ATLAS Level-1 Topological Processor Project specifications

B. Bauß, U. Schäfer, E. Simioni

Institut für Physik
Johannes Gutenberg Universität
Mainz, Germany

Version 2.0

April 29, 2014

List of Figures

List of Tables

Abstract

In this document the design of the L1Topo module for the ATLAS Level-1 Phase-0 upgrade is described. This document is an update of the specifications for the L1Topo prototype: `http://www.staff.uni-mainz.de/uschaefe/browsable/L1Calo/Topo-Proto/topospecs.pdf`

The L1Topo project was presented so far in three reviews:

- L1Calo internal review, May 2012, `http://www.staff.uni-mainz.de/uschaefe/browsable/L1Calo/Topo/TopoReview.htm`
- PDR, Feb 2013, `https://edms.cern.ch/file/1225618/1`
- FDR, May 2013, `https://edms.cern.ch/document/1289737/1`

1 Introduction

This document describes the specification for the Level-1 topology processor module (L1Topo). The specification covers the processor main board as well as the internal and external interfaces.

- Sect. 1: Overview of the module.
- Sect. 2: Functional requirements.
- Sect. 3: Implementation in details.
- Sect. 4: Describe the algorithmic firmware and tests.
- Sect. 5: Internal and external interfaces.

1.1 Related Projects

L1Topo is a processor crate within the future Cluster, Jet and Muon Processor scheme of the ATLAS Level-1 trigger. L1Topo will from 2013/14 be located between the CMX and the CTP in the Level-1 architecture. The muon trigger will supply a small amount of signals into L1Topo in an initial phase. Additional connectivity will be available in upgrade phase 1. Propedeuticity: TTC [1], L1Calo modules [2], TTCDec [3], CTP [4], CMX [5], MUCTPI[6].

1.2 L1Topo processor board

The Topological Processor will be a single processor crate equipped with two (or more) processor modules. The processor modules will be identical copies, with firmware adapted to the specific topologies to operate on.
L1Topo will receive the topological output data of the sliding window processors from L1Calo and data from the L1Muon system. A preliminary version of the data formats transmitted into L1Topo is defined (see Appendix H). It comprise of TOB data (Trigger Object data) for jets, clusters and muons. The data will consist of a description of the position of an object (jet, e/m cluster, tau and muons) along with some qualifying information, like the energy sum within the object. Data are transmitted on optical fibres. After conversion to electrical representation, data are received and processed in

FPGAs equipped with on-chip Multi-Gigabit Transceivers (MGT). Results are sent to the Central Trigger processor (CTP). The L1Topo module will be designed in AdvancedTCA (ATCA form factor).

1.2.1 Real-time data path

ATCA Backplane zone 3 of L1Topo is used for real-time data transmission. The input data enter L1Topo optically through the backplane. The fibres are fed via four ~~to five~~ blind-mate backplane connectors that ~~can~~ carry ~~up~~ to ~~72~~ fibres each. ~~In the baseline design 48-way connectors will be used.~~ The optical signals are converted to electrical signals in 12-fibre receivers. For reason of design density miniPOD receivers ~~will be~~ used. The electrical high-speed signals are routed into two FPGAs, where they are de-serialized in MGT receivers; the parallel data are presented to the FPGA fabric. The two FPGAs operate on their input data independently and in parallel. High bandwidth, low latency parallel data paths allow for real-time communication between the two processors. The final results are transmitted towards the CTP on both optical fibres and electrical cables. The electrical signals are routed via an extension mezzanine module.

Fig.1 shows a conceptual drawing of L1Topo. The real-time processor FPGAs are labelled A and B. They are surrounded by the optical receivers. Non-real-time module control functionality is implemented on FPGA C, as well as on Mezzanine modules (extension module "X", IPMC module "I"). More detailed diagrams are shown in Sect.3.

Data reception The optical data arrive on the main board on 12-fibre ribbons. Since the backplane connectors support multiples of 12 fibres, the optical signals will be routed via "octopus" cable sets, splitting 48 fibres into groups of 12. It should be noted that un-armed, bare fibre bundles are very flexible and can easily be routed to their destination, even in high-density designs. However, they need to be protected from mechanical stress. The opto-electrical conversion will be performed in Avago miniPOD 12-channel devices. The opto receivers exhibit programmable pre-emphasis so as to allow for improvement on signal integrity for given track length.

