eXchange (FELIX). FELIX is the core of the new ATLAS Trigger/DAQ architecture. Functioning as a router between custom serial links and a commodity network, FELIX is implemented by server PCs with commodity network interfaces and PCIe cards with large FPGAs and many high speed serial fiber transceivers. By separating data transport from data manipulation, the latter can be done by software in commodity servers attached to the network. Replacing traditional point-to-point links between Front-end components and the DAQ system by a switched network, FELIX provides scaling, flexibility, uniformity and upgradability and reduces the diversity of custom hardware solutions in favour of software.

THE data-acquisition, detector control and timing systems of the ATLAS experiment will be interfaced to new or upgraded detector and trigger systems via commodity

networking and a server PC based system: FELIX (Front-End Link eXchange). The FELIX system connects to on-detector radiation tolerant GBTx ASICs or to off-detector FPGAs via PCIe cards equipped with large FPGAs (e.g. Xilinx Ultrascale) and optical links. Fig. 1 shows an overview of the system, as it should be installed for the ATLAS Phase 1 upgrade during the next long shutdown (2019-2020) of the LHC. About 2000 optical links will be serviced by the Phase 1 system. After the third long shutdown of the LHC, foreseen for the mid-2020's, the FELIX approach will be used to interface to all detector and trigger systems. FELIX is designed to be able to easily evolve from this first generation, taking advantage of future developments of FPGAs, commodity computers and networks.

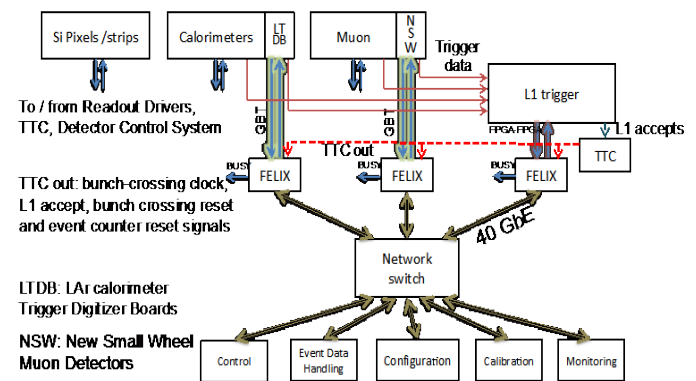


Fig. 1. The FELIX system and its connections, as it should be installed during the next long shut down of the LHC. The network type is indicated as the current 40 Gb/s Ethernet.

FELIX functions as a switch, agnostic to both the nature of the data it transfers and to the state of the data-acquisition system. The GBT ASIC aggregates up to 41 slow serial links (80, 160 or 320 Mb/s), called E-links, onto one 5 Gb/s serial link. FELIX allows packets transmitted on these E links (typically from Front-end ASICs) to be individually routed from/to different network end-points. Event data, slow control data, configuration, control and calibration commands can be routed according to FELIX's configurable routing tables. Most data streams will consist of 8b/10b encoded data, with command symbols indicating packet boundaries. FELIX FPGAs handling these data (from up to 32 links per FPGA) will decode the data streams, mark fragment boundaries and forward the data under DMA control via its PCIe interface. FELIX also supports the HDLC protocol needed by the GBT-

Manuscript received May 30 2016.

J. Anderson is with Argonne National Laboratory, Lemont, IL 60439, USA.

A. Borga is with Nikhef National Institute for Subatomic Physics / University of Amsterdam.

H. Boterenbrood is with Nikhef National Institute for Subatomic Physics / University of Amsterdam.

H. Chen is with Brookhaven National Laboratory, Brookhaven National Laboratory, P.O. Box 5000, Upton, NY 11973-5000, USA.

K. Chen is with Brookhaven National Laboratory, Brookhaven National Laboratory, P.O. Box 5000, Upton, NY 11973-5000, USA.

G. Drake is with Argonne National Laboratory, Lemont, IL 60439, USA.

M. Donszelmann is with Radboud University Nijmegen, Comeniuslaan 4, 6525 HP Nijmegen, Netherlands

D. Francis is with CERN, CH-1211 Geneva 23, Switzerland.

B. Gorini is with CERN, CH-1211 Geneva 23, Switzerland

D. Guest is with the University of California Irvine.

F. Lanni is with Brookhaven National Laboratory, Brookhaven National Laboratory, P.O. Box 5000, Upton, NY 11973-5000, USA.

G. Lehmann Miotto is with CERN, CH-1211 Geneva 23, Switzerland.

L. Levinson is with the Department of Particle Physics, The Weizmann Institute of Science, Rehovot 76100, Israel.

J. Narevicius is with the Department of Particle Physics, The Weizmann Institute of Science, Rehovot 76100, Israel.

A. Roich is with the Department of Particle Physics, The Weizmann Institute of Science, Rehovot 76100, Israel.

S. Ryu is with Argonne National Laboratory, Lemont, IL 60439, USA.

F. Schreuder is with Nikhef National Institute for Subatomic Physics / University of Amsterdam.

J. Schumacher is with CERN, CH-1211 Geneva 23, Switzerland and the Department of Computer Science, University of Paderborn, Germany

W. Vandelli is with CERN, CH-1211 Geneva 23, Switzerland.

J. Vermeulen is with Nikhef National Institute for Subatomic Physics / University of Amsterdam.

W. Wu is with Brookhaven National Laboratory, Brookhaven National Laboratory, P.O. Box 5000, Upton, NY 11973-5000, USA.

J. Zhang is with Argonne National Laboratory, Lemont, IL 60439, USA.

SCA, Slow Control Adapter, ASIC. Off-detector devices with FPGAs will typically connect to FELIX via FPGA-to-FPGA links with a single 8b/10b encoded data stream.

GBT E-links transfer with a fixed latency and low jitter. FELIX is therefore able to forward a stable LHC clock, trigger accept, reset and other beam-synchronous signals with fixed latency to its Front-end GBT clients. Other data transmitted via the Timing, Trigger and Control (TTC) system can also be forwarded. A new TTC system will be installed during the third long shutdown. Since the FELIX PCIe card has an interface to the TTC system, FELIX will make this upgrade transparent to the existing Front-ends.

A block diagram of the firmware is presented in Fig. 2. The firmware is, or will be, deployed on three platforms with a PCIe Gen3 interface: the Xilinx VC 709 evaluation board, the

Hitech-Global HTG 710 and a board developed in-house by BNL, the BNL 711 (with 4, 20 and 48 input links respectively). A block diagram of the BNL 711 is shown in Fig. 3. It is currently being tested.

All data received by FELIX will be forwarded to the commodity network, which today is 40 Gb/s Ethernet, but could be 100 Gb/s by Phase 1. Event data will be sent to server PCs where sub-detector or trigger electronics specific data handling and formatting will take place in software. In the current systems this is done by firmware in dedicated hardware (the ReadOut Drivers or RODs). The FELIX approach therefore introduces uniform interfacing. Functionality that is currently implemented in firmware in the RODs will be implemented in software, reducing hardware dependence as well as easing the development and maintenance effort.




Fig.2. Block diagram of the firmware for FELIX (FH: “From Host”, TH: “To Host”).




Fig. 3. Block diagram of the FELIX Phase 1 prototype, the PCIe Gen3 16x BNL-711