

ATLAS Level-1Trigger

L1Topo Module

- V1.0 (Prototype) -

Project Specification

Version 1.0

Date: 10/01/2013 09:16:27

Bruno, Eduard, Uli

Universität Mainz

Table of Contents

1	Introduction	3
1.1	Related projects	3
1.2	L1Topo processor board	3
1.2.1	Real-time data path	3
	Data reception	4
	Data processing	4
1.2.2	Clock distribution	5
1.2.3	Configuration, monitoring, and control	5
	Pre-configuration access / ATCA compliant slow control	5
	Embedded controller	5
	FPGA configuration	6
	Environmental monitoring	6
	Module control	6
	DAQ and ROI	6
2	Functional requirements	6
2.1	Signal paths	6
2.2	Real-time data reception	7
2.3	Real-time data processing	7
2.4	Clock distribution	7
2.5	Configuration and JTAG	8
2.6	Module control	8
2.7	DAQ and ROI	9
2.8	Extension mezzanine (X)	9
2.9	Zynq module (Z)	9
2.10	IPMC dimm	10
2.11	ATCA power brick	10
3	Implementation	10
3.1	Scalable design	11
3.2	Clock	11
3.3	Control	12
	Module control	13
3.4	Board level issues: signal integrity, power supplies, and line impedances	14
4	Firmware, on-line software and tests	14
4.1	Service- and infrastructure firmware	15
4.2	Algorithmic firmware	15
5	Interfaces : connectors, pinouts, data formats	16
5.1	Internal interfaces	16
5.2	Front panel	16
5.3	Backplane connector layout	16
5.4	Interfaces to external systems	16
5.5	Data formats	18
6	Appendix	18
6.1	Checklist for detailed design	18
6.2	Glossary	20
6.3	Change log	21

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 mezzanine modules. Section 1 of this document gives an overview of the module. Section 2 describes the functional requirements. Section 3 details the implementation. Section 5 summarizes the external interfaces.

1.1 Related projects

L1Topo is a major processor module 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.

TTC	http://www.cern.ch/TTC/intro.html
L1Calo modules	http://hepwww.rl.ac.uk/Atlas-L1/Modules/Modules.html
TTCDec	http://hepwww.rl.ac.uk/Atlas-L1/Modules/Components.html
CTP	http://atlas.web.cern.ch/Atlas/GROUPS/DAQTRIG/LEVEL1/ctpttc/L1CTP.html
CMX	http://ermoline.web.cern.ch/ermoline/CMX
Muon Trigger	http://

1.2 L1Topo processor board

The Topological Processor will be a single processor crate equipped with one or several 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. The details of the data formats are not yet finally defined. However, the information transmitted into L1Topo will be basically comprised of the ROI data that are currently transmitted to the 2nd level trigger. The data will consist of a description of the position of an object (jet, e/m cluster, and tau) along with some qualifying information, basically the energy sum within the object. Preliminary data formats have been devised. Data are going to be 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-serialised 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.

Figure 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, Zynq module Z, IPMC module I). More detailed diagrams are shown in section 3.

ATLAS Level-1 Trigger




Figure 1: L1Topo module

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 24, 36, 48, or 72 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 centimetres of electrical trace length, the multi-gigabit signals are de-serialised in the processor FPGAs. They allow for programmable signal equalization on their inputs. The exact data formats are as yet undefined, though standard 8b/10b encoding is envisaged for purpose of run length limitation and DC balance. The processors are supplied with required bias and 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. Since this data path adds approximately one bunch tick of latency, it might be more attractive to fan out data electrically or optically at the source, should both processors need to be supplied with the same 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 resulting trigger data are expected to exhibit a rather small volume. They will be transmitted to the CTP optically.

ATLAS Level-1 Trigger

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 40.08 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 LVDS or CML.

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

The current L1Calo modules receive their LHC bunch clock and associated data through the TTCDec module, based on a TTCRx chip. Future LHC bunch clock distribution might differ from this scheme. L1Topo will have to operate with the existent clock distribution scheme initially. At a later stage it will probably be driven in a different way. For backward compatibility a TTCDec module will be mounted on L1Topo. It will be wired to the control FPGA. Alternative optical LHC clock paths will be possible via the extension mezzanine module. The spare optical inputs of DAQ and ROI SFP modules (see below) are available to route LHC clock and data into L1Topo via the extension mezzanine.

1.2.3 Configuration, monitoring, and control

While the L1Topo prototype is meant to be very close to future production modules, there are still optional schemes required on the prototype, regarding configuration and module control. Part of the circuitry described below will be removed from the module before final production.

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 (see below). 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. The signals are routed to the FPGA components via the extension module.

Embedded controller

Due to uncertainties about the IPMC standardization and the usability of the IPMC for module control, there is a Zynq based embedded controller available on L1Topo. A commercial mezzanine module (MARS ZX3) will be used on the prototype. The option of a Zynq based controller had been recommended during the PDR. The module will boot off an SD flash card.

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 write the flash cards through some network connection to the FPGAs, once they are configured.

A local SPI memory will allow for an alternative configuration method for the control FPGA. This FPGA does not contain any algorithmic firmware and is therefore not expected to be updated particularly often. The processor FPGAs will be configured off an SD flash memory, which in turn is sequenced by the control FPGA. In-situ configuration update will be possible via IP access to the control FPGA. When the MARS module is mounted on L1Topo, it will configure the FPGAs from the flash card contents.

Environmental monitoring

Board level temperature and voltage monitoring will be a requirement. The current L1Calo modules are monitored via CAN bus. 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. 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. A CAN microcontroller might be added on a future version of the extension module, if required.

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, the MARS module, or to an MGT link on the control FPGA via an SGMII Phy device.

A dedicated SFP based control link into the control FPGA will transport module control signals initially. It will be connected into an optical VME bus extender, so as to allow controlling L1Topo from a VME crate CPU within the current L1 online software empire.