After just a few centimeters of electrical trace length, the multi-gigabit signals are de-serialized in the processor FPGAs. They allow for programmable signal equalization on their inputs. The exact transmission protocol is defined (see Sect.5.2.1) The processors are supplied with required bias and

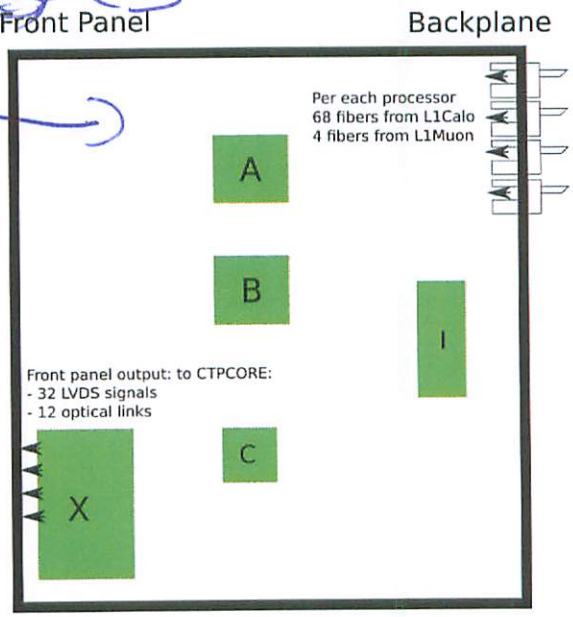

Figure 1: L1Topo module basic outline: 144 input fibers from the backplane, o/e converted to the A,B (processor FPGAs); C (control FPGA), I (IPMC mezzanine to interface with the ATCA crate), front panel output to the CTP (electrical and optical) and X (extension mezzanine).

termination voltages, as well as suitable reference clocks.

Data processing Topology data are processed in two FPGAs. There is no data duplication implemented at PCB level. Therefore two different schemes can be employed. The two processors can communicate via their fabric interface to get access to data that cannot be received directly via the multi-gigabit links. Though according to the device data sheets higher data rates should be possible, a maximum bit rate of 1Gb/s per differential pair is anticipated for the inter-FPGA link. That will limit parallel connectivity to 238 Gb/s of aggregate bandwidth. This would correspond to 24×238 bits per BX (5712 bits) which allow for sharing more than 250 generic trigger objects (TOBs). This is more than the outputs of all of the sort trees combined. Since this data path adds approximately one bunch tick of latency, input data is fanned out optically at the source, so that both processors are supplied with the same input data.

Due to the large amount of logic resources in the chosen FPGAs, a significant number of algorithms is expected to be run on the real-time data in parallel. The expected output to the CTP consists of individual bits indicating whether a specific topology algorithm passed or not plus an overflow bit. The resulting trigger data are expected to exhibit a rather small volume. They will be transmitted to the CTP optically or electrically. A single fibre-optical ribbon connection per processor FPGA, running through the front panel of the module is provided for this purpose. A mezzanine board will be required to interface L1Topo to the CTPCORE module electrically via 32 LVDS signals at low latency.

1.2.2 Clock distribution

The operation of the real-time data path requires low-jitter clocks throughout the system. For synchronous operation, data transmitters will have to be operated with clean multiples of the LHC bunch clock. Receiver reference clocks may as well be derived from local crystal oscillators, though tight limits on the frequency range will have to be observed. The L1Topo module will be designed for 40.0789 MHz operation of the real-time data path only.

The fabric clocks run at the LHC bunch clock frequency. The MGTs are clocked off multiples of the LHC clock. The jitter on the MGT bunch clock path is tightly controlled with help of a PLL device. The clock fan-out chips

chosen are devices with multiple signal level input compatibility, and output levels suitable for their respective use, either LVDSor CML

There are two independent clock trees for the fabric clocks into all FPGAs. There is one common MGT clock tree into all FPGAs supplying one of the two reference clock into each MGT quad. It carries the LHC based reference clock. Moreover all MGT quads receive an additional segmented crystal clock, to allow for use of a multitude of bit rates on the incoming MGT links. The control FPGA receives additional local clocks since it handles DAQ, ROI, and control links as well.

For backward compatibility a TTCDec module will be mounted on L1Topo. The current L1Calo modules receive their LHC bunch clock and associated data through such TTCDec module, based on a TTCRx chip. It will be is wired to the control FPGA. Spare optical inputs is available via a miniPOD receiver and four additional links wired to the backplane or the extension mezzanine for alternative optical LHC clock paths. *

1.2.3 Configuration, monitoring, and control

Pre-configuration access / ATCA compliant slow control L1Topo is a purely FPGA-based ATCA module. All communications channels are controlled by programmable logic devices and become functional only after successful device configuration. An initial step of module initialization is performed by an IPMC device. It communicates to the shelf via an I2C port (IPMB) available on all ATCA modules in zone 1. The prospective ATLAS standard (LAPP IPMC / mini-DIMM format) will be mounted on L1Topo. Additionally, the JTAG port in zone 1 will allow for limited access to the module at any time.

Embedded controller There is a IPMC based controller available on L1Topo.

FPGA configuration The baseline (legacy) FPGA configuration scheme on L1Topo is via a CompactFlash card and the Xilinx System ACE chip. This configuration scheme has been successfully employed on several L1Calo modules so far, including the JEM. On the JEM the required software and firmware has been devised to allow for in-situ update of the CompactFlash images. The required board-level connectivity will be available on L1Topo, to

* Phase-1 compatible clocking is accomplished via the ATCA fabric interface

⊕ Data into ATLAS data acquisition and 2nd level trigger is made through an S-Link compatible optical interface via an embedded ROD (firmware only) located on the control FPGA.

write the flash cards through a network connection to the FPGAs, once they are configured. In-situ configuration update will be possible via IP access to the control FPGA. ×
A local SPI memory will allow for an alternative configuration method for the control FPGA. The control FPGAs can be also configured off an SD flash memory, which in turn is sequenced by the control FPGA. This FPGA does not contain any algorithmic firmware and is therefore not expected to be updated particularly often.

Monitoring The default ATCA monitoring and control path is via I2C links in zone 1. The backplane I2C port (IPMB) is connected to the IPMC DIMM to provide monitoring information to the DCS system. On L1Topo configured FPGAs can be monitored for die temperature and internal supply voltage. This is achieved by the use of their built-in system monitor. These measurements will be complemented by board-level monitoring points connected via I2C.

Module control On standard ATCA modules, IP connectivity is mandatory for module control. This is generally provided by two 10/100/1000 Ethernet ports on the backplane in zone 2 (redundant base interface). The respective lines are wired to the extension module, where the required connectivity can ~~be~~ made ~~to the IPMC, or~~ to an MGT link on the control FPGA via an SGMII Phy device.

⊕ **DAQ and ROI** Since L1Topo would be used in the existent pre-phase-1 environment initially, a scheme for transmitting event data into the ATLAS data acquisition and 2nd level trigger was made compatible to the existent L1Calo modules. It is assumed that DAQ/ROI wise L1Topo would live in the L1Calo ecosystem and compatibility to L1Calo RODs is mandatory. An optical channel into DAQ and two to the ROI RODs is provided. The optical interface is made via a miniPOD device wired to the control FPGA (S-Link emulation). The existent hardware will allow for an entirely firmware based embedded ROD to be added [12].

× the SystemACE configurator is located on a separate, small mezzanine and can be replaced by an advanced configurator at a later stage.

2 Functional requirements

This section describes the functional requirements only. Specific rules need to be followed with respect to PCB design, component placement and decoupling. These requirements are detailed in Sect. 3 and Appendix A. For requirements on interfaces with other system components, see Sect.5.

2.1 Signal paths

L1Topo is designed for high speed differential signal transmission, both on external and internal interfaces. Two differential signalling standards are employed: LVDS and CML. For requirements regarding signal integrity see Appendix A.

2.2 Real-time data reception

Real-time data are received optically from the back of the ATCA shelf, converted to electrical representation, transmitted to the processors and deserialized in on-chip MGT circuits on the processors. The requirements with respect to data reception and conditioning are:

- Provide a minimum of four MPO/MTP compatible blind-mate fibre-optical backplane connectors in ATCA zone 3.
- Route bare fibre bundles to 12-channel opto receivers.
- Run the signal paths into the processors.
- Provide suitable coupling capacitors for multi-Gigabit links.
- Connect single ended CMOS level control lines to the control FPGA.
- Supply opto-electrical components with appropriately conditioned supply voltages.

2.3 Real-time data processing

The L1Topo processing resources are partitioned into two processors. This is a consequence of potential limitations on MGT link count and processing resources on currently available FPGAs. The requirements on processing

What do you mean with stages?

resources depend on the physics algorithms and are currently under investigation (see Sect.4).