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. A single optical channel into each the DAQ and ROI RODs is provided. The optical interface is made via SFP sockets. There is additional optical connectivity available on the control FPGA via a miniPOD device. The existent hardware will allow for an entirely firmware based embedded ROD to be added.

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 sections 3 and 6. For requirements on interfaces with other system components, see section 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 sect. 6.

2.2 Real-time data reception

Real-time data are received optically from the back of the ATCA shelf; they are converted to electrical representation, transmitted to the processors and de-serialised in on-chip MGT circuits.

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
- Supply opto-electrical components with appropriately conditioned supply voltages
- Connect single ended CMOS level control lines to the control FPGA
- Provide suitable coupling capacitors for multi-Gigabit links
- Run the signal paths into the processors

2.3 Real-time data processing

The L1Topo processing resources are partitioned into two processors. This is a consequence of limitations on MGT link count and processing resources on currently available FPGAs. The requirements on processing resources depend on the physics algorithms and are not currently known.

The requirements with respect to the processors are:

- Provide an aggregate input bandwidth of 160 GB/s (payload) into the processors
 - 160 channels of up to 10 Gb/s line rate
- Process the real-time data in a 2-stage (maximum) processing scheme (data might need to cross chip boundary between the two processors)
- Allow for footprint compatible medium- and high-end FPGAs for scalability of processing resources
- Minimise latency on chip-to-chip data transmission
- Maximise bandwidth between the two processors
 - Send bulk output data to CTP on MGT links, so as to maximise low-latency inter-processor communication bandwidth
 - Use higher latency channels for non-real-time links where possible
- Provide an aggregate bandwidth of up to 24 GB/s on MGT outputs towards the CTP
 - 24 channels, 6.4 Gb/s (capable of 10Gb/s) line rate.
 - Additional 32-channel low latency electrical (LVDS), capable of 160 Mb/s

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 40.00/40.08 MHz operation is impossible. Therefore a requirement for 40.08 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 40.08MHz 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.

ATLAS Level-1 Trigger

- Connect the MGT clocks to the processor FPGAs such that the central quad of 3 quads is supplied. This rule must be followed such that the requirements are met, whether a smaller or larger device is mounted.
- Provide two fabric clocks to all FPGAs (40.08 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 40.08 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 via the extension mezzanine

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 both the control FPGA and the Zynq based controller
- 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.

2.6 Module control

On ATCA modules serial protocols are used to achieve board level configuration and control. Typically Ethernet would be used to control module operation. On L1Topo any serial protocol compatible to Xilinx MGT devices can be used to communicate to the control FPGA, once it has been configured. All board level control goes via the control FPGA. The control FPGA is in turn connected to a Zynq based processor (MARS bar) and an ATLAS standard IPMC.

The requirements with respect to general board control are:

- Provide an optical bidirectional channel from the front panel to a control FPGA MGT (control link)
- Provide eight-lane access from zone 2 to the extension mezzanine, compatible to 10/100/1000 Ethernet, so as to allow for an Ethernet Phy to be mounted on the extension mezzanine module

ATLAS Level-1 Trigger

- Provide for Ethernet connectivity onwards to IPMC and MARS
- Provide two-lane access from the mezzanine on to the control FPGA (one MGT, bidirectional)
- Provide bi-directional connectivity between processors and control FPGA
- Provide a control bus from the control FPGA to all opto-electrical transceivers (I2C and static control)
- Provide a single ended and a differential control bus from the control FPGA to the mezzanine module
- Provide an interconnect between control FPGA and control CPLD (via extension module)

The CPLD is in charge of mainly static controls that need to be asserted at power-up, before the FPGAs are configured.

The requirements with respect to the CPLD are:

- Communicate to the general board control system via a bus to the control FPGA (on the extension module).
- Communicate to the System ACE sub-system so as to control FPGA configuration and allow for in-situ update of the CompactFlash card
- Control the static FPGA configuration control lines

2.7 DAQ and ROI

A single fibre for each DAQ and ROI transmission will be provided on L1Topo.

The requirements with respect to DAQ and ROI interface are:

- Provide an optical ROI output channel to the front panel
- Provide an optical DAQ output channel to the front panel
- Use standard SFP modules
- Connect them to MGTs on the control FPGA
- Provide a separate 40 MHz (or multiple) clock to the MGT quads driving DAQ and ROI fibres
- Provide up to a full miniPOD worth of bandwidth for use with an embedded ROD core

2.8 Extension mezzanine (X)

The extension mezzanine provides some connectivity and real estate for electrical real-time signalling and general control purposes. The requirements with respect to auxiliary controls on the mezzanine board are:

- Breakout and signal conditioning for electrical trigger signals (32 lanes to CTP)
- Handle optical and electrical input of LHC bunch clock (TTC and future protocols)
- Route TTC clock towards TTCdec module
- Receive two 4-pair Ethernet signals from backplane zone 2 (base interface)
- Connect to two Ethernet ports on the IPMC
- Connect the mezzanine to the control FPGA via a single MGT for purpose of module control
- Connect the mezzanine to the control FPGA via an LVDS level bus
- Connect the mezzanine to the control FPGA via a CMOS bus for purpose of static and slow controls
- Connect the mezzanine to the control CPLD (16 lines)
- Connect the Zynq module via 32-way CMOS level control bus
- Drive the three MiniPOD control buses (I2C)

2.9 Zynq module (Z)

The Zynq mezzanine is a commercial module (MARS ZX-3)

The requirements with respect to the mezzanine board are:

ATLAS Level-1 Trigger

- Wire the module to Ethernet via the extension module
- Connect the module to an SD flash card acting as processor boot device and configuration storage
- Connect the parallel buses to the control FPGA for purpose of module control

2.10 IPMC dimm

A “standard” ATLAS IPMC controller in mini-dimm format is used.

The requirements with respect to the IPMC mezzanine board are:

- Connect to the ATCA zone-1 control lines
- Connect to the zone-1 JTAG pins via a cable loop
- Connect to Ethernet via the extension module
- Connect power control / alarm on the ATCA power brick
- Supply the IPMC with the management voltage generated on the power bridge

2.11 ATCA power brick

A “standard” ATCA power brick is to be used to simplify ATCA compliant power handling.

The requirements with respect to the power brick are:

- 48/12V power brick
- ATCA power control / monitoring on-brick
- Internal handling of hot swap / in-rush control
- Internal management voltage generation

3 Implementation

This section describes some implementation details of the L1Topo module. This section will be expanded before module final production.

L1Topo is built in ATCA form factor. The main board is built as a 20 layer PCB. The PCB is approximately 2mm thick and should fit ATCA standard rails. If required, the board edges would be milled down to the required thickness. A detailed block diagram of the real-time data path is shown in Figure 2.

ATLAS Level-1 Trigger

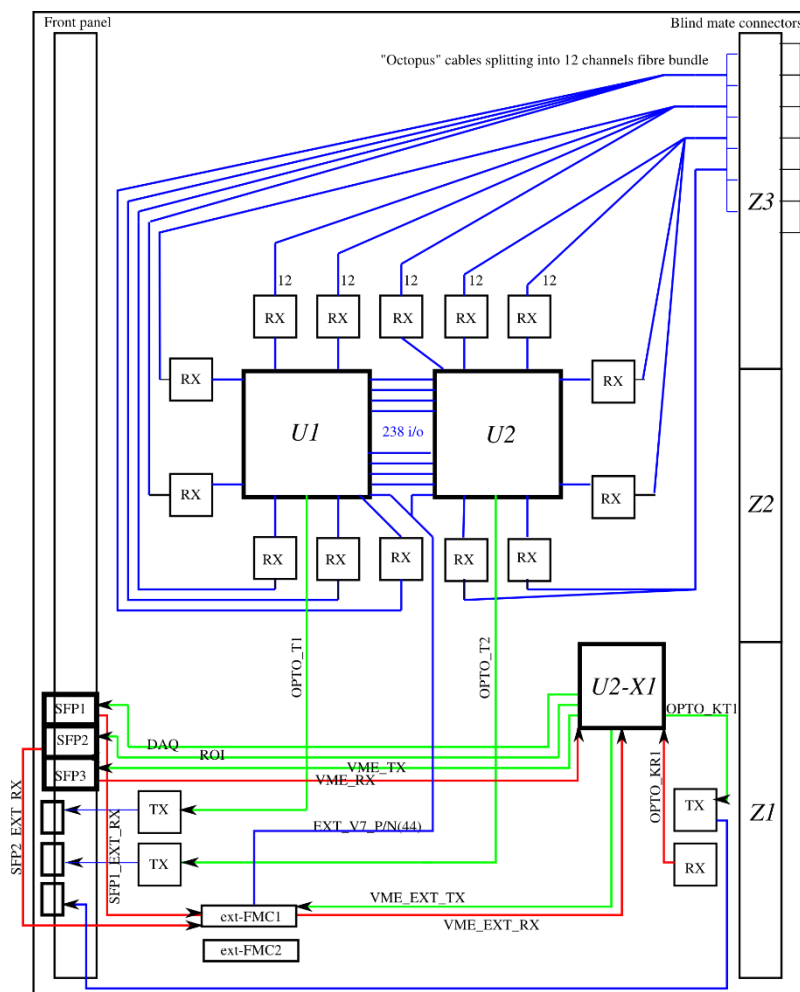


Figure 2 – L1Topo real-time data path and MGT links

3.1 Scalable design

L1Topo can be assembled with a choice of footprint compatible devices. Initially XC7VX485T will be mounted due to component availability. The most powerful device XC7VX690T will be mounted on later copies of the module. L1Topo needs to be designed such that all vital signals are run on signal pins available on the smaller device. The extension mezzanine module socket will provide some spare connectivity to make sure that L1Topo functionality can be upgraded at a later stage by replacing the relatively inexpensive mezzanine module only.

The back panel optical connectors provide a capacity of 360 fibres maximum, if five shrouds are mounted. At the envisaged density of 48 fibres per connector the 160 input channels will populate four connectors. The back panel connection scheme has been chosen to simplify maintenance on the module. However, in case of problems with the high density fibre connectors, fibre bundles could alternatively be routed to the front panel.

3.2 Clock

The clock circuitry comprises various crystal clocks and a TTCdec clock mezzanine module for clock generation, a jitter cleaner for signal conditioning, and several stages of electrical fan-out. Various Micrel device types are used to drive and fan out clocks of LVDS and CML type at low distortion. All Micrel devices are sink terminated on-chip. The jitter cleaner used on L1Topo is a Silicon Labs [5326](#). It allows for jitter cleaning and frequency synthesis up to multiples of the bunch clock frequency.

A detailed block diagram of the clock path is shown in Figure 3.