The requirements with respect to the processors are:

- Provide an aggregate input bandwidth up to 208 GB/s (payload) into the processors:
  - 160 channels, 6.4 Gb/s baseline speed (capable of 13Gb/s) line rate.
- Process the real-time data in a 2-stage (maximum 3) processing scheme (data might need to cross chip boundary between the two processors).
- Minimize latency on chip-to-chip data transmission.
- Maximize bandwidth between the two processors:
  - Send bulk output data to CTP on MGT links, so as to maximize low-latency inter-processor communication bandwidth.
  - Use higher latency channels for non-real-time links where possible.
- Provide an aggregate bandwidth of up to 12 GB/s on MGT outputs towards the CTP:
  - The aggregate bandwidth to CTP corresponds to 12×128 bits per BX (1536 bits), at the baseline link speed (See. 5.2.1).
  - The low-latency LVDS links are 32 bits/BX at 40 MHz or 64 bits at 80 MHz.

2.4 Clock distribution

Both the FPGA fabric and the MGT links need to be supplied with suitable clock signals. Due to the synchronous, pipelined processing scheme most of the FPGA-internal clocks need to be derived from the LHC bunch clock or a device emulating it. Due to requirements on MGT reference clock frequency accuracy, a mixed 40.00/40.0789 MHz operation is impossible. Therefore a requirement for 40.0789 MHz operation has to be introduced.

The requirements with respect to the clock distribution on the main board are:

- Provide the FPGAs with clocks of multiples of either a 40.0789 MHz crystal clock, or the LHC bunch clock.
- Receive an optical TTC signal from the front panel.
- Recover TTC clock and data on a TTCdec mezzanine module.
- Allow for clock conditioning hardware in the clock path (jitter control, multiplication).
- The TTC-based clock must be used for reference on all real-time transmitters (CTP port).
- Provide the common TTC-based MGT clock to all FPGAs.
- Provide additional crystal-based MGT reference clocks to both processor FPGAs for use on the receiver sections. ~~Allow for segmentation of the clock trees to cope with possibly different line rates from various parts of Level-1.~~
- Connect the MGT clocks to the processor FPGAs such that all quads are supplied with two reference clocks.
- Provide two fabric clocks to all FPGAs (40.0789 MHz crystal, bunch clock).
- Provide a separate crystal based MGT clock to the control FPGA for use on the control link.
- Provide a separate crystal based MGT clock to the control FPGA for use on the DAQ and ROI link outputs (~~40.00 MHz or multiple~~).
- Provide a separate crystal based clock to the control FPGA for use by an embedded ROD (S-link compatible).
- Provide a separate crystal based 40.0789 MHz (or multiple) MGT receive reference clock to the control FPGA on the receive section of the ROI link, for use on future LHC bunch clock recovery circuitry (TTCDec replacement), and thus:
  - Allow for the input portion of the DAQ and ROI transceivers to be used for optical reception of LHC clock and data.

not sure about frequency. defined by S-Link requirement

2.5 Configuration and JTAG

JTAG is used for board level connectivity tests, pre-configuration access, and module configuration. During initial board tests and later hardware failure analysis, JTAG access will be required to connect an automatic boundary scan system, generally when the module is on a bench, not in an ATCA shelf. Also the initial configuration of non-volatile CPLDs will be performed at that stage.

The requirements with respect to boundary scan and CPLD configuration are:

- Allow for the connection of a boundary scan system to all scannable components of L1Topo: FPGAs, CPLDs, System ACE, and mezzanine modules via standard JTAG headers, following the standard rules on pull-up, series termination, and level translation between supply voltage domains.
- Allow for CPLD (re)configuration, following the boundary scan tests, and occasionally on the completed system.
- There is currently no requirement nor possibility known regarding integration of devices sourcing or sinking MGT signals externally, into the boundary scan scheme.

The requirements with respect to FPGA configuration are:

- Employ the standard System ACE configuration scheme to configure the FPGAs upon power-up.
- Connect the System ACE external JTAG port to the extension mezzanine for further routing.
- Allow for static control of the FPGA configuration port settings and read-back of the status via the control CPLD.
- Provide SPI configuration of the control FPGA according to the KC705 configuration scheme.
- Connect an SD card to the control FPGA.
- Connect the serial configuration lines of the processor FPGAs to user I/O of the control FPGA to allow for configuration off flash card data.

Is the configuration memory ?