ATLAS Level-1 Trigger

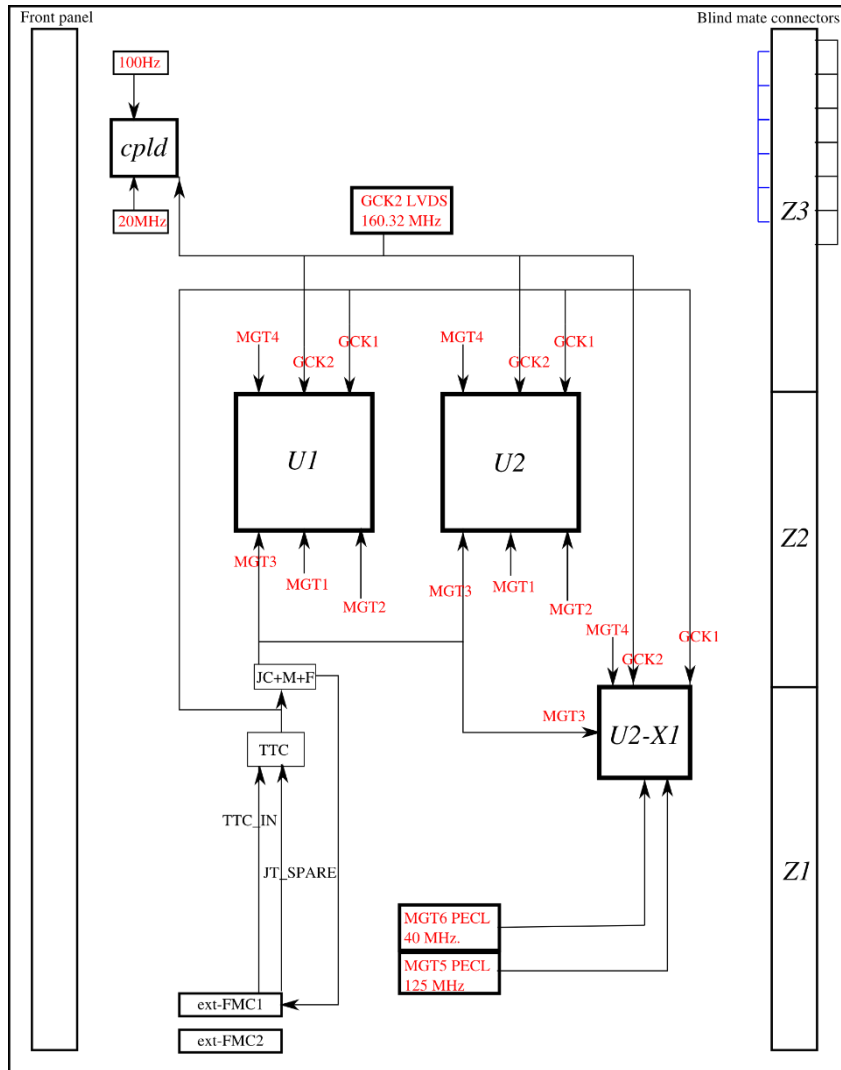


Figure 3 – L1Topo clock distribution

3.3 Control

A detailed block diagram of control paths is shown in Figure 4.

ATLAS Level-1 Trigger

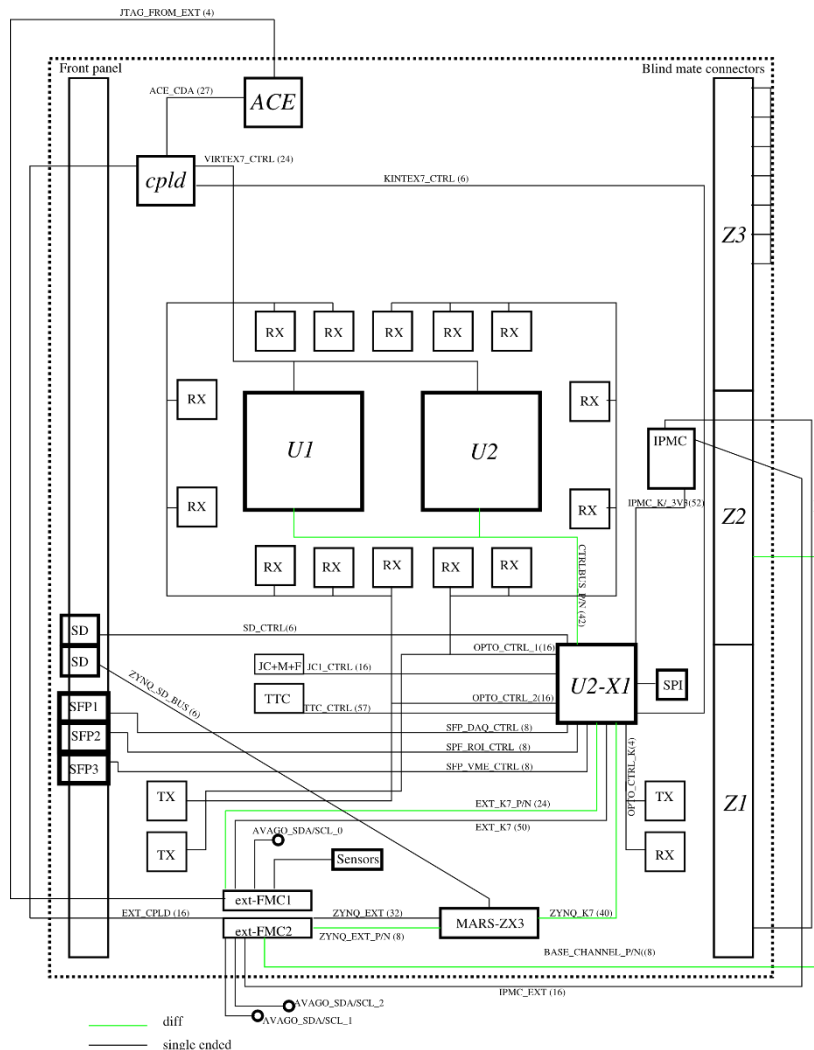


Figure 4 – L1Topo control paths

Module control

L1Topo module control is initially done via a serially extended VME bus, as outlined above. This choice was made since such an operation will not require any driver software to be written. The module is seen from a VME crate CPU as if it were a local VME module.

As soon as the required software environment is available, this access mode will be converted to standard Ethernet access. The required hardware components are available on L1Topo: The optical (SFP) link can be directly connected to any 1000BASE-SX Ethernet port. Also, there is the required space and connectivity available on the extension module to connect the control FPGA to the electrical Ethernet port located on the backplane, via an SGMII Phy (M88E1111 or similar).

The use of Ethernet for module control has been extensively explored in a previous Mainz project. Both UDP and bare Ethernet schemes have been used. A similar scheme devised by the University of Bristol (IPbus) is currently being considered by the L1Calo group for ATCA based control on future L1Calo modules. The hardware provided on L1Topo is compatible to that scheme and IPbus will be adopted if L1Calo decides in favour of that approach. Specific firmware and software is required to that end (see below).

3.4 Board level issues: signal integrity, power supplies, and line impedances

L1Topo is a large, high-density module, carrying large numbers of high-speed signal lines of various signal levels. The module relies on single-ended CMOS (3.3V), and differential (some or all of LVDS, PECL2.5, CML3.3, and CML2.5) signalling. System noise and signal integrity are crucial factors for successful operation of the module. Noise on operating voltages has to be tightly controlled. To that end, large numbers of decoupling capacitors are required near all active components. FPGAs are particularly prone to generating noise on the supply lines. Their internal SERDES circuitry is highly susceptible to noisy operating voltages, which tend to corrupt their high-speed input signals and compromise the operation of the on-chip PLL, resulting in increased bit error rates. To suppress all spectral components of the supply noise, a combination of distributed capacitance (power planes) and discrete capacitors in the range of nF to hundreds of μ F are required. On the FPGAs there are capacitors included in the packages for decoupling of noise up to highest frequencies.

Unlike previous modules, L1Topo will receive its MGT supply voltages from switching regulators. Current drawn on the MGT supply lines was considered too high for linear regulators, if all links are operated at full speed. The converters are equivalent to the ones used on Xilinx evaluation modules, and the ripple, according to the specifications, is well below the specified limits for the MGTs.

L1Topo base frequency is 40.08 MHz. Parallel differential I/O operates at multiples up to 1Gb/s. Multi-Gigabit (MGT) links operate at 6.4 or 10Gb/s. This is only possible on matched-impedance lines. Differential sink termination is used throughout. All FPGA inputs are internally terminated to 100 Ω or to 50 Ω ||50 Ω , according to the manufacturer's guidelines. All lines carrying clock signals must be treated with particular care. In the appendix there is a checklist for the detailed module design.

4 Firmware, on-line software and tests

L1Topo is an entirely FPGA based module. For both hardware commissioning and operation a set of matching firmware and software will be required. These two phases are well separated and requirements differ considerably. Hardware debug and commissioning will require intimate knowledge of the hardware components and will therefore be in the hands of the hardware designers. Both firmware and software will be restricted to simple, non-OO code. Hardware language is VHDL, software is plain C. GUI based tools are not required and will not be supplied by the hardware designers. Module commissioning from a hardware perspective is considered complete once the external hardware interfaces, board level connectivity, and basic operation of the hardware devices have been verified. The hardware debug and commissioning will involve JTAG/boundary scan, IBERT/ChipScope tests, and firmware/software controlled playback/spy tests with any data source/sink device available. Initially the GOLD will be available for test vector playback on a small number of 12-fibre ports. At a later stage a CMX prototype module will be used as a data source.

Module control is initially via an opto fibre connection to a VMEbus module carrying SerDes devices and optical transceivers. The opto-electrical interface will be a SFP module with LC-type opto connectors. Control software will be based upon calls to the CERN VME library. SFP outputs provided for DAQ and ROI data transmission will be tested for bit errors and link stability only. No specific data formats will be tested on these links.

In parallel to the hardware debug and commissioning, higher level software and firmware will be developed for later operation of L1Topo. As far as control software is concerned, a board level control scheme for future trigger modules needs to be devised. This will most likely not follow the initial approach of VME serialisation. A likely candidate for module control is the IPbus firmware (VHDL based UDP stack) and software suite developed at the University of Bristol. Alternatively the MARS module can be used to communicate to the controller PC via TCP protocol. It will be possible to re-use existent firmware and software to a large extent.

ATLAS Level-1 Trigger

The test environment available in the home lab will allow for simple real-time data path tests only. There is no hardware, software installation, nor expertise available to run any tests involving DAQ/RODs/ROSeS. Therefore all system level tests will have to be done in an appropriately equipped L1(Calo) test lab. Currently the L1Calo CERN test rig seems to be the only available option, before the L1Topo prototype can be connected up at point 1.

4.1 Service- and infrastructure firmware

The correct operation of high-speed (MGT) links and FPGA-internal clock conditioning circuitry is prerequisite for any “physics” firmware running on the recovered data streams. This service- and infrastructure firmware is closely tied to the device family used. To a large extent the code is generated with IP core generator tools.

The bare MGT link instantiations are complemented with code for link error detection and monitoring, as well as playback and spy circuitry for diagnostics purposes. These functions are controlled by the crate controller CPU, initially via VMEbus serialisation, as described above.

Real-time data are copied to the DAQ via ROD links upon reception of a Level-1 trigger. These data offer a further means of monitoring the correct functionality of L1Topo. There exists an additional data path to the LVL-2 ROI builders for the purpose of region-of-interest data transmission, if required.

All infrastructure described above has been successfully implemented and operated for all L1Calo processor modules. The detailed description is found in the collection of documents referenced in section 1. All functionality will be re-implemented on L1Topo. The JEM VHDL code is a purely behavioural description and can therefore be re-used with just minor modifications. Some optimization for the chosen FPGA devices might be desirable, but isn't strictly required.

4.2 Algorithmic firmware

There are several classes of physics algorithms currently being explored with help of physics simulations. Most algorithms consist of identification of energetic objects and subsequent cuts on angular distribution. This class of algorithms can be most efficiently implemented as a comparison of incoming objects in a highly parallelised sort tree, followed by a look-up table or a set of comparators, yielding a very small number of result bits which are transmitted to the CTP. The algorithms need to be pipelined at the LHC bunch crossing rate. The pipeline depth must be kept to the minimum to match latency requirements.

A $d\phi$ cut between the two most energetic jets of 128 has been coded as a reference implementation. In a first example it has been targeted to a Virtex-6 device on the GOLD demonstrator. Preliminary results suggest a latency of the parallel sort tree alone of about one LHC bunch tick. The angular cut in the lookup-table will take less than half a bunch tick. The reference implementation takes about 14 % of a Virtex XC6VLX240T device¹.

An initial implementation of a 240-wide sort tree for determination of the two largest objects on a Virtex XC7VX690T yields a latency of approximately one bunch tick. This is a preliminary result in a standalone implementation. The sort tree has not yet been tied to real MGT input circuits. This sort tree takes about 6 % of the logic resources available on the target device.

Further algorithms will be coded and simulated in several participating institutes. There will be a separate document on software and on algorithmic firmware produced at a later date.

¹ An FPGA based Topological Processor prototype for the ATLAS Level-1 Trigger upgrade, B. Bauss et al. 2012 JINST 7 C12007 doi:10.1088/1748-0221/7/12/C12007

5 Interfaces : connectors, pinouts, data formats

5.1 Internal interfaces

The L1Topo module carries four daughter modules: an IPMC module, a TTCdec card, a MARS bar equipped with a Zynq processor, and an extension mezzanine.

The extension card provides 44 differential real-time signal pairs and additional control connectivity. It is connected via two 400-way Samtec connectors (SEAM on the mezzanine module, SEAF on L1Topo).

The Zynq based processor is a commercial module (Enclustra MARS ZX3). Documentation is found on the manufacturer's website.

The TTCdec clock mezzanine module is an L1Calo standard card. It is connected via two 60-pin connectors (Samtec Q-Strip).

The IPMC module is a 244-pin mini-dimm, mechanically compatible to DDR3 memory modules. Documentation, including pinout is available from LAPP, Annecy.

5.2 Front panel

The front panel shows is carrying connectivity for both real-time output and non real-time control:

- Standard ATCA handle switches and LEDs (from IPMC)
- Compact Flash card slot for SystemACE
- SD card slot
- 3 MPO connectors (from MiniPODs, to CTP and embedded ROD links)
- 3 SFP connectors (module control and DAQ/ROI G-Links)
- 1 TTC optical AMP input
- 1 TTC differential lemo input
- 1 electrical (SCSI VHDCI) connector (to CTP)
- 1 Lemo-00 (Busy from embedded ROD)

5.3 Backplane connector layout

The backplane connector is made to standard ATCA layout in zones 1 and 2. Zone 3 is populated with four MTP/MPO connectors that connect onto a RTM with hermaphroditic blind-mate shrouds (MTP-CPI).

5.4 Interfaces to external systems

L1Topo interfaces to external systems mainly via optical fibres. The physical layer is defined by the fibre-optical components used. On the real-time data path fibre-bundles with multiples of 12 multimode fibres are used. The external connections are made by MTP/MPO connectors. L1Topo is equipped with male MTP connectors.

The calorimeter and muon trigger processors are connected via 48-way fibre assemblies, for the CTP outputs 12-way assemblies will be used. Opto-electrical translation on the real-time path is made with miniPOD 12-channel transceivers. Ideally the far end transceivers should be of same type as on L1Topo. Thus the optical power budget can be maximised. For short fibres and in absence of any optical splitting, other transceiver types are assumed to be compatible. However, optical power budget should be checked carefully.

Data rates and formats are defined by the FPGAs used for serialization and deserialization. For the L1Topo prototype 6.4 Gb/s or 10 Gb/s input data rates will be supported. For the production module

ATLAS Level-1 Trigger

any rate up to 13 Gb/s will be supported by the processor FPGAs, though maximum data rate will depend on the exact device type being purchased in 2013/14. At current the miniPOD transceivers will limit the data rate to a maximum of 10 Gb/s. A possible road map to higher data rates is not known. For data output to the CTP 6.4 Gb/s is used.

All real-time data streams are used with standard 8b/10b encoding. The data links are assumed to be active at all time. There are neither idle frames nor any packetizing of data. For reason of link stability and link recovery an option of transmitting zero value data words encoded into comma characters is being considered. This might also simplify initial link start-up. A first attempt has been made to define an encoding scheme (see <https://indico.cern.ch/getFile.py/access?contribId=5&resId=0&materialId=slides&confId=215417>).

The baseline deserialization scheme is relying on a de-multiplexing the 6.4Gb/s data stream to 8 sub-ticks (word slices) of 16 bits wide, transmitted on a 320MHz clock. This de-multiplex factor will most likely yield lowest latency. It should be noted, however, that alternatively the data could be de-multiplexed into four sub-ticks of 32 bits wide. At marginally higher latency this scheme might simplify further de-multiplexing to the bunch clock rate. It is anticipated that data will need to be fully de-multiplexed to the LHC base rate, since the algorithms require correlating all input data from all incoming sub-ticks. Further information might have to be transmitted to identify sub-ticks of the 40.08 MHz LHC bunch ticks. The de-multiplexing and word slice alignment scheme is a firmware option and will not affect the design of the interfacing processors. Compatibility to the data sources will be guaranteed.

All non-real-time external links are SFP. Data rates and encoding scheme need to be kept within the capabilities of the control FPGA (Kintex-7, up to 6.6 Gb/s). Since the use of L1Calo RODs is envisaged in an initial phase from 2013/14, the DAQ and Level-2 links need to be compatible to the legacy data formats (G-Link). Firmware has been devised at Stockholm University to generate compliant signals in Xilinx MGTs. The optical control link is assumed to be run at 1.25 Gb/s line rate and will therefore be compatible to Ethernet (SGMII). Two spare optical inputs (up to 6.6 Gb/s) will be available on the DAQ/Lvl-2 SFP devices. They will be routed such that the LHC bunch clock can be received that way. It should be noted that the optical wavelength of those devices is not compatible to the current TTC system and an external converter needs to be employed. Additional miniPOD based connectivity is available for DAQ data, using an embedded ROD. This will be a firmware plugin on the control FPGA. For purpose of busy handling a Lemo socket will be made available via the extension mezzanine.

A limited amount (32 differential lanes) of electrical I/O is available from the processor FPGAs. These data lines are reserved for the most latency critical paths to the CTP. The signals are routed via a mezzanine module, so as to allow for signal conditioning. The maximum data rate on the electrical path will depend on the signal conditioning scheme chosen, and the cable length. The FPGAs themselves would allow for more than 1 Gb/s per signal pair. The electrical port width is up to 22 signal pairs per processor FPGA. A VHDCI style connector has been suggested for the electrical port. Connector mechanics are affecting the mezzanine module only and are not within the scope of this review.

I/O	From/to	bandwidth		
Real-time input	various	160 * up to 13 Gb/s	Opto fibre / MTP 48	miniPOD, 8b/10b
Real-time output	CTP	24 * up to 13 Gb/s	Opto fibre / MTP	miniPOD, 8b/10b
Real-time electrical	CTP	32 * up to 160Mb/s	via mezzanine	LVDS
Control		2*1Gb/s	Ethernet electrical, zone 2	
Control		1Gb/s	SFP	
DAQ	D-ROD	~1Gb/s	SFP	
ROI	R-ROD	~1Gb/s	SFP	
IPMB			Zone 1	
LHC clock	TTC etc.		Optical	front panel

Table 1 : external interfaces

5.5 Data formats

The L1Topo module is entirely based on programmable logic devices. Therefore the detailed data formats do not affect the hardware design of the module. At the time of writing this documents data formats on the real-time interfaces are being defined and written up. A preliminary input format for data coming from the CMX modules is shown in Table 2.

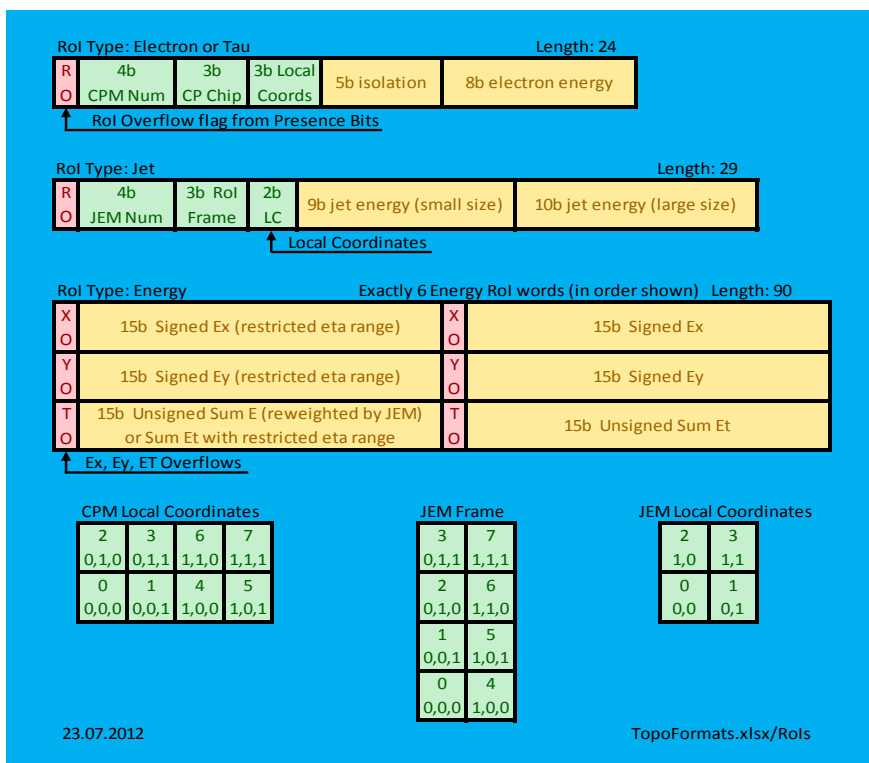


Table 2: Data format for input from CMX

6 Appendix

6.1 Checklist for detailed design

Detailed rules regarding signal integrity are to be followed so as to make sure the high density/high speed module can be built successfully. In addition a few details on signal wiring for FPGA control pins are listed. This list might be expanded for a detailed design review.

The rules with respect to power supply are:

- Use low-noise step-down converters on the module, placed far from susceptible components.
- Use local POL linear regulators for MGT link supplies
- According to the device specifications the following supply voltages need to be applied to the FPGAs: $V_{ccint}=1.0\pm 0.05V$, $V_{ccaux}=1.8V$
- On all FPGA supply voltages observe the device specific ramp up requirement of 0.2ms to 50ms.
- Run all supply voltages on power planes, facing a ground plane where possible, to provide sufficient distributed capacitance
- Provide at least one local decoupling capacitor for each active component

ATLAS Level-1 Trigger

- For FPGAs, follow the manufacturer's guidelines on staged decoupling capacitors (low ESR) in a range of nF to mF
- Observe the capacitance limitations imposed by the voltage convertors
- Minimise the number of different V_{CCO} voltages per FPGA to avoid fragmentation of power planes
- avoid large numbers of vias perforating power and ground planes near critical components

The rules with respect to general I/O connectivity are:

- Tie V_{CCAUX} and most bank supplies to 1.8V. A given FPGA is supplied by only one 1.8 V plane.
- Use all processor FPGA banks for LVDS (1.8V) only
- Use HP banks on the control FPGA for LVDS connections to the processor FPGAs and mezzanine modules
- For the control FPGA only: wire a small number of banks for 3.3V single ended operation (HR banks)
- Neither reference voltages nor DCI termination are required on the processor FPGAs. Use respective dual-use pins for I/O purposes
- For the control FPGA HR banks allow for DCI termination on single ended lines

The rules with respect to single ended signalling are:

- Run FPGA configuration and FPGA JTAG clock lines on approximately 50 Ω point-to-point source terminated lines
- Observe the requirements on overshoot and undershoot limitation, in particular for System ACE and FPGA JTAG and configuration lines. Use slew rate limited or low current signals and/or series termination

The rules with respect to differential signalling are:

- For discrete components, use internally sink-terminated devices throughout. Any non-terminated high-speed devices need to be documented in a separate list.
- Use LVDS on all general-purpose FPGA-FPGA links
- Use LVDS on all GCK clock lines
- Use DC coupling on all LVDS lines
- Design all LVDS interconnect for 1Gb/s signalling rate
- Use CML signalling on all MGT lines, for both data and clocks
- Design all MGT data links for 10Gb/s signalling rate
- Generally use AC coupling on all MGT differential inputs and outputs, for both data and clocks
 - SFP devices might be internally decoupled, microPOD transmitters might have a sufficient common mode range to allow for direct connection
- Use CML on all common clock trees; rather than using AC coupling, observe the signalling voltage and input compatibility rules as outlined by the device manufacturers
- Use AC coupling or suitable receivers when crossing voltage or signal standard domains, except on LVDS
- Place coupling capacitors close to source
- Use bias networks on AC coupled inputs where required
- Route all differential signals on properly terminated, 100 Ω controlled-impedance lines
- Have all micro strip lines face a ground plane
- Have all strip lines face two ground planes or one ground plane and one non-segmented power plane
- avoid sharply bending signal tracks
- minimise cross talk by running buses as widely spread as possible
- Avoid in-pair skew, in particular for MGT links and clocks
- Make use of device built-in programmable signal inversion to improve routing
- Avoid impedance discontinuities and stubs, in particular on MGT links and clocks

The rules with respect to processor FPGA pre-configuration and configuration control pins are:

- Wire configuration lines for optional JTAG or slave serial configuration
- Allow mode lines M0, M2 to be jumpered to either V_{CC} or GND. Pre-wire to V_{CC}
- Connect M1 to the CPLD (GND=JTAG mode, V_{CC} =slave serial)

ATLAS Level-1 Trigger

- Connect PROGRAM, INIT and DONE lines to the CPLD
- Pullup DONE 330Ω, INIT 4k7 PROGRAM 4k7
- Connect Vccbatt to GND
- Wire DIN, DOUT and CCLK (series terminated) configuration lines to the CPLD

The rules with respect to system monitor pins are:

- Connect DXN, DXP to I2C based monitoring circuits
- Decouple analog power and GND according to UG370 with ferrite beads and wire the system monitor for internal reference (both Vref pins to analog GND)
- Do not use analog sense lines Vn and Vp and connect to analog GND

6.2 Glossary

1000BASE-SX	Ethernet optical (multimode) physical layer
8b/10b	Industry standard data encoding scheme for purpose of DC balance and run length limitation (bit change enforcement)
ATCA	Advanced TCA, Advanced Telecommunications Computing Architecture
Avago	Manufacturer of 12-channel opto-electrical transceivers. The Avago transceivers used on L1Topo are of miniPOD type
Base interface	ATCA compliant backplanes provide pre-defined redundant IP connectivity via Ethernet 10/100/1000 from each slot to two modules in the crate (dual star)
BLT	Backplane And Link Tester module, fits in CMM slot of the L1Calo processor backplane
CAN	Controller Area Network, a differential serial bus (automotive)
CML	Current Mode Logic, a high-speed differential signalling standard
CMX	Future replacement of the L1Calo Common Merger Module
CTP	The Central Trigger Processor of ATLAS
DAQ	Data Acquisition (link). Physical layer G-link compatible (L1Calo standard)
G-link	Pre-historic HP/Agilent Phy chip. ~ 1Gb/s, proprietary link encoding scheme
GOLD	Generic Opto Link Demonstrator
IBERT	Xilinx automated bit error rate test tool for MGTs
IPMB	Intelligent Platform Management Bus (redundant, IPMB-A and IPMB-B), located in ATCA zone 1
JEM	Jet – and Energy module, L1Calo module delivering trigger signals via CMX
LVDS	Low-Voltage Differential Signaling standard
MGT	Multi-Gigabit Transceiver
miniPOD	High density 12-channel opto-electric transceiver
MPO/MTP	Industry standard optical connector for fibre bundles, here 12-72 fibres
Phy	A device implementing the physical level (electrical representation) of a network protocol
Quad	The Virtex Serialiser/Deserialiser circuits are arranged in tiles of four MGT links each
ROI	Region of Interest, as signaled to 2 nd level trigger. ROI link to Level-2 has G-link data format (L1Calo standard)
RTDP	Real-time data path, i.e. the data path going to the CTP. Latency critical path.
RTM	Rear Transition Module (note: ATCA RTMs mate with the front modules directly in Zone 3, not via the backplane)
SGMII	Serial Gigabit Media Independent Interface (for Ethernet Phy), 1.25Gb/s
SFP	Small Form factor Pluggable, electromechanical standard for opto transceiver
TTCDec	L1Calo specific mezzanine module for connection to the ATLAS Timing, Trigger and Control system, based on the TTCrx chip
VME(bus)	Versa Module Eurocard (bus). L1Topo is optionally connected to a VME module for purpose of module control, via a bidirectional serial link initially
Zone	ATCA connector zones 1 (mainly power), 2 (backplane data links), 3 (RTM)

	connections)
--	--------------

6.3 Change log

2012

May 30 – add drawings

May 30 – initial